CrowdStrike Falcon Devices Add-on for Splunk

Installation and Configuration Guide v3.1.5+

Table of Contents

Introduction

This guide covers the deployment, configuration and usage of the CrowdStrike Falcon Devices Technical Add-on (TA) for Splunk version 3.1.5 and up.

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

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. It also provides an input type that can be used by Falcon Flight Control customers.

Important Reminder

It's important to keep in mind that a 'device' in CrowdStrike is identified by the sensor ID. This value is referred to in different key values such as 'Device ID' and 'AID' (Agent ID). The purpose of the Device/AID is to act as an unalterable unique identifier for the device within the CrowdStrike environment and remains constant even when other identifiable characteristics such as IP address, MAC Address and Hostnames are changed.

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, input Data Manager (IDM) or Splunk Cloud instance that supports modular input data ingestion.
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 base URL for the CrowdStrike Cloud environment that the Falcon instance resides in.

Major Modifications

The following are some of the major modifications made to this version of the add-on that differ from previous versions:

1. Based on customer feedback the timestamp field value used for both the API query and the Splunk event timestamp is now configurable. Customers can configure the event timestamp value to be based on either the modified_timestamp or last_seen field value.

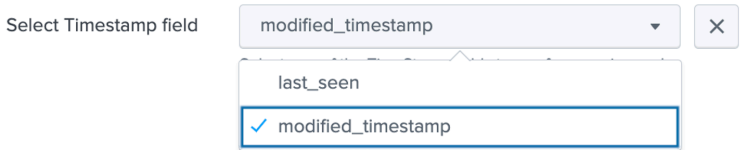

2. The ta_data section of the event has been enhanced with 3 new fields:

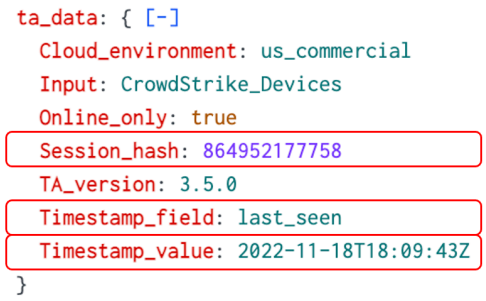

- **Session_hash**: A randomly generated hash that is created when the input starts up and is assigned to the events that were processed during that collection
- **Timestamp_field**: Indicates the fieldname selected in the configuration to use as the event timestamp
- **Timestamp_value**: Indicates the timestamp value associated with the TimeStamp_field selected and will be used for the event timestamp in Splunk (expressed in UTC)

Getting Started

API Endpoint(s), Filter(s) and Timestamp(s):

The TA will make API calls to some or all of the following endpoints. Some API calls may leverage different filter fields depending on the selected options.

API Endpoint	Filter fields(s)
/oauth2/token	
/devices/queries/devices-scroll/v1	modified_timestamp or last_seen
/devices/entities/online-state/v1	
/devices/entities/devices/v1 (3.1.1)	
/devices/entities/devices/v2 (3.1.6+)	

The default event timestamp used by the TA is the '**modified_timestamp**' value but can be configured in the input configuration page. The value selected under the 'Select Timestamp Field' in the configuration setting is what will be used in Splunk for the _time value. The field selection and the value can be located in the 'ta_data' section of the event as 'Timestamp_field' and 'Timestamp_value'. Note the 'Timestamp_value' will be reported in UTC.

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:

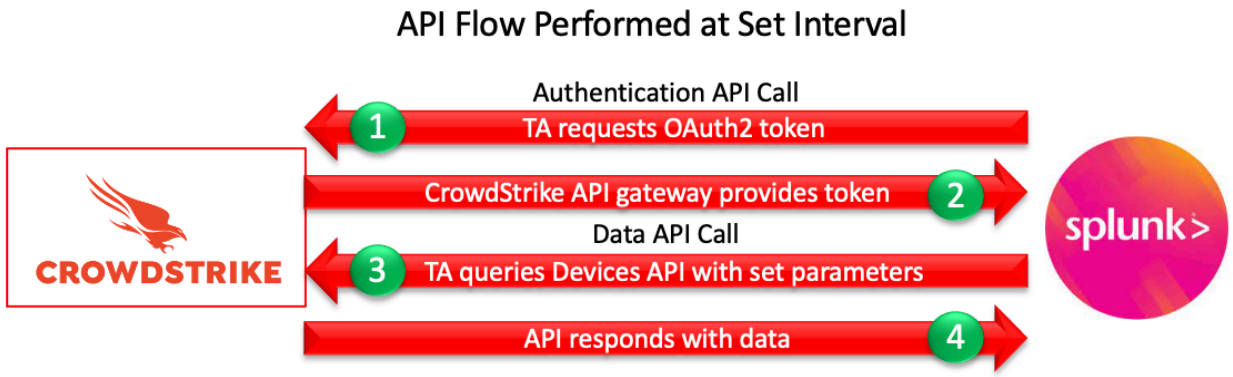

1. The TA will call the CrowdStrike API gateway with the configured credentials and request an OAuth2 authentication token that is valid for 30 minutes.
2. If the API credentials are valid the API gateway will respond to the TA with an Oauth2 token.
3. The TA will use the OAuth2 token to call the Devices API with the configured parameters.
4. The API will respond with whatever appropriate data matches the configured parameters.

Technical Add-On Layout

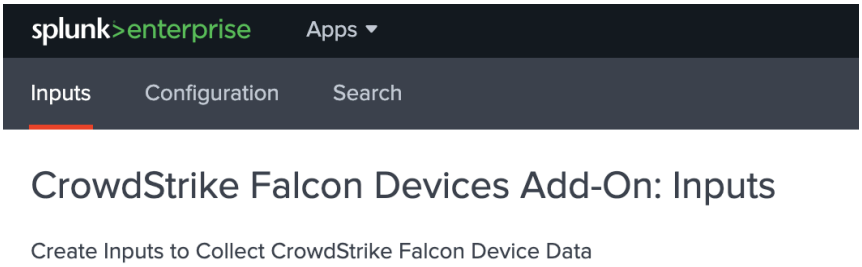

The CrowdStrike Falcon Devices TA has 3 tabs associated with it:

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

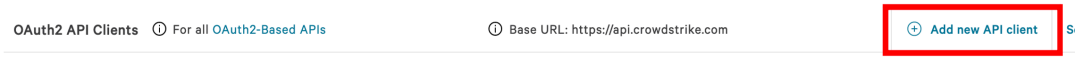

4. Give the new API client a name and description (recommended) and under 'API Scopes' select the 'Hosts' scope and the 'Read' capability

