

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

1. Overview

AI SIEM provides a unified security event management interface for AWS environments. It correlates events from CloudTrail, VPC Flow Logs, GuardDuty, Security Hub, and AWS Config into a single threat timeline with ML-driven behavioral analysis.

The Editor UI is a web application served via ALB with Cognito authentication. It provides:

- Threat Dashboard — real-time severity breakdown, source health, trend graphs
- Threat List — active and resolved threats with severity scores
- Timeline Viewer — chronological event narrative per threat
- Source Management — configure and monitor data source health
- Detection Rules — custom static rules and ML tuning
- Playbooks — AI-generated incident response guidance
- Compliance — posture scoring against SOC2, PCI, HIPAA

2. Threat Dashboard

The dashboard is the default view when you open the editor. It shows:

- Active threats by severity (critical, high, medium, low, info)
- Threat trend graph (7-day rolling)
- Data source health (last event received timestamp per source)
- Top affected principals (IAM users/roles with most threat activity)
- MITRE ATT&CK coverage map (advanced+ tiers)

3. Threat List

3.1 Viewing Threats

The threat list shows all detected threats sorted by severity and recency. Each row shows:

- Severity score (0–100)
- Threat type (e.g., Impossible Travel, Credential Compromise, Port Scan)
- Affected principal and resources
- First and last seen timestamps
- Status: active, acknowledged, investigating, resolved, false_positive

3.2 Managing Threats

Click a threat to see its full timeline. From the detail view:

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

- Acknowledge — marks the threat as seen (stops repeat notifications)
- Investigate — marks as under active investigation
- Resolve — closes the threat (with notes)
- False Positive — flags a false alarm and removes it from the active list (recorded for audit). It does not auto-tune detection; to stop the pattern recurring, edit the rule or use Analyze on the group
- Generate Playbook — AI creates step-by-step remediation guidance
- Remediate — execute auto-remediation actions (enterprise only)
- Add a note — record a free-text note as you work the threat; notes are append-only and appear in Audit History and the audit trail
- Download — export a self-contained HTML incident report (facts, evidence, timeline, correlated identities, audit trail, playbook) with live links back to the threat and to any custom rule that fired

The detail view has two tabs: Summary (key facts, description, assignment, and a quick overview for fast triage) and Details (the full evidence, timeline, correlated events, audit history, and playbook). In the evidence, a custom rule that fired is a link that opens it under Detection Rules for tuning, with a "Back to threat" link to return.

3.3 Filtering and Assigning Threats

The Threats toolbar filters the list by status, severity, type, and a time window (last 24 hours, 48 hours, week, month, 2 months, or older), plus free-text search. Filters combine, and a bulk selection applies an action (acknowledge, investigate, resolve, false positive, dismiss, or assign/unassign) to every matching threat at once.

Assign a threat from its detail view: an admin can assign it to any admin or analyst, or claim it; an analyst can claim it for themselves. The detail-view assignee box also accepts any user's full email address (useful for someone who has not yet signed in). To assign many at once, select threats and use the toolbar's Assign controls: admins pick a user from the dropdown, analysts claim the selection for themselves, and either can Unassign. Before threats that already have an owner are reassigned, you are told how many will change hands. Marking a single unassigned threat as Investigating also prompts you to claim it. Tick "Assigned to me" to see only your queue, and (admins) use the "Assigned to" filter to view a specific owner's queue or find unassigned threats. Every assignment is recorded individually in the audit trail.

4. Timeline Viewer

The unified timeline shows all correlated events for a threat in chronological order:

- CloudTrail API calls (who did what, from where)
- VPC Flow Log connections (network-level activity)
- GuardDuty findings (AWS-detected threats)
- Config changes (infrastructure modifications)
- Security Hub findings (aggregated compliance issues)

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

Each event is tagged with its source, MITRE ATT&CK tactic/technique (if mapped), and threat relevance score.

5. Data Sources

Navigate to Sources in the sidebar to view and manage data source configuration.

