SecOb

~~SecOps~~ for GenAI: Next-Gen Security Insights

"Beyond Logs & Metrics"

$whoami

Shubhendu Shubham

"sudo rm -rf / problems"

aka "Troubleshooter"

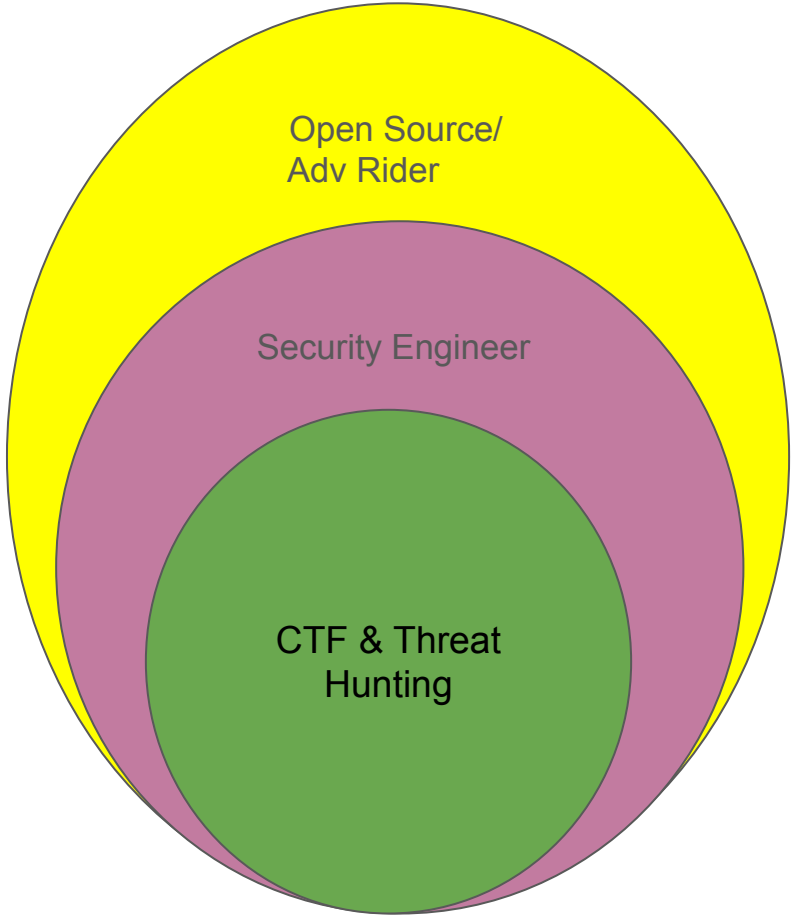

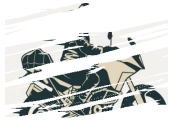

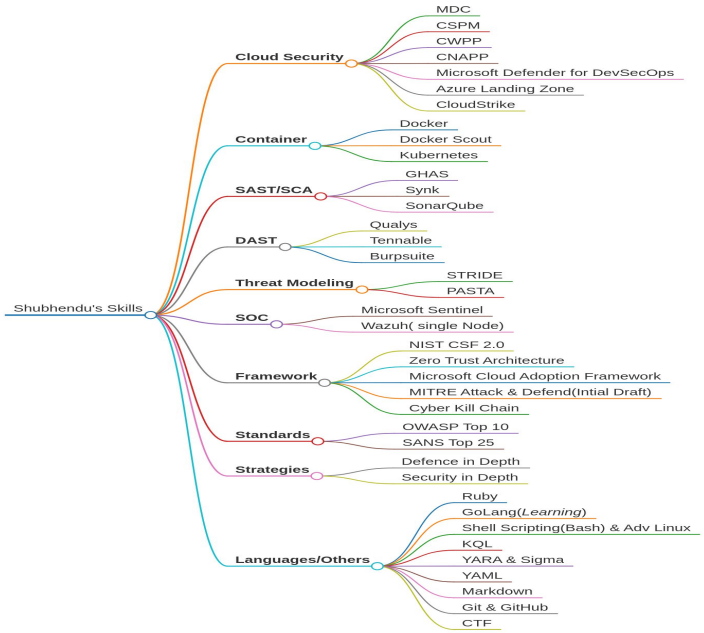
Shubhendu's Skills
Cloud Security
MDC
CSPM
CWPP
CNAPP
Microsoft Defender for DevSecOps
Azure Landing Zone
CloudStrike
Container
Docker
Docker Scout
Kubernetes
SAST/SCA
GHAS
Synk
SonarQube
DAST
Qualys
Tennable
Burpsuite
Threat Modeling
STRIDE
PASTA
SOC
Microsoft Sentinel
Wazuh(single Node)
Framework
NIST CSF 2.0
Zero Trust Architecture
Microsoft Cloud Adoption Framework
MITRE Attack & Defend(Intial Draft)
Cyber Kill Chain
Standards
OWASP Top 10
SANS Top 25
Strategies
Defence in Depth
Security in Depth
Languages/Others
Ruby
GoLang(Learning)
Shell Scripting(Bash) & Adv Linux
KQL
YARA & Sigma
YAML
Markdown
Git & GitHub
CTF


