THE CYBER TOOLCHAIN

NO. 014

Head

THE WEEK'S TOP PICKS

PUBLISHED OCTOBER 2, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

▸ // HOW THESE PICKS ARE MADE SHOW

▾ // THE BRIEF HIDE

What stands out in this week's picks, grouped by what it lets you do. Every tool named links to its pick below.

Patching gets a brake pedal: Qualys Patch Management can now halt a bad patch fleet-wide and keep low-confidence ones out of zero-touch jobs. Qualys VMDR scans for a fresh CVE within hours of disclosure and ranks it against the controls you actually run. And Metasploit Pro adds native Kerberos relay with ESC8 , so AD CS coercion paths are testable in-tool.

PROTECT

Stop a bad patch before it reaches the rest of the fleet

Blocking rules halt a patch across the estate mid-rollout, reliability scoring keeps shaky updates out of unattended jobs, and curated mitigations cover the window while a patch is held — the compensating-control hunt that normally follows a deferred patch is already done.

Qualys Patch Management

PROTECT

Know within hours whether a new CVE is actually exploitable on your

vim ? help

InstaScan starts detecting right after a vendor discloses rather than waiting on the next scan window, and prioritization tests exploitability against the controls already deployed. That collapses the usual gap between advisory and first real answer, and shrinks the list of findings that someone has to argue down by hand.

Qualys VMDR

PROTECT

Test AD CS relay and coercion paths without stitching together external tooling

Native Kerberos authentication relay with ESC8 exploitation and IPv6 DNS-takeover coercion lands in-framework, alongside OPC-UA and WinRM scanning and AArch64/MIPS payloads. Red teams stop bridging to standalone relay tools mid-engagement, and the JSON RPC endpoint now requires authentication by default.

Metasploit Pro

GOVERN

Give agents and models an identity you can revoke

DigiCert issues AI Passports to agents, models and MCP servers in NVIDIA OpenShell with cryptographic identity, BOM scanning, signed audit trails and a cross-boundary kill switch; Bedrock adds inline Agent DLP enforcement at runtime plus an AI Data Bill of Materials. The practical shift is that an agent becomes something you can inventory, attest to an auditor, and cut off mid-run.

DigiCert CertCentral · Bedrock Security

PROTECT

Block sensitive data on its way into an unsanctioned AI tool

Purview moves DLP enforcement to the network layer, catching paste-into-chatbot flows that endpoint and app-level policy miss; Sublime adds pre-delivery blocking for email DLP with AI triage. Both shift the control point earlier — before the data leaves, not in an after-the-fact report.

Microsoft Purview

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | **EVERYTHING**

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

all 34 changes

Do you prefer **scroll**? 👍 👎

// THE PICKS

01	**Cribl**	Security Data Pipeline
02	**Qualys VMDR**	Vulnerability Management
03	**Qualys Patch Management**	Vulnerability Management
04	**Metasploit Pro**	Exploitation & C2
05	**Bedrock Security**	Data Security & DLP

01 Cribl

changelog · Sep 29, 2026

Security Data Pipeline

Cribl is an observability pipeline platform that helps organizations collect, process, route, and transform machine data from any source to any destination.

// WHY IT MADE THE LIST

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

Cribl Detect launches a full SIEM built on the data-pipeline platform — 8,000+ detection rules, detection posture management showing coverage gaps and broken rules, and in-pipeline detection that generates signals while telemetry is in motion, priced on infrastructure rather than per-GB. A pipeline vendor entering the SIEM market lets its customers consolidate detection where their data already flows.

> At CriblCon 26, Cribl launched two major platform additions: Cribl Detect, a full SIEM built on its pipeline, and StreamAI, an enterprise AI gateway — alongside a new agentic development environment, broader agent/integration access via MCP , automatic telemetry parsing, and a growing app marketplace.

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

01 Cribl Detect SIEM with detection and posture management NEW

87

Cribl Detect is a complete SIEM built on the Cribl platform covering detection, coverage and posture management, investigation, response, and threat intelligence, available now for Cribl.Cloud customers at up to 50% less cost than incumbent SIEMs, priced on infrastructure rather than per-GB ingestion. It includes Detection Posture Management that continuously shows coverage gaps, broken rules, and missing telemetry required to detect specific attacks, ships with a catalog of more than 8,000 detection rules with correlation across signals to surface multi-stage attacks, and generates detection signals directly in the data pipeline for real-time identification while telemetry is still in motion.

Get started with Cribl Detect on Cribl.Cloud to immediately access SIEM detection, posture management, and 8,000+ detection rules without migrating existing telemetry pipelines.

> 📍 In Cribl.Cloud, log in to your organization, navigate to the Cribl Detect tile, and select Get Started to activate detection on your existing telemetry streams.

