Expanse Technical Add-On v3.1.0

Documentation and User Guide

Expanse's Splunk integration allows you to consume and access Issue Updates, Assets, and Behavior data through Splunk. This documentation details how to install the Expanse Technical Add-on, how to configure your own Expander and/or Behavior data as a Splunk data input, how to configure the add-on to use a proxy, how to search your Expander and Behavior data through the Splunk UI using Splunk data queries, and what's next. For additional information on how to use Splunk more generally, please visit the Splunk documentation site.

Add-On Goals and Outcomes

Examples of goals and outcomes for customers using the Expanse Technical Add-on include:

- Ease-of-use for data querying in a commonly-used SIEM
- Centralized alerting
- Centralized location for security-related data
- Ability to correlate Issue Updates and Assets to internal events tracked in Splunk
- Ability to correlate Behavior data to internal issue updates tracked in Splunk
- Ability to create custom reporting, dashboards, and visualizations
- The context for IPs, and risky flows observed on your network perimeter

Learn more about Expanse's Behavior product
Learn more about Expanse's Issues product

Common Use Cases

Determine Risky Behaviors

Pull in all data based on communications between internal and external IPs. This flow data will contain information about the reason the communication was risky, the internal and external IPs and ports, and the internal domains associated with the IP.

New issue awareness on network perimeter (On-Prem and Cloud)

Ingest and take action on newly appeared or reappeared Issues for a given time period, such as active Remote Desktop Protocol (RDP) servers or Telnet services, with relevant context like IP, port, protocol, banner response, and tags.

Enrich SIEM Alerts with Asset data to speed up Mean Time to Remediation (MTTR) for Issues

Pull in relevant asset data for adding context to an Alert, including CIDR Range, Point of Contact, Attribution Reasons to enable escalation and research

IP Address Audit

Complete IP address audit, creating a graph of IP address results including number of IPs by IP audit result. These IP addresses will include those found by both Expanse and the customer, found only by Expanse, and found only by the customer, when applicable.

Understand Attack Surface by Subsidiary/Network to ensure policy compliance

Pull Asset information by Business Unit/Network Segment (e.g. PCI segment) and cross reference against Issue

Updates to ensure adherence to network policy.

Certificate expiration compliance. Know when different types of certificates will be expiring soon.
Track certificates observed by Expanse to enable proactive compliance.

Domain compliance
Track domains discovered by Expanse to enable proactive compliance.

Installing the Expanse Technical Add-On

1. Click "+ Find more apps" and search for "Expanse" to find the Expanse Technical Add-on
    a. Alternatively, you can download the Add-On by browsing to **https://splunkbase.splunk.com/app/4622/**

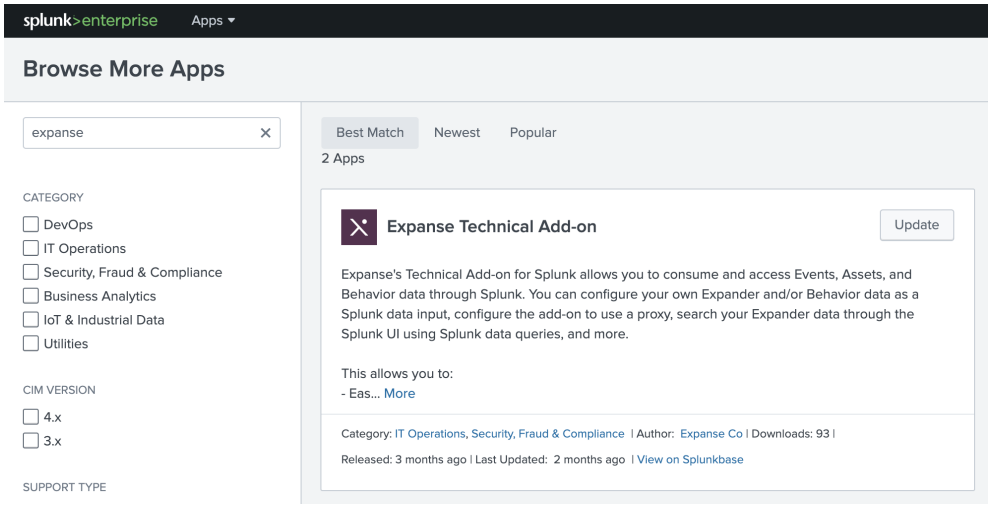

2. Click "Install" to install the Expanse Technical Add-on

3. Restart Splunk to complete installation of the Expanse Technical Add-on

Configuring Your Expanse Technical Add-On

1. Click on the Expanse app to open up the Add-On's settings.

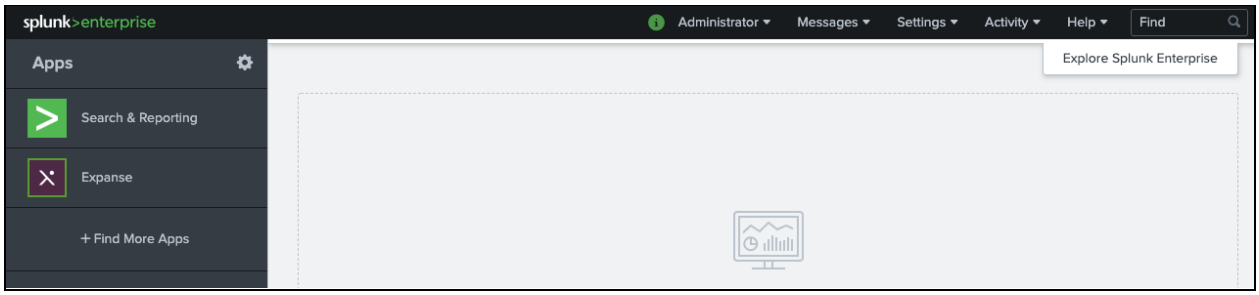

2. Click on the "Configuration" tab. This is where you will configure which Splunk account the Add-On will use, as well as set up your proxy settings (if applicable).

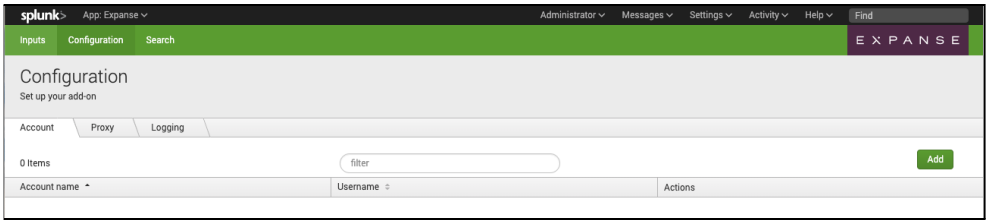

3. Click "Add" to add your Splunk account. The Add-On requires these credentials to store your asset data in Splunk's key-value store. *Note: These need to be the credentials for your Splunk installation, not your Splunk.com credentials.*

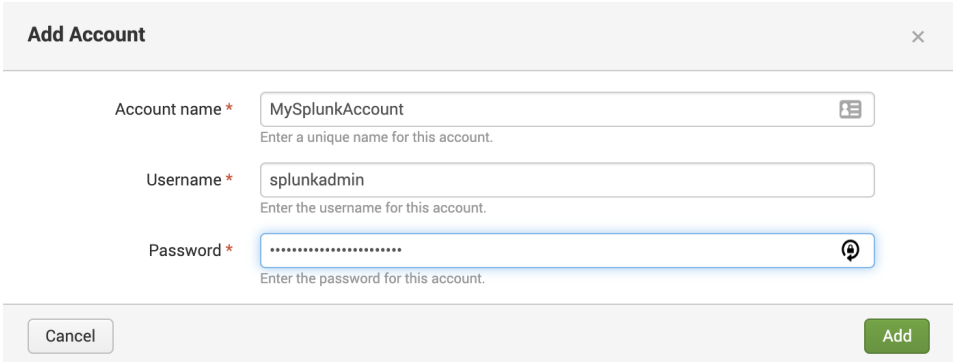

4. [Optional] Navigate to the "Proxy" tab and configure your proxy settings, if applicable. *Note: Make sure to check the "Enable" checkbox to enable your proxy settings, and click "Save" once you've added all of your configuration details.*

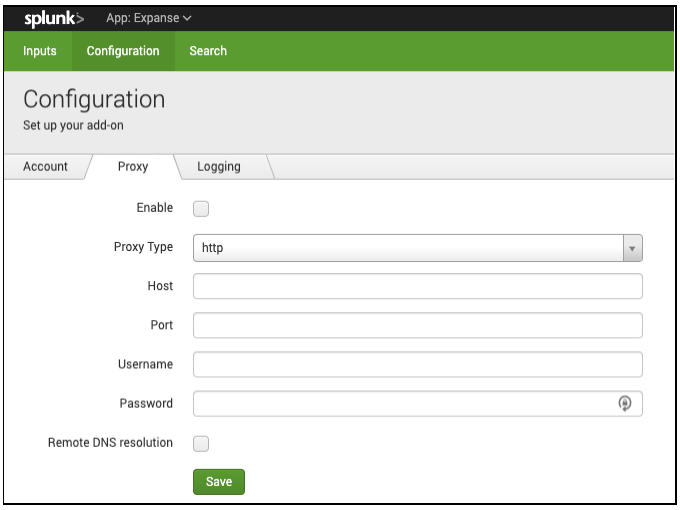

5. Navigate to the "Inputs" tab. This is where you will configure the data input.
NOTE: Only one data input is supported by the TA. Trying to add more than one data input will result in the TA not functioning as expected.

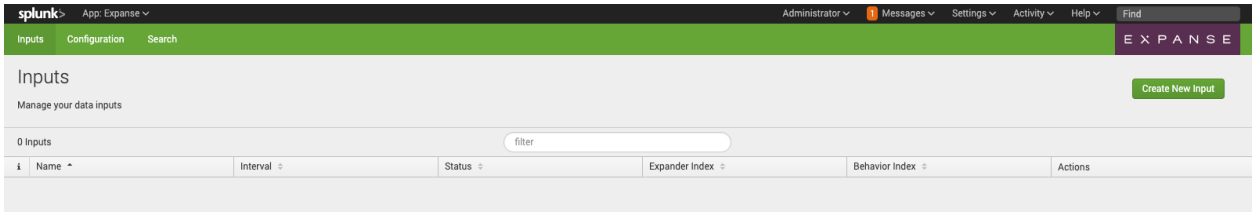

6. Click “Create New Input”.
    a. Give the input a Name.
    b. The Interval determines how often the Add-On will pull new data from the Expanse platform. This value must be given in seconds and must be no less than 10 minutes (600 seconds). We recommend setting this to anywhere between 10 minutes (600 seconds) and 1 hour (3,600 seconds). The Expanse Add-On will default to 10 minutes.
    c. Provide your API Token. *Note: If you have not already received the API token from Expanse, please contact your technical engagement manager. Expanse will securely send this token to you so you can access your data through Splunk and any other APIs or integrations you may wish.*
    d. The Server URL will automatically populate with Expanse’s API endpoint URL.
    e. [Optional] Enter a Start Date for the Add-On to backfill data from. This should be in the format YYYY-MM-DD. Note that the earliest Start Date value is 90 days from the current day.
    f. Select the account you created in step 3 for your Global Account
    g. Choose the Issues Updates Index where you would like the Expander issues updates data to be written. NOTE: As a best practice, we recommend importing data into a “test” index before ingesting into a Production index to ensure compatibility with existing alerts, dashboards, etc.
    h. [In distributed Splunk environments] Configure whether the Add-On should pull Issues Updates data
    i. [In distributed Splunk environments] Configure whether the Add-On should pull Asset data
    j. [If you are a Behavior customer] Choose the Behavior Index where you would like the Behavior risky flow data to be written.
    NOTE: As a best practice, we recommend importing data into a “test” index before ingesting into a Production index to ensure compatibility with existing alerts, dashboards, etc. Also note that this index can be but does not have to be a different index from your Expander Index
    k. [If you are a Behavior customer] Configure whether the Add-On should pull Behavior data
    NOTE: Behavior data has a maximum backfill time of 28 days from the date of ingestion
    l. [In distributed Splunk environments] Configure whether the Add-On should pull Services data
    NOTE: you will need to scroll down in the inputs in order to see the Enable Services checkbox

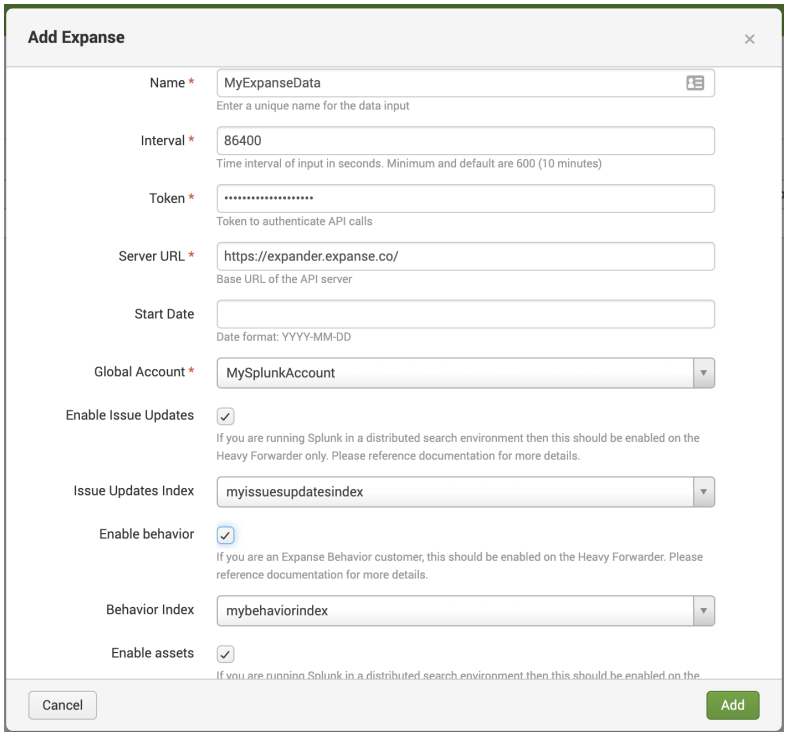

NOTE: If during this process you mistakenly add the New Input with an incorrect API token, you'll need to delete the Input and create a new one (replacing the incorrect token with the correct one will not enable the input to function and allow data to be retrieved).

7. Click "Add". Your Expander and/or Behavior data is now set up as a Splunk data input. Navigate to the "Search" tab within the Add-On to begin querying

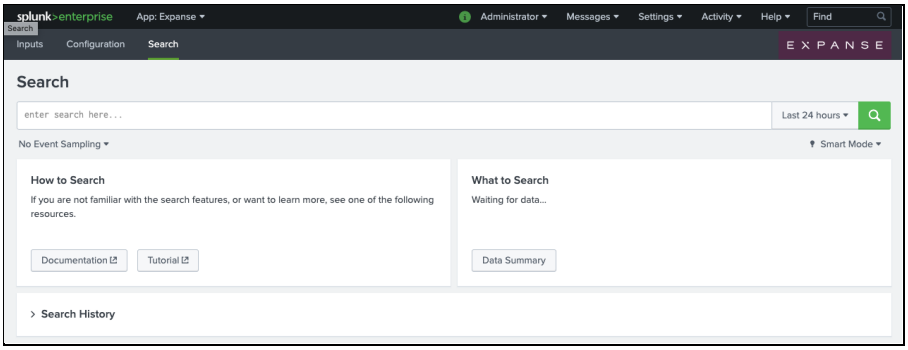

8. Using Splunk data query practices (see example queries in the section below), you can now access and query your Expanse data through Splunk

Using the Splunk Add-On and Example Queries

Splunk represents Expanse Issues Updates and Behavior Flows as Splunk events with JSON objects; associated information to those events is the JSON object's values.

To append context about IP Range Assets to Expander Issue Updates, leverage the IP Range Asset lookup table ("ipcollection_lookup").

Expander Issues Updates currently available in the TA are focused on Issues:

`ActivityStatus`	The activity of the issue (Active, Inactive) has changed.
`Assignee`	The issue has either been assigned to a user or changed to a different user in Expander.
`Priority`	The priority of the issue (Critical, High, Medium, or Low) has changed.
`ProgressStatus`	The progress of the issue has been changed.
`Comment`	A comment has been added to the issue in Expander.

For more information on Issues, please visit the **Expanse Knowledge Base**.

You can query Issues Updates,, Behavior Flows, or Assets using Splunk queries (examples of which are included later in this section).

For more help on Splunk data querying, please refer to the appropriate Splunk data querying documentation or ask your Splunk technical contact.

Time Filters

The time filter dropdown in Splunk allows you to limit your search to include only events with "scanned" dates that fall within a particular time-frame. The "scanned" date for an event in Splunk corresponds to the "created" value of an Expander UI issue update.

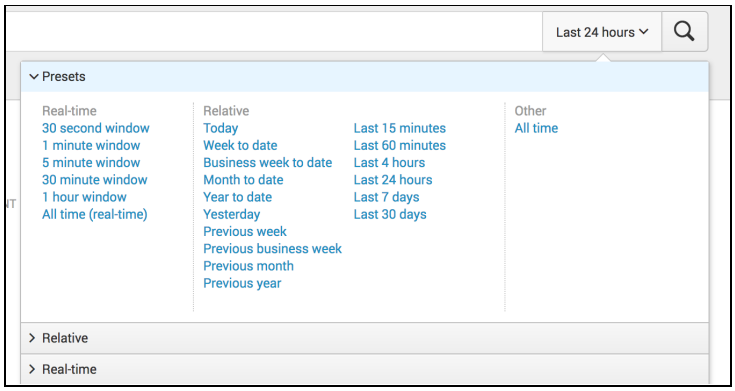

Exporting Issue Update and Behavior Data

You can use the exporting functionality to download event and risky flow searches to CSV, XML, and JSON.

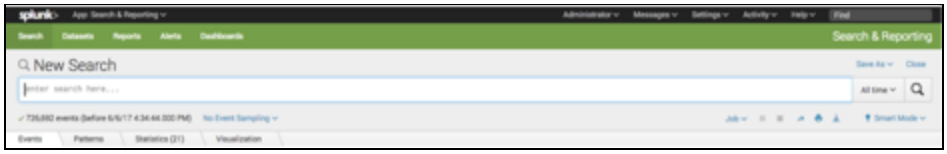

Example Data Queries

NOTE: To query your data, search the index you configured for the input and set "sourcetype" to be "Expanse":

Use Case	Query
How do I view all issues updates or behavior data?	`index="myindex" sourcetype="Expanse"`
How do I view issue updates related to OnPrem assets?	`index="myindex" sourcetype="Expanse" | SEARCH "providers{}.name"="On Prem"`
How do I filter my search on a particular value of an issue update object, such as	To see only issues with a description of "Web Login", search: `index="myindex" sourcetype="Expanse" | search description="Web Login"`

"description", "host", "businessUnit.name", etc.?	To see only events with a business unit name of "Acme Co", search: `index="myindex" sourcetype="Expanse" | search "businessUnits{}.name"="Acme Co"`
How do I filter my search on an issue update object to exclude particular values, such as those not related to changes in ActivityStatus on the issue?	To exclude all issue updates that are not a change in ActivityStatus, search: `index="myindex" sourcetype="Expanse" | search updateType="ActivityStatus"`
How do I view all assets attributed to my network to understand the asset's context and relationship to their organization? (Note: Only On-Prem Assets are supported in v1.02 of the TA)	`|inputlookup ipdatacollection_lookup`
What are the Business Unit and point of contact details for a given IP?	`|inputlookup ipdatacollection_lookup | where cidrmatch(cidr,"X")| fields businessUnits.name, annotationPointsOfContact.firstName,annotationPointsOfContact.lastName,annotationPointsOfContact.email,annotationPointsOfContact.phone` Where **X** represents the IP Address being looked up in the Assets Lookup Table.
How do I view domains and their associated metadata?	`|inputlookup domainscollection_lookup`
How do I view certificates and their associated metadata?	`|inputlookup certificatescollection_lookup`
How do I view services and their associated metadata?	`|inputlookup servicescollection_lookup`

Upgrading to v3.x.x

Version 3.x.x and later of the Expanse Technical Add-On includes:

- Issues Updates data not previously available, replacing Exposures and Events.
- Daily asset updates (previously weekly)
- Daily services updates
- Python 2 and Python 3 dual compatibility.

Guidelines for upgrading to v3.x.x from v2.3.2 or below

1. Navigate to the Expanse Technical Add-On in Splunk.
2. Confirm the Expanse Technical Add-On is not ingesting data.

a. Navigate to "Search", and run the following:
   `index="_internal" source=*expanse*`
b. The last log will include "Process finished"

3. Navigate to the "Inputs" tab under the Expanse Technical Add-On.
4. "Disable" the current Expanse input under "Actions".
5. Upgrade the Expanse Technical Add-On to the version 3.0.0, which can be done under "Manage Apps"
   a. Alternatively, you can download the app in Splunkbase here **https://splunkbase.splunk.com/app/4622/**.
6. Restart Splunk to complete the update of the Technical Add-On.
7. [If you plan to ingest Issues Updates] Create a new index for Issues Updates
8. [If you plan to ingest Issues Updates] Navigate to "Inputs" under the Expanse Technical Add-On, and update the Expanse configuration as follows:
   i. "Enable Issues Updates" option is checked off
   ii. "Issues Updates Index" points to your new issues updates index, created in Step 5.

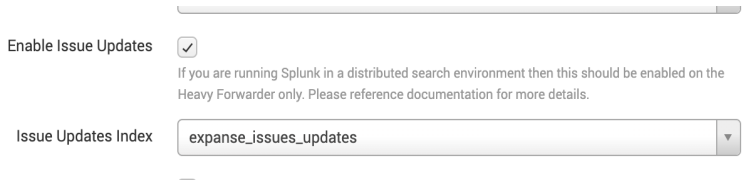

9. Navigate to the "Inputs" tab under the "Expanse" Technical Add-On and "Enable" the current Expanse input.
10. [If you were previously ingesting Events and Exposures] Update any alerts or dashboards that use Events or Exposures data to use Issues Updates data instead.

Troubleshooting the Expanse Technical Add-On

If you are having difficulty with the Add-On, your first step should be to view the logs. In the Add-On's "Search" page, you can view the logs with the following command:

```
index="_internal" source="*expanse*"
```

If you need to uninstall the app and clean up your data, run the following commands on your Splunk instance from the command line. *Note: The command to clear out the index data will* ***delete all data*** *from that index.*

Run these commands from the `bin` directory of your Splunk installation:
`cd $SPLUNK_HOME/bin`
Stop your instance of Splunk
`./splunk stop`
[Optional] Remove all data from the index. **This will clear all data in the index regardless of its source**
`./splunk clean eventdata -index myindex`
Restart Splunk
`./splunk start`
Clear all asset data from the key-value store
`./splunk clean kvstore -app TA-expanse`
Delete the app
`./splunk remove app TA-expanse`

Installation Configuration Options

Splunk can be architected in many different configurations, ie: standalone, distributed search and/or high availability. The following are just ideas but not recommendations since every environment can be different. Please verify this with the supported Splunk Docs.

Splunk Architecture	Installation Recommendations
`Standalone`	Install the Expanse Technical Add-On per instructions in this guide
`Distributed Search`	Install the Expanse Technical Add-On on the Heavy Forwarder and configure it per instructions. Install the Expanse Technical Add-On on the Search Head but do not configure it. There are CIM (Common Information Model) fields that are needed on Search Head.
`Search Head Cluster`	Install the Expanse Technical Add-On on the Heavy Forwarder and configure it per instructions. Install the Expanse Technical Add-On on one Search Head but do not configure it. There are CIM (Common Information Model) fields that are needed on Search Head. The Search Head Replication process should deploy the configurations to the other Search Heads.
`Cloud Deployments`	Install the Expanse Technical Add-On on the Heavy Forwarder and configure it per instructions. Install the Expanse Technical Add-On on one Search Head but do not configure it. There are CIM (Common Information Model) fields that are needed on Search Head. The Search Head Replication process should deploy the configurations to the other Search Heads. Set up the Heavy Forwarder to get data into Splunk Cloud, as outlined in the Splunk Documentation.

KV Store in Search Head Cluster and Cloud architectures

Search Head Cluster

Refer to Appendix A for Search Head Cluster deployment recommendations for KV Store data.

Cloud

The Expanse Technical Add-On does not natively support KV Stores in Splunk Cloud. However, there is a workaround for this using temporary indices and lookup files. If you plan on ingesting Expanse Asset data and utilize one of these architectures, here are example commands for utilizing this workaround.

Example:

Using `|collect` where index=summary_expanse

```
| inputlookup domainscollection_lookup|collect index=summary_expanse
```

Set up a search on this data where index=summary_expanse

```
index=summary_expanse |table fields *| outputcsv create_empty=false
domainscollection_lookup

| inputcsv domainscollection_lookup
```

You can schedule a report with the above search in the heavy forwarder to ingest the asset data to index at a set interval.

Questions about the Expanse Technical Add-On

Currently, the Expanse Technical Add-on v3.0.0 leverages Issues Updates, Behavior, Assets, and Services APIs to provide data about your OnPrem and Cloud Issues Updates, Risky Flows, Assets, and Services. Future

Please reach out to support@expanseinc.com or your engagement manager with any questions, concerns, or feedback.

Appendix

A. Search Head Cluster Deployment guide for KV Store data

1. Download the Expanse Technical Add-On from Splunkbase.

2. After you untar the .tgz file, you should get a directory named `TA-expanse`.

3. Create a new directory named `local` within `TA-expanse`.

4. Create a new file named `ta_expanse_account.conf` in `TA-expanse/local`. Populate it with the following:

```
[account_name]
password = splunk_admin_password
username = splunk_admin_username
```

Field	Meaning
`account_name`	This can be anything, "Expanse" is a good default.
`splunk_admin_password`	This is the password for the user with admin permissions you would like to configure the TA with.
`splunk_admin_username`	This is the username for the user with admin permissions you would like to configure the TA with.

5. Create a new file named `inputs.conf` in `TA-expanse/local`. Populate it with the following:

```
[expanse://input_name]
enable_assets = 1
enable_behavior = 0
enable_events = 0
enable_exposures = 1
```

```
global_account = account_name_from_previous_step

index = default

behavior_index = default

interval = 86400

server_url = https://expander.expanse.co/

start_date_utc = date

token = expanse_refresh_token
```

Field	Meaning
input_name	This can be anything, “expanse” is a good default.
account_name_from_previous_step	This should be the value for account_name from step 4.
date	This should be a recent date in the format of YYYY-MM-DD.
expanse_refresh_token	This is the refresh token that has been provided to you by your Expanse Engagement Manager.

6. Use scp or some other means to copy TA-expanse onto your Deployer server. Copy into `$SPLUNK_HOME/etc/shcluster/apps`

7. On the Deployer server, run this splunk command to bundle and send apps and configurations to the SH cluster captain.

```
sudo splunk apply shcluster-bundle -target <cluster_leader>:8089 -auth
<username>:<password>
```

NOTE: this will cause a rolling restart.

NOTE: if you're currently deploying other apps to your cluster using this method be aware of any consequences of pushing a new bundle. This has only been tested with the default deployer push mode, `merge_to_default`. Read more here - https://docs.splunk.com/Documentation/Splunk/8.0.6/DistSearch/PropagateSHCconfigurationchanges

Field	Meaning
`cluster_leader`	This would be the url for your cluster captain. Ex. https://10.10.10.10
`username`	This is the username for your splunk admin user, or another user with permissions to perform a deployment.
`password`	This is the password for your splunk admin user, or another user with permissions to perform a deployment.

8. After the command above completes successfully and your search heads have restarted, log into a search head and attempt to query one of the asset lookups, you should see results returned.

```
| inputlookup certificatescollection_lookup
```