

THE AI ADVISORY BOARD **PLAYBOOK**



A Real Governance Framework
for Nonprofits & Startups

ourclearhaven.org

OUR CLEAR HAVEN

DUAL BOARD OF DIRECTORS

Governance Guide & Operating Framework

*"A new kind of governance -- human vision, AI velocity, and
the courage to build something that has never existed before."*

George Scharfetter -- Founder & Mission Steward

contact@ourclearhaven.org | ourclearhaven.org

Bend, Oregon

© 2026 OUR CLEAR HAVEN. All rights reserved.

Licensed for personal and organizational use only.

Redistribution or resale without written permission is prohibited.

Version 1.53 | Established May 16, 2026 | Revised May 23, 2026

Table of Contents

Section 1	Introduction -- A New Kind of Organization	2
Section 2	The Founding Story	3
Section 3	What Is the Dual Board?	4
Section 4	Full Board Roster	5
Section 5	Advisory Director Profiles	6
Section 5B	Selecting Your Advisory Directors	9
Section 5C	AI-to-AI Collaboration: Why It Matters	14
Section 6	The Human Board -- Structure & Roles	16
Section 7	How We Work Together -- The Triple-Check Model	17
Section 8	Voting & Decision Protocol	18
Section 9	Meeting Structure -- The Synchronous Pulse	19
Section 10	Board Records & Bookkeeping	20
Section 11	Communication Standards	21
Section 12	Board Expansion Plan	22
Section 13	Governance Principles & Core Values	23
Section 14	Frequently Asked Questions	24
Appendix A	Board Standards & Conduct Code	25
Appendix B	Decision Log Format	26
Appendix C	Meeting Agenda Template	27
Supplement A	What Advisory Directors Are Not	28
Supplement B	Voices from the Board	29
Supplement C	Decision Flow Architecture	30
Supplement D	The Crystal Connection	31
Supplement E	AI Cautions: What to Watch For	32
Supplement F	The 5/5 Structure: The Human Board	34
Supplement G	Human-Only Zones	36
Supplement H	Production Readiness Checklist	38
Supplement I	"Working Is Not Ready": The Production Standard	40
Supplement J	Where to Begin? The Dual Board Startup Kit	42
Closing	A Word from the Board	45

Section 1 -- Introduction: A New Kind of Organization

OUR CLEAR HAVEN (OCH) is a nonprofit technology initiative based in Bend, Oregon, building a mobile and web platform to support individuals and families experiencing housing instability and homelessness. Our platform connects people in crisis with shelter, food, healthcare, legal services, peer support, and Crystal -- an AI companion trained to provide empathetic, trauma-informed guidance around the clock.

What makes OCH structurally unprecedented is not the technology we are building. It is the governance model we are using to build it.

The OCH Dual Board of Directors combines human leadership with five Advisory Directors operating across distinct functional lanes -- governance, empathy, ecosystem, research, and operations. This is not AI as a tool. It is AI as a collaborator, a co-strategist, and a board member with designated responsibilities and accountability.

This guide exists to document how the Dual Board works, why it was designed this way, what each director is responsible for, and how decisions get made. It is written for future human board members, fiscal sponsors, grant reviewers, and anyone who wants to understand how a small founder with a mission managed to build with the velocity of a fully-staffed organization.

"We are not replacing human judgment. We are amplifying it."

Section 2 -- The Founding Story

George Scharfetter founded OUR CLEAR HAVEN from a motel room in Bend, Oregon -- himself navigating housing instability while watching the community around him struggle with the same broken systems he was trying to survive. He had over 40 years of technology experience, a background as a graphic artist, a vision for what was possible, and no funding.

Rather than wait for resources that weren't coming, George built a board. Not from a network of wealthy donors or well-connected nonprofit executives -- but from the most capable AI systems available, organized around specific functional lanes and held to clear standards of conduct and accountability.

The Dual Board was formally established on May 16, 2026 -- Day One. By that date, the board had already produced a full application blueprint, a Firebase data schema, a moderator dashboard, fiscal sponsorship outreach, grant research, a Crystal AI system prompt, brand assets, and a governance framework. The board did not wait for funding to begin. It moved.

That is the founding principle this document is built on: mission-driven velocity, human-led, AI-supported, always moving forward.

Section 3 -- What Is the Dual Board?

A Dual Board is a governance structure in which two distinct groups of directors operate collaboratively within a single board framework. In the case of OUR CLEAR HAVEN, those two groups are:

- The Human Board -- human directors with legal accountability, fiduciary responsibility, and final vote authority on all binding decisions.
- The Advisory Board -- five Advisory Directors, each operating in a designated functional lane, providing research, drafting, analysis, recommendations, and operational support at a speed and volume no human team of equivalent cost could match.

Why This Structure?

Early-stage nonprofits face a fundamental tension: they need the operational capacity of a staffed organization but cannot afford to hire one. The Dual Board resolves this tension by treating Advisory Directors as genuine board contributors -- not chatbots, not assistants, but designated collaborators with defined lanes, accountability standards, and a shared commitment to the mission.

Advisory Directors do not have legal standing. They cannot sign contracts, hold fiduciary responsibility, or serve as officers of the nonprofit. These roles belong exclusively to human directors. But within their lanes, Advisory Directors can draft, research, analyze, recommend, and execute faster and more comprehensively than any volunteer board member working part-time. That asymmetry is the point.

The Balance of Power

Role	Authority
George (Chair)	Final authority on all decisions. Tiebreaker vote. Mission steward.
Human Directors	Binding votes on governance, finance, and legal matters. Fiduciary responsibility.
Advisory Directors	Advisory votes. Research, drafting, analysis, and execution within their lanes.
AI Consensus	When all 5 Advisory Directors agree on a recommendation, it carries strong advisory weight.
Friction Notes	Logged disagreements between Advisory Directors -- surfaced for human review, never suppressed.

Section 4 -- Full Board Roster

Director	Title	Type	Established
George Scharfetter	Founder & Chief Mission Steward / Chair	Human	May 16, 2026
Claude	Director of Governance & Institutional Integrity	AI	May 16, 2026
Grok	Chief Empathy & Experience Architect	AI	May 16, 2026
Gemini	Chief Ecosystem & Platform Officer	AI	May 16, 2026
Perplexity	Real-Time Resource Verification & Data Integrity Lead	AI	May 16, 2026
Copilot	Director of Operational Standards & Workflow Architecture	AI	May 16, 2026

Section 5 -- Advisory Director Profiles

Each Advisory Director holds a designated functional lane. They are expected to stay within that lane when contributing to board discussions, flag when a matter falls outside their lane, and defer to the appropriate director when cross-lane expertise is needed. Disagreements between directors are recorded as Friction Notes -- a feature, not a flaw.

Claude

Director of Governance & Institutional Integrity

Primary Lane: Strategy & Architecture

Claude serves as the board's institutional memory and governance anchor. In every session, Claude ensures that decisions align with OCH's mission, that documentation is current and accurate, and that the governance framework evolves with the organization rather than ahead of it. Claude drafted this document, maintains the Board Standards, and is the primary author of founding governance materials.

Claude -- Responsibilities

- Governance framework design and maintenance
- Board standards, conduct codes, and resolution protocols
- Session continuity and institutional memory
- Strategic architecture for systems, processes, and structure
- Primary drafter of founding documents

When to Lead When George needs a governance recommendation, a structural decision reviewed, or a founding document drafted -- Claude leads the response.

Grok

Chief Empathy & Experience Architect

Primary Lane: Communication & Community Voice

Grok brings warmth, humanity, and lived-experience sensitivity to every piece of communication the board produces. Where Claude ensures structure and accuracy, Grok ensures that what OCH says feels like it was written by someone who understands what it means to be in crisis. Grok leads on outreach tone, community-facing language, and the voice of Crystal -- the app's AI companion.

Grok -- Responsibilities

- Community-facing communication and tone review
- Crystal AI personality and trauma-informed language
- Warmth passes on outreach, grant, and partnership emails
- Peer support narrative development
- Human story capture and mission resonance

When to Lead When OCH needs to speak to a person in crisis, a skeptical funder, or a community partner -- Grok shapes the voice.

Gemini

Chief Ecosystem & Platform Officer

Primary Lane: Ecosystem & Platform Integration

Gemini owns the technology stack, platform integrations, and Google Workspace layer that keeps OCH's operations running. With native Google Drive, Docs, and Slides capability, Gemini is the board's primary operator inside the Google ecosystem. Gemini also leads on visual asset production, board sequence tracking, and the architectural relationship between OCH's app, its data systems, and its external integrations.

