7AI

A BUYER'S GUIDE FOR SECURITY LEADERS

THE DFIR EVALUATION GUIDE

20 questions to ask any digital forensics and incident response provider before you sign or renew a retainer. Written for the AI SOC world.

Four categories. A listening guide for every question. A scoring worksheet to compare providers side by side.

01 **Speed and activation**	02 **Context and environment knowledge**
03 **Forensic rigor and investigation depth**	04 **Commercial terms and ongoing value**

THE RETAINER YOU SIGN IN PEACETIME DECIDES HOW THE WORST WEEK GOES.

Most DFIR retainers are evaluated on brand, hourly rate, and a reassuring conversation with a sales engineer. Then the incident happens, and the questions that actually mattered surface for the first time: how long until real traction, what does the response team know about your environment, and what happens to the hours you paid for and never used.

The AI SOC changes what you should expect from all of them. Attacks now move at machine speed, and agentic security platforms mean a responder can inherit context on day zero instead of building it mid-incident. A retainer conversation that ignores that is a 2020 conversation.

These 20 questions separate a provider built to respond from one built to orient. Ask all of them, of every provider, including us.

01

SPEED AND ACTIVATION

Time to a phone call is marketing. Time to meaningful investigative traction in your environment is the number that decides how the incident goes.

02

CONTEXT AND ENVIRONMENT KNOWLEDGE

A responder who cannot tell your normal from your abnormal spends the most important hours of the incident learning what your team already knows.

03

FORENSIC RIGOR AND INVESTIGATION DEPTH

Collection scope, malware analysis, root cause discipline, and evidence handling that stands up when legal proceedings follow the incident.

04

COMMERCIAL TERMS AND ONGOING VALUE

How scope is set, how hours are consumed and reported, what the engagement report contains, and whether the retainer produces value in the quarters when nothing goes wrong.

TWO OPERATING MODELS. ONE QUESTION: WHO SHOWS UP ALREADY KNOWING YOUR ENVIRONMENT?

Every DFIR provider promises senior people and fast response. The structural difference is what the response team has before the incident starts. A traditional retainer buys you access to experts who begin learning your environment after activation. A context-native retainer, built on an AI SOC platform, buys you experts whose tooling connects when the retainer is signed, and who start every engagement with your telemetry, baselines, and investigation history already in hand. The strongest version of this model is a retainer from the provider whose platform already runs in your environment: the tooling, the baselines, and the investigation history are in place before anything goes wrong.

DIMENSION	TRADITIONAL DFIR RETAINER	CONTEXT-NATIVE DFIR (AI SOC BASED)
TIME TO TRACTION	24 to 48 hours. Initiation, scoping, orientation, and tool deployment all happen after activation.	The first hour. Collection connects at retainer signing, so tooling and context are in place before any incident.
ENVIRONMENT KNOWLEDGE	Learned mid-incident, from your team, while the clock runs.	Learned from the day the retainer starts: baselines, integrations, and a growing investigation history.
TOOLING	Forensic collectors deployed with your IT team during the response.	Connected through agentless, read-only APIs at signing, extended only where the incident requires.
SCOPING	Fixed engagement blocks sized for the provider's staffing model.	Advisory, focused, or full-scope, matched to severity and re-scoped as the incident unfolds.
BETWEEN INCIDENTS	The retainer sits idle. Unused hours expire.	Hours fund readiness work, and environment context keeps accruing.
THE RECORD YOU KEEP	A PDF report and a memory of some phone calls.	Evidence-linked investigations preserved in the platform, plus the written engagement report.

Neither column is automatically right for every organization. The next four sections establish which one you are actually being sold.

SPEED AND ACTIVATION 01

Containment value decays by the hour. These five questions establish the real timeline between "we have an incident" and "someone is doing useful forensic work in our environment," which is rarely the timeline on the datasheet.

QUESTION 01

"How long from activation until your team has meaningful investigative traction in our environment?"

Listen for: an honest number that includes orientation and tool deployment, not just the time to a phone call. If the answer is a response-time SLA, ask what happens in the hours after the phone is answered.

QUESTION 02

"What are the activation channels, and who answers at 2 AM on a holiday?"

Listen for: a direct 24/7 incident response line staffed by responders, not a ticket queue or an account manager who pages someone else.

QUESTION 03

"Walk me through the sequence between the activation call and forensic work starting."

Listen for: a defined sequence measured in minutes and hours: scoping call, priorities agreed, a named lead assigned. "We assemble the right team" usually means the team does not exist yet.

QUESTION 04

"When do your collection tools reach our systems: before the incident or during it?"

Listen for: tooling that is pre-positioned or already integrated. Mid-incident deployment burns the exact hours when containment matters most, and it burns them from your retainer.

QUESTION 05

"Who exactly is assigned to us, and can that person make judgment calls at 3 AM?"

Listen for: a senior Investigative Lead who owns the engagement end to end as your single point of contact, empowered to scope, re-scope, and act without waiting for a committee.

CONTEXT AND ENVIRONMENT KNOWLEDGE

02

The most expensive thing a response team can do is learn your environment on your retainer hours. This category is where the AI SOC world has most changed what a buyer should expect, because context can now exist before the incident does.

QUESTION 06

"What will your responders know about our environment before an incident ever happens?"

Listen for: specifics. Telemetry, integrations, baselines, and investigation history, not a promise to "get up to speed quickly." Quickness is the problem, not the answer.

QUESTION 07

"How do your responders tell our normal from our abnormal?"

Listen for: baselines learned from your environment over time. Generic threat heuristics find generic threats; incidents live in the gap between your normal and everyone else's.

QUESTION 08

"Can you use our alert and investigation history during a response?"

Listen for: an investigation record the response team can query on day zero. The alert that preceded the incident by three weeks is often the root cause in disguise.

QUESTION 09

"What happens when the evidence lives in systems you are not connected to?"

Listen for: a concrete method for cloud audit logs, SaaS platforms, and unmanaged infrastructure, with specialized tooling deployed at your discretion, not a shrug and a scope exclusion.

QUESTION 10

"If we work with you across multiple engagements, does the context carry forward?"

Listen for: persistence. Each engagement should start smarter than the last. If every incident starts from zero, you are paying the orientation tax repeatedly.

FORENSIC RIGOR AND INVESTIGATION DEPTH

03

Speed without rigor produces a fast, wrong answer. These questions test whether the provider can establish what actually happened, prove it, and handle the evidence in a way that survives the lawyers.

QUESTION 11

"What do you collect and analyze: endpoints, logs, network, memory?"

Listen for: all four, scoped to the incident, with a clear method for each. Providers that only read logs are doing incident interpretation, not incident response.

QUESTION 12

"Do you perform malware analysis in-house, or does it go to a third party?"

Listen for: in-house analysis of binaries and suspected malicious code, tied directly to the investigation's conclusions, not a two-week round trip to an external lab.

QUESTION 13

"How do you establish root cause, and how is the attack chain documented?"

Listen for: adversary TTPs mapped to MITRE ATT&CK and a findings narrative where every claim traces to evidence. "The attacker was sophisticated" is not a root cause.

QUESTION 14

"Will your evidence handling stand up in a legal proceeding?"

Listen for: chain of custody discipline, practiced coordination with counsel, and practitioners who have actually supported litigation, not just a paragraph about it in the contract.

QUESTION 15

"Can you contain, or only advise?"

Listen for: the ability to execute pre-authorized containment actions directly, with humans on the loop and a full audit trail. Advice at machine-attack speed is a press release waiting to happen.

COMMERCIAL TERMS AND ONGOING VALUE 04

Most retainer economics are designed around one assumption: you will probably never call. These questions establish what you get in the years you do not, and what you actually receive when you do.

QUESTION 16

"What happens to the hours we do not use?"

Listen for: whether unused hours can fund proactive readiness, such as IR plan reviews, tabletop exercises, and playbook development, or whether they simply vanish at renewal.

QUESTION 17

"How is engagement scope set, and can it change mid-incident?"

Listen for: severity-matched engagement types, from a rapid consultation to full-scope response, with scope treated as a living document. One fixed engagement size means you overpay on small incidents and under-respond on big ones.

QUESTION 18

"How do we see what our hours were actually spent on?"

Listen for: itemized consumption and a regular balance summary, quarterly at a minimum. If hour accounting is opaque before you sign, it will not improve during an incident.

QUESTION 19

"What do we get in writing when the engagement closes, and how fast?"

Listen for: a committed delivery window and a report that serves all three audiences: an executive summary, evidence-backed technical findings with affected inventory, and a prioritized remediation plan.

QUESTION 20

"How does what you learn in our incident make us harder to hit next time?"

Listen for: detection improvement recommendations fed back into your program, and ideally into the platform watching your environment, not a report that goes into a drawer.

SCORING WORKSHEET

SCORE EVERY PROVIDER. INCLUDING US.

Score each answer from 0 to 2. **0** = no real answer, or the answer contradicts the contract. **1** = a credible answer with gaps or caveats. **2** = a specific, verifiable answer the provider will put in writing. Maximum score: 40.

#	QUESTION	PROVIDER A	PROVIDER B	PROVIDER C
01 · SPEED AND ACTIVATION				
1	Time to meaningful investigative traction			
2	24/7 activation channels, staffed by responders			
3	Defined activation-to-forensics sequence			
4	Collection tooling in place before the incident			
5	Senior Investigative Lead as single point of contact			
02 · CONTEXT AND ENVIRONMENT KNOWLEDGE				
6	Environment knowledge before an incident			
7	Your normal vs. abnormal, from learned baselines			
8	Alert and investigation history usable on day zero			
9	Method for evidence in unconnected systems			
10	Context persists across engagements			
03 · FORENSIC RIGOR AND INVESTIGATION DEPTH				
11	Endpoint, log, network, and memory collection			
12	In-house malware analysis			
13	Root cause and ATT&CK-mapped attack chain			
14	Evidence handling that supports legal proceedings			
15	Pre-authorized containment, humans on the loop			
04 · COMMERCIAL TERMS AND ONGOING VALUE				
16	Unused hours fund readiness work			
17	Severity-matched, re-scopable engagements			
18	Transparent hour consumption reporting			
19	Committed, complete engagement report			
20	Findings improve your detection program			
Total (out of 40)				

32 to 40

A provider built for the AI SOC world. Verify the answers land in the contract, then negotiate from strength.

20 to 31

Credible, with structural gaps. Probe the zeros: they are usually where the first 48 hours go.

Below 20

You are buying a phone number, not a response capability. Keep looking.

7AI

ABOUT 7AI

WE WROTE THE QUESTIONS BECAUSE WE ARE READY TO ANSWER THEM.

7AI is the foundational AI security company. The 7AI platform gives enterprises the data, context, intelligence, and action layer for AI-native security operations: AI agents investigate every alert to an evidence-backed conclusion, with humans on the loop.

7AI DFIR is built on that foundation, offered as an add-on for 7AI platform and PLAID ELITE customers. The service is delivered through the platform already running in your environment, so responders start with your tooling connected, your baselines learned, and your investigation history on file. Senior incident response practitioners activate 24/7 through a direct line, run forensic work on tooling already in place, and close every engagement with a board-ready report. Up to 20% of unused retainer hours fund readiness work, so the retainer earns its keep even in a quiet year.

Founded by Cybereason co-founders Lior Div and Yonatan Striem-Amit. $166 Million Total Funding. A force multiplier, not a replacement.

PUT A NUMBER ON YOUR TIME TO TRACTION

Talk to the 7AI DFIR team, walk through activation step by step, and score us against this worksheet.

7ai.com/contact