SENTIENT REVIEW

VOLUME 01 · PUBLICATION № 11 · FLAGSHIP RESEARCH PUBLICATION

The Enterprise Agent Control Plane

Architecting Authority, Execution Continuity, Assurance, and Economics Across Models and Agents

Version 1.0 · Published November 5, 2026

Author & Editor: Fritzie M. Alex

https://sentient-review.com/perspectives/the-enterprise-agent-control-plane

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Executive Summary

An architectural examination of enterprise agent control planes, execution continuity, delegated authority, assurance, and AI operating economics.

The Framework

FW 26 · Execution Continuity Framework · Original Framework · Version 1.0 · Proposed Framework

A framework for preserving enterprise authority, execution evidence, resource integrity, and recovery obligations as AI models, agents, and runtimes change.

https://sentient-review.com/frameworks/execution-continuity-framework

Developed Through

FN 37 · Fallback Is an Architectural Transition

https://sentient-review.com/field-notes/fallback-is-an-architectural-transition

FN 38 · Replace the Harness. Preserve the Session.

https://sentient-review.com/field-notes/replace-the-harness-preserve-the-session

FN 39 · Authority Is Enforced at the Action Boundary

https://sentient-review.com/field-notes/authority-is-enforced-at-the-action-boundary

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

CONTENTS

Contents

CHAPTER 01 The Enterprise Control Problem

CHAPTER 02 The Architecture of Agentic Systems

CHAPTER 03 The Emerging Enterprise Architecture

CHAPTER 04 Where Architectural Substitution Creates Risk

CHAPTER 05 Four Enterprise Control Surfaces

CHAPTER 06 The Execution Continuity Framework

CHAPTER 07 Architectural Judgment: When Complexity Earns Its Place

Research Foundations and Connected Work

Citation

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

CHAPTER 01 OF 07

The Enterprise Control Problem

Enterprise AI is entering a different architectural phase.

The first wave concentrated on model capability. Organizations evaluated reasoning quality, context capacity, inference performance, and the economics of individual model interactions.

The next phase concerns the systems through which intelligence becomes enterprise execution.

A business workflow can engage several models. An agent can delegate work to specialized participants. Enterprise tools can execute transactions. Human reviewers can authorize consequential decisions. Execution can continue across different infrastructure environments.

These capabilities create new architectural possibilities.

They also change the responsibilities of enterprise technology architecture.

A model contributes inference. A harness organizes execution. A runtime provides operational continuity. Enterprise systems establish authority and record business consequences.

The architecture must connect these responsibilities into a coherent operating environment.

This is the emerging domain of the **enterprise agent control plane**.

From Model Capability to Operating Capability

Model capability determines what an AI system can interpret or produce.

Operating capability determines what the enterprise can safely delegate, execute, verify, and economically sustain.

The distinction becomes visible when an agent participates in consequential work.

Consider a procurement workflow.

An agent evaluates supplier proposals and identifies a preferred vendor. Another agent reviews contractual terms. A finance service verifies budget availability. A human executive authorizes the commitment.

The resulting purchase order creates an enterprise obligation.

Each participant contributes to the process. The enterprise remains accountable for the outcome.

The operating architecture must preserve the relationship between the decision and its authorization. It must retain the evidence supporting the commitment. It must also establish the actual transaction result.

These requirements extend beyond the capabilities of any individual model.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

The Emerging Control Plane

An enterprise agent control plane is the collection of shared management and control capabilities through which an organization directs agentic execution across applications and operating environments.

It establishes visibility into participating agents and their owners.

It defines approved execution conditions.

It coordinates policy and lifecycle management.

It connects operational evidence with intervention decisions.

The control plane can operate through distributed services. Its architecture may combine capabilities supplied by cloud platforms and enterprise applications. Specialized security and observability systems may contribute additional functions.

The important architectural distinction is between **defining operating requirements and enforcing them during execution**.

A management service can establish approved tool permissions. An authorization gateway can enforce those permissions. The business application ultimately determines the result of the transaction.

The enterprise control environment must account for all three responsibilities.

The Central Thesis

As enterprise AI systems become more compositional, the durable architectural responsibility shifts toward preserving coherent control across changing execution configurations.

Models can be replaced.

Agents can be delegated new work.

Harnesses can evolve.

Runtimes can be migrated.

Enterprise obligations continue through these changes.

The control plane provides a shared operating capability for managing those obligations.

The architecture must also support the execution mechanisms that enforce them.

This distinction creates the central proposition of this publication:

Enterprise AI maturity is measured through the quality of governed execution, the integrity of accepted outcomes, and the economics of sustained operation.

Model capability remains essential.

Enterprise value emerges when that capability participates in a dependable system of work.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

The Architecture of Agentic Systems

Enterprise AI architecture contains several distinct forms of composition.

Understanding their differences is essential for determining which operating controls are required.

The Single-Model Harness

A single-model harness provides the execution machinery surrounding one primary task-performing model.

It may assemble context and manage prompts. It may coordinate tool calls. It may also maintain workflow state and enforce execution limits.

The harness translates model capability into a repeatable execution process.

A sophisticated single-model harness can support consequential enterprise work. Its operating requirements depend on the actions it performs and the systems it accesses.

The Multi-Model Harness

A multi-model harness coordinates two or more task-performing models within a defined execution environment.

Different models may perform different tasks.

A routing mechanism may select a model according to capability or cost. A fallback mechanism may select an alternative provider when execution conditions change.

The architecture introduces model heterogeneity.

Models may differ in their tool interfaces and output behavior. They may operate under different provider conditions. Their performance characteristics may vary across workloads.

The harness must manage these differences according to the requirements of the workflow.

The Multi-Agent System

A multi-agent system distributes work across multiple agent instances.

Each agent may maintain a distinct task context or execution loop. Agents may coordinate through an orchestrator or through defined communication mechanisms.

Several agents can use the same model.

A single agent can also use several models.

Model diversity and agent topology are therefore separate architectural properties.

The number of participating agents provides limited information about the system's operating maturity.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

The important questions concern coordination, authority, state, and accountability.

The Enterprise Agent Control Plane

The enterprise control plane addresses a broader management scope.

It connects agent identities with accountable owners. It manages approved configurations and operating policies. It supports fleet visibility and operational intervention.

A control plane may govern a single-model application or a heterogeneous agent environment.

Its defining characteristic is shared management across an operational domain.

The control plane also depends on runtime enforcement services. Those services carry policy decisions into the execution path.

Four Independent Architectural Dimensions

A more useful taxonomy examines four dimensions.

Dimension	Architectural question
Model repertoire	Which models can participate in execution?
Agent topology	How is work distributed among agents?
Trust-domain topology	Which organizational and technical boundaries does execution cross?
Management scope	Which controls operate at workflow, application, or enterprise level?

These dimensions can vary independently.

A single-model application can require rigorous authorization because it executes financial transactions.

A multi-model research workflow can operate within a limited authority envelope because it produces advisory analysis.

A multi-agent system can use one provider while crossing several enterprise security boundaries.

The control requirements follow the actual execution topology and business consequence.

Complexity as an Architectural Choice

Multi-model and multi-agent systems offer opportunities for specialization.

They can distribute independent tasks and match model capabilities to workload requirements.

They also introduce coordination overhead.

Additional execution paths create more states to observe. Delegation increases the importance of explicit ownership. Model substitution requires compatibility testing.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Research on agent-system scaling demonstrates that coordination benefits depend on task structure. Some parallelizable tasks benefit from additional agents. Sequential tasks can experience substantial performance degradation from coordination overhead.

The architectural objective is therefore disciplined composition.

The enterprise should select the execution structure that produces accepted outcomes with the required reliability and economic performance.

The control plane supports that judgment through evidence.

Its purpose is to make architectural choices observable, accountable, and revisable.

EXHIBIT 01. ENTERPRISE AGENT ARCHITECTURE TAXONOMY

Dimension	Architectural question
Model repertoire	Which models can participate in execution?
Agent topology	How is work distributed among agents?
Trust-domain topology	Which organizational and technical boundaries does execution cross?
Management scope	Which controls operate at workflow, application, or enterprise level?

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

The Emerging Enterprise Architecture

The enterprise agent control plane is becoming visible through several distinct market developments.

Frontier-model companies are expanding their execution infrastructure. Hyperscalers are developing shared agent-management capabilities. Enterprise platforms are introducing model routing and operational monitoring.

Security providers are also demonstrating the value of coordinating specialized models within controlled workflows.

These developments reveal an expanding architectural domain.

They also reveal different approaches to ownership and implementation.

Frontier Models and Durable Execution

Anthropic's Managed Agents architecture provides a useful engineering example.

The architecture separates the session record from the agent harness and execution sandbox.

The session preserves the history of work. The harness manages model interaction and tool coordination. The sandbox provides an environment for executing code and interacting with resources.

This separation allows components to be replaced or recovered independently.

The architectural value lies in durable interfaces and explicit component boundaries.

A model improvement may change the behavior required from the harness. A sandbox failure may require replacement of the execution environment. The session history can remain available throughout these changes.

The design illustrates an important principle.

Execution state should have an ownership model that survives changes in the machinery carrying the work.

Hyperscalers and Fleet Management

Microsoft Foundry Control Plane represents another architectural direction.

Its documented capabilities include agent inventory and lifecycle management. It connects fleet monitoring with security and compliance services. It also supports management across specified agent platforms.

This is an example of enterprise control extending beyond an individual application.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

The enterprise gains a consolidated view of supported agents and their operating conditions.

The distinction between visibility and enforcement remains important.

A fleet-management interface can expose operating status. Effective control also requires enforcement mechanisms along the execution path.

AWS AgentCore provides a complementary example.

Its gateway and policy capabilities support authorization for mediated tool interactions.

The architecture illustrates how policy management and action enforcement can be separated into explicit responsibilities.

Together, these examples demonstrate the development of shared operating infrastructure for agentic systems.

Their product boundaries and supported integrations remain vendor-specific.

Enterprise Model Gateways

Uber's GenAI Gateway provides an early production example of centralized model access.

In July 2024, Uber Engineering reported that the gateway served close to 30 customer teams and processed 16 million queries per month.

The architecture provided a common interface for external and Uber-hosted models. It incorporated authentication and authorization. Operational metrics and audit logging supported reporting, security oversight, and cost attribution.

This established a documented production example of shared model-access infrastructure operating across enterprise teams.

The example concerns model-access infrastructure. Broader agent execution requires complementary capabilities for delegated authority, durable workflow state, and consequential business actions.

Source: Uber Engineering, July 11, 2024.

https://www.uber.com/us/en/blog/genai-gateway/

Databricks Unity Gateway demonstrates a related development.

Its documented Smart Routing capability, identified as beta, selects models for supported coding-agent workflows. Delegated subagent tasks can receive independent model selections.

This makes model choice an operational decision inside the workflow.

The capability also creates requirements for route visibility and cost attribution.

Security as a Multi-Model Operating Environment

Palo Alto Networks' Unit 42 Continuous Frontier AI Defense provides a specialized market signal.

The service combines frontier models with proprietary multi-model harnesses and human security expertise.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Its architecture routes offensive-security tasks according to model capabilities.

The resulting workflow connects vulnerability discovery with expert validation and remediation activity.

Unit 42 reports improved coverage and operating performance from its implementation.

Those performance claims are vendor-reported and specific to its testing environment.

The architectural observation remains useful.

Specialized models can participate in a common operating workflow where orchestration and validation determine how their outputs become actionable.

The Market Direction

These implementations support three conclusions.

First, agent infrastructure increasingly extends beyond inference endpoints.

Second, shared operational capabilities are appearing across model providers and enterprise platforms.

Third, the market is developing several architectural approaches to agent management and execution.

The emerging pattern is **greater separation between model capability and enterprise operating responsibility**.

The precise boundaries of a future enterprise control plane remain an architectural and commercial choice.

Enterprises can compose these capabilities through existing platforms and specialized services.

The durable requirement is coherent control across the resulting execution environment.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Where Architectural Substitution Creates Risk

An enterprise workflow can change its execution configuration while retaining the same business objective.

The change may improve performance or resilience.

It may also alter the conditions under which the work remains authorized and verifiable.

Architectural substitution therefore requires examination at the point where the execution path changes.

Model Routing and Fallback

Model routing selects an execution model according to defined criteria.

Those criteria may include task capability and inference cost. Availability and latency may also influence the decision.

A fallback mechanism introduces an alternative execution configuration when the original path becomes unsuitable.

The replacement model may support different tools or output formats. It may operate through another provider. It may also require different context-management behavior.

These differences create compatibility requirements.

A valid fallback decision must account for the complete workflow.

Model eligibility is one component.

The enterprise must also examine whether the replacement can perform the required actions under approved conditions.

Delegation and Authority

Agent delegation transfers responsibility for part of a workflow to another execution participant.

The delegated task may require access to enterprise data or tools.

Authority must be represented through explicit permissions.

The receiving agent requires a defined execution identity and appropriate scope.

The originating workflow retains accountability for the overall business outcome.

This creates two related control responsibilities.

The enterprise must determine which participant can act.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

It must also establish which owner remains accountable for the resulting consequence.

Delegation becomes manageable when both responsibilities are explicit.

Context and Evidence Lineage

Agent execution depends on context.

A replacement model may receive a summary of earlier activity. A delegated agent may receive selected business records. A recovered workflow may reconstruct context from durable event history.

Each approach changes the information available to the next participant.

The system must preserve the provenance and validity of consequential information.

A summary may support efficient reasoning. Authoritative business records remain necessary for establishing completed external actions.

This distinction becomes critical during financial execution and operational recovery.

State and Recovery

A workflow may invoke an external service and experience an interruption before receiving confirmation.

The external action may already have completed.

The execution system must reconcile the uncertain result before continuing.

Durable checkpoints support reconstruction of workflow state.

Transaction identifiers support correlation with external operations.

Idempotency mechanisms help manage repeated requests.

Compensating actions address effects that require correction.

These mechanisms serve different responsibilities within recovery.

Evaluation Across Configurations

A replacement model changes the behavior of the deployed system.

The enterprise therefore requires evaluation of the resulting configuration.

Evaluation should include the model and its harness. It should also examine tool behavior and the resulting environment state.

A coherent explanation of successful execution is one form of evidence.

The accepted business outcome is another.

Anthropic's evaluation guidance distinguishes execution transcripts from the final state produced in the environment.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

That distinction is fundamental to enterprise assurance.

The evaluation system must establish whether the intended result occurred under the applicable operating conditions.

The Control Requirement

Routing, delegation, and recovery are legitimate architectural capabilities.

Each creates a transition requiring appropriate verification.

The enterprise must determine which obligations remain satisfied and which conditions require renewed authorization or reconciliation.

Architectural flexibility becomes an enterprise capability when execution transitions preserve accountable control.

This is the problem addressed by the Execution Continuity Framework.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Four Enterprise Control Surfaces

An enterprise agent control plane must connect technical execution with business accountability.

Four control surfaces organize the operating decisions.

They represent distinct management responsibilities supported by shared execution infrastructure.

Control Surface 01. Outcome Economics

The economic unit of agentic execution is the accepted business outcome.

Inference cost measures one component of the operating system.

A complete economic assessment also includes tool execution and infrastructure consumption. Human review and remediation contribute additional costs.

The enterprise must establish whether the workflow produces accepted results at an economically sustainable cost.

A useful measure is:

Cost per Accepted Outcome = Total Attributable Workflow Cost / Number of Accepted Outcomes

The Accounting Boundary

The measurement follows the canonical economic model established in Agentic Execution, Chapter 12.

Total attributable workflow cost includes human execution and AI consumption. It includes infrastructure and integration expenditure. Review and recovery costs form part of the same accounting boundary. Allocated transformation investment completes the measurement.

Costs and accepted outcomes must share a defined measurement period and outcome class.

The numerator includes failed attempts and rejected work. Retry expenditure and recovery effort remain attributable to the workflow.

Shared costs require a consistent allocation basis. Transformation investment is allocated over a disclosed useful-life horizon.

Each cost is assigned once. Human execution labor and human review labor are classified consistently to preserve accurate attribution.

Each distinct accepted outcome enters the denominator once after satisfying its acceptance criteria.

A period with zero accepted outcomes reports the incurred expenditure separately. The unit-cost measure remains undefined for that period.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

This accounting discipline enables meaningful comparison across execution configurations while preserving the distinction between operating efficiency and realized business value.

Connected reference: Agentic Execution, Chapter 12, The Economics of Agentic Execution.

https://sentient-review.com/perspectives/agentic-execution/12-the-economics-of-agentic-execution

The numerator includes expenditure associated with failed attempts and recovery.

The denominator includes outcomes that satisfy defined business acceptance criteria.

The measure should be accompanied by quality and workload coverage.

A workflow that processes only simple cases may achieve attractive unit economics while transferring difficult work to human teams.

Economic accountability therefore requires a complete view of the operating process.

The control plane supports this responsibility through attribution and operating thresholds.

Control Surface 02. Authority and Action Safety

Authority determines which participants may perform consequential actions.

An agent may receive permission to read records or prepare recommendations.

A separate authorization may permit a financial commitment or a change to an enterprise system.

These permissions belong to explicit identities and scopes.

The control environment establishes delegated authority and approval requirements.

Enforcement services apply those requirements at the relevant action boundaries.

The enterprise must also preserve intervention capabilities.

A workflow may require suspension when permissions change. A compromised credential may require revocation. An uncertain transaction may require reconciliation.

The control surface connects delegated capability with accountable decision rights.

Control Surface 03. Evaluation and Assurance

Agentic execution requires continuous evidence about operating performance.

Predeployment evaluation establishes an initial qualification baseline.

Production monitoring identifies changes in behavior and operating conditions.

Incident analysis connects observed failures with corrective action.

Evaluation must examine the deployed configuration and the resulting business state.

A model comparison may establish differences in task quality.

An execution assessment must also examine tool effects and recovery behavior.

The enterprise should maintain evaluation suites appropriate to the workflow's consequence.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Higher-impact actions require stronger evidence and more explicit acceptance conditions.

Assurance becomes an ongoing operating process.

Control Surface 04. Capacity and Capital Discipline

AI execution consumes compute and supporting infrastructure.

Routing decisions influence inference expenditure.

Concurrency and serving architecture influence capacity utilization.

Infrastructure placement affects latency and operational exposure.

These decisions connect workload design with infrastructure economics.

The control plane can support operating decisions through resource accounting and capacity policy.

Longer-term capital decisions require an additional perspective.

An enterprise may purchase model services through APIs. It may reserve capacity or operate dedicated infrastructure.

Each arrangement creates a different economic exposure.

Infrastructure ownership introduces commitments concerning hardware and utilization. Power availability and equipment lifecycle can also affect investment economics.

These decisions belong to enterprise technology and capital planning.

The relationship between execution and capital becomes visible when workload demand informs capacity commitments.

The Shared Execution Foundation

The four control surfaces depend on common technical capabilities.

Durable state supports recovery.

Identity and policy services support authorization.

Execution traces support operational investigation.

Authoritative business records establish completed external effects.

Resource accounting connects consumption with operating limits.

These capabilities may be supplied through several platforms.

The enterprise architecture must establish how they operate together.

The control plane creates coherence across the four decision domains. The execution infrastructure provides the mechanisms through which those decisions become effective.

This relationship connects AI architecture with enterprise operating performance.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

EXHIBIT 02. FOUR ENTERPRISE CONTROL SURFACES

1. Outcome Economics
2. Authority and Action Safety
3. Evaluation and Assurance
4. Capacity and Capital Discipline

Shared Execution Foundation

Durable state supports recovery.

Identity and policy services support authorization.

Execution traces support operational investigation.

Authoritative business records establish completed external effects.

Resource accounting connects consumption with operating limits.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

The Execution Continuity Framework

The emerging control-plane architecture creates a specific design requirement.

Enterprise obligations must remain satisfied as execution changes across models and operating environments.

The **Execution Continuity Framework** provides a proposed method for examining this requirement.

It extends the Delegated Action Architecture through transition-specific verification.

The Execution Envelope

Every consequential workflow operates within an execution envelope.

The envelope defines the conditions governing a unit of work.

It establishes five obligations.

Outcome Integrity

The workflow retains a valid definition of accepted completion.

Authority Integrity

Consequential actions remain subject to current authorization.

Evidence Integrity

Execution decisions and resulting effects remain attributable through a coherent evidence history.

Resource Integrity

Operating consumption remains within the applicable limits and approved adjustments.

Recovery Integrity

Execution preserves sufficient state and external evidence to support safe continuation.

These obligations form the continuity requirements.

They can be represented through workflow specifications and policy configuration. Durable records provide supporting evidence.

Execution Transitions

An execution transition changes a consequential element of the operating configuration.

Examples include model replacement and provider fallback.

Agent delegation creates another transition.

Human handoff changes the participant responsible for a decision.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Runtime migration changes the environment carrying execution.

Recovery introduces a transition between interrupted and resumed work.

Each transition requires verification appropriate to its consequence.

Continuity Verification

The framework defines a five-stage verification method.

Stage 1. Identify the proposed transition.

Record the current configuration and the intended replacement.

Stage 2. Establish the current execution state.

Retrieve durable state and relevant external evidence.

Stage 3. Evaluate the continuity obligations.

Examine the five invariants against current operating requirements.

Stage 4. Determine the operational disposition.

Authorize continuation or require an appropriate control response.

Stage 5. Record the resulting decision.

Preserve the transition evidence and updated execution state.

The disposition has four possible forms.

Proceed authorizes continuation under the governing conditions.

Reauthorize requires an accountable decision.

Reconcile requires confirmation of execution state or external effects.

Contain places execution into a controlled condition while the unresolved issue is addressed.

A Payment Execution Example

Consider an approved supplier payment.

An agent submits the payment request.

The payment service accepts the request. The connection then times out before confirmation reaches the agent.

The workflow selects another model to continue.

The replacement model may be capable of issuing the same payment instruction.

The continuity requirement concerns the actual transaction state.

The system retrieves the original transaction identifier and checks the authoritative payment service.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

If the payment completed, the workflow records the confirmed result.

If the result remains uncertain, the workflow enters reconciliation.

Current authorization governs any subsequent action.

The model can change while the payment obligation retains its identity and execution history.

Architectural Contribution

Execution Continuity combines established distributed-systems principles with a specific enterprise AI design method.

Its contribution is the relationship between the governing execution envelope and architectural transitions.

The framework provides a structured way to test whether an execution change preserves enterprise obligations.

Its effectiveness requires implementation and empirical evaluation.

The method can be assessed through controlled failures and transition testing.

Relevant measures include recovery correctness and evidence completeness. Authorization accuracy and accepted-outcome economics provide additional measures.

The Principle

The enterprise can evolve its execution architecture while preserving the obligations that govern the work.

This is the foundation of execution continuity.

It connects the flexibility of agentic systems with the discipline required for enterprise operation.

EXHIBIT 03. EXECUTION CONTINUITY ARCHITECTURE

Execution Envelope → Execution Transition → Continuity Verification → Operational Disposition

- I1 · Outcome Integrity
- I2 · Authority Integrity
- I3 · Evidence Integrity
- I4 · Resource Integrity
- I5 · Recovery Integrity

Proceed | Reauthorize | Reconcile | Contain

---

https://sentient-review.com/frameworks/execution-continuity-framework

FW 26 · Proposed Framework

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

Architectural Judgment: When Complexity Earns Its Place

The development of enterprise agent control planes creates opportunities for greater architectural flexibility.

It also raises a question of engineering judgment.

Which operating requirements justify additional architectural complexity?

Multi-model routing can improve workload economics when tasks exhibit meaningful differences in model suitability.

Multi-agent coordination can create value when work benefits from specialization or parallel execution.

Shared control infrastructure can improve visibility across distributed applications.

Each capability introduces operating responsibilities.

The enterprise must evaluate the resulting system against its intended business outcome.

The Economics of Coordination

Coordination consumes resources.

An additional agent may require separate inference and context processing.

A delegated task may create further tool calls.

Verification and reconciliation can increase execution time.

These costs belong within the assessment of the complete workflow.

Research on agent-system scaling demonstrates that performance varies according to task structure and coordination architecture.

Parallelizable work can benefit from coordination.

Sequential tasks can experience additional overhead and degraded performance.

The result is a practical architectural principle.

Complexity earns its place through measurable improvements in accepted outcomes and operating performance.

Control Plane Architecture as a Composition Decision

An enterprise control plane can emerge through several implementation approaches.

A hyperscaler may supply shared management services.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

An enterprise platform may provide native agent controls.

A specialized gateway may enforce model-access policies.

Existing workflow infrastructure may provide durable execution.

Security services may establish action authorization and intervention.

These capabilities can be composed into an operating architecture.

The enterprise must establish clear ownership of each control responsibility.

The appropriate design depends on the operating environment and the consequences of delegated action.

Architectural Independence

The value of an enterprise control plane lies in the stability of its governing responsibilities.

Model capability will continue to evolve.

Execution frameworks will change.

New providers will introduce different operating models.

Enterprise applications will embed increasingly capable agents.

A well-designed control architecture can accommodate these developments through explicit interfaces and accountable execution boundaries.

This requires deliberate choices about portability and operational dependency.

It also requires evaluation of the commercial conditions under which infrastructure services are supplied.

Architectural independence is a design objective with economic trade-offs.

Its value should be established through the enterprise's actual requirements.

From Agent Deployment to Enterprise Operating Capability

The strategic opportunity extends beyond deploying additional agents.

It concerns the development of an enterprise environment in which machine execution can participate in accountable work.

The organization gains the ability to evaluate execution configurations and manage delegated authority.

It gains stronger evidence about operational outcomes.

It can connect workflow economics with infrastructure and investment decisions.

The enterprise control plane supports these capabilities through shared management and operating discipline.

The execution infrastructure carries them into business systems.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

The Enterprise Architecture Imperative

The enduring challenge is the relationship between changing technical capability and stable enterprise obligations.

Models will improve.

Agents will become more specialized.

Execution environments will evolve.

The enterprise must preserve coherent authority and accountable outcomes across those developments.

The Execution Continuity Framework provides one proposed method for examining this responsibility.

Its five invariants establish the governing obligations.

Its transition method establishes the verification process.

Its operating dispositions connect architectural change with controlled execution.

Together, these elements support a broader principle of enterprise AI architecture.

The enterprise agent control plane establishes the conditions under which intelligence becomes accountable execution.

Its value is realized through the quality of work completed and the economic results achieved.

That is the architecture of enterprise AI control.

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

REFERENCES

Research Foundations and Connected Work

Primary Technical and Market References

01. Anthropic. Scaling Managed Agents: Decoupling the Brain from the Hands. April 2026.

Engineering account of durable sessions and replaceable execution components.

https://www.anthropic.com/research/managed-agents

02. Microsoft. Microsoft Foundry Control Plane.

Technical documentation covering supported agent fleet management and operational visibility.

https://learn.microsoft.com/en-us/azure/foundry/control-plane/overview

03. Amazon Web Services. AgentCore Gateway and Policy.

Technical documentation describing authorization responsibilities for mediated tool execution.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/policy-permissions.html

04. Uber Engineering. Navigating the LLM Landscape: Uber's Innovation with GenAI Gateway. July 11, 2024.

Production engineering account of shared model access and operational management.

https://www.uber.com/us/en/blog/genai-gateway/

05. Databricks. Select Models or Use Smart Routing.

Documentation of model selection and routing in supported agent workflows. Smart Routing is identified as a beta capability.

https://docs.databricks.com/aws/en/ai-gateway/coding-agent-model-selection

06. Palo Alto Networks. Introducing Unit 42 Continuous Frontier AI Defense. September 2026.

Vendor account of a specialized multi-model security service. Performance claims are vendor-reported.

https://www.paloaltonetworks.com/blog/2026/09/introducing-unit-42-continuous-frontier-ai-defense/

07. Anthropic. Demystifying Evals for AI Agents. January 2026.

Engineering guidance covering agent evaluation and the distinction between execution traces and resulting environment state.

https://www.anthropic.com/research/demystifying-evals-for-ai-agents

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

08. Kim et al. Towards a Science of Scaling Agent Systems. 2025–2026.

Research examining coordination architecture and performance across different task structures.

https://arxiv.org/abs/2512.08296

09. Ong et al. RouteLLM: Learning to Route LLMs with Preference Data. 2024.

Research on model routing and cost-performance trade-offs.

https://arxiv.org/abs/2406.18665

Connected Sentient Review Work

Execution Continuity Framework

The original framework developed for this publication.

The Delegated Action Architecture

Foundational framework for responsible delegated enterprise execution.

https://sentient-review.com/frameworks/the-delegated-action-architecture

Agentic Execution: The Architecture of Delegated Action in the Enterprise

Research publication covering enterprise agent architecture and operating responsibilities.

https://sentient-review.com/perspectives/agentic-execution

Computational Coordination Architecture

Framework addressing computational coordination and enterprise commitments.

https://sentient-review.com/frameworks/computational-coordination-architecture

Publication Classification

This publication combines established engineering evidence with an original architectural synthesis.

The Execution Continuity Framework is presented as a proposed method for transition-specific verification.

The enterprise agent control plane is examined as an emerging architectural capability with multiple implementation approaches.

The research evidence supports the architectural analysis. Future implementation and testing will strengthen the framework's empirical foundation.

© 2026 Sentient Review. All rights reserved.

CITATION

Citation

Sentient Review. (November 2026). The Enterprise Agent Control Plane. Volume 01, Publication 11. Version 1.0.

https://sentient-review.com/perspectives/the-enterprise-agent-control-plane

© 2026 Sentient Review. All rights reserved. · Original Intellectual Property · sentient-review.com

SENTIENT REVIEW

Intellectual Property Notice

This publication and its original frameworks, models, structures, diagrams, terminology, and analytical materials are intellectual property of Sentient Review.

© 2026 Sentient Review. All rights reserved.

No part of this artifact may be reproduced, republished, distributed, adapted, incorporated into commercial materials, or represented as another party’s work without prior written permission, except as permitted by applicable law.

Published by Sentient Review

sentient-review.com