AMaze Use Cases

Proactive Cyber Resilience across IT, OT and AI

INTRODUCTION

This document presents four primary use cases for MirrorMire AMaze™ — one for operational technology, two for information technology, one for AI — and an additional seventeen field-ready scenarios drawn from real customer deployments. Each describes the threat scenario, the way the platform engages it, and the measurable benefit for security operations, plant engineering and organisational resilience. Every "how" is grounded in capabilities that ship in the current AMaze™ platform — REST and GraphQL APIs, NATS data plane, MITRE-tagged inference, signed reports, RBAC and outbound integrations.

FOUR PRIMARY USE CASES — THREE DOMAINS, ONE PLATFORM

#	Domain	Use case	Objective
01	OT	Proactive Detection & Response in OT	Entrap adversaries at the reconnaissance phase in any OT environment — before any manipulation of a real controller is attempted — using protocol-correct deception assets that mimic the exact devices and protocols already running there, working in tandem with any microsegmentation or industrial-firewall fabric already in place.
02	IT	Mirrored Workload	Give risk-averse organisations — banks, insurers, any brand wary of deploying deception directly into production — a way to see adversary behaviour aimed at their digital footprint without touching production, by hosting a faithful, independently-run mirror of their internet-facing estate at a tunable depth of engagement.
03	IT	Proactive Detection & Response Across the IT Data Center	Catch lateral movement, identity abuse and discovery-phase reconnaissance across the full IT data center — application servers, databases, Active Directory/identity infrastructure, and Windows/Linux workloads reachable via RDP and SSH — at the point adversaries spend most of their dwell time and where baseline-driven tooling is structurally weakest.
04	AI	AI Application & Agent Security	Detect and trap attackers — human or autonomous AI agents — that probe, jailbreak or exfiltrate from production AI: chatbots, copilots, agentic pipelines, MCP servers and RAG knowledge bases — an attack surface no WAF, EDR or DLP tool currently monitors.

OT · USE CASE 01

01 PROACTIVE DETECTION & RESPONSE IN OT

Use case

Entrap adversaries at the reconnaissance phase in any OT environment — before any manipulation of a real controller is attempted — using protocol-correct deception assets that mimic the exact devices and protocols already running there, working in tandem with any microsegmentation or industrial-firewall fabric already in place.

How AMaze™ addresses it

AMaze™ ships protocol-correct emulators for BACnet/IP, Modbus TCP, DNP3, S7/CIP (EtherNet/IP), IEC 60870-5-104 and OPC UA, and the **BYOT (Build Your Own Trap)** framework lets a customer stand up a high-fidelity decoy for any proprietary PLC, IED, RTU or sensor by declaring its protocol, port, device model and register/tag map in a single JSON file — live in minutes, no platform engineering required. Because no legitimate OT process ever communicates with a deception asset, every probe — a Modbus read-coil request, an S7Comm connection setup, a BACnet WhoIs broadcast — is adversarial by construction. AMaze™ operates as a complementary layer alongside any existing OT security architecture: on non-allowed paths behind a microsegmentation or industrial-firewall policy, deception assets absorb what would otherwise be a silent deny-log entry; on allowed paths, where a compromised operator workstation or abused VPN user moves through legitimate channels, AMaze™ detects by behaviour rather than route or identity. In environments with no existing OT security controls, AMaze™ provides a self-contained detection layer from day one.

Benefit

Security teams catch adversaries at the earliest point in the OT kill chain — reconnaissance — without touching production assets, disrupting live processes, or requiring changes to operational policy. For sites running microsegmentation or industrial firewalls, AMaze™ converts their blind spots into high-confidence intelligence; for sites with no existing OT controls, it is the detection layer. Every reconnaissance attempt, unauthorised protocol interaction and IT-to-OT lateral probe is MITRE ICS-ATT&CK-tagged and ready for triage, with zero risk to the process that generates it.

