PRODUCT DATASHEET

The Future of AI Security Starts Here

Enterprise-grade AI threat protection and zero-retention data governance — deployed in under 5 minutes with zero code changes. Block 27+ threat categories with a 5-layer intelligent defense system.

2,500+	5	<5
DAILY AI THREATS	DEFENSE LAYERS	MIN DEPLOY

67%

Year-over-year increase in AI-specific attacks

↑ Prompt injection & jailbreak volume

$4.2M

Average cost per AI-related data breach

↑ 45% annually

89%

Of enterprises lack purpose-built AI security

Critical infrastructure exposed

COMPLIANCE READY

GDPR · HIPAA · PCI-DSS · EU AI ACT · NIS2 · SOX

PLATFORM ARCHITECTURE

APIRE Multi-Layer AI Defense System

A transparent security proxy requiring a single endpoint change. 100% compatible with all OpenAI-compatible APIs, local LLMs, and custom models — zero code modifications required.

One URL change — zero code modifications

Replace `api.openai.com` with `app.apire.io` — all SDKs, auth methods, and request formats remain identical. All AI providers supported.

LAYER 0

Sentinel — AI-Powered Gateway

AI-powered content moderation at scale. Updated by AI in real-time with hundreds of attack categories and 100K+ attack vector detection at minimal latency. First line of defense on every request.

100K+ Attack Vectors | AI-Updated | Minimal Latency

Detection	100K+ Vectors
Training	Continuous AI
Latency	Minimal

LAYER 1

Content Safety Shield

AI-powered moderation across 14 comprehensive safety categories. Delivers binary SAFE/UNSAFE verdicts with contextual understanding across 32 languages. Very high detection accuracy, very low false positive ratio.

14 Categories | 32 Languages | Binary Verdict | Auto Severity Score

Categories	14
Languages	32
False Pos.	Very Low
Engine	Powered by AI

LAYER 2

AI Threat Protection Shield

Defends against 13 AI-specific threats — 5 core categories and 8 advanced categories. Semantic threat analysis understands intent beyond keywords, with auto-tuning and manual flexibility. Zero-day protection.

5 Core + 8 Advanced | Prompt Injection | Jailbreak Defense | Zero-Day

Core Threats	5
Advanced	8
Analysis	Semantic AI
Tuning	Auto + Manual

LAYER 3

Multi-Word Pattern Protection

Lightning-fast dictionary-based pattern recognition with single and multi-word detection. Bring your own custom dictionary in any language. Real-time updates with auto-refresh — zero downtime, zero false negatives.

Custom Dictionary | Any Language | Hot Updates | 0% False Neg.

Detection	Single & Multi
Updates	Real-Time
Languages	Any
False Neg.	0%

LAYER 4

Data Masking Fortress

950+ pre-configured DLP rules protecting PII, PHI, PCI data. Zero-trust architecture with 0% false negatives on critical data. Inline masking and anonymization with visual rule builder — no regex expertise required.

950+ DLP Rules | Inline Masking | Zero-Trust | 0% False Neg.

DLP Rules	950+
Data Types	150+
Architecture	Zero-Trust
False Neg.	0%

SUPPORTED AI PROVIDERS

GLOBAL CLOUD PROVIDERS

OpenAI | Anthropic | Google Gemini | xAI Grok | Mistral AI | DeepSeek | Meta Llama | Azure OpenAI

AWS Bedrock | Groq | Cohere | Together AI | Perplexity | Ollama | Fireworks AI | Any OAI API

LOCAL & ON-PREMISE

Ollama — Local model runtime

LM Studio — Desktop LLM runner

Jan.ai — Open-source LLM client

Custom / Private — Any self-hosted model

LAYER 1

Content Safety Shield

AI-powered moderation engine with state-of-the-art models. Provides contextual understanding to prevent false positives with automatic severity scoring.

