

SECTOR PACKAGE · AI AND AUTONOMOUS AGENTS

AI agents that act, actions that can be proven

Control, human approval and verifiable evidence for artificial-intelligence agents, automations and autonomous workflows.

Your agents no longer just recommend: **they use tools, modify information, execute code and trigger entire workflows**, frequently without continuous human intervention. A valid credential is no longer enough as a control. *Who limits what each agent can do, who approves what is sensitive, and what evidence remains of each action?*

WHAT VELANTIA BRINGS TO THIS SEGMENT

AUTONOMY — governs each agent: **capabilities, scopes, conditions, human approval, revocation and traceability**, also for robots and automated processes.

GATE — decides each critical action an agent requests: **allow, condition, reduce, route to approval or block**, based on policy and context.

SENTINEL — observes behavior, **detects anomalies** and attributes each action to its agent, identity and tool.

EVIDENCE — preserves a **verifiable history** of each decision and outcome, ready for audit, compliance and investigation.

WHAT YOU CAN DO WITH IT

- ▶ **Define limits per agent and tool:** rules, permissions and execution policies per agent, tool, data and action.
- ▶ **Require human approval** before sensitive, irreversible or privileged actions — without slowing the rest of the automation.
- ▶ **Monitor in real time** the activity of agents, automated workflows and integrations, with their context.
- ▶ **Introduce agents with less risk:** the damage of a compromised agent is bounded to what its permission allows, with a record of everything it does.

The agent can act — **within explicit limits and with an attributable history of what it does.**

DEPLOYMENT

Software, virtual machine, private cloud or Trust Appliance, with unified governance in **VELANTIA ONE**. For dedicated, local control of agents, the **AUTONOMY GATEWAY 500** appliance (available for pre-order). It integrates with your AI, identity and security platforms **without replacing the existing stack**.

WHY NOW

Every month, **more critical actions run without continuous human intervention**, and non-human identities already outnumber human ones. Regulators and customers are starting to **demand proof** of which control was applied to each autonomous action. Whoever does not govern the agent today will only reconstruct what it did tomorrow.

Next step: an initial meeting and a demonstration.

Full flow in a controlled environment: agent action → policy → approval or block → evidence. A scoped AI-agent security pilot proposal is available, with metrics.

Agustín López Mayor — Inventor · Co-founder · **Stephan Hoof** — Co-founder · International Expansion

contact@velantia.io · velantia.io · Villajoyosa (Alicante, Spain) · Weikersheim (Germany) · Ruse (Bulgaria)

Nine patent applications filed with the Spanish Patent and Trademark Office (OEPM) · pending · Strategic presentation document