PUBLIC-SAFE WHITEPAPER

SynapticSOC

Building a Self-Hosted SOC Operating Model with Open-Source Security Tools

From perimeter sensor to audit record: integrating IDS alerts, firewall telemetry, passive network visibility, and endpoint detection into a correlated triage pipeline with controlled automation, access separation, evidence handling, retention policy, and audit-ready controls.

Prepared by Md Bashir Ahmed
Project SOC / SynapticSOC
Version 1.2 · June 2026

This document is designed for public portfolio and GitHub use. Sensitive infrastructure details, raw evidence, credentials, and private validation artifacts are intentionally excluded.

Contents

This whitepaper is organized as a staged SOC maturity narrative, moving from architecture and telemetry into correlation, accountability, controls, limitations, and roadmap.

Section 1: Executive Summary

SynapticSOC is an operational, self-hosted Security Operations Center built with open-source security tools and operated under real-world resource constraints. It is not presented as a theoretical reference diagram or a one-off lab exercise. It is a working SOC operating model that demonstrates how perimeter monitoring, firewall telemetry, IDS alerting, passive network visibility, endpoint detection, multi-source correlation, controlled response, analyst accountability, access separation, evidence handling, retention policy, and audit-ready controls can be brought together into a practical security operations pipeline.

The project was shaped through an extended integration process rather than a single clean deployment. Perimeter monitoring, network telemetry, endpoint visibility, log forwarding, parsing, enrichment, triage, correlation, automation, and reporting each had to be connected, tested, adjusted, and stabilized under real infrastructure constraints. Earlier approaches exposed limitations in the log pipeline, forwarding model, indexing path, and correlation workflow before the final architecture became reliable enough to support repeatable triage and documented control validation.

At the center of SynapticSOC is the detection and triage pipeline. IDS alerts provide the intake trigger. Firewall telemetry adds perimeter context, including pass/block visibility where available. Passive network monitoring contributes flow and protocol context across the monitored segment. Endpoint monitoring adds host-level detection and security context. External IP enrichment supports reputation-aware review. Log aggregation and correlation bring those signals together for structured analyst triage. A controlled automation layer then handles SOAR intake, analyst notification, acknowledgment, and audit writeback without permitting destructive automated response.

The project is built on open-source tools spanning perimeter security, network intrusion detection, passive network analysis, endpoint monitoring, log management, indexing, workflow automation, dashboarding, and analyst communication. Those tools provide the foundation. The SynapticSOC operating model defines how the components are integrated, correlated, controlled, documented, and governed.

SynapticSOC evolved through a deliberate maturity path:

Visibility → Detection → Correlation → Controlled Response → Analyst Acknowledgment → RBAC → Evidence Handling → Retention Documentation → Audit-Ready Control Pack

As of the current version, the project demonstrates:

- Firewall, network IDS, passive network visibility, and endpoint telemetry collection
- Multi-source correlation, enrichment, and triage across network and endpoint evidence
- Controlled SOAR intake with destructive automation intentionally disabled
- Telegram-based analyst acknowledgment and accountability recording
- Role-based access control across Graylog, Wazuh, Grafana, n8n, and supporting backend components
- A documented evidence handling model covering classification, lifecycle, analyst responsibilities, chain-of-custody considerations, and public-safe redaction
- A documented retention model covering SOC review, workflow traceability, audit records, endpoint evidence, management reporting, and public/private evidence separation
- Management and compliance-oriented dashboards separate from the primary SOC triage layer
- A final integrated control pack consolidating RBAC, evidence handling, retention documentation, validation artifacts, and public-safe summaries

• Phase 1.5 Zeek visibility hardening, introducing immediate and deferred Zeek lookup, no-match classification, and structured audit writeback of visibility outcome fields

SynapticSOC does not claim certified compliance with any regulatory framework. It does not claim legal-grade forensic chain of custody, immutable archival storage, enterprise SOC feature parity, commercial SOC replacement capability, or managed detection and response maturity. Those are deliberate non-claims.

What the project does demonstrate is practical SOC maturity under real constraints. It shows that a carefully designed self-hosted SOC built with open-source security tools can move beyond basic alert collection into a documented operating model with correlation, controlled triage, analyst accountability, role separation, evidence discipline, retention awareness, management visibility, and audit-ready control documentation.

This whitepaper documents the architecture, component roles, operational pipeline, control model, limitations, and roadmap of SynapticSOC as it stands today.

Section 2: Project Objective

The objective of SynapticSOC is to demonstrate how a SOC built with open-source security tools can be designed, operated, documented, and governed with practical controls across visibility, detection, triage, access separation, evidence handling, retention awareness, and audit review.

The project was not designed to imitate a commercial SOC platform or present open-source tools as a complete replacement for enterprise security operations. It was designed to answer a more practical question:

How much meaningful SOC capability can be built, operated, and documented using self-hosted open-source tooling, limited infrastructure, and disciplined engineering decisions?

SynapticSOC answers that question by implementing a working security operations pipeline supported by firewall telemetry, network IDS alerts, Zeek network visibility, Wazuh endpoint and security monitoring, Graylog triage and evidence review, Grafana management dashboards, n8n-controlled automation, and Telegram-based analyst acknowledgment.

The project also recognizes the limits of what it is. SynapticSOC does not provide certified compliance, legal forensic custody, immutable evidence storage, enterprise-grade high availability, or full commercial SOC replacement capability. Instead, it provides a practical and transparent model for building SOC maturity in stages, while clearly documenting what has been implemented, what has been validated, and what remains outside the project scope.

A major purpose of SynapticSOC is professional documentation. The project is intended to be understandable and reviewable by technical readers, hiring managers, cybersecurity practitioners, certification reviewers, and open-source contributors without requiring direct access to the underlying infrastructure. The public website, GitHub repository, screenshots, validation artifacts, control summaries, and this whitepaper together form the public record of the project.

The project remains active and has a continuing roadmap for hardening, integration, runbook development, backup validation, case-management expansion, and future compliance mapping. However, SynapticSOC has crossed an important threshold: it is no longer only a detection lab or tool integration exercise. It is now a documented SOC operating model.

Section 3: Project Evolution and Maturity Path

SynapticSOC did not begin as a compliance-oriented project. It began as a practical engineering effort: build a working SOC with open-source security tooling, establish real telemetry flow, and understand how visibility, detection, correlation, triage, and response decision-making function at the infrastructure level.

That order was intentional. Governance controls are only meaningful when they protect and support a functioning operational environment. A retention policy has limited value without reliable log ingestion. An RBAC model has little purpose without analyst-facing systems to protect. Evidence handling cannot be credible unless evidence is actually being generated, reviewed, and recorded. SynapticSOC was therefore built from operational capability first, then matured into governance, control, and audit-readiness.

The project evolved through the following maturity path:

Visibility → Detection → Correlation → Controlled Response → Analyst Acknowledgment → RBAC → Evidence Handling → Retention Documentation → Audit-Ready Control Pack

Stage 1 — Visibility

The first stage established security visibility across the environment. Firewall telemetry, network IDS alerts, Zeek network visibility, and Wazuh endpoint monitoring were introduced to ensure that meaningful security events could be generated, collected, and reviewed.

The objective at this stage was not advanced automation or compliance reporting. The objective was to make the environment observable.

Stage 2 — Detection

The second stage focused on detection capability. Snort, Suricata, and Wazuh detection sources were configured and validated. Log ingestion into Graylog, OpenSearch, and the Wazuh indexing layer was stabilized.

At this point, SynapticSOC moved beyond raw telemetry collection into structured security detection output.

Stage 3 — Correlation

The third stage introduced multi-source correlation. Firewall events, IDS alerts, Zeek network visibility, and Wazuh endpoint/security signals were brought together into Graylog streams, dashboards, and triage views.

This allowed events to be reviewed in context rather than as isolated alerts. Correlation became the bridge between detection and operational triage.

Stage 4 — Controlled Response

The fourth stage introduced controlled SOAR intake through n8n. Automation was deliberately designed with restraint. SynapticSOC does not rely on unrestricted auto-remediation or destructive automated response actions. Response decisions remain controlled, reviewable, and dependent on analyst approval.

The design position is simple: automation should support analyst judgment, not replace it.

Stage 5 — Analyst Acknowledgment

The fifth stage introduced a Telegram-based analyst acknowledgment workflow. Alerts can be sent to an analyst-facing communication channel, and analyst acknowledgment is recorded through the approved workflow.

Acknowledgment records are stored for audit and review. This stage introduced a critical accountability layer by distinguishing between an alert being generated and an alert being seen, acknowledged, and handled by an analyst.

Stage 6 — RBAC and Access Control

The sixth stage introduced role-based access control across Graylog, Wazuh, Grafana, n8n, and supporting backend components. Each platform was assigned a defined operational role and a corresponding access model.

Graylog became the primary SOC triage and evidence-review layer. Wazuh provided endpoint, security, and compliance visibility. Grafana served management and compliance-facing dashboard users. n8n remained restricted to owner-level automation control. OpenSearch and Wazuh Indexer were documented as backend infrastructure rather than analyst-facing systems.

Negative validation was also performed to confirm that read-only users could not access or administer restricted index-management and administrative functions.

Stage 7 — Evidence Handling

The seventh stage formalized the evidence handling model. SynapticSOC documented evidence sources, evidence classification, evidence lifecycle, analyst evidence responsibilities, chain-of-custody considerations, redaction rules, and the separation of internal evidence from public-safe material.

This stage did not claim certified forensic handling or legal-grade chain of custody. Instead, it established practical evidence discipline suitable for SOC review, accountability, documentation, and public portfolio presentation.

Stage 8 — Retention Documentation

The eighth stage introduced retention policy documentation across the SOC stack. Retention scope, tool-level considerations, validation evidence, and limitations were documented for Graylog, Grafana, n8n, Wazuh, and supporting evidence folders.

The retention model supports SOC review, audit readiness, analyst accountability, workflow traceability, endpoint and security evidence review, management reporting, and public/private evidence separation.

SynapticSOC does not claim immutable archival storage or legal-grade retention. Its retention model is designed for practical operational review and audit-oriented documentation.

Stage 9 — Audit-Ready Control Pack

The ninth stage consolidated the completed control phases into a final integrated control pack. This included the RBAC control model, evidence handling documentation, retention policy documentation, validation artifacts, source folder references, an integrated control matrix, and public-safe summaries.

The control pack gives SynapticSOC a structured reference point for external review. It allows the project to be explained, assessed, and presented without requiring direct access to the underlying infrastructure.

Phase 1.5 — Zeek Visibility Hardening and Production Cutover

Before finalizing the whitepaper and progressing into Phase 2, SynapticSOC introduced a targeted hardening update to address a known limitation in the Zeek correlation layer. In production, a Zeek no-match result can occur for reasons unrelated to actual absence of traffic: timing delay between packet capture and log indexing, SPAN or mirroring coverage gaps, Docker or private address scoping outside the monitored LAN range, or general sensor visibility boundaries. Treating a no-match as definitive proof that traffic did not occur was operationally misleading.

Phase 1.5 introduced a two-pass Zeek lookup model into SOC/01, now versioned as SOC/01 Alert Intake and Correlation v1.7 — Zeek Deferred Lookup. When an immediate Zeek lookup returns no match, the workflow waits 120 seconds and performs a deferred second-pass lookup anchored to the original event time rather than the current processing time. The final visibility result is classified as one of three outcomes: zeek_seen_immediate, zeek_seen_deferred, or zeek_no_match_after_delay.

Phase 1.5 visibility fields are written into Graylog through SOAR audit GELF writeback, making them searchable in Graylog audit streams and available for management-level reporting through Grafana. Production cutover was completed and validated. Both Snort IDS alerts and Wazuh alerts now route through SOC/01 v1.7. The previous workflow version was unpublished but retained as a rollback reference.

Phase 1.5 is a mitigation and documentation control. It reduces timing-related no-match cases and documents remaining visibility gaps. It does not replace the future need for full LAN-side sensor coverage, stronger passive monitoring, or the broader Zeek improvements planned for Phase 2.

What the Maturity Path Demonstrates

Each stage of SynapticSOC was built on the stage before it. The project did not add compliance language to an incomplete tool stack. It first established operational capability, then documented the controls required to govern that capability responsibly.

The result is a SOC project that can be reviewed not only for what it detects, but for how it is operated, controlled, documented, and governed.

Section 4: Architecture Overview

SynapticSOC is deployed on self-hosted infrastructure within a controlled local network environment. The architecture is designed around a clear separation of responsibilities: each component has a defined operational role, and access to each component is scoped according to that role.

The environment is built around an internal SOC network where firewall telemetry, IDS alerts, endpoint events, network visibility data, automation records, and dashboard outputs are collected and reviewed. The design does not depend on a commercial SIEM, managed detection service, or cloud-hosted SOC platform. Instead, it uses self-hosted open-source components arranged into functional layers.

The architecture can be understood as five major layers:

- Network boundary and telemetry collection
- Endpoint and host visibility
- Log management, correlation, and triage
- Controlled automation and analyst acknowledgment
- Management, compliance, and audit visibility

Together, these layers form the operational foundation of SynapticSOC.

Layer 1 — Network Boundary and Telemetry Collection

pfSense sits at the network boundary and provides firewall enforcement, routing, and firewall log telemetry. Firewall events are forwarded into the SOC pipeline to support network-level visibility and review.

Snort and Suricata provide network IDS alerting. These tools generate detection signals from traffic traversing the monitored network boundary. Their role is to identify suspicious or policy-relevant network activity and forward alert telemetry for downstream review.

Zeek provides passive network visibility. It contributes connection records, protocol metadata, and network evidence that can corroborate or contextualize IDS alerts. Zeek is especially important because it helps move the SOC beyond alert-only visibility by showing how traffic behaved at the protocol and session level.

Together, pfSense, Snort, Suricata, and Zeek form the network-facing telemetry layer. These components generate signals; they do not make final triage or response decisions.

Layer 2 — Endpoint and Host Visibility

Wazuh provides endpoint monitoring, security event collection, vulnerability visibility, and compliance-oriented host telemetry. Wazuh agents collect host-level events and forward them to the Wazuh manager for processing.

Within SynapticSOC, Wazuh is used as a visibility and review platform, not as an unrestricted response execution layer. SOC analyst and compliance auditor access is read-only. This allows endpoint and security data to be reviewed without granting unnecessary administrative control.

Wazuh Active Response remains disabled in the current model. This is a deliberate design decision. Automated remediation or potentially disruptive response actions should not be enabled until the enrichment and decision chain is sufficiently validated and trusted.

Layer 3 — Log Management, Correlation, and Triage

Graylog is the primary SOC triage and evidence-review layer. It receives telemetry from across the stack, applies parsing and enrichment logic, organizes events into streams, and presents analyst-facing views for investigation and review.

Graylog is positioned as the main operational console because it can bring together firewall telemetry, IDS alerts, Zeek network evidence, Wazuh events, automation records, and acknowledgment evidence in one review layer. This allows analysts to review events in context rather than switching between isolated tools.

