ANOMALI

Guide

OPERATIONALIZING THREAT INTELLIGENCE:

Actionable SOC Workflows for Detection, Triage, Investigation, and Response

Security teams do not need more data. They need to turn data into decisions.

That challenge sits at the center of modern SOC operations. Alert volume continues to outpace analyst capacity. Telemetry, intelligence, and investigation context remain distributed across multiple systems. Even when the right signals exist, teams still struggle to answer the questions that matter most: what requires attention, what justifies escalation, and what can be safely closed.

At the same time, the operating environment has evolved. Data has limits. Detection logic changes continuously. Automation has become required. Emerging models such as agentic workflows are beginning to shift how decisions are made, not just how they are supported.

A strong threat intelligence (TI) program addresses these challenges when embedded directly into SOC workflows. In that model, intelligence does not sit as enrichment or reporting. It operates as a layer that improves how teams prioritize, investigate, and act.

This guide translates that model into practical, repeatable workflows aligned to how modern SOCs operate.

WHY SOC TEAMS NEED STRUCTURED THREAT INTELLIGENCE WORKFLOWS

Most SOC teams already have the core components of a security program: detections, alerts, enrichment, and case management. The missing element is a consistent workflow that connects those components into a reliable decision system.

Without that structure, the same issues appear repeatedly:

- Detections generate volume without sufficient context
- Alerts are evaluated in isolation rather than as part of broader patterns
- Threat hunting produces one-time insights instead of improving detection coverage
- Indicator feeds introduce noise when not curated for operational use
- Vulnerabilities are prioritized based on severity instead of exploitability and asset relevance
- Intelligence requirements exist conceptually but remain disconnected from operations

A structured threat intelligence workflow closes these gaps by creating a repeatable path from question to evidence to action.

The objective remains straightforward: improving decision quality and consistency across the SOC.

THE INTELLIGENCE-LED SOC WORKFLOW

Effective SOC operations follow a continuous, intelligence-driven process. Detection, triage, investigation, and response operate as connected stages within a single workflow.

1. Start with a clear intelligence question
2. Build a shared operating view for telemetry, intelligence, and context
3. Use detections to surface meaningful signals
4. Add entity context before prioritization
5. Triage alerts through a structured decision process
6. Investigate using surrounding telemetry and evidence
7. Apply historical context to validate significance
8. Turn threat hunting into a repeatable process
9. Promote validated logic into ongoing detections
10. Curate indicators before operationalizing them
11. Use IOC-driven alerting when confidence is high
12. Prioritize vulnerabilities with threat and asset context
13. Treat intelligence requirements as continuous workflows
14. Escalate only when evidence supports it
15. Build a feedback loop

This workflow operates as a closed loop. Each investigation strengthens detection quality, prioritization, or response speed over time.

1. START WITH A CLEAR INTELLIGENCE QUESTION

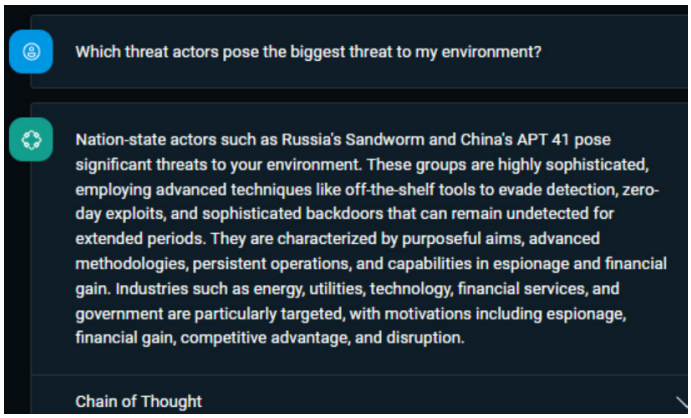

The most effective workflows begin with a focused question rather than a general review of alerts or dashboards.

Examples include:

- Review detections and alerts in a single environment
- Pivot into entities, raw logs, and surrounding telemetry
- Correlate activity with intelligence and known techniques
- Move from triage to investigation without losing context

The question defines:

- What the team needs to determine
- What evidence supports that determination
- What output the workflow produces
- Who owns the outcome

In more mature environments, this becomes a priority intelligence requirement (PIR), allowing the same question to run as a recurring workflow.

Analyst actions

- Write the question in one sentence
- Define the objective (detection, investigation, prioritization, or response)
- Assign ownership
- Set a review window
- Define the expected output

2. BUILD A SHARED OPERATING VIEW FOR TELEMETRY, INTELLIGENCE, AND CONTEXT

Once the question is defined, analysts need a complete operational view.

Workflows degrade when alerts, telemetry, and intelligence remain separated. Analysts require the ability to:

- Review detections and alerts in a single environment
- Pivot into entities, raw logs, and surrounding telemetry
- Correlate activity with intelligence and known techniques
- Move from triage to investigation without losing context

Fragmentation introduces delay and increases the chance of missed evidence.

