


Travtus+

AI for Property Management

What it actually is, where it already sits in your business, and where it gets you into trouble.



Most housing enterprises are further into AI than they realize, and less in control of it than they think.

This industry has been buying AI for a decade under other names. Screening scores. Revenue management. Leasing chat. Maintenance triage. Invoice coding. Automated valuation. Then the technology changed sharply in the space of about three years, the vocabulary around it got considerably worse, and the result is that very few leadership teams can currently say where AI sits in their business, what it decides on their behalf, or who is accountable when it gets a decision wrong.

That is a solvable problem, and this is the plain-English version of the solution: what the word actually covers, where it already sits in your business, what your team can do with it this quarter, what it still cannot do, and the four ways it goes wrong in housing specifically. Every example is a job someone in your business does on a Tuesday. And because the rules that matter here are federal, state, and increasingly municipal, we say plainly which is which.

A note on pace. AI moved faster between the writing of this guide and your reading of it than any technology we have covered before. Treat the shape of the argument as durable and the specifics, especially the regulatory specifics in Part 6, as a snapshot to verify.

CONTENTS

01 What is AI, actually?

Five things the same two letters are used for, and why the difference matters to you

02 Where is AI already used in property management?

You bought AI years ago. An honest audit of where it sits today

03 What can property teams do with AI right now?

One worked morning, then five starting points that need no budget

04 What can AI not do?

The limits worth knowing before you put it in front of a resident

05 How does AI go wrong in housing?

Fair housing, confidentiality, pricing, and who signs the decision

06 What are the AI rules?

What applies to you now, and the AI policy you should already have

07 What to buy, and what to build

Where we come in, and what we think you should build yourself

A Seven questions property teams ask

The short answers in one place


SCOPE

Roles. Regional and community managers, leasing, resident services, asset and revenue managers, finance and compliance, and the executives above them.

Sectors. Multifamily, single family and build to rent, student, and affordable housing.

Jurisdictions. United States, federal through municipal.

You are already an AI business.
You are probably just not a
deliberate one.



01.

What is AI, actually?

The same two letters are used for five quite different things. Telling them apart is most of what AI literacy amounts to, and it changes which vendor claims you should believe.

Start with the least useful definition, because it is the honest one: artificial intelligence is a moving label for whatever computers have most recently become able to do that we assumed required a person.

That definition sounds evasive. It is actually the most practical thing in this guide, because it explains why the word is so slippery in a sales meeting. Spell check was AI once. Now it is spell check. The label keeps sliding forward, which means a vendor can say “AI powered” about a system built in 2009 and not technically be lying.

So ignore the label and ask which of the following five things is under the hood. They have very different capabilities, very different failure modes, and very different regulatory exposure.

1. Rules

Someone wrote down the logic. If the rent is more than five days late, send this letter. If the work order is coded as no heat, escalate to priority one. Rules are not AI by any modern standard, but a great deal of software marketed as AI in this industry is rules with a friendlier front end. Rules are predictable, auditable, and completely unable to handle a situation nobody anticipated.

2. Machine learning

Instead of writing the logic, you show the system thousands of past examples and it works out the pattern itself. Feed it ten years of applications and outcomes and it learns which combinations of features tended to precede a default. This is the engine inside tenant screening scores, delinquency prediction, churn models, and most revenue management. It is genuinely powerful, and it has one defining property you must never forget: **it learns the patterns in the history you gave it, including the ones you are not proud of.** Part 5 returns to this.

3. Deep learning

Machine learning at a scale that changed what was possible. Very large mathematical networks, loosely inspired by neurons, trained on enormous quantities of examples. Deep learning is why a phone can recognize a face, why a camera can pick out a cracked roof tile, and why speech to text stopped being a joke around 2015. In your business it powers image based inspection triage, call transcription, and document scanning.

