

Normora Agent Expenditure Governance

Pre-Execution Resource Governance for AI Agents

PUBLIC CONCEPT BRIEF / v1.1 / AUGUST 2026

Govern the execution before managing the expenditure.

A public framework for defining why an AI agent may act, how far it may proceed, what resources it may consume, and when execution must stop.

	Introduction
1	System Definition
2	Why Now
3	The Next Control Boundary
4	Strategic Distinction
5	Governance Architecture
6	Control Domains
7	Authority Levels
8	Gate Decisions
9	Contract & Receipt
10	Controlled Execution Example
11	Control Model
12	System Relationship
13	Public & Protected
14	Controlled Validation
	Closing Principle

Introduction

Every cost begins with an execution decision.

AI agents do more than generate answers. They interpret intent, inspect context, call tools, modify systems, verify outcomes, retry failures, and recover from errors. Each step can be reasonable in isolation, while the total execution grows beyond the user's original purpose.

Normora moves the control point upstream - before execution begins, and again before it expands.

Execution Expenditure Is a Governance Decision

Normora Agent Expenditure Governance is a pre-execution governance framework and control architecture for AI agents. It structures purpose, scope, authority, iteration, evidence, and stop conditions before resource consumption expands. It does not promise an exact token forecast. It governs the sequence of decisions that creates expenditure.

Why Execution Governance Is Needed Now

As agents move from response generation to tool use and system modification, expenditure no longer comes from a single model call. It accumulates across execution paths.

Unpredictable Cost Expansion

Consumption grows through analysis, tool use, verification, and recovery.

Retry Loops

The same unresolved hypothesis is executed repeatedly without material new evidence.

Tool-Call Expansion

Agents invoke more tools than the intended outcome requires.

Context Expansion

More files, history, and external information enter the task.

Agent Drift

Execution gradually moves beyond the original purpose or authorized scope.

These are not only efficiency failures. They are failures of execution governance.

The Next Control Boundary Is Execution

Firewalls established boundaries around networks.

IAM assigned identity and access.

Zero Trust made every access request conditional.

AI agents introduce a new control problem: a system can consume resources, invoke tools, alter state, and continue acting after the initial request.

EXECUTION GOVERNANCE

A control architecture for determining whether an agent may act, how far it may proceed, and what must remain true for execution to continue.

Normora structures the execution boundary.

Not Another Token Optimization Layer

Optimization improves the economics of an authorized execution. Normora governs the authorization itself.

EXISTING CONTROL

Token Limits

Cost Tracking

Tool Permissions

Retry Limits

Context Compression

Usage Logs

Post-Use Reporting

Technical Possibility

NORMORA CONTROL

Execution Boundaries

Expenditure Governance

Authority Assignment

Evidence-Driven Retry

Context Admission

Contract-to-Receipt Traceability

Pre-Execution Decision

Operational Justification

From Request to Controlled Execution

- 1 Task Intent**
Define the intended outcome and operating value.
- 2 Scope Envelope**
Define what may be inspected, changed, or protected.
- 3 Authority Assignment**
Set autonomy and approval requirements.
- 4 Expenditure Gate**
Determine whether execution is justified.
- 5 Controlled Execution**
Permit only authorized actions and iterations.
- 6 Evidence Receipt**
Record what occurred and what proves completion.

Each transition is a decision boundary.

Six Domains Govern the Execution Boundary

1**Outcome Justification**

Why the task should exist and what defines completion.

2**Scope Envelope**

Which systems, files, tools, and data are authorized.

3**Authority**

What the agent may inspect, propose, modify, verify, or deploy.

4**Iteration**

How many analysis, modification, verification, and retry cycles are permitted.

5**Evidence**

What must change before continuation, retry, or completion.

6**Stop & Re-entry**

When execution must stop and what approval allows it to resume.

No New Evidence, No New Expenditure.

Autonomy Must Be Assigned, Not Assumed

A0	Observe	Inspect and report only.
A1	Propose	Recommend without changing state.
A2	Constrained Act	Modify within a fixed boundary.
A3	Act & Verify	Execute and validate within scope.
A4	Conditional Autonomy	Proceed under explicit policy conditions.

The correct level is the minimum authority required to achieve the intended outcome. More autonomy is not evidence of greater maturity.

Execution Is a Conditional State

READY	Conditions are sufficiently defined for controlled execution.
--------------	---

CONDITIONAL	Proceed only with specified limits, evidence, or approvals.
--------------------	---

HOLD	Material information, authority, or controls remain unresolved.
-------------	---

NOT JUSTIFIED	Expected operational value does not justify the expenditure.
----------------------	--

A Gate Decision does not guarantee technical success. It determines whether responsible execution conditions are sufficiently structured.

Control Requires a Before-State and an After-State

Agent Expenditure Contract

The pre-execution record of authorized purpose, protected state, scope, authority, iteration, evidence, approval, and stop conditions.

Agent Expenditure Receipt

The post-execution record of actual actions, deviations, retries, evidence, expenditure signals, and completion state.

Together, the two records make agent execution reviewable as an operating decision - not an opaque usage event.

A Controlled Execution Example

USER REQUEST	Review this PDF, identify the material risks, and prepare a management summary.
OUTCOME	A concise executive summary covering confirmed risks, unresolved issues, and decision requirements.
SCOPE	Uploaded PDF only. No external research. No modification of source files. No assumptions beyond the document.
AUTHORITY	A1 - Analyze and propose only.
ITERATION	Initial analysis: 1. Verification pass: 1. Retry only if a material omission is identified.
EVIDENCE	Every stated risk must reference a source section or page.
GATE	READY
STOP CONDITION	Stop if the document is incomplete, corrupted, or requires external evidence.

Analyze → Verify → Produce Summary

A Model-Agnostic Control Layer

The framework can begin as a manual protocol and progress toward machine-enforced controls.

CURRENT

Manual Protocol

Controlled prompts, contracts, approvals, and receipts.

NEXT

Policy Wrapper


Contract compilation, scope checks, retry controls, and approvals.

FUTURE

Control Layer

Runtime enforcement across agent tools and operating environments.

Connected to Normora Operating Architecture



BEOO defines business and operating purpose. A-WBS structures scope and work. AI Pre-Execution Gate evaluates readiness. Agent Expenditure Governance governs the conditions under which agent resources may be consumed.

DHEX is Normora's protected judgment architecture for structuring decision states before execution. Its internal encoding and implementation logic are not disclosed in this public brief.

The Concept Is Public. The Decision Logic Is Protected.

PUBLIC

Purpose and category definition

Six control domains

Authority-level model

Gate vocabulary

Contract and Receipt concepts

Implementation direction

PROTECTED

Scoring and weighting

Thresholds and critical conditions

Evidence Sufficiency logic

Contract Compiler rules

Enforcement mappings

Automated decision logic

The public document establishes the category and architecture without disclosing the operating logic required to reproduce the system.

Test the Framework Before Making Performance Claims

The first objective is not to prove a percentage reduction in tokens. It is to determine whether pre-execution control reduces unjustified exploration, scope expansion, repeated hypotheses, rework, and recovery cycles across comparable tasks.

User request count	Agent runs	Files touched
Scope deviations	Retries	Rollbacks
Completion time	Observable credit use	Result satisfaction

No token-reduction claim should be made until repeated controlled trials produce defensible evidence.

Govern the Decision Before Managing the Cost.

Every token is the consequence of an execution decision.

Normora Agent Expenditure Governance does not begin with the question of how cheaply an agent can act. It begins with whether the action is justified, how far it may proceed, and what must be true before it continues.