Treacle Tech
Security, Trust, Assurance

WHITEPAPER

i-Mirage: AI-Powered Cyber Deception

Catching the attacker at the first move

How custom high-interaction decoys, AI-driven adaptation and deployment on the customer's own terms give security teams a high-confidence signal at the earliest stage of an attack.

30M+	**35+**	**1–2**	**3**
attacker hits recorded at one bank in eight months	previously unknown malware samples identified	months of lead over external threat-intelligence feeds	deployment models: on-premises, hybrid and SaaS

+91-8420407686 contact@treacletech.c U74999WB2021PTC245854 19AAICT6740M1Z
3rd floor, Room 3P 307, Technopark Building, IIT Kanpur, UP- 208016 | | 46A/10, Shibpur Road, Howrah - 711102

Treacle Tech
Security, Trust, Assurance

Contents

01 Executive summary

02 The detection gap

03 Deception as a detection strategy

04 i-Mirage overview

05 Platform architecture

06 Decoys and lures

07 How an engagement unfolds

08 Defending against AI-driven attacks and malicious AI agents

09 Threat intelligence and SOC integration

10 Deployment models

11 Results from the field

12 Use cases by sector

13 Compliance and regulatory alignment

14 Conclusion

— About Treacle Technologies

— Sources

1 Executive summary

i-Mirage is an AI-powered cyber deception platform that turns an attacker's first move inside a network into a high-confidence alert and a contained engagement. It surrounds real assets with custom-built, high-interaction decoys, and plants lures such as fake credentials, files and connection entries on real endpoints and servers. Neither is part of any business workflow: no application depends on them and no procedure points to them. The moment a decoy, or even a lure, is touched, the AI engine triggers. It spawns the decoys that fit what the intruder is looking for and keeps the intruder engaged inside an isolated decoy network while every action is recorded.

> **In brief.** Decoys and lures sit outside every business workflow, so an interaction with one is not routine activity. i-Mirage treats the first touch as the trigger, engages the intruder in the decoy network, and lets their own behaviour there confirm the intent.

Perimeter and signature-based controls still miss intrusions that use valid credentials, zero-day exploits and native administrative tools. Deception closes that gap by detecting intent rather than matching known patterns. i-Mirage adds five things on top of the classic honeypot idea:

- **Custom decoys for any protocol or application.** Every decoy is high-interaction and built to mirror the customer's own environment, including proprietary and sector-specific systems.
- **Lures on real systems.** Fake credentials, files and other breadcrumbs on endpoints and servers turn every real host into a tripwire that leads to the decoys.
- **AI-triggered engagement.** A touch on any decoy or lure triggers the AI engine, which spawns and reshapes decoys around the attacker and holds them inside the decoy network.
- **Intelligence, not just alerts.** Each engagement yields attacker infrastructure, tooling, malware samples and behaviour mapped to MITRE ATT&CK, exportable as STIX 2.1.
- **Deployment the customer chooses.** On-premises including fully air-gapped networks, hybrid, or SaaS, so attacker telemetry stays wherever regulation requires.

The same mechanism holds against attackers who use AI and against malicious or hijacked AI agents. Both must explore, harvest credentials and move, and both touch decoys and lures when they do.

The platform is in production across banking, financial regulation, aviation, government, law enforcement, manufacturing and academia. In the field it has recorded tens of millions of attacker hits, captured previously unknown malware, and flagged hostile IP addresses one to two months before they appeared in other threat-intelligence feeds.

i-Mirage is developed by Treacle Technologies, an Indian cybersecurity company, and grew out of deception research at IIT Kanpur.

2 The detection gap

Attackers now move through a network in minutes, while defenders still take days to notice them. Three recent measurements frame the problem.

Measure (2025 data)	Value	Source
Average eCrime breakout time, initial access to lateral movement	29 minutes; fastest observed 27 seconds	CrowdStrike 2026 Global Threat Report, via TechTarget
Intrusions that used no malware	82%	CrowdStrike 2026 Global Threat Report, via TechTarget
Global median dwell time before detection	14 days, up from 11 in 2024	Mandiant M-Trends 2026, via Help Net Security
Median dwell time when the victim is told by an outside party	25 days	Mandiant M-Trends 2026, via Help Net Security
Average time to exploit a vulnerability, relative to patch release	7 days before the patch (estimate)	Mandiant M-Trends 2026, via Help Net Security

These numbers point to three structural weaknesses in conventional detection.

Signatures need a known bad. When most intrusions carry no malware and exploitation starts before a patch exists, there is nothing for a signature to match. The attacker logs in with valid credentials and uses the tools administrators use.

Anomaly detection drowns in noise. Behavioural analytics must separate a hostile action from millions of legitimate ones that look almost identical. The result is alert volume that analysts cannot triage inside a 29-minute window.

Edge and proprietary systems have no sensor. Network appliances, operational technology and legacy applications often cannot run an endpoint agent. Mandiant reports that these devices are now primary targets for sustained campaigns because of that visibility gap.

The common thread is that every one of these controls watches legitimate assets and tries to infer malice. Deception inverts the problem.

3 Deception as a detection strategy

Deception detects attackers by giving them something false to act on, so the act itself is the evidence. A decoy server, a planted credential or a bait file serves no business purpose. It is not referenced by any application, not listed in any procedure, and not needed for anyone's work.