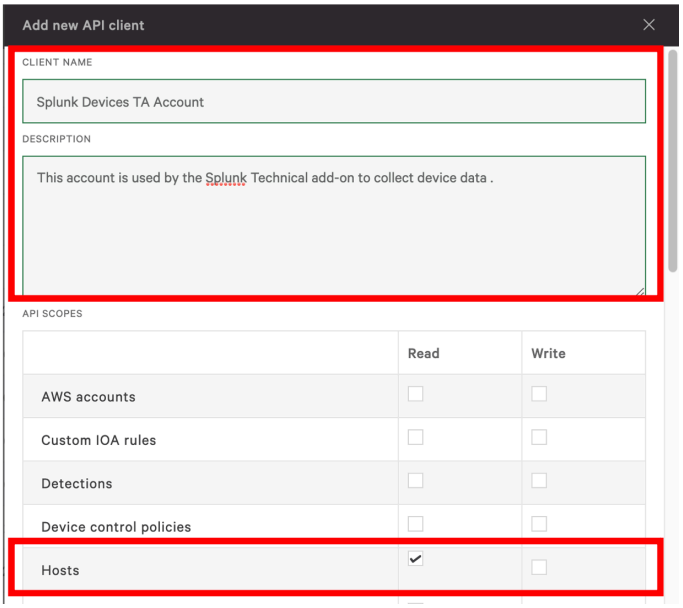

5. Select ‘Add’ once completed and a window will appear with the Client ID, Secret and the Base URL. NOTE: This is the only time the Secret will be visible – ensure it is recorded in a protected location.

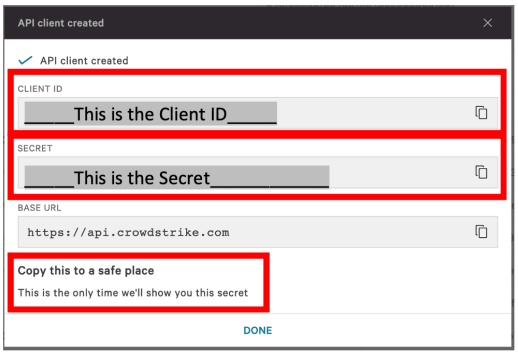

6. In addition, make note of the ‘BASE URL’ value in either the API client created window or the ‘OAuth2 API client’ area as this will be used to determine the CrowdStrike Cloud the instance is in

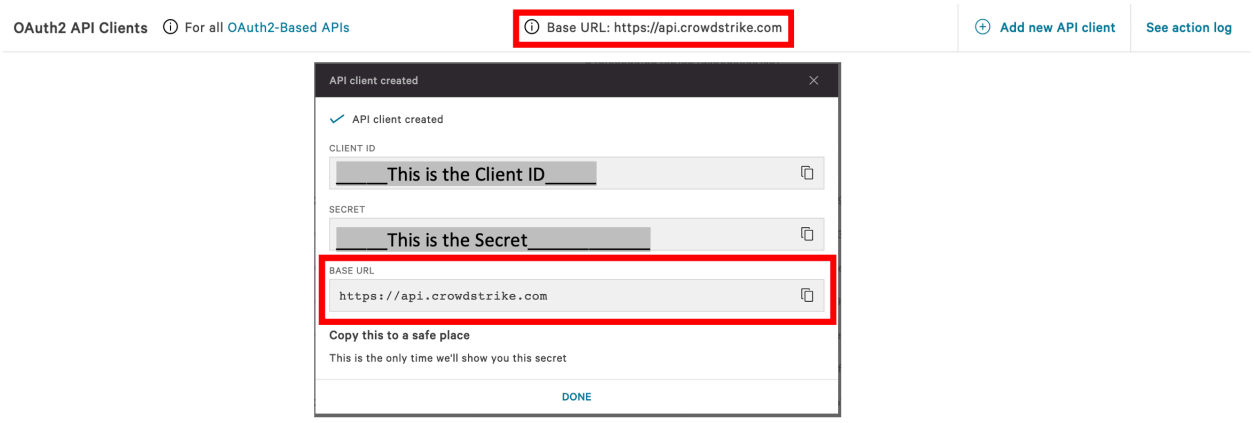

7. Select ‘Done’ to close the window and finish creating the credential

Proxy Considerations

The CrowdStrike Devices Technical Add-On establishes a secure connection with the Falcon cloud platform. In some environments network devices may impact the ability to establish and maintain a secure 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:

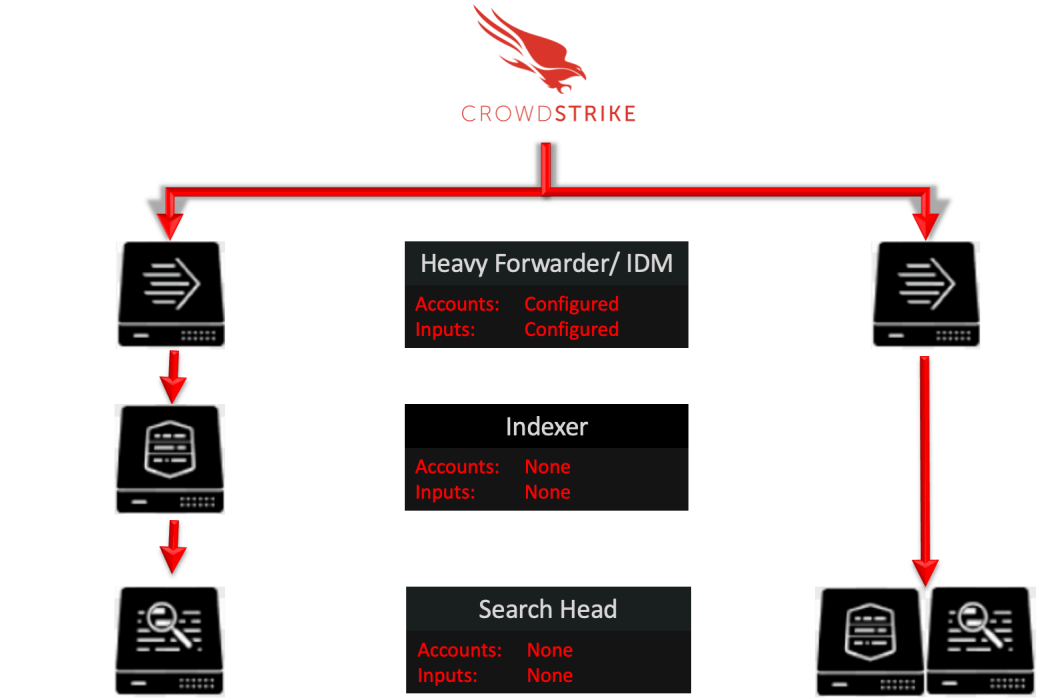

Splunk Enterprise: Distributed

Splunk Cloud

Configuring the TA

TA Layout

The TA contains 3 sections:

