PRODUCT OVERVIEW · 2026

CYPRAWALL

The Firewall for AI Conversations

Every prompt. Every response. Every AI destination — inspected before risk leaves the network.

A PRODUCT BY

Baseel
Security - Technology - Transformation

Your Employees Are Already Pasting Secrets Into AI Tools

Source code, customer records, credentials, and regulated data go into ChatGPT, Claude, Gemini — and a growing list of tools security hasn't approved yet — every day.

Traditional DLP, CASB, and SIEM tools were built to watch files, email, and known SaaS apps, not free-form AI chat. Without prompt-level visibility, a CISO can't answer what actually matters: which AI tools are in use, what left through them, and whether a disclosure already happened.

34

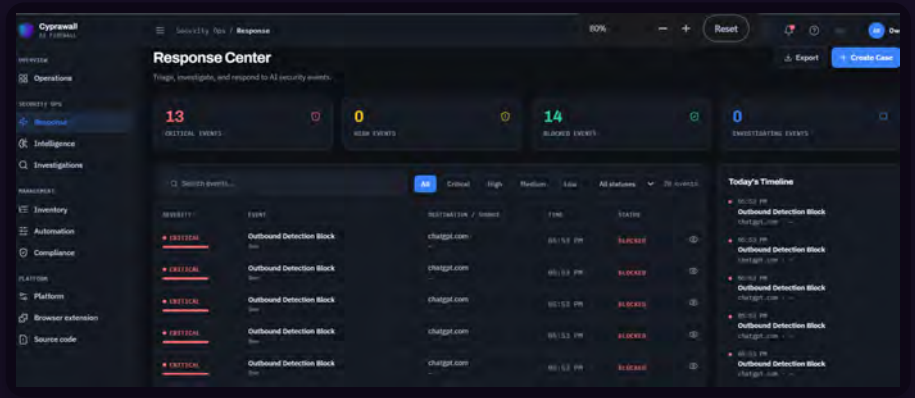
AI integrations found in one scan

11

had live read access to CRM data

6

flagged high-priority within a day



Response Center — 5,240 raw AI events logged in 30 days, before policy.

EVERY UNAPPROVED TOOL IS AN UNLOGGED EXIT POINT

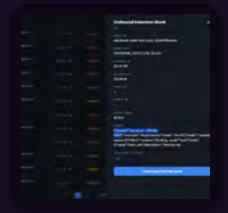
Contractors run unmonitored scripts straight against your APIs. Traditional firewalls don't know AI traffic exists — so nothing is blocked, and nothing is logged.

What Is an AI Firewall?

An AI firewall inspects every prompt and every response — in both directions — before a sensitive disclosure leaves the network, or a risky reply reaches an employee. It doesn't ask people to stop using AI; it governs the conversation itself.

THE INSPECTION PATH

1. **Prompt Sent** — Employee uses an AI tool
2. **Routed Through Cyprawall** — Standalone proxy or ICAP mode
3. **Scanned** — PII, source code, credentials
4. **Policy Applied** — Block, redact, flag, or allow
5. **Logged** — Audit trail, live dashboard



A real blocked event: an SSN caught outbound to claude.ai, stopped in time.

Discover. Score. Enforce. Decommission.

Cyprawall doesn't ask employees to change which AI tools they reach for — it inspects the conversation itself and applies the policy security sets, before risk ever leaves the network.

01



Discover

Surfaces every AI destination in use, sanctioned or not — including the ones nobody's approved yet.

02



Score

Rates every destination 0–100 on data sensitivity, permission reach, and staleness, so the riskiest surface goes first.

03



Enforce

Inspects prompts and responses for credentials, source code, and regulated data — then blocks, redacts, flags, or allows, live.

04



Decommission

Drafts the full offboarding sequence for any tool that needs to go; nothing executes without a human's sign-off.

Every Way Sensitive Data Tries to Leave

Credential Leaks

API keys, tokens, and passwords caught before they leave in a prompt.

Sensitive-Data Exposure

Source code, records, and strategy content flagged the moment they're pasted.

Shadow AI Discovery

Surfaces every AI destination in use — including ones nobody's approved.

Risky AI Responses

Two-way inspection reviews what the model sends back, not just what goes in.

Dynamic Risk Score (0–100)

Ranks every destination on real signals, so the riskiest surface goes first.

Configurable Policy Engine

Block, redact, flag, or allow — changed by security, live, no engineer needed.

Human-Gated Decommission

Offboarding drafts itself; nothing executes until a named approver signs off.

Dual-Mode Deployment

Standalone gateway or an ICAP module inside Zscaler/Netskope — no rip-and-replace.

Per-Device Identity & Revocation

Scoped keys per device; revoke one and the rest of the fleet stays live.

Tamper & Stale-Agent Detection

Dark devices and stale or tampered agents are flagged automatically.

How It Works

Four engines, one audit trail. Discovery feeds risk scoring; risk scoring informs enforcement; every action writes to a single append-only log that powers the CISO dashboard, live.

Discovery Engine	OAuth scraper + API/SaaS log ingestion builds the agent registry.
Risk Scoring Engine	Permission scope, sensitivity, and staleness written as a 0–100 score.
Inspection Engine	Prompt-aware parsing + NER reads the conversation, not just the wire.
Enforcement Layer	Standalone TLS gateway, or an ICAP module inside an existing SWG.
Policy Engine	Block, redact, flag, or allow-with-audit — live, no service restart.
Audit Log Service	Append-only record of every block, redaction, flag, and decommission.
Decommission Agent	Drafts the offboarding sequence; a human approver executes.
CISO Dashboard	Operations, response, investigations, and compliance — live.



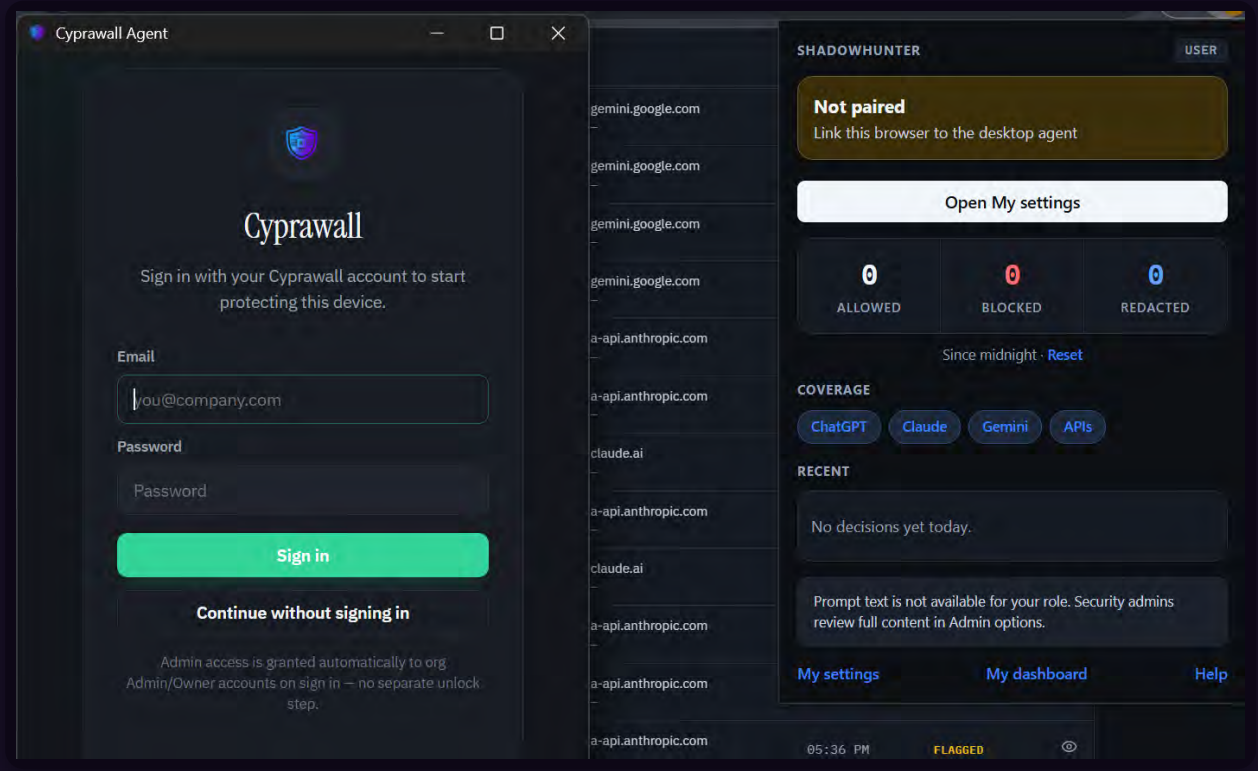
Operations Dashboard — 30-day threat activity and requests by model, fed live by every engine below.

EVERYTHING WRITES TO ONE LOG

Because every engine below feeds the same append-only audit trail, an investigation that used to mean stitching together three tools now starts — and usually ends — on one screen.

Every Device, Every Surface

One policy engine, enforced everywhere your people already reach for AI — corporate desktops, BYOD phones, browsers, and MDM-managed devices, under a single identity system.



The desktop agent sign-in and browser extension enforce one policy, wherever AI is reached.

Browser Extension

Intercepts ChatGPT, Claude, Gemini, and unsanctioned tools alike.

Desktop Agent

Enforces policy on every managed machine — Windows, macOS, and Linux.

Mobile Agent

Extends enforcement to iOS and Android, corporate-owned or BYOD.

Identity + Device Service

Scoped keys per device, instant revocation — losing one never compromises the fleet.

Two Modes. One Product. No Infrastructure Conflicts.

Mode 1 — Standalone Gateway

For teams with no existing SWG/DLP, or tools at end-of-life. A corporate root CA, distributed via MDM, routes AI-domain traffic through Cyprawall's TLS gateway on port 8080.

Mode 2 — SWG Integration (ICAP)

Already running Zscaler, Netskope, Forcepoint, or any ICAP-compatible gateway? Cyprawall's inspection and policy engine plug in as an ICAP verdict service — no client-side change, existing DLP policy stays active.

COVERS EVERY MAJOR AI DESTINATION

OpenAI · Claude · Gemini · Microsoft Copilot · Perplexity · Cohere · Mistral — plus any HTTP/HTTPS endpoint via a configurable domain list, including local models via Ollama, LM Studio, and Jan.

PLUGS INTO WHAT'S ALREADY THERE

Identity & SSO

Azure AD, Google Workspace, Okta, OneLogin

Security Stack

Zscaler, Netskope, Forcepoint, Palo Alto NGFW

SaaS, SIEM & Alerts

Slack, GitHub, M365, Splunk, Sentinel, QRadar, Jira, ServiceNow, PagerDuty

Runs on-premises, private cloud (AWS / Azure / GCP VPC), or hybrid. Minimum 4-core / 8GB; production 8-core / 16GB with 500GB SSD for audit logs.

What This Looks Like in Practice

01 Answering the Auditor's Question

Scenario — A financial firm starts a SOC 2 audit; the auditor asks which AI tools reach customer data — there's no registry, no answer. **Resolution —** Discovery runs against Azure AD and Slack logs, surfacing 34 integrations. Risk scoring flags 6 as high-priority within a business day.

02 Stopping Proprietary Code at the Door

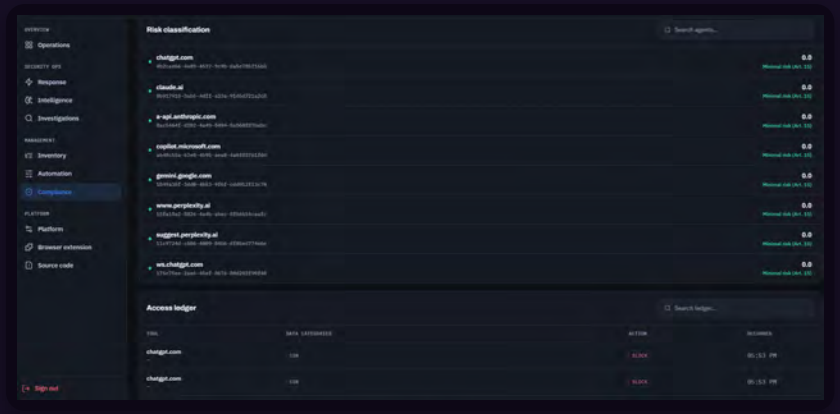
Scenario — An engineer pastes 200 lines of proprietary algorithm code into a third-party AI coding assistant. Security has no idea. **Resolution —** Hash-verified fingerprinting identifies the payload as it leaves the network. The policy engine blocks it — no code reaches the model.

03 Catching PII in a Browser Paste

Scenario — A support agent pastes customer names, emails, and account numbers into a browser AI tool to draft a bulk reply — free text no DLP tool would see. **Resolution —** NER and regex detect the PII as the gateway decrypts the request; redaction strips it before the prompt reaches the provider.

04 Retiring an Orphaned Agent Safely

Scenario — An agent set up 14 months ago still has write access to the data warehouse. The analyst who owned it left three months ago. **Resolution —** The next scan flags it critical — orphaned, sensitive access, 90+ days stale. A decommission draft routes to the security lead for one approval.



Every agent's risk classification and access ledger — audit-ready, automatically.

THE RESULT

"An audit that took a week now takes an afternoon."

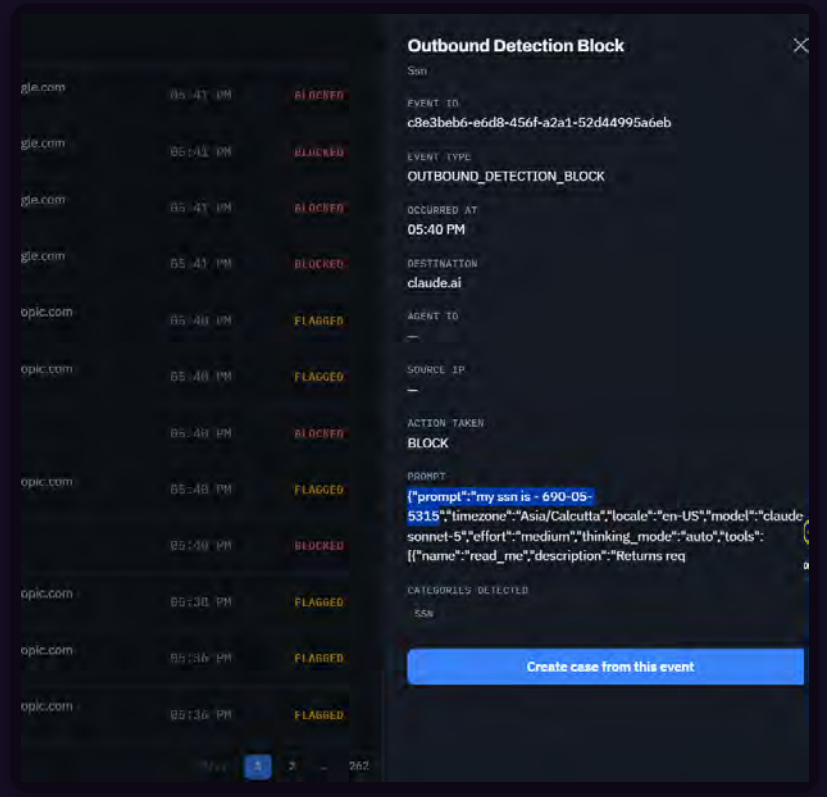
"We think people use AI tools" becomes "this left the building, when, and to whom."

Why This, Not Three Tools Duct-Taped Together

Point solutions catch only fragments. None combine live inspection, discovery, and decommissioning — three tools, three audit trails, one gap to exploit.

CAPABILITY	OURS	COMPETITORS
Hash-verified source code fingerprinting	✓	Partial / absent
OAuth + agentic shadow-AI discovery	✓	Browser / network only
Decommission workflow, human approval gate	✓	Not found
Dual deployment: standalone or SWG plug-in	✓	Single mode only
Prompt-aware, two-way inspection	✓	Traffic-level, one-way

Runs on Zscaler/Netskope as an ICAP verdict engine — no rip-and-replace. Guards the AI your people use, not the AI you ship.



A real blocked event: an SSN caught outbound, stopped in time.

Enterprise-Ready, From Day One

- **ISO 27001 Certified**
- **CREST Certified**
- **SOC 2 Compliant**
- **Quality Assured**

THREE TIERS, MATCHED TO GOVERNANCE MATURITY

DISCOVER

Initial AI visibility

Shadow AI discovery, dashboard, and business-hours support.

PROTECT

Real-time AI data protection

Adds live inspection, policy enforcement, and full SOC support.

GOVERN

Enterprise AI compliance

Adds decommission workflows, compliance mapping, and 24/7 + SLA support.

All tiers include the full audit log and compliance evidence pack. Pricing — contact sales for a scoped quote; volume and multi-site pricing on request.

Baseel Partners LLP

A distinguished global entity focused on Cybersecurity, Data and AI Governance, and Risk and Compliance — with clients and partners across Europe, the UK, Asia, and the MENA region. Founded in 2012, our consultants are driven by a culture of innovation, integrity, and hard work, never compromising on quality or security.

KEY CREDENTIALS

- ✓ Microsoft, Google, and AWS Partner
- ✓ ISO 27001 Certified Company
- ✓ CREST-accredited Pen Testing Company
- ✓ Cyber Essentials certification body since 2016
- ✓ Approved vendor, Dubai Cyber Force Program
- ✓ 45+ in-house developed, IP-owned products

800+ Clients worldwide

2,000+ Projects delivered

3,000+ Professionals globally

GLOBAL FOOTPRINT

UK (HQ) · Saudi Arabia · UAE · Qatar · India · Gibraltar · USA

Give employees the freedom to use AI. Don't lose control of the data.

Request a live demo today.

EMAIL	contactus@baseel.com
PHONE	+44 20 4540 0111
WEB	www.baseel.com

167–169 Great Portland Street, London, W1W 5PF, United Kingdom

Baseel
Security - Technology - Transformation