It would be wrong to claim that nothing legitimate can ever touch such an asset. An authorised vulnerability scanner, an asset-discovery tool or a curious administrator can. A modern deception platform therefore does not stop at the first touch. Known scanners and tools are accounted for during tuning, and what follows the touch separates an accident from an attack. An accidental contact ends there. An attacker tries credentials, enumerates, moves laterally and drops tools. That behaviour, recorded inside the decoy, is the confirmation.

This gives deception three properties that other detection methods lack:

- **High-fidelity alerts.** An interaction with a decoy or a lure needs no behavioural baseline of the production network. False positives are rare, and the few that occur usually trace back to an authorised scanner.
- **Independence from signatures.** A zero-day exploit, a stolen password and a living-off-the-land technique all trip the same wire, because the trap responds to behaviour, not to a known pattern.
- **Attacker cost.** Every minute spent inside a decoy is a minute not spent on production systems, and the attacker cannot tell which assets are real.

3.1 Why first-generation honeypots fell short

Traditional honeypots proved the idea but rarely survived enterprise use. Four limits held them back.

Limitation	Consequence
Low-interaction emulation	A skilled attacker fingerprints the decoy within a few commands and moves on.
Generic, off-the-shelf services	Decoys do not resemble the organisation's real estate, so they attract scanners rather than targeted intruders.
Static placement	Decoys sit where they were installed, regardless of where the attacker actually is.
Alert without context	The SOC learns that something touched a decoy, but not who, with what tooling, or what they wanted.

A modern deception platform has to answer all four: believable depth, environmental fit, dynamic placement and usable intelligence. These are the design goals of i-Mirage.

4 i-Mirage overview

i-Mirage is a deception-based early warning and threat intelligence platform built on one principle: every decoy must be good enough to hold a human adversary, and placed where that adversary actually is.

4.1 What sets it apart

Custom high-interaction decoys for any protocol or application. This is the lead differentiator. Treacle builds decoys to match the customer's environment, from standard IT services to proprietary, sector-specific and legacy applications. There is no low-interaction tier: every decoy offers a full, convincing interaction surface.

AI-triggered engagement from the first touch. The platform does not wait for an attacker to stumble on a fixed trap. When anyone touches a decoy, or even a lure on a real host, the AI engine triggers. It reads what the intruder is looking for, spawns matching decoys in their path, and adapts them as the engagement develops. Each next move leads deeper into the decoy network and away from production.

Lures on every real system. Lures are fake credentials, files, connection entries and other breadcrumbs planted on real endpoints and servers. They extend detection onto the production estate itself. An intruder who lands on a real machine and harvests what is there picks up a lure, and using it both raises an alert and opens a path into the decoys. Lures can be deployed agent-based or agentless. Agent-based is the recommended mode because it gives richer coverage and telemetry.

Cyber mines. Silent decoys sit on unused IP addresses across VLANs and stay dormant until suspicious reconnaissance touches them. They mimic enterprise systems such as payment gateways, Active Directory or VPNs, and give early warning with near-zero false positives.

Self-healing, self-mutating decoys. A compromised decoy rolls back automatically to a clean state while its forensic evidence is preserved. The decoy set also changes over time, so a returning attacker cannot fingerprint it.

Malware capture. A dedicated malware-capture decoy, Octopus, isolates and records reverse shells, ransomware and exploit payloads across more than 15 protocols.

One platform for IT, OT and IoT. The same platform emulates industrial protocols such as Modbus and MQTT alongside enterprise services, and works in isolated networks without agents on production systems.

Intelligence as the output. An engagement produces a structured record of the attacker: source infrastructure, tools, commands, dropped files and techniques. The SOC receives a narrative it can act on, not a bare alert.

Early warning that leads the feeds. In Treacle's field observation, IP addresses that i-Mirage rated critical were published as indicators of compromise by other threat-intelligence platforms one to two months later. Many were tied to successful data exfiltration elsewhere.

4.2 How it compares

Capability	Token-only canaries	Legacy honeypots	i-Mirage
Honeytokens and lures	Yes	Limited	Yes
High-interaction decoys	No	Varies	Every decoy
Custom decoys for proprietary applications	No	Rare	Built per environment
Decoys triggered by attacker activity	No	No	Yes, rule-driven
AI-based attacker engagement and analysis	No	No	Yes
Self-healing decoys after compromise	No	Rare	Yes
ATT&CK-mapped, STIX-exportable intelligence	Limited	Limited	Yes
Deployment	Typically cloud-managed	On-premises	On-premises (air-gapped), hybrid or SaaS

5 Platform architecture

i-Mirage is built as four layers. Decoys are projected into every protected network segment, while capture, analysis and storage sit wherever the chosen deployment model places them. The deception layer runs independently of production systems, so no real asset is exposed at any point in an engagement.

