

Nurse AI OS™ — Media Kit

Directive v1.1 canonical status

Canonical English source-control statement: Nurse AI OS is the **governed professional environment**. The public Community experience is available now. Florence-X and EDENA capabilities become binding only where implemented and evidenced. Institutional authority, credentials, conformance, PHI eligibility, and high-consequence use require separate formal gates. Agents propose. Humans judge. Nurses steward.

- **Available now:** browser-first Community experience; D0/D1 only; no PHI; human review; no external action.
- **Capability-gated:** Florence-X orchestration, EDENA mechanical enforcement, memory, tools, connectors, and automation only where a named implementation and evidence exist.
- **Formal-gate programs:** credentials, institutional or sovereign environments, D2–D4 eligibility, high-consequence use, and assurance require separately published authorization and program gates.
- **Conformance:** NAO Conformant is unavailable until the assurance program exists and the specific release passes it.
- **Architecture evidence:** pre-Directive reports are historical implementation evidence, not current canonical architecture or conformance claims.
- **Licensing:** licensing is artifact-specific. Existing Apache and CC BY grants in this repository remain valid and are not silently narrowed or revoked. Check terms attached to separately maintained artifacts; the public open-core boundary remains pending.

Current through July 28, 2026

Website: nurse-ai-os.org · **Media contact:** robert@nurse-ai-os.org

In one sentence: Nurse AI OS is a nurse-designed, browser-first kit that helps nurses use ChatGPT or Claude with practical workflows, a personal values-and-boundaries file, and explicit human-review rules—without handing over judgment.

This kit provides current, quotable language for journalists, nursing organizations, educators, technology reviewers, conference producers, and collaborators. Product and governance claims are deliberately

bounded. Please use the accuracy notes before publication.

The story in brief

Nurses are already using general-purpose AI for studying, planning, writing, organization, career development, and everyday administration. The central question is not whether they will use AI. It is whether they will use it with clear boundaries, practical judgment, and a professional voice in how the technology is governed.

Nurse AI OS gives nurses a concrete starting point. Today, it is a sophisticated configuration and workflow kit for AI tools people may already have. It includes guided prompts, practical templates, local folders, role-based workflows, safety practices, and a personal **SOUL file** that records the user's values, boundaries, and preferred way of receiving help.

The default path runs in a web browser with ChatGPT or Claude. No separate application or Hermes installation is required to begin. Nurses who later want an AI agent to work with local folders, reusable skills, and explicitly approved tools may add **Hermes Agent**, a separate open-source desktop runtime from Nous Research.

Nurse AI OS does not provide clinical authority, institutional approval, uniform enforcement across AI hosts, or a finished standalone operating system. Its direction is a portable, nurse-governed control layer in which policies, human gates, and evidence can travel across tools.

Why Nurse AI OS matters across healthcare roles

The strategic value of Nurse AI OS is not that it gives every healthcare professional more AI tools. It is that it helps people learn how to direct, question, supervise, and improve AI-assisted work before they encounter higher-consequence institutional systems. Personal Edition is designed as a no-PHI workforce-development environment: a flight simulator for intelligent healthcare work, not an EHR-connected clinical platform or autonomous decision engine.

Its purpose is to build five forms of human capital: stronger reasoning and synthesis; reusable workflows; demonstrated professional development; better preparation and coordination; and the stewardship capacity to recognize risk, preserve dignity, question AI, and stop when human authority is required.

SOUL, Mission Control, and MCP do three different jobs

- **SOUL tells the system who it serves.** A user-reviewed SOUL records purpose, role constellation, working and learning preferences, approved memory choices, governance boundaries, and decisions

the AI must not make. It is a human constitution for assistance—not a psychological dossier, credential, scope determination, or authority grant.

- **Mission Control turns intention into governed work.** A shared core and separate role dashboards organize priorities, projects, learning, reusable workflows, approval queues, evidence, and progress. Missions follow an Assess → Define → Plan → Implement → Evaluate discipline. Progress is based on reviewed evidence, not clicks or AI-use volume.
- **MCP defines what Hermes may eventually reach and do.** A connection must specify the role that may use it, allowed information and actions, prohibited data classes, approval points, logging, failure behavior, revocation, and reversibility. Read, Draft, and Act are distinct permission levels. A connector never grants professional or institutional authority.

