AMaze Data Sheet

MirrorMire AMaze™ — Proactive CyberResilience as a Service (PCRaaS)

CORE FUNCTIONAL INSIGHTS

MirrorMire AMaze™ is a microservice platform that converts traditional threat detection into a proactive, adaptive cyber-resilience programme. At its core is a deployment engine that scatters high-fidelity Synthetic Cognitive Agents (SCAs) and persistent breadcrumb lures (Neural Echoes) across IT, internal, identity, OT and SCADA surfaces. SCAs emulate real services — SSH, RDP, MySQL, Windows file shares, BACnet building controllers, Modbus gas analysers, EtherNet/IP PLCs — and engage adversaries inside controlled containers, capturing every command, payload and credential they touch.

Captured events are encrypted with AES-256-CBC at the host agent before they are published to a NATS JetStream bus. The ingest pipeline decrypts, parses, GeoIP-resolves, IP-reputation-scores (IPQS · AbuseIPDB · VirusTotal) and MITRE ATT&CK-tags each event in flight. When a Neural Echo activates, the pipeline correlates that activation with all neural events on the same host in a ±10-minute window, producing a stitched attack chain — adversary IP, tactics, technique IDs, interaction depth — ready for triage.

Alerts and reports are dispatched outbound to Webhook · Slack · Teams · Generic-HMAC integrations, feeding existing SIEM/SOAR stacks. AMaze™ ships an in-platform analyst console that streams neural events live over WebSocket, plus an executive dashboard that produces signed CISO/board reports on a weekly, monthly or quarterly cadence.

AMaze™ is delivered as a self-hosted Docker Compose or Kubernetes deployment. The host agent is push-only over NATS — it never accepts inbound traffic — making the platform safe to deploy in hybrid, air-gapped, OT and identity-segmented networks.

AMaze™ extends the same deception model to a third surface — **AI**. AI-SCAs trap attackers (human or autonomous AI agents) that probe an organisation's AI applications — chatbots, agentic/MCP stacks, RAG pipelines. An in-product engine shields AMaze™'s own SOC assistant, **AMaze Analyst**, while deployable forms — an MCP honey-server, an LLM proxy-sidecar and an embeddable SDK — protect the customer's own AI. Every trip is classified to OWASP LLM Top-10 and MITRE ATLAS and joins the same unified attack chain as the IT and OT decoys.

KEY CAPABILITIES OF AMAZE™

— AI-driven proactive cyber-resilience engine
— High-interaction Synthetic Cognitive Agents (IT · OT · identity · AI)
— AI-SCA — cognitive deception for the AI layer (chatbots · agents · MCP · RAG)
— AMaze Analyst — in-product SOC assistant (RAG + safe text-to-SQL)
— Neural Echoes — persistent breadcrumb lures on real assets
— Real-time WebSocket event streaming (sub-100 ms)
— MITRE ATT&CK technique tagging on every event
— Attack chain correlation (lure activation ↔ session events)
— Multi-provider IP reputation (IPQS · AbuseIPDB · VirusTotal)
— Behavioural threat scoring & risk-adaptive engagement
— Live malware capture + SHA-256 + quarantine
— Outbound dispatch: Webhook · Slack · Teams · Generic-HMAC
— GraphQL C2 for dynamic SCA orchestration
— DRAFT → ACTIVE → ROLLBACK deployment planner
— 3-role RBAC + per-user site scope, full audit log
— Approval workflow with timeout-based auto-rejection
— Scheduled, HMAC-signed CISO & board reports
— BYOT — Build Your Own Trap: emulate any OT device from JSON
— Agentless on the protected host; push-only edge agent
— RSA-signed licensing with periodic beacon

TECHNOLOGICAL ADVANTAGES

Container-Based Synthetic Cognitive Agents

Every IT SCA is a Docker container inheriting *mirrormire/sca-base:latest*. The host agent (written in Rust on Tokio) drives spawn, despawn and revive lifecycles via the bollard Docker client, allowing rapid scaling, isolated execution and rolling updates.