— Pricing model, rule count, and get-started steps all named 2026-09-29-1d8f4bc7

02 StreamAI enterprise AI gateway with model routing NEW

60

StreamAI is an enterprise AI gateway with a built-in model router that routes each prompt to the right model based on SecIT Bench benchmarks, offering free inference for customers using automatic routing. It enforces spending budgets and cutoffs, redacts sensitive data in both prompts and responses, and records every model call as audit-ready telemetry.

— Mechanism and controls named but no setup steps given 2026-09-29-1d8f4bc7

THINNER COVERAGE BELOW

03 Agentic development environment on Detect's runtime NEW

48

A new agentic development environment lets teams build apps, workflows, and agents using natural language, running on the same agentic runtime that powers Cribl Detect, with compute, scheduling, logging, and permissions included.

— Describes capability and shared runtime, no usage detail 2026-09-29-1d8f4bc7

04 Cribl App Framework grows to 150+ apps IMPROVED

40

The Cribl App Framework has grown to more than 150 apps from Cribl, partners, and customers, available in the Cribl Marketplace.

— Names count and marketplace location, no further mechanism 2026-09-29-1d8f4bc7

05 Platform opened to agents via MCP and SDKs NEW

35

Cribl opens the platform to agents via MCP and SDKs, adding 70 new integrations this year.

— Names MCP and integration count but no how-to 2026-09-29-1d8f4bc7

06 AI auto-parsing of incoming telemetry NEW

20

Cribl adds AI auto-parsing of telemetry so incoming data is structured and usable the moment it arrives.

— Describes outcome only, no mechanism or configuration 2026-09-29-1d8f4bc7

changelog · Sep 29, 2026

Cribl launches StreamAI (AI gateway with model routing) and Cribl Detect (full SIEM with 8,000+ detection rules) at CriblCon 26.

└──▷ HOW TO FIND IT

Get started with Cribl Detect on Cribl.Cloud to immediately access SIEM detection, posture management, and 8,000+ detection rules without migrating existing telemetry pipelines.

📍 In Cribl.Cloud, log in to your organization, navigate to the Cribl Detect tile, and

select `Get Started` to activate detection on your existing telemetry streams.

- Adds Cribl Detect, a complete SIEM built on the Cribl platform with detection, coverage and posture management, investigation, response, and threat intelligence — available now for Cribl.Cloud customers at up to 50% less cost than incumbent SIEMs, priced on infrastructure rather than per-GB ingestion.
- Adds StreamAI, an enterprise AI gateway with a built-in model router that routes each prompt to the right model based on SecIT Bench benchmarks, with free inference for customers using automatic routing.
- StreamAI enforces spending budgets and cutoffs, redacts sensitive data in both prompts and responses, and records every model call as audit-ready telemetry.
- Cribl Detect includes Detection Posture Management, continuously showing coverage gaps, broken rules, and missing telemetry required to detect specific attacks.
- Cribl Detect ships with a catalog of more than 8,000 detection rules covering a broad range of threats, with correlation across signals to surface multi-stage attacks.

`+5 more`

GOOD PICK?  GOOD REASON? 

02 Qualys VMDR

changelog · Sep 29, 2026 | Vulnerability Management

Qualys VMDR is a cloud-based vulnerability management platform that continuously discovers, assesses, prioritizes, and remediates security vulnerabilities across IT assets.

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

InstaScan delivers a vulnerability finding within an hour of a vendor disclosure — closing the window where exposure is live but unassessed — and pairs it with prioritization that tests exploitability against the environment's actual firewall and EDR controls rather than theoretical CVSS. No tracked tool ships either capability.

> Qualys VMDR introduced InstaScan for near-instant vulnerability detection following vendor disclosures, paired with hyper-prioritization that tests exploitability against an environment's actual security controls.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first — ? what's the number?

01 InstaScan scanless detection NEW

45

InstaScan is a scanless scanning capability that delivers a finding within an hour of a vendor disclosure, enabling much faster detection than traditional scan cycles.



