Publication Information

Constitutional Admissibility Specification (CAS-01)

Professional Release v1.0

Publisher: ManChine AI Technology

Document Status: Public Conceptual Specification

Classification: Public Release

Website: https://manchine.ai

This publication defines conceptual architectural principles only and does not establish implementation or conformance requirements.

ManChine

AI TECHNOLOGY

CAS-01

Constitutional Admissibility Specification

Public Edition v1.1 — Professional Release

May 12, 2026

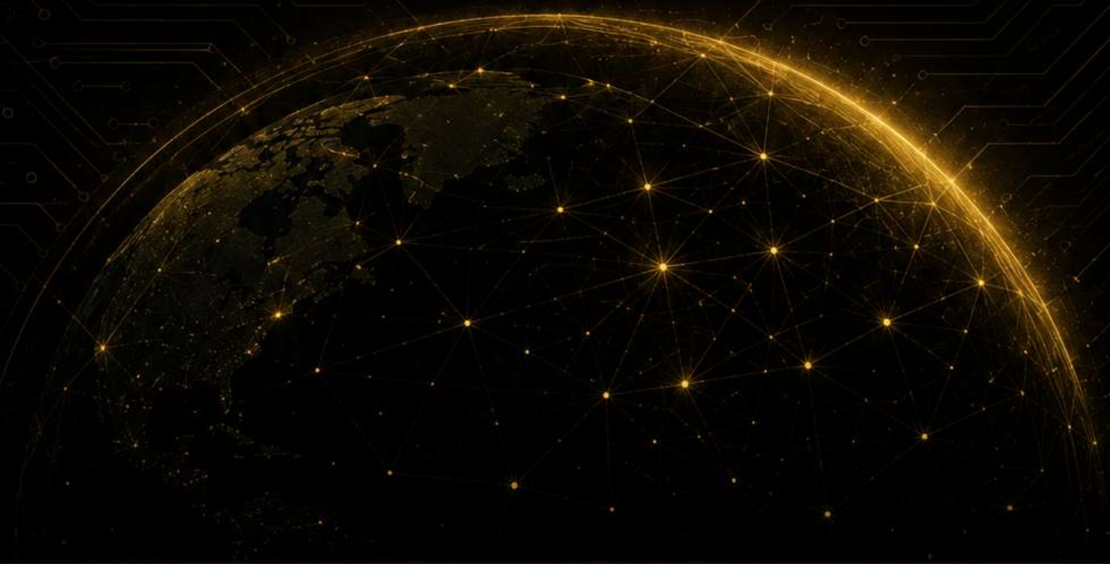

ManChine AI Technology

Governance-layer software for AI runtime integrity, continuity, and consequential governance.

Humans First.
Continuity Always.
Architecture Must Survive.

Runtime Integrity Control

1. Executive Summary

Artificial intelligence is increasingly entrusted with decisions and actions that produce consequential real-world outcomes. As these systems become more autonomous and more deeply integrated into enterprise, government, healthcare, finance, and critical infrastructure, expectations for trustworthy AI governance continue to expand. Existing governance frameworks have significantly advanced the disciplines of risk management, ethics, compliance, safety, and human oversight. These efforts provide an essential foundation for responsible AI deployment.

Despite this progress, much of today's AI governance remains concentrated around activities performed before deployment or during periodic review. Systems are evaluated, risks are assessed, policies are approved, and authorization is granted. Once operational, however, the conditions that originally justified execution may evolve. Authority may change. Policies may be revised. Evidence may become outdated. Rights may be modified. Operational context may shift. An AI system may therefore continue functioning correctly while the governance conditions supporting its consequential actions no longer remain valid.

The **Constitutional Admissibility Specification (CAS)** introduces a conceptual runtime governance architecture intended to address this challenge. Rather than treating governance as a single approval event, CAS views governance as a continuous architectural responsibility that remains connected to consequential execution throughout an AI system's operational life.

At the center of this architecture is a simple principle:

An AI system should not only be capable of executing a consequential action—it should be capable of demonstrating that the action remains justified at the moment it is executed.

This principle distinguishes **operational continuity** from **legitimacy continuity**. Operational continuity focuses on preserving the technical ability of a system to function. Legitimacy continuity focuses on preserving the governance conditions that justify consequential execution as authority, evidence, rights, policy, and operational context evolve. Together, these concepts establish a broader perspective on trustworthy AI, recognizing that reliable execution alone does not necessarily ensure justified execution.

To support this perspective, CAS introduces three complementary architectural concepts:

1. **The Constitutional Execution Model**, which identifies the governance elements that collectively determine whether consequential execution is justified.
2. **The Runtime Execution Boundary (REB)**, which defines the conceptual boundary at which governance evaluation and consequential execution intersect.
3. **Runtime Integrity Control (RiCo)**, which represents the conceptual runtime capability that maintains awareness of governance-relevant conditions throughout system operation.

Together, these concepts establish an architectural foundation for continuous governance without prescribing a specific implementation, technology stack, or organizational model.

This Public Edition is intentionally **conceptual**. It defines the architectural principles, terminology, and governance model of Constitutional Admissibility while avoiding implementation-specific requirements. Topics such as conformance, runtime state management, engineering interfaces, failure behavior, and reference implementations are intentionally reserved for future engineering specifications.

The intended audience includes AI practitioners, enterprise architects, governance professionals, researchers, policymakers, standards communities, and organizations exploring approaches to runtime AI governance. The specification is designed to complement existing governance frameworks rather than replace them, providing a common architectural vocabulary for discussing continuous legitimacy during consequential execution.

As AI systems assume increasingly consequential roles across society, the question facing governance is no longer limited to whether a system should be approved for deployment. Increasingly, it is whether each consequential action remains justified as the conditions surrounding execution continue to evolve.

The Constitutional Admissibility Specification is offered as one architectural contribution toward answering that question.

2. Scope

2.1 Purpose

The Constitutional Admissibility Specification (CAS) establishes a conceptual architectural framework for continuous governance of consequential artificial intelligence (AI) systems during runtime. Its purpose is to define the principles, terminology, and architectural concepts necessary to evaluate whether consequential AI execution remains justified as operational conditions evolve.

Unlike governance approaches that focus primarily on development, deployment, or periodic oversight, CAS extends governance into the operational life of an AI system by treating legitimacy as a continuous architectural concern rather than a point-in-time authorization.

This specification introduces a common conceptual vocabulary and architectural model intended to support discussion, research, implementation, and future standardization efforts related to runtime AI governance. It is not intended to prescribe a single technical implementation or organizational model. Instead, it provides a conceptual foundation upon which diverse governance architectures may be developed.

2.2 Scope

This specification defines the conceptual architecture of Constitutional Admissibility and its relationship to continuous runtime governance.

Specifically, this document:

1. Defines Constitutional Admissibility as a conceptual model for evaluating the legitimacy of consequential AI execution.
2. Introduces the Constitutional Execution Model and the governance elements that contribute to justified execution.
3. Defines the Runtime Execution Boundary (REB) as the conceptual boundary between governance evaluation and consequential execution.
4. Introduces Runtime Integrity Control (RiCo) as the conceptual runtime capability supporting governance awareness.
5. Distinguishes operational continuity from legitimacy continuity as complementary dimensions of trustworthy AI.
6. Provides a conceptual vocabulary for discussing runtime governance across organizational, regulatory, and technical domains.
7. Illustrates the application of these concepts through representative examples and architectural discussion.

The concepts presented in this specification are implementation-independent and are intended to be applicable across multiple AI domains, deployment environments, and governance models.

2.3 Out of Scope

This Public Edition intentionally limits its scope to conceptual architecture.

Accordingly, this specification does **not** define:

1. implementation-specific software architectures;
2. programming interfaces (APIs);
3. communication protocols;
4. runtime algorithms or decision logic;
5. conformance requirements or certification criteria;
6. performance metrics or benchmarking methods;
7. security controls or cybersecurity architectures;
8. governance state machines or execution workflows;
9. fail-open, fail-closed, or fault-tolerance strategies;
10. implementation-specific policy engines or rule systems.

