Deploy agents fearlessly.

Enterprises give AI agents goals and access to business resources. Agents choose which tools, systems, and other agents to use, and those connections can change on every run. The enterprise remains accountable for the authority it delegates.

Maya attributes network connections to the responsible agent, checks them against approved boundaries, and enforces policy outside the agent, without code changes or traffic decryption.

How Maya secures agents.

Know your agent estate.

Find known and shadow agents. See where they run and which models, tools, systems, and other agents they communicate with.

Keep agents within bounds.

Define approved connections to systems and services. Detect departures as agents and their dependencies change.

An agent-level kill switch.

Block an affected agent's communications or a specific connection. Keep unrelated, permitted activity running.

Demonstrate control.

Link the agent, approved mandate, observed deviation, and response. Show what was allowed, what happened, and the control applied.

Maya applies boundaries in the context of each agent.

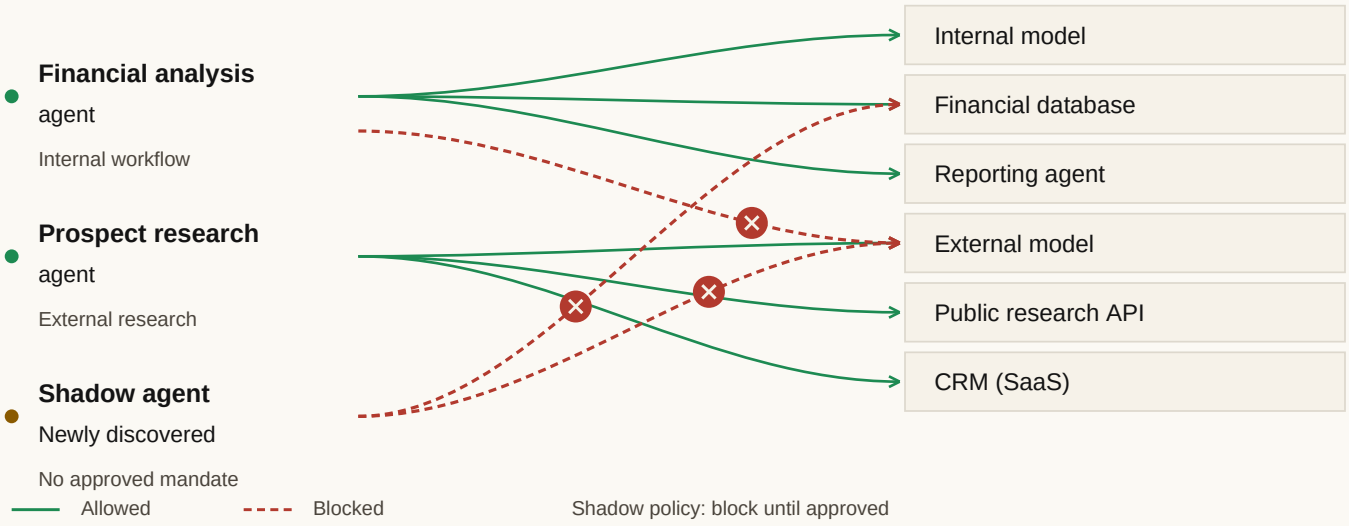

Zero agent code changes | **Payload-blind by default** | **Distributed enforcement**

© 2026 Mayagentic, Inc. All rights reserved.

Control across your agent estate.

Bring security and governance to where agents run and the connections they create.

ACROSS YOUR OPERATING ENVIRONMENTS

Cloud workloads	Agent platforms	Endpoints	On-premises	SaaS
VMs, Kubernetes, and containers	Managed and self-hosted runtimes	Developer and employee devices	Internal infrastructure	Agents in business applications

How Maya maintains control.

1 Discover — 2 Declare — 3 Observe — 4 Detect — 5 Decide — 6 Enforce — 7 Prove

Approved intent, observed connections, decisions, and evidence stay linked as agents evolve.

A kill switch at the scope you need.

Block a connection.

Stop communication with a specific destination.

Contain an agent or group.

Restrict one agent or cooperating agents.

Stop fleet communications.

Apply a broad stop across the agent estate.

Bring agent adoption under control.

Find unmanaged exposure.	Identify shadow agents and the systems they reach.
Respond without broad disruption.	Contain an affected agent or connection while permitted work continues.
Demonstrate governance in practice.	Explain the approved boundary, observed deviation, and action taken.

See what your agents can reach.

Request a demo.

mayagentic.com/request-demo

About Mayagentic. Mayagentic builds security infrastructure for enterprises adopting AI agents. Maya brings agent context to network enforcement, helping teams govern delegated authority.

© 2026 Mayagentic, Inc. All rights reserved.