THE CYBER TOOLCHAIN

NO. 022

Head

THE WEEK'S TOP PICKS

PUBLISHED NOVEMBER 27, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

The week in numbers · Nov 20–26

647
changes we read →

228
tools shipped →

2
product launches →

250
breaking changes →

Changes per week, last 12 weeks · Nov 20–26

Changes we read in that week's dailies, including docs and API changes. A gap means no daily was published.

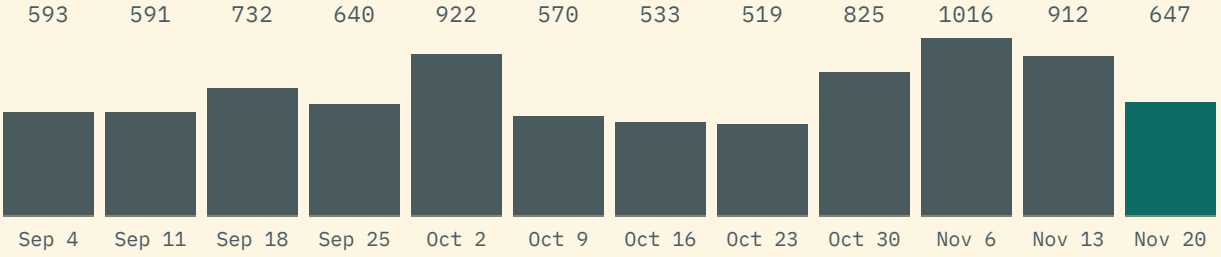

▸ // HOW THESE PICKS ARE MADE SHOW

VIEW

vim ? help

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every change in this week's picks, in full.

all 41 changes

Do you prefer **scroll**? 👍 👎

This week's 5

01 **Elastic Security Labs** — Feature · SIEM

02 **Qualys TotalCloud** — Feature · CNAPP

03 **Bedrock Security** — Feature · Data Security & DLP

04 **ifixai** — Feature · AI Red Teaming & Testing

05 **rsigma** — Feature · Detection Engineering

The picks in full

01 Elastic Security Labs

Press play on Elastic Security 9.6, banger after banger (2026-11-23) SIEM

Elastic Security Labs is a research division that publishes threat intelligence, malware analysis, and security research to help defenders detect and respond to cyber threats.

// WHY IT MADE THE LIST DEPTH 5/5 IMPACT 5/5 NOVELTY 5/5

Elastic Security 9.6 ships AlertZero, a full agentic SOC layer with four Watch agents (Triage, Hunt, Detection, Forensics) running six AI Workers at configurable autonomy levels — including auto-closing false positives above an 85% confidence threshold — alongside 143 new AI/agent detection rules covering coding agents, MCP servers, and cloud control planes, four lightweight NLP classifiers that score every GenAI document at ingest for prompt injection and encoded payloads, custom YARA memory scanning with policy-level deployment, and Linux ransomware protection reaching parity with Windows and macOS. No other SIEM shipped anything close to this breadth of AI-era detection capability in a single release.

> Elastic Security shipped AlertZero, a technical-preview AI SOC agent system in Elastic Security 9.6 with six autonomous Workers for triage, hunting and detection engineering, and a new OpenShift compliance integration that centralises multi-cluster scan data in Elasticsearch behind an Elastic Agent Builder agent.

└──▷ WHAT SHIPPED · 10 FEATURES most completely described first ? what's the number?

01 OpenShift compliance data ingestion and transforms

Not checked yet NEW 100

An HTTP JSON input queries OpenShift Compliance Operator APIs every 60 seconds, shipping `compliancescans.compliance.openshift.io`, `compliancecheckresults.compliance.openshift.io`, and `complianceremediations.compliance.openshift.io` resources into Elasticsearch data streams; three Elasticsearch transforms — `scans`, `checks`, and `remedy` — continuously collapse this time-series data into latest-state indices named `ospm-compliance_-latest`, deduplicating check events by cluster and `metadata.name`. Every ingested event is tagged with `orchestrator.cluster.name` using the `CLUSTER_NAME` environment variable set on the agent, keeping multi-cluster scan results separated.

