

AUGUST 2025

Architecture: The Secret to Consistent Success in Agentic AI

Torsten Volk, Principal Analyst

Executive Summary: Unlocking Trustworthy Agentic AI

Agentic AI is more than a technology trend. It comes with the potential to fundamentally transform business processes, enhance decision-making, and unleash innovation potential across all parts of the organization. Therefore, frameworks that enable organizations to successfully implement agentic AI-driven capabilities in a consistently trustworthy and successful manner will ultimately gain significant advantages in the marketplace.

As a result, enterprises have software vendors banging down their doors to talk about agentic AI, AI factories, Model Context Protocol (MCP), and other key technologies that help customers harness the business value of AI. While the promise of agentic AI is real and unparalleled in its potential business impact, at the same time, it comes with several critical risk factors in the areas of compliance, cost, reliability, security, and accuracy. These risk factors are amplified by the fact that most organizations lack the experience and staff skills to successfully plan, architect, implement, and operate AI capabilities that integrate with their existing technology infrastructure without adding significant operational risk and uncertainty.

Red Hat OpenShift AI aims to address these challenges through its extensive approach to enterprise AI workloads. As a key contributor to open source AI projects including vLLM, llm-d and llama-stack, Red Hat focuses on delivering a combination of innovation and enterprise stability. The platform seeks to deliver fast, flexible, and efficient inferencing across diverse hardware accelerators, AI models, and deployment targets, from on-premises infrastructure to all major public clouds. This flexibility is essential as enterprises navigate the inevitable diversity of AI workloads while maintaining consistent operations across their hybrid environments.

This paper will explore how organizations can use their current OpenShift footprint in combination with Red Hat and KPMG LLP's Agentic AI Architecture framework to benefit from the advantages of agentic AI in a well-governed, secure, and cost-effective manner.


Challenges and Solutions on the Way to Consistently Successful Agentic AI Implementation at Scale

According to research from Enterprise Strategy Group, the top three goals for enterprise-ready AI are strong governance, accurate results, and easy integration with existing systems.¹ Achieving these goals consistently across the entire organization requires a disciplined approach that combines strategic planning, cross-functional collaboration, and iterative implementation.

¹ Source: Enterprise Strategy Group Research Report, [Navigating Build-versus-buy Dynamics for Enterprise-ready AI](#), January 2025, All Enterprise Strategy Group research references and charts in this Showcase are from this report.

Figure 1. Key Factors Delaying the Value of AI Solutions

Which of the following factors do you feel most significantly delayed your organization from seeing value from AI in the time you indicated (please rank the top 3 factors from 1 to 3, with 1 being most impactful)? (Percent of respondents, percent ranked #1 displayed)



Source: Enterprise Strategy Group, now part of Omdia

Red Hat-KPMG Agentic AI Reference Architecture

Red Hat and KPMG use a clear, structured framework to support AI implementation and overcome common barriers to achieving value. This framework includes standardized reference architectures, strong governance practices, and secure, scalable technology foundations. It includes Red Hat's OpenShift AI platform as the technical infrastructure, orchestration, and deployment capabilities required for the deployment and operations of agentic AI. KPMG LLP's Agentic AI Reference architecture categorizes and defines clear roles and responsibilities of AI agents, enabling organizations to implement agentic AI more effectively. In addition to KPMG LLP's Agentic AI Reference Architecture, KPMG LLP's Trusted AI Framework brings clearly defined governance processes, risk management controls, security measures, compliance standards, and business alignment practices to the table.

Addressing the Critical Challenges of Agentic AI Success

The joint approach aims to help organizations proactively prevent common implementation pitfalls while also addressing reactive issues, such as unclear objectives, skill gaps, integration challenges, data quality problems, and security concerns, and aiming to ensure the safe, secure, reliable, and complaint rollout of agentic AI capabilities across the enterprise.

Integration Automation

Organizations often struggle with enabling secure and reliable communication and data exchange between AI agents from different vendors and with integrating the resulting agents with existing complex enterprise systems and application environments. According to Enterprise Strategy Group research, 41% of organizations reported that difficulties integrating a new AI solution with their existing systems was a significant roadblock on the way toward realizing the expected value of their AI projects.

