



PRODUCT BROCHURE

TrackMe for Splunk

Visibility & operational excellence

TrackMe discovers, maintains and monitors the availability and quality of your data, and of the platform around it: feeds, hosts, licensed volume, fields, scheduled workload, custom objects, the people using Splunk, the HEC endpoints receiving your data and the Cribl pipelines in front of it. One Splunk deployment or many, at any scale, entirely inside your Splunk environment.



TWO WAYS OF WATCHING, ONE VOCABULARY

ENTITY MONITORING seven components, one entity model: a state, an impact score, stateful alerts

DSM Data Source Monitoring Is the data arriving?	DHM Data Host Monitoring Who is sending, and what?	MHM Metric Host Monitoring Are my metrics arriving?	VOL Volume Outliers Is my licensed volume normal?
FQM Field Quality Monitoring Is my data well-formed?	WLK Workload Knowledge Are my searches healthy?	FLX Flex Objects Anything else, from your SPL	

FINDINGS MONITORING four dedicated components: policies raise findings you triage

UAM User Activity Monitoring Who is doing what?	HEC Splunk HEC Are my endpoints receiving?
CS Cribl Stream Is the pipeline in front healthy?	CE Cribl Edge Is every Edge Node where it should be?



Topology Studio: a public web estate built from TrackMe entities, and the flow topology of a Cribl Stream deployment, Routes to Destinations, coloured by their findings.

TrackMe Limited

London, United Kingdom

trackme-solutions.com · docs.trackme-solutions.com

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

In a large Splunk environment, things break silently. 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. The platform around the data fails just as quietly: an index doubles its licensed volume, an account runs all-time searches on a schedule nobody owns, a HEC token goes silent, a Cribl Destination backs up in front of Splunk. Splunk itself will not tell you.

TrackMe is a control plane on top of Splunk. It continuously reads what Splunk already holds and turns it into something you can operate on: health states and impact scores for your data, findings for the platform around it, policies, stateful alerts, live topology maps and, optionally, AI-assisted investigation and triage. 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 and metric feed is tracked for delay, latency and volume, with thresholds that follow the day and the hour.

Is my data still good?

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

Is my licensed volume normal?

Every index as an entity: drops, spikes, trend, month projection, cost and chargeback.

Are my searches healthy?

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


Who is doing what on Splunk?

Accounts, their searches, their cost, their privileges and the knowledge objects they own.

Is the pipeline in front healthy?

HEC endpoints, pools and tokens; Cribl Stream and Cribl Edge, from their own metrics and Leader.

And the one behind all of them: who do I tell, and how urgently? Priority, SLA and impact scoring decide what is worth an alert and what is noise. Stateful alerting and finding notifications reach the right people once, not every five minutes.



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, plain-English routines and an AI Agent that triages findings. Opt-in, inert until you configure a provider, audited end to end.

Seven components, one entity model

Entity monitoring is the historical core of TrackMe. Seven components turn data feeds, hosts, metrics, the licensed volume of each index, field quality, scheduled searches and any SPL into 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. Classification, protection, stateful alerting, scoring and ML outliers are learned once and apply to all seven.

DSM Data Source Monitoring

Index and 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, read incrementally from the license usage log: 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. Global license usage answers who consumes, the trend, days over quota, costs and chargeback.

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, host OS, HTTP heartbeats and more.



VOL: global license usage of a tenant, trend, month projection, quota, cost and the insights of the period.



FQM: one field as an entity, its coverage and success rate, the sampled values and the processing status.



WLK: a scheduled search as an entity, its SLA and its execution metrics over 24 hours.

The platform around your data, monitored out of the box

Some things are not entities with a health score: an account running all-time searches, a HEC token gone silent, a Cribl Destination in backpressure, an Edge Fleet left on an old version. Four dedicated components share one findings model. Built-in policies evaluate a condition and raise a finding about one object, with its evidence and a severity from Informational to Critical. All findings about one object form a Finding group: one status, one assignee, one severity, one notification thread. Findings recover by themselves when the condition clears; a person decides the rest. Nothing is installed on the platforms being watched, and every component ships its own AI Agent.

UAM User Activity Monitoring

Who is doing what on your Splunk. People and service accounts, through the audit trail and the introspection of up to ten search head tiers: who owns which schedule, who ran all-time searches, who holds high privileges and why, who eats the search CPU. Sixteen built-in policies in four families, custom policies in SPL, catalogs of every knowledge object, and reviewed emails to the account concerned. It never changes anything on the platform.