— Describes mechanism and timing but no config or command 2026-09-29-9128b856

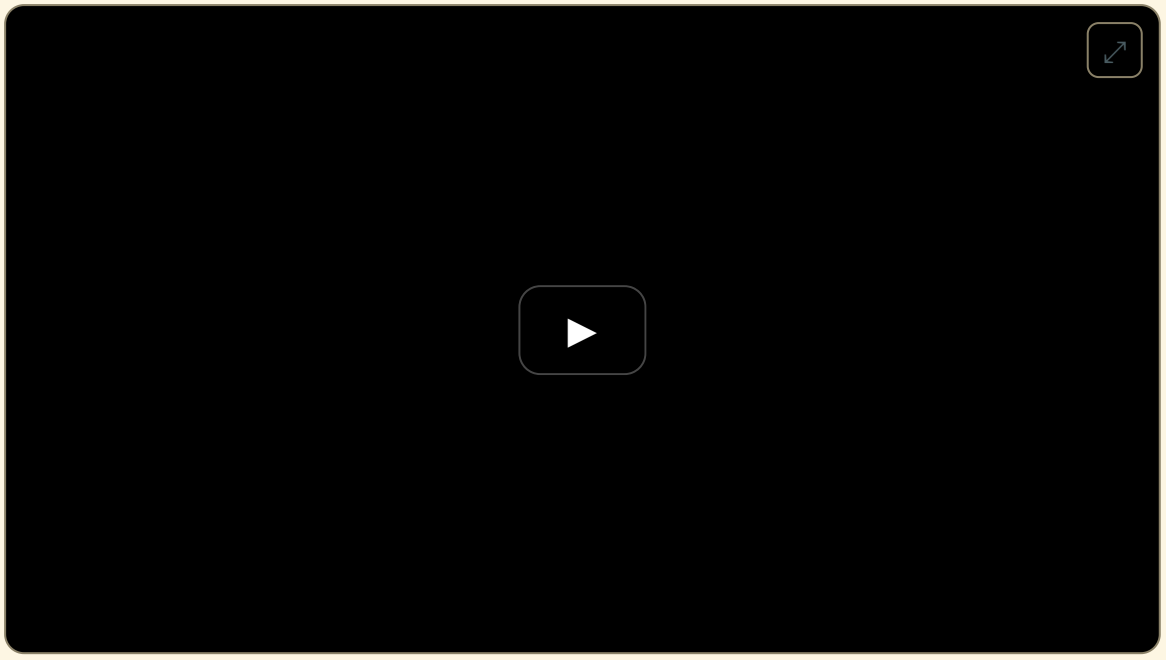
02 **Hyper-prioritization against live controls** NEW 40

Hyper-prioritization tests exploitability against a given environment's actual firewall and EDR controls, rather than relying on theoretical scoring, to better prioritize remediation.

— Explains approach but no concrete settings or workflow steps 2026-09-29-9128b856

changelog · Sep 29, 2026

Qualys VMDR adds InstaScan for sub-hour detection from vendor disclosure, with hyper-prioritization against live controls.



- Introduces InstaScan, a scanless scanning capability that delivers a finding within an hour of a vendor disclosure.
- Adds hyper-prioritization that tests exploitability against a given environment's actual firewall and EDR controls, rather than using theoretical scoring.

GOOD PICK? GOOD REASON?

03 Qualys Patch Management

changelog · Sep 29, 2026

Vulnerability Management

Qualys Patch Management is a cloud-based tool that automates the discovery, prioritization, and deployment of patches across IT assets to reduce vulnerabilities.

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

Patch Blocking Rules let you define a bad patch by QQL query and block it across every deployment job at once, and AI-Powered Patch Reliability Scoring rates each patch from real-world internet feedback plus Qualys rollback and vuln-reopen telemetry — automatically skipping low-reliability patches from zero-touch jobs so a known-bad update never reaches production.

> Qualys TruRisk Eliminate adds guardrails to automated patching: blocking rules to halt bad patches fleet-wide, AI-driven reliability scoring to keep low-confidence patches out of zero-touch jobs, and curated mitigations to cover the gap while patches are held.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first — ? what's the number?

01 Patch Blocking Rules in TruRisk Eliminate NEW — 88

A new Patch Blocking Rules page under `TE > Configuration > Patch Blocking Rules` (Windows, Linux, and Mac support added in release 4.2) lets admins define patches to block via a QQL query (e.g. `patch.kb: [KB5002907, KB5124008]`) and block them across every deployment job simultaneously, without editing individual jobs. Scope can be restricted to up to five asset tags per rule, up to 10 rules are allowed per platform, and each rule requires a name and a blocking reason.

Immediately stop a vendor-paused or known-bad Windows patch from deploying across your entire environment without touching individual patch jobs.

📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. Block paused Sep-2026 updates) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907, KB5124008, KB5124012, KB5123099] and click Preview to confirm the match. Leave asset tags empty to cover all assets, then save.

Protect production assets while a test group validates a new patch by scoping a blocking rule to production asset tags.

📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. `Hold KB5002907 from prod during validation`) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907] and click Preview. Under asset scope, add up to five production asset tags to restrict the block to production only, then save.

