CrowdStrike Falcon Devices Add-on for Splunk

Installation and Configuration Guide v2.1

Introduction

This guide covers the deployment, configuration and usage of the CrowdStrike Falcon Devices Technical Add-on (TA) for Splunk.

The CrowdStrike Falcon Devices Technical Add-on for Splunk allows CrowdStrike customers to retrieve device data from the CrowdStrike Hosts API and index it into Splunk. The specific endpoint that is used is:

/devices/queries/devices-scroll/v1

To get more information about this API, please refer to the API documentation which can be found in the CrowdStrike Falcon UI: https://falcon.crowdstrike.com/support/documentation/84/host-and-host-group-management-apis.

Multitenancy - This TA is able to have multiple independent inputs enabled at the same time, each collecting data from different Falcon Instances and storing it in independent indexes.

Requirements

The following are the requirements to leverage this technical add-on:

1. An active subscription to the CrowdStrike Insight or Prevent modules
2. A Splunk Heavy forwarder or Input Data Manager (IDM)
3. A Splunk account with proper access to deploy and configure technical add-ons
4. A properly scoped API credential or proper access to the CrowdStrike Falcon instance to create one
5. The CrowdStrike Cloud environment that the Falcon instance resides in

Getting Started

High Level API Call Flow

The CrowdStrike Falcon Devices TA performs the same API calls at each time interval that's configured within the TA input:

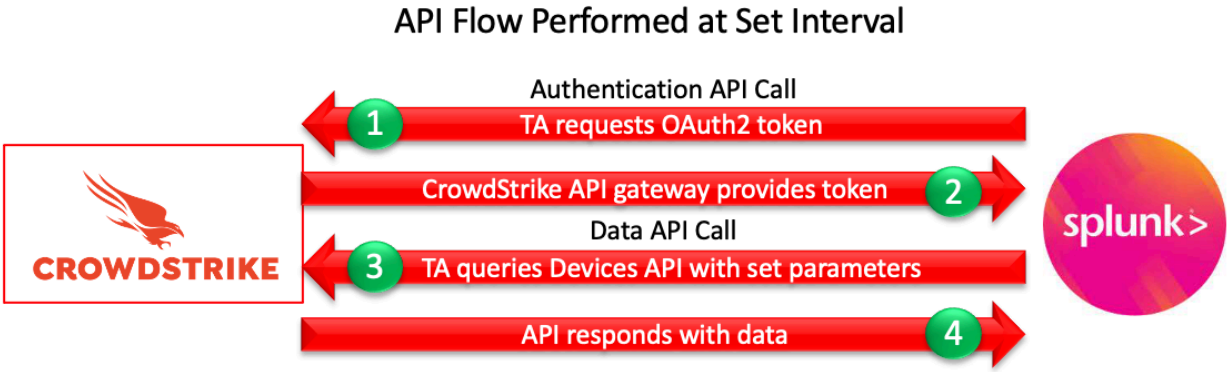

1. The TA will call the CrowdStrike API gateway with the configured credentials and request an OAuth2 authentication token that is valid for 30 minutes.

2. If the API credentials are valid the API gateway will respond to the TA with an Oauth2 token.

3. The TA will use the OAuth2 token to call the Devices API with the configured parameters.

4. The API will respond with whatever appropriate data matches the configured parameters.

Technical Add-On Layout

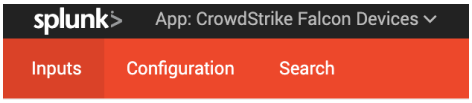

The CrowdStrike Falcon Devices TA has 3 tabs associated with it:

1. **Inputs** – The Inputs tab (only configure on a Splunk Heavy Forwarder or IDM) contains the connection configuration(s) that the TA uses to communicate with the API.
2. **Configuration** – The Configuration tab contains the API credential information, proxy server configuration information and logging level.
3. **Search** – A link to Splunk search that's specific to the TA.

Creating/Validating the API Credential Scope

While the CrowdStrike Falcon Devices TA can leverage an existing OAuth2 based API credential, it is recommended that a dedicated credential be created and used. This can be accomplished by the following:

1. Access the CrowdStrike Falcon user interface (UI) with an account that is able to create API clients and keys
2. Navigate to 'Support'>'API Client and Keys' page
3. Create a new API client by selecting 'Add new API client' in the OAuth2 API client's area

4. Give the new API client a name and description (recommended) and under 'API Scopes' select the 'Hosts' scope and the 'Read' capability