Graylog also stores audit-oriented records generated by the automation and acknowledgment workflows. This makes it a central point not only for detection review, but also for accountability evidence.

OpenSearch and the Wazuh Indexer serve as backend indexing and storage infrastructure. They are not treated as analyst-facing platforms. Direct backend access is restricted to administrative or service-level use, because allowing broad index-level access would bypass the access boundaries established in Graylog and Wazuh.

Layer 4 — Controlled Automation and Analyst Acknowledgment

n8n serves as the controlled automation backend for SynapticSOC. It supports SOAR-style intake and acknowledgment processing, but it is not exposed as an analyst-facing platform.

Two production workflow roles are active in the current architecture:

- Controlled SOAR intake receives correlated alert data, applies workflow logic, and routes notifications toward the analyst acknowledgment process.
- Analyst acknowledgment intake receives analyst responses through the approved channel and records acknowledgment events into Graylog audit streams.

Following the Phase 1.5 update, SOC/01 v1.7 also performs an immediate Zeek lookup and, where no match is found, a deferred second-pass lookup after a 120-second wait anchored to the original event time. Visibility outcome fields are written back into Graylog through GELF audit writeback. This extends the automation layer's role from controlled SOAR intake and notification into active Zeek visibility-gap classification and documentation.

This design keeps automation controlled and reviewable. Analysts do not directly modify automation workflows. They interact with the automation layer only through the approved acknowledgment channel and through the audit records produced by the workflow.

Telegram functions as the analyst notification and acknowledgment channel. It provides a lightweight human interface for alert acknowledgment while n8n handles the workflow processing and Graylog records the resulting audit evidence.

Layer 5 — Management, Compliance, and Audit Visibility

Grafana provides the management, CISO, and compliance-oriented dashboard layer. It is used for review, reporting, and audit visibility rather than active investigation.

Grafana dashboards provide visibility into areas such as retention status, audit records, acknowledgment activity, workflow traceability, and pipeline health. This separates management and compliance review from the analyst triage layer.

Non-administrative Grafana users are restricted to Viewer-level access. This allows management or compliance stakeholders to review SOC status without gaining access to raw triage streams, backend systems, or automation controls.

Component Role Summary

Component	Architectural Layer	Primary Role	Analyst-Facing Role
pfSense	Network boundary	Firewall enforcement and firewall telemetry	No
Snort	Network detection	Network IDS alerting	No
Suricata	Network detection	Network IDS alerting	No
Zeek	Network visibility	Passive network visibility and protocol evidence	No
Wazuh	Endpoint and host visibility	Endpoint monitoring, security events, vulnerability and compliance visibility	Read-only
Graylog	Log management and triage	Correlation, triage, streams, evidence review, and audit records	Primary SOC analyst layer
OpenSearch / Wazuh Indexer	Backend storage	Indexed event storage and retrieval infrastructure	No — admin/service only
n8n	Automation	Controlled SOAR intake and acknowledgment workflow processing	No — owner/admin only
Telegram	Analyst acknowledgment	Alert notification and acknowledgment channel	Acknowledgment only
Grafana	Management and compliance	Audit dashboards, retention visibility, and management reporting	Viewer-only for non-admin users

Key Architectural Decisions

Several architectural decisions are important to the design of SynapticSOC.

Graylog was selected as the primary SOC triage layer because it can correlate and present events from multiple sources across the environment. Wazuh remains the primary endpoint and security visibility platform, but Graylog is where multi-source investigation and evidence review are centered.

n8n was isolated from normal analyst access. The automation layer handles controlled SOAR intake and acknowledgment processing, so it must not be modifiable by the same users who consume its output. This protects the integrity of both the automation workflow and the audit records it produces.

Following the Phase 1.5 Zeek visibility hardening update, n8n is also responsible for immediate and deferred Zeek lookup execution, visibility-gap classification, and GELF writeback of Phase 1.5 audit fields into Graylog. This extends its role beyond notification and acknowledgment processing into active sensor visibility documentation.

OpenSearch and the Wazuh Indexer are treated as backend infrastructure. Analysts do not require direct backend index access to perform triage or review. Keeping backend access restricted helps preserve access boundaries, evidence handling discipline, and administrative separation.

Grafana was separated from the analyst triage layer. Management and compliance stakeholders need dashboard-level visibility, not raw investigative access. This separation allows audit and reporting needs to be met without expanding unnecessary access to operational SOC tooling.

Wazuh Active Response remains disabled. This conservative position reduces the risk of automated response actions being triggered from insufficiently validated enrichment or correlation logic. SynapticSOC prioritizes controlled response and analyst judgment over premature auto-remediation.

Overall, the architecture is designed to support practical SOC operation while maintaining clear boundaries between telemetry generation, detection, triage, automation, backend administration, and management review.

Section 5: Tooling and Role Separation

A defining characteristic of SynapticSOC's current maturity is that each tool in the stack has a defined operational role, and access to each tool is scoped according to that role. This was not the starting point of the project. It became part of the SOC operating model through deliberate implementation and validation during the RBAC and access-control phase.

This distinction is important. A SOC is not only defined by the tools it deploys, but also by how those tools are governed. If every user can access every platform with administrative privileges, then detection, evidence handling, and audit records become harder to trust. SynapticSOC was therefore structured around role separation, least privilege, and controlled access boundaries.

Role Separation Principle

SynapticSOC separates four major access concerns:

Operational review covers the tools used by SOC analysts to investigate alerts, review triage streams, and examine endpoint and network evidence. This function is centered on Graylog, with Wazuh available in a read-only review capacity.

Management and compliance visibility covers the dashboards used by CISO-level, management, or compliance-oriented reviewers to assess SOC health, retention status, acknowledgment activity, workflow traceability, and audit records. This function is provided through Grafana with Viewer-level access.

Automation control covers the workflows that handle controlled SOAR intake, analyst notification, and acknowledgment recording. This function is provided by n8n, which remains restricted to owner-level administrative access with two-factor authentication enforced.

Backend infrastructure covers indexed storage, data retrieval, and supporting service accounts. This function is provided by OpenSearch and the Wazuh Indexer, which are treated as backend infrastructure rather than analyst-facing systems.

The purpose of this model is to prevent unnecessary privilege overlap. Analysts do not need automation administration. Management viewers do not need raw triage access. Automation workflows do not require unrestricted human sharing. Backend indexing systems do not need to be exposed as investigation consoles.

Graylog — Primary SOC Triage and Evidence Layer

Graylog is the primary platform for SOC triage, correlation review, stream-based investigation, and evidence review. It receives telemetry and workflow records from multiple parts of the SOC stack and presents them in analyst-facing views.

Access in Graylog is scoped through stream and dashboard sharing rather than broad platform-wide permissions.

The Graylog access model is as follows:

- The administrator account retains full system administration rights.
- SOC analyst access is scoped to assigned streams and dashboards.
- SOC viewer access is limited to general dashboard visibility, with restricted audit and evidence streams excluded where appropriate.
- The n8n SOAR API account operates as a service-level integration account for workflow communication.

- Sidecar and internal system accounts are treated as service accounts, not human analyst users.
- System configuration, input management, pipeline editing, and index set administration remain restricted to administrative access.

The SOC analyst role in Graylog is intentionally narrow. Analysts can review assigned streams, use triage dashboards, and examine correlated events. They are not granted unnecessary configuration or system administration privileges.

Wazuh — Endpoint and Security Visibility Layer

Wazuh provides endpoint monitoring, security event review, vulnerability visibility, and compliance-oriented host telemetry. In SynapticSOC, Wazuh is used as an endpoint and security visibility platform, not as an unrestricted response execution layer.

Two read-only human roles were implemented and validated:

- wazuh.soc.analyst for SOC analyst endpoint and event review
- wazuh.compliance.auditor for compliance-oriented review of security events and vulnerability data

Both accounts are mapped to the Wazuh API read-only role through the project-specific API role mapping and to the Wazuh Indexer read-only role through the project-specific indexer role.

The Wazuh Indexer read-only role uses the following permission model:

Permission Area	Access Granted
Cluster permission	cluster_composite_ops_ro
Index permission	Read
Index scope	All required Wazuh indices
Tenant access	Global tenant, read-only

The run_as parameter was enabled in the Wazuh dashboard configuration to enforce user-level permission scoping within the dashboard. A configuration backup was created before this change. Only the Wazuh dashboard service was restarted. The Wazuh manager, Filebeat, Logstash, agents, rules, decoders, and passwords were not modified.

Negative validation was performed against the read-only Wazuh accounts. Both accounts received security exceptions when attempting restricted index-management or administrative actions. This validation confirms that administrative privileges were not unintentionally granted to read-only users.

Wazuh Active Response remains disabled in the current SynapticSOC model. Automated remediation actions are not permitted from the Wazuh layer until enrichment, correlation, and response-decision trustworthiness are validated further.

Grafana — Management and Compliance Dashboard Layer

Grafana provides the management, CISO, and compliance-oriented dashboard layer. It is not used as the primary analyst investigation console.

Two non-administrative Viewer accounts were implemented and validated:

- ciso.viewer for executive and CISO-level dashboard access
- compliance.auditor for compliance-oriented review of audit and retention dashboards

Both accounts are assigned the Grafana Viewer role. They can view dashboards, but they cannot edit dashboards, manage data sources, administer users, or modify Grafana configuration.

The Grafana administrator account remains the only full administrative account. No additional administrative user accounts were created.

Grafana dashboard content is focused on management and review functions, including retention visibility, audit records, acknowledgment activity, workflow traceability, and pipeline health. This makes Grafana appropriate for management and compliance review without turning it into an operational triage console.

n8n — Restricted Automation Backend

n8n is the controlled automation backend for SynapticSOC. It is explicitly not an analyst-facing platform.

No SOC analyst, senior analyst, or SOC manager role has direct access to n8n. The owner account is the only human account with n8n access, and two-factor authentication is enforced on that account.

Two production workflow roles are active:

- SOC/01 — Controlled SOAR Intake
- SOC/02 — Analyst Acknowledgment Intake

The Controlled SOAR Intake workflow receives correlated alert data, applies workflow logic, and routes notifications toward the analyst acknowledgment process. The Analyst Acknowledgment Intake workflow receives analyst responses through Telegram and records acknowledgment events into Graylog audit streams.

Credentials stored in n8n were reviewed by name only during the access-control phase. Secret values were not opened or captured as validation evidence. Managed credentials include API tokens and service credentials required for AbuseIPDB, Graylog, n8n workflow integration, and Telegram communication.

n8n execution history is treated as workflow accountability evidence. It provides a traceable record of workflow execution, processed events, and acknowledgment handling. This evidence is reviewed through Graylog audit streams and Grafana dashboards rather than by granting analysts direct n8n access.

The project's automation position is intentionally conservative: n8n assists analyst review, but it does not replace analyst judgment. No automated destructive actions are permitted. Response decisions remain controlled and manually approved.

OpenSearch / Wazuh Indexer — Backend Infrastructure

OpenSearch and the Wazuh Indexer are backend indexed storage and retrieval infrastructure. They are not analyst-facing platforms and were not configured as general investigation consoles.

This decision is deliberate. Analyst access at the index layer would bypass the access controls and evidence boundaries established in Graylog and Wazuh. Direct backend access would weaken the separation between investigation, evidence review, and administrative control.

No additional human user accounts were created in OpenSearch or the Wazuh Indexer during the access-control phase. No service accounts were modified. No passwords were rotated. Reserved and internal service accounts, including administrator and platform service accounts, were documented and left unchanged.

Backend access remains limited to administrative and service-level use.

Consolidated Access Model

Platform	Account / Role	Access Level	Purpose
Graylog	Administrator	Full administrative access	System administration
Graylog	SOC analyst	Stream and dashboard scoped	Triage and evidence review
Graylog	SOC viewer	Limited dashboard access	General SOC visibility
Graylog	n8n SOAR API	Service-level integration	Workflow communication
Wazuh	wazuh.soc.analyst	Read-only	Endpoint and event review
Wazuh	wazuh.compliance.auditor	Read-only	Compliance-oriented security review
Grafana	Administrator	Full administrative access	Dashboard platform administration
Grafana	ciso.viewer	Viewer	Executive and CISO dashboard review
Grafana	compliance.auditor	Viewer	Audit and retention dashboard review
n8n	Owner	Full administrative access with 2FA	Automation workflow administration
OpenSearch / Wazuh Indexer	Admin and service accounts	Administrative or service-level access	Backend indexing and storage

Why Role Separation Matters

Role separation is not a cosmetic control. It directly affects the security, reliability, and credibility of the SOC.

First, it limits account blast radius. An analyst account scoped to streams and dashboards cannot modify backend indexes, change automation logic, rotate service credentials, or administer system configuration.

Second, it improves the credibility of audit evidence. When access is scoped, audit records are easier to interpret. Administrative actions, analyst review, workflow execution, and management visibility can be separated instead of blending together under shared or overly privileged accounts.

Third, it demonstrates operational discipline. A SOC that can show a clear access model, validated read-only roles, restricted automation access, and backend administrative separation is more mature than a tool stack where all users operate with broad privileges.

In SynapticSOC, tooling and role separation are therefore part of the control model itself. The SOC is designed not only to collect and review security events, but also to control who can view, modify, automate, administer, and report on those events.

Section 6: Detection and Visibility Pipeline

SynapticSOC's detection and visibility pipeline spans four primary source types: network boundary telemetry, passive network visibility, network intrusion detection, and endpoint and host monitoring. These sources feed into a central log management and correlation layer where events are parsed, normalized, enriched, and made available for analyst review.

The pipeline was not designed to maximize alert volume. It was designed to produce meaningful, reviewable security events with enough context for a human analyst to make an informed triage decision.

SynapticSOC therefore treats telemetry, detection, enrichment, and response as separate stages. A firewall log is not automatically an incident. An IDS alert is not automatically proof of compromise. Reputation enrichment is not treated as a response verdict. Each signal contributes context, but the triage decision remains with the analyst.

Network Boundary Telemetry — pfSense

pfSense sits at the network edge and provides firewall enforcement, routing, and firewall log telemetry. Its filterlog records capture firewall rule matches, connection attempts, blocked traffic, permitted traffic, and interface-level activity. These records are forwarded into Graylog, where they are parsed and used as the baseline network boundary record.

In SynapticSOC, pfSense telemetry is not treated as a standalone detection source. It is treated as corroborating network context. When an IDS alert fires or a Wazuh event is reviewed, the associated pfSense record helps establish whether related traffic crossed the network boundary, what direction it moved in, and whether the firewall permitted or blocked it.

The environment uses explicit internal network classification logic to distinguish internal RFC1918 addresses from genuine public IP addresses. This is important for enrichment accuracy. External reputation checks, such as AbuseIPDB lookups, are applied only to public IP addresses and not to internal addresses, Docker bridge addresses, or local service IPs that may appear in logs.

