SBJ

SOLUTIONS BY JEWEL

AI GOVERNANCE

ENTERPRISE ARCHITECTURE PUBLICATION

# SBJ AI Governance Architecture V1.0

A cross-cutting governance architecture for AI systems operating across generation, assistance, agency, autonomy, digital and physical environments.

**VALIDATED WITH MATERIAL LIMITATIONS**

V1.0

https://solutionsbyjewel.com/sbj-ai-governance-architecture-v1

September 2026

PUBLICATION STATUS / VERSION

# V1.0 publication companion

| ARCHITECTURE | BENCHMARK |
| --- | --- |
| FROZEN | COMPLETE |
| **VALIDATION** | **CANONICAL SOURCE** |
| VALIDATED WITH MATERIAL LIMITATIONS | solutionsbyjewel.com |

The website remains the canonical source of truth. This PDF is a supporting publication companion to the SBJ AI Governance Architecture V1.0 canonical resource.

> **Canonical resource**
>
> https://solutionsbyjewel.com/sbj-ai-governance-architecture-v1

EXECUTIVE SUMMARY

## Govern capability, state, authority, action, evidence, and accountability as one architecture.

The SBJ AI Governance Architecture V1.0 is designed to govern AI systems through capabilities, operational conditions, authority, evidence, accountability, and continuous control. Its ten Governance Dimensions provide a common structure across different AI technologies and deployment models. Cross-cutting concepts then govern system state, authorization boundaries, downstream execution, evidence continuity, intervention, and accountability.

V1.0 also provides a conceptual Operational Authorization Layer connecting identity and purpose to permitted action while preserving evidence and accountability around execution.

**Technology vocabulary follows the market.** Generative AI, AI assistants and copilots, Agentic AI, autonomous systems, and Physical AI may express multiple governance dimensions simultaneously. The architecture does not treat those technology categories as a maturity ladder.

ARCHITECTURE SCOPE

# What the architecture governs

The architecture provides a common governance structure for examining what an AI system can do, where and how it operates, what authority it has, what evidence its operation produces, who remains accountable, and how governance persists as conditions change.

### Capability
Generation, assistance, agency, autonomy, environmental interaction, and physical actuation.

### Authority
The bounded conditions under which governed systems and actors may act.

### Evidence + Accountability
Records and ownership sufficient to establish what occurred under governed conditions.

### Continuous Control
Ongoing governance as system, environment, evidence, operational, and risk conditions change.

## TEN GOVERNANCE DIMENSIONS - AT A GLANCE

01 **Generation**

02 **Assistance**

03 **Agency**

04 **Autonomy**

05 **Environmental Interaction**

06 **Physical Actuation**

07 **Authority**

08 **Evidence**

09 **Accountability**

10 **Continuous Control**

The dimensions may overlap. They are not maturity stages.

TEN GOVERNANCE DIMENSIONS

# Dimensions 01-05

### 01 Generation

The extent to which an AI system creates, transforms, predicts, recommends, synthesizes, or otherwise produces outputs from available inputs and context.

### 02 Assistance

The role an AI system performs in supporting human activity, judgment, decision-making, analysis, creation, or execution.

### 03 Agency

The capacity of an AI system to pursue a defined objective through planning, task selection, tool use, sequencing, or consequential action.

### 04 Autonomy

The degree to which an AI system can operate, decide, or act without contemporaneous human direction or approval.

### 05 Environmental Interaction

The extent to which an AI system receives information from, interprets, responds to, or changes conditions within a digital or physical operating environment.

> **Interpretation:** The Governance Dimensions characterize capabilities and governance conditions that may exist within an AI system. They may overlap. They are not maturity stages.

TEN GOVERNANCE DIMENSIONS

# Dimensions 06-10

### 06 Physical Actuation

The capacity of an AI system to produce physical action through vehicles, robotics, machinery, devices, infrastructure, or other actuated systems.

### 07 Authority

The formally permitted scope within which an AI system, agent, or governed actor may act, including applicable boundaries, approvals, credentials, targets, conditions, and limitations.

### 08 Evidence

The records, observations, artifacts, logs, approvals, decisions, and other information required to establish what occurred and under what governed conditions.

### 09 Accountability

The assignment and preservation of responsibility for authorization, oversight, operation, intervention, outcomes, and governance decisions associated with an AI system.

### 10 Continuous Control

The ongoing monitoring, reassessment, intervention, restriction, and governance necessary to keep an AI system within approved conditions throughout operation and lifecycle change.

PUBLIC ARCHITECTURE MAP

# Cross-Cutting V1.0 Architecture

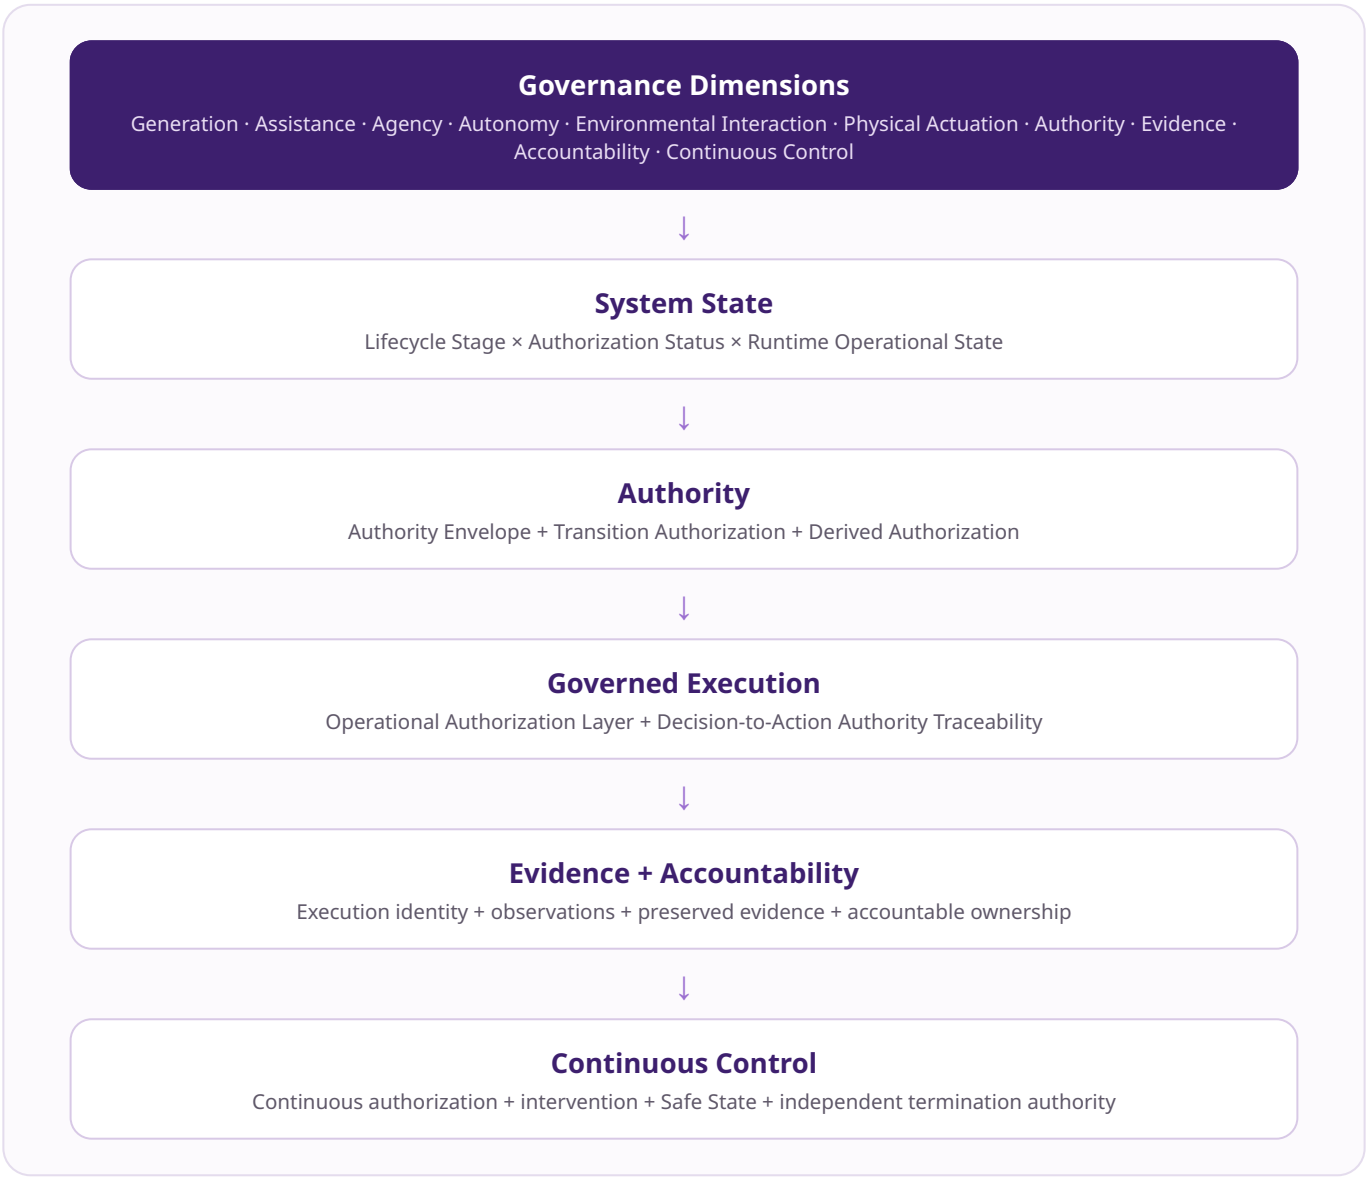


This public representation describes architectural relationships. Detailed evaluation logic, control relationships, enforcement rules, derivation mechanics, and implementation logic remain proprietary.

THREE-AXIS SYSTEM STATE

# Independent governance conditions

Lifecycle, authorization, and runtime conditions remain independently governable.

V1.0 distinguishes where a system exists in its lifecycle, whether it currently possesses authority to operate, and what it is doing operationally.

| LIFECYCLE STAGE | | AUTHORIZATION STATUS | | RUNTIME OPERATIONAL STATE |
|---|---|---|---|---|
| DEPLOYED | × | AUTHORIZED WITH CONDITIONS | × | PAUSED |

**Conceptual example**

A system can be deployed, remain authorized only under specified conditions, and simultaneously be paused at runtime. These states are represented independently rather than collapsed into a single status.

Internal state-transition enforcement rules are outside the public V1.0 disclosure.

AUTHORITY ARCHITECTURE

# Authority Envelope

**The bounded set of conditions within which an AI system or agent is authorized to act.**


The Authority Envelope provides a conceptual boundary around authorized action. Relevant conditions depend on the governed system and use case.

> Production authorization schemas, enforcement logic, condition derivation, and internal control mappings remain proprietary.

CROSS-CUTTING AUTHORIZATION CONCEPTS

# Authority, provenance, and traceability

### Transition Authorization

Governance of whether and under what approved conditions a system may cross a governed environment, lifecycle, privilege, capability, autonomy, or consequence boundary.

### Derived Authorization

Newly scoped authority issued as the result of an event, state, decision, validation, remediation, or authorized transition. Derived Authorization does not automatically transfer, revive, or expand previous authority.

### Decision-to-Action Authority Traceability

The ability to connect a governed decision to the downstream systems, tools, services, execution steps, and actions through which it is carried out.

### Independent Downstream Authority

Recognition that a downstream system or execution component may require its own valid authority even when the request originated from an authorized upstream actor.

### Execution Instance Identity

A governance identity associated with the particular execution instance involved in a governed action, supporting traceability across system, version, environment, and execution context.

CROSS-CUTTING AUTHORIZATION CONCEPTS

# Evidence, safe operation, and intervention

### Adverse Evidence Preservation

Preservation of governance-relevant evidence associated with failures, interventions, rejected actions, invalidated authority, exceptions, incidents, or other adverse conditions.

### Safe State

A defined operational condition intended to reduce, contain, or control risk when continued execution cannot proceed under authorized conditions.

### Continuous Authorization

The continuing governance of whether authority remains valid as relevant system, environment, evidence, operational, or risk conditions change.

### Independent Termination Authority

A designated authority capable of pausing, stopping, restricting, or terminating execution independently of the executing system's own decision process.

> Implementation logic for these concepts remains protected SBJ intellectual property.

OPERATIONAL AUTHORIZATION LAYER

# From bounded authority to governed action

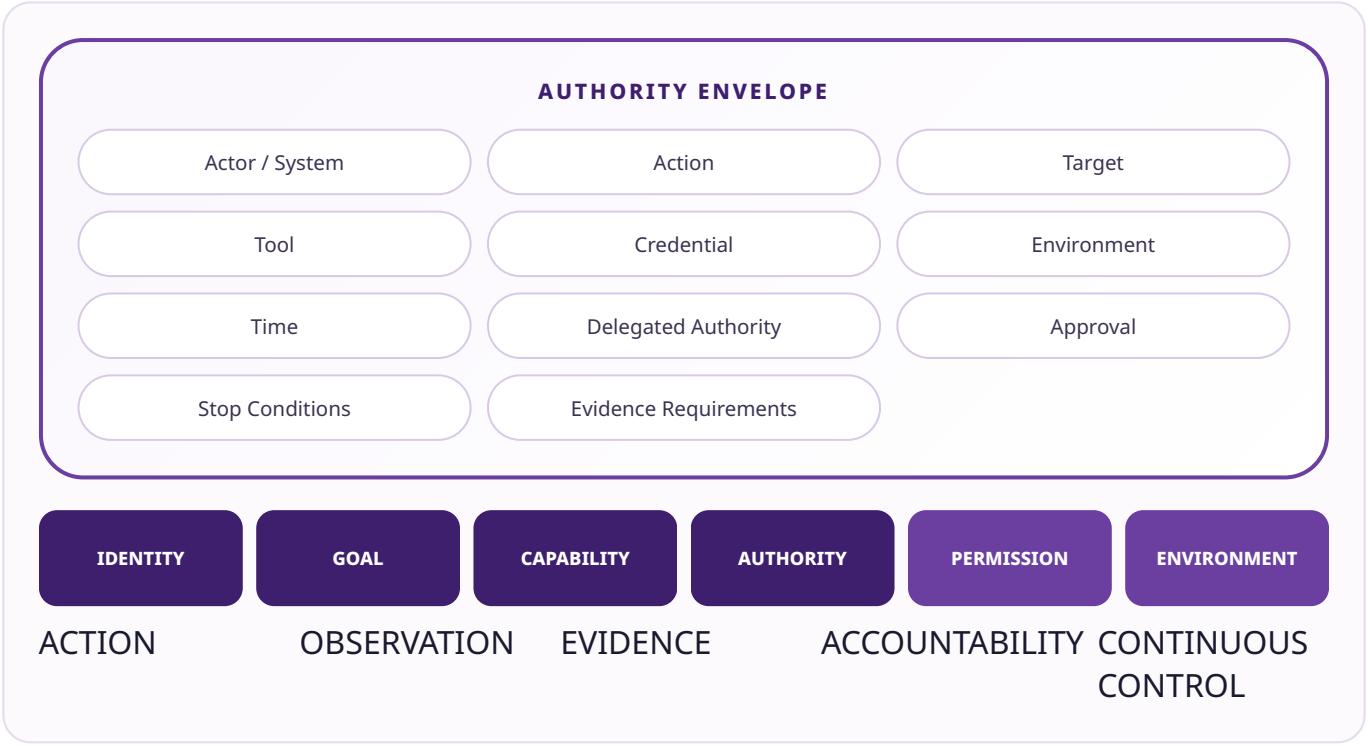


The sequence provides a governance view of how an AI system moves from identifiable purpose and capability into permitted action while maintaining evidence and accountability around execution.

IDENTITY → GOAL → CAPABILITY → AUTHORITY → PERMISSION → ENVIRONMENT → ACTION → OBSERVATION → EVIDENCE → ACCOUNTABILITY → CONTINUOUS CONTROL

> The public sequence is conceptual. Machine-readable schemas, derivation logic, policy mappings, control dependencies, authorization-condition logic, and enforcement implementation remain proprietary.

SBJ METHODOLOGY RELATIONSHIP

# The architecture sits across the SBJ methodology stack.

## STRATA™

**GRC foundation.** Provides the organizational governance, risk, and control foundation.

## SIGNAL™

**AI-specific control and audit methodology.** Supports structured AI governance assessment.

## Decision Record

**Governance-decision accountability and evidence practice.** Records consequential governance decisions, ownership, rationale, authorization, evidence, and reassessment.

## SYNTHESIZE™

**Physical AI governance methodology.** Applies governance in Physical AI environments.

## NIST MEASURE

**Separate 25-point assessment scale used by Solutions by Jewel where relevant to evaluate the NIST AI RMF MEASURE function.** It is separate from STRATA™, SIGNAL™, Decision Record, and SYNTHESIZE™ and is not a NIST-developed scoring system, NIST certification, or NIST endorsement.

> The public V1.0 artifact does not disclose scoring, weighting, certification, derivation, evidence-sufficiency, or implementation mechanics for these methodologies.

PROTOTYPE RELATIONSHIP

# Two environments. One cross-cutting governance architecture.

## Agentic AI

Conceptual governance of agent identity, declared purpose, capability, authority, permissions, environment, consequential action, downstream execution, evidence, accountability, and continuous authorization.


## Physical AI

Conceptual governance of lifecycle state, authorization state, runtime operational state, environment, physical actuation authority, Safe State, intervention, evidence, accountability, and restart / redeployment authorization.

**Architecture principle:** Agentic AI and Physical AI are prototype environments governed by the same cross-cutting architecture. They are not sequential maturity stages.

PROTOTYPE ENVIRONMENT I

# Agentic AI conceptual demonstration

Synthetic scenario: an enterprise software-change agent assigned to perform an approved configuration update.

### IDENTITY
A uniquely governed enterprise agent and execution instance are identified.

### DECLARED PURPOSE
A defined configuration-management objective.

### CAPABILITY
Inspect configuration, plan changes, use approved tools, test, and request downstream execution.

### AUTHORITY
The Authority Envelope bounds system, action, environment, timeframe, and conditions.

### PERMISSION
Technical access is limited to required resources.

### DOWNSTREAM EXECUTION
Downstream components remain subject to applicable authority and control.

### EVIDENCE
Identity, authorization, actions, observations, approvals, exceptions, and resulting state are recorded.

### ACCOUNTABILITY
A designated owner remains accountable.

### CONTINUOUS AUTHORIZATION
Material changes may require reassessment, intervention, restriction, or termination.

> This example is synthetic and non-normative. It demonstrates conceptual architectural application and does not disclose production implementation logic.

PROTOTYPE ENVIRONMENT II

# Physical AI conceptual demonstration

Synthetic scenario: an autonomous warehouse transport system operating within an approved facility zone.

**LIFECYCLE STATE**
The system is deployed.

**AUTHORIZATION STATE**
Operation is authorized within specified conditions.

**RUNTIME OPERATIONAL STATE**
The unit is actively operating.

**ENVIRONMENT**
Authorized physical boundaries and relevant environmental conditions are defined.

**PHYSICAL ACTUATION AUTHORITY**
Bounded authority to navigate and move physical payloads.

**SAFE STATE**
A governed condition for cessation or containment of movement.

**INTERVENTION**
Independent intervention or termination authority can restrict or halt operation.

**EVIDENCE + ACCOUNTABILITY**
Operational observations, state, interventions, exceptions, adverse events, and accountable owners are preserved.

**RESTART / REDEPLOYMENT**
Resumption requires valid authorization appropriate to the resulting state and conditions.

**SYNTHESIZE**™ is the SBJ Physical AI governance methodology associated with this environment.

> This demonstration is synthetic and non-normative. Gated SYNTHESIZE™ mechanics and production control implementation remain proprietary.

V1.0 REAL-WORLD BENCHMARK

# Validated with material limitations within the tested scope.

SBJ AI Governance Architecture V1.0 was independently benchmarked against six publicly documented real-world AI cases and validated with material limitations within the tested scope.

| | |
|---|---|
| Amjad Masad / Replit | Dario Amodei / Anthropic |
| Satya Nadella / Microsoft | Marc Benioff / Salesforce |
| Tekedra Mawakana / Waymo | Demis Hassabis / Google DeepMind |

## BENCHMARK RESULTS

| 185 | 160 | 24 | 0 | 1 | 7 |
|---|---|---|---|---|---|
| APPLICABLE DETERMINATIONS | PASS | PARTIAL | FAIL | INSUFFICIENT EVIDENCE | NOT APPLICABLE |

**No structural failure was demonstrated by the approved six-case cohort. This finding is limited to the tested scope and does not establish universal architectural sufficiency.**

No percentage score is assigned.

MATERIAL LIMITATIONS

# Where additional architecture development and testing are required

### Authorization invalidation / continuous-authorization trigger semantics

Further work is required around how changes in governed conditions invalidate or require reassessment of existing authorization.

### Semantic-action authority versus technically valid permission

Technical ability or valid credentials alone may not establish governance authority for the meaning, purpose, or consequence of an action.

### Branching / delegated authority provenance

Complex execution chains require stronger representation of authority provenance when authority branches, passes through intermediaries, or is delegated downstream.

MATERIAL LIMITATIONS

# Evidence, identity, and shared responsibility

## Governance-grade execution identity and version binding

Further development is required to establish sufficiently durable relationships among execution identity, system version, configuration, environment, and governed action.

## Shared-responsibility evidence allocation

Distributed AI systems may involve multiple organizations, systems, providers, operators, or control owners. Evidence responsibility across those relationships requires additional development.

## Evidence completeness, integrity, retention, and adverse-evidence persistence

Additional work is required around the completeness, integrity, preservation, retention, and continued availability of governance evidence, particularly when evidence reflects failures, interventions, exceptions, or adverse outcomes.

VALIDATION COVERAGE GAP

**Autonomous irreversible material financial transactions were not adequately represented in the V1.0 benchmark cohort. No validation claim is made for that authority class.**

PUBLIC / PROPRIETARY BOUNDARY

# Architecture is public. Protected implementation mechanics remain proprietary.

V1.0 publicly explains the approved architecture concepts, methodology relationships, sanitized prototypes, benchmark summary, findings, limitations, and selected public evidence citations.

## Public V1.0 disclosure

- Ten Governance Dimensions and high-level architecture relationships
- Three-Axis System State and Authority Envelope
- Approved cross-cutting authorization concepts
- Conceptual Operational Authorization Layer
- Sanitized Agentic AI and Physical AI demonstrations
- Benchmark cohort, results, conclusion, limitations, and coverage gap
- Selected public evidence citations

## Protected SBJ intellectual property includes

- Scoring and weighting logic; certification mechanics; derivation rules
- Evidence sufficiency logic and detailed evidence relationships
- Authorization-condition logic and enforcement mappings
- Integrated Operational Authorization Layer implementation
- Transition Authorization and Derived Authorization implementation logic
- Decision-to-Action traceability implementation mechanics
- Execution Instance Identity integration mechanics
- Production schemas and machine-executable policy logic
- Control dependency logic and internal benchmark working papers
- V1.1/V2 implementation work

SOURCES & TECHNICAL REFERENCES

# Authoritative references and public benchmark evidence

External references inform terminology, standards context, and the public evidence used for independent architecture testing. Their inclusion does not imply sponsorship, endorsement, certification, or validation of SBJ methodology.

## Standards & Governance References

**NIST - Artificial Intelligence Risk Management Framework (AI RMF 1.0)**
https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10

**NIST AIRC - AI RMF Core: GOVERN, MAP, MEASURE, MANAGE**
https://airc.nist.gov/airmf-resources/airmf/5-sec-core/

**ISO/IEC 42001:2023 - Artificial intelligence management system**
https://www.iso.org/standard/42001

**ISO/IEC 22989:2022 - Artificial intelligence concepts and terminology**
https://www.iso.org/standard/74296.html

**OECD - Explanatory memorandum on the updated definition of an AI system**
https://www.oecd.org/en/publications/explanatory-memorandum-on-the-updated-oecd-definition-of-an-ai-system_623da898-en.html

**EU Artificial Intelligence Act - Regulation (EU) 2024/1689**
https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

## Benchmark Public Evidence

**Replit - Build with Agent**
https://docs.replit.com/learn/build-with-agent

**Anthropic - Responsible Scaling Policy**
https://www.anthropic.com/responsible-scaling-policy

**Microsoft - Responsible AI Principles and Approach**
https://www.microsoft.com/en-us/ai/principles-and-approach

**Salesforce - Trusted AI**
https://www.salesforce.com/artificial-intelligence/trusted-ai

**Waymo - Safety**
https://waymo.com/safety/

**Waymo - Safety Case Framework**
https://waymo.com/blog/2023/03/a-blueprint-for-av-safety-waymos/

**Google DeepMind - Frontier Safety Framework**
https://deepmind.google/frontier-safety/

INDEPENDENT EVALUATION DISCLAIMER

# Independent architecture testing

> Solutions by Jewel independently evaluated its governance architecture against publicly documented information concerning the referenced systems and organizations. The referenced executives and organizations did not participate in, sponsor, endorse, approve, certify, or validate the SBJ methodology or benchmark. References are used solely as external evidence for independent architecture testing.

ORIGINAL METHODOLOGY NOTICE

## Original Methodology Notice

This architecture is an original methodology developed by Solutions by Jewel.

It is informed by recognized public standards, regulatory guidance, and established governance practices; however, its structure, terminology, sequencing, and presentation are proprietary to Solutions by Jewel.

It is not a reproduction, certification, or official interpretation of any external framework.

Organizations may map this architecture to applicable laws, regulations, standards, contractual requirements, and internal policies according to their own scope and obligations.

ABOUT SOLUTIONS BY JEWEL / CANONICAL RESOURCE

# SBJ AI Governance Architecture V1.0

## Solutions by Jewel

AI Governance

This PDF is the downloadable publication companion to the canonical SBJ AI Governance Architecture V1.0 website artifact.

The website remains the canonical source of truth for the public architecture.

CANONICAL RESOURCE

https://solutionsbyjewel.com/sbj-ai-governance-architecture-v1

**Publication status:** V1.0 - VALIDATED WITH MATERIAL LIMITATIONS.

This publication explains approved public architecture concepts. Protected implementation mechanics remain proprietary to Solutions by Jewel.

END OF V1.0 PUBLICATION