

ARISE: AGENTIC RUNTIME IDENTITY SECURITY ENFORCEMENT FOR AGENTS IN MOTION

LAWRENCE PINGREE

PAUL WEBBER





We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, our timely, opinionated insights help modern security leaders make smarter, faster decisions.

About Us

Software Analyst Cyber Research (SACR) is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over **80,000 readers and followers**, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.

About the Authors

Lawrence Pingree is the Head of Research at SACR, where he leads research on data protection, AI security, identity, cloud security, and agentic security models. With over 16 years of industry analyst experience, including more than a decade at Gartner, he has authored over 300 research notes and provides research driven insights to enterprise security leaders on emerging cyber threats, cloud security, AI security, endpoint defense (EDR), and SD-WAN.

Paul Webber has over 28 years of experience in cybersecurity and enterprise technology, with expertise spanning cybersecurity strategy, threat detection and response, SOC operations, cloud transformation, and cyber resilience across enterprise environments. A former Gartner Senior Director, he advised global CISOs and security vendors on product strategy, go-to-market alignment, and emerging cybersecurity trends. He brings extensive practitioner experience, advising enterprise security leaders on emerging security technologies, operational strategy, ransomware defense, and the evolving cyber threat landscape.

TABLE OF CONTENTS

Executive Summary	3
The ARISE Control Gap	6
The Machine-Speed Problem	7
Market Definition and ARISE Control Model	12
What ARISE Is Not	16
ARISE Control Depth	19
Vendor Landscape and Use-Case Trends	21
What the Vendor Study Revealed	23
Where Vendors Differentiate	26
Six Vendor Patterns in the ARISE Market	29
Market Direction and CISO Guidance	33
Buyer Guidance and Practical Recommendations	34
Conclusion and Research Methodology	38
Vendor Profile and Assessment	40

Executive Summary

AI agents are becoming operational actors inside the enterprise. These agents do not merely generate text. They call tools, request credentials, query data, invoke APIs, modify records, create tickets, generate code, trigger workflows, and they can act as insider threats or collaborative threat actors.

This progression changes the identity security problem. Traditional IAM, IGA, PAM, NHI governance, and gateway controls remain necessary, but they were not designed to decide whether a live agent action should continue at the moment of execution.

In this report, SACR defines ARISE, Agentic Runtime Identity Security Enforcement, as the runtime enforcement layer for agents in motion. It evaluates whether a live agent action, tool call, data access request, credential use, or workflow step still makes sense before the action completes.

The report builds upon SACR's earlier work on AIAP and Runtime Security for AI Agents. Earlier this year, AIAP was the authentication layer for AI agents; and ARISE is the runtime and authorization layer.

The crux of ARISE was built around: should this agent be allowed to take this action, with this tool, credential, data, and delegated authority, for this purpose, right now?

Thanks for reading Software Analyst Cyber Research! Subscribe for free to receive new posts and support my work.

Thesis: The Imperative Now

A new market is developing centered on the substantially growing enterprise problem, where probabilistic AI agents are beginning to act with real credentials, tool access, delegated human authority, and machine-speed execution, while [legacy identity, endpoint, gateway, and governance controls were not designed to supervise agent actions continuously](#).

While traditional static identity controls remain foundational, the rapid, [autonomous execution and delegated authority of AI agents](#) necessitate a shift toward Agentic Runtime Identity Security Enforcement (ARISE), a runtime checkpoint layer designed to evaluate live actions, execute pre-completion interventions, and preserve auditable evidence.

ARISE is needed because enterprise identity has moved through four security eras. Each era solved the prior control problem, but each left a new gap as enterprise actors became more distributed, more automated, and finally more autonomous.



Figure 1: Four Eras of Enterprise Identity Security

Era 1: Human identity and perimeter access.

As referenced, early enterprise IAM assumed that the primary actor was a human employee, contractor, partner, or administrator. The control model centered on authentication, directory membership, role assignment, manager approval, and periodic access review. A person requested access, a manager approved it, the user logged in, systems generated activity, and security reviewed logs or certifications later. This was already imperfect, but it matched the speed of the actor: humans worked at human speed.

Era 2: Cloud, SaaS, and Zero Trust.

As work moved to cloud services, SaaS applications, APIs, and remote access, the perimeter weakened. [Zero Trust shifted the industry toward per-request decisions](#) over users, devices, resources, workflows, and policy enforcement points rather than static network location. NIST's Zero Trust architecture framed access as a policy decision and enforcement problem: the subject must be authentic, the request must be valid, and risk signals can change whether access should be allowed. This was a major advance, but it was still mostly built around access to a resource, not autonomous reasoning over a chain of actions.

Era 3: Non-human identities, secrets, and automation.

DevOps, cloud infrastructure, CI/CD, service accounts, OAuth grants, API keys, workload identities, bots, and machine-to-machine integrations created a much larger non-human identity population. PAM, secrets management, cloud IAM, NHI governance, and identity posture tools emerged to discover, scope, rotate, broker, and review these identities. This reduced credential risk, but most automation still followed a relatively predictable path: a job, script, pipeline, or service executed known steps.

Era 4: Agentic execution.

AI agents change the control problem again. They can reason, plan, use tools, maintain memory, call APIs, browse, retrieve data, write code, modify records, create tickets, and trigger workflows. [OWASP distinguishes agents from traditional LLM applications](#) because agents can take actions and use tools, creating risks beyond prompt injection alone. [NIST's AI Agent Standards Initiative](#) similarly frames agents as systems capable of autonomous actions on behalf of users, requiring secure interoperability and trust.

The practical historical lesson is simple: every major shift in enterprise computing created a new control plane. Networks created perimeter controls. Cloud created identity-centric and posture controls. DevOps created secrets and workload identity controls. SaaS created CASB, SSPM, and API governance. Agentic AI now requires runtime identity enforcement for agents in motion.

Across SACR vendor briefings, vendors converge on four control themes:

- 1 Agent discovery and inventory:** Identifying AI agents, MCP servers, tools, service accounts, owners, delegated users, and connected resources.
- 2 Runtime enforcement:** Blocking, allowing, masking, routing, or escalating agent actions before the action completes.
- 3 Identity and entitlement context:** Connecting agents to humans, non-human identities, OAuth tokens, service accounts, groups, roles, and resource-level permissions.
- 4 Behavioral and intent analysis:** Using graph analytics, prompt/tool-call context, LLM judges, risk scoring, red teaming, or specialized models to infer whether agent behavior matches expected intent.

The market is not yet uniform. Some vendors approach ARISE from identity governance and access management, some from AI security gateways, some from secrets and privileged access, and others from runtime AI application security. This fragmentation is useful for market definition because it shows that ARISE is less a single product category today than a control plane pattern: runtime identity security for agents in motion.

The ARISE Control Gap

The problem is not merely that agents are non-human. The problem is that they are non-human actors with delegated authority, dynamic plans, tool access, memory, and machine-speed execution. Traditional automation usually followed a predefined path. Agentic automation can generate the path at runtime. The agent may choose different tools, retry failed steps, summarize context, infer missing information, combine permissions, or act on untrusted content. A valid credential can still be used for an invalid action.