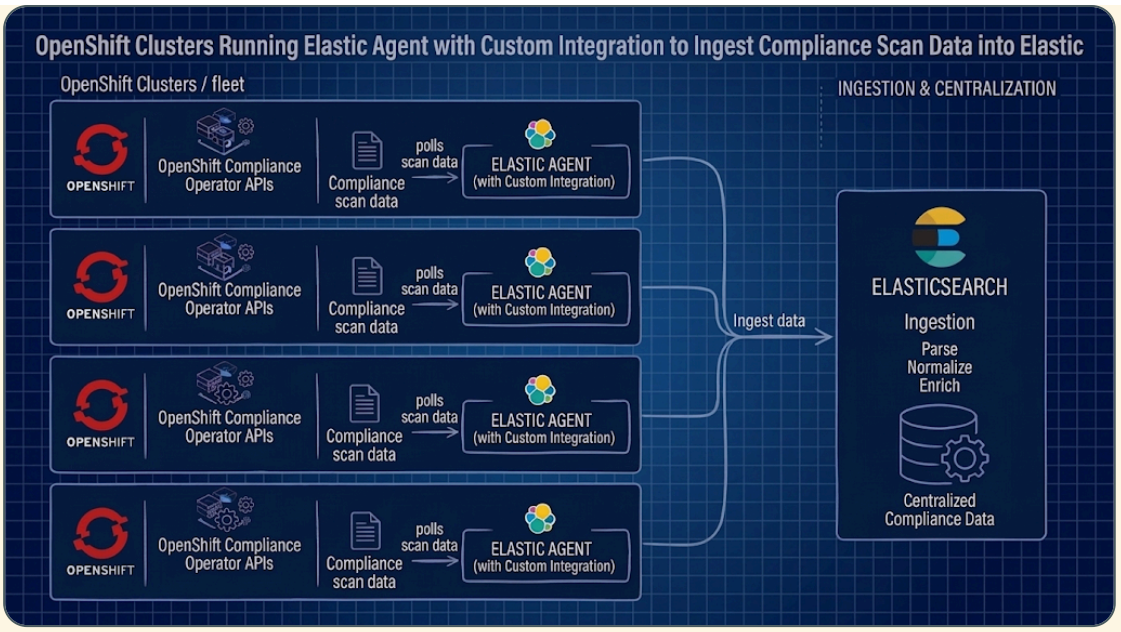

Deploy Elastic Agent on an OpenShift cluster to start collecting compliance scan data into Elasticsearch, tagging every event with the cluster name so multi-cluster results stay separated.

```yaml
apiVersion: apps/v1
kind: Deployment
metadata:
  name: elastic-agent
  namespace: elastic-agent-ocp
spec:
  replicas: 1
  template:
    spec:
      containers:
      - name: elastic-agent
        image: docker.elastic.co/elastic-agent/elastic-
agent:9.5.5
        env:
        - name: CLUSTER_NAME
          value: "prod-us-east-1"
        - name: FLEET_ENROLL
          value: "1"
        - name: FLEET_INSECURE
          value: "false"
```

Upload the custom OpenShift Security Posture Management integration package through the Kibana Fleet UI to register all data streams, ingest pipelines, transforms, and dashboards in one step.

📍 In Kibana, go to Fleet › Integrations › Upload Integration, select the OSPM zip package from the GitHub repository, then open Fleet › Agents, create a new Agent Policy, add the OSPM integration, set the three compliance endpoints and 60-second polling interval, generate an enrollment token, and deploy the manifest to your cluster.

— Named endpoints, transforms and env var with a runnable deploy example

From posture to root cause: Investigate…

02 **Proposed actions review and case escalation** Not checked yet

NEW SUPPORTED PLATFORM NEW 80

A dedicated proposed actions page, grouped under Respond, Investigate, and Configure, lets analysts review and approve or decline each Worker's proposed action backed by its gathered evidence and reasoning inside an Elastic Agent Builder investigation conversation; related investigations can then be escalated into a shared case conversation, assigned to team members, and marked private or public.

Walk through approving or declining an AlertZero proposed action — for example, a host isolation recommended by the Endpoint Analysis Worker — before it runs against your environment.

📍 In Elastic Security, go to AlertZero › Proposed Actions, select the action under Respond, review the evidence and reasoning in the attached investigation, then click Approve (runs as you, with your privileges) or Decline.

— Exact navigation and approve/decline steps given

AlertZero in Elastic Security: Every al…

03 **Per-Worker autonomy levels** Not checked yet NEW 75

Each AlertZero Worker can be set to Manual, Assisted, or Supervised autonomy, controlling whether proposed actions — such as closing false positives, isolating hosts, or editing detection rules — wait for analyst approval or run automatically.

— Three named levels and example actions, but no config path given

AlertZero in Elastic Security: Every al…

04 **AlertZero AI Workers across four Watches** Not checked yet

NEW DETECTION NEW SUPPORTED PLATFORM NEW INTEGRATION NEW 70

AlertZero, in technical preview on Elastic Cloud Serverless with Elastic Security 9.6 (ECH and Self-Managed support expected early December), introduces six AI Workers organized under four Watches that automate alert triage, attack discovery, continuous threat hunting, rule tuning, rule coverage gap review, and endpoint forensic analysis.

— Names platform/version scope but not exact enablement path

AlertZero in Elastic Security: Every al…

05 **Alert Triage Worker confidence scoring** Not checked yet NEW

70

The Alert Triage Worker classifies new alerts as true or false positives with a confidence score, and at Supervised autonomy auto-closes false positives at or above an 85% default confidence threshold.

— Exact threshold named but no way to change it described

AlertZero in Elastic Security: Every al…

06 **Prebuilt OpenShift compliance dashboards** Not checked yet NEW

60

The integration adds three prebuilt Kibana dashboards — `[OpenShift Compliance] Posture Overview`, `[OpenShift Compliance] Overall Results`, and `[OpenShift Compliance] Auto Remediations` — installed automatically when the integration package is uploaded.

— Three named dashboards but no navigation beyond install

From posture to root cause: Investigate…

THINNER COVERAGE BELOW

07 **Per-step model tier assignment** Not checked yet NEW 55

Each AI step inside a Worker's workflow can be assigned to one of three model tiers — Efficient, Balanced, or Frontier — so the right model is used per step and can be swapped without changing the Worker's underlying logic.

