DUKE INFORMATION SECURITY · CODE+ 2025

Threat enrichment and blocked-IP feedback with STINGAR.

MICHAEL LEI · RUDHMILA HOQUE · REESE PAGTALUNAN · JAMES WRIGHT

Under STINGAR, we built four systems that turn raw attacker noise into intelligence Duke can *act on,* from can *act on,* from first-touch triage to a self-defending honeypot that learns from every block. block.

SCANNER-FIRST TRIAGE

Scanner Intel Lite

- Right now, every honeypot event lands in Attack Analysis untriaged – an analyst has to manually check reputation, ASN, and behavior for each one before knowing it matters.
- Scanner Intel Lite checks a known-scanner inventory first (free, instant), then honeypot behavior signals, then a ranked external cascade – resolving most events before a single paid API is touched

KEY RESULT

Max 3 paid API calls per event · $0 for known-scanner matches · real-time 4-way classification (benign / suspicious / malicious / unknown). Enriched sessions land in Attack Analysis with no analyst action required.

CATCH-ALL SINKHOLE

Queen Bee

- When Duke blocks a known-bad attacker today, it vanishes into a black black hole and Duke learns nothing
- Queen Bee answers every **address** and every **port** the attacker probes, keeps them on the line, and hands them to STINGAR's honeypots, which which record everything

42k TCP handshakes/sec for one listener

KEY RESULT

One trap answers thousands of scanned addresses. Every session (who, what target, how many bytes) is logged and routed to STINGAR's honeypot, so Duke finally learns from attacks it used to to throw away.

INTELLIGENT HONEYPOTS

Adaptive Honeypot

- When Duke blocks a known-bad attacker today, it vanishes into a black hole and Duke learns nothing
- Seven-protocol honeypot adapts to unseen commands
- 82-rule intent gate and live shared world keep sessions consistent
- Five personas and 22 red-team probes reduce fingerprinting
- Three-VPC GCP deployment sustained 64 cache misses and 30 SSH sessions with 100% success

KEY RESULT

7 protocols · 82 intent rules · 5 coherent personas · 22 red-team probes. 100% success with 0 fallbacks at 64 concurrent adaptive requests (5.75 s p95); 100% at 30 concurrent SSH sessions (347 ms command p95).

ENRICHING MALWARE, IPS & PLAYBOOKS

Threat Enrichment

Attack sessions used to rot unread in an index — the only output was the attacker's IP, a signal that dies when they rotate addresses. Threat Enrichment folds every session into a durable intelligence graph of IPs, command playbooks, malware, and client fingerprints, then layers on reputation, behavior analysis, and LLM-written explanations to separate real threats from scanner noise — closing the loop from bait to intelligence.

KEY RESULT

4 entity types · 279-tag taxonomy · STIX/TAXII export to partner tools · live on Duke attack traffic. Federation matches malware and playbook hashes across universities — campaigns flagged days before IP reputation catches up.

ACKNOWLEDGEMENTS

We thank the Duke Code+ program for the opportunity, our project lead Hugh Thomas for guidance and technical expertise, and Alex Merck for deep insight into the Duke network. Special thanks to Duke Network Architect Will Brockelsby for mentorship on Queen Bee, and to Microsoft for their generous support.

THE HIVE, DEFENDED