

STRATEGIC VISION + DECISION BRIEF

PSA / IHPF HEALTHCARE VISION 2.0

From Abundant Intelligence to Person-Directed Healthcare

A high-level reconciliation of Gary Gomes's 10-year vision with John Griffith/IHPF research, clearly attributable GPT-assisted work maintained in Kris Witte's healthcare archive, documented PSA lived experience, and authoritative external evidence current through 28 August 2026.

THREE DELIVERABLES

- 1 • One-page executive vision 2 • 10–15-page strategic reconciliation 3 • Decision appendix: explore, test, monitor, reject or defer

Prepared for PSA Board consideration

Evidence cutoff: 28 August 2026 | Prior platform research cutoff: 10 August 2026

DELIVERABLE 1**Executive Vision**

Destination. Within ten years, every New Mexican should be able to direct a portable, understandable plan for obtaining healthcare—not merely receive disconnected services. A persistent support layer should help the person and chosen caregivers understand what happened, decide what matters next, assign work, overcome coverage, scheduling, transport and pharmacy barriers, and verify that each important next step actually closed. It should connect—not replace—clinicians, hospitals, payers, pharmacies, community organizations, families and public systems. [G1][P1]

Core judgment. Gary Gomes’s vision is strategically strong and unusually clear: intelligence can become broadly available; essential services can be redesigned around individuals; geography may constrain fewer tasks; trusted community and accountable human relationships become more important; and AI may eventually recede from the user’s description of good service. Those propositions are best treated as a destination and design hypothesis—not as evidence that the destination will occur or that AI alone produces abundance. [G1]

What PSA/IHPF adds. Later work supplies the operating model Gary intentionally leaves open: a publicly accountable connective layer; navigation through verified closure; paid and supervised community health workers/promotoras; named human authority; bounded AI; consent and caregiver delegation; interoperability adapters; bilingual and offline access; safety queues; review-capacity measurement; braided payment; burden measures; independent evaluation; and explicit scale/stop gates. [P1][P2][P4][K1–K6]

Scarcity correction. AI can make first-pass summarization, translation, drafting, education, pattern finding and coordination suggestions far cheaper. It does not make accountable clinical judgment, trusted relationships, physical care, specialist appointments, clean interoperable data, review time, implementation capability, transportation or financing abundant. Indeed, abundant AI output can worsen the scarcest resource: competent human attention. A 2026 pragmatic randomized trial found better documentation with LLM assistance but no significant reduction in treatment failure; that is precisely why PSA should test completion and burden outcomes rather than infer them from capability. [E3][K4]

First strategy. For the next one to three years, PSA should pursue person-directed care continuity, not autonomous diagnosis. The first New Mexico pilot should be a staged, bilingual “Next Step Closure” service in Doña Ana County for a tightly defined Turquoise Care population after a visit, discharge or specialty referral. A paid promotora/CHW should conduct teach-back, translate the plan into assigned tasks, coordinate appointment/transport/pharmacy/community supports, escalate clinical questions to a named

professional, and verify completion. AI may draft, translate approved content, classify barriers and prioritize queues; a named human remains accountable for consequential output. [P1][P4]

Board posture. Authorize only a 90-day Stage 0: partner commitments, resident/caregiver co-design, baseline workflow and capacity measurement, clinical/privacy/economic protocols, and an evaluation plan. Require a payer/provider financing term sheet before build. Begin in manual-plus-shadow mode, then enroll 150–300 people; expand toward the previously proposed 800–1,200 only if closure improves without unacceptable safety events, inequitable performance, or added resident/caregiver/staff burden. [P1][K5]

DISPOSITION

Keep Gary’s paper unchanged as the signed/original vision essay and use it as the introductory vision section of a larger PSA/IHPF document. Adopt this Healthcare Vision 2.0 as the strategic and evidentiary reconciliation around it.

DELIVERABLE 2

Strategic Reconciliation

Gary + John/IHPF + Kris + external evidence | Decision-quality synthesis

1. Executive assessment of Gary's vision

Gary's paper is a compelling 10-year vision and a useful provocation for PSA. It reframes AI from a tool inserted into today's institutions into a reason to reconsider inherited service boundaries. Its most durable contribution is not a prediction about model capability; it is the proposition that services should organize around people, with technology extending skilled humans and leaving judgment, empathy, relationships, physical presence and accountability unmistakably human. [G1 | VISION]

The paper is also incomplete by design. It does not specify an accountable operating entity, clinical boundaries, data rights, payment flows, review capacity, interoperability strategy, implementation burden, empirical milestones or stop conditions. Those omissions are not fatal in a vision essay. They become fatal if the essay is treated as a pilot blueprint. Later PSA/IHPF work has gone substantially further on nearly every one of these operating questions.

[P1][P2][K1–K6 | MODEL-GENERATED SYNTHESIS]

The correct reconciliation is therefore additive. Gary supplies the destination, memorable language and economic test. John/IHPF supplies the person-directed and connective-layer architecture, grounded in the difference between healthcare services and the work of obtaining them. Clearly attributable documents maintained in Kris Witte's GPT healthcare archive supply unusually concrete controls for role definition, supervision, review bottlenecks, semantic meaning, integrity checks and stop rules. External evidence makes the opportunity plausible but sharply qualifies any claim that capability automatically improves outcomes. [L2][E1–E9]

BOTTOM LINE

Approve the vision; qualify the abundance thesis; operationalize the connective layer; test one closed-loop workflow; and refuse to scale on output, engagement or enthusiasm alone.

2. Evidence discipline: what kind of statement is this?

Label	Meaning	How this report uses it
VISION	Normative destination or design aspiration	Gary's 10-year propositions; not a forecast or outcome claim.

