Extend the Identity Control Plane With *Aembit IAM and CrowdStrike* AI Detection and Response

Aembit's Identity and Access Management (IAM) platform gives every AI agent its own verified identity and enforces run-time, least privilege access to enterprise MCP servers, tools, and data.

This integration with CrowdStrike AI Detection and Response (AIDR) centralizes agent traffic, content inspection, and policy enforcement so risky content is caught without bypassing identity and access controls.

Agentic AI Risks

- Without enforced policies, agents can access MCP servers and tools that they shouldn't.
- Malicious requests can come from unauthorized agents.
- Trusted agents can carry sensitive data or return malicious instructions.
- Content inspection deployed per client or in agent code requires piece-by-piece configuration.
- Identity, access, and content decisions live in separate, disconnected logs.

With Aembit + CrowdStrike AIDR

- Access policies control exactly which MCP servers and tools agents can reach and use.
- Unauthorized agents are blocked from accessing sensitive resources.
- Tool inputs and outputs are inspected in real time to block or mask sensitive or malicious information.
- Centralized content inspection, with no per-client AIDR proxy or custom application code required.
- Every access event and agent action is correlated in one auditable record.

Why Aembit + *CrowdStrike*

Centralized enforcement in a single control plane.

Customers can apply zero-trust access control and content inspection in the same enforcement path, with no per-client proxy or application-level code required.

Strengthen AIDR with verified identity for every agent.

Aembit attaches a blended human-agent identity, plus the access policy decision behind it, to everything CrowdStrike AIDR evaluates.

Built on a complete identity foundation.

Every AI agent already has its own identity, secretless access, and a full audit trail under Aembit. CrowdStrike AIDR adds content security on top of that foundation.

Request a demo at www.aembit.io

aembit | CROWDSTRIKE

How It *Works*

1 Apply by access policy

Attach a CrowdStrike AIDR Content Security configuration to the access policies protecting your most sensitive MCP servers, using the Aembit UI, API, or Terraform.

2 Inspect three checkpoints

AIDR evaluates tool listings, tool inputs, and tool outputs against detection rules you define in the CrowdStrike console.

3 Enforce the decision

Aembit applies CrowdStrike's outcome at its enforcement point: pass content through unchanged, block it, or forward a transformed version.

4 Record everything

Workload events capture the AIDR decision, an explanation when CrowdStrike provides one, correlation data, and any inspection errors alongside Aembit's identity and access record.

