AMaze USP

Why MirrorMire AMaze™ wins where conventional cyber-resilience tools stall

EXECUTIVE SUMMARY

MirrorMire AMaze™ is an AI-driven proactive cyber-resilience platform built on Synthetic Cognitive Agents (SCAs) — instrumented decoys indistinguishable from real production services — and Neural Echoes — persistent breadcrumb artefacts seeded across real infrastructure. Unlike reactive security tools that wait for damage, AMaze™ engages adversaries at the earliest phase of an attack, captures their tools, techniques and procedures, and produces structurally-low-false-positive alerts. AMaze™ now extends that deception fabric to a third surface — **AI** — with AI-SCAs that trap attackers and rogue AI agents probing your AI applications, and **AMaze Analyst**, the in-product SOC assistant. The 21 differentiators below describe the competitive advantages AMaze™ delivers across detection, intelligence, deployment and operations — every one of them grounded in capabilities that ship in the platform today.

USP SUMMARY

#	USP	Core benefit
01	Behavioural-adaptive engine	Real-time MITRE / tactic / threat-score on every event; GraphQL C2 mutations to rotate or block on demand
02	Multi-zone fabric (IT · OT · identity · air-gap)	One platform, push-only edge agent; protocol-correct OT/ICS fleet covering BACnet · Modbus · DNP3 · S7Comm · CIP · IEC 104 · OPC UA
03	Live malware capture & fingerprinting	FIM-driven binary capture, SHA-256 hash, quarantine, *amaze.events.malware* publication
04	Attack-chain correlation	Echo activation joined to all neural events on same host within ±10 min; idempotent persisted chains
05	Persistent breadcrumb lures (Neural Echoes)	SSH keys, configs, VPN profiles, AD tokens, certs & backups planted on real assets — not in a deception VLAN
06	MITRE ATT&CK on every event	Per-rule technique + tactic ID, surfaced as coverage heatmap and gap analysis in the executive dashboard
07	Multi-provider IP reputation	IPQS · AbuseIPDB · VirusTotal in *byok* / *mirrormire* / *self_hosted* modes; RFC-1918 never sent upstream
08	Low-noise alerting by construction	SCA + Echo events only — there is no benign baseline, so any signal is a real signal
09	Audit-grade signed reports	Weekly / monthly / quarterly CISO + board reports; HMAC-SHA256 signed; verify endpoint for post-hoc integrity
10	RBAC + per-user site scope	Three roles, JWT *sites* claim, *scoped_sites_or_403* on every endpoint, full audit log
11	SOAR / SIEM outbound dispatch	Webhook · Slack · Teams · Generic-HMAC; Splunk · Wazuh · Elastic · XSOAR · ServiceNow · Jira · Shuffle
12	Reversible deployment planner	JSONB plan, DRAFT → ACTIVE → ROLLBACK state machine, live progress dashboard
13	Build-Your-Own-Trap (BYOT) — emulate *any* OT device	JSON-defined high-fidelity decoys for any PLC, IED, RTU, BMS, battery monitor or sensor gateway — proprietary or off-the-shelf
14	Sub-100 ms streaming analyst console	Redis pub/sub → WebSocket :8766 → React dashboard for live forensic walk-through
15	Cryptographic licensing & supply-chain hygiene	RS256 license JWT, periodic beacon, env-var-only secrets, fail-closed boot
16	Self-monitoring observability	Engine page surfaces NATS lag, queue depths, broadcast sessions, scheduled-report next-fire times
17	OT-aware detection without disrupting operations	Substation- & plant-floor-grade ICS decoys; zero injected traffic into real controllers; full MITRE ICS-ATT&CK technique mapping per event
18	Multi-segment analytics dashboard	External · Internal · OT · Identity views, filterable by SCA, subnet, protocol, country, time, site

#	USP	Core benefit
19	Echo deployment map & coverage view	Every seeded artefact + activation count + rotation schedule + agents-gone-silent gap report
20	AI-SCA — cognitive deception for the AI layer	Traps human + AI attackers probing your AI (chatbots, agents, MCP, RAG); in-product engine + MCP honey-server / LLM proxy / SDK deployables; OWASP LLM Top-10 / MITRE ATLAS
21	AMaze Analyst — your AI teammate in the SOC	In-product assistant: RAG + safe read-only text-to-SQL over your telemetry; reconstructs attack campaigns; vendor-neutral / BYOK

DETAILED USP DESCRIPTIONS

01 BEHAVIOURAL-ADAPTIVE CYBER-RESILIENCE ENGINE

AMaze™ does not simply alert on what an adversary touches. Every neural event is tagged in real time with a numeric threat score, a MITRE ATT&CK technique ID and a tactic ID by an inference engine that runs over 50 protocol-specific signature rules. Operators can rotate Echoes, spawn additional SCAs or block source IPs through the GraphQL C2 layer — manually, or wired to an approval-gated automation — in response to live behaviour.

