EXECUTIVE BRIEFING

# WorkingAgents

## Executive Guide

### The Execution Control Layer for AI Agents

Control at the point of action, inside your infrastructure.

SECTION 01

# Executive summary

AI agents are already operating inside the enterprise. They read, write, approve, and transact across systems that were built to be operated by people.

Guardrails shape what a model says. They do not decide what an agent is permitted to do.

Identity, data loss prevention, logging, and orchestration were built for human users and static applications. They do not govern autonomous systems acting at machine speed.

The gap is at the action layer. Every approval, record update, API call, and transaction an agent attempts is a decision. Most are accepted by default.

WorkingAgents is the Execution Control Layer. It sits between AI agents and enterprise systems and enforces control at the point of action, inside the customer's infrastructure, with zero data egress and full chain of custody.

The question is no longer whether agents will operate inside the enterprise. It is whether their actions will be mediated and proven, or accepted by default.

## Control is not a model problem. It is an action problem.

SECTION 02

# The shift

Enterprise controls were built for a human pace. Access was provisioned. Behavior was audited after the fact. The model held because the actor was a person or a known application.

*Agents break the model.*

An agent is not a user. An agent is not an application. It chains decisions, invokes tools, and acts on requests it only partially understands. It can read, write, approve, and transact in seconds, in parallel, across systems that were never designed to be operated by non-humans.

The risk is not that agents are unsafe. The risk is that no layer decides, in real time, whether a specific action should be permitted. Every downstream system is asked to trust the agent by default.

Trust by default is not a policy. It is the absence of one.

*Enterprises control access. They do not control actions.*

SECTION 03

# Why existing approaches fail

Each existing control has a role. None enforce at the point of action.

### Identity and access management

IAM validates identity. It does not authorize behavior. A valid token is not a decision about an action.

### Data loss prevention

DLP inspects content. It does not evaluate actions. A clean payload can still trigger a harmful operation.

### Logging and observability

Logs record events. They do not prevent them. By the time a log exists, the action has been accepted.

### Orchestration frameworks

Orchestration coordinates steps. It does not mediate them. Control inside the agent is control the agent can bypass.

*Systems validate identity. They do not evaluate intent.*

SECTION 04

# What an Execution Control Layer is

An Execution Control Layer is where an agent's intended action is mediated before it reaches the system that would carry it out.

It is not a filter on language. It is a decision on behavior.

It classifies the action, applies policy, enforces the outcome, and produces evidence.

It is the first point where action is decided, not observed.

*Without a layer, actions are accepted by default. With a layer, actions are mediated.*

SECTION 05

# Where WorkingAgents sits

WorkingAgents sits at the action boundary — between AI agents and the enterprise systems they operate.

Not at the model. Not at the user. Not at the network edge. At the point where an action becomes a consequence.

The boundary covers SaaS tools, internal APIs, databases, and MCP-exposed services. Policy is enforced in line. The action is permitted, modified, blocked, or routed for approval. Every decision is recorded.

It runs inside the customer's infrastructure. Data does not leave. Chain of custody is preserved end to end.

*Not a layer above the stack. A layer inside it.*

SECTION 06

# What changes when control moves to the point of action

Approvals become deterministic, not discretionary.

Policy becomes enforceable, not aspirational.

Audit becomes evidence, not reconstruction.

Delegation becomes a governance decision, not a leap of faith.

Agent scope expands, because the risk surface is contained where it matters.

*Leadership gains a single place to answer: what is an agent permitted to do here, and how do we prove it?*

SECTION 07

# What this is not

The category is defined as much by what it excludes as by what it enforces.

### Not an agent builder

Customers build and run agents. WorkingAgents governs what those agents are permitted to do.

### Not a model provider

Model-agnostic by design. Neutral on provider.

### Not a monitoring tool

Monitoring observes. Control decides.

### Not a dashboard

A dashboard reports outcomes. A control layer determines them.

### Not a SaaS overlay

Inside the customer's infrastructure. Not above it, not beside it.

SECTION 08

# Executive decision checklist

Eight questions a leadership team should be able to answer before scaling agent deployment.

**01.** Who can state, today, what actions an AI agent is permitted to take in our production systems?

**02.** When an agent takes an action, where is that decision evaluated, and by what?

**03.** If auditors asked us to prove a specific agent action was authorized, what record would we produce?

**04.** Are our current controls evaluating actions, or only identities, content, and logs?

**05.** If an agent were compromised tomorrow, what contains the blast radius at the action layer?

**06.** Which systems in our stack accept agent input without independent mediation?

**07.** Do we have a single layer that enforces policy across all agent-to-system interactions, or is policy scattered across orchestration code?

**08.** Is our agent strategy constrained by our ability to govern agent actions, or by something else?

SECTION 09

# Call to action

Three paths. Each is a decision, not a discussion.

### For a leadership conversation

Align your executive team on the category and the control gap. One hour. No technical prerequisites.

### For a security review

Scoped architectural review with your CISO or head of platform security.

### For a technical review

Direct engagement with the team running agent deployments and MCP integrations.

## Contact

liem@workingagents.ai

workingagents.ai