ClearSkies™

CORE FUNCTION BRIEF

Detection Factory

A core function of the ClearSkies iISOC platform

Detection content as a governed, measured and continuously improving lifecycle.

Contents

SECTION 1

Executive Summary

IN ONE SENTENCE.

The Detection Factory is the core function that governs the whole life of detection content, authoring, validating, deploying, measuring, tuning and retiring the logic that decides what the platform recognizes as suspicious, so coverage is a maintained and provable asset rather than a rule library that quietly decays.

Detection engineering decides, in advance, what a security operation is able to recognize. In most operations that work ends up as a rule library: logic written by an engineer, added to a console, and left alone until an incident, an audit or a noisy alert forces someone back to it. The library grows by accumulation. Nobody owns its decay, because nobody measures it.

The Detection Factory replaces that model with a production process. Threat intelligence, validated hunt findings and coverage gaps enter at one end. Generative and agentic AI draft candidate logic, tagged with its MITRE ATT&CK technique identifier at the moment of creation. Every candidate is tested against the tenant's own historical telemetry and against known-benign traffic before it goes live. Once live, every detection is measured on precision, fidelity and duplicate rate, and a rule that falls below the governed threshold is tuned or formally retired rather than left running on reputation.

Two consequences follow. Coverage becomes something a buying committee can query rather than take on faith, because every rule carries its technique identifier and its measured quality. And the return on one piece of detection engineering compounds, because a validated detection deploys once against the shared entity model and applies to every tenant sharing the relevant exposure profile.

The function is precisely scoped. The Detection Factory decides what is detected. It does not decide how urgent a firing detection is, which is the Risk Scoring Formula; it does not decide who handles it, which is Intelligent Alert Routing; and it does not decide what happens next, which is SOAR and Playbooks. Each of those is a separate core function with its own brief.

Six-stage lifecycle

Source, author, validate, deploy, measure, tune or retire

Coverage that is provable

A queryable ATT&CK matrix per tenant, not a claim

Compounding return

One authoring event serves every relevant tenant

SECTION 2

Why Detection Content Decays

Detection engineering is underinvested compared with 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 specific stage of the lifecycle in section 3.

2.1 Detection debt

A rule is written once, by a particular engineer, for a particular threat, and then left alone. The environment changes around it, adversaries shift technique variants, and the person who understood the original intent moves on before anyone revisits it. The rule keeps firing, or keeps not firing, and neither outcome is examined.

2.2 Coverage that cannot be proven

Most operations cannot say with confidence which ATT&CK techniques they detect well, which they detect poorly and which they miss entirely. Rule intent lives in prose, in ticket history and in memory rather than in a structure that supports a query. Asked what the organization actually detects, the answer is assembled from recall rather than from evidence.

2.3 Tuning at the pace of headcount

Every noisy or duplicate detection found in triage needs an engineer to open the rule, reconstruct its intent, edit it, test it and redeploy it. That work competes for the same scarce hours as new coverage, and it loses the argument against the next incident. The backlog grows, and the noise it produces is absorbed by analysts instead of being fixed at the source.

2.4 A one-way pipeline from intelligence to content

Threat intelligence and threat hunting produce findings, but a finding has to be turned into a deployable rule by hand, usually in a different tool, on a different schedule, by whichever engineer is free. Days or weeks pass between proving that a technique is present and catching the next occurrence automatically. The organization stays exposed to something it has already confirmed.

2.5 The cost of the workaround

The usual response is more people rather than a different process: engineers hired to keep pace with decay, coverage requests and the tuning backlog. Cost scales with the rule base and with the rate of change in the threat landscape, while the underlying problem, content nobody measures, is untouched.

SECTION 3

How the Detection Factory Works

The Detection Factory is one continuously operating lifecycle native to the TDIR engine, not a set of separate tools. Each stage feeds the next, and the last stage feeds the first.

> *Figure 1. The six-stage lifecycle. Measurement is what closes the loop, because a detection nobody measures cannot be improved or retired on evidence.*

3.1 The six stages

STAGE	WHAT HAPPENS	FAILURE MODE IT ANSWERS
Source	Threat intelligence indicators, validated threat hunting findings and coverage gaps surfaced from normalized telemetry enter as the raw material for new or revised content.	The one-way pipeline
Author	Generative and agentic AI draft candidate logic from that material, each candidate tagged with its ATT&CK technique identifier as native metadata at the moment of creation.	Coverage that cannot be proven
Validate	Candidate logic is tested against historical normalized telemetry from every layer the platform ingests, and against known-benign traffic, so both detection and false-positive behavior are known before promotion.	Detection debt
Deploy	Validated logic is promoted into production across every tenant sharing the relevant exposure profile at once, because every tenant shares the same normalized schema.	The cost of the workaround
Measure	Every live detection carries native quality metrics, tracked per rule rather than for the rule base as a whole.	Coverage that cannot be proven
Tune or retire	A detection below the governed quality threshold is flagged automatically, then tuned and revalidated, or formally retired with its audit trail intact.	Tuning at the pace of headcount

3.2 How a detection is judged

Quality is measured rather than assumed, on three metrics that a provider can report and a client can audit.