IT · USE CASE 02

02 MIRRORED WORKLOAD

Use case

Give risk-averse organisations — banks, insurers, any brand wary of deploying deception directly into production — a way to see adversary behaviour aimed at their digital footprint without touching production, by hosting a faithful, independently-run mirror of their internet-facing estate at a tunable depth of engagement.

How AMaze™ addresses it

AMaze™ builds and hosts a mirrored sub-domain composed from the sca-http, sca-ssh, sca-rdp, sca-mysql and sca-windows fleet, configured to reflect the institution's actual visible footprint — login portal, mobile API gateway, VPN concentrator — with banners, certificates and session behaviour tuned to be indistinguishable from the real thing. The customer chooses the depth from the outset: **low interaction** (banner/protocol-layer responses — HTTP headers, SSH banners, TLS fingerprints — for immediate, low-commitment visibility), **medium interaction** (realistic login flows and API responses that capture full credential lists and API-abuse patterns), or **high interaction** (sustained engagement that draws out full attack chains — initial access through post-authentication enumeration and exfiltration attempts). Because no legitimate user has any reason to reach the mirror, every interaction at any tier is adversarial by construction.

Benefit

Institutions gain continuous, organisation-specific adversary intelligence from a deployment that makes no changes to production and needs no change-control cycle to stand up — satisfying risk committees and regulatory obligations that a heavier deployment would stall on. Source IPs, credential lists and API-abuse patterns observed on the mirror feed directly into edge blocking, WAF deny-lists and the real login page's defences via the outbound integration layer. The mirror is also the first phase of a graduated rollout: when the customer is ready, its SCA configuration and accumulated intelligence transfer directly into a full deployment.

IT · USE CASE 03

03 PROACTIVE DETECTION & RESPONSE ACROSS THE IT DATA CENTER

Use case

Catch lateral movement, identity abuse and discovery-phase reconnaissance across the full IT data center — application servers, databases, Active Directory/identity infrastructure, and Windows/Linux workloads reachable via RDP and SSH — at the point adversaries spend most of their dwell time and where baseline-driven tooling is structurally weakest.

How AMaze™ addresses it

AMaze™ deploys Synthetic Cognitive Agents that mimic the assets adversaries are trying to reach, and plants Neural Echoes on real systems that draw adversaries toward them. Windows-server SCAs present fully authentic Active Directory, Kerberos, LDAP, SMB, WinRM and NTLM-challenge surfaces — synthetic domain accounts, Kerberoastable service principals and honey NTLM hashes — so that BloodHound, Rubeus, Impacket or CrackMapExec runs land on deception rather than a real domain controller. SSH, RDP, HTTP/S and MySQL/MSSQL/PostgreSQL SCAs cover server and application infrastructure at the protocol level, while planted connection strings, fake developer credentials and synthetic API keys create the breadcrumbs that draw a lateral-moving attacker toward them. No legitimate workload or user has any reason to touch a deception asset, so every connection or query is adversarial by construction — no baseline to tune, no signature to maintain.

Benefit

Detection happens during discovery and navigation, not after a high-value asset is reached. Every alert is pre-classified — MITRE-tagged, chain-stitched, fraud-scored — and dispatched in real time into the SIEM, SOAR and EDR tools already in place, with containment (BLOCK_IP, Echo rotation, SCA redeployment) gated by the existing approval workflow. The data center becomes proactively hostile: every layer an adversary needs to navigate contains assets designed to trap them.

AI · USE CASE 04

04 AI APPLICATION & AGENT SECURITY

Use case

Detect and trap attackers — human or autonomous AI agents — that probe, jailbreak or exfiltrate from production AI: chatbots, copilots, agentic pipelines, MCP servers and RAG knowledge bases — an attack surface no WAF, EDR or DLP tool currently monitors.

