INTEZER

# The 2026 MDR Renewal Checklist

## Financial Services Industry

This checklist is built for CISOs and security leaders at banks, insurance companies, investment firms, credit card companies, and fintech organizations who are evaluating or renewing their MDR.

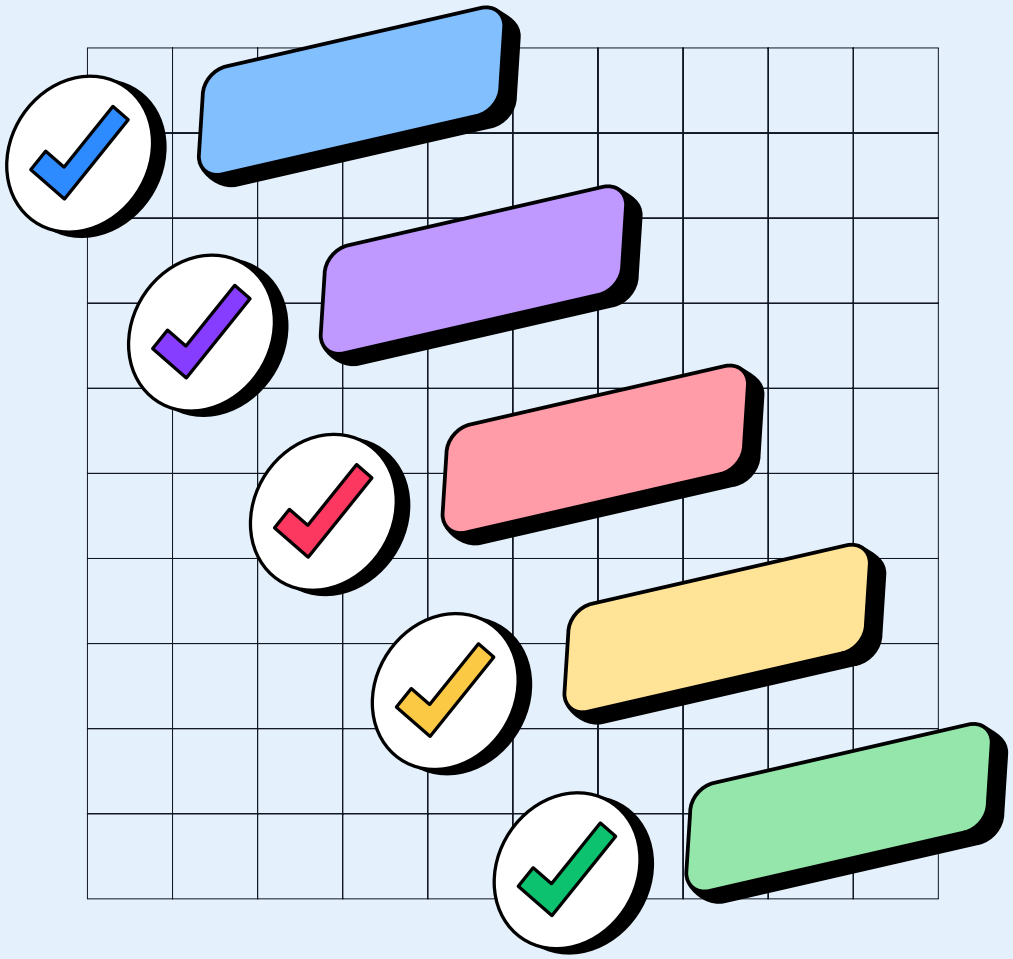

Your MDR investigated 40% of your alerts but your regulator holds you accountable for 100%. Use this checklist before your next renewal to assist you in evaluating your MDR with a focus on coverage gaps, regulatory risk accountability, and the case for AI SOC in financial services.

# Is your MDR still solving the right problem?

Financial services is the most targeted industry in the world. On a global scale, one in six cyber investigations involves a financial services organization. Industry reports indicate that ransomware demands against FSI firms average $900,000 with median payouts in this high-stakes sector often climbing as high as $2 million to $3 million.

That threat exposure, combined with regulatory pressure and security teams too small to monitor alerts around the clock, is exactly why mid-market financial services firms outsourced MDR. It was sound logic to let a managed service handle triage and investigation 24/7 so your team can focus on what actually needs a human decision.