5.1 CloudTrail

Ingests from the S3 bucket specified at deploy time. Processes management events and optionally data events. Source health shows last file processed timestamp.

5.2 VPC Flow Logs

Ingests from S3 if configured at deploy time. Parses accepted/rejected flows for port scanning, lateral movement, and data exfiltration detection.

5.2.1 VPC Flow Panel (Essential+)

Flow logs are the highest-volume source AI SIEM handles, so the Configure – VPC Flow panel puts their cost and behavior on one screen. The panel opens with a live snapshot: current mode, estimated scan volume and spend for the trailing window, the top talking sources, and recent sweep activity.

Mode is the main lever. Athena (the default) leaves flow logs in the datalake and runs scheduled sweeps against them – low, predictable cost that scales with query cadence rather than traffic. Off stops flow detection entirely without touching your other sources. Either way flow logs are always retained in S3 and remain fully queryable via Investigations.

The spend shown is an estimate derived from the bytes scanned by the sweeps and per-gigabyte rates, not a billed figure from AWS – treat it as a directional guide, not an invoice. Detection from flow logs (port scanning and data exfiltration sweeps) runs on a schedule: roughly every 15 minutes for port-scan and hourly for exfiltration. If flow-log delivery is not yet configured, the panel and its dashboard cards will read zero until logs begin arriving.

5.2.2 Per-VPC Detection Thresholds (Essential+)

The default port-scan and large-egress thresholds under Detection tunables apply to every VPC. When one setting does not fit a mixed fleet – a busy production VPC that trips a global threshold with normal traffic, or a quiet management VPC where real activity would hide under it – use Per-VPC thresholds to override the defaults for specific VPCs. These overrides apply to the Athena sweeps.

Click Add VPC, give the profile a name, and pick the VPC. The VPC list is populated from the VPCs that have actually sent flow logs recently, so you can select rather than type an ID (you can still type one that has not reported yet). Set a port-scan threshold, an exfil-bytes threshold, or both; leave either blank to

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

inherit the global default for that VPC. Each saved profile shows as a row with its effective thresholds and Edit and Delete controls; the list scrolls when it grows. As with all configuration, changes take effect after you click Save.

The Scan unpinned VPCs checkbox controls VPCs that have no profile. Checked (the default), they are scanned with the global thresholds. Unchecked, only the VPCs you have listed are scanned and all others are ignored – turning the profile list into an allowlist. Flow-based threats record the VPC they came from, shown in the threat detail, so you can tell which network a detection belongs to.

5.3 GuardDuty

Receives findings automatically via EventBridge. No configuration needed. Findings are enriched with behavioral context from other sources.

5.4 Security Hub

Receives findings via EventBridge. Deduplicated against direct GuardDuty findings. Compliance findings contribute to posture scoring.

5.5 AWS Config

Receives configuration changes via EventBridge. Tracks infrastructure modifications during threat windows (e.g., security group changes, IAM policy changes).

AWS SIEM reads Config “compliance-change” events, which are only emitted when a Config RULE evaluates a resource – so enabling the Config recorder alone is not enough; you also need a delivery channel and at least one managed rule. If the source shows “AWS service not enabled,” Config is simply off in your account, and that status is harmless if you do not need it. To turn it on, the simplest path is the AWS Console: open AWS Config, choose Get started (this creates the recorder, the delivery bucket, and its permissions for you), then add at least one managed rule – for example s3-bucket-public-read-prohibited – on the rules screen. Once a rule evaluates a resource, the “AWS Config Changes” source turns healthy on its own; no setup is needed on the AI SIEM side. Candidly, this is the most setup-heavy source for the least incremental detection value in a typical deployment – CloudTrail and GuardDuty already provide strong coverage – so enable Config specifically when you want compliance-change detection, not by default.

