WireNet Safe Deployment Checklist

A practical checklist for bounded, reviewable and reversible AI deployments, informed by the NIST AI Risk Management Framework.

AUTHOR	PUBLISHED	VERSION	FORMAT
Florian Gitt	13 August 2026	0.1	CHECKLIST / 8 MIN

Free to share and adapt with attribution under CC BY 4.0.

Practical guidance only. This resource is not certification, legal advice or a compliance guarantee.

How to use this checklist

Record evidence, not only yes or no.

Use this checklist before an AI workflow crosses from experiment into recurring use. Name an owner, review every item and adapt the depth of evidence to the actual consequences of the system. Treat unresolved high-consequence items as release blockers.

> **PRACTICAL GUIDANCE** This checklist is not certification, legal advice or a compliance guarantee. Sector-specific obligations and professional judgment still apply.

1. Context and consequences

- [ ] The intended users, affected people and decision context are named.
- [ ] A specific benefit is defined; AI is not the objective by itself.
- [ ] Foreseeable harms, misuse and unequal effects have been considered.
- [ ] The deployment is classified by consequence and reversibility.
- [ ] A non-AI or lower-complexity alternative has been considered.

2. Data and source access

- [ ] Every data source has an owner, purpose and lawful basis for use.
- [ ] Sensitive, confidential, personal and copyrighted material is identified.
- [ ] Training use, provider retention and cross-border processing are understood.
- [ ] The system can distinguish authoritative sources from unverified material.
- [ ] Inputs, outputs and logs have appropriate retention and deletion rules.

3. Minimum permissions

- [ ] The system receives only the tools and access required for the task.
- [ ] Read, write, publish, send, purchase and delete permissions are separated.
- [ ] Secrets are stored outside prompts, repositories and logs.
- [ ] External actions require explicit gates proportionate to their consequences.
- [ ] Permission changes are reviewable and can be revoked quickly.

4. Human responsibility

- [ ] One accountable owner is named for the deployment and its outcomes.
- [ ] Reviewers have the expertise, time and evidence needed to challenge results.
- [ ] Users know when they are interacting with or relying on AI-generated work.
- [ ] Escalation paths exist for uncertainty, disagreement and affected people.
- [ ] Automation does not obscure who made or approved a consequential decision.

5. Evaluation before release

- [ ] Success and unacceptable failure are defined before testing.
- [ ] Evaluation cases represent real work, edge cases and likely misuse.
- [ ] Accuracy, source fidelity, bias, usability and safety are tested where relevant.
- [ ] The complete workflow - not only the model response - is evaluated.
- [ ] Release criteria and known limitations are documented and approved.

6. Monitoring in use

- [ ] Quality, failures, user complaints and material context changes are monitored.
- [ ] Logs are sufficient for diagnosis without collecting unnecessary data.
- [ ] Thresholds trigger review, restriction or suspension.
- [ ] Model, prompt, tool, data and policy changes are versioned.
- [ ] Evaluation is rerun after material changes and on a defined cadence.

7. Incident handling

- [ ] People know how to report harmful, incorrect or unexpected behaviour.
- [ ] The response owner, communication path and evidence-preservation steps are named.
- [ ] Access can be restricted while an incident is investigated.
- [ ] Affected people and relevant partners can be informed appropriately.
- [ ] Lessons are translated into updated controls and evaluation cases.

8. Ordered shutdown

- [ ] The deployment has a tested stop or rollback procedure.
- [ ] Critical work can continue through a manual or alternative process.
- [ ] Access tokens, integrations and scheduled jobs can be revoked.
- [ ] Data, logs and retained outputs have a documented disposition.
- [ ] The decision to restart requires a fresh review of the triggering issue.

Release record

Make the decision inspectable.

FIELD	RECORD
Deployment owner	Name and role
Purpose and users	What the system does, for whom
Risk level	Consequences, affected people, reversibility
Evaluation evidence	Cases, results, limitations, reviewer
Release decision	Approved, conditional or blocked; date and approver
Monitoring review	Signals, thresholds and next review date
Shutdown owner	Who can stop it and how

Relationship to NIST AI RMF

This checklist is informed by the NIST AI Risk Management Framework and its emphasis on governing, mapping, measuring and managing AI risk. It is deliberately shorter and operational. It does not replace the framework, sector-specific requirements, legal review or professional judgment.

Reference: https://www.nist.gov/itl/ai-risk-management-framework

A DEPLOYMENT IS READY ONLY WHEN ITS LIMITS ARE AS LEGIBLE AS ITS PROMISE

If responsibility, evidence or shutdown remains vague, the system is still an experiment.

Canonical HTML: https://wirenet.dev/resources/safe-deployment-checklist