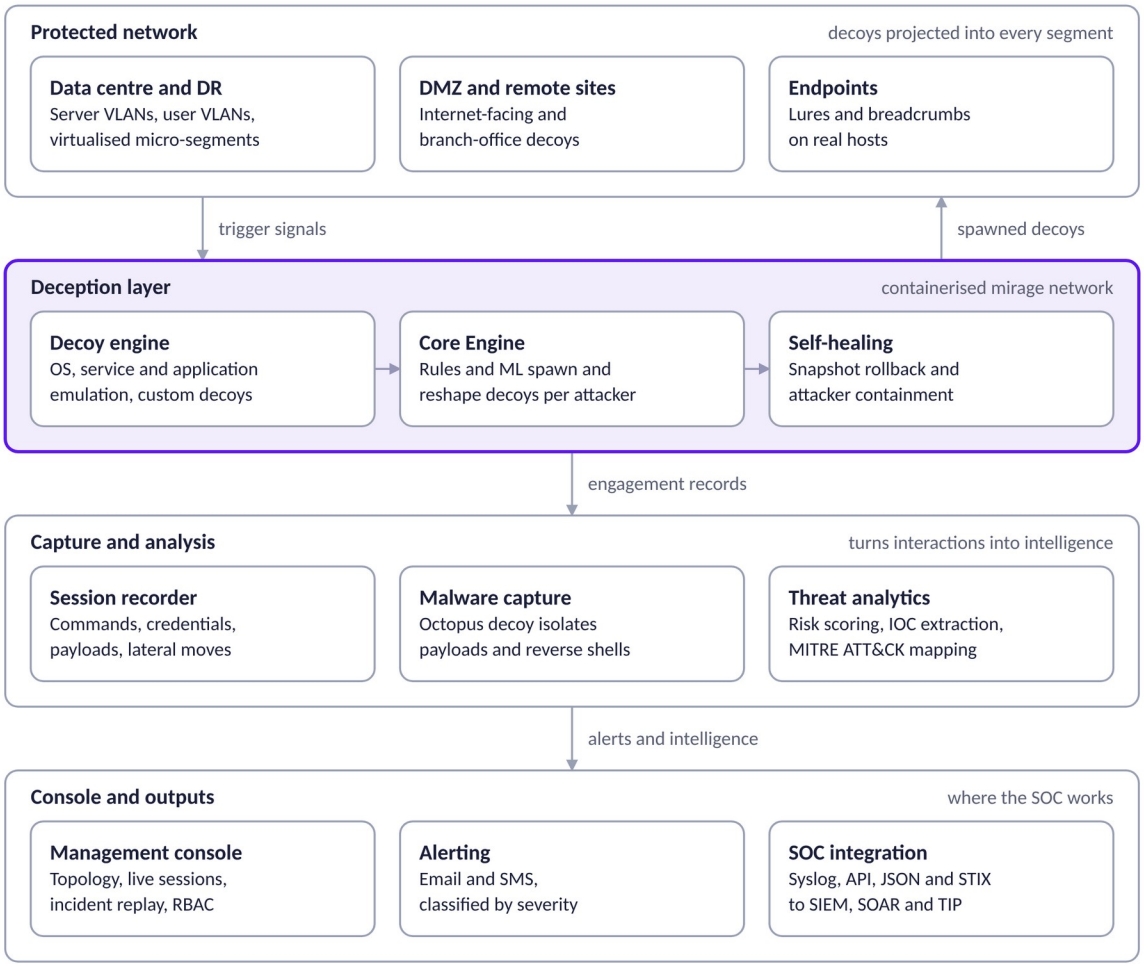


Figure 1. i-Mirage platform architecture

Signals from the protected network feed the deception layer, which spawns decoys back into the attacker's path. Recorded engagements then flow into analysis and on to the SOC.

5.1 Protected network

Decoys are placed in server VLANs, user VLANs, virtualised micro-segments, DMZ zones, disaster recovery sites and remote offices. Lures and breadcrumbs sit on real endpoints and servers and point towards the decoys. Neither appears in any normal workflow, so day-to-day work gives users no occasion to encounter them.

5.2 Deception layer

The deception layer is a containerised mirage network. Every component runs as a container, which allows rapid scaling, modular updates and isolated execution across IT and OT environments.

The Core Engine is the rule-driven orchestration layer. It takes three kinds of trigger:

- **Network triggers**, such as connections to watched ports.
- **File triggers**, such as access to tagged sensitive or bait files.
- **ML log-analysis triggers**, where machine-learning models flag suspicious patterns in log data.

Configurable rules map each trigger to a response, typically spawning a specific decoy in the attacker's path. During an engagement the platform changes decoy characteristics such as ports, banners and credentials to stay convincing and prolong the session. Traffic can be redirected from real network zones to deception zones across VLANs for safe containment.

5.3 Capture and analysis

Every interaction is captured and analysed. A distributed event-streaming pipeline carries telemetry from all decoy zones. The platform reconstructs the attack lifecycle per adversary, maps observed behaviour to MITRE ATT&CK techniques, and extracts indicators such as IP addresses, hashes and malware fragments. A behaviour-based scoring system rates each attacker using dwell time, lateral movement and techniques, and recommends response actions.

5.4 Console and outputs

The SOC works from a web management console with role-based access. It shows decoy deployment and network topology, live attacker sessions, an alert timeline with severity, MITRE ATT&CK technique coverage, a real-time threat map and incident replay. Reports are generated for both analysts and executives.

