The Age of Agentic: Welcoming Agentic AIOps with a Comprehensive Framework for Hybrid Cloud, On-Premises, and IBM Z® Mainframe Ecosystems

The future is no longer on its way. Now, in 2026, it has arrived, and it's brought with it a fundamental shift in the way IT operations function. As IT operations shift away from insight-driven automation and move toward a paradigm of autonomous, agentic operations, traditional AIOps continue to be a critical piece of the puzzle, applying advanced machine learning to spot issues, correlate events, and deliver actionable insights at scale.

In today's enterprise IT landscape, what we think of as "traditional" AIOps primarily serves a supportive purpose, informing intelligent decision-making and helping human operators by tracking patterns and recommending corrective actions. Rather than replacing this familiar foundation, Agentic AIOps builds on it.

Agentic AIOps extends the capabilities of traditional AIOps by introducing goal-driven agents that function with total autonomy and have the ability to reason over AIOps insights, apply policy and context, and perform multi-step workflows completely on their own.

In other words: If Traditional AIOps answers the question, "What is happening, and why?", Agentic AIOps determines "What should be done," coordinating action across domains, and then safely executing within the guardrails you've set for it.

What results is a move away from remediation that requires human intervention, and a push toward machine-executed autonomy. AIOps still has a place in the picture, providing a key intelligence layer, while agentic systems connect us with the decisioning, orchestration, and closed-loop execution needed to keep complex hybrid, multi-cloud, and mainframe environments operational at scale.

Capgemini

IBM Platinum Partner

The Evolution of Operational Intelligence

Moving from Pattern Recognition to Goal-Driven Agency

Understanding the difference between traditional AIOps and the agentic model that's quickly taking hold means understanding the capabilities of the underlying artificial intelligence model.

Traditional AIOps systems, which primarily involve prediction and classification, are ideal for time-series analysis, clustering, and using past observations to actively spot the well-known signs indicating that failure is on the horizon. While this might work well for goals like noise reduction and alert correlation, they leave us with a system that's essentially reactive, relying on static rules or predetermined models to do its work. This requires human intervention to make the manual adjustments necessary to keep these systems relevant, especially as their environments evolve.

Agentic AI changes the game by replacing symbolic logic and rigid scripts with a more flexible form of generative reasoning, utilizing Large Language Models (LLMs) and advanced orchestration frameworks to work more autonomously and make decisions more proactively.

Agentic AI changes the game by replacing symbolic logic and rigid scripts with a more flexible form of generative reasoning, utilizing Large Language Models (LLMs) and advanced orchestration frameworks to work more autonomously and make decisions more proactively.

In this new, agentic ecosystem, intelligence is actively distributed across a network of specialized AI "agents," rather than centralized in a single place. This model allows each AI agent to possess not just a local knowledge base, but also an expertise related to their specific domain.

Communication, negotiation, and action coordination occur concurrently without the need for constant, centralized oversight, moving toward what we like to call a "zero-maintenance" operational paradigm.

Functional Attribute	Traditional AIOps Systems	Agentic AIOps Frameworks
Operational Philosophy	Augmented human capability; tool-centric	Autonomous entity; goal-centric
Logic and Reasoning	Predictive ML; static rules; if-then scripts	Generative AI; dynamic reasoning; LLM-based
Decision-Making Model	Centralized; requires human approval	Distributed; autonomous within policy bounds
Autonomy Level	Low to Moderate; reactive	High; proactive and adaptive
Workflow Complexity	Linear; predetermined; brittle	Non-linear; multi-step; self-correcting
Resource Management	Static threshold-based scaling	Adaptive, context-aware optimization
Maintenance Profile	High manual tuning and rule updates	Low; self-learning and self-optimizing

To illustrate why this change is necessary, we must understand the "transactional economy" that characterizes hybrid cloud environments today.

As organizations expand their IT capabilities—adopting microservices, containers, and architectures that can be distributed across public clouds, on-premises data centers, and mainframes—a breaking point is reached, where the sheer number of dependencies and failure modes begins to exceed the degree to which manual or even semi-automated systems can effectively manage them.

The solution? Agentic systems, which give us the ability to move from observation to action.

Reference Architecture for Multi-Environment Agentic AIOps

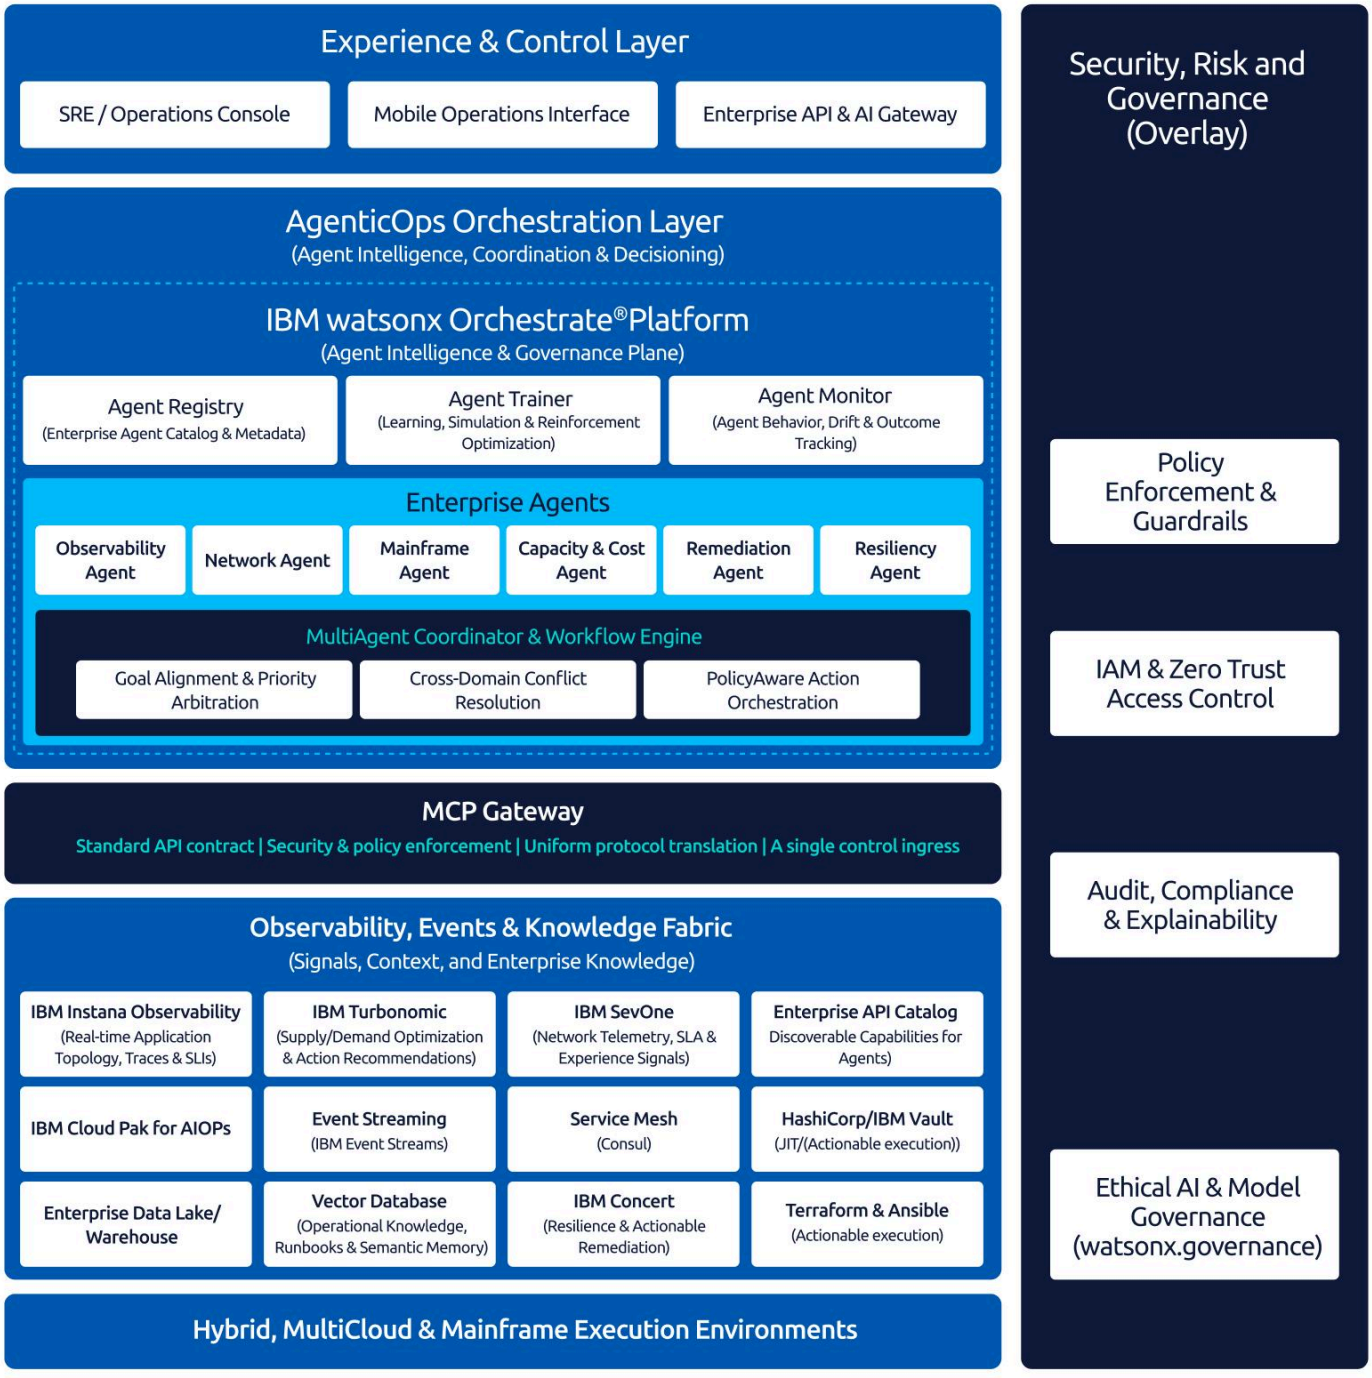

The current age of Agentic AIOps calls for a modern, governed, and extensible framework that makes autonomous operations possible across a variety of environments. IBM's AgenticOps makes this kind of autonomy possible across hybrid cloud, network, application, and mainframe environments by combining three key components:

1. The intelligence of enterprise-grade, agentic AI
2. The level of governance made possible by orchestration by IBM's watsonx
3. The extensibility of MCP (Model Context Protocol)

The result is an architecture that unlocks closed-loop, intent-driven operations, all of which are kept in full alignment with enterprise policy, risk controls, and human oversight.

Finally, the MCP Gateway allows agents to securely access the capabilities, tools, and enterprise systems they need to complete their work through a series of standardized, policy-governed interfaces.

Experience & Control Layer

Where humans and systems interact with AgenticOps - This layer provides the interfaces and gateways through which operators, engineers, and external systems engage with autonomous operations. It provides controlled, human-governed access to autonomous operations across the enterprise.

Components:

- SRE / Operations Console -Unified visibility into agent decisions, actions, recommendations, and system health.
- Mobile Operations Interface -On-call notifications, approvals, and real-time operational insights.
- Enterprise API & AI Gateway (API Connect) - Allows ITSM tools, CI/CD pipelines, and external applications to interact programmatically.
- Conversational Interface (watsonx Orchestrate) - Natural-language interaction for querying insights and approving/triggering actions.

AgenticOps Orchestration Layer

Agent Intelligence, Coordination, and Decisioning - This is the cognitive core of the architecture, powered by IBM watsonx Orchestrate providing reasoning, governance, and explainability for operational decisions.

Agent Registry - Central catalog of enterprise agents, capabilities, and permissions.

Agent Trainer - Simulation and reinforcement learning to continuously refine agent behavior.

Agent Monitor - Tracks drift, safety, compliance, and execution outcomes.

Enterprise Agents

Observability Agent	Network Agent
Mainframe Agent	Capacity & Cost Agent
Remediation Agent	Resiliency Agent

Each agent encodes a domain-specific operational intent and reasons using enterprise signals.

Multi-Agent Coordinator & Workflow Engine

- Goal alignment & prioritization
- Cross-domain conflict resolution
- Policy-aware and risk-aware action orchestration

MCP Gateway (Model Context Protocol Gateway)

Standardized, secure access to enterprise capabilities - The MCP Gateway provides a centralized control and mediation point through which all Enterprise Agents access downstream capabilities using the Model Context Protocol (MCP). It enforces security, compliance, access control, rate-limiting, and usage policies, while abstracting the underlying tools and platforms from agent logic.

MCP Connectors are registered and managed through the MCP Gateway and expose downstream systems as agent-ready capabilities via standardized MCP APIs. This approach enables agents to consume tools and platforms in a secure, normalized, and governable manner—without direct coupling to implementation details.

Supported MCP Connectors include:

IBM Instana Observability
IBM Turbonomic
IBM SevOne
IBM Z Operations
IBM Concert
HashiCorp / IBM Vault
Terraform & Ansible

Observability, Events & Knowledge Fabric

Signals, context, and operational knowledge - This layer aggregates the real-time data, historical insights, enterprise knowledge , and actionable capabilities required for agents to reason effectively.

Core Components:

IBM Instana Observability
real-time application telemetry

IBM SevOne
network performance and SLA signals

Event Streaming (IBM Event Streams)
high-volume telemetry

HashiCorp/IBM Vault
JIT secrets and identity for execution

IBM Concert
resilience analytics and automated remediation insights

Vector Database
semantic knowledge, runbooks, and prior decisions

IBM Turbonomic
supply/demand optimization & actions

IBM Cloud Pak for AIOps
event correlation and RCA

Service Mesh (Consul)
service-to-service visibility & resilience data

Terraform & Ansible
declarative and procedural automation

Enterprise Data Lake / Warehouse
historical and analytical data

Enterprise API Catalog
discoverable enterprise operations APIs

Hybrid, Multi-Cloud & Mainframe Execution Environments

Where autonomous actions are executed - This layer represents the platforms where actions initiated by agents are actually performed. It provides consistent, governed execution environments for autonomous workflows

Supported Environments:

- Public Cloud
- Private Cloud (Red Hat OpenShift)
- IBM Z® Platform (z/OS, Linux on Z, z/VM)
- Other On-Prem Environments

Security, Risk & Governance (Overlay)

Continuous trust, compliance, and control - This layer applies across all layers of the architecture. It ensure that autonomy is safe, transparent, compliant, and accountable.

Capabilities:

- Policy Enforcement & Guardrails
- IAM & Zero Trust Access Control
- Audit, Compliance & Explainability
- Ethical AI & Model Governance (watsonx.governance)

watsonx.governance integrates seamlessly into the Security, Risk & Governance overlay, which spans all layers of the reference architecture to ensure continuous trust, compliance, and control. Positioned as a key component for Ethical AI & Model Governance, it acts as a centralized platform for monitoring and managing AI models used by enterprise agents.

In the architecture diagram, it overlays the AgenticOps Orchestration Layer (where agents reason and decide) and the MCP Gateway (for secure tool access), enforcing policies on AI-driven actions. It also connects to the Observability, Events & Knowledge Fabric for real-time data ingestion and the Hybrid Execution Environments to audit autonomous executions.

By providing end-to-end visibility into model performance, bias, and drift, watsonx.governance ensures that agentic decisions remain aligned with enterprise risk thresholds, regulatory standards, and ethical guidelines, preventing unchecked autonomy in critical operations.

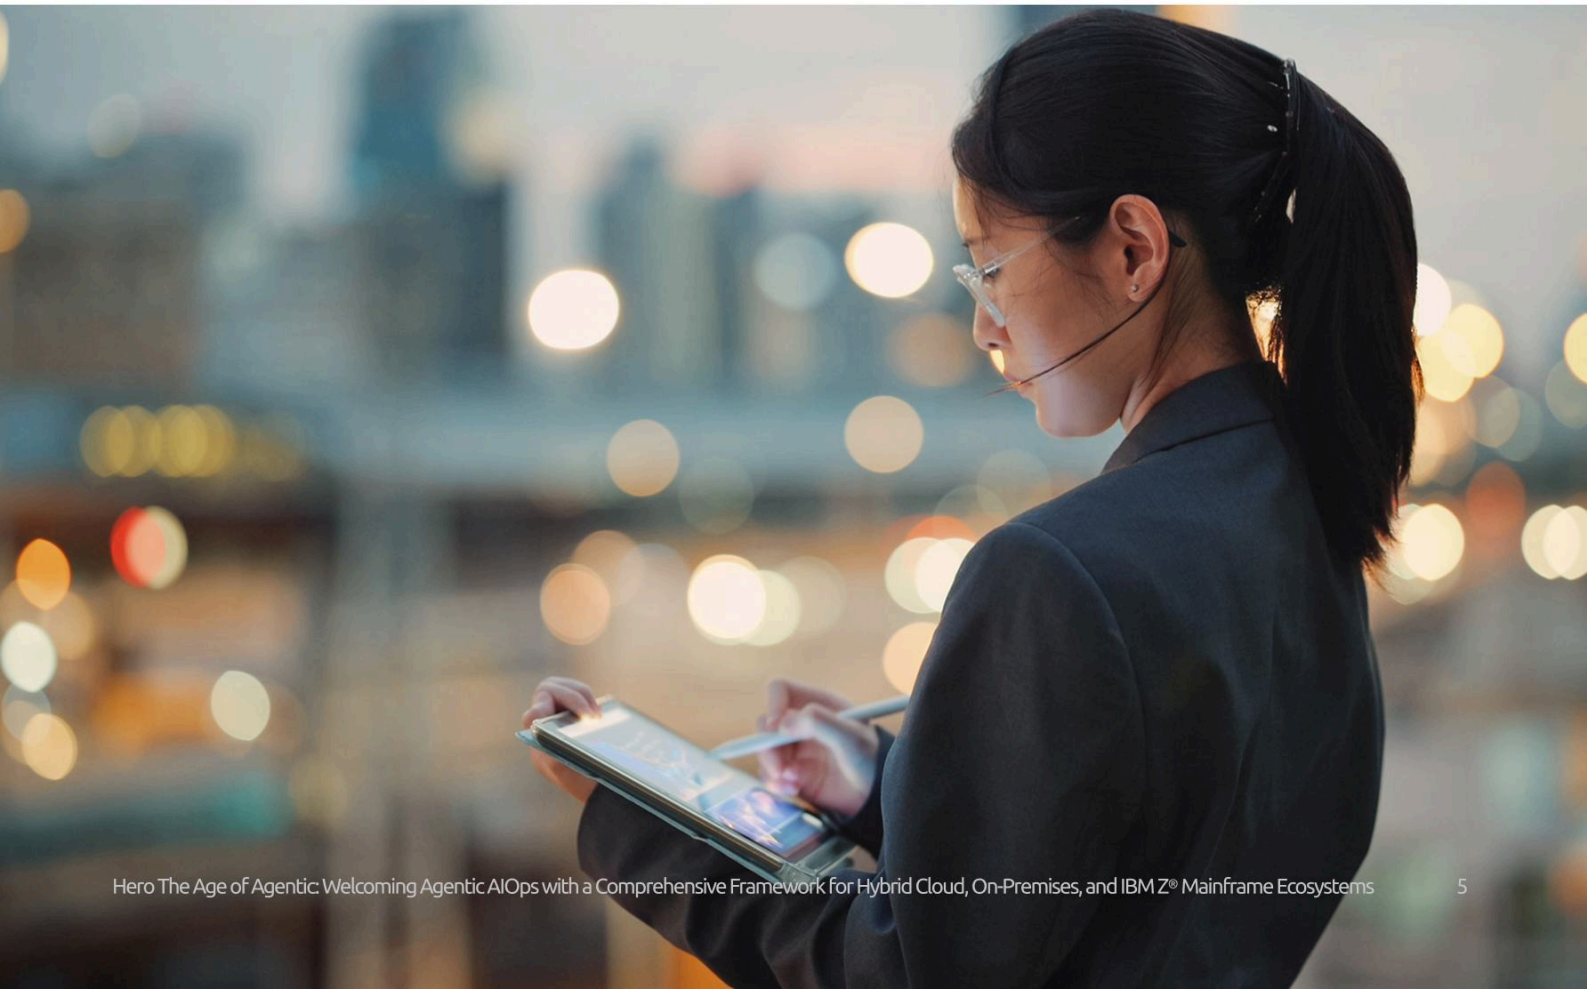

The AgenticOps Closed-Loop Operating Model

1. Signals and Telemetry Collection - Operational signals and telemetry from cloud platforms, networks, applications, and IBM Z systems are first collected into the Observability, Events, and Knowledge Fabric, where the raw data is correlated, enriched, and contextualized with enterprise knowledge.

2. Observability, Events & Knowledge Fabric - Within this fabric, all incoming data is normalized, infused with contextual insights, and prepared for secure downstream processing through the AgenticOps pipeline.

3. MCP Connector and Gateway Activation - The normalized signals are then transferred through MCP Connectors and exposed to the platform via the MCP Gateway, which applies the required policy controls, security measures, and access governance.

4. Enterprise Agent Reasoning - Enterprise Agents, powered by Watson Orchestrate, consume the governed context and use it to reason over system goals, policies, operational constraints, and current conditions.

5. Multi-Agent Coordination - The Multi-Agent Coordinator aligns the priorities of different agent , resolves conflicts among their recommended actions, and determines the most appropriate course of action.

6. Policy-Governed Execution - Once actions are approved, they are executed across hybrid and mainframe environments using policy-governed automation tools such as Vault, Terraform, Ansible, and Concert.

7. Hybrid and Mainframe Execution Paths - These governed actions run consistently across cloud, on-prem, and IBM Z environments, ensuring unified execution across all infrastructure layers.

8. Outcome Monitoring and Feedback Loop - Execution results are continuously monitored and fed back into the Agent Monitor and Vector Database, enabling agents to learn from outcomes, adapt their reasoning, and improve future decisions—ultimately closing the operational loop.

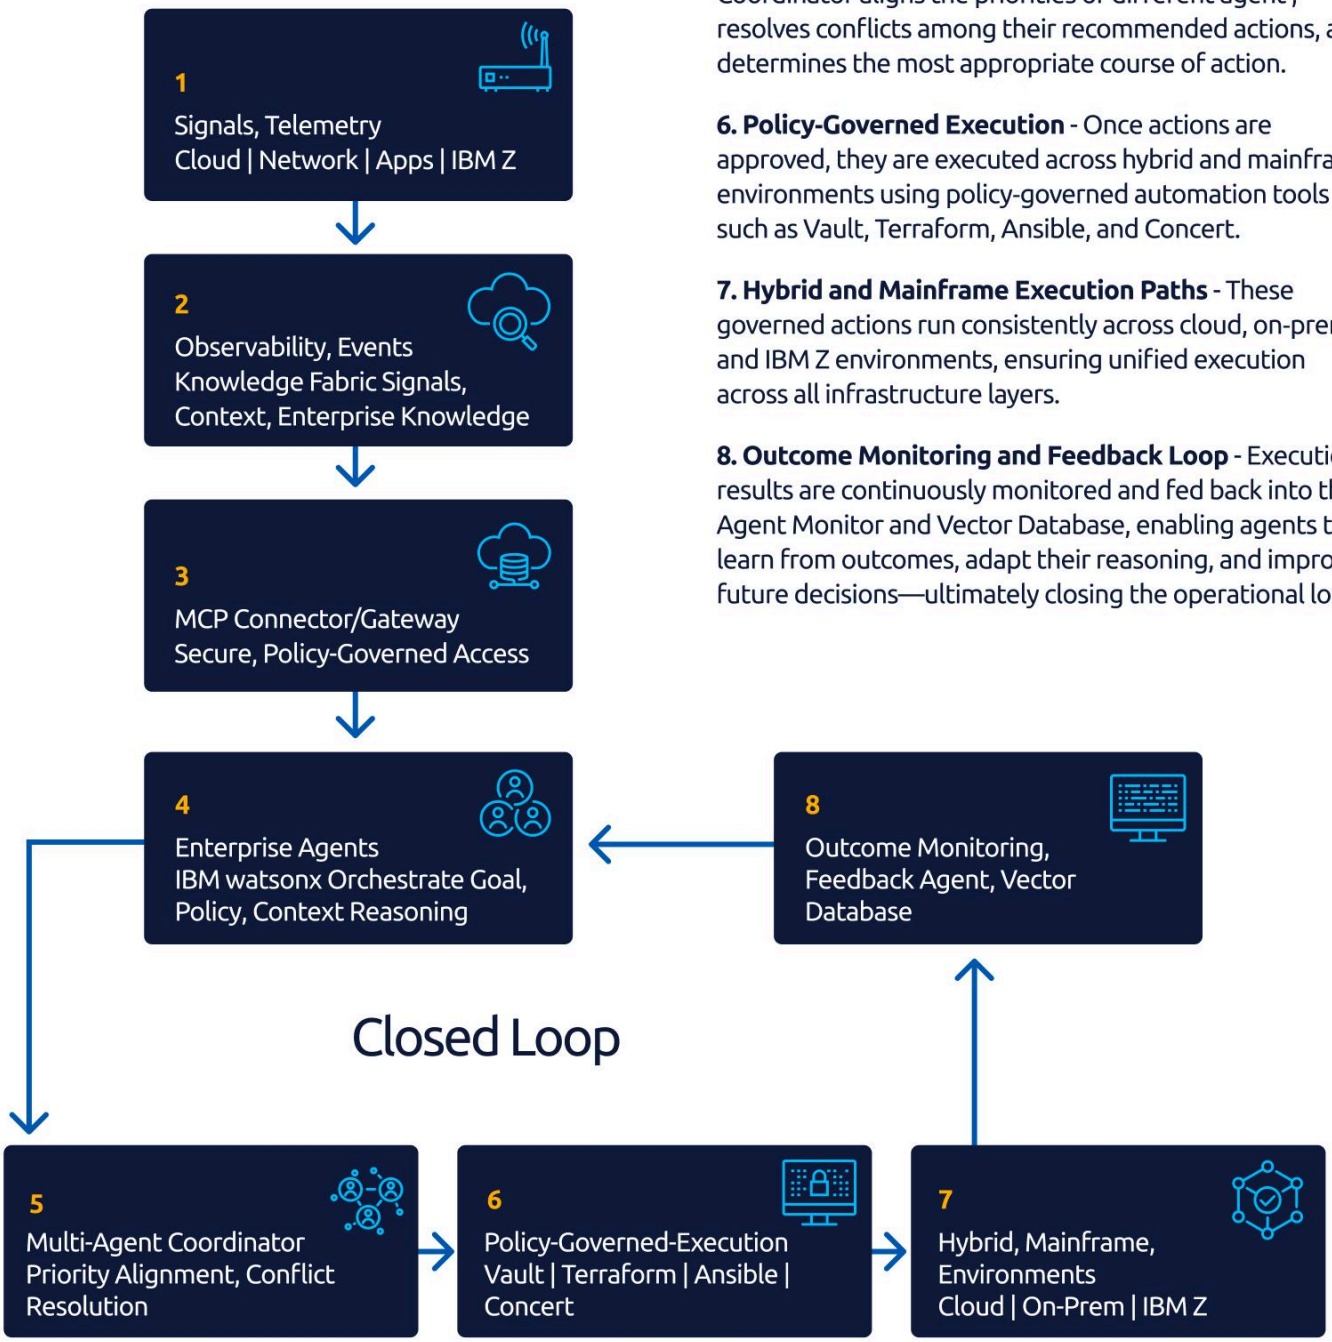

IBM Tooling Ecosystem for Agentic AIOps Implementation

IBM's revolutionary reference architecture is made possible by a specialized suite of IBM products, each of which helps it excel at reasoning, observability, and cross-platform management.

watsonx.ai® and the Granite® Model Family

If IBM's reference architecture forms an exciting new home for the future of enterprise IT, Wasonx.ai is the cognitive foundation upon which it sits. By providing access to frontier foundation models—including the IBM Granite family—watsonx.ai offers higher performance in code generation, an understanding of natural language prompts, and the ability to perform logical reasoning.

By including a unified developer studio for managing "AgentOps," watsonx.ai also provides tools for tracing, troubleshooting, and evaluating agent performance…so the entire lifecycle can be monitored, from experimentation to production.

watsonx Orchestrate®: The Orchestration Backbone

As an "all-in-one" solution for both building and managing AI agents, watsonx Orchestrate is a truly future-forward solution. Its "Agent Builder" function offers no-code interfaces for task-specific agent creation, alongside a pro-code Agent Development Kit (ADK) that makes it possible to build agents capable of performing more complex tasks.

With a catalogue containing hundreds of pre-built tools and domain-specific agents, watsonx Orchestrate can be applied to areas like IT support, procurement, and HR, giving you the power to scale your agent efforts without sacrificing the time it takes to build every integration from scratch.

IBM Concert®: Resilience and Proactive Risk Management

In the world of Agentic AIOps, IBM Concert is your AI-powered operations center. By actively unifying data that's typically siloed, IBM Concert provides you with an app-centric view of your overall resilience profile.

IBM Concert collects information from various sources—including everything from observability tools like Instana to ITSM platforms like ServiceNow, and everything in between—to generate a Resilience Score. The resulting shift from reactive firefighting to proactive governance means you spend less time dealing with outages by identifying the issues that lead to them before they even occur.

IBM Cloud Pak® for AIOps

IBM Cloud Pak for AIOps offers IT teams a modern platform that proactively detects, diagnoses, and resolves issues across hybrid environments. By leveraging machine learning and natural language processing, IBM Cloud Pak makes it easier than ever to correlate events, reduce alert noise, and pinpoint root causes…all in a matter of seconds.

Even better, the platform integrates seamlessly with both DevOps pipelines and ITSM systems, so you can enjoy a continuous feedback loop between development and operations.

watsonx.governance®: Ensuring Responsible AI and Compliance

As a comprehensive governance platform for the future of Agentic AIOps, IBM watsonx.governance accelerates responsible, transparent, and explainable AI workflows across an entire ecosystem.

By directing, managing, and monitoring AI activities, this platform more effectively governs generative AI and ML models from both IBM and third-party vendors…such as watsonx.ai, Amazon SageMaker, Google Vertex, and Microsoft Azure.

With core features like automated monitoring for model health, accuracy, drift, and bias—along with other key quality metrics like hallucination and relevance—watsonx Orchestrate brings a variety of powerful governance, risk, and compliance (GRC) tools to the table, streamlining audits and making regulatory adherence more efficient with customizable workflows featuring approvals, risk scorecards, dashboards, and automated reports.

By integrating with watsonx Orchestrate, watsonx.governance is able to apply key guardrails during agent coordination and decision-making, and can guarantee that multi-agent workflows are kept in compliance with your organization's policies. OpenPages can be leveraged for scalable GRC to enable risk assessment for agentic actions like remediation and scaling, while mainframe integration monitors AI inference on IBM Z® via the Spyre Accelerator to keep your sensitive data secure at all times.

In addition, a connection to IBM Cloud Pak for AIOps provides incident summarization governance, which effectively prevents the inadvertent use of biased or inaccurate AI outputs in operational resolutions.

Given that trust is a key component whenever AI or machine learning is in use, this tooling empowers organizations to both build and maintain trust in their autonomous systems, keeping compliance checks automated and providing traceable AI lineages that markedly reduce any risk for ethical lapse in today's hybrid environments. These automation features also extend to data sourcing and protection, integrating directly with the Knowledge Fabric layer to make sure data privacy policies are enforced during agent training and inference alike.

When these ethical guardrails are implemented by design, watsonx.governance is able to proactively manage risks in IBM Concert, incorporating governance insights thanks to accurately generated resilience scores. This effectively transforms Agentic AIOps from reactive intelligence to governed autonomy so organizations can foster innovation with increased confidence.

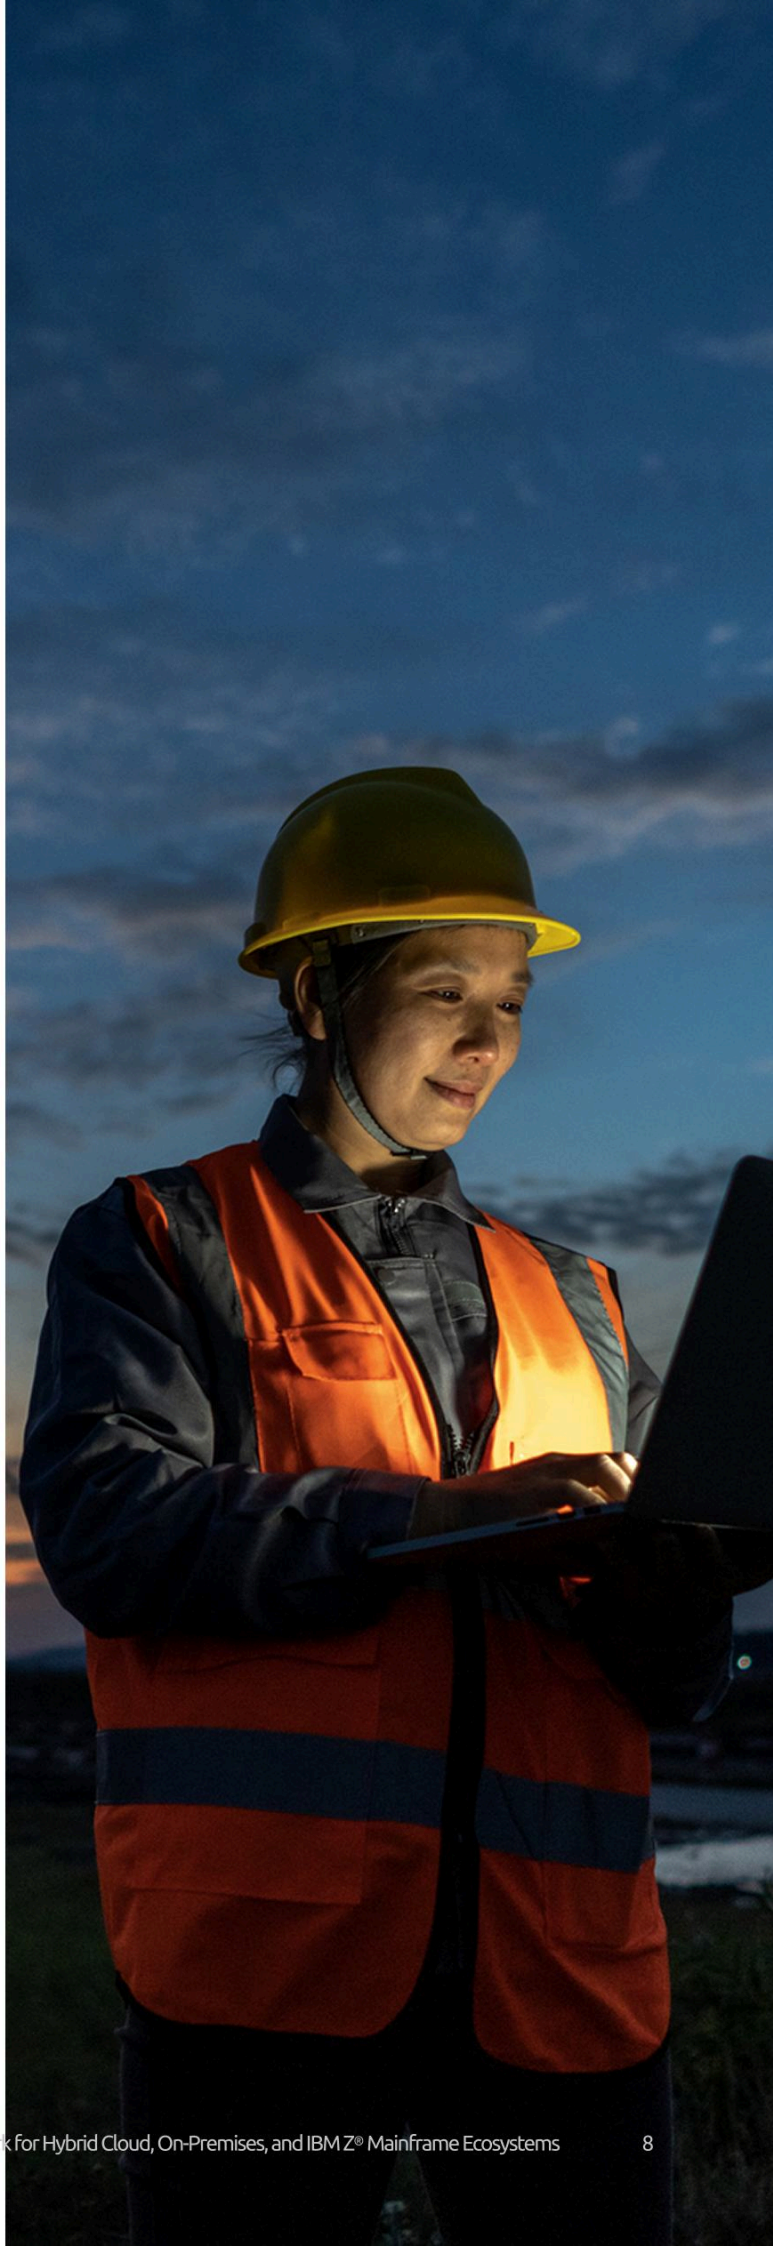

Integrating the IBM Z® Mainframe into the Agentic Ecosystem

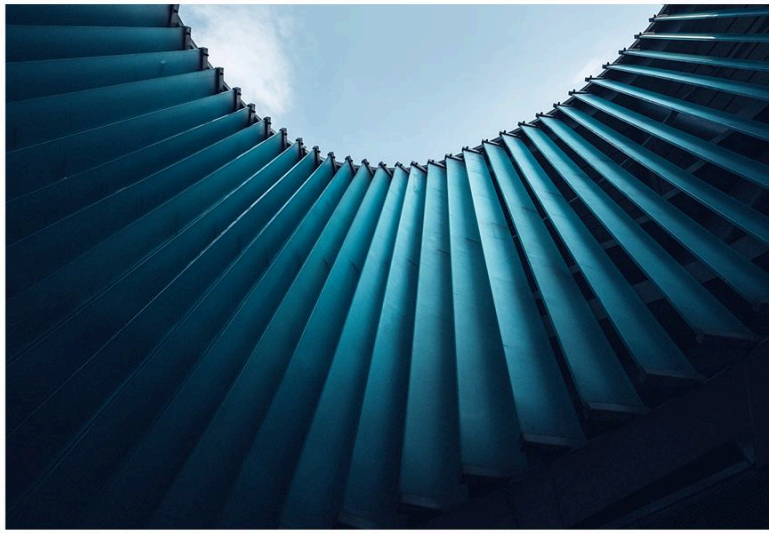

Of the many advantages that come with IBM's agentic framework, one of the most significant is its ability to bring modern AI capabilities directly into the mainframe. In the agentic ecosystem, IBM Z goes from a siloed component to a full participant in the hybrid cloud operational model, offering the same benefits that would be found with intelligent automation as a distributed system.

watsonx Assistant® for Z and Specialized Mainframe Agents

By providing a conversational AI interface, watsonx Assistant for Z makes operation and system programming simpler and easier than it was in the past. With its ability to interpret natural language queries and provide actionable insights into the system's overall health and performance, the Assistant is able to guide junior staff through complex troubleshooting processes with ease.

The Assistant facilitates these interactions with a suite of specialized mainframe agents:

- The **Support Agent** enables the execution of Ansible playbooks for tasks like taking and sending z/OS dumps.
- The **OMEGAMON Insights Agent** provides access to and an analysis of various performance metrics directly through a chat interface.
- The **CICS Agent** offers direct assistance with the troubleshooting of transaction error codes and understanding CICS topology.
- The **Upgrade Agent** leverages z/OSMF APIs to provide precise responses and guidance for z/OS version migrations.
- **IBM Z IntelliMagic Vision Agent** – Provides advanced RMF/SMF-driven performance analytics, I/O health scoring, storage path intelligence, and early-warning insights to enhance mainframe operational decision-making.

On-Platform AI Inferencing with IBM Spyre™ Accelerator

In the interest of maintaining data privacy and minimizing latency, IBM has given organizations the ability to run agentic AI workloads natively on the mainframe. Powered by the Telum II processor, the Spyre Accelerator makes high-performance, on-platform AI inference possible, allowing agents to access sensitive operational data like logs, system metrics, and configuration details within the highly secure boundaries of the IBM Z environment.

Component	Role in Mainframe AIOps	Integration Point
zRAG	Augments LLMs with IBM Z documentation	watsonx Assistant for Z
System Topology	Interacts with z/OSMF APIs for environment mapping	z/OSMF, watsonx Assistant for Z
Identity Management	Enforces secure access to mainframe APIs	z/OSMF, OMEGAMON
OMEGAMON Integration	Streams mainframe metrics to open AIOps platforms	Instana, Cloud Pak for AIOps
IBM Concert for Z	Manages APAR tracking and maintenance risk	IBM Z Software stack
IBM Z IntelliMagic Vision	Advanced RMF/SMF analytics, I/O scoring, capacity insights, and performance modeling	Mainframe Agent, Cloud Pak for AIOps

Shifting Left: Agentic AIOps in Development and QA

Production environments aren't alone in their ability to enjoy the benefits of the age of agentic AI. By "shifting left," organizations are now able to plug intelligent automation directly into the earliest stages of the software development lifecycle (SDLC), from coding to unit testing, functional verification, and beyond. With this proactive approach, developers can now identify and address defects earlier in the cycle, when they're significantly less costly and easier to resolve.

The Economic Rationale for Shift-Left AIOps

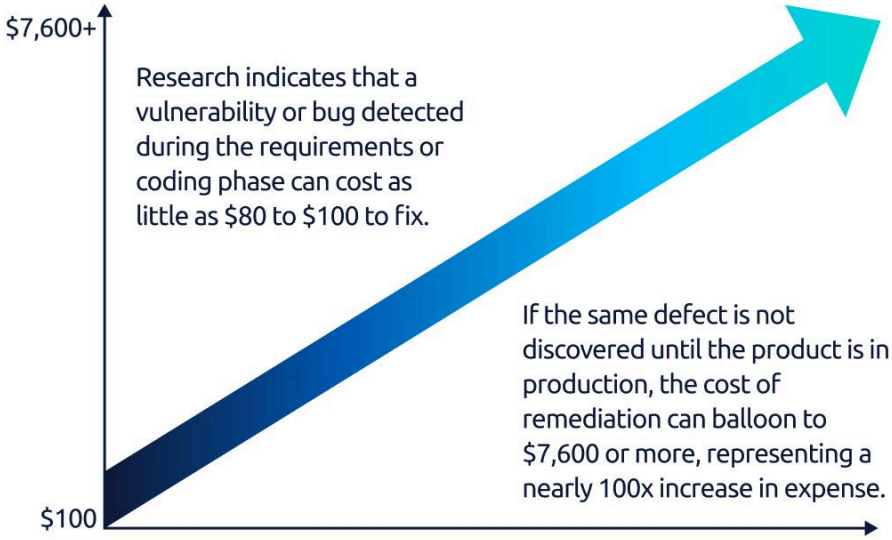

Agentic AI facilitates this by acting as a "crystal ball" for developers, analyzing code patterns and historical data to predict where problems are most likely to emerge.

When watsonx.governance is incorporated into shift-left practices, organizations are able to ensure that the AI models being used in development and QA are properly governed from the very beginning of the process. During code analysis with Granite models or test generation via the Functional Testing Agent, for example, watsonx.governance actively monitors for any bias in training data, evaluating model accuracy and proactively preventing defects from propagating.

Integration with Instana further applies compliance checks to CI/CD pipelines, while also flagging risks like data drift as early as possible…all of which not only cuts down on the cost of remediation but also ensures that ethical AI standards are embedded throughout the pipeline, aligning shift-left automation with your organization's established governance policies.

Agentic Test Automation with IBM Test Accelerator for Z®

With the IBM Test Accelerator for Z, organizations can now introduce AI to the functional testing process. This "Functional Testing Agent" interprets natural language to automatically generate executable Java-based Galasa tests. By scanning both MFS (Message Format Service) and BMS (Basic Mapping Support) project files, the agent can identify potential use cases and bulk-generate test cases that reflect your intended business logic, cutting back on the need for manual scripting while simultaneously accelerating the speed at which comprehensive test buckets can be created.

Continuous Testing and CI/CD Observability

Thanks to an integration with Instana Observability, organizations can now enjoy real-time visibility into the behavior of various applications throughout the testing phase. While legacy monitoring tools relied on sampling, Instana captures snapshots with a single-second granularity, giving developers the ability to see immediate upstream and downstream effects created by their code changes.

This ability to enjoy "continuous testing" ensures that regressions can be identified well before they get the chance to impact the broader system.

The Economic Rationale for Shift-Left AIOps

SDLC Phase	Agentic AIOps Activity	IBM Tooling involved
Requirements/Design	Automated validation of business logic in test cases	watsonx Assistant, Test Accelerator
Coding/Unit Testing	Static code analysis and predictive bug detection	watsonx.ai, Granite models
Functional Testing	Automated generation of Galasa test cases	Functional Testing Agent, Galasa
Integration/Staging	Real-time observability and anomaly detection	Instana, Turbonomic
Production	Autonomous remediation and performance optimization	Concert, Cloud Pak for AIOps

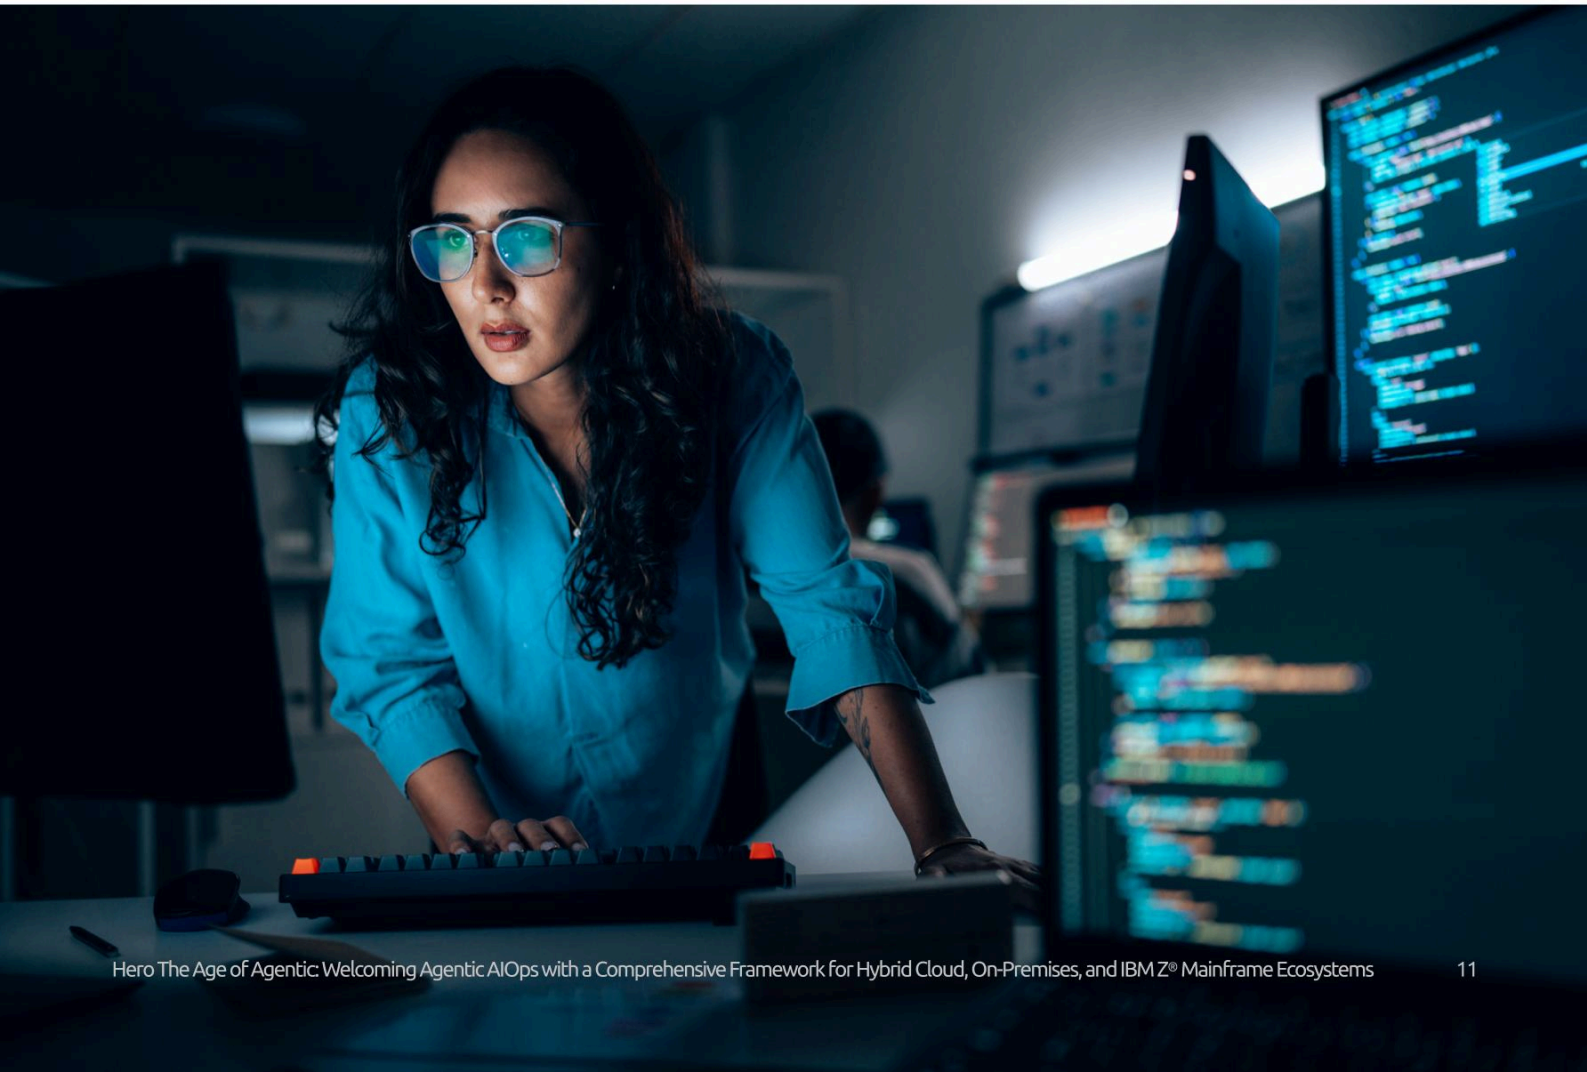

The Human-in-the-Loop Angle: Governance and Accountability

The future of autonomous AI is an exciting one, but we still need to rely on the occasional human for various tasks. A "Human-in-the-Loop" (HITL) framework enables operators to maintain trust, safety, and accountability…which is required for many mission-critical systems in an enterprise context. With this framework, however, the human operator's role goes from fighting fires and performing manual tasks to that of a "trusted approver" and "AI trainer," providing strategic oversight and handling complex scenarios when they arise.

Mechanisms for Human Intervention

Within the IBM framework, HITL is implemented through structured architectural patterns that incorporate key human feedback at critical stages of the AI workflow with LangGraph. This tool makes persistent execution states possible, so a workflow can be paused until human input is received.

There are two primary approaches to HITL intervention:

Static Interrupts

In this instance, the system is designed to pause at predetermined points in a given workflow—such as before executing a major configuration change or after generating a complex remediation plan—to wait for a human's approval.

Mechanisms for Human Intervention

In these instances, the agent itself will decide to interrupt the workflow, requesting human assistance according to the current state of the environment or its own level of confidence in a proposed action.

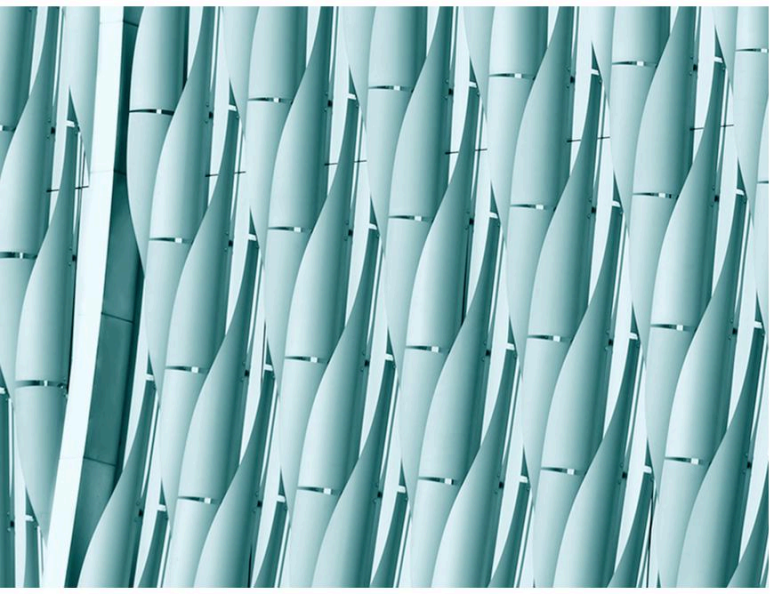

An integration with watsonx.governance enhances this framework even further, ensuring accountability by keeping all interventions logged in customizable dashboards to support compliance audits and the continuous improvement of agent behavior.

The SRE as the Strategic Governor

While the agentic ecosystem is the ideal vessel for sailing into the future of AIOps, it's still good to have a few human hands on deck to help steer the ship. In an agentic ecosystem, Site Reliability Engineers (SREs) interface with the AI through simple, conversational instant messaging platforms like Slack or Microsoft Teams.

Any time an incident occurs, the agent will summarize its details in plain language, presenting a "remediation card" that lays out a list of proposed actions for the SRE to review. Tools like Instana validate the AI's reasoning path, which can then be reviewed by the SRE before a fix is either approved or modified. By making sure humans remain the final authority, AI is able to handle these cognitive burdens with minimal oversight and maximum accuracy.

Business KPIs and Outcomes of Agentic AIOps

The transition to Agentic AIOps changes more than how the development and production pipelines operate: It also radically reimagines the way operational success is measured. Traditional Key Performance Indicators (KPIs) like uptime and latency are still relevant, but they don't adequately capture the strategic value of autonomous intelligence. To keep pace with the future of agentic AIOps, we need outcome-driven metrics that help us understand strategic ability, innovation, and autonomous value creation.

watsonx.governance makes it possible to track a series of AI compliance rates whose metrics complement both autonomy and success rates by quantifying the trust inherent to your organization's agentic operations.

Technical Performance and Efficiency Metrics

At an operational level, organizations now find themselves focused on tracking the quality of agent decision-making and execution.

- **Autonomy Rate:** The percentage of operational issues resolved by agents without requiring human intervention.
- **Mean Time to Resolution (MTTR) Reduction:** A primary goal of Agentic AIOps is to dramatically accelerate incident response. Tools like Instana have demonstrated the potential to reduce MTTR by up to 80%.
- **Agent Success Rate:** The ratio of successful goal completions to total attempts, with a benchmark of over 85% typically required for production-grade agents.
- **Hallucination Rate:** A critical safety metric, specifically for generative AI agents, which must be kept below 2% through continuous validation and guardrail testing.
- **Deflection Rate:** In an internal support context, the percentage of queries handled fully by the agent, freeing up human teams for more strategic work.

Developer Productivity and Experience (DX) KPIs

Shifting AIOps left introduces a new set of metrics, which now focus on the efficiency of the software delivery pipeline.

- **Lead Time for Changes:** The interval from the first commit to production deployment. High-performing teams use agentic automation to identify and remove bottlenecks in CI/CD pipelines.
- **Deployment Frequency:** A measure of organizational responsiveness and the degree of automation in the release process.
- **Change Failure Rate (CFR):** The percentage of deployments that result in production issues, serving as a proxy for the effectiveness of agentic testing and quality checks.
- **Flow Efficiency:** The ratio of active work time to the total elapsed time of a task, highlighting where work is stalled by manual handoffs or context switching.

Strategic Business Outcomes

Ultimately, organizations must be sure that their investment in Agentic AIOps is justified by its contribution to the organization's bottom line and long-term competitiveness.

- **Operational Cost Savings:** Tangible reductions in staffing needs, third-party service spend, and emergency patching costs.
- **Revenue Lift:** Opportunities that would have been lost due to downtime or bad data but were preserved by proactive agentic intervention.
- **Risk and Compliance Reduction:** Fewer regulatory penalties and audit flags attributable to the always-on governance and audit-readiness of the agentic framework.
- **Time-to-Market Acceleration:** The speedup in product delivery enabled by automated testing and environment provisioning, allowing the business to respond faster to market demands.
- **Enhanced Risk Mitigation through Governance:** organizations achieve measurable reductions in compliance risks, such as fewer audit failures and regulatory fines. It contributes to revenue lift by ensuring AI-driven optimizations are ethical and reliable, preserving customer trust. In shift-left scenarios, it accelerates time-to-market by automating governance approvals, potentially reducing SDLC cycles by 20–30% while embedding accountability.

Specific information regarding each of the preceeding pages was provided by the following authors:

Prabhakaran Arivalagan

SRE Lead Architect @ Capgemini

Prabhakaran Arivalagan is a SRE Architect at Capgemini, specializing in site reliability engineering, AIOps, cloud and platform engineering. With 19 years of experience in the field of IT operations and artificial intelligence, Prabhakaran has a deep understanding of AIOps, machine learning, observability, telemetry, DevOps, platform engineering, automation, and its applications in modern IT environments.

Prabhakaran has been instrumental in implementing and optimizing SRE solutions using tools such as IBM Instana, Splunk, AppDynamics, Dynatrace, Azure AppInsight, DevOps, Terraform and other related technologies. He has developed solution accelerators such as Dashboard as a Code (DaaC), Synthetic Monitoring Framework (SMF) and Infrastructure as Code (IaC) framework for rapid adoption.

John Jabez John

SRE Lead Architect @ Capgemini

An accomplished SRE Lead Architect with over 18 years of comprehensive IT experience, encompassing 7 years in Site Reliability Engineering (SRE), 8 years in performance engineering, and 3 years dedicated to development and production support. Renowned for exceptional organizational, communication, management, leadership, and problem-solving skills, the author specializes in SRE enablement, maturity assessments, transformation strategies, and operational excellence.

Proficient in leveraging advanced Application Performance Management (APM) tools such as Dynatrace, NewRelic, Instana, AppDynamics, and Splunk, as well as cloud platforms including IBM, Azure, AWS, PCF/Tanzu, and GCP. The author has extensive database experience that includes DB2, SQL DB, Oracle and other cloud DB solutions. Skilled in multiple programming languages, including Java (Spring Boot) and COBOL, the author also has expertise in mainframe environments, utilizing technologies such as IMS DC, REXX, JCL, and CICS.

With a strong focus on observability and resilience engineering, the author drives innovative SRE GenAI solutions that enhance system reliability and efficiency.

Ajay Walgude

Global Head of FS IBM Hybrid Cloud CoE @ Capgemini

Ajay Walgude who envisions big ideas and takes action to bring them to life. He leads IBM and Red Hat Cloud COE for Financial Services globally, and is the Global Senior Executive, IBM Champion, an Architect, a Tester, and Practice Leader helping customers achieve large outcome driven transformations across Banking and Insurance customers.

Anand Awasthi

AIOps & Automation Architect @ IBM

Anand Awasthi is an AIOps and Automation Architect at IBM, specializing in transforming enterprise IT operations using AI-driven and agentic approaches. With strong experience across AIOps, observability, and automation, he focuses on building intelligent, resilient, and self-healing systems in hybrid cloud environments.

He has deep expertise in leveraging platforms such as Instana, Turbonomic, Concert, Cloud Pak for AIOps, and watsonx Orchestrate, along with tools like Terraform, Vault, Ansible, Event Streaming, and AI & MCP Gateway to drive operational efficiency and reliability. Anand has been instrumental in designing and implementing solutions that integrate telemetry, event correlation, and AI-powered insights to enable proactive issue detection and faster root cause analysis.

Actively contributing to Agentic AIOps initiatives, Anand works on building autonomous workflows and AI agent-driven solutions that enhance system resilience and reduce operational overhead, helping enterprises evolve from reactive operations to predictive and self-optimizing ecosystems.

Key Takeaways for Enterprise Leadership

The journey toward Agentic AIOps is less a case of simple tool implementation and more a strategic transformation that is best realized with a series of concrete steps.

1.
A Foundation of Unified Observability

Autonomous agents are only as effective as the data they consume. Investing in high-fidelity, real-time observability via Instana and SevOne is a non-negotiable first step.

2.
Contextual Grounding via RAG

To ensure continued reliability and trust, agents must be grounded in enterprise-specific knowledge. Leveraging IBM's zRAG and vector database capabilities is essential for minimizing AI hallucinations and keeping reasoning accurate.

3.
Governance by Design

Human-in-the-loop safeguards must be architected into the system from the beginning. Utilizing LangGraph for structured interrupts and maintaining transparent reasoning paths keeps the organization in control of its autonomous systems.

4.
A Holistic Shift-Left Strategy

The greatest economic gains of AIOps are found in the early stages of the SDLC. Organizations should prioritize the deployment of agentic testing tools like IBM Test Accelerator for Z to dramatically reduce the cost of quality.

5.
Standardized Integration Surfaces

Adopting the Model Context Protocol (MCP) ensures that agentic systems can scale across the entire infrastructure stack without being hindered by brittle, proprietary integrations.

As enterprises move into the future, Agentic AIOps are rapidly becoming the central nervous system of the digital business, capable of not only maintaining the health of the infrastructure but also proactively optimizing it to drive superior customer experiences and business growth.

By combining IBM's specialized AI hardware, frontier foundation models, and comprehensive orchestration software, organizations can now enter the age of Agentic AIOps with a secure, scalable, and open environment for this next era of IT operations.

Capgemini

IBM Platinum Partner