

Generative AI Roadmap: Priorities, Milestones, and Communication



Roadmap = portfolio of bets, platform enablers, governance gates, funded owners



Sequence GenAI initiatives with defensible priorities and evidence-based milestones



Most enterprises run several GenAI initiatives; average holds 3.4 at once



Governance maturity and ownership lag adoption – the bottleneck is organizational



You leave with: roadmap canvas, milestone gate definition, executive narrative template



Why GenAI Roadmaps Break Traditional Planning Assumptions



Model capability, token economics, and regulation shift faster than annual release cycles



78% of enterprises now have GenAI in production; sequencing, not starting, is the constraint



Average organization runs 3.4 concurrent AI initiatives, up from 1.8 a year ago



Pilots stall on escalating cost, unclear value, and weak risk controls — rarely model quality



Manage uncertainty with staged commitments and kill criteria written before work starts




Agentic systems take actions, not just return text, raising the governance bar



The GenAI Use-Case Portfolio:

Quick Wins, Strategic Bets, Platform Enablers, Experiments


- Four archetypes: quick wins, strategic bets, platform enablers, experiments
- Quick wins: low-risk productivity; experiments: learning with kill criteria
- Strategic bets: differentiated, higher risk; platform enablers: shared infrastructure
- Score every use case on five dimensions, including value, feasibility, data readiness, risk, readiness
- Balance horizons: near-term productivity, medium-term customer impact, long-term capability
- Start with a Tier-1 low-risk use case to build earned capability; avoid pet projects
- Publish scoring rationale and set a fixed review cadence; guard against vendor-driven prioritization



Risk Tiering as a Prioritization Input


HIGHER CLARITY
SAFER INNOVATION
GREATER IMPACT

- Tier by consequence, not technology:
low, moderate, high, critical
- Four risk categories: technical, operational, ethical, regulatory
- Score on output consequence, data sensitivity, visibility, reversibility
- High on any two dimensions moves the use case up a tier
- Common misclassification: automated decisions about people
- When in doubt, classify up — strict where it matters, light elsewhere



Milestone Architecture: From Proof of Value to Scaled Adoption

DISCIPLINE
PROGRESS
IMPACT



-  Six gates: proof of value, bounded pilot, production hardening, governance approval, staged rollout, adoption
-  Write entry and exit criteria before work starts — quality, cost, latency, safety, adoption
-  Align each gate to funding release and executive review cycles
-  Fewer, harder, evidence-based gates beat many soft checkpoints
-  Two-directional contract: thresholds to go live and thresholds to stop

Production Readiness: The Five Gates Before Scaling

ASSESS > HARDEN > SCALE


- Score five capabilities 0–3 before committing the scaling budget
- Monitoring, evaluation, staffing, integration, funded data pipeline
- Below 10 of 15 means pilot, not production plan
- Hardening a focused workflow runs 6–14 weeks on top of pilot spend
- Budget roughly a third integration, a third AI logic, a third monitoring
- Red flags: no named owner, no multi-quarter budget, engagement-only metrics

FIVE PRODUCTION READINESS GATES SCORE 0–3 BEFORE SCALING



Sequencing Dependencies and Platform Enablers

- Identify shared enablers early: data access, retrieval, model gateway, evaluation harness
- Reuse retrieval and eval capability across multiple bets; avoid duplicate spend
- Balance platform investment against visible business value; platform-only roadmaps lose sponsorship
- Map dependencies across data, legal, security, procurement, change management
- Security review and data access stall work more than model selection
- Route across a model portfolio by cost, sensitivity, latency, control



Operating Model, Ownership, and Decision Rights



Define roles: programme lead, workflow owner, product owner, engineering, data, risk, change, sponsor



Settle decision rights in writing: use-case selection, vendor choice, risk acceptance, funding release



Most-cited stall reason: unclear ownership; the accountable owner sits in the business unit



Cadence: weekly delivery stand-up, monthly portfolio review, quarterly strategy refresh