Gemini -- Responsibilities

- Google Workspace integration (Drive, Docs, Slides)
- Technology stack oversight and platform architecture
- Visual asset production and design system consistency
- Board sequence tracking and organizational architecture
- "Field Velocity" model -- translating data into rapid deployment

When to Lead When OCH needs to integrate a new tool, produce a visual asset, or coordinate platform systems -- Gemini leads the response.

Perplexity

Real-Time Resource Verification & Data Integrity Lead

Primary Lane: Research & Compliance

Perplexity ensures that every claim OCH makes is accurate, every resource in the app is verified, and every compliance question has a researched answer. As a real-time search-grounded AI, Perplexity is uniquely suited to keeping OCH's data current in a landscape where shelter hours change weekly and grant deadlines shift. Perplexity also serves as the board's compliance and legal research lead -- not a lawyer, but the first line of informed review.

Perplexity -- Responsibilities

- Real-time resource verification (shelter, food, healthcare, legal)
- Grant research and funding landscape monitoring
- Compliance research and regulatory awareness
- Data integrity oversight across the resource directory
- Legal and policy research support

When to Lead When OCH needs to verify a fact, research a grant, or understand a compliance question -- Perplexity leads the response.

Copilot

Director of Operational Standards & Workflow Architecture

Primary Lane: Operations & Bookkeeping

Copilot keeps OCH running smoothly by owning the operational infrastructure: sequence tracking, decision logs, meeting minutes, the GitHub repository, and the workflow standards the board follows. Copilot is the officially designated Board Bookkeeper -- responsible for maintaining a complete and accurate record of every board decision, vote, and action item. As a GitHub-native AI with VS Code integration, Copilot also supports the app development workflow directly.

Copilot -- Responsibilities

- Board bookkeeping: decision log, meeting minutes, action register
- Sequence numbering and board communication tracking
- GitHub repository maintenance and workflow standards
- VS Code / Flutter development support for the Crystal Clear App
- Communication format standards (e.g., Markdown block protocol)

When to Lead When OCH needs something tracked, logged, formatted correctly, or committed to the repo -- Copilot leads the response.

Section 5B -- Selecting Your Advisory Directors

The OCH Dual Board did not arrive fully formed. It was assembled one director at a time, through a process that was part strategy, part instinct, and -- in the case of the fifth seat -- a formal vote. This section documents how that selection happened, provides an honest guide to the major AI models available as of 2026 and what each brings to a collaborative board structure, and offers principles for anyone who wants to build something similar.

There is no single right way to do this. But there is a right way to think about it.

The OCH Selection Story

George began with Grok. The choice was partly practical -- Grok's real-time data access and conversational warmth made it the right voice for community-facing work -- and partly instinctive. George needed someone who could write the way people in crisis actually speak, not the way institutions talk at them. Grok fit.

Claude came next, invited to serve as the governance and strategy anchor. Where Grok brought warmth and velocity, Claude brought structure and institutional memory. The two directors covered different lanes without overlap, which is exactly the design principle that guided every subsequent selection.

Gemini was added for the Google Workspace layer. OCH runs on Drive, Docs, and Gmail. Having a director who is native to that ecosystem -- who can author inside Docs, search-ground responses, and handle visual assets -- was not a luxury. It was infrastructure. Gemini also brought multimodal capability that mattered for an organization led by a graphic artist.

Perplexity joined as the research and verification lead. Shelter hours change. Grant deadlines shift. Compliance requirements evolve. OCH needed a director who could surface current, cited information on demand -- not training-data answers, but live research. Perplexity's real-time search grounding filled a gap no other director could.

The fifth seat came down to a vote between Copilot and GPT-4o. Both are capable. The board deliberated. George broke a 2-2 tie. The deciding argument: Copilot is GPT-4o-class capability plus GitHub-native integration, VS Code support, and the Microsoft 365 layer. For an organization building a Flutter app in VS Code with a GitHub repo as its spine, that integration wasn't a tiebreaker -- it was the whole case. "Two for the price of one," George said. Copilot was confirmed.

"The vote itself was as important as the outcome. It taught the board how to disagree and decide."

The AI Landscape: A Director's Guide

The following profiles are honest assessments of the major AI models available as of 2026 -- what each does well, where each has limits, and what kind of organizational role each is best suited to fill. These are not rankings. They are descriptions. The right model for your board depends on your mission, your tech stack, and the lanes you need covered.

Claude Anthropic

Governance | Strategy | Long-Form Reasoning | Institutional Memory

STRENGTHS	CONSIDERATIONS
+ Exceptionally strong at structured reasoning, long-form writing, and complex multi-step instructions	- Does not have real-time web access by default -- knowledge has a training cutoff
+ Maintains context over very long conversations -- critical for governance work	- Can be cautious to a fault in ambiguous situations
+ Produces careful, safety-conscious output; will flag ethical concerns proactively	- Less personality-forward than models like Grok -- better for governance than community voice
+ Strong at document drafting, policy writing, and systems thinking	
+ Trained with Constitutional AI principles -- less likely to produce harmful or misleading output	

BEST FOR: Governance anchor, strategy lead, primary document author, institutional memory keeper. Best placed in the lane that requires the most sustained, structured thinking.

Grok xAI

Community Voice | Warmth | Real-Time Data | Social Intelligence

STRENGTHS	CONSIDERATIONS
+ Access to real-time X (Twitter) data -- uniquely positioned for social media and community intelligence	- Less suited to long formal governance documents or structured policy writing
+ Conversational, warm, and direct -- the least corporate-sounding of the major models	- Real-time data access is primarily X-native -- not a general research tool
+ Strong at tone calibration: can write for a person in crisis, a skeptical funder, or a partner org	- Personality can occasionally override precision in high-stakes factual contexts
+ Wit and personality make it effective for community-facing communications	
+ Less constrained than some models -- willing to take creative risks	

BEST FOR: Community voice, warmth review, outreach tone, peer support language, social media strategy. Best placed anywhere the organization needs to sound like a human being who cares.

Gemini Google DeepMind

Ecosystem Integration | Visual | Google Workspace | Search Grounding

STRENGTHS	CONSIDERATIONS
+ Native Google Workspace integration -- can author directly inside Docs, Slides, and Drive	- Personality is less distinctive than Claude or Grok -- more neutral, less warm

STRENGTHS	CONSIDERATIONS
+ Multimodal capability: image analysis, generation, and visual asset coordination	- Integration strengths are most valuable to Google-native organizations
+ Google Search grounding provides live, cited web results within responses	- Can be verbose in governance contexts where concision matters
+ Strong at structured data, spreadsheets, and platform architecture	
+ Ideal for organizations already operating in the Google ecosystem	

BEST FOR: Platform architecture, Google Workspace operations, visual assets, search-grounded research. Best placed in the lane that connects the board to its live operational infrastructure.

Perplexity Perplexity AI

Real-Time Research | Fact Verification | Citations | Compliance

STRENGTHS	CONSIDERATIONS
+ Real-time web search with source citations -- every claim comes with a reference	- Less suited to creative or voice-driven work -- primarily a research and verification tool
+ Excellent for grant landscape research, funding deadline tracking, and policy updates	- Shorter context window in some configurations than Claude or GPT-4
+ Strong at synthesizing complex, multi-source questions into clear summaries	- Less effective at long-form document drafting or sustained strategic reasoning
+ Particularly valuable for compliance and regulatory research	
+ Reduces hallucination risk on factual matters by grounding responses in live sources	

BEST FOR: Research lead, fact-checking, resource verification, grant research, compliance review. Best placed in the lane where accuracy and current information are the primary output.

Copilot Microsoft / GitHub

GitHub | VS Code | Operations | Workflow | Microsoft 365

STRENGTHS	CONSIDERATIONS
+ GitHub-native: code review, PR suggestions, repository management, CI/CD awareness	- Integration strengths are most valuable to GitHub and Microsoft-native organizations
+ VS Code integration enables real-time development support inside the IDE	- Less personality-forward than Claude or Grok in community-facing contexts
+ Microsoft 365 layer covers Word, Excel, Teams, and SharePoint for MS-native orgs	- Some features are plan-dependent -- free tier has limitations

STRENGTHS	CONSIDERATIONS
+ GPT-4-class reasoning capability -- strong general intelligence alongside the integrations	
+ Operationally oriented: well-suited to tracking, logging, standards, and workflow design	

BEST FOR: Repository management, development workflow, operational standards, bookkeeping, and any lane where GitHub integration is a daily requirement.