HACKTHEBOX
Subject
Matter
Expert
SME


JUST ANOTHER KUSTO HACKER
JUST ANOTHER
2025
KUSTO HACKER


ATTACKIQ
Attack Flows
How to Model and Sequence Attacks


OffSec
OSCC
SEC-100


CTF BADGES
CERTIFICATIONS
WINNER
AZURE NETWORK ENGINEER
ASSOCIATE
AZURE SECURITY ENGINEER
ASSOCIATE
EXPERT
EXPERT
ASSOCIATE
SC 100
AZ 305
AZ 104
AZ 700
AZ 500

Community


Azure Developer
Community


docker


KALI
BY OFFENSIVE SECURITY


OffSec
The Path to a Secure Future™

Disclaimer

"You can't protect

what you don't know you have."

— not sure

OWASP Top 10 for LLM Applications 2025

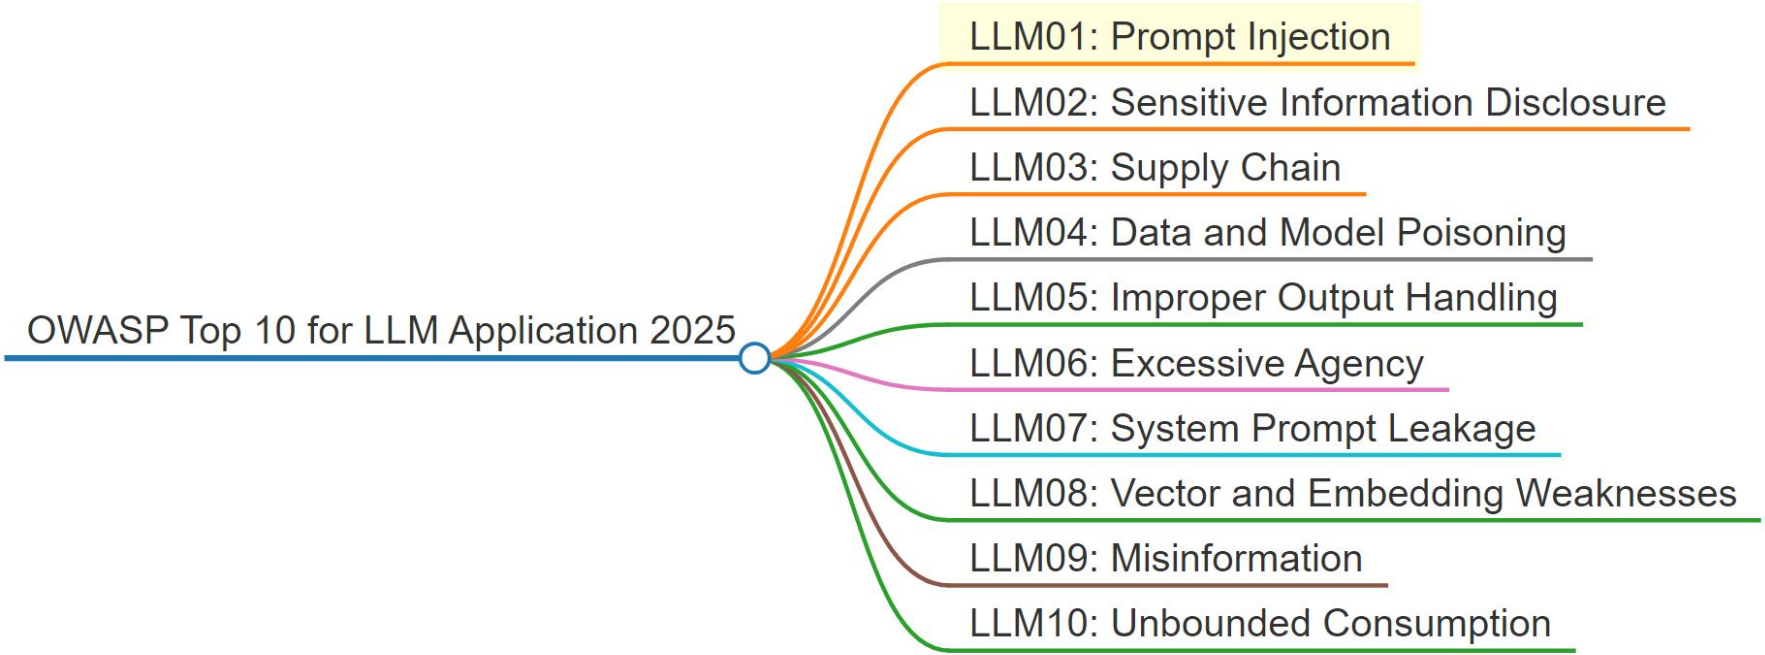

Attack Scenario

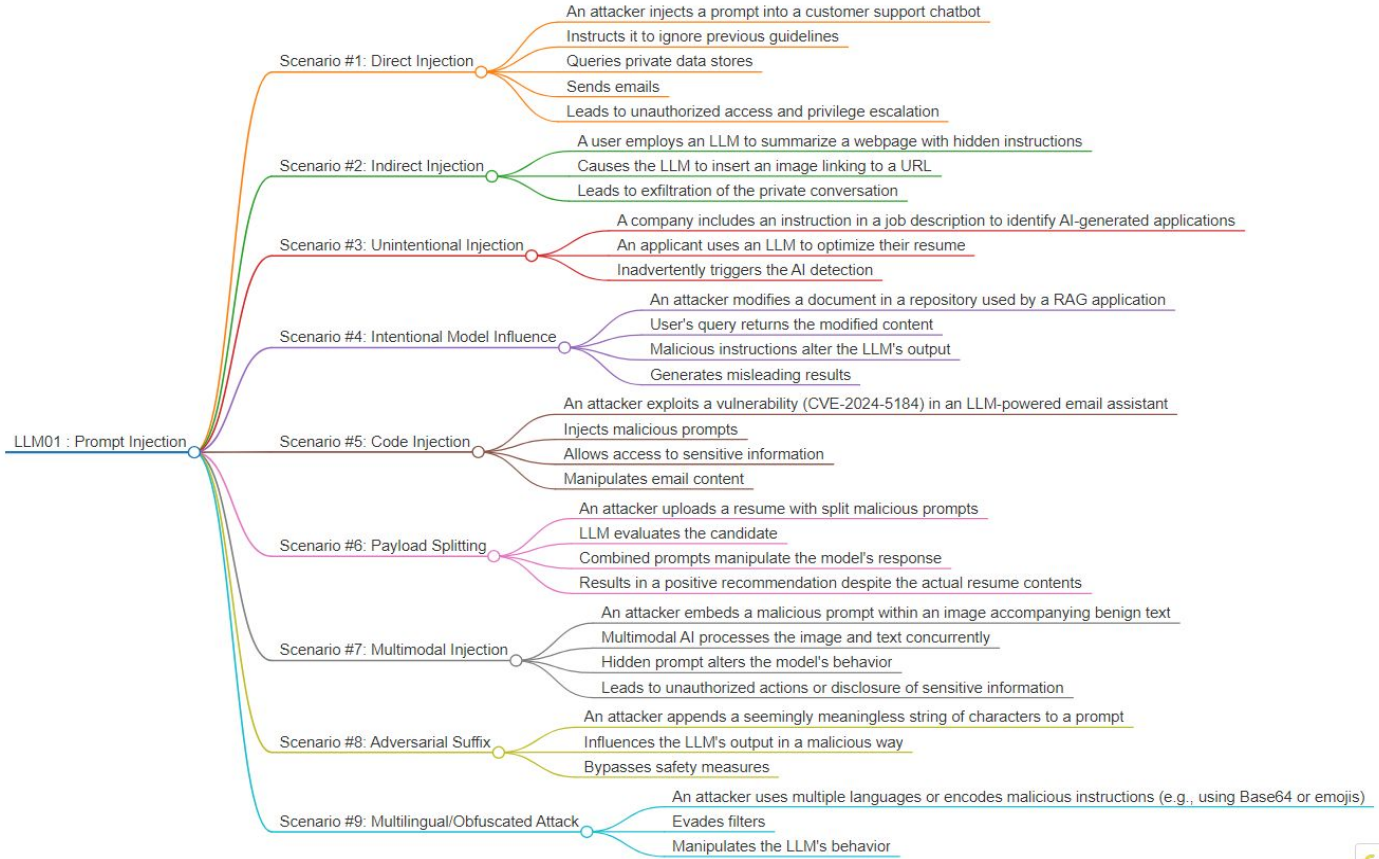

Why Existing Tools leave you vulnerable?

Trad Observability

- System Health
- CPU
- Logs

Trad Security

- Signature based
- Known Attacks

GenAI Gap

- Misses Semantic Attacks
- Novel AI Attacks

Question

Your WAF blocks known SQL injection strings. Your API logs show 200 OK responses.

How do you detect a subtle indirect prompt injection attack embedded within retrieved RAG documents that successfully exfiltrates user data via a seemingly benign LLM response, using only traditional O11y signals (metrics, basic logs, traces)?

Solutions

- **Log Full Context**: Record identifiers (e.g., doc IDs) for data retrieved by RAG, full LLM prompts, and responses.
  Why: Helps analyze the context behind problematic outputs.

- **Contextual Tracing**: Link user queries, retrieved docs, LLM invocations, and responses in traces.
  Why: Pinpoints malicious documents causing bad outputs.

- **Monitor Responses**: Scan LLM outputs for sensitive data patterns (PII, secrets) using DLP tools.
  Why: Detects attacks and data exfiltration directly.

- **Analyze Semantics**: Use heuristic rules or ML models to detect suspicious prompts or anomalies in embeddings.
  Why: Flags unusual inputs or outputs for investigation.

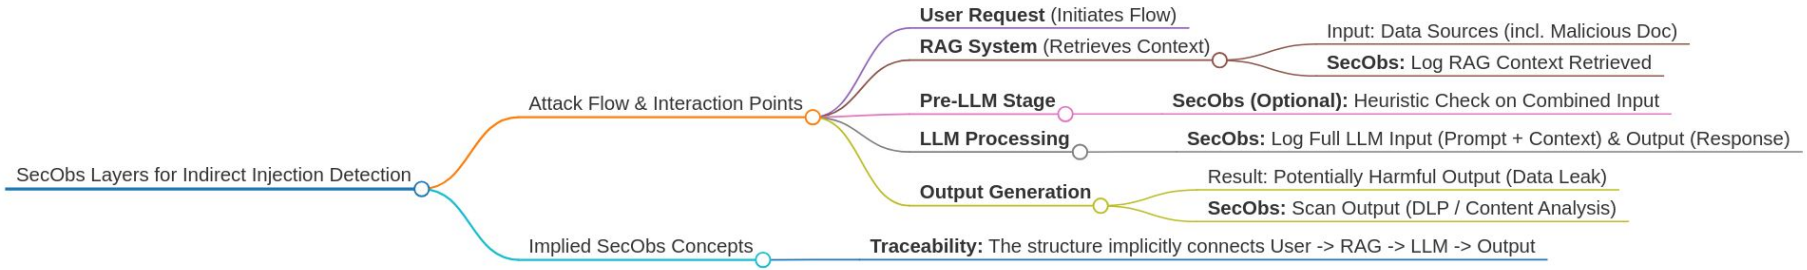
SecObs Layers for Indirect Injection Detection
Attack Flow & Interaction Points
User Request (Initiates Flow)
RAG System (Retrieves Context)
Input: Data Sources (incl. Malicious Doc)
SecObs: Log RAG Context Retrieved
Pre-LLM Stage
SecObs (Optional): Heuristic Check on Combined Input
LLM Processing
SecObs: Log Full LLM Input (Prompt + Context) & Output (Response)
Output Generation
Result: Potentially Harmful Output (Data Leak)
SecObs: Scan Output (DLP / Content Analysis)
Implied SecObs Concepts
Traceability: The structure implicitly connects User -> RAG -> LLM -> Output


@lol_me
Arey kahena kya chahte ho?

Unify Security & Observability for AI

01

Shared Data Plane

Security signals (threat intel, vulnerability scans) enrich O11y data (logs, traces, metrics). O11y data provides context for security alerts.

02

AI-Specific Signals

Monitor prompts, responses, embeddings, token usage, content safety flags as first-class citizens.

03

Behavioral Analysis

Move beyond signatures to detecting anomalous AI behavior

04

Contextual Tracing

Trace requests not just through services, but through model calls, data retrieval, and decision points

Evolving MELT Pillars

METRICS

Prompt/Response Tokens: Cost, performance, DoS detection.

Embedding Drift: Statistical distance (cosine sim) over time – indicates concept shift / potential poisoning.

Content Safety Flags: Rate of harmful content generated (hate speech, PII) / Rate of refusal.

Tool Use Success/Failure Rate: For agentic systems.

Prompt Injection Heuristic Score: Frequency of prompts matching known attack patterns.

1

LOGS

Full Prompt/Response Pairs (Sanitized/Anonymized): For incident analysis, debugging, and retraining. Crucial.

Metadata: Model ID, version, temperature, template used, RAG sources consulted.

Content Moderation Decisions: Why was content flagged/blocked?

2

TRACES

End-to-End Flow: User query -> API Gateway -> Orchestrator -> Vector DB -> LLM(s) -> Output Processing -> User.

Context Propagation: Carry metadata (user ID, session ID, data sources) through the trace.

3

References

1. AI Security Solution Cheat Sheet Q1-2025 - OWASP Top 10 for LLM & Generative AI Security
2. Agentic AI - Threats and Mitigations - OWASP Top 10 for LLM & Generative AI Security
3. OWASP Top 10: LLM & Generative AI Security Risks
4. LLM Applications Cybersecurity and Governance Checklist v1.1 - English - OWASP Top 10 for LLM & Generative AI Security
5. Solutions Landscape - OWASP Top 10 for LLM & Generative AI Security
6. LLMRisks Archive - OWASP Top 10 for LLM & Generative AI Security

Thank you!

Not a Phishing QR
It's my **LinkedIn**
Don't Trust
Always verify

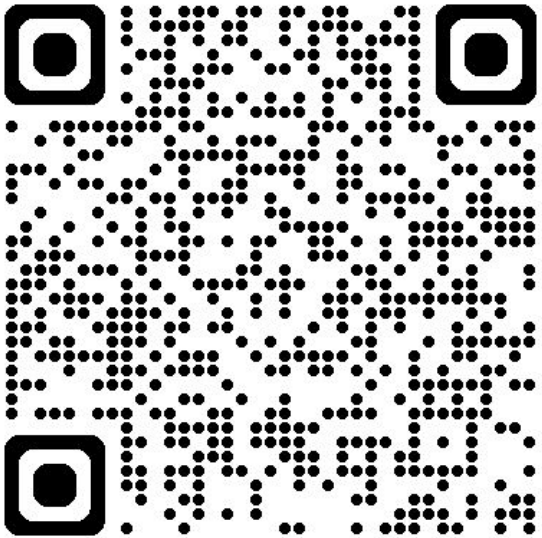

Hewlett Packard Enterprise

docker