The intended progression is **personal identity** → **governed role workspace** → **reviewed capability connection**. These layers are complementary, not interchangeable.

What the architecture can bring to each role

Role lane	Why Nurse AI OS matters	What SOUL and Mission Control organize	What optional, separately approved connections could add
Prelicensure student or nursing assistant	Builds disciplined learning and AI habits before clinical authority is at stake.	Learning preferences, academic-integrity boundaries, study maps, simulation preparation, skills tracking, reflection, and portfolio evidence.	Approved personal notes, files, calendar, learning platform, tasks, and references; no restricted course content or authority by pathway selection.
Bedside or advanced-practice nurse	Creates a governed space for professional growth and workflow reflection without becoming a clinical system.	Certification planning, policy learning, professional portfolio, shift-life organization, and synthetic or non-sensitive workflow analysis.	Approved personal calendar, email, files, documents, tasks, and public references; no patient cases, charts, diagnosis, treatment, or institutional action.
Medical resident or physician	Supports evidence study, research preparation, administrative organization, and longitudinal formation while preserving supervision.	Board preparation, journal review, research and quality-improvement planning, milestones, wellbeing limits, and portfolio evidence.	Approved public literature, citation tools, personal files, documents, tasks, and calendar; no charts, sign-out, patient-specific reasoning, evaluation, or entrustment.
Nurse educator or preceptor	Helps educators design learning while keeping faculty judgment, learner dignity, and academic authority human.	Curriculum drafts, lessons, simulations, rubrics, learner pathways, educational media, assessment drafts, and course evaluation.	Approved learning platform, documents, slides, forms, calendar, and design tools; no student records, autonomous grading, hidden ranking, or accreditation claims.
Advanced Studies overlay	Keeps specialization, certification, graduate work, and lifelong learning connected to one coherent professional identity.	Reading synthesis, thesis or capstone milestones, certification plans, learning evidence, and development priorities across roles.	Approved research databases, citation manager, calendar, files, learning platform, and writing tools; no credential or eligibility claims.

Role lane	Why Nurse AI OS matters	What SOUL and Mission Control organize	What optional, separately approved connections could add
Charge nurse or nurse manager	Converts leadership pressure into clearer priorities, accountable plans, and visible follow-through without creating a shadow HR system.	Daily briefs, meeting preparation, action tracking, policy learning, project charters, education plans, and synthetic or aggregate improvement work.	Approved files, documents, calendar, spreadsheets, forms, and project tools; no workforce surveillance, personnel decisions, staffing action, or organizational authority.
Hospital administrator or clinic manager	Provides a rehearsal space for safer strategy and pilot design before proposals touch live people or systems.	Strategic plans, pilot charters, policy drafts, implementation roadmaps, meeting intelligence, and non-sensitive operational evidence.	Approved files, forms, documents, calendar, spreadsheets, and project tools; no live intake, patient scheduling, workforce data, or institutional-system access.
Quality, safety, governance, research, or innovation leader	Makes evidence, uncertainty, risk, provenance, and decision ownership visible throughout improvement work.	Evidence maps, PDSA drafts, risk registers, policy comparisons, protocol preparation, traceability, evaluation plans, and claim-versus-proof review.	Approved public research sources, documents, spreadsheets, forms, project systems, repositories, and analytics; no QI/research determination, participant data, or official reporting authority.
Wellness and personal-life manager	Protects sustainable workload, relationships, learning, and recovery as part of professional capability.	Personal goals, routines, family logistics, reflection, digital wellbeing, and learning-life balance.	Personal calendar, notes, tasks, files, and optional wellness tools chosen by the user; no medical monitoring or employer visibility by default.
Entrepreneur, consultant, creator, or founder	Turns ideas into governed offers and reusable operations while preserving consent, claims, and release authority.	Market research, service design, content plans, proposals, customer workflows, product roadmaps, and business evidence.	Approved email, calendar, documents, project and design tools, website systems, repositories, or CRM; no automatic outreach, publishing, contracting, or financial authority.