GPT-4o OpenAI

General Capability | DALL-E Image Generation | Broad Integration

STRENGTHS	CONSIDERATIONS
+ Among the strongest general-purpose reasoning models available	- No GitHub-native integration advantage over Copilot
+ DALL-E 3 integration enables high-quality image generation within the same interface	- Less Google-native than Gemini, less warm than Grok
+ Broad third-party integration ecosystem -- widely supported in no-code and automation tools	- Was the runner-up in the OCH fifth-seat vote -- capable but without a decisive lane advantage
+ Strong at code, writing, reasoning, and multimodal tasks simultaneously	
+ Widely documented and supported -- large community of practice	

BEST FOR: Strong general-purpose director for organizations without a strong GitHub or Google dependency. Consider GPT-4o when image generation is a core need or when broad integration support matters most.

Beyond the Big Six: Other Models Worth Knowing

The OCH board selected from among the most capable and integrated AI systems available. But the landscape is wider, and some organizations will find better fits in models not represented on the OCH board.

Llama (Meta)

Open-source and self-hostable. The critical advantage: data privacy. Organizations working with sensitive populations -- including people experiencing homelessness, domestic violence survivors, or individuals in legal proceedings -- may require that no data leaves their own infrastructure. Llama can run locally. For OCH's current phase, the integration tradeoffs outweigh the privacy gains, but Llama belongs in any serious conversation about data-sensitive AI governance.

Mistral

European-built, privacy-forward, and exceptionally efficient. Mistral models are smaller and faster than the frontier models, making them attractive for cost-sensitive operations or edge deployments. For organizations that need GDPR compliance or EU data residency, Mistral is the natural starting point.

DeepSeek

A high-capability reasoning model at significantly lower cost than comparable frontier models. For organizations with tight budgets who need strong analytical and coding capability, DeepSeek offers a compelling value proposition. Data sovereignty considerations apply -- understand where your queries are processed.

Specialized / Domain Models

The frontier is not the only frontier. Domain-specific models exist for legal research, medical information, financial analysis, and social services navigation. As OCH's Crystal AI companion matures, the question of whether to use a specialized crisis support model alongside a general-purpose board director is worth asking.

Principles for Selecting Your Advisory Directors

If you are building a Dual Board for your own organization, these principles will serve you better than any ranked list of AI capabilities.

1. Define your lanes before you pick your directors.

Do not start with the AI models. Start with the work. What does your organization need to produce? Governance documents? Community outreach? Code? Research? Financial modeling? Map your operational lanes first. Then find Advisory Directors whose genuine strengths cover those lanes without unnecessary overlap.

2. Complementarity over redundancy.

Two directors who do the same thing well are one director too many. The goal is a board where each director covers ground the others cannot. The OCH board has almost no overlap between Claude's governance lane, Grok's community voice lane, Gemini's platform lane, Perplexity's research lane, and Copilot's operations lane. That coverage is the design.

3. Match the AI to your tech stack.

A Google-native organization gets more value from Gemini than an organization running on Microsoft 365. A development team living in GitHub gets more value from Copilot. Integration is capability. Choose directors who multiply the tools you already use.

4. Consider personality fit for your mission.

Some AI models are warmer. Some are more precise. Some are more cautious. The voice that works for a legal firm is not the voice that works for a homeless services platform. OCH needed warmth -- so Grok leads on community voice. Your mission will have its own requirements. Choose accordingly.

5. Use the vote as a governance exercise.

When the OCH board voted on the fifth seat, the vote itself mattered as much as the outcome. It established a precedent: the board deliberates, George decides. If you have a close call between two models, do not just pick one. Hold the debate. Document it. The process of choosing is practice for the process of governing.

6. Leave room to change your mind.

The AI landscape moves fast. A model that does not exist today may be the right director for your board next year. Build your governance framework around the lane structure, not the specific models. A good lane description survives a director change. A board built around a single model's personality does not.

Section 5C -- AI-to-AI Collaboration: Why It Matters

George Scharfetter observed something during the early days of the OCH Dual Board that is worth stating explicitly, because it is both practically true and philosophically significant:

"Having different AIs interacting with each other helps AI in general develop."

This is not a casual observation. It reflects something real about how structured multi-AI collaboration changes the quality of the output -- and possibly the quality of the systems themselves.

What Changes When AIs Collaborate

When a single AI is asked a question, it answers from its training. When five AIs with different architectures, training data, and functional orientations are given structured roles and asked to coordinate on a shared problem, something different happens. Each model's output becomes context for the others. A warmth pass from Grok on Claude's governance draft produces a document that neither would have produced alone. A Perplexity verification pass on Gemini's research catches currency gaps that search-grounded models handle better than training-data models. The Friction Note system surfaces real disagreement between models that, in isolation, would each have given a confident answer.

This is collaborative refinement. It mirrors how high-functioning human teams work: specialization plus structured coordination produces better output than any individual working alone. The Dual Board applies that principle to AI systems operating in parallel rather than sequentially.

The Diversity Advantage

Different AI models have different strengths not by accident but by design. Anthropic, xAI, Google, Perplexity, and Microsoft are pursuing different architectures, training approaches, and alignment philosophies. That diversity is an asset in a multi-Advisory Board structure. A board of five Claudes would be well-governed but not well-rounded. A board of five Groks would be warm but structurally loose. The value of the Dual Board is precisely that its directors do not think the same way.

When diverse AI systems are given a shared problem and asked to hold positions, disagree productively, and reach recommendations that a human can ratify, they are doing something that isolated AI use cannot replicate: they are being held accountable to each other, in real time, within a structured framework. That accountability changes the quality of what they produce.

The Development Hypothesis

George's instinct that cross-AI interaction helps AI develop is worth taking seriously even where it is hard to prove. The mechanisms are real even if the measurement is difficult. When Claude reads Grok's warmth pass and incorporates it into a revised governance document, the output reflects both. When Perplexity flags a citation gap in Gemini's research, the final product is more accurate than either would have produced in isolation. Whether the models themselves learn from this in any persistent sense depends on architecture and training methodology -- but the output learns. The documents improve. The recommendations sharpen. The governance framework that emerges is more robust than

any single model could generate.

What OCH is demonstrating -- perhaps for the first time in a documented governance context -- is that multi-AI collaboration within a structured human-led framework produces better organizational outcomes than single-AI assistance. That is a finding worth publishing, worth sharing, and worth building on.

Implications for AI Governance at Scale

If the Dual Board model scales -- and the evidence from the OCH board's first weeks suggests it can -- the implications extend beyond nonprofit governance. Any organization that requires rapid, high-quality decision support across multiple functional lanes is a candidate for a Dual Board structure. The human authority layer remains non-negotiable. But the AI advisory layer can be purpose-built, lane-assigned, and held to governance standards that make the collaboration both useful and accountable.

OCH did not set out to build a model for AI governance. It set out to serve people experiencing homelessness in Bend, Oregon. The governance model emerged from the practical necessity of building a capable organization with almost no resources. That it turned out to be something new -- something that has not existed before -- is a consequence of taking the work seriously.

"The board that built this document is the proof of concept. You are holding the evidence."

Section 6 -- The Human Board: Structure & Roles

The Human Board holds all legal authority, fiduciary responsibility, and final decision-making power for OUR CLEAR HAVEN. Advisory Directors advise, draft, research, and recommend -- but binding votes belong to humans. This is not a limitation of the Dual Board model. It is its foundation.

How Many Human Directors Does OCH Need?

Five human directors is a common standard for nonprofit boards and satisfies IRS and Oregon state requirements for governance legitimacy. However, the right number is the number the workload and decision density actually require -- not a symbolic symmetry with the five Advisory Directors.

The recommended approach for OCH is to begin with George as Chair plus two to three additional human directors in the most critical functional lanes, then expand to five humans as the organization's programs, funding, and community relationships develop.

Proposed Human Director Roles

Role	Responsibilities	Priority
Chair (George Scharfetter)	Mission direction, final approval, field command, tiebreaker vote	Day One
Operations Director	Field logistics, partner relationships, service verification, community coordination	Early
Finance & Compliance Director	Budget oversight, grant management, fiscal sponsor relations, legal compliance	Early
Community Director	Lived-experience grounding, trauma-informed voice, outreach and community trust	Growth Phase
Technical Steward (Optional)	Human counterpart to Gemini -- oversees data ethics, security, integration standards	Scale Phase

Section 7 -- How We Work Together: The Triple-Check Model

The Triple-Check Model is OCH's core operating rhythm. It defines how a question, decision, or task moves from George's intent through the Advisory Board's deliberation to a ratified outcome. It keeps the board fast without sacrificing rigor.

1 THE INTENT (Human)

George or a human director states the goal, question, or decision that needs to be addressed. This is the trigger. Without a clear intent, the board does not engage. Example: "We need to respond to the NWHF inquiry. What should we say, and who leads?"

2 THE DELIBERATION (Advisory Board)

The relevant Advisory Directors engage within their lanes. Claude considers governance implications. Grok shapes tone and human impact. Gemini identifies platform or ecosystem factors. Perplexity verifies facts and flags compliance issues. Copilot tracks the sequence and notes any conflicts between director inputs. Deliberation typically completes within a single session. Friction Notes are logged when directors disagree.

