A BUYER'S GUIDE

FOR SECURITY LEADERS EVALUATING MANAGED DETECTION, OR CONSIDERING SOMETHING DIFFERENT

The MDR evaluation guide

20 questions to ask your current provider and every alternative before you renew or replace: what to ask, why each question matters, and what a good answer sounds like.

FORMAT	USE IT ON	PUBLISHED BY
Four categories, five questions each, with a scoring worksheet	Your current provider and every alternative, including 7AI	7AI, the Foundational AI Security Company

BEFORE YOU START

Score the service, not the SLA

Managed detection is bought on trust and renewed on inertia. The renewal is the one moment the leverage is yours: the moment to ask what share of your alerts is actually investigated, who does the work after an escalation, and what you would keep if you left. This guide exists to make that conversation specific.

TRAP 01

The SLA trap

Most SLAs measure time to acknowledge, not time to conclude. A 15-minute acknowledgment followed by hours of workup on your desk is not a 15-minute service. Score the conclusion, not the clock.

TRAP 02

The filter trap

A quiet queue is not the same as a covered environment. If the noise went away because most alerts were suppressed rather than investigated, the risk did not go away; it went dark. Ask where the alerts went.

How to score

Score each of the 20 questions from 0 to 2, for a maximum of 40 points per provider.

0 POINTS	1 POINT	2 POINTS
Cannot demonstrate it today. A roadmap answer, however sincere, scores zero. You are buying what exists.	Partially demonstrated, or documented but not shown working on a real case in front of you.	Demonstrated live, on real cases, in a way you could verify yourself. Reserve 2s for what you saw working.

Three ground rules

01 **Ask your current provider first**

The most useful baseline in your evaluation is the service you already pay for. Put these questions to your incumbent in writing before you look at any alternative.

02 **Ask for evidence from your own environment**

Your provider holds months of your data. There is no reason to accept an anonymized case study when your own cases exist. Alternatives should show live investigations, not slides.

03 **Count the work you still do**

For one week, log the hours your team spends on escalations, tuning requests, and re-investigation. That number belongs in the cost column of every option you compare.

THE COMPARISON

Two operating models

Traditional managed detection is people-led: a shared bench of analysts triages a subset of your alerts against playbooks, and capacity scales by hiring. The agentic model inverts it: AI agents do the investigation on every alert, and expert humans stay on the loop for judgment, response, and oversight. Neither is a feature difference. They are different answers to the question of who does the work.

	THE TRADITIONAL MDR MODEL	THE AGENTIC MODEL
OPERATING MODEL	People-led. Analysts triage a sample of alerts against playbooks; capacity scales by hiring.	**AI agents investigate every alert end to end; humans on the loop for judgment, response, and oversight.**
COVERAGE	A fraction of alerts fully investigated; the rest filtered, enriched, or suppressed.	**Every alert investigated, across every source, at every hour.**
SPEED	SLAs measure time to acknowledge, in hours.	**Median time to a concluded investigation, in minutes.**
ESCALATIONS	A verdict and a severity; your team does the workup.	**A complete investigation: root cause, timeline, evidence, recommended actions.**
CONTEXT	Generic detection rules; the customer adapts to the provider.	**Customized to your environment, exceptions, and policies; your knowledge codified into the service.**
CONSISTENCY	Shared, rotating analysts; institutional knowledge walks out the door.	**Named experts backed by agents that never rotate off your account.**
TIME TO VALUE	Onboarding measured in months.	**Production in days.**

ALERT COVERAGE

44%

of daily security alerts go uninvestigated in the average organization. Cisco security outcomes research.

THE HUMAN COST

71%

of SOC analysts report burnout, and 64% are considering leaving within a year. Tines, Voice of the SOC Analyst.

THE MARKET

600+

providers claim to offer MDR services. Gartner, Market Guide for Managed Detection and Response.

The 20 questions that follow apply to both models. They are the questions the traditional model finds hardest to answer well, and the ones any credible alternative should welcome.

Coverage and what actually gets investigated

10 POINTS

Filtering is not investigating. These questions establish how much of your alert volume receives a real investigation, and what happens to the rest.

01 **What share of our alerts do you fully investigate, versus filter, enrich, or suppress?**

WHY IT MATTERS Most services keep noise away from you by investigating a fraction and suppressing the rest. The suppressed alerts are where the quiet risk lives.

LISTEN FOR A specific number, per alert source, that you can verify yourself in the console.

02 **Do you investigate across identity, cloud, email, and SaaS, or primarily endpoint?**

WHY IT MATTERS Modern attacks cross domains in minutes. A service built for endpoint leaves the rest of the kill chain to your team.

LISTEN FOR Live investigations shown in each domain, not a connector list.

03 **What happens to coverage at 2 a.m. on a Saturday, or when your queue backs up?**

WHY IT MATTERS Shift gaps, holidays, and queue depth are where things get missed. Coverage that depends on staffing is coverage that varies.

LISTEN FOR An honest account of how capacity scales, and coverage that does not depend on who is on shift.

04 **Do you hunt in our telemetry, or only in the data you collect?**

WHY IT MATTERS A threat hunter who hunts only in the provider's own telemetry leaves everything else dark, and the potential for missing something is high.

LISTEN FOR Hunting inside your data, in your SIEM or data lake, with examples.

05 **How do you keep a real threat from being filtered out with the noise?**

WHY IT MATTERS False positive reduction achieved by suppression trades noise for blind spots. The question is not the false positive rate; it is the false negative rate.

LISTEN FOR A defensible account of how benign verdicts are checked, and how you would know if one was wrong.

Escalations and the work you still do

10 POINTS

The most common complaint about managed detection is quiet and structural: the customer still does the work. These questions measure the handoff.

06 **When you escalate, what arrives: a verdict with a severity, or a complete investigation we can act on without redoing it?**

WHY IT MATTERS Forwarding an alert with a label is not an escalation; it is a handoff of work. If your team re-investigates every escalation, you are paying twice.

LISTEN FOR Root cause, timeline, evidence, and recommended actions in every escalation, shown on a real case.

07 **Who closes the case: you, or us? Walk us through the last mile of response.**

WHY IT MATTERS Detection without response leaves the hardest part on your desk at the worst time.

LISTEN FOR Response executed under a policy you approved, with clear boundaries, not ticket tennis.

08 **What do we need to staff on our side to make the service work?**

WHY IT MATTERS Every managed service has a quiet internal cost. The honest ones can tell you what it is.

LISTEN FOR A straight answer in hours per week, backed by references who confirm it.

09 **When we report a recurring false positive, what changes, and how fast?**

WHY IT MATTERS Recurring false positives that never get fixed are the single most cited reason customers leave a provider.

LISTEN FOR A tuning workflow with an owner and a date, and proof the change held over time.

10 **Can we assign, track, and own cases in your system, or is it a one-way feed?**

WHY IT MATTERS A filtered queue you cannot work in adds a system of record without adding a way to operate.

LISTEN FOR Case ownership, assignment, and collaboration your team can use day one.

Customization and organizational context

10 POINTS

Generic detection treats your business as noise. These questions test whether the service learns your organization, or asks your organization to adapt to it.

11 **How does the service learn our environment: our exceptions, our policies, our VIPs, our maintenance windows?**

WHY IT MATTERS Without organizational context, every scheduled job looks like an anomaly and every executive login looks like a threat.

LISTEN FOR Context applied at investigation time, visible in the investigation itself.

12 **Who tunes detections for us, how often, and does a change require a ticket?**

WHY IT MATTERS Tuning cycles measured in weeks mean your environment is always being defended by last month's assumptions.

LISTEN FOR Named people responsible for your tuning, and changes that take effect in days.

13 **Can we codify our own knowledge into the service: playbooks, past incidents, internal documentation?**

WHY IT MATTERS A service that ignores your institutional knowledge starts from zero every morning.

LISTEN FOR Demonstrated ingestion of your knowledge, with an example of it changing an outcome.

14 **Will the same people know our environment, or do we get a rotating shared queue?**

WHY IT MATTERS Analyst turnover is structural in this market. When your named analyst leaves, their knowledge of your environment leaves with them.

LISTEN FOR A named team, plus an account of what persists when people change.

15 **Is our data used to train shared models, and what stays ours?**

WHY IT MATTERS Your incident data is among the most sensitive information your company holds.

LISTEN FOR A written commitment, plus an architectural explanation of how your context is applied without training on it.

CATEGORY D / QUESTIONS 16 TO 20

Transparency, speed, and proof

10 POINTS

Trust in a managed service is not a feeling; it is a property of the system. These questions establish whether the service shows its work, and what you keep if you leave.

16 **Is your SLA time to acknowledge, or time to conclude? What is your median time to a concluded investigation?**

WHY IT MATTERS An SLA clock that stops at acknowledgment measures responsiveness, not resolution. The work you care about happens after the clock stops.

LISTEN FOR A median time to conclusion they already track, in a dashboard, not an anecdote.

17 **Can we see the full reasoning behind every verdict, including the ones closed as benign?**

WHY IT MATTERS A verdict you cannot inspect is a verdict you cannot defend to an auditor, a regulator, or your own leadership.

LISTEN FOR Step-level reasoning on every investigation, including closures, inside the platform.

18 **What do we keep if we leave: data, detections, tuning, institutional knowledge?**

WHY IT MATTERS Vendor lock-in in managed detection is rarely contractual. It is the years of tuning and context that live only in the provider's system.

LISTEN FOR Your data and your customizations, exportable, with the exit path in writing.

19 **How long until we see production value, and how will we measure it together?**

WHY IT MATTERS Onboarding measured in months is a cost that never appears on the invoice.

LISTEN FOR A dated onboarding plan measured in days, with the success metrics agreed up front.

20 **Can we speak with customers who moved to you from another provider, on the record?**

WHY IT MATTERS References are the only claim a vendor cannot manufacture, and switchers can tell you what the transition actually costs.

LISTEN FOR Named customers with named security leaders, arranged in days, not weeks.

THE WORKSHEET

Scorecard

Score 0, 1, or 2 per question. 40 points possible per provider. Score only what you saw demonstrated. Provider A is your incumbent.

#	QUESTION	PROVIDER A	PROVIDER B	PROVIDER C
A. Coverage and what actually gets investigated				
01	What share of our alerts do you fully investigate, versus filter, enrich,...			
02	Do you investigate across identity, cloud, email, and SaaS, or primarily...			
03	What happens to coverage at 2 a.m. on a Saturday, or when your queue backs...			
04	Do you hunt in our telemetry, or only in the data you collect			
05	How do you keep a real threat from being filtered out with the noise			
B. Escalations and the work you still do				
06	When you escalate, what arrives: a verdict with a severity, or a complete...			
07	Who closes the case: you, or us? Walk us through the last mile of response.			
08	What do we need to staff on our side to make the service work			
09	When we report a recurring false positive, what changes, and how fast			
10	Can we assign, track, and own cases in your system, or is it a one-way feed			
C. Customization and organizational context				
11	How does the service learn our environment: our exceptions, our policies,...			
12	Who tunes detections for us, how often, and does a change require a ticket			
13	Can we codify our own knowledge into the service: playbooks, past...			
14	Will the same people know our environment, or do we get a rotating shared...			
15	Is our data used to train shared models, and what stays ours			
D. Transparency, speed, and proof				
16	Is your SLA time to acknowledge, or time to conclude? What is your median...			
17	Can we see the full reasoning behind every verdict, including the ones...			
18	What do we keep if we leave: data, detections, tuning, institutional...			
19	How long until we see production value, and how will we measure it together			
20	Can we speak with customers who moved to you from another provider, on the...			
Total (of 40)				

32 AND ABOVE	24 TO 31	BELOW 24
A service worth keeping, or an alternative worth a proof of value on your live alert volume.	Promising with gaps. Get every gap answered in writing, with dates, before you renew or commit.	You are staffing the gaps yourself. Price that in, and widen the evaluation.

AN INVITATION

Ask us all twenty.

We published this guide because we believe these are the right questions, and because we are prepared to answer all 20 of them: live, on real cases, on the record. 7AI agents investigate every alert end to end, across endpoint, identity, cloud, email, and SaaS, with full reasoning one click from every conclusion and humans on the loop.

PLAID ELITE is our fully managed service: dedicated AI Security Engineers who customize the platform to your environment, and a 24x7 elite team that handles escalations on your behalf. Your team is the hero of that story, freed from false positives and cleanup work, not handed another queue.

In the words of one customer CISO: "7AI is faster than our MDR, more accurate, and has found what they missed."

Bring this scorecard, your hardest alerts, and your most skeptical analyst. That is who we built it for.