



FIREFLY.AI

MCP Secure Fetcher + AI Shield

A Safety Layer for AI Agents That Visit Websites

Protecting AI systems from malicious websites while safeguarding your bottom line.



The Real Cost of Unsafe AI Browsing

When LLMs browse the web without protection, they expose enterprises to severe financial and operational risks:

Data Breach Exposure

Prompt injection attacks can leak sensitive internal data.
Average cost per incident: **\$4.45M** (IBM 2023)

System Downtime

Malicious websites manipulate agents, causing operational failures. Cost: **\$300K–\$1M per day**

Regulatory Penalties

SSRF attacks expose internal services, triggering compliance violations. GDPR fines: **up to 4% of revenue**

Trust Erosion

Automated phishing and data exfiltration damage reputation in ways that are **impossible to quantify**



Our Solution: Secure Web-to-LLM Pipeline

We intercept and validate every URL before content reaches your AI agents, creating a comprehensive safety barrier:

01

URL Validation

02

Safe Fetch

03

HTML Sanitization

04

Text Extraction

05

AI Shield Detection


06

Decision & Delivery



System Architecture

A multi-layered defense pipeline transforms untrusted web content into verified, safe context for your AI agents:





Financial Impact: Quantifying the Savings

What You're Risking Today

- **Wrong AI decisions:** Manipulated financial data leads to costly trades or strategic errors
- **Security false positives:** Misleading threat intelligence causes expensive downtime
- **Data leaks:** Customer information exposed, triggering \$4M+ in legal costs
- **Service outages:** Internal network exposure costs \$300K per hour in enterprise environments

\$500K

Minimum Annual Savings

For organizations with moderate AI agent usage

\$5M

Maximum Annual Savings

For enterprises with heavy AI-driven operations

1

Prevented Attack

Pays for the entire system many times over


What We Protect

- **Block malicious content:** Prevent multimillion-dollar data breach incidents
- **Stop SSRF attacks:** Secure internal networks from unauthorized access
- **Ensure accuracy:** Eliminate wrong model outputs that drive operational mistakes
- **Preserve reputation:** Maintain customer trust and brand integrity long-term



Where Companies Lose Money

Unprotected AI browsing creates vulnerabilities across your entire operation. Understanding these exposure points reveals the true cost of inadequate protection:



Operational Disruption

Compliance Violations


Incident Response

Customer Attrition



ROI: Where Our Solution Reduces Cost

MCP Secure Fetcher + AI Shield delivers measurable value across four critical dimensions:



Attack Prevention

Block threats before they reach production systems, eliminating incident response costs, breach notifications, and emergency remediation expenses.



Operational Continuity

Operational Continuity

Maintain uninterrupted AI operations by preventing manipulation that causes downtime, false alerts, and cascading system failures.



Compliance Assurance

Demonstrate proactive security controls to auditors and regulators, avoiding penalties while accelerating compliance certification processes.



Safe Scaling

Deploy AI agents confidently across more use cases without proportionally increasing security risk or requiring additional security team resources.



Future Roadmap & Opportunities

Technical Evolution

-  **Transformer-Based Detection**
-  **JavaScript Sandbox Analysis**
-  **Domain Reputation Scoring**
-  **Adaptive Sanitization**
-  **Continuous Learning**

Market Expansion

-  **Enterprise Middleware**
-  **Browser Extensions**
-  **RAG Protection**
-  **Regulated Industries**
-  **Strategic Partnerships**