Utility Agentic AI Governance Checklist

A practical governance checklist for local-first, human-reviewed agentic AI pilots in transformer APM, condition-based maintenance, and critical infrastructure evidence workflows.

Use this with a controlled GridAPM pilot request. Keep AI draft output human-reviewed before it becomes reportable.

1. AI workflow purpose

[ ] Name the exact transformer APM workflow: evidence pack, CBM review, maintenance-window package, event handoff, data contract, or governance review.

[ ] State the intended AI assistance: summarize approved evidence, identify gaps, draft questions, prepare rationale, or format reviewer-ready packages.

[ ] State the prohibited use: autonomous control, protection changes, switching, operating-limit approval, interconnection approval, final diagnosis, or maintenance approval.

2. Source boundaries and provenance

Control	Governance question	Status
Approved sources	Which DGA, PRPD, SFRA, thermal, inspection, maintenance, event, CMMS/EAM, historian, or SCADA exports are allowed?	[ ]
Source owners	Who owns each source and who confirms it is valid for pilot use?	[ ]
Units and time	Are timestamps, units, sampling dates, instrument context, and asset identity preserved?	[ ]
Sensitivity	Which fields require masking, local-only handling, or restricted access?	[ ]
Export boundary	Which outputs may be exported to a report, evidence pack, EAM/CMMS note, or meeting pack?	[ ]

3. Human review and authority

[ ] Name the human reviewer for each pilot output before it becomes reportable.

[ ] Require review states: draft, reviewed, edited, rejected, approved, escalated.

[ ] Require reviewer comments explaining accepted, rejected, or modified AI draft language.

[ ] Require an explicit statement that GridAPM output is review material until approved by qualified personnel.

4. Audit trail and monitoring

[ ] Record source files, extracted facts, prompt context, draft rationale, missing evidence, reviewer edits, and final exported output.

[ ] Track repeated missing evidence, rejected draft statements, unsupported assumptions, and reviewer corrections for pilot learning.

[ ] Define retention, deletion, export, and access expectations before the pilot starts.

[ ] Review whether additional cybersecurity, privacy, OT, legal, or procurement review is needed before moving beyond pilot scope.

5. Go / no-go questions

[ ] Is the workflow narrow enough for a first pilot?

[ ] Are approved evidence sources available without live integration?

[ ] Are prohibited AI actions written down?

[ ] Are reviewer roles and approval paths named?

[ ] Can the team measure preparation effort, traceability, missing evidence, and work-package quality?

[ ] Can the team stop or revise the pilot if evidence quality or governance controls are insufficient?

Boundary and disclaimer

This checklist is a planning aid for human-reviewed utility AI governance. It is not legal advice, cybersecurity assessment, compliance certification, operating procedure, procurement advice, engineering diagnosis, or maintenance approval. It does not authorize autonomous control, switching, protection changes, interconnection approval, operating-limit decisions, or final transformer condition conclusions.

Recommended official references

NIST AI Risk Management Framework: https://www.nist.gov/itl/ai-risk-management-framework

NIST Artificial Intelligence: https://www.nist.gov/artificial-intelligence

NIST SP 800-82 Rev. 3 Guide to Operational Technology Security: https://csrc.nist.gov/pubs/sp/800/82/r3/final

NSA/CISA AI in operational technology guidance: https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/4347041/nsa-cisa-and-others-release-guidance-on-integrating-ai-in-operational-technology/

Anthropic Building Effective Agents: https://www.anthropic.com/engineering/building-effective-agents

OpenAI How Agents Are Transforming Work: https://openai.com/index/how-agents-are-transforming-work/