30 October 2022

Horizon NDR

Indicator Management

User Guide

Check Point
SOFTWARE TECHNOLOGIES LTD.

Introduction

The Check Point Horizon NDR operational concept is composed of the following flows:

- Network traffic is analyzed by sensors, which generate analytical results in the form of logs
- Logs are transmitted to the NDR cloud for storage and analysis
- Behavioral Analytics AI engines process the logs and generate analytical conclusions
- Human analysts are provided with event visualization tools for further data comprehension
- Data anomalies are incriminated through correlation with ThreatCloud intelligence and application risk scoring
- Analytical conclusions are published in the form of threat indicators and tags
- **Threat indicators are also received from third party threat intelligence sources**
- **Indicators are applied by enforcement points and matched to network traffic, taking DETECT or PREVENT action**

This article focuses on the Intel capabilities of the Horizon NDR application.

Table of Contents

Indicators

An indicator is a typed pattern that is used to characterize network traffic flows or elements, in a specific context. For example, traffic to or from a host might be identified using the host's IP address. File hashes can be used to identify files transferred over the network.

The pattern (also referred to as an observable or match condition) is defined by type and value. For example, an IP type indicator is defined using IPv4 dotted notation (e.g. 10.0.3.8), whereas an IP_RANGE type could support either CIDR (e.g. 10.0.3.0/24) or address range (10.0.3.0-10.0.3.255).

In addition to the match condition, the indicator defines context as a set of metadata attributes. These attributes describe the indicator's source, the time when it was created and last updated, and the flows and entities it is intended to describe. In addition it may prescribe actionable instructions to enforcement devices that will match the indicator to network traffic, or to network visualization tools such as NDR's Threat Topology analytics view.

Indicator Value

In the example below, we see four indicators (out of 151 defined on the Horizon NDR domain). The first two were created manually by a Horizon NDR user, the third pulled automatically from an external Threat Intelligence Platform (Anomali), and the fourth was received from the Behavioral input feed, which represents the output of the Horizon NDR Behavioral Analytics AI engines.

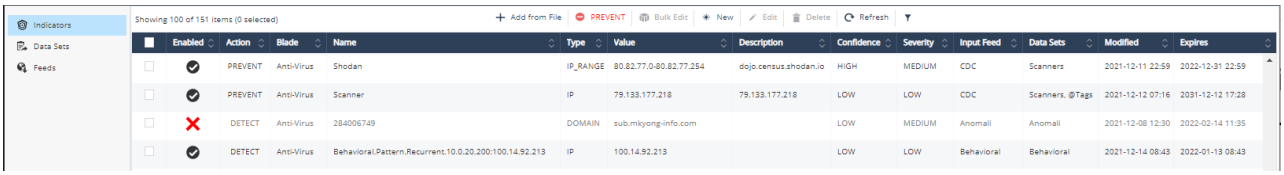

Indicators are keyed by **Input Feed**; **Type**; and **Value**. Thus if the same match condition (i.e. Type, Value) is received from multiple sources (identified as Input Feeds), NDR Intel will retain multiple copies, each with its associated metadata.

In the example below, filtering on a single value brings up two copies; the second instance was automatically generated using the Behavioral.Port AI engine, which identified an anomalous burst of emails from a source with unknown reputation. The analyst disabled that indicator, and tagged the host with its domain name. Both copies are retained, as they associate different metadata attributes with the identified host. The next time the Behavioral engine detects an anomaly on this IP, it will refresh the modification timestamp for the indicator whose input feed is "Behavioral".

Indicator types are intended to be interpreted by enforcement points, and are therefore mapped to Check Point "software blades". Each blade provides support for a set of indicator types. For example, hash types (MD5, HASH_SHA1, HASH_SHA256) are only applicable to the "Anti-Virus" blade, whereas IP can be matched by both "Anti-Virus" and "Anti-Bot", designating pre-infection and post-infection behavior, respectively.

Currently, the following indicator types are supported (Restrictions apply for some Check Point gateway versions):

- Indicator types supported by Anti-Virus and Anti-Bot blades, defined in sk132193 (Custom Intelligence Feeds)
- Snort - supported by IPS blade on some Check Point gateway versions
- Experimental (YARA, PCAP, REGEX, TEXT) - available only by request
- None - used in cases where NDR Intel is used as a repository, e.g. for defining threat visualization tag values

Indicator Metadata

Each indicator is associated with a **Name** attribute, linking the indicator to its semantics. Multiple indicators may be tagged with the same Name, supporting a search for all indicators with that Name via the filtering facility:

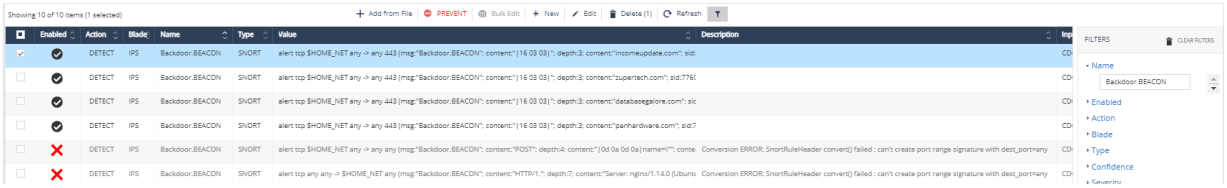

For Anti-Virus and Anti-Bot blade matches, the Name is incorporated in the “protection_name” and “observable_name” log fields.

The (optional) **Description** attribute provides additional free-text context to the indicator. The description is also incorporated in Anti-Virus and Anti-Bot blade logs.

Additional metadata attributes that can be entered for an indicator may include:

- **Status** (ENABLED/DISABLED) - controls whether the indicator is included in data set output feeds (assuming it’s not expired)
- **Action** (DETECT/PREVENT) - a "hint" to the enforcement engine; determines the URL on which a data set serves the indicator
- **Expiration** - date/time after which the indicator is implicitly disabled
  (Note: Expired indicators are automatically deleted from NDR Intel after 14 days.)
- **Confidence/Severity** - included in log records for some blades' indicator match logs
- **CVE/Mitre Tactic/Mitre Technique** - used for linking the indicator to public references
- **Data Sets** - zero, one, or more output feeds that the indicator is published on (assuming it's not disabled or expired)

In addition, the system maintains the following attributes which cannot be edited by the user:

- **Created/Created_by/Modified/Modified_by** - user name and timestamp

Indicator Management

NDR Intel implements a simple indicator management concept that allows analysts to manage large quantities of indicators. These operations can be invoked from the portal, or via API:

- **New** - create a new indicator, specifying match condition and metadata
- **Add from File** - bulk-add import into NDR Intel from a file
- **Delete** - delete the selected indicator(s) from the NDR Intel repository
- **Edit** - modify the attributes of a single indicator
- **Bulk Edit** - modify attributes for multiple selected indicators
- **PREVENT** - easy to use "red button" for changing indicators' Action to PREVENT

A paging paradigm is used to display indicators. You can select the page size, the default is to display 20 indicators at a time. Clicking the "Select All" check box will select only the displayed indicators. The bottom of the screen shows which page is currently displayed, and allows you to select a different page.

