FortiNDR Best Practice Guide

FortiNDR 7.6.0

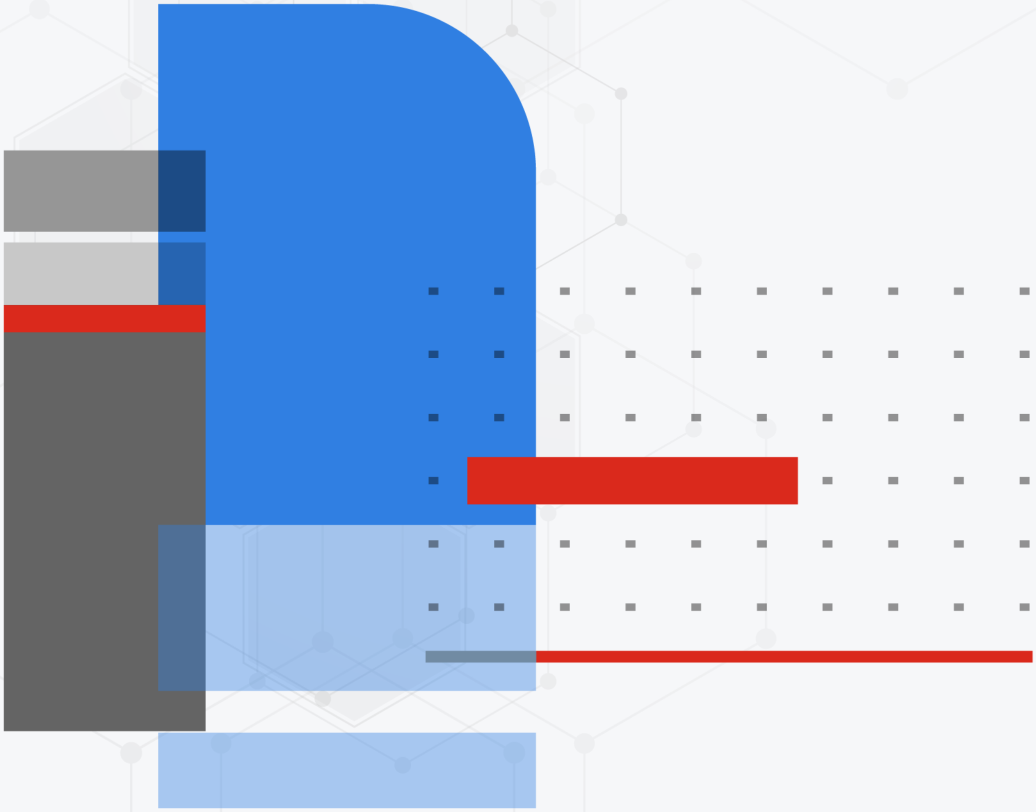

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

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

October 16, 2025
FortiNDR 7.6.0 FortiNDR Best Practice Guide
55-760-1138571-20251016

TABLE OF CONTENTS

Change Log

Date	Change Description
2025-05-30	Initial release.
2025-06-11	Updated Dual Center Management on page 27.
2025-07-17	Added ERSPAN and RSPAN Support on page 34
2025-09-25	Updated Hardening on page 7.
2025-10-16	Added Understanding FortiNDR on-premise machine learning on page 11.
2026-01-12	Added Separating ports for large scale deployments on page 39.

Overview

This guide is a collection of best practices guidelines for using FortiNDR. Use these best practices to help you get the most out of your FortiNDR products, maximize performance, and avoid potential problems.

Hardening

System hardening reduces security risk by eliminating potential attack vectors and shrinking the system's attack surface.

- Register your product with Fortinet Support
- Physical security on page 7
- Vulnerability: PSIRT monitoring on page 7
- Firmware on page 8
- Encrypted protocols on page 8
- Trusted Host
- FortiGuard databases on page 9
- Penetration testing on page 9
- Password policies
- Disable Unnecessary Services
- Configuration backup
- Logging

Physical security

Install the FortiNDR in a physically secure location. Physical access to the FortiNDR can allow it to be bypassed, or other firmware could be loaded after a manual reboot.

Vulnerability: PSIRT monitoring

The Product Security Incident Response Team (PSIRT) continually tests and gathers information about Fortinet hardware and software products, looking for vulnerabilities and weaknesses. The findings are sent to the Fortinet development teams, and serious issues are described, along with protective solutions, in advisories listed at https://www.fortiguard.com/psirt.

Firmware

Keep the FortiNDR firmware up to date. The latest patch release has the most fixed bugs and vulnerabilities and should be the most stable. Firmware is periodically updated to add new features and resolve important issues.

- Read the release notes. The known issues may include issues that affect your business. See, *FortiNDR Release Notes* in the Fortinet Document Library.
- Do not use out-of-support firmware. Review the Product Life Cycle > Software page and plan to upgrade before the FortiNDR End of Support (EOS) date, which is when Fortinet Support services for the firmware version expire.
- Enable *Restrict login to trusted hosts* in the *Administrator* settings to restrict admins to log in using a trusted host. For information, see Administrators in the *FortiNDR Administration Guide*.

Encrypted protocols

Use encrypted protocols whenever possible, for example:

- LDAPS instead of LDAP
- SNMPv3 instead of early SNMP versions
- SSH instead of telnet
- SCP instead of FTP or TFTP

> ⚠
>
> When configuring an LDAP connection to an Active Directory server, an administrator must provide Active Directory user credentials.
>
> - To secure this connection, use LDAPS on both the Active Directory server and FortiGate. See Configuring an LDAP server and Configuring client certificate authentication on the LDAP server.
> - Apply the principle of least privilege. For the LDAP regular bind operation, do not use credentials that provide full administrative access to the Windows server when using credentials. See Configuring least privileges for LDAP admin account authentication in Active Directory.

Trusted Host

For administrative access to FortiNDR please use trusted hosts configuration via CLI. For information, see `config system admin` in the FortiNDR CLI Reference Guide.

FortiGuard databases

Ensure that FortiGuard databases, such as IPS, AV, ANN and other NDR are updated punctually by setting updates to automatic.

Penetration testing

Test your FortiNDR to try to gain unauthorized access or use internal tools or third-party tools and companies to verify FortiNDR access and configuration.

Password policies

Create a secure password policy to ensure user passwords meet the minimum number of characters, numbers, symbols and letters. For information, see config system password-policy.

Disable unnecessary services

To protect FortiNDR from unnecessary exposure, consider disabling the following features when not in use:

- Interface connectivity (ping/snmp/telnet etc.)
- Netflow
  Run CLI: `execute netflow <on/off>`
- For pure malware scanning deployment, NDR daemon can be disabled:

  Run CLI: `execute ndrd <on|off>`
- If the deployment does not require malware scanning by AV/ANN, you can disable sniffer malware detection. Manual submission, HTTP2 and OFTP will still work as file input sources.
  Run CLI: `execute snifferd <on|off>`
- Disable ICAP server configuration if not required. This feature is disabled by default.
  See ICAP Connectors in the FortiNDR Administration Guide.

Configuration backup

The FortiNDR configuration file has important information that should always be kept secured, including details about your network, users, credentials, etc. There are many reasons to back up your configuration, such as disaster recovery, preparing for migrating to another device, and troubleshooting. Evaluate the risk involved if your configurations were exposed and manage your risk accordingly. Store the configuration file in a secure location. Delete old configuration files that are no longer needed.

Logging

Logging generates system event, traffic, user login, and many other types of records that can be used for alerts, analysis, and troubleshooting. The records can be stored locally (data at rest) or remotely (data in motion). Due to the sensitivity of the log data, it is important to encrypt data in motion through the logging transmission channel. When logging into third party devices, make sure that the channel is secure. If it is not secure, it is recommended that you form a VPN to the remote logging device before transmitting logs to it.

Logging options include FortiAnalyzer, Syslog, and a local disk. Logging with Syslog only stores the log messages. Logging to FortiAnalyzer stores the logs and provides log analysis. If a Security Fabric is established, you can create rules to trigger actions based on the logs. For example, sending an email if the FortiNDR configuration is changed, or running a CLI script if a host is compromised.

FortiSIEM (Security Information and Event Management) and FortiSOAR (Security Orchestration, Automation, and Response) both aggregate security data from various sources into alerts and supports logging from FortiNDR.

Understanding FortiNDR on-premise machine learning

This topic describes how FortiNDR On-Premises leverage Machine Learning (ML) and Artificial Intelligence (AI). FortiNDR On-Premises utilizes the following ML and AI techniques:

- **Patented ANN** for real-time malware scanning, classification and tracing sources of malware attacks combined with Fortinet AntiVirus engines.
- **One State Vector Machine** for traffic profiling and anomalies detection.

The following sections explain how FortiNDR On-Premises performs traffic profiling and detects anomalies.

Standalone vs Center deployments

Machine Learning (ML) functions in Standalone, Sensor, and Center deployments. While the core principles are the same across these setups, the deployment details vary slightly.

Center deployment

In a Center deployment, sensors transmit session data in real time to the FortiNDR Center, where Machine Learning (ML) baselining and anomaly detection are performed. The FortiNDR Center offers a centralized view of each sensor's ML configuration and detected anomalies. FortiNDR hardware, such as the FNR-3600G, includes a GPU (Graphics Processing Unit) to accelerate ML processing and configuration tasks.

Dual center deployment

In a Dual Center deployment, Machine Learning is performed independently at each Center. This setup offers flexibility, allowing each center to apply different ML profiles to the same sensor(s) for anomaly detection. This approach is best suited for advanced users who have a deep understanding of their network. For most users, a consistent ML configuration across both Centers is recommended. For more information, see Dual Center Management on page 27.

Understanding the network

In most cases, FortiNDR ML is deployed by one of two types of administrators:

Administrators with limited network knowledge

For administrators who are still learning about their network environment, it is recommended to use the default ML settings, typically a 7-day baseline with the default features enabled. This configuration helps uncover network behavior and detect anomalies. Fine-tuning can be done using the *User Feedback* feature and *NDR Muting*.

For more information about these features, see the User Feedback and Advanced Muting sections of this topic.

Administrators with advanced network insight

Experienced administrators who have a strong understanding of their network’s traffic profile can customize ML settings to suit their needs. This includes setting a specific baseline duration, enabling selected features, and using ML to detect anomalies more precisely. This approach is especially relevant for highly regulated or sensitive environments such as finance, telecommunications, and operational technology (OT) networks.

Understanding baseline configuration

Anomaly detection is only as effective as the quality of its baseline. The baseline can vary depending on the configured machine learning features and the duration over which it has been established. Both of these can be configured in the GUI and the CLI.

Machine Learning features

These features range from network attributes like source and destination ports, TLS version, and VLAN ID, to device-specific details such as IP address, MAC address, model, OS, vendor, category, and geolocation, as well as behavioral indicators like protocol actions and time of week. Collectively, they help establish a baseline for detecting anomalies in network activity.

To configure these features, see ML Configuration in the *FortiNDR Administration Guide*.

Duration of baseline learning

The default baseline learning is seven days. However, this can be customized with the CLI command: `execute reset-ml-baseline-time`. For more information, see `execute reset-ml-baseline-time` in the *FortiNDR CLI Reference Guide*.

Baseline configuration recommendations

Understand your network's normal behavior	To determine the ideal baseline length for your network segment, it's important to first understand what "normal" traffic looks like in that environment.
Configure ML profiles per Center in dual deployment	In Center/Sensor mode, each Center can define its own Machine Learning profile to monitor anomalies independently. For configuration details, refer to Dual Center Management on page 27.

Use separate ML profiles for distinct sniffer interfaces	If FortiNDR monitors multiple networks using different sniffer interfaces, configure separate ML profiles for each. This allows anomaly detection to be tailored to the unique characteristics of each network or subnet.

Deployment examples

Standalone deployment

In this example the administrator has a strong understanding of the network.

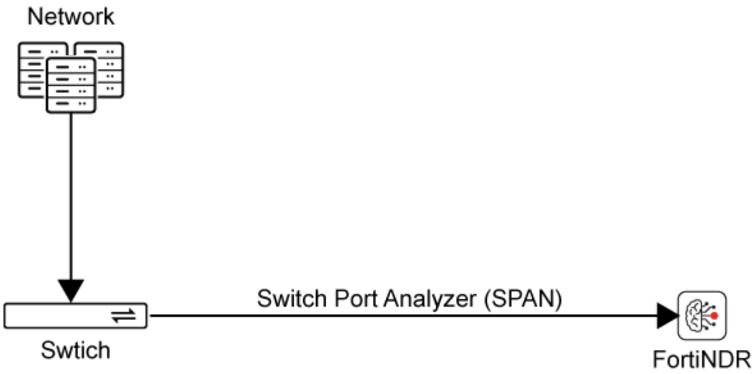

If the server subnet *10.10.1.0/24* only communicates with *192.168.1.0/24*, and other user subnets are primarily used for internet browsing, with interest focused only on detecting new applications and geographic anomalies, then machine learning profiles can be configured accordingly for each segment as follows:

The default Machine Learning profile is applied to the user’s subnet:

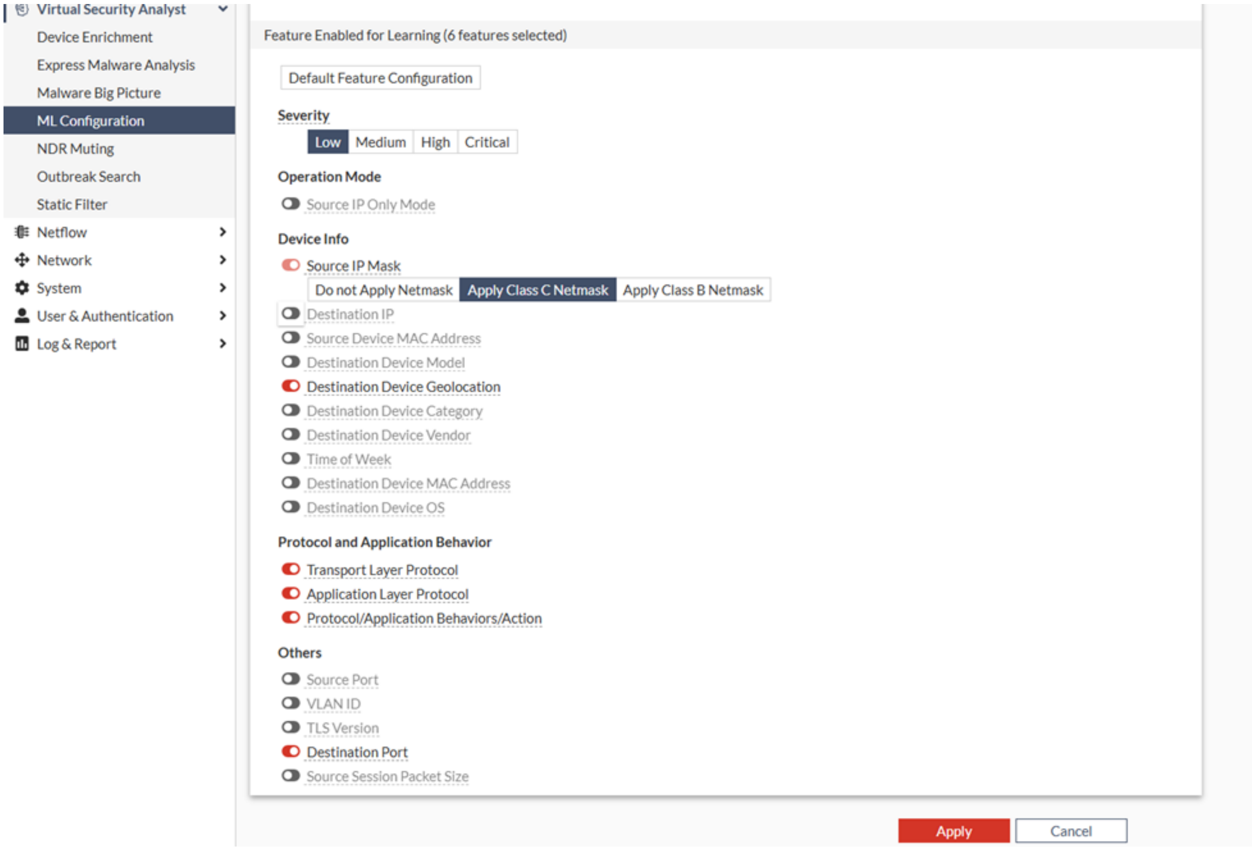

Feature Enabled for Learning (6 features selected)

Default Feature Configuration

Severity

[✓ Low] Medium High Critical

Operation Mode

○ Source IP Only Mode

Device Info

● Source IP Mask

Do not Apply Netmask [✓ Apply Class C Netmask] Apply Class B Netmask

○ Destination IP
○ Source Device MAC Address
○ Destination Device Model
● Destination Device Geolocation
○ Destination Device Category
○ Destination Device Vendor
○ Time of Week
○ Destination Device MAC Address
○ Destination Device OS

Protocol and Application Behavior

● Transport Layer Protocol
● Application Layer Protocol
● Protocol/Application Behaviors/Action

Others

○ Source Port
○ VLAN ID
○ TLS Version
● Destination Port
○ Source Session Packet Size

The *Source IP* is for the servers:

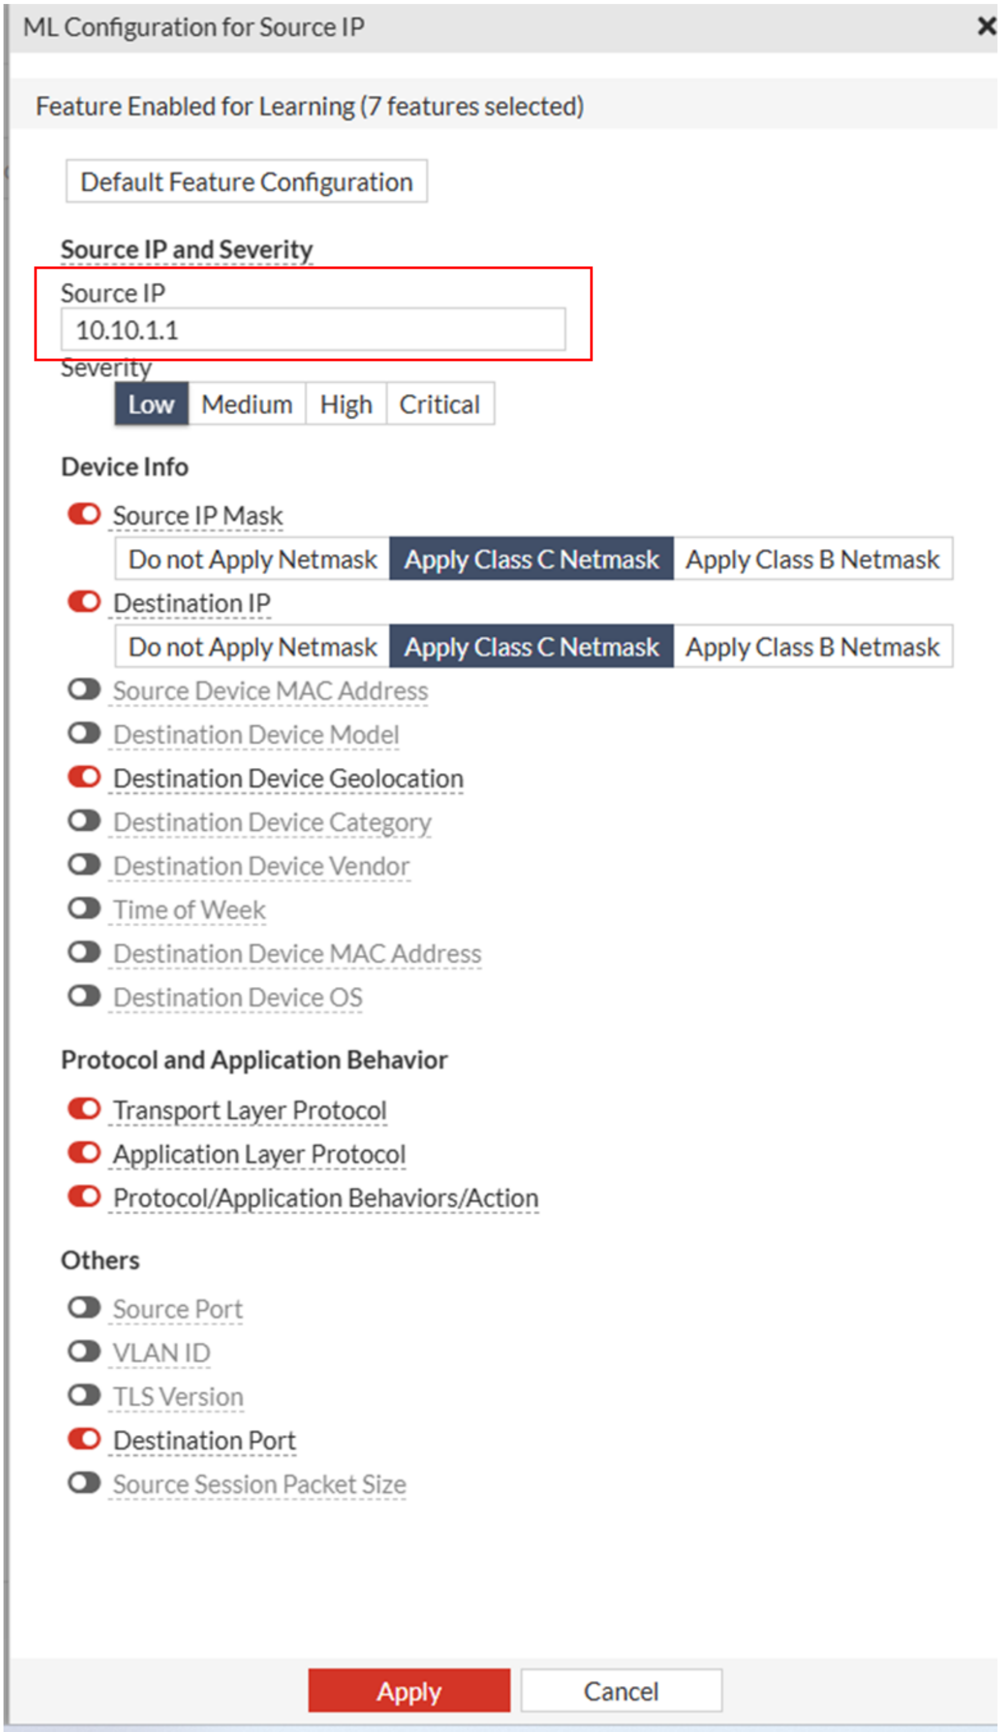

ML Configuration for Source IP ✕

Feature Enabled for Learning (7 features selected)

Source IP and Severity

Source IP: 10.10.1.1

Severity: [✓ Low] Medium High Critical

Device Info

● Source IP Mask: Do not Apply Netmask [✓ Apply Class C Netmask] Apply Class B Netmask

● Destination IP: Do not Apply Netmask [✓ Apply Class C Netmask] Apply Class B Netmask

○ Source Device MAC Address

○ Destination Device Model

● Destination Device Geolocation

○ Destination Device Category

○ Destination Device Vendor

○ Time of Week

○ Destination Device MAC Address

○ Destination Device OS

Protocol and Application Behavior

● Transport Layer Protocol

● Application Layer Protocol

● Protocol/Application Behaviors/Action

Others

○ Source Port

○ VLAN ID

○ TLS Version

● Destination Port

○ Source Session Packet Size

If the server's subnet is sensitive to the traffic time, the *Time of Week* option can be selected.