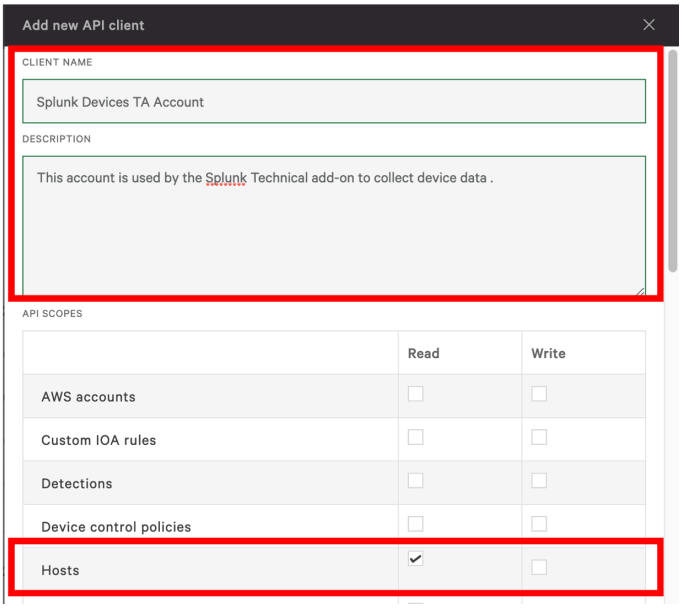

5. Select ‘Add’ once completed and a window will appear with the Client ID, Secret and the Base URL. NOTE: This is the only time the Secret will be visible – ensure it is recorded in a protected location.

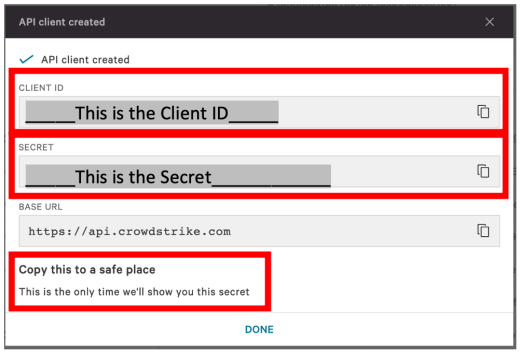

6. In addition, make note of the ‘BASE URL’ value in either the API client created window or the ‘OAuth2 API client’ area as this will be used to determine the CrowdStrike Cloud the instance is in

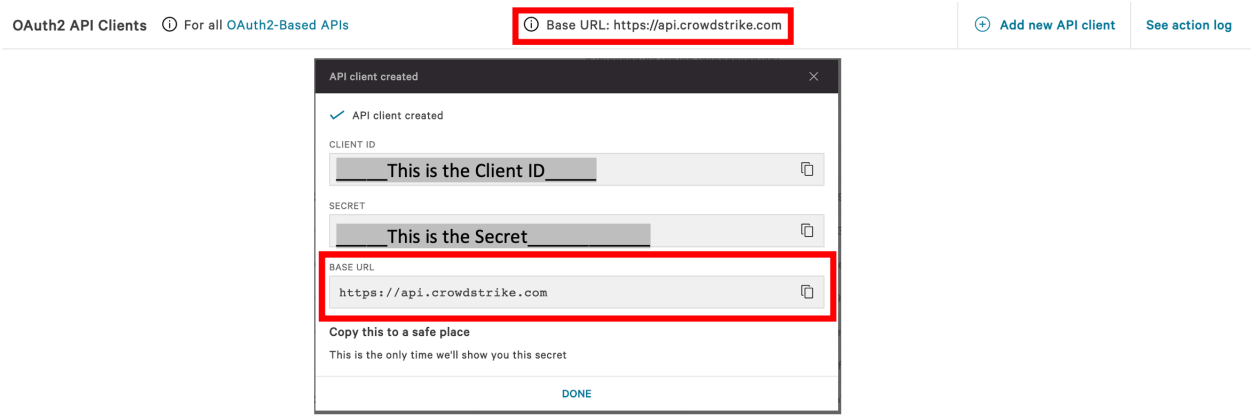

7. Select ‘Done’ to close the window and finish creating the credential

Proxy Considerations

The CrowdStrike Devices Technical Add-On establishes a secure persistent connection with the Falcon cloud platform. In some environments network devices may impact the ability to establish and maintain a secure persistent connection and as such these devices should be taken into account and configuration modifications should be done when necessary.

Ensure that the API URLs/IPs for the CrowdStrike Cloud environment(s) are accessible by the Splunk Heavy forwarder. For a complete list of URLs and IP address please reference CrowdStrike's API documentation.

The current base URLs for OAuth2 Authentication per cloud are:

US Commercial Cloud	: https://api.crowdstrike.com
US Commercial Cloud 2	: https://api.us-2.crowdstrike.com
US GovCloud	: https://api.laggar.gcw.crowdstrike.com
EU Cloud	: https://api.eu-1.crowdstrike.com

Splunk Architecture

Splunk Search Head(s) and Splunk Cloud: The TA should be installed to provide field mapping and search macro support. These are often required to support CrowdStrike Apps. The TA should be deployed without any accounts or inputs configured and any search macros should be properly configured for use.

