

Communicating Generative AI to Stakeholders

- GenAI communication is a distinct discipline: probabilistic outputs, rapid change, high expectations
- Tailored for product, operations, L&D, and governance audiences
- Roadmap: shared vocabulary, stakeholder concerns, realistic expectations, governance, messaging, measurement
- Outcomes: explain in plain language, set realistic expectations, address concerns, align decisions



Why GenAI Communication Fails: Three Root Causes



Probabilistic systems:
same prompt, different
outputs; demos aren't
proof of reliability



Rapid change:
model-specific
judgements and policies
age badly within months



Expectation gap:
hype and unmanaged
risk narratives both
erode trust



**Communication
Fails**



Avoid:

feature tours instead of decisions, false certainty, silence after setbacks

Shared Vocabulary in Plain Language



Key terms: foundation and frontier models, prompts, fine-tuning, RAG, agents, hallucination



Grounding ties outputs to verifiable sources; improves accuracy, never guarantees it



Vendors, regulators, and academics differ on terms—agree on one internal glossary




Plain-English definitions first; technical detail as a second layer for specialists



Capability, Reliability, and the Metaphors We Use

Probabilistic (GenAI)

Composes new content



Deterministic (Traditional Software)

Retrieves or transforms



-  Probabilistic vs. deterministic: GenAI composes new content; traditional software retrieves or transforms
-  Strong at drafting, summarizing, classifying, and adapting language at scale
-  Fails at exact figures, real-time lookups, and high-stakes output without human review
-  Hallucination is how these models work, not a bug awaiting a patch
-  Better metaphor: a brilliant generalist who sometimes confuses what it read with what is true
-  Avoid misleading metaphors: “it thinks,” “it knows”

Stakeholder Map: Who Needs What, and Why



Product

Needs:

model quality evidence, evaluation results, and pilot-to-production gaps

Decision:

What to build and release



Operations

Needs:

workflow impact, SOPs, escalation paths, and human-in-the-loop checks

Decision:

How to operate and scale



L&D

Needs:

skill-shift clarity, safe experimentation permission, and role-specific application

Decision:

How to enable and train



Governance

Needs:

traceable controls, documentation, and clear regulatory mapping

Decision:

How to approve and oversee



Map influence and interest to tailor depth, cadence, and channel

Top Concerns by Role and How to Address Them

PEOPLE • RISKS • BETTER CONVERSATIONS • BRIGHTER OUTCOMES

Survey: privacy and security (73%), compliance (50%), governance (46%)




Operations worry about data leakage, output reliability, and limited skills




Stakeholders ask:
Is it relevant?
Valuable? Risky?
Trustworthy?



Match channels to decisions: briefings, demos, FAQs, office hours



Setting Realistic Expectations: Demos vs. Production



Demos show possibility; production demands latency, cost, evaluation, monitoring, and maintenance



Frame pilots as learning: clear success criteria, stop criteria, named owners



Counter 'AI will do everything' and 'AI will replace us' — both crowd out useful decisions



Model choice and agentic workflows drive spend; tie AI consumption to value created

Communicating Uncertainty Without Losing Credibility



Segment confidence: name where the system performs reliably and where caution is required



Three-part frame: what is validated, what is conditional, what stays under human review



Explain the consequence of each error type and the review rule that mitigates it



Drop jargon: say "needs more checking," not "calibration drift" or "confidence attenuation"



Validated



Conditional

Needs more checking



Human-reviewed


Risk, Governance, and Regulation in Plain Language

- Six core risks: accuracy, privacy, security, bias, IP, vendor dependency
- EU AI Act: Article 50 transparency live since August 2, 2026
- Pre-existing systems on transition must comply by December 2, 2026
- US: no federal AI law; 280 state bills enacted as of October 2026
- Map state obligations; route questions to legal, privacy, security, risk

EU AI Act Timeline



US State AI Obligations



US: no federal AI law;
280 state bills enacted
as of October 2026

Six Core Risks



accuracy



privacy



security



bias



IP



vendor
dependency



Map state obligations;
route questions to legal,
privacy, security, risk

Messaging Patterns That Work

CLEARER MESSAGES. BRIGHTER OUTCOMES.



Situation, capability, limits, control, ask: end every message with a decision or action



Match evidence to audience: trade-offs for leaders, workflow guidance for operators, traceable evidence for auditors




Launch rules: short announcements, measurable intent over superlatives, link policy, training, and help channel



Pre-approve answers on data use, job impact, liability, and vendor lock-in with a clear escalation path

Communicating to L&D and Operational Audiences



Communicating When Things Go Wrong

BE CLEAR. BE TIMELY. BUILD TRUST.



Standard incident template: what happened, impact, root cause, actions, prevention, next update



Speed over perfection: say "we are investigating" rather than staying silent



Own the issue plainly: "our model made an error"



Separate shared facts from function-specific implications to prevent competing interpretations



Follow through visibly: close the loop with stakeholders who raised concerns

Incident Template



What happened



Impact



Root cause



Actions



Prevention



Next update

Measuring Understanding, Trust, and Adoption



Comprehension



Calibrated Trust



Adoption



Escalation Quality

- Measure comprehension, calibrated trust, adoption, and escalation quality—not logins alone
- Lightweight tools: three-question pulse surveys, help-channel question logs, short interviews
- Leading indicators: time-to-first-successful-workflow, repeat question rate, manager-reported clarity
- Tie communication cadence to model, vendor, and regulatory change
- Turn these metrics into a repeatable rhythm



Workshop: Crafting a One-Page Stakeholder Brief



Scenario exercise: draft a one-page GenAI brief for one stakeholder group



Peer review against situation, capability, limits, control, ask



Revision round: add risk framing, controls, and question-routing guidance



Commitment: one communication action within two weeks, one metric to report



Close the loop: share results with the group at the next session

One-Page Stakeholder Brief



Situation



Capability



Limits



Control



Ask

Peer-Review Checklist











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/6135d8bb/public