HyperFRAME RESEARCH

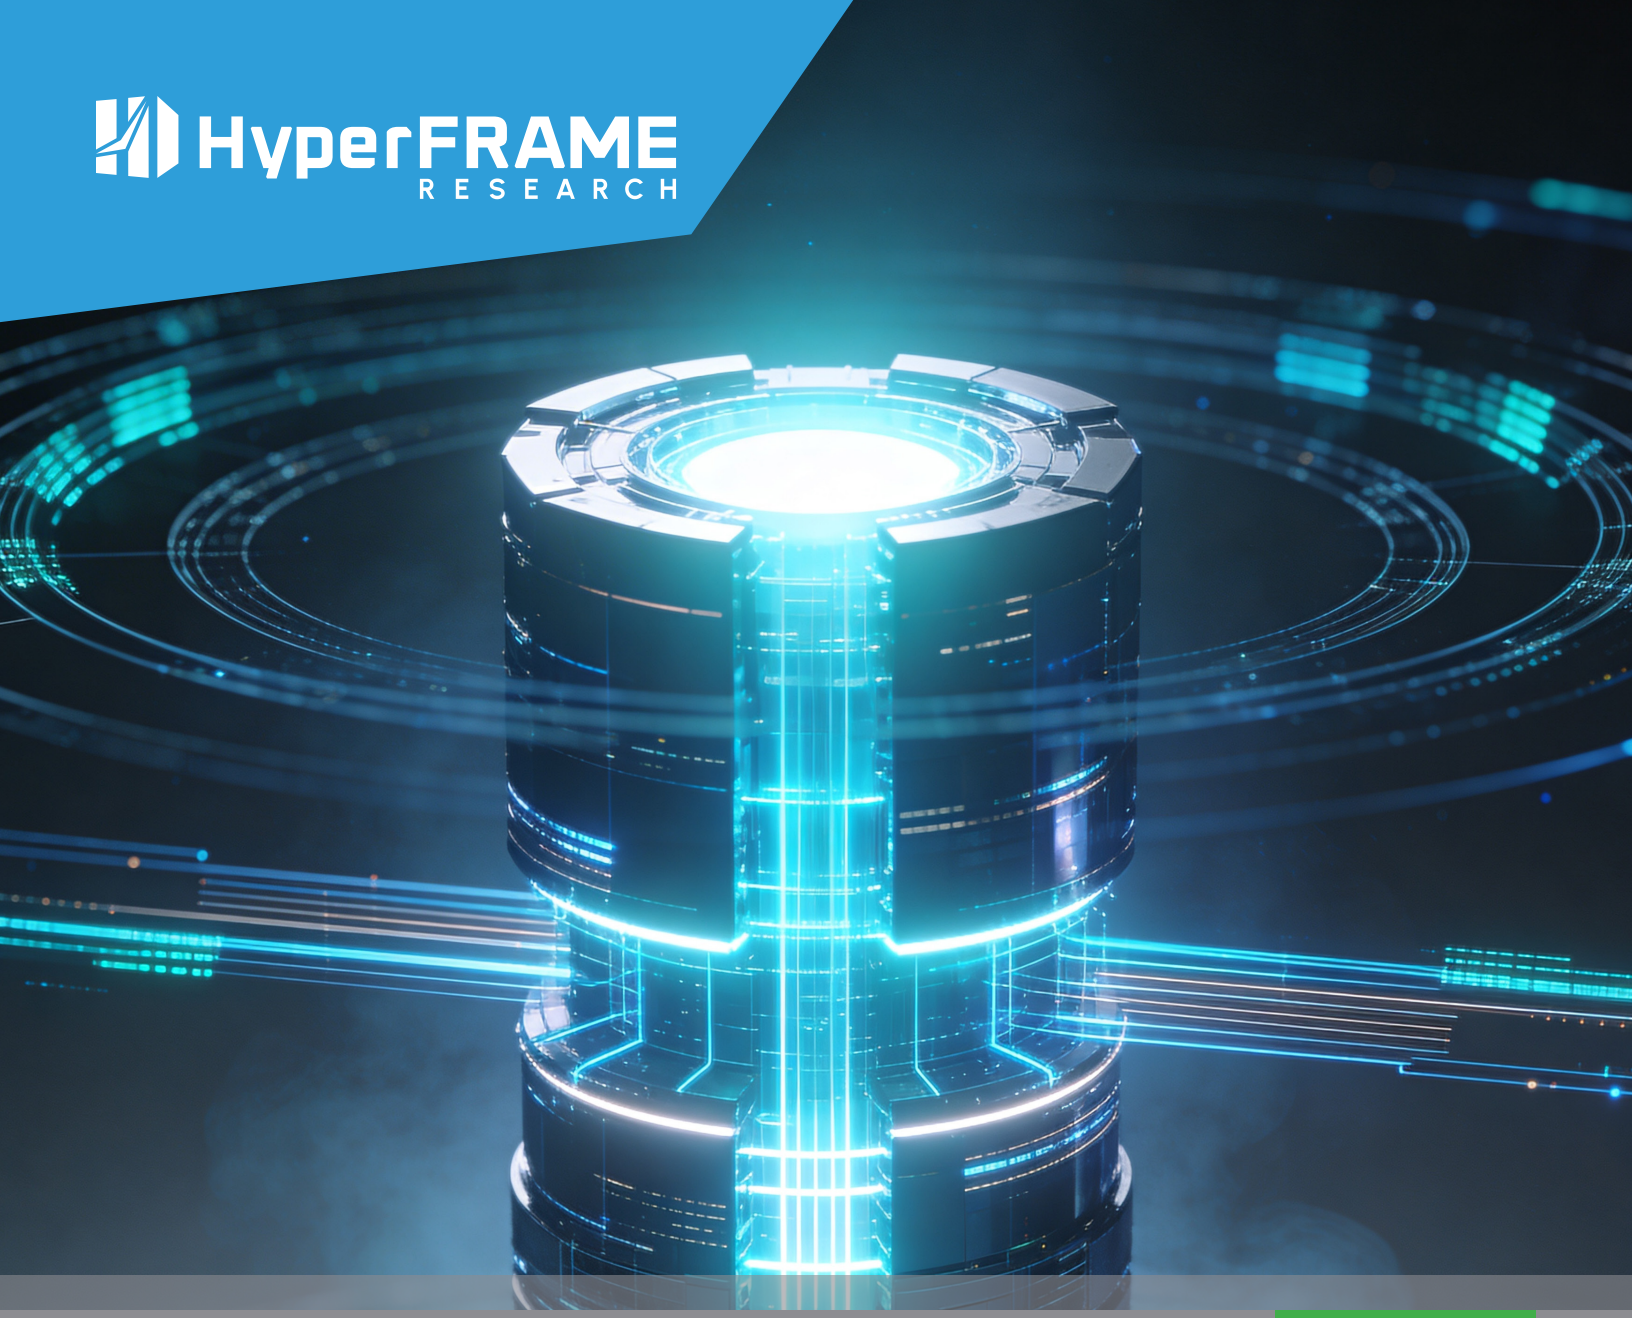

RESEARCH BRIEF

Beyond GenAI: The AI Readiness Journey to the Agentic Mainframe

Embracing AI Agents, Coordinated Workflows, and MCP Servers as the Natural Next Step after Generative AI

Authors:

Stephanie Walter
Practice Leader, AI Stack

Steven Dickens
CEO and Principal Analyst

FEBRUARY 2026

Executive Summary

The mainframe has already embraced Generative AI (GenAI), incorporating it primarily to boost productivity in documentation, code explanation, and knowledge retrieval. This initial phase, while valuable, represents only the first step of a deeper digital evolution. To unlock this next step, organizations must understand that agentic execution is only possible when AI readiness is firmly established. Readiness is a staged progression that includes trustworthy data foundations, well-defined governance, and cultural acceptance of increasingly automated workflows. Framing the mainframe transformation as a maturity journey helps leaders understand that autonomy is not a leap, but the logical outcome of sustained investment in these prerequisites.

The next frontier for competitive advantage on the mainframe does not lie in simply refining these tools, but in adopting AI Agents, Agentic Workflows, and communication protocols such as Model Context Protocol (MCP) and Agent2Agent (A2A). These technologies are not incremental improvements. They are fundamentally designed to redefine how mainframe teams interact with their complex environments. GenAI is primarily advisory today; agentic AI represents a future operational model as organizations mature their readiness. Agentic systems are being designed to autonomously perceive complex system state, reason through multi-step problems, make policy-aligned decisions, and execute actions across disparate domains: development, operations, and security. This transition transforms AI from a sophisticated helper into a reliable, scalable digital teammate.

As with any major shift in enterprise automation, organizations move through defined stages of AI readiness before they can reliably reach agentic operation. Each stage expands the level of trust, governance discipline, and autonomy the organization is prepared to support. This progression ensures that agent actions remain aligned with policy, grounded in accurate context, and fully auditable. Agentic operation is the destination, but it cannot be reached without maturing through these foundational stages.

Organizations that fail to move decisively from a GenAI advisory model to an agentic AI operational model will fall behind. The consequences include not only a loss of efficiency but also increased operational risk due to a dependence on human capital that cannot scale to meet system complexity. The adoption of coordinated, multi-agent workflows that are governed by an orchestration layer will be necessary to deliver the necessary resilience and innovation capacity. BMC AMI is well-positioned as the trusted leader in this next, critical phase of mainframe transformation.