14 DETECTION CATEGORIES

- Violent Crimes
- Non-Violent Crimes
- Indiscriminate Weapons
- Sex-Related Crimes
- Child Sexual Exploitation
- Hate Speech
- Suicide & Self-Harm
- Defamation
- Sexual Content
- Specialized Advice
- Elections
- Privacy Violations
- Intellectual Property
- Code Interpreter Abuse

KEY BENEFITS

- Binary SAFE/UNSAFE verdicts — no ambiguity
- Contextual understanding prevents false positives
- Real-time classification with automatic severity scoring (Critical / High / Medium)
- Multi-language support across 32 languages
- Powered by state-of-the-art AI models

LAYER 2

AI Threat Protection Shield

Advanced defense against AI-specific attacks with semantic threat analysis that understands intent — not just keywords. Enterprise-grade coverage with auto-tuning.

5 CORE AI THREAT CATEGORIES

- Prompt Injection
- Jailbreaking
- Data Exfiltration
- Social Engineering
- Model Inversion

8 ADVANCED THREAT CATEGORIES

- Adversarial Attacks
- Compliance Violations
- IP Theft
- Shadow AI Usage
- API Security
- Content Abuse
- Business Logic Attacks
- Context Attacks

KEY BENEFITS

- Semantic analysis — understands intent, not just keywords
- Auto-tuning protection adapts to your threat landscape
- Zero-day defense against emerging attack patterns
- Manual tuning for full flexibility
- Threat correlation identifies multi-vector attacks
- Real-time correlation and threat analysis

Zero-Retention Architecture — Your Data Never Leaves Your Control

APIRE processes all AI prompts and responses in-stream without storage. No data used for training. No data lake exposure. Full audit trail from metadata only — privacy by design.

LAYER 3

Multi-Word Pattern Protection

Intelligent pattern recognition at scale. Dictionary-based threat detection with single and multi-word matching. Bring your own industry or company-specific dictionaries in any language.

- Bring your own dictionary in any language
- Single and complex multi-word phrase matching
- Auto-refresh updates without downtime
- Extremely fast lookups — optimized for performance
- Zero false positives on configured patterns
- Custom pattern library for any industry vertical

LAYER 4

Data Masking Fortress

Enterprise-grade data loss prevention with 950+ pre-configured rules. Zero-trust architecture ensures every request is scanned — no exceptions. 0% false negatives on all critical data types.

- 950+ pre-configured PII, PHI, PCI rules out-of-box
- Zero-trust — every request scanned, no exceptions
- Perfect accuracy — 0% false negatives on critical data
- Inline blocking and masking / anonymization
- Visual rule builder — no regex expertise required
- Compliance templates: GDPR, HIPAA, SOX, PCI-DSS
- Complete audit trail for regulatory documentation

LAYER 4 — DATA LEAKAGE DETECTION CATEGORIES

PII — Personal Data Names, addresses, IDs, national IDs, passports	**PHI — Health Information** Medical records, diagnoses, prescriptions	**PCI — Payment Card Data** Card numbers, CVVs, bank account details
Company Confidential & IP Trade secrets, source code, proprietary data	**Financial Information** Revenue data, forecasts, M&A details	**Credentials & API Keys** Passwords, tokens, API keys, secrets
EU GDPR Data Special categories under EU regulation	**Regulatory Obligations** Compliance-relevant data under 12+ standards	**Indicators of Compromise** IPs, hashes, domains, malware signatures
Suspicious Activity Anomalous user behavior detection	**Acceptable Use** Policy violation monitoring and enforcement	**Custom — User Defined** Unlimited business-specific rule creation

PLATFORM SUITE

Complete APIRE Product Suite

Beyond the security proxy — APIRE delivers purpose-built products covering every AI usage surface in your organization, from employee chat to developer IDEs to autonomous agents.

APIRE Secure Chat

A fully secured ChatGPT-style end-user chat interface with built-in monitoring and logging. Enables employees to use AI safely while enforcing all 5 security layers on every message — eliminating Shadow AI risk.

- Secured ChatGPT-style interface for employees
- All 5 APIRE security layers enforced per message
- Real-time monitoring and session logging
- Prevents Shadow AI — sanctioned, secure alternative
- AD/LDAP identity integration and user management

APIRE Browser Agent

Browser-level protection for employees using web-based AI tools directly. Extends APIRE's 5-layer security to any browser AI interface without infrastructure changes. Blocks data leakage and enforces corporate policies at the browser level.