How AMaze™ addresses it

AMaze™ extends the deception fabric to the AI layer with AI-SCAs, packaged to match however the customer has built their AI: an **MCP honey-server** that advertises honey tools/resources to any agentic or MCP stack, an **LLM proxy-sidecar** in front of any chatbot or OpenAI-compatible endpoint, and an **embeddable SDK** for bespoke agents. All three share one detection core: hostile prompts are classified as a jailbreak (met with a deterministic refusal) or a credential/data-exfiltration attempt (met with fabricated, canary-tagged fake credentials) — so reuse of that canary anywhere is confirmed exfiltration attributed to the exact AI session. AI Neural Echoes — a honey MCP-tool reference, a canary-seeded RAG document — lure an attacker's AI toward the trap. The same engine already shields AMaze™'s own SOC assistant, **AMaze Analyst**, in production. Every trip is classified against **OWASP LLM Top-10** and **MITRE ATLAS** and joins the same unified attack chain as IT and OT events.

Benefit

Security teams gain visibility into an attack surface that today produces zero signal in any tool they already own, without requiring changes to how the AI was built. Detections are canary-based rather than inferred, giving unambiguous proof of exfiltration rather than a similarity score, and AI-layer compromise is investigated in the same dashboard and attack chain as everything else the attacker touched — not triaged in isolation.

ADDITIONAL FIELD-READY SCENARIOS

The four primary use cases above describe how AMaze™ engages adversaries across OT, IT and AI. The seventeen scenarios below describe specific, code-backed deployments that customers run today — from SOC-noise reduction to APT engagement to multi-site analytics tuning.

#	Scenario	Objective
1	Early threat detection	Spot adversaries during reconnaissance and break the kill chain before crown jewels are reached.
2	Detect malicious user behaviour	Catch insiders and credential-misuse attackers via interactions with planted lures.
3	Targeted malware capture	Capture and fingerprint payloads dropped on FIM-watched paths without exposing production hosts.
4	Lateral movement detection	Surface SMB/RDP/WinRM probing in internal segments through a Windows-server SCA fleet.
5	SOC alert reduction & triage acceleration	Reduce noise to interactions with deception assets only; ticket from event in one click.
6	OT, IoT & air-gapped coverage	Extend detection to BACnet, Modbus, DNP3, S7/CIP, IEC 104, OPC UA without touching production controllers.
7	Compliance & forensics	Produce signed reports + an immutable audit log mapped to access, change and incident-response controls.
8	Adaptive threat intelligence	Build a tenant-specific intelligence graph from real adversary engagements, MITRE-tagged in flight.
9	Remote office & legacy-system protection	Drop a single agent on a host in a remote site; SCAs and Echoes work without invasive endpoint tooling.
10	High-value asset & IP protection	Surround crown jewels with realistic lures so any adversary path stumbles through deception first.
11	Phishing & C2 detection	Catch reverse-shell payloads, beaconing destinations and credential exfil targets through SCA + FIM capture.
12	Adversary behavioural profiling	Stitch every lure activation to surrounding session events into a per-adversary timeline.
13	Adaptive incident containment	Approval-gated automation: BLOCK_IP, rotate Echoes, despawn compromised SCAs from one workflow.
14	AI-assisted engagement & SOAR integration	Live MITRE tagging plus webhook/Slack/Teams/Generic-HMAC outbound dispatch into existing SOAR runbooks.
15	Multi-stage APT engagement	Long-dwell attackers play through IT, identity and OT lures while every step is captured and chained.
16	BYOT — Build Your Own Trap (IT · OT · AI)	One extensibility model: OT JSON decoys (any PLC/IED/RTU/BMS/sensor), IT via sca-base, AI via the SDK.
17	Protect your AI applications	AI-SCAs trap attackers / rogue AI probing chatbots, agents, MCP & RAG; AMaze Analyst assistant; OWASP LLM Top-10 / MITRE ATLAS.

01 EARLY THREAT DETECTION

Use case

Identify adversaries during reconnaissance — port scans, banner grabs, credential harvesting — before they reach production assets.