Analyst actions

- Confirm required log sources for the use case
- Validate identity, endpoint, network, and cloud context
- Ensure pivoting from alert to raw data
- Document minimum data requirements for repeatable workflows

3. USE DETECTIONS TO SURFACE MEANINGFUL SIGNALS

Detections translate telemetry into operational signals. Their purpose centers on identifying activities that warrant attention.

Effective detections:

- Identify suspicious or noteworthy behavior
- Align to known tactics or risk themes
- Support prioritization rather than noise generation

Detection logic requires continuous iteration, tuning, and refinement based on investigation outcomes and intelligence.

Analyst actions

- Review detections by severity and trend
- Map high-value detections to known techniques
- Analyze rule logic and underlying data sources
- Tune detections that generate volume without investigative value
- Promote reliable logic into repeatable workflows

4. ADD ENTITY CONTEXT BEFORE PRIORITIZATION

Alert-centric workflows overlook the broader context of the user, host, or application behind the activity.

Entity-centric analysis evaluates behavior in context:

- Comparison to baseline activity
- Comparison to peer groups
- Contributing detections to current risk
- Pattern development over time
- Impact based on role or access level

This step determines whether an alert reflects routine behavior or meaningful risk.

Analyst actions

- Review the associated entity and its current risk profile
- Examine contributing detections
- Compare behavior to baseline and peer group
- Evaluate anomalies in authentication, access, or data transfer
- Assess impact based on role or privilege

5. TRIAGE ALERTS THROUGH A STRUCTURED DECISION PROCESS

Triage consumes most SOC time. Consistency drives efficiency.

Each alert review should determine:

- What occurred
- Who or what is affected
- Whether the behavior repeats
- Whether it aligns to known techniques
- Whether intelligence increases confidence
- Whether escalation is warranted

Analyst actions

1. Confirm severity, status, and timing
2. Review associated entities and risk
3. Check related detections and technique alignment
4. Pivot into underlying event data
5. Expand to surrounding telemetry
6. Decide whether to close, investigate further, or escalate

Structured triage reduces variability between analysts and improves throughput.

6. INVESTIGATE WITH SURROUNDING TELEMETRY AND EVIDENCE

When triage indicates deeper review, investigation remains within the same operational context.

Analysts should move from alerts into:

- Parsed events for rapid understanding
- Raw logs for validation and detail
- Related entities and behaviors
- Historical activity

Investigation establishes:

- What occurred
- When it began
- Whether it spread or repeated
- What evidence supports the hypothesis
- What response follows

Analyst actions

- Review the underlying event set
- Start with summaries, then validate with raw data
- Search for related activity across entities and time
- Capture supporting evidence
- Define criteria for escalation

7. USE HISTORICAL CONTEXT TO DISTINGUISH ROUTINE ACTIVITY FROM RISK

Many behaviors only become meaningful when viewed over time.

Examples:

- Large data transfers align with specific roles
- Repeated low-level activity indicates persistence
- New techniques appear in historical data

Historical analysis allows teams to determine whether activity represents:

- New behavior
- Persistent activity
- Anomaly
- Broader trend

Analyst actions

- Compare current activity to historical behavior
- Evaluate recurrence over time
- Compare against peer activity
- Re-run detections against historical telemetry

8. TURN THREAT HUNTING INTO A REPEATABLE PROCESS

Threat hunting produces value when it generates repeatable outcomes.

A practical workflow:

- Define a behavior or technique
- Test against telemetry
- Validate findings
- Determine significance

The objective centers on improving detection coverage.

Analyst actions

- Define a scenario or hypothesis
- Execute queries against retained data
- Review detailed evidence
- Determine whether activity appears benign or suspicious
- Identify coverage gaps

9. PROMOTE VALIDATED LOGIC INTO ONGOING DETECTIONS

When investigations or hunts produce reliable logic, that logic moves into operations.

This includes:

- Converting queries into detection rules
- Assigning severity and thresholds
- Mapping to relevant techniques
- Integrating into alert workflows

Analyst actions

- Save validated logic
- Define thresholds and scheduling
- Assign ownership
- Route results into triage workflows

This ensures past effort reduces future workload.

10. CURATE INDICATORS BEFORE OPERATIONALIZING THEM

Indicators vary in quality and require filtering before use.

Effective curation includes:

- Filtering by recency and active status
- Applying confidence thresholds
- Validating source quality

A practical model:

- High-confidence indicators move to enforcement
- Medium-confidence indicators require analyst review
- Lower-confidence indicators remain for context

Without curation, indicator-driven workflows introduce noise and reduce trust.

Analyst actions

- Filter indicators by age and status
- Apply confidence thresholds
- Validate sources
- Convert curated sets into rules
- Monitor operational impact

11. USE IOC-DRIVEN ALERTING WHEN CONFIDENCE IS HIGH

IOC-based alerting delivers value when intelligence remains reliable and current.

Alerts must include sufficient context:

- Host and process details
- File paths and hashes
- Associated entity

Analyst actions

- Verify indicator confidence and relevance
- Review supporting context
- Determine response
- Maintain or suppress indicators as needed

12. PRIORITIZE VULNERABILITIES WITH THREAT AND ASSET CONTEXT

Vulnerability prioritization improves when aligned to real-world risk.

Relevant factors include:

- Active exploitation
- Exploit probability
- Asset exposure
- Business impact

This approach focuses remediation on meaningful risk.

Analyst actions

- Validate asset context
- Review threats associated with vulnerabilities
- Rank based on exploitability and impact
- Prioritize remediation accordingly

13. TREAT INTELLIGENCE REQUIREMENTS AS CONTINUOUS WORKFLOWS

Priority intelligence requirements operate as active workflows rather than static documentation.

They define:

- The question
- Relevant data sources
- Analytical method
- Expected output
- Ownership and frequency

Operationalizing these workflows reduces manual effort and standardizes analysis.

14. ESCALATE ONLY WHEN EVIDENCE SUPPORTS IT

Not every alert should become an incident. Escalation depends on evidence.

Promote to incident when:

- Multiple detections align
- Entity risk rises to a meaningful level
- Behavior shows progression
- Intelligence increases confidence

Close or suppress when:

- Activity reflects normal behavior
- Detection logic remains too broad
- Indicators lack relevance
- Evidence does not support the hypothesis

Analyst actions

- Document evidence reviewed
- Record decision and rationale
- Identify follow-up actions

15. BUILD THE FEEDBACK LOOP

Threat intelligence workflows improve through iteration.

Each investigation contributes to:

- Detection tuning
- New detection logic
- Threshold adjustment
- Improved data requirements
- Refined intelligence workflows

After-action questions

- Should detection logic change?
- Should this logic move into an alert?
- Should thresholds adjust?
- Do additional data sources need to be onboarded?

The goal centers on continuous improvement in coverage, prioritization, and response.

A NOTE ON AUTOMATION AND EMERGING MODELS

Automation supports well-defined, high-confidence tasks such as enrichment, indicator enforcement, and case management. Its effectiveness depends on the quality of underlying data and decision logic.

Emerging models, including agentic workflows, extend this capability by assisting with query generation, summarization, and identification of relevant pivots across telemetry and intelligence.

These models operate across different execution patterns:

- Human in the loop (HITL): analysts validate decisions before action
- Human on the loop (HOTL): analysts supervise automated decisions and intervene when required
- Human out of the loop (HOOTL): systems execute decisions autonomously within tightly controlled scenarios

Most SOC environments operate across all three. Decision delegation depends on confidence, data quality, and operational maturity.

TURNING INTELLIGENCE INTO DECISIONS

A threat intelligence program creates value when it improves daily SOC operations.

That includes:

- Identifying stronger signals
- Triaging more efficiently
- Investigating with full context
- Validating risk using historical data
- Operationalizing high-confidence findings
- Prioritizing action based on real-world risk

The objective remains unchanged: consistent, evidence-based decisions executed at scale.

HOW ANOMALI MAKES THIS OPERATIONAL

Implementing effective threat investigation workflows like these only works when threat intelligence, security telemetry, investigation context, and analyst actions come together in the same operating model. That is where Anomali fits. Anomali's Agentic SOC Platform brings together a unified security data lake, threat intelligence, and agentic AI in a single experience so teams can move more easily from detection to investigation to response.

Anomali Unified Security Data Lake gives teams a place to centralize and retain large volumes of security telemetry while keeping that data fast and searchable for investigation and threat hunting, enabling the ability to search years of data quickly enough to support real investigations rather than passive storage.

On the intelligence side, ThreatStream Next-Gen is designed to continuously enrich security telemetry with curated, confidence-scored threat intelligence so analysts can connect activity to attackers, campaigns, infrastructure, and intent. That helps sharpen detections, accelerate investigations, and make intelligence more useful in day-to-day triage rather than treating it as a separate research function.

AI also plays a practical role in making workflows more usable at scale. AI-guided workflows, natural language interaction, and analysis acceleration across detection, investigation, threat hunting, and intelligence distribution.

Anomali Agentic AI supports tasks like creating and refining detections, summarizing alert context, exploring historical data more quickly, and structuring recurring intelligence workflows around analyst priorities.

That operational model extends into intelligence distribution and prioritization as well. Anomali Integrator is built to automate and orchestrate the distribution of relevant intelligence across the security stack, which supports indicator curation and enforcement workflows. ThreatStream Next-Gen PIRs are positioned as a way to turn priority intelligence requirements into ongoing intelligence workflows, which maps directly to repeatable question-driven processes.

TURNING INTELLIGENCE INTO DECISIONS

A threat intelligence program creates the most value when it improves how the SOC works every day.

That means helping analysts identify stronger signals, triage faster, investigate with better context, validate threats across historical data, operationalize high-confidence intelligence, and prioritize action based on measurable risk. It also means building workflows that can be repeated, tuned, and improved over time.