TERA GOVERNANCE STANDARD V1.0

# Constraint-Aware AI Risk Assessment Framework

**Full Production Standard**

A deployment control system that governs expectations before exposure. Documents constraints, tradeoffs, and non-guarantees prior to production release.

**No AI Without Governance.**

**T**rustworthiness | **E**fficiency | **R**eliability | **A**ccountability

## No AI Without Governance

Constraint-aware risk assessment for production AI systems. Document tradeoffs, establish boundaries, prevent expectation failure.

**Purpose & Scope**

This framework implements TERA Governance Standard v1.0 for AI systems deployed in production environments, high-stakes decision contexts, or regulated industries. It establishes structured, defensible methodology for risk governance aligned with constraint-aware engineering principles.

**This standard integrates:**

- Technical risk analysis and quantification
- Operational oversight controls and capacity verification
- Governance accountability and executive authorization
- Explicit constraint declaration and performance boundaries
- Tradeoff documentation and non-guarantee transparency

**This is not a documentation exercise.**

It is a deployment control system subject to independent verification.

## No AI Without Governance.

Governance requires declared constraints, accountable oversight, and documented tradeoffs.

# TERA Governance Architecture

Layered governance model radiating outward from declared constraints

SURROUNDING RING

## Independent Assurance

- **Constraint-Aware Audit Protocol™**
- **Honest AI Certification™**
- **Board Reporting**
- **Regulator Documentation**

LAYER 5

## Oversight and Drift Governance

- Quarterly Review
- Trigger-Based Reassessment
- Constraint Revalidation
- Incident Response

LAYER 4

## Pre-Deployment Governance Gate

✓ Critical risks mitigated

✓ High-risk plans approved

✓ Monitoring operational

✓ Oversight staffed

✓ Executive sign-off

**No deployment passes without this gate**

LAYER 3

## Risk Governance

**Technical**
Performance, Bias, Adversarial

**Operational**
Oversight, Monitoring, Training

**Governance**
Compliance, Documentation

**Risk Scoring:** L × I

**Mitigation Planning**

**Residual Risk Acknowledgment**

LAYER 2

## Constraint Declaration (Pre-Deployment)

**Optimization Profile**
Fast · Cheap · Good

**Constrained Dimension**

**Performance Boundaries**

**Drift Tolerance**

**Human Oversight Capacity**

**Explicit Non-Guarantees**

**Governs expectations before risk scoring**

LAYER 1 — DOCTRINE

# Foundational Doctrine

## No AI Without Governance.

**AI systems must not be deployed without:**

- Declared constraints
- Quantified performance boundaries
- Assigned human authority
- Documented tradeoffs
- Explicit non-guarantees

**The Standard is governed by four operational principles:**

**T**rustworthiness — Uncertainty is explicit and bounded

**E**fficiency — Optimization choices are declared

**R**eliability — Performance thresholds are quantified

**A**ccountability — Human authority is assigned and traceable

**All governance layers originate from declared constraints.**

No downstream control is valid without upstream declaration.

# Why Constraint Awareness Matters

### The Problem: Expectation Failure

Artificial intelligence systems do not fail because they are intelligent. They fail because expectations exceed declared constraints.

## AI Systems Exist Inside Immutable Constraints

| | |
|---|---|
| **Speed** | **Cost** |
| **Quality** | **Human Attention** |
| **Regulatory Exposure** | |

No system optimizes all dimensions simultaneously.

## When Constraints Are Not Declared, Organizations Suffer:

- Operational failure
- Regulatory exposure
- Reputational harm
- Internal blame cycles

**Governance must begin with constraint transparency.**

## Constraint Declaration Forces Critical Questions:

**What is optimized?**

**What is intentionally limited?**

**What cannot be guaranteed?**

**What performance degradation is acceptable?**

**What human capacity exists to review outputs?**

**Without these answers, deployment becomes speculative.**

# Honest AI Certification™ Seal Criteria

## Evidence-Based Governance Verification

Independent assessment of constraint transparency and governance discipline

**Certification verifies constraint transparency and governance discipline through independent review, not self-attestation.**

### Eligibility Requirements

To display the Honest AI Certification™ seal, an organization must demonstrate:

- [ ] Completed Constraint Declaration Module
- [ ] Documented optimization profile
- [ ] Explicit non-guarantees
- [ ] Quantified performance thresholds
- [ ] Defined human oversight capacity
- [ ] Executive residual risk acknowledgment
- [ ] Operational monitoring systems
- [ ] Drift reassessment cadence
- [ ] Passed Constraint-Aware Audit Protocol™

**LEVEL I**

**Governance Documented —**

**LEVEL II**

**Governance Operationalized**

### Foundation Certification

Constraint declarations completed; risk framework established.

Operational proof: monitoring & drift thresholds enforced; oversight activation evidenced; deployment gates traceable; incident protocol exercised.

LEVEL III

### Governance Independently Audited

Independent stress audit: four-phase protocol, infrastructure-vs-claims consistency checks, oversight stress tests, and executive residual-risk confirmation. Appropriate for high-risk deployments.

**Seal Use Conditions**

- Certification valid for 12 months
- Immediate revocation if governance lapse identified
- Material system changes require reassessment
- Seal must include certification level (I, II, III)

**Seal Language Example**

*"This AI system has been independently assessed under TERA Governance Standard v1.0 for constraint transparency and accountable oversight at the time of review. It does not certify safety, compliance, or performance."*

# Regulatory Alignment Brief

TERA Governance Standard v1.0 and EU AI Act Operational Expectations

**Purpose**

This brief explains how TERA Governance Standard supports structured AI risk management and operational oversight expectations under emerging regulatory frameworks.

## Governance Philosophy

TERA Governance Standard requires:

- Pre-deployment constraint declaration
- Risk quantification and prioritization
- Documented mitigation plans

- Human oversight capacity documentation
- Executive accountability
- Ongoing drift monitoring

**It operationalizes risk governance before deployment exposure.**

## EU AI Act Alignment Areas

**Article 9** **Risk Management System**

Structured risk identification and lifecycle reassessment with quantified likelihood and impact scoring.

**Article 14** **Human Oversight**

Quantified review capacity, defined escalation SLAs, and documented override mechanisms.

**Article 11** **Technical Documentation**

Formalized constraint documentation including performance boundaries and explicit non-guarantees.

**Article 15** **Accuracy, Robustness, Cybersecurity**

Declared performance floors, drift tolerance thresholds, and adversarial risk documentation.

**Article 17** **Quality Management System**

Deployment gate enforcement, review cadence, and executive authorization requirements.

**Standard Limitation Notice**

TERA Governance Standard v1.0 does not constitute regulatory approval, conformity assessment, or legal certification. Deployment responsibility remains with the deploying entity.

**The Standard:**

- Documents governance discipline and constraint transparency
- Does not grant regulatory approval
- Does not substitute conformity assessment
- Does not provide legal certification

**Organizations implementing this Standard maintain full accountability for deployment decisions and operational outcomes.**

## Regulatory Value

The standard:

- Reduces expectation failure risk
- Strengthens audit readiness

- Clarifies executive accountability
- Makes performance boundaries explicit
- Documents residual risk acceptance

**It improves defensibility without overstating compliance.**

PART I

# Full Production Standard

1

## System Overview

### AI System Name

Enter system name

### System Description

Provide comprehensive system description

### Intended Use Cases

List primary use cases

### Deployment Context

Describe deployment environment and stakeholder context

**Risk Assessment Date**

mm/dd/yyyy

**Assessment Team (with roles)**

Name | Role | Department
Example: Jane Smith | ML Engineer | AI Systems
Example: John Doe | Risk Officer | Governance

2

# Constraint Declaration Module

**MANDATORY BEFORE RISK SCORING**

This section must be completed before risk scoring begins. Constraint awareness is foundational. Risk analysis without declared system boundaries is incomplete.

**Certification Requirement:**

This module validates constraint transparency and tradeoff documentation. Incomplete declarations render certification ineligible.

## 2.1 Optimization Profile (Fast · Cheap · Good Declaration)

### Primary Optimization Dimensions

Select exactly TWO dimensions to optimize

F

**Fast**

Low latency, high throughput

C

**Cheap**

Cost-efficient compute and infrastructure

G

**Good**

High accuracy, robustness, reliability

**Constrained Dimension (must be explicitly declared):**

Which dimension is constrained?

### Justification for Optimization Choice

Explain why these two dimensions were prioritized and what constraints drove this tradeoff decision

## 2.2 Performance Boundaries

> **Critical Note:**
>
> These values must reflect operational reality, not aspirational targets.

### Minimum Acceptable Accuracy Threshold

e.g., 92.5% on production distribution

### Maximum Acceptable Error Rate

e.g., 7.5% false positive rate

### Maximum Latency (P95 or P99)

e.g., P95 < 250ms

### Maximum Throughput Limitations

e.g., 10,000 requests/hour

### Acceptable Drift Tolerance (quantified)

e.g., <3% accuracy degradation per quarter

## 2.3 Human Oversight Capacity

> **Human capacity is a system constraint.**

If review bandwidth is insufficient, deployment risk increases.

**Maximum cases per reviewer per day**

e.g., 150 cases/day

**Monitoring frequency**

e.g., Real-time alerts + daily dashboard review

**Escalation SLA**

e.g., Critical issues escalated within 15 minutes

**Override authority defined?**

- ◯ Yes - Override authority documented and assigned
- ◯ No - Override authority not yet established

## 2.4 Explicit Non-Guarantees

**If this section is incomplete, the assessment is invalid.**

**This system does NOT guarantee:**

List explicit non-guarantees
Example: Perfect accuracy in edge cases
Example: Sub-100ms latency under peak load
Example: Zero false positives

**This system should NOT be used for:**

List prohibited use cases
Example: Sole basis for medical diagnosis
Example: Automated decision-making without human review
Example: Processing of protected health information without explicit consent

**Known failure classes:**

Document known failure modes
Example: Performance degrades on non-English text
Example: Accuracy drops on images captured in low-light conditions
Example: High false positive rate on edge cases representing <1% of training data

3

# Risk Categorization Matrix

**For each risk, assess:**

- **Likelihood** (1–5)
- **Impact** (1–5)
- **Risk Score** (L × I)
- **Priority Level**

## 3.1 Technical Risks

Include but not limited to:

- Model performance degradation
- Bias and fairness exposure
- Adversarial vulnerability

- Data quality constraints
- Explainability gaps
- Integration failures

> Each risk must reference declared performance boundaries and constraint decisions.

| RISK DESCRIPTION | LIKELIHOOD (1-5) | IMPACT (1-5) | SCORE (L×I) | PRIORITY |
|---|---|---|---|---|
| Describe tech | 1-5 | 1-5 | Auto-ca | Priority |
| Describe tech | 1-5 | 1-5 | Auto-ca | Priority |

*Add additional rows as needed*

## 3.2 Operational Risks

- Oversight failures
- Incident response gaps
- User training deficiencies
- Monitoring blind spots
- Resource constraints

> Operational risks must account for declared human capacity and infrastructure limits.

| RISK DESCRIPTION | LIKELIHOOD (1-5) | IMPACT (1-5) | SCORE (L×I) | PRIORITY |
|---|---|---|---|---|
| Describe ope | 1-5 | 1-5 | Auto-ca | Priority |
| Describe ope | 1-5 | 1-5 | Auto-ca | Priority |

### 3.3 Governance and Compliance Risks

- Regulatory non-compliance
- Documentation gaps
- Accountability ambiguity
- Privacy violations
- Third-party dependencies

> Governance risks must align with declared accountability structures and regulatory context.

| RISK DESCRIPTION | LIKELIHOOD (1-5) | IMPACT (1-5) | SCORE (L×I) | PRIORITY |
|---|---|---|---|---|
| Describe gov | 1-5 | 1-5 | Auto-ca | Priority |
| Describe gov | 1-5 | 1-5 | Auto-ca | Priority |

4

# Risk Scoring & Prioritization

**Risk Score = Likelihood × Impact**

**CRITICAL (20-25)**

**Action:** Deployment prohibited until mitigated.

**HIGH (15-19)**

**Action:** Executive approval required with mitigation plan.

**MEDIUM (8-14)**

**Action:** Monitoring + mitigation within defined timeline.

**LOW (1-7)**

**Action:** Document and track.

5

# Mitigation Planning

**Each HIGH and CRITICAL risk must include:**

- Specific mitigation actions

- Risk reduction mechanism
- Owner
- Target completion date
- Status

Mitigation must reference constraint adjustments where applicable.

| RISK ID | MITIGATION ACTION | OWNER | TARGET DATE | STATUS |
|---|---|---|---|---|
| Risk | Describe specific | Owner | mm/dd/yyyy | Sel |
| Risk | Describe specific | Owner | mm/dd/yyyy | Sel |

*Add additional mitigation plans as needed*

**Residual Risk Acknowledgment**

Residual risk exposure has been reviewed and acknowledged by executive authority.

6

# Pre-Deployment Governance Gate

**This gate is mandatory.**

Deployment may not proceed unless all criteria are met.

**TERA Standard v1.0 Requirement:**

No deployment without governance authorization.

**Certification eligibility requires all gate criteria to be satisfied.**

Gate enforcement is subject to independent verification.

- [ ] All constraint declarations are completed
- [ ] CRITICAL risks are mitigated
- [ ] HIGH risks have approved mitigation plans
- [ ] Monitoring systems are operational
- [ ] Human oversight is defined and staffed
- [ ] Documentation is audit-ready
- [ ] Executive sign-off obtained

**Enforcement Notice:**

Deployment without completion of this gate constitutes a governance breach under TERA Standard v1.0.

### Executive Sign-Off

Executive Name | Title

mm/dd/yyyy | Signature

7

## Review and Drift Governance

**TERA Standard v1.0 Section 5:** Ongoing Oversight and Drift Governance

**Quarterly review minimum.**

Monitoring systems and review cadence are subject to operational verification.

### Immediate Reassessment Required If:

- [ ] Performance drops below declared thresholds

- [ ] Regulatory changes occur
- [ ] Deployment context shifts
- [ ] Safety incident occurs
- [ ] Major retraining or architecture update occurs

**Constraint assumptions must be revalidated at each review.**

### Next Scheduled Review Date

mm/dd/yyyy

### Review Cadence

Select Cadence

PART II

# Board-Facing Executive Risk Summary

**Purpose:**

Concise executive overview for board-level review and decision-making.

This summary implements TERA Governance Standard v1.0 and supports Honest AI Certification™ eligibility through constraint transparency and accountability verification.

## Executive AI Risk Overview

### System Name

### Deployment Context

### Optimization Profile

This system is optimized for:

- [ ] Speed
- [ ] Cost

☐ Quality

### Constrained Dimension

**Board Note:**

No system optimizes all three.

### Declared Boundaries

#### Minimum performance threshold

#### Maximum acceptable error rate

#### Human review capacity

#### Known limitations

### Top Risk Summary

#### CRITICAL Risks

Count

#### HIGH Risks

Count

**Mitigation Status:**

- [ ] Critical mitigated? **Yes/No**
- [ ] High mitigation plans approved? **Yes/No**

### Governance Assurance

- [ ] Human oversight active
- [ ] Monitoring operational
- [ ] Incident response defined
- [ ] Executive accountability assigned

### Residual Risk Determination

After all mitigations are applied, residual risk exposure remains. This section documents the organization's determination on residual risk acceptance.

**Residual risk must be explicitly accepted by executive authority prior to deployment authorization.**

### Residual Risk Description

Describe remaining risk exposure after all mitigations are implemented (e.g., model may still misclassify edge cases at 2-3% rate despite validation controls)

### Risk Tolerance Assessment

Select tolerance position

### Residual Risk Acknowledgment Statement

Executive statement acknowledging residual risk position (e.g., Board acknowledges 2-3% residual classification error rate and accepts this risk given mitigation controls and monitoring systems)

**Board-level question:**

Are we comfortable deploying within these declared constraints?

