



Monitoring the great value in your Big Data

HANDS-ON WALKTHROUGH · VOLUME OUTLIERS

Volume outliers on the license usage

One tenant, a few clicks, and every index of the deployment is watched for abnormal drops, spikes and silence — with a month of history backfilled, a cost trend and a projection, and the whole environment's license usage on one screen. From an empty tenant to 43 monitored indexes.

Captured live on TrackMe 2.4.18 · tenant volume-detection



The feed is green — and half of its volume is gone

Data source monitoring sees freshness and latency. The volume of what arrives is a different signal — and the one the license bill, the detections and the capacity plan all depend on.

A drop — the pipeline is broken

A forwarder stops, a relay loses a route, a modular input silently dies, a parsing change discards events. The feed is still alive; half of it is missing.

A spike — cost and noise

A debug level left on, a runaway application, a duplicate forwarder, a new source pointed at the wrong index. Found at month end, it is an overage; found within the hour, a ticket.

A silent index — nobody correlates


The index still exists, dashboards still run over it. Its last event is a week old and no one noticed.

What Volume Outliers does

VOL reads the license usage log — the volume Splunk counts against your license, one source for every index of the deployment — incrementally, in 5-minute buckets. Every index becomes a TrackMe entity: an ML outliers model on its rolling volume (drops and / or spikes), an inactivity policy, the daily licensed volume, a 30-day trend and a month projection. State, score, history and alerting work exactly like a data source in DSM.

Splunk Enterprise or Splunk Cloud, local or remote deployment — and part of the Foundation edition.

One source, one tracker, one entity per index



- 1 Source — the license usage log**
What Splunk bills: one log for every index of the deployment, on Enterprise and on Cloud.
- 2 Collect — one incremental tracker**
Every 5 minutes, only the new events, into 5-minute buckets per index. Never a rescan.
- 3 KPIs — per-index trends**
Rolling 60 min / 4 h / 12 h / 24 h volume and the daily licensed volume per license day, as trackme.splk.vol.* metrics.
- 4 Detect — outliers and inactivity**
One ML model per index on the rolling volume, drops and / or spikes; a silent index turns red.
- 5 Act — score, alert, cost**
Impact score to state, stateful alerts, the 30-day trend and the month projection.

The diagram is built into the wizard — expand "How it works" on the first step.

What we are going to do — one tenant, four parts

PART 1 • SETUP

Create the tenant, let it backfill

- The wizard: basics, volume, inactivity, RBAC
- 30 days of history replayed at creation
- Day one: pending → complete, models trained

PART 2 • READ

What the tenant tells you

- The VOL table: volume, yesterday, trend, projection
- An index: Volume & cost, rolling volume, hourly profile
- The outliers model and its bounds

PART 3 • TUNE

Make the detection yours

- Drops, spikes or both — per index, in bulk
- Inactivity: business hours vs nights, per index
- Default KPIs and the outliers scope


PART 4 • OPERATE

Analytics, alerting, AI

- Global license usage: figures, insights, pools
- SPLK-VOL alerts and stateful emails
- The AI ML Advisor reviews a volume model


Everything runs in one tenant created for the occasion: 43 indexes, one license pool, Splunk Enterprise. The same walkthrough applies to Splunk Cloud and to remote deployments.

A dedicated tenant — pick Splunk Volume (Outliers)



- 1

Splunk Volume (Outliers)
Outliers on the licensed volume of every index, daily volume and month projection, local or remote, history backfilled on day one.



- 2

Why VOL, in three lines
Minimal compute cost · outliers that scale safely · signal, not noise.
- 3

Tenant id — volume-detection
Immutable, checked for availability as you type.
- 4

Alias and description
What people see on the Virtual Tenants card.

Step 2 — Volume: where to collect, what to detect, how far back

Create a new TrackMe Virtual Tenant: Splunk Volume (Outliers)

✓ Basics

2 Volume

3 Inactivity

4 Indexes & RBAC

5 Review

VOL tracks the licensed volume of every index from the license usage logs, in 5-minute buckets. Outliers on the rolling 24 hours volume detect abnormal drops and / or spikes, and the licensed volume is recorded per license day with a month projection. There are no use cases to choose: pick the environment and how the detection behaves.

Target environment **1**

local

The Splunk deployment whose license usage is tracked: this instance, or a remote account.

