ClearSkies™

CORE FUNCTION BRIEF

Intelligent Alert Routing

A core function of the ClearSkies iISOC platform

The right alert, to the right analyst, without a human triaging the queue by hand.

Contents

SECTION 1

Executive Summary

IN ONE SENTENCE.

Intelligent Alert Routing is the core function that turns a scored alert into a governed assignment, matching analyst tier to alert difficulty and balancing workload in real time, so who works an alert is decided by policy rather than by whoever happens to be triaging the queue.

In a conventional security operation a person decides who works each alert. That single dependency creates a predictable pattern: experienced analysts spend their day sorting rather than investigating, high-risk items sit in a shared queue behind noise, and when volume spikes there is no elastic capacity to absorb it. Routing by hand scales with neither alert volume nor analyst headcount, and the same alert can land differently depending on who is on shift.

Intelligent Alert Routing replaces the person with a policy. A scored alert is mapped to the analyst tier its difficulty requires, matched against the least-loaded eligible analyst at that tier, and assigned under rules that resolve ties the same way every time. When every eligible analyst at a tier saturates against its service-level threshold, eligible lower to medium complexity alerts route to the platform's autonomous handling capacity instead of waiting, and route back to human analysts as capacity recovers.

Every decision the function makes is visible and reconstructable: which tier an alert was mapped to, which analyst was eligible, and why the assignment landed where it did. Nothing about the decision depends on memory or on which lead happened to be on shift when the alert arrived.

THE FUNCTION IS PRECISELY SCOPED.

Intelligent Alert Routing decides who works an alert. It does not decide how urgent the alert is, which is the Risk Scoring Formula. It does not investigate the alert, which is Generative and Agentic AI. And it does not decide or execute a response, which is SOAR and Playbooks. Each of those is a separate core function with its own brief.

Five stages
Tier, workload, assignment, overflow, state

Three tiers
Matched to alert difficulty

Elastic
Surge absorbed without added headcount

SECTION 2

Why Manual Triage Falls Short

Assignment is underinvested relative to the rest of the security operation, and the consequences compound quietly rather than surfacing in an incident review. Four failure modes explain why, and each is answered by a stage of the mechanism in section 3.

2.1 Skilled analysts consumed by sorting

A senior analyst who triages the queue by hand spends the first minutes of every shift deciding what to work rather than working it. That time is not recovered, and it scales with alert volume rather than with the value the analyst produces once actually investigating.

2.2 High-risk alerts wait behind noise

A shared queue does not differentiate by difficulty or urgency once an alert lands in it. A critical alert filed a minute after a routine one can wait behind it if nobody has yet looked at either, and the delay is invisible until someone happens to notice.

2.3 No headroom under surge

When volume spikes, a fixed analyst headcount has nowhere to send the overflow. Every alert still needs a human decision about who works it, so the same triage bottleneck that slows a quiet day becomes a missed commitment on a busy one.

2.4 Assignment that varies by shift

Two identical alerts, triaged by two different leads, can reach different tiers or different analysts, because the decision rests on individual judgment rather than on a consistent policy. Reviewing why a specific alert reached a specific analyst becomes a question about who was on duty rather than a question the system can answer on its own.

2.5 The cost of the workaround

The usual response is a dedicated triage lead or a rotating duty roster, a role whose entire function is to keep the queue moving rather than to investigate anything in it. The cost scales with alert volume and with the number of shifts covered, while the underlying problem, an assignment decision that depends on a person rather than a policy, is untouched. The workaround also caps how far the operation can grow: adding tenants or alert sources means adding triage capacity in lockstep, so growth is bounded by how fast people can be hired and trained to sort rather than to investigate.

Industry research on security operations consistently ties manual triage overhead to lost investigation time and to analyst attrition, though the specific figures are not yet sourced for this brief.

SECTION 3

How Intelligent Alert Routing Works

Intelligent Alert Routing is one continuously operating mechanism native to the TDIR core, not a queue watched by a person. A scored alert enters at one end and an auditable assignment leaves at the other, with the load picture updated in real time throughout.

> *Figure 1. The five-stage assignment mechanism.*

3.1 The five stages

Each stage is named once here and referenced by name throughout the rest of this brief, paired with the failure mode from section 2 that it answers.

STAGE	WHAT HAPPENS	THE FAILURE MODE IT ANSWERS
Tier mapping	The score band an alert carries from the Risk Scoring Formula is resolved against the tenant's analyst roster and skill levels, so low-score alerts map to L1, mid-score to L2, and top-score or critical alerts to L3.	Assignment that varies by shift
Workload awareness	The pending alert count for every analyst eligible at the required tier is tracked continuously, rather than assumed from a snapshot.	Skilled analysts consumed by sorting
Assignment decision	The eligible analyst with the fewest pending alerts receives the assignment. Where loads are equal, a defined tie-break rule, round robin or longest idle, resolves it the same way every time.	High-risk alerts wait behind noise
SLA overflow routing	When every eligible analyst at a tier saturates against its service-level threshold, eligible lower to medium complexity alerts route to the platform's autonomous handling capacity instead of a human queue.	No headroom under surge
State tracking	Pending counts and pool composition update in real time as alerts are assigned, resolved or reassigned, so the next decision is made on the current picture rather than a stale one.	Skilled analysts consumed by sorting

The mechanism resolves one decision at a time, tier then load then tie-break, rather than weighing every factor at once. That ordering is deliberate: difficulty match is never traded away to balance load, and load is never ignored once difficulty is matched.

3.2 How an assignment is judged

Quality is measured rather than assumed, on definitions a provider can report and a customer can audit. The measures below are reported natively and per tenant. They are defined here because the same metric name is measured differently by different vendors, and target values are agreed per engagement rather than asserted.

MEASURE	THE DEFINITION USED	WHAT A POOR VALUE TRIGGERS
Time to assignment	Elapsed time from a score and band arriving to an assignment being made.	Investigation of the specific stage where the latency is introduced.
Tier match accuracy	Assignments requiring no manual override to reach the correct tier, as a share of all assignments.	A review of the score-to-tier thresholds for the affected alert type.
Load balance variance	The spread of pending-alert counts across eligible analysts within a tier.	A review of tie-break rules, or of roster capacity at that tier.
SLA attainment under surge	Alerts meeting the service-level response target during a saturation event, as a share of all alerts in that event.	A review of tier thresholds and roster size against the observed peak.
Autonomous overflow share	Alerts routed to autonomous handling during saturation, as a share of all alerts in that period.	Nothing on its own. It is read as evidence that overflow is engaging as designed.

Overflow is a designed state, not a failure state. ***A tenant whose queues never saturate is not necessarily healthier than one whose overflow routing engages occasionally. Overflow engaging as designed, absorbing eligible alerts and releasing them back to human analysts as capacity recovers, is the mechanism protecting the service commitment while it works.***

SECTION 4

Intelligent Alert Routing in the Platform

4.1 What it draws on, what it writes back, and what it does not do

What it draws on. The score and urgency band the Risk Scoring Formula publishes for each alert, the tenant's analyst roster and skill levels, and the service-level thresholds the KPIs and SLAs core function defines. Each is an input to a specific stage rather than a general capability the function generates itself.

What it writes back. The assignment decision and its audit trail, published into the shared model. The assigned analyst tier or the platform's autonomous handling capacity acts on it, and KPIs and SLAs consumes the queue and assignment timing for reporting. The relationship is one way: nothing returned to Intelligent Alert Routing changes an assignment already made, and the measures in section 3.2 come from the mechanism's own decision history rather than from a signal a downstream function engineers for the purpose.

