

From AI Pilot to Production

IBM's Open AI Operating Model
A Practical Path to Production, Governance, and
Business Impact



Li9 & IBM Webinar
April 22, 2026 | 1:00 PM EST



Armando Arias
Li9 Technology Solutions
President & CEO
armando.arias@Li9.com
480.712.1812



Matthew Kelm
Product Manager
IBM

Agenda

- Evolution of Enterprise AI
- Why Enterprise AI Pilots Stalled
- What Production-Ready AI Requires
- AI Standards. Finally!
- IBM's Agentic AI Operating Model
- IBM's Open Ecosystem for Agentic AI
- OpenShift: A Unified Platform for Modern IT
- Fusion HCI for On-Premise Agentic AI

Consider this slide or cover the Li9

WHY AI PILOTS DON'T REACH PRODUCTION



Poor Data Quality

Fragmented, low-quality, or inaccessible enterprise data.



Governance & Risk Concerns

Limited oversight, control, and accountability.



Platform Complexity

Disconnected tools, processes, and environments.



Hard to Scale

Difficulty moving from pilot to production with repeatability and consistency.



Poor Workflow Fit

Pilots that don't align with real business and IT processes.



WHAT DRIVES ENTERPRISE AI SUCCESS AT SCALE



Governed Data Access

Trusted enterprise content for retrieval and decision-making.



Centralized AI Governance

Policy, risk, and monitoring across the AI landscape.



Production-Ready Platform

Resilient infrastructure with consistent operations and security.



Unified Operating Model

A common foundation across AI, apps, and data.



Workflow-Specific Design

AI embedded into business processes with measurable outcomes.

One-Off AI Solutions to Standard AI Infrastructure

2024–2025

The Custom Integration Era (No Standards)



Custom-Built Everything

Moving pilots to production required custom-building every layer (APIs, RAG, auth, governance) for every tool/LLM combination.



One-Off Engineering

Deployments were manual, creating “interoperability chaos” and fragmented liability chains.



Scalability Hurdles

Lack of shared protocols made AI agents difficult to audit, secure, or reuse across teams, contributing to a **high** pilot failure rate.

2025–2026

Standardized AI Infrastructure



MCP (Model Context Protocol)

The “USB-C for AI.”
Standardizes how agents connect to external tools, systems, and enterprise data sources.



AGENTS.md

The “README for bots.”
A standardized Markdown format for providing coding agents with project-specific instructions, build rules, and constraints.



goose

The “Autonomous Developer”
Unlike traditional AI assistants that only suggest code, Goose owns the entire execution loop: read, plan, edit, run commands, test, and iterate until the task is done.

Agentic AI Foundation (AAIF)

Open Governance



Founded Dec 9, 2025

Established under the Linux Foundation to provide neutral stewardship for agentic infrastructure.



Unified Projects

Brings together MCP (Anthropic), AGENTS.md (OpenAI), and goose (Block) under a single open-source governance model.



Enterprise Strategy

De-risks AI adoption by reducing vendor lock-in and providing a standardized path from experimentation to production-grade implementation.

Why Enterprise AI Pilots Stalled



**Fragmented
data and
content**



**Weak
governance
and risk
controls**



**Too many
disconnected
tools**



**Difficult to
scale and
repeat**



**Not tied to
business
outcomes**


Requirements for Production AI at Scale




Governed access to current, trusted data




Centralized AI governance and lifecycle management



Production ready Day-2 operations



Standardized operating models



AI embedded in business workflows

Agentic AI Standards across the watsonx Platform



1 Model Context Protocol (MCP)

Standardized access to data, tools, and enterprise systems



watsonx.data

Local and remote MCP servers for natural language access to lakehouse data, metadata, and document libraries.



Content-Aware Storage (CAS)

MCP server for semantic search, vector stores, and retrieval of enterprise content using natural language.



watsonx Orchestrate

Natively integrates MCP endpoints so agents can collaborate and connect to 80+ enterprise applications.



IBM's Sterling Application

MCP integration extends Order Management Systems (OMS) with external AI agents—with role-based access and audit logging.



2 AGENTS.md

Contextual governance for coding agents



IBM Bob

Uses AGENTS.md files to provide project-specific context: purpose, build and test steps, coding conventions, security rules, and human-review checkpoints.



Enterprise Value

Gives coding agents the institutional judgment they need to work reliably within repositories—reducing errors, increasing auditability, and improving SDLC governance.



3 goose

Open framework interoperability across enterprise infrastructure



Cross-Platform Execution

Goose is model-agnostic and can be powered by Open models, including IBM Granite.



Infrastructure Interoperability

Goose can interact with Red Hat OpenShift via MCP servers to retrieve system information, detect anomalies, and post updates.



Agent Catalog

Agents built on open-source frameworks like Goose can be onboarded into watsonx Orchestrate's curated Agent Catalog and operate as first-class citizens in governed enterprise workflows.



IBM is turning open agent standards into governed enterprise infrastructure.



Accelerate innovation with open standards

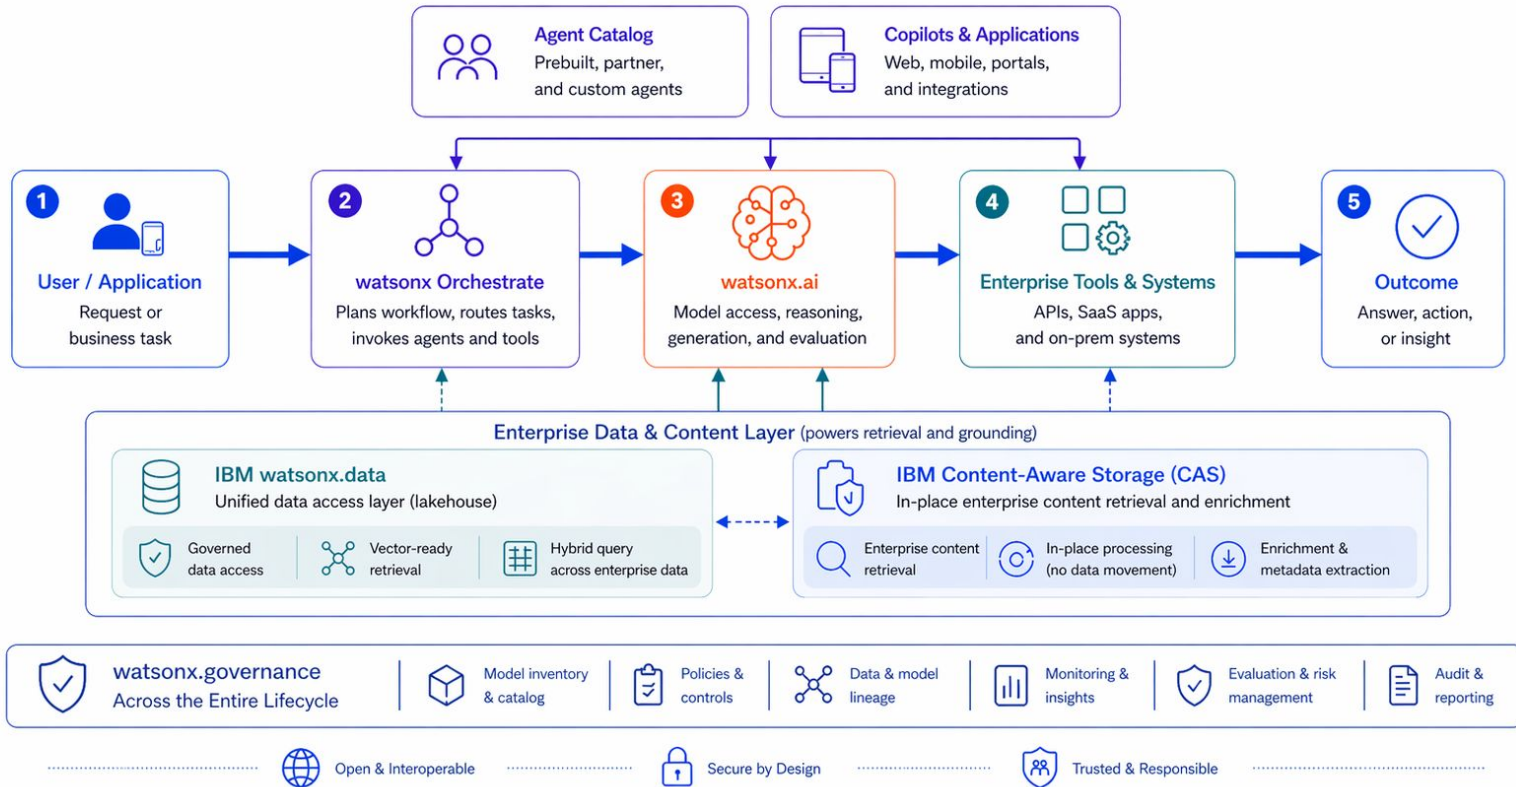



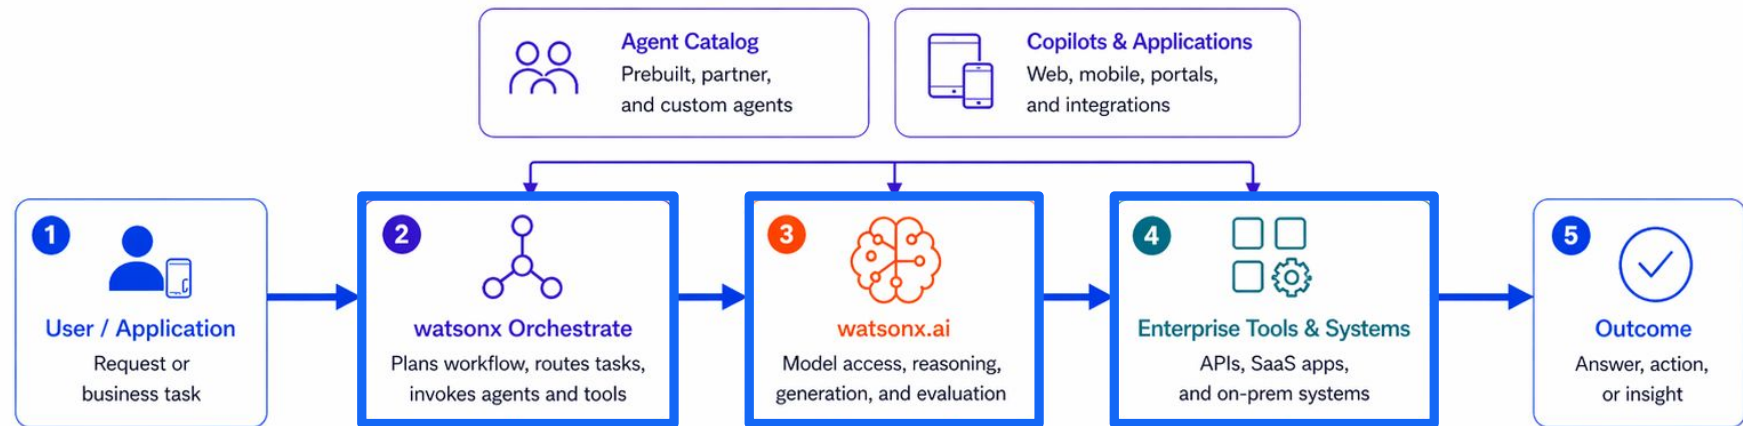
Reduce risk with built-in governance and security



Move from pilots to production at scale

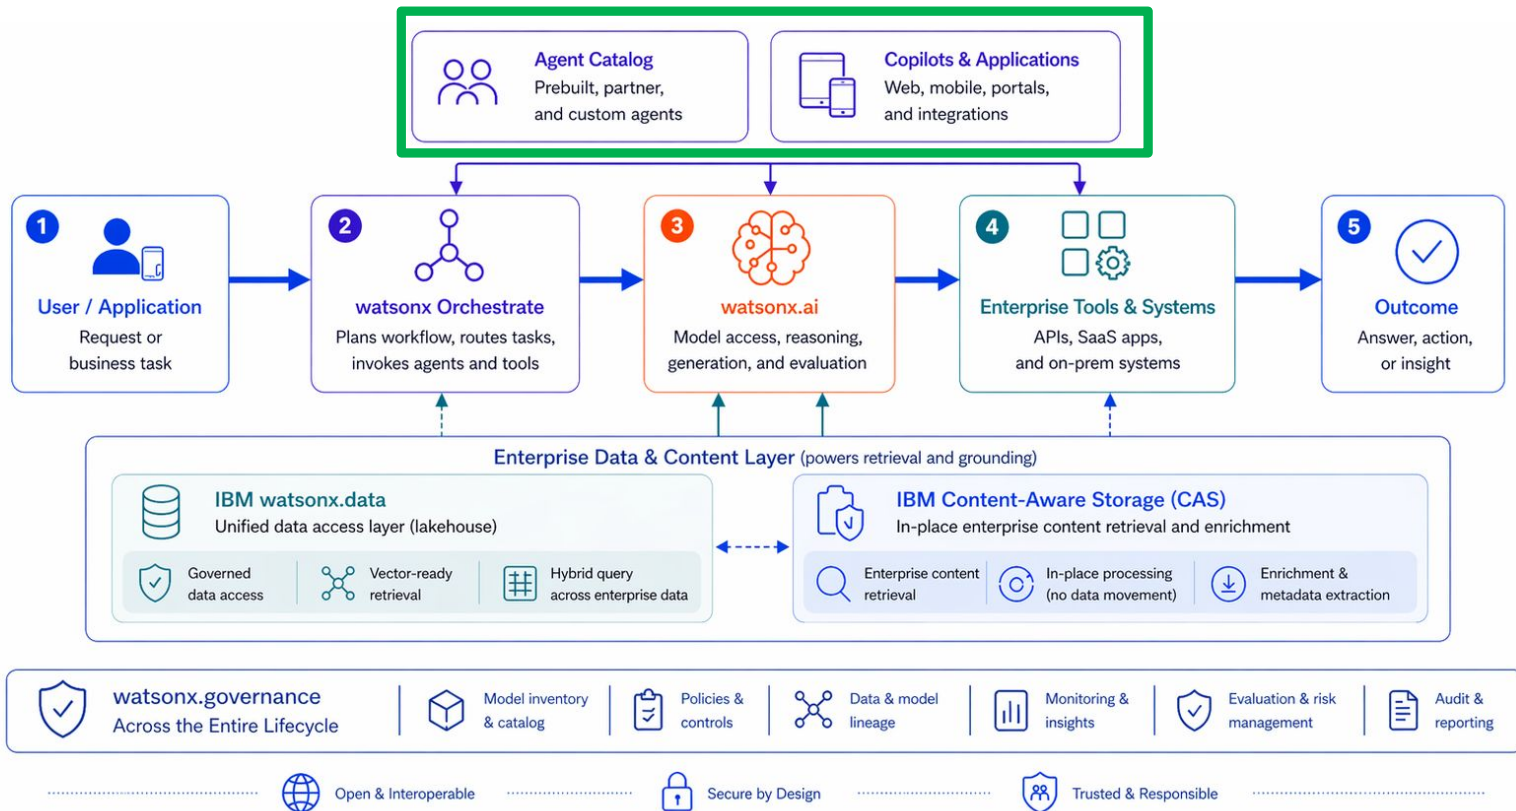
IBM's Agentic AI Operating Model





..... Open & Interoperable Secure by Design Trusted & Responsible

IBM's Agentic AI Operating Model



IBM Open Ecosystem for Agentic AI

An open, secure, and governed platform that lets you build, run, and scale AI with choice and control.



Model & Ecosystem Partners

- Hugging Face**
Model repository and community hub
- Meta**
Llama open model ecosystem
- Mistral AI**
Enterprise-ready open foundation models
- AI Anthropic**
Leading frontier models
- Other Third-Party LLMs**
Additional model choice and flexibility

Builders, Frameworks & Agent Tools

- LangChain**
Popular framework for LLM and agent applications
- CrewAI**
Multi-agent orchestration and workflow building
- Flowise**
Visual builder for LLM and agent flows
- Composio**
Tool and app integrations for AI agents
- Other Agent Tools**
Flexible builder and runtime choices

Data & Platform Partners

- Snowflake**
Enterprise cloud data platform
- Databricks**
Lakehouse and AI/ML workflows
- MongoDB**
Operational and vector data support
- Cloudera**
Hybrid enterprise data platform
- DataStax**
Vector database and retrieval support
- Other Enterprise Data Sources**
Flexible integration across platforms



Open & Interoperable
Built on open standards.
Works with your ecosystem.

Secure by Design
Enterprise security, privacy,
and compliance built in.

Trusted & Responsible
Built on governance, transparency,
and responsible AI principles.

OpenShift: A Unified Platform for Modern IT

AI Workloads

Models • Agents • AI Services



watsonx AI • Orchestrate • Data • CAS • Governance

Containerized Applications

Microservices • APIs • Cloud-native Apps



Build • Deploy • Scale

Virtual Machines

Existing Enterprise / Legacy Workloads



Modernize • Protect • Run



Red Hat OpenShift

One Platform. Consistent capabilities. Any workload. Any environment.



Security



Monitoring



Logging



Automation



Backup &
Disaster Recovery



Cluster
Management



Access Control



Public Cloud

AWS • Azure • Google Cloud



On-Premises

Private Data Centers

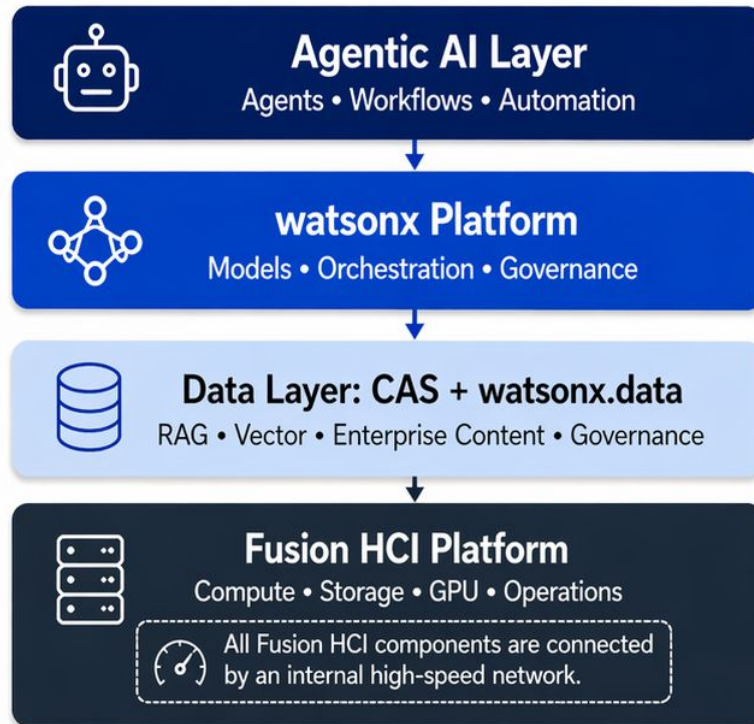


IBM Fusion HCI





Factory integrated Network, Servers,
Storage, & GPU Platform

How Fusion HCI Enables watsonx, CAS, and Agentic AI

Turning models, data, and workflows into production-ready enterprise systems





What Fusion HCI Provides

-  **Data Access & Memory for Agents**
CAS + watsonx.data provide governed, real-time access to enterprise knowledge
-  **Runtime for Agentic Execution**
Fusion delivers low-latency inference and always-on execution
-  **Integrated AI + Data Platform**
Co-locates compute and data to improve performance and reduce cost
-  **Enterprise Operations & Governance**
Ensures reliability, security, and lifecycle management at scale

IBM Fusion HCI

Fusion HCI is an integrated infrastructure designed to reduce deployment time and eliminate integration risk.

- It's delivered as a standard 42U rack, pre-cabled for both network and power, and fully integrated and tested at the factory.
- That means it arrives ready to deploy, not something that has to be assembled and validated on-site.
- It can expand up to 3 racks, supporting a mix of CPU servers and GPU nodes.
- In total, this can scale to dozens of servers and up to 96 NVIDIA GPUs, depending on configuration
- Supported end-to-end by IBM




IBM Fusion HCI Components


IBM and Customer monitoring of Fusion HCI Hardware & Software Infrastructure



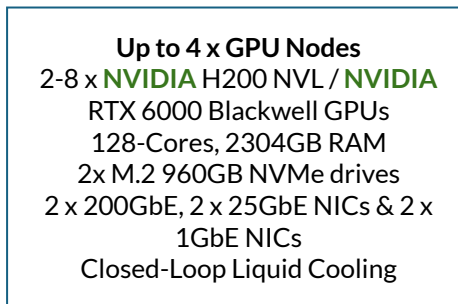


OpenShift Management Cluster & Storage (Control Nodes)



OpenShift Containers, VMs, and AI & Servers (Worker Nodes)




Integrated Fusion HCI Networking

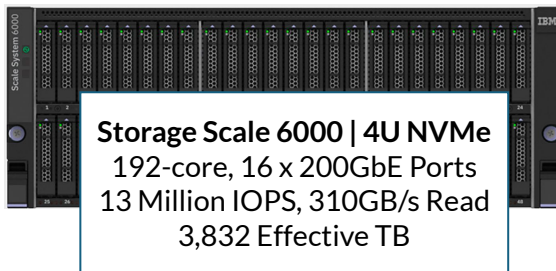



Fusion HCI External Storage





Storage Network



IBM Storage Scale System



IBM FlashSystem



Entry-Level Fusion HCI AI Solution

The key point is that organizations can start with a compact, production-ready configuration and scale from there as demand grows.

- At the top is the service node, followed by integrated networking for both data and management.
- Next are the OpenShift control nodes, which manage the cluster.
- Below that are the worker nodes, where applications and services run.
- And at the bottom are GPU-enabled worker nodes, which support AI and agentic workloads.
- This shows how everything is packaged into a single rack—from infrastructure to AI execution.



Next Steps

Explore Your AI Readiness with Li9 and IBM

Let's have a conversation about your goals, challenges, and opportunities so we can help you plan a path to enterprise AI success.

Schedule a working session with Li9 and IBM to:



Learn about IBM's solutions

Review IBM's enterprise AI, watsonx, OpenShift, and Fusion HCI solutions.



Share your priorities

Discuss your current AI, automation, data, and infrastructure priorities.



Identify future opportunities

Explore where Agentic AI may create value in the next 6–12 months.



Plan with confidence

Understand what preparation is needed before production AI becomes urgent.

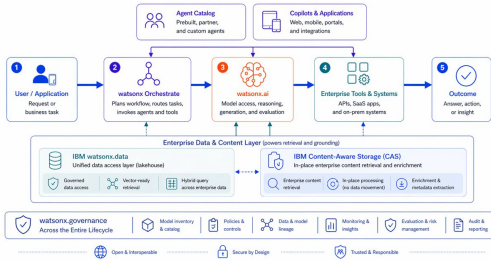


Agentic AI is moving quickly.

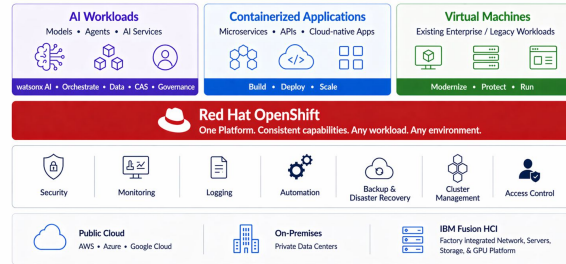
The best time to assess your platform readiness is before the requirement becomes immediate.

Questions?

IBM's Agentic AI Operating Model



OpenShift: A Unified Platform for Modern IT



Fusion HCI – The Platform for Enterprise AI & Agentic Workloads