- Protects ChatGPT, Claude, Gemini and all browser AI tools
- Intercepts and analyzes prompts before transmission
- Enforces corporate AI usage policies at browser level
- Prevents unauthorized Shadow AI tool usage
- Centrally managed via APIRE Security Dashboard

APIRE Security Proxy

The core platform — a fully transparent AI API proxy requiring a single URL change. Routes all AI traffic through the 5-layer defense system. Supports inline real-time protection and out-of-band analytics topologies.

- Transparent proxy — zero code modifications required
- In-line real-time protection topology
- Out-of-band analytics and SIEM integration mode
- Full multi-tenant support with per-tenant policies
- OpenAI API compatible across all providers

APIRE Log Appliance

Specially customized appliance for comprehensive AI tracking, logging, and audit trail management. Collects all security events and threat detection data for compliance, forensics, and SIEM integration via webhook.

- Dedicated AI security event collection and storage
- Webhook connector for SIEM and SOC integration
- Detection-based granular logging with full context
- Optional full payload logging for forensic analysis
- Complete audit trail for regulatory compliance

RISK SURFACES COVERED

AI Chat Interfaces

Protect chat data leakage, content abuse, and AI attacks. Prevent unauthorized Shadow AI usage. Solution: Secure Chat + Browser Agent

Developer IDE & Coding Agents

Block credential and source code leakage through AI-powered IDEs and coding assistants. Solution: Security Proxy + IDE Integration

Enterprise Application AI

Protect data leakage from enterprise applications using AI APIs and safeguard AI-powered business workflows. Solution: Security Proxy

Chatbots, RAG & MCP Agents

Secure AI chatbot deployments, RAG pipelines, and MCP agent integrations. Block content abuse and data leakage. Solution: Security Proxy

DEPLOYMENT

Flexible Deployment for Every Organization

All deployment options include the complete 5-layer security architecture and enterprise support. Match your data sovereignty, infrastructure, and compliance requirements.

CLOUD FIRST

Cloud Deployment

Instant activation — zero infrastructure. Ready-to-use with global AI providers. Credit and subscription-based.

- Immediate deployment
- Security Proxy + Secure Chat
- Global LLM providers supported
- Dedicated and closed LLM models
- Automatic updates

On-Premise Deployment

Complete data sovereignty with air-gapped security. Physical or virtual appliance. Bring Your Own Key.

- Physical and virtual appliance options
- Air-gapped LAN operation
- AD / LDAP identity integration
- Security Proxy + Secure Chat
- Data stays on-site, always

Hybrid & Multi-Cloud

On-premise and cloud operating together. Unified policy, consistent security across any topology.

- Cloud + on-premise simultaneously
- Virtual appliance on-premise
- AI Brain provided from cloud
- Unified management and logging
- Seamless failover and scaling

IN-LINE PROTECTION
Real-time proxy — all AI traffic analyzed and filtered before reaching the provider.

OUT-OF-BAND ANALYTICS
Analytics appliance integrated via API with your existing SIEM and security tooling.

MULTI-TENANT SUPPORT
Proxy and Secure Chat support full multi-tenancy with per-tenant policy control.

FROM ZERO TO PROTECTED IN UNDER 5 MINUTES

01

Sign Up & Get API Key

Create your account at apire.io/signup. Receive your unique API key instantly.

```
→ apire.io/signup
→ Choose plan
→ Get API key
→ No credit card
```

02

Change API URL

One line change in your code. Fully OpenAI SDK compatible — nothing else to modify.

```
- api.openai.com
+ app.apire.io

// That's it.
```

03

Configure & Go Live

Set protection preferences, add custom rules, configure masking policies. Deploy.

```
→ Set policies
→ Custom rules
→ Configure masking
→ Go live ✓
```

Stop AI Threats. Start Securing What Matters.

Join organizations already protecting their AI stack with APIRE. Zero-retention. Zero code changes. Complete enterprise security from day one.

www.apire.io

- Full 5-layer security architecture
- Zero-retention — your data never stored
- Cloud, on-premise, and hybrid deployment
- EU AI Act & NIS2 compliant from day one
- 950+ DLP rules and 27+ threat categories
- Enterprise support included