Label	Meaning	How this report uses it
HYPOTHESIS	A testable causal or operating proposition	Example: verified closure will improve completed next steps without increasing burden.
SYNTHESIS	Model-generated or GPT-assisted analysis	Useful for conceptual triangulation; never counted as independent empirical validation.
LIVED	Documented first-person experience	Design evidence about burdens, sequences and failure modes; not population proof.
AUTH	Authoritative primary policy, regulatory or agency evidence	Defines rules, programs, standards, funding and official data.
EMPIRICAL	Measured outcome evidence	Trials or evaluated implementations; confidence depends on design and transferability.
FORECAST	Inference about future feasibility or adoption	Explicitly uncertain and monitored, not presented as fact.

Evidence rule: agreement among Gary’s paper, PSA GPT syntheses and Kris-archive GPT documents is conceptual triangulation. Confidence rises only when independent primary evidence, implementation experience or measured outcomes support the proposition.

3. Gary’s strongest ideas—and why they matter

Gary contribution	Strategic value	Qualification
Scarcity → abundance	Directs attention from automating queues to redesigning the service model.	Abundance applies to selected informational tasks, not care capacity as a whole.
Organize around individuals	Names the governing direction: institutions become resources in a person’s plan.	“Person-directed” must not become “person abandoned.”
Episodic → continuous support	Matches chronic care, recovery and navigation realities.	Continuous support must be consented, proportionate and not continuous surveillance.
Geography matters less; community more	Combines global knowledge with local trust and presence.	Physical care, transport, broadband and specialist capacity remain geographic.
Humans become more important	Prevents substitution of relationship and accountability with fluency.	Roles, authority, staffing, supervision and payment must be explicit.

Gary contribution	Strategic value	Qualification
Economics is inseparable	Forces “who pays, saves, loses and captures?” before scale.	Savings will not accrue automatically to the entity doing the work.
Institutions change, not vanish	Avoids a false disruption-versus-status-quo choice.	Incumbents may resist revenue loss, data sharing and redistributed authority.
AI disappears from the description	Good services should feel coherent, not like technology demonstrations.	AI may fade in UX but never in provenance, audit, disclosure or accountability.

4. Theme-by-theme finding

Theme	Finding	Why / PSA implication
Useful intelligence becomes cheap	Qualified agreement	Summaries, translation, drafting and suggestions scale; validated judgment and review do not. [E3][E9]
Essential-service capacity becomes abundant	Disagreement with the strong form	Information may scale; accountable judgment, physical care, trusted relationships, review, interoperability and financing remain scarce. [E3][E5]
Redesign, do not merely automate	Strong agreement	Automation can accelerate broken handoffs. Start with task ownership and closure. [G1][P1]
Services around the person	Strong agreement	Make the person principal; institutions and tools serve a portable plan. [G1][P3]
Persistent intelligent assistant	Unproven	Potentially valuable, but agency, consent, memory, error recovery and liability remain unresolved.
Continuous, personalized support	Qualified agreement	Useful for follow-through; guard against surveillance, alert fatigue and digital exclusion.
Geography matters less	Qualified agreement	Information travels; people, specimens, procedures and specialist slots often do not. [E5]

Theme	Finding	Why / PSA implication
Community matters more	Qualified agreement	Community must be plural, chosen, culturally safe and paid—not assumed volunteer capacity. [L2]
Human work shifts upward	Unproven at system level	Some tasks shift; review, exception handling and emotional labor may expand. Measure net burden. [K4][E3]
Earlier intervention improves outcomes	Unproven as a general AI claim	Plausible by use case; requires prospective outcome evidence, not model performance.
Institutions evolve	Strong agreement	The connective layer should federate across institutions, not become another silo. [P1]
Outcome-aligned economics	Qualified agreement	CMS ACCESS shows a payment mechanism is possible; it is not a funding answer for a Turquoise Care pilot. [E1]
AI disappears	Qualified agreement	Ambient at the surface; conspicuous in governance, provenance and audit.
Global intelligence + local humans	Strong agreement	Consistent with paid CHW/promotora models and rural hub-and-spoke designs. [E4][P4]
Professional productivity rises	Unproven/variable	Documentation may improve without clinical outcome gains. [E3]
A 10-year system transformation	Forecast	Directionally plausible; timing and institutional response cannot be inferred from capability.

Finding categories answer the requested comparison: strong agreement; agreement with important qualification; disagreement; unproven; and areas where later PSA/IHPF work has gone substantially further. No core Gary theme is rejected outright; several strong forms are rejected below.

5. Where PSA/IHPF has gone substantially further

Later PSA/IHPF work converts a person-centered vision into a care-continuity operating model. Its central object is not a chatbot, record repository or directory. It is a shared task-and-closure engine: each important next step has an owner, due date, status, escalation path and verified completion state. Resident, caregiver, CHW/navigator, clinical supervisor and governance users see different surfaces over the same governed workflow. [P1 | SYNTHESIS]

