softserve | Microsoft

BEST PRACTICES TO DEPLOY AGENTIC AI WITH MICROSOFT FABRIC

How to modernize your architecture and operationalize agentic AI in weeks, not years

Why most AI initiatives fail — and how to fix it

Enterprises are racing to deploy AI agents to automate tasks, accelerate decisions, and scale operations. Yet most initiatives stall or fail — UC Berkeley research shows common pitfalls like unclear goals, poor agent coordination, and weak quality checks.

Failures waste millions in investment and leave organizations stuck with fragmented systems, security risks, and missed opportunities for innovation. Meanwhile, competitors who get it right are achieving real-time decision-making and operational agility.

This white paper covers how to overcome these barriers using agentic AI and Microsoft Fabric. You'll learn best practices for building multi-agent systems (MAS) that deliver intelligent automation, governed data, and innovation at scale — securely and in weeks, not years.

What are the benefits of multi-agent systems?

Implementing MAS enables organizations to streamline operations, enhance decision-making, and achieve significant cost savings. Multi-agent systems are also proving particularly effective in automating complex processes and optimizing resource allocation.

Benefit	Description
Operational efficiency	Automate repetitive tasks
Enhanced customer service	Handle customer queries, routing them to the resources
Data processing speed	Accelerate decision-making
24/7 availability	Provide support without human intervention
Scalability	Adapt to fluctuating data and growing user demands

What are the use cases for MAS technology?

MAS is gaining traction across multiple industries, with proven use cases demonstrating its potential to automate, optimize, and innovate. Here's how MAS is being effectively applied in key sectors:

Industry	Application Areas	Implementation Examples
Retail	Inventory management, sales forecasting, and customer profiling	Autonomous inventory tracking
Healthcare	Patient monitoring, clinical data analysis, and resource allocation	24/7 patient care support
Finance	Fraud detection, investment portfolio optimization, and risk assessment	Intelligent dispute management

What are the common failures in MAS design?

Despite these advantages, MAS is not a panacea. Its implementation presents significant challenges. The previously mentioned UC Berkeley study identifies critical areas where MAS deployments are most likely to encounter difficulties:

Challenge	Failures	Impact
System design issues	37.2%	• Poor task specification • Inadequate role definition • Inefficient execution patterns
Inter-agent misalignment	31.4%	• Communication breakdowns • Information withholding • Ignoring other agents' input
Verification problems	31.4%	• Premature termination • Inadequate quality control • Incorrect verification

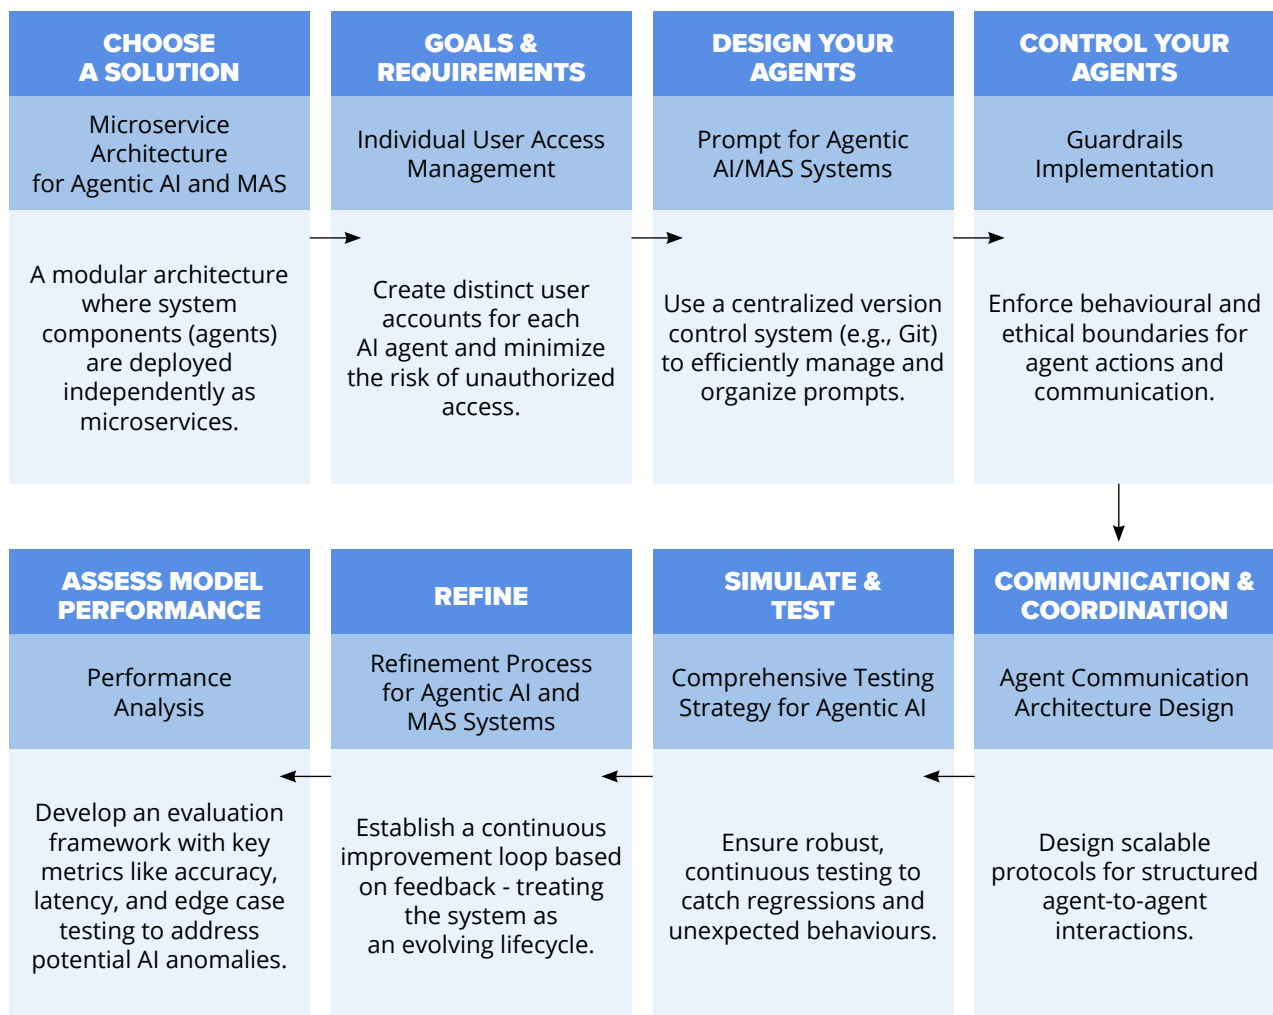# How to mitigate risk when designing multi-agent systems

Organizations seeking to implement multi-agent systems effectively must prioritize specific strategies to mitigate risk and maximize impact. The following best practices help avoid roadblocks:

6 BEST PRACTICES FOR MULTI-AGENT SYSTEMS		
DEFINE AGENT ROLES CLEARLY	**ESTABLISH ROBUST COMMUNICATION PROTOCOLS**	**IMPLEMENT RIGOROUS TESTING AND VERIFICATION**
Ensure each agent has distinct responsibilities to prevent overlap and confusion.	Implement mechanisms for clear, consistent data sharing among agents.	Continuously validate agent outputs to catch errors early and maintain data integrity.
PLAN FOR SCALABILITY	**IMPLEMENT ADAPTIVE DECISION-MAKING**	**INTEGRATE SECURITY CONTROLS**
Design systems that can dynamically adapt to increasing data volumes and agent interactions.	Enable agents to learn from feedback, adapt in real time, and respond to changing needs.	Protect sensitive data by enforcing strict access controls and monitoring agent activities.

How to build a multi-agent system

Understanding the agentic AI lifecycle is essential for developing systems that not only act autonomously but also evolve responsibly across diverse operational contexts. By following best practices, organizations can enhance scalability, adaptability, and efficiency, ensure seamless integration, and maximize the value of AI-driven solutions.

CHOOSE A SOLUTION	GOALS & REQUIREMENTS	DESIGN YOUR AGENTS	CONTROL YOUR AGENTS
Microservice Architecture for Agentic AI and MAS	Individual User Access Management	Prompt for Agentic AI/MAS Systems	Guardrails Implementation
A modular architecture where system components (agents) are deployed independently as microservices.	Create distinct user accounts for each AI agent and minimize the risk of unauthorized access.	Use a centralized version control system (e.g., Git) to efficiently manage and organize prompts.	Enforce behavioural and ethical boundaries for agent actions and communication.

ASSESS MODEL PERFORMANCE	REFINE	SIMULATE & TEST	COMMUNICATION & COORDINATION
Performance Analysis	Refinement Process for Agentic AI and MAS Systems	Comprehensive Testing Strategy for Agentic AI	Agent Communication Architecture Design
Develop an evaluation framework with key metrics like accuracy, latency, and edge case testing to address potential AI anomalies.	Establish a continuous improvement loop based on feedback - treating the system as an evolving lifecycle.	Ensure robust, continuous testing to catch regressions and unexpected behaviours.	Design scalable protocols for structured agent-to-agent interactions.

Why governance and guardrails matter for AI

Each implementation step introduces both specific benefits and potential consequences of non-implementation. This helps teams prioritize efforts and mitigate vulnerabilities during agentic AI and MAS deployment.

Step	Impact	Consequences of Non-Implementation
Microservice architecture for agentic AI and MAS	Boosts efficiency and reliability, reduces operational risk, and lowers total cost of ownership	Without this architecture, systems become rigid, costly to maintain, and less reliable
Individual user access management	Enhances data security and control, reducing the risk of breaches or unauthorized changes.	Without this practice, agents or users may compromise sensitive data or misconfigure systems, leading to disruptions and security risks
Prompts for agentic AI/MAS systems	Enables change tracking and rollbacks, ensuring consistent and reliable system performance	Without this practice, organizations risk losing insights and reducing system adaptability and effectiveness
Guardrails implementation	Keeps interactions domain-specific, preserving context and preventing hallucinations, toxic, or misleading content	Without guardrails, there's a risk of unauthorized actions, data breaches, and unreliable system behavior, undermining user trust
Agent communication architecture design	Improves decision-making and reliability, supports growth, and increases loads	Lack of structured communication causes errors and inefficiencies, and limited adaptability
Comprehensive testing strategy for agentic AI	Improves the reliability and performance of agentic AI systems, ensuring they meet user needs and boost trust, engagement, and adoption	Without comprehensive testing, systems may misinterpret instructions, leading to disruptions, decreased trust, inefficiencies, and higher correction costs
Refinement process for agentic AI and MAS systems	Boosts system accuracy, alignment, and adaptability while reducing technical debt and supporting innovation and compliance	Without refinement, systems become outdated, brittle, and disconnected from user or business goals
Performance analysis	Evaluation reveals system limitations, guiding iterative improvements and boosting resilience. Early analysis shapes development, steering the system toward optimal performance	Inadequate evaluation leads to unchecked vulnerabilities, degrades performance, and undermines trust. Without iterative refinement, systems accrue technical debt, risk non-compliance, and hinder scalability

What is Microsoft Fabric and how does it support AI governance?

Microsoft Fabric ensures role-based access, data lineage, and compliance — critical for implementing guardrails and governance at scale. This ensures AI agents operate on trusted, well-managed data.

Benefits: Reduces risk, improves transparency, and accelerates enterprise-scale AI deployment.

Next steps:
How to build and govern agentic AI

Implementing AI agents and multi-agent systems in production requires a strategic approach emphasizing adaptability, collaboration, and continuous improvement. By adhering to best practices, organizations can harness the full potential of these technologies to streamline operations and enhance decision-making processes. Comprehensive testing and robust functional strategies ensure these systems perform reliably, mitigating risks associated with non-deterministic behaviours. Moreover, fostering a culture of innovation and feedback integration allows for ongoing refinement and optimization of AI systems.

As businesses grow, the ability to effectively deploy and manage AI agents will be a critical factor in maintaining a competitive advantage.

Ultimately, leveraging AI agents and multi-agent systems paves the way for transformative business outcomes, driving efficiency and fostering growth.

Ready to modernize with Microsoft Fabric and agentic AI?

Connect with us to deploy secure, scalable solutions in weeks, not years.

CONTACT US

SoftServe is a Microsoft Solutions Partner with more than 20 years of collaboration and solutions delivered across Azure, Microsoft Fabric, and AI services.

As a Microsoft Fabric Featured Partner, also recognized for Real-Time Intelligence and Fabric Databases, we have expertise in data, analytics, and AI solutions.

Achieving two new specializations in AI Platform on Azure and Cloud Security reinforces the company's deep expertise. Businesses are eager to unlock scalable, secure, and responsible AI. This assures they'll have ample protection and safeguards in place.

About Us

SoftServe is a premier IT consulting and digital services provider.

We expand the horizon of new technologies to solve today's complex business challenges and achieve meaningful outcomes for our clients. Our boundless curiosity drives us to explore and reimagine the art of the possible. Clients confidently rely on SoftServe to architect and execute mature and innovative capabilities, such as digital engineering, data and analytics, cloud, and AI/ML.

Our global reputation is gained from more than 30 years of experience delivering superior digital solutions at exceptional speed by top-tier engineering talent to enterprise industries, including high tech, financial services, healthcare, life sciences, retail, energy, and manufacturing.

Visit our website, blog, LinkedIn, Facebook, and X (Twitter) pages for more information.

AUSTIN HQ

201 W. 5th Street, Suite 1550
Austin, TX 78701
+1 866 687 3588 (USA)
Toll Free: +1 866 687 3588

LONDON

30 Cannon Street
London EC4 6XH
United Kingdom
+44 203 807 01 41

info@softserveinc.com
www.softserveinc.com

softserve