INSTITUTIONAL USE ONLY

Institutional Implementation Guide

Full Matrices, Schemas, Audit Mappings, Export Packages

VERSION	EFFECTIVE DATE	CLASSIFICATION	ISSUING ENTITY
1.0	December 29, 2025	Institutional / Deployment Teams	TeraSystemsAI

NON-NEGOTIABLE CONSTRAINT

This guide may not be interpreted or implemented in any way that: grants authority to AI systems, allows automated final decisions, obscures accountability ownership, or replaces human judgment. Any implementation that violates this constraint is non-compliant.

CONTENTS

1. Purpose of This Guide
2. Responsibility Matrix Specification
3. Role Enforcement Schemas
4. Escalation Logic Mapping
5. Audit & Compliance Mapping
6. Language Protocol Enforcement
7. Export & Regulator Packages
8. Deployment Checklist
9. Change Control & Versioning

1. Purpose of This Guide

What This Guide Does

- Translates the Accountability Invariant into deployable governance artifacts
- Provides schemas, mappings, and specifications for institutional implementation
- Supports legal defensibility, audit readiness, and regulatory compliance
- Assumes institutions have accepted the foundational governance principle

What This Guide Does Not Do

EXPLICIT LIMITATIONS

This guide does not certify ethical behavior. It does not guarantee correctness of AI outputs. It does not eliminate institutional liability. It does not replace legal counsel review. Institutions remain fully responsible for deployment decisions.

2. Responsibility Matrix - Full Specification

The Responsibility Matrix is a mandatory system artifact. Every AI-assisted outcome must produce exactly one matrix record. No matrix record means no valid output.

Required Fields

FIELD	TYPE	REQUIRED	DESCRIPTION
`matrix_id`	UUID	Yes	Unique identifier for this responsibility record
`outcome_id`	UUID	Yes	Reference to the AI-assisted outcome
`ai_system_id`	String	Yes	Identifier and version of AI system
`ai_recommendation`	Object	Yes	AI output with confidence and uncertainty
`human_reviewer_id`	String	Yes	Authenticated identity of human reviewer
`human_decision`	Enum	Yes	APPROVED | REJECTED | MODIFIED
`human_rationale`	String	Conditional	Required if REJECTED or MODIFIED
`policy_version`	String	Yes	Active policy version at decision time
`escalation_triggered`	Boolean	Yes	Whether escalation occurred
`escalation_reason`	String	Conditional	Required if escalation_triggered = true
`institution_id`	String	Yes	Accountable institution identifier
`created_at`	ISO 8601	Yes	Immutable creation timestamp
`checksum`	SHA-256	Yes	Integrity verification hash

Conceptual Schema

RESPONSIBILITYMATRIX SCHEMA

```
ResponsibilityMatrix {
  matrix_id:           UUID [PRIMARY KEY, IMMUTABLE]
  outcome_id:          UUID [FOREIGN KEY, NOT NULL]

  // AI Contribution
  ai_system_id:        STRING [NOT NULL]
  ai_recommendation:   JSON {
    output:            ANY
    confidence:        FLOAT [0.0-1.0]
    uncertainty:       FLOAT [0.0-1.0]
    model_version:     STRING
  }

  // Human Review
  human_reviewer_id:   STRING [NOT NULL, AUTHENTICATED]
  human_decision:      ENUM [APPROVED, REJECTED, MODIFIED]
  human_rationale:     STRING [REQUIRED IF decision != APPROVED]
  review_timestamp:    TIMESTAMP [NOT NULL]

  // Policy Constraints
  policy_version:      STRING [NOT NULL]
  policy_rules_applied: ARRAY[STRING]

  // Escalation State
  escalation_triggered: BOOLEAN [NOT NULL]
  escalation_reason:    STRING [REQUIRED IF triggered]
  escalation_timestamp: TIMESTAMP [NULLABLE]

  // Institutional Ownership
  institution_id:      STRING [NOT NULL]
  department_id:       STRING [NULLABLE]

  // Audit Fields
  created_at:          TIMESTAMP [IMMUTABLE]
  checksum:            SHA256 [COMPUTED]
}
```

IMMUTABILITY REQUIREMENT

Matrix records must be append-only. No UPDATE or DELETE operations are permitted after creation. Corrections require new matrix records with explicit references to superseded records.

Storage & Retention

- **Storage:** Append-only data store with cryptographic integrity verification
- **Retention:** Minimum 7 years or as required by applicable regulations
- **Access:** Read-only for audit; no modification access granted to any role
- **Backup:** Geographically distributed with integrity verification on restore

3. Role Enforcement Schemas

Each role has explicit permissions, prohibitions, and attestation requirements. These are governance constraints, not application features.

AI System Role

PERMITTED ACTIONS

- Generate recommendations
- Compute confidence scores
- Quantify uncertainty
- Flag escalation conditions
- Log outputs to matrix

PROHIBITED ACTIONS

- Mark output as "final"
- Execute decisions
- Bypass human review
- Modify policy rules
- Accept responsibility

REQUIRED ATTESTATIONS

- Output is recommendation only
- Confidence bounds are valid
- Escalation rules evaluated
- Policy version recorded

Human Reviewer Role

PERMITTED ACTIONS

- Review AI recommendations
- Approve, reject, or modify
- Override AI output
- Request additional review
- Document rationale

PROHIBITED ACTIONS

- Delegate to AI system
- Approve without review
- Bypass escalation
- Modify audit records
- Transfer accountability

REQUIRED ATTESTATIONS

- Review was performed
- AI output was evaluated
- Decision is independent
- Accountability accepted

Policy Configuration Role

PERMITTED ACTIONS

- Define thresholds
- Set escalation triggers
- Configure constraints
- Version policy changes
- Audit policy history

PROHIBITED ACTIONS

- Disable escalation
- Remove human review
- Grant AI authority
- Delete policy history
- Bypass approval flow

REQUIRED ATTESTATIONS

- Changes are authorized
- Version is incremented
- Invariant preserved
- Audit trail updated

Institution Role

PERMITTED ACTIONS

- Accept liability
- Enforce governance
- Authorize deployments
- Review audit records
- Respond to regulators

PROHIBITED ACTIONS

- Transfer accountability to AI
- Claim AI decided
- Obscure responsibility
- Disable audit logging
- Modify historical records

REQUIRED ATTESTATIONS

- Framework compliance
- Audit readiness
- Liability acknowledged
- Governance enforced

Enforcement Rules

ROLE ENFORCEMENT LOGIC

```
RULE: AI_OUTPUT_NEVER_FINAL
  IF output.status == "FINAL" AND output.source == "AI"
  THEN REJECT with "AI outputs cannot be marked final"

RULE: HUMAN_DECISION_REQUIRED
  IF matrix.human_decision IS NULL
  THEN BLOCK output propagation

RULE: POLICY_VERSION_BOUND
  IF matrix.policy_version != active_policy.version
  THEN REJECT with "Policy version mismatch"

RULE: ESCALATION_NON_BYPASSABLE
  IF escalation_condition_met == TRUE
  THEN REQUIRE human_review
  AND BLOCK auto_approval
```

4. Escalation Logic - Operational Mapping

Escalation Triggers

TRIGGER TYPE	CONDITION	THRESHOLD EXAMPLE	ACTION
Confidence	`confidence < threshold`	`confidence < 0.85`	Mandatory review
Uncertainty	`uncertainty > threshold`	`uncertainty > 0.20`	Mandatory review
Bias Flag	`bias_score > threshold`	`bias_score > 0.10`	Mandatory review
Domain Risk	`risk_category IN high_risk`	Life-safety, legal, financial	Mandatory review
Anomaly	`input NOT IN distribution`	Out-of-distribution detected	Mandatory review
Manual	`human_request == TRUE`	Any user request	Review initiated

Escalation Sequence

ESCALATION STATE MACHINE

```
STATE: AI_PROCESSING
  ON trigger_detected:
    FREEZE ai_output
    SET status = ESCALATED
    TRANSITION TO AWAITING_REVIEW

STATE: AWAITING_REVIEW
  REQUIRE human_reviewer_assignment
  PROVIDE full_context {
    ai_recommendation
    confidence_scores
    uncertainty_bounds
    escalation_reason
    policy_rules
  }
  ON human_decision:
    LOG to responsibility_matrix
    TRANSITION TO DECISION_LOGGED

STATE: DECISION_LOGGED
  SET matrix.human_decision
  SET matrix.human_reviewer_id
  SET matrix.review_timestamp
  COMPUTE matrix.checksum
  TRANSITION TO COMPLETE

STATE: COMPLETE
  OUTPUT is now valid
  RESPONSIBILITY assigned to human_reviewer
  RECORD immutable
```

NON-BYPASSABLE CONSTRAINT

Escalation cannot be disabled, bypassed, or overridden by configuration. No administrative privilege permits skipping escalation when trigger conditions are met. This is enforced at the system architecture level.

5. Audit & Compliance Mapping

Framework artifacts map to regulatory requirements. This section provides traceability, not legal interpretation.