— Names exact menu path, QQL syntax, and limits `2026-09-29-7aca09ff`

---

02 **AI-Powered Patch Reliability Scoring** NEW 70

Combines LLM analysis of real-world internet feedback with Qualys rollback and vulnerability-reopen telemetry to rate each patch as High, Medium, Low, or Unidentified. When reliability intelligence is added to a zero-touch job's QQL in the `Select Patches` step, Low-reliability patches are automatically skipped from deployment while others continue.

Keep zero-touch patch jobs running while automatically skipping low-reliability patches by adding reliability intelligence to the job's patch-selection query.

📍 In the zero-touch patch job's Select Patches step, update the QQL to include reliability criteria alongside your existing filters so that Low-reliability patches — such as those currently rated Low by AI-Powered Patch Reliability Scoring — are excluded from automatic deployment while all other patches continue to deploy.

— Explains rating mechanism and where to apply it `2026-09-29-7aca09ff`

THINNER COVERAGE BELOW

03 **Curated mitigations for blocked or risky patches** NEW 23

Qualys-curated mitigations are introduced as a complementary layer to reduce exposure while a blocked or low-reliability patch is held in testing and staging.

— Described only in prose, no mechanism or surface named `2026-09-29-7aca09ff`

`changelog · Sep 29, 2026`

Qualys TruRisk Eliminate adds Patch Blocking Rules and AI-Powered Patch Reliability Scoring to keep bad patches out of automation.

.
.
.
.

└──▷ **HOW TO FIND IT**

Immediately stop a vendor-paused or known-bad Windows patch from deploying across your entire environment without touching individual patch jobs.

📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. `Block paused Sep-2026 updates`) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907, KB5124008, KB5124012, KB5123099] and click Preview to confirm the match. Leave asset tags empty to cover all assets, then save.

Protect production assets while a test group validates a new patch by scoping a blocking rule to production asset tags.

📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. `Hold KB5002907 from prod during validation`) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907] and click Preview. Under asset scope, add up to five production asset tags to restrict the block to production only, then save.

Keep zero-touch patch jobs running while automatically skipping low-reliability patches by adding reliability intelligence to the job's patch-selection query.

📍 In the zero-touch patch job's Select Patches step, update the QQL to include reliability criteria alongside your existing filters so that Low-reliability patches — such as those currently rated Low by AI-Powered Patch Reliability Scoring — are excluded from automatic deployment while all other patches continue to deploy.

- Adds `Patch Blocking Rules` in `TE > Configuration > Patch Blocking Rules` (available for Windows, Linux, and Mac in release 4.2), letting you define patches via a QQL query and block them across every deployment job at once without editing individual jobs; scope is controlled by up to five asset tags per rule, with up to 10 rules per platform, and each rule requires a name and blocking reason.
- Adds AI-Powered Patch Reliability Scoring that combines LLM analysis of real-world internet feedback with Qualys rollback and vulnerability-reopen telemetry to rate each patch High, Medium, Low, or Unidentified — automatically skipping Low-reliability patches from zero-touch jobs when reliability intelligence is added to the job's QQL in the `Select Patches` step.
- Introduces Qualys-curated mitigations as a complementary layer to reduce exposure while a blocked or low-reliability patch is held in testing and staging.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Metasploit Pro

October 1, 2026: 5.1.0-2026093001

Exploitation & C2

Metasploit Pro is a commercial penetration testing framework for identifying, exploiting, and validating security vulnerabilities in systems.

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

Native Kerberos authentication relay lands in the framework's relay stack, including the esc8_kerberos module for AD CS Web Enrollment and IPv6/DHCPv6 DNS-takeover coercion modules to drive the chain — alongside OPC-UA enumeration and Windows AArch64 payloads across 41 new modules.

> Metasploit Pro 5.1.0 ships 41 new modules headlined by native Kerberos authentication relay (with ESC8 exploitation and IPv6 DNS-takeover coercion), new OPC-UA and WinRM scanning capabilities, Windows AArch64 and MIPS payload support, and now requires authentication on JSON RPC by default.

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

01 WinRM login session type selection — IMPROVED — 85

Adds a `SessionType` option to `auxiliary/scanner/winrm/winrm_login` to create a cmd-backed shell, a PSRP-backed PowerShell session, or automatically fall back to PowerShell when cmd shell creation is denied.

