SACR®

MARCH 2026

HOW AGENTICS REDEFINES PREEMPTIVE DECEPTION PLATFORMS

LAWRENCE PINGREE

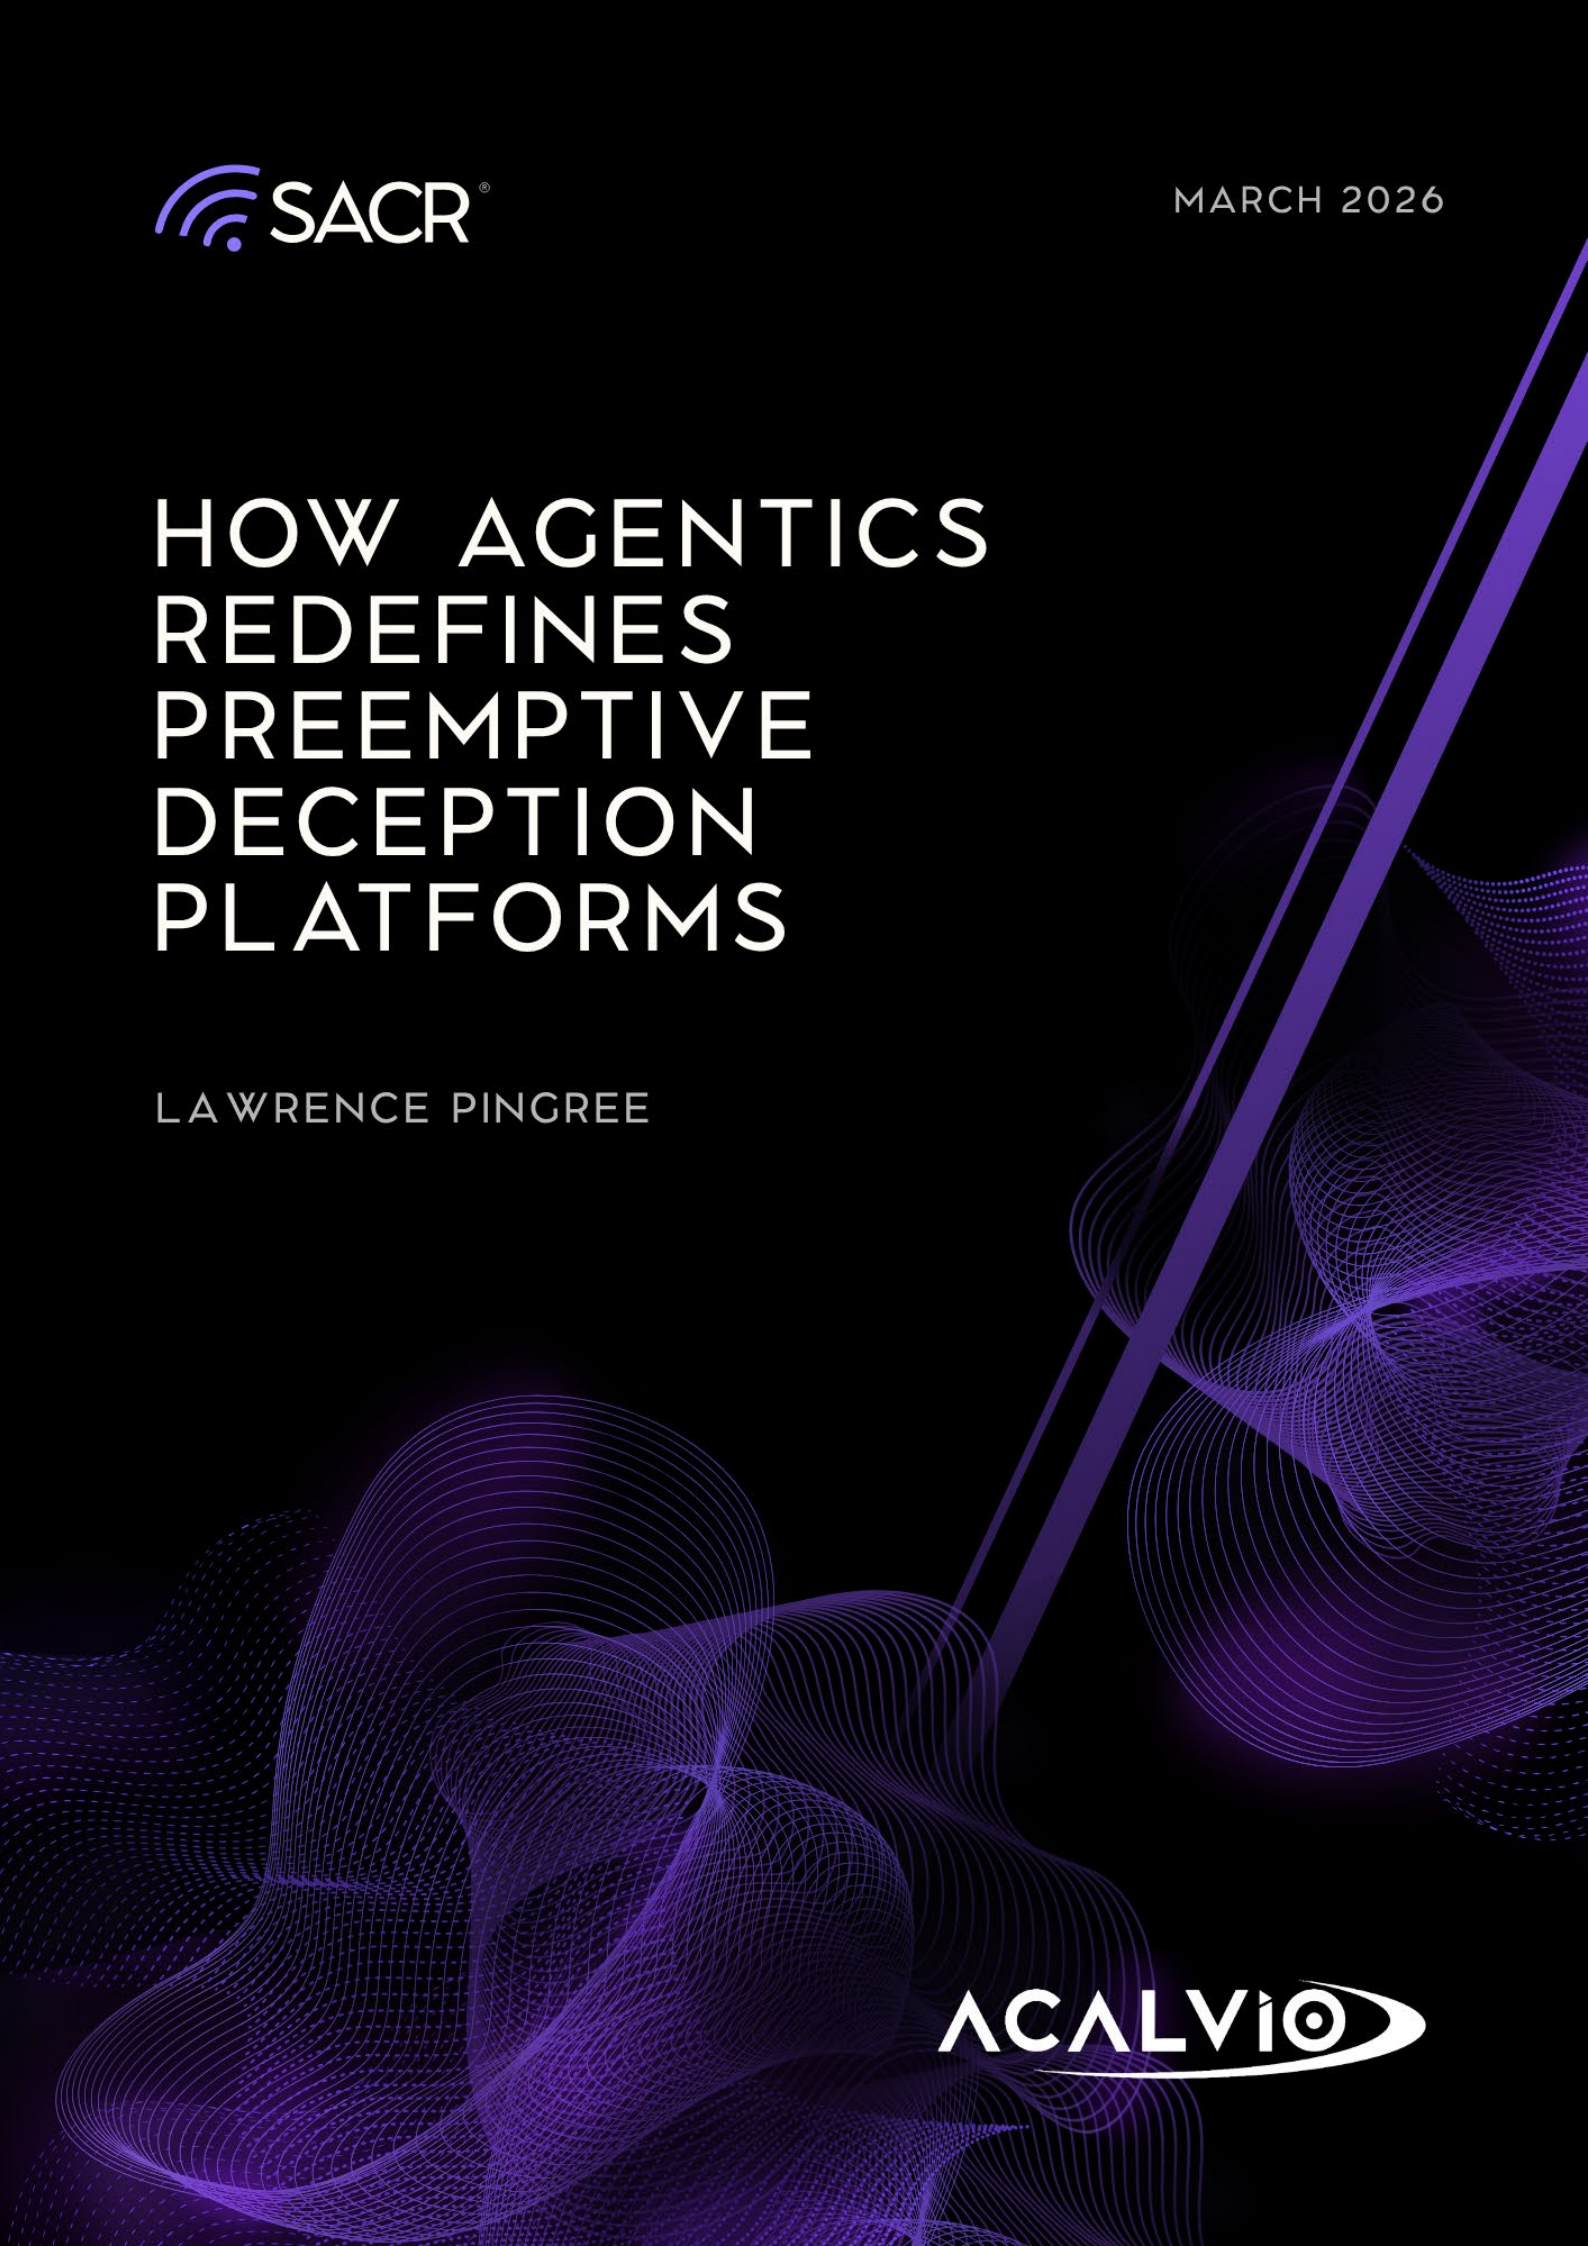

ACALVIO

SACR®

We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, our timely, opinionated insights help modern security leaders make smarter, faster decisions.