Missing or thin in Gary	PSA/IHPF development	Design consequence
Patient/consumer direction	Person is principal; caregiver proxies are permissioned; support does not equal abandonment. [P3]	Portable goals, choices, consent and correction—not merely personalization.
Connective layer	Federated civic continuity layer across visits, homes and community services. [P1]	Do not build another EHR, HIE or standalone app.
Navigation + verified closure	Completion, not referral sent or information delivered, is the unit of value. [P1][K6]	Capacity confirmation, exception queues and closure evidence are required.
Human accountability	Named human approves consequential output and owns escalation. [K1][K2]	“Human in the loop” becomes an authority map, service level and funded role.
Bounded AI	Narrow tasks, source-controlled content, logs, audits, change control and kill switch. [K2][K5]	No diagnosis, prescribing, test interpretation or unrestricted agent action.
CHWs/promotoras	Paid, trained, supervised meaning-and-trust layer. [P2][P4][E4]	Budget labor, supervision, clinical integration and emotional workload.
Rural access	Bilingual/offline channels, transport, hub-spoke-pop-up delivery. [P2][P4]	Digital reach supplements local capacity; it does not erase distance.
Interoperability	Adapters to SYNCRONYS, payer and provider APIs; semantic closure definitions. [P1][K6]	Use standards and minimum data; do not wait for perfect integration.
Payment/reimbursement	Braided funding and pre-negotiated value capture. [P1][E1]	No scale without a payer/provider contract and workforce funding.
Safety/privacy/governance	Clinical escalation, incident review, purpose limits, audit, revocation and vendor exit. [P1][K2]	Governance is a product function, not a policy appendix.
Review/workforce capacity	Measure reviewer minutes, queue age, caseload and authority. [K4]	AI volume is capped by safe review capacity.
Burden + evaluation	Resident, caregiver, staff and implementation burden are outcomes; hard stop gates. [K5]	Scale only when benefit is net of shifted work and risk.

6. What Gary adds that PSA/IHPF should preserve

- 1. A memorable 10-year destination that prevents PSA from reducing the project to a workflow tool.
- 2. A selective abundance lens: ask which inherited structures exist only because knowledge and coordination were scarce.
- 3. The proposition that organizations—not people—should do the adapting.
- 4. A division-of-labor vision in which human judgment, relationships and accountability become more central, not less.
- 5. The geography/community paradox: globally available intelligence can increase the value of locally trusted people.
- 6. An explicit political-economy test: redesign fails if the party that pays cannot capture value, or if threatened revenue blocks adoption.
- 7. A realistic role for institutions: transformed and recombined, not wished away.
- 8. The aspiration that AI disappears from ordinary service language once it is reliable—while remaining visible to governance.

7. Lived experience: the process of obtaining healthcare

Documented PSA accounts reveal a system burden that conventional service inventories miss. A service may exist, a clinician may perform appropriately, and every institutional component may be paid—yet the person still lacks a coherent answer to “what should I do next, who owns it, and how will I know it happened?” John Griffith’s formulation is decisive: there is healthcare, and there is the process of getting healthcare. [L2 | LIVED]

Documented experience	Observed burden / failure	Design signal—not proof
Gary: subdural hematoma, June 2026	Weeks of symptoms; rural distance; hospital transfer and bed wait; discharge dependencies; large record volume; driving/walking limits shifted work to Karen. [L1]	Emergency choice support, transfer status, discharge dependency map, caregiver capacity check, record synthesis and closure.

Documented experience	Observed burden / failure	Design signal—not proof
Gary: cardiology/testing, August 2026	Prior false positive shaped refusal; alternative CT caused severe reaction and failed intended measurement; incidental findings, fragmented interpretation, clinician departure after acquisition, disengagement risk. [L2]	Preference/history memory, test intent and result reconciliation, adverse-experience capture, named interpretation owner, clinician continuity and “stop the rabbit hole” option.
John: heart attack in 2011	Alone, delayed emergency decision, avoided ambulance, drove himself and improvised aspirin. [L1]	Pre-authored emergency plan and simple decision tree; technology must never normalize delay or self-transport in emergencies.
Margot: chronic back pain	Limited relief, repeated care and self-directed research created work without a coherent pathway. [L1]	Goal-based plan, treatment history, response tracking and explicit reassessment/stop criteria.
Karen: caregiver role	Mobility/back limitations collided with Gary’s temporary dependence and household obligations. [L1]	Caregiver burden and capacity must be measured before assigning tasks; support cannot assume an able family proxy.

No sufficiently clear, documented personal healthcare experience from Kris Witte was located in the accessible project material. Kris’s contribution is therefore represented only through GPT-assisted strategy documents whose provenance is clear from the named project archive; the report does not infer personal views or lived experience. [PROVENANCE NOTE]

8. Gary’s recent experience as a design case

Gary’s recent experience is an excellent case for journey mapping and a poor basis for claiming population effectiveness. It is useful because it spans emergency recognition, rural routing, interfacility transfer, procedures, discharge constraints, caregiver capacity, follow-up, diagnostic history, testing risk, incidental findings, clinician turnover and patient disengagement. It also shows that high education, willingness to research, financial flexibility and AI access do not eliminate system complexity. [L1][L2 | LIVED]

CASE-DERIVED PROTOTYPE QUESTION

Could a person-directed support layer maintain Gary’s goals and history, explain the intended decision at each step, surface prior false positives and reactions, identify the accountable human, assign logistics, include Karen’s capacity, and verify closure—without diagnosing, directing emergency delay, or creating another queue?