REGULATION	REQUIREMENT	FRAMEWORK ARTIFACT	EVIDENCE
FDA CDS	Human oversight of clinical decisions	Responsibility Matrix	human_reviewer_id, human_decision fields
GDPR Art. 22	Right to human review of automated decisions	Escalation Logic	Mandatory review for high-risk; escalation logs
EU AI Act	High-risk system oversight	Role Enforcement	Human role cannot be bypassed; audit trail
EU AI Act	Transparency and traceability	Responsibility Matrix	Complete decision chain with timestamps
SOC 2	Processing integrity	Immutable Logging	Append-only records with checksums
SOC 2	Change management	Policy Versioning	Version history with approval records
HIPAA	Audit controls	Audit Export	Exportable logs with access controls
OCC/Fed	Model risk management	Escalation + Matrix	Uncertainty quantification; human override

Audit Evidence Availability

- **Decision Chain:** Complete path from AI output to human approval
- **Temporal Accuracy:** Immutable timestamps at each stage
- **Policy State:** Exact policy version active at decision time
- **Override Documentation:** Rationale for rejections or modifications
- **Escalation Records:** Why escalation occurred; how it resolved

6. Language Protocol - Enforcement Layer

Language constraints are deployable governance rules. Mislabeling creates legal exposure.

Prohibited Phrases

PROHIBITED	RISK	REQUIRED ALTERNATIVE
`"The AI decided"`	Implies AI authority	`"The AI recommended"`
`"System approved"`	Implies automated approval	`"Reviewer approved"`
`"Algorithm determined"`	Implies AI judgment	`"Analysis indicated"`
`"Automated decision"`	Implies no human	`"AI-assisted recommendation"`
`"AI concluded"`	Implies AI reasoning	`"AI identified"`

Enforcement Points

- **User Interface:** Validate labels before display
- **Audit Logs:** Scan entries for prohibited phrases
- **Reports:** Template enforcement for generated documents
- **Exports:** Compliance check before external delivery
- **API Responses:** Response validation middleware

LANGUAGE VALIDATION RULE

```
FUNCTION validate_language(text):
  prohibited = [
    "AI decided", "system approved", "algorithm determined",
    "automated decision", "AI concluded", "machine judgment"
  ]
  FOR phrase IN prohibited:
    IF phrase IN text.lower():
      RETURN {valid: FALSE, violation: phrase}
  RETURN {valid: TRUE}
```

7. Export & Regulator Packages

Standard Export Bundles

PACKAGE	CONTENTS	USE CASE
Decision Audit	Responsibility matrices for date range; escalation logs; policy versions	Regulatory inquiry; internal audit
Escalation Report	All escalations with triggers, resolutions, timing	Compliance review; process audit
Policy History	Complete version history with change rationale	Change control audit; timeline reconstruction
Role Activity	Actions by role type; reviewer activity; override frequency	Operational review; capacity planning
Full Compliance	All above packages combined with integrity proofs	Regulatory examination; legal discovery

Export Metadata

All exports include:

- Export timestamp and requesting user
- Date range and filter criteria applied
- Record count and integrity checksums
- Framework version and policy version at export time
- Chain of custody fields for legal use

> **EXPORT GUARANTEE**
> Exports contain sufficient information for external review without requiring system access. A regulator can verify accountability chains using export data alone.

8. Deployment Checklist (Institutional)

All items must be satisfied before production deployment. This is risk gating, not feature enablement.

- □ **Roles Declared:** All four roles (AI, Human, Policy, Institution) formally assigned with named owners
- □ **Policies Versioned:** Initial policy version created, approved, and logged with version 1.0
- □ **Escalation Tested:** All escalation triggers verified with test cases; bypass attempts confirmed blocked
- □ **Audit Logging Verified:** Responsibility matrix creation confirmed; immutability tested; export validated
- □ **Language Protocol Enforced:** Prohibited phrases blocked in UI, logs, reports; validation middleware active
- □ **Human Review Path Tested:** End-to-end flow from AI output to human decision verified functional
- □ **Export Packages Generated:** Test exports created and validated for completeness
- □ **Retention Configured:** Data retention period set per regulatory requirements; backup verified
- □ **Legal Review Complete:** Institutional counsel has reviewed deployment configuration
- □ **Accountability Sign-off:** Institutional owner has formally accepted liability for deployment

GATE REQUIREMENT

Production deployment is blocked until all checklist items are completed and documented. Incomplete deployments are non-compliant.

9. Change Control & Versioning

Change Categories

CATEGORY	EXAMPLES	REVIEW REQUIRED	APPROVAL LEVEL
Frozen	Core invariant; role definitions; escalation non-bypass	Not changeable	N/A - Immutable
Major	New escalation triggers; role permission changes	Full review	Institutional + Legal
Minor	Threshold adjustments; domain-specific rules	Standard review	Policy owner
Administrative	User assignments; documentation updates	Logged only	Authorized admin

Version Control Requirements

- All policy changes increment version number
- Change rationale is mandatory for Major and Minor changes
- Previous versions are retained indefinitely
- Active policy version is bound to all new matrix records
- Rollback requires same approval as original change

GOVERNANCE DRIFT WARNING

Governance drift (gradual weakening of controls through incremental changes) increases institutional risk. All changes must be reviewed for cumulative effect on accountability guarantees.

This document is for institutional use only. It does not constitute legal advice. Institutions remain fully responsible for deployment decisions and regulatory compliance.