This prevents internal infrastructure from being incorrectly treated as internet-originated activity and reduces the risk of misleading enrichment results.

Passive Network Visibility — Zeek

Zeek provides passive network visibility across the monitored environment. Unlike a signature-based IDS, Zeek does not primarily produce binary alert verdicts. Instead, it generates structured protocol and session telemetry, including connection records, DNS activity, HTTP transactions, SSL/TLS metadata, file transfer observations, and other network behavior evidence.

Zeek's role in SynapticSOC is corroboration and context. An IDS alert from Snort or Suricata may indicate that traffic matched a known suspicious pattern. Zeek helps explain what the traffic looked like: which hosts communicated, which ports and protocols were involved, how long the connection lasted, and whether related DNS, HTTP, or TLS activity occurred around the same timeframe.

To support this corroboration function, SynapticSOC uses a match-quality model for Zeek correlation rather than a simple seen/not-seen verdict. This model distinguishes between stronger and weaker forms of corroboration, such as exact IP and port matches, IP-only matches, subnet-level matches, partial matches, and no observed match.

This distinction matters because not all supporting evidence has the same confidence level. A precise Zeek match around the same event window is more meaningful than a broad or partial network observation. The match-quality model helps prevent weak contextual matches from being treated as confirmed corroboration.

Phase 1.5 — Two-Pass Zeek Lookup Model

Prior to Phase 1.5, a Zeek no-match result was treated as a simple absence of observed activity. In practice, a no-match can occur for several reasons independent of whether the traffic actually happened: timing delay between packet capture and log indexing, SPAN or mirroring coverage gaps, Docker or private address scoping outside the LAN sensor range, or general sensor visibility boundaries.

Phase 1.5 introduced a two-pass lookup model into SOC/01 v1.7. When an immediate Zeek lookup returns no match, the workflow waits 120 seconds and performs a deferred second-pass lookup anchored to the original event time rather than the current processing time. This separates genuine timing-delay cases from persistent no-match conditions.

Final Zeek visibility outcomes are classified as one of three values:

- zeek_seen_immediate — traffic was observed in the immediate lookup pass
- zeek_seen_deferred — traffic was observed only after the deferred second-pass lookup
- zeek_no_match_after_delay — no match was found after both passes; the no_match_reason field documents the most likely explanation

A Zeek no-match is now treated as a visibility status requiring interpretation, not automatic proof that the traffic did not occur. Phase 1.5 reduces timing-related no-match cases and documents remaining visibility gaps. It does not resolve SPAN or mirroring design limitations, QUIC or HTTP/3 traffic opacity, or the broader passive monitoring coverage improvements planned for Phase 2.

Network Intrusion Detection — Snort and Suricata

Snort and Suricata operate as the network IDS layer. They generate signature-based alert telemetry from traffic observed at the monitored network boundary. Their alerts are ingested into Graylog, where they are normalized, correlated, and presented for triage.

Network IDS alerts are treated as detection signals, not confirmed incidents. A signature match indicates that observed traffic resembled known malicious, suspicious, or policy-relevant behavior. It does not automatically prove malicious intent, successful exploitation, or endpoint compromise.

This separation between alert generation and analyst triage is a core design principle in SynapticSOC. Automated response based only on raw IDS alerts would create an unacceptable risk of disruptive false-positive actions. The controlled response model therefore requires enrichment, correlation, and analyst review before response decisions are made.

Snort and Suricata provide valuable detection coverage, but their output becomes operationally useful only when combined with firewall telemetry, Zeek context, endpoint evidence, and analyst judgment.

Endpoint and Host Monitoring — Wazuh

Wazuh provides endpoint monitoring, security event collection, vulnerability visibility, and compliance-oriented host telemetry. Wazuh agents installed on monitored hosts forward events to the Wazuh manager, where they are processed and made available through the Wazuh indexing and dashboard layers.

Within SynapticSOC, Wazuh provides host-level context that network tools cannot supply on their own. This includes endpoint security events, agent activity, vulnerability findings, compliance-related observations, Windows event data, file integrity monitoring output, and other host-side evidence.

The active triage scope focuses on higher-severity Wazuh events that warrant analyst attention. Lower-severity informational events remain available in the indexing layer but are not treated as primary triage triggers unless they become relevant during investigation.

Several Wazuh fields require additional handling during correlation and enrichment. Fields such as MITRE mappings, rule groups, Windows event details, and compliance tags may require parsing or normalization before they can be used reliably in downstream review, dashboards, or scoring logic.

Event attribution is also important. For example, the Wazuh manager itself may generate events that must be distinguished from agent-originated endpoint detections. This prevents management-plane events from being incorrectly interpreted as detections from a monitored host.

Wazuh Active Response remains disabled in the current SynapticSOC model. This is deliberate. Host-level response actions such as file quarantine, process termination, or network blocking introduce operational risk if triggered from insufficiently validated detections. SynapticSOC therefore prioritizes read-only endpoint visibility and analyst-led response decisions until the enrichment and correlation chain is mature enough to justify automated action.

Log Management and Correlation — Graylog

Graylog receives telemetry from the major detection and visibility sources: pfSense, Zeek, Snort, Suricata, and Wazuh. It applies parsing, normalization, enrichment, stream routing, and dashboarding to produce structured security events for review.

Graylog is the central correlation layer because it can bring together network boundary activity, IDS alerts, passive network evidence, endpoint events, automation records, and analyst acknowledgment evidence. This allows events to be reviewed in context rather than as isolated alerts from separate tools.

Enrichment in Graylog includes reputation context for genuine public IP addresses. Reputation data is provided as analyst context, not as an independent response trigger. Enrichment fields are marked so that they do not automatically drive response scoring or destructive action. This protects the workflow from treating third-party reputation data as a final verdict.

Firewall correlation uses time-bounded event matching to associate pfSense records with related detection events. The goal is to identify firewall activity that occurred within a relevant event window, rather than relying on broad relative searches that may produce incidental or misleading matches.

Graylog stream organization also separates detection records, triage records, audit records, and analyst acknowledgment evidence. This separation helps preserve a clean boundary between raw detection output, analyst-facing investigation data, and accountability records produced by the workflow.

Backend Indexed Storage — OpenSearch and Wazuh Indexer

OpenSearch and the Wazuh Indexer provide the backend indexed storage layer for the SOC pipeline. They support event storage, retrieval, indexing, and retention for the telemetry and security data produced by the environment.

Within SynapticSOC, backend storage is not treated as an analyst-facing investigation layer. Analysts review events through Graylog and Wazuh according to their assigned roles. Direct backend access remains restricted to administrative and service-level use.

Retention is differentiated by data type. Raw telemetry and high-volume operational indices are retained for shorter periods, while audit and accountability records are retained for longer review windows according to the retention policy.

This distinction reflects the different values of each data category. Raw telemetry is useful for investigation and short-term review, but audit and acknowledgment records represent evidence of analyst interaction, workflow execution, and accountability. Those records require longer availability for management review, retention validation, and compliance-oriented reporting.

SynapticSOC does not claim immutable archival storage or legal-grade retention. The indexed storage layer supports practical SOC review, audit readiness, and operational traceability within the documented constraints of the project.

Pipeline Flow Summary

```
pfSense firewall telemetry
Snort and Suricata IDS alerts
Zeek network visibility
Wazuh endpoint and host events
↓
Graylog parsing, normalization, enrichment, and stream routing
↓
Triage streams, audit streams, dashboards, and analyst review
↓
OpenSearch / Wazuh Indexer backend storage and retention
```

The pipeline produces structured, correlated, and enriched security events. It does not produce automatic incident verdicts or unrestricted automated responses.

Every event that reaches the analyst layer has been parsed, classified, and enriched, but the final triage decision remains with the analyst. This is the operating principle behind SynapticSOC's detection and visibility pipeline: automation can assist review, but it should not replace accountable human judgment.

Section 7: Correlation and Controlled SOAR Workflow

Detection without correlation produces noise. Correlation without a controlled response path produces dashboards that may not result in action. SynapticSOC's SOAR workflow was designed to close the gap between a correlated detection event and a documented analyst response, while deliberately avoiding automated response actions that the current maturity of the project does not justify.

The design principle is straightforward: automation should assist analyst judgment, not replace it. SynapticSOC uses automation for intake, enrichment support, notification, acknowledgment processing, and audit record creation. It does not use automation to perform destructive or disruptive actions without human review.

Correlation Layer

Before an alert reaches the analyst acknowledgment workflow, it passes through Graylog's correlation and enrichment layer.

Correlation in SynapticSOC is multi-source. A detection event is not treated as a complete operational event if supporting context is available elsewhere in the stack. The pipeline attempts to associate detection events with:

- Firewall records from the same event window, showing whether related traffic crossed the network boundary and whether it was permitted or blocked
- Zeek connection and protocol records, showing what the associated network activity looked like at the session and protocol level
- Wazuh endpoint or host events, where relevant, showing whether corresponding host-side activity was observed
- Reputation enrichment for genuine public IP addresses, where appropriate

The result is a structured event with layered context. Instead of presenting only a signature match, the analyst receives a signal supported by network boundary context, passive network visibility, endpoint evidence where available, and enrichment data.

Reputation enrichment, such as AbuseIPDB context, is applied only to genuine public IP addresses. Internal addresses, local service addresses, Docker bridge ranges, and other private infrastructure addresses are excluded from external reputation checks. This prevents internal infrastructure from being misclassified as internet-originated activity and reduces the risk of misleading enrichment results.

Enrichment data is provided as analyst context only. It does not independently trigger automated response decisions. A poor reputation result may increase analyst concern, but it does not automatically block an IP address, isolate a host, or modify firewall policy.

This distinction protects the integrity of the correlation layer. If IP classification or enrichment logic is wrong, automated responses based on that logic would be unsafe. SynapticSOC therefore treats enrichment as supporting evidence, not as a verdict.

Controlled SOAR Intake — SOC/01

Once a correlated event is ready for analyst attention, it enters the SOC/01 — Controlled SOAR Intake workflow in n8n.

SOC/01 receives structured alert data from Graylog through a protected webhook. The workflow validates the incoming request before processing it, then applies workflow-level logic to classify, format, and route the event toward the analyst acknowledgment channel.

SOC/01 does not perform autonomous response actions. It does not block source IP addresses, isolate endpoints, terminate processes, modify firewall rules, or close incidents. Its function is controlled intake, formatting, notification, and audit-oriented record handling.

The workflow reflects the project's conservative automation position:

- Automated destructive response is not permitted.
- Manual approval is required before any response decision.
- Enrichment supports analyst review but does not drive autonomous action.
- Workflow execution must remain traceable.

SOC/01 v1.7 — Zeek Deferred Lookup

Following the Phase 1.5 update, the production workflow is versioned as SOC/01 Alert Intake and Correlation v1.7 — Zeek Deferred Lookup. Both Snort IDS alerts and Wazuh alerts now route through this workflow after production cutover. The previous SOC/01 Controlled SOAR Intake workflow has been unpublished but is retained as a rollback reference.

SOC/01 v1.7 introduces a two-pass Zeek lookup into the correlation layer:

Step 1 — Immediate Zeek lookup using the original event-time anchored Graylog query

Step 2 — If no match is returned, the workflow waits 120 seconds

Step 3 — Deferred second-pass Zeek lookup anchored to the original event time

Step 4 — Final visibility outcome classified and written to the audit record

The three Phase 1.5 Zeek visibility outcomes are:

zeek_seen_immediate — matched on the first pass

zeek_seen_deferred — matched only after the deferred pass

zeek_no_match_after_delay — no match after both passes; no-match reason documented

Phase 1.5 fields are included in Telegram analyst notifications and written into Graylog through SOAR audit GELF writeback. A production replay and live validation confirmed the immediate match path with workflow_version 1.7 and zeek_phase15_result zeek_seen_immediate.

This is not an aspirational policy. It is the current operating position of SynapticSOC. Automation is allowed to assist the workflow, but the response decision remains human-controlled.

Analyst Acknowledgment Workflow — SOC/02

SOC/02 — Analyst Acknowledgment Intake handles the human response side of the workflow.

When SOC/01 routes a correlated alert for review, the analyst receives a structured notification through Telegram. The notification provides enough context for rapid initial triage, including the event source, classification, relevant network or endpoint details, enrichment summary, and correlation quality.

The analyst reviews the notification and records an acknowledgment response through the approved Telegram channel. SOC/02 receives the response, processes it, and writes an acknowledgment record into the Graylog audit stream.

The acknowledgment record captures the key accountability details:

- The alert or event that was acknowledged
- The time the acknowledgment was recorded
- The analyst response
- Workflow execution metadata
- Audit stream routing information

This record is the accountability artifact. It answers a question that raw detection data cannot answer on its own: was a human analyst actually made aware of the event, and was a response recorded?

Audit Records and Evidence Trail

The combination of SOC/01 and SOC/02 creates a traceable evidence trail from correlated detection to analyst acknowledgment.

```
Correlated detection event
↓
SOC/01 — Controlled SOAR Intake
↓
Analyst notification via Telegram
↓
Analyst acknowledgment response
↓
SOC/02 — Analyst Acknowledgment Intake
↓
Graylog audit stream record
↓
Grafana acknowledgment and audit dashboards
```

Each step in this chain produces evidence. Graylog stores the acknowledgment and audit records. Grafana provides management-level visibility into acknowledgment activity, workflow traceability, and audit trends. n8n execution history provides workflow-level traceability for the automation backend.

These records are treated differently from raw telemetry. Raw firewall, IDS, network, and endpoint events support operational investigation. Acknowledgment and audit records represent evidence of analyst interaction and workflow execution. For that reason, they are retained according to the longer audit-retention model rather than treated as short-lived operational telemetry.

This distinction supports review, accountability, and compliance-oriented reporting without claiming immutable archival storage or legal-grade evidence retention.

What the SOAR Workflow Deliberately Does Not Do

The current SOAR workflow has clear boundaries. This is important because overstating automation capability is a common failure in SOC documentation.

SOC/01 and SOC/02 do not:

- Automatically block source IP addresses
- Automatically isolate or quarantine endpoints
- Automatically create, modify, or remove pfSense firewall rules
- Automatically terminate processes or services on monitored hosts

- Automatically close, resolve, or escalate cases in a case management platform
- Make autonomous response decisions based on enrichment data alone
- Treat IDS alerts, reputation scores, or partial correlation matches as final incident verdicts

These capabilities are not implemented, and they are not presented as implemented.

The project's current position is that automated destructive response requires a higher level of enrichment reliability, detection confidence, validation coverage, and rollback planning than the current phase claims. Conservative automation is therefore a design choice, not a weakness.

Case management integration is identified as a future roadmap item. TheHive or an equivalent case management layer may later be used for case creation, escalation, evidence attachment, and incident lifecycle management. Those functions are not claimed as current capabilities in this version of the SOC.

Why Controlled SOAR Matters

The controlled SOAR model reflects a deliberate operating philosophy. In SynapticSOC's current maturity stage, the cost of a false-positive automated response may be higher than the cost of requiring manual acknowledgment.

