User Guide

FortiNDR Cloud 26.1.0

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO LIBRARY
https://video.fortinet.com

FORTINET BLOG
https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT
https://support.fortinet.com

FORTINET TRAINING & CERTIFICATION PROGRAM
https://www.fortinet.com/training-certification

FORTINET TRAINING INSTITUTE
https://training.fortinet.com

FORTIGUARD LABS
https://www.fortiguard.com

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com

February 3, 2026
FortiNDR Cloud 26.1.0 User Guide
78-261-1243239-202600203

TABLE OF CONTENTS

Change Log **9**
Getting started **10**
Logging into the portal 10
Navigating the portal 12
Configuring global search 13
Overview **15**
Network entity 16
Network events 16
Key terms and concepts 17
Dashboard **19**
Shared dashboards 21
Observation details 21
Frequency of observation graph 21
Observation instances table 21
Observation selector 22
MITRE ATT&CK 22
Detection indicators 23
Viewing the MITRE ATT&CK Matrix 23
DPI dashboards 24
DPI - Threats 24
DPI - AppCtrl 25
DPI - OT 26
DPI - Gen AI 27
Gen AI dashboard 28
Gen AI custom dashboard 30
Creating custom dashboards 31
Customs dashboards 32
Detections **34**
Detector Categories 36
Triage detections 37
Search for detections with the detector description 40
Impacted Devices 41
Adding custom filters to a detector query 41
Search for a device hostname in detections 42
Muting 43
Mute all detectors for a device 43
Mute a detector 44
Mute a device 44
Viewing muted devices 45
Excluding devices 45
Disabling detectors 46
Resolving detections 46
Creating a detector 48

Detector settings 49
Start an investigation 50
Viewing related investigations 51
Running queries in a detection 51
Entity Panel 52
Entity information tabs 53
Adding annotations and viewing malicious files 54
Date ranges 54
Detections visualizer 55
Nodes 56
Filtering the Visualizer 58
Action buttons 58
Detections device timeline 58
Detections table 61
Filtering events 62
Identified Assets 64
Detections context 64
Detection context page 65
Behavioral observations 66
Behavioral Observations page 67
Working with Behavioral Observations 68
Behavioral Observation fields 69
Assigning detections 70
Assigning detections from the Detections Table 70
Assigning detections from the Triage detections page 71
Viewing assigned detections 72
Statistics 73
Managing detectors 73
Response configuration 75
Creating column profiles 76
Risk score calculation 77
Scoring Matrix 77
Maximum Score Limits 77
Investigations 79
Entity lookup 79
Passive DNS 80
Investigate 82
Filtering investigations 82
Creating investigations 83
Viewing investigation details 84
Adding queries to an investigation 87
Adding notes to an investigation 89
Watch an investigation 90
Facet Search 90
Tag and comment events 92
Investigation tooltip 96
Share investigations 96

Packet capture 98
Packet capture tasks 98
Reviewing a task 99
Creating a packet capture 99
Terminating and deleting packet captures 100
BPF resources 101
PCAP encryption 104
Managing encryption keys 105
Encryption key settings 107
Private search 108
Creating IQL queries in Private Search 109
Guided queries 112
Adding a guided query to an investigation 114
Running a guided query of event records 114
Threat intelligence 116
Example query: 116
Search for intel 117
Example search for intel 118
Reports 119
Generating reports 119
FortiNDR Cloud Network Traffic Usage of a Sensor Report 119
FortiNDR Cloud Network Traffic Usage Report 119
FortiNDR Cloud Detections Report 120
FortiNDR Cloud Network Security Posture Report 121
Report history 121
View investigations 121
Pending queries in reports 123
Settings 124
Profile settings 124
My profile 124
Authentication 124
API Tokens 125
Email notifications 126
Manage annotations 127
Viewing annotations in the portal 128
Automatic critical asset identification 128
Managing annotations 129
Mutes and excludes 132
Mutes tab 132
Excludes tab 133
Subnets tab 133
Sensors 134
Sensors toolbar 135
Account telemetry 136
Sensor details 136
Account telemetry 141
Sensor settings 142

Device view 143
Account management 145
Creating users and assigning roles 146
SAML SSO 151
PCAP encryption keys 154
Multi-factor authentication 155
User activity timeout 156
Disabling an account 156
Sensor email alerts 157
Device enrichment 157
Add or edit subnets 159
Light/Dark Mode 160
Sensors deployment 162
Sensor specifications 162
Sensor Types 162
Network interfaces for physical sensors 163
Minimum virtual sensor (ESX) host requirement 163
Network data sources 163
SPAN (mirror) port 163
Network TAP 164
Network aggregator 165
ERSPAN 165
Complex or combination deployments 165
Sensor deployment strategy 165
Sensor data source configuration 168
Collector interface 168
Configuring the collector Interface 168
NetFlow 170
Configuring NetFlow for FortiNDR Cloud 171
ERSPAN 174
Enabling ERSPAN on the FortiNDR Cloud sensor 174
Zscaler ingestion 177
Zscaler setup 177
Zscaler events 179
Sensor provisioning 181
Generate a registration code 181
Register a sensor 182
FortiAI 185
Licensing 185
Tokens 185
Enabling FortiAI 185
Using FortiAI 186
FortiAI Data privacy 186
FortiNDR Cloud Integrations 188
Fortinet Automation Service 189
FortiNDR Essentials Solution Pack 189
Getting started with the Fortinet Automation Service 190

Provisioning the service 190
Installing the agent 191
Installing and configuring connectors 192
Running playbooks 193
Solution pack versions 197
FortiNDR Cloud APIs 198
Available APIs 198
Metastream 198
IQL reference guide 199
Purpose of this reference guide 199
Using guided queries 199
Sample queries 199
Core IQL concepts 200
IQL Clause 200
Fields 201
Value Types 204
Object Types 205
Fields and field types 208
Field types 208
Enriched object field types 209
Common fields 218
Event fields 220
IQL operators 263
Comparison operators 263
Logical operators 264
Exclude operators 264
Pattern operators 265
Units 265
Supported units 265
Fields with units 266
Advanced Query Concepts 266
Putting it all together 266
Array matching 267
Aggregations 268
De Morgan's Law 269
Field reference 270
Schema and field references 270
Event-type expansion 270
Field expansion 271
Synthetic fields 272
IQL query examples 272
Insight Query Language Basics 272
DCE-RPC Examples 273
DNS Examples 273
HTTP Examples 273
Flow Examples 273
FTP Examples 274

Kerberos Examples 274
NTLM Examples 274
PE Examples 274
RDP Examples 274
SMB Examples 275
SMTP Examples 275
SSH Examples 275
SSL Examples 275
X509 Examples 276
Natural Language queries 277
Supported event types and languages 277
Feature Constraints 278
Running Natural Language queries 278
NL Query guidelines 279
Best Practices 279
Example queries 279
Flow Examples 279
DNS Examples 279
HTTP Examples 280
SSL Examples 280
X509 Examples 280

Change Log

Date	Change Description
2026-02-03	Updated Event fields on page 220.
2026-01-19	Updated Fortinet Automation Service on page 189.
2026-01-08	Initial release of version 26.1.0.

Getting started

The following table provides a list of tasks to help you get started with FortiNDR Cloud:

Enable Multi-Factor Authentication (MFA)	Require all users to enter an MFA token when they log into the FortiNDR Cloud portal. To enable MFA, see Multi-factor authentication on page 155.
Configure email notifications	By default you will receive an email notification for every detection in your account and a daily digest summarizing all of the detections from the past 24 hours. To customize your email notifications, go to *Settings* > *Email Notifications*. For more information, see Email notifications on page 126.
Review the data available to you	• Network entity on page 16 • Network events on page 16 • Enriched object field types on page 209
Perform an Entity Lookup	An *Entity Lookup* is the starting point for an investigation. For more information, see Entity lookup on page 79.
View the Entity Panel	The *Entity Panel* displays the contextual information collected for an entity from within and outside the network. For more information, see Entity Panel on page 52.

Logging into the portal

Users can log into the FortiNDR Cloud portal using either a FortiNDR Cloud account or Single Sign-On (SSO).

The following table provides an overview of how user accounts are managed in FortiNDR Cloud, including user creation, multi-factor authentication, and permission management.

Account	User creation	Multi-Factor Authentication	Permission management
FortiNDR Cloud	Admin creates user in FortiNDR Cloud	Managed by FortiNDR Cloud	• Admin assigns permissions • Training access included automatically
SSO enabled	Admin creates user, or user logs in with SSO	• Managed by SSO provider if logging in with SSO	• Managed at the time the user is created in the portal (or later)

Account	User creation	Multi-Factor Authentication	Permission management
		• Manged by FortiNDR Cloud if logging in with FortiNDR Cloud username and password.	• Training access included automatically
SSO enabled with SSO only	User logs in with SSO	Managed by SSO provider	• May be added only after user logs in once • Training access only when first logged in
SSO not enabled	Admin creates user	Managed by FortiNDR Cloud	• Managed at time user is created (or later) • Training access included automatically

- FortiNDR Cloud only supports IdP initiated SAML.
- Assertions must always be signed.

You can log into the FortiNDR Cloud portal with an email address or with a FortiCloud sub-user account.

To log into the portal:

1. Go to https://portal.fortindr.forticloud.com/.
2. Do one of the following:

Log in with	Description
Email	Enter your email address and password, then click *Login*.
FortiCloud	1. Click *FortiCloud*. The FortiCloud login page opens. 2. Enter your FortiCloud email address, password and token to login. You can only login in with a FortiCloud sub-user account. The FortiNDR Cloud portal does not support IAM users at this time. For information, see *User permissions* in the FortiCloud Services Guide.

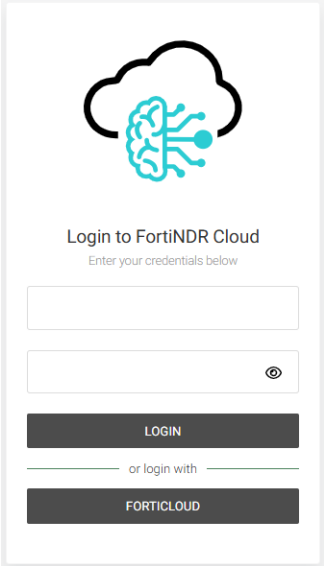

Navigating the portal

This topic provides an overview of the navigation tabs and menus in the FortiNDR Cloud portal.

The portal is organized into tabs located in the navigation menu at the top of the portal. Links to the product documentation and *Settings* pages are located in the top-right corner of the page.

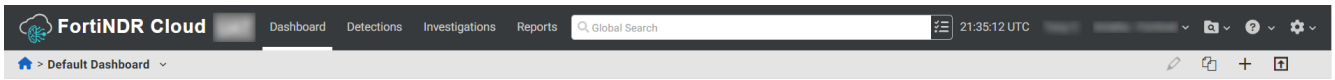

Dashboard	This is the landing page for the FortiNDR Cloud portal and provides high-level summary information. For more information, see Dashboard on page 19.
Detections	This tab shows detections that have fired in your account. For more information, see Detections on page 34.
Investigations	This is where you perform queries or run guided queries for forensic analysis and hunting over your network data. For more information, see Investigations on page 79.
Reports	Use this tab to run the *FortiNDR Cloud Network Security Posture Report* and the *FortiNDR Cloud Detections Report*. For more information, see Reports on page 119.

Global Search	Use the Global Search function to search FortiNDR Cloud with a text string, IP address or domain. Search results are organized by *Detections*, *Detections Coverage*, *Investigations*, *Search Timeline* and *Entity Lookup*. You can enter multiple IPs or domains separated by a comma or a space. However, if you are performing a bulk search for IPs FortiNDR Cloud will stop the search after it finds the first IP in the list.
Settings	This icon located in the top-right provides access to auxiliary pages related to user and account settings and management. For more information, see Settings on page 124.

Configuring global search

The Global Search function allows you to search FortiNDR Cloud using a text string, IP address, or domain. You can enter multiple IPs and domains, separated by a comma or space.

You can configure Global Search to:

- Show or hide categories
- Limit the number of results
- Arrange the order of results on the page

To configure global search:

1. Click the dropdown menu at the right side of the search field. The *Configure Global Search* dialog opens.

2. Configure the search settings.

Include	Select/Deselect the categories to appear in the results.
Limit	Select 5, 10, or 50 results to be displayed.

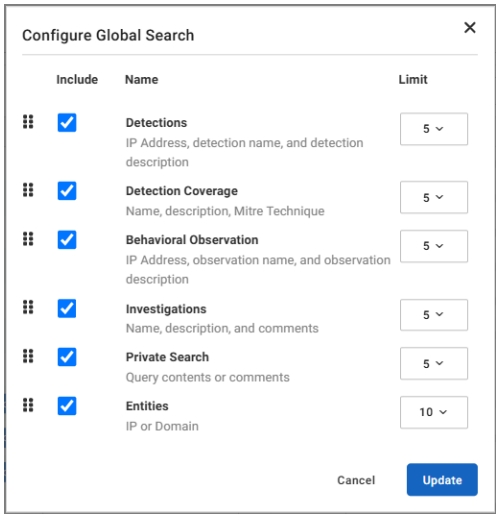

3. To arrange the order the results are displayed, drag a heading up or down in the dialog.
4. Click *Update*.

Overview

FortiNDR Cloud is a cloud-native network detection and response solution built for the rapid detection of threat activity, investigation of suspicious behavior, proactive hunting for potential risks, and directing a fast and effective response to active threats.

The following diagram illustrates the components and benefits of the solution at a high level:

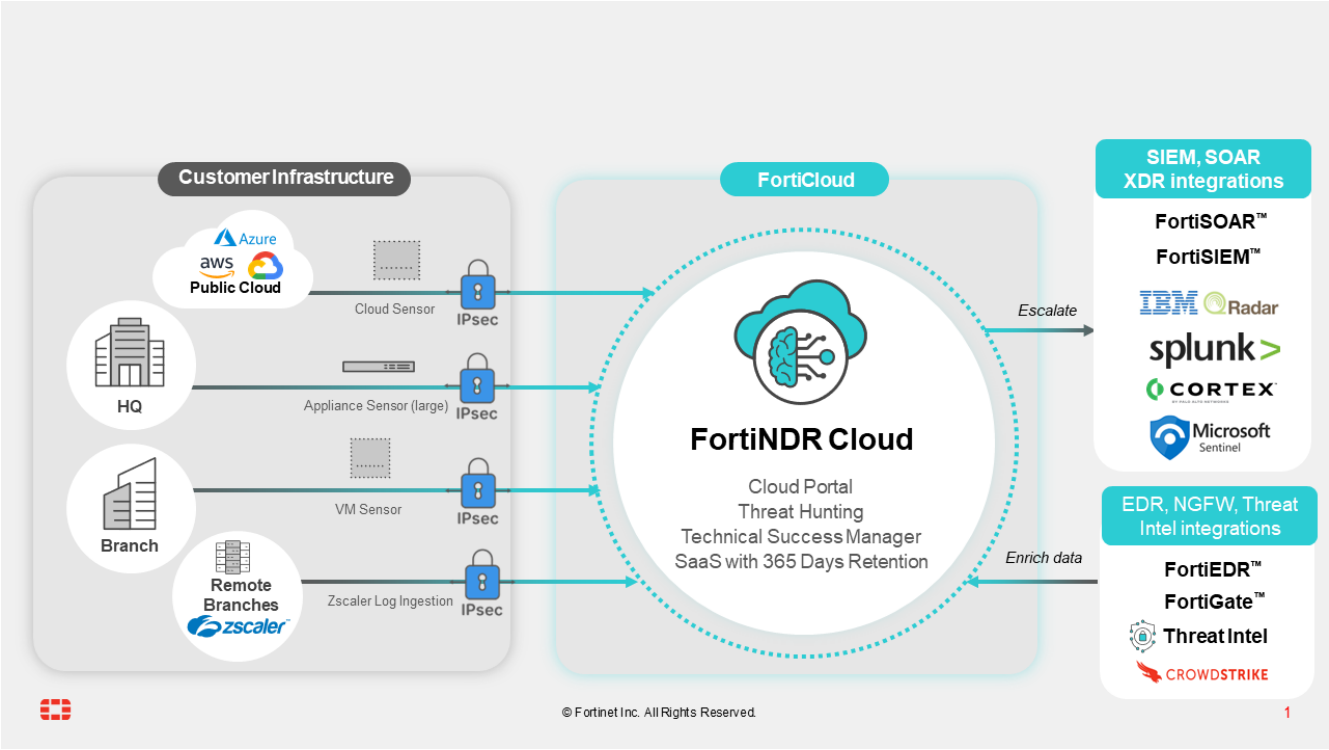

Key notes relating to architecture and securing customer data:

- Data from customer and/or public cloud sensors encrypts network meta data collected to SaaS solution with strong IPSEC encryption. This encryption is end-to-end to ensure customer network metadata is not compromised (data in transit).
- Network data from customers is encrypted at rest in FortiNDR Cloud.
- Customers will have a portal which enable access to illustrate detection, conduct investigations, and threat hunting.
- Third-party integrations such as EDR, NGFW, SIEM and SOAR products are enabled via APIs available from FortiNDR Cloud.
- FortiNDR Cloud data are enriched with different threat and network feeds to make data useful to comprehend.
- Network metadata collected do not contain PCAPS (despite it being possible to collect PCAPS on sensors for forensic analysis), please see further chapters on enabling PCAPs
- Fortinet data security and privacy practices are documented here: Data Privacy Practices

Network entity

An *Entity* is a unique identifier on the network. At this time, IP addresses and domains are supported entities. Entities are extracted from the event data and catalogued in their own data store. Contextual information is then added to the entities when applicable such as:

- First seen / last seen timestamps
- Associated hostnames and usernames from DNS, DHCP, Kerberos, and NTLM events
- WHOIS and Registration information
- VirusTotal intelligence
- Associated software

Entities observed in your account are stored for one year. This allows analysts to determine who is interacting with the network and answer questions such as:

- Which / how many of my hosts are interacting with this entity?
- Who is responsible for this entity?
- What other entities are associated with this entity?
- What does everyone else know about this entity?

Working with entity information

You can perform an *Entity Search* (or Lookup) by simply entering an IP address or domain in the *Search* field at the top navigation menu. An Entity Search is an excellent starting point for an investigation if you have very little information to work with, because the entity record may contain important contextual information. For more information about entity searches, see Entity lookup on page 79

The *Entity Panel* displays all of the information collected for an entity from both within and outside of the network. You can access the *Entity Panel* for an entity by left-clicking any entity anywhere in the portal. For more information, see Entity Panel on page 52

Network events

FortiNDR Cloud network sensors perform deep packet inspection of all observed network traffic and extract key protocol metadata for processing by the FortiNDR Cloud data pipeline. This metadata is organized into records called *Events*.

Flow

A *flow* is how FortiNDR Cloud organizes traffic for parsing and tying together events. A flow is a unique session between two hosts. Specifically, a flow is a collection of continuous packets having the same unique five-tuple (source IP, source port, destination IP, destination port, transport protocol) within a short time frame.

Every flow is identified with a unique `flow_id`. Multiple events can be produced from a single flow and are assigned the same `flow_id`.

There three categories of events:

- *Flow events*: The *Flow* event type, contains metadata from the lower layers of the OSI model (IPs, ports, byte counts, transport protocol, etc).
- *Protocol events*: Most event types such as DNS, HTTP, and SSL, contains metadata from the upper layers of the OSI model.
- *Synthetic events*: The *Suricata* and *Software* event types, contains metadata produced by processes that scan or analyze traffic rather than metadata taken directly from network traffic.

Every flow will have exactly one *Flow* event, zero or more protocol events, and zero or more synthetic events. There can only be one *Flow* event because FortiNDR Cloud can summarize all the networking/flow data in one record. There can be zero or more protocol events because the flow could be a raw network socket with no known application, an HTTP connection with numerous HTTP requests over the same connection, an RDP connection over SSL with an X.509 certificate exchanged, or anything else. Similarly, one flow could trigger twelve Suricata queries just as easily as zero queries.

Regardless of how many events are produced from a single flow, FortiNDR Cloud assigns them the same unique `flow_id`, which provides a bigger picture surrounding other events in the session.

Working with events and flows

Running a query will return a list of events. If an event in the list stands out for some reason, you can run a separate query for that event's `flow_id` to see what other events were produced during that session/connection/conversations/flow.

Protocols are parsed regardless of port or service. Events are normalized for time and enriched with Geo-IP information and Threat Intelligence for additional context. Once this processing and enrichment is finished, events are surfaced through the FortiNDR Cloud portal and APIs.

For a complete list of supported field types, go to *IQL reference guide* > *Fields and field types on page 208*.

Key terms and concepts

Term	Definition
ATR	FortiGuard Applied Threat Research
Behavioral Observation	A *Behavioral Observation* is an output from a system that analyzes events and behaviors to identify potentially malicious activity (e.g., *Domain Similar to Malware DGA Domain* and *Malicious PE File*). Depending on your environment, not all Behavioral Observations indicate malicious activity. For example, if you recently created a new SSH server, then the *New SSH Server* observation is not malicious. See, Behavioral observations on page 66.
Detection	An alert mechanism that notifies you when a unique pair of events satisfy a detector. Detections allow you to quickly identify and respond to suspicious or known malicious activity in your network.

Term	Definition
Detection lifecycle	The status states of a detection (*Active*, *Muted*, or *Resolved*).
Detector	A query and other parameters used to detect something.
Dwell	Average time (in seconds) between when an incident was first seen and when it was resolved. See the *FortiNDR Cloud Detections Report* section in FortiNDR Cloud Detections Report on page 120.
Example	Example dashboards are custom dashboards created by Fortinet and shared with all customers, allowing users to view and use them within their own environments.
Five-tuple (5-tuple)	The source IP, source port, destination IP, destination port, and transport protocol. For more information, see Network events.
Flow	A collection of continuous packets having the same unique five-tuple (source IP, source port, destination IP, destination port, transport protocol) within a short time frame.
Indicators	An *indicator* is a field value extracted from a detection's event(s) as defined by the detector. This information is useful for identifying related activity and tracking indicators over time. Detectors can define up to five fields to extract indicators from, and each detection can store up to five unique indicators for each indicator field.
Mean Time To Detect (MTTD)	Average time (in seconds) between when an incident was first seen and when it was created in the system. See the *FortiNDR Cloud Detections Report* section in FortiNDR Cloud Detections Report on page 120.
Mean Time To Resolve (MTTR)	Average time (in seconds) between when an incident was created and when it was resolved. See the *FortiNDR Cloud Detections Report* section in FortiNDR Cloud Detections Report on page 120.
MITRE ATT&CK	*MITRE ATT&CK* is a knowledge base of threat behaviors relied upon by security professionals worldwide. You can map FortiGuard Lab detectors to MITRE ATT&CK, to enable visibility into the threat coverage provided by FortiNDR Cloud.
Tuning	The process of hiding known behaviors in a detector using one of the following three mechanisms: • *Muting*: Hides a detection but allows it to be created. Muted detections can be reviewed in bulk on a recurring basis. See Muting detectors. • *Excluding*: Prevents detections from ever being created. Excluded detections cannot be reviewed in bulk on a recurring basis. See Excluding devices. • *Filtering*: Tuned out everything else, (such as external entities and non-entity fields) by adding your own logic to detectors authored by FortiGuard Labs to customize the detector to your network. See Adding filters to detectors.

Dashboard

The *Default Dashboard* organizes detection activity into interactive widgets that display high-risk devices, detection trends, severe threats, and resolved issues over time, helping analysts quickly review threats and prioritize their next steps.

The global date picker (located above the *Detections by MITRE Tactic* widget) applies a selected time range across multiple widgets, ensuring consistent context when reviewing historical data. Analysts can identify high-risk devices, examine detection severity by confidence level, and monitor notable patterns such as new or spiking detections. Additional widgets present detections by MITRE ATT&CK® tactics, behavioral observations, and resolution trends, providing structured insights for threat analysis and response.

The default dashboard features a clean, modern layout with enhanced functionality and rich visualizations. Its design is centered around analyst workflows and risk-based prioritization. All widgets (both default and custom) adopt a simplified style for a streamlined appearance. Widgets load in a structured sequence to ensure visual consistency during page loading.

Core functionality:

- Identify high-risk devices by reviewing risk scores and crown icons, and open the entity panel for detailed context.
- Review detection severity using confidence-level filters and switch between severity levels with dropdown controls.
- Monitor notable detections with tags such as *New* or *Spike* to spot emerging or unusual activity.
- Explore detections by MITRE tactic and track resolved detections with visualizations.
- Access behavioral observations and other widgets for additional context.
- Adjust the timeline to view historical data and hover over detections for quick access to details or context.

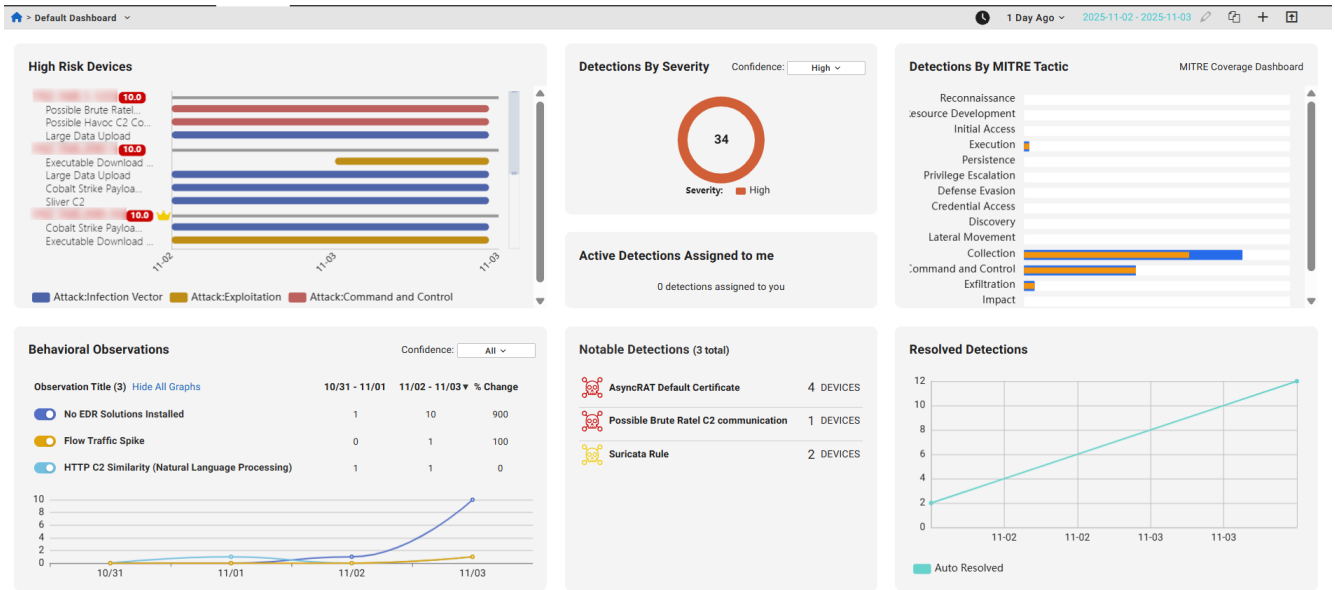

Widget	Description
High Risk Devices	Displays only devices with high risk and active detections. It highlights the top 5 high-risk devices, similar to those shown on the Detection Device Timeline page. Devices are sorted by risk score (displayed next the device IP). Click the device IP to open the Entity Panel. Hover over a bar in the chart to view details about the detection. Click the *Detection Detail* button to quickly navigate to the detection detail page for the selected detector. Click the *Detection Context* button to view the detections and observations related to this IP on the *Detection Context* page. Click the widget top open the *Detection Device Timeline* page.
Detections By Severity	Displays the number of detections by confidence level (High, Moderate, Low or All). Use the dropdown menu to change the confidence level. Click the widget name to open the *Detections Table*.
Active Detections Assigned to me	Displays the number of detections assigned to you. Click the widget name to open the *Detections Table*.
MITRE ATT&CK	Displays detections organized by the MITRE ATT&CK® framework. Each detection activity includes two bars: orange shows the previous time period, and blue shows the current range. • Hover over bars to view detection counts. • Click bars to open the *Detections Table*. • Click *MITRE Coverage Dashboard* to open the *MITRE ATT&CK Matrix*. Row names may vary depending on account coverage.
Behavioral Observations	Shows behavioral observations from the past two time ranges (1 day, 1 week, 2 weeks, 30 days). • Click the widget title to open the *Behavioral Observations* page • Click an observation title to view details • Hover over graph data points for details • Use *Hide All Graphs* and toggles to filter observations • Use the *Confidence* dropdown to filter by level (*All*, *High*, *Moderate*, *Low*)
Notable Detections	Displays active detections with the highest severity and detection count. The *New* and *Spike* labels highlight new detections and spikes in detection activity. • *New* indicates that there were no active detections during the baseline period (defined as 30 to 7 days ago), but at least one detection has occurred in the past 7 days. • *Spike* indicates that the number of active detections in the past 7 days is more than three times higher than the baseline count.
Resolved Detections	Shows daily counts of resolved detections over time, including *Total*, *Average*, and *Maximum*. • Click a data point or the *Total* count to view resolved detections in the *Detections Table*

Shared dashboards

When a user opens a shared dashboard with query charts, a new investigation is created in their own account. This ensures that:

- The query results shown are based on the current account's data, not the dashboard creator's.
- Clicking the chart title also opens the query inside the investigation specific to the current account.

When a user clones a dashboard that contains query charts, a new investigation is automatically created in the user's account for each query chart widget. This ensures that the cloned dashboard runs fresh queries and displays results based on the current account data. The investigation is independent of the original dashboard and tailored to the account.

Users with only the *Admin* role (and no additional roles like *User*) will not see dashboards that contain query charts. This ensures that only users with the appropriate permissions can access dashboards with query-based data.

Observation details

The *Observation Details* page provides detailed information about a selected observation. It includes a graph showing the frequency of occurrences and a table listing up to 1,000 recent instances.

- Use the date picker to view behavioral observations for any 90-day period within the past year.
- To view observation details for a device, enter the IP in the *Device to search* field.
- To filter the page by *Confidence*, select *All*, *H*, *M*, or *L* (Low, Moderate, or High).

Frequency of observation graph

The *Frequency of Observation* graph shows how often a specific observation has occurred over time, categorized by confidence level.

- Hover over the graph to view the number of instances by confidence level.
- To filter the table, click a confidence level (Low, Moderate, or High).
- Click on a bar in the graph to apply its time range and confidence filter to the page.
- Hover over a confidence level at the top of the graph to isolate it.

Observation instances table

The *Observation Instances* table displays the most recent instances for the selected observation, up to 1,000 entries.

- Click any column header to sort the table by that column.
- To refine the table, enter a search term in the *Filter current observation results* field and click *Filter*.

Observation selector

Use the observation selector at the top-center of the page to switch between different observations available for your account.

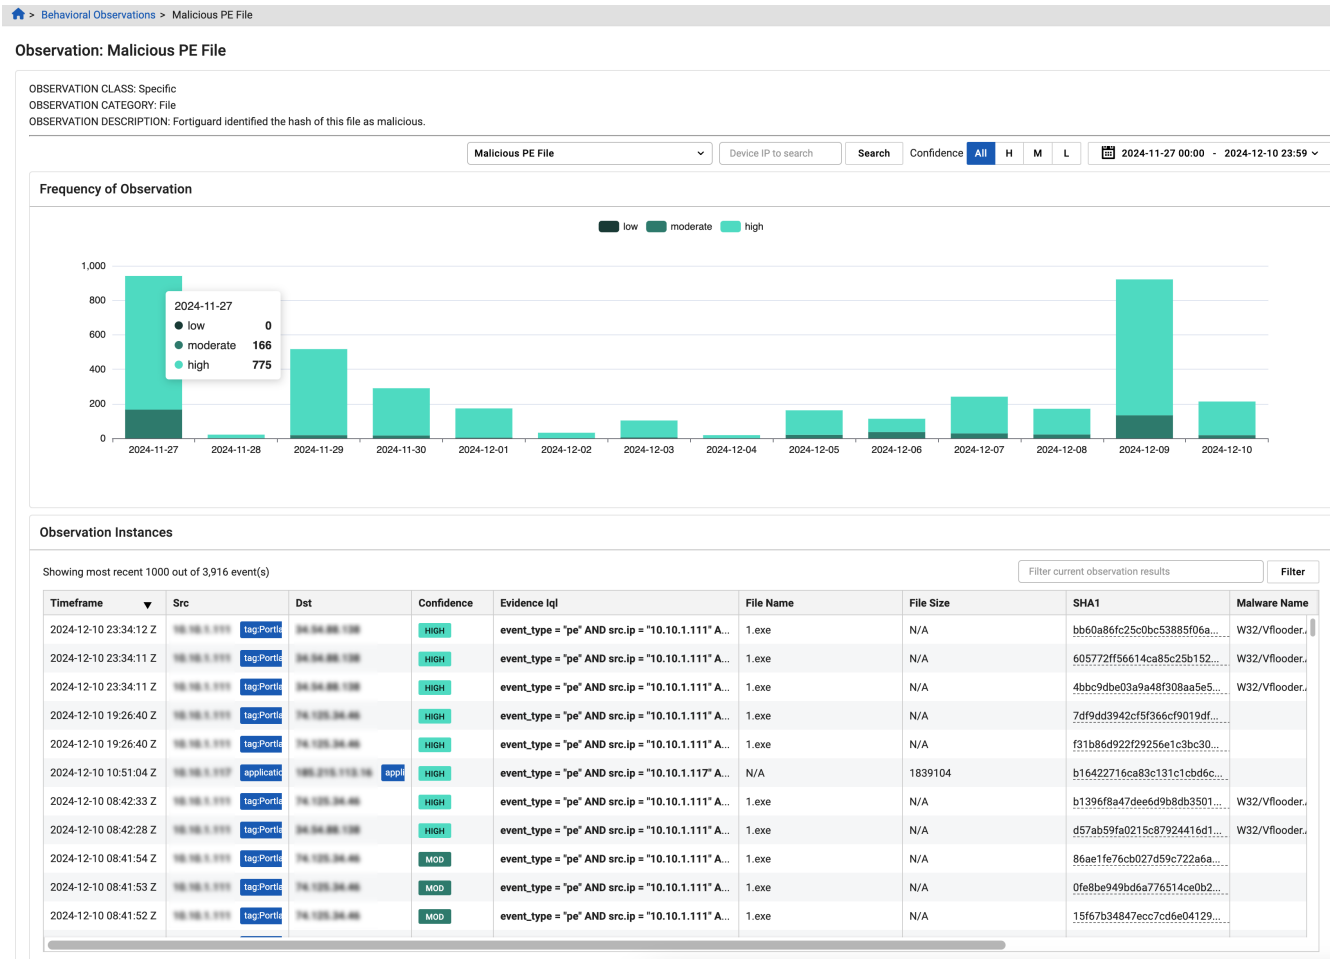

MITRE ATT&CK

The *MITRE ATT&CK Matrix* dashboard displays detection coverage based on detectors developed by FortiGuard Labs.

MITRE ATT&CK is a globally recognized knowledge base of threat behaviors and techniques used by security professionals to understand and respond to threats. FortiGuard Lab detectors can be mapped to MITRE ATT&CK to provide visibility into the threat coverage offered by FortiNDR Cloud.

The dashboard presents detections by behavior type (behavioral and non-behavioral) and by technique type (primary and secondary):

- *Primary Technique*: The main technique used to detect the behavior.
- *Secondary Technique*: A related technique that may not be directly observed on the network but is associated with the threat. This is not displayed in most cases.
  To view the secondary technique, click the plus (**+**) symbol in the bottom-right corner of a Primary Technique box.

Detection indicators

- A blue shield icon indicates active detections for a technique or sub-technique, and that you have permission to view them on the *Detections* page.
- An empty shield icon indicates that detections are resolved, but still viewable.
- Techniques shown as plain text either have no detections or you lack permission to view them.

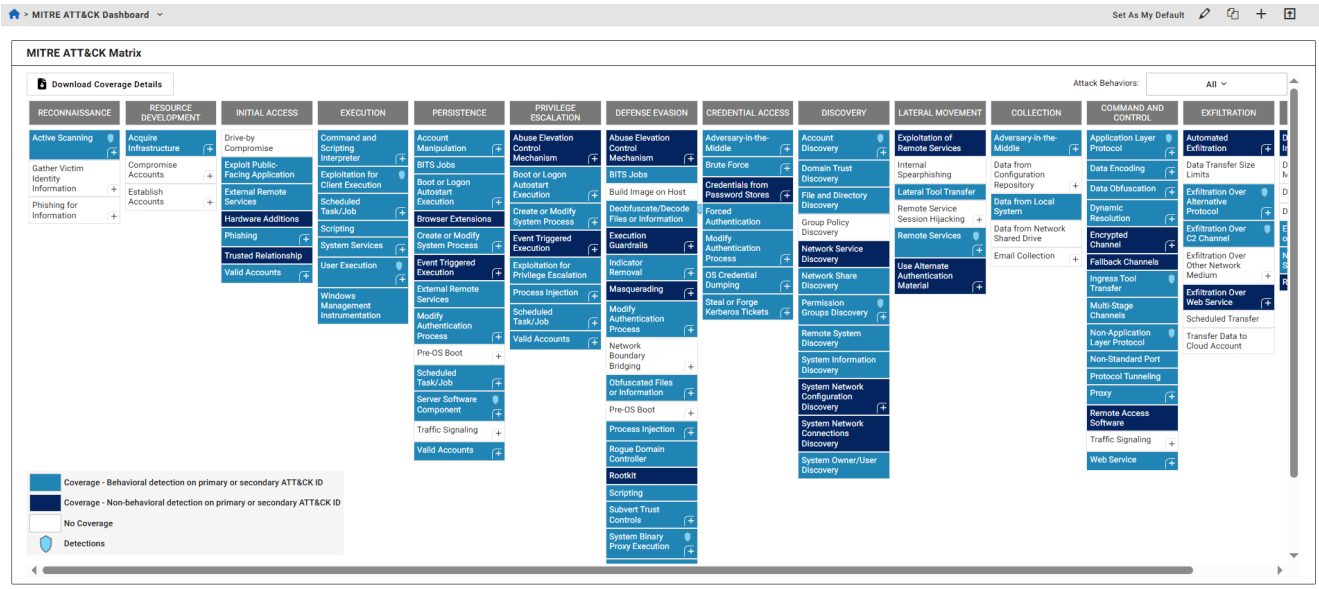

Viewing the MITRE ATT&CK Matrix

To view the MITRE ATT&CK Matrix:

1. Click the *Dashboard* tab. Do one of the following:
   - At the top left-side of the page, click *Default Dashboard > MITRE ATT&CK Dashboard*.
   - In the *MITRE ATT&CK* widget, click *Go to MITRE Coverage Dashboard*.
2. Click the *Attack Behaviors* drop-down at the top-right of the dashboard to filter the dashboard by behaviors:
   - *All*
   - *Ransomware*
   - *Insider Threat*
   - *Cyber Espionage*

3. Click a technique in the table. A summary of the technique is displayed.

Column	Description
Tactic	The tactic of the behavior.
Coverage	The coverage status of the technique and the sub-techniques.
Name	The behavior name.
ID	ID number of the technique and the sub-techniques. For techniques and sub-techniques with active detections (indicated by a blue shield icon), the ID number is a hyperlink that directs you to the *Detections* page.

To download the coverage details:

- Click the *Download Coverage Details* button to download the coverage details as a CSV file which contains the *Date Updated*, *Name*, *Primary Attack ID*, *Secondary Attack ID* and *Description*.

DPI dashboards

DPI dashboards are available from the *Dashboard* menu but will only display data when *Fortinet DPI* is enabled in the *Sensor Settings* page (see, Sensor settings). The dashboards display DPI events from either the past 24 hours or the past 7 days, depending on the dashboard. The data can be refreshed at any time. You can view the dashboards as a chart, pie chart, or table, and export the data as a CSV file. DPI dashboards are useful when starting an investigation. For example, if an IP address is flagged in one of the dashboards, you can enter it in the *Global Search* field or use it to create a query in *Private Search*.

The following dashboards are available:

- DPI - Threats
- DPI - AppCtrl
- DPI - OT
- DPI - Gen AI

DPI - Threats

The DPI - Threats dashboard displays detected threats and their corresponding counts. The dashboard provides a summary of the most frequently detected threats and highlights the IP addresses that are triggering the highest number of signatures. When an IP address triggers a large number of IPS signatures, it's a strong indicator that the IP should be investigated further.

This dashboard contains three monitors:

Widget	Description
Top Threats	This monitor queries high-severity IDS alerts (severity level 4) detected by DPI, where either the source or destination is internal. It excludes alerts triggered by device tagged as *Scan* and *Nessus* and filters out two noisy Apache-related signatures. The results are grouped by alert signature, helping identify which threat signatures are most frequently triggered.
Top IP	This monitor retrieves high-severity IDS alerts (severity level 4) detected by DPI, where either the source or destination is internal. It excludes alerts triggered by devices tagged as *Scan* or *Nessus*, and filters out two noisy Apache-related signatures. The results are grouped by source IP, helping identify which internal hosts are generating the most IDS alerts.
Botnet from internal	This monitor identifies outbound botnet-related DPI alerts where the source IP is internal. It groups the results by both the internal source IP and the specific botnet signature that was triggered, helping pinpoint which internal hosts are attempting to communicate with known botnets.

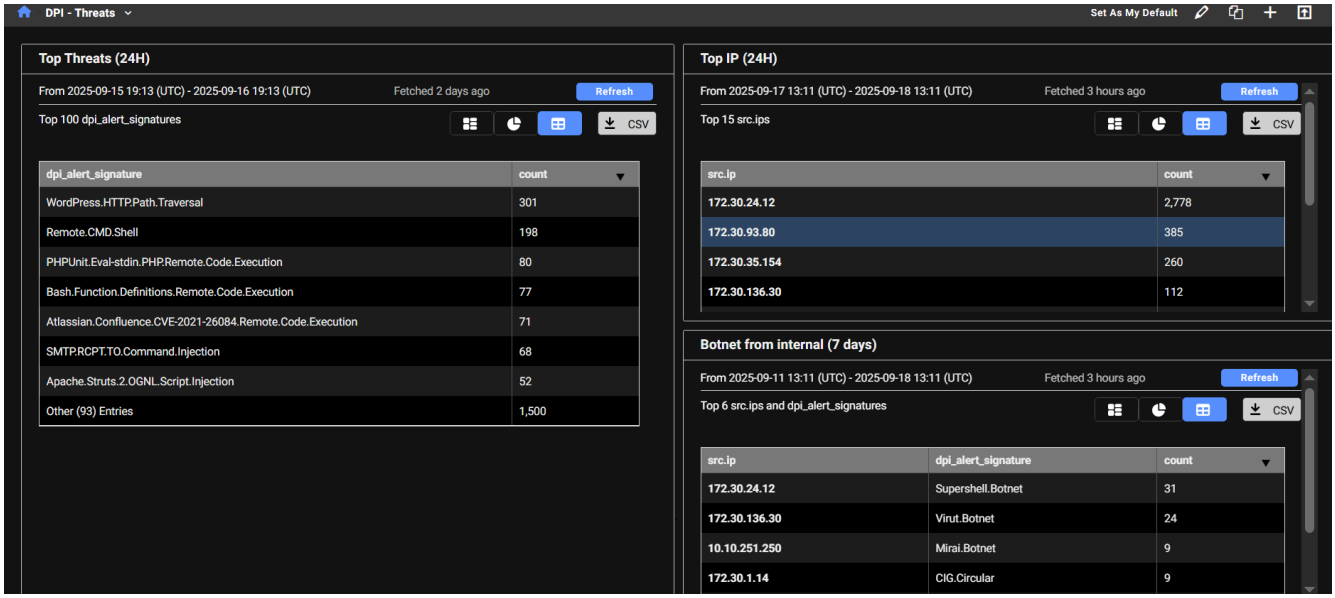

DPI - AppCtrl

The *DPI - AppCtrl* dashboard displays detections of applications and protocols used by IP addresses, such as DNS, HTTP, and other common services. This provides insight into the types and volume of traffic an IP address is generating.

Widget	Description
Top Application - Exclude Popular (24H)	This monitor filters out common or expected traffic (such as DNS, ICMP, ping, and browser activity) to highlight less typical application usage.

Widget	Description
	The results are grouped by application signature, helping identify less common or potentially suspicious applications being used internally
Top Application - All (24H)	This monitor includes all detected application types, including browser activity, offering a complete view of application traffic. The results are grouped by application signature, allowing you to see which applications are being detected across internal traffic, without the noise from automated scanners. This helps focus on legitimate or potentially suspicious application usage within the network.

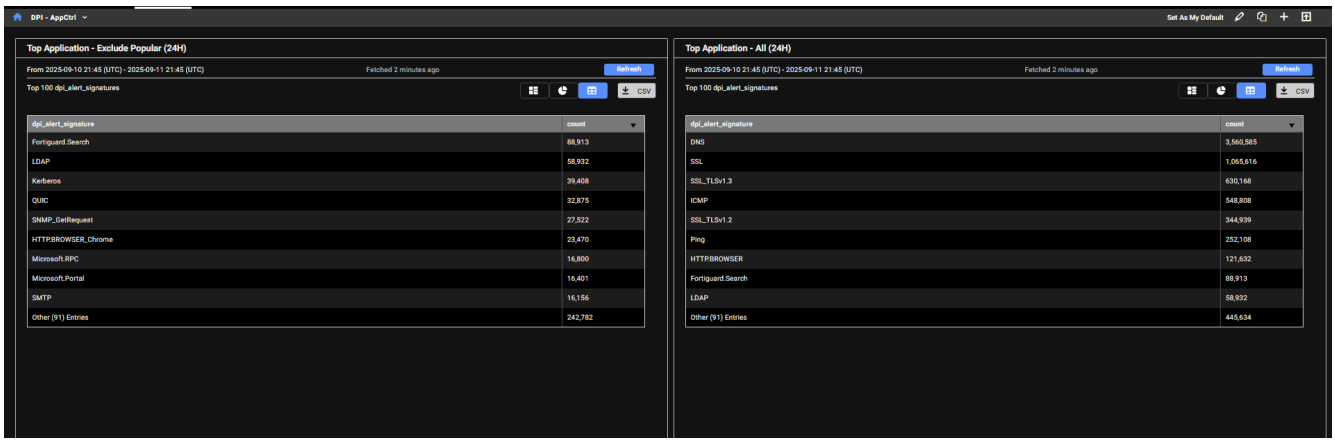

DPI - OT

The *DPI - OT* dashboard provides visibility into OT (Operational Technology) protocols used in industrial control systems. Any OT-related activity detected on the network will be tracked and displayed here. The dashboard highlights specific OT protocols (such as Bacnet, Profinet, and DNP3) with MP3 being one of the more commonly observed.

Widget	Description
OT Protocol	This monitor displays DPI alerts categorized as *OT - Protocol*, which relate to industrial control system protocols, where either the source or destination IP is internal. It excludes alerts triggered by device tagged as *Scan* and *Nessus*. The results are grouped by both the OT protocol signature and the source IP, allowing you to: • See which internal IPs are generating OT protocol traffic. • Identify which specific OT protocols are being used or triggered by each IP. This helps in monitoring legitimate OT activity and detecting unusual or unauthorized use of industrial protocols.

Widget	Description
OT Threats	This monitor displays DPI alerts categorized as *OT - Threats*, which indicate suspicious or malicious activity targeting Operational Technology (OT) systems. It filters for alerts where either the source or destination IP is internal and excludes alerts triggered by device tagged as *Scan* and *Nessus*. The results are grouped by both the OT threat signature and the source IP, allowing you to: • Identify which internal IPs are involved in OT-related threat activity. • See which specific OT threat types are being detected per IP. This helps in monitoring and investigating potential compromises or unauthorized access attempts within industrial environments.

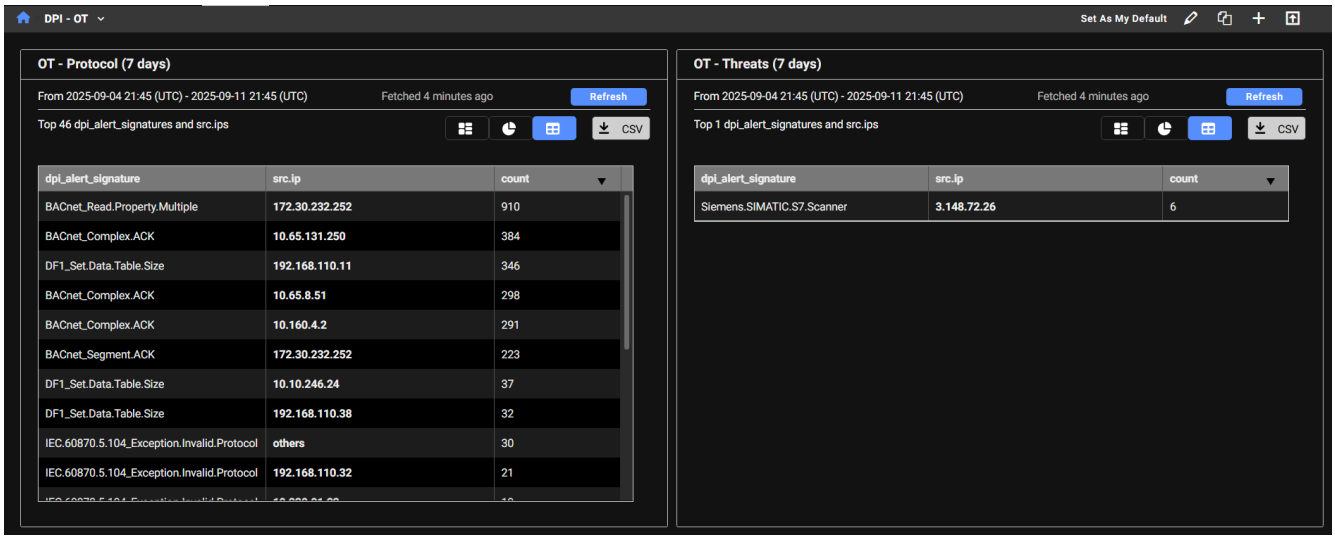

DPI - Gen AI

The *Gen AI* Dashboard provides visibility into organizational usage of Generative AI applications by tracking DPI (Deep Packet Inspection) events.

Widget	Description
Gen AI Usage by DPI Alert Signature	Displays counts of DPI events triggered by specific Gen AI alert signatures (e.g., Microsoft Copilot, Google Gemini).
Gen AI Usage by Vendors	Shows the distribution of DPI event counts across different vendors, helping identify which providers are most used.
Browser vs Client/Server Usage	Compares DPI event counts for Gen AI accessed via browser-based technologies versus client-server connections.
Gen AI Usage by Host (External)	Lists external hosts associated with Gen AI traffic, ranked by DPI event counts, providing insight into endpoints most frequently accessed.

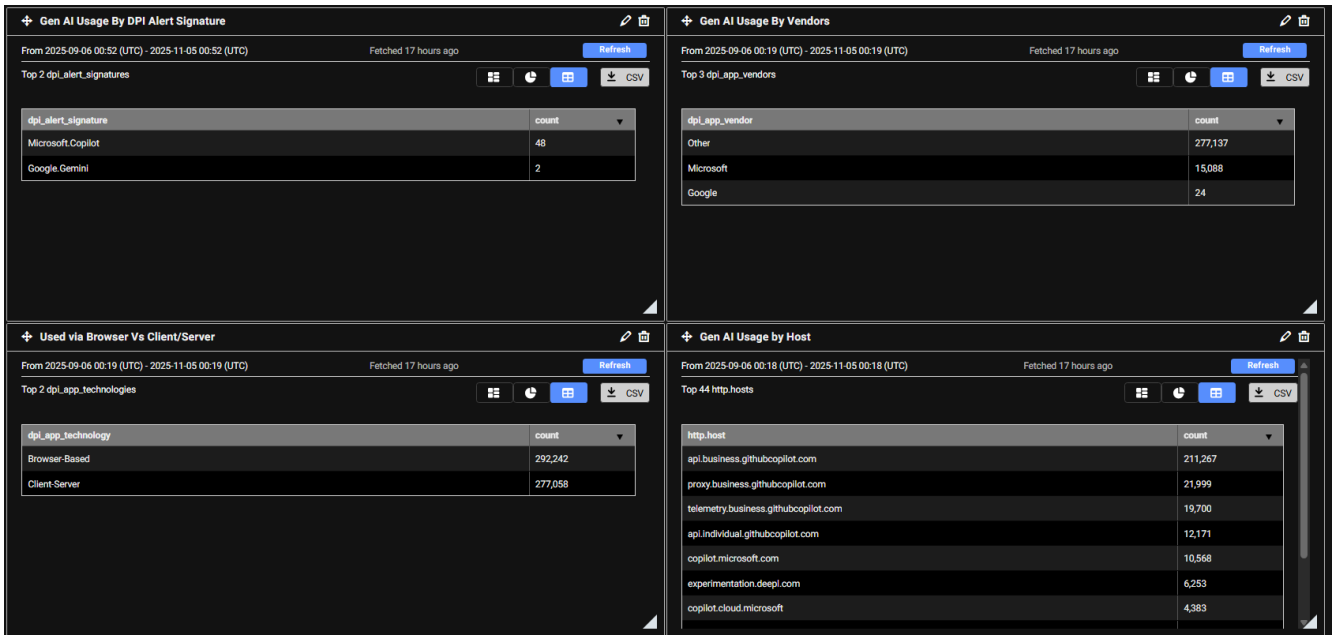

Gen AI dashboard

The *Gen AI Dashboard* provides a centralized view of generative AI usage within your organization, helping analysts identify unauthorized activity, detect potential data exfiltration, investigate anomalies with full device and detection context, and align AI activity with the organization's security policies.

The widgets consolidate all observation details from NDR Cloud into a single view, giving analysts visibility into AI usage across the organization. This unified perspective is useful for identifying patterns and anomalies by turning raw observations into actionable insights, enabling faster, more informed investigations.

You can use the widgets in this dashboard to track which AI providers are accessed and by which devices and identify large uploads and frequent connections that may indicate risky behavior. For example:

- **Detect unusual activity**: Beyond monitoring AI usage, investigators can spot anomalies such as large uploads or file transfers that might include sensitive content.
- **Drill down into context**: Clicking an IP address opens the entity panel, showing related detections and annotations. For example:
  - If the IP belongs to a guest network, large data transfers may be acceptable.
  - If the IP is tied to crown assets (e.g., domain controllers or DNS servers), a large upload to an AI service could signal a serious issue.
- **Access source details**: Hovering over an IP reveals observation details, including device and detection context. Clicking on an IP with large uploads surfaces associated detections and other alerts for that device.
- **Identify policy violations**: Unauthorized provider usage stands out. For instance, if an organization prohibits Cursor but sees activity from it, investigators can trace why it is being used and verify compliance with internal policies.

The dashboard contains the following widgets:

Widget	Description
No. of Connections to Gen AI Providers	Tracks the number of connections and provider type for the last 14 days. The widget also identifies which generative AI is a General Assistant and a Coding Assistant.
No. of Connections from Source	Shows which IPs connect to which providers and the number of connections to AI providers for the last 14 days. Click on the IP address to open the Entity Panel. Hovering over the bar opens a tooltip that connects to the *Observation Details* and *Detection Context*.
Large outbound traffic to Gen AI	Displays size of large outbound transfers for the last 14 days. Click the IP to open the Entity Panel. Hover over the lines in the chart to see details about the upload.

Widget	Description
Sum of Bytes Sent to Gen AI	Aggregates data volume per IP and provider for the last 14 days. The data is sorted highest to lowest to identify top uploader IP address. Hovering over the bar opens a tooltip that connects to the *Observation Details* and *Detection Context*.
No. of Files Uploaded by Source	Counts files sent to AI services for the last 14 days.
Large outbound traffic sessions	Monitors connection frequency of large outbound connections. Hovering over the bar opens a tooltip that connects to the *Observation Details* and *Detection Context*.

Gen AI custom dashboard

The *Gen AI Dashboard* is also available as a configurable custom dashboard.

Data source	Header	
Gen AI usage	• No. of Connections to Gen AI Providers • No. of Connections from Source • Sum of Bytes Sent to Gen AI	• Past 24 hours • Past 1 week • Past 2 weeks
Gen AI File Upload	• No. of Files Uploaded by Source	
Large Outbound Data Transfer to Gen AI	• Large outbound traffic to Gen AI • Large outbound traffic sessions	
AI/ML API Request Spike	• No. of AI/ML API Requests	

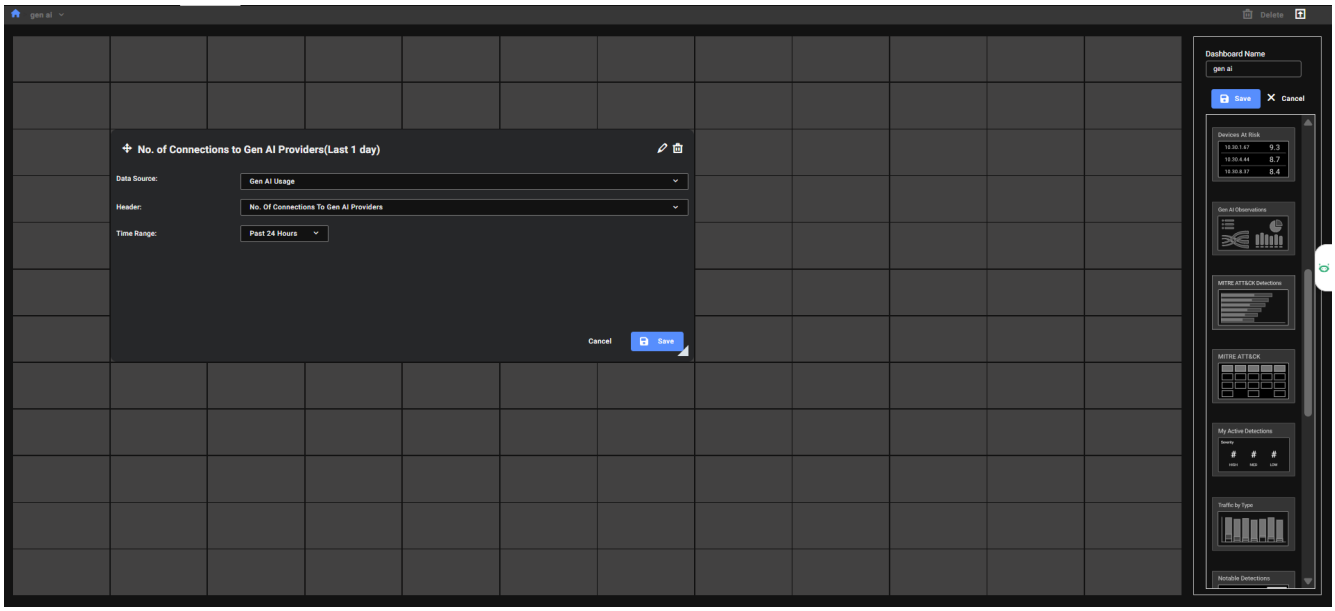

Creating custom dashboards

Combine widgets to build custom dashboards tailored to your needs. These dashboards automatically refresh approximately every five minutes.

You can also set a custom dashboard as your default view. To switch between dashboards, click the *Default Dashboard* dropdown in the toolbar at the top-left corner of the page.

To create a custom dashboard:

1. Click the *Dashboard* tab.
2. In the toolbar at the top- right corner of the page, click the plus symbol (+). The *Create Dashboard* dialog opens.
3. In the *Name* field, enter a name for the dashboard and click *Create*.
4. Drag and drop the widgets onto the dashboard.
5. Arrange the widgets on the dashboard and click *Save*.
   - To move a widget, use the handle at the top-left of the widget to drag it.
   - To change the widget name, click the pencil icon.
   - To remove the widget from the dashboard, click the delete icon.

   - Each widget uses a different amount of space on the dashboard. Some widgets may not fit onto one dashboard.
   - Each widget has a default size. Some widgets cannot be condensed smaller than their default size.

6. Click *Save*.

To edit a custom dashboard

1. Click the *Dashboard* tab.
2. Click the *Default Dashboard* dropdon at the top-left corner of the page and select a dashboard from the list.
3. In the toolbar, click the edit icon.
4. Edit the dashboard and click *Save*. The dashboard is added to the *Default Dashboard* drop down.

You cannot edit the default dashboard.

To copy a dashboard:

1. Click the *Dashboard* tab.
2. Click the *Default Dashboard* dropdown at the top-left corner of the page and select a dashboard from the list.
3. In the toolbar, click the copy icon. The *Copy Dashboard* dialog opens.
4. In the *Name* field, enter a new name for the dashboard.
5. In the *Account* drop down, select where the dashboard will appear in the menu.
6. Click *Copy*.

To set a custom dashboard as the default:

1. Click the *Dashboard* tab.
2. Click the *Default Dashboard* menu at the top-left corner of the page and select a dashboard from the list.
3. In the toolbar, click *Set as My Default*.

Customs dashboards

Dashboard	Description
Devices At Risk	Displays a list of device IPs in ascending order by Risk Score. For information about how the Risk Score is calculated, see Risk score calculation on page 77.
Detections By Category	Displays detections by category and attack as a bar chart.
Detections By Severity	Displays the number of active detections and the severity as a pie chart.
Detections Over Time	Displays the number of detectors, and a graph of the active detections over time.
Detections Summary	Displays the number detections as a graph by severity.
Devices	Displays the total number of devices, external and internal traffic as a percentage, and a graph of visible devices.
Investigations	Displays investigations as list by *Name*, *Status*, *Days Open*, and *Last*

Dashboard	Description
	Modified by.
Gen AI Observations	Displays generative AI usage within your organization. See Gen AI dashboard on page 28.
MITRE ATT&CK Detections	Displays the MITRE ATT&CK detections activity.
Mitre Attack	Displays the MITRE ATT&CK matrix.
Notable Detections	Displays the notable detections and descending order by number of devices affected. The *New* and *Spike* labels highlight new detections and spikes in detection activity. • *New* indicates that there were no active detections during the baseline period (defined as 30 to 7 days ago), but at least one detection has occurred in the past 7 days. • *Spike* indicates that the number of active detections in the past 7 days is more than three times higher than the baseline count.
Observations	Displays a list of the observations for the previous two weeks as a scrollable table.
Query Chart	Displays data from saved *Group By* queries created in *Investigations*. You can customize the widget by selecting a time range, choosing a chart type or table view, and assigning a custom name. To update the data, click the *Refresh* button. You can also download the displayed data as a CSV file.
Sensors	Displays the number of online and offline sensors, as well the number of errors and degraded sensors. A graph displays the *Total Traffic Captured*, as well as the number of *Events Per Second*, *Visible Devices* and *Bandwidth Usage*.
Sensors Throughput	Displays the sensors throughput as bar chart that can be downloaded.
Traffic by Type	Displays the data in the *Events* tab in the *Sensor telemetry* page. You can click the widget header to pivot to the *Sensor telemetry* page.. All the filters applied to the widget will be transferred to the Sensor Telemetry page.

Detections

FortiNDR Cloud *Detections* is an alert mechanism that notifies you when events matching a specific criteria appear in your account. Detections allow you to quickly identify and respond to suspicious or known malicious activity in your network.

The *Detections* page displays a list of *Detectors* with active *Detections* in your account.

- A *Detector* is the query and parameters used to identify activity in the network.
- A *Detection* is the actual occurrence of activity satisfying a detector.

Each row in the page displays a single detector with at least one active detection.

A Detection is created when an event matches a detector's query. Detections are identified based on both the IP address and the Sensor ID to avoid issues with overlapping IP space. A duplicate detection is not generated if a detection already exists for the IP address and sensor ID pair. Instead, the *Last Seen* timestamp is updated and the event is added to the detector's latest events. This also resets the counter for the detection's *Resolution Period* if detections for the detector are set to resolve automatically.

By default the *Detections* page displays all *Active* detectors in your account. Once all detections for a detector are resolved or muted, the detector's status is automatically updated from *Active* to *Idle*. You can create a filter to view all detectors and detections regardless of their status.

The *Detections* page displays the following information:

<table>
<tr><td>Name</td><td>The detector name.</td></tr>
<tr><td>Category</td><td>There are three categories for detectors: Attack, Potentially Unwanted Application (PUA), and Posture. Each category contains a more detailed subcategory. For more information, see Detector Categories.</td></tr>
<tr><td>Severity</td><td>The severity measures the potential impact to the confidentiality, integrity, or availability of information systems and resources if the activity is confirmed to be a true positive. Severity can be assigned to one of the following values:
<table>
<tr><th>Severity</th><th>Description</th><th>Examples</th></tr>
<tr><td>High</td><td>Significant to fair impact with the potential to spread or escalate</td><td>Malicious code execution, C2 communications, lateral movement, data exfiltration.</td></tr>
<tr><td>Moderate</td><td>Fair impact with minimal potential to spread or escalate</td><td>Activity that could indicate malicious intent, untargeted attacks with unknown success, data leakage, subversion of security or monitoring tools</td></tr>
</table></td></tr>
</table>

Severity	Description	Examples
Low	Little to no impact expected	Potentially unauthorized software, devices, or resource use, untargeted adware or spyware, compromise of a personal device or device on an untrusted network, insecure configurations

Confidence

Confidence measures how likely events matching the detector's query are indicative of the activity specified in the detector description. A detector's confidence indicates its minimum true-positive detection rate.

Confidence	Minimum True-Positive Rate
High	90%
Moderate	75%
Low	50%

FortiGuard Lab assigns a detector's initial confidence based on its performance during testing. Once deployed, detectors are monitored for changes in their true-positive detection rate, which is based on the resolution state chosen by an analyst when resolving a detection. Once a detector crosses a higher or lower threshold, it is reviewed to determine whether it should be tuned or whether the confidence should be modified.

Last Seen

The UTC date and time when the last known event tied to the detector was observed. This is useful when determining when the most recent change to a detector has occurred.

Author

The account that authored the detector.

Impacted Devices

The internal IP address in the `src.ip` or `dst.ip` fields used to generate detections. This field is configurable.

Status

By default, every detection is in an *Active* state upon creation. *Active* detections generate a notification (see Email notifications on page 126), but *Muted* detections will not. Detections remain *Active* until they are resolved manually by an analyst or automatically based on the detector's *Resolution Period*. Once resolved, their status changes to *Resolved*.

Detection State	Description
Active	When an event matching a detector is observed, a detection is generated and set to *Active* by default. A notification is triggered for Active detections.
Muted	When an event matching a detector is observed, but some aspect of it is muted. A notification is *not* triggered for

Detection State	Description
	muted detections.
Resolved	When a detection is resolved, either manually by an analyst or automatically, and is no longer Active.

Detector Categories

Category	Subcategory	Description
Attack	Infection Vector	Attacks in the initial stages before an exploit attempt has been made or malicious code has been executed. Examples include downloading a malicious executable file, navigating to a web site that is known to redirect to exploitation servers, or an attempt to authenticate to an SSH server from a malicious host.
Attack	Exploitation	Attacks in the process of exploiting known vulnerabilities such as those listed in MITRE's Common Vulnerabilities and Exposures (CVE) list. While FortiNDR Cloud may be unable to determine the success of a launched exploit, any hosts attempting exploits (that are not approved internal scanners) should be investigated for signs of compromise.
Attack	Installation	Installation of malicious software (staging) for persistence in an environment. For example, the Cobalt Strike staging tool downloading a Beacon backdoor over HTTP in order to provide persistence on a compromised host and run further post-exploitation commands.
Attack	Lateral Movement	Tools and techniques commonly used by attackers to pivot from a compromised host to other assets within the environment. Such tools may also be legitimately used by system administrators but should be investigated, especially for hosts from which this activity has not be observed before.
Attack	Command and Control	Command and control traffic between compromised hosts and attacker infrastructure.
Attack	Exfiltration	Data exfiltration from compromised assets to external entities.
Attack	Discovery	Tools and techniques commonly used by attackers to identify accesible hosts and services. Such tools may also be legitimately used by system administrators but should be investigated, especially for hosts from which this activity has not be observed before.
Attack	Impact	Malware or behavior intended to disrupt the business, such as distributed denial of service (DDoS) and ransomware attacks.
PUA	Adware	Malware characterized by its use of advertisements to generate revenue

Category	Subcategory	Description
		for the author. Adware is often installed alongside third-party applications and remains on a system as a browser add-on or self-proclaimed optimization software. Most adware is considered low risk due to its innocuous nature.
PUA	Spyware	Malware characterized by its focus on gathering device and user information without the user's knowledge. This information is usually sent back to the authors for a variety of purposes, ranging from market research to targeted monitoring. Spyware is usually installed alongside third-party applications and persists on a system as a backdoor or as software that purports to be useful. Most spyware is considered low risk due to its historical use for low-impact data collection and advertising.
PUA	Unauthorized Resource Use	Applications that utilize system resources without a user's knowledge or consent. Such applications are usually installed alongside third-party applications or as a component of malware in order to monetize a successfully compromised host (for example, via click fraud or cryptocurrency mining).
Posture	Potentially Unauthorized Software of Device	Applications or devices that circumvent organizational policies or increase the attack surface of an organization. These detectors cover various applications that may be used to bypass monitoring tools and access controls, or store sensitive information in unauthorized locations. This category also includes tools that may be legitimately used for system administration, development, or penetration testing, but are also commonly used by attackers to enumerate access and pivot within a compromised environment.
Posture	Insecure Configuration	Configurations within an environment that make it more vulnerable to exploitation or post-exploitation techniques used by attackers. Such configurations include outdated software, use of deprecated cryptographic standards, or configurations resulting in data leakage.
Posture	Anomalous Activity	Network activity that is abnormal and should be investigated to determine its cause. The activity may be malicious in nature or a misconfiguration that may or may not have security implications.

Triage detections

The *Triage detections* view is the landing page for the *Detections* tab. Use this view to review and respond to detections triggered by the detector.

To view the Triage detections page:

1. Go to *Detections* > *Triage detections*. The *Detections* page opens.
2. (Optional) Filter the detections on the page.

Search	Enter the technique ID, technique name or technique description. Detectors are filtered based on the prefix matching the selected technique ID. If Technique T1234 is entered, the detectors returned include its sub-techniques T1234.001, T1234.002, T1234.003, etc.
Severity	Select High (**H**), Medium (**M**), or Low (**L**).
Additional Filters	Click the filter icon to view additional filters.

Filter	Description
Category	Filter the detectors by category. See, Detector Categories.
Assigned to	Filter by assigned detections. See, Assigning detections on page 70.
Created By	Filter by the account that created the detector.
Sensors	This filter displays all sensors in the account. The dropdown is divided into two groups: online sensors appear at the top, while other statuses are listed below.
Technique	Filter by the technique used for the detection.
Confidence	Select High (**H**), Medium (**M**), or Low (**L**).
Detection Status	Select *All*, *Active* or *Idle*.
	All — Returns all detections the user has access to regardless of whether or not it was triggered in the current account.
	Active — Returns all active detections.
	Idle — Returns all detections that have been triggered in the current account but are not currently active.
Muted	Select *Unmuted* or *Muted*. See, Muting on page 43.
Disabled	Select *Enabled* or *Disabled*. See, Disabling detectors on page 46.

Order By	Order the detectors by *Impacted Devices*, *Muted Devices*, *Severity*, *Confidence*, *Category*, or *Last Seen*.

3. Click a detector to open the *Details* page. The following information is displayed:

Category	The attack category.
First Seen	The UTC date and time the first event associated with the detection occurred.
Last Seen	The UTC date and time of the last known event tied to the detector was observed.
Updated	The UTC date and time the detector was modified.
Resolution Method	• *Automatic*: The detection will be resolved if events containing the same host and sensor ID are not observed for the specified time period. • *Manual*: The detection will remain active until an analyst resolves the detection.
MITRE ATT&CK	The MITRE ATT&CK ID.
Primary Technique	The primary attack name and ID.
Specificity	
Behaviors	The behavior coverage.
Description	A description of the detection. You can use this description to search for detections. See, Search for detections with the detector description on page 40
Next Steps	Recommendations to resolve the detection.
Show Matching Events	Click to view the *Entity Lookup*.
Author	The detector author.
Impacted Device Field	The fields used to generate the detection. The internal IP address in the `src.ip` or `dst.ip` fields is the default.
Indicator Fields	The indicators the detector uses to generate the detection. This information is useful for identifying related activity and tracking indicators over time. Detectors can define up to five fields to extract indicators from, and each detection can store up to five unique indicators for each indicator field.
Impacted devices	The active detections for the detector. All Active defections are displayed by default. You can create a filter to view Muted or Resolved detections. See, Impacted Devices on page 41. You can use this tab to resolve detections or to search for a device by IP.

Query	This tab displays the IQL query defined for the detector. You can use a query string to create a custom detector. See, Adding custom filters to a detector query on page 41.
Events	This tab displays all of the events that have matched the detector's query. • Left-click on an entity to open the *Entity Panel*. • Right-click a field to open its menu (for example, *Search Events*, *Targeted Search* and *Copy to Clipboard*). • Hover a column header to lock, sort or arrange the columns. These events are duplicates of the original matching event. When an event matches a detector's query, a copy is created and added to the detector's list of Latest Events so the event remains associated with the detector. This list can display up to the last 1000 matching events. Events could remain in the list in perpetuity if the detector rarely fires.
Indicators	This tab displays the field value extracted from a detection's event(s) as defined by the detector. This information is useful for identifying related activity and tracking indicators over time. Detectors can define up to five fields to extract indicators from and each detection can store up to five unique indicators for each indicator field.
Detections Graph	The *Detections Graph* plots a detector's detection volume over time. If a posture-related detector fires constantly, the graph will help show whether the issue is improving or worsening over time.

Search for detections with the detector description

You can use text of the detector description to search for detections. Copy and paste the description text into and *Global Search* field and click Enter. Search results will be highlighted in the *Detection Description* column of the in the Detections section of results.