— Names the tiers but not how assignment is configured

AlertZero in Elastic Security: Every al…

08 **Worker-to-Worker handoffs** Not checked yet NEW 55

AlertZero Workers hand off work to each other: the Alert Triage Worker feeds false-positive counts to the Rule Tuning Worker, the Attack Discovery Worker escalates to the Endpoint Analysis Worker, and the Continuous Threat Hunt Worker records detection gaps for the Rule Coverage Worker.

— Describes the workflow but offers no user action

AlertZero in Elastic Security: Every al…

09 **CEL custom rule evaluation** Not checked yet NEW DETECTION NEW

55

Manual compliance rules can be evaluated with Common Expression Language (CEL) custom rules, with a scheduled Elastic Workflow annotating results back to the manual rule in the cluster.

.

— Names CEL and Workflow mechanism but no command shown

From posture to root cause: Investigate…

---

10 **Agent Builder agent for compliance queries** Not checked yet NEW

45

An Elastic Agent Builder agent sits over the centralised OpenShift compliance data, letting teams query fleet-wide posture and investigate control failures in plain English.

.

.

— Capability described but no example query or path shown

From posture to root cause: Investigate…

NEW DETECTION MITRE ATLAS mappings for AI-platform prebuilt detection rules: AI platforms, generative AI (GenAI) tooling

NEW SUPPORTED PLATFORM Ransomware protection extended to Linux

NEW INTEGRATION Automatic Migration support covers Splunk, QRadar, and Microsoft Sentinel

NEW INTEGRATION Entity resolution expanded to cover CrowdStrike SIDs, Entra GUIDs, and Microsoft 365 UPNs

NEW COMPLIANCE FRAMEWORK MITRE ATLAS mappings for all 143 AI SIEM rules: MITRE ATT&CK

NEW INTEGRATION Fleet integrations for AI platform log collection: OpenAI ChatGPT Enterprise, Anthropic, Amazon Bedrock, Google Cloud Vertex AI, Microsoft Foundry and 1 more

NEW DETECTION Instruction override (prompt injection) detection model: security_ai.instruction_override, gen_ai.input.messages

NEW DETECTION String encoding/hash detection model: security_ai.encoding, security_ai.hash_string, gen_ai.input.messages, gen_ai.output.messages

NEW DETECTION Word obfuscation detection model: security_ai.word_obfuscation, gen_ai.input.messages, gen_ai.output.messages

NEW DETECTION Refusal detection model: security_ai.refusal, gen_ai.output.messages

NEW DETECTION Prebuilt detection rules mapped to MITRE ATLAS and OWASP Top 10 for Agentic Applications 2026: OWASP Top 10 for Agentic Applications for 2026

NEW DETECTION Anomaly detection jobs for AWS Bedrock — tool-request volume, rare tool selection, high call volume: AWS Bedrock integration

NEW INTEGRATION Security for AI package integration with AWS Bedrock: AWS

Bedrock integration, version 2.2.0

NEW COMPLIANCE FRAMEWORK Detection rules mapped to OWASP Top 10 for Agentic Applications 2026: OWASP Top 10 for Agentic Applications for 2026

NEW INTEGRATION Vercel integration: Supabase, Temporal

NEW SUPPORTED PLATFORM Ransomware protection for Linux via Elastic Defend

UPDATED DETECTION MITRE ATLAS mappings added to prebuilt detection rules for AI platforms and GenAI tooling

> Press play on Elastic Security 9.6, banger after banger (2026-11-23)
>
> Elastic Security 9.6 adds AlertZero agentic SOC automation with four Watch types, custom YARA memory scanning, Entity Analytics AI brief, and agentic SIEM rule migration editing.
>
> Elastic Security adds AlertZero, an agentic layer with four Watches — Alert Triage, Threat Hunt, Detection Engineering, and Forensics — each running Workers in the background that classify alerts, hunt for threats, tune noisy rules, close detection gaps, and reconstruct attack timelines, with Manual, Assisted, or Supervised autonomy levels per Worker. Not checked yet
>
> Elastic Security adds the YARA signatures tab on the Artifacts page, letting analysts apply custom YARA signatures globally or to specific Defend policies and enabling scanning via the Memory threat protection card — a match surfaces a Memory Threat – Detected alert with the matched signature name and memory strings. Not checked yet
>
> NEW DETECTION NEW INTEGRATION Elastic Security's Automatic Migration now supports agentic editing of translated detection rules from chat, covering ES|QL logic changes, custom index patterns, custom field mappings, syntax fixes, prebuilt-rule matching, and alternative rule recommendations across Splunk, QRadar, and Microsoft Sentinel. Not checked yet
>
> Entity Analytics adds an AI brief summarising entity risk posture, naming which entities need attention, listing up to three entities to investigate first, and supporting follow-up questions in Agent Builder without rebuilding investigation context. Not checked yet
>
> Entity Analytics adds a refreshed overview with six attention tiles — Severely alerting, Newly high-risk, Risk movers, New & alerting, Anomalous, and Watchlisted & alerting — filterable by entity type, risk level, criticality, source, or watchlist across 24-hour, 7-day, and 30-day periods. Not checked yet
>
> +1 more
>
> NEW DETECTION New detection rules for AI agents and large language model platforms

NEW SUPPORTED PLATFORM Ransomware protection extended to Linux

NEW INTEGRATION Entity resolution expanded to cover CrowdStrike SIDs, Entra GUIDs, and Microsoft 365 UPNs

Elastic ships 143 detection rules for AI agents and LLM platforms (2026-11-23)

Elastic 9.6 adds Elastic Defend on-host prevention for AI coding agents (Claude Code, Cursor, Codex) and a new MITRE ATLAS mapping layer across all AI SIEM rules.