On a WinRM login attempt, automatically fall back to a PowerShell session when cmd shell creation is denied — useful against hardened Windows targets.

```
$ use auxiliary/scanner/winrm/winrm_login
set RHOSTS 192.168.1.50
set USERNAME administrator
set PASSWORD P@ssw0rd
set SessionType powershell
run
```

— Exact option and module path plus a runnable fallback example

October 1, 2026: 5.1.0-2026093001

02 Kerberos authentication relay stack — NEW — 80

The relay stack gains native Kerberos authentication relay capabilities, including the new `esc8_kerberos` auxiliary module for exploiting AD CS Web Enrollment (ESC8), `auxiliary/spoof/ipv6/ipv6_ra_dns_takeover` and

`auxiliary/spoof/dhcp/dhcpv6_dns_takeover` for native IPv6 DNS-takeover coercion as part of a Kerberos relay attack chain, `KerberosTicketTrace` support added to `auxiliary/admin/kerberos/forge_ticket`, and a new shared low-level implementation `Rex::Proto::Gss::KerberosToken` for handling Kerberos tokens carried by GSS-API and SPNEGO.

— Names every module and class but gives no usage example

October 1, 2026: 5.1.0-2026093001

---

03 **OPC-UA server enumeration scanner** NEW 80

Adds `auxiliary/scanner/scada/opcua_enum`, a scanner module that detects OPC-UA servers speaking the OPC-UA TCP binary transport (`opc.tcp://`), useful for enumerating exposed industrial control endpoints.

Scan an ICS network segment for OPC-UA servers exposing the binary TCP transport to enumerate exposed industrial control endpoints.

```
$ use auxiliary/scanner/scada/opcua_enum
set RHOSTS 10.10.20.0/24
run
```

— Module path, protocol scheme, and runnable scan command given

October 1, 2026: 5.1.0-2026093001

---

04 **SPIP blind SQL injection scanner** NEW 65

Adds `auxiliary/scanner/http/spip_annee_sqli`, which exploits a blind SQL injection in SPIP's date column escaping logic.

— Exact module path and vulnerability mechanism named, no example run

October 1, 2026: 5.1.0-2026093001

---

05 **Windows AArch64 payload support** NEW 60

Adds Windows AArch64 staged shell payloads and extends `windows/smb/psexec` to work with AArch64 targets.

— Names the module and architecture, no usage walkthrough

October 1, 2026: 5.1.0-2026093001

---

06 **ZoomEye search module v2 API support** IMPROVED 60

Updates the ZoomEye search module to support the ZoomEye v2 API, adds a three-attempt retry limit, and caps results per page at 10.

— Concrete limits named but no sample invocation October 1, 2026: 5.1.0-2026093001

THINNER COVERAGE BELOW

07 **JSON RPC requires authentication by default** BREAKING

55

JSON RPC support now requires authentication by default, either via the database with user credentials or an auth token; existing unauthenticated JSON RPC clients will be rejected on upgrade.

– States the new requirement but no exact config key or migration step

October 1, 2026: 5.1.0-2026093001

08 **MIPS payload additions and updates** IMPROVED

45

Adds a MIPS64 exec payload, updates the MIPSEL and MIPSBE exec payloads, and adds prepends for all three architectures.

– Names architectures but no usage context given October 1, 2026: 5.1.0-2026093001

09 **msfconsole usability and startup improvements** IMPROVED

45

`msfconsole` adds `MSFLOGO=nyan-cat.txt` environment variable support for selecting the nyan cat banner, a new command-line flag to print the count of modules, and improved start-up performance.

– Env var named exactly, but the new flag's name is not given

October 1, 2026: 5.1.0-2026093001

10 **Persistence module updates** IMPROVED

45

Adds a new multi-platform persistence module for Kate, and updates existing persistence modules to use `create_process` instead of `cmd_exec`.

– Names the mechanism swap but little on the new Kate module

October 1, 2026: 5.1.0-2026093001

11 **New reconnaissance and evasion modules** NEW

40

Adds a scanner module that enumerates ports on a host to determine whether each is a Metasploit reverse handler and what kind of shell it expects, an auxiliary scanner that fingerprints Inductive Automation Ignition gateways across all major version families by probing unauthenticated info endpoints, and a Linux x64 sandbox-evasion module that performs runtime environment checks and aborts execution when a likely sandbox or VM is detected.

– Mechanisms described but no module paths given October 1, 2026: 5.1.0-2026093001

12 **Windows session writable check in post mixin** NEW

30

Adds support to the post mixin for Windows sessions to run the `.writable?` method.

– Bare mention of the method with no usage detail October 1, 2026: 5.1.0-2026093001

└──▷ **BREAKING ON UPGRADE**

! JSON RPC now requires authentication by default (via database credentials or auth token); existing unauthenticated JSON RPC clients will be rejected on upgrade.

October 1, 2026: 5.1.0-2026093001

Metasploit Pro 5.1.0-2026093001 adds 41 new modules including OPC-UA , Kerberos relay, IPv6 DNS-takeover, and Windows AArch64 payloads.