What it does not do. It does not compute a risk score, and it does not investigate or act on an alert once assigned. Its contribution is not a new signal or a new action but a consistent, auditable answer to who should work an existing, already-scored alert.

4.2 A worked example

The comparison below is illustrative and technically representative rather than drawn from a named engagement.

WHAT MANUAL TRIAGE DOES	WHAT INTELLIGENT ALERT ROUTING RESOLVES
A shift lead scans the queue periodically and assigns alerts by impression, in the order noticed rather than in order of urgency.	Tier mapping resolves the alert's score band the moment the score arrives, before any human has looked at the queue.
The lead checks who seems free, often from memory rather than from a current count.	Workload awareness tracks the pending count for every eligible analyst continuously.

WHAT MANUAL TRIAGE DOES	WHAT INTELLIGENT ALERT ROUTING RESOLVES
Two analysts with similar loads both seem available, and the lead picks one, sometimes the one who is easier to reach.	Assignment decision selects the least-loaded eligible analyst and applies a defined tie-break rule when loads are equal.
A surge arrives. The lead has no capacity to create and starts a waiting list, while the clock keeps running.	SLA overflow routing detects saturation against the threshold and routes eligible alerts to autonomous handling immediately.
The queue picture is only as current as the lead's last scan.	State tracking updates the load picture on every assignment, resolution and reassignment, in real time.

4.3 Delivery across many tenants

Tier thresholds, roster composition and service-level targets are configured per tenant, so the same five-stage mechanism serves every tenant's policy without a bespoke rebuild. What travels between tenants is the routing logic and its configuration; what never travels is one tenant's roster, queue contents or assignment history. A provider can therefore explain any assignment to any customer, from the same mechanism, without a manual audit or a per-tenant routing engine.

Because the function is switched on with the platform rather than deployed separately, time to value is a matter of tuning thresholds to a roster rather than an onboarding project. In the first week, default tier thresholds apply against the imported roster, so assignment begins immediately even before thresholds are tuned. Through the first month, tier match accuracy and load balance variance surface any threshold that needs adjustment for that tenant's alert mix and skill distribution. At steady state, SLA attainment under surge is tracked as an ongoing measure, reviewed whenever roster size or alert volume shifts materially.

SECTION 5

ATT&CK Alignment and Supported Frameworks

Intelligent Alert Routing does not map detection content and does not compute technique relevance itself. Where a technique identifier is present, it travels with the score and band the Risk Scoring Formula supplies, and it is carried into the assignment audit trail so an analyst can see, at the moment of

assignment, the adversary behavior the routing decision was made against.

FRAMEWORK OR STANDARD	ITS ROLE IN INTELLIGENT ALERT ROUTING
MITRE ATT&CK, Enterprise, Cloud and Identity	Carried through from the Risk Scoring Formula's output into the assignment audit trail. It is not computed or validated by this function.
NIST Cybersecurity Framework 2.0	Outcome language used where platform reporting discusses assignment and service-level performance to an executive or board audience.

Mapping is inherited, not asserted. Because this function neither produces nor validates a technique mapping, it makes no coverage claim of its own. Control-level mapping to the frameworks a customer reports against is held by the Regulatory Frameworks core function.

SECTION 6

Questions Raised in Evaluation

The questions below recur in buyer evaluations, each answered directly against the mechanism described in section 3.

QUESTION, AS A BUYER ASKS IT	ANSWER
Does Intelligent Alert Routing decide how urgent an alert is?	No. It consumes the score and band the Risk Scoring Formula computes. Its own decision is which tier and which analyst, not how urgent the alert is.
Does the autonomous handling capacity investigate alerts on its own authority?	Routing decides that an eligible alert moves to autonomous handling when human capacity saturates. The investigation itself is performed by Generative and Agentic AI, under its own governance, not by this function.
What happens when every analyst at a tier is at capacity?	Eligible lower to medium complexity alerts route to autonomous handling. Higher-tier cases wait for the next available analyst under policy rather than being downgraded to a tier they do not fit.
Can routing rules be customized per tenant?	Yes. Tier thresholds, roster composition and service-level targets are configured per tenant, and the same mechanism applies the configuration consistently.

