

AGENTIFYING THE MAINFRAME



Throughout its history, the mainframe has adapted to every major shift in enterprise computing. The client-server model expanded how applications connect to the platform, and hybrid cloud computing connected it to the broader distributed world.

AI feels different.

Not because it changes how systems connect, but because it changes how people access knowledge, understand what is happening, make decisions, and increasingly, how work gets done.

The question I keep coming back to is not whether AI belongs on the mainframe. That seems inevitable. The more interesting question is what kind of AI organizations will ultimately trust with some of their most critical workloads.

This creates a new opportunity, one to **Agentify the Mainframe**, ensuring that AI on the platform is grounded in the expertise, context, and governance on which business depends. It means rethinking outdated perceptions and reimagining workflows so that mainframe teams can innovate more quickly while maintaining the resilience, security, and operational excellence that have always defined the platform.

In conversations with customers, I've found that the discussion quickly moves beyond AI capabilities. The real conversation is about confidence. Can AI understand operational context? Can it explain its reasoning? Can it act within clear boundaries? These questions have shaped our approach at BMC.

To achieve these goals, the BMC vision goes beyond general-purpose tools that simply access mainframe data. It makes AI *mainframe-native* with purpose-built AI that brings intelligence directly into the workflow, combining decades of BMC expertise, your organization's institutional knowledge, and live operational context with infused governance. The goal is not simply to bring AI to the

mainframe. It is to extend the value, expertise, and operational knowledge of the mainframe into the broader enterprise AI ecosystem.

In my time with BMC, I've seen firsthand the passion our customers have for the mainframe and the critical role it continues to play in their businesses. One thing they have consistently made clear is that the mainframe isn't going anywhere. They are not looking to replace it. They are looking to optimize it, extend its value, and ensure it evolves alongside the rest of the enterprise.

Agentifying the Mainframe is one way of supporting that evolution. It places the mainframe alongside enterprise AI initiatives while helping organizations scale hard-won expertise across more people and workflows. The result is the ability to onboard employees faster, make decisions more quickly, reduce dependency on a limited number of specialists, and create more capacity for innovation.

In short, **Agentifying the Mainframe** is about making mainframe expertise more accessible, scalable, and actionable through purpose-built digital workers, increasing confidence and capacity without compromising resilience, security, or control.

How do we make any AI mainframe-native?

If AI is going to earn trust on the mainframe, I believe a few characteristics become especially important.

The BMC vision of **Agentifying the Mainframe** is defined by six characteristics:

- **Mainframe-native expertise** that applies specialized mainframe knowledge to critical workflows and decisions.
- **Integrated intelligence** that connects expertise, institutional knowledge, and live operational context.
- **Enterprise openness** that securely integrates with enterprise tools, models, agents, and ecosystems.
- **Governed execution** that applies policy, oversight, transparency, and control to agentic actions.
- **Customer data** that keeps enterprise data governed by established security, access, and control policies.
- **Investment protection** that extends existing mainframe, BMC AMI, and enterprise AI investments rather than creating another isolated platform or collection of point solutions.

Each of these characteristics is valuable on its own. Together, they form the basis of BMC's approach to agentic AI on the mainframe: purpose-built, mainframe-native AI employing specialized agents that understand the context of the workflows in which they operate and coordinate with other agents and tools to extend human expertise by investigating, analyzing, explaining, recommending, and advancing defined work within governed boundaries.

One finding from the [2026 BMC Mainframe Survey](#) particularly stood out to me. As organizations become more familiar with AI, they are not necessarily becoming more comfortable handing over complete control. In several areas, preference is shifting toward AI that recommends actions rather than executes them autonomously.

When asked about AI problem diagnosis, willingness to have AI complete actions fell from 32% to 23% year-over-year, while preference for AI-generated recommendations increased from 36% to

41%.

I don't view this as resistance to AI.

I view it as a sign that organizations are thinking carefully about where trust belongs.

This suggests that many mainframe leaders see AI first as an advisor, and only later as an autonomous executor.

Agentifying the Mainframe addresses this need with governed boundaries and human-in-the-loop controls, creating a trusted advisor that assists mainframe professionals and enhances their work rather than replacing them. Organizations may eventually choose higher levels of automation, but trust must precede autonomy.

