

Meet Us Where We Work

Agentic coworkers, governed participation, and the case for voice

AN XSI WHITEPAPER · JUNE 2026



Abstract

Most enterprise AI agents fail their teams not on capability but on **participation**. They run in a window off to the side, operated by hand, returning results a human has to carry back into the systems where work actually happens.

This paper argues that an agentic coworker must meet a team where it already works — on Microsoft Teams and Slack, in Azure DevOps and Jira, in the repo and the wiki, in the tools each function already lives in — and must be assigned, scored, and managed by the same discipline applied to any other resource on the team.

It describes how **XSI-AIMS**, XSI's agent-governance model, makes that discipline concrete: an Orchestrator that routes the right task to the right resource, a human Sovereign at the top of the authority chain, and a deterministic oversight layer that holds every agent accountable.

It closes with the first surface XSI is delivering on this principle — **conversational voice support, now in preview**, for XSI LodeStone and XSI-AIMS for the Microsoft Agent Framework — and is explicit about the surfaces still on the roadmap.



Meet the Team

On Teams, Slack, DevOps, Jira — where work already happens



Manage the Agent

Assigned, routed, scored, and held accountable like any resource



Start with Voice

Conversational presence, now in preview for LodeStone and XSI-AIMS for the Microsoft Agent Framework.

Key Themes



Presence over capability



Agents as managed resources



Governed participation



The XSI-AIMS authority chain



Voice as the first surface



Honest gaps and roadmap

1. The Presence Problem

Consider the most capable hire a team has ever made — faster than anyone, never tired, never bored — who then refuses to use any of the team's tools. No standup, no board, no wiki. Finished work arrives as text a human re-keys into the system by hand. That hire does not last, and not because of any deficit in skill.

That is the shape of the AI agent most organizations have deployed. The capability is real. The integration is not. The agent is a thing you operate from the outside — you bring it the work and you carry back the result — and so it never becomes part of how the team runs.



No Standup

Never shows up where the team coordinates

No Board

Work never enters the system of record

No Wiki

Knowledge stays locked in a private window



2. An Agent Is a Resource on the Team

The discipline of management predates software, and it has always been a single act — assigning the right task to the right resource. Knowing who is suited to what. Knowing what to delegate and what to keep. Knowing when to step in and when to stay out of the way. Done well, it is most of what distinguishes a functioning team from a group of capable individuals.

An agent is a resource subject to exactly this discipline. It will excel at some tasks, be adequate at others, and need a human closely involved on the rest. Treating it otherwise — as a magic box that either works or doesn't — wastes the capability and creates risk.



Present

The agent must be where work happens — on the team's own surfaces, not off to the side.



Routed

Work must be directed to it by something that knows what it is good for.



Accountable


It must be held to the same standard as every other resource on the team.

The useful frame is the management frame: agents are, in the words of the XSI-AIMS model, **virtual colleagues with definitions, behaviors, drift modes, and improvement needs**. They get roles. They get reviews. They get held to the work.



3. Meet Us Where We Work

Coordination does not happen in a chat box bolted to the side of the organization. It happens in Microsoft Teams and Slack, where conversations and decisions live. It happens in Azure DevOps and Jira, where work is assigned and tracked. It happens in Confluence and the wiki, where knowledge is written down, and in the repository, where it is built. These are not incidental tools — they are the surfaces on which a team is a team.



A Teammate

- Is **present** on the team's surfaces
- **Participates** in the workflow directly
- Operates under the **same governance**




A Tool

- **Waits** to be picked up and operated
- Requires a human to **carry results back**
- Creates **context switches** and friction

4. The Right Task to the Right Resource

Routing is where the management frame becomes architecture. In the XSI-AIMS model, a team of agents is structured the way a team of people is — with distinct roles and an explicit chain of authority — rather than as a single undifferentiated assistant. AIMS defines ten agent **archetypes**, each carrying one essential function, so that "send this work to the right resource" has a concrete answer instead of a hope.



AIMS Archetype	Team-Role Analogue	What It Does
Sovereign	The principal / accountable owner	Origin of authority; carries the human's intent; the one everything answers to
Visionary	Ideation lead	Generates candidate approaches and reformulations
Architect	Standards / policy owner	Defines the regulatory, ethical, and procedural structure
Orchestrator	The manager / dispatcher	Routes each task to the resource suited to it; coordinates delegation
Enforcer	Quality and compliance gate	Holds the registry of approved agent definitions; gates what may run
Provider	Resourcing / operations	Allocates budget, capability, and access
Relay	Intake / front desk	Sanitizes and gates inbound traffic from outside the team
Executor	Doer / individual contributor	Invokes and manages the agents that perform the work
Articulator	Delivery / packaging	Composes work into the form each target environment needs
Sustainer	Project continuity	Manages long-running tasks, retries, and working memory

5. Scored and Managed the Same Way

You do not manage a resource you cannot observe. XSI-AIMS gives every agent the same accountability surface a person has: a defined role, known failure modes, an audit trail, and correction when it drifts.

The Witness Layer

A cross-cutting audit protocol woven across the framework, under which framework operations are **observed and recorded** rather than trusted on faith. Every action leaves a trace.

Deterministic Shadow Governance

A rule-based oversight layer that runs **underneath the agents** and, critically, *cannot be modified by the agentic layer it watches*. Each archetype is monitored against its own characteristic failure signature.

Because the shadow layer is separate from and immune to the agents it governs, the oversight does not degrade as the agents adapt. When behavior drifts from the role the agent was given, the deterministic layer flags it.

The purpose is not surveillance. It is **parity**. An agent on the payroll should be as accountable as anyone else on it — assigned work, reviewed output, a record of what it did and why. A teammate that cannot be held accountable is not a teammate; it is unmonitored output.

AIMS makes the accountability structural, so that adding agents to a team adds capacity without subtracting oversight.

100%

Every Action Traced

No operation runs without a recorded audit trail in the Witness Layer

10x

Role-Specific Monitoring

Each of the ten archetypes is monitored against its own characteristic failure signature

0

Drift Detection

Deterministic shadow governance flags deviations the moment behavior diverges from assigned role

6. The Function Decides the Tool

Tools change with the job, and they always have. Marketing works in Canva, in Figma, in the Adobe suite. Sales works in Dynamics 365 and Salesforce. Engineering works in the repository and the issue tracker; finance, in the ledger.

A human teammate meets the function where the function already operates — no one asks a designer to lay out a campaign in a spreadsheet. An agentic coworker has to meet the same standard: the integration target is not one surface but many, chosen by the function the agent serves.

This is deliberately ambitious. It is also what the word *coworker* demands: a coworker adapts to the team's tools, rather than asking the team to adapt to it.



Marketing

Operates in Canva, Figma, and the Adobe suite — where creative work already lives.



Sales

Works inside Dynamics 365 and Salesforce, meeting the pipeline where it runs.



Engineering

Operates in the repository and the issue tracker — the surfaces engineers already own.



Finance

Works in the ledger and financial tooling, where the numbers are kept and reviewed.



Field Service

Works in field service management platforms — scheduling, dispatch, and on-site resolution where the work physically happens.



Customer Service

Operates in service desks and support platforms, meeting customers and agents where cases are opened, tracked, and resolved.



7. Voice Comes First

One tool predates every other on this list. The voice. Teams talk — they interrupt, they redirect, they say 'not like that' before work goes too far in the wrong direction. A large share of how a team coordinates is never typed anywhere; it is spoken, in a meeting or on a call.

An agent reachable only by typing into a box is an agent held permanently at arm's length from that channel. So that is where XSI is starting.



Join a Meeting

An agent can participate by speech, joining a meeting the way a colleague would.



Take a Huddle

Answer and be interrupted naturally, in real-time conversational flow.



On Your Infrastructure

XSI LodeStone runs on the team's own infrastructure — a sovereign agentic appliance built on kernel-isolated orchestration.



Governed by Default

Same archetypes, same authority chain, same witness and shadow layers. What changes is the surface.

8. Gap: What Is Not Yet Here

Stating where the work is honest about its own boundaries is part of taking it seriously.

Voice Is Preview, Not GA


Conversational voice support is in preview, not general availability. It is an early surface with a defined scope, offered to a limited audience. It is a starting point, deliberately, and not a finished claim.

Broad Tool Participation Is Roadmap

A teammate that lives natively inside Teams, Slack, Azure DevOps, Jira, Confluence, the repository, and the function-specific tools of marketing, sales, and finance is the destination the category demands. XSI is building toward it; XSI is not claiming it is complete. Connector-by-connector participation is roadmap, and this paper marks it as such rather than implying it is done.

Governance Maturity Leads Surface Maturity

The accountable, routed, audited substance in Sections 4 and 5 is the strongest part of the model today. The surfaces that substance reaches are catching up to it. That ordering is intentional — the governance has to be real before the participation is safe.



9. The Team You Actually Wanted

Return to the capable hire who would not use the tools. Nothing about that hire's ability had to change. What had to change was presence — standing up, walking over, joining the meeting, using the board, taking the call. Becoming, at last, a member of the team.

That is the whole aim. Not agents smarter than the people they work with, but agents present alongside them — on the same surfaces, routed by the same management discipline, accountable through the same oversight — so the work a team and its agents do together comes out better than what either would produce alone.

Voice is where that presence begins, because it is the surface teams have used the longest. The rest is what XSI is building toward, in the open, with the gaps named.

XSI-AIMS

The Agent Instrumentation and Management Specification — XSI's open agent-governance standard. Archetypes, four operating layers (L-Intent / L-Plan / L-Form / L-Exec), three functional modes (M-Generative / M-Constraining / M-Integrative), the witness layer, and shadow governance are defined there.

XSI LodeStone

XSI's sovereign agentic appliance and operating environment for running governed agentic coworkers on a team's own infrastructure.

XSI-AIMS for Microsoft Agent Framework

XSI's runtime governance for agents built on the Microsoft Agent Framework — bringing governed, voice-capable resources to teams already invested in that ecosystem.

 **Preview status:** Conversational voice support is offered in preview / limited availability as of June 2026. Scope and availability are subject to change. Microsoft Agent Framework, Microsoft Teams, Azure DevOps, and Dynamics 365 are products of Microsoft. Slack, Jira, Confluence, Salesforce, Canva, Figma, and the Adobe suite are products of their respective owners; references here are descriptive of where teams work and imply no endorsement or integration beyond what is stated.