Splunk Indexer(s): The TA can be installed to provide field mapping and search macro support. The TA should be deployed without any accounts or inputs configured and any search macros should be properly configured for use. If a custom index is going to be used, then it should be created here.

Splunk Heavy Forwarder(s) & Information Data Managers (IDMs): The TA is required to be installed here as this is where the data from the Devices API will be collected. The appropriate accounts and inputs should be properly configured for data collection. Ensure that if a customer index is being used, which is highly recommended, that the index has been created on the indexer tier. If the Heavy Forwarder is storing events (not required but is an optional Splunk configuration) prior to forwarding them to the Indexer and a custom index is being used, ensure that the index has been created on both the Heavy Forwarder as well as the Indexer(s).

Note:

Due to python requirements the TA can only be configured for data collection on Heavy Forwarders and IDMs.

The following diagram shows the flow of data from the Devices API and the Falcon Device TA configuration within a distributed Splunk Enterprise and Splunk Cloud environment:

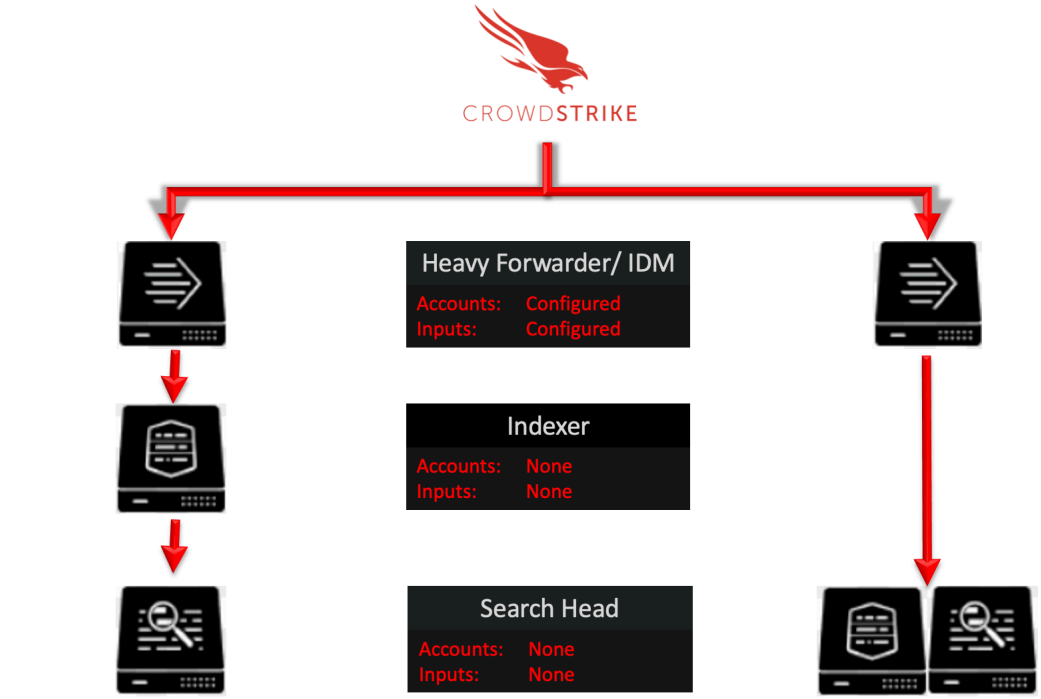

Configuring the TA

TA Layout

The TA contains 3 sections:

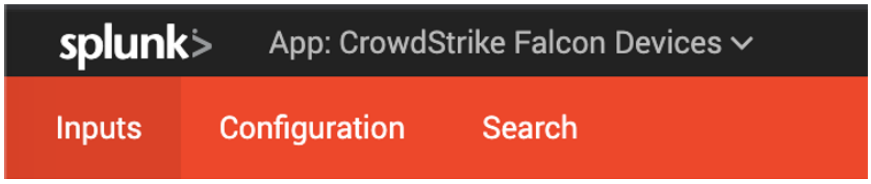

- The Inputs section
- The Configuration section
- The Search section

Inputs Section

The Inputs section is where inputs are configured, modified and listed. Prior to configuring any inputs an account needs to be created under the Configuration section (see below). The Inputs section contains a 'button' that will create a new input configuration in the far-right corner.

