



CROWDSTRIKE

CrowdStrike Falcon Event Streams Add-on for Splunk

Installation and Configuration Guide V3+

Table of Contents

Major Release Tracking	4
Overview	5
Requirements	6
Major Modifications	6
Getting Started	1
<i>Enable Access to the Event Streams API</i>	1
<i>Proxy Considerations</i>	3
<i>Splunk Architecture</i>	3
<i>Add-On and Alert Action Logging</i>	4
Comparing and Upgrading v2.x to v3.x	5
<i>Overview</i>	5
<i>Comparison Matrix</i>	5
<i>Preparing to Upgrade v2.x to v3.x</i>	6
Initial Installation / Re-Installation / Manual Update	7
Heavy Forwarder/ Information Data Manager Configuration	9
<i>Proxy Configuration (Optional)</i>	10
<i>Event Streams TA Account Configuration</i>	11
<i>Event Streams TA Inputs Configuration</i>	13
Search Macro Configuration	16
Modify, Remove or Clone Existing Settings	19
<i>Inputs</i>	19
<i>Configuration: Accounts</i>	21
<i>Configuration: Logging</i>	23
Custom and Calculated Fields	24
<i>Custom Fields: ta_data</i>	24
<i>Calculated Fields</i>	25
Understanding the Event Streams API and Offset Values	26
<i>Understanding Multiple Data Feeds and Applds</i>	28

Token Refresh Check Alert & Restart Input Alert Action	29
<i>Configuring the custom alert to restart an input</i>	30
<i>Enabling the custom alert to restart an input</i>	32
Scheduled Reports and Alerts	34
<i>Locating the Reports and Alerts</i>	34
<i>Understanding the Reports and Alerts</i>	35
Recommendations	36
<i>Custom Indexes</i>	36
<i>Dedicated API Credential</i>	36
<i>Interval Setting</i>	36
<i>Event Filtering</i>	36
Troubleshooting	37
<i>Troubleshooting Overview</i>	37
How TA gets data into Splunk	37
Check the connection in the Falcon UI:	37
<i>Troubleshooting typical situations</i>	38
Support	40
<i>Prior to Contacting CrowdStrike Support</i>	40
<i>Contacting CrowdStrike Support</i>	41
About CrowdStrike	42

Major Release Tracking

V 2.0.9 – Released May 2021:

- Coding modifications for improved network communications
- Customer alert for inputs
- Custom action for inputs
- Additional search macros

V 2.1.0 – Released June 2021:

- A new search macro `cs_es_reset_action_details` added to collect more information about the Restart Input alert action
- Streamlined the search macro `cs_es_tc_input(1)` by removing and addresses a typo in a log statement
- Addressed a typo in a log statement related to the refreshing of OAuth2 tokens
- Modified the "CrowdStrike Token Refresh Check" saved search trigger condition to be: "Trigger alert when "Custom" search count < 2 "

V 3.1.0 – Released Jun 2022:

- Coding migration to FalconPy SDK (<https://github.com/CrowdStrike/falconpy>)
- Introduction of event filtering
- Ability to designate historic or current start time for event ingestion

Overview

This document outlines the deployment and configuration of the technology add-on for CrowdStrike Falcon Event Streams.

This technical add-on (TA) facilitates establishing a connecting to the CrowdStrike Event Streams API to receive event and audit data and index it in Splunk for further analysis, tracking and logging. The Event Streams Add-on v3.x+ represents a significant update to v2.x in regards to configuration, capabilities and codebase.

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, Prevent and or Horizon 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:

Feature	v2.x	v3.x
Create Dedicated Offset File	Automatic	Not Present/Leveraged
Filter by Event Types	No	Yes
Initial Starting Point	No	Yes

Due to changes in Splunk's use of SSL certificates in the most recent release, input restart alert actions may not function as designed or as they had in previous versions.

Getting Started


Prior to deploying the CrowdStrike Falcon Event Streams Technical Add-on (TA) ensure the following:

1. The latest version of the TA has been downloaded from Splunkbase
2. All Splunk systems that the TA will be deployed to have been identified
3. An account with proper access to identified Splunk systems is available
4. Properly scoped API credentials have been created and recorded from the Falcon UI
5. Any custom indexes being used have been created on the appropriate systems
6. (optional) – If the communication between Splunk and the Falcon platform will traverse a proxy server then appropriate configurations should be taken into account. If the connection will need to authenticate to the proxy, then appropriate credentials should be created and available.

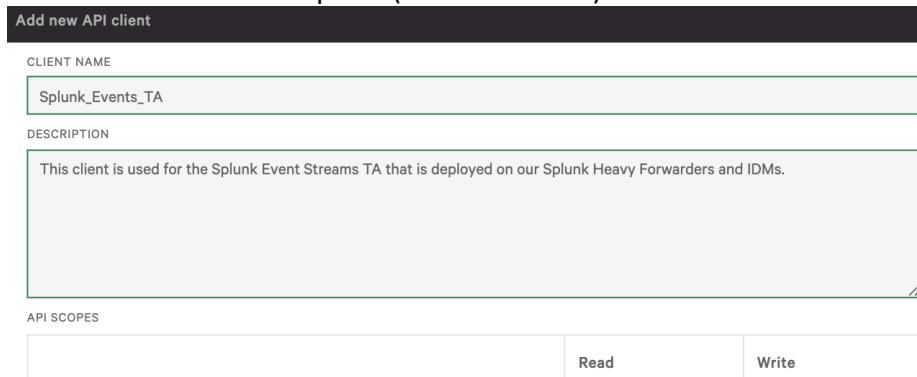
Enable Access to the Event Streams API

***NOTE: This process is not required if there is an existing API client with proper access but it is recommended to leverage a dedicated account for the TA.**

1. Log into the Falcon UI with an account that has administrator level permissions
2. Navigate to 'Support', 'API Clients and Keys' in the Falcon menu:
3. Select 'Add new API Client' to the right of 'OAuth2 API Clients':



4. Provide a client name and description (recommended):

A screenshot of the "Add new API client" form. The form has a dark header with the text "Add new API client". Below the header, there are three sections: "CLIENT NAME" with a text input field containing "Splunk_Events_TA"; "DESCRIPTION" with a larger text area containing the text "This client is used for the Splunk Event Streams TA that is deployed on our Splunk Heavy Forwarders and IDMs."; and "API SCOPES" with a table. The table has two columns, "Read" and "Write", each with a checkbox. The "Read" checkbox is checked.

5. Under 'API Scopes' select the 'Read' check box next to 'Event streams':

Spotlight vulnerabilities	<input type="checkbox"/>	—
Event streams	<input checked="" type="checkbox"/>	—
User management	<input type="checkbox"/>	<input type="checkbox"/>

6. Click 'ADD' to create the client:

7. A pop-up window will appear with the newly created Client ID and Secret
Ensure to record the secret correctly and store it in a safe place as this is the only time it will be visible/accessible

API client created ✕

✓ API client created

CLIENT ID
this15justasamplecliend1d 📄

SECRET
this15justasamples3cr3t007 📄

Copy this to a safe place
This is the only time we'll show you this secret

DONE

8. Once the credentials have successfully copied to a safe and secure location click 'DONE' to close the window:

API client created ✕

✓ API client created

CLIENT ID
this15justasamplecliend1d 📄

SECRET
this15justasamples3cr3t007 📄

Copy this to a safe place
This is the only time we'll show you this secret

DONE

Proxy Considerations

The CrowdStrike 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

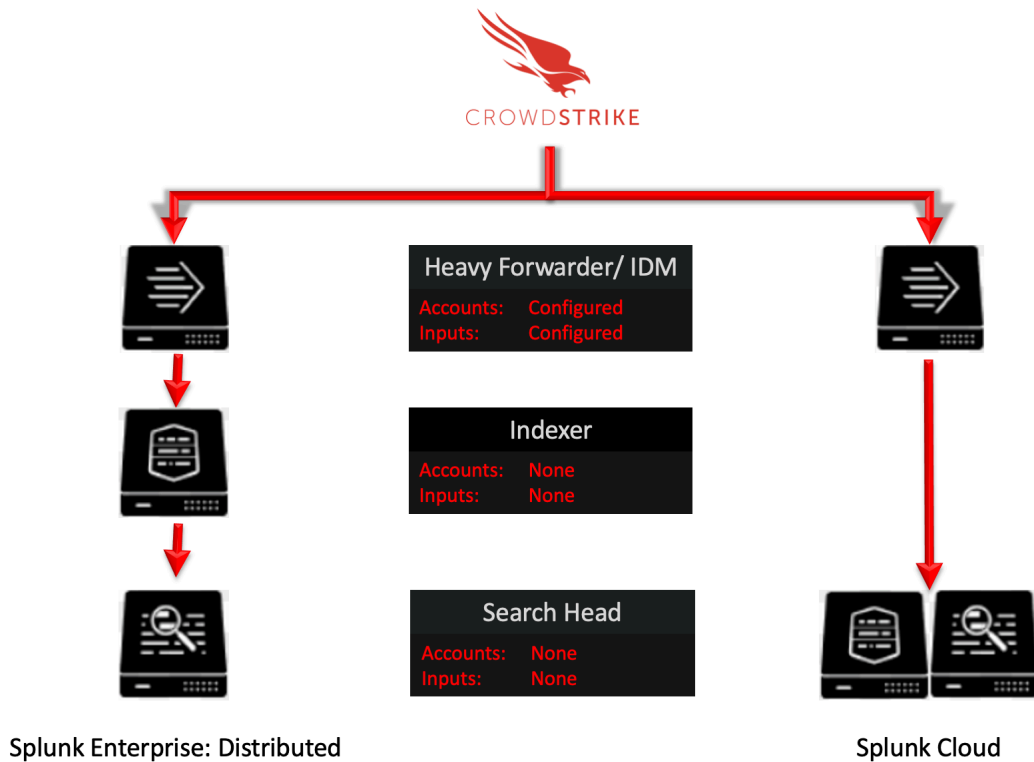
Splunk Search Head(s) and Splunk Cloud: The TA should be installed to provide field mapping and search macro support. These are often required to support CrowdStrike Apps. The TA should be deployed without any accounts or inputs configured and any search macros should be properly configured for use.