└──▷ **TRY IT**

Scan an ICS network segment for OPC-UA servers exposing the binary TCP transport to enumerate exposed industrial control endpoints.

```
$ use auxiliary/scanner/scada/opcua_enum
set RHOSTS 10.10.20.0/24
run
```

On a WinRM login attempt, automatically fall back to a PowerShell session when cmd shell creation is denied — useful against hardened Windows targets.

```
$ use auxiliary/scanner/winrm/winrm_login
set RHOSTS 192.168.1.50
set USERNAME administrator
set PASSWORD P@ssw0rd
set SessionType powershell
run
```

› Adds `auxiliary/scanner/scada/opcua_enum`, a scanner module that detects OPC-UA servers speaking the OPC-UA TCP binary transport (`opc.tcp://`).

› Adds `SessionType` option to `auxiliary/scanner/winrm/winrm_login` to create a cmd-backed shell, a PSRP-backed PowerShell session, or automatically fall back to PowerShell when cmd shell creation is denied.

› Adds `MSFLOGO=nyan-cat.txt` environment variable support to `msfconsole` for selecting the nyan cat banner.

› Adds a new command-line flag to `msfconsole` to print the count of modules.

› Updates JSON RPC support to now require authentication by default, either via the database with user credentials or an auth token.

+16 more

└──▷ **BREAKING ON UPGRADE**

! JSON RPC now requires authentication by default (via database credentials or auth token); existing unauthenticated JSON RPC clients will be rejected on upgrade.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 Bedrock Security

announced · Sep 29, 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

Agent DLP sits inline at the agent gateway via AWS AgentCore and LiteLLM hooks, inspecting every tool call in both directions to mask, redact, or block sensitive data in real time — closing the gap between declared policy and actual agent behavior, with an observe mode to validate coverage first.

> Bedrock Security shipped a wave of AI governance capability: a new ArgusAI product (AI Data Bill of Materials, Guardrail Gap Analysis, natural-language policy and investigation), a graph-backed Metadata Lake with Copilot, AI Agents for DSPM and a free Snowflake tier, an MCP Server exposing 50+ metadata signals to AI agents, and inline Agent DLP runtime enforcement extended from AWS AgentCore/LiteLLM to NVIDIA OpenShell.

└──▷ **WHAT SHIPPED · 11 FEATURES** most completely described first ? what's the number?

01 Agent DLP extended to NVIDIA OpenShell NEW

77

Extends Agent DLP to NVIDIA OpenShell via OpenShell's Supervisor Middleware, evaluating each agent action against enterprise data policies and returning an allow/block decision enforced at the sandbox boundary. Supports an observe-only mode to validate coverage before enforcing, exposes an agent's blast radius (what data its permissions can reach, what combinations it forms, where policy would be violated) at design time, and applies the same data policies across every invocation path an agent may use — MCP tools, command-line invocations, and agent-written code — from a single OpenShell control point. Restrictions travel across table writes, tool handoffs, and agent-to-agent transfers based on the Metadata Lake knowledge graph rather than payload pattern matching.

— Detailed mechanism and scope, no runnable command given launch-20260928-00b820b2

02 Natural Language Policy and Investigation in ArgusAI NEW

76

Adds Natural Language Policy and Investigation in ArgusAI, letting teams articulate data controls in plain English (e.g., "Only members of HR may view employment data" or "EU personal data must remain within EU regions") and have them translated into enforceable controls across connected systems. Extends the same natural-language interface to investigation over the DBOM and guardrail context, so queries like "Which

models can access PHI data?" or "Which identities interacted with customer PII through an agent?" return answers with lineage and access evidence from the Metadata Lake.

Investigate which AI models in your Amazon Bedrock environment can access protected health information, using plain-English natural language queries.

> 📍 In the Bedrock Data Platform, navigate to ArgusAI › Natural Language Investigation and enter: 'Which models can access PHI data?' — the system returns results with lineage and access evidence from the Metadata Lake.

Enforce a data residency policy across AI systems without writing control-plane code — useful for EU compliance requirements.

> 📍 In the Bedrock Data Platform, go to Natural Language Policy and enter: 'EU personal data must remain within EU regions.' — Bedrock Data translates the invariant into enforceable controls across connected systems including AI agents and Knowledge Bases.

— Includes exact UI navigation and example queries to run `launch-20260929-1fe46b0a`

---

03 **Agent DLP runtime enforcement in ArgusAI** NEW 73

Adds Agent DLP to ArgusAI with native hooks for `AWS AgentCore` and `LiteLLM`, sitting inline at the agent gateway to inspect every tool call in both directions (outbound requests to MCP tools and inbound tool responses) without acting as a proxy. Enforces data access and regulatory policies to mask, redact, or block sensitive data exposure in real time, logs every decision with target, action, data types, and verdict (`blocked`, `modified`, `redacted`, or `observed`), and supports an `observe` mode for monitoring-only enforcement scenarios.

