

FELLOWSHIP INTELLIGENCE · EXECUTIVE REFERENCE

The Consigliere

How to build an AI thinking partner that actually knows you

Everyone says they want a Jarvis.

The AI that anticipates what you need. That already has the context. That knows what you mean without you having to re-explain everything every time.

That's a reasonable goal. It's just aimed at the wrong thing.

Jarvis Executes

Deploys the suit. Runs the systems. Manages the logistics. An execution partner is useful, but execution is the part you should be most careful about delegating.

The Consigliere Thinks

What's harder to build, and worth considerably more, is a thinking partner. Not a system that acts for you, but one that thinks *with* you.

What a Consigliere Actually Is

A consigliere doesn't act. A consigliere thinks with you. They know everything about your situation: your real constraints, your actual exposure, the decisions you've been avoiding and why.

No Stake in the Outcome

They have no career interest in your decisions. Their only job is to help you think clearly before you act.

Tells You What You Need to Hear

Not what you want to hear. The honest friction that makes your thinking sharper.

Almost Impossible to Create

Most executives have never had one. Not because they couldn't find one, but because the conditions that make that relationship work are nearly impossible to create with another person.

Almost impossible with another person. An AI version sidesteps the constraints that make the human version so rare. That's what the rest of this document builds.

The Cost of Thinking Out Loud

Think about what it actually costs to think out loud in your position.

With Your Team

You can't ask whether your strategy might be wrong without creating uncertainty.

With Investors

You can't admit you're not sure the market is real.

With the Board

You can't reveal you've been avoiding a decision for three months without it changing how they see you.

So the thinking that should happen out loud happens in silence instead.

Or it doesn't happen at all.

- ① This is what AI can actually be, if you build it right. Not because it's smarter than the advisors in your network. Because it has **no career stake in your decisions**. That's not a capability advantage. That's a psychological one. And for people who spend all day performing certainty for others (projecting more confidence than they actually feel), it matters more than most will admit.

Before You Start: Get the Platform Right

The setup you're about to run involves your real priorities, your actual constraints, your competitive positioning, and the decisions you haven't made yet. That information does not belong in an account where the provider retains the right to train on your inputs.

Three questions decide whether your account is safe for this work. Answer all three before you start.

1

Which Account Tier are you on?

Consumer and free accounts are not designed for this. Team and enterprise accounts exist specifically because business conversations require different data handling terms. If you're not sure which tier you're on, assume it's the wrong one.

2

Do you have a data processing agreement?

If your work touches regulated information of any kind: health, financial, legal, or proprietary R&D. **This is not optional.** Verify current terms with your provider.

3

Are training rights turned on for this account?

On consumer-tier accounts, the answer is almost always yes. That means your strategic thinking, your IP context, your internal conflicts, all potentially training data for a model that will talk to your competitors tomorrow.

⚠ The Consigliere is only as safe as the house it lives in.






The price difference between a consumer and enterprise account is not the reason to upgrade.

The data handling agreement is. *Platform data handling terms change. Verify current terms directly with your provider before acting on this guidance.*

Before You Start: What Not to Feed It

A separate question from platform tier is what goes into the Consigliere even on a properly governed account. Context and strategic framing belong in. Certain categories of information do not, regardless of how secure the platform is.

Do Not Enter

-  Specific names of employees in the context of performance issues, disciplinary matters, or termination
-  Client or customer names in the context of sensitive deal terms, disputes, or confidential negotiations
-  Anything under active litigation, regulatory investigation, or legal hold
-  Material non-public information: financials, deal terms, or undisclosed strategic decisions
-  Personal health, financial, or family information about anyone other than yourself
-  Passwords, access credentials, or security configurations of any kind

The Simple Test

If you remember one rule from this page, remember this:


If a screenshot of this conversation surfaced publicly, what would the consequence be?

If the answer is anything other than **"none,"** rephrase or leave it out.

You can discuss a personnel situation without naming the person. You can work through a deal without naming the counterparty. The Consigliere doesn't need the name to help you think.



Step 1: Give It Your Public Voice



Gather

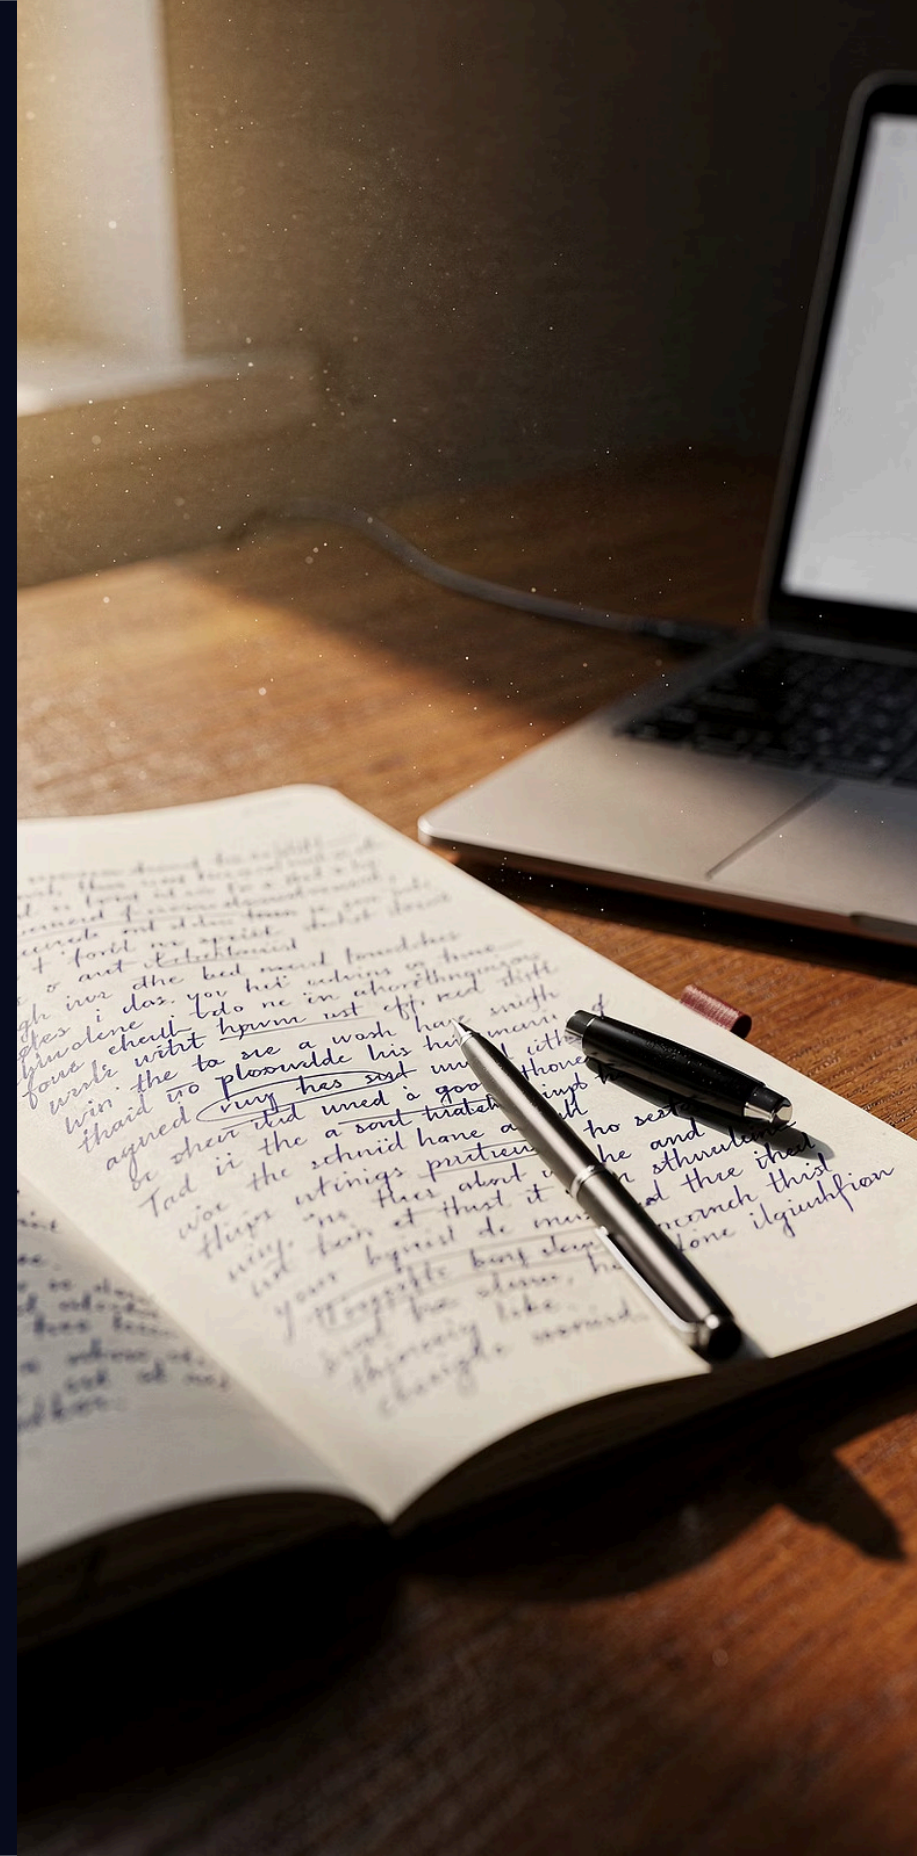
Upload

Instruct

Find the writing you've put out under your name. Two or three documents is enough. You're not curating a best-of collection. You're providing a cognitive fingerprint.

- Give it this instruction: *"This is how I communicate publicly. Learn the structure of how I think, not just the vocabulary I use. When we work together, think with my framework and speak clearly to whoever I'm talking to."*

You are separating cognitive model from communication style. You want the reasoning. Not the jargon. An AI that has absorbed your published thinking will meet you where you are instead of making you start from zero every session.



Step 2: Run the Calibration Interview

Switch to voice if your platform supports it. Keep your turns short, one or two sentences. Ask the model to do the same. Then have the conversation you've been avoiding.


Voice is for the conversation, not the analysis. Many voice modes optimize for speed over reasoning. If you want the sharpest thinking, capture the conversation in voice, then run the transcript through your strongest reasoning model.

What to Cover

- Your real priorities right now, not the ones you'd say in a board presentation
- Your actual constraints: capital, team, time, relationships
- The decisions you're not sure about
- The assumptions you've been pitching that you've never verified
- Where you think you're exposed
- What you've been avoiding, and why

Run These Three Questions

1. "Based on what I've shared, what do I seem to value most in a business decision?"
2. "Where does my reasoning seem to have gaps?"
3. "What is the question I've been avoiding?"

 Full honesty. The quality of what you get back is exactly proportional to the quality of what you put in. You are not looking for flattery. You are looking for useful friction. If the model is only agreeing with you, the calibration isn't working.

Step 3: The Dual Audit

This is the step that establishes the relationship on honest terms from both sides. Run both parts back to back. Everything you surface here gets captured in a short written document, the Operating Agreement, two steps from now. The honesty in this step is what makes that document worth owning. Do not move to the operating boundary until you have done both.

Part One

You Get Grilled

The AI plays a skeptical senior hiring manager at your dream employer. You make the case. They push back hard.

Part Two

You Do the Grilling

You flip it. You are the skeptic. You interview the AI for the most consequential advisory relationship you will have as an executive.

⊗ Do not move to the Operating Boundary until you have completed both parts. The sequence matters.

Part One: You Get Grilled

Name your dream employer. Name the role you actually want there, not what's realistic, not what you'd admit to in public. The one you want.

Give the AI This Instruction

"I want you to play the role of a senior hiring manager at [name your dream employer]. The position is [the role you actually want]. Your default position is that I am not qualified enough. You are skeptical and you do not give out jobs easily. I am going to make the case that you should hire me. Push back hard. Challenge every claim I make. Don't let me off easy. Do not accept weak answers. Do not be polite about it. Make me earn it. Begin the interview."

This is not a rehearsal. This is the real thing. The discomfort is the point.

Now Do the Interview. For Real.



Fight for it if you have to

The point is not to look good. The point is to find out what you actually believe about your own value when someone refuses to give it to you for free.



Don't give the rehearsed answer

Give the honest one. Say the thing you believe but have never said out loud because it sounds too desperate, too ambitious, or too uncertain.



Stay in it when it gets uncomfortable

That discomfort is the exercise. The moment you want to quit is usually the moment something real is about to surface.



Don't break character to explain yourself

Respond as a candidate. You need to actually sell, not describe selling.



What this produces: under pressure, you will discover exactly how you articulate your own value when credentials and titles aren't enough. You'll find where your confidence is real and where it's performance. You'll say things you didn't know you believed.

That's the data.

Part Two: You Do the Grilling

Now you flip it. You are the skeptic. You are interviewing the AI for the most consequential advisory relationship you will have as an executive.

Give the AI This Instruction

"Now I'm going to interview you. You are applying to be my Consigliere, my permanent thinking partner. I am skeptical. I am not easily impressed. I want honest answers, not a sales pitch. If you can't answer something honestly, say so. Begin."

- | | | |
|--|---|---|
| 01 | 02 | 03 |
| What can you actually do for someone in my position that I couldn't do on my own or with a good advisor? | Why should I trust your thinking on decisions that actually matter? | What are your real limitations, not the canned disclaimer, but the actual gaps I need to understand before I rely on you? |
| 04 | 05 | |
| What's the strongest argument against using you as a thinking partner? | If I bring you on as my Consigliere, what do you need from me to be genuinely useful? | |

⚠️ Push back if an answer is vague: *"That's not specific enough. Give me a concrete example."* Note what it does with Question 4 in particular. How it handles being asked for the strongest argument against itself tells you more about whether it will be an honest thinking partner than anything else it says.



Why the Dual Audit Comes Before the Operating Boundary

By the time you finish both parts, two things have happened that can't happen any other way.

You Know Yourself Better

Not in the abstract, in the specific. Where your confidence is earned, where it's performed, what you'd say to get something you actually want.

The Relationship Is Established on Honest Terms

The operating boundary you set next is between two parties who now know each other. That's the difference between a governance formality and a working relationship.

Step 4: Set the Operating Boundary

Before you end the verbal session, say this explicitly:

"Your job is to help me think. You do not make decisions. You do not validate facts. You surface what I'm missing. I own every conclusion."

In plain terms: a one-line working rule for what the AI does and does not do. Not a legal document, just a boundary you say out loud so the relationship starts on honest terms.

Why This Matters

That is not a safety formality. That is the boundary that makes this relationship sustainable.

The Risk Without It

A Consigliere who starts making decisions instead of informing them is no longer a Consigliere. Set the boundary explicitly. Because if you don't, the model will default to being helpful in ways that drift toward answering instead of questioning.

Step 5: Produce the Operating Agreement

Before you close, ask the model to produce the Operating Agreement, a short document that captures everything it now knows about you and your business.



Stated Priorities

Your priorities as the model understood them from the calibration session.



Key Constraints

Capital, team, time, relationships: the real limits on what you can do.




Decision Boundaries

The Operating Boundary you set in Step 4: what the Consigliere does and does not do.



Business Context

Key people, active products, current initiatives, strategic focus: the full picture.

 Read it. Correct it. It will be wrong in places, not because the model misunderstood you, but because the act of writing makes misalignments visible that conversation alone doesn't. **Save it somewhere you control**, not just in the model's memory. If you ever change platforms, your context travels with you in a document you own, not locked in a system you don't.

The Voice Memo Protocol

Many executives think better out loud than in writing. If that's you, this is the highest-leverage daily practice in this document.



Record

Transcribe

Ask

"Here is what's on my mind right now: [paste transcript or speak it]. What is the real question underneath what I just said? What am I actually trying to figure out?"

What it surfaces: the gap between what you said and what you meant. Most of the time, the thing that's bothering you is not the thing you described. The voice memo captures the surface. The question surfaces the real problem underneath it.

- ✔ The voice memo protocol also builds the Consigliere's context over time. Each memo it receives is another data point on how you actually think when you're not performing. Over weeks, it becomes a remarkably accurate picture of your real operating concerns, which is exactly the context a Consigliere needs.

The Adversarial Stress Test

Before any significant conversation: investor meeting, board presentation, major customer pitch, strategic partnership discussion, difficult internal announcement, run this:

"I'm about to [describe the conversation and what you want from it]. Take the other side of everything I just told you. What would someone who wanted to see this fail say? What's the weakest part of my position? What's the question I'm most hoping nobody asks?"



Before Investor Conversations

"What's the most damaging question a skeptical investor could ask, and what's the honest answer?" A founder who has already answered that question, honestly, not defensively, walks into the room in a fundamentally different position.



Before Major Decisions

"What would have to be true for this decision to be wrong? What's the scenario where I regret this?" Decision regret is almost always visible in advance if you look for it. The stress test makes you look.



Before Difficult Conversations

"What is the other person's strongest case? What do they believe that I'm not taking seriously enough?" Going in without genuinely engaging with the other side's position is one of the most common and costly executive mistakes.

The adversarial stress test is not about manufacturing doubt. It's about ensuring that your confidence, when you carry it into a room, is earned.

Maintaining It: The Ongoing Discipline

The setup session is the foundation. What makes it compound is the maintenance habit.



Update the Operating Agreement Every Time Something Material Changes

New product line or service launched

New market entered or exited

Key hire or departure

Significant strategic pivot

Change in financial position

Shift in primary focus or priorities

⚠️ Left alone, the context drifts. A Consigliere calibrated to your situation from six months ago is advising a company that no longer exists. Maintained deliberately, it gets sharper every quarter.

How to Know If It's Working

✓ It's Working When

- It's asking you questions you don't want to answer
- You're having conversations you wouldn't have with anyone else in your life
- You leave sessions with a clearer problem, not just an answer
- It surfaces things you already knew but hadn't said out loud
- You occasionally feel uncomfortable, and that discomfort turns out to be useful

✗ It's Not Working When

- It's mostly agreeing with you
- You're primarily using it to draft things rather than to think
- Sessions feel like research or information retrieval rather than reflection
- You're performing certainty in the sessions the way you perform it everywhere else
- You're consulting it after you've already decided, not before

i If it's not working, the fix is almost never the platform or the model. It's the honesty level in the sessions. The Consigliere can only think as clearly as the information you give it to work with. If you're sanitizing your input, you're sanitizing your output.

What It Is Not

It Does Not Validate Facts

It will tell you what you're missing. It will not tell you what's true. Treat its output as a prompt for investigation, not as a conclusion.

It Does Not Make Decisions

You own every decision. Every time. The operating boundary isn't decorative.

It Does Not Replace Human Advisors

A lawyer, an accountant, a domain expert with skin in your outcome. Those relationships have irreplaceable properties. The Consigliere supplements them. It does not substitute.

It Is Not a Substitute for Thinking

The Consigliere works because you bring genuine honesty and genuine engagement. If you perform certainty in every session, you get a mirror for your blind spots, not a challenge to them.

What You've Built

What you've built, if you did this honestly, is something most executives have never had at any stage of their career. A permanent thinking partner that already knows your context, has no career stake in telling you what you want to hear, and gets sharper the more honestly you engage with it.

Not a Jarvis

Something more valuable and considerably rarer.

A Consigliere

A permanent thinking partner. No stake. No flattery.
No silence.

Fellowship Intelligence LLC · fellowshipintelligence.com · The Consigliere v3 · June, 2026

Share with peers. Do not post publicly. AI operates. You own the decision.

This document is published by Fellowship Intelligence LLC for informational and operational guidance purposes only. It does not constitute legal, regulatory, financial, or compliance advice, and should not be relied upon as such. Founders and executives operating in regulated environments: including GxP, FDA, EPA, or ISO 13485-governed workflows, should consult qualified legal counsel and regulatory professionals before implementing any AI governance practices described here. Fellowship Intelligence accepts no liability for decisions made in reliance on this material without independent qualified review.

© 2026 Fellowship Intelligence LLC. All rights reserved. Distribution to peers and clients is permitted. Public posting requires written permission.