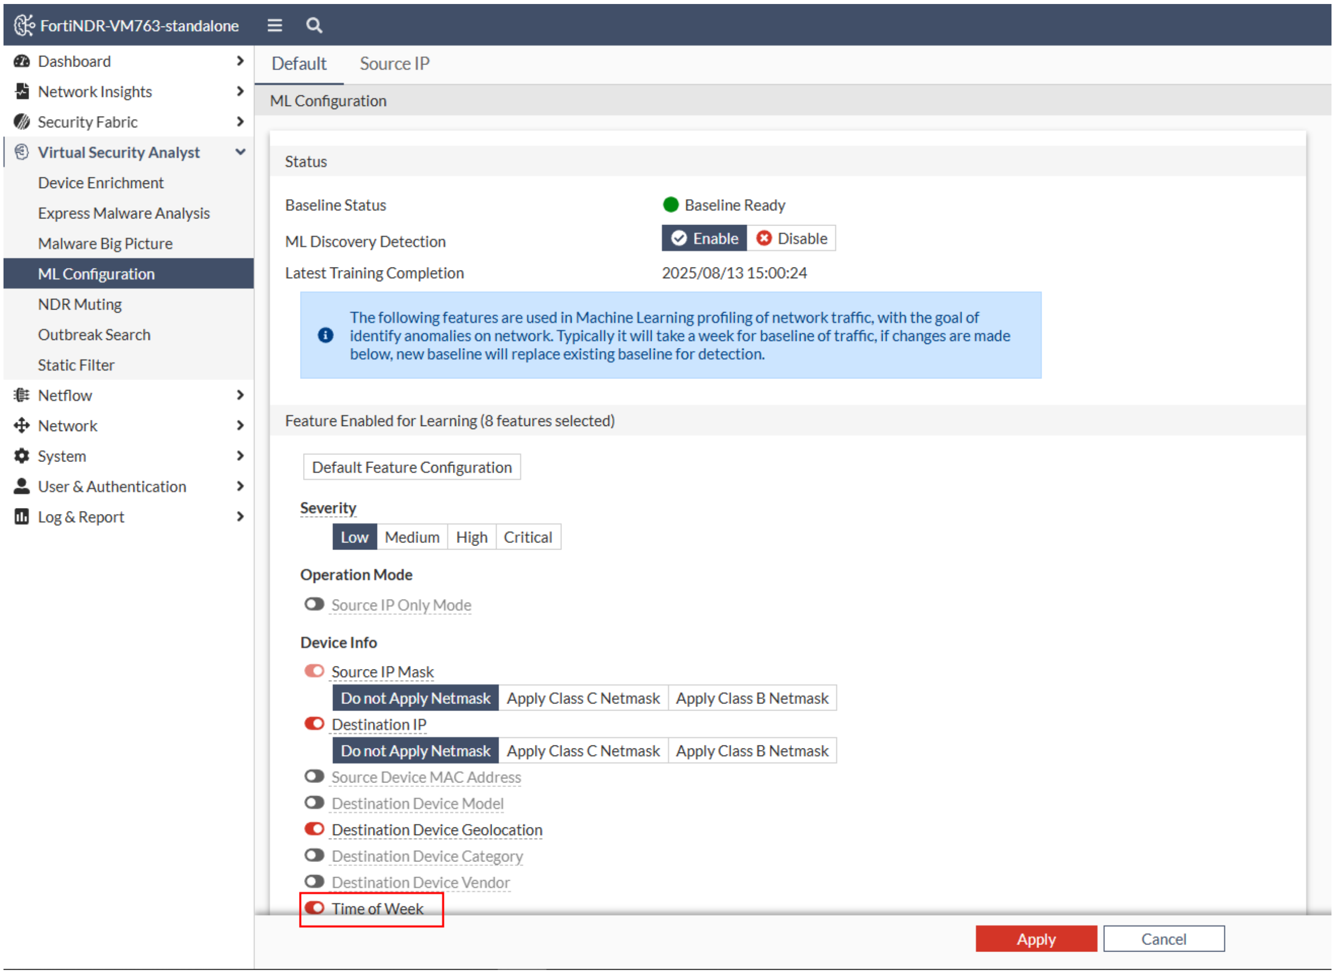

ML Configuration

Status

Baseline Status: ● Baseline Ready

ML Discovery Detection: [✓ Enable] Disable

Latest Training Completion: 2025/08/13 15:00:24

The following features are used in Machine Learning profiling of network traffic, with the goal of identify anomalies on network. Typically it will take a week for baseline of traffic, if changes are made below, new baseline will replace existing baseline for detection.

Feature Enabled for Learning (8 features selected)

Default Feature Configuration

Severity: [✓ Low] Medium High Critical

Operation Mode
- Source IP Only Mode: off

Device Info
- Source IP Mask: on — [✓ Do not Apply Netmask] Apply Class C Netmask Apply Class B Netmask
- Destination IP: on — [✓ Do not Apply Netmask] Apply Class C Netmask Apply Class B Netmask
- Source Device MAC Address: off
- Destination Device Model: off
- Destination Device Geolocation: on
- Destination Device Category: off
- Destination Device Vendor: off
- Time of Week: on

Center deployment

In this deployment example, three distributed branches each have their own FortiNDR sensors. From the dashboard, administrators can view and manage machine learning profiles for each sensor.

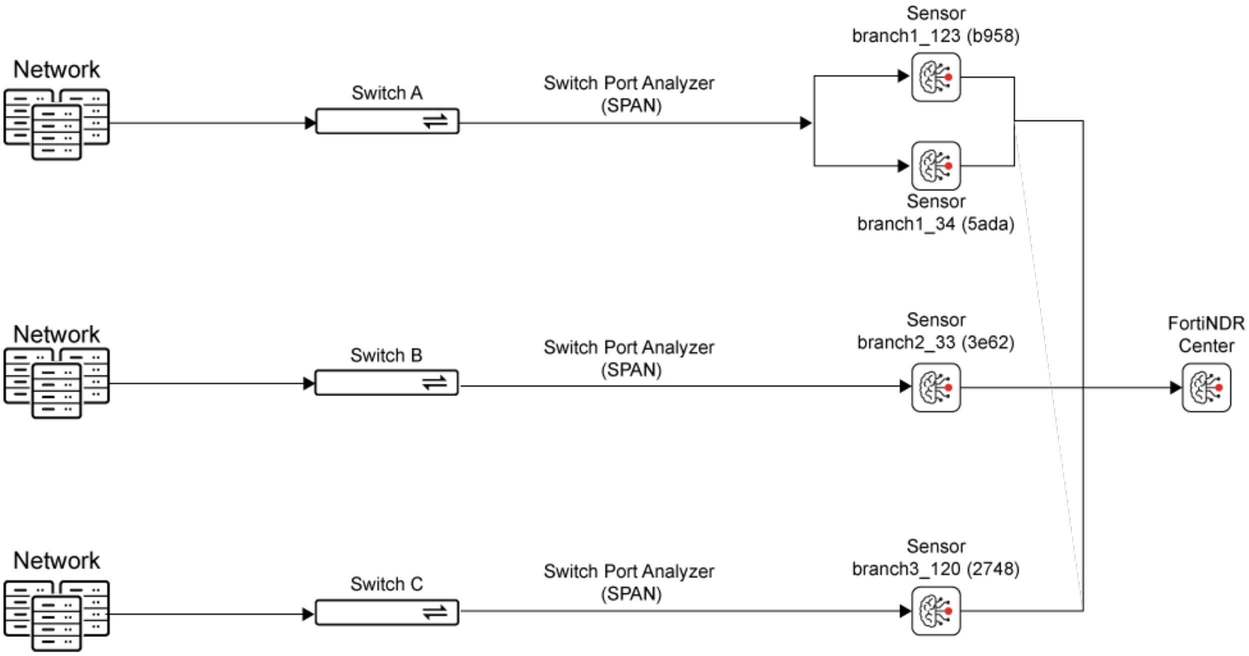
Network
Switch A
Switch Port Analyzer (SPAN)
Sensor branch1_123 (b958)
Sensor branch1_34 (5ada)
Network
Switch B
Switch Port Analyzer (SPAN)
Sensor branch2_33 (3e62)
FortiNDR Center
Network
Switch C
Switch Port Analyzer (SPAN)
Sensor branch3_120 (2748)

The administrator can create three sensor groups for the three distributed branches. All the sensors in each branch are in one group. The profiles can be configured separately.

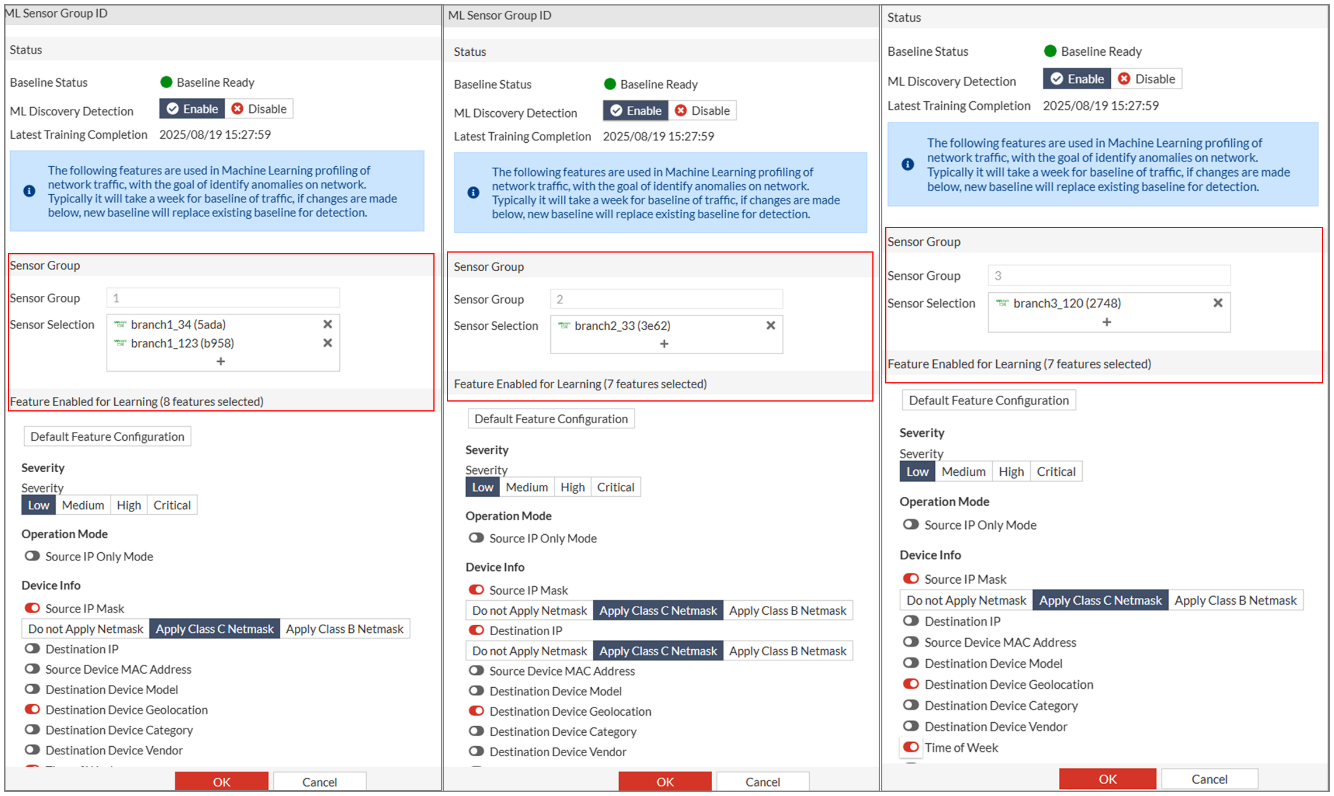

ML Sensor Group ID

Status

Baseline Status ● Baseline Ready

ML Discovery Detection [✓ Enable] Disable

Latest Training Completion 2025/08/19 15:27:59

The following features are used in Machine Learning profiling of network traffic, with the goal of identify anomalies on network. Typically it will take a week for baseline of traffic, if changes are made below, new baseline will replace existing baseline for detection.

Sensor Group

Sensor Group 1

Sensor Selection branch1_34 (5ada) ✕
branch1_123 (b958) ✕
+

Feature Enabled for Learning (8 features selected)

Default Feature Configuration

Severity

Severity [✓ Low] Medium High Critical

Operation Mode

○ Source IP Only Mode

Device Info

● Source IP Mask

Do not Apply Netmask [✓ Apply Class C Netmask] Apply Class B Netmask

○ Destination IP

○ Source Device MAC Address

○ Destination Device Model

● Destination Device Geolocation

○ Destination Device Category

○ Destination Device Vendor

OK Cancel

ML Sensor Group ID

Status

Baseline Status ● Baseline Ready

ML Discovery Detection [✓ Enable] Disable

Latest Training Completion 2025/08/19 15:27:59

The following features are used in Machine Learning profiling of network traffic, with the goal of identify anomalies on network. Typically it will take a week for baseline of traffic, if changes are made below, new baseline will replace existing baseline for detection.

Sensor Group

Sensor Group 2

Sensor Selection branch2_33 (3e62) ✕
+

Feature Enabled for Learning (7 features selected)

Default Feature Configuration

Severity

Severity [✓ Low] Medium High Critical

Operation Mode

○ Source IP Only Mode

Device Info

● Source IP Mask

Do not Apply Netmask [✓ Apply Class C Netmask] Apply Class B Netmask

● Destination IP

Do not Apply Netmask [✓ Apply Class C Netmask] Apply Class B Netmask

○ Source Device MAC Address

○ Destination Device Model

● Destination Device Geolocation

○ Destination Device Category

○ Destination Device Vendor

OK Cancel

Status

Baseline Status ● Baseline Ready

ML Discovery Detection [✓ Enable] Disable

Latest Training Completion 2025/08/19 15:27:59

The following features are used in Machine Learning profiling of network traffic, with the goal of identify anomalies on network. Typically it will take a week for baseline of traffic, if changes are made below, new baseline will replace existing baseline for detection.

Sensor Group

Sensor Group 3

Sensor Selection branch3_120 (2748) ✕
+

Feature Enabled for Learning (7 features selected)

Default Feature Configuration

Severity

Severity [✓ Low] Medium High Critical

Operation Mode

○ Source IP Only Mode

Device Info

● Source IP Mask

Do not Apply Netmask [✓ Apply Class C Netmask] Apply Class B Netmask

○ Destination IP

○ Source Device MAC Address

○ Destination Device Model

● Destination Device Geolocation

○ Destination Device Category

○ Destination Device Vendor

● Time of Week

OK Cancel

The baselines between the three branches are isolated. The detection is based according to the baseline.

Use the *Source IP* field when building baselines for different networks. This helps tailor anomaly detection to specific network segments. For more information, see the *Source IP tab* section of ML Configuration in the *FortiNDR Administration Guide*.

The following image is from a Center deployment and shows the view of multiple branches as seen from the Center device.

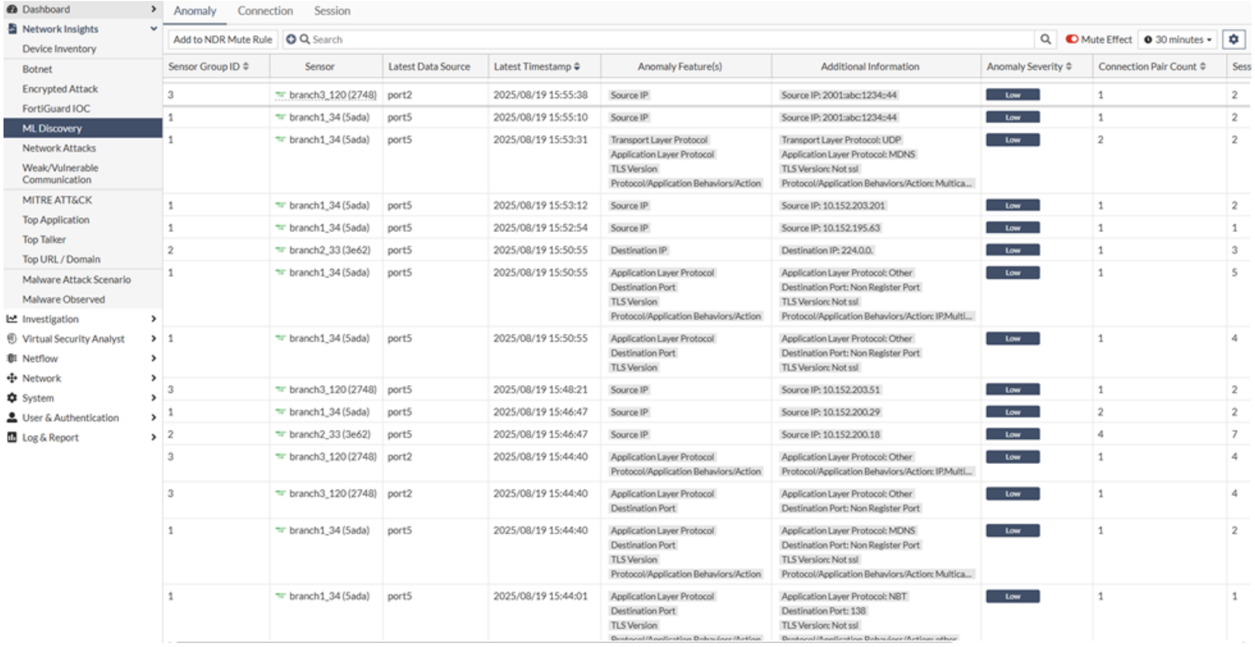

Sensor Group ID	Sensor	Latest Data Source	Latest Timestamp	Anomaly Feature(s)	Additional Information	Anomaly Severity	Connection Pair Count	Sess
3	branch3_120 (2748)	port2	2025/08/19 15:55:38	Source IP	Source IP: 2001:abc:1234::44	Low	1	2
1	branch1_34 (5ada)	port5	2025/08/19 15:55:10	Source IP	Source IP: 2001:abc:1234::44	Low	1	2
1	branch1_34 (5ada)	port5	2025/08/19 15:53:31	Transport Layer Protocol Application Layer Protocol TLS Version Protocol/Application Behaviors/Action	Transport Layer Protocol: UDP Application Layer Protocol: MDNS TLS Version: Not ssl Protocol/Application Behaviors/Action: Multica...	Low	2	2
1	branch1_34 (5ada)	port5	2025/08/19 15:53:12	Source IP	Source IP: 10.152.203.201	Low	1	2
1	branch1_34 (5ada)	port5	2025/08/19 15:52:54	Source IP	Source IP: 10.152.195.63	Low	1	1
2	branch2_33 (3e62)	port5	2025/08/19 15:50:55	Destination IP	Destination IP: 224.0.0.	Low	1	3
1	branch1_34 (5ada)	port5	2025/08/19 15:50:55	Application Layer Protocol Destination Port TLS Version Protocol/Application Behaviors/Action	Application Layer Protocol: Other Destination Port: Non Register Port TLS Version: Not ssl Protocol/Application Behaviors/Action: IP.Multi...	Low	1	5
1	branch1_34 (5ada)	port5	2025/08/19 15:50:55	Application Layer Protocol Destination Port TLS Version	Application Layer Protocol: Other Destination Port: Non Register Port TLS Version: Not ssl	Low	1	4
3	branch3_120 (2748)	port5	2025/08/19 15:48:21	Source IP	Source IP: 10.152.203.51	Low	1	2
1	branch1_34 (5ada)	port5	2025/08/19 15:46:47	Source IP	Source IP: 10.152.200.29	Low	2	2
2	branch2_33 (3e62)	port5	2025/08/19 15:46:47	Source IP	Source IP: 10.152.200.18	Low	4	7
3	branch3_120 (2748)	port2	2025/08/19 15:44:40	Application Layer Protocol Protocol/Application Behaviors/Action	Application Layer Protocol: Other Protocol/Application Behaviors/Action: IP.Multi...	Low	1	4
3	branch3_120 (2748)	port2	2025/08/19 15:44:40	Application Layer Protocol Destination Port	Application Layer Protocol: Other Destination Port: Non Register Port	Low	1	4
1	branch1_34 (5ada)	port5	2025/08/19 15:44:40	Application Layer Protocol Destination Port TLS Version Protocol/Application Behaviors/Action	Application Layer Protocol: MDNS Destination Port: Non Register Port TLS Version: Not ssl Protocol/Application Behaviors/Action: Multica...	Low	1	2
1	branch1_34 (5ada)	port5	2025/08/19 15:44:01	Application Layer Protocol Destination Port TLS Version Protocol/Application Behaviors/Action	Application Layer Protocol: NBT Destination Port: 138 TLS Version: Not ssl Protocol/Application Behaviors/Action: other	Low	1	1

Viewing the anomaly results

Single feature violation

A single feature violation occurs when one attribute of the traffic deviates from the established baseline. For example, this could include changes in:

- Source IP address
- Destination port
- Application type

Example:

During the training phase, traffic originates from the subnet 172.3.2.0/24 and uses the YouTube application. Once the baseline is established, if new traffic appears with a different source IP, different application, or different destination port, FortiNDR will trigger an alert. In the image below, the anomaly is detecting an application (Netflix) Source IP (192.168.1.2)and Destination Port (808) have never been seen before.