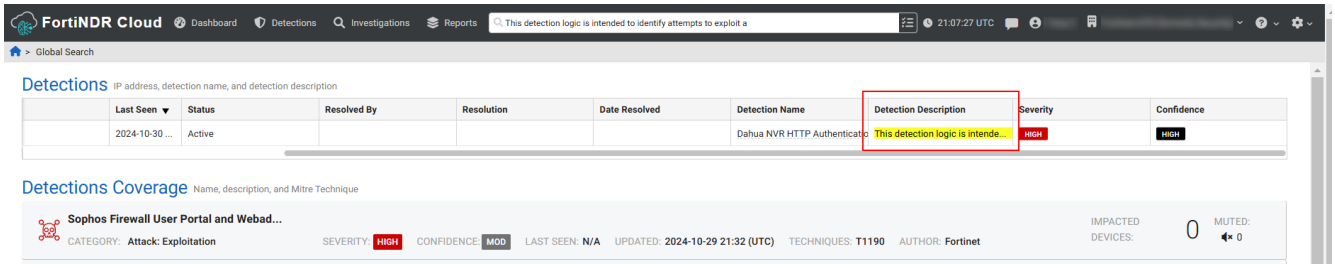

Impacted Devices

Column	Description
Device IP	The device IP address.
DHCP Hostname	The DHCP lease hostname.
Username	The device username.
Hostname	The device hostname.
MAC Address	The device MAC address
Lifetime Events	The number of events over the device lifetime. Click the link to drill down to the earliest events.
Indicators	The number of indicators of compromise. Click the link to view the indicators associated with the device IP.
First Seen	The date the event was first seen.
Last Seen	The date the event was last seen.
Created	The date the event was created.
Updated	The date the event was updated.
Sensor ID	The sensor ID. Hover over the ID to view the sensor information and annotations. Tags associated with the sensor are displayed within the column. Click the ID to open the *Sensor Details* page.
Account	The account the device belongs to.
Status	The detection status (*Active*, *Muted* or *Resolved*). See Detections on page 34.
Muted by	The user who muted the detector.
Date Muted	The date the detector was muted.
Resolved by	The user who resolved the detection.
Resolution	The resolution description.
Date Resolved	The date the detection was resolved.

Adding custom filters to a detector query

You can customize a detector authored by FortiGuard Labs by adding an additional layer of logic to a query. Filters extend the detection logic to account for differences specific to your network that muting and excluding do not account for.

To add a custom filter to a query:

1. Go to *Detections* > *Triage detections* and open the detector.
2. Click the *query* tab.
3. Click *Add a Customer Filter*.
4. In the *Custom Filter* pane, enter a valid IQL string.

The query string needs to be true in addition to FortiGuard Labs's logic for a detection to be created. Similar to excluding, no detection will be created if an event is filtered by your custom logic.

The example below excludes traffic using a custom, internally defined `UserAgent` string.

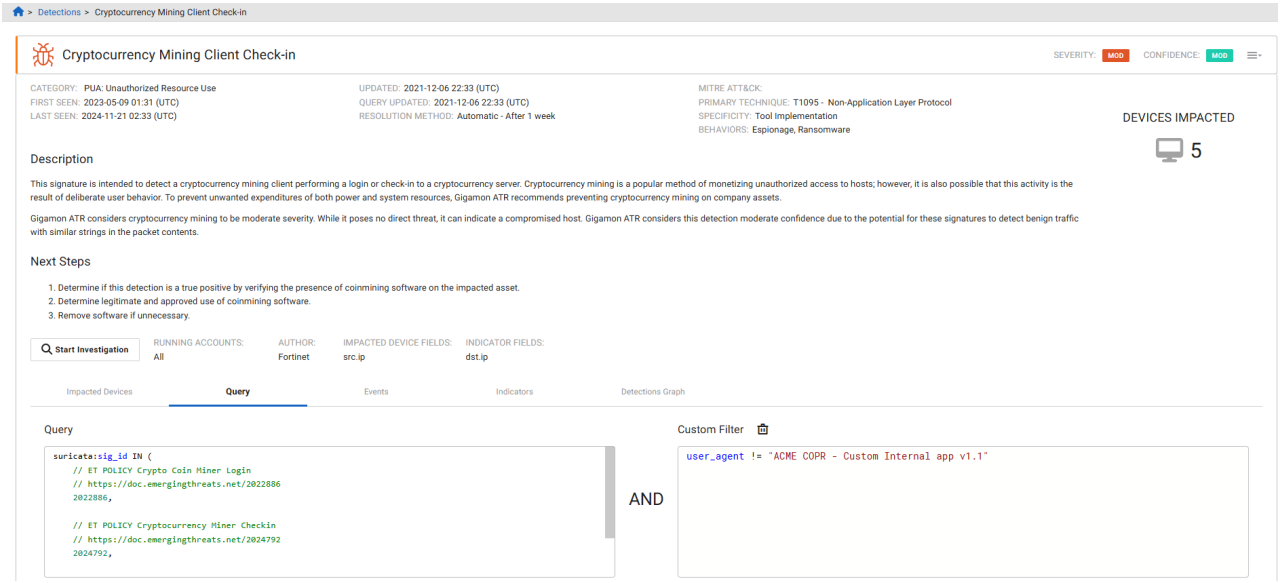

5. Click *Test Filter*.
6. Click *Save Filter* to apply your logic to the detector.

To modify a custom query, click *Update Custom Filter* or click the delete icon above the *Custom Filter* pane.

Search for a device hostname in detections

A detector query does not allow for the inclusion of a device hostname in the detector logic. However, you can use a custom filter to search for a device by its hostname. For example, if there is a particular device hostname of interest in can be incorporated into a detector by creating a custom filter as shown below.

```
http:uri.path matches ". *W/[wN][iT][nN][nN][tT]V[ss][y~][ss][t T][eE](mM]32W/.(1,6}\. [eE][XX]
     [eE].** and uri.path matches ". {0,4 0}?[\/]([ss][cC][rR][it][pP][tT][sS]|[cc][gG][iTI\-
     (bIiI1ΓnN]| [mM][s5] [aA] [dD][cC]|_[W][tT][iI]_[bB][it][nN]|\.(2})[\V/]L.[2.*
```

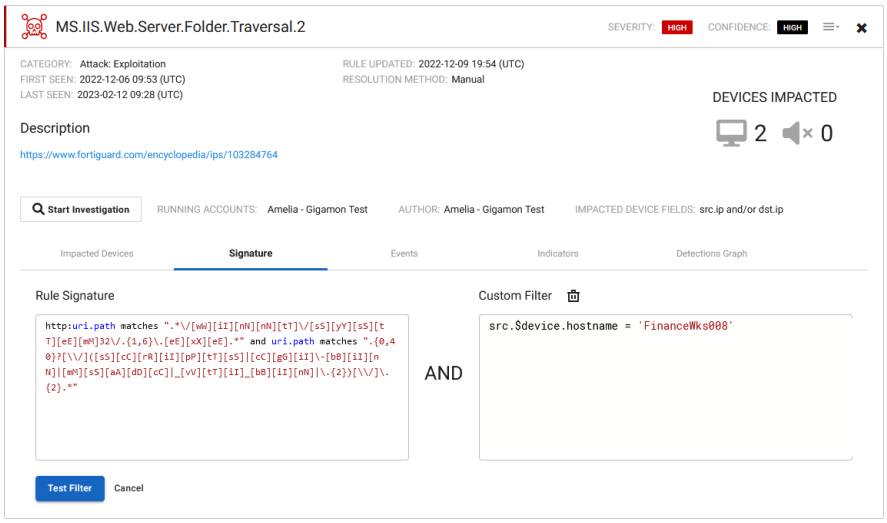

Only the `"="`, `"!="`, and `"IN"` filter conditions are supported for device hostname filters. Filter conditions such as `"LIKE"` and `"MATCH"` are unsupported.

The current Entity Tracking System only analyzes DHCP records. A custom filter leveraging a device hostname will only be as accurate as the available DHCP information.

Muting

Muting allows you to ignore authorized and expected behaviors to identify anomalies for the specific host. When a detector is muted, any related detection will have a status of *Muted*. This means a notification will not be generated for the detection. A muted detection will auto-resolve after the specified time frame or can be resolved manually.

To view all muted devices, detectors, and detections, go to the Mutes and excludes on page 132.

Mute all detectors for a device

Muting a device for all detectors is most commonly used for devices like sandboxes and vulnerability scanners, which routinely trigger detections as part of their normal operation. Since these alerts are expected, muting such devices is often one of the first steps when configuring FortiNDR Cloud.

To mute a device for all detectors:

1. Click the *Detections* tab.
2. In the toolbar, click the gear icon at the right side of the page and select *Muted Devices*. The *Muted Devices* page opens.

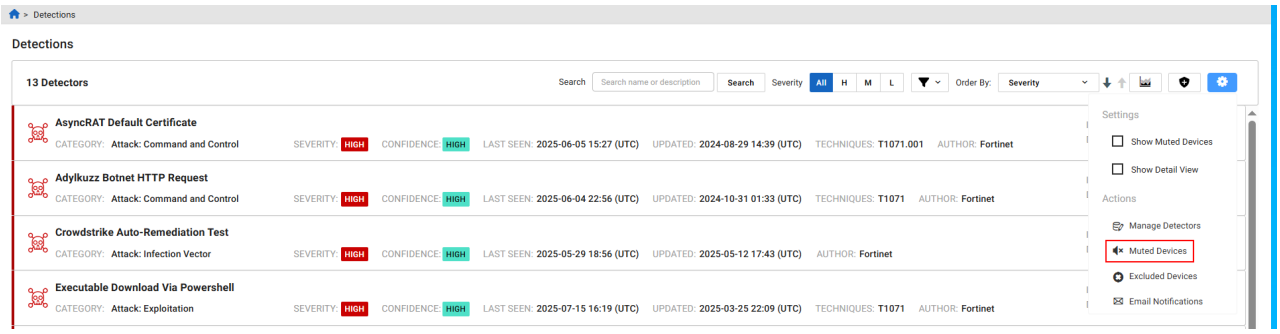

3. Click *Add New device Range*.
4. In the *Device IP or Range* field, enter an IP address or CIDR range.
5. Click *Add Device(s)*.

Mute a detector

Muting a detector will cause all its future detections to be muted, regardless of which device triggered the detector. This is commonly used for posture-aware detectors that identify approved or expected behavior.

To mute a detector:

1. Click the *Detections* tab.
2. Click the menu icon in the last column at the right-side of the page, and select *Mute Detector*.

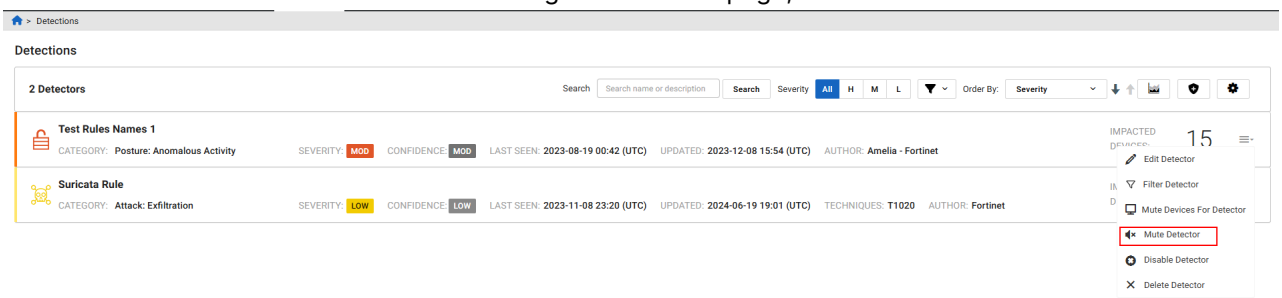

3. In the dialog that opens, enter a comment in the *Comments* field, and click *Mute Detector*.

Mute a device

You can mute a device for a detection, detector or an account. This is commonly used for suspicious behaviors from approved devices, such as remote access from an administrator workstation. Detections that contain a muted detector are appended with *Muted* in the *Status* of column of the *Detections Table*.

To mute a device:

1. Click the *Detections* tab and open a detector in the list.
2. In the *Impacted Devices* tab, select the detection that contains the device and detector.
3. Click the *Actions* menu at the right side of the page and select one of the following options.
   - *Mute Device for Detection*
   - *Mute Device for Detector*
   - *Mute Device for Account*

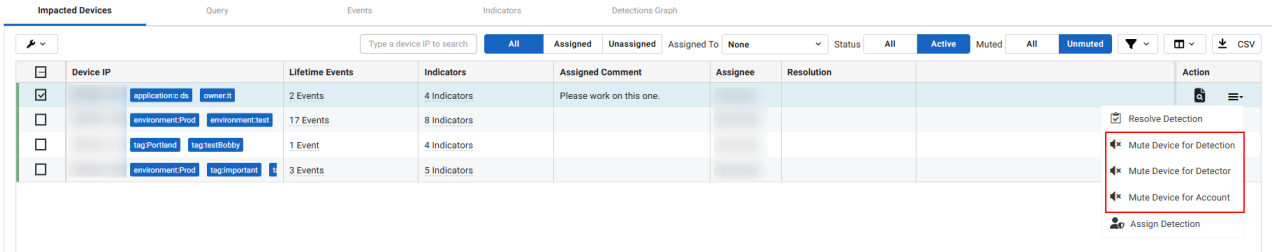

4. In the dialog that opens, enter a comment in the *Comments* field, and click *Mute Detector*.

Alliteratively, you can go to *Detections > Detections Table*. In the *Action* column, click the menu and select *Mute device for detector*.

Viewing muted devices

Option	Description
Mutes and Excludes	**1.** Click the gear icon at the top-right of the page and select *Mutes and Excludes*. **2.** Scroll down to the Muted Devices
Detections	**1.** Go to *Detections*. **2.** Click the *Settings* menu ⚙ at the top-right of the page. **3.** Under *Actions* select *Muted Devices*.
Detections Table	**1.** Go to *Detections > Detections Table*. **2.** Click the column selector ▥ and show the *Device Muted* column

Excluding devices

You can exclude a device across all detectors. This is useful in devices that are meant to perform functions that look suspicious out of context.

We recommend muting devices rather than excluding to allow for auditing and to have detections to reference if needed.

To exclude devices:

1. Click the *Detections* tab.
2. In the toolbar, click the gear icon at the right side of the page and select *Excluded Devices*.

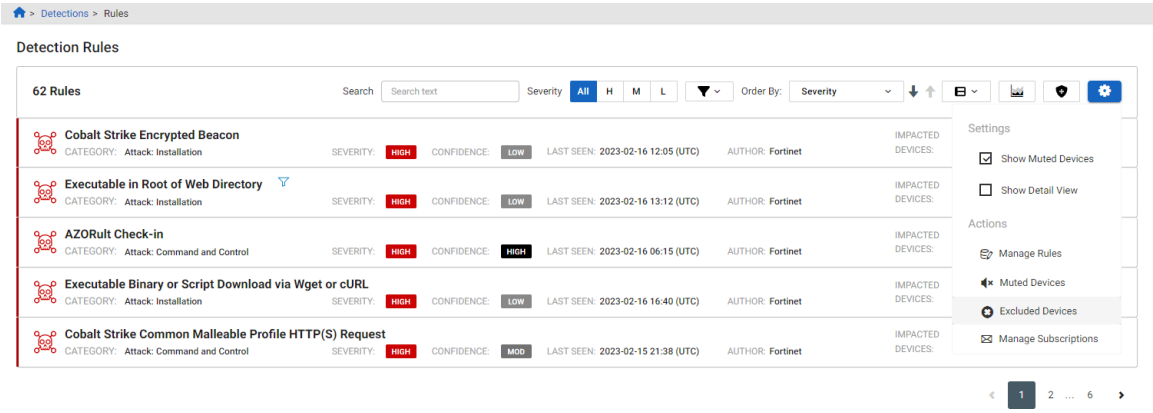

3. Click *Add New device Range*.
4. In the *Device IP or Range* field, enter an IP address or CIDR range.
5. Click *Add Devices*

Disabling detectors

Disable a detector to exclude it from matching events. Disabling detectors is useful for posture-focused detectors that detect approved behavior

To disable a detector:

1. Go to *Detections*.
2. In the toolbar, click the gear icon at the right-side of the page and select *Manage Detectors*. The *Manage My Detectors* page opens.
3. In the *Actions* column, click the menu dropdown and select *Disable Detector*. A confirmation dialog opens.
4. Click *OK*.

Resolving detections

You can resolve a detection to change its state from *Active* and remove it from the default view.

FortiGuard Labs curates detection logic over time. When the resolution ratio shows a high rate of False Positives, FortiGuard Labs will take steps to determine what changes are necessary in order to increase detector performance.

Detection resolutions are your direct feedback line to FortiGuard Labs. We recommend resolving detections to improve the quality of the detectors you see.

To resolve a detection:

1. Click the *Detections* tab and open a detector in the list.
2. In the *Impacted Devices* tab, select the detection you want to resolve.
3. Click the *Actions* menu at the right side of the page and select *Resolve Detection*. The *Resolve <IP address>* dialog opens.
4. From the *Resolution* drop down, select one of the following options.

Resolution State	Description	Example
True Positive: Mitigated	The threat was investigated and resolved, contained, or removed.	Malware was discovered on a host.
True Positive: No Action	The threat has been acknowledged, however no action was taken to resolve it.	An analyst ran a post-exploit tool for testing purposes.
False Positive	The matched events don't represent the reported activity.	A query for malware C2 instead flagged web browser traffic to a common site.
Unknown	The status or veracity of the detection is unknown.	You have no idea what you're even looking at, nor what to do with it.

5. (Optional) In the *Comments* field, enter brief description of the resolution.
6. Click *Resolve detection*.
7. (Optional) To unresolve a detection, select *Unresolve Detection* from the action menu.

Resolving a detection does not delete the detection, it is simply removes it from the default view. Detections remain in your account in perpetuity and can be viewed or pulled via the API at any time.

To view resolved deflections, click the *Filter* button in the *Impacted Devices* tab on the detector page and select *Resolved Detections*.

To bulk resolve detections:

1. Click the *Detections* tab and open a detector in the list.
2. In the *Impacted Devices* tab, click the select all box in the first column of the table. The *Bulk Resolve* icon is displayed.
3. Click *Bulk Resolve Detections*.

4. In the Impacted Devices tab, click Bulk Resolve Detections. the Resolve X Detections dialog opens.
5. From the *Resolution* drop down, select one of the following options.

Resolution State	Description	Example
True Positive: Mitigated	The threat was investigated and resolved, contained, or removed.	Malware was discovered on a host.
True Positive: No Action	The threat has been acknowledged, however no action was taken to resolve it.	An analyst ran a post-exploit tool for testing purposes.
False Positive	The matched events don't represent the reported activity.	A query for malware C2 instead flagged web browser traffic to a common site.
Unknown	The status or veracity of the detection is unknown.	You have no idea what you're even looking at, nor what to do with it.

6. (Optional) In the *Comments* field, enter brief description of the resolution.
7. Click *Resolve detections*.

Creating a detector

Create custom detectors using a unique query or from a saved query. Each account can store up to 50 detectors. If you reach this limit, an error message will appear. We recommend regularly reviewing your detectors to ensure they are still in use and deleting any that are no longer needed. To increase the detector limit for your account, please contact Customer Support.

Before you create a detector, consider using a detector filter to customize a detector created by Fortinet. detector filters save time creating a new detector and help manage the number of detectors in your account. For information, see Adding custom filters to a detector query on page 41.

To create a new detector:

1. Click the *Detections* tab.
2. In the toolbar at the top-right of the page, click the shield icon. The *Create A Detector* page opens.
3. Enter a query in the text field and click *Test Query*.
4. Resolve any errors flagged by the system.
5. Configure the detector settings and click *Save Detector*.

To create a detector from an existing query:

1. Click the *Detections* tab.
2. In the toolbar at the top-right of the page, click the shield icon. The *Create A Detector* page opens.

3. Under Detector Query, click the hyperlinked text, *select a previously run query*. The *Select a New Query* page opens.

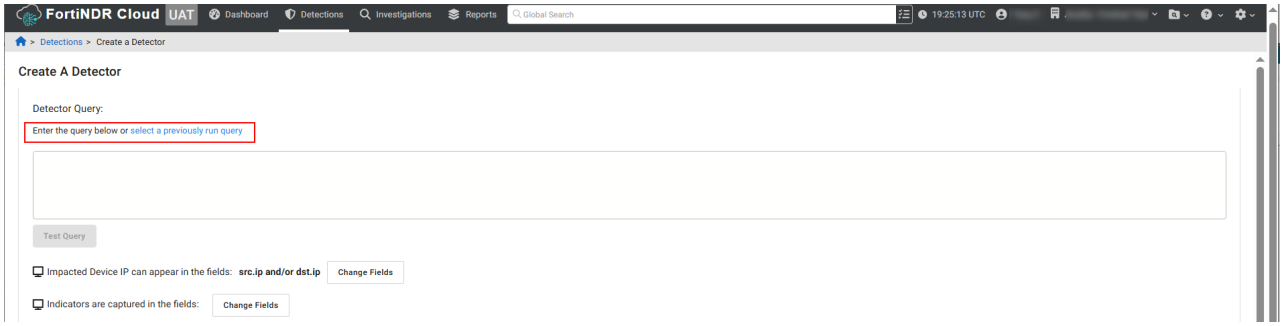

4. Select a query from the *Saved Queries* or *Query History* tab and click *Select*. The query is added to the text field.
5. If necessary, edit the query, and click *Test Query*.Resolve any errors flagged by the system.
   You do not need to test the query if you do make any edits.
6. Configure the detector settings and click *Save Detector*.

Detector settings

Impacted Device IP can appear in the fields	Click Change Fields to select the specific fields you want to use to generate a detection. By default, any internal IP address in the `src.ip` or `dst.ip` fields will be used to generate detections.
Indicators are captured in the fields	Click *Change Fields* to add or remove an Indicator Field for a detector. You can choose up to five fields.
Name	Enter a name for the detector.
Severity	Choose *High*, *Moderate* or *Low*.
Confidence	Choose *High*, *Moderate* or *Low*.
Category	Click the drop down to select a category from the list.
Primary Technique	Enter the Primary Technique ID.
Secondary Technique	Enter the Secondary Technique ID.
Run on Accounts	When creating a detector on a parent account, enable *Current account and all children account* to run the detector on the current account and child accounts. When creating a detector on a child account, select *Move to parent (Account1) and run on parent and all children accounts* to run on the detector on all accounts (current, parent and children). This option is only available to customers with parent and child accounts.

	These selections cannot be undone.
Data Sources	Enable/disable *Zeek*, *Fortinet*, *Zuricata*, or *Zscaler*.
Resolution Style	Select *Auto* or *Manual*.
Automatic Resolution Period	Select *6 hours* to *1 Month*.

Start an investigation

To start an investigation:

1. Go to *Detections* > *Triage detections*. The *Detections* page opens.
2. Click a detector to open the *Details* page.
3. Click *Start Investigation*. The *Add Query* to Investigation dialog opens.

Query Name	Enter a name for the query.
Search Query	Enter the query string.
Last 7 Days	Click to set the data range to *Last Hour*, *Last 24 Hours*, *Last 7 days*, *Last 30 days*, *Last 60 days* or last *90 days*.
Sort by timestamp	Select *Ascending* or *Descending*.
Retrieve up to	Click to set the number of rows retrieved (*100*, *500*, *1000*, or *10,000*).
Create a New Investigation	Click to create a new investigation.
Add to Existing Investigation	The *Choose Investigation* dropdown is displayed. Select an investigation from the list.
Run a Private Query	Select this option to add a query to an adhoc search.
Investigation Name	Enter a name for the new investigation.
Description	Enter a short description of the new investigation.
Choose Investigation	

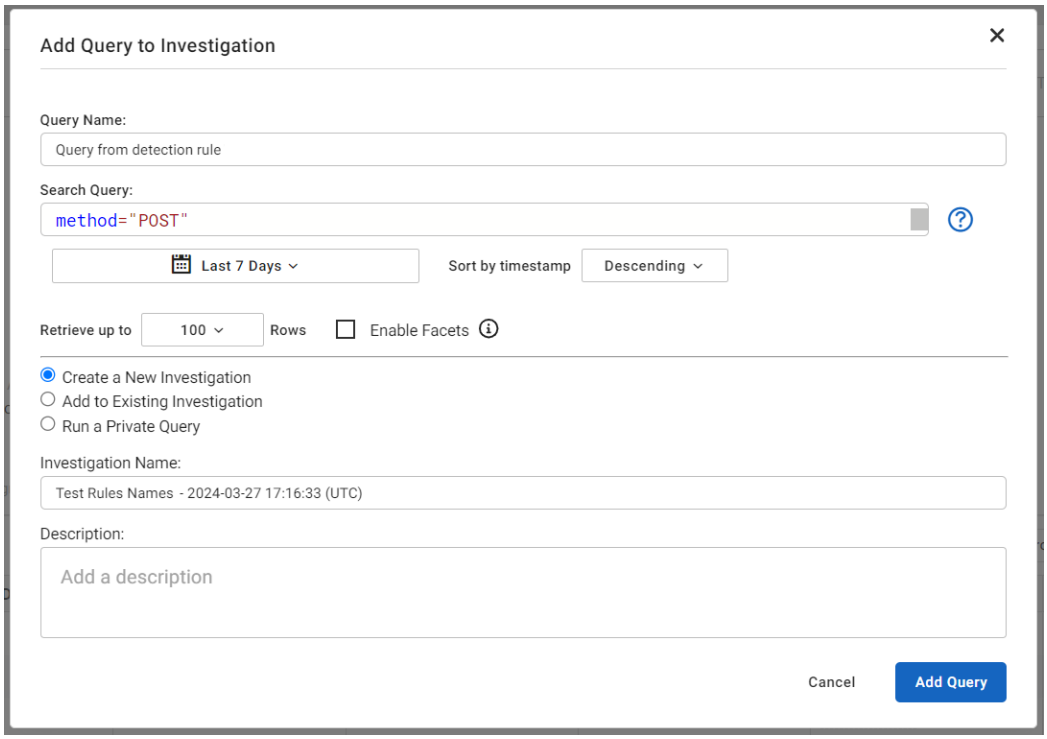

4. Click *Add Query*.

Viewing related investigations

To view related investigations.

1. Click the *Detections* tab and select a detector from the list.
2. Click *View Related Investigations*. The Investigations page opens.

Running queries in a detection

Run a queryused by the detector for a detection.

To view a query in a detector:

1. Click the *Detections* tab and open a detector in the list.
2. Click the *Events* tab.
3. In the *Timestamp* column, right-click an entry and select *Guided Queries*. The *Add Guided Query* page opens.
4. Click *Select* next to a query in the list.

Entity Panel

The *Entity Panel* provides contextual information about specific identifiers, known as *entities*, to aid in security analysis. The entities currently supported are IP addresses and domains. Entities are extracted from event data and stored in their own data store. The Entity Panel displays all contextual information collected for an entity, drawing data from both within and outside the network.

The Entity Panel supports investigations and threat hunting by consolidating data types, providing historical context, external intelligence, and relevant network activity related to the entity.

To access the Entity Panel:

- Left-click any entity (such as an IP address) anywhere in the portal.
- Click an IP address in the detector details tabs.
- Click *View Device Details* in the *Actions* menu.
- Click a device IP in the *High Risk Devices* dashboard widget.
- Click the IP label on the *Detections Device Timeline*.

The Entity Panel is organized into tabs on the right side of the page, providing focused views of the gathered data:

- **Tabs:** Display contextual and activity information, along with external threat intelligence and enrichment details.
- **Pin Option:** Users can click the pin icon to keep the Entity Panel open and visible when navigating to other pages where it is available.
- **Date Range:** The date range displayed in the Entity Panel defaults to the time range of the page from which it was opened. See, Date ranges.
- **History:** Displays up to 50 of the most recent actions performed on the device, including manual isolation, automatic isolation, and removal of isolation. History is only applicable to CrowdStrike, EDR, and FortiManager integrations. Actions performed by the Fortinet Automation Service will appear in the *Fortinet Automation Service* tab.

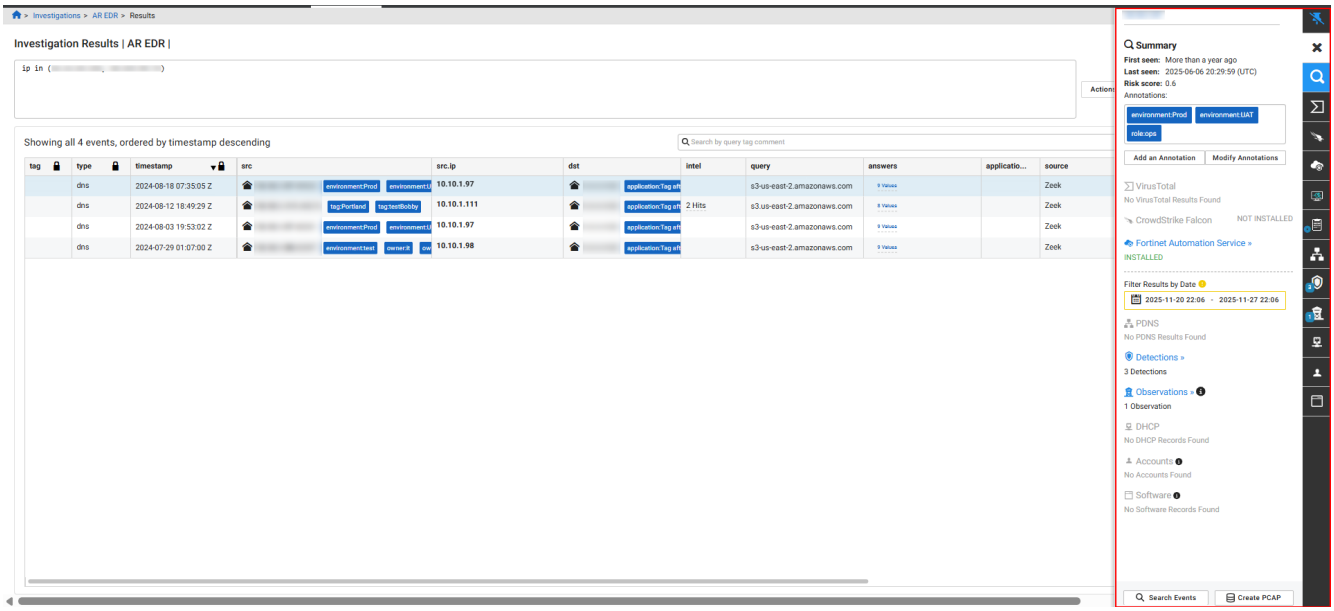

Entity information tabs

The tabs on the right side of the panel display contextual and activity information, along with external threat intelligence and enrichment details.

Tab	Description
Summary	This tab displays timestamps for when the entity was first and last seen, applied tags, and a summary of records found in other tabs. The summary also includes a button to initiate immediate response actions, such as Contain, Isolate, or Ban an endpoint.
VirusTotal	Contains details from the FortiNDR Cloud integration with VirusTotal, including information on Detected URLs, Resolved URLs (passive DNS resolution), and hashes of files (Communicating, Downloaded, or Referrer Samples) related to the entity.
Crowdstrike	This tab appears when the Crowdstrike integration is enabled. For more information see, CrowdStrike Falcon integration for FortiNDR Cloud.
Fortinet Automation Service	This tab appears when the Fortinet Automation Service is provisioned. See, Fortinet Automation Service on page 189.
FortiEDR	This tab appears when the FortiEDR integration is enabled. For more information see, FortiEDR integration for FortiNDR Cloud.
FortiManager	This tab appears when the FortiManager integration is enabled. For more information see, FortiNDR Cloud for FortiManager.
PDNS (Passive DNS)	Displays all passive DNS records observed for the entity for the entire life of the account.

Tab	Description
Detections	Lists all *Active* detections associated with the entity within the selected time range, along with their Last Seen and Created dates. This allows analysts to quickly view specific threats targeting that entity.
Observations	Displays a list of any *Behavioral Observations* associated with the entity. Behavioral Observations offer context for threat hunting and investigating network activities. Clicking an observation title here opens the *Observation Details* page. See Behavioral observations on page 66.
DHCP	Displays all DHCP records for the entity for the life of the account.
Accounts	Shows Kerberos and NTLM records observed for the entity over the past 30 days. This is useful for identifying the users of an internal asset.
Software	Lists all software associated with the entity, observed from any network protocol.
FortiGuard	Indicates if a malicious file is detected with the message *File identified as malicious* and provides a hyperlink to view attributes about the malicious file.
WHOIS	Provides registration information populated by FortiNDR Cloud WHOIS lookups.

Adding annotations and viewing malicious files

To add an annotation:

1. In the *Summary* tab click *Add an Annotation*. The *Create an annotation* dialog opens.
2. From the *Select an annotation type* drop-down, select the annotation type.
3. In the *Enter an annotation name* field, enter a name for the annotation.
4. In the *Enter a description* field, enter the annotation.
5. Click *Save*. The annotation is added to the *Summary* tab.

To modify annotations:

1. In the Entity Panel, click *Modify Annotations*. The *Manage Annotations for <IP_address>* dialog opens.
2. (Optional) In the search field, enter an annotation name.
3. Select or deselect an annotation and click *Update*.

To view malicious files with FortiGuard:

1. In the investigation results, click the link in the *File* column.
2. Click a link in the *Files* dialog.
3. The *FortiGuard* area displays the *File identified as malicious* flag.

Date ranges

Keep the following considerations in mind when view viewing results with the date range picker.

Summary tab	• The date range picker is displayed In the *Summary* tab. The results in each section above the dashed line (*Detections*, *DHCP*, *Account* and *Software*)is captured within this date range. The information below the dashed line is independent from this date range. • Sections in the Summary tab that use the date picker (such as DHCP) will also display the date picker in the corresponding tab. • The date range picker in any tab is global. If you change the start and end date in one tab it will change the date range everywhere in the panel.
Date out of range	• The *Account* and *Software* tabs only display results for last 90 days. If the date picker end date exceeds 90 days, *Date out of range* is displayed.

Default time range	• The date range on Entity Panel defaults to the time range based on the page the panel is opened in.
	  • The time range in the Entity Panel matches range when opened from the following pages:
	  • Entity Lookup
	  • Visualizer
	  • Detection Table
	  • Sensor Visibility
	  • Investigate Results
	  • Adhoc Search
	  • Observation Detail
	  • Detections is default to last 7 days when opened from the following pages:
	  • Detection page
	  • Detection-Indicator page
	  • Detection-Triage Page

Detections visualizer

Go to *Detections* > *Detections Visualizer* to view detections data from existing APIs in a graphical interface. You can use the visualizer to view the relationship between the detectors and devices, inspect detectors and impacted device details, and navigate to the node view from the list of impacted nodes.

The visualizer will initially display all active, unmuted detections over the past 14 days in graphical form with nodes representing impacted devices and detectors.

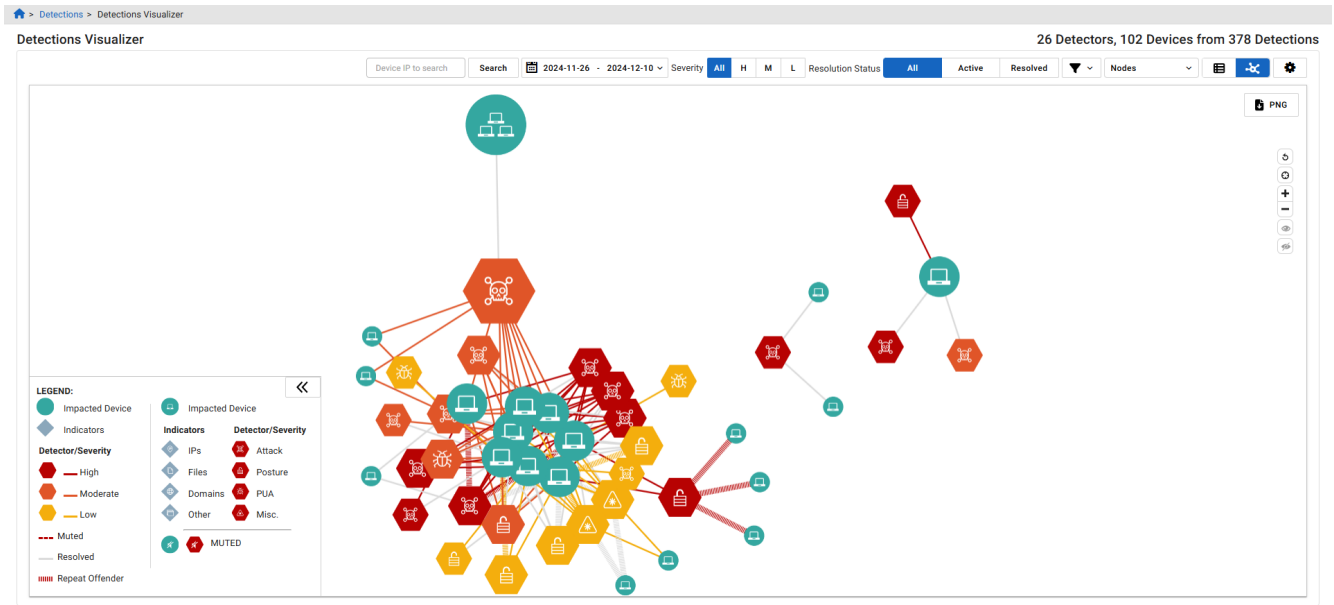

Nodes

You can hover over the nodes in the Visualizer to view summary information about a detector, device, indicator or connector line. Click a node to open the *Quick View* panel on the right side of the page. Right-click a node to open a context menu.

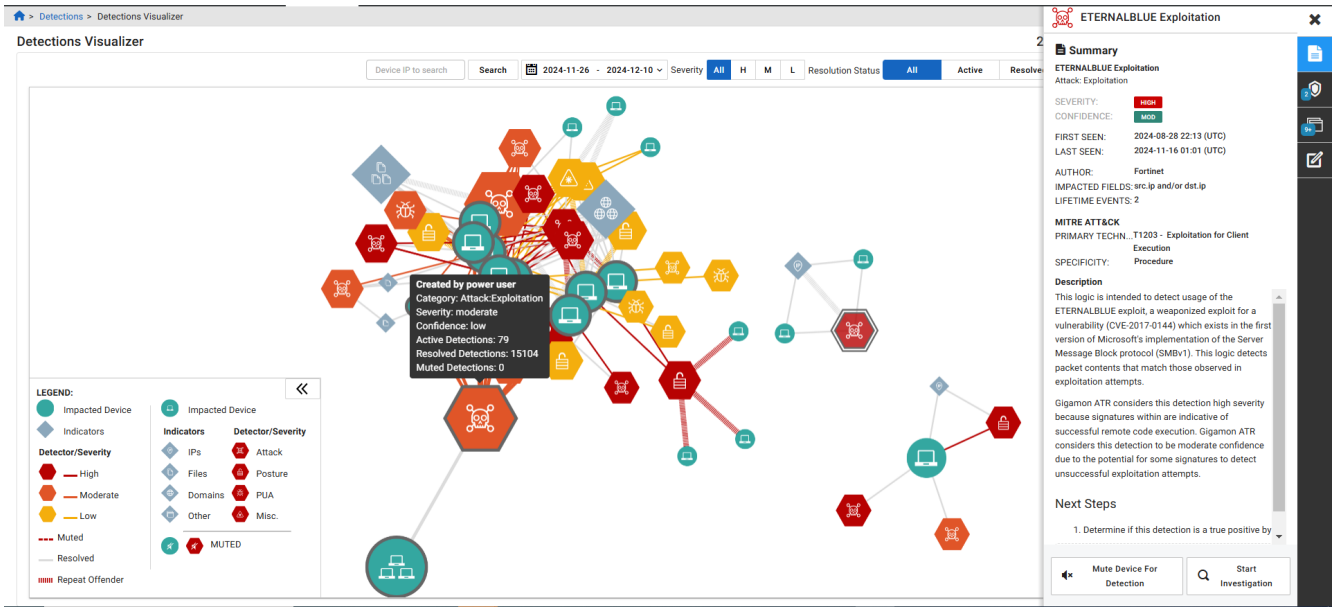

Detector nodes	Hover over a detector node to view related information about the detection such as the detector's *Category*, *Severity*, *Confidence* rating as well as the number of *Active* and *Resolved Detections*. The detector and its impacted devices are also highlighted.

Device nodes	Hover over a device node, to view the device IP address. If you hover over a device group, the list of IP addresses is shown. The device group and related detections will be highlighted. Right-click a device node to show/hide the label or the node, add an annotation, or mute the device
Indicator node	Hover over an indicator node to view the indicator and to highlight related detections and devices. Right-click an Indicator node to show/hide the label or the node, or add an annotation.
Connector lines	Hover over the connector lines to view summary information pertaining to what the line connects, such as the indicators, device IPs, and/or detections. Related devices, detections, or indicators will be highlighted. Right-click a connector line to resolve the detection or mute the device for that detector. If any node is a group or can be grouped, you will have an option to *Expand* (ungroup) or *Collapse* (regroup) the set of nodes.
Quick views	Click a node in the Visualizer to open the *Quick View* panel at the right side of the screen. Quick Views display summary information as well as a series of detail-view options and actions. The available options and actions will vary depending on the type of node selected. **Summary** — Provides a summary of the detection and corresponding devices along with options to access further details: **Software** — Displays the *Version*, *Events*, *First Seen* and *Last Seen* for the software detected on the device. **Indicators** — Displays the Indicators list. **Accounts** — Displays the Account, User, First Seen, Last Seen and Service detected on the device. **DHCP** — Displays the Dynamic Host Configuration Protocol. **Detections** — Shows a list of detections, each citing the date and time it was last seen and the impacted account; • Click an item to open the detector view • Click the options drop-down on an item to resolve the detection or mute the device for the specified detector or account **PDNS** — Displays the Passive DNS/ **Query** — Displays the query. **Virus Total** — Displays the total number of viruses detected. **WHOIS** — Provides registered domain information.

Filtering the Visualizer

Use the filters at the top of the visualizer to change the content displayed in the canvas. Some filter options are static, others are dynamic based on the criteria selected elsewhere. When you modify the filter, the graph will be redrawn per the selected options. The Visualizer can retrieve up to 10,000 detections from the API regardless of the filter criteria.

Use the *Nodes* filter to select the types of nodes to display. There are three types of nodes:

- Indicators
- Impacted Devices
- Detectors

When the *Indicators* option is selected, groups of indicators and impacted devices related to the same detector may be clustered together on the graph. While any combination can be selected, omitting *Detection Name* will usually result in a disjointed graph.

Action buttons

PNG	Export the current graph as a PNG file.
	Reset the graph (resets all filters, reloads data, and generates a new graph).
	Recenter the graph (fits all existing data in the screen).
+ −	Zoom in or out.
	Reveal hidden nodes. This option is available after one or more nodes have been hidden. To hide a node, right-click it and select *Hide node*.
	Hide hidden nodes. This option is available after one or more nodes have been hidden. to hide a node, right-click on it, and click Hide node.

Detections device timeline

Go to *Detections* > *Detections Device Timeline* to view all detections sorted by device risk score.

A solid background color in each bar on the chart represents a detection category, as indicated in the legend at the bottom of the page. If a bar is striped, it means all detections within that range have been resolved. A single bar does not correspond to one detection; instead, it may represent multiple detections that occurred within the same time range.

Hover over a bar in the chart to view details about the detection. Click the *Detection Context* button to view the detections and observations related to this IP on the *Detection Context* page. Click the *Detection Detail* button to quickly navigate to the detection detail page for the selected detector.

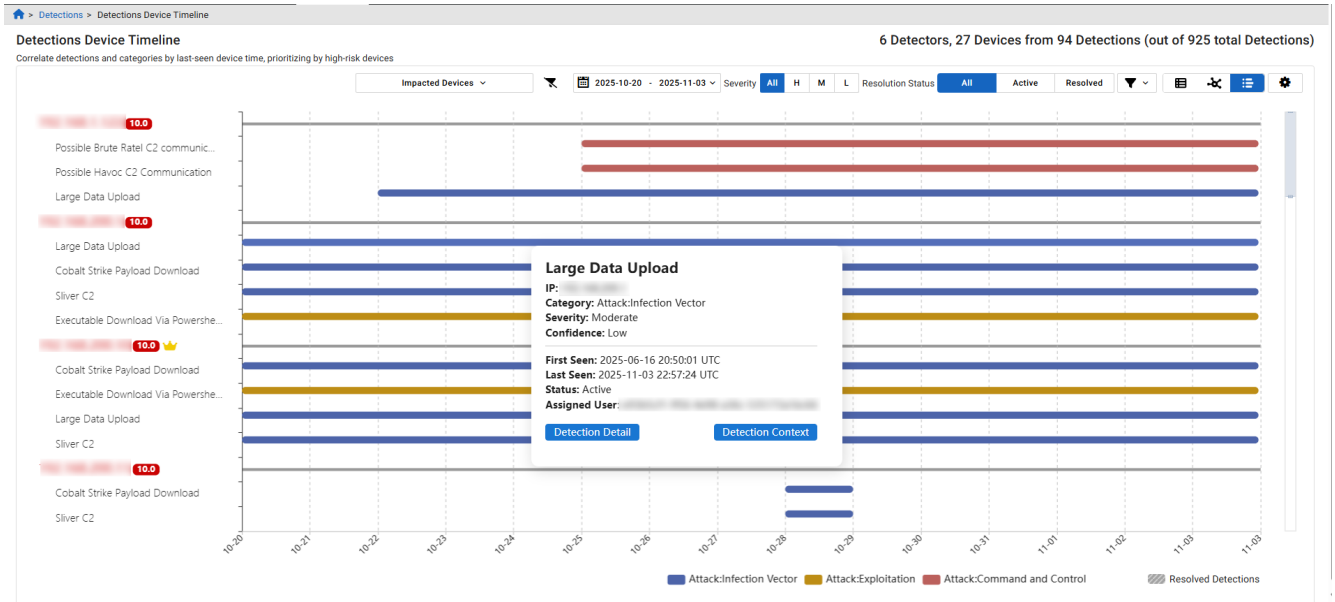

Hover over the line next to the IP label to view its risk score. Any annotations related to the IP will be displayed here. The risk score is also displayed next to the IP label.

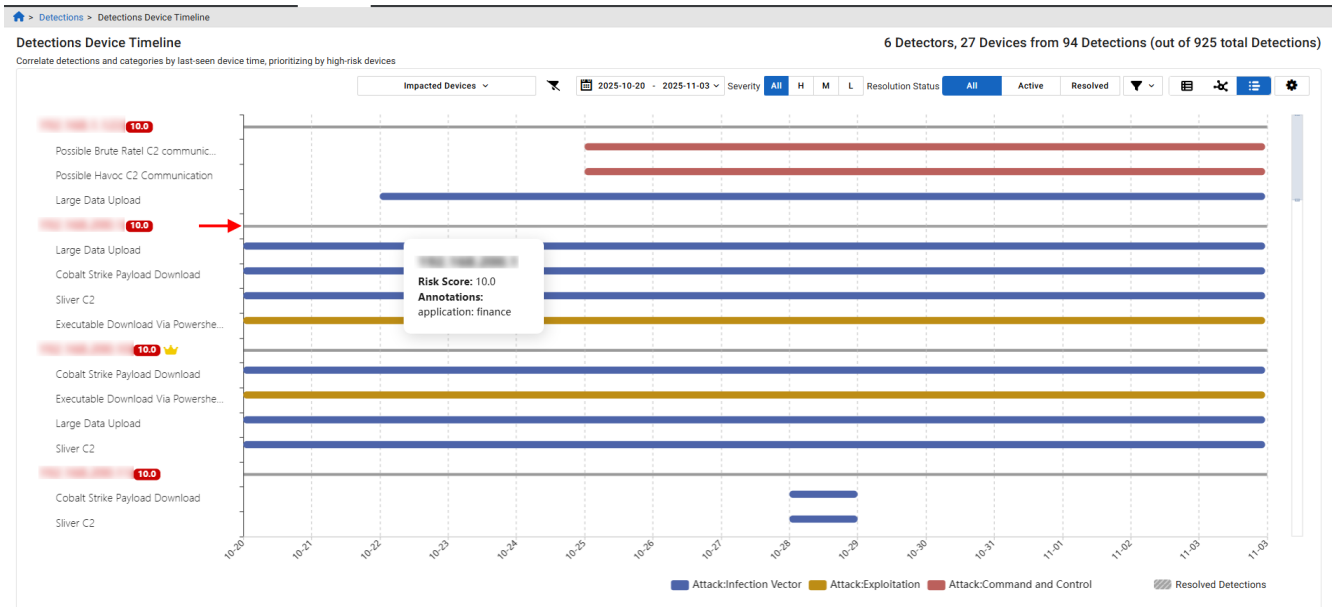

Click the IP label to open the *Entity Panel*. Right-click the IP label to open the context menu.

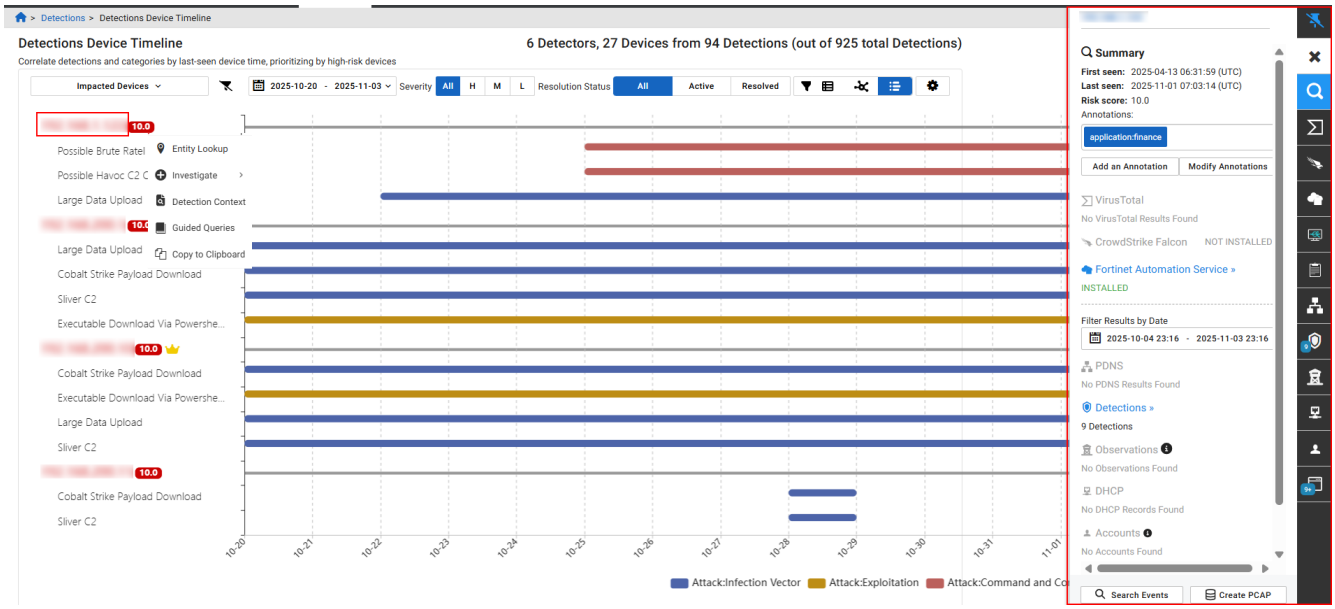

You can filter the view to hide detections that have no associated events during the selected time range. Use the toggles on the right side of the page to switch between the *Detections Table* and *Detections Visualizer* views. Both views also support the *Detections Device Timeline* toggle. To filter the table by a specific detector, click its name below the IP label.

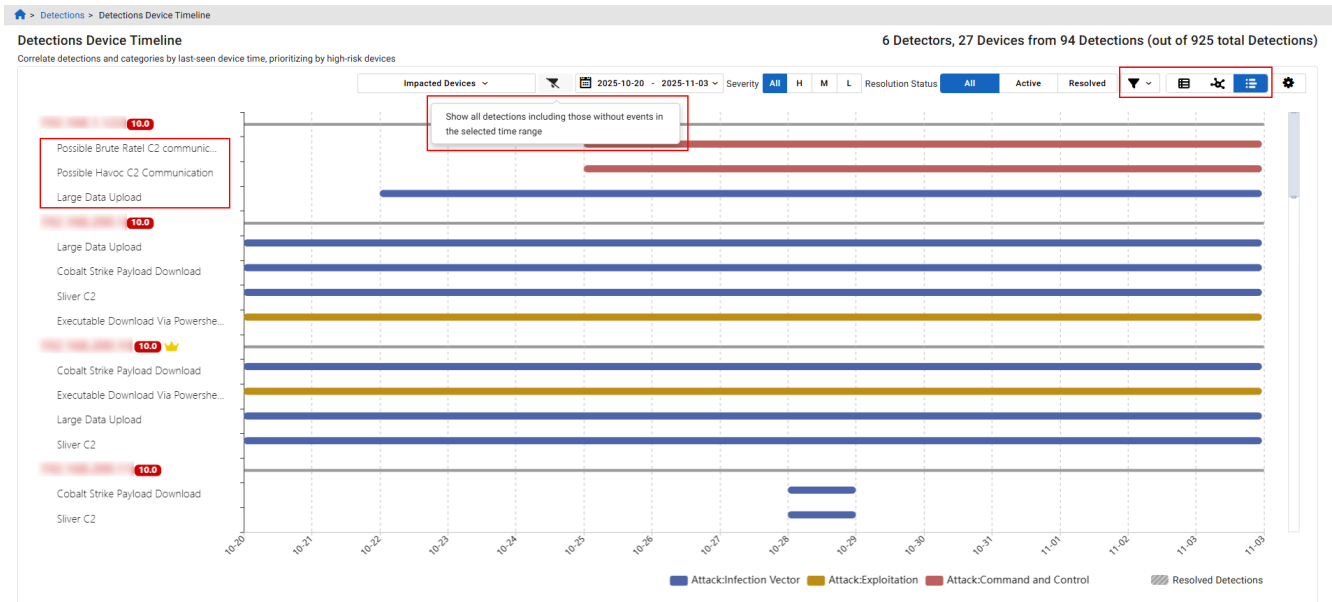

The *Detections Device Timeline* is available as a dedicated dashboard widget. By default, it displays the top five IPs with the highest risk scores from the past seven days. These settings are customizable.

Detections table

The *Detections Table* is where you can view all detections. Whereas the *Triage Detections* and *Detections Triage* views show detections by detector or device, the *Detections Table* shows detections by detector and device over time. By default, the table displays detections for the last two weeks. A color-coded bar at the left side of the table indicates active and resolved detections. A green bar indicates an active detection. A red bar indicates a resolved detection.

To access the Detections Table:

- Go to *Detections* > *Detections Table*.
- On the *Dashboard*:
  - In the *MITRE ATT&CK* widget, click a bar in the chart.
  - In the *Resolved Detections* widget, click *Total* or click a data point in the chart.

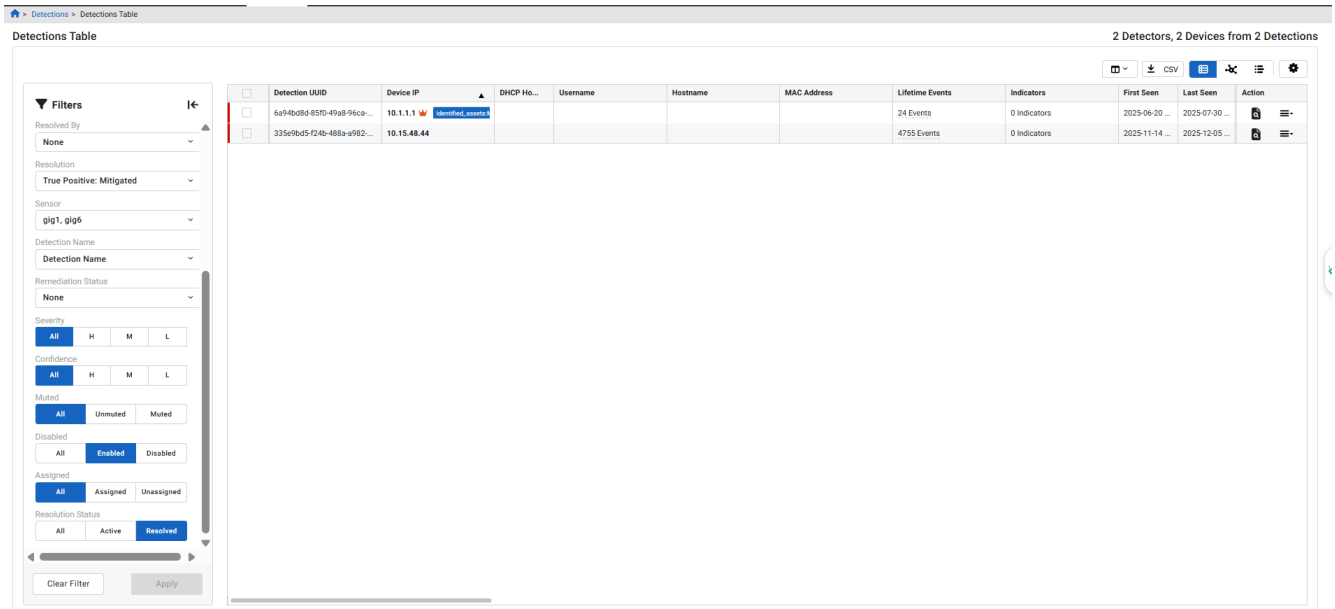

Filtering events

By default, the *Detections Table* displays detections by all severities and detection statuses for the previous two weeks ending on the current date. You can use any column header to sort the detections. The Filters pane on left side of the page allows you to view detections for a specific IP, refine the list by *Severity* and *Detection Status*. You can also toggle between table and graph view.

Impacted Devices	Click the dropdown to view the list of impacted devices. Use the search field to enter an IP address to locate a specific device. You can also select one or more devices from the list to filter the view .	
Time range	Click to open the date picker. Use the calender to set the start and end date or select an option from the *Quick Ranges* (*Last Hour* to *Last 90 days*). Click the *Date Range Type* dropdown to display detections by *Active Date*, *Creation Date*, and *Resolution Date*. The date displayed in the date picker will mirror the dates in the *Entity Panel*.	
Severity	Select High (**H**), Medium (**M**), or Low (**L**).	
Detection Status	**All**	Detections that were active during time range and are still active or resolved now. For example, a detection that was active on May 5 and resolved on May 10 is counted as *ALL*.
	Active	Detections that were active during time range and are still active.
	Resolved	Detections that were active during time range and are resolved now.

Additional filters	**Category**	Select a category from the list. See, Detections > *Detector Categories on page 36*.
	Created By	Select and account that created the detector from the list.
	MITRE ATT&CK	Select the detection by behavior from the list. See, MITRE ATT&CK on page 22.
	Assigned to	Select a user assigned to a detection.
	Resolved by	Select a user from the list.
	Resolution	Select *All*, *True Positive: Mitigated*, *True Positive: No Action*, *False Positive*, or *Unknown*.
	Sensor	Select one or more sensors from the list.
	Detection Name	Select a parameter used for the detection from the list.
	Confidence	Select *All*, High (*H*), Medium (*M*), or Low (*L*).
	Muted	Select *All*, *Unmuted* or *Muted*. See, Muting on page 43.
	Disabled	Select *All*, *Enabled* or *Disabled*. See, Disabling detectors on page 46.
	Assigned	Select *All*, *Assigned*, or *Unassigned*.
Columns selectors	**Individual Columns**	Select one of the following options: • Show all columns • Hide All Columns • Reset to default • Select columns to show or hide in the table.
	Column Profiles	Select one of the following options: • Click a profile in the list to view the layout. • Save the profile • Create a new profile. For more information, see Creating column profiles on page 76
CSV	Click to export the list as a CSV file.	
Table View	Click for table view (default).	
Graph View	Click to open the Visualizer.	
Action	Select one of the following options: • Create Detectors	

	• Manage Detectors • Muted Devices • Excluded devices • Manage Subscriptions Click the *Detections Context* icon to view the detection in a timeline along with its behavioral observations. See Detections context on page 64.

Identified Assets

A crown icon appears only on assets annotated by FortiGuard ATR. It is color-coded to indicate severity levels:

- Red for high risk
- Orange for moderate risk
- Yellow for low risk

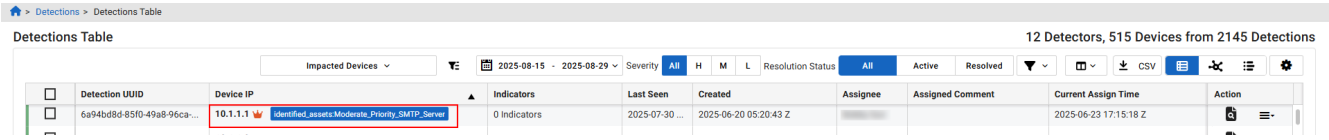

Detections context

The *Detections Context* page allows you to view detections and observations for a device within a specified time range, and provides detailed insights that includes a timeline, detections, and behavioral observations tables. You can use this page to filter, mute, or exclude devices, and navigate to detailed information pages.

The device timeline only supports detections that are less than a year old.

You can pivot to the *Detection Context* page from any page that displays an IP address, this includes:

- *Detections Table*:
  - Right-click an IP that was last seen within the year and select *Detections Context*.
  - Right-click the *Indicators* column.
  - Click the *Detections Context* icon in the *Actions* column.
  - Click the *Actions* menu in the *Entity Panel* and select *Detections Context*.
- The *Events table* > *Investigation* results page. Note that the page will not display a selected detection because you are pivoting from an event.
- The *Private Search* page.
- The *Triage Detection* page > *Events* tab.
- *Detections* details > *Lifetime Events* column.
- The *Behavioral Observations* details page

- The *Aggregation* table including the table in a report. When you pivot from the *Aggregation* table in a report, the *Detection Context* page will always show the last 90 days.
- The *Entity lookup* table. This includes the *Entity Lookup* table in *Global Search* results.
- The *Manage Annotations* page. This is limited to valid IPs for the last 90 days.
- The *Entity Panel*. You can pivot to the *Detection Context* page when the Entity Panel title is an IP address.
- Detections Table > *Indicators* column.

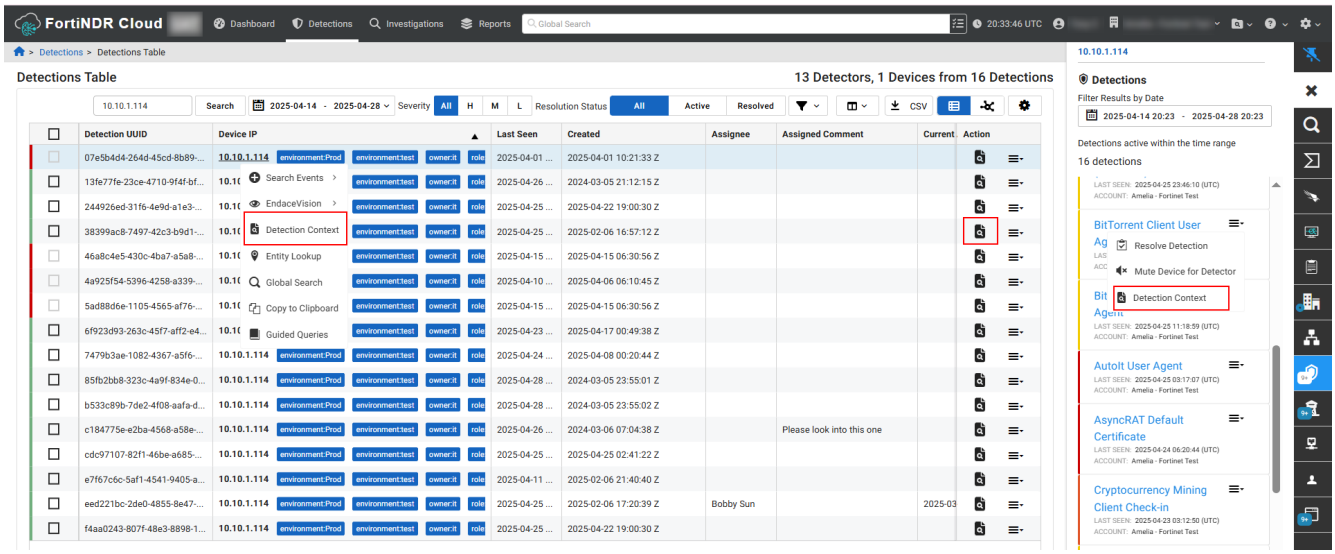

Detection context page

The *Detection Context* page displays the detections and observations timeline, as well as *Detections* and *Behavioral Observations* tables. The tables are sorted by *Last Seen* in descending order. The *Detection Context* page will display a message indicating that there are no detections or observations when none are present.

The detection you pivoted from in the *Detections table* will appear as the *Selected Detection* in the center of the timeline and display details about the detection. The timeline is sorted by *Last Seen* in ascending order. To change the *Selected Detection*, click a row in the *Detections* table. To change the selection to an observation, click a row in the *Behavioral Observations* table. You can also use the scroll bar to navigate back and forth in the timeline.

To pivot to the *Detections* or *Behavioral Observations* pages, click the *Detection Name* or observation *Title* in the table, or click a tile in the timeline.

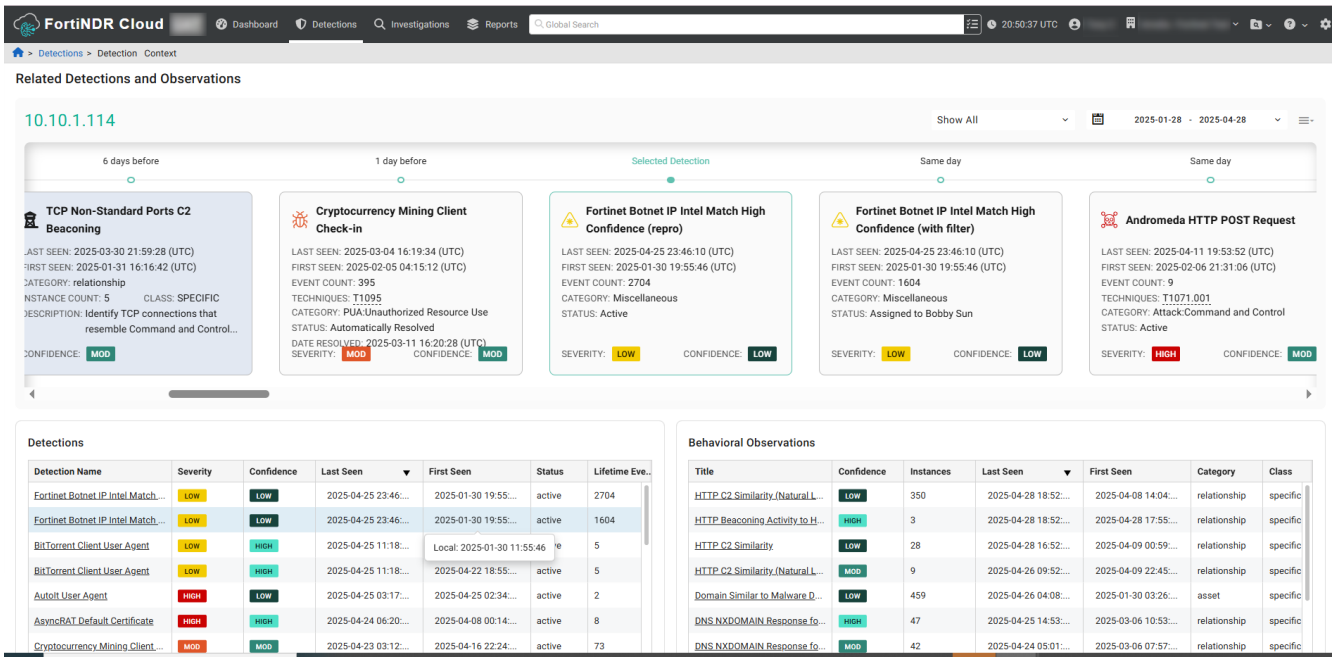

To view the *Entity Panel* for the device, click the IP address at the top-left side of the page or click the *Actions* menu next to the date picker and select *View Device Details*. You can use this menu to *Mute Device for Account*, *Exclude Device* and copy the device *Permalink*.

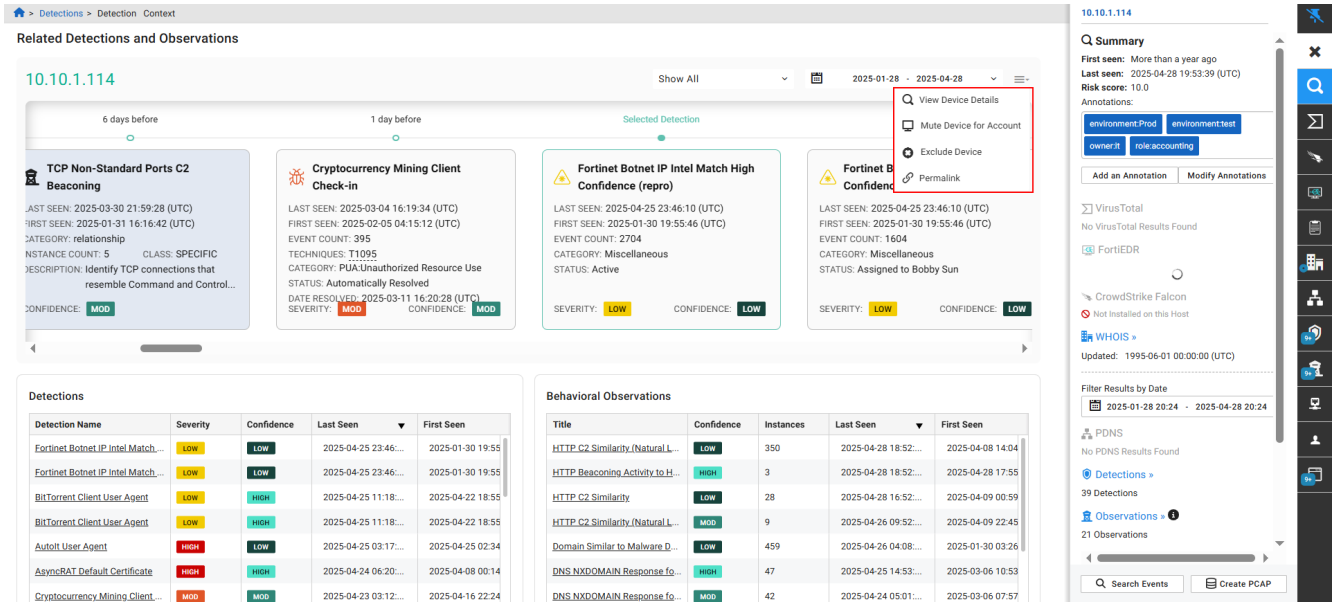

Behavioral observations

A *Behavioral Observation* is an output from an expert system or machine learning-based model that considers one or more event types and historical events. These observations are produced by analyzing threat actors'

behaviors, profiling various aspects to identify unknown malicious activity. Not every observation is malicious on its own, but those deemed detection-worthy will have detections created by the Fortinet team, typically for high and some moderate-level observations.

FortiNDR Cloud's power comes from combining detections and observations, which can be viewed in various sections like the *Entity Panel* and *Observations* page.

How Behavioral Observations are different from Detections:

Behavioral Observations	Detections
• Non-malicious observations provide context for threat hunting, investigations, and detection triage. • Observations do not have severity levels. • Observations cannot be assigned or resolved in workflows. • Observations cannot be muted	• Suspicious or malicious behavior is usually flagged as detections. • Detections can be based on single network events, Suricata events or observations. • Detections can be assigned or resolved in work flows. • Detections can be muted.

Behavioral Observations page

The *Behavioral Observations* page shows observations for a selected time range and filters. By default, the page shows observations for the previous two weeks and all confidence levels. This is also the landing page for the *Behavioral Observations* widget in the default *Dashboard*.

You can use the search field to find observations that contain instances of a specific IP address, *Observation UUID* or text in the *Observation Title* and *Description* columns. Use the date picker to create a custom time frame. Behavioral Observations can be retrieved for up to the last 90 days.