Role lane	Why Nurse AI OS matters	What SOUL and Mission Control organize	What optional, separately approved connections could add
Builder or AI orchestrator	Gives healthcare innovation an accountable path from problem definition to testable artifact and supervised agent workflow.	Requirements, agent specifications, prototypes, tests, evaluation evidence, version history, stop rules, rollback, and release gates.	Approved repositories, development environments, documentation, project systems, and test tools; no secret access, hidden delegation, merge, deployment, or production release without fresh authorization.

One person may use several coordinated role dashboards beneath one approved SOUL—for example, bedside nurse, graduate student, educator, wellness manager, and emerging founder. The dashboards keep contexts distinct; they do not create fragmented identities or transfer authority between roles.

What happens when a user gives these layers to Hermes

Hermes Agent is a separate optional runtime. Giving Hermes a SOUL export, Mission Control build kit, or MCP companion is a request for inspection—not permission to install or connect.

1. Hermes verifies the exact files, checksums, source inventory, compatibility, target location, existing work, and rollback plan in read-only preflight.
2. The user previews and approves the minimum derived SOUL information. Raw reflective answers are not retained by default, and the user chooses what may or may not be remembered.
3. Hermes proposes how the approved SOUL would guide a role-specific Mission Control while keeping credentials, scope, employment, enrollment, and institutional permission unverified.
4. For a self-install build kit, Hermes displays the exact Implementation Activation Card and stops. Without approval of that exact card, it makes no installation mutation.
5. If approved, Hermes builds only the written local scope, preserves compatible work, keeps agents and external capabilities inactive by default, and produces visible evidence and rollback information.
6. Any MCP companion is reviewed separately. The current published role companions are blocked, disabled, and expose zero tools; their candidate dependency records are not complete transitive locks or activation-ready runtimes.
7. Each future connection requires a new card covering exact tools, data, paths, credentials, provider terms, destinations, logging, retention, revocation, failure behavior, reversibility, and expiry.
8. Read permission does not imply Draft permission; Draft does not imply Act. Every consequential external action remains subject to a destination preview and explicit approval by an authorized human.

Downloading, opening, unzipping, selecting a role, or handing files to Hermes does not import a SOUL, build Mission Control, install a profile, authenticate an account, activate a connector, create persistent memory, schedule work, send, publish, or act. Current role packages vary in maturity and availability; this architecture describes the governed Personal Edition direction, not a claim that every role dashboard or connection is complete today.

Fast facts

- **What users receive today:** a browser-first guided setup, 2-minute SOUL Quiz, editable values-and-boundaries file, prompts, templates, local folder structure, role-based workflows, and safety practices.
 - **What users need to begin:** access to ChatGPT or Claude in a browser. A phone, tablet, Chromebook, Mac, or Windows computer can be used for the browser-first path.
 - **What is optional:** Hermes Agent on a recent Mac or Windows computer for advanced local-folder, reusable-skill, persistence, and approved-tool workflows.
 - **Who it serves:** nursing students, bedside nurses, nurse leaders, nurse educators, nurse practitioners, nurse entrepreneurs, and other nurses exploring responsible AI use.
 - **Core cost: \$0, free forever** for core browser-first setup and self-customization. No credit card and no automatic upgrade.
 - **Separate optional offerings:** a tailored profile/setup package is **\$10 one time**; developed-resource access is **\$29.90 per year** as resources are released; advanced courses include free and paid options.
 - **Summit:** free and paid advanced courses are planned for the virtual summit on **September 15–17, 2026**. Course-specific prices must be shown before enrollment.
 - **Founding-year recognition:** community recognition and related founding-year benefits continue through **December 31, 2026**. That date does not end the free core or create future automatic billing.
 - **Privacy boundary:** users are told not to enter patient information, workplace credentials, secrets, or identifiable operational records. Material pasted or uploaded to an AI host is handled under that provider's privacy and data settings.
 - **Authority boundary:** Nurse AI OS is not an EHR, clinical decision-support system, autonomous clinician, certification body, or source of institutional approval.
 - **Doctrine:** *Agents propose. Humans judge. Nurses steward.*
-

Quotable descriptions

Short description

Nurse AI OS™ is a browser-first, nurse-designed kit for using ChatGPT or Claude with practical workflows, personal boundaries, and human review. It helps nurses learn, organize, plan, and build without treating AI output as judgment. The core setup and self-customization are free forever; Hermes is a separate optional desktop runtime.

