

# THE NEXT EVOLUTION IN ENTERPRISE AI IS PURPOSE-BUILT



Mainframe organizations continue to run some of the most critical systems in the global economy while managing increasing complexity and rising expectations for innovation. Artificial intelligence (AI) provides a practical way to scale expertise, reduce operational friction, and help teams move work forward with greater confidence.

BMC believes organizations will increasingly leverage AI to broaden access to mainframe capabilities and integrate these systems more fully into their broader enterprise IT ecosystems. In many cases, AI will also act as a catalyst for modernization in place—helping organizations evolve applications, workflows, and operational practices while continuing to run critical workloads on the platform that already powers their business.

## **The Direction: From AI assistance to coordinated intelligence**

Enterprise AI is moving beyond generating insights and explanations. Organizations now expect intelligence that can safely participate in execution across critical systems. **Our direction is clear:** [BMC AMI](#) solutions are evolving into intelligent participants in a governed ecosystem of AI agents that collaborate autonomously across development, operations, data, and security workflows to accelerate innovation on the mainframe.

Supporting this direction requires enterprise AI on the mainframe to operate across three coordinated layers:

**Intelligence layer** — Domain knowledge and reasoning grounded in mainframe expertise, operational telemetry, and institutional knowledge.

**Coordination layer** — Orchestrated AI agents collaborating across workflows to connect

understanding with action.

**Governance layer** — Policy-aware controls ensuring AI-driven actions remain secure, transparent, auditable, and subject to human oversight.

Together, these layers enable AI to extend beyond analysis and participate responsibly in execution. This foundation now enables the next phase of our direction: extending AI beyond knowledge and insight toward coordinated, governed execution across the BMC AMI platform.

## AI in production: Trust earned, not claimed

In July 2024, we published a [Statement of Direction](#) focused on [infusing intelligence](#) directly into BMC AMI product experiences. Our objective was straightforward: reduce complexity, close knowledge gaps, and make it easier for teams to work confidently on the mainframe.

That direction was grounded in three core principles:

- **Consistent and contextual intelligence**, with a common AI foundation infused across the AMI portfolio
- **LLM freedom**, allowing organizations to choose the models that align with their enterprise AI strategies
- **In-the-moment intelligence**, embedding AI directly into workflows to surface answers and insights where work happens

Since then, we have delivered.

[BMC AMI Assistant](#) was introduced beginning with [GenAI-powered code explanation](#) and expanded across consecutive quarterly releases into a cohesive intelligence layer embedded throughout the BMC AMI portfolio. Using [AI as an advisor](#), developers gained faster understanding of unfamiliar applications, while operators gained explainable insight into system and operational issues. Investigation that once required manual analysis increasingly includes guided insight and next steps directly inside the workflow.

The addition of [Knowledge Expert Chat and Knowledge Hub](#) further expanded this intelligence layer. Together, these capabilities combine curated product intelligence, decades of BMC mainframe expertise, and each organization's own institutional knowledge, making trusted guidance available at the moment of decision. **This approach helps less-experienced staff work with greater confidence while allowing senior experts to focus on higher-value work, turning institutional knowledge into a shared operational asset.** Watch this [short video](#) for a demo of how Knowledge Expert Chat provides answers at the moment of need.

Instead of searching across documentation, runbooks, and tickets, teams can surface answers through natural-language questions using a [knowledge expert](#) chat inside their workflow. Early usage shows that this approach not only accelerates troubleshooting and decision-making but also helps users discover capabilities and documentation they did not previously know existed.

Trust in AI is earned through real outcomes. This intelligence foundation now supports the next phase of our direction.

# Customers are redefining what AI must do

Through advisory boards, design programs, beta participation, and enterprise [AI readiness](#) discussions, a clear message has emerged: customers want AI that does more than explain what happened. They want orchestrated intelligence that helps move work forward.

They want to reduce repetitive analysis, move faster from insight to action, and make it easier for teams to build expertise and focus on higher-value priorities. They want to realize the full promise of [BMC AMI as automated mainframe intelligence](#).

Generative AI laid the foundation by helping teams understand, explain, and accelerate work. Agentic AI builds on that foundation by enabling intelligence to plan, decide, and safely carry work forward.

To meet these expectations at enterprise scale, our new Statement of Direction is grounded in five core principles:

- **Establish the core first:** Build an orchestrated agentic intelligence foundation.
- **Specialized agents:** Reflect domain expertise with focused, autonomous actions.
- **Human-in-the-loop:** Governed, transparent, and observable autonomy.
- **Open ecosystem:** Open standards-based, composable, and extensible.
- **Outcome-driven adoption:** Deliver measurable value early and expand over time.

The next phase extends AI beyond insight toward trusted participation in execution, within enterprise guardrails.

And that shift defines our next statement of direction.

## From AI assistance to orchestrated intelligence

Building on the foundation of the AI-driven knowledge and guidance we have delivered with BMC AMI Assistant, our next statement of direction is clear: extending that foundation toward proactive, coordinated intelligence that can safely execute actions within defined governance policies across the BMC AMI portfolio.

This evolution introduces coordinated AI agents with specialized, policy-aware capabilities operating across development, operations, data, and security. Each agent is domain-specific, working together within enterprise guardrails to deliver governed, collaborative outcomes.

In this next phase, AI agents for operations, development, data, and security will operate as orchestrated participants across domains—supporting actions such as system and performance diagnostics, development workflows, security validation, and operational recovery. They will learn across past incidents and participate in execution with transparency, verifiability, enterprise governance, and human oversight built in by design.

We will expand beyond explanation and recommendation to enable validated action inside workflows. Agentic workflows will move from detection to investigation to explanation to resolution within a single, governed AI experience.

AI will not replace expertise. Human validation remains essential. Principal engineers and system programmers will continue to define policy, validate outcomes, and shape how intelligence operates. AI scales expertise rather than removing it.

The result is faster problem identification, more contextual decision-making, and reduced operational friction—while ensuring the mainframe participates fully in broader enterprise AI strategies. The mainframe will operate as a connected, policy-aligned participant in the enterprise AI ecosystem across hybrid environments, maintaining the security, reliability, and trust on which organizations depend.

## Delivering agentic workflows across the BMC AMI portfolio

Over the coming releases, this direction will take shape through **agentic workflows that coordinate knowledge, reasoning, and action across BMC AMI solutions**. Initial workflows could include agentic AIOps incident resolution, agentic diagnostics, AI-assisted application insights, development troubleshooting, and test-case generation.

These workflows build on the **intelligence layer** already established across the BMC AMI portfolio, combining development and operations telemetry, mainframe domain expertise, and organizational knowledge to provide the context required for responsible execution.

From there, [coordinated AI agents operate across solutions](#), aligning development, operations, data, and security activities that were previously siloed. Instead of isolated AI features, intelligence becomes part of the workflow itself, helping teams connect system understanding with the next operational step.

Every action remains governed and transparent. **Human validation, enterprise policy, and operational guardrails remain central**, ensuring AI participation strengthens reliability rather than introducing risk.

The objective is clear: translate insight into trusted action while preserving the control and discipline enterprise systems require.

## Establishing the foundation for governed AI execution

As AI becomes more operational, governance and trust become increasingly essential. Organizations expect AI not only to inform decisions, but to operate safely, predictably, and transparently within enterprise policy boundaries.

Across the industry, fragmented approaches are already emerging: isolated AI integrations, independently managed execution layers, and disconnected tool endpoints. While intended to accelerate innovation, these models often introduce new complexity, inconsistent policy enforcement, and operational sprawl.

Enterprise mainframe environments cannot afford that fragmentation, especially as AI Agent orchestration becomes a core capability for coordinating multi-step agentic workflows. This can result in agents giving conflicting recommendations and teams lose confidence and revert to manual work. Without a unified and governed approach, orchestration itself becomes fragmented, leading to unpredictable behavior, loss of insight, and reduced trust in AI-driven operations.

As part of our next phase, our direction includes establishing an **MCP Gateway** as a shared, governed access layer across the BMC AMI portfolio. Rather than creating multiple independently governed AI entry points, the MCP Gateway will provide a centralized, secure, and policy-aware interface through which AI agents interact with BMC AMI solutions. Every AI-driven action will be visible, governed, and aligned to enterprise policy with consistent controls across the platform. This

enables AI agents not only to access BMC AMI systems, but to safely execute actions across them within defined policy boundaries.

Supporting this architecture, we will introduce an **Agent Gateway** to facilitate how AI agents communicate and collaborate. Instead of agents interacting independently, interactions will flow through the Agent Gateway—where the agent interactions are visible—ensuring governance, auditability, logging, and policy enforcement across agentic workflows.

Together, the MCP Gateway and Agent Gateway extend the BMC AMI Platform into a governed AI foundation that enables coordinated intelligence and trusted execution across the portfolio (see Figure 1). The BMC AMI Platform serves as the enterprise foundation for this direction—providing a unified layer of core capabilities and services that connects intelligence, orchestration, and governance across the BMC AMI portfolio. It brings together innovations such as BMC AMI Assistant, the Agent Gateway, and the MCP Gateway to simplify mainframe transformation, accelerate innovation, and enable AI to operate consistently and securely at scale.

Industry standards define how AI systems connect to tools and agents. Our direction focuses on how those connections are governed, secured, and operationalized at enterprise scale—without introducing fragmentation or operational risk.