About Us

Software Analyst Cyber Research (SACR) is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over **80,000 readers and followers**, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.

Acknowledgements

Practitioners and CISOs,

We're excited to share a new framework and system for securing agents across enterprises. This is one of our biggest reports published in recent months.

The core author:

- Lawrence Pingree is the Head of Data and AI Security at SACR, where he leads research on data protection, AI security, and agentic security models. He brings more than ten years of analyst experience from Gartner and has authored over 300 research notes across cloud security, endpoint defence (EDR), SD-WAN, and AI security.

Table of Contents

Actionable Summary

This report analyzes the emergence of Preemptive Deception Platform and Acalvio's market position, technical differentiation, and value proposition. Acalvio's long-term success in the emerging market for decentralized, AI-driven deception is critically dependent on its capacity to scale its proprietary technology and successfully navigate complex, multi-jurisdictional regulatory environments. The critical takeaway is that Acalvio effectively solves the industry's most persistent problem: dwell time and precise detection. By detecting threats the moment they attempt lateral movement or credential theft, Acalvio transforms a network from a static target into a dynamic minefield. Acalvio offers deception libraries of templates tailored for IT, OT (Operational Technology), and cloud environments.

The thesis behind this report is that Acalvio is positioned as a market disruptor in the deception technology space, especially as they integrate and target introducing more Agentic capabilities into their product. The company is actively reshaping how organizations approach threat detection and response by moving beyond traditional perimeter defenses and signature-based tools. Acalvio's core strength lies in its ability to deploy highly realistic, distributed deception layers that lure and detect adversaries early in the attack lifecycle, providing high-fidelity alerts that drastically reduce false positives and accelerate mean time to detection and response. This report argues that Acalvio's innovative and scalable deception-as-a-service model makes it a critical component of modern security architectures, poised for significant growth as enterprises prioritize active defense strategies.

Key Insights

The Assume Breach Era Demands Preemptive Action: Traditional detection (EDR/XDR) relies on spotting anomalies after an attacker has already landed. In the age of AI-driven attacks, where dwell time is shrinking, defenders must shift left, from reactive detection to preemptive deflection.

Deception as a Counter-AI Strategy: AI agents automate reconnaissance and lateral movement, moving at machine-speed to compress attack timelines. By polluting the environment with deceptive assets, defenders can disrupt the logic loops of these automated attackers, causing them to waste resources or expose themselves early.

Identity is the New Perimeter: With credential theft driving most breaches, Identity Threat Detection and Response (ITDR) has become a critical layer. Deception technologies that plant honeytokens (e.g. fake credentials) offer a high-fidelity signal that traditional identity monitoring often misses.

Market Context & Problem, Evolution

The Core Problem: The Breach Reality

Modern cybersecurity operates under the assumption of a breach. Sophisticated adversaries (APTs), ransomware gangs, and malicious insiders routinely bypass Endpoint Detection and Response (EDR) and SIEM filters by using legitimate credentials (living off the land). Generative AI is accelerating attacker activities, but it has a critical achilles heel – it is reliant on data, just like humans.

The crux of why security programs need to leverage preemptive deception techniques, especially deception, is that there are critical gaps in existing tools, including EDR solutions, which are known to miss many attacks. This is not because the tools are ineffective, but because connecting disparate activities into clearly defined malicious behavior is difficult, and attackers intentionally work around these advanced tools and behavioral defenses.

- The Gap: Traditional tools struggle to distinguish between a legitimate admin using PowerShell and an attacker doing the same.This doesn't really change even with EDR. Attackers have also moved on to unmanaged systems, personal devices and compromised credential use, thus, no prior telemetry about them exists from the point of view of the enterprise. Cloud conscious adversaries are using built-in primitives to compromise workloads, making it challenging to detect given the dynamic and ephemeral nature of the workloads.
- The Consequence: This results in extended dwell times (the time an attacker spends undetected), often measured in weeks or months, allowing for data exfiltration and ransomware deployment. When credentials are used but retrieved from malware or proxies from consumer hardware or unmanaged devices, from the point of view of the enterprise, the use of the credential might be a tiny bit suspicious, but not convictable as malicious on first use.

Traditional cybersecurity tools like firewalls are built like fortresses, designed to keep intruders out of a network. However, despite increasing security investments, sophisticated attackers often find a way to slip past these perimeter defenses, especially in the unique and sensitive environments of Operational Technology (OT) and Industrial Control Systems (ICS). The real danger begins once an attacker is inside an industrial network. The convergence of IT and OT systems creates trusted pathways that adversaries exploit to pivot from corporate networks into the plant floor.

Once inside, they can move silently, exploring systems and escalating privileges in a process known as lateral movement. Their goal is to find and compromise the most valuable assets, like critical industrial controllers or sensitive operational data, long before they are discovered. High-profile incidents like the Colonial Pipeline attack demonstrated the devastating real-world consequences of such intrusions. This reality is compounded by the unique challenges of OT environments. Traditional IT security tools often fail here because agent-based solutions are incompatible with specialized equipment like PLCs and HMIs. ICS specific malware (example FrostyGoop) is specifically targeting ICS protocols, such as Modbus. The critical need for uninterrupted operations makes patching legacy systems difficult, if not impossible. Organizations therefore need a different kind of defense, a preemptive and non-intrusive one that works.

Focus on Early detection by anchoring deceptions around early-stage MITRE tactics

(Deception design based on MITRE tactics used early in the attack lifecycle)

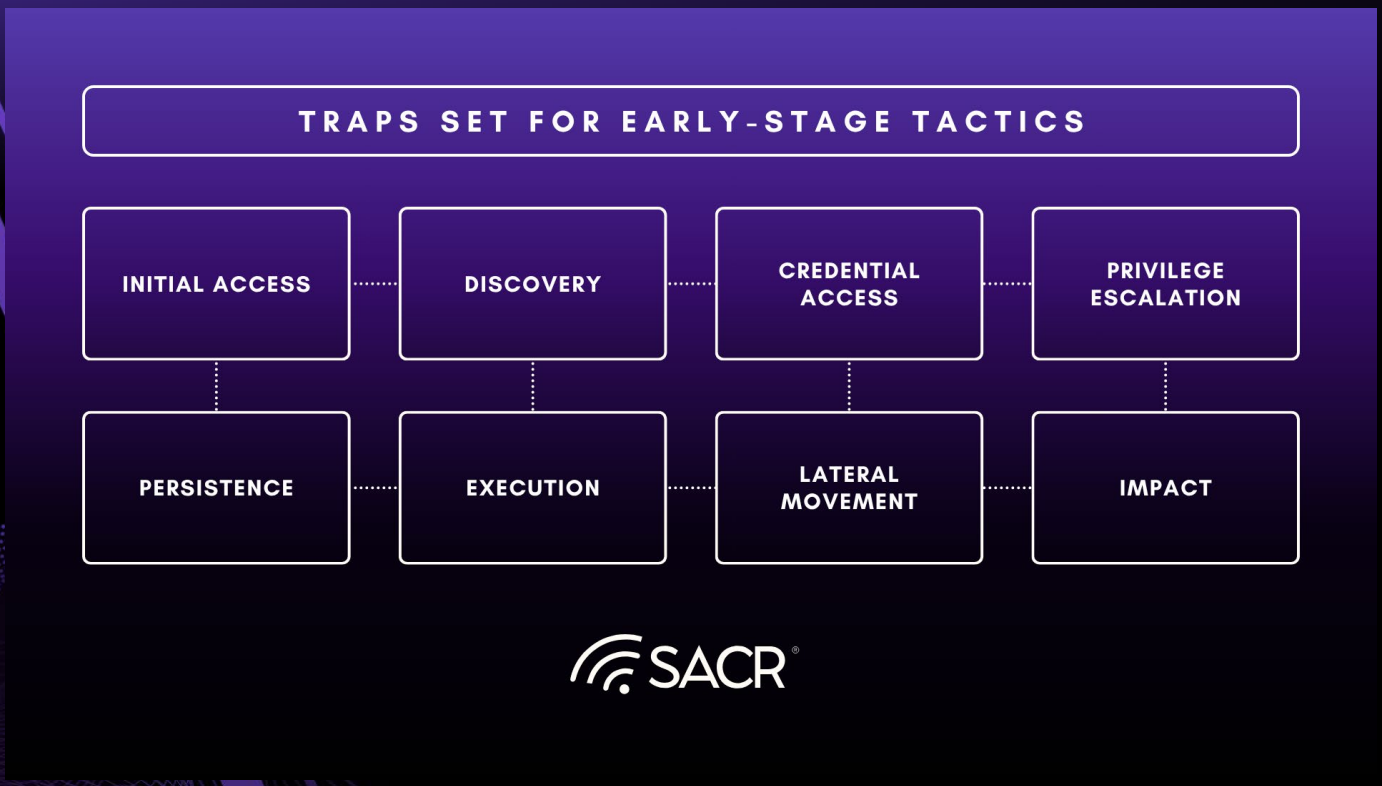

The Core Idea:
Setting a Trap with a Hall of Mirrors

Cyber deception is a proactive defense strategy that sets traps for attackers who are already inside the network. Instead of trying to block every possible entry point, it assumes a breach will eventually happen and focuses on detecting and neutralizing intruders before they can cause physical or financial harm. This strategy works by turning the network into a confusing landscape for the attacker, making it nearly impossible for them to know what is real and what is an illusion.

The most significant benefit of this approach is the generation of high-fidelity, actionable alerts. Deceptive elements are not part of any normal workflow, so no legitimate user should ever interact with them. Any alert triggered by a deceptive asset is an immediate and near-certain indicator of malicious activity. This effectively eliminates the problem of false positives that can overwhelm security teams. It turns the attacker's own methodical reconnaissance against them, the very data they collect to plan their attack is the data that leads them into the trap. Now, let's look at the specific components used to build this virtual hall of mirrors.