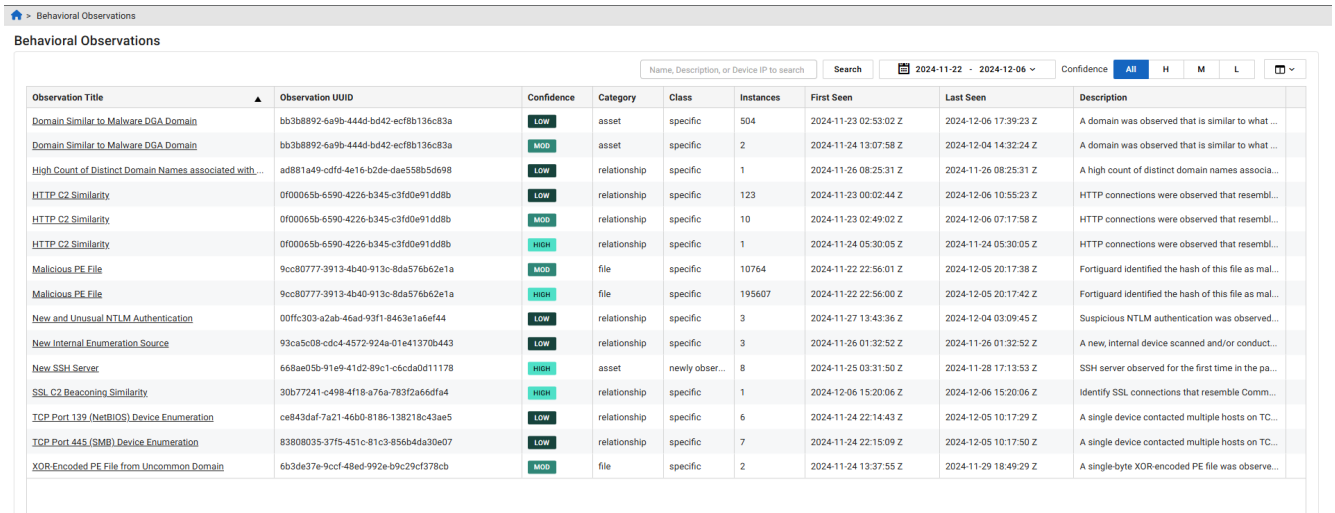

Observation Title	Observation UUID	Confidence	Category	Class	Instances	First Seen	Last Seen	Description
Domain Similar to Malware DGA Domain	bb3b8892-6a9b-444d-bd42-ecf8b136c83a	LOW	asset	specific	504	2024-11-23 02:53:02 Z	2024-12-06 17:39:23 Z	A domain was observed that is similar to what ...
Domain Similar to Malware DGA Domain	bb3b8892-6a9b-444d-bd42-ecf8b136c83a	MOD	asset	specific	2	2024-11-24 13:07:58 Z	2024-12-04 14:32:24 Z	A domain was observed that is similar to what ...
High Count of Distinct Domain Names associated with ...	ad881a49-cdfd-4e16-b2de-dae558b5d698	LOW	relationship	specific	1	2024-11-26 08:25:31 Z	2024-11-26 08:25:31 Z	A high count of distinct domain names associa...
HTTP C2 Similarity	0f00065b-6590-4226-b345-c3fd0e91dd8b	LOW	relationship	specific	123	2024-11-23 00:02:44 Z	2024-12-06 10:55:23 Z	HTTP connections were observed that resembl...
HTTP C2 Similarity	0f00065b-6590-4226-b345-c3fd0e91dd8b	MOD	relationship	specific	10	2024-11-23 02:49:02 Z	2024-12-06 07:17:58 Z	HTTP connections were observed that resembl...
HTTP C2 Similarity	0f00065b-6590-4226-b345-c3fd0e91dd8b	HIGH	relationship	specific	1	2024-11-24 05:30:05 Z	2024-11-24 05:30:05 Z	HTTP connections were observed that resembl...
Malicious PE File	9cc80777-3913-4b40-913c-8da576b62e1a	MOD	file	specific	10764	2024-11-22 22:56:01 Z	2024-12-05 20:17:38 Z	Fortiguard identified the hash of this file as mal...
Malicious PE File	9cc80777-3913-4b40-913c-8da576b62e1a	HIGH	file	specific	195607	2024-11-22 22:56:00 Z	2024-12-05 20:17:42 Z	Fortiguard identified the hash of this file as mal...
New and Unusual NTLM Authentication	00ffc303-a2ab-46ad-93f1-8463e1a6ef44	LOW	relationship	specific	3	2024-11-27 13:43:36 Z	2024-12-04 03:09:45 Z	Suspicious NTLM authentication was observed...
New Internal Enumeration Source	93ca5c08-cdc4-4572-924a-01e41370b443	LOW	relationship	specific	3	2024-11-26 01:32:52 Z	2024-11-26 01:32:52 Z	A new, internal device scanned and/or conduct...
New SSH Server	668ae05b-91e9-41d2-89c1-c6cda0d11178	HIGH	asset	newly obser...	8	2024-11-25 03:31:50 Z	2024-11-28 17:13:53 Z	SSH server observed for the first time in the pa...
SSL C2 Beaconing Similarity	30b77241-c498-4f18-a76a-783f2a66dfa4	HIGH	relationship	specific	1	2024-12-06 15:20:06 Z	2024-12-06 15:20:06 Z	Identify SSL connections that resemble Comm...
TCP Port 139 (NetBIOS) Device Enumeration	ce843daf-7a21-46b0-8186-138218c43ae5	LOW	relationship	specific	6	2024-11-24 22:14:43 Z	2024-12-05 10:17:29 Z	A single device contacted multiple hosts on TC...
TCP Port 445 (SMB) Device Enumeration	83808035-37f5-451c-81c3-856b4da30e07	LOW	relationship	specific	7	2024-11-24 22:15:09 Z	2024-12-05 10:17:50 Z	A single device contacted multiple hosts on TC...
XOR-Encoded PE File from Uncommon Domain	6b3de37e-9ccf-48ed-992e-b9c29cf378cb	MOD	file	specific	2	2024-11-24 13:37:55 Z	2024-11-29 18:49:29 Z	A single-byte XOR-encoded PE file was observe...

Working with Behavioral Observations

Behavioral Observations can be used in threat hunting and as additional evidence for analyzing network activities. They can be viewed at the device level within the Entity Panel. You can use Behavioral Observations to create custom detectors and as evidence in IQL to initiate investigations.

Behavioral Observations Widget

When you log into the FortiNDR Cloud Portal, the *Default Dashboard* displays the *Behavioral Observations* widget. This widget shows a list of the *Behavioral Observations* for the previous two weeks. Click an *Observation Title* to pivot to the *Behavioral Observation Details* page.

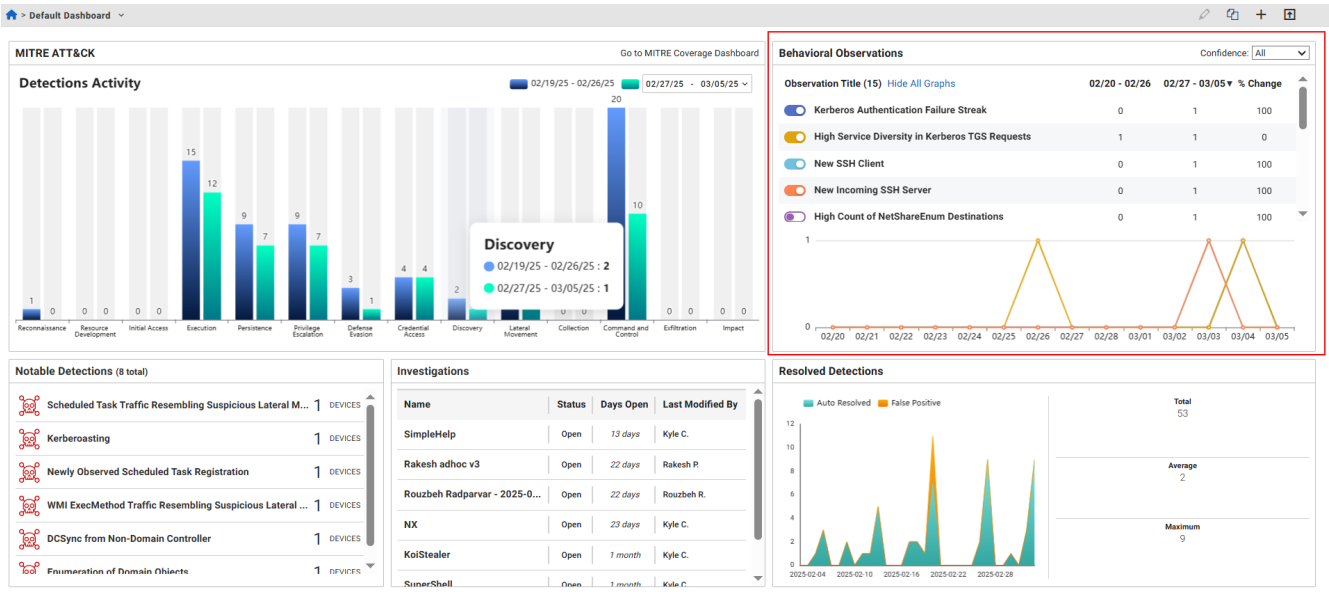

Behavioral Observation Details

The observation class, category and description appear at the top-left of the page. You can view Behavioral Observations for an individual entity in the Entity Panel by clicking the IP in the *Src* column.

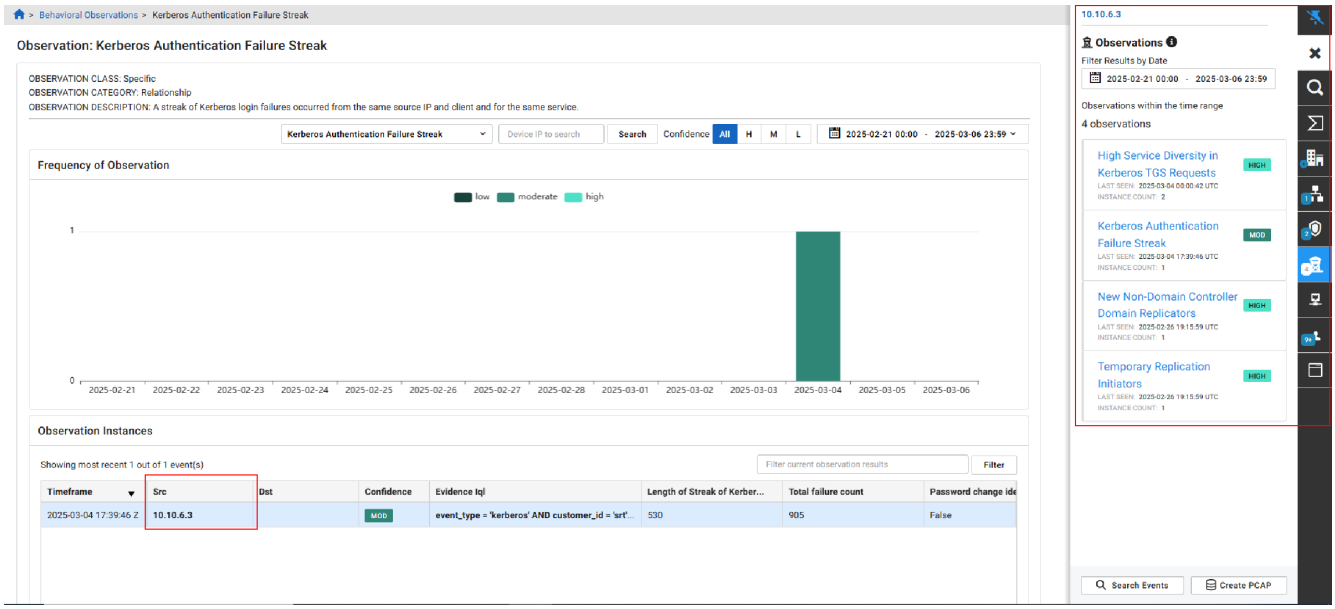

Right-click the Source IP to *Search Events* by field, or launch an *Entity Lookup*, *Global Search* or *Guided Query*. Click the *CSV* button in the *Observation Instances* section, to download the data in the page as CSV file.

You can use queries based on the observation details to create a new detector. For more information, see Creating a detector on page 48.

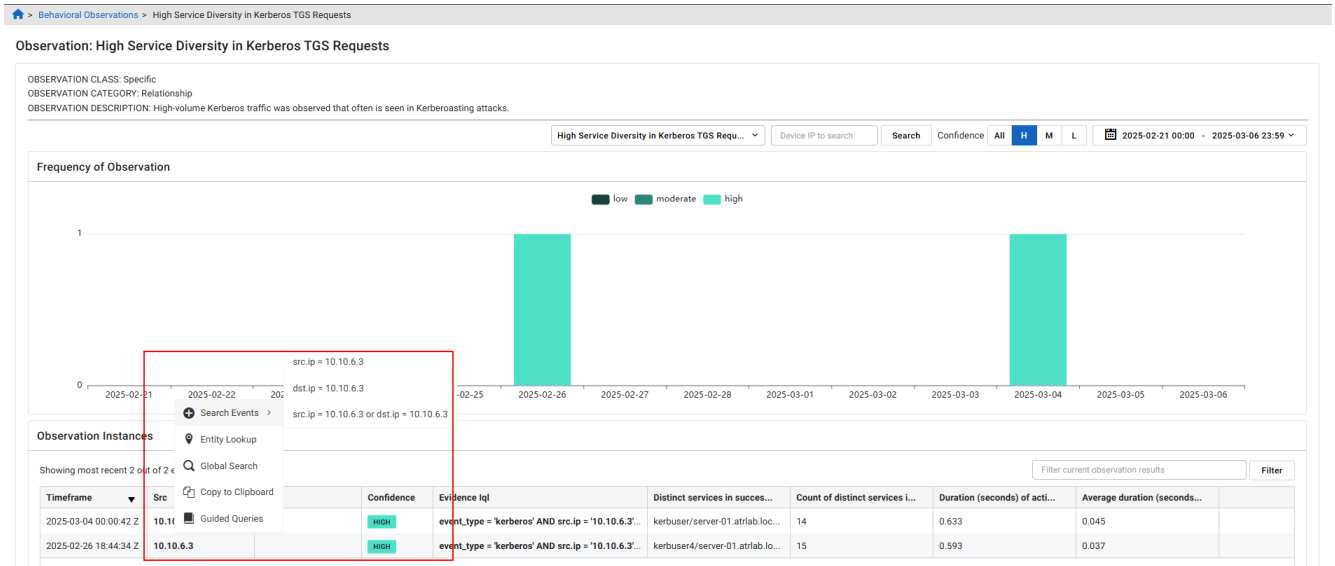

Behavioral Observation fields

Property	Description
`category`	Category of the observation: `asset`, `account`, `software`, `flow`, file, `relationship`

Property	Description
`class`	Class of the activity: `anomalous`, `newly observed`, `specific`
`dst_ip`	The destination IP of the impacted device. There may be observations with no destination device.
`src_ip`	The source IP of the impacted device. There may be observations with no source device.

Assigning detections

Assigning detections in FortiNDR Cloud allows you to manage and delegate security tasks efficiently. This topic provides instructions on how to assign, unassign, and view detections across various pages within the portal.

Assigning detections from the Detections Table

To assign a detection from the Detections Table:

1. Go to *Detections* > *Detections Table*.
2. Click the *Actions* menu at the right side of the page and select *Assign Detection*. The *Assign* dialog opens.

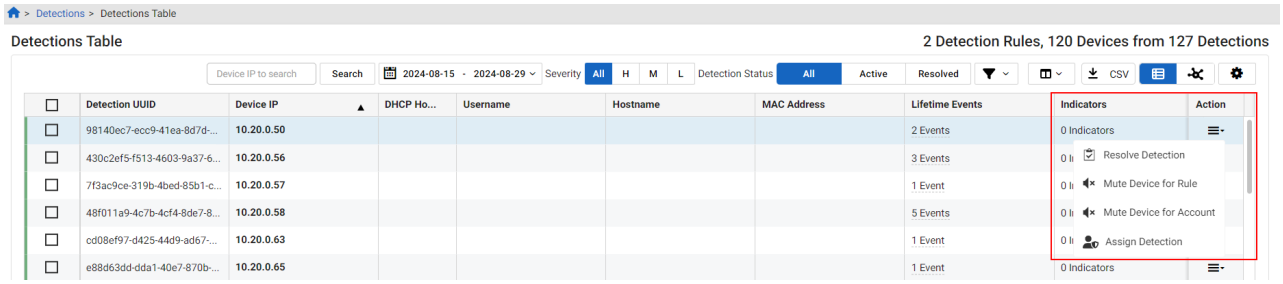

3. From the *Assignee* dropdown, select a user from the list. You have the option of assigning the detection to yourself.
4. (Optional) Enter a comment in the *Comments* field.
5. Click *Confirm*. A confirmation appears at the top of the page.

To bulk assign detections:

1. Go to *Detections* > *Detections Table*.
2. Select the detections you want to assign. The *Tools* menu appears.

3. Select *Assign <#> Detections*. The *Assign* dialog opens.

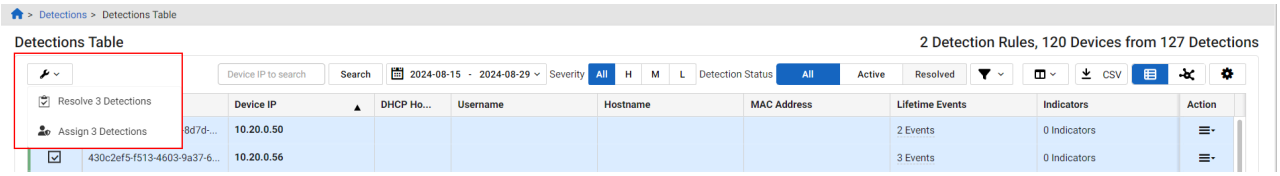

4. From the *Assignee* dropdown, select a user from the list. You have the option of assigning the detection to yourself.
5. (Optional) Enter a comment in the *Comments* field.
6. Click *Confirm*. A confirmation appears at the top of the page.

To unassign detections:

1. Go to *Detections > Detections Table*.
2. Click the Actions menu at the right side of the page and select *Assign Detection*. The *Assign* dialog opens.
3. From the *Assignee* dropdown, select *Unassigned*.
4. (Optional) Enter a comment in the *Comments* field.
5. Click *Confirm*. A confirmation appears at the top of the page.

Assigning detections from the Triage detections page

To assign a detection from the Triage detections page:

1. Go to *Detections > Triage detections*. The *Triage detections* page opens.
2. Open a detector in the list.
3. Click the *Actions* menu on the right side of the page and select *Assign Detection*. The *Assign* dialog opens.

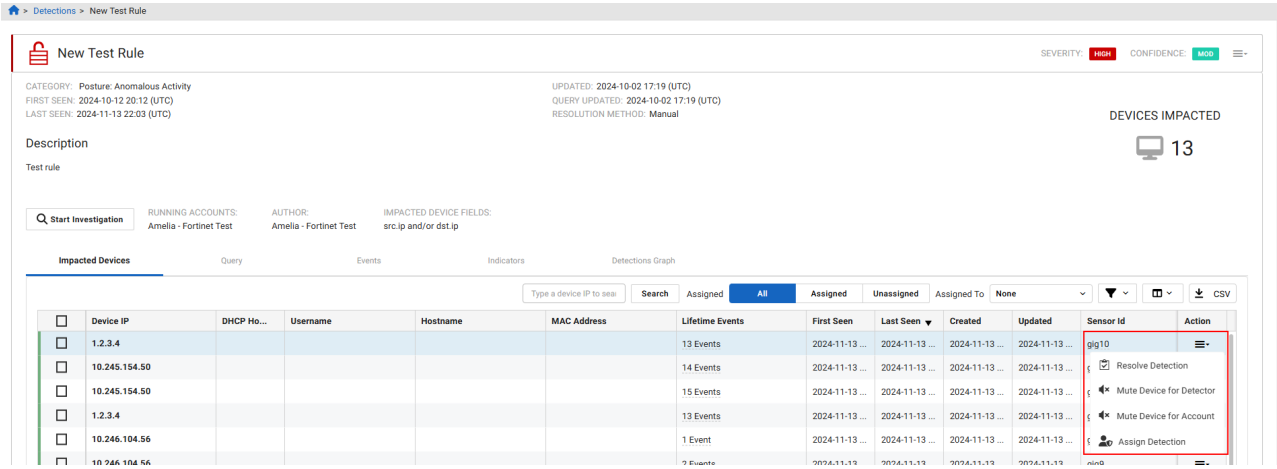

4. From the *Assignee* dropdown, select a user from the list. You have the option of assigning the detection to yourself.
5. (Optional) Enter a comment in the *Comments* field.
6. Click *Confirm*. A confirmation appears at the top of the page.

Viewing assigned detections

Detections Table

The Detections Table contains four columns with assignment information:

Assigned Comment	Notes about the detection to the assignee.
Assignee	The name of the user assigned to the detection.
Current Assign Time	The date and time the assignment was updated.
Initial Assign Time	The date and time the detection was assigned.

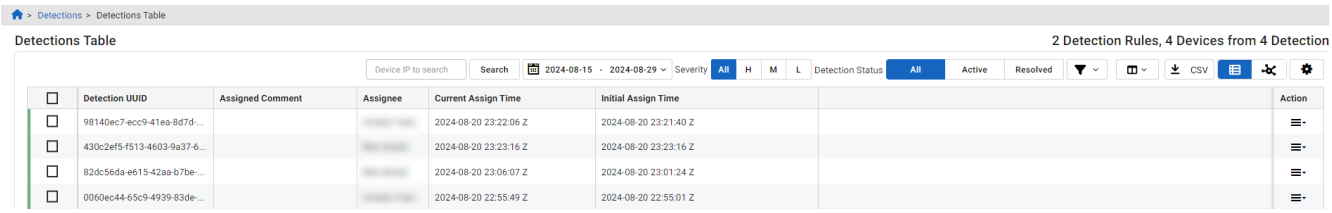

You can also use the filter to show *Assigned* and *Unassigned* detections.

Triage detections and Triage Devices

The *Assigned Comment*, *Assignee*, *Current Assign Time* and *Initial Assign Time* columns appear in the detections table of the Triage detections and Triage Devices pages. To filter the table, use the *Assigned*, *Unassiagned* and *Assigned to* filters at the top of the table.

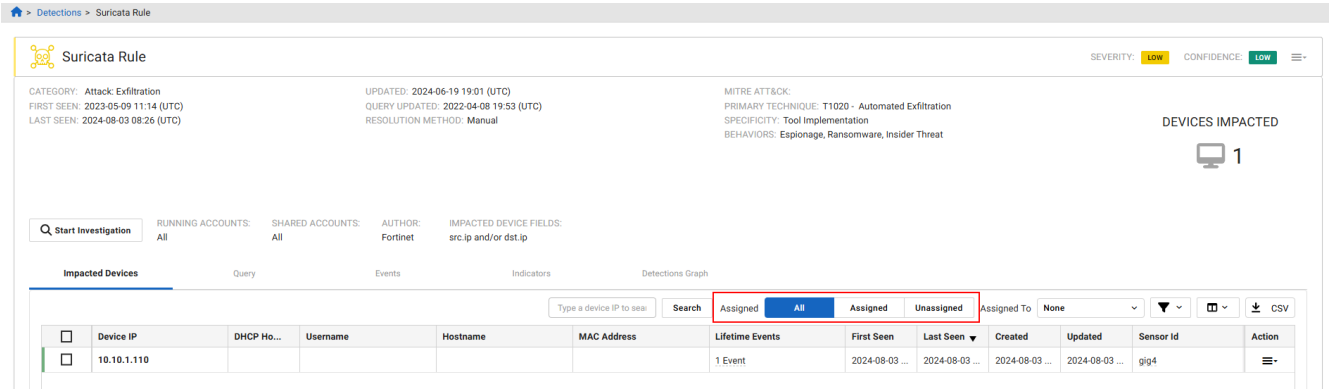

Viewing detections *Assigned to me*:

To quickly view detections that are assigned to you, in the *Triage Detections* or *Triage Devices* pages, click the filter icon and from the *Assigned to* dropdown, select *Assigned to me*.

Statistics

The *Statistics* page shows the *Active Detections Over Time* graph. Hover a line in the graph to view the defections for specific day. You can group the statistics by *Detector*, *Category* or *Severity*.

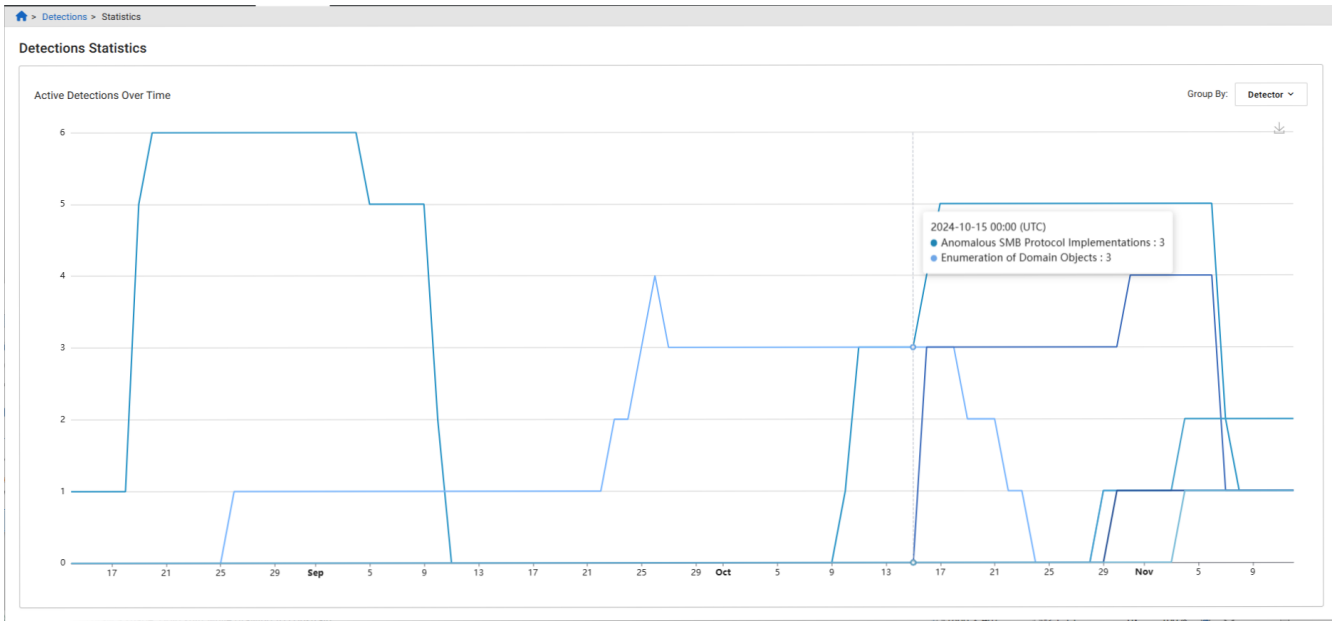

Managing detectors

The *Manage My Detectors* page allows you to view, edit and create detectors. You can also mute, disable and delete detectors.

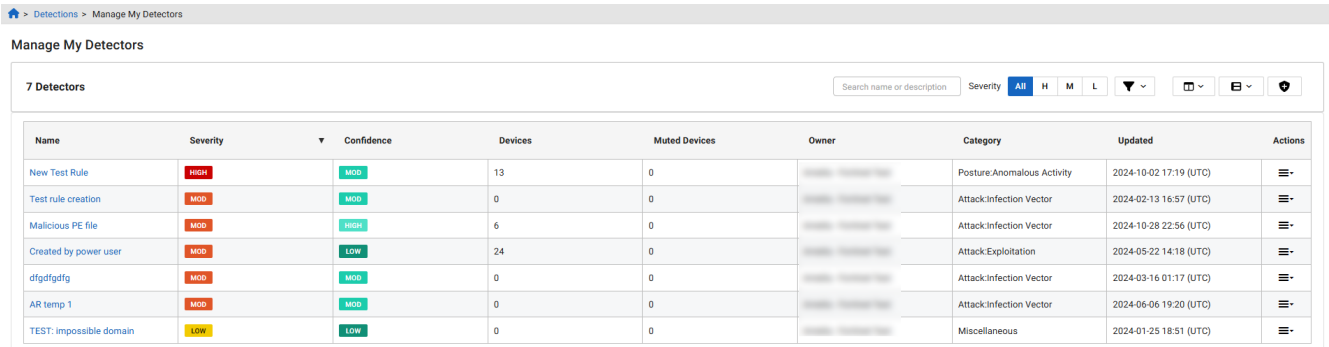

The *Manage my Detectors* page displays the following information:

Name	Click to view the detector details. An icon is displayed with the detector is disabled (⚙) or muted (🔇)

Muted	Displays an icon that indicates the detector is muted ( ) or unmuted ( ).
Enabled	Displays an icon that indicates the detector is enabled ( ) or disabled ( ).
Severity	The FortiGuard ATR severity level (Low, Moderate or High).
Confidence	The FortiGuard ATR confidence level (Low, Moderate or High).
Devices	The number of devices impacted by the detector. To view the devices, click the link in the *Name* column and review the details in the Impacted *Devices* and *Events* tab.
Muted Devices	The number of devices muted for the detector.
First	The date the detector was first detected.
Last	The date the detector was last detected.
Owner	The account name.
Category	The detector category.
Updated	The date the detector was updated.
Actions	Click the dropdown menu to: • Edit • Mute detector • Mute Device for detector • Enable detector • Delete detector

The following tools are available in the toolbar

Search titles	Filter the table by the detector name.
Severity All H M L	Filter the table by the FortiGuard ATR confidence level (Low, Moderate or High).
	Additional filters. Filters persist until you refresh the page (except for *Search title*). An indicator ( ) is added when you change a filter from the default. A number indicates the number of changes that were applied. Click *Reset to Default* to clear the filters.

Filter	Description
Category	Click to select a category from the dropdown.
Technique	Click to select a technique from the dropdown.
Confidence	Filter by FortiGuard ATR confidence level (All, H, M or H). *All* is the default.
Detection Status	Filter by detection status (*All*, *Active* or *Idle*). *All* is the default.
Muted	Select *Unmuted* or *Muted* . *All* is the default.

	Filter	Description
	Disabled	Select *Enabled* or *Disabled*. *All* is the default.
		Show or hide all columns in the table, or select the columns you want to view.
		Set the page height.
		Create a new detector. See Creating a detector on page 48.

Response configuration

The *Response Configuration* feature allows you to automatically ban an IP address when a high-severity and high-confidence detection occurs.

Automated integration response is available for FortiEDR, CrowdStrike Falcon EDR and FortiGate via FortiManager at this time. Only a single integration can be set to *Auto-Remediate* at a time. Other integrations may be configured, but must be set up to respond manually.

To enable automated response configuration:

1. Go to *Detections* > *Response Configuration*. The *Integration Response Configuration* dialog opens.
2. In the *Action* column, click *Edit* next to the integration.
3. In the *Configure* dialog, select *Auto-remediate* and click *Save*.

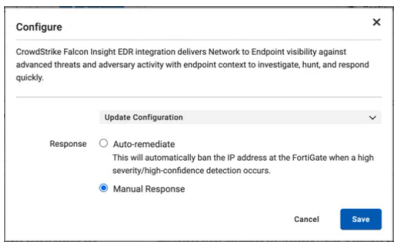

You can also enable *Response Configuration* in the *Account Management* > *Modules* page by clicking *Configure* in the integration's tile.

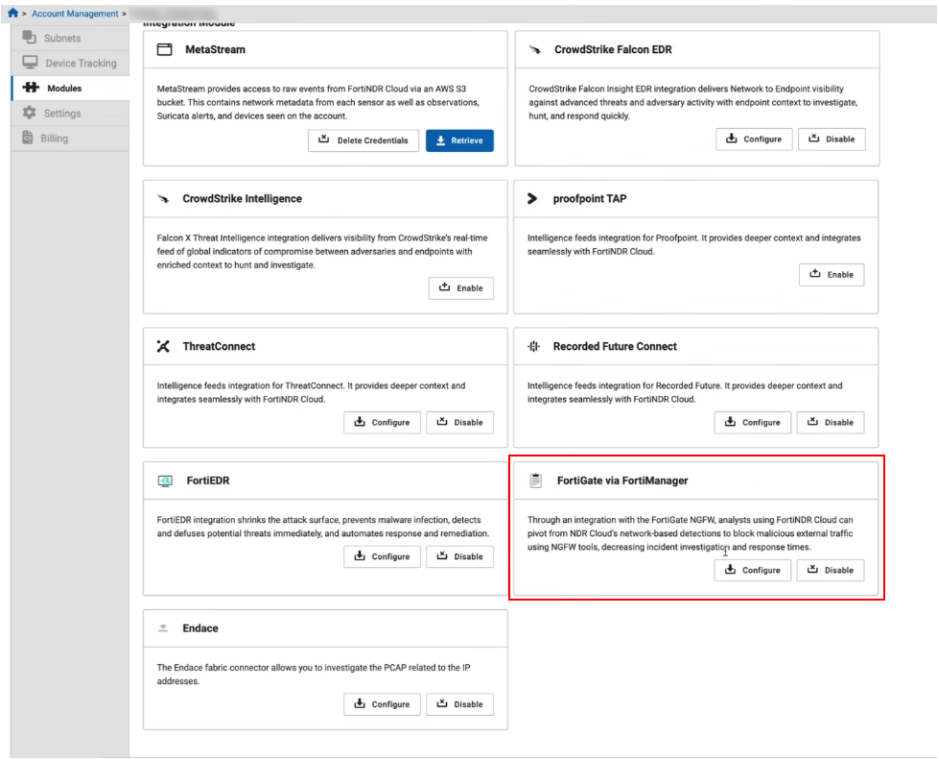

Creating column profiles

You can create and manage column profiles to help you organize and customize your data views. Custom profiles can also be shared with other users in your organization.

To create a column profile:

1. Go to:
   - *Detections > Detections table.*
   - *Investigation results*
2. Select the columns you want to include in the profile, apply filters, and adjust the column width.
3. Click the column selector icon.

4. Under *Column Profiles*, click *Create New Profile*. The Column Profile dialog opens.
5. Configure the column profile settings and click *Save*.

Name	Enter a name for the column profile.
Include time range	Select one of the following: • *Absolute (xxx-xx-xx - xxxx-xx-xx)*: This is the date range in the date picker. • *Relative (last xx Days)*: The value of *xx* is the difference between the

	start date and the end date. This option only applies to the Detections Table.
Include other filters	Enable to include any filters you applied to the table. This option only applies to the Detections Table.
Shared	Enable to share the column profile with other members of your organization.

To view and edit column profiles:

1. Select the columns you want to include in the profile, apply filters, and adjust the column width.
2. Click the column selector icon.
3. Under *Column Profiles > My profiles*, click the profile you want to view.
4. Click *Save this Profile* to update any changes you made.

Risk score calculation

The risk score for a device is calculated as a weighted sum of individual detection scores, based on a predefined matrix. This sum is capped at a maximum score, ensuring it does not exceed a defined ceiling. If a device has multiple detections with varying severities, the ceiling is determined by the highest severity level among those detections.

If a detection is muted or resolved, its score is 0. Otherwise, the score is calculated using the following matrix:

Scoring Matrix

Severity	Low Confidence	Moderate Confidence	High Confidence
Low	0.1	0.3	0.5
Moderate	0.5	1	2.5
High	1	2.5	5

Maximum Score Limits

To prevent extreme values, the score is capped based on severity:

Severity	Max Points
Low	2.5
Moderate	5

Severity	Max Points
High	10

This scoring system helps prioritize detections based on how confident and severe they are, while also allowing flexibility for high-severity cases.

Investigations

Use the tools in the *Investigations* module to respond to detections and to hunt for malicious activity on you network.

Entity lookup

An *Entity Lookup* (or search) is the starting point for an investigation if you have very little information to work with.

You can start an Entity Search by entering an IP address or domain in the *Search* field in the navigation menu at the top of the portal.

To perform an entity lookup:

1. Go to *Investigations* > *Entity Lookup*.
2. Enter an IP address or a domain name in the search field. Separate Multiple IP addresses and domain names by spaces.
3. Click the date picker to select the time range. The default is *Last Seven Days*. The maximum is 90 days.

   If you are pivoting to the *Entity Lookup* from a page with a time range of more than the last 90 days, the date range picker will display a yellow border around the date field and default to the *Last Seven Days*.

4. Click *Search*. The following results are returned.

Network Intelligence	Network traffic by service, by device, and source addresses interacting with the entity
Entity Intelligence	WHOIS, IP History, Registrar History, Passive DNS
Security Intelligence	Associated VirusTotal Detections, VirusTotal Detections Over Time, Detections, and Observations,

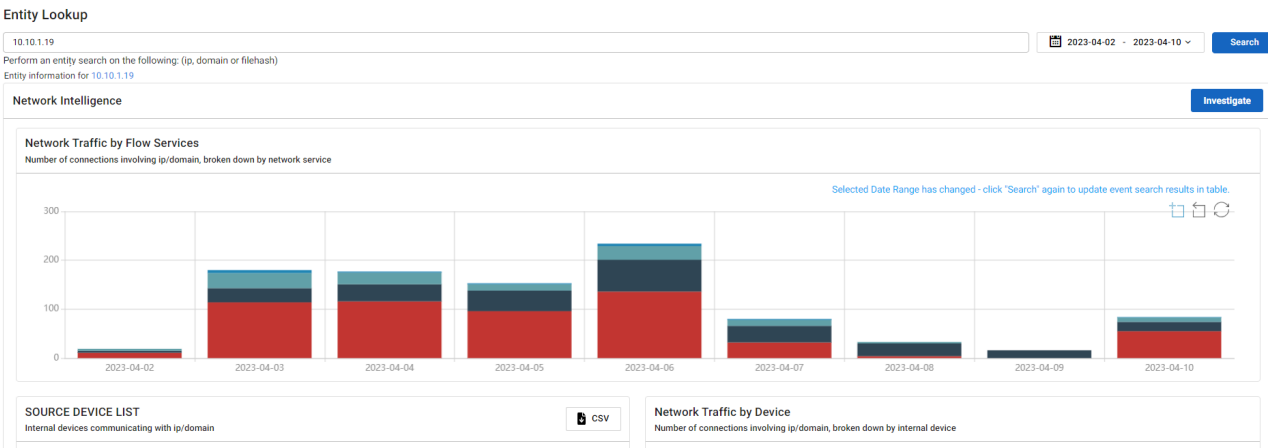

You can view the *Entity Panel* by clicking the IP address at the top-left of the page next to *Entity information for <IP address>*.

5. (Optional) If multiple IP addresses or domain names are looked up, right-click on a result and select *Entity Lookup* to view the intelligence panes.
6. (Optional) Click *Investigate* to launch the new investigation.

To perform a bulk entity export:

1. In the search field, enter IP addresses or a domain names separated by spaces.
2. Click *Search*.
3. Click the *CSV* button. A CSV file with the *timestamp*, *action*, *param*, *user_uuid*, *account_uuid*, and *account* are downloaded to your device.

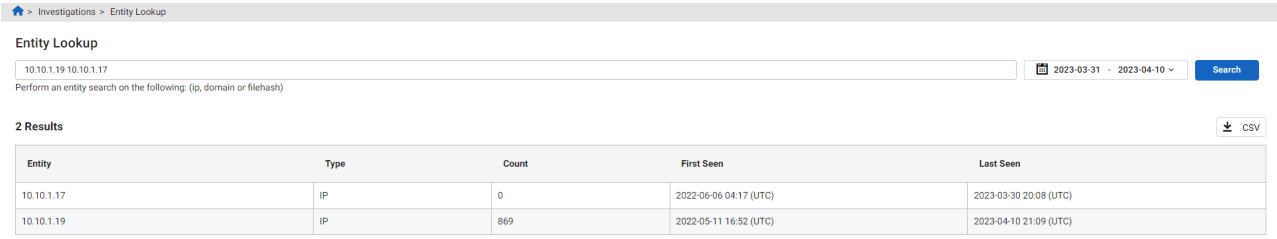

Passive DNS

Passive DNS links on the entity panel function like normal links. Clicking the link replaces the entity panel with the panel for the clicked on element.

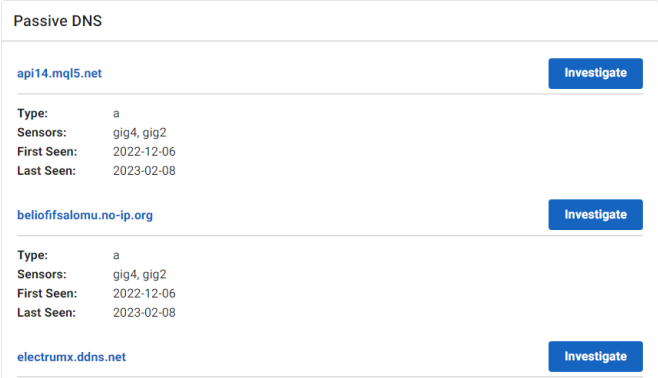

Right-clicking opens a context menu.

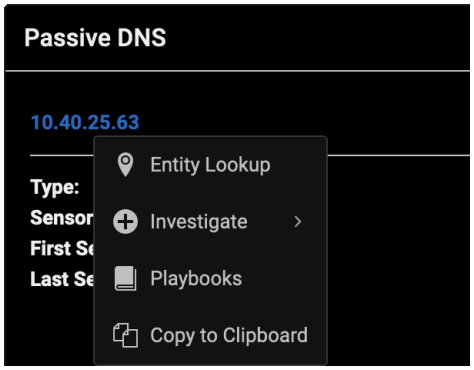

Option	Description
Entity Lookup	Open the entity lookup page for the item.
Copy to Clipboard	Copy the item to the clipboard.
Guided Queries	Launch Guided Queries. This options is not available for ad-hoc search result items
Investigate	Show appropriate pivots for the item type. This options is not available for ad-hoc search result items.

Search Events	Show the event searches appropriate for the type. The text in the search box is replaced, but the search will not run automatically. This options is only available for ad-hoc search result items.
	Types include:
	• IP:
	  • ip='IP'
	  • dst.ip='IP'
	  • src.ip='IP'
	• domain:
	  • domain='domain'

Investigate

Investigations allow you to quickly obtain details required in investigations via search queries and/or Guided Queries.

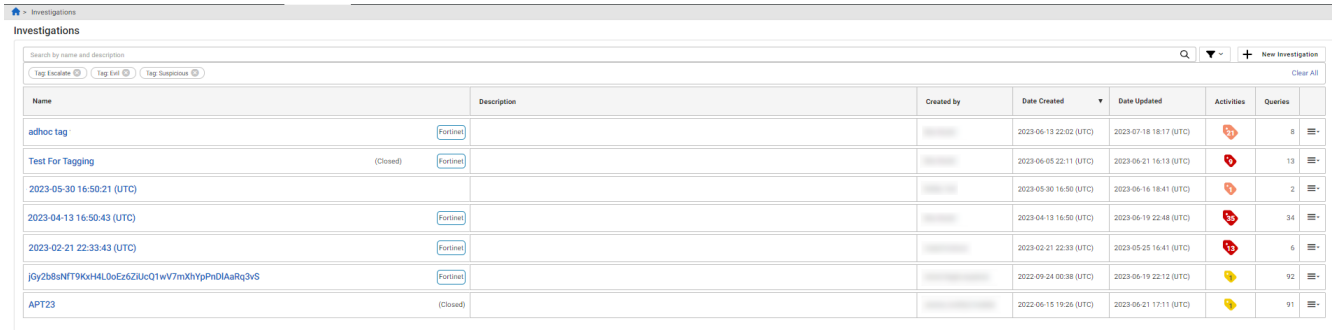

The Investigations page displays the following information:

Name	The investigation name.
Description	The description of the investigation.
Created by	The user who created the investigation.
Date Created	The date the investigation was created.
Date Updated	The date the investigation was updated.
Queries	The number of queries added to the investigation.

Filtering investigations

Click the filter icon next to the *Search* button to view by following attributes:

Created by	Select FortiNDR Cloud user from the list.
Relates to	Select a related investigations from the list.
Tag	You have the option of viewing only tagged or untagged investigations. You can also filter by a specific tag.
Investigation Status	Select All , Open or Closed investigations.
Investigation Type	Select *All*, *Standard* or *Report*.

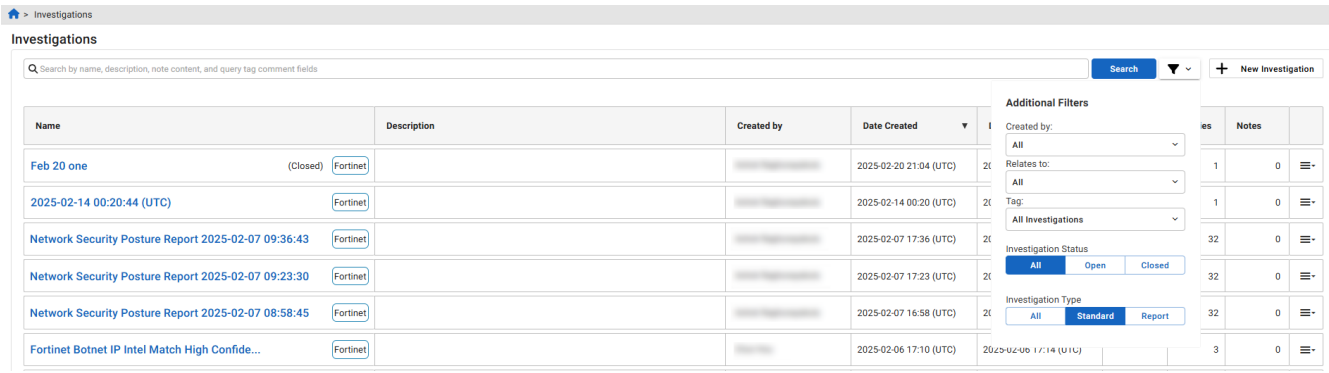

The selected filters are persistent. For example, if you sort the table by *Date Updated* and then browse to a different page in the GUI, the investigations table will still be sorted by *Date Updated* when you return to the *Investigations* page.

When you add filters, the filter chips will be shown under search bar.

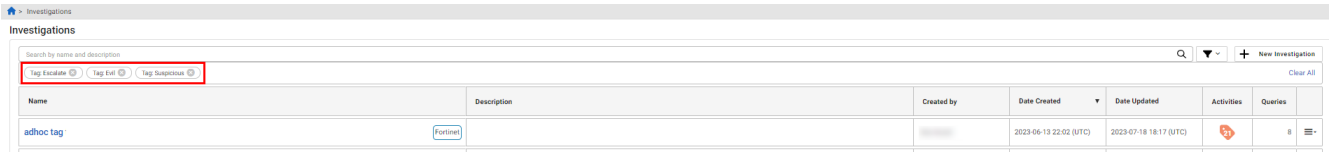

Creating investigations

An investigation is run against the account shown in the account picker. The account name that owns the investigation appears to the right of the investigation name if it differs from your primary account.

- If you have access to multiple accounts and the account shown in the account picker is different from the account that contains your user, then the account is listed.
- If you have access to multiple accounts, and the account shown in the account picker is the same as the account that contains your user, then the account is not shown in the investigation list. The investigation created is run against the account shown in the account picker.

To create an investigation:

1. Go to *Investigations* and click *New Investigation* at the top-right corner of the page. The *New Investigation* dialog opens.

   The default investigation name is the first and last name of the user creating the investigation with the time stamp of when the investigation was created.

2. Enter an *Investigation name* and *Description*, then click *Create Investigation*. The investigation is created. The system now automatically adds a summary note that is visible above any subsequent query entries.

3. Add the following to your investigation:
   - Query: Adding queries to an investigation on page 87
   - Guided query: Adding a guided query to an investigation on page 114
   - Notes: Adding notes to an investigation on page 89

To close an investigation:

1. Go to *Investigations* and click the investigation you want to close.
2. Click the gear icon at the top-right side of the page and select *Close Investigation*. A confirmation dialog opens.
3. Click *Close Investigation*.

To delete an investigation:

1. Go to *Investigations* and click the investigation you want to delete.
2. Click the gear icon at the top-right side of the page and select *Delete Investigation*. A confirmation dialog opens.
3. Click *Confirm*.

Deleting an investigation is irreversible and will remove everything in the investigation

To edit an investigation name:

1. Go to *Investigations* and click the investigation you want to edit.
2. Click the gear icon at the top-right side of the page and select *Edit Investigation*. A dialog opens.
3. Update the *Investigation name* and *Description* and click *Save*.

Viewing investigation details

To view the investigation details, go to *Investigations*, and click an investigation name. The investigations details page displays the following information:

- Investigation Creator
- Link to single or multiple related detections
- IQL query
- Notes (if any)
- Date/time the query was added
- Number of events (if complete)
- Executed Guided Queries that are part of that investigation
- Close date (if investigation was closed)

If the investigation contains more than one related detection, the *MORE>>* link appears. You can click the link to view all the related detections.

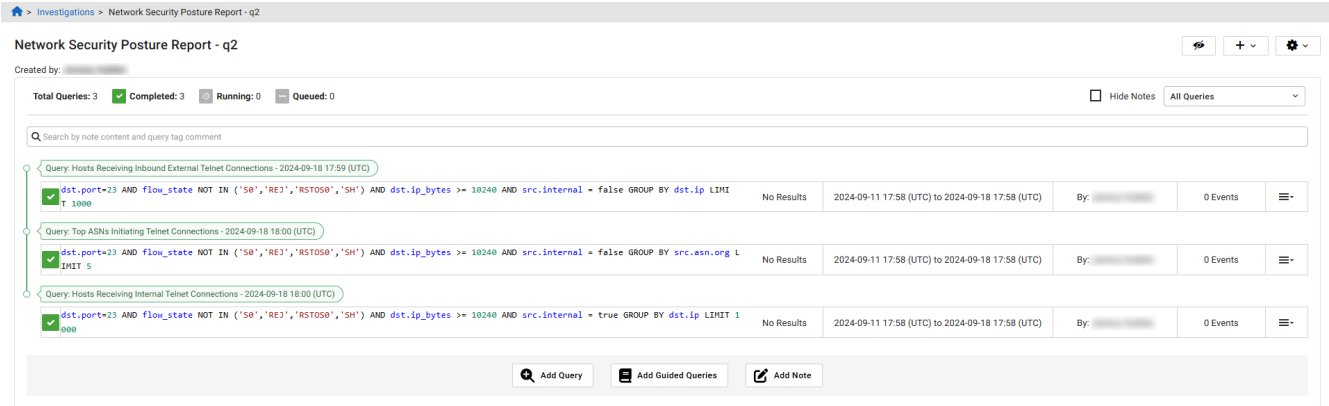

Query Status Icons

	Query completed successfully. Results (if any) are available.
	Query is currently running.
	Query is queued to run. It will run automatically when resources are available.
	Query failed due to an internal error. If problem persists, please contact Fortinet support.

Viewing results

To view the investigation results, click the *View Results* button in the investigation details.

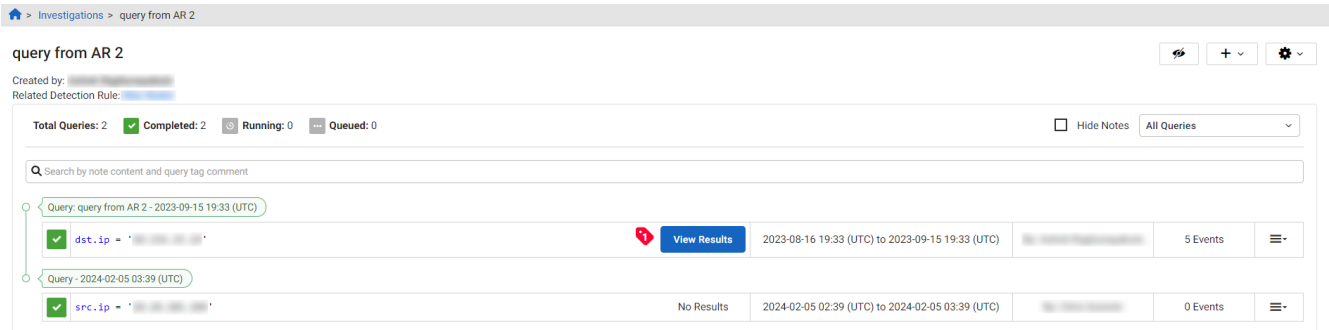

The following information is displayed:

- IQL Query string
- Date Range
- Number of events
- A table of the events where you can:
  - Click on column filter to change the visible columns in the way that the current event search does including column visibility sets.
  - Click the *CSV* button to export the results as a CSV file

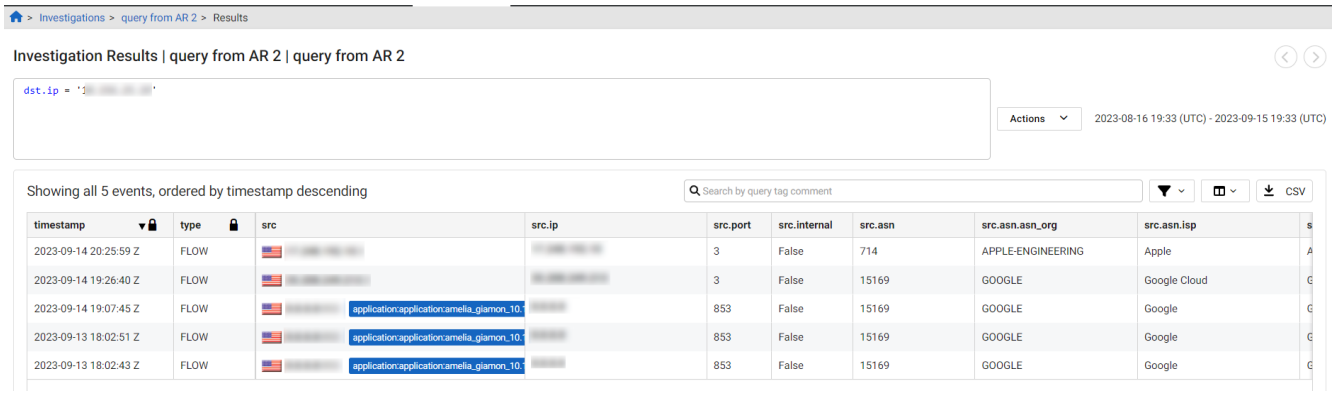

Viewing column data

To quickly scroll through the column headings, hold down the Shift key and use the scroll wheel on your mouse.

To adjust the columns to fit the widest cell in the table or to hide a column, right-click the column header.

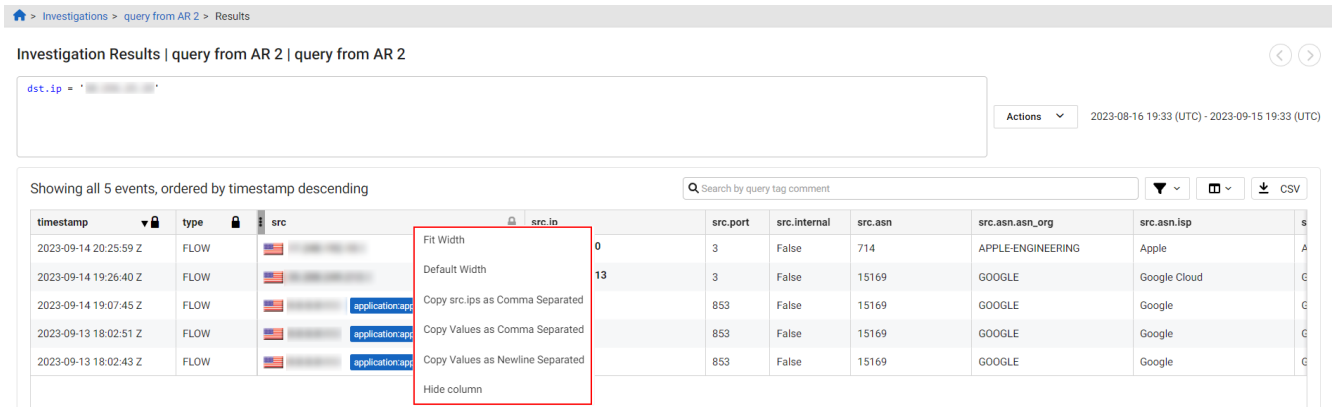

Single event view

You can view all details for a single event by double-clicking a blank area within the event row. This opens a pop-up displaying the full row data in JSON format. To copy the JSON, click the copy icon next to the first line. This saves time by eliminating the need to scroll through individual cells in the investigation results table.

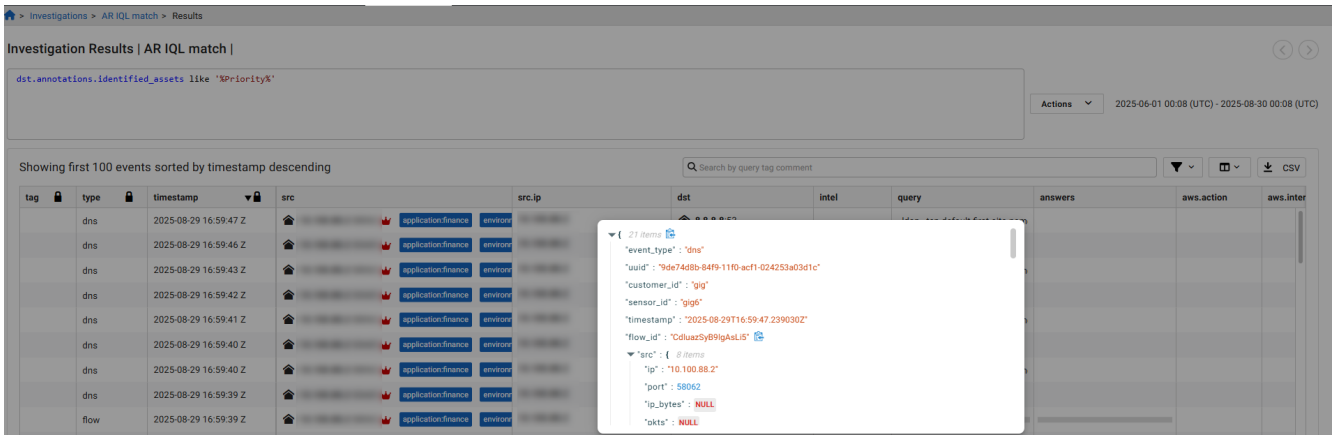

Chart types

Investigations results can be viewed as a *Tree Map*, *Pie* chart, or *Graph*. A *Sankey* chart is available for aggregations where the *Group By* includes two IP fields or when there are two dimensions and a measure. The Sankey chart type appears only when there are 50 or fewer dimensions.

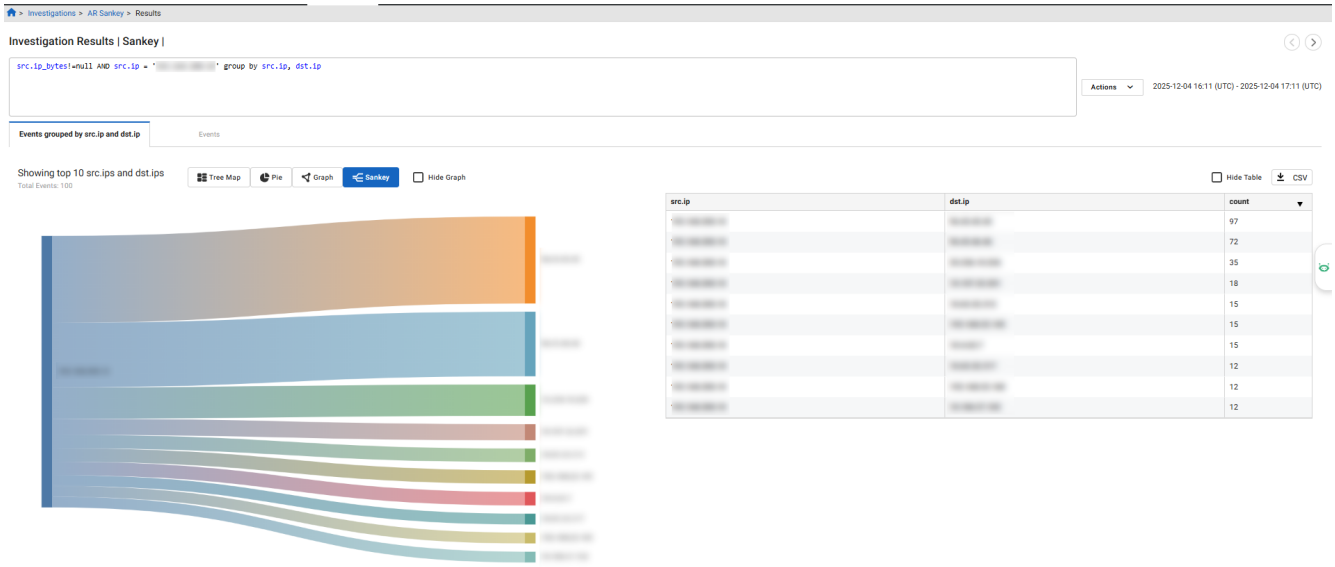

Adding queries to an investigation

You can add one or more queries to an investigation.

To add a query to an investigation:

1. Go to *Investigations* and click an investigation the list.
2. Click *Add Query*. The *Add a New Query* page opens.

3. Configure the query settings.

Name	Enter a name for the query.
Select Saved Query	Click to base the new query on a saved query.
Query	Enter the query string.
Actions	Options are: • *Bulk Add Indicators* • *Create a Detection*
Sort by timestamp	Select *Ascending* or *Descending*.
Last 7 Days	Use the date picker to update the date range and click *Apply*.
Retrieve up to *xxx* rows	Select between 100 to 10,000 rows.
Enable Facets	Select to return the panel that allows narrowing the search. This may make the query longer to complete. For more information, see Facet Search on page 90.

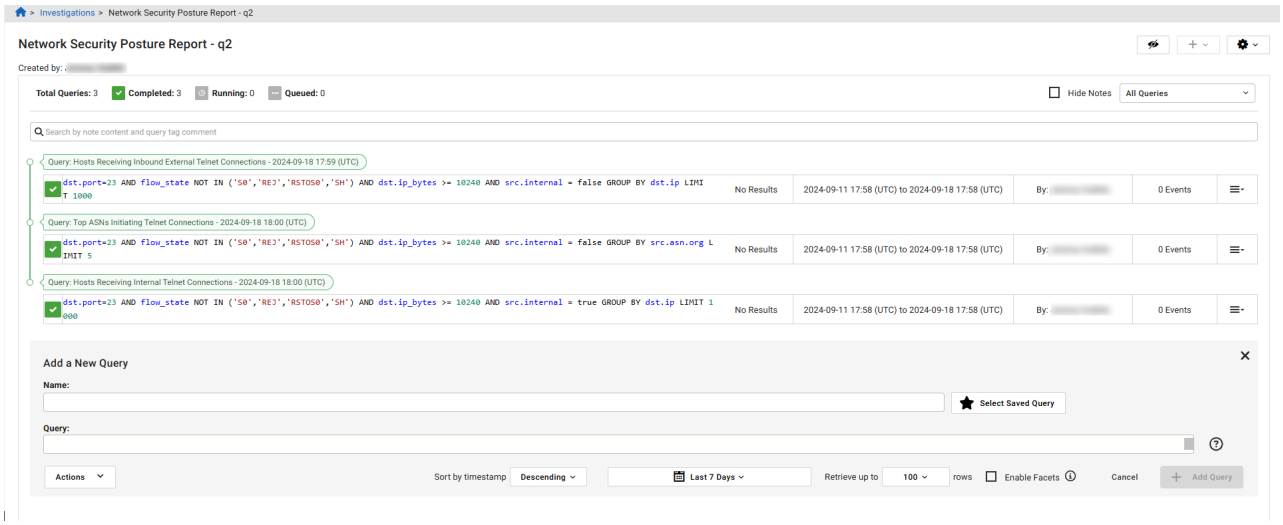

4. Click *Add Query*.
5. (Optional) To add another query to the investigation, click *Add Query*.

To rename a query:

1. From the Investigation Detail page, locate the query you want to rename.
2. Click the *Actions* menu on the right side of the page and select *Rename*.

3. Enter the name in the *Query name* field.
4. Click *Rename*.

To clone a query:

You can clone a query in a closed investigation. However, the cloned query must be added to a different investigation.

1. Click *Investigations*.
2. Click the investigation that contains the query you want to clone.
3. Click the *Actions* menu on the right side of the page and select *Clone*. The *Add Query to Investigation* dialog opens.
4. Configure the query settings.
5. Create a new investigation or save the query to an existing investigation.

Create a New Investigation	Enter an *Investigation Name* and *Description*.
Add to Existing Investigation	From the *Choose Investigation* dropdown, select an investigation. By default the cloned query is added to current investigation.
Run a Private Query	Select this option to add a query to an adhoc search.

6. Click *Add Query*.

To delete a query:

1. Click *Investigations*.
2. Click the investigation that contains the query you want to delete.
3. Click the *Actions* menu on the right side of the page and select *Delete*. The *Delete Query* dialog opens.
4. Click *Confirm*.

To save a query:

1. Click *Investigations*.
2. Click the investigation that contains the query you want to save.
3. Click the *Actions* menu on the right side of the page and select *Save*. The *Save Query* dialog opens.
4. Enter a *Query Name* and *Description*.
5. Click *Save*.

Adding notes to an investigation

To add a note to investigation:

1. Go to *Investigations* > *Investigate*.
2. Click Select to open an investigion.
3. Click *Add Note*. Optionally, you can click the Add menu (**+**) in the top-right of the page and select *Add Note*.

4. In the *Notes* field enter the details in plain text or markdown. Rendered markdown text will be visible. The note contents will be displayed along with the timestamp of when it was created.

To update a note:

1. Click the Actions menu on the right side of the note and select *Update*.
2. Update the note and click *Update Note*.

To delete a note:

1. Click the Actions menu on the right side of the note and select *Delete*. The *Delete Note* dialog opens.
2. Click *Confirm*.

Watch an investigation

You can check the status of your query by clicking the *Notification* icon to the right of the account name in the top navigation. A panel displays the list of queries being watched, along with the number of queries completed and running.

When the query is complete, you will see a green check mark in the top right corner.

To watch an investigation:

1. Go to *Investigations* and click *Select* to open the investigation you want to watch.
2. Click the *Not Watching* icon.

To unwatch an investigation:

1. Go to *Investigations* and click *Select* to open the investigation you want to watch.
2. Click the *Watching* icon.

Facet Search

A *Facet* filters results of an IQL query in a pane adjacent to the main results table of an IQL query. A facet is an automatic filter that saves time configuring a search with the GUI.

The facet options are results-based attributes from a sample of the events found in the initial search. The facets will change based on the data in the records found by the search.

Faceted Searches are useful for getting a quick multidimensional view of the results to identify the most or least common elements.

You can enable Facets when:

- Adding queries to an investigation on page 87
- Adding a guided query to an investigation on page 114

Enabling facet search, may increase the time to process the query.

Refine results using facet search

You can further refine your search on the results from the original query using facet search.

To refine the results in a facet search:

1. Click *Investigations*.
2. Click *Select* next to the investigation you want to open.
3. Click *View Results* for the facet search query you want to refine. The *Refine Search* pane displays a breakdown of the query results.

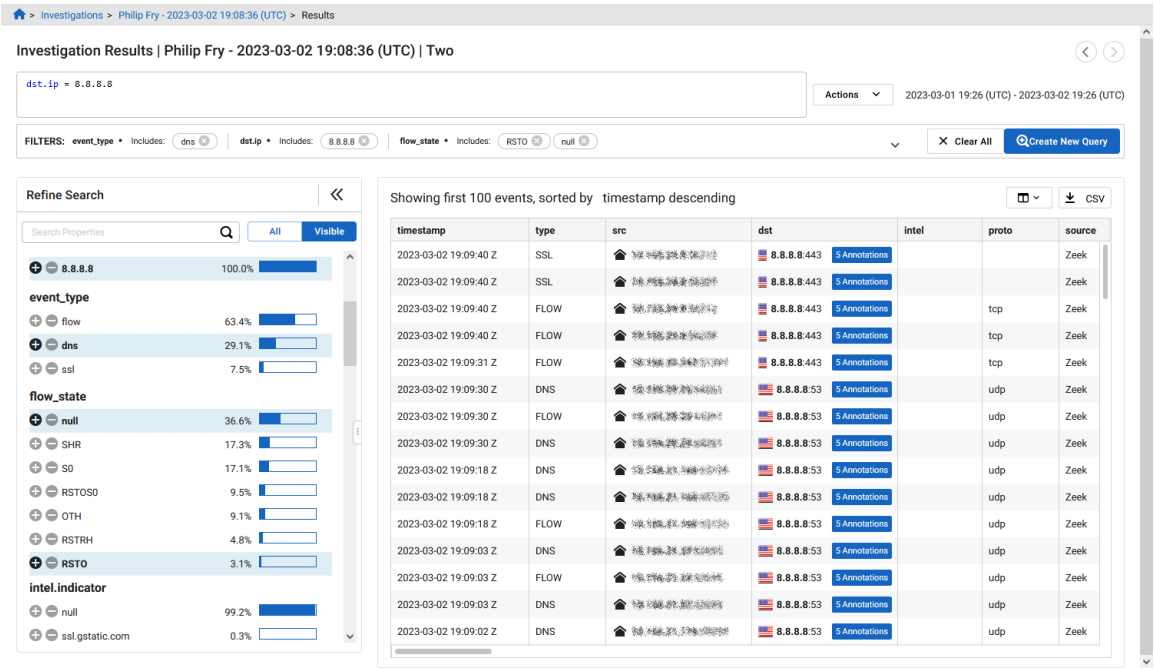

4. Add or remove the filters based on your requirement. The selected filters appear under the original search query. You can also clear the selected filters by clicking *Clear All* .
5. Click *Create New Query*.

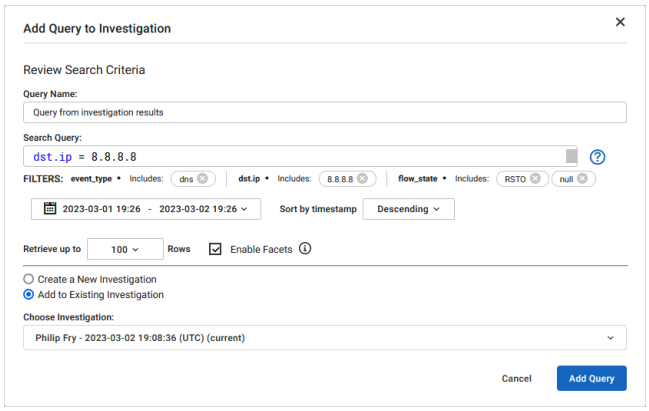

6. Create a new investigation or add the query to an existing investigation. By default, the new query is added to the current investigation.

Create a New Investigation	Select this option to create a new investigation. Enter the *Investigation Name* and *Description*. The default name for new investigations is the first and last name of the user creating the investigation as well as a date stamp of when the investigation was created.
Add to Existing Investigation	From the *Choose Investigation* dropdown, select and investigation.

7. Click *Add Query*. The query and all the included and excluded facets will be shown in the investigation details page.

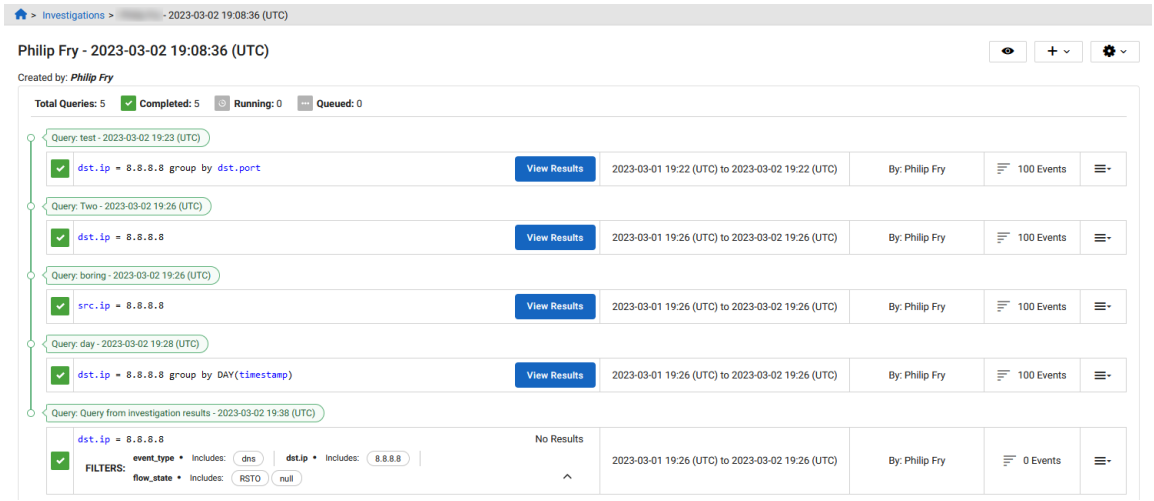

Tag and comment events

Use the *tag* column to communicate with members of the security team about an event in an investigation. Tags and comments are viewable to any user with access to the investigation. You can use filters to view only tagged investigations or use the *Search* function to search for text in notes and comments.

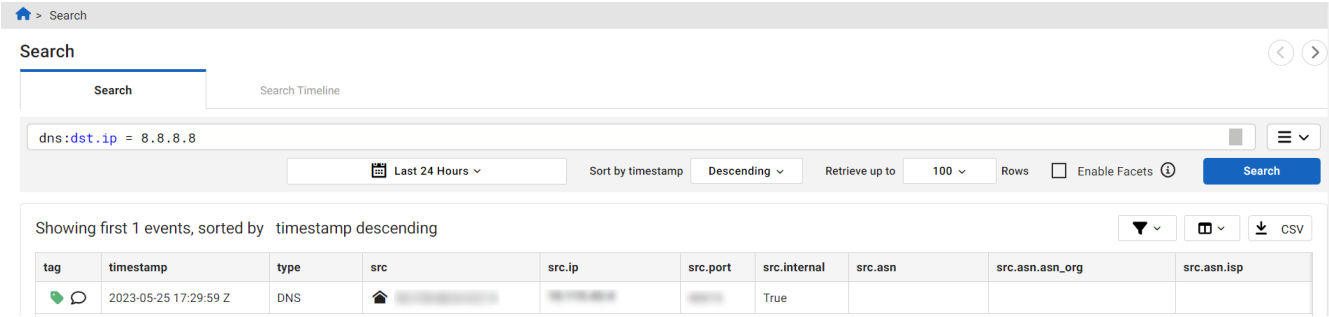

To add a tag to an event:

1. Do one of the following:
   - Click the *Investigations* tab, open an investigation and click *View Results*.
   - Go to *Investigations > Private Search*. In the *Private Search* tab, click *View Results*.
2. Click the *tag* column next to the event. The *Tag and Comment* dialog opens.

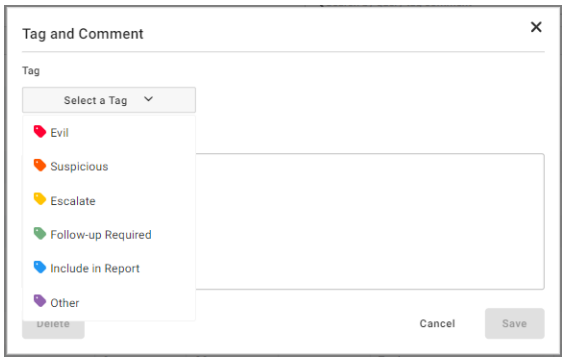

3. Select a tag from the dropdown.
4. (Optional) Add a comment to the event.
5. Click *Save*. The tag and comment icons are displayed in the *tag* column.

To remove a tag from an event:

1. Click the *tag* column next to the event. The *Tag and Comment* dialog opens.
2. Click *Delete* and then click *Confirm* in the dialog that opens.

Viewing and filtering tagged events

Tagged events are displayed in the *Investigations* and *Private Search* tabs. Hover over a tag to see an overview of the tagged events in the investigation.