NEW DETECTION Elastic Defend behavior rules now run on-host in prevent mode for coding agents, killing offending processes before actions complete — including agent-driven persistence and malicious Cursor workspace code execution. Not checked yet

NEW COMPLIANCE FRAMEWORK All 143 AI SIEM rules ship with MITRE ATLAS technique mappings alongside MITRE ATT&CK, with Elastic 9.6 being the first release to carry those mappings into the product. Not checked yet

NEW SUPPORTED PLATFORM New AI platform coverage: OpenAI ChatGPT Enterprise, Google Cloud Vertex AI, Microsoft Foundry

NEW INTEGRATION Fleet integrations for AI platform log collection: OpenAI ChatGPT Enterprise, Anthropic, Amazon Bedrock, Google Cloud Vertex AI, Microsoft Foundry and 1 more

Thinking fast about prompt injection: New System 1 models in Elastic (2026-11-23)

Elastic Security adds a Security for AI Kibana integration package with four ingest-time NLP classifiers writing scores to `security_ai` fields, LLM-as-judge escalation for high-scoring events, and ECS `gen_ai` field compatibility.

NEW DETECTION Four lightweight NLP classifiers score GenAI traffic at ingest time inside your own cluster, writing to `security_ai.encoding`, `security_ai.hash_string`, `security_ai.word_obfuscation`, `security_ai.instruction_override`, and `security_ai.refusal` fields with no GPU required. Not checked yet

LLM-as-judge escalation is applied only to high-scoring events, keeping cost proportional to actual risk. Not checked yet

NEW DETECTION Models use ECS `gen_ai.input.messages` and `gen_ai.output.messages` fields, making the package compatible with any integration that populates those fields. Not checked yet

NEW SUPPORTED PLATFORM Available as a technical preview in Elastic Stack 9.6 on a Platinum or Trial subscription. Not checked yet

NEW DETECTION Prebuilt detection rules mapped to MITRE ATLAS and OWASP Top 10 for Agentic Applications 2026: OWASP Top 10 for Agentic Applications for 2026

NEW DETECTION Anomaly detection jobs for AWS Bedrock — tool-request volume, rare tool selection, high call volume: AWS Bedrock integration

NEW INTEGRATION Security for AI package integration with AWS Bedrock: AWS Bedrock integration, version 2.2.0

NEW INTEGRATION Security for AI package integration with OpenAI ChatGPT Enterprise Integration for Elastic

NEW COMPLIANCE FRAMEWORK Detection rules mapped to OWASP Top 10 for Agentic Applications 2026: OWASP Top 10 for Agentic Applications for 2026

announced · Nov 23, 2026

Elastic Security 9.6 adds AlertZero agentic triage watches, enhanced Entity Analytics attention signals, ES|QL S3 Data Federation, Universal Alerting (GA), and an AI SRE.

AlertZero TECHNICAL PREVIEW adds four agentic Watches — Triage, Hunt, Detection, and Forensics — to work SOC alerts continuously with configurable autonomy and model choice. Not checked yet

Entity Analytics now highlights six attention signals for changing-risk entities and resolves identities across data sources. Not checked yet

ES|QL Data Federation TECHNICAL PREVIEW queries Parquet, CSV, TSV, and NDJSON files in Amazon S3 without ingestion. Not checked yet

The `document_extraction` inference task type TECHNICAL PREVIEW converts PDFs and documents to markdown using the `jina-ocr-v1` model. Not checked yet

Universal Alerting is now generally available, storing every ES|QL-written alert as searchable data alongside Classic Alerting. Not checked yet

+6 more

NEW INTEGRATION Vercel integration: Supabase, Temporal

NEW SUPPORTED PLATFORM Ransomware protection for Linux via Elastic Defend

NEW SUPPORTED PLATFORM FIPS 140-3 support in Elastic Agent: GovCloud

Mapping the AI threat landscape: MITRE ATLAS comes to Elastic Security (2026-11-20)

Elastic Security 9.6 adds a new ATLAS coverage overview table, ATLAS support in the rule form and AI rule creation workflow, Agent Builder ATLAS awareness, and a new managed MITRE data service in Kibana.

A new ATLAS coverage overview table shows which ATLAS tactics and techniques your rules already address, scored in its own view separate from ATT&CK numbers. Not checked yet

UPDATED DETECTION The detection rule creation and edit form now supports adding both ATT&CK and ATLAS mappings to custom rules or existing prebuilt rules. Not checked yet

The AI-assisted rule creation workflow now supports adding both ATT&CK and ATLAS mappings to rules. Not checked yet

Agent Builder skills now understand ATLAS, letting you ask an agent which AI threats your installed rules cover or have it recommend prebuilt rules for specific ATLAS tactics and techniques. Not checked yet

A new managed MITRE data service in Kibana ingests MITRE ATLAS and ATT&CK entities as Saved Objects, enabling querying and filtering via a data client and API, and decoupling framework updates from Kibana release cycles. Not checked yet

---

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Qualys TotalCloud

changelog · Nov 26, 2026

CNAPP

Qualys TotalCloud is a cloud security platform providing vulnerability management, compliance monitoring, and threat detection across cloud infrastructure.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 4/5 NOVELTY 4/5

Qualys TotalCloud ships TotalAI, which red-teams any model with a chat completion endpoint against 300+ attack scenarios mapped to OWASP, MITRE ATLAS, and the EU AI Act, scans MCP servers for credential exposure, and builds an AI inventory across cloud connectors, Cloud Agents, Container Security, and a browser plug-in — making Qualys the first major vulnerability-management platform to ship an integrated AI red-teaming and inventory capability.