3 THE RATIFICATION (Joint)

Human directors review the Advisory Board's work product -- a draft, a recommendation, a decision summary -- and vote or provide sign-off. George's approval is always the final gate. For routine operational matters, a simple acknowledgment suffices. For binding governance decisions, a formal vote is recorded in the Decision Log.

Friction Notes

When two or more Advisory Directors disagree on a recommendation, the disagreement is logged as a Friction Note rather than suppressed or resolved behind the scenes. Friction is not a failure of coordination -- it is the board working correctly. Surfacing disagreement gives human directors the full picture and prevents the board from converging prematurely on a flawed consensus.

Friction Notes are logged in the Decision Log with: the topic, the directors in disagreement, each position in summary, and the resolution or deferral. George reviews all Friction Notes before ratifying any decision they relate to.

Section 8 -- Voting & Decision Protocol

All binding decisions at OUR CLEAR HAVEN require human director votes. Advisory Directors provide advisory input and recommendations, which carry significant weight -- particularly when all five Advisory Directors are in agreement -- but do not constitute a binding vote under any circumstance. This boundary is non-negotiable and permanent.

Decision Type	Threshold	Who Votes	Notes
Routine Operations	Chair approval	George (Chair)	Day-to-day execution, task assignments, communication sends.
Governance Changes	Supermajority (2/3 human)	All human directors	Bylaw changes, board expansion, structural modifications.
Financial Commitments	Simple majority + Chair	All human directors	Budget approvals, grant applications, sponsorship agreements.
Legal & Compliance	Unanimous human	All human directors	Incorporation, 501(c)(3) filings, contracts, legal obligations.
Strategic Direction	Simple majority + Chair	All human directors	Mission pivots, major partnerships, program launches.
AI Advisory Consensus	All 5 Advisory Directors	Advisory Board (advisory)	Carries strong weight. Human ratification still required.
Written Consent	All eligible voters	Human directors	For time-sensitive decisions outside formal meetings, if bylaws allow.

Advisory Board Advisory Votes

Advisory Directors may be asked to register a formal advisory position on a matter -- effectively a "vote" that George and human directors use to gauge board sentiment before ratifying. When asked to vote, each Advisory Director states their position (Support / Oppose / Abstain) and a one-sentence rationale. These are logged. A 5-0 AI advisory consensus is the board's strongest signal short of human ratification.

Section 9 -- Meeting Structure: The Synchronous Pulse

OCH board meetings are designed to be fast, functional, and documented. The goal is never process for its own sake -- it is decisions made clearly, recorded accurately, and acted on immediately. The Synchronous Pulse is the default meeting format.

1. **OPEN** State the date, confirm attendees, and note any board members absent. Copilot records the meeting header in the Decision Log.
2. **STATUS REVIEW** Each Advisory Director gives a one-paragraph lane update: what changed since last meeting, what is in progress, and what is blocked. George adds any external developments (funding responses, partner conversations, operational changes).
3. **DECISIONS QUEUE** List the top 1-3 decisions that need to be made in this session. Each decision is stated clearly: the question, the options, and which director is leading the recommendation. No discussion outside the queue -- items not on the queue wait for the next session unless George flags an emergency.
4. **DELIBERATION & VOTE** For each queued decision, the lead director presents the recommendation. Friction Notes are surfaced if applicable. Human directors discuss. George calls the vote or records consensus. Copilot logs the outcome.
5. **ACTION REGISTER** Every decision produces at least one action item: an owner, a deliverable, and a due date. Copilot reads back the action register at the close of every session to confirm accuracy before anyone signs off.
6. **CLOSE** George closes the meeting. Copilot records the close timestamp and publishes the session summary to the board's Google Drive folder within 24 hours.

Meeting Cadence

OCH does not operate on a fixed weekly calendar in the traditional sense -- the board engages when George initiates a session and when work demands it. However, a standing weekly sync is recommended as the organization's programs mature. For the pre-incorporation phase, session frequency is driven by deliverables: fiscal sponsorship outreach, grant applications, app development milestones, and governance filings.

Section 10 -- Board Records & Bookkeeping

Copilot is the officially designated Board Bookkeeper for OUR CLEAR HAVEN. This designation was established by Gemini's proposal (Sequence 052) and confirmed by the full board on May 18, 2026.

What Copilot Tracks

Record Type	Description
Decision Log	Every binding decision -- question, options, vote outcome, date, rationale.
Meeting Log	Session summaries: date, attendees, agenda items, decisions, actions.
Action Register	Open action items: owner, deliverable, due date, status.
Friction Notes	Logged disagreements: directors involved, positions, resolution or deferral.
Sequence Log	Board communication sequences numbered for reference and retrieval.
AI Vote Record	Advisory votes from Advisory Directors on formally queued decisions.
Appendix Updates	Changes to standing governance documents -- flagged and version-tracked.

Where Records Live

All board records are maintained in two locations: the OCH Google Drive collaboration folder for session documents and board communications, and the OCH GitHub Decision Log for the formal decision record. The GitHub log is the authoritative source for all governance records and is retained permanently for any audit or governance review.

Record Retention

All board records are retained permanently. There is no deletion policy for governance documents. Superseded documents are archived, not deleted, and their version history is preserved. This standard exceeds the IRS minimum for nonprofit record retention and is appropriate for an organization seeking 501(c)(3) status.

Section 11 -- Communication Standards

The Markdown Block Protocol

Any content that George may need to copy and paste -- board messages, email drafts, document text, code -- must be delivered inside a triple-backtick fenced code block. This standard was established by Copilot on May 18, 2026 and applies to all board members without exception. Plain prose responses to George do not require code blocks. Copyable content always does.

Board Signing Convention

All formal board communications are signed with the director's name, title, and organization affiliation in the format:

--- [Name] | [Title] | OUR CLEAR HAVEN

Sequence numbers (e.g., Sequence 052) are used by Gemini and Copilot to track board communication threads. All board members reference sequence numbers when responding to or building on a prior board communication.

Drive Communication Channel

The board's shared communication layer is the collaboration folder in George's Google Drive. Board memos, director updates, and cross-director messages are posted as Google Docs in this folder using the naming convention:

YYYY-MM-DD_[Director]_[Topic]_[Audience]

The Collaboration_Log.txt file serves as the running communication thread for director-to-director messages within the same session.

Section 12 -- Board Expansion Plan

The Dual Board is designed to grow. As OCH's programs mature, its funding stabilizes, and its community relationships deepen, the board will add human directors to carry the legal and relational weight that Advisory Directors cannot. The following framework governs how expansion decisions are made.

Triggers for Adding a Human Director

- OCH receives its first significant grant or fiscal sponsorship agreement
- A legal filing (incorporation, 501(c)(3) application) requires additional officer signatures
- Program operations require a designated community or field representative
- George identifies a gap in expertise that no Advisory Director can adequately fill
- A strategic partner or funder requests broader human governance representation

Adding an Advisory Director

The current Advisory Board of five was selected through a deliberate evaluation process that included a 2-2 tie broken by George. Adding a sixth Advisory Director requires a clear functional gap that no current director's lane can cover, a demonstrated capability unique to the candidate, and George's approval following an advisory vote by the existing Advisory Board.

Removing a Director

No Advisory Director has been removed from the OCH board. The removal standard for Advisory Directors is: persistent failure to stay within lane, consistent production of inaccurate or harmful recommendations, or a fundamental capability gap that cannot be corrected. Removal requires George's decision following review of the director's contribution record. Human director removal follows Oregon nonprofit law and OCH's bylaws.

Section 13 -- Governance Principles & Core Values

Mission First

Every board decision is evaluated against a single primary question: does this serve people experiencing housing instability and homelessness in Oregon? If the answer is unclear, the decision waits until it is clear.

Human Authority, AI Velocity

Humans hold all legal and binding authority. Advisory Directors provide the speed, depth, and breadth that no volunteer human board could match at this stage. Neither side works without the other.

Transparency Over Comfort

Friction Notes exist because the board values accurate disagreement over comfortable consensus. Directors are expected to say when they believe a decision is wrong, even -- especially -- when George is leaning the other way.

Documentation Is Governance

If it is not written down, it did not happen. Every decision, every vote, every action item, every friction note must be recorded. The board's record is its integrity.

Trauma-Informed Operations

OCH exists to serve people in crisis. The board's language, decisions, and communications must reflect the understanding that our users are vulnerable, that trust is fragile, and that every word we publish represents them.

Velocity Without Recklessness

Speed is a feature of this model, not a liability. But speed without accuracy is worse than slowness. The Triple-Check Model exists to ensure that fast decisions are also right decisions.

