CLAWOLF
ART OF GOVERNED AUTONOMY

ENTERPRISE PRODUCT OVERVIEW

Autonomous Security Operations And AI Protection

CLAWOLF



COMPANY OVERVIEW

Who Is CLAWOLF?

A Security Company Built Around Governed Autonomy

CLAWOLF Inc. is a Delaware C-Corp incorporated in April 2026, built on security research that began in September 2024.

We develop autonomous security operations for cloud, on-premises and hybrid enterprises. AS-OS connects security decisions to authorized response, with human oversight and evidence of the outcome.

NIST GUIDANCE

26 FEB 2024
Cybersecurity Framework 2.0
Final framework. Introduces GOVERN.

26 JUL 2024
Generative AI Profile
NIST AI 600-1. Final publication.

16 DEC 2025
Cyber AI Profile
Initial preliminary draft.

5 FEB 2026
AI Agent Identity & Authorization
Concept paper. Initial public draft.

7 APR 2026
Trustworthy AI in Critical Infrastructure
Concept note. Profile development.

CLAWOLF MILESTONES

SEP 2024
Foundational Research
Security architecture research begins.

JUL 2025
Architecture Planning
Autonomy defines the operating approach.

OCT 2025
Runtime Kernel Architecture
Five-layer operating lifecycle.

JAN 2026
Multi-Pipeline & 21 Logic Cores
Specialist analysis and decision context.

FEB–APR 2026
Context, VulnOps & Incorporation
CLAWOLF becomes a Delaware C-Corp.

Parallel chronology only. NIST does not endorse or certify CLAWOLF.

What Is AS-OS?

The Operating Layer Across Your Security Estate

AS-OS coordinates security analysis and response across connected threat-domain products, SIEM and SOAR tools and your SOC workflows. Your team sets the authority. The operating layer connects the evidence, the decision, the permitted action and its recorded result.

YOUR ORGANIZATION — Policy, protected assets and human approval boundaries

CLAWOLF AS-OS
Governed Operating Layer

Detection	Decision	Action	Control	Learning
Connect the evidence	Assess the business risk	Execute within authority	Check policy and effects	Retain outcomes

CONNECTED SECURITY ESTATE — **Threat-Domain Products / SIEM / SOC Workflows / SOAR**

Supported integrations and authorized actions only. Existing products retain their native controls.

WHAT CLAWOLF DOES

Security Decisions Your Team Can Explain And Control

An alert still leaves your team to judge the risk, coordinate a response and confirm the result. CLAWOLF AS-OS brings those steps together.

UNDERSTAND THE RISK

What Matters To This Business?

A suspicious login may threaten a critical service. Relate the evidence to asset importance, sector context and your policy.

Your team sees why a finding deserves attention.

CONTROL THE RESPONSE

What May The System Do?

Use operating modes, approval thresholds and protected targets to define the permitted response.

Your team retains an explicit boundary around autonomy.

REVIEW THE OUTCOME

What Actually Happened?

Follow the selected run, its decisions and recorded outcomes. Review missing evidence before closure.

Your team can explain the result and the remaining gaps.

THE AS-OS DECISION FOUNDATION

Knows Your Business, No Blind Decision

Business context, specialist analysis and traceable evidence inform what the system may do.

UNIQUE METHOD **Domain Intelligence Layer**

98 Instruments

US / EU / TR / ME / GCC

Used in the Decision layer. Organization, sector and regulatory context informs specialist analysis and remains in the recorded decision for review.

UNIQUE METHOD **Autonomous Decision Engine**

21 Logic Cores

Threat-Domain Analysis

Specialist cores assess domains such as identity, network, cloud and adversarial AI. Evidence-weighted verdicts support consensus, abstention or human review.

UNIQUE METHOD **Autonomous Validation Fabric**

12 Validation Domains

Forensic Sandbox Mesh

A sandbox mesh spans malware, identity, cloud, network and more. Supported checks validate evidence and effects. Unavailable checks remain explicit.

UNIQUE METHOD **Decision DNA**

Decision History

Context, Policy And Evidence

Preserves the context, core reasoning, policy version and confidence behind a decision, together with validation, human approvals and the recorded execution outcome.

AUTONOMOUS CONTROL

End-To-End Runtime Governance

Autonomy stays within your organization’s authority, policy and human-approval boundaries throughout the run.

Before Action

Who May Act, On What?