METRIC	WHAT IT MEASURES	WHAT A POOR SCORE TRIGGERS
Precision	The share of firings that turn out to be true positives.	Tuning of the logic or its scope, then revalidation before redeployment.
Fidelity	The share of firings that arrive with correct urgency and usable context.	A metadata or enrichment correction, so downstream scoring and routing behave.
Duplicate rate	The share of firings already covered by another rule.	Consolidation, and retirement of the weaker of the two rules.

Retirement is an outcome, not a failure. ***A rule that is closed formally, with its provenance and change history preserved, leaves the coverage matrix accurate. A rule quietly disabled in a console leaves the matrix wrong, which is worse than the noise it made.***

SECTION 4

The Detection Factory in the Platform

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

What it draws on. Normalized telemetry from the iCollector and the entity resolution the Centric-AI Fabric performs, so a candidate rule is written against the same objects every other layer resolves to. Threat intelligence indicators, validated threat hunting findings, and the coverage gaps the platform surfaces from telemetry no existing rule addresses.

What it writes back. Validated, ATT&CK-tagged detection content, and the per-rule quality metrics that come with it, published into the shared model. Risk scoring, alert routing, SOAR and playbooks consume that output. The relationship is one way: nothing is returned to the Detection Factory that changes what it authors, and the quality metrics come from a detection's own firing history rather than from a signal engineered by a downstream function.

What it does not do. It does not decide urgency, ownership or response, and it holds no telemetry of its own. A native add-on contributes cross-layer visibility by adding a signal no other layer can see. The Detection Factory contributes something different in kind: it decides what any of that telemetry is capable of revealing.

4.2 A worked example

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

WHAT A CONVENTIONAL OPERATION DOES	WHAT THE DETECTION FACTORY DOES
A hunter confirms a technique is present. The finding is written up and queued for an engineer.	The finding enters the lifecycle at Source, with no separate queue and no hand-off.
An engineer drafts a rule weeks later, scoped to the exact indicator, and tests it against whatever sample data is at hand.	Logic is drafted against the underlying technique rather than the single indicator, tagged to its ATT&CK identifier, and validated against ninety days of normalized cross-layer telemetry.
The rule is deployed to the tenant where the finding originated. Other tenants stay exposed until a similar finding turns up there.	Deploy promotes the validated rule across every tenant sharing that exposure profile in a single event.
The resulting alert reaches triage unscored and unrouted, waiting on a person to establish urgency and owner.	The alert reaches risk scoring and alert routing already ATT&CK-tagged and carrying the rule's precision history.
A noisy or duplicate rule is found months later in a rule-base review, if a review happens at all.	Measure flags the rule on its own metrics within weeks, and Tune or retire acts on it as routine work.

4.3 Delivery across many tenants

A detection validated in one tenant is generalized to the technique and deployed to every tenant whose exposure profile it covers, so detection engineering capacity stops tracking headcount and starts tracking the size of the client base. Tenant separation is unaffected: what travels between tenants is the

logic and its metadata, never the telemetry that produced it, and no tenant sees, holds or infers the data of another. A provider can therefore answer the coverage question per client, from the same matrix, without a manual audit per client.

SECTION 5

ATT&CK Alignment and Supported Frameworks

MITRE ATT&CK is the detection spine, and the Detection Factory is the function that owns the distinction the platform draws throughout its documentation: a technique listed in a coverage matrix is claimed, and a technique exercised in a purple-team run or against MITRE Engenuity ATT&CK Evaluations is demonstrated. Both are reported, and they are never presented as the same thing.

Detection content is mapped across the Enterprise, Cloud and Identity matrices. A technique is cited by identifier and name on first use, for example T1566.002 Spearphishing Link, and by identifier alone thereafter. [CLEARSKIES INPUT: confirm the matrix version in force and the refresh cadence]

LIFECYCLE STAGE	HOW ATT&CK IS USED
Author	Every candidate carries its technique identifier, or identifiers, as native metadata rather than as documentation added afterwards.
Validate	The mapping is checked against the matrix version in force before a candidate is promoted.
Measure	Precision and fidelity are reported per technique as well as per rule, so a weak area of the matrix is visible rather than averaged away.
Tune or retire	A change of rule state updates the coverage matrix immediately, so the matrix describes what is live rather than what was once written.

FRAMEWORK OR STANDARD	ROLE IN THE DETECTION FACTORY
MITRE ATT&CK, Enterprise, Cloud and Identity	Technique mapping for every detection, and the basis of the per-tenant coverage matrix.
MITRE Engenuity ATT&CK Evaluations	Independent validation of detection behavior against an emulated adversary, and the basis of the demonstrated claim.
MITRE D3FEND	Countermeasure vocabulary, cross-referenced where a detection maps to a defensive technique.