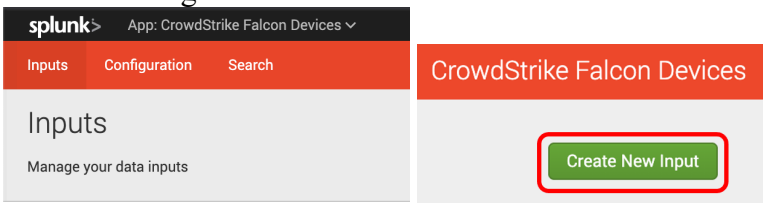

Configuration Section

The Configuration section contains 3 configuration tabs:

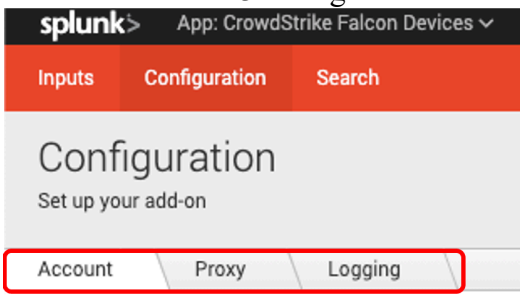

- **Account Tab**: This is where the OAuth2 API credentials are entered.
- **Proxy Tab**: This is where proxy server configurations are entered.
- **Logging Tab**: This is where the logging level is configured.

Search Section

The Search section opens a standard Splunk search page but within the context of the TA.

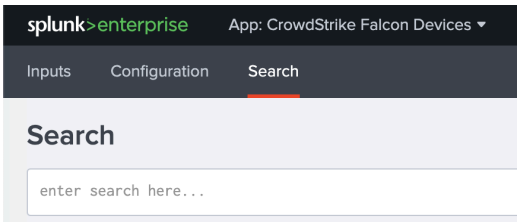

Configuring the TA to collect data

NOTE This action should only be performed on a Splunk Heavy Forwarder or Splunk IDM

Configure Proxy Settings (optional)

1. Proxy settings are configured under the Configuration section, Proxy tab. Proxies can cause authentication issue if not configured correctly, the proxy should not perform SSL/TLS proxying on any API calls.

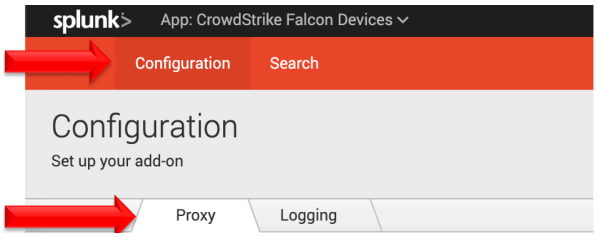

2. Configure the following fields as appropriate:

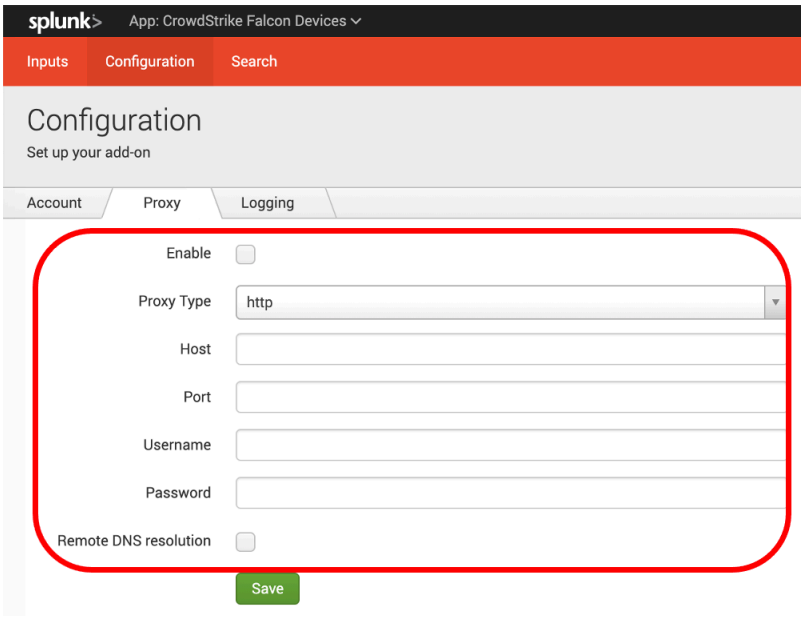

- **Enable**: This checkbox is used to enable/disable the proxy settings
- **Proxy Type**: This dropdown is used to select the proxy type
- **Host**: The hostname/IP address for the proxy server
- **Port**: The communication port for the proxy server
- **Username**: The authentication username for the proxy (optional)
- **Password**: The authentication password for the proxy (optional)
- **Save**: This button is used to safe the configuration

Configure an Account

1. An account is configured using a properly scoped OAuth2 API credential.
2. An account is created under the Configuration section, Account tab:

