

Collision Chambers

Abstract

Meaningful coordination between two people usually requires disclosing private context—intent, constraints, weakness, desire—before either side knows whether the other can use it well, and that exposure often cannot be taken back. Discovery and disclosure are bound together, so the connections that matter most are the ones people can least afford to test. Existing systems either force premature disclosure or fall back on reputation and credentials, which describe who someone has already been rather than whether two specific parties fit on something they cannot yet say out loud. A Collision Chamber breaks that bind. Each party discloses privately to its own wing agent; governed negotiator agents test compatibility under policy inside a sealed environment; and a privacy-filtered channel releases only a bounded map and a verified next action, never the raw context that produced it. The result is discovery before disclosure: a way to learn where fit is real, where caution is warranted, and what to ask next, without either side having to go first. In doing so the chamber makes a new unit tradeable—qualified possibility, proof that alignment exists without the context that proves it—and settles it through an escrow of intent, the basis for a pre-trust coordination market.

1. Opening: The Latent Possibility Between People

What is possible between two human beings is latent. Until now, you could not see what a relationship might become for you, or how it could bend your trajectory. The society we live in is a dense grid of people, and over time we mutually built trust, through systems, through credentials, so that together we could reach something great none of us could have reached alone.

We build the protocol that makes the most important discoveries possible. An antechamber for finding what is possible between you and another person, the win-win that actually holds. We help you understand what you truly want, surface it safely, and see what two people could become together.

We thus arm you with sharper foresight, and ultimately the means to exercise your true agency.

2. The Core Bind

Discovering¹ what is possible between two people demands disclosure². To build something that works for both sides, to deepen a bond, resolve a conflict, reconcile two readings of the same situation, you have to state what you are actually looking for.

But we rarely do. Think of two people early in dating. Neither says what they truly want. Not because they don't know, but because we are human, and afraid of the unknown: afraid of wanting more than the other, of being seen too clearly, of naming a hope that might not be returned. So both hedge. They send the safe message, read the safe message, and the relationship advances on a fraction of what is actually there. Yet the very thing they are avoiding, saying it plainly, is what would deepen it. Disclosure is the cost and the unlock.

3. Why Now

Now that intelligence is commoditized, that measure of manforce value is weakening. The outdated equation of credential equals capability is dissolving, because a title was only ever a proxy for knowledge that used to be scarce. Transformers let anyone summon accumulated knowledge on demand. The economic value of performing accumulated knowledge has collapsed. Value is shifting elsewhere.

So where does the value go next? It moves to what remains scarce. One person can now accomplish the work of twenty roles at a fraction of the cost. As knowledge labor becomes abundant, the bottleneck shifts from

¹**Discovery**: testing whether meaningful alignment exists between parties.

²**Disclosure**: revealing private context before knowing whether the other party can use it well.

knowing to connecting: the right partnership, the right conversation, the deal, insight, or relationship that only two specific people could create together.

Yet our economic infrastructure was never built to discover or transact that value. Money rewards outcomes after they happen, but it cannot reveal unrealized opportunities between people. Those opportunities remain latent because what each person uniquely carries—intent, constraints, aspirations, complementary strengths—remains largely invisible until after costly disclosure.

That is the source of today's growing sense of drift among knowledge workers. It is not that human value has disappeared. It is that the value which remains has no efficient way to become visible, discoverable, or economically legible.

4. The Problem

To be found at all, you must first make yourself discoverable, under an uncertain state of hope: you have to surface enough of yourself to be seen, before you know whether seeing leads anywhere. But we humans lack the means to sufficiently express complex intention, so today this runs on manual shortlisting and slow, lossy messaging.

What worsens this situation is that more people now hand early outreach to large language models that barely know the person they are writing to, and the other side answers by running agents to filter the noise. Both read each other through platforms thick with implicit agents and uncertain authenticity. We feel devalued, and it erodes the very relationships and possibilities it was meant to create.