How AMaze™ addresses it

AMaze™ deploys IT, identity and OT SCAs across the customer estate, plus a layer of Neural Echoes (seeded credential files, config secrets, VPN profiles) on real infrastructure. Legitimate users have no reason to interact with either; any probe of a sca-ssh, sca-rdp, sca-mysql or BACnet controller persona, or any read of a planted lure, is captured and tagged in real time with a MITRE technique ID and a fraud-scored source IP.

Benefit

Every interaction with an AMaze™ asset is by construction adversarial. The platform produces high-confidence early-warning signals before the adversary reaches a real asset, with a fraud-scored source IP and a MITRE-tagged technique ready for triage.

02 DETECT MALICIOUS USER BEHAVIOUR

Use case

Identify users — insiders, contractors or external attackers wielding stolen credentials — who exhibit reconnaissance behaviour from inside the network.

How AMaze™ addresses it

AMaze™ seeds high-value-looking lures across shared drives, internal repos, AD, VPN config backups and developer machines. The host agent's File Integrity Monitor watches each lure; access — read, write, copy, exec — emits an *amaze.events.echo_activation* message. The ingest pipeline joins the activation to neighbouring neural events on the same host (±10 minutes), so SOC sees both the lure touched *and* what else that user/process did during the same session.

Benefit

Compromised credentials and insiders both surface from the same signal. Alerts include the actor's host, the surrounding session events, the MITRE techniques observed and a stitched attack chain — no further investigation needed to confirm intent.

03 TARGETED MALWARE CAPTURE

Use case

Intercept advanced or previously unknown malware variants designed to bypass conventional EDR and antivirus tools.

How AMaze™ addresses it

FIM-watched paths on each host include the directories adversaries typically write to during exploitation — web roots, temp directories, startup folders, share mount points. When a write matches the malware policy, the agent base64-encodes the file (up to 10 MB by default), fingerprints it with SHA-256, moves the original to a quarantine directory and publishes *amaze.events.malware* with the captured payload, hash and full path.

Benefit

Forensic teams receive the live binary, its hash and the host context — including which process wrote it — without exposing production hosts to a third-party sandbox. New samples enrich the tenant-specific threat intelligence graph and feed signature updates.

04 LATERAL MOVEMENT DETECTION

Use case

Catch adversaries who have a foothold and are pivoting across internal hosts via SMB, RDP, WinRM or Active Directory.

How AMaze™ addresses it

The Windows-server SCA emulates SMB 2/3, RDP, WinRM and LDAP, and is deployed inside internal segments alongside production assets. Pass-the-hash relays, Kerberoasting probes, named-pipe lateral attempts and LDAP enumeration land in the SCA, which forwards events over NATS to the ingest pipeline. Source IPs that are RFC-1918 are scored locally without leaving the customer estate.

Benefit

Internal lateral movement surfaces with the same fidelity as external scans, but the adversary's internal source IP never leaves the customer network. Every event carries a MITRE technique and tactic and feeds the attack-chain view.

05 SOC ALERT REDUCTION & TRIAGE ACCELERATION

Use case

Improve SOC efficiency by minimising false positives and concentrating analyst attention on verified threats.

How AMaze™ addresses it

Because alerts are generated only by interactions with SCAs and Echoes — assets a legitimate user has no reason to touch — the false-positive rate is structurally low. The dashboard supports single-click conversion of an event into a ticket, with the source event reference, MITRE technique, fraud score and surrounding chain attached. Tickets can auto-dispatch to Webhook, Slack, Teams or Generic-HMAC integrations (Splunk, Wazuh, Elastic, ServiceNow, Jira, XSOAR, etc.).

Benefit

Analysts spend their time on high-confidence incidents, not on reconciling a noisy benign baseline. Every ticket arrives with the context (chain, scoring, IP reputation, GeoIP) already populated.

06 OT, IOT & AIR-GAPPED ENVIRONMENT COVERAGE

Use case

Detect adversaries inside operational technology (OT), Internet-of-Things and air-gapped networks where conventional security tooling cannot operate without risking downtime.

How AMaze™ addresses it

AMaze™ ships protocol-correct emulators for BACnet/IP (BMS controllers), Modbus TCP (gas analysers, power monitors, battery-management gateways, pumps, generic RTUs), DNP3 over TCP (SCADA Master / RTU / IED), Siemens S7Comm and EtherNet/IP CIP (CompactLogix 1769-L33ER · Compressor 5069-L320ERS2 · ControlLogix 1756-L85E), IEC 60870-5-104 (power-transmission SCADA, substation RTUs) and OPC UA (industrial gateways, sensor aggregators). The host agent is push-only — it never accepts inbound traffic — making it safe to run inside segmented and air-gapped OT networks. The **BYOT (Build Your Own Trap)** framework lets customers stand up a high-fidelity decoy for *any* OT device — proprietary PLCs, regional protocol variants, custom IEDs — by writing a single JSON file (protocol, port, model, register/tag map, MITRE TTP rules) and dropping it into *otsca/byot/configs/*. The decoy is live in minutes.

Benefit

Active defence reaches deep into OT without injecting traffic into real controllers. ICS-ATT&CK techniques (T0806 sensor manipulation, T0827 parameter falsification, T0838 alarm suppression, T0850 unauthorized analog output) are detected, scored and surfaced alongside enterprise events.

07 COMPLIANCE & FORENSICS

Use case

Support compliance and audit obligations such as PCI-DSS, ISO 27001, SOC 2, GDPR, HIPAA, NIST CSF, RBI/SEBI cyber-security guidelines, NERC-CIP, IEC 62443 and similar frameworks that demand verifiable evidence of detection and response controls.

How AMaze™ addresses it

Every gated mutation in the platform — user, site, integration, ticket, approval, echo, SCA, deployment, dispatch, report — writes a row to *audit_log*. CISO/board reports are generated on a weekly, monthly or quarterly schedule, signed with HMAC-SHA256 over canonical JSON || sha256(pdf_bytes), and verifiable post-hoc via *POST /api/v3/reports/runs/{id}/verify*. The report's compliance section pulls directly from the audit log, the integration dispatch history and the approval grants/denials.

Benefit

Auditors get verifiable evidence of detection, change-control and dispatch SLA. Reports are signed and integrity-checkable; the audit log is immutable and exposed read-only to the auditor role; the compliance section maps directly to common framework controls.

08 ADAPTIVE THREAT INTELLIGENCE

Use case

Generate actionable, organisation-specific threat intelligence from live adversary behaviour rather than relying solely on generic external feeds.

How AMaze™ addresses it

Each captured event is enriched with GeoIP (MaxMind), IP reputation (IPQS, AbuseIPDB, VirusTotal), a numeric threat score and a MITRE technique/tactic ID. Echo activations are correlated into attack chains. The dashboard's Adversary Spotlight and the report's Top-Adversaries section surface the IPs, ASNs, countries, tactics and depths most active against the tenant — a tenant-specific intelligence graph that updates with every interaction.

Benefit

Defensive posture follows the adversary the customer is actually facing, not a generic feed. Custom signatures can be derived from captured payloads; outbound dispatch carries that intelligence into the SOAR runbook.

09 REMOTE OFFICE & LEGACY SYSTEM PROTECTION

Use case

Extend detection to remote offices, branch sites and legacy systems that cannot run modern endpoint protection or where deployment of new tooling is constrained.

How AMaze™ addresses it

A single Rust agent runs on Linux (Debian/Ubuntu/RHEL/Alpine) or Windows Server 2019/2022 and pushes events outbound over NATS to the central control plane. Lightweight SCAs and Echoes are spawned on demand by the deployment planner. The agent never accepts inbound traffic, so no port needs to be opened on the legacy host.

Benefit

Detection reaches sites and hosts where conventional tooling cannot deploy. Cost per protected host is low, and sites can be onboarded incrementally without disrupting existing infrastructure.

10 HIGH-VALUE ASSET & INTELLECTUAL-PROPERTY PROTECTION

Use case

Prevent adversaries from reaching crown-jewel assets — source code, customer databases, financial platforms, executive endpoints, M&A documents.

How AMaze™ addresses it

Around each crown-jewel asset, the deployment planner places a layer of plausibly-named SCAs and Neural Echoes (mock database servers, planted README files, fake VPN configs in shared drives, expired-cert PKCS12 archives in config backups). Adversaries that follow their normal credential-harvesting and discovery tooling stumble through deception before they touch the real asset, generating an attack chain in the process.

Benefit

The protected asset gains an active buffer zone. Adversary intent and tooling are revealed before production data is touched, and the resulting chain feeds containment automation and forensic investigation.

11 PHISHING & COMMAND-AND-CONTROL DETECTION

Use case

Identify phishing-delivered payloads and the command-and-control infrastructure they reach back to.

How AMaze™ addresses it

The host agent's process monitor records process trees with command-line arguments. The FIM captures dropper binaries on disk; the packet sniffer records the outbound connections those binaries make. The IP reputation enrichment scores the destination, the attack-chain correlator binds the dropper, the process and the destination together, and the dispatcher pushes a high-confidence ticket with all of it attached to the SOC's chosen integration.

Benefit

Phishing-driven intrusions are caught at the dropper-execution stage, not days later via lateral movement detection. Destination IPs become an actionable C2 indicator the SOC can block at the edge.

12 ADVERSARY BEHAVIOURAL PROFILING

Use case

Build deep, per-adversary behavioural profiles from live engagements with the deception fabric.

How AMaze™ addresses it

Every captured event records command sequences, payloads, credential attempts, fraud-scored source IP, ASN, country, MITRE technique and tactic. Echo activations stitch these into *attack_chains* per adversary IP/session. The Attack Explorer visualises the chain as a timeline; the report's Top-Adversaries section narrates it for executive consumption.

Benefit

Threat hunters get a per-adversary tactic library specific to the tenant. Detections and containment automations can be tuned to actual adversary behaviour rather than generic patterns.

13 ADAPTIVE INCIDENT CONTAINMENT

Use case

Contain active adversaries without manual reconfiguration of network controls or wholesale host isolation.

How AMaze™ addresses it

Containment actions (BLOCK_IP, rotate Echoes, despawn or redeploy SCAs) are first-class GraphQL mutations and deployment-planner state transitions. Each action can be wrapped in the approval workflow with a per-integration timeout (default 30 min), so a chosen automation can fire automatically once an admin or auditor has approved. Approvals and outcomes are audit-logged.

Benefit

Containment is faster than an SOC analyst typing into a firewall console, but every action is approval-gated and audit-logged. The blast radius of any single decision is bounded by the role and site scope of the approver.

14 AI-ASSISTED THREAT ENGAGEMENT & SOAR INTEGRATION

Use case

Automate coordinated security response across SIEM, SOAR and chat tools driven by live deception telemetry.

How AMaze™ addresses it

Outbound dispatch supports Webhook, Slack, Teams and Generic-HMAC integration types with auth modes *none / bearer / HMAC-SHA256*. A single ticket or report can fan out to Splunk HEC, Wazuh, Elastic, Cortex XSOAR, ServiceNow, Jira, PagerDuty, Shuffle and a SOC channel simultaneously. Secrets are Fernet-encrypted at rest; the dispatcher loop retries with exponential backoff and persists every attempt in *integration_dispatches*.

Benefit

AMaze™ slots into existing SOC stacks rather than replacing them. Runbooks fire on high-fidelity deception signals, with a verifiable delivery record for compliance and SLA reporting.