The analyst uses the Filtering option to constrain the set of indicators displayed. Sorting can be performed by clicking on an individual column heading. It is then possible to select the currently displayed indicator subset and perform bulk edit or bulk delete operations.

The following example illustrates this paradigm. Some customers desire a staging process, whereby only vetted indicators are “Enabled” and thereby published to the enforcement points. Suppose an input feed policy specifies that newly created indicators should default to DISABLED status. The analyst will filter on the specific Input Feed, Status DISABLED, sort in descending order by modification date, select the indicators to be enabled, and then click "Bulk Edit":

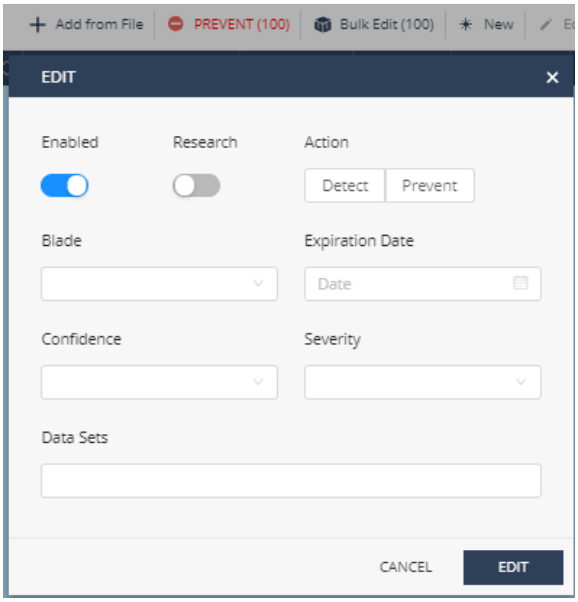

Notes

- The Bulk Edit feature applies changes to the selected indicators only on attributes that are touched on this pane. In this example, the only attribute to be modified is the “Enabled” status.
- In order to bulk-disable, click Enabled twice and then click EDIT.
- The modified data overwrites the existing data. Since Data Sets is a multiple-value field, you must specify all indicator data sets associated with the selected indicators, not just new data sets

Similarly, indicator confidence, severity, and action (DETECT/PREVENT) are defined per Input Feed via Policy. New indicators can be created in DETECT action, and upgraded to PREVENT after a suitable triage period in which the false positive ratio is deemed acceptable. Such triage processes can be performed manually by the analyst, and can also be automated.

Filtering

Filters determine the indicators that are displayed on the Indicators tab.

The indicators that match the defined filters are displayed in a sort order that defaults to descending on Modified timestamp. Clicking on a column heading sets sorting to apply based on the selected column. Clicking again reverses the sort order.

Clicking the filter button opens the FILTERS tab. By default, no filters are defined, and all indicators are displayed. CLEAR FILTERS reverts to that state.

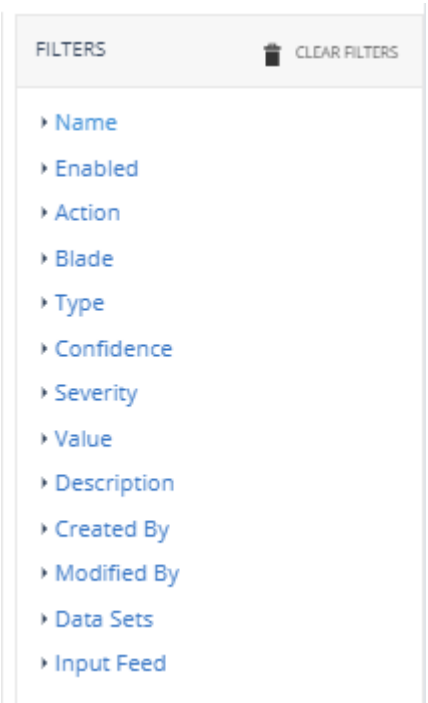

- Name – exact match on indicator Name attribute
- Enabled – allows display of only Enabled or only Disabled indicators
- Action – allows display of only Prevent or only Detect indicators
- Blade – allows selection of a single Blade for displaying indicators
- Type – multi-select (via check boxes) for indicator Type
- Confidence, Severity – allows display of a single Confidence and/or Severity attribute
- Value – exact match on indicator Value attribute
- Description – prefix match on indicator Description attribute
- Created By, Modified By – exact match on user name (up to @ character)
- Data Sets – multi-select (via check boxes) on Data Sets
- Input Feed – allows display of indicators from a selected (from drop down menu) input feed

Data Sets

Indicators may be associated with Data Sets. The data set is used both to group indicators with a common semantic, and to publish output feeds of the indicators, intended to be consumed by enforcement points and other external platforms.

For example, consider the @Tags data set, previously mentioned in the context of qualification of indicators by input feed. @Tags is created automatically on the domain when the analyst tags a host on the Threat Topology Analytics view, binding the host IP with the analyst-entered Name. Threat Topology will use the IP indicators on the @Tags data set for tag display. This means that an alternative to right-click "Add Tag" is to simply add an indicator with the host's IP, and associate it with @Tags. Using indicators as a storage repository for tags provides the ability to integrate with CMDB systems using indicator input feeds; loading tags in bulk with "Add from File"; and exporting tags via data set output feeds.

In addition, associating indicators with a data set allows the analyst to easily focus on specific data sets' indicators by filtering on the data set(s).

Data Set Attributes

Each data set is qualified by **Name**, and an optional free-text **Description**. In addition, the following attributes control output feed publication of the indicators associated with the data set:

- **Public** – enables data set output feed publication outside of the NDR system for consumption by Check Point and 3rd party enforcement points
- **Apply on NDR Sensors** – CDC-managed dedicated NDR sensors automatically pull and apply the data set output feeds

Data Set Output Feeds

Public data set output feeds are published on URLs. Select and click "Show Data Set URLs":