The joint Red Hat and KPMG framework aims to address the integration challenge through a structured, standards-based approach leveraging OpenShift's open architecture and KPMG LLP's enterprise-grade implementation methodologies, including:

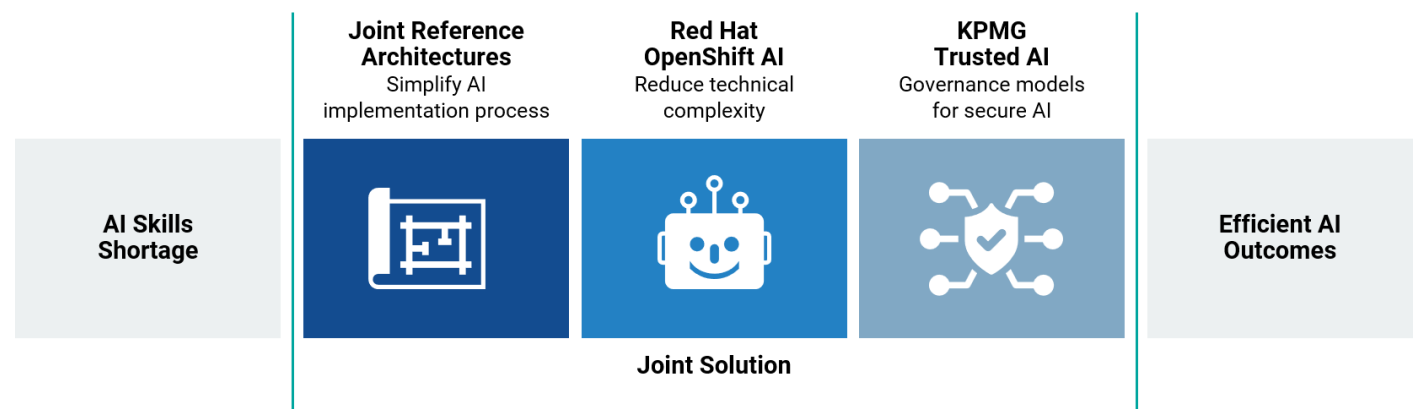
- **Open and standardized frameworks.** The solution is built upon open standards, such as Agent2Agent and MCP, to enable secure, structured, and reliable communication and data exchange between AI agents, enterprise systems, APIs, and existing applications.
- **Scalable container-based infrastructure:** Leveraging Red Hat OpenShift's container orchestration capabilities and Ansible Automation, the platform aims to facilitate rapid deployment and efficient integration of agentic AI components with diverse internal infrastructure, legacy software systems, and cloud-based services.
- **Enhanced inference across diverse environments:** Red Hat's platform supports the full spectrum of inference workloads, accelerators, and models that enterprises inevitably encounter. Through strategic partnerships with major semiconductor vendors, OEMs, and NVIDIA NVAIE, Red Hat aims to ensure optimal performance regardless of the underlying infrastructure, providing the flexibility enterprises need without sacrificing efficiency.
- **Structured data integration and management.** Through MCP-enabled structured input/output mechanisms, the joint framework provides a consistent, standardized way to connect AI agents to existing data sources, APIs, middleware, and business workflows, streamlining integration efforts.
- **Enterprise-grade governance and controls.** KPMG provides governance best practices, compliance frameworks, risk management controls, and secure implementation methodologies that proactively address integration complexity, helping ensure alignment with organizational standards and objectives.

By combining these elements into a cohesive reference architecture, the joint Red Hat and KPMG platform can significantly reduce the complexity of effectively integrating AI agents with the rest of the corporate infrastructure.

Lack of Staff Skills

Insufficient staff skills remain a critical challenge that hinders many organizations from successfully realizing the expected value of AI initiatives. Agentic AI solutions require specialized knowledge across multiple disciplines, including data science, AI modeling, integration architecture, cybersecurity, and governance. Many organizations find themselves lacking in-house expertise to effectively design, deploy, and manage these advanced AI capabilities. Without adequate training, knowledge transfer, and support structures to build these skill sets internally, organizations face prolonged implementation cycles, increased reliance on external resources, and heightened risks of delays or failures in delivering anticipated business outcomes.

