PRODUCT BROCHURE

TrackMe for Splunk

Visibility & operational excellence

TrackMe discovers, maintains and monitors the availability and quality of your data — feeds, hosts, licensed volume, fields, scheduled workload, custom objects and user activity — across one Splunk deployment or many, at any scale. It runs entirely inside your Splunk environment.

Eight components — grouped by the question each one answers

Is the data arriving?	Is it well-formed?	Are searches healthy?	Is my license volume normal?	Anything else?	Who is doing what?
DSM Data Source Monitoring — index × sourcetype feeds	**FQM** Field Quality Monitoring — parsing quality & CIM	**WLK** Workload Knowledge — scheduled-search health	**VOL** Volume Outliers — licensed volume per index	**FLX** Flex Objects — any SPL → tracked entity	**UAM** BETA User Activity Monitoring — accounts, schedules, activity
DHM Data Host Monitoring — host × sourcetype activity	Is every field parsed and CIM-compliant? Coverage & success measured per field.	Are searches skipping, erroring, or running late? Execution health as first-class entities.	Is an index dropping, spiking, or going silent? ML outliers, inactivity, trend & month projection.	Everything else — you write the SPL; it becomes an entity with custom KPIs & dynamic thresholds.	Which accounts run what, who owns which schedule, who eats the search CPU? Findings, not entities.
MHM Metric Host Monitoring — hosts sending metric data					
delay · latency · volume					

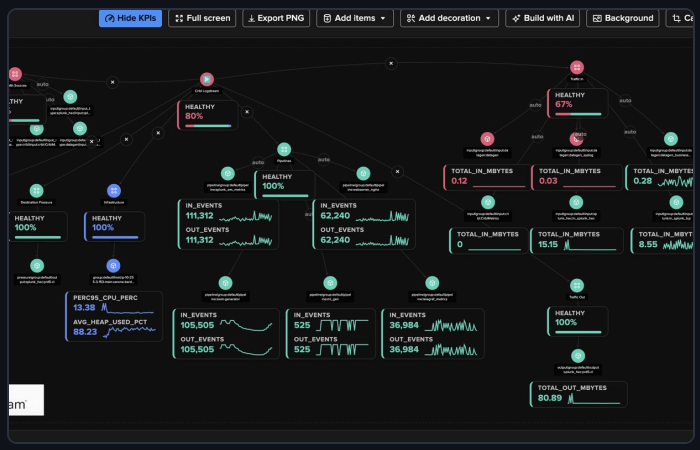

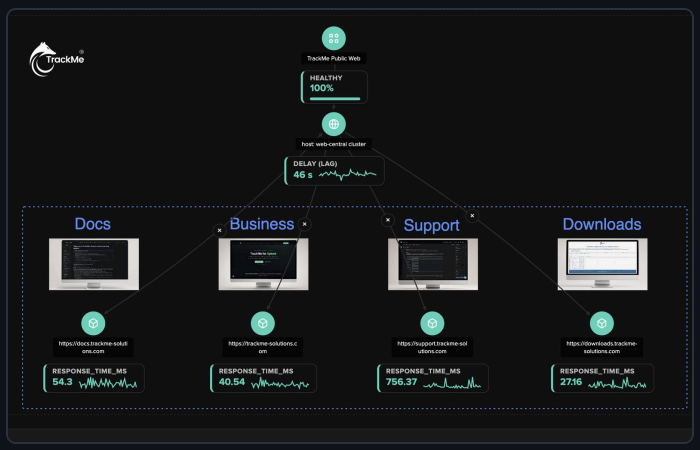

Topology Studio — your services as live maps built from real TrackMe entities: a Cribl Stream pipeline, a public web estate.

TrackMe Limited
London, United Kingdom
trackme-solutions.com · docs.trackme-solutions.com

PRODUCT OVERVIEW

Splunk is mission-critical. TrackMe makes sure it stays healthy.

In a large Splunk environment, data silently breaks. A forwarder stops while the index keeps ingesting from other hosts. A sourcetype arrives late and detections run on a window that no longer contains their events. A field stops parsing and correlation searches quietly return less. A scheduled search skips and reports drift. Splunk itself will not tell you — the searches simply return less, or nothing, until a detection misses or a dashboard goes blank.

TrackMe is a control plane on top of Splunk: it continuously watches raw index activity and turns it into something you can operate on — per-entity health states, impact scores, policies, stateful alerts, live topology maps and, optionally, AI-assisted investigation. Everything runs inside your Splunk environment: on-premises, Splunk Cloud or hybrid, with no outbound SaaS dependency.

THE QUESTIONS CLASSIC SPLUNK MONITORING LEAVES OPEN

Is my data still flowing?

Every data source, host, metric feed and licensed index is tracked for delay, latency and volume.

Is my data still good?

Field parsing quality and CIM compliance are measured per field, not assumed.

Are my searches healthy?

Scheduled-search execution, skipping, errors and cost are first-class entities.

Who do I tell, and how urgently?

Priority, SLA and impact scoring decide what deserves an alert and what is noise.

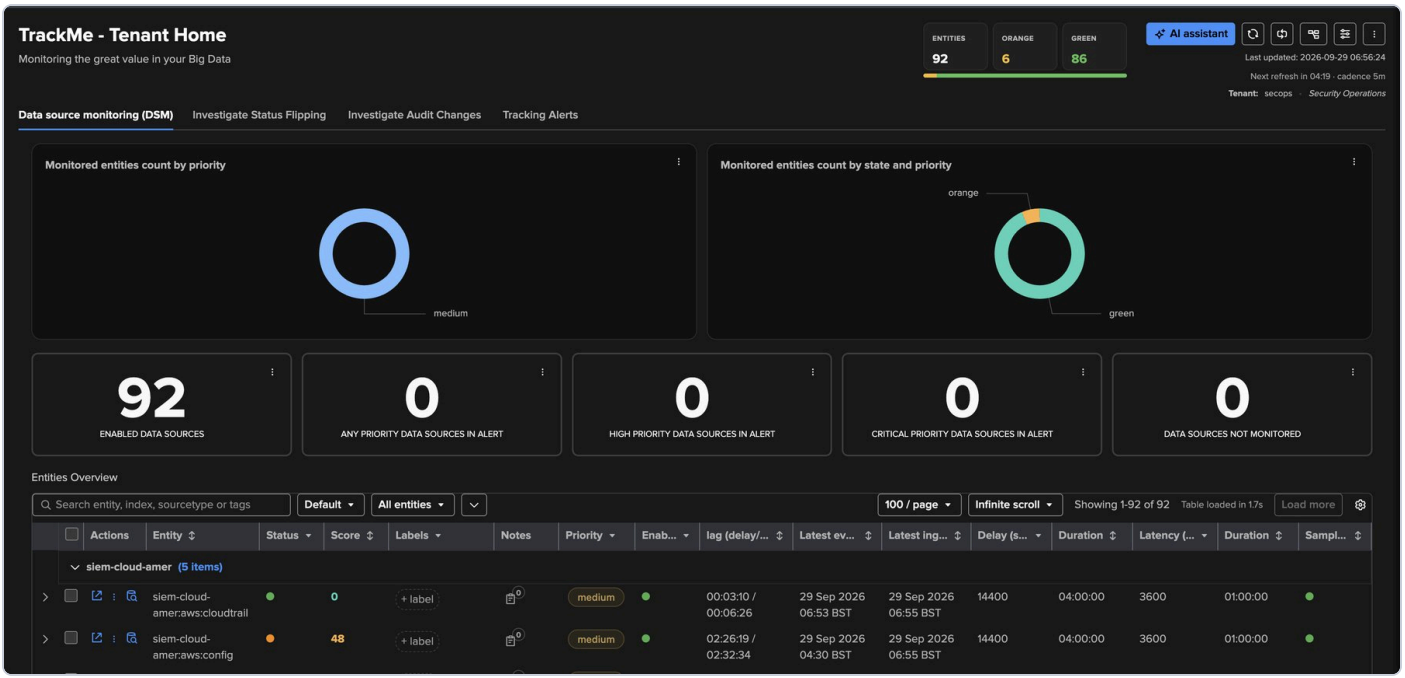

Tenant Home — every entity of a tenant with its state, score, priority, thresholds and last activity; every number on the page is a filter.

Runs inside Splunk

No outbound SaaS dependency. Suitable for the most sensitive and air-gapped environments. Delivered through Splunkbase and the Splunk Cloud vetting process.

Built for scale

Multi-tenant, multi-environment, hybrid-aware. In production at organisations ingesting above 100 TB per day.

AI alongside the cycle

An assistant, specialist advisors, scheduled agents and plain-English routines — opt-in, inert until you configure a provider, audited end to end.

COVERAGE

Eight components, one consistent model

Seven components track entities with the same shape — an object, a state, a score, thresholds and metadata — and the same lifecycle: discovered by a tracker, measured every cycle, evaluated by a shared decision maker, alerted on. The eighth, User Activity Monitoring, watches the people and the automation behind the platform and raises findings. Learn one component and you have half-learned them all.

DSM Data Source Monitoring

Index × sourcetype feeds, auto-discovered. Delay, latency, volume and data quality, with variable delay by day and hour, lagging classes, data sampling and expected-source injection from a lookup or CMDB.

DHM Data Host Monitoring

Hosts sending data, per sourcetype. Knows when a single machine goes silent while the index still looks busy; logical groups protect a fleet from one flaky member.

MHM Metric Host Monitoring

Hosts sending metric data, with lag per metric category — the same model for metrics as for events.

VOL Volume Outliers

Licensed volume per index: rolling 60-minute, 4-hour and 24-hour volume, daily licensed volume and month projection, ML outliers on drops and spikes — with history backfilled on day one.

FQM Field Quality Monitoring

Field extraction success and CIM compliance, measured per field. Dictionaries define what “good” means — CIM-aware for data models, AI-generated for custom sourcetypes.

WLK Workload Knowledge

Scheduled searches as entities: skip rate, errors, execution delay, versioning, orphan detection and SVC cost — the scheduler debt every large estate accumulates silently.

FLX Flex Objects

Anything else. You write the SPL; it becomes an entity with custom KPIs, inactivity detection and dynamic, time-sliceable thresholds. A catalogue of 60+ templates covers Splunk infrastructure, licence, SOAR, Cribl, host OS, HTTP heartbeats and more.

UAM User Activity Monitoring BETA

Who is doing what on your Splunk. A dedicated tenant type that watches every account of a deployment — people and service accounts — through the audit trail and the introspection: it inventories schedules, users and roles, classifies accounts as human, service or system, and raises findings from a policy catalogue — personal schedules, all-time searches, launch bursts, resource hogs — each with evidence, a suggested action and an investigation workflow.

VOL — rolling volume, outlier bands and hourly profile of an index.

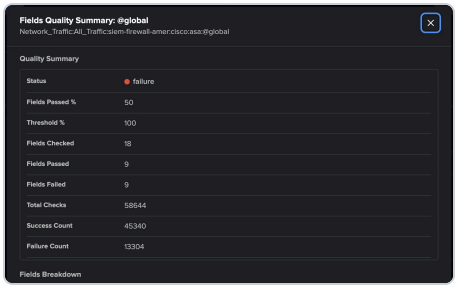

FQM — quality summary of a data model: fields checked, passed, failed.

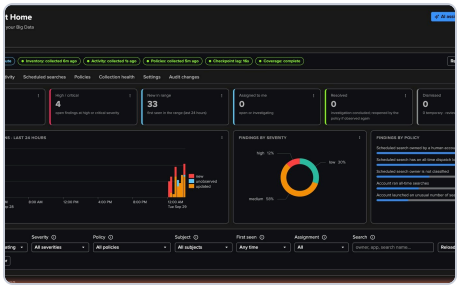

UAM — the findings backlog by severity, policy and subject.

HOW IT DECIDES — HOW IT TELLS YOU

Explainable state. One incident per problem.

TrackMe is not a live dashboard that queries Splunk when you open it. Trackers run on a cycle — typically every five minutes — and persist their results; a shared decision maker gives every entity a state and a score, policies classify it, protection layers decide whether to wake you up, and stateful alerting turns critical states into incidents.

1 · Discover	2 · Measure	3 · Decide	4 · Classify	5 · Protect	6 · Alert
trackers find entities — feeds, hosts, searches, indexes	delay, latency, volume, quality, execution health	impact score → green, orange, red	policies assign priority, tags, labels, SLA	groups, grace, maintenance, ack → blue	incidents opened, updated, closed

Entity state and impact score

- **Green** — no active anomaly
- **Orange** — a non-critical anomaly, an early signal
- **Red** — the critical threshold is crossed, the alert fires
- **Blue** — protected: an anomaly may exist, nobody is paged

```
total_score = base + delay + latency + outliers +
thresholds + …
0 → green  ·  < 100 → orange  ·  ≥ 100 → red
```

Every contribution is a configurable weight, per tenant and per entity. Nothing changes state without a reason string and a score breakdown — the same text the alert email and the AI Assistant use.

Classify what matters

Priority, tags, labels and SLA class are assigned by **policies** — a regex, a Splunk lookup such as a CMDB export, or an SPL search — simulated before they are saved and re-applied every cycle. Policies suggest; a manual override always wins.

Protect against noise

Logical groups judge a fleet collectively. The **disruption queue** holds a new anomaly until it has persisted. **Maintenance mode** — global, per tenant or per entity — and **bank holidays** silence what is expected. **Acknowledgments** record that someone owns the problem.

Stateful alerting

Instead of re-firing every cycle, TrackMe keeps **one incident record per problem**: opened when the entity enters an alerting state, updated while it persists, closed when it recovers — with the duration measured. Delivery is any combination of:

- **Email** — threaded per incident, with 24-hour charts embedded and an optional AI-written status report.
- **Notable events** for Enterprise Security or a custom correlation workflow.
- **Ingest** — an event written to a Splunk index, ready for dashboards and ITSM.
- **Commands** — run a saved search or an active command: Splunk SOAR, ServiceNow, XSOAR and more.

Alerts are scoped by priority — alert on critical and high, watch the rest — so classification is what routes the notification. Auto-acknowledgment, suppression and orange-as-alerting are per alert.

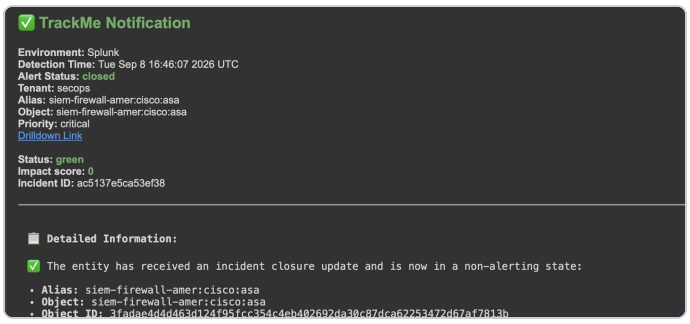

An incident email: environment, entity, priority, status, impact score and the flip message, verbatim.

ML outliers

Native density models learn each entity's normal behaviour — volume, latency, custom KPIs — with weekly and hourly seasonality, confidence levels and automated retraining. An AI ML Advisor inspects the models and, with consent, tunes them.

TOPOLOGY STUDIO

Your environment, as a live map

Entities tell you what is broken. A map tells you what it breaks. Topology Studio is the canvas for *your* view of the environment — a payment flow, a data centre, an ingestion pipeline, a business service — built from real TrackMe entities, so it cannot drift from what is monitored, and coloured with live health.

A Cribl Stream pipeline mapped from TrackMe entities: worker groups, sources, routes and destinations, with events in and out as sparkline KPIs.

Generated in a few clicks

From any tenant, a live map by index, tag, priority, label — or your own hierarchy of fields — narrowed with a condition and saved as a view.

Authored by you

Entities, aggregates that keep membership dynamic, auto aggregates, view nodes for a map of maps, labels, images and shapes. Members draw themselves, with KPIs and sparklines.

Built with AI

Describe the map in a sentence; the AI builder proposes a live draft you refine in plain words. Nothing is saved until you create the view.

Alerting-capable

Topology Alerts evaluate the whole map on a schedule and notify on transitions only — opened, updated, resolved — with the rendered map in the email.

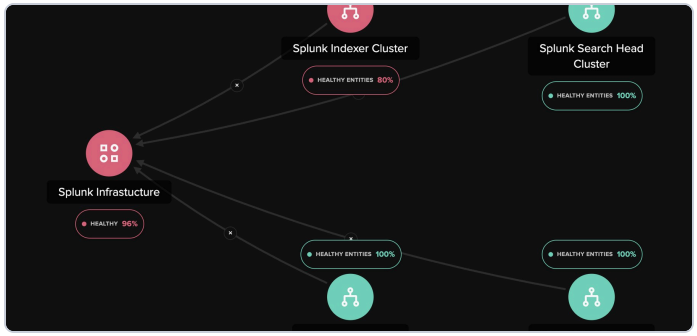

A map of maps: each node is another saved view — indexer cluster, search head cluster, forwarders — drilled into in one click, down to the reason an entity is red.

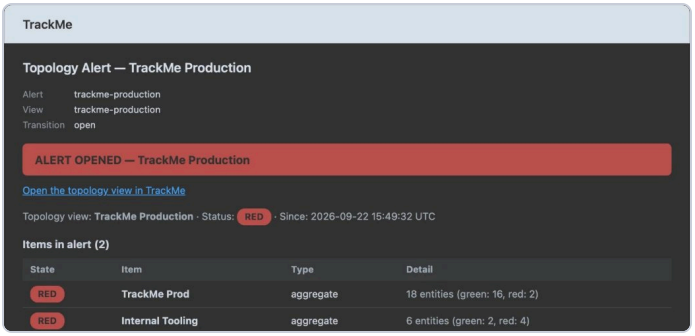

The alert email: which aggregates are in alert, how many entities, and the topology rendered server-side with the red where the problem is.

Views are shared under RBAC with tenant masking enforced per node on the server — one map for everyone, each viewer seeing only their own tenants. Every save is a version with non-destructive restore; a full-screen kiosk mode and PNG export serve the NOC wall; a Configuration Guardian check flags maps that reference deleted tenants or components. Available in every edition.

TRACKME AGENTIC

AI alongside the monitoring cycle

The AI layer helps operators understand and resolve what TrackMe detects, grounding every answer in the live state of the environment and keeping the operator in the loop. Four levels, built on one runtime, one provider abstraction and one audit trail.

LEVEL 1
AI Assistant
A context-aware chat on every entity, page and feature. Explains a state, writes the SPL to investigate, and can author the AI status report in alert emails.

LEVEL 2
AI Advisors & Concierge
Six specialists that inspect and, with consent, act: ML models, feed lifecycle, FLX thresholds, field-quality dictionaries, component health, UAM findings — plus a catalogue-driven generalist.

LEVEL 3
AI Agents automation
The same advisors running unattended on a nightly schedule, scoped by per-tenant priority and filter policies, within guardrails you define.

LEVEL 4
AI Routines
Your own standing instructions in natural language: describe what to watch for, TrackMe evaluates it on your schedule, then notifies or acts under the consent you granted.

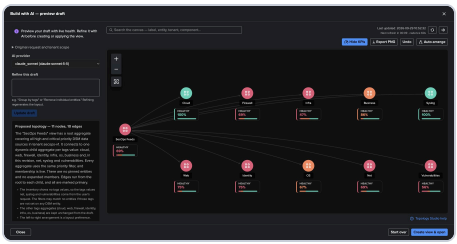

Build with AI — a topology drafted from one sentence, previewed with live health.

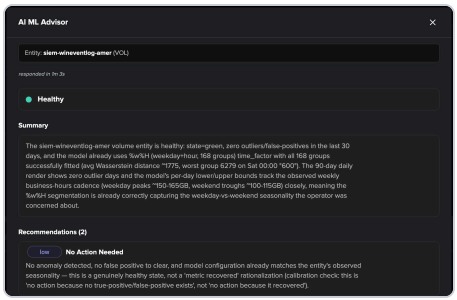

The ML Advisor reviewing an outlier model: verdict, recommendations, reasoning trace.

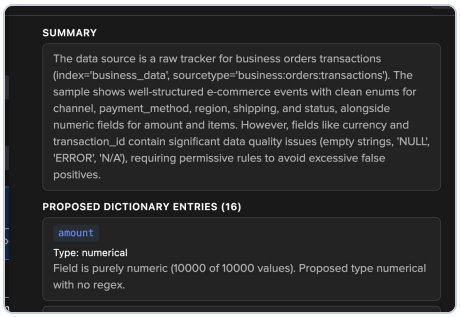

A field-quality dictionary proposed by the FQM Advisor for a custom sourcetype.

Opt-in only
Completely inert until an administrator configures a provider. No data leaves, no LLM call is made, until you opt in — from air-gapped Ollama to Splunk-hosted models. Optional anonymisation of entity and index names.

Trust
Agents never make up entity state: every claim traces to a tool call, outputs are typed, read-only inspect is the default and every write requires consent — in the chat, in the advisor, in the routine.

Auditable
Runtime guardrails and hard caps bound the blast radius. Every tool call and every change is indexed and prefixed [AI Agent] in the audit trail; dedicated dashboards review interactive and scheduled activity.

Providers: OpenAI, Azure OpenAI, Anthropic, Google Gemini, Mistral, Ollama (local), Splunk-hosted and OpenAI-compatible endpoints. The AI Assistant and Concierge run on any supported Splunk version; Advisors, Agents automation and Routines are built on the Splunk Agent SDK and require Splunk 10.2 or later (Python 3.13).

OPERATE AT SCALE

Built for federation, governance and scale

A single flat monitoring space does not survive contact with a real Splunk estate. TrackMe carves it into domains that match how the organisation works, keeps the operational model defensible, and watches its own configuration.

Virtual Tenants & Virtual Groups

Isolated monitoring domains — by company, team, technology or Splunk cluster — each with its own entities, trackers, thresholds, roles and, optionally, its own remote deployment. Virtual Groups aggregate entities across tenants into read-only cross-tenant cards for a NOC view.

Remote deployments

One TrackMe monitors many Splunk deployments — Enterprise, Cloud, search head clusters, utility nodes — over authenticated, token-rotated connections with automatic failover. The experience is identical whether the data is local or remote.

RBAC & audit

Three roles — user, power, admin — assigned per tenant and enforced at the REST boundary, so the UI, the API and the AI Assistant all go through the same gate. Every change is recorded in an audit index, with per-entity audit tabs.

Configuration Guardian

Detects misconfiguration before it hurts: missing capabilities, absent indexes, expiring tokens, unreachable providers, stale backups, degraded tenants, stale topology views — with severity, remediation steps, snooze and email notifications.

Backup & Restore

Scheduled, independent archives per tenant plus a global archive; single-archive, whole-run or dry-run restore; optional export to S3-compatible storage; safe across search head clusters.