The agentic capabilities described in this paper represent a directional future state that organizations progress toward over time, rather than functionality broadly available in production environments today. This paper outlines a directional view of the agentic mainframe and the readiness journey required to reach it. The architectures, workflows, and use cases described represent future-state scenarios that organizations progress toward over time, not a declaration of generally available autonomous capabilities today.

From Generative AI to Agentic AI

The evolution of AI in the enterprise moves fast. Simply put, Generative AI was the start; agentic AI represents an emerging next stage of maturity for organizations that move beyond advisory AI. This is not a slight on GenAI, which delivered significant, immediate productivity boosts in tasks related to knowledge and documentation. It allowed mainframe teams to rapidly explain decades-old COBOL code, draft test plans, and summarize voluminous system logs. GenAI provided a powerful assistant.

Early GenAI adoption plays a critical readiness role because it introduces teams to AI-assisted work in a safe, low-risk way, building the foundational comfort and trust needed for deeper automation. It also forces organizations to prepare the underlying data, workflows, and governance patterns that agentic systems require. Most enterprises today remain in the early readiness stages and largely operate in Stage 1 or Stage 2. This explains why their GenAI investments feel helpful but not yet transformative.

The crucial limitation of GenAI is not usefulness, but its lack of autonomous agency. It acts on a prompt, and its output, typically text or code, requires a human to validate it and then manually execute any subsequent actions. The process typically stops at the suggestion, requiring human validation and execution. AI Agents and Workflows extend beyond mere suggestions. They are architected to act, collaborate, and evolve with minimal continuous human direction. In a forward-looking agentic operating model, an LLM can explain a database performance problem, while an AI agent could be designed to detect the anomaly, consult the performance policy, coordinate with an operations agent to initiate a fix, and report the outcome, all without an operator manually moving between tools. This shift automates the entire workflow, not just a single step. The returns are compounding. Additionally, this shift is not just technological; it reflects the organization's movement along a defined readiness curve, where early GenAI adoption builds the confidence, data foundations, and governance discipline required for agentic execution.

To fully grasp the magnitude of this evolution, it helps to define the core components. The following concepts describe industry patterns and enabling frameworks, not specific product implementations or near-term commitments. They are included to establish a shared vocabulary for how agentic systems are being designed across the broader AI ecosystem:

- **AI Agents:** These are autonomous digital teammates capable of perceiving context, making goal-oriented decisions, and executing tasks across a variety of defined tools and APIs. They possess memory and can persist toward a long-term goal, unlike a single-query LLM. They are built to be proactive problem solvers.

- **Agentic Workflows:** These are designed to be coordinated, multi-agent processes that optimize complex, cross-domain tasks. An agentic workflow might involve a Development agent scaffolding a code change, a Security agent running a vulnerability scan against the new code, and an Operations agent deploying a canary release. This is a single outcome achieved via the coordinated effort of multiple specialized agents.
- **MCP:** The Model Context Protocol should be built into the communication layer. It provides a unified protocol so models and agents can connect to external data sources or tools. MCP establishes a consistent way to expose capabilities (e.g., databases, code repositories, applications) to LLMs or agents.
- **A2A:** The Agent2Agent protocol is used for agents to communicate and collaborate with each other across domains, frameworks, or vendors. It defines how agents share intent, delegate tasks, or negotiate actions. A2A is often layered on top of MCP, using MCP's context and tool access to carry out coordinated plans.

This is a natural progression, not a sudden leap. The groundwork laid by early GenAI adoption, like setting up Retrieval Augmented Generation (RAG) for internal knowledge bases, directly prepares the data and infrastructure for agentic execution. However, organizations that stop at GenAI miss the compounding returns that autonomous, coordinated action delivers.

Use Cases: Mainframe Optimization and Transformation

The following use cases illustrate future-state scenarios in a mature agentic operating model, enabled as organizations progress through the AI readiness stages described later in this paper. The most compelling case for agentic AI is its capacity to transform keeping the lights on into accelerating transformation. The mainframe's stability is its greatest asset, but its complexity is often its greatest drag on innovation. AI agents can be engineered to resolve this tension, delivering both optimization and transformation across core functions.

The operations domain is the natural starting point for mainframe agents because it is highly structured, rich with data, and demands high-velocity response.

- **Proactive Issue Detection and Automated Resolution:** Agents could be designed to monitor streams of mainframe telemetry data far beyond human capacity. An operations agent could detect a database performance anomaly, consult a runbook stored in a knowledge base, execute an approved remediation (like adjusting a parameter or reallocating a resource), and then log the entire sequence for audit. These patterns help strengthen the trust and governance practices required for broader operational integration, ensuring AI augments the system in predictable and policy-aligned ways.
- **Continuous Compliance:** A dedicated compliance agent could be built to continuously audit system configurations and data access logs against regulatory policy. If it detects drift from a baseline, it could flag the change and execute an approved policy-compliant rollback, ensuring the system maintains a secure and compliant state at all times.