Red Hat and KPMG LLP's joint architecture addresses the lack of staff skills by providing clearly defined reference architectures, prebuilt accelerators, and standardized frameworks that simplify and streamline AI implementation. Red Hat OpenShift AI offers preconfigured, container-based infrastructure and development tools that aim to reduce technical complexity and shorten learning curves for internal teams. KPMG LLP's structured Trusted AI Framework brings established governance models, risk management practices, and guided methodologies for secure and compliant AI deployments. Together, these capabilities aim to significantly reduce the need for deep in-house expertise, enabling organizations to leverage best practices, rapidly upskill their teams through structured guidance, and efficiently achieve desired AI outcomes.

Figure 2. KPMG and Red Hat's Joint Solution Addresses Staff Skills Gaps

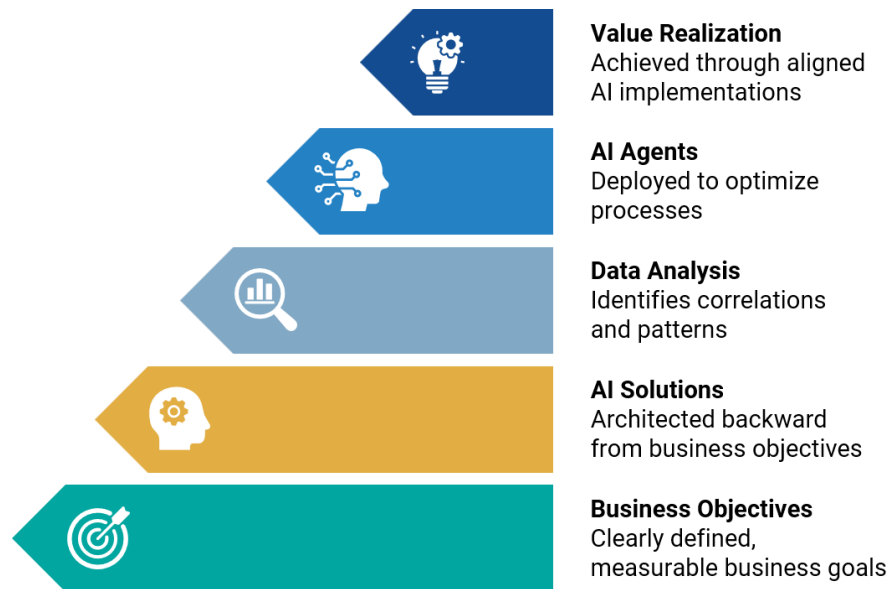
Source: Enterprise Strategy Group, now part of Omdia

Business-centric Design

As in every technology project, clearly defined business objectives are essential from the outset. Over one-third (38%) of organizations cited unclear business goals as a primary roadblock to realizing value from their AI initiatives. Agentic AI deployments frequently lose focus when they begin with technology rather than defined business outcomes. Organizations must first establish measurable business objectives, then architect their AI solutions backward from these goals. For example, consider customer onboarding abandonment: Data analysis might reveal strong correlations between form completion rates and factors such as process duration, error frequency, and application latency. A clear and measurable business objective emerges: Reduce abandonment rates by 50% through faster, more accurate form completion. To achieve this, AI agents can be deployed to retrieve customer data from multiple sources, automatically pre-populate forms, and proactively identify potential completion barriers. These agents can analyze historical patterns and regulatory requirements to suggest workflow optimizations before issues arise.

Red Hat's approach to connecting models with enterprise data and use cases seeks to enable organizations to rapidly prototype and deploy these agent-based solutions. The platform's wide spectrum of capabilities accelerate agent development while seeking to ensure efficient integration with existing applications and data estates—a critical requirement as enterprises move beyond isolated AI experiments to production-scale deployments.

The joint Red Hat and KPMG solution specifically addresses the issue of unclear business objectives by combining KPMG LLP's structured approach to defining measurable business outcomes with Red Hat's flexible, scalable AI infrastructure. KPMG provides proven methodologies, governance frameworks, and industry-specific skills to guide organizations in establishing clear, outcome-based goals aligned to their strategic priorities. Red Hat OpenShift AI complements this approach with an adaptable technology foundation that enables rapid prototyping, iterative experimentation, and agile deployment. Together, this combined approach ensures that AI implementations remain closely aligned with clearly articulated business objectives, accelerating value realization and reducing the risk of project delays or failures due to unclear or shifting priorities.

Figure 3. Clear Business Goals Are the Foundation for Realizing the Value of AI Agents


Source: Enterprise Strategy Group, now part of Omdia

Critical AI Implementation Challenges

Cost, data management and quality issues, and concerns over data privacy and security are key AI-related implementation challenges. This underscores the necessity for organizations to adopt a structured and holistic approach to deploying and scaling agentic AI solutions.

Figure 4. Top 3 Challenges Encountered While Implementing AI

What are the top challenges your organization has encountered while implementing AI? (Percent of respondents, N=376, three responses accepted)



Source: Enterprise Strategy Group, now part of Omdia

Cost

High implementation costs frequently hinder successful AI adoption, as organizations encounter substantial expenses related to infrastructure, licensing, integration complexity, specialized talent, and ongoing maintenance. Enterprise Strategy Group research underscores these financial pressures, highlighting cost as a major barrier to realizing anticipated AI value. Without a carefully managed strategy, AI projects can quickly become financially unsustainable, leading to stalled initiatives, delayed timelines, or outright cancellations.

Red Hat's focus on fast, flexible, and efficient inferencing aims to directly address cost concerns by optimizing resource utilization across diverse accelerators and deployment models. This efficiency, combined with the platform's ability to leverage existing infrastructure investments across hybrid cloud environments, creates a financially sustainable path to enterprise-scale AI deployment.

KPMG LLP's structured methodologies, architecture blueprint, prebuilt accelerators, and proven governance frameworks streamline implementation efforts, shorten project timelines, and help minimize resource-intensive tasks. By combining Red Hat's technology foundation with KPMG LLP's specialization in deployment and management, the joint platform can deliver a financially sustainable path to deploying agentic AI capabilities at enterprise scale and with lower risk to delivery.

Figure 5. Implementation Expenses Can Quickly Derail AI Projects

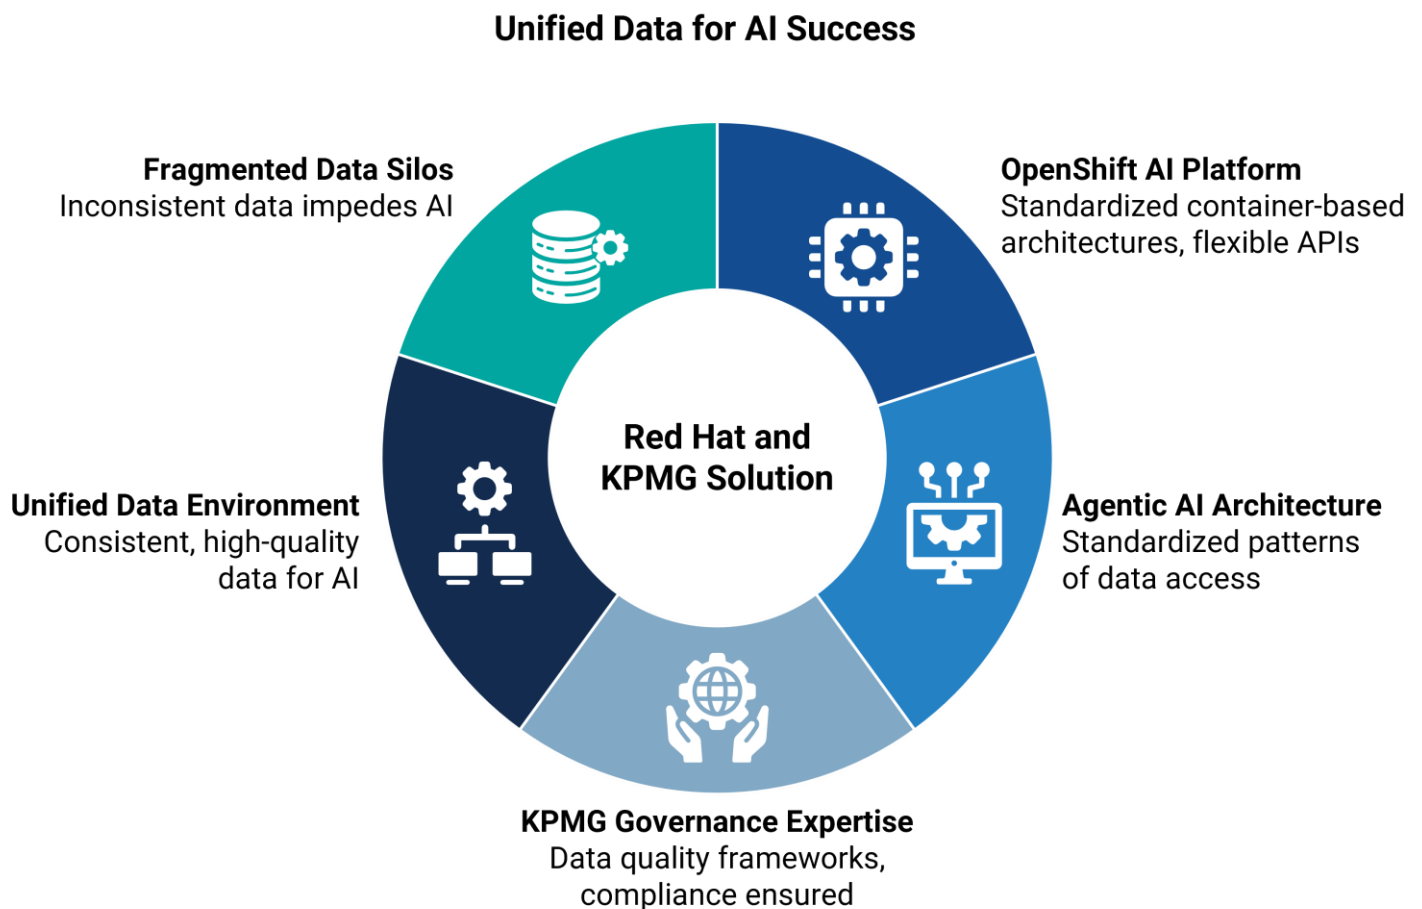


Source: Enterprise Strategy Group, now part of Omdia

Data Management and Quality

Data management and quality issues represent the second most critical barrier organizations face when implementing AI solutions. Effective AI models depend fundamentally on timely, accurate, and complete data. However, many organizations struggle with fragmented data environments, inconsistent data formats, outdated information, and poor-quality records spread across multiple systems and silos. These data challenges significantly impede AI implementation, reducing model accuracy, undermining trust in insights, and ultimately limiting the business value organizations can realize from their AI initiatives.

The joint Red Hat and KPMG solution aims to address data access and quality issues by providing a structured and unified approach to enterprise data integration and management. Red Hat's OpenShift AI platform leverages standardized container-based architectures, flexible APIs, structured methodologies such as retrieval-augmented generation, and integrated data pipelines to enable organizations to realize greater business value out of unstructured data. KPMG and Red Hat's Agentic AI Reference Architecture provides standardized patterns of data access and integration. Complementing this, KPMG LLP's specialization in robust governance practices, data quality frameworks, and compliance ensures data is consistently managed, validated, and monitored across the AI lifecycle. Together, these capabilities enable organizations to overcome fragmentation, improve data consistency, and enhance overall data quality, accelerating successful AI deployment and ensuring reliable, high-quality results from AI-driven solutions.

Figure 6. Data Integration Is the Foundation for Realizing Optimal Value From AI Initiatives

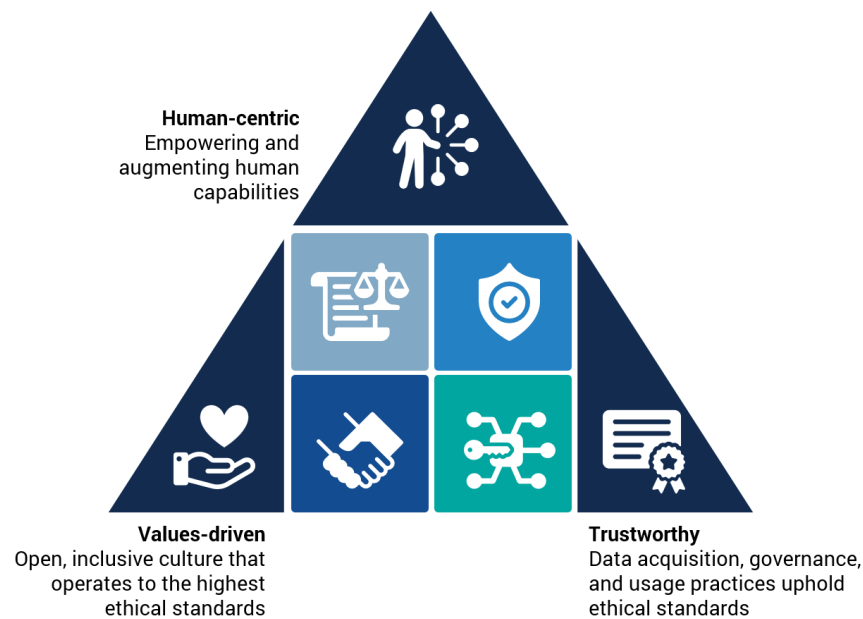
Source: Enterprise Strategy Group, now part of Omdia

Data Privacy, Protecting Intellectual Property, and Security

Data privacy, protecting intellectual property, and security concerns represent critical challenges organizations face when implementing AI solutions. Without robust security controls and clearly defined compliance frameworks, organizations risk exposure to data breaches, unauthorized access, regulatory penalties, and reputational damage, which can severely derail AI initiatives.

The joint solution from Red Hat and KPMG aims to address these challenges by embedding rigorous governance, security measures, and compliance standards directly into the AI implementation process. Red Hat OpenShift AI provides a secure and enterprise-ready infrastructure platform, offering built-in features such as identity management, secure container orchestration, and extensive monitoring capabilities. KPMG complements this with its Trusted AI Framework, which delivers structured governance models, risk management controls, compliance tracking mechanisms, and security best practices tailored specifically for AI solutions. As a regulated entity, KPMG brings the same amount of rigor to helping clients deploy AI solutions in a regulated and compliant environment as they are required to do with other aspects of their business operations. Together, Red Hat's secure infrastructure and KPMG LLP's extensive governance approach aim to enable organizations to confidently deploy AI solutions that consistently meet stringent security requirements and regulatory standards.

Figure 7. KPMG LLP's Trusted AI Framework



Source: Enterprise Strategy Group, now part of Omdia

Conclusion

Achieving the ability to consistently glean business value from agentic AI projects across the enterprise has become a key differentiator for most organizations. The collaboration between KPMG and Red Hat aims to pick up organizations where they currently are and provide them with the processes, tools, and staff skills they need to effectively harness agentic AI. This structured and holistic approach helps organizations navigate common pitfalls, such as integration complexity, staff skill gaps, and unclear business objectives.

Red Hat's strategic focus on enhanced inference, efficient enterprise data connectivity, accelerated agent development, and hybrid cloud scalability—combined with its investments in open source AI innovation—provides a stable yet flexible foundation for enterprises. This approach, enhanced by KPMG LLP's governance skills, aims to enable organizations to move beyond experimental AI toward production-ready, enterprise-scale deployments that deliver consistent value. Ultimately, adopting this joint reference architecture could position enterprises to confidently leverage agentic AI at scale, turning AI-driven innovation into tangible business outcomes.

©2025 TechTarget, Inc. All rights reserved. The Informa TechTarget name and logo are subject to license. All other logos are trademarks of their respective owners. Informa TechTarget reserves the right to make changes in specifications and other information contained in this document without prior notice.

Information contained in this publication has been obtained by sources Informa TechTarget considers to be reliable but is not warranted by Informa TechTarget. This publication may contain opinions of Informa TechTarget, which are subject to change. This publication may include forecasts, projections, and other predictive statements that represent Informa TechTarget's assumptions and expectations in light of currently available information. These forecasts are based on industry trends and involve variables and uncertainties. Consequently, Informa TechTarget makes no warranty as to the accuracy of specific forecasts, projections or predictive statements contained herein.

Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of Informa TechTarget, is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact Client Relations at cr@esg-global.com.