

FACTSHEET



An In Flow Governance Control Plane
for Autonomous AI Agents



CONTENTS

1

From Language
Models to Governed
Autonomous Execution

2

The Execution
Control Gap

3

AgentIQ Architecture

4

In Flow Governance
and Runtime Control

5

Verifiable and Enterprise
Ready Autonomous
Systems

Organizations are beginning to operationalize AI across core business functions. Instead of simply assisting users, AI agents are increasingly expected to retrieve enterprise data, interact with business applications, execute operational workflows, and coordinate decisions across complex enterprise environments. **As AI moves from informing decisions to executing decisions, autonomous execution becomes a new enterprise capability.**

That capability introduces a fundamentally different governance challenge. Every autonomous action must be authorized, governed, observable, and fully auditable before it reaches enterprise systems. The challenge is no longer whether AI can execute the work. It is whether enterprises can trust AI to execute the work safely, consistently, and in accordance with organizational policy. **Trust is no longer an outcome of AI adoption. It is a prerequisite for enterprise deployment.**



From Language Models to Governed Autonomous Execution

Enterprise AI is moving from systems that generate information to systems capable of executing work. Language models can now retrieve information, invoke tools, interact with applications, and coordinate activities across specialized agents. As these systems increasingly take actions on behalf of users, governance requirements extend beyond model outputs and into runtime execution itself.

Tool calling, reasoning, memory, and agent orchestration have transformed language models into systems capable of planning and executing Multi-Step Workflows. As agents gain the ability to access enterprise data, invoke tools, and make decisions autonomously, governance must evolve from monitoring outputs to controlling actions during execution.

Each advancement increases operational capability, but it also expands the governance surface. Software is no longer limited to generating responses. It can retrieve records, invoke tools, modify systems, and initiate business processes. The consequences of autonomous systems therefore arise not only from what they say, but increasingly from what they do.

AgentIQ™ was built specifically for this transition. It introduces an In Flow Governance Control Plane that participates directly in execution and establishes policy enforcement, human oversight, sensitive data protection, and auditability at the point where decisions become actions.

The Evolution of Enterprise AI

01**LARGE LANGUAGE MODELS**

Generate and interpret language.

02**TOOL AUGMENTED MODELS**

Interact with APIs and enterprise applications.

03**AUTONOMOUS AGENTS**

Reason, plan, act, observe, and adapt during execution.

04**MULTI-AGENT SYSTEMS**

Coordinate specialized agents across complex workflows.

Traditional software executes instructions.
Autonomous systems execute decisions.

AgentIQ™ governs those decisions while they are being executed.

2 The Execution Control Gap

The AI tooling ecosystem has matured rapidly for development and observability. It remains comparatively immature for governance and end-to-end control of autonomous systems. LLM gateways provide prompt visibility and model routing. Observability platforms reconstruct traces and support debugging. Output guardrails inspect responses for safety concerns. Model governance platforms address model risk and explainability. These technologies solve important problems, yet all remain largely external to runtime execution itself.

They cannot intercept tool calls before execution, pause actions for approval, redact sensitive information in flight, or enforce policy before actions occur.

Autonomous systems therefore require governance that operates from inside the execution flow.

AgentIQ addresses this architectural gap through an In Flow Governance Control Plane capable of denying actions, redacting sensitive information, pausing execution, and producing regulatory grade evidence in real time. Unlike proxies and callbacks that observe execution from outside the agent workflow, AgentIQ participates directly in runtime execution and establishes authority over actions while they are taking place.

Why Existing Controls Fall Short

Capability	Existing Governance Approaches	AgentIQ
Prompt and Response Visibility	✓	✓
Execution Tracing	✓	✓
Policy Enforcement Before Actions Execute	✗	✓
Real-Time Action Denial	✗	✓
Human Approval During Execution	✗	✓
PII Protection Before Model Exposure	✗	✓
Governance Inside Tool Execution	✗	✓
Tamper-Evident Auditability	Limited	✓

Why AgentIQ Is Different

1.

Inside the Graph, Not Outside It

Governance participates directly in execution.

2.

Governance Before Execution, Not After Execution

Policies are evaluated before actions occur.

3.

Framework Agnostic by Design

The governance engine remains independent of the underlying framework.

4.

Regulatory Grade by Design

Every execution event becomes reconstructable and verifiable.


By the time an event is observed, the action has already happened.

AgentIQ establishes authority over actions while execution is taking place.

3 AgentIQ Architecture

AgentIQ consists of two integrated components that together establish a centralized governance layer for enterprise AI agents.

- 1. LIGHTWEIGHT SDK** Embedded directly into agent code and responsible for In Flow Governance enforcement.
- 2. AgentIQ CONSOLE** Provides Real-Time Monitoring, approval workflows, policy management, compliance reporting, audit reporting, and cost visibility across autonomous systems.



Core Governance Engine

- 1. PolicyEngine**
Evaluates actions against organizational policies and runtime conditions.
- 2. PIIDetector**
Identifies and masks sensitive information before it reaches the model context window.
- 3. HITLGate**
Introduces approval workflows and escalation paths for High-Risk Actions.
- 4. AuditLogger**
Captures prompts, responses, tool executions, and governance decisions.
- 5. TraceAssembler**
Reconstructs complete execution histories across agents, tools, and approvals.

Framework Agnostic by Design

AgentIQ launches with LangGraph because of its graph-based workflows, native checkpointing, first class Human In The Loop capabilities, and state persistence. Each adapter is approximately one hundred and fifty lines of code, allowing new frameworks to be added without rewriting governance services.


→ New frameworks can be integrated without modifications to the governance engine, preserving a single governance model across heterogeneous agent ecosystems.

AgentIQ is not an observability platform layered around agents.
It is a governance control plane embedded directly into execution.

In Flow Governance and Runtime Control

AgentIQ participates directly in runtime execution by inserting governance into the tool execution path. Through SDK hooks and a framework-specific adapter, every tool invocation passes through policy evaluation, sensitive data protection, human oversight, and audit logging before actions are executed.

Rather than observing activity after execution has completed, AgentIQ evaluates whether actions should proceed, determines when approvals are required, and records every governance decision as part of the execution itself.



AgentIQ governs every tool call before it executes, enforcing policy, approvals, sensitive data protection, and audit requirements in real time. Every tool invocation passes through four levels of control.

Runtime Controls

1. ALLOW

Low-Risk Operations execute immediately while remaining fully visible and auditable.

2. CONFIRM

Medium-Risk Actions pause and require explicit user confirmation.

3. ESCALATE

High-Risk Actions route to designated approvers while execution state is preserved.

4. DENY

Prohibited actions never execute and are blocked by policy.

Runtime Capabilities

Role Based Policy Enforcement • PII Masking Before Model Exposure • Native Pause and Resume Workflows • Human Oversight Using Framework Interrupts • Token and Cost Visibility • Immutable Audit Records • Execution Tracing • Policy Decision Visibility

The defining property of governed autonomy is the ability to intervene while execution is in progress.

Verifiable and Enterprise Ready Autonomous Systems

Enterprise deployment ultimately depends on whether autonomous actions can be reconstructed, explained, and defended. AgentIQ extends governance across six execution surfaces and applies consistent policy evaluation, sensitive data protection, human oversight, and audit logging throughout the execution lifecycle.



Every prompt, response, tool invocation, approval, and policy decision becomes part of a continuously assembled execution record that can be reconstructed, independently verified, and presented as regulatory evidence.

Governance Control Points

Input Validation | Model Selection | Tool Execution | Data Retrieval | Agent Handoff | Output Delivery

Don't Block the Tool. Govern the Action.



AgentIQ evaluates each action at runtime and can dynamically allow, confirm, escalate, or deny a tool call based on its parameters and risk. This enables agent autonomy where it is safe and targeted human oversight where it matters.

Control point	Risk Addressed	Governance capability needed	Regulatory driver
Input Validation	Adversarial Inputs, Rate Abuse	Pii Detection, Injection Blocking, Rate Limiting	GDPR Art.25, CCPA
Model Selection	Unapproved Models, Cost Overrun	Model Allowlisting, Version Pinning, Cost Ceilings	EU AI Act Art.9
Tool Execution	Unauthorized Actions	Role-Based Access Control, Runtime Parameter Validation, Conditional Authorization, Human Approval	SOC 2 CC6.1, OCC Guidance
Data Retrieval	Unauthorized Data Access	Scope Enforcement, Data Classification, Access Logging	GLBA, HIPAA, GDPR
Agent Handoff	Privilege Escalation	Delegation Approval, Tier-Based Routing	EU AI Act Art.14
Output Delivery	Harmful or Non-Compliant Responses	Disclaimer Injection, Compliance Review, PII Scanning	Fair Lending, UDAP

Enterprise Capabilities

Prompt and Response Logging • Tool Execution Visibility • Role-Based Access Control • PII Detection and Redaction • Human Oversight • Hash Chained Audit Trails • Compliance Reporting • Regulatory Mapping for SOC 2, EU AI Act, OCC, HIPAA, and GDPR • Visual Policy Builder with Role Mapping • Agent Registry • Model Allow Lists • Prompt Injection Guard • Maximum Iteration Limits • Token Usage Limits • Out of the Box Audit Report Generation • Compliance Reporting Templates • Kill Switch • Additional Framework Adapters

Every execution event incorporates the previous event's SHA 256 hash, producing Tamper-Evident Records that can be independently verified and used as regulatory evidence.

Built for High Assurance Environments

✓
Financial Services

✓
Healthcare

✓
Insurance

✓
Government

Current Future Roadmap



Estimated Agentic AI Risk Rate



Global Policy



Dynamic Tool Governance



LDAP Support



Policy Import/Export



Support for AWS Strands, Google ADK, and Microsoft Agent Framework

 **AgentIQ** transforms autonomous systems from ungoverned execution engines into **observable, controllable, and verifiable enterprise systems** that can be deployed with confidence in regulated environments.



meshIQ.com

