

Five questions any AI SOC should be able to answer.

A public, fair way for MSPs and MSSPs to check any AI SOC claim, CYREBRO AI's included.

OPEN DRAFT · VERSION 0.1

PUBLISHED OCTOBER 12, 2026

COMMENT UNTIL NOVEMBER 11, 2026

- | | | |
|----|-------------------------------|---|
| 01 | Zero-touch rate | What share of alerts close with no human touching them? |
| 02 | Alert time | How long from log ingestion to verdict? |
| 03 | Human role | Who steps in, when, and to do what? |
| 04 | Build burden | Who builds and maintains the detection logic? |
| 05 | Multi-tenant readiness | Was it built to serve many clients from one place? |

Disclosure: CYREBRO AI created the AI SOC Truth Test and is one of the platforms it measures. This is an open draft, not independent or neutral third-party research. It publishes measurements and gives no grades, rankings, badges or endorsements. soctruthtest.com

WHY THIS EXISTS

“AI-powered” is on every SOC. What it means isn’t.

Every SOC vendor now says “AI-powered,” and many say “agentic.” A service provider building its business on someone else’s SOC has no simple way to tell real automation from a new label on the same work.

GARTNER CALLS THE GAP “AGENT WASHING”

“the rebranding of existing products, such as AI assistants, robotic process automation (RPA) and chatbots, without substantial agentic capabilities.”

Gartner, Inc., “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027,” press release, June 25, 2025.

The AI SOC Truth Test proposes five questions any platform should be able to answer. CYREBRO AI publishes its own answers first and invites every other vendor to do the same. A company that calls out agent washing has to be the first to show its work.

Who it’s for: MSP and MSSP owners and security leads comparing SOC platforms; security practitioners who want a method worth critiquing; SOC vendors who want a fair, published standard applied the same way to everyone; and reporters and analysts who want a clear, checkable story.

WHAT THE TEST IS, AND ISN’T

IT IS	IT IS NOT
An open draft (version 0.1), clearly labeled as a draft	A finished standard, or something a platform can pass
Measurements, sourced from each vendor’s own reporting or public claims	Grades, scores, rankings, badges, or endorsements
Created by CYREBRO AI, and says so plainly	Presented as independent or neutral third-party research
Applied the same way to every vendor, including CYREBRO AI	A comparison chart naming competitors
Open to public comment and to an independent review panel	A gated lead form. The whole method is readable without signing up.

How draft 0.1 defines each measure

Every answer comes from the vendor's own reporting or public claims, with the source stated. Each measure ends with an open question for public comment.

01 Zero-touch rate

Of all the alerts the platform handles, what share does it take from first signal to closed verdict with no human touching them?

WHY IT MATTERS

It's the plainest test of "agentic." If people still review most alerts, the AI is assisting, not operating.

WHAT A VENDOR PUBLISHES

The rate, the period it covers, the number of alerts in that period, and what the vendor counts as a touch.

HOW DRAFT 0.1 MEASURES IT

The share of alerts, over a stated period, that reach a final verdict and are closed without any human action. Any human review, approval, edit or reopening counts as a touch.

SOURCE

The vendor's own reporting or public claims, with the source stated.

Open question: Should alerts that the platform closes on its own, and a person later reopens, count against the rate?

02 Alert time

How long does it take from the moment a log arrives to the moment the platform reaches a verdict?

WHY IT MATTERS

Speed claims are easy to make and hard to compare. Fixing where the clock starts and stops makes them comparable.

WHAT A VENDOR PUBLISHES

The typical figure, the period it covers, and how the two timestamps are recorded.

HOW DRAFT 0.1 MEASURES IT

Elapsed time from log ingestion to verdict, as the platform records it. The clock starts when the log is ingested, not when an alert is raised, and stops at the verdict, not at a later notification.

SOURCE

The vendor's own reporting or public claims, with the source stated.

Open question: Should version 1.0 ask for a median and a slow-case figure (such as the 90th percentile) instead of a typical figure?

03 Human role

When a person is involved, who is it, when do they step in, and what do they do?

WHY IT MATTERS

"AI-powered" can still mean an analyst behind every alert. A partner needs to know whose people do the work, and when.

WHAT A VENDOR PUBLISHES

The description, and whether it changes by service tier.

HOW DRAFT 0.1 MEASURES IT

A plain description, not a number. Who: the vendor's analysts, the partner's analysts, or the client. When: every alert, a sample, or escalations only. What: investigation, approval, or remediation.

SOURCE

The vendor's own reporting or public claims, with the source stated.

Open question: Should the test separate people who approve an action from people who carry it out?

04 Build burden

Who builds and maintains the detection logic, and what else does a partner have to buy, deploy or run?

WHY IT MATTERS

If the partner has to write the rules or bring a separate SIEM, the work and the cost move to the partner.

WHAT A VENDOR PUBLISHES

The description, and a list of anything a partner must supply.

HOW DRAFT 0.1 MEASURES IT

A plain description of who writes and tunes the detections, and which components, such as a SIEM, are included and which a partner must supply.

SOURCE

The vendor's own reporting or public claims, with the source stated.

Open question: Should the test also ask how much work it takes to bring a new client on board?

05 Multi-tenant readiness

Was the platform built to serve many clients from one place, with each client's data kept separate?

WHY IT MATTERS

MSPs and MSSPs run many clients at once. A single-tenant platform stretched across clients adds work and risk.

WHAT A VENDOR PUBLISHES

The description, and whether multi-tenancy was built in or added later.

HOW DRAFT 0.1 MEASURES IT

A plain description of how clients are kept separate, and how a partner manages them together and one at a time.

SOURCE

The vendor's own reporting or public claims, with the source stated.

Open question: What evidence of client separation should version 1.0 ask vendors to point to?

We go first

CYREBRO AI's answers under draft 0.1, from our own reporting, applying the same rules we ask of every vendor. Our complete answers, including the zero-touch rate, come with version 1.0.

CYREBRO AI		Self-reported · draft 0.1 · not independently reviewed	
02 · ALERT TIME		01 · ZERO-TOUCH RATE	
Typically under one second		To be published with version 1.0	
from log ingestion to verdict.		Measured the way the final methodology defines it.	
03	Human role	CYREBRO AI's escalation desk (or the partner's analyst) on escalation only, with remediation by the partner or client, guided by step-by-step recommendations.	
04	Build burden	Detection logic built and maintained by CYREBRO AI, with CYREBRO Hyper-SIEM built in.	
05	Multi-tenant readiness	Multi-tenant by design.	
Source: CYREBRO AI's own reporting, as published with draft 0.1 on October 12, 2026.			

DISCLOSURE AND GOVERNANCE

Who made it CYREBRO AI created the AI SOC Truth Test and is one of the platforms it measures. The test grades, ranks and endorses no one, CYREBRO AI included.	The review panel An independent panel reviews the method for version 1.0 and audits the draft log. Members and their conflict disclosures are published with version 1.0. Not on the panel: • Paul Warnagiris (board member and investor) • CYREBRO AI staff • The partners quoted at launch
--	--

Moderation during the comment period

CYREBRO AI moderates the public draft log at soctruthtest.com until the panel is seated. We publish every on-topic comment whose author agrees to it, usually within two business days, edit only to remove other vendors' names and personal details, and mark every edit. The panel then audits the full log, including what wasn't published and why.

WHAT HAPPENS WHEN

From draft to version 1.0

MON, OCT 12, 2026

Draft 0.1 published

The five measures and CYREBRO AI's initial answers go public. Comments open.

WED, NOV 11, 2026

Public comment closes

Every comment is read and considered for version 1.0.

WEEK OF NOV 30, 2026

Version 1.0

The named review panel, CYREBRO AI's complete answers including the zero-touch rate, and vendor self-reports.

TAKE PART

Shape version 1.0

Comment on the draft

On any measure, the method, or CYREBRO AI's answers. Until November 11, 2026. You choose whether it's published, and under what name.

Join the review panel

Open to practitioners, MSP and MSSP leaders, and researchers. Every applicant discloses conflicts of interest.

Submit a platform

Vendors submit their own answers and sources for version 1.0. They're published as submitted, not graded or ranked.

All three are at soctruthtest.com, or email comments@soctruthtest.com.

USING THE FIVE QUESTIONS

Anyone can use the five questions with any vendor, without signing up or telling us. Ask each vendor to answer from its own reporting or public claims, and to state the source for each answer.

CITING THE TEST

The AI SOC Truth Test, draft 0.1. Created by CYREBRO AI. Published October 12, 2026. soctruthtest.com

Please include: CYREBRO AI created the AI SOC Truth Test and is one of the platforms it measures.