CLAWOLF AIS
GOVERNED FRONTIER GUARD

CNDA PORTFOLIO / FLAGSHIP PRODUCT

AI Trust. Security. Governance.

CLAWOLF AIS

Protect AI. Constrain Harmful Agent Actions.
Preserve Trust. Govern The Outcome.

PLANNED PRODUCT CAPABILITIES
TARGET RELEASE: 30 NOVEMBER 2026



AIS Core and AIS Signature

Two standalone editions. Both include AI Trust, AI Governance and CLAWOLF Immunity.

Cloud annual list	AIS Core $229	AIS Signature $299
Per billable asset / year	All eight Core feature groups	All Core capabilities plus the six Signature mechanisms
Protection scope	Discovery, identity, prompts, memory, tools, models, data, red teaming and governed response	Core protection plus causal integrity, joint effects, outcome assurance, decontamination, safe continuity and assurance invalidation
Shared foundations	AI lifecycle governance, accountable authority, evidence/reporting and governed guardrail engineering	The same complete Core foundations, with advanced Signature mechanisms

One AIS Product, Two Capability Packages

Signature includes Core for the same licensed scope. Standalone AIS needs no AS-OS runtime or licence. Paid annual ACD is separate, including for cloud. Deployment and final scope follow the authorized quotation.

AIS governs AI systems and agents. Optional AS-OS integration adds enterprise-wide security decisions and coordinated response.

AIS Core Features

Full-stack protection for AI models, agents, memory, tools and connected enterprise resources.

01 AI Discovery & Security Posture

Discover AI models, apps, agents, tools and dependencies. Map accountable owners, AI bills of materials, permissions, trust boundaries and attack paths.

02 Agent Identity & Delegated Authority

Trace human-to-agent and agent-to-agent authority. Expose excessive privileges, token misuse and escalation. Bind actions to current, accountable permissions.

03 Prompt Injection & Multimodal Protection

Detect direct and indirect injection, jailbreaks and task hijacking across supported text, documents, images, audio, code and tool responses.

04 RAG & Memory Protection

Authorize retrieval sources, preserve integrity and lineage, and contain poisoned memory, manipulated instructions and dependent context.

05 MCP, Tools, Skills & AI Supply Chain

Inspect tool definitions, skills, permissions, arguments and changed behavior. Check model/component provenance, integrity and supply-chain risk.

06 Models, Sensitive Data & Output Protection

Protect model and package integrity, personal data, secrets and intellectual property. Govern outputs, data access and exfiltration destinations.

07 Continuous Red Teaming & Resource Defense

Exercise bounded multilingual, multi-turn, multimodal and multi-agent attacks. Constrain recursive loops, runaway delegation and resource abuse.

08 Runtime Detection & Governed Containment

Detect purpose drift, task hijacking and cross-agent abuse. Block calls, revoke access, quarantine context and verify the effects of containment.

AIS Signature Features

Six mechanisms protect the mission, its evidence and the real-world consequences of agent actions.

01 Causal Evidence Integrity

Trace origins across copies, summaries and agents. Repeated derivatives of one source do not become independent confirmation.

Five agents quoting one document still equal one source.

02 Semantic Effect Assurance

Compare authorized target and scope with the separately observed effect of an action. A success claim alone does not establish the result.

An allowed read must not silently export customer data.

03 Joint-Outcome Defense

Evaluate cumulative effects across agents and concurrent branches. Splitting a harmful objective does not reset its limits.

Two allowed reads must not combine into prohibited disclosure.

04 Causal Decontamination

Trace and quarantine poisoned memory, summaries and dependent branches while preserving trustworthy context.

Remove the contaminated path while preserving unaffected work.

05 Mission-Preserving Defense

Continue only work that remains verifiably safe within its existing authority. Stop or escalate when no safe continuation exists.

Contain the risky branch while a safe workflow continues.

06 Continuous Assurance Invalidation

Bind assurance to model, tool, permission and context versions. Recheck only evidence affected by a relevant change.

Yesterday's approval does not certify today's changed agent.

AI Governance Across the Lifecycle

Accountable ownership and policy become ongoing control of AI use, change and retirement.

Governance responsibility	AIS capability	Business value
Ownership and Purpose	Register use cases, accountable owners, AI assets, intended purpose and risk boundaries.	Know who is responsible and which AI activity the organization has authorized.
Risk, Approval and Exceptions	Assess risk, record approvals and time-bound exceptions, and require review after relevant changes.	Keep authorization current as agents, data and permissions evolve.
Trustworthiness Evaluation	Evaluate applicable technical trust requirements and preserve results, limits and supporting evidence.	Make trust decisions on the relevant system version and actual evidence.
Policy-to-Runtime Control	Connect approved policy to actual enforcement and monitoring, with visible coverage and gaps.	Distinguish a policy on paper from a control that is operating.
Change and Retirement	Invalidate stale assurance, re-evaluate changes and govern evidence retention and decommissioning.	Avoid carrying old approvals into a changed or retired AI use case.

Governance, trust and security work together. Framework mapping supports governance design and does not itself constitute certification.

CLAWOLF Immunity

Governed guardrail engineering converts observed weaknesses into qualified, recoverable protection.

01 Discover and Reproduce

Identify a genuine weakness and reproduce it within an authorized, bounded scope.

Understand the attack before changing protection.

02 Generate a Protection Candidate

Derive a constrained guardrail candidate from the finding. Generation never grants permission to activate it.

Automate protection engineering while retaining control.

03 Qualify Against Attacks and Valid Work

Test the original and held-out attacks, together with legitimate tasks and the intended operating scope.

A protection must preserve the work the business needs.

04 Approve and Promote

Apply the assigned policy, approval and delegated authority to the exact protection version and target.

Activate only the protection the organization authorized.

05 Verify and Recover

Observe the actual protective effect independently. Restore and verify the last known good version when necessary.

Prove what changed and retain a safe recovery path.

06 Refresh Trust and Evidence

Re-evaluate affected assurance and retain the finding, qualification, authorization and outcome lineage.

Protection improves without treating old evidence as current.

CLAWOLF
ART OF GOVERNED AUTONOMY

AIS / COMPETITIVE LANDSCAPE

AI-Agent Security: Current Market State

Market capabilities and their functional counterparts across CLAWOLF AIS, native AS-OS and the CNDA portfolio.

Capability	Palo Alto Prisma AIRS	Zenity	Noma	HiddenLayer	F5	CLAWOLF Combined Coverage
AI Discovery & Posture	AI Inventory & AI-SPM	Agent Inventory & Posture	Agents, Apps & MCP Inventory	AI Asset Discovery	AI App, Agent & API Discovery	**AI Estate + Attack-Path Context**
Agent Identity & Context	Identity-Aware Agent Policies	Intent, Memory & Handoffs	Identity, Sessions & Permissions	Agent / Harness Runtime Context	Runtime Access & Traffic Context	**Identity + Purpose + Observed Effects**
Tools & Runtime Enforcement	AI Runtime Inspection	Workflow-Aware Intervention	Inline MCP & Tool Controls	Tool, Code & Action Inspection	API / AI Gateway Guardrails	**Effect-Aware Governed Actions**
Adversarial & Supply-Chain Testing	Red Teaming + Model Scanning	Agent Exposure Validation	Red Teaming + AI Component Risk	Model Attacks + Supply-Chain Risk	AI Red Team + AI Remediate*	**Red Teaming + Guardrail Engineering**
Evidence & Decision Context	Policy / Platform Telemetry	Agent Intent & Behavior Traces	Session & Policy Context	AI Attack & Runtime Events	Explainable Risks & Security Logs	**Causal Evidence + Decision Lineage**
Response & Recovery	Runtime Blocking + Case Response	Block / Steer Risky Agent Actions	Block, Mask or Seek Approval	Runtime Blocking + Integrations	Runtime Guardrails + Risk Mitigation	**Mission-Safe Defense + Causal Cleanup**
AI Governance & Trust Lifecycle	AI Governance, Identity & Runtime Policies	Agent Governance & Runtime Boundaries	Agent Access Policies & Runtime Governance	AI Security Policies & Runtime Enforcement	AI Policies, Audit & Access Controls	**AI Lifecycle + Change-Bound Trust**
Adaptive Protection / Immunity	Red Team Findings & Mitigation Guidance	Exposure Validation & Runtime Rules	Red Team Findings & Runtime Hardening	AI Red Teaming & Runtime Protection	AI Remediate* Guardrail Generation	**Immunity + Qualified Protection & Recovery**
Operating Architecture	AI Security + SOC Orchestration	Agent Policies & Runtime Control	Central Policy & Runtime Control	AI Security + SOC Integration	AI Policy + Runtime Enforcement	**Native Five-Layer CLAWOLF AS-OS**
Cross-Domain Contribution	PANW Security Ecosystem	Agent / Business Application Context	AI, Data & Access Context	Model / Agent Dependencies	AI / API Security Context	**Native CLAWOLF CNDA Portfolio**

CLAWOLF AIS | **CNDA Portfolio Contribution** | **Native AS-OS Contribution**

Colors identify CLAWOLF functional counterparts. Combined coverage adds the advantages in the right column. Functional crosswalk, not whole-product parity.

* F5 AI Remediate: early access. AIS planned capabilities / target release: 30 November 2026. CNDA contributions follow the product roadmap.

Why Choose CLAWOLF AIS?

Core provides standalone AI protection and governance. Signature adds the six advanced mechanisms.

Customer need	AIS capability	Customer outcome
Protect the AI Estate	Core discovery, identity, prompt, memory, tool, model, data and runtime controls.	See the protected AI estate and constrain harmful agent activity.
Govern AI Use	Core ownership, purpose, risk, approval, exception and policy-to-runtime workflows.	Keep AI decisions accountable and authorization current.
Improve Guardrails Safely	CLAWOLF Immunity in both editions qualifies protection against attacks and legitimate work.	Improve protection with approval, observed effects and recovery.
Understand Combined Consequences	Signature adds causal integrity, semantic effect assurance and joint-outcome analysis.	Expose false corroboration and harmful combinations across agents.
Recover the Mission	Signature adds causal decontamination, mission-preserving defense and assurance invalidation.	Remove affected context, preserve safe work and refresh affected trust.
Extend to Enterprise Operations	Optional AS-OS integration links AI-domain evidence to assigned enterprise decisions and response.	Investigate and coordinate beyond the AI domain while AIS retains its own governance.

Both editions remain independently operable. AS-OS expands enterprise coordination. Separately adopted CNDA products add their qualified domain context.

Native Integration: Customer Gains

AIS governs AI models, agents, tools and data. AS-OS governs the assigned enterprise response.

Native integration value	Combined mechanism	Customer benefit
1 Attributable Incident Context	Connect AI findings, identities and affected assets to the enterprise investigation.	Investigate the AI incident with business and security context while retaining its evidence origin.
2 Authority Across Handoffs	Bind supported actions to the tenant, target, current permissions and mission/rate limits.	Keep approvals and resource limits intact as work crosses product boundaries.
3 Scoped, Coordinated Response	Combine AI-local containment with separately authorized identity, infrastructure or data actions.	Contain the supported attack path while limiting the scope of the enterprise response.
4 Observed Outcomes and Recovery	Link the decision and action receipt to independent effect observation and recovery evidence.	Distinguish an accepted command from a protective change that actually occurred.
5 Change-Aware Trust	Carry relevant versions and causal evidence into requalification and affected-context recovery.	Avoid relying on stale assurance and preserve unaffected, authorized work.

Native integration design scope. Supported actions, deployment profiles and release qualification determine available coverage. Additional CNDA products follow their own roadmap.

CLAWOLF AIS + CLAWOLF CNDA

Six additional, separately adoptable products connect AI activity to enterprise risk.

Roadmap target	CLAWOLF product	Combined customer value
December 2026	Identity Defense Intelligence	Connect agents to accountable owners, delegated authority and reachable resources. Constrain stolen or excessive access.
December 2026	Data & Database Intelligence	Connect agent actions to sensitive records and database reachability. Expose authorized-query exfiltration paths.
January 2027	Behavioural Intelligence	Combine agent, user and workload baselines with task purpose. Expose stealthy misuse and harmful sequences.
January 2027	Insider Defense Intelligence	Relate employee behavior, privileges and data access to delegated agents. Expose insider and compromised-account abuse.
February 2027	Supply Chain Exposure	Relate suppliers, SaaS, MCP services and tokens to agent dependencies. Trace compromised third-party attack paths.
February 2027	Cloud Exposure & Attack In Motion Intelligence	Combine agent permissions, cloud exposure and active evidence to prioritize reachable business risk.

Portfolio roadmap targets. AIS retains AI-domain authority. AS-OS governs the assigned enterprise response across adopted domains.

Market & Taxonomy Alignment

Governance, trust, security and agent-risk frameworks mapped to the CLAWOLF product architecture.

Authority / framework	Verified market concept	CLAWOLF architecture mapping
Gartner	AI TRiSM: AI governance, AI runtime inspection and enforcement, and information governance.	AI estate and lifecycle governance, runtime controls, data protection and evidence.
IDC	AI agent security and governance, identity, platform context and sovereignty.	Accountable agent authority, customer control and institutional risk context.
Forrester	People define policy and risk boundaries. Controls continuously apply and monitor them as agents act on data.	External authority, current policy, observed activity and evidence-based review.
NIST AI RMF	GOVERN, MAP, MEASURE and MANAGE throughout the AI lifecycle.	Use-case ownership, risk evaluation, protection, change and evidence lifecycle.
MITRE ATLAS / OWASP Agentic Top 10	Agent goals, tools, identity, memory, supply chain, cascading effects and AI adversarial techniques.	Prompt/tool/identity controls, causal memory integrity, joint effects and adversarial qualification.
CSA MAESTRO	Threat modeling across interacting agent-system layers and trust boundaries.	Model, data, tool and operating boundaries, assurance and ecosystem authority.

Security and governance design mapping by CLAWOLF. This is not certification, endorsement or a claim of full framework conformity.

Gartner IDC Forrester NIST MITRE OWASP CSA