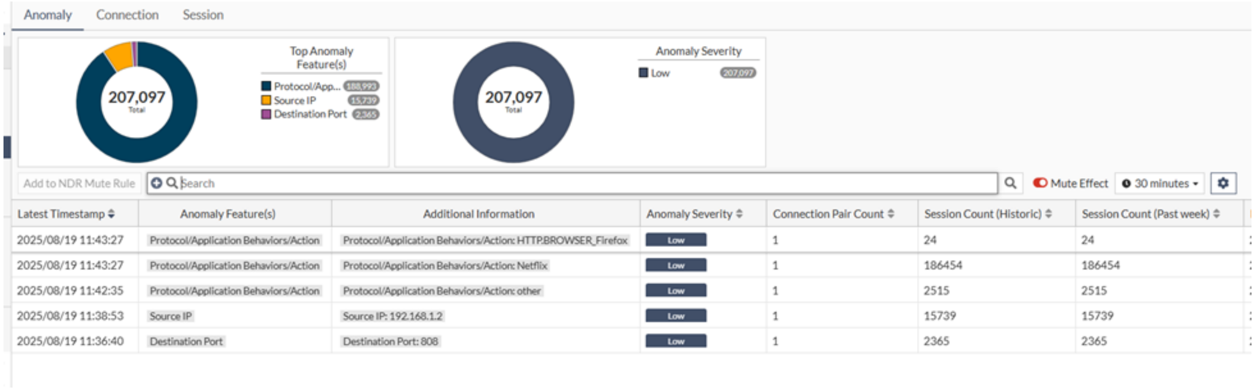

Latest Timestamp	Anomaly Feature(s)	Additional Information	Anomaly Severity	Connection Pair Count	Session Count (Historic)	Session Count (Past week)
2025/08/19 11:43:27	Protocol/Application Behaviors/Action	Protocol/Application Behaviors/Action: HTTP.BROWSER_Firefox	Low	1	24	24
2025/08/19 11:43:27	Protocol/Application Behaviors/Action	Protocol/Application Behaviors/Action: Netflix	Low	1	186454	186454
2025/08/19 11:42:35	Protocol/Application Behaviors/Action	Protocol/Application Behaviors/Action: other	Low	1	2515	2515
2025/08/19 11:38:53	Source IP	Source IP: 192.168.1.2	Low	1	15739	15739
2025/08/19 11:36:40	Destination Port	Destination Port: 808	Low	1	2365	2365

Multiple feature violation

A multiple feature violation involves a combination of attributes that, while possibly benign on their own, have not been observed together during the baseline period. This could include:

- Protocol (e.g., SMB)
- Port number
- Time of day or day of the week

Example:

If the baseline has never observed FTP traffic on destination port 21 or SMB traffic on port 443, the system will raise an alert.

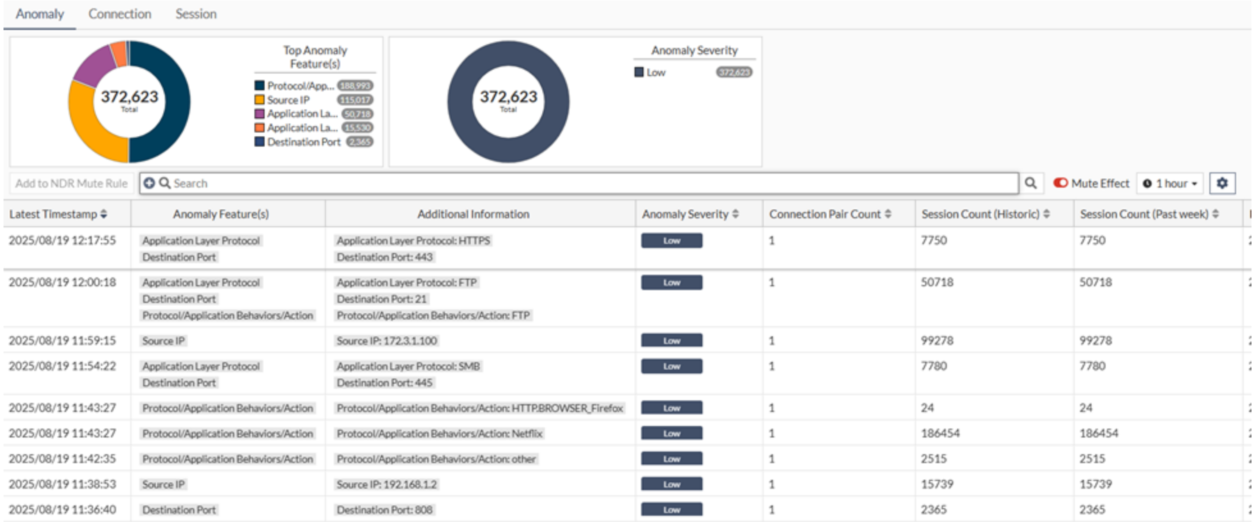

Latest Timestamp	Anomaly Feature(s)	Additional Information	Anomaly Severity	Connection Pair Count	Session Count (Historic)	Session Count (Past week)
2025/08/19 12:17:55	Application Layer Protocol Destination Port	Application Layer Protocol: HTTPS Destination Port: 443	Low	1	7750	7750
2025/08/19 12:00:18	Application Layer Protocol Destination Port Protocol/Application Behaviors/Action	Application Layer Protocol: FTP Destination Port: 21 Protocol/Application Behaviors/Action: FTP	Low	1	50718	50718
2025/08/19 11:59:15	Source IP	Source IP: 172.3.1.100	Low	1	99278	99278
2025/08/19 11:54:22	Application Layer Protocol Destination Port	Application Layer Protocol: SMB Destination Port: 445	Low	1	7780	7780
2025/08/19 11:43:27	Protocol/Application Behaviors/Action	Protocol/Application Behaviors/Action: HTTP.BROWSER_Firefox	Low	1	24	24
2025/08/19 11:43:27	Protocol/Application Behaviors/Action	Protocol/Application Behaviors/Action: Netflix	Low	1	186454	186454
2025/08/19 11:42:35	Protocol/Application Behaviors/Action	Protocol/Application Behaviors/Action: other	Low	1	2515	2515
2025/08/19 11:38:53	Source IP	Source IP: 192.168.1.2	Low	1	15739	15739
2025/08/19 11:36:40	Destination Port	Destination Port: 808	Low	1	2365	2365

Negativing to the session information page

Each insight page contains three tabs: *Anomaly*, *Connection* , and *Session*. Users can navigate to see the sessions that deviate from the baseline, depending on the tab they are viewing. The following steps describe how to access session details from the *ML Discovery* page.

From the Anomaly tab:

1. Go to *Network Insights > ML Discovery*.
2. Click the *Anomaly* tab.
3. Click an anomaly in the list.
4. In the *Anomaly Information* pane, double-click a *Connection Pair*.
5. In the *Sessions with selected condition* pane, click *View Session* or double-click a detection.

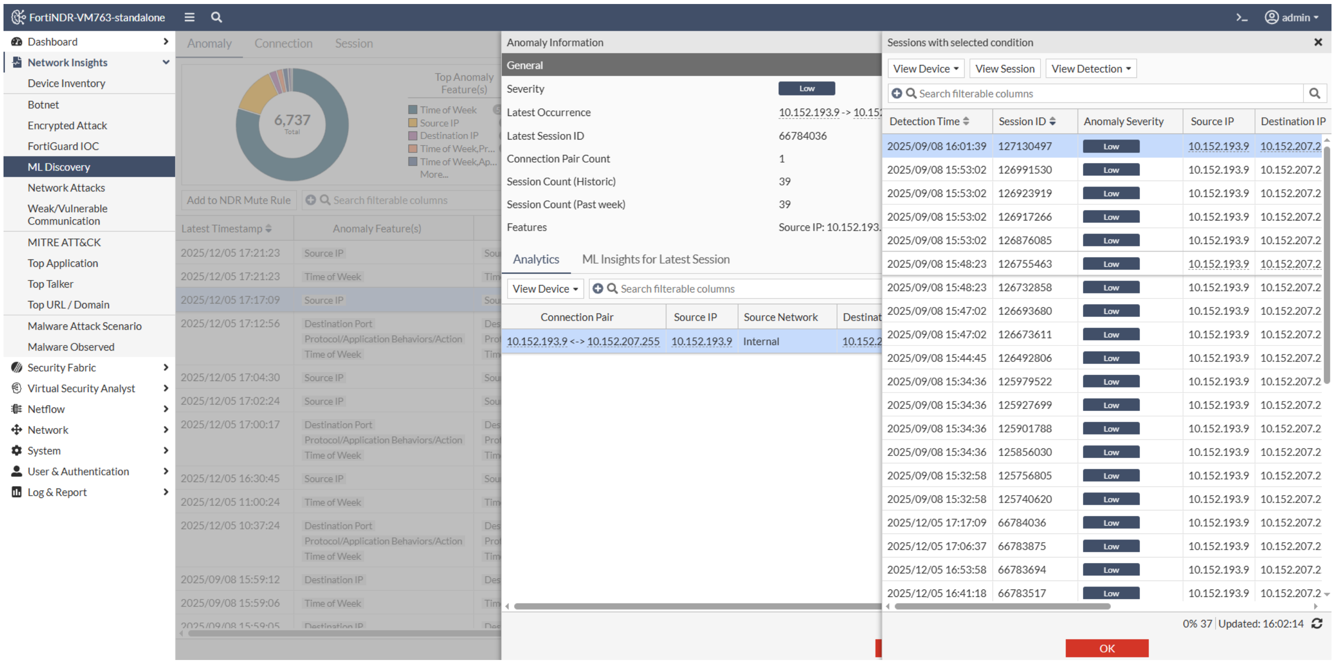

Anomaly Information

General

Severity	Low
Latest Occurrence	10.152.193.9 -> 10.152
Latest Session ID	66784036
Connection Pair Count	1
Session Count (Historic)	39
Session Count (Past week)	39
Features	Source IP: 10.152.193.

Analytics | ML Insights for Latest Session

Connection Pair	Source IP	Source Network	Destinat
10.152.193.9 <-> 10.152.207.255	10.152.193.9	Internal	10.152.2

Sessions with selected condition

Detection Time	Session ID	Anomaly Severity	Source IP	Destination IP
2025/09/08 16:01:39	127130497	Low	10.152.193.9	10.152.207.2
2025/09/08 15:53:02	126991530	Low	10.152.193.9	10.152.207.2
2025/09/08 15:53:02	126923919	Low	10.152.193.9	10.152.207.2
2025/09/08 15:53:02	126917266	Low	10.152.193.9	10.152.207.2
2025/09/08 15:53:02	126876085	Low	10.152.193.9	10.152.207.2
2025/09/08 15:48:23	126755463	Low	10.152.193.9	10.152.207.2
2025/09/08 15:48:23	126732858	Low	10.152.193.9	10.152.207.2
2025/09/08 15:47:02	126693680	Low	10.152.193.9	10.152.207.2
2025/09/08 15:47:02	126673611	Low	10.152.193.9	10.152.207.2
2025/09/08 15:44:45	126492806	Low	10.152.193.9	10.152.207.2
2025/09/08 15:34:36	125979522	Low	10.152.193.9	10.152.207.2
2025/09/08 15:34:36	125927699	Low	10.152.193.9	10.152.207.2
2025/09/08 15:34:36	125901788	Low	10.152.193.9	10.152.207.2
2025/09/08 15:34:36	125856030	Low	10.152.193.9	10.152.207.2
2025/09/08 15:32:58	125756805	Low	10.152.193.9	10.152.207.2
2025/09/08 15:32:58	125740620	Low	10.152.193.9	10.152.207.2
2025/12/05 17:17:09	66784036	Low	10.152.193.9	10.152.207.2
2025/12/05 17:06:37	66783875	Low	10.152.193.9	10.152.207.2
2025/12/05 16:53:58	66783694	Low	10.152.193.9	10.152.207.2
2025/12/05 16:41:18	66783517	Low	10.152.193.9	10.152.207.2

0% 37 | Updated: 16:02:14

From the Connections tab:

1. Go to *Network Insights > ML Discovery*.
2. Click the *Connection* tab.
3. Click a connection pair in the list.
4. Select an entry in the list. The session information page opens.

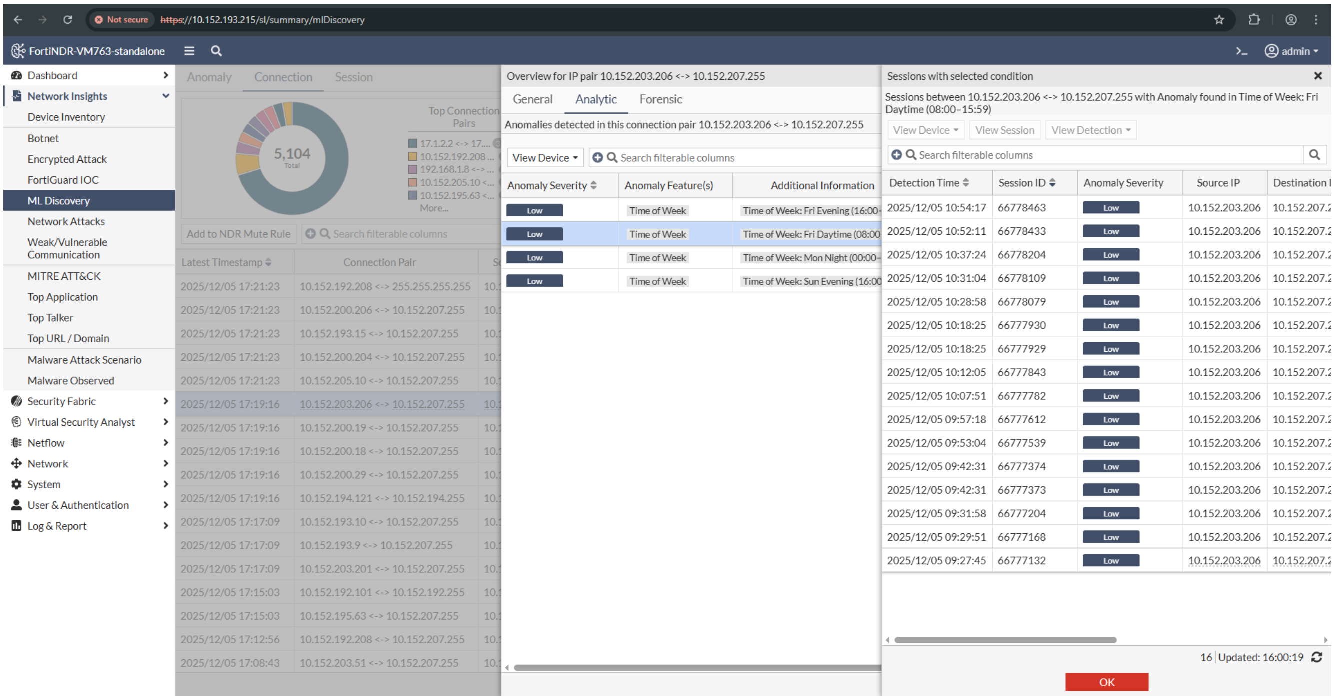

Latest Timestamp	Connection Pair
2025/12/05 17:21:23	10.152.192.208 <-> 255.255.255.255
2025/12/05 17:21:23	10.152.200.206 <-> 10.152.207.255
2025/12/05 17:21:23	10.152.193.15 <-> 10.152.207.255
2025/12/05 17:21:23	10.152.200.204 <-> 10.152.207.255
2025/12/05 17:21:23	10.152.205.10 <-> 10.152.207.255
2025/12/05 17:19:16	10.152.203.206 <-> 10.152.207.255
2025/12/05 17:19:16	10.152.200.19 <-> 10.152.207.255
2025/12/05 17:19:16	10.152.200.18 <-> 10.152.207.255
2025/12/05 17:19:16	10.152.200.29 <-> 10.152.207.255
2025/12/05 17:19:16	10.152.194.121 <-> 10.152.194.255
2025/12/05 17:17:09	10.152.193.10 <-> 10.152.207.255
2025/12/05 17:17:09	10.152.193.9 <-> 10.152.207.255
2025/12/05 17:17:09	10.152.203.201 <-> 10.152.207.255
2025/12/05 17:15:03	10.152.192.101 <-> 10.152.192.255
2025/12/05 17:15:03	10.152.195.63 <-> 10.152.207.255
2025/12/05 17:12:56	10.152.192.208 <-> 10.152.207.255
2025/12/05 17:08:43	10.152.203.51 <-> 10.152.207.255

Overview for IP pair 10.152.203.206 <-> 10.152.207.255

Anomalies detected in this connection pair 10.152.203.206 <-> 10.152.207.255

Anomaly Severity	Anomaly Feature(s)	Additional Information
Low	Time of Week	Time of Week: Fri Evening (16:00-
Low	Time of Week	Time of Week: Fri Daytime (08:00
Low	Time of Week	Time of Week: Mon Night (00:00–
Low	Time of Week	Time of Week: Sun Evening (16:00

Sessions with selected condition

Sessions between 10.152.203.206 <-> 10.152.207.255 with Anomaly found in Time of Week: Fri Daytime (08:00–15:59)

Detection Time	Session ID	Anomaly Severity	Source IP	Destination I
2025/12/05 10:54:17	66778463	Low	10.152.203.206	10.152.207.2
2025/12/05 10:52:11	66778433	Low	10.152.203.206	10.152.207.2
2025/12/05 10:37:24	66778204	Low	10.152.203.206	10.152.207.2
2025/12/05 10:31:04	66778109	Low	10.152.203.206	10.152.207.2
2025/12/05 10:28:58	66778079	Low	10.152.203.206	10.152.207.2
2025/12/05 10:18:25	66777930	Low	10.152.203.206	10.152.207.2
2025/12/05 10:18:25	66777929	Low	10.152.203.206	10.152.207.2
2025/12/05 10:12:05	66777843	Low	10.152.203.206	10.152.207.2
2025/12/05 10:07:51	66777782	Low	10.152.203.206	10.152.207.2
2025/12/05 09:57:18	66777612	Low	10.152.203.206	10.152.207.2
2025/12/05 09:53:04	66777539	Low	10.152.203.206	10.152.207.2
2025/12/05 09:42:31	66777374	Low	10.152.203.206	10.152.207.2
2025/12/05 09:42:31	66777373	Low	10.152.203.206	10.152.207.2
2025/12/05 09:31:58	66777204	Low	10.152.203.206	10.152.207.2
2025/12/05 09:29:51	66777168	Low	10.152.203.206	10.152.207.2
2025/12/05 09:27:45	66777132	Low	10.152.203.206	10.152.207.2

16 Updated: 16:00:19

From the Sessions Tab

1. Go to *Network Insights > ML Discovery*.
2. Click the *Sessions* tab and double-click an ML anomaly.
3. In the *Anomaly Information* pane, double-click an entry in the list.