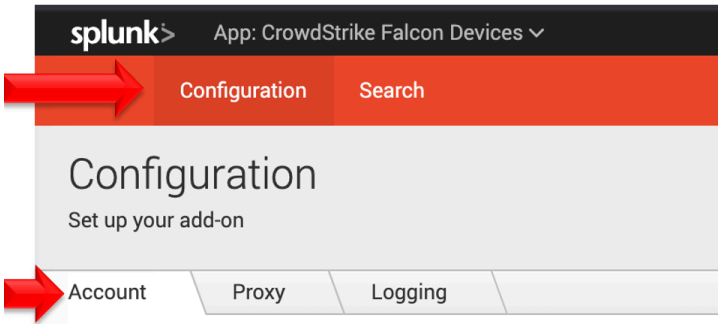

3. On the right side of the screen click the “Add” button:

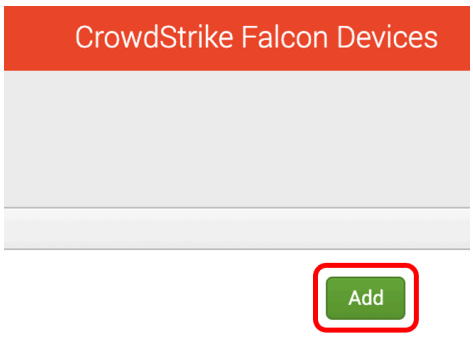

4. Configure the following fields:

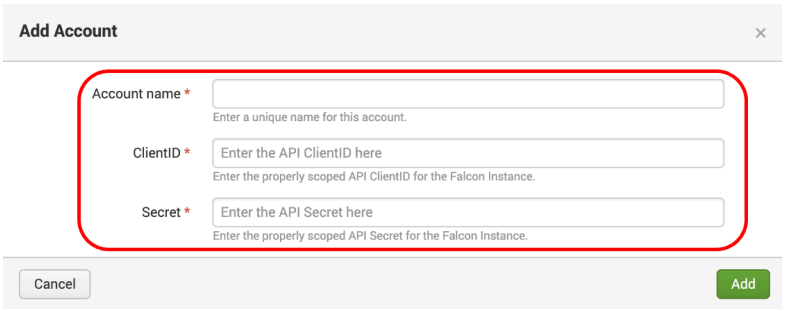

- **Account Name**: A name unique for the Splunk instance
- **ClientID**: The ClientID of the API credential created in the CrowdStrike Falcon UI.
- **Secret**: The Secret of the API credential created in CrowdStrike Falcon UI.

5. Click the ‘Add’ button in the bottom right corner to save the account.

Configure an Input

1. An input will require a valid account be created already.
2. An input is created under the Inputs section:

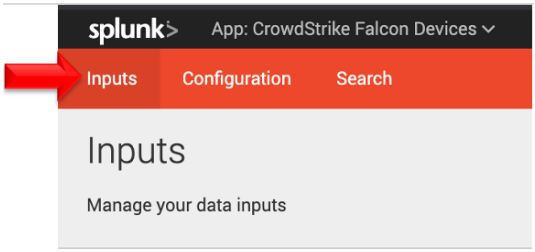

3. Configure the appropriate fields:

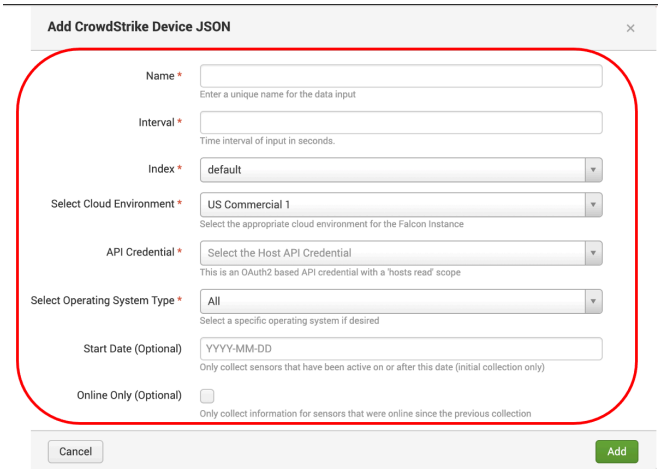

- **Name**: (required) A name unique to the Splunk Environment
- **Interval**: (required) How often the TA will collect data, expressed in seconds (default is 300)
- **Index**: (required) The Splunk Index that the data will be stored in
- **Cloud Environment**: (required) The CrowdStrike cloud environment that that API call will be made to (match the URL indicated in the Falcon UI 'API Client and Keys' page):
    - **US Commercial 1**: https://api.crowdstrike.com
    - **US Commecial 2**: https://api.us-2.crowdstrike.com
    - **GovCloud**: https://api.laggar.gcw.crowdstrike.com
    - **EUCloud**: https://api.eu-1.crowdstrike.com
- **API Credential**: (required) The account that will be used to authenticate to the CrowdStrike API
- **Operating System Type**: (required) Indicates the operating system types to collect