One source of truth for roadmap status, risks, dependencies, and the AI system inventory

Role	Select use cases	Choose models	Accept risk	Release funding	Stop a deployment
Programme lead	☒	☐	☐	☐	☐
Workflow owner	☒	☐	☐	☐	☐
Product owner	☒	☒	☐	☐	☐
Engineering	☐	☒	☐	☐	☐
Data governance	☐	☐	☒	☐	☐
Risk	☐	☐	☒	☐	☐
Change management	☐	☐	☐	☐	☐
Sponsor	☐	☐	☐	☒	☒



Governance and Responsible AI as Roadmap Enablers

BUILT IN. MOVING FORWARD TOGETHER.



- Embed governance into milestone gates, not final approval
- Six pillars: risk classification, standards, compliance, security, lifecycle, oversight
- Launch prerequisites, not afterthoughts: model card, data card, impact assessment, audit trail
- Map once across EU AI Act, NIST AI RMF, ISO/IEC 42001, SOC 2 — not parallel programs
- Mature governance reports materially fewer production AI incidents

Metrics and Evidence That Survive Executive Review



Delivery



Model Quality



Adoption



Business Value



Four metric families — delivery, model quality, adoption, business value — never substitute one for another



Core decision metric: fully loaded cost per verified successful workflow outcome vs current process



Minimum pilot scorecard: task success, human acceptance, cycle time saved, incident rate, cost per success



Every metric needs a pre-pilot baseline and automatic rollback thresholds set before launch



Evaluation harness: real cases with known-good outcomes, run in CI on every change

Communication Strategy: Audiences, Narrative, and Evidence



Key Guidance

-  Segment audiences: exec sponsors, business units, technical teams, risk, end users
-  Tailor each cut: executives need value and risk; engineers need architecture and trade-offs
-  Lead with evidence — pilot results, user metrics, cost data, risk posture
-  One narrative spine: why now, what we bet on, what we prove next, what we need
-  Avoid four pitfalls: overpromising, jargon, ignoring risk, certainty instead of staged commitment

CLEAR
AUDIENCES.
STRONGER
ALIGNMENT.
REAL PROGRESS.

Stakeholder Scenarios: Funding, Kill Decisions, and Bad News

SAME PRINCIPLES. CLEARER CONVERSATIONS. STRONGER OUTCOMES.



Funding ask

Funding ask: present the portfolio, the gates, and cost per successful outcome — not a status tour



Kill decision

Kill decision: killing a pilot on pre-agreed criteria is the programme working, not a failure



Bad news

Bad news: report regressions with detection method, containment action, and the eval case added



Scope creep

Scope creep: expansion waits until production works on the same single workflow the pilot proved



Practice technique

Practice technique: one-page format, the ask stated first, risk disclosed honestly, next milestone dated

Roadmap Workshop:

Build, Stress-Test, and Communicate Your Sequence



Apply prioritization and risk tiering to your chosen GenAI portfolio



Write milestone gates with entry/exit criteria and one named owner each



Stress-test the sequence against dependency, risk, funding, staffing



Confirm platform enablers are funded, not assumed



Commit metrics and stop-side thresholds before launch



Deliver two outputs: one-page executive narrative and 90-day plan



90-Day Action Plan, Common Pitfalls, and Next Steps

FROM ACTION TO ACCOUNTABILITY



Days 0–30

- Days 0–30: publish interim policy, inventory GenAI in flight, assign owners



Days 30–90

- Days 30–90: pilot risk tiering and evaluation on two real use cases



Month 6

- Month 6: 80% of deployments in inventory; top use cases continuously evaluated



Common Pitfalls

- **Avoid:** pilot-forever, no named owner, platform-only, engagement metrics, paperwork governance



Next Steps

- **Commit:** roadmap canvas filled, narrative drafted, next gate scheduled

PersonWise.

PersonWise • AI digital-human course platform



Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/1fc2dc30/public