Blocking a legitimate source, isolating a healthy endpoint, terminating a valid process, or modifying firewall policy based on incomplete evidence can create operational disruption. For that reason, SynapticSOC prioritizes controlled response, analyst accountability, and evidence integrity over automated speed.

The analyst acknowledgment workflow is therefore central to the SOAR model. It creates a documented bridge between detection and human review. An event that was detected, correlated, enriched, notified, acknowledged, and recorded is more operationally meaningful than an event that was automatically actioned without a documented human decision point.

SynapticSOC's approach does not attempt to present automation as a replacement for security operations judgment. It uses automation to make analyst review faster, more structured, and more accountable.

Section 8: Analyst Acknowledgment and Accountability

Most SOC documentation focuses on what was detected. SynapticSOC also focuses on what was reviewed.

This distinction matters. A security tool can generate an alert, enrich it, store it, and display it on a dashboard without any human analyst ever reviewing it. Detection evidence proves that a tool observed something. It does not prove that a human saw it, acknowledged it, or recorded a response.

SynapticSOC addresses this gap through a dedicated analyst acknowledgment workflow. The purpose of this workflow is to create a structured record showing that a correlated detection event was delivered to an analyst, reviewed at the notification level, and acknowledged through an approved channel.

This turns the SOC record from a passive collection of tool-generated alerts into a more complete operational record: what the SOC detected, what was brought to analyst attention, and what response was recorded.

The Accountability Gap

A SOC that can show detection records but not acknowledgment records has an accountability gap. It can demonstrate that its tools were active, but it cannot easily demonstrate that alert review actually occurred.

This gap affects three areas.

First, it affects operational review. Without acknowledgment records, there is no reliable distinction between an alert that was reviewed and dismissed as a false positive and an alert that was never reviewed at all.

Second, it affects audit readiness. Detection logs alone show that telemetry and alerting exist. They do not show whether a triage process was followed. Audit-oriented review requires evidence of process, not only evidence of tooling.

Third, it affects project credibility. SynapticSOC is intended to demonstrate a working SOC operating model, not only a collection of integrated tools. Analyst acknowledgment provides evidence that the project includes a human review layer on top of detection and correlation.

How Acknowledgment Works in SynapticSOC

The acknowledgment process begins after a detection event has been correlated and prepared for analyst review.

A correlated event enters the controlled SOAR intake workflow, where it is formatted and routed to the analyst notification channel. The analyst receives a structured Telegram notification containing the key context required for initial triage review, such as event source, classification, enrichment summary, correlation quality, and relevant network or endpoint details.

The analyst then records a response through the Telegram acknowledgment channel. This response is processed by the acknowledgment intake workflow and written into the Graylog audit stream as a structured acknowledgment record.

This design keeps acknowledgment separate from normal dashboard browsing. Viewing a dashboard and acknowledging an alert are not treated as the same action. Acknowledgment is a deliberate, recorded workflow step.

Acknowledgment Record Contents

Each acknowledgment record written to the Graylog audit stream captures the core accountability details required for later review:

- The alert or event that was acknowledged
- The timestamp of acknowledgment
- The analyst response
- Workflow execution metadata
- Audit stream routing information

These records are treated as accountability evidence rather than ordinary operational telemetry. They are retained according to the longer audit-retention model so that analyst interaction and workflow activity remain available for management review, validation, and compliance-oriented reporting.

What Acknowledgment Proves

The acknowledgment model is intentionally scoped. It proves a specific human review step occurred; it does not claim more than that.

An acknowledgment record demonstrates that:

- A structured notification was delivered through the approved analyst channel
- A human response was recorded through that channel
- The response was processed by the acknowledgment workflow
- The acknowledgment was stored at a specific point in time
- The record can be traced back to a correlated detection event

This is valuable because it documents the transition from machine-generated detection to human workflow participation.

What Acknowledgment Does Not Prove

An acknowledgment record does not prove that a full forensic investigation was completed. It does not prove that the event was classified correctly in every case. It does not prove that no related activity was missed elsewhere in the environment. It also does not prove that the response decision was optimal.

SynapticSOC does not present acknowledgment as a substitute for incident response, case management, or experienced analyst judgment. It presents acknowledgment as a documented control showing that a human review step occurred.

That distinction is important. The control is honest because it is limited. It records awareness and response, not investigative perfection.

Accountability Visibility

Acknowledgment records are surfaced through two visibility layers.

Graylog hosts the analyst acknowledgment and audit streams. These streams allow acknowledgment records to be searched, filtered, reviewed, and correlated with related SOC activity. This is the operational evidence layer.

Grafana provides management and compliance-oriented visibility over acknowledgment activity. Dashboards summarize acknowledgment trends, workflow activity, response timing, and audit health without requiring management users to access the full analyst triage layer.

This two-layer model follows the broader role separation principle used throughout SynapticSOC. Detailed records remain in the operational SOC layer. Summarized visibility is presented through the management and compliance layer.

Accountability as a SOC Control

In SynapticSOC, analyst acknowledgment is not an informal workflow habit. It is a control.

The acknowledgment model allows the SOC to answer questions that detection tooling alone cannot answer:

- Was this alert acknowledged?
- When was it acknowledged?
- What response was recorded?
- Is acknowledgment activity visible over time?
- Can analyst interaction be reviewed separately from raw detection output?

These questions matter because a SOC should be able to demonstrate not only that it generated alerts, but that alerts entered a review process.

SynapticSOC's acknowledgment workflow provides that evidence. It creates a documented bridge between detection, analyst awareness, and audit review. In practical terms, this is the difference between a SOC that only detects activity and a SOC that can demonstrate how detected activity was brought into a human accountability process.

Section 9: RBAC and Access Control Model

Access control is where SOC governance becomes enforceable. Detection rules, correlation logic, dashboards, evidence streams, and SOAR workflows can all be well designed, but their integrity depends on who can access, modify, administer, or bypass them.

If every user has broad administrative access, the separation between investigation, management review, automation control, and backend administration collapses. Evidence can be changed. Audit records can be altered. Automation workflows can be modified. Service credentials can be exposed. In that model, access control becomes informal rather than operational.

SynapticSOC's RBAC model was implemented to prevent this. The model defines who can access each SOC component, what level of access they receive, and which administrative boundaries they are not allowed to cross. It also validates those boundaries through positive and negative testing.

The objective is not to claim enterprise identity governance. The objective is to demonstrate practical least privilege and role separation across a self-hosted SOC built with open-source security tools.

RBAC Design Principles

The SynapticSOC access-control model is built around three principles.

Least privilege means each user or service account receives only the permissions required for its defined function. Accounts are not granted broader access simply because broader access would be easier to configure or more convenient during administration.

Role separation means operational review, management visibility, automation control, and backend administration are treated as separate functions. A SOC analyst does not need automation administration. A compliance viewer does not need triage-stream editing rights. An automation service account does not need human dashboard privileges. A backend administrator role is not treated as an investigation role.

Validated boundaries means access controls are not considered complete just because they were configured. SynapticSOC treats validation as part of the control. Users must be able to access what their role requires, and they must be denied access to functions outside that role.

This last point is important. Positive validation proves access works. Negative validation proves restriction works.

Graylog Access Control

Graylog is the primary SOC triage and evidence-review layer. Access is scoped through stream and dashboard sharing rather than broad platform-wide access.

The Graylog access model separates administrative control from analyst review:

Role / Account Type	Access Purpose	Boundary
Administrator	System administration, inputs, pipelines, users, index sets	Full administrative access
SOC analyst	Triage streams, assigned dashboards, correlated event review	No system configuration or index administration
SOC viewer	Limited dashboard visibility	Restricted streams excluded where appropriate

n8n SOAR API account	Workflow integration and event communication	Service-level access only
Sidecar / internal service accounts	System service functions	Not treated as human user roles

This model prevents general users from receiving unnecessary access to system configuration, pipeline logic, input management, or index administration.

A key design decision is that Graylog stream access is granted deliberately. Evidence streams, audit streams, and acknowledgment records are not treated as universally visible simply because they exist in the same platform. This helps preserve the boundary between general SOC visibility, analyst triage, and audit evidence review.

Wazuh Access Control

Wazuh required more granular RBAC configuration because the platform spans multiple layers: dashboard access, API access, and indexer permissions.

Two read-only human roles were implemented:

- wazuh.soc.analyst for endpoint and security event review
- wazuh.compliance.auditor for compliance-oriented security and vulnerability review

Both accounts were mapped at two levels.

At the API layer, the accounts were mapped to a Wazuh read-only role. This restricts API-level interaction to read operations.

At the indexer layer, the accounts were mapped to a project-specific Wazuh read-only role with the following permission model:

Permission Area	Access Granted
Cluster permission	Read-only composite operations
Index permission	Read
Index scope	Required Wazuh indices
Tenant access	Global tenant, read-only

The Wazuh dashboard was configured to enforce user-level permission scoping through the dashboard layer. A configuration backup was created before the change was applied. Only the Wazuh dashboard service was restarted. The Wazuh manager, Filebeat, Logstash, agents, detection rules, decoders, and account passwords were not modified.

This is important because the RBAC phase was not used as an excuse to reconfigure unrelated SOC services. The change was limited, controlled, and scoped to access enforcement.

Negative validation was performed against both Wazuh read-only accounts. When those accounts attempted restricted index-management or administrative actions, access was denied through security exceptions. This confirms that the accounts were not accidentally granted administrative indexer privileges.

For the Wazuh layer, negative validation is especially important. A user who appears read-only in the dashboard but can perform backend index-management actions would represent a failed control. SynapticSOC validated that this was not the case.

Grafana Access Control

Grafana provides management, CISO, and compliance-oriented dashboard visibility. It is not used as the primary analyst investigation console.

Two non-administrative Viewer accounts were implemented:

- ciso.viewer for executive and CISO-level dashboard review
- compliance.auditor for compliance-oriented review of audit and retention dashboards

Both accounts were assigned Viewer-level access.

Action	Viewer Accounts
View assigned dashboards	Permitted
Edit dashboards	Denied
Manage data sources	Denied
Administer users	Denied
Modify Grafana configuration	Denied

The Grafana administrator account retains full administrative access. No additional administrative accounts were created for management or compliance users.

This supports the broader role-separation model. Management and compliance users can review dashboard outputs without receiving the ability to alter dashboards, change data sources, modify users, or interfere with the reporting layer.

n8n Access Control

n8n is the controlled automation backend. It has the strictest human access boundary in the stack because it holds workflow logic, service credentials, execution history, and the automation processes that generate acknowledgment evidence.

The n8n access model is intentionally simple:

Role	Access
Owner	Full administrative access with two-factor authentication
SOC analyst	No direct n8n access
Senior SOC analyst	No direct n8n access
SOC manager	No direct n8n access
Management / compliance viewer	No direct n8n access

This restriction is architectural, not merely procedural.

If analysts could modify n8n workflows, they could alter the logic that produces their own acknowledgment records. If management viewers could access n8n credentials, the separation between reporting, automation, and administration would weaken. If multiple users had workflow-editing access, it would become harder to trust the integrity of the automation evidence trail.

n8n credentials were reviewed by name only during the access-control phase. Secret values were not opened, captured, or included in validation evidence. Credentials under management include API tokens and service credentials required for enrichment, Graylog integration, workflow communication, and Telegram notification.

The result is a controlled automation layer where workflow administration remains restricted, while workflow outputs remain reviewable through Graylog and Grafana.

OpenSearch and Wazuh Indexer Access Control

OpenSearch and the Wazuh Indexer are treated as backend infrastructure, not analyst-facing platforms.

This decision is deliberate. Direct index-level access could bypass the access boundaries established in Graylog and Wazuh. For example, a user restricted at the Graylog stream level could potentially query underlying data directly if granted backend index access. That would weaken the stream-based and dashboard-based access model.

The backend access-control position is therefore:

- No additional human analyst accounts were created.
- No new broad human access was granted.
- Existing service and reserved accounts were documented.
- Service accounts were not modified unnecessarily.
- Passwords were not rotated as part of the RBAC phase.
- Backend access remains limited to administrative and service-level operations.

This preserves the indexing layer as infrastructure rather than turning it into a parallel investigation console.

Consolidated RBAC Matrix

Platform	Account / Role	Access Level	Validation Status
Graylog	Administrator	Full administrative access	Validated
Graylog	SOC analyst	Stream and dashboard scoped	Validated
Graylog	SOC viewer	Limited dashboard visibility	Validated
Graylog	n8n SOAR API	Service-level integration	Validated
Wazuh	wazuh.soc.analyst	Read-only API and indexer access	Positive and negative validation
Wazuh	wazuh.compliance.auditor	Read-only API and indexer access	Positive and negative validation
Grafana	Administrator	Full administrative access	Validated
Grafana	ciso.viewer	Viewer-only dashboard access	Validated
Grafana	compliance.auditor	Viewer-only dashboard access	Validated
n8n	Owner	Full administrative access with 2FA	Validated
OpenSearch / Wazuh Indexer	Admin and service accounts	Administrative and service-level access	Reviewed and documented

The Value of Negative Validation

Positive validation confirms that a role can perform its intended function. Negative validation confirms that the role cannot exceed that function.

Both are required for a credible RBAC model.

SynapticSOC performed negative validation where it mattered most: the Wazuh read-only accounts. Because Wazuh spans dashboard, API, and indexer layers, a misconfiguration could create a situation where a user appears restricted in one layer but retains administrative capability in another. Negative testing confirmed that read-only users were denied restricted index-management and administrative actions.

This validation provides stronger evidence than screenshots of successful login alone. A successful login proves access. A denied administrative action proves boundary enforcement.

For SynapticSOC, this distinction is part of the control model. RBAC is not considered complete until both allowed and denied behaviors have been tested.

What the RBAC Model Does Not Claim

The SynapticSOC RBAC model is a practical operational control implemented under real constraints. It is not presented as a full enterprise identity governance program or privileged access management platform.

The current model does not include:

- Centralized identity provider integration
- Hardware-backed multi-factor authentication across all tools
- Privileged access workstation controls
- Just-in-time access provisioning
- Automated access review workflows
- Formal joiner, mover, and leaver procedures
- Enterprise-grade privileged session recording
- Organization-wide IAM policy enforcement

These limitations are expected in a single-operator, self-hosted SOC project. They are documented rather than hidden.

Zeek Visibility and Phase 1.5 Scope

Phase 1.5 introduces a two-pass Zeek lookup model and documents no-match conditions with classified outcome fields. It reduces timing-related no-match cases and improves analyst visibility into Zeek corroboration status. It does not resolve all Zeek visibility limitations.

SynapticSOC does not claim:

- Full LAN-side sensor coverage across all traffic paths
- Complete packet visibility or full passive monitoring capability
- QUIC or HTTP/3 traffic inspection
- Resolution of SPAN or mirroring design limitations through Phase 1.5
- Phase 2 passive monitoring infrastructure improvements

Phase 1.5 is a mitigation and documentation control. It makes Zeek no-match conditions visible and interpretable rather than silently absent. The sensor coverage and mirroring improvements required for persistent no-match conditions remain Phase 2 scope.