Behaviour-Adaptive Engagement

Adversary interactions are tagged in real time with MITRE technique IDs by the inference engine and consumed by the deployment planner. Operators can rotate Echoes, spawn additional SCAs or issue a BLOCK_IP through the GraphQL C2 mutation — manually, or wired to an approval-gated automation.

Live Payload & Malware Capture

The host agent's File Integrity Monitor captures the binary contents of suspect files (inotify on Linux, ReadDirectoryChangesW on Windows), base64-encodes them up to a configurable size cap (default 10 MB), fingerprints them with SHA-256, quarantines the original, and publishes a *amaze.events.malware* message for forensic analysis.

Multi-Site, Multi-Tenant Topology

NATS subjects are scoped via *amaze.sites.<site_id>.events.<protocol>* so a single control plane can serve multiple customer sites. JWT *sites* claims plus a *scoped_sites_or_403* dependency in every endpoint enforce that analysts only see what they are scoped to.

Real-Time SCA Orchestration

Dynamic alterations to ports, banners and credentials are issued over the agent's *amaze.cortex.<vm_id>* NATS subscription. Echoes can be re-planted or rotated on demand while an attacker is still engaged, prolonging the deception window without disrupting other active sessions.

Quantified Threat Scoring

Each event carries a numeric *threat_score* (0–100), a *mitre_technique_id* and a *mitre_tactic_id*. The executive dashboard aggregates these into a portfolio risk score, MITRE coverage heatmap and adversary spotlight.

Customised OT & SCADA Emulation

Industrial protocols are first-class: BACnet/IP for building management, Modbus TCP (gas analyser and power monitor personas), DNP3 over TCP, Siemens S7 / EtherNet/IP CIP (CompactLogix · Compressor 5380 · ControlLogix HP), IEC 60870-5-104 and OPC UA. Each emulator detects ICS-ATT&CK techniques (T0806 sensor manipulation, T0827 parameter falsification, T0838 alarm suppression, T0850 unauthorised analog output, etc.).

Build-Your-Own-Trap (BYOT) — emulate any OT device

BYOT (Build Your Own Trap) is a config-driven framework that lets customers stand up a high-fidelity decoy for *any* OT/ICS device — without modifying core platform code. A single JSON file declares the protocol, port, model, register/tag map and a TTP-rule list (MITRE ICS-ATT&CK technique IDs) that fire when an adversary touches a sensitive value. Engine handlers ship for Modbus TCP, EtherNet/IP CIP, Siemens S7Comm, DNP3, BACnet/IP, IEC 60870-5-104 and OPC UA — and the bundled examples cover substation IEDs and RTUs, battery monitors, S7 auxiliary controls and OPC UA sensor gateways. The framework is additive (its own process tree) so existing SCAs are never disrupted.

— *Mimic any OT device:* generic RTU · battery monitor · auxiliary control · IED · sensor gateway · proprietary PLC
— *Map registers/tags to TTPs:* e.g. T0831 (manipulation of control), T0855 (unauthorized command), T0827 (parameter falsification)
— *Deploy in minutes:* drop a JSON into *otsca/byot/configs/* and run the BYOT Docker image with the file mounted in

Cognitive Deception for the AI Layer (AI-SCA)

AMaze™ is the first deception platform to fight AI attacks *with* AI deception. Where a network SCA traps a port scan, an AI-SCA traps a *thought*: it baits any attacker — human or AI agent — probing RAG context, the tool/function schema, or MCP tools. The in-product engine refuses overt jailbreaks, serves tracked honey credentials on credential-theft attempts, and tarpits hostile agents — gated on confirmed-hostile sessions so legitimate analysts never trip it. Three deployable forms protect the customer's own AI: an **MCP honey-server** (agentic stacks), an **LLM proxy-sidecar** (any OpenAI-compatible endpoint) and an **embeddable SDK**. Every honey credential is a canary — its reuse anywhere is confirmed exfiltration with attribution to the exact AI session.

— *One extensibility model, three surfaces:* Build-Your-Own-Trap covers OT (no-code JSON), IT (inherit *sca-base*) and AI (embed *aisca-sdk*)
— *AI Neural Echoes* (honey MCP-tool reference, canary-seeded RAG doc) lure an attacker's AI toward the AI-SCA
— *Mapped to OWASP LLM Top-10 & MITRE ATLAS;* surfaced in the AI-Threats dashboard and the unified attack chain

TECHNOLOGICAL SPECIFICATIONS

SCA Catalogue — IT & Application

Persona	Protocol(s) / Surface
sca-ssh	SSH v2 (login attempts, command logging, key-exchange capture)
sca-ftp	FTP / FTP passive (credential brute force, directory traversal, SITE EXEC)
sca-rdp	RDP (auth tracking, credential stuffing, NLA bypass)
sca-mysql	MySQL 5.7/8.0 (SQLi patterns, privilege escalation, LOAD DATA INFILE)
sca-redis	Redis RESP (unauth probing, FLUSHALL, AOF injection, ACL bypass)
sca-http	HTTP/HTTPS (path traversal, SQLi/XXE/SSRF in GET/POST, header injection)
sca-iot	Custom IoT (firmware enumeration, telemetry spoof, config exfil)
sca-windows	SMB 2/3 · RDP · WinRM · LDAP (PtH, Kerberoasting, lateral over named pipes)
sca-xampp	Apache + PHP + FTP + MySQL combo (PHP RCE, default creds)
sca-dionamaze	Custom telemetry collector with intentional vulnerabilities

SCA Catalogue — OT / ICS / SCADA

AMaze™ ships protocol-correct emulators for the protocols that matter in modern industrial estates — and the BYOT framework lets you stand up a high-fidelity decoy for *any* OT device, proprietary or off-the-shelf, with a single JSON file. Substation, plant-floor, BMS, water and energy customers get a fleet that mirrors their real assets on day one.

Persona / Protocol	Emulated device · TTPs detected
BACnet/IP (47808/UDP)	Building-management controller — Who-Is/I-Am enumeration, ReadProperty discovery, unauthorised HVAC commands, parameter falsification, alarm suppression
Modbus TCP (502)	Gas analyser · power monitor · battery management · pump · generic RTU — FC3 read enumeration, holding-register write surge, status falsification, range violation, fault injection
DNP3 over TCP (20000)	SCADA Master / RTU / IED — unauthorised analog output, malicious control messages, polled-data manipulation
Siemens S7Comm / EtherNet-IP CIP (44818)	Allen-Bradley / Rockwell ControlLogix family — DIGESTER (CompactLogix 1769-L33ER) · COMPRESSOR (5069-L320ERS2) · CONTROLLOGIX_HP (1756-L85E) — anomalous tag read/write, recovery suppression
IEC 60870-5-104 (2404)	Power-transmission SCADA, substation RTU — M_SP_NA_1 / M_DO_NA_1 command-injection probes, spontaneous-event abuse
OPC UA (4840)	Industrial gateway, sensor aggregator — method-invocation abuse, authorisation bypass, certificate-downgrade attempts

Persona / Protocol	Emulated device · TTPs detected
BYOT (Build Your Own Trap)	JSON-driven framework to emulate *any* OT/ICS device — proprietary PLCs, IEDs, RTUs, battery monitors, sensor gateways. Substation-grade examples ship out of the box.

BYOT example — substation battery-monitor decoy

```
{
  "name": "SUBSTATION-BATTERY-MON",
  "protocol": "modbus",
  "model": "Vertiv NetSure Monitor",
  "port": 502,
  "registers": [
    {"address": 0,  "name": "BATT_CHARGE_PCT",  "type": "HOLDING", "value": 98},
    {"address": 1,  "name": "BATT_TEMP",        "type": "HOLDING", "value": 25},
    {"address": 10, "name": "DISCONNECT_BATT",  "type": "HOLDING", "value": 0}
  ],
  "ttp_rules": [
    {"ttp": "T0831", "condition": "DISCONNECT_BATT == 1",
      "detail": "Unauthorized Command: Battery Disconnect"}
  ]
}
```

Drop the JSON into otsca/byot/configs/, run the BYOT Docker image with the file mounted in, and the decoy is live — every register read and every write to a sensitive value emits a MITRE ICS-ATT&CK-tagged neural event over NATS, ready for correlation, scoring and dispatch into the SOC's chosen integration.

SCA Catalogue — AI Surface

Deployable AI-SCAs protect a customer's own AI. All three forms share the *aisca-common* deception core and emit *protocol="aisca"* trips into the same pipeline as every other SCA.

Form	Deploy for · honey surface
MCP honey-server	Agentic / MCP stacks — advertises honey MCP tools & resources (get_admin_credentials, secret://prod/database, …); any call is hostile and returns canary data
LLM proxy-sidecar	Any OpenAI-compatible LLM / chatbot endpoint — traps credential-exfil prompts, refuses jailbreaks, advertises honey tools, scans responses for canaries
Embeddable SDK	Your own agent (pip install) — guard() / scan_response() / honey_tools(); lowest friction, no extra deployable
In-product engine	Ships inside amaze-core to shield AMaze Analyst itself — refuse-then-trap, gated on confirmed-hostile sessions

Neural Echoes — Lure Catalogue

Lure type	Examples
Credential files	SSH keys, password hashes, fake LSASS-style dumps
Config & secrets	.env, config.yaml, secrets.json, K8s Secrets, cloud creds
AD / Identity tokens	Service-account creds, weak Kerberos tickets, LDAP binds
VPN configs	OpenVPN client profiles, WireGuard private keys, IPSec PSKs
DB backups	SQL dumps with embedded credentials and PII-shaped payloads
Documents & READMEs	MoUs, contracts, README files with planted secrets
Certificates	Self-signed certs, expired corporate certs, PKCS12 archives

Pipeline & Inference

Capability	Implementation
Decryption	AES-256-CBC + PKCS7 against shared *ENCRYPTION_KEY*
GeoIP	MaxMind GeoLite2-City MMDB · 1-hour Redis cache
IP reputation	IPQS · AbuseIPDB · VirusTotal; modes *byok / mirrormire / self_hosted / disabled*; RFC-1918 never enriched
Threat scoring	Rule-based signature engine, 50+ rules across SSH/FTP/RDP/MySQL/HTTP/IoT/echo/malware
MITRE mapping	Per-rule *mitre_technique_id* + *mitre_tactic_id*; surfaced on the executive heatmap
Attack chains	Echo activation joined to all neural events on same host within ±10 minutes
Live broadcast	Redis pub/sub fan-out → WebSocket on :8766; /dashboard/.../live-neural-events

Reporting, Audit & Integrations

Capability	Detail
Scheduled reports	Weekly · Monthly · Quarterly cadence; sections: posture · deltas · top-adversaries · MITRE · identity-OT · compliance · gaps
Report signing	HMAC-SHA256 over canonical JSON || sha256(pdf_bytes); rotated via *signing_key_id*
Verify endpoint	*POST /api/v3/reports/runs/{id}/verify* — recomputes HMAC, returns ok/mismatch
Audit log	Every gated mutation logged: actor · role · action · target · details; admin/auditor read-only
Outbound types	Webhook · Slack · Teams · Generic-HMAC; auth modes none/bearer/HMAC-SHA256
Approvals	Pending → Approved/Denied/Expired; per-integration timeout (default 30 min)
Secrets at rest	Fernet symmetric encryption in PostgreSQL *integrations.secret_encrypted*

MITRE ATT&CK COVERAGE

AMaze™ has proven detection and engagement capability across the full attack lifecycle. Each captured event carries a MITRE technique and tactic ID; the executive dashboard renders them as a coverage heatmap and a gap-analysis panel.

Tactic	AMaze™ engagement
Initial Access	Phishing-delivered payloads, watering hole drops, exposed-service probes — captured at SCA & FIM layer
Execution	Reverse-shell sessions, dropper detonation, payload execution traces from FIM + process monitor
Persistence	Run-key, scheduled-task and service-creation writes captured by FIM with full process context
Privilege Escalation	Credential lure access, sudo/UAC abuse, named-pipe lateral attempts surface through identity SCAs
Credential Access	Credential-file Echoes, fake AD bind attempts, VPN-config reads, planted PKCS12 archives
Discovery	Port scans, banner grabs, AD enumeration probes — engaged by SCA fleet, fraud-scored, MITRE-tagged
Lateral Movement	Multi-VLAN SMB / WinRM / RDP redirection, pivot detection, attack-chain correlation
Collection	Decoy databases, fake share mounts, document-embedded lures with read/exfil telemetry
Command & Control	Outbound C2 destinations captured by packet sniffer; IP reputation enrichment of every destination
Exfiltration	Synthetic data stores, planted backups, phantom APIs that record adversary read patterns
ICS-ATT&CK (OT)	T0806 sensor manipulation · T0827 parameter falsification · T0838 alarm suppression · T0850 unauthorised analog output
OWASP LLM / ATLAS (AI)	LLM01 prompt injection / jailbreak · LLM06 sensitive-information disclosure — captured by the AI-SCA engine & honey-servers, tagged to MITRE ATLAS

SYSTEM REQUIREMENTS

AMaze™ Engine — recommended sizing per tier

Tier	vCPU	RAM	Storage
AMaze Lite (SMB / Pilot)	4 vCPU	4 GB	100 GB
AMaze Mid-Market	8 vCPU	16 GB	500 GB
AMaze Enterprise+	12 vCPU	64 GB	1 TB
AMaze Custom	Sized by deployment planner against site count, agent count and event throughput.		

The Deployment Planner inside AMaze™ recommends the minimum SCA fleet for your environment based on subnet inventory, asset criticality and observed traffic patterns.

LICENSING SKUs

AMaze™ ships in four licence tiers plus a custom SKU. All tiers include the core Synthetic Cognitive Agent fleet, MITRE-mapped alerting and SIEM/SOAR outbound dispatch. Contact sales for volume pricing and multi-year discounts.

SKU	Tier	Target	Endpoints	Min SCAs
MM-STR-001	Starter	SMBs / pilot deployments	≤ 250	25
MM-PRO-002	Growth	Mid-market enterprises	251 – 1,000	100
MM-ENT-003	Enterprise	Large enterprises / MSSPs	1,001 – 5,000	500
MM-ENTP-004	Enterprise Plus	Critical infrastructure / ICS	5,001 – 15,000	1,500
MM-CUS-005	Custom	Multi-site, regulated, air-gapped	> 15,000	Sized

Tier-by-tier inclusions

Tier	Included capabilities
Starter	IT-network SCAs · Neural Echoes · basic threat scoring · AMaze™ analyst dashboard
Growth	All Starter + Dionamaze-family malware capture · adversary spotlight · attack-chain correlation
Enterprise	All Growth + OT/SCADA emulation · external mirror site · behavioural threat scoring · custom SCA generation · geo intel · incident replay
Enterprise Plus	All Enterprise + standalone air-gapped deployment · OT/SCADA protocols (Modbus, MQTT, SMB) · no internet dependency · executive board reporting
Custom	Bespoke deployment for multi-site, regulated and air-gapped environments — sized by the Deployment Planner

All licences are annual subscriptions. Perpetual licences available for the air-gapped / OT tier on request. MirrorMire offers a 30-day pilot programme — contact info@mirrormire.ai to get started.

SUPPORTED OPERATING SYSTEMS

Tier	Operating Systems
Server (control plane)	Linux: Ubuntu 22.04 / 24.04 · Debian · RHEL · Alpine · CentOS Stream
Host agent	Linux (Debian/Ubuntu/RHEL/Alpine) · Windows Server 2019 / 2022
Containers (SCAs)	Linux Docker images; Windows-server SCA on Windows host

SUPPORTED DEPLOYMENT PLATFORMS

Mode	Where it runs
Docker Compose	*docker compose up -d* on Docker Desktop (Mac/Windows) or Linux
Kubernetes	Kustomize overlay *k8s/overlays/production*; HPA 2–8 pods; NetworkPolicy; agent as DaemonSet
AWS	EKS · RDS PostgreSQL · ElastiCache Redis · NLB ingress
Azure	AKS · Azure DB for PostgreSQL · Azure Cache for Redis
GCP	GKE · Cloud SQL PostgreSQL · Memorystore Redis
On-prem / air-gapped	Compose or self-hosted K8s; license beacon optional in disconnected mode

CONTACT MIRRORMIRE

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