6. Detection Rules

6.1 ML Behavioral Baselines

The system automatically learns normal behavior per IAM principal:

- Which APIs and API families are typically called

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

- What time of day activity occurs
- Which source IPs are used
- Which regions the principal operates in
- Which resources (by service) are accessed
- The identity type of the principal (service, role, user, root)

Deviations across these nine behavioral dimensions generate anomaly scores as a distance from the learned per-principal norm. Baselines are rebuilt on a configurable schedule (default: daily), and learn continuously from new activity.

6.2 Static Rules

Built-in rules fire immediately without baseline learning:

- Impossible travel (same credential from geographically distant IPs)
- Credential stuffing (high rate of failed auth attempts)
- Root account usage
- CloudTrail disabled or modified
- Security group opened to 0.0.0.0/0
- IAM policy with * permissions attached
- S3 bucket made public

6.3 Context-Aware Severity (Advanced+)

On advanced tiers and above, a matched rule can score higher in riskier context. Two fixed modifiers apply on top of a rule's severity, and both only ever raise a score – never below the rule's stated severity, and never over a higher ML score:

- Root principal – a rule matched by the account root identity is scaled up. No configuration required.
- Untrusted source IP – a match from an address outside your trusted networks is scaled up. This modifier is inert until you list your trusted IPs (`severityModifiers.trustedIps` under Detection Rules); with no list configured it does nothing, so enabling the tier never silently inflates scores.

The modifiers key only on identity and network facts, not on whether the activity is anomalous (the ML baseline already measures that), so a finding is never double-counted. A root action from an unknown IP can reach maximum severity; the same action from a trusted network scores at the rule's base.

6.4 Custom Detection Rules (Advanced+)

Beyond the built-in rules, you can author your own under Configure → Detection Rules → Custom rules. A custom rule fires when all its conditions match and raises the threat score to at least its severity (it never lowers a higher score), then rides the same timeline, alerts, reports, and audit trail as any built-in detection. There are three ways to create one:

- Start from a template – click a ready-made detection (secret read, access-key creation, MFA

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

disabled, snapshot shared, and more) to fill the builder, then tweak. Works on any tier that includes custom rules, no AI required.

- Describe it in plain English – type what you want to detect (for example, “someone reads a secret from a new IP”) and AI drafts the rule for review. Requires the AI (Bedrock) feature. You can also Refine an existing rule the same way, or use More ideas for AI-suggested detections you do not already have.
- Import a Sigma rule – paste a Sigma rule (YAML) and it is translated into a custom rule. Import is a fast start for a rule you already have – a bounded translator, not a full Sigma engine, and it never silently mistranslates: anything it cannot map is named so you can add it by hand. It handles a selection of AND-ed fields with the modifiers equals / in-list / contains / starts-with / ends-with / greater-than / less-than and their or-equal forms / cidr / exists, plus level, log source, and attack.* tags. An OR condition (selA or selB, or “1 of ...”) is imported as several rules – one per branch – since AI SIEM rules already fire independently; review each before saving. Not translated (listed for you to add by hand): regular expressions (excluded by design – use contains / starts-with / ends-with or an in-list instead), the |all and base64 modifiers, and nested/AND-across-selection or negated conditions. For anything the importer cannot take, describe the detection in plain English and let AI draft it.

Conditions use a fixed operator set (equals / not-equals, in-list / not-in-list, contains / does-not-contain, starts-with / ends-with, greater-than / less-than and their or-equal forms, in-CIDR / not-in-CIDR for IP fields, and exists / does-not-exist) on any normalized event field – there is no free-form expression language or regex, so a rule cannot slow or break detection. The negated operators make exclusions easy – for example, exclude AWS service-to-service calls with principal-type not-equals AWSService and principal does-not-contain :role/aws-service-role/. A rule can also require a count within a time window (for example, N matches in 5 minutes), optionally counting distinct values of a field.