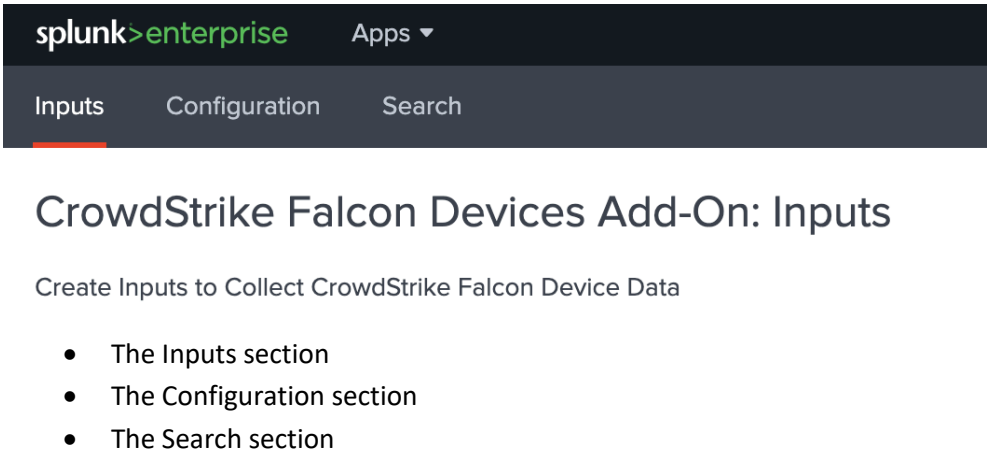

- The Inputs section
- The Configuration section
- The Search section

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.

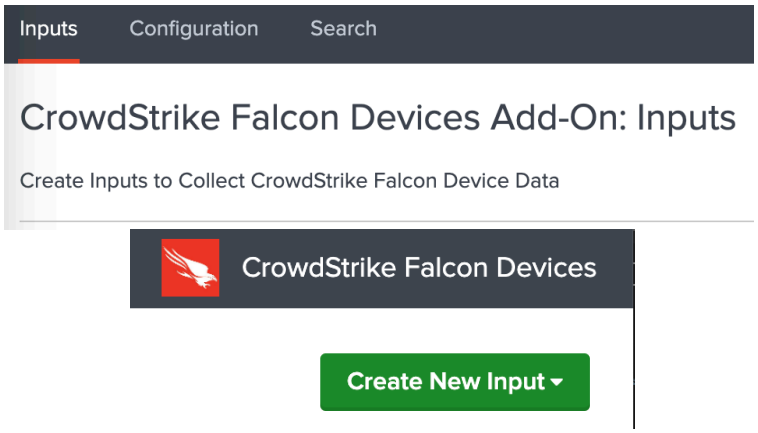

Configuration Section

The Configuration section contains 3 configuration tabs:

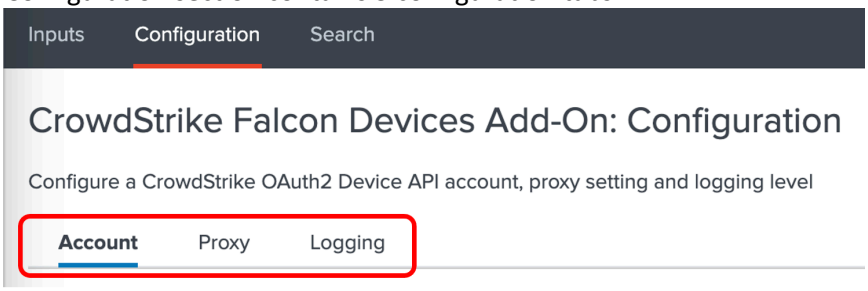

- **Account Tab**: This is where the OAuth2 API credentials are entered.
- **Proxy Tab**: This is where proxy server configurations are entered.
- **Logging Tab**: This is where the logging level is configured.

Search Section

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

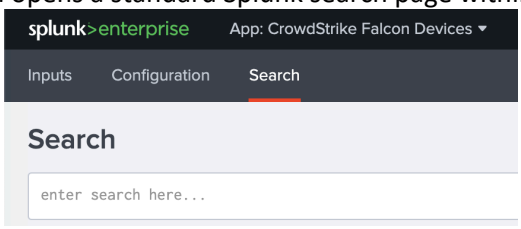

Configuring the TA to collect data

***NOTE* This action should only be performed on a Splunk instance designed for collecting data**

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.

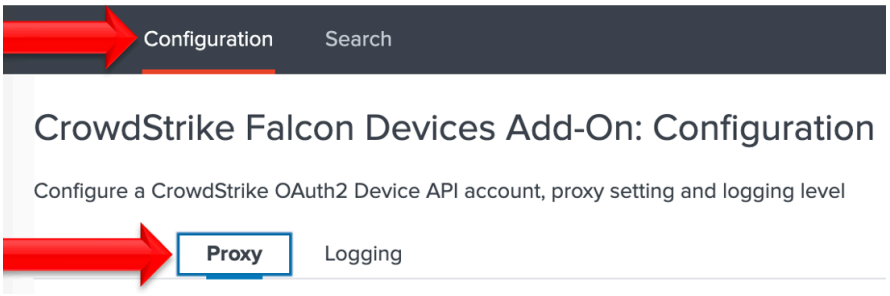

2. Configure the following fields as appropriate:

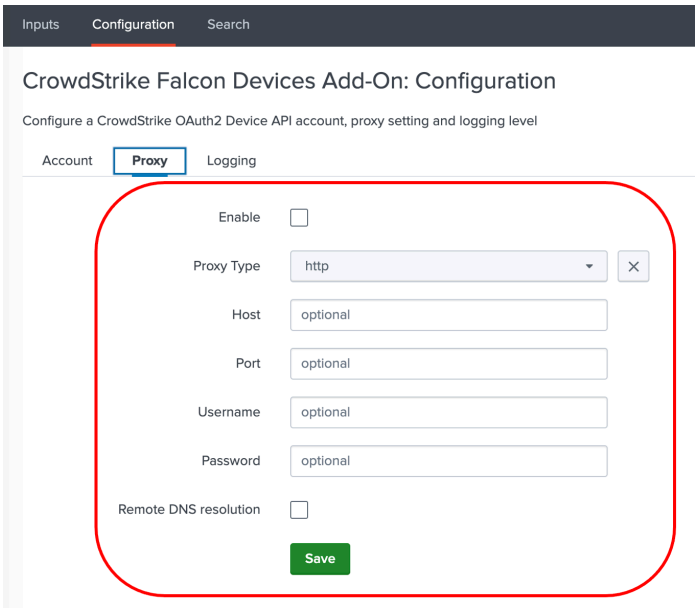

- **Enable**: This checkbox is used to enable/disable the proxy settings
- **Proxy Type**: This dropdown is used to select the proxy type
- **Host**: The hostname/IP address for the proxy server
- **Port**: The communication port for the proxy server
- **Username**: The authentication username for the 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:

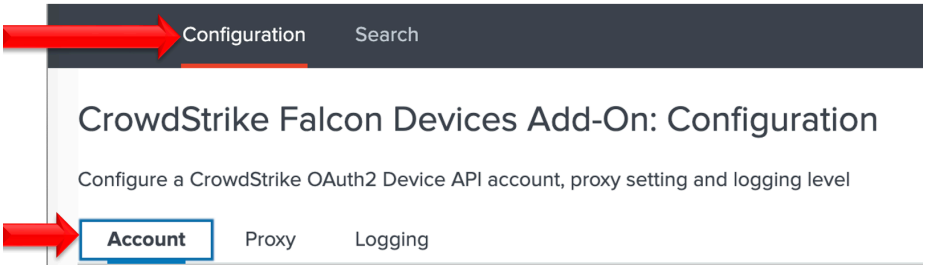

3. On the right side of the screen click the “Add” button:

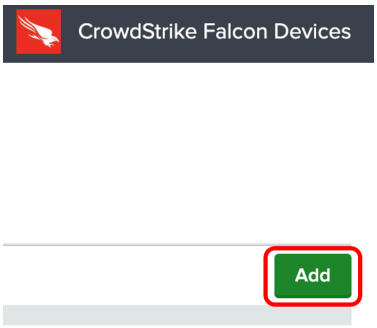

4. Configure the following fields:

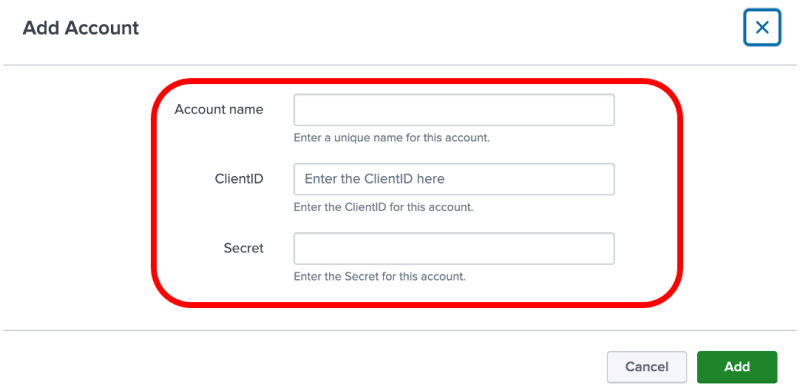

- **Account Name**: A name unique for the Splunk instance
- **ClientID**: The ClientID of the API credential created in the CrowdStrike Falcon UI
- **Secret**: The Secret of the API credential created in CrowdStrike Falcon UI

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 already created.
2. An input is created under the Inputs section:

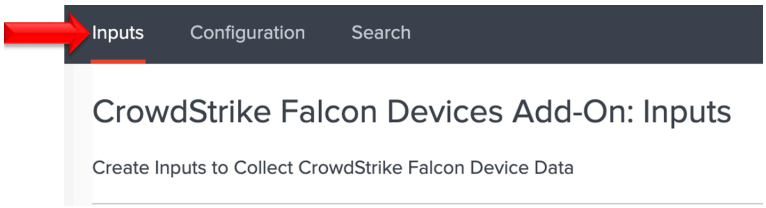

3. From the dropdown‘Create New Input’ in the top right corner select ‘CrowdStrike Device JSON’

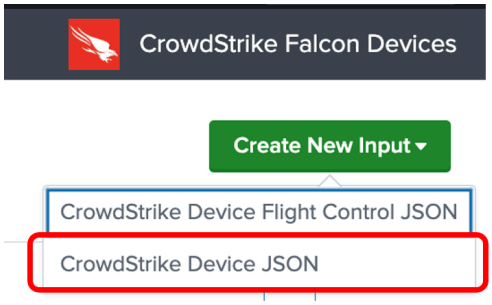

4. Configure the appropriate fields:

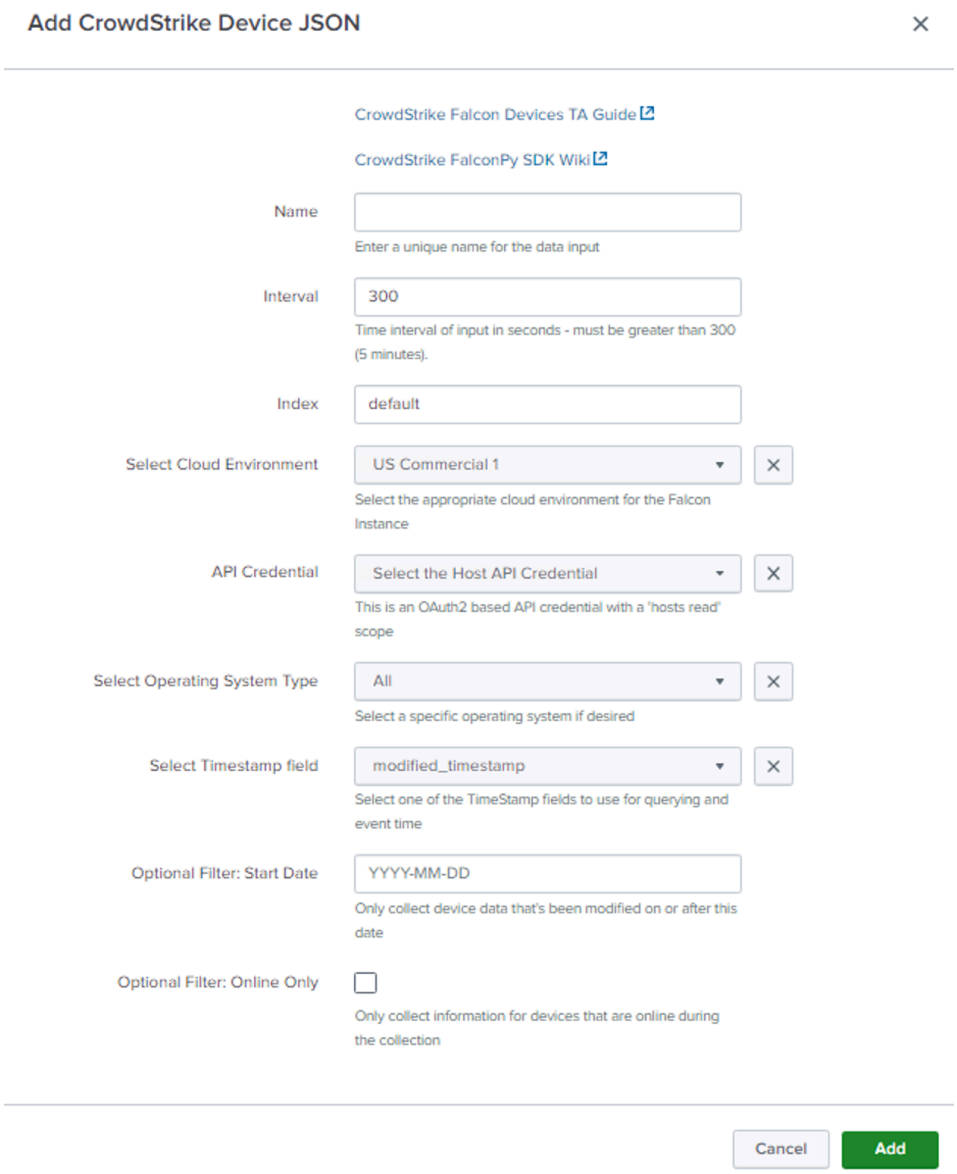

- **Name**: (required) A name unique to the Splunk Environment
- **Interval**: (required) How often the TA will collect data, expressed in seconds (the default and minimum is 300)
- **Index**: (required) The Splunk Index that the data will be stored in
- **Cloud Environment**: (required) The CrowdStrike 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
- **Timestamp Field**: (required) The field value used for API query filtering as well as the event time within Splunk
- **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 data for devices that are online at the time of the collection will be collected

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

Configure an Input: Flight Control

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

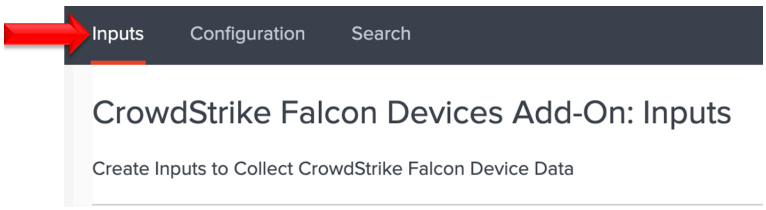

3. From the dropdown'Create New Input' in the top right corner select 'CrowdStrike Device JSON'

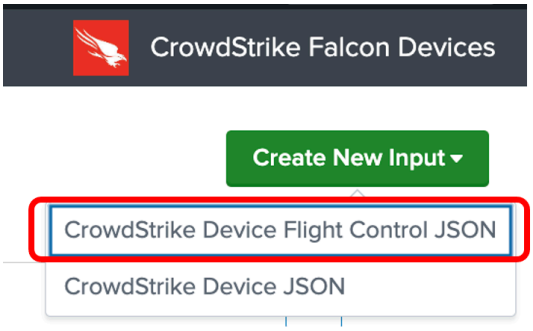

4. Configure the appropriate fields:

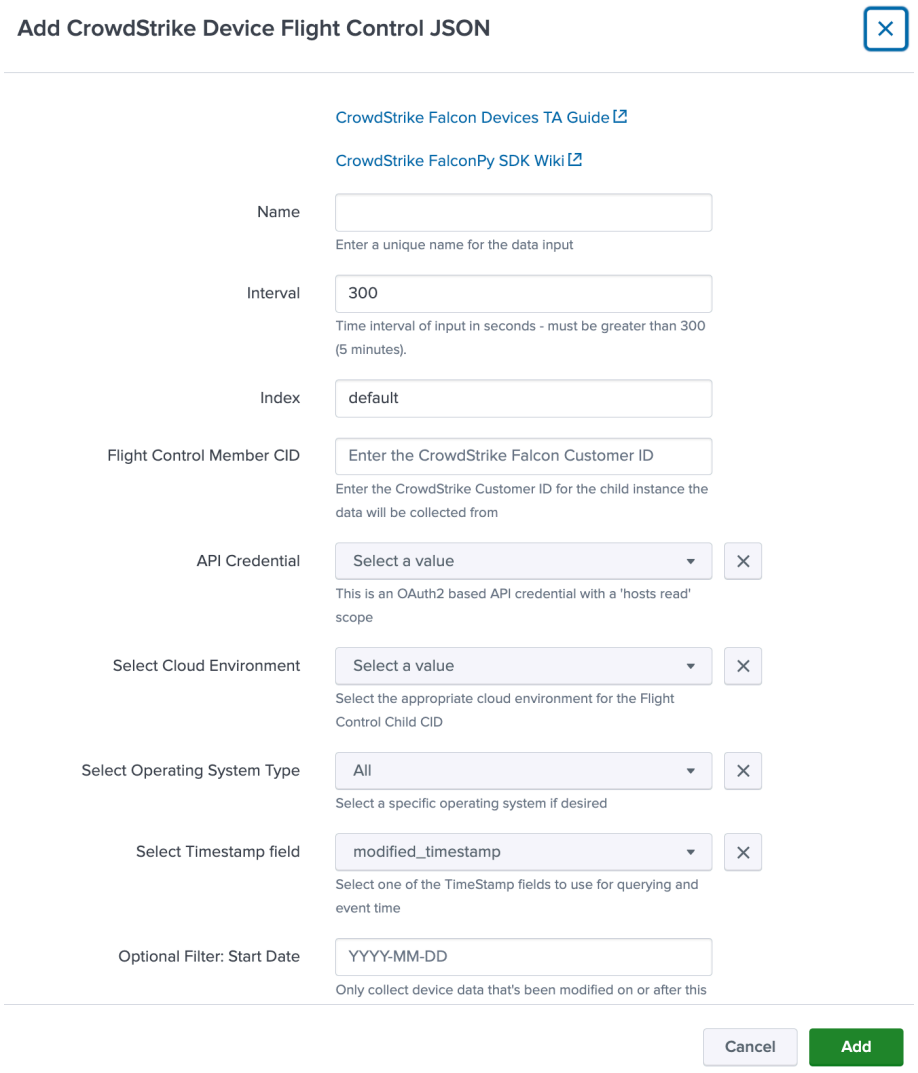

- **Name**: (required) A name unique to the Splunk Environment
- **Interval**: (required) How often the TA will collect data, expressed in seconds (the default and minimum is 300)
- **Index**: (required) The Splunk Index that the data will be stored in
- **Flight Control Member CID**: Falcon Customer ID for the Falcon Flight Control child instance
- **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 of the Falcon Flight Control Parent instance

- **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
- **Timestamp Field**: (required) The field value used for API query filtering as well as the event time within Splunk
- **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 data for devices that are online at the time of the collection will be collected

5. 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 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 :