    - **All**: Collects all operating system types
    - **MAC**: Collects only Apple OSX based systems
    - **Windows**: Collects only Windows based systems
    - **Linux**: Collects only Linux based systems
- **Start Date**: (optional) A date in YYYY-MM-DD format that serves as a starting point from which to collect device data. Devices must have been online on or after the date to be collected
- **Online Only**: (optional) Only devices that have a 'falcon_device.last_seen' field value older than the timestamp of the previous device pull are collected.

4. Click the 'Add' button in the bottom right corner to save and active the input.

Search Macros

Search macros are reusable chunks of Search Processing Language (SPL) that you can insert into other searches. Search macros can be any part of a search, such as an eval statement or search term, and do not need to be a complete command. You can also specify whether the macro field takes any arguments.

https://docs.splunk.com/Documentation/Splunk/8.0.3/Knowledge/Definesearchmacros

This TA contains 4 search macros to assist in examining and associating the data collected as quickly and efficiently as possible.

Name	Definition
cs_fd_device_hostname(1)	`cs_fd_get_index` falcon_device.hostname=$hostname$ | dedup falcon_device.device_id
cs_fd_device_id(1)	`cs_fd_get_index` falcon_device.device_id=aid | dedup falcon_device.device_id
cs_fd_device_ip(1)	`cs_fd_get_index` falcon_device.external_ip IN ($ip_address$) OR falcon_device.local_ip IN ($ip_address$) | dedup falcon_device.device_id
cs_fd_get_index	index=falcon_devices

Locating the Search Macros

The search macros can be located by navigating to:

1. Select that ‘Settings’ dropdown in the Splunkbar:

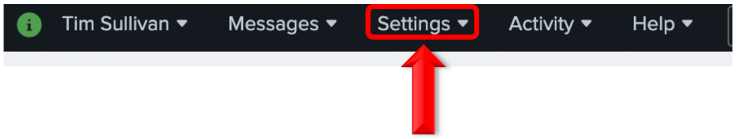

2. Select ‘Advanced Search’ from the dropdown menu:

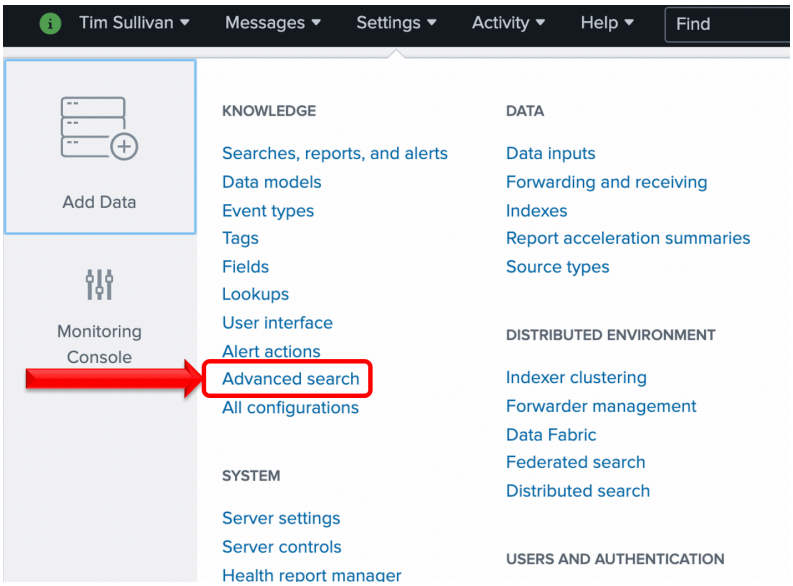

3. Under ‘Advanced Search’ select ‘Search Macros:

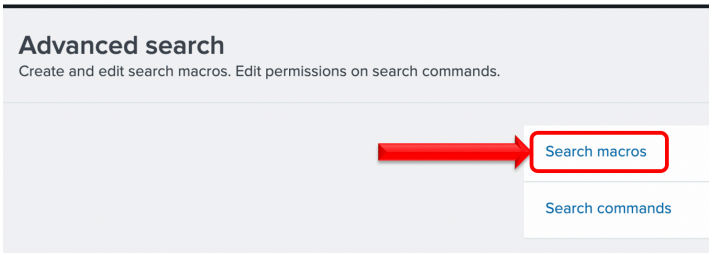

4. Ensure the ‘CrowdStrike Falcon Device Technical Add-on’ is selected in the ‘App’ selection and ‘Created in App’ is selected from the pulldown selection:

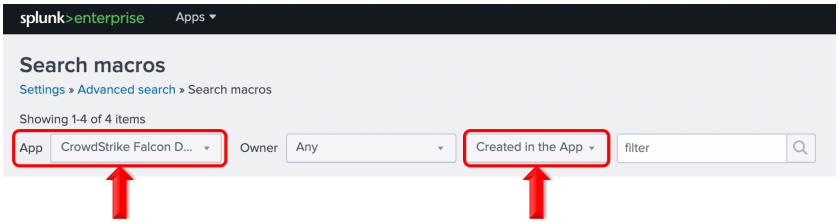

Configuring and Leveraging the Search Macros

There currently is only 1 search macro that requires configuration but this search macro must be configured for the other search macros to function.

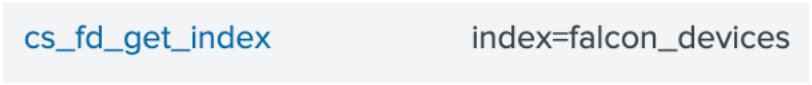

- **`cs_fd_get_index`** (CrowdStrike Falcon Device get index) A search macro that points to the index(es) that contain the data received by the TA inputs. The default for this search macro is to point to an index named 'falcon_devices' and should be adjusted to reflect the specific index(es) that the Heavy Forwarder/IDMs are pushing the data to.

cs_fd_device_hostname(1)	`cs_fd_get_index` falcon_device.hostname=$hostname$ | dedup falcon_device.device_id	hostname
cs_fd_device_id(1)	`cs_fd_get_index` falcon_device.device_id=aid | dedup falcon_device.device_id	aid
cs_fd_device_ip(1)	`cs_fd_get_index` falcon_device.external_ip IN ($ip_address$) OR falcon_device.local_ip IN ($ip_address$) | dedup falcon_device.device_id	ip_address

- **`cs_fd_device_hostname(1)`** A search macro that takes 1 input (a hostname). This search macro will search the configured index(es) for the hostname contained in the parenthesis.
- **`cs_fd_get_device_id(1)`** A search macro that takes 1 input (a Falcon device/agent ID). This search macro will search the configured index(es) for the Falcon device/agent ID contained in the parenthesis.
- **`cs_fd_get_ip(1)`** A search macro that takes 1 input (a IP address). This search macro will search the configured index(es) for the IP address contained in the parenthesis. This search will include both the internal and external addresses.

NOTE: These search macros will not function correctly if the `cs_fd_get_index` search macro is not configured correctly and all search macros must be contained in backticks (not single quotations).