Whatever the source, the rule opens in the builder for review – nothing is applied until you click Apply and then Save in the header. Configuration edits accumulate into one working copy across every panel and are committed together by that single Save; the header also shows an unsaved-changes indicator and an Undo that discards pending edits, and Settings → Configuration History lets you review the exact field-level changes before saving. Custom rules can carry an ATT&CK tactic and technique, so they extend your MITRE coverage, and (advanced+) they receive the same context-aware severity modifiers as built-in rules. Every save is recorded in the audit trail with a per-rule diff of what was added, changed, or removed. Each detection – built-in and custom – also has an Email alerts toggle: turn it off to keep the detection scoring threats, building timelines, and appearing on the dashboard while silencing its email. When several rules raise the same threat the email is muted only if all of them are muted, and a threat raised purely by the ML baseline still alerts on its score – so muting a rule narrows notifications without ever creating a blind spot.

7. Incident Response

7.1 AI Playbooks (Advanced+)

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

Click Generate Playbook on any threat to receive Bedrock-generated step-by-step guidance tailored to:

- The specific threat type
- Affected resources in your account
- Timeline context (what happened before/after)
- Best practices for containment and remediation

Identifier privacy is always on and cannot be turned off, on every tier. Before any prompt is sent to Bedrock, real identifiers (IAM principals and ARNs, IP addresses, account and resource IDs) are replaced with type-preserving placeholders; the model never sees a literal value. The playbook is cached with its identifiers redacted and the lookup map encrypted with a dedicated KMS key in your account, and real values are restored only when the result is shown to you. If a real identifier ever escapes redaction into the model output, that playbook is discarded and a safe deterministic template is used instead. The badge on each playbook shows which protection applied.

7.2 Auto-Remediation (Enterprise)

Available actions (opt-in per type):

- Disable IAM access key
- Revoke active IAM sessions
- Attach deny-all inline policy
- Modify security group to isolate instance

Every action is recorded in DynamoDB with full audit trail (who approved, when executed, result).

7.3 Investigations (Advanced+)

The Investigations view answers "what was detected between these two dates?" for a post-mortem, without writing any query language. Pick a From and To date and AI SIEM searches the long-retention forensic archive – which reaches further back than the live Threats queue – and lists every threat detected in that window, newest first. Click any result to open the full threat (current status, timeline, and playbook); if a threat has aged out of the live store, its detection-time record from the archive is shown instead. Queries run in the background: a fast one returns in seconds, while a wide one keeps working and appears under Previous investigations to reopen later, so the screen never hangs. Saved investigations are retrievable and expire on the same retention window as reports. Available to admins and analysts.

8. Compliance Posture (Advanced+)

The Compliance view scores your environment against security frameworks based on events observed:

- SOC 2 — Trust Service Criteria mapping
- PCI DSS — Cardholder data environment controls
- HIPAA — Technical safeguards for ePHI

AI SIEM — User Guide

Version: 26h4 | **Product:** AI SIEM | Perfware.cloud

Daily use of the AI SIEM Editor UI.

Scores update continuously as events flow. Weekly posture reports can be emailed via SNS.

9. Monitoring & Alarms

The stack publishes a CloudWatch dashboard (named <stack>-monitoring) and a set of CloudWatch alarms in your own account, so operators can watch the pipeline's health without opening the editor. The dashboard leads with an alarm-status panel summarizing every AI SIEM alarm at a glance, followed by graphs for events ingested vs. promoted, threats by severity, Lambda invocations/duration/errors, SQS queue and dead-letter depth, the flow-log Athena sweeps, and the daily retention/archive pass.

Alarms notify the configured SNS topic and cover: ingestion and detector Lambda errors, ingestion and threat dead-letter queues (which should stay at zero), threat-queue backlog, high-severity threat spikes, WORM audit-write failures, the retention/archive job (volume, active evictions, and errors), and sustained OpenSearch write failures. The last is worth understanding: the OpenSearch dashboards are a best-effort projection of the authoritative data the editor reads, so a brief OpenSearch outage only makes those views momentarily stale (the daily reconcile then self-heals them) – detection, alerting, and the audit trail are unaffected. Point the SNS topic at your email or ticketing system to receive these.

10. Support

- Documentation: <https://perfware.cloud/aisiem>
- Support: support@perfware.cloud
- Security: security@perfware.cloud