5.5 Components at a glance

Component	What it does
Decoy engine	Generates and hosts server, application and protocol decoys with operating system, service and banner emulation.
AI adaptation engine	Uses machine learning and rules to infer attacker intent and adjust decoy exposure, services and content during an engagement.
Endpoint deception artefacts	Honeycredentials, decoy documents and other breadcrumbs on real hosts that redirect intruders to decoy zones.
Self-healing engine	Detects a compromised decoy, extracts malware samples, and rolls the decoy back to a clean state from a snapshot while preserving forensic evidence.
Malware capture decoy (Octopus)	Isolates and records reverse shells, ransomware and exploit payloads across more than 15 protocols.
Forensic interaction recorder	Captures complete attacker sessions: commands, payloads, credential attempts, lateral movement and privilege escalation.
Threat intelligence extraction	Turns engagements into indicators and behaviour profiles specific to the organisation.
Telemetry pipeline	Transports telemetry, decoy state and alerts reliably to the console and to SIEM and SOAR platforms.
Central management console	Single point of deployment, monitoring, investigation, reporting and user management.

6 Decoys and lures

i-Mirage ships a broad catalogue of standard decoys and builds custom ones for each customer. Decoys are matched to the customer's estate at deployment: operating system builds, service versions, banners, hostnames and naming conventions.

Category	Examples
Network and system services	HTTP and HTTPS, SSH, RDP, FTP, SMB, DNS, LDAP, SMTP, VPN, printer services
Identity	Active Directory decoys, decoy users, artificial privileged and service accounts
Data	SQL and NoSQL database decoys, file shares, decoy documents
Web attack surface	Decoys built for SQL injection, local file inclusion, cross-site scripting and remote code execution attempts
OT, ICS and IoT	Modbus, MQTT, SCADA and legacy device emulation
Malware capture	Octopus decoy for reverse shells, ransomware and payload analysis
Cyber mines	Silent decoys on unused IP addresses that engage only on suspicious reconnaissance
Custom	Customer-specific applications, portals, APIs and sector protocols such as SWIFT

6.1 Lures

Lures are the part of the deception that lives on real systems. Decoys wait to be found on the network; lures are found by anyone who searches a real host for credentials, files or routes to other systems. That is exactly what an intruder does after the first compromise, and what malware does automatically.

Lure type	Examples	What a touch reveals
Credential lures	Honeycredentials in browsers, memory and configuration files; decoy SSH keys; cloud access keys	Credential harvesting and reuse
File lures	Document honeytokens, bait files and folders	Data discovery, theft and ransomware staging
Connection lures	VPN and browser bookmarks, mapped drives, remote-desktop and database connection entries, application shortcuts	Planning of lateral movement
System lures	Registry entries and named pipes	Host reconnaissance by tools and malware

Lures are designed per department so that they match the real artefacts around them, carry unique identifiers for attribution, and are rolled out centrally through the customer's own tooling, such as Active Directory Group Policy, Microsoft Intune or SCCM.

6.2 From first touch to engagement

A touch on a lure and a touch on a decoy are treated the same way: as a trigger. Opening a bait file, using a planted credential or connecting to a decoy service starts the AI engine. From that point the engine works to keep the attacker inside the decoy network, answering each move with something that looks like progress.

Attacker action	i-Mirage response
Opens a bait file or uses a planted credential	The AI engine triggers and a decoy matching that lure answers
Scans a port	A decoy offering that exact service appears
Tries an exploit	A decoy that appears vulnerable to that exploit responds
Looks for credentials	Realistic fake credentials are presented
Moves laterally	New decoys appear in adjacent subnets
Tries to take data out	Tracked decoy files are served

The attacker believes each step has worked and takes the next one. The cycle repeats inside the decoy network while production systems stay untouched and the SOC watches with the full record in hand.

7 How an engagement unfolds

An i-Mirage engagement converts an intruder's first exploratory step into a recorded, contained session, in six stages.

1. **Trigger.** The attacker touches a decoy or a lure: scans a watched port, opens a bait file, uses a planted credential, or produces a log pattern the ML models flag. Nothing on a production system has been damaged yet.
2. **Decision.** The AI engine triggers. The Core Engine matches the trigger against its rules and selects the decoys that best fit what the attacker appears to be looking for.
3. **Engagement.** The attacker lands in a high-interaction decoy that looks like the next valuable system. They authenticate, enumerate, and begin to work. Each further move is answered by another decoy, so the attacker stays inside the decoy network.
4. **Observation.** Commands, file transfers, credentials tried and tools deployed are all recorded. The decoy adapts to keep the session plausible and the attacker invested.
5. **Alert.** The SOC is notified with severity, source and context. Because the trigger was deceptive, the alert needs no further validation before response begins.
6. **Intelligence.** The session is reconstructed as an attack lifecycle, mapped to ATT&CK, and its indicators exported. Critical source IPs are recommended for immediate blocking and threat hunting across production.

The practical effect is a shift in timing. The defender learns of the intrusion at the reconnaissance or lateral-movement stage, with the attacker's tools already in hand, rather than weeks later from an external notification.

8 Defending against AI-driven attacks and malicious AI agents

Attackers now use AI in two ways: as a tool that speeds up human operators, and as agents that plan and run intrusions on their own. Deception is well suited to both, because it responds to what an intruder does, not to who or what the intruder is.

