HOW CONVERSATIONAL AND CONTEXTUAL AI POWER MAINFRAME TRANSFORMATION



The hum of critical systems never sleeps. Somewhere, a mainframe operator stares at a spike in CPU usage that could ripple across millions of transactions by morning. A developer scrolls through lines of JCL, trying to figure out why a batch job failed and what to fix before business hours. They know the stakes, but the answers they need aren't in front of them. They're buried in decades of code, lost in static documentation, or locked in the mind of a senior expert who's already retired.

This is the reality for too many mainframe teams today. Not because the platform is outdated, far from it. The mainframe remains one of the world's most powerful, secure, and resilient engines for mission-critical workloads. But the weight of complexity, institutional knowledge gaps, and disconnected tools can slow down even the most modern environments.

What if the mainframe could talk back? What if it could surface the answers you need, in context, in the moment, not as static text, but as a dynamic, clear, conversational guide that helps teams solve problems, optimize performance, and modernize safely? This isn't a distant dream. It's happening now.

A Challenge hiding in plain sight

The pressure on mainframe teams has never been higher. Organizations are pushing for 24/7 availability, faster innovation, and tighter integration with modern cloud and hybrid environments. Regulations get stricter, workloads get heavier, and new demands show up overnight. Meanwhile, the people who know the mainframe best—the experts who built, tuned, and optimized these environments for decades—are retiring, and they're taking a vast store of institutional knowledge

with them.

This is where the cracks begin to show. For years, enterprises relied on tight-knit groups of specialists who carried decades of context in their heads—how a particular job control language script works, which performance quirks to watch out for, what business rules live inside thousands of lines of COBOL. As those stewards move on, newer developers and operators step up, only to discover that the roadmaps they need are scattered across aging manuals, static PDFs, tribal wikis, and half-remembered hallway conversations.

Take the story of a global financial institution that faced overnight escalations due to batch processing issues hidden deep inside critical JCL scripts. Without clear explanations of how those scripts actually worked, every incident meant hours of detective work—hunting through old configurations, cross-checking log files, hoping someone on the team might recall why a step was set up the way it was. When they brought contextual AI into the picture, those marathon troubleshooting sessions shrank to minutes. By recognizing historical job patterns and proactively suggesting corrective actions, AI turned guesswork into guided resolution.

This is the real pain point: not just missing documentation, but missing understanding. When teams lack confidence in what the code does—and what could happen if they change it—optimization stalls. Innovation takes a back seat to caution. The mainframe keeps running, but the people running it spend too much energy working around uncertainty instead of pushing forward.

When tools and context don't connect

Of course, lack of institutional knowledge isn't the only challenge. Many modern mainframe shops have invested heavily in tools like dashboards, consoles, APM solutions, log analyzers, ticketing systems. Each does its job well. But the human glue holding them together is under strain.

Let's go back to the operator staring at that CPU spike at 2 a.m. The alert comes in, but the context doesn't. They jump between screens: logs in one window, a dashboard in another, a knowledge article in a third, maybe a colleague's archived notes in a fourth. Every new tab means friction. Every switch means time lost. And when systems are mission-critical, time is the one thing no one can waste.

Contextual artificial intelligence (AI) changes that. Instead of static dashboards and manual hunting, teams can get explanations and next best actions right where they're needed. No more toggling back and forth, no more pasting error strings into different consoles. The answers surface inside the tools teams already rely on — with enough context to make sense of what's happening and what to do next.

This shift doesn't just save time. It changes how people work. It lets operators focus on resolution instead of research. It frees developers to act instead of second-guessing. And it keeps the mainframe running at peak performance without making human teams perform contortionist routines to get the insight they need.

The hidden weight of process debt

For some teams, the hesitation to change is not about what the code does, but what might break if they touch it. The mainframe itself is modern and robust. Its core languages—COBOL, PL/I, assembler—are proven and still remarkably fit for purpose when used correctly. What drags teams

down is the accumulated tangle of undocumented tweaks, one-off customizations, brittle integrations, and long-forgotten workarounds that create invisible risk.

Too often, these hidden layers of process debt create a climate of fear. A developer tasked with refactoring a section of high-impact workload knows the potential upside: faster performance, smoother integrations, lower resource costs. But what if a small change triggers an unexpected dependency? What if a workaround from ten years ago quietly fails and no one notices until something mission-critical grinds to a halt?

This fear of breakage is paralyzing. Work that should move forward stays parked. Opportunities to optimize get deferred, year after year, because teams can't see what's safe to change.

Conversational and contextual AI offer a new path. When AI can explain unfamiliar code, highlight dependencies, and even recommend safe next steps, teams gain the confidence they need to modernize without fear. Clarity replaces uncertainty. Smart suggestions replace second-guessing. And the mainframe stays robust, not just because of its engineering, but because the people running it know exactly what they're doing and why.

Breaking free from the static knowledge trap

Maybe the deepest frustration for many teams is that so much critical knowledge exists yet stays locked away. Every day, mainframe developers and operators become digital archaeologists. They sift through decades-old manuals, dense PDFs, cryptic support tickets, and scattered SME notes just to answer a single, urgent question.

Finding the right piece of insight often feels like an excavation project. Is that config tweak documented somewhere? Did someone explain this logic five years ago in a PDF buried in an internal wiki? Was the answer buried in a support ticket two managers ago?

Conversational and contextual AI rewrite this script. Instead of static repositories, knowledge becomes dynamic. AI understands context, intent, and nuance. Ask a question in plain language, like "Why is this job spiking CPU?" or "What does this piece of code actually do?" and get an answer that's clear, relevant, and tailored to the real problem at hand.

This is the difference between a dusty archive and an intelligent guide. Teams stop digging and start doing.

The bigger vision: A conversational mainframe

This is the evolution: a mainframe that doesn't just power mission-critical workloads, but explains itself, protects itself, and guides your teams forward. Decades of institutional knowledge, available at your fingertips. Fewer barriers between people and the answers they need. More confidence to modernize safely without slowing down what works. And with conversational and contextual AI woven in, organizations can also strengthen governance and security by ensuring the right answers reach the right people, policies stay consistent, and sensitive operational knowledge is safeguarded while still being accessible to those who need it most.

A vision that's already here

This isn't just theoretical; this vision is already taking shape inside BMC AMI Assistant. With new

capabilities like chat with a mainframe knowledge expert, teams can ask questions in natural language and get immediate, contextual answers, whether they're trying to understand complex code, resolve performance issues, or plan safe changes.

Powering this is the mainframe knowledge hub, the intelligence layer within BMC AMI Assistant that continuously ingests, organizes, and activates institutional knowledge from manuals and support tickets to internal best practices, all without tedious manual tagging.

Together, these capabilities help turn decades of static, scattered expertise into living knowledge that teams can tap into anytime, right where they work.

The path forward: A future of connected intelligence

The journey toward an AI-powered mainframe isn't about adding flashy tools for the sake of it. It's about rethinking how teams work, how knowledge flows, and how systems and people collaborate.

This transformation requires more than just smart algorithms. It demands platforms built for seamless interoperability. Frameworks that standardize intelligence exchange. Autonomous agents ready to act on insight, not just surface it.

As we expand what's possible with conversational and contextual AI through BMC AMI Assistant, we're not just enhancing technology, we're fundamentally changing how knowledge is captured, connected, and leveraged. The mainframe doesn't just process transactions, it becomes an active partner in resilience and innovation.

The question isn't whether conversational and contextual AI will shape the future of mainframes, it's how swiftly organizations will embrace this intelligent shift to build resilient, optimized, and future-ready environments.

What could your teams deliver if the mainframe didn't just run, but responded? The future is already talking back. The only thing left is to listen.

Learn more about how BMC AMI Assistant was recognized with the AI Breakthrough Award for Conversational AI Innovation.