QUESTION, AS A BUYER ASKS IT	ANSWER
Does routing ever override the risk score?	No. Tier mapping follows the score as published. Routing does not re-score an alert or adjust the urgency it was given.
Can a specific analyst be favored or bypassed?	No. Assignment follows tier eligibility, then load, then a defined tie-break rule, which is what removes shift-dependent variation from the decision.

SECTION 7

The Benefits

Policy-driven assignment is where a security operation stops paying for sorting, and the outcomes are countable: how much analyst time goes to investigation rather than triage, whether the commitment holds when volume spikes, and whether any assignment can be explained after the fact.

7.1 At a glance

BENEFIT	WHAT PRODUCES IT
Analyst time goes to investigation rather than sorting	Tier mapping and assignment decision run the moment a score arrives, so no one triages the queue by hand.
A high-risk alert does not wait behind a routine one	The score band, not queue order, determines the tier an alert is mapped to.
Commitments hold when volume spikes	SLA overflow routing absorbs eligible alerts into autonomous handling at saturation, and releases them back as capacity recovers.
The same alert reaches the same class of decision on any shift	A defined tie-break rule and continuous state tracking, rather than a lead's judgment and last scan.
Assignment quality becomes a reportable number	Five measures defined in section 3.2, reported natively and per tenant.

BENEFIT	WHAT PRODUCES IT
Any assignment can be explained after the fact	An audit trail carrying the score, tier and workload in force at the moment the decision was made.

7.2 What changes, and for whom

Policy-driven assignment changes what each role spends time on, and what each role can rely on being consistent.

ROLE	WHAT CHANGES
SOC analyst	Time goes to investigation rather than to sorting, and assignment matches difficulty to skill without waiting on a shift lead.
SOC manager	Assignment quality becomes a reportable, auditable number rather than a judgment call that varies by shift.
CISO and risk	Every assignment carries a defensible rationale, reconstructable after the fact rather than dependent on who was on duty.
Service provider practice lead	Routing capacity scales with the customer base rather than with the number of triage leads on staff, and service commitments hold under surge without over-staffing for peak.

7.3 What sets it apart

▸ **Policy, not memory.** Tier and analyst selection follow a defined mechanism, so the same alert reaches the same class of decision regardless of who is on shift.

▸ **Workload-aware, not tier-blind.** The least-loaded eligible analyst is selected automatically, rather than the first one a triage lead happens to think of.

▸ **Elastic without added headcount.** Saturation routes eligible alerts to autonomous handling instead of breaching a commitment, and releases capacity back as load recovers.

THE VALUE IN ONE LINE.

A security operation without this function pays a person to decide who works each alert, and gets a different answer depending on who that person is. With it, the answer follows a policy, holds under surge, and can be explained afterwards.

SECTION 8

Summary

Intelligent Alert Routing is how the platform decides who works an alert. It is one core function among twelve, scoped to a five-stage assignment mechanism that runs continuously inside the TDIR core rather than as a decision made by whoever is triaging the queue that shift.

▸ Tier mapping, workload awareness, assignment decision, SLA overflow routing and state tracking replace a person deciding who works each alert.

▸ The score and band the function consumes come from the Risk Scoring Formula. Routing never re-scores an alert.

▸ Saturation routes eligible alerts to autonomous handling instead of breaching a commitment, and releases capacity back as load recovers.

▸ Every assignment is auditable, traceable to the score, tier and workload in force at the moment it was made.

▸ The function decides who works an alert, and leaves urgency and response to the core functions built for those decisions.

Forward-looking statement: any statement about future capability is forward-looking and non-binding, and capability described in the present tense is capability available at the publication date, subject to the deployment configuration in force.