

The Plain-Language Safety Rules

Nurse AI OS · Safety & Governance. You memorized the rights of medication administration. These are shorter — and they protect your patients, your license, and you, at every level of AI use.

Agents propose. Humans judge. Nurses steward.

Free for nurses and nursing students · Founding Year 2026 · nurse-ai-os.org

The five rules

The one-minute test before entering anything sensitive: if you cannot name the approved environment, the data type, the AI feature, and the accountable human owner — don't put it in the tool.

1 · No patient data. Ever.

No names, charts, screenshots, room numbers, record numbers, dates, or stories that could identify a patient — in any AI tool, at any level of this program. If you're unsure whether something counts: it counts. Leave it out.

2 · AI drafts. You decide.

Every AI output is a draft until a human reviews it. Nothing gets sent, published, installed, or acted on without your judgment. Human gates are checkpoints, not suggestions.

3 · Your license is not outsourceable.

No AI here supports clinical decisions, assessments, or patient-specific care questions. The system carries your busywork — never your judgment, never your accountability.

4 · Employer systems stay at work.

No EHR screens, no confidential policies, no identifiable staff or student data, no workplace secrets. Public, fictional, de-identified, or personal material only — unless your organization's compliance team has explicitly approved a tool and workflow.

5 · When in doubt, ask a human.

Bring gray areas to the monthly Lamp Huddle or email robert@nurse-ai-os.org. A question asked early beats an incident reported late.

How EDENA applies

EDENA is the governance framework used by Nurse AI OS. Its full expansion is not recorded in the current NIN–NAIO Master Directive v1.1, so this poster does not invent one. The current model classifies three independent dimensions before work begins:

- **Risk:** Green, Yellow, Orange, Red-P (prohibited), or Red-E (exceptional high-risk).
- **Data:** D0–D4. This public poster and Governance Kit support D0 and D1 only; no PHI, employer-confidential data, or patient-adjacent data.
- **Action:** Observe, Draft, Recommend, Prepare Action, Act With Approval, or Constrained Autonomy. Unrestricted action is prohibited.

Green supports **Observe or Draft. Recommend begins at Yellow** and requires evidence appropriate to the claim plus a named human reviewer. Orange, Red-E, D2–D4, and every action mode above Recommend require separate formal authorization and implemented controls; this poster and the public Governance Kit do not grant either. Red-P is refused.

These are advisory instructions, not mechanical enforcement. Humans remain responsible for every decision and action.

Before installing any new routine — the seven questions

Purpose · Data involved · PHI? (if yes: stop) · Autonomy level · Human approval point · What failure looks like · Go or no-go. If you can't answer all seven, the routine isn't ready.

Carry the lamp. Keep the ledger.

© 2026 Robert Domondon · Not medical advice · Not for PHI · Not a replacement for licensed judgment · nurse-ai-os.org/directive.html

nurse-ai-os.org · No PHI · Agents propose. Humans judge. Nurses steward. · This PDF is generated from assets/safety-rules-edena.md — the web copy is canonical.