4. Generative AI

The one that arrived in public in late 2022 and reset everyone's expectations. Instead of classifying or scoring, these systems produce new material: text, images, audio, code. The text variety, the **large language model**, works by repeatedly predicting what should come next given everything before it. That sounds trivial and is not, because to predict the next clause of a lease well you have to have absorbed an enormous amount about how leases work.

This is the shift that matters for property teams, and it is worth being precise about why. Every earlier form of AI needed structured, tidy inputs. Your business is not structured or tidy. It is fifteen years of emails, call recordings, inspection notes, maintenance comments, lease clauses, and board packs. Generative AI was the first technology that could read all of that, in the messy form it already exists in, and answer a question about it.

5. Reasoning models and AI agents

The current frontier, and the least well understood outside technical teams. **Reasoning models** work through a problem in steps before answering rather than responding immediately, which makes them markedly better at multi step questions where an early mistake ruins the answer. An **agent** is a model given a goal, access to tools, and permission to take steps on its own: look up the record, check the policy, draft the notice, log the outcome, and flag the exception for a person.

This is where the interesting operational ground is, and where the honest caution belongs. An agent that can act is an agent that can act wrongly, at speed, across your whole portfolio. Everything in Parts 4 and 5 becomes more important, not less, the more capable these systems get.

THE SHORT VERSION

Ask which of the five you are being sold. Rules are predictable and rigid. Machine learning is powerful and inherits your history. Deep learning handles images and audio. Generative AI reads the messy material your business actually runs on. Reasoning and agents can carry out multi step work, which is exactly why they need supervision. Most “AI” pitched at this industry is one of the first two wearing the language of the last two.

PLAIN ENGLISH GLOSSARY

model

The trained thing itself. Not a program someone wrote line by line, but a very large set of numbers adjusted until it produced good answers. You cannot open it up and read the logic, which is why explainability is hard.

prompt

What you ask. The single biggest quality lever available to a non technical user. Vague question, vague answer.

grounding

Deliberately giving a model your own material to answer from, and requiring it to cite what it used. The single most effective defense against invented answers. Sometimes sold under the acronym **RAG**.

agent

A model given a goal and the tools to pursue it over several steps, rather than answering one question.

model agnostic

A platform not welded to one AI provider, so you can move to a better or cheaper model without rebuilding. Worth asking about, given how fast the underlying models change.

training

The expensive, one time process of building a model from examples. Nothing you type into a well governed business tool should be training anything. Verify that in the contract.

context

What the model can see while answering: your question plus whatever material has been put in front of it. A model with no access to your portfolio can only give you generic answers, however clever it is.

hallucination

A confident, fluent, entirely invented answer. Not a bug being fixed, but a property of how these systems work. See Part 4.

guardrail

A constraint placed around a model: what it may see, what it may do, what it must refuse, and when it must hand over to a person.

automated decision

A legal term of art, not a technical one. A decision made about a person with no meaningful human involvement now carries specific obligations under a growing number of state laws. Part 6.

02.

Where is AI already used in property management?

The question is not whether to start using AI. You have been buying it for a decade. The question is whether anyone in the business can currently tell you where it sits and what it decides.

Most leadership teams we meet believe they are at the beginning of this. They are not. They are somewhere in the middle of it, without a map.

Before deciding what to adopt, do the audit. Go through the list below and, for each row, write down the vendor, the person accountable, and whether a human reviews the output before it affects a resident or applicant. It takes an afternoon and it is the most useful afternoon your leadership team will spend on this subject.

Where it sits	What it is doing	The question to ask
Applicant screening	Scoring or ranking applicants against learned patterns of past outcomes	Can we explain a decline to the applicant, and to a regulator?
Rent and renewal pricing	Recommending prices from market signal and comparable properties	What inputs does it use, and does it see competitors' non public numbers?
Leasing chat	Answering enquiries, qualifying, booking viewings	What is it authorized to promise on our behalf?
Call handling	Transcribing, routing, summarizing, sometimes answering	Are residents told, and where do the recordings live?
Maintenance triage	Classifying and prioritizing work orders, sometimes from photos	What happens when it misclassifies an urgent safety issue?

