



Illumio App for Splunk

# Illumio App for Splunk 3.2.0

## Document History

Date	Changes
May 29, 2017	Initial draft
June 6, 2017	Added CIM information
October 18, 2017	Addition of new features with V1.1 of the App
December 12, 2017 - Jan 8 2018	Editing and style changes
May 5, 2018	Cosmetic prep for Illumio ASP 18.1
Sept 21, 2018	<ul style="list-style-type: none"> <li>• Updates for Illumio ASP 18.1 and 18.2</li> <li>• Illumio App for Splunk v2.0.0</li> </ul>
Jan 3, 2019	<ul style="list-style-type: none"> <li>• Cosmetic update to table of contents</li> <li>• Version stamp for Illumio ASP 18.2.1</li> </ul>
July 22, 2019	<ul style="list-style-type: none"> <li>• Added new configuration, dashboards for App version 2.2.0</li> <li>• Added more Troubleshooting sections</li> </ul>
Sept 9, 2019	Updated the guide version from 2.2.0 to 2.2.1 to match the shipping version of the Illumio Splunk App
Dec 16, 2019	Updated for version 2.3.0. Added the Alert Configuration screen, Alerts screen, updated screen shots, support for Illumio ASP 19.3 and API v2, drill down on Audit Events in Workload Investigation Dashboard, traffic flow timestamp = VEN timestamp, improved support for S3 data, example Splunk queries, added <code>ip_lists</code> and <code>services</code> to list of REST API endpoints called

February 21, 2020	Updated for version 3.0.0. Added: Splunk 8.0 is supported; app is compatible with Python 2/3; how to edit <code>summariesonly</code> macro for faster UI performance; after upgrade, rebuild data model and remove local customizations; new label filters in Workload Investigation dashboard; new Allowed option in Security Operations dashboard.
August 27, 2021	Correct instructions for configuring Splunk Add-on for AWS so that key prefix uses 'illumio' instead of <UUID>.  Add clarification on quarantine workload using Enterprise Security Suite
October 28, 2021	Correct error in PCE Secure Cloud log description—all traffic flows are logged.
November 4, 2021	Updated for version 3.2.0. Added new events in CIM Mapping Table, and three new dashboards—PCE Authentication Events, Change Monitoring, and Traffic Explorer

## Table of Contents

<b>Architecture</b> .....	<b>1</b>
App Components .....	2
<b>TA-Illumio (Technology Add-on for Illumio)</b> .....	<b>3</b>
Splunk Index, Source, and Source Types .....	4
Index .....	4
Sourcetype .....	5
Field Extractions .....	5
Purpose of TA-Illumio at Different Splunk Components .....	5
Data Model and Data Model Acceleration .....	6
CIM mapping .....	7
<b>Illumio App for Splunk</b> .....	<b>8</b>
Dashboards .....	9
Security Operations Dashboard .....	9
PCE Operations Dashboard (On-Prem Only) .....	12
PCE Authentication Events Dashboard .....	14
Workload Operations Dashboard .....	14
Workload Investigation Dashboard .....	16
Traffic Explorer Dashboard .....	18
Alert Configuration Page .....	20
Alerts Page .....	20
Change Monitoring Dashboard .....	20
<b>Installation</b> .....	<b>22</b>
Installation Prerequisites .....	22
Splunk Single Server Deployment .....	22
Splunk Distributed Deployment .....	22
Installing TA-Illumio .....	22
Application of TA-Illumio to Splunk Components .....	22
Installing Illumio App for Splunk in a Distributed Splunk Environment .....	23
Using Splunk Heavy Forwarder .....	24
Using Splunk Universal Forwarder .....	24
Splunk Cloud Instance Deployment .....	25
Installation Commands .....	25
<b>Configuration</b> .....	<b>26</b>
Splunk App Configuration .....	27
On-Premise PCE Configuration .....	34

PCE runtime_env.yml Configuration .....	34
Illumio Cloud PCE Configuration .....	34
Configure Amazon S3 Bucket .....	35
Configure Splunk Add-on for AWS .....	39
Speeding Up UI Rendering.....	44
Configuring Alerts .....	45
<b>Post-Installation Required Settings .....</b>	<b>50</b>
Accelerate Data Model .....	50
Update Search Macros for Custom Index .....	51
<b>Upgrade the App .....</b>	<b>51</b>
<b>Alerting Actions and Adaptive Response Framework .....</b>	<b>52</b>
Quarantine Workload Using Splunk Core Alert Actions .....	53
Quarantine Workload using Enterprise Security Suite .....	53
Quarantine Workload from Illumio Splunk App.....	57
Access to Quarantine Workload Action .....	57
<b>Example Splunk Queries .....</b>	<b>59</b>
<b>Troubleshooting.....</b>	<b>61</b>
Data collection not working.....	61
Can't use same port in new Data Input (Modular Input).....	61
Data not available immediately after configuring data input (modular inputs).....	62
Authentication failure on Data Input (Modular Input) page.....	62
Quarantine button is grayed out or does not work as expected .....	62
Invalid Certificate File error on Data Input (Modular input) page.....	63
PCE labels are not updated in Security Operations dashboard .....	63
Security Operations shows "Search is waiting for input" .....	64
Path for the custom certificate: invalid certificate file.....	64
Authentication Failed: Invalid PCE URL or API key id or API Secret.....	68
Sankey diagram is not displayed in Traffic Explorer dashboard.....	71
Label filters (i.e.App, Env and Loc) are not populated.....	72
<b>Known Limitations .....</b>	<b>72</b>
<b>Compatibility Matrix.....</b>	<b>72</b>

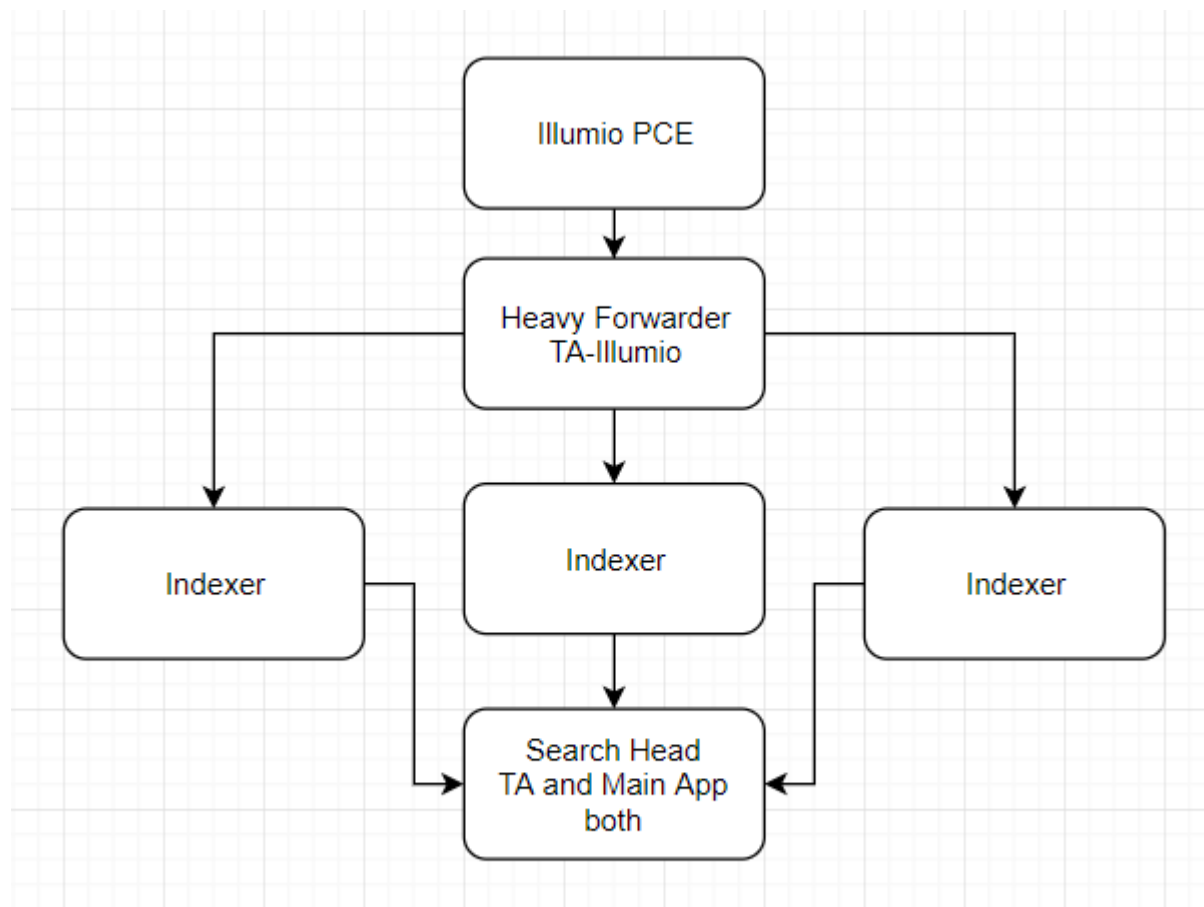
## Architecture

The Illumio App for Splunk integrates Splunk with the Illumio Policy Compute Engine (PCE). Using the app, you can conveniently access PCE data through Splunk, and gain security and operational insights into your Illumio-secured data center.

The Technology Add-on for Illumio (TA-Illumio) performs data collection, data normalization, and data visualization using data that comes from the Illumio Policy Compute Engine (PCE) through REST API calls and syslog.

The following diagrams show a typical data collection architecture from PCE to Splunk in distributed and standalone environments.

### Splunk Distributed Environment



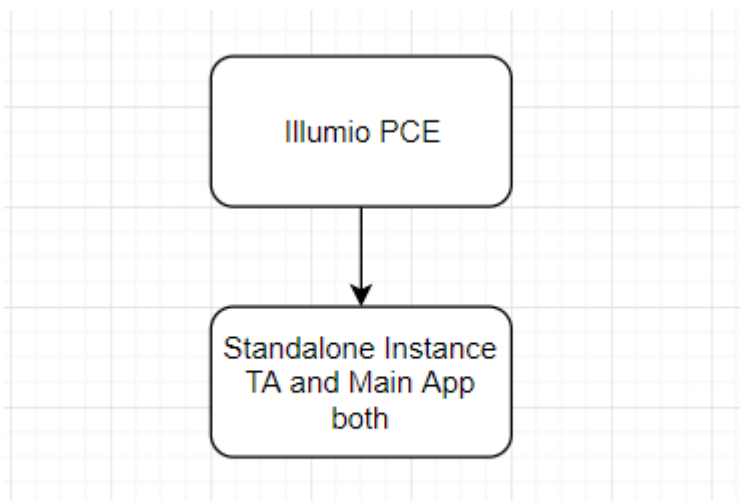
**Topology in Distributed Environment**

For information about how to install each component in a Splunk distributed environment, see [“Application of TA-Illumio to Splunk Components.”](#)

If you use Splunk Universal Forwarder on a dedicated data collection node, see [“Using Splunk Universal Forwarder.”](#)

### Splunk Standalone Environment

In a standalone environment, the PCE forwards data directly to the Splunk instance. The Splunk heavy forwarder is not involved.



### Topology in Standalone Environment

## App Components

The Illumio app for Splunk is comprised of two parts:

- TA-Illumio (Technology Add-on for Illumio)
- Illumio App for Splunk

TA-Illumio and the Illumio App for Splunk are typically deployed together in the search head. TA-Illumio receives and transforms data, and enriches events with CIM fields; it can also be deployed at the indexer or forwarder. The Illumio App for Splunk uses the data enriched by TA-Illumio to display informational dashboards.

## TA-Illumio (Technology Add-on for Illumio)

The Technology Add-on for Illumio (TA-Illumio) is a Splunk module that receives PCE data for Splunk and performs data normalization. TA-Illumio collects data from the PCE and enriches the data according to the Common Informational Model (CIM). CIM is the native data representation used by Splunk. Illumio data in CIM format can be used easily with Splunk applications such as Splunk Enterprise Security and Splunk App for PCI Compliance.

Data collection from the PCE is accomplished in two ways: through the Illumio ASP REST API and the Illumio PCE syslog.

The Adaptive Response Framework components that are used by Splunk Enterprise Security Suite are packaged with TA-Illumio.

### **Illumio ASP REST API**

TA-Illumio pulls data using the Illumio ASP REST API. For data collection to work, you must set up the API configuration in TA-Illumio to use Data Input, also known as modular input. Data collected from API calls is used to create metadata for workloads, labels, and services. The API data is used to enrich syslog data, such as traffic flow summaries and auditable events.

The following Illumio ASP REST API endpoints are called:

- GET /api/v2/orgs/1/workloads/
- GET /api/v2/orgs/1/labels/
- GET /api/v2/orgs/1/health/
- GET /api/v2/product\_version
- GET /api/v2/orgs/1/sec\_policy/draft/ip\_lists
- GET /api/v2/orgs/1/sec\_policy/draft/services

### **Illumio PCE Syslog**

TA-Illumio receives and processes messages directly from the PCE using the TCP port configured in Data Input (modular input). The types of messages are:



- Events, which are structured JSON messages representing auditing information.
- Traffic flow summaries, which are structured JSON messages representing enriched traffic flows. Traffic flow summaries contain flows, Illumio labels, and other data about the flow.
- PCE System Health messages in syslog format (key-value pairs).
- Other syslog messages.

## Splunk Index, Source, and Source Types

Index and source type are default Splunk fields used to categorize and filter the indexed data to narrow down search results.

### Index

In Splunk, raw syslog data is stored in indexes, classified by source type. With TA-Illumio, you can select an index while creating Data Input (modular input). Data collected from that modular input will be collected into the selected index.

If you chose the default index in Data Input, no further configuration is needed.

If you chose a non-default index, you must also update the search macros as follows to use the custom index. Otherwise, the dashboards will not display charts.

To modify the search macro:

1. In **Settings > Advanced Search > Search Macros > App: Illumio App for Splunk**, select `Illumio_get_index`.
2. In **Definition**, do one of the following:
  - If you use the default index, enter open and close parentheses:  
`()`
  - If you have created a custom index, enter the name of your index in parentheses:

```
(index=custom_index_name)
```

Search macros New Search Macro

Advanced search > Search macros

Showing 1-7 of 7 items

App: Illumio App for Splunk (I) | Owner: Any | Visible in the App | filter

Name	Definition	Arguments	Owner	App	Sharing	Status	Actions
illumio_get_index	0		No owner	IllumioAppforSplunk	Global   Permissions	Enabled	Clone
illumio_get_time()	\$field\$=strtime('\$field\$', '%b %d %H:%M')	field	No owner	IllumioAppforSplunk	App   Permissions	Enabled	Clone

## Sourcetype

The following table shows how Illumio data is classified by source types.

Source type	Description
illumio:pce	Events collected from the Illumio PCE through syslog.
illumio:pce:metadata	Workloads, labels, iplists, and services collected from the PCE using REST API calls.
illumio:pce:collector	Traffic flow summaries collected from the Illumio PCE through syslog. Note: The time stamp for traffic flow summaries is the stamp in the message itself (as reported by the VEN), and is not the time when the message is received by the PCE or relayed to Splunk. Effectively, the timestamp of traffic flow summaries is the time when the traffic actually occurred.

## Field Extractions

TA-Illumio extracts fields from various source types using regular expressions.

## Purpose of TA-Illumio at Different Splunk Components

TA-Illumio has different purposes at different Splunk components

**Heavy Forwarder:** TA-Illumio at the Heavy Forwarder is used to do data collection. You configure the modular input as described in the Installation section below. Heavy forwarder does not have to be a separate component. It could be the same as an indexer or Search head.

**Indexer:** TA-Illumio may be installed at the indexer for the following purpose: sometimes the PCE sends invalid JSON data which is not supposed to be indexed. Logic to send such events to the null queue is implemented in TA. Users can choose to not install TA on the indexer if they are not worried about indexing additional events.

**Search head:** TA-Illumio at search head is used to do search time field extractions, which will be used by Illumio App for Splunk in visualizations.

## Data Model and Data Model Acceleration

The app consists of one data model named "Illumio". The data model used in this application is not accelerated by default. If you wish to improve the responsiveness of the dashboards, you should enable data model acceleration with a 1-week period. Accelerated data models help improve the performance of the dashboard, but also increase the disk usage on the indexer node.

To enable acceleration:

1. On the Splunk menu bar, click **Settings > Data models**.
2. From the list of data models, click **Edit** in the Action column of the row for the Illumio data model.
3. From the list of actions, select **Edit Acceleration**. The Edit Acceleration menu is displayed.
4. Check the **Accelerate** checkbox to enable data model acceleration.
5. Select the summary range and specify an acceleration period of 1 week.
6. Click **Save**.

If you don't need to use the already indexed accelerated data model, the data model can be configured to rebuild from scratch for the specified acceleration period.

To rebuild the data model:

1. On the Splunk menu bar, click **Settings > Data models**.
2. From the list for Data models, expand the Illumio row by clicking the > arrow in the first column. Additional details appear.
3. From the Acceleration section, click **Rebuild**.

4. Monitor the status of Rebuild in the Status field of the Acceleration section.  
Reload the page to get the latest rebuild status.

## CIM mapping

PCE events are mapped to multiple Common Information Model (CIM) data models as shown in the following table.

Event type	CIM data model	CIM field	Illumio field
sourcetype="illumio:pce"  category = "auditable" event_type="user.sign_in" OR event_type="user.login"	Authentication	src	src_ip
		user	created_by.user.username
		app	"Illumio"
		action	"failure" OR "success"
sourcetype="illumio:pce"  category = "auditable" event_type="agent.tampering" OR event_type="agent.firewall_config"	Network Changes	action	"modified"
		status	status
		vendor_product	"illumio:pce"
		change_type	change_type
		src	src_ip
		user	created_by.user.username
sourcetype="illumio:pce"  category = "auditable" (event_type="*.create" OR event_type="*.delete" OR event_type="*.update") (event_type!="user.*")	Auditing Changes	action	"created" OR "deleted" OR "modified"
		src	src_ip
		status	status
		vendor_product	"illumio:pce"
		user	created_by.user.username
		change_type	change_type

sourcetype="illumio:pce" category = "auditable" event_type="user.create" OR event_type="user.update" OR event_type="user.delete"	Account Management Changes	action	"created" OR "deleted" OR "modified"
		src	src_ip
		status	status
		vendor_product	"illumio:pce"
		src_user	created_by.user.username
		change_type	change_type
		user	resources_changes.resource.username
sourcetype="illumio:pce:collector"	Network Traffic	action	pd
		bytes	tbi + tbo
		bytes_in	tbi
		bytes_out	tbo
		dest	dst_ip
		dest_ip	dst_ip
		dest_port	dst_port
		src	src_ip
		protocol	proto

## Illumio App for Splunk

The Illumio App for Splunk integrates Splunk with the Illumio PCE to provide security and operational insights into your Illumio-secured data center. Multiple dashboards display an overview of your data center while monitoring the PCE and Illumio Virtual Enforcement Nodes (VENs) installed in your data center.

With improved visibility of east-west traffic, your Security Operations Center (SOC) staff can detect unauthorized activity and potential attacks from traffic blocked by Illumio segmentation policies on workloads in the "Enforced" policy state (policy is enforced). Additionally, the Illumio App for Splunk provides visibility into potentially blocked traffic for workloads in the "Test" policy state (policy is visualized but not enforced). This enables SOC staff to quickly pinpoint potential attacks and remedy those situations.

