

AI Readiness for CRE Brokerages

A step-by-step preparation playbook for brokerages getting their people, process, and data ready for AI

Prepared by CRE OneSource · 2026

WHAT'S IN THIS PAPER

- **The Foundation** — what stays human, why AI runs on data, and what readiness is worth
- **The Playbook** — five pillars, in order, each with an owner and a finish line
- **The Brokerage Advantage** — why a firm's combined data outperforms any single broker's
- **Appendices** — data standards, habit checklists, readiness metrics, and the assessment

Here is a test any brokerage leader can run this afternoon. Point an AI tool at the firm's records and ask it two questions: which deals are at risk, and who hasn't been contacted in sixty days?

The tool will answer quickly, in clean formatting, with names and dates. Then someone who actually knows the pipeline reads the list.

At most firms, the answer is a confident, well-formatted list that is partly wrong — and nobody can tell which part.

That is not a failure of the model. The model did exactly what it was asked with exactly what it was given. The deals marked active that died months ago, the contacts entered three different ways, the calls that happened but were never logged, the prospect sitting in one broker's spreadsheet and nowhere else — all of it went in, and all of it came out as an answer.

Our earlier paper, *CRE Brokerage in 2030*, argued that the administrative half of a broker's week is going away for the firms that pursue it, and that **the tools will be ready before most firms' data is**. That moment has largely arrived. The capabilities exist today; what most brokerages lack is the foundation to use them.

This paper is the how: a specific, sequenced playbook for getting a brokerage's people, process, and data ready for AI, built on five years of close work with thousands of brokers and on published research into what makes data usable. The heart of it is the five-pillar playbook in Sections 7 through 13. It is written primarily for brokerage leaders, but every core principle applies to individual brokers as well, and there is no reason an individual cannot run the same playbook in their own business.

KEY FINDINGS

76%

of CRM users say less than half of their organization's CRM data is accurate and complete (Validity, 2025)

~9 in 10

CRE investors and occupiers piloting AI; 5% report achieving all their AI goals (JLL, 2025)

63%

of organizations lack AI-ready data practices or are unsure whether they have them (Gartner, 2025)

47%

of newly created business records contain at least one critical error (Harvard Business Review)

AI adoption is outrunning the records it depends on. The gap closes when firms set the standards, build the habits, and use software that makes both easier to keep.

Why Readiness Stalls

Most brokerage leaders will agree with the paragraphs above, and most will not act on them this year. The reasons deserve to be stated at full strength.

- **Data entry is unpaid work for the people who generate the data.** Brokers are paid on closed deals, most are independent contractors, and every minute spent in a system is a minute not spent with a client. Asking a top producer to type more is asking them to earn less.
- **At many firms, a broker's book belongs to the broker.** The relationships and history are the broker's life's work, often the recruiting pitch that brought them in, and sometimes written into their agreement. A firm-wide data program can look like a land grab.
- **The cleanup is enormous, and it belongs to no one.** Years of duplicates, stale contacts, and half-entered deals, scattered across an old system, a dozen spreadsheets, and every broker's inbox.
- **The AI will clean it up for us.** Models already deduplicate records, standardize formats, and pull structured data out of emails. Waiting for the tools to do the cleanup looks cheaper than doing it now.
- **Nothing is broken.** Deals are closing, brokers find what they need, and no one can point to a deal lost because a field was blank.

The first is true, and the playbook is designed around it: software now does more of the recording every year, and the broker's role is shifting from typing to approving. The second is a legitimate cultural position, and Section 15 treats it as one. The third is exactly why the playbook refuses to start with a cleanup. The fourth is partly true, but software cannot set standards the firm never set or recover a conversation that never passed through any system. The fifth holds right up until someone asks the AI a question and gets a confident, wrong answer back.

THE FOUNDATION

1. The Client-Centric Broker

Start with the broker, because every other part of this playbook exists to serve them.

The best brokers in this business share a recognizable profile. They are client-facing. They have strong people skills, and they use them: they listen well, read a room, and tell clients things they don't want to hear in a way the client can accept. They are client-centric rather than deal-centric: they think about the client's portfolio over ten years, not the commission on this transaction. And they build deep, long-term relationships, which is why their business compounds through retention and referral.

None of that is going away. As *CRE Brokerage in 2030* argued, the opposite is happening: as document production becomes universally available, it stops being a differentiator, and the human half of the job (relationships, negotiation, market knowledge, trust) becomes the part that is scarce.

That has a practical consequence. A firm cannot protect the human work until it has named it, and most brokers have never separated their week into the work only they can do and the work that merely has to get done. Doing so is the first step of the playbook, because it tells the firm what the rest of the effort is for.

THE UNCHANGED JOB

AI changes how a broker's week is spent, not what the job is. Readiness starts by naming the work that will stay human, so that everything the firm builds gives that work more hours.