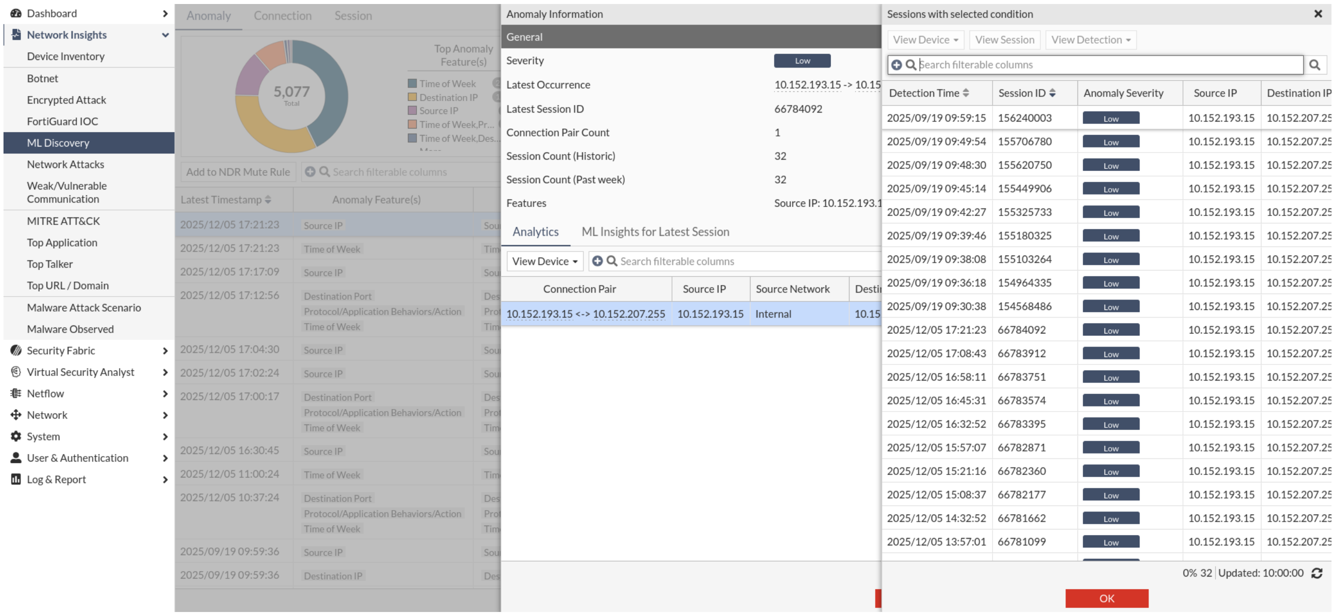

Latest Timestamp	Anomaly Feature(s)
2025/12/05 17:21:23	Source IP
2025/12/05 17:21:23	Time of Week
2025/12/05 17:17:09	Source IP
2025/12/05 17:12:56	Destination Port Protocol/Application Behaviors/Action Time of Week
2025/12/05 17:04:30	Source IP
2025/12/05 17:02:24	Source IP
2025/12/05 17:00:17	Destination Port Protocol/Application Behaviors/Action Time of Week
2025/12/05 16:30:45	Source IP
2025/12/05 11:00:24	Time of Week
2025/12/05 10:37:24	Destination Port Protocol/Application Behaviors/Action Time of Week
2025/09/19 09:59:36	Source IP
2025/09/19 09:59:36	Destination IP

Anomaly Information

General

Severity	Low
Latest Occurrence	10.152.193.15 -> 10.15
Latest Session ID	66784092
Connection Pair Count	1
Session Count (Historic)	32
Session Count (Past week)	32
Features	Source IP: 10.152.193.1

Connection Pair	Source IP	Source Network	Desti
10.152.193.15 <-> 10.152.207.255	10.152.193.15	Internal	10.15

Sessions with selected condition

Detection Time	Session ID	Anomaly Severity	Source IP	Destination IP
2025/09/19 09:59:15	156240003	Low	10.152.193.15	10.152.207.25
2025/09/19 09:49:54	155706780	Low	10.152.193.15	10.152.207.25
2025/09/19 09:48:30	155620750	Low	10.152.193.15	10.152.207.25
2025/09/19 09:45:14	155449906	Low	10.152.193.15	10.152.207.25
2025/09/19 09:42:27	155325733	Low	10.152.193.15	10.152.207.25
2025/09/19 09:39:46	155180325	Low	10.152.193.15	10.152.207.25
2025/09/19 09:38:08	155103264	Low	10.152.193.15	10.152.207.25
2025/09/19 09:36:18	154964335	Low	10.152.193.15	10.152.207.25
2025/09/19 09:30:38	154568486	Low	10.152.193.15	10.152.207.25
2025/12/05 17:21:23	66784092	Low	10.152.193.15	10.152.207.25
2025/12/05 17:08:43	66783912	Low	10.152.193.15	10.152.207.25
2025/12/05 16:58:11	66783751	Low	10.152.193.15	10.152.207.25
2025/12/05 16:45:31	66783574	Low	10.152.193.15	10.152.207.25
2025/12/05 16:32:52	66783395	Low	10.152.193.15	10.152.207.25
2025/12/05 15:57:07	66782871	Low	10.152.193.15	10.152.207.25
2025/12/05 15:21:16	66782360	Low	10.152.193.15	10.152.207.25
2025/12/05 15:08:37	66782177	Low	10.152.193.15	10.152.207.25
2025/12/05 14:32:52	66781662	Low	10.152.193.15	10.152.207.25
2025/12/05 13:57:01	66781099	Low	10.152.193.15	10.152.207.25

0% 32 Updated: 10:00:00

ML Discovery and ML Insights

Depending on the anomaly, the ML Discovery and Insights sections display either newly detected source and destination IP addresses, or charts that illustrate how the anomaly was detected.

- The *ML Discovery* section appears in the Session Information page.
- The *ML Insights for Latest Session* section appears in the *Anomaly Information* panel when an anomaly is double-clicked in the *Anomaly* tab.

Source IP, Destination IP

Only the new IP address will be displayed, since its appearance alone constitutes the irregularity.

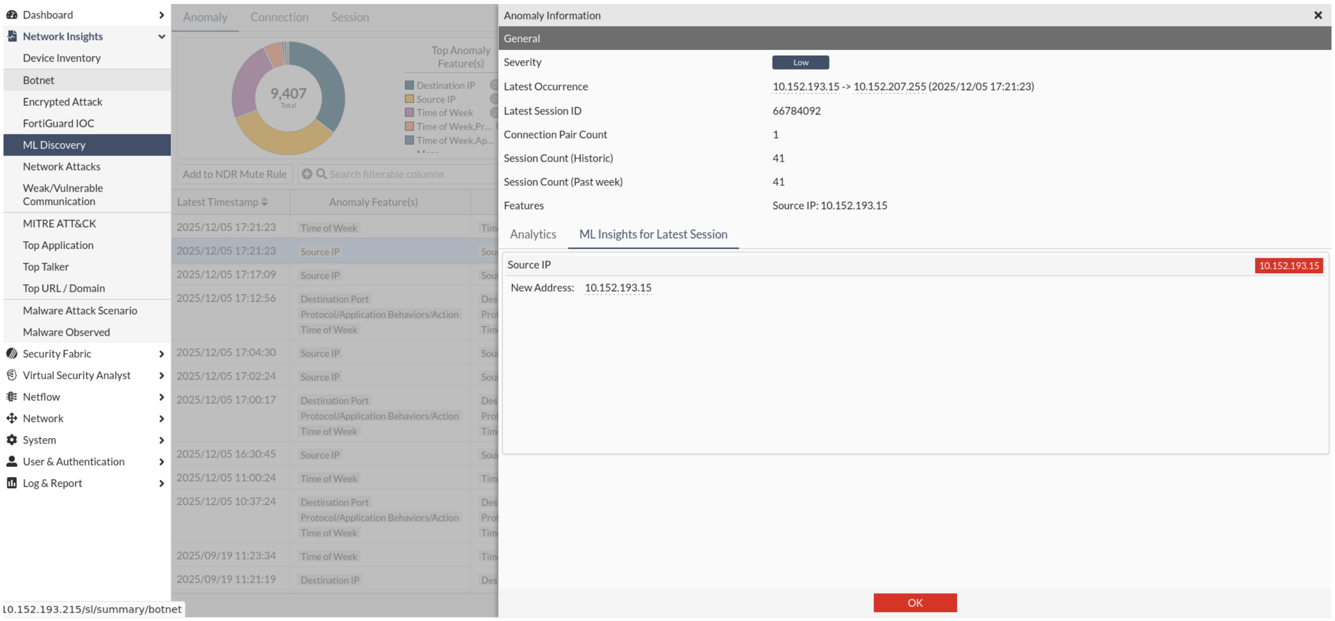

Latest Timestamp	Anomaly Feature(s)
2025/12/05 17:21:23	Time of Week
2025/12/05 17:21:23	Source IP
2025/12/05 17:17:09	Source IP
2025/12/05 17:12:56	Destination Port Protocol/Application Behaviors/Action Time of Week
2025/12/05 17:04:30	Source IP
2025/12/05 17:02:24	Source IP
2025/12/05 17:00:17	Destination Port Protocol/Application Behaviors/Action Time of Week
2025/12/05 16:30:45	Source IP
2025/12/05 11:00:24	Time of Week
2025/12/05 10:37:24	Destination Port Protocol/Application Behaviors/Action Time of Week
2025/09/19 11:23:34	Time of Week
2025/09/19 11:21:19	Destination IP

Anomaly Information

General

Severity	Low
Latest Occurrence	10.152.193.15 -> 10.152.207.255 (2025/12/05 17:21:23)
Latest Session ID	66784092
Connection Pair Count	1
Session Count (Historic)	41
Session Count (Past week)	41
Features	Source IP: 10.152.193.15

Source IP 10.152.193.15

New Address: 10.152.193.15

Time of week (7.6.3 and higher)

The Timetable displays session activity for the current source IP, segmented by day of the week and three distinct time periods:

- *Night*: 00:00–07:59
- *Daytime*: 08:00–15:59
- *Evening*: 16:00–23:59

Currently, these time settings are defined by the system and cannot be customized.

Each time slot displays the cumulative number of sessions observed since the start of training. You can hover over any slot to view the exact session count.

- Red-colored slots indicate anomalous time periods. These had zero sessions during training but now show activity.
- Empty (uncolored) slots represent time periods with no sessions recorded since training began.

You can also view machine learning-based anomalies for each session on the *Session Profile* page. Use this chart to understand what the typical behavior looks like for the Source IP. It also helps provide context for identifying anomalies.

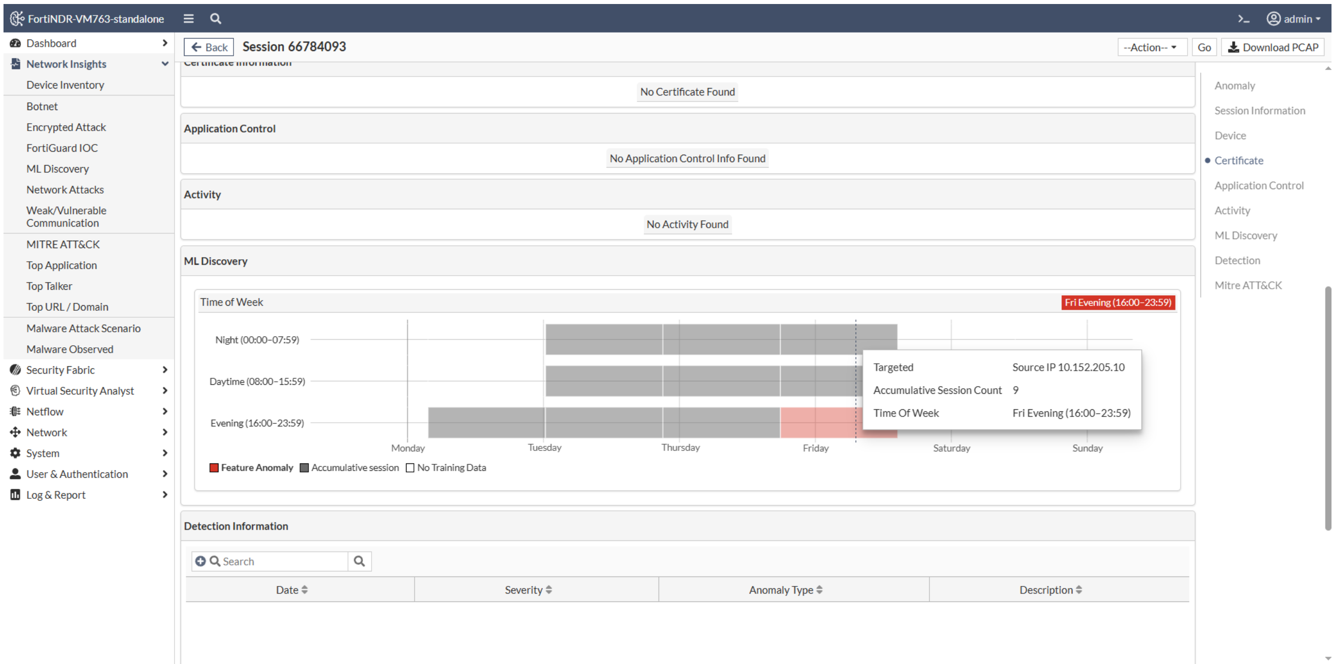

Session 66784093

No Certificate Found

Application Control

No Application Control Info Found

Activity

No Activity Found

ML Discovery

Time of Week — Fri Evening (16:00–23:59)

Targeted	Source IP 10.152.205.10
Accumulative Session Count	9
Time Of Week	Fri Evening (16:00–23:59)

Detection Information

Date	Severity	Anomaly Type	Description

The following image shows the violating IP pairs, the date and time of their last occurrence, and the associated features (e.g., *Time of Week: Friday daytime*). This indicates that FortiNDR has not observed any traffic during this time period for these IP pairs in the baseline.

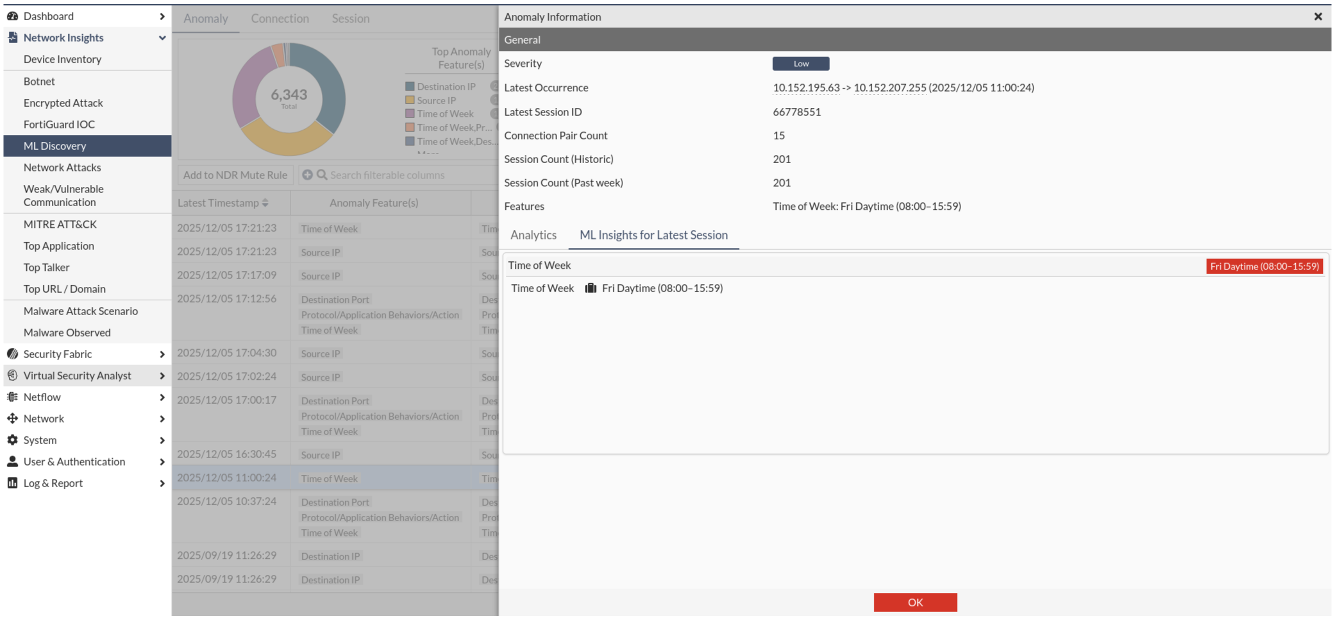
Dashboard
Network Insights
Device Inventory
Botnet
Encrypted Attack
FortiGuard IOC
ML Discovery
Network Attacks
Weak/Vulnerable Communication
MITRE ATT&CK
Top Application
Top Talker
Top URL / Domain
Malware Attack Scenario
Malware Observed
Security Fabric
Virtual Security Analyst
Netflow
Network
System
User & Authentication
Log & Report
Anomaly
Connection
Session
6,343
Total
Top Anomaly Feature(s)
Destination IP
Source IP
Time of Week
Time of Week,Pr...
Time of Week,Des...
Add to NDR Mute Rule
Search filterable columns
Latest Timestamp
Anomaly Feature(s)
2025/12/05 17:21:23 Time of Week
2025/12/05 17:21:23 Source IP
2025/12/05 17:17:09 Source IP
2025/12/05 17:12:56 Destination Port Protocol/Application Behaviors/Action Time of Week
2025/12/05 17:04:30 Source IP
2025/12/05 17:02:24 Source IP
2025/12/05 17:00:17 Destination Port Protocol/Application Behaviors/Action Time of Week
2025/12/05 16:30:45 Source IP
2025/12/05 11:00:24 Time of Week
2025/12/05 10:37:24 Destination Port Protocol/Application Behaviors/Action Time of Week
2025/09/19 11:26:29 Destination IP
2025/09/19 11:26:29 Destination IP
OK

Anomaly Information

General

Severity	Low
Latest Occurrence	10.152.195.63 -> 10.152.207.255 (2025/12/05 11:00:24)
Latest Session ID	66778551
Connection Pair Count	15
Session Count (Historic)	201
Session Count (Past week)	201
Features	Time of Week: Fri Daytime (08:00–15:59)

Analytics | ML Insights for Latest Session

Time of Week — Fri Daytime (08:00–15:59)

Time of Week: Fri Daytime (08:00–15:59)

Combination feature ML anomaly detection

FortiNDR's ML algorithms automatically detect both single-feature violations and combinations of feature violations. A *Combination Feature ML Anomaly* refers to a session where the specific combination of feature values has never been observed in the ML baseline. In other words, this exact set of feature values was not seen during the training period or at any point before the time of detection, making it an outlier in the learned behavior.

A pie chart displays the number of detections for each value within a specific feature category, based on the current session’s Source IP. It reflects the total session count accumulated from the beginning of the training period up to the point of detection.

Each time period displays the cumulative number of sessions observed since the start of training. You can hover over any time period to view the exact session count.

- A red tile indicates anomalous time periods. These had zero sessions during training but now show activity.
- An empty tile represent time periods with no sessions recorded since training began.

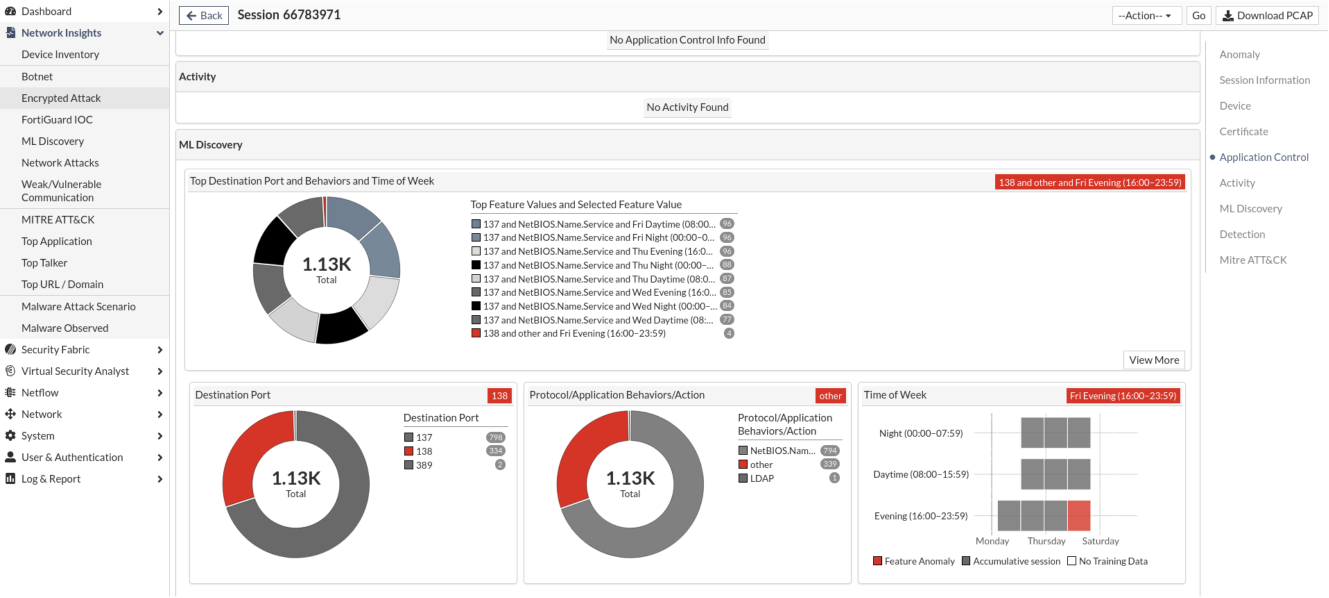
Dashboard
Network Insights
Device Inventory
Botnet
Encrypted Attack
FortiGuard IOC
ML Discovery
Network Attacks
Weak/Vulnerable Communication
MITRE ATT&CK
Top Application
Top Talker
Top URL / Domain
Malware Attack Scenario
Malware Observed
Security Fabric
Virtual Security Analyst
Netflow
Network
System
User & Authentication
Log & Report
Back
Session 66783971
--Action--
Go
Download PCAP
No Application Control Info Found
Activity
No Activity Found
ML Discovery
Top Destination Port and Behaviors and Time of Week
138 and other and Fri Evening (16:00–23:59)
1.13K
Total
Top Feature Values and Selected Feature Value
137 and NetBIOS.Name.Service and Fri Daytime (08:00... 96
137 and NetBIOS.Name.Service and Fri Night (00:00–0... 96
137 and NetBIOS.Name.Service and Thu Evening (16:0... 96
137 and NetBIOS.Name.Service and Thu Night (00:00–... 88
137 and NetBIOS.Name.Service and Thu Daytime (08:0... 87
137 and NetBIOS.Name.Service and Wed Evening (16:0... 85
137 and NetBIOS.Name.Service and Wed Night (00:00–... 84
137 and NetBIOS.Name.Service and Wed Daytime (08:... 77
138 and other and Fri Evening (16:00–23:59) 4
View More
Destination Port
138
1.13K
Total
Destination Port
137 798
138 334
389 2
Protocol/Application Behaviors/Action
other
1.13K
Total
Protocol/Application Behaviors/Action
NetBIOS.Nam... 794
other 339
LDAP 1
Time of Week
Fri Evening (16:00–23:59)
Night (00:00–07:59)
Daytime (08:00–15:59)
Evening (16:00–23:59)
Monday
Thursday
Saturday
Feature Anomaly
Accumulative session
No Training Data
Anomaly
Session Information
Device
Certificate
Application Control
Activity
ML Discovery
Detection
Mitre ATT&CK

Top behaviors

To learn more about the reasoning behind ML Discovery detections, refer to the *ML Discovery* section of each session’s details page. The top behavior chart summarizes the top combinations of detection values, including the one currently flagged as the ML detection. These combinations and their counts help you understand which patterns are most common for the Source IP that is in question.