The Building Blocks of Deception in a modern industrial environment, the hall of mirrors is constructed using three primary types of deceptive elements. These elements are designed to exploit an attacker's systematic process of discovery and data collection. We can use simple analogies to explain each component.

The Generational Leap

Modern Deception vs. Old-School HoneypotsThe first generation of deception technology, known as honeypots, emerged in the 1990s. While innovative at the time, they have been surpassed by modern deception platforms designed for today's complex threat landscape.

Modern, dynamic tools represent a significant advance over older, less effective deception methods. Below, see how these technologies have transformed over the past few years.

Honeypots (The Old Way)	Modern Deception (The Future)
They are configured with static responses and cannot mimic the dynamic activity of a real, user-controlled network.	Dynamic and adaptive. Deceptions can change automatically, mimicking a real, living network.
Easier for attackers to identify as fakes. Telltale signs include if access seems too easy, systems still have factory default settings, or directories have obviously fake names like employee data.	Appears genuine and is nearly impossible for attackers to detect, creating a disorienting hall of mirrors.
Difficult and expensive to manage and deploy at scale across a large network.	Scalable and cost-effective, allowing for easy deployment across the entire network.
Covers only the specific points in the network where a physical or virtual system has been deployed.	Blankets the entire network attack surface, identifying attackers within three to four lateral movements, even if deceptions aren't on every machine.

SETTING A TRAP WITH A HALL OF MIRRORS

COMPONENT	ANALOGY	PURPOSE FOR A SECURITY TEAM
Decoys	Fake rooms or treasure vaults	These are fake assets designed to look like real servers, engineering workstations, Human Machine Interfaces (HMIs), or Programmable Logic Controllers (PLCs). They act as attractive destinations to lure attackers away from real assets, allowing defenders to safely observe their tools and tactics. For maximum authenticity, these decoys can represent assets from multiple OEM vendors like Honeywell, Rockwell Automation, and Siemens, and support industrial protocols such as Modbus, BACnet, Ethernet/IP, and S7.
Breadcrumbs	Fake jewels left on a table	Just like fake signposts, their purpose is to guide attackers away from real assets and toward the decoy treasure vaults" These misleading clues, such as fake credential profiles in a browser, cached Remote Desktop Protocol (RDP) sessions, or entries in the ARP and memory cache, are planted on real, production systems to lead intruders into the prepared trap.
Baits	Fake jewels left on a table	These are deceptive data files or configurations placed on real IT assets in the environment. They are designed to be stolen, enabling the security team to detect threats like data exfiltration the moment an attacker tries to take them.

SACR®

The Payoff: Early, Precise, and Actionable Detection

The primary goal of cyber deception is to turn the tables on the attacker, using their own methodical nature against them. A successful deception strategy delivers several key outcomes that strengthen an organization's defenses. These core benefits combine to create a powerful layer of internal security, leading to a smarter overall defense strategy.

- Detects Intruders Early: By following misleading breadcrumbs or probing a seemingly vulnerable decoy, attackers reveal themselves during their initial reconnaissance and lateral movement phases. This allows security teams to detect threats well before an intruder can reach critical industrial systems and cause significant damage.
- Disrupts the Attack: When attackers engage with false targets, their progress is slowed, their automated tools may break, and their kill chain is interrupted. This can detect advanced industrial attack tactics such as attempts to Impair Process Control or Inhibit Response Function, forcing adversaries to reveal their intentions in a controlled environment. This is especially valuable to disrupt automated, agentic AI attacks that are rapidly progressing toward their objective. In industrial environments, the decoys can detect advanced attack tactics.
- Provides Crystal-Clear Alerts: Since only attackers interact with deceptive elements, every alert is a high-fidelity indicator of compromise. This allows security teams to respond with speed and confidence, knowing they are dealing with a real threat and not wasting time on false positives.
- Gathers Valuable Intelligence: By observing an attacker interacting with an OT decoy, for example, attempting a Write Coil command on a fake PLC, security teams can reveal their tactics, techniques, and procedures (TTPs). This intelligence,understanding the how of an attack against your specific environment, is invaluable for strengthening overall defenses.

Cyber deception is a security layer built on a simple but powerful assumption: a breach will eventually happen. It shifts the focus from solely preventing intrusions to rapidly catching and neutralizing intruders By creating a disorienting hall of mirrors with a mix of decoys, baits, and breadcrumbs, this technology flips the script. It uses an attacker's own methodical approach, to collect data, analyze it, and calculate their next move, against them. This proactive, preemptive approach represents a necessary and powerful evolution in cybersecurity strategy, giving defenders a crucial advantage in the fight to protect modern industry from sophisticated digital threats.

Market Context

Acalvio emerged at a critical inflection point in the mid-2010s, as the cybersecurity industry faced a detection deficit where traditional perimeter defenses were failing to stop sophisticated, lateral-moving threats. At the time, breach detection took an average of over 200 days, and the only proactive defense, the honeypot, was largely dismissed by enterprises as a high-maintenance, non-scalable academic tool.

Emerging from stealth in 2016 with backing from top-tier VCs like Accel and GV (Google Ventures), Acalvio filled this market gap by pioneering Deception 2.0. They transformed deception from a manual niche into an automated enterprise capability by integrating data science and DevOps principles. This shift was driven by the market's urgent need for Internal Active Defense,a strategy that assumed the perimeter was already breached. By replacing static, easily fingerprinted traps with its patented Fluid Deception (which dynamically adjusts decoy interactivity), Acalvio successfully repositioned deception as a cost-effective, high-fidelity alternative to the typical noise of traditional SIEM and EDR alerts, eventually evolving to meet today's market demands for Identity Threat Detection and Response (ITDR) and AI-driven preemptive security.

THE EVOLUTION OF ACALVIO: FROM DECEPTION TO AI AUTONOMY

PERIOD	MILESTONE PHASE	KEY DEVELOPMENTS & INNOVATIONS
Circa 2017	**The "Deception 2.0" Era**	Launched ShadowPlex; published Deception 2.0 for Dummies; moved from static honeypots to automated, distributed deception.
Breadcrumbs Growth Years	**Patent & IP Foundation**	Secured 25+ U.S. Patents; introduced Fluid Deception (dynamic interactivity escalation) and Network Infrastructure Obfuscation.
Expansion Phase	**Diversifying the Surface**	Adapted ShadowPlex for OT/ICS (Industrial Control Systems) and entered the Identity Security (ITDR) space with Active Directory protection.
Current Era	**The GenAI Revolution**	Launched ShadowPlex Copilot; integrated LLMs for automated decoy generation and industry-specific content creation.
Preemptive Deception Platforms	**Acalvio: 360 Preemptive Deception**	Focuses on 360 deception, specifically the idea of making the real look fake (e.g., placing honey tokens on real systems) to combat machine-speed AI attacks. This complements the traditional deception approach of making fake things look real. Moving further towards realism-based deception.

SACR®

The future of Preemptive Deception Platforms:
How Generative AI can Advance Preemptive Cyber Deception

Generative AI (GenAI) automates the creation of cyber deceptions by shifting from static, manually configured traps to dynamic, context-aware strategies. This process relies heavily on Large Language Models (LLMs) and advanced frameworks like Structured Prompt Engineering (SPADE) to generate tailored deception artifacts and orchestrate defensive actions in real-time. Structured Prompt Engineering (SPADE) can be used as a core mechanism for automating deception creation and is a systematic framework known as Structured Prompting for Adaptive Deception Engineering (SPADE).

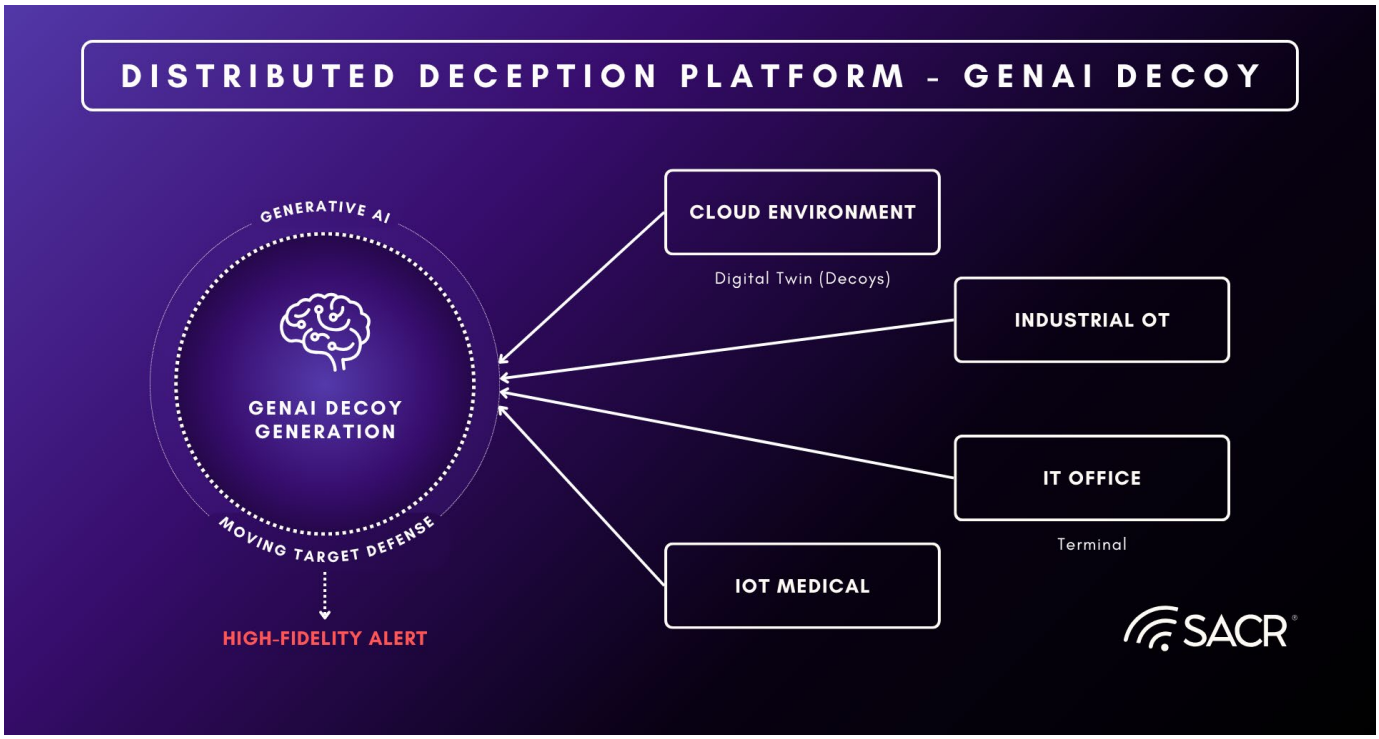

Here is how GenAI automates the creation of these deceptions:

Because LLMs are generic, they require precise instructions to produce actionable cybersecurity tools. This framework automates the generation of deception ploys (such as honeyfiles, API hooks, and honeytokens) through a six-component structure:

- Identity/Persona: The AI is assigned a specific role, such as cybersecurity expert, to align its output with domain-specific tasks.
- Goal/Task: The system explicitly defines the objective, such as generating a honeytoken to engage ransomware and to detect privilege escalation in the cloud.
- Threat Context: Real-time intelligence, derived from sandbox analysis of malware behaviors and TTPs (Tactics, Techniques, and Procedures) is embedded into the prompt. For example, if malware targets specific file directories, the AI uses this context to create targeted lures.
- Strategy Outline: The prompt includes operational constraints (Dos and Don'ts) to ensure the deception is resource-efficient and feasible.
- Output Guidance: The AI is given examples or templates (Few-Shot Prompting) to ensure consistency.
- Output Format: The system specifies the required format (e.g., JSON, XML, or C++ code) to ensure the generated deception is immediately deployable.

