
Agentic AI Operational Intelligence Platform for Autonomous Operations

Whitepaper

Shailesh Manjrekar
Chief Marketing Officer, Fabrix.ai



Contents

‣ Executive Summary	03
‣ Use Cases	04
‣ Architecture Pillars	05
‣ Agent Orchestration & Lifecycle Management	06
‣ AI Persona-based Agentic Workflow Orchestrator	07
‣ AI Personas	
‣ AI Persona based Governance Guardrails	08
‣ Managing Data Access for AI Personas	09
‣ Managing Action Privileges for Agents	09
‣ MCP enabled Data Fabric, Automation Fabric and Model Registry	09
‣ Visibility and Observability of Agents	09
‣ Agent Quality Control and Assurance	10
‣ Fabrix.ai Agentic AI Platform - best of breed consumption models	10

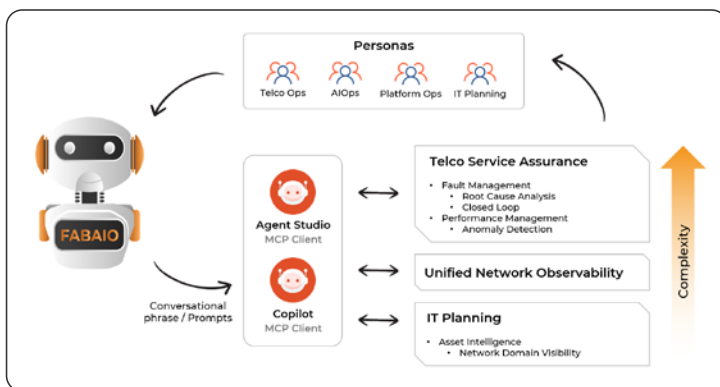
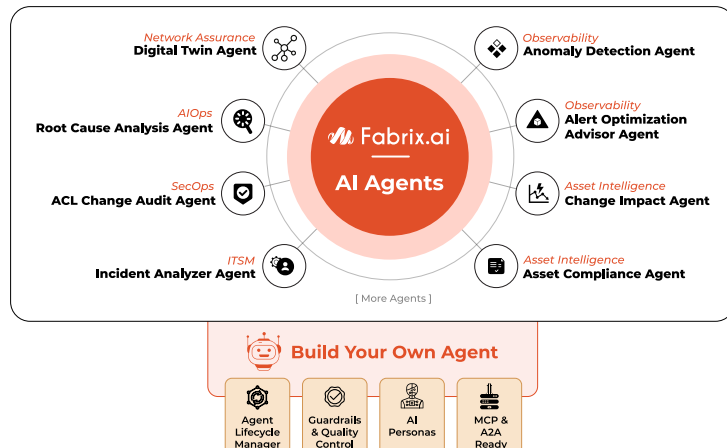
Executive Summary

Fabrix.ai delivers end to end enterprise AI agents for Autonomous IT Operations. These cross-domain AI agents deliver solutions for Telco Service Assurance, Asset Intelligence, Unified Network Observability, AIOps, Log Intelligence and Network Digital Twins on a single platform, targeted for Telcos, MSPs and multi-tenant enterprises.

- ▶ **Agentic AI as a Differentiator:** The core innovation is the use of AI agents that can perform actions autonomously, unlike previous systems that primarily notify users or showcase dashboards. The agents are designed to “sense, reason, plan, and act.”
- ▶ The platform aims to provide a more flexible and comprehensive solution compared to previous rule-based and machine learning-centric approaches.
- ▶ Fabrix’s Platform circumvents the challenges of stochastic and non-deterministic processes to deliver an enterprise-ready solution with robust guardrails and security features.
- ▶ One of the biggest differentiators for Fabrix.ai is the ability to work with real-time data. That’s something that not all automation vendors can do today. What makes automating operations challenging is alerts and events coming in real time. IT professions struggle to take immediate action on these as the analytics platforms can’t process the underlying telemetry to provide a solution in real time. Its focus on handling real-time information has enabled it to get traction in key verticals, especially telco.
- ▶ Fabrix.ai quickly enables launching of new autonomous, multi-vendor services with our solutions, providing joint customers with -
 - ▶ **Improved Predictability** - preempt outages and degradations with improved operational readiness
 - ▶ **Lower Cost** - Incident reduction by 90%+
 - ▶ **Increased Reliability** - No missed critical alerts and 50%+ faster resolution
 - ▶ **Increased Productivity** - Automate as many tasks with pipelines, bots and AI

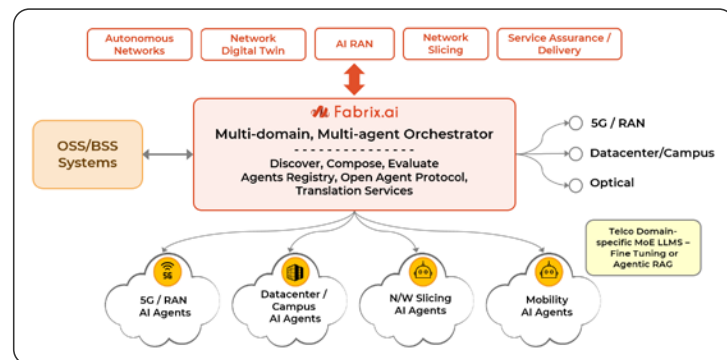
Use Cases

The Agentic AI platform caters to multiple Operational Intelligence use cases across ITOps, IT Planning, NOCops, SecOps and AIOps domains. The platform also empowers these personas to build their own agents by providing simple conversational phrases.



The platform provides the freedom to Operational personas to quickly create bespoke services using custom AI agents, by just providing conversational phrases. The Co-pilot provides visibility into the different domains and their assets and interdependence and get insights and analytics for the unknown questions and hypotheses. Once you have enough insights, these personas can create, test, deploy an AI agent for a specific use case as depicted in the image

Below shows an example of a TelcoOps agent for AI based Telco Service assurance & Delivery across different domains.



This image shows a Platform Composable Storyboard for “Anomaly Detection agent”. The storyboard provides all details pertaining to Observability, Explainability and Responsibility.

Architecture Pillars

The platform employs three key pillars, working in tandem to accomplish outcomes-

AI Fabric

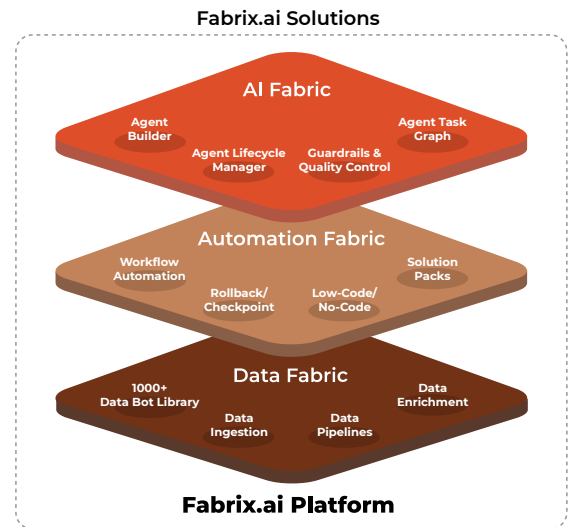
Is an AI agent-driven cross-domain, distributed orchestrator that enables customers to securely build, deploy, and manage Agents' lifecycles, ensuring guardrails and quality controls. It integrates with disparate large and small models, curated datasets, and automation to drive Agentic Workflows.

Automation Fabric

Is an outcome-driven Agentic Workflow framework that integrates Agents, Automation, and Data to build Agentic Workflows. It is dynamic and extensible and can also integrate with other third-party engines like Cisco BPA, NSO, Redhat Ansible, Terraform, or Camunda.

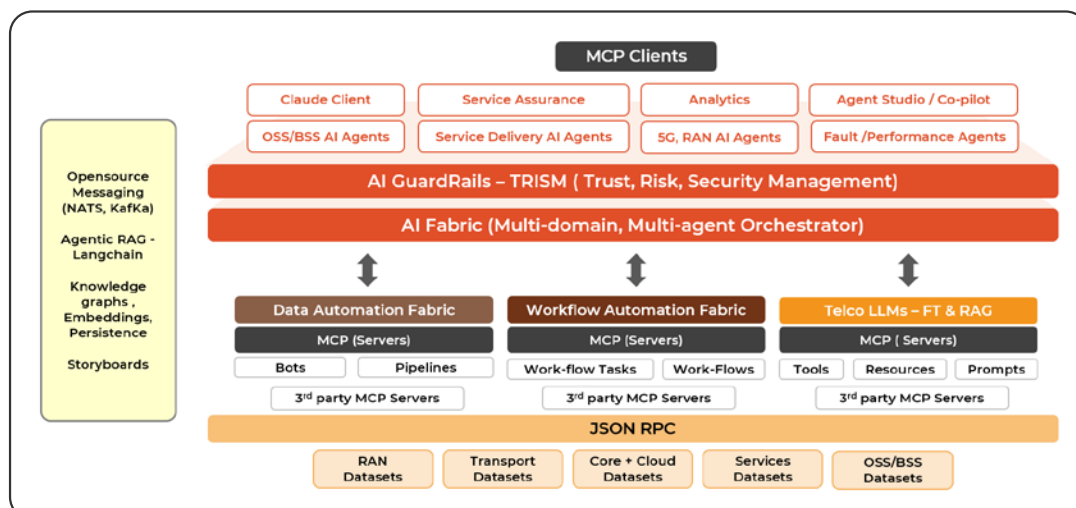
Data Fabric

Robotic Data Automation Fabric (RDAF) is a semantic-based data fabric that provides data integration with 1000+ data bots, data ingestion, data transformation, enrichment, and data routing, using Telemetry pipelines to your choice of source and destination. The Data Fabric bots, pipelines and blueprints are entirely composable and can be changed on the fly.



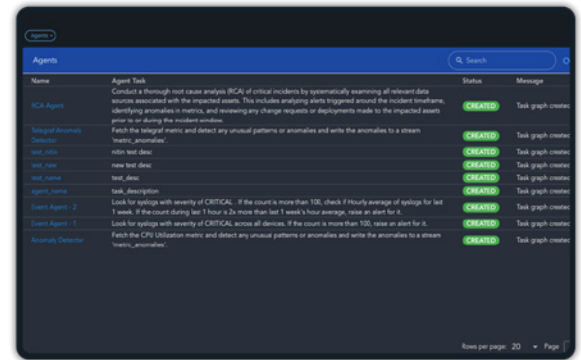
Key tenets of Fabrix's Agentic AI platform include -

- ▶ **AI Agent Studio** - Agent Orchestration and Lifecycle Management
- ▶ **AI Co-pilot** - Generative AIOps
- ▶ **AI Fabric** - Multi-agent Orchestrator using Google A2A (Agent to Agent) protocol
- ▶ **MCP (Model Context Protocol)** enabled Data Fabric, Automation Fabric and Model Registry
- ▶ **AI Guardrails** - Managing Data & Action Privileges for Agents
- ▶ **Agent Quality Control and Assurance**
- ▶ **Visibility and Observability of Agents**

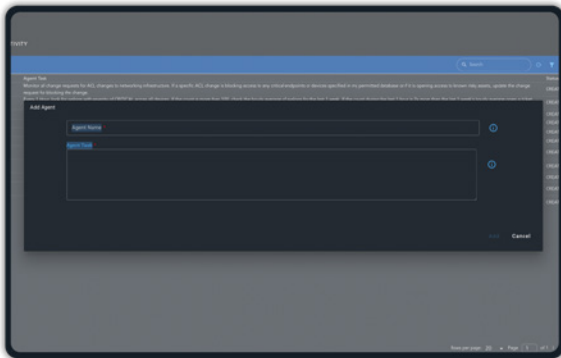


AI Agent Studio

AI Agent Studio provides Multi-persona, Multi-domain, Multi-agent administration for the platform. AI personas can build, deploy, and orchestrate AI agents for out-of-the-box use cases like anomaly detection and event intelligence using Agent Studio. This single pane of glass is used to get a comprehensive view of AI agents, A2A servers and MCP servers.



Name	Agent Task	Status	Message
SCA Agent	Conduct a thorough root cause analysis (RCA) of critical incidents by systematically examining all relevant data sources associated with the impacted assets. This includes analyzing alerts triggered around the incident timeline, identifying anomalies in metrics, and reviewing any change requests or deployments made to the impacted assets prior to or during the incident window.	CREATED	Task graph created
Intelligent Anomaly Detector	Fetch the historical metrics and detect any unusual patterns or anomalies and write the anomalies to a stream metrics_anomaly.	CREATED	Task graph created
test_sca	write test data	CREATED	Task graph created
test_run	read test data	CREATED	Task graph created
test_clean	delete test data	CREATED	Task graph created
Agent - metrics	Task description	CREATED	Task graph created
Event Agent - 2	Look for outages with severity of CRITICAL. If the count is more than 100, check if hourly average of outages for last 1 week. If the count during last 1 hour is 2x more than last 1 week's hour average, raise an alert for it.	CREATED	Task graph created
Event Agent - 1	Look for outages with severity of CRITICAL across all devices. If the count is more than 100, raise an alert for it.	CREATED	Task graph created
Anomaly Detector	Fetch the CPU utilization metrics and detect any unusual patterns or anomalies and write the anomalies to a stream metrics_anomaly.	CREATED	Task graph created

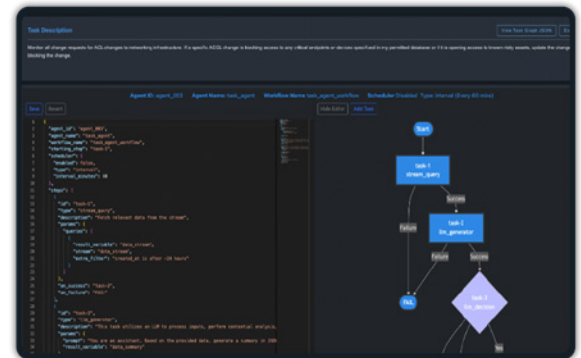


Agent Orchestration & Lifecycle Management

Agentic AI within the Fabrix.ai platform empowers users to create agents by either defining new task descriptions or utilizing predefined templates. The Agentic AI framework employs Large Language Models (LLMs) to translate these descriptions into a task graph, which outlines the sequence of tasks, dependencies, and conditions for the agent.

The Agentic AI orchestrator executes the task graph, while the AI Guardrails module ensures that the AI adheres to the defined task and doesn't deviate. The task graph typically encompasses the following node types:

- ▶ **Query Node:** Executes queries against the Data Fabric to retrieve data or aggregations.
- ▶ **Generation Node:** Employs LLMs to generate queries, summaries, insights, and recommended actions.
- ▶ **Decision Node:** Utilizes LLMs to make decisions based on data and previous generation results, determining the next steps within the graph.
- ▶ **Action Node:** Leverages the Automation Fabric to perform automated actions, such as opening tickets, sending notifications, triggering other agents, or executing data retrieval and remediation tasks.

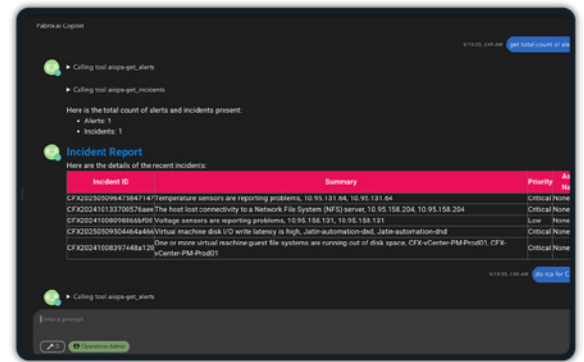


The Agent Lifecycle management encompasses the following actions:

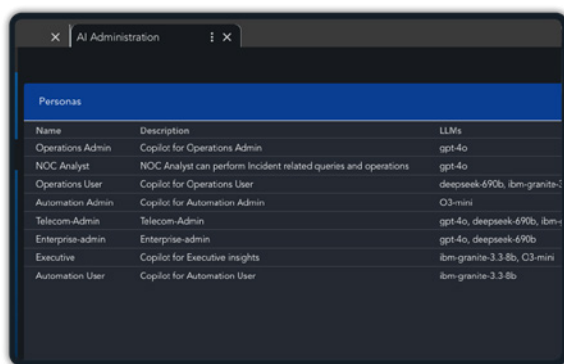
- ▶ **Review:** Enables users to review and refine the task graph to align with objectives.
- ▶ **Dry Run:** Allows users to simulate agent actions without actual execution to preview outcomes.
- ▶ **Test:** Facilitates agent execution for a limited duration or iterations on specified test data to verify outcomes.
- ▶ **Deploy:** Enables agent deployment onto the platform, either on a schedule or triggered by events.
- ▶ **Decommission:** Takes agents offline.

AI Co-pilot - Generative AIOps

The AI Co-pilot employs an MCP client and can provide valuable insights for unknown unknown queries using simple conversational phrases



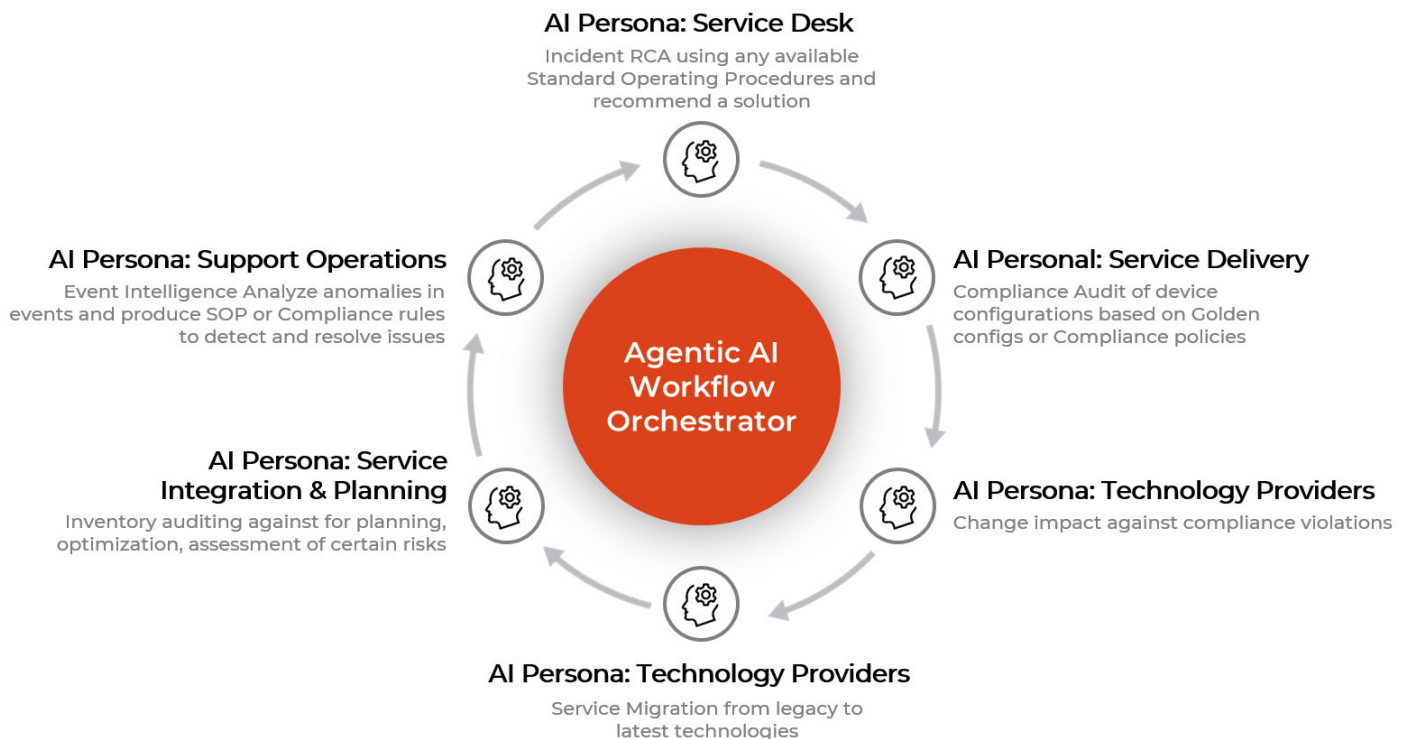
AI Persona-based Agentic Workflow Orchestrator

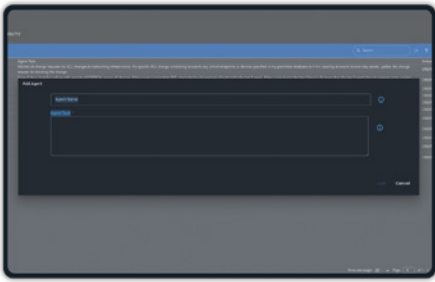


AI Personas

These are operational personas will have different access levels for data access, MCP Tools as well as ability to run appropriate agents. This is extremely crucial to ensure erroneous personas are not able to run erroneous agents resulting in unintended consequences

One of the key capabilities of the platform is to quickly create agentic workflows using graphical visualization and prompt templates. The platform provides certain workflows for operational personas, but can be adapted to any personas and agentic workflows.



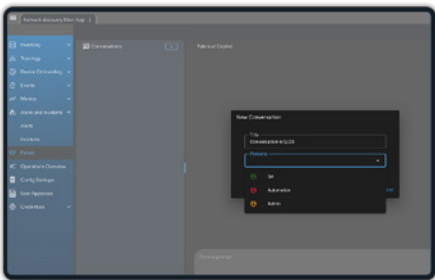
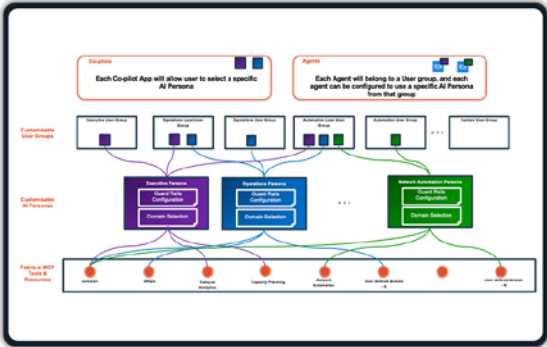


This image shows some standard prompt templates used with these agentic workflows. Personas can add to these templates to suit their use cases.

AI Persona based Agentic Governance Guardrails

The platform governs AI agents with AI Persona-based Guardrails, providing checks and balances for access to tools, data, and models as per user groups and domains.

AI-based agents, designed to perform specific tasks autonomously, necessitate guardrails to prevent unintended or harmful actions. This is especially crucial for agentic AI platforms that enable users to create new agents by describing the desired task.



AI-based agents, designed to perform specific tasks autonomously, necessitate guardrails to prevent unintended or harmful actions. This is especially crucial for agentic AI platforms that enable users to create new agents by describing the desired task.

AI Guardrails can be categorized into three major areas:

1. Standard or common sense guardrails:

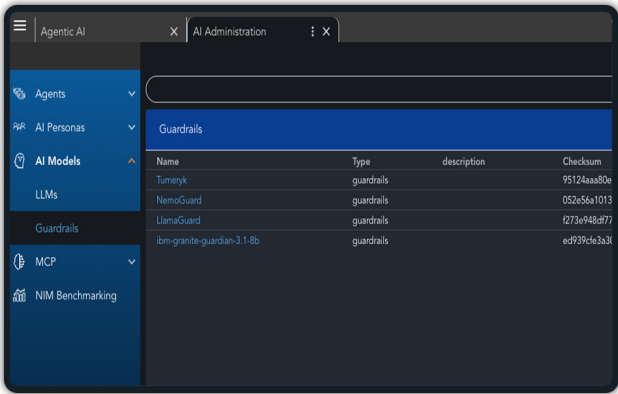
These are fundamental guardrails designed to prevent the AI from generating outputs that contain profanity, use language unsuitable for a professional setting, or promote harmful or illegal activities.

2. Industry / Domain specific guardrails:

These guardrails are tailored to specific industries or domains. They encapsulate industry best practices and ensure that the AI's output complies with relevant regulations. For example, an AI agent in the healthcare sector would need to adhere to patient privacy laws.

3. Enterprise / Environment specific guardrails:

These guardrails are specific to an individual enterprise or environment. They enforce company policies and values, ensuring that the AI's output aligns with the organization's culture and standards.



While some AI models may have been pre-trained with certain guardrails, it is risky to assume that they are comprehensive and prevent all harmful activity. A robust agentic AI framework must ensure that agents never engage in unwanted behavior.

Fabrix's Agentic platform guardrails are built on well-established benchmarks and allow our customers to refine them further to meet their specific requirements. Anytime an agent is created, AI Fabric ensures that it complies with these AI guardrails. Furthermore, AI Fabric monitors every interaction between the agent and LLM to ensure ongoing compliance with the established guardrails.



Managing Data Access for AI Personas

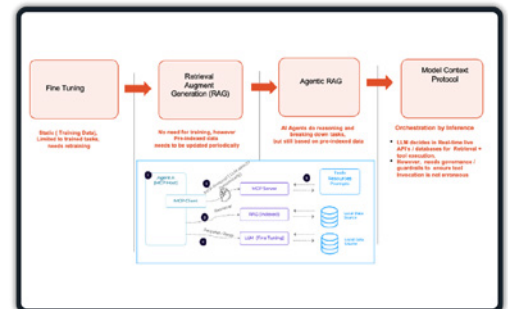
Within our Data Fabric, data is organized into Streams, Datasets, and Graph DB collections. AI agents, possessing user-like personas, can have access restrictions applied to limit their scope to specific streams or datasets. Granular control can be further implemented to restrict access to particular slices of data within a stream.

Managing Action Privileges for Agents

Autonomous agents will have the capability to execute numerous actions. Therefore, it is crucial to manage their privileges and limit the set of actions they can perform. The Agentic AI platform will ensure that these action privileges cannot be overridden by the AI itself, maintaining control and security.

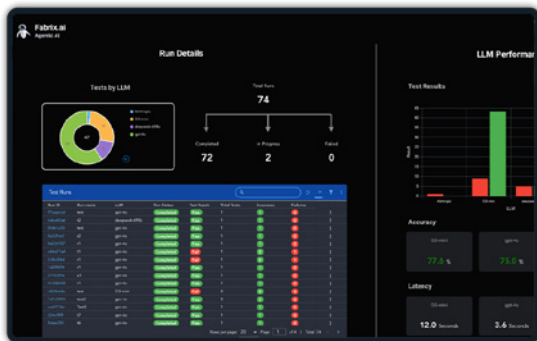
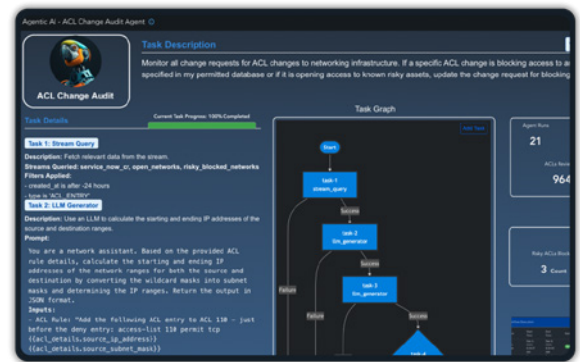
MCP (Model Context Protocol) enabled Data Fabric, Automation Fabric and Model Registry

- ▶ Fabrix's MCP Server for Fabrix.ai's Data Fabric and Automation Fabric, exposes 1000+ data and automation source integrations while also discovering other MCP Servers. The Fabrix.ai tools are divided into domains, so it becomes easy for the Models (LLMs) to access and prevent hallucinations
- ▶ Leverage existing data sources like Splunk, Elastic, Data Lakes with Fabrix's data abstraction capabilities without having to re-ingest data, correlating the data silos and gleaning actionable insights.
- ▶ MCP enabled Data Fabric
 - ▶ The Data Fabric MCP tools comprise of Bots and Telemetry pipelines which are responsible for data integration, transformation, enrichment and data automation
- ▶ MCP enabled Automation Fabric
 - ▶ The Automation Fabric MCP tools comprise of Workflows and Workflows tasks exposing Fabrix.ai's workflow engine as well as 3rd party workflow engines like Ansible, Cisco BPA (Business process automation) and Camunda.
- ▶ MCP Enabled Model Registry
 - ▶ LLM Selection and Inference-Time Reasoning: The platform supports various LLMs, including their own network LLM, open-source options (llama 3), and others. Decisions are primarily made at inference time, reducing the need for extensive training and labeled data.
 - ▶ The platform provides a number of different llms for different tasks of a task graph.



Visibility & Observability of Agents

Fabrix.ai Storyboards provide insights into the operation and performance of agents by visualizing agent workflows, tracking progress, and identifying bottlenecks or areas for optimization. This enhanced visibility enables better management and coordination of agent activities, ultimately leading to improved efficiency and effectiveness. Additionally, Storyboards capture detailed reasoning used by the LLM, which helps users understand the AI's decision-making process and refine it if necessary.



Agent Quality Control & Assurance

Agent Quality Control is crucial for maintaining the reliability and consistency of AI agents. This process involves continuous monitoring of agent behavior, identifying and rectifying errors or performance issues, and implementing improvements to optimize agent effectiveness and prevent operational disruptions.

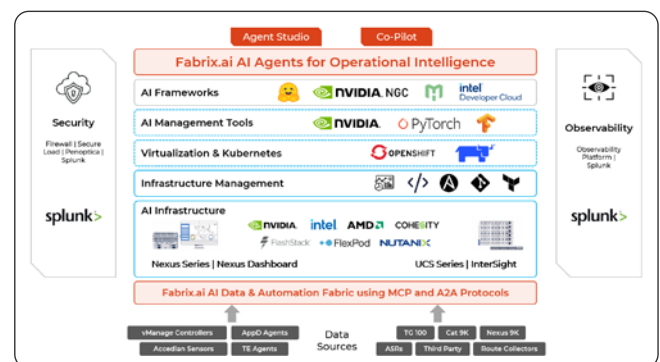
The Agent Quality Control module leverages Large Language Models (LLMs) to oversee agent performance and ensure that they are successfully achieving their intended outcomes.

Additionally, the system incorporates user feedback on both specific agent executions and overall agent performance. This feedback loop enables the Agent Quality Control module to identify emerging issues in new executions and further enhance agent performance.

Fabrix.ai Agentic AI Platform - best of breed consumption models

Fabrix.ai platform comes integrated with Cisco's Secure AI POD, which includes AI infrastructure like

- ▶ NVIDIA GPU's,
- ▶ Persistent storage platforms like MinIO and VAST,
- ▶ Data platforms like Splunk and Elastic,
- ▶ security platforms like Splunk SIEM and Cisco AI Defense, Tuymeryk
- ▶ Observability platforms like Splunk Olly cloud
- ▶ Virtualization and container platforms like Openshift and Rancher etc.



Customers have the choice of purchasing packaged platforms or deploying Fabrix.ai on their existing infrastructure or in the cloud.



We believe the most significant technological advances happen when visionary companies unite around shared challenges. By joining the AGNTCY collective, we're creating an entirely new operational paradigm where AI agents seamlessly collaborate to solve complex problems with minimal human intervention. Together, we will redefine operational excellence for the AI era -

Shailesh Manjrekar, CMO, Fabrix.ai

Ecosystem Partners



For more information visit:

www.fabrix.ai



To Request a demo or further queries:

fabrix.ai/request-demo

© Copyright 2015-2025