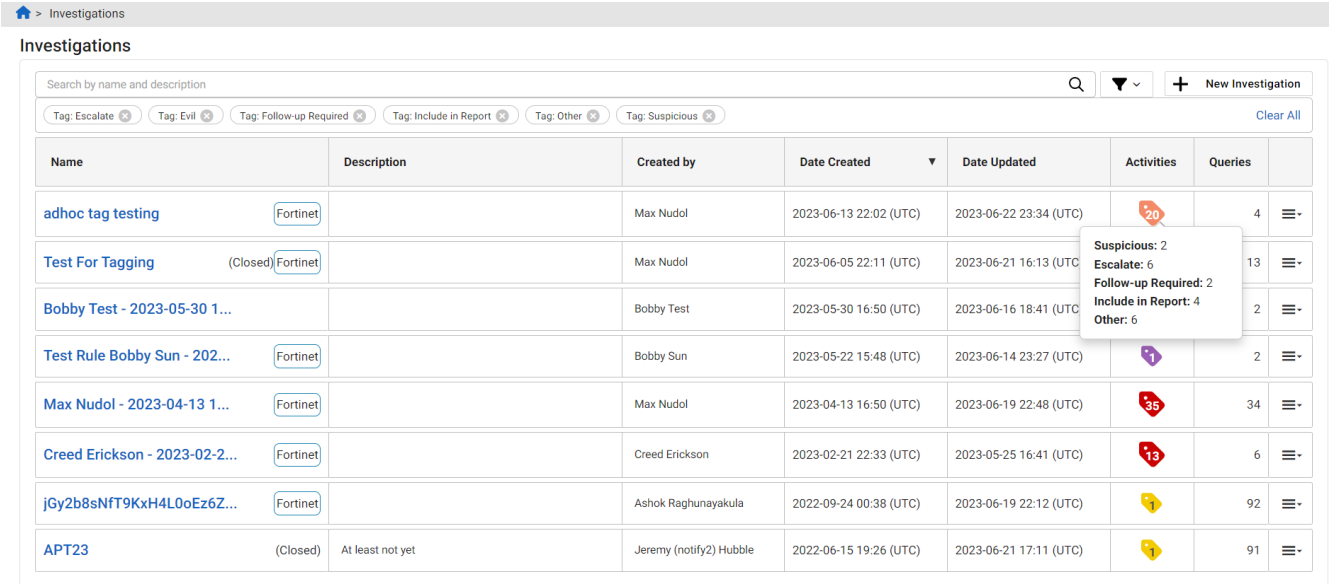

To use tags and notes to filter investigations:

Option	Description
Go to *Investigations > Investigate*	1. Click the *Filter* icon. 2. In the *Tag* section, select *Tagged Investigations*. 3. (Optional) To refine results, select a tag label from the list (such as *Evil*). 4. Click the investigation name. 5. (Optional) Click *Hide Notes* to only see the tags. 6. Click *View Results*.
Go to *Investigations > Private Search*	1. Click the *All Queries* dropdown. 2. In the *Tag* section, select *Tagged Investigations*. 3. (Optional) To refine results, select a tag label from the list (such as *Evil*). 4. Click *View Results*.
Go to *Investigations*	1. Enter keywords in the *Search* field to search for text in comments and notes. Matching results are highlighted in yellow. 2. Hover over the results in the *Activities* and *Notes* column. • Click a matched note to open the results table displaying the matched results. • Click *View Details* to open the investigation. The matched text will be highlighted.

After you filter the investigations, you can copy the URL to send the filtered view a member of your team.

Using tags to pivot to the events table

You can use a tagged event in the investigation dialog to quickly pivot to the *Events* table in an investigation. This function is available in the *Investigations* and *Private Search* pages. Click the tag icon in the tooltip.

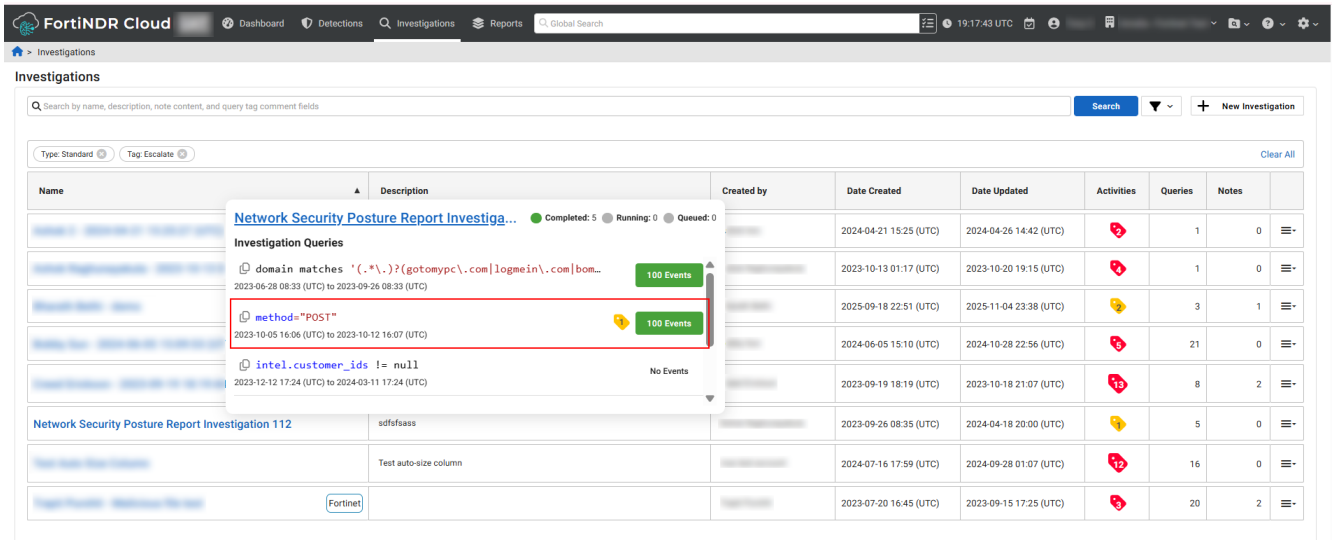

The *Events* table will display the same number of events tagged in the investigation tooltip.

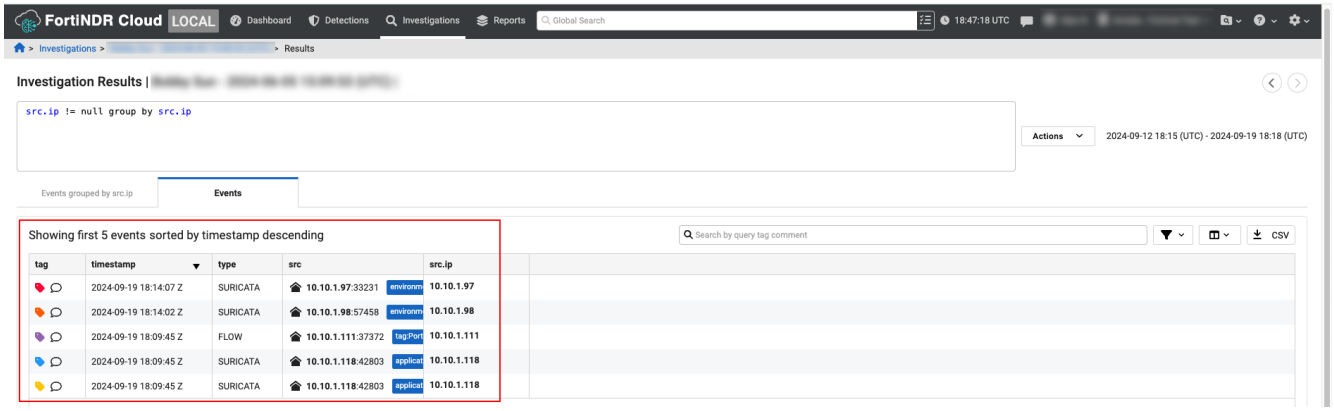

Locking the tags column

To lock the *tag* column to the left side of the table, in the *Individual columns* filter under *Individual Columns*, select *Reset to Default* or *tag*. You can also lock the column by selecting *Default* under the *Column Profile* in the same menu.

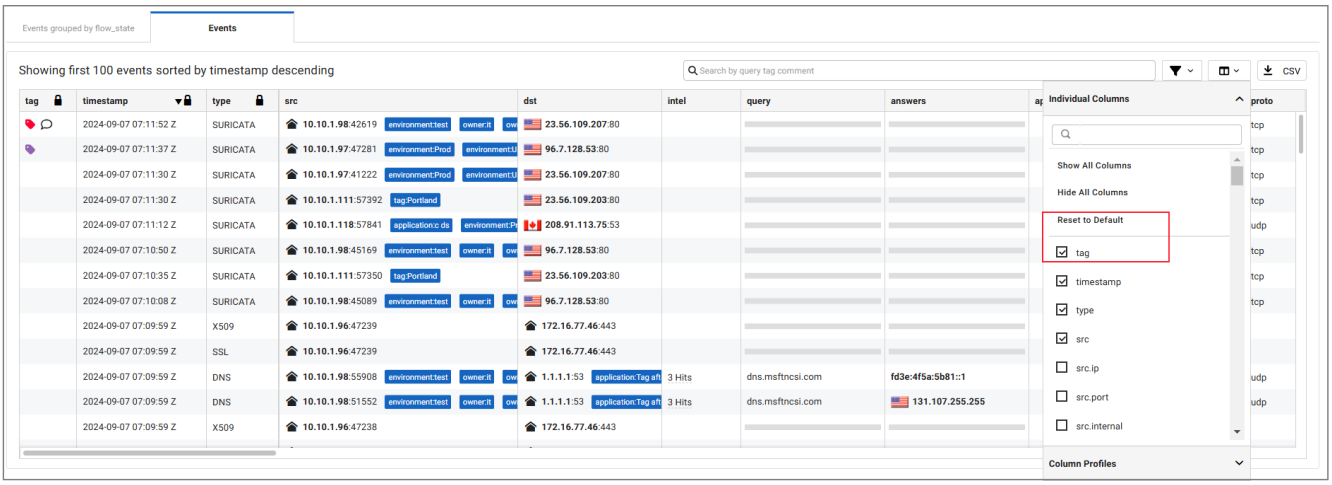

Investigation tooltip

The investigation tooltip is available in the *Investigations* page, the *Investigation* widget in the default dashboard, and global search results. The tooltip can be disabled from the *Profile Settings* page.

Hover over the investigation name to view a summary of the query status. The investigation tooltip shows the number of queries that are *Completed*, *Running* and *Queued* as well as the tags associated with the query. To view the query parameters at a glance, hover over the query string in the tooltip. To copy the query string, click the *Copy Query* icon at the left-side of the string.

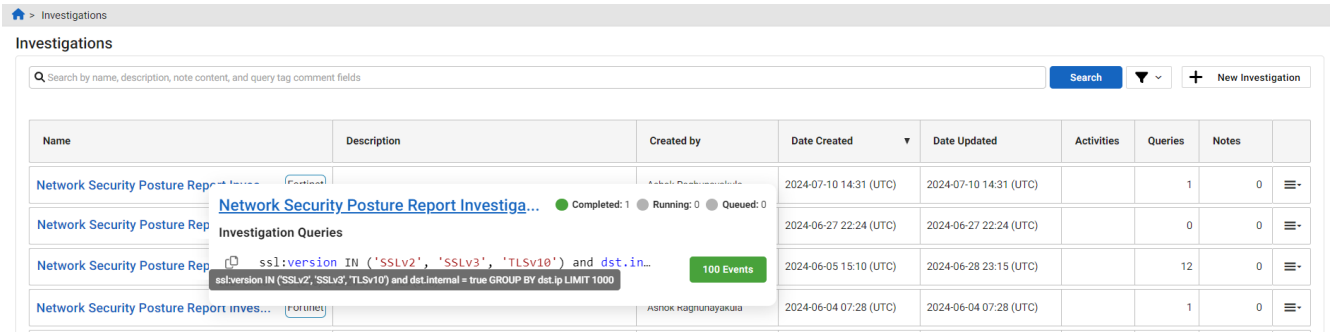

To disable the investigation tooltip:

1. Click gear icon at the top-right of the page and select *Profile Settings*. The *My Profile* page opens.
2. Under *User Information*, disable *Tooltip*.

Share investigations

Users with multiple accounts can share investigations in their primary account with users in their secondary account. Sharing an investigation will allow all users with access to the secondary account to see and make changes to the investigation. Once the investigation is shared, it cannot be undone.

Requirements:

- The user must have multiple accounts
- The investigation must be active.
- Users in the primary and secondary accounts must have a *User* role.

To share an investigation:

1. Go to *Investigations* and do one of the following:
   - Create a new investigation.
   - Open an investigation in the list. Investigations in your primary account are indicated with an account icon.
2. Click the gear icon at the top-right of the page and select *Share With Account*. The *Share with Account* dialog opens.

3. Click *Confirm*. The *Investigations* page opens. The account icon is removed from the investigation and a confirmation message appears at the top of the page.

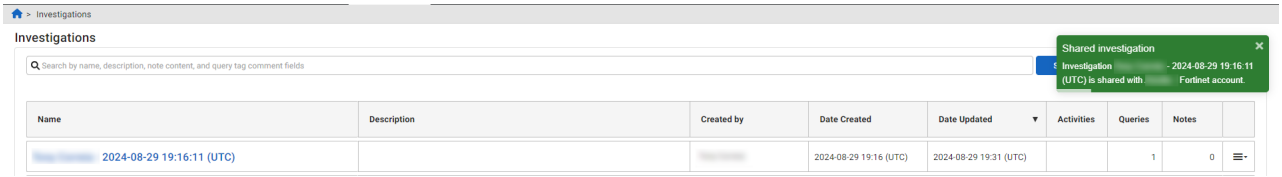

To share an investigation with the action menu:

1. Go to *Investigations*.
2. Locate an investigation in your primary account. Investigations in your primary account are indicated with an account badge.

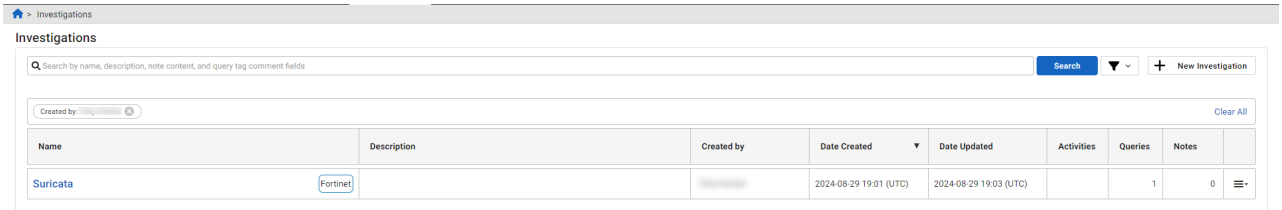

3. From the actions menu, select *Share*. The *Share with Account* dialog opens.

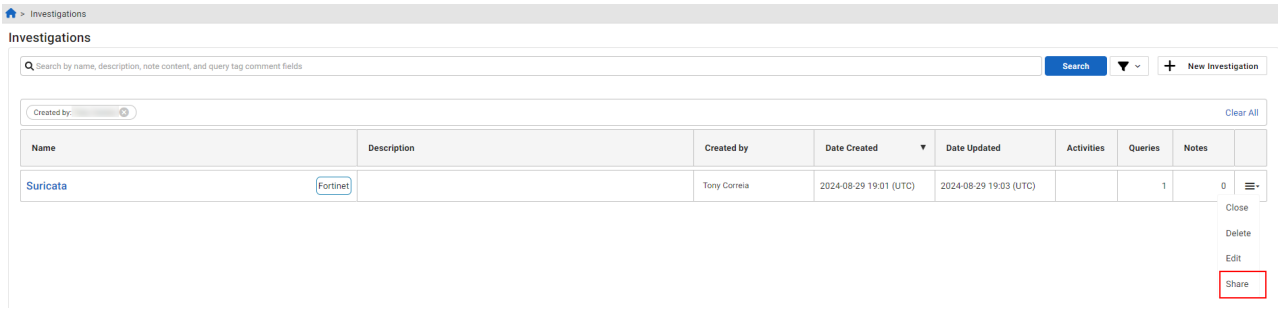

4. Click *Confirm*. The account icon is removed from the investigation and confirmation message appears at the top of the page.

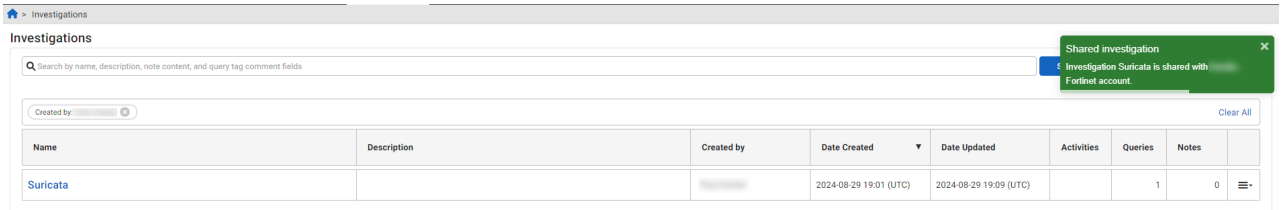

Packet capture

Packet capture tasks

Packet Capture tasks are defined and deployed on a per-sensor basis. A single task can be deployed to one, all, or any combination of sensors. Each sensor can spool up to four individual tasks, but only one task may run at a time.

The active task will execute for 60 minutes or until it captures 1 MB of data, whichever comes first. Once either of those conditions are met, the active task will pause, and the next spooled task will execute. The same task will begin again if it is the only one spooled. Tasks will continue to be spooled until they pass the specified expiration time or are terminated manually.

Packet capture tasks can have one of two states:

State	Description
Active	The task is currently in rotation for execution.
Inactive	The task has reached the requested end time or has been terminated by a user.

Packet capture tasks can be created, viewed, or terminated from the *Packet Capture* page. All tasks, both *Active* and *Inactive*, are displayed by default.

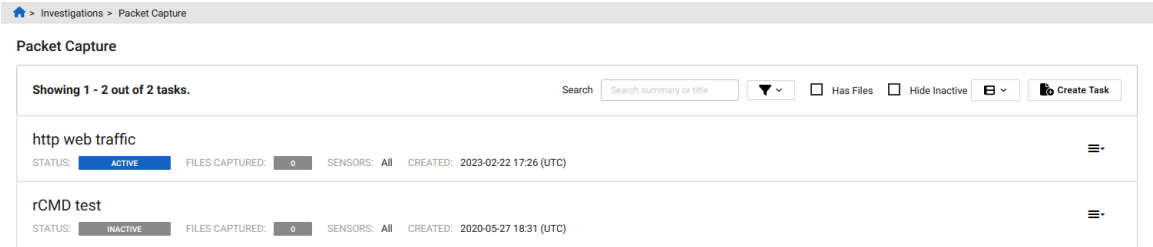

Reviewing a task

Click a task on the page to view metadata for the task and any PCAP data captured. Each execution of a task will produce exactly one log file and one PCAP.

- The log file will specify the start and end times of the respective execution .
- The PCAP file will contain any captured traffic.

The PCAP file will be empty if no traffic matched the BPF. Each file collected as part of the PCAP task can then be downloaded and viewed within WireShark or another preferred PCAP analysis tool. You can adjust which files are displayed (only PCAP, all PCAP, only non-empty PCAP) by checking or unchecking the respective options on the task page.

PCAP files are retained for 180 days. They can be deleted earlier by deleting the PCAP task.

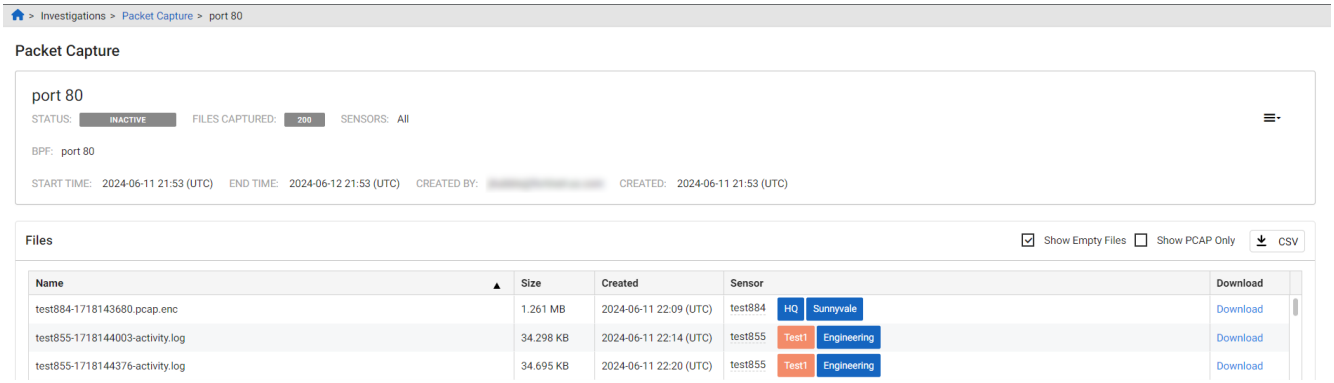

Creating a packet capture

To create a new task, the selected account should have one or more sensors with the PCAP feature enabled.

To create a packet capture task:

1. Go to *Investigations > Packet Capture*.
2. Click *Create Task*. The *Create New Packet Capture Task* window opens.

3. Configure the task settings.

Field	Required	Description
Title	Yes	The name of the task.
BPF	Yes	The BPF for traffic to match.
Date Range	Yes	The interval that the task will be active for, default = the next 24 hours.
Sensors	No	The sensors that the task will run on, default = All Sensors.
Description	No	A description of the task.

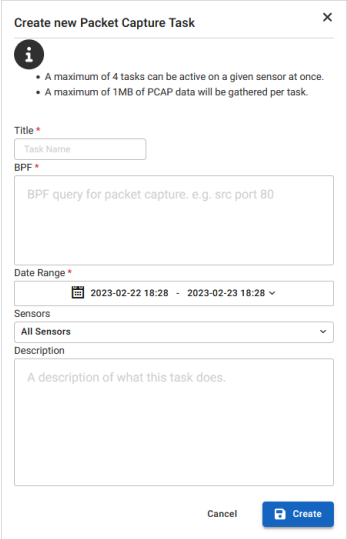

Sensors can only spool four (4) tasks at once, so only specify sensors that the task is relevant to. For example, if you are trying to troubleshoot one particular host in a particular data center, you probably only need to deploy the task to one sensor.

4. Click *Create*.

Terminating and deleting packet captures

To terminate a packet capture task:

1. Go to *Investigations* > *Packet Capture*.
2. Click the *Actions* menu at the right side of the task and click *Terminate Task*. A confirmation dialog opens.

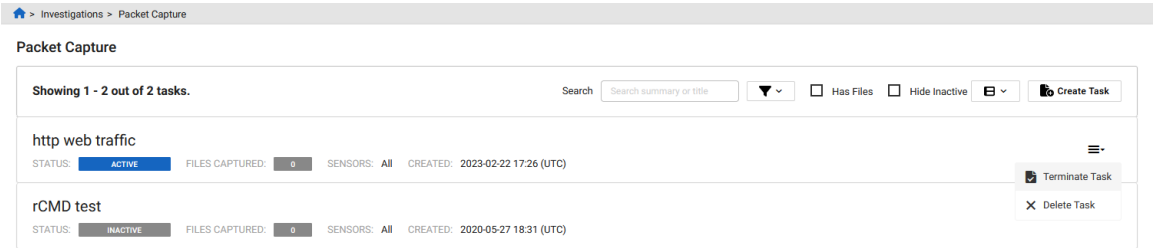

3. Click *Confirm*. The task changes to *Inactive*.

To delete a packet capture:

1. *Go to Investigate > Packet Capture.*
2. Click the *Actions* menu at the right side of the task and click *Delete*. A confirmation dialog opens.

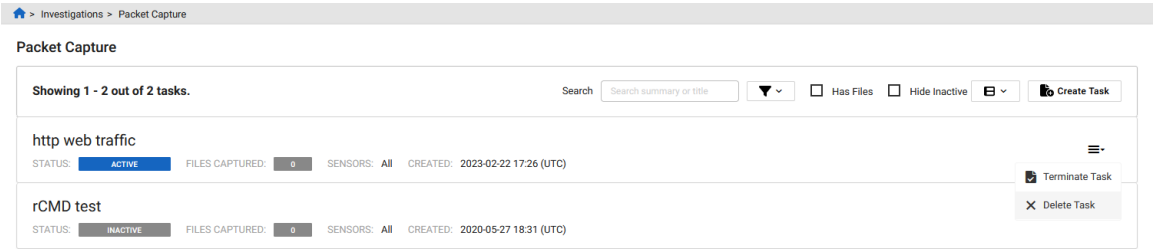

3. Click *Confirm*.

BPF resources

For in-depth information on Berkeley Packet Filters (BPFs), see The Linux Kernel Archives web site at https://www.kernel.org/. You can also download the BPF reference guide from here.

SYNTAX					
[Protocol] [Direction] [Type] {ip/subnet/port/portrange}					
PROTOCOL		DIRECTION		TYPE	
Limit the match to a specific protocol. If no protocol is supplied, all protocols consistent with the type are assumed.		*Transfer direction to and/or from the type. If no direction is supplied, 'src or dst' is assumed.*		*Type of entity, port, or range of ports. If no type is supplied, host is assumed.*	
ether	*ethernet*	src or dst (default)	*source or destination*	host (default)	*ip address*
fddi	*alias for ether*	src and dst	*source and destination*	net	*ip address or subnet*
icmp	*internet control message protocol*	src	*source only*	port	*tcp/udp port number*
wlan	*wireless lan; alias for ether*	dst	*destination only*	portrange	*range of tcp/udp ports (xxxx-xxxx)*
ip	*ipv4*	[proto] broadcast	proto must be ip or ether		
ip6	*ipv6*	OPERATORS			
arp	*address resolution protocol*	'='	*equal to*	'||' 'or'	*logical or*
tcp	*transmission control protocol*	'!' or 'not'	*not equal to*	'<' 'less'	*less than*
udp	*user datagram protocol*	'&&' 'and'	*logical and*	'>' 'greater'	*greater than*

COMMON EXPRESSIONS	
host xxx.xxx.xxx.xxx	*all packets to/from a host*
src host xxx.xxx.xxx.xxx && dst host xxx.xxx.xxx.xxx	*all packets from a source host to a destination host*
dst port 23	*all packets to port 23 (telnet)*
udp src net xxx.xxx.xxx && dst host xxx.xxx.xxx.xxx	*only udp packets from a dotted pair subnet to destination host*
ip6 && not net xxx.xxx.xxx	*only IPv6 packets outside of a dotted triple subnet*
src host xxx.xxx.xxx.xxx && (dst portrange xxxx-xxxx && dst net xxx.xxx.xxx)	*all packets from a source host to a destination port range in a dotted triple subnet*
dst portrange 49152-65535 && gateway xxx.xxx.xxx.xxx	*all packets to non-standard ports on a gateway*
host xxx.xxx.xxx.xxx || host xxx.xxx.xxx.xxx	*all packets to/from host A or host B*

BYTE LEVEL FILTERING	
ip[9]!=47	*all packets where IP protocol field is GRE (tunnel)*
ip[8]<64	*all packets where IP time-to-live (TTL) is less than 64*
icmp[0]=3	*all packets with ICMP message type 3 (destination unreachable)*
tcp[13]=32 || tcp[13]=8	*all packets with TCP flags set to PSH or URG*

HOW TO READ PACKET HEADERS

Word 0																															
Byte Offset 0								Byte Offset 1								Byte Offset 2								Byte Offset 3							
Nibble 0				Nibble 1				Nibble 2				Nibble 3				Nibble 4				Nibble 5				Nibble 6				Nibble 7			
BIT 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

TCP HEADER – RFC 793

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
20 BYTES	Offset 0								Offset 1								Offset 2								Offset 3							
	Source Port Number																Destination Port Number															
	Offset 4								Offset 5								Offset 6								Offset 7							
	Sequence Number																															
	Offset 8								Offset 9								Offset 10								Offset 11							
	Acknowledgement Number																															
	Offset 12								Offset 13								Offset 14								Offset 15							
	Hacker Length				Reserved				CWR	ECE	URG	ACK	PSH	RST	SYN	FIN	Window Size															
	Offset 16								Offset 17								Offset 18								Offset 19							
	Checksum																Urgent Pointer															
VARIABLE	Offset 20								Offset 21								Offset 22								Offset 23							
	TCP Options																															
	Data																															

UDP HEADER – RFC 768

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
8 BYTES	Offset 0								Offset 1								Offset 2								Offset 3							
	Source Port Number																Destination Port Number															
	Offset 4								Offset 5								Offset 6								Offset 7							
	Length																Checksum															
VAR	Offset 8								Offset 9								Offset 10								Offset 11							
	Data																															

ICMP HEADER – RFC 792

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Offset 0								Offset 1								Offset 2								Offset 3							
Message Type								Message Code								Checksum															
Offset 4								Offset 5								Offset 6								Offset 7							
(Variable Contents Depending on Type and Code)																															

VAR × 4 BYTES

IPv4 HEADER – RFC 791

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Offset 0								Offset 1								Offset 2								Offset 3							
Version				IP Header Length				Type of Service								Total Length (in Offsets)															
Offset 4								Offset 5								Offset 6								Offset 7							
IP Identification Number																x	D	M	Fragment Offset												
Offset 8								Offset 9								Offset 10								Offset 11							
Time to Live (TTL)								Protocol								Header Checksum															
Offset 12								Offset 13								Offset 14								Offset 15							
Source IP Address																															
Offset 16								Offset 17								Offset 18								Offset 19							
Destination IP Address																															
Offset 20								Offset 21								Offset 22								Offset 23							
IP Options																															
Data																															

20 BYTES

FLAGS
x = Reserved D = Do Not Fragment M = More Fragments Follow

IPv6 HEADER – RFC 2460

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Offset 0								Offset 1								Offset 2								Offset 3							
Version				Traffic Class								Flow Label																			
Offset 4								Offset 5								Offset 6								Offset 7							
Payload Length																Next Header								Hop Limit							
Offset 8								Offset 9								Offset 10								Offset 11							
Source IP Address																															
Offset 12								Offset 13								Offset 14								Offset 15							
Source IP Address (continued)																															
Offset 16								Offset 17								Offset 18								Offset 19							
Source IP Address (continued)																															
Offset 20								Offset 21								Offset 22								Offset 23							
Source IP Address (continued)																															
Offset 24								Offset 25								Offset 26								Offset 27							
Destination IP Address																															
Offset 28								Offset 29								Offset 30								Offset 31							
Destination IP Address (continued)																															
Offset 32								Offset 33								Offset 34								Offset 35							
Destination IP Address (continued)																															
Offset 36								Offset 37								Offset 38								Offset 39							
Destination IP Address (continued)																															
Offset 40								Offset 41								Offset 42								Offset 43							
Net Header								Extension Header Information																							
Extension Header																															
Data																															

40 BYTES

VARIABLE

PCAP encryption

FortiNDR Cloud requires the encryption of all PCAP data captured and stored on the platform, backed by public key cryptography. Adding a PEM-encoded RSA key to an account on the Account management page will enable this feature.

Activation of the PCAP encryption feature prevents FortiNDR Cloud analysts from reviewing the contents of any captured packet data, and renders that data unrecoverable should the private key associated with the uploaded public key be lost.

Generating a key

Be sure to only upload the contents of the `public.pem` file and keep the `private.pem` file safe. In the event that `private.pem` is lost, FortiNDR Cloud is unable to recover either it or the contents of any PCAP encrypted with the matching public key

For instructions on how to upload the generated public key, see the Settings on page 124 page.

Windows

To generate a key pair on Windows, we recommended using the PCAPUtil program. You can download the binary here or from the *Settings* inAccount management on page 145 .

You must be logged in to FortiNDR Cloud to download the binary.

Generate a key pair with files named `public.pem` (public key) and `private.pem` (private key) in the current directory. PCAPUtil supports overriding all file names and locations via command line arguments.

```
bash
pcaputil generate
```

macOS and Linux

Generate a public/private key pair using the built-in OpenSSL library.

```
bash
openssl genrsa -out private.pem 4096
openssl rsa -in private.pem -outform PEM -pubout -out public.pem
```

Decrypting a PCAP

Unencrypted PCAP files are denoted with an extension of `.pcap`, and encrypted PCAP files are denoted with the extension `.pcap.enc`.

Windows

Encrypted PCAP files can be decrypted with the FortiNDR CloudPCAPUtil binary.

You must be logged in to FortiNDR Cloud to access this file.

```
pcaputil decrypt -private private.pem -src sen1-1502499443.pcap.enc -dst sen1-1502499443.pcap
```

macOS and Linux

Use the following script to extract and decrypt the PCAP:

```
#!/usr/bin/env bash
show_help () {
echo "Usage: $0 private_key encrypted_pcap decrypted_pcap"
}
if [ -z $3 ]; then
show_help
exit 0
fi
tar zxf $2
openssl pkeyutl -decrypt -inkey $1 -in session.key.enc -out session.key
#openssl rsautl -decrypt -inkey $1 -in session.key.enc -out session.key
key=$(xxd -p -c 96 session.key | cut -c 1-64)
iv=$(xxd -p -c 96 session.key | cut -c 65-96)
openssl enc -aes-256-cbc -d -in data -out $3 -nosalt -K $key -iv $iv
rm data
rm session.key
rm session.key.enc
```

Managing encryption keys

Any PCAP captured and stored in FortiNDR Cloud will be encrypted by adding the associated keys to the account.

FortiNDR Cloud requires the encryption of all PCAP data captured and stored on the platform, backed by public key cryptography.

Encryption key requirement impact on existing sensors

If you do not have a PCAP-enabled sensor	The encryption key will be required to enable PCAP on sensors
If you have a PCAP-enabled sensor	• There is no change in behavior for existing PCAP-enabled sensors. • After the encryption key is provided, the PCAP-enabled sensor will upload encrypted PCAP files. • For existing PCAP-enabled sensors that are capturing without a key, you should still be able to disable them without a key. • Encryption keys can be updated directly without needing to delete an existing key. Existing behaviors and PCAP-enabled sensors will not be impacted.
When deleting the encryption key	• PCAP will be disabled on all the sensors for this account. • All PCAP upload requests for those sensors will be silently ignored. • When the encryption key is provided again after it's been deleted, you will need to enable PCAP on the sensor manually.

Enabling PCAP on a sensor requires encryption

When enabling PCAP on an individual sensor, the *PCAP Enabled* option is disabled unless you have encryption enabled and display a note advising that you must enable encryption before enabling PCAP.

Warning appears on Sensor Update dialog accessed from the list of sensors:

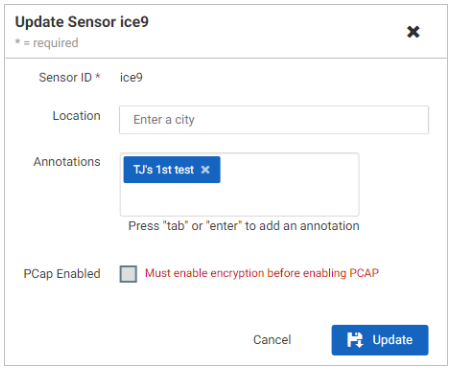

Warning appears on the detailed Sensor Settings page:

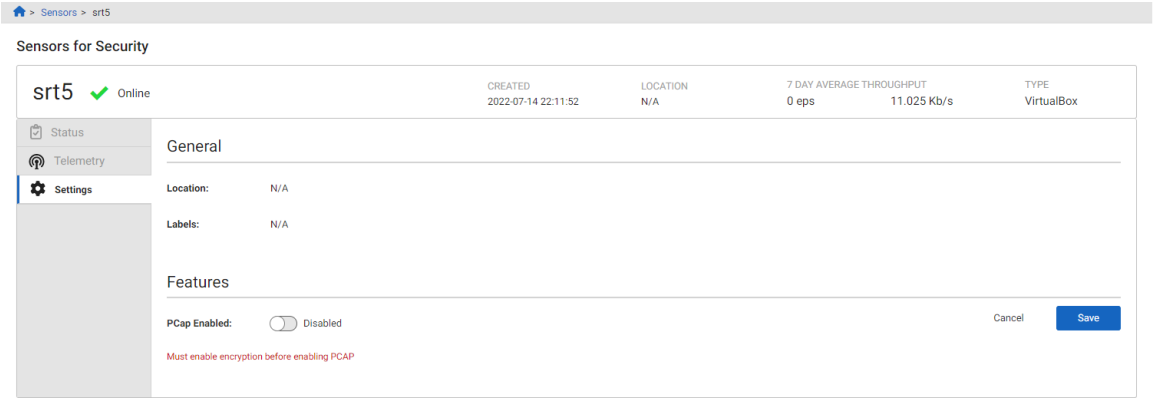

Deleting a PCAP encryption key

When deleting a PCAP key for an account, a warning will appear advising that PCAP will be disabled for sensors associated with that account.

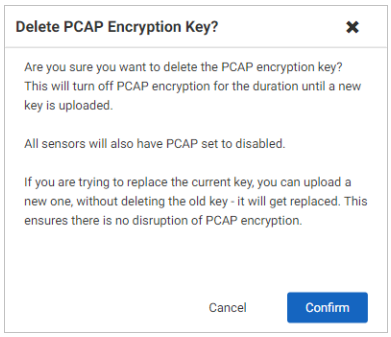

Click *Confirm* to acknowledge the message and proceed.

Encryption key settings

To access PCAP Encryption Keys settings:

1. Click on the gear icon on the top right and select *Account Management*.
2. Select an account.

3. On the left navigation, select *Settings*.

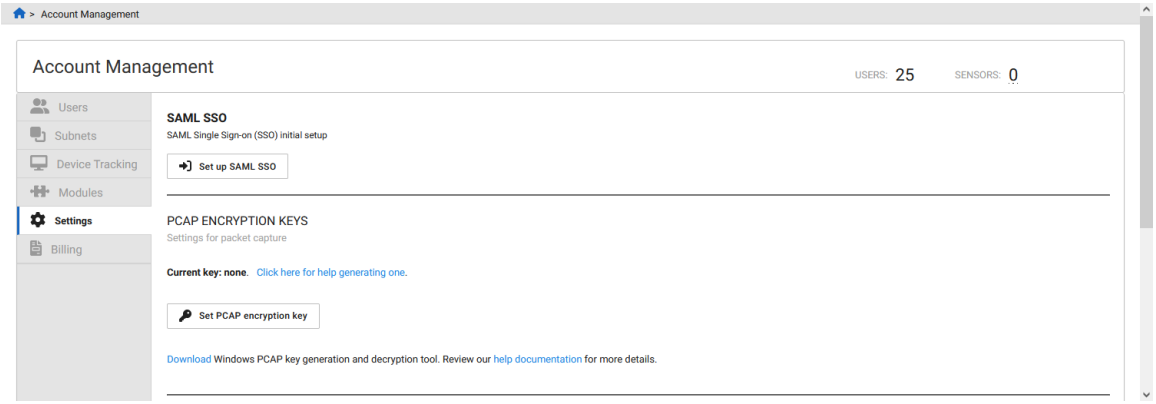

The *Set PCAP encryption key* button will only appear for the Admin role.

Private search

The *Private Search* page shows the history of Adhoc queries. Use this page to view the query status, past query results, delete query, and create detection out of the selected Adhoc query.

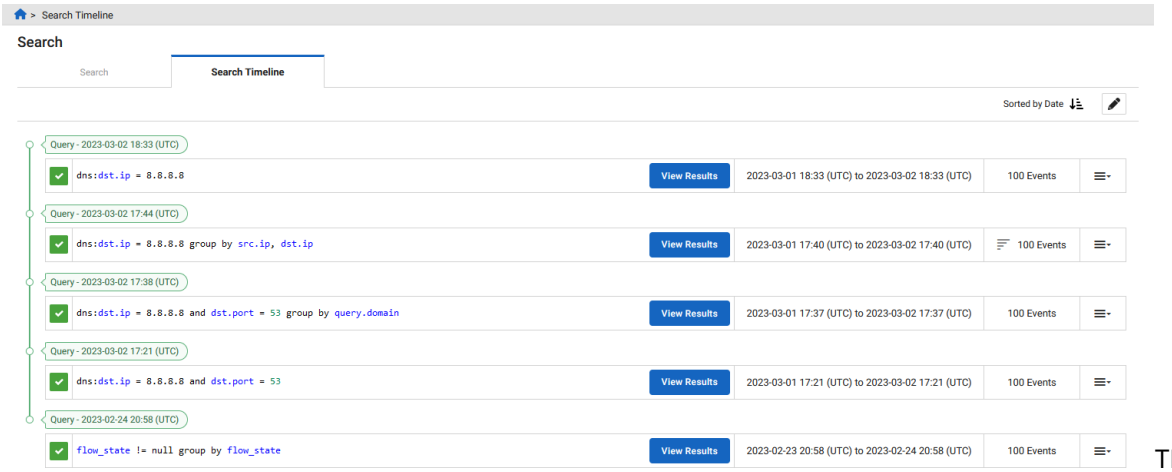

The Search tab contains example queries of topics such as Flow, DNS, X.509, RDP, HTTP, SSH, SMTP, FTP, SSL, Kerberos,

SMB, NTLM, DCE-RPC and PE are added. You can click any of the example queries, modify them, and then perform the search operation.

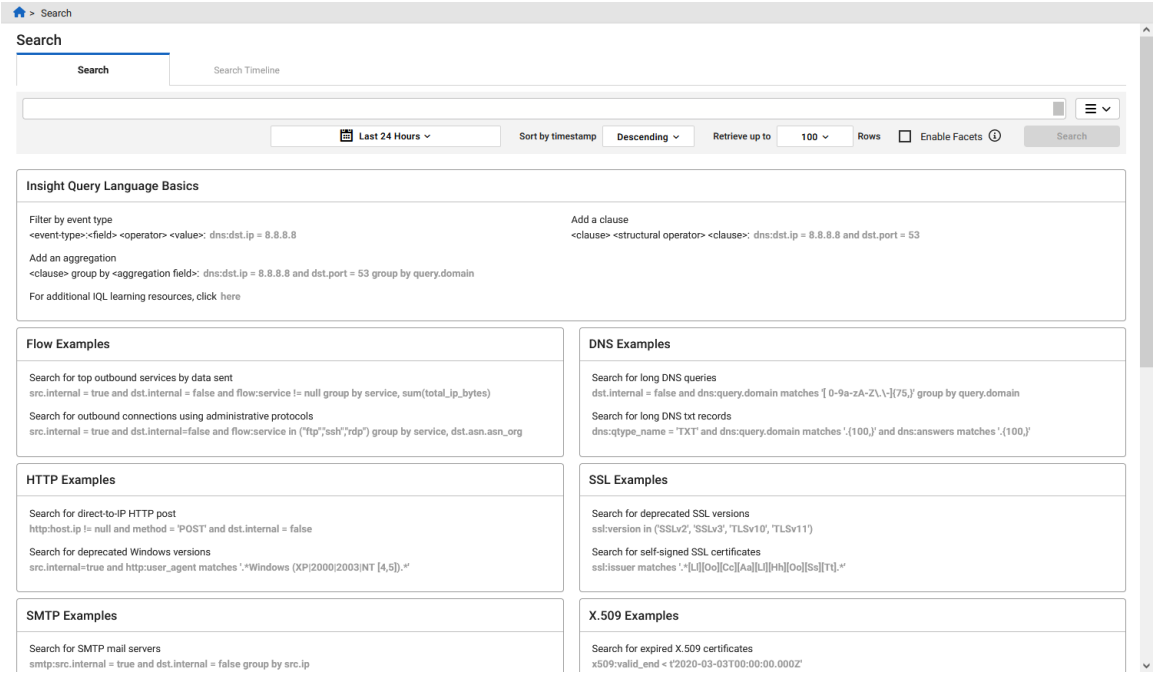

Creating IQL queries in Private Search

IQL queries provide a high level of precision and control over investigations. They support detailed filtering, complex logic, and advanced operators, making them ideal for handling large datasets and intricate conditions. IQL queries can be saved, reused, added to investigations, and used to create detectors.

FortiNDR Cloud also supports Natural Language (NL) queries that allow analysts to write queries in plain language without learning specialized syntax. For more information, see Natural Language queries on page 277

To perform a Private Search with IQL queries:

1. Go to *Investigations > Private Search*.
2. In the *Search* tab, enter the query string in the search field. For example queries see, IQL query examples on page 272.
3. Configure the search settings.

Date range	Use the date picker to configure the date range or select *Last Hour*, *Last 24 Hours*, or *Last 7 days* and click *Apply*. You can select any time period within the last 365 days as long as it is limited to seven days.
Sort by timestamp	Select *Ascending* or *Descending*.

Retrieve up to *xxx* Rows	Select *100*, *500* or *1,000* rows.
Enable Facets	Select to return the panel that allows narrowing the search. This may make the query longer to complete. For more information, see Facet Search on page 90.

4. Click *Search*.

To move Private Search queries to Investigations:

1. Click Investigations > Private Search.
2. Click the *Private Search* tab.

To move a query	Click the *Actions* menu ≡ at the end of the row and select *Move to an Investigation*.
To move multiple queries	1. Click the *Edit* icon ✎ and select the queries to be moved. 2. Click *Actions* > *Move to an Investigation* .

3. In the *Move Query to Investigation* dialog, create a new investigation or add the query to an existing investigation.

Create a New Investigation	Select this option to create a new investigation. Enter the *Investigation Name* and *Description*. The default name for new investigations is the first and last name of the user creating the investigation as well as a date stamp of when the investigation was created.
Add to Existing Investigation	From the *Choose Investigation* dropdown, select and investigation.

4. Click *Move Query*.

To delete queries in the Private Search tab:

1. Click *Investigations* > *Private Search*.
2. Click the *Private Search* tab.

To delete a query	Click the *Actions* menu ≡ at the end of the row and select *Delete Query*.
To delete multiple queries	1. Click the Edit icon ✎ and select the queries to be deleted. 2. Click *Actions* > *Delete Query* .

3. In the confirmation dialog, click *Confirm*.

To create a detection from an adhoc query:

1. Click the *Private Search* tab.
2. Click the *Actions* ≡ menu at the end of the row and click *Create Detection*.The *Create A Detector* page opens.

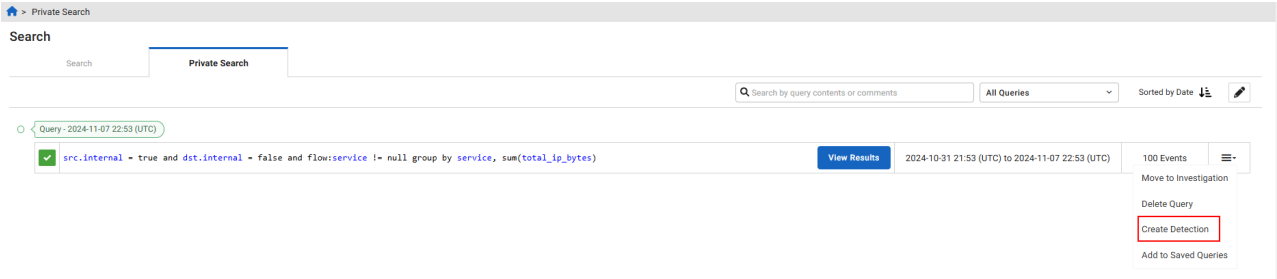

3. Configure the detector and click *Save Detector*.

Detector Query	You have the option of selecting a new query or using the query parameters the results are based on. • The query field displays the facet filters used in the query. • Click *Select a new Query* to select a saved query or a query from your history.
Impacted Device IP can appear in the fields	Click *Change Fields* to select the specific fields you want to use to generate a detection. By default, any internal IP address in the *src.ip* or *dst.ip* fields will be used to generate detections.
Indicators are captured in the fields	Click *Change Fields* to add or remove an Indicator Field for a detector. You can choose up to five fields.
Name	Enter a name for the detector query.
Severity	Select *High*, *Moderate* or *Low* from the dropdown list.
Confidence	Select *High*, *Moderate* or *Low* from the dropdown list.
Category	Select the detector category from the dropdown list.
Primary Technique	Select the Primary Technique from the dropdown list.
Secondary Technique	Select the Secondary Technique from the dropdown list.
Specificity	Select *Campaign*, *Tool Implementation*, *Procedure*, *Technique*, or *Tactic* from the dropdown.
Description	Enter a description of the new detector.
Run on Accounts	Click *Manage Run List* to choose which accounts the new detector should run in. In the dialog that opens, choose an account and click *Save*. This is applicable only if you have access to multiple accounts. For example, if your organization acquired another organization, once you deploy sensors in their network, it might be easier to ingest that data into a separate account and give your team access to it. If you were to write a detector targeting specific subnets in your account, that detector wouldn't be applicable to the acquired company's network, so you would only want to deploy it in your account.
Data Sources	Enable *Zeek*, *Fortinet*, *Suricata*, or *Zscaler*.
Resolution Settings	**Resolution Style** Select *Auto* or *Manual*.

Automatic Resolution Period	Select between *6 hours* and *1 Month*. The default is *1 Week*.

To save a query:

1. Click the *Private Search* tab.
2. Click the *Actions* menu at the end of the row and click *Add to Saved Queries*.The *Save Query* dialog opens.
3. Enter the query details and click *Save*.

Query Name	Enter a name for the query.
Search Query	This field cannot be edited.
Description	Enter a description of the query.

You can use a saved query when you create a new detector or investigation.

Guided queries

Use *Guided Queries* to start a new investigation, add queries to expand upon an existing one, or run event queries. The pre-defined queries on this page have been created by FortiGuard Labs with a focus on identifying potential security vulnerabilities or suspicious activities within a network.

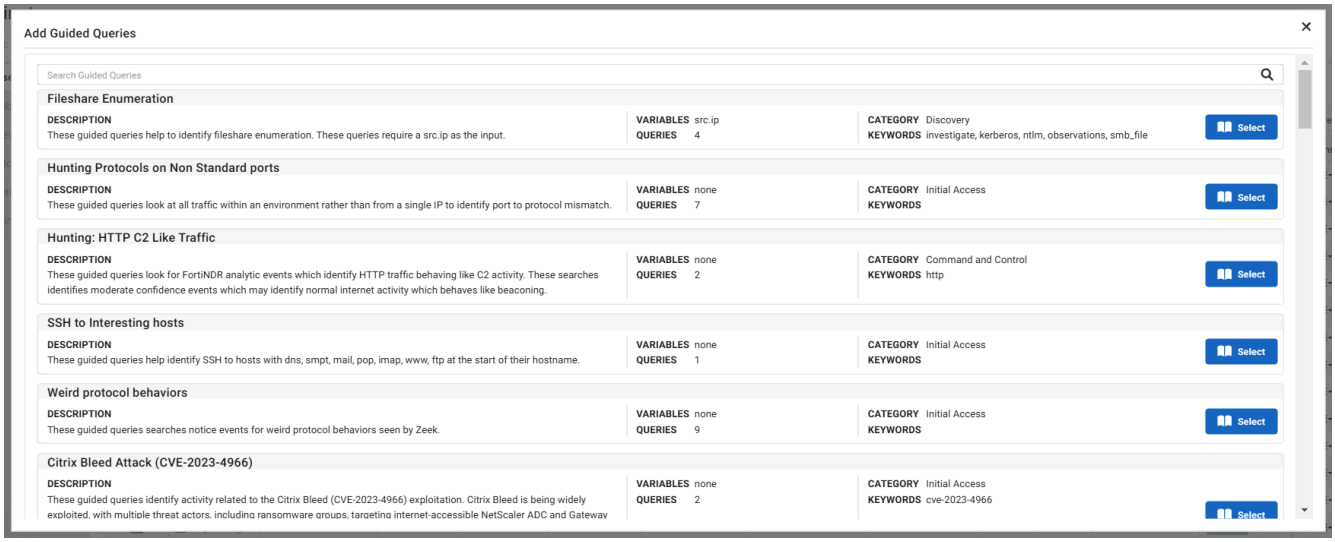

To run a guided query:

1. Go to *Investigations > Guided Queries*.
2. Scroll through the list of guided queries, or use the search field to find a query by keyword. Click *Select*. The query details page opens.

   If this is your first query, we suggest running the query named *Example Hunt* to start.

3. Configure the query settings:

Date range	Use the date picker to configure the date range.
Enable Facets	Select to return the panel that allows narrowing the search. This may make the query longer to complete. For more information, see Facet Search on page 90.
Variables	Enter the required variable(s) for the queries. Multiple variables are supported. Values can be entered either as: • Individual items, followed by the tab or enter key. The value appears as a pill that can then be deleted, if required. • *Bulk indicator* icon. This brings up an entry screen. Pasting the text is supported. After pressing the button, FortiNDR Cloud extracts the applicable indicators from the text and adds them as variables. You can also delete the unneeded variables.
Create a New Investigation	Select this option to create a new investigation. Enter the *Investigation Name* and *Description*. The default name for new investigations is the first and last name of the user creating the investigation as well as a date stamp of when the investigation was created.
Add to Existing Investigation	From the *Choose Investigation* dropdown, select and investigation.

   Not all guided queries use variables.

4. (Optional) In the *Investigation Name* field, enter a unique name for the query.
5. Click *Run Guided Queries*. The query starts to run.
6. After the query has run, go to *Investigations* and click the query name. The Investigation details page opens.
7. Click *View Results*. The query results are displayed.

Adding a guided query to an investigation

To add a guided query to an investigation:

1. Go to *Investigations* > *Investigate* and open an investigation in the list.
2. Click the *Add Guided Queries* button. Alternatively, click on Add menu (**+**) in the top-right corner of the page and select *Add Guided Queries*. The *Add Guided Queries* page opens.
3. Click *Select*.
4. Configure the query settings.

Date range	Use the date picker to configure the date range.
Enable Facets	Select to return the panel that allows narrowing the search. This may make the query longer to complete. For more information, see Facet Search on page 90.
Variables	Enter the required variable(s) for the queries. Multiple variables are supported. Values can be entered either as: • Individual items, followed by the tab or enter key. The value appears as a pill that can then be deleted, if required. • *Bulk indicator* icon. This brings up an entry screen. Pasting the text is supported. After pressing the button, FortiNDR Cloud extracts the applicable indicators from the text and adds them as variables. You can also delete the unneeded variables.
Create a New Investigation	Select this option to create a new investigation. Enter the *Investigation Name* and *Description*. The default name for new investigations is the first and last name of the user creating the investigation as well as a date stamp of when the investigation was created.
Add to Existing Investigation	From the *Choose Investigation* dropdown, select and investigation.

5. Click *Run Guided Query*.

Running a guided query of event records

Query event records to retrieve specific information from event logs during an investigation.

Query event records

To run a guided query of event records:

1. Go to *Investigations* > *Investigate* and select an investigation from the list.
2. Click *View Results* to view the investigation results.
3. Right click on an entity to open the context menu and select *Guided Queries*.

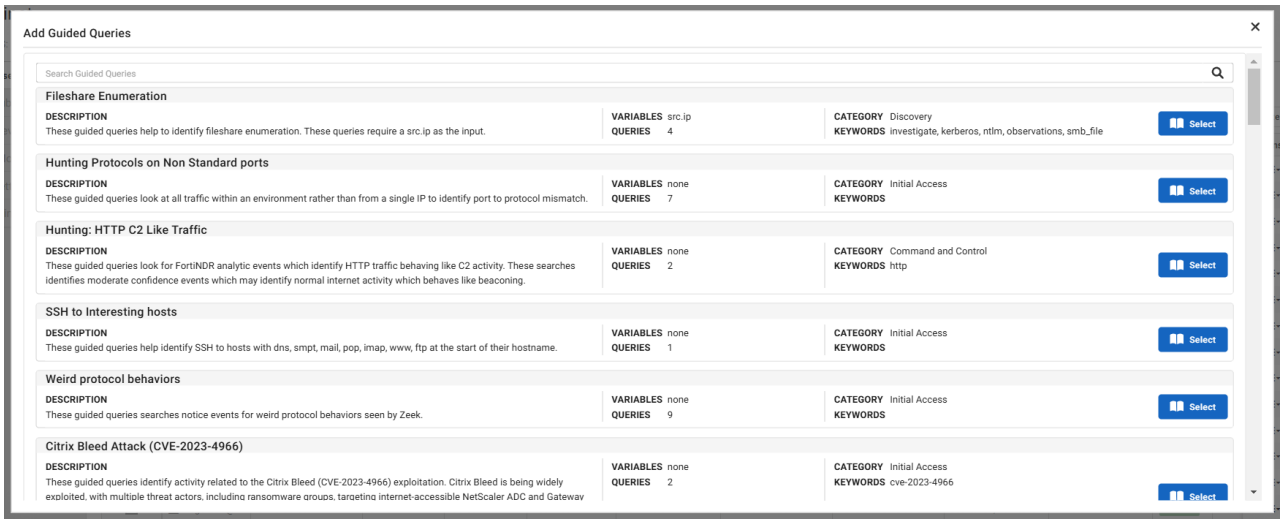

4. Select a guided query from the list. If the event record has matching variables in the query , then the variables will be populated with values from the event record.

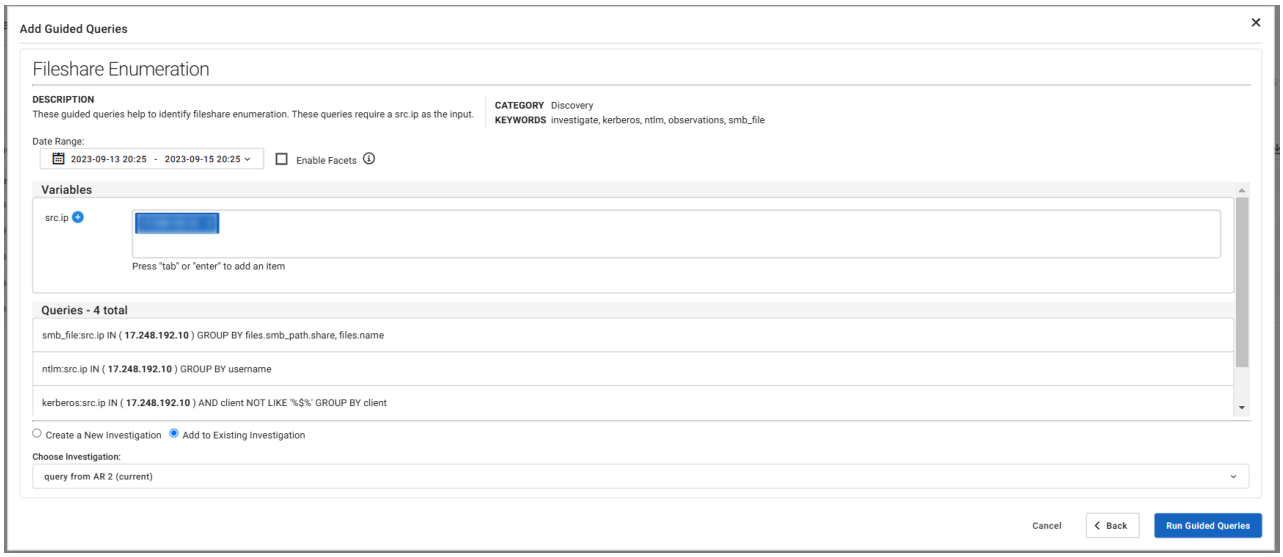

5. Add or modify the values for the variables.
6. Create a new investigation or add the guided query to an investigation.

Create a New Investigation	Select this option to create a new investigation. Enter the *Investigation Name* and *Description*. The default name for new investigations is the first and last name of the user creating the investigation as well as a date stamp of when the investigation was created.
Add to Existing Investigation	From the *Choose Investigation* dropdown, select and investigation.

7. Click *Run Guided Queries*.

Threat intelligence

FortiNDR Cloud ingests threat intelligence from a wide variety of sources, including commercially purchased feeds, open source threat intelligence data, vertical/industry/government information sharing organizations, and closed trust-based communities. This threat intelligence is reviewed and curated by the Fortinet FortiGuard Labs team, and allows for real-time matching of network traffic against known indicators.

Events are enriched with ingested threat intelligence by matching indicators from the data to entities within an event. All matched intel records are contained within the `intel` field, which is a common field across all event types. The intel records are then searchable with IQL.

Contact your TSM if you have access to an intel source or feed that you would like integrated with FortiNDR Cloud.

Example query:

The following query is a simple way to determine whether or not network traffic has matched with threat intelligence data in your network. When the results load, you will notice the `intel` column shows whether or not an event has a match against a threat intelligence source.

```
// show events that have at least one matched intel record
```

```
intel.indicator != null
```

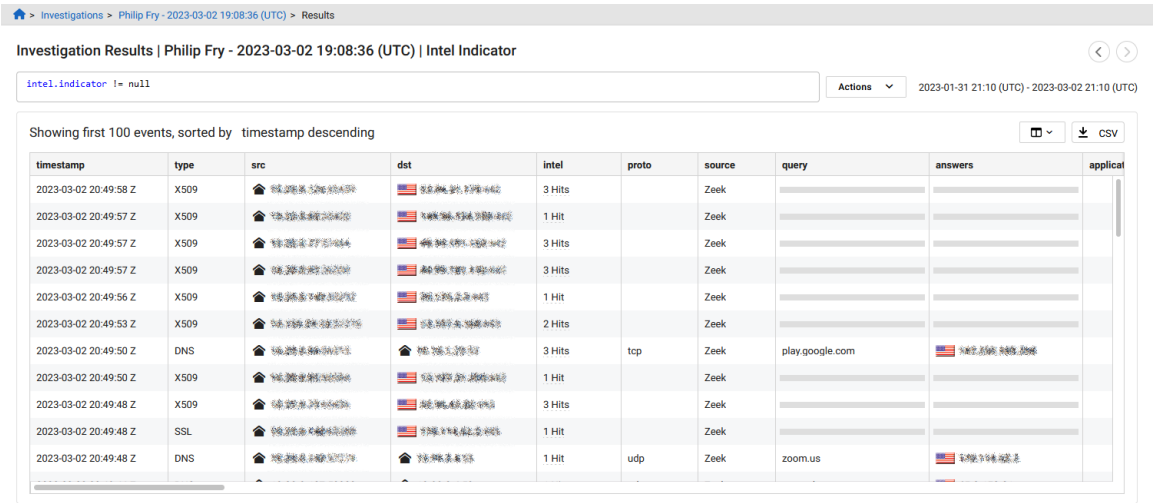

Click the number of *hits* in the *Intel* column to view the matched *intel* records.

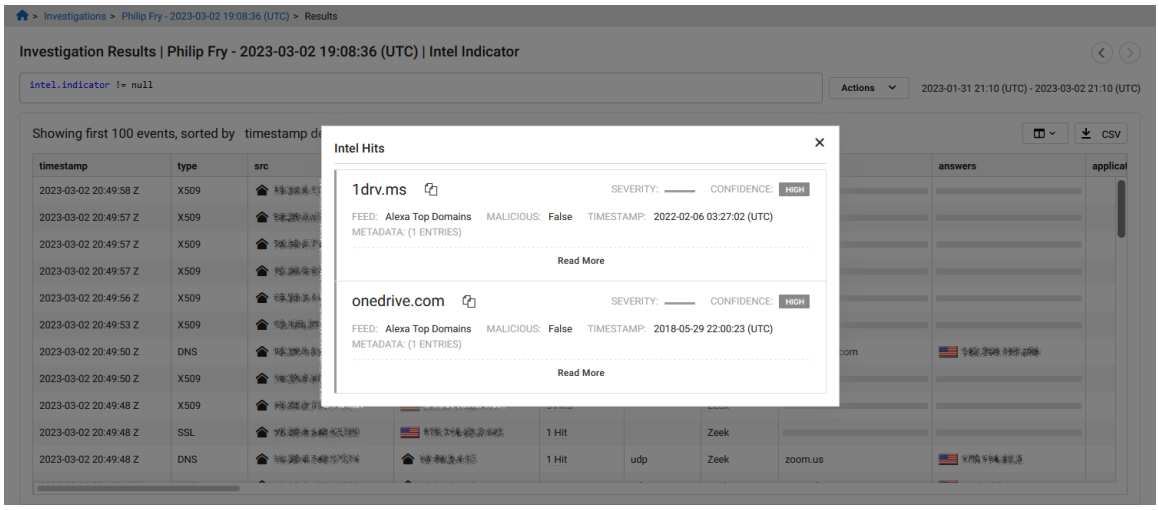

Search for intel

The `intel` field is an array of *intel-objects*, meaning there could be multiple records for a given event. When a query is applied to an event with multiple intel records, the values for each field are flattened into individual arrays before the query logic is applied to the values.

The following table lists the fields contain in *intel-objects*:

Field	Type	Description	Example
`confidence`	String	The overall confidence rating of the intel source	`high`
`feed`	String	The name of the intel source	`Sinkholes`
`indicator`	String	The matched entity	`131.253.18.12`
`indicator_type`	String	The entity type	`ip_address`
`is_malicious`	Boolean	Indicates whether the indicator is believed to be malicious	`false`
`meta`	String	A JSON string of all metadata provided by the intel source	`{"description":"Observed C2 Activity","references": ["Fortinet FortiGuard Labs"]}`
`severity`	String	The overall severity rating of the intel source	`high`
`timestamp`	Timestamp	The creation time of the intel record	`2019-01-01T00:00:00.000Z`

Example search for intel

In this example, we will create two queries to search for the following events:

- **Event 1**: `[{confidence: high, severity: low}, {confidence: low, severity: high}]`
- **Event 2**: `[{confidence: high, severity: high}, {confidence: low, severity: low}]`

Example 1:

In this example we will use a query to compare an array of records in *Event 1* and *Event 2*.

Query string:

```
intel.confidence = high & intel.severity = high
```

What the query will do:

1. The two records are flattened into arrays of values for each field, so the query logic is applied to all values all at once and not to records individually.
2. The query is compared to the array of records in *Event 1* and *Event 2*.

Response:

This query will return Event 1 and 2 because at least one inner object contains `confidence=high` and at least one inner object contains `severity=high`.

- Event 1: `confidence =[high,low]` and `severity = [high,low]`
- Event 2: `confidence =[high,high]` and `severity = [high,low]`

Example 2:

In this example, we will create a query to match individual objects of a nested field (such as intel, path, files, etc.).

Query string:

```
intel {confidence=high & severity=high}
```

Response:

This query will only return Event 2 because at least one of the objects in the event meets both criteria.

- Event 2: `confidence =[high,high]` and `severity = [low,low]`

Reports

The following reports are available: *FortiNDR Cloud Network Security Posture Report*, *FortiNDR Cloud Network Traffic Usage Report*, *FortiNDR Cloud Network Traffic Usage of a Sensor Report* and *FortiNDR Cloud Detections Report*.

Generating reports

To generate a report:

1. From the top navigation, select *Reports*. The *Reports* page opens.
2. Select the date range and click *Apply*.
3. Click *Run Report*.The browser will transition from the template list to the report page while retrieving data to complete the report. Each section will update individually as data is retrieved. Sections will appear as data is ready.
4. Click *Print*. The *Print* dialog opens.
5. Click *Save*. Select a location to save your report and click *Save* again.

FortiNDR Cloud Network Traffic Usage of a Sensor Report

This report provides daily insights into network traffic patterns for an individual sensor by identifying top source and destination IPs, high-volume IP pairs (Top Talkers), busiest destination ports, and ports with unidentified protocols.

FortiNDR Cloud Network Traffic Usage Report

The *Network Traffic Usage* report provides a comprehensive view of traffic distribution across multiple dimensions. It includes visualizations such as bar charts, pie charts, and detailed data tables to help analyze network activity. A Sankey chart is available for aggregations where *Group By* includes two IP fields or when there are two dimensions and a measure. The Sankey chart type appears only when there are 50 or fewer dimensions.

Traffic usage can be grouped by source and destination IP addresses, applications, protocols, and other relevant categories. Each chart displays comparative usage trends, while the accompanying data tables show

exact values and percentages for each category. This report enables analysts to quickly identify high-traffic sources, dominant applications, and overall usage patterns for effective monitoring and optimization.

FortiNDR Cloud Detections Report

The *FortiNDR Cloud Detections Report* is useful for analysis and threat hunting. It provides key security performance metrics, focusing exclusively on resolved detections. This approach ensures that analysts can accurately measure the effectiveness of their detection and response processes.

This report provides an overview of the number detections within a specific time range and can be useful for threat hunting. The report includes only resolved detections when calculating metrics. Any active detections within the selected time range are excluded from these calculations.

When calculating results, FortiNDR Cloud applies the following filters:

- *Created Date* within the selected time range
- *Muted* and *Disabled* state set to *All*
- *Detection Status* set to *Resolved*

The report criteria includes detections observed, attack category and severity. For each detection there is an overview and the number of events that satisfy the detector. The Executive Summary displays:

Total Detections	Number of detections within the specified time range.
Devices with detections	Number of devices with detections within the specified time range.
Mean Time to Detect (MTTD)	Average time in seconds between when an incident was first seen and when it was created in the system. *Mean Time to Detect* is calculated by averaging the time difference (in seconds) between the *FirstSeen* and *Created* timestamps for all detections with a status of *Resolved*.
Mean Time to Resolve (MTTR)	Average time in seconds between when an incident was created and when it was resolved. *Mean Time to Resolve* is calculated by averaging the time (in seconds) between the *Created* timestamp and the *ResolutionTimestamp* for all detections marked as *Resolved*.
Mean Dwell Time (Dwell)	Average time in seconds between when an incident was first seen and when it was resolved. Dwell Time is calculated by averaging the time (in seconds) between the *FirstSeen* and *ResolutionTimestamp* for all detections with a status of *Resolved*.
Devices with Detections	Total of unique device IP from all detections

FortiNDR Cloud Network Security Posture Report

This report analyzes 10 aspects related to your overall security posture. This report allows you to view an investigation and the results for an event. The report also provides a list a of generated reports in the *Report History*. Please a allow a few minutes for the report to generate. To run the report, select one or more sensors from the dropdown menu, set the time range, and then click *Run Report*.

You can navigate away from the Posture Report page after clicking *Run Report*. However, the report will remain incomplete indefinitely if you close your browser tab or log out of the portal. When this occurs, the following error message is displayed: *The report is incomplete. Please run it again*.

Report history

The *Report History* panel in the *FortiNDR Cloud Network Security Posture Report* tile, displays a log or previous reports by time range. You can click a range in the list to regenerate the report.

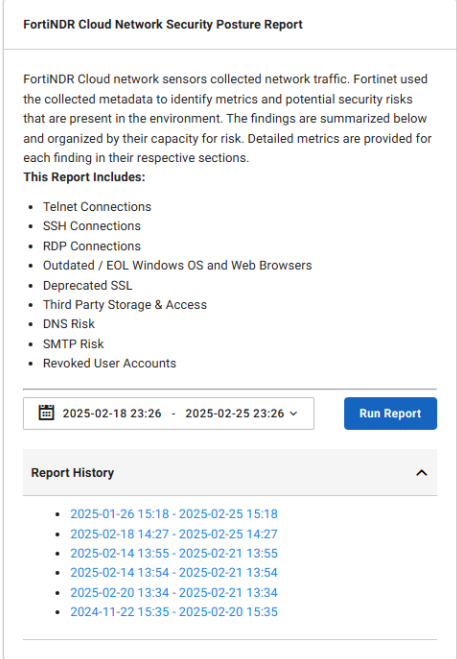

View investigations

After the report is generated, click the *View Investigations* button in the report to view the investigation in *Read-Only* mode.

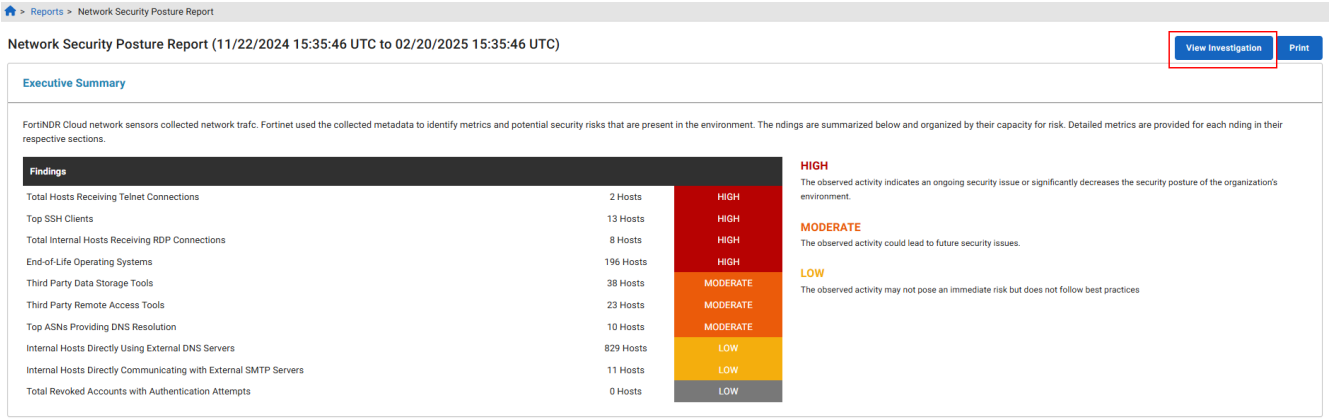

In the report, click *View Results* to view individual results for an event, or click *Show Report* to return to the report.

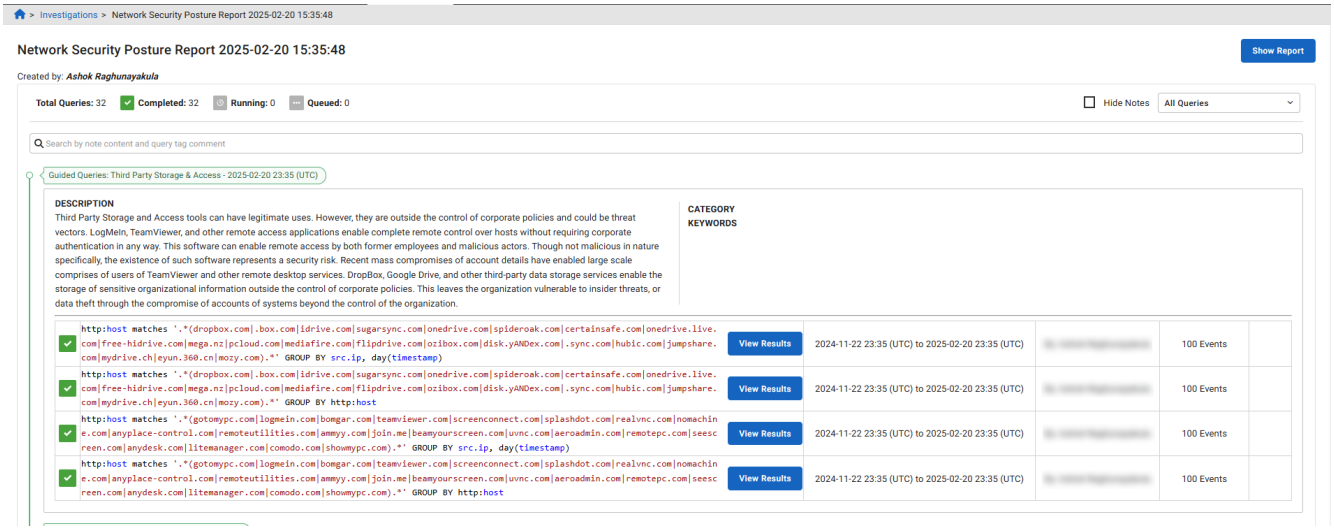

In the *Investigations* page, you can use the *Report* filter to search for *FortiNDR Cloud Network Security Posture Report* investigations.

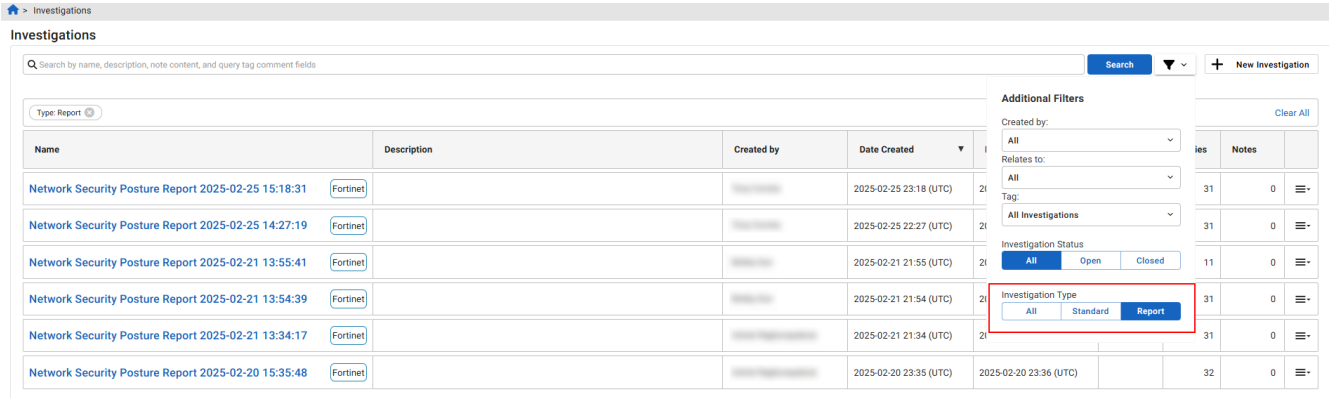

Pending queries in reports

FortiNDR Cloud can support up to 35 pending queries simultaneously. To prevent system overload, a tooltip will appear across all of your accounts advising users to wait before running another report.

Settings

You can apply global settings FortiNDR Cloud by clicking on the gear in the top-right corner of the portal.

Profile settings

Use *Profile Settings* to configure your profiles such as your account and configure authentication.

My profile

User Information	
User Email	The email the user logs into the application with.
User Name	The user's first and last name.
User UUID	The user's unique ID.
User MFA	Indicates if Multifactor Authentication is disabled or enabled.
Tooltip	Disables the investigation tooltip in the *Investigations* page, the *Investigations* widget in the in the default dashboard, and global search results. For more information, see *Investigate* > .
Account Information	
Account Name	The name of the account the user belongs to.
Account UUID	The account's unique ID. The Account UUID is useful when interacting with the APIs. Most APIs allow you to specify an account UUID to pull data for; this is equivalent to setting the Account Selector to a specific account. If you do not specify an account UUID, you receive data from all accounts you have access to.
Subscription Serial Number	The serial number for the account.

Authentication

Password	Click Change my password to update your FortiNDR Cloud password.

	Passwords must be a minimum of eight characters and are valid for 180 days. FortiNDR Cloud will notify you when your password is about to expire. If you attempt to log in after your password has expired, you will be prompted to create a new password.
Multi-Factor Authentication	Click Enable MFA to enter a token each time you log into FortiNDR Cloud. Multi-Factor Authentication requires a Time-based one-time password (TOTP) such as FortiToken. You will be required to configure an MFA token as soon as you long in.

API Tokens

API tokens are used to access FortiNDR Cloud cloud APIs. The token is only shown when it is created. With the exception of the token description, the actual token will not be visible in the portal. Older tokens may be revoked.

For integrations or scenarios where multiple users will rely on the token, a token tied to an API-only user is highly recommended. See, Creating users and assigning roles on page 146.

API Tokens	Click *Create New Token* to create permanent authentication tokens for authenticating API calls. These tokens never expire, and remain valid until revoked.

To create an API token:

1. Go to *Profile Settings* and scroll down to *API Tokens*.
2. Click *Create new token*. The *Create New API Auth Token* dialog opens.
3. In the Description field, enter a description of the token. The description will be visible in the *API Tokens* columns of the *Users* page and the *User Details*.
4. Click *Create*.

To revoke an API token:

1. Go to *Profile Settings* and scroll down to *API Tokens*.
2. In the last column of the table click, *Revoke token*. The *Revoke API Token?* dialog opens.

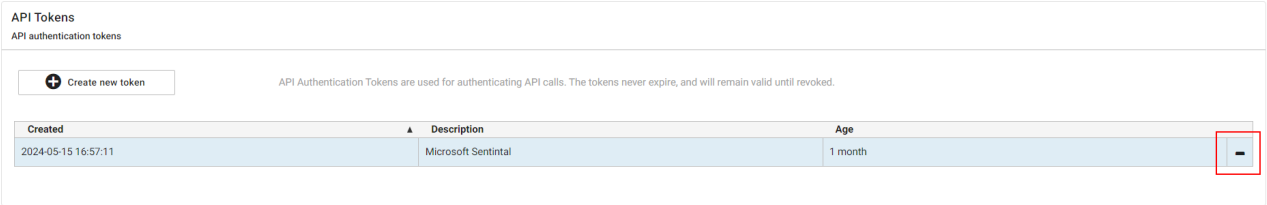

3. Click *Confirm*.

Email notifications

Receive an email notification when a detector triggers a detection. Notifications are configured and applied on a per-user basis using the email address tied to a user's account. If you are logging in for the first time or have never updated your notifications, you will see the *Default Notification* created for every user.

The *Email Notifications* page displays the notifications for the account. You can filter the page by *Email Type* (*Assigned Detections* or *New Detections*) and by *Status* (*All*, *Enabled*, or *Disabled*).

To create a notification:

1. In the toolbar, click the gear icon and select *Email Notifications*. The *Notifications* page opens.
2. Click the *Create Notification* button at the top right-side of the page. The *Create a New Notification* dialog opens.

3. Enter the *Notification Name*.
4. From the *Account* dropdown, select an account.
5. Select the *Detection Type*.
   - *Assigned Detections*: Select to send an email notification to the user the detection is assigned to.
   - *New Detections*: Select to create and configure a new notification.
6. Configure the new notification:

Severities

Select one of the following:

Severity	Description	Examples
High	Significant to fair impact with the potential to spread or escalate	Malicious code execution, C2 communications, lateral movement, data exfiltration
Moderate	Fair impact with minimal potential to spread or escalate	Activity that could indicate malicious intent, untargeted attacks with unknown success, data leakage, subversion of security or monitoring tools
Low	Little to no impact expected	Potentially unauthorized software, devices, or resource use, untargeted adware or spyware, compromise of a personal device or device on an

	Severity	**Description**	**Examples**
			untrusted network, insecure configurations

Confidences	Select one of the following:

Confidence	Minimum True-Positive Rate
High	90%
Moderate	75%
Low	50%

Categories	Select a category from the list. For information, see *Detections* > *Detector Categories*.
Email Type	• *Individual*: Sends an email for each individual detector that becomes active. • *Digest*: Sends you a single email each day at the specified time (default 08:00 Eastern) summarizing detectors that became active and/or were resolved during the previous day. Select *Include Resolved Details* to include detection resolution information in the email The *Email Notifications* page will display *Digest with Resolve Details* next to the email when enabled.

7. Click *Create*.

To edit a notification:

1. Click the *Actions* menu at right side of the notification.

2. Click *Edit Notification* . The *Edit Notification* dialog opens.
3. Edit the notification details and click *Save*.

To delete or disable a notification:

1. Click the *Actions* menu at right side of the notification.
2. Click *Delete Notification* or *Disable Notification*. A confirmation dialog opens.
3. Click *Confirm*.

Manage annotations

Annotations help you categorize and enrich entities (such as IP addresses, domains, or usernames) with additional context for easier investigation and filtering.

Viewing annotations in the portal

Annotations appear against a blue background wherever an IP address is displayed. They can also be searched using IQL when the annotations are present in the data. When you hover over an annotation, a tooltip displays the annotation name and description. Clicking the annotation opens a pop-up with full details.

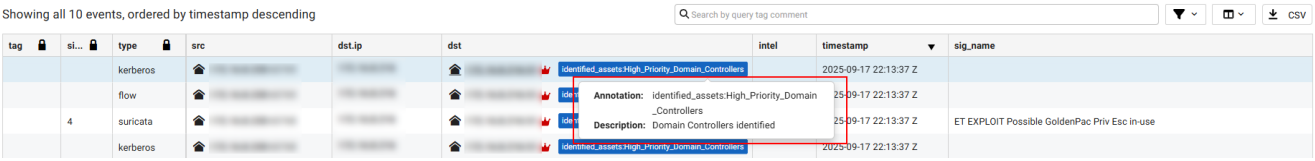

On the *Manage Annotations* page, you can open the *Entity Panel* by clicking the *Entity Name* when it is a valid IP, CIDR, domain, or URL. Right-click an entity with a valid IP to:

- Search Events
- View Detection Context
- View/Create Annotations
- Perform an Entity Lookup
- Perform a Global Search
- Open a Guided Queue

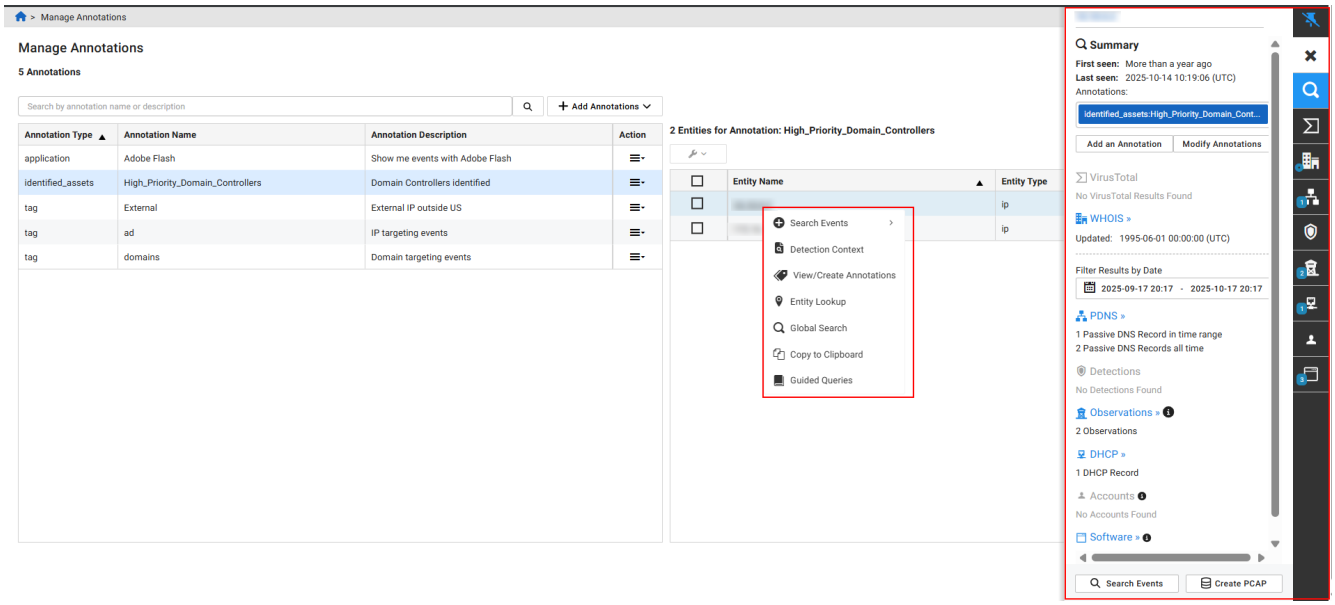

Automatic critical asset identification

FortiGuard ATR uses rich network metadata to automatically discover and classify critical assets across enterprise environments. This process identifies high-value infrastructure components such as Domain Controllers, DNS servers, SSH servers, FTP servers, SMTP servers, and other core services by analyzing behavioral patterns, protocols, and role-based traffic correlations.

Once identified, these assets are annotated within the FortiNDR Cloud platform to enhance network visibility and analytical context. This enrichment helps security teams distinguish routine activity from potential threats targeting essential systems. By adding this layer of asset intelligence, FortiNDR Cloud improves detection accuracy, prioritization, and relevance for high-impact business systems.

A crown icon appears next assets annotated by FortiGuard ATR in detection tables. The crown is color-coded to indicate its severity level:

- Red for high risk
- Orange for moderate risk
- Yellow for low risk

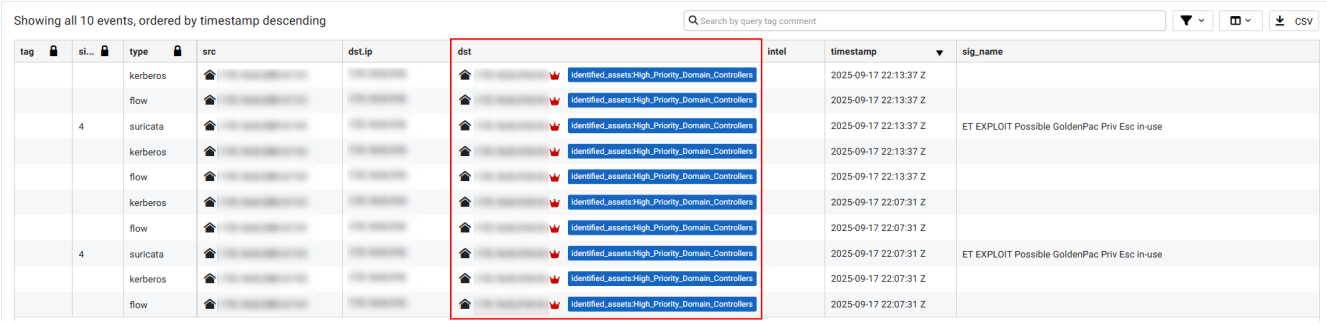

Managing annotations

To manage annotations, click the gear icon at the top-right of the page and select *Manage Annotations*. The *Manage Annotations* page allows you to create, edit, and organize annotations and associate them with specific entities.

The annotations table lists all existing annotations with their type, name, description, and actions, and provides an option to add new annotations. The entities table shows all entities linked to a selected annotation, including their names, types, and management actions. The search function allows you to search for any text in the Annotation Name and Annotation Description columns.

You can create the following annotation types: *Application*, *Environment*, *Location*, *Owner*, *Role*, *Tag* and *Identified Assets* (system-generated). *Identified Assets* annotations are automatically created by FortiGuard ATR and cannot be manually added.

To create an annotation:

1. Click *Add Annotations* > *Create Annotation*.
2. Configure the annotation settings:

Select an annotation type	Select *Application*, *Environment*, *Location*, *Owner*, *Role*, *Tag*, or *Identified Assets*.

	Identified Assets only applies to FortiGuard ATR. See, Automatic critical asset identification. A color-coded crown icon will appear only on assets annotated by FortiGuard ATR in the events and detections tables. See Detections table on page 61.
Enter an annotation name	Enter a name for the annotation.
Enter a description	Enter the annotation.

3. Click *Save*.

To add annotations with a CSV file:

1. Create the CSV file. The file must contain the following : *annotation type*, *annotation name*, *description*, *entity*, *entity_type*.

   The *annotation type* must begin with a lower case letter, and the *annotation name* must be unique within the same type.

	A	B	C	D	E
1	location	USA	us head	1.1.1.1	ip
2	environment	Prod	prod	1.1.1.1	ip
3	owner	test owner	owner description	test	application
4	tag	test tag		1.1.1.1	ip

2. Click *Add Annotations* > *Upload CSV*.
3. Upload the CSV file.
4. Click *Save*.

To edit an annotation:

1. Click the *Actions* menu at the right side of the annotation and select *Edit Annotation*.

2. Update the annotation and click *Save*.

To delete an annotation:

1. Click the *Actions* menu at the right side of the annotation and select *Remove Annotation*.

2. Click *Confirm*.

To add entities:

1. Click *+Add Entity*. The *Add Entities* dialog opens.

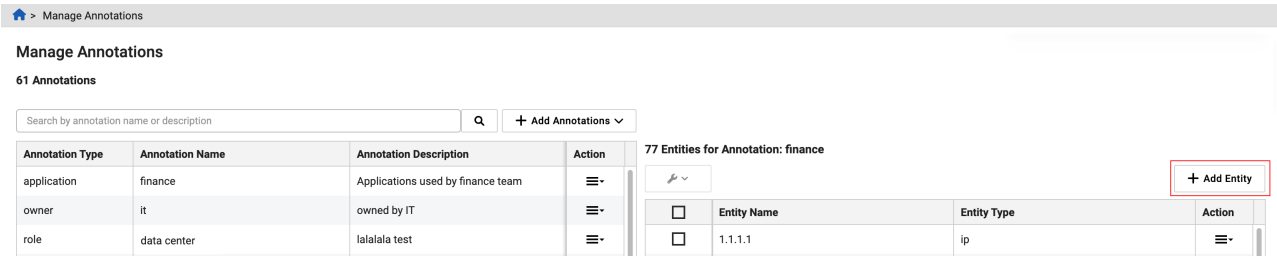

2. Enter one or more entities (IP Address, CIDR, domain or username) separated by a comma, space, or return.

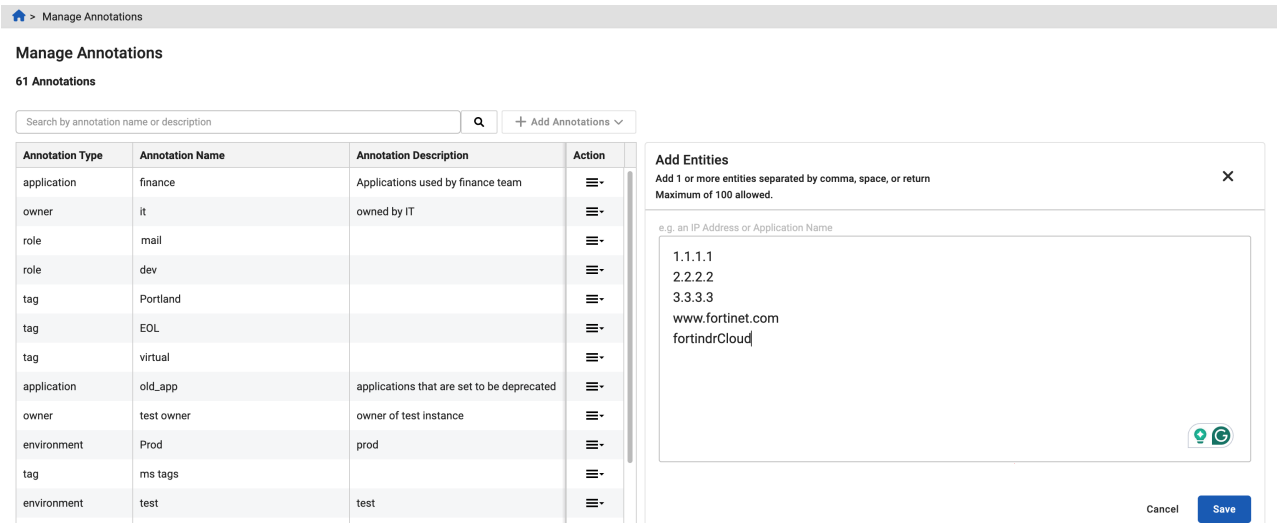

3. Click *Save*. FortiNDR Cloud validates the fields and identifies any errors.

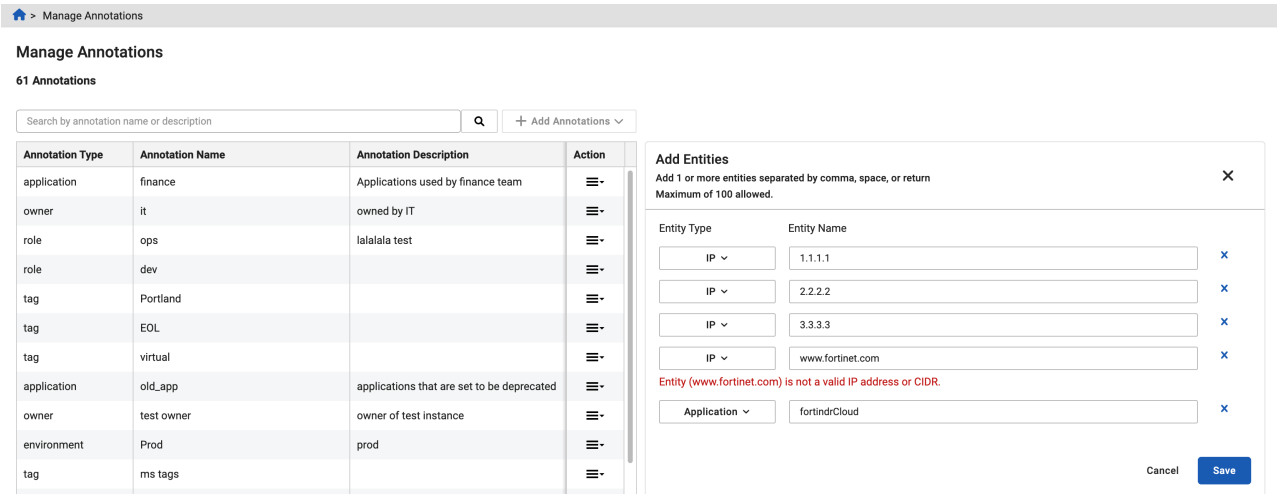

To bulk remove entities:

1. Above the entities table, Click *Remove bulk entities*.

2. Click *Confirm*.

Mutes and excludes

The *Mutes and Excludes* page summarizes all muted and excluded devices, including device-level mutes for detectors. It contains three tabs: *Mutes*, *Excludes* and *Subnets*. To open the page, click the gear icon and select *Mutes and Excludes*.

Mutes tab

The Mutes tab displays four categories:

- *Muted Detectors*: Detectors that are muted across all devices.
- *Muted Devices*: Devices that have muted all detectors.
- *Muted Devices for Detectors*: Devices muted for specific detectors.
- *Muted Detections*: Detections that are muted either at the account level or for a specific detector.

You can use the menu in the *Actions* column of the tables to add a muted device for the whole account, unmute or edit existing muted devices, add or update a muted device for a specific detector. For more information, see Muting on page 43

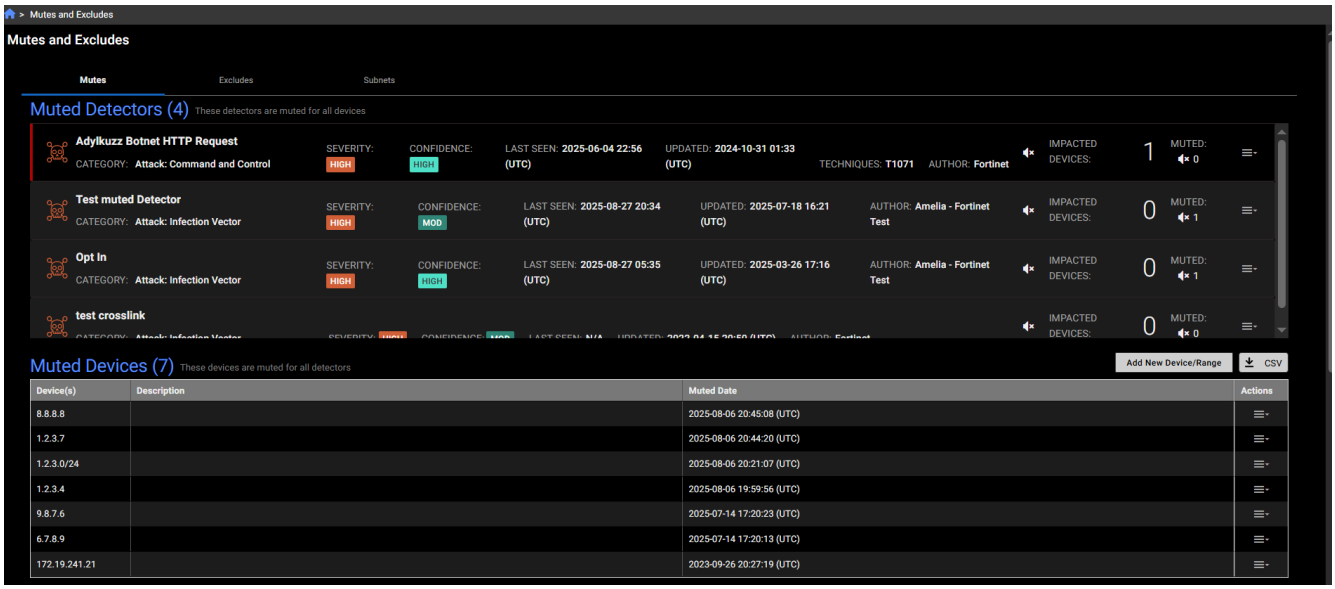

Excludes tab

The *Excludes* tab lists devices excluded at the account level, meaning they will not trigger any detections. It also includes disabled detectors.

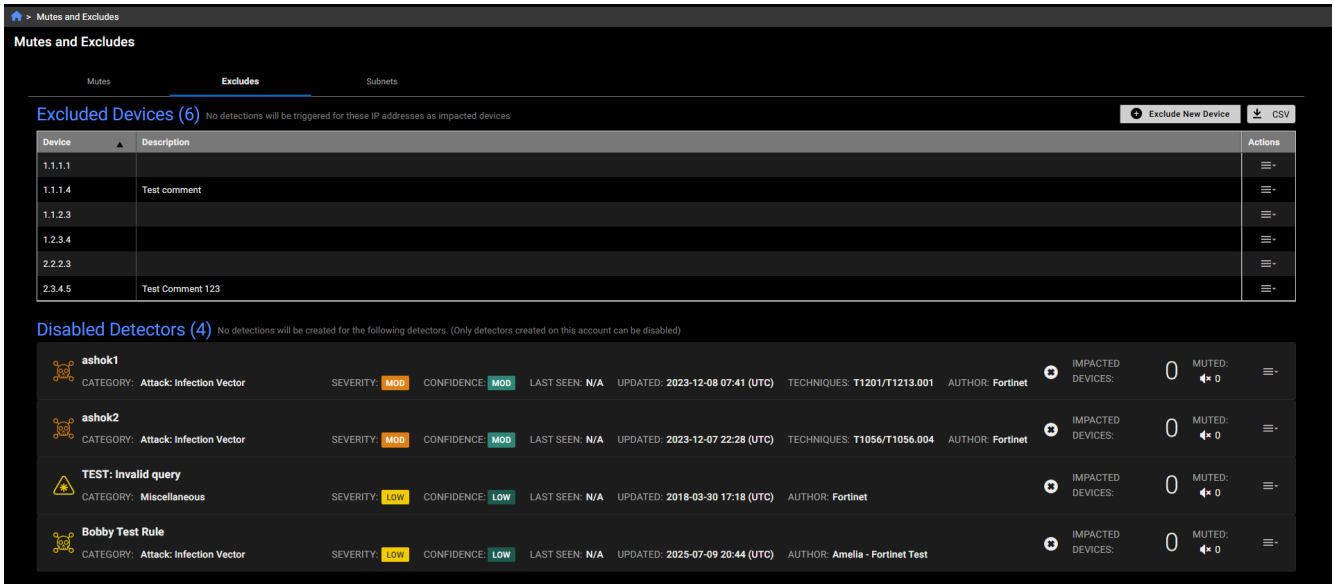

Subnets tab

This tab displays all internal subnets for the account. Detections will only be created when the impacted device is within an internal subnet. It allows you to view and modify settings related to muted devices, excluded devices, and subnets all in one place.

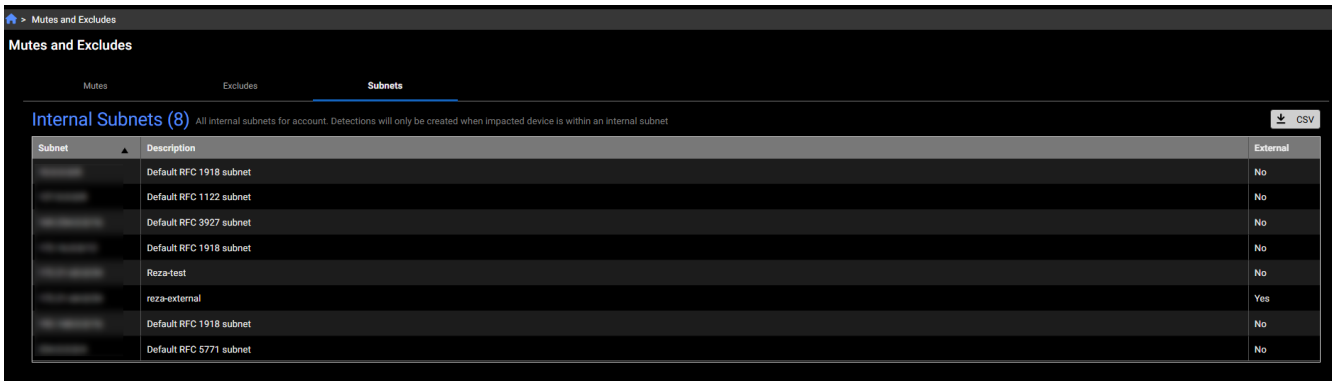

=

Sensors

The *Sensors* page shows the sensors deployed in your account, both in the aggregate and individually. Use this page to generate provisioning codes, check the status of individual sensors, and view telemetry data.

To access to the Sensors page, click the gear icon at the top-right of the page and select *Sensors* from the menu.

Sensor ID	Click the Sensor ID to view the sensor *Status*, *Telemetry* and *Settings* pages. For information, see Sensor details on page 136 You can pivot to the *Sensor Details* page from the *Sensor ID* column in the *detections Details*. Go to *Detections > Triage detections* and open a detector. Click a sensor in the *Sensor ID* column. If the sensor is available, the *Sensor Details* page opens.
Status	The sensor connection status. **Online** — Sensor is connected to FortiNDR Cloud within last hour. **Offline** — No telemetry data received by the sensor for at least an hour. **Provisioning** — Provisioning code has been created and made initial connection but provisioning process is not complete. **Decommissioned** — Sensor has been factory reset (only applicable for 1.12 or above).

	Decommissioned (legacy)	A sensor earlier than 1.12 has been marked as decommissioned and has not sent any additional data. If sensor sends data to FortiNDR Cloud, status will change to *Online*.
	Decommissioned (auto)	A sensor 1.12 or later has been marked as decommissioned, but has not communicated with FortiNDR Cloud in the last 7 days. If the sensor later connects to FortiNDR Cloud, it should factory reset itself and switch to *Decommissioned* status.
	Decommission Pending	The sensor decommissioning has been initiated.
	Paused	The sensor is not receiving traffic and can be enabled later.
	Pausing	The sensor is in the process of being paused. You cannot resume a sensor while it is in this state.
	Resuming	The sensor is in the process of being resumed. You cannot pause a sensor while it is in this state.
	Shutdown	A Zscaler virtual sensor is no longer active.
	All other statuses are written by the sensor itself.	
Version	The sensor version. *Unknown* is displayed when there is no data for the version.	
Labels	Annotations that are applied to the sensor. See, Manage annotations on page 127.	
Location	The sensor location.	
EPS (7 Day Average)	The average throughput over last 7 days as Events Per Second.	
BITS/S (7 Day Average)	The average throughput over last 7 days as Bits Per Second.	
Type	The platform the sensor was deployed on.	
Features	Lists the enabled tools used to analyze network traffic and detect anomalies, such as Suricata, PCAP or DPI.	

Sensors toolbar

The following table describes the toolbar options available on the page and their functions:

	Filter the page by *Status*, *Type*, *Version* or *Features* .
	View the sensor analytics. See Account telemetry.
	View network devices detected by sensors, with options to filter, drill down into subnets and hosts, and analyze traffic patterns over time. See Device

	view on page 143.
Actions	Download the sensor image or provision a sensor.
	Show or hide columns on the page.
CSV	Down load the sensor data a CSV file.

Account telemetry

The *Telemetry* page displays aggregated telemetry data from all sensors in your account. The legend at the right side of the page lists the entries in descending order from highest to lowest. You can use the toggles in the legend to show or hide lines in the graph.

To view the telemetry for each sensor, click the *Telemetry* tab in the *Sensor Stat*us page. See Sensor details on page 136.

To view the Account Telemetry page:

1. Click the gear icon at the top-right of page select *Sensors*.
2. Click the *Telemetry* button at the top-right of the page. The *Throughput* page opens.
3. (Optional) Click *Chart Type* to switch between *Line* and *Bar* views.

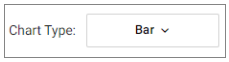

4. (Optional) Filter the page.

Group by	View the telemetry data by *Sensor*, *Event Type*, or *Interface* when available.
Interval	Select *Day*, *Hour* or *5 minutes*.
Date Range	Click to configure the date range using the date picker, or choose a value from the *Quick Ranges* list.

5. Click the *CSV* button to export the data as a CSV file. The CSV file will download everything in the graph. You can use the legend to select the sensor data you want to download.

Sensor details

This Senor details page provides an overview of each sensor's status, resource usage, connectivity, and configuration details.

The top of the page displays the following information:

Statistic	Description
Created	The date and time when the sensor was provisioned or first connected.
Location	The physical or logical site associated with the sensor.
Type	The deployment platform for the sensor. For example, *ESXi*, indicating it runs as a virtual appliance on VMware ESXi.
CPU	The current processor utilization reported by the sensor. The usage value is color-coded: • Green if usage is 0-59% • Yellow if usage exceeds 60% • Red if usage exceeds 90%
Memory	The current memory utilization reported by the sensor. The usage value is color-coded: • Green if usage is 0-59% • Yellow if usage exceeds 60% • Red if usage exceeds 90%
EPS	The average number of events per second generated over the last seven days.
Bits/s	The average throughput in bits per second over the last seven days.

Status Tab

The *Status* tab shows the current state of the sensor, including whether it is online, offline, or in a transitional phase. It also provides hardware details and connectivity indicators so you can quickly assess sensor health.

Connection Status

Field	Description
Status	Indicates whether the sensor is currently online and connected to the management system.
Serial Number	The unique identifier for the sensor hardware or virtual instance.
Management IP	The IP address used for managing and communicating with the sensor.
Last Updated	The timestamp of the most recent status update from the sensor.

Device Enrichment Status

This section appears when Device Enrichment is enabled. See Device enrichment on page 157.

Field	Description
Last Run Time	Displays the timestamp of the most recent device enrichment cycle executed, whether scheduled or manually triggered.
Last Upload Time	Shows the timestamp of the most recent upload of enrichment results from the sensor. This indicates when updated device information was last synchronized.
Message	Displays raw status or progress information provided by the sensor during the device enrichment process. This field is intended for troubleshooting and provides visibility into the sensor's activity.

Interfaces

Each interface displays its IP address when that information is available from the API response. This is especially helpful if the interface is configured as a NetFlow collector.

A green interface indicates that a cable is connected, while a gray interface means no connection. Click the interface label to view its MAC address.

Interface	Description
ensxxx (mgmt)	Displays traffic rate and IP address for the management interface.
ensxxx	Shows traffic rate for the secondary interface used for data capture.

The following table details the naming convention for interfaces on FortiNDR Cloud sensors.

Label	Sensor Type	Interface Type	Purpose	Max Bandwidth
em4	Physical	Ethernet	Management	1 Gb/s
em3	Physical	Ethernet	Monitoring	1 Gb/s
em2	Physical	Ethernet	Monitoring	10 Gb/s
em1	Physical	Ethernet	Monitoring	10 Gb/s
p#p##	Physical	Fiber	Monitoring	10 Gb/s
eth0	Virtual	Virtual	Management	N/A
eth1+	Virtual	Virtual	Monitoring	N/A

Hardware

Field	Description
Processor(s)	Details the CPU model and specifications used by the sensor.
Number of Cores	Indicates how many CPU cores are available for processing tasks.
Memory	The total amount of RAM allocated to the sensor for operations and the percentage currently in use.

Field	Description
Total Disk Space	The total storage capacity available for logs and system files.

CPU & Memory Usage

Shows the percentage of CPU currently in use by the sensor. The graph tracks CPU and memory usage for the last 7 days. Hover over the graph to view usage at a specific point in time.

The chart can be viewed by hour or by day. The tooltip displays time for hourly intervals and date for daily intervals. A 5 minute interval is available only for ranges within 24 hours; if the range is larger, the system automatically switches to a compatible interval such as hourly or daily.

The maximum selectable range is 14 days. The last time shown may end at `xx:55` instead of `xx:59` because of interval rounding.

Software

Field	Description
Operating System	The OS running on the sensor, including version details.
ZEEK Version	The installed version of ZEEK used for network traffic analysis.
Suricata Version	The installed version of Suricata for intrusion detection and analysis.
Sensor Version	The current version of the sensor software package.

History

The *History* table is sorted in descending order by timestamp. A message appears if there is no history to display.

Field	Description
Timestamp	The date and time the action was recorded in UTC.
Action	The operation performed on the sensor, such as Provision, Pause, or Resume.
User Account Name	The account associated with the user who performed the action.
User Name	The individual user or service identity that triggered the action.
Comment	An optional note explaining the reason or context for the action.

Telemetry Tab

The *Telemetry* tab displays traffic and performance metrics such as throughput and event rates over time. This tab helps you monitor sensor activity and detect trends that may affect network visibility or performance.

The Telemetry page includes three tabs that provide different types of sensor data:

Tab	Description
Throughput	Displays measurements of total throughput across the sensor's interfaces in bits per second. You can view the data as a line or bar graph for any time period, group by interface name, set the interval to Day, Hour, or 5 Minutes, and download the data as a CSV file. The legend displays the total throughput count for each individual sensor from highest to lowest. Use the toggles in the legend to show or hide a line in the graph. You also have the option of showing or hiding all entries.
Events	Shows the number of events produced by the sensor. The data can be displayed as a line or bar graph and grouped by event type. The legend lists sensors from highest to lowest event count and includes toggles to show or hide individual lines or all entries.
Visibility	Displays observed devices for the sensor, similar to a simplified version of the Devices page. This helps you identify endpoints seen by the sensor over time.
CPU & Memory Usage	Shows the percentage of CPU currently in use by the sensor. The graph tracks CPU and memory usage for the last 7 days. Hover over the graph to view usage at a specific point in time. The chart can be viewed by hour or by day. The tooltip displays time for hourly intervals and date for daily intervals. A 5 minute interval is available only for ranges within 24 hours; if the range is larger, the system automatically switches to a compatible interval such as hourly or daily. The maximum selectable range is 14 days. The last time shown may end at `xx:55` instead of `xx:59` because of interval rounding.

The *Traffic by Type* custom dashboard displays the data in the *Events* tab in the *Sensor telemetry* page. When you click the widget header it opens the Sensor telemetry page. All the filters applied to the widget will be transferred to the *Sensor Telemetry* page. See, Creating custom dashboards on page 31

Settings Tab

Update sensor information such as the name, location, and labels, and configure options such as packet capture.

For assistance with these settings, contact your Technical Success Manager.

Enabling PCAP has security and privacy complications. Before enabling PCAP, consult with your Technical Success Manager.

For example, networks with data that is subject to regulatory requirements may require certain controls to be in place before enabling this feature. Enabling this feature may also require uploading a public key to encrypt any PCAPs. See, Account management on page 145 or contact Customer Support for more information on public keys.

Statistic	Description
General	Click *Edit General Settings* to edit these settings.
Location	The physical location of the sensor.
Labels	arbitrary labels (hostname, site/building code, etc.)
Features	Click *Edit Feature Settings* to edit these settings.
PCAP Enabled	Toggle ON to enable PCAP capture.
Packet Inspection Engine	The network packets used to identify traffic and enforce security policies.

Account telemetry

The *Telemetry* page displays aggregated telemetry data from all sensors in your account. The legend at the right side of the page lists the entries in descending order from highest to lowest. You can use the toggles in the legend to show or hide lines in the graph.

To view the telemetry for each sensor, click the *Telemetry* tab in the *Sensor Status* page. See Sensor details on page 136.

To view the Account Telemetry page:

1. Click the gear icon at the top-right of page select *Sensors*.
2. Click the *Telemetry* button at the top-right of the page. The *Throughput* page opens.
3. (Optional) Click *Chart Type* to switch between *Line* and *Bar* views.

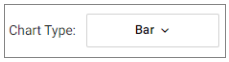

4. (Optional) Filter the page.

Group by	View the telemetry data by *Sensor*, *Event Type*, or *Interface* when available.
Interval	Select *Day*, *Hour* or *5 minutes*.
Date Range	Click to configure the date range using the date picker, or choose a value from the *Quick Ranges* list.

5. Click the *CSV* button to export the data as a CSV file. The CSV file will download everything in the graph. You can use the legend to select the sensor data you want to download.

Sensor settings

Use the sensor *Settings* page to update the sensor location, make annotations and enable or disable Packet Capture. You can also access the sensor settings from the *Actions* menu on the *Sensors* page.

Requirements:

- You must have Admin privileges to edit the sensor settings.

To edit the sensor settings:

1. Click the gear icon at the top-right of page select *Sensors*. The *Sensor* page opens.
2. Click the *Sensor ID*. The sensor *Status* page opens.
3. Click the *Settings* tab. The *General* page displays the sensor *Location*, *Labels* and *PCAP* status.
4. Click *Edit General Settings* to edit the sensor *Location* and *Labels*.

Location	Update the sensor location.
Labels	Enter keywords about the sensors. To add annotation, type the phrase or keyword and press Tab or Enter. Annotations with an orange background are internal an cannot be edited. Annotations with a blue background can be added or deleted.

5. Click *Edit Features Settings* to enable/disable *Packet Capture*.

PCAP Enabled	Enable packet capture. For more information, see Packet capture on page 98.

To edit the settings from the Sensors page:

1. On the *Sensors* page, click the actions menu at the right side of the page and click *Edit*.
2. Update the Sensor details and click *Update*.

Location	Update the sensor location.
Annotations	Enter keywords about the sensors. To add annotation, type the phrase or keyword and press Tab or Enter. Annotations with an orange background are internal an cannot be edited. Annotations with a blue background can be added or deleted.
PCAP Enabled	Enable packet capture. For more information, see Packet capture on page 98.

Packet Inspection Engine	• *Suricata*: A Suricata event is created when Suricata (an intrusion detection tool) alerts or metadata are integrated into Zeek logs, highlighting threat detection signatures and behaviors. See, Suricata fields on page 258. • *Fortinet DPI*: A DPI (Deep Packet Inspection) event is created by the Fortinet IPS (Intrusion Prevention System) engine running on the sensor which logs informative and pattern matching based events. The IPS engine logs AppID (Applications seen by the engine for software and protocols), IDS (signatures for vulnerabilities), OT Protocols/Threats (Operational Technology based protocol parsing and signatures), Botnet (Botnet based traffic patterns), and Info (informational events about protocols). See,DPI fields on page 229.

Packet Inspection Engine Guidelines

VM Sensors	To run *Suricata* and *Fortinet DPI* engines concurrently, the following minimum recommended system resources are required: • CPU: 24 cores • RAM: 48 GB
Physical Sensors	Enabling DPI on physical sensors depends on available system resources. Customers should contact the support team for confirmation and guidance.

Device view

FortiNDR Cloud continuously collects data about devices on the network. This data is gathered on a per-sensor basis because multiple sensors may report the same IP address. This can occur due to IP address reuse within the environment or because traffic from a single IP crosses multiple monitoring points.

You can use Device View to:

- Quantify FortiNDR Cloud sensor visibility coverage over time.
- Verify that FortiNDR Cloud sees both internal and external traffic from network devices.

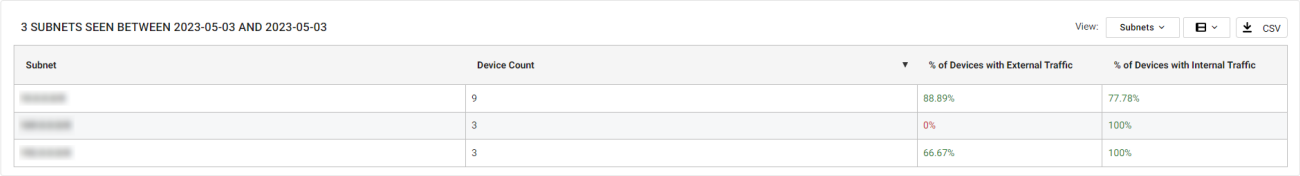

Viewing visible devices

To view the visible devices:

1. Click the gear icon at the top-right of the page and select *Sensors*.

2. In the toolbar, click *Visible Devices* . The page is organized into three sections:

All Subnets	**Search**	Enter a subnet or prefix to view a specific device.
	Date	Click to open the date picker to view devices within a specif date range.
	Additional Filters	Click the filter icon to view devices by sensor and Internal and External traffic directions.
Devices by Subnet	**Highlight by**	Select *External Traffic %* or *Internal Traffic %* to change the colors in the box-plot chart to show the percentage of assets. Use this view to verify FortiNDR Cloud is seeing both internal (East-West) and external (North-South) traffic on a specific subnet.
	View	• *By Subnet*: This the default view. • *Over Time*: Shows how many devices were seen within the selected subnet over time.

		This graph is if sensor coverage is experiencing issues or to debug problems with missing events for a certain time period.
	Box-plot chart	Click the box-plot chart to drill down into the selected subset of the network.
# SUBNETS SEEN BETWEEN YYYY-MM-DD AND YYYY-MM-DD	Shows either a summary of subnets or a list of discrete devices. This table is useful for reviewing the traffic on a per device basis.	

Account management

Use the *Account Management* page to create new users and manage global settings for your account. You must have Admin privileges for one or more accounts to view the *Account Management* page.

To view the Account Management page:

Click the gear icon at the top-right of the page and select *Account Management*.

- If you have access to only one account, you will see the *Account Management* page for your account
- If you have Admin privileges for more than one account, you will see the *Account Inventory* page. From there, click an account to view it's *Account Management* page.
  You can filter page by *Account Name*, *Created Date*, *Account Code*, *Number of Users*, *Number of Sensors* and *Last Login*. You can also sort the *Accounts* page by *Last Login* to view which accounts are in use to help determine if they should be removed.

The top of the page will display descriptive parameters for the account, namely the account's UUID and sensor code, as well as the number of users and sensors provisioned in the account. A banner is displayed when your account is set to expire in less than 90 days.

The *Account Management* page contains the following tabs:

Users	Create new users and assign roles.
Subnets	Lists all internal IP address ranges for the account. This list will always include the ranges defined in RFC 1918, link local addresses (169.254.0.0/16), and multicast addresses (224.0.0.0/4). We recommend adding a public IP space owned by your organization, such as post-NAT, egress, or externally-accessible IP addresses, to this list. Doing so better characterizes the directionality of your network's traffic. Contact your TSM with any public IP addresses or ranges that you would like to add to this list. Admin users can add, edit or delete subnets in an account. See Add or edit subnets on page 159
Modules	Displays the available integrations for FortiNDR Cloud.

Settings	Enable SAML SSO, mulit-factor authentication, and generate PCAP encryption keys.
Billing	The billing summary provides three top-level views: **Bandwidth**: Bandwidth data is shown for all customers. Accounts are billed on the 95th percentile of aggregate bandwidth usage across all sensors, measured over 10-second intervals. Data volume transferred is expressed in KB, MB, or bytes. **Log Ingestion**: This tab is shown for all customers but only applies to customers with a log ingestion license. Accounts are billed on the average EPS (Events Per Second). Log processing rate is expressed in EPS. **Fortinet Automation Service**: This tab is shown only when the *Fortinet Automation Service* is enabled. It provides the account's License Details , including the serial number, description, API key status, and the number of concurrent users allowed. The tab also displays the daily playbook limit, storage usage, and license expiry status, along with a warning about the upcoming expiration. Additionally, it lists playbook metrics, such as concurrent execution limits, execution errors, the number of agents, playbooks created, playbook execution rate, and polling frequency. *Bandwidth* and *Log Ingestion* tabs show your account's usage for the current date compared to your available license. Both have three sections: *Current month usage*, *Daily stats* (configurable to any range up to 90 days), and *Monthly History* (configurable for any range). For customers with multiple accounts, toggle between parent and child accounts using *View <Account Name> alone* or *View <Account Name> and child*.

Creating users and assigning roles

Go to *Account Management > Users* to add users and assign roles. You also have the option of creating *API Only* users. The User Management table displays all the users with access to the portal.

The *Account Management > Users* page displays the following information:

Column	Description
Email	The user's email address ADMIN — Indicates the user has Admin privileges. Indicates the user is a Portal user. Indicates the user is API Only user.
Full Name	The user's full name.
First Name	The user's first name.
Last Name	The user's last name.
UUID	The user's unique ID.
Last Login	The date and time the user last logged into the account.
Created	The date the user was crated.
Updated	The date and time the user's details were updated.
Status	The user's current status (*Enabled*/*Disabled*).
Locked Out	Indicates the user has been locked out of the account.
MFA	Indicates Mufti-Factor Authentication is enabled or disabled.
Roles	The user role. This column is not displayed by default.
Actions	Use the menu in this column to: • Edit the user details • Move the user between accounts • Email/reset the password. • Disable the user.

To create a new user:

1. Click the gear icon at the top-right of the page and select *Account Management*. (Click the *Users* tab if it is not already open.)
2. Click *Create User*. The *Create New User* dialog opens.
3. Enter the user's details. Required fields are indicated with an asterisk (*).

Email	Enter the user's email address.
First name	Enter the user's first name.
Last name	Enter the user's last name.
Assign role	Select the user role. The following descriptions are also displayed in the portal when you hover over the role name.

Role	Description
User	This role grants permission to perform all non-administrative functions within the portal, including the ability to manage all features for the Detections function of the product Most users will utilize this role for their duties within the product.
Limited User	This role grants permissions to perform the most basic functions within the portal, however it limits a user's ability to manage Detectors, Mutes, and Exclusions within Detections. This role is primarily designed for teams utilizing a multi-teir SOC in which lower-tier analysts should not be able to prevent future detections from firing without review from an upper-tier analyst.
Admin	This role has permissions to configure account-level settings (such as PCAP encryption, enforcing MFA requirements, and so on) and allows grantees the ability to manage users within the account. **Note**: Admins must also have a *User* permission to perform actions in the portal such as viewing Detections or running queries. When the *Admin* role is selected, the system automatically checks for the *User* role. This is because *Admins* need the *User* role for full functionality. If the *User* role is not selected, a warning will appear. You can still create the user if you choose to ignore the warning.

API Only	*API-only users* are primarily designed for integration configurations. They cannot have passwords or multi-factor authentication enabled, they do not receive emails, and their keys are managed entirely by those with *Admin* privileges for the account. API-only users do not appear in the user list by default, but can be displayed by adjusting the page filters. See, To filter the user list. *API Only* is the user role when mandatory SSO is enabled.

4. Click *Create*.

New users are automatically assigned the *Training User* role on the Training Modern account, even if the administrator has not assigned any roles to the user. If the account is a parent account, and the administrator has access to child accounts, then a checkbox is available to include child accounts.

To view user details:

Double-click a user in the list. The user details pane opens.

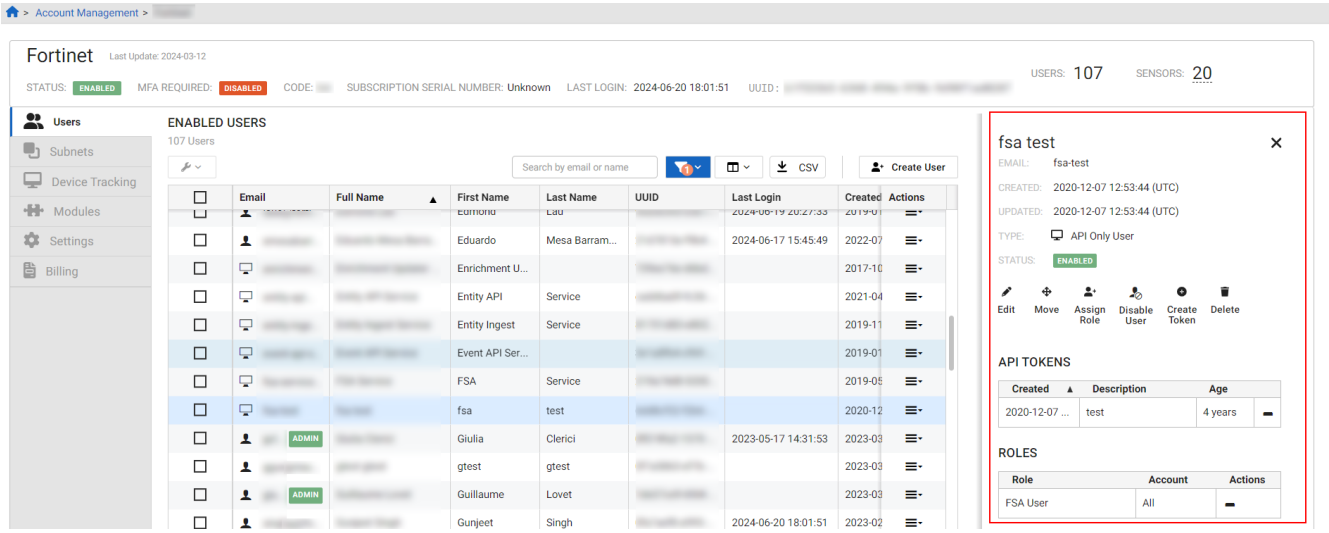

- The icon indicates the role assigned to the user also belongs to child accounts.
- *Edit* and *Reset Password* are disabled with mandatory SSO is enabled.

To filter the user list:

1. Click the Filter icon.

2. Select the filter type.

Status	Select *All*, *Enabled* or *Disabled*.
User Type	Select *All*, *Portal* or *API Only*.
Account Access	Select an account from the dropdown list.
User Role	Select a user role from the dropdown list.
Oldest API Tokens Age	Select *Any Token Age*, *No Token* or a value between 3 - 12 months.

To update a user's details:

1. Click a user in the list. The *User Details* pane opens.

Option	Purpose
Edit	Modify the email or name for the user account.
Move	Assign the user to a different account.
Assign Role	Assign a role to a user. • *User* • *Limited User* • *Admin*
Reset Password	Send an email with a password reset link to the user.
Disable MFA	Disable the requirement for an MFA token for the user. If *Require MFA* is enabled for the account, the user will be required to re-establish an MFA token on next log in.
Unlock	Unlock the user account. User accounts are locked after five failed password attempts in 10 minutes.
Disable User	Disable log in access to the user account and any of its API tokens.

Optionally, you can use the menu in the *Actions* column to quickly *Edit User*, *Move User*, *Email Password Reset* or *Disable User*.

The *Edit User*and *Email Password Reset* are disabled when mandatory SSO is enabled.

2. Click close (**X**) to close the pane.

To perform bulk actions:

1. Select the users in the lists or select all. The tools icon is activated.

2. Click the tool icon and select *Move Users*, *Enable Users*, *Disable Users*, *Assign Role* or *Revoke Role*.

To export the user list as a CSV file:

- In the toolbar, click the *CSV* button. The list is saved to your device.

In the *user_role* column, if the user has:

- No account name in front of the role, this indicates the user belongs to the current account (Admin, User, Limited User).
- The same role in two or more accounts, the account name is displayed followed by a colon (:) followed by the user role.
- A child account, the *user_roles* column will indicate *includes children*.
- A role in a different account, the role it displayed in a separate *user_role* column for the account.

SAML SSO

FortiNDR Cloud translates SAML authentication from the identity provider into the native authentication scheme. User login is the same regardless of whether the user has logged in using SAML or a password. The session state in FortiNDR Cloud is independent of the SAML session. Logging out of SAML does not log the user out of FortiNDR Cloud.

When enabling SAML SSO keep the following considerations in mind:

- First time FortiNDR Cloud users will have a user record created automatically when they first authenticate using SAML. Users are required to have a first name, but the last name is optional. These users will initially have no permissions. An Admin will need to grant roles to these users using the normal Account Management UI.
- When existing users authenticate using SAML, any changes to their first and last name will be updated in FortiNDR Cloud as well.
- FortiNDR Cloud identifies users from SAML by their email address. If the user's email address has changed in the SAML SSO Provider, FortiNDR Cloud will create a new user record for that user the next
- Disabling a user in FortiNDR Cloud also disables SAML authentication for that user. However, disabling a user in the SAML SSO Provider does not disable the user in FortiNDR Cloud. The user will still have access if they have a password or API token. Users need to be manually disabled in FortiNDR Cloud as well.
- Users authenticating with SAML are also allowed to authenticate using passwords as well. Typically, at least one Admin in the account should have a password as a backup in case SAML authentication fails.

Failure Scenarios

There are a variety of reasons why SAML authentication may fail.

- SAML has not been configured for the account.
- SAML has been configured, but disabled.
- The user is attempting to authenticate with the wrong account. For example, the user belongs to the Acme account but is trying to authenticate with the Acme Subsidiary account.
- The user has been disabled in FortiNDR Cloud.
- The user does not have a first name.

For security reasons, FortiNDR Cloud may not provide the exact reason for the failure. Please make sure that SAML is configured correctly for the account and the user.

To enable SAML login:

1. Click the gear icon at the top-right of the page and select *Account Management*.
2. Select an account.
3. Click the *Settings* tab.
4. Click *Set up SAML SSO*. The *SAML Single Sign-on (SSO) Initial Setup* dialog opens.

5. Copy the values from the *Single Sign-On URL* and *Entity ID* fields and paste them into the general settings of your SAML Provider configuration.

> *Entity ID* may also be called *Audience URI* or *SP Entity ID*.

6. Set the application's subject or username to *Email*.
7. Add an attribute statement, `first_name`, with the value for a user's first name.
8. Add an attribute statement, *last_name*, with the value for a user's last name.
9. Enter the following information from your SAML SSO Provider into the *SAML Single Sign-on (SSO) Initial Setup* dialog:
   - *IdP Entity ID*
   - *X.509 Certificate (IdP Public Key)*
10. Click *Save*.

To login with SAML SSO:

1. Navigate to your SAML SSO Provider's dashboard
2. Click the ThreatINSIGHT or FortiNDR Cloud button from the SAML SSO Provider's dashboard

> - FortiNDR Cloud only supports IdP (identity-provider) initiated logins where the user will need to initiate login from their SAML SSP Provider's dashboard.
> - If you are a new user logging into FortiNDR Cloud for the first time, you will see a message indicating that you do not have permission to use this application. This means that your roles have not yet been granted. Contact your administrator to assign your roles.

To disable SAML SSO:

1. Click the gear icon at the top-right of the page and select *Account Management*.
2. Select an account.
3. Click the *Settings* tab and click *Disable SAML Settings*.
4. In the Confirmation Dialog, click *Confirm*.

OneLogin SAML Configuration

Requirements:

- SAML *Single Sign-On URL*.

To configure OneLogin:

1. Add a new application using the *SAML Custom Connector (Advanced)*. For more information, see the product documentation.
2. In the *Configuration* section, use your FortiNDR Cloud*Single Sign-On URL* for the following fields:
   - *Audience (EntityID)*
   - *Recipient*
   - *ACS (Consumer) URL Validator*
   - *ACS (Consumer) URL*

   The ACS (Consumer) URL Validator is a regular expression. Replace the beginning of the URL:
   `https://portal.fortindr.forticloud.com/v1/saml/`
   With the following:
   `^https:\/\/portal.fortindr.forticloud.com\/v1\/saml\/`
3. Make sure the *SAML initiator* field is set to *OneLogin*.
4. Change the *SAML signature element* to *Both*.
5. In the *Parameter* section, add the following fields and select the *Include in SAML assertion* flag for each:

Name	Value
first_name	First Name
last_name	Last Name

6. In the *SSO* section, copy the *Issuer URL* and the *X.509 Certificate*. You will need these later.

To configure FortiNDR Cloud:

Update the *SSO SAML Setup* fields with the OneLogin values you copied earlier.

Field	Value
IdP Entity ID	OneLogin *Issuer URL*.
X.509 Certificate	OneLogin *X.509 Certificate*.

Mandatory SSO

You can require all users to log into FortiNDR Cloud using SSO. Before enabling mandatory SSO, keep the following considerations in mind:

- Multi-Factor Authentication (MFA) is disabled.
- You can only edit API users
- *Change my password* and *Enable MFA* are disabled in *Profile Settings > My Profile > Authentication*
- *Edit User* and *Email Password Reset* are disabled in *Account Management > Users > Actions*.

Requirements:

- SAML SSO must be enabled.
- User must have *account.sso_required.update* permissions

To enable mandatory SSO:

1. Click the gear icon at the top-right of the page and select *Account Management*.
2. Select an account.
3. Click the *Settings* tab.
4. Under *SAML SSO* enable *Require SSO Login (disable login with username/password)*. The *Confirm enabling mandatory SSO login* dialog opens.
5. Click *Confirm*.

PCAP encryption keys

PCAP Encryption Keys are used in conjunction with Packet Capture. If an encryption key is uploaded, all PCAP files will be encrypted with the provided key. This prevents FortiNDR Cloud from having any visibility into the raw PCAP data that was captured. For more information, see Packet capture on page 98.

The *Uploaded by* field displays the full name and UUID of the user who uploaded the encryption key as well as the *Uploaded date*. If the user does not belong to the account, *Unknown User* is displayed.

The corresponding private key will be required to decrypt any downloaded PCAP files. If the private key is lost, the encrypted PCAP files cannot be recovered.

To upload an encryption key:

1. Click the gear icon at the top-right of the page and select *Account Management*.
2. Select an account.
3. Click the *Settings* tab.
4. Under *PCAP ENCRYPTION KEYS*, click *Set PCAP Encryption Key*. The *Set PCAP Encryption Key* dialog opens.
5. Paste the public key and click *Set Key*. The encryption key is validated for errors.

The key will take effect for any new PCAP files generated. Existing PCAP files are not retroactively encrypted.

Multi-factor authentication

Enable Multi-factor authentication (MFA) to require all users to enter an MFA token the next time they log in to FortiNDR Cloud. Users will not be able to navigate to any FortiNDR Cloud page until they confirm their MFA token.

To enable Multi-factor authentication:

1. Click the gear icon at the top-right of the page and select *Profile Settings*.

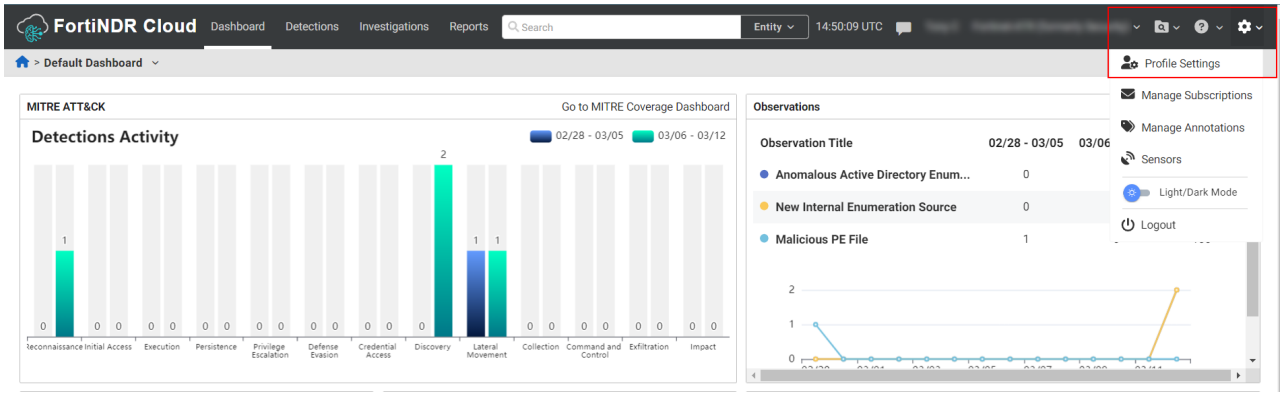

2. Under *Authentication*, click *Enable MFA*.

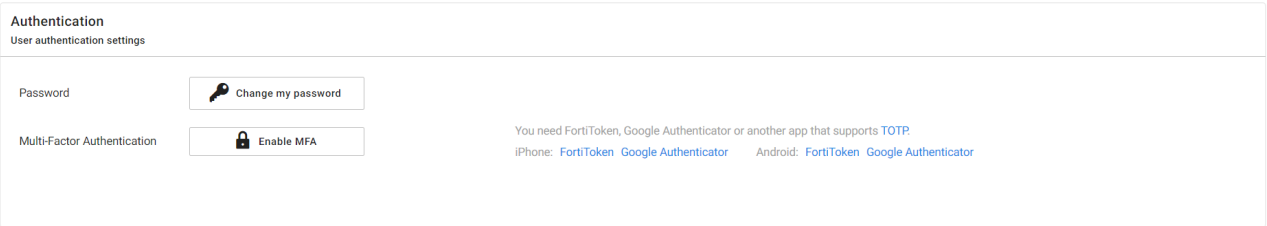

3. Scan the QR code with a token application to validate and enable MFA.

User activity timeout

Automatically log out users who belong to the account you are in. Users who only have access to the account are not affected by this setting.

1. Click the gear icon at the top-right of the page and select *Account Management*.
2. Select an account.
3. Click the *Settings* tab and scroll down to *User Activity Timeout*.
4. Enter a value between 15 and 480 minutes.
5. Click *Update*.

Disabling an account

Technical Success Managers can disable accounts that are either no longer in use or should no longer be in use. This option has the following effects:

- Disables login for all users in the account.
- Disables all notifications to those users.
- Stops ingest of all data.
- Removes the account from default account lists.

This can be completed by clicking the option icon in *Account Management* for a given account and then clicking on *Disable*.

Sensor email alerts

Administrators can create email notifications to alert you when sensor is offline or the event rate is low.

To create a sensor email alert:

1. Click the gear icon at the top-right of the page and select *Account Management*.
2. Select an account.
3. Click the *Settings* tab and scroll down to *Notification Emails*.
4. In the *Email* field, enter a recipient's email address.
5. Select *Sensor Offline Alert* and/or *Event Rate Low Alert*.
6. Click *Update*.
7. Click *Add Record* to add another email address.
8. Click **X** to delete an email address.

Device enrichment

You can enhance device identification using *Device Enrichment*. When configured, it retrieves hostname information from Windows Active Directory (AD) and DNS servers in the target network. Once enabled, the enrichment process runs on the schedule defined in the enrichment settings.

After a cycle completes, the process schedules the next cycle based on the profile settings. If the current cycle is still running when the next scheduled cycle is due, the system skips that cycle.

Once the profile is configured, it retrieves a list of devices and their names, performs DNS queries to resolve corresponding IP addresses, and sends detailed information for each device, including its name, IP address, operating system, and other attributes.

Only one sensor can be used for Device Enrichment per account.

Sensor requirements:

- Sensors running 2.4.0 or higher.

Network Requirements:

Active Directory and DNS queries originate from the sensor's management interface. Ensure that firewall policies allow the sensor's management IP to access the following:

- The LDAP server on port 389 or 636
- The DNS server on port 53 (or the configured DNS port)

Failure to allow these connections will prevent Active Directory synchronization.

To configure Active Directory:

1. Click the gear icon at the top-right of the portal and select *Account Management*.
2. Click *Settings*.

3. Scroll down to *Device Enrichment Configuration* and click *Configure*.
4. Configure the *Basic Settings*.

Setting	Description
Sensor ID	Select a Sensor from the dropdown.
Enabled	Toggle ON to enable.
LDAP Server	The IP address of the LDAP server.
Use SSL	Toggle ON to enable Secure Sockets Layer (SSL) encryption.
LDAP Port	The port used to communicate with the LDAP server. By default, FortiNDR Cloud uses port 636 . • If *Use SSL* is enabled, FortiNDR Cloud communicates over port 636 (LDAPS). • If *Use SSL* is disabled, it uses port 389 (LDAP).
Base DN	Enter the distinguished name (DN) of the part of the LDAP directory tree within which FortiNDR Cloud will search for user objects, such as `ou=People`, `dc=example,dc=com`.
Search DN	The base distinguished name (DN) in the LDAP directory where search operations start.
Bind DN	The LDAP user and its LDAP directory tree location for binding. For example, `CN=fndr_svc,CN=testUser`, `DC= example-domain`, `DC= com`.
Bind Password	The password for the LDAP user account for binding. For example, `DC= example-domain`, `DC= com`.
Search Scope	The method of retrieving the information from the tree: • *Base*: Only retrieve information from the base level of the directory tree specified in search base • *One Level*: Only retrieve information from the search base and one level down. • *Sub_tree*: Retrieve everything underneath the specified search base.
DNS Address	The IP address of the DNS server used to resolve domain names during Active Directory queries. By default, DNS queries use port 53. If the DNS server uses a non-standard port, specify it in the following format *"DNS_IP:DNS_PORT"*.
Run Start Time	The time of day when the Active Directory synchronization or data collection process begins.
Run Interval Hours	The frequency (in hours) at which the Active Directory synchronization or data collection process repeats.

5. Configure the *Advanced Settings*.

Setting	Description
Allow Insecure LDAP	Toggle ON to enable LDAP connections without SSL/TLS encryption.

Setting	Description
LDAP Timeout Seconds	The maximum time (in seconds) FortiNDR Cloud waits for an LDAP query to complete before timing out.
DNS Lookup Interval Seconds	The frequency (in seconds) DNS lookups are performed during Active Directory operations.
DNS Lookup Size	The number of DNS records retrieved in a single lookup.
DNS Lookup Retries	The number of retry attempts for DNS lookups if the initial attempt fails.
DNS Timeout Seconds	The maximum time (in seconds) FortiNDR Cloud waits for a DNS query to complete before timing out.

6. Click *Save*.

Device Enrichment Status

To view the Device Enrichment Status, go the Sensor's details page. See Sensor details on page 136.

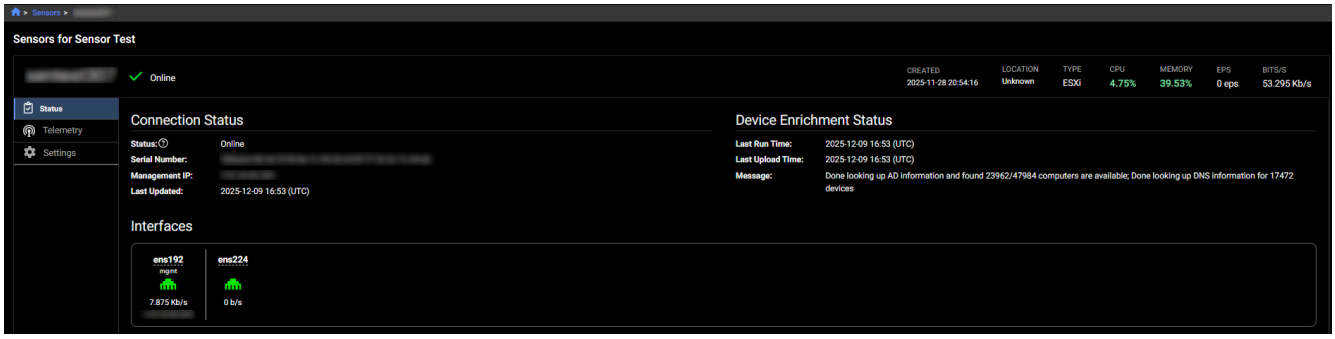

Add or edit subnets

The *Subnets* page lists all internal IP address ranges for the account. Admin users can add, edit or delete subnets in an account.

To add a subnet:

1. Click the gear icon in the top navigation and select *Account Management*.
   - If you have access to one account, the account page will appear.
   - If you have access to multiple accounts, select an account.
2. Click the *Subnets* tab and click *Add Subnet*. The *Add a Subnet* dialog opens.
3. Configure the subnet and click *Add Subnet*.

Subnet	Enter the IP address for the subnet.
Description	(Optional) Enter a description of the subnet.

Exteral	Select if this is an internal subnet that will be treated as external by Suricata.

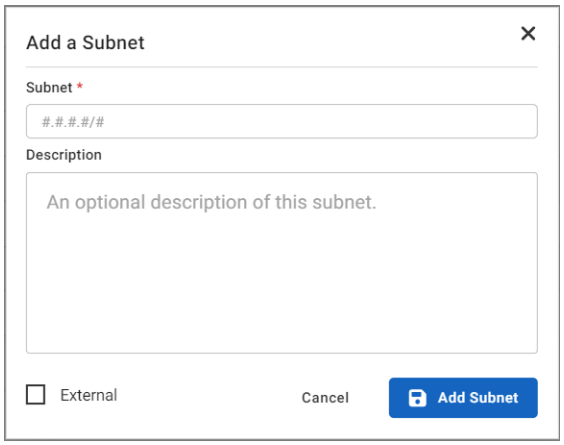

To edit a subnet:

1. Click the gear icon in the top navigation and select *Account Management*.
2. Click the *Subnets* tab.
3. In the *Actions* column, click the dropdown and select *Edit*. The *Update Subnet* dialog opens.
4. Edit the subnet and click *Update Subnet*.

To delete a subnet:

1. Click the gear icon in the top navigation and select *Account Management*.
2. Click the *Subnets* tab.
3. In the *Actions* column, click the dropdown and select *Delete*. The *Delete xx.xx.xxx.x/xx?* dialog opens.
4. Click *Confirm*.

To perform a bulk import:

1. Click the *CSV* button to download the current subnets.
2. Add or remove entries in the file and save it.
3. Click the *Import Subnets* button and upload the file. and re-upload the file.
4. Click the *Reset to Default* button to delete all subnets except the default.