Where it sits	What it is doing	The question to ask
Marketing	Generating listing copy, images, and ad targeting	Could the targeting or the imagery narrow who sees our homes?
Collections	Predicting delinquency, sequencing contact	Is the contact strategy defensible for a vulnerable resident?
Finance and accounting	Reading invoices, coding spend, flagging anomalies	Who checks what it coded?
Valuation and acquisition	Automated valuation, deal screening, comparables	How much of the investment committee paper is model output?

What is shadow AI, and why is it your biggest exposure?

There is a tenth row, and it is almost never on the audit because nobody procured it. Your teams are already using consumer AI tools, on their own accounts, to get through the day. A leasing consultant is pasting a difficult resident email into a free chatbot to get help with the reply. A regional manager is asking one to summarize a twelve page inspection report. A finance analyst is dropping a rent roll into one to reformat a table.

They are not being reckless. They are being resourceful, because the tool in front of them is slow and the deadline is real. But every one of those interactions may have taken resident information, or commercially sensitive numbers, outside your control and into a system whose terms you have never read.

DO THIS IN THE NEXT THIRTY DAYS

Two actions, in this order. **First, ask.** Not as a disciplinary exercise, framed as research: which tools are you using, and what for? You will learn more about where your real friction is than any process review will tell you. **Second, give them somewhere sanctioned to do it.** A ban you cannot enforce simply moves the activity onto personal phones, where you have no visibility at all. The organizations handling this well replaced the consumer tools with a governed alternative rather than issuing a prohibition.

THE SHORT VERSION

You are already an AI business. You are probably just not a deliberate one. Name every place it already sits, name the person accountable for each, and find out what your teams are quietly doing without you. That audit is the real starting line, and it costs nothing but an afternoon.

VENDOR QUESTIONS THAT GET REAL ANSWERS

which model?

Ask which underlying model, from which provider, and what happens when a better one is released. Vagueness here usually means either a very thin layer over a public model, or rules.

grounded how?

When it answers a question about our portfolio, what exactly is it reading, and will it show us?

what does it not do?

A vendor who cannot readily name the limits of their own system either does not know them or is not telling you. Both are useful to establish early.

trained on what?

Whose material built it, and does our material become part of it? Get the answer in the contract, not the demo.

who owns the logic?

If we want to change how something is scored or handled next quarter, do we change it, or do we raise a ticket and wait for your roadmap?

who is accountable?

When it gets a decision wrong about a real applicant, whose name is on that? The answer is almost always yours. Price that in.

03.

What can property teams do with AI right now?

Not a transformation program. One worked example of a job you already do, then five starting points that need no budget, no integration project, and no data team.

The reason most AI initiatives in this industry stall is that they begin with a strategy deck instead of a task. So let us begin with a task: a regional manager's Monday morning.

Nine properties. A leadership call at eleven. Roughly two hours to work out where the problems are. Traditionally that means opening four systems, exporting three reports, and forming a view based on whichever numbers loaded fastest. Here is the same two hours with a capable, well grounded AI system that can see the portfolio.

- 1. Open with the vague question, not the precise one.** "Which of my nine properties had the worst week, and why?" You are not expected to know the metric that matters in advance. That is the shift. The answer comes back naming two properties, one on rising work order backlog, one on a cluster of complaints about the same contractor.
- 2. Follow the thread.** "Show me those complaints." Six residents, five of them in the same stack of apartments, all referencing the same failed repair over eleven days. This is the piece that no dashboard would have surfaced, because the signal was in the wording of six separate conversations rather than in a field anyone counts.
- 3. Ask the question you would have needed an analyst for.** "Has that contractor been slower than the others across the whole region this quarter?" Yes, on second visits specifically. That is now a procurement conversation with evidence behind it rather than a hunch.
- 4. Get the brief you would otherwise have written yourself.** "Give me one page on each of those two properties for the eleven o'clock: what happened, what I have done, what I need." Ten minutes of drafting rather than fifty.
- 5. Push back when it is wrong.** One of the two properties looks worse than it is, because a bulk work order upload skewed the backlog. Say so. "Exclude the bulk upload on the fourteenth and recalculate." The willingness to argue with the output is a skill, and it is the one that separates people who get value from this from people who do not.
- 6. Make it repeat.** The genuinely valuable step, and the one most people miss. That was useful, so it should not be artisanal. Ask for it every Monday at seven, with the two worst properties and the reasoning, waiting in your inbox before you open your laptop.

