

How Shannon Works

The intelligence underneath your engagement: what it reads, what it scores, and why it gets sharper every month you work with us.

Everything FBED Labs delivers to you runs on one system. It is called Shannon, after Claude Shannon, the mathematician of information: the job is managing uncertainty in commercial conversations. This document explains, in plain language, what Shannon is, how the three instruments you will meet (RWV, ICPX, and Cascade) fit together, and what happens to your data. You will receive documents, scores, and analyses throughout your engagement; this is the one that explains where they all come from. Keep it. It will still be accurate in a year, and the system it describes will be better.

4 NERVOUS SYSTEMS READ PER CALL	8 BUYER STATES TRACKED	107 SCORING MECHANISMS	1 CLOSED LEARNING NETWORK
--	----------------------------------	----------------------------------	-------------------------------------

OWNER: J. BARDSLEY, CEO

CLASSIFICATION: CLIENT · UNIVERSAL ISSUE

SIBLING: FBED PRODUCT CANON v1.0

HOW SHANNON WORKS · v1.0

REVISE ON CHANGE

01 What Shannon is, and what it is not

Shannon is a domain-specific intelligence system built from real commercial conversations and the outcomes that followed them. It did not start as software. It started as a decade of live selling, tens of thousands of recorded conversations, each one carrying the moment trust formed (or failed to form) between previously unconnected people, cross-referenced against what actually happened to the deal. The pattern in that material was unmissable: conversations that closed exhibited the same trust mechanic, and conversations that died exhibited the same set of failures. Shannon is that mechanic, derived mathematically and then trained against outcomes.

Two things are worth being precise about. First, the predictions you will see are grounded in FBED's own deterministic behavioural mathematics, validated against real deal outcomes; the language you read is the explanation layer over that mathematics, not the source of it. Second, Shannon is not a chatbot, a content generator, or a general-purpose AI wrapper. It is narrow on purpose. It does one job: reading how commercial trust forms, and telling you what to do about it.

THE BOUNDED AGENT

Shannon does not fabricate claims, does not return a score without evidence, and does not act on ambiguity. Every score carries its reasoning, grounded in something observable. Where the system is uncertain, it says so and escalates to a human with the reason named.

02 The four-brain model: why deals actually die

Every commercial conversation involves four nervous systems, not two. You bring a cortex (your plan, your framework, your questions) and a limbic system (your urgency, your need, your nerves). Your buyer brings the same pair: a cortex that states criteria and compares options, and a limbic system that decides before the criteria are ever stated.

Deals do not die in the cortex. The evaluation looks fine, the process is followed, the CRM stage is correct, and the deal dies anyway, because a limbic signal fired that nothing was scoring. The polite "sounds good" that closes nothing. The thread that goes quiet. The nodding that is compliance, not buying. Shannon scores all four systems, which is why it sees the deal turning before the pipeline report does.

Across a conversation, Shannon tracks the buyer through eight behavioural states, from first alertness through curiosity, threat, validation, and hope, to trust, readiness, and commitment. The state is not a label; it is a position on a trajectory, with a direction. Most of what Shannon will tell you to do next is about moving that trajectory, not about what to say.

03 RWV: what your score means

RWV is the measurement framework, and it begins with a distinction that drives everything we will do together: the difference between what is **latently** true about you and what is **reachably** true. Your expertise, your results, your testimonials, your authority: these may all genuinely exist. The question RWV asks is whether a stranger, arriving cold (a buyer, a researcher, or increasingly an AI agent doing the buyer's homework), can actually find and verify them. The gap between latent and reachable is where qualified demand quietly leaks away. You can be excellent and invisible at the same time, and the market will treat you as invisible.

Your RWV report scores both sides of that gap across the components of your commercial surface: your influence signals, digital presence, trust assets, social proof, offer clarity, transaction experience, onboarding visibility, and authority position. Where the gap is wide, the report names it and the engineering brief tells you (or us) exactly what closes it. When you are rescored, you are measured against your own baseline: the score is a delta instrument, built to show movement.

04 ICPX: who you should be talking to

ICPX scores populations the way Shannon scores conversations. It maintains your list (we call it the Bucket): every individual in it gated first against your ideal client profile, then scored for how reachable and winnable they actually are. People who fail the profile gate are not pursued regardless of how attractive they look; people who pass are ranked, so effort lands where conversion is likeliest.

The part that matters most is the refinement loop. Every completed conversation generates conversion signal: did the next call happen, did pipeline emerge. ICPX feeds that signal back into the scoring. A static list decays from the day it is built; your Bucket does the opposite. Over the months of your engagement it becomes a progressively higher-fidelity reflection of what actually converts in your market, which is knowledge neither you nor we could have written down on day one.

05 Cascade: reading every conversation

Cascade is the surface where you meet Shannon directly, and it works in three passes around every conversation. **Before:** preparation built from what is publicly knowable about the person you are meeting, your own positioning, and patterns from comparable conversations. **During:** structure, so the conversation follows the sequence that converts rather than the sequence that feels natural. **After:** a reflection pass: what worked, what did not, where the trajectory points, and what to do next. The reflection feeds the next preparation, which is the small loop inside the big one.

A Cascade read is specific. It will not tell you the call "went well." It will tell you where the buyer pulled back, which commitment is live and which is slipping, and what the next move is. The first time you receive one on a call you remember differently, you will understand the product.

06 The loop: why month twelve is smarter than month one

Each instrument feeds the others. ICPX decides who; Cascade reads what happened; the outcome refines ICPX; RWV measures whether your commercial surface is converting the attention all of it generates. Underneath, the model trains continuously on real outcomes across the whole network. This is why the engagement compounds rather than repeats: the system you are using in month twelve has learned from everything that happened since month one, including things that happened outside your engagement entirely.

THE CLOSED NETWORK PRINCIPLE

Shannon learns from every conversation across every FBED Labs engagement. Learnings are anonymised and aggregated. No client-identifying information, no conversation content, and no specific intelligence outputs ever cross between engagements. The network learns; your engagement stays yours.

07 Where you will meet Shannon

- **Your document shelf.** Your RWV report, your ICPX deep dive, and your playbooks are all Shannon-built. As you add context and assets, the shelf fills; the documents are living, not snapshots.
- **Your conversations.** Every recorded conversation in your engagement receives the three-pass Cascade treatment.
- **The Live rooms.** Podcast Live and Cascade Live run every two weeks: live analysis and live sales work with Josh and Shannon, on real data, for every client in the line.
- **Your reviews.** Your scores, your Bucket health, and your trajectory are the standing inputs to every review in your operating rhythm. Decisions in this engagement are made against measurement, including the decision to continue it.

One operating rule sits under everything above, and it is ours as much as the system's: nothing here asks you to take the read on faith. Every score arrives with its reasoning, every recommendation with its evidence, and every rescore against your own baseline. If a number ever appears without its why, ask your named CSM for it; producing the why is the product.