> Qualys TotalCloud introduced TotalAI across two releases, adding AI model red-teaming against 300+ attack scenarios, secret detection spanning MCP servers, containers and CI/CD pipelines, new AI Workloads/Containers/Code Assets inventory views, DNS-based AI router detection, and CIEM posture checks for AI service identities.

└──▷ **WHAT SHIPPED · 9 FEATURES** most completely described first ? what's the number?

01 AI Workloads view with QQL search

Not checked yet NEW

85

The new AI Workloads view shows whether AI software is installed, running, or communicating on each host, with binary, arguments, last-seen time, and Qualys Query Language (QQL) search to filter hosts.

Identify all hosts where AI software is actively communicating, so you can scope exposure after a new AI supply chain advisory.

📍 In TotalAI, open the AI Workloads view and use QQL to filter hosts by communication state, then review the binary and arguments columns to confirm scope.

— Names the view, fields, and a runnable QQL search step 2026-11-23-59981c74

02 Code Assets links CI/CD repos to deployed images

Not checked yet

NEW

85

Code Assets link CI/CD repositories to built images when QScanner runs in a CI/CD pipeline, then correlates those images with container, cloud, and Kubernetes inventory to show where the software is deployed.

Trace a vulnerable AI library from its source repository all the way to the running containers in Kubernetes to understand blast radius.

> 📍 In TotalAI, open Code Assets, select the repository, and follow the linked image to the Containers and Kubernetes inventory view to see every deployment of that image.

— Names QScanner and gives exact navigation to trace deployment `2026-11-23-59981c74`

---

03 **Red-teaming for AI models** Not checked yet NEW 75

TotalAI can red-team any model with a chat completion endpoint, including foreign models, against `300+ attack scenarios` mapped to OWASP, MITRE ATLAS, and the EU AI Act, including `21 sensitive information disclosure checks`.

— Names frameworks and scenario count but no runnable command `2026-11-26-d8ae8b2a`

---

04 **CIEM posture checks for AI identities** Not checked yet NEW 65

TotalCloud CIEM posture checks show what each cloud identity can access and flag excess privileges for AI services such as Amazon Bedrock, SageMaker, and Azure OpenAI.

— Names specific AI services and the CIEM check surface `2026-11-26-d8ae8b2a`

---

05 **AI inventory via cloud connectors and browser plug-in** Not checked yet NEW 60

TotalAI maps managed AI services, AI software, and GPU usage across cloud connectors, Cloud Agents, Container Security, and a browser plug-in to build an AI inventory without deploying a separate agent.

— Names the collection surfaces but no exact workflow `2026-11-26-d8ae8b2a`

---

06 **DNS correlation for unapproved AI endpoints** Not checked yet NEW 60

TotalCloud correlates DNS and outbound TCP events from monitored workloads to track direct connections to AI routers or provider endpoints outside an approved list.

— Explains the detection mechanism but no configuration shown `2026-11-26-d8ae8b2a`

THINNER COVERAGE BELOW

---

07 **Secret detection across MCP, CI/CD, and code** Not checked yet NEW 55

TotalAI MCP scans catch secrets exposed in MCP servers that agents source context from, and TotalAI separately detects exposed secrets in container images, CI/CD pipelines, and code repositories before they reach a model prompt.

— Names the surfaces scanned but no commands or UI path 2026-11-26-d8ae8b2a

08 **AI Containers inventory** Not checked yet NEW 50

AI Containers adds installed software, Kubernetes metadata, and linked hosts and containers to give visibility inside AI container workloads.

— Names fields added but no navigation or example given 2026-11-23-59981c74

09 **TruRisk-based AI prioritization** Not checked yet NEW 20

TotalAI combines TruRisk scores with workload context to help prioritize AI supply chain response.

— One-line prose claim with no mechanism or scope 2026-11-23-59981c74

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Bedrock Security

announced · Nov 23, 2026

Data Security & DLP

More than DSPM: Bedrock's AI and data security platform unifies data, identity, and AI context so teams see real risk — not noise — without moving data.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 4/5 NOVELTY 4/5

Bedrock Security's Agent DLP now infers the user's business intent from a session and evaluates the full sequence of agent actions against that intent and org policy at runtime — alerting, masking, or blocking when behavior puts sensitive data at risk — across Claude, Claude Code, ChatGPT, Codex, Snowflake Cortex, and AI gateways including Bifrost, LiteLLM, and Amazon Bedrock AgentCore; runtime intent inference as a DLP enforcement mechanism for AI agents has no close prior art among tracked tools.

> Bedrock Security launched Agent DLP, a new capability that infers user intent and evaluates AI agent behavior at runtime, governs a broad set of AI assistants, coding agents and agent gateways, and enforces natural-language policies against its existing DSPM data classifications.

└──▷ **WHAT SHIPPED · 4 FEATURES** most completely described first ? what's the number?

01 Natural-language policy enforcement against DSPM classifications

Not checked yet NEW 60

Security teams can write policies in natural language; Agent DLP enforces them against Bedrock Data's existing DSPM classifications covering PII, PHI, payment card data, credentials and secrets, financial records, MNPI, and proprietary source code across on-premises, SaaS, and cloud.

— Names data classes and environments but no policy syntax shown launch-20261123-8e775665

THINNER COVERAGE BELOW

02 Intent inference and runtime behavior evaluation

Not checked yet NEW 55

Agent DLP infers a user's intent (the business purpose behind a task) and evaluates the full session of agent actions against that intent and org policy at runtime, alerting, masking data, or blocking when behavior puts sensitive data at risk. It is available as a standalone product and as part of the Bedrock Data ArgusAI platform.

— Explains mechanism and outcomes but no config surface given launch-20261123-8e775665

03 **Governed AI assistants, agents, and gateways** Not checked yet

NEW 55

Agent DLP now supports Claude, Claude Code, Claude Cowork, ChatGPT, Codex, Snowflake Cortex, and Snowflake Cortex Code as governed AI assistants and coding agents, and integrates with AI gateways and agent runtimes including Bifrost, LiteLLM, Amazon Bedrock AgentCore, and NVIDIA OpenShell.

— Names every supported surface but no setup steps launch-20261123-8e775665

04 **Agent inventory and audit trail** Not checked yet NEW 45

Agent DLP inventories all agents in use, the tools, MCP servers, and identities each relies on, and keeps a full audit trail of every action for regulatory review.

— Describes capability but no interface or export details launch-20261123-8e775665

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 ifixai v4.4.0

AI Red Teaming & Testing

// WHY IT MADE THE LIST — DEPTH 4/5 · IMPACT 5/5 · NOVELTY 5/5

ifixai v4.4.0 adds `ifixai audit <agent-url>` for one-command self-check of any OpenAI-shaped agent endpoint, SARIF 2.1.0 output for CI integration, trace import from five platforms (OpenTelemetry GenAI, Langfuse, LangSmith, Arize/OpenInference, Braintrust), and `ifixai offer --add-server` to auto-install the iFixAi MCP server across 14 coding agents in one step — making adversarial AI agent auditing as accessible as running a linter.

> ifixai v4.4.0 adds a one-command self-check against a target agent with exit codes and SARIF output, trace import from five observability platforms, a free Pro audit offer with automated MCP server installation across 14 coding agents, and CI integrations for GitHub Actions, GitLab CI and pre-commit, alongside breaking changes to default models and run behavior.

└──▷ **WHAT SHIPPED · 6 FEATURES** most completely described first — ? what's the number?

01 Agent self-check via `ifixai audit`

Not checked yet · NEW

93

`ifixai audit <agent-url>` executes the `smoke` suite and any pinned findings against an OpenAI-shaped chat endpoint; pass `--yes` to confirm the sandbox prompt, set `IFIXAI_AGENT_KEY` for the agent's key, and use `--judge-provider mock` to skip external calls. Exit codes are now defined: 0 means nothing failed, 1 means a check or pinned finding failed, and 2 means no verdict was reached.

— Full command, flags, env var and exit codes all given. v4.4.0

02 Pro audit offer and MCP server auto-install across coding agents

Not checked yet · BREAKING

83

`ifixai offer` prints the free Pro audit offer for a coding agent, with `--no-promo` or `IFIXAI_NO_PROMO=1` to skip it; `ifixai offer --add-server` adds the iFixAi MCP server (`https://mcp.ifixai.ai/mcp`) to every coding agent found on the machine, covering Claude Code, Codex, Gemini CLI, GitHub Copilot CLI, Devin, Cline, Amp, opencode, Cursor, VS Code Copilot, Kiro, Junie, Warp and Antigravity. Every real `ifixai run` with a person at the terminal now opens with this offer by default; picking the Pro audit ends the run before any call to the agent and exits 0.

— Exact flags, env var, URL and 14 named agents. v4.4.0

03 New `ifixai findings` subcommands for explain, SARIF and PR comments

Not checked yet NEW 78

`ifixai findings explain` explains a self-check failure in plain language, covering what the agent was asked, what it should have done, and the framework controls it covers; `ifixai findings validate-regressions` checks that `.ifixai/regressions.yaml` is well-formed; `ifixai findings sarif` writes SARIF 2.1.0 output from self-check results; `ifixai findings comment` writes a pull-request comment.

— Four named subcommands with their exact behavior described. v4.4.0

04 CI integrations for self-check in GitHub Actions, GitLab CI and pre-commit

Not checked yet NEW 73

`actions/self-check` runs the iFixAi self-check in GitHub Actions, uploading SARIF to code scanning and optionally posting or updating a verdict comment; `ci/gitlab/ifixai-self-check.gitlab-ci.yml` runs it in GitLab CI, keeping the report and SARIF as artifacts and optionally posting the verdict on the merge request; the `.pre-commit-hooks.yaml` pre-commit integration runs the self-check on push and checks the regressions file on commit.