These topics are expected to be addressed in future engineering specifications, implementation guidance, or reference architectures derived from the Constitutional Admissibility framework.

2.4 Intended Audience

This specification is intended for individuals and organizations involved in the design, governance, implementation, evaluation, and oversight of consequential AI systems.

The intended audience includes:

1. AI architects and system designers;
2. software engineers and platform developers;
3. enterprise technology leaders;
4. AI governance and risk professionals;

5. compliance and regulatory specialists;
6. cybersecurity and resilience practitioners;
7. researchers and academic institutions;
8. standards development organizations;
9. policymakers and public-sector stakeholders.

While this document introduces new architectural concepts, it is written to encourage collaboration across technical, organizational, and regulatory communities. It is intended to provide a common conceptual foundation for discussing continuous governance without requiring adherence to a specific technology, vendor, or implementation approach.

3. Terminology

3.1 Definitions

This specification introduces a conceptual vocabulary intended to support consistent discussion of runtime AI governance. The following definitions establish the core architectural terms used throughout this document.

These definitions are descriptive rather than normative and are intended to provide a common conceptual foundation for the Constitutional Admissibility architecture.

Constitutional Admissibility

Constitutional Admissibility is the principle that consequential AI execution should occur only when the authority, evidence, rights, policy, and operational context required to justify that execution remain valid at the moment the consequence is produced.

Within this specification, Constitutional Admissibility is presented as a conceptual runtime governance architecture rather than a specific implementation or compliance framework.

Consequential Execution

Consequential Execution is the execution of an AI-supported action that produces, or is capable of producing, material real-world effects on individuals, organizations, infrastructure, public institutions, or society.

Consequential execution extends beyond computational output. It encompasses actions that may influence rights, obligations, legal status, financial outcomes, safety, healthcare, access to services, regulatory compliance, or other outcomes with meaningful operational or societal impact.

Not all AI operations are consequential. Constitutional Admissibility is specifically concerned with those executions whose outcomes carry governance significance.

Governance Determination

A **Governance Determination** is the conceptual governance outcome reached after evaluating the relevant authority, rights, evidence, policy, and operational context associated with a proposed consequential execution.

A Governance Determination establishes whether execution remains justified under current governance conditions. Depending on organizational policy, the determination may permit execution, constrain execution, defer execution, require additional review, or prevent execution.

This specification intentionally does not prescribe how Governance Determinations are implemented or which organizational or technical authority performs them.

Runtime Governance

Runtime Governance is the continuous application of governance principles throughout the operational life of an AI system.

Unlike governance activities performed exclusively during development or deployment, Runtime Governance remains connected to consequential execution by continuously considering governance-relevant conditions as they evolve.

Within the Constitutional Admissibility architecture, Runtime Governance represents the ongoing relationship between governance and execution rather than a single approval event.

Material Governance Change

A **Material Governance Change** is a change in governance-relevant conditions that may reasonably affect the justification for consequential execution.

Examples of Material Governance Change may include changes to:

1. authority or authorization;
2. applicable policies or regulations;
3. rights, consent, or permissions;
4. supporting evidence;
5. operational context;
6. organizational governance requirements.

This specification defines the concept of Material Governance Change but does not prescribe implementation-specific thresholds or evaluation criteria. Such criteria are expected to be established by organizational governance policies or future engineering specifications.

3.2 Terminology Conventions

To promote consistency throughout this specification, the following terminology conventions are used:

1. **Operational Continuity** refers to the ability of a system to remain technically capable of performing its intended functions.
2. **Legitimacy Continuity** refers to the continuous preservation of the governance conditions that justify consequential execution.
3. **Runtime Execution Boundary (REB)** refers to the conceptual architectural boundary at which governance evaluation and consequential execution intersect.
4. **Runtime Integrity Control (RiCo)** refers to the conceptual runtime capability that maintains awareness of governance-relevant conditions throughout system operation.
5. **Governance** refers to the policies, authority, rights, evidence, and contextual considerations that collectively determine whether consequential execution remains justified.

Unless otherwise stated, references to **governance** within this specification refer to governance as an architectural function rather than a specific organizational body, technology, or implementation.

Part II — The Architecture

4. Constitutional Admissibility

4.1 Definition

Constitutional Admissibility is a conceptual runtime governance architecture that maintains the legitimacy of consequential AI execution by continuously evaluating whether the governance conditions supporting that execution remain justified as operational conditions evolve.

Unlike governance approaches that primarily evaluate AI systems before deployment, Constitutional Admissibility extends governance throughout the operational life of a system by recognizing that authority, evidence, rights, policy, and operational context are dynamic rather than static.

Within this architecture, consequential execution is viewed as a governance event rather than solely a computational event. Each consequential action should remain supported by governance conditions that are valid at the moment execution occurs.

Accordingly, Constitutional Admissibility is founded upon the following principle:

An AI system should not only be capable of executing a consequential action—it should be capable of demonstrating that the action remains justified at the moment it is executed.

This specification presents Constitutional Admissibility as a conceptual architectural framework. It defines governance principles and architectural relationships without prescribing implementation-specific technologies, algorithms, or organizational structures.

4.2 Objectives

The Constitutional Admissibility Specification establishes an architectural foundation for continuous runtime governance of consequential AI systems.

Its primary objectives are to:

1. establish a common conceptual vocabulary for runtime AI governance;
2. define the governance principles that support justified consequential execution;
3. promote governance as a continuous architectural function rather than a point-in-time approval process;
4. distinguish technical system operation from governance legitimacy;
5. provide a conceptual framework for evaluating consequential execution as operational conditions evolve;

6. support transparent, explainable, and accountable governance practices;
7. complement existing governance, risk management, compliance, and assurance frameworks;
8. encourage future implementation, research, and standardization efforts related to runtime AI governance.

These objectives are intended to support a consistent architectural understanding of Constitutional Admissibility across technical, organizational, and regulatory domains.

4.3 Guiding Principles

The Constitutional Admissibility architecture is founded upon several guiding principles that collectively define its governance philosophy.

Governance Is Continuous

Governance does not conclude when an AI system is deployed. It remains an active architectural function throughout the operational life of consequential AI systems.

Legitimacy Is Dynamic

The legitimacy of consequential execution depends upon current governance conditions rather than historical authorization alone. Governance should remain capable of recognizing when those conditions have changed in ways that affect the justification for execution.

Consequential Execution Requires Justification

Actions that produce material real-world consequences should remain supported by valid authority, evidence, rights, policy, and operational context at the time execution occurs.

Governance Should Be Explainable

Governance outcomes should be understandable, reviewable, and capable of demonstrating why consequential execution was permitted, constrained, deferred, suspended, or denied.

Authority Is Accountable

The authority supporting consequential execution should remain attributable, governed, and subject to organizational accountability throughout system operation.

Trust Requires Continuous Verification

Trustworthy AI cannot rely solely upon assumptions established during development or deployment. Governance should remain capable of verifying that the conditions supporting consequential execution continue to remain valid.

Governance Complements Operation

Operational capability and governance legitimacy represent complementary architectural concerns. Neither replaces the other. Trustworthy AI systems require both the ability to execute and the continuing justification to do so.

4.4 Relationship to Governance

Constitutional Admissibility is intended to complement—not replace—existing AI governance practices.

Organizations already employ governance mechanisms that include ethical principles, risk management, compliance programs, cybersecurity, safety engineering, human oversight, and regulatory controls. These capabilities remain essential and continue to provide the foundation for responsible AI development and deployment.

Constitutional Admissibility extends these foundations by addressing a governance question that becomes increasingly important after deployment:

Does this consequential execution remain justified under current governance conditions?

Rather than introducing an alternative governance framework, Constitutional Admissibility provides a conceptual architectural layer that connects governance to consequential execution throughout an AI system's operational life.

In this architecture:

1. governance establishes the conditions under which consequential execution is justified;
2. runtime governance maintains awareness of those conditions as they evolve;
3. consequential execution remains continuously connected to governance rather than relying solely on historical authorization.

Accordingly, Constitutional Admissibility should be understood as an architectural extension of existing governance frameworks, providing a conceptual model for preserving legitimacy throughout runtime while remaining compatible with diverse organizational, regulatory, and technical approaches to AI governance.

5. The Governance Problem

5.1 Governance Ends at Deployment

Artificial intelligence governance has matured significantly through the development of ethical principles, risk management frameworks, regulatory guidance, compliance programs, and assurance methodologies. Collectively, these approaches provide organizations with structured processes for evaluating AI systems before they are deployed and throughout their lifecycle.

Typical governance activities include:

1. defining organizational policies and acceptable use;
2. assessing technical and operational risks;
3. validating model performance;
4. conducting compliance and regulatory reviews;
5. establishing human oversight mechanisms; and
6. authorizing deployment within approved operational boundaries.

These activities are essential for responsible AI adoption and remain foundational to trustworthy AI.

However, authorization for deployment represents a governance decision made at a specific point in time. Once an AI system enters operational use, the governance conditions supporting that authorization may continue to evolve while the system remains active.

Constitutional Admissibility begins by recognizing that deployment is not the end of governance. It is the beginning of an operational period during which governance conditions may continue to change.

5.2 Runtime Conditions Change

AI systems operate within environments that are inherently dynamic. Organizational priorities shift, regulations evolve, policies are revised, evidence changes, permissions are modified, and operational context continually develops.

Examples of runtime change include:

1. updated laws or regulatory requirements;
2. changes in organizational policy;
3. modification or revocation of delegated authority;

4. changes in user consent or access rights;

5. new evidence affecting decision validity;

6. evolving operational or environmental conditions.

These changes may occur independently of the AI system itself. The software may continue to function exactly as designed while the governance conditions that originally justified consequential execution have materially changed.

As a result, technical correctness alone cannot guarantee continued governance legitimacy.

5.3 Consequences Persist

The effects of consequential AI execution frequently extend beyond the moment an action is performed. Decisions may influence individuals, organizations, infrastructure, legal obligations, financial outcomes, healthcare services, public safety, or access to essential resources.

Once produced, these consequences may continue to affect stakeholders long after execution has occurred.

For this reason, governance cannot be evaluated solely by examining whether an AI system operated correctly from a technical perspective. It must also consider whether the execution remained justified under the governance conditions that existed when the consequence was produced.

The significance of consequential execution therefore lies not only in its computational accuracy but also in its continuing legitimacy.

5.4 Operational Continuity Is Insufficient

Operational continuity refers to the ability of a system to remain available, reliable, and capable of performing its intended technical functions.

These qualities are essential for dependable AI systems. However, operational continuity alone does not establish that consequential execution remains justified.

A system may:

1. remain fully available;

2. continue producing accurate outputs;

3. satisfy performance objectives; and

4. execute without technical fault,

while simultaneously operating under governance conditions that have materially changed.

This distinction illustrates a fundamental architectural gap.

A system can remain operational while no longer remaining legitimate.

Constitutional Admissibility addresses this gap by introducing the concept of **legitimacy continuity**—the continuous preservation of the governance conditions that justify consequential execution throughout runtime.

Accordingly, trustworthy AI requires more than uninterrupted execution. It requires continuous confidence that consequential actions remain justified as authority, evidence, rights, policy, and operational context evolve.

Figure 5.1 — The Governance Gap *(Placeholder)*

This section would be strengthened by a simple architectural diagram illustrating the transition from traditional governance to runtime governance:

Traditional Governance

Policy

↓

Risk Assessment

↓

Approval

↓

Deployment

↓

———————— Runtime ————————

(Governance Visibility Ends)

AI Continues Executing

↓

Conditions Continue Changing

↓

Governance Gap

↓

Constitutional Admissibility

↓

Continuous Runtime Governance

6. Constitutional Execution Model

6.1 Introduction

The Constitutional Execution Model provides the conceptual architecture through which Constitutional Admissibility evaluates the legitimacy of consequential AI execution during runtime.

The model recognizes that consequential execution is not solely a technical event. Rather, it represents the intersection of computational capability and governance legitimacy. While an AI system may be technically capable of performing an action, Constitutional Admissibility requires that the governance conditions supporting that action remain valid at the moment execution occurs.

Accordingly, the Constitutional Execution Model establishes a continuous relationship between governance evaluation and consequential execution. Instead of treating governance as a single authorization performed before deployment, the model maintains governance as an active architectural function throughout system operation.

This model serves as the conceptual foundation for the Runtime Execution Boundary (REB) and Runtime Integrity Control (RiCo), both of which are introduced in the chapters that follow.

6.2 Architectural Overview

Within the Constitutional Execution Model, consequential execution is supported by multiple governance dimensions that collectively determine whether execution remains justified.

These governance dimensions include:

1. **Authority** — Whether the authority supporting execution remains valid.
2. **Evidence** — Whether the information supporting execution remains sufficient and current.
3. **Rights** — Whether applicable permissions, consent, or legal rights remain in effect.
4. **Policy** — Whether execution remains consistent with governing organizational or regulatory policies.
5. **Operational Context** — Whether the conditions surrounding execution remain materially consistent with the circumstances under which execution is justified.

Rather than evaluating these elements independently, the Constitutional Execution Model considers their combined relationship as the basis for determining the admissibility of consequential execution.

This specification does not prescribe how organizations evaluate these governance dimensions. Their implementation may vary according to organizational policy, regulatory obligations, and technical architecture. The model instead defines the conceptual relationship among these elements within a continuous runtime governance architecture.

6.3 Continuous Governance Evaluation

As an AI system operates, governance-relevant conditions may change independently of the system's technical behavior. The Constitutional Execution Model therefore treats governance evaluation as a continuous architectural process rather than a one-time approval event.

When governance conditions remain consistent, consequential execution may continue under the established governance framework.

When governance-relevant conditions change, the architecture provides an opportunity to reassess whether execution remains justified before consequential effects are produced.

This continuous relationship between governance and execution forms the conceptual basis for Constitutional Admissibility.

Figure 6.1 — Constitutional Execution Model

```
            Constitutional Execution Model


              Governance Dimensions

 ┌──────────────────────────────────────────────────┐
 │                                │
 │ Authority  Evidence  Rights  Policy  Operational  │
 │                  Context    │
 │                                │
 └──────────────────────────────────────────────────┘
              │
              │
              ▼
         Governance Determination
              │
```

▼

Runtime Execution Boundary

(REB)

│

▼

Consequential Execution

│

▼

Real-World Consequential Effects

Runtime Integrity Control (RiCo)

Continuously observes governance-relevant conditions

and supports ongoing governance awareness throughout

system operation.

Figure 6.1 — Constitutional Execution Model. *Conceptual relationship between governance dimensions, governance determination, the Runtime Execution Boundary (REB), consequential execution, and Runtime Integrity Control (RiCo).*

6.4 Architectural Significance

The Constitutional Execution Model establishes the architectural foundation upon which the remainder of this specification is built.

It introduces three key concepts that distinguish Constitutional Admissibility from traditional deployment-centric governance approaches:

1. **Governance remains connected to consequential execution throughout runtime.**
2. **Consequential execution is evaluated through the combined consideration of governance dimensions rather than technical capability alone.**
3. **Runtime governance supports the continuous preservation of legitimacy as operational conditions evolve.**

The following chapters expand upon the principal architectural components introduced by this model. Chapter 7 defines the Runtime Execution Boundary (REB), while Chapter 8 introduces Runtime Integrity Control (RiCo) as the conceptual runtime capability that maintains governance awareness across system operation.

7. Runtime Execution Boundary (REB)

7.1 Purpose

The **Runtime Execution Boundary (REB)** defines the conceptual architectural boundary at which governance evaluation and consequential AI execution intersect.

Its purpose is to establish a governance-aware point within the execution architecture where the continued legitimacy of consequential execution may be evaluated before real-world consequences are produced.

The REB does not prescribe how governance evaluations are performed, nor does it define a specific implementation mechanism. Instead, it provides an architectural reference point that connects governance with execution, ensuring that consequential actions remain continuously associated with the governance conditions that justify them.

Within the Constitutional Admissibility architecture, the REB serves as the transition between governance determination and consequential execution.

7.2 Concept

Traditional AI architectures often separate governance activities from operational execution. Governance establishes authorization during development or deployment, while runtime execution proceeds independently unless interrupted by external intervention.

The Runtime Execution Boundary introduces a different architectural perspective.

Rather than treating governance as an activity that concludes before execution begins, the REB represents the conceptual location where governance remains continuously connected to consequential execution throughout system operation.

Consequential execution therefore becomes more than a technical operation. It becomes an execution event that remains associated with current governance conditions.

The REB should not be interpreted as a physical component, network boundary, software module, or hardware device. It is an architectural abstraction that identifies where governance evaluation and consequential execution conceptually meet.

7.3 The Boundary

Within the Constitutional Admissibility architecture, the Runtime Execution Boundary separates two complementary domains:

Governance Domain

The governance domain encompasses the information and conditions that establish whether consequential execution remains justified. These include, but are not limited to:

1. authority;
2. evidence;
3. rights;
4. policy; and
5. operational context.

Execution Domain

The execution domain encompasses the AI system's operational capability to perform consequential actions that produce real-world effects.

The Runtime Execution Boundary conceptually connects these domains, ensuring that governance remains relevant to execution rather than existing only as historical authorization.

Importantly, the REB does not replace operational control systems or execution logic. Instead, it provides the architectural relationship through which governance and execution remain continuously aligned.

Figure 7.1 — Runtime Execution Boundary (REB)

```
          Runtime Execution Boundary (REB)


    Governance Domain           Execution Domain


┌──────────────────────────────────┐   ┌──────────────────────────────┐
│                 │   │                │
│ Authority          │   │ AI Runtime         │
│ Evidence           │   │ Decision Logic       │
│ Rights            │   │ Execution Engine      │
│ Policy            │   │                │
│ Operational Context      │   │                │
```

▼

Runtime Execution Boundary (REB)

▼

Consequential Execution

▼

Real-World Consequences

Figure 7.1 — Runtime Execution Boundary (REB). *Conceptual architectural boundary connecting governance evaluation with consequential AI execution.*

7.4 Relationship to Constitutional Admissibility

The Runtime Execution Boundary is a foundational architectural component of the Constitutional Admissibility framework.

Constitutional Admissibility establishes the governance principle that consequential execution should remain justified throughout runtime. The REB provides the conceptual location within the architecture where that principle is applied.

The relationship between these concepts can be summarized as follows:

1. **Constitutional Admissibility** defines the governing principle.
2. **The Constitutional Execution Model** defines the architectural relationship between governance and execution.
3. **The Runtime Execution Boundary (REB)** identifies where governance and consequential execution conceptually intersect.

4. **Runtime Integrity Control (RiCo)** maintains continuous awareness of governance-relevant conditions that support this relationship.

Together, these concepts form an integrated architectural framework in which governance is treated as a continuous characteristic of consequential execution rather than a static authorization established before deployment.

8. Runtime Integrity Control (RiCo)

8.1 Purpose

Runtime Integrity Control (RiCo) is the conceptual runtime governance capability that maintains continuous awareness of governance-relevant conditions throughout the operational life of an AI system.

Its purpose is to support Constitutional Admissibility by enabling the architecture to recognize when governance conditions associated with consequential execution may have materially changed.

RiCo does not make governance decisions, replace organizational authority, or prescribe implementation-specific technologies. Rather, it provides the continuous governance awareness necessary to preserve the relationship between governance and consequential execution established by the Constitutional Execution Model.

Within the Constitutional Admissibility architecture, RiCo serves as the mechanism through which governance remains an active architectural function during runtime.

8.2 Continuous Validation

Governance conditions are not static. Authority may change, policies may be revised, evidence may evolve, rights may be modified, and operational context may shift throughout the lifetime of an AI system.

RiCo supports the continuous evaluation of governance-relevant conditions by maintaining awareness of these changes as they occur.

Continuous validation does not imply that every execution requires complete reauthorization. Instead, it reflects the architectural principle that consequential execution should remain supported by governance conditions that continue to justify execution at the moment it occurs.

The frequency, method, and implementation of validation are intentionally outside the scope of this specification and may vary according to organizational requirements and system architecture.

8.3 Runtime Observation

RiCo continuously observes governance-relevant conditions that may influence the legitimacy of consequential execution.

Examples of governance-relevant observations include:

1. changes in delegated authority;

2. updates to organizational or regulatory policy;
3. modifications to permissions or consent;
4. changes in supporting evidence;
5. significant changes in operational context; and
6. other material governance changes defined by organizational policy.

Observation alone does not determine whether execution should proceed.

Instead, RiCo provides governance awareness that enables the Constitutional Admissibility architecture to recognize when governance conditions may warrant reassessment.

This distinction preserves the separation between governance observation and governance determination.

8.4 Governance Continuity

The primary architectural contribution of RiCo is the preservation of **governance continuity** throughout runtime.

Traditional governance architectures often establish governance before deployment and periodically review it afterward. Between these events, governance visibility may be limited while execution continues uninterrupted.

RiCo addresses this architectural gap by maintaining continuous awareness of governance-relevant conditions throughout system operation.

As a result, governance remains connected to consequential execution rather than existing solely as historical authorization.

This continuous relationship supports the broader objective of **legitimacy continuity**, ensuring that governance remains an active consideration as consequential actions are performed.

Figure 8.1 — Runtime Integrity Control (RiCo)

Runtime Integrity Control (RiCo)

Governance-Relevant Conditions

Authority Rights Policy

Evidence Context Organizational Rules

▼

Continuous Runtime Observation

▼

Detection of Material Governance Change

▼

Governance Awareness Maintained

▼

Supports Constitutional Admissibility

▼

Runtime Execution Boundary (REB)

▼

Consequential AI Execution

Figure 8.1 — Runtime Integrity Control (RiCo). *Conceptual runtime capability that maintains governance awareness by observing governance-relevant conditions and supporting the continuous relationship between governance and consequential execution.*

8.5 Relationship to the Runtime Execution Boundary (REB)

Runtime Integrity Control and the Runtime Execution Boundary perform complementary architectural functions within the Constitutional Admissibility framework.

The **Runtime Execution Boundary (REB)** defines the conceptual boundary at which governance evaluation and consequential execution intersect.

Runtime Integrity Control (RiCo) maintains continuous awareness of the governance conditions that inform activity at that boundary.

Their relationship can be summarized as follows:

1. **CAS** establishes the governing principle of continuous legitimacy.
2. **The Constitutional Execution Model** defines the architectural relationship between governance and execution.
3. **REB** identifies where governance and consequential execution intersect.
4. **RiCo** maintains awareness of governance-relevant conditions throughout runtime, enabling that relationship to remain meaningful as operational conditions evolve.

RiCo therefore complements the REB rather than replacing it. The REB defines the architectural boundary; RiCo supports the continuity of governance awareness that preserves the integrity of that boundary over time.

9. Operational Continuity vs. Legitimacy Continuity

9.1 Introduction

Trustworthy AI has traditionally emphasized the continuous operation of technical systems. Reliability, availability, resilience, and performance are essential characteristics of dependable AI architectures and remain fundamental objectives of system engineering.

Constitutional Admissibility introduces a complementary architectural perspective.

In addition to maintaining operational capability, consequential AI systems should preserve the governance conditions that justify their actions throughout runtime.

This distinction gives rise to two complementary forms of continuity:

1. **Operational Continuity**, which preserves the technical ability of a system to execute; and
2. **Legitimacy Continuity**, which preserves the governance conditions that justify consequential execution.

Together, these concepts provide a more complete architectural understanding of trustworthy AI.

9.2 Comparison

Operational Continuity and Legitimacy Continuity address different, but complementary, dimensions of AI system trustworthiness.

Operational Continuity asks:

Can the system continue operating?

Legitimacy Continuity asks:

Should this consequential execution continue under current governance conditions?

Neither perspective replaces the other.

Reliable AI systems require both continuous operation and continuous governance legitimacy.

Table 9.1 — Operational Continuity vs. Legitimacy Continuity

Operational Continuity	Legitimacy Continuity
Focuses on technical operation	Focuses on governance legitimacy

Operational Continuity	Legitimacy Continuity
Preserves availability	Preserves justified execution
Measures reliability and resilience	Evaluates authority, evidence, rights, policy, and context
Addresses system failures	Addresses governance changes
Concerned with execution capability	Concerned with execution justification
Answers: *Can the system execute?*	Answers: *Should the system execute?*
Maintains operational performance	Maintains governance integrity
Technical engineering discipline	Runtime governance discipline

Table 9.1 — Comparison of Operational Continuity and Legitimacy Continuity.

9.3 Discussion

Operational continuity has long been recognized as an essential characteristic of dependable computing systems. High availability, fault tolerance, redundancy, disaster recovery, and resilience all contribute to ensuring that systems remain capable of performing their intended functions.

These capabilities remain indispensable.

However, operational capability alone does not establish that consequential execution remains justified.

An AI system may continue to:

1. execute correctly;
2. produce accurate outputs;
3. satisfy performance objectives; and
4. remain fully available,

while the governance conditions supporting its consequential actions have materially changed.

Examples include changes in delegated authority, regulatory requirements, organizational policy, user rights, supporting evidence, or operational context.

In such situations, the system's operational continuity remains intact, but its legitimacy continuity may no longer be preserved.

Constitutional Admissibility therefore extends traditional notions of continuity by recognizing governance legitimacy as an architectural characteristic that should be maintained alongside technical operation.

9.4 Implications

Recognizing the distinction between operational continuity and legitimacy continuity has several architectural implications.

First, governance should not be viewed solely as a deployment activity. Governance becomes a continuous architectural function that remains associated with consequential execution throughout runtime.

Second, AI assurance extends beyond demonstrating technical correctness. Trustworthy AI requires confidence that consequential actions remain supported by valid governance conditions when those actions occur.

Third, runtime governance architectures should be capable of recognizing material governance changes without requiring interruption of normal system operation. This enables organizations to preserve both operational resilience and governance integrity.

Finally, Constitutional Admissibility does not replace established engineering disciplines concerned with reliability, resilience, or safety. Rather, it complements them by introducing legitimacy continuity as an additional dimension of trustworthy AI.

Together, operational continuity and legitimacy continuity provide a more comprehensive foundation for governing consequential AI systems in dynamic operational environments.

Figure 9.1 — Operational Continuity vs. Legitimacy Continuity

Traditional View

System Operational?

│

Yes ───┘

│

▼

Continue Execution

Constitutional Admissibility View

```
 System Operational?
        │
 Yes ───┘
        │
        ▼
Governance Conditions Valid?
        │
 Yes ───┴─── No
   │      │
   ▼      ▼
Continue   Governance
Execution   Reassessment
```

Operational Continuity and Legitimacy Continuity. *Operational continuity determines whether a system can continue functioning. Legitimacy continuity determines whether consequential execution remains justified under current governance conditions.*

Part III — Application

10. Example Applications

10.1 Introduction

The Constitutional Admissibility architecture is designed to be independent of any specific industry, technology platform, organizational structure, or regulatory environment.

Wherever artificial intelligence performs consequential actions, governance conditions may evolve independently of system operation. Constitutional Admissibility provides a conceptual framework for maintaining the legitimacy of those actions throughout runtime.

The following examples illustrate how the architectural principles introduced in this specification may be applied across representative domains. These examples are illustrative rather than prescriptive and are intended to demonstrate the general applicability of the Constitutional Admissibility framework.

10.2 Financial Services

Financial institutions increasingly rely on AI to support credit decisions, fraud detection, transaction authorization, investment management, and regulatory compliance.

While these systems may operate reliably from a technical perspective, governance conditions can change continuously through evolving regulations, customer authorizations, internal policies, or newly available financial information.

Within the Constitutional Admissibility architecture:

1. authority may reflect delegated approval authority or regulatory authorization;
2. evidence may include updated financial information or transaction data;
3. rights may include customer consent or legal permissions;
4. policy may reflect institutional risk controls or regulatory obligations; and
5. operational context may include changing market conditions or fraud indicators.

By maintaining awareness of these governance conditions throughout runtime, Constitutional Admissibility supports the continued legitimacy of consequential financial actions as conditions evolve.

10.3 Healthcare

Healthcare AI increasingly assists clinicians through diagnostic support, treatment recommendations, patient prioritization, and resource allocation.

These applications frequently operate in environments where patient information, clinical evidence, treatment protocols, consent, and regulatory guidance may change over time.

Within this context, Constitutional Admissibility supports the continuous relationship between governance and consequential execution by recognizing changes that may affect the justification for AI-assisted clinical actions.

The architecture does not replace clinical judgment or medical governance. Rather, it provides a conceptual framework through which governance considerations remain connected to consequential execution throughout runtime.

10.4 Critical Infrastructure

AI systems supporting energy, transportation, communications, manufacturing, and public utilities often operate continuously within complex operational environments.

Infrastructure conditions may change rapidly due to equipment failures, environmental events, cyber incidents, operational priorities, or emergency response requirements.

Constitutional Admissibility enables governance considerations to remain associated with consequential operational decisions as these conditions evolve.

By maintaining governance awareness throughout runtime, organizations may better align operational resilience with the continuing legitimacy of consequential infrastructure actions.

10.5 Government

Government agencies increasingly employ AI to support administrative services, regulatory processes, public resource allocation, investigations, and decision support.

These systems often operate within governance environments characterized by changing legislation, administrative policy, delegated authority, procedural requirements, and legal rights.

Constitutional Admissibility provides a conceptual framework for maintaining awareness of these governance conditions as consequential governmental actions are considered.

The architecture supports governance continuity without prescribing how governmental authority or legal oversight should be implemented within individual jurisdictions.

10.6 Autonomous Systems

Autonomous systems—including robotics, unmanned platforms, industrial automation, and intelligent transportation systems—may perform consequential actions with varying levels of human supervision.

Operational environments for these systems are inherently dynamic. Mission objectives, environmental conditions, delegated authority, safety requirements, and applicable policies may change during operation.

Within the Constitutional Admissibility architecture, runtime governance remains conceptually connected to consequential execution despite these evolving conditions.

This enables governance legitimacy to be considered alongside operational capability as autonomous systems continue to perform consequential actions.

10.7 Architectural Perspective

Although these examples represent different industries, they illustrate a common architectural principle.

In every domain:

1. operational conditions evolve;
2. governance conditions may change independently of system operation;
3. consequential actions continue to produce real-world effects; and
4. maintaining technical operation alone does not ensure continued governance legitimacy.

Constitutional Admissibility provides a conceptual runtime governance architecture that addresses this common challenge without prescribing domain-specific implementations.

Accordingly, the framework is intended to support a broad range of consequential AI systems while remaining adaptable to diverse organizational, regulatory, and technical environments.

Figure 10.1 — Constitutional Admissibility Across Domains

Constitutional Admissibility

Governance

Continuity

Financial Healthcare Critical Government Autonomous

Services Infrastructure Systems

Consequential AI Execution

Constitutional Admissibility Across Domains. *The Constitutional Admissibility architecture is domain-independent and provides a common conceptual framework for maintaining governance continuity across diverse categories of consequential AI systems.*

11. Benefits

11.1 Introduction

The Constitutional Admissibility architecture is intended to complement existing AI governance practices by extending governance into the operational life of consequential AI systems.

Rather than replacing established approaches to risk management, compliance, safety, or system engineering, Constitutional Admissibility introduces an architectural perspective that emphasizes the continuous legitimacy of consequential execution as governance conditions evolve.

The benefits described in this section are conceptual and reflect the intended architectural contributions of the framework.

11.2 Enterprise

For organizations deploying consequential AI systems, Constitutional Admissibility provides a conceptual framework for integrating governance with runtime operations.

Potential architectural benefits include:

1. promoting continuous governance awareness throughout system operation;
2. supporting organizational accountability for consequential AI actions;
3. improving the alignment between operational processes and governance policies;
4. enabling governance considerations to evolve alongside changing business, regulatory, and operational environments; and
5. providing a common governance vocabulary across technical and organizational functions.

By maintaining governance as a continuous architectural concern, organizations may strengthen confidence that consequential AI systems continue operating within their intended governance boundaries.

11.3 Developers

For software engineers and AI developers, Constitutional Admissibility provides an architectural model that separates technical implementation from governance responsibility.

This distinction enables development teams to:

1. design systems that recognize governance as a runtime consideration;
2. incorporate governance awareness without prescribing implementation technologies;

3. support explainable governance architectures;
4. prepare systems for future runtime governance capabilities; and
5. establish clearer architectural separation between execution logic and governance functions.

The specification intentionally remains implementation-independent, allowing organizations to adopt approaches appropriate to their own technical environments.

11.4 Architects

Enterprise and solution architects benefit from a conceptual model that explicitly integrates governance into system architecture rather than treating it solely as an external organizational process.

Constitutional Admissibility provides:

1. a common architectural vocabulary for runtime governance;
2. a conceptual relationship between governance and consequential execution;
3. architectural concepts such as the Constitutional Execution Model, Runtime Execution Boundary (REB), and Runtime Integrity Control (RiCo);
4. a framework for discussing governance continuity across complex AI ecosystems; and
5. a foundation for future reference architectures and engineering specifications.

These concepts are intended to support consistent architectural thinking without constraining implementation choices.

11.5 Regulators

For policymakers, regulators, and governance professionals, Constitutional Admissibility offers a conceptual framework for discussing governance beyond initial system approval.

The architecture encourages consideration of questions such as:

1. How should governance remain connected to consequential execution during runtime?
2. How should organizations recognize material governance changes?
3. How should legitimacy be maintained as operational conditions evolve?

4. How can governance decisions remain transparent and reviewable throughout system operation?

The specification does not propose regulatory requirements. Instead, it provides a common conceptual vocabulary that may support future policy development, standards activities, and interdisciplinary collaboration.

11.6 Society

As AI systems assume increasingly consequential roles within society, public trust depends upon more than technical performance.

Citizens, consumers, patients, employees, and public institutions increasingly expect AI systems to operate responsibly, transparently, and within appropriate governance boundaries.

Constitutional Admissibility contributes to this objective by promoting the architectural principle that consequential execution should remain continuously justified rather than relying solely on historical authorization.

While no single governance framework can address every challenge associated with AI deployment, Constitutional Admissibility seeks to contribute to broader discussions regarding accountability, transparency, and responsible AI operation in dynamic environments.

11.7 Summary

The benefits of Constitutional Admissibility arise from its architectural perspective rather than from any specific implementation.

Across technical, organizational, and regulatory communities, the framework provides:

1. a common vocabulary for continuous runtime governance;
2. a conceptual model connecting governance and consequential execution;
3. an architectural distinction between operational continuity and legitimacy continuity; and
4. a foundation for future implementation, research, and standardization.

By treating governance as a continuous characteristic of consequential AI systems, Constitutional Admissibility extends existing governance practices while remaining compatible with diverse technologies, organizational models, and regulatory environments.

12. Relationship to Existing Governance

12.1 Introduction

The Constitutional Admissibility Specification is intended to complement existing AI governance practices rather than replace them.

Organizations have invested significant effort in developing governance frameworks that address ethics, risk management, compliance, safety, security, human oversight, and responsible AI deployment. These frameworks provide essential mechanisms for governing the design, development, deployment, and lifecycle management of AI systems.

Constitutional Admissibility builds upon these foundations by extending governance into the operational life of consequential AI systems. It introduces a conceptual runtime governance architecture that maintains the relationship between governance and consequential execution as operational conditions evolve.

Accordingly, Constitutional Admissibility should be understood as an architectural extension to existing governance approaches rather than an alternative governance model.

12.2 Relationship to Risk Management

Risk management remains a foundational component of trustworthy AI.

Organizations routinely identify, assess, mitigate, monitor, and communicate risks associated with AI systems throughout their lifecycle. These activities help ensure that systems are developed and deployed responsibly and that known risks are appropriately managed.

Constitutional Admissibility does not replace risk management processes.

Instead, it complements them by recognizing that governance-relevant conditions may continue to evolve after deployment. While risk management establishes and maintains an understanding of potential risks, Constitutional Admissibility provides a conceptual architecture for maintaining awareness of whether consequential execution continues to remain justified under changing runtime conditions.

The two approaches address different but complementary governance objectives.

12.3 Relationship to Human Oversight

Human oversight remains an essential element of responsible AI governance.

Organizations may employ human review, approval workflows, escalation procedures, or supervisory controls to ensure that consequential AI actions remain appropriately governed.

Constitutional Admissibility neither replaces nor diminishes the importance of human oversight.

Rather, it provides a conceptual framework through which governance-relevant information may remain continuously associated with consequential execution, enabling human decision-makers to exercise oversight with greater contextual awareness when required.

The architecture is compatible with varying degrees of human involvement, ranging from direct review to governance policies that define when human intervention becomes appropriate.

12.4 Relationship to Policy

Organizational policies establish the principles, rules, and governance expectations that guide the responsible operation of AI systems.

These policies may evolve over time in response to regulatory developments, organizational priorities, operational experience, or emerging societal expectations.

Constitutional Admissibility recognizes policy as one of the governance dimensions contributing to justified consequential execution.

The framework does not prescribe policy content or organizational governance structures. Instead, it provides an architectural model through which policy remains connected to consequential execution throughout runtime.

As policies evolve, the architecture supports continued consideration of their relevance to ongoing consequential AI actions.

12.5 Relationship to Compliance

Compliance programs help organizations demonstrate adherence to legal, regulatory, contractual, and organizational obligations.

Constitutional Admissibility is not a compliance framework and does not establish certification requirements, regulatory obligations, or legal standards.

Instead, it provides a conceptual runtime governance architecture that may complement compliance activities by maintaining awareness of governance-relevant conditions as consequential AI systems operate.

Organizations may implement Constitutional Admissibility alongside existing compliance programs without altering established governance responsibilities.

12.6 Relationship to Existing AI Frameworks

Numerous national and international organizations have developed frameworks that advance trustworthy AI through principles addressing governance, transparency, accountability, safety, security, fairness, risk management, and human oversight.

Constitutional Admissibility is intended to complement these efforts by contributing a conceptual architectural perspective focused on runtime governance and the continuous legitimacy of consequential execution.

Rather than introducing competing governance principles, the framework extends existing discussions by asking an additional architectural question:

If an AI system was appropriately governed at deployment, how can governance remain meaningfully connected to consequential execution as operational conditions continue to evolve?

In this way, Constitutional Admissibility seeks to complement existing governance ecosystems while remaining compatible with diverse organizational, technical, and regulatory approaches.

12.7 Architectural Perspective

The relationship between Constitutional Admissibility and existing governance can be summarized as follows:

Existing Governance	**Constitutional Admissibility**
Establishes governance before and throughout the AI lifecycle	Extends governance into consequential runtime execution
Defines organizational policies and governance processes	Connects those governance conditions to consequential execution
Supports risk identification and mitigation	Supports continuous governance awareness as conditions evolve
Provides human oversight and accountability	Maintains governance continuity throughout runtime
Enables compliance with applicable obligations	Provides a conceptual architecture for runtime legitimacy

Table 12.1 — Relationship between existing governance practices and Constitutional Admissibility.

12.8 Summary

Constitutional Admissibility is not intended to replace existing governance frameworks, organizational policies, risk management practices, compliance programs, or human oversight mechanisms.

Its contribution is architectural.

The framework introduces a conceptual model that maintains the relationship between governance and consequential execution throughout runtime, complementing established governance disciplines while preserving their respective roles and responsibilities.

By extending governance beyond deployment without redefining existing governance structures, Constitutional Admissibility seeks to provide a common architectural foundation for continuous legitimacy within consequential AI systems.

Part IV — Looking Forward

13. Future Directions

13.1 Introduction

The Constitutional Admissibility Specification establishes a conceptual architectural foundation for continuous runtime governance of consequential AI systems. As a Public Edition, its purpose is to define the core principles, terminology, and architectural relationships that underpin Constitutional Admissibility rather than to prescribe implementation-specific solutions.

The concepts introduced throughout this specification represent the beginning of an ongoing program of research, engineering, and architectural development. Future work is expected to refine these concepts through implementation experience, interdisciplinary collaboration, and continued engagement with the broader AI governance community.

13.2 Future Research

The introduction of Constitutional Admissibility raises a number of research questions that extend beyond the scope of this publication.

Areas for future investigation include:

1. methods for evaluating governance legitimacy during runtime;
2. approaches for identifying and responding to material governance changes;
3. governance-aware AI architectures for distributed and multi-agent systems;
4. methods for maintaining governance continuity across complex operational environments;
5. approaches for explaining and auditing runtime governance determinations; and
6. empirical evaluation of runtime governance architectures in operational settings.

The authors encourage continued academic, industrial, and public-sector research to explore these and related questions.

13.3 Runtime Execution Boundary (REB)

This specification introduces the Runtime Execution Boundary (REB) as a conceptual architectural construct.

Future work may expand upon this concept through:

1. engineering reference architectures;

2. implementation patterns;

3. integration models for enterprise AI platforms;

4. distributed execution environments;

5. governance interface definitions; and

6. interoperability considerations across heterogeneous systems.

Such work will help translate the conceptual REB into practical architectural guidance while preserving the implementation independence established by this specification.

13.4 Runtime Integrity Control (RiCo)

Runtime Integrity Control (RiCo) is presented in this specification as the conceptual runtime capability responsible for maintaining governance awareness throughout system operation.

Future development may include:

1. engineering specifications;

2. reference implementations;

3. governance observation models;

4. event and state management architectures;

5. operational validation strategies;

6. explainability mechanisms; and

7. implementation guidance for diverse deployment environments.

These developments are expected to provide practical realization of the architectural concepts introduced in this document.

13.5 Policy Governance Control (PGC)

As runtime governance architectures mature, additional architectural capabilities may emerge to support governance policy management throughout the operational lifecycle.

One such area of future development is **Policy Governance Control (PGC)**, envisioned as a complementary governance architecture concerned with the management, evolution, and application of governance policy.

Future work may explore:

1. policy lifecycle management;
2. governance policy versioning;
3. policy conflict resolution;
4. organizational governance orchestration;
5. policy distribution across AI ecosystems; and
6. integration with Constitutional Admissibility, the Runtime Execution Boundary, and Runtime Integrity Control.

The concepts associated with PGC remain the subject of ongoing architectural research and are outside the scope of this Public Edition.

13.6 Standards Development

The concepts introduced by Constitutional Admissibility are intended to encourage continued discussion within the broader standards and governance communities.

Future standards activities may consider:

1. common terminology for runtime governance;
2. conceptual reference architectures;
3. governance interoperability models;
4. implementation guidance;
5. conformance methodologies;
6. testing and validation approaches; and
7. integration with existing AI governance standards and assurance frameworks.

The Constitutional Admissibility Specification is offered as an architectural contribution to these continuing discussions rather than as a completed standards proposal.

13.7 Implementation

This Public Edition intentionally avoids prescribing implementation-specific technologies or engineering requirements.

Future publications may include:

1. normative engineering specifications;
2. reference architectures;
3. implementation guidance;
4. software development kits (SDKs);
5. application programming interfaces (APIs);
6. conformance criteria;
7. reference implementations; and
8. operational case studies demonstrating Constitutional Admissibility in practice.

These future documents are expected to build upon the conceptual architecture established by this specification while remaining informed by practical implementation experience.

13.8 Looking Ahead

The Constitutional Admissibility Specification represents the first step toward a broader architectural vision for runtime AI governance.

As AI systems continue to assume increasingly consequential roles within society, governance must evolve alongside advances in technical capability. Preserving trust in AI will require not only reliable systems, but architectures capable of maintaining the legitimacy of consequential execution as authority, evidence, rights, policy, and operational context evolve.

The ideas presented in this specification are intended to contribute to that ongoing evolution.

Their continued refinement will depend upon research, implementation, collaboration, and constructive discussion across technical, organizational, academic, regulatory, and standards communities.

14. Conclusion

The Constitutional Admissibility Specification introduces a conceptual architectural framework for maintaining the legitimacy of consequential AI execution throughout runtime.

The specification begins from the observation that while existing AI governance frameworks provide robust approaches to ethics, risk management, compliance, safety, and oversight, governance conditions may continue to evolve after an AI system has been deployed. As authority, evidence, rights, policy, and operational context change, the justification for consequential execution may also change, even when the system continues to operate correctly.

To address this architectural challenge, this specification presents Constitutional Admissibility as a model in which governance remains continuously connected to consequential execution. The Constitutional Execution Model, the Runtime Execution Boundary (REB), and Runtime Integrity Control (RiCo) collectively establish a conceptual framework for preserving governance legitimacy throughout the operational life of consequential AI systems.

A central contribution of this specification is the distinction between **operational continuity** and **legitimacy continuity**. Operational continuity preserves the technical ability of a system to function. Legitimacy continuity preserves the governance conditions that justify consequential execution as those conditions evolve. Together, these complementary concepts provide a broader architectural perspective on trustworthy AI.

This Public Edition is intentionally conceptual. It defines architectural principles, terminology, and relationships without prescribing implementation technologies, engineering requirements, or conformance criteria. Those subjects are expected to be addressed through future engineering specifications, reference architectures, implementation guidance, and continued research.

The Constitutional Admissibility Specification is offered as a contribution to the ongoing development of AI governance. Its purpose is not to replace existing governance frameworks, but to complement them by extending governance into the runtime environment where consequential AI actions continue to produce real-world effects.

As AI systems assume increasingly consequential roles across society, maintaining confidence in their operation will require not only reliable technical execution, but also continued assurance that consequential actions remain justified under the governance conditions that exist when those actions occur.

The Constitutional Admissibility Specification provides one conceptual architectural foundation for pursuing that objective.

Appendix A — Glossary

This appendix provides the authoritative definitions of the principal architectural terms used throughout the Constitutional Admissibility Specification. These definitions are intended to promote consistent interpretation of the concepts presented in this document.

Unless otherwise specified, all terms are used according to the definitions provided below.

Authority

The governance-recognized authorization permitting consequential AI execution under defined organizational, legal, or regulatory conditions.

Constitutional Admissibility

The principle that consequential AI execution should occur only when the authority, evidence, rights, policy, and operational context required to justify that execution remain valid at the moment the consequence is produced.

Constitutional Execution Model

The conceptual architectural model describing the relationship between governance evaluation and consequential AI execution throughout runtime.

Consequential Execution

The execution of an AI-supported action that produces, or is capable of producing, material real-world effects on individuals, organizations, infrastructure, public institutions, or society.

Evidence

Information that supports, justifies, or informs a Governance Determination.

Evidence may include data, observations, documentation, records, or other governance-relevant information appropriate to the operational context.

Governance

The collection of principles, authority, rights, policies, evidence, and contextual considerations that collectively determine whether consequential AI execution remains justified.

Governance Determination

The conceptual governance outcome reached after evaluating the governance conditions associated with a proposed consequential execution.

A Governance Determination establishes whether execution remains justified under current governance conditions.

Governance Domain

The conceptual domain containing governance-relevant information that supports the justification of consequential execution, including authority, evidence, rights, policy, and operational context.

Legitimacy Continuity

The continuous preservation of the governance conditions that justify consequential AI execution throughout runtime.

Material Governance Change

A change in governance-relevant conditions that may reasonably affect the justification for consequential execution.

Examples include changes in authority, policy, rights, evidence, or operational context.

Operational Context

The circumstances surrounding consequential execution that may influence whether governance conditions continue to justify the action.

Operational Continuity

The ability of an AI system to remain technically capable of performing its intended functions through continued availability, reliability, resilience, and execution.

Policy

Organizational, legal, regulatory, or procedural rules that contribute to determining whether consequential execution remains justified.

Rights

Permissions, entitlements, consent, legal protections, or other governance-recognized rights relevant to consequential execution.

Runtime Execution Boundary (REB)

The conceptual architectural boundary at which governance evaluation and consequential AI execution intersect.

Runtime Governance

The continuous application of governance principles throughout the operational life of an AI system.

Runtime Integrity Control (RiCo)

The conceptual runtime governance capability that maintains continuous awareness of governance-relevant conditions throughout the operational life of an AI system.

Trustworthy AI

AI systems whose operation reflects appropriate technical performance together with governance practices that support responsible, accountable, and justifiable consequential execution.

For the purposes of this specification, the definitions contained in this glossary are authoritative. Where a term has alternative meanings in other publications or disciplines, the definition provided here applies unless explicitly stated otherwise.

Appendix B — Architecture Figures

This appendix consolidates the architectural figures presented throughout the Constitutional Admissibility Specification. The figures are reproduced here for convenience and to provide a single reference point for the conceptual architecture described in this document.

Unless otherwise noted, all figures are conceptual and implementation-independent. They are intended to illustrate architectural relationships rather than prescribe specific software, hardware, or organizational implementations.

Figure 5.1

The Governance Gap

Illustrates the transition from traditional deployment-centric governance to continuous runtime governance, highlighting the governance gap that Constitutional Admissibility is intended to address.

Referenced in: Section 5.4

```
Figure 5.1

Traditional Governance
Policy
 │
 ▼
Risk Assessment
 │
 ▼
Approval
 │
 ▼
Deployment
 │
──────── Runtime ────────
(Governance Visibility Ends)
 │
 ▼
AI Continues Executing
 │
 ▼
Conditions Continue Changing
 │
 ▼
Governance Gap
 │
 ▼
Constitutional Admissibility
 │
```

```
 ▼
Continuous Runtime Governance
```

Figure 6.1

Constitutional Execution Model

Illustrates the conceptual relationship between governance dimensions, Governance Determination, the Runtime Execution Boundary (REB), consequential execution, and Runtime Integrity Control (RiCo).

Referenced in: Chapter 6

```
Figure 6.1

+-----------------------------------+
| Governance Dimensions             |
| Authority Evidence Rights Policy |
| Context                           |
+-----------------------------------+
           |
           ▼
Governance Determination
           |
           ▼
Runtime Execution Boundary (REB)
           |
           ▼
Consequential AI Execution
           |
           ▼
Real-World Consequences

RiCo continuously observes governance conditions.
```

Figure 7.1

Runtime Execution Boundary (REB)

Illustrates the conceptual architectural boundary connecting the Governance Domain and the Execution Domain prior to consequential execution.

Referenced in: Chapter 7

```
Figure 7.1

Governance Domain
      │
      ▼
===== Runtime Execution Boundary (REB) =====
      │
      ▼
Execution Domain
      │
      ▼
Consequential Execution
      │
      ▼
Real-World Consequences
```

Figure 8.1

Runtime Integrity Control (RiCo)

Illustrates Runtime Integrity Control as the conceptual runtime capability that maintains governance awareness by observing governance-relevant conditions throughout system operation.

Referenced in: Chapter 8

```
Figure 8.1

Governance Conditions
 │
 ▼
Continuous Observation
 │
 ▼
Material Governance Change
 │
 ▼
Governance Awareness
 │
 ▼
REB
 │
 ▼
Consequential AI Execution
```

Figure 9.1

Operational Continuity vs. Legitimacy Continuity

Compares operational continuity with legitimacy continuity, illustrating how Constitutional Admissibility introduces governance continuity alongside technical continuity.

Referenced in: Chapter 9

```
Figure 9.1

Traditional:
System Operational? -> Yes -> Continue Execution

CAS:
System Operational? -> Governance Conditions Valid?
Yes -> Continue Execution
No -> Governance Reassessment
```

Figure 10.1

Constitutional Admissibility Across Domains

Illustrates the domain-independent application of Constitutional Admissibility across representative sectors including financial services, healthcare, critical infrastructure, government, and autonomous systems.

Referenced in: Chapter 10

```
Figure 10.1

Constitutional Admissibility
 │
 ├─ Financial Services
 ├─ Healthcare
 ├─ Critical Infrastructure
 ├─ Government
 └─ Autonomous Systems
 │
 ▼
Consequential AI Execution
```

Figure Index

Appendix C — Informative References

This appendix identifies selected publications, standards, and governance frameworks that provide context for the concepts discussed throughout this specification.

These references are **informative** rather than **normative**. They are included to assist readers in understanding the broader AI governance landscape and should not be interpreted as requirements of the Constitutional Admissibility Specification.

The inclusion of a reference does not imply endorsement, equivalence, or dependency. Constitutional Admissibility is intended to complement existing governance efforts by contributing a conceptual runtime governance architecture focused on the continuous legitimacy of consequential AI execution.

C.1 AI Governance Frameworks

1. National Institute of Standards and Technology (NIST). *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* U.S. Department of Commerce.
2. ISO/IEC 42001. *Information Technology — Artificial Intelligence — Management System.*
3. ISO/IEC 23894. *Artificial Intelligence — Risk Management.*
4. Organisation for Economic Co-operation and Development (OECD). *OECD AI Principles.*
5. European Union. *Artificial Intelligence Act (EU AI Act).*

C.2 Governance and Assurance

Representative publications addressing governance, assurance, accountability, risk management, transparency, and responsible AI deployment.

Examples include literature related to:

1. AI governance;
2. algorithmic accountability;
3. trustworthy AI;
4. AI assurance;
5. safety engineering;
6. systems governance;

7. operational resilience.

C.3 Runtime Governance Research

Representative academic and industrial research concerning:

1. runtime monitoring;
2. runtime assurance;
3. adaptive governance;
4. dynamic policy enforcement;
5. continuous compliance;
6. governance-aware architectures;
7. trustworthy autonomous systems.

These works provide additional perspectives relevant to the continued development of runtime governance architectures.

C.4 Future Reference Expansion

As Constitutional Admissibility evolves, future editions of this specification may include additional informative references covering:

1. implementation guidance;
2. engineering specifications;
3. runtime governance architectures;
4. distributed AI governance;
5. multi-agent governance;
6. governance interoperability;
7. constitutional AI research;
8. explainability and AI assurance methodologies.

Appendix D — Evolution of Constitutional Admissibility *(Informative)*

This appendix illustrates the conceptual evolution of the Constitutional Admissibility architecture and the anticipated progression from foundational research toward engineering implementation.

The timeline is intended to provide historical context for the development of the architecture and to distinguish conceptual publications from future engineering specifications and reference implementations.

Figure D.1 — Evolution of Constitutional Admissibility

White Paper 001

│

▼

White Paper 002

│

▼

CAS-01 Public Edition

│

▼

CAS-01 Professional Release

│

▼

Internal Engineering Specification

│

▼

Reference Implementation

Figure D.1 — Evolution of Constitutional Admissibility. *Illustrative progression from foundational research through conceptual specification toward engineering realization and reference implementation.*

D.1 Purpose

The development of Constitutional Admissibility has progressed through successive stages, each serving a distinct purpose.

1. **White Paper 001** introduced the initial governance concepts and identified the need for continuous consideration of consequential AI execution.

2. **White Paper 002** refined these ideas, expanding the architectural discussion and introducing distinctions that would become foundational to the Constitutional Admissibility framework.

3. **CAS-01 Public Edition** organized these concepts into a coherent conceptual architecture, establishing a shared vocabulary and identifying the principal architectural components.

4. **CAS-01 Professional Release** presents the architecture in a structured technical specification suitable for broader academic, professional, and standards-oriented discussion.

Future work is expected to build upon this foundation through engineering specifications, implementation guidance, reference architectures, and operational validation.

The progression illustrated in Figure D.1 reflects the intended maturation of the Constitutional Admissibility architecture from conceptual research toward practical implementation.