2. Habits, Standards, and the Software That Supports Them

Great brokers also have great habits. They communicate proactively. They follow through on what they said they would do, when they said they would do it. They respond quickly. They hold high standards for their own work and notice when something is sloppy. None of these are talents. They are routines, repeated until they stop requiring thought.

The habits that make a broker great are the same habits that produce great data. A logged call is follow-through made visible. A next step with a date on every active prospect is responsiveness written down. A deal stage updated the day it changes is a high standard applied to the data instead of the flyer. The questions a brokerage struggles to answer about its data, such as why the pipeline report is wrong or why no one knows who last spoke to a client, are almost always questions about habits wearing a technology costume.

THE HABIT RECORD

A brokerage's data is a record of its brokers' habits. Firms do not have a data problem separate from their habits problem. They have one problem, visible in two places.

Until recently, that meant the best data came from the best brokers, because keeping data current depended entirely on personal discipline. The average broker — busy, well-intentioned, inconsistent — left gaps, and no amount of training closed them for long.

Software is changing that equation, and it helps to name the software this paper means. A brokerage's **operating system** is the platform that holds the firm's contacts, properties, deals, and activity, and that the rest of the firm's work runs through. It is also called the CRM, the system of truth, or the source of truth; this paper uses "operating system" throughout.

A modern operating system can read a broker's inbox and calendar, recognize a new contact, a meeting that happened, or a deal that moved, and propose the update as a suggestion the broker approves, edits, or rejects in seconds. The direction of travel is toward eliminating manual data entry almost entirely, with the broker's role shifting from typing records to approving them. Software does not replace standards; someone still has to decide what good looks like. But it can make meeting them the default rather than the exception.

THE LIFTED AVERAGE

Top producers already keep good records. The largest gain from software that captures and suggests updates goes to the average broker, whose habits it reinforces every day, and whose data starts to look like the best in the firm.

3. Context Is the Input

An AI model's usefulness to a brokerage depends on two things: how capable the model is, and how much it knows about the specific situation. The first is largely settled. Today's widely available models are capable enough for the large majority of brokerage work, and they improve every few months without the firm doing anything. The second is where firms differ, and it improves only when the firm does.

A general-purpose model knows a great deal about commercial real estate in the abstract and nothing about this market, this client, or this deal. Ask it to draft a follow-up and it writes a competent, generic email. Give it the full history — the three tours, the client's objection about parking, the lease expiration eighteen months out, the competing building they are also considering — and it writes the email the broker would have written. Give it the broker's email history and it picks up the broker's voice and writing style as well. The difference between those drafts is not the model. It is the context.

CRE Brokerage in 2030 called this the context tax: when AI is used one task at a time, the broker pays it by assembling the context by hand, every time, and that assembly is most of the work. The tax

disappears only when the context already lives in a system the AI can reach. The quality of a firm's AI output is capped by the quality of its data, not eventually but from the first day it tries.

A simple self-test: if an AI assistant were handed your contact and deal records today and asked “who should I call this week, and why,” how well would it do? The honest answer is a direct measure of readiness, and it has almost nothing to do with which AI you use.

4. What Clean Means

“Clean data” is said so often it has lost its meaning. Here it means four properties, in order of importance.

Property	What it means	What breaks without it
Consistent	Everyone records the same things the same way: one set of deal stages, one list of property types, one format for names, dates, and sizes.	The AI can't compare what doesn't line up. “Office,” “OFC,” and “Office – Class B” become three categories.
Clean	Correct and current. One record per person, company, and property; dead deals closed out; job changes updated.	Duplicates split one relationship across several records. Stale deals inflate the pipeline.
Thorough	Complete for the fields a decision depends on, not every field.	A deal with no next step looks the same as a deal with no problem.
Centralized	One operating system, not a system plus spreadsheets, inboxes, and notebooks.	What isn't in the operating system doesn't exist to the AI.

Volume is not on the list. A smaller database that is consistent, clean, and thorough beats a larger one that isn't: an AI can reason over the first and can only guess over the second. Firms that respond to “AI needs data” by importing purchased lists are working on the least important property first.

THE SMALLER, CLEANER DATABASE

Five hundred complete, consistent records outperform five thousand partial ones. Quantity is the last property of good data to matter, and the first one firms reach for.

5. Bad Data Scales

BY THE NUMBERS

76%

of CRM users and administrators say less than half of their organization's CRM data is accurate and complete, in a 2025 survey of 602 respondents. 37% reported losing revenue as a direct result of poor data quality (Validity, *State of CRM Data Management*, 2025).

Measured studies point the same way. When executives reviewed the last 100 records their own departments had created, 47% contained at least one critical error (Harvard Business Review, 2017). These were new records, made by people who knew they would be reviewed. Errors are mostly created at the moment of entry, not over time.

Decay adds to it. Contacts change jobs, companies move, and email addresses stop working. Email verification services that process billions of addresses a year report that roughly a quarter of the addresses on a typical list go bad within twelve months.

Before AI, bad data was a tax paid slowly. A broker hit a wrong phone number, sighed, and looked it up. AI changes the rate. **A model fed incomplete or stale records does not produce slightly worse work slowly. It produces confident, wrong work quickly, across every deal at once**, and hands it to the person least likely to catch the error — the one who asked the AI precisely so they wouldn't have to check.

THE 1-10-100 RULE — LABOVITZ AND CHANG

It costs \$1 to verify a record when it is entered, \$10 to clean it up later, and \$100 in consequences if it is never fixed at all.

The rule is more useful for its shape than its exact numbers. Every dollar of effort moved to the point of entry saves ten later.

That is why this playbook puts standards, habits, and capture at the source ahead of any cleanup project. Data quality is decided the moment a record is created. A cleanup fixes the past; only a standard, followed every day, fixes the future.

6. What Readiness Is Worth

Readiness work is not free. It costs leadership attention, some broker patience, and a few quarters of discipline before the results show. The fair question is whether it is worth it. It is, by a wide margin.

CRE Brokerage in 2030 put the document and administrative work in a full-time commercial broker's week at about 24 hours, or roughly 1,150 hours a year. At the \$100 to \$500 effective hourly rates we see for broker and operations work, that is \$115,000 to \$575,000 of the firm's most expensive time, per broker, every year. We projected that a broker on modern systems can bring that load down to two hours a week or less by 2030. None of that recovery is available to a firm whose data is not ready, because every drafted proposal, tour package, and prospect report is only as good as the data it is built from.

What readiness returns	What it looks like in practice
Time back	The document and data work that fills about half a broker's week is drafted, assembled, and updated by software, and the broker approves it. The hours go back to clients.
Revenue that stops leaking	Prospects stop going quiet because someone lost track. In one 2025 survey, 37% of CRM users reported losing revenue to poor data quality.
A pipeline leadership can see	Forecasts come from current data rather than from asking around, so leaders know what is closing and what is at risk.
Faster client service	Tour packages, BOVs, and prospect reports in hours instead of days: the turnaround we expect clients to treat as standard before the end of the decade.
Continuity	When a producer leaves, the firm keeps what its ownership policy says it keeps, instead of starting over.
AI that compounds	Every clean record makes the next AI use case faster to launch and more accurate. Unready firms start every use case from scratch.

The cost of waiting runs the other way. The models improve every month whether a firm acts or not, and they cost roughly the same for everyone. Clean, consistent data does not appear on its own, and building it takes quarters, not days. A firm that waits until AI is obviously ready and only then starts on its data begins a long job at the moment it needed to be finished, while competitors who started earlier are already giving their hours back to clients. And almost all of this work would be worth doing even if AI stalled tomorrow: a firm with documented workflows, current data, and one operating system is a better-run firm today.

THE DOUBLE RETURN

Readiness pays twice. It makes the firm better run now, with cleaner pipelines, better forecasts, and fewer lost prospects. And it is the price of admission for everything AI will do for brokerages over the next five years.

THE PLAYBOOK

7. How the Playbook Works

The playbook has five pillars. They run in order, because each depends on the one before it: a firm cannot write data standards before it knows its workflows, cannot choose an operating system before it knows its standards, and should not layer on AI before the data it will run on is ready.

Pillar	The question it answers	Owner	Done when	Typical time
1. People & Relationships	What work stays human, and how does each broker win business?	Principal or managing broker	Human work is named, the ownership policy is written, and an owner is chosen.	1 week
2. Process & Workflows	Which firm-wide workflows and requirements keep the data current?	Managing broker with operations lead	Workflows are documented and every firm meeting and report runs from the operating system.	2–4 weeks
3. Data	How consistent, clean, thorough, and centralized is the firm's data, and how does it stay that way?	Operations lead	Standards are in use, the active set is clean, and quarterly re-measurement is scheduled.	1–2 months
4. Systems	Which operating system, and which partner, does the firm commit to?	Principal with operations lead	One operating system; side sources retired.	1–2 months, alongside Pillar Three
5. AI	How do we layer AI onto workflows so it scales good data, not bad?	Principal with operations lead	One use case is live, verified, and measured.	30–90 days per use case

The first two pillars are mostly decisions and cost little beyond leadership attention. That order is deliberate: the highest-value readiness work is also the least expensive, and it is the work most firms skip on the way to buying something.

The timeframes are what we typically see at firms that commit to the work, and Pillars Three and Four usually run side by side. Anyone can turn on an AI tool tomorrow, but that does not make it useful, sustainable, or impactful. A committed firm that seriously works through this entire playbook can usually have its first AI use case live, and worth keeping, within three to six months.

Two kinds of ownership work together. The **AI Readiness Owner**, named in Pillar One, owns the playbook as a whole: keeping the pillars in sequence, reviewing the readiness metrics in Appendix C with leadership each month, and deciding when a pillar is actually done. Each **pillar owner**, shown in the table above, does the work of that pillar. At a small firm, the principal may hold every one of these roles. What does not work is assigning readiness to “the team,” or to whoever is best with technology. Readiness fails when it belongs to everyone.

8. Pillar One — People & Relationships

What work will always remain human, and how does each broker build relationships best?

1. **Inventory the week.** Each broker lists the work in a normal week and marks each item client-facing, production, or administrative. The workload table in Section 1 of *CRE Brokerage in 2030* is a reasonable starting template. Most brokers are surprised by the result.
2. **Name the human work.** Negotiating, advising, being in the market, reading the room, delivering hard news, structuring the deal. These are what recovered time is for. Put them on the calendar first.
3. **Name each broker’s relationship engine.** Referrals, cold outreach, a deep sphere, a specialty niche, landlord relationships. Every broker wins business a particular way. The processes and systems that follow should strengthen that engine, not impose a different one.
4. **Make the culture decision.** Leadership decides where the firm stands on data ownership and writes it down, so every broker knows, without guessing, whether they could take their history with them if they left tomorrow. There are three models, each detailed in Section 15:
 - **The broker’s book.** Each broker owns their data and relationships.
 - **The firm’s database.** The firm owns the data.
 - **Joint ownership.** The broker can take their data when they leave, and the firm keeps a copy.
5. **Name the AI Readiness Owner.** One person accountable for the playbook end to end, usually the principal, the managing broker, or an operations lead. Section 7 describes the role.

Done when: every broker can name the three activities that most drive their business, the firm’s data-ownership policy is written and communicated to every producer, and an AI Readiness Owner has been chosen.

9. Pillar Two — Process & Workflows

Which firm-wide workflows and requirements keep the data current without slowing brokers down?

Habits belong to individuals. Workflows and requirements belong to the firm, and they are the firm's main tool for shaping habits across every producer.

1. **Write down the core workflows as actually performed.** A new prospect. A tenant requirement through tours. A listing from signed to closed. A deal from LOI to close. A closing to invoice and commission splits. Step by step, as done today, not as imagined. A process nobody has written down cannot be improved, and cannot be automated.
2. **Set the firm's requirements.** A short list the firm can actually enforce: every meaningful client touch recorded the same day; every active prospect has a next step and a date; deal stages updated the day they change; closed-deal terms entered at close. Appendix B turns these into daily and weekly habits for brokers and leaders.
3. **Let the operating system do the capturing.** A modern operating system can do much of this automatically: syncing email and calendar activity; reading inbound messages and proposing record updates for the broker to approve, edit, or reject; and turning a voice note into a logged meeting. The broker's job shifts from typing records to approving them. Cut any field that no decision depends on.
4. **Run every firm-wide process through the operating system.** The weekly pipeline meeting, monthly reports, the forecast, production tracking: all of it runs from the operating system, not from a spreadsheet compiled from several sources. Every side spreadsheet is a second version of the truth, and a signal that the operating system doesn't need to be current.
5. **Principals and top producers go first.** Brokers watch what the principals and the top producers do, not what anyone announces. Getting top producers on board (visibly using the operating system and leading by example) is the fastest route to cultural change. This cannot be "do as I say, not as I do."

6. **Give the requirements teeth.** Requirements shape habits only when following them matters. The levers below are the ones we see work, and most firms combine several.

Lever	How it works
The commission gate	Commissions are not paid until the deal's required information is in the operating system and every requirement is met. It is the strongest lever a brokerage has.
Firm resources	Marketing support, admin help, research, and lead routing flow through the operating system, so the brokers who use it get the most from the firm.
Meetings and reports	Every meeting runs from the operating system where applicable, and every report comes from it (step 4).
Training and culture	Onboarding, training, and the firm's way of working are built around the operating system from a new broker's first day.
Recognition	Leadership calls out the brokers whose data is most current, in meetings and in front of their peers.
Incentives	Rewards are tied to participation and results in the operating system, such as activity recorded, next steps kept, and data completeness.

THE PIPELINE MEETING TEST

If the firm's meetings and reports run off the operating system, it stays current. If they run off spreadsheets and memory, it becomes an archive.

Done when: the core workflows are documented, the requirements are written and communicated, and every firm-wide meeting and report now runs from the operating system.

10. Pillar Three — Data

How can we centralize our data and maximize its consistency, cleanliness, and thoroughness?

1. **Measure the baseline.** Use the method from the study in Section 5: pull the last 100 records created in each major category — contacts, properties, deals — and have someone who knows the business mark every record with an obvious error. Count the clean ones. It takes an afternoon, and it replaces opinions about the firm's data with a number.
2. **Write the standards.** Definitions first: what counts as a prospect versus a contact versus a client; the firm's deal stages, each with a clear entry point; property types and classes; naming and formatting rules; the required fields for each record type. Appendix A is a starting checklist. If the standards don't fit on a few pages, they won't be followed.
3. **Fix the flow before the stock.** Apply the standards to every new record starting now. That stops the problem from growing, and it is the \$1 end of the 1-10-100 rule (Section 5).

4. **Clean the active set, not the archive.** Next, clean what is in use: live deals, active prospects, current listings, the clients a broker spoke to this year. The ten-year-old contact who has never been called can wait. AI can carry much of this work: flagging likely duplicates and suggesting merges, spotting deals with no activity in months, standardizing formats, and pulling terms from executed documents. A person approves each change.
5. **Assign an owner to every data set.** Contacts, properties and availability, deals, commissions, comps. For each, name who keeps it current and who can see it. This is where the permissions discussed in Section 15 get set.
6. **Re-measure every quarter.** The same 100-record test, the same categories. The trend matters more than the starting number.

THE ACTIVE SET

Do not start by cleaning the archive. Standardize what is new, clean what is in use, and let the rest age out.

Done when: *written standards are in use for every new record, the active set is clean, every data set has an owner, and the quarterly re-measurement is on the calendar.*

11. Pillar Four — Systems

Which operating system should we commit to and invest in for the long haul?

The distinction that matters here is between tools and an operating system. A tool is software that performs a specific task. The operating system is the platform the firm runs on: it holds the firm's data, carries its workflows, and determines how and when every other tool is used. Brokerages that have added tool after tool (email, spreadsheets, a CRM, a deal tracker, research products, a separate back office, now AI) often have more software and a weaker operating system than they did five years ago. More tools do not automatically create a better operating system.

1. **Consolidate onto one operating system, and as few other tools as possible.** Contacts, properties, deals, and the back office work best in one place, because the most valuable questions cross all of them, and because AI is most effective when all of the data and context it needs sits in one operating system.
2. **Evaluate against the workflows written in Pillar Two, not against a demo.** Run each documented workflow through the candidate operating system, start to finish.
3. **Commit to the right partner.** Switching operating systems every two or three years resets the data, the habits, and the firm's patience, so pick one the firm intends to be married to. That raises the stakes of the choice. When AI capabilities change every few months, committing to a system that has stopped improving is one of the most expensive mistakes a firm can make. Commit to a company that shares the firm's values and is continuously adopting new technology, frameworks, and best practices.

4. **Retire the side sources.** Migrate what is active, archive the rest, and set a date after which the spreadsheet is no longer the source of truth.

Question to ask of any operating system	Why it matters
Does it hold contacts, properties, and deals together?	The context AI needs crosses all three. Separate systems mean separate slices.
Is activity captured as a byproduct of work?	Every field typed by hand is a habit that will eventually slip.
Does the deal flow through to the back office?	Re-entry between closing and accounting is where errors and delays enter.
Is AI built where the data already lives?	AI bolted on beside the system will keep asking brokers for context the system already has.
Is the company continuously innovating?	A system that isn't adopting new AI capabilities, models, and best practices falls further behind every year the firm stays on it.
Does the company share the firm's values?	A long-term commitment is a partnership. The firm should believe in where the company is going.
Can the data get out, and can permissions match the firm's obligations?	Export and departure terms follow from the Section 15 decision. Sharing without permissions is a liability.
Where does the data sit, and is it used to train anything?	A vendor should be able to answer plainly.

Every vendor says it innovates. These questions test whether it actually does:

- What AI capabilities has it shipped in the last twelve months, and are customers using them every day?
- How often does it release updates, and do existing customers get them without paying for a re-implementation?
- When better AI models arrive, does it adopt them, or is it locked into one?
- Will it share insights on its product roadmap, and has it delivered what it said was coming?
- Can it name firms like yours that are using its newest capabilities today?

Ask these questions of every vendor, including us.

Pillars Two and Three matter more than which software a firm runs, and no firm should switch operating systems for novelty's sake. But the choice is not neutral. A firm committed to an operating system that has stopped innovating will be exposed over the next five to ten years, as competitors on modern platforms put capabilities to work that its own cannot offer. The question is not only whether the current operating

system works today, but whether it will keep getting better.

Done when: *the firm has committed to one operating system and a partner it trusts, its active data lives there, and every side source has a retirement date that has passed.*

12. Pillar Five — AI

How do we layer AI onto our workflows so it scales good data, not bad?

AI scales whatever it is given. Layered onto documented workflows and clean data, with a person approving the output, it scales good work across every broker at once. Layered onto messy data with no checkpoint, it scales errors just as efficiently. That is why AI comes last in the playbook: not because it matters least, but because it depends on everything else.

1. **Start with a workflow, not a tool.** Pick the one or two workflows where AI would have the largest impact; the table below lists good first candidates. How will the firm know it is working? What does success look like, in numbers?
2. **Check that the data behind that workflow is ready.** Run the Pillar Three measurement on the specific data set the use case depends on. A forecasting use case needs clean deal stages. A follow-up use case needs current activity. If the data isn't ready, fix that data set first.
3. **Keep a human in the loop.** Decide in advance who approves the output before it reaches a client, and what they check. The AI should prepare the action. The human should verify the action. The system should keep an audit trail.
4. **Run it inside the workflow, not beside it.** A use case that requires brokers to open a separate tool and paste in context will be abandoned. The work should happen where the data already is.
5. **Measure against the baseline, then expand.** Add the next use case only after the first one demonstrably works. One well-designed automation is worth a hundred poorly designed ones.

Good first use cases share three traits: they come up often, they run on data the firm already records, and a person can check the output in minutes.

Use case	Data it depends on	Where a person approves
Deals at risk and quiet prospects	Current deal stages, last-touch dates, and next steps	The managing broker reviews the list in the pipeline meeting
Weekly call list: “who should I call this week, and why”	Activity history, next steps, lease expirations	The broker reorders or edits the list
Follow-up emails in the broker’s voice	Activity history and the broker’s own email history	The broker edits and sends
Landlord prospect reports	Tours, inquiries, and activity on the listing	The broker reviews and adds commentary before it goes out

Use case	Data it depends on	Where a person approves
Tour packages	The requirement, current availability, and comps	The broker cuts properties that fit on paper but not in practice
Comp sets and lease abstracts	Closed-deal terms and executed documents	The broker judges relevance and verifies the abstract

The first row is the test from page 1 of this paper. Once deal stages and activity are current, it is also one of the best places to start, because leadership sees the value in the first pipeline meeting.

BY THE NUMBERS

66% / 5%

of commercial real estate professionals use AI weekly or daily, against those who trust it with an actual deal decision, in a survey of 255 CRE professionals fielded in April 2026 (First American Data & Analytics and DealGround).

That gap is not closed by asking brokers to trust more. It is closed by step three. Trust follows from being able to see what the AI drew on and knowing that a person checked it — which is only possible when the data underneath it is data the firm can stand behind.

Done when: one use case is live inside a documented workflow, has a named verification point, and has been measured against the baseline it was meant to improve.

13. The Individual Broker’s Version

Everything in this playbook applies to a single broker or a small team, at a smaller scale and in the same order.

Name the human work and protect it. Know the firm’s policy on who owns the data, and if there isn’t one, ask. Then keep four habits, and let software capture as much of them as possible:

- Record every meaningful client touch the same day.
- Keep a next step and a date on every active prospect.
- Update deal stages the day they change.
- Enter closed-deal terms at close.

Keep everything in one operating system rather than a system plus a spreadsheet plus a notebook. Clean the active set, not the whole history. Then pick one AI use case (“who should I call this week, and why” is a good first one) and judge it honestly against what you would have done without it.

A broker who does this is not just readier for AI. By every measure that shows up on a commission statement, they are a better-organized broker today.

THE BROKERAGE ADVANTAGE

14. One Broker's Notebook, One Firm's Database

An individual broker using AI is limited to their own data: their contacts, their deals, their comps, their read of the market. For a strong producer that is a lot, but it is still one person's slice of a market.

A brokerage has something no individual broker can have: the combined data of every producer at the firm. Every lease the firm negotiated, every sale it closed, every tour it ran, every requirement it saw, every piece of market intelligence that crossed any broker's desk. Pointed at that data, AI becomes a different tool. Comps drawn from all of the firm's closed deals instead of one broker's handful. A forecast built on how the firm's deals actually progress, stage by stage. A pattern — which tenants in which submarkets tend to start moving well before expiration — that no single broker has enough examples to see.

This is the part of the AI transition where brokerages, rather than individual brokers, hold the advantage. The models cost roughly the same for everyone. The firm's combined data does not exist anywhere else, cannot be bought, and takes years to build. Yet at many firms, the comp and market data a team generates lives with whoever generated it, and leaves when they do.

Nor is the advantage reserved for large firms. A ten-broker shop in which everyone records the same way holds several times the deal history of any one of its brokers, and in a local market that can be a meaningful share of all the activity there is to see.

THE COMBINED DATABASE

A broker's data describes a book of business. A firm's data describes a market. AI is far more useful pointed at the second — and only the firm can build it.

15. Two Cultures, Both Legitimate

Whether a brokerage builds a combined database is not a technology decision. It is a values decision, and there are two established positions.

Some brokerages believe strongly in each broker's individuality. The broker's data and relationships are their own; the firm provides the platform, the brand, and the support. Other brokerages hold a team or company perspective, and encourage — or require — company-wide systems and shared data.

Taken to its extreme, the firm-wide model looks like some of the large corporate brokerage firms. Over the past two decades, several have invested heavily in brand, research, and shared resources, and the firm sits at the center of the client relationship. Broker splits at those firms have generally run lower than at

independent brokerages, in exchange for those resources. It is a different model, not a better or worse one, and many independent firms exist precisely because brokers wanted the other side of that trade.

Dimension	The broker's book	The firm's database
Recruiting independent top producers	Stronger. Many established brokers will only join a firm where their book stays theirs.	Harder. Some top producers will not accept it.
Broker autonomy	Stronger. Brokers run their business their way.	Weaker. Shared standards and visibility constrain how brokers work.
Adoption effort	Lower. The firm asks less of each broker.	Higher. Requires standards, enforcement, and leadership attention indefinitely.
Continuity when a producer leaves	Weak. The history leaves with the broker.	Strong. The firm keeps what it has built.
Firm-wide pipeline and forecasting	Limited to what brokers choose to report.	Available in the system, current, and comparable.
Training juniors	Limited. Each broker's history is private.	Stronger. Juniors have more to learn from: more brokers, more data, more examples.
AI across the firm	Limited to each broker's slice.	Draws on all of the firm's data.

Between the two sits a hybrid many firms adopt: **joint ownership**. It arises when a firm offers, encourages, or requires a company-wide system but still wants its brokers to feel independent. The broker keeps the right to take their data and relationships with them if they leave; the firm retains a copy, and ownership of that copy. The firm gets the data advantages of a shared system, and the broker keeps both their independence and a portable book. Like the other two models, it is a cultural choice. Whichever model a firm chooses, it should be put in writing and explicitly agreed to by the firm and every broker before anyone leaves, not negotiated afterward.

Our position has always been that the right model is whatever reflects the firm's culture and core values, and we have seen each of them work. But objectively, more data, more consistent data, and more context make AI more effective, and a firm-wide database is the largest source of all three. Our projection, based on what we observe across the industry, is that over the next five years the firms that lean into data sharing will be positioned to use AI across the firm in ways individual-first firms cannot. That is a trend, not a verdict on which culture is better.

Under any model, sharing is not the same as everyone seeing everything. Landlord positions, client financials, and terms under NDA are confidential by contract or by trust. Permissions should be designed alongside the data: deal terms can feed the firm's comps without the client's name being visible, and

leadership can see the pipeline without every broker seeing every other broker's prospects.

THE CULTURE DECISION

Data ownership is a values decision before it is a technology decision. Make it on purpose, write it down, and tell brokers before they join — not the day one of them leaves.

Conclusion

AI readiness is a people, process, and data project first. Software's role in it is growing fast, capturing activity, suggesting updates, and keeping data current with a click instead of a keyboard. But software amplifies whatever standards and habits a firm already has. The brokers who get the most from AI will be the ones who were already client-centric and disciplined. The brokerages that get the most will be the ones that turn those habits into consistent, clean, thorough data, and, where the firm's culture supports it, combine that data across every producer.

The resistance to this is real. Brokers are right that data entry is unpaid, right that their book is their life's work, and right that nothing appears broken today. The playbook respects all three: software that turns typing into approving, an ownership decision made openly, and a cleanup limited to what matters now. The return on that work is far larger than its cost. What the playbook does not accept is waiting. The models are ready. Most databases are not.

Appendix A — Data Standards Checklist

The checklist follows the order the paper ranks the properties of good data: consistency first, then cleanliness, then thoroughness. Volume is deliberately absent. Adapt the specifics to the firm's specialties; keep the structure.

Firm-wide consistency standards

- One written definition each for *contact*, *prospect*, and *client*.
- One set of deal stages for each deal type (lease, sale, tenant rep, landlord rep), each with a written entry point: the event that moves a deal into that stage.
- One controlled list of property types, classes, and submarkets, selected from a list rather than typed.
- One format each for names, company names, addresses, dates, square footage, and rates.
- Picklists instead of free text wherever the answer comes from a known set.
- A written rule for who creates the record when two brokers know the same person or property.

Record type	Consistent	Clean	Thorough
Contacts and companies	Every contact linked to a company record; titles and roles from the firm's list.	One record per person and company; duplicates merged, not deleted; job changes updated when learned.	Name, company, role, one working contact method, owning broker, relationship status, last-touch date.
Properties and availabilities	Type, class, and submarket from the controlled lists; one address format.	One record per property; availabilities marked leased or withdrawn the day they are.	Address, type, size, owner or landlord contact, current availability status.
Deals in motion	Stage from the firm's defined list; deal type set.	Dead deals closed out with a reason, not left open.	Client, property or requirement, stage, estimated value, expected close, next step and date.
Activities	Activity types from a short list: call, meeting, tour, email, proposal.	Logged against the right contact and deal.	Date, type, and one line on the outcome or next step.
Closed deals and comps	The same set of terms captured for every closed deal of a given type.	Entered once, at close, from the executed documents, not reconstructed later.	The terms comps depend on (rate, term, size, concessions, escalations, commencement), plus confidentiality flags.

Appendix B — Habits Checklist

Cadence	Broker	Leadership
Daily	Record every meaningful client touch the same day, or approve the system's suggested update. Set or update the next step and date on every prospect touched. Update any deal whose stage changed.	Resolve data questions and exceptions the same day. Keep own records current.
Weekly	Review every active prospect with no activity in the last 30 days. Confirm every active deal's stage, value, and next step before the pipeline meeting.	Run the pipeline meeting from the system. Raise stale deals and missing next steps in the meeting.
Monthly	Close out dead deals with a reason. Merge duplicates found during the month.	Run monthly reports from the system. Review the readiness metrics in Appendix C.
Quarterly	Review top relationships for job changes.	Re-run the 100-record measurement. Revise any standard that isn't working.
At every close	Enter closed-deal terms from the executed documents.	Confirm the deal moved from broker to accounting without re-entry.

Appendix C — Readiness Metrics

Establish a baseline before setting any target. The trend matters more than any single reading.

Metric	How to measure	Direction
Record error rate	The 100-record test from Section 10, by category.	Down
Duplicate rate	Duplicate records found in a sample, as a share of the sample.	Down
Next-step coverage	Share of active prospects with a next step and a date.	Up
Stage currency	Share of active deals updated in the last 30 days.	Up
Quiet active prospects	Active prospects with no logged touch in 60+ days.	Down
Producers current	Share of producers whose records pass the weekly review.	Up
Forecast accuracy	90-day forecast from the system against actual closings.	Closer
Close-to-invoice time	Days from executed documents to invoice and splits calculated.	Down
Reports from the system	Share of firm-wide meetings and reports run from the system.	All of them

Appendix D — The AI Readiness Assessment

CRE OneSource's AI Readiness Assessment measures where a brokerage or an individual broker stands today on the ground this paper covers. It takes a few minutes, and a first question routes each respondent to one of two tracks:

- **Individual broker or small team.** Where your contacts and deal history live and whether they are portable, how you decide who to follow up with, how clearly you can see your pipeline and expected commissions, and how well an AI assistant could work from your records today.
- **Brokerage leadership.** Where the firm's data lives and what it keeps when a producer leaves, how many producers keep their records current, how the pipeline and forecast come together, how closed deals reach accounting, and what AI would find if pointed at the firm's data today.

Take the assessment at creonesource.com/ai-readiness-assessment.

Related Reading

CRE Brokerage in 2030 — why the administrative half of the broker's job is going away, and who captures the gains. Available at creonesource.com/whitepapers.

Sources

1. CRE OneSource — “CRE Brokerage in 2030” (2026), available at creonesource.com/whitepapers. Source of the context-tax concept; the 24-hour weekly workload, \$100–\$500 effective hourly rates, and 2030 projections cited in Section 6; the workload inventory referenced in Section 8; and the preparation steps this paper expands.
2. Validity — “The State of CRM Data Management in 2025” (July 10, 2025). Survey of 602 CRM users and administrators across the U.S., U.K., and Australia; 76% said less than half of their organization's CRM data is accurate and complete; 37% reported losing revenue as a direct consequence of poor data quality.
3. Tadhg Nagle, Thomas C. Redman, and David Sammon — “Only 3% of Companies’ Data Meets Basic Quality Standards,” Harvard Business Review (September 11, 2017). Executives reviewed the last 100 records created in their departments; on average 47% of newly created records had at least one critical error. Source of the 100-record measurement method in Sections 5 and 10 and Appendix C.
4. JLL — “2025 Global Real Estate Technology Survey” (October 28, 2025). More than 1,500 senior CRE investor and occupier decision-makers across 16 markets; 88% of investors, owners, and landlords and 92% of occupiers piloting AI; 5% reporting all AI program goals achieved.
5. Gartner — “Lack of AI-Ready Data Puts AI Projects at Risk” (February 26, 2025). 63% of organizations lack or are unsure of AI-ready data practices.
6. ZeroBounce — “2026 Email List Decay Report” (2026). Approximately 23% of email addresses on a typical list go bad each year, based on more than 11 billion verified addresses.
7. George Labovitz and Yu Sang Chang — the 1-10-100 rule, quality-management research (1992); widely applied to data management.
8. First American Data & Analytics and DealGround — “CRE Industry Pulse Check” (May 2026). Survey of 255 commercial real estate professionals fielded March 31 – April 8, 2026; source of the 66% weekly-or-daily usage and 5% deal-decision trust findings.
9. CRE OneSource internal observations. The profile of client-centric brokers and their habits in Sections 1 and 2; the returns described in Section 6; the playbook, timeframes, and enforcement levers in Sections 7–13; the data-ownership models and the description of corporate and independent brokerage models in Section 15; the five-year data-sharing projection; and the checklists in Appendices A–C are drawn from five years of close work with thousands of brokers and from the forty-year brokerage career of the firm’s founder. No individual firm or broker is described.

This is a perspective paper. Its conclusions and recommendations are CRE OneSource's opinion, stated as such. Survey results and study findings are measurements as published on the dates cited and were current as of September 2026. The 2030 figures in Section 6 and the statements in Sections 11 and 15 about the next five to ten years are projections, not measurements. The timeframes, checklists, requirements, and metrics in the playbook and appendices are practical guidance drawn from observation, not industry benchmarks; individual firms will vary and should establish their own baselines before setting targets. CRE OneSource develops brokerage software, and the assessment in Appendix D is its own. Research and survey publishers are named only as sources. No brokerage firm and no competing software product is named, and no product comparison is intended or implied.