Light/Dark Mode

The Light/Dark mode setting is saved to your browser. When you switch accounts, you will see the same theme as the previous user account. The mode you select does not affect other users with the same account.

To switch between light and dark mode:

Click the gear icon at the top-right of the page and toggle between *Light* and *Dark* mode.

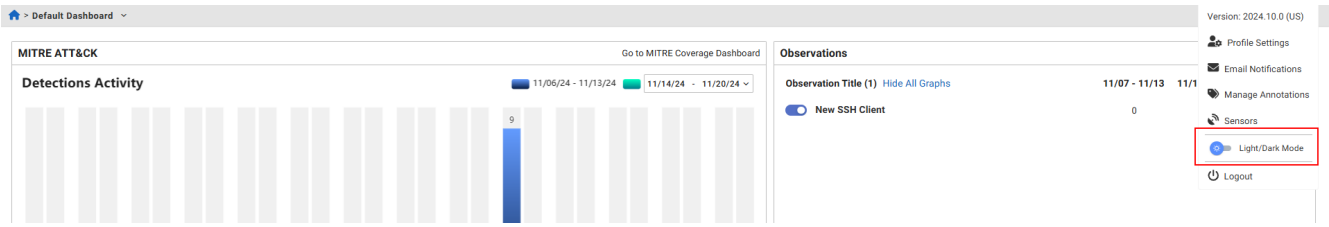
> Default Dashboard
Version: 2024.10.0 (US)
Profile Settings
Email Notifications
Manage Annotations
Sensors
Light/Dark Mode
Logout
MITRE ATT&CK
Go to MITRE Coverage Dashboard
Detections Activity
11/06/24 - 11/13/24
11/14/24 - 11/20/24
9
Observations
Observation Title (1)
Hide All Graphs
11/07 - 11/13
New SSH Client
0

Sensors deployment

FortiNDR Cloud deploys network sensors to monitor your virtual and physical on-premises infrastructure. Once deployed and configured, network metadata is collected and sent to FortiNDR Cloud for security analysis, threat detection, and indexing. A web application and application programming interface (API) are provided for analysis of security events. FortiNDR Cloud is delivered as a Software-as-a-Service (SaaS) and is fully managed by Fortinet, including network sensors.

The maximum size of the folder that stores the logs is 10G. Sensors are designed to retain logs for seven days. In the event of an issue affecting the upload, logs that are seven days and older will expire and are no longer available. Cleanup scripts are in place to automatically clean up the files when the log directory exceeds a certain size to prevent excessive disk usage.

Sensor specifications

Sensor Types

The following table lists the available sensor types and the maximum sustained throughput each type can consume.

Sensor Type	Form	Interfaces	NDR Sniffer Throughput*
FNDR Cloud 500F **Small sensor**	1U Server	2x 1G Copper 2x 10G SFP+ 2x 10G Copper	6 Gbps (metadata processing) across all ports
FNDR Cloud 900F **Large sensor**	1U Server	2x 1G Copper 2x 10G SFP+ 2x 10G Copper	13 Gbps (metadata processing) across all ports
FNDR Cloud 2540G **Extra large sensor**	2U Server	2x 10/25GbE SFP28 and 4x 1GbE RJ45 2x 10GbE RJ45 (breakout cable supported)	38 Gbps (metadata processing) across all ports
FNDR Cloud Virtual Sensors	OVF File	1 mgmt + min 1 TAP	Hypervisor dependent

*Using FortiTester default Enterprise Profile

Network interfaces for physical sensors

- 1 x 1Gbps Ethernet interface for management
- 1 x 1Gbps Ethernet interface for monitoring
- 2 x 10Gbps Ethernet interfaces for monitoring
- 2 x 10Gbps SFP (fiber) interfaces for monitoring

Minimum virtual sensor (ESX) host requirement

For details, the *ESXi Sensor Installation Guide*.

Network data sources

A network data source must be configured for the sensor. Sensors collect and process network data using standard network packet capture sources such as a network switch Switched Port Analyzer (SPAN) port or Test Access Port (TAP) device connected to a monitoring interface on the sensor.

SPAN (mirror) port

A SPAN port (sometimes called a mirror port) is a software feature built into a switch that creates a copy of selected packets passing through the device and sends them to a designated SPAN port. Using software on the network switch, an administrator can easily configure what data is monitored by a FortiNDR Cloud sensor connected to the SPAN port.

If the switch CPU is already heavily utilized prior to configuring a SPAN, SPAN data will likely be given a lower priority on the switch. The SPAN also uses a single egress port to aggregate multiple links, so it may become oversubscribed.

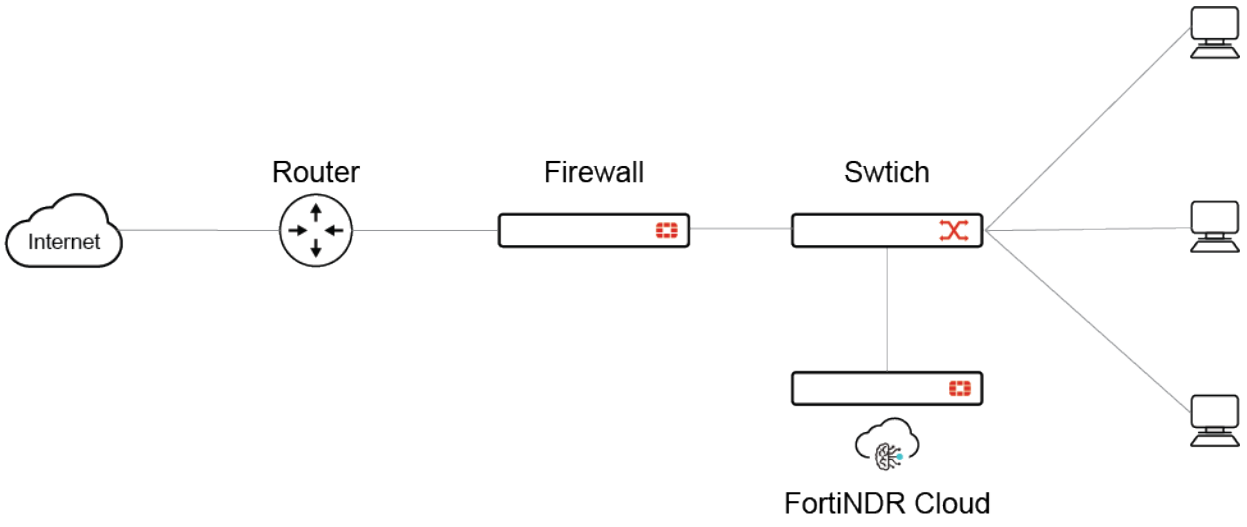

When to consider a SPAN port

- Limited ad hoc monitoring in locations with SPAN capabilities where a network TAP does not currently exist.
- Production emergencies where there is no maintenance window in which to install a TAP.
- Remote locations with modest traffic that cannot justify a full-time TAP on the link.
- Access to traffic that either stays within a switch or never reaches a physical link where the traffic can be TAPed.
- Locations with limited light budgets where the split ratio of a TAP may consume too much light.

Network TAP

A network TAP (Test Access Point) is a device that connects directly to the cabling infrastructure. Instead of two switches or routers connecting directly to each other, the network TAP sits between the two devices and all data flows through the TAP. Using an internal splitter, the TAP creates a copy of the data for monitoring while the original data continues unimpeded through the network.

This ensures every packet of any size will be copied. This technique also eliminates any chance of subscription overage. Once the data is TAPed, the duplicate copy can be sent to a FortiNDR Cloud sensor.

Inserting a TAP into an existing network link requires a brief cable disconnect. TAPs are typically installed during a maintenance window.

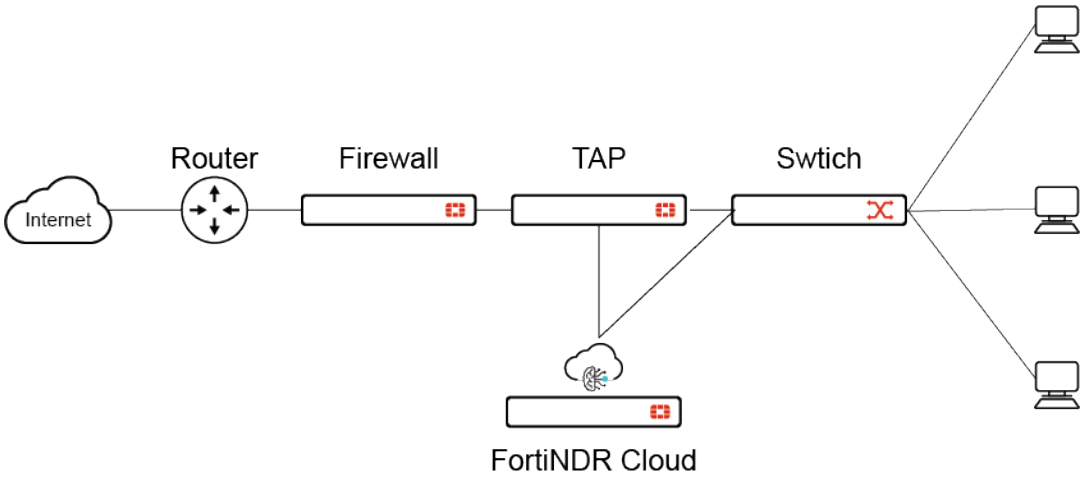

When to consider a network TAP

- Switch CPU already highly utilized and may drop packets.
- When additional load on the switch could impact network performance.
- No ports available on the switch.
- Hardware does not support SPAN functionality.

- When legal regulations or corporate compliance mandate that all traffic for a particular segment be monitored.

Not sure which data source(s) to use? Ask your FortiNDR Cloud representative.

Network aggregator

For many organizations, a network aggregator is configured to monitor traffic at several key locations within the network. FortiNDR Cloud sensors can deploy off a network aggregator if one is available within the network. Some network aggregation appliances also have the ability to decrypt network traffic, which can greatly increase the fidelity and visibility of the FortiNDR Cloud sensor.

Network aggregators are also commonly used to monitor traffic from networks with 40Gbps links. In this case, an aggregator is utilized to split traffic from a 40Gbps line to four separate FortiNDR Cloud appliances monitoring up to 10Gbps per sensor.

ERSPAN

Starting with version 2.4.0, ERSPAN Type II and Type III can be used to forward packets from network devices that support ERSPAN to the sensor for analysis.

Complex or combination deployments

Multiple FortiNDR Cloud sensors can be deployed to obtain full visibility across the environment. Each sensor reports back to the FortiNDR Cloud, providing cross-enterprise visibility through a single, unified platform. Queries can be executed against data from all sensors, or a subset as specified by an analyst.

Sensor deployment strategy

Sensor placement is prioritized for network locations where security events are most likely to occur. Data collected from multiple locations provides a complete and accurate picture of potential security threats. Below is a prioritized list of data source locations in a typical network environment.

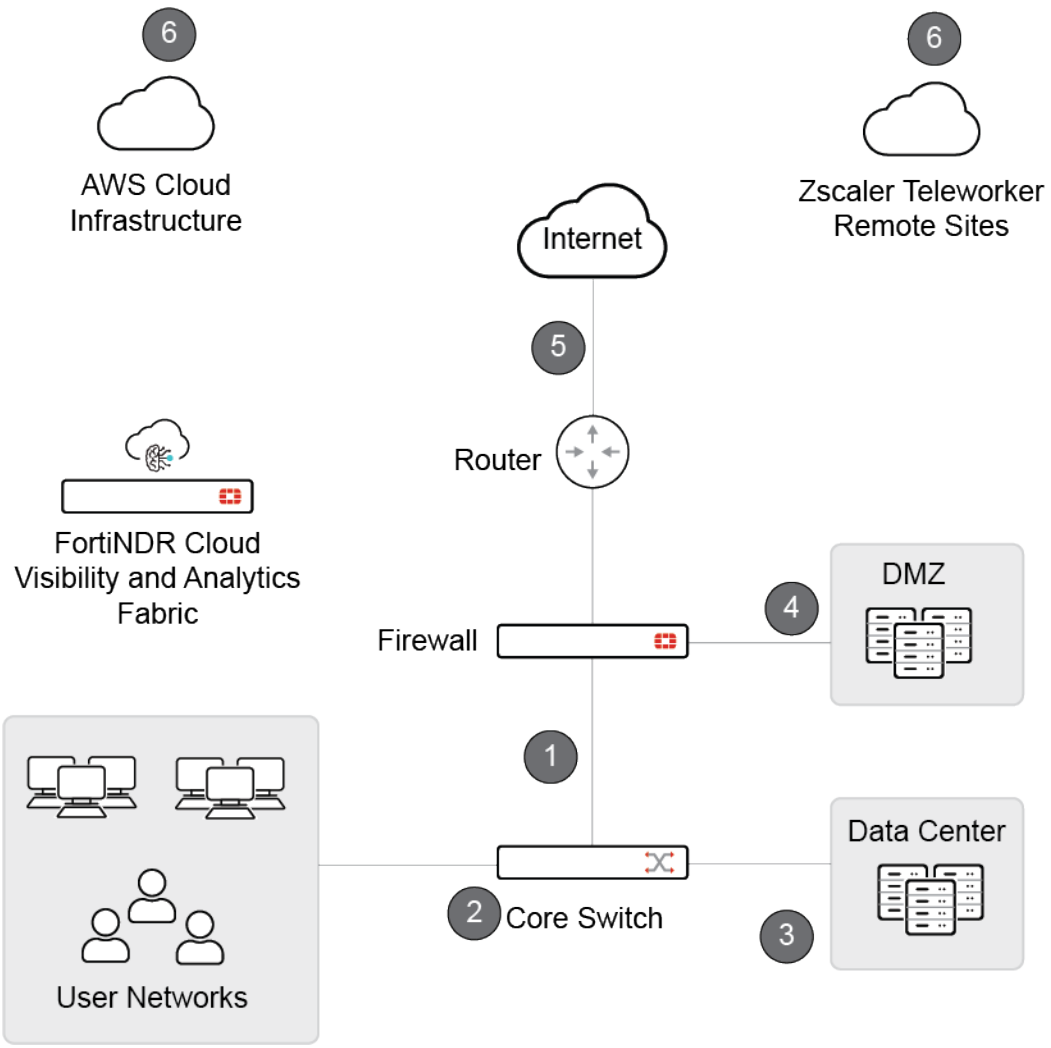

Number	Location	Description
1	**Egress Points**	Monitoring activity between your network environment and the Internet provides visibility of security events related to malware beaconing, command and control, network tunneling and data exfiltration activity. Benefits: • Captures north/south traffic from clients and servers • Enables detection of exfiltration, C2, tunneling, beaconing
2	**Core Switch**	Activity within your network can include security events related to lateral movement and staging of attacks between workstations and important internal resources such as internal web applications, file servers or your system infrastructure.

Number	Location	Description
		Benefits: • Captures east/west traffic between clients and servers • Enables detection of lateral movement, staging, internal threats
3	**Data Center**	Your data center infrastructure is where your valuable information is stored, making it a target for theft and unauthorized access. Sensors placed between these servers and virtual hosts provide visibility of security events related to this activity. Benefits: • Captures east/west traffic between servers (including virtual) • Enables detection of data theft, unauthorized access
4	**DMZ**	Public facing applications such as mail services, web sites and business-to-business applications are constantly attacked. Monitoring network zones that host these applications provides visibility of security events related to unauthorized access and data exfiltration. Benefits: • Captures north/south traffic between DMZ and external clients • Enables detection of unauthorized access, vulnerability exploitation, exfiltration
5	**External Link**	Benefits: • Captures north/south traffic between external clients and the internal networks. Provides visibility to traffic even if it is blocked by the firewall • Enables detection of exploitation attempts
6	**Cloud Visibility**	Benefits: • Cloud infrastructure workload traffic analysis via AWS/Azure Machine Images or VM/KVM. • Teleworker and Remote Sites not backhauled to VPN via Zscaler integration. • Enables detection of un-managed and IoT devices and access to cloud infrastructure

Sensor data source configuration

For instructions on sensor data source configuration for VMware ESX, see the *ESXi Sensor Installation Guide*.

Collector interface

The collector interface is a special type of TAP/monitoring sensor interface with the IP stack enabled. It is designed for packet-forwarding solutions that require a destination IP address, such as NetFlow and ERSPAN.

Configuring the collector Interface

Prerequisites

Before configuring a collector interface, ensure that the following prerequisites are met regarding the number of sensor ports required for each feature you plan to enable:

Feature / Use Case	Azure, OCI, AWS (VXLAN Monitoring)	Other Platforms
SPAN Only	1 × Management	1 × Management, Min 1 × TAP
NetFlow Only	1 × Management, 1 × Collector	1 × Management, 1 × Collector
ERSPAN Only	1 × Management, 1 × Collector	1 × Management, 1 × Collector, Min 1 × TAP
SPAN + ERSPAN	1 × Management, 1 × Collector	1 × Management, 1 × Collector, Min 1 × TAP
SPAN + ERSPAN + NetFlow	1 × Management, 1 × Collector	1 × Management, 1 × Collector, Min 1 × TAP

Interface requirements:

Interface	Requirment
TAP	Requires a physical uplink connection to receive mirrored traffic.
Collector	The IPv4 stack must be enabled, and an uplink is required to receive ERSPAN and NetFlow traffic.

To configure the collector interface:

It is best practice to keep the collector and management interfaces on separate subnets.

1. From the config menu, select *Set Collector Interface* (Press c).
   - Highlight the monitoring interface you want to use as the collector.
   - Ensure this interface has an IP stack enabled on the network.
2. If DHCP is available on the collector subnet, choose *Configure Using DHCP* and select *Submit*.

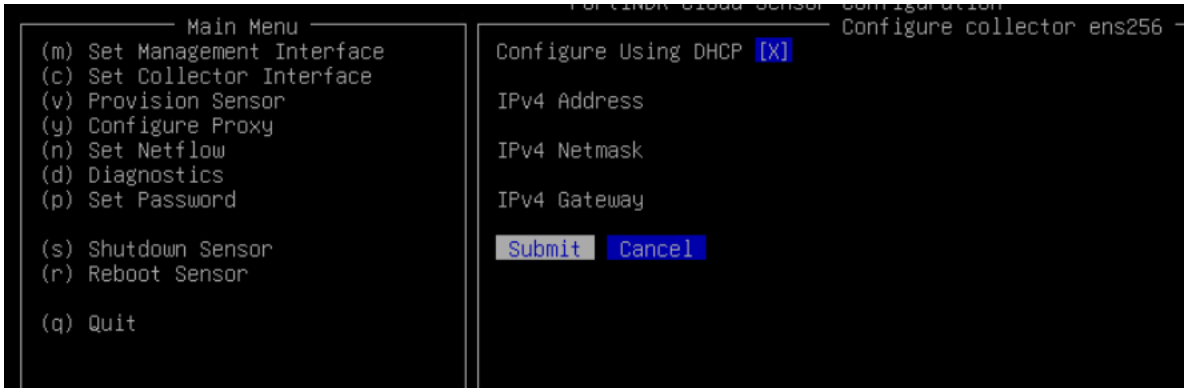

3. To configure a static IP on collector interface:
   a. Deselect the DHCP by pressing the space bar.
   b. Enter the desired *Address*, *Netmask*, and default *Gateway*.
   c. Select *Submit*.

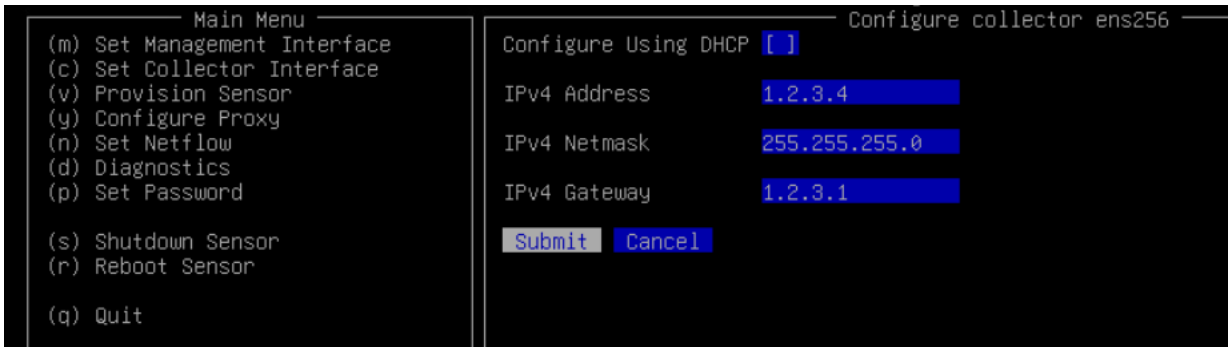

4. The menu will redirect to the *Interfaces* section with the collector interface reflecting your settings.

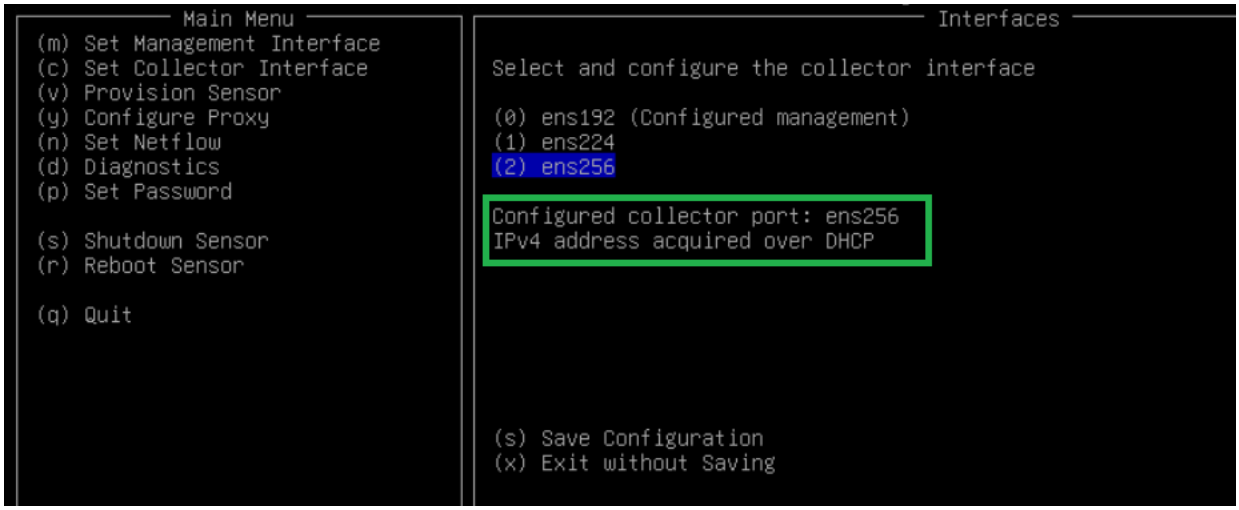

   a. Select *Save Configuration*. A confirmation dialog box will appear, requesting a sensor restart to apply the interface changes. Select *Yes* to proceed with the restart.

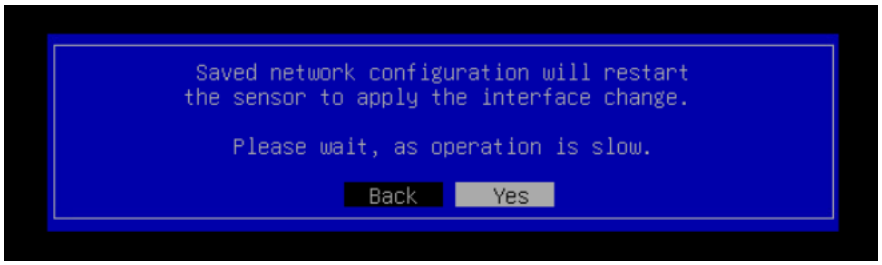

b. Wait a few minutes until the restart completes and the message *Successfully restarted sensord* appears.

c. Press *Enter* to return to the main menu.

The collector's IP address will now appear in the status pane. Allow a few minutes for the sensor to update its status to *Online*.

NetFlow

NetFlow is a network monitoring protocol widely used for collecting and analyzing IP traffic. It provides visibility into network usage, application behavior, and potential threats by exporting flow records to a collector.

Starting from version 2.3.0, FortiNDR Cloud sensor can operate as a NetFlow collector, enabling network devices to send flow data for behavioral analysis and threat detection.

To use this feature, point your flow exporters to FortiNDR Cloud sensor collector's IP and port. The sensor listens on UDP/2055 (NetFlow v5, v9, IPFIX) and UDP/6343 (SFlow) by default, with ports configurable as needed.

To view the complete list of NetFlow fields, see NetFlow fields.

A separate Log Ingestion license is required to collect NetFlow data. Without this license, the data will not be visible in the portal.

Refer to your NetFlow exporter configuration to verify supported transport protocols (UDP) and ensure inbound firewall rules allow traffic on the configured NetFlow(s) Flow ports.

Configuring NetFlow for FortiNDR Cloud

To verify the sensor status:

1. Log into the sensor console using:
   - Username: `config`
   - Password: (The password set during initial installation)
2. Confirm that the sensor is *Online* and the collector interface is up with an IP address. See Collector interface on page 168.

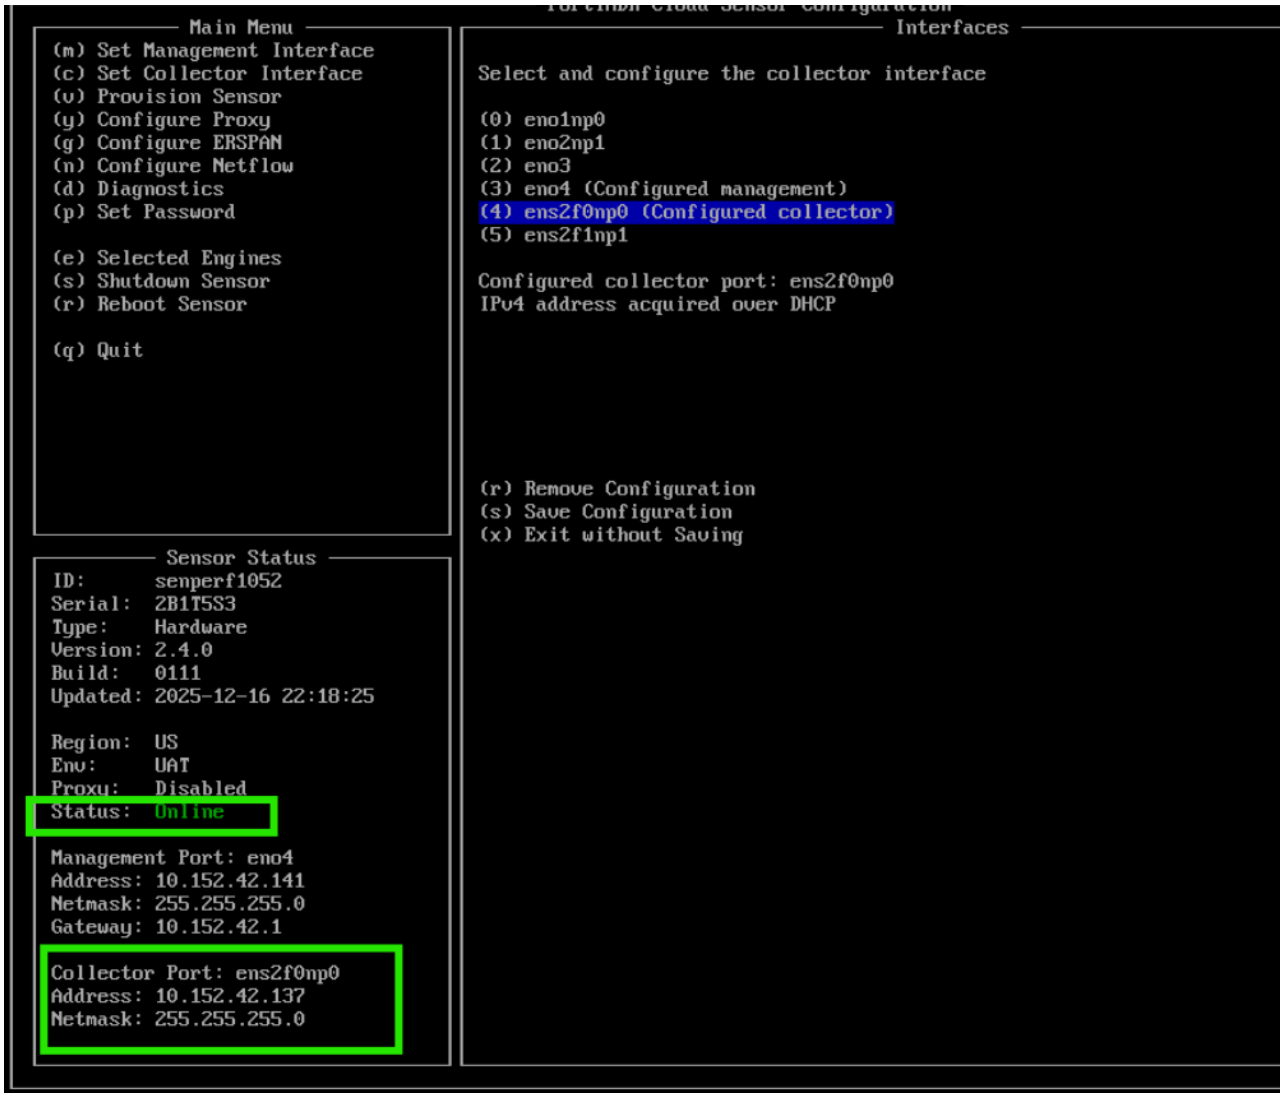

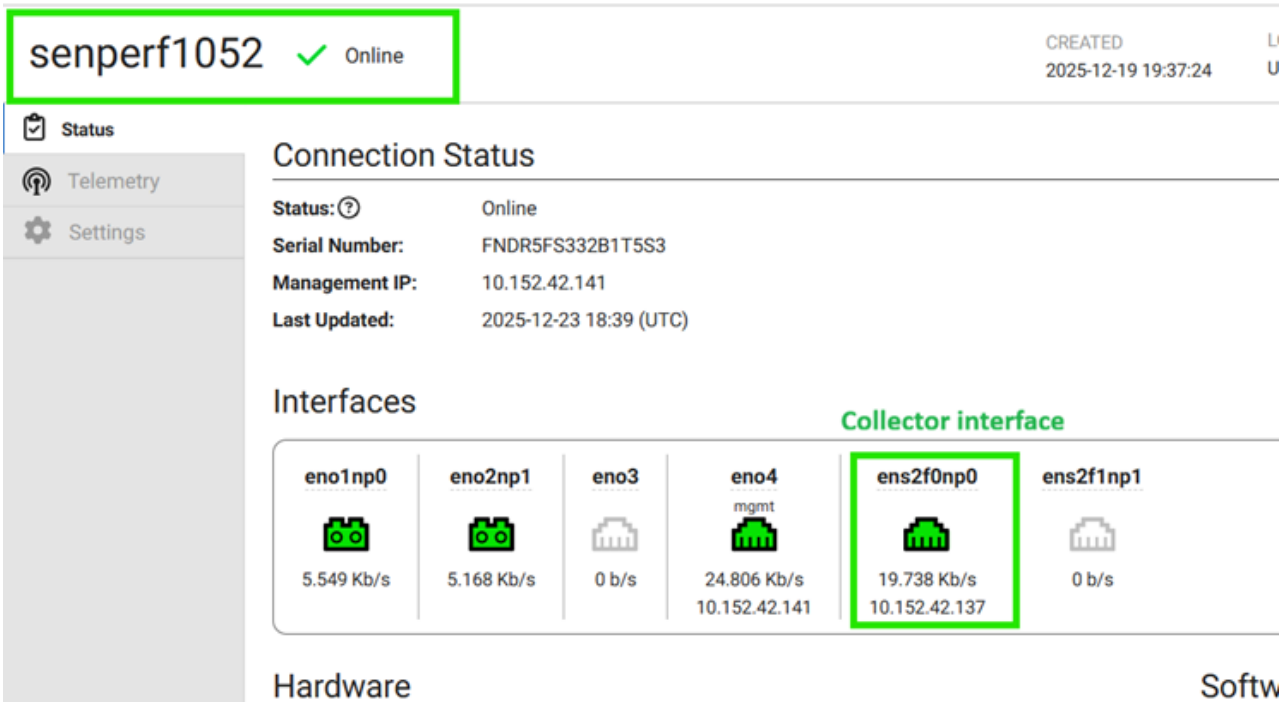

3. From the sensor config menu, select *Configure Netflow* (or press n). Then select *Configure* (or press c).

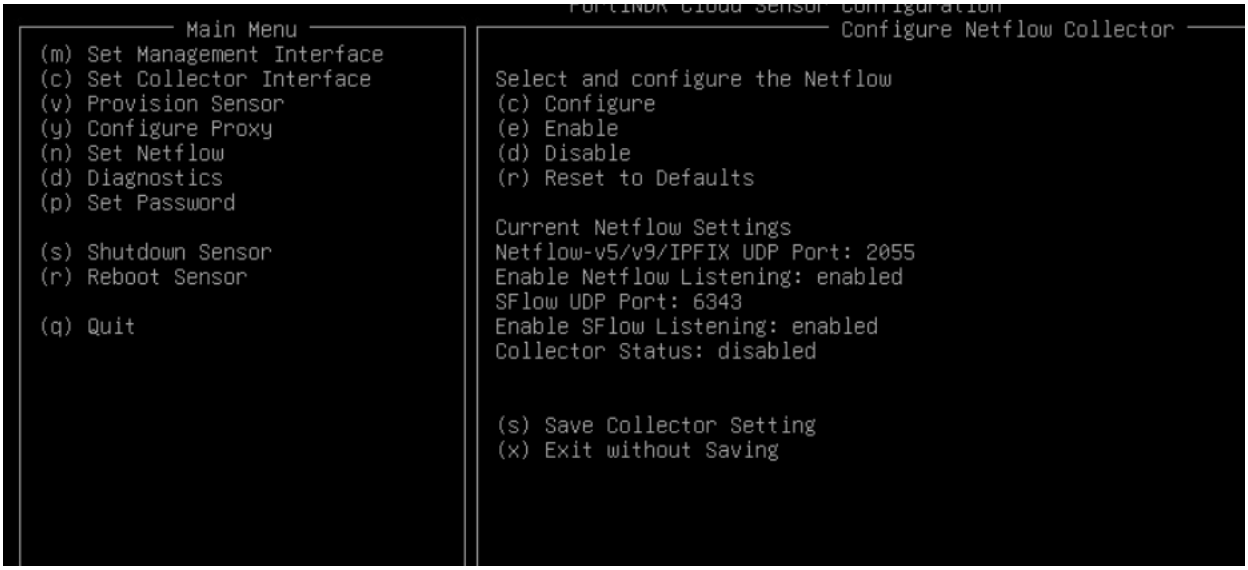

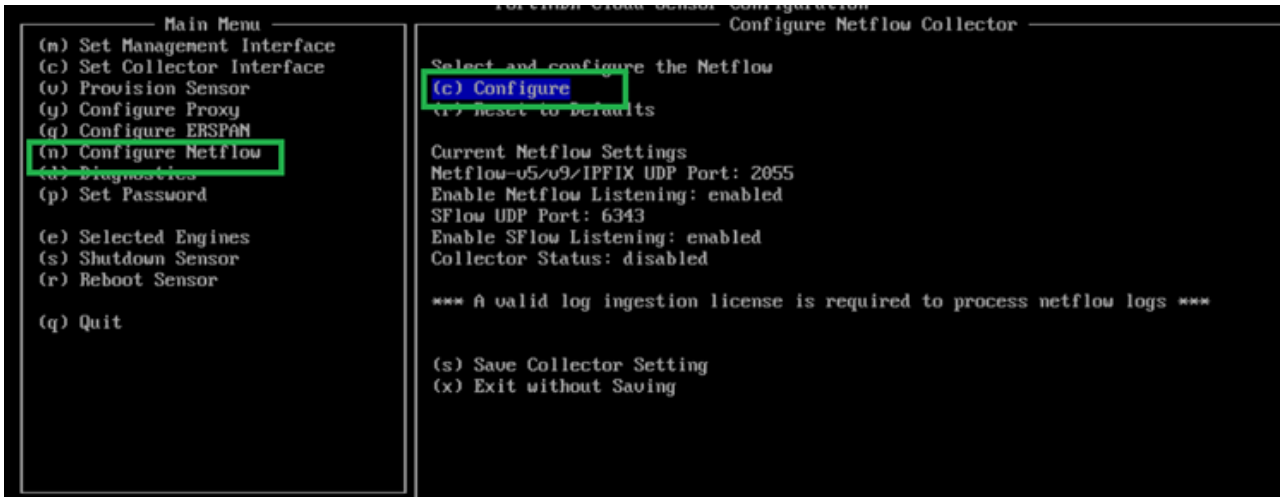

4. Review the default NetFlow settings.
   - UDP/2055: Netflow v5, v9, IPFIX
   - UDP/6343: SFlow
5. If changes are required, select Configure (press c), adjust the port or listening status, and select Submit. Ensure *Collector Enabled* is checked

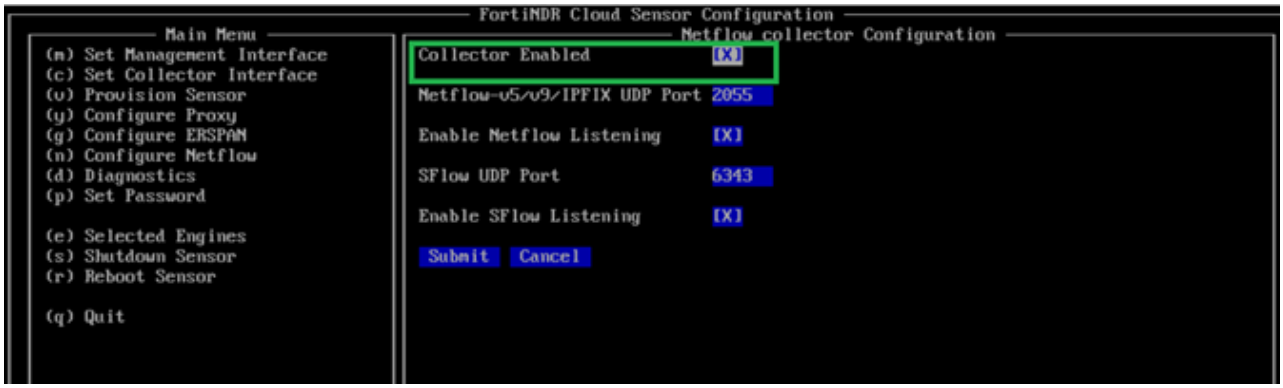

6. The menu will redirect back to the Netflow config menu. To save changes, select *Save Collector Setting* (press s).

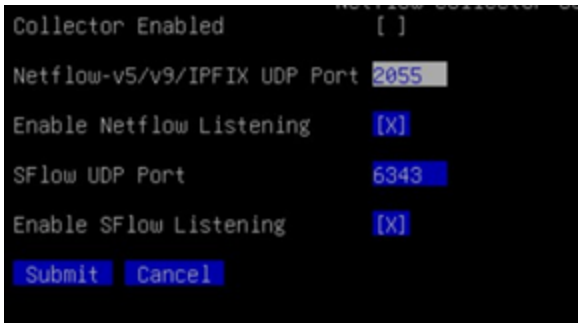

ERSPAN

ERSPAN (Encapsulated Remote SPAN) mirrors traffic from one or more source interfaces and encapsulates it using GRE so it can traverse a routed IP network.

Starting from version 2.4.0, the FortiNDR Cloud sensor can forward ERSPAN packets for threat analysis.

FortiNDR Cloud supports ERSPAN Type II and Type III.

Refer to your switch, router, or firewall documentation for instructions on forwarding ERSPAN packets to the sensor. Ensure that your firewall allows inbound GRE traffic to the sensor.

Enabling ERSPAN on the FortiNDR Cloud sensor

To enable ERSPAN on the cloud sensor:

1. Log into the sensor console using:
   - Username: `config`
   - Password: (The password set during initial installation)
2. Confirm that the sensor is online and the collector interface is up with an assigned IP address (see the Collector Port Configuration section in this document).

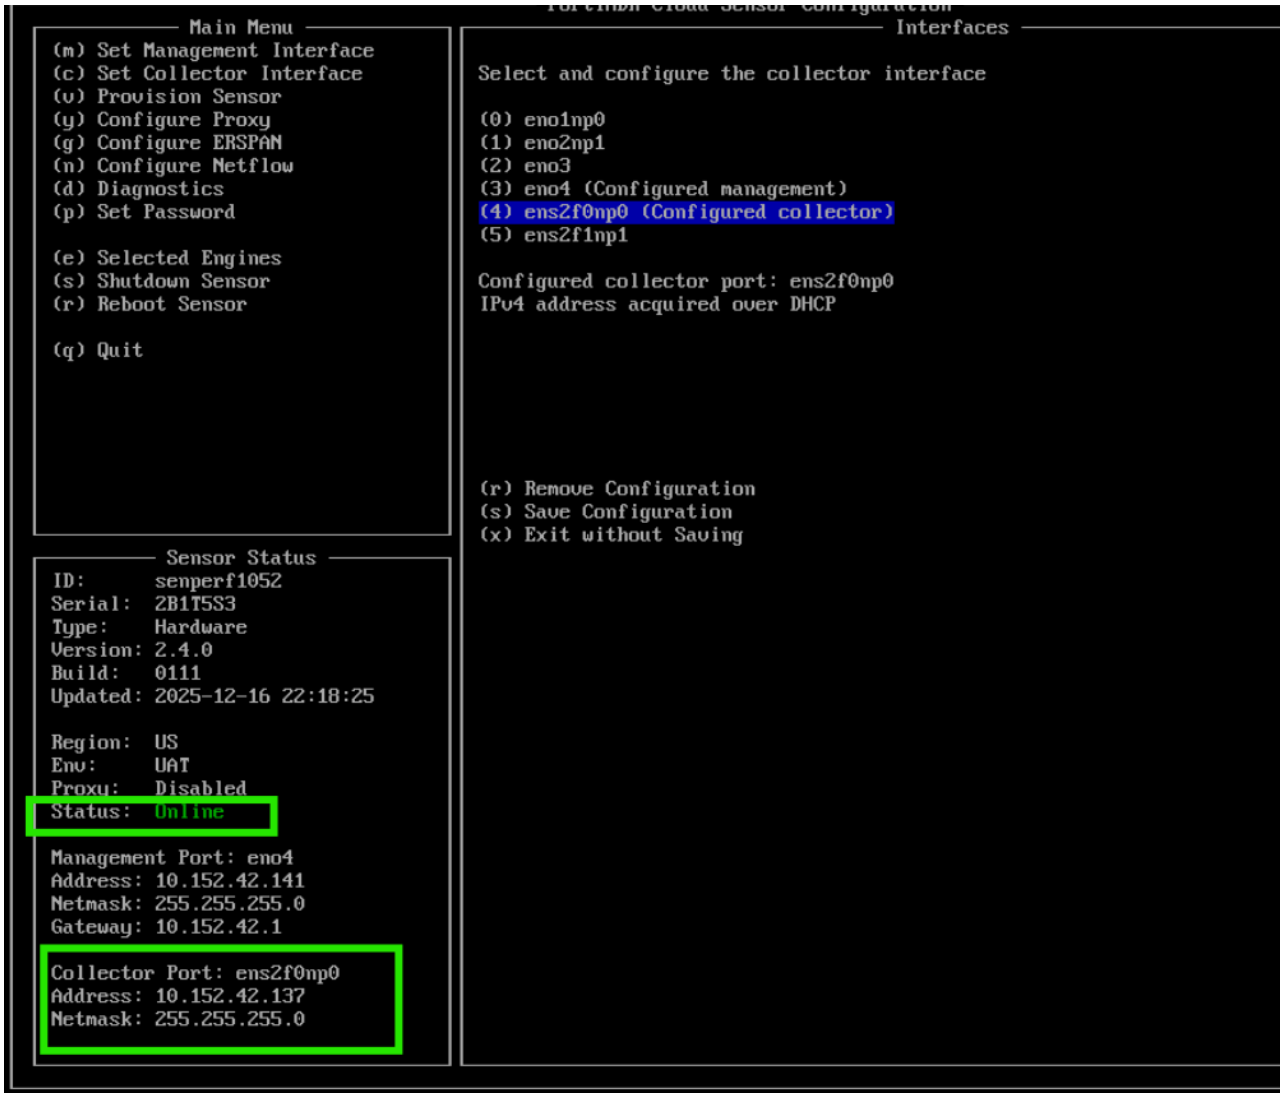
Main Menu
(m) Set Management Interface
(c) Set Collector Interface
(v) Provision Sensor
(y) Configure Proxy
(g) Configure ERSPAN
(n) Configure Netflow
(d) Diagnostics
(p) Set Password
(e) Selected Engines
(s) Shutdown Sensor
(r) Reboot Sensor
(q) Quit
Sensor Status
ID: senperf1052
Serial: 2B1T5S3
Type: Hardware
Version: 2.4.0
Build: 0111
Updated: 2025-12-16 22:18:25
Region: US
Env: UAT
Proxy: Disabled
Status: Online
Management Port: eno4
Address: 10.152.42.141
Netmask: 255.255.255.0
Gateway: 10.152.42.1
Collector Port: ens2f0np0
Address: 10.152.42.137
Netmask: 255.255.255.0
Interfaces
Select and configure the collector interface
(0) eno1np0
(1) eno2np1
(2) eno3
(3) eno4 (Configured management)
(4) ens2f0np0 (Configured collector)
(5) ens2f1np1
Configured collector port: ens2f0np0
IPv4 address acquired over DHCP
(r) Remove Configuration
(s) Save Configuration
(x) Exit without Saving

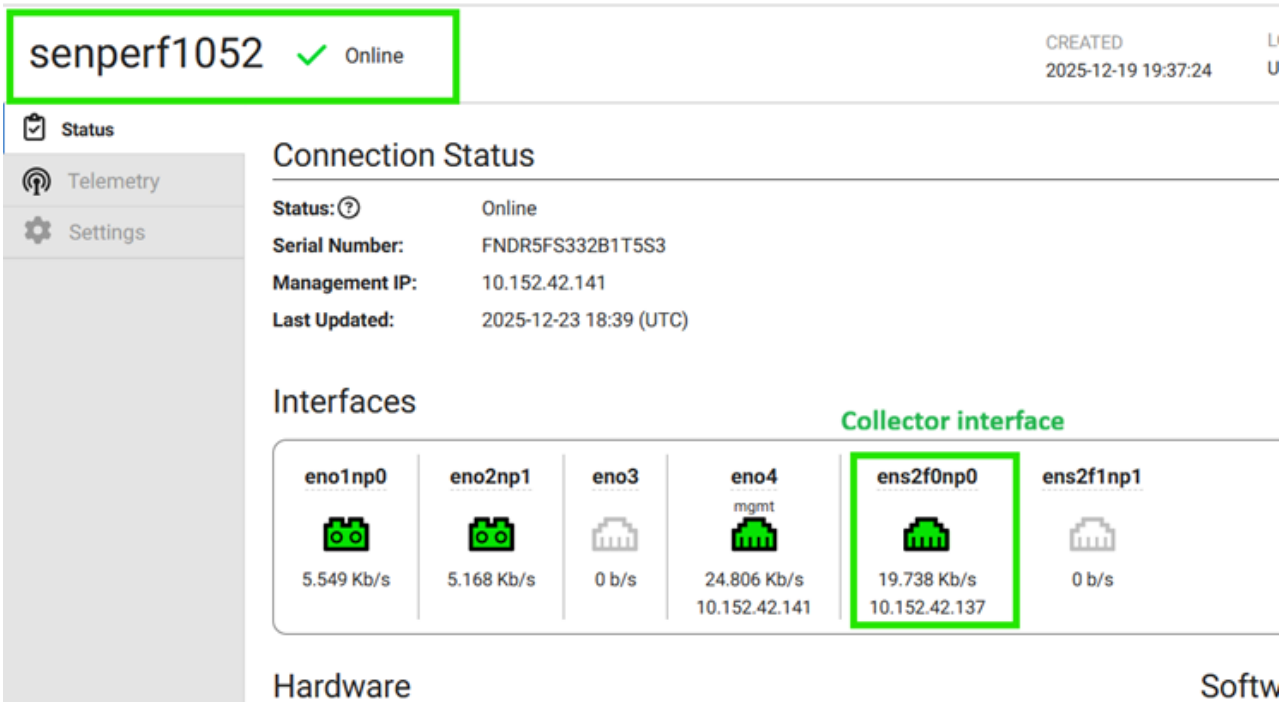

3. Set up the collector interface with DHCP or static IP. See Collector interface on page 168.
4. From the sensor config menu, select *Configure ERSPAN* (or press g).
5. Select *Enable* (or press e).
6. Select *Yes* in the next menu.

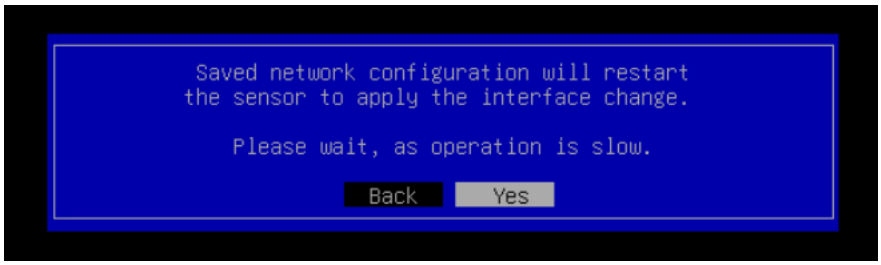

7. Once the sensor status pane is populated, sensor is ready to receive ERSPAN packets.
8. Configure your switch/router/firewall to send ERSPAN packets to the sensor's collector IP.
9. On the FortiNDR cloud portal, verify that the *veth_erspan* exists.

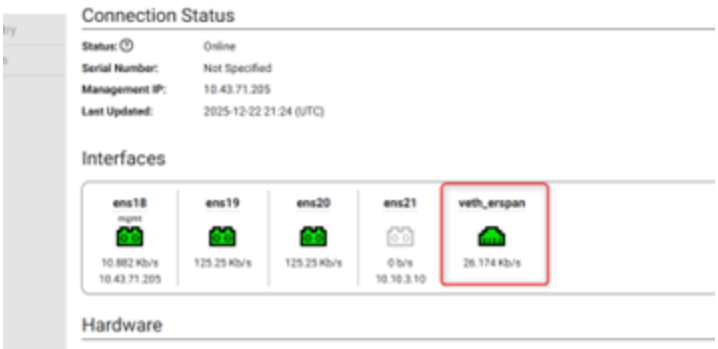

Zscaler ingestion

Zscaler ingestion provides FNDRC with remote access activity logs. When enabled, FortiNDR Cloud can identify threats impacting remote users. FortiNDR Cloud users can use the detection data to conduct investigations and entity searches to identify the threat source and mitigate attacks on the network.

- Zscaler setup on page 177
- Zscaler events on page 179

Zscaler setup

Cloud NSS

Zscaler Cloud NSS is a managed service from Zscaler. When using Cloud NSS, you do not need to deploy the NSS Virtual Machines. Cloud NSS sends logs to a HTTP endpoint or an S3 bucket. The integration with FortiNDR is through the S3 bucket path. Check with your Zscaler Account team to ensure you have this subscription enabled.

Cloud NSS Setup for S3

Ensure that you have the following to configure Zscaler Cloud NSS. Contact Fortinet Support to obtain these values.

- AWS Access Id
- AWS Secret Key
- S3 Folder URL

Using S3 requires the correct set of permissions and configuration. To learn more, see the Zscaler and S3 Deployment Guide, section Zscaler Cloud NSS with Amazon S3, on setting up S3 to work with Cloud NSS.

Configuring Cloud NSS for Web Logs

The following configuration information was adapted from the *Zscaler and Fortinet Deployment Guide*.

To configure Cloud NSS for Web Logs:

1. Log in as an administrator and go to *Administration > Nanolog Streaming Service*.
2. Go to *Cloud NSS Feeds* and click *Add Cloud NSS Feed*.
3. In the *Add Cloud NSS Feed* dialog, configure the following:

Feed Name	Enter a Feed Name.
NSS Type	Select *NSS for Web*.

Status	*Enabled*
SIEM Rate	*Unlimited*
SIEM Type	*S3*
AWS Access Id	Enter the access ID.
AWS Secret Key	Enter the secret key.
S3 Folder URL	Enter the folder URL.
HTTP Headers	Enter a dummy HTTP key and value pair. This is required.
Log Type	Select *Web Log*.
Feed Output Type	Select *Custom*.
Feed Escape Character	Enter `,\"`
Feed Output Format	`zscaler_log_type=web\ttimestamp=%d{yyyy}-%02d{mth}-%02d {dd}T%02d{hh}:%02d{mm}:%02d{ss} Z\tzscaler_recordid=%d {recordid}\tzscaler_proto=%s{proto}\tsrc_ip=%s{cip}\tdst_ ip=%s{sip}\tstatus_code=%s{respcode}\tmethod=%s {reqmethod}\tuser_agent=%s{ua}\ treferrer=%s {ereferer}\trequest_length=%d{reqsize}\tresponse_length=%d {resp- size}\turi=%s{eurl}\tfile_md5=%s{bamd5}\tcontent_ type=%s{contenttype}\tclient_ci- pher=%s {clientsslcipher}\tclient_version=%s {clienttlsversion}\tserver_cipher=%s{s- rvsslcipher}\tserver_ version=%s{srvtlsversion}\tzscaler_username=%s{login}\ tzscaler_hostname=%s{devicehostname}`

Configuring Cloud NSS for Firewall Logs

To configure Firewall logs, follow the steps in Configuring Cloud NSS for Web Logs on page 177 with the following exceptions.

NSS Type	Select *NSS for Firewall*.
Log Type	Select *Firewall Logs*.
Firewall Log Type	Both Session and Aggregate Logs
Feed Output Format	`zscaler_log_type=firewall\ttimestamp=%d{yyyy}-%02d{mth}-%02d{dd} T%02d{hh}:%02d{mm}:%02d{ss}Z\tzscaler_recordid=%d {recordid}\tsrc_ip=%s{c- sip}\tsrc_port=%d{csport}\tdst_ip=%s {cdip}\tdst_port=%d{cdport}\tdura- tion=%d {durationms}\tprotocol=%s{ipproto}\tservice=%s{nwsvc}\trequest_ bytes=%ld{outbytes}\tresponse_bytes=%ld{inbytes}\tzscaler_ username=%s{login}\`

Configuring Cloud NSS for DNS Logs

To configure DNS logs, follow the steps in Configuring Cloud NSS for Web Logs on page 177 with the following exceptions.

NSS Type	Select *NSS for Firewall*.
Log Type	Select *DNS Logs*.
Feed Output Format	`zscaler_log_type=dns\ttimestamp=%d{yyyy}-%02d{mth}-%02d{dd}T%02d{hh}:%02d{mm}:%02d{ss} Z\tzscaler_recordid=%d{recordid}\tsrc_ip=%s{cip}\tdst_ip=%s{sip}\tdst_port=%d{sport}\ tquery=%s{req}\tqtype_name=%s{reqtype}\tresponse=%s{res}\tzscaler_username=%s{login}\`

Zscaler events

Zscaler logs are mapped to the following FortiNDR Cloud event types. Events from Zscaler can be identified by `source="Zscaler"`.

- DNS
- Flow
- HTTP
- SSL

DNS

Field	Comments
answers	Zscaler provides a single answer.
qtype	This is derived from `qtype_name`, so it may be missing for unexpected values.
rcode	This is derived from `rcode_name`, so it may be missing for unexpected values.
rcode_name	Zscaler also uses this as an error field, so it may contain unexpected values that are passed through.
src.ip	

Flow

Field	Comments
dst.ip	

Field	Comments
dst.ip_bytes	
dst.port	
duration	
proto	The values are mostly passed through from Zscaler. Some values will match and others will not.
service	The values are mostly passed through from Zscaler. Some values will match and others will not.
src.ip	
src.ip_bytes	
src.port	
total_ip_bytes	
upload_percent	

HTTP

Field	Comments
headers.content_type	Zscaler may be translating some values into human-readable forms (for example, *Flash*).
method	Zscaler provides a value of *CONNECT* for *HTTPS*.
referrer	Zscaler does not provide the scheme (for example., `http://`).
request_len	
response_len	
src.ip	
status_code	
uri	
user_agent	

SSL

Every HTTPS request will have both an HTTP and SSL event. SSL events are only available for HTTPS. Also, Zscaler documentation suggests that it can be configured to intercept SSL. In that case, the cipher and version field represents the server, which may be different from the values for the client.

Field	Comments
cipher	Zscaler values are passed through without conversion.
dst.ip	
src.ip	
server_name	
server_name_indication	
version	Zscaler values are converted, but unexpected values will be passed through.

Sensor provisioning

FortiNDR Cloud sensors are self-provisioning appliances that require a registration code from the portal.

To provision a sensor:

1. Generate a registration code on page 181
2. Register a sensor on page 182

Once these steps are complete, the sensor will call home, provision itself, and then be ready to ingest raw mirrored traffic. By default, a sensor will use DHCP but a static IP address can be set if desired.

FortiNDR Cloudsupports unlimited sensors. For deployments involving more than 10 sensors, we recommend customers work with their TSM to ensure best practices are followed and the configuration is optimized.

Generate a registration code

Registration codes can be generated on the *Sensors* page\ within FortiNDR Cloud. If you do not have access to this page, please contact your Fortinet representative.

- Codes expire 24 hours after creation
- Codes may be used to provision multiple sensors prior to expiration
- Codes work for both physical and virtual sensors
- Each account is limited to ten (10) sensors by default. To expand this limit, contact your Technical Success Manager

To generate a registration code:

1. Click the *Settings* icon at the top right of the page and select Senors. The *Sensors* page opens.

2. In the toolbar, click , *Actions* > *Provision Sensor*. The *New Registration Code* dialog displays a randomly generated registration code prepended with the sensor code for its respective account.
3. If you have access to multiple accounts, verify that the generated code begins with the three- letter sensor code of the proper account.
4. Register the sensor.

Be sure to write the code down or copy it locally as it will not be shown again after the pop-up box is closed. If you accidentally close the pop-up box before copying down the code, simply generate another code.

Register a sensor

Registering the sensor takes place within the sensor console. Once registered, the sensor will call home, provision itself, and then be ready to ingest raw mirrored traffic.

See **Verifying Network Connectivity** to troubleshoot connectivity issues.

Registering a sensor requires an Internet connection. =Please ensure that the appliance is connected before proceeding.

To register a sensor:

1. Log in to the sensor console.
2. Select *Provision Sensor* or type `v`.

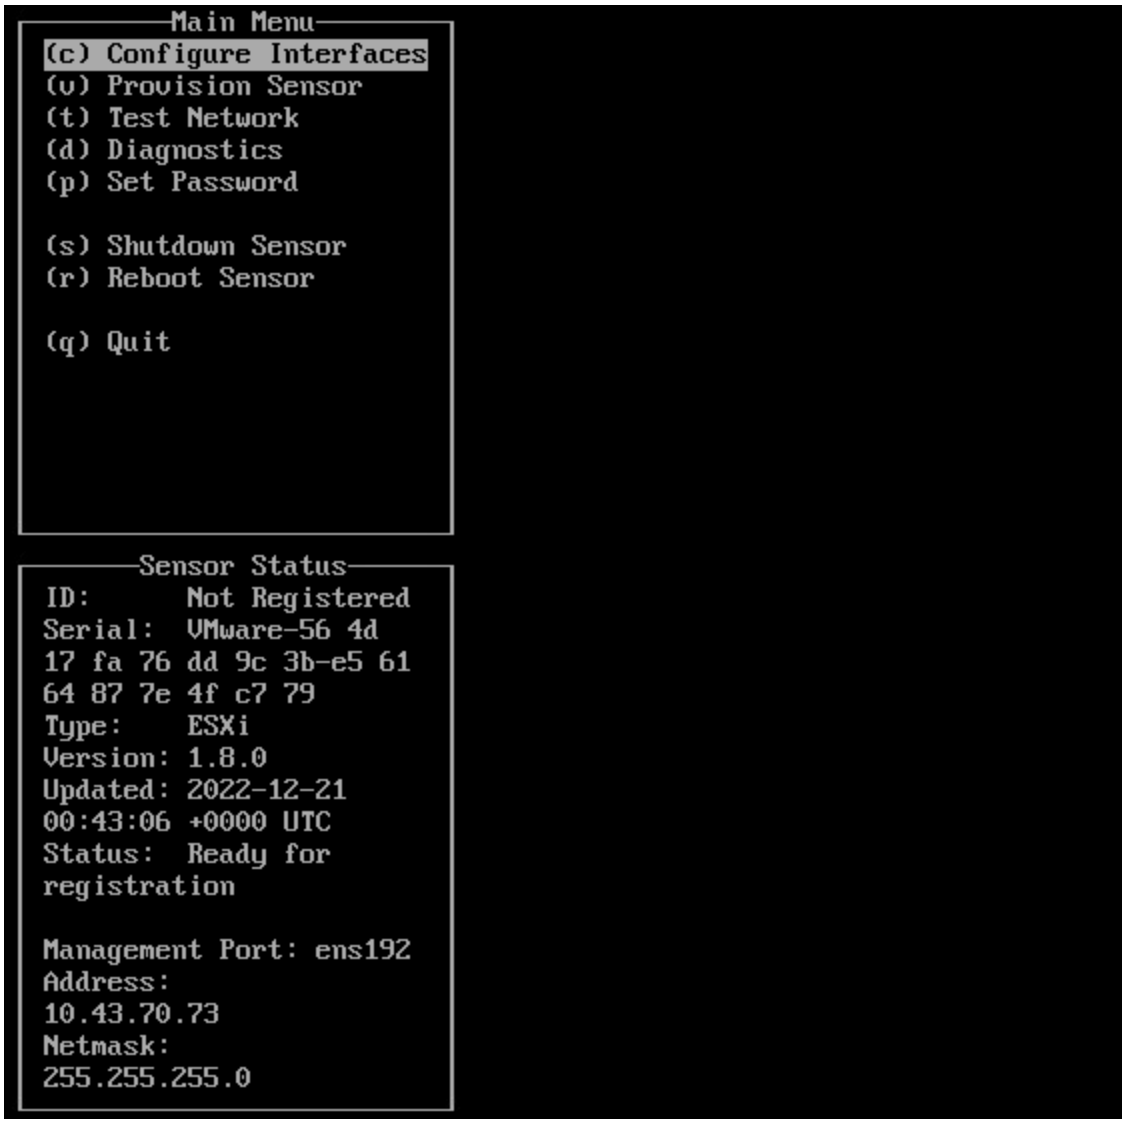

3. Enter the registration code in the text box. See Generate a registration code on page 181.

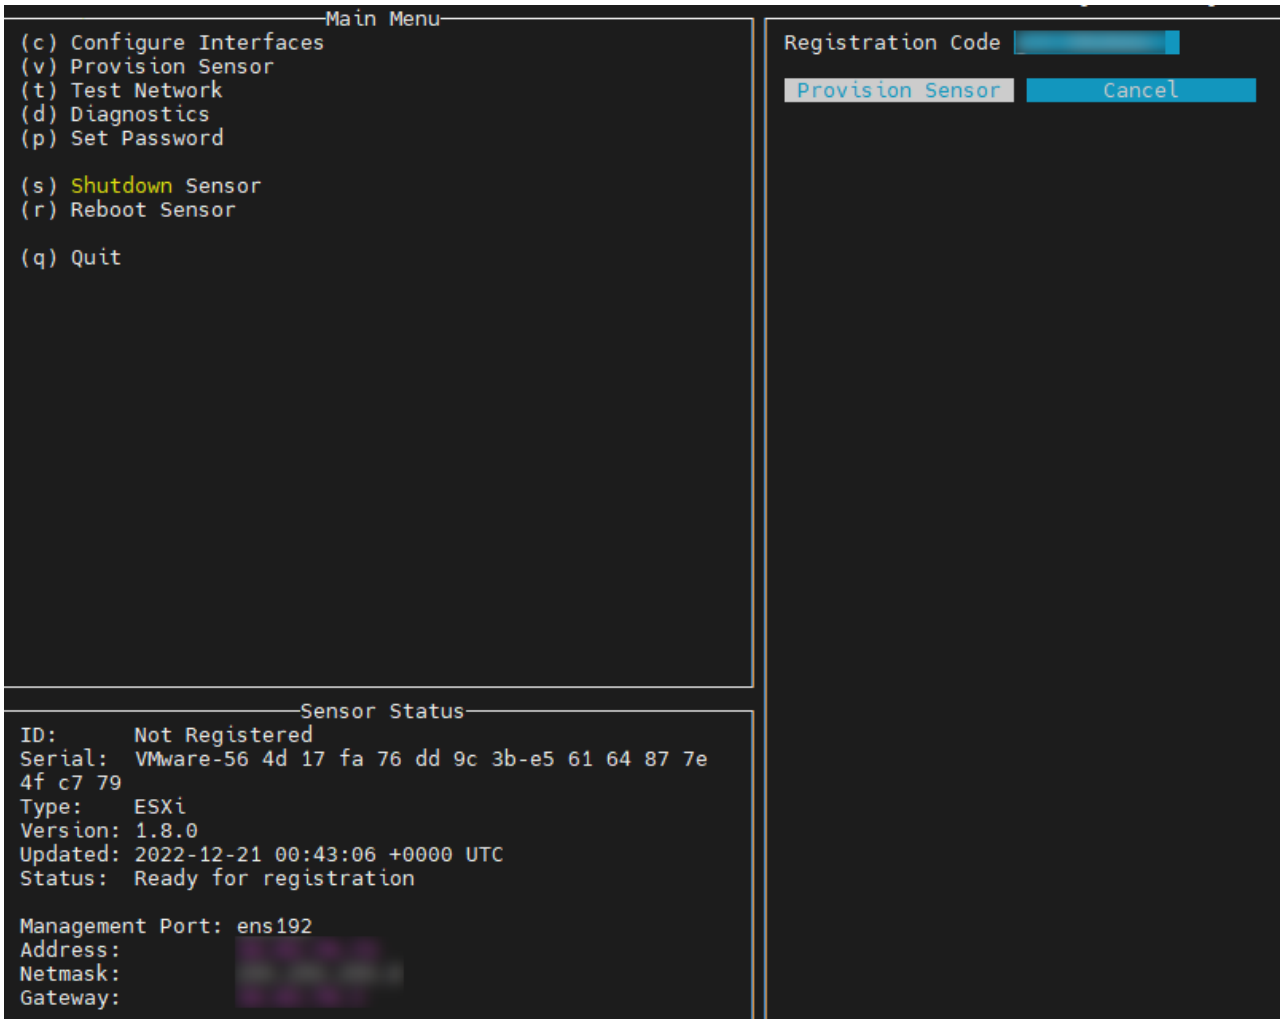

4. Select *Provision Sensor* to begin the registration process. The `Status` changes to `Sensor is provisioning`.
5. Wait for the `Status` to change to `Online`.

Troubleshooting

To troubleshoot connectivity issues:

1. Go to *Settings* > *Sensors*.
2. Click *Visible Devices*.
3. Next to *View*, click *Over Time*.

FortiAI

FortiAI in FortiNDR Cloud is Fortinet's generative AI assistant designed to accelerate threat investigation and improve visibility into network activity. Integrated into the FortiNDR portal, FortiAI enables administrators to interact in natural language and receive precise, context-aware insights about detections, understand which threats and behaviors are being monitored, and generate detailed entity reports with historical context through intuitive, conversational queries.

Licensing

FortiAI requires a valid FortiAI license.

FortiAI uses a token-based entitlement system, where each query consumes tokens based on complexity and response length. When the token limit is reached, access to FortiAI is temporarily suspended until the usage period resets. The number of available tokens is displayed at the bottom of the FortiAI chat box.

Tokens

Customers can purchase a one-year subscription that provides a fixed number of tokens each month (for example, 1,000,000 or 10,000,000 tokens, depending on the license). Unused tokens expire at the end of the month and do not roll over. Each month starts with the same allocation. Tokens can be used across Fortinet products.

Starter tokens are available to allow customers to use FortiAI features without purchasing a license. These trial-based tokens provide access for evaluation before committing to a paid plan. Starter tokens are a one-time allocation and apply only to FortiNDR. They do not renew or roll over.

For more information or to request a trial of FortiAI, contact your TSM or account manager.

Enabling FortiAI

Fortinet requires customers to opt in to use FortAI features. Once enabled, administrators must assign the *FortiAI* role users who need access to FortAI. If the organization has child accounts, FortAI must be enabled individually for each child account.

FortiAI is an account-level setting that is enabled on the *Account Management* settings page. Users cannot access FortiAI when it is disabled, even if the FortiAI role is assigned to them.

To enable FortiAI:

1. Click the gear icon and select *Account Management > Settings*.
2. Scroll down to *FortIAI Generative AI Assistant* and click *Enable*. The token balance is displayed.
3. Assign the FortiAI role to users:
   a. Go to Account *Management > Users*.
   b. Select a user and click *Assign Role*.
   c. From the *Assign Role* dropdown, select *FortiAI*.
   d. Click *Assign*.

Using FortiAI

FortiAI can be accessed by clicking the FortiAI icon at the right-side of the portal. When the Entity Panel is open, click the icon at the bottom-right corner of the page. You can expand the width of the FortiAI panel to fit the page.

Enter your prompt into the text field or select an example prompt using the horizontal scroll bar. To start a new request, click the *Start a new session* icon. FortiAI retains up to 10 previous queries, which you can access through the horizontal scroll bar.

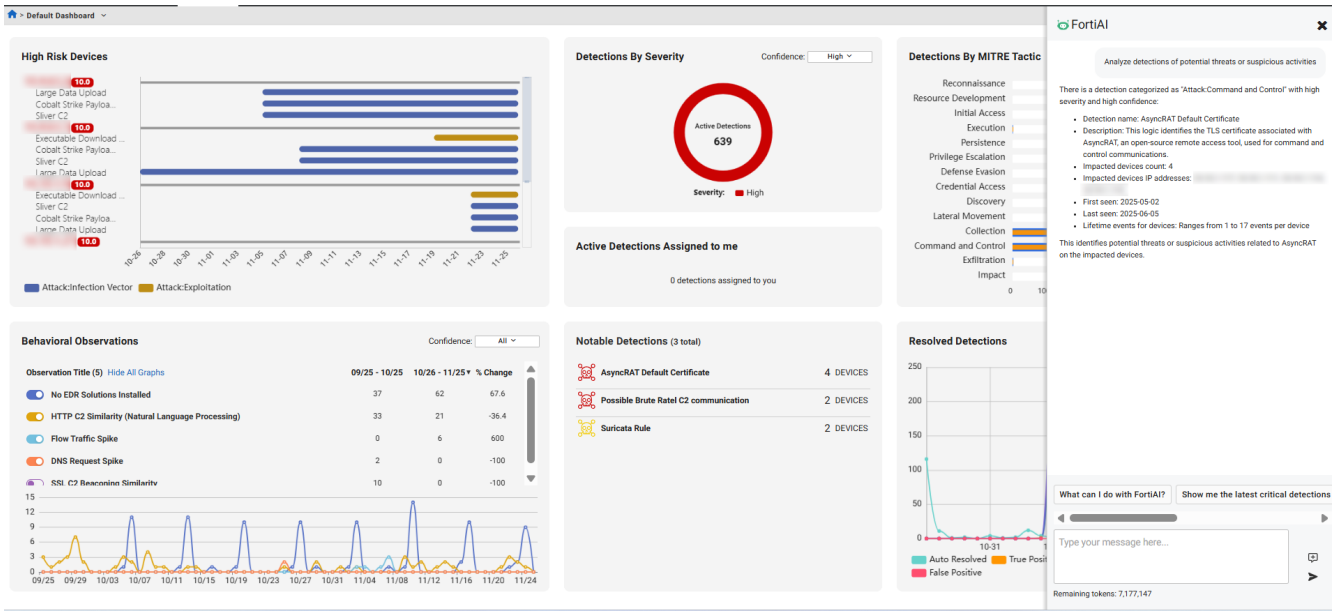

FortiAI Data privacy

FortiAI integrates with FortiNDR Cloud using a secure proxy-based architecture that protects sensitive information when communicating with the large language model (LLM). Before any data leaves FortiNDR Cloud,

Ipv4 and ipv6 addresses in user prompts are automatically masked. This ensures that IP address data is never exposed to the external FortiAI inference engine.

Currently, FortiNDR Cloud supports automatic masking for Ipv4 and ipv6 addresses only. The LLM receives only the masked prompt and does not retain or process original values directly. Once a response is returned to FortiNDR Cloud, the system re-associates the result with the original, unmasked values locally.

FortiNDR Cloud Integrations

FortiNDR Cloud natively supports integrations with multiple security tools and intelligence feeds. It also provides an open framework for creating custom integrations.

The following integrations are currently supported:

Category	Integration	Supported Version/Notes
Deception	FortiDeceptor	Requires Automation Service
SIEM	CrowdStrike	Tested with Parser 1.0.2
	FortiSIEM	7.1.0 or higher
	Microsoft Sentinel	Integration supported via API-based ingestion.
	QRadar	IBM QRadar SIEM version 7.3.3 or higher
	Splunk	Splunk Cloud versions: 9.3, 9.2, 9.1
SOAR	Cortex-XSOAR	Tested on: 6.6
	FortiSOAR	Tested on: 7.3.2-2150
	Splunk SOAR	7.3.2-2150 or higher
EDR / Firewall	FortiEDR	Manager 6.2.0 or higher Collector 5.2.0 or higher
	FortiClientEMS	Requires Automation Service
	FortiManager	7.4.2 or higher
	FortiGate	7.4.2 or higher
	CrowdStrike EDR	Requires latest Falcon EDR APIs
	SentinelOne	Requires Automation Service
Intelligence Feeds	CrowdStrike Falcon Intel	License required
	Fortinet Botnet IP List	Included with FortiNDR Cloud
	Internet Scan Data B (Shodan)	Included with FortiNDR Cloud
	Known Sinkholes	Included with FortiNDR Cloud
	PhishTank	Included with FortiNDR Cloud
	Proofpoint TAP	License required
	Recorded Future connect	License required
	Threat Connect	License required
	Tor Nodes	Included with FortiNDR Cloud

Category	Integration	Supported Version/Notes
	URLHaus	Included with FortiNDR Cloud
Other	Endace	7.2.2 or higher
	ERSPAN	Type II and Type III
	Netskope	Integration via Cloud TAP Stitcher.
	Netflow	NetFlow v5, v9, IPFIX and UDP/6343 (SFlow)
	Zscaler	Integration supported through NSS for traffic and threat logs.

For additional integrations, the SIEM/SOAR integration guide contains details for integrating with other tools. See, SIEM and SOAR Integration Guide.

For network data ingestion, FortiNDR Cloud supports hardware sensors as well as virtual sensors on various platforms, including AWS and ESXi.

- AWS Sensor Installation Guide
- ESXi Sensor Installation Guide
- Azure Sensor Installation Guide

FortiNDR Cloud also supports ingesting NSS log data from Zscaler. See, Zscaler ingestion on page 177.

Automated integration response

Automated integration response modules are available for FortiEDR and CrowdStrike Falcon EDR. Only a single integration can be set to *Auto-Remediate* at a time; others may be configured, but must be set up to respond manually.

Fortinet Automation Service

The *Fortinet Automation Service* integration streamlines and automates security operations within FortiNDR Cloud. This service enables security teams to execute predefined playbooks that perform specific actions based on connector configurations and conditional logic. A playbook can range from a simple API call to a complex, multi-step process involving several queries. Users can trigger playbooks without needing to understand the underlying logic, allowing them to focus on the intended outcome rather than the implementation details. This service enhances operational efficiency by simplifying tasks such as isolating devices, retrieving deployment network details, and executing other automated actions.

FortiNDR Essentials Solution Pack

The *FortiNDR Essentials Solution Pack* provides a set of automation playbooks that simplify incident response and security operations across multiple connectors. It connects with Fortinet products to perform common

tasks automatically. These tasks include retrieving endpoint and agent details, isolating or restoring network connectivity for compromised systems, and managing IP blocks on firewalls.

The latest version of the service pack is downloaded to your account when the automation service is provisioned. Service pack updates must be applied manually. To view the contents of the latest solution pack, see Solution pack versions.

The Fortinet Automation Service is based on the FortSOAR platform. Fortinet will regularly release solution packs that include updated connectors and new playbooks.

To request any new connectors and actions, please contact your designated TSM or log a support ticket.

Getting started with the Fortinet Automation Service

Follow these steps to begin using the Fortinet Automation Service:

Task	Description
Provision the service	Contact your TSM to provision the service.
Install the Local Agent (if needed)	Only one agent is needed for all integrations within the same network.
Configure the connectors	Set up the connectors you intend to use.
Run the playbooks	Once the connectors are configured, related playbooks will appear in the *Entity Panel*.

Provisioning the service

When the Fortinet Automation Service is provisioned, the *FortiNDR Essentials Solution Pack* is installed automatically. This service pack contains both connectors and playbooks.

Log into the FortiNDR Cloud portal and go to *Account Management > Modules*. The *Fortinet Automation Service* module will appear near the top of the page.

If you have purchased the module but do not see it listed, please contact your account team or TSM.

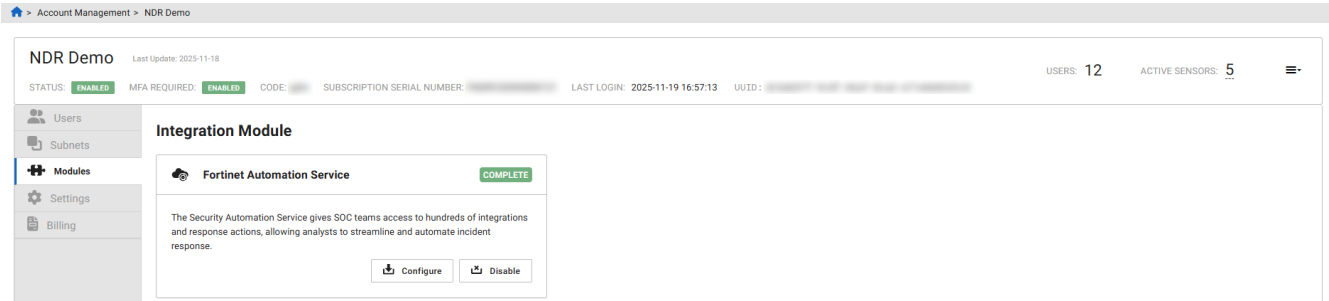

Installing the agent

The Fortinet Automation Service agent is a lightweight software component deployed within your environment. Its primary role is to facilitate secure communication between FortiNDR Cloud and the target systems or infrastructure. Agents execute playbook actions locally, such as running scripts, collecting data, or interacting with third-party tools, based on instructions received from the automation service. This allows for real-time automation while maintaining control and visibility within your network.

Only one agent is needed for all integrations within the same network.

- If you plan to use on-premise integrations, install the local agent.
- Cloud-only integrations do not require a local agent; they use a Cloud agent.

Connector	Agent Required
FortiClientEMS	Yes
FortiDeceptor	Yes
FortiEDR	No
FortiGate	Yes
FortiProxy	Yes
SentinelOne	No

Install only one agent at a time. The automation service does not support installing both the cloud and on-premises agents simultaneously.

Recommended resource requirements

- 1 GB RAM
- 1 vCPU
- 16 GB of available disk space
- Rocky Linux 9.3/9.4/9.5 or Red Hat Enterprise Linux (RHEL) Server 9.3/9.4/9.5.

Agent requirements

- Ensure that `repo.fortisoar.fortinet.com` is reachable or resolvable from the VM where you plan to install the agent.
- Ensure that the device where you plan to install the agent has outbound access to FortiNDR Cloud on ports 443 and 5671.
- Ensure connectivity to the RabbitMQ server.

To install the automation service agent:

1. Click the gear icon at the top-right of the portal and select *Account Management*.
2. Click *Modules*. The *Modules* page opens.
3. In the *Fortinet Automation Service* module, click *Configure*.
4. Click the *Agents* tab.
5. On the *Agents* page, click *Create New agent*.
6. Click *Download Installer*.

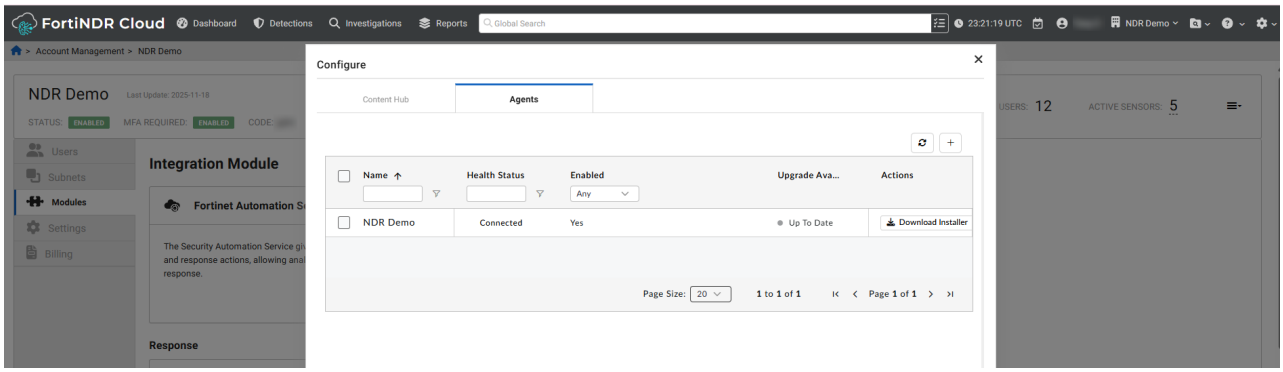

7. Choose the connectors you want to include while installing the Agent. You can choose from the following options:
   - *Do not install connectors by default*
   - *Custom*
   - *All connectors installed on the current node*
   - *Include pre-existing connectors on agent*
8. Set the *Installer type* to *Bash Script*.
9. Copy the downloaded installer script on the Agent device.
10. Run the installer script to install the Agent.

Troubleshooting agent installation

Incorrect installed connector list displayed after reconfiguring the Agent on a new VM

When reconfiguring an existing Agent on a new device, the connector list from the previous agent may incorrectly be displayed on the new Agent. This occurs when the *Do not install connector by default* option is selected during reconfiguration.

Resolution

To resolve this, select the *Include pre-existing connectors on Agent* option when reconfiguring the agent on the new VM.

Installing and configuring connectors

A connector allows the FortiNDR Cloud to interact with external systems, applications, or endpoints. It executes specific actions such as data collection, enrichment, or remediation as part of automated workflows that are triggered by playbooks and depend on network connectivity to the target systems.

To install and configure a connector:

1. Click the gear icon at the top-right of the portal and select *Account Management*.
2. Click *Modules*.
3. In the *Fortinet Security Automation Service* module, click *Configure*.
4. In the *Content Hub* tab, click the connector that you want to install.
5. In the *Connector* pop-up, click *Install*.
6. In the Confirmation dialog, click *Yes, Confirm*. If successful, a confirmation message appears and the *Configuration(s)* tab opens.
7. Configure the required fields for the connector and click *Save*. When configuring a connector, make sure to set the configuration you want to use as *Mark as default configuration*. You can create multiple configurations, but only the default configuration will be used to run playbooks.

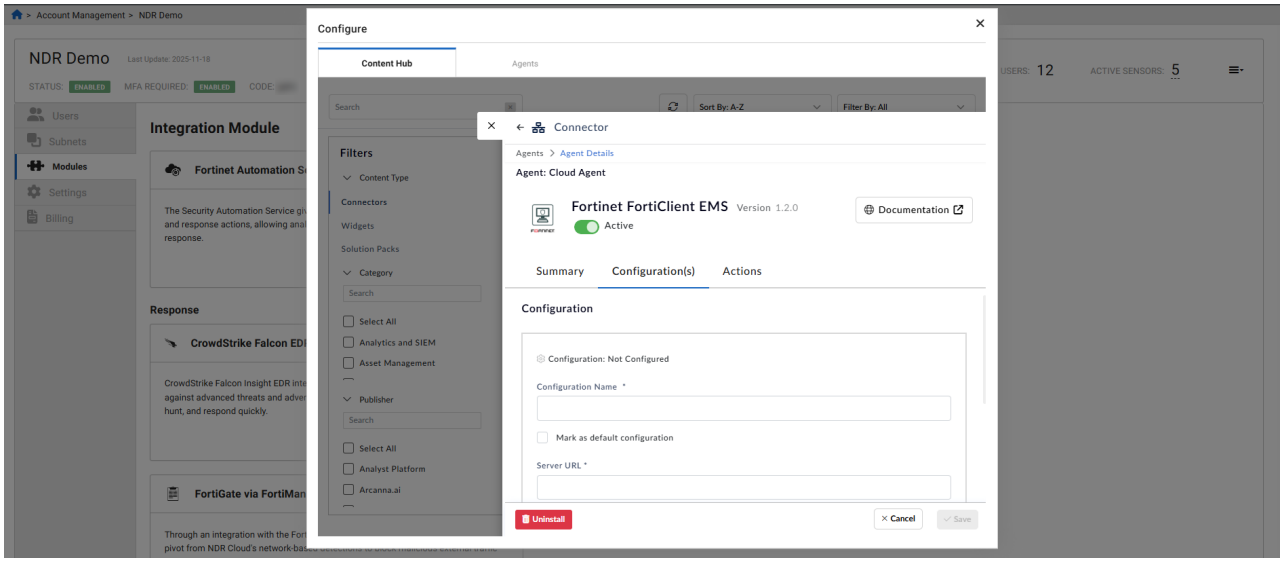

For detailed information to configure the connector, click the *Documentation* button.

Running playbooks

Playbooks are executed from the *Entity Panel*. When the Fortinet Automation Service is enabled, a link and a corresponding tab will appear in the Entity Panel, allowing you to access and execute playbooks.

The following playbooks are available:

Connector	Playbook	Description
FortiClientEMS	*Get Endpoint Details via FortiClient*	Show information the FortiClientEMS has on the endpoint, including user information, security posture and configuration.
	Quarantine Endpoint via FortiClient	Block all network traffic to or from the endpoint via FortiClientEMS.
	Unquarantine Endpoints via FortiClient	Restore network connectivity to and from the endpoint via FortiClientEMS.
FortiDeceptor	*Show All FortiDeceptor Decoys*	Get details on all decoys from FortiDeceptor.
FortiEDR	*Get Collector Details from FortiEDR*	Get Collector details from FortiEDR including user details, discovered assets and vulnerabilities.
	Unisolate Collector via FortiEDR	Restore normal network connectivity for the endpoint using FortiEDR.
	Isolate Collector via FortiEDR	Restrict the endpoint from accessing the internet via FortiEDR.
FortiGate	*Unblock IP address on FortiGate*	Unblocks IP address on Fortinet FortiGate and removes IP from the banned IP list.
	Block IP address on FortiGate	Blocks IP address on Fortinet FortiGate by Quarantine based and adds IP into the banned IP list.
FortiProxy	*Ban User by IP*	Bans IP address on Fortinet FortiProxy and adds IP into the banned users by IP list.
	Unban User by IP	Unbans IP address on Fortinet FortiProxy and removes IP from the banned users IP list.
SentinelOne	*Reconnect Agent via SentinelOne*	Restore normal network connectivity for the agent using SentinelOne (reconnect agent).
	Disconnect Agent via SentinelOne	Restrict the agent from accessing the internet via SentinelOne (isolate agent).
	Get Agent Details from SentinelOne	Show the information from SentinelOne including agent information, security posture, and configuration.

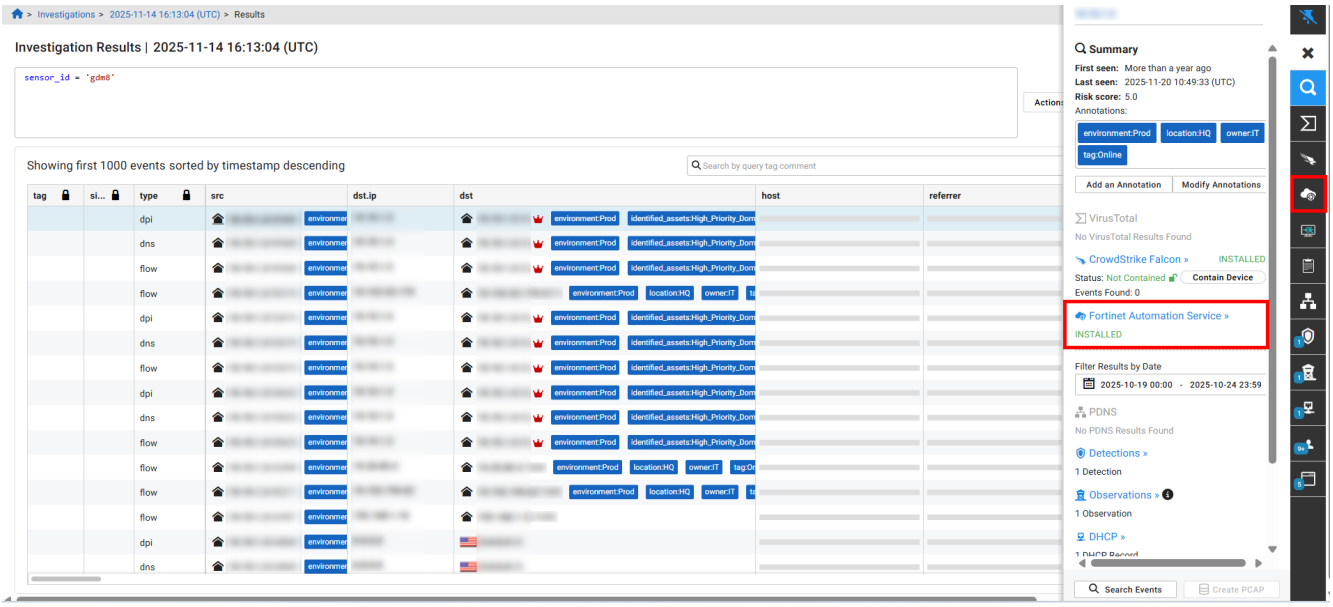

To run a playbook:

1. Open the *Entity Panel* by doing one of the following:
   - Click any entity (such as an IP address) anywhere in the portal.
   - Click an IP address in the detector details tabs.
   - Click *View Device Details* in the *Actions* menu.
   - Click a device IP in the *High Risk Devices* dashboard widget.
   - Click the IP label on the *Detections Device Timeline*.
2. In the *Entity Panel*, click the *Fortinet Automation Service* link or tab. The *Playbook List* opens.
   - For information about the playbook, hover over the information icon (`i`).
   - Click the *View* icons to view the playbooks as a list or categories.
   - Enter a keyword in the *Search* field to find a playbook by name.
   - Click the filter icon to filter based on a tag.

Open the *Entity Panel* by doing one of the following:

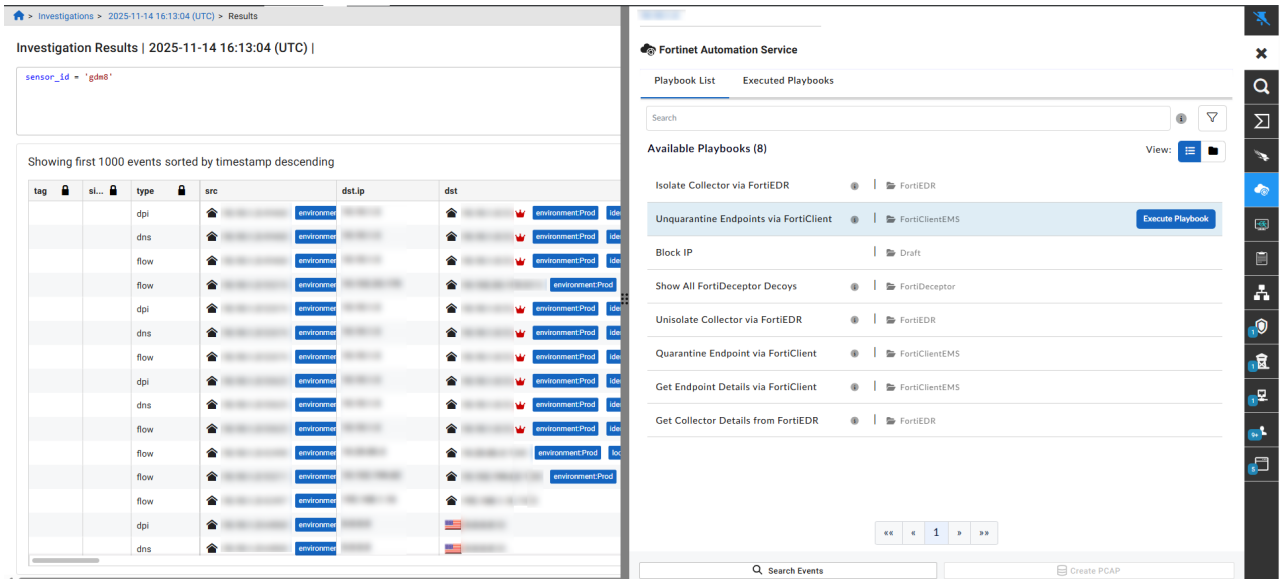

3. Hover over the playbook and click *Execute Playbook*.

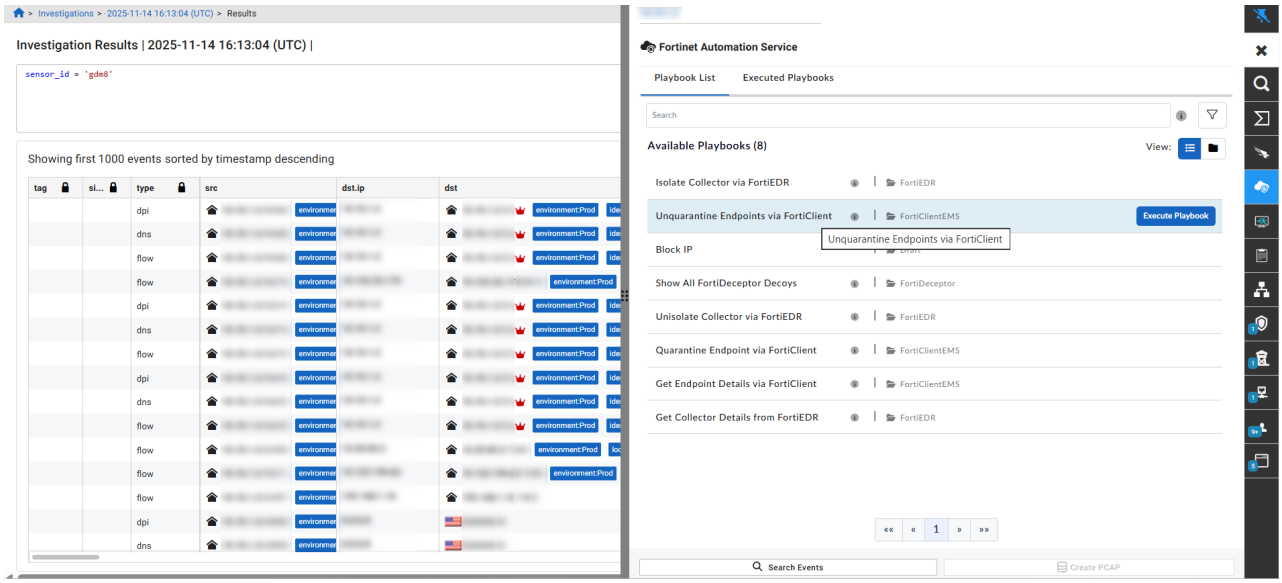

After the playbook is executed the results are displayed.

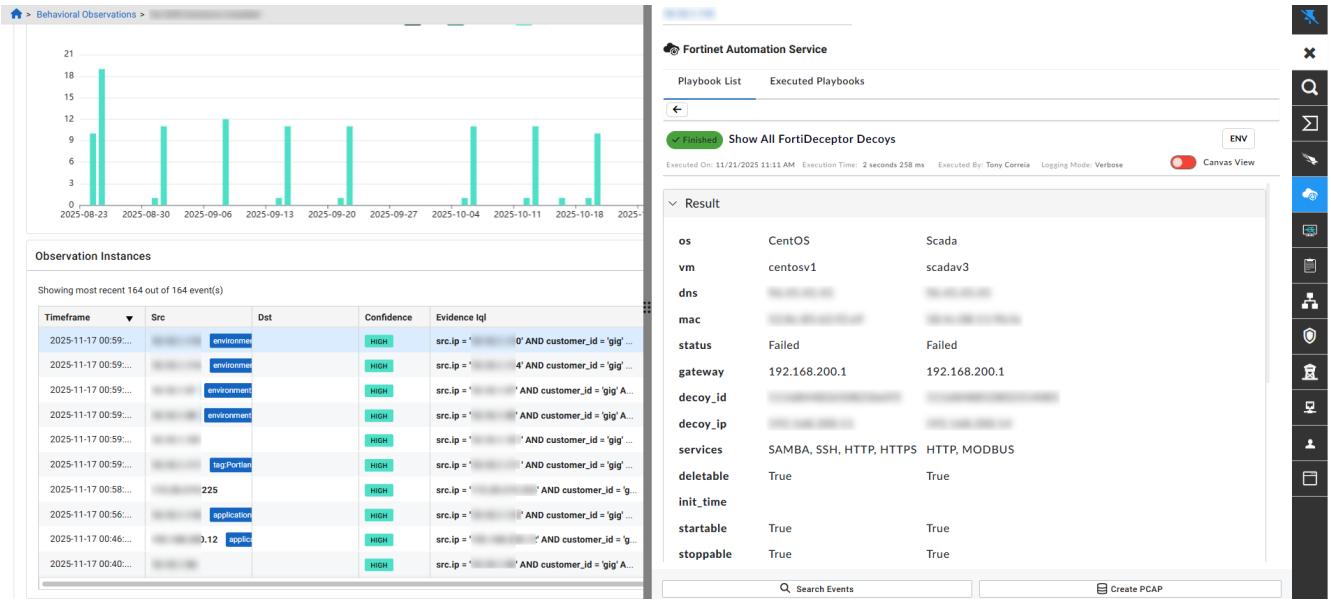

Enable *Canvas View* to visualize playbook actions as a topology.

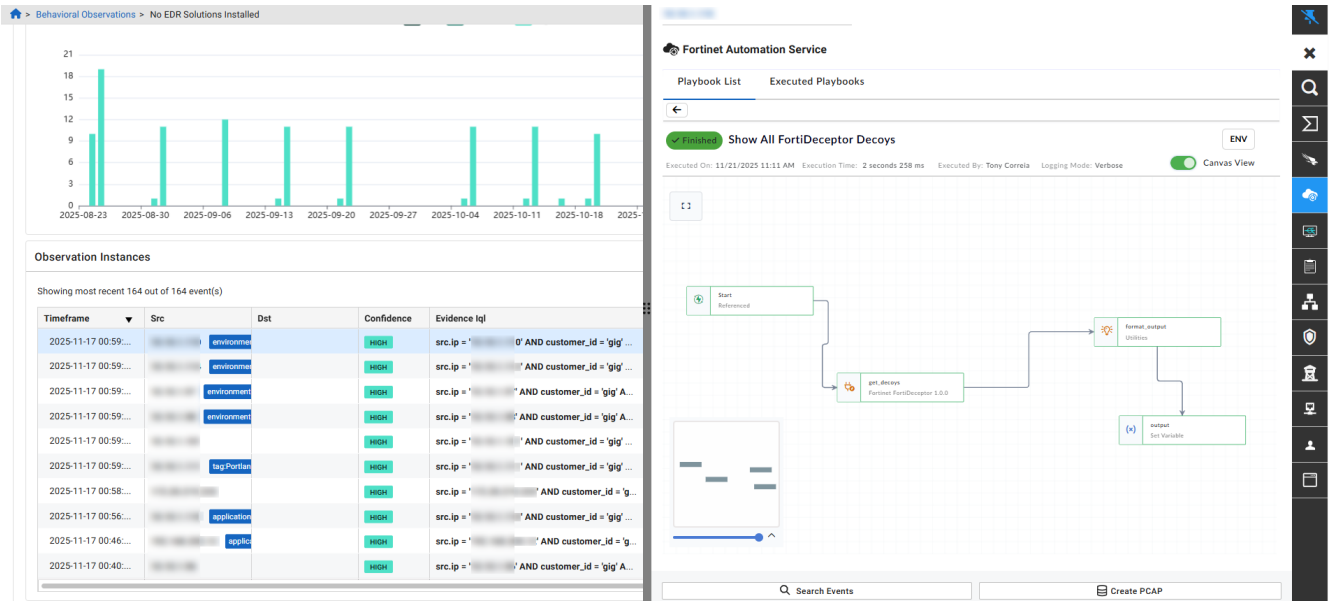

Solution pack versions

Solution Pack Version	Connectors and Playbooks
1.0.0	FortiClientEMS, FortiEDR, FortiDeceptor
1.0.1	FortiClientEMS, FortiEDR, FortiDeceptor, FortiGate, Sentinel One
1.0.2	FortiClientEMS, FortiEDR, FortiDeceptor, FortiGate, FortiProxy, Sentinel One

FortiNDR Cloud APIs

FortiNDR Cloud API documentation is available on the Fortinet Developer Network (FNDN).

Available APIs

- **Entity API**: Obtain details on individual entities such as IPs, domains, file hashes. This API supports providing details on an entity such as DHCP and DNS information and when it was first and last seen. For information about Entities, see Entity Panel on page 52.
- **Detections AP**I: Provides details on malicious events that were detected. SeeDetections on page 34
- **Sensor API**: Provides APIs for interacting with sensors.
- **Investigations API**: APIs for managing investigations and running queries.

Metastream

FortiNDR Cloud also provides access to the most recent seven days of events on Metastream. A python client is available to facilitate interacting with the most used events.

- Metastream documentation is available on the Fortinet Developer Network (FNDN).
- Client library documentation is available in the Document library. See, FNC Python Client Library.

IQL reference guide

Internal Query Language (IQL) is used in FortiNDR Cloud for identifying, querying, filtering, and analyzing various network events such as *flow*, *HTTP*, *and* SSL events. It supports detections, behavioral observations, guided queries, and investigations. The results of an IQL query include enriched events, which are enhanced with intelligence indicator matches from FortiNDR Cloud's threat intelligence database. Additionally, IP enrichments such as ASN, internal/external status, and geographical attributes are included to provide comprehensive insights into network activities.

Purpose of this reference guide

This reference guide is intended as an introduction to creating IQL queries in FortiNDR Cloud. Where possible, we have provided example queries and short exercises to help you get started.

Using guided queries

If this is your first time creating queries, we recommend running a few Guided Queries to start. These will help familiarize you with query strings and their results. You can also use the results to add new queries to experiment with. For more information, see Guided queries on page 112.

Sample queries

The portal also offers a library of sample queries for common searches. To access these samples, log into the portal and navigate to *Investigations* > *Private Search*.

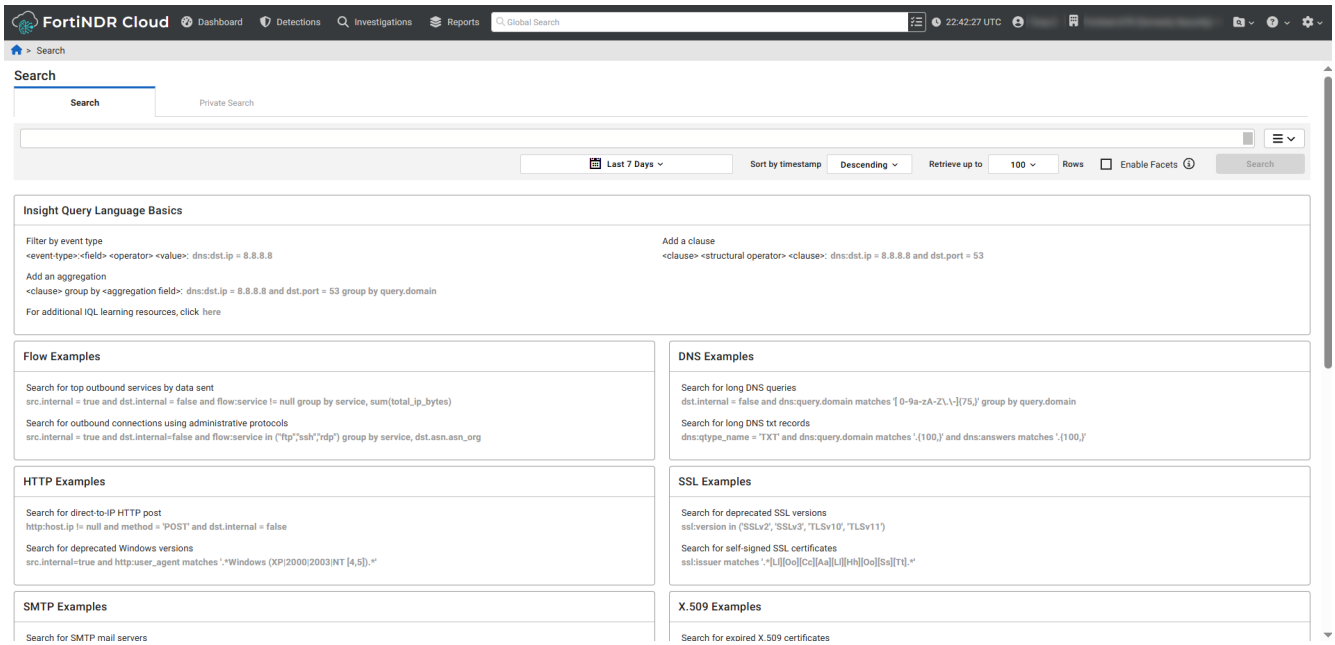

Core IQL concepts

IQL Clause

IQL clauses follow the format `<field> <operator> <value>` and can be combined using logical operators like `AND` and `OR`. Parentheses can be used to control the order of these logical operators in a query.

Example:

```
ip = 8.8.8.8 AND host LIKE "%.google.com".
```

<field>	`ip = 8.8.8.8`
<operator>	`AND host LIKE`
<value>	`"%.google.com"`

Exercise:

1. Go to *Investigations* > *Guided Queries*.
2. Run the *Example Hunt* query. For more information, see Guided queries on page 112

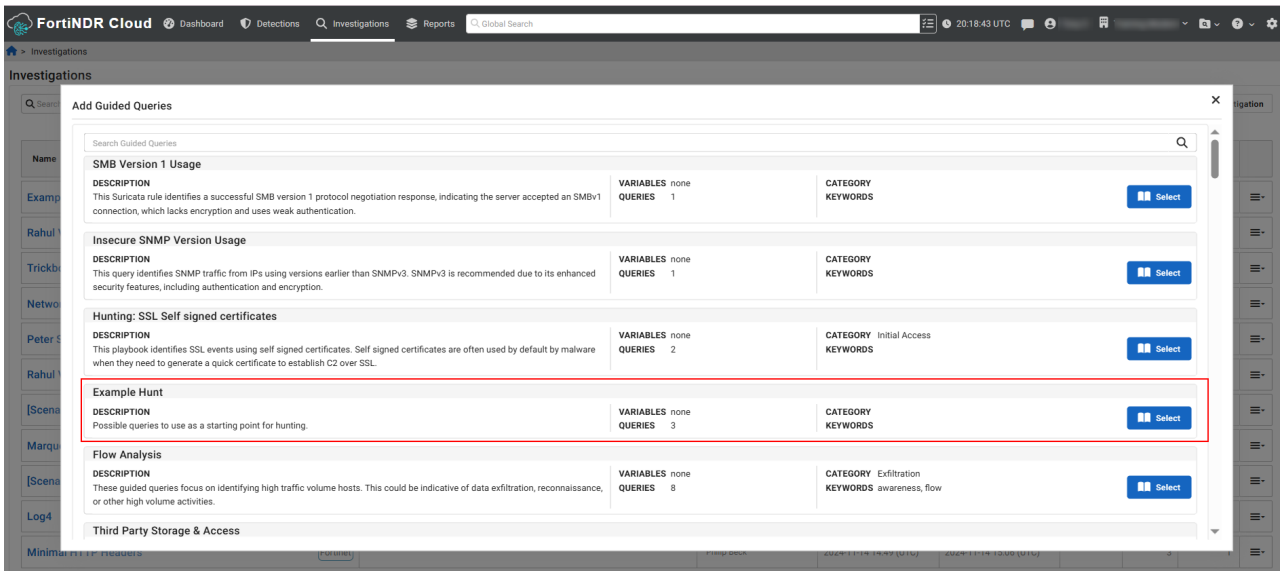

3. After the query is completed, go to *Investigations* and click the query name in the list, then click *View Results*.
4. In the *Investigations Results* page, click the *Events* tab.
5. In the `src`column, click an IP address to open the *Entity Panel* and then copy the IP at the very top of the pane.
6. Use the IP and the fields in the *Events* tab to create a new query.
   Example: `ip = 10.10.31.101 AND dst.geo.country LIKE "FR"`
7. To add this query to your investigation, click the investigation name in the bed crumb at the top-left of the page , and click *Add Query* at the bottom of the investigation details page.

Fields

Fields are used to specify and limit event types for querying and analyzing network events.

Event Type

An event type specifies the category of network events you want to query or analyze. The *event_type* field applies to all events, allowing you to filter and focus on particular types of network activities.

By using `<event_type>:<field>`, you can focus your query to a specific type of event. This helps make your search more precise and relevant to the data you are interested in.

- Flow
- HTTP
- DNS

- SSL

Example event types:

Flow

A `flow` event refers to a record of network traffic between two endpoints. It typically includes information such as the source and destination IP addresses, source and destination ports, protocol used (e.g., TCP, UDP), the amount of data transferred, and the duration of the connection.

This information helps you monitor and analyze traffic patterns, detect anomalies, and identify potential security threats.

Exercise:

1. Using the results to from the investigation, click *Add Query*and name it *Flow Events*.
2. In the *Query* field, type `event_type = 'flow'`.
3. Click *Add Query* to run the query and then view the results.
4. (Optional) Click the Individual columns dropdown to show and hide the columns to view the data.

HTTP

An HTTP event type refers to a record of HTTP traffic between a client and a server. It typically includes details about the HTTP request and response, such as the method used, the URL accessed, headers, and status codes.

This information helps you monitor web traffic, detect malicious activities such as web attacks, and ensure compliance with security policies.

Exercise:

1. Using the results to from the investigation you created earlier, click *Add Query* and name it *HTTP Events*.
2. In the *Query* field, type `event_type = 'http'`. Use lowercase for `http`.
3. Click *Add Query* to run the query and then view the results.

DNS

A DNS event type refers to a record of DNS (Domain Name System) queries and responses between a client and a DNS server. It typically includes details about the DNS request and the corresponding response.

This information helps you monitor DNS traffic, detect anomalies such as DNS spoofing or tunneling, and ensure the integrity and security of domain name resolutions within the network.

Exercise:

1. In *Investigation Results* page for the HTTP query, click an IP address in `src` column to open the *Entity Panel*.
2. At the bottom of the *Entity Panel*, click *Search Events*. The *Add Query to Investigation* dialog opens.
3. In the *Query Name* field, type *DNS*.
4. In the *Search Query* field, type `event_type = 'dns'`.
5. Click *Add Query* to run the query and then view the results.

SSL

An SSL event type refers to a record of SSL/TLS (Secure Sockets Layer/Transport Layer Security) traffic between a client and a server. It typically includes details about the SSL handshake, certificates, and encrypted data transfer.

This information helps you monitor encrypted traffic, ensure the security of SSL/TLS connections, and detect potential issues such as expired certificates, weak ciphers, or SSL/TLS vulnerabilities.

Exercise:

1. Using the results to from the investigation, click *Add Query* and name it *SSL Events*.
2. In the *Query* field, type `event_type = 'ssl'`.
3. Click *Add Query* to run the query and then view the results.

Sub-fields

A sub-field is a more specific field within a broader parent field. When you search for a sub field without specifying the parent field, the search will include all subfields with that name.

Examples:

Parent field	Sub-field search
IP	`src.ip`, `dst.ip`, `host.ip`, `answers.ip`, `referrer.host.ip`, `headers.location.ip`, etc.
Domain	`host.domain`, `query.domain`, `helo.domain`, `san_dns.domain`, etc.
URI	`uri.uri`, `referrer.uri`, etc.
Query	`uri.query`, `referrer.query` (but not `dns:query`; use `query.domain` instead).

Exercise

This is exercise is based on the *Example Hunt* investigation you ran earlier.

1. Click *View Results* next the first query in the list.
2. Click the *Events* tab. The columns to the right of the type column represent the sub fields for the parent event.

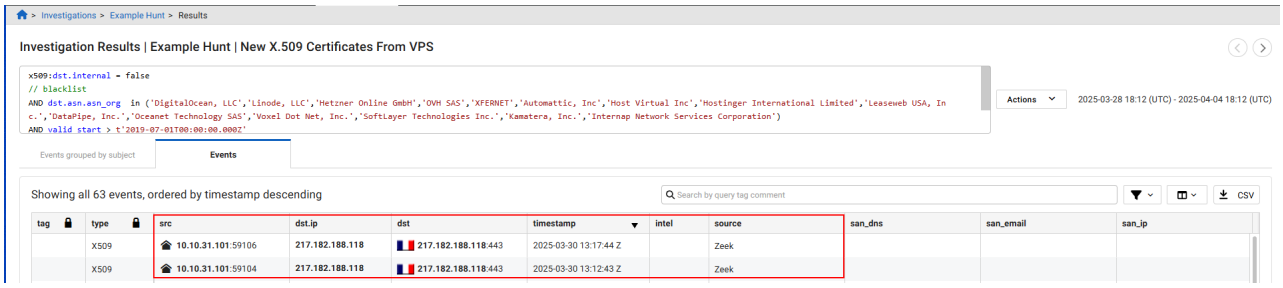

3. Record a column header and its value. For the purpose of this exercise, we will use `dst.ip`.
4. Go back to your investigation and click *Add Query*.
5. In the query field, create a new query based on the event type and sub field. If you need help with an operator, see Operators.
   Example: `event_type = 'flow'AND dst.ip = "10.10.1.5"`

Commonly Confused Fields

Field	Example
URI	`uri.uri` vs. *uri.path* and *uri.query*
MIME	`request_mime` vs. `request_mimes`
File	`file.*` vs. `files.*`

Some fields cannot be searched, such as `account` and `observation:context`.

Value Types

A value type refers to the specific data or value that you are querying or filtering for within a field. It is the actual content you are looking for in your search. For example, in the clause `<field> <operator> <value>`, the content you are looking for is `<value>`.

Value types are used in conjunction with fields and operators to form complete IQL clauses, allowing you to perform precise and targeted searches within your data.

Integer	A number such as , 9, 54458, -8 (`snmp:snmp_version != '3'`)
Float	A number with decimal points, such as 4.5, 125.5554
Boolean	True, false, or null (`dns:src.internal = true and dns:dst.internal = false`)
String	Alphanumeric characters contained in single or double-quotes (`kerberos:error_msg = 'KDC_ERR_CLIENT_REVOKED'`).
Timestamp	In the format `t”2023-02-28T00:00:00.000Z”` contained in single or double-quotes, 'millisecond- or microsecond-precision '(`valid_start > t'2019-07-01T00:00:00.000Z`)
IP	Single IP or CIDR, quoted or unquoted (`ip =8.8.8.8`)
Object	Anything with a sub-field, such as: `IP-objects`, `Domain-objects`, `Host-objects`, `URI-Objects`, `File-Objects`, `Email-Objects`
Array	IQL clause is satisfied if any value in the array satisfies the clause. (`suricata:sig_id IN (10098240,10099368)`)

Object Types

An object type is the category or class of data that you are looking for. It helps you define what kind of information you want to find and makes your search more specific and accurate.

By specifying an object type, you can focus your search on particular kinds of data. This makes your queries more precise and helps you find exactly what you need.

The table below lists the available object types along with their descriptions and examples. Click on an object type in the *Object Type* column to view a sample query.

Object Type	Description	Example
IP	Information related to internet protocol addresses.	ASN (Autonomous System Number), geo (geographical location), internal, port. *Flow Events*: ip_bytes, pkts (packets).
ASN	Details about the Autonomous System Number.	ASN, asn_org (organization), ISP (Internet Service Provider), org (organization).
Geo	Geographical information.	City, country, location, subdivision.
Domain	Information about domain names.	City, country, location, subdivision.
URI	Uniform Resource Identifier details.	Fragment, host, path, port, query, scheme, uri.
File	Information about files.	Bytes, MD5 (hash), MIME type, name, SHA1 (hash), SHA256 (hash).
email	Information related to email addresses.	Domain, email, name.
host	Combines IP and domain information.	

Sample object queries

The following example queries are intended to help you get started with query objects. Each example uses curly braces `{}` for multiple conditions.

IP

This query will return results that match both the specified IP address and the country within the IP object.

```
ip {
   address = "203.0.113.5"
   AND geo.country = "Canada"
}
```

ip	This specifies that you are querying the IP object.
address = "203.0.113.5"	This condition filters the query to include only IP addresses that match 203.0.113.5.

AND geo.country = "Canada"	This additional condition ensures that the query also matches IP addresses located in Canada.

ASN

This query will return results that match both the specified ASN and organization within the ASN object.

```
asn {
   asn = "12345"
   AND org = "Example Organization"
}
```

asn	This specifies that you are querying the ASN object.
asn = "12345"	This condition filters the query to include only ASNs that match 12345.
AND org = "Example Organization"	This additional condition ensures that the query also matches ASNs associated with the organization named *Example Organization*.

Geo

This query will return results that match both the specified city and country within the geo object.

```
geo {
   city = "Vancouver"
   AND country = "Canada"
}
```

geo	This specifies that you are querying the geo object.
city = "Vancouver"	This condition filters the query to include only geographical locations in the city of Vancouver.
AND country = "Canada"	This additional condition ensures that the query also matches locations within Canada.

URI

This query will return results that match both the specified domain name and country within the domain object.

```
domain {
   name = "example.com"
   AND geo.country = "Canada"
}
```

domain	This specifies that you are querying the domain object.
name = "example.com"	This condition filters the query to include only domains that match example.com.

AND geo.country = "Canada"	This additional condition ensures that the query also matches domains located in Canada.

File

This query will return results that match both the specified file name and SHA-256 hash within the file object.

```
file {
   name = "example.txt"
   AND sha256 = "d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2d2"
}
```

file	This specifies that you are querying the file object.
name = "example.txt"	This condition filters the query to include only files named example.txt.
AND sha256 = "d2d2d2d2d2d2d..."	This additional condition ensures that the query also matches files with the specified SHA-256 hash.

Domain.

This query will return results that match both the specified domain name and country within the domain object.

```
domain {
   name = "example.com"
   AND location.country = "Canada"
}
```

domain	This specifies that you are querying the domain object.
name = "example.com"	This condition filters the query to include only domains that match example.com.
AND location.country = "Canada"	This additional condition ensures that the query also matches domains located in Canada.

Email

This query will return results that match both the specified domain and name within the email object.

```
email {
   domain = "example.com"
   AND name = "John Doe"
}
```

email	This specifies that you are querying the email object.
domain = "example.com"	This condition filters the query to include only emails from the domain example.com.
AND name = "John Doe"	This additional condition ensures that the query also matches emails associated with the name John Doe.

Host

The following query will return results that match both the specified IP address and domain within the host object.

```
host {
   ip = "192.168.1.1"
   AND domain = "example.com"
}
```

host	This specifies that you are querying the host object.
ip = "192.168.1.1":	This condition filters the query to include only hosts with the IP address 192.168.1.1.
AND domain = "example.com":	This additional condition ensures that the query also matches hosts with the domain example.com.

Fields and field types

This document provides information about event types, field types and enriched object field types used in FortiNDR Cloud for network event analysis.

- Field types on page 208
- Enriched object field types on page 209
- Common fields on page 218

Field types

Most fields are atomic, meaning they cannot be broken down further. However, FortiNDR Cloud fields can also be a structured object, either an object or an array. See Enriched object field types on page 209.

Fields in FortiNDR Cloud can be one of the following types.

Field Type	Description	Example
int	An integer value (port, bytes, packets, etc.)	`1`
float	A decimal value (distance, entropy, etc.)	`1.0`
Boolean	true of false	`True`
string	A sequence of arbitrary characters	`hello world`
timestamp	A RFC3339 timestamp value	`2019-01-01T00:00:00.000Z`
ip	A single IP address or valid CIDR-notation	`8.8.8.8`, `10.0.1.0/24`
object	An arbitrary JSON structure containing nested subfields	N/A
array	An array of values of the same type	N/A

Enriched object field types

A field that is of type object simply means the field is actually a collection of sub-fields. Some of those sub-fields could also be another collection of sub-fields. Think of an *object* as a JSON block, or a dictionary for the Python users, or a map for the C/C++ users. Sub-fields are then referenced using dot notation, (for example, `dst.geo.country`).

Some object types are very common and are used over and over again, such as an `ip-object`. An *ip-object* refers to a field with the structure shown in the ip-object table. These field types are used throughout the different event types, so you should be familiar with them.

Deprecation notice:

The `asn.isp` and `asn.org` fields are no longer supported. Please use `asn.asn_org` or `asn.asn` fields instead. This change applies to all IP-related fields.

The following topics provide a description of each object field type and the sub-fields it contains:

- IP-Objects on page 209
  - Active Directory (AD) objects on page 211
- Domain-Objects on page 216
- Host-Objects on page 216
- URI-Objects on page 217
- URL-Objects on page 217
- File-Objects on page 217
- Email-Objects on page 218

Back to top.

IP-Objects

The following table describes the fields that contain enriched information for an IP address:

Field	Type	Description
`asn`	asn-object	ASN information for the IP address Example: See table below
`$device`	synthetic field	Enables querying devices by hostname or MAC address. Note: this field is only available for the `src` and `dst` fields.
`geo`	geo-object	Geographic information for the IP address Example: See table below
`internal`	Boolean	Indicates whether the IP address is internal to the network Example: `true`

Field	Type	Description
`ip`	ip	The IP address Example: `10.10.10.10`
`ip_bytes`	int	The number of bytes transmitted by the IP address within the flow (only populated in Flow events) Example: `458 Bytes`
`pkts`	int	The number of packets transmitted by the IP address within the flow (only populated in Flow events) Example: `8`
`port`	int	The port used by the IP address Example: `52843`
`username`	int	The user name from Zscaler used in device detections (only populated in DNS, Flow, HTTP, and SSL events). Example: `john.smith@fortinet.com`
`hostname`	int	The host name from Zscaler used in device detections (only populated in DNS, Flow, HTTP, and SSL events). Example: `F09NQJM1ABC`

The `asn` field contains the following subfields.

Field	Type	Description
`asn`	int	The Autonomous System Number Example: `16509`
`asn.asn_org`	string	The organization name associated with the ASN (they actually use the ASN) Example: `Amazon.com, Inc.`
`asn.asn`	string	The upstream ISP for the ASN Example: `Amazon.com`
`org`	string	The upstream owner of the ASN - may differ from `asn_org` Example: `Amazon.com`

The `geo` field contains the following subfields.

Field	Type	Description
`city`	string	The city of record Example: `Boardman`
`country`	string	The country of record Example: `US`

Field	Type	Description
`location`	object	The longitude and latitude of record Example: `(45.8491,-119.7143)`
`subdivision`	string	The segment of the country (states in the US) Example: `OR`

Back to Enriched object field types.

Active Directory (AD) objects

Active Directory enrichment enhances device identification by collecting hostname information from Windows AD on a scheduled basis. See Device enrichment on page 157.

When Active Directory enrichments are enabled, IP addresses are enriched with the following fields:

The following table describes the fields that contain enriched information for *IP_enrichments*:

Property	Type	Description
`annotations`	annotations object	User- and system-generated metadata associated with an entity or event. Annotations are typically added during investigations to capture analyst notes, contextual observations, or links to related findings, and are used to enrich the interpretation of security data.
`asn`	asn object	ASN information for the IP address
`device_data_timestamp`	string	Time which device enrichment data was determined Example: 2025-01-16T23:26:28.000000Z
`device_hostnames`	array	A collection of hostnames associated with a device, derived from observed network activity or integrated identity sources.
`device_last_logoff`	string	Timestamp device was last recorded logging out Example: 2025-01-16T22:15:08.000000Z
`device_last_logon`	string	Timestamp device was last recorded logging in Example: 2025-01-16T22:13:08.000000Z
`device_os_name`	string	Device operating system name Example: Windows
`device_os_name_with_version`	string	Device operating system with version Example: Windows 10 Pro
`device_os_version`	string	Device operating system version Example: 10.0 (54983)
`device_os_version_major`	string	Device operating system major version number Example: 10

Property	Type	Description
`device_os_version_minor`	string	Device operating system minor version number Example: 0
`geo`	geo object	Geographic information associated with an IP address.
`internal`	boolean	Indicates whether the IP address is internal to the network Example: true

The following table describes the fields that contain enriched information for `ip_enriched_with_port 1`:

Property	Type	Description
`annotations`	annotations object	User- and system-generated metadata associated with an entity or event. Annotations are typically added during investigations to capture analyst notes, contextual observations, or links to related findings, and are used to enrich the interpretation of security data.
`asn`	asn object	ASN information for the IP address
`device_data_timestamp`	string	Time which device enrichment data was determined. Example: 2025-01-16T23:26:28.000000Z
`device_hostnames`	array	A collection of hostnames associated with a device, derived from observed network activity or integrated identity sources.
`device_last_logoff`	string	Timestamp device was last recorded logging out. Example: 2025-01-16T22:15:08.000000Z
`device_last_logon`	string	Timestamp device was last recorded logging in. Example: 2025-01-16T22:13:08.000000Z
`device_os_name`	string	Device operating system name. Example: Windows
`device_os_name_with_version`	string	Device operating system with version. Example: Windows 10 Pro
`device_os_version`	string	Device operating system version. Example: 10.0 (54983)
`device_os_version_major`	string	Device operating system major version number. Example: 10
`device_os_version_minor`	string	Device operating system minor version number. Example: 0
`geo`	geo object	Geographic information associated with an IP

Property	Type	Description
		address.
`internal`	boolean	Indicates whether the IP address is internal to the network. Example: true
`ip`	string	An IP address. Example: 8.8.8.8
`port`	integer	A port number. Example: 443

ip_enriched 1

The following table describes the fields that contain enriched information for `ip_enriched 1`:

Property	Type	Description
`annotations`	annotations object	User- and system-generated metadata associated with an entity or event. Annotations are typically added during investigations to capture analyst notes, contextual observations, or links to related findings, and are used to enrich the interpretation of security data.
`asn`	asn object	ASN information for the IP address
`device_data_timestamp`	string	Time which device enrichment data was determined. Example: 2025-01-16T23:26:28.000000Z
`device_hostnames`	array	A collection of hostnames associated with a device, derived from observed network activity or integrated identity sources.
`device_last_logoff`	string	Timestamp device was last recorded logging out. Example: 2025-01-16T22:15:08.000000Z
`device_last_logon`	string	Timestamp device was last recorded logging in. Example: 2025-01-16T22:13:08.000000Z
`device_os_name`	string	Device operating system name. Example: Windows
`device_os_name_with_version`	string	Device operating system with version. Example: Windows 10 Pro
`device_os_version`	string	Device operating system version. Example: 10.0 (54983)
`device_os_version_major`	string	Device operating system major version number. Example: 10

Property	Type	Description
`device_os_version_minor`	string	Device operating system minor version number. Example: 0
`geo`	geo object	Geographic information associated with an IP address.
`internal`	boolean	Indicates whether the IP address is internal to the network. Example: true
`ip`	string	An IP address. Example: 8.8.8.8

The following table describes the fields that contain enriched information for `ip_enriched 2`:

Property	Type	Description
`annotations`	annotations object	User- and system-generated metadata associated with an entity or event. Annotations are typically added during investigations to capture analyst notes, contextual observations, or links to related findings, and are used to enrich the interpretation of security data.
`asn`	asn object	ASN information for the IP address
`device_data_timestamp`	string	Time which device enrichment data was determined. Example: 2025-01-16T23:26:28.000000Z
`device_hostnames`	array	A collection of hostnames associated with a device, derived from observed network activity or integrated identity sources.
`device_last_logoff`	string	Timestamp device was last recorded logging out. Example: 2025-01-16T22:15:08.000000Z
`device_last_logon`	string	Timestamp device was last recorded logging in. Example: 2025-01-16T22:13:08.000000Z
`device_os_name`	string	Device operating system name. Example: Windows
`device_os_name_with_version`	string	Device operating system with version. Example: Windows 10 Pro
`device_os_version`	string	Device operating system version. Example: 10.0 (54983)
`device_os_version_major`	string	Device operating system major version number. Example: 10
`device_os_version_minor`	string	Device operating system minor version number.

Property	Type	Description
		Example: 0
`geo`	geo object	Geographic information associated with an IP address.
`internal`	boolean	Indicates whether the IP address is internal to the network. Example: true
`ip`	string	An IP address. Example: 8.8.8.8

The following table describes the fields that contain enriched information for an `interface_enriched`:

Property	Type	Description
`annotations`	annotations object	User- and system-generated metadata associated with an entity or event. Annotations are typically added during investigations to capture analyst notes, contextual observations, or links to related findings, and are used to enrich the interpretation of security data.
`asn`	asn object	ASN information for the IP address
`device_data_timestamp`	string	Time which device enrichment data was determined. Example: 2025-01-16T23:26:28.000000Z
`device_hostnames`	array	A collection of hostnames associated with a device, derived from observed network activity or integrated identity sources.
`device_last_logoff`	string	Timestamp device was last recorded logging out. Example: 2025-01-16T22:15:08.000000Z
`device_last_logon`	string	Timestamp device was last recorded logging in. Example: 2025-01-16T22:13:08.000000Z
`device_os_name`	string	Device operating system name. Example: Windows
`device_os_name_with_version`	string	Device operating system with version. Example: Windows 10 Pro
`device_os_version`	string	Device operating system version. Example: 10.0 (54983)
`device_os_version_major`	string	Device operating system major version number. Example: 10
`device_os_version_minor`	string	Device operating system minor version number. Example: 0

Property	Type	Description
`geo`	geo object	Geographic information associated with an IP address.
`internal`	boolean	Indicates whether the IP address is internal to the network. Example: true
`ip`	string	An IP address. Example: 8.8.8.8
`mac`	string	A MAC address. Example: 00:1A:2B:3C:4D:5E
`port`	integer	A port number. Example: 443

The following table describes the fields that contain enriched information for an `device_hostname`:

Property	Type	Description
`domain_name`	string	Domain of the device's fully qualified domain name. Example: apps.google.com
`fqdn`	string	Device's fully qualified domain name. Example: server1.apps.google.com
`name`	string	Hostname of the device's fully qualified domain name. Example: server1
`secondary_level_domain_name`	string	Secondary level domain of the device's fully qualified domain name. Example: google.com

Domain-Objects

The following table describes the fields that contain enriched information for a domain:

Field	Type	Description
`domain`	string	The domain Example: `portal.fortindr.forticloud.com`
`domain_entropy`	float	The computed Shannon entropy of the domain Example: `3.5`

Back to Enriched object field types

Host-Objects

Host-Objects fields contain enriched information for both IP addresses and domains because the field could be either one. For example an HTTP Host header or a DNS answer.

Host-Objects contain the combined sub-fields in:

- IP-Objects on page 209
- Domain-Objects on page 216

Back to Enriched object field types

URI-Objects

Fields that contain a URI are broken up into its different components.

Field	Type	Description
`fragment`	string	The fragment identifier component Example: `#`
`host`	host-object	The content of the Host header Example: `portal.fortindr.forticloud.com`
`params`	object-array	The HTTP parameters as an array of key-value pairs **Example**:
`path`	string	The path of the requested resource Example: `search`
`port`	integer	The specified port Example: `443`
`query`	string	The full parameter string Example: `query=8.8.8.8&sort_dir=desc`
`scheme`	string	The specified scheme Example: `https`
`uri`	string	The full URI Example: `https://portal.fortindr.forticloud.com:443/search?query=8.8.8.8&sort_dir=desc#`

URL-Objects

Fields that contain both a *host-object* and a *uri-object* are referred to as a *url-object*.

URL-Objects contain the combined sub-fields in:

- IP-Objects on page 209
- Domain-Objects on page 216
- URI-Objects on page 217

Back to Enriched object field types

File-Objects

File-Objects fields contain enriched information for an observed file.

Field	Type	Description
`bytes`	int	The file's size in bytes Example: `145922`
`md5`	string	The computed MD5 hash Example: `92a4d0aeede3ce110b4121342df48496`
`mime_type`	string	The fingerprinted MIME-type Example: `application/x-dosexec`
`name`	string	The observed name Example: `2487ff63fb4e79.gif`
`sha1`	string	The computed SHA1 hash Example: e63932430d4028b51fa25dae13d9e0188e9a02a5
`sha256`	string	The computed SHA256 hash Example: `227193160a2448dfa8bbbd2cf125afa9cca0d1a718b109a3adae5df8a24cdf6e`

Back to Enriched object field types

Email-Objects

Email-Objects fields contain an email address broken up into its different components.

Field	Type	Description
`domain`	string	The domain Example: `gmail.com`
`email`	string	The entire email address Example: `jdoe@gmail.com`
`name`	string	The name Example: `jdoe`

Back to Enriched object field types

Common fields

Several fields are common across all event types. Some serve administrative purposes (such as a unique event identifier or the originating sensor) while others are essential for interpreting network traffic, including timestamps and source/destination IP addresses. Each of the following fields is present in every event, with a few exceptions noted in the table below.

Field	Type	Description
`account`	string	The name of the account that owns the event

Field	Type	Description
		Example: `Training`
`customer_id`	string	The code of the account that owns the event Example: `chg`
`dst`	ip-object	The responder to the connection Example: `8.8.8.8`
`event_type`	string	The type of event recorded Example: `smp`
`flow_id`	string	A unique identifier for a flow shared by all events produced from that particular flow Example: `CtjvJR1nIzN4WFSuc7`
`geo_distance`	float	The difference between `src` and dst geo values Example: `1410.373826280689`
`intel`	intel-array	An array of intel-objects matching entities in the event
`sensor_id`	string	The sensor that created the event Example: chg1
`source`	string	The source of the event. Example: `Zeek`
`src`	ip-object	The initiator of the connection Example: `10.10.10.10`
`timestamp`	timestamp	The time at which traffic for the event began Example: `2019-01-01T00:00:00.000Z`
`uuid`	string	A unique identifier for the event Example: `1ca116cb-9262-11e9-b5bf-02472fee9a4a`

The `intel` field is an array of values of type *intel-object*. The table below lists the sub-fields contained within the `intel` field.

Field	Type	Description
`confidence`	string	The overall confidence rating of the intel source Example: `high`
`feed`	string	The name of the intel source Example: `Sinkholes`
`indicator`	string	The matched entity Example: `131.253.18.12`
`indicator_type`	string	The entity type Example: `ip_address`

Field	Type	Description
`is_malicious`	Boolean	Indicates whether the indicator is believed to be malicious Example: `false`
`meta`	string	A JSON string of all metadata provided by the intel source Example: `{"description":"Observed C2 Activity","references":["Fortinet FortiGuard Labs"]}`
`severity`	string	The overall severity rating of the intel source Example: `high`
`timestamp`	timestamp	The creation time of the intel record Example: `2019-01-01T00:00:00.000Z`

Exceptions to common fields

Event type	Exception
DPI	The `flow_id` is not included in the `dpi` events.
Netflow	In NetFlow events, the `src` (source) and `dst` (destination) fields are replaced with `interface_enriched`, a type based on `ip-object`. This enriched type includes everything in `ip-object`. Unique to Netflow, the `src` and `dst` also include the `mac` (MAC address) field
Software	The Software event type does not have `src` and `dst` fields because it is not extracted from raw network traffic. Instead, the record is inferred based on the contents of one or more fields.
Suricata	The Suricata event type does not have a `flow_id` field because it is generated by a completely different process than the other event types. You must match `suricata` events to their associated flows using the IP address and ports of the event.
VPC	VPC Flow fields do not include the common `flow_id` field.

Event fields

The following topics describe the fields unique to each event type.

- BACnet Device control fields
- BACnet Discovery fields
- DNP3 Object fields on page 227
- DNS fields on page 228
- LDAP fields on page 237
- LDAP search fields
- Profinet event
- QUIC fields on page 248
- RDP fields on page 249
- Software fields on page 255
- SSH fields on page 256

- BACnet Property fields
- BACnet header fields
- DCE RPC fields on page 225
- DHCP fields on page 225
- DNP3 fields
- DNP3 Control fields
- DPI fields on page 229
- Flow fields on page 231
- FTP fields on page 233
- HTTP fields on page 233
- Kerberos fields on page 236
- Modbus fields on page 239
- Netflow fields
- Notice Fields on page 242
- NTLM fields on page 244
- Observation fields on page 244
- PE fields on page 245
- SMB file fields on page 251
- SMB mapping fields on page 252
- SMTP fields on page 252
- SNMP fields on page 254
- SSL fields on page 257
- Suricata fields on page 258
- Tunnel fields on page 259
- VPC Flow fields
- x509 fields on page 261

Back to Event Fields.

BACnet Device control fields

A BACnet device control event occurs when BACnet messages like Reinitialize-Device or Device-Communication-Control are detected. These events log administrative actions that affect device availability and behavior.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Property	Type	Description
`bacnet_device_ctrl_ignore_time`	integer	Time in minutes that the device should obey the control command; e.g., in Device-Communication-Control, how long the device is to suppress or enable communications per the request. Example: 5
`bacnet_device_ctrl_invoke_id`	integer	Unique identifier used to correlate a confirmed APDU request (such as Device-Communication-Control or Reinitialize-Device) with its acknowledgment or response. Example: 1
`bacnet_device_ctrl_pdu`	string	The specific BACnet APDU service invoked for device control (e.g., “ReinitializeDevice” or “DeviceCommunicationControl”). Example: reinitialize_device
`bacnet_device_ctrl_pwd_hash`	string	The SHA-256 hash of the password supplied in the Device-Communication-Control or Reinitialize-Device request if required by the device for authentication or to execute the control command. Example: ba7816bf8f01cfea414140de5dae2223b00361a396177a9cb410ff61f20015ad
`bacnet_device_ctrl_result`	string	Outcome of the control operation: one of Success, Error, Reject, or Abort. Example: ERROR

Property	Type	Description
`bacnet_device_ctrl_result_code`	string	If the result was Error, Reject, or Abort, this is the specific error/reject/abort code returned by the device; otherwise often “OK” or similar success-indicator. Example:
`bacnet_device_ctrl_state`	string	The state to which the device is being set by the control service (for instance, the state in Reinitialize-Device such as “coldstart”, “warmstart”, etc.). Example: coldstart
`is_orig`	boolean	True if the message is sent from the originator. Example: true

Back to Event Fields.

BACnet Discovery fields

A BACnet discovery event is created when `Who-Is/I-Am/Who-Has/I-Have` messages are observed, recording device/object identifiers and vendor information for rapid inventory. This log focuses on unconfirmed services used for discovery.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Property	Type	Description
`bacnet_discovery_id`	integer	Numerical part of the device’s identifier (the instance number) used in discovery to uniquely address the device on the network. Example: 1
`bacnet_discovery_instance_num`	integer	The instance number of the object being discovered, combined with object_type to uniquely identify that object within the device. Example: 930101
`bacnet_discovery_object_name`	string	The name property of the object discovered (Object_Name BACnet property), e.g. a human-readable name for the device or object as configured on the BACnet device. Example: FLR12_DEMAND
`bacnet_discovery_object_type`	string	The type of BACnet object that is announced/discovered in the discovery process (for example, Device, Analog-Input, Binary-Output, etc.). Example: device
`bacnet_discovery_pdu`	string	The specific BACnet discovery service in use (for example, “Who-Is” or “I-Am”) Example: who-is

Property	Type	Description
`bacnet_discovery_range`	string	The “Who-Is” discovery range that was used (e.g. “0-4194303”) indicating lower and upper limits of device instance numbers requested to announce themselves; helps scope discovery messages. Example: 1944802-1944802
`bacnet_discovery_type`	string	Type of identifier used to represent the device’s identity (often the “Device” object identifier or its subtype) Example:
`bacnet_discovery_vendor`	string	The vendor identifier or vendor name of the device responding to the discovery, per the BACnet Vendor ID registry. Example: Schneider Electric
`is_orig`	boolean	True if the message is sent from the originator. Example: true

Back to Event Fields.

BACnet Property fields

A BACnet property event is created when Read-Property-Request, Read-Property-ACK, or Write-Property-Request messages are observed, capturing object type, instance number, property identifier, array index, and value. This log focuses on confirmed services used for reading and writing properties.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Property	Type	Description
`bacnet_property_index`	integer	If the property is an array, this is the index of the element being accessed; if omitted or zero, it often means the whole array or default behavior per spec. Example: 1
`bacnet_property_instance_num`	integer	The instance number of the object within the device. Example: 111
`bacnet_property_invoke_id`	integer	The unique identifier used to correlate confirmed APDU property requests (Read-Property or Write-Property) with their acknowledgments in BACnet traffic. Example: 232
`bacnet_property_object_type`	string	The type of BACnet object (e.g. Analog Input, Binary Output, Device, etc.) whose property is being accessed or modified.

Property	Type	Description
		Example: device
`bacnet_property_pdu`	string	The specific BACnet APDU service invoked for device control (e.g. “ReinitializeDevice” or “DeviceCommunicationControl”)) Example: read-property-ack
`bacnet_property_type`	string	The property identifier within the object (e.g. Present_Value, Status_Flags, Description, etc.) being read or written. Example: object-list
`bacnet_property_value`	string	The value of the property (for Read-Property-ACK or Write-Property-Request) as represented in the BACnet message; could be numerical, enumeration, string, etc. Example: device: 111
`is_orig`	boolean	True if the message is sent from the originator. Example: false

Back to Event Fields.

BACnet header fields

A BACnet header event is created when any BACnet/IP packet is seen; the log captures header information for both APDU and NPDU messages. BACnet is a building automation/control protocol used for device discovery, property access, and supervisory functions.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Property	Type	Description
`bacnet_bvlc_func`	string	The BVLC (BACnet Virtual Link Control) function for this BACnet/IP packet (identifies how the packet is being used, e.g. Original-Unicast-NPDU, Forwarded-NPDU, etc.) Example: BVLC_Result
`bacnet_invoke_id`	integer	The unique identifier (invoke ID) used to track Confirmed APDU/NPDU requests and their acknowledgements/responses. Example: 215
`bacnet_pdu_service`	string	The Bacnet service (which service is being invoked or replied to, e.g. ReadProperty, WriteProperty, WhoIs, etc.) Example: read_property

Property	Type	Description
`bacnet_pdu_type`	string	The Bacnet service type (the APDU PDU type, e.g. Confirmed-Request, Unconfirmed-Request, Simple-ACK, Error, etc.) Example: CONFIRMED_REQUEST
`bacnet_result_code`	string	The Error/reject/abort code or reason if the APDU is an Error, Reject, or Abort. This field is not applicable for NPDU context, it will be null. Example: Successful_completion
`is_orig`	boolean	True if the packet is sent from the originator.

DCE RPC fields

A `dce_rpc` event is created when a Distributed Computing Environment / Remote Procedure Call message is observed over a connection, capturing RPC operations like bind, request, or response. This protocol enables clients to execute procedures on remote servers.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`dce_rpc_endpoint`	string	The remote service targeted by the command Example: `samr`
`dce_rpc_operation`	string	The command submitted to the remote service Example: `SamrOpenDomain`
`named_pipe`	string	The name of the target pipe (or the destination port if not named Example: `\pipe\lsass`
`round_trip_time`	float	The time in seconds between command execution and results returned Example: `0.01`

Back to Event Fields.

DHCP fields

A `dhcp` event is created when a Dynamic Host C onfiguration Protocol exchange occurs, such as a client requesting or receiving network addressing from a DHCP server. This protocol is used to dynamically assign IP addresses and other network configuration settings.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`assignment`	ip-object	The IP assigned to the client Example: `10.0.0.10`
`dhcp_msg_type`	string	Shows whether a lease is being requested or acknowledged Example: `Request`
`hostname`	string	The client hostname Example: `bob-pc`
`lease_duration`	float	Number of seconds that the lease is valid Example: `1800`
`lease_end`	timestamp	The time at which the lease expires Example: `2019-06-24T07:31:35.012Z`
`mac`	string	The client MAC address Example:
`trans_id`	int	The transaction ID, ties together requests and acknowledgments. Example: `1191705957`

Back to Event Fields.

DNP3 fields

A `dnp3` event is created when DNP3 (Distributed Network Protocol), commonly used in industrial control systems, logs requests or replies. The protocol enables master-to-outstation communication for monitoring and control.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`dnp3_function_reply`	string	The name of the function message in the reply. Example: `RESPONSE`
`dnp3_function_request`	string	The name of the function message in the request. Example: `CONFIRM`
`dnp3_indication_number`	integer	The response's "internal indication number". Example: `0`

Back to Event Fields.

DNP3 3 Control fields

A `dnp3_control` event is generated when DNP3 control messages—specialized commands for remote control or configuration are observed. It supports supervisory control operations in DNP3 networks.

The following table shows fields unique to the dnp3_control event type:

Field	Type	Description
`dnp3_function_code`	string	Function code (READ or RESPONSE) Example: `RESPONSE`
`dnp3_object_count`	integer	DNP3 object type Example: `32-Bit Binary Counter`
`dnp3_object_type`	string	DNP3 object type Example: `32-Bit Binary Counter`
`dnp3_range_high`	integer	Range (high) of object Example: `9`
`dnp3_range_low`	integer	Range (low) of object Example: `0`
`is_orig`	boolean	True if the packet is sent from the originator Example: `true`

DNP3 Object fields

A `dnp3_object` event is generated when DNP3 object-level constructs (such as analog or binary inputs/outputs) are seen in the traffic, facilitating insight into SCADA-style data models. It reflects structured data exchanged via DNP3.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`dnp3_function_code`	string	Function code (READ or RESPONSE) Example: `RESPONSE`
`dnp3_object_count`	integer	DNP3 object type Example: `32-Bit Binary Counter`
`dnp3_object_type`	string	DNP3 object type Example: `32-Bit Binary Counter`
`dnp3_range_high`	integer	Range (high) of object Example: `9`
`dnp3_range_low`	integer	Range (low) of object Example: `0`

Field	Type	Description
`is_orig`	boolean	True if the packet is sent from the originator Example: `true`

Back to Event Fields.

DNS fields

A `dns` event is created when a Domain Name System query or response message is captured over the network. DNS enables the resolution of human-friendly domain names to IP addresses.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`answers`	host-object-array	The answers returned by the DNS server for the query Example: `[103.2.116.79, 103.2.116.83]`
`proto`	string	The transport layer protocol used Example: `udp`
`qtype`	int	The numeric code of the query type Example: `1`
`qtype_name`	string	The string name of the query type Example: `A`
`query`	domain-object	The domain being queried Example: `www.google.com`
`rcode`	int	The numeric code of the result Example: `0`
`rcode_name`	int	The string name of the result Example: `NOERROR`
`rejected`	Boolean	Indicates whether the query was rejected by the server Example: `false`
`ttls`	int-array	An array of TTL values, one per result Example: `[299, 299]`

Back to Event Fields.

DPI fields

A `dpi` (Deep Packet Inspection) event is created by the Fortinet IPS (Intrusion Prevention System) engine running on the sensor which logs informative and pattern matching based events. The IPS engine logs AppID (Applications seen by the engine for software and protocols), IDS (signatures for vulnerabilities), OT Protocols/Threats (Operational Technology based protocol parsing and signatures), Botnet (Botnet based traffic patterns), and Info (informational events about protocols).

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field Name	Type	Description
`dpi_alert_category`	string	Type of category of the IPS signature. • `Info`: IDS with informational severity • `AppID`: Application control • `IDS`: Intrusion Detection • `Botnet`: IDS's botnet specific signature • `OT - Threats`: IDS for Operational Technology • `OT - Protocol`: AppCtrl for Operational Technology Example:`IDS`
`dpi_alert_severity`	integer	Severity of the triggered IPS signature. • Info: 0 • Low: 1 • Medium: 2 • High: 3 • Critical: 4 Example:`0`
`dpi_alert_signature`	string	The triggered IPS signature name. Example:`ITCM.Class.D_Wayside.Status.Message.WIUStatus.Timed.Beacon`
`dpi_alert_signature_id`	integer	Attack ID or ID of the IPS signature. Example:`12343`
`dpi_app_behavior`	array	Possible behavior for the application in which the triggered IPS signature refers to. Example:`Evasive`
`dpi_app_category`	string	The application category for the triggered IPS signature, if there is any. Example:`Operational.Technology`
`dpi_app_language`	string	Language used in the application in which the triggered IPS signature refers to. Example:`N/A`
`dpi_app_name`	array	Name of the application. Example:`Other`

Field Name	Type	Description
`dpi_app_os`	array	OS of the application or vulnerable system/devices. Example:`All`
`dpi_app_technology`	array	Technology group or type for the application in which the triggered IPS signature refers to. Example:`Client-Server`
`dpi_app_vendor`	string	Vendor of the application in which the triggered IPS signature refers to. Example:`Other`
`dpi_expected_port`	string	Default port and protocol for the application in which the triggered IPS signature is referring to. Example:`UDP/1900`
`dpi_parent_vuln_id`	integer	ID of the IPS signature that link to the triggered IPS signature. Example:`56843`
`dpi_rulegroup`	string	Which group the triggered IPS signature belongs to. Example:`SCADA`
`dpi_ruleset_rev`	integer	Version number for the triggered IPS signature. Example:`13401`
`dpi_session_id`	integer	Session ID for the traffic. Example:`0`
`dpi_sig_cve`	array	ID for the CVE reference. Example:`20050380`
`dpi_ssl_decrypt_req`	boolean	Does the current IPS signature need SSL decryption to work. Example:`False`
`dpi_vuln_id`	integer	Vulnerablity ID or Applicatioin ID for the IPS signature (Note: One VID could contain multiple AID). Example:`33456`
`dpi_vuln_type`	string	Type of vulnerability this IPS signature is related to. Example:`Other`

The common field of `flow_id` is not included in the `dpi` events.

Back to Event Fields.

Flow fields

A `flow` event is created when a unidirectional or bidirectional network flow is identified, summarizing traffic between endpoints over time, such as packet count, byte count, and states. A network flow is defined by a unique combination of `src.ip`, `src.port`, `dst.ip`, `dst.port`, and `proto`.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`community_id`	string	An additional flow identifier for joining Flow events. Example: `1:f69i+MdCEA8QnAKnKVE0Pyyta24=`
`duration`	float	The number of seconds the flow lasted Example: `7s`
`flow_state`	string	Lifecycle summary of the connection observed for this flow. Includes standard Zeek/Bro connection states and periodic P* states. Example: `SF` Supported values:

flow_state	Description
`S0`	Connection attempt seen, no reply.
`S1`	Connection established, not terminated.
`SF`	Normal establishment and termination.
`REJ`	Connection attempt rejected.
`S2`	Connection established and close attempt by originator seen (but no reply from responder).
`S3`	Connection established and close attempt by responder seen (but no reply from originator).
`RSTO`	Connection established, originator aborted (sent a RST).
`RSTR`	Responder sent a RST.
`RSTOS0`	Originator sent a SYN followed by a RST, we never saw a SYN-ACK from the responder.
`RSTRH`	Responder sent a SYN ACK followed by a RST, we never saw a SYN from the (purported) originator.

Field	Type	Description
		<table><tr><th>flow_state</th><th>Description</th></tr><tr><td>SH</td><td>Originator sent a SYN followed by a FIN, we never saw a SYN ACK from the responder (hence the connection was “half” open).</td></tr><tr><td>SHR</td><td>Responder sent a SYN ACK followed by a FIN, we never saw a SYN from the originator.</td></tr><tr><td>OTH</td><td>No SYN seen, just midstream traffic (a “partial connection” that was not later closed).</td></tr></table> Additionally, FortiNDR Cloud logs *P** flow states. These states are logged for long-lived connections once every 24 hours, with the *flow_state* reflecting the current state of the TCP/UDP state machine. The byte totals logged are cumulative since connection start, rather than incremental since the previous log entry. In practice, this typically results in *PS0* (one-sided connection, retry under the 5-minute timeout), *PS1* (two-sided connection where the start was observed), and *POTH* (the start of the connection was not observed).
`proto`	string	The transport layer protocol used Example: `tcp`
`service`	string	The application(s) observed in the flow, if any Example: `http`
`total_ip_bytes`	int	The total combined bytes transmitted over the connection Example: `927 bytes`
`total_pkts`	int	The total combined packets transmitted over the connection Example: `11`
`upload_percent`	int	The percentage of bytes transmitted by the src for the flow (56% == 56) Example: `56%`

Back to Event Fields.

FTP fields

An `ftp` event is created when File Transfer Protocol commands or responses are observed during an FTP session. This protocol is used for transferring files between client and server.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`data_channel.dst`	ip-object	The destination of the data channel Example: `10.0.0.2`
`data_channel.geo_distance`	float	The distance (in miles) between the IP addresses of the data channel Example: `5077.89`
`data_channel.passive`	Boolean	Indicates whether the session is in passive mode Example: `True`
`data_channel.src`	ip-object	The source of the data channel Example: `10.0.0.10`
`files`	file-array	Files transferred over the session Example: N/A
`ftp_arg`	string	The full argument string supplied to the command Example: `ftp://10.0.0.2/secrets.zip`
`ftp_command`	string	The client command Example:RETR
`reply_code`	int	The server response code to the command Example: `227`
`reply_msg`	string	The server response string to the command Example: `Entering Passive Mode (10,0,0,2,197,36)`
`username`	string	The username used to establish the connection Example: `Admin101`

Back to Event Fields.

HTTP fields

An `http` event is created when HTTP requests or responses—including headers and message boundaries are processed over HTTP connections. HTTP is the foundational protocol for web communications.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`cookie_vars`	string-array	Variable names extracted from all cookies. Example: `disp.prefs,_utmz ,_utmc,_utma, TS01f95106, _utmb`
`files`	file-object-array	Files downloaded over the HTTP connection
`headers.accept`	string-array	The content of the Accept header Example: `[image/webp, image/apng, image/*, */*;q=0.8]`
`headers.client_header_names`	string-array	The vector of HTTP header names sent by the client. Example: `Cache-Control, Connection, Pragma, Content-Type, User-Agent, X-Havoc, X-Havoc-Agent, Content-Length, Host`
`headers.content_md5`	string	The computed MD5 hash of the headers content Example: `d41d8cd98f00b204e9800998ecf8427e`
`headers.content_type`	string-array	The contents of the Content Type header Example: `[text/xml; charset="utf-8"]`
`headers.cookie_length`	int	The length of the cookie in bytes Example: `194`
`headers.location`	url-object	The content of the Location header Example: `http://amupdatedl3.microsoft.com/server/amupdate/metadata/UniversalManifest.cab`
`headers.origin`	url-object	The content of the Origin header Example: `http://go.com`
`headers.proxied_ip_clients`	ip-object-array	The sequence of IPs the HTTP connection is proxied through Example: `[172.16.0.1, 172.16.0.2]`
`headers.refresh.refresh`	string	The full content of the Refresh header Example: `1;URL=http://travelingtravelerhome.wordpress.com/`
`headers.refresh.timeout`	int	The timeout period in seconds Example: `1`
`headers.refresh.uri`	uri-object	The URI of the Refresh header Example: `http://travelingtravelerhome.wordpress.com/`
`headers.server`	string	The web server software Example: `Microsoft-IIS/6.0`
`headers.server_header_names`	string-array	The vector of HTTP header names sent by the server.

Field	Type	Description
		Example: `VIA, DATE SERVER, CONNECTION, X-2SENDPT1, X-WSENDPT2, CONTENT-LENGTH`
`headers.x_powered_by`	string	The application software running on the server Example: `ASP.NET`
`host`	host-object	The content Host header Example: `www.google.com`
`info_msg`	string	The message returned with a 100-level response code Example: `Continue`
`method`	string	The HTTP method selected Example: `GET`
`proxied`	string-array	A list of proxy steps Example: `PROXY-CONNECTION -> Keep-Alive`
`referrer`	url-object	The content of the Referrer header Example: `http://au.search.yahoo.com/search?p=planetside.co.uk&fr=sfp&fr2=sb-top-search`
`request_len`	int	The length in bytes of the request Example: `0`
`request_mimes`	string-array	The fingerprinted MIME-type(s) of the request content, use instead of `request_mime` Example: text/plain
`response_len`	int	The length in bytes of the response Example: `24`
`response_mimes`	string-array	The fingerprinted MIME-type of the response content, use instead of `response_mime` Example: `text/html`
`status_code`	int	The numeric code of the server's response Example: `200`
`status_msg`	string	The string name of the server's response Example: `OK`
`trans_depth`	int	The depth of redirects Example: `4`
`uri`	uri-object	The full URI of the request Example:`/index.php`
`user_agent`	string	The content of the UserAgent header Example: `Mozilla/5.0 (Windows NT 6.1; WOW64; Trident/7.0; rv:11.0) like Gecko`

Field	Type	Description
`username`	string	The username used with Basic Auth, if any Example: `dave`

Back to Event Fields.

Kerberos fields

A `kerberos` event is generated when Kerberos authentication messages (e.g., AS or TGS requests/replies) are detected. Kerberos is a network authentication protocol that uses tickets to allow nodes to prove their identity.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`cipher`	string	The cipher suite used to encrypt the ticket Example: `aes256-cts-hmac-sha1-96`
`client`	string	The client that requested the ticket; machine accounts have a `$` at the end of their name but user accounts do not. Example: `jane.doe/ACME.CORP`, `financewks008$/ACME.CORP`
`client_cert_fuid`	string	Client certificate file unique ID Example: `Xbtku3TdsfdsdfasdfA8VNsk`
`client_cert_subject`	string	Client certificate Subject field Example: `CN=C865433`
`error_msg`	string	The error message returned for failed requests Example: `KDC_ERR_CLIENT_NAME_MISMATCH`
`forwardable`	Boolean	Indicates whether the ticket's forwardable flag is set Example: `True`
`renewable`	Boolean	Indicates whether the ticket's renewable flag is set Example: `True`
`request_type`	string	The type of ticket requested, either a ticket-granting ticket from the authentication server (AS) or a service ticket from the ticket-granting server (TGS) Example: `AS`, `TGS`
`server_cert_fuid`	string	Server certificate file unique ID Example: `FvAdJGsjeXuhSvE9m`
`server_cert_subject`	string	Server certificate Subject field Example: `CN=dc09.google.com`

Field	Type	Description
`service`	string	The service for which a ticket is being requested Example: `krbtgt/ACME.CORP`
`success`	Boolean	Indicates whether the request was successful Example: `True`
`ticket_duration`	float	The ticket duration in seconds Example: `86400`
`ticket_from`	timestamp	Time the ticket is good from Example: `2015-09-13T02:48:05.000Z`
`ticket_till`	timestamp	Time the ticket is good until Example: `2037-09-13T02:48:05.000Z`

Back to Event Fields.

LDAP fields

An `ldap` event is generated when LDAP (Lightweight Directory Access Protocol) messages—such as authentication, search, or directory operations—are observed. This protocol provides directory services, like querying user or organizational data.

The following table shows the fields contained in this event type excluding the previously identified common fields:

This event type is supported in Sensor version 2.2.0 and later.

The following table shows fields unique to the `ldap` event type:

Field	Type	Description
`ldap_argument`	string	Additional arguments this message includes. Example: `REDACTED`
`ldap_diagnostic_message`	string	Diagnostic message if the LDAP message contains a result.
`ldap_message_id`	integer	The unique identifier that is used to correlate requests and responses. Example: `2`
`ldap_object`	string	The objects names this message refers to. Example: `ATRLAB\\Administrator`
`ldap_opcode`	string	The operation code indicating what type of message it is.

Field	Type	Description
		Example: `bind, simple`
`ldap_result`	string	The result code if the message contains a result. Example: `success`
`ldap_version`	integer	LDAP version. Example: `3`

Back to Event Fields.

LDAP Search fields

A `ldap_search` event is created when a client performs an LDAP search operation.

This event type is supported in Sensor version 2.2.0 and later.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`ldap_diagnostic_message`	string	Diagnostic message if the LDAP message contains a result.
`ldap_message_id`	integer	The unique identifier that is used to correlate requests and responses. Example: *2*
`ldap_result`	string	Result code of search operation. Example: `success`
`ldap_search_attribute`	string	A list of attributes that were returned in the search. Example: 2
`ldap_search_base_object`	string	Base search objects. Example: 2
`ldap_search_deref_aliases`	string	Set of deref alias. Example: 2
`ldap_search_filter`	string	A string representation of the search filter used in the query. Example: 2
`ldap_search_result_count`	integer	Number of results returned. Example: 2

Field	Type	Description
`ldap_search_scope`	string	Set of search scopes. Example: `2`
`source`	string	The source of the event. Example: *Zeek*

Back to Event Fields.

Modbus fields

A `modbus` event is created when Modbus protocol commands or responses—typically used in industrial automation systems—are captured. This allows reading or writing of registers or coil values in connected devices.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`is_orig`	boolean	Example: `true`
`modbus_address`	integer	Starting address of value(s) field.
`modbus_function`	string	The name of the function message that was sent. Example: `READ_INPUT_REGISTERS`
`modbus_quantity`	integer	Number of addresses/values read or written to.
`modbus_request_response`	string	REQUEST or RESPONSE
`modbus_tid`	integer	Modbus transaction identifier
`modbus_unit`	integer	Modbus terminal unit identifier.
`modbus_values`	string[]	Value(s) of coils, discrete_inputs, or registers read/written to. Example: `555,0,100`

Back to Event Fields.

Netflow fields

A `netflow` event is created when IP traffic flow data—typically collected by routers or switches—is captured and exported for analysis. This allows visibility into network usage patterns, including source and destination IPs, protocols, ports, and byte counts.

- A NetFlow annual subscription license is required for FortiNDR Cloud to ingest third-party logs for anomaly detection.
- Only NetFlow-based botnet detections are currently displayed. Detections for spam, phishing, Tor, and proxy traffic are available at this time. Additionally, an IOC (Indicator of Compromise) risk score may not be shown for every IP address.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`netflow_bytes`	integer	Number of bytes in a flow. Example: `106`
`netflow_dst_net`	string	Destination network address associated with a particular network flow with the mask. Example: `0.0.0.0/0`
`netflow_dst_vlan`	integer	Virtual LAN identifier associated with egress interface. Example: `0`
`netflow_etype`	string	Ethernet type (0x0800 for IPv4). Entire list is here: https://en.wikipedia.org/wiki/EtherType Example: `IPv4`
`netflow_forwarding_status`	integer	Forwarding status is encoded on 1 byte with the 2 left bits giving the status and the 6 remaining bits giving the reason code.Status is either unknown (00), Forwarded (10), Dropped (10) or Consumed (11). Example: `0`
`netflow_frag_id`	integer	The fragment ID. Example: `19093`
`netflow_frag_offset`	integer	The fragment-offset value from fragmented IP packets. Example: `0`
`netflow_icmp_code`	integer	Code of the ICMP message. Example: `0`
`netflow_icmp_type`	integer	ICMP flags Example: `0`
`netflow_input_interface`	integer	Input interface. Example: `512`
`netflow_ip_flags`	integer	IP flags Example: `0`
`netflow_ip_tos`	integer	IP Type of Service.

Field	Type	Description
		Example: `0`
`netflow_ip_ttl`	integer	TTL value observed for packets of the flow. Example: `64`
`netflow_ipv6_flow_label`	integer	IPv6 flow label as in RFC 2460 definition. Example: `0`
`netflow_layer_size`	array	Size of protocols seen in the flow. Example: `[14, 4, 20, 8]`
`netflow_layer_stack`	array	Protocols seen in this flow. Example: `[Ethernet, MPLS, IPv4, ICMP]`
`netflow_output_interface`	integer	Output interface. Example: `0`
`netflow_sampled`	integer	Denominator of how frequently data is collected. Meaning a sampling rate of 100 means one out of every 100 packets is sampled. Helps reduce the load on network devices and collectors by only exporting a portion of the traffic. Example: `1`
`netflow_sampler_address`	string	The IP address of the network device (typically a router) that is performing packet sampling and exporting NetFlow data. Example: `169.254.0.2`
`netflow_seq_num`	integer	A cumulative counter that increments with each exported datagram to detect and account for any missing or dropped NetFlow datagrams. Example: `766`
`netflow_source`	string	Type of netflow Example: `IPFIX`
`netflow_src_net`	string	Source network address associated with a particular network flow with the mask. Example: `0.0.0.0/0`
`netflow_src_vlan`	integer	Virtual LAN identifier associated with ingress interface. Example: `0`
`netflow_tcp_flags`	integer	TCP flags Example: `0`
`netflow_timestamp_end`	string	Time the flow ended in nanoseconds.
`netflow_timestamp_received`	string	Timestamp in nanoseconds when the flow message

Field	Type	Description
		was received by the NetFlow collector or analysis system.
`netflow_vlan_id`	integer	Allows you to associate network traffic flows with their respective VLANs. Example: `0`
`proto`	string	Protocol used in the traffic. Example: `TCP`
`tag`	string	The type of event Example: `flow`
`total_pkts`	integer	Number of packets in a flow. Example: `1`
`dst.ip_bytes`	integer	The number of bytes transmitted by the IP address
`dst.pkts`	integer	The number of packets transmitted by the IP address
`switched`	boolean	If the source and destination IPs are switched due to port values. Example: `false`
`src.ip_bytes`	integer	The number of bytes transmitted by the IP address
`src.pkts`	integer	The number of packets transmitted by the IP address

In NetFlow events, the `src` (source) and `dst` (destination) fields are replaced with `interface_enriched`, a type based on `ip-object`. This enriched type includes everything in `ip-object`. Unique to Netflow, the `src` and `dst` also include the `mac` (MAC address) field

Back to Event Fields.

Notice Fields

A `notice` event is raised when unusual or noteworthy activity is detected and logged as a security or policy notification. It flags anomalies or policy-triggered events across Zeek's analysis.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`application`	application	The classified application for a flow
`dst_ip`	string	The IP of the responder to the connection Example: `8.8.8.8`

Field	Type	Description
`dst_ip_enrichments`	ip_enrichments	Enrichments for an IP
`dst_port`	integer	The port of the responder to the connection Example: `53`
`file_desc`	string	Description of a file to provide more context. For example, if a notice was related to a file over HTTP, the URL of the request would be shown.
`file_mime_type`	string	If the notice event is related to a file, this will be the mime type of the file.
`fuid`	string	A file unique ID if this notice is related to a file.
`msg`	string	Description of activity noticed. Example: `10.1.0.47 appears to be guessing SSH passwords (seen in 30 connections).`
`n`	integer	Associated count, or perhaps a status code.
`note`	string	Notice type Example: `SSH::Password_Guessing`
`notice_actions`	string	The actions which have been applied to this notice. Example: `[Notice::ACTION_LOG]`
`peer_descr`	string	Textual description for the peer that raised this notice, including name, host address and port.
`proto`	string	The transport protocol.
`src_ip`	string	The IP of the initiator of the connection Example: `10.10.10.10`
`src_ip_enrichments`	ip_enrichments	Enrichments for an IP
`src_port`	integer	The port of the initiator of the connection Example: `52843`
`sub`	string	Technical details of the activity. Example: `Sampled servers: 10.1.0.86, 10.1.0.86, 10.1.0.86, 10.1.0.86, 10.1.0.86`
`suppress_for`	number	This field indicates the length of time that this unique notice should be suppressed.
`tag`	string	The type of event Example: `flow`

Back to Event Fields.

NTLM fields

An `ntlm` event is generated when NT LAN Manager authentication exchanges are seen, including domain, username, hostname, and whether authentication succeeded. This is a Microsoft authentication protocol.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`auth_domain`	string	The domain used to authenticate the client Example: `ACME`
`hostname`	string	The client hostname used Example: `FINANCEWKS008`
`ntlm_status`	string	String indicating the result of the authentication Example: `SUCCESS`
`success`	Boolean	Indicates whether the authentication succeeded Example: `True`
`username`	string	The client username used Example: `sqlservice`

Back to Event Fields.

Observation fields

An `observation` event is created when the FortiNDR Cloud analytics backend identifies a correlation of information of interest. See below for valid values for the following fields:

You can view the list of observations in the *Observations* widget in the *Default Dashboard* . For more information, see:

- `observation_category`: asset, account, software, flow, file, relationship
- `observation_class`: anomalous, newly observed, specific

Observations run independently from the metadata extraction process, and are not tied to `flow` events with a `flow_id`. Additionally, an `observation` event may only have one of `src.ip` or `dst.ip`, although it could contain both.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`evidence_end_timestamp`	timestamp	The timestamp for which the flagged activity ended.

Field	Type	Description
		Example: `2019-01-01T00:00:00.000Z`
`evidence_iql`	string	An IQL statement that attempts to identify the events used to generate the observation. Example: `src.ip = '10.10.10.10' AND customer_id = 'chg' AND dce_rpc:dce_rpc_operation = 'NetrSessionEnum' AND timestamp >= t'2019-01-01T22:00:00.000000Z' AND timestamp <= t'2019-01-01T22:10:00.000000Z'`
`evidence_start_timestamp`	timestamp	The timestamp for which the flagged activity began. Example: `2019-01-01T00:00:00.000Z`
`observation_category`	string	The subject of an observation. Example: `relationship`
`observation_class`	string	The class of what was observed about the subject. Example: `specific`
`observation_confidence`	string	The confidence (high, medium, or low) in the model output to what was attempted to be observed. Example: `high`
`observation_title`	string	The title of what was attempted to be detected - similar to a suricata sig name. Example: `High Count of NetSession Destinations`
`observation_uuid`	string	A unique identifier for the model used to generate the observation. Multiple models may exist for the same title. Example: `ac33189b-ee31-4f5e-b6a1-dcb63d9a7295`
`sensor_ids`	string array	A list of sensors from which activity was used as part of the observation. Example: `[chg1,chg2,chg3]`

Back to Event Fields.

PE fields

A `pe` event is created when a Portable Executable file (e.g., Windows .exe or .dll) is transferred or extracted during file analysis. The PE format is the executable file format for Windows binaries.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`compile_timestamp`	timestamp	The compile timestamp extracted from the file Example: `2015-11-12T10:23:51.000Z`
`file`	file-object	The enriched file properties (hashes, size, MIME-type) Example: N/A
`has_cert_table`	Boolean	Indicates whether the file has an attribute certificate table Example: `True`
`has_debug_data`	Boolean	Indicates whether the file has a debug table Example: `True`
`has_export_table`	Boolean	Indicates whether the file has an export table Example: True
`has_import_table`	Boolean	Indicates whether the file has an import table Example: `True`
`id`	string	An internal unique identifier for the file Example: `FrkSk6Y0mqKGxMBF6`
`is64_bit`	Boolean	Indicates whether the file is 64-bit Example: `True`
`is_exe`	Boolean	Indicates whether the file is executable or just an object Example: `True`
`machine`	string	The architecture the file was compiled for Example: `I386`
`os`	string	The OS the file was compiled for Example: `Windows XP`
`section_names`	string-array	An array of section names extracted from the file Example: `[.text, .rdata, .data, .rsrc]`
`subsystem`	string	The subsystem the file was compiled for Example: `WINDOWS_GUI`
`uses_aslr`	Boolean	Indicates whether the file supports ASLR Example: `True`
`uses_code_integrity`	Boolean	Indicates whether the file enforces code integrity checks Example: `True`
`uses_dep`	Boolean	Indicates whether the file supports DEP Example: `True`

Field	Type	Description
`uses_seh`	Boolean	Indicates whether the file uses SEH Example: `True`

Back to Event Fields.

Profinet event

A profinet event is created by the use of PROFINET an Ethernet protocol for communication between devices in industrial automation systems.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field Name	Type	Description
`profinet_activity_uuid`	string	Identifies communication relationships
`profinet_auth`	integer	Authentication protocol - set to 0 for no authentication
`profinet_broadcast`	boolean	Flag 1 Bit 6 Broadcast (is the call a broadcast)
`profinet_cancel_req`	boolean	Flag 2 Bit 1 Cancel was pending at call end (a cancellation request was received from the client for a specific remote procedure call (RPC), but the call completed before the cancellation could be processed.)
`profinet_char_encoding`	string	Character encoding: ASCII, EBCDIC
`profinet_fack`	string	Version Fack
`profinet_float_encoding`	string	Floating point representations: IEEE, VAX, CRAY, etc.
`profinet_frag`	boolean	Flag 1 Bit 2 Fragment
`profinet_frag_num`	integer	Fragment number set to the number of the current fragment.
`profinet_hint`	integer	Activity hint
`profinet_idempotent`	boolean	Flag 1 Bit 5 Idempotent
`profinet_int_endian`	string	Integer encoding: Big Endian or Little Endian
`profinet_interface_hint`	integer	Interface hint
`profinet_interface_major`	integer	Interface version major
`profinet_interface_minor`	integer	Interface version minor
`profinet_interface_uuid`	string	Identifies the interface of an IO device, controller, etc.
`profinet_last_frag`	boolean	Flag 1 Bit 1 Last Fragment

Field Name	Type	Description
`profinet_length`	integer	Length of body set to the number of octets of NDRDdata in the current frame.
`profinet_max_frag`	integer	Maximum fragment size
`profinet_max_tsdu`	integer	Maximum Tsdu
`profinet_maybe`	boolean	Flag 1 Bit 4 Maybe (the client sends a request but does not wait for a response)
`profinet_no_frag`	boolean	Flag 1 Bit 3 No fragment acknowledge requested
`profinet_object_uuid`	string	Object instance within a physical device
`profinet_operation_num`	string	Operation number identifies the PNIO service supported by the PNIO interfaces.
`profinet_request_type`	string	Packet Type: Request, Response, Fault, etc.
`profinet_reserved_bit0`	boolean	Flag 1 Bit 0 Reserved for implementation
`profinet_reserved_bit7`	boolean	Flag 1 Bit 7 Reserved for implementation
`profinet_rpc_version`	integer	Used RPC version
`profinet_sel_ack`	array	Array of selective ACK
`profinet_sel_ack_len`	integer	Selective ACK length
`profinet_seq_num`	integer	Used with activity_UUID to uniquely identify a RPC call.
`profinet_serial_high`	integer	The high octet of the fragment number of the call
`profinet_serial_low`	integer	The low octet of the fragment number of the call
`profinet_serial_num`	integer	Serial number
`profinet_server_boot_time`	integer	Server boot time
`profinet_win_size`	integer	Window size
`proto`	string	Transport protocol

Back to Event Fields.

QUIC fields

A `quic` event is generated when QUIC protocol activity—Google's transport layer network protocol combining UDP and TLS—is detected, providing performance and security for web traffic.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`quic_client_initial_dst_conn_id`	string	Destination Connection ID (DCID). This DCID is used for routing and packet protection by client and server. Example: `95412c47018cdfe8`
`quic_client_protocol`	string	QUIC Application-Layer Protocol Negotiation (ALPN) extension. This is the extension's first entry. Example: `h3`
`quic_client_src_conn_id`	string	Source Connection ID chosen by the client in its INITIAL packet. This ID is used for packet protection and is typically random and unpredictable. Example: `4823dfc5a047e6acd230b5c5e047ced9b0a6b542`
`quic_history`	string	Provides a history of QUIC protocol activity in a connection, similar to the history field in Conn. Example: `ISisH`
`quic_server_src_conn_id`	string	A QUIC-supported server responds to a DCID by selecting a Source Connection ID (SCID). Occurs within the server's first INITIAL packet. Example: `0130dfc5a047e6acd230b5c5e047ced9b0a6bbf0`
`quic_version`	string	A string interpretation of the QUIC version number, usually "1" or "quicv2" Example: `1`
`server_name_indication`	ip_or_domain_enriched	An IP or domain with its enrichments

Back to Event Fields.

RDP fields

An `rdp` event is created when Remote Desktop Protocol sessions are observed, capturing details like client build, keyboard layout, desktop size, and security negotiation. It tracks remote Windows desktop connections.

Authentication cannot always be determined as the necessary data may be encapsulated within an encrypted tunnel. Therefore, the `result` field may contain a "best-guess" based on available data.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`cert_count`	int	The number of certificates seen Example: `0`
`cert_permanent`	Boolean	Indicates if the provided certificate or certificate chain is permanent Example: `True`
`cert_type`	string	The type of certificate used if the connection is encrypted with native RDP encryption Example: `RSA`
`client_build`	string	The client RDP version Example: `RDP 5.1`
`client_dig_product_id`	string	The client product ID Example: `715e03e8-6eef-4c53-b022-rbcd967`
`client_name`	string	The client hostname Example: `bob-PC`
`cookie`	string	The truncated account name used by the client Example: `bob`
`desktop_height`	int	The client desktop height Example: `1080`
`desktop_width`	int	The client desktop width Example: `1920`
`encryption_level`	string	The encryption level used Example: `Client compatible`
`encryption_method`	string	The encryption method used Example: `128bit`
`keyboard_layout`	string	The client keyboard layout (language) Example: `English -United States`
`requested_color_depth`	string	The color depth requested by the client in the high_color_depth field Example: `32bit`
`result`	string	The result for the connection, derived from a mix of RDP negotiation failure messages and GCC server create response messages Example: `Succeed`
`security_protocol`	string	Security protocol chosen by the server Example: `RDP`

Back to Event Fields.

SMB file fields

An `smb_file` event is generated when files transferred over SMB/CIFS are observed, logging file-related actions like creation, modification, renaming, with metadata like paths and timestamps. This monitors file-level operations in SMB sessions.

The following table shows the fields contained in this event type excluding the previously identified common fields:.

Field	Type	Description
`files`	file-array	Files transferred over the SMB connection Example: N/A
`files.accessed_timestamp`	timestamp	The last time the file was accessed Example: `2018-04-08T22:48:07.958Z`
`files.bytes`	int	The file's size in kilobytes Example: `145922`
`files.changed_timestamp`	timestamp	The last time the file's metadata changed Example: `2018-04-08T22:48:07.958Z`
`files.created_timestamp`	timestamp	The time the file was created Example: `2018-04-08T22:48:07.958Z`
`files.modified_timestamp`	timestamp	The last time the file's content changed Example: `2018-04-08T22:48:07.958Z`
`files.name`	string	The post-transfer name of the file (can be renamed before writing to disk) Example: `secrets.zip`
`files.previous_name`	string	The pre-transfer name of the file Example: `exfil.zip`
`files.smb_path.path`	string	The full network path to the target share Example: `\\DYNACCOUNTIC-DC.dynaccountic.com\sysvol`
`files.smb_path.share`	string	The target network share Example: `sysvol`
`files.smb_path.system`	string	The target host Example: `DYNACCOUNTIC-DC.dynaccountic.com`
`smb_action`	string	The action taken on the files Example: `SMB::FILE_OPEN`

Back to Event Fields.

SMB mapping fields

An `smb_mapping` event is created when an SMB share is mapped, capturing tree paths, share types (disk, printer, pipe), and native file system info. It tracks resource sharing mappings over SMB.

The following table shows the fields contained in this event type excluding the previously identified common fields:.

Field	Type	Description
`native_file_system`	string	The file system type on the target host (for Disk shares) Example: `NTFS`
`share_type`	string	The type of share established Example: `DISK`
`smb_path.path`	string	The full network path to the target share Example: `\\DYNACCOUNTIC-DC.dynaccountic.com\sysvol`
`smb_path.share`	string	The target network share Example: `sysvol`
`smb_path.system`	string	The target host Example: `DYNACCOUNTIC-DC.dynaccountic.com`
`smb_service`	string	The service used to establish a connection to the share Example: `IPC`

Back to Event Fields.

SMTP fields

An `smtp` event is created when Simple Mail Transfer Protocol messages—such as MAIL FROM, RCPT TO, HELO/EHLO—are observed during an email session. This protocol is used to send email between servers.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
`date`	string	The content of the Date header Example: `Thu, 12 Jul 2015 17:59:01 -0400 (EDT)`
`files`	file-object-array	An array of the files attached to the email
`first_received`	string	The full content of the first Received header

Field	Type	Description
		Example: `from JIM@GMAIL.COM ([198.51.100.1]) by SALLY@GMAIL.COM ([101.9.210.120]) with mapi id 14.01.1039.013; Thu, 12 Jul 2015 18:09:44 -0500`
`from`	email-object	The content of the From header Example: `jdoe@gmail.com`
`helo`	host-object	The argument supplied to the HELO command Example: `client.example.com`
`in_reply_to`	string	The Message-ID in the In-Reply-To header Example: `<b8bba2baae4c2a08fdff4e223458577d@gmail.com>`
`is_webmail`	Boolean	Indicates whether the message was sent through a webmail interface Example: `true`
`last_reply`	string	The last message the server sent to the client Example: `250 Message accepted for delivery`
`mailfrom`	string	The argument supplied to the MAIL FROM command Example: `support@acme.corp`
`msg_id`	string	The Message-ID of the message Example: `<b8bba2baae4c2a08fdff4e223458577d@gmail.com>`
`path`	ip-object-array	The message transmission path extracted from the Received headers Example: `[192.161.0.200, 204.148.78.113]`
`rcptto`	string	The argument supplied to the RCPT TO command Example: `jdoe@gmail.com`
`reply_to`	email-object	The content of the Reply-To header Example: `jdoe@gmail.com`
`second_received`	string	The content of the second Received header Example: `from JIM@GMAIL.COM ([198.51.100.1]) by SALLY@GMAIL.COM ([101.9.210.120]) with mapi id 14.01.1039.013; Thu, 12 Jul 2015 18:09:44 -0500`
`subject`	string	The content of the Subject header Example: `Click this link!`
`tls`	Boolean	Indicates whether the connection switched to using TLS Example: `true`

Field	Type	Description
`to`	email-object-array	The content of the To header Example: `[jdoe@gmail.com, kdoe@gmail.com]`
`trans_depth`	int	The depth of this message transaction where multiple messages were transferred in a single connection Example: `1`
`urls`	string-array	A list of URLs extracted from the message Example: `[http://malware.pwn//root.ps1, https://www.google.com]`
`user_agent`	string	The content of the client's User-Agent header Example: `SquirrelMail/1.4.22`
`x_originating_ip`	ip-object	The content of the X-Originating-IP header Example: `8.8.8.8`

Back to Event Fields.

SNMP fields

An `snmp` event is created when Simple Network Management Protocol messages—used for monitoring and managing network devices—are detected, including version, community string, and request types. It supports network device telemetry.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
`snmp_community`	string	Community string of the first packet associated with the session Example: `public`
`snmp_display_string`	string	A system description of the SNMP responder endpoint Example: `Roma v1.9 v9.5.0_W EQ`
`snmp_duration`	number	Amount of time between the first in the session and the latest one seen in seconds Example: `12.209241`
`snmp_get_bulk_requests`	integer	Number of variable bindings in GetBulkRequest PDUs seen for the session Example: `3`
`snmp_get_requests`	integer	Number of variable bindings in GetRequest/GetNextRequest PDUs seen for the session

Field	Type	Description
		Example: `7`
`snmp_get_responses`	integer	Number of variable bindings in GetResponse/Response PDUs seen for the session Example: `2`
`snmp_set_requests`	integer	number of variable bindings in SetRequest PDUs seen for the session Example: `10`
`snmp_up_since`	string	Time at which the SNMP responder endpoint claims it's been up since Example: `2024-09-19T00:00:49.536262Z`
`snmp_version`	string	Version of the protocol being used Example: `2c`

Back to Event Fields.

Software fields

A `software` event is generated when software metadata—such as client or server software versions—is detected via protocol-specific exchanges (e.g. DHCP client, HTTP user-agent).

Software events do not have a `src` or `dst` column like all other event types because they only refer to behavior observed from one host and not the underlying connection.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
`host`	ip-object	The host from which the software was observed Example: `10.0.0.10`
`software_name`	string	The name of the observed software Example: `Wget`
`software_type`	string	The category of the observed software Example: `HTTP::BROWSER`
`software_version.additional`	string	Arbitrary notes about the software Example: `linux-gnu`
`software_version.major`	int	The major version number Example: `1`

Field	Type	Description
`software_version.minor`	int	The first minor version number Example: `19`
`software_version.minor2`	int	The second minor version number Example: `1`
`software_version.minor3`	int	The third minor version number Example: `0`
`software_version.version`	string	The full version string Example: `Wget/1.19.1 (linux-gnu)`
`software_version.version_number`	string	The full version number Example: `1.19.1`

Back to Event Fields.

SSH fields

An `ssh` event is created when SSH connection metadata or authentication results—like client/server version strings or auth success/failure—are captured. SSH provides secure remote shell and file transfer capabilities.

Authentication cannot be accurately determined because the necessary data is encapsulated within the encrypted tunnel. Therefore, the `auth_success` field contains a "best-guess" based on available data.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
`auth_success`	Boolean	The inferred authentication result Example: `True`
`cipher_alg`	string	The encryption algorithm used Example: `aes128-ctr`
`client`	string	The client version string Example: `SSH-2.0-OpenSSH_7.6`
`compression_alg`	string	The compression algorithm used Example: `none`
`direction`	string	The direction of the connection, `Outbound` if the client was a local host logging into an external host and `Inbound` in the opposite situation Example: `Inbound`

Field	Type	Description
`hassh`	string	Network fingerprinting which can be used to identify specific Client and Server SSH implementations. Example: `98ddc5604ef6a1006a2b49a58759fbe6`
`hassh_server`	string	Network fingerprinting which can be used to identify specific Server SSH implementations. Example: `cd77a550c195f0e4ea637e367fa499e8`
`host_key`	string	The server fingerprint Example: `a1:a2:79:80:6d:b1:77:82:d8:6c:aa:ee:25:19:23:42`
`host_key_alg`	string	The server's key algorithm. Example: `ssh-rsa`
`kex_alg`	string	The key exchange algorithm used Example: `ecdh-sha2-nistp256`
`mac_alg`	string	The signing (MAC) algorithm used Example: `hmac-sha1`
`server`	string	The server version string Example: `SSH-2.0-OpenSSH_7.4`
`ssh_version`	int	The SSH major version (1 or 2) Example: `2`

Back to Event Fields.

SSL fields

An `ssl` event is generated when secure session negotiations are observed, logging details like cipher suite, certificate chain, server name, and session resume status. It provides insight about encrypted communications by parsing and logging the connection's metadata.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
`cipher`	string	The cipher suite selected by the server Example: `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`
`client_issuer`	string	The Issuer field of the client's certificate Example: `CN=Google Internet Authority G2,O=Google Inc,C=US`
`client_subject`	string	The Subject field of the client's certificate Example: `CN=*.google.com,O=Google Inc`

Field	Type	Description
`issuer`	string	The Issuer field of the server's certificate Example: `CN=Google Internet Authority G2,O=Google Inc,C=US`
`ja3`	string	The computed JA3 hash for the client Example: `4d7a28d6f2263ed61de88ca66eb011e3`
`ja3s`	string	The computed JA3 hash of the server Example: `4d7a28d6f2263ed61de88ca66eb011e3`
`ja4`	*string*	The computed JA4 hash for the client hello packet Example: `t13d1516h2_acb858a92679_e5627efa2ab1`
`server_name_indication`	domain-object	The enriched Server Name Indication set by the client Example: `www.google.com`
`session_id`	string	The ID used for session resumption (deprecated) Example: N/A
`subject`	string	The Subject field of the server's certificate Example: `CN=*.google.com,O=Google Inc`
`validation_status`	string	Result of certificate validation for this connection (deprecated) Example: `Success`
`version`	string	The SSL/TLS version being used (period omitted) Example: `TLSv10`

Back to Event Fields.

Suricata fields

A `suricata` event is created when Suricata (an intrusion detection tool) alerts or metadata are integrated into Zeek logs, highlighting threat detection signatures and behaviors.

Suricata runs independently from the metadata extraction process, and thus is not tied to `flow` events with a `flow_id` even though both a `suricata` and `flow` event will exist for the traffic. Additionally, directionality is not maintained by Suricata, so the `src.ip` and `dst.ip` fields for a `suricata` event may be reversed from the related `flow`.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
community_id	string	An additional flow identifier for joining Suricata and Conn events.
payload	byte-array	Payloads are generated by the sensor's IDS engine. This field displays the raw payload from traffic that matched a detection signature. This ASCII representation helps you determine whether the traffic is malicious or benign. Payloads are disabled by default due to the potential exposure of sensitive or personally identifiable information (PII). When enabled, you can click the field to view the payload in FortiNDR Cloud. Payloads can be enabled upon request through Fortinet Support.
proto	string	The transport layer protocol used. Example: `tcp`
sig_category	string	The query's category. Example: `A Network Trojan was Detected`
sig_id	int	The query's ID. Example: `2024290`
sig_name	string	The query's name. Example: `ET TROJAN Jaff Ransomware Checkin M1`
sig_rev	float	The query's revision number. Example: `2`
sig_severity	int	The query's severity rating (1 = high, 3 = low) Example: `1`

Back to Event Fields.

Tunnel fields

A `tunnel` event is generated when tunneled sessions—such as VPN, SSH tunnels, or other encapsulations—are detected, noting tunnel types and actions. This event helps trace encapsulated traffic flows.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
`tunnel_action`	string	The action taken on the tunnel Example: `Tunnel::DISCOVER`
`tunnel_type`	string	The protocol/application running over the tunnel Example: `Tunnel::HTTP`

Back to Event Fields.

VPC Flow fields

A VPC Flow fields event occurs when raw VPC Flow Log data is parsed and its individual fields are extracted and normalized into a structured event. These events are only visible when the VPC feature is enabled. To enable it, contact your TSM or Customer Support.

The following table shows the fields contained in this event type excluding the previously identified common fields:

Field	Type	Description
`proto`	string	Protocol used in the traffic. Example: TCP
`switched`	boolean	If the source and destination IPs are switched due to port values. Example: false
`tag`	string	The type of event. Example: flow
`total_ip_bytes`	integer	The number of bytes transferred during the flow. Example: 76
`vpc_account_id`	string	AWS account ID owning the source network interface. Example: 123456789101
`vpc_action`	string	The action associated with the traffic. Example: ACCEPT
`vpc_availability_zone`	string	Availability Zone ID of the network interface. Example: usw2-az1
`vpc_dst_ip`	null or ip_enriched object	The preserved original `dstaddr` field when dst.ip was overridden with `pkt-dstaddr`.
`vpc_end_timestamp`	string	The time when the last packet was received in the aggregation interval. Example: 2019-01-01T00:00:00.000000Z
`vpc_flow_direction`	string	The direction of the flow relative to the interface. Example: ingress
`vpc_id`	string	ID of the VPC containing the network interface. Example: vpc-123f7d9bb71e45e11
`vpc_instance_id`	string	ID of the associated instance. Example: i-123b3953f10184bde
`vpc_interface_id`	string	ID of the network interface.

Field	Type	Description
		Example: eni-0ff7168c44159f431
`vpc_ip_version`	string	Type of traffic IP version. Example: IPv4
`vpc_log_status`	string	Logging status of the flow log. Example: OK
`vpc_pkt_dst_ip`	null or ip_enriched object	Packet-level original destination IP address. Example: 10.1.0.1
`vpc_pkt_dst_subnet_name`	string	Subnet name for packet destination IP. Example: AMAZON
`vpc_pkt_src_ip`	null or ip_enriched object	Packet-level original source IP address. Example: 10.1.0.2
`vpc_pkt_src_subnet_name`	string	Subnet name for packet source IP. Example: S3
`vpc_proto`	integer	IANA protocol number of the traffic. Example: 6
`vpc_region`	string	AWS region containing the network interface. Example: us-west-2
`vpc_reject_reason`	string	Reason the traffic was rejected. Example: BPA
`vpc_src_ip`	null or ip_enriched object	The preserved original srcaddr field when `src.ip` was overridden with `pkt-srcaddr`
`vpc_subnet_id`	string	ID of the subnet containing the interface. Example: subnet-12356986a7885a583
`vpc_tcp_flags`	integer	Bitmask value for TCP flags. Example: 2
`vpc_total_pkts`	integer	Number of packets transferred during the flow. Example: 1
`vpc_version`	integer	VPC Flow Logs version. Example: 8

x509 fields

An `x509` event is created when X.509 certificates exchanged in TLS/SSL sessions are parsed and logged, capturing certificate metadata, fingerprints, extensions, and alternate names.

The following table shows the fields contained in this event type excluding the previously identified common fields::

Field	Type	Description
`ca_constraints`	Boolean	Indicates whether the CA flag is set Example: `False`
`ca_constraints_len`	int	The maximum path length Example: `10`
`cert_id`	string	The file ID of the certificate Example: `FNbDqq2ZxjNk10D7ie`
`issuer`	string	The content of the Issuer field Example: `O=Internet Widgits Pty Ltd,ST=Some-State,C=AU`
`key_len`	int	The length of the key Example: `2048`
`key_type`	string	The type of key used Example: `rsa`
`san_dns`	host-array	The list of DNS entries in the SAN Example: `[*.outlook.com, *.office365.com]`
`san_email`	email-array	The list of email entries in the SAN Example: `[dave@email.corp]`
`san_ip`	ip-array	The list of IP entries in the SAN Example: `[169.254.1.1]`
`san_uri`	uri-array	The list of URI entries in the SAN Example: `[https://169.254.1.1]`
`serial`	string	The serial number of the certificate Example: `E3BD4F4F884EADDA`
`subject`	string	The content of the Subject field Example: `O=Internet Widgits Pty Ltd,ST=Some-State,C=AU`
`valid_end`	timestamp	The time before the certificate became valid Example: `2018-01-11T14:35:34.000Z`
`valid_start`	timestamp	The time once the certificate becomes invalid Example: `2018-01-11T14:35:34.000Z`
`version`	string	The X.509 version Example: `3`

Back to Event Fields.

IQL operators

The following operators are supported in IQL.

- Comparison operators on page 263
- Logical operators on page 264
- Exclude operators on page 264
- Pattern operators on page 265
- Units on page 265
- Supported units on page 265
- Fields with units on page 266

Comparison operators

Comparison operators are used to compare fields to values. The following comparison operators are supported by IQL.

Operator	Description	Example
`=`, `==`	Equals	`ip = 8.8.8.8`
`!=`, `<>`	Does not equal	`ip != 8.8.8.8`
`IN`	Set/list operator - the field matches any of the listed values	`ip IN (8.8.8.8, 8.8.4.4)`
`>`	Greater than	`ip_bytes > 100`
`<`	Less than	`ip_bytes < 100`
`>=`	Greater than or equal to	`ip_bytes >= 100`
`<=`	Less than or equal to	`ip_bytes <= 100`

Most comparison operators are standard and intuitive. However, the `IN` operator has two behaviors worth mentioning:

- The values in the list must all be of the same type
- The values in the list will all be treated as exact matches
  - Fuzzy matches in lists are not supported

Also, the absence of a property can be tested by comparing the desired field to the `null` keyword.

```
// Returns HTTP requests that did not receive a response
```

```
http:status_code == null
```

Logical operators

Logical operators are used to chain clauses together to form a more complex query.

Operator	Description	Example
`AND`	Both clauses must be satisfied	`ip = 8.8.8.8 AND port = 53`
`OR`	Only one clause must be satisfied	`ip = 8.8.8.8 OR port = 53`
`NOT`	The inverse must be true (applied to other operators)	`ip NOT IN (10.0.0.10, 8.8.8.8)`

Logical operators allow chaining of multiple clauses. However, in the case of `AND`, all field comparisons must apply, which means all event-types involved must support all fields referenced. For example, the following query is illegal because `flow` events don't have a `qtype_name` field and `dns` events don't have a `service` field. In other words, no single event can have both a `flow`-specific field and a `dns`-specific field.

```
// invalid no single event can be both FLOW and DNS
```

```
dns:qtype_name = 'A' AND flow:service = 'dns'
```

The above example does not apply to the `OR` operator because a single event could be either a `dns` event or a `flow` event.

```
// This is ok, because a single event could match just one clause
```

```
dns:qtype_name = 'A' OR flow:service = 'dns'
```

Exclude operators

The `exclude` operator, for example, A exclude B, provides relative complement filtering that allows all items matching a criteria to be excluded from the result set.

For example, `event_type = 'flow' and ip != 10.30.0.3` may return an event with `src.ip = 10.30.0.1` and `dst.ip = 10.30.0.3` because `src.ip` satisfies the constraint that the event has an ip field that is not `10.30.0.3`. This may not be the desired intention. In comparison, `event_type = 'flow' exclude ip = 10.30.0.3` would not return the event previously described. It will only return flow events excluding those events that match `ip = 10.30.0.3`.

Syntax:

The exclude operator is a low precedence, infix operator with left associativity. For example, with A, B, and X below representing complex expressions:

- A exclude X ## base example of matching everything in A except what matches X
- A and B exclude X ## this is the same as (A and B) exclude X
- A or B exclude X ## this is the same as (A or B) exclude X
- A exclude X and Y ## this is the same as A exclude (X and Y)
- A exclude X or Y ## this is the same as A exclude (X or Y)

- A exclude X exclude Y ## this is the same as (A exclude X) exclude Y which is the same as A exclude (X or Y)
- (A exclude X) and (B exclude Y) ## example of using exclude in a restricted context
- exclude X ## This is a special case and interpreted as * exclude X

Pattern operators

Pattern operators allow you to identify strings that contain certain patterns. The `LIKE` operator provides simple fuzzy matching, while the `MATCHES` operator provides access to Regex for more complex pattern matching.

Operator	Description	Example
`LIKE`	The `LIKE` operator supports simple pattern matching using `%` (any number of characters) and _ (a single character).	`domain NOT LIKE "%.google.com"`
`MATCHES`	Regex matching	`domain MATCHES ".*\.(com|net|org|edu)"`

Strings must be provided to pattern operators, meaning the characters must be surrounded by quotes. For the `LIKE` operator, the exact string will be matched if no wildcards exist in the provided string.

Units

IQL supports units for several numeric fields. Units are optional but can greatly increase readability of queries that use time, size, or distance values. Here are some examples:

```
dst.ip_bytes > 5MB  // will convert 5MB to 5242880 bytes
```

```
dst.ip_bytes > 5.5mb  // will convert 5.5mb to 5767168 bytes
```

Unit labels are case insensitive.

Supported units

Name	Type	IQL Label
`bytes`	*size*	b
`kilobytes`	*size*	kb
`megabytes`	*size*	mb
`gigabytes`	*size*	gb

Name	Type	IQL Label
`terabytes`	*size*	tb
`petabytes`	*size*	pb
`miles`	*distance*	mi
`kilometers`	*distance*	km
`nanoseconds`	*time*	ns
`microseconds`	*time*	us
`miliseconds`	*time*	ms
`seconds`	*time*	s
`minutes`	*time*	m
`hours`	*time*	h
`days`	*time*	d

Fields with units

Fields	Units
`geo_distance`	miles
`lease_duration`	seconds
`ip_bytes`	bytes
`duration`	seconds
`total_ip_bytes`	bytes
`request_len`	bytes
`file.bytes`	bytes

Advanced Query Concepts

Putting it all together

The following query searches for outbound traffic from an internal network to external destinations. It also looks for traffic that is going through a proxy server that acts as an intermediary between a device and the internet.

```
// Outbound traffic
(
   http:src.internal = true
```

```
   OR http:source IN ("Zscaler")
)
AND (
   dst.internal = false
   OR (
         // Not internal IP address
         host.internal != true
         // Proxied traffic
         AND uri.scheme != null
   )
)
```

Outbound Traffic	The query is looking for traffic that is leaving the internal network.
Conditions for Source	• `http:src.internal = true`: The source of the traffic is internal. • `http:source IN ("Zscaler")`: The source is from Zscaler, a cloud security company.
Conditions for Destination	• `dst.internal = false`: The destination is external (not internal). • `host.internal != true`: The host is not internal.

Array matching

Array matching

The following table provides example queries for array matching where `answers.ip` is the array field:

Show me events where at least one answer value is:	Query
8.8.8.8	`answers.ip = 8.8.8.8`
not 8.8.8.8	`answers.ip != 8.8.8.8`
8.8.8.8 and one is not 8.8.8.8	`answers.ip = 8.8.8.8` `AND answers.ip != 8.8.8.8`

Excluding values in DNS queries

The `!=` operator will not exclude values in a DNS query (*answers.ip != 8.8.8.8*). Instead, you will need to use the EXCLUDE condition:

```
event_type = "dns"
EXCLUDE answers.ip = 8.8.8.8
```

Using Curly Braces for multiple Conditions

Curly braces `{}` are used to group multiple conditions together in an array of objects, such as `intel` and `files`. This helps to specify detailed criteria for your queries.

Format:

```
<array of objects field> {
    <subfield> <operator> <value>
    …
}
```

Examples:

In the following example, the query will return results for both confidence and severity.

High-Confidence and High-Severity Intel Matches	
Query	Show me events with a high-confidence intel match and a high-severity intel match.
Syntax	`intel.confidence = "high"` `AND intel.severity = "high"`

In the following example, the curly braces `{}` help to group the conditions together, making it clear that both conditions must be met within the same intel object.

High-Confidence and High-Severity Intel Matches	
Query	Show me events with a high-confidence intel match and a high-severity intel match.

Syntax	`intel {`
	`  confidence = "high"`
	`  AND severity = "high"`
	`}`

Aggregations

An aggregation is achieved by adding GROUP BY at the end of the query, this allows for summarizing data, typically resulting in event counts by default. The portal provides both visual and tabular representations of these results.

Aggregations can include up to two unambiguous fields, with default limits of 100 and 10, respectively, which can be adjusted but must not exceed a product of 10,000. High-entropy fields like *uuid* and *flow_id* cannot be used. Functions such as SUM, MIN, and MAX are available.

Example:

The following query identifies Kerberos authentication errors where the client's credentials have been revoked and groups the results by the client name.

```
kerberos:error_msg = 'KDC_ERR_CLIENT_REVOKED' AND client NOT LIKE '%$/%.%' GROUP BY client LIMIT 10
```

Kerberos Error Message	`kerberos:error_msg = 'KDC_ERR_CLIENT_REVOKED'`: This looks for Kerberos errors with the specific message `KDC_ERR_CLIENT_REVOK`.

	This error indicates that the client's credentials have been revoked by the Key Distribution Center (KDC).
Client Filtering	`client NOT LIKE '%$/%.%'`: This condition filters out clients that have a dollar sign ($) or a period (.) in their name. Typically, in Kerberos, machine accounts end with a dollar sign, so this filter is excluding machine accounts.
Grouping and Limiting	`GROUP BY client`: This groups the results by the client name. `LIMIT 10`: This limits the output to the top 10 results.

The query results will look like this:

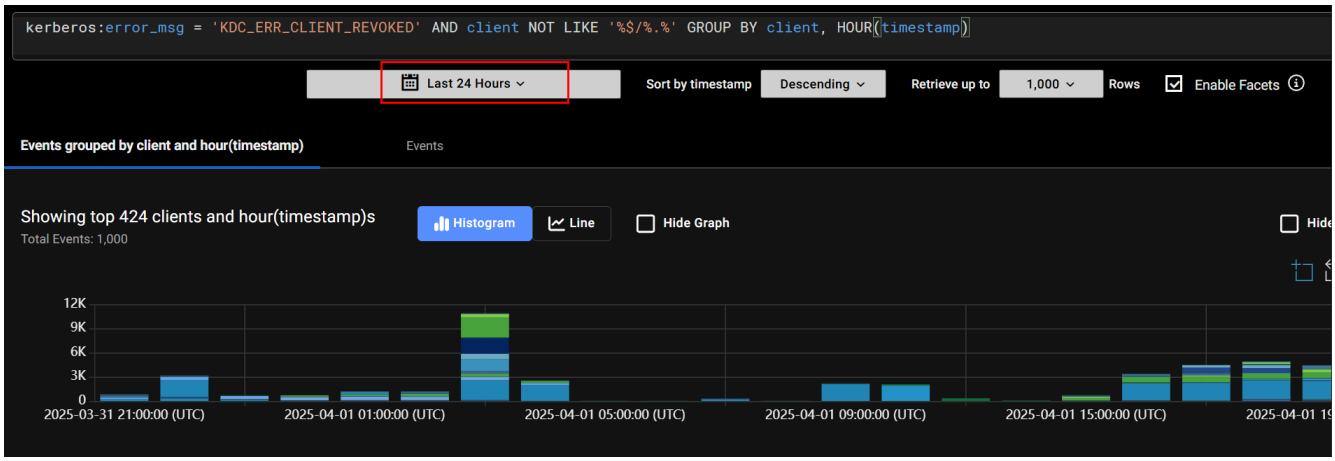

De Morgan's Law

You cannot use the NOT operator to negate a group of clauses directly. This means you cannot write a query like:

```
NOT (
    dst.ip = 8.8.8.8
    AND host = "dns.google.com"
)
```

Instead, you need to apply the NOT operator to each clause individually and then combine them using the OR operator. The equivalent query would be:

```
NOT dst.ip = 8.8.8.8
OR NOT host = "dns.google.com"
```

This way, the query will return results where either the `dst.ip` is not `8.8.8.8` or the `host` is not `dns.google.com`. This ensures that at least one of the conditions is not met.

De Morgan's Law does not apply to array fields. An array field is an array of values as opposed to a number or a string. The `answers.ip` field is an example is an example of an array.
In the following example, the two conditions can both be true at the same time:

```
answers.ip = 8.8.8.8
answers.ip != 8.8.8.8
```

Field reference

This section describes how to use fields including where flexibility exists and the implications of that flexibility.

- Schema and field references on page 270
- Event-type expansion on page 270
- Field expansion on page 271
- Synthetic fields on page 272

Schema and field references

Queries are evaluated against the events datastore. Every event type has a set of properties – we refer to them as **fields** – that carry data of a defined primitive type. For instance, every event has a `sensor_id` property that is of type `string` and a `timestamp` property of type `timestamp`. The full schema for all available event types and their properties is available within the Event Types page.

All queries consist fundamentally of matching an event field against a value; for instance, "Show me all events for which the destination IP is 8.8.8.8." However, there is some room for flexibility. Do you really want *all* event types, or is there one in particular you're interested. Do you really want to restrict results to cases where 8.8.8.8 is the *destination* IP address, or would any involvement of that IP address be interesting?

Each field involved in a query must be resolved to a specific field of a specific event type. A fully-specified field is of the format `event-type:field`; for instance, `flow:sensor_id` and `dns:dst.geo.country` are both fully specified. For a field that's not fully specified, either by omitting the event type or part of the field, the system will expand the field to include all fully-qualified fields that fit the ambiguity.

The next two subsections will show how these expansions work and what their implications are.

Event-type expansion

A field without a specifed event type will infer all valid event types. For example, `dns` and `flow` events both have a `proto` field, so a query containing just `proto` without an event-type prefix will expand to include both event types. Effectively, the query on the first line below is rewritten by the query engine on the backend to the query on the second line.

```
// original query
```

```
proto = 'udp'
```

```
// rewrite produced by the query engine on the backend
```

```
dns:proto = 'udp' OR flow:proto = 'udp'
```

If a field only belongs to one event type, then the event type does not need to be specified since the results would be the same. For example, the `qtype_name` field is unique to the `dns` event type, so only one event type can be inferred. This means that the two queries below are equivalent.

```
// original query
```

```
qtype_name = 'A'
```

```
// the rewrite is equivalent
```

```
dns:qtype_name = 'A'
```

Field expansion

Some fields hold values of a structural type (Event Type and Fields), meaning they contain subfields that must be referenced. To make this clear, let's use the `src` field as an example. The `src` field is of the type *ip-object*, i.e. a JSON structure. Looking at the following code block, we couldn't compare `src` to an IP address because we'd have to specify the entire JSON structure for them to match on structure. Instead, we must compare the `ip` subfield to an IP address.

```
// invalid because src is type ip-object and we're comparing it to an ip
```

```
src = 10.0.0.10
```

```
// valid because src.ip is type ip and we're comparing it to an ip
```

```
src.ip = 10.0.0.10
```

If a subfield is used without the parent field, the query will be expanded to include all valid parent fields. For instance, the subfield `ip` could expand to `dst.ip`, `src.ip`, and a number of others. The block below shows the complete expansion for the `ip` field in a `dns` event.

```
// original query
```

```
dns:ip = 10.0.0.10
```

```
// rewritten to expand the unspecified parent field
```

```
dns:src.ip = 10.0.0.10 OR dns:dst.ip = 10.0.0.10 OR dns:answers.ip = 10.0.0.10
```

Event-type and field expansion can be applied to the same query. For example, if we simply specified the `ip` field, the query engine would expand to all possible parent fields in all possible event types.

```
// original query
```

```
ip = 10.0.0.10
```

```
// complete expansion of event type and parent field (truncated)
```

```
dns:src.ip = '10.0.0.10' OR dns:dst.ip = '10.0.0.10' OR dns:asnwers.ip = 10.0.0.10 OR flow:src.ip
= '10.0.0.10' OR flow:dst.ip = '10.0.0.10'
```

Synthetic fields

A **synthetic field** is a field that doesn't exist in an event record, i.e. it isn't static. Synthetic fields are dynamically evalutated and converted into static values before your IQL query is run against the event data store. This enables more robust capabilities that aren't possible with a simple query of static values.

Synthetic fields begin with a `$`. The example query below demonstrates the `$device` synthetic field, which enables a user to search for a source or destination device by hostname or MAC address instead of just the observed IP address. The hostname is evaluated behind the scenes to produce a large array of IP addresses and valid time ranges, which are then used to query the event data store.

```
src.$device.hostname = 'FinanceWks008' and dst.internal = false
```

IQL query examples

Insight Query Language Basics

Query	IQL

Filter by event type	`<event-type>:<field> <operator> <value>:`
	
	`dns:dst.ip = 8.8.8.8`

Add an aggregation	`<clause> group by <aggregation field>:`
	
	`dns:dst.ip = 8.8.8.8 and dst.port = 53 group by query.domain`

Add a clause	`<clause> <structural operator> <clause>:`
	
	`dns:dst.ip = 8.8.8.8 and dst.port = 53`

DCE-RPC Examples

Query	IQL
Search for RPC service creation	`dce_rpc:dce_rpc_operation like '%CreateService%'`
Search for RPC scheduled task registration	`dce_rpc:dce_rpc_operation = 'SchRpcRegisterTask'`

DNS Examples

Query	IQL
Search for long DNS queries	`dst.internal = false and dns:query.domain matches '[ 0-9a-zA-Z\.\-]{75,}' group by query.domain`
Search for long DNS txt records	`dns:qtype_name = 'TXT' and dns:query.domain matches '.{100,}' and dns:answers matches '.{100,}'`

HTTP Examples

Query	IQL
Search for direct-to-IP HTTP post	`http:host.ip != null and method = 'POST' and dst.internal = false`
Search for deprecated Windows versions	`src.internal=true and http:user_agent matches '.*Windows (XP|2000|2003|NT [4,5]).*'`

Flow Examples

Query	IQL
Search for top outbound services by data sent	`src.internal = true and dst.internal = false and flow:service != null group by service, sum(total_ip_bytes)`
Search for outbound connections using administrative protocols	`src.internal = true and dst.internal=false and flow:service in ("ftp","ssh","rdp") group by service, dst.asn.asn_org`

FTP Examples

Query	IQL
Search for executable files over FTP	`ftp_arg matches '.*[Ee][Xx][Ee]'`

Kerberos Examples

Query	IQL
Search for revoked Kerberos login attempts	`kerberos:error_msg = 'KDC_ERR_CLIENT_REVOKED'`

NTLM Examples

Query	IQL
Search for admin NTLM user accounts	`ntlm:username matches '.*[Aa]dmin.*' group by username, src.ip`
Search for revoked NTLM user accounts	`ntlm_status in ('ACCOUNT_DISABLED', 'ACCOUNT_EXPIRED', 'ACCOUNT_LOCKED_OUT', 'ACCOUNT_RESTRICTION', 'INVALID_WORKSTATION', 'NO_SUCH_USER')`

PE Examples

Query	IQL
Search for recently compiled, non-GUI executable files	`pe:compile_timestamp > t'2020-01-01T00:00:00.000Z' and subsystem != 'WINDOWS_GUI' group by file.name`

RDP Examples

Query	IQL
Search for unencrypted RDP traffic	`rdp:src.internal=true and result!='Success' and result != 'encrypted' group by dst.ip`

SMB Examples

Query	IQL
Search for SMB access to temporary paths	`smb_file:files.smb_path.path matches '.*[Tt][Ee][Mm][Pp].*'`
Search for SMB access to filenames containing "password"	`smb_file:files.name matches ".*\\[Pp][Aa][Ss][Ss][Ww][Oo][Rr][Dd][^\\]+\.[a-zA-Z]{2,4}" group by files.name`
Search for hosts with accessible C$ shares	`smb_mapping:smb_path.share = 'C$' group by smb_path.system`

SMTP Examples

Query	IQL
Search for SMTP mail servers	`smtp:src.internal = true and dst.internal = false group by src.ip`

SSH Examples

Query	IQL
Outbound SSH to rare or unknown SSH server versions	`ssh:auth_success = true and dst.internal = false and src.internal = true group by server`

SSL Examples

Query	IQL
Search for deprecated SSL versions	`ssl:version in ('SSLv2', 'SSLv3', 'TLSv10', 'TLSv11')`
Search for self-signed SSL certificates	`ssl:issuer matches '.*[Ll][Oo][Cc][Aa][Ll][Hh][Oo][Ss][Tt].*'`

X509 Examples

Query	IQL
Search for expired X.509 certificates	`x509:valid_end < t'2020-03-03T00:00:00.000Z'`

Natural Language queries

Natural Language (NL) Queries allow analysts to use simple statements for investigations and private searches as an alternative to Internal Query Language (IQL). You can start an NL Query from the *Investigations* or *Private Search* pages. NL queries appear with regular queries, and their results appear the same. In the query history, `NL Query` is shown next to the query name (for example: `NL Query - 2025-11-18 06:04 (UTC)`).

The *Explain* button next to the query results displays how the system interpreted your request and the queries it executed. You can copy this explanation and edit it to refine your queries. The Explain button is also available in *Private Search* and the *Visualizer*.

- NL queries are enabled by default. To disable NL Queries, contact your TSM.
- Users can make up to 300 NL queries per day for the account they belong to. This limit is subject to change.

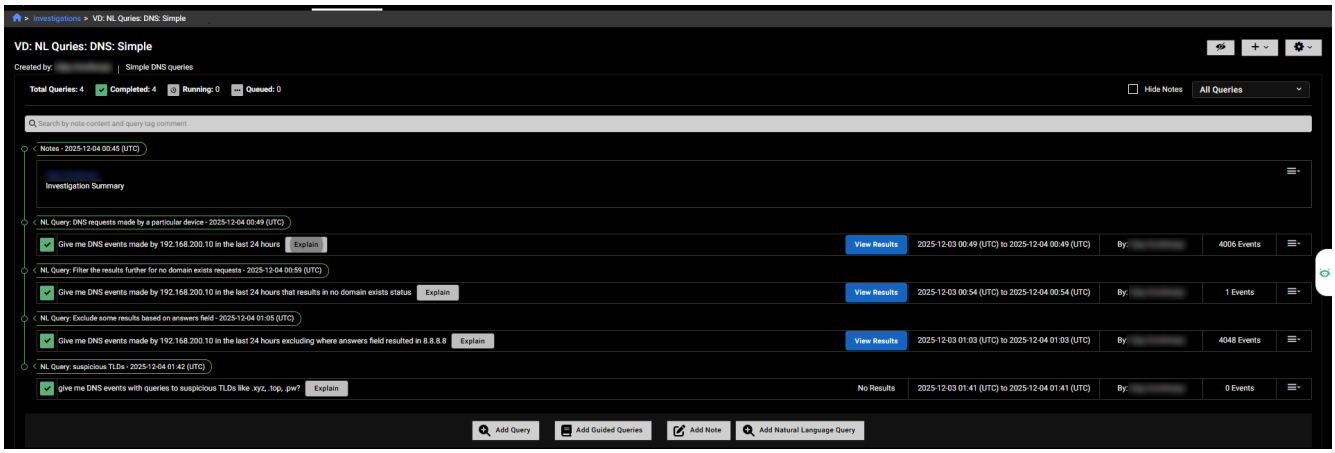

Supported event types and languages

NL Queries support the following event types and languages at this time:

Category	Support
Event Types	DCE RPC, DHCP, DNS, DPI, Flow, flow_state, FTP, HTTP, Kerberos, NTLM, SSH, SSL, Suricata, x509
Languages	Arabic, Chinese, Croatian, Czech, Dutch, English, Finnish, French, German, Greek, Hebrew, Hindi, Italian, Japanese, Korean, Norwegian, Polish, Portuguese, Romanian, Russian, Spanish, Swedish, Turkish

Feature Constraints

- The *Actions* menu and *Facets* options are disabled for NL queries.
- NL Query text does not appear in *Global Search* results.
- NL Query results cannot be used to create detectors.
- Querying using annotations or tags is not supported.
- Not all NL aggregation queries generate a chart; some display only a table.
- Raw events are not shown for NL aggregation queries.

Running Natural Language queries

To add a Natural Language query to an investigation:

1. Go to *Investigations* and click an investigation in the list.
2. Click *Add Query*. The *Add a New Query* page opens.
3. Click *Add Natural Language Query*. The *Add a New Natural Language Query* dialog opens.
4. In the *Name* field, enter a name for the query.
5. In the *Natural Language Query* field, type your query.

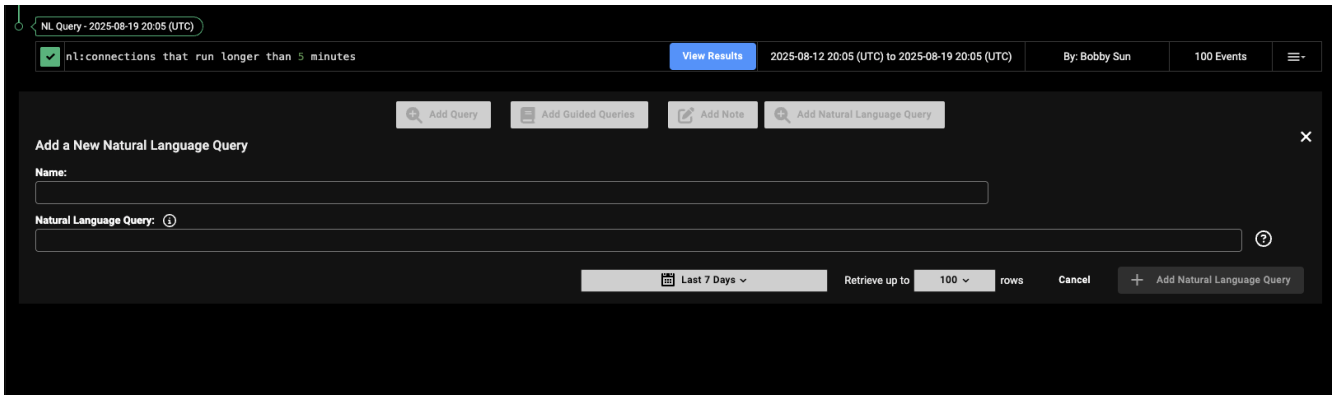

To use a Natural Language queries in Private Search:

1. Go to *Investigations* > *Private Search*.
2. Enable *Natural language Query*.
3. In the *Name* field, enter a name for the query.
4. In the *Natural Language Query* field, type your query.
5. Click *Search*.

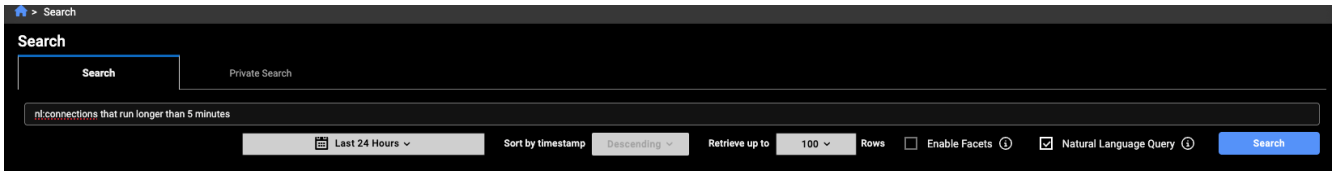

NL Query guidelines

Best Practices

- Use short time ranges for faster queries and fewer timeouts. A week or a month is ideal.
- Limit your queries to supported event types.
- Search within a single account at a time. Natural language queries do not support multi-account searches.
- Use shorter time ranges in queries to avoid exceeding the 5-minute execution limit, which can cause queries to fail.
- Clearly define the desired output format. For example:
  - *Provide the result as group by source IP, issuer.*
  - *Include the summary with counts.*

Example queries

Flow Examples

Query	Natural Language
Search for top outbound services by data sent	Show me flow events with services that transferred most outbound data from internal hosts. Provide the results as group by service name where service is not null and total ip bytes sent.
Search for outbound connections using administrative protocols	Show me flow events with internal hosts connecting to external destinations where service field contains ftp or ssh or rdp. Provide the results as group by service and destination asn org.

DNS Examples

Query	Natural Language
Search for long DNS queries	Show me DNS events with external destinations where the length of query field is more than 75 characters.

Query	Natural Language
Search for long DNS txt records	Show me dns events where query type ='TXT' and length of query and response fields more than 100.
Search for DNS requests made by a device	Show me DNS events made by 192.168.200.10 in the last 24 hours

HTTP Examples

Query	Natural Language
Search for direct-to-IP HTTP post	Show me http events with POST method and host IP field is not null and destination is external.
Search for deprecated Windows versions	Show me events with internal hosts where http user agent indicates a deprecated windows version.

SSL Examples

Query	Natural Language
Search for deprecated SSL versions	Find internal hosts using SSL versions in ('SSLv2', 'SSLv3', 'TLSv10', 'TLSv11'). provide the results as group by source IP and SSL version.
Search for self-signed SSL certificates	Search for self signed SSL certificates. Look for issuer like localhost using a case insensitive search. Provide the result as group by source IP, issuer.

X509 Examples

Query	Natural Language
Search for expired X.509 certificates	Search for expired X.509 certificates. Provide the result as group by source IP, validity end date.

www.fortinet.com

Copyright© 2026 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's Chief Legal Officer, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.