FIELD CISO ADVISORY | FIELD SECURITY INSIGHTS

FIELD SECURITY REPORTS

The Promise of AI SOC: From Detection to Prevention

A practitioner's framework for judging AI SOC platforms, the gaps nobody advertises, and how three vendors hold up once you push past detection and triage.

FEATURING

BLUSAPPHIRE ENTERPRISE SECURITY SIMPLIFIED

SPHARAKA™

Exaforce

AUTHOR
Prateek Bhajanka

PEER REVIEWED BY
Vatsal Sharma, Parth Khullar

PUBLISHED
September 2, 2026

CONTACT
insights@fieldciso.com

The Promise of AI SOC: From Detection to Prevention

Abhijit Chakravarty's DDRRRP framework, the gaps the category still has, and how BluSapphire, Spharaka and Exaforce hold up against it.

The DDRRRP Framework

Abhijit Chakravarty, EVP for Networks and Cyber Security at Kotak Mahindra Bank, treats AI SOC as six distinct stages rather than one capability[1]: detect, diagnose, report, recommend, remediate, prevent. Both the engineering effort and the business value rise as you move right along the curve. Most products on the market today work in the first two stages only. Governance becomes a serious question at remediation, and prevention is the hardest stage to sell because the payback takes longest to demonstrate.

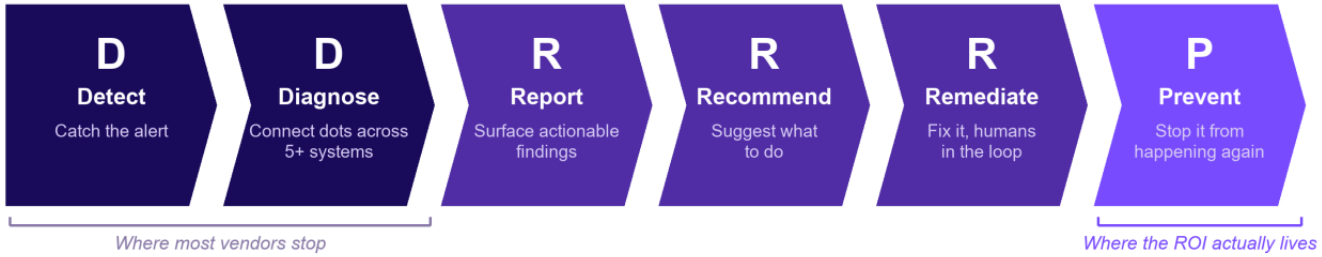

What We Should Expect From an AI SOC

- **Volume without fatigue.** An analyst cannot run ten investigations at once at three in the morning. A machine can, and the tenth should be handled as well as the first.
- **Noise into signal.** Correlate EDR, SIEM and threat intelligence so that genuine attacks stop hiding inside false positives.
- **The engineering calculator model.** Fast, accurate and tireless, but the analyst still owns the logic. The tool should extend judgement rather than substitute for it.
- **Data economics as a feature.** Keep security-critical data hot, archive the rest to object storage and federate when you need it. Not every log deserves online retention.

Where AI SOC Falls Short Today

- **Still firefighting.** A hundred alerts today and the same hundred in six months means nothing was fixed. Root-cause work rarely happens.
- **AI does not do grey.** It handles binary logic and coded rules well. Context, competing priorities and business judgement stay with people.
- **Vendors do not understand SOC workflow.** They ship products rather than platforms, then wonder why onboarding stalls.
- **Cost is ignored.** SIEM bills climb, storage grows and compliance asks for seven years of logs. Very little of that needs to sit online, and almost nobody models the difference.

> *"Knowledge is not about answering more questions. It is about having fewer questions to begin with."*
>
> Abhijit Chakravarty, on why prevention matters more than faster detection

Responsibility and Auditability

- Every AI action should be logged, attributed and explained: why the detection fired, which sources it drew on, which agent acted and who approved it.
- Anything touching a crown-jewel asset should go to a human approver. "The model decided" is not an answer a regulator will accept.
- If a vendor cannot show you the audit trail during the POC, assume it does not exist.

From Detection to Prevention

- Prevention means moving past patching towards architecture and process redesign, so that a whole class of attack disappears rather than one instance of it.
- The scoreboard has to change with it: stop counting new detections and start counting the fall in repeat alerts, and whether the fixes hold.
- Compliance reporting should be automated too, because banking, pharma and finance each carry their own KPIs and nobody should be assembling those by hand at midnight.

Top Questions to Ask Any AI SOC Platform

- **Speed and accuracy.** Alerts per day, false-positive rate, and can you show it improving?
- **Integration depth.** What genuinely connects, and who builds the connector when it does not?
- **Data cost.** Retention, tiering and federated search performance. What does year three cost?
- **Compliance.** Are industry-specific reports generated, or is that still someone's weekend?
- **Custom rules.** Can I write detection logic in plain language, or do I need a SIEM specialist?
- **Prevention.** Does alert volume fall over time, and will it recommend architectural change?
- **Auditability.** Is every decision logged, and can I trace where the data came from?

Vendor Profile: BluSapphire (Garuda / AR²)

BLUSAPPHIRE ENTERPRISE SECURITY SIMPLIFIED **blusapphire.com**

Kiran Vangaveti, founder and CEO, walked through Garuda's agentic product, AR², live rather than with slides.[2,3]

- **Company background & onboarding.** Nine years old heading into its tenth, 500+ customers overall, 300+ on the agentic platform, including the SEBI-regulated Bombay Stock Exchange managed SOC. New connectors can be auto-built from a vendor's API documentation, right on the first run in most cases, and the platform accepts inbound webhooks as well as outbound pulls.
- **Use case.** Point it at an alert and it correlates across SIEM, cloud and identity sources on its own, scores a recommendation, and queues one-click remediation behind human approval. The same engine also runs on demand: ask it to scan S3 buckets for public exposure across Azure, AWS and GCP on a weekly schedule, and it hands back one consolidated report with no separate CSPM tool required.
- **Model approach.** A proprietary gateway rather than commitment to one frontier LLM, a Bedrock-hosted default, and bring-your-own-model support. Kiran tied this directly to the risk of frontier-model access being restricted, citing the recent Fable/Mythos suspension.
- **Data architecture.** Detection@Edge keeps logs local and forwards only security signals, about 0.02% of traffic, solving both egress cost and residency; Kiran expects industry pricing to shift from per-GB toward compute-based licensing.
- **Kiran's advice for evaluators.** Ask how a vendor handles frontier-model gatekeeping, whether you can switch models without lock-in, and whether the platform ever fully cuts off access when usage runs out.

Vendor Profile: Spharaka Networks (Sphere)

spharaka.com

Rajeev Raj, CEO, and Ravi Nalluri, CTO, co-founders with nineteen years each in security operations, presented Sphere together, alternating between vision and live demo.[4,5]

- **Platform.** Sphere unifies SIEM, UEBA, SOAR, EDR, XDR and threat hunting on one data lake, powered by SAGE, Spharaka's own security-trained model, and AuraXP, 40+ collaborative agents coordinated by a master agent.
- **Use case.** Full autonomous investigation with no internet connection at all: point it at a raw alert and it reconstructs the whole attack chain, from reconnaissance to exfiltration, with MITRE mapping and supporting evidence, in seconds rather than the 30 to 45 minutes a skilled analyst would typically need, entirely air-gapped with no pre-built playbook involved.
- **Remediation, positioning & sovereignty.** Remediation is configurable in three tiers by asset risk score: autonomous, approval-gated, or recommend-only. Rajeev was explicit that Spharaka isn't replacing or sitting as a thin layer on a SIEM, since bolting agentic AI onto a SIEM alone only gets you alert enrichment. SAGE itself runs fully on-premises and air-gapped, with no data leaving the network and no per-token billing.
- **Rajeev and Ravi's advice for evaluators.** Confirm whether the model is dedicated to you rather than shared in someone else's cloud, whether the platform is genuinely autonomous rather than a co-pilot, whether data ever leaves the network, and whether the vendor is actually investigating an alert or just explaining it.

Vendor Profile: Exaforce

exaforce.com

Aksa Taylor, chief security evangelist, opened, then Arif Huq, co-founder and head of product, ran the live demo.[6,7]

- **Company background & platform.** Founded in 2023, about $200M raised across two rounds, in production with customers for over a year and a half, including Replit and larger enterprises. Exaforce built its own data platform first, so agents work from ingested identity, permission and configuration data, not alerts alone; 100+ integrations, agentic auto-onboarding, and every alert gets an automatic disposition and priority score, with about a 95% cut in false positives.
- **Use case.** Ask it something like "list every AI application in my environment and its permissions into Google Workspace and Microsoft 365," an identity and permissions question a SIEM alone can't answer, and it writes the queries, cross-references identity data, and returns a plain-language answer with the full reasoning shown, drivable from Slack, MCP clients, or the mobile app.
- **Pricing & advice for evaluators.** Outcome-based (alerts triaged, investigations run) rather than metered by token, except for custom agents you build yourself. Arif's advice: ask how a vendor reaches its conclusions, whether you can inspect the underlying data, and whether you can define custom response actions rather than relying only on prebuilt ones.

Conclusion

- Model quality is not what will decide this market. It will come down to which vendors understand how a SOC actually runs, take a clear position on data cost, and can move a team from firefighting towards prevention.
- All three platforms demonstrated real strength at detect and diagnose live, and all three currently keep a human in the loop before anything executes on the remediation side.
- Prevention is where the category is thinnest today, and audit-trail or escalation commitments belong in the contract, not the demo.
- Whatever platform you choose, review it after six months against one measure: are the same alerts still arriving.

References

1. Kotak Mahindra Bank. Abhijit Chakravarty, EVP, Networks and Cyber Security. in.linkedin.com/in/abhijit-chakravarty-1740105

2. BluSapphire. "Garuda by BluSapphire, AI Native Agentic Security Platform." www.blusapphire.com/garuda

3. FieldCISO. "BluSapphire: Agentic AI SOC, SIEM, SOAR." fieldciso.com/field-security-insights/vendors/blusapphire

4. Spharaka Networks. "Autonomous Cyber Defence Platform." spharaka.com

5. Spharaka Networks. "The End of the Human-Speed SOC." spharaka.com/blog

6. Exaforce, Inc. "Exaforce Raises $125M Series B." Press release, May 12, 2026.

7. Exaforce, Inc. "Exaforce Unveils $75M Series A." Press release, April 17, 2025.

Author: Prateek Bhajanka | Peer Reviewed by: Vatsal Sharma, Parth Khullar | Contact: insights@fieldciso.com