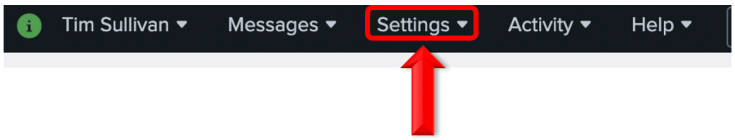

2. Select 'Advanced Search' from the dropdown menu:

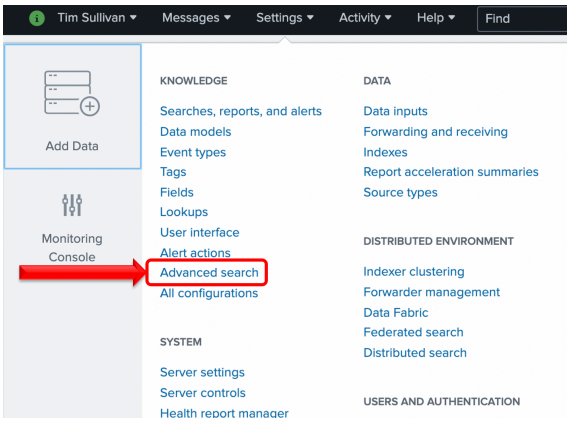

3. Under 'Advanced Search' select 'Search Macros:

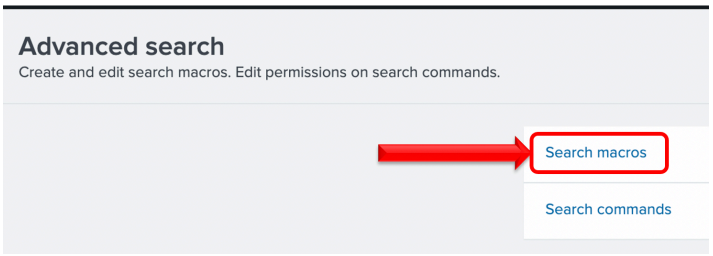

4. Ensure the 'CrowdStrike Falcon Device Technical Add-on' is selected in the 'App' selection and 'Created in App' is selected from the pulldown selection:

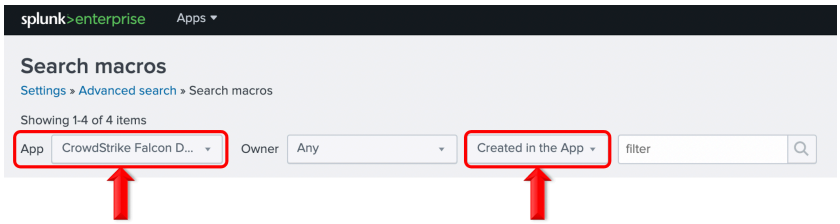

Configuring and Leveraging the Search Macros

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

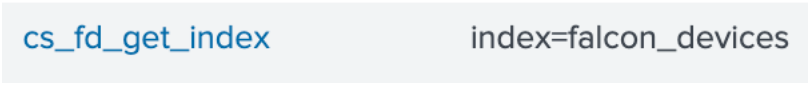

- **`cs_fd_get_index`** (CrowdStrike Falcon Device get index) A search macro that points to the index(es) that contain the data received by the TA inputs. The default for this search macro is to point to an index named 'falcon_devices' and should be adjusted to reflect the specific index(es) that the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

The other search macros that are available once the required search macro has been properly configured:

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 (they are not single quotation marks).

Search Macro Examples

`cs_fd_get_index`

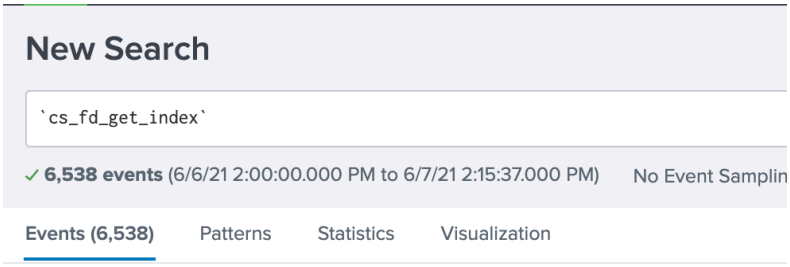

`cs_fd_device_hostname(Win10-001)`

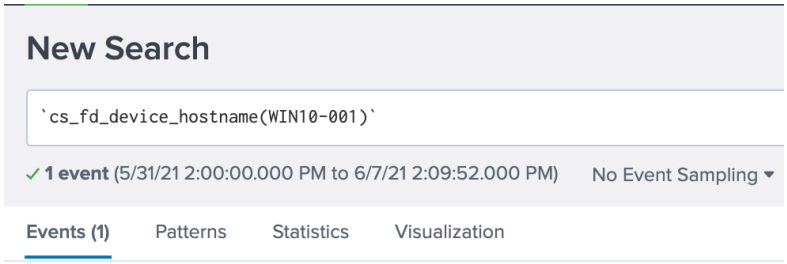

`cs_fd_get_device_id(c8b6q5716xa440408a29637ae244a0p1)`

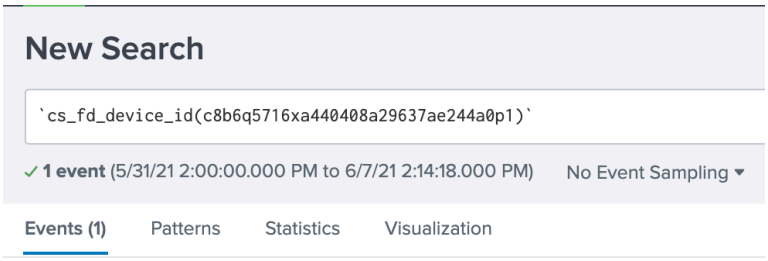

`cs_fd_get_ip(192.168.67.22)`

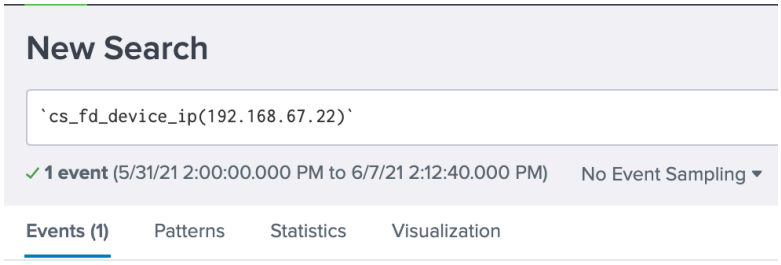

Reports

Locating the Reports

The reports included in the TA can be found by navigating to:

1. Open the Settings menu, under ‘knowledge’ select ‘Searches, reports and alerts’:

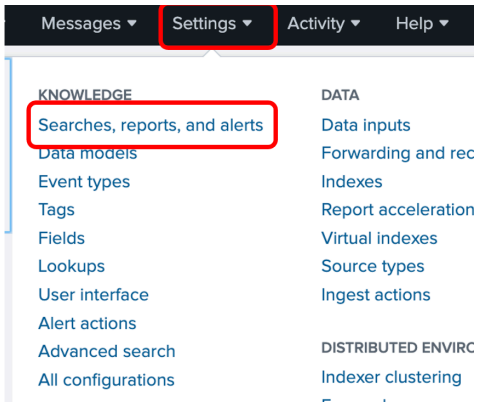

2. In the ‘App’ dropdown ensure that the ‘CrowdStrike Falcon Devices’ is selected and under the ‘Owner’ dropdown ensure that ‘All’ is selected

