

COMMERCIAL REAL ESTATE AI STRATEGY

CRE AI Stack Decision Brief

A practical buyer's guide for \$10M-\$100M commercial real estate operators who need faster underwriting, cleaner lease data, better market intelligence, and less software waste.

BOTTOM LINE

**Do not buy AI.
Buy friction
removal.**

SOURCE BASE	BEST FIRST BET	RISK CONTROL	BUYER RULE
<p>20</p> <p>CRE AI tools reviewed across underwriting, data, leases, sourcing, and reporting.</p>	<p>1</p> <p>Pilot one painful workflow before signing a platform-wide contract.</p>	<p>30</p> <p>Days is enough to validate speed, accuracy, integration fit, and team adoption.</p>	<p>0</p> <p>No vendor gets approved without testing on your actual OM, rent roll, T-12, and lease files.</p>

01 EXECUTIVE THESIS

The winning CRE AI stack is not the tool with the loudest demo. It is the stack that removes the most analyst hours, underwriting uncertainty, document risk, and follow-up chaos with the least operational drag.

For a \$10M-\$100M operator, the correct sequence is simple: fix the data workflow first, then automate underwriting, then standardize lease abstraction, then improve sourcing and investor reporting.

The danger: buying a shiny AI platform before the firm has clean sample files, decision rules, owners, and a pilot scorecard. That turns software into another abandoned operating burden.

Start where errors are expensive and repetition is high.

The best use cases are not generic "AI assistant" tasks. They are places where CRE teams repeatedly convert messy documents and market inputs into investment decisions.

WORKFLOW	BEST STARTING TOOLS	WHY IT MATTERS	PILOT STANDARD
Underwriting	Smart Capital Center, Cactus, Archer	Turns OMs, rent rolls, and T-12s into faster models, memos, and decision-ready outputs.	Compare model output to your existing underwriting on one live deal.
Lease Data	Prophia, Docugami, LeaseLens	Bad lease data quietly contaminates NOI, risk review, tenant obligations, and buyer confidence.	Use real leases with amendments and confirm every critical field back to source language.
Market Comps	HelloData AI, CoStar, CommercialEdge	Rent, concession, supply, and expense assumptions drive valuation quality.	Test speed and accuracy against a recent comp set your team trusts.
Deal Sourcing	Reonomy, CommercialEdge, Cherre	Owner intelligence improves prospecting, pipeline focus, and off-market targeting.	Build a 50-owner list and measure match quality, stale records, and outreach usefulness.
Investor Ops	Juniper Square, Rentlytics	Professional LP communication and reporting reduce manual follow-up and credibility risk.	Recreate one investor update and capital event workflow before committing.

Consultant-level filter: if a tool cannot improve one of four measurable constraints - time, accuracy, risk visibility, or decision throughput - it is probably a demo expense, not infrastructure.

03 RECOMMENDED STACK FOR \$10M-\$100M OPERATORS

Build a stack in layers, not in slogans.

The right stack depends on asset class, transaction volume, and internal discipline. Most firms should not start institutional. They should start with the smallest system that proves ROI on live files.

LEAN STACK

Best for: small owners, lean acquisition teams, early AI adoption.

- HelloData AI for multifamily comps and market surveys.
- Smart Capital Center or Cactus for underwriting automation.
- Prophia per document for lease abstraction risk.
- Reonomy for owner research and prospecting.

GROWTH STACK

Best for: recurring acquisitions, active portfolio management, investor reporting.

- CoStar plus HelloData for broader coverage and multifamily precision.
- Smart Capital Center for connected underwriting and reporting.
- Prophia plus Docugami for higher document complexity.
- Juniper Square for LP communication and administration.

INSTITUTIONAL STACK

Best for: data-heavy firms with repeatable governance and dedicated owners.

- Cherre or similar data layer for unified property intelligence.
- Argus plus AI underwriting for formal valuation workflows.
- Custom internal assistant only after source data is governed.
- BI dashboards for portfolio KPIs and risk alerts.

04 BUY, BUILD, OR AVOID

DECISION	USE THIS RULE	CRE EXAMPLES
Buy	When the vendor has proprietary data, repeatable extraction, security controls, and better coverage than you can build.	Lease abstraction, market comps, owner records, investor portals.
Build Around	When your edge is the decision process, assumptions, narrative, or brand-specific output.	IC memo templates, branded investor updates, acquisition screening scorecards.
Avoid	When the product is a generic chatbot, a no-code toy, or a deck generator pretending to be underwriting.	Unverified custom GPTs, no-data copilots, manual app builders with high upkeep.

A pilot should answer one question: does this tool make better decisions faster on our real work?

Do not let vendors control the test. Bring your documents, your assumptions, your old model, your timing benchmark, and your decision criteria.

WEEK 1

WORKFLOW AUDIT

- Pick the highest-friction workflow.
- Collect recent OM, rent roll, T-12, lease, loan quote, model, and memo samples.
- Define baseline hours, defects, handoffs, and decision delays.

WEEK 2

VENDOR DEMO

- Force a live walkthrough using representative files.
- Ask about data security, ownership, export rights, integrations, and implementation burden.
- Document every manual workaround.

WEEK 3

PARALLEL TEST

- Run the same deal through the old process and the AI tool.
- Score output accuracy, speed, explainability, and team trust.
- Find where humans still need to review.

WEEK 4

DECISION

- Approve, extend, or kill the tool.
- Negotiate contract terms around proven usage only.
- Assign an owner, KPI, file standard, and quarterly review date.

Minimum pilot scoreboard: analyst hours saved, model accuracy delta, source traceability, export quality, integration friction, training burden, security fit, and cost per useful decision.

Good vendors welcome hard questions. Weak vendors hide behind demos.

The buyer's job is not to be impressed. The buyer's job is to reduce uncertainty.

SCORING STANDARD

Score every vendor 1-5 on each question. Any tool below 36 out of 50 should not move to contract without a narrow, low-cost extension.

KILL CRITERIA

Kill the pilot if the tool cannot export usable outputs, explain its assumptions, preserve source traceability, or beat the current workflow on a live file.

- 01 Which asset classes, geographies, and deal types does the tool actually cover?
- 02 What data does the AI use, and what is deterministic versus model-generated?
- 03 Can every extracted lease, rent, expense, or assumption be traced back to a source?
- 04 How does the system handle poor PDFs, amendments, missing data, and contradictory records?
- 05 Who owns the uploaded files, extracted data, and exported work product?
- 06 What integrations exist for Excel, Argus, CRM, property management, VDR, and reporting?
- 07 What can be customized: templates, assumptions, models, workflows, and branding?
- 08 What implementation work is required from our team before value appears?
- 09 How are accuracy claims measured, audited, and corrected?
- 10 What contract minimums, usage caps, export limits, and renewal traps apply?

07 WHERE THE AI HYPE IS THICK

The fastest way to waste money is to buy AI that is detached from CRE data and workflow ownership.

GENERIC CHATBOTS

Useful for drafting and summarizing, weak for investment decisions unless they are connected to governed data, source documents, and review controls.

NO-CODE AI APP BUILDERS

Often require more configuration than buyers expect and usually do not solve underwriting, lease data, or market intelligence by themselves.

PRETTY DECK GENERATORS

Helpful for output polish, but dangerous when the team confuses a sharp presentation with tested assumptions and defensible underwriting.

ENTERPRISE DATA PLATFORMS

Powerful when a firm has owners, governance, and budget. Heavy when a smaller operator just needs cleaner files and faster decisions.

08 WHITE-SPACE OPPORTUNITIES

GAP	WHY IT MATTERS	WHO CARES
Development and Construction AI	Most teams still use static budgets and schedule assumptions; real-time cost and schedule intelligence remains underdeveloped.	Developers, lenders, owners
Capex and Reserve Forecasting	AI can help convert building age, systems, climate exposure, and historical repairs into better reserve planning.	Asset managers, lenders
Broker Outreach Automation	Owner intelligence is improving, but CRE-specific outreach still needs cleaner data, better personalization, and CRM discipline.	Brokers, acquisition teams

If your CRE team is considering AI, audit the workflow before buying the software.

ShelOR Group helps owners, investors, operators, and capital partners move from software curiosity to operational clarity: where AI actually fits, what to pilot, what to avoid, and how to turn tools into repeatable business infrastructure.

OFFER

CRE AI Stack Readiness Review

Designed for \$10M-\$100M CRE operators evaluating underwriting automation, lease abstraction, market intelligence, sourcing tools, and reporting systems.

- Workflow friction map
- Vendor shortlist and kill list
- 30-day pilot design
- Implementation risk review
- Executive-ready recommendation memo

CONTACT

Get the stack right before the contract gets expensive.

ShelOR Group
Property Advisory and AI-enabled
operating systems for commercial real
estate.

Office@shelorgroup.com
(404) 354-3417
shelorgroup.com

Source basis: ShelOR Group - CRE AI State.pdf, reviewed May 24, 2026. Public vendor pages checked May 24, 2026: smartcapitalcenter.com, hellodata.ai, prophia.com, junipersquare.com. This brief is an advisory screening tool, not a substitute for vendor due diligence, legal review, cybersecurity review, or contract negotiation.