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 should be installed here as this is where the data from the Streaming API will be collected. The appropriate accounts or inputs should be properly configured for data collection. If the Heavy Forwarder is storing events 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 installed on Heavy Forwarders, Splunk Cloud IDMs and Splunk Cloud Victoria.

The following diagram shows the flow of data from the Streaming API and the Event Streams TA configuration within a distributed Splunk Enterprise and Splunk Cloud environment:



Add-On and Alert Action Logging

The Add-On logs can be found under: `$Splunk/var/log/splunk/` and begin with `'ta_crowdstrike_falcon_event_streams'`. These logs contain information about the configuration of the Add-On, API calls made to both CrowdStrike's API as well as the internal Splunk API's and other functionality

The Alert Action logs are separate from the Add-On logs but are also located under: `$Splunk/var/log/splunk/` and begin with `'crowdstrike_event_streams_restart_input_modalert'`

The search macros `'cs_es_reset_action_logs'` and `'cs_es_ta_logs'` can also be used to collect log data provided the `_internal` index data is available.

Comparing and Upgrading v2.x to v3.x

Overview

The Event Streams v3.X represents a significant update to the features, functionality and code base when compared to v2.x. As such it is important to examine the differences between the two and to properly prepare for upgrading from v2.x to v3.x.

Comparison Matrix

Feature	v2.x	v3.x
Create Dedicated Offset File	Automatic	Not Present/Leveraged
Filter by Event Types	No	Yes
Initial Starting Point	No	Yes

Create Dedicated Offset File: The dedicated offset file was introduced in v2.x as a backup for the data that's stored in the Splunk KVStore and to provide the option to manually configure the file with offset values to force new TA inputs to start at specific values. As this file is not accessible in Splunk Cloud instances and with the addition of the initial starting point option in v3.x, this feature has been phased out. (This feature is also not eligible to be added to the UI since there is no way for the TA to identify the number of datafeed URLs associated with a feed.)

Filter by Event Types: This feature is similar to the filter feature found in the CrowdStrike SIEM connector and allows customers to select predefined Event Types to include in their Event Stream data. This allows customers to create more custom data feeds based on event types and also allows for distributing data between multiple inputs to reduce the amount of data that a single input has to send into Splunk.

Initial Starting Point: This selection enables customer to indicate if they would like the input to collect all the available historical data available or start collection at the first available event when the input is the initially enabled.

Due to changes in Splunk's use of SSL certificates in the most recent release, input restart alert actions may not function as designed or as they had in previous versions.

Preparing to Upgrade v2.x to v3.x


This TA upgrade was designed to be an in-place upgrade, meaning that v3.x should be able to installed over v2.x without issue. However, there are steps that should be taken to prior to the upgrade to increase the probability of a successful upgrade.

1. **Ensure that Previously Collected Data is Backed Up:** As with any major upgrade, it's recommended that all data previously collected by the TA be backed up and/or placed in a location that will ensure that it will not be impacted by any issues that may arise from the upgrade.
2. **Disable Current Input(s):** In order to help prevent issue and/or data loss, it's recommended to disable any currently enabled input(s).
3. **Record Input Names:** The offset values recorded in Splunk's KVStore are recorded using the name of the input. In the event that there's an issue with an existing input, creating a new input with the exact same name (this is case sensitive) should result in the input retrieving the previously recorded offset values.

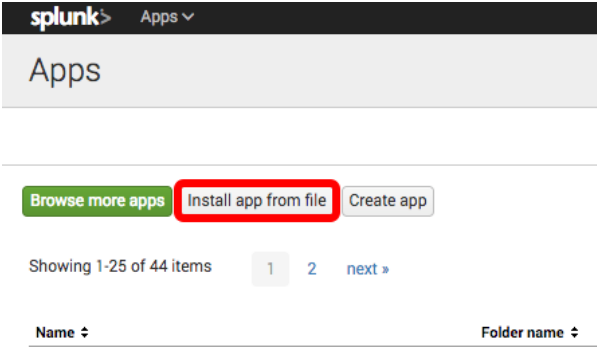
Initial Installation / Re-Installation / Manual Update

PERFORMING THIS ACTION REQUIRES A SYSTEM RESTART

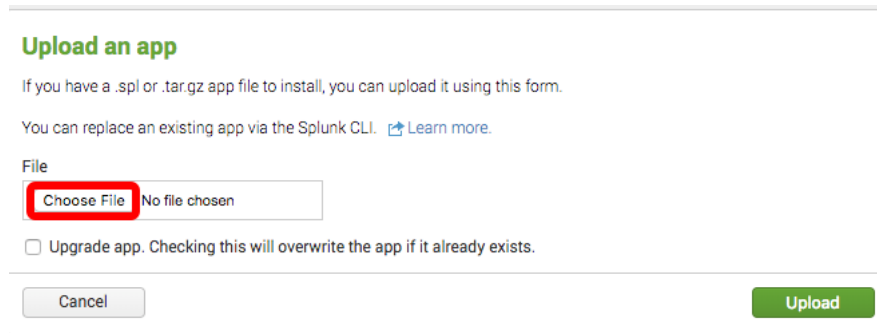
- 1. From the Splunk home page or dropdown menu select 'Manage Apps':



- 2. From the Manage Apps menu select 'Install app from file'



- From the 'Upload an app' window, select 'Choose File' *
*if upgrading or reinstalling an existing installation check the 'Upgrade app' checkbox



Upload an app

If you have a .spl or .tar.gz app file to install, you can upload it using this form.

You can replace an existing app via the Splunk CLI. [Learn more.](#)


File

Choose File No file chosen

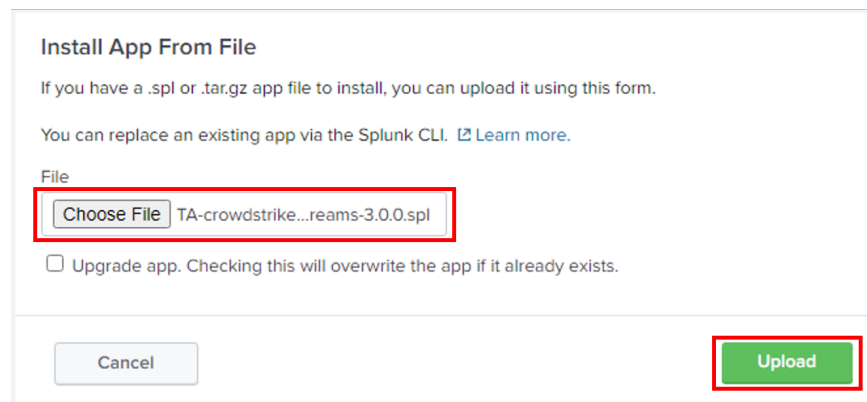
Upgrade app. Checking this will overwrite the app if it already exists.

Cancel Upload

- Select the downloaded Falcon Event Streams add-on file

 TA-crowdstrike-falcon-event-streams-3.0.0.spl

- Once the file is selected click 'Upload' to upload the add-on to system



Install App From File

If you have a .spl or .tar.gz app file to install, you can upload it using this form.

You can replace an existing app via the Splunk CLI. [Learn more.](#)

File

Choose File TA-crowdstrike...reams-3.0.0.spl

Upgrade app. Checking this will overwrite the app if it already exists.

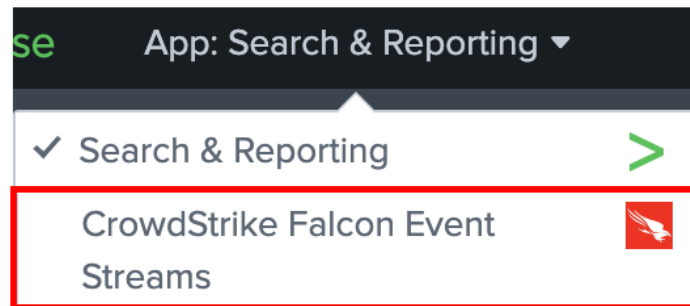
Cancel Upload

- Once the add-on has been installed the system will require a restart for the add-on to complete installation.

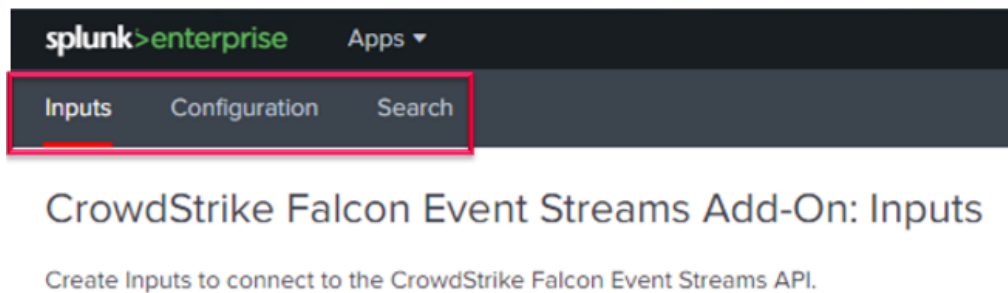
-----This concludes the Initial Installation / Re-Installation / Manual Update process-----

Heavy Forwarder/ Information Data Manager Configuration


1. From the Splunk drop down menu select 'CrowdStrike Falcon Event Streams'



2. There are three sub menus within the add-on: 'Inputs', 'Configuration' and 'Search'



3. Select the submenu 'Configuration'



Proxy Configuration (Optional)

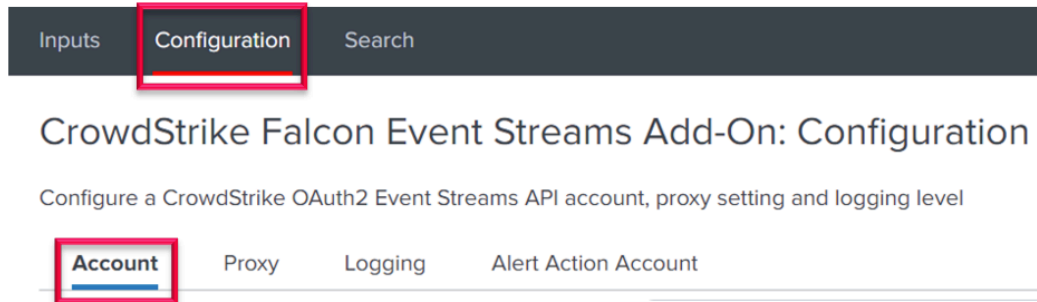
Select the 'Proxy' tab under 'Configuration' - Check the 'Enable' checkbox, select the Proxy Type from the drop down, enter the proxy host name, the proxy port and the credentials to allow communication.

The screenshot shows the 'CrowdStrike Falcon Event Streams Add-On: Configuration' page in Splunk Enterprise. The 'Proxy' tab is selected and highlighted with a red box. Below the tabs, there are several configuration options:

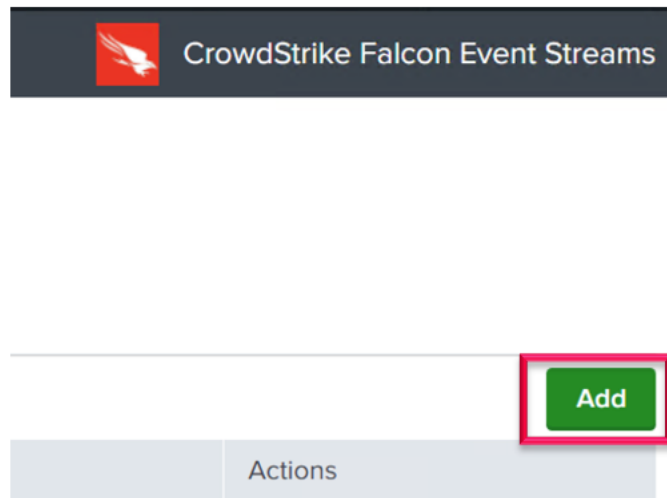
- Enable:** A checkbox that is checked, with a red arrow pointing to it.
- Proxy Type:** A dropdown menu set to 'http', with a red arrow pointing to it.
- Host:** A text input field containing 'myoutboundproxy.com', with a red arrow pointing to it.
- Port:** A text input field containing '8080', with a red arrow pointing to it.
- Username:** A text input field containing 'CSEventStreams', with a red arrow pointing to it.
- Password:** A password input field with masked characters, with a red arrow pointing to it.
- Remote DNS resolution:** An unchecked checkbox.
- Save:** A green button at the bottom.

Event Streams TA Account Configuration

1. Select the 'Account' tab under 'Configuration'



2. On the right-hand side select 'Add'



3. Configure the account for the Event Stream by providing the following:

- **Account Name** – This is a unique name for the account within Splunk
- **ClientID** – This is the ClientID for the API credential created
- **Secret** – This is the Secret for the API credential created

The screenshot shows a dialog box titled "Add Account" with a close button (X) in the top right corner. It contains three input fields, each with a red border:

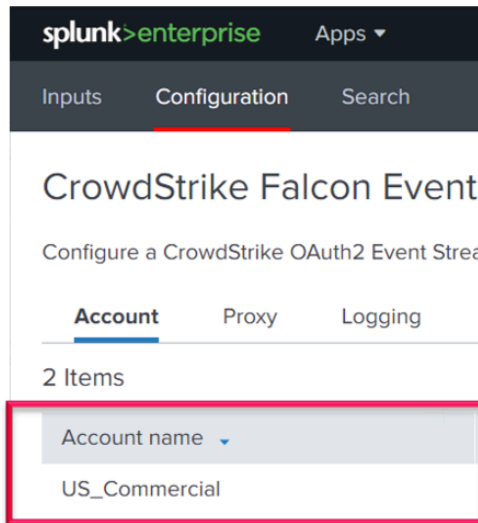
- Account name:** The input field contains "US_Commercial". Below it is the placeholder text "Enter a unique name for this account."
- ClientID:** The input field contains "This_is_where_the_ClientID_goes". Below it is the placeholder text "Enter the ClientID for this account."
- Secret:** The input field contains a masked value ".....". Below it is the placeholder text "Enter the Secret for this account."

At the bottom right of the dialog, there are two buttons: a grey "Cancel" button and a green "Add" button.

4. Once the information has been entered correctly click 'Add'



5. Validate that the account was saved successfully




Event Streams TA Inputs Configuration

1. From the Splunk drop down menu select 'CrowdStrike Falcon Event Streams'




2. There are three sub menus within the add-on: 'Inputs', 'Configuration' and 'Search'




CrowdStrike Falcon Event Streams Add-On: Inputs

3. Select the "Inputs" sub menu:



CrowdStrike Falcon Event Streams Add-On: Inputs

4. On the right-hand side, click 'Create New Input'



5. Configure the input for the Event Stream by indicating the following:


The screenshot shows a configuration window titled "Add CrowdStrike Event Streams" with a close button (X) in the top right corner. The window contains the following fields and options:

- Name:** A text input field containing "US_Commercial_ALL". Below it is the instruction: "Enter a unique name for the data input".
- Interval:** A text input field containing "0". Below it is the instruction: "Time interval must be 0".
- Index:** A text input field containing "event_streams_ALL".
- Select Cloud Environment:** A dropdown menu showing "US Commercial" with a clear button (X) to its right. Below it is the instruction: "Select the appropriate cloud environment for the Falcon Instance".
- API Credential:** A dropdown menu showing "US_Commercial" with a clear button (X) to its right. Below it is the instruction: "This is an OAuth2 based API credential with Event Streams scope".
- Application ID:** A text input field containing "US_Comm_ALL". Below it is the instruction: "Application IDs must be a unique value per CrowdStrike Instance".
- Event Types:** A text input field containing "All x". Below it is the instruction: "Select specific event type(s) to collect using this input".
- Initial Starting Point:** A radio button group with two options: "Historical" (selected) and "Now". Below it is the instruction: "Select an event collection starting point option (only used on the initial collection)".

At the bottom right of the window, there are two buttons: "Cancel" and "Add".

