

The AI Pilot-to-Production Framework

Public edition • A stage-gated, evidence-driven method for productionizing enterprise AI and agentic applications - also known as the Netwoven AI Productionization Bridge™.

Netwoven helps enterprises move copilots, RAG applications and agentic workflows from pilot to secure, integrated, measurable and supportable production - and then operates them to sustain business value.

Why a framework is necessary

Enterprises are not short of AI pilots; they are short of a repeatable path from promising demonstration to governed business system. MIT's research found 95% of GenAI pilots deliver no measurable P&L impact. S&P Global reports the share of companies abandoning most of their AI initiatives jumped to 42% in a single year, with 46% of proofs-of-concept scrapped before production. Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027.

The cause is rarely the model. Production requires a funded outcome, a redesigned workflow, permission-aware data, safe integration, measurable quality, security controls, adoption - and an operating model after go-live.

What production-ready means

Dimension	Production-ready condition
Business	An accountable owner, measurable target, adoption plan and approved investment case
Functional	Completes the intended task across representative scenarios; handles exceptions predictably
Data & permissions	Sources are current, governed, and accessed according to user and agent permissions
Architecture	No more complex than necessary; meets enterprise non-functional requirements
Trust	Quality, grounding, safety and human oversight measured against accepted thresholds
Security & compliance	Identity, secrets, network, logging, data protection and regulatory obligations addressed
Operations	Telemetry, runbooks, change control, support ownership and cost governance in place

The eight production gaps

Every candidate - new idea, active pilot, or stalled initiative - is assessed against eight gaps, each with evidence-based exit criteria.

#	Gap	The question it answers
1	Strategy-to-use-case	Is the workflow frequent, measurable and valuable enough to fund production?
2	Pilot-to-workflow	Is AI embedded where work actually happens, with roles and exceptions designed?
3	Content & data	Are sources current, governed and permission-ready for production-scale use?
4	Architecture	Is this the simplest pattern that meets the need - scalable, resilient, testable?
5	Integration	Can the solution safely read from and write to systems of record and approvals?
6	Trust & evaluation	Can quality be proven with datasets, thresholds and regression - not anecdotes?
7	Governance	Are ownership, risk controls, auditability and human oversight explicit?
8	Run-model	Who operates, evaluates, refreshes and improves the system after launch?

The lifecycle: evidence gates, not enthusiasm

Seven stages, each ending in a gate where required evidence is reviewed by accountable owners before investment or release proceeds.

Stage	Focus
1. Assess & prioritize	Score the eight gaps; produce readiness scorecard, value case and recommended pattern
2. Architect & plan	Target architecture, non-functionals, security model, evaluation and release plan
3. Build, harden, & integrate	Landing zone, identity and permissions, safe system-of-record integration, telemetry
4. Validate & release	Evaluation, security and resilience evidence; support readiness; go-live approval
5. Launch & adopt	Controlled rollout, hypercare, enablement and baseline value measurement
6. Operate & optimize	Managed AgentOps: monitoring, versioning, evaluation, access and cost governance
7. Scale	Replicate proven patterns to adjacent workflows - or retire what underperforms

Right-sized autonomy

Not every AI application should act autonomously. The framework classifies agent behavior on a five-level ladder - from assistance and drafting through bounded execution to orchestration - with approval expectations and human-override requirements defined at each level. Autonomy is earned with evidence, not assumed.

How engagements are structured

Module	What it delivers
Production Readiness Assessment	2-4 weeks fixed-fee assessment: scorecard, gap map, target pattern, roadmap and investment case - fee credited toward delivery
AI Landing Zone & Production Hardening	Identity, network, secrets, logging, policy guardrails, environments and release controls
Workflow Integration & Reliability	System-of-record integration, action safety, evaluation gates, adoption and controlled launch
Managed AgentOps	Ongoing monitoring, incidents, change control, content and access operations, cost and monthly value reporting

Proven production patterns

- **Contract intelligence & drafting**- grounded contract Q&A with clause citations, compliant drafting and deviation detection - in production for enterprise legal and procurement
- **Sales operations agent** - forecast hygiene and pricing guardrails delivered in Teams with safe CRM write-back - live at a Nasdaq-listed semiconductor leader; ~85% of requests resolved without an LLM call
- **Enterprise knowledge & policy agent** - permission-aware answers with citations across governed content - a replicable pattern for HR, IT and compliance knowledge

Grounded in recognized guidance

The framework aligns to the Microsoft Cloud Adoption Framework for AI. The NIST AI Risk Management Framework, Generative AI Profile, and Microsoft Responsible AI practices - combined with Netwoven's AI production delivery experience as a Microsoft Solutions Partner across all six solution areas.

Start the conversation

The best first step is a complimentary 2-hour AI Pilot-to-Production workshop: your team inventories its AI initiatives, scores the strongest against the eight gaps, and leaves with a red/amber/green heatmap and a 30-day plan - yours to keep, whatever you decide next.

netwoven.com • info@netwoven.com • +1 877 638 9683

© Netwoven Inc. Public edition 1.0, July 2026. Statistics: MIT NANDA, *The GenAI Divide* (2025); S&P Global Market Intelligence (2025); Gartner (June 2025).