MANAGED DETECTION & RESPONSE

AI-Accelerated. Human-Validated. MDR you can trust.

CRITICALSTART® is the only Managed Detection and Response provider with contractual per-alert commitments, a deterministic-first architecture, and human validation behind every security decision.

See how Critical Start is raising the bar for MDR.

Every MDR vendor will tell you how fast their AI responds. **Ask them what happens when it's wrong.**

4
THINGS TO ASK EVERY VENDOR

01 ACCOUNTABILITY · NOT JUST SPEED

Speed is the table stake. Accountability is the difference.

The market races to show how fast AI can respond. Your real question is: what happens when it's wrong? At Critical Start, every threat alert, low, medium, or high severity, is validated by a qualified SOC analyst.

96% of CISOs now own **AI governance & risk management.** The board question is already in the room.

SOURCE: GARTNER CISO SURVEY, 2025

02 HUMAN GOVERNANCE · AUDIT-READY BY DESIGN

"Fully-autonomous" is a risk word. Not an aspiration word.

Every one of our AI recommendations carries a full audit trail: what the agent proposed, what the analyst decided, and why.

Gartner estimates **60% of AI incidents by 2026 will stem from governance failure.** Our answer isn't a policy document, it's the architecture itself.

03 DATA ISOLATION · BOUNDED BY ARCHITECTURE

Your data stays in your environment. Full stop.

Most AI security platforms treat cross-customer data sharing as a network effect. For regulated industries, that's a data governance risk. Our architecture keeps telemetry inside your environment. Agents operate on your signal, within your Rules of Engagement and never share investigation outputs across customers.

78% of CISOs name **data leakage** as their #1 concern with AI in security operations.

SOURCE: GARTNER PEER INSIGHTS, 2025

04 CONTRACTUAL OUTCOMES · NOT BENCHMARKS

We put our SLA in writing. Ask if your other vendors will.

Every AI vendor has a benchmark: response in five minutes, triage in seconds, investigation cut by 68%, measured in vendor-controlled environments, against vendor-selected workloads, with no contractual commitment behind them.

Critical Start offers **per-alert SLAs**, audit-ready obligations on how every signal is handled. The question isn't whose AI is fastest in a demo. **It's who will put it in writing.**

Ready to see SOC AI in action?

REQUEST A DEMO →

HOW IT WORKS

From your first signal to a resolved verdict in one flow.

AI accelerates every step. Your analysts and ours validate the moments that matter. The result is a contractual, audit-ready outcome.

Five agents : one governance model

Every agent **proposes**. A qualified human **approves**. Each one has a tightly scoped charter: what it can do, what it cannot, and what it must hand off.

01 TBR Agent

Examines false positives in production and proposes new entries to the Trusted Behavior Registry — the deterministic rules engine that has filtered ~99% of events for more than a decade.

02 Investigation Agent

Accelerates triage by pre-populating investigative workflows, running OSINT enrichment, and suggesting verdicts on alerts for analyst review.

03 Case Agent

Detects, links, and proposes closure of related alerts during triage so analysts work a single coherent case instead of duplicates.

04 Response Agent

Suggests existing automations from the catalog to run on true-positive verdicts and proposes new workflows for human review before publication.

05 Threat Hunting Agent

Runs hypothesis-first proactive hunts against all ingested events to surface threats before they become alerts and suggests new detections for the catalog.

02 Detect · Investigate · Respond · Improve