Below the chart, individual information blocks display each ML feature type detected for the current source IP.

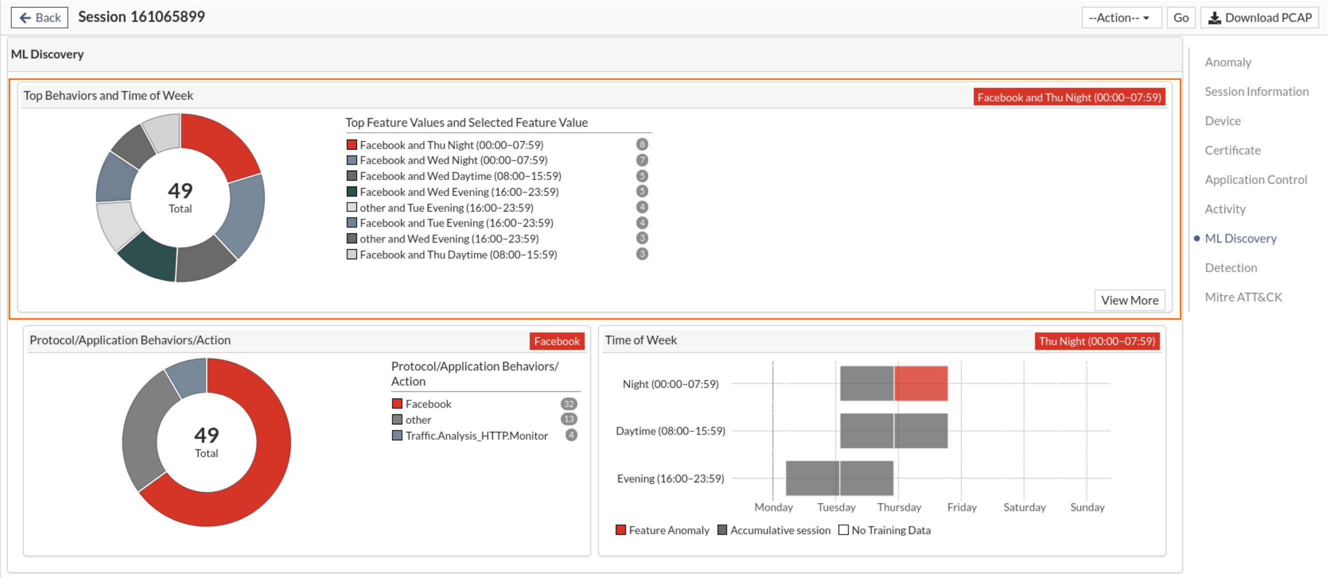
Back
Session 161065899
--Action--
Go
Download PCAP
ML Discovery
Top Behaviors and Time of Week
Facebook and Thu Night (00:00–07:59)
49
Total
Top Feature Values and Selected Feature Value
🟥 Facebook and Thu Night (00:00–07:59) 8
▨ Facebook and Wed Night (00:00–07:59) 7
▨ Facebook and Wed Daytime (08:00–15:59) 5
▨ Facebook and Wed Evening (16:00–23:59) 5
□ other and Tue Evening (16:00–23:59) 4
▨ Facebook and Tue Evening (16:00–23:59) 4
▨ other and Wed Evening (16:00–23:59) 3
□ Facebook and Thu Daytime (08:00–15:59) 3
View More
Protocol/Application Behaviors/Action
Facebook
49
Total
Protocol/Application Behaviors/Action
🟥 Facebook 32
▨ other 13
▨ Traffic.Analysis_HTTP.Monitor 4
Time of Week
Thu Night (00:00–07:59)
Night (00:00–07:59)
Daytime (08:00–15:59)
Evening (16:00–23:59)
Monday
Tuesday
Thursday
Friday
Saturday
Sunday
🟥 Feature Anomaly
▨ Accumulative session
□ No Training Data
Anomaly
Session Information
Device
Certificate
Application Control
Activity
ML Discovery
Detection
Mitre ATT&CK

Click the View More button to drill down into the information.

The following image shows that the combination of the Facebook and HTTP Monitor applications for this session and IP address was not observed in the baseline. Since the time feature is enabled, sessions are categorized into different time segments.

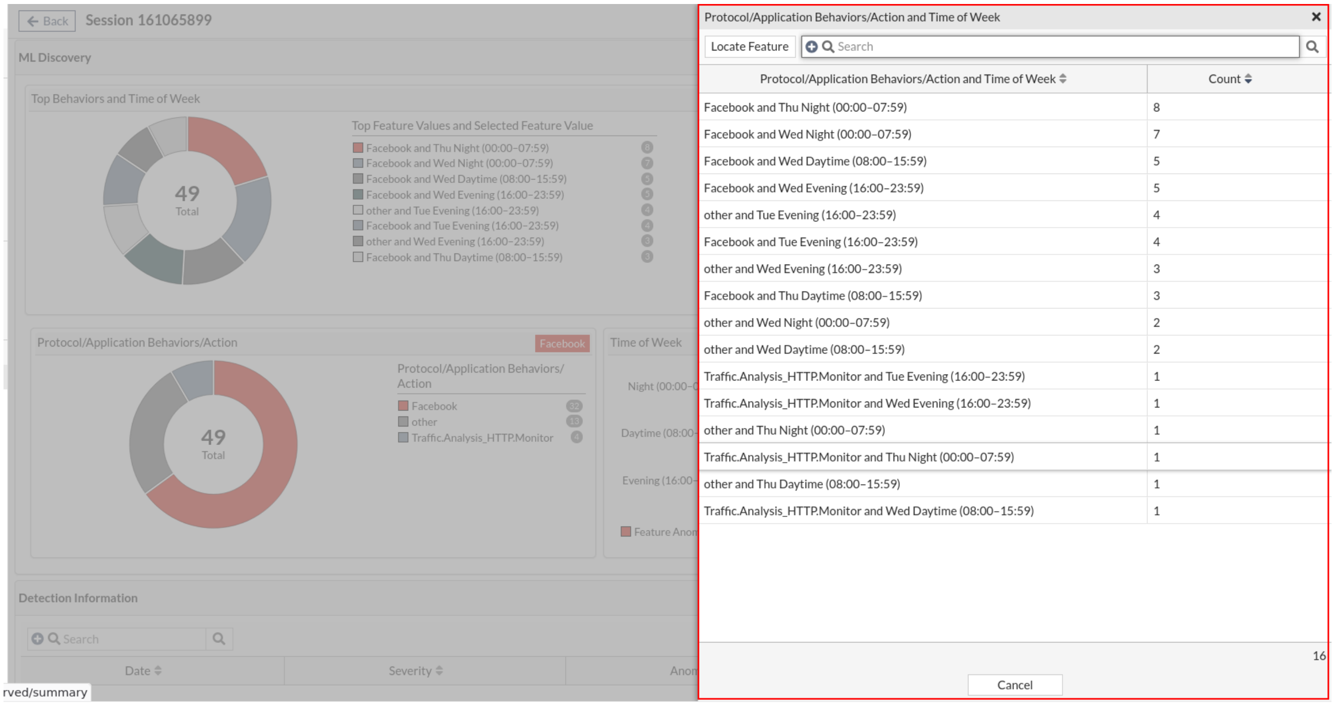

Protocol/Application Behaviors/Action and Time of Week

Protocol/Application Behaviors/Action and Time of Week	Count
Facebook and Thu Night (00:00–07:59)	8
Facebook and Wed Night (00:00–07:59)	7
Facebook and Wed Daytime (08:00–15:59)	5
Facebook and Wed Evening (16:00–23:59)	5
other and Tue Evening (16:00–23:59)	4
Facebook and Tue Evening (16:00–23:59)	4
other and Wed Evening (16:00–23:59)	3
Facebook and Thu Daytime (08:00–15:59)	3
other and Wed Night (00:00–07:59)	2
other and Wed Daytime (08:00–15:59)	2
Traffic.Analysis_HTTP.Monitor and Tue Evening (16:00–23:59)	1
Traffic.Analysis_HTTP.Monitor and Wed Evening (16:00–23:59)	1
other and Thu Night (00:00–07:59)	1
Traffic.Analysis_HTTP.Monitor and Thu Night (00:00–07:59)	1
other and Thu Daytime (08:00–15:59)	1
Traffic.Analysis_HTTP.Monitor and Wed Daytime (08:00–15:59)	1

16

Advance Muting

NDR Muting helps streamline your security monitoring by focusing on the most critical threats. Muting allows you to hide anomalies that are not relevant to your network. By muting anomalies, FortiNDR ensures that these detections are no longer visible in insight pages and prevents related alerts and enforcement actions.

Once an attack is muted, any future information related to this anomaly will be hidden from the insight pages, however, the information is not deleted. Muted entries will be omitted from Syslog and will not be quarantined by the Automation Framework. FortiNDR does not mute any historical violations. See NDR Muting.

User Feedback

The feedback feature in the *ML Discovery* page allows users to provide input on anomalies detected by the system's Machine Learning engine to help correct false positives. By submitting feedback, users can influence how the system interprets future anomalies, improving the accuracy of ML-based detections over time. Feedback is shown in the *Current Feedback Status* column of the *ML Discovery* monitor.

The main difference between *User Feedback* and *Advanced Muting* is that *User Feedback* updates the baseline, preventing similar anomalies from appearing again.*Advanced Muting does not affect the baseline*. Instead, it controls the display of anomalies, which can be toggled on or off in the GUI settings, including pages like *Network Insights*.

For more information, see ML Discovery (Center - Standalone).

Dual Center Management

Dual Center Machine Learning ensures data redundancy by synchronizing sensor data to two Center devices simultaneously. This configuration allows SOC operators to log onto either the Primary or Secondary Center FortiNDR, providing flexibility and reliability. Importantly, there is no communication between the two Centers, maintaining their independence and ensuring data integrity.

The following topology illustrates the principles of Dual Center Machine Learning:

1. The SOC Operator logs onto either the Primary or Secondary device.
2. The sensors synchronize the data between the Primary or Secondary devices.
3. There is no communication between the Center devices.
4. Both sensors receive updates from FortiGuard.

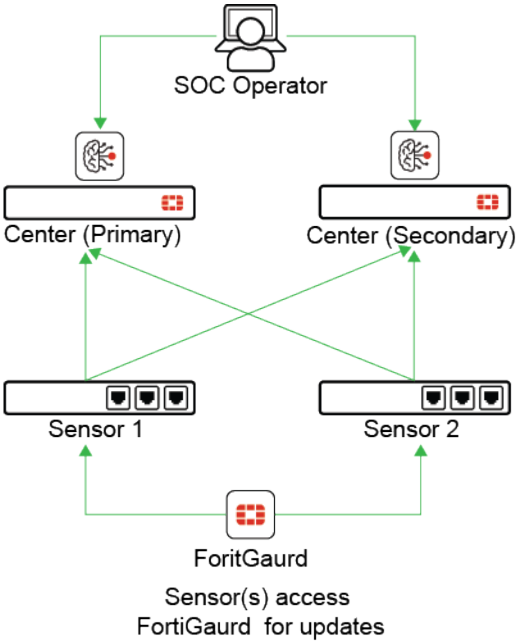

The FortiNDR system supports multiple Centers, designed for redundancy and offering centralized views of sensor detections. However, due to the lack of synchronization between the Centers, it does not function like traditional *High Availability* setups. As a result, log views may be similar across Centers but not identical.

Centers that operate independently are useful when:

- Traffic profiling uses different machine learning (ML) baseline traffic from Centers, based on features and timing, to detect anomalies. For Example:
  - Center 1: Evaluates source/destination IPs and geographical profile settings from sensors 1, 2, and 3.
  - Center 2: Evaluates applications and destination ports from sensors 1, 4, and 5.
- In less sophisticated environments, only one Center's ML can be enabled to log detections (optionally forwarded to FortiAnalyzer), especially when the administrator is not familiar with network traffic. This helps reduce alert fatigue.

- ML baseline profile timing: Since Centers operate independently, the learning period of the baseline (start and stop time) can differ. Even if the baseline timing is the same across multiple centers, based on log synchronization principles, the ML detection log views may be similar but not exactly the same.

Dual center Synchronization

Sensors (regardless of how many) synchronize logs in batches, alternating between Primary and Secondary Center devices. Each Center processes the log batches independently. Log batches include detections (e.g., IPS/malware) and session logs for Machine Learning.

If network connectivity is interrupted, a backlog may occur; however, log synchronization will eventually catch up once connectivity is restored.

The following topology illustrates the Dual Center synchronization process:

1. Sensors synchronize logs in batches.
2. When Batch 1 is finished synchronizing with the Primary Center, it will synch with the Secondary Center.
3. Each Center processes logs independently.

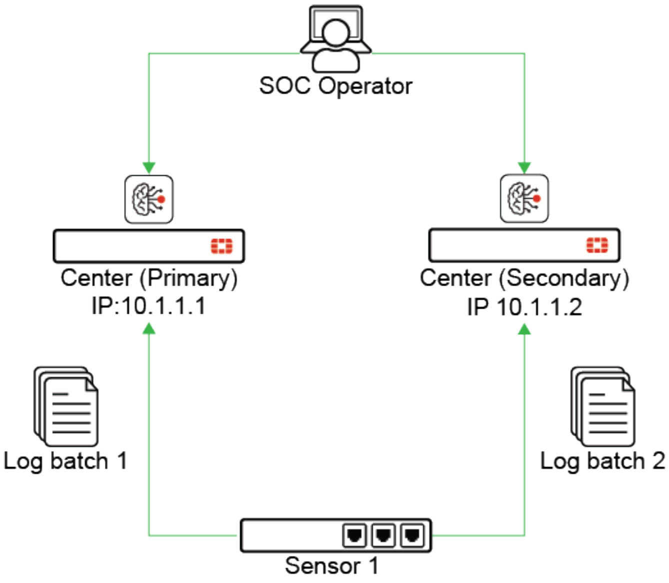

Batch log synchronization example

In this example, an SOC Operator connects a Primary and Secondary Center device to a sensor.

The sensor sends the logs in batches. Once the Primary Center is connected to the sensor, the sensor starts synchronizing the batches with the Primary Center.

After the sensor has completed synchronizing the first batch of logs with the Primary Center, it begins synchronizing with the Secondary Center. If, for instance, a network connectivity issue interrupts the synch causing a backlog, once the connection issue is resolved, the synchronization resumes. As a result, the views in the two Centers may be similar but not identical due to independent synchronization and processing of session IDs.

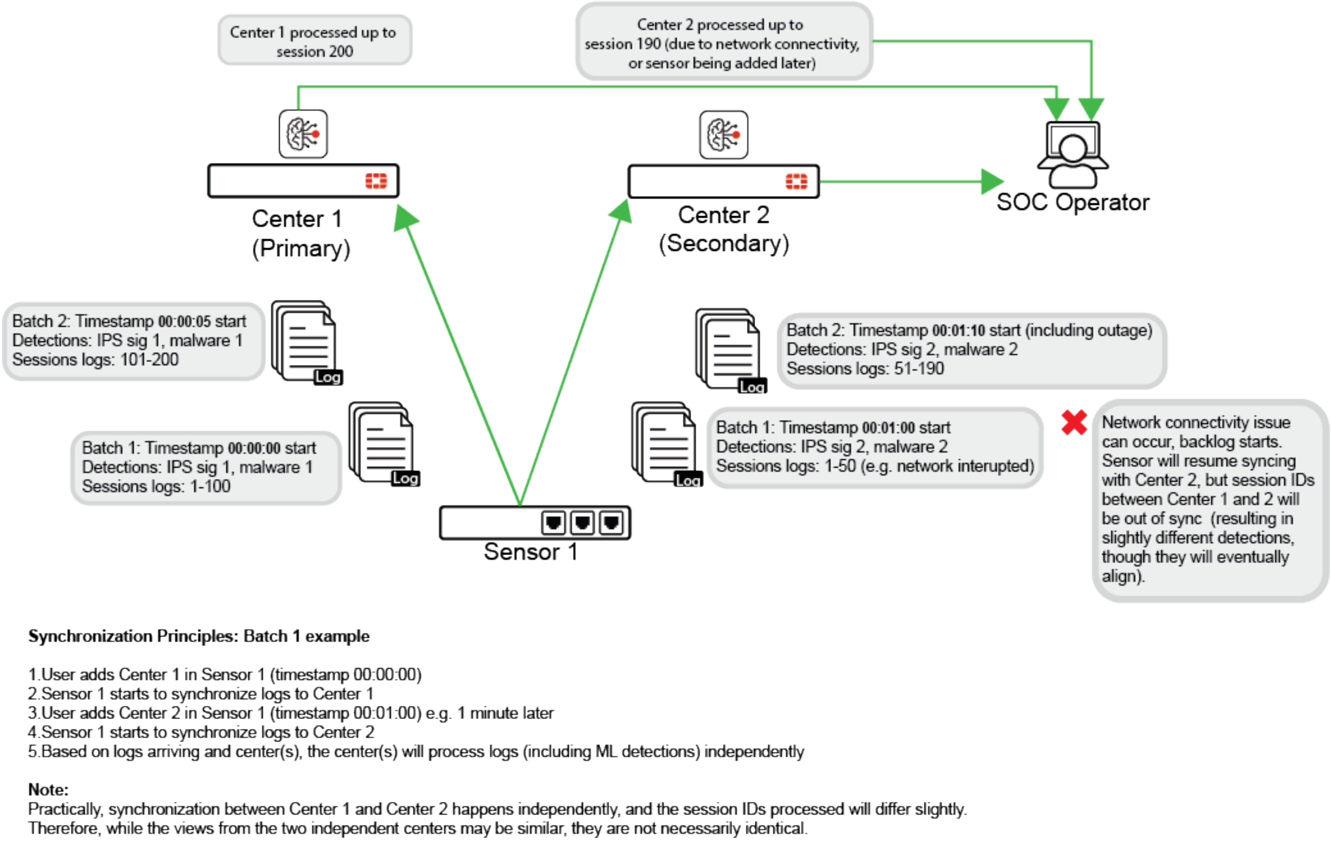
Center 1 processed up to session 200
Center 2 processed up to session 190 (due to network connectivity, or sensor being added later)
Center 1 (Primary)
Center 2 (Secondary)
SOC Operator
Batch 2: Timestamp 00:00:05 start Detections: IPS sig 1, malware 1 Sessions logs: 101-200
Batch 1: Timestamp 00:00:00 start Detections: IPS sig 1, malware 1 Sessions logs: 1-100
Batch 2: Timestamp 00:01:10 start (including outage) Detections: IPS sig 2, malware 2 Sessions logs: 51-190
Batch 1: Timestamp 00:01:00 start Detections: IPS sig 2, malware 2 Sessions logs: 1-50 (e.g. network interupted)
Network connectivity issue can occur, backlog starts. Sensor will resume syncing with Center 2, but session IDs between Center 1 and 2 will be out of sync (resulting in slightly different detections, though they will eventually align).
Log
Log
Log
Log
Sensor 1

Synchronization Principles: Batch 1 example

1.User adds Center 1 in Sensor 1 (timestamp 00:00:00)
2.Sensor 1 starts to synchronize logs to Center 1
3.User adds Center 2 in Sensor 1 (timestamp 00:01:00) e.g. 1 minute later
4.Sensor 1 starts to synchronize logs to Center 2
5.Based on logs arriving and center(s), the center(s) will process logs (including ML detections) independently

Note:
Practically, synchronization between Center 1 and Center 2 happens independently, and the session IDs processed will differ slightly. Therefore, while the views from the two independent centers may be similar, they are not necessarily identical.

Independent view of ML detections in dual center setup

Because logs are synchronized in batches, detections based on session logs synchronized to Centers in larger-scale deployments may be similar but not exactly the same, even when the ML baseline profile is identical in both features and timing. However, Centers can catch up based on these batches.

In smaller-scale deployments, detections can be identical due to fewer logs needing synchronization.

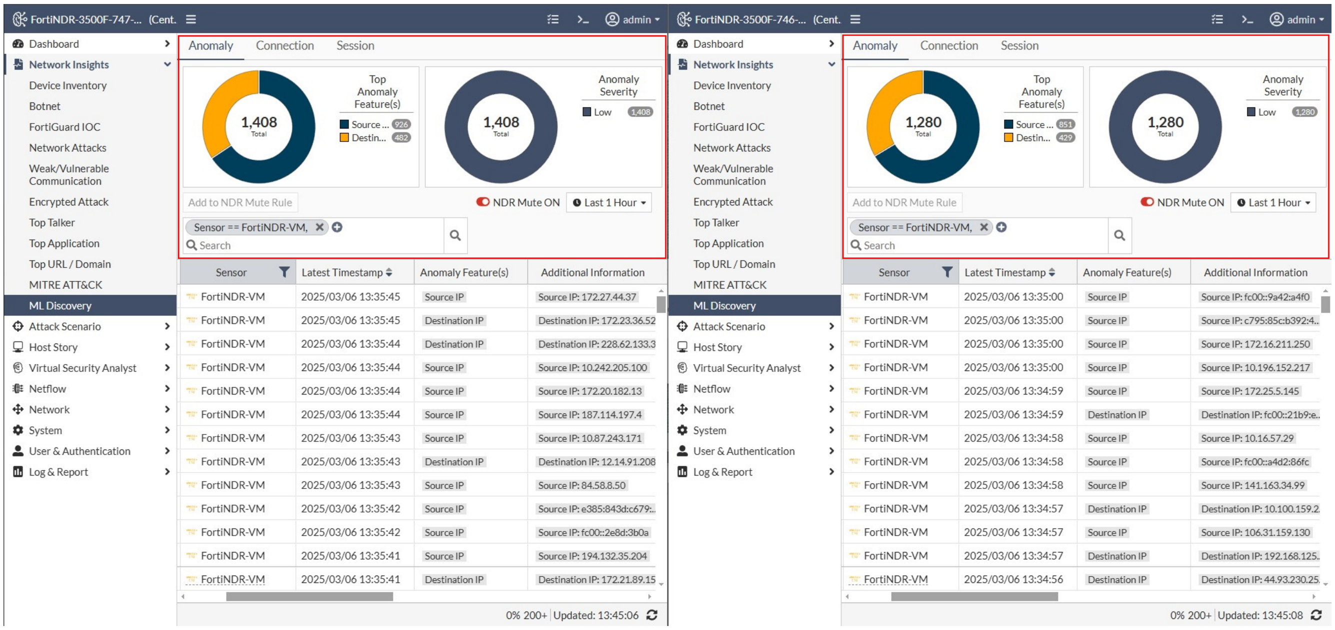

Add to NDR Mute Rule

NDR Mute ON

Last 1 Hour

Sensor == FortiNDR-VM,

Sensor	Latest Timestamp	Anomaly Feature(s)	Additional Information
FortiNDR-VM	2025/03/06 13:35:45	Source IP	Source IP: 172.27.44.37
FortiNDR-VM	2025/03/06 13:35:45	Destination IP	Destination IP: 172.23.36.52
FortiNDR-VM	2025/03/06 13:35:44	Destination IP	Destination IP: 228.62.133.3
FortiNDR-VM	2025/03/06 13:35:44	Source IP	Source IP: 10.242.205.100
FortiNDR-VM	2025/03/06 13:35:44	Source IP	Source IP: 172.20.182.13
FortiNDR-VM	2025/03/06 13:35:44	Source IP	Source IP: 187.114.197.4
FortiNDR-VM	2025/03/06 13:35:43	Source IP	Source IP: 10.87.243.171
FortiNDR-VM	2025/03/06 13:35:43	Destination IP	Destination IP: 12.14.91.208
FortiNDR-VM	2025/03/06 13:35:43	Source IP	Source IP: 84.58.8.50
FortiNDR-VM	2025/03/06 13:35:42	Source IP	Source IP: e385:843d:c679:..
FortiNDR-VM	2025/03/06 13:35:42	Source IP	Source IP: fc00::2e8d:3b0a
FortiNDR-VM	2025/03/06 13:35:41	Source IP	Source IP: 194.132.35.204
FortiNDR-VM	2025/03/06 13:35:41	Destination IP	Destination IP: 172.21.89.15

0% 200+ | Updated: 13:45:06

Add to NDR Mute Rule

NDR Mute ON

Last 1 Hour

Sensor == FortiNDR-VM,

Sensor	Latest Timestamp	Anomaly Feature(s)	Additional Information
FortiNDR-VM	2025/03/06 13:35:00	Source IP	Source IP: fc00::9a42:a4f0
FortiNDR-VM	2025/03/06 13:35:00	Source IP	Source IP: c795:85c:b392:4..
FortiNDR-VM	2025/03/06 13:35:00	Source IP	Source IP: 172.16.211.250
FortiNDR-VM	2025/03/06 13:35:00	Source IP	Source IP: 10.196.152.217
FortiNDR-VM	2025/03/06 13:34:59	Source IP	Source IP: 172.25.5.145
FortiNDR-VM	2025/03/06 13:34:59	Destination IP	Destination IP: fc00::21b9:e..
FortiNDR-VM	2025/03/06 13:34:58	Source IP	Source IP: 10.16.57.29
FortiNDR-VM	2025/03/06 13:34:58	Source IP	Source IP: fc00::a4d2:86fc
FortiNDR-VM	2025/03/06 13:34:58	Source IP	Source IP: 141.163.34.99
FortiNDR-VM	2025/03/06 13:34:57	Destination IP	Destination IP: 10.100.159.2.
FortiNDR-VM	2025/03/06 13:34:57	Source IP	Source IP: 106.31.159.130
FortiNDR-VM	2025/03/06 13:34:57	Destination IP	Destination IP: 192.168.125..
FortiNDR-VM	2025/03/06 13:34:56	Destination IP	Destination IP: 44.93.230.25.

0% 200+ | Updated: 13:45:08

Independent view of other detections in dual center setup

Because logs are synchronized in batches, detections will be matched. However, the detection counts will be close but not exactly the same, particularly when there is a backlog caused by network outages or other issues. Eventually, the centers will catch up with the backlog.

Figure: Dual center view comparing Primary and Secondary center - mismatch in Top URL/Domain counts

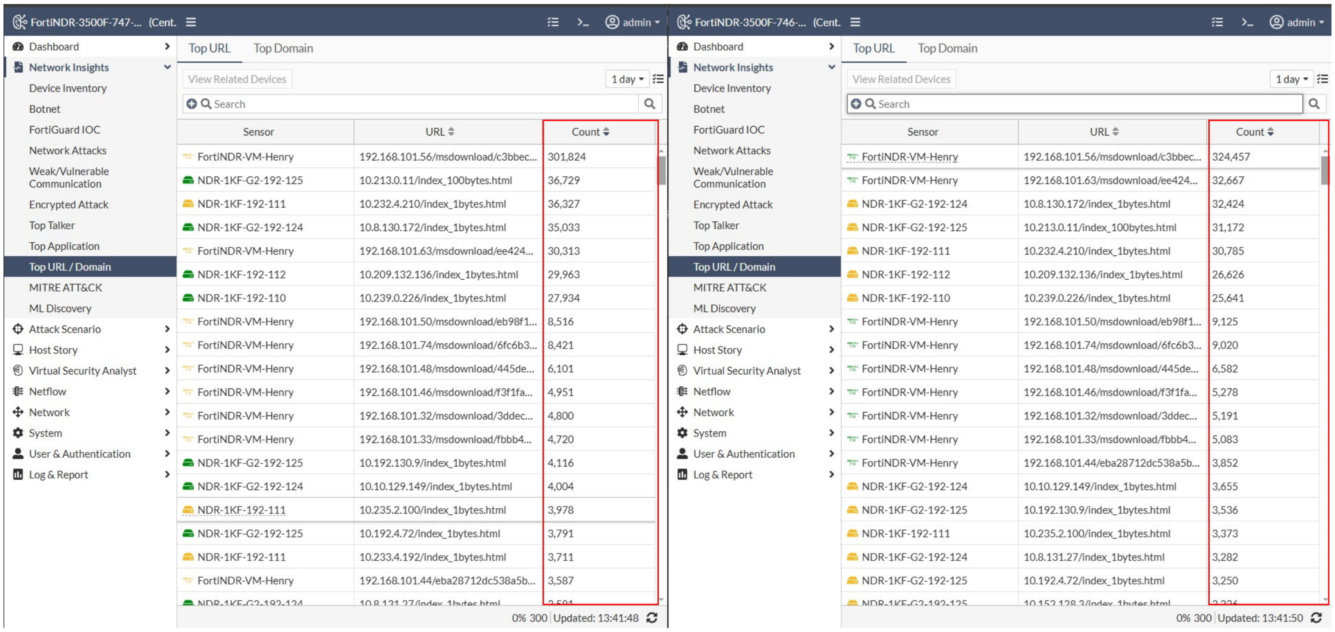

Sensor	URL	Count
FortiNDR-VM-Henry	192.168.101.56/msdownload/c3bbec...	301,824
NDR-1KF-G2-192-125	10.213.0.11/index_100bytes.html	36,729
NDR-1KF-192-111	10.232.4.210/index_1bytes.html	36,327
NDR-1KF-G2-192-124	10.8.130.172/index_1bytes.html	35,033
FortiNDR-VM-Henry	192.168.101.63/msdownload/ee424...	30,313
NDR-1KF-192-112	10.209.132.136/index_1bytes.html	29,963
NDR-1KF-192-110	10.239.0.226/index_1bytes.html	27,934
FortiNDR-VM-Henry	192.168.101.50/msdownload/eb98f1...	8,516
FortiNDR-VM-Henry	192.168.101.74/msdownload/6fc6b3...	8,421
FortiNDR-VM-Henry	192.168.101.48/msdownload/445de...	6,101
FortiNDR-VM-Henry	192.168.101.46/msdownload/f3f1fa...	4,951
FortiNDR-VM-Henry	192.168.101.32/msdownload/3ddec...	4,800
FortiNDR-VM-Henry	192.168.101.33/msdownload/fbbb4...	4,720
NDR-1KF-G2-192-125	10.192.130.9/index_1bytes.html	4,116
NDR-1KF-G2-192-124	10.10.129.149/index_1bytes.html	4,004
NDR-1KF-192-111	10.235.2.100/index_1bytes.html	3,978
NDR-1KF-G2-192-125	10.192.4.72/index_1bytes.html	3,791
NDR-1KF-192-111	10.233.4.192/index_1bytes.html	3,711
FortiNDR-VM-Henry	192.168.101.44/eba28712dc538a5b...	3,587
NDR-1KF-G2-192-124	10.8.131.27/index_1bytes.html	3,581

0% 300 Updated: 13:41:48

Sensor	URL	Count
FortiNDR-VM-Henry	192.168.101.56/msdownload/c3bbec...	324,457
FortiNDR-VM-Henry	192.168.101.63/msdownload/ee424...	32,667
NDR-1KF-G2-192-124	10.8.130.172/index_1bytes.html	32,424
NDR-1KF-G2-192-125	10.213.0.11/index_100bytes.html	31,172
NDR-1KF-192-111	10.232.4.210/index_1bytes.html	30,785
NDR-1KF-192-112	10.209.132.136/index_1bytes.html	26,626
NDR-1KF-192-110	10.239.0.226/index_1bytes.html	25,641
FortiNDR-VM-Henry	192.168.101.50/msdownload/eb98f1...	9,125
FortiNDR-VM-Henry	192.168.101.74/msdownload/6fc6b3...	9,020
FortiNDR-VM-Henry	192.168.101.48/msdownload/445de...	6,582
FortiNDR-VM-Henry	192.168.101.46/msdownload/f3f1fa...	5,278
FortiNDR-VM-Henry	192.168.101.32/msdownload/3ddec...	5,191
FortiNDR-VM-Henry	192.168.101.33/msdownload/fbbb4...	5,083
FortiNDR-VM-Henry	192.168.101.44/eba28712dc538a5b...	3,852
NDR-1KF-G2-192-124	10.10.129.149/index_1bytes.html	3,655
NDR-1KF-G2-192-125	10.192.130.9/index_1bytes.html	3,536
NDR-1KF-192-111	10.235.2.100/index_1bytes.html	3,373
NDR-1KF-G2-192-124	10.8.131.27/index_1bytes.html	3,282
NDR-1KF-G2-192-125	10.192.4.72/index_1bytes.html	3,250
NDR-1KF-G2-192-125	10.152.128.3/index_1bytes.html	3,226

0% 300 Updated: 13:41:50

Display statistics from the FortiNDR sensor

This section explains how system profiles and privileges are required to display sensor statistics on dashboard.

If you are using a wildcard LDAP administrator, please refer to *Creating remote wildcard administrators* in the *FortiNDR Administration Guide*.

To display the sensor statistics:

1. Ensure the sensor is added to the admin profile.

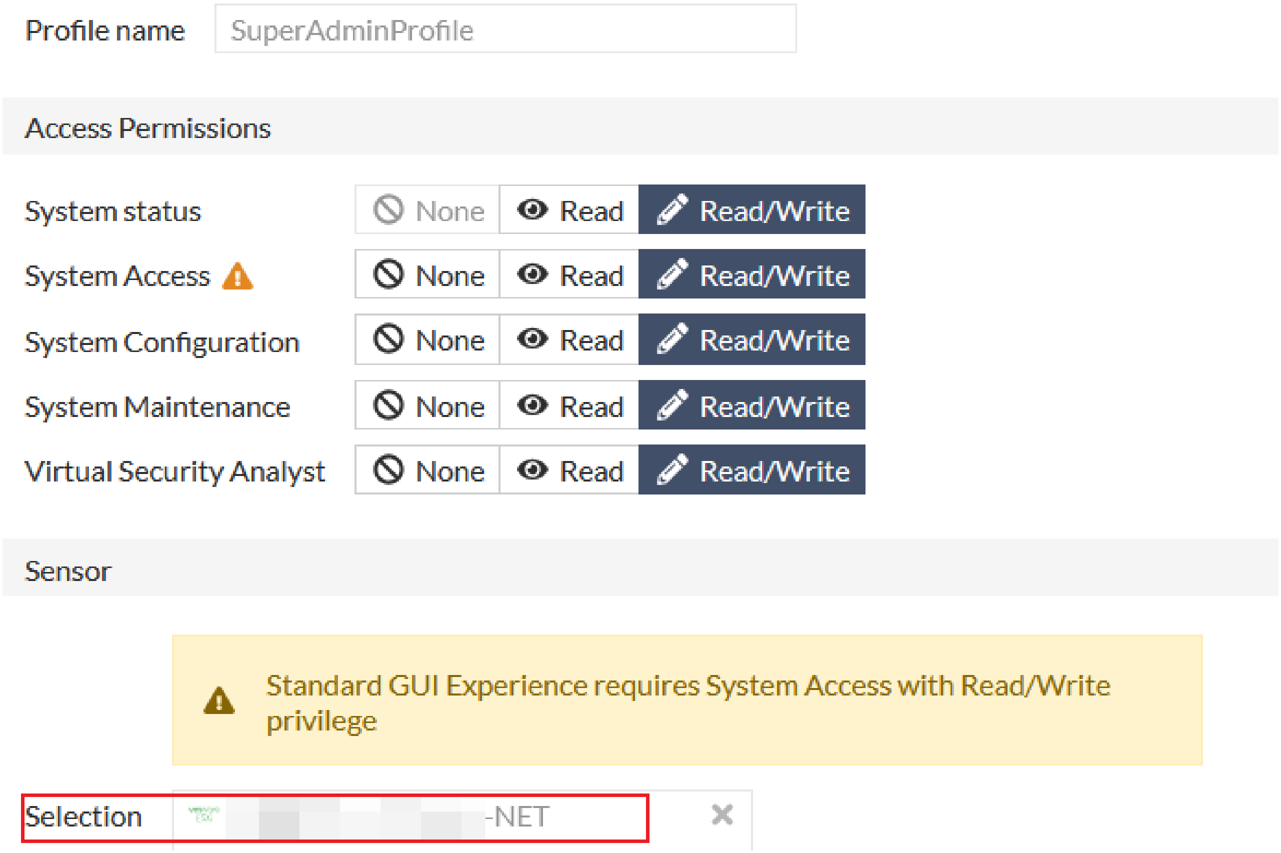

Profile name: SuperAdminProfile

Access Permissions

System status	None	Read	Read/Write
System Access ⚠	None	Read	Read/Write
System Configuration	None	Read	Read/Write
System Maintenance	None	Read	Read/Write
Virtual Security Analyst	None	Read	Read/Write

Sensor

⚠ Standard GUI Experience requires System Access with Read/Write privilege

Selection: [redacted]-NET ✕

2. Access the widget settings.

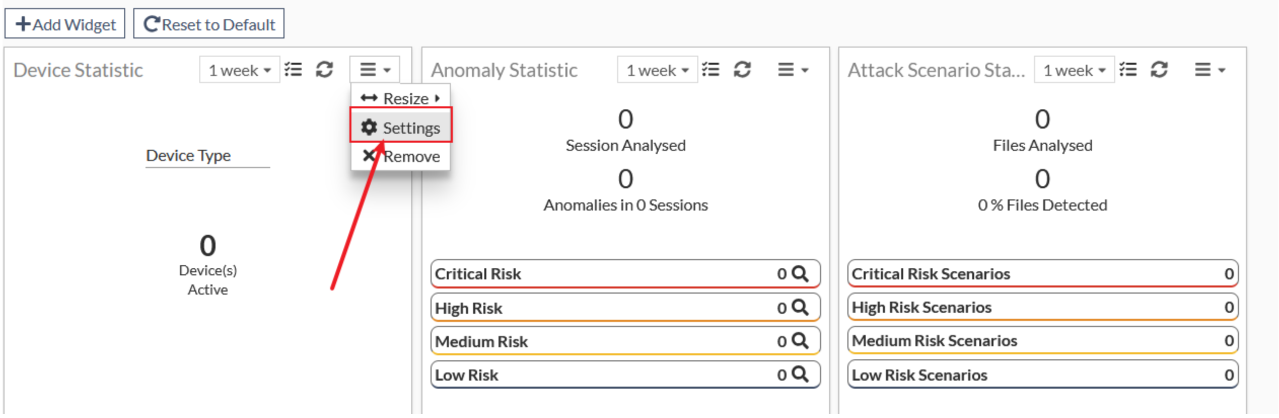

3. Select the sensor.

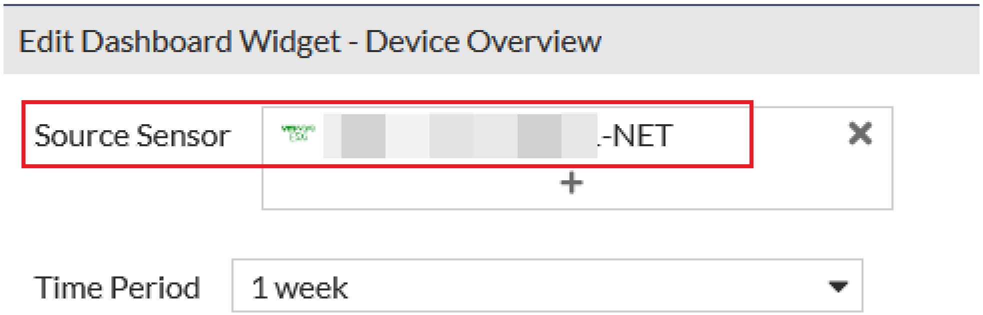

Edit Dashboard Widget - Device Overview

Source Sensor [redacted]-NET

Time Period 1 week

The statistics will appear once the sensor is properly configured in the widget settings.

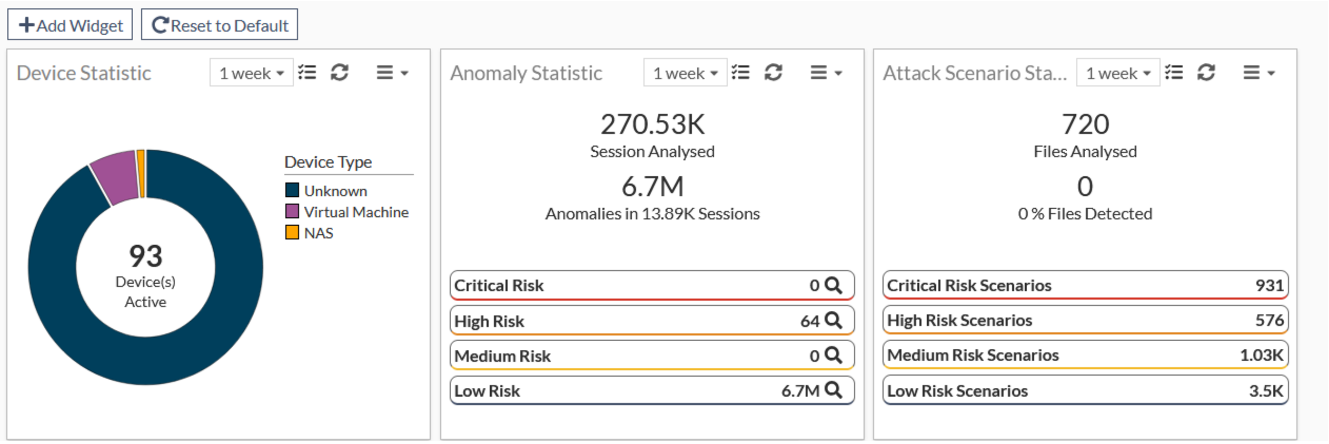

ERSPAN and RSPAN Support

ERSPAN (Encapsulated Remote SPAN) is a method used to monitor network traffic by copying data from one device and sending it to another device for analysis. It uses a special type of packaging (called GRE) to send this copied traffic across different networks, even if they are not directly connected.

RSPAN (Remote Switched Port Analyzer) also mirrors traffic but uses VLANs to send the copied data across switches.

ERSPAN support

FortiNDR On-Premise (version 7.6.2 and above) supports ERSPAN Type II traffic from VMware environments. It uses the IPS sniffer engine to receive and analyze this traffic. The sniffer interface is configured with a destination IP to receive ERSPAN traffic.

Once traffic arrives, FortiNDR unpacks it and logs session data, performs device inventory, and detects anomalies—just like it would with regular SPAN traffic.

To configure ERSPAN on VMWare ESXi:

1. Go to *Edit Port Mirroring Session > Edit properties*.
2. Set *Encapsulation type* to *ERSPAN Type II* and configure the rest of the settings as required.

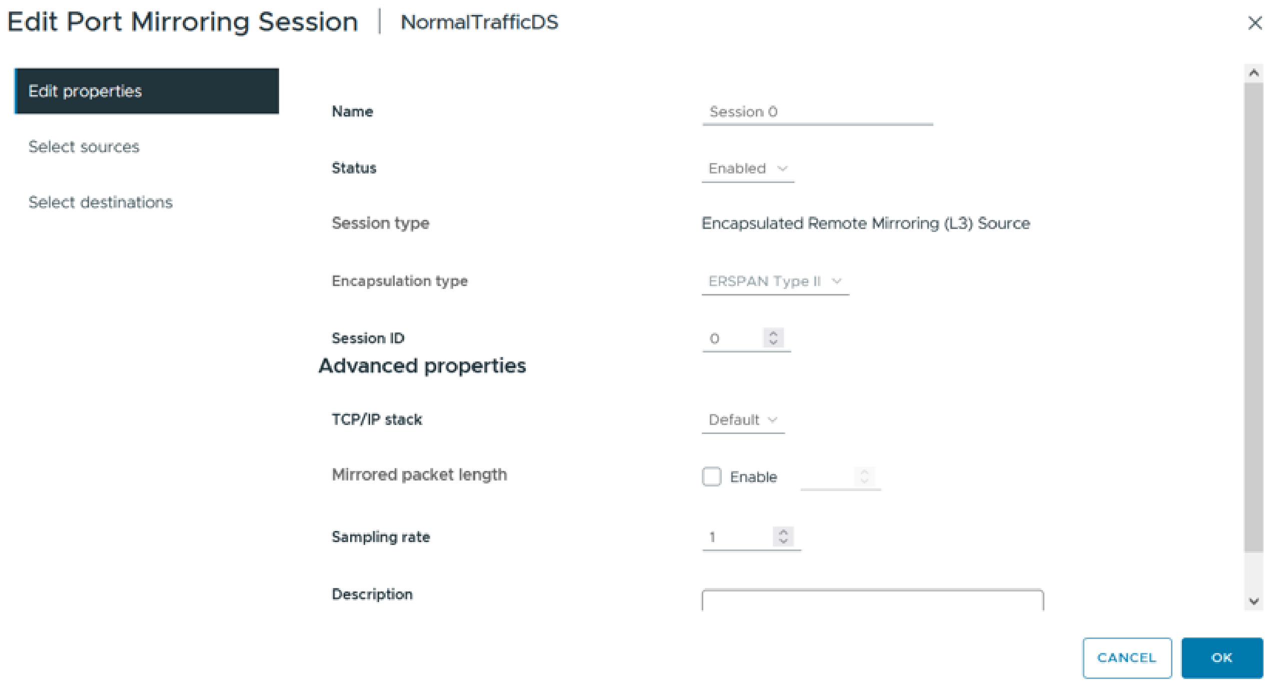

Edit Port Mirroring Session | NormalTrafficDS

Name	Session 0
Status	Enabled
Session type	Encapsulated Remote Mirroring (L3) Source
Encapsulation type	ERSPAN Type II
Session ID	0

Advanced properties

TCP/IP stack	Default
Mirrored packet length	□ Enable
Sampling rate	1
Description	

3. Configure the destination IP address.

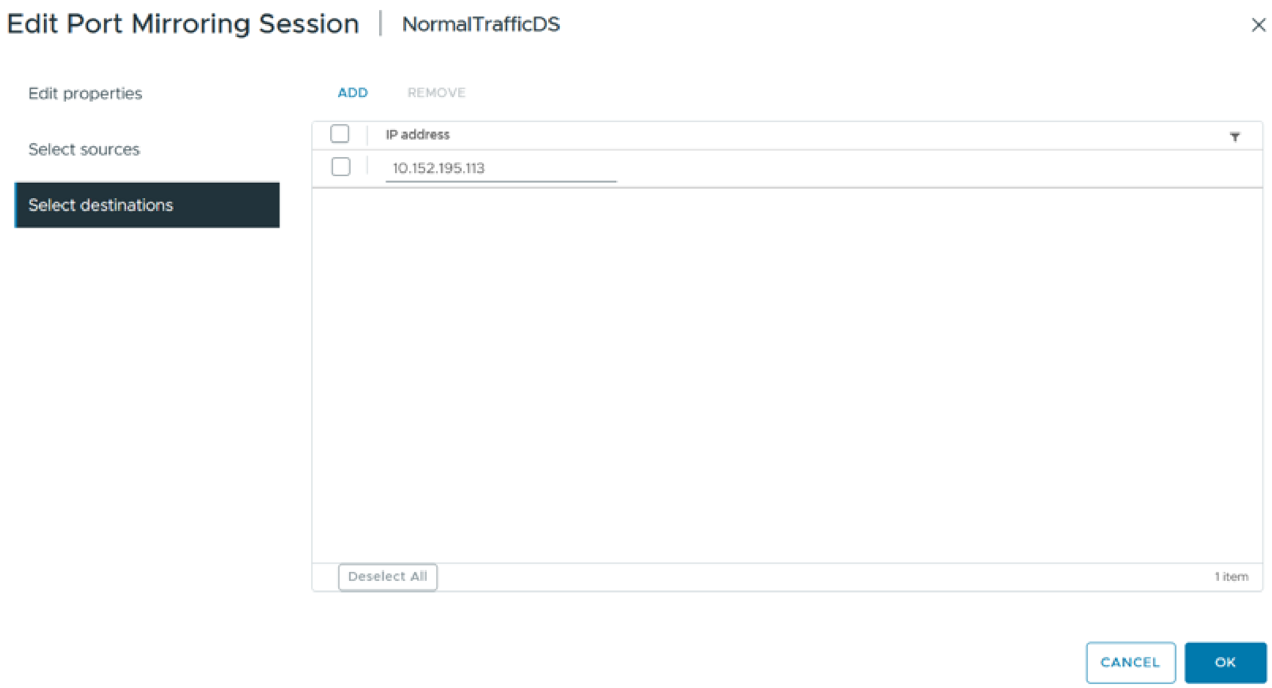

Edit Port Mirroring Session | NormalTrafficDS

	IP address
□	10.152.195.113

To configure the sniffer port:

1. In FortiNDR go to *Network* > *Interfaces*.
2. Configure the *port2(SNIFFER)* interface to the destination IP address you configured on VMWare ESXi.

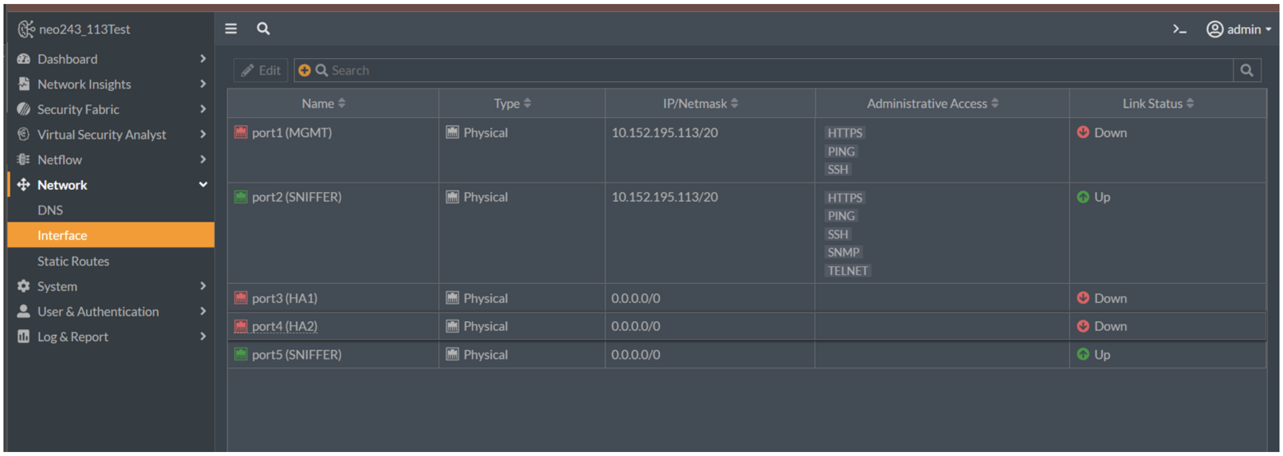

Name	Type	IP/Netmask	Administrative Access	Link Status
port1 (MGMT)	Physical	10.152.195.113/20	HTTPS PING SSH	Down
port2 (SNIFFER)	Physical	10.152.195.113/20	HTTPS PING SSH SNMP TELNET	Up
port3 (HA1)	Physical	0.0.0.0/0		Down
port4 (HA2)	Physical	0.0.0.0/0		Down
port5 (SNIFFER)	Physical	0.0.0.0/0		Up

3. To view the logs, go to *Log & Report > NDR Log*.

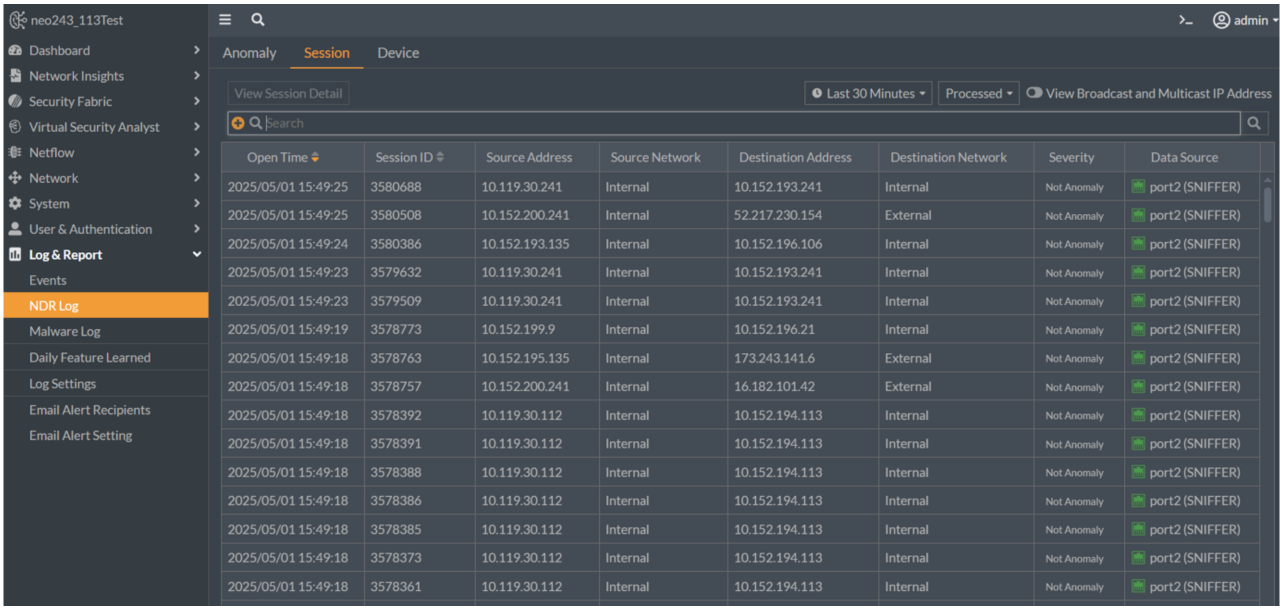

Open Time	Session ID	Source Address	Source Network	Destination Address	Destination Network	Severity	Data Source
2025/05/01 15:49:25	3580688	10.119.30.241	Internal	10.152.193.241	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:25	3580508	10.152.200.241	Internal	52.217.230.154	External	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:24	3580386	10.152.193.135	Internal	10.152.196.106	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:23	3579632	10.119.30.241	Internal	10.152.193.241	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:23	3579509	10.119.30.241	Internal	10.152.193.241	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:19	3578773	10.152.199.9	Internal	10.152.196.21	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578763	10.152.195.135	Internal	173.243.141.6	External	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578757	10.152.200.241	Internal	16.182.101.42	External	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578392	10.119.30.112	Internal	10.152.194.113	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578391	10.119.30.112	Internal	10.152.194.113	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578388	10.119.30.112	Internal	10.152.194.113	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578386	10.119.30.112	Internal	10.152.194.113	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578385	10.119.30.112	Internal	10.152.194.113	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578373	10.119.30.112	Internal	10.152.194.113	Internal	Not Anomaly	port2 (SNIFFER)
2025/05/01 15:49:18	3578361	10.119.30.112	Internal	10.152.194.113	Internal	Not Anomaly	port2 (SNIFFER)

You can monitor the device in *Network Insights > Device Inventory*.

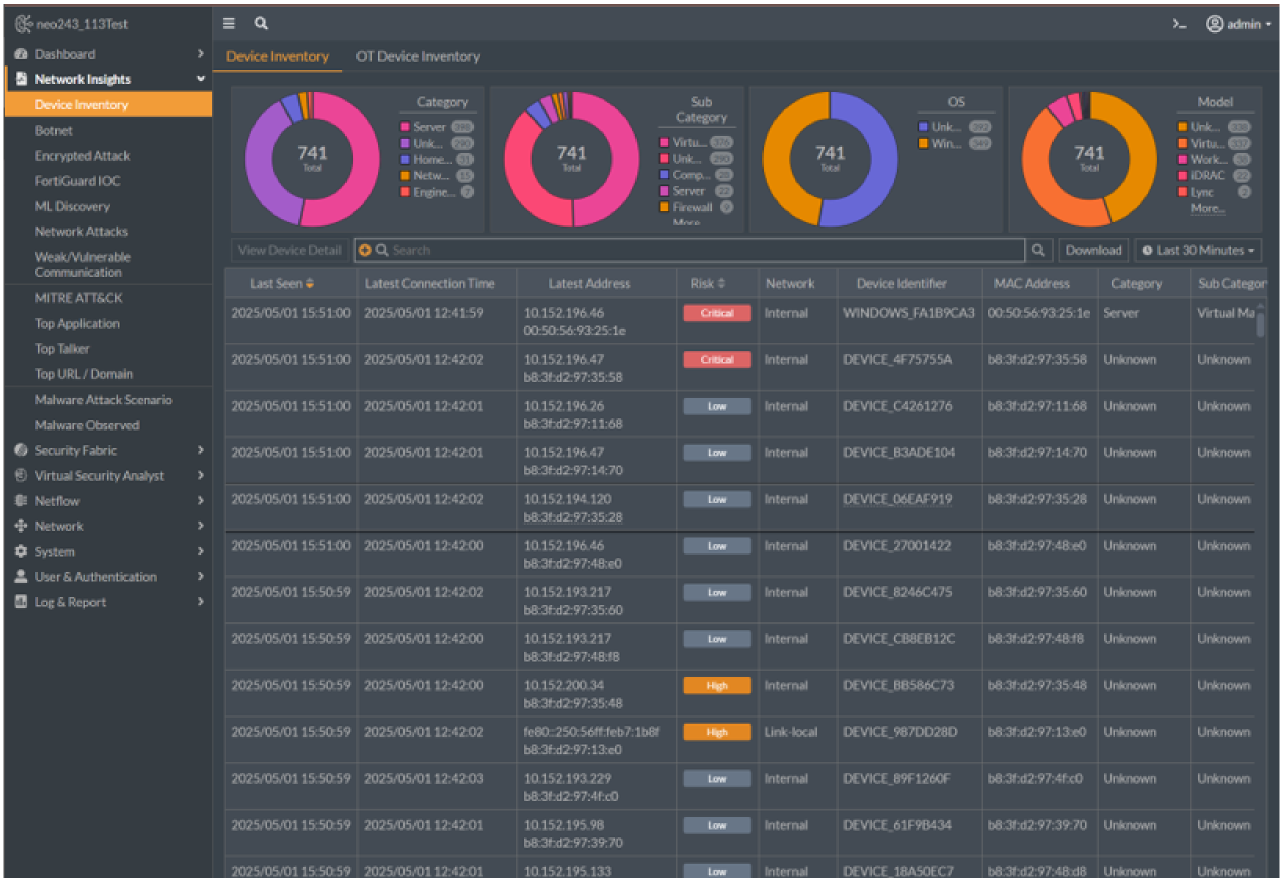

Last Seen	Latest Connection Time	Latest Address	Risk	Network	Device Identifier	MAC Address	Category	Sub Categor
2025/05/01 15:51:00	2025/05/01 12:41:59	10.152.196.46 00:50:56:93:25:1e	Critical	Internal	WINDOWS_FA1B9CA3	00:50:56:93:25:1e	Server	Virtual Ma
2025/05/01 15:51:00	2025/05/01 12:42:02	10.152.196.47 b8:3f:d2:97:35:58	Critical	Internal	DEVICE_4F75755A	b8:3f:d2:97:35:58	Unknown	Unknown
2025/05/01 15:51:00	2025/05/01 12:42:01	10.152.196.26 b8:3f:d2:97:11:68	Low	Internal	DEVICE_C4261276	b8:3f:d2:97:11:68	Unknown	Unknown
2025/05/01 15:51:00	2025/05/01 12:42:01	10.152.196.47 b8:3f:d2:97:14:70	Low	Internal	DEVICE_B3ADE104	b8:3f:d2:97:14:70	Unknown	Unknown
2025/05/01 15:51:00	2025/05/01 12:42:02	10.152.194.120 b8:3f:d2:97:35:28	Low	Internal	DEVICE_06EAF919	b8:3f:d2:97:35:28	Unknown	Unknown
2025/05/01 15:51:00	2025/05/01 12:42:00	10.152.196.46 b8:3f:d2:97:48:e0	Low	Internal	DEVICE_27001422	b8:3f:d2:97:48:e0	Unknown	Unknown
2025/05/01 15:50:59	2025/05/01 12:42:02	10.152.193.217 b8:3f:d2:97:35:60	Low	Internal	DEVICE_8246C475	b8:3f:d2:97:35:60	Unknown	Unknown
2025/05/01 15:50:59	2025/05/01 12:42:00	10.152.193.217 b8:3f:d2:97:48:f8	Low	Internal	DEVICE_CB8EB12C	b8:3f:d2:97:48:f8	Unknown	Unknown
2025/05/01 15:50:59	2025/05/01 12:42:00	10.152.200.34 b8:3f:d2:97:35:48	High	Internal	DEVICE_BB586C73	b8:3f:d2:97:35:48	Unknown	Unknown
2025/05/01 15:50:59	2025/05/01 12:42:02	fe80::250:56ff:feb7:1b8f b8:3f:d2:97:13:e0	High	Link-local	DEVICE_987DD28D	b8:3f:d2:97:13:e0	Unknown	Unknown
2025/05/01 15:50:59	2025/05/01 12:42:03	10.152.193.229 b8:3f:d2:97:4f:c0	Low	Internal	DEVICE_89F1260F	b8:3f:d2:97:4f:c0	Unknown	Unknown
2025/05/01 15:50:59	2025/05/01 12:42:01	10.152.195.98 b8:3f:d2:97:39:70	Low	Internal	DEVICE_61F9B434	b8:3f:d2:97:39:70	Unknown	Unknown
2025/05/01 15:50:59	2025/05/01 12:42:01	10.152.195.133	Low	Internal	DEVICE_18A50EC7	b8:3f:d2:97:48:d8	Unknown	Unknown

As well as *Dashboard > NDR Overview*.

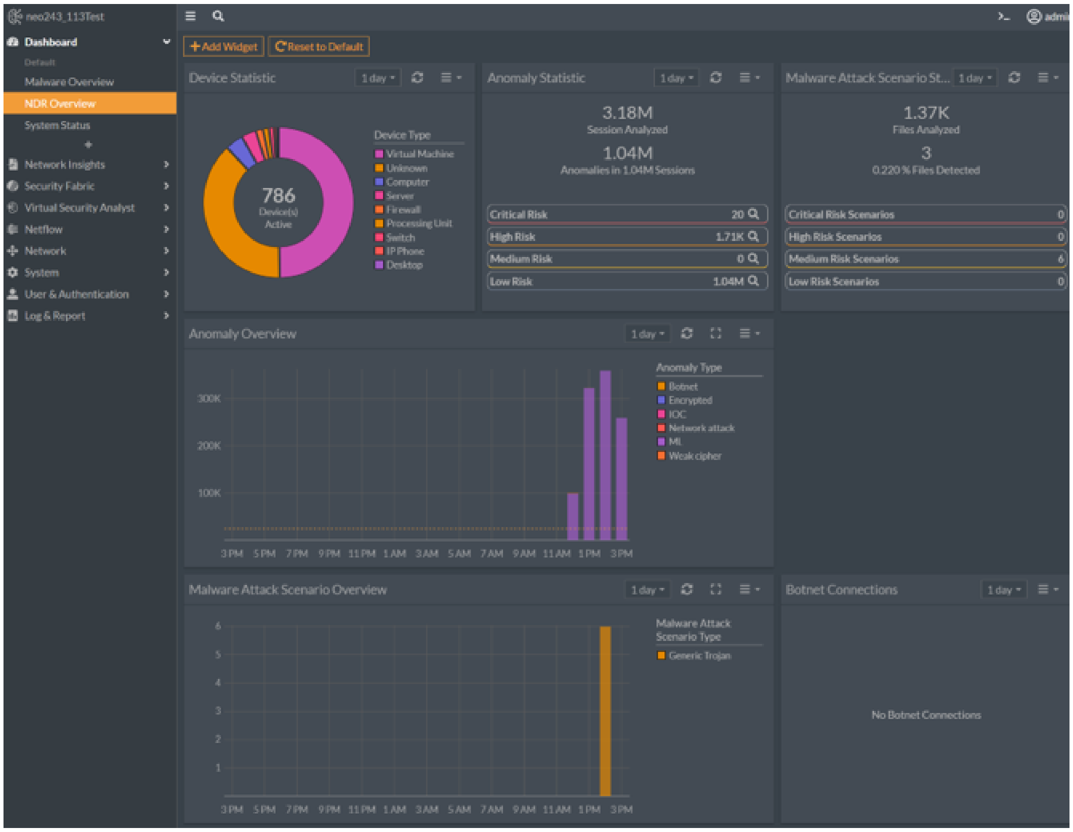

FortiNDR RSPAN support

FortiNDR also supports RSPAN, which uses VLAN IDs to encapsulate mirrored traffic. The IPS engine in FortiNDR can parse VLAN IDs and analyze the traffic accordingly.

FortiNDR OnPrem uses Fortinet's IPS engine, which supports GRE, VLAN, and VXLAN encapsulation.

Example:

RSPAN uses VLAN to encapsulate traffic.

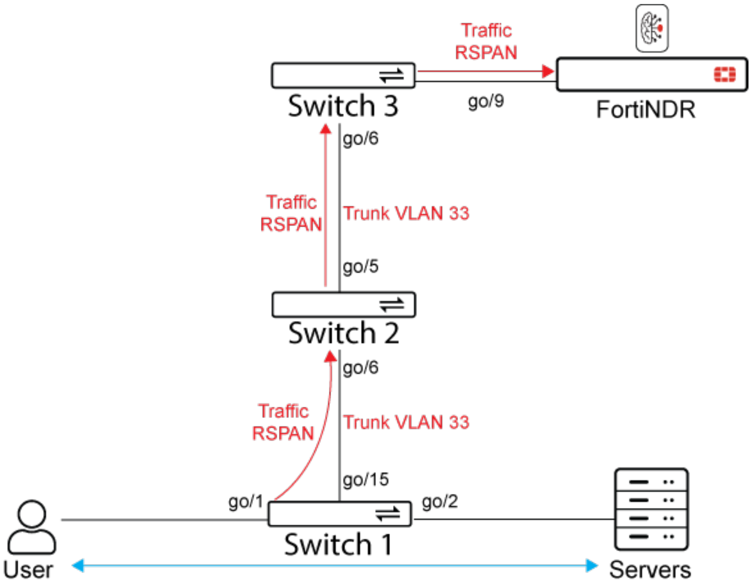

FortiNDR IPS engine parses the VLAN IDs.

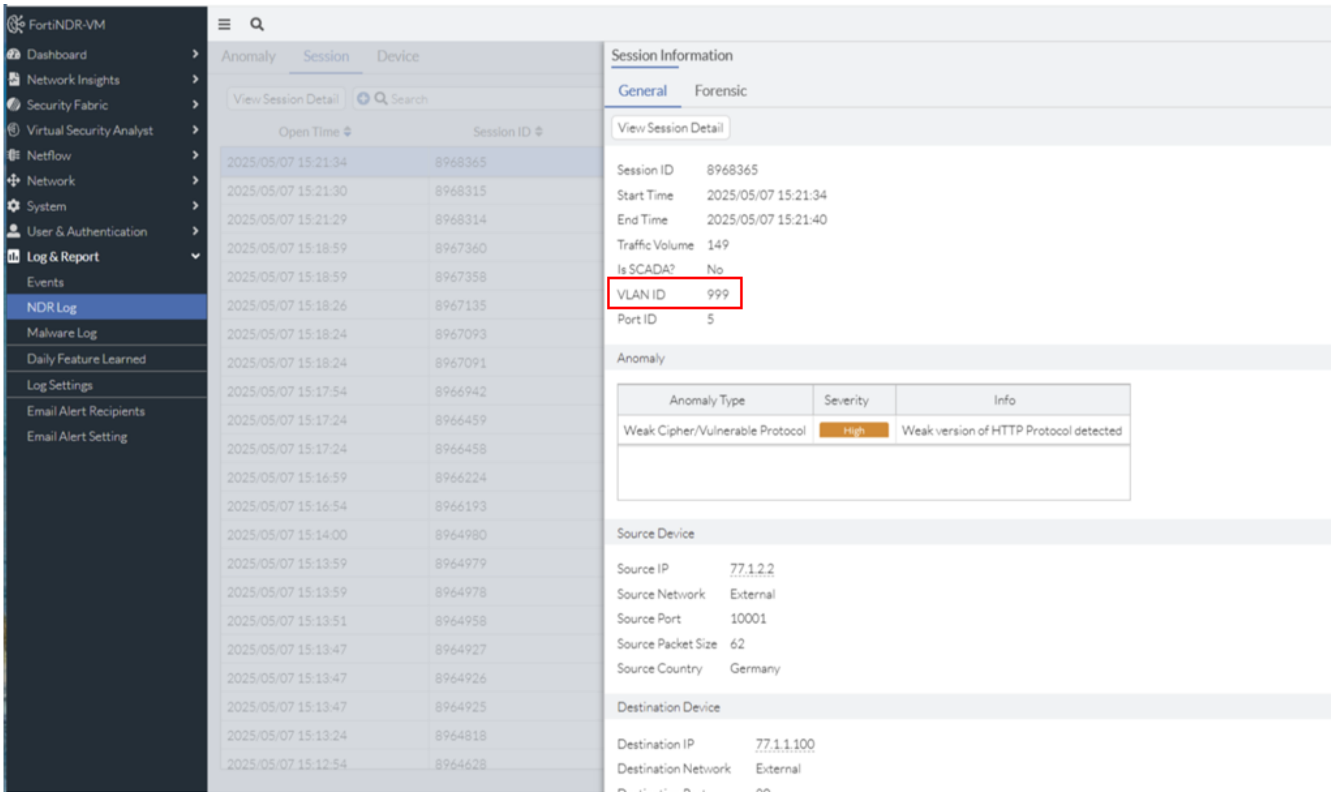

Open Time	Session ID
2025/05/07 15:21:34	8968365
2025/05/07 15:21:30	8968315
2025/05/07 15:21:29	8968314
2025/05/07 15:18:59	8967360
2025/05/07 15:18:59	8967358
2025/05/07 15:18:26	8967135
2025/05/07 15:18:24	8967093
2025/05/07 15:18:24	8967091
2025/05/07 15:17:54	8966942
2025/05/07 15:17:24	8966459
2025/05/07 15:17:24	8966458
2025/05/07 15:16:59	8966224
2025/05/07 15:16:54	8966193
2025/05/07 15:14:00	8964980
2025/05/07 15:13:59	8964979
2025/05/07 15:13:59	8964978
2025/05/07 15:13:51	8964958
2025/05/07 15:13:47	8964927
2025/05/07 15:13:47	8964926
2025/05/07 15:13:47	8964925
2025/05/07 15:13:24	8964818
2025/05/07 15:12:54	8964628

Session ID	8968365
Start Time	2025/05/07 15:21:34
End Time	2025/05/07 15:21:40
Traffic Volume	149
Is SCADA?	No
VLAN ID	999
Port ID	5

Anomaly

Anomaly Type	Severity	Info
Weak Cipher/Vulnerable Protocol	High	Weak version of HTTP Protocol detected

Source Device

Source IP	77.1.2.2
Source Network	External
Source Port	10001
Source Packet Size	62
Source Country	Germany

Destination Device

Destination IP	77.1.1.100
Destination Network	External

Separating ports for large scale deployments

This best practice outlines recommended port utilization configurations for FortiNDR Center (FNDR-3600G) and Sensor (FNDR-1000F) in large-scale deployments. For most environments, reusing the management port for both management and service functions (such as Center-Sensor communication on the Center node and flow collection on the Sensor node) is sufficient. For larger or high-load deployments, you can separate management and service functions across different ports instead of reusing the same port for both.

Recommendations

- Use the default port configuration as specified in the documentation. See Initial setup.
- If you need to separate management traffic from Sensor/Center synchronization traffic, use Port2 for management on models 1000F and 3600G.
- Avoid using data ports for Netflow and logging. Ideally, use port1 for simplicity and alignment with documentation.

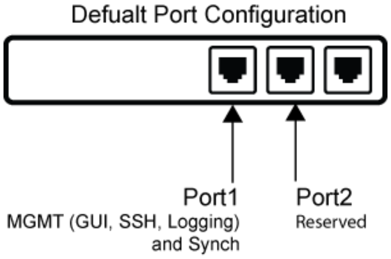
Defualt Port Configuration
Port1
MGMT (GUI, SSH, Logging) and Synch
Port2
Reserved

Center to Sensor Communication

Center-to-Sensor synchronization traffic must use port1 by default on both FNR-3600G and Sensor 1000F. If you need to separate management traffic, configure a static route pointing to port2 for management and ensure that port1 and port2 use different subnets.

In this setup, synchronization traffic remains on port1 while management traffic is routed through port2. Internal testing has confirmed that both the Center and Sensor can use port2 for management and port1 for synchronization, provided the subnets differ.

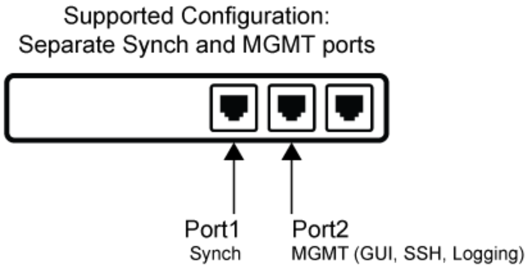

In NDR 1000F models, turn the sniffer function off with the CLI: `config system interface`.

Sensor

The following image shows the Sensor configured with two active interfaces (port1 and port2), enabling the option to separate management traffic from other functions by using port2 instead of relying solely on port1.

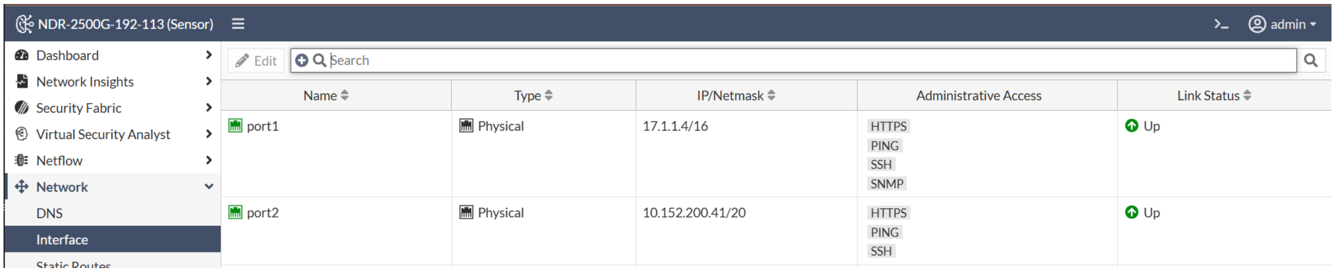

Name	Type	IP/Netmask	Administrative Access	Link Status
port1	Physical	17.1.1.4/16	HTTPS PING SSH SNMP	Up
port2	Physical	10.152.200.41/20	HTTPS PING SSH	Up

This image shows the Sensor configured with a static route that directs management traffic through port2, separating management traffic from center-to-sensor synchronization that remains on port1.

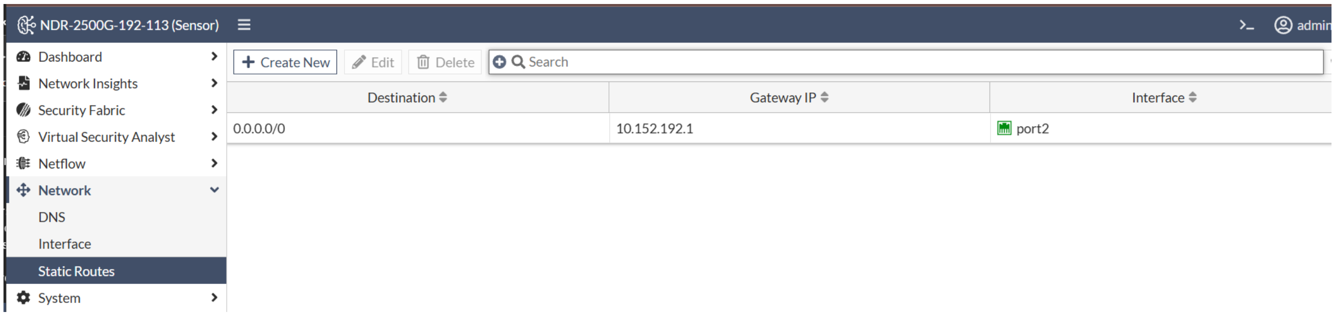

Destination	Gateway IP	Interface
0.0.0.0/0	10.152.192.1	port2

The Sensor is successfully registered to the Center, which is necessary after configuring separate ports for management and synchronization.

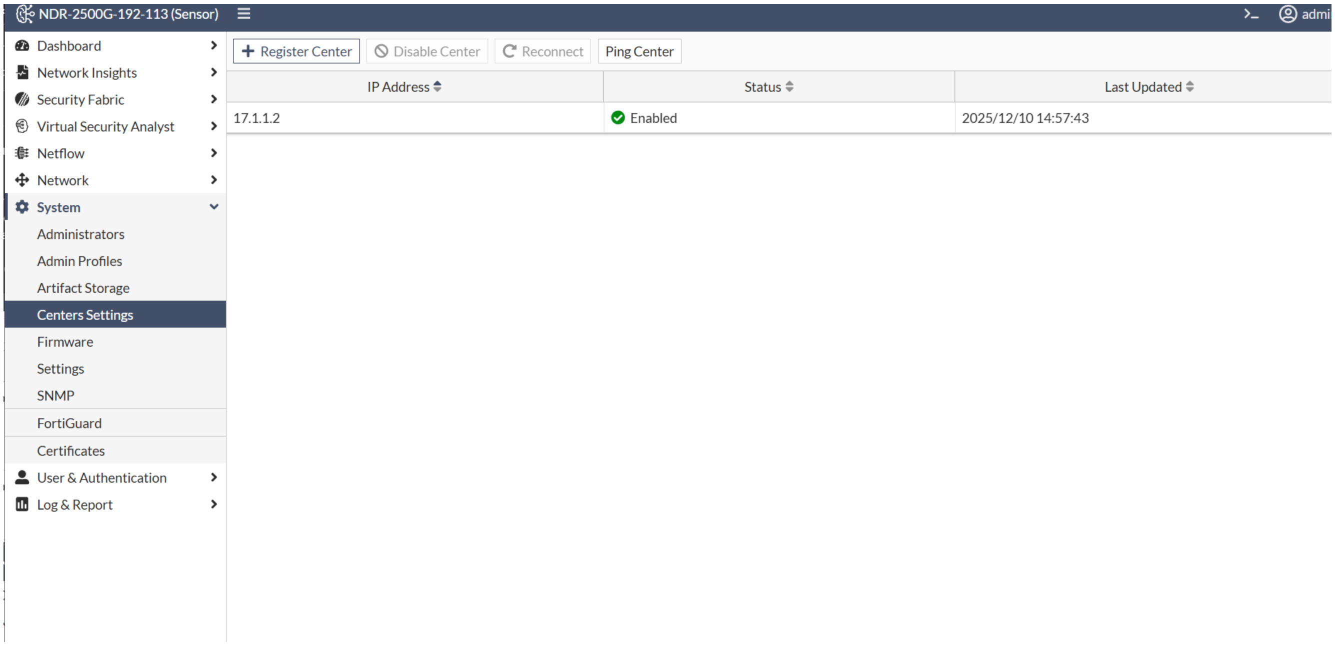

IP Address	Status	Last Updated
17.1.1.2	✓ Enabled	2025/12/10 14:57:43

Center

This image shows the Center with multiple interfaces enabled, where port1 is used for management and port2 can be configured for separate management traffic, isolating management from synchronization.

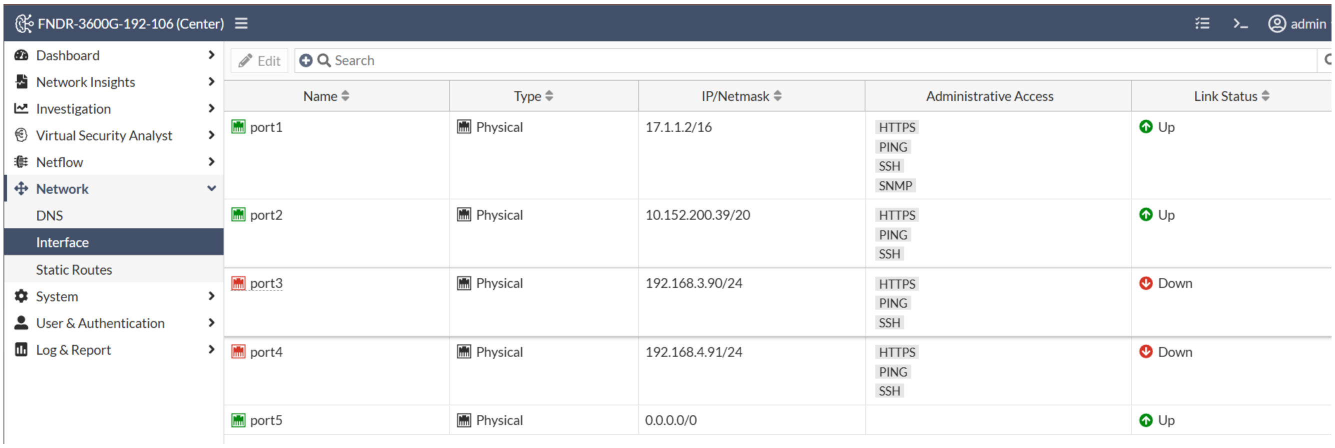

Name	Type	IP/Netmask	Administrative Access	Link Status
port1	Physical	17.1.1.2/16	HTTPS PING SSH SNMP	Up
port2	Physical	10.152.200.39/20	HTTPS PING SSH	Up
port3	Physical	192.168.3.90/24	HTTPS PING SSH	Down
port4	Physical	192.168.4.91/24	HTTPS PING SSH	Down
port5	Physical	0.0.0.0/0		Up

This image shows the Center configured with a static route that uses port2 for management traffic, separating management traffic from Center-to-Sensor synchronization that remains on port1.

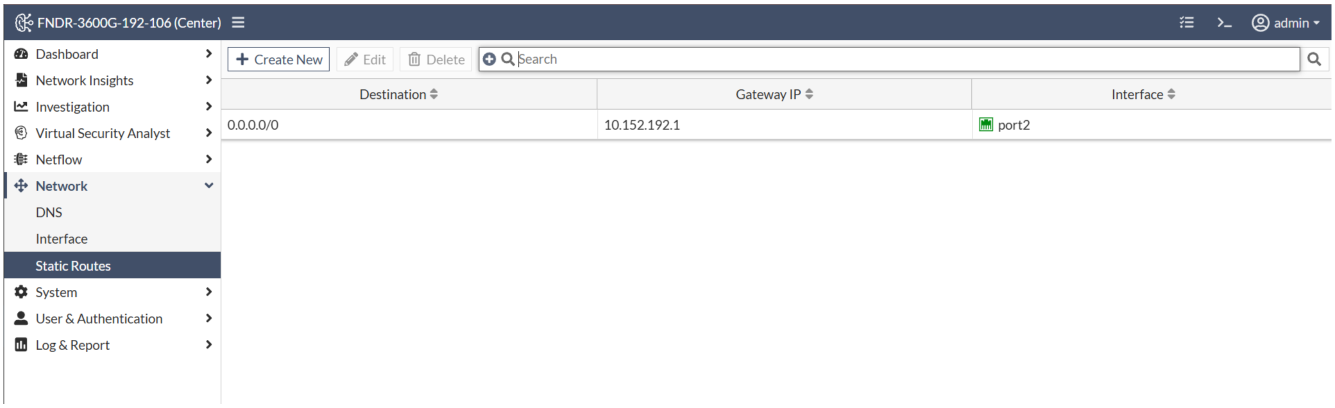

Destination	Gateway IP	Interface
0.0.0.0/0	10.152.192.1	port2

This image shows the Center confirming that the Sensor is connected and operational, which verifies successful communication after configuring separate ports for management and synchronization.

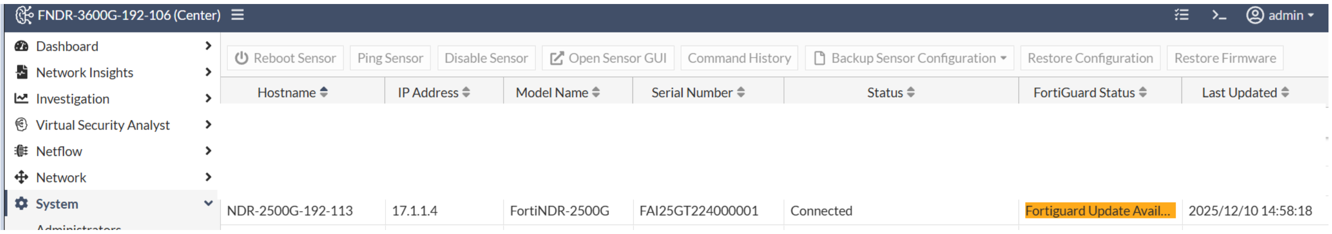

Hostname	IP Address	Model Name	Serial Number	Status	FortiGuard Status	Last Updated
NDR-2500G-192-113	17.1.1.4	FortiNDR-2500G	FAI25GT224000001	Connected	Fortiguard Update Avail...	2025/12/10 14:58:18

Sensor Netflow and logging collection

The default recommendation is to use port1 for NetFlow telemetry. Alternatively, the Sensor can receive NetFlow on port2, port3, or port4, and packets sent to these ports will be accepted and processed. However, these data ports also run a sniffer process that attempts to parse UDP packets. If NetFlow data is collected on sniffer ports, you must disable the sniffer function on that port.

For sending logs to external sources (e.g., Syslog or FortiAnalyzer/FortiSIEM), the management port is used to transmit external log traffic. The traffic will follow the system's route table, so Syslog will use the same gateway as management.

Similar to disabling the sniffer function when a port is used for management, you must also disable the sniffer function on any port receiving NetFlow data. You can do this using the CLI command: execute ndrd off.

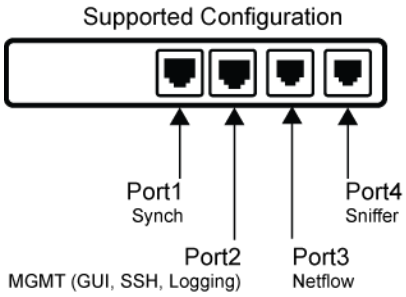

www.fortinet.com

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.