



CROWDSTRIKE

CrowdStrike

Falcon Spotlight Vulnerability Data

Add-on for Splunk

Installation and Configuration Guide v3.2+

Table of Contents

INTRODUCTION	2
REQUIREMENTS	3
GETTING STARTED	4
API Endpoints, Filters and Timestamps	4
Spotlight Data Communication Flow	4
Data Volume Considerations	5
Understanding and Comparing UI and API Based Data	5
High Level Data Flow	7
Validating that Spotlight is Enabled	8
Identifying the CrowdStrike Cloud:	9
Generating and Scoping Credentials	10
Generating New API Credentials	10
Modifying Existing Credentials	10
Scoping the Credentials	11
Proxy Considerations	12
Splunk Architecture	12
CONFIGURING THE TA	14
TA Layout	14
Inputs Section	14
Configuration Section	15
Search Section	15
Configuring the TA to collect data	16
Configure Proxy Settings (optional)	16
Configure an Account	16
Creating an Input	18
Configure an Input	18
SEARCHES, REPORTS, AND ALERTS	22
SEARCH MACROS	23
RECOMMENDATIONS	24
Custom Indexes	24
Interval Settings	24
Optional Filters & Data	24
EVENT DATA STRUCTURE	26
TROUBLESHOOTING	28
Configuring the TA to collect log data	28
Change Logging Level	28
Review Log Data in Splunk	29
Examples of Troubleshooting Situations and Remediation Steps	30
Prior to Contacting CrowdStrike Support	33

Introduction

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

The CrowdStrike Falcon Spotlight Vulnerability Data Technical Add-on for Splunk allows CrowdStrike customers to retrieve CrowdStrike Spotlight Vulnerability data from CrowdStrike Falcon instance that have the Spotlight module enabled via API.

To get more information about this CrowdStrike Falcon Spotlight please refer to the documentation for the Spotlight module located in the CrowdStrike Falcon UI:

<https://falcon.crowdstrike.com/documentation/43/falcon-spotlight>

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.

This Technical Add-On does not currently support Falcon Flight Control Architectures. API access is direct to the Falcon Instance.

Requirements

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

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

If you do not have a current CrowdStrike Spotlight subscription:

1. Contact your CrowdStrike sales team to acquire one
2. Navigate to the CrowdStrike store in your falcon instance and request a trial: [Click Here to See the CrowdStrike Spotlight App in the CrowdStrike Store](#)

This Technical Add-On does not currently support Falcon Flight Control Architectures. API access needs to be direct to the Falcon Instance.

Getting Started

API Endpoints, Filters and Timestamps

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	Required Filter Fields(s)	Optional Filter Field(s)
/oauth2/token		
/spotlight/combined/vulnerabilities/v1	updated_timestamp	facets status host_info.platform_name cve.severity cve.exprt_rating

The event timestamp used by the TA is the **'updated_timestamp'** value. This is because when a vulnerability event is created the **'created_timestamp'** and the **'updated_timestamp'** are identical. As the vulnerability event is updated the **'updated_timestamp'** value will change to reflect when the change took place. This value will be what Splunk uses for the **_time** value and can be found in the event data under **'falcon_spotlight.updated_timestamp'**.

Spotlight Data Communication Flow

The CrowdStrike Falcon Spotlight Vulnerability Technical Add-on for Splunk leverages the **'combined'** Spotlight API endpoint to collect vulnerability data. The TA communication process is as follows:

1. The TA will authenticate to the CrowdStrike API gateway for the configured CrowdStrike Cloud environment to collect an OAuth2 token
2. The OAuth2 token will then be used by the TA to connect and collect Spotlight vulnerability data from the CrowdStrike Spotlight API combined endpoint:
/spotlight/combined/vulnerabilities/v1 *
3. Up to 4000 vulnerabilities will be called per API call **
4. Once all relevant data has been retrieved the TA will process the data
5. The TA will post the event data to Splunk to be indexed, the TA cannot and does not send the data to the indexer

* The credential used to acquire the OAuth2 token must be scoped correctly to be able to connect to the API

** multiple API calls maybe required to collect all available data

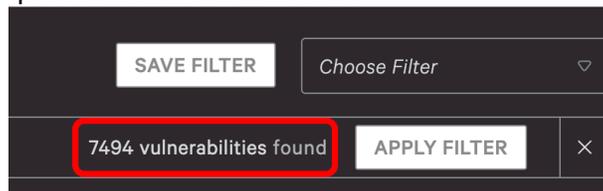
Data Volume Considerations

Spotlight Vulnerability data can be a large amount of data depending on the size of the environment and the time range of the data being collected. Some key considerations to take into account when collecting this data:

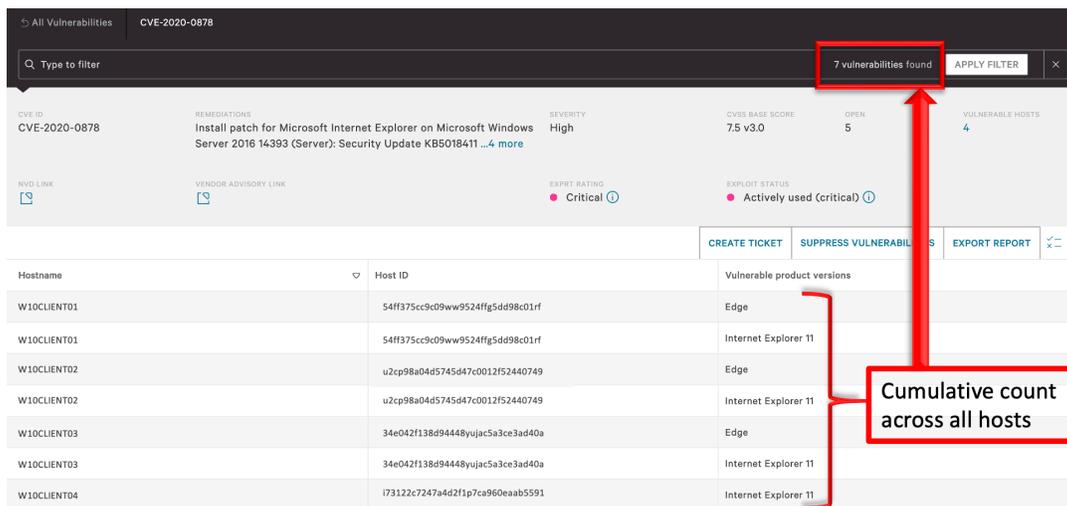
1. The amount of data that will be ingested
2. The available resources to collect and process the data
3. The time range or frequency of the data collection