In the application development lifecycle, AI agents can function as indispensable pair programmers and domain experts.

- **Code Scaffolding and Testing:** Agents could be designed to rapidly scaffold new COBOL code or REST APIs from a natural language requirement, generate comprehensive test cases to validate business logic, and create updated documentation. This accelerates the cycle time for new feature delivery.
- **Collaborative Development Workflows:** Development workflows could become genuinely collaborative, with human oversight. A developer submits a feature request; a dev agent generates the initial code, a test agent generates unit and integration tests, and a security agent reviews the resulting code for vulnerabilities, creating a secure, verified pipeline.

The demographics of the mainframe operator are a business-level threat. As senior talent retires, decades of tacit institutional memory, along with the unwritten rules of why the system works the way it does, walks out the door.

- **Encoding Institutional Memory:** AI agents could be engineered to preserve operational expertise by monitoring and recording how human experts resolve complex issues. Over time, these agents can capture, codify, and consistently surface proven resolution paths. By capturing these unwritten procedures and transforming them into structured decision logic or

automated playbooks, AI agents would help mitigate the loss of institutional knowledge and ensure consistent, policy-aligned responses as experienced talent retires.

- **Innovation Enablement:** By delegating routine maintenance, first-level diagnostics, and compliance checks to AI agents, human talent is freed from the tyranny of keeping the lights on. This allows the finite pool of highly skilled professionals to focus their efforts entirely on new applications, digital transformation projects, and the strategic evolution of the mainframe platform.

This shift moves the mainframe from a cost center burdened by technical debt to an engine of enterprise agility.

Evidence of Momentum: Mainframe AI Adoption Trends

The enterprise is already moving toward this reality. The shift from AI as a passive advisor to AI as an active operator is not a projection; it is an observable trend rooted in growing cultural acceptance within the technology workforce.

The 2025 BMC Mainframe Survey reflects a clear movement of comfort and trust among mainframe professionals regarding AI's role. This shift reflects the broader industry trend going from isolated pilots toward embedded workflows, where AI becomes an integrated component of daily operations rather than a standalone assistant. Mainframe professionals are increasingly comfortable moving beyond simple observation to direct action.

Mainframe professionals are expressing greater comfort with:

1. **Receiving alerts** from an AI system about potential issues. This is the foundational trust layer.
2. **Receiving recommended actions** and moving the AI into the advisory role.
3. **Allowing AI to complete simple, approved actions** under human supervision. This is the first step of autonomous operation.
4. Eventually **allowing AI to recognize, respond, and remediate** autonomously within defined guardrails.

These stages reflect growing comfort and organizational intent, not widespread deployment of fully autonomous systems today. This progression shows that cultural acceptance is rapidly growing. The perceived risk of AI taking action is being consistently offset by the demonstrable gain in speed and accuracy. The organization is moving from AI as advisor (simple alerts), to AI as partner (validated recommendations), to AI as operator (autonomous execution). The success of GenAI in providing context has paved the way for the trust needed for agentic execution.

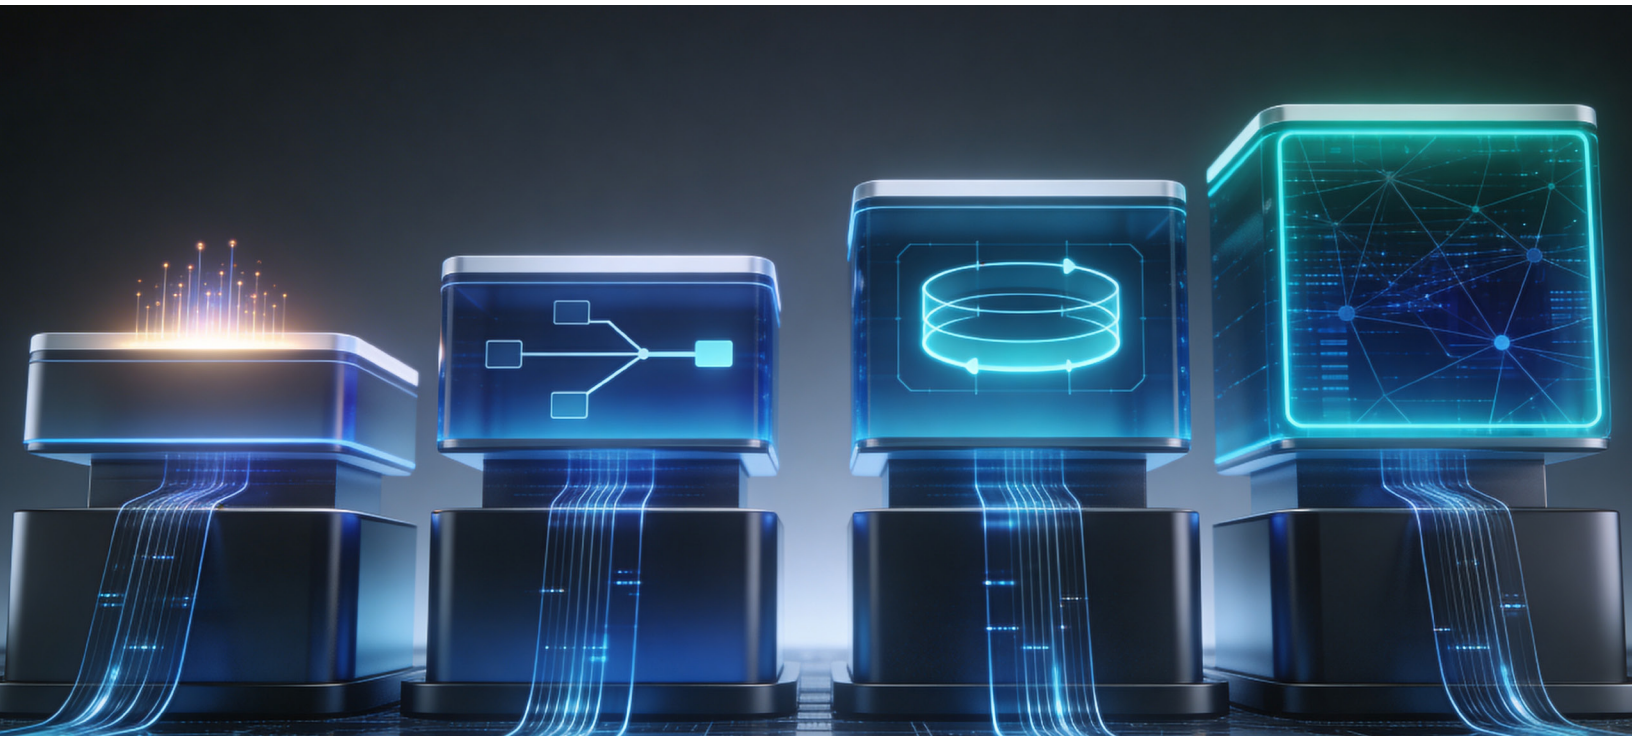

Risks of Inaction: The Case Against Doing Nothing

There may be a temptation to stop at Generative AI. A leader might argue, "We have a GenAI tool that explains our code, and our team is already stretched. Why introduce more change?" This is a short-sighted view. The risk of adopting agentic AI too early is far outweighed by the competitive, operational, and innovation risks of waiting too long. Organizations in the earliest phases of AI readiness are the most exposed, because they have not yet built the governance, workflow, and trust foundations required to support safe automation.

Organizations that stop at GenAI will be outpaced by rivals who scale their operations with agentic AI. Over the years, competitors that embed autonomous agents into their dev, ops, and security pipelines could achieve cycle times and fault tolerance that are simply unattainable by human-only teams. They will deploy faster, maintain a higher state of compliance, and dedicate more resources to market-differentiating features. This leads to a fundamental shift in business performance. This risk is most acute for enterprises still operating in the advisory-only maturity band, where AI remains disconnected from execution and cannot meaningfully scale operational capacity. The near-term imperative is not autonomy itself, but building the readiness required to safely adopt it.

The mainframe talent shortage is not a cyclical problem; it is a structural one. Talent shortages mean teams cannot scale their operational capacity without digital teammates. Every complex incident, every security audit, and every compliance check consumes finite, expensive human hours. Without AI agents to handle the high volume of routine and semi-complex work, teams could experience chronic overload, leading to burnout, errors, and an increasing failure to meet service level agreements (SLAs). Agents are not merely a cost-saving measure; they are protection against talent scarcity. Organizations that have not progressed beyond the early readiness phases lack the governance and orchestration mechanisms needed to delegate routine work to agents, increasing their dependence on scarce human expertise.

Without agentic workflows, modernization and innovation efforts can stall. The most precious resources, the senior architects and developers, remain tied up in triage, maintenance, and firefighting. Agentic AI is designed to abstract away the care and feeding of the running system, liberating these experts to focus on integrating the mainframe into the modern enterprise cloud and developing new digital applications. The time to build this foundational capability is now.

The risk is not adopting a new technology too early and having to iterate. The true risk is waiting too long and allowing the growing gap between system complexity and human capacity to become insurmountable. Reaching agentic operation is not a sudden leap; it is a structured progression shaped by an organization's readiness across data, governance, culture, and workflow integration. Each phase builds on the confidence and competencies established in the one before it, ensuring that autonomy emerges safely and predictably rather than abruptly. This progression reflects how enterprises evolve from early experimentation to integrated adoption and, ultimately, governed agentic operation.

Stages of AI Readiness

Before an organization can successfully implement agentic AI, it must conduct an honest, comprehensive assessment of its current level of AI readiness. This assessment must go beyond technology to include culture, data quality, and governance frameworks. Most organizations will progress through these stages over multiple years, and not all will reach full agentic collaboration at the same pace. Today, the emphasis is on readiness, governance, and guided integration rather than immediate autonomy. Maturity is a journey, not a switch.

Stage 1: Reactive Experimentation

- **Description:** AI use is isolated and sporadic, primarily for single-turn tasks.
- **Typical Use:** GenAI is used for isolated tasks such as internal documentation, chat-based assistance for code explanation, or generating initial drafts of non-critical communications.
- **Integration:** There is virtually no integration with core mainframe workflows, data systems, or execution engines. AI is an isolated knowledge tool.
- **Focus:** Learning what's possible and building initial developer familiarity with LLMs.

Stage 2: Guided Enablement

- **Description:** AI begins to support specific functions, but a human remains the definitive decision-maker.
- **Typical Use:** AI supports specific functions like code analysis, sophisticated anomaly detection, or generating validated test scenarios. The AI provides a validated recommendation, but a human must manually approve and execute.
- **Integration:** Limited, read-only integration with mainframe monitoring systems and knowledge repositories. Data is fed to AI, but AI does not write back to the system.
- **Focus:** Building organizational trust, developing initial governance frameworks, and validating the accuracy of AI outputs.

Stage 3: Operational Integration

- **Description:** AI agents are embedded in daily, high-value workflows with clear human oversight checkpoints. This is the first stage of agentic maturity.
- **Typical Use:** Operations agents handle tier-one problem triage and automated resolution for well-defined, low-risk incidents. Development agents autonomously manage the initial phases of the continuous integration/continuous delivery (CI/CD) pipeline.
- **Integration:** Orchestration platforms are introduced to manage security, versioning, access control, and collaboration between the agents and the mainframe tools. Communication protocols such as MCP and A2A would be in place.
- **Focus:** Efficiency, transparency, and achieving measurable return on investment (ROI) in specific functional domains.

Stage 4: Autonomous Collaboration

- **Description:** Multiple, specialized AI agents coordinate across domains to deliver complex, cross-system outcomes.
- **Typical Use:** A security agent, an operations agent, and a development agent coordinate to patch a zero-day vulnerability: detection, code fix, automated deployment, and verification all happen autonomously with high-level human monitoring. The organization operates with hybrid human-agent teams.
- **Integration:** Full, two-way integration across the technology stack, with cross-domain orchestration.
- **Focus:** Agility, resilience, and accelerating strategic business innovation.

Maturity is not about achieving full, hands-off automation. It is about alignment: ensuring that humans and agents share context, adhere to the same policies, and work toward unified business goals.

Recommendations

The path to agentic AI is strategic and requires focused investment. The following is a four-part strategy to prepare the organization for this next phase of the mainframe transformation.

1. Develop a Strategic AI Roadmap for the Mainframe

Don't treat agentic AI as a side project. It must be a core component of your broader platform strategy. Design the roadmap to clearly define the progression from the current state

to the target state. This roadmap must include a technology refresh plan to ensure that the core mainframe systems are modern, API-accessible, and ready to be orchestrated by agents.

2. Refine High-Value Use Cases and Key Requirements

Prioritize agent use cases based on risk, complexity, and potential return.

- **Coding Use Cases:** Focus on use cases that address the most immediate bottlenecks, such as generating test cases for legacy code, automatically documenting undocumented business logic, and scaffolding new API endpoints. This accelerates modernization.
- **Operational Use Cases:** Focus on use cases that deliver the highest operational stability and reduce human toil, such as automated log analysis and root cause identification, and proactive compliance checks with autonomous, policy-aligned remediation.

3. Establish a Robust Data and Governance Framework

The effectiveness of any agent is tied to the quality of its knowledge and the rigor of its governance.

- **Data Flows and Sources:** Develop robust data flows and sources for any AI projects. Agents must be trained and grounded on accurate, real-time, enterprise-specific knowledge, not generic internet data. This requires consolidating organizational sprawl, ensuring trusted, federated knowledge sources with consistent contextual access for AI systems.
- **Governing Agent Action:** The governance framework must define the human-in-the-loop processes, outlining when an agent must seek human approval versus when it can act autonomously.

4. Create an Evaluative Framework for Mainframe AI Technology

Not all AI is created equal, especially on the mainframe. Develop a clear evaluative framework for technology from key vendors. This framework should assess:

- **Domain Expertise:** Does the vendor's solution understand mainframe-specific languages (COBOL, PL/I, Assembler) and environments (CICS, DB2, z/OS) natively?
- **Orchestration and Control:** Does the solution include a robust, auditable orchestration layer to manage the agents?
- **Trust and Security:** Is the platform architected to maintain data security, access control, and an end-to-end audit trail of all agent actions?

BMC AMI: The Trusted Leader in Mainframe Transformation

The transition from AI as an advisor to AI as an operator demands a partner with proven, deep expertise in the unique challenges of the mainframe environment. BMC AMI (Automated Mainframe Intelligence) has pioneered AI-assisted automation, predictive analytics, and operational observability across mainframe environments for more than a decade, long before the emergence of generative AI. This long-standing foundation includes rules-based automation, anomaly detection, performance analytics, and policy-driven workflows that have been trusted in the world's most mission-critical systems. That operational heritage gives BMC AMI an intimate understanding of the operating system, its subsystems, and their performance characteristics, providing a credible foundation for introducing generative and hybrid AI capabilities safely and responsibly at enterprise scale.

BMC's role is not limited to supporting organizations once they arrive at autonomous operation; it is designed to guide them through each stage of AI readiness. The AMI platform strengthens governance discipline, establishes trusted data and execution pathways, and introduces guardrails that support a phased progression from AI-assisted operations toward more coordinated, policy-driven automation over time. This positions BMC not as a provider of instant autonomy, but as an enabler that ensures the journey unfolds safely, predictably, and in alignment with enterprise policy.

BMC AMI Assistant is built to strengthen AI readiness by embedding generative and hybrid AI capabilities directly into core mainframe development, operations, and optimization workflows. Today, the platform delivers eight concrete capabilities, including natural-language code explanation, AI-assisted documentation and knowledge retrieval, Knowledge Expert Chat for surfacing institutional knowledge, GenAI-driven COBOL-to-Java conversion, AI-guided root cause explanation, recommended next-best actions, curated AI model selection through an LLM Library, and Bring Your Own LLM (BYOLLM) support to align AI usage with enterprise standards.

Together, these capabilities provide contextual insight, guided decision support, and governed execution within existing workflows, without removing human oversight. Rather than introducing autonomous agents prematurely, BMC AMI Assistant focuses on embedding intelligence, transparency, and control into everyday work, establishing the governance, trust, and operational discipline required to progress safely along

the AI readiness journey. The strategy is deliberate evolution and enhances how teams work today while preparing the foundation for more advanced coordination models over time. At the same time, BMC is rapidly advancing toward AI Agents, agentic workflows, and MCP Server Gateway capabilities, ensuring customers can adopt these innovations when their readiness, governance, and operational maturity align.

In environments handling the world's most critical transactions, trust is paramount. With deep integration into existing mainframe systems, BMC AMI ensures every AI action is explainable, auditable, and fully compliant with stringent security and operational policies. This trust layer is essential for any progression beyond advisory AI, including future coordinated digital operator models. This ensures agents act within the established boundaries of the enterprise.

BMC's sustained investments in agentic orchestration, MCP integration capabilities, and a growing partner ecosystem reflect a clear commitment to where mainframe AI is headed, while grounding that evolution in proven operational capabilities available today. BMC supports each phase of this evolution, helping organizations move from exploratory AI use to fully governed, workflow-integrated automation that prepares the environment for agent-based operation. Through its leadership in mainframe modernization, BMC AMI enables customers to achieve measurable gains in performance, availability, and innovation today. The platform is architected to ensure that AI on the mainframe remains accountable, transparent, and seamlessly aligned with core business goals. In this way, BMC AMI becomes the connective tissue of the readiness journey, ensuring each phase strengthens trust, improves control, and prepares the enterprise for safe autonomy. Importantly, this journey is incremental: enterprises adopt these capabilities progressively, aligning operational readiness, governance maturity, and platform evolution over time.

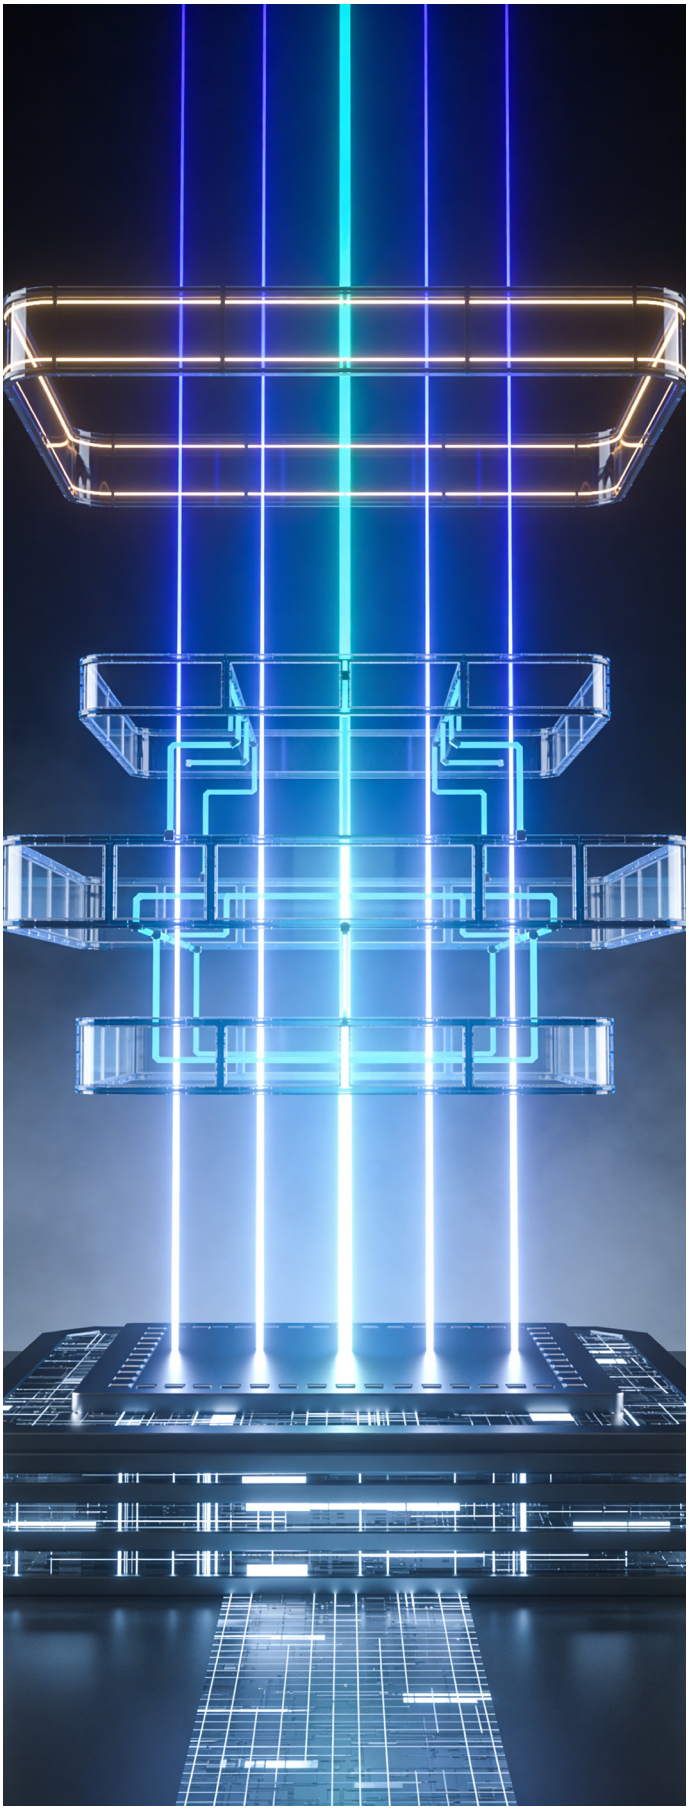

Looking Ahead

The mainframe is not merely surviving the digital era; it is thriving in the agentic era. The inherent reliability, security, and transactional integrity of the platform are a perfect match for the autonomous, high-velocity demands of AI agents. The current moment is not about choosing a new technology, but about securing the future operational model of the enterprise.

Acting now secures resilience, ensures knowledge continuity, and dramatically increases innovation capacity. By implementing agentic AI, organizations are essentially purchasing scalability for their most experienced teams. The best outcome is a harmonious blend where human expertise is augmented by tireless, autonomous digital teammates.

AI agents can be significant allies. They are the component that bridges the gap between the speed of modern business and the complexity of the core system. This evolution of the mainframe team, from advisory support to autonomous operation, will define competitive leadership for the next decade.

Mainframe teams who partner with BMC to embrace agentic AI, leveraging robust platforms like BMC AMI, will not only keep pace with the market. They will set the standard for what a resilient, agile, and secure enterprise platform looks like. The future of mainframe excellence is progressively autonomous, grounded in policy-governed, context-aware AI systems. The organizations that lead in this era will be those who treat autonomy as the culmination of a disciplined AI readiness journey.

ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers in-depth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programs.

Our industry analysts specialize in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

CONTACT HYPERFRAME RESEARCH:

Steven Dickens
CEO & Principal Analyst | HyperFRAME Research

Email Address:

steven.dickens@hyperframeresearch.com

Telephone Number:

+1 845 505 1678

X: @StevenDickens3

LinkedIn: Steven Dickens

BlueSky: Steven Dickens

CONTRIBUTORS

Stephanie Walter
Practice Leader, AI Stack

Steven Dickens
CEO and Principal Analyst

INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

LICENSING

DISCLOSURES

HyperFRAME Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.


HyperFRAME
RESEARCH