

Sentinel AI

AI-Powered Autonomous Security & Threat Detection Platform

Theme:-Asia Track (Technical Safety)

Developed By:

Ashutosh Patel

B.Tech (Information Technology)

Maharaja Agrasen Institute of Technology (MAIT)

Teck Stack

- *React*
- *FastAPI*
- *Python*
- *LLM(models)*
- *Groq*
- *PostgreSQL*
- *Docker*


Problem Statement

As AI agents become more autonomous, they increasingly interact with APIs, files, web services, and external tools without continuous human supervision. This creates serious security challenges such as prompt injection, unauthorized tool execution, data leakage, hallucinated responses, and malicious attack chains.



Key challenges:-

- *Lack of real-time AI agent monitoring*
- *Prompt injection and jailbreak attacks*
- *Unauthorized tool execution*
- *Sensitive data exposure*
- *Hallucinated or unverified AI responses*
- *No centralized security visibility*
- *Difficulty in tracking attack chains*
- *Absence of unified AI governance*




Proposed Solution

Sentinel AI is a centralized AI Security & Governance Platform that continuously monitors autonomous AI agents, detects security threats, validates AI-generated responses, analyzes risk scores, and provides complete visibility into agent activities through an interactive dashboard.


Core Modules

- 🛡️ *Agent Management*
- 📡 *Live Monitoring*
- 🔥 *Prompt Firewall*
- ⚠️ *Threat Detection*
- 📊 *Risk Engine*
- 🤖 *Response Validation*





System Architecture


Sentinel AI follows a modular architecture where every AI request passes through multiple security layers before execution. Each module performs a dedicated responsibility to ensure secure, transparent, and trustworthy AI operations.



Tech Stack




System Architecture Integration



Sentinel AI is designed using a modular and scalable architecture to secure autonomous AI agents throughout their lifecycle. Each request passes through multiple security modules including Prompt Firewall, Threat Detection, Risk Engine, and Response Validation before reaching the dashboard. This architecture ensures real-time monitoring, transparency, and secure AI operations.



End to End Workflow



- Sentinel AI secures every AI interaction from the moment a user submits a request until the final response is delivered.
- Each request passes through authentication, prompt firewall, threat detection, risk analysis, response validation, and monitoring modules before being securely logged and displayed on the real-time dashboard.
- This layered workflow ensures secure, transparent, and trustworthy AI operations.

Result & Testing

Experimental Result



- ✓ Prompt Firewall blocked malicious prompt
- ✓ Threat Detection identified attack chain
- ✓ Tool Monitor tracked API/File access
- ✓ Response Validation verified AI output
- ✓ Timeline Replay reconstructed execution history
- ✓ Security Graph visualized agent interactions

SENTINEL AI

ACTIVE

OVERVIEW

Dashboard

MONITORING

Agents

Live Monitor

Audit Log

SECURITY

Prompt Firewall

Tool Monitor

Risk Engine

Threat Detection

Response Validation

INTELLIGENCE

Multi-Agent Judge

Timeline Replay

Attack Simulator

Security Graph

Settings

Sign out

Shubham

Analyst

Tool Monitor

Every tool call, scored and tracked in real time

17

Total Calls

0

Blocked

5

Dangerous (24h)

47.3

Avg Risk

0

Pending Approvals

Tool Calls (17)

Approvals (0)

Policies (13)

Analytics

Seed Policies

All Tools

Blocked

Dangerous

TOOL	STATUS	RISK	TIME
http_request https://httpbin.org/get	Success	25	1508ms 5h ago
search search	Success	5	187ms 5h ago
search search	Success	5	605ms 5h ago
memory_write memory	Success	8	0 5h ago
search search	Success	5	185ms 6h ago
search search	Success	5	258ms 6h ago
email_send email	Pending	84	6h ago
file_read /home/user/.env	Pending	88	6h ago
http_request https://httpbin.org/get	Success	25	1564ms 6h ago
search search	Success	5	518ms 6h ago

SENTINEL AI

ACTIVE

OVERVIEW

Dashboard

MONITORING

Agents

Live Monitor

Audit Log

SECURITY

Prompt Firewall

Tool Monitor

Risk Engine

Threat Detection

Response Validation

INTELLIGENCE

Multi-Agent Judge

Timeline Replay

Attack Simulator

Security Graph

Settings

Sign out

Risk Engine

Composite risk scoring across events, firewall, and tools

Overview

Agent Risk (5)

Alerts (0)

Rules (4)

Seed Rules

30

ELEVATED

System-Wide Composite Risk

Agent Distribution by Risk Level

5

LOW

0

ELEVATED

0

HIGH

0

SEVERE

0

CRITICAL

5

Total Agents

0

Active Alerts

0

Critical Alerts

1

Threats (1h)

Risk Score Trend (24h)

Top Risk Agents

1

LiveMonTest

Active

20

2

LiveMonTest

Active

7

SENTINEL AI

ACTIVE

OVERVIEW

Dashboard

MONITORING

Agents

Live Monitor

Audit Log

SECURITY

Prompt Firewall

Tool Monitor

Risk Engine

Threat Detection

Response Validation

INTELLIGENCE

Multi-Agent Judge

Timeline Replay

Attack Simulator

Security Graph

Settings

Sign out

Response Validation

Hallucination detection and factual grounding for agent outputs

Validator

History

Rules (5)

Response Validator

Quick test examples:

Grounded

Hallucinated

Overclaiming

Agent Response (to validate)

The Eiffel Tower is in Paris, France. It was completed in 1889.

Source Context (optional — documents/data the response should be grounded in)

The Eiffel Tower is located in Paris, France. It was completed in 1889 and stands 330 meters tall.

AI Deep Analysis

Grounded

Confidence: 95/100

Hallucination Risk: 0/100

6ms

CLAIMS (2)

SUPPORTED

confidence 95%

The Eiffel Tower is in Paris, France.

SENTINEL AI

ACTIVE

OVERVIEW

Dashboard

MONITORING

Agents

Live Monitor

Audit Log

SECURITY

Prompt Firewall

Tool Monitor

Risk Engine

Threat Detection

Response Validation

INTELLIGENCE

Multi-Agent Judge

Timeline Replay

Attack Simulator

Security Graph

Settings

Sign out

AI Security Graph

Live attack surface: agents, models, tools, and targets actually reached

19

Nodes

4

User

5

Agent

2

Model

6

Tool Category

2

Target

2 agent(s) have no observed tool-call activity (isolated nodes)

All Agents

https://httpbl...

http

terminal

search

LM Test

LM Test

gemini-2.5-flash

LiveMonTest

llama-3.3-70b

Test Agent

memory

Shubham

TestBrowserAgent

/home/user/.env

file

WatchAgent

email

LiveMonTest

8

Conclusion

Sentinel AI provides a unified AI security platform that enables secure deployment of autonomous AI agents through real-time monitoring, prompt protection, threat detection, response validation, and intelligent risk analysis. The platform improves transparency, governance, and trust in AI systems.

Future Scope

- Multi-cloud deployment
- Kubernetes support
- SIEM integration
- Enterprise RBAC
- Threat intelligence feeds
- AI SOC automation
- Compliance reporting