Understanding and Comparing UI and API Based Data

In Spotlight, a vulnerability is defined as the combination of Host, Product and CVE IDs. The total cumulative count is presented in the filter bar:



An example of how this calculation is made can be seen by the examining details of a specific CVE:

A screenshot of the Spotlight interface for CVE-2020-0878. The top filter bar shows '7 vulnerabilities found' (highlighted with a red box). Below, a table lists vulnerable hosts and their product versions. A red bracket groups the product versions for W10Client01, W10Client02, and W10Client03, with a callout box stating 'Cumulative count across all hosts'.

Hostname	Host ID	Vulnerable product versions
W10CLIENT01	54ff375cc9c09ww9524ffg5dd98c01rf	Edge
W10CLIENT01	54ff375cc9c09ww9524ffg5dd98c01rf	Internet Explorer 11
W10CLIENT02	u2cp98a04d5745d47c0012f52440749	Edge
W10CLIENT02	u2cp98a04d5745d47c0012f52440749	Internet Explorer 11
W10CLIENT03	34e042f138d94448yujac5a3ce3ad40a	Edge
W10CLIENT03	34e042f138d94448yujac5a3ce3ad40a	Internet Explorer 11
W10CLIENT04	i73122c7247a42f1p7ca960eab5591	Internet Explorer 11

W10Client01, W10Client02 and W10Client03 each have 2 vulnerable product versions, while W10Client04 has one – the total vulnerability product version count is 7 (2+2+2+1), which matches the count for the vulnerabilities found in the filter bar.

NOTE: That while the Hostname is shown above, the Spotlight API guide states: “*when calculating the number of hosts vulnerable to a CVE, Spotlight searches for unique host IDs*”. In the example above, while ‘Hostname’ may seem like the field being leveraged for uniqueness it is actually the ‘Host ID’ field. (Host ID and Agent ID (AID) are synonymous)

The method to identify a unique event in the Spotlight API data is to use the 'id' field, which is defined by the API guide as the "Unique system-assigned ID of the vulnerability". In the Spotlight Vulnerability Data TA for Splunk this field is the 'Falcon_Spotlight.id' field.

The screenshot shows a Splunk search interface. On the left, a search bar contains 'falcon_spotlight.id'. Below it, there are options for 'Reports', 'Top values', 'Top values by time', and 'Events with this field'. A 'Selected' dropdown is set to 'Yes'. On the right, a JSON response is displayed. The 'id' field is highlighted with a red box. The JSON structure is as follows:

```

{
  "falcon_spotlight": {
    "aid": "redacted",
    "apps": [
    ],
    "cid": "redacted",
    "closed_timestamp": "2022-10-24T16:16:00Z",
    "created_timestamp": "2022-10-24T15:12:01Z",
    "cve": [
    ],
    "host_info": {
    },
    "id": "59a632f3e211401ca8d7f7bf167bd282_00d2b5d8c7283f459185fe8b44292c23",
    "remediation": {
    },
    "status": "closed",
    "suppression_info": {
    },
    "updated_timestamp": "2022-10-24T16:16:00Z"
  }
}

```

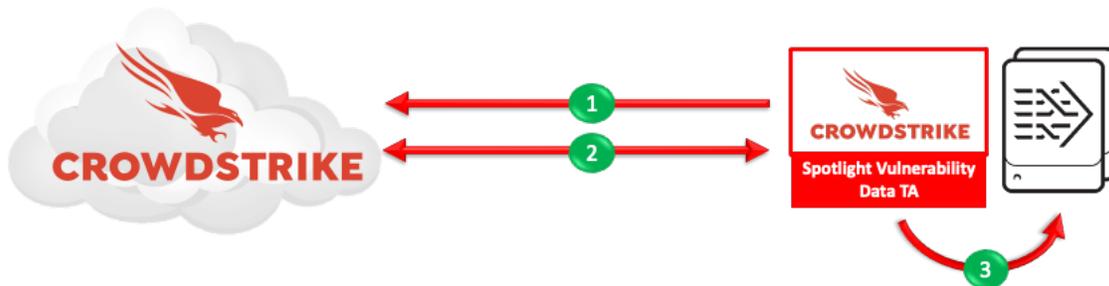
The data collected from the Spotlight Vulnerability API does differ in some regards to the data presented in the Falcon UI. For example:

- **The EXPIRED status:** The Spotlight API guide states: "If a host is deleted or inactive for 45 days, the status of vulnerabilities on that host changes to expired. Expired vulnerabilities are removed from Spotlight after 3 days. Expired vulnerabilities are only visible in API responses and are not included in reports or the Falcon console."
- **The host_last_seen_timestamp:** The Spotlight API guide states: "UTC timestamp of device's most recent connection to Falcon in ISO 8601 format. A host_last_seen_timestamp value is only provided for hosts that have been offline for 3 or more days. The timestamp value will not change until the host comes back online, at which point the value resets to null. If the host has been online during the last 3 days, the host_last_seen_timestamp is null." This should be taken into account if the data comparison is using the 'Last Seen Within' filter in the UI to compare to the API data.