The case supports design requirements: persistent but permissioned memory; goal and preference capture; a distinction between test ordered, test completed, intended measurement obtained and result interpreted; named responsibility; caregiver delegation; “what would change the plan?” prompts; explicit escalation; and a person-controlled option to pause low-value cascades. The case does not show that such a layer will reduce utilization, improve outcomes or save money. Those remain hypotheses for a prospective pilot. [HYPOTHESIS]

9. Scarcity to abundance—reconciled

AI can make more abundant	Likely to remain scarce
First-pass summaries and visit-plan drafts	Accountable clinical judgment and liability-bearing decisions
Approved-language translation and reading-level adaptation	Trusted relationships, cultural legitimacy and conflict repair
Navigation knowledge and personalized education	Physical examination, procedures, home care and transport
Task extraction, reminders and barrier classification	Specialist slots, beds and appointment capacity
Pattern detection and queue prioritization	Competent review time and exception handling
Simulation, training and protocol rehearsal	Implementation leadership, workflow redesign and change capacity
Monitoring signals and administrative drafting	Clean, semantically aligned, permissioned interoperable data
Low-cost software replication	Sustainable financing and fair value capture
Potentially continuous informational support	Caregiver time, patient attention and freedom from surveillance

The decisive qualification is that abundant outputs can intensify scarcity. If a model drafts ten times more messages, flags or plans, safe review may become the bottleneck. In the 2026 Kenya cluster-randomized trial, LLM support improved documentation quality but did not significantly reduce 14-day treatment failure; 4.2% of reviewed high-severity outputs were

rated somewhat unsafe/inappropriate or unsafe/inappropriate. Transfer to New Mexico is limited, but the lesson is direct: cheap generation is not the same as outcome-producing capacity. [E3 | EMPIRICAL]

Thus PSA should define “abundance” operationally: a resource is abundant only when a person can obtain it at low marginal cost, with acceptable quality, without creating uncompensated work or safety risk elsewhere. By that definition, trustworthy care coordination is not yet abundant. [MODEL INFERENCE]

10. Person-directed healthcare—not consumer abandonment

Gary’s proposition that services should organize around individuals is strongly supported as a governing principle. PSA’s earlier “consumer-directed” synthesis adds ownership, portability, transparent choices and an alternative front door. Yet “consumer” can imply a marketplace transaction, purchasing power or personal responsibility for overcoming structural failures. Gary’s concern that the phrase carries a political history is well taken. “Person-directed” is more accurate: the person defines goals, permissions, preferences and acceptable burden; professionals retain duties; institutions retain obligations; and the connective layer does the organizing work. [G1][P3][L2]

Person direction does not require each person to become a clinical expert or project manager. It means that no institution may treat its local workflow as the whole plan. The support layer should show the person one coherent set of next steps, make responsibility visible, accept any channel, involve chosen caregivers, and provide a human route when technology fails. The person may delegate extensively, decline an AI function, use paper or phone, or ask a navigator to act. [P1 | DESIGN PRINCIPLE]

11. Geography should matter less; community should matter more

The proposition is directionally right and easily overstated. Geography should matter less for access to knowledge, translation, pre-visit preparation, monitoring review, benefits navigation and some specialist consultation. It will continue to matter for emergency response, hands-on diagnosis, imaging, procedures, specimens, rehabilitation, pharmacy access, transportation, broadband reliability and the location of scarce clinicians. HRSA’s live shortage-area data remain the clearest warning that digital intelligence does not create licensed capacity. [E5 | AUTH]

New Mexico’s FY26 Rural Health Transformation award—\$211.5 million—and the state’s proposed initiatives for rural access, workforce, community-led innovation, provider sustainability and a data hub create a timely partnership environment, but funds and initiatives are not proof of operational capacity or a commitment to PSA. [E6 | AUTH]

“Community” may mean	Design implication
Geographic place: village, colonia, county, frontier region	Local delivery, transport, broadband and capacity mapping matter.
Tribal nation or pueblo	Sovereignty, data authority and governance cannot be collapsed into state rules.
Cultural or language community	Bilingual text is insufficient; meaning, trust and norms need local stewards.
Family or chosen caregivers	Permission, delegation and burden/capacity checks are required.
Socioeconomic circumstances	Time, devices, transport, housing and money shape what “choice” is real.
Clinical or condition community	Peer support may help, but should be opt-in and evidence bounded.
Digital/affinity network	Useful for mobile people and rare conditions; may not provide physical presence.
Individual preference	Some people want little community involvement; person direction includes privacy and refusal.

The pilot should therefore ask each participant which people and communities are useful to them. It should not assign an idealized “community” or treat family/volunteer labor as a free substitute for funded services. [L2][K3 | LIVED + SYNTHESIS]

12. Economics: who pays, benefits, saves, loses and captures?

Question	Likely answer for a first pilot	Board implication
Who pays?	Turquoise Care MCO administrative or value-based funds; provider care-management resources; eligible Medicaid CHW services; RHT/grant capital; PSA/philanthropy for independent evaluation.	Do not rely on patient fees or a grant-only operating model.
Who benefits?	Residents/caregivers through less work and more completion; clinicians through cleaner handoffs; providers/payers through fewer failures; CHWs/CBOs through paid roles.	Measure each party’s benefit and burden separately.

Question	Likely answer for a first pilot	Board implication
Who saves?	Possibly payer/provider from avoided low-value utilization and admin rework; person from time/travel; caregiver from fewer crises.	Savings are hypotheses; specify baseline, time horizon and counterfactual.
Who loses revenue?	Potentially facilities, testing entities, vendors or intermediaries paid by fragmented volume; some departments lose local control.	Map resistance and transition, not merely total savings.
Who captures value?	Often the payer, integrated provider or platform vendor—not automatically the community operator doing closure work.	Negotiate attribution, payment and data access before launch.
What supports redesign?	Outcome-based payment, access standards, payer APIs, shared-savings opportunity, rural funding, clinician burden reduction. [E1][E6–E8]	Use policy tailwinds as infrastructure, not as evidence of impact.
What resists it?	Fee-for-service volume, fragmented budgets, weak data, procurement cycles, liability, workforce shortages, review cost and vendor lock-in.	Require modular procurement, exit rights and a no-build financing gate.