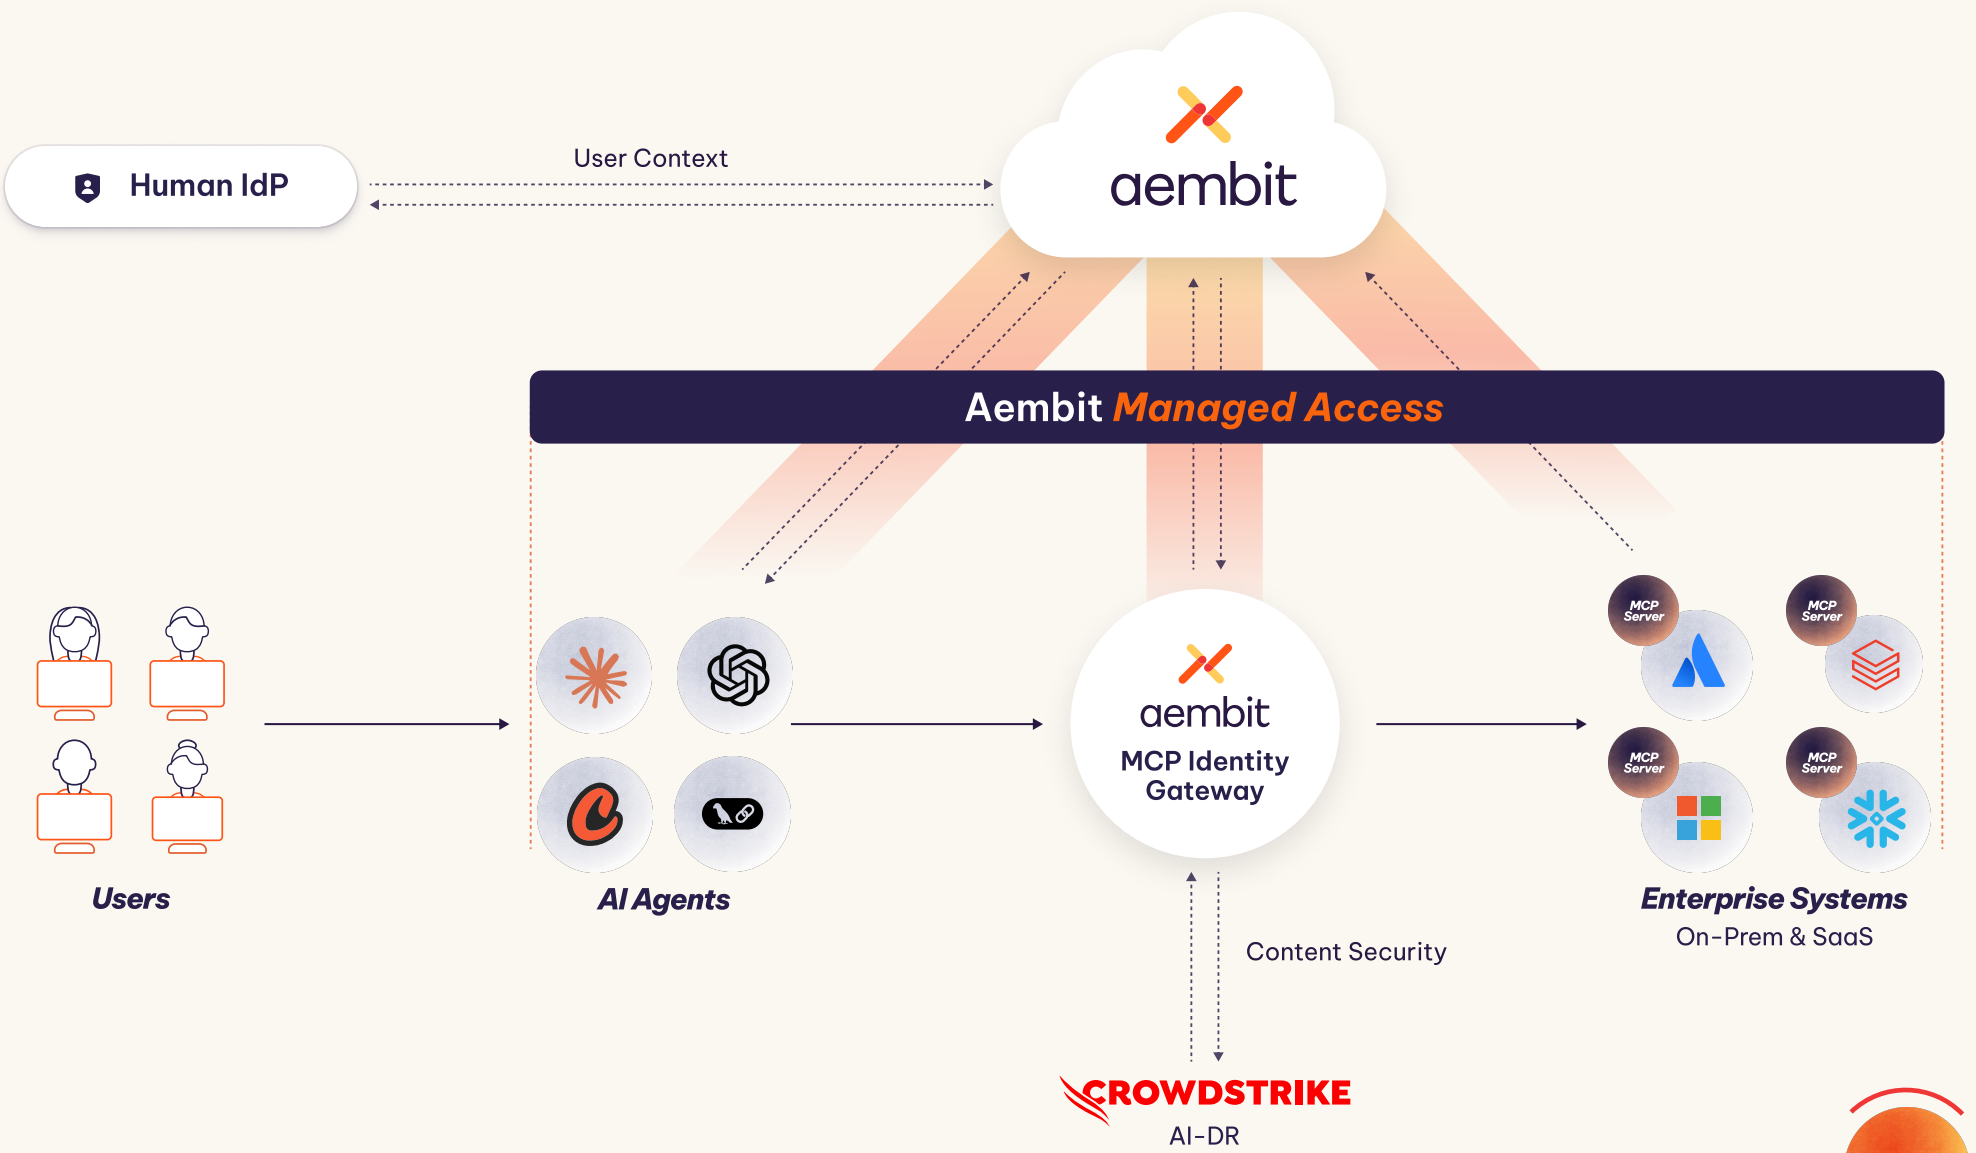

Requires a CrowdStrike AIDR for Agents subscription.