Standard description

Nurse AI OS™ helps nurses use the AI they may already have without handing over professional or personal judgment. The browser-first kit combines a 2-minute SOUL Quiz, an editable values-and-boundaries file, guided prompts, practical templates, local folders, role-based workflows, and safety practices. Users can start from a phone, tablet, Chromebook, Mac, or Windows computer with ChatGPT or Claude. Nothing requires Hermes to begin. Nurses who want advanced work with local folders, reusable skills, persistence, and approved tools may later add Hermes Agent as a separate desktop runtime. The project prohibits patient information in its intended workflows, keeps clinical and institutional decisions outside the tool, and distinguishes guidance from enforceable controls. Core setup and self-customization remain free forever.

Founder quote

“Nurses should not be treated as passive end users of AI. We should be among the people who set the boundaries, notice when a system fails in real workflow, and decide when technology must stop.”
— **Robert Domondon, founder**

Three message pillars

1. **Practical agency:** nurses can start with one useful browser workflow rather than waiting for a new hospital platform or mastering an agent stack.
 2. **Human primacy:** AI can draft, organize, compare, and propose; accountable people still review, decide, escalate, and own consequences.
 3. **Nurse-led stewardship:** bedside and operational knowledge should shape AI governance, not arrive only after procurement or deployment.
-

Product maturity: Today / Emerging / Direction

State	Public description	Do not infer
Today — practical product	A nurse-centered configuration and workflow kit for ChatGPT or Claude: SOUL file, prompts, templates, folders, role workflows, and safety practices.	Not a new model, EHR, autonomous clinician, or finished standalone OS.
Emerging — technical layer	Public policy artifacts, schemas, tests, release verification, evidence controls, EDENA review language, governed build kits, Florence-X research components, and a tested Integration Contract reference implementation with Mission Control role packets and a synthetic healthcare sandbox.	Not uniform host enforcement, clinical validation, compliance certification, or a deployed institutional control plane.
Direction — fuller OS	A portable, enforceable, nurse-governed control layer in which policies, human gates, provenance, and evidence can travel across tools.	Not a claim that this future architecture is complete or operational today.

New in the public repository: the Integration Contract (July 2026)

In July 2026 the emerging technical layer gained a working, tested reference implementation, published as open code in the project repository. The **Nurse AI OS Integration Contract** defines six stable interfaces — orchestration, policy decision, privacy transformation, memory, knowledge retrieval, and observability — so external AI engines plug in behind nurse-governed boundaries instead of becoming vendored dependencies. Behind those interfaces sit a deny-by-default, fail-closed EDENA policy-gateway prototype; an ADPIE workflow runtime with a hard human-authorization gate; an evidence ledger in which source-backed claims require citations that exist (“no citation, no claim”); five installable Mission Control role packets — pre-licensure student, staff nurse, educator, leader, and licensed clinician — that ship with empty default permissions, integrations off, and propose-only starter agents behind a data-boundary notice; a fully synthetic healthcare sandbox whose layered admission boundary refuses real-looking records; and research governance gates that require a named, recorded IRB-or-equivalent approval before research execution. The committed Integration Contract and Mission Control artifacts are deterministic and are rebuilt and verified in continuous integration.

These artifacts are software-publication evidence for the control-layer direction. They are reference implementations: they do not create uniform host enforcement, compliance certification, institutional authorization, or a deployed institutional control plane, and EDENA’s cross-host posture remains shadow/observe-only.

Safety and governance boundaries

The no-PHI participation rule

Nurse AI OS is designed as a no-PHI learning and stewardship space. Users are instructed not to enter patient names, dates of birth, medical-record numbers, images, recordings, workplace credentials, secrets, or details that could identify a patient or expose an employer system.

This is a participation and design boundary—not a claim of HIPAA certification. A user's selected AI provider may process anything the user chooses to paste or upload under that provider's own settings and terms.