Detect **2**

Drops and spikes

The default detection direction for every index; each index can be changed afterwards.

Backfill history (days) **3**

30 days

History replayed at creation: the rolling volume the outliers train on and the daily licensed volume, so detection and the cost trend work on day one. Capped by the _internal retention.

Outliers on these priorities **4**

critical x high x medium x low x

Indexes of these priorities are in scope for volume outliers. All priorities by default: volume outliers are the point of this component.

Outliers minimum history (days) **5**

15

Metrics history an outliers model needs before its confidence is normal: a low-confidence model does not add to the impact score. A tenant setting (the system default is 30 days); it applies to every component of this tenant and can be changed later in the tenant configuration. 1 to 365.

Outliers KPIs **6**

Rolling 24h volume x

Every index gets one outliers model per KPI selected here. The rolling 24h volume follows the daily pattern; the shorter windows react faster to a sudden drop or spike. Changeable later in Tenant Home (Manage: Outliers default KPIs).

✓ With 30 days of backfill, the outliers models reach a normal confidence on day one. **7**

> Advanced

< Back

Next >

- 1 Target environment — local**
This Splunk, or any remote account.
- 2 Detect — drops and spikes**
The default direction of every new index.
- 3 Backfill history — 30 days**
Replayed at creation, capped by retention.
- 4 Outliers on every priority**
The VOL scope: all priorities by default.
- 5 Minimum history — 15 days**
Before a model reaches normal confidence.
- 6 Outliers KPIs — rolling 24h**
One model per index, on the daily pattern.
- 7 Covered on day one**
The backfill exceeds the minimum history.

TrackMe for Splunk • Hands-on: volume outliers

docs.trackme-solutions.com

Advanced — two settings you normally never change

Create a new TrackMe Virtual Tenant: Splunk Volume (Outliers)

Backfill history (days)

30 days

History replayed at creation: the rolling volume the outliers train on and the daily licensed volume, so detection and the cost trend work on day one. Capped by the _internal retention.

Outliers on these priorities

critical X high X medium X low X

Indexes of these priorities are in scope for volume outliers. All priorities by default: volume outliers are the point of this component.

Outliers minimum history (days)

15

Metrics history an outliers model needs before its confidence is normal: a low-confidence model does not add to the impact score. A tenant setting (the system default is 30 days); it applies to every component of this tenant and can be changed later in the tenant configuration. 1 to 365.

Outliers KPIs

Rolling 24h volume X

Every index gets one outliers model per KPI selected here. The rolling 24h volume follows the daily pattern; the shorter windows react faster to a sudden drop or spike. Changeable later in Tenant Home (Manage: Outliers default KPIs).

With 30 days of backfill, the outliers models reach a normal confidence on day one.

Advanced

Settle margin (seconds)

300

License usage events newer than now minus this margin are not read yet, so late events are counted. 60 to 3600, default 300.

License search constraint

The base search giving access to the license usage events: every license search of the tracker is formed from it. Indexed fields only (index, source, sourcetype, host). Change it only to target another data set with the license_usage.log format, whose key=value fields (type, idx, b, pool) are extracted at search time. Default: index=_internal source="*/license_usage.log"

< Back

Next >

1 Settle margin — 300 s

Events newer than now minus the margin wait for the next run, so late license events are counted. 60 to 3600.

2 License search constraint

The base search every license search of the tracker is built from — indexed fields only.

3 The default

index=_internal source="*/license_usage.log". Change it only to read a copy of the license usage events.

Leave the defaults

Both settings exist for the rare deployment that forwards its license usage events elsewhere, or delivers them late. Everything else in VOL is tuned per index after creation.

Step 3 — Inactivity: when a silent index turns red

Create a new TrackMe Virtual Tenant: Splunk Volume (Outliers)

An index whose last licensed volume is older than the threshold in force turns red. This global policy applies to every index without its own; each index can override it later (entity menu or bulk edit).

Inactivity policy

Variable (per day and hour)

1

Timezone notice

Threshold slot days and hours are evaluated in the Splunk server's local time, not your browser's time.

Splunk server: Mon 28 Sep 2026, 19:40:34 (UTC, +00:00) Your browser: Mon 28 Sep 2026, 20:40:34 (BST, +01:00)

Your browser is 1h ahead of the server. The hours below are shown in your local browser time — TrackMe converts them to the server's time automatically on save. Weekday labels stay on the server's calendar, so a slot whose hours cross midnight keeps the server's weekday.