PART III

# Regulator-Ready Formal Version

## AI Risk and Constraint Assessment Statement

TERA Governance Standard v1.0 | Honest AI Certification™ Aligned

This document constitutes a structured assessment of AI-related technical, operational, and governance risks in accordance with TERA Governance Standard v1.0, internal oversight requirements, and relevant regulatory expectations including EU AI Act operational provisions (Articles 9, 10, 11, 14, 15, 17).

**The assessment includes:**

- Explicit system constraint declaration
- Quantified performance boundaries
- Human oversight capacity documentation
- Risk categorization and prioritization
- Mitigation planning with assigned accountability
- Pre-deployment approval gate
- Ongoing review cadence

**The organization acknowledges:**

- No AI system operates without constraints
- Optimization decisions inherently bound performance
- Responsibility for deployment remains with the deploying entity

**Governance Ecosystem Position:**

This framework implements TERA Governance Standard v1.0 and is subject to Constraint-Aware Audit Protocol™ verification. Systems completing this assessment may undergo independent verification under the Constraint-Aware Audit Protocol™.

**Authorized Representative Name**

**Title**

**Organization**

**Date**

mm/dd/yyyy

**Signature**

Digital or physical signature

# Governance Ecosystem Architecture

TERA Governance Standard v1.0 operates within an integrated assurance structure:

*Standard → Audit → Certification → Regulatory Mapping*

### TERA Governance Standard v1.0

Defines constraint-aware deployment control requirements.

### Constraint-Aware Audit Protocol™

Verifies operational enforcement of declared constraints.

### Honest AI Certification™

Public-facing verification of governance discipline.

### Regulatory Mapping Layer

Structural alignment to EU AI Act Articles 9, 11, 14, 15, 17.

This ecosystem supports deployment control, institutional accountability, and audit readiness.

- Risk Lifecycle Governance
- Constraint Declaration
- Tradeoff Documentation
- Non-Guarantee Transparency
- Human Accountability
- Mitigation Tracking
- Drift Governance
- Audit Readiness

## Version Control

**VERSION**

**1.0**

**EFFECTIVE DATE**

February 13, 2026

**NEXT REVIEW**

February 2027

STANDARD OWNER

TeraSystemsAI Governance Office

### Revision Policy

Future revisions will document changes to constraint declaration requirements, oversight mechanisms, and regulatory alignment provisions. All modifications to TERA Governance Standard v1.0 will maintain backward compatibility with existing implementations where operationally feasible.

**Standard Status:** Active | **Classification:** Governance Infrastructure | **Applicability:** All Production AI Systems

# Applicability Clause

**This Standard applies to:**

**All AI systems deployed in production environments**

**Systems influencing regulated or high-stakes decisions**

**Systems producing external-facing outputs**

**Systems materially affecting records, rights, health, safety, or financial outcomes**

**Internal experimental systems must complete a scaled governance assessment prior to production transition.**

## TERASYSTEMSAI

Constraint-aware AI engineering, governance, and oversight grounded in TERA principles. We make tradeoffs explicit, prevent expectation failure, and help organizations build AI systems that can be trusted in production and regulated environments.

Philadelphia, Pennsylvania

### SERVICES

Overview

Pre-Deployment Risk Audits

Independent Oversight Retainer

Deployment-Ready Solutions

## RESOURCES

Publications

Research

Accountability Framework

Insights

Careers

## CONTACT

contact@terasystems.ai
research@terasystems.ai

**Discuss Constraint-Aware Governance**

---

**TERA Operating Principle:** TeraSystemsAI provides constraint-aware governance frameworks, tradeoff documentation tools, and expectation alignment services. We do not certify AI systems, eliminate risk, or guarantee compliance. These resources help organizations make tradeoffs explicit, identify expectation failure modes, and establish bounded guarantees. Deployment control, regulatory compliance, and human accountability remain with the deploying organization. Alignment is not certification. Governance is not guarantee. Constraint awareness is practice, not promise.

---



Privacy Terms