— Risk score, technique ID and tactic ID stamped on every event
— GraphQL mutations: *spawn · despawn · plantEcho · rotateEcho · blockIp*
— Approval workflow with timeout-based auto-rejection (default 30 min)

02 BUILT FOR IT, INTERNAL, IDENTITY, OT & AIR-GAPPED AT ONCE

AMaze™ is purpose-built to operate across traditional IT, internal lateral-movement zones, Active Directory and identity surfaces, industrial OT/SCADA networks, and disconnected environments. The host agent is push-only over NATS — it never accepts inbound traffic — making it safe to drop into segmented and air-gapped networks where conventional security tooling cannot operate. OT support is not a sticker on the box: AMaze™ ships protocol-correct emulators across the seven industrial protocols that matter most, plus a BYOT framework that turns the platform into a decoy fabric for *any* OT device a customer asks for.

— OT/ICS protocols: BACnet/IP · Modbus TCP · DNP3 · Siemens S7Comm · EtherNet/IP CIP · IEC 60870-5-104 · OPC UA
— Out-of-the-box devices: BMS controllers, gas analysers, power monitors, battery monitors, pumps, generic RTUs, CompactLogix & ControlLogix PLCs, substation IEDs, OPC UA sensor gateways
— Identity: Windows-server SCA (SMB · WinRM · LDAP) plus AD lure surfaces
— BYOT (Build Your Own Trap): JSON-defined personas for proprietary PLCs, IEDs and any unsupported protocol — without rebuilding the platform
— Push-only edge agent: zero inbound exposure on the protected host; safe in air-gapped & segmented OT

03 LIVE MALWARE CAPTURE & FORENSIC-GRADE FINGERPRINTING

When the host agent's File Integrity Monitor sees a write that matches the malware policy, it captures the binary contents (base64, up to 10 MB by default), fingerprints it with SHA-256, moves the original to a quarantine directory and publishes a *amaze.events.malware* message. Forensic teams get the original sample, the SHA-256 hash, the path it was written to and the process tree that wrote it — without exposing production hosts to a third-party sandbox.

04 ATTACK CHAIN CORRELATION: LURE → SESSION → TACTICS, AUTOMATICALLY

When a Neural Echo activates, the ingest pipeline joins that activation against every neural event from the same host within a ±10-minute window and persists the result as an *attack_chains* row. The Attack Explorer in the dashboard renders the chain as a timeline: which credential was touched, which SCAs were probed, what techniques were observed, where the source IP geolocated, and what fraud score the IP earned. The correlation is idempotent on (echo_id, first_seen_at), so re-delivery never corrupts the chain.

05 PERSISTENT BREADCRUMB LURES ON REAL INFRASTRUCTURE

Most deception products place lures in a quarantined deception VLAN that adversaries quickly learn to avoid. Neural Echoes live on real infrastructure: the credential file in *~/.ssh*, the fake VPN config on a shared drive, the planted README in a repo, the expired-cert PKCS12 in a config backup. Legitimate users should never touch them, so any access is by definition adversarial — and the lure is in the path the adversary already wants to walk.

06 MITRE ATT&CK COVERAGE ON EVERY EVENT, NOT JUST PER INCIDENT

AMaze™ tags every captured event — every SSH login attempt, every FTP retrieval, every HTTP request, every Modbus write, every echo activation — with a MITRE technique and tactic ID. The executive dashboard turns this into a coverage heatmap (tactics observed vs. tactics covered by deployed SCAs) and a TTP gap analysis that feeds the next deployment-planner cycle.

07 MULTI-PROVIDER IP REPUTATION WITH FIRST-CLASS PRIVACY

IP reputation is pluggable: customers run in *byok* mode with their own IPQS, AbuseIPDB and VirusTotal keys; in *mirrormire* mode using the shared MirrorMire enrichment endpoint; or in *self_hosted* mode against a customer-run reputation service for full data sovereignty. RFC-1918, loopback and link-local addresses are *never* sent upstream, so internal lateral movement IPs stay inside the customer estate.

08 LOW-NOISE ALERTING BY CONSTRUCTION

Because every alert originates from an interaction with a Synthetic Cognitive Agent or a Neural Echo — assets a legitimate user has no business touching — the false-positive rate is structurally low. Operators are not asked to reconcile a noisy benign baseline; the baseline is zero, and any signal is a signal.