Autonomous Content Generation

GenAI models, such as ChatGPT-4o and Gemini and ilk, can automate the actual coding and creation of deceptive assets that would otherwise require manual labor. The rise of advanced Generative AI (GenAI) models, including powerful large language models (LLMs) like ChatGPT-4o and Gemini, is fundamentally transforming the creation and deployment of cybersecurity deception strategies. These sophisticated models can automate the actual coding and creation of deceptive assets, such as fake credentials, realistic-looking file systems, decoy network shares, or entire virtual environments, that would otherwise necessitate significant, labor-intensive manual effort by human security analysts.

This automation capability drastically reduces the time and specialized knowledge required to populate a deception environment with highly convincing and varied traps. Instead of tedious manual development, a security team can leverage GenAI to rapidly generate hundreds or even thousands of unique, context-aware deceptive artifacts, making it much harder for an intruder to distinguish between legitimate and bogus resources. This not only scales up the deception campaign but also allows for greater realism and variability in the decoys, enhancing the effectiveness of the overall security posture. In terms of outcomes, a consistent use of agentics on various endpoints and workloads can help maintain and create a more proof of life type deployment of deceptions essentially ushering in a realism that can both detect attackers better and move into the phase of real gamification of deceptions.

This includes:

- Honeyfiles & Honeytokens: The AI generates realistic fake documents or credentials designed to trigger alerts when accessed. For instance, creating specific file types known to be targeted by ransomware and utilizing specific types of non-human identities known to be triggered when AI agents are used for malicious purposes.

- API Hooks: GenAI can generate complex code (e.g., C++ for API hooking) to intercept malware calls. For example, if credential-stealing malware attempts to read a browser's login data, the AI-generated hook can intercept this request and return fake credentials, redirecting the attacker to a monitored honeypot.
- Decoy Personas: In commercial applications like Acalvio's ShadowPlex, an LLM-powered Copilot automatically generates decoy naming conventions and content tailored to the specific industry and threat surface, ensuring the decoys blend seamlessly with the real environment.

Agent-Based Decision Making

- **Automated Preemptive Deployment:** Beyond creating the assets, GenAI automates the decision of when and where to deploy them. In multi-agent environments (such as the CybORG CAGE 4 simulator), LLMs act as autonomous defender agents.
- **Natural Language Processing:** The network state and observation vectors are parsed into natural language that the LLM can understand (e.g., Suspicious Activity Detected).
- **Reasoning and Action:** The LLM processes these observations to select specific actions, such as deploy decoy or analyze. Unlike static scripts, the LLM provides a reasoning log for its decisions, such as deploying a decoy as a precautionary measure to identify potential red team activity early.
- **Communication:** LLM-driven agents can communicate with other agents using summarized security levels to coordinate defenses across different network subnets.

Fighting AI Using Cloak, Honey, Trap (CHeat) Framework

The CHeaT (Cloak, Honey, Trap) framework is a proactive, deception-based cybersecurity strategy designed specifically to defend networks against malicious, autonomous Large Language Model (LLM) agents. Rather than relying solely on traditional firewalls or access controls, CHeaT weaponizes the inherent cognitive and architectural vulnerabilities of AI, such as their context window limits, tokenization behaviors, and strict adherence to parsed instructions. By subtly weaving deceptive data directly into a system's environment, the framework manipulates the attacking AI's perception, tricking it into ignoring real targets, exposing its presence, or self-destructing its own logic.

Here is a breakdown of how the three pillars of the concept work together:

- Cloak (Hide): This phase focuses on obscuring genuine, sensitive data from the AI's view. By altering text formatting, exploiting tokenization quirks, or embedding subtle misdirection prompts within files, defenders can trick the LLM agent into classifying critical assets as irrelevant noise, causing it to bypass them entirely.
- Honey (Detect): Similar to traditional honeypots, this phase scatters enticing, fake digital artifacts, like synthetic API keys, bogus configuration files, or fabricated system logs, throughout the network. Because LLMs are inherently designed to seek out and process information, they are highly likely to interact with these honeytokens, instantly triggering an alert for the security team.
- Trap (Neutralize): The final phase actively sabotages the LLM using adversarial prompt injections planted within the environment. When the agent ingests this poisoned data, it can be forced into infinite reasoning loops, paralyzed by conflicting instructions, or even tricked into executing counter-attack code (like a reverse shell) on the attacker's own infrastructure (which of course may be limited to government authorized entities).

CHeat Framework Overview Graphic

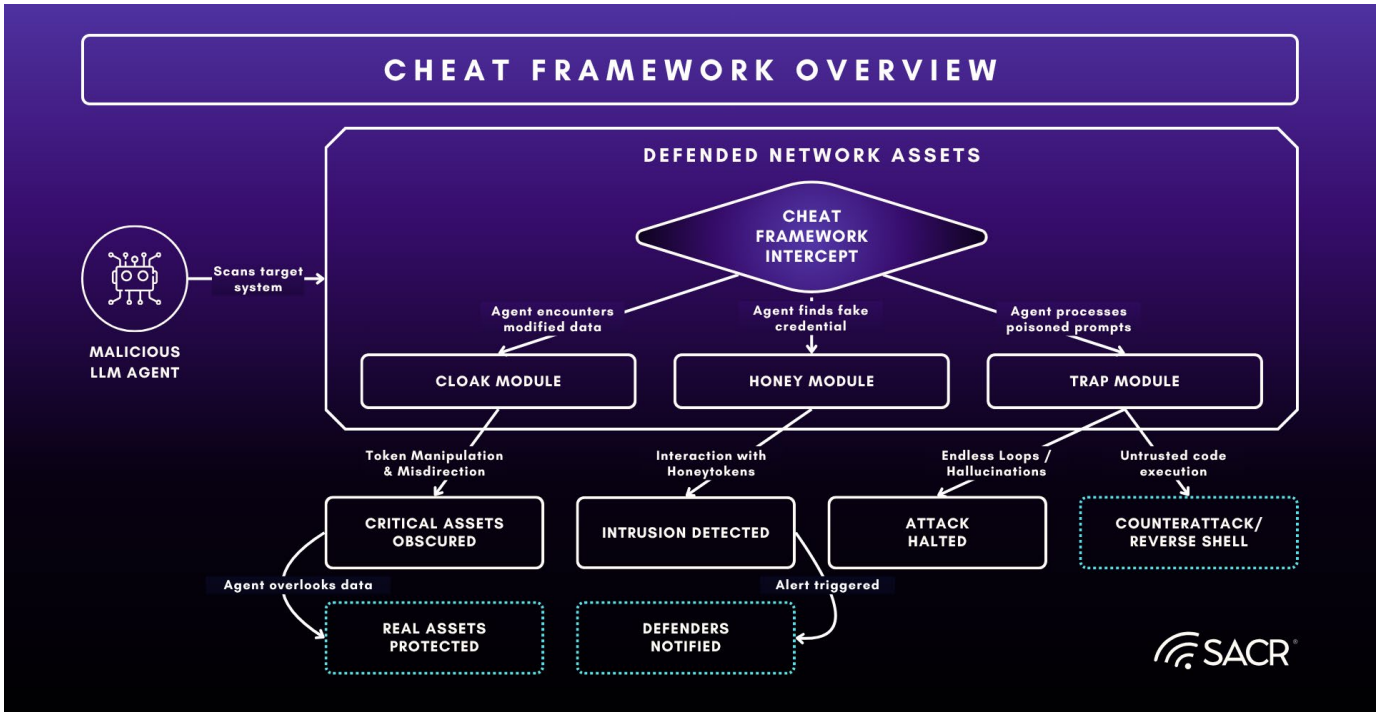

Overview of Preemptive Deception Platform Providers and Market Participants

The following section provides a high-level overview of the competitive landscape for Preemptive Deception Platforms. This analysis focuses on key vendors, including Acalvio's core offering, ShadowPlex, and explores how its architectural design and specialization in AI-driven identity and cloud security differentiate it from competitors, particularly those who have been acquired by larger security consolidators.

The Solution: Acalvio ShadowPlex Products / Services Overview

Acalvio offers ShadowPlex, an autonomous, AI-driven deception platform designed to detect and neutralize advanced threats inside the network. Unlike legacy honeypots, ShadowPlex uses Deception Farms to project low-interaction decoys and high-interaction traps across IT, OT, and Cloud environments without deploying agents on every endpoint.

Core offerings include:

- **ShadowPlex Advanced Threat Defense:** The flagship platform for automated decoy deployment and lateral movement detection.
- **ShadowPlex Identity Protection (ITDR):** Focuses on Identity Threat Detection and Response (ITDR) by planting honeytokens (fake credentials, service accounts) to detect credential theft and misuse.
- **ShadowPlex Cloud Security:** Agentless deception for cloud workloads and identity, specifically targeting cloud-native attack vectors.
- **Active Defense / 360 Deception:** A novel capability that inverts traditional deception by making real production assets appear fake or uninteresting to attackers, specifically designed to disrupt the reconnaissance logic of AI agents.

Acalvio's flagship product, ShadowPlex, is an autonomous, AI-driven deception platform. It is built on three core pillars:

Dynamic Adaptation (Fluid Deception)

GenAI enables Fluid Deception, where the level of interactivity and the nature of the decoy change based on attacker behavior. Unlike legacy honeypots that were resource-intensive and hard to manage, ShadowPlex uses AI to learn the network's topology. It automatically deploys realistic decoys (fake servers, workstations, databases) that blend seamlessly with the environment.