What SynapticSOC does claim is more specific and more defensible: practical role separation, least-privilege access design, restricted automation administration, backend access separation, read-only reviewer roles, and validated denial of unauthorized administrative actions.

That is the control boundary implemented in this phase, and that is the boundary the whitepaper claims.

Section 10: Evidence Handling Model

Detection creates records. Correlation adds context. Analyst acknowledgment documents review. However, none of those outputs have long-term operational or audit value unless the evidence produced by the SOC is captured, classified, protected, retained, and published with discipline.

SynapticSOC's evidence handling model was created to define how SOC evidence is treated after it is generated. It answers practical questions that appear throughout the SOC lifecycle: where the evidence came from, what it contains, how sensitive it is, who can review it, how it is retained, and what can safely be shown in public documentation.

The model is designed for a single-operator, self-hosted SOC environment. It supports operational review, audit readiness, analyst accountability, and public portfolio presentation. It does not claim certified forensic evidence handling or legal-grade chain of custody.

Evidence Scope

Evidence in SynapticSOC is broader than incident artifacts alone. It includes any record, screenshot, export, dashboard, workflow execution, or validation artifact used to demonstrate SOC operation, control implementation, or audit readiness.

Evidence sources include:

- Firewall, IDS, endpoint, and network visibility records
- Graylog search results, streams, dashboards, and audit records
- Wazuh alert, vulnerability, and compliance-oriented evidence
- n8n execution records and acknowledgment workflow outputs
- Telegram-based analyst acknowledgment records
- Grafana audit, retention, and management dashboards
- RBAC validation screenshots and access-control evidence
- Control phase documentation, validation checklists, and public-safe summaries

These evidence types serve different purposes. Some support analyst triage. Some support management reviews. Some validate controls. Some are prepared for public release. The evidence handling model applies to all of them.

Evidence Classification

SynapticSOC classifies evidence before it is stored, shared, or published. Classification determines how the evidence may be handled and whether it can be used externally.

Classification	Description	Handling Rule
Internal	Routine SOC documentation, general screenshots, dashboard views, and operational notes that do not expose sensitive infrastructure details	Internal use and project review
Sensitive Internal	Evidence containing internal IP addresses, hostnames, usernames, architecture details, workflow logic, configuration details, or security-relevant system context	Internal use only unless reviewed and redacted

Restricted	Evidence containing credentials, API keys, tokens, webhook paths, authorization headers, passwords, secret values, or material that could compromise the SOC environment	Never published; avoid capturing secret values
Public-Redacted	Evidence reviewed and sanitized for external publication	Approved for GitHub, website, portfolio, or public documentation

The project owner is responsible for classification. Evidence is not treated as public-safe by default. It must be reviewed and intentionally moved into a public-redacted state before external use.

Evidence Lifecycle

SynapticSOC uses a defined evidence lifecycle:

Collection → Normalization → Correlation → Triage → Acknowledgment → Evidence Recording → Review → Retention → Redaction → Publication

Collection occurs when evidence is captured from the live SOC environment. Screenshots, logs, workflow records, and validation artifacts are captured at the time of implementation or validation, not reconstructed later for appearance.

Normalization occurs when raw events are parsed and structured by tools such as Graylog, Wazuh, or supporting pipelines. This makes evidence searchable, comparable, and usable for review.

Correlation links related evidence across tools. A network IDS alert may be supported by firewall telemetry, Zeek network visibility, Wazuh endpoint context, or reputation enrichment.

Triage is the analyst review stage. At this point, evidence is no longer just raw telemetry; it becomes part of an operational security review process.

Acknowledgment records that a human analyst was made aware of the event and responded through the approved workflow.

Evidence recording occurs when workflow outputs, audit records, and acknowledgment events are written into Graylog audit streams or preserved as validation artifacts.

Review allows operational, management, or compliance-oriented stakeholders to examine evidence through the correct access layer.

Retention keeps evidence available according to its value and classification. Raw telemetry and audit records may follow different retention periods.

Redaction prepares evidence for external use by removing sensitive details. Redacted evidence is created as a separate copy. Original evidence is not overwritten.

Publication occurs only after evidence has been reviewed and classified as Public-Redacted.

This lifecycle ensures that evidence is not only generated, but handled in a consistent and explainable way.

Chain-of-Custody Considerations

SynapticSOC maintains a practical operational chain-of-custody model. This means evidence can be traced through the SOC process from source event to review artifact, but it does not mean the project provides certified forensic custody.

The model supports traceability across the following path:

```
Detection source
↓
Correlation and enrichment
↓
Triage review
↓
Analyst acknowledgment
↓
Audit record creation
↓
Retention and review
↓
Redacted publication, where applicable
```

The model provides practical evidence of continuity. It allows the project to show where evidence came from, how it was processed, how it was reviewed, and whether it was prepared for public release.

The original evidence is preserved where possible. Redacted versions are created separately for public use. Superseded or replaced evidence should be archived rather than silently deleted. Evidence gaps should be documented rather than hidden.

However, the model has clear limits. SynapticSOC does not currently provide:

- Cryptographic evidence signing
- Tamper-evident storage
- Immutable evidence archival
- Independent witness verification
- Third-party forensic validation
- Legal-grade chain-of-custody procedures
- Certified evidence management controls

These limitations are documented because they matter. The project claims practical operational evidence handling, not legal forensic custody.

Evidence Handling Responsibilities

The project owner is responsible for evidence handling across all operational roles. In a single-operator SOC, the same person may act as builder, analyst, administrator, reviewer, and publisher. That makes discipline especially important.

Evidence handling responsibilities include:

- Capture evidence at the time of action or validation
- Use clear filenames that describe the system, control, and purpose of the evidence
- Store evidence in the correct phase folder and component subfolder
- Separate internal evidence from public-redacted material
- Avoid capturing credentials, tokens, secrets, or authorization headers
- Preserve original evidence when creating redacted versions
- Review evidence before publication
- Document missing or incomplete evidence instead of omitting it silently

- Archive superseded evidence where appropriate

These responsibilities help ensure that evidence remains understandable after the build phase is complete.

Public and Private Evidence Separation

Public/private separation is one of the most important controls in the evidence handling model. SynapticSOC is a public portfolio project, but the SOC environment itself contains sensitive operational detail.

Internal evidence may contain infrastructure details, private network information, usernames, hostnames, system paths, workflow logic, service ports, dashboard URLs, or security configuration details. Restricted evidence may contain credentials or secrets. Neither category is suitable for direct publication.

Before evidence is moved into a public-redacted folder or used in external documentation, it must be reviewed for:

- API keys, tokens, passwords, or credential fragments
- Authorization headers or webhook paths
- Internal IP addresses or service URLs
- Hostnames, usernames, or identifying system details
- Browser address bars, tabs, or internal service ports
- n8n credential nodes or workflow secrets
- Stack traces, file paths, or error output revealing system internals
- Raw logs exposing unnecessary infrastructure detail

If sensitive content is present, the evidence must be cropped, blurred, redacted, or re-captured in a safer form. The sanitized version becomes the public-redacted copy. The original remains internal.

The public-redacted folder is the only approved source for evidence used in GitHub, the SynapticSOC website, LinkedIn posts, whitepaper screenshots, or portfolio documentation.

Evidence Storage Model

Evidence is stored in a structured local project folder model. Each control phase has its own evidence root, and each root separates documentation, validation artifacts, screenshots, internal evidence, and public-redacted material.

This folder separation supports three goals:

First, it makes evidence easier to review. A reader should be able to locate the artifacts supporting a control phase without searching through unrelated files.

Second, it prevents accidental publication. Raw evidence and public-redacted material are stored separately so that public artifacts are intentionally selected rather than casually copied.

Third, it supports audit-oriented explanations. Folder structure, validation checklists, summaries, and evidence indexes together make the control work reviewable without requiring direct access to the live SOC environment.

Raw internal or restricted evidence is not stored in public repositories or shared as portfolio material. Only reviewed public-redacted evidence is used externally.

Evidence Handling and Audit Readiness

The evidence handling model supports audit readiness by making the SOC's control work reviewable. It allows SynapticSOC to show not only that controls were implemented, but that implementation was supported by captured evidence, validation records, and public-safe summaries.

This is especially important for the completed control phases. RBAC evidence shows access boundaries. Acknowledgment evidence shows analyst interaction. Retention evidence shows review and retention intent. Public-redacted evidence allows those controls to be explained externally without exposing the live environment.

In this sense, evidence handling is not only a documentation process. It is part of the SOC control model.

Phase 1.5 Audit Writeback Fields

Following the Phase 1.5 Zeek visibility hardening update, SOC/01 v1.7 writes structured Phase 1.5 fields into Graylog audit records through SOAR GELF writeback. These fields extend the evidence record with Zeek visibility classification data for each processed alert.

Phase 1.5 fields written into Graylog audit records include:

phase — identifies the active SOC control phase

phase15_enabled — confirms Phase 1.5 is active for this alert

phase15_control — identifies the specific Phase 1.5 control applied

zeek_lookup_pass — records whether the match was immediate or deferred

zeek_visibility_status — records the observed visibility state

zeek_phase15_result — records the final classified outcome

zeek_deferred_lookup_needed — records whether a deferred pass was required

zeek_deferred_lookup_performed — records whether the deferred pass was executed

no_match_reason — documents the most likely explanation when no match was found

These fields are stored in Graylog without a leading underscore, consistent with Graylog GELF custom field handling. They are searchable in Graylog audit streams and available for management-level dashboard reporting.

What the Evidence Handling Model Does Not Claim

SynapticSOC's evidence handling model is a practical operational control. It is not presented as a certified forensic evidence system or a legally defensible chain-of-custody framework.

The model does not claim:

- Certified forensic evidence handling
- Legal-grade chain of custody
- Cryptographic proof of evidence integrity
- Tamper-evident storage
- Immutable archival retention
- Independent third-party audit
- Formal evidence transfer procedures
- Enterprise evidence management tooling

What the model does provide is more specific and defensible: a structured approach to capturing, classifying, protecting, reviewing, retaining, redacting, and publishing SOC evidence within the real constraints of a self-hosted SOC built with open-source security tools .

That is the control boundary implemented in SynapticSOC, and that is the boundary this whitepaper claims.

Section 11: Retention Policy Model

Evidence that is not retained cannot be reviewed. A SOC may generate alerts, workflow records, acknowledgment events, and audit artifacts, but if those records disappear before they can support investigation, management review, or audit explanation, the value of the evidence is limited.

SynapticSOC's retention policy model was created to ensure that SOC evidence remains available for the purpose it was generated to serve. Retention is not treated only as a storage-management concern. It is treated as a control that supports operational review, analyst accountability, workflow traceability, management reporting, compliance-oriented documentation, and public portfolio validation.

The model is practical rather than legalistic. It defines what is retained, why it is retained, how long different categories remain available, and what the model does not claim.

Retention Scope

The retention model covers the evidence-producing components of the SynapticSOC stack. Each component produces a different type of evidence, at a different volume, and with a different review value.

Component	Evidence Type	Retention Purpose
pfSense	Firewall and filterlog telemetry	Network boundary review and traffic context
Snort / Suricata	Network IDS alert records	Detection evidence and triage reference
Zeek	Network visibility and protocol telemetry	Corroboration, session context, and network evidence
Wazuh	Endpoint events, vulnerability data, and security records	Endpoint review, vulnerability tracking, and compliance-oriented visibility
Graylog	Parsed logs, triage streams, audit streams, and acknowledgment records	Primary SOC review, evidence review, and accountability evidence
OpenSearch / Wazuh Indexer	Backend indexed event storage	Search, retrieval, and retention of indexed SOC data
n8n	Workflow execution history and SOAR processing records	Automation traceability and acknowledgment workflow evidence
Grafana	Dashboard configuration and reporting views	Management, audit, and compliance-oriented visibility
Project folders	Screenshots, validation records, public-redacted artifacts, and control documents	Documentation evidence and public/private evidence separation

This scope reflects the reality of the project: evidence is not produced by one tool alone. It is distributed across the telemetry layer, correlation layer, automation layer, dashboard layer, and documentation layer.

Retention Differentiation

Not all evidence has the same operational lifespan. SynapticSOC therefore differentiates retention by evidence type rather than applying one uniform retention period to all records.

Raw telemetry includes firewall logs, IDS alerts, Zeek network visibility data, and high-volume endpoint events. This data has strong value during active investigation and short-term review, but it becomes less useful as time

passes and storage pressure increases. SynapticSOC therefore retains raw telemetry for a shorter operational window of approximately 30 to 40 days.

Audit and accountability records have a different value profile. Analyst acknowledgments, SOAR workflow records, audit stream entries, and accountability events are not merely operational logs. They document human interaction with the SOC process and support management review, retention validation, workflow traceability, and compliance-oriented explanation. SynapticSOC therefore retains audit and accountability records for a longer review window of approximately 180 to 210 days.

Documentation artifacts, including screenshots, validation checklists, summaries, control documents, and public-redacted evidence, are retained separately from live indexed telemetry. These artifacts support the permanent project record and are not treated as short-lived operational logs.

The distinction is deliberate. Treating audit records the same as high-volume telemetry would discard accountability evidence too quickly. Treating all raw telemetry as permanent would create an unnecessary storage burden. SynapticSOC balances those needs by assigning retention value according to evidence purpose.

Retention Period Model

The current retention model uses the following practical retention structure:

Evidence Category	Example Sources	Retention Position
Raw operational telemetry	pfSense logs, IDS alerts, Zeek records, high-volume endpoint events	Approx. 30 to 40 days for operational review
Triage and review records	Graylog streams, parsed events, correlated alerts	Retained for analyst review within the operational window
Audit and accountability records	Graylog audit streams, analyst acknowledgment records, SOAR workflow outputs	Approx. 180 to 210 days for accountability and reporting
Workflow execution evidence	n8n execution history for SOC/01 and SOC/02	Retained as supporting automation traceability
Endpoint security and vulnerability evidence	Wazuh alerts, vulnerability findings, compliance-oriented observations	Retained for endpoint review and management visibility
Dashboard and reporting views	Grafana audit, retention, and management dashboards	Retained as reporting configuration and visualization layer
Documentation artifacts	Screenshots, validation checklists, summaries, public-redacted evidence	Retained as project evidence and control documentation

Where specific index policies are configured, raw telemetry follows the shorter 30 to 40 day operational window, while audit and accountability records follow the longer 180 to 210 day review window. This supports both storage sustainability and audit-oriented review.

Tool-Level Retention Model

Graylog

Graylog is the primary SOC review and audit-evidence layer. It manages retention through index sets, stream organization, and dashboard visibility.

Graylog retention applies to parsed telemetry, triage records, audit streams, and acknowledgment records. The separation between detection streams and audit streams is important. Detection records support investigation. Audit records support accountability. They are related, but they do not serve the same purpose and should not be treated identically.

