

# TAKING A RIDE ON THE MODERN MAINFRAME TRIP WITH YOUR AI BEST FRIEND



If you've been working in enterprise software most of your life like me, then artificial intelligence (AI) might seem like the next inevitable disruptive wave that is going to wash over our industry whether we want it or not, leaving good people's careers in its wake.

But just like factory automation hasn't replaced the need for skilled auto workers, and self-driving trucks still can't replace skilled drivers from first mile to last, we'll wind up in a situation where advanced systems will only increase the scale and complexity of IT work. Humans will need to learn new skills and stay in the loop, to keep applications as performant as possible ahead of increasing demand.

No aspect of enterprise software reflects this trend more than the mainframe modernization journey we're on right now. The mainframe is still the beating heart of the enterprise, responsible for our most critical business logic and transactional capabilities.

At the same time, we are experiencing a talent attrition crisis, as skilled mainframers—with their decades of experience and understanding—are retiring and moving on. Enterprises need to bring forward a new generation of mainframe talent and continue to innovate on the mainframe to meet ever-expanding business challenges.

In this arena, we're not looking for another disruption. As it turns out, AI might just be our best friend on this ride, so we don't have to go it alone.

Fortunately, BMC has been steadily and pragmatically working on AI-driven functionality across their Automated Mainframe Intelligence (AMI) portfolio to help ease the mainframe transformation

journey, culminating in the release of their [BMC AMI Assistant](#). Here are several new ways generative AI (GenAI) with specialized AI models could become our perfect traveling guide and companion.

## Bringing SME knowledge forward with a GenAI knowledge expert

Let's start by addressing the skills gap before we embark. As long-tenured SME subject matter experts (SMEs) are leaving the traveling party, we need to do everything we can to impart their institutional knowledge to newer engineers. And if GenAI systems were ideally suited to do one thing well, it's documentation and knowledge transfer.

However, a large language model (LLM) is only as good as the data that feeds it. To avoid irrelevance, we'll need much more than another chatbot that provides manually canned answers or information scraped from the internet.

While much of the existing codebase may be poorly commented and not clearly mapped out, BMC AMI Assistant can trace work on mainframe modules over time, documenting the accumulated knowledge gained through changes made by SMEs over the preceding years and decades, to help newer team members understand how to safely unlock procedural dependencies and connect new business services to the mainframe.

## Navigating difficult passages with AI-guided issue resolution

Monitoring and observability tools have their place within any enterprise software estate. To move forward with confidence, we need systems that can flag problems and alert teams to take action. Unfortunately, IT Ops teams and site reliability engineers (SREs) who are used to dealing with web-centric architectures will throw issues "over the wall" to the mainframe team for resolution if a problem appears to originate in a back-end system.

Mainframes can be opaque to traditional observability and remediation tools, so it's best to start with best-of-breed solutions that can glean telemetry signals from IBM Z<sup>®</sup> architectures, IBM<sup>®</sup> CICS<sup>®</sup> regions, IBM<sup>®</sup> Db2<sup>®</sup> databases and so on, such as [BMC AMI Ops Monitoring](#). But simply detecting a system issue isn't enough if there's not enough context in cryptic failure codes and abends for mainframers to reach remediation.

Here's where AI agents really shine. BMC AMI Assistant agents act like private investigators, marrying mainframe telemetry data with root cause analysis workflows, providing mainframers with GenAI-guided issue resolution help—explaining complex interactions between services, and zeroing in the exact locations and root causes of errors and failures by interacting with [BMC AMI Ops Insight](#), which provides telemetry and rules-based logic to contribute to machine learning (ML).

### **This "Hybrid AI" approach is core to BMC's strategy for simplifying mainframe transformation:**

The GenAI of BMC AMI Assistant is infused into BMC AMI Ops Insight, which provides rules-based logic and machine learning as it observes telemetry data with a deep understanding of mainframe operations and incident root causes. Then, the LLM in BMC AMI Assistant communicates findings to teams, parsing that telemetry into natural language explanations of root causes and step-by-step recommendations of the next best actions to take to achieve resolutions.

Empowered with expert AI guidance and next-step instructions, mainframe teams can reduce mean time to detect (MTTD) and mean time to resolve (MTTR) for both minor and severe incidents, as well

as kicking back issues that originated in front-ends, API services and networks that have nothing to do with the mainframe—i.e. “reframing the mainframe blame game.”

## **Pick the right LLM for the trip, or bring your own**

As AI makes its way from the top of the hype cycle to genuine productivity, there will never be “one model to rule them all.” That is why I always advocate for true composite AI approaches that use the right AI model for the job, rather than expecting an off-the-shelf LLM trained on generalized data to fulfill enterprise needs.

In this sense, BMC AMI Assistant acts like a concierge for selecting from multiple flavors of GenAI LLMs and specialized language models (SLMs), and orchestrating AI workloads across them, based on the particular use case or mainframe work to be done.

End customers can choose from a pre-curated set of open source LLMs and SLMs, each vetted as fit-for-purpose for specific work. For instance, you might check out Mixtral to explain PL/I and COBOL code, Granite for Assembler and JCL, and a Llama 3 running on GPU for operational insights.

Or better yet, bring your own LLM that is tuned to your own business needs and policies—an especially valuable capability for secure or proprietary mainframe environments where internal data security and sovereignty are important. There's no longer any reason to compromise or lock in one particular GenAI approach.

## **The Intellyx Take**

There's no sense in going it alone on the mainframe journey anymore when AI can be your guide.

The modern mainframe really is one of the bright spots where we can see Hybrid AI improving team collaboration and productivity, while future-proofing the mainframe's critical functionality against some of the AI-generated risk we see emerging in other areas of enterprise software.

To get there, we need the flexibility of GenAI assistants that can understand the context of your own internal business workflows and knowledge, paired with a “never fail” approach that we have always expected from our core transactional system investments.

*©2025 Intellyx B.V. Intellyx is editorially responsible for this document. At the time of writing, BMC is an Intellyx subscriber. None of the other organizations mentioned here are Intellyx customers. No AI bots were used to write this content.*