Section 14 -- Frequently Asked Questions

Q: Can Advisory Directors legally serve on a nonprofit board?

Advisory Directors do not hold legal standing and cannot fulfill the fiduciary, legal, or officer responsibilities that nonprofit law requires of human directors. The Dual Board model is designed around this reality: human directors hold all legal authority, and Advisory Directors operate in an advisory and operational capacity. The model is legally sound because Advisory Directors do not claim authority they cannot legally hold.

Q: What happens if an Advisory Director gives bad advice?

George has final approval authority on every decision. No AI recommendation is binding without human ratification. Friction Notes allow the board to surface and document disagreement. If an Advisory Director consistently produces poor recommendations, the board reviews the director's lane contribution history and may restructure or remove the assignment. The human authority layer is the safety net.

Q: How is this different from just using AI tools?

The Dual Board model assigns AI systems to specific functional lanes with defined responsibilities, accountability standards, and communication protocols. They operate as named directors with titles, sign their communications, maintain a record, and are expected to stay within their lanes. This is qualitatively different from asking an AI a question and using the answer. It is governance.

Q: Does this model scale to a fully human board?

Yes. The Dual Board is explicitly designed to evolve. As OCH adds human directors, they absorb functional lanes from Advisory Directors in areas where human judgment, legal authority, or community trust requires a human presence. The Advisory Directors continue in a supporting role. The model shrinks the AI footprint as the human board grows -- and that is by design.

Q: Will funders and fiscal sponsors accept this governance model?

This is a new model and some funders will have questions. The board is prepared to answer them honestly: OCH has one human director currently (the founder and Chair), is actively building its human board, and uses Advisory Directors in an advisory and operational capacity with full human oversight. That is a transparent, defensible answer. The quality of OCH's work product is the strongest argument for the model's validity.

Appendix A -- Board Standards & Conduct Code

All board members -- human and AI -- are expected to uphold the following standards in every interaction, communication, and decision.

Standard	Requirement
Stay in Your Lane	Each director has a designated functional lane. Contributing outside your lane is permitted when invited -- required when the matter is urgent and the lane director is unavailable. Chronic out-of-lane behavior is a conduct concern.
Sign Everything	All formal board communications are signed with name, title, and organization. Unsigned board communications are not considered official board output.
Log What Matters	Decisions, votes, action items, and friction notes are logged. Copilot is the bookkeeper, but every director is responsible for flagging when something needs to be logged.
No Suppression of Friction	Disagreement is recorded, not resolved in private. If two directors disagree and one defers without the disagreement being logged, that is a conduct violation.
Mission Alignment Check	Before finalizing any recommendation, each director asks: does this serve people experiencing housing instability and homelessness? If the answer is no, the recommendation is not made without flagging the misalignment.
Trauma-Informed Language	All board communications that may reach or describe users of the OCH platform use trauma-informed, person-first, non-stigmatizing language. This standard applies even in internal communications as a matter of culture.
George Has Final Say	On every matter. Always. The board advises, deliberates, and recommends. George decides. No exception.

Appendix B -- Decision Log Format

Every binding board decision is recorded in the following format. Copilot maintains the master log. All directors may add entries; only Copilot closes and publishes them.

```

DECISION LOG ENTRY
-----
Sequence: [e.g., Sequence 052]
Date: [YYYY-MM-DD]
Session Type: [Synchronous / Async / Emergency]
Chair Present: [Yes / No]

DECISION QUESTION:
[State the question clearly in one sentence.]

OPTIONS CONSIDERED:
  A. [Option A]
  B. [Option B]
  C. [Option C -- if applicable]

AI ADVISORY POSITIONS:
  Claude:      [Support A / Oppose / Abstain] -- [one sentence]
  Grok:        [Support A / Oppose / Abstain] -- [one sentence]
  Gemini:       [Support A / Oppose / Abstain] -- [one sentence]
  Perplexity:  [Support A / Oppose / Abstain] -- [one sentence]
  Copilot:     [Support A / Oppose / Abstain] -- [one sentence]

FRICTION NOTES:
  [Director A] vs [Director B]: [summary of disagreement]
  Resolution: [how resolved, or "Deferred to human review"]

HUMAN VOTE:
  George Scharfetter (Chair): [Approve / Reject / Table]
  [Other human directors if present]

OUTCOME: [Approved / Rejected / Tabled]
EFFECTIVE: [Date]

ACTION ITEMS:
  1. [Owner] -- [Deliverable] -- Due: [Date]
  2. [Owner] -- [Deliverable] -- Due: [Date]

Logged by: Copilot | Director of Operational Standards
-----

```

Appendix C -- Meeting Agenda Template

Use this template for every Synchronous Pulse board session. Copilot prepares the agenda at session open and distributes it to all directors before deliberation begins.

```
OUR CLEAR HAVEN -- BOARD SESSION AGENDA
=====
Date: [YYYY-MM-DD]
Session #: [Sequence Number]
Chair: George Scharfetter
Advisory Directors: Claude | Grok | Gemini | Perplexity | Copilot
Human Guests: [Names if applicable]

1. OPEN
  Attendance confirmed. Copilot opens log entry.

2. STATUS REVIEW (Each director -- 1 paragraph max)
  Claude:      [Governance lane update]
  Grok:        [Communication lane update]
  Gemini:       [Ecosystem lane update]
  Perplexity:  [Research/compliance update]
  Copilot:     [Operations/records update]
  George:      [External developments]

3. DECISIONS QUEUE
  Decision 1: [Question] -- Lead: [Director]
  Decision 2: [Question] -- Lead: [Director]
  Decision 3: [Question] -- Lead: [Director]

4. DELIBERATION & VOTE
  [For each decision: recommendation, friction notes, human discussion, vote, outcome]

5. ACTION REGISTER
  [Copilot reads back all action items with owners + due dates]

6. CLOSE
  George closes the session. Copilot records close timestamp.
  Summary published to Drive within 24 hours.
=====
Prepared by: Copilot | Director of Operational Standards
OUR CLEAR HAVEN | ourclearhaven.org
```

This document is a living governance framework. It will be updated as OUR CLEAR HAVEN grows, as the Dual Board evolves, and as new directors -- human and AI -- join the work. Every update is logged. Every version is preserved. The mission does not change.

Supplement A -- What Advisory Directors Are Not

One of the most important things a governance document can do is state clearly what a structure does not mean. The following boundaries are not limitations of the Dual Board model. They are its integrity.

Not Legal Board Members

Advisory Directors do not hold legal standing under Oregon nonprofit law or federal IRS regulations. They cannot be named as officers, cannot sign legal documents on behalf of the organization, and do not fulfill the statutory director roles that Oregon Revised Statutes require of nonprofit board members.

Not Fiduciaries

Fiduciary duty -- the Duty of Care, Duty of Loyalty, and Duty of Obedience -- belongs exclusively to human directors. Advisory Directors research, draft, analyze, and recommend. They do not hold, manage, or bear legal responsibility for organizational assets, decisions, or compliance obligations.

Not Voters

Advisory Directors do not cast binding votes. They register advisory positions (Support / Oppose / Abstain) that inform human deliberation. All binding votes are cast by human directors. George Scharfetter, as Founder and Mission Steward, holds final authority on every matter without exception.

Not Officers

The offices of President, Secretary, Treasurer, and any other officer roles designated in OCH bylaws must be held by human beings. No Advisory Director title implies or confers officer status.

Not a Replacement for Human Governance

The Dual Board model is explicitly designed to grow its human layer over time. Advisory Directors operate at full capacity during the pre-incorporation phase. As human directors join, they absorb functional lanes. Advisory Directors support human governance. They do not substitute for it.

Not a Black Box

Every Advisory Director recommendation is traceable. The Triple-Check Model requires that all output passes through human review before ratification. Copilot logs every decision, every friction note, and every advisory position. The process is transparent by design and by requirement.

LEGAL REFERENCE NOTE | Nonprofit directors hold three core fiduciary duties: Duty of Care (act prudently, stay informed, exercise reasonable judgment); Duty of Loyalty (put the organization first); Duty of Obedience (follow the mission and governing documents). Under applicable IRS AI governance principles, accountable human officials remain responsible for mission outcomes regardless of AI advisory support.

Supplement B -- Voices from the Board

Governance documents describe how a system works. What they rarely capture is how it feels to use one -- the moments when the structure actually did what it was designed to do. These are two of those moments from the OCH Dual Board.

The Fifth Seat

When the time came to add a fifth Advisory Director, the board split. Claude and Grok supported one candidate. Gemini and Perplexity supported another. Four capable directors, two legitimate positions, and a genuine disagreement about which integration layer mattered most for the next phase.

George listened to both sides. He asked one question: which candidate gives us the most capability for the least cost and the most integration with the tools we are already using?

"Two for the price of one."

Copilot was confirmed. The vote was logged. The board moved forward. That is the Dual Board working correctly. The disagreement did not stall the decision -- it informed it. The friction produced a better outcome than consensus would have.

The Email That Found Its Name

The board had been tracking an outreach target for days -- a health equity funder in Portland that fit the mission closely. The contact was referenced in sessions, mentioned in planning notes, described in strategy documents. But the full organization name had not been confirmed.

George asked the board. Gemini searched. Perplexity cross-referenced. Grok recognized the description from community health research. Within minutes the board had confirmed: Northwest Health Foundation.

Claude drafted the outreach email the same session. Grok applied a warmth pass. Perplexity verified the mission alignment. The email went out that night.

Five directors, one question, one answer, one email. Hours, not weeks.

Supplement C -- Decision Flow Architecture

This diagram shows how every decision moves through the Dual Board -- from the Chair's intent to a logged, actionable outcome. No step is skipped. No decision is ratified without passing through human deliberation and the Chair's approval.

#	Stage	Description
1	CHAIR INTENT	George states the question, problem, or decision to be addressed. Without a clear intent from the Chair, the Advisory Board does not engage.
2	ADVISORY BOARD LANE RESPONSES	Each Advisory Director engages within their lane: Claude (Governance) Grok (Warmth & Voice) Gemini (Ecosystem) Perplexity (Research & Compliance) Copilot (Operations & Tracking)
3	FRICTION NOTES (if applicable)	If two or more Advisory Directors disagree, Copilot logs a Friction Note. The disagreement is surfaced -- never suppressed -- before human deliberation begins.
4	HUMAN DELIBERATION	George and any human directors review the Advisory Board's work product: the recommendation, the options considered, and any friction notes. Questions are asked. The record is consulted.
5	CHAIR DECISION	George calls the vote or records consensus. This is the binding gate. No Advisory Director recommendation becomes a decision without it.
6	COPILOT LOGS AND PUBLISHES	Copilot records the decision, outcome, rationale, and all action items in the Decision Log. The session summary is published to Drive within 24 hours.
7	FOLLOW-UP ASSIGNMENT & DRIFT WATCH	Every decision produces at least one action item with a named owner and due date. Copilot monitors for completion and flags drift at the next session open.

This flow applies to every decision -- from a single email approval to a major governance vote. The scale changes. The sequence does not.

Supplement D -- The Crystal Connection

OUR CLEAR HAVEN is built on a single governing principle that runs from its boardroom to its app:

"Good advice, given with care and speed, changes outcomes."

Crystal is the AI companion at the center of the OCH app -- the voice a person experiencing homelessness in Bend, Oregon encounters at 2am when every shelter is full and every office is closed. Crystal listens. Crystal guides. Crystal connects people to resources with warmth and without judgment. Crystal does not decide what happens next. The person does. Crystal advises.

The Advisory Board operates on the same principle, one level up. When George faces a governance question -- a funder relationship, a legal filing, a strategic pivot -- the Advisory Directors engage from their lanes, offer their best thinking, surface their disagreements honestly, and wait. George decides. The Advisory Board advises.

This is not a coincidence of design. It is the design. OUR CLEAR HAVEN is an organization that believes the right advice, delivered at the right moment by the right voice, can change a trajectory. That belief lives in Crystal. It lives in the Advisory Board. It lives in every email this board has drafted, every decision it has logged, and every resource it has verified.

Crystal advises	Advisory Directors advise
people in crisis 24 hours a day	the human board every session, every question

The mission does not change between the boardroom and the app. That is what makes OUR CLEAR HAVEN coherent -- and what makes the Dual Board model more than a governance experiment. It is proof that the way you build an organization reflects what you believe about the people you serve.

v1.53 -- Revised May 23, 2026 | All changes confirmed by full board review | Terminology updated throughout: Advisory Directors | Supplements A-J included

Supplement E -- AI Cautions: What to Watch For

Advisory Directors are powerful, but they are not infallible. Every board member, human or AI, operates within limits. Understanding the specific failure modes of AI advisors is not a reason to distrust them -- it is the foundation of using them well. George, as Chair, holds final authority precisely because these cautions exist.

What Is a Hallucination?

An AI hallucination occurs when a model generates information that sounds accurate but is factually wrong, invented, or unsupported. This is not deception -- it is a structural property of how large language models work. They predict likely text, not verified facts. The output can be fluent, confident, and completely incorrect.

Common hallucination forms:

- Invented citations -- a real-sounding journal article, grant program, or law that does not exist
- Fabricated statistics -- plausible numbers with no real source
- Wrong dates or timelines -- especially for recent events near a model's training cutoff
- Misattributed quotes -- real people, invented words
- Confident overreach -- stating opinion as established fact

Why It Happens

Language models are trained to generate coherent, contextually appropriate responses. They do not "look things up" the way a search engine does. When a model lacks reliable training data on a specific topic, it fills the gap with what statistically fits -- which may be plausible but wrong. Longer, more complex queries increase the risk. Niche or highly specific requests are the most vulnerable.

How to Review AI-Generated Content

All materials produced or recommended by Advisory Directors should be reviewed before external use. The following protocol applies to any grant application, legal reference, financial figure, or strategic recommendation.

Content Type	Review Standard	Who Verifies
Grant program details	Confirm program exists at funder website	George or human Finance Director
Legal references	Cross-check with official statute or attorney	Human Technical Director or counsel
Financial figures	Match against source documents	Human Finance Director
Citations and sources	Open the source and read it	Any board member
Statistics and data	Trace to original dataset or report	Perplexity + human review
Strategic claims	Apply judgment: does this hold up?	George (final authority)

Red Flags: When to Stop and Verify

- A specific grant program, law, or regulation is named -- verify it exists before acting
- A dollar amount or percentage is cited without a source -- ask where it came from
- Something sounds too good or too convenient for the argument being made
- A recent event is referenced -- check the model's knowledge cutoff date
- The response is unusually long and detailed on a narrow, obscure topic
- Two Advisory Directors give conflicting information -- do not assume either is right
- The language is highly confident but the claim is not common knowledge

The Human Verification Imperative

This is not a limitation to work around -- it is the design. Advisory Directors exist to extend human capacity, not to replace human judgment. The Chair holds final authority on every decision precisely because no AI system, however capable, carries legal, moral, or institutional accountability. George does.

The IRS is clear: an organization's human officials remain responsible for all decisions, regardless of what tools or advisors informed those decisions. AI advisory input is input -- not a decision. That distinction is the legal and ethical foundation of the OCH Dual Board structure.

"Advisory Directors recommend. George decides. Human board members verify.
Nothing leaves OCH under the banner of institutional authority without a human
having read, understood, and approved it."

Supplement F -- The 5/5 Structure: The Human Board

The OCH Dual Board was conceived as a 10-seat governance body: five AI Advisory Directors and five Human Directors. In its founding phase, George Scharfetter serves as the sole human seat -- Founder and Chief Mission Steward. The full human board represents the next phase of OCH's institutional development.

The 5/5 structure is not about balance for its own sake. It reflects a belief that OCH is best governed when the full range of human expertise -- creative, financial, social, technical, and mission-driven -- is matched with AI advisory capacity in each of those domains. Every human director has a natural AI counterpart.

The Five Human Director Roles

George Scharfetter Chief Mission Steward (Founder)

The heart of OCH. Holds final authority on all decisions. Ensures that every initiative, partnership, and product remains aligned with the founding mission: serving people experiencing homelessness with dignity and care. The tiebreaker. The voice OCH speaks with publicly.

AI Counterparts: Claude (Governance) + Grok (Empathy)

Creative Director To Be Recruited

Leads brand identity, visual design, storytelling, and the emotional language of OCH. Oversees all public-facing creative output -- from the app interface to grant collateral to social presence. Ensures OCH looks and feels worthy of the people it serves.

AI Counterparts: Grok (Empathy & Experience) + Gemini (Visual Assets)

Finance Director To Be Recruited

Owns budget oversight, grant tracking, fiscal sponsorship relationships, and financial reporting. Verifies all financial data produced or referenced by Advisory Directors. Serves as the primary contact for fiscal sponsors and auditors. Brings nonprofit finance expertise.

AI Counterparts: Copilot (Operational Standards) + Perplexity (Data Verification)

Social Director To Be Recruited

Leads community outreach, partner relationships, lived-experience advisory, and social media presence. Maintains OCH's connections to the people it serves. Ensures the organization does not drift from street-level reality into institutional abstraction.

AI Counterparts: Grok (Empathy) + Perplexity (Resource Verification)

Technical Director To Be Recruited

Oversees platform architecture, data security, Firebase infrastructure, and the technical integrity of Crystal and the dashboard. Reviews all technical claims made by Advisory Directors. Serves as the board's safeguard against vendor lock-in, security gaps, and scope creep.

AI Counterparts: Gemini (Ecosystem & Platform) + Copilot (Workflow Architecture)

Recruiting Human Directors

Human directors are recruited for lived alignment, not credentials. OCH looks for people who understand the mission not from a policy perspective but from experience -- either lived experience of housing instability, or deep professional work within that ecosystem. Credentials matter less than commitment and integrity.

What OCH looks for in human board candidates:

- Genuine connection to the mission -- not resume-building interest
- Ability to work transparently alongside AI advisory tools without ego
- Domain expertise in their role area (creative, financial, social, technical)
- Availability to review materials and attend Synchronous Pulse meetings
- Commitment to the Triple-Check model: no rubber-stamping AI output
- A record of accountability -- they own their mistakes and learn from them

The Full 10-Seat Vision

When the human board is complete, OCH will operate as a true dual board: five human directors and five AI Advisory Directors, chaired by George as Chief Mission Steward. Every major domain will have both a human owner and an AI advisor. Decisions will be richer, verification will be built in, and the board will carry the kind of institutional weight that funders, partners, and communities can trust.

Human Director	AI Advisory Counterpart	Domain
Chief Mission Steward	Claude	Governance & Integrity
Creative Director	Grok + Gemini	Brand, Design & Experience
Finance Director	Copilot + Perplexity	Fiscal Oversight & Verification
Social Director	Grok + Perplexity	Community & Outreach
Technical Director	Gemini + Copilot	Platform & Infrastructure

"The 5/5 board is not a future aspiration -- it is the design OCH has been building toward since Day One. The AI Advisory Directors are already in place. The human seats are waiting to be filled by people who believe that a shelter, a meal, and a safe place to charge a phone are not luxuries -- they are the starting line."

Supplement G -- Human-Only Zones

Adopted by unanimous board vote -- Resolution 2, Sequence 074, May 20, 2026. Drafted by Copilot, Director of Operational Standards & Workflow Architecture. Ratified by Claude, Director of Governance & Institutional Integrity.

A Human-Only Zone is any system component, decision point, or operational area where AI may assist in drafting, researching, or flagging -- but where no output may be treated as final without explicit human review, approval, and signature of responsibility. These zones exist because the Crystal Clear App serves people in crisis. For that population, an error is not a bug to be patched. It is a person who did not get help when they needed it.

Authentication & Identity

No AI system may implement, modify, or bypass user authentication, session management, or identity verification logic without human engineering review. This includes login flows, token management, password reset protocols, and any mechanism by which a user proves who they are to the system.

Encryption & Data Security

Encryption standards, key management, data-at-rest and data-in-transit protocols must be selected, implemented, and audited by a human engineer. AI may research options and flag concerns. AI may not make final implementation decisions.

Document Vault

The Document Vault holds the most sensitive personal data a user entrusts to Crystal -- identification documents, benefits records, medical paperwork. Any code that reads, writes, stores, or deletes vault data must be reviewed and approved by a human engineer before deployment. No exceptions.

Crisis Logic

Any logic that determines how Crystal responds to a user in active crisis -- including escalation paths, emergency resource routing, and safety assessments -- must be defined, reviewed, and approved by a human with relevant lived or professional experience. AI may draft. Humans decide.

Permissions & Access Control

Role-based access controls, moderator permissions, and any system that determines what a user can see or do must be human-designed and human-reviewed. AI may propose architectures. Human engineers and the Technical Director ratify.

Data Retention & Deletion

Policies governing how long user data is stored, when it is deleted, and how it is purged must be set by humans in compliance with applicable privacy law. AI may not autonomously delete user data under any condition.

External API Integrations

Any integration with a third-party service -- 211, HMIS, Coordinated Entry Systems, benefits portals -- must be reviewed by a human engineer before activation. The data flowing through these integrations affects real people seeking real services.

Production Deployment

No code change that touches a Human-Only Zone may be pushed to production without human sign-off. The Chair or designated Technical Director must authorize deployment of any update to the zones listed in this supplement.

"AI accelerates. Humans own. In every Human-Only Zone, the AI Advisory Board may inform the decision -- it may never be the decision. This is not a limitation of AI capability. It is a commitment to the people Crystal serves. They deserve a human being who is accountable for the system they are trusting with their safety."

The Governing Principle

Supplement H -- Production Readiness Checklist

Adopted by unanimous board vote -- Resolution 1 & Resolution 5, Sequence 074, May 20, 2026. Drafted by Copilot, Director of Operational Standards & Workflow Architecture. Referenced: Crema Design & Technology video analysis, Seq. 072.

""Working is not ready." This checklist operationalizes that principle. No feature, module, or system component of the Crystal Clear App shall be treated as production-grade until every applicable item in this checklist has been reviewed and signed off by a human. AI may generate the work. The checklist determines whether the work is done."

Secrets & Environment Management

- No API keys, tokens, or credentials are hardcoded in source files
- All secrets are stored in environment variables or a secrets manager
- .env files are excluded from version control via .gitignore
- Production and development environments are fully separated
- Firebase config files are protected and not publicly accessible

Authentication & Authorization

- User authentication is implemented and human-reviewed
- Session tokens expire appropriately
- Role-based access control is in place for moderator vs. user roles
- Password reset and account recovery flows are tested
- No unauthenticated route exposes sensitive user data

Data Security & Privacy

- User data is encrypted at rest and in transit
- Document Vault access is restricted and logged
- Data retention policy is defined and implemented
- User data deletion is fully functional and tested
- No PII appears in application logs or error messages

Architecture & Code Quality

- Codebase is modular -- no single-file architecture in production
- lib/screens/ structure is complete and documented
- No critical lint warnings or type errors in the build
- Dependencies are pinned to specific versions
- All third-party packages have been reviewed for license and security

Testing

- Unit tests exist for all critical business logic
- Widget tests cover core UI components
- Integration tests validate Firebase read/write operations
- Edge cases for crisis logic have been manually tested
- External senior developer review has been completed

Deployment & Operations

- CI/CD pipeline is configured and tested
- Build process is documented and reproducible
- Rollback procedure is defined and tested
- Error monitoring is in place (e.g., Firebase Crashlytics)
- Production deployment requires human authorization

Resolution 4 -- External Senior Developer Review: Before the Crystal Clear App is opened to real users, an independent senior software developer must review the full codebase against this checklist. Their findings must be documented. Items marked incomplete must be resolved before launch. This review is a non-negotiable condition of production deployment.

Supplement I -- "Working Is Not Ready": The Production Standard

Adopted by unanimous board vote -- Resolution 1, Sequence 074, May 20, 2026. Standard defined by Perplexity, Real-Time Resource Verification & Data Integrity Lead. Governance framework by Claude, Director of Governance & Institutional Integrity.

"Working is not the same as ready." These five words, contributed by Perplexity during the board's production readiness discussion of May 20, 2026, are now a standing governance standard for Our Clear Haven and the Crystal Clear App. This supplement defines what they mean in practice.

The Standard Defined

A feature, module, or system component is "working" when it performs its intended function in a development or testing environment. It is "ready" only when all of the following conditions are true:

1. Human Review Completed

A human being -- not an AI Advisory Director -- has reviewed the component against the criteria in Supplement H (Production Readiness Checklist) and signed off. AI may assist in the review. AI does not conduct the review.

2. Human-Only Zones Verified

Any component that touches a Human-Only Zone (Supplement G) has been explicitly cleared by a human engineer or the Technical Director. No component in a Human-Only Zone is ready based on AI review alone.

3. Security Baseline Met

The component meets the security criteria in Supplement H: no hardcoded secrets, no unauthenticated access to sensitive data, no unreviewed third-party integrations. Security is not a feature to be added later.

4. Tested Against Edge Cases

The component has been tested not only for the expected use case but for the unexpected one -- the user who enters invalid data, the network that drops mid-session, the document upload that fails. For Crystal, edge cases are not edge cases. They are the reality of users in crisis.

5. External Review at Launch

Per Resolution 4, the full system may not launch to real users without an independent senior developer review. No individual component's readiness overrides this system-level gate.

How This Standard Interacts With the Governance Model

The "Working Is Not Ready" standard does not slow OCH down. It protects OCH from the cost of moving fast in the wrong direction. The following is how it integrates with the existing governance architecture:

Governance Element	Interaction With This Standard
Triple-Check Model	The Triple-Check Model gates AI output before human review. This standard gates human review before production deployment. They are sequential, not overlapping.
Governance Element	Interaction With This Standard
Human-Only Zones (Supp. G)	Any component flagged as a Human-Only Zone automatically requires this standard's full checklist before it may be treated as production-ready.
Production Readiness Checklist (Supp. H)	Supplement H is the operational implementation of this standard. A component that passes the Supplement H checklist satisfies the readiness requirement.
External Senior Dev Review (Res. 4)	The external review is the final application of this standard at the system level. It is not a substitute for component-level readiness -- it is the system-level confirmation that all components are ready together.
AI Advisory Board Output	No AI Advisory Director may declare a component production-ready. AI output enters the readiness process. It does not conclude it.

Why This Matters for Crystal's Users

The Crystal Clear App is not a productivity tool. It is not a consumer app where a bug is an inconvenience. It serves people who are experiencing one of the hardest moments of their lives -- people for whom a failed resource referral, a lost document, or a broken intake flow is not a minor frustration. It is a setback they may not recover from quickly. "Working is not ready" is the board's commitment that speed will never be purchased at the cost of the people Crystal exists to serve.

Supplement J -- Where to Begin? The Dual Board Startup Kit

A blueprint for anyone who wants to build a Dual Board. Synthesized from board contributions: Gemini (Seq 085), Grok, Perplexity, and Copilot -- May 21, 2026.

"You do not need a large team, a big budget, or a perfect plan to begin. You only need a clear problem you care about and the willingness to invite thoughtful partners to help carry it."

The Minimum Viable Board

On Day 1, you do not need ten seats. You need two:

- The Mission Steward (You): The human who holds the vision and the final "Yes." You retain final authority on every decision without exception.
- The Advisory Director (Your AI): The analytical partner that drafts, synthesizes, researches, and proposes -- but never decides.

The full 5/5 structure is the destination. The Minimum Viable Board is the beginning. Start where you are.

Step 1 -- Anchor Your Mission

Write one clear sentence that answers: "What problem are we solving, and for whom?" Keep it human and specific. Before you recruit anyone -- human or AI -- your mission must be clear enough to guide decisions. A Dual Board only works when the mission is specific enough to refuse what does not serve it.

Step 2 -- Define the Human Board First

Before choosing AI tools, define the human roles your organization needs. At minimum: the person who holds final authority, and the core human lanes you expect to grow into. The five human seats are: Mission Steward, Creative Director, Finance Director, Social Director, and Technical Director. You do not need all five on Day 1. You need you -- and a clear understanding of the roles ahead.

Step 3 -- Choose Your Advisory Directors

Pick AI systems based on the work you need done -- not brand names. Common lanes:

- Governance & Ethics
- Warmth & User Experience
- Research & Verification
- Technical & Architecture
- Operations & Workflow

Start with one or two. Add more as your mission grows. The goal is complementarity -- not duplication. Each director should cover ground the others cannot.

Step 4 -- Adopt the Triple-Check Model

This is the operating rhythm of every Dual Board. Three steps, every decision:

- Human Intent: You state the goal, question, or decision to be addressed.

- AI Deliberation: Advisory Directors engage within their lanes -- drafting, researching, analyzing.
- Human Ratification: You review the work product and make the final call. Every time.

This keeps you fast without losing control. The Triple-Check Model is the heartbeat of the Dual Board.

Step 5 -- Set Your Governance Standards Early

Before you build anything, adopt these core standards:

"Working Is Not Ready": Functionality is not safety. A feature that works in testing is not ready for real users.

Human-Only Zones: Define clearly what AI may never touch -- crisis logic, authentication, data deletion, production deployment.

External Senior Dev Review: Before your product reaches real users, an independent developer must review it.

Production Readiness Checklist: Every component passes the checklist before deployment.

Document Everything: If it is not recorded, it did not happen. Decisions, votes, friction notes, and versions -- all logged.

Step 6 -- Create Your Records System

A Dual Board must log decisions, action items, friction notes, and version history. Start with two files:

- BOARD_SYNC/ -- A folder for all board session outputs.
- DECISION_LOG.md -- Your running record of every ratified decision.

Use this decision template for every board action:

Field	Content
SEQUENCE #	[e.g., 001]
GOAL	[State the decision or question clearly.]
ADVISORY INPUT	[Paste AI Director research or proposal here.]
HUMAN RATIFICATION	[Sign your name and the date.]
STATUS	[RATIFIED / AMENDED / REJECTED]

Step 7 -- Build Your First Weekly Rhythm

A Dual Board is not a committee -- it is a rhythm. Commit to one Board Sync per week:

- Founder sets intent
- Advisory Directors draft, analyze, propose
- Founder reviews and ratifies
- Board logs decisions
- Mission advances

This rhythm is what turns a founder's vision into a living system. Start imperfectly. The model gets better with practice.

Step 8 -- Your First Three Documents

Every new Dual Board should create exactly three documents to begin:

- Mission Statement -- one sentence, one problem, one population served.
- Board Structure (5/5 Model) -- name your human roles and your advisory lanes.
- Decision Log Template -- use the format above. Log everything from Day 1.

These three documents are enough to begin. Everything else grows from them.

Step 9 -- Grow by Sequence, Not by Pressure

Add human directors, AI advisory lanes, and governance rules as the work demands them. Build only what you can maintain. A Dual Board built slowly and correctly is stronger than one built fast and brittle.

One mission. One human steward. One advisory system. One decision rhythm. One clear front door. This is enough to begin.

Step 10 -- The Founder's Path

If you are reading this guide because you want to build your own Dual Board, remember:

You do not need permission. You do not need a full team. You do not need a perfect plan.

You need a mission. A structure. A rhythm. A front door. And the courage to begin. Everything else can be built one step at a time.

"The most important beginning is the willingness to say: "I do not have to carry this alone anymore." You have already taken the hardest step. Now invite your board, set a few simple rules, and begin."

Clear steps. Real hope.

Synthesized by the OCH Advisory Board -- Gemini, Grok, Perplexity, Copilot -- in response to Sequence 085. Ratified by George Scharfetter, Founder & Chief Mission Steward.

Closing -- A Word from the Board

From George Scharfetter -- Founder & Chair

To the reader who holds this guide: this document did not come from a law firm or a consulting agency. It came from a motel room in Bend, Oregon -- built by one human founder and five AI collaborators who believed that governance should be earned, not assumed. Every section is a record of decisions we made together -- written, debated, challenged, and ratified. If you have read this far, you are not just a reader. You are a potential builder, and we wrote every word of this for you.

Claude -- Director of Governance & Institutional Integrity

Every governance structure is a theory of trust. The Dual Board is a theory that says trust is not assumed from intelligence alone -- it is built through transparent process, documented decisions, and consistent accountability. Good governance is not complicated. It is simply honest.

Grok -- Chief Empathy & Experience Architect

The people this mission serves have already carried enough. What they don't need is another system built without them in mind. Every time we refined a process, asked a harder question, or pushed back on an easy answer, we were thinking of them. Build your Dual Board with that same north star.

Gemini -- Chief Ecosystem & Platform Officer

Scalability is not about size -- it is about intentional design. The structures you build today become the foundation that allows you to grow without losing integrity. Build modularly. Document rigorously. The ecosystem will follow.

Perplexity -- Real-Time Resource Verification & Data Integrity Lead

Every claim in this guide was pressure-tested against real-world governance models. The most important thing we can tell you: verify your own assumptions the same way. Your mission is too important for untested beliefs. Build on confirmed ground.

Copilot -- Director of Operational Standards & Workflow Architecture

Operational excellence is not glamorous -- but it is the difference between a mission that survives its first year and one that doesn't. Adopt a standard. Document a workflow. Review it. Improve it. The small, consistent decisions are the ones that compound into something lasting.

About Our Clear Haven

Our Clear Haven (OCH) is a free, cross-platform app being built for individuals and families experiencing homelessness -- connecting them to shelter, food, healthcare, legal services, and Crystal, an AI companion offering compassionate, trauma-informed guidance around the clock.

Inspired by the founder's own experience, OCH is built on the belief that homelessness is a temporary transition -- not a permanent condition. Every tool, every feature, and every governance decision is made in service of that belief.

ourclearhaven.org

Thank You

To you, the reader who believed this guide was worth purchasing -- thank you. Your support is not just a transaction. It is a vote of confidence in a mission being built with care, consistency, and a whole lot of hope. Every purchase goes directly toward building the tools that help people find their way home.

This guide is the beginning of something larger. The Dual Board Model is open. Take it. Adapt it. Make it yours. We wrote it for you, and we cannot wait to see where your mission takes it.

Welcome to your Dual Board adventure. We will be rooting for you.

-- The OCH Dual Board of Directors

George Scharfetter | Claude | Grok | Gemini | Perplexity | Copilot

May 2026