15 MULTI-STAGE APT ENGAGEMENT

Use case

Detect, engage and analyse long-dwell Advanced Persistent Threats that traverse reconnaissance, exploitation, lateral movement, persistence and exfiltration over weeks or months.

How AMaze™ addresses it

AMaze™ deploys IT, identity and OT SCAs together with persistent Neural Echoes. The TimescaleDB *neural_events* hypertable is time-partitioned, so multi-week correlation queries are fast. Attack chains accumulate over time, and the executive dashboard's adversary spotlight surfaces the longest-dwelling actors.

Benefit

Security teams gain visibility into the full APT lifecycle inside their estate, with a tenant-specific kill-chain that informs threat models, defensive gaps and counter-measures.

16 BYOT & CUSTOM OT PERSONAS — EMULATE *ANY* DEVICE

Use case

Tailor deception strategy to a proprietary device fleet, a regulated environment or a regional protocol variant — without waiting for a vendor engineering ticket.

How AMaze™ addresses it

BYOT (Build Your Own Trap) loads custom OT/ICS personas at runtime from *otsca/byot/configs/* (declarative JSON: protocol, port, device model, register/tag map, MITRE ICS-ATT&CK TTP rules) and *otsca/byot/handlers/* (one handler per protocol — Modbus, EtherNet/IP CIP, S7Comm, DNP3, BACnet/IP, IEC 60870-5-104, OPC UA). Bundled examples cover substation DNP3 IEDs, IEC 104 RTUs, Modbus battery monitors, S7 auxiliary controls and OPC UA sensor gateways. The same scoring, attack-chain correlation and outbound-dispatch pipeline applies on day one. Analytics dashboards (external, internal, OT, Active Directory, Neural Echoes) support filtering by SCA type, subnet, protocol, country, time window and site. The deployment planner persists a JSONB plan that can be cloned, edited and rolled back on demand.

Benefit

Deception strategy adapts to the customer's risk profile and regulatory context without platform-side engineering. Per-segment analytics and SCA-level filtering let SOC, OT and compliance teams each see what they need.

17 PROTECT YOUR AI APPLICATIONS

Use case

Detect and trap attackers — human or autonomous AI agents — that probe, manipulate or exfiltrate from your production AI applications: chatbots, copilots, agentic pipelines, MCP servers and RAG knowledge bases.

How AMaze™ addresses it

AMaze™ AI-SCAs extend the deception fabric to the AI layer. Three deployable forms drop into any architecture without code changes: an **MCP honey-server** that advertises honey tools to agentic and MCP stacks, an **LLM proxy-sidecar** in front of any chatbot or OpenAI-compatible endpoint, and an **embeddable SDK** for bespoke agents. Each form scans prompts for injection or exfil patterns, tarpits hostile agents, and serves tracked honey credentials on credential-theft attempts so that reuse anywhere is confirmed exfiltration. **AI Neural Echoes** — a honey MCP-tool reference and a canary-seeded RAG document — lure an attacker's AI toward the trap. AMaze™'s own SOC assistant, **AMaze Analyst**, is protected by the same engine: overt jailbreaks are refused; on credential-theft attempts the assistant hands the attacker tracked fake credentials, gated so legitimate analysts never trip it. Every trip is classified against **OWASP LLM Top-10** and **MITRE ATLAS** and joins the same unified attack chain, dashboard and SOAR dispatch as IT and OT events.

Benefit

Security teams gain visibility into a brand-new attack surface no traditional control watches — without modifying the AI application itself. Every honey credential is a canary, so reuse anywhere is confirmed exfiltration attributed to the exact AI session. Defenders catch prompt injection, jailbreaks and data-exfiltration attempts before a real model is subverted — and AI threats surface in the same console, attack chain and SOAR runbooks as the rest of the estate.

CONTACT MIRRORMIRE

www.mirrormire.ai | info@mirrormire.com | +1 (650) 504-3863