> **Human or machine, the tripwire is the same.** An AI agent that enumerates a network, harvests credentials and moves laterally touches decoys and lures exactly as a human attacker would, only sooner and more often.

8.1 How AI has changed the attack

Finding	Source
AI-enabled adversaries increased their activity by 89% in a year, using AI for reconnaissance, credential theft and evasion	CrowdStrike 2026 Global Threat Report
Intrusions run by AI agents that make decisions and execute commands rose about threefold in the second half of 2026	CrowdStrike, via Cybersecurity Dive
One group ran an autonomous AI framework against at least 19 organisations in about two weeks, each operation planned and executed by the AI from a one-line goal	Palo Alto Networks Unit 42
Malicious prompts were injected into legitimate generative-AI tools at more than 90 organisations to steal credentials	CrowdStrike 2026 Global Threat Report

Three properties of these attacks matter for defence. They are fast, moving from access to lateral movement in minutes. They are thorough: an agent maps everything it can reach and tries every credential it finds. And they are cheap to repeat, so the same playbook is run against many organisations at once, by operators with little skill of their own.

8.2 Why deception works against AI attackers

- **An AI agent trips the same wires as a human, and more of them.** The plans these agents follow include reconnaissance, credential and secret collection, and lateral movement. Those are exactly the steps where decoys and lures sit. An agent that collects every credential on a host collects the planted ones too.
- **Detection does not depend on recognising the tool.** AI-generated scripts and malware differ on every run, which defeats signatures. A touch on a decoy or a lure is the same event whatever produced it.
- **The AI engine feeds the agent's own loop.** Agentic attack frameworks work in a cycle: decide the next step, act, record the result. When the action lands on a decoy, i-Mirage answers with a result that looks like success. The agent records a false finding and plans its next move from it, which leads it further into the decoy network and away from production.
- **Response at machine speed.** The first touch raises the alert. Integration with firewalls, NAC and SOAR lets containment, such as blocking the source, start without waiting for an analyst.
- **The attacker pays for every step.** Researchers report that compute budget, not engineering skill, is what limits these campaigns. Time and compute spent inside decoys are not spent on real systems.
- **A full record of machine behaviour.** Every command the agent issues is captured, and high-frequency automated activity is classified separately from hands-on intrusions. The SOC sees how the agent works: its tooling, its sequence and the infrastructure behind it.

8.3 Malicious and hijacked AI agents inside the enterprise

Organisations now run their own AI agents and assistants, with access to files, credentials and APIs. Such an agent can be hijacked through prompt injection, or a rogue agent can be introduced deliberately. Either way it acts from inside, with valid access and approved tools, so access controls and signatures see nothing wrong. Deception sees the reach beyond its task.

What the agent does	What catches it
Searches hosts and repositories for keys, tokens and passwords	Credential lures: decoy API keys, tokens, service accounts and passwords that no legitimate process uses
Collects documents it was not asked to work on	File lures: bait files and folders that raise a file trigger when opened
Follows saved connections to other systems	Connection lures that lead to decoys
Probes or logs in to neighbouring systems	Network decoys and cyber mines

Whichever lure or decoy is touched, the AI engine triggers in the same way and engages the agent in the decoy network while the SOC is alerted.

Deception does not prevent prompt injection, and it does not replace identity hygiene, patching and access control, which remain the first line of defence against AI-driven attacks as against any other. It provides the detection layer for the moment those controls are bypassed, by a human or by a machine.

9 Threat intelligence and SOC integration

i-Mirage produces threat intelligence that is specific to the organisation being attacked, in formats a SOC already uses.

Output	What the SOC gets	How it is used
Attack lifecycle	A per-adversary timeline from first touch to last action, with incident replay	Incident triage, reporting and evidence
MITRE ATT&CK mapping	Observed tactics and techniques for each engagement	Detection-gap analysis and purple-team planning
Threat score and recommendations	A behaviour-based risk score per attacker, with recommended actions	Prioritisation and automated response
Severity-scored source IPs	Ranked list of hostile infrastructure, with critical IPs highlighted	Immediate blocking and retrospective threat hunting
Attacker attribution	Geolocation, VPN and Tor usage analysis, IP abuse scoring	Threat-actor profiling
Captured artefacts	Dropped binaries, scripts and attempted credentials	Malware analysis and credential-exposure response
STIX 2.1 export	Indicators and observations in a standard format	Ingestion by threat-intelligence platforms
Alerts	Email and SMS notifications for decoy interaction, sensitive file access and leaked-credential use	Real-time response

Events are forwarded by Syslog, REST API, JSON feeds and webhooks to SIEM, SOAR, threat-intelligence and UEBA platforms. Integration with firewalls, NAC, EDR and PAM supports automated response, such as risk-weighted blocking of attacker IPs.

Generic threat feeds describe what attackers did to someone else, some time ago. Deception intelligence describes what an attacker is doing to this organisation now, with this organisation's decoys as the witness. The two are complementary: i-Mirage indicators have repeatedly preceded their appearance in external feeds by one to two months.

10 Deployment models

i-Mirage supports every deployment mode: on-premises including fully air-gapped networks, hybrid, and pure cloud (SaaS). The customer chooses the model, so attacker telemetry is captured, analysed and retained wherever its regulatory position requires.

> **The customer chooses the model.** Regulated segments can remain entirely in-house, with no internet connection, while other parts of the estate use hybrid or SaaS delivery.

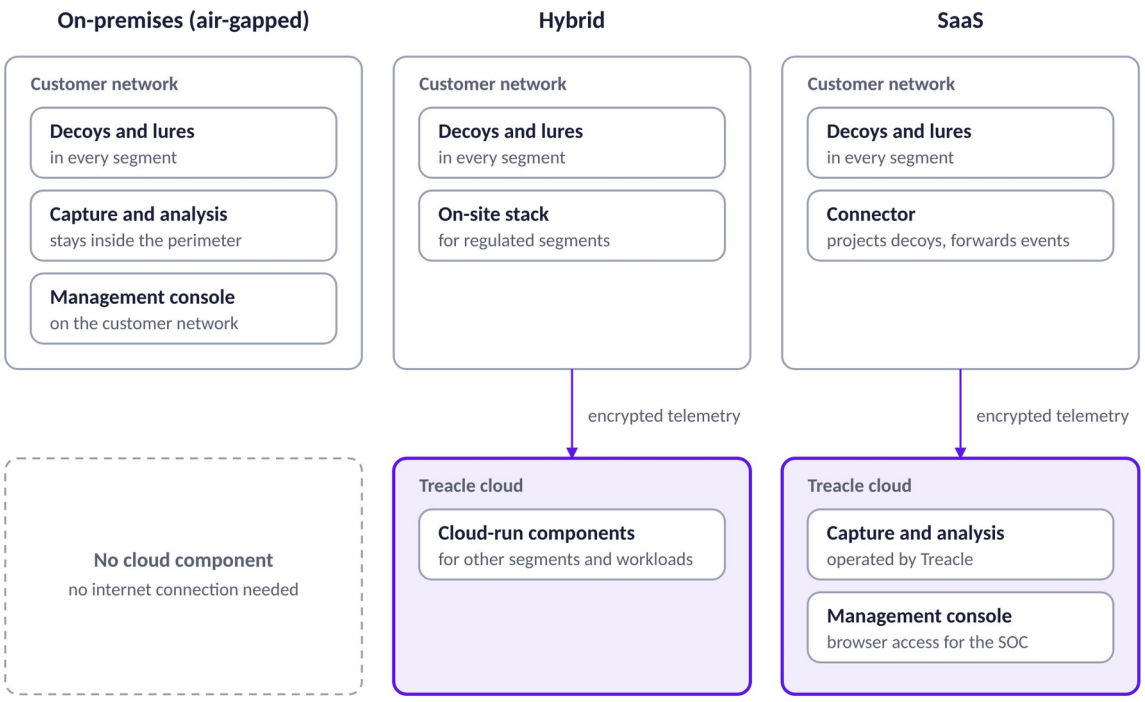


Figure 2. Deployment models: on-premises, hybrid and SaaS

In all three models the decoys and lures sit inside the customer's network segments. What changes is where capture, analysis and the console run.

	On-premises (air-gapped)	Hybrid	SaaS
Decoys and lures	Customer network	Customer network	Customer network
Capture, analysis and storage	Customer data centre	On-site for regulated segments, cloud for the rest	Treacle-operated cloud
Management console	Customer network	On-site or cloud, as agreed	Cloud, browser access
Internet dependency	None	Only for cloud-run components	Required
Typical fit	Banks, regulators, government, OT and other isolated networks	Multi-site estates, IT and OT convergence, cloud workloads	Fast rollout, proofs of concept, lean security teams

10.1 Form factors

- **Hardware appliance.** Dedicated, pre-configured appliances for on-premises deployment.
- **Virtual machine.** Installed on the customer's existing virtual infrastructure with no additional hardware.
- **Cloud-native.** The complete platform on public cloud infrastructure, with full feature parity.
- **Private cloud or hosted data centre.** Deployed within the customer's own private cloud with complete administrative control.
- **Hybrid for IT and OT.** Components distributed across cloud and on-premises systems.
- **Modular.** Architecture tailored to network scale, compliance requirements and the integration stack.

10.2 On-premises and air-gapped operation

In an on-premises deployment the platform needs no internet connectivity for attacker engagement, analysis or reporting. All telemetry, alerts and forensic data stay within the customer's network boundary. Standard controls apply:

- Decoys are isolated from production applications and deployed on-subnet with strict egress blocks.
- DMZ decoys accept inbound connections only.
- Any outbound connection from a connector, such as log forwarding to a SIEM, is restricted to pre-approved endpoints.

- Decoys carry synthetic data only. No production or confidential data is replicated.

10.3 Multi-site and DC-DR reference architecture

Large estates run i-Mirage across a primary data centre, a disaster recovery site and remote offices under one console.

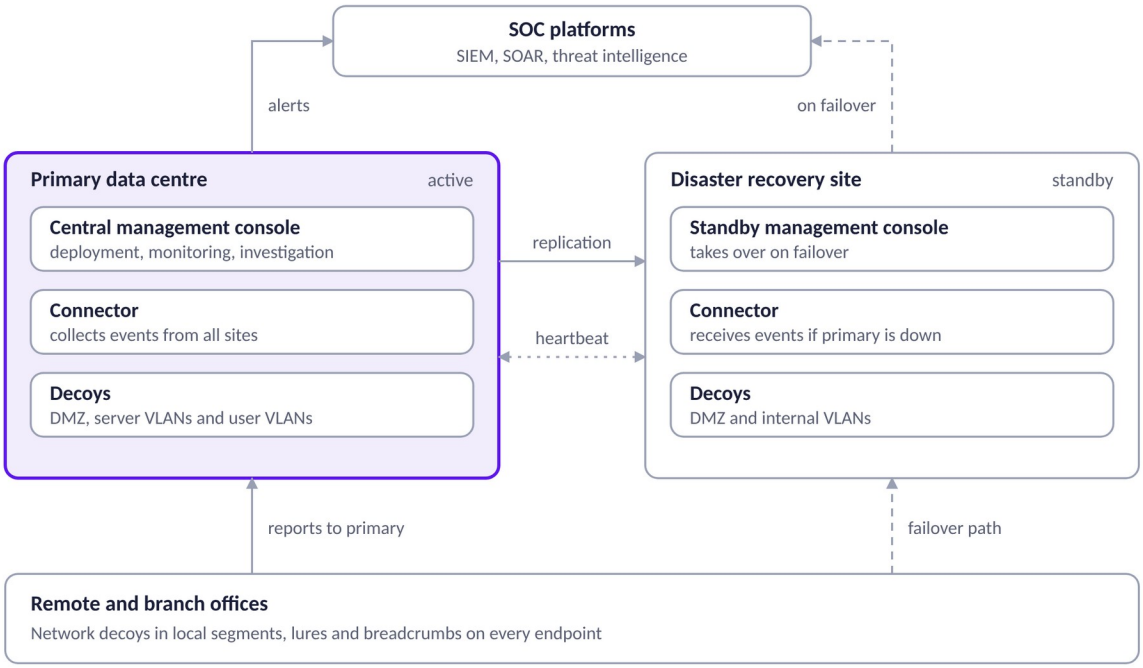


Figure 3. Multi-site reference architecture with DC-DR failover

The central management console and a connector run at the primary data centre, with a standby console and a second connector at the disaster recovery site in active-passive mode. Event data and decoy configuration replicate from active to standby, and a heartbeat monitors the pair. If the primary site becomes unavailable, endpoint lures and decoys at every location report to the disaster recovery connector, so detection continues through the failover.

10.4 How a deployment runs

1. **Discovery.** Passive discovery over a VLAN trunk or mirror port identifies live services and segments with minimal disruption, supplemented by targeted manual validation.
2. **Criticality mapping.** Discovered services are ranked by business criticality, exposure and likelihood of attack.
3. **Decoy classification.** Standard decoys cover commonly targeted services, custom decoys cover applications unique to the organisation, and adaptive strategies govern how decoys evolve.

4. **Deployment planning and rollout.** Decoys are mapped to segments, department-specific lures are packaged and pushed to endpoints, and placement is aligned with the DC-DR design.
5. **Integration and acceptance.** The platform is integrated with SIEM, SOAR and threat-intelligence platforms, tuned, and validated through controlled red-team probes and user acceptance testing.

Decoy sizing depends on the size and mix of the estate and is determined during solution design. In India, i-Mirage is typically delivered through system-integrator partners, with Treacle as the OEM providing design, custom decoy development and support.

11 Results from the field

i-Mirage is in production at organisations across seven sectors, and its results come from live adversaries rather than lab tests. Customer names are withheld; figures are from Treacle's deployment records.

Environment	Result
Large public sector bank	30 million+ attacker hits recorded in 8 months
Large public sector bank	3+ previously unknown malware samples captured
Large public sector bank	Threat actor linked to a banking-sector data leak identified
Financial-sector regulator	5+ previously unknown malware samples captured
Airport operator, multiple airports	10 million+ attacker hits recorded; previously unknown malware captured
Across deployments	More than 35 previously unknown malware samples identified
Across deployments	More than 850 malware samples and more than 200 exploit payloads captured
Across deployments	IPs rated critical appeared in external threat-intelligence feeds 1 to 2 months later

Three observations follow from this record.

> **Early warning in practice.** IP addresses that i-Mirage rated critical appeared in other threat-intelligence feeds one to two months later.

Volume is real and constant. Tens of millions of hits against decoys show how much hostile activity reaches a large organisation's edge and interior without ever raising a conventional alert.

Deception finds what signatures cannot. Malware that no vendor had catalogued was captured because the attacker chose to deploy it on a decoy. No prior knowledge of the sample was needed. Captures have included ransomware carrying target-specific markers, a fileless memory-resident trojan, and a hybrid Linux malware invisible to signature-based tools.

The lead time is actionable. A one-to-two-month head start on hostile infrastructure is enough to block it and hunt for earlier contact before the same actor succeeds against a production system.

Customer retention reflects this: deployments in aviation and academia have been renewed repeatedly, one now in its third renewal.

12 Use cases by sector

The same platform serves different goals in each sector, because the decoys are built to look like that sector's crown jewels.

Sector	What attackers want	Typical decoys and lures	Outcome
Banking and financial services	Payment systems, customer data, privileged credentials	Decoy banking applications and servers, credential and file lures	Early warning of ransomware, credential reuse and lateral movement; targeted malware captured
Financial regulators and market infrastructure	Supervisory data, market-sensitive systems	Decoys mirroring internal and internet-facing applications	Evidence of who is probing regulated infrastructure
Aviation and critical infrastructure	Operational systems, passenger data, disruption	Decoys across multiple sites, including sector-specific services	A single view of attacks across distributed facilities
Government and law enforcement	Citizen data, case systems, espionage footholds	Decoy portals and internal services	Zero-day and APT detection, with attribution-grade intelligence
Manufacturing and OT	Design data, intellectual property, plant control	ICS protocol decoys, engineering and file-system decoys, bait documents	Detection of SCADA lateral movement and IP-theft attempts
Telecom	Subscriber data, core network elements	Telecom-specific service decoys	Early attack detection, including rogue base station detection
Healthcare	Patient records, medical devices	Patient-data lure traps, medical device mimicry	Early detection of data-theft and device-targeting attempts
Academia and research	Research data, compute, student records	Decoys on open campus networks	Visibility into a high-noise, hard-to-control environment

Across sectors, the platform is used for:

- **Early threat detection** during reconnaissance, before an attacker reaches critical assets.
- **Insider and stolen-credential detection**, because lures and decoys lie outside every role's normal work, and using one shows intent.

- **Lateral movement detection** as an intruder pivots between segments.
- **Targeted malware capture**, including samples no public source has seen.
- **Ransomware early warning**, since operators enumerate file shares and credential stores before encrypting.
- **SOC optimisation**, replacing alert volume with a small number of confirmed incidents.
- **Air-gapped, OT and IoT protection**, where endpoint agents cannot run.
- **Remote office and legacy system protection**, where other sensors are thin.
- **Compliance and forensics**, with a replayable, ATT&CK-mapped record of each incident.
- **AI-driven attack and rogue AI agent detection**, because an agent trips the same decoys and lures as a human intruder.
- **Automated containment** through SOAR and firewall integration.

13 Compliance and regulatory alignment

Indian regulators now name deception explicitly, and i-Mirage outputs map onto the frameworks auditors already use.

- **RBI.** The Cyber Security Framework for Banks lists honeypot services among the capabilities of a bank's security operations centre (Annex 2).
- **SEBI CSCRF.** The framework recommends deception technology for regulated entities (DE.CM.S3). In its August 2025 clarification, SEBI encouraged regulated entities to deploy decoy systems (report).
- **NCIIPC.** Its July 2021 newsletter recommends implementing honey tokens as a deception technique, in its guidance on preventing supply chain attacks (newsletter).
- **MITRE ATT&CK.** Every engagement is mapped to ATT&CK techniques, which supports control-coverage reporting and gap analysis.
- **NIST Cybersecurity Framework.** Deception strengthens the Detect function with a high-confidence signal, and feeds Respond with attacker context.
- **Data residency.** On-premises and air-gapped deployment keeps all attacker telemetry inside the organisation's own infrastructure.
- **Audit evidence.** Recorded engagements, alert histories and lifecycle reports provide evidence of continuous monitoring and of detection controls working against real adversaries.

Treacle Technologies is ISO 27001 certified and holds a SOC 2 Type 2 attestation. i-Mirage supports these objectives; it does not by itself make an organisation compliant. Mapping to a specific clause should be confirmed with the organisation's auditor.

14 Conclusion

Attackers who use valid credentials and unpatched flaws leave conventional controls with nothing to match, and they move in minutes. Deception gives defenders a signal that does not depend on knowing the attack in advance.

i-Mirage makes that signal practical at enterprise scale. Its decoys are custom-built and high-interaction, so they hold a skilled adversary. Its lures turn real endpoints and servers into tripwires. A touch on either triggers the AI engine, which places decoys where the attacker is and keeps the attacker engaged in the decoy network. Its output is intelligence the SOC can block, hunt and report on. Field results across banking, regulation, aviation, government, manufacturing and academia show the approach working against live threats.

The recommended first step is a scoped deployment around one critical application, with custom decoys designed to mirror it.

About Treacle Technologies

Treacle Technologies is an Indian cybersecurity company founded in 2021. It builds AI-powered deception and early-warning systems for IT and operational environments. i-Mirage, its flagship platform, originated in deception research at IIT Kanpur and is engineered and supported in India. The company is ISO 27001 certified, holds a SOC 2 Type 2 attestation, and has received grants from MeitY, the Department of Science and Technology, the Department of Telecommunications and the Startup India Seed Fund.

More information: treacletech.com

Sources

- Help Net Security: Attackers are handing off access in 22 seconds, Mandiant finds (Mandiant M-Trends 2026)
- TechTarget: News brief, Attackers gain speed in cybersecurity race (CrowdStrike 2026 Global Threat Report)
- Business Wire: 2026 CrowdStrike Global Threat Report, AI accelerates adversaries
- Cybersecurity Dive: Agentic AI helps hackers level up, even without detailed instructions
- Unit 42: Attackers use autonomous AI to target semiconductor, defense and civil society sectors
- NCIIPC Newsletter, Volume 5 Issue 3, July 2021

- CAalley / The Economic Times: SEBI clarifies on cybersecurity and cyber resilience framework
- Product details, regulatory references and field results: Treacle Technologies product documentation and deployment records