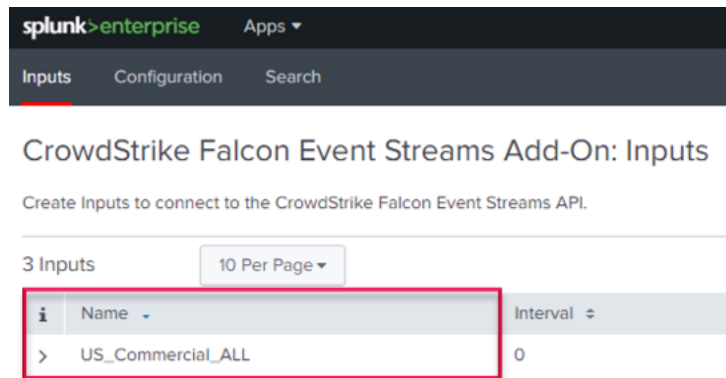
- **Name** – The Splunk unique name for the input being configured
- **Interval** – This field is not used in this TA and will always be set to 0
- **Index** – The index that the data will be stored in (must be an existing index)
- **Cloud Environment** – The CrowdStrike cloud environment the Falcon instance being connected to resides in
- **API Credential** – The corresponding API credential for the Falcon instance in the select Cloud Environment
- **Application ID** – An identifier for the API calls being made back to CrowdStrike (15 character maximum)
- **Event Types** – Identifies the Event Type categories that this input will collect *
- **Initial Starting Point** – Identifies if collection should start as far back as possible or from the first API connection forward (initial collection only)

6. Once the Input parameters have been correctly configured click 'add'*



***Newly created inputs are enabled by default**

7. Ensure that the Input has been successfully saved



*For information on the events types please refer to the Event Streams Event Guide located in the Falcon UI - <https://falcon.crowdstrike.com/documentation/62/streaming-api-event-dictionary>

This concludes the Heavy Forwarder/Information Data Manager Configuration process

Search Macro Configuration

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>

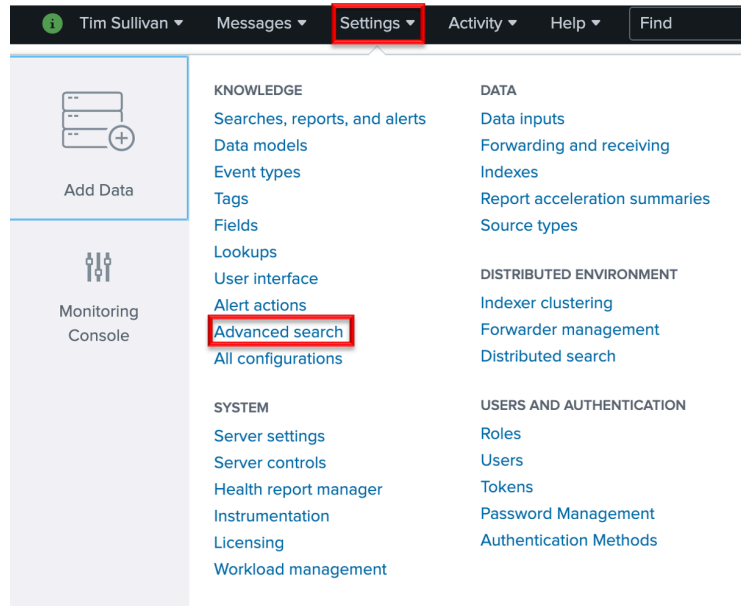
The Event Streams TA contains the following search macros:

Name ↕	Definition ↕	Arguments ↕
<code>cs_es_get_index</code>	<code>index=***</code>	
<code>cs_es_reset_action_logs</code>	<code>index=_internal sourcetype=splunkd component=sendmodalert action=restart_CS_ES_input</code>	
<code>cs_es_ta_logs</code>	<code>index=_internal sourcetype=tacrowdstrikefalconeventstreams:log</code>	
<code>cs_es_tc_input(1)</code>	<code>index=_internal sourcetype=tacrowdstrikefalconeventstreams:log "OAuth2 Token was successfully" *for input: \$input\$*</code>	input


- **cs_es_get_index** (CrowdStrike Event Streams get index) A search macro that points to the index(es) that contain the data received from the Event Stream API. The default for this search macro is to point to all indexes to search for data but should be adjusted to reflect the specific index(es) that the Heavy Forwarder/IDMs are pushing the data to.
- **cs_es_reset_action_logs**: A search macro that provides access to the 'CrowdStrike Event Streams – Restart Input' alert action logs.
- **cs_es_ta_logs**: A search macro that provides access to the CrowdStrike Event Streams TA logs.
- **cs_es_tc_input(1)**: A search macro that's designed to work in conjunction with the 'CrowdStrike Event Streams – Restart Input' alert action. This search macro requires that an input name be declared.
 - The (1) in the search macro indicates that 1 variable needs to be provided – this would be the name of the CrowdStrike Event Stream TA input that is to be monitored.

The search macros can be modified as follows:

1. Select the 'Settings' dropdown menu in the Splunk bar and select 'Advanced Search'




2. Select 'Search macros'




3. Configure the search settings as follows:

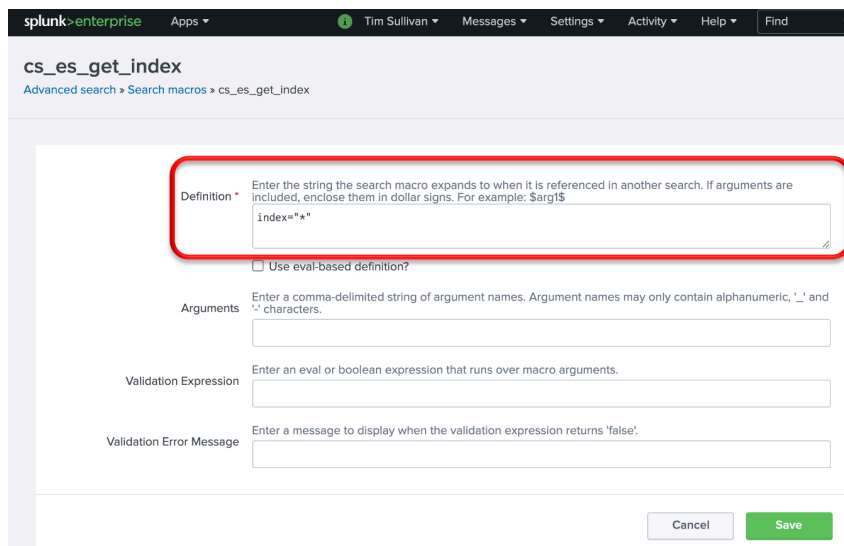
- **App** - CrowdStrike Falcon Event Streams
- **Owner** - Any
- Created in the App




- To modify a search macro, click on the name of the macro (`cs_es_get_index` for example):



- Under 'Definition' enter the index or indexes that contain the Event Stream data to the right of "index=" – separate multiple indexes with the "OR" Boolean.



- To leverage a search macro, open a search window within Splunk and enter the search macro enclosed with backquotes: ``cs_es_get_index`` (the backquote key is the same key as a tilde on a US keyboard layout and should not be confused with a single quote)

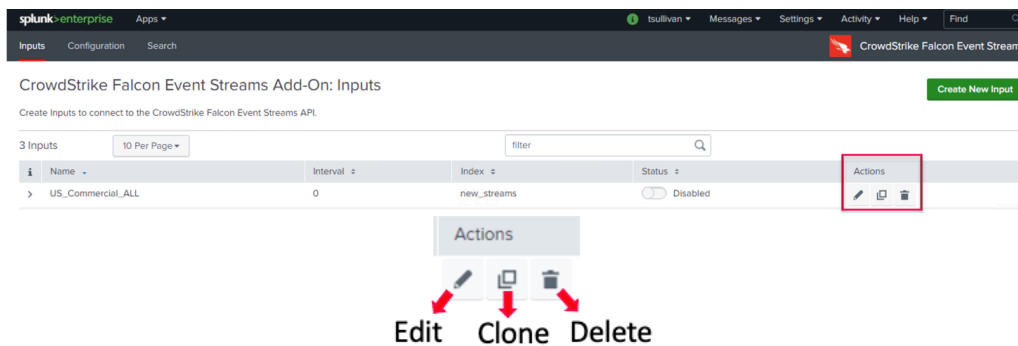


This concludes the Search Macro Configuration process


Modify, Remove or Clone Existing Settings

Inputs

1. Under the “Inputs” tab, under the “Action” column for an input, there are 3 options: “Edit”, “Clone” or “Delete”



2. **Editing:** allows for changing all the input fields with the exception of the input’s original name



3. **Deleting:** allows for the input to be deleted

4. **Cloning:** allows all the settings of the input to be replicated with the exception of the “Name” field*

*Note the Application ID must be unique per CrowdStrike Instance

Clone CrowdStrike Event Streams

Name
Enter a unique name for the data input

Interval
Time interval must be 0.

Index

Select Cloud Environment
Select the appropriate cloud environment for the Falcon Instance

API Credential
This is an OAuth2 based API credential with Event Streams scope

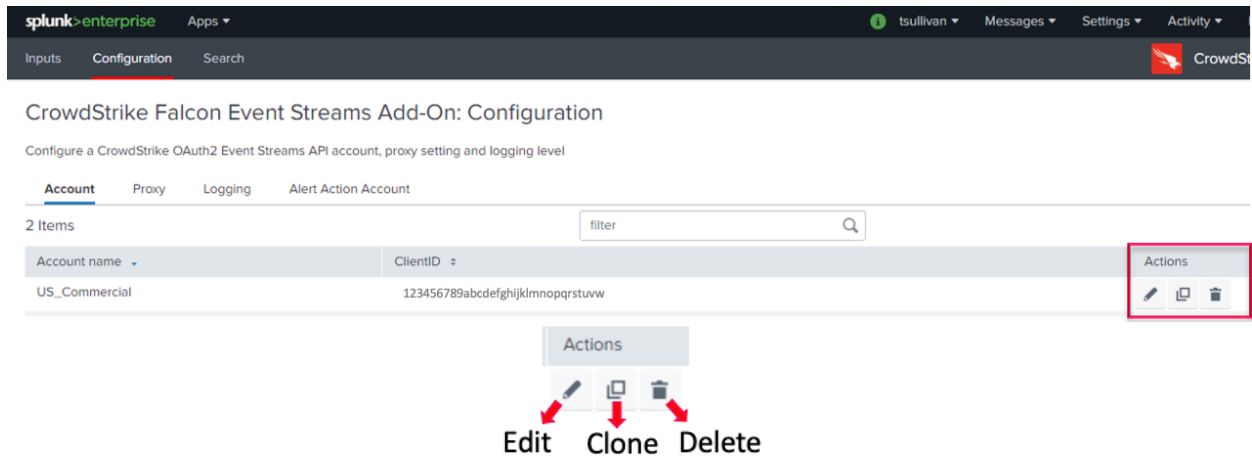
Application ID * SEE NOTE ABOVE
Application IDs must be a unique value per CrowdStrike Instance

Event Types
Select specific event type(s) to collect using this input

Initial Starting Point
Select an event collection starting point option (only used on the initial collection)

Configuration: Accounts

1. Under the “Configuration” sub menu, “Account” tab, and the “Actions” column for an account, there is a graphical menu with the options for “Edit”, “Clone” or “Delete”



2. **Editing:** allows for the changing of the ClientID and Secret - the name is NOT able to be edited once created

The 'Update Account' dialog box contains the following fields and buttons:

- Account name * (Note: Enter a unique name for this account.)
- ClientID * (Note: Enter the ClientID for this account.)
- Secret * (Note: Enter the Secret for this account.)
- Buttons: Cancel, Update

3. **Deleting:** allows a configuration to be deleted however it has to be removed from all inputs before this can be accomplished

The 'Delete Confirmation' dialog box displays the following message and button:

- Message: "US_Commercial_1*" cannot be deleted because it is in use
- Button: OK

4. **Cloning:** allows for a second account to be created with the same ClientID as the original but requires a new Account Name and Secret to be entered

Clone Account ×


Account name *
Enter a unique name for this account.

ClientID *
Enter the ClientID for this account.

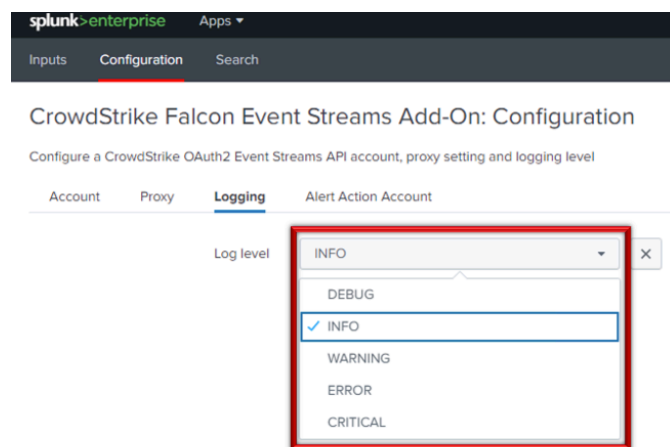
Secret *
Enter the Secret for this account.

Configuration: Logging

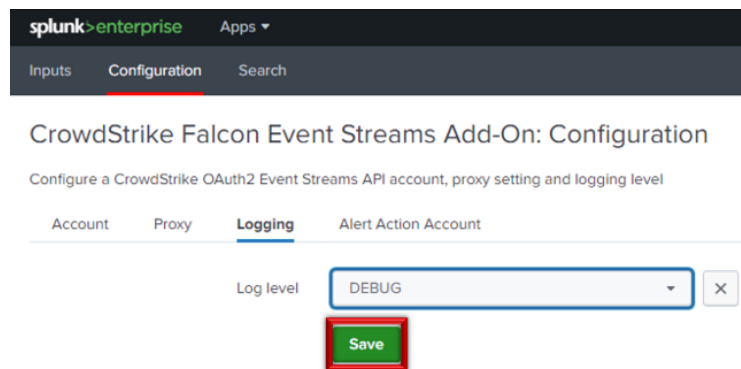
1. Under the “Configuration” sub menu, “Account” tab, and the “Actions” column for an account, there is a pull-down menu for setting the Log Level – which is ‘INFO’ by default



2. The TA provides the typical log levels available for a modular input. Those levels are (from most to least verbose): DEBUG, INFO, WARN, ERROR, FATAL.



3. To select a new Log Level, select it from the drop down and click “Save”



Custom and Calculated Fields

Custom Fields: ta_data

The Event Streams TA creates a custom information section and adds into all events to provide valuable information on the origin of the data and to assist in troubleshooting.

```
ta_data: { [-]
  App_id: tls_testing
  Cloud_environment: us_commercial
  Event_types: ['All']
  Feed_id: 0
  Initial_start: historic
  Input: Demo_2
  Multiple_feeds: False
  TA_version: 3.1.5
}
```

- **ta_data** - The name of the data section that provides the custom TA data
- **App_id** - The AppID configured for the input
- **Cloud_environment** – The cloud environment selected for the Input
- **Event_types** – The API event types that the input is configured to collect
- **Feed_id** – The id number for the data URL feed
- **Initial_start** – Identifies is the initial start included historic events or not
- **Input** – The name of the configured Input that received the data
- **Multiple_feeds** – Indicates if the Event Stream contains multiple data URLs
- **TA_version** – Data pulled from the TA configuration file and indicates the version of the TA

Calculated Fields

There are two calculated fields created in this TA:

The screenshot shows the Splunk Enterprise interface for the 'CrowdStrike Falcon Event Streams' table. The 'Calculated fields' section displays two fields:

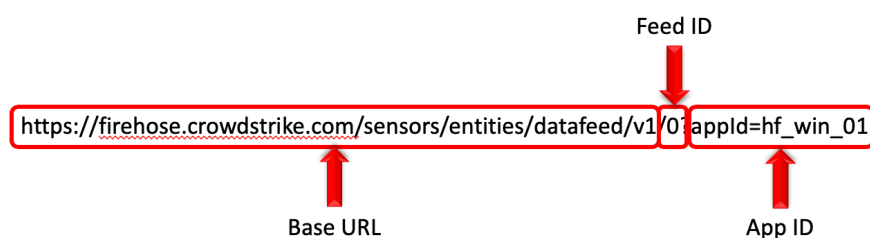
Name	Field name	Eval expression	Owner	App	Sharing	Status	Actions
CrowdStrike.Event.Streams:JSON : EVAL-action	action	case(event.PatternDispositionValue== 0, "allowed",event.PatternDispositionValue' == 128,"allowed", 'event.PatternDispositionValue' == 272, "allowed",event.PatternDispositionValue' == 768, "allowed",event.PatternDispositionValue' == 1280, "allowed",event.PatternDispositionValue' == 2304, "allowed",event.PatternDispositionValue' == 4638, "allowed",event.PatternDispositionValue==16,"blocked",event.PatternDispositionValue==512,"blocked", 'event.PatternDispositionValue==1024,"blocked", 'event.PatternDispositionValue==2048, "blocked", 'event.PatternDispositionValue==2176, "blocked", 'event.PatternDispositionValue==4096, "blocked", 'event.PatternDispositionValue==4112, "blocked", 'event.Success=="false", "failure", 'event.Success=="true", "success")	No owner	TA-crowdstrike-falcon-event-streams	Global Permissions	Enabled	Clone
CrowdStrike.Event.Streams:JSON : EVAL-vendor_product	vendor_product	"CrowdStrike Falcon"	No owner	TA-crowdstrike-falcon-event-streams	Global Permissions	Enabled	Clone

- **Action** – This field is calculated to be able to map to the 'Action' field in both the authentication and malware CIM (Common Information Model) tables
 - Authentication – The 'event.Success' field for authentication events is evaluated to provide the correct value
 - Malware – the numerical value of the 'event.PatternDispositionValue' is evaluated to provide the correct value
- **Vendor_product** – is calculated based on the source type to indicate that it was from CrowdStrike's Falcon platform

Understanding the Event Streams API and Offset Values

The CrowdStrike Event Streams API provides a substantial amount of data. In some instances, the amount of data is large enough that it is not feasible for a single URL to provide it all and the information is broken up into multiple data URL feeds. This is transparent to the end user for the most part and takes place during the API authentication process. Once the credential is authenticated the API will provide a list of data URL feeds that the client needs to connect to for data collection. All data URL feed connection must be successfully established and maintained to ensure all the appropriate data is being collected.