Quick templates

Working hours vs. off-hours

Weekdays vs. week-ends

Three-tier (working hours / nights / week-ends)

Time slots (1)

+ Add slot

The first slot matching the current day and hour applies. Thresholds accept seconds or a unit suffix (30m, 12h, 1d, 1w); 0 disables the check in that slot.

working_hours

1d

Remove

Days:

Select all

Mon

Tue

Wed

Thu

Fri

Sat

Sun

Hours:

Select all

0

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

Outside the slots:

7d

< Back

Next >

1

Variable (per day and hour)
Or static: one threshold at all times.

2

Timezone notice
Slots run in the server's time; the editor converts your browser hours on save.

3

Quick templates
Working vs off-hours, weekdays vs week-ends, three tiers.

4

working_hours — 1d
Mon–Fri 08:00–19:59 server time: a silent index is red the next day.

5

Outside the slots — 7d
Nights and week-ends wait longer.

Tenant-wide, overridable

The policy applies live to every index without its own threshold. Any index can override it later — from its menu, the bulk edit, or a click in the table.

TrackMe for Splunk • Hands-on: volume outliers

docs.trackme-solutions.com

Steps 4-5 — Indexes and RBAC, review, created in seconds

Create a new TrackMe Virtual Tenant: Splunk Volume (Outliers)

Basics

Volume

Inactivity

Indexes & RBAC

Review

Choose indexes for TrackMe charts and assign ownership and roles.

Summary index

trackme_summary

Audit index

trackme_audit

Metric index

trackme_metrics

Notable index

trackme_notable

Splunk owner for created objects

nobody

Admin roles

trackme_admin

Power roles

trackme_power

Readonly roles

trackme_user

< Back

Next >

- 1

TrackMe indexes

Summary, audit, metric and notable — the defaults are fine.
- 2

Splunk owner — nobody

A service account instead if the tenant will run AI automation.
- 3

Roles

trackme_admin, trackme_power, trackme_user — or your own roles.

Create a new TrackMe Virtual Tenant: Splunk Volume (Outliers)

Basics

Volume

Inactivity

Indexes & RBAC

Review

Please review the configuration before confirmation.

Update comment (optional)

Add a note for this creation

Action Progress Events (auto-refresh 1s)

Time	Event
9/28/2026 8:41:44.531 PM	2026-09-28 19:41:44,531 INFO trackme.rest.vtenants_admin trackme_rest_handler_vtenants_admin.py post_add_tenant 5049 tenant_id="volume-detection", create vtenant account was operated successfully, response.status_code="201" index = _internal source = /opt/splunk/var/log/splunk/trackme_rest_api_vtenants_admin.log sourcetype = trackme_rest_api
9/28/2026 8:41:42.906 PM	2026-09-28 19:41:42,906 INFO trackme.rest.vtenants_admin trackme_libs.py trackme_create_report 3159 tenant_id="volume-detection", action="success", report_name="trackme_health_tracker_tenant_volume-detection" with kwargs={"owner": "nobody", "sharing": "app", "perms.write": "trackme_admin,trackme_power", "perms.read": "trackme_admin,trackme_power,trackme_user" }

< Back

Close >

- 4


Action progress events

The creation, live, as it happens.
- 5

Tenant created

The account, its reports and trackers — in seconds.

Created — 43 indexes collected at once, the backfill pending



- 1 43 enabled entities
One per index, from the first collection at creation.
- 2 Backfill pending
The card follows the backfill until it completes.
- 3 The Tenant Home vignette
Pending, then the percentage — a click opens the volume collection screen.
- 4 All green from the start
Every index had licensed volume in the last 24 hours.
- 5 Volume 24h — already live
The seed reads the last 24 hours of the license usage log.
- 6 The cost columns wait
Licensed yesterday, average, trend and projection need the replayed history.

A month of history, replayed at the first hourly run

Volume tracker & backfill

1

Live collection

Every 5 minutes the tracker reads the license usage logs incrementally (5-minute buckets behind the settle margin) and derives the rolling volumes and the daily licensed volume.

Collected up to

28/09/2026, 20:40:00

Validated history since

27/09/2026, 20:35:00

License manager timezone

UTC+0000

Deployment

Splunk Enterprise (pool usage reported)

Environment

local

Run tracker now

2

Backfill

The backfill replays the license usage history chosen at creation: the rolling volume the outliers models train on, and the daily licensed volume. It runs hourly in day chunks until it reaches the start of the live collection, then triggers the models training.

3

Pending: a 30-day backfill was requested at creation. The backfill job plans and starts it at its next scheduled run (hourly, next at 2026-09-28 21:17:00 BST); its progress shows here then.

Refresh

Close

Volume tracker & backfill

Live collection

Every 5 minutes the tracker reads the license usage logs incrementally (5-minute buckets behind the settle margin) and derives the rolling volumes and the daily licensed volume.

Collected up to

28/09/2026, 21:10:00

Validated history since

27/09/2026, 20:35:00

License manager timezone

UTC+0000

Deployment

Splunk Enterprise (pool usage reported)

Environment

local

Run tracker now

Backfill

The backfill replays the license usage history chosen at creation: the rolling volume the outliers models train on, and the daily licensed volume. It runs hourly in day chunks until it reaches the start of the live collection, then triggers the models training.

Complete

4

5

100.0%

History request

30 days

Window

31/08/2026, 01:00:00 → 27/09/2026, 20:35:00

Daily volumes replayed (index × day)

1247

Metric points written

358104

First live day trends

pending: replayed once the live history has a full day

Pause

Resume

Cancel

Restart

6

Refresh

Close

Left: right after creation. Right: the same screen once the backfill has run.

1

Live collection

Every 5 minutes, incremental, behind the settle margin.

2

Run tracker now

Collect on demand.

3

Pending

Planned at the next hourly run of the backfill job.

4

Complete — 100%

Replayed, then every model trained.

5

What was replayed

30 days: 1,247 index-days, 358,104 metric points.

6


Controls

Pause, resume, cancel; restart extends history backward.

TrackMe for Splunk · Hands-on: volume outliers

docs.trackme-solutions.com

The Tenant Home, complete on day one



- 1 Two tabs**
The indexes, and the global license usage.
- 2 Header counters**
Entities and green; the vignette is gone.
- 3 By priority and by state**
Every index medium and green.
- 4 Key counts**
43 enabled, none in alert.
- 5 Entities overview**
Grouped by license pool.
- 6 Cost columns filled in**
Yesterday, average, trend, projection.

01

PART 2 • READ

What the tenant tells you

One row per index with its volume, its trend and its projection; one screen per index with the charts behind the figures; and the outliers model that turns a deviation into a state.

One row per index — the volume, the cost and the detection

Entities Overview

Q Search entity, object or tags

Default

All entities

100 / page

Infinite scroll

Showing 1-43 of 43

Table loaded in 0.6s


Load more

Actions	Entity	Status	Score	Labels	Notes	Priority	Enab...	De	2	3	License...	Avg	4	Trend ...	5	Inac	6	7
▼ auto_generated_pool_enterprise (43 items)																		
>	business_data		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	49.52 GB	52.83 GB	51.90 GB	↗ +1.9%	1.53 TB	35d			
>	cribl_datagen		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	4.11 GB	4.30 GB	4.36 GB	↗ -1.1%	130.7 GB	7d			
>	cribl_signals		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	5.36 GB	5.44 GB	3.94 GB	↗ +113.6%	121.6 GB	7d			
>	em_metrics		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	24.41 GB	24.99 GB	26.02 GB	↗ -2.0%	781.0 GB	7d			
>	prod_trackme		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	17.36 GB	17.38 GB	17.39 GB	↗ -0.0%	522.7 GB	7d			
>	siem-cloud-amer		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	32.70 GB	31.12 GB	33.02 GB	↗ +3.1%	1003.3 GB	7d			
>	siem-cloud-apac		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	13.30 GB	11.60 GB	14.40 GB	↗ -4.4%	436.1 GB	7d			
>	siem-cloud-emea		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	10.61 GB	9.60 GB	14.62 GB	↘ -86.0%	437.7 GB	7d			
>	siem-cloud-latam		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	5.39 GB	4.43 GB	5.19 GB	↗ +0.8%	157.9 GB	7d			
>	siem-firewall-amer		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	5.57 GB	5.00 GB	5.66 GB	↗ +3.0%	171.5 GB	7d			
>	siem-firewall-apac		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	5.12 GB	4.07 GB	2.91 GB	↗ +52.6%	89.66 GB	7d			
>	siem-firewall-emea		0	+ label		medium			↓ Drops ↑ Both ↑ Spikes	2.56 GB	2.40 GB	2.55 GB	↗ +7.7%	107.3 GB	7d			

- 1 Grouped by license pool
The pool is the entity group.
- 2 Detect — one click
Drops, both or spikes, in place.
- 3 Volume 24h • licensed
Rolling volume, last license day.
- 4 Avg daily • trend 30d
Baseline and direction, coloured.
- 5 Month projection
At the weekday pattern.
- 6 Inactivity
The threshold in force: 7d, 35d.
- 7 Outliers
The model state of the index.

em_metrics detects spikes only; business_data has its own 35-day inactivity threshold.

Volume & cost — the figures and the daily licensed volume



1 The header
Volume, trend, month, inactivity.

2 Key figures
24h, yesterday, average, trend, month to date, projection, pool used.


3 Detect
Drops only here, applied at once.

4 Daily licensed volume
One bar per license day, 60 days.

5 Projection
End of this month or of next.


6 Projected days
Lighter bars, weekday pattern.

Rolling volume, the outliers model and the hourly profile



- 1 Rolling volume — 7 days**
The KPIs a model can train on.
- 2 24h, 12h, 4h, 60m**
One unit per chart, here GB.
- 3 Outliers — rolling 24h**
The model, its direction, anomalies.
- 4 The learned band**
Outside it: a candidate outlier.
- 5 Hourly profile**
Business hours peak, nights low.

The outliers model — trained on day one, normal confidence



- 1 The model**
Rolling 24h volume, one per KPI.
- 2 Time range**
The last 7 days of history.
- 3 AI ML Advisor · Manage**
Review it, or tune it by hand.
- 4 Impact score**
What an outlier adds to the score.
- 5 Outliers by bound**
Accepted, rejected, corrected.
- 6 Value against the band**
The KPI and its thresholds.

02

PART 3 • TUNE

Make the detection yours

Which direction matters for each index, when silence becomes an incident, which KPIs carry a model, which indexes are in scope — all live, per index or in bulk.

Drops, spikes or both — per index, in one click or in bulk

Entities Overview

Search entity, object or tags

Default All entities

100 / page Infinite scroll Showing 1-43 of 43 Table loaded in 0.6s Load more

Actions	Entity	Status	Score	Labels	Notes	Priority	Enab...	Detect	Volum...	License...	Avg dal...	Trend ...	Month ...	Inactl...	Outl...
auto_generated_pool_enterprise (43 items)															
> [icon] [icon] [icon]	business_data	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	2 GB	52.83 GB	51.90 GB	++1.9%	1.53 TB	35d	●
> [icon] [icon] [icon]	cribl_datagen	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	4.11 GB	4.30 GB	4.36 GB	+-1.1%	130.7 GB	7d	●
> [icon] [icon] [icon]	cribl_signals	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	5.36 GB	5.44 GB	3.94 GB	↑113.6%	121.6 GB	7d	●
> [icon] [icon] [icon]	em_metrics	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	1 GB	24.99 GB	26.02 GB	+-2.0%	781.0 GB	7d	●
> [icon] [icon] [icon]	prod_trackme	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	17.36 GB	17.38 GB	17.39 GB	+-0.0%	522.7 GB	7d	●
> [icon] [icon] [icon]	siem-cloud-amer	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	32.70 GB	31.12 GB	33.02 GB	++3.1%	1003.3 GB	7d	●
> [icon] [icon] [icon]	siem-cloud-apac	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	13.30 GB	11.60 GB	14.40 GB	+-4.4%	436.1 GB	7d	●
> [icon] [icon] [icon]	siem-cloud-emea	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	10.61 GB	9.60 GB	14.62 GB	↓86.0%	437.7 GB	7d	●
> [icon] [icon] [icon]	siem-cloud-latam	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	5.39 GB	4.43 GB	5.19 GB	++0.8%	157.9 GB	7d	●
> [icon] [icon] [icon]	siem-firewall-amer	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	5.57 GB	5.00 GB	5.66 GB	+-3.0%	171.5 GB	7d	●
> [icon] [icon] [icon]	siem-firewall-apac	●	0	+ label	[icon]	medium	●	↓ Drops ↑ Both ↑ Spikes	5.12 GB	4.07 GB	2.91 GB	↑52.6%	89.66 GB	7d	●

Bulk edit entities (4 selected)

Category

Volume Detection Direction

Action

Set detection direction

Configuration

Detect *

Which abnormal licensed volume changes raise an outlier for the selected indexes.

