

AI Production Readiness Evidence Pack

A structured evidence inventory for system boundaries, representative tests, human authority, failure taxonomy, release criteria, and unresolved risk.

Buyer problem

Teams need a practical inventory of evidence before approving an AI production release or enterprise review.

System and authority

- Name the business problem and decision owner
- Define users, affected parties, and consequences
- Inventory models, prompts, retrieval, vendors, tools, and downstream actions
- Map data sources, trust states, access boundaries, and retention expectations
- Name human reviewers, approvers, execution roles, and escalation owners

Representative evaluation

- Golden cases for expected work
- Edge and ambiguous cases
- Adversarial and prompt-injection cases
- Refusal and insufficient-evidence cases
- Access-control and tenant-boundary cases
- Cost, latency, timeout, and dependency cases

Release evidence

- Versioned test fixtures and system configuration
- Observed results and failure taxonomy
- Human calibration and disagreement records
- Release, exception, rollback, and blocked-action criteria
- Known limitations, unknowns, owners, and review dates

Related engagement

Service	AI Production Readiness & Evidence Sprint
Duration	3 weeks
Starting investment	\$20,000
Payment	50% at start; 50% at delivery
Decision supported	Proceed, narrow, remediate, or stop - with a measured baseline and explicit release conditions.

Usage boundary: A planning and review aid, not a formal assessment, audit, certification, or legal opinion.

Last reviewed 2026-07-23 | HTML version: <https://longtermcapabilities.com/resources/ai-production-readiness-evidence-pack/> | Contact: 1 (464) 274-1476 / consulting@longtermcapabilities.com