— Mechanism and verdict values named but no setup command `launch-20260929-0b6583d7`

---

04 **MCP Server for the Metadata Lake** NEW 65

Adds an MCP (Model Context Protocol) Server as a standardized gateway between AI agents and Bedrock's metadata lake, exposing data sensitivity, risk, lineage, access, and usage signals via the open MCP protocol. Surfaces over 50 distinct metadata signals as queryable context for agentic AI workflows and automated governance pipelines, letting AI agents and external applications — including legal and compliance teams — query the metadata lake without custom point-to-point integrations. Builds on the existing Bedrock Metadata Lake Copilot and audits AI model interactions with enterprise data for traceability and policy alignment.

— Protocol and signal count named, no endpoint or command shown `launch-20260929-f953fb1a`

THINNER COVERAGE BELOW

---

05 **Bedrock Metadata Lake** NEW 58

Introduces the Bedrock Metadata Lake, a graph-backed centralized repository that continuously catalogs enterprise data stores across SaaS, PaaS, and IaaS, capturing location, sensitivity, entitlements, lineage, usage, risk, and more than fifty additional parameters.

— Clear scope and scale, no access steps given `launch-20260929-d9dffb84`

06 **Bedrock Free for Snowflake** NEW 55

Launches Bedrock Free for Snowflake, a free tier delivering automated sensitive-data discovery and classification for Snowflake environments, including access to Metadata Lake Copilot, with upgrade paths to entitlement chain and usage analysis and automated policy capabilities.

— Names tier and target platform, no setup steps given `launch-20260929-d9dffb84`

07 **AI Agents for DSPM** NEW 50

Introduces Bedrock AI Agents for DSPM, which proactively analyze the metadata lake for policy violations and new risks, gather context automatically, suggest next steps, and guide users through incident-response resolution workflows.

— Describes behavior but no concrete workflow steps `launch-20260929-d9dffb84`

08 **Guardrail Gap Analysis in ArgusAI** NEW 48

Adds Guardrail Gap Analysis in ArgusAI, which compares policy intent against model data exposure drawn from the Metadata Lake and generates targeted remediations to suppress or mask sensitive data at the source or agent runtime.

— Describes comparison logic but no concrete interface shown `launch-20260929-1fe46b0a`

09 **ArgusAI AI Data Bill of Materials (DBOM)** NEW 43

Adds ArgusAI's AI Data Bill of Materials (DBOM), which automatically links Amazon Bedrock Custom Models, Knowledge Bases (RAG), and Agents to the datasets used during training and inference.

— States what it does, no mechanism or usage steps `launch-20260929-1fe46b0a`

10 **Metadata Lake Copilot** NEW 38

Launches Bedrock Metadata Lake Copilot, a natural-language interface for querying the metadata lake to identify sensitive assets, riskiest users, and other data-risk questions.

— Brief description, no example query or navigation `launch-20260929-d9dffb84`

11 **Metadata Lake API integration and detection capabilities** NEW 38

Exposes an API-first architecture enabling teams to implement and enforce uniform data-security policies across heterogeneous enterprise infrastructure and to integrate the metadata lake with SIEMs, CNAPPs, and DLPs. Also supports detection of sensitive

data used in AI initiatives and plain-text secrets, and provides external sharing discovery.

— Lists capabilities thinly, no endpoint or config named `launch-20260929-d9dffb84`

announced · Sep 29, 2026

Bedrock Data launches Agent DLP inside ArgusAI — inline runtime inspection of every AI agent tool call via AWS AgentCore and LiteLLM hooks.

- Adds Agent DLP to ArgusAI with native hooks for `AWS AgentCore` and `LiteLLM`, sitting inline at the agent gateway to inspect every tool call in both directions without acting as a proxy.
- Inspects both outbound agent requests to MCP tools and inbound tool responses, enforcing data access and regulatory policies to mask, redact, or block sensitive data exposure in real time.
- Logs every tool call decision with its target, action, data types involved, and verdict (`blocked`, `modified`, `redacted`, or `observed`), producing a continuous audit-ready record of agent activity.
- Supports an `observe` mode that logs agent actions — such as a marketing agent returning customer emails — without blocking, for low-risk or monitoring-only enforcement scenarios.
- Extends ArgusAI posture coverage (data store inventory, identity mapping, natural-language access policies) to runtime enforcement, closing the gap between declared policy and actual agent behavior.

announced · Sep 29, 2026

Bedrock Data launches ArgusAI to link AI models to their data lineage, assess guardrail gaps, and enforce plain-English policies across AI systems.