Drops only (lower bound)

Update Comment (Optional)

Security indexes: a drop is the concern, a spike is welcome data

Enter a note for this update...

Cancel Apply

1 Both — the default

business_data: a drop and a spike both matter.

2 Spikes only

em_metrics: a drop costs nothing, a spike costs license.

Drops only

Drops only fits security indexes: a drop is the incident. Four siem-wineventlog indexes are switched at once on the right.

3 Volume Detection Direction

Category, then Set detection direction.

4 Drops only (lower bound)

Applied to the 4 selected indexes.

5 Update comment

Why — kept in the audit trail.

When silence becomes an incident — tenant policy, index override

Global inactivity threshold

Inactivity policy

Variable (per day and hour)

Timezone notice

Threshold slot days and hours are evaluated in the Splunk server's local time, not your browser's time.
Splunk server: Mon 28 Sep 2026, 19:52:57 (UTC, +00:00) Your browser: Mon 28 Sep 2026, 20:52:58 (BST, +01:00)
Your browser is 1h ahead of the server. The hours below are shown in your local browser time — TrackMe converts them to the server's time automatically on save. Weekday labels stay on the server's calendar, so a slot whose hours cross midnight keeps the server's weekday.

Quick templates

Working hours vs. off-hours Weekdays vs. week-ends Three-tier (working hours / nights / week-ends)

Time slots (1)

The first slot matching the current day and hour applies. Thresholds accept seconds or a unit suffix (30m, 12h, 1d, 1w); 0 disables the check in that slot.

working_hours

1d

Remove

Days: Select all

Mon Tue Wed Thu Fri Sat Sun

Hours: Select all

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23

Outside the slots:

7d

Comment

Cancel Save

Inactivity threshold — business_data

An index whose last licensed volume is older than the threshold in force turns red. Set a static or variable threshold for this selection, or let it follow the tenant global policy.

In force now: 35d — set for this index

Inactivity policy

Static

Threshold

35d

35d — an index whose last licensed volume is older than this turns red. 0 disables the check.

Comment

Optional audit comment

Cancel Save

- 1 Variable (per day and hour)
Or static: one threshold.
- 2 Server time zone
Your browser's hours are converted on save.
- 3 Quick templates
Three common shapes.
- 4 working_hours — 1d
Mon–Fri, business hours.
- 5 Outside the slots — 7d
Nights and week-ends.

- 6 In force now: 35d, set for this index
The table and the status message name the slot and its source.
- 7 Static override
Or variable, or back to the tenant policy.
- 8 35d
s, m, h, d, w units or [D+]HH:MM:SS.

Which KPIs carry a model, which indexes are in scope

Manage: Outliers default KPIs

Every index gets one outliers model per KPI selected here. The rolling 24h volume follows the daily pattern of an index; the shorter windows react faster to a sudden drop or spike, at the cost of more models to train and evaluate.

Default KPIs

Rolling 24h volume X

Existing indexes

☐ Also apply to the existing indexes

Comment

Recorded on the audit event of this change

Cancel

Save

1

Default KPIs
Rolling 24h, 12h, 4h, 60m — one model per KPI per index.

2

Also apply to the existing indexes
Adds the new models and removes the deselected ones; never touches the others.

Manage: ML Outliers tenant scope

Configure the Volume Outliers scope: only the indexes whose priority is selected here, and that match the optional filter expression, are trained and evaluated. VOL has its own scope: the tenant-wide filters of the other components do not apply; every index is in scope by default.

Priority filter

☒ Critical

☒ High

☒ Medium

☒ Low

Filter expression (optional)

Syntax: field=value — wildcards supported (* any sequence, ? one char). Space = implicit AND. Use OR for alternatives. Parentheses for grouping. All comparisons are case-insensitive.

tags=prod
labels=critical OR labels=billing
object=siem-* NOT tags=lab

Cancel

Save

3

Every priority
Volume outliers are the point.

4

Filter expression
object=siem-* NOT tags=lab ...

Independent
The tenant-wide ML filters of the other components do not apply to VOL: its scope is its own.


03

PART 4 • OPERATE

The whole environment's license, on one screen


Global license usage (VOL) reads what the tenant already collected — no monitoring, no alerting, no extra search on the license usage log — and answers the questions every Splunk administrator gets asked.

Key figures and insights, in plain words




- 1 **Period**
30, 60 or 90 closed license days.
- 2 **Yesterday · average · peak**
What the license is measured on.
- 3 **Trend**
The slope of the environment.
- 4 **Month to date · projection**
The sum of the per-index figures.
- 5 **Quota used yesterday**
Against the pool quota.
- 6 **Insights**
Stable or growing, movers, quota.
- 7 **Where the license goes**
One index is 15%, the top 5 53%.

Who consumes the license, day by day — and where the month ends





- 1 **Top consumers slider**
5 to 100 series, the rest together.
- 2 **The daily quota line**
Every pool's quota together.
- 3 **Stacked per index**
Each license day, by consumer.
- 4 **The whole environment**
Projected to the end of the month.
- 5 **Projected days**
Lighter bars, weekday pattern.

Consumers, movers and pools




Alerting — the SPLK-VOL alert and stateful emails



- 1 SPLK-VOL (Volume Outliers)**
The component alert: notables and acknowledgments.
- 2 volume-alert**
Defaults: any priority, red, all indexes, outliers on.

- 3 Emails and Ingest**
Emails plus stateful events.
- 4 Delivery account**
Any email account in TrackMe.
- 5 Environment name**
Stamped in the email header.
- 6 Recipients**
Charts and AI report further down.

The AI ML Advisor reviews an index's model



1 AI provider
Any provider configured in TrackMe.

2 Inspect or act
Review, or remediate after consent.

3 Additional instructions
The operator's intent, in words.

4 Start Analysis
A background agent job.

5 Live progress
Each tool call streamed as it runs: outlier context, render history, score history, training details.

The verdict — the model already fits the index

AI ML Advisor

×

Entity: **siem-wineventlog-amer** (VOL)

responded in 1m 3s

Healthy 1

Summary 2

Recommendations 3

low

No Action Needed

No anomaly detected, no false positive to clear, and model configuration already matches the entity's observed seasonality — this is a genuinely healthy state, not a 'metric recovered' rationalization (calibration check: this is 'no action because no true-positive/false-positive exists', not 'no action because it recovered').

low

Configuration Change

Detection direction is already 'lower' (drops only) as the operator requires for this security Windows event log index — no change needed here, confirming operator's monitoring intent is already correctly configured. Listed for completeness/verification only.

> Reasoning Trace

Close

1

Healthy

Green, no outlier and no false positive in the last 30 days.

2

Summary

The model uses weekday × hour seasonality — 168 groups, all fitted — so the week-end lows are expected, not drops.

3

Recommendations — two, low

No action needed; the drops-only direction confirmed, as the operator asked.

Inspect first, act on consent

Act mode would apply period exclusions, retraining or bound changes after consent. The automated batch runs the same review nightly over the in-scope indexes.

Cheap by design — what VOL costs to run

Collection

One bounded search per tenant every 5 minutes, only the new license usage events — whatever the number of indexes. No per-index, no per-sourcetype scan.

Metrics

About 216 KB per index per day of metric-index licensing; the backfill writes the same series over the chosen depth (the wizard shows the estimate).

Models

One model per index on the rolling 24 hours, trained on the standard schedule; the VOL scope and the default KPIs control how many.

Integrity

Live and backfill never overlap; no point is written twice; a failed search never advances the cursor, so missing volume is never recorded as a drop.

The same component serves Splunk Enterprise and Splunk Cloud, on this deployment or on a remote account.

Where we are

In one walkthrough

● Created

a Volume Outliers tenant covering 43 indexes, with a month of history backfilled at creation

● Read

the volume, the cost and the detection of every index, and the model behind each state

● Tuned

the direction per index, the inactivity policy with an override, the KPIs and the scope

● Operated

the global license usage view, SPLK-VOL alerting with emails, and the AI Advisors

What made it work

The right source

The license usage log: the volume Splunk bills, one log for every index, read incrementally.

Useful on day one

The backfill replays the history, so the models are trusted and the trend is complete immediately.

Signal, not noise

One entity per index, a direction per index, an inactivity policy by business hours — and impact scoring.