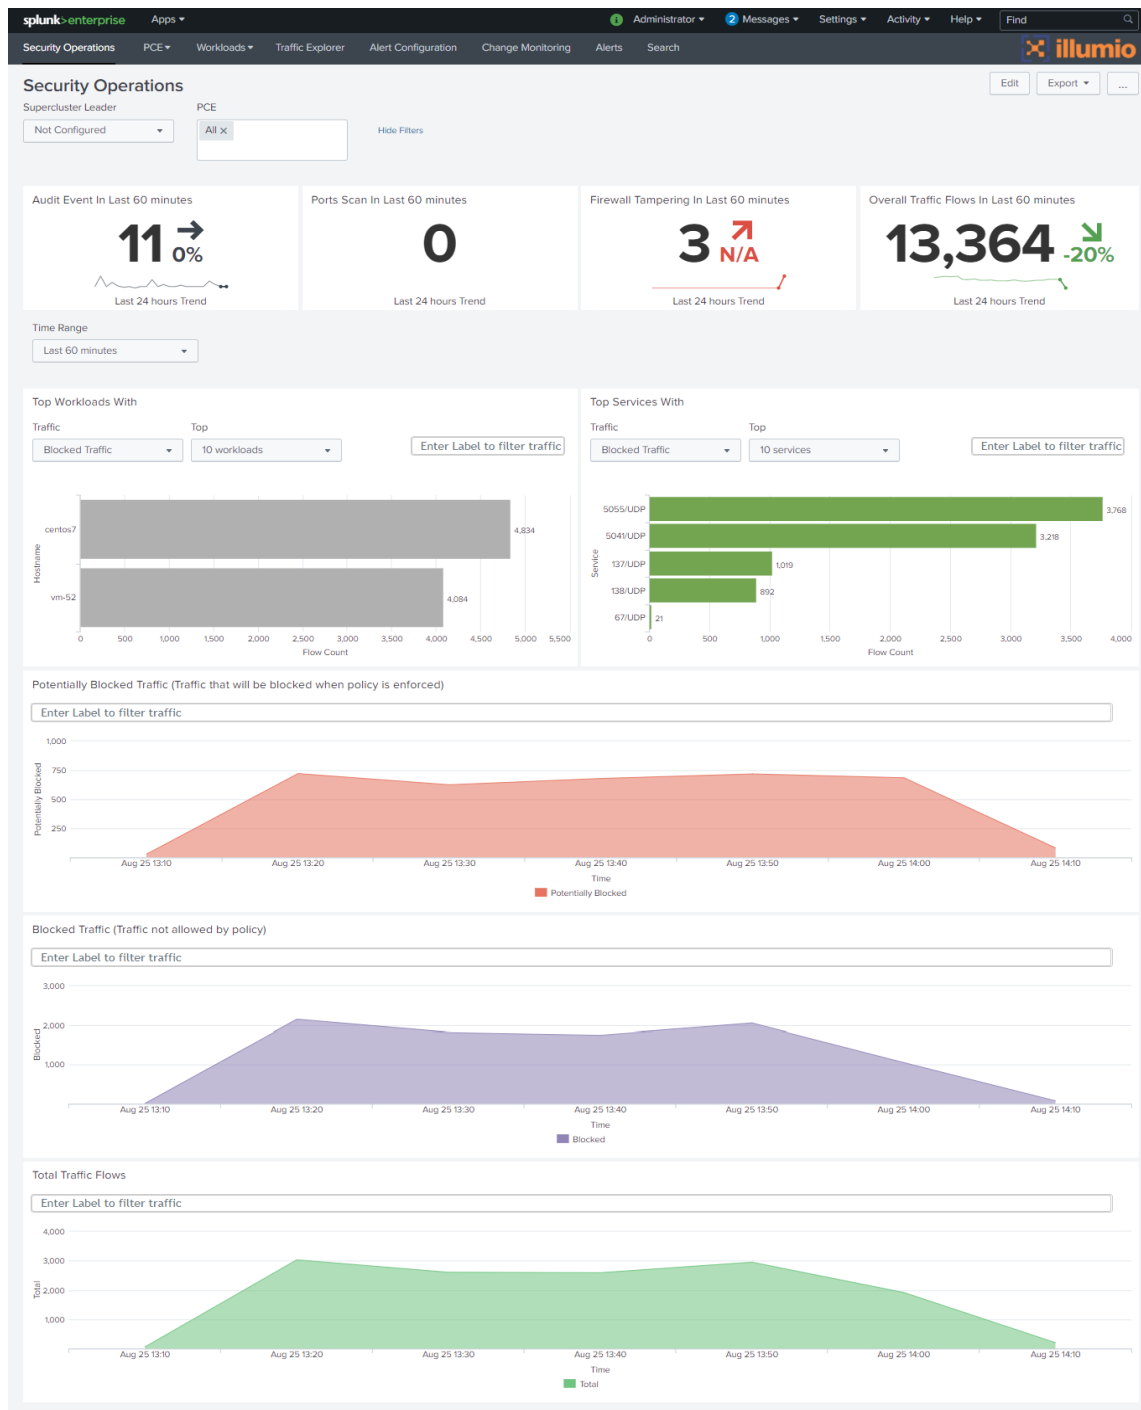
## Dashboards

The Illumio App for Splunk has multiple dashboards to display system activities associated with the PCE instance. The following dashboards can be accessed from the top row of the app:

- Security Operations Dashboard
- PCE Operations (On-Prem Only) Dashboard
- PCE Authentication Events Dashboard
- Workload Operations Dashboard
- Workload Investigation Dashboard
- Traffic Explorer Dashboard
- Alert Configuration Page
- Change Monitoring Dashboard
- Alerts Page

### Security Operations Dashboard

The Security Operations dashboard provides an overview which allows Splunk administrators to monitor the overall security state of the network, as determined from traffic flows reported by PCE instances. Top Blocked, Potentially Blocked, and Allowed traffic is displayed by host and by service. To see Allowed traffic, choose it in the dropdown list under Top Workloads With or Top Services With. In most panels, flows can be filtered using Illumio labels. You can also drill down to investigate notable events, such as Port Scans and Firewall Tampering.



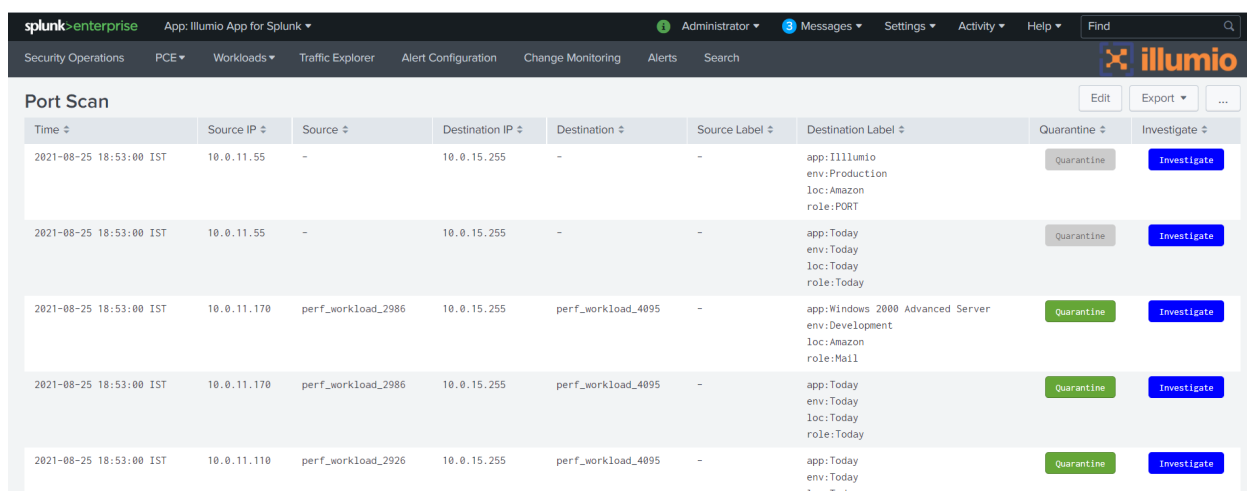
The Security Operations dashboard is built using data from the following sources:

- Traffic flow summaries
- REST API calls made to the PCE

- Events

## Investigate Workload from Illumio Splunk App

When you are viewing a list of workloads, such as through the Port Scan or Firewall Tampering screens, you can click **Investigate** to view the Workload Investigation dashboard for the selected workload. For details, see “[Workload Investigation Dashboard](#).”



Time	Source IP	Source	Destination IP	Destination	Source Label	Destination Label	Quarantine	Investigate
2021-08-25 18:53:00 IST	10.0.11.55	-	10.0.15.255	-	-	app: Illumio env: Production loc: Amazon role: PORT	Quarantine	Investigate
2021-08-25 18:53:00 IST	10.0.11.55	-	10.0.15.255	-	-	app: Today env: Today loc: Today role: Today	Quarantine	Investigate
2021-08-25 18:53:00 IST	10.0.11.170	perf_workload_2986	10.0.15.255	perf_workload_4095	-	app: Windows 2000 Advanced Server env: Development loc: Amazon role: Mail	Quarantine	Investigate
2021-08-25 18:53:00 IST	10.0.11.170	perf_workload_2986	10.0.15.255	perf_workload_4095	-	app: Today env: Today loc: Today role: Today	Quarantine	Investigate
2021-08-25 18:53:00 IST	10.0.11.110	perf_workload_2926	10.0.15.255	perf_workload_4095	-	app: Today env: Today loc: Today	Quarantine	Investigate

Depending on the results of the investigation, you might want to quarantine the workload.



### PCE Operations Dashboard (On-Prem Only)

The PCE Operations dashboard enables Splunk administrators to monitor the health of multiple on-prem PCE instances from one Splunk server. This includes the overall PCE cluster status, service status summary, per-node service status, CPU, Memory and Disk utilization metrics. If multiple PCE instances are connected to Splunk, you can use the dropdown list at the top of the dashboard to choose which PCE to monitor.

The PCE Operations dashboard is built using data from the following source:

- REST API calls made to the PCE (PCE 17.2 and later)

**illumio** App

Security Operations PCE Workbooks Traffic Explorer Alert Configuration Change Monitoring Alerts Search

**PCE Operations (On-Prem Only)**

Time Range: PCE  
 Between Date-From: 2022-09-01 00:00:00.000Z Date-To: 2022-09-01 00:00:00.000Z Refresh

Cluster Status: **Normal** PCE Run Level: **5**

PCE Service Status: **Not Running: 0** **Optional: 1** **Partial: 0** **Running: 44** **Unknown: 0**

Policy Database Summary: Database size: 11.2514 GB Database Disk Utilization: 15.2851 % Transaction ID Max Age: 29397822 Vacuum Backlog: 3.3324 %

Cluster Claws

Node	Status	CPU Utilization	Memory Utilization	Disk Utilization	Policy Disk Latency	Traffic Disk Latency
casf-cs2mnc	Normal	8% (-33%)	26% (0%)	8% (0%)	1ms (-90%)	N/A ms
casf-cs2mnc	Normal	6% (-50%)	24% (0%)	8% (0%)	0ms (-100%)	N/A ms
casf-cs2mnc	Normal	4% (-80%)	17% (0%)	16% (0%)	51ms (750%)	51ms (750%)
casf-cs2mnc	Normal	2% (-71%)	18% (0%)	13% (0%)	0ms (-100%)	0ms (-100%)

Parents below are available for on-premise Splunk PCE 19.3.2 and 20.2.0 or higher versions.

VDR Parent: VDR Health Latency (Average) No results found

Type: Collector No results found

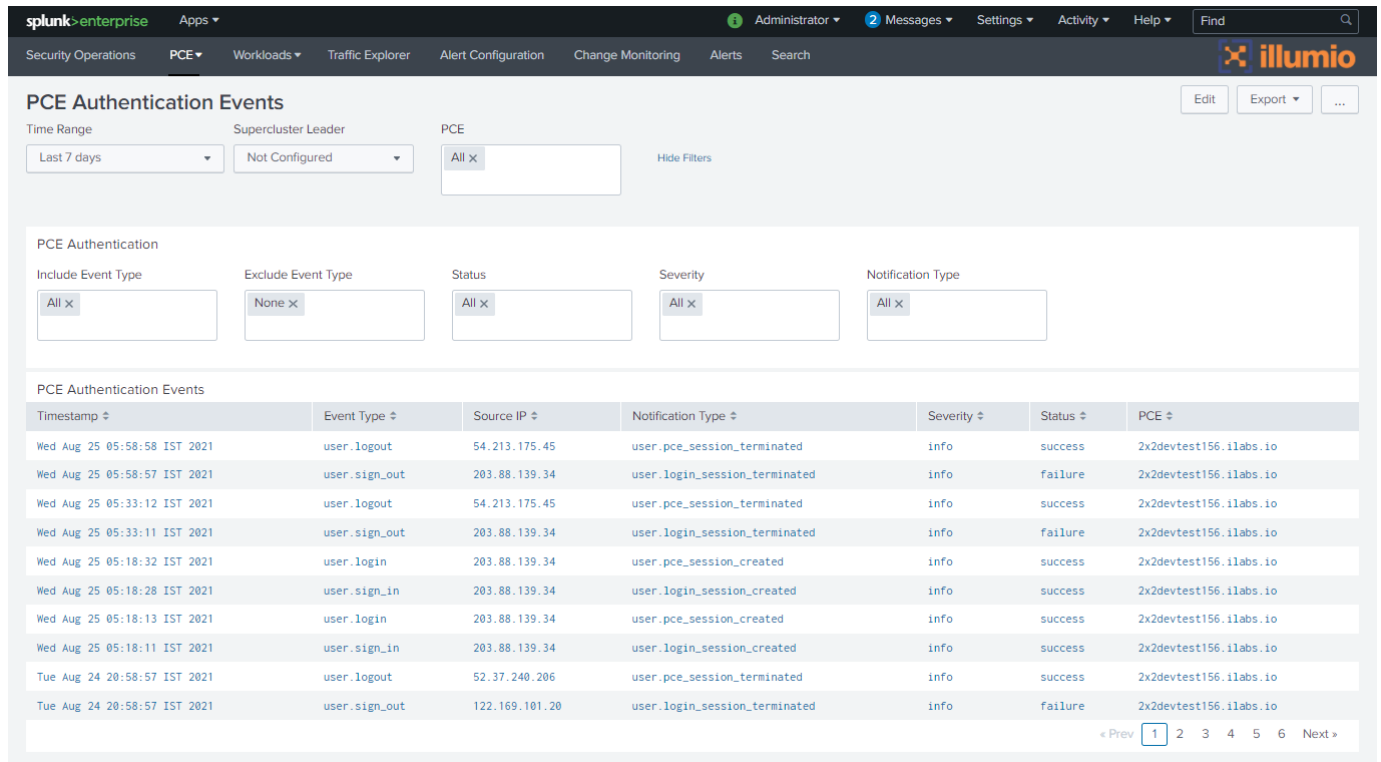
Collector Flow Rate (Average) No results found

Data Ingestion Volume in the Last 30 Days

Time	Index	Source Type	GB
2022-09-01	artfact	[[DataIngestion-artfact]]	8.213
2022-09-01	artfact	[[DataIngestion-artfact]]	14.828
2022-09-01	data	[[DataIngestion-artfact]]	8.888
2022-09-01	artfact	[[DataIngestion-artfact]]	8.888

## PCE Authentication Events Dashboard

The PCE Authentication Events dashboard enables Splunk administrators to search and filter types of user authentication data.



**PCE Authentication Events**

Time Range: Last 7 days | Supercluster Leader: Not Configured | PCE: All x

Include Event Type: All x | Exclude Event Type: None x | Status: All x | Severity: All x | Notification Type: All x

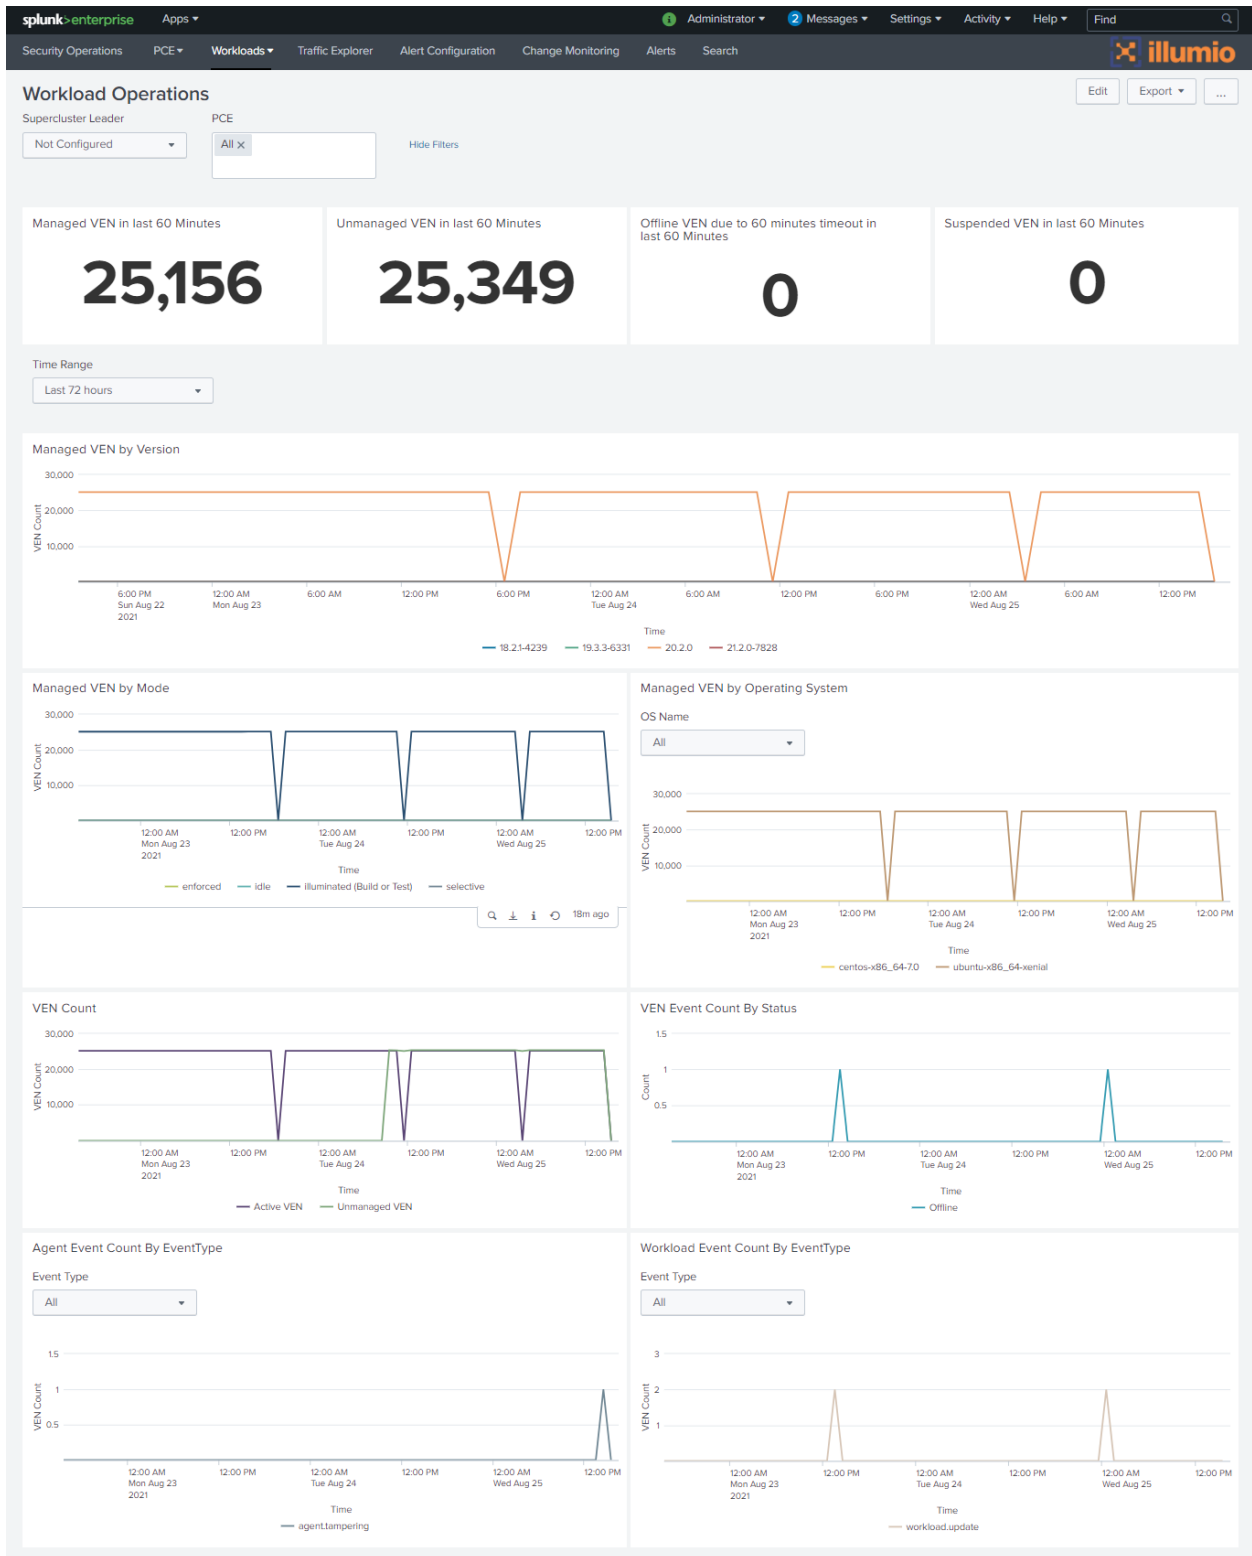
Timestamp	Event Type	Source IP	Notification Type	Severity	Status	PCE
Wed Aug 25 05:58:58 IST 2021	user.logout	54.213.175.45	user.pce_session_terminated	info	success	2x2devtest156.1labs.io
Wed Aug 25 05:58:57 IST 2021	user.sign_out	203.88.139.34	user.login_session_terminated	info	failure	2x2devtest156.1labs.io
Wed Aug 25 05:33:12 IST 2021	user.logout	54.213.175.45	user.pce_session_terminated	info	success	2x2devtest156.1labs.io
Wed Aug 25 05:33:11 IST 2021	user.sign_out	203.88.139.34	user.login_session_terminated	info	failure	2x2devtest156.1labs.io
Wed Aug 25 05:18:32 IST 2021	user.login	203.88.139.34	user.pce_session_created	info	success	2x2devtest156.1labs.io
Wed Aug 25 05:18:28 IST 2021	user.sign_in	203.88.139.34	user.login_session_created	info	success	2x2devtest156.1labs.io
Wed Aug 25 05:18:13 IST 2021	user.login	203.88.139.34	user.pce_session_created	info	success	2x2devtest156.1labs.io
Wed Aug 25 05:18:11 IST 2021	user.sign_in	203.88.139.34	user.login_session_created	info	success	2x2devtest156.1labs.io
Tue Aug 24 20:58:57 IST 2021	user.logout	52.37.240.206	user.pce_session_terminated	info	success	2x2devtest156.1labs.io
Tue Aug 24 20:58:57 IST 2021	user.sign_out	122.169.101.20	user.login_session_terminated	info	failure	2x2devtest156.1labs.io

## Workload Operations Dashboard

The Workload Operations dashboard enables Splunk administrators to monitor the Workloads managed by the PCE instances. The dashboard displays VEN deployment statistics as well as VEN-reported events. If multiple PCE instances are connected to Splunk, you can use the dropdown list at the top of the dashboard to choose which PCE to monitor.

The Workload Operations dashboard is built using data from the following sources:

- REST API calls made to the PCE
- Events



## Workload Investigation Dashboard

The Workload Investigation dashboard enables Splunk administrators to search detailed information about one or more workloads. If multiple PCE instances are connected to Splunk, you can use the dropdown list to choose which PCE to monitor. You can use the Time Range dropdown list to filter the display. Wildcards or IP addresses can be used to select multiple workloads. Instead of using hostnames or IP addresses to select workloads, you can define a workload scope using the App Label, Env Label, and Loc Label dropdown lists.

The Workload Investigation dashboard has two panels:

- Workload Details – Hostname, IP, Operating System, Status of policy, PCE
- Audit Events – Events recorded for the workloads. You can click an event in the list to drill down for more details about the event.

splunk-enterprise Apps Administrator Messages Settings Activity Help

Security Operations PCE Workloads Traffic Explorer Alert Configuration Change Monitoring Alerts Search

### Workload Investigations

Time Range:

Supercluster Leader:

PCE:

Hostname/IP:

App Label:

Env Label:

Loc Label:  Hide Filters

Please specify either hostname(s) or scope labels. If both hostname and scope labels are specified, then the filter uses AND condition.

**Active VEN**  
  

# 25,241

**Suspended VEN**  
  

# 0

**Stopped VEN**  
  

# 22

#### Policy Enforcement State

#### Policy Synchronization Status

#### Workload Details

Workload	Host Name	Interfaces	IPs	OS Name	Workload labels - RHEL	Updated at	Online	Policy State	Status	Policy Sync	Policy applied at	Log Traffic	PCE
k8s--role--env--number	k8s--role--env--number	docker0 ens192	10.0.9.155 172.17.0.1	centos-x86_64-7.0	PORT Illumio Production Amazon	Mon Aug 23 08:37:12 IST 2021	true	illuminated	active	applied	Mon Aug 23 12:07:32 IST 2021	false	2x2devtest156.11abs.io
vm-157	vm-157	docker0 ens192	10.0.9.157 172.17.0.1	centos-x86_64-7.0	- CrestDastSystem ABCD Location	Thu Aug 12 11:17:35 IST 2021	true	illuminated	active	applied	Mon Aug 23 12:07:45 IST 2021	false	2x2devtest156.11abs.io
vm-52	vm-52	docker0 ens192	10.0.9.52 172.17.0.1	centos-x86_64-7.0	Today Today Today	Thu Jul 29 13:45:02 IST 2021	true	enforced	active	applied	Mon Aug 23 12:10:57 IST 2021	true	2x2devtest151.11abs.io
vm-new-45	vm-new-45	ens192	10.0.9.45	centos-x86_64-7.0	Today Today Today	Thu Jul 29 13:23:57 IST 2021	false	enforced	stopped	syncing	Thu Jul 29 12:30:10 IST 2021	false	crest-mnc.11abs.io
vm-new-44	vm-new-44	ens192	10.0.9.44	centos-x86_64-7.0	- ABC -	Thu Jul 29 13:23:57 IST 2021	false	enforced	stopped	syncing	Thu Jul 29 12:29:13 IST 2021	false	crest-mnc.11abs.io

< Prev 1 2 3 4 5 6 7 8 9 10 Next >

#### Audit Events

Include Event Type:

Exclude Event Type:

Status:

Severity:

Notification Type:

Timestamp	Event Type	Host Name	Source IP	Notification Type	Severity	Status	PCE
Wed Aug 25 09:14:48 IST 2021	request.authentication_failed	-	122.169.101.20	request.authentication_failed	err	failure	2x2devtest79.11abs.io
Wed Aug 25 09:14:48 IST 2021	request.authentication_failed	-	122.169.101.20	request.authentication_failed	err	failure	2x2devtest79.11abs.io
Wed Aug 25 09:14:48 IST 2021	request.authentication_failed	-	122.169.101.20	request.authentication_failed	err	failure	2x2devtest79.11abs.io
Wed Aug 25 09:14:46 IST 2021	request.authentication_failed	-	122.169.101.20	request.authentication_failed	err	failure	2x2devtest79.11abs.io
Wed Aug 25 09:13:48 IST 2021	system_task.prune_old_log_events	-	52.42.94.101	system_task.event_pruning_completed	info	success	2x2devtest79.11abs.io
Wed Aug 25 09:13:48 IST 2021	system_task.prune_old_log_events	-	52.42.94.101	system_task.event_pruning_completed	info	success	2x2devtest79.11abs.io
Wed Aug 25 08:43:47 IST 2021	request.authorization_failed	-	203.88.139.34	request.authorization_failed	err	failure	2x2devtest151.11abs.io
Wed Aug 25 08:43:47 IST 2021	request.authorization_failed	-	203.88.139.34	request.authorization_failed	err	failure	2x2devtest151.11abs.io
Wed Aug 25 08:43:47 IST 2021	request.authorization_failed	-	203.88.139.34	request.authorization_failed	err	failure	2x2devtest151.11abs.io



### **Traffic Explorer Dashboard**

The Traffic Explorer Dashboard helps Splunk administrators to visualize traffic data which are coming from syslog, and enables them to search and filter traffic events.



# Illumio App for Splunk

**Traffic Explorer**

Time Range: Last 60 minutes | Supercluster Leader: Not Configured | PCE: All x | Hostname/IP: | App Label: All | Env Label: All

Loc Label: All | Hide Filters

Policy Decision: All x | Port: All x | Protocol: All

Please specify either hostname(s) or scope labels. If both hostname and scope labels are specified, then the filter uses AND condition.

**Policy Decision**

**Flows by Port**

Port	Flows
5855	436
138	4144
5841	3674
137	2821
123	124

**Flows by Policy Decision**

Policy Decision	Flows
blocked	1661
potentially-blocked	3658
allowed	868
unknown	35

Please be aware that you will need to install the Sankey Diagram App in order to see the following panel.

**Communications Map between Labeled Workloads**

Number of links to chart: All

**Traffic Events**

Timestamp	Source IP	Source Host	Source labels - RAEI	Direction	Port	Protocol	Flows	Policy Decision	Destination labels - RAEI	Destination IP	Destination Host	PCE
Wed Aug 25 14:44:55 IST 2021	10.0.8.250	perf_workload_2298	--	I	138	UDP	1	allowed	Today Today Today	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:55 IST 2021	10.0.8.198	perf_workload_2246	--	I	138	UDP	1	allowed	Today Today Today	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:55 IST 2021	10.0.8.110	perf_workload_2158	--	I	138	UDP	1	allowed	Today Today Today	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:55 IST 2021	10.0.6.238	perf_workload_1774	--	I	137	UDP	1	allowed	Today Today Today	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:55 IST 2021	10.0.6.69	perf_workload_1665	--	I	138	UDP	1	allowed	Today Today Today	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:55 IST 2021	10.0.6.69	perf_workload_1665	--	I	137	UDP	1	allowed	Today Today Today	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:36 IST 2021	10.0.1.99	perf_workload_355	--	I	138	UDP	1	blocked	Web QuarantineApp Production 30 Katharinenstr, Hamburg	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:27 IST 2021	10.0.9.194	perf_workload_2498	--	I	138	UDP	1	blocked	Web QuarantineApp Production 30 Katharinenstr, Hamburg	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:21 IST 2021	10.0.14.280	perf_workload_3784	--	I	138	UDP	1	blocked	Web QuarantineApp Production 30 Katharinenstr, Hamburg	10.0.15.255	centos7	2x2devtest79.illabs.io
Wed Aug 25 14:44:21 IST 2021	10.0.8.110	perf_workload_2158	--	I	138	UDP	1	blocked	Web QuarantineApp Production 30 Katharinenstr, Hamburg	10.0.15.255	centos7	2x2devtest79.illabs.io

**Note:** This dashboard leverages the Splunk Sankey Diagram app for visualization. You will need to install this [app](#) .



## Alert Configuration Page

See “[Configuring Alerts](#)” in the Configuration section later in this document.

## Alerts Page

Click the Alerts link to view the Splunk Alerts Page. Use this page to view all alerts for the Illumio for Splunk app. Use the links in this page, such as Edit and Open in Search, to further work with the alerts. For example, using the Edit link, you can set up email notifications for alerts. See Splunk documentation for more information about this page.

**Alerts**

Alerts set a condition that triggers an action, such as sending an email that contains the results of the triggering search to a list of people. Click the name to view the alert. Open the alert in Search to refine the parameters.

7 Alerts All Yours This App's

i	Title ^	Actions	Owner ⇅	App ⇅	Sharing ⇅	Status ⇅
>	Illumio_Check_PCE_Collector_Data	<a href="#">Open in Search</a>	nobody	IllumioAppforSplunk	App	Enabled
▼	Illumio_PCE_Health_Alert Illumio_PCE_Health Enabled: ..... Yes. Permissions: ..... Shared in App. Owned by nobody. Modified: ..... Jan 1, 1970 12:00:00 AM Alert Type: ..... Scheduled. Cron Schedule. Trigger Condition: .. Number of Results is > 0. Actions: ..... ▼ 1 Action 🔔 Add to Triggered Alerts	<a href="#">Open in Search</a>	nobody	IllumioAppforSplunk	App	Enabled
>	Illumio_Policy_Provisioning_Alert	<a href="#">Open in Search</a>	nobody	IllumioAppforSplunk	App	Enabled
>	Illumio_Rule_Update_Alert	<a href="#">Open in Search</a>	nobody	IllumioAppforSplunk	App	Enabled
>	Illumio_VEN_Inactivity_Timer_Alert	<a href="#">Open in Search</a>	nobody	IllumioAppforSplunk	App	Enabled
>	Illumio_Workload_Labeling_Alert	<a href="#">Open in Search</a>	nobody	IllumioAppforSplunk	App	Enabled

## Change Monitoring Dashboard

The Change Monitoring Dashboard helps Splunk administrators to search detailed level information of changes performed by users.



**Change Monitoring**

Time Range: Last 7 days | Supercluster Leader: Not Configured | PCE: 2x2devtest156.ilabs.io x

Daily Changes  
**8<sup>↑</sup>**<sub>8</sub>

Daily Creates  
**2<sup>↑</sup>**<sub>2</sub>

Daily Updates  
**3<sup>↑</sup>**<sub>3</sub>

Daily Deletes  
**1<sup>↑</sup>**<sub>1</sub>

Changes by Object				Changes by User			
Object Type	Count	Actual User	Count	Actual User	Count	Actual User	Count
syslog_destination	10	vatsal.halpara@crestdatasys.com	10	vatsal.halpara@crestdatasys.com	13		
workload	5	smit.belani@crestdatasys.com	12				
workload_service_report	5	kartavya.patel@crestdatasys.com	2				
label	4						
permission	2						
labels	1						

Timestamp	Actual User	Source IP	PCE	Object Type	Change Type	Resource Href	Change Details
Wed Aug 25 05:22:16 IST 2021	smit.belani@crestdatasys.com	203.88.139.34	2x2devtest156.ilabs.io	syslog_destination	update	-	["node_status_included":{"before":false,"after":true}]
Tue Aug 24 19:33:59 IST 2021	vatsal.halpara@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	labels	delete	-	null
Tue Aug 24 19:33:36 IST 2021	vatsal.halpara@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload_service_report	-	-	-
Tue Aug 24 19:33:35 IST 2021	vatsal.halpara@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload	update	/orgs/1/workloads/412e7465-986a-46cc-9b32-ac959e8487e0	["labels":{"deleted":[{"href":"/orgs/1/labels/29","key":"role","value":...}]}]
Tue Aug 24 19:33:29 IST 2021	vatsal.halpara@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	label	create	/orgs/1/labels/30	["key":{"before":null,"after":"role"},"value":{"before":null,"after":"Illumio_CREST"}]
Tue Aug 24 19:33:02 IST 2021	vatsal.halpara@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload_service_report	-	-	-
Tue Aug 24 19:33:01 IST 2021	vatsal.halpara@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload	update	/orgs/1/workloads/412e7465-986a-46cc-9b32-ac959e8487e0	["labels":{"deleted":[{"href":"/orgs/1/labels/18","key":"role","value":...}]}]
Tue Aug 24 19:32:52 IST 2021	vatsal.halpara@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	label	create	/orgs/1/labels/29	["key":{"before":null,"after":"role"},"value":{"before":null,"after":"Supercluster"}]
Mon Aug 23 08:37:13 IST 2021	smit.belani@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload_service_report	-	-	-
Mon Aug 23 08:37:12 IST 2021	smit.belani@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload	update	/orgs/1/workloads/8e00f52-2986-4308-9268-9962f519c4ad	["enforcement_mode":{"before":"idle","after":"visible"}]
Mon Aug 23 08:35:35 IST 2021	smit.belani@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	label	create	/orgs/1/labels/28	["key":{"before":null,"after":"role"},"value":{"before":null,"after":"Supercluster"}]
Mon Aug 23 08:34:52 IST 2021	smit.belani@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload_service_report	-	-	-
Mon Aug 23 08:34:51 IST 2021	smit.belani@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	workload	update	/orgs/1/workloads/8e00f52-2986-4308-9268-9962f519c4ad	["labels":{"deleted":[{"href":"/orgs/1/labels/17","key":"app","value":...}]}]
Mon Aug 23 08:33:13 IST 2021	smit.belani@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	syslog_destination	create	-	["pce_scope":{"before":null,"after":["2x2devtest156.ilabs.io"]},"before":null,"after":{"RemoteSyslogDestination"},"address":{"before":null,"after":"203.88.139.43"},"port":{"before":null,"after":9182},"protocol":{"before":null,"after":17},"tls_ca_bundle":{"before":null,"after":"FILTERED"},"tls_verify_certificate":{"before":null,"after":true},"node_status_included":{"before":null,"after":false}]
Mon Aug 23 07:38:29 IST 2021	kartavya.patel@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	permission	delete	/orgs/1/permissions/1c31c86b-8435-4f21-b6ce-74e6a1d36072	null
Mon Aug 23 07:38:29 IST 2021	kartavya.patel@crestdatasys.com	122.169.101.20	2x2devtest156.ilabs.io	permission	create	/orgs/1/permissions/826af3fa-0a68-4252-8e1b-27e93ef052ee	["role":{"before":null,"after":{"href":"/orgs/1/roles/owner"}},"auth_security_privilege":{"before":null,"after":{"href":"/orgs/1/auth_security/3bca-4cc6-b288-a87d29f1ee9"}}]
Mon Aug 23 07:23:37 IST 2021	vatsal.halpara@crestdatasys.com	203.88.139.34 203.88.139.34	2x2devtest156.ilabs.io	syslog_destination	-	-	-
Mon Aug 23 07:22:49 IST 2021	vatsal.halpara@crestdatasys.com	203.88.139.34 203.88.139.34	2x2devtest156.ilabs.io	syslog_destination	-	-	-
Mon Aug 23 07:22:37 IST 2021	vatsal.halpara@crestdatasys.com	203.88.139.34 203.88.139.34	2x2devtest156.ilabs.io	syslog_destination	-	-	-
Mon Aug 23 07:10:18 IST 2021	smit.belani@crestdatasys.com	122.169.101.20 122.169.101.20	2x2devtest156.ilabs.io	syslog_destination	-	-	-

1 2 Next

## Installation

This section tells how to install the Illumio App for Splunk and TA-Illumio.

### Installation Prerequisites

- Environment variable SPLUNK\_HOME set to the Splunk directory.
- Splunk Enterprise 7.3.x, 8.0.x, 8.1.x, or 8.2.x.
- Illumio PCE installed. For compatible PCE versions, see “[Compatibility Matrix](#).”

### Splunk Single Server Deployment

In a single server deployment, a single instance of Splunk Enterprise works as a data collection node, indexer, and search head. In such scenarios, install both TA-Illumio and Illumio App for Splunk applications on this node. Then complete the setup of TA-Illumio to start data collection.

### Splunk Distributed Deployment

In a distributed deployment, install Splunk Enterprise on at least on two instances. One node works as the search head, and the other node works as the indexer and data collection node. In a Splunk distributed deployment, the data collection node and indexer are deployed on separate servers. In this environment, install the Illumio App for Splunk application on each search head node and TA-Illumio on each indexer/forwarder and search head node.

### Installing TA-Illumio

This section tells how to install the TA-Illumio add-on.

### Application of TA-Illumio to Splunk Components

This section describes how TA-Illumio is applied to various Splunk components.

**Splunk Heavy Forwarder:** On the heavy forwarder, which is a Splunk Enterprise instance, TA-Illumio is used for data collection. TA-Illumio is required because the Illumio App for Splunk depends on both API and syslog data from Illumio. TA-Illumio provides both.

To make TA-Illumio data collection work, you must configure Data Input (modular input) as described in the “[Installation](#)” section of this document.

Depending on the Splunk deployment, the heavy forwarder might not be a separate component. It can be deployed on the same node as the indexer or search head.

**Splunk Indexer:** TA-Illumio has a special purpose on the indexer. The PCE might send invalid JSON data which does not need to be indexed. TA-Illumio filters out invalid JSON events. If invalid JSON events are not a concern, TA-Illumio does not need to be installed on the indexer. On the Splunk indexer, you can manually create the index in which the data is stored.

**Splunk Search Head:** TA-Illumio is used with the Splunk search head to extract time fields, which are then used by the Illumio App for Splunk in dashboard visualizations.

## Installing Illumio App for Splunk in a Distributed Splunk Environment

The following table describes the apps to deploy when installing within a Splunk distributed environment.

App Name	Search Head	Indexer	Heavy Forwarder/Data Collection Node
Data Input (aka Modular Input or REST Modular Input)	Configure data input with API keys and data collection disabled (not checked)	Configure data input with API keys and data collection disabled (not checked)	Configure data input with API Keys and data collection enabled
Illumio App for Splunk	Yes	Not Applicable	Not Applicable

Illumio Technology Add-on (Illumio-TA) for Splunk	Yes	Optional (if filtering of invalid JSON is desired)	Yes
---	-----	--	-----

The deployment procedure varies depending on whether you are using Heavy Forwarder or Splunk Universal Forwarder.

### Using Splunk Heavy Forwarder

In a distributed environment with Splunk Heavy Forwarder:

- On the search head, install the Illumio App for Splunk and the TA-Illumio add-on.
- On the Splunk Heavy Forwarder, install TA-Illumio.

### Using Splunk Universal Forwarder

In a distributed environment with Splunk Universal Forwarder:

- Set up a data collection node with Splunk Universal Forwarder.
- Configure the PCE to forward data from all nodes to the Splunk Universal Forwarder.
- Configure the Splunk Universal Forwarder to send the data to Splunk Indexer or Splunk Heavy Forwarder.

The installation steps are as follows.

1. Configure the Splunk Universal Forwarder to collect data from the Illumio PCE.
  - a. Create a TCP stanza in `$(SPLUNK_HOME)/etc/system/local/inputs.conf` file:

```
[tcp://<PORT>]
index=<INDEX-NAME>
sourcetype=illumio:pce
```

- b. Configure the Splunk Universal Forwarder to send the data to the Splunk Indexer. Execute this command on the Splunk Universal Forwarder (for `<IP>:<PORT>`, fill in the Splunk Indexer IP and Listening Port):

```
$(SPLUNK_HOME)/bin/splunk add forward-server <IP>:<PORT>
```

2. Configure the Splunk Indexer to receive data from SUF. Create the following stanza in the file `$SPLUNK_HOME/etc/system/local/inputs.conf`:

```
[splunktcp://<PORT>]
```

In a distributed environment:

- If you have a separate data collection node, be sure it is running a full Splunk Enterprise version.
- Complete the Data Input configuration on the data collection node (Heavy Forwarder) with API keys and data collection enabled.
- On all other nodes, configure the data input with the API keys and data collection disabled.
- In setups where a non-default index is used, it may be necessary to configure the search macro `illumio_get_index` with a definition of “`index=illumio`”. Use the steps in [“Splunk Index, Source, and Source Types.”](#)

## Splunk Cloud Instance Deployment

In the Splunk Cloud, data indexing takes place in a cloud instance. The data collection can take place in an on-premise Splunk instance in your environment that will work as heavy forwarder.

## Installation Commands

The Illumio App for Splunk and TA-Illumio can be installed either through a command line or from the Splunk UI.

Use these commands for a fresh installation. If you are upgrading from a previous version, see [“Upgrade the App.”](#)

- To install from the UI, log in to Splunk. Go to **App > Manage Apps** and click **Install app from a file**. Then choose the SPL file to install and click **Upload the SPL**.
- To install from the command window, go to the folder `$SPLUNK_HOME/bin` and execute the following command. For the XXs, substitute the rest of the actual file name.

```
./splunk install app TA-Illumio-XX-XXXX-XX.spl  
./splunk install app IllumioAppForSplunk-XX-XXXX-XX.spl
```

## Configuration

This section describes how to configure the following after installing the Illumio App for Splunk:

- The Splunk app itself
- Your on-premise or cloud PCE
- Alerts

You may also view video instructions at <https://support.illumio.com/knowledge-base/articles/How-to-Configure-Illumio-to-Send-Event-logs-to-Splunk.html>

## Splunk App Configuration

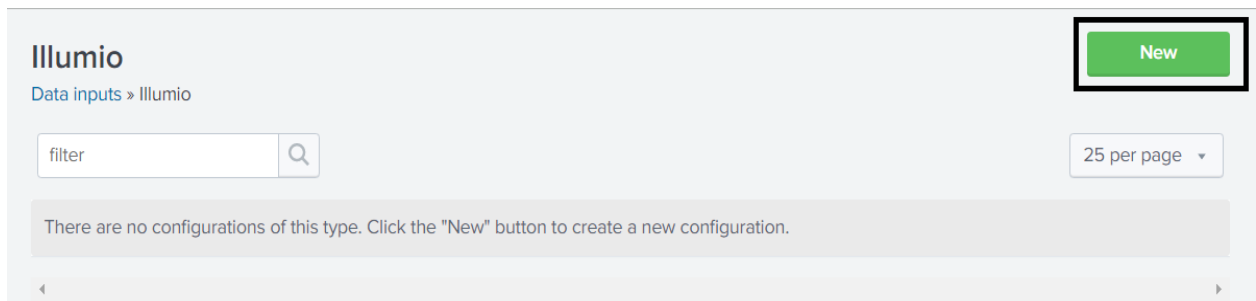
Once the installation of TA-Illumio has completed successfully, follow the steps below to configure Splunk to receive the data from the Illumio PCE syslog and to get workload and label information indexed into the Splunk App using the Illumio ASP REST API.

1. Log in to the Splunk web app.
2. Go to **Settings > Data inputs**.
3. Locate "Illumio" and click it.

<p><b>UDP</b></p> <p>Listen on a UDP port for incoming data, e.g. syslog.</p>	0	+ Add new
<p><b>Scripts</b></p> <p>Run custom scripts to collect or generate more data.</p>	5	+ Add new
<p><b>Input</b></p> <p>Go to the add-on's configuration UI and configure modular inputs under the Inputs menu.</p>	0	+ Add new
<p><b>Illumio</b></p> <p>Enable data inputs for splunk add-on for Illumio</p>	0	+ Add new

4. Click **New** to create Data Input (Modular Input) for ingesting data from Illumio PCE to Illumio App for Splunk.

**NOTE:** If you have multiple PCEs sending data to a single Splunk, then a different Data Input with different TCP ports is needed for each PCE.



5. In the Modular Input page, enter the configuration details using the following table:



Input Parameter	Mandatory or Optional	Description
Name	Mandatory	Identifying name of the Illumio PCE.
Supercluster Leader/ PCE URL	Mandatory	<p>Enter the PCE URL including HTTPS and the port number.(Make sure that if the provided PCE is part of the supercluster then it must be the leader of the supercluster)</p> <p>For example: https://illumiofce.com:443/</p>
API Authentication Username	Mandatory	<p>API Authentication Username used to authenticate with the Illumio PCE. To generate the API key, log in to the Illumio PCE web console, then click <b>Username &gt; My API Keys &gt; Add New.</b></p> <p>For example: api_16175f6af766fcd7b</p> <p>If API Username and Secret are not specified, the PCE Operations and Workload Operations dashboards will not work.</p>
API Authentication Secret	Mandatory	<p>The API Secret is the password for an API Key that is used to authenticate with the PCE. The API Secret is generated by the API key.</p> <p>For example: 4ed8ff8a5c40201dc52c89a59936f7b1003b950e0027204b2aaaa633ba040d22</p>

Input Parameter	Mandatory or Optional	Description
TCP Port Number for incoming syslog from PCE	Optional	Splunk server port on which the Splunk App should listen for syslog messages from the PCE. The PCE in turn should be configured to forward syslog to this port on the Splunk server. If creating multiple Data Inputs, use a different TCP port for each PCE.  For example: 5014
Port Scan configuration: Scan interval in seconds	Mandatory	Minimum time duration of connections between two workloads to determine a port scan.  For example: If two workloads show flows between 10 unique ports within 60 seconds, then a port scan is registered.  Default: 60 seconds.
Port Scan Configuration: Unique ports threshold	Mandatory	Minimum threshold of unique ports between two workloads to determine a port scan.  For example: If two workloads show flows between 10 unique ports within 60 seconds, then a port scan is registered.  Default: 10 ports.
Labels to quarantine Workloads	Optional	Comma-separated list of labels of type App, Environment, and Location. In the comma-separated list, whitespace is not allowed. Labels must be supplied in the exact order Application,Environment,Location. These labels should exist on the PCE with

Input Parameter	Mandatory or Optional	Description
		<p>appropriate policy to quarantine workloads. These labels will be applied while quarantining the workloads using the App or AR action.</p>
Organization ID	Optional	<p>For Illumio Data Center (on-premise) customers, the organization ID is 1.</p> <p>For Illumio Cloud customers, the organization ID can vary. The administrator can determine the organization ID from the Illumio PCE Web Console by logging in, clicking the administrator’s name at the top right, and clicking <b>My API keys &gt; Add New</b>. The Create New API dialog shows the Organization ID.</p>
IP addresses of PCE Nodes	Optional	<p>Comma-separated IP addresses (private, public) of all the nodes managed by this PCE instance. All IP addresses must be provided. Use only commas. Do not add space characters.</p>
Data Collection	Mandatory	<p>When enabled, the TA will collect data on this instance. If using a Splunk Cluster, this should be enabled on indexer node but disabled on search head nodes.</p> <p>Default: enabled.</p> <p>Note: When invoking “Quarantine Workload” with Splunk Cluster, the search head node TA needs to be configured with data</p>

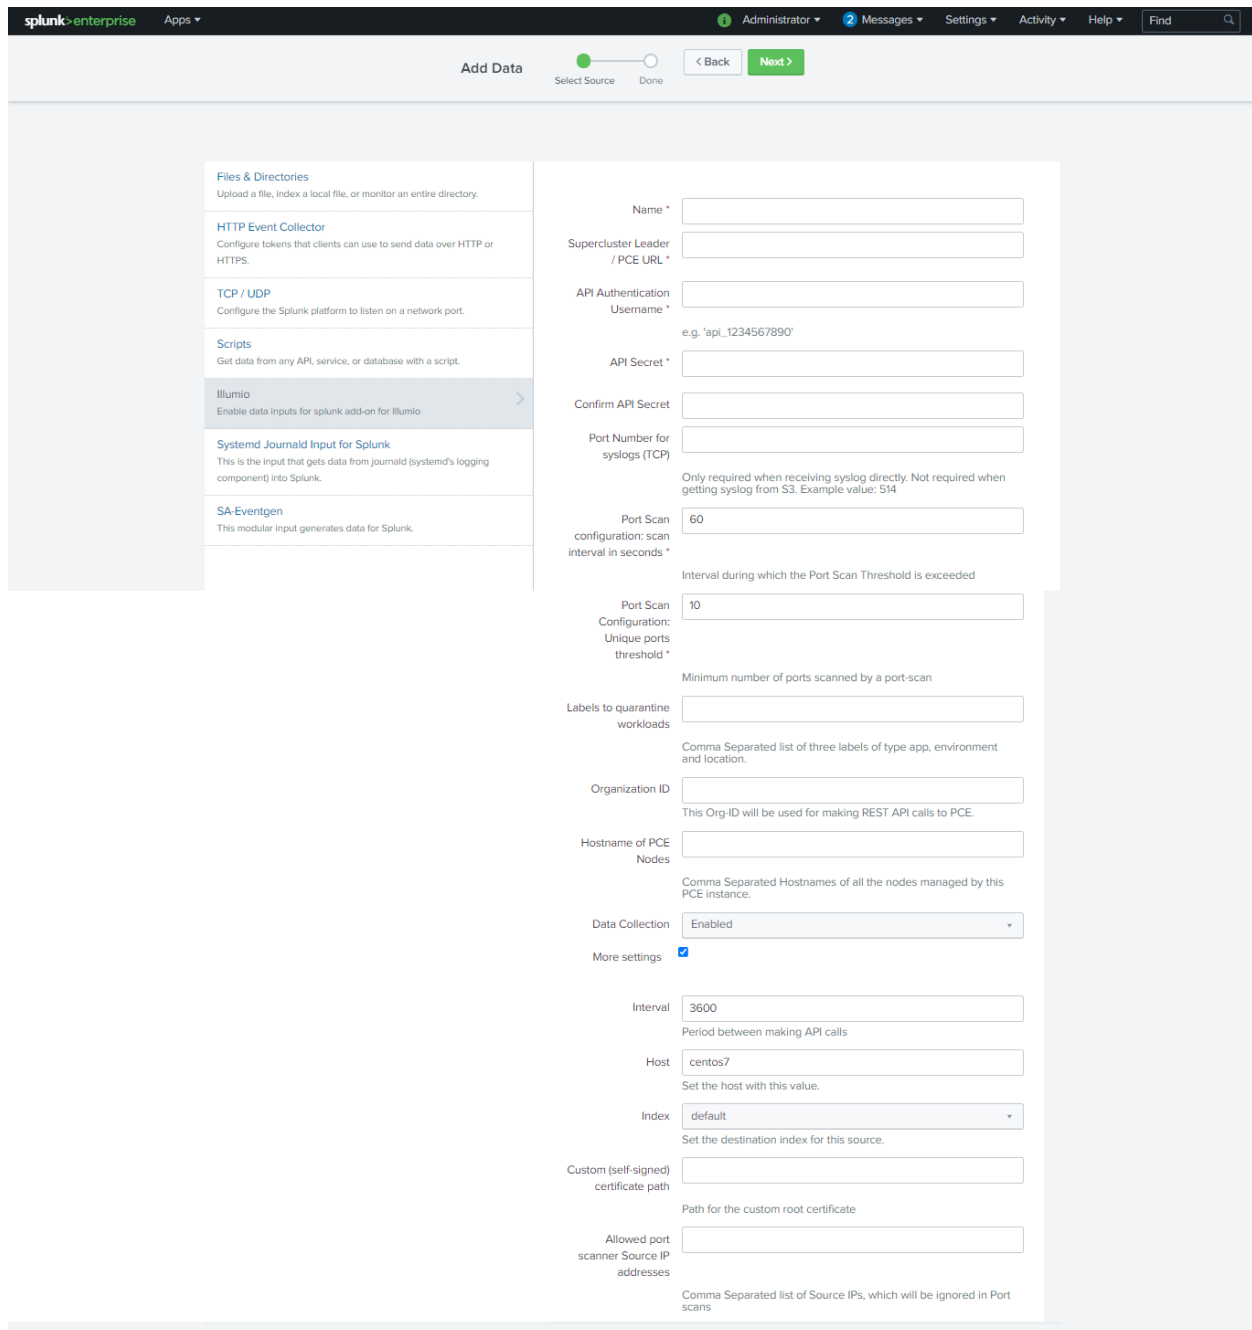
Input Parameter	Mandatory or Optional	Description
		collection disabled (and indexer node TA data collection enabled).

6. Make any desired additional settings from the following table:

Input Parameter	Mandatory or Optional	Description
Interval	Optional	Interval between REST API calls made by the Splunk App to refresh data from PCE. The minimum value is 3600 seconds.  Default: 3600.
Host	Optional	Host information added into events to be indexed by Splunk. Illumio recommends using the FQDN of the Splunk server.
Index	Optional	For use by advanced Splunk users. Change the index name under which received events are categorized. If you use a non-default (custom) index such as “Illumio”, the index should be created manually, and the search macros modified to return “index=illumio”. See <a href="#">“Splunk Index, Source, and Source Types.”</a>
Custom (Self-Signed or Local CA Authority) Certificate Path	Optional	When a local CA Authority issued SSL certificate or a self-signed SSL certificate is used with the PCE, its SSL Certificate needs to be uploaded onto the Splunk Server and the full path to the directory

Input Parameter	Mandatory or Optional	Description
		<p>containing the certificate should be provided here.</p> <p>For correct SSL operation, the Splunk server must be able to fully trust the PCE's certificate. If using a local CA Authority or a certificate issued by a secondary CA, the Splunk server CA trust chain must be updated to verify the certificate presented by the PCE. For example, on Linux, use the update-ca-trust tool.</p>
Allowed port scanner IP addresses	Optional	White-list IP addresses of known port scanners, such as Qualys hosts. These addresses are excluded when determining port scans, which avoids false positives in the Port Scans panels.

For example:



The screenshot shows the 'Add Data' configuration page for the Illumio app in Splunk Enterprise. The page is titled 'Add Data' and has a progress indicator showing 'Select Source' is complete and 'Done' is pending. The 'Next >' button is highlighted in green. The sidebar on the left lists several data sources, with 'Illumio' selected and highlighted in grey. The main content area contains a form with the following fields and options:

- Name \***: Text input field.
- Supercluster Leader / PCE URL \***: Text input field.
- API Authentication Username \***: Text input field with example value 'e.g. 'api\_1234567890''.
- API Secret \***: Text input field.
- Confirm API Secret**: Text input field.
- Port Number for syslogs (TCP)**: Text input field.
- Port Scan configuration: scan interval in seconds \***: Text input field with value '60'.
- Interval during which the Port Scan Threshold is exceeded**: Text input field with value '10'.
- Port Scan Configuration: Unique ports threshold \***: Text input field.
- Labels to quarantine workloads**: Text input field.
- Organization ID**: Text input field.
- Hostname of PCE Nodes**: Text input field.
- Data Collection**: Dropdown menu set to 'Enabled'.
- More settings**: Checked checkbox.
- Interval**: Text input field with value '3600'.
- Host**: Text input field with value 'centos7'.
- Index**: Dropdown menu set to 'default'.
- Custom (self-signed) certificate path**: Text input field.
- Allowed port scanner Source IP addresses**: Text input field.

7. Click **Next** after adding each value for data input (modular input).
8. Look for a success message displayed as a header in the setup page. This indicates correct credentials and validation passed. For incorrect credentials or errors in validation, a failure message is displayed. For more information, see [“Troubleshooting.”](#)

## On-Premise PCE Configuration

You must make configuration changes on the PCE so that data is forwarded to the Splunk server.

### Syslog Configuration

Follow the steps in "Configure Events Forwarding to External Syslog Server" in the *Illumio ASP Web Console User Guide*.

### PCE Runtime Configuration

(For PCE versions before 18.2.1. If you have PCE 18.2.1 or later, skip this section.)

To generate and send traffic flow summaries to the PCE syslog and forward to Splunk, you need to make the following changes to the PCE Runtime Environment file `runtime_env.yml`. Changes must be made to the `runtime_env.yml` file on all PCE nodes in the cluster, and the PCE must be restarted to put the changes into effect.

### PCE `runtime_env.yml` Configuration

```
export_flow_summaries_to_syslog:
```

- accepted
- potentially\_blocked
- blocked

For more information about `runtime_env.yml` and the setting `export_flow_summaries_to_syslog`, see the *Illumio ASP PCE Deployment Guide*.

## Illumio Cloud PCE Configuration

If you are using Illumio Secure Cloud, follow these steps:

1. Perform the configuration steps from "[On-Premise PCE Configuration](#)".

2. In addition, two components are needed so your PCE data can be relayed to the Illumio App for Splunk:
  - Amazon S3 bucket, which permits reliably storing events from Illumio Cloud.
  - Splunk Add-on for AWS, which permits reading events from an Amazon S3 bucket.

Illumio PCE Secure Cloud logs all traffic flows, including allowed traffic, blocked traffic, potentially blocked traffic, and auditable events to your Amazon S3 bucket. You may choose to disable specific types of events in Illumio Cloud by filing a support ticket. The Splunk Add-on for AWS reads the data from Amazon S3, enriches the data with source type, and enables data to be processed by TA-Illumio. The data is then visualized with the Illumio App for Splunk.

Starting with the Illumio App for Splunk 2.3.0, the consumption of data from S3 is more robust than in earlier versions.

### **Configure Amazon S3 Bucket**

To implement the Illumio App for Splunk with your Illumio PCE Secure Cloud, you must provide an AWS S3 bucket. The S3 bucket can be created and configured using an Illumio-provided CloudFormation template, which you can get from [Flow Logs for Illumio Secure Cloud PCE](#) in the Illumio Knowledge Base.



Load the template into CloudFormation as follows:

1. Select a template.

### Select Template

Select the template that describes the stack that you want to create. A stack is a group of related resources that you manage as a single unit.

**Design a template** Use AWS CloudFormation Designer to create or modify an existing template. [Learn more.](#)

**Choose a template** A template is a JSON/YAML-formatted text file that describes your stack's resources and their properties. [Learn more.](#)

Select a sample template

Upload a template to Amazon S3  
 illumio-flow-logs.json

Specify an Amazon S3 template URL

2. Specify the details. The value of **Stack name** can be any label you like; it is for convenience and display only.

### Specify Details

Specify a stack name and parameter values. You can use or change the default parameter values, which

**Stack name**

### Parameters

**Bucketname**

**Externalid**

3. Specify options.

### Options

#### Tags

You can specify tags (key-value pairs) for resources in your stack. You can add up to 50 unique key-value pairs for each stack. [Learn more.](#)

Key (127 characters maximum)	Value (255 characters maximum)	
1 <input type="text"/>	<input type="text"/>	<input type="button" value="+"/>

#### Permissions

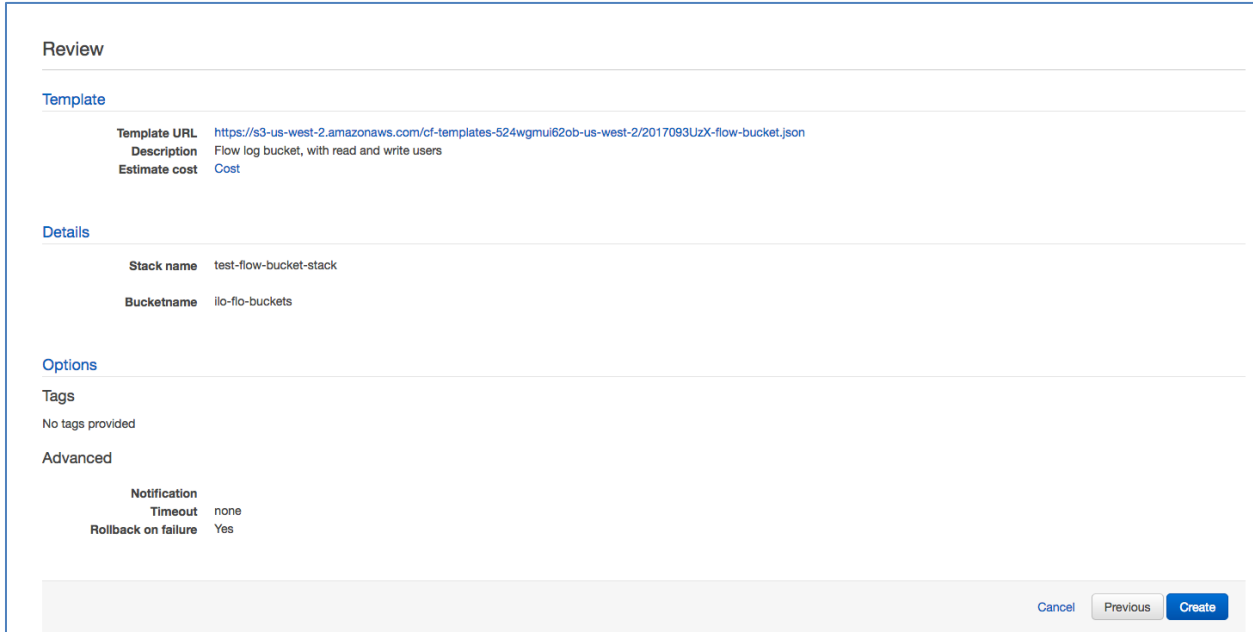
You can choose an IAM role that CloudFormation uses to create, modify, or delete resources in the stack. If you don't choose a role, CloudFormation uses the permissions defined in your account. [Learn more.](#)

**IAM Role**    
Enter role arn

▸ **Advanced**

You can set additional options for your stack, like notification options and a stack policy. [Learn more.](#)

4. Review the configuration.
5. Click Create.



The screenshot shows the 'Review' page for an AWS CloudFormation stack. The page is divided into several sections:

- Review**: The main title of the page.
- Template**: A section containing:
  - Template URL**: <https://s3-us-west-2.amazonaws.com/cf-templates-524wgmui62ob-us-west-2/2017093UzX-flow-bucket.json>
  - Description**: Flow log bucket, with read and write users
  - Estimate cost**: Cost
- Details**: A section containing:
  - Stack name**: test-flow-bucket-stack
  - Bucketname**: ilo-fio-buckets
- Options**: A section containing:
  - Tags**: No tags provided
  - Advanced**:
    - Notification Timeout**: none
    - Rollback on failure**: Yes

At the bottom right of the page, there are three buttons: 'Cancel', 'Previous', and 'Create'.

The above example contains the Illumio AWS account ID. “Externalid” is an extra password to ensure that root access to the Illumio production account is not enough to access your S3 bucket, to prevent a poorly functioning third-party service. For more information, see [How to Use External ID When Granting Access to Your AWS Resources](#) in the Amazon Blog.

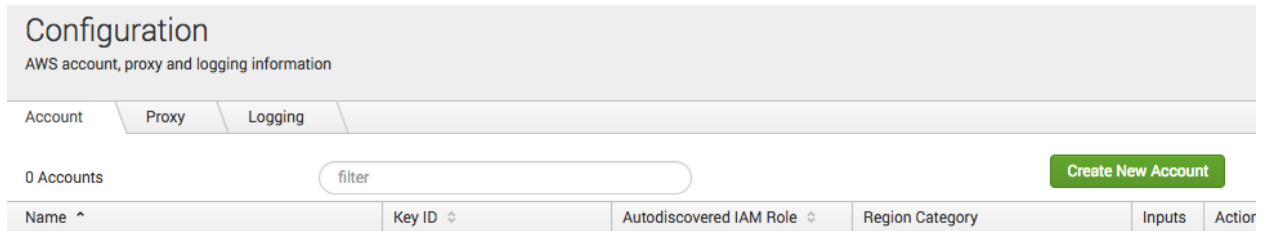
Provide the following information to Illumio:

- The AWS S3 bucket name that you have chosen.
- Your AWS account ID. This is available under "My Account", or <https://console.aws.amazon.com/billing/home?#/account> in the AWS console.

## Configure Splunk Add-on for AWS

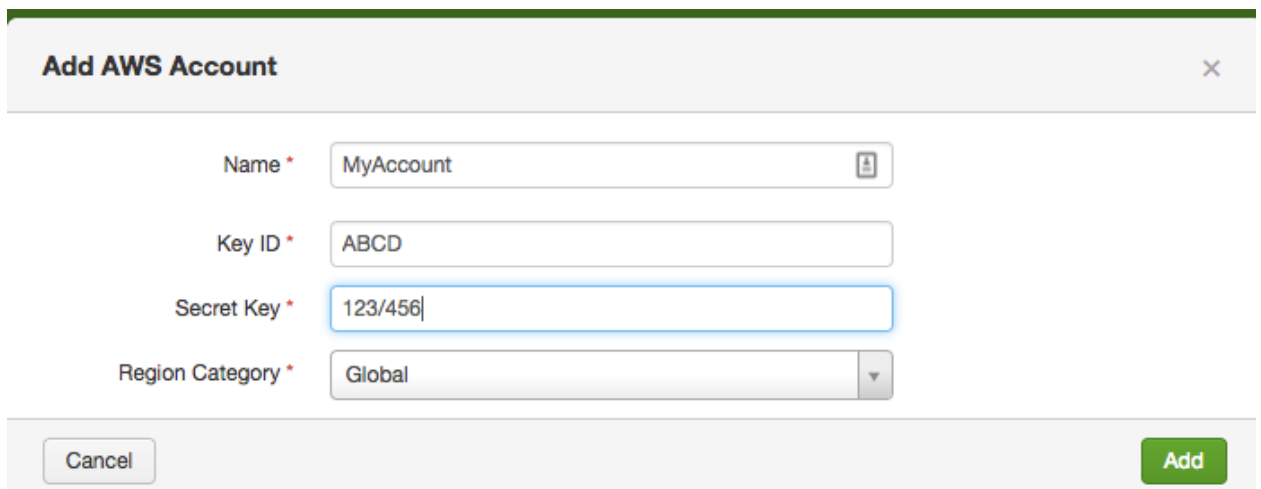
Install the [Splunk Add-On for AWS](#). This is not to be confused with the Splunk App for AWS, which is an entirely different app.

1. Enter your account into the Splunk Add-On for AWS App:



The screenshot shows the 'Configuration' page for the Splunk Add-On for AWS. It has tabs for 'Account', 'Proxy', and 'Logging'. Below the tabs, it shows '0 Accounts' and a search filter. A table lists account details with columns: Name, Key ID, Autodiscovered IAM Role, Region Category, Inputs, and Action. A 'Create New Account' button is visible in the top right.

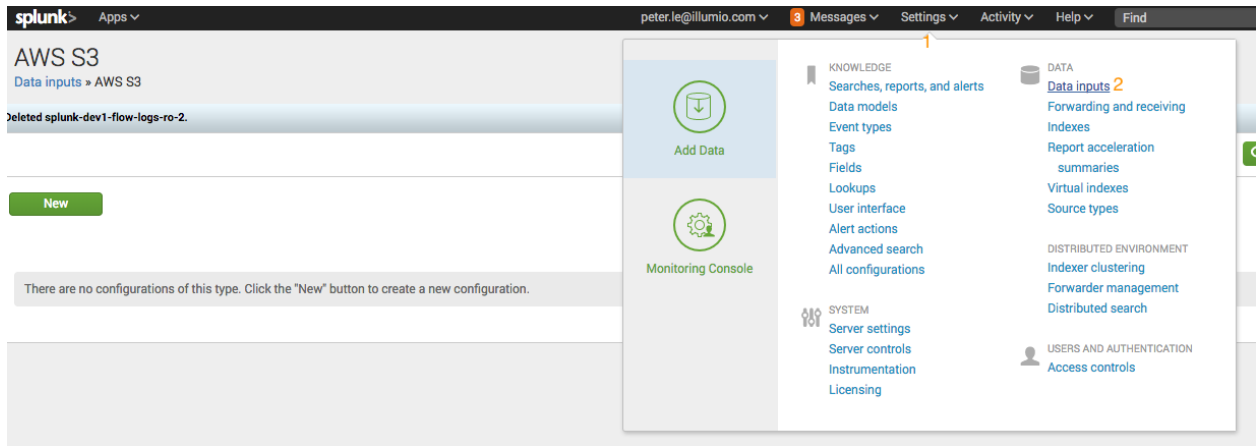
2. Enter Name, KeyID, Secret Key, and select the "global" Region Category:



The screenshot shows the 'Add AWS Account' dialog box. It contains four input fields: 'Name' (MyAccount), 'Key ID' (ABCD), 'Secret Key' (123/456), and 'Region Category' (Global). There are 'Cancel' and 'Add' buttons at the bottom.

3. Be sure to create an IAM S3 Bucket Policy that allows for Splunk to access the S3 Bucket. For instructions, see ["Configure S3 permissions"](#) in the Splunk documentation.


4. Choose **Settings > Data inputs**:



Next, you will create two data inputs for AWS S3:

- One data input for events, with source type set to `Illumio:pce`.
- One data input for traffic flow summaries, with source type set to `Illumio:pce:collector`.

5. Find AWS S3, then click **Add New**:



6. Add configuration data for Events. Fill in the following:

- Name
- AWS Account
- S3 Bucket name
- Polling interval (1800 seconds, optional)
- Key prefix: `Illumio/auditable_events`

You can accept the defaults for everything else on this screen.

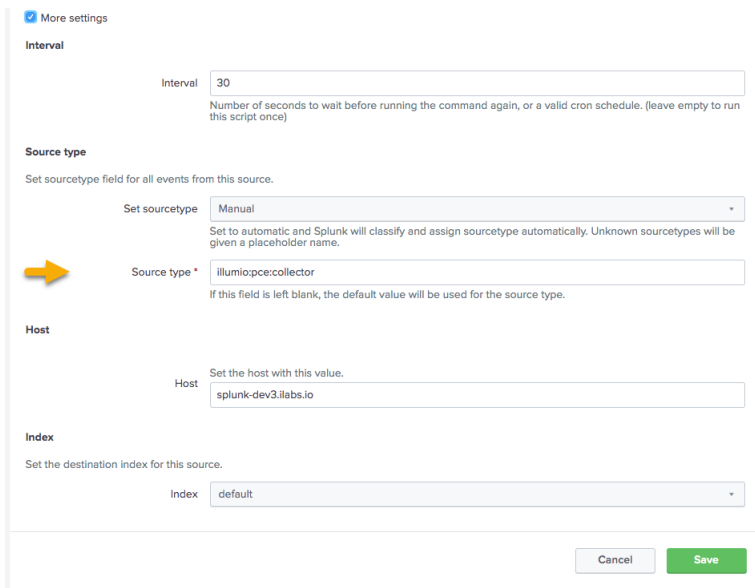
Select Source   
  Done   

Collect and index log files stored in AWS S3.

	Unique data input name	<input style="width: 95%;" type="text" value="Illumio Auditable Events"/>
	Secure S3 connection	<input type="text" value="True"/>
	S3 host name	<small>For example: s3-ap-south-east-1.amazonaws.com</small> <input type="text" value="s3.amazonaws.com"/>
	AWS Account *	<input type="text" value="corp-aws-splunk"/>
	Bucket Name *	<input type="text" value="audit-data"/>
	Polling interval	<input type="text" value="1800"/>
	Key prefix	<input type="text" value="illumio/auditable_events/"/>
	For folder keys	<input type="text" value="-1"/>
	Start datetime	<small>Only S3 keys which have been modified after this datetime will be considered</small> <input type="text" value="default"/>
	End datetime	<small>Only S3 keys which have been modified before this datetime will be considered</small> <input type="text"/>
	Max trackable items	<input type="text" value="100000"/>
	Max number of retry attempts to stream incomplete items	<input type="text" value="3"/>
	Whitelist Regex	<small>S3 key names which match this regex will be indexed</small> <input type="text"/>
	Blacklist Regex	<small>S3 key names which match this regex will be ignored, but whitelist dominates</small> <input type="text"/>
	The encoding used in your S3 files	<input type="text" value="auto"/>
	Blacklist for CloudTrail Describe events	<small>Only valid when manually set sourcetype=aws:cloudtrail. PCRE regex for specifying event names to be excluded. Leave blank to use the default set of read-only event names</small> <input type="text" value="^\$"/>
	index for the excluded CloudTrail events	<input type="text"/>
	Assume Role	<input type="text"/>
	More settings	<input type="checkbox"/>

7. Click the More Settings checkbox. Fill in the following:

- Source type (illumio:pce)
- You can accept the defaults for everything else.



More settings

**Interval**

Interval

Number of seconds to wait before running the command again, or a valid cron schedule. (leave empty to run this script once)

**Source type**

Set sourcetype field for all events from this source.

Set sourcetype

Set to automatic and Splunk will classify and assign sourcetype automatically. Unknown sourcetypes will be given a placeholder name.

Source type \*

If this field is left blank, the default value will be used for the source type.

**Host**

Set the host with this value.

Host

**Index**

Set the destination index for this source.

Index

8. Click **Save**.

9. Add configuration data for traffic flow summaries. Fill in the following:

- Name
- AWS Account
- S3 Bucket name
- Polling interval (1800 seconds, optional)
- Key prefix: Illumio/summaries
- Source type (illumio:pce:collector) – click More Settings to access this field

You can accept the defaults for everything else.

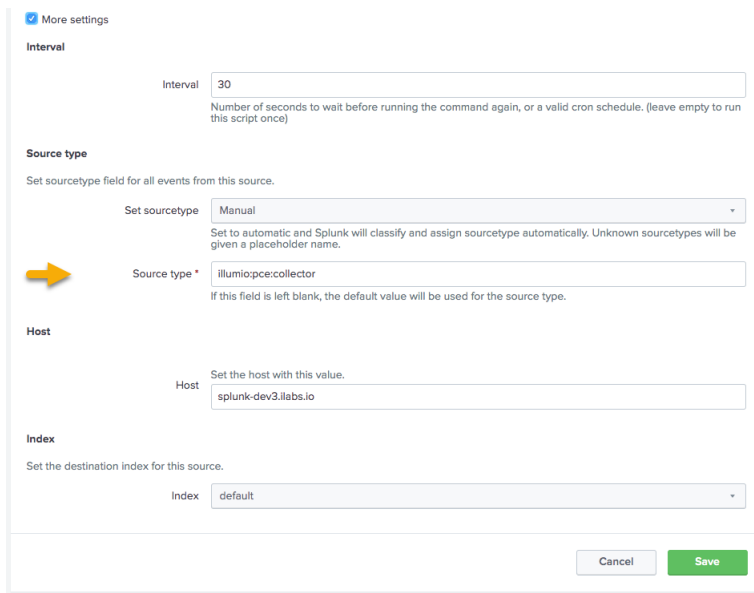
Select Source
  Done

Collect and index log files stored in AWS S3.

Name *	Unique data input name <input type="text" value="Traffic Flows"/>
Secure S3 connection	<input type="text" value="True"/>
S3 host name	For example: s3-ap-south-east-1.amazonaws.com <input type="text" value="s3.amazonaws.com"/>
➔ AWS Account *	<input type="text" value="corp-splunk"/>
➔ Bucket Name *	<input style="border: 2px solid #0070C0;" type="text" value="traffiq"/>
Polling interval	<input type="text" value="1800"/>
➔ Key prefix	<input type="text" value="illumio/summaries"/>
For folder keys	<input type="text" value="-1"/>
Start datetime	Only S3 keys which have been modified after this datetime will be considered <input type="text" value="default"/>
End datetime	Only S3 keys which have been modified before this datetime will be considered <input type="text"/>
Max trackable items	<input type="text" value="100000"/>
Max number of retry attempts to stream incomplete items	<input type="text" value="3"/>
Whitelist Regex	S3 key names which match this regex will be indexed <input type="text"/>
Blacklist Regex	S3 key names which match this regex will be ignored, but whitelist dominates <input type="text"/>
The encoding used in your S3 files	<input type="text" value="auto"/>
Blacklist for CloudTrail Describe events	Only valid when manually set sourcetype=aws:cloudtrail. PCRE regex for specifying event names to be excluded. Leave blank to use the default set of read-only event names <input type="text" value="^\$"/>
index for the excluded CloudTrail events	<input type="text"/>
Assume Role	<input type="text"/>
More settings	<input type="checkbox"/>



After clicking More settings, fill in the source type:



More settings

**Interval**

Interval: 30  
Number of seconds to wait before running the command again, or a valid cron schedule. (leave empty to run this script once)

**Source type**

Set sourcetype field for all events from this source.

Set sourcetype: Manual

Source type: **illumio:pce:collector**  
If this field is left blank, the default value will be used for the source type.

**Host**

Host: splunk-dev3.illabs.io

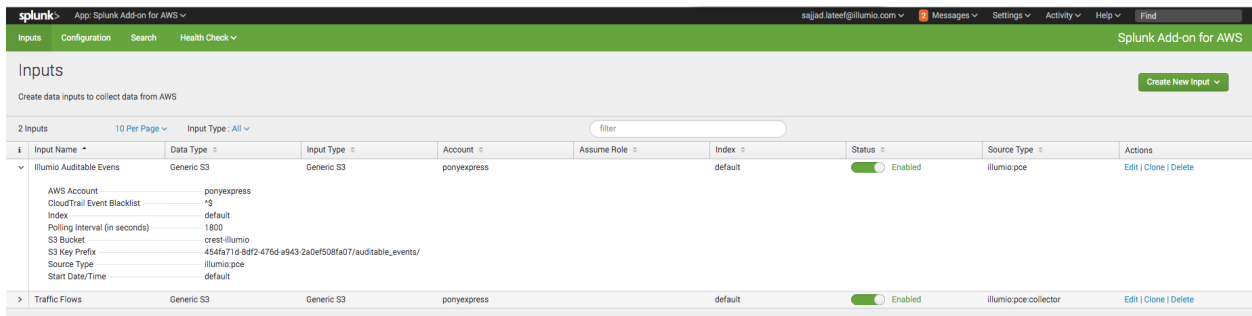
**Index**

Index: default

Buttons: Cancel, Save

10. Click **Save**.

Back in the Inputs screen, you should now see your two new inputs. Click the arrow next to the input name to see some details about the input:



#	Input Name	Data Type	Input Type	Account	Assume Role	Index	Status	Source Type	Actions
1	illumio Auditable Events	Generic S3	Generic S3	ponyexpress		default	Enabled	illumio:pce	Edit   Clone   Delete
<ul style="list-style-type: none"> <li>AWS Account: ponyexpress</li> <li>CloudTrail Event Blacklist: *</li> <li>Index: default</li> <li>Polling Interval (in seconds): 1800</li> <li>S3 Bucket: crest-illumio</li> <li>S3 Key Prefix: 454fa71d-8df2-476d-a943-2a0ef508fa07/auditable_events/</li> <li>Source Type: illumio:pce</li> <li>Start Date/Time: default</li> </ul>									
2	Traffic Flows	Generic S3	Generic S3	ponyexpress		default	Enabled	illumio:pce:collector	Edit   Clone   Delete

You should now have access to your VEN flow data and auditable event logs. See [“Example Splunk Queries”](#) for examples of how to access the data. Illumio can provide additional sample Splunk queries if needed. Contact Illumio Technical Support for assistance.

## Speeding Up UI Rendering

If most of your searches will cover a time period of 7 days or less, you can make the panels in the app respond more quickly by modifying the `summariesonly` macro.

1. Choose Settings > Advanced Search > Search Macros.
2. Click the `summariesonly` macro. The macro editor is displayed.
3. Change the definition of the macro to "`summariesonly=true`".

## Configuring Alerts

If you have administrator privileges on the Illumio App for Splunk, you can create or update alert configurations using the Alert Configuration page. By using alert configurations, you can watch for events that are of interest related to a variety of Illumio PCE entities such as rules and workloads.

To display the Alert Configuration page, click Alert Configuration in the top-level navigation menu. This link only appears if your user account has the "admin" role.

After creating alert configurations, use the Alerts page to set up the usage of the alerts, such as sending emails whenever an alert is triggered. See Splunk documentation for details about alert configuration and the Alerts page.

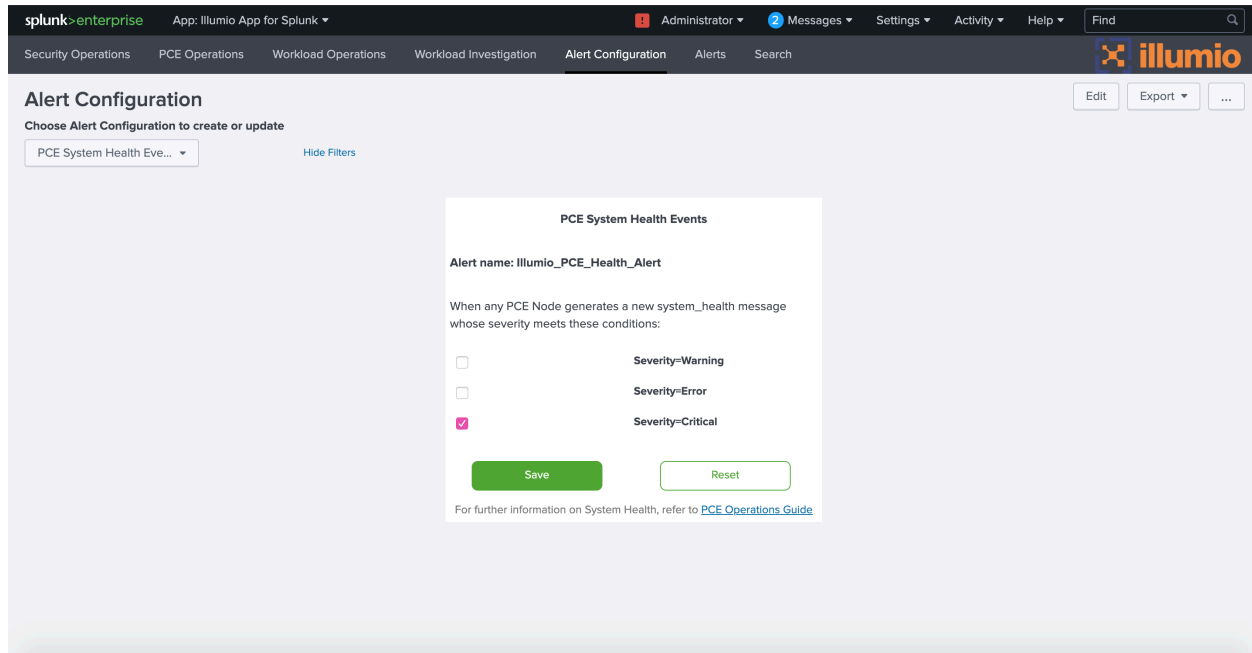
In the Illumio App for Splunk, you can configure five different types of alerts. Choose the desired alert type in the dropdown list on the Alert Configuration screen.

The options in the dropdown are:

- PCE System Health Events
- Rule Set Writing/Update
- Rule Writing Update
- Policy Provisioning
- Workload Labeling

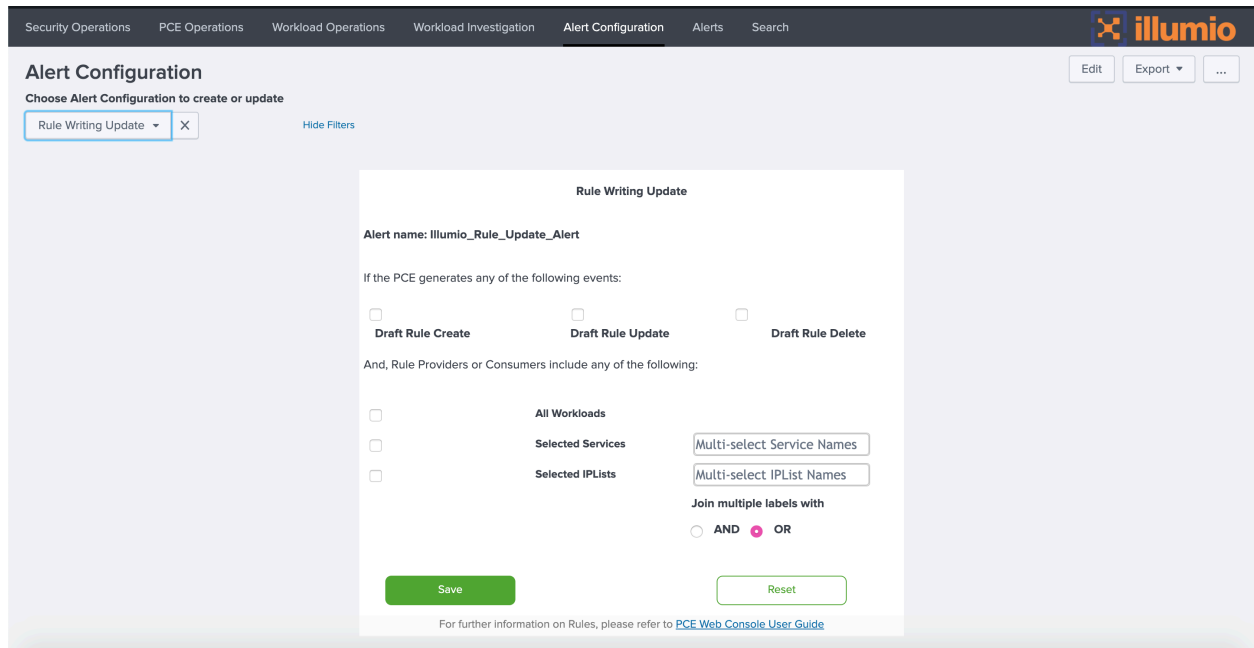
To configure alerts about system health events, choose PCE System Health Events from the dropdown, then choose which level of event severity (warning, error, or critical) to include. For details on conditions that trigger severity of warning or above, see PCE


Operations Guide's section on [PCE Health Monitoring with Syslog](#) in the *PCE Operations Guide*.



You can configure alerts about changes to rules on the PCE. For example, a draft rule might be created that affects all workloads. Since this is a very wide-ranging effect, which might have been unintentional, you might want to be alerted so you can confirm the rule is correct.

To configure alerts about changes to draft rules, choose Rule Writing Update in the dropdown, then choose which type of rule change to include (create new rule, update a rule, or delete a rule, or any combination) and which rule providers or consumers to include (all workloads, or a subset based on service names or IP lists). Choose the AND operator if all the selected rule providers/consumers must be matched. Choose the OR operator to match any one provider/consumer from the selected list.



Security Operations PCE Operations Workload Operations Workload Investigation **Alert Configuration** Alerts Search 

**Alert Configuration** Edit Export ...

Choose Alert Configuration to create or update

Rule Writing Update X Hide Filters

**Rule Writing Update**

Alert name: illumio\_Rule\_Update\_Alert

If the PCE generates any of the following events:

Draft Rule Create  Draft Rule Update  Draft Rule Delete

And, Rule Providers or Consumers include any of the following:

All Workloads  Selected Services   
 Selected IPLists

Join multiple labels with

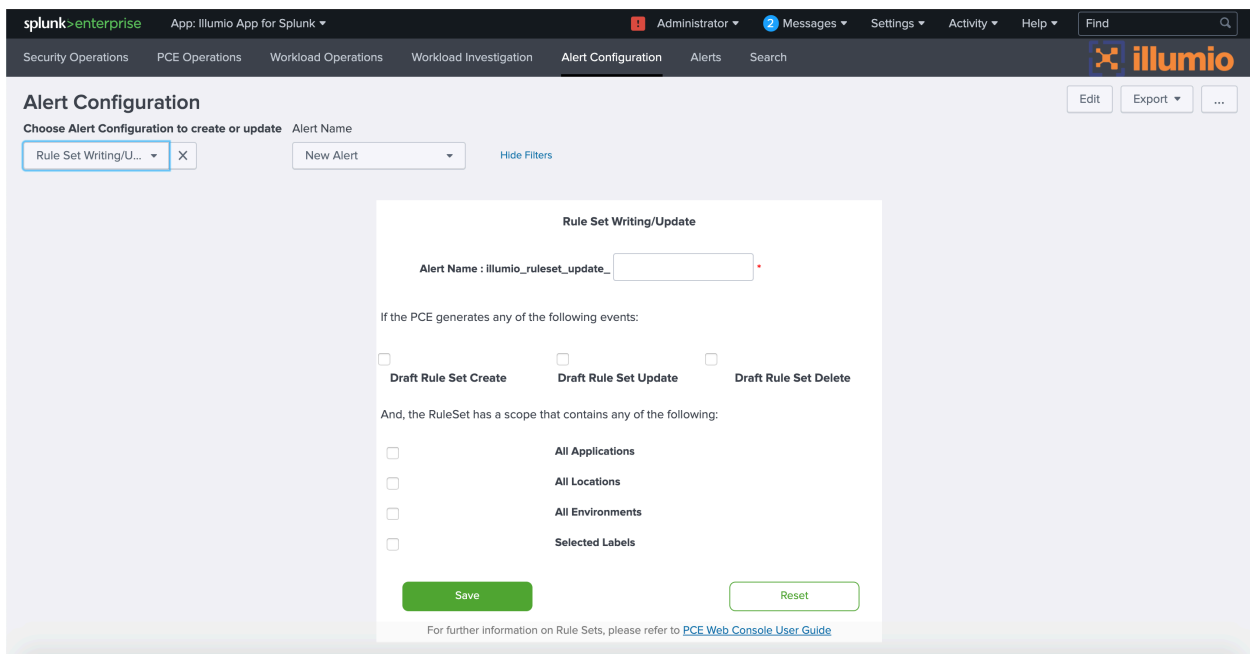
AND  OR

Save Reset

For further information on Rules, please refer to [PCE Web Console User Guide](#)

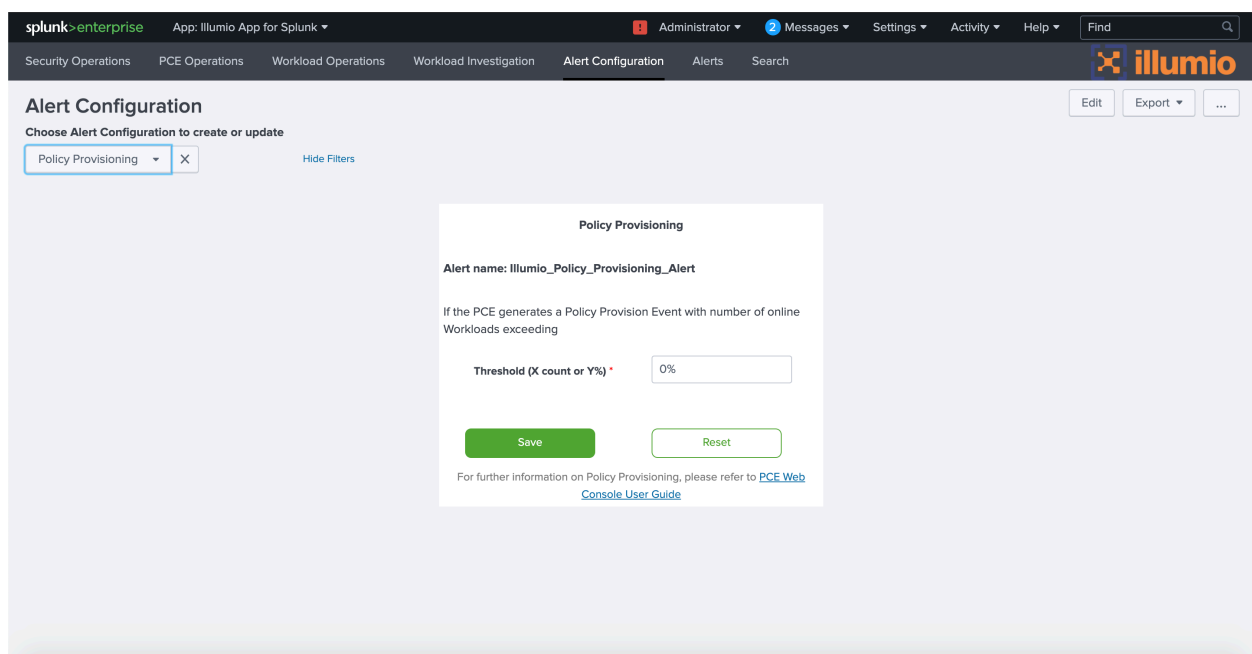
Similarly to rules, you can configure alerts about changes to draft rulesets on the PCE. For example, a draft ruleset might be created that has a broad scope. When provisioned, the ruleset might affect too many workloads unintentionally. It is useful to be alerted so you can confirm the ruleset's scope is correct.

To configure alerts about new, changed, or deleted rulesets, choose Rule Set Writing/Update from the dropdown. In the Alert Name dropdown, choose New Alert if you are setting up a new alert, or choose the name of an existing alert if you want to make changes to its configuration. If you are creating a new alert, give it a name in the Alert Name field. Choose which type of ruleset change to include (create new ruleset, update a ruleset, or delete a ruleset) and which ruleset scopes to include (based on applications, locations, environments, or labels).



In PCE 19.1.0 and later, you can configure alerts to be triggered when new policies are provisioned. For example, you might want to know if a new policy is being provisioned to a large number of workloads.

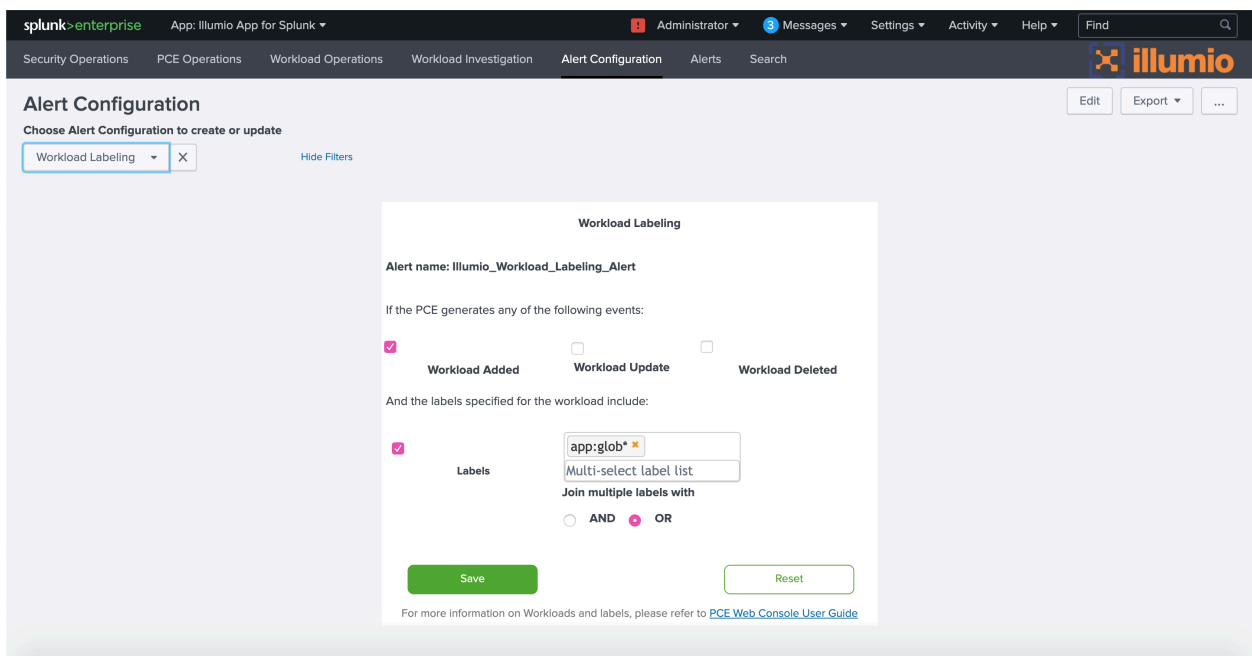
To configure alerts about provisioning of new policies, choose Policy Provisioning from the dropdown, then set the minimum number of workloads that must receive the provisioning. The number can be specified as an absolute number, such as 100, or a percentage, such as 10% of the workloads. To trigger the alert no matter how many workloads are involved, set the threshold to 0.



The screenshot shows the Splunk Alert Configuration interface for the 'Policy Provisioning' alert. The interface includes a navigation bar with 'splunk>enterprise' and 'App: Illumio App for Splunk'. The main content area is titled 'Alert Configuration' and shows a dropdown menu set to 'Policy Provisioning'. The alert name is 'Illumio\_Policy\_Provisioning\_Alert'. The configuration text reads: 'If the PCE generates a Policy Provision Event with number of online Workloads exceeding'. Below this, there is a 'Threshold (X count or Y%)' field set to '0%'. There are 'Save' and 'Reset' buttons. At the bottom, there is a link to the 'PCE Web Console User Guide'.

You can configure alerts about changes to workload labels. For example, it might be a reason for concern if a workload label is changed in a way that reduces the workload’s security posture, such as changing from “Production Top Secret” to “Internal Testing.”

To configure alerts about changes to workload labels, choose Workload Labeling from the dropdown, then choose which type of change to include (add, update, or delete label) and which labels the workload must have. Choose the AND operator if a workload must have all the selected labels. Choose the OR operator to match any one label from the selected list.



## Post-Installation Required Settings

After installing and configuring the Illumio App for Splunk, make the following settings.

### Accelerate Data Model

The Illumio App for Splunk ships with the Illumio data model acceleration disabled, as required for Splunk App Certification. After installation, you must enable the acceleration of the data model, which is used for visualizations in the dashboards. Use the steps in [“Data Model and Data Model Acceleration.”](#)

## Update Search Macros for Custom Index

If you choose a non-default index, you must update the search macros to use the custom index. Use the steps in [“Index, Source, and Source Types.”](#)

## Upgrade the App

You can upgrade the Illumio App and TA using the CLI or UI.

Upgrade through the CLI:

1. Download the tarball of App or TA from Splunk base.
2. Stop the Splunk server.
3. Run this command:

```
$SPLUNK_HOME/bin/splunk install app APP_NAME.tgz -update 1  
-auth username:password
```

4. Start the Splunk Server.
5. If you are upgrading to Splunk 3.0 from a previous version of the app, rebuild the data model after you install the app. See [“Data Model and Data Model Acceleration.”](#)
6. If you are upgrading to Splunk 3.0 from a previous version of the app, remove any customizations you have made to the app in the `local` directory.

Upgrade through the UI:

1. Click **Manage Apps**.
2. Find the Illumio app and TA entry from the list.
3. If a newer version is available, an entry “Update to ...” is displayed.
4. Click the link of the newer version under the version column.
5. If you are upgrading to Splunk 3.0 from a previous version of the app, rebuild the data model after you install the app. See [“Data Model and Data Model Acceleration.”](#)
6. If you are upgrading to Splunk 3.0 from a previous version of the app, remove any customizations you have made to the app in the `local` directory.

Upgrade by installing files:

1. Click **Manage Apps**.
2. Click **Install App from File**.



3. In the dialog, upload the SPL file corresponding to the newer version of TA or App.
4. Select the **Upgrade App** check box.
5. Click **Upload**.
6. Restart Splunk after uploading both the TA and app.
7. If you are upgrading to Splunk 3.0 from a previous version of the app, rebuild the data model after you install the app. See “[Data Model and Data Model Acceleration](#).”
8. If you are upgrading to Splunk 3.0 from a previous version of the app, remove any customizations you have made to the app in the `local` directory.

## Alerting Actions and Adaptive Response Framework

This section tells how the Alerting Actions and Adaptive Response Framework work with the Illumio App for Splunk. This section covers features where the Splunk App takes action by invoking update APIs on the Illumio PCE.

There are two types of quarantine provided by Illumio:

- A custom alert action provided for Splunk Enterprise, also called Splunk Core, the base Splunk product. See [Using custom alert actions](#) in the Splunk documentation.
- An adaptive response action provided for Splunk Enterprise Security (ES), which is different from the Splunk core product. See [Set up adaptive response actions in Splunk Enterprise Security](#) in the Splunk documentation.

The Splunk core already provides standard alert actions such as sending emails, notable events, and calling a Webhook URL. Modular actions on top of standard alert actions are nothing but custom alert actions. These custom alert actions let you invoke Python scripts that use APIs external to Splunk.

The Enterprise Security Suite app provides support for Correlation/Saved Searches with notable actions. When a Splunk Enterprise Security Correlation/Saved Search (with notable event mapped) is executed and gets at least one event in the results, notable events will be created through a standard notable action. These notable events are visible in the Incident Review dashboard of Splunk Enterprise Security App. No other alert action (other than the notable action) is executed automatically, because none are mapped.

Splunk provides the Adaptive Response Framework in the Enterprise Security Suite by leveraging the modular action functionality provided in Splunk\_SA\_CIM.

Using Splunk Enterprise Security's Adaptive Response Framework, Illumio PCE administrators can quarantine workloads managed by the PCE directly from Splunk Apps whenever the events are detected in Splunk, based on data sent by any source of alerts in Enterprise Security.

There are two ways to invoke actions on the workloads:

- Quarantine workloads using Splunk Core Alert Actions.
- Quarantine workloads using Splunk Enterprise Security Suite's Adaptive Response Framework.

## Quarantine Workload Using Splunk Core Alert Actions

If Splunk Enterprise Security Suite (ESS) is not installed in your Splunk infrastructure, the Illumio App for Splunk offers a way to monitor and take action on the events reported by analytics on Illumio PCE logs.

To achieve this, the Illumio Add-on for Splunk leverages the custom alert action to quarantine the workload. These actions are available on the drill downs from the main dashboards.

## Quarantine Workload using Enterprise Security Suite

Splunk provides the Splunk Enterprise Security Suite (ESS), which leverages Splunk's Adaptive Response Framework and allows administrators to monitor and manage threats and incidents directly from Splunk apps. It has rich dashboards that help monitor incidents and take actions on these incidents.

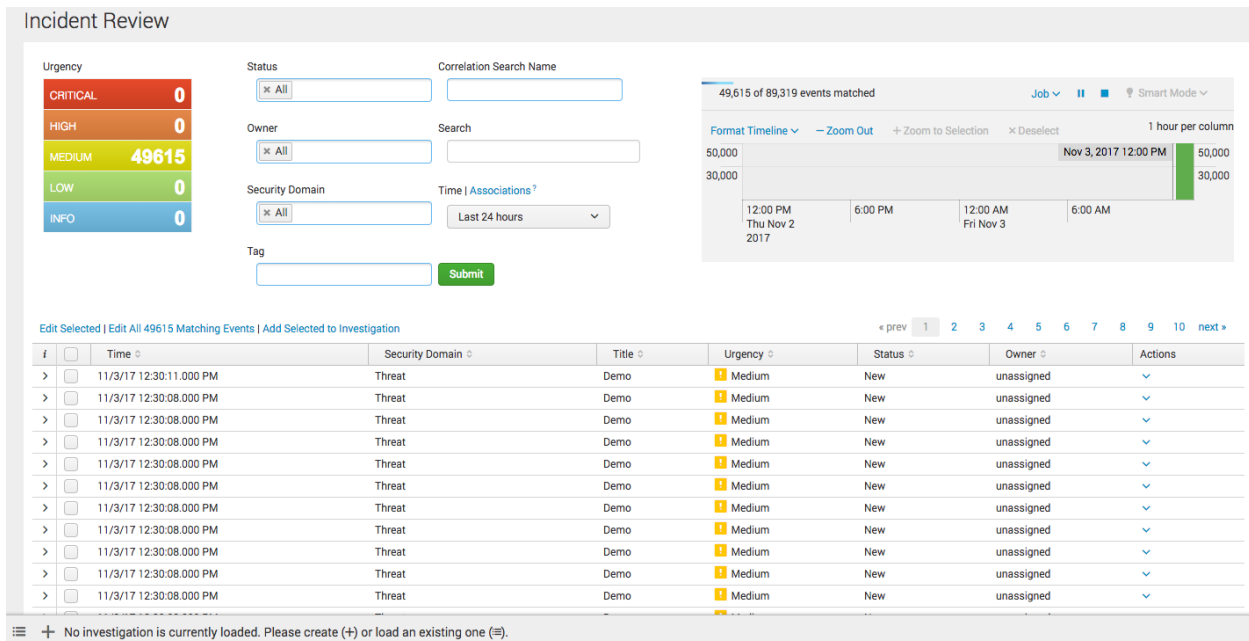
Splunk Enterprise Security Suite is extendable by adding a compatible Module App (Adaptive Response Add-ons) for a particular domain or technology. The Suite detects configurations in these Adaptive Response Add-ons and helps monitor and take actions on the incidents reported by these Add-ons.

The Illumio Add-on for Splunk (TA-Illumio) is one such module for Splunk Enterprise Security Suite. It leverages Splunk Adaptive Response Framework and empowers

System Administrators to monitor and take actions on incidents reported by analytics on Illumio PCE events or logs from the Splunk Enterprise Security Suite dashboards.

When using the Splunk Enterprise Security (ES) suite, the Illumio Splunk TA can be installed on a single ES Search Head (SH), or on both an ES SH and an associated ES Search Head Cluster (SHC). This allows the Adaptive Response to be invoked from any installed TA location. The Illumio data is stored on the indexers only, and not on the search head nodes, so the data is not duplicated. If the TA is installed only on a single ES SH, the data is normalized for the associated SHC.

The Incident Review dashboard:



**Incident Review**

Urgency: CRITICAL 0, HIGH 0, MEDIUM 49615, LOW 0, INFO 0

Status: [All] Correlation Search Name: [ ]

Owner: [All] Search: [ ]

Security Domain: [All] Time | Associations: Last 24 hours

Tag: [ ] [Submit]

49,615 of 89,319 events matched

Format Timeline | Zoom Out | Zoom to Selection | Deselect | 1 hour per column

#	Time	Security Domain	Title	Urgency	Status	Owner	Actions
>	11/3/17 12:30:11.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵
>	11/3/17 12:30:08.000 PM	Threat	Demo	Medium	New	unassigned	⌵

⌵ No investigation is currently loaded. Please create (+) or load an existing one (=).

Splunk, as a part of the Adaptive Response Framework, has enhanced this Incident Review dashboard in the Enterprise Security Suite app, which provides the option to take actions on these notable events.

To view the notable event details, expand the left arrow for that notable event. To execute alert actions manually for each of the notable events, click **Run Adaptive Response Actions** for the notable event and select the specific Alert Action.

The screenshot shows the Splunk interface for a notable event. The event details include:

- Description:** unknown
- Additional Fields:** Value, Action
- Event Details:**
  - event\_id: 2096BB9A-ECCF-499C-9833-F50811DE7F49@notable@9b40e5a1f42d6c3bd91a75485028b0fe
  - event\_hash: 9b40e5a1f42d6c3bd91a75485028b0fe
  - eventtype: modnotable\_results
  - notable
  - Short ID: Create Short ID

The dropdown menu is open, showing the following options:

- Add Event to Investigation
- Create notable event
- Build Event Type
- Extract Fields
- Run Adaptive Response Actions** (highlighted)
- Share Notable Event
- Suppress Notable Events
- Show Source

Below the menu, the 'Adaptive Responses' table shows:

Response	Mode	Time	User	Status
Notable	saved	2017-11-03T12:30:07+0000	nobody	success

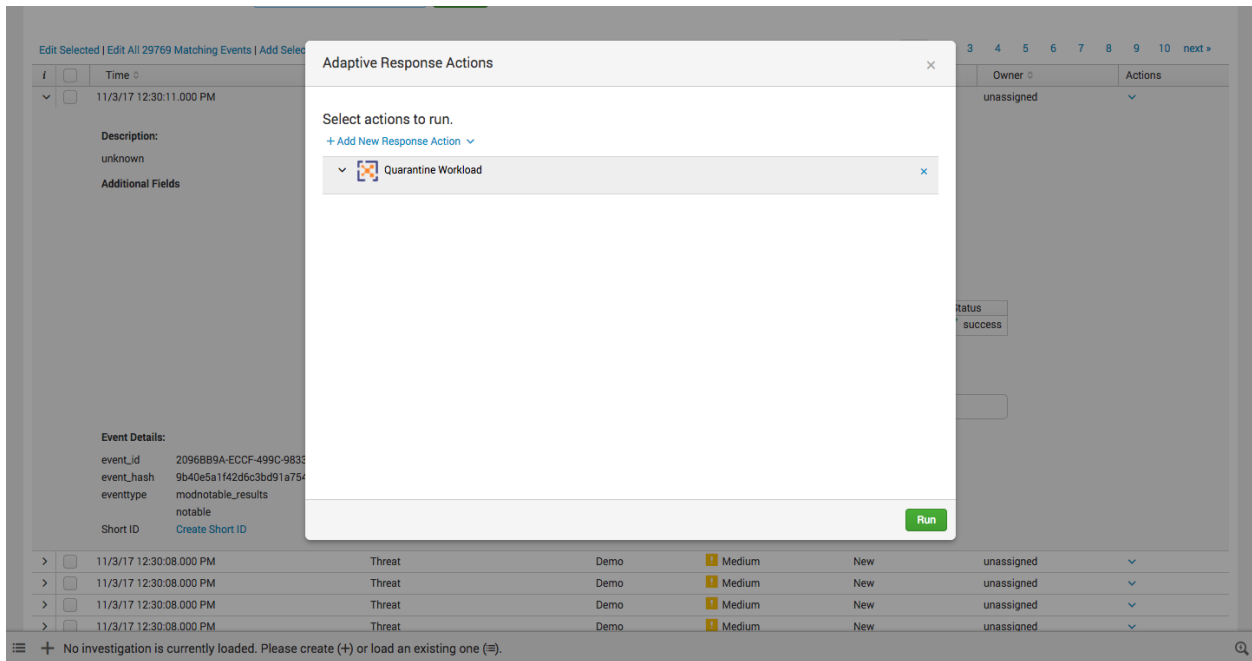
The 'Next Steps' section indicates: No Next Steps defined.

When you click Run Adaptive Response Actions for a notable event, a menu appears that lists all of the standard and custom actions. This list is created by reading the alert\_actions.conf files of all the installed apps on the Splunk instance. Users can select multiple actions on this popup menu and run them for that notable event.

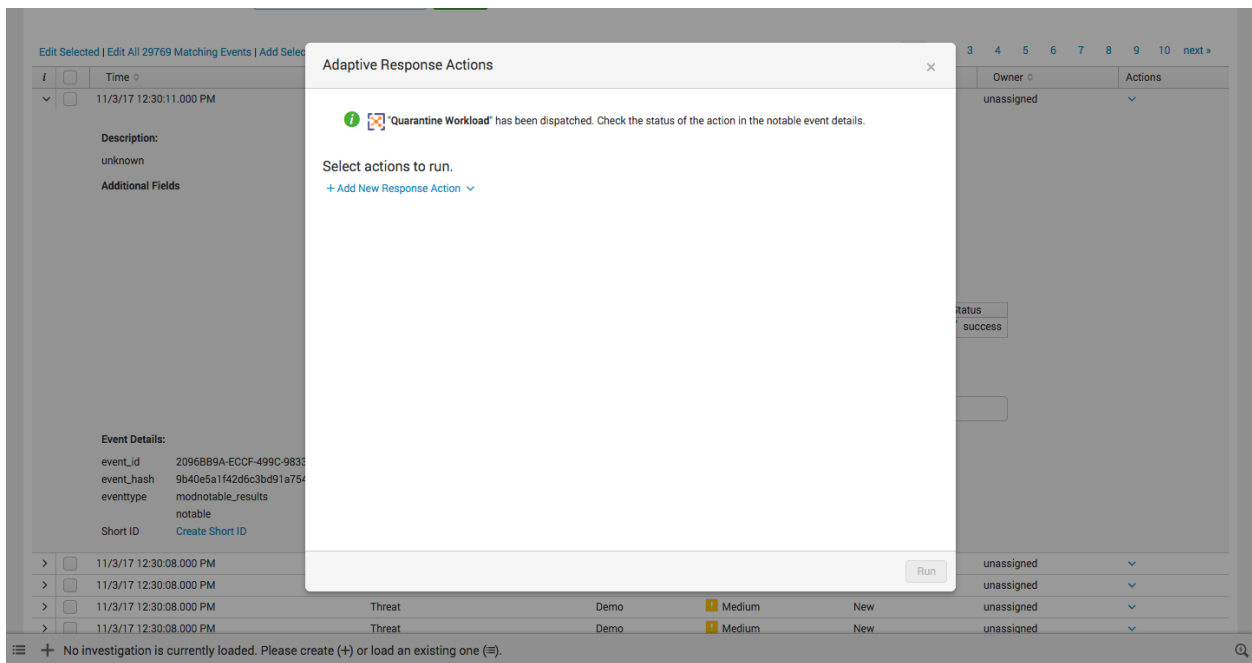
The screenshot shows the 'Adaptive Response Actions' popup menu. The title is 'Adaptive Response Actions'. Below the title, it says 'Select actions to run.' and '+ Add New Response Action'. There is a search bar and a category dropdown set to 'All'. The list of actions includes:

- Stream Capture** (STM icon): Creates stream capture. Category: Information Gathering | Task: create | Subject: network capture | Vendor: Splunk
- Quarantine Workload** (Quarantine icon): Custom action for marking a workload as quarantine. Category: Information Gathering | Task: Update | Subject: Workload | Vendor: Illumio (highlighted with a red box)
- Nbtstat** (Nbtstat icon): Runs the nbtstat command. Category: Information Gathering | Task: scan | Subject: device | Vendor: Operating System
- Nslookup** (Nslookup icon): Runs the nslookup command. Category: Information Gathering | Task: scan | Subject: device | Vendor: Operating System
- Ping** (Ping icon): Runs the ping command.

A 'Run' button is located at the bottom right of the popup menu.



The screenshot shows the Splunk interface with a modal dialog titled "Adaptive Response Actions". The dialog contains the text "Select actions to run." and a link "+ Add New Response Action". Below this, a dropdown menu shows "Quarantine Workload" selected. At the bottom right of the dialog is a green "Run" button. The background shows a list of events with columns for Time, Description, Additional Fields, Event Details, and a table with columns for Owner and Actions.



The screenshot shows the same Splunk interface, but the modal dialog now displays a green information icon and the message: "Quarantine Workload" has been dispatched. Check the status of the action in the notable event details." Below this message, the "Select actions to run." section is still visible, but the "Run" button is now disabled (greyed out).

When these actions are run, each selected corresponding action is invoked from alert\_actions.conf.

## Quarantine Workload from Illumio Splunk App

If you have both the “admin” role and “Illumio\_quarantine\_workload” role, you can quarantine Workloads from the Illumio Splunk App by clicking the Quarantine button, which appears on the following dashboards:

- Port Scan (on Security Operations dashboard)
- Firewall Tampering (on Security Operations dashboard)

If the Quarantine button is greyed out, then you do not have adequate permissions to quarantine workloads. See [“Access to Quarantine Workload Action.”](#)

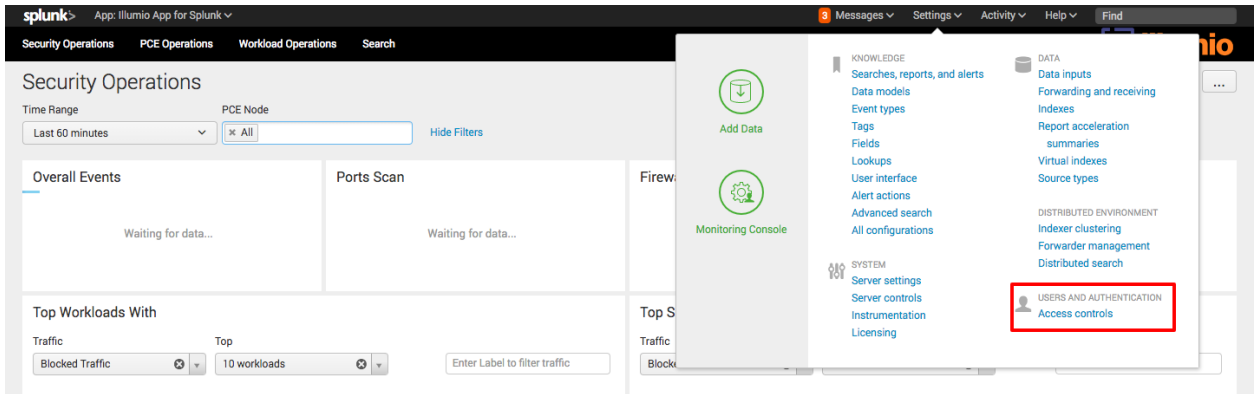
Time	Source IP	Source	Destination IP	Destination	Source Label	Destination Label	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:011b	-	f000:0000:0000:0000:0200:000a:0000:015b	-	app:Point-of-Sale env:PCI loc:CA role:Database	app:HREnrollment env:Production loc:CA role:Processing	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:011a	-	f000:0000:0000:0000:0200:000a:0000:015b	-	app:Point-of-Sale env:PCI loc:CA role:Database	app:HREnrollment env:Production loc:CA role:Processing	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:0116	-	f000:0000:0000:0000:0200:000a:0000:015b	-	app:Point-of-Sale env:PCI loc:CA role:Web	app:HREnrollment env:Production loc:CA role:Processing	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:0111	-	f000:0000:0000:0000:0200:000a:0000:015b	-	app:HRM env:Staging loc:NY role:Database	app:HREnrollment env:Production loc:CA role:Processing	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:0110	-	f000:0000:0000:0000:0200:000a:0000:015b	-	app:HRM env:Staging loc:NY role:Database	app:HREnrollment env:Production loc:CA role:Processing	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:016c	-	f000:0000:0000:0000:0200:000a:0000:0156	-	app:ShoppingCart env:Production loc:AWS role:Database	app:eCommerce env:Production loc:Azure role:Database	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:0114	-	f000:0000:0000:0000:0200:000a:0000:0156	-	app:HRM env:Staging loc:NY role:Web	app:eCommerce env:Production loc:Azure role:Database	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:0172	-	f000:0000:0000:0000:0200:000a:0000:0152	-	app:Catalog env:Production loc:AWS role:Database	app:eCommerce env:Production loc:Azure role:Web	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:0147	-	f000:0000:0000:0000:0200:000a:0000:0152	-	app:CoreServices env:Production loc:CA role:DomainController	app:eCommerce env:Production loc:Azure role:Web	Quarantine	Investigate
2019-12-12 23:03:00 UTC	f000:0000:0000:0000:0200:000a:0000:0132	-	f000:0000:0000:0000:0200:000a:0000:0152	-	app:Ordering env:Production loc:CA role:Load Balancer	app:eCommerce env:Production loc:Azure role:Web	Quarantine	Investigate

## Access to Quarantine Workload Action

By default, users do not have access to the Quarantine Workload action either in the Splunk App or in Adaptive Response Action.

To enable a Splunk user to take quarantine actions on Workloads, grant the user the “Illumio\_quarantine\_workload” role and the “admin” role. Only local users can be granted this role. SAML users cannot, as their roles are controlled by an external system.

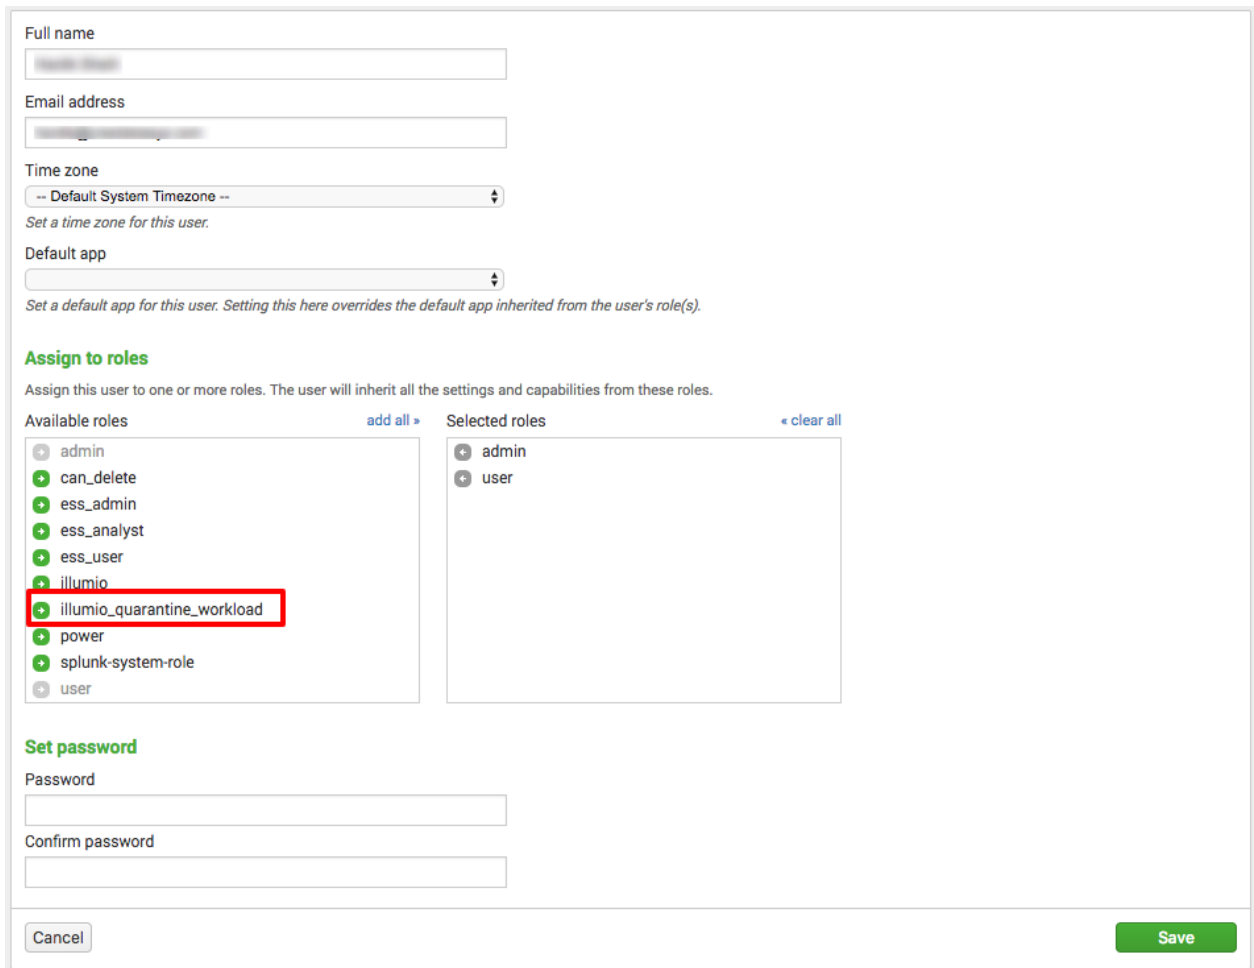
1. Click **Settings** > **Access Control**.



2. Click **Users**.



3. Click the username to which the role needs to be granted.
4. In the Role section of the edit screen, grant the required roles.



The image shows a user configuration form in Splunk. It includes fields for 'Full name', 'Email address', 'Time zone' (set to 'Default System Timezone'), and 'Default app'. Below these is the 'Assign to roles' section, which has two columns: 'Available roles' and 'Selected roles'. In the 'Available roles' list, the role 'illumio\_quarantine\_workload' is highlighted with a red box. The 'Selected roles' list contains 'admin' and 'user'. At the bottom, there are 'Set password' fields for 'Password' and 'Confirm password', and 'Cancel' and 'Save' buttons.

5. Click **Save**.

## Example Splunk Queries

This section provides sample queries to help you get started writing your own Splunk queries using Illumio data.

### Workload report

This is a fairly complicated query, but it can be used to generate a workload report showing the labels associated with each workload. The query results can be exported as CSV for reporting purposes.



```
`illumio_get_index` sourcetype="illumio:pce:metadata"
(illumio_type="illumio:pce:workload") | search "agent.href"="*" fqdn="*" |
rex field=href "orgs\\/\\d+\\/workloads\\/ (?<workload_uuid>\\S+)" | fields
labels{}.href uuid hostname os_id public_ip agent.config.mode
agent.config.log_traffic agent.status.status workload_uuid | mvexpand
labels{}.href | rename labels{}.href as href | lookup
illumio_workload_mapping_lookup href workload_uuid OUTPUTNEW type label |
eval {type}_label=label | stats values(*) as * by workload_uuid | table
hostname, public_ip, os_id, agent.config.mode, agent.config.log_traffic,
agent.status.status, role_label, app_label, env_label, loc_label
```

### Top events

```
`illumio_get_index` sourcetype="illumio:pce" | top event_type
```

### Top outgoing connections

```
sourcetype="illumio:pce:collector" | search dir=O | top dst_ip
```

### Top incoming connections

```
sourcetype="illumio:pce:collector" | search dir=I | top src_ip
```

### Most active machines

```
sourcetype="illumio:pce:collector" | search dir=I | top dst_ip
```

### Top source ports

```
sourcetype="illumio:pce:collector" | top dest_port
```

### Top machines making connections to machines with a given label

This example query returns the top machines making connections to machines with the Production label:

```
sourcetype="illumio:pce:collector" dest_env_label="Production" | top src_ip
Available label search keys:
dest_role_role, dest_app_label, dest_env_label, dest_loc_label
src_role_label, src_app_label, src_env_label, src_loc_label
```

## Top machines with connections in a given network

This example query returns the top machines with connections in 10.0.0.0/8:

```
sourcetype="illumio:pce:collector" | search dst_ip=10.0.0.0/8 | top dst_ip
```

## Geolocate destination IPs (plot on a map)

```
sourcetype="illumio:pce:collector" | search dst_ip!=10.0.0.0/8 | iplocation  
dst_ip | geostats count latfield=lat longfield=lon
```

## Troubleshooting

This section provides tips to diagnose and fix common issues.

### Data collection not working

After installing the application, all dashboards should start populating data. If you don't see data in the dashboards, use the following steps for troubleshooting:

- If you specified a non-default index in the Data Input, confirm that you have modified the macro `illumio\_get\_index` with indexes selected while creating Data Input (Modular input) and that you have modified the Advanced search for either the default index or your custom index. See "[Index, Source, and Source Types](#)".
- Run the following query to verify that data is being indexed into Splunk:

```
search `illumio_get_index` | stats count by sourcetype
```

- Verify that SPLUNK\_HOME is pointing to the correct Splunk directory.
- Look for errors in the file ta-illumio.log. This file is in the \$SPLUNK\_HOME/var/log/TA-Illumio/ folder.
- Check to see whether the selected time range covers the time when the traffic flow summaries were generated.
- If a data model acceleration is disabled, graphs will not display. Enable data model acceleration. See "[Data Model and Data Model Acceleration](#)" earlier in this document.

### Can't use same port in new Data Input (Modular Input)

**Symptom:** After deleting an existing configured data input (modular input), can't create a new modular input on the same port.

**Cause:** The port number is still in the Data Input TCP ports list.

**Fix:** Remove the port from the Data Input “TCP” ports list before you try to use the same port again. This enables you to re-use the port which was configured in the previous data input (modular input).

## Data not available immediately after configuring data input (modular inputs)

**Symptom:** Upon successful creation of data inputs (modular input), 5 minutes elapse before data starts indexing.

**Cause:** 5 minutes is the default time interval configured.

**Fix:** No action is required. This is expected behavior.

## Authentication failure on Data Input (Modular Input) page

**Fix:**

- Check the network connectivity in between applications to ensure that there are no connectivity issues.
- Ensure that the API Key and API Secret stored in the setup page are in sync with the API key generated by the Illumio PCE. As required by App certification, these secrets are stored in a secure key store. If the data input (modular input) is modified, they will need to be entered again, as they cannot be read back from the secure key store.
- If the Illumio Application is deployed on a non-trusted CA or using a self-signed certificate, the certificate directory path has to be provided. The certificate also needs the correct certificate validation trust chain.

## Quarantine button is grayed out or does not work as expected

**Cause:** To use the Quarantine button, the Splunk user needs to have both the ‘admin’ role and the ‘illumio\_quarantine\_workload’ role. If the button is green but does not work as expected, it is likely because of a missing ‘admin’ role. To investigate the cause, check the file ta-illumio.log for error messages.

**Fix:** To grant the required roles, use the steps in “[Access to Quarantine Workload Action](#).”

## Invalid Certificate File error on Data Input (Modular input) page

**Cause:** The PCE may present an SSL certificate issued by different authorities such as a primary CA, a secondary CA, Local CA authority or even a self-signed certificate. For correct SSL operation, the Splunk server must be able to fully trust the PCE’s certificate and verify the certificate’s trust chain.

**Fix:** When a local CA Authority issued SSL certificate or a self-signed SSL certificate is used with the PCE, the CA Certificate bundle needs to be uploaded onto the Splunk Server and the full path to the directory containing the certificate should be provided in the data input.

If using a local CA Authority or a certificate issued by a secondary CA, the Splunk server CA trust chain must be updated to verify the certificate presented by the PCE.

For example, on Linux, use the update-ca-trust tool. Copy the certificate chain to `/etc/pki/ca-trust/source/anchors/` and then run the following commands:

```
update-ca-trust force-enable
update-ca-trust extract
update-ca-trust check
```

See Splunk documentation for further information.

## PCE labels are not updated in Security Operations dashboard

**Symptom:** If there are new labels, or workloads are added to the PCE, the new labels and workloads will not be visible right away.

**Cause:** The default interval to sync workloads and labels from the PCE is set to a minimum of 60 minutes. This period is configurable through the data input. The newly added labels or workloads on the PCE should be available in the Splunk App after an interval of 60 minutes.

**Fix:** To force a resync of labels to the PCE, you can disable and enable the TA. This forces the TA to make API calls to the PCE. This operation should be used with caution, due to the additional API calls to the PCE.

## Security Operations shows “Search is waiting for input”

**Symptom:** When upgrading an older version of the app, the Security Operations panels do not display graphs. Instead, they display "Search is waiting for input ...".

**Cause:** Old dashboard files stored locally on the Splunk server.

**Fix:**

1. Delete the folder `$SPLUNK_HOME/etc/apps/IllumioAppforSplunk/local`.
2. Restart Splunk.

## Path for the custom certificate: invalid certificate file

**Symptom:** An error is generated in the Policy Compute Engine (PCE) `ta-illumio.log` file when attempting to add Illumio Data Inputs. Saving any information for the Data Inputs is not allowed.

```
Error from file /opt/splunk/var/log/TA-Illumio/ta-illumio.log :  
2018-10-24 16:33:48,844 - Illumio_MODINPUT - ERROR - Path for the  
custom certificate: Invalid certificate file
```

A Splunk error, due to PCE certificate trust, is also displayed:

splunk>enterprise Apps 3 Messages Settings Activity Help

knpce1  
Data inputs > Illumio > knpce1

Encountered the following error while trying to update: Path for the custom certificate: Invalid certificate file

PCE URL \*

API Authentication Username \*   
e.g. 'api\_1234567890'

API Secret \*

Port Number for syslogs (TCP) \*

Port Scan configuration: scan interval in seconds \*   
Interval during which the Port Scan Threshold is exceeded

Port Scan Configuration: Unique ports threshold \*

**Cause:** This error is an indication that a PCE certificate was not trusted, even though the certificate has already been added the local system certificate store.

**Fix:** Adding Illumio Data Inputs allows the Illumio App for Splunk to connect to a configured PCE to extract data for PCE health and workloads information. When the Illumio App for Splunk attempts a connection to the PCE, it can fail due to a certificate trust even when a local browser trusted the PCE certificate, since it was already added to the local system certificate store. Splunk uses a Python library that is local to the Splunk application; therefore, it carries its own local certificate authority file it trusts.

There are two ways to add a secure trust to the PCE:

- Add both intermediate and root certificate authority to the local Python cacert.pem file:
  - In Windows: C:\Program Files\Splunk\Python-2.7\Lib\site-packages\requests\cacert.pem
  - In Linux: /opt/splunk/lib/python2.7/site-packages/requests/cacert.pem

OR,

- Create a certificate file that includes the PCE server certificate, intermediate certificate, and root CA certificate in that order, then place the file in the Splunk home directory. The certificate should be in PEM format. Added the certificate path to the Illumio Data Inputs. Use the following steps:

1. Use a text editor to cut and paste the certificate chain, and avoid extraneous characters. The Splunk home directory is as follows:
  - Windows Splunk home directory: C:\Program Files\Splunk\
    - Linux Splunk home directory: /opt/splunk/
2. Export the certificates using any browser, then cut and paste them together. The following is an example of what should be in a certificate file:

```
-----BEGIN CERTIFICATE-----  
  
< Sever Certificate base64 encoded >  
  
-----END CERTIFICATE-----  
  
-----BEGIN CERTIFICATE-----  
  
< Intermediate Certificate base64 encoded >  
  
-----END CERTIFICATE-----  
  
-----BEGIN CERTIFICATE-----  
  
< Root CA Certificate base64 encoded >  
  
-----END CERTIFICATE-----
```

3. Set the path in the Illumio Data Inputs under 'Settings > Data Inputs > Illumio > (select input) > checkbox "more settings" > Custom (self-signed) certificate path > (path of certificate)'.
  - Windows Splunk home directory: C:\Program Files\Splunk\
    - Linux Splunk home directory: /opt/splunk/

Data Collection

More settings

Interval   
Period between making API calls

Host   
Set the host with this value.

Index   
Set the destination index for this source.

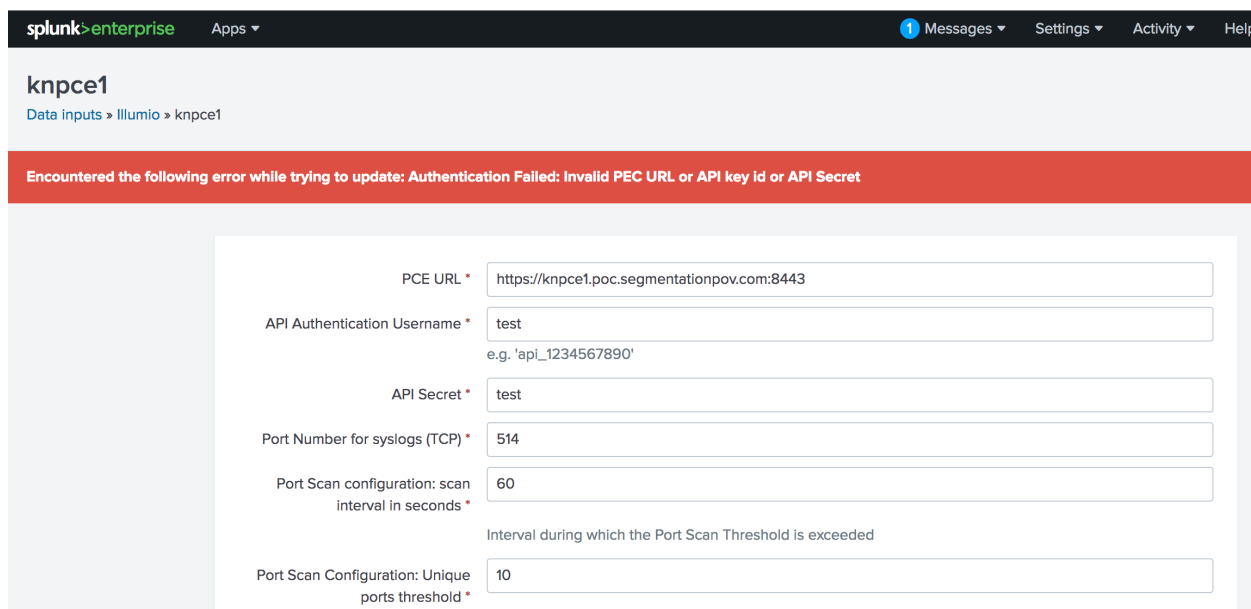
Custom (self-signed) certificate path   
Path for the custom certificate

Allowed port scanner IP addresses   
Comma Separated list of hosts, which will be ignored in Port scans



## Authentication Failed: Invalid PCE URL or API key id or API Secret

**Symptom:** When applying data inputs in Splunk for the Illumio App for Splunk, you receive the following error from the Splunk UI: “Authentication Failed: Invalid PEC URL or API key id or API Secret.”



The screenshot shows the Splunk Enterprise interface for the 'knpce1' data input. A red error banner at the top states: "Encountered the following error while trying to update: Authentication Failed: Invalid PEC URL or API key id or API Secret". Below the banner, the configuration fields are as follows:

PCE URL *	<input type="text" value="https://knpce1.poc.segmentationpov.com:8443"/>
API Authentication Username *	<input type="text" value="test"/> <small>e.g. 'api_1234567890'</small>
API Secret *	<input type="text" value="test"/>
Port Number for syslogs (TCP) *	<input type="text" value="514"/>
Port Scan configuration: scan interval in seconds *	<input type="text" value="60"/> <small>Interval during which the Port Scan Threshold is exceeded</small>
Port Scan Configuration: Unique ports threshold *	<input type="text" value="10"/>

The error also appears in the file `/opt/splunk/var/log/TA-Illumio/ta-illumio.log`.

**Cause:** This error is caused by an authentication issue to the Policy Compute Engine (PCE). When this occurs, data inputs will not be saved until a valid API response is received from the PCE with the correct API Authentication username and API secret.

**Fix:** To validate an authentication failure, look at the PCE core node haproxy logs, which will show a 401 auth failure HTTP response (highlighted in yellow in the example below):

```
Mar 11 11:20:36 level=info host=core0.domain.com
program=illumio_pce/agent[12311]:

sec=328436.974 sev=INFO pid=12389 tid=35453020 rid=92e0c733-
b06f-4619-9624-7e1dbf515eb6 XStarted

GET /api/v1/product_version/ 10.6.7.40

Mar 11 11:20:37 level=warning host=core0.domain.com
program=illumio_pce/agent[12311]:

sec=328437.741 sev=WARN pid=12389 tid=35453020 rid=92e0c733-
b06f-4619-9624-7e1dbf515eb6

{"category":"auditable","event_type":"authn_failure","severity":
"warning","timestamp":"2019-03-11T18:20:36+00:00",

"href":"/orgs/0/audit_log_events/4e28d281-7cf9-4046-97f1-
fb78060c3b4c","created_by":{"system":{}}},

"data":{"uri_path":"/api/v1/product_version/","username":"api_1f
1ec61c67e853576","src_ip":"10.6.7.40"}}

Mar 11 11:20:37 level=info host=core0.domain.com
program=illumio_pce/agent[12311]: sec=328437.745

sev=INFO pid=12389 tid=35453020 rid=92e0c733-b06f-4619-9624-
7e1dbf515eb6 XCompleted 401

GET /api/v1/product_version/ 10.6.7.40 0.099828328

Mar 11 11:20:37 level=info host=core0.domain.com
program=haproxy[2624]: 10.6.7.40:56152

[11/Mar/2019:11:20:36.969] https~ agent/agent0 3/0/0/103/106 401
304 - - ---- 1/1/0/1/0 0/0

{|keep-alive} "GET /api/v1//product_version/ HTTP/1.1"
```

To see whether the API username/secret is correct, use the cURL command below and validate it with the logs from PCE core nodes:

```
Copy and paste the curl command below with the correct API
username/secret:

- - - Begin copy (change api username/secret) - - -

curl \

-u \

api_1f1ec61c67e853576:2a0bfa6e81965e27a6ce668df8b3022c051b7a6c6b
0868c5df4b94035562f05b \

-H Content-Type:application/json \

-X GET \

'https://pcecore0.domain.com:8443/api/v1//product_version/' \

| python -mjson.tool

- - - End copy - - -
```

Successful curl request logs from PCE core nodes with 200 http response code:

```
Mar 11 11:45:44 level=info host=core0.domain.com
program=illumio_pce/agent[23340]: sec=329944.610 sev=INFO

pid=23408 tid=24064620 rid=063accb6-7036-46f0-96a1-8726f14436ea
XStarted GET /api/v1/product_version/ 10.6.7.40

Mar 11 11:45:44 level=info host=core0.domain.com
program=illumio_pce/agent[23340]: sec=329944.734 sev=INFO

pid=23408 tid=24064620 rid=063accb6-7036-46f0-96a1-8726f14436ea
XCompleted 200 GET /api/v1/product_version/ 10.6.7.40
0.124496975

Mar 11 11:45:44 level=info host=core0.domain.com
program=haproxy[2624]: 10.6.7.40:56446
[11/Mar/2019:11:45:43.966]

https~ agent/agent0 643/0/0/126/769 200 442 - - ---- 2/2/0/1/0
0/0 {115|keep-alive} "GET /api/v1//product_version/ HTTP/1.1"
```

On the Splunk server, a successful request will allow the data inputs to be saved without any errors in the PCE web console or the `/opt/splunk/var/log/TA-Illumio/ta-illumio.log` file. Tail the `ta-illumio.log` when configuring the data inputs to see the latest logs. Enabling and disabling the data inputs will trigger the request to the PCE, which is a good way to test it.

The data input information should be saved in the location below without the API username/password:

```
/opt/splunk/etc/apps/IllumioAppforSplunk/local/inputs.conf

[illumio://knpce1]
api_key_id =
api_secret =
cnt_port_scan = 10
enable_data_collection = Enabled
interval = 3600
pce_url = https://pce.domain.com:8443
port_number = 514
self_signed_cert_path = /opt/splunk/custom_certificate.cer
time_interval_port = 60
disabled = 0
```

## Sankey diagram is not displayed in Traffic Explorer dashboard

- You need to install the [Sankey Diagram App](#) in order to visualize the diagram in the “Communications Map between Labeled Workloads” panel.

## Label filters (i.e.App, Env and Loc) are not populated

- Try to run the "Illumio\_Workload\_Mapping" saved search via expanding time range.
- Make sure that interval configuration for input is less than 24 hours.

## Known Limitations

- In case of multiple input configuration, port scan will be done based on the last configured input's port scanner threshold value for all the inputs.
- In Splunk version 7.1, real-time search options will not be disabled, because of a known Splunk issue. Reference: [SPL-76798]  
<https://docs.splunk.com/Documentation/Splunk/7.1.0/ReleaseNotes/KnownIssues>
- Due to certification requirements, the TA only supports TCP for syslog.
- Due to certification requirements, data model acceleration is disabled by default. Without data model acceleration, some visualizations will not work. You can enable data model acceleration, if needed, using the steps in "[Data Model and Data Model Acceleration](#)."
- Editing data input (modular input) while it is disabled can lead to exposing the Key\_ID and the Secret of the user.

## Compatibility Matrix

PCE Versions	Splunk Version	Illumio App for Splunk Version	Illumio TA Version
17.1, 17.2, 17.3, 18.1	6, 7	1.x	1.1.3
18.2.0*, 18.2.x, 18.3, 19.1, 19.3	7	2.x	2.2.1, 2.3.0
18.2.0*, 18.2.x, 18.3, 19.1, 19.3	7, 8**	3.x	3.x

PCE Versions	Splunk Version	Illumio App for Splunk Version	Illumio TA Version
18.2.0*, 18.2.x, 18.3, 19.1, 19.3, 20.1, 21.2.x	7.3, 8.0, 8.1, 8.2	3.2	3.2

\* Special configuration is needed with version 18.2.0. Contact Illumio Support.

\*\* The Illumio App for Splunk 3.x is compatible with Python 2 and Python 3. Python 3 is available starting in Splunk 8.0.