It is highly recommended that customers review the published API guide for the Spotlight APIs to ensure that they are aware of the most up to date information: <https://falcon.crowdstrike.com/documentation/98/spotlight-apis>

High Level Data Flow

The CrowdStrike Falcon Spotlight Vulnerability Data TA leverages CrowdStrike API calls to collect data:



1. The TA acquires an OAuth2 token from the CrowdStrike API gateway
2. The TA uses the OAuth2 token to query and collect CrowdStrike Falcon Spotlight Vulnerability data
3. The TA posts the data to the internal Splunk API

This process will continue until all the vulnerability data has been collected as long as the input's interval setting allows. In the event that the interval setting is too short the running process will be interrupted and a new process will be started.

***Note – The API Client must be scoped for access to the Spotlight API to be granted access**

Validating that Spotlight is Enabled

The CrowdStrike Falcon Spotlight Vulnerability Data TA requires a valid Spotlight subscription and that Spotlight has been enabled on the CrowdStrike instance.



Spotlight

Dashboard

Vulnerabilities

Installed Patches

Custom Filters

Reports

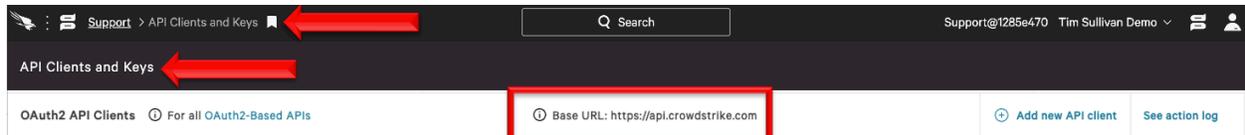
1. Access the CrowdStrike Falcon user interface (UI)
2. Ensure that 'Spotlight' is listed in the dropdown menu

If you have a Spotlight subscription but are not able to access it or create an API credential, please submit a support ticket through the support portal:

<https://supportportal.crowdstrike.com/>

Identifying the CrowdStrike Cloud:

The Spotlight Vulnerability Data TA requires the CrowdStrike Cloud environment be identified when being configured. The cloud environment can be located in the CrowdStrike Falcon UI under the 'Support' > 'API Clients and Keys'

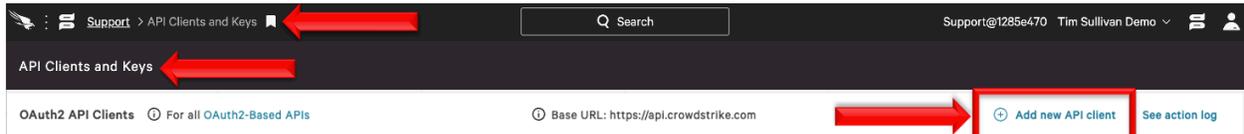


- US Commercial 1: <https://api.crowdstrike.com>
- US Commercial 2: <https://api.us-2.crowdstrike.com>
- US GovCloud: <https://api.laggar.gcw.crowdstrike.com>
- EU Cloud: <https://api.eu-1.crowdstrike.com>

Generating and Scoping Credentials

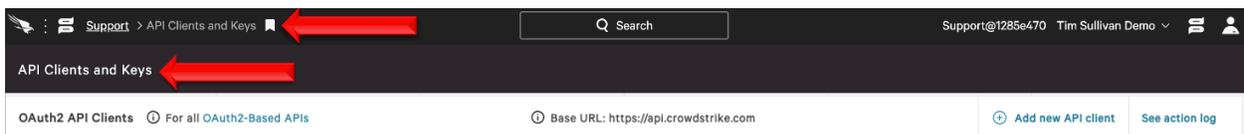
The Spotlight Vulnerability Data TA requires API credentials that are located in the CrowdStrike Falcon UI in order to access the APIs. These can be existing API credentials with the Spotlight scope or can be newly generated credentials.

Generating New API Credentials

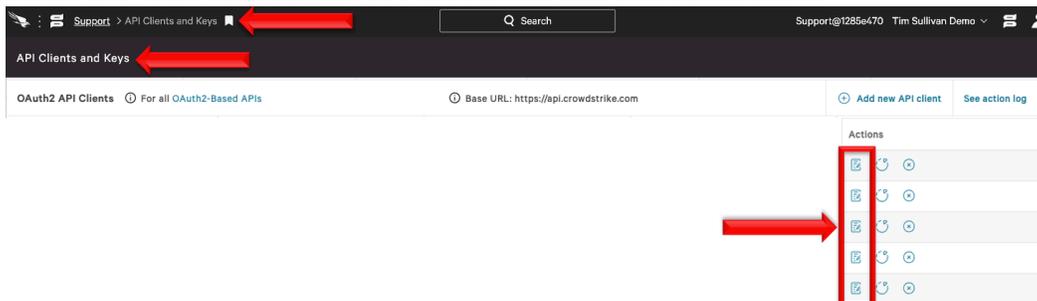


1. Access the CrowdStrike Falcon user interface (UI) with an account that is able to view/create the API clients and keys page
2. Navigate to 'Support' > 'API Client and Keys' page
3. On the same line as 'OAuth2 API Clients' select 'Add new API client'
4. Give the client a name and assign the appropriate scope(s)
5. Select Add and save the credentials *NOTE: This is the only time the Secret appears, failure to capture it at this time will require that it be regenerated

Modifying Existing Credentials



1. Access the CrowdStrike Falcon user interface (UI) with an account that is able to view/create the API clients and keys page
2. Navigate to 'Support' > 'API Client and Keys' page
3. Select the 'edit' icon for the API credential that will be modified



Scoping the Credentials

The API credential, either new or existing, must be assigned the correct scope to be able to access the data.

Sensor Download	<input type="checkbox"/>	—
Sensor update policies	<input type="checkbox"/>	<input type="checkbox"/>
Sensor Visibility Exclusions	<input type="checkbox"/>	<input type="checkbox"/>
Spotlight vulnerabilities	<input checked="" type="checkbox"/>	—
Event streams	<input type="checkbox"/>	—
User management	<input type="checkbox"/>	<input type="checkbox"/>
Zero Trust Assessment	<input type="checkbox"/>	—

CANCEL ADD

1. Under the 'API SCOPES' selection located 'Spotlight vulnerabilities'
2. Check the 'READ' scope for Spotlight vulnerabilities (Note: there is no 'WRITE' scope for this API).
3. Select 'ADD' to save this assignment

Proxy Considerations

The CrowdStrike Falcon Spotlight Vulnerability Data Add-On communicates with the CrowdStrike's APIs and any proxy systems in the environment should be configured to allow this communication.

Since the CrowdStrike Falcon Spotlight Vulnerability Data Add-On has been transitioned to the CrowdStrike FalconPy SDK, the proxy configurations must now follow proper proxy configurations. **This means that 'HTTPS' must be selected from the 'Proxy Type' drop down.** This is because all calls to the CrowdStrike API involve a URL beginning with 'https' and when the TA looks for the appropriate proxy entry that is the protocol it will look for.

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.

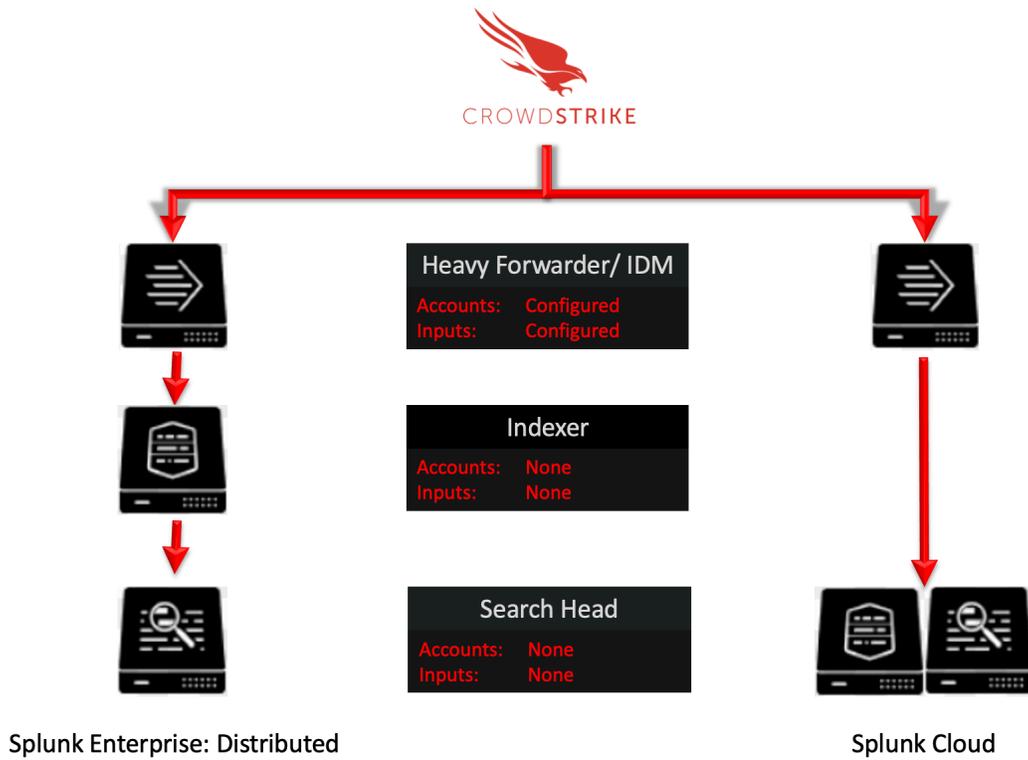
Splunk Indexer(s): The TA can be installed to provide field mapping and search macro support. The TA should be deployed without any accounts or inputs configured and any search macros should be properly configured for use. If a custom index is going to be used, then it should be created here.

Splunk Heavy Forwarder(s) & Information Data Managers (IDMs): The TA is required to be installed here as this is where the data will be collected. The appropriate accounts and inputs should be properly configured for data collection. Ensure that if a custom 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, IDMs or Splunk Cloud instances that support data ingestion.

The following diagram shows the flow of data from CrowdStrike and CrowdStrike Falcon Spotlight Vulnerability TA configuration within a distributed Splunk Enterprise and Splunk Cloud environment:



Configuring the TA

TA Layout

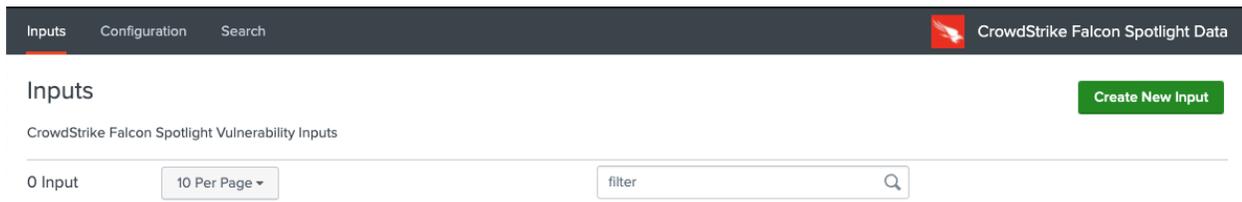
The TA contains 3 sections.



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

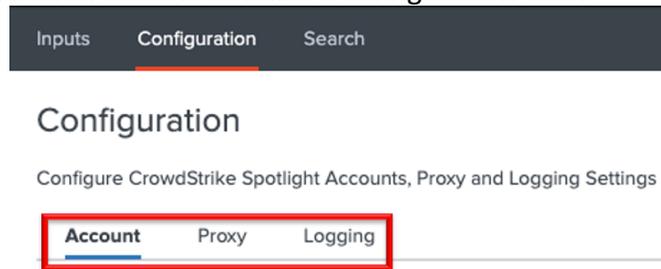
Inputs Section

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



Configuration Section

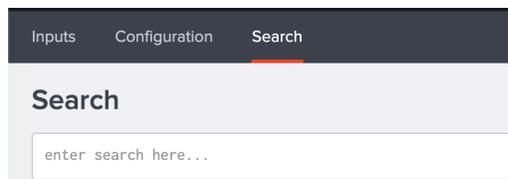
The Configuration section contains 3 configuration tabs:



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

Search Section

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



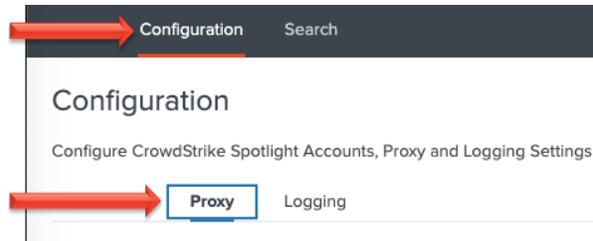
Configuring the TA to collect data

NOTE

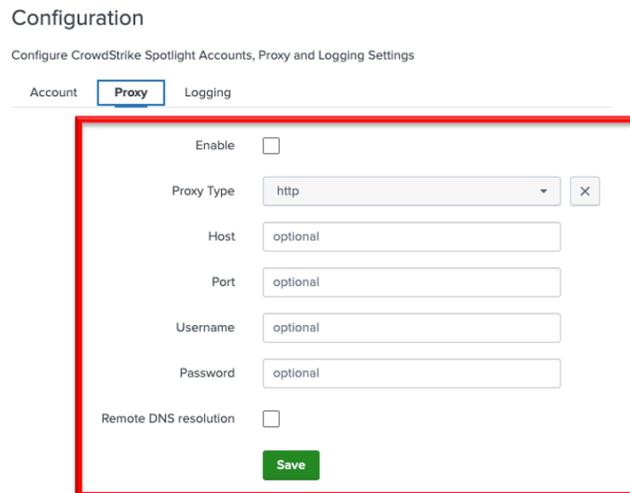
These actions should only be performed on Splunk systems ingesting data

Configure Proxy Settings (optional)

1. Proxy settings are configured under the Configuration section, Proxy tab. Proxies can cause authentication issue if not configured correctly, ensure that the proxy does not interfere with communication between the TA and the CrowdStrike APIs



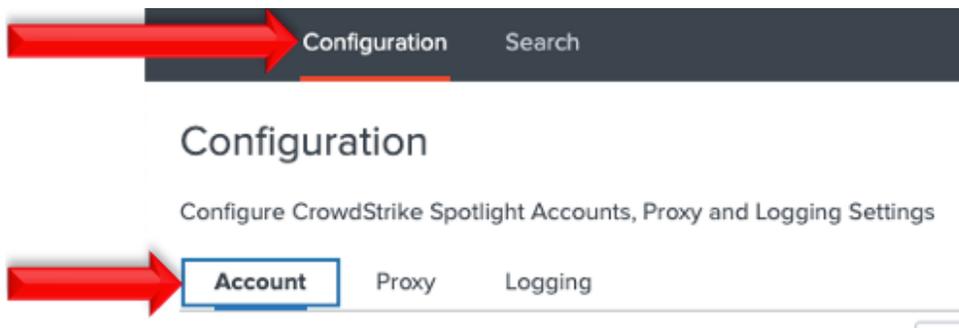
2. Configure the following fields as appropriate:

A screenshot of the Proxy configuration form in Splunk. The form is titled 'Configuration' with a subtitle 'Configure CrowdStrike Spotlight Accounts, Proxy and Logging Settings'. There are three tabs: 'Account', 'Proxy', and 'Logging'. The 'Proxy' tab is selected. The form contains the following fields: 'Enable' (checkbox), 'Proxy Type' (dropdown menu with 'http' selected), 'Host' (text input field with 'optional' placeholder), 'Port' (text input field with 'optional' placeholder), 'Username' (text input field with 'optional' placeholder), 'Password' (text input field with 'optional' placeholder), and 'Remote DNS resolution' (checkbox). A green 'Save' button is located at the bottom of the form.

- **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 save the configuration

Configure an Account

1. An account is configured using an API client credential from the CrowdStrike Falcon UI.
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:

Add Account

Account name
Enter a unique name for this account.

ClientID
Enter the ClientID here
Enter the ClientID for this account.

Secret
Enter the Secret for this account.

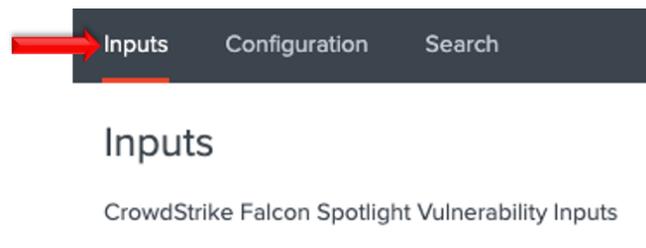
Cancel Add

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

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

Creating an Input

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



3. In the top right corner select the 'Create New Input' dropdown to display the available input types.



Configure an Input

The Spotlight Vulnerability Data TA can be configured with multiple inputs. These can be for the same or different CrowdStrike Falcon environments.

1. Select the 'Create New Input'



CrowdStrike Falcon Spotlight Data



Create New Input

2. Configure the appropriate fields:

Add CrowdStrike Falcon Spotlight Vulnerability
✕

[CrowdStrike Spotlight TA Guide](#)

[CrowdStrike FalconPy SDK Wiki](#)

Name

Enter a unique name for the data input

Interval

Time interval of input in seconds - must be greater than 3600 (1 hour)

Index

API Credential Select the Appropriate Spotlight API Creden... ✕

This is an OAuth2 based API credential with a 'Spotlight vulnerabilities read' scope

Select Cloud Environment Select the appropriate CrowdStrike cloud ✕

Select the appropriate cloud environment for the Falcon Instance

Optional Filter: Start Date

Only collect data that was updated on or after this date(UTC) for the initial collection

Optional Data: Facets

Include the additional contextual data specific to the facet type or types

Optional Filter: Status All ✕

Only collect data that currently has the specific status or statuses

Optional Filter: Platform Name All ✕

Collect data with a specific platform name value

Optional Filter: CVE Severity All ✕

Collect data with specific CVE severity value(s)

Optional Filter: ExPRT Rating All ✕

Collect data with specific CrowdStrike Exprt Rating value(s)

Cancel
Add

Required Settings

Optional & Additional Settings

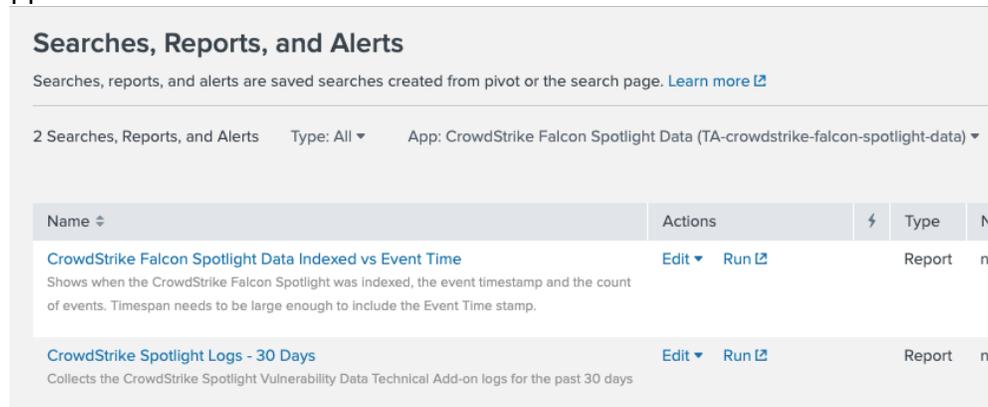
Required Settings	
Name	A name unique to the Splunk Environment
Interval	How often the specific input will run, expressed in seconds
Index	The Splunk Index that the data will be stored in
API Credentials	The appropriate account from the configuration tab
Cloud Environment	The CrowdStrike Cloud the instance resides in

Optional Filters & Data	
Start Date	Filters the data collected to be only those with the created or updated after this date
Facets	Provides the option to include host, CVE, remediation data and evaluation logic
Status	Filters the data collected to be only those with the selected status(es)
Platform Name	Filters the data collected to be only those with the selected platform name(s)
CVE Severity	Filters the data collected to be only those with the selected CVE severity(s)
ExpRT Rating	Filters the data collected to be only those with the selected ExpRT Rating(s)

- **Facets** is a multiselect field that can be configured with multiple options, by default none are selected.
- **Status, Platform Name, CVE Severity and ExpRT Rating** are all multiselect fields that can be configured with multiple options. By default, these fields are set to the 'All' selection, which must be removed to enter individual selections. The input will not save if 'All' is not the only selection present in any of these filters.

Searches, Reports, and Alerts

The Spotlight Vulnerability Data TA contains saved reports that can be found under the 'Searches, Reports, and Alerts' section. Ensure that 'CrowdStrike Falcons Spotlight Data' is the selected app:



The screenshot shows a web interface titled "Searches, Reports, and Alerts". Below the title, there is a subtitle: "Searches, reports, and alerts are saved searches created from pivot or the search page. [Learn more](#)".

At the top of the list, it says "2 Searches, Reports, and Alerts". To the right, there are filters: "Type: All" and "App: CrowdStrike Falcon Spotlight Data (TA-crowdstrike-falcon-spotlight-data)".

Name ↕	Actions	⚡	Type	N
CrowdStrike Falcon Spotlight Data Indexed vs Event Time Shows when the CrowdStrike Falcon Spotlight was indexed, the event timestamp and the count of events. Timespan needs to be large enough to include the Event Time stamp.	Edit Run		Report	n
CrowdStrike Spotlight Logs - 30 Days Collects the CrowdStrike Spotlight Vulnerability Data Technical Add-on logs for the past 30 days	Edit Run		Report	n

- **CrowdStrike Falcon Spotlight Data Indexed vs Event Time:** This report will show when the CrowdStrike Falcon Spotlight data was indexed by Splunk, the event timestamp and the number of events for that index time. This report helps show when the Spotlight data was actually indexed into Splunk
- **CrowdStrike Spotlight Logs – 30 days:** This report shows the TA's logs for the past 30 days. This information show be provided, at a minimum, when requesting support from CrowdStrike.

Search Macros

The Spotlight Vulnerability Data TA contains configurable search macros:

Search macros

Settings > Advanced search > Search macros

Successfully updated "cs_spotlight_get_index" in TA-crowdstrike-falcon-spotlight-data.

Showing 1-3 of 3 items

App: CrowdStrike Falcon S... Owner: Any Created in the App filter

Name	Definition	Arguments	Owner
CrowdStrike_Spotlight_AppName(1)	`cs_spotlight_get_index` falcon_spotlight.apps[].product_name_version=*\$app_name\$*	app_name	tsullivan
CrowdStrike_Spotlight_CVE(1)	`cs_spotlight_get_index` falcon_spotlight.cve.id=\$cve\$	cve	tsullivan
cs_spotlight_get_index	index=*		tsullivan

- **cs_spotlight_get_index**: This is configured to point to the index that is storing the CrowdStrike Spotlight Vulnerability data. The default configuration is to '*'. **This search macro must be configured for the other search macros and reports to work**
- **CrowdStrike_Spotlight_AppName(1)**: This search macro allows for an app name to be entered in the variable position (replacing the '1') and will search the 'product_name_version' for that name.
- **CrowdStrike_Spotlight_CVE(1)**: This search macro allows for a CVE number to be entered in the variable position (replacing the '1') and will search the 'cve.id' field for that CVE number.

Ensure the following:

1. Search Macros must be enclosed by 'back ticks', not single quotes. This key is located above the 'Tab' key, to the left of the number 1 on most US style keyboards.
2. Ensure that the account leveraging the macro has the correct permissions to use the macro or adjust the permission of the macro accordingly.
3. Ensure that the account leveraging the macro has the correct permissions to access the Spotlight Vulnerability data.
4. Ensure that the index(es) have been designated correctly.

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 Falcon Spotlight Vulnerability TA data.

Some examples of benefits that leveraging custom indexes can provides:

- Allows multiple teams to reference the data without exposing other data sets that may be more sensitive.
- Allows data collection types to be assigned to different Heavy Forwarders/IDM for access and resource allocation considerations.
- Improves searching response times and reduces resources needed.

Interval Settings

The amount of data that Spotlight can produce depend on the size of the environment and the associated update hygiene. Even in an environment with hosts that are relatively up to date the amount of data that's collected can be considerable, especially on the initial data collection. The interval setting should be configured to provide for a large amount of time to ensure the proper data collection and to this end it is recommended that customers begin with an interval equivalent to 6-12 hours. The TA log files can be examined to determine the average time frame that data collection need and the interval can be adjusted as needed. However, it is advisable to ensure that any adjustment allows for large spikes in the data and any resource constraints that maybe encountered during collection.

Optional Filters & Data

Avoid Duplicate Data - Facets: The CrowdStrike Falcon Spotlight Vulnerability TA allows customer to add additional information by selecting different facets in the input configuration. These can help enhance the information that's already present in a standard event. It's recommended that all facets of interest be enabled on a single input to avoid duplicating the baseline event information within Splunk.

Avoid Duplicate Data - Filters: The CrowdStrike Falcon Spotlight Vulnerability TA allows customer to select several different filters when configuring inputs. This can help reduce the amount of data a single input will have to process by allowing multiple inputs to collect the data at the same time. However, it is important to ensure that the filters are configured so as to not collect the same event and cause duplicative data in the same index.

Event Data Structure

The CrowdStrike Falcon Spotlight Vulnerability TA events contain three main data groupings:

```

{ [-]
  falcon_spotlight: { [+]
  }
  meta: { [+]
  }
  ta_data: { [+]
  }
}

```

- **falcon_spotlight:** This section contains the event data as it was returned by the API.
- **meta:** This section contains information from the API return.
- **ta_data:** This section is constructed by the TA and contains information about the specific input’s configuration. It was specifically designed to assist customer and CrowdStrike support with troubleshooting potential issues.

META:

This data section contains information that was collected from the API’s return response and will look similar to this:

```

meta: { [-]
  pagination: { [-]
    after: WzE2NjY2NDM4MjEwMDAsIjJlMmUzZjVlYjMyNTRhODZiMmMxYzI4MTA1ZDFjOGZhXzAyMjU1M2U5NjNmNDMzODk5NjBkZGUzZTl1ZjI4MTgxI10=
    limit: 4000
    total: 525538
  }
  powered_by: spapi
  query_time: 0.294382575
  trace_id: 1f68601b-7861-4e48-926c-81698553606d
}

```

Field Name	Field Description
pagination.after	A ‘next page’ value that the TA uses to continue collecting data when the total amount of events exceeds the limit of a single API call
limit	The value that the limit filter was set to for the API call, the TA sets this to 4,000 to avoid potential timeouts (the default maximum for the API is 5,000)
total	The total number of events that match the API query
powered_by	n/a
query_time	The amount of time it took to query the API
trace_id	A unique ID that allows CrowdStrike support to search for information about the specific API query

TA_DATA:

This data section contains data that was collected and created about the TA and the specific inputs configuration and will look similar to this:

```
ta_data: { [-]
  Cloud_environment: us_commercial
  Collection_hash: 8982200704605723028600136942955009540005465667266353777616985247673570217298
  Facets: [ [+]
]
  Filters: { [-]
    CVE_ExPRT_rating: [ [+]
]
    CVE_Severity: [ [+]
]
    Platform_name: all
    Status: [ [+]
]
}
Input: Example_All
Start_date: 2022-09-15T18:43:14
TA_version: 3.2.0
```

Field Name	Field Description
Cloud Environment	The CrowdStrike Cloud Environment configured in the input configuration
Collection_hash	A 256 character dynamically calculated hash for a particular input pull
Facets	A list of facets that were configured for this specific input at the time of the collection
Filters	A list of the optional filters and their settings at the time of the collection
Input	The name of the input in Splunk
Start_date	The start date that's configured for the input, which is only used if configured when the input is initially enabled
TA_version	The version of the TA that collected this data

Troubleshooting

CrowdStrike only provides support for:

- TA code-based functionality errors
- CrowdStrike API based access errors

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

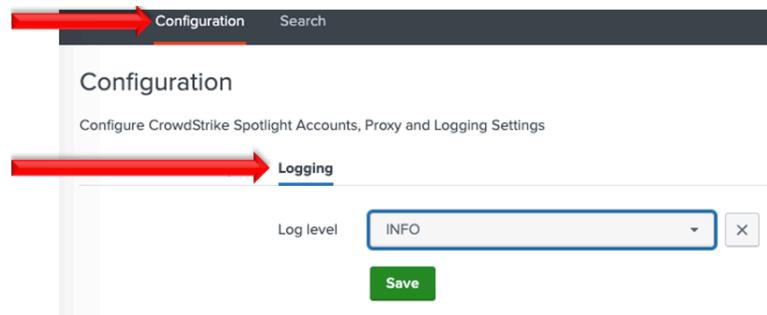
- Proxy based issues
- Firewall based issues
- Network connectivity issues
- Authentication issues (based on misconfigured credentials)
- Splunk CIM field mapping

Configuring the TA to collect log data

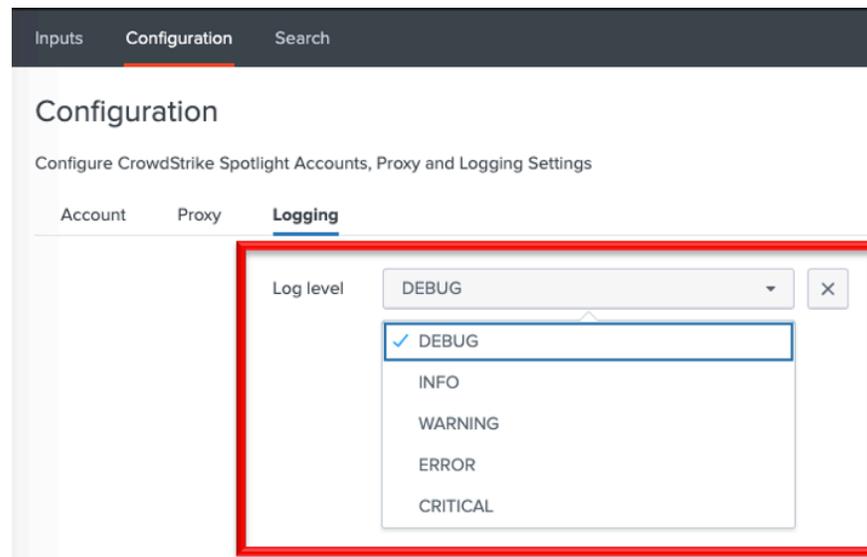
The TA logging level is set to 'info' by default and will only log a minimal amount of information. To properly troubleshoot issues with the TA the logging level should be set to 'debug'.

Change Logging Level

1. Navigate to the Configuration section, Logging tab:



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



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

Review Log Data in Splunk

1. Run and review the **CrowdStrike Spotlight 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.

Compare the Total to the Event Count

As of version 3.2, it's possible to count the number of events that were indexed by Splunk and compare it to the total that was reported by the API return. This is accomplished by using the counting the number of events associated with a single 'ta_data.Collection_hash' field value and comparing it to the 'meta.pagination.total' field value of the same collection.

For Example:

```
meta: { [-]
  pagination: { [-]
    after: WzE2NjY2NTAxODIwMDAsIjVmOWFjMTcxN2U0MzQ2OTk4ZDUzNGVhMzJkZjI4Y2JkXzAyMjU1M2U5NjNmNDMzODk5NjI
    limit: 4000
    total: 871
  }
  powered_by: spapi
  query_time: 5.06953326
  trace_id: 6987608b-9dfc-4d5b-828c-28d38bd56e69
}
ta_data: { [-]
  Cloud_environment: us_commercial
  Collection_hash: 12321135320991956297294923938745763004992849458153822741274987601218027607546
```

Total Number of Events

Field to count by

Here's there's a total of 871 events that matches the API query. Counting by the value of the ta_data.Collection_hash:

ta_data.Collection_hash	count
12321135320991956297294923938745763004992849458153822741274987601218027607546	871

We can see that the two numbers align.

NOTE: In the event that the count is less than the total ensure the following:

1. The input has completed collecting
2. The search timeframe is set to encompass all possible timestamps
3. Validate the number indicated in the TA log as being processed, in the event that the TA process number matches the total but not the count the issue is most likely something in Splunk happening between ingestion and indexing

Differences in these counts does not automatically indicate an issue with the TA's functionality. If the TA logs show the correct number of events being sent to Splunk for processing then it is most likely an internal Splunk issue.

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. The event count in Splunk does not align with the count in the UI:

- 4.1. Ensure that any filtering on the data collection is taken into account such as the use of a start date filter, which will only collect data that has been updated from that date moving forward
- 4.2. Review situations above and ensure that none are having an impact on the data collection
- 4.3. Ensure that the Splunk search and the associated time window will encompass all the potential data
- 4.4. Keep in mind that Spotlight is continually evaluating the hosts status and reporting that information to the UI, as such there will most likely always be a delta and depending on the size of the environment this may cause some significant numbers
- 4.5. Data with systems that have been removed from the platform will not be visible in the UI but will remain in the API for proximately 72 hours and in addition the data will still be retained within Splunk itself

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
 - 6.1. Run the **CrowdStrike Spotlight Logs – 30 days** Report with the time picker set to 'All Time' and export all the results in the RAW format
 - 6.2. (If possible) Download the all-log files containing 'ta_crowdstrike_falcon_spotlight-data' under the \$Splunk/var/log/splunk/ directory
 - 6.3. 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

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

(Access to the CrowdStrike Falcon UI Required)

[CrowdStrike Spotlight Guide](#)

[Spotlight API Guide](#)

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.