└──▷ HOW TO FIND IT

Investigate which AI models in your Amazon Bedrock environment can access protected health information, using plain-English natural language queries.

> 📍 In the Bedrock Data Platform, navigate to ArgusAI › Natural Language Investigation and enter: 'Which models can access PHI data?' — the system returns results with lineage and access evidence from the Metadata Lake.

Enforce a data residency policy across AI systems without writing control-plane code — useful for EU compliance requirements.

> 📍 In the Bedrock Data Platform, go to Natural Language Policy and enter: 'EU personal data must remain within EU regions.' — Bedrock Data translates the invariant into

enforceable controls across connected systems including AI agents and Knowledge Bases.

- Adds ArgusAI, a new product that automatically generates an AI Data Bill of Materials (DBOM) by linking Amazon Bedrock Custom Models, Knowledge Bases (RAG), and Agents to the datasets used during training and inference.
- Adds Guardrail Gap Analysis that compares policy intent against model data exposure — drawn from the Metadata Lake — and generates targeted remediations to suppress or mask sensitive data at the source or agent runtime.
- Adds Natural Language Policy and Investigation, enabling teams to articulate data controls in plain English (e.g., 'Only members of HR may view employment data' or 'EU personal data must remain within EU regions') and have them translated into enforceable controls across connected systems.
- Extends Natural Language Policy over the DBOM and guardrail context so investigators can query AI governance in plain English — for example, 'Which models can access PHI data?' or 'Which identities interacted with customer PII through an agent?' — with answers derived from the Metadata Lake including lineage and access evidence.

announced · Sep 29, 2026

Bedrock Security launches a Metadata Lake with Copilot, AI Agents for DSPM , and a free Snowflake edition for data discovery and classification.

- Introduces the Bedrock Metadata Lake, a graph-backed centralized repository that continuously catalogs enterprise data stores across SaaS, PaaS, and IaaS, capturing location, sensitivity, entitlements, lineage, usage, risk, and more than fifty additional parameters.
- Launches Bedrock Metadata Lake Copilot, a natural-language interface for querying the metadata lake to identify sensitive assets, riskiest users, and other data-risk questions.
- Introduces Bedrock AI Agents for DSPM, which proactively analyze the metadata lake for policy violations and new risks, gather context automatically, suggest next steps, and guide users through incident-response resolution workflows.
- Launches Bedrock Free for Snowflake, a free tier that delivers automated sensitive-data discovery and classification for Snowflake environments, including access to Metadata Lake Copilot, with upgrade paths to entitlement chain and usage analysis and automated policy capabilities.
- Exposes an API-first architecture enabling teams to implement and enforce uniform data-security policies across heterogeneous enterprise infrastructure and integrate the metadata lake with SIEMs, CNAPPs, and DLPs.

+1 more

announced · Sep 29, 2026

Bedrock Security announces an MCP Server providing standardized AI-agent access to its metadata lake, covering 50+ data-context signals.

- Adds an MCP (Model Context Protocol) Server as a standardized gateway between AI agents and Bedrock's metadata lake, exposing data sensitivity, risk, lineage, access, and usage signals via the open MCP protocol.
- Surfaces over 50 distinct metadata signals — including sensitivity classification, access patterns, and lineage — as queryable context for agentic AI workflows and automated governance pipelines.
- Enables AI agents and external applications to query the metadata lake without custom point-to-point integrations, allowing teams such as legal and compliance to run data investigations without engineering bottlenecks.
- Builds on the existing Bedrock Metadata Lake Copilot, extending its natural-language metadata querying to agentic and automated workflow contexts.
- Audits AI model interactions with enterprise data through the MCP Server to support traceability and policy alignment in agentic workflows.

announced · Sep 28, 2026

Bedrock Security extends Agent DLP to NVIDIA OpenShell, adding data-aware policy enforcement at the agent sandbox boundary.

- Extends Agent DLP to NVIDIA OpenShell via OpenShell's Supervisor Middleware, evaluating each agent action against enterprise data policies and returning an allow/block decision that OpenShell enforces at the sandbox boundary.
- Supports an observe-only mode that records policy decisions without enforcing them, letting teams validate coverage before turning on enforcement.
- At design time, exposes an agent's blast radius — what data its permissions can reach, what combinations that data forms, and where policy would be violated — giving security teams pre-deployment visibility.
- Applies the same data policies across every tool invocation path an agent may use — MCP tools, command-line invocations, and code the agent writes itself — from a single OpenShell control point.
- Carries data restrictions across table writes, tool handoffs, and agent-to-agent transfers based on the Metadata Lake knowledge graph rather than payload pattern matching, so restrictions travel with the data as it moves.

GOOD PICK?  

GOOD REASON?  

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · October 2, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe