RAYLEE HAWKINS

Detection Engineer | Security Automation | Detection-as-Code
Gadsden, AL • raylee@hawkinsops.com • hawkinsops.com • github.com/HawkinsOps • linkedin.com/in/raylee-hawkins

SUMMARY

Detection engineer who thinks in data planes. Built a production-grade SOC from zero computer experience in eight months by treating the telemetry pipeline as the product, not the SIEM. Operate a Wazuh → Cribl Stream → Splunk HEC pipeline end-to-end, validated GREEN with 39.7% measured volume reduction and 100% high-severity preservation (LSASS rule 100312: 7/7). Treat detection engineering as code: every rule versioned, every change tested in CI, every disposition logged to an atomic ledger. Prior Tier 1 automotive supervision (30+ operators, 24/7 shifts, IATF 16949) established the change-control and audit-trail thinking I apply to every pipeline stage I ship.

CORE OPERATING EVIDENCE

Cases Processed	Detections	CI Assertions	Pipeline Validation
324,074 0 reconciliation mismatches	**210 total** 103 Sigma | 79 Splunk | 28 Wazuh	**208 / 208** Sigma gate green	**Wazuh→Cribl→Splunk** 39.7% reduction | 100% preserved

METHODOLOGY & APPROACH

- **Reduction with preservation.** Validated a Wazuh → Cribl Stream → Splunk HEC pipeline with 39.7% measured volume reduction at the Cribl layer while preserving 100% of level-10+ detections. Volume is a cost lever. Signal is not. Operated the pipeline through seven revisions to green with full field-extraction validation on the Splunk side.
- **Measure before you tune.** Fixed agent queue saturation (3.3x buffer, 2x drain rate) before classifying 151K alerts, because suppression decisions built on dropped-event data are unreliable by design.
- **Preserve detection when in doubt.** Classified 110K of 147K alerts as noise but deliberately did not suppress the largest category (~83K Git/bash subprocess churn) because those binaries have dual-use attack potential. Deferred pending path-based re-tuning.
- **Scope over blanket exclusion.** All 28 suppression rules scoped to the single noisy host. Same binary remains fully alerted on every other agent where its execution would be anomalous.
- **Reversibility as a design constraint.** Suppressions deployed as level-0 child rules, never as parent-rule edits. Each guard is independently revertible. Parent detection logic and tuning logic remain decoupled.
- **Minimum viable change under fire.** Patched a TOCTOU race condition in the live pipeline at 505,836-file queue depth with 16 net new lines across 2 files. Zero behavioral changes. Zero data loss. Verified against 1,056 vanishing files with 0 unhandled exceptions.
- **Multi-model AI consensus.** Every non-trivial architectural decision triangulated across Claude, ChatGPT, and Gemini before execution. Never single-source, always verify, same discipline I apply to detection signals.
- **Atomic ledger and reconciliation.** 324,074 cases processed, 0 reconciliation mismatches across all dispositions, even through the race-condition incident. Integrity is a property of the pipeline, not a report generated about it.

TECHNICAL SCOPE

Telemetry Pipeline & SIEM: Cribl Stream 4.13 (Wazuh→Cribl→Splunk validated GREEN, 39.7% volume reduction with 100% high-severity preservation, seven revisions to green, full field-extraction validation), Splunk Enterprise (SPL, HEC, correlation, threat hunt), Wazuh 4.14 (manager + 10-agent fleet, API, active response, FIM, SCA, path-based Layer B suppression), OpenSearch 7.10, Sysmon, Windows Security Event Log.

Detection Engineering: Sigma (103 rules, UUIDv4-enforced), Splunk SPL (79 rules), Wazuh (28 rule blocks), 10 IR playbooks, MITRE ATT&CK mapping (60 technique IDs / 12 tactics), false-positive tuning, hypothesis-driven threat hunting, detection-as-code CI gate (208/208 assertions).

AI & Automation: Custom MCP server authoring in TypeScript (Wazuh 4.14 + OpenSearch 7.10), LLM-driven triage workflow design, Python pipeline engineering (35 scripts, TOCTOU-safe state, atomic ledger), SOAR-pattern orchestration, GitHub Actions CI/CD for detection validation.

Security Engineering: Windows audit policy hardening (27 subcategories Success+Failure), Microsoft Security Baseline exceedance (12 settings, zero regressions across 60), CommandLine logging, Security log 20MB→1GB, CIS, NIST 800-53.

Infrastructure: Proxmox VE (10-VM cluster, 72 cores, 2TB RAM, ~96TB ZFS), Ubuntu/Debian, Windows Server, Docker, Tailscale mesh, PowerShell 7, Bash.

Frameworks: MITRE ATT&CK, NIST CSF, NIST 800-53, CIS, PCI DSS, HIPAA, financial-sector controls awareness.

EXPERIENCE

Detection Engineer / Security Automation — HawkinsOps (Independent) — hawkinsops.com ***Sept 2025 – Present***

- Operate **SignalFoundry**, a 7-stage detection pipeline (35 Python scripts) processing 324,074 cases at 88% auto-close with 0 reconciliation mismatches. Every disposition atomically ledgered; every stage idempotent; state machine TOCTOU-safe.
- Built a **Wazuh** → **Cribl** → **Splunk** reduction pipeline end-to-end on a Cribl Stream 4.13 homelab node from bare VM. Validated GREEN with 39.7% measured reduction of level-10-and-above traffic while preserving 100% of high-severity detections (LSASS rule 100312: 7/7 preserved).
- Authored and maintain **210 detections** (103 Sigma, 79 Splunk SPL, 28 Wazuh) plus 10 IR playbooks across 60 MITRE ATT&CK technique IDs / 12 tactics. Deployed a Sigma detection-as-code CI gate (208/208 assertions) enforcing UUIDv4 IDs, required-field presence, and global uniqueness before merge.
- Authored a custom **Wazuh MCP server** in TypeScript against Wazuh 4.14.3 + OpenSearch 7.10.2, exposing SIEM search, agent state, and rule metadata as first-class tools for LLM-driven triage workflows.
- Diagnosed and patched a **TOCTOU race condition** in the live queue at 505,836-file depth. 4-site fix across poll-alerts.py and triage.py, 16 net new lines, 0 behavioral changes. Verified against 1,056 vanishing files during sort with 0 unhandled exceptions.
- Ran a 24-hour overnight **tuning sprint**: classified 151,384 process-creation alerts into infrastructure noise, signal, and ambiguous. Fixed queue saturation (42 warnings → 0), deployed 28 scoped suppressions as level-0 child rules, preserved Git/bash category pending path-based Layer B tuning.
- Hardened Windows audit policy: 27 subcategories moved from No Auditing to Success+Failure, CommandLine logging enabled, Security log expanded 20MB → 1GB. 12 settings verified exceeding Microsoft Security Baseline with zero regressions across 60 subcategories.
- Operate a 10-agent Wazuh fleet with role-based groups (windows_workstations, linux_servers, infrastructure, honeypot), group-specific FIM and SCA, and active response for automated containment with documented blast-radius scope.

AI Model Evaluator — Outlier AI (Remote) ***Dec 2025 – Present***

- Evaluate frontier AI model outputs for factual accuracy, reasoning quality, and safety across cybersecurity, detection logic, and infrastructure architecture. Produce structured assessments under strict rubric and SLA; flag hallucination, logic gaps, and policy failures to improve model alignment.

Production Supervisor — Unipres Alabama — Gadsden, AL ***Mar 2025 – Dec 2025***

- Supervised 30+ operators across 24/7 rotating 12-hour shifts in Tier 1 automotive stamping (Honda, Toyota, Nissan, Subaru) under IATF 16949 / TISAX / ISO 9001. Change-control, audit-trail, and failure-mode habits translate directly to regulated-sector detection engineering.

Production Team Lead — Fehrer Automotive North America — Gadsden, AL ***2024 – Mar 2025***

- Promoted from second-shift operator to third-shift Team Lead within ~5 months based on performance and reliability.
- Led A, B, and I production lines totaling roughly 5,000+ parts per shift, balancing staffing, workload by operator skill, shift continuity, downtime response, and throughput under production pressure.

SELECTED CASE STUDIES (AVAILABLE ON REQUEST)

- **Wazuh → Cribl → Splunk Reduction Pipeline.** End-to-end homelab build: seven pipeline revisions to green, 39.7% volume reduction, 100% high-severity preservation, full field-extraction validation.
- **AutoSOC Pipeline Race Condition.** Post-incident engineering analysis: 11-section writeup of a TOCTOU defect at 505K queue depth, four-site hotfix, verification method, rollback procedure, and lessons learned. Integrity evidence: 321,351 = sum of dispositions, 0 mismatches.
- **Overnight Process-Telemetry Tuning Sprint.** Detection-engineering decision record: 151K-alert classification, queue-saturation root cause, 28-rule scoped suppression deployment, and explicit reasoning for why the largest noise category was deliberately preserved.

CERTIFICATIONS & EDUCATION

- **CompTIA Security+** — In progress, WIOA/ITA funded.
- **IATF 16949 / TISAX / ISO 9001 / SQF Ed. 9** — Working knowledge from Tier 1 automotive supervision.
- High school education completed.