In my experience, organizations rarely become comfortable with automation all at once. Confidence is built gradually, one successful interaction at a time.

The appropriate level of autonomy will always depend on workflows, risk, and governance requirements. Organizations should be able to expand AI responsibility as trust grows and business needs evolve, while maintaining the oversight required for critical operations.

Governance by Design

The BMC approach includes governance embedded directly into the operating model itself, rather than bolted on later as an additional feature or service.

It is governance by design, practiced across four pillars: agent performance, resource utilization, policy compliance, and auditable activity tracking.

Here, too, a human-in-the-loop approach remains important. Human oversight is incorporated wherever intervention or approval is required. With this governance foundation, mainframe-native AI can move thoughtfully from intelligence to action.

Making It Real

Ideas only matter when they improve the way people work.

That is why we have focused on bringing these concepts into practical workflows today.

Agentifying the Mainframe is not simply a future vision. BMC is making it a reality through new BMC AMI Agentics offerings. These solutions use agents to investigate, analyze, explain, recommend, and help move defined work forward.

Rather than replacing operators, developers, or data professionals, these capabilities are designed to help them investigate faster, understand complex situations more clearly, and move work forward with greater confidence.

[BMC AMI Agentics for Ops](#) uses purpose-built agentic workflows that advance defined operational investigations, explain findings, and recommend next steps for proactive anomaly detection and AI-guided root-cause analysis, helping teams investigate faster and act with confidence without leaving their workflow. Agentic CPU investigation in [BMC AMI Ops Monitoring](#) increases the speed with which operations teams can investigate CPU spikes; incorporation of agentic AI into [BMC AMI Ops Insight](#) provides Db2 probable cause insights and recommends next steps; and a new explain agent

within [BMC AMI Ops Automation](#) explains REXX scripts in plain language to help operators understand existing automation scripts more quickly, especially when the original authors or subject matter experts are not available.

[BMC AMI Agentics for Data](#) uses AI agents to investigate database issues, explain findings, and guide optimization. AI Reorg Advisor adds a new agentic workflow to [BMC AMI Utilities for Db2®](#), which investigates REORG indicators and custom criteria, then explains recommendations in plain language. The SQL Workload Advisor explains application and SQL performance and recommends tuning actions.

AI-Assisted Application Development

Making AI mainframe-native also means rethinking how developers work.

With the new BMC AMI DevX Code Pipeline MCP server, now in beta, developers can connect familiar AI coding assistants such as Cursor, GitHub Copilot, and Claude Code to mainframe build and deployment processes.

Instead of treating the mainframe as a separate development environment, we are making it part of the AI-assisted workflows developers increasingly use every day.

The goal is simple: allow developers to use the tools they already know while making the mainframe a natural part of the modern development experience.

A Vision of Agentifying the Mainframe

System complexity continues to grow, and so does the pressure to innovate faster. As AI adoption accelerates across the enterprise, expectations are increasing while specialized mainframe expertise remains limited.

For me, that is the real opportunity behind **Agentifying the Mainframe**.

Combining mainframe-native expertise, integrated intelligence, enterprise openness, and governed execution can help organizations make AI more useful, more trusted, and more closely aligned with the realities of enterprise operations. The outcome is not simply greater automation, but greater capacity, resilience, and confidence.

As AI continues to evolve, I believe the most important question is not how quickly organizations can introduce intelligence into their systems.

It is how thoughtfully they do so.

The mainframe has earned trust over decades because of its reliability, resilience, and ability to support business processes that matter most.

As we enter the agentic era, those qualities become even more important.

In many ways, **Agentifying the Mainframe** is not about changing what makes the platform valuable.

It is about extending those strengths into a new generation of human and AI collaboration.

Learn more about BMC's vision for agentic AI on the mainframe on the [Agentic Mainframe page](#) of BMC.com.

See all of the new features included in the October release of BMC AMI enhancements on the [What's New in Mainframe Solutions page](#).

Read more about BMC's approach to agentic AI governance in [Delivering Predictable Outcomes](#) in the AI Era, by Chief Technology Officer Ram Chakravarti.