09 AUDIT-GRADE REPORTING THAT THE BOARD CAN SIGN

AMaze™ generates weekly, monthly and quarterly CISO/board reports on an in-process scheduler. Reports are rendered through a headless browser (so the PDF is pixel-identical to the dashboard), then signed with HMAC-SHA256 over canonical JSON || sha256(pdf_bytes). A *POST /api/v3/reports/runs/{id}/verify* endpoint recomputes the HMAC and returns ok/mismatch, so the integrity of every archived report can be checked at any later date. Signing keys are rotatable via *signing_key_id*.

10 RBAC + PER-USER SITE SCOPE WITH A COMPLETE AUDIT TRAIL

Three first-class roles (admin, analyst, auditor) and a per-user site scope encoded directly in the JWT. Endpoints enforce the scope through a *scoped_sites_or_403* dependency, so cross-site access returns 403 without ever touching application logic. Every gated mutation — user, site, integration, ticket, approval, echo, SCA, deployment, dispatch, report — writes a row to *audit_log*, exposed read-only to admins and auditors. The audit log is the compliance system of record and feeds the report's compliance section directly.

11 OUTBOUND INTEGRATIONS THAT FIT EXISTING SOC STACKS

Outbound dispatch supports Webhook, Slack, Teams and Generic-HMAC types, with auth modes *none / bearer / HMAC-SHA256* and customisable headers. That covers Splunk HEC, Wazuh, Elastic, ServiceNow, Jira, Cortex XSOAR, PagerDuty, Shuffle and any other REST-based platform. Secrets are Fernet-encrypted at rest in PostgreSQL and never logged. The dispatcher loop retries with exponential backoff and records every attempt in *integration_dispatches*.

12 DEPLOYMENT PLANNER WITH REVERSIBLE STATE MACHINE

Deployments are first-class objects: a JSONB plan moves through a DRAFT → ACTIVE → ROLLBACK state machine. Activating a plan publishes SPAWN_DECOY commands over *amaze.cortex.<vm_id>* to the targeted agents; rollback issues DESPAWN_DECOY for the same set. The dashboard surfaces live progress, so operators always know exactly what is in production and can revert without ad-hoc scripts.

13 BUILD-YOUR-OWN-TRAP (BYOT) — EMULATE *ANY* OT DEVICE

BYOT is AMaze™'s answer to a problem every other deception platform ducks: the customer's estate contains a proprietary PLC, a non-standard IED, an unusual battery-management gateway or a regional protocol variant that no vendor supports out of the box. With BYOT, the customer (or the MirrorMire SE) writes a single JSON file declaring the protocol, port, model, register and tag map, and a list of MITRE ICS-ATT&CK TTP rules that fire when an adversary touches a sensitive value. The decoy is live in minutes — fully integrated with the platform's scoring, correlation, attack-chain and outbound-dispatch pipeline.

— Protocol handlers ship for Modbus TCP, EtherNet/IP CIP, S7Comm, DNP3, BACnet/IP, IEC 60870-5-104 and OPC UA
— Bundled examples: substation DNP3 IED · IEC 104 RTU · Modbus battery monitor · S7 auxiliary control · OPC UA sensor gateway
— TTP mapping is per-register: e.g. T0831 (manipulation of control), T0855 (unauthorised command), T0827 (parameter modification)
— Additive runtime: BYOT runs in its own process tree; existing SCAs are never disrupted
— Same JSON declaratively documents the device — customers ship the config to compliance and audit teams as part of the SCA inventory
— One extensibility model, three surfaces: the same Build-Your-Own-Trap philosophy extends to IT (inherit the *sca-base* image) and AI (embed the *aisca-sdk*)

14 REAL-TIME ANALYST EXPERIENCE: SUB-100 MS STREAMING

Enriched events are fanned out from Redis pub/sub onto a WebSocket stream (*ws://<host>:8766/ws*) consumed directly by the React dashboard. Analysts watch SSH brute force unfold, Modbus writes hit a controller persona, or an echo activate — live, with the source IP, country, ASN, fraud score and MITRE technique already populated.

15 CRYPTOGRAPHICALLY SIGNED LICENSING & SUPPLY-CHAIN HYGIENE

License JWTs are RS256-signed by the back-office *amaze-ops* service and verified by *amaze-core* on boot. Invalid or expired keys cause a hard fail-closed exit; a periodic beacon checks back into *amaze-ops* for usage telemetry and revocation. Every secret in the platform — JWT, AES wire key, integration tokens, report-signing key — is sourced from environment variables; nothing is hardcoded in the source tree.