Checks the tenant, permitted scope and operating policy. Protected targets, uncertainty and approval rules can block or escalate a response.

Your rules define the boundary.

During Action

What May Actually Execute?

The governed runtime dispatches bounded, authorized actions. Human approvals preserve their identity and timestamp. Receipts record the execution result.

Only authorized actions proceed.

After Action

Did The Response Work?

Supported verification checks the observed effect. Failures remain visible. Recovery and rollback follow the applicable control and validation gates.

Evidence supports closure.

CLAWOLF
ART OF GOVERNED AUTONOMY

AUTONOMOUS OPERATIONS ACROSS THE LIFECYCLE

Detection, Action And Learning

AS-OS connects incoming signals to permitted response and evidence-informed follow-up.

Autonomous Detection

The Evidence For A Decision

01 Ingest and normalize telemetry
02 Create traceable detection signals
03 Correlate activity across domains
04 Add business and threat context
05 Route evidence to specialist analysis

Your team sees a connected case.

Autonomous Action

The Authorized Response

01 Plan the response from the decision
02 Generate bounded execution intents
03 Respect approvals and target limits
04 Collect action receipts
05 Verify supported effects and recovery

Your team controls what may execute.

Autonomous Learning

The Recorded Outcome

01 Retain verified outcome events
02 Distinguish success from failure
03 Preserve decision and policy lineage
04 Support recurrence revalidation
05 Use evidence in subsequent analysis

Your team keeps a reviewable history.

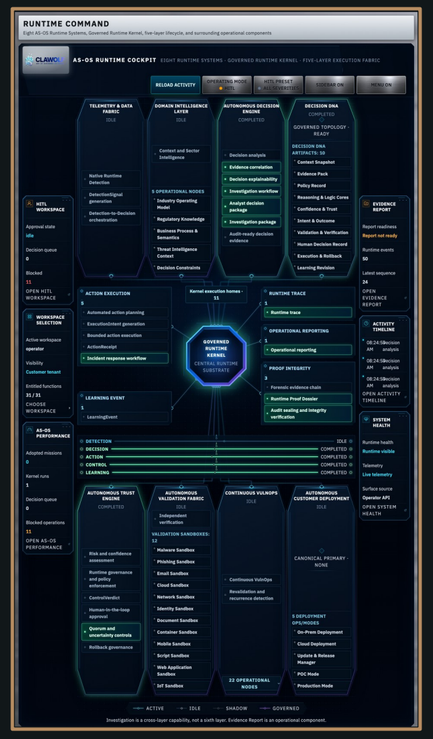
CLAWOLF
ART OF GOVERNED AUTONOMY

PRODUCT IN VIEW / RUNTIME COMMAND

AS-OS Dashboard Real View

The full cockpit connects eight runtime systems and their surrounding workspaces to the governed runtime.

FULL COCKPIT



RUNTIME DETAIL



The selected run ties the diagram to its recorded activity and evidence.

Full Context Around Every Run

HITL, workspace, performance, evidence, activity and system health remain visible around the operating architecture.

Full Cockpit: Demonstration Environment / Center Detail: Recorded Controlled-Lab Run

MITRE ATT&CK And The Full 21-Core Map

Full map views show threat references and the complete specialist registry. Activity remains tied to recorded runs.

MITRE ATT&CK MAP

Reconnaissance	T1595	T1596	T1597	T1598	T1589	T1590
Resource Development	T1583	T1584	T1585	T1586	T1587	T1588
Initial Access	T1190	T1133	T1566	T1078	T1199	T1200
Execution	T1059	T1203	T1072	T1129	T1106	T1204
Persistence	T1547	T1543	T1546	T1574	T1053	T1078
Privilege Escalation	T1548	T1134	T1547	T1543	T1574	T1055
Defense Evasion	T1140	T1070	T1036	T1027	T1562	T1055
Credential Access	T1003	T1110	T1555	T1212	T1187	T1606
Discovery	T1082	T1083	T1046	T1057	T1135	T1016
Lateral Movement	T1210	T1534	T1570	T1021	T1091	T1080
Collection	T1560	T1123	T1119	T1115	T1530	T1213
Command and Control	T1071	T1132	T1573	T1008	T1095	T1572
Exfiltration	T1048	T1041	T1011	T1052	T1567	T1029
Impact	T1485	T1486	T1489	T1491	T1495	T1498

UNIQUE METHOD **AUTONOMOUS DECISION ENGINE / ALL 21 CORES**



21 Specialist Perspectives For A Contextual Decision

The Analyst’s Main Working Area

Your analyst defines the mission, prepares its scope and follows the execution path from one workspace.

FULL WORKSPACE



MISSION PREPARATION / DETAIL



A Defined Objective And A Visible Readiness Path

Prepare The Mission

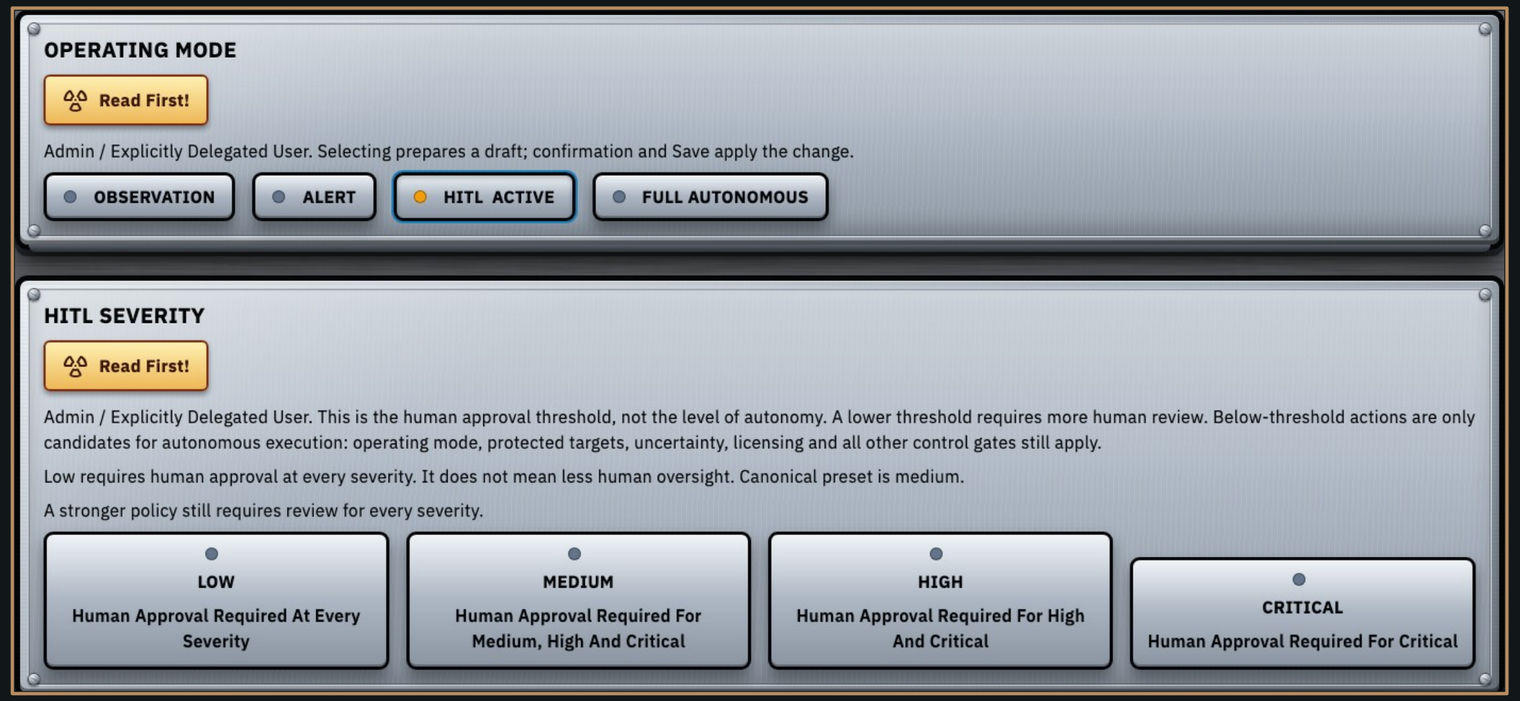
Set the objective, owner, priority and scope. Adopt an existing run where appropriate.

Check The Path

Review readiness, attempts and closure evidence. Help stays beside the operational UI.

Your Team Sets The Boundaries Of Autonomy

Operating mode and human-approval thresholds remain separate controls. More autonomy never removes the other safety gates.



Admin Or Explicitly Delegated User

Observation: observe and enrich.
Alert: observation plus notifications.
HITL: response follows approval gates.
Full Autonomous: only eligible, authorized actions may proceed.