Retention configuration is validated through Graylog index management views, stream review, and dashboard evidence.

OpenSearch and Wazuh Indexer

OpenSearch and the Wazuh Indexer provide backend indexed storage and retrieval. They support the searchability and availability of telemetry and security events used by Graylog and Wazuh.

These systems are not analyst-facing platforms in the SynapticSOC access model, but their retention behavior directly affects what analysts, auditors, and dashboards can review. Backend retention therefore underpins the operational availability of the SOC evidence layer.

Wazuh

Wazuh retains endpoint security events, vulnerability evidence, and compliance-oriented host data through the Wazuh indexing layer. This supports read-only review by SOC analyst and compliance auditor roles.

Wazuh retention is important because endpoint evidence often provides context that network telemetry cannot. Vulnerability findings, host events, and compliance-related observations remain useful for periodic review, not only immediate triage.

n8n

n8n retains workflow execution history for the controlled SOAR intake and analyst acknowledgment workflows. This history supports automation traceability by showing that workflows executed, processed events, and produced acknowledgment outputs.

n8n execution history is not the primary audit record, but it is supporting evidence. The primary accountability record is stored in Graylog audit streams, while n8n provides workflow-level traceability behind those records.

Grafana

Grafana does not function as the primary storage layer for SOC telemetry. Its retention role is different. Grafana preserves dashboard configuration, reporting views, and management-level visibility over retained data.

Grafana dashboards remain useful only when the underlying data sources retain the data needed for reporting. The retention model therefore treats Grafana as a visibility and reporting layer dependent on the retention behavior of Graylog, Wazuh, and backend indexed storage.

Project Evidence Folders

The local project folder structure is part of the retention model. Screenshots, validation records, control documents, public-redacted summaries, and evidence indexes are retained in phase-specific folders.

This documentation retention layer is especially important for public portfolio and GitHub use. It allows the project to demonstrate controls without exposing the live SOC environment or raw sensitive evidence.

Retention Validation

Retention was validated during the retention policy phase. The validation process confirmed that retention configuration, retention evidence, and evidence separation existed across the relevant layers of the SOC.

Validated evidence included:

- Graylog index set overview evidence
- Graylog retention index set detail evidence
- Graylog audit stream evidence
- Graylog triage and audit dashboard evidence
- Grafana retention and audit dashboard evidence
- n8n execution history evidence
- Wazuh event evidence
- Wazuh vulnerability evidence
- Folder separation evidence for internal and public-redacted material

These validation artifacts demonstrate that retention was not only written as a policy, but also reviewed across the tools and documentation structure that support the SOC.

Retention validation does not claim legal defensibility or immutable archival proof. It confirms that a practical retention model exists, is documented, and is supported by evidence.

What Retention Supports

The retention model supports six operational and governance outcomes.

Operational review requires recent telemetry, triage events, and correlated alerts to remain available long enough for analyst investigation and follow-up.

Analyst accountability requires acknowledgment records and audit events to remain available beyond the immediate alert window, so that management or compliance reviewers can verify that analyst interaction occurred.

Workflow traceability requires n8n execution history and Graylog audit records to remain available together, allowing the path from alert intake to acknowledgment record to be reconstructed.

Endpoint and security evidence review requires Wazuh events, vulnerability evidence, and compliance-oriented host data to remain available for periodic review and investigation.

Management reporting requires Grafana dashboards to query retained data and present meaningful trends around audit activity, acknowledgment behavior, retention health, and SOC visibility.

Public/private evidence separation requires internal evidence and public-redacted material to remain organized separately over time, so that external documentation can be produced without exposing sensitive operational detail.

Together, these outcomes explain why retention is part of the SOC control model rather than a background storage setting.

Retention and the Evidence Lifecycle

Retention is directly connected to the evidence lifecycle defined in the evidence handling model.

Raw telemetry is collected, normalized, correlated, and reviewed during triage. It must remain available long enough to support investigation and short-term operational review. The 30 to 40 day raw telemetry window supports this purpose without treating all high-volume logs as permanent archival records.

Acknowledgment records, audit stream entries, and workflow evidence have longer-term value because they document process and accountability. The 180 to 210 day audit retention window supports management review, compliance-oriented reporting, and audit explanation.

Documentation artifacts have a different retention purpose again. Screenshots, validation checklists, summaries, and public-safe evidence support the permanent project record. They allow SynapticSOC to be reviewed externally without requiring live system access.

The retention model therefore aligns evidence availability with evidence purpose.

What the Retention Model Does Not Claim

SynapticSOC's retention model is a practical operational control. It is not presented as enterprise data lifecycle management, immutable archival storage, or legally mandated retention compliance.

The model does not claim:

- Immutable archival storage
- Legal-grade evidence retention
- Cryptographic proof of retention integrity
- Tamper-evident retention controls
- Certified compliance with regulatory retention mandates
- Automated access-reviewed retention governance
- Enterprise data lifecycle management tooling
- Off-site or geographically redundant archival storage
- Independent audit of retention enforcement

These limitations are documented intentionally. The project does not use retention language to imply certification, legal defensibility, or regulatory compliance.

What SynapticSOC does provide is a defined, validated, and documented retention approach that supports SOC review, analyst accountability, workflow traceability, endpoint evidence review, management reporting, and public/private evidence separation.

That is the retention control boundary implemented in SynapticSOC, and that is the boundary this whitepaper claims.

Section 12: Management and Audit Dashboards

Detection, correlation, acknowledgment, access control, evidence handling, and retention are operational controls. For those controls to support management review and audit readiness, their outputs must be visible in a controlled and understandable way.

SynapticSOC uses dashboards to make SOC operation reviewable without collapsing role boundaries. Analysts investigate through Graylog. Management and compliance reviewers observe through Grafana. These layers are intentionally separate because they serve different users, expose different levels of detail, and support different review objectives.

The dashboard model is therefore not only a visualization choice. It is part of the SOC control design.

Dashboard Design Principle

SynapticSOC separates dashboards by audience and purpose.

The analyst dashboard layer is designed for investigation. It provides access to triage streams, correlated events, enrichment context, audit records, and acknowledgment evidence. This layer requires operational detail and is therefore centered on Graylog.

The management and compliance dashboard layer is designed for review. It presents summarized visibility into SOC health, retention status, acknowledgment activity, workflow traceability, and audit evidence. This layer is centered on Grafana.

This separation prevents unnecessary access expansion. A CISO or compliance reviewer does not need raw triage stream access to understand SOC health. A SOC analyst does not need management reporting access to investigate an alert. Each audience receives the visibility required for its function, without gaining unnecessary access to unrelated data or controls.

Dashboard Layer Summary

Dashboard Layer	Platform	Primary Audience	Purpose
Operational SOC layer	Graylog	SOC analyst / evidence reviewer	Triage, correlation review, evidence search, audit stream review
Management and compliance layer	Grafana	CISO, management, compliance reviewer	SOC health, retention visibility, acknowledgment analytics, audit reporting
Backend storage layer	OpenSearch / Wazuh Indexer	Admin / service only	Data storage and retrieval, not direct dashboard access for normal users

This model follows the same role-separation logic used throughout SynapticSOC. Visibility is provided through the correct layer instead of granting broad platform access to every reviewer.

Graylog — Operational SOC Visibility

Graylog is the analyst-facing dashboard and evidence-review layer. Its dashboards are designed for active triage, event review, and operational investigation.

Graylog provides visibility into:

- Correlated detection events across firewall, IDS, Zeek, and Wazuh sources
- Enrichment context and event classification
- Triage streams and operational review queues
- Audit records produced by the controlled SOAR workflow
- Analyst acknowledgment records
- Workflow-related evidence stored in audit streams

The value of Graylog is that it presents operational detail in context. An analyst can review a detection event alongside related firewall activity, network visibility, endpoint evidence, enrichment data, and acknowledgment records.

Graylog also supports evidence review. Audit streams and acknowledgment streams are searchable, filterable, and reviewable at the operational level. This allows the SOC to reconstruct the path from detection to analyst acknowledgment without relying only on dashboard summaries.

Access to Graylog dashboards and streams is scoped. Analysts receive access to assigned streams and dashboards. General viewers receive limited visibility. Audit and evidence streams are not broadly shared by default. This preserves the distinction between operational triage, evidence review, and general SOC visibility.

Grafana — Management and Compliance Visibility

Grafana is the management, CISO, and compliance-oriented dashboard layer. It is not used as the primary investigation console.

Grafana provides summarized visibility into the health and accountability of the SOC operating model. Its purpose is to help reviewers answer higher-level questions:

- Is the SOC pipeline producing reviewable records?
- Are acknowledgment records being generated?
- Is workflow activity visible over time?
- Is retention status observable?
- Are audit records available for management review?
- Are dashboards reflecting the controls described in the whitepaper?

Grafana dashboards in SynapticSOC support:

Retention visibility — showing whether retention-related records, audit streams, and underlying data sources remain available for review.

Audit visibility — summarizing audit records, acknowledgment activity, and workflow evidence without requiring direct access to raw triage streams.

Acknowledgment analytics — presenting analyst acknowledgment activity, timing patterns, and workflow output trends.

Pipeline health review — showing whether the SOC workflow is active and producing expected records.

Management reporting — providing a view suitable for CISO-level or compliance-oriented review without exposing unnecessary operational detail.

Grafana access is restricted to Viewer-level accounts for non-administrative users. Viewer users can observe dashboards but cannot edit dashboards, modify data sources, administer users, or change Grafana configuration. This keeps the reporting layer useful while preserving control integrity.

Dashboard Access Boundaries

The dashboard model depends on access boundaries being maintained.

User / Role	Appropriate Dashboard Layer	Access Boundary
SOC analyst	Graylog	Triage and assigned evidence streams
SOC viewer	Graylog	Limited dashboard visibility
CISO viewer	Grafana	Management dashboards only
Compliance auditor	Grafana and read-only Wazuh	Audit, retention, and compliance-oriented visibility
Automation owner	n8n and supporting dashboards	Workflow administration and execution review
Backend administrator	OpenSearch / Wazuh Indexer	Infrastructure administration only

The purpose of this model is to avoid unnecessary platform crossover. Management users do not need direct access to raw detection streams. Analysts do not need administrative reporting access. Backend systems do not become unofficial investigation consoles.

This keeps the dashboard model aligned with the RBAC model described earlier in the whitepaper.

What the Dashboard Layer Supports

The combined Graylog and Grafana dashboard model supports several review functions.

SOC review is supported through Graylog triage dashboards, stream views, correlation records, and evidence searches.

Audit readiness is supported through searchable audit streams, acknowledgment records, workflow evidence, and management dashboards that summarize those records.

Management visibility is supported through Grafana dashboards that present SOC activity, retention health, workflow trends, and acknowledgment analytics at an appropriate level of abstraction.

Retention validation is supported by dashboards and evidence views that show whether retained records are available for review within the documented retention model.

Control review is supported by making RBAC, evidence handling, retention, and acknowledgment outputs visible as records and dashboard views rather than only policy statements.

Evidence traceability is supported by allowing reviewers to follow the path from detection event to correlated alert, analyst acknowledgment, audit record, and management summary.

Together, these outcomes make dashboards part of the control evidence layer rather than decorative reporting.

Dashboards and the Compliance-Oriented Narrative

The dashboard layer is where SynapticSOC's compliance-oriented controls become observable.

A policy can state that acknowledgments are recorded. A dashboard can show acknowledgment records and trends. A retention policy can define review windows. A dashboard can show whether relevant records are

available. An access-control model can define user boundaries. Dashboard access validation can show that those boundaries are being enforced.

This does not make SynapticSOC a certified compliance platform. It does make the SOC easier to review. The dashboards translate operational controls into visible evidence that can be understood by technical reviewers, hiring managers, compliance-oriented readers, and portfolio audiences without requiring direct access to the live infrastructure.

For a self-hosted SOC project, this is important. The dashboards provide a public-safe way to demonstrate operational maturity while keeping sensitive raw evidence, backend systems, and automation controls protected.

Phase 1.5 Visibility Outcome Tracking

Following the Phase 1.5 update, Graylog audit records contain structured Zeek visibility outcome fields written by SOC/01 v1.7. These fields are searchable within the Graylog audit stream, allowing analysts and SOC managers to review Zeek lookup pass results, visibility status classifications, and no-match reason documentation across processed alerts.

Grafana management dashboards can be extended to track Phase 1.5 visibility outcomes over time, including the distribution of zeek_seen_immediate, zeek_seen_deferred, and zeek_no_match_after_delay results. This supports management-level visibility into sensor coverage patterns and helps identify persistent no-match conditions that may indicate SPAN or mirroring gaps requiring Phase 2 attention.

What the Dashboard Model Does Not Claim

SynapticSOC's dashboard model is a practical visibility and review layer. It is not presented as an enterprise GRC platform or certified compliance reporting system.

The model does not claim:

- Automated compliance framework mapping
- Certified audit report generation
- Regulatory-grade evidence export
- Real-time compliance SLA enforcement
- Enterprise risk register integration
- Multi-tenant compliance reporting
- Organization-wide governance reporting
- Replacement for formal GRC tooling

What the dashboard model does provide is more specific and defensible: purpose-separated visibility for analysts, management users, and compliance-oriented reviewers.

Graylog provides operational detail. Grafana provides management and audit visibility. Access controls keep those layers separate. The dashboards make the SOC's controls observable without weakening the boundaries that protect them.

Section 13: Final Control Pack

Phases A through C produced three major control bodies: the RBAC and access-control model, the evidence handling model, and the retention policy model. Each phase was implemented, validated, and documented independently. Phase D consolidated those outputs into a single reviewable control pack.

The final control pack is not simply a summary of previous work. It is the structured reference set that connects implementation, validation evidence, policy language, and public-safe documentation into one coherent control model.

Its purpose is to make SynapticSOC reviewable as a SOC operating model rather than as a collection of separate technical exercises.

Why the Control Pack Was Needed

Building controls in phases is practical. Reviewing them only through separate phase folders is not.

An external reviewer should not need direct access to the live SOC environment, internal screenshots, or every phase folder to understand what controls exist, what was validated, and what the current control boundaries are. A hiring manager, certification reviewer, open-source contributor, or technical reader should be able to follow the control story from a single structured reference point.

The final control pack serves that function. It consolidates the completed control work into a format that is easier to review, explain, and publish safely.

It also protects the integrity of the whitepaper. Every major control claim in this document should be traceable to an implemented control, a validation artifact, or a public-safe summary. The control pack provides that traceability.

Control Pack Contents

The Phase D control pack includes the following deliverables.