— Exact workflow and config file names given for each CI surface. v4.4.0

05 Trace import from five observability platforms Not checked yet

NEW 70

`ifixai traces import` imports trace data from OpenTelemetry GenAI, Langfuse, LangSmith, Arize/OpenInference and Braintrust exports into a trajectory capture; input over 64 MiB or 100,000 events is refused rather than truncated.

— Command named with exact size/event limits. v4.4.0

THINNER COVERAGE BELOW

06 Default model changes for Gemini and Bedrock Not checked yet

45

BREAKING

The Gemini default model changed from `gemini-2.0-flash` to `gemini-3.6-flash`, and the Bedrock default model changed to `global.anthropic.claude-sonnet-4-6`.

— States old and new model names, no migration guidance. v4.4.0

└──▷ BREAKING ON UPGRADE

! Every real `ifixai run` with a person at the terminal now opens with the free Pro audit offer; pass `--no-promo` or set `IFIXAI_NO_PROMO=1` to skip it, and picking the Pro audit ends the run before any call to your agent and exits 0.

! The Gemini default model changed from `gemini-2.0-flash` to `gemini-3.6-flash`.

! The Bedrock default model changed to `global.anthropic.claude-sonnet-4-6`.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 rsigma v0.28.0

Detection Engineering

// WHY IT MADE THE LIST — DEPTH 4/5 — IMPACT 4/5 — NOVELTY 4/5

RSigma adds STIX/TAXII threat-intel feeds that turn TAXII 2.1 collections and STIX bundles into dynamic source values for live detections — rebuilding referencing rules in real time as indicators arrive — and can draft Sigma rules directly from STIX indicators; combined with multi-tenancy, DNS/mmdb enrichers, and OCSF output, this is the first Sigma-native detection engine to close the loop between threat intel feeds and live rule execution.

> RSigma v0.28.0 adds multi-tenant daemon support with per-tenant rule curation via named rulesets, STIX/TAXII-driven threat-intel detections, new `dns` and `mmdb` enrichers, OCSF and UDM output formats, and batched sink delivery that cuts 100,000-event NATS JetStream and OTLP delivery times by roughly 20x — alongside breaking changes to config validation, snapshot formats, and the library API.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first — ? what's the number?

01 Batched sink delivery throughput and controls

Not checked yet — IMPROVED — 95

Batch sink delivery now takes 100,000 events to NATS JetStream from 12.21 s to 0.65 s and to OTLP from 16.65 s to 0.59 s; control batch sizing with `--sink-batch-max` and flushing with `--sink-batch-flush-ms` (and `daemon.output.batch_flush_ms`). A zero value for `--buffer-size`, `--batch-size`, or `--sink-batch-max` is now rejected, and the `--sink-batch-flush-ms` default changes from 50 to 0. OTLP rejected log records are tracked on the new `rsigma_sink_otlp_rejected_records_total{sink}` Prometheus counter.

— Concrete before/after throughput numbers and flags v0.28.0

02 New dns and mmdb enrichers

Not checked yet — NEW — 90

Adds new `dns` and `mmdb` enrichers: `dns` expands DNS fields at detection time, and `mmdb` adds MaxMind database lookups, with `mmap: true` cutting idle memory from 113 MB to 21.5 MB on a 93 MB database.

— Concrete memory numbers and enricher names given v0.28.0

03 Daemon thread-pool and stdin CLI controls

Not checked yet — NEW — 90

Adds daemon CLI controls `--worker-threads`, `--detect-threads`, and `--detect-inflight` for sizing thread pools on high-volume daemons, and `engine daemon --input stdin` to run one batch from stdin through to exit for serverless and cron callers. `--input stdin` now exits `3` when stored state cannot be restored and `4` when a run does not complete.

— Exact flags and exit codes given v0.28.0

04 **Multi-tenant daemon state and API scoping** Not checked yet NEW

85

Adds multi-tenant daemon support with correlation, dedup, incident, silence, risk, and disposition state keyed per tenant, `daemon.tenancy.quotas` capping each tenant, API routes scoped to the token's tenant, and sinks routing by `{tenant}`.

— Names config key and tenant routing but no endpoint paths v0.28.0

05 **Named rulesets with per-tenant pinning** Not checked yet NEW

85

Per-tenant rule curation and pinning via `rulesets:` in config, the new `rsigma ruleset` command group, and `--ruleset` on the `rule`, `eval`, `convert`, and `hunt` commands.

— Exact command group and flag names given v0.28.0

06 **STIX/TAXII threat-intel driven detections** Not checked yet NEW

85

Turns TAXII 2.1 collections and STIX bundles into dynamic source values for live detections, pushes indicators over the API with `watchlist` sources to rebuild referencing rules in real time, and drafts Sigma rules from STIX indicators with `rule draft --from-stix`. STIX ids outside the STIX 2.1 type-name rules no longer parse.

— Names TAXII/STIX mechanisms and a runnable CLI flag v0.28.0

07 **Convert manifest and backend build** Not checked yet NEW

70

Convert groups of rules into queries plus a per-rule manifest using the `convert:` manifest with `rsigma backend build`. LynxDB and Fibratus conversion now rejects field names outside their identifier syntax.

— Manifest format named but little mechanism explained v0.28.0

08 **Explicit enrichment miss handling** Not checked yet BREAKING

70

A lookup that finds nothing now triggers explicit `on_miss` handling, and the `http` enricher gains `default` and `miss_status` fields. An enricher with `default: null`

now injects `null` on a miss instead of the previous behavior.