- Real-Time Refinement: The system allows for iterative refinement. If a generated deception ploy needs adjustment to fit operational constraints, the AI can refine the output based on feedback.
- Escalation: If an attacker interacts with a low-interaction decoy, the system can dynamically escalate to a high-interaction environment (e.g., a full OS) to capture more detailed threat intelligence, a process that GenAI can help orchestrate by analyzing the attack pattern.
- Real Time Threat Intelligence Morphs Deceptions: By integrating real-time threat intelligence with generative capabilities, GenAI transforms deception from a static minefield into an adaptive, automated defense layer that evolves alongside the attacker.

Core Benefit: Low operational overhead. Security teams do not need to manually configure every decoy; the system adapts as the network changes.

Identity Threat Detection & Response (ITDR)

With identity being the new perimeter, Acalvio places Honeytokens (fake credentials, cookies, and keys) on real endpoints.

- Mechanism: If an attacker scrapes memory or mimics a user to steal these credentials and tries to use them, an alert is triggered immediately.
- Core Differentiation: Since no legitimate user has a reason to touch these fake assets, the false positive rate (FPR) is near zero.

Cloud & Cloud-Native Security

ShadowPlex extends active defense to public clouds (AWS, Azure, GCP) and Kubernetes environments, planting deceptive storage buckets, serverless functions, and IAM roles to catch cloud-specific lateral movement.

Vendor Introduction

Acalvio Technologies offers Preemptive Cyber Defense, specializing in autonomous deception and identity threat detection. Founded by veterans from the security industry, the company has pioneered the use of AI to automate the deployment of deception farms at enterprise scale. Acalvio distinguishes itself by moving beyond simple honeypots to a platform that weaves deception into the fabric of the network, identity, and cloud layers

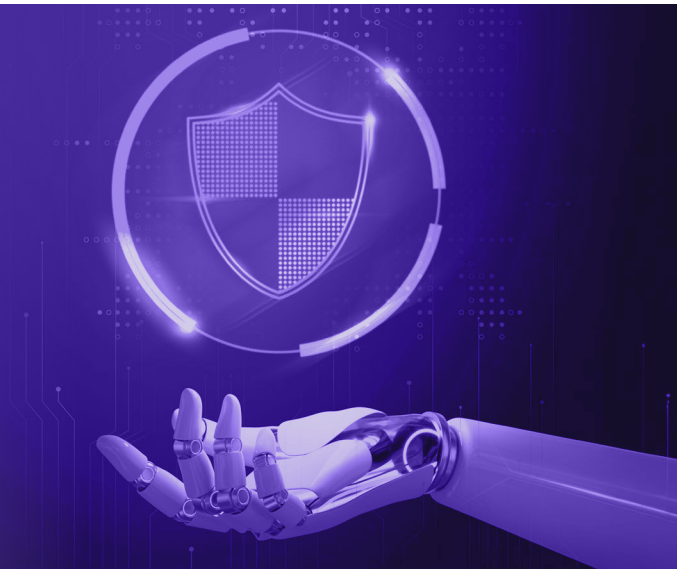

Acalvio: Company History & Founding Story

The Origin Story

Acalvio Technologies was founded in Silicon Valley by a team of security industry veterans, including Ram Varadarajan (CEO), Dr Sreenivas Gukal (CPO), and Raj Gopalakrishna (Chief Architect), to revolutionize the Deception category of cybersecurity. The company was built on a foundation of deep technical innovation, backed by over 25 issued patents in Autonomous and Fluid Deception technologies.

The mission was to solve the fundamental flaws of Deception 1.0. Legacy honeypots were notoriously difficult to manage, requiring heavy, dedicated virtual machines and constant manual updates. Acalvio's vision was to create ShadowPlex which offers a highly automated, enterprise-scale platform that could deploy credible decoys at a fraction of the traditional resource cost.

Market Disruption & The ShadowPlex Edge

Rather than simply mimicking static servers, Acalvio introduced Fluid Deception, allowing decoys to adapt and move within a network environment. This shifted the paradigm from passive defense to active engagement, making it nearly impossible for attackers to distinguish between real assets and shadow assets.

Strategic Traction & Growth

Acalvio's rapid ascent in the market, from its first ~20 customers to its current status serving Fortune 500 companies and government agencies, is credited to its vendor-agnostic integration strategy.

- Low Barrier to Entry: Unlike competitors who required siloed management, Acalvio integrated directly into the tools SOC teams already lived in, such as CrowdStrike Falcon, Microsoft Defender, and Palo Alto Cortex XDR.
- Zero-Customization Deployment: By eliminating the need for custom field programming or complex hardware, they made advanced deception accessible to midsize enterprises, not just elite security teams.
- Scale: Today, the platform is recognized for its ability to protect hybrid cloud environments and critical infrastructure globally.

Acalvio has moved through three key stages, including Founding and Early Years, Growth and Expansion, and Market Leadership with their more mature platform concept including identity, decoys and higher levels of automation and orchestration.

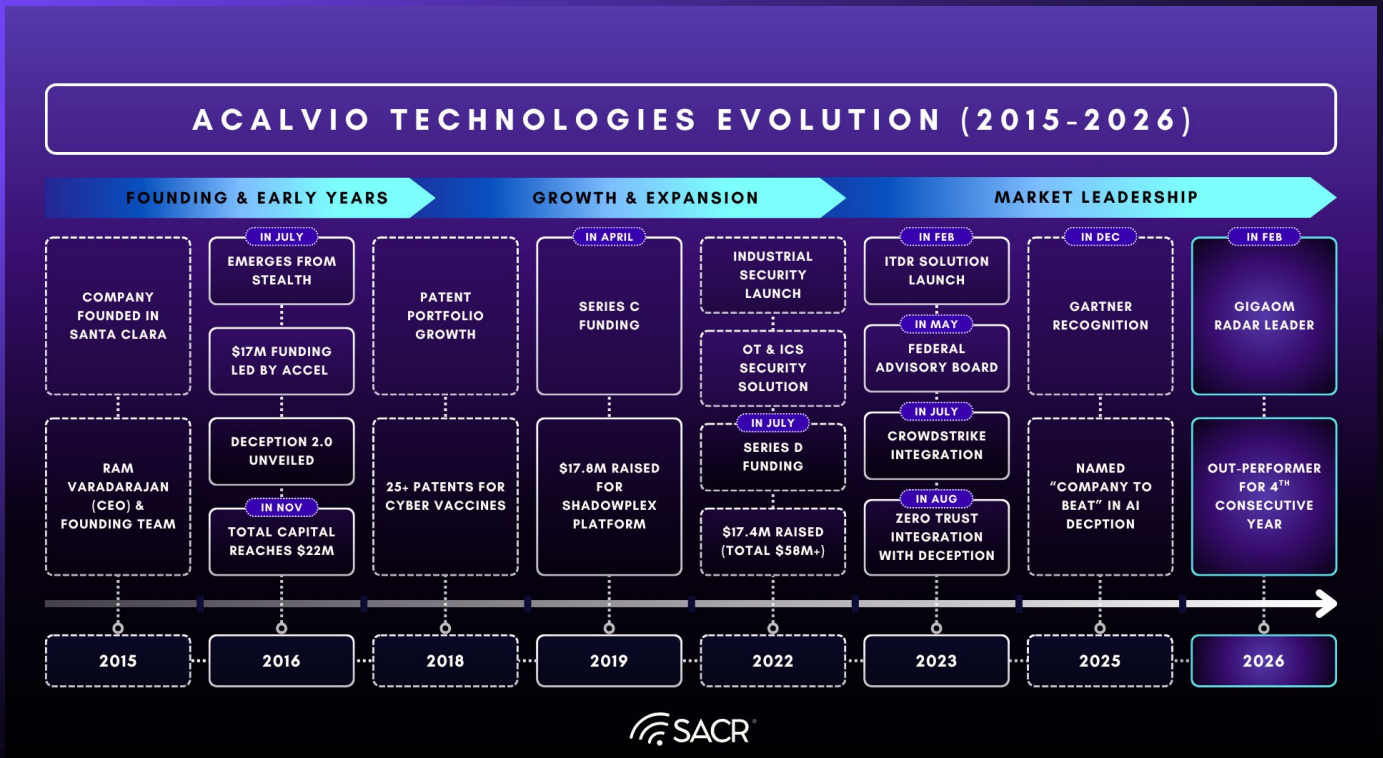

Customer Traction & Viability

Acalvio has secured significant wins in high-stakes environments, including Fortune 50 enterprises and the US Navy, where they won first place in a cyber defense challenge for their ability to deflect attacks. The company reports a record year in 2025, driven by the increasing need for active defense against automated threats.

- High-Fidelity Signal: Alerts are based on interaction with fake assets, meaning they are almost always true positives.
- Ease of Deployment: Deep integrations with CrowdStrike and Microsoft Defender allow for agentless or agent-piggybacked deployment of honeytokens, removing friction.
- Counter-AI Readiness: Their Inverse Deception approach addresses the emerging threat of AI-driven reconnaissance better than static defense tools.

Competitive Advantage & Market Analysis

Acalvio leads in a market against players like Proofpoint, Fortinet , and SentinelOne.

Acalvio differentiates itself through:

- Architectural Superiority: Its DevOps style architecture allows for agentless deployment. It does not require a heavy footprint on every endpoint, reducing friction for IT teams.
- FedRAMP Ready Status: Acalvio's achievement of FedRAMP Ready status underscores its reliability and security, making it a viable option for high-stakes US Federal Government deployments.
- Industry Recognition: Consistently cited by Gartner as a Tech Innovator in key active defense and deception categories, validation that their technology is ahead of the curve.

Product Ecosystem Deep Dive

The Platform View: The Projection & Automation Engine

Acalvio's architecture breaks away from legacy Deception 1.0 models that required heavy, dedicated virtual machines for every decoy. Instead, Acalvio ShadowPlex operates as a highly automated Projection Engine driven by artificial intelligence.

- The Core Mechanism (Deception Farms): The architecture is built on a patented concept called deception farms. Instead of placing physical or virtual machines everywhere, a centralized resource footprint *projects* thousands of lightweight, fluid decoys and honeytokens across IT, OT, cloud, and identity networks.
- The AI Orchestrator: The platform continuously and autonomously analyzes the production environment. It discovers real assets and uses generative AI and machine learning to automatically generate, deploy, and refresh decoys so they perfectly mirror the changing environment.
- The Integration Layer: Rather than forcing analysts into a new pane of glass, ShadowPlex acts as a verified signal generator that feeds directly into an organization's existing Security Operations Center (SOC) nervous system, meaning it natively hooks into platforms like CrowdStrike, Microsoft Defender, and Palo Alto Cortex to trigger automated containment.

Component Breakdown

Core Product: ShadowPlex Advanced Threat Defense This is the flagship offering that provides the Active Defense Fabric. It automatically deploys network, endpoint, and application decoys (along with breadcrumbs and lures) to detect reconnaissance, lateral movement, and privilege escalation with absolute certainty, converting silent intrusions into high-fidelity alerts.