Notice what did not happen. Nobody built a data warehouse. Nobody wrote a specification. The regional manager did not learn a query language, and did not delegate the thinking to the machine either. They stayed in charge of the judgment and handed over the retrieval, the reading, and the drafting.

Five starting points, in order of how easily they land

If you want something moving before the next board meeting, start here. None of these require a project.

01

The reply nobody wants to write

The three paragraph response to an angry, partly justified resident complaint. Give the system the thread and the relevant policy, ask for a draft that acknowledges what went wrong without conceding what did not happen, then edit it. Saves twenty minutes and, more importantly, raises the floor on the worst version of that letter your team sends.

02

The catch up brief

Somebody is about to walk into a call about a resident, a community, or an asset they have not touched in six months. A short, made for the job summary of everything already known, assembled in seconds, is worth more to that call than any dashboard. Different roles need genuinely different summaries of the same resident, which is the point.

03

The report that has been on the list for a year

Every business has one. It is not built because it is nobody's priority and it takes a week. Describing it in a sentence and getting a first version back the same morning changes the economics of asking. Expect the second and third versions to be the good ones.

04

The process that only lives in one person's head

Ask the person who does it to talk through it while the system listens, then have it produce the written procedure. This is the single most valuable and least glamorous use in this industry, because tribal knowledge walking out of the door is a bigger operational risk for most portfolios than anything on the technology roadmap.

The pile nobody has read

Two hundred survey responses. A quarter of exit comments. Eighteen months of maintenance notes on one building. Reading unstructured material at volume and telling you what is in it is the thing this technology is unambiguously better at than a person with a deadline.

THE SHORT VERSION

Start with a task, not a strategy. Pick one job that is currently slow, boring, and done by someone senior enough to judge the output. Do it three times by hand with an AI system before you commit to anything. Then, and only then, make it repeat. The organizations getting real value are not running the biggest programs; they are the ones where a lot of people have quietly stopped doing a lot of small painful tasks.

GETTING BETTER ANSWERS

say who you are

"I am a regional manager preparing for a board call" gets a materially different answer than the same question asked cold. Role and purpose are context.

ask for the reasoning

"Show me what you used to conclude that." If it cannot show you, do not act on it. This one habit prevents most of the damage described in Part 4.

set the constraints

Length, tone, audience, what to leave out. "Under 200 words, for a resident, no legal admissions" does more work than any clever phrasing.

give it the source

Paste the policy, attach the report, point it at the record. An answer grounded in your own material is worth ten answers from general knowledge.

iterate out loud

The first answer is a draft. Tell it what is wrong. Three rounds of correction beats one perfectly engineered question, and requires no training.

keep the judgment

Use it for retrieval, reading, and drafting. Keep the decision. That division of labor is not a limitation to be engineered away; it is the design.

04.

What can AI not do?

Fluency is not accuracy, and confidence is not competence. The limits below are not temporary gaps waiting on the next release. Some of them are structural.

The most dangerous property of current AI systems is not that they get things wrong. All software gets things wrong. It is that they get things wrong in beautifully written, entirely plausible prose, at exactly the same confidence level they use when they are right.

Why does AI make things up?

A system asked about a repair obligation it has never actually seen will not say so. It will produce a well formed answer, occasionally with a clause number and a statutory reference, that is simply not real. We have watched this happen with an invented notice period, a fabricated section of a management agreement, and a case citation that does not exist.

This is not a defect being patched out. It follows from how the technology works: these systems are built to produce the most plausible continuation, and a plausible sounding answer is exactly what they are optimized for whether or not the underlying fact is available to them. The practical defenses are three, and they are all within your control: **ground it** in your own documents rather than its general knowledge, **require citations** so a person can check the source in seconds, and **never let it be the last reader** on anything with legal or financial consequence.

Can AI tell cause from correlation?

It is extremely good at noticing that two things occur together and structurally unable to tell you which one caused the other. A model may correctly observe that residents who contact the office more often renew less often. It cannot tell you whether frustration causes the contact, the contact causes the frustration, or both are downstream of a building with a failing elevator. Every one of those has a different intervention, and picking the wrong one is expensive.

This is the limit most likely to cost you money quietly, because the output looks like insight. Treat correlation from a model exactly as you would treat correlation from a spreadsheet: as a question worth investigating, never as a finding.

What about the knowledge that was never written down?

A property has a reputation in its local market. A contractor is tolerated because the alternative was worse. A regional manager knows which site will not cope with a staff change. None of that is in a system, so none of it is in the answer. AI can see the recorded business with a completeness no person can match, and it is blind to the unrecorded one, which in this industry is substantial.

The corollary is more encouraging than it sounds. The organizations getting the most from this are the ones deliberately writing more down, because context is now the thing that compounds.

It cannot tell you what it does not know

Ask a person a question outside their expertise and you usually get a signal: hesitation, a caveat, a referral. These systems have almost no reliable equivalent. The answer arrives in the same register whether it rests on ten thousand well documented examples or on nothing at all. Any process you build has to supply that missing signal from outside, through grounding, citations, confidence thresholds, and a person who owns the outcome.

Who is accountable when AI gets it wrong?

Worth stating plainly because it is a governance point rather than a technical one. There is no jurisdiction in which “the model decided” is a defense. If an applicant is declined, a notice is served, or a safety report is missed, the accountable party is your organization and, quite often, a named individual within it. That does not argue against using these systems. It argues for knowing precisely where in each process a human being takes responsibility, and being able to show it.

THE SHORT VERSION

Fluent is not the same as correct, and these systems are relentlessly fluent. They invent with confidence, they confuse correlation with cause, they cannot see what was never recorded, they cannot flag their own ignorance, and they cannot carry the consequence. Design around all five. Ground everything in your own material, demand citations, and keep a person accountable at every point where the output touches a resident, an applicant, or a dollar.

three question test

Before acting on any AI output that matters: **What did you use? Show me the passage. What would change your answer?** If any of the three cannot be answered, treat the output as a hypothesis.

the named human

Every process that involves AI should have one person who can be asked why it did what it did. Not a committee, not a vendor. If nobody can be named, the process is not ready.

the numbers rule

Never let a figure reach a board pack, an investor report, or a resident letter without one person tracing it back to the source system. Generated prose around verified numbers is fine. Generated numbers are not.

quotes, not summaries

For anything contractual or statutory, ask for the exact wording rather than a paraphrase. Paraphrase is where the drift happens.

05.

How does AI go wrong in housing?

Housing is not a low stakes place to experiment. These systems touch who gets a home, at what price, and how they are treated once they have one. Four failure modes, and what each one costs.

Every general AI guide has a risks chapter. This one is different because our industry carries an obligation most do not: the decisions in question are protected by housing discrimination law, and the consequences of getting them wrong land on people who are trying to find somewhere to live.

1. Can AI discriminate without ever seeing race?

A model trained on past decisions learns the patterns in those decisions. If your approvals historically skewed against applicants from certain zip codes, and those areas correlate with race, the model will reproduce that skew with no field anywhere in it labeled race. This is the crucial mechanism to understand: **you do not need to feed a system a protected characteristic for it to discriminate on one.** Proxies are everywhere. Address, name, school, employer, benefit income, banking history, and phone type all carry demographic signal.

And the exposure does not require intent. A neutral looking rule that produces a disparate outcome for a protected group can be unlawful on its effect alone.

The **Fair Housing Act** covers this, and disparate impact liability applies whether or not a human wrote the rule. Federal guidance issued in 2024 addressed tenant screening and advertising specifically, and made two points worth internalizing: the housing provider does not escape responsibility by outsourcing the score to a vendor, and ad targeting that narrows who sees a listing can itself be a violation.

Screening also sits under the **Fair Credit Reporting Act**, which means accuracy obligations, adverse action notices, and a real dispute route for the applicant. A score you cannot explain is a score you cannot lawfully act on.

Then there is the state and local layer, which is where operators get caught out most often. A majority of states and a long list of cities protect classes the federal statute does not, with source of income and housing voucher status the most consequential for screening.

A model trained on decisions your business made before those protections applied in a given market will carry on making them.

The practical control is unglamorous and effective: test the outcomes, not the intentions. Take last year's decisions, group them by protected characteristic where you lawfully can, and look at the approval and pricing rates. If the pattern is uncomfortable, you have found it before a regulator or a claimant did.

7

protected classes under the federal Fair Housing Act. Many states and cities add more, and a model does not need to see any of them to discriminate on one.

Fair Housing Act, 42 U.S.C. §3604

2. What happens to resident information you put into AI?

When someone pastes a resident's complaint, a rent roll, or an applicant's financial details into a consumer AI tool, that information has left your perimeter. Depending on the service and the plan, it may be retained, reviewed by humans for quality, or used to improve the underlying model. Under the growing set of state privacy laws, that can be an unlawful disclosure, a reportable incident, or both.

The fix is not a memo. It is the combination described in Part 2: a sanctioned tool with contractual guarantees on retention and training, and a genuine reason for staff to prefer it. People route around controls that make their job harder, every time.

3. Is algorithmic rent pricing legal?

This is the risk that has moved fastest and is least reflected in most operators' governance, so it deserves its own heading.

Recommending rents from a model that is informed by competitors' non public pricing and occupancy is being treated by enforcers as a coordination problem rather than a technology one. The theory is straightforward: if several competing landlords in a submarket

delegate pricing to a shared system fed by each other's confidential numbers, the outcome may look a great deal like an agreement not to compete, whether or not anyone in the room ever spoke.

Federal and state antitrust action has advanced on precisely this basis, and a growing list of cities and states have passed or proposed ordinances restricting algorithmic rent setting outright. Assume this is live in at least one of your markets, and check the rest.

THE THREE PRICING QUESTIONS TO ASK YOUR VENDOR IN WRITING

1. Does the recommendation draw on non public pricing, occupancy, or concession information from competing landlords, in any aggregated form? **2.** Can we deviate from the recommendation without limitation, penalty, or contractual pressure, and do you record our deviations? **3.** Which jurisdictions we operate in currently restrict algorithmic rent setting, and how does the product comply? Get the answers on letterhead, and route them past counsel rather than procurement.

4. Who owns an AI decision inside your business?

The quietest failure, and the most common. A score appears in a workflow. It was configured by someone who has since left, on assumptions nobody wrote down, and it now shapes hundreds of decisions a month. Nobody is doing anything wrong. Nobody is accountable either.

This is why the audit in Part 2 matters more than any policy document. A named owner for every AI touchpoint, a written statement of what it does and does not decide, and a review date. It is the least interesting slide in the deck and it is the one that holds up under scrutiny.

THE SHORT VERSION

Test outcomes, not intentions. A system can discriminate without ever seeing a protected characteristic, and neither jurisdiction accepts a vendor as the responsible party. Give your teams a sanctioned place to work so information stops leaking to consumer tools. Treat algorithmic pricing as an antitrust exposure and get written answers. And put a name against every AI touchpoint in the business, because an unowned decision is the one that will surprise you.

TERMS YOUR COUNSEL WILL USE

disparate impact

A neutral practice that produces a materially worse outcome for a protected group. Unlawful on effect, without any need to prove intent. The US framing.

proