Decision Memo — CLEAR's Borrower-Facing Mortgage Platform

To	CLEAR Home Loans leadership
From	Christa Votaw, National Sales Leader
Analysis by	**Willis** — Christa's AI research and analysis assistant
Date	August 9, 2026
Decision requested	Authorize a two-week decision sprint. **No build commitment today.**

> Every finding, figure and recommendation in this memo was produced and independently source-verified by **Willis**, and is labeled as such throughout. Christa is bringing it forward for a decision.

The ask, in one paragraph

CLEAR's loan officers each run their own copy of Floify. It works. It is also the same tooling every competitor has, it prices **per seat** — which means the bill grows in direct proportion to how well our recruiting strategy works — and it cannot do the two things that would actually distinguish us. **I am asking for two weeks and a small amount of leadership attention** to determine whether we buy, build, or do something in between. **I am not asking for a build commitment today**, because Willis's review found that the information required to make that decision responsibly does not yet exist.

What we'd be trying to win

Two capabilities, neither available from any vendor at any price:

1. The borrower's own loan officer, present from the first tap. An AI assistant carrying the LO's name, photo and voice, docked on every screen, answering the questions borrowers get stuck on at 10pm — what's a gift letter, why do you need two months of statements, where am I in this process. It never quotes a rate, never implies an approval, never gives advice. It explains, reassures, and hands off to the real human the instant anything touches loan terms. **No competitor has this.**

2. A realtor experience that doesn't require a realtor to log in. Every lender builds a realtor portal. Realtors don't use them, because a busy agent will not maintain a password for one lender's dashboard. Our approach pushes updates to where agents already are — a text they can forward to their client with zero clicks — and offers a portal only as an optional extra. **Realtors are the referral engine of this business and every platform on the market treats them as spectators.**

Both are additive to how we already work. Neither requires replacing the systems that manufacture the loan.

Why Willis says the decision can't be made yet

The specification behind this memo — *The Borrower Experience Atlas* — runs 13 parts and 121 features, and it recommends building. **Willis then re-checked that recommendation against primary sources and concluded it is not yet supportable on the information we have**, for three reasons:

We have no build cost. There is no team size, no schedule and no dollar figure anywhere in the specification. We asked vendors for quotes and never asked our own development team for an estimate with the same rigor. *Comparing a priced vendor to an unpriced build is not a comparison.*

We have no baseline. The specification's targets were anchored to an industry statistic — roughly 68% of online applications abandoned — which Willis traced to a payments vendor's marketing blog about **general credit applications, not mortgages.** Meanwhile our actual completion rate is sitting inside Floify, unpulled. **Without our own number, we cannot prove this worked even if it does.** That correction has been applied to the Atlas.

We framed the choice as all-or-nothing. Build everything, or buy everything. Willis identified a third structure the original document never considered, and it may be the right one.

The option we hadn't considered

Of the 121 features specified, **roughly 96 are commodity** — the application form, disclosure timing, e-signature, document collection, closing. Every vendor has them. None differentiate us. They are also the source of the single largest hidden cost of building: **once we own that code, we own the compliance maintenance forever.** Every TRID revision, every URLA change, every state-law shift becomes our engineering backlog, permanently.

The remaining **~25 features are the differentiators.**

Option C: buy the commodity layer, build only the differentiators. The AI assistant and the realtor push layer both consume loan status through an API. They do not require owning the application form. If a vendor's API exposes what we need — and their terms permit it — we could have the moat in months instead of years, without taking on permanent compliance engineering.

Two questions decide whether Option C works, and both are answerable in two weeks with a handful of vendor calls:

1. **Can our AI assistant be properly grounded in live loan status through each vendor's API?**
2. **Do their terms permit us to put our own assistant inside their borrower experience?**

> **Willis's caveat, stated plainly:** Willis expects the answer to question 2 is **no** — vendors are unlikely to permit a lender's own AI assistant inside their borrower UI — and that question 1 is harder than it sounds, because the assistant's value depends on *condition-level* needs-list state in near-real time, not just loan status. **If that prediction holds, the case for building the whole thing gets considerably stronger.** But a prediction is not evidence, and two weeks buys evidence.

The financial frame, honestly

The most-cited number in mortgage technology is what it costs to originate a loan: **$11,898** as of MBA's Q1 2026 report, and **$12,209** on a fully loaded annual basis for independent lenders in 2025. That is the number any platform investment is measured against. A system that improves how many applications become **funded loans**, even slightly, is measured against roughly $12,000 per loan — not against its own license fee.

On the other side, per-seat pricing:

LO count	Floify @ $79/user/mo	Annual
~25 today	$1,975/mo	**$23,700**
100	$7,900/mo	**$94,800**
300 (recruiting target)	$23,700/mo	**$284,400**

And that is the **entry-tier document tool.** Blend and nCino/SimpleNexus don't publish pricing and negotiate per lender; expect materially higher at the same headcount. **This is the growth tax in plain terms — the platform bill scales with precisely the thing our strategy is built to maximize.**

(Willis's note on sourcing: an earlier draft of the Atlas cited Floify at $49/month from a secondary aggregator. That figure was wrong; the correct list price is $79 per user per month, verified against Capterra and GetApp on Aug 9, 2026. The correction makes the argument stronger, not weaker.)

What we still don't know: what a build costs. That is the gap this sprint closes.

What I'm asking leadership to authorize

A two-week decision sprint, producing four things:

Deliverable	Owner
Written quotes from Blend, nCino/SimpleNexus and Floify at 100 and 300 seats	Christa
A costed build estimate — team, months, dollars — against a re-scoped pilot	Development team
API feasibility: can the assistant be grounded through vendor APIs, and do terms permit it	Development team
Our current Floify baseline — completion rate, cycle time, mobile share	CLEAR ops (**can start today**)

Then a one-hour decision meeting with all four numbers on the table.

Cost of the sprint: roughly two weeks of part-time effort across three people. No engineering commitment, no vendor commitment, no contract signed.

Cost of skipping it: we commit to a multi-quarter build on a preference rather than a comparison — which is the specific failure pattern that has cost us before.

What happens after

Whatever the decision, the rollout is the same and it is deliberately conservative:

1. **Measure first.** Pull the Floify baseline before anything is built. Without it, success is unprovable.
2. **Pilot with 3–5 volunteer loan officers.** Chosen for enthusiasm, not seniority. At least one high producer and one recent hire.
3. **Run in parallel.** Pilot LOs keep Floify for loans already in flight. No borrower is ever migrated mid-loan — that is where data loss and compliance gaps happen.
4. **Expand only on the number.** If the pilot doesn't beat the Floify baseline on measured **pull-through**, we fix the product. We don't mandate adoption.

The pilot scope is roughly **31 features, not 121.** It exists to answer one question: *does the AI assistant plus the realtor push layer produce more funded loans than what we run today?*

> **A metrics correction Willis made to the Atlas, worth one line here:** the original specification's north star was *application completion rate*. That is the wrong goal — a system optimized for completion rate is optimized to let **anyone** finish an application, which raises LO workload and adverse-action volume while lowering what actually pays for the system. The measures are now **pull-through** and **cost per funded loan**, with completion rate retained as a diagnostic.

The risks I want on the record

If we build: we own compliance maintenance permanently. TRID changes, URLA revisions, state-law shifts become our engineering backlog. This is real, it never ends, and it must be staffed and budgeted deliberately rather than discovered in year two.

If we buy: we compete on the same tooling as everyone else, with the same borrower experience and the same conversion rates, and we pay a fee that grows every time recruiting succeeds.

If we do nothing: the status quo — which is working, and which is also the reason this memo exists.

Regardless of path, two items need attention now, because they are slow:

- **Counsel review.** Willis flagged a question the original specification missed entirely: whether an AI assistant that *guides a borrower through completing a loan application* implicates mortgage-originator licensing under the SAFE Act. It is likely defensible — the borrower is completing their own application — but it is a novel posture and it deserves its own written opinion **before** launch, not after an exam. Related questions cover our realtor co-marketing structure and the federal trigger-lead law effective this past March.
- **The knowledge base.** The assistant is only as good as the several hundred approved answers behind it, and those must be written and compliance-reviewed **by CLEAR**, not by the development team. This is on the critical path under every option and can start immediately.

Recommendation

Authorize the two-week sprint. Start the Floify baseline pull and the knowledge-base work this week, since both are useful under every outcome. **Decide in three weeks with real numbers.**

Supporting detail: ***The Borrower Experience Atlas — Revision 2*** *(13 parts, 121 features; supersedes the Aug 9 draft) and the* ***Build Kit*** *— 28 screens, the journey map, the UI kit, the integration map and the filterable backlog.*

All figures independently verified by Willis against primary sources as of August 9, 2026: MBA Quarterly Performance Report Q1 2026; Floify pricing via Capterra and GetApp; Homebuyers Privacy Protection Act via Hunton and McGlinchey. MBA's Q2 2026 data publishes mid-August and should be refreshed before external circulation. ***This memo is a business recommendation, not legal advice*** *— all compliance positions require counsel review.*