Search Macro Examples

`cs_fd_get_index`

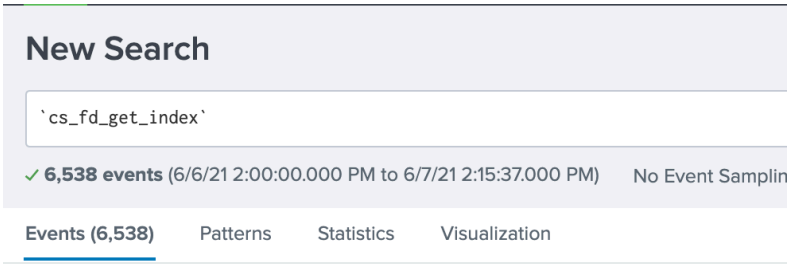

`cs_fd_device_hostname(Win10-001)`

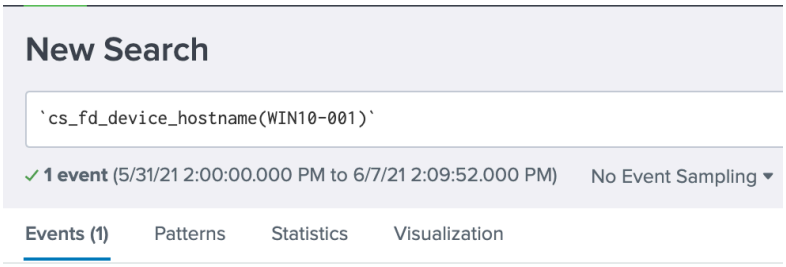

`cs_fd_get_device_id(c8b6q5716xa440408a29637ae244a0p1)`

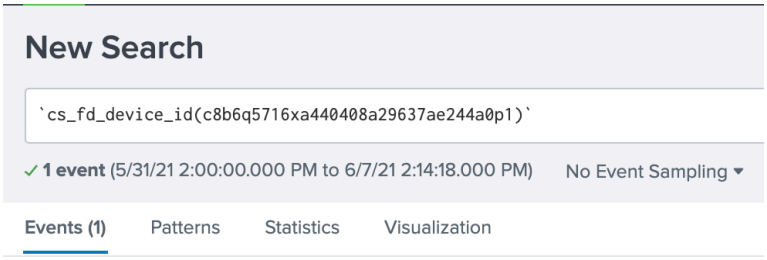

`cs_fd_get_ip(192.168.67.22)`

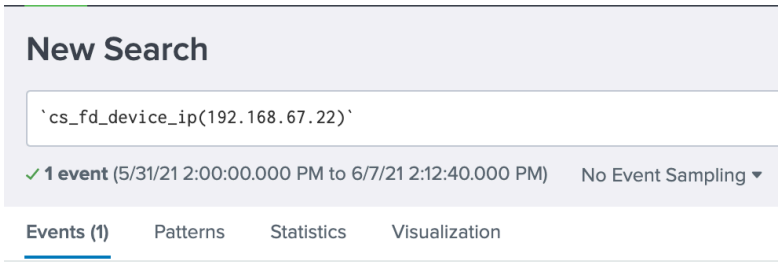

Recommendations

The following are general recommendations. They may not be optimal in all situations and should be evaluated on an environment by environment basis.

Custom Indexes

The use of a dedicated custom index is strongly recommended for the CrowdStrike device data. If data inputs are configured for specific operating system types, it is also recommended that this data be put into dedicated customer indexes for the specific operating system. ie Windows devices would be stored in index ‘W’, MAC OSX devices would be stored in index ‘A’ and Linux devices would be stored in index ‘L’.

This enables the index to be queried specifically as part of either an individual search or a more complex search. It also allows multiple teams to reference the data without exposing other data sets that may be more sensitive.

Troubleshooting

CrowdStrike only provides support for:

- TA code-based functionality errors
- API/Gateway based errors

Examples of issues that are outside the scope of CrowdStrike support:

- Proxy based issues
- Firewall based issues
- Network connectivity issues
- Authentication issues (based on misconfigured credentials)
- Splunk CIM field mapping

Configuring the TA to collect log data

The TA logging level is set to 'info' by default and will only log a minimal amount of information. To properly troubleshoot issues with the TA the logging level should be set to 'debug'.

Change Logging Level

1. Navigate to the Configuration section, Logging tab:

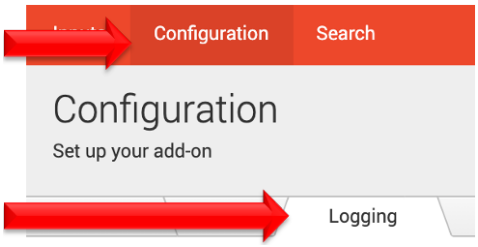

2. Select the logging level from the drop-down menu:

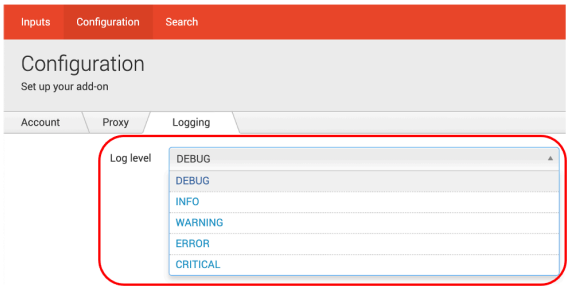

3. Click 'Save' to save the logging level.

Contacting Support

1. Ensure that the OAuth2 credential has been scoped correctly
2. Set the TA log level to 'DEBUG'
3. Repeat and record the action(s) that are associated with the issue you are reporting
4. Download the all log files containing 'ta_crowdstrike_falcon_devices' under the $Splunk/var/log/splunk/ directory
5. Record the following information about the Splunk system:
    - Splunk environment type
    - Splunk version
    - TA version
6. Identify the types of networks devices that the connection will traverse and ensure that they have been properly configured
7. Collect API audit logs from the Falcon instance for the time frame when the issue is occurring
8. Navigate to https://supportportal.crowdstrike.com/
9. Provide (at a minimum) the information from steps 4-7

Additional Resources

CrowdStrike Host and Host Group Management API Documentation

About CrowdStrike

CrowdStrike® Inc. (Nasdaq: CRWD), a global cybersecurity leader, is redefining security for the cloud era with an endpoint protection platform built from the ground up to stop breaches. The CrowdStrike Falcon® platform's single lightweight-agent architecture leverages cloud-scale artificial intelligence (AI) and offers real-time protection and visibility across the enterprise, preventing attacks on endpoints on or off the network. Powered by the proprietary CrowdStrike Threat Graph®, CrowdStrike Falcon correlates over 3 trillion endpoint-related events per week in real time from across the globe, fueling one of the world's most advanced data platforms for security.

There's only one thing to remember about CrowdStrike: **We stop breaches**.