Blast Radius Clock

When buyers are evaluating agentic security, many are considering the blast radius effects of agentic misbehavior or misalignment, as recent examples of agentics breaching other organizations become a clear and present danger and a new insider threat. The race to get as much telemetry and data to make the best decisions continues to evolve as vendors build greater control around agents.




Figure 2: The Blast Radius Clock

This history explains why existing categories are necessary but incomplete:

- 1 IAM and IGA can say who or what has access, but not always whether a live agent action still matches the business purpose.
- 2 PAM and secrets management can broker credentials, but not always evaluate the intent and sequence of tool use after a credential is issued.
- 3 NHI governance can inventory and govern machine identities, but not every machine identity is an autonomous agent.
- 4 Zero Trust can **continuously evaluate access**, but agentic systems require the same discipline to extend from access requests into action chains, tool calls, delegated authority, and data movement.
- 5 AI gateways can inspect prompts and responses, but agent risk also lives in identity context, tool invocation, API behavior, data sensitivity, and workflow outcomes.

The need for ARISE therefore comes from a timing mismatch. Identity governance historically asked questions before access was granted or after activity was logged. Agents create risk while the action chain is unfolding. A risky tool call, credential use, data movement, SaaS change, or workflow trigger may complete before a human reviewer sees an alert.

The Machine-Speed Problem

ARISE exists because after-the-fact review is too slow for autonomous execution. The issue is that not only do agents act faster than people, but the surrounding threat environment is accelerating.

Vulnerability discovery, exploit development, reconnaissance, credential misuse, data movement, and workflow execution are all compressing from days or weeks into hours, minutes, or seconds.




Figure 3: Human-Speed Governance vs Agentic Execution

This speed of human governance versus agentic execution creates a structural mismatch. Most enterprise governance assumes a sequence of data collection, request, approval, access, activity, logging, alerting, investigation, and remediation with each of these steps taking significant time to execute and subject to delays.

Agentic systems compress the activity portion of that chain. A single agent run can select tools, request or use credentials, query data, invoke APIs, write code, modify records, and trigger follow-on workflows before a human reviewer has enough context to intervene.

The same compression is visible in the broader cyber market:

- AI-assisted cyber operations are moving from advice to execution. [Anthropic reported disrupting a cyber espionage campaign](#) in which a threat actor used Claude Code’s agentic capabilities not merely for guidance, but to execute parts of the attack flow against roughly thirty targets, with successful intrusions in a smaller number of cases. Anthropic characterized the activity as a large-scale cyberattack executed without substantial human intervention.

- Frontier cyber models are changing vulnerability economics. [Anthropic’s Project Glasswing announcement](#) stated that Mythos Preview had already found thousands of high-severity vulnerabilities, including vulnerabilities in every major operating system and web browser. The [UK AI Security Institute’s evaluation of Claude Mythos Preview](#) found that the model could autonomously attack small, weakly defended and vulnerable enterprise systems in a cyber range once network access had been gained. Even when constrained to defensive release channels, these capabilities show that vulnerability discovery and exploit construction can become more automated, scalable and repeatable.
- The disclosure-to-exploitation window is collapsing. [Google Cloud Security reported](#) that the window between vulnerability disclosure and active exploitation collapsed from weeks to days in the second half of 2025, alongside AI-assisted probing and continued data-focused theft. Google concluded that organizations should move toward more automatic defenses.
- Attackers are using AI for vulnerability exploitation and initial access. [Google Threat Intelligence Group has documented](#) adversaries leveraging AI for vulnerability exploitation, augmented operations and initial access. This does not mean every exploit is fully autonomous, but it does mean that AI is accelerating the attack workflow at multiple stages.
- Zero-day and rapid-exploitation pressure is increasing. [CISA’s Known Exploited Vulnerabilities Catalog](#) reflects the need to treat actively exploited vulnerabilities as urgent priorities rather than leave them to ordinary patch cycles. The [CrowdStrike 2026 Global Threat Report](#) separately reported that average eCrime breakout time fell to 29 minutes and that zero-day vulnerabilities exploited before public disclosure increased.

These examples matter for ARISE because an enterprise agent can become part of the same speed problem from the inside. A compromised, misdirected, over-permissioned, or poorly governed agent does not need to “break in” like an external attacker.

It may already have access to tools, SaaS connectors, cloud roles, code repositories, ticketing systems, data stores, and workflow engines. If it is manipulated through [prompt injection, poisoned context, tool abuse, excessive delegation](#), or a vulnerable integration, it can convert legitimate authority into harmful action very quickly.

[Prompt-injection-resistant design](#) therefore requires restricting sensitive-data access and placing additional confirmation or authorization around high-impact actions.

A realistic machine-speed failure pattern looks like this:




Figure 4: Agent Attack chain: A Second-by-Second View

At human speed, this is one suspicious workflow. At agent speed, it is already a completed action chain.

This is why “detect and alert” is not enough for production agent governance. Detection still matters, but the key ARISE question is whether the system can make or influence a decision before the risky action completes. Specifically, ARISE asks:

- Can it stop a tool call when the agent’s intent has drifted?
- Can it narrow or revoke a credential during a session?
- Can it block data movement to an unapproved destination?
- Can it require human review before code, SaaS, cloud, or ticketing changes are committed?
- Can it quarantine an agent that is chaining tools in an unusual sequence?
- Can it preserve evidence linking the prompt, delegated authority, tool, data, policy decision, and outcome?

Lineage of SACR Identity Security Research

The ARISE report is intrinsically and historically related to SACR’s own back catalogue of research from 2025 and 2026 in both the Identity Security and Cybersecurity Governance domains which espoused the concepts of “runtime as the new perimeter” and the convergence of AI and Data security controls and frameworks.

ARISE is also a direct response to the rapid adoption of AI-agentic workflows and the need to retrofit agentic runtime controls and governance.




Figure 5: SACR’s Three-Tier Intelligent AI Security Stack

Suggested Further Reading: SACR Identity Research

- [AIAP \(Emerging Agentic Identity Access Platforms \(AIAP\)\)](#): This report pre-dates the ARISE report and originally charted the emergence of Agentic Identity platforms.

- [Runtime Security for AI Agents: An Identity Governance Perspective](#). A report that set the scene for ARISE, published by SACR, including academic research from Kevin He, Shachar Ram and Lauren Place - it introduced the concept of a move from Deterministic Access to Intent-Aware Governance and provided a framework for managing agents.
- [The Convergence of AI and Data Security: An Industry-Wide Technoscope of Unified Agentic Defense Platforms](#). The initial report by Lawrence Pingree of SACR that introduced the UADP (Unified Agentic Defense Platforms) definition and convergence of AI and Data security disciplines.

Taken together, this prior work establishes the architectural context for ARISE: AIAP addresses agent identity and access, while ARISE extends that control into live runtime action governance.

Emerging Market Definition for ARISE

ARISE, or Agentic Runtime Identity Security Enforcement, is the security discipline and vendor category for governing AI agent actions while they are executing. It combines identity context, delegated authority, tool governance, behavioral analysis, organizational context, data sensitivity, policy decisions, and audit evidence to decide whether an agent action should proceed.

