


# Grow your managed services. Not your headcount.

Evolvable.ai is the AI agent platform for MSSPs. Run AI agents and agentic workflows across every client's security stack from one platform. Agents do the repetitive work, your analysts make the decisions, and every action is approved, logged and reported per client.



- 
**Grow revenue, not headcount**  
 Agents take the tier-1 load: triage, enrichment, containment prep and reporting. Your analysts handle more clients and higher-value work.
- 
**Standardise every delivery**  
 Turn your best analyst's playbook into an agentic workflow. Every client gets the same quality, at 3 p.m. or 3 a.m.
- 
**One platform, every client**  
 Each client is an isolated tenant with its own connectors, model, policies and approvers. Build once, roll out to all.

## MORE THAN AN AI SOC TOOL A complete AI platform under every client.

- 
**Agentic workflows, not static playbooks**  
 80 node types with agents inside the flow, plus approval and decision steps. Logic that adapts, beyond SOAR.
- 
**Dual RAG over each client**  
 Runbooks, past incidents and threat intel as vector search plus a knowledge graph, separate per tenant.
- 
**AI-assisted playbook building**  
 The Builder agent and builder MCP API let your engineers ship a new playbook in hours, not sprints.
- 
**Agents that are secured themselves**  
 AI firewall with 60+ detectors on every agent, so an agent that can isolate devices cannot be hijacked.
- 
**Oversight per client**  
 Five gate modes, per-client approvers and SLAs, and a full audit trail for every action taken.
- 
**Proof per tenant**  
 ROI, quality and response KPIs per client, ready for QBRs and renewals.

NEW SERVICES YOU CAN SELL


AI security & Shadow AI for clients

Agentic SOC tiers

Compliance & executive reporting

# Where agents work *in your SOC.*

ONE ALERT, FIVE MINUTES, ONE DECISION BY A PERSON



12 SECURITY USE CASES, READY TO RUN PER CLIENT

| Detect & respond                                                                                 | Exposure & intel                                                                           | Report & comply                                                                         |
|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| <b>Alert triage &amp; enrichment</b><br>Every SIEM and EDR alert enriched, scored and summarised | <b>Threat intel briefings</b><br>Client-specific intel from your feeds, every morning      | <b>Client reporting</b><br>Monthly and executive reports drafted per tenant             |
| <b>Phishing response</b><br>Reported emails analysed; quarantine on approval                     | <b>Vulnerability prioritisation</b><br>Findings ranked by exposure and active exploitation | <b>Compliance evidence</b><br>Control evidence collected for audits                     |
| <b>Identity threat response</b><br>Risky sign-ins checked; disable or revoke on approval         | <b>SAST &amp; code findings triage</b><br>False positives filtered, fixes suggested        | <b>Client service desk</b><br>Clients ask about tickets and status in Teams or WhatsApp |
| <b>Incident containment</b><br>Isolate devices and block indicators, gated by an analyst         | <b>Third-party risk</b><br>Vendor questionnaires reviewed against evidence                 | <b>AI security for clients</b><br>Shadow AI, AI firewall and governance as a service    |

|                                                            |                                                            |                                            |                                                 |
|------------------------------------------------------------|------------------------------------------------------------|--------------------------------------------|-------------------------------------------------|
| ILLUSTRATIVE · 20-CLIENT MSSP<br><i>What it adds up to</i> | ≈ 1,000<br>analyst hours a month returned from tier-1 work | +8<br>clients onboarded with the same team | 1<br>playbook library delivered to every client |
|------------------------------------------------------------|------------------------------------------------------------|--------------------------------------------|-------------------------------------------------|

Directional example: 12,000 tier-1 alerts a month closed by agents with evidence at ~5 minutes of analyst time saved each. Your numbers are baselined in the pilot.

WORKS WITH THE STACKS YOUR CLIENTS ALREADY RUN

|                                                                                                         |                                                                                                    |                                                                                                     |
|---------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| <b>SIEM</b><br>Microsoft Sentinel · IBM QRadar · Google SecOps · Elastic · Exabeam · Securonix · Rapid7 | <b>EDR / XDR</b><br>CrowdStrike · Defender XDR · SentinelOne · Cisco XDR                           | <b>Identity</b><br>Entra ID · Okta · CyberArk · Ping · Duo                                          |
| <b>Threat intel</b><br>Recorded Future · Mandiant · Anomali · MISP · OpenCTI · GreyNoise                | <b>Response &amp; ops</b><br>Cortex XSOAR · PagerDuty · HashiCorp Vault · ticketing via MCP or API | <b>Collaboration</b><br>Microsoft Teams · Slack · email · WhatsApp for approvals and client updates |

Connector availability varies by version; anything else connects through MCP or REST.

BUILT FOR MULTI-CLIENT DELIVERY

|                                                                                                                       |                                                                                                               |
|-----------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Tenant per client</b><br>Schema-per-tenant isolation. Each client's data, keys, connectors and logs stay separate. | <b>Template library</b><br>Publish agents, workflows, guardrails and skills once; assign them to any tenant.  |
| <b>Clone &amp; roll out</b><br>Copy a proven agent to a new client in minutes, then point it at their tools.          | <b>Client-specific approvals</b><br>Per-client approver groups, SLAs and escalation in Teams, Slack or email. |
| <b>Proof of value per client</b><br>ROI by tenant: hours saved, AI cost and outcomes, ready for QBRs.                 | <b>Sovereign deployment</b><br>On your infrastructure or the client's, any LLM, air-gapped if required.       |

*Pilot it on two of your clients.  
Prove it before you scale it.*

We stand up your tenants, connect their security tools and launch two playbooks with your analysts in the loop. Partnership and revenue-share models available.

BECOME A PARTNER →

info@evolvable.ai · www.evolvable.ai