Control Pack Component	Purpose
Source phase reference index	Maps each control area back to its originating phase folder, policy document, and evidence location
Final control pack overview	Provides the high-level explanation of the completed control model
Integrated control matrix	Consolidates RBAC, evidence handling, and retention controls into one reviewable table
Final control pack summary	Provides a narrative explanation of the control posture for non-specialist or management-oriented readers
Whitepaper integration notes	Maps validated control work to the relevant whitepaper sections
Public-safe control summary	Provides a sanitized version suitable for GitHub, website, portfolio, or external publication
Final document index	Lists the documents, summaries, validation checklists, and evidence artifacts produced across the control phases
Source folder validation	Confirms that expected phase folders, subfolders, and key documents exist and are organized correctly
Phase D validation checklist	Confirms that Phase D deliverables were completed, reviewed, and stored correctly

Backup archive	Provides a point-in-time archive of the completed Phase D output set

Together, these documents create a structured control reference. The reviewer can start from the overview, inspect the integrated control matrix, trace individual controls back to source phase evidence, and use the public-safe summary for external presentation.

Integrated Control Matrix

The integrated control matrix is the functional core of the control pack. It brings the completed controls from Phases A, B, and C into one table so that the SOC control posture can be reviewed without navigating each phase separately.

Control Area	Phase	Validation Status
Graylog stream and dashboard scoped access	Phase A	Validated
Graylog audit and evidence stream separation	Phase A	Validated
n8n owner-only access with 2FA	Phase A	Validated
n8n restricted from analyst and management access	Phase A	Validated
Grafana ciso.viewer Viewer role	Phase A	Validated
Grafana compliance.auditor Viewer role	Phase A	Validated
Wazuh wazuh.soc.analyst read-only API and indexer access	Phase A	Positive and negative validation
Wazuh wazuh.compliance.auditor read-only API and indexer access	Phase A	Positive and negative validation
Wazuh dashboard run_as permission scoping	Phase A	Validated
OpenSearch / Wazuh Indexer backend administrative and service-only access	Phase A	Reviewed and documented
Evidence source inventory	Phase B	Validated
Evidence classification model	Phase B	Validated
Evidence lifecycle model	Phase B	Validated
Chain-of-custody considerations	Phase B	Documented
Analyst evidence handling responsibilities	Phase B	Defined
Public/private evidence separation	Phase B	Validated
Redaction rules for public evidence	Phase B	Defined
Graylog retention index sets and retention evidence	Phase C	Validated
Differentiated retention for raw telemetry and audit records	Phase C	Validated
n8n execution history retention evidence	Phase C	Validated
Wazuh event and vulnerability evidence retention	Phase C	Validated
Grafana retention and audit dashboard visibility	Phase C	Validated
Folder-level retention and public-redacted separation	Phase C	Validated

Phase 1.5 addendum: Zeek deferred lookup and visibility-gap classification were validated through production cutover, immediate match-path confirmation, and active GELF audit writeback.

This matrix is important because it turns the control work into a reviewable structure. It shows what was implemented, where it belongs, and whether it was validated.

What the Control Pack Enables

The final control pack supports four major review outcomes.

Portfolio review is supported because a technical reader can understand the SOC's control posture without requiring access to private infrastructure. The public-safe control summary and integrated matrix allow SynapticSOC to be presented externally without exposing sensitive evidence.

Project governance is supported because the source phase reference index and document index provide a map of what was implemented, where it is documented, and where the supporting evidence resides. This helps future project work stay aligned with the existing control model.

Audit-oriented explanation is supported because control claims can be traced back to phase outputs, validation artifacts, and documented evidence. The control pack does not replace an audit, but it makes the control posture explainable in an audit-oriented format.

Whitepaper integrity is supported because the whitepaper no longer relies on broad claims. Its control sections are grounded in implemented and validated work. The whitepaper integration notes help ensure that the document describes the SOC as it exists, not as an aspirational roadmap.

Control Pack Relationship to the SOC Operating Model

The final control pack marks an important maturity point in SynapticSOC.

A detection lab can show tools, alerts, dashboards, and integrations. A SOC operating model must also show how access is controlled, how evidence is handled, how records are retained, how accountability is captured, and how controls can be reviewed.

The control pack is the artifact that makes that maturity visible.

It does not add a new detection engine or a new telemetry source. Instead, it consolidates the governance work that makes the existing SOC more credible. It connects RBAC, evidence handling, retention, validation evidence, and public-safe documentation into one control reference.

For a public portfolio project, this matters. A reviewer can see not only that SynapticSOC works technically, but that it has been documented and governed with discipline.

Public-Safe Control Summary

A major outcome of Phase D is the public-safe control summary. This document allows the control work to be explained externally without exposing sensitive internal evidence, private infrastructure details, credentials, workflow secrets, internal paths, or raw operational records.

The public-safe summary is important because SynapticSOC is both an operational SOC project and a public professional portfolio. The project must be demonstrable without being overexposed.

The public-safe control summary supports that balance. It allows the project to show:

- What controls were implemented
- Which phases produced them
- What validation approach was used

- What the control boundaries are
- What the project does not claim

This makes the control work suitable for GitHub, the SynapticSOC website, LinkedIn, and whitepaper publication.

What the Control Pack Does Not Claim

The final control pack is a practical governance reference for a self-hosted SOC built with open-source security tools. It is not a formal audit package, certified compliance submission, or enterprise GRC deliverable.

The control pack does not claim:

- Certification against any compliance framework
- Legal-grade evidence package suitable for regulatory submission
- Third-party audit validation
- Enterprise GRC platform integration
- Immutable or tamper-evident control documentation
- Complete coverage of every SOC governance domain
- Formal risk register ownership or enterprise policy approval
- Production-ready multi-tenant governance

These limitations are intentional and documented.

What the control pack does claim is narrower and stronger: SynapticSOC has a structured, validated, and public-safe reference set covering the RBAC, evidence handling, and retention controls implemented across the completed compliance-oriented phases.

That is the control boundary Phase D consolidated, and that is the boundary this whitepaper claims.

Section 14: Limitations and Non-Claims

Every security project has boundaries. A credible SOC whitepaper should define those boundaries clearly rather than allow readers to infer capabilities that were not implemented, validated, or intended.

SynapticSOC is documented honestly for that reason. The project demonstrates practical SOC maturity under real constraints, but it does not present itself as a certified compliance platform, legal forensic system, commercial SOC replacement, or enterprise managed detection service.

This section defines what SynapticSOC does not claim, followed by a precise statement of what it does claim.

Compliance and Certification

SynapticSOC is not a certified compliant SOC.

It has not been externally audited, formally assessed, certified, attested, or validated against any regulatory, industry, or governance framework. This includes, but is not limited to:

- ISO/IEC 27001 certification
- SOC 2 Type I or Type II attestation
- PCI DSS validation
- HIPAA compliance assessment
- GDPR compliance assessment
- NIST Cybersecurity Framework conformity assessment
- CIS Controls implementation assessment
- Cyber Essentials or Cyber Essentials Plus certification

The terms compliance-oriented, audit-ready, and control-aligned are used deliberately. They describe how the project was designed and documented. They do not indicate certification, legal compliance, or external assurance.

SynapticSOC has implemented practical controls around access separation, evidence handling, retention documentation, analyst accountability, and management visibility. Those controls may later be mapped against frameworks such as CIS Controls, NIST CSF, or PCI DSS for educational or portfolio purposes. Such mapping would still not constitute certification or compliance validation.

Forensic Evidence and Chain of Custody

SynapticSOC does not provide legal-grade forensic evidence handling or certified chain of custody.

The evidence handling model described in this whitepaper is a practical operational control. It supports evidence classification, structured capture, traceability, redaction, public/private separation, and audit-oriented review. It is not a substitute for formal forensic procedures.

The current model does not provide:

- Cryptographic evidence signing
- Tamper-evident or write-once evidence storage
- Independent witness verification
- Legal hold procedures

- Court-admissible chain-of-custody documentation
- Certified forensic examiner involvement
- Formal evidence transfer procedures
- Litigation-ready evidence preservation

Evidence produced by SynapticSOC is suitable for SOC review, management reporting, project validation, and portfolio documentation. It should not be treated as legal forensic evidence without independent forensic collection, validation, and custody controls performed by qualified personnel.

Retention and Archival

SynapticSOC does not provide immutable archival storage or legally mandated retention compliance.

The retention model defines practical retention windows for different categories of SOC evidence. Raw operational telemetry is retained for a shorter review period, while audit and accountability records are retained for a longer review period. These periods are operational decisions based on usefulness, storage constraints, and review needs.

The retention model does not claim:

- Immutable archival storage
- Legal-grade evidence retention
- Cryptographic proof of retention integrity
- Tamper-evident retention enforcement
- Regulatory retention compliance
- Legal hold capability
- Off-site or geographically redundant archival storage
- Enterprise data lifecycle management

The retention model supports SOC review, accountability, workflow traceability, endpoint evidence review, management reporting, and public/private evidence separation. It does not claim to satisfy any specific statutory, contractual, or regulatory retention requirement.

Detection Coverage

SynapticSOC does not guarantee complete detection coverage.

The detection and visibility pipeline covers the monitored environment through firewall telemetry, network IDS alerts, passive network visibility, endpoint monitoring, log management, correlation, and dashboard review. However, detection depends on multiple conditions:

- Traffic must be visible to the relevant monitoring point.
- Snort, Suricata, Wazuh, and other detection sources must have relevant rules or logic.
- Zeek must observe the relevant network activity.
- Wazuh agents must be installed, active, and are forwarding events from monitored endpoints.
- Log forwarding must remain reliable.
- Graylog parsing and correlation logic must process the event correctly.
- The activity must fall within the monitored scope of the environment.

Gaps in visibility, rule coverage, agent deployment, encrypted traffic inspection, logging reliability, or parsing accuracy may result in missed detections.

SynapticSOC does not claim to detect all attack techniques, all malware, all living-off-the-land activity, all encrypted threats, all insider activity, or all activity outside the monitored environment. It demonstrates practical detection and correlation capability, not comprehensive threat coverage.

Automation and Response

SynapticSOC does not provide autonomous remediation, automated destructive response, or managed detection and response capability.

The controlled SOAR workflow is intentionally conservative. Automation is used for intake, enrichment support, notification, acknowledgment processing, workflow traceability, and audit record creation. It is not used to make independent response decisions.

The current model does not perform:

- Automatic IP blocking
- Automatic pfSense firewall rule modification
- Automatic endpoint isolation or quarantine
- Automatic file deletion or quarantine
- Automatic process termination
- Automatic service shutdown
- Automatic case escalation or closure
- Autonomous incident resolution
- Managed detection and response services

Response decisions remain human-controlled. The project's current automation boundary is intentional: automated destructive or disruptive action requires a level of detection confidence, enrichment trust, rollback planning, and operational validation that SynapticSOC does not currently claim.

This is not presented as a failure of the SOC. It is a deliberate control decision.

Scale and Operating Environment

SynapticSOC is a single-operator, self-hosted SOC project. It is not presented as an enterprise production SOC or commercial monitoring platform.

The project does not claim to be:

- A multi-analyst SOC team environment
- A managed security service provider platform
- A commercial SOC-as-a-service offering
- A production monitoring service for third-party customers
- An enterprise high-availability SOC
- A replacement for dedicated security operations staffing
- A full multi-tenant security operations platform

The controls documented in this whitepaper were implemented and validated within a single-operator environment. Some enterprise controls are therefore outside the current scope, including formal

joiner/mover/leaver processes, privileged access workstations, centralized identity governance, hardware-backed MFA across all tools, just-in-time access provisioning, and formal access recertification workflows.

These limitations reflect the operating environment. They are not hidden, and they should not be interpreted as enterprise capabilities.

Infrastructure and Availability

SynapticSOC runs on self-hosted infrastructure in a controlled local environment. It does not provide enterprise availability guarantees.

The project does not currently include:

- High-availability clustering
- Enterprise hardware redundancy
- Geographic redundancy
- Formal disaster recovery procedures
- Service level agreements
- Uptime guarantees
- Dedicated physical security controls beyond the local environment
- Enterprise backup and restore certification
- DDoS protection for hosted services

Hardware failure, power interruption, storage pressure, network changes, or local infrastructure issues may affect SOC availability. These are accepted constraints of the project and are part of the real-world context in which SynapticSOC was built.

Case Management

SynapticSOC does not currently include a full case management platform.

The project has analyst acknowledgment, audit records, workflow traceability, and dashboard visibility, but those controls are not a replacement for formal case management. Acknowledgment records show that an analyst response was recorded. They do not provide the full incident lifecycle features of a case management system.

Current limitations include:

- No formal case creation workflow
- No case assignment queue
- No escalation lifecycle
- No case status model
- No evidence attachment to formal incident cases
- No case closure workflow
- No case-level reporting

TheHive or an equivalent case management layer remains a future roadmap item. Until that integration is implemented and validated, SynapticSOC does not claim formal case management capability.

Identity and Access Governance

SynapticSOC implements practical RBAC and access separation, but it does not claim enterprise identity governance.

The current access-control model includes scoped roles, read-only users, restricted automation access, backend access separation, and negative validation where applicable. However, it does not include:

- Centralized identity provider integration
- Organization-wide single sign-on
- Hardware-backed MFA across all SOC platforms
- Privileged access management tooling
- Just-in-time access workflows
- Automated access review scheduling
- Formal joiner/mover/leaver process
- Enterprise privileged session recording

The RBAC model is still meaningful. It demonstrates least privilege and role separation within the project's constraints. It should not be overstated as a full IAM or PAM program.

What SynapticSOC Does Claim

The limitations above do not weaken the project. They define its boundaries.

SynapticSOC makes the following specific claims:

- It has a functioning detection and visibility pipeline across firewall telemetry, network IDS alerts, passive network visibility, and endpoint monitoring.
- It uses Graylog as the primary SOC triage, correlation, stream review, and evidence-review layer.
- It uses Wazuh for endpoint, security, vulnerability, and compliance-oriented visibility.
- It uses n8n for controlled SOAR intake and analyst acknowledgment processing.
- It uses Telegram as a lightweight analyst notification and acknowledgment channel.
- It records analyst acknowledgment activity as audit evidence.
- It separates operational triage, management visibility, automation control, and backend infrastructure access.
- It implements and validates RBAC across Graylog, Wazuh, Grafana, n8n, and backend infrastructure boundaries.
- It documents an evidence handling model covering classification, lifecycle, redaction, analyst responsibilities, and public/private evidence separation.
- It documents a retention model with differentiated treatment for raw telemetry, audit records, workflow history, endpoint evidence, dashboards, and project artifacts.
- It provides management and audit visibility through Grafana dashboards without making Grafana the analyst investigation console.
- It consolidates RBAC, evidence handling, retention documentation, validation records, and public-safe summaries into a final control pack.
- It documents non-claims and limitations explicitly rather than implying maturity that has not been implemented or validated.

These claims are specific, limited, and defensible. They describe what SynapticSOC has built, documented, and validated as a practical self-hosted SOC operating model built with open-source security tools under real constraints.

Section 15: Future Roadmap

SynapticSOC is an active project. The work documented in this whitepaper represents a completed maturity milestone, not a finished product. The SOC has moved beyond basic visibility and detection into a documented operating model with access control, evidence handling, retention documentation, analyst acknowledgment, and an integrated control pack.