FRAMEWORK OR STANDARD	ROLE IN THE DETECTION FACTORY
Sigma	Portable detection format for exchange with other tooling, in both directions.
STIX and TAXII	Structured intelligence intake at the Source stage.

Alignment to OCSF and OpenTelemetry, and the syslog, CEF, LEEF and REST API ingestion formats, are properties of the collection layer and are described in the platform brief rather than restated here.

SECTION 6

Business Impact and Delivery

The Detection Factory reports the measures below natively and per tenant. Definitions are published because the same metric name is measured differently by different vendors. Target values are agreed per engagement rather than asserted here.

MEASURE	DEFINITION USED	REPORTED AS
Technique coverage	Techniques with at least one live, validated detection, as a share of the techniques in scope for the tenant.	A matrix per tenant, queryable at any time.
Precision per rule	True positives as a share of all firings for that rule.	Trended per rule and per technique.
Duplicate rate	Firings already covered by another live rule, as a share of firings.	Per rule, with the consolidation candidates named.
Time from source to deployment	Elapsed time from a finding or indicator entering Source to the validated detection being live.	Trended, and reported per tenant.
Retirement rate	Rules formally retired in the period, as a share of the live rule base.	Trended, with the reason recorded.

6.1 What changes, and for whom

ROLE	WHAT CHANGES
Detection engineer	Authoring and routine tuning move into the lifecycle, so the engineer's time goes to custom requirements and to reviewing what the lifecycle proposes.

ROLE	WHAT CHANGES
SOC analyst	Fewer duplicate and low-precision alerts reach triage, because the rules that produce them are flagged and fixed at the source.
SOC manager	Coverage and detection quality become reportable numbers rather than an assessment gathered by asking the team.
CISO and risk	The coverage question has an evidenced answer, per technique, with claimed and demonstrated distinguished.
MSSP practice lead	Detection engineering capacity scales with the client base rather than with the number of engineers on staff, and per-tenant coverage supports contractual reporting.

6.2 What sets it apart

▸ **Authored, not imported.** Logic is drafted and validated against the tenant's own normalized telemetry rather than adapted from a generic content pack written for an environment nobody involved has seen.

▸ **ATT&CK mapping as native metadata.** The technique identifier is attached when the rule is created, which is what makes the coverage matrix a query rather than a documentation exercise.

▸ **Coverage measured, not asserted.** Precision, fidelity and duplicate rate are tracked per rule, so a claim about detection quality can be produced from the platform rather than from a conversation.

▸ **One authoring event, many tenants.** A validated detection reaches every tenant sharing the exposure profile without separate engineering per tenant, and without any tenant data moving.

6.3 Objection handling

QUESTION, AS A BUYER ASKS IT	ANSWER
Is the Detection Factory a separate product?	No. It is one of the twelve core functions of the TDIR engine, scoped to the lifecycle of detection content. There is no separate license, no separate console and no separate deployment.
Does it replace a detection-engineering team?	It changes what that team spends time on. Authoring, validation and routine tuning run in the lifecycle; the organization's engineers remain the authority for custom requirements and for reviewing what the lifecycle proposes.
How is a claimed detection different from a demonstrated one?	A technique listed in the coverage matrix is claimed. A technique exercised in a purple-team run or validated against MITRE Engenuity ATT&CK Evaluations is demonstrated. Reporting states which of the two applies.

QUESTION, AS A BUYER ASKS IT	ANSWER
What happens to a detection that stops performing?	Measure tracks precision, fidelity and duplicate rate per rule. A rule below the governed threshold is flagged for tuning or formal retirement rather than left live on the strength of having once been useful.
Does a detection written for one tenant help other tenants?	Where tenants share the relevant exposure profile, yes. The logic and its metadata deploy across them in one event. No tenant telemetry moves, and no tenant sees another tenant's data.
Does it receive anything back from risk scoring, routing, SOAR or playbooks?	No. The relationship is one way. The Detection Factory publishes content and quality metrics; those functions consume them. Quality measurement comes from a detection's own firing history.

SECTION 7

Summary

The Detection Factory is how the platform produces and maintains detection content. It is one core function among twelve, scoped to the lifecycle that decides what the platform is capable of recognizing, and it runs continuously inside the TDIR engine rather than as a bundle of separately branded capabilities.

▸ Content is authored and validated against the tenant's own normalized telemetry, not imported from a generic file.

▸ Every rule carries its ATT&CK technique identifier as native metadata from the moment it is created.

▸ Coverage and quality are measured per rule and per technique, so both can be queried rather than asserted.

▸ A validated detection deploys once and applies to every tenant sharing the relevant exposure profile.

▸ The function decides what is detected, and leaves urgency, ownership and response to the core functions built for those decisions.