AI Security, Side by Side

Four companies are selling security for AI agents to large enterprises right now. Each one covers a different side of the same problem: agents can read data, use tools, and take action, and someone has to keep that safe. This page lays the four across the same rows, with Agentic Armor beside them, so the difference in ground is easy to see.

Since this comparison was first prepared, the house has filed its designs. Six U.S. provisional patent applications stand filed, on November 16, 17, and 21, 2026, Patent Pending: The Backbone, No. 64/200,228. The Helix Registry, No. 64/201,147. Armorgram, No. 64/201,166. The Agentic Keys, No. 64/201,184. The Agentic Dome, No. 64/204,100. The Agentic Teller, No. 64/204,118. The four companies' rows stand exactly as the market note states them; what changed is the Armor column, which now names the filed designs behind its rules.

What they do

Check Point Governance

An enterprise platform that watches how AI is used across a whole organization and applies policy at machine speed. It covers shadow AI employees adopted on their own, agents acting inside the business, sensitive data flowing to models, AI applications, and everyday workforce use. The promise is visibility and control from the inside: see what every agent is doing, and decide what may pass, what needs conditions, and what must stop.

Aikido Security Offense

Autonomous agents that penetration test what a company ships. One set of agents maps and attacks web apps, APIs, networks, mobile apps, and AI agents and MCP setups. A separate set re-exploits every finding to prove it is real before it is reported, so false positives are dropped. Reports come out audit grade for SOC 2 and ISO work within hours. An airgapped on premises version exists for regulated buyers.

Salt Security Discovery

An outside in assessment that maps the AI agents a company has exposed to the internet, the MCP servers those agents can reach, and the APIs behind them. It delivers the exposed surface, the attack paths an attacker could use today, and a prioritized fix list.

Britive Runtime control

Control of the privilege an agent holds while it works. Every action runs one arc: verify the identity of the agent, assess what the request asks for, authorize at the moment it happens, observe while the work runs, and revoke when the task ends. Privilege is created for the single action rather than granted in advance, the default is deny, sensitive actions can wait for a person to approve before any access exists, and every decision is recorded as it happens, down to the prompt and the arguments used.

Agentic Armor Household runtime control

Security for the agents working inside a person's daily life. The house runs runtime control as plain rules: approvals bound to the exact transaction, no standing grants for money, credentials, exports, sends, or deletes, and out of band confirmation before a consequential act. The Agentic Dome is the consumer shape of the same argument: an agent living under a consent roof the family itself set, in code the family can read, where what is private stays private until its owner lets it out.

Who they sell to

Check Point Governance

Large enterprises with security departments, procurement processes, and compliance deadlines.

Aikido Security Offense

Large enterprises, with an airgapped on premises version for regulated buyers.

Salt Security Discovery

Large enterprises that need to know what they have exposed.

Britive Runtime control

Large enterprises buying runtime control as architecture.

Agentic Armor Household runtime control

Households, solo founders, and small businesses, and the agents now working inside their daily lives, who face the same class of risk with no security team at all.

How the work starts

Check Point Governance

From the inside, across the whole organization, watching AI use as it happens.

Aikido Security Offense

As an attack on what the company ships, run by autonomous agents.

Salt Security Discovery

From a domain alone. No internal access, and nothing installed. The assessment usually lands inside 24 hours.

Britive Runtime control

At runtime, action by action, while the agent works.

Agentic Armor Household runtime control

Inside the home and the person’s own work, under consent rules the family or the founder set for their own agent.

What it costs

Check Point Governance

The market note states no price.

Aikido Security Offense

A flat rate. A typical assessment runs around $4,000, and the standing term is that no high or critical findings means no pay.

Salt Security Discovery

The market note states no price.

Britive Runtime control

The market note states no price.

Agentic Armor Household runtime control

The market note states no price.

What they do not do

Check Point Governance

It governs enterprise AI use. It is not built for a household or a small organization, and it does not attack or test what a company ships.

Aikido Security Offense

It tests and reports. It does not run the day to day control of an agent while the agent works, and it is not built for a household.

Salt Security Discovery

It looks from the outside and delivers an exposed surface and a fix list. It does not go inside, and it does not control an agent at runtime.

Britive Runtime control

It sells control to enterprises as platform architecture. It is not a product a family can read, set, and run at the kitchen table.

Agentic Armor Household runtime control

Armor is not built for large enterprises with security departments and procurement processes. That is not its ground.

Where Armor stands against each one

Check Point Governance

Same inside view, different ground. Check Point gives a security department sight and control over a company. Armor gives a person sight and control over their own agent, with no security department required.

Aikido Security Offense

Opposite direction. Aikido attacks what companies ship, for a fee, and proves the holes. Armor stands on the protection side, guarding the person’s own daily life rather than testing an enterprise.

Salt Security Discovery

Outside in versus inside the home. Salt finds what a company exposed without ever going in. Armor works where the agent actually lives and acts, inside the household’s own consent rules.

Britive Runtime control

The closest cousin. The runtime control Britive sells as enterprise architecture is, in substance, the doctrine this house already runs as plain rules. The difference is the buyer and the shape: a platform for a security team there, household rules and The Agentic Dome here.

Agentic Armor Household runtime control

The four prove the market agrees the problem is real. None of them is building for the kitchen table. That ground is Armor’s.

The four, one line each

Check Point governs from the inside

See what every agent in the company is doing, and decide what may pass, what needs conditions, and what must stop. The doctrine named here is the subject of the house’s six filed applications, listed at the head of this comparison, each Patent Pending.

Aikido attacks

Agentic penetration testing that proves each finding before it is reported, around $4,000 a typical assessment, and no high or critical findings means no pay.

Salt discovers

From a domain alone, with no internal access, it maps the exposed agentic surface and hands back the attack paths and a prioritized fix list.

Britive controls privilege at runtime

Privilege created for the single action, never granted in advance, with every decision recorded as it happens.

Where Agentic Armor stands

The four sell to large enterprises with security departments. Agentic Armor’s ground is households, solo founders, and small organizations, the people those platforms do not serve, who face the same class of risk with no security team at all.

The runtime control Britive sells as enterprise architecture is the doctrine this house already runs as plain rules: approvals bound to the exact transaction, no standing grants for money, credentials, exports, sends, or deletes, and out of band confirmation before a consequential act.

The Agentic Dome is the consumer shape of that same argument: an agent living under a consent roof the family itself set, in code the family can read, where what is private stays private until its owner lets it out. The large platforms prove the market agrees the problem is real. None of them is building for the kitchen table.

First prepared November 5, 2026, from the house market note of the same date. Prices and terms are as the note states them. Updated November 23, 2026, to the filed state: six U.S. provisional patent applications, Patent Pending.