The data URL feed format is as follows:



- **Base URL** – The cloud environment’s base URL for the CrowdStrike Event Stream API gateway
- **Feed ID** – The numerical count of the data feed (count starts a ‘0’)
- **App ID** – The App ID assigned in the TA Input configuration

The TA will examine the API response to determine the number of URL feeds and attempt to create and maintain an independent connect to each one. As events are processed from the URL feed(s) the TA will include the associated Feed_id (single URL feeds will always be ‘0’) and if there were multiple feeds presented in the ‘ta_data’ section:

```
ta_data: { [-]  
  Cloud_environment: us_commercial  
  Feed_id: 0  
  Initial_start: historic  
  Input: TS_Demo  
  Multiple_feeds: False  
  TA_version: 2.9.0  
}
```


Each event within a URL feed contain a unique numerical value called an 'offset' value. This value is used as a unique identifier for event within that URL feed. It is visible in the 'metadata' section of the Splunk event:

```
metadata: { [-]
  customerIDString: REDACTED
  eventCreationTime: 1590723450557
  eventType: IncidentSummaryEvent
  offset: 17676670
  version: 1.0
}
```

In the event that the network connection is disrupted the TA will leverage this information as the marker to determine the last event processed. Since the TA is able to support multiple inputs it uses the name of the Input as the unique identifier and then relates the data feed URL and offset values with it.

Input_name{datafeedURL:offset}

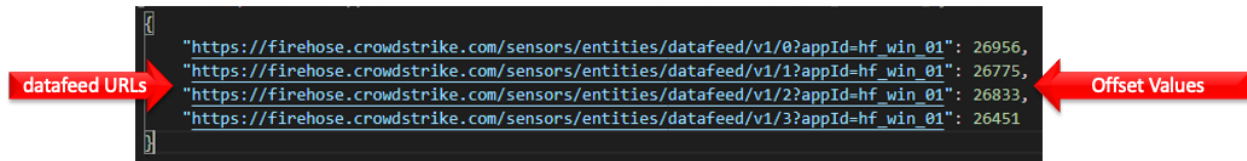
This information is then stored by the TA in:

- Splunk KV Store: The first location is within the Splunk KV (key:value) store. This is an internal Splunk location that the TA will call via API to both read and write data. (for more information please reference: <https://docs.splunk.com/Documentation/Splunk/8.0.4/Admin/AboutKVstore>)

Understanding Multiple Data Feeds and Applds

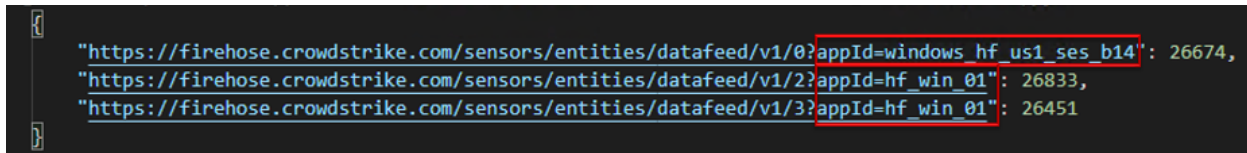
(The Offset data is only stored in the Splunk KV Store as of version 3.x , the images in this section are from an external file and not a visual representation of the Splunk KV store)

If the Falcon instance is leveraging multiple datafeed URLs to deliver data to the client, each of these datafeeds will be represented by a specific URL and each will have a dedicated offset value associated with them. The following depicts an Event Streams API that has 4 datafeed URLs and their respective offset values.



```
[{"url": "https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/0?appId=hf_win_01": 26956, "offset": 26956}, {"url": "https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/1?appId=hf_win_01": 26775, "offset": 26775}, {"url": "https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/2?appId=hf_win_01": 26833, "offset": 26833}, {"url": "https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/3?appId=hf_win_01": 26451, "offset": 26451}
```

The API gateway is able to accept, maintain and track multiple connections to the Event Streams API by leveraging an 'appId' value. The 'appId' serves as a unique name for that Falcon instance and the Event Streams gateway. This name is created by the customer but needs to be unique within the Falcon instance. If a client attempts to connect to the Event Streams API with the same 'appId' of another connection that's already established the connection will be refused. The 'appId' is, in fact, part of the datafeed URL itself as shown below:



```
[{"url": "https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/0?appId=windows hf_us1_ses_b14": 26674, "offset": 26674}, {"url": "https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/2?appId=hf_win_01": 26833, "offset": 26833}, {"url": "https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/3?appId=hf_win_01": 26451, "offset": 26451}
```

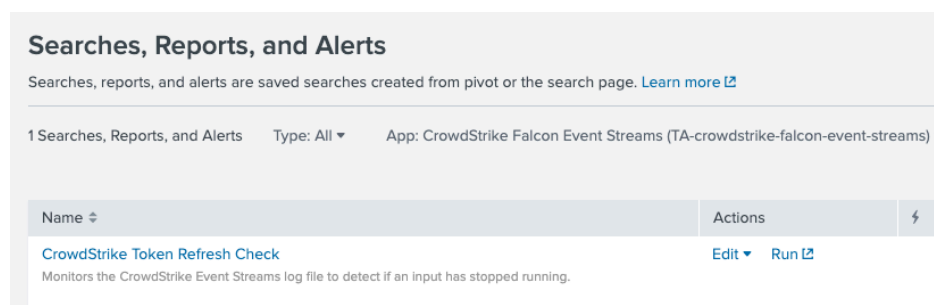
Token Refresh Check Alert & Restart Input Alert Action

Note: Due to architectural changes to SSL usage and requirements introduced in Splunk Enterprise v9.x and the associated Splunk Cloud version(s) this alert action may not function as designed.

In order to help mitigate issues with API connectivity and data transfers, v2.0.9 introduced a custom search and alert action that can be configured to identify if the connection can become blocked or unstable and data is no longer able to be collected. A few examples of situation that can cause this to occur are:

- Network congestion
- Network devices that block persistent connections (and reconnect attempts) after a prolonged period of time
- Malformed connection communications such as a data connection within the API being closed but the connection to the API gateway remains active
- Internal Splunk errors
- Input accidentally disabled

A properly functioning input should attempt to refresh its OAuth2 token every 20 minutes. The 'CrowdStrike Token Refresh Check' alert is designed to look for OAuth2 issue and refresh logs within the TA to help determine if the input is still processing data correctly. This is accomplished by leveraging the **cs_es_tc_input(1)** search macro to look for OAuth2 issue/refresh events within a 60 minute time window and ensure that there are at least 2 events. The search macro takes an input name by default so an alert is considered specific to that input. In an environment with multiple inputs, it's recommended to configure alerts for all active inputs. If there are not at least 2 events the alert should fire and take the alert actions that have been properly configured.

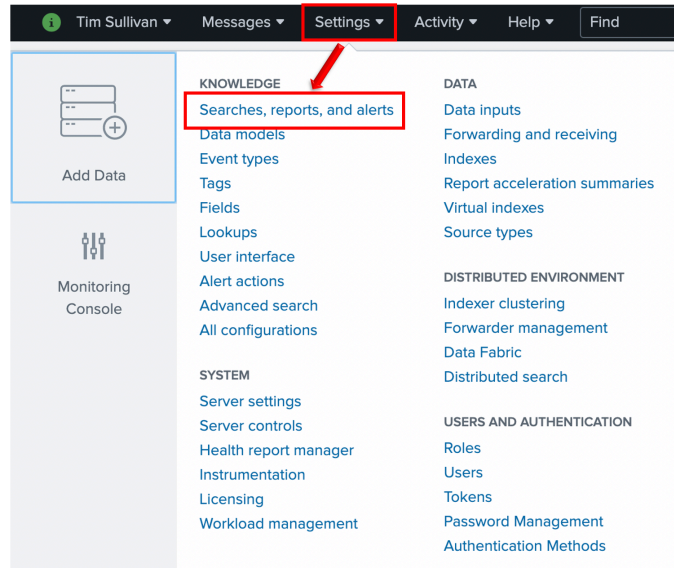


- **CrowdStrike Token Refresh Check:** This alert is designed to detect if there have been more than 2 token refresh/issue logs within the past 60 minutes (default settings)


The combination of this alert and the alert action is designed to cycle an input through the 'disable' and 'enable' actions via the Splunk REST API. In order for this to be successful the account that is performing this action should have the proper level of access to those REST endpoints. Typically, this is an account with 'admin' or 'system' level access.

Configuring the custom alert to restart an input


1. Under 'Settings' select 'Searches, reports, and alerts'




2. Ensure that the CrowdStrike Falcon Event Streams TA is selected as the 'App':



3. Locate the 'CrowdStrike Token Refresh Check' alert:



4. Under 'Actions' select 'Edit':



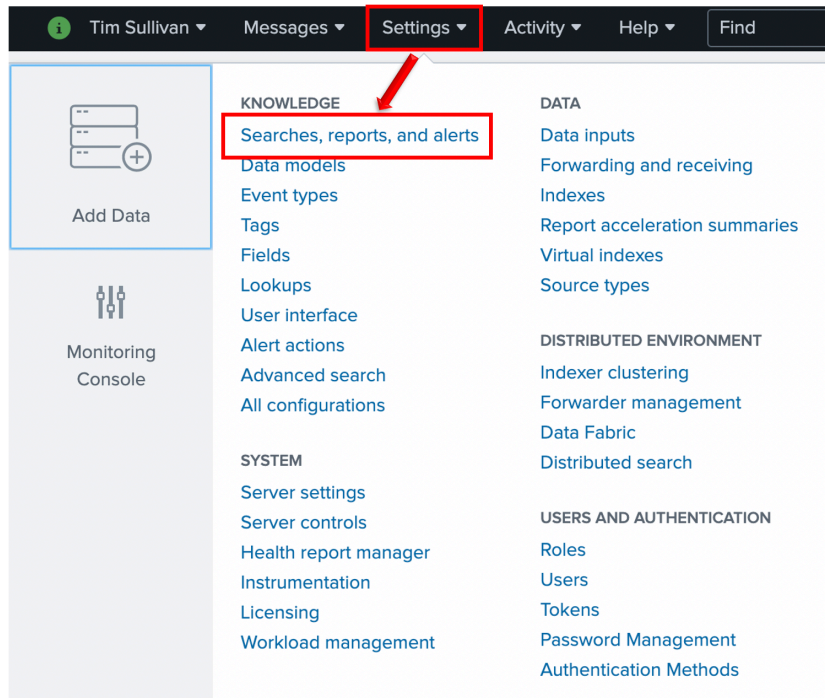
5. Configuring the Alert and associated actions:

A screenshot of the 'Edit Alert' configuration page in Splunk. The page is titled 'Edit Alert' and contains several sections: 'Settings', 'Trigger Conditions', and 'Trigger Actions'. Red circles with numbers 1 through 10 are overlaid on the page to indicate key configuration steps. Step 1 points to the 'Search' field containing the query: `'cs_es_tc_input(demo)' | stats count`. Step 2 points to the 'Alert type' section, which has 'Scheduled' selected. Step 3 points to the 'Run on Cron Schedule' dropdown. Step 4 points to the 'Time Range' dropdown, which is set to 'Last 60 minutes'. Step 5 points to the 'Cron Expression' field, which contains `* /60 * * * *`. Step 6 points to the 'Expires' field, which contains '999'. Step 7 points to the 'Trigger alert when' dropdown, which is set to 'Custom'. Step 8 points to the 'Trigger alert when' text field, which contains `search count < 2`. Step 9 points to the 'Trigger' dropdown, which is set to 'For each result'. Step 10 points to the 'Trigger Actions' section, which shows a list of actions: 'Add to Triggered Alerts' and 'CrowdStrike Event Streams - Restart Input'. The 'Add to Triggered Alerts' action has a 'Severity' dropdown set to 'Critical'. At the bottom of the page, there are 'Cancel' and 'Save' buttons.


1. **Search:** This is the search that will be used to look for the log information
2. **Alert Type:** This alert is designed to be a scheduled alert
3. **Schedule configuration** – the default configuration is a Cron Schedule
4. **Time Range:** The default time range is 60 minutes which should have at least 2 logs
5. **Cron Expression:** The default Cron expression sets the search to run every 60 minutes
6. **Expires:** The alert is set to expire in 999 days
7. **Trigger alert when:** Custom
8. **Number of results evaluation** – The results count value is less than 2
9. **Trigger:** The default configuration is 'Once' *
10. **When Triggered:** There are two default actions:
 - a. **'Add to Triggered Alerts'**
 - b. **'CrowdStrike Event Streams Restart Input'**

Enabling the custom alert to restart an input


1. Under 'Settings' select 'Searches, reports, and alerts'



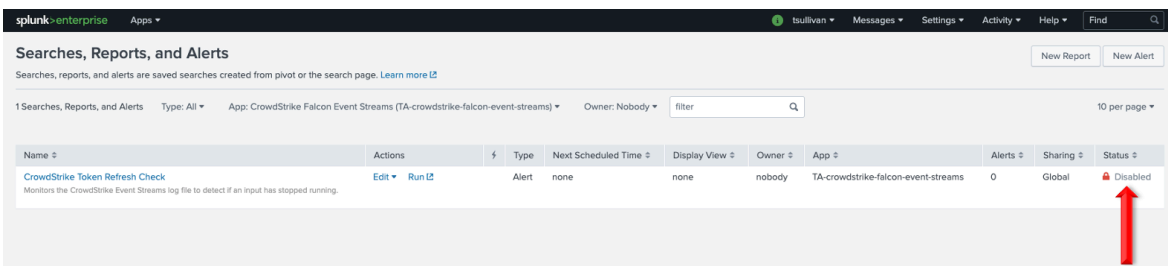
2. Ensure that the CrowdStrike Falcon Event Streams TA is selected as the 'App':



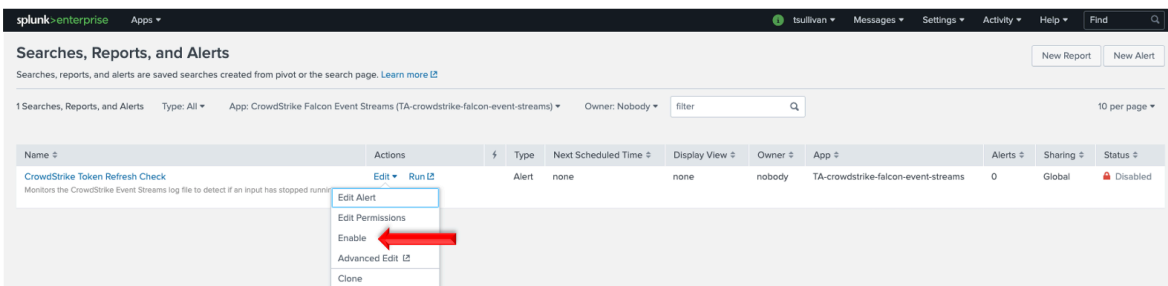
3. Locate the 'CrowdStrike Token Refresh Check' alert:



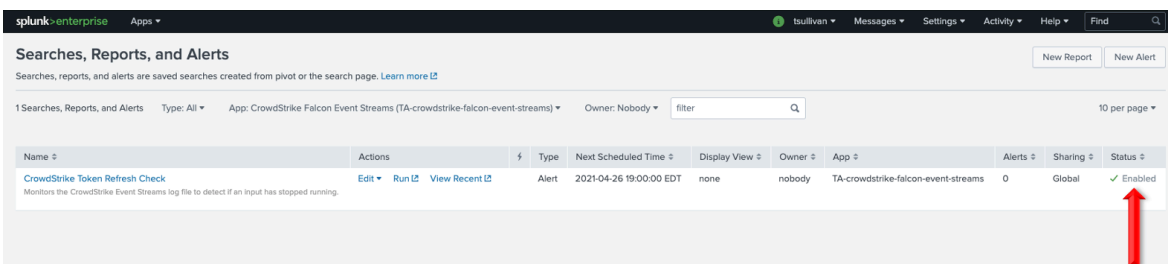
4. Verify that the alert is not currently enabled:



5. Under 'Actions' select 'Enable':



6. Ensure that the alert status has changed to 'Enabled':

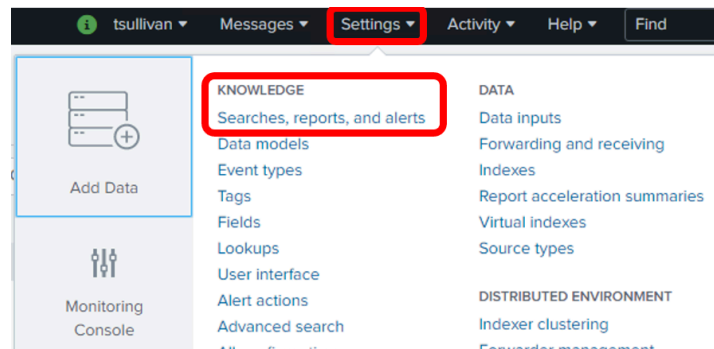


Scheduled Reports and Alerts


Locating the Reports and Alerts

The TA contains searches that can be run to collect specific data sets. These reports are particularly useful in troubleshooting and providing TA logs to CrowdStrike support.

To locate the TA searches, select 'Settings' and under 'Knowledge' select 'Searches, Reports, and alerts':



In the 'Searches, reports, and alerts' section ensure that the appropriate TA is selected under 'App' and that 'All' is selected under 'Owner':



This will display the searches that are provided by the add-on:

NOTE: THESE SEARCHES REQUIRE SEARCH MACRO(S) TO BE PROPERLY CONFIGURED

A screenshot of the 'Searches, Reports, and Alerts' page showing a table of searches. The table has columns: Name, Actions, Type, Next Scheduled Time, Display View, Owner, App, Alerts, Sharing, and Status. There are four rows of searches listed.

Name	Actions	Type	Next Scheduled Time	Display View	Owner	App	Alerts	Sharing	Status
CrowdStrike Event Stream Data Indexed vs Event Time Shows when the CrowdStrike Event Streams data was indexed, the event timestamp, the source type, the metadata.eventType and the count of events. Timespan needs to be large enough to include the event's time stamp.	Edit Run	Report	none	none	nobody	TA-crowdstrike-falcon-event-streams	0	Global	Enabled
CrowdStrike Event Streams TA Logs - 30 Days Collects the CrowdStrike Event Streams Technical Add-on logs for the past 30 days	Edit Run	Report	none	none	nobody	TA-crowdstrike-falcon-event-streams	0	Global	Enabled
CrowdStrike Event Streams: Indexing Shows the amount of CrowdStrike Event Streams Data that was indexed within a specific timeframe.	Edit Run	Report	none	none	nobody	TA-crowdstrike-falcon-event-streams	0	Global	Enabled
CrowdStrike Token Refresh Check Monitors the CrowdStrike Event Streams log file to detect if an input has stopped running.	Edit Run	Alert	none	none	nobody	TA-crowdstrike-falcon-event-streams	0	Global	Disabled

Understanding the Reports and Alerts

These reports require the index search macro be correctly configured

- **CrowdStrike Event Streams Indexing**: Shows the amount of CrowdStrike Event Streams Data that was indexed within a specific time frame.
- **CrowdStrike Event Stream Data Indexed vs Event Time**: Shows when the CrowdStrike Event Streams data was indexed, the event timestamp, the sourcetype, the metadata.eventType and the count of events. Timespan needs to be large enough to include the event's time stamp.

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

- **CrowdStrike Event Streams TA Logs – 30 Days**: Collects events from the internal index for sourcetype of the TA log files 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.

This alert is disabled by default:

- **CrowdStrike Token Refresh Check**: Monitors the CrowdStrike Event Streams log file to detect if an input has stopped running and attempts to disable and re-enable it*.

*Due to changes in Splunk's use of SSL certificates in the most recent release, input restart alert actions may not function as designed or as they had in previous versions.

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 data. If data inputs are configured for event types, it may also be prudent to index that data into dedicated indexes. One of the key factors that can influence this configuration if the contents of the data dictate that specific access controls be applied.

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 has no impact of the Event Streams connection as it is a persistent connection. As such changing the interval value has zero impact on the way the TA functions and should not be done.

Event Filtering

The Event Streams API data is processed an event at a time since that is how the data is received and in order to maintain an accurate offset value. This will typically result in the TA being able to ingest (on average) 1700 events per minute per input. In the overall number of events exceeds this number the data processing may be unable to process the incoming data in a timely manner. In these types of situations and where it makes numerical sense, it's recommended that event types of higher importance/concern/interest (such as detection events) be moved to a filtered input(s). For example: The number CSPMSearchStreamingEvent events is averaging 1500 events per minute and causes detection events ingestion to be delayed. The detection events could be collected using a new input that filtered for just those events. Another option would be that all events except the CSPMSearchStreamingEvent events could be collected in the existing input and a new dedicated input would be deployed for just the CSPMSearchStreamingEvent events. The appropriate solution depends on a customer's environment as well as the processes and procedures.

Troubleshooting

CrowdStrike provides support for the TA's code, the functionality of that code and authentication to the API endpoint(s). The following topics fall outside of that scope:

1. Network connectivity issues unrelated to authentication response from the CrowdStrike API endpoint
2. Tagging and CIM mapping (these are considered feature requests and will be evaluated by the integrations team)

Troubleshooting Overview

How TA gets data into Splunk

In the simplest terms, the TA acts as a 3rd party Python based data broker that's collecting information from one API, processing it and then sending it to another API. In this case the TA is establishing a persistent connection to the CrowdStrike Event Streams API to collect data as it's available. That data is collected by the python code, the event structure is created within the code and then that event is pushed to a Splunk API that will accept that event for processing and indexing. The TA has no 'knowledge' of the Splunk infrastructure outside what's been configured when an input is created.

Check the connection in the Falcon UI:

The connection to the Falcon system can be validated by using the API Audit log within the Falcon UI. The specific API Client ID, User IP and Company can be indicated if desired but the API action should be set to 'Stream Started'. Once this is submitted there should be audit details available for review. Under 'Activity Details' validate the AppID and the ClientID listed match the ones configured in the input configuration.

API Audit Trail
* Review audit logs for actions taken via our OAuth2-based APIs.

API Client ID: [] User IP: [] API action: Stream Started [X] Company: Tim Sullivan Demo [X] Time range: Last 14 days [] [Submit] [Hide Filters]

Time(UTC)	API action	Activity Details	User IP	City	Country	API Client ID	Company
2022-05-18 14:12:12	Stream Started	App ID: v3.1 Partition: #0 Offset: 2983372	Redacted	Ashburn	United States	Redacted	Tim Sullivan Demo

It's important to keep in mind that the Event Streams API is a persistent connection and as such if the TA is functioning properly the audit time range needs to be adjusted properly so that it encompasses when the input was initially enabled. For example, if the input was configured and enabled on a Friday and troubleshooting is taking place the following Monday, the time range for an API Audit search should be far back enough to include all of Friday's events.

Troubleshooting typical situations

- **The TA isn't collecting data or doesn't appear to be connecting to the API:**
 1. If the TA was previously connected the connection may have been interrupted, disable and re-enable the input to allow the connection to attempt to be re-established.
 2. Ensure that the API credential has been input correctly, cutting and pasting the secret can cause issues in some cases and may require that the secret to be manually inputted
 3. Ensure that the API credential being selected for the input is the appropriate credential for the selected CrowdStrike Cloud. The TA does not support redirection.
 4. If using a proxy, ensure that the proxy is not interfering with the network connection or the authentication process.
 5. Ensure that there is not an internal Splunk issue preventing the TA sending data into Splunk.
 6. Determine if there are any network configurations that are preventing a stable, persistent connection to the Event Streams API.
 7. Examine the API audit log in the Falcon UI to determine if there is any log of a connection being attempted.
 8. Ensure that the AppID that's being used is unique for the Falcon Instance. If a connection is already established with that AppID, all follow on connections attempting to use that AppID will not be established.
- **The TA appears to be working correctly but there's no data in the index:**
 1. Ensure that there's actually data available to be collected. In some cases there may not be data within the time frame being viewed.
 2. Validate that the index exists, the TA only validates the existence of the index when the input is initially configured.
 3. Validate that the index is enabled. If the index exists but is disabled Splunk will simply drop the data that the TA passes to it.
 4. Ensure that the search timeframe being used encompasses the time the event took place in Falcon and not the time it was expected to have been indexed in Splunk.
 5. Ensure that the AppID that's being used is unique for the Falcon Instance. If a connection is already established with that AppID, all follow on connections attempting to use that AppID will not be established.

- **The TA is ‘behind’ or the data ‘is coming in slowly’:**

1. The Splunk architecture that the TA runs in is such that it’s been determined that the average amount of events that can be processed by a single input in around 1,700 events per minute. If the event volume of the input exceeds this average there may be delays in processing.
2. If there are particular event types that maybe generating high volumes or would significantly reduce the overall volume, consider creating an input just for these event types and remove from being processed by another input(s).
3. The TA does not post data to the index but rather it posts the data to an API exposed on the heavy forwarder/IDM which then sends the data to the indexer to be indexed. This means that the time that the TA receives the event is not synonymous with Splunk’s ‘_indextime’ stamp. If the heavy forwarder/IDM or indexer are being heavily utilizes the indexing process may slow down as a result. The easiest way to determine when an event was received by the TA is by looking in the TA logs for the recording of the offset value of that input’s event into the Splunk KVstore (if there are multiple datafeed URLs ensure that you’re locating the correct one) and comparing it to the _indextime value.

```
> 5/18/22 2022-05-18 14:12:16,618 INFO pid=5900 tid=Thread-1 file=base_modinput.py:log_info:295 | CrowdStrike Event Streams TA 3.1.0 demo: OFFSET KV STORE: Offset r
10:12:16.618 AM ecoreded to KV Store: demo_feed_num_0 {'https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/0?appId=v3.1': 2903373}
host = ip-10-0-10-28.ec2.internal source = /opt/splunk/var/log/splunk/ta_crowdstrike_falcon_event_streams_crowdstrike_... sourcetype = tacrowdstrikefalconeventstreams.log
```

```
demo_feed_num_0 {'https://firehose.crowdstrike.com/sensors/entities/datafeed/v1/0?appId=v3.1': 2903373}
```

Support

This TA is designed to help facilitate the collection of data provided by the CrowdStrike API(s). CrowdStrike provides support for the TA code functionality as it was designed. Example 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
- Support requests without the appropriate data outlined below
- Splunk CIM field mapping or custom data modification requests

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 Event Streams TA 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 'crowdstrike_event_streams' 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 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 above as well as any modifications that have been made to the TA outside of the processes 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.

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.**

© 2022 CrowdStrike, Inc. All rights reserved.