Judgment stays with people

Nurse AI OS materials are for learning, drafting, organizing, reflection, and governed preparation. They are not for patient-specific diagnosis, treatment, triage, prescribing, clinical orders, or replacing qualified supervision. Employer policies, faculty rules, professional scope, and institutional approval still apply.

EDENA

EDENA is the governance, risk, safety, and assurance control plane. It provides review language for equity and ethics, dignity and data, environmental and external effects, nursing relevance and nurse wellbeing, and human agency and action. Its Green/Yellow/Orange/Red-P/Red-E tiers help people classify risk and identify human gates. Its official long-form expansion has not yet been recorded by the Directive owner.

The current public posture is **shadow/observe-only**. EDENA can structure review and produce inspectable policy artifacts; it does not guarantee that ChatGPT, Claude, Hermes, or another host will enforce every instruction identically.

Evidence and release integrity

Selected public build kits and technical artifacts use manifests, checksums, tests, explicit activation boundaries, and human approval gates. Those controls provide software-publication evidence. They do not prove clinical effectiveness, regulatory compliance, institutional authorization, or patient outcomes.

For nursing organizations and educators

Boilerplate: Nurse AI OS™ is a nurse-led, browser-first kit for practical AI learning and stewardship. It helps nurses and students use general-purpose AI with explicit values, boundaries, role-based workflows, and human review. The project's no-PHI rule excludes patient information and workplace credentials from intended use. Educational and governance materials do not confer competence, certification, accreditation,

clinical readiness, or institutional approval. Schools and employers retain their own policies, supervision, assessment, privacy, and authorization responsibilities.

Suggested discussion angles:

- What nurses should ask before AI enters a workflow.
- Why AI literacy must include escalation, refusal, and stop authority—not only prompting.
- How educators can assess process evidence and disclosure without relying on automated accusation.
- Why governance that does not reach real nursing workflow remains incomplete.
- How nurse leaders can begin with low-risk administrative burden while protecting professional judgment.

For technology and AI-governance media

Boilerplate: Nurse AI OS™ is a nurse-centered configuration and workflow layer for existing AI hosts. The browser-first product uses ChatGPT or Claude; Hermes Agent is a separate optional open-source desktop runtime and is not vendored or renamed by Nurse AI OS. Public materials distinguish instruction-following from mechanically enforced controls, and current EDENA governance remains shadow/observe-only. Florence-X is an emerging research and build architecture—not a deployed institutional mission-control service.

Technical angles:

- **Control above capability:** models provide cognition; human authority, policy, provenance, and review determine what may happen.
 - **Guidance versus enforcement:** values files and prompts can guide behavior, while schemas, checks, human gates, and runtime controls can enforce only specific tested boundaries.
 - **Portable governance:** durable files, skills, policies, and evidence are intended to survive model and provider changes; automatic provider failover is not claimed.
 - **Agent supply-chain risk:** AI-authored tools and skills require provenance, scanning, review, and explicit promotion before trusted use.
 - **Nursing as a governance discipline:** assessment, coordination, escalation, continuity, documentation, dignity, and stop authority are transferable design patterns for accountable AI.
-

Founder

Short bio

Robert Domondon, RN, CCRN, CSC, CMC, is a critical-care nurse with 22 years at the ICU bedside and 13 years in healthcare management. His background includes physician training and practice in the Philippines and hospital administration on a former U.S. naval base. He founded Nurse AI OS, the NAIO Institute, and the Nurse Intelligence Network.

Standard bio

Robert Domondon's career spans clinical medicine, hospital administration, healthcare management, AI adoption, and more than two decades of critical-care nursing in the United States. He trained and practiced as a physician in the Philippines, administered a hospital on a former U.S. naval base, and later returned deliberately to bedside nursing. That combination—the physician's clinical lens, the administrator's systems view, and the nurse's lived workflow—shapes his work on nurse-led AI governance. Robert founded Nurse AI OS, the NAIO Institute, and the Nurse Intelligence Network to help nurses build practical AI capability while preserving human judgment, dignity, and accountability.

Ecosystem glossary

- **Nurse AI OS™** — the browser-first nurse workflow kit and developing nurse-governed control-layer direction.
 - **NAIO Institute** — an initiative developing nurse-led AI-governance education, methods, standards proposals, and stewardship infrastructure. It does not currently confer accreditation, licensure, certification, or institutional approval.
 - **EDENA** — the governance, risk, safety, and assurance control plane; a stewardship lens and risk-review language currently operating in shadow/observe posture.
 - **Florence-X** — an emerging research and build architecture exploring orchestration, verification, evidence, and visible human authorization.
 - **Nurse Intelligence Network (NIN)** — the professional community, summit, podcast, and collaboration network.
 - **Florence Media Network** — the public content and distribution program supporting the ecosystem.
 - **Hermes Agent** — a separate open-source desktop AI agent from Nous Research that can serve as an optional advanced runtime.
-

Accuracy guide

Preferred language

- “A browser-first, nurse-designed AI workflow and stewardship kit.”
- “Core setup and self-customization are free forever.”
- “Hermes is a separate optional desktop runtime.”
- “No patient information in intended workflows.”
- “EDENA provides review language in shadow/observe posture.”
- “Florence-X is an emerging research and build architecture.”
- “Agents propose. Humans judge. Nurses steward.”

Please do not say

- “HIPAA-compliant,” “HIPAA-certified,” or “safe for PHI.”
- “Clinical decision support,” “autonomous nurse,” or “diagnostic AI.”
- “A finished standalone operating system.”
- “Automatically enforces every rule across ChatGPT, Claude, and Hermes.”
- “Institutionally approved,” “clinically validated,” “certifies nurses,” or “improves patient outcomes.”
- “Hermes is Nurse AI OS” or “Nurse AI OS owns the underlying model/runtime.”
- “The free core ends after the founding year.”

Affiliation note

ChatGPT is a product of OpenAI. Claude is a product of Anthropic. Hermes Agent is developed by Nous Research. References identify compatible or optional hosts and do not imply endorsement, partnership, or ownership unless separately announced by the relevant organizations.

Suggested interview questions

1. What did 22 years in intensive care teach you about how AI should be governed?
2. Why make the first Nurse AI OS experience browser-first instead of requiring a new application?
3. What is a SOUL file, and what can it guide without guaranteeing enforcement?
4. Why is “no PHI” a participation rule rather than a compliance claim?
5. What would meaningful nurse stop authority look like when AI enters a hospital workflow?
6. How does EDENA distinguish low-risk assistance from work that requires review or refusal?
7. What is practical today, what is emerging, and what remains a long-term direction?

8. How can schools teach AI literacy while protecting academic integrity and student agency?
 9. What should a CNO or CNIO ask before adopting an AI workflow?
 10. Why do you believe nursing could become a custodial discipline for clinical intelligence?
-

Assets and links

- **Primary site and current product:** nurse-ai-os.org
 - **2-minute SOUL Quiz:** nurse-ai-os.org/soul-quiz.html
 - **Browser-first setup:** nurse-ai-os.org/start-here.html
 - **FAQ and current pricing:** nurse-ai-os.org/faq.html#pricing
 - **Founder and ecosystem:** nurse-ai-os.org/about.html
 - **Architecture Report:** nurse-ai-os.org/architecture-report.html
 - **Licensing and trademark:** nurse-ai-os.org/licensing.html
 - **Care Intelligence White Paper:** PDF
 - **Nurse Intelligence Network:** nurseintelligence.com
 - **YouTube:** [@NurseIntelligence](https://www.youtube.com/@NurseIntelligence)
 - **Founder image:** available at nurse-ai-os.org/about.html or in higher resolution on request.
 - **Media, interviews, speaking, and collaboration:** robert@nurse-ai-os.org
-

Licensing and trademark note

Nurse AI OS™ is identified publicly as a trademark of Robert Domondon; the registration application is pending, and no federal-registration claim is made. Original project components use the licenses published with the repository. Trademark rights, copyright, upstream component licenses, criticism, and permitted open-source use remain distinct. Review the current terms at nurse-ai-os.org/licensing.html.

Carry the lamp. Keep the ledger. Agents propose. Humans judge. Nurses steward.