This is how we intend to evolve agentic AI into an *enterprise-grade foundation* across the BMC AMI portfolio (see Figure 1).

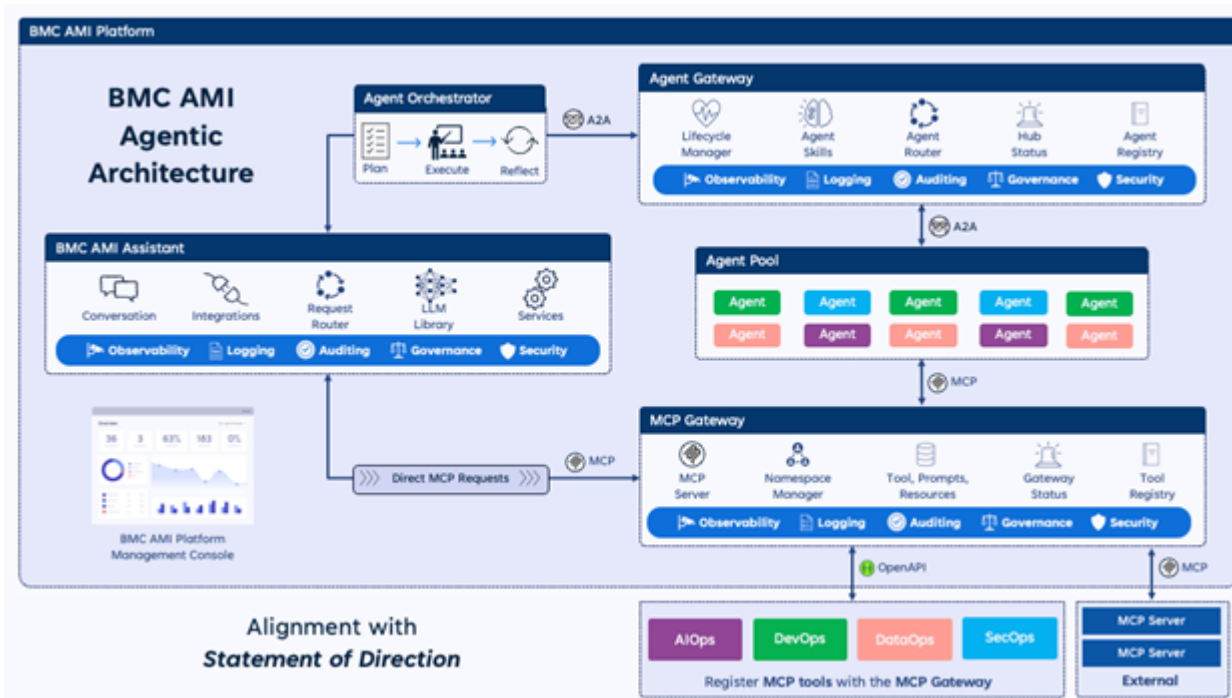


Figure 1: High-

level overview of BMC AMI Platform's agentic architecture supporting enterprise-grade AI.

## What you can expect — and how to shape what comes next

Over the coming releases, you will experience a meaningful shift in how AI participates in the BMC AMI environment. Intelligence will extend beyond explanation and guidance to support the safe execution of operational tasks. Capabilities that once appeared as isolated features will increasingly operate as coordinated agentic workflows across the BMC AMI portfolio, connecting development, operations, data, and security activities.

AI will participate responsibly within enterprise guardrails, operating under the policies, validation models, and controls defined by the teams who run these systems. This direction represents more than a single capability release, it establishes the foundation for a new way of working with BMC AMI solutions.

We invite customers to help us shape this next phase. You can engage with us in several ways: join our [Customer Design Partner program](#) to help refine and validate the most impactful agentic workflows, participate in early access initiatives focused on execution-oriented capabilities, and work with us to shape the governance and policy models that will guide enterprise-scale AI execution.

## Moving forward with confidence

We delivered on our commitments. We established a foundation of trust. And we are now stepping into the next phase with clarity.

The future of AI on the mainframe centers on purpose-built enterprise intelligence, where your choice of AI models, your institutional knowledge and operational practices, your people, and your chosen platforms work together to drive intelligent execution across the enterprise. In this model, the mainframe is not an isolated environment. It is a fully governed participant in enterprise AI strategies, capable of supporting intelligent execution across hybrid systems.

Organizations that embrace this direction will not simply modernize their mainframe environments. They will unlock new ways to operate complex workloads with greater visibility, control, and intelligence.

And this is only the beginning. Before you run AI on your most essential platform, BMC First.

Are you truly ready to move from AI insight to AI-driven execution on the mainframe?

Assess your organization's AI readiness and maturity level—register for an [AI Readiness Discovery Workshop](#) to take the next step.