ARISE focuses on the point where agent intent becomes enterprise action:

- Tool calls
- MCP / A2A interactions
- API invocation
- Credential issuance or use
- SaaS actions
- Cloud actions
- Data access, movement, or transformation
- Workflow execution
- [Agent-to-agent handoff](#)
- [Browser-agent or coding-agent behavior](#)
- High-risk autonomous remediation

ARISE is a narrower framework than generic AI security. Prompt and response controls help, but ARISE also needs identity, tool, data, policy, and runtime action context.

Breaking Down the ARISE Control Model

Recent [runtime-governance research](#) similarly argues that enforceable policy must operate outside the model itself, separating deterministic controls from probabilistic model judgment.




Figure 6: Runtime Identity Security Pyramid Model

Layer 1: Deterministic governance

Question: What is this agent allowed to access?

Layer 1 defines baseline boundaries. It includes [agent identity, ownership, scoped credentials, allowed tools, approved APIs](#), RBAC/ABAC policies, JIT access, least privilege, MCP authorization, and human approval requirements for known high-risk actions. MCP deployments can [expose credentials, tokens, and secrets](#) when authentication material is over-scoped, improperly stored, logged, or passed through agent workflows without adequate controls.

This layer is necessary but not sufficient. It tells the enterprise what should be allowed under normal conditions, but it does not fully answer whether a live, context-specific action remains appropriate.

Layer 2: Behavioral and intent analysis

Question: What is the agent doing, and does that behavior make sense?

Layer 2 observes runtime behavior. It includes prompt-to-action reconstruction, tool-call telemetry, credential-use monitoring, intent drift detection, sequence analysis, behavior baselining, anomaly scoring, task-alignment analysis, and session-level evidence.

This is where ARISE begins to diverge from static identity security. A system must understand not only whether the agent has access, but whether the agent's behavior still matches the task, policy, data context, and business purpose.

Layer 3: Dynamic runtime governance

Question: Should this action continue, be narrowed, escalated, or stopped right now?

Layer 3 uses live signals to change the outcome. It includes allow, block, redact, pause, step-up authentication, human-in-the-loop review, scope reduction, credential revocation, quarantine, termination, ticket creation, and SOC escalation. Layer 3 is where visibility turns into control.




Figure 7: ARISE Core Use Cases

1**Runtime enforcement**

Can the product block, allow, redact, pause, approve, revoke, quarantine, terminate, or restrict agent actions before completion?

2**Agent identity and delegation**

Can the product tie agent activity to a human owner, NHI, credential, entitlement, delegated authority, session, and task?

3**Tool / MCP governance**

Can the product inspect and govern MCP servers, tool calls, tool risk, tool authorization, and action-level policy? Is there a facility to manage token cost and recommend/enforce model choices?

4**Behavioral/intent analysis**

Can the product detect intent drift, anomalous tool use, risky sequences, and deviations from expected task behavior?

5**Evidence/audit / response**

Can the product reconstruct prompt, decision, tool call, data touched, action, outcome chains with audit-grade evidence?

6**Agentic brains / organizational context**

Can the product build or use an organization-aware context graph across owners, agents, roles, workflows, policies, data, tools, identities, and prior decisions?

7**AI-powered policy engines and scoring**

Can the product use AI or reasoning systems to author, recommend, simulate, tune, explain, or score policy and action risk?

8**Model and third-party ecosystem support**

Can the product govern multiple foundation models, third-party LLMs, open-source models, self-hosted models, SaaS agents, browser agents, coding agents, MCP ecosystems, and customer-selected deployment patterns?

What ARISE is Not

ARISE is a runtime enforcement layer, not a repackaging of existing categories.

IAM and IGA govern identity, entitlement, access, and review as relatively static disciplines, whereas ARISE governs live action. This distinction places ARISE downstream of provisioning and policy, focused instead on the moment-by-moment validity of what an agent actually does.

PAM addresses privileged session brokering and credential management, functions that remain necessary but insufficient on their own. ARISE extends this by asking whether a specific agent action remains valid given its current context, a question that session-level controls are not designed to answer. Similarly, NHI governance tools discover and manage non-human identities at the inventory and policy level, yet coverage gaps persist: not every NHI functions as a reasoning agent, and not every NHI platform exercises control over runtime behavior.

ISPM and IVIP contribute meaningful visibility through identity posture assessment and relationship graphing, but visibility alone does not constitute enforcement. ARISE closes this gap by operating in-path, where decisions about agent actions are actually made. CNAPP and CDR similarly protect cloud workloads at the runtime layer, but agentic

risk extends beyond workload protection into intent, delegated authority, tool invocation, and multi-step action chains that these platforms are not built to evaluate.

DLP and DSPM assess data sensitivity as a standalone concern, while ARISE incorporates data risk as one input within a broader agent action decision that also weighs identity, context, and intent. AI gateways offer prompt inspection and response filtering at the LLM layer, but this filtering operates without visibility into identity, delegation chains, tool usage, or business-process context, leaving a structural blind spot that ARISE is designed to address.

ARISE should be understood as a sub-discipline within the broader Unified Agentic Defense Platform (UADP) architecture, rather than a standalone category. It occupies a specific functional layer, responsible for runtime action governance, that complements and depends on the other capabilities described above.

Why Identity Context is a Central Theme

Agents act through authority. Identity security must therefore evolve from “who has access?” to “what is the agent doing with delegated authority right now?” ARISE adds the runtime identity question that IAM alone does not answer: a valid credential does not make a live agent action valid.

The enterprise needs to know which agent acted, who owned it, whose authority it used, which credential was issued, which tool was called, what data was touched, what policy applied, what decision was made, and what evidence survived.

That authority may come from:

- A human user
- A service account
- A workload identity
- An API token
- An OAuth grant
- A cloud role
- A SaaS connector
- A privileged session
- A short-lived credential
- Another agent

Why Organizational Context Matters

The next phase of ARISE will depend on organizational context. Static rules and prompt filters cannot understand business purposes on their own. This context layer is what allows an enforcement engine to distinguish a legitimate high-privilege action from an unusual or excessive one.

The implication is straightforward: agent security is becoming a reasoning problem over enterprise context, not only a policy matching problem. Strong ARISE platforms will increasingly need context graphs that map:

- Agents and owners
- Tools and MCP servers
- Identities and delegated authority
- Business workflows
- Data boundaries
- Policies and exceptions
- SaaS and cloud resources
- Prior decisions and approvals
- Normal vs. abnormal action sequences

Why Runtime Visibility, Real-Time Analysis, and Control Matter So Much

Runtime visibility, real-time analysis, and continuous control represent the critical threshold where security transitions from passive observation into active enforcement. In autonomous agentic environments, relying solely on static posture, discovery, or post-event logging leaves an enterprise exposed to high-velocity risk. Because agents operate by calling tools, requesting credentials, invoking APIs, and executing multi-step workflows in seconds, traditional after-the-fact detection creates a timing mismatch where audit events land long after the blast radius has already been established.

Achieving effective agentic security requires instrumenting and correlating three core runtime capabilities:

- **Runtime Visibility (Continuous Observability):** Provides granular telemetry across prompt-to-action chains, tool calls, data access requests, credential utilization, and agent-to-agent interactions. It captures complete execution context to establish audit-grade evidence and session lineage.
- **Real-time Analysis (Behavioral & Intent Analysis):** Dynamically evaluates live telemetry against organizational context, behavioral baselines, and policy boundaries. It detects subtle intent drift, anomalous tool sequences, privilege misuse, and unsafe data movement before execution completes.
- **Real-time Control (Dynamic Enforcement):** Sits directly in the execution path to execute pre-completion interventions such as allowing, blocking, redacting, pausing, stepping up authentication, restricting scope, or terminating an action, ensuring delegated authority remains safe in production.




Figure 8: ARISE Runtime Checkpoint

ARISE Control Depth

ARISE Control Depth (ACD) Framework

To objectively evaluate vendor capabilities and enterprise readiness, SACR establishes the ARISE Agentic Control Depth (ACD) Maturity Model. This framework measures the operational effectiveness of agentic security controls across six distinct levels of capability, ranging from passive observation to fully autonomous, self-healing runtime governance.

The three-layer ARISE model and the ACD framework describe different aspects of the same control problem. The three-layer model describes the functional architecture of ARISE: deterministic governance, behavioral and intent analysis, and dynamic runtime governance. ACD measures the maturity and depth with which those capabilities are implemented, from blind execution at ACD 0 through autonomous adaptive governance at ACD 5.

The Six Levels of ARISE Control Depth:

ACD 0 — Ungoverned / Blind Execution: The organization lacks visibility into agent deployment, tool usage, or delegated permissions. Agent actions occur without logging or policy boundaries.

ACD 1 — Static Discovery & Inventory: The system inventories non-human identities (NHIs) and static agent definitions. It maps static access rights but lacks visibility into real-time tool calls or intent.

ACD 2 — Post-Event Telemetry & Alerting: Agent activity, tool invocations, and data queries are captured in logs or SIEM/SOAR platforms. Alerts are generated after the action completes, providing auditability but no active prevention.

ACD 3 — Perimeter & Deterministic Guardrails: Static gateway rules, prompt-filtering guardrails, and role-based access control (RBAC) enforce hard boundaries. The system blocks known bad inputs or unauthorized APIs but cannot evaluate subtle intent drift or complex action sequences.

ACD 4 — Inline Pre-Completion Interventions (Production Baseline): The control plane sits in-path during tool selection, credential brokering, and API execution. It can evaluate live contextual signals such as data sensitivity, intent drift, and delegated authority to allow, block, redact, pause, step-up, or restrict the action before impact.

ACD 5 — Autonomous Self-Healing Governance: Real-time contextual graphs dynamically adjust policy boundaries, revoke compromised delegation, isolate rogue agents, and continuously recalibrate least-privilege scoping across multi-agent workflows without manual intervention.

For enterprise security teams preparing to deploy autonomous agents in production, the minimum operational threshold should generally be **ACD 4** for any target use case involving sensitive data, credential usage, or critical workflow execution. Anything below ACD 4 leaves the enterprise reliant on post-event remediation after a machine-speed blast radius has already been established.




Figure 9: Agentic Control Depth (ACD) Maturity Model

Important: No single control point covers the entire ARISE problem. Identity, gateways, endpoints, MCP, secrets, graph analytics, and runtime AI security each capture part of the risk surface. The minimum threshold for credible production ARISE should generally be **ACD 4**.

Vendor Evaluation Framework

SACR evaluated and scored ARISE vendors using evidence-weighted criteria. Vendors did not receive full credit for marketing claims, dashboard visibility, or after-the-fact logging alone.

Primary Evaluation Dimensions

Dimension	What Strong Buyer Evidence Looks Like
Runtime Control Depth	Demonstrated pre-completion block, pause, step-up, restrict, revoke, quarantine, or terminate
Identity and Delegation Context	Clear linkage among agent, human owner, NHI, credential, entitlement, delegated authority, session, and task
Tool / MCP Governance	Tool registry, MCP gateway, broker control, per-tool policy, action-level allow/deny, and tool-risk scoring
Behavior and Intent Analysis	Drift detection, anomalous tool-use detection, sequence analysis, and task-alignment scoring
Evidence and Auditability	Prompt/action/tool/identity/policy timeline with SIEM/SOAR export or audit-ready logs
Agentic Brain / Context Graphing	Organization-aware context graph, reasoning layer, memory, and business-process mapping
AI Policy Engine / Scoring	AI-assisted policy generation, simulation, adaptive scoring, explainability, and vendor/control scoring
Model / Third-Party Ecosystem Support	Support for multiple models, third-party LLMs, agent frameworks, MCP/A2A, SaaS agents, coding agents, browser agents, and customer-selected models



ARISE Vendor Landscape and Use-Case Trends

The ARISE market is forming across several existing security categories rather than emerging from one established product segment. The vendors assessed by SACR approach the problem from different architectural starting points, including identity governance, privileged access and secrets management, AI gateways, MCP and tool governance, runtime AI security, and agent lifecycle assurance.

These starting points determine what each platform can observe, which context it can use, where it can intervene, and how directly it can affect an agent action. Identity-oriented platforms typically provide stronger ownership, entitlement, credential, and delegation context. Gateway-oriented platforms sit closer to prompts, model interactions, and tool requests. Privileged-access platforms concentrate on credential brokering and time-bound authority. Runtime AI-security platforms tend to provide richer behavioral telemetry, sequence analysis, and agent-specific threat detection.

No single starting point automatically produces complete ARISE coverage. The relevant buyer question is whether the platform can connect identity, behavioral, organizational, and operational context to a live decision and influence the action before its consequences become irreversible.




Figure 10: ARISE Market Map

The Market Is Converging Around Three Control Layers

Despite their different origins, the assessed vendors are converging around the three layers of the ARISE model:

- 1 **Deterministic governance:** Establishing agent identity, ownership, permitted tools, approved resources, credential scope, access conditions, and predefined policy boundaries.
- 2 **Behavioral and intent analysis:** Determining whether an agent's live behavior remains consistent with its assigned task, delegated authority, expected sequence, and organizational context.
- 3 **Dynamic runtime governance:** Using live signals to allow, block, redact, pause, narrow, revoke, quarantine, escalate, or require approval before an action completes.

The first layer is the most consistently represented across the cohort because it extends familiar IAM, PAM, gateway, and policy practices. The second layer is less uniform. Vendors differ in how deeply they reconstruct agent behavior, identify intent drift, correlate tool sequences, and incorporate identity, data, and business context.

The third layer is the clearest point of differentiation. Many vendors can discover agents, display activity, generate risk scores, or create alerts. Fewer can demonstrate that those signals change the outcome of a live action. Buyers should therefore distinguish between observing an agent, recommending a response, initiating an external workflow, and enforcing a decision directly in the action path.




Figure 11: ARISE 3 Layer Alignment - Cohort Scores (0-5)

What the Vendor Study Revealed

Discovery and inventory are approaching table stakes

Most vendors can identify at least part of the agentic environment, including agents, MCP servers, tools, service accounts, credentials, owners, and connected applications. The depth of this discovery still varies, particularly across local agents, browser agents, SaaS-embedded agents, coding environments, multi-agent workflows, and shadow AI.

Discovery is foundational, but it does not establish ARISE alignment on its own. Inventory tells an organization what exists. ARISE must also help determine whether the agent’s current action is appropriate and what should happen next.

Identity context and runtime behavior are converging

Identity-oriented vendors are adding activity context, tool governance, and runtime enforcement. Runtime-oriented vendors are adding agent ownership, credential lineage, delegated authority, and entitlement context. This convergence is central to the ARISE market.

Identity data establishes who or what is acting and which authority is being used. Runtime telemetry shows what the agent is doing with that authority. Either signal in isolation can be misleading. A credential may be valid while the action is inappropriate, and anomalous behavior may be difficult to assess without knowing the agent's owner, task, permissions, and business purpose.

The strongest ARISE implementations connect these forms of context at decision time.

Enforcement granularity is a meaningful differentiator

A broad allow-or-block decision at the application or tool level is less precise than controlling a particular action against a particular resource. For example, a platform may allow an agent to use GitHub while preventing it from deleting a repository, or permit database access while restricting writes to a sensitive table.

Buyers should test whether policy can operate at the level of:

- Individual tools and tool functions
- Named resources and data stores
- Read, write, delete, export, or administrative actions
- Particular credentials, scopes, and delegated identities
- Specific workflow stages
- Data classifications and destinations
- Agent-to-agent handoffs

This level of granularity determines whether ARISE can reduce risk without unnecessarily stopping legitimate agent workflows.

Human approval remains the dominant pattern for consequential actions

Human-in-the-loop review is widely used for sensitive or ambiguous decisions. Vendors commonly support approval, step-up authentication, escalation, ticket creation, and guided remediation. Fully autonomous response is less consistently available and is often constrained to narrowly defined, high-confidence scenarios.

This is not necessarily a market weakness. Enterprises need enforcement that is explainable, reversible, and calibrated to operational risk. The more important evaluation question is whether the platform can introduce the appropriate intervention before completion, whether that intervention is an automated block, reduced scope, temporary pause, or human decision.

Evidence quality remains uneven

Vendors varied in the strength of the evidence supporting their ARISE claims. Some demonstrated live controls and provided technical documentation. Others relied more heavily on architectural descriptions, screenshots, roadmap statements, or adjacent capabilities.

SACR therefore distinguishes between capabilities that were:

- Demonstrated during the research process
- Supported by technical or product documentation
- Described by the vendor but not directly demonstrated
- Identified as planned or under development

Runtime-enforcement claims should be tested in a live workflow. Buyers should ask the vendor to show the complete sequence from agent request through policy decision, intervention, evidence capture, and final outcome.




Figure 12: ARISE Use Cases Average Scores

Where Vendors Differentiate

Across the cohort, the most meaningful differences appear in six areas.

1. Runtime insertion point

The insertion point determines which actions a platform can see and control. Current approaches include identity providers, credential brokers, gateways, MCP proxies, agent frameworks, SaaS integrations, endpoints, browsers, cloud workloads, and application APIs.

No insertion point provides complete coverage. Buyers must determine whether a vendor's enforcement position matches the agent workflows that create the greatest risk in their environment.

2. Agentic Brains and organizational context

Agentic Brains provide the organizational intelligence needed to evaluate an agent's actions in context. This capability connects owners, agents, identities, credentials, roles, entitlements, tools, data, workflows, policies, prior decisions, and business objectives. It enables the platform to assess not only whether an action is technically permitted, but whether it is appropriate for the agent's assigned purpose and current operating context.

Vendors differ considerably in the depth and operational use of their Agentic Brains. Some provide a visual context graph or searchable inventory. Others use organizational context directly in policy decisions, behavioral analysis, risk scoring, and runtime enforcement.

Buyers should distinguish between platforms that merely map relationships and those that use the resulting intelligence to affect live actions. The key questions are what context the Agentic Brain contains, how frequently it is updated, whether it informs live policy decisions, and whether the platform can explain which organizational factors affected an outcome.

3. Behavioral and intent analysis

Vendors use different combinations of deterministic rules, behavior baselines, sequence analysis, graph analytics, risk scoring, specialized models, and LLM-assisted judgment. The important distinction is whether the analysis can identify a meaningful change in context or intent during execution.

Buyers should ask what evidence informs the decision, how the system handles uncertainty, whether expected behavior is defined manually or learned from activity, and whether the reasoning can be inspected after an intervention.

4. Action-level enforcement

Platforms differ in whether they can alert, recommend, initiate a workflow, or enforce directly. They also differ in the actions available to them, including allow, block, redact, pause, narrow access, revoke a credential, require approval, quarantine an agent, terminate a session, or escalate to security operations.

These options should be evaluated for latency, reliability, reversibility, exception handling, and production impact.

5. Audit and reconstruction

A useful ARISE evidence trail should connect the agent, owner, task, delegated authority, credential, prompt or request, tool call, data accessed, policy decision, intervention, and outcome. This allows security teams to investigate incidents, explain automated decisions, validate policy, and demonstrate control effectiveness.

Buyers should determine whether evidence is retained at the event, session, workflow, and agent levels and whether it is sufficiently complete for security operations, compliance, and post-incident review.

6. Token cost and model selection

As agent adoption grows, organizations must govern not only what agents can do, but also which models they use and how much they consume. Token-cost management and model-selection controls help enterprises manage operating costs, performance, data-handling requirements, approved-model policies, and regional or regulatory restrictions.

Relevant capabilities include:

- Token metering by agent, user, team, application, or workflow
- Usage limits and spending thresholds
- Rate restrictions
- Approved-model policies
- Model routing based on cost, risk, performance, or data sensitivity
- Controls over externally hosted, private, and self-hosted models
- Reporting and attribution of model consumption
- Alerts or enforcement when an agent exceeds approved parameters

Maturity remains uneven. Some vendors provide native metering, budget controls, and model routing. Others integrate with AI gateways or consume this information as additional policy context.

Buyers should determine whether these capabilities are directly enforced by the ARISE platform, supplied through an integration, or identified only as roadmap functionality.




Figure 13: ARISE Layers Average Alignment of Cohort

Six Vendor Patterns in the ARISE Market

The assessed vendors cluster around six overlapping patterns. These groupings describe their primary architectural orientation, not mutually exclusive categories or overall rankings.

Identity-first runtime governance

Identity-first vendors extend identity governance, NHI management, entitlement analysis, and access policy into live agent activity. Their strongest contribution is connecting agent actions to owners, identities, credentials, roles, and delegated authority.

These platforms are most relevant when the primary question is whose authority the agent is using and whether that authority can be constrained or revoked during execution.

Representative vendors: Akeyless, Britive, Silverfort, aizome, Andromeda Security, Oak, Permiso Security, Orchid, and Xage Security.

Privileged access, secrets, and runtime authority

PAM and secrets-oriented vendors concentrate on preventing agents from holding standing credentials. They broker credentials, issue temporary access, enforce least privilege, and connect authority to particular tools, resources, and actions.

These platforms are particularly relevant where agents interact with privileged infrastructure, databases, cloud resources, development environments, or administrative workflows.

Gateway and prompt control

Gateway-oriented vendors inspect model traffic, prompts, responses, data movement, and tool requests. Many are expanding toward MCP governance, identity-aware policy, redaction, model selection, cost control, and action-level enforcement.

Their effectiveness depends on how much relevant traffic passes through the gateway and whether the platform can connect model activity to downstream identity, credential, application, and resource context.

Representative vendors: Cato, Check Point, NeuralTrust, SentinelOne, Lumia Security, Aurascape, TrojAI, Cequence Security, WitnessAI, AIR Security, Akeyless, Britive, and QuilrAI.

MCP and tool governance

MCP and tool-governance vendors operate close to the point where model intent becomes an external action. They register tools, inspect calls, apply authorization, enforce action-level policy, and retain evidence about how agents interact with enterprise systems.

These platforms provide a practical ARISE insertion point, but buyers should test whether they evaluate only predefined permissions or can also respond to behavioral changes during execution.

Representative vendors: Operant AI, PointGuard, Cequence Security, TrojAI, VirtueAI, Zenity, WitnessAI, SentinelOne, Check Point, Silverfort, Britive, aizome, AIR Security, NeuralTrust, and Aembit.

Runtime behavior and agent detection

Runtime-oriented vendors focus on behavior baselines, anomalous sequences, intent drift, agent activity reconstruction, and threat detection. Their strongest contribution is determining whether an agent's activity remains consistent with its expected task.

The central evaluation question is whether detection changes the outcome in time or produces an alert after the action has already completed.

Representative vendors: Noma Security, HiddenLayer, Ent.ai, Permiso Security, Straiker, Pillar, Zenity, WitnessAI, Operant AI, Lasso Security, aizome, and ClearVector.

Agent lifecycle assurance and adjacent controls

Some vendors begin with agent discovery, posture management, red teaming, MCP scanning, prompt-injection testing, development guardrails, or predeployment validation. These controls can reduce runtime risk, but they should not be treated as runtime enforcement unless the product can govern live actions.

Representative vendors: Pillar, Lasso Security, VirtueAI, Noma Security, HiddenLayer, Straiker, and NeuralTrust.

A smaller group of vendors combines several of these patterns. Their platforms bring together gateway controls, MCP and tool governance, runtime behavior analysis, agent identity and delegation context, Agentic Brains, audit evidence, and response.

Breadth can improve coverage, but buyers should still verify the depth and maturity of each claimed control rather than assuming that a broad platform provides complete ARISE functionality.

Representative vendors with broad ARISE alignment: AIR Security, Akeyless, AuthMind, NeuralTrust, Oak, Okta, Silverfort, and Zenity.

Common Capability Gaps




Figure 14: Current Market: Control Depth Maturity

The study identified several areas where market maturity remains uneven:

- Detecting behavioral drift after a tool or MCP server has already been approved
- Connecting agent actions consistently to human and machine owners
- Applying policy at the individual action and resource level
- Combining identity, behavioral, data, and business context during a live decision
- Using Agentic Brains directly in policy and enforcement rather than only for visualization
- Intervening before completion across distributed and multi-agent workflows
- Safely automating remediation without creating unacceptable operational risk
- Attributing token consumption and model usage to particular agents, owners, and workflows
- Enforcing model-selection and spending policies across heterogeneous AI environments
- Preserving an audit trail that explains both the agent action and the enforcement decision
- Distinguishing generally available functionality from demonstrations, limited previews, and roadmap plans

These gaps reinforce the need to evaluate ARISE through demonstrated outcomes rather than category labels or feature counts.




Figure 15: ARISE Core Use Cases

SACR Analyst Take

ARISE is currently a control-plane pattern in the process of becoming a more recognizable market category. Buyers should not shortlist vendors based only on whether they use agent-security language. Nor should they assume that identity, gateway, secrets, or runtime origins determine which platform is most mature.

The best fit depends on where the organization’s agent risk is concentrated and which intervention points are technically and operationally available. A company concerned primarily with privileged agent access may favor a different architecture from one focused on SaaS agents, coding agents, MCP ecosystems, adversarial manipulation, or model-cost governance.

Across these scenarios, the strongest ARISE implementations combine identity context, runtime control, and an Agentic Brain capable of interpreting an action within the organization’s operating context. Token-cost and model-selection controls add a further governance layer by determining whether agents are using approved models and resources within defined operational boundaries.

Buyers should require vendors to demonstrate who or what is acting, whose authority is being used, what action is being attempted, why the action is considered safe or unsafe, which model and resources are being used, and how the product can change the outcome before completion.

Market Direction

The current market patterns are likely to converge. Identity platforms are adding runtime telemetry and enforcement. Gateway and runtime-security platforms are adding identity and delegation context. PAM and secrets platforms are extending credential controls into agent and tool workflows. MCP-security vendors are becoming action-level policy enforcement points. Agentic Brains are increasingly connecting these controls through shared organizational context, while token-cost and model-selection capabilities are extending ARISE into the operational governance of enterprise AI consumption.

This convergence may occur through product development, partnerships, integrations, or acquisition. It does not require every ARISE vendor to become an end-to-end agentic-security platform. The nearer-term measure of market leadership is whether a vendor can connect relevant identity, behavioral, organizational, and operational context to a safe, timely, and auditable runtime decision.

Future View: ARISE Market Destination: Unified Agentic Defense Platforms (UADP)

Although we believe ARISE is an early emerging and rapidly growing category, SACR views ARISE as an emerging sub-category of Unified Agentic Defense Platforms from a market storyline perspective, corresponding to the rise of short term exponential adoption of AI Agents in the enterprise and as runtime monitoring requirements and implementation urgency has evolved increasing demand from enterprises to manage the emerging risks of Agentic Identities and the users that command AI agents.

Over the longer-term horizon, SACR views UADP as the ultimate destination of end-to-end agentic security. ARISE provides the immediate-term necessary runtime identity enforcement for agents in motion; we believe ARISE providers will evolve towards UADP capabilities alongside comprehensive data security (DSPM/DLP) features, AI governance lifecycle management, shadow AI discovery, and end-to-end threat detection across the entire agentic attack surface.

See [\(The Convergence of AI and Data Security: An Industry-Wide Technoscope of Unified Agentic Defense Platforms\)](#)




Figure 16: Unified Agentic Defense Platform

CISO Guidance

CISOs should use ARISE as a procurement, architectural, and governance lens when deploying and scaling production agentic AI. Beyond evaluating basic functional capabilities, security leaders must establish rigorous criteria across access boundaries, operational visibility, and real-time intervention points to ensure agent autonomy remains secure and [aligned with enterprise risk tolerance](#).

Key Questions Buyers Should Ask

1. Which agents exist, and who owns them?
2. Whose authority does each agent use?
3. What tools, APIs, MCP servers, SaaS apps, and data sources can the agent touch?
4. Can access be scoped, short-lived, and revoked?
5. Can the product evaluate whether a live action matches the declared task?
6. Can the product detect intent drift or risky tool chains?
7. Can the product block, pause, restrict, revoke, terminate, or escalate before completion?
8. Can the product reconstruct the action chain for audit, incident response, and board reporting?
9. Can policies be simulated and tuned before enforcement?
10. Does the platform support the organization's mix of first-party, third-party, open-source, SaaS, browser, coding, and self-hosted agent/model environments?

CISOs should evaluate each vendor against the following criteria:

- 1. Agent identity model:** Does the product distinguish autonomous agents, delegated agents, NHIs, service accounts, and human users? Can it map ownership and responsibility?
- 2. Runtime enforcement point:** Where does enforcement occur (identity provider, MCP gateway, AI gateway, endpoint, API connector, SaaS integration, workload identity broker)? Can the product block before action completion?
- 3. Granularity of control:** Does policy apply at model, app, tool, API, resource, field, action, or data-classification level? Can it differentiate read/write/delete/export actions?
- 4. Behavior and intent analysis:** Does the system correlate prompts, tool calls, outputs, sessions, and multi-turn behavior? Does it use deterministic rules, ML, LLM judges, or graph analytics?
- 5. Human-in-the-loop and JIT:** Can sensitive actions trigger approvals, step-up authentication, or just-in-time access? Are approvals auditable and integrated into collaboration tools?
- 6. Credential and delegated authority handling:** Are agents prevented from directly holding secrets? Does the system handle OAuth delegation, service accounts, and workload identities?
- 7. MCP and tool ecosystem coverage:** Can it discover, score, approve, block, and monitor MCP servers and tools? Does it detect tool drift or only pre-launch risk?
- 8. Audit, reconstruction, and evidence:** Can the buyer reconstruct what happened, who/what initiated it, what credentials or tokens were used, what data was touched, and why policy allowed or denied it?
- 9. Operational integration:** Does it integrate with SIEM, SOAR, ITSM, Slack/Teams, IGA, IDP, XDR, MDM, cloud, SaaS, and developer platforms?
- 10. Maturity and roadmap:** Which capabilities are GA, beta, preview, roadmap, or custom/professional services?

Buyer Decision Rules

Avoid treating visibility-only tools as runtime enforcement. For production agent governance, prioritize vendors that can prove the following are viable options for security enforcement for AI agents:

- An in-path enforcement point
- Clear identity and delegation mapping
- Tool/MCP-level control
- Behavioral or intent signals
- Pre-completion intervention
- Evidence-grade audit
- Operational tuning and false-positive management
- Coverage across the buyer's real agent surfaces

Practical Recommendations

Actionable steps for leaders

- Examine ARISE use cases and select and prioritize those that are most important to the individual organization/industry vertical. Make these the key criteria for the ongoing evaluation and select candidate solutions based on their ability to fulfil these specific needs.
- Use the 6 listed patterns that define distinct groups of vendors to understand their specific strategic alignment and raison d'être – as this may be a more significant factor than the otherwise small differences in actual capability between vendors in this cohort.
- Since the solutions are mostly capable of meeting at least the first 2 layers of ARISE, assess whether the layer 3 dynamic runtime governance is a requirement now or in future and evaluate specific capabilities that provide the necessary real-time governance and enforcement that is needed now and projected for the future.
- Determine whether the provided integrations and workflow connectors are able to provide strong 2 way actions between key business applications and operations toolsets out of the box. For any not listed or supported by a given vendor, understand the possibilities to build these integrations and evaluate this as part of a POC.
- Time to value and ROI are as important to consider as the technical capabilities of ARISE vendors, and the product selection process should also take care to select solutions that can be readily managed by the target organization without the need for protracted training, additional resources or 3rd party managed services.

- Where vendors are unable to accurately forecast the licence costs, management overhead and especially any model/token costs, make sure that the procurement process and eventual contractual stipulations protect against unforeseen or punitive charging models that penalise heavy use of the system.
- Work with candidate vendors to prepare a projected deployment plan, integration and migration schedule and identify key success criteria and time-bound project objectives as part of a 30/60/90-day view.
- Identify metrics to track and the metrics that will be used to assess this – particularly important if the underlying objectives and success criteria are part of a wider AI transformation program or any other process that has pre-assumed goals relating to consolidation aims, resource reduction and efficiency gains.
- Use SACR research and associated control frameworks and measurement criteria to inform the project and business of the compliance and risk management aspects of any underlying threats or risk exposure that the company is tracking and to verify whether the specific controls and governance facilities of ARISE vendors can meet those objectives.

Business Impact

ARISE should be framed as an enablement category, not only a defensive control. Enterprises want the productivity benefits of AI agents, but they need a governance model that allows safe autonomy. The business case is straightforward: ARISE lets the business use agents without asking security to accept blind autonomy.

ARISE supports five business outcomes:

- 1 Move agents into production safely**
Security teams can approve agent use with enforceable controls rather than blanket prohibition.
- 2 Preserve productivity while reducing risk**
Low-risk actions proceed automatically; high-risk actions receive friction, escalation, or denial.
- 3 Reduce blast radius**
Agents can be constrained before risky chains propagate across tools, data, APIs, and SaaS apps.
- 4 Improve accountability**
Every action can be tied to an agent, owner, delegated authority, policy, tool, data context, and outcome.
- 5 Support compliance and board reporting**
Agentic workflows become auditable, explainable, and defensible.

Conclusion

AI agents will force identity security to operate in real time. The enterprise cannot rely only on provisioning, entitlement review, prompt inspection, or after-the-fact audit when agents can act through delegated authority at machine speed. The security control must move closer to the action. ARISE gives security leaders a practical framework for that shift.

It defines the runtime enforcement layer that evaluates agent identity, delegated authority, tool risk, data context, behavior, policy, and organizational context before the action completes. It also gives vendors a clearer standard: do not claim ARISE based on visibility alone. Prove where the product can decide, intervene, and generate evidence. As agents move into production, the strongest platforms will govern more than access. They will decide whether the agent's live action still makes sense.

SACR final takeaway: The future of agent security is runtime action governance, not access approval. Traditional access control models evaluate static permissions and entitlement grants before an interaction begins, but autonomous AI agents operate dynamically across multi-step action chains, tool invocations, and delegated authority.

As agents execute tasks at machine speed, static authorization and post-event logging fail to prevent high-risk operations before impact occurs. Effective security requires continuous, in-path evaluation of every live tool call, API request, credential usage, and data movement in real time to determine whether an action should be allowed, constrained, or blocked before completion.

Research Methodology

This research by Software Analyst Cyber Research (SACR) is based on proprietary analyst interpretation of vendor briefings, market research, product materials, public technical sources, practitioner discussions, and SACR's emerging ARISE evaluation framework. Findings are directional and should be validated against product demonstrations, GA capability status, customer evidence, and implementation details.

The agentic AI security market is changing rapidly. Vendor capabilities, terminology, and category boundaries may evolve after publication.

The background is a solid purple color with a series of thin, white, curved lines that flow from the bottom left towards the top right, creating a sense of movement and depth.

britive

Vendor Overview

Britive is aligned with the privileged access, cloud entitlement and identity security market rather than as an ARISE-only point product. The relevant ARISE capabilities are delivered through Britive ARC™ (Agentic Runtime Control), an extension of its existing platform model: human users, service or non-human identities, and agentic identities are treated as first-class identity types with common access, ownership and policy controls.

The current organizational fit is strongest where the client already uses just-in-time access, zero-standing privilege, credential checkout, cloud access governance and service-account governance, and then wants to apply the same operating model to agents and MCP tools. This is, however, not framed as a separate agent security console. The typical customer profile is a cloud-heavy enterprise, or its security engineering team, platform engineering group, or any regulated technology organization that wants identity and privilege controls to remain close to the point of execution.

Market messaging is centered on consolidation: one access model across human, NHI and AI identity rather than another stand-alone control plane.

Product Architecture and Hosting Options

Britive provides a hosted MCP server for Britive platform actions, such as profile checkout, and a separate MCP Gateway intended to be deployed inside the customer environment. The gateway aggregates backend MCP servers, publishes tools back to the platform where needed, and enforces per-user tool visibility and call-time authorization. OAuth-based MCP servers can be discovered through standard metadata, while non-OAuth systems such as MySQL can use ephemeral credentials checked out from Britive and injected downstream as headers. The downstream MCP server does not need to hold standing credentials.

The architecture is a hybrid one. The Britive SaaS platform stores OAuth tokens and evaluates policy. The gateway runs in the customer network and reaches back to the platform for token retrieval, policy decisions and audit sync. Bridge, a separate optional deployable component, provides wire-level session recording and protocol interception for SSH, MySQL, RDP, HTTP proxy and other database protocols. Britive also exchanges Shared Signals Framework events in both directions: inbound CAEP and RISC events from identity providers and security tools can revoke an active checkout mid-task.

Identity and credential issuance are core to the design. Britive supports workload federation including AWS STS, OIDC sources such as GitHub Actions and Bitbucket Pipelines, SPIFFE and static-token fallback. Agentic identities carry ownership, business-use and on-behalf-of controls.

Open integrations include MCP and REST APIs, SDK/Terraform-style automation, OAuth-backed MCP servers and cloud applications such as Figma, GitHub, HubSpot, Snowflake and Atlassian-style tool namespaces.

Market Context, Competition and Positioning

Britive is best positioned as a privilege and identity control vendor extending into agentic runtime enforcement. The technical focus is not prompt security alone, nor broad SaaS posture management.

Instead, the emphasis is on scoped access, credential issuance, policy enforcement and runtime mediation across users, service identities and agents.

This puts Britive on shortlists where buyers already care about cloud privileged access, ephemeral credentials, service identity governance and MCP gateway control.

It is less likely to be the first choice for organizations primarily looking for DSPM-style data classification, model-cost governance, LLM red-team management or a broader agent behavior analytics layer.

The obvious competitive set includes privileged access management, cloud infrastructure entitlement management and newer MCP gateway or agent security products.

Britive states it does not require endpoint agents for the core discovery facility, apart from selective proxies such as Gateway and Bridge. That suits engineering environments where forcing all traffic through a new security path would be impractical.

Alignment to the ARISE Core Layers

ARISE layer	Alignment	Evidence and gaps
Layer one: deterministic governance	Strong alignment	Human, service and agentic identities are first-class. Ownership, orphan reporting, scoped profiles, allowed tools, MCP authorization, JIT access, least privilege, default-deny policy and human-in-the-loop controls are supported through the common policy model.
Layer two: behavioral and intent analysis	Partial alignment	The gateway injects LLM session ID, user prompt and LLM-stated intent into every tool call. Audit logs can reconstruct tool calls, arguments, outcomes and sessions. Native intent-drift detection, behavior baselining, anomaly scoring and task-alignment analytics are not yet strongly evidenced.
Layer three: dynamic runtime governance	Strong alignment for access and session controls	Britive can allow, deny, require human approval, trigger ephemeral privilege elevation, terminate live sessions and block Bridge-mediated sessions without active checkout. Response inspection can operate in audit or enforcement mode. Quarantine and partial redaction were not shown.

Britive aligns most in deterministic governance and agent identity and delegation. The Layer 2 position is more nuanced. There is better telemetry than a basic access gateway would provide, particularly because prompt and stated intent are captured alongside the tool call.

Moat and Differentiation

Britive's clearest moat is the just-in-time, zero-standing privilege model applied across identity types. This is where the ARISE fit is also strongest.

Credential checkout, ephemeral access, OAuth token storage on the platform rather than in the LLM client, and downstream injection for non-OAuth systems give the control model substance.

A second differentiator is the unified identity model. Human users, non-human identities and agentic identities use the same platform policy machinery.

That reduces tool sprawl and is a sensible enterprise architecture point, especially for organizations already dealing with NHI and cloud access governance.

The MCP Gateway design is another important pillar. Per-user tool filtering, server-name prefixing, OAuth discovery and call-time authorization give Britive powerful action-level governance.

Bridge adds a separate but relevant control surface. Wire-level session recording and blocking for MySQL, SSH, RDP and HTTP proxy are not strictly MCP controls, yet they matter when an agent connects to a database or server through mediated access.

The low-friction deployment story is noteworthy. The vendor claims engineers can keep using familiar tools and that the security layer is mostly invisible.

While not a distinct feature by itself, it is often the difference between a control being used and bypassed. DevOps teams in particular now favor using their chosen toolset direct-to-harness.

Use Case Alignment

Use case	Alignment	Notes
Agent identity and delegation	Very strong	Agentic identities have ownership, business-use fields, on-behalf-of delegation, profile-level impersonation control and workload/OIDC/SPIFFE/AWS STS authentication.
Runtime enforcement	Strong	Default deny, policy pass-through, JIT/profile checkout, human-in-the-loop approval, Bridge blocking, SQL command blocking and response inspection in audit or enforce mode. Not yet a universal agent kill-switch or quarantine layer.
Tool / MCP governance	Strong	Hosted Britive MCP server, customer-deployed gateway, backend MCP server management, OAuth discovery, tool publishing, dynamic per-user tool filtering, call-time authorization and audit logging.
Evidence / audit / response	Strong	Tool logs include session ID, user prompt, LLM-stated intent, tool arguments, outcome and user. Logs sync to the SaaS platform and can be sent to SIEM via webhooks.
Credential issuance or use	Very strong	Ephemeral credential/profile checkout is central. OAuth tokens stay on the Britive platform, and temporary credentials can be passed to downstream MCP servers without creating standing credentials there.

The best-aligned areas are cloud actions and API invocation. Cloud access is a natural strength because Britive's common policy model already covers cloud providers and privileged profile checkout. API invocation is supported through REST and automation interfaces.



business



personal



softwareanalyst.io



Trusted research. Sharp insights. Real conversation.

CISO

VENDOR

SECURITY
TEAMS

INVESTORS