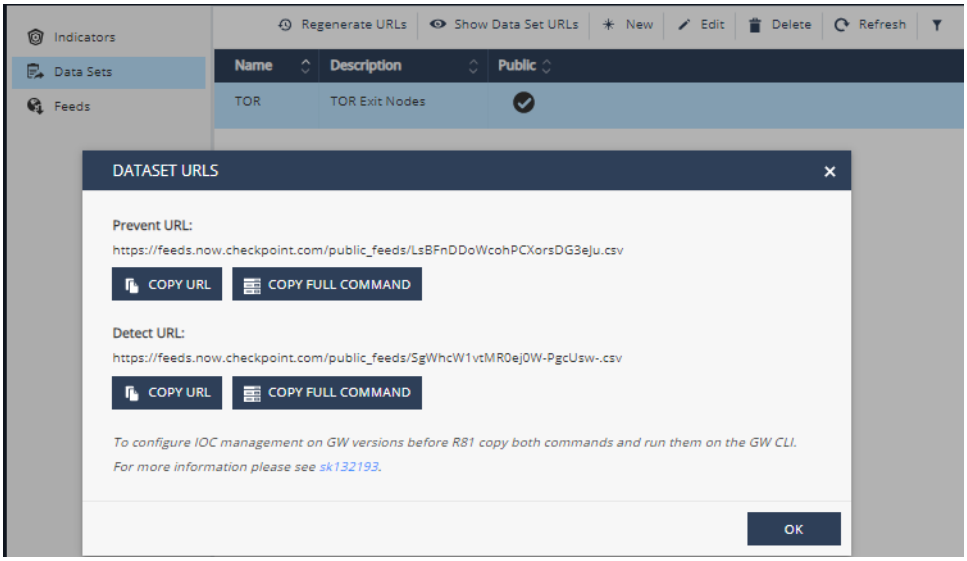

Each data set is published on two output feed URLs: Prevent and Detect. Each of these feeds includes the indicators associated with the data set that have the corresponding action (Prevent or Detect). This separation is useful because Check Point Security Gateways associate action with an individual feed configured on the gateway.

The URL for each output feed is displayed on the portal, and may also be easily copied to the clipboard via the “COPY URL” button. For convenience, a “COPY FULL COMMAND” option is provided, which enters a complete gateway ioc_feeds command that can be entered on a gateway for pulling the output feed.

The criteria for inclusion of an indicator in an output feed include:

- Indicator must be ENABLED
- Indicator must not be expired
- Indicator action must match the specific output feed (DETECT/PREVENT)
- Indicator type must be compatible with the output feed (see below)

Regenerate URLs

On public data sets, the feed URLs are randomly generated, and constitute both the feed identifier and authentication. All gateways that consume the feed are configured with these URLs. If there is a concern that these URLs might have been exposed to unauthorized parties, the “Regenerate URLs” command can be used to create new ones. This will pop up a prompt such as the below:

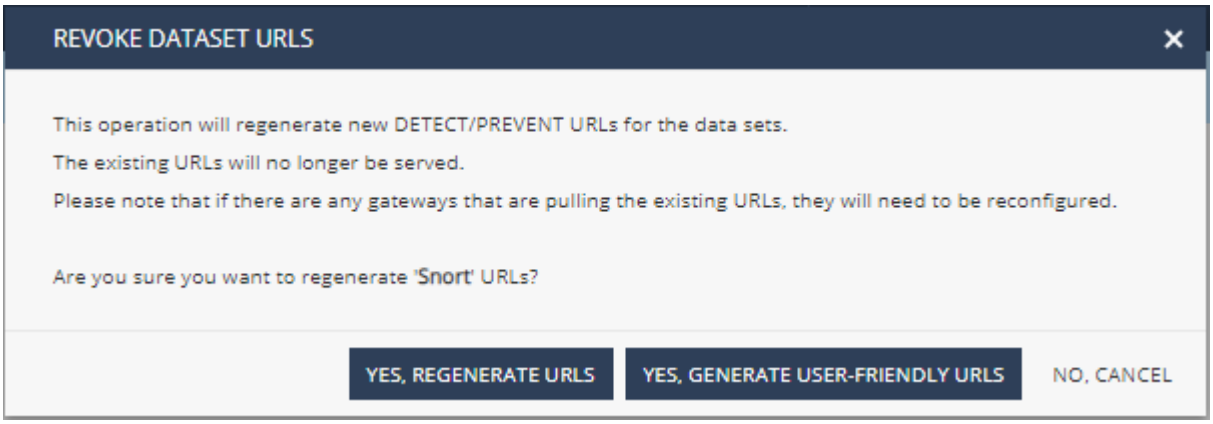

The second option, “GENERATE USER-FRIENDLY URLS” is used for sharing data set output feeds with other organizations. Instead of a randomly-generated URL, the feed will be served on a URL constructed from the domain name, detect/prevent action, and the data set name. For example, for data set Snort, on domain test1, the DETECT feed will be served at: https://feeds.now.checkpoint.com/public_feeds/test1-Snort-detect.csv.

Output Feed Formatting

By default, a data set's output feed is published in a CSV file format compatible with sk132193 (Custom Intelligence Feeds). This allows Check Point Quantum and CloudGuard Security Gateways (version R80.30 and up) to consume the feeds natively with the Anti-Virus and Anti-Bot blades.

Many enforcement products and versions have limitations on the type, number, and volume of threat indicators and feed formats that they can consume.

If you expand the "Advanced" section of the new data set / edit data set dialog, you will see more options that enable the customization of the output feed format to accommodate more enforcement solutions. These options include:

- **Format** – File (default) or Text
- **Content format** – "Check Point CSV" (default) or "Values only"
- **Indicators limit** (0 for unlimited) – truncates output feeds to the given number of indicators
- **Indicator types** - only indicators with the selected types will be published on output feeds

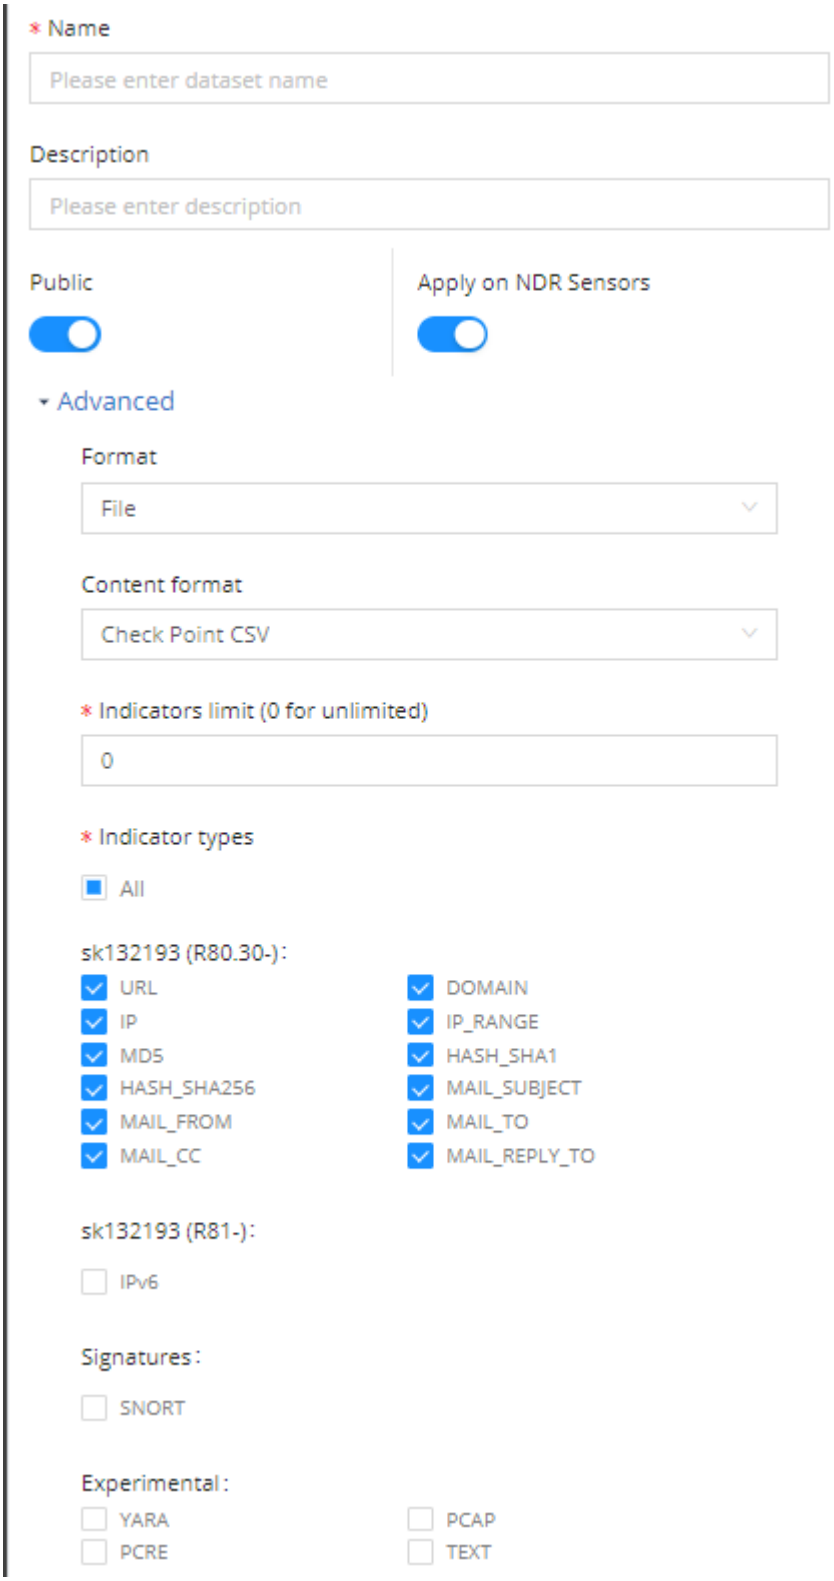

Examples:

- For Security Gateway R80.30 (and Check Point Quantum Spark appliances), de-select HASH_SHA1 and HASH_SHA256 as these indicator types are only supported from R80.40. Doing so will protect these gateways from receiving these indicator types on the data set, even if such indicators are associated with the data set
- IPv6 indicators are allowed by default, however they are supported only from version R81
- For Fortinet and Palo Alto Networks firewalls, configure separate data sets per supported indicator type, configured with “Values only” and a single indicator type each (e.g. IP)

Note: The default Check Point CSV format used for sk132193 compatibility (as well as for input feeds and Add from File) uses the following columns in the following order:

"# UNIQ-NAME","VALUE","TYPE","CONFIDENCE","SEVERITY","PRODUCT","COMMENT"

Input Feeds

Input feeds are a mechanism for automating indicator ingestion into NDR Intel. The input feed determines a source for the indicators, what initial metadata attributes they receive (via policy assignment), and the initial data sets they are published on. The feed source is automatically polled by NDR and the indicators from the feed created and/or updated in NDR Intel. After an indicator is created, its metadata and data set assignments may be further modified by the user.

The following attributes characterize an input feed:

- **Name** – used to identify the indicators that are generated off this feed
- **Active** – Controls whether indicators are being generated or not
- **Periodicity** – determines the poll period for the feed (default: once a day)
- **Type** – Characterizes the input feed source. Supported types include:
    - **Check Point CSV** – Same file format as used for CSV output feeds. A header is mandatory as it determines column order. Only TYPE and VALUE are mandatory.
    - **Single-type CSV/list** – A file containing only indicator values, one per line (Further qualified by has-header, indicator type, column and delimiter.)
    - **Multi-type CSV/list** – A file containing only indicator values, one per line. Indicator types may be intermixed, with the type inferred from the indicator value. Supported types include: IP, DOMAIN, URL (with protocol specifier), MD5, SHA1, SHA256. (Further qualified by has-header, column, and delimiter.)
    - **STIX/TAXII** – Mitre STIX/TAXII v1.0 or v1.1
    - **Behavioral Analytics** – Horizon NDR Behavioral Analytics AI engines
    - **IOC Harvester** – IOC extraction from Check Point Threat Prevention logs
    - **API** – Used for managing indicators via NDR Intel API
    - Third Party Threat Intelligence Platforms (TIPs) via REST API
        - **Anomali** (IOCs)
        - **Cybernet** (IOCs)
        - **ThreatConnect – Signatures** (Snort signatures)
- **URL** – Source of the input feed (not needed for Behavioral Analytics) (Optional server credentials can be defined – username and password.)
- **Collection** – optionally qualifies the pulled indicators (supported by some feed types)
- **Policy** – initial metadata attributes for indicators created by the input feed
- **Data Sets** - The initial data set associations for indicators created by the feed

Indicator Updates

New indicators received on an input feed are added to the indicators repository, and the Created timestamp set. When an indicator value is pulled from an input feed, and that indicator already exists in NDR Intel, the indicator's Modified timestamp is refreshed. If the input feed provides metadata attributes in addition to the type and value, these metadata attributes overwrite the existing attribute values (unless the "Override" setting has been enabled in the policy). Other indicator metadata attributes will persist.

For example, on a Check Point CSV input feed, the source can override confidence, severity, product, and comment (i.e. description). However, action and status are not affected as they are not included in the input feed. An Anomali feed will update indicators deactivated on Anomali to state DISABLED, immediately removing them from any output feed.

Delta Feeds and Indicator Expiration

Depending on the input feed Type, NDR Intel might be pulling all of the source's indicators on every poll ("full feed"), or only the newly added or updated indicators since the previous poll ("delta feed").

The following types are full feeds: Check Point CSV, Single-type CSV/list, Multi-type CSV/list, ThreatConnect – Signatures. The following types are delta feeds: STIX/TAXII, Anomali, Cybernet, Behavioral Analytics, IOC Harvester.

Indicator expiration provides the primary mechanism for synchronizing with delta feed sources. Whenever the indicator is updated by the source, the indicator's expiration attribute is refreshed, adding the Time To Live value defined in the input feed's policy object to the current time.

On a delta feed, if the source deactivates or deletes an indicator, it will not be refreshed on NDR Intel (with the exception of Anomali), and will eventually expire based on the policy expiration setting. Once expired, it is immediately removed from all data set output feeds it is associated with, and will be deleted from NDR Intel 14 days later (unless its expiration is refreshed by a subsequent update).

Expired indicators' expiration is displayed in red font in the Indicators view:

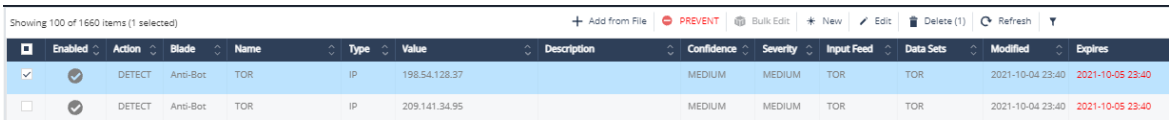

Policies

A Policy object determines the default metadata attribute values that an indicator receives when it is generated from an input feed or from an Add from File operation. The following policy attributes are supported:

- **Name** – The Policy name, also used as the Indicator name where not available
- **Enabled** – The initial status (ENABLED/DISABLED) of newly-generated indicators
- **Override** – Determines feed behavior if the feed source provides metadata attributes (e.g. Confidence) in addition to indicator name, type, and value. By default, the feed-provided value takes precedence. If Override is enabled, all indicators on the input feed will receive the metadata attribute values defined in the Policy, regardless of feed-provided values.**Description** – A description of the Policy's objective
- **Prevent** – The initial action (DETECT/PREVENT) of newly-generated indicators
- **Minimum confidence** – If action is "Prevent", specifies the minimum confidence for which this is applied. Lower confidence levels will remain in "Detect".
- **Blade, Confidence, Severity** – The initial attribute values for newly-generated indicators
- **Expiration In Days** – The Time To Live for a newly-generated indicator, from current time. Relevant only for delta feeds and Add from File. Ignored on full feeds

Add from File

The input feed mechanism can also be invoked manually by specifying a file as the source for the indicators. A policy object determines the metadata attribute values as with automated input feeds. The indicators read from the file can be assigned to one or more data sets:

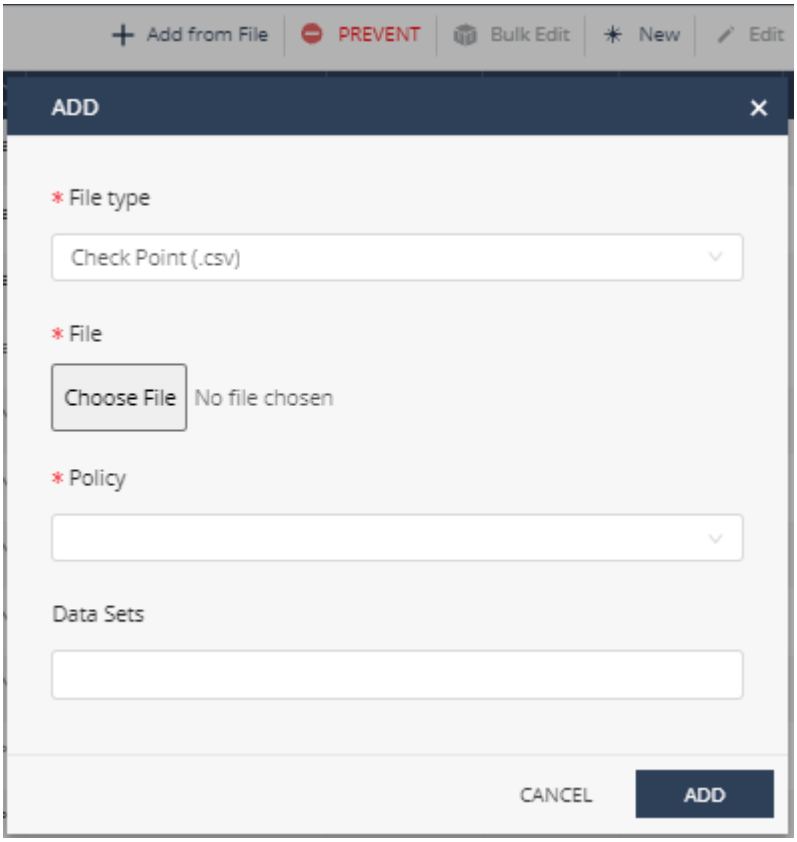

The following file types are supported by Add from File:

- **Check Point (.csv)** - Same file format as used for CSV output feeds
- **Single-type CSV/list (.txt .csv)** – Typed list, one value per line. Configuration:

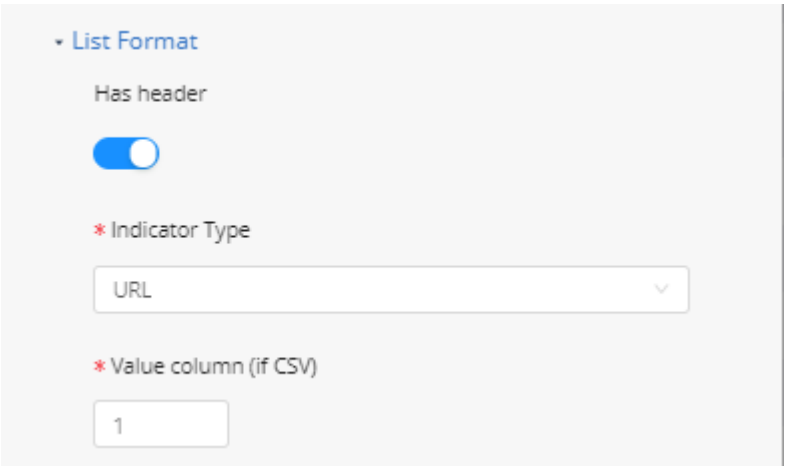

- **Multi-type CSV/list (.txt .csv)** – Untyped list, one value per line
- **SNORT (.rules)** – One Snort rule per line
- **STIX 1.x (.xml)** – STIX indicator file

Behavioral Analytics

The NDR platform includes a set of Behavioral Analytics AI engines that process logs, and identify different types of anomalies (Geo, Port, User, etc.). These anomalies are then correlated with Check Point ThreatCloud reputation services and with application categorization and risk scoring in order to incriminate them as suspicious events that should trigger an alert.

In some cases, Behavioral Analytics will also sign an IP or URL that is involved in the anomaly, by publishing a threat indicator on NDR Intel. These indicators are mapped to an input feed of type "Behavioral Analytics". When available, confidence is also associated to some of these indicators.

Whenever a new NDR domain is created, NDR will automatically create an active input feed, a policy object, and a data set – all named "Behavioral":

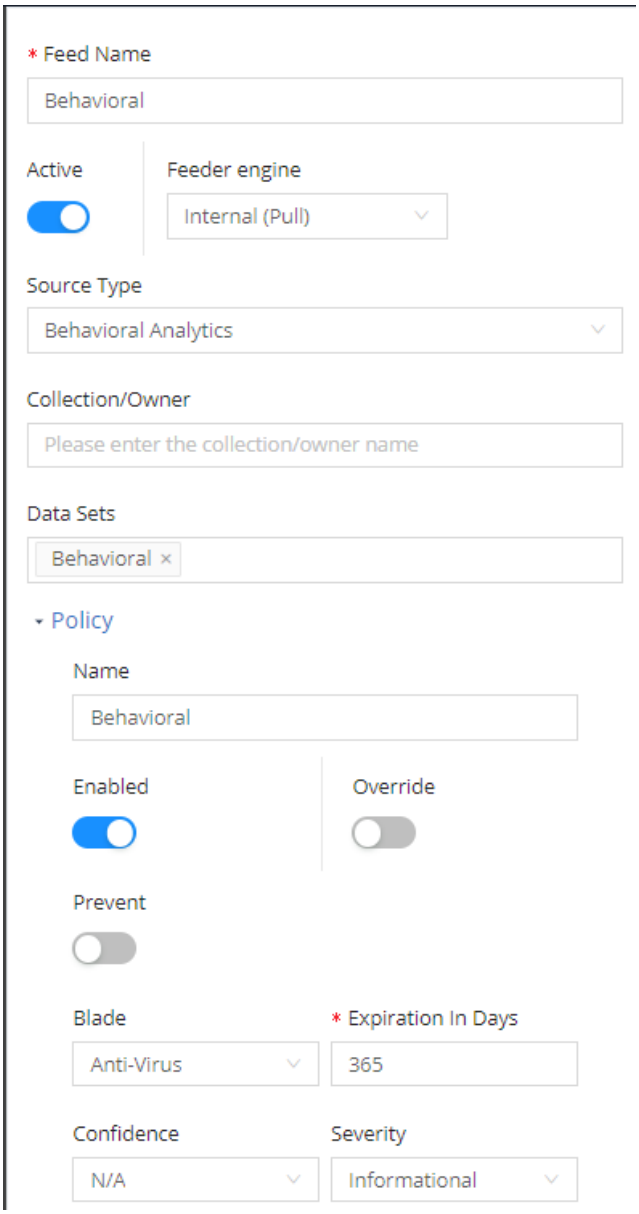

This scheme allows the user to control whether the indicators should be generated (Active on the input feed), and whether they are created in ENABLED state, DETECT or PREVENT (qualified by confidence level), etc.

The input feed "Collection" attribute provides finer-grained control over indicator generation. Currently supported collections include: Geo, Protocol, Users, and RecurrentConnections. Specifying a collection is optional – an empty value matches all Behavioral Analytics engines.

Snort Indicators

Snort rules may be added to NDR Intel as indicators via the NDR portal; using 'Load From File'; via the NDR Intel API; or using an automated input feed. They may be edited via the portal or the API. Only IPS blade is supported for Snort indicators.

When using a Single type list of SNORT rules as input feed, or loading a file manually using 'Load From File', the input file must contain a set of Snort rules, one on each line. Comment lines starting with '#', as well as empty lines, are ignored. The text on each line is entered as a separate Snort indicator value. Description is set to the imported file name. The indicator name is set to the value of the msg field in the rule.

If a rule includes CVE reference attributes, e.g. "reference:cve,2014-6271", these references will be copied into the "CVE" metadata attribute for the rule. This attribute is used to correlate Snort signatures with native Check Point IPS signatures.

When an entered Snort indicator is not compatible with Check Point supported rule syntax:

- The indicator will automatically be set to the "Disabled" state, so that it will not be published on any associated data set
- The SnortConverter error message will override the indicator's Description field on the UI:

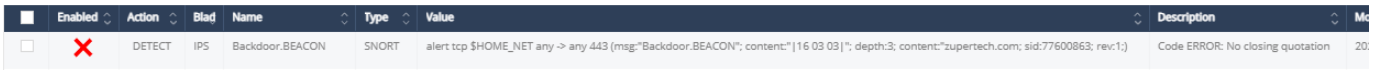

Snort Output Feeds

It is recommended not to mix both Snort and other types in a single data set. A Snort-only data set is defined as follows:

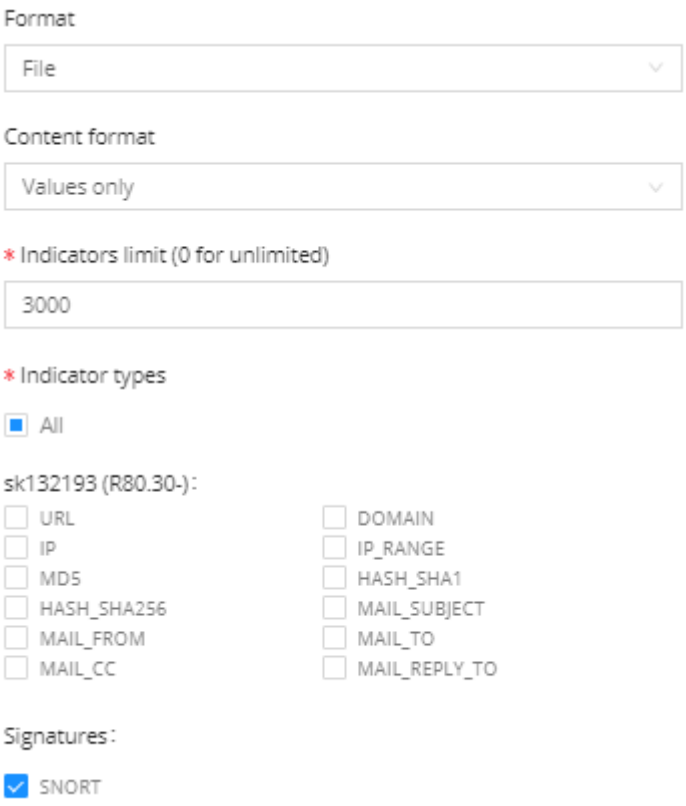

When thus configured, the data set's output feed is compatible with gateway-side ioc_feeds capability on some limited-release Check Point versions. It can also be imported into the IPS database on a Check Point Security Management Server.

In the output feed, a unique ID is inserted as a suffix to the rule's msg, in order to guarantee uniqueness and allow correlation of logs to rules in case of multiple rules with the same msg string. This unique ID is visible in logs generated by IPS signature matches.

Confidence and Severity Metadata Attributes

When pulling Snort indicator feeds from NDR Intel, Check Point Gateways and Security Management Servers currently support only "Values only" format. Meta data attributes including confidence, severity, and performance impact are assigned on the gateway to the feed, not to an individual indicator.

These feed attributes also determine the action upon indicator match – based on the IPS threat prevention policy. Thus the 'Detect' and 'Prevent' distinction on NDR Intel is only a convention that determines which feed URL the indicator is published on, and the action should be aligned with the gateway feed configuration. The NDR Intel 'Confidence' and 'Severity' are ignored by the gateway in relation to Snort indicators.

The following recommendations may alleviate this discrepancy:

- It is suggested to hide the 'Confidence' and 'Severity' columns on NDR Intel for users that only manage Snort indicators, as they are not carried over to the gateway implementation
- On each gateway, define two feeds for each applicable data set. The DETECT feed should be assigned severity 3 and confidence 2. The Prevent feed should be assigned severity 4 and confidence 5. Define both with performance impact 1.

  These values were selected to correspond to widely used threat prevention profile settings in order to trigger the corresponding action (Detect/Prevent).

Known Limitations

- The number of indicators supported on an individual feed file is limited to 3000.
- Maximum supported length of a Snort rule is 2048 characters.
- The Snort rule syntax (for the indicator value) supported by Check Point is documented in: https://sc1.checkpoint.com/documents/R81.10/WebAdminGuides/EN/CP_R81.10_ThreatPrevention_AdminGuide/Topics-TPG/SNORT-Signature-Support.htm.

Use Case Examples

The Onion Router (TOR) Exit Nodes Input Feed

Check Point publishes a list of TOR IPs at https://secureupdates.checkpoint.com/IP-list/TOR.txt. This list is refreshed on a daily basis.

Step 1 - Create a data set that can be used to publish these indicators to Check Point gateways:

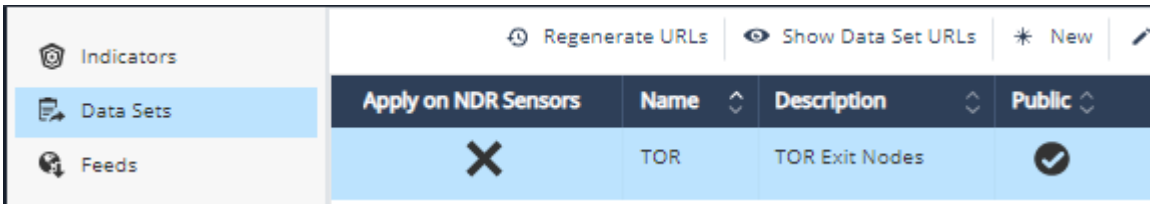

Step 2 - Configure the input feed, referencing the data set. The feed type is "LIST", because the source is simply a text file with an indicator per line.

Make sure to uncheck "Has header" (because there's none on the feed). Leave "Feeder engine" as Internal – this means that the system polls the feed instead of a push paradigm:

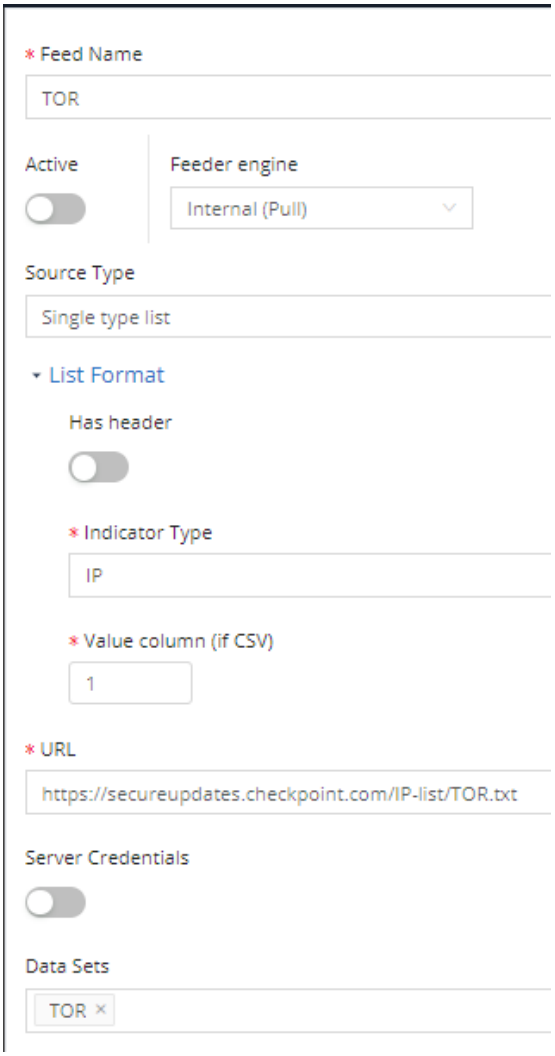

Here's what it looks like on the Feeds tab:

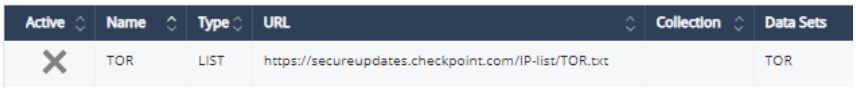

Polling interval is about 10 minutes, after which the indicators should appear on the Indicators tab:

Note that once these indicators are created, you can update the meta data (status, action, blade, confidence, severity, data set assignments) for each indicator individually and the change will persist, even when the indicator is updated on the feed.

Also note that the TOR data set was configured as "Public", meaning that the data set URLs can be pulled from outside Horizon NDR. Simply "Show Data Set URLs" to access them.

Publishing IOC Output Feed to Fortigate

Fortigate supports configuration of an external indicator feed. A Check Point Horizon NDR sensor can be deployed passively on the customer's network, inspecting network traffic and publishing threat indicators to signal the Fortigate to block attack traffic.

Step I – Create a file data set with a single type list in "Values only" content format, for example:

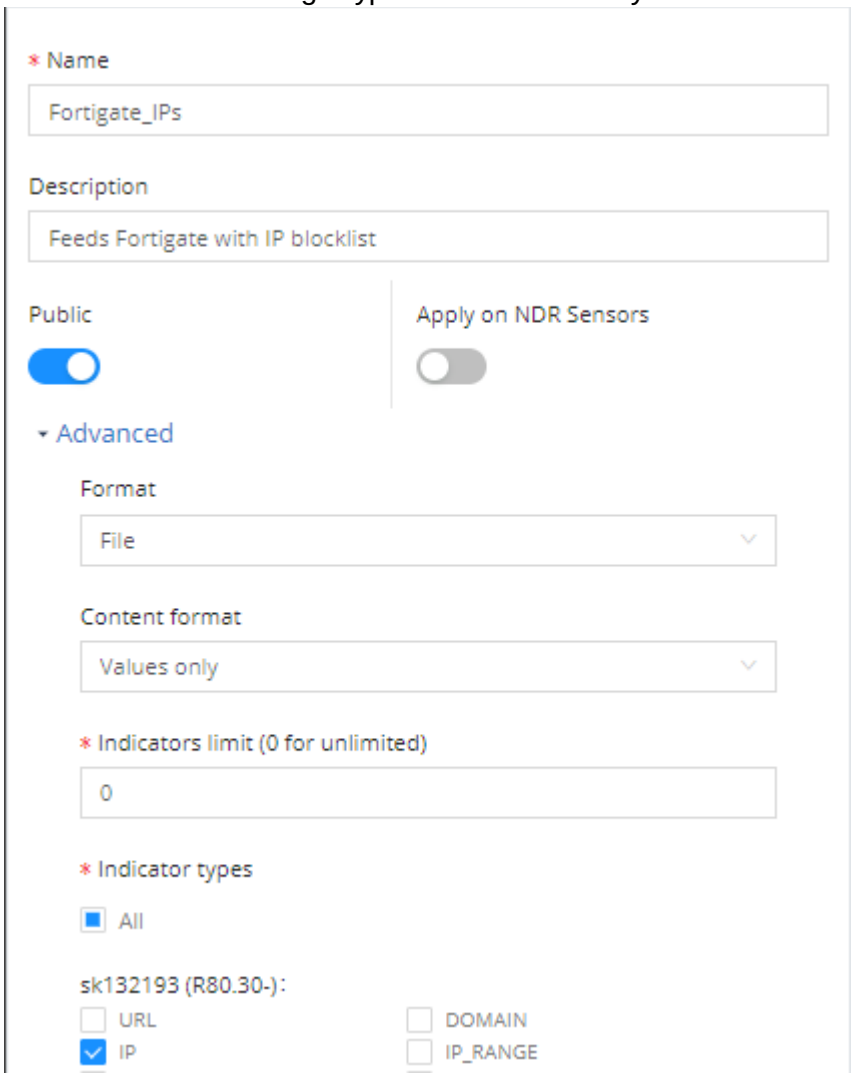

Step II – Select the new data set and Show Data Set URLs. Copy the relevant output feed's URL.

Step III – On the Fortigate user interface, create a connector, specifying the feed's URL:

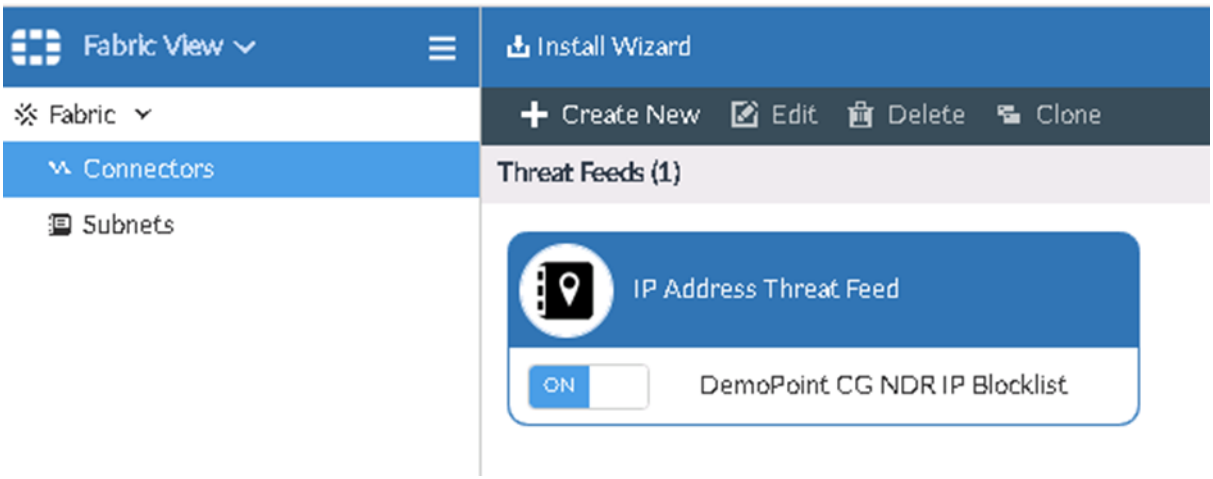

Step IV – Define the connector object as source or destination in the rule base.

Below is an example Fortigate log, matched on an NDR-delivered IP indicator:

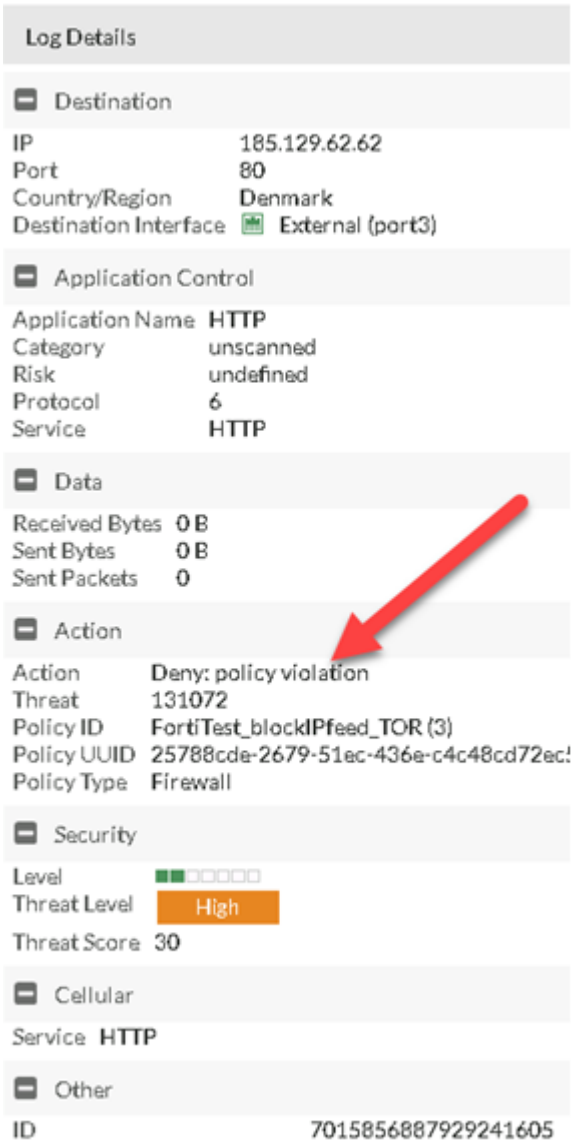

Check Point Quantum Spark (SMB) Appliances

Check Point Quantum Spark version R80.20.x support a variant of the sk132193 ioc_feeds command with the following usage pattern:

```
ioc_feeds <action> [options]
Actions:
      set <feed name> - set an external feed. options:"
            --resource    - set the remote URL for the feed"
            --transport   - specify the transport protocol [http|https]"
            --action      - specify the action [detect|prevent]"
            --state       - specify whether the feed is active [true|false]"
      delete <feed name> - deletes the feed <feed name>"
      delete_all         - delete all the feeds"
      show               - show configured feeds"
      sched [interval]   - set periodic pull interval (in seconds, minimum 30)
      enable [on|off]    - enables/disables external IOCs
```

For example: `ioc_feeds set main_Detect --transport https --resource "https://feeds.now.checkpoint.com/public_feeds/****.csv" --action Detect`