MDRs are still doing what they always did: triaging higher-priority alerts to the extent their human analysts could scale, typically around 40% of all alerts received. But attackers are using AI to craft activity that generates mostly low-severity or informational alerts, knowing those are the ones no one opens. Intezer's research across 25 million alerts found that nearly 1% of real threats originate in low-severity alerts. For an average enterprise, this translates to approximately 54 real threats per year. More than one per week, every week, going uninvestigated.

Furthermore, most MDR contracts are scoped around endpoint detection. Identity, cloud, email, and network telemetry often are not included. However, in 2026, credential-based attacks, account takeovers, and cloud misconfigurations are where many breaches have been starting. Not just on the endpoint.

Financial services firms carry an additional weight here. NCUA, NAIC, FFIEC, and GLBA all hold the institution accountable for what was investigated and documented, regardless of what an MDR vendor was contracted to do. When your examiner asks what your team investigated last quarter, which alerts were reviewed, what evidence was retained, and whether your board approved the detection program, an MDR escalation report is not the answer they are looking for. Examiners are specifically looking for granular detail, something most MDRs do not provide you with.

And that accountability gap is only part of the problem. The data your SOC generates is becoming as strategically important as the detections themselves. As organizations build their AI approach to security and IT, AI agents need a foundation to work from, including accumulated verdicts, tuned detection rules, case history, and institutional memory.

Outsourcing investigation to an MDR means that foundation belongs to your vendor, not you. It lives inside their platform and disappears when the contract ends. Before your next renewal, it is worth asking whether that is a constraint you want to carry into the next phase of how your organization runs security.

**This guide is a practical tool for exploring what your MDR is and isn't doing, and what that means for your risk posture today.**

# 01 The MDR coverage audit

Ask your MDR provider for this data before your next QBR. If they cannot provide it, that is itself worth noting.

### Coverage

- ☐ For each severity tier (critical, high, medium, low, informational), what percentage of alerts were fully investigated in the past 90 days?
- ☐ For alerts that were not investigated, what happened? Were they deprioritized, skipped, or never reached the queue?
- ☐ What data sources are being monitored? Endpoint only, or does coverage include cloud, identity, email, network, and SaaS?
- ☐ What percentage of alerts from identity sources (SSO, MFA, Entra ID, Okta) received individual investigation versus batch processing?

### Investigation quality

- ☐ When an investigation surfaces a finding, does it feed back into detection rule improvement? Or does the same alert pattern keep firing because nothing was tuned?
- ☐ Is there a defined investigation methodology every analyst follows, or does depth depend on who is on shift?
- ☐ When a verdict is reached, can your team see the actual evidence chain, the specific telemetry that led to the conclusion, or do you only receive a high-level summary?

### Detection ownership

- ☐ Where do detection rules live? In your SIEM or in the MDR's proprietary
- ☐ If you cancel the contract today, what detection content comes with
- ☐ How many new detection rules were deployed into your environment in the past 12 months? What triggered each of them?

### Regulatory documentation

- [ ] Which alerts were reviewed last quarter and which were not? Can you produce that breakdown by severity tier?

  *Note: NCUA and NAIC examiners will ask this directly. If your MDR has no record of what was skipped, you cannot answer it.*

- [ ] What evidence was retained from each investigation, and does that evidence belong to your institution or to your MDR provider?

  *Note: NAIC Section 5 requires the licensee to investigate and not use a vendor on their behalf. If the forensic evidence lives inside the MDR's platform, you cannot produce it for a regulator.*

- [ ] Has your board reviewed and approved the detection program itself, or only the MDR

  *Note: FFIEC requires annual board approval of the information security program. NCUA Letter 24-CU-02 requires board-level cybersecurity oversight. Approving a vendor contract is not the same thing. Examiners are starting to ask about the difference.*

- [ ] What would your team find if it looked at the alerts your MDR skipped?

  *Note: Every FSI framework holds the institution accountable for the outcome, not the effort. If a breach originated in an alert the MDR never opened, the regulatory clock started without you knowing.*

## 02 Is it the right time to evaluate MDR alternatives??

The time to evaluate alternatives is 12 months before your MDR contract renews. That gives you enough runway to run a real parallel evaluation and make a decision on data, not urgency. If one or more of the below apply, it's the right time to begin evaluation of MDR alternatives.

- [ ] Your MDR contract renews within the next 12 months
- [ ] In the past 12 months you had a missed threat or near-miss, a significant price increase at renewal, a new CISO or VP Security with a mandate to re-evaluate, or board pressure to justify MDR ROI
- [ ] You have started asking internally whether what you are paying for is reducing risk or managing analyst workload
- [ ] A regulatory examination raised questions about your investigation documentation or detection program visibility

# 03 Seven signs your MDR has hit its ceiling

The following are structural constraints of the human-scaled managed service model.

**If 3 or more are relevant,**
you have likely outgrown what MDR can deliver.

**If 5 or more are relevant,**
you have almost certainly hit the ceiling of the human-scaled MDR model.

**1. You cannot measure whether you are actually safer.** Instead **you have SLA data on response time.** You do not have data on threats missed in uninvestigated alerts, detection posture over time, or actual breach risk reduction. You are measuring the service, not the security outcome.

**2. Alert fatigue has not improved.** Your analysts still feel overwhelmed. The MDR is handling escalations, but the total noise reaching your team has not come down. The volume moved but it did not go away.

**3. Escalations feel like noise.** You get tickets for things that turn out to be nothing. You miss things that turn out to be real. Quality varies depending on analyst experience and who is on shift that day.

**4. You cannot audit the investigation.** When you ask why a verdict was reached, you get a summary. Not a chain of forensic evidence, not the telemetry behind the conclusion. For a regulated financial institution, this creates a specific exposure when an examiner asks you to demonstrate your investigation process.

**5. You do not know what your MDR is not looking at.** There is no reporting on coverage by severity tier. You have no visibility into what percentage of your alert stream goes uninvestigated. You could not answer that question if asked.

**6. Detection engineering is reactive.** Rules get tuned when you complain about volume. New detections take weeks to build and deploy. Your MITRE ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) coverage has not materially improved in the past year.

**7. AI appears in the marketing but not in the outcomes.** Your MDR has started talking about AI-powered detection. Your alert coverage, investigation speed, and costs are unchanged.

# 04 The constraint that made MDR necessary is gone with AI SOC

MDR made sense because no in-house security team could investigate every alert. The volume was too high, the hours too long, the talent too scarce. Outsourcing triage was the only practical answer, even though MDRs face the same issue with humans unable to handle high volumes of alerts.

But those constraints no longer matter.

Intezer AI SOC autonomously triages and investigates all alerts from security detection tools, including EDR, NDR (Network Detection and Response), cloud, phishing, identity, and SIEM, including the low-severity and informational alerts MDR providers skip. This is not enrichment or summarization, which is what most agentic AI SOC tools do. It is a forensic-level investigation with evidence collection, behavioral correlation, memory forensics, kill-chain mapping, and code analysis. Every alert. Every source. Around the clock.

Less than 2% of alerts reach a human analyst. When they do, the evidence is already assembled and the decision trail is fully visible to your team. Your analysts supervise real threats. They do not triage noise.

Every investigation feeds back into detection engineering. When Intezer surfaces a real threat, that finding triggers rule improvement directly in your SIEM. Detection coverage grows with every case. Your MITRE ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) coverage expands continuously rather than waiting for a quarterly review or a complaint about volume. The rules are yours permanently, whether you stay on the platform or not.

Every investigation also produces a documented evidence trail that belongs to your institution, not a vendor. Your board can review and approve an actual detection program, not a service contract.

Ownership matters for another reason that is becoming harder to ignore. As security teams begin deploying frontier AI agents like Anthropic Claude or OpenAI Codex, those agents need a foundation to work from: accumulated verdicts, tuned detection rules, case history, and institutional memory. If your investigation layer is outsourced to an MDR, none of that foundation belongs to you. It lives inside the vendor's platform. Bringing investigation in-house with an AI SOC is not only a security operations decision, but the prerequisite for any effective AI agent strategy in your SOC.

While the staffing gap that necessitated MDR remains, Intezer bridges it without the need for outsourcing. Your team retains full ownership of detections, evidence, and critical decision-making.

# 05 Case study: The account takeover the MDR never reported

A mid-market life and health insurance firm ran Intezer AI SOC alongside their existing MDR in a 30-day evaluation on live production traffic across 11 connected data sources.

**The situation before the evaluation:** Hundreds of escalations per month, fewer than five actionable. Detection rules had never been tuned for their environment. Coverage was endpoint only; identity, cloud, and network telemetry were unmonitored. When asked to produce investigation documentation for a regulatory review, the security team could not.

Results after 30 days:

| Metric | Incumbent MDR | Intezer AI SOC |
|---|---|---|
| Alert coverage | Endpoint only | 100% across all 11 sources |
| Mean time to | ~18 hours | 1 minute 30 seconds |
| Alerts requiring analyst review | Hundreds per month | 1.9% of total |
| Detection rules owned by institution | 0 | 200+ deployed in institution's SIEM |
| MITRE ATT&CK coverage | 34% | 68% at 30 days |

**The incident that defined the evaluation:**
During the evaluation, an employee received a phishing email, surrendered their credentials and MFA token, and an attacker authenticated and registered trusted devices in the company's identity platform. Intezer correlated the phishing event with the authentication anomaly and device registration across three sources in approximately four minutes. The response was immediate.The account was disabled, MFA wiped, attacker-registered devices removed, endpoint quarantined and reimaged.

The incumbent MDR did not escalate this incident. The attack entered through email and identity, two sources it was not monitoring. Under NAIC, the 72-hour notification clock would have started without the security team knowing.

# 06 Before you sign anything

**The five things that separate real AI SOC from rebranded MDR**

Every MDR is adding AI to its pitch. Dozens of startups are claiming to deploy AI SOC agents. These five tests separate real autonomous investigation from marketing claims. Test each in a POC (proof of concept) before signing.

**1. Full coverage, including low and informational severity.** The platform must investigate, not just enrich or summarize, 100% of alerts including low-severity signals. Ask specifically what happens to them. Ask for the data.

**2. Evidence-based verdicts, not LLM summaries.** Real investigation means automated evidence collection, memory forensics, binary analysis, and behavioral correlation. Not a natural language restatement of what the SIEM already said. You should be able to see the evidence trail behind every verdict.

**3. Predictable pricing that does not penalize full coverage.** Ask exactly how the platform is priced. Per alert, per data volume, per agent, or per endpoint. Pricing models that charge by alert volume or data ingestion create a negative incentive to limit what you send to the platform, which defeats the purpose of full coverage. When cost scales with alert volume, low-severity alerts become expensive to investigate, and you are back to the same coverage gap you had before. Intezer is priced per endpoint, making cost predictable, scalable, and decoupled from alert volume. Full coverage stays the default regardless of how noisy your environment gets.

**4. Human experts at escalation, not a queue or chatbot**. When a confirmed high-confidence incident is escalated, you should reach a security expert with full case context already in hand. Test this in the POC.

**5. Closed-loop detection engineering**. Investigation outcomes must feed back into detection rule improvement. Ask how many detection updates were made in the past 30 days, what triggered them, and where those rules now live.

**What you need to own if you decide to end any MDR relationship**
Whether you stay with MDR or move to an AI SOC, these are terms worth insisting on. Any vendor that cannot agree to them is building dependency, not capability.

- [ ] Detection rules are deployed to your SIEM, not stored in a vendor platform
- [ ] Investigation data is accessible to your team, not just escalations but all
- [ ] Tuning history and feedback loops are preserved in your environment
- [ ] You can measure risk reduction over time, not just response time metrics

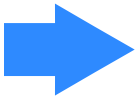

## Two paths forward

If this checklist has surfaced real gaps, you have two practical paths depending on where you are with your MDR.

**Option 1**

**Replace at renewal**
Use the 12-month window before your MDR contract renews to run a POC with an AI SOC platform on live alert traffic. A 30-day pilot on your own data will tell you more than any vendor briefing.

**Option 2**

**Augment first, replace later.**
Add an AI SOC as the investigation layer while your MDR continues. Your team will see in real time what was being missed. By renewal, the comparison will be based on your data, not a vendor's case study.

**Either path starts with Part 1.** If you do not know what percentage of your alerts are actually being investigated right now, that is the first thing to find out.