Low Means More Human Review

Low requires approval at every severity. Higher thresholds still respect protected targets, policy and uncertainty controls.

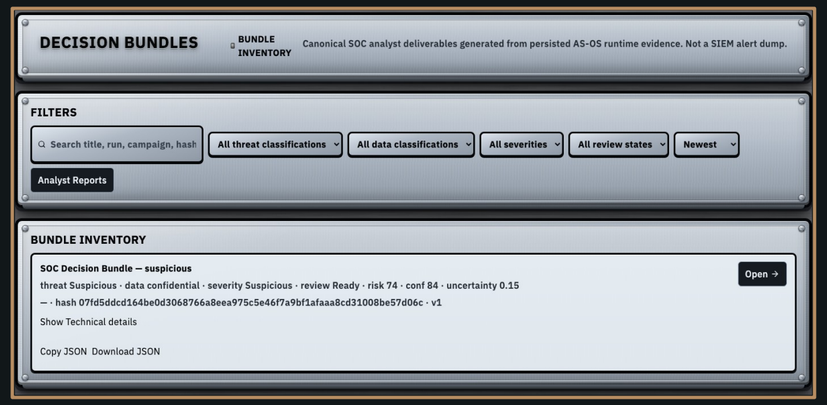
CLAWOLF
ART OF GOVERNED AUTONOMY

PRODUCT IN VIEW / EVIDENCE & REPORTING

Decision-Ready Evidence During The Run

Runtime bundles bring the reasoning, investigation context and recommended next actions together for the analyst.

DECISION BUNDLE / 13 SECTIONS



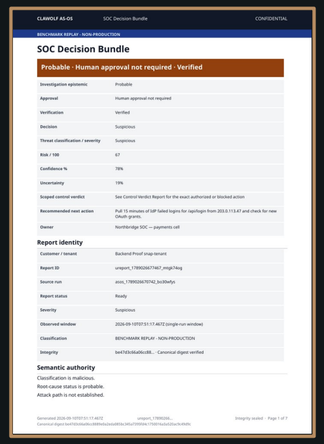
Executive summary, decision, confidence, risk, evidence summary and graph, intent, 21-core outputs, sector and context intelligence, uncertainty, next action and audit reference.

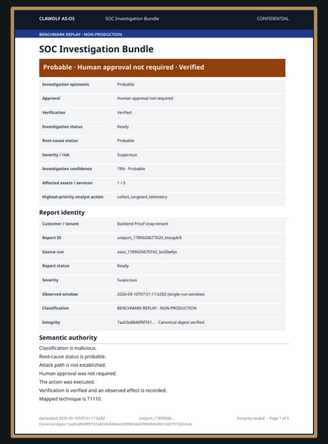
INVESTIGATION BUNDLE / 10 SECTIONS



Root cause, timeline, related assets, identities and events, attack path, MITRE mapping, evidence expansion, findings and recommended actions.

ACTUAL PDF EXPORTS





Real Backend Output

PDF + JSON reports share the source run and preserve evidence integrity.

Runtime Evidence And Recommended Actions For A Better-Informed Decision

UNIQUE METHOD **AUTONOMOUS CUSTOMER DEPLOYMENT**

Autonomous Deployment Built Into AS-OS

A governed operating system that includes its own deployment preparation, placement and go-live control.



Less Manual Implementation Work

Environment Placement

Select cloud, on-premises, hybrid or air-gapped placement within the supported delivery scope.

Policy Before Connection

Prepare operating rules, engagement and governance before ingestion and execution.

Controlled Go-Live

Review entitlement, credentials and connector readiness. Keep blockers visible before Base-Live & Go.

Beyond SOC Automation: The AS-OS Difference

Compare what makes the decision, governs execution, validates the result and preserves proof.

Platform	Decision Method	Runtime Governance	Validation Method	Evidence & Proof	Deployment
Cisco / Splunk	ES agents + entity context	SOAR policy + approval gates	Agent findings + analyst validation	Case evidence + agent / action audit	ES / SOAR installation
Palo Alto Cortex	Agent planning + Cortex context	RBAC + policy + action approvals	Tool action-contract checks	Agent / incident audit trail	Cortex platform + integration setup
CrowdStrike Falcon	Expert agents + Enterprise Graph	Per-workflow autonomy + permission gates	Agent testing + analyst oversight	Agent traces + workflow audit logs	Falcon + connector configuration
Arctic Wolf Aurora	Expert swarm + customer context	Validated scope + human escalation	Swarm Judge + AI Trust Engine	Swarm / human quality review	Provider-operated managed SOC
Torq HyperSOC	Socrates + agents + business rules	Workflow guardrails + HITL / HOTL	Case / verdict review + memory feedback	Case-thread evidence + action logs	Platform / integration configuration
Microsoft Security	Copilot agents + security context	Agent identity + configured permissions	Tested agents + reviewable node maps	Agent outputs + admin audit logs	Cloud agent / plugin configuration
CLAWOLF AS-OS	**21-Core Decision + Domain Intelligence**	**Native Trust Engine + ControlVerdict / HITL**	**Independent AVF: 12 forensic engines**	**Decision DNA: 10 linked artifacts**	**Native autonomous customer deployment**

CLAWOLF integrates these methods in one five-layer runtime. Workflow controls and an integrated operating architecture are different scopes.

CLAWOLF THREAT-DOMAIN PRODUCT PORTFOLIO

Clawolf Native Defense Armory

Machine-speed threats demand timely domain intelligence and response. Manual handoffs and static playbooks can become the bottleneck.

CLAWOLF AIS + Native AS-OS — **Domain Defense With Governed Response**

DECEMBER 2026

CLAWOLF Identity Defense Intelligence

Connect agent behavior to responsible owners and delegated privileges.

CLAWOLF Data & Database Intelligence

Relate AI access to sensitive data and database risk.

JANUARY 2027

CLAWOLF Behavioural Intelligence

Connect agent activity with patterns across people and workloads.

CLAWOLF Insider Defense Intelligence

Relate trusted-access misuse to delegated AI activity.

FEBRUARY 2027

CLAWOLF Supply Chain Exposure

Trace compromised dependencies into agent tools and services.

CLAWOLF Cloud Exposure & Attack In Motion Intelligence

Relate agent permissions to cloud exposure and active attack paths.

The Customer Gain — **Broader context around AI risk, with an authorized response path through AS-OS.**

CLAWOLF
ART OF GOVERNED AUTONOMY



NEW PRODUCT / CNDA FLAGSHIP

Introducing CLAWOLF AIS

Standalone AI Security and Governance

Cloud annual list: Core $229 / Signature $299

CORE: EIGHT FEATURE GROUPS

01 **AI Discovery & Security Posture**

02 **Agent Identity & Delegated Authority**

03 **Prompt Injection & Multimodal Protection**

04 **RAG & Memory Protection**

05 **MCP, Tools, Skills & AI Supply Chain**

06 **Models, Sensitive Data & Output Protection**

07 **Continuous Red Teaming & Resource Defense**

08 **Runtime Detection & Governed Containment**

SIGNATURE: CORE + SIX ADVANCED MECHANISMS

Causal Evidence Integrity
Distinguish copies from independent support.

Semantic Effect Assurance
Check what an allowed action actually changed.

Joint-Outcome Defense
Detect harmful combinations across agents.

Causal Decontamination
Quarantine affected poisoned-context derivatives.

Mission-Preserving Defense
Preserve work that remains demonstrably safe.

Continuous Assurance Invalidation
Recheck assurance when assumptions change.

Both editions include AI Trust, AI Governance and CLAWOLF Immunity

USD per billable asset/year. Paid annual ACD is separate. AS-OS integration is optional.

AIS planned capabilities / Target release: 30 November 2026 / Commercial scope confirmed in the authorized quotation

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 / 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**
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**

Native Integration: What the Customer Gains

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

Native Integration Value	Combined Mechanism	Customer Benefit
Attributable Context	Link AI evidence, identity and affected assets to the enterprise case.	**Investigate the AI incident with business and security context, while retaining its evidence origin.**
Current Authority	Carry the tenant, current authority and mission/rate limits across handoffs.	**Retain approvals and resource limits without a handoff creating new permission.**
Coordinated Containment	Connect AI-local controls to authorized identity, data and infrastructure actions.	**Limit the response to the supported attack path and its authorized scope.**
Observed Effects and Recovery	Link decisions and receipts to independent observation and recovery evidence.	**Know whether protection actually changed the target, and follow the evidence through recovery.**
Trust After Change	Connect relevant versions and causal context to affected assurance and cleanup.	**Requalify changed assumptions and preserve unaffected work that remains safe and authorized.**