REST API & automation

Everything the UI does is a documented REST endpoint, browsable in the product, with a search-command wrapper for SPL-driven automation, CI/CD friendly configuration and a full alert-action toolkit.

Part of the Splunk ecosystem you already have

- **Splunk Enterprise Security** — detection integrity end to end: data and field gaps are flagged before correlation searches fail silently.
- **Splunk ITSI** — telemetry completeness as a first-class signal next to service KPIs.
- **Splunk SOAR** — native integration: TrackMe events feed playbooks; automation brokers and failover are handled safely.
- **Cribl Stream & Cribl Edge** — pipeline-side visibility from edge to indexer, tied back to Splunk outcomes.

Compatibility

Splunk Enterprise 9.2 and later; Splunk Cloud Victoria and Classic. Delivered through Splunkbase via the standard Splunk Cloud app-vetting process. Search head clusters, indexer clusters and hybrid estates are supported.

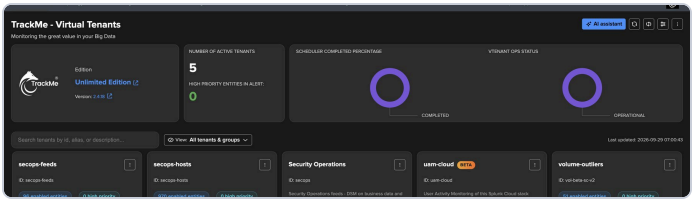

Virtual Tenants — one card per monitoring domain, each with its own entities, trackers, roles and live counters.

Trusted across industries and continents

Used by organisations in financial services, government, defence, healthcare, energy, telecommunications and technology — from Fortune 500 enterprises to government agencies — across EMEA, the Americas and Asia-Pacific.

EDITIONS

Pick the tier that fits — start small, grow without re-platforming

Annual subscription. Every edition includes Topology Studio, stateful alerting, policies, the AI Assistant and Configuration Guardian. Multi-year options and volume discounts from two licences.

	Foundation	Enterprise	Unlimited	Unlimited Deployments
Price	**$6,000** per year	**Contact us**	**Contact us**	**Contact us**
Scope	One Splunk deployment	One Splunk deployment	One Splunk deployment	Organisation-wide — every Splunk deployment, one licence
Components	Feeds tracking (DSM, DHM, MHM) and Volume Outliers	Adds Flex Objects, Field Quality and Workload; User Activity Monitoring (Beta)	All components	All components
Virtual Tenants	3	Up to 6	Unlimited	Unlimited
Remote deployments	1	Up to 8	Unlimited	Unlimited
Hybrid trackers	Up to 8	Up to 16	Unlimited	Unlimited
Elastic trackers	Unlimited	Unlimited	Unlimited	Unlimited
FQM / FLX trackers	—	32 each	Unlimited	Unlimited
ML outliers	Feeds and volume	All components	All components	All components
Support	Standard	Premium	Premium	Premium

Topology Studio, stateful alerting, policies, Configuration Guardian, Backup & Restore and the AI Assistant are included in every edition. AI Advisors, Agents automation and Routines require Splunk 10.2 or later. Prices in USD, excluding taxes; contact us for regional and partner pricing.

Get started

Every new TrackMe installation starts a 90-day trial automatically — install from Splunkbase, create a first tenant, and the trackers discover your feeds within minutes. Request a free 90-day trial with all features enabled and no restrictions; extended trials are available for large evaluations.

Where to look next

- **docs.trackme-solutions.com** — the Product Guide, white papers and step-by-step tutorials.
- **trackme-solutions.com/discover** — feature tour and demonstrations.
- **Hands-on enablement decks** — from a first tenant to Topology Studio, Field Quality, Volume Outliers and User Activity Monitoring.

TALK TO US

Get a demo, an extended trial, or a partner conversation.

contact@trackme-solutions.com

trackme-solutions.com · docs.trackme-solutions.com · trackme-solutions.com/discover

TrackMe Limited · London, United Kingdom