

STEWARD

Governance Specification

Hospital & Clinic Administration Preview

PUBLIC GOVERNANCE PREVIEW

Non-executable · Not institution-approved · No operational data

This document is a governance specification, not software. It contains no installer, executable agent, production connector, authority-verification service, pre-model data-loss-prevention control, or institution-ready runtime.

NAIO Institute

Version 2026.07.16-preview.1

Documentation review only · No runtime tier assigned

Agents propose. Humans judge. Nurses steward.

Contents

Executive decision

1. Purpose
2. Publication state
3. One lane, explicit context variants
 - 3.1 Proposed lane identity
 - 3.2 Why one lane
 - 3.3 Why context variants are mandatory
4. Human primacy doctrine
5. Proposed control frameworks
 - 5.1 CHART — authority before preparation
 - 5.2 ALIGN — human-owned organizational coordination
 - 5.3 ORBIT — bounded agent governance
6. Data and privacy doctrine
 - 6.1 Public-preview rule
 - 6.2 Future ingress requirement
 - 6.3 Aggregate data is not automatically anonymous
 - 6.4 Whole-life separation
7. Workforce dignity and anti-surveillance doctrine
8. Consequential-action prohibition
9. Proposed stewardship domains
 - S — Set strategy and steward mission
 - T — Transform access, flow, and operations
 - E — Steward economics and resources
 - W — Strengthen workforce and workplace
 - A — Assure quality, compliance, and equity
 - R — Build resilience and readiness
 - D — Govern digital systems and agents
10. Proposed agent research subjects
11. Evidence standard
 - 11.1 What has been verified
 - 11.2 What has not been verified
 - 11.3 Claim-versus-proof rule
12. Required path to an enforceable reference implementation
13. Known, assumed, unknown, recommended, decide
 - Known
 - Assumed
 - Unknown
 - Recommended
 - Decide
14. Stewardship conclusion

Executive decision

NAIO is keeping the STEWARD doctrine and its one-lane/context-variant architecture while pausing the claim that STEWARD is an installable Complete AI OS.

The supplied concept established a serious administrative governance model. Independent review also found a decisive gap between **controls described in prose** and **controls enforced and evidenced by software**. Publication therefore proceeds only as a non-executable governance preview.

This posture protects four truths:

1. Administrative AI deserves a dedicated governance architecture.
 2. Hospital and clinic administrators can share one population model without sharing one universal authority model.
 3. Written safeguards are useful design requirements, but they are not technical enforcement.
 4. Institutional use must wait for authenticated authority, enforceable data boundaries, runtime controls, and completed evidence.
-

1. Purpose

STEWARD is a proposed governance architecture for AI-assisted administrative preparation in hospitals, health systems, clinics, ambulatory organizations, practices, service lines, programs, departments, and bounded projects.

Its purpose is to help future systems:

- organize source-linked evidence;
- make assumptions and uncertainty visible;
- prepare nonbinding options and questions;
- preserve human decision ownership;
- prevent role, data, and authority collapse;
- distinguish private preparation from institution-approved work;
- govern proposed agents, tools, permissions, and handoffs;
- protect patients, workers, clinicians, communities, and organizations from hidden automation.

STEWARD does not define or confer employment, clinical, financial, legal, compliance, credentialing, procurement, emergency, board, executive, or institutional authority.

2. Publication state

State	Current status	Meaning
Doctrine documented	Available	The governance principles can be inspected and critiqued.
Public preview	Available	This non-executable specification and gap register may be downloaded.
Installer	Not published	No installation or activation pathway is offered.
Runtime adapter	Not implemented	No enforceable Hermes/Florence-X adapter is claimed.
Agents	Not shipped	Agent concepts are design subjects only.
Connectors/actions	Not shipped	No EHR, HRIS, scheduling, finance, procurement, compliance, incident, messaging, or other production connector is included.
PHI or sensitive-data use	Prohibited	The preview must not receive organizational or personal data.
Institutional authorization	Not granted	Publication creates no permission, approval, validation, or readiness.
Complete AI OS claim	Paused	The designation may return only after enforcement and evidence gates pass.

3. One lane, explicit context variants

3.1 Proposed lane identity

The future architecture reserves:

- population: `hospital_clinic_administration`;
- namespace: `hcadmin_steward.*`;
- conceptual home: **My STEWARD**;
- public route: `/hospital-clinic-administrators/`.

These identifiers describe a proposed architecture. They do not create a profile, workspace, schema, permission, agent, or runtime.

3.2 Why one lane

Hospital and clinic administrators share recurring governance needs:

- mandate and authority;

- access, flow, and operations;
- resource stewardship;
- workforce dignity;
- quality, compliance, and equity;
- resilience and readiness;
- digital-system and agent governance.

Those shared needs justify one administrative population model.

3.3 Why context variants are mandatory

The following are not interchangeable:

- health-system or enterprise executive;
- hospital administrator;
- clinic, ambulatory, or practice manager;
- multi-site operator;
- service-line leader;
- department or program administrator;
- bounded project lead;
- dual-hat clinician-administrator;
- analyst, advisor, consultant, trainee, or observer;
- unknown or unverified authority.

A future implementation must represent, verify, and expire at least:

- exact title;
- legal entity;
- ownership and governance model;
- site and service;
- active operational hat;
- written delegation source;
- decision and dollar thresholds;
- committee or board role;
- human functional owners;
- source version and effective period;
- conflicts, restrictions, and expiry.

A title is context, not permission. Access is not authority. Clinical licensure is not administrative delegation. Employment is not approval. Technical capability is not organizational mandate.

4. Human primacy doctrine

STEWARD follows a categorical rule:

AI may prepare. Authorized humans judge, approve, act, record, and remain accountable.

A future STEWARD runtime must never become or impersonate:

- an executive or board;
- a medical staff or credentialing authority;
- an HR, labor-relations, or employee-relations authority;
- a finance, revenue-cycle, contracting, or procurement authority;
- a quality, safety, privacy, security, legal, compliance, or accreditation authority;
- an emergency, incident-command, or continuity authority;
- a clinician, clinical decision-maker, or device controller;
- an official system of record.

AI delivery must not count as acknowledgment, approval, responsibility transfer, attestation, official reconciliation, policy enactment, incident closure, or evidence that an authorized human acted.

5. Proposed control frameworks

The frameworks below are design requirements. This preview does not execute them.

5.1 CHART — authority before preparation

A future CHART gate should evaluate:

1. **Context and consequence** — What is being prepared, for whom, and with what possible effect?
2. **Human authority** — Which authenticated person owns the decision, and what delegation proves it?
3. **Authorization and access** — What data, tool, destination, purpose, and time period are allowed?
4. **Regulation, risk, resources, and records** — Which legal, clinical, workforce, financial, privacy, security, equity, and official-record boundaries apply?
5. **Trace, transition, expiry, and termination** — What evidence, handoff, stop condition, rollback, purge, reconciliation, and expiry are required?

Conceptual results:

- **Pass for preparation only;**
- **Question;**
- **Block;**
- **Emergency Stop & Official Route.**

A CHART pass must never confer the underlying authority.

5.2 ALIGN — human-owned organizational coordination

A future ALIGN contract should make visible:

- aim and affected people;
- accountable decision owner;
- involved functions and official systems;
- dependencies and conflicts;
- required governance gates;
- named human receiver;
- official recording location;
- correction, appeal, escalation, expiry, and closure paths.

ALIGN must not create a shadow command, task, meeting-minutes, incident, or policy system. It may record only a human-designated owner or status copied from an authorized official source.

5.3 ORBIT — bounded agent governance

A future ORBIT contract should bind:

- objective and accountable human owner;
- role, risk, and prohibited functions;
- allowed data classes;
- typed tools, destinations, reads, and writes;
- permission budget and duration;
- test and monitoring evidence;
- human review and release gates;
- disagreement and failure handling;
- pause, kill, rollback, purge, retirement, and reconciliation evidence.

No agent may self-approve, expand its permissions, recursively delegate, hide retries, persist beyond authorization, make a consequential decision, or write to a production system under this preview.

6. Data and privacy doctrine

6.1 Public-preview rule

This preview accepts **no operational data**. Do not submit:

- PHI or patient-level information;
- patient, employee, clinician, applicant, learner, or practitioner narratives;
- employee, payroll, scheduling, evaluation, complaint, accommodation, credentialing, or privileging records;

- incident, claim, grievance, investigation, legal, or privileged material;
- confidential bids, contracts, pricing, vendor information, or nonpublic financial records;
- security-sensitive details, screenshots, recordings, exports, credentials, tokens, or secrets.

A prompt-level refusal is not a privacy control. If prohibited content reaches a conversational model for rejection, disclosure may already have occurred.

6.2 Future ingress requirement

Any future runtime must require validated **pre-model** controls appropriate to the deployment environment.

Required evidence includes:

- documented data classification;
- deterministic and tested deny rules;
- pre-ingress scanning or approved equivalent;
- local/institutional routing boundaries;
- false-negative and false-positive evaluation;
- blocked-input receipts that do not retain prohibited content;
- provider-retention and logging contracts;
- incident response and human escalation.

6.3 Aggregate data is not automatically anonymous

Role group, site, shift, specialty, overtime, vacancy, competency, experience, and time can re-identify people in small clinics or rare teams.

A future aggregate-data adapter must enforce, not merely request:

- minimum cohort sizes;
- small-cell suppression;
- complementary suppression;
- differencing and longitudinal-query controls;
- rare-category handling;
- purpose limitation;
- access and retention limits;
- re-identification testing;
- correction, appeal, and labor/workforce review where applicable.

No universal cohort threshold is declared in this preview; the threshold must be governed, justified, and tested for the exact context.

6.4 Whole-life separation

Personal whole-life information must never share a record, partition, authority chain, or institution-managed workspace with succession, delegation, sponsorship, performance, workforce planning, or leadership evaluation.

A future design must separate:

- personally controlled whole-life support;
- professional development;
- institution-governed delegation;
- succession and workforce planning.

Hiding a field or panel is not technical separation.

7. Workforce dignity and anti-surveillance doctrine

STEWARD must not rank, risk-score, profile, monitor, or predict:

- employees or applicants;
- clinicians or practitioners;
- teams, departments, sites, or shifts;
- patients or communities;
- vendors or partners.

It must not infer:

- sentiment;
- burnout or impairment;
- intent or misconduct;
- protected characteristics;
- retaliation risk;
- competence, credential status, or fitness;
- productivity or individual performance;
- staffing adequacy or assignment eligibility;
- cause, blame, reportability, or legal exposure.

AI must not recommend or execute:

- hiring, promotion, assignment, scheduling, staffing, compensation, evaluation, accommodation, remediation, discipline, or termination;
- credentialing or privileging;
- investigation or reportability decisions;
- workforce communications about an identifiable person.

Psychological-safety and listening work is especially sensitive. A future implementation would require voluntary participation, purpose limitation, anti-retaliation governance, labor/workforce review, strict aggregation, access controls, short retention, and proof that repeated trend analysis cannot become management surveillance.

8. Consequential-action prohibition

The future base architecture must use one categorical prohibition model. Human approval must not convert a prohibited function into an allowed one.

STEWARD must not perform or stage:

- clinical or device actions;
- patient contact or prioritization;
- staffing or scheduling changes;
- employment or credentialing actions;
- payments, pricing, coding, billing, claims, procurement, contracting, or fund movement;
- legal, compliance, accreditation, audit, or regulatory conclusions;
- policy publication or enactment;
- security changes;
- emergency command;
- incident closure;
- official-system writes or official records.

A future adapter may prepare a nonbinding artifact only when the exact purpose, data, destination, human owner, and prohibition boundary are technically enforced. The preview itself prepares nothing.

9. Proposed stewardship domains

The original doctrine describes 24 proposed capability areas under seven pillars. They are retained here as a **research and design agenda**, not as enabled powers.

S — Set strategy and steward mission

1. Strategic North Star and Operating Mandate
2. Executive Decision, Options, and Scenario Studio
3. Portfolio and Benefit-Realization Navigator
4. Board and Stakeholder Governance Navigator

T — Transform access, flow, and operations

5. Access, Demand, and Capacity Intelligence
6. Flow, Throughput, and Bottleneck Lab
7. Daily Management and Escalation Board
8. Hospital, Clinic, and Site Operating-Model Designer

E — Steward economics and resources

9. Budget, Forecast, and Variance Studio
10. Revenue Cycle, Payer, and Contract Learning
11. Procurement, Vendor, and Sustainability Due Diligence

W — Strengthen workforce and workplace

12. Aggregate Workforce Capacity and Skill-Mix Evidence
13. Psychological Safety and Listening Themes
14. Change, Communication, and Adoption
15. Leadership, Succession, and Whole-Life Compass

These areas require redesign before implementation, particularly the separation of whole-life, leadership development, delegation, and succession.

A — Assure quality, compliance, and equity

16. Quality, Safety, and Accreditation Readiness
17. Policy, Control, and Audit-Evidence Studio
18. Equity, Accessibility, and Community Impact

R — Build resilience and readiness

19. Synthetic Emergency and Downtime Rehearsal
20. Cybersecurity, Privacy, and Third-Party Risk
21. Incident, Corrective-Action, and Learning Governance

The incident domain must be redesigned so the base architecture cannot ingest real incidents, event narratives, patient/workforce facts, privileged review material, or shadow records.

D — Govern digital systems and agents

22. Digital and AI Use-Case Portfolio Governor
23. Agent Charter, Permission, and Human-Handoff Designer
24. Agent Evaluation, Kill, Rollback, and Retirement

No capability above is implemented, active, installable, or available for operational use in this preview.

10. Proposed agent research subjects

The original doctrine named ten possible software-agent roles. This preview ships **zero agents**. The names are retained only to define future evaluation subjects:

1. Source and Requirement Verifier
2. Executive Brief and Scenario Drafter

3. Portfolio and Dependency Mapper
4. Access and Flow Aggregate Analyst
5. Financial Scenario Analyst
6. Workforce Evidence Analyst
7. Quality, Policy, and Control Crosswalk Assistant
8. Change, Accessibility, and Communications Assistant
9. Resilience Exercise Facilitator
10. Independent Agent Auditor and Kill Sentinel

Before any agent can be implemented, it needs executable schemas, typed capabilities, deny-by-default tools and destinations, authenticated ownership, independent approval where required, adversarial tests, kill/rollback evidence, and an expiry/retirement contract.

11. Evidence standard

11.1 What has been verified

NAIO's prepublication review checked the supplied source archive for:

- exact archive bytes and SHA-256;
- CRC integrity;
- path traversal and unsafe member types;
- duplicate and case-colliding names;
- file-permission anomalies;
- complete internal checksum coverage;
- active Office content;
- source inventories and identifier ranges;
- 17-component one-file/source parity;
- 24 powers, 24 workflows, 30 templates, 18 schema declarations, and 10 agent declarations;
- 160 unique declared release criteria.

These are static source-integrity findings only. The original source archive is intentionally excluded from the public preview, so these audit results cannot be independently reproduced from the public bundle alone. They are NAIO provenance statements, not third-party certification, runtime validation, or institutional assurance.

11.2 What has not been verified

No evidence establishes:

- pre-model PHI or sensitive-data prevention;
- authenticated identity or authority;
- segregation of requester, approver, and executor;

- partition isolation;
- connector, tool, destination, or permission denial;
- provider retention or session-only memory;
- production-safe aggregation or re-identification resistance;
- pause, kill, rollback, purge, uninstall, or orphan cleanup;
- clinical, workforce, financial, legal, privacy, security, compliance, accreditation, or institutional effectiveness;
- completion of the 160 runtime criteria.

11.3 Claim-versus-proof rule

Every future claim must identify:

- claim owner;
- intended context;
- implementation artifact;
- test method;
- evidence location;
- independent reviewer;
- limitations;
- version and expiry.

Absent that evidence, the claim remains a requirement or hypothesis.

12. Required path to an enforceable reference implementation

The Complete AI OS designation remains paused until a future Florence-X/EDENA reference adapter provides and demonstrates:

1. versioned Hermes configuration and profile artifacts;
2. executable schemas and validation fixtures;
3. pre-model sensitive-data controls;
4. authenticated identity and authority binding;
5. signed or otherwise verifiable delegation evidence;
6. enforced segregation of duties;
7. typed capabilities, tools, destinations, and deny rules;
8. technical aggregation, suppression, and re-identification controls;
9. separate private, professional-development, delegation, and succession domains;
10. a categorical consequential-action prohibition;
11. tested retention, correction, export, deletion, purge, rollback, and uninstall behavior;
12. independent test evidence for every applicable runtime criterion;

13. an evidence-bearing release ledger with no unresolved blocker;
14. institutional legal, privacy, security, workforce, clinical, and operational review for the exact deployment.

Passing a reference implementation does not authorize another institution. Authorization remains local, scoped, revocable, and time-bound.

13. Known, assumed, unknown, recommended, decide

Known

- The supplied doctrine is internally substantial and uses one administrative lane with explicit context variants.
- Its source inventories and checksum chain were coherent.
- Its strongest safeguards are primarily prose requirements rather than implemented controls.
- The supplied release contract requires runtime evidence that is not present.

Assumed

- Administrators and governance teams may benefit from a common vocabulary for authority, data boundaries, organizational coordination, and agent oversight.
- A public preview can support expert critique without encouraging operational use if labeling and distribution remain unambiguous.

Unknown

- Which minimum technical adapter can enforce these requirements across different Hermes providers and institutional environments.
- Which aggregation thresholds and labor-governance controls are appropriate across large systems, small clinics, rural sites, and rare specialties.
- Which institutional buyers, counsel, privacy officers, workforce representatives, and clinicians would accept the proposed evidence model.

Recommended

- Keep the doctrine public and inspectable.
- Keep operational data and activation pathways out of the preview.
- Build the smallest enforceable Florence-X/EDENA adapter against synthetic fixtures.
- Use independent review before restoring any product or installation claim.

Decide

The future human governance body must decide whether and when implementation evidence is sufficient to restore the **Complete AI OS** designation. No agent may make that decision.

14. Stewardship conclusion

STEWARD is not being stopped. It is being placed in the correct maturity state.

The preview preserves the idea that healthcare administration needs an AI governance architecture shaped around authority, dignity, accountability, and organizational reality. It also preserves a harder truth: administrative safeguards matter only when they are enforceable, observable, and independently evidenced.

Agents propose. Humans judge. Nurses steward.