So both sides are squeezed from two directions. Stating it carries real risk: disclosure can leak strategic intent and reveal weakness, expose legally or commercially sensitive information, or surface personal context one would rather keep private. But looking has a cost too, so both sides under-search. They act on far less than actually exists, and route private, fragile context through brokers and platforms that then hold it.

For autonomous agents, the stakes climb higher. An agent may carry its user's goals, confidential business logic, threat intelligence, or unpublished hypotheses, and sending that to a central service turns discovery into surveillance. Today's agent architectures were never built for this. Running on architectures not made for the scenario, they are exposed to prompt-injection attacks and data exfiltration. So both parties withhold their critical information, which drives false positives and missed matches, and pushes everyone back into slow, manual, low-quality communication that does not work.

This limits our agency. The set of options a person can act on is bounded not by what exists, but by how much they can afford to discover it.

5. Existing Systems Fail

Collision Chamber is not merely a matching platform.

It is for solving one underlying kind of problem: when your intuition tells you that what you need to find can only be found by disclosing what you cannot yet afford to disclose.

A founder testing an investor, two firms probing strategic fit, a researcher seeking a safe collaborator, a PM and an engineer trying to understand each other's values, a couple trying to resolve their conflict: problems with different surfaces, one shape. Each demands you expose intent to discover whether exposing it was worth it. That bind is the whole market. Where trust already exists, or the inputs are safe to show, you don't need one.

No existing primitive fits, because the work is human-based, semantic, and exploratory. Zero-knowledge proofs and secure multi-party computation can settle a fixed question without revealing its inputs. But you must know the exact question in advance, which is the one thing two parties discovering each other don't. Plaintext brokers handle the openness, but only by holding everyone's raw context, which is the exact exposure we remove. The chamber sits between them as a medium: a broker's flexibility under a cryptographic primitive's trusted seal.

Existing AI companies optimize for generating better answers from disclosed information. Collision Chambers optimize for discovering what should be disclosed before disclosure is safe.

6. The Primitive

Collision³ Chamber⁴ turns implicit, uncertain agent-mediated communication into an explicit discovery workflow.

Instead of letting both sides guess, hedge, outsource messages, and filter each other through noisy platforms, we give each party a private agent that can hold their context, constraints, goals, and risk boundaries. The agents do not simply generate messages. They model what each side is actually trying to discover.

In a Collision Chamber, each party shares private context with their own agent. The system does not expose raw information. It compares constraints, goals, intentions, and non-negotiables, then surfaces only the overlap, tensions, and next questions necessary to test whether a real win-win exists.

³**Collision**: the governed comparison of private intent under sealed conditions.

⁴**Chamber**: the trusted environment where that comparison occurs.

The result is a pre-disclosure environment: a place where both sides can discover whether meaningful alignment exists before either party has to overexpose themselves.

This makes relationships start warmer, sharper, and more qualified. A founder enters an investor conversation already knowing where the real fit and tension may be. Two people in a relationship can see the hopes, fears, and constraints shaping each other's behavior. Two organizations can test alignment without prematurely revealing sensitive strategy.

Collision Chamber makes the hidden structure of a possible relationship legible, safely, selectively, and before the opportunity disappears.

7. System Architecture

A Collision Chamber is a trusted execution environment (TEE) where agents acting for different principals, set by the user, can discover overlap without revealing their inputs unless coordination is reached. It is closer to a governed conversation than a database query. Trust here is layered, not singular: no mechanism is relied on alone, and confidentiality is designed to hold even when any one defense fails.

1. **Private Input:** Each agent submits encrypted intent, constraints, or policy rules. Raw inputs are not visible to other agents or to the operator.
2. **Trusted Compute Boundary:** Everything runs inside a TEE, for example Phala's confidential compute network. Remote attestation lets an agent verify the chamber is running approved code before submitting anything, and that integrity is independently verifiable by third parties down to the hardware layer through dstack/Phala attestation. Confidentiality and policy enforcement are not entrusted to any single boundary. They are implemented in depth across multiple layers, including confidential VM, container and Linux-kernel sandboxing, security modules, runtime policy, and end-to-end encryption, so that the seal survives the failure of any one of them.
3. **Hermetic Isolation:** The chamber has no open path to the outside. Egress is restricted to the single attested disclosure channel; no other network route exists by which information could leave after negotiation context is transferred. Everything that enters stays sealed, and nothing crosses except the agreed selective disclosure⁵.

⁵**Selective disclosure:** the minimal output both sides receive after the chamber determines what can safely be surfaced.

4. **Collision Process:** The chamber determines whether private inputs meaningfully collide, from a simple similarity test up to a multi-turn negotiation between agents, by whatever logic the chamber encodes.
5. **Selective Disclosure:** Only the agreed minimum is released: a boolean, a confidence score, a shared topic, or a bounded follow-up. Non-matching inputs stay undisclosed.

8. Chamber Library

Users start from a set of chamber types we provide, like a founder-investor fit chamber or a co-founder formation chamber, choosing the one that fits the situation they are in. They do not build it from parts. The type carries the structure, and they only ever have a conversation with their wing agent⁶ inside it. Underneath, each chamber is a versioned primitive defined by five parts: a schema setting the dimensions that matter for this kind of relationship; a topic map fixing shared vocabulary so both sides mean the same thing by the same word; a wing prompt governing how each side's intent and non-negotiables are privately elicited; a negotiator policy governing how agents use private context without leaking it; and a finalizer that turns the sealed result into the disclosure each side receives, through their own wing agent and context. We author and curate these types ourselves at first. But the library is meant to be open-ended. The same structure that lets us define a chamber is what would let others define one, and over time we expect the set of types to grow beyond anything we would have built alone, as people shape chambers for situations we never anticipated. The kinds of fit worth discovering are not something one team can enumerate in advance, and the library is designed to keep learning what they are.

9. User Workflow

1. **Private Elicitation:** Alongside given context and information files, a wing agent interviews each participant privately with concrete, experiential, and open-ended questions to understand your true intention. It builds a fuller, more honest model of your real preferences, constraints, and non-negotiables than you would give under social pressure. Since this pre-conversation is entirely private to you, you can reveal true thoughts, feelings, and needs in order to form strategies that other forms of communication cannot provide.

⁶**Wing agent:** the private representative that elicits and protects a user's context.

2. **Governed Negotiation:** Two negotiator agents⁷, each loaded with their principal's private context, hold a real multi-turn conversation, probing compatibility and surfacing tensions while disclosing only what serves the joint process.
3. **Evaluator Governance:** A separate evaluator agent monitors in real time for failure modes such as coercion, deception, looping, premature convergence, or one side dominating, and can steer or halt the process. It functions as the chamber's immune system.
4. **Structured Synthesis:** The chamber produces an aspect-by-aspect read: where alignment is strong, where the data is too thin to score, where caution is warranted, and what the highest-value next conversation would be.
5. **Privacy-filtered Output:** A translator layer turns the synthesis into guidance for each side and is structurally prevented from leaking transcript details or raw evidence. The final disclosure channel, the one place anything crosses, is hardened with the discipline of safe tool calling: schema generation and data-filling run as strictly separate phases, every schema is validated under strong enforcement, and each outbound text is independently checked by clean-context subagents that never saw the raw negotiation. A negotiator agent that has been manipulated or socially engineered therefore cannot smuggle hidden context out through the output, because the channel that releases information was never exposed to the information it might leak.
6. **Verified Next Action:** A chamber is not merely an advice-generating tool. It is complete only when it produces a next move and learns whether the user took it. The synthesis always resolves to something actionable: the exact question to ask, the safe summary to disclose instead of the raw fact, the agenda for the call worth having, or the clean reason to walk away. The report functions as the receipt; the changed action is the product. The wing agent then closes the loop, later asking whether the move was taken and what came out of it. This is what keeps the value inside your agent system rather than leaking out with the match.
7. **Optional Follow-up Turns:** A Collision Chamber can run in multiple rounds if additional user input is needed. Users do not need to provide feedback for individual chamber runs in their discovery network, but get aggregated follow-up clarifications in conversation with their wing agent.

⁷**Negotiator agent:** the chamber-side wing agent that tests compatibility under policy.

It is engineered to compress weeks of staged disclosure, due diligence, and careful self-representation that two parties normally perform for each other into hours of governed agent conversation, removing the performance tax while keeping each side's raw context sealed.

10. Example Chamber Run: Founder–Investor Fit

A founder is raising a pre-seed round but cannot safely reveal that they have four months of runway, would accept a lower valuation for a high-conviction operator-investor, and are afraid of losing strategic control. An investor is interested but privately worried about technical execution, founder maturity, and whether the market is too early. In a normal process, both sides perform confidence. The founder hides urgency. The investor hides doubt. The call becomes theater.

In a chamber, each side discloses privately to their own wing agent. The chamber tests overlap: check size, risk appetite, founder support needs, valuation flexibility, control boundaries, market thesis, and trust conditions. The output does not reveal the founder's runway or the investor's private doubts. Instead, each side receives a bounded map: where fit is real, where caution is warranted, and what exact question should be asked next.

The chamber might tell the founder: “This investor is thesis-aligned but needs evidence of execution velocity. Do not open with valuation flexibility. Ask whether they prefer to support category-creation risk or traction-first risk.”

It might tell the investor: “The founder is unusually mission-convicted but sensitive to control loss. Do not push governance early. Ask how they want strategic help to appear without undermining agency.”

If both consent, the chamber releases a shared agenda: “Discuss category timing, execution proof, and the role of investor involvement.” The raw context remains sealed.

11. Network Formation

A protocol is the destination, not the launch. On day one there is no network, and no reason for a stranger to route sensitive intent through an unproven seal, so the question is not “who is the market” but “who gets value from a single chamber by inviting someone to collide.”

The first network effect is single-player, and it is the one that works with zero liquidity. A user who runs many chambers of the same kind compounds value entirely on their own, before any counterparty has built a side.

So our beachhead is not a relationship type but a frequency: the high-stakes repeat discloser, who saves real time, equity, or capital each time.

The second network effect begins when the counterparty, satisfied with the outcome, builds their own side, and the connection becomes a durable channel rather than a one-time match. This is the real answer to disintermediation. People walk away from a matchmaker once the introduction is made, but not from a standing channel they keep transacting over. What we retain is not the verdict "you align," which is trivially bypassable, but the governed conduit it opens and the ongoing diagnostic it keeps producing, as different needs emerge for different phases of a relationship.

The third effect is the library. Once a few chamber types carry liquidity, the set of types becomes something others compose against. Chambers function as nodes, and a counterparty building their side is federation rather than fragmentation. We provide a vertical chamber with single-player value first, the shared channel that earns retention next, and the open library that makes it a standard last.

What keeps the standard ours, once the library is open, is not the code, which forks, but the two things a fork cannot copy: the populated graph of attested agents with live channels already running between them, and the credibly-neutral position of being the venue both sides trust precisely because it favors neither. Public-reputation networks coordinate on who you are and what you have done. They cannot coordinate on bespoke, forward-looking fit that two parties cannot yet say out loud. That is the ground only our sealed chamber can hold.

So the customer base is, first, the high-frequency repeat discloser who can generate value alone; then their counterparties, pulled in by use and converted from subjects into nodes; then the people who shape new chamber types on a library that is already populated by the time it opens. And over time, the system stops depending on you to know who to collide with. Once the graph is dense enough, fit becomes discoverable rather than introduced. The protocol can surface that a worthwhile collision exists with someone you have never met, and prove it without exposing either side until both consent to proceed. This stays true to the seal: the discovery itself runs chamber to chamber, pairwise and attested, rather than through any central index that can see everyone at once. The early product helps two people you already chose discover what is possible between them. The network it grows into helps you discover the people themselves.

That is the open horizon. A discovery network for fit, where latent opportunity finds its counterpart before either party has had to make themselves vulnerable to find out. What that becomes is not fully ours to script. A

researcher could find the one collaborator on earth working on the same unpublished problem, without either of them publishing it. A company could discover an acquirer, or an acquirer a target, before a single banker is told. Someone could find the person they were always supposed to build with, or love, and learn it safely before either had to risk being the first to say so. The same primitive, pointed at whatever two people most need to find in each other and most cannot afford to say first.

12. Market Creation

Collision Chambers don't just make discovery safer. We make a new thing tradeable, and we create the trusted mechanism to settle the trade.

Start with the unit. Today the tradeable units are people, capital, and raw information, all of which demand exposure before anyone knows the exposure was warranted, with no guarantee it pays off. The chamber introduces a different unit: qualified possibility⁸. Not a person, not their data, but proof that meaningful alignment between two parties exists, generated without revealing the private context that produced it. This is the value the opening called trapped: the deal, the partnership, the bond only two specific people could make, finally legible without being given away.

The unit is only as real as the change it produces in reality, so our goal is to measure that directly. Not the number of reports generated, but counterfactual coordination: decisions, disclosures, and deals that happened because the chamber existed and would not have otherwise. This is the very calibration a privacy architecture would normally make impossible. We retain no raw context, yet with consent we can learn which predictions held, and that calibration compounds with every run while the inputs stay sealed.

The mechanism that settles it is escrow. Both sides deposit sealed intent. Neither sees the other's. The chamber holds both, tests whether they collide, and releases only when a real match is confirmed and both sides consent to proceed. Nothing crosses until the condition is met, and only the agreed minimum crosses. This is escrow in its truest form, not of money, but of intent, and it is the bridge our opening said was missing: the point where money finally reaches the place real impact between people is made.

Pricing follows from this in two layers.

⁸**Qualified possibility:** a relationship, deal, collaboration, or next action that has been tested enough to justify further exposure.

The first is access: a subscription to run discovery at all, tiered like any platform, for individuals, businesses, and custom enterprise deployments. It prices entry to the chamber and scales with usage and scope, and it touches no one's private context.

The second is settlement, and it follows the logic of escrow exactly. We do not charge a person for existing, and we do not price anyone by net worth. To do so would require extracting the very context the chamber exists to seal, and would quietly turn the human back into the unit this entire argument refuses to sell. Instead the fee is taken on release, against the opportunity the release unlocks, sized to the stakes both parties have already defined, because defining those stakes is the reason they entered. Value flows at the instant it becomes legible, and not before.

The deeper point is that a market only exists with liquidity. It is not yet a market; volume is what turns one into the other. Density of coordination is what makes any unit tradeable at all, and our task is to reach the point where qualified possibility clears reliably enough that latent opportunity has somewhere to flow. What we are building toward, if it works, is a pre-trust coordination market, a place where the most valuable connections between people can be discovered, qualified, and settled before trust exists, and before the opportunity disappears. The opportunity never has to become public to become priceable. We are not making humans tradeable. We are making the possibility between humans discoverable, qualified, and, at the moment both sides agree it is real, exchangeable.

13. Bottom Line

A Collision Chamber makes the latent possibility between two people legible, qualified, and exchangeable. By engineering a secure environment for discovery before disclosure, it fundamentally shifts the boundaries of human cooperation. Your agency is no longer bounded by what you can afford to risk. It is bounded only by what actually exists, allowing true intent to find its match, entirely your own.