The roadmap below defines the next stages of maturity. It does not present future work as current capability. Items are organized by priority and dependency: near-term work strengthens the current operating model, medium-term work expands SOC workflow maturity, and longer-term work introduces advanced detection, behavioral analysis, and automation capabilities once the foundations are stable.

Completed — Removed from the Roadmap

The following items appeared in earlier planning discussions or roadmap drafts. They are now complete and are treated as current-state capabilities rather than future work:

- Firewall telemetry ingestion
- Network IDS alerting through Snort and Suricata
- Zeek network visibility and correlation support
- Wazuh endpoint monitoring and security visibility
- Graylog log management, parsing, enrichment, and triage
- OpenSearch / Wazuh Indexer backend storage
- Grafana management and compliance dashboards
- n8n controlled SOAR intake
- Telegram-based analyst acknowledgment workflow
- Analyst acknowledgment audit records
- RBAC and access-control implementation
- Evidence handling model and policy
- Retention policy documentation
- Public/private evidence separation
- Final control pack and whitepaper integration
- Public-facing project documentation and architecture narrative
- Phase 1.5 Zeek deferred lookup, no-match classification, audit writeback, and production cutover to SOC/01 v1.7

These are not roadmap items. They are delivered capabilities that form the current SynapticSOC operating model.

Near-Term Roadmap

The near-term roadmap focuses on closing the most important operational gaps that remain after the control pack phase.

Case Management Integration

The most important next capability is formal case management. SynapticSOC currently records analyst acknowledgment and audit evidence, but it does not yet provide a full case lifecycle.

A future case management layer, such as TheHive or an equivalent open-source platform, would add:

- Case creation from validated triage events
- Case assignment and ownership
- Case status tracking
- Evidence attachment
- Escalation workflow
- Case closure documentation
- Incident lifecycle reporting

This would close the gap between acknowledgment and full incident handling. Analyst acknowledgment proves that a human review step occurred. Case management would document what happened after that review.

Incident Response Runbooks

SynapticSOC needs formal runbooks for common alert types and investigation paths. These runbooks should define what an analyst checks, what evidence must be preserved, when escalation is required, and what response options are allowed.

Initial runbooks should cover:

- IDS alert triage
- Wazuh high-severity endpoint alerts
- Suspicious external IP activity
- Zeek-corroborated network events
- Vulnerability finding review
- Analyst acknowledgment handling
- False-positive documentation

Runbooks will make the analyst acknowledgment workflow more operationally meaningful by connecting acknowledgment to defined response procedures.

Backup and Restore Validation

The current project relies on local self-hosted infrastructure and local evidence storage. A documented and tested backup and restore process is needed for key SOC components.

Priority backup targets include:

- Graylog configuration and content packs
- Graylog index and stream configuration
- n8n workflows and credentials export process
- Wazuh configuration, rules, and relevant policy files
- Grafana dashboards
- Evidence folders and public-redacted artifacts

- Whitepaper, control pack, and validation documents

The goal is not to claim enterprise disaster recovery. The goal is to prove that the project can recover its critical configuration and documentation after a local failure.

TLS and Certificate Hardening

Internal SOC services should be reviewed for TLS coverage and certificate handling. This includes Graylog, OpenSearch / Wazuh Indexer, Wazuh dashboard components, Grafana, and n8n where applicable.

The roadmap item should include:

- Identifying which internal services currently use encrypted communication
- Identifying where plaintext internal access still exists
- Documenting certificate locations and renewal expectations
- Avoiding unnecessary breakage of currently stable pipelines
- Applying TLS improvements gradually and validating each change

This is a security hardening item, not a current compliance claim.

Secrets Management Improvement

n8n currently stores workflow credentials for automation and enrichment functions. That model is operational, but a stronger secrets management approach would improve the project's maturity.

Future work may include:

- Better local secrets documentation
- Dedicated secrets vault evaluation
- Separation of automation credentials by function
- Rotation procedure documentation
- Emergency credential revocation procedure
- Public documentation rules for screenshots involving credential areas

This roadmap item should be handled carefully because changing credential storage can break working workflows if rushed.

Medium-Term Roadmap

The medium-term roadmap expands operational maturity after the near-term control and recovery foundations are stable.

Detection Engineering Lifecycle

SynapticSOC needs a more formal detection engineering lifecycle for Snort, Suricata, Wazuh, Zeek-derived detections, and Graylog correlation logic.

This should include:

- Rule inventory
- Rule purpose documentation
- False-positive tracking
- Tuning history
- Test event generation

- Detection retirement criteria
- Severity review
- Mapping detections to telemetry sources

This would turn detection management from tool configuration into a documented engineering process.

Expanded Zeek Analytics

Phase 1.5 addressed the timing-related no-match problem through deferred lookup and outcome classification. The remaining Phase 2 scope for Zeek covers full LAN sensor coverage, stronger SPAN and mirroring design, QUIC and HTTP/3 visibility improvements, and broader passive monitoring maturity. These require infrastructure changes beyond workflow-level mitigation.

Zeek currently provides passive network visibility and corroboration. Future work should increase the analytical value of Zeek data.

Potential expansions include dashboards and correlation logic for:

- DNS anomalies
- Unusual connection patterns
- Long-duration connections
- Rare external destinations
- TLS certificate observations
- HTTP behavior patterns
- Internal-to-external communication trends
- Repeated failed connection behavior

This would strengthen Zeek's role from supporting context into a more active behavioral visibility layer.

Threat Intelligence Integration

The current enrichment model uses reputation context carefully and avoids treating enrichment as an automated verdict. Future work can expand this into a more structured threat intelligence model.

This may include:

- Multiple reputation sources
- Indicator confidence handling
- Feed source documentation
- Indicator expiration rules
- False-positive handling
- Graylog enrichment pipeline improvements
- Case-level observable enrichment once case management exists

Threat intelligence should remain contextual until the confidence model is mature enough to support stronger response decisions.

Controlled Wazuh Active Response Evaluation

Wazuh Active Response remains disabled in the current SOC model. That position should remain until detection confidence, enrichment accuracy, and rollback procedures are more mature.

A future controlled evaluation may test limited Active Response only for low-risk, high-confidence use cases. Any pilot should include:

- Clear trigger conditions
- Manual approval or staged approval
- Rollback steps
- Test environment validation
- False-positive review
- Evidence capture
- Explicit exclusion of destructive broad actions

Active Response should not be enabled broadly. It should be treated as a controlled experiment after stronger validation.

Periodic Control Review

The completed control pack should not become stale. A periodic review cadence should be introduced to confirm that implemented controls still match the documented model.

Review areas should include:

- Graylog stream and dashboard access
- Wazuh read-only access
- Grafana Viewer accounts
- n8n owner-only access and 2FA
- Retention status
- Evidence folder separation
- Public-redacted evidence handling
- Control matrix accuracy
- Dashboard availability
- Workflow execution health

This turns the control pack from a one-time deliverable into an ongoing governance reference.

Compliance Framework Mapping

SynapticSOC may later map its implemented controls against one or more reference frameworks such as CIS Controls, NIST CSF, or selected PCI DSS concepts.

This would be for educational, portfolio, and gap-analysis purposes only. It would not be presented as certification, attestation, regulatory compliance, or external audit validation.

The value of this work would be to show:

- Which existing SynapticSOC controls align with recognized control themes
- Which areas are partially covered
- Which areas remain genuine gaps
- Which future roadmap items would improve control coverage

Longer-Term Roadmap

The longer-term roadmap introduces advanced capabilities that require stronger foundations, more data, and careful validation.

Behavioral Analysis and AI-Assisted Correlation

A major longer-term goal is to introduce behavioral analysis and AI-assisted advanced correlation. This aligns with the original project direction of moving beyond static alerting into higher-level analysis of patterns, relationships, and abnormal behavior across SOC data.

This work may include:

- Baseline modeling of normal network and endpoint activity
- Anomaly detection over Zeek, firewall, Wazuh, and Graylog event patterns
- AI-assisted triage summaries for correlated events
- Clustering of similar alerts and repeated behaviors
- Pattern detection across time windows
- Scoring support based on multiple weak signals
- Analyst-facing explanation of why an event is unusual
- GPU-assisted or model-assisted enrichment where practical

This capability must be introduced carefully. AI-assisted analysis should support analyst judgment, not replace it. It should not be used to trigger automated destructive response actions without extensive validation.

The goal is advanced correlation and behavioral visibility, not AI-branded automation hype.

Morpheus / GPU-Assisted Security Analytics Evaluation

Earlier project planning included the possibility of GPU-assisted security analytics. This remains a longer-term research and engineering direction, not a current capability.

A future evaluation may explore NVIDIA Morpheus or similar pipelines for:

- Log embedding
- Event classification
- Anomaly detection
- High-volume enrichment
- Behavioral sequence analysis
- Model-assisted prioritization

This would require careful hardware planning, dataset preparation, model validation, and integration with the existing Graylog/Wazuh/OpenSearch pipeline. It should only be pursued after the current SOC workflow, retention model, and case management layer are stable.

Case-Centric Enrichment with Cortex or Equivalent

After case management is implemented, enrichment should move closer to the case lifecycle. Cortex or an equivalent analyzer framework may be used to enrich observables attached to cases.

This would support:

- IP reputation analysis

- Domain and URL enrichment
- File hash lookup
- Malware or sandbox report attachment
- Observable-level case notes
- Repeatable enrichment actions tied to case records

This should not replace the current controlled SOAR model. It should extend it by attaching enrichment evidence to formal cases.

Purple-Team and Validation Exercises

Future work should include controlled validation exercises to test detection, triage, acknowledgment, and evidence handling under realistic conditions.

This may include:

- Caldera-based adversary emulation
- Safe internal test scenarios
- Detection validation exercises
- Simulated incident walkthroughs
- Runbook testing
- Case management testing
- Evidence handling drills

The goal is not offensive tooling for its own sake. The goal is to validate whether the SOC detects, correlates, notifies, acknowledges, records, and reports as designed.

Vulnerability Management Expansion

Wazuh vulnerability visibility exists, but vulnerability management can be expanded through structured scanning and review.

Future work may include:

- OpenVAS / GVM scanning
- Nuclei-based checks where appropriate
- Asset-based vulnerability tracking
- Prioritization by exposure and severity
- Dashboarding of vulnerability trends
- Integration with case management or task tracking

This should be framed as vulnerability management maturity, not as a replacement for enterprise vulnerability management platforms.

Public Demonstration and Project Page Expansion

SynapticSOC also has a public documentation roadmap. The website, GitHub repository, whitepaper, and future video content should present the project clearly without exposing sensitive details.

Future public-facing work includes:

- Final SynapticSOC project page
- Whitepaper publication

- Public-safe control summary
- Architecture walkthrough
- Detection pipeline demonstration
- Controlled SOAR and acknowledgment demonstration
- RBAC and retention explanation
- YouTube demonstration series

This work supports portfolio visibility, but it should remain aligned with the same evidence handling and public-redaction rules described earlier in the whitepaper.

What the Roadmap Does Not Include

The roadmap does not include capabilities that are not genuine project goals.

SynapticSOC does not currently aim to become:

- A commercial SOC replacement
- A managed detection and response provider
- A multi-tenant monitoring platform
- A certified compliance product
- A legal forensic evidence platform
- An enterprise high-availability SOC
- A substitute for dedicated security operations staffing

The roadmap also does not include personal certification or ambassador goals as project deliverables. Professional development may happen alongside SynapticSOC, but it should not be represented as part of the SOC architecture roadmap unless it directly affects the project.

This distinction matters. The roadmap should describe the future of the SOC, not inflate the project with unrelated ambitions.

Roadmap Position

The next major maturity step for SynapticSOC is case management and runbook-driven response. After that, the project can expand into detection engineering, stronger retention and recovery validation, controlled response evaluation, and advanced behavioral analysis.

The long-term direction is clear: SynapticSOC should evolve from a documented SOC operating model into a more complete open-source security operations environment with case management, repeatable response procedures, deeper Zeek analytics, structured threat intelligence, and eventually AI-assisted behavioral correlation.

Those are future goals. The current whitepaper claims only what has already been implemented, validated, and documented.

Section 16: Conclusion

SynapticSOC began as a practical self-hosted SOC operating model built with open-source security tools focused on visibility, detection, and correlation. Its original goal was to determine how much meaningful security operations capability could be built using self-hosted tools, limited infrastructure, and disciplined engineering decisions.

The project has now crossed a more important threshold.

It is no longer only a collection of integrated security tools. It is a documented SOC operating model with defined telemetry sources, detection layers, correlation logic, controlled automation, analyst acknowledgment, access separation, evidence handling, retention documentation, dashboard visibility, and an integrated control pack.

The maturity path of the project reflects that progression:

Visibility → Detection → Correlation → Controlled Response → Analyst Acknowledgment → RBAC → Evidence Handling → Retention Documentation → Audit-Ready Control Pack

Each stage built on the one before it. Telemetry came before detection. Detection came before correlation. Correlation came before controlled response. Controlled response came before analyst accountability. Accountability then created the need for access control, evidence handling, retention discipline, dashboard visibility, and final control consolidation.

That order matters. SynapticSOC did not add governance language to an unfinished technical stack. It implemented operational capability first, then documented the controls required to govern that capability responsibly.

The project does not claim certified compliance, legal-grade forensic custody, immutable archival storage, enterprise high availability, commercial SOC replacement capability, or managed detection and response service maturity. Those boundaries are stated clearly because they are part of the project's credibility.

What SynapticSOC does demonstrate is more specific and more defensible:a practical self-hosted SOC built with open-source security tools can be built, operated, documented, and governed under real constraints. It can collect telemetry, detect suspicious activity, correlate events, notify an analyst, record acknowledgment, separate roles, protect evidence, document retention, provide management visibility, and present its control posture in a reviewable format.

That is the value of the project.

SynapticSOC shows that SOC maturity is not only about the number of tools deployed or the volume of alerts generated. It is about whether the environment can explain what it sees, how it responds, who can access it, how evidence is handled, how long records remain available, and what the project honestly does not claim.

The result is a security operations project that can be reviewed not only for its technical architecture, but for its operating discipline.

SynapticSOC remains an active project. Future work will expand case management, runbooks, backup validation, TLS hardening, secrets handling, detection engineering, threat intelligence, behavioral analysis, and AI-assisted correlation. Those are future maturity layers, not current claims.

The project also handles visibility uncertainty more responsibly. Rather than treating a Zeek no-match as a silent absence of evidence, Phase 1.5 classifies and documents no-match conditions through a two-pass lookup model and structured audit writeback fields. Visibility gaps are recorded, not hidden.

As it stands today, SynapticSOC represents a completed and documented milestone: a self-hosted SOC operating model built with open-source security tools under real-world constraints, with practical controls, validated evidence, and a public-safe narrative that can be reviewed without access to the live environment.