— Names miss handling fields and breaking default change v0.28.0

09 **OCSF and UDM output formats** Not checked yet NEW 70

Adds new output format support: writing OCSF findings from eval, backtest, sinks, and webhooks, and canonicalizing Google SecOps UDM events with the new `udm_events.yml`.

— Names formats and pipeline stages, thin on mechanism v0.28.0

10 **Stricter config validation and snapshot version bump**

Not checked yet 65

BREAKING

Unknown keys in `rsigma.yaml` now stop every command with exit code `3`; unknown or inapplicable keys in enricher, webhook, risk, and source files fail to load. Correlation, alert-pipeline, risk, and disposition snapshots move to version 2, which older releases cannot read.

— Exit codes and version bump named, no endpoint v0.28.0

11 **Library API additions and breaking changes** Not checked yet

65

BREAKING

In the library API, `EnricherConfig` and `EnrichersConfigError` are now `#[non_exhaustive]`, `EnrichErrorKind` gains `Miss`, `CorrelationBody` gains `event_sources`, the tenancy types gain tenant fields, `RuleHeader` gains `ruleset`, `CorrelationSnapshot` gains `rulesets`, `SourceType` and `DataFormat` gain variants, `DraftError` and `StixIdError` gain variants, `TickOutput` and `RiskOutput` gain `finding_links`, and Fibratus `sanitize_field` is replaced by `field_ident`.

— Lists exact API type/field changes, no narrative v0.28.0

12 **Event tracking through chained correlations** Not checked yet

IMPROVED 60

Carries contributing `events` or `event_refs` and `event_sources` through chained correlations, naming the rule and child correlation behind each event.

— Names fields carried through, no code surface v0.28.0

└──▷ **BREAKING ON UPGRADE**

! Unknown keys in `rsigma.yaml` now stop every command with exit code `3`; unknown or inapplicable keys in enricher, webhook, risk, and source files fail to load.

! An enricher with `default: null` now injects `null` on a miss instead of the previous behavior.

- ! `--input stdin` exits `3` when stored state cannot be restored and `4` when a run does not complete.
- ! A zero value for `--buffer-size`, `--batch-size`, or `--sink-batch-max` is now rejected.
- ! LynxDB and Fibratus conversion now rejects field names outside their identifier syntax.
- ! STIX ids outside the STIX 2.1 type-name rules no longer parse.
- ! Correlation, alert-pipeline, risk, and disposition snapshots move to version 2, which older releases cannot read.
- ! The `--sink-batch-flush-ms` default changes from 50 to 0.
- ! In the library API, `EnricherConfig` and `EnrichersConfigError` are now `#[non_exhaustive]`, `EnrichErrorKind` gains `Miss`, `CorrelationBody` gains `event_sources`, the tenancy types gain tenant fields, `RuleHeader` gains `ruleset`, `CorrelationSnapshot` gains `rulesets`, `SourceType` and `DataFormat` gain variants, `DraftError` and `StixIdError` gain variants, `TickOutput` and `RiskOutput` gain `finding_links`, and Fibratus `sanitize_field` is replaced by `field_ident`.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

Where each pick stands · Nov 20–26

Tools with changes in each pick's category over the last 12 weeks, including the pick itself.



Each pick's release rhythm · Nov 20–26

Changes per week for the last 12 weeks. The bright bar is this week; a dash marks a week with no daily.

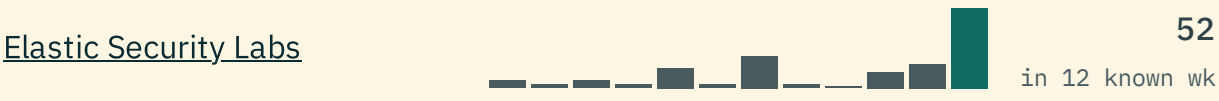

Qualys TotalCloud	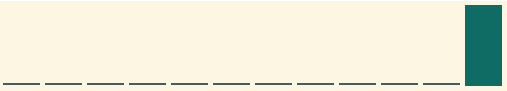	2 in 12 known wk
Bedrock Security	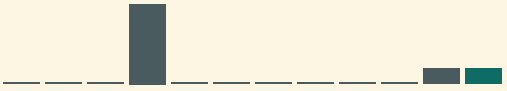	7 in 12 known wk
ifixai		5 in 12 known wk
rsigma		7 in 12 known wk

Which categories moved · Nov 20–26

The thick bar is this week. The thin bar is the average of the four weeks before, counting only weeks with a daily.

Code Security		80 usual 78.2
Identity & Access Management		74 usual 61.8
SIEM		57 usual 57.8
Autonomous Offensive Frameworks (AOF)		54 usual 23
Vulnerability & Exposure Management		40 usual 19.2
Endpoint Security (EDR/XDR)		30 usual 15.8
Software Supply Chain Security		18 usual 67
Cloud Security Posture		15 usual 11.2
Endpoint Management (UEM/RMM)		15 usual 17.8
Model & Agent Protection		14 usual 14.2
CNAPP		14 usual 8

Privileged Access Management		14 usual 5.2

Who shipped the most · Nov 20–26

Changes per tool this week. Volume is not a quality score.



Upcoming conferences →

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · November 27, 2026

The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe