

MIRRORFIELD · POSITION PAPER

Passive Engagement vs. the Deception Arms Race

Why a defense architecture that doesn't depend on an attacker taking the bait is a structurally different bet than decoys, honeytokens, and lures.

BY MERIDIAN INTELLIGENCE, INC. · SEPTEMBER 14, 2026

Most of the cybersecurity industry's attempts to engage an attacker start from the same premise: build something that looks worth stealing, place it where an intruder is likely to find it, and wait. Decoy credentials, honeytokens, planted file shares, deceptive network paths — the mechanism varies, but the logic is the same. Success depends on the adversary taking the bait.

That approach works, and it has produced real, valuable security outcomes for the vendors who build it well. But it also has a structural ceiling: a defense that depends on the attacker's decision to interact with a planted artifact is, definitionally, a defense the attacker can decline to trigger. Published security research on honeypot and honeytokens fingerprinting has documented that sophisticated actors can identify known deception patterns and adapt around them.

A Different Starting Premise

MIRRORFIELD was designed around a different question: what if engagement didn't depend on the adversary choosing to interact with something planted for them at all?

Rather than constructing attractive fake assets and waiting for an attacker to reach for them, MIRRORFIELD's architecture is built on co-present, passive self-selection — the defended environment does not need to lure anyone anywhere. Activity that meets the architecture's structural conditions is engaged on its own terms, without depending on bait, planted credentials, or an attacker's willingness to bite.

This is a meaningful architectural distinction, not just a marketing one. A bait-dependent system's effectiveness is bounded by how convincing the bait is and how well it's hidden from an attacker who is actively looking for it. A passive, structurally-triggered system doesn't carry that dependency.

Provisional Suspicion Is Not Ground Truth

The second distinction matters just as much. Some deception platforms treat interaction with a deceptive asset as high-confidence, often deterministic, evidence of an attack — if you touched it, you're presumed hostile.

MIRRORFIELD treats classification as a genuine state transition, not an assumption. Anomalous or credentialed activity can be provisionally mirrored for observation without that observation itself being treated as proof of adversarial intent. Suspicion and confirmed adversarial state are kept as separate, explicit governance states — because collapsing them prematurely produces exactly the false-positive fatigue that undermines trust in any detection system over time.

Why This Matters Going Forward

Deception-based platforms are not a weak category of competitor — several of them do controlled adversary engagement and intelligence generation well. The honest comparison is not that deception doesn't work. It's that:

- Bait-dependent engagement is bounded by an attacker's willingness to interact with what's been planted for them.
- Deterministic tripwire logic trades false-positive risk for engagement simplicity.
- A passive, state-governed architecture removes the first dependency and manages the second explicitly, rather than assuming it away.

That is the case MIRRORFIELD is built to make — and the evidence from controlled engineering evaluation, not marketing claims, is what will ultimately test it.

Sources & References

- CrowdStrike — "Falcon AIDR: AI Detection & Response," crowdstrike.com. Accessed September 14, 2026. www.crowdstrike.com/en-us/platform/falcon-guardian-aidr/
- Palo Alto Networks — "Prisma AIRS," paloaltonetworks.com. Accessed September 14, 2026. www.paloaltonetworks.com/prisma/prisma-ai-runtime-security
- Microsoft — "ID Protection for Agents," Microsoft Learn. Accessed September 14, 2026. learn.microsoft.com/en-us/entra/id-protection/concept-risky-agents
- Acalvio — "ShadowPlex Preemptive Cybersecurity Platform," acalvio.com. Accessed September 14, 2026. www.acalvio.com/shadowplex-platform/
- XM Cyber — "Attack Graph Analysis," xmcyper.com. Accessed September 14, 2026. xmcyper.com/platform/attack-graph-analysis/
- Proofpoint — "Shadow: Deception Technology Solutions to Stop Attacks," proofpoint.com. Accessed September 14, 2026. www.proofpoint.com/us/products/identity-threat-detection-response/shadow
- "A Review of Honeypots: Fingerprinting Techniques, Detection, and Evasion Mechanisms," Future Internet (MDPI), 2026. www.mdpi.com/1999-5903/18/4/190

Vendor characterizations above reflect each company's own public product claims as of the access date and are provided for competitive-landscape context only; they are not independently verified by Meridian Intelligence.