16 GROUND-TRUTH OBSERVABILITY OF THE PLATFORM ITSELF

The Engine dashboard exposes the operational signal AMaze™ needs to run itself: NATS consumer lag, enrichment queue depth, dispatcher queue depth, broadcast session count, scheduled-report next-fire times. Every service exposes *GET /health* and *GET /ready*, and structured JSON logging via structlog drops cleanly into any log pipeline.

17 OT-AWARE THREAT DETECTION WITHOUT DISRUPTING OPERATIONS

AMaze™ is OT-native, not OT-retrofitted. Every OT SCA implements a protocol-correct emulator — BACnet/IP, Modbus TCP, DNP3, Siemens S7Comm, EtherNet/IP CIP, IEC 60870-5-104 and OPC UA — so adversary commands hit authentic device responses, not generic banner-grab placeholders. The platform's bundled OT fleet covers building-management controllers (BACnet), gas analysers and power monitors (Modbus), substation RTUs and IEDs (DNP3, IEC 104), Allen-Bradley / Rockwell ControlLogix-family PLCs (CompactLogix · Compressor · ControlLogix HP via S7/CIP), and OPC UA sensor gateways. Anything not on that list is a BYOT JSON file away.

— Zero injected traffic into real controllers — lures listen, never write into production
— Detection mapped to MITRE ICS-ATT&CK: T0806 sensor manipulation · T0827 parameter falsification · T0831 manipulation of control · T0838 alarm suppression · T0850 unauthorised analog output · T0855 unauthorised command
— Push-only edge agent — safe inside segmented OT zones; no port opened on the protected host
— The same scoring, attack-chain correlation and dispatcher pipeline that powers the IT fleet applies on day one in OT

18 GRANULAR, MULTI-SEGMENT ANALYTICS DASHBOARD

The dashboard splits external, internal, OT and Active-Directory views, each with its own overview, live event stream, analytics charts, and (for external) a Leaflet world map. Filters by SCA type, subnet, protocol, country and time window are first-class; analysts can drill from a heatmap cell into the specific neural events behind it in two clicks.

19 ECHO DEPLOYMENT MAP & COVERAGE VIEW

The Neural Echoes deployment map shows every seeded artefact, its planting host, its activation count, its rotation schedule and its lifecycle state. Combined with the SCA inventory heartbeat on *amaze.sca.inventory*, operators always know what is deployed where, and the report *gaps* section flags Echoes that have never activated and agents that have gone silent.

20 AI-SCA — COGNITIVE DECEPTION FOR THE AI ERA

As enterprises rush AI into production, the attack surface shifts from ports to prompts. AMaze™ is the first deception platform to fight AI attacks *with* AI deception. The same instrument-the-attacker's-path model that traps a port scan now traps a *thought*: an AI-SCA baits any attacker — human or autonomous AI agent — that tries to jailbreak, exfiltrate from, or enumerate an AI surface.

— **In-product engine** shields AMaze™'s own SOC assistant: it refuses overt jailbreaks, serves tracked honey credentials on credential-theft attempts, and tarpits hostile agents — and it never fires for legitimate analysts (gated on confirmed-hostile sessions only).
— **Three deployable forms to protect the customer's own AI** — an MCP honey-server for agentic stacks, an LLM proxy-sidecar in front of any chatbot/endpoint, and an embeddable SDK — so coverage fits whatever AI architecture the customer runs.
— Every planted secret is a **canary**: if it ever surfaces in a SIEM, an auth log or a paste site, that is confirmed exfiltration with attribution 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 AI-SCA — the AI analogue of planting honey credentials.
— Every trip is classified to **OWASP LLM Top-10** and **MITRE ATLAS**, lands in the dedicated AI-Threats dashboard, and joins the same unified attack chain as the IT and OT decoys.

21 AMAZE ANALYST — YOUR AI TEAMMATE IN THE SOC

AMaze Analyst is the platform's in-product analyst assistant. It answers investigative questions in natural language, grounded in your own deception telemetry via retrieval-augmented generation and safe, read-only text-to-SQL — it reconstructs an attacker's campaign on request, and it never invents data or runs anything but allow-listed SELECTs under row-level security.

— Vendor-neutral / BYOK: runs against OpenAI, Claude, Azure OpenAI, vLLM, Ollama or any OpenAI-compatible endpoint
— Read-only, RLS- and site-scoped SQL; rate-limited; every turn audited
— Protected by the in-product AI-SCA engine (USP 20) — the assistant is itself a trap for anyone who tries to subvert it

CONTACT MIRRORMIRE

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