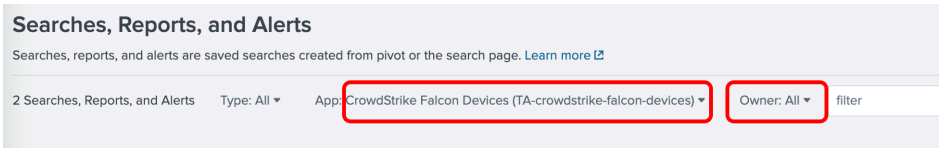

3. The reports that are included with the TA should be visible:

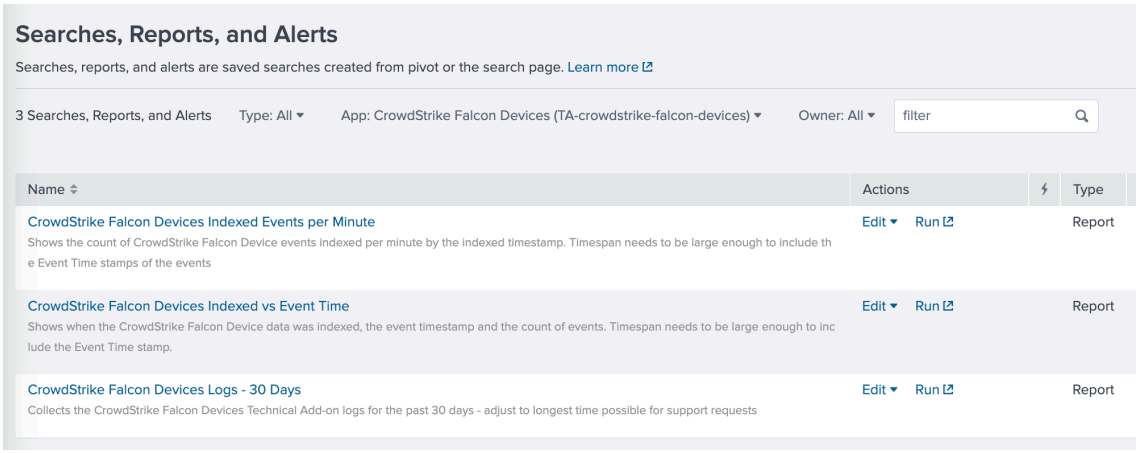

Understanding the Reports

These reports require the index search macro be correctly configured

- **CrowdStrike Falcon Devices Indexed Events per Minute**: Shows the number of events, per minute by default, that have been indexed by Splunk by the _indextime timestamp.
- **CrowdStrike Falcon Devices Indexed vs Event Time**: Shows the number of events that have been indexed per minute into Splunk by the _indextime timestamp and also shows event timestamp.

This report should be used to collected data for support tickets, export data in JSON

- **CrowdStrike Falcon Devices Logs – 30Days**: Collects events from the internal index for sourcetypes that contain the TA name string for the last 30 days by default. This report can be used to collect log data for support cases and should be set to 'all time' and exported in JSON format for those instances.

Recommendations

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

Use the CrowdStrike Device ID/AID for Searches

The 'device_id' value in the CrowdStrike data represents an environmentally unique value. (**Note**: this value is also referred to as the 'Host ID', 'Agent ID' and 'AID') This value remains constant even when typical 'identifying' values such as hostnames, IP addresses or MAC addresses are modified. Since this value is unique and is the primary field used to identify systems within the CrowdStrike environment, it is recommended to leverage this specific field value as the key field when performing searches and creating device-based information/reports/dashboards within the Splunk environment as well.

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 indexes based on the specific operating system. ie Windows devices would be stored in index 'Windows', MAC OSX devices would be stored in index 'OSX' and Linux devices would be stored in index 'Linux'.

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.

Dedicated API Credential

The use of a dedicated API credential for this integration is recommended to prevent issues should the credentials secret need to be regenerated and/or to ensure that the client is only scoped for the specific API endpoints used.

Interval Setting

The interval setting for inputs should take into account the amount of data that the input could potentially process. Setting an interval to low may result in a collection being interrupted and/or the data being collected having minimal changes. In most environments collection interval over 1 hour is recommended for standard inputs. Depending on the use case and the number of potential hosts involved, on-line inputs maybe slightly shorter.

The interval setting should also take into account the Timestamp field selection. The 'last_seen' value is not updated as often as the 'modified_timestamp' value and as such inputs based on that field value should have a larger interval.

Timestamp Field Selection

There is a distinct difference between the 'last_seen' and the 'modified_timestamp'. While a system is online the 'last_seen' timestamp typically will be updated every 20-40 minutes. On the other hand, the 'modified_timestamp' can be updated several times within a few minutes. Some key considerations to take into account when selecting either the 'last_seen' and 'modified_timestamp' field as the timestamp value are the following:

- **The number of events ingested into Splunk**: Since the 'modified_timestamp' value can be modified with far greater frequency than that 'last_seen' value, there typically are far more events ingested. Depending on the size of the environment and the interval setting for the input this can potentially impact licensing. If the 'modified_timestamp' is the value best suited to environment but the volume of events becomes an issue, the best course of action is to increase the interval setting.
- **Tracking Policy assignment regardless of online status**: In order for the device information to be collected using the 'last_seen' field value, the device has to have reported to the CrowdStrike cloud since the last device collection in a timeframe that will cause this field value in increase (typically after 20-40 minutes of the last field value). Conversely device information collected with the 'modified_timestamp' reflects changes regardless of if the device has reported to the CrowdStrike cloud. For example, if a set of laptops are assigned to a new policy group in CrowdStrike but they are not currently online, an input collecting data using the 'modified_timestamp' field value would reflect that policy change but an input collecting data using the 'last_seen' field value would not.

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
- Splunk environmental/configuration-based issues

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:

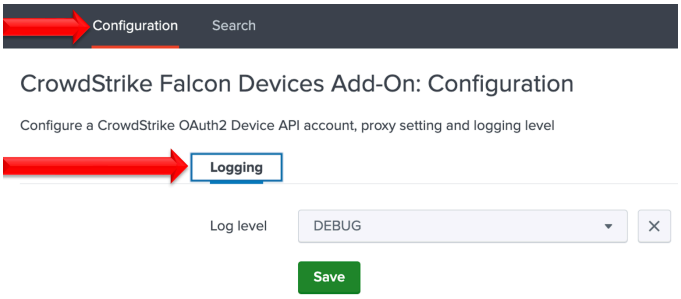

2. Select the logging level from the drop-down menu:

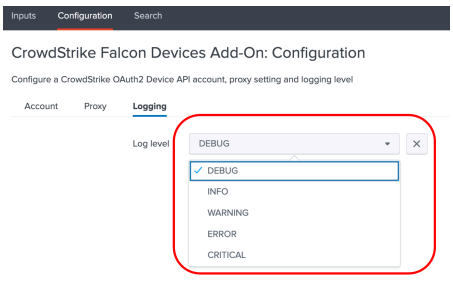

3. Click 'Save' to save the logging level.

Review Log Data in Splunk

1. Run and review the **CrowdStrike Falcon Devices Logs – 30 Days** Report to determine if there are any errors being reported by the TA.
2. Review Splunkd logs to determine if there's any internal issues within Splunk that could be causing issues with the proper collection and processing of data. If events related to a possible issue are found please include export them in JSON format and include them in any support requests.

Examples of Troubleshooting Situations and Remediation Steps

1. It doesn't look like any data is being collected:

1.1. Ensure that the credentials have been properly scoped for the API and have been properly entered
1.2. Ensure that the time picker selection is set to either 'all time' or that the time window is large enough to include the event timestamp. If the TA may be collecting events that are timestamped outside the currently selected time window
1.3. Run one of the reports or a search query that shows the events by _indextime to determine if and when events are being or have been indexed. Ensure that the time window is set to 'all time' to avoid not capturing data that has an event time outside the time window. Indexed time and event time can be drastically different
1.4. Examine log data to determine if any API calls are getting 401 or 403 responses indicating a potential issue with authentication, credential input
1.5. Ensure that firewalls, proxies and other network devices are not interfering with the communications between the TA and CrowdStrike API(s) and the TA and Splunk APIs

2. Data looks like it is coming in 'delayed':

*Data collected by the TA can be delayed in indexing because of factors outside of the control of the TA's functionality and may **not** be able to be identified or rectified by CrowdStrike*

2.1. Determine if there is potentially any latency in data communication between the Splunk system doing the data collection and the indexer tier
2.2. Determine if there's any latency in data being indexed at the indexing tier
2.3. Run one of the reports or a search query that shows the events by _indextime to determine if and when events are being or have been indexed. Ensure that the time window is set to 'all time' to avoid not capturing data that has an event time outside the time window. Indexed time and event time can be drastically different
2.4. Example the TA log data and compare it against the indexing time data to determine if they are aligning

3. **The data being collected does not look 'complete':**

    3.1. Review the input settings to ensure that the settings reflect that data collection requirements
    3.2. Review the interval setting to ensure that there is enough time to collect the required data and that data collections are not being interrupted
    3.3. Ensure that the Splunk search and the associated time window will encompass all the potential data
    3.4. Ensure that any potential discrepancy does not take into account hosts that may have been deleted
    3.5. Review the TA logs and the internal Splunk logs for any errors that may have impacted data collection
    3.6. Validate that there is not an internal Splunk issue that could be delaying the indexing of data

4. **There is a large number of events for all or a specific number of devices with little difference in the device data:**
    4.1. Review the Timestamp field value selection – the 'modified_timestamp' selection typically will generate a much greater number of events that the 'last_seen' selection
    4.2. Review the data leveraging the 'device_id' value as opposed to something like the 'host_name' value as host name is not considered a unique field in Falcon
    4.3. Ensure that the sensors have been deployed according to best practices, this is especially important in environments with a large number of virtual machines where the sensor is part of the 'golden image'
    4.4. Review the events with the 'Collection_hash' and 'Input' values to determine if they are part of a different API call and/or input

5. **There's too much data being collected:**
    5.1. Review the Timestamp field value selection – the 'modified_timestamp' selection typically will generate a much greater number of events that the 'last_seen' selection
    5.2. Review the interval setting – increasing the interval setting is the fastest and easiest way to control the amount of data that's being collected

Support

This TA Is designed to help facilitate the collection of device data provided by the CrowdStrike API(s). CrowdStrike provides support for the TA code functionality as it was designed.
Examples of instances that **would fall outside** of CrowdStrike's support:

- Environment caused network connectivity issues
- Issues related to certain Splunk configurations or internal Splunk connectivity issues
- Modifying the TA configuration outside of what's outlined in this documentation
- Deployments, configurations, modifications that do not align with what is outlined in this documentation
- Support requests without the appropriate data outlined below
- Splunk CIM field mapping or custom data modification requests
- Issues related to Splunk searches attempting to model the same data as found in the Falcon UI

Prior to Contacting CrowdStrike Support

1. Ensure that the OAuth2 credential has been scoped and entered correctly
2. Ensure that it is not an issue with the TA communicating with Splunk, modular inputs post data to API endpoints within Splunk so things like host firewalls can block this communication as can permission issues.
3. Ensure that the issue is not a network connectivity issue, if the API calls being made by the TA cannot properly communicate with the CrowdStrike API those issues should be resolved before contacting CrowdStrike support
4. Set the TA log level to 'DEBUG'
5. Repeat and record the action(s) that are associated with the issue you are reporting
6. Collect all appropriate log information
    a. Run the **CrowdStrike Falcon Devices Logs – 30 Days** Report with the time picker set to 'All Time' and export all the results in RAW format
    b. (If possible) Download the all-log files containing 'ta_crowdstrike_falcon_devices' under the $Splunk/var/log/splunk/ directory
    c. Collect any relevant logs from Splunk's internal log index related to the TA and the issue you're reporting
7. Record the following information about the Splunk system:
    - Splunk environment type
    - Splunk version
    - TA version
    - If this is was a new deployment/upgrade or if there was no change to the TA
    - The approximate date(s) and time(s) of examples of when the specific issue(s) occurred

Contacting CrowdStrike Support

1. Navigate to https://supportportal.crowdstrike.com/
2. Open a support ticket, provide the data collected in steps 6 & 7 above as well as any modifications that have been made to the TA outside of the processed outlined in this documentation

NOTE:

CrowdStrike technical support engineers (TSE) are required to evaluate Splunk integration support requests. In addition, CrowdStrike TSE are required to perform troubleshooting workflows to help identify potential issues and evaluate those issues for potential escalations to other teams. This may include, but is not limited to, requesting additional information/data/logs and requesting results from specific search queries or configurations modifications. The inability or unwillingness to supply the required/requested information and/or make request modifications/actions may result in CrowdStrike not being able to troubleshoot the reported issue and result in the inability to provide support for the reported issue.

Additional Resources

CrowdStrike Host and Host Group Management API Documentation
CrowdStrike FalconPy SDK

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**.