CMS ACCESS is important external validation of feasibility, not of PSA's exact model. It pays recurring, outcome-aligned amounts for clinician-guided, technology-supported chronic care in Original Medicare and explicitly supports PCP co-management. It does not directly finance a Turquoise Care adult pilot. PSA should explore ACCESS-aligned partners for a parallel Medicare pathway while obtaining an independent Medicaid/MCO financing agreement for Doña Ana. [E1 | AUTH]

13. Should AI disappear?

Yes at the level of ordinary experience; no at the level of rights and accountability. A mature service should be described as “I understood my plan, someone helped me complete it, and the right professional responded when needed,” not as an AI product. That is Gary's strongest usability insight. But opacity is unacceptable when AI affects prioritization, translation, summaries, escalation or decisions. AHRQ emphasizes that realized benefits depend on adoption, calibrated trust, monitoring and governance; over-trust can create harm, while exhaustive review can erase efficiency. [G1][E9 | VISION + AUTH]

May recede from the user's foreground	Must remain visible and inspectable
--	--

May recede from the user's foreground	Must remain visible and inspectable
Branding, model names and novelty language	Notice that AI was used in a consequential workflow
Routine drafting and low-risk automation	Source/provenance, model/version, prompt/policy and output log
Technical architecture details	Named accountable human and appeal/correction route
Repeated consent screens for unchanged low-risk functions	Material changes, new purposes and consent/opt-out choices
Developer-centric "AI features"	Performance by language, geography, disability and subgroup
Invisible background plumbing	Incidents, overrides, review time, audit results and kill-switch authority

RECOMMENDED LANGUAGE

AI should become ambient, not invisible: quiet in the service experience; explicit in provenance, governance, auditability and accountability.

14. August 2026 platform research: comparison and judgment

The 10 August Personal + Community Health Support Platform research is the most complete PSA/IHPF operating synthesis. Its recommended "publicly accountable connective layer" translates Gary's person-organized services into a bounded first strategy: bilingual post-visit and referral closure, not autonomous diagnosis. It specifies user surfaces, task/closure semantics, CHW workflows, human review, governance, architecture, financing, evaluation and stop gates. [P1 | SYNTHESIS]

Dimension	Gary's 10-year vision	August platform strategy	Vision 2.0 reconciliation
Unit of design	Individual across institutions	Next step with owner, due date, escalation and closure	Person-directed portable plan composed of verified next steps
Role of AI	Abundant intelligence; largely disappears	Draft/translate/classify/prioritize within boundaries	Ambient support; visible provenance and human authority
Human layer	Trusted local professionals and caregivers	Paid CHW/promotora + clinical/supervisor queues	Funded meaning, navigation and accountability layer

Dimension	Gary's 10-year vision	August platform strategy	Vision 2.0 reconciliation
Institutional stance	Change, not disappear	Federate; adapt through APIs; do not replace EHR/HIE	Connect and hold seams accountable
First move	Not specified	Doña Ana post-visit/referral closure, 800–1,200	Stage 0 + 150–300 launch; expand only through gates
Evidence	Vision/hypothesis	Completion, safety, equity and burden metrics	Add external comparator, independent evaluation and financing gate

The report affirms the August strategic choice but modifies its launch sequence. An 800–1,200-person pilot may be an appropriate evaluated scale after workflow validation. It is too large for the first operational exposure if partner response capacity, baseline closure, review time and payment are unknown. Vision 2.0 therefore adds a 90-day Stage 0, manual-plus-AI-shadow testing and a 150–300-person initial cohort. Coverage continuity remains a valuable module and monitored alternative, but running two “first” pilots would dilute learning. [MODEL INFERENCE]

15. What has changed since the 10 August cutoff?

Development	Effect on Gary / PSA	Pilot consequence
FDA GenAI medical-device discussion paper, 18 Aug 2026	Strengthens long-term feasibility while confirming unique risks, competency assessment, clinical confirmation and postmarket monitoring for GenAI-enabled devices. [E2]	Keep first use case outside diagnosis/treatment; document intended use; consider an FDA docket comment by 19 Oct.
CMS ACCESS rolling implementation, including 17 Aug start	Makes outcome-aligned payment and paid PCP co-management more concrete for Original Medicare. [E1]	Explore ACCESS participant partnership; do not misstate it as Turquoise Care funding.
HRSA shortage data refreshed through 28 Aug	Reaffirms persistent workforce scarcity; no evidence that digital capability erased rural capacity. [E5]	Measure responder capacity and escalation delays before enrollment.

Development	Effect on Gary / PSA	Pilot consequence
Current randomized LLM evidence	A large pragmatic trial found better documentation but no significant treatment-failure benefit. [E3]	Power the pilot for completion and burden; avoid clinical outcome promises in Stage 1.
FHIR/USCDI and payer API trajectory	Standards and 2027 payer API requirements make federated data access more feasible, though implementation is uneven. [E7][E8]	Build adapters and minimum viable data now; do not make perfect interop a prerequisite.
NM Rural Health Transformation implementation environment	Creates plausible partners and capital for rural workforce, remote care, community-led innovation and data infrastructure. [E6]	Seek alignment; do not assume award dollars or duplicate the state data hub.

No authoritative post-cutoff development justifies changing the first use case to autonomous clinical AI, device-first diagnostics or a statewide platform. The net change is greater feasibility for infrastructure and payment, plus greater clarity about regulatory and review obligations. [MODEL INFERENCE]

16. Integrated PSA/IHPF framework

Layer	Integrated position
10-year destination	Every person can direct a portable care plan across institutions; most informational support is inexpensive; local human relationships, physical care and accountable judgment are strengthened.
Design principles	Person-directed, not person-abandoned; completion over content; human authority; bounded AI; plural community; any-channel access; federated interoperability; burden-aware economics; visible governance; reversible scale.
1–3 year strategy	Build the governed connective layer around one care-continuity wedge; establish the paid human network, semantics, adapters, evidence system and financing contract before adding scope.
First practical uses	Visit-plan teach-back; referral and appointment closure; transport/pharmacy coordination; caregiver delegation; benefits/coverage support; record/test-intent reconciliation; human escalation.
First NM pilot	Doña Ana “Next Step Closure” for a defined Turquoise Care diabetes/hypertension cohort after visit/discharge/referral; bilingual, phone/SMS/paper/human, paid promotora/CHW, clinical backstop.

Layer	Integrated position
Evidence requirements	Baseline and comparator; 30-day completion; time to first action; understanding; queue age; review minutes; escalation appropriateness; safety events; resident/caregiver/staff burden; equity; implementation fidelity; total cost of delivery.
Scale conditions	Meaningful completion gain; no unacceptable safety signal; no material burden transfer; equity noninferior or better; response capacity and funding secured; auditability and vendor exit tested.
Stop conditions	Serious preventable harm; unsafe queue delay; subgroup harm; uncompensated labor; privacy breach beyond threshold; inability to identify accountable human; no improvement after prespecified sample; no sustainable payer.
Defer conditions	Any use whose evidence, governance, responder capacity, payment or regulatory classification cannot be made explicit before exposure.

17. Healthcare Vision 2.0—recommended outline

1. Gary's original essay: "A Possible 10-Year Vision" (unchanged, clearly attributed).
2. The governing commitment: from abundant intelligence to person-directed healthcare.
3. The real problem: healthcare services versus the process of obtaining healthcare.
4. Selective abundance: what AI can commoditize and what remains scarce.
5. The connective layer: one person-directed plan, verified next steps and accountable seams.
6. Human architecture: person, caregiver, CHW/promotora, navigator, clinician, supervisor and steward.
7. Bounded AI: allowed functions, red lines, evidence tiers, change control and kill switch.
8. Plural community and rural New Mexico: place, culture, tribe, family, socioeconomic circumstance and preference.
9. Interoperability and semantics: adapters, consent, provenance and definitions of completion.
10. Economic architecture: payer, beneficiary, savings, losses, value capture and sustainability.
11. Learning strategy: first pilot, evidence plan, burden measures and scale/stop gates.

12. Ten-year roadmap: expand only from verified care continuity toward distributed diagnostics or broader service redesign.

18. Document disposition

RECOMMENDATION

Choose: “become the introductory vision section of a larger PSA/IHPF document.” Preserve Gary’s paper unchanged and attributed as the original vision essay. Do not revise it into an operating plan and do not supersede it. Place this reconciliation immediately after it, with a short editorial note separating Gary’s vision from PSA/IHPF’s collective strategy and evidence claims.

This protects authorship and rhetorical force while preventing later model syntheses from being mistaken for Gary’s views. It also gives PSA a versioned structure: Gary’s vision can remain stable; the strategy, evidence register and decision appendix can be updated as implementation and external policy change.

19. Top 12 strategic conclusions

1. Gary’s paper is the right opening vision and the wrong standalone implementation plan.
2. “Person-directed” is the governing term: stronger than institution-centered care and safer than consumer abandonment.
3. The connective layer—not a chatbot, new EHR or general AI clinician—is PSA’s highest-value strategic wedge.
4. Verified closure is the unit of value; information delivered, referral sent and appointment listed are intermediate states.
5. AI abundance is selective. Generation becomes cheap; accountable judgment, attention, physical capacity, trust and financing remain scarce.
6. Abundant output can increase review scarcity; reviewer minutes, authority and queue age are core capacity metrics.
7. Human accountability must be a named, paid role with response standards—not the phrase “human in the loop.”
8. CHWs/promotoras are not an inexpensive substitute for clinicians; they are a paid meaning, trust and continuity layer with defined scope and supervision.
9. Geography can matter less for information and more for logistics; community must be plural, chosen and never assumed free labor.

10. Interoperability should be adapter-based and standards-aligned, while semantic agreement on “closed” is as important as data exchange.
11. A pilot cannot claim savings until payer attribution and value capture are contracted; measure costs shifted to people, caregivers and staff.
12. AI may disappear from the foreground only if it remains fully visible in provenance, performance, audit, appeal and accountability.

20. Top 10 unresolved questions

1. Which exact event and population create enough repeatable next-step failure to justify the first wedge?
2. Which provider, MCO and clinical leader will own response service levels and consequential decisions?
3. What is the baseline 30-day completion rate, by step type, language, geography and subgroup?
4. Who pays for CHW/promotora, supervision, review and evaluation after grant capital ends?
5. What data can be accessed from providers, SYNCRONYS, pharmacies and payers, with what latency and consent?
6. What does “verified closure” mean for each referral, medication, transport, benefit and follow-up task?
7. How many reviewer minutes and escalation slots are available per 100 enrolled people?
8. Which AI functions are regulated medical-device functions, and how will intended use remain bounded?
9. What resident, caregiver and workforce burden thresholds trigger redesign or stop?
10. Which governance entity can hold data, contracts, community authority, incident responsibility and vendor exit rights?

DELIVERABLE 3**Decision Appendix**

What PSA should explore, test, monitor, reject or defer

A. Decision portfolio

Decision posture	What belongs here now
EXPLORE	Turquoise Care/MCO outcome contract; RHT alignment; SYNCRONYS/payer/provider adapters; public-benefit governance; resident/tribal/caregiver authority; partner capacity; accessible incumbent components; ACCESS-aligned Medicare pathway.
TEST	Manual baseline; teach-back; task ownership; verified closure semantics; CHW workload; clinical response SLAs; caregiver delegation; bilingual/offline access; AI shadow drafting/classification; audit and vendor exit; burden measurement.
MONITOR	FDA GenAI framework; CMS ACCESS results and participants; 2027 payer APIs; NM RHT procurements; HRSA shortages; NM CHW reimbursement; peer-reviewed real-world AI outcomes; privacy/Part 2/consumer-protection rules; coverage policy changes.
REJECT	Autonomous diagnosis/treatment; EHR/HIE replacement; smartphone-only design; directories without live capacity; black-box prioritization; unpaid community/caregiver labor; surveillance as default; cost-savings claims without counterfactual; scaling on engagement alone.
DEFER	Neighborhood diagnostics; digital twins; unrestricted EHR agents; unsupervised mental-health support; statewide launch; direct-pay marketplace; continuous monitoring; shared-savings guarantees; broad chronic-care expansion before the closure wedge passes gates.

B. First New Mexico pilot—board-ready specification

Element	Recommended decision
Name	Doña Ana Next Step Closure Pilot
Purpose	Test whether a person-directed human+AI support layer improves completion of agreed care next steps without increasing safety risk or burden.

Element	Recommended decision
Population	Stage 1: 150–300 consenting Turquoise Care adults with diabetes and/or hypertension after a defined visit, discharge or specialty referral. Expand only after gates.
Channels	English/Spanish phone, SMS, web, paper and in-person navigator support; disability access; permissioned caregiver proxy.
Human workflow	Paid promotora/CHW teach-back, barrier discovery, task assignment, appointment/transport/pharmacy/community coordination, follow-up and verified closure; named clinical backstop.
Allowed AI	Draft resident summary; translate approved material; extract candidate tasks; classify barriers; prioritize nonclinical work queue; prepare human-review handoff.
Prohibited AI	Diagnosis, prescribing, interpretation of tests, autonomous triage, denial of service, final eligibility determination, unsupervised mental-health care, emergency delay.
Primary outcome	Proportion of eligible agreed next steps verified complete within 30 days.
Co-primary safety/burden	Serious safety events; escalation appropriateness and latency; unresolved-task age; reviewer minutes; resident/caregiver/staff burden; subgroup performance.
Stage 0 gate	Partner response map, baseline, data/consent plan, manual workflow, clinical protocol, economic term sheet, independent evaluation and stop rules approved before enrollment.
Expansion gate	Meaningful completion gain; acceptable confidence bounds; no material subgroup harm or burden increase; safe queues; workforce and payer commitment; vendor exit tested.

C. Evidence and stop-condition scorecard

Domain	Minimum evidence before scale	Stop / pause trigger
Completion	Prespecified improvement over credible comparison	No meaningful gain after target sample or effect driven only by easy tasks
Clinical safety	Prospective event review; tested escalation drills	Serious preventable harm, delayed urgent escalation or no accountable reviewer

Domain	Minimum evidence before scale	Stop / pause trigger
Equity	Outcomes by language, rurality, disability, age and connectivity	Materially worse completion or burden for a protected/priority subgroup
Resident/caregiver burden	Time, contacts, confusion, privacy concern and task load	Net burden increases or family work becomes required/unpaid
Workforce/review	Caseload, queue age, reviewer minutes, turnover and overtime	SLA breaches, unsafe queue growth or uncompensated scope expansion
Implementation	Fidelity, training, data latency, partner response and exception volume	Core workflow cannot operate without manual heroics
Economics	Full delivery cost, payer attribution, avoided/shifted costs, signed pathway	No sustainable payer; savings accrue elsewhere with no capture mechanism
Governance/privacy	Consent, access logs, incidents, correction, revocation, audit and exit	Material breach, undisclosed reuse, inability to correct or vendor lock-in
Technical	Source accuracy, uptime, drift, override and audit completeness	Untraceable output, unacceptable error/drift or failed rollback

D. 90-day research and action agenda

Timing	Work	Concrete output
Days 1–14	Authority + evidence control	Board sponsor; source/provenance register; confirmed Gary introductory status; conflict/decision log.
Days 8–28	Partner and capacity discovery	Interviews with HCA/RHT, one MCO, FQHC/hospital, clinicians, SYNCRONYS, pharmacy, CHW/promotora organizations, tribal/rural representatives; responder-capacity map.
Days 15–35	Person/caregiver co-design	10–15 compensated journey interviews; Gary case map as one exemplar; plural-community segmentation; accessibility requirements.

Timing	Work	Concrete output
Days 22–45	Economics	Cost model; payer/provider value map; attribution logic; draft financing term sheet; explicit losers/resistance analysis.
Days 29–55	Clinical/privacy/governance	Intended-use statement; red lines; escalation protocol; consent/proxy model; incident process; AI disclosure; kill-switch authority.
Days 36–65	Workflow + semantics	Manual closure workflow; definitions for sent / accepted / scheduled / completed / verified / failed; response SLAs; burden map.
Days 43–70	Baseline + time study	Current completion baseline; call/contact burden; CHW and reviewer minutes; unresolved queue age; data availability.
Days 50–76	Component/vendor diligence	Build/buy/partner matrix; interoperability test; security review; model/data terms; portability and exit test.
Days 60–84	Evaluation protocol	Comparator, power assumptions, outcomes, subgroup plan, safety adjudication, stop rules, independent evaluator and preregistration decision.
Days 85–90	Board gate	Pilot charter, partner letters, financing term sheet, budget, risk register and GO / REVISE / STOP decision.

E. Specific decisions for the PSA Board

1. Adopt “person-directed healthcare” as the governing term and define it to exclude person abandonment.
2. Preserve Gary’s vision paper unchanged and approve it as the attributed introductory vision section of the larger PSA/IHPF document.
3. Approve the connective-layer / verified-closure strategy as the 1–3 year focus; decline a general healthcare AI platform mandate.
4. Authorize only a 90-day Stage 0, with no patient deployment or vendor commitment beyond reversible discovery work.
5. Select post-visit/referral closure as the presumptive first wedge and require Stage 0 evidence before final population/site approval.

6. Require a named clinical sponsor, response SLA and funded CHW/promotora/supervisor/reviewer capacity before enrollment.
7. Require a payer/provider financing term sheet and value-capture logic before software build or scale.
8. Adopt AI red lines, provenance/audit requirements, consent and caregiver-proxy rules, incident reporting and kill-switch authority.
9. Approve completion, safety, equity and total burden—not activity or model accuracy alone—as the governing outcomes.
10. Approve the scale/stop scorecard and empower an independent safety/evaluation function to pause the pilot.
11. Authorize partner outreach to NM HCA/RHT, a Turquoise Care MCO, Doña Ana clinical partners, SYNCRONYS, CHW/promotora organizations and an independent evaluator—without implying commitment.
12. Decide whether PSA should submit a focused public comment to FDA docket FDA-2026-N-7874 before 19 October 2026, emphasizing person-facing GenAI, human accountability, postmarket monitoring and burden.

F. Top 10 external resources PSA should retain

E1 CMS ACCESS Model. Outcome-aligned, technology-supported Original Medicare care and PCP co-management. [Open source](#)

E2 FDA: GenAI-Enabled Medical Devices Discussion Paper / docket. Current regulatory direction; competency, clinical confirmation, monitoring and agentic/foundation-model questions. [Open source](#)

E3 Nature Medicine: pragmatic cluster-randomized LLM trial. Real-world outcome evidence: better documentation, no significant treatment-failure improvement. [Open source](#)

E4 Journal of General Internal Medicine: standardized CHW evaluation. Measured CHW outcomes, implementation fidelity, integrated payer-provider incentives and limits. [Open source](#)

E5 HRSA Health Workforce Shortage Areas dashboard. Continuously refreshed source for the clinical-capacity constraint. [Open source](#)

E6 CMS New Mexico Rural Health Transformation abstract + award. Official initiatives and FY26 award context; potential alignment, not an assumed commitment. [Open source](#)

E7 ASTP/ONC 2026 SVAP standards. USCDI v6, FHIR US Core and updated Da Vinci prior-authorization standards. [Open source](#)

E8 CMS Interoperability and Prior Authorization Final Rule. 2026 process rules and 2027 payer APIs relevant to patient/provider data and authorization closure. [Open source](#)

E9 AHRQ: Understanding the AI Wave. Safety, trust calibration, bias, hallucination, monitoring and workforce governance. [Open source](#)

E10 WHO guidance on large multi-modal models for health. International governance principles for generative/multimodal health AI. [Open source](#)

G. PSA/IHPF source and provenance register

Code	Source	Evidence status
G1	Gary Gomes, “A Possible 10-Year Vision: AI and the Redesign of Essential Services,” August 2026.	Original vision; authorship supplied by project context; not empirical evidence.
P1	PSA/IHPF Personal + Community Health Support Platform Research, 10 August 2026.	Model-generated synthesis grounded in cited PSA and external sources.
P2	PSA Healthcare Access & Affordability Master Compendium, binder dated 10 March 2026.	Model-generated PSA research synthesis.
P3	Consumer-Directed Healthcare Model.	Model-generated conceptual synthesis; used with qualification.
P4	PSA NM Non-Professional Care Layer Playbook, January 2026.	PSA planning synthesis; cites external evidence; not itself an outcome study.
K1	HC-AI Master / Operational Blueprint, Kris Witte healthcare GPT archive.	Clearly located in Kris-named GPT project archive; GPT-assisted; not assumed personal opinion.
K2	HC-AI Governance Charter / AI Acceleration Control.	Same provenance; governance synthesis.
K3	New Mexico HC-AI Hidden Assumptions.	Same provenance; critical synthesis.
K4	AI New Workforce Bottleneck / Blended Workforce Coordination.	Same provenance; capacity/accountability synthesis.
K5	PSA Integrity Check for Healthcare Pilots.	Same provenance; evaluation and stop-gate synthesis.

Code	Source	Evidence status
K6	AI Agents Meaning Layer.	Same provenance; semantic/closure synthesis.
L1	John Griffith–Gary Gomes conversation transcript, 3 June 2026.	Documented lived experience; anecdotal design evidence.
L2	John Griffith–Gary Gomes conversation transcript, 26 August 2026.	Documented lived experience and conceptual discussion; anecdotal design evidence.
N1	New Mexico Healthcare Expenditures Data Needs, prepared for NM Legislative Council Service, January 2025.	State-commissioned analysis; authoritative context with stated data limits.

H. Final synthesis

Gary's vision deserves to survive intact because it states the destination more clearly than the later operating documents. PSA/IHPF's responsibility is to prevent clarity of destination from becoming false confidence about the route. The next generation document should therefore hold two ideas together: intelligence may become abundant, and care will still depend on scarce human attention, trust, judgment, physical capability, institutional response and financing.

The practical opportunity is not to manufacture another intelligent answer. It is to make the process of obtaining healthcare legible, person-directed and accountable—to ensure that an agreed next step has an owner, that the person and caregiver are not carrying invisible project-management work, that a qualified human responds when meaning or safety is at stake, and that completion is verified. That is the connective layer. It is both less glamorous and more consequential than autonomous care.

PSA should proceed—but only as a learning institution: with an attributed vision, an explicit causal hypothesis, bounded tools, paid humans, real partners, independent evidence, transparent economics, and the willingness to stop.