HEC Splunk HEC

Every HTTP Event Collector endpoint, pool and token of a Splunk deployment, discovered from the telemetry Splunk already writes: no agent, no network path, no credential. Eighteen built-in policies cover availability, refused senders, back pressure (naming the stage that cannot keep up), load-balancer imbalance, token silence and ML traffic outliers. Tokens are known by name only, never by value. Splunk Cloud included.

CS Cribl Stream

One Cribl Stream deployment end to end: the Leader, Worker Groups, Workers and their machines, and every Source, Destination, Pipeline and Route, from the metrics Cribl already sends to Splunk and the Leader API, read-only. A traffic ML model on every data object, sixteen built-in policies, and flow topology drawn with zero configuration. Nothing installed on Cribl.

CE Cribl Edge

The Fleets of a Cribl Edge Leader and every Edge Node in them, read in one bulk call so the cost does not grow with the fleet. Fleets are judged as a share of their nodes; version drift after a rollout, configuration drift after a deploy, failing Sources and Destinations and full disks become findings. Eleven policies, explicit and explainable rules, no machine learning by design.

Policy

a built-in or custom rule evaluates one condition

Finding

about one object, with evidence and a severity

Finding group

every finding of an object, one decision

Triage


status, assignee, snooze; routed by assignment rules

Notify & review


email, webhooks, indexed events, AI Agent




UAM: the findings backlog by status, severity and policy, ready for the AI Agent review.



Splunk HEC: the live topology of pools and endpoints, availability, events and the member that stopped.



Cribl Stream: the tenant workspace, its connection, Leader and models, and the open findings.



Cribl Edge: the Fleet at a glance, 17 of 17 nodes connected, versions and the 24-hour trend.

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. Findings tenants run the same idea with their own jobs: collect, evaluate the policies, raise or recover findings, notify.

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, then green, orange, red	policies assign priority, tags, labels, SLA	groups, grace, maintenance, acknowledgment: 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 is green • below 100 orange • 100 and above 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.

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.

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**: 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: the AI status report, then the incident, flipping and performance charts of the last 24 hours.

For findings, the same discipline

A severity from Informational to Critical, one notification thread per Finding group, assignment rules that route findings to teams, webhooks and findings actions that push them anywhere, and every transition indexed for Enterprise Security or SOAR.

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 and findings objects, so it cannot drift from what is monitored, and coloured with live health.



Generated in one click from the Splunk HEC tenant: pools, endpoints and tokens with live status, events, CPU and memory as sparkline KPIs, refreshed every minute.

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. Findings tenants draw their pools, Workers, Routes and Fleets the same way.

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.

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 and one agent for findings, 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, 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.


FOR FINDINGS

The AI Agent


User Activity Monitoring, Splunk HEC, Cribl Stream and Cribl Edge share one agentic operator. Asked about a finding, or during its hourly scheduled triage, it reads the evidence, the history and the facts TrackMe holds about the object, assesses the finding and recommends one decision with its confidence. In act mode it applies only what the tenant allows, resolves only with proof, never overrides a person, and logs every decision. Administrators give it custom instructions in plain English, what to treat as expected, what to look at first, and evaluate them side by side with the default behaviour before enabling.



Custom instructions, evaluated: the same finding reviewed with and without them, writing nothing.



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



The ML Advisor documenting an outlier investigation and the remediation it applied, as a note on the entity.

Opt-in only

Completely inert until an administrator configures a provider. No data leaves and 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, in the agent.

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, Routines and the AI Agent for findings are built on the Splunk Agent SDK and require Splunk 10.2 or later.

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 or findings, 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 and 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, stalled collections, 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:** dedicated tenants for the pipeline in front of Splunk, from Cribl's own metrics and Leader API, 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.

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
Entity components	Feeds tracking (DSM, DHM, MHM) and Volume Outliers	Adds Flex Objects, Field Quality and Workload	All components	All components
Dedicated components User Activity, Splunk HEC, Cribl Stream, Cribl Edge	Not included	Included	Included	Included
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	Not included	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, Routines and the AI Agent for findings 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 walkthroughs: from a first tenant to Topology Studio, Field Quality, Volume Outliers, User Activity Monitoring, Splunk HEC, Cribl Stream and Cribl Edge.

TALK TO US

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

contact@trackme-solutions.com

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

TrackMe Limited · London, United Kingdom