



MEET TOKENVISOR SPACES

AI Agents on the job. Humans in control.

AI agents that work across files, code, browsers and tools.
Governed at every step.

SUPERMICRO

Builds the infrastructure.

AMD

Powers the compute.


EMBEDDED LLM

Governs models & agents.

Your agents' workspace.

TokenVisor Spaces • Persistent. Isolated. Governed.

YOU STAY IN CONTROL



Parallel agents
Multiple agents.
Work together.

Agent-to-agent
Coordinate task
hand-offs.
Work as a team.

**Visible Task
State**
See what every
agent is doing.

Human Approval
Sensitive actions.
Wait for your go-
ahead.

Isolate agent workspaces

TokenVisor Spaces provides isolated agent workspaces to securely execute code, manage files, and browse the web.

Any model. Your rules.

Deploy your preferred models on Supermicro while TokenVisor automatically enforces access policies, budgets, and routing.

Human-controlled actions

Sensitive actions require human approval, while side-by-side activity logs make auditing effortless.

Real work. Clear Boundaries.


SUPERMICRO • AMD • EMBEDDED LLM

EXPLORE THE GOVERNED AI STACK →



INFRASTRUCTURE · COMPUTE · GOVERNANCE

Three partners. One governed AI stack.



Four steps. Zero blind spots.

- | | | | |
|---|---|--|---|
| 1 | 2 | 3 | 4 |
| TASK | POLICY | APPROVAL | AUDIT |
| An authorized user opens a Space and briefs an agent. | Every model call passes through TokenVisor for access control, routing, and usage limits. | Sensitive actions require human approval before execution. | Every agent action and model request is logged for audit and oversight. |

Bring one workflow. Leave with a production plan.

Sit down with Supermicro, AMD, and Embedded LLM to map your target models, users, data requirements, and infrastructure.

Hardware	system_aisolutions@supermicro.com
Software	info@embeddedllm.com

