


AIRMDR

Grade Your MDR

Case Quality Scoring Toolkit

A practical framework for evaluating
the quality of your MDR's investigations.



The Hidden Foundation of the SOC

This document and the companion spreadsheet describe and use a scorecard and checklist to make case quality measurable and actionable. This toolkit is the manual version of the scoring approach that our open-source tool, SOC Grader applies automatically and at scale.

Toolkit Contents



Implementation Guide

How to turn case reviews into a repeatable quality loop: sample → score → roll up → act.

Checklist

The observable things a high-quality case should contain so that a case reviewer can independently verify the work.

Scorecard

A method that converts the checklist into a repeatable 0–10 score for each case, plus flags to note when key evidence is missing.

Interactive Scorecard spreadsheet

This document explains what gets scored and what “good” looks like. Use the accompanying spreadsheet template to automate calculations described in this document.

Use the spreadsheet locally with your own cases; no case data needs to be uploaded.

Anatomy of a quality SOC case

A quality SOC case turns noisy telemetry data into a decision you can trust – and a next step you can take. The four key components of a quality case include:

1

Investigation plan

What you set out to validate and why – your hypothesis.

2

Evidence + context

The specific telemetry and context that supports (or refutes) the hypothesis.

3

Conclusion + confidence

A clear disposition and why you are confident.

4

Next action

What should happen next (contain, monitor, notify, remediate) and who owns it.

Implementation Guide: How to use this Toolkit

Score a sample of cases. For example, select 10 recent cases a single MDR provider delivered or that are destined for a single SOC queue. Do not pick only your best cases – true near-random sampling reveals a more accurate quality level.

Goal Case reviews become a repeatable quality-assurance loop. Consistent scoring reveals recurring gaps and provides evidence to improve internal processes or hold your MDR provider accountable.

Method

1 Pick a sample

Choose 10 recent cases from the same provider or workflow.

2 Score each case

For each item in the Case Quality Checklist, mark it: Complete / Partial / Missing and note where the evidence appears in the case.

Important Note: You'll want to adapt generic labels in the Case Quality Checklist below to specific alerts. The checklist uses universal cybersecurity concepts (e.g., "Infrastructure Trust"), but the specific evidence required changes based on the alert type. You must adapt these categories to the scenario.

- ↳ Example 1: Login Alerts. For these alerts, "Infrastructure Trust" in the Case Quality Checklist means checking VPN logs, Geolocation, and ASN reputation.
- ↳ Example 2: Malware/Network Alerts: For these alerts, "IoCs & threat context" means checking for Command & Control traffic patterns and DNS requests.

3 Roll up

Compute an average score and list the most common critical and minor gaps.

4 Act

Hold your MDR provider accountable. Determine whether the gaps point to an internal process change, a provider remediation plan, or a provider change.

Case Quality Checklist

Use this table to review each case and determine whether it contains the minimum auditable decision context.

Investigation Check	Case Requirement	Tier	Weight
Alert trigger stated	States the specific condition/event pattern that caused the alert to fire.	Critical	5
Trigger validated with telemetry	Validates the trigger using primary telemetry (log excerpts, query output, screenshots, or equivalent).	Critical	5
Actor identity verified	Identifies the actor (user/service/device/workload) and confirms legitimacy (expected owner, role, and context).	Critical	5
Compromise check performed	Checks for direct compromise indicators for the actor/source (auth anomalies, prior alerts, suspicious sequences) and records results.	Critical	5
Scope enumerated	Enumerates affected assets/accounts/resources and confirms scope boundaries (what is and is not impacted).	Critical	5
Impact checked	Confirms whether sensitive actions/data access occurred (or documents that it was checked and not found).	Critical	5
Baseline comparison	Compares activity to historical baseline (new vs normal) for the entity involved.	Important	3
Timeline/sequence provided	Provides a short timeline (before/during/after) so the sequence of events is verifiable.	Important	3
Infra trust validated	Validates source location/path trust (VPN/corp ranges, geo/ASN, known infra) when network/access is part of the alert.	Important	3
Access path legitimacy	Confirms access method and posture (app, auth method, device posture, MFA) match expected behavior.	Important	3
IoCs/TI checked	Checks relevant artifacts (IPs/domains/files/URLs/hashes) against reputation/TI and records outcomes.	Important	3
Disposition tied to evidence	States a clear disposition/conclusion and explains why the available evidence supports it.	Important	3
Next action & owner	Documents the next step (contain/monitor/notify/remediate) and who owns it; includes deadlines if time-sensitive.	Important	3
Reproducible sources/queries	Includes references to data sources or queries run (enough for another analyst to reproduce key validations).	Nice-to-have	1
Additional context	Adds secondary enrichment (peer comparisons, broader hunting, related cases) when it materially increases confidence.	Nice-to-have	1

Some checks may not apply to every alert type. When a case explicitly establishes that a check is not applicable, mark it complete rather than leaving it unscored.

Legend

Investigation Check: The specific validation performed in the case.

Case Requirement: Must be explicit and verifiable in the write-up.

Tier / Weight: Critical=5, Important=3, Nice-to-have=1

Action Item: Yes only if the Grade is Partial (i.e., 1) or Missing (i.e., 0) and explicitly called out as a next step.

Scorecard: How to Grade

The scorecard helps convert the checklist into a repeatable numeric score, while still preserving evidence-based review.

STEP 1

Grade each Case Quality Checklist item using the evidence in the case

Label	Meaning	Value
Complete	The case explicitly addresses the Case Requirement with verifiable evidence.	2
Partial	The case touches on the Case Requirement but misses key validation or evidence.	1
Missing	The Case Requirement is not addressed in the case.	0
Action item	If Partial or Missing, record whether it is explicitly listed as an Action Item/Next Step. That's because a case can be incomplete but still operationally responsible if it clearly documents what remains to be done (especially for items that are MDR customer responsibilities, or where the analyst lacks access).	Yes/No

STEP 2

Compute the weighted score and normalize to 0-10

These calculations are explained here, and captured as formulas in the associated spreadsheet. Weighted points earned = Item Weight × (Grade / 2). Final score = weighted points earned ÷ total possible weighted points × 10.

Total points possible	SUM (Checklist Item Weights)
Weighted points earned	SUM (Checklist Item Weight × (Grade Value / 2))
Normalized score (0-10)	(Weighted points earned / Total points possible) × 10

Next Steps

Use the spreadsheet included with this kit to grade 10 recent cases. The spreadsheet applies the case-quality methodology described in this document – and it's the same methodology used by SOC Grader, AirMDR's automated scoring tool.

Prefer automated scoring?

AirMDR built the tool to help accelerate and scale the scoring process, and **it's free to use**.

Try the Tool