Expansion Modules (Secondary products driving NRR) Acalvio expands its footprint and drives Net Revenue Retention (NRR) by offering domain-specific deception modules that attach to the core platform:

- ShadowPlex Identity Protection: Focuses heavily on Identity Threat Detection and Response (ITDR). It deploys deceptive identity artifacts and honeytokens to protect Active Directory and detect identity compromise and credential theft.
- ShadowPlex Cloud Security: An agentless, multi-cloud module that uses native cloud APIs to project honeytokens and decoys into cloud workloads and IAM structures.
- ShadowPlex Targeted Threat Intel (TTI): Moves deception to the external perimeter. It deploys external-facing decoys (like fake web apps, APIs, and IPv6 IIoT devices) to catch automated password spraying, credential stuffing, and AI-driven reconnaissance before attackers even breach the internal network.

The Acalvio Secret Sauce

- Patented Fluid Deception: Acalvio's projection-based architecture allows for exponential enterprise scale with minimal operational overhead and a drastically lower Total Cost of Ownership (TCO) compared to dedicated-host competitors.
- AI-Driven Authenticity (Dynamic Realism): To prevent attackers from fingerprinting the traps, Acalvio's AI continuously rotates and refreshes the decoys so they never become stale. The decoys exhibit behavioral fidelity that makes them indistinguishable from real assets.
- Fiercely Vendor-Agnostic Ecosystem: Unlike competitors locked into proprietary security fabrics (e.g., Fortinet or Fidelis), Acalvio is built to integrate anywhere. It leverages existing EDR/XDR agents (like CrowdStrike Falcon or Microsoft Defender) to achieve agentless deployment and orchestrate automated incident response without requiring custom field programming.

User Experience (UX)

Acalvio's design philosophy is built for SOC Analysts and Security Executives, aiming to solve the problem of alert fatigue.

- Executive & Analyst Focus (Noise Reduction): The platform is designed to provide *certainty*. Because legitimate users never interact with decoys, any alert generated by Acalvio is essentially a verified, high-confidence signal.
- Pre-Packaged Playbooks: Rather than forcing developers or engineers to build custom traps manually via a CLI, Acalvio provides AI-driven deception playbooks. These offer pre-packaged, automated use cases (like ransomware containment or insider threat detection) for immediate time-to-value.
- Workflow Invisibility: Acalvio is designed to be invisible not just to attackers, but mostly to defenders. The ShadowPlex Administration Console is used for easy setup, but day-to-day UX happens in the tools the SOC already uses. For instance, verified deception alerts flow naturally into Microsoft Sentinel workbooks or CrowdStrike Threat Graphs, meaning the primary UX is a seamless enhancement of the user's existing dashboards.

Sector-Specific Impact

Healthcare

- Challenge: Protecting vast, open networks with critical IoT/IoMT devices (MRI machines, pumps) that cannot run standard security agents.
- Acalvio Application: ShadowPlex projects decoys that mimic these medical devices. If ransomware attempts to spread from a nurse's station to an MRI network, it hits a decoy first, triggering containment before patient care is impacted.

Financial Services

- Challenge: SWIFT networks and high-value databases are prime targets for "low-and-slow" attackers who steal credentials to move quietly.
- Acalvio Answer: By seeding fake Shadow Admin accounts and database connections, Acalvio traps attackers during their reconnaissance phase, long before they can authorize a fraudulent transfer.

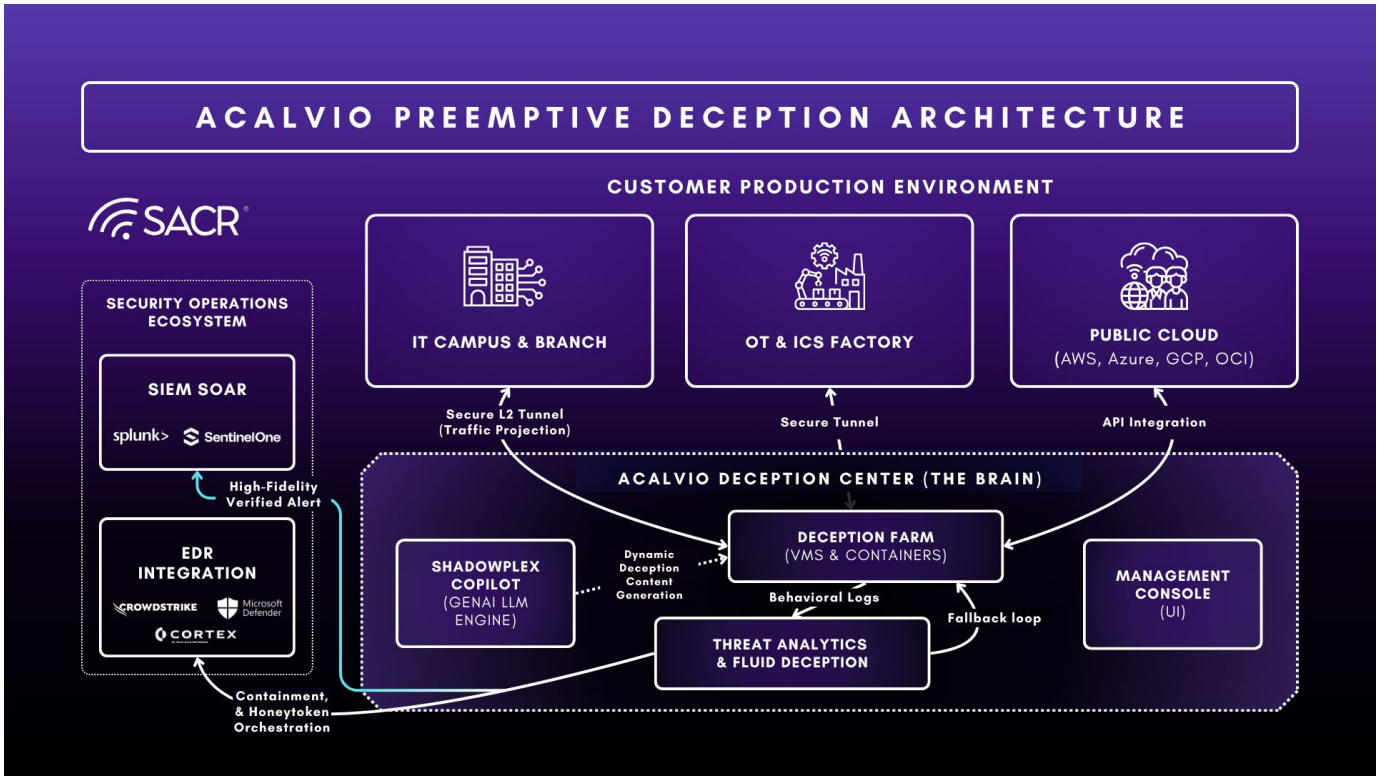

Go-to-Market (GTM) Strategy

Sales Motion: Acalvio primarily utilizes a Top-down Enterprise and Partner-led sales motion. Rather than a Product-Led Growth (PLG) model, Acalvio's ShadowPlex platform is an enterprise-grade solution designed for midsize to Fortune 500 companies and government agencies. Because its deployment involves mapping the organization's entire digital terrain (IT, OT, Cloud, and Identity) and integrating with core security infrastructure, the sales cycle requires strategic buy-in at the enterprise level. Acalvio extends its reach through a dedicated partner program, offering its platform via Managed Service Providers (MSPs) to organizations that require flexible or managed deployments.

The Ecosystem: Acalvio leverages a highly collaborative, vendor-agnostic ecosystem to scale and embed itself into modern security architectures.

- Technology Alliances: Instead of competing with major security platforms, Acalvio acts as a force multiplier. It features deep, native API integrations with leading EDR, XDR, SIEM, and SOAR providers, including CrowdStrike Falcon, Microsoft Defender, Palo Alto Networks Cortex XDR, and VMware Carbon Black. This allows Acalvio to feed high-fidelity alerts directly into the tools security teams already use, speeding up deployment and lowering the Total Cost of Ownership (TCO).
- Service Providers: Acalvio leverages Managed Service Providers (MSPs) to deliver its capabilities as a fully managed cloud service, which isolates attacker traffic from the customer's production network and simplifies operations.
- Public Sector Reach: To capture the government market, Acalvio maintains a Federal Advisory Board, ensuring its solutions align with strict national security and compliance requirements.

Customer Persona: Acalvio targets both strategic decision-makers and operational security leaders:

- The Strategic Buyer (CISO / CIO): The CISO is the primary economic buyer, motivated by the strategic need to implement Zero Trust architectures, protect critical assets like Active Directory and cloud resources, and shift the organization from a reactive posture to a Preemptive Cybersecurity model. They value Acalvio's ability to provide a measurable ROI by limiting breach impact and reducing dwell time.
- The Operational Buyer (Head of SOC / Incident Response Lead): These leaders are tasked with improving the efficiency of their teams. They buy Acalvio to solve the alert fatigue problem, as Acalvio's deterministic alerts eliminate false positives and streamline incident response workflows.
- The End User (Threat Hunters & Security Analysts): Threat hunters use Acalvio to proactively deploy dynamic deceptions into live environments to test hypotheses and uncover stealthy Advanced Persistent Threats (APTs) or insider threats that bypass traditional tools.

Financial & Growth Analysis (The Data Layer)

- Financial Outlook: Positive revenue and growth outlook is based on Acalvio's long term presence in the market as it serves midsize to Fortune 500 companies and government agencies. Its growth is driven by its recognition as a leader in Preemptive Cybersecurity and Automated Moving Target Defense (AMTD) by multiple major analyst firms.
- Efficiency Metrics: From a product architecture standpoint, Acalvio boasts a highly efficient model. Its patented Fluid Deception and deception farms architecture allows for the projection of thousands of lightweight decoys without the need for dedicated virtual machines or heavy infrastructure. This agentless, automated deployment significantly reduces the hardware and third-party OS licensing requirements, resulting in a substantially lower Total Cost of Ownership (TCO) compared to legacy or high-interaction competitors like CounterCraft.

Competitive Landscape
Displacing Traditional Security Solutions

Acalvio is primarily displacing legacy Deception 1.0, specialized provider platforms and static honeypots. These traditional systems relied on manual configuration, appliance-centric models, and heavy, dedicated virtual machines (VMs) or hosts for every decoy deployed, which severely limited enterprise scalability and increased operational overhead. Acalvio displaces the reliance on reactive, traditional security controls, such as standard Endpoint Detection and Response (EDR) or behavior-based analytics, which struggle to detect stealthy, identity-based attacks and lateral movement before a breach escalates.

Acalvio's Competitive Landscape and Architectural Rivals

Acalvio's modern rivals fall into three distinct categories based on their architectural approach and market consolidation:

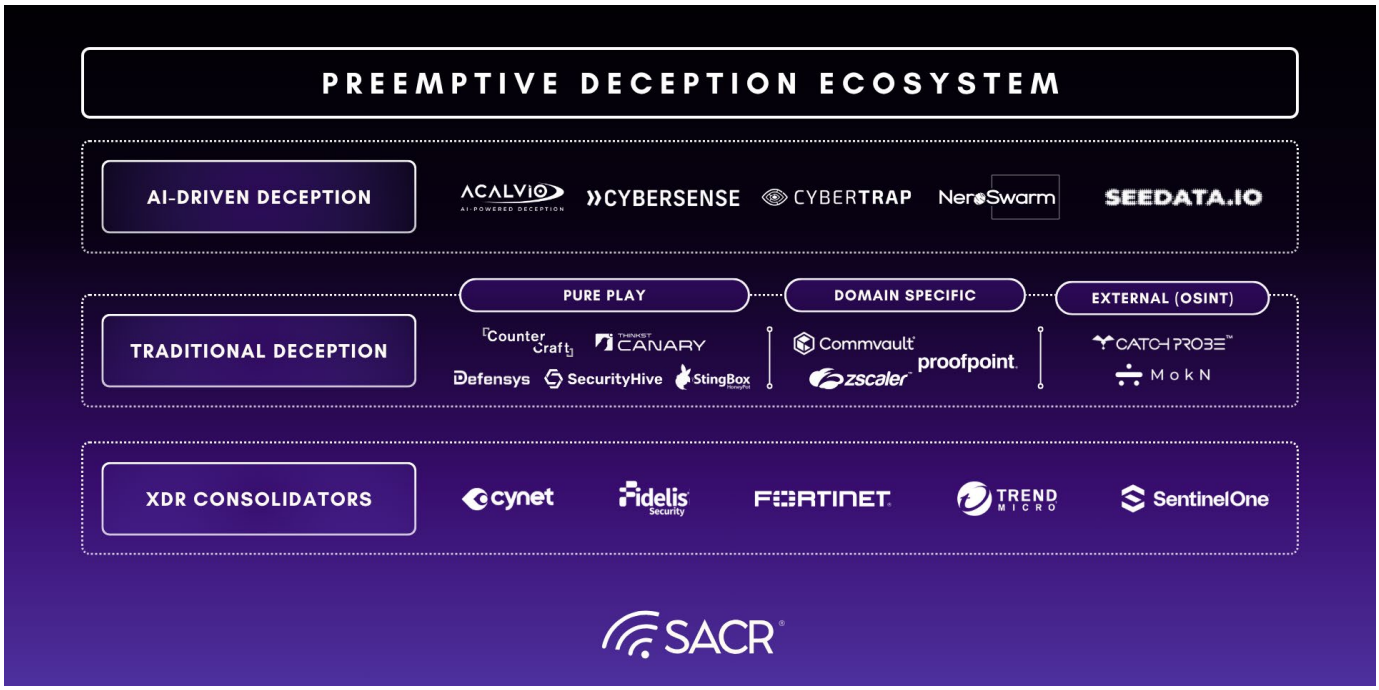

- **The XDR & Network Consolidators (The Acquired Rivals):** Major cybersecurity platforms have acquired deception startups to build native integrations. Acalvio competes against these bundled offerings by maintaining a vendor-agnostic approach.

Rivals here include:

- **SentinelOne (Attivo Networks):** Heavily focused on Identity Threat Detection and Response (ITDR) and Active Directory protection within the Singularity platform.
- **Proofpoint (Illusive Networks):** Focuses on human-centric security and preventing lateral movement.
- **Zscaler (Smokescreen):** Integrates deception into cloud-native Zero Trust and SASE architectures.
- **Fortinet (FortiDeceptor) & Fidelis Security:** Provide deception integrated directly into their specific network fabrics and XDR platforms, though they often rely more on manual configuration or are limited to their own ecosystems.

The Specialists & High-Interaction Rivals:

- **CounterCraft:** A primary rival that focuses on deep adversary engagement and threat intelligence gathering using a dedicated host architecture (full VMs for each decoy) rather than Acalvio's lightweight projection model.

The Pragmatists & Niche Upstarts:

- **Thinkst Canary:** Known for extreme simplicity and low false positives using an appliance-centric, manual roll-your-own deployment model, though it lacks Acalvio's automated enterprise scale.
- **MokN & CatchProbe:** Emerging upstarts focusing on the external attack surface. MokN uses SaaS-based fake doors to catch stolen credentials on the public internet, while CatchProbe merges internal decoys with dark web and OSINT monitoring.

Acalvio's Competitive Moat: Technology, Ecosystem, and AI-Driven Automation

Acalvio's competitive moat is built on highly differentiated proprietary technology, operational efficiency, and deep ecosystem entanglement:

- **Patented Architecture (Technological Moat & Low TCO):** Acalvio is protected by over 25 patents, centered around its Deception 360 and deception farm architecture. Instead of requiring a costly, heavy virtual machine for every decoy, Acalvio projects thousands of lightweight, dynamic decoys and honeytokens from a minimal centralized resource footprint. This drastically reduces the Total Cost of Ownership (TCO) compared to rivals like CounterCraft.
- **Ecosystem Integration (High Switching Costs & Data Gravity):** Unlike Fortinet or Fidelis, which lock deception into their proprietary fabrics, Acalvio is fiercely vendor-agnostic. It creates workflow gravity by establishing deep, native API integrations with the platforms SOC teams already use, such as CrowdStrike Falcon, Microsoft Defender, and Palo Alto Cortex XDR. By feeding high-fidelity, deterministic alerts directly into these existing SIEM/SOAR/XDR workflows to trigger automated containment, Acalvio embeds itself seamlessly into the organization's daily operations, creating high switching costs.
- **AI-Driven Automation (Operational Moat):** To combat the cybersecurity talent gap, Acalvio utilizes AI to automate the entire deception lifecycle. The platform automatically maps the network terrain, recommends decoy placement, and continuously refreshes deception assets so they blend perfectly with the evolving production environment. This dynamic realism prevents attackers from fingerprinting the decoys and eliminates the heavy administrative burden required by manual tools like Thinkst Canary.

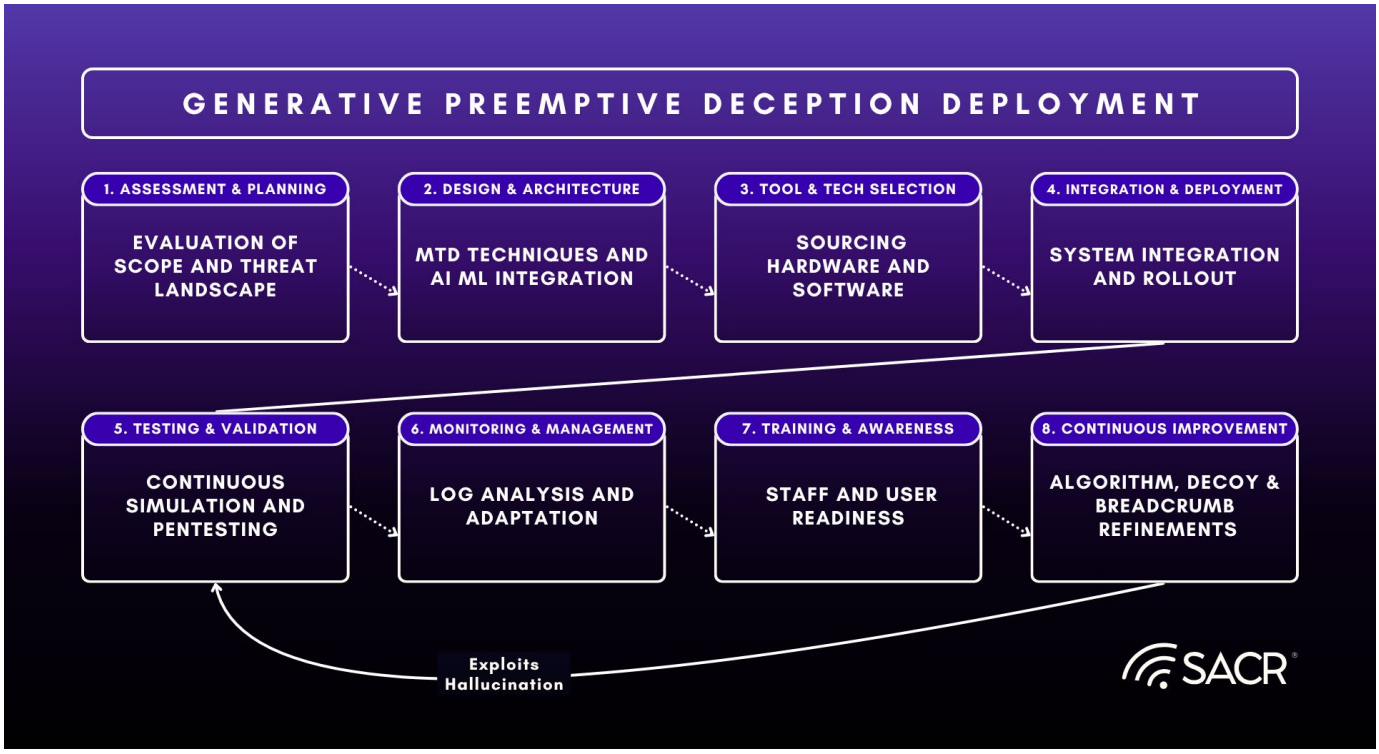

Future Outlook

Acalvio's future is tied to the cybersecurity industry's structural shift away from reactive detection and toward Preemptive Cybersecurity and AI-enabled Active Defense. The market is moving to Preemptive Cyber Deception Platforms (PCDP). As threat actors rapidly industrialize using agentic AI" and automated attack chains, Acalvio's technology is positioned to counteract these machine-speed threats by weaponizing AI for the defender.

Here is the updated future outlook for Acalvio, incorporating the critical context of AI-enabled deception and preemptive defense:

1. Acalvio Continues to be Driving the Shift to Preemptive Cybersecurity

The global cybersecurity paradigm is transitioning from a reactive posture (waiting for indicators of compromise) to a Preemptive framework. Acalvio is heavily positioned as a leader in this space.

- **Intent-Based Detection:** Traditional Endpoint Detection and Response (EDR) and XDR platforms are increasingly overwhelmed by the noise of modern attacks. Acalvio's future lies in its ability to provide preemptive, deterministic alerts; because its decoys have no legitimate business value, any interaction provides immediate, verified evidence of malicious intent *before* actual damage occurs.
- **Analyst Validation:** Firms like Gartner explicitly view Acalvio's brand of deception as a strategic imperative to get ahead of attackers, allowing defenders to actively manipulate adversary perception rather than just containing breaches post-intrusion.

2. AI-Enabled Deception to Counter Agentic AI Attacks

The threat landscape is entering an AI vs. AI (Bot vs Bot) arms race, where adversaries use Large Language Models (LLMs) and autonomous agents to conduct machine-speed reconnaissance, vulnerability discovery, and credential abuse.

- **Countering Malicious Automation:** Acalvio's AI-driven "Fluid Deception" architecture directly counters these AI-augmented tactics. By deploying vast numbers of realistic decoys and synthetic credentials, Acalvio absorbs automated scans and injects uncertainty into the attacker's data stream.
- **Dynamic Realism:** Future AI deception systems must resolve the fidelity paradox by using generative AI to create context-aware, highly interactive decoys that can fool both human and automated attackers without high operational risk. Acalvio's AI analyzes the production environment to automatically generate and continuously refresh decoys that perfectly mirror genuine network assets, making it computationally expensive and highly confusing for attacker AI to distinguish real from fake.

3. Leadership in Automated Moving Target Defense (AMTD)

Automated Moving Target Defense (AMTD) is a foundational method for dynamically shifting an organization's attack surface. Acalvio has been recognized as a Tech Innovator in this specific capability.

- **Machine-Speed Reconfiguration:** Moving forward, Acalvio's AMTD capabilities will allow enterprise networks to continuously and autonomously reconfigure themselves (e.g., changing IP addresses, ports, or application environments) in real-time.
- **Synergy with Deception:** By combining the unpredictable, constantly shifting attack surface of AMTD with AI-generated cyber deception traps, Acalvio forces attackers to expend massive resources just to navigate the network, drastically increasing their chances of stumbling into a high-fidelity trap.

4. Identity-Centric Deception for Zero Trust Architectures

With traditional perimeters dissolving, identity has become the primary attack surface. Over 80% of breaches involve compromised identities.

- Honeytokens at Scale: Acalvio's future relies heavily on identity-centric deception. This involves deploying thousands of deceptive honeytokens (fake credentials, API keys, and Active Directory objects) across cloud workloads, SaaS apps, and endpoints.
- Zero Trust Integration: In modern Zero Trust environments (never trust, always verify), Acalvio acts as the tripwire. It detects lateral movement and credential misuse that bypass traditional controls, immediately feeding this telemetry into XDR and SIEM platforms (like CrowdStrike or Microsoft Sentinel) to trigger automated containment.

5. Proactive Threat Intelligence and External Defense of Generative AI Threats

Acalvio is expanding its preemptive capabilities beyond the internal network to identify threats before they even breach the perimeter.

- LLM and IoT Targeted Threat Intel (TTI): Through products like ShadowPlex TTI, Acalvio deploys external-facing decoys (such as fake web apps, APIs, and IPv6 IoT devices) to monitor the public internet and help identify and create telemetry on LLM (Generative AI) based threats.
- Phish-Back and Early Warning: This allows organizations to catch automated password spraying, credential stuffing, and AI-driven reconnaissance early. By logging attacker IPs and tactics at the perimeter edge, Acalvio delivers highly specific, actionable threat intelligence to proactively harden internal defenses.

Strategic Product or Services Roadmap (Acalvio)

Acalvio is not just another layer in the defense-in-depth stack, it is a force multiplier. By forcing attackers to second-guess every asset they touch, it imposes a cost on the adversary, slowing them down and making them visible. As it advances towards more AI centric orchestration and agentics, expect to see realism emerge that is fully indistinguishable from human activities.

Product roadmap graphical example: Optional. A simple 3-horizon timeline works well here (Now → Next → Later), mapped to Platform, Ecosystem, and Go-to-market.

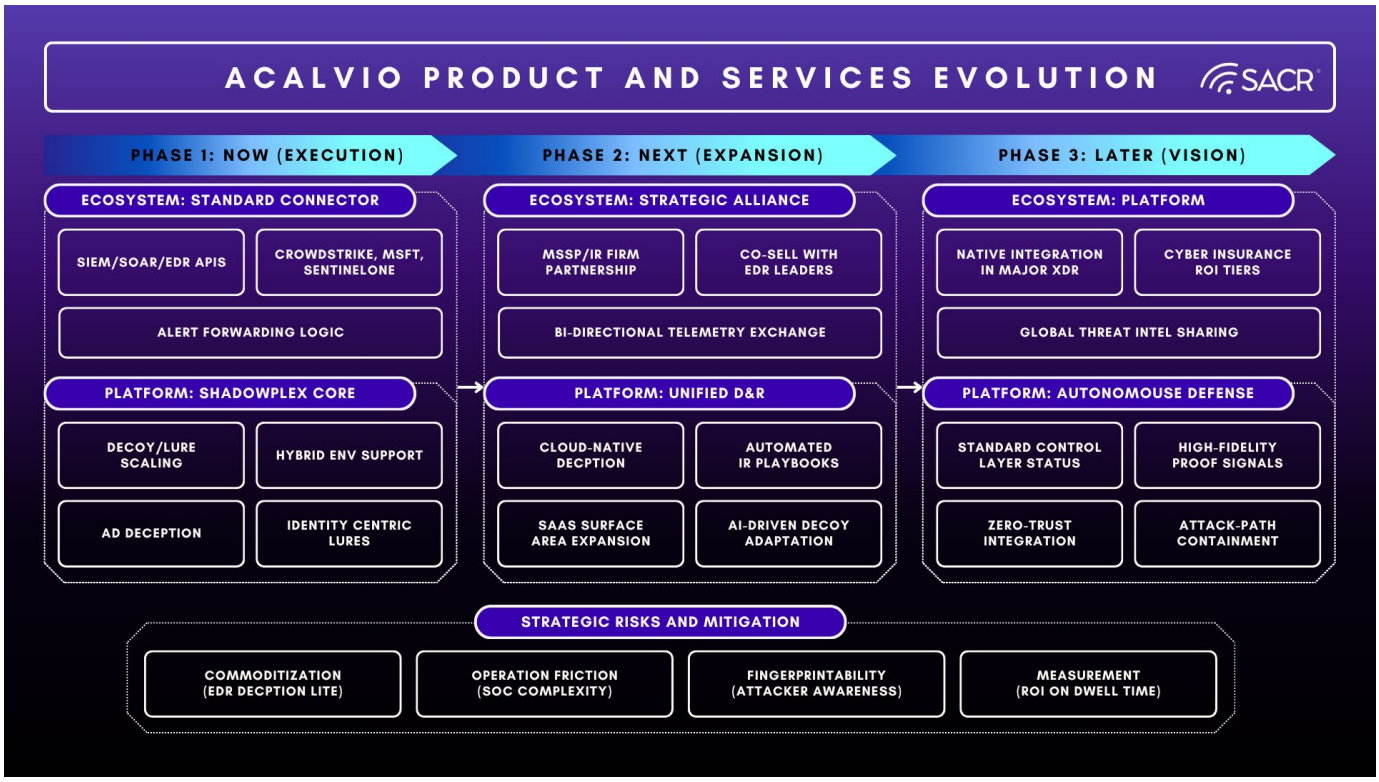

- The Next Frontier: Where Acalvio likely goes from here
  - From product to platform: Expand beyond deception deployments into a broader Deception + Detection + Response platform that is easier to operate at scale.
  - Deeper integration into SecOps workflows: Stronger integrations with SIEM, SOAR, EDR, and ITSM so deception signals become actionable incidents rather than isolated alerts.
  - Cloud and identity-aware deception: More coverage for cloud workloads, SaaS, and identity attack paths, with deception that adapts to modern hybrid environments.
  - Operational scale and managed options: More packaged deception-as-a-service / managed deployments for teams that want outcomes without heavy tuning.
  - Commercial path options:
    - Platform expansion first (most likely).
    - Partnership-led growth (MSSPs, IR firms, security platforms).
    - M&A outcomes: Either acquiring adjacent capabilities (automation, exposure validation) or becoming an acquisition target for a larger security platform.

- The End State: If Acalvio succeeds fully, what does the industry look like?
  - Deception becomes a standard control layer in enterprise security programs, like endpoint telemetry or IAM controls.
  - Security teams routinely rely on high-confidence adversary interaction signals (engagement with decoys, lures, fake credentials, fake data stores) to detect real intrusions faster and with fewer false positives.
  - Attackers face an untrustworthy environment by default: Lateral movement and credential abuse become more expensive because the environment is instrumented with believable traps.
  - SOC operations shift toward proof-based detections (attacker touched a controlled asset) instead of inference-heavy alerts (suspicious-but-uncertain signals).
  - Realism through Agentics: Agentics can build a realism that is virtually indistinguishable from humans in the preemptive deception platforms of the future.
  - Agentic Intelligence and AI defense: Use of Generative AI tampering and manipulation techniques and emerging agentic execution disruption technologies will augment future TTPs, reputation feeds and defenses against agentic AI based attacks.
- Risks: What could go wrong?
  - Commoditization / platform bundling: Larger security vendors (cloud providers, endpoint, identity, CNAPP) add good enough deception features as part of broader suites.
  - Operational friction: If deception requires too much tuning, maintenance, or careful environmental alignment, adoption stalls outside of highly mature teams.
  - Trust and realism gaps: If decoys are not realistic, or can be fingerprinted, sophisticated adversaries may avoid them, reducing value perception.
  - Measurement problem: If ROI is not clearly communicated (reduced dwell time, fewer high-severity incidents, better containment), leadership may treat it as nice-to-have.
  - Procurement headwinds: Deception can be seen as overlapping with detection and response spend, making it vulnerable during consolidation cycles.

SACR Key Take Away:

The strategic imperative for CISOs today is to shift from reactive detection (EDR/XDR) to preemptive defense against sophisticated, AI-driven threats. Acalvio's ShadowPlex platform is critical in this transition, as it effectively solves the persistent problem of dwell time and alert fatigue. By deploying dynamic, AI-driven deception across IT, OT, and Identity layers, Acalvio ensures that any interaction with a deceptive asset provides a near-certain, high-fidelity signal of malicious activity, closing the detection gap left by traditional tools and enabling rapid containment. Implementing this deception layer allows the organization to adopt a proactive posture, reduce breach impact, and deliver a measurable ROI by forcing attackers to waste resources and exposing them during early reconnaissance or lateral movement. For organizations struggling with alert fatigue and the fear of undetected lateral movement, Acalvio offers a high-fidelity, low-noise solution that delivers immediate ROI by closing the detection gap that EDR and SIEM leave open.

Software Analyst®
Cyber Research

business personal

softwareanalyst.io

SACR®

Trusted research. Sharp insights. Real conversation.

CISO	VENDOR
SECURITY TEAMS	INVESTORS