INTEZER

The 2026 MDR renewal checklist for CISOs and security leaders

60% of alerts go uninvestigated by MDR providers.
Use this guide before your next renewal to assist you in evaluating your MDR as well as how an AI SOC closes MDR gaps.

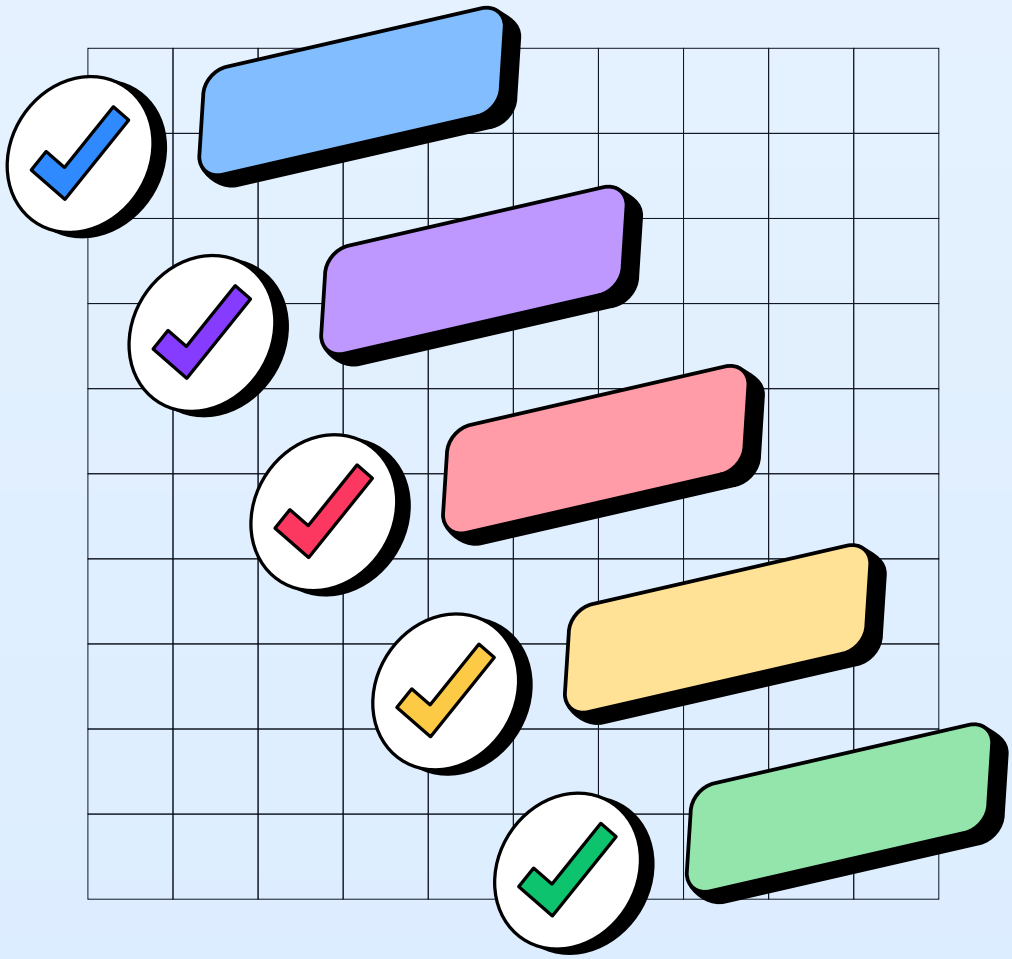

01 Are you in the window to switch?

The best time to evaluate alternatives is 12 months before your MDR contract renews. If you answer yes to any of the questions below, you are in the window.

- [ ] Your MDR contract renews within the next 12 months.
- [ ] You have experienced a missed threat or near-miss in the past 12 months.
- [ ] You have seen a significant price increase at renewal.
- [ ] You are new to the role, or have been tasked to re-evaluate your security stack.
- [ ] You are facing board pressure to justify your MDR return on investment (ROI).
- [ ] You have begun asking internally whether what you are paying for is actually reducing risk or just managing workload.
- [ ] You realize that AI may reduce the need to offload alert handling to a third party.
- [ ] You want to bring security capabilities and knowledge in-house instead of relying on an external vendor.

02 The MDR coverage audit

These questions expose the gap between what your MDR promises and what it actually investigates. Request this data from your provider before your next QBR.

- [ ] For each severity tier, what percentage of alerts were fully investigated in the past 90 days?
- [ ] For those that weren't investigated, what was the outcome: deprioritized, skipped, or never reached?
- [ ] When an investigation turns up a finding, does it feed back into detection engineering, or does the same alert pattern keep firing because nothing gets tuned?
- [ ] Is there a defined investigation methodology every analyst follows for consistent coverage, or does the depth of an investigation depend on who happens to be on shift?

Industry benchmark

Over 60% of enterprise alerts go unreviewed by SOC and MDR teams. For an enterprise generating 450,000 alerts per year, approximately 54 of those are real threats sitting in uninvestigated low-severity alerts, roughly one per week.

Based on Intezer's research of over 25M alerts, which shows that nearly 1% of real threats originate in low-severity alerts

03 Seven signs your MDR has hit its ceiling

Check if any apply to your current experience.

3 or more

You have likely outgrown what MDR can deliver. The constraints you are experiencing are structural, not specific to your provider.

5 or more

You have almost certainly hit the ceiling of the human-scaled MDR model. The question is what comes next.

1. You cannot measure whether you are actually safer. You have SLA data on response time. But most critically, you do not have data on threats missed in uninvestigated alerts, detection posture trend, or real breach risk reduction.

2. Alert fatigue has not improved. Your analysts still feel overwhelmed. The MDR is handling escalations, but the noise has not meaningfully decreased.

3. Escalations feel like noise. You are getting tickets for things that turn out to be nothing, and missing things that turn out to be real. Escalation quality is inconsistent.

4. You cannot audit the investigation. When you ask why a verdict was reached, you get a summary, not a chain of forensic evidence. You are trusting the outcome without being able to verify it.

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 is going uninvestigated.

6. Detection engineering is reactive and slow. Rules get tuned when you complain about volume. New detections take weeks. Your MITRE ATT&CK coverage has not materially improved in the past year.

7. AI is mentioned in marketing but not in your outcomes. Your MDR talks about AI-powered detection, but your coverage, speed, and costs have not improved.

04 Five things an AI SOC must prove before you commit

Not all AI SOC platforms are equivalent. These five criteria separate real autonomous investigation from AI-assisted human workflows with new branding.

1. **Full coverage, including low and medium severity.** The platform must investigate, not just enrich or summarize, 100% of alerts including P3 and P4. Ask for the data. Ask specifically what happens to low-severity alerts.

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 on every verdict.

3. **Predictable and scalable pricing.** Per-alert pricing forces you to cherry-pick which alerts to send, leaving low-severity signals unexamined. Per-endpoint pricing makes 100% coverage the default and decouples cost from alert volume spikes. Ask about this in a POC.

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

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

05 What you need to own at the end of any MDR relationship

Whether you stay with MDR or switch to an AI SOC, these are the terms you should insist 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-side platform.
- [ ] Investigation data is accessible to your team. Not just escalations, all verdicts.
- [ ] Tuning history and feedback loops are preserved in your environment.
- [ ] You can measure risk reduction over time, not just response time metrics.

What comes next

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 renewal to run a POC with an AI SOC platform. A rigorous 30 day pilot with your alert feeds will give you the data to make the decision confidently.

Option 2

Augment first, replace later

Add an AI SOC as the investigation layer while your MDR continues. Let your team see in real time what was being missed. By renewal, the comparison will be data-driven rather than theoretical.

Either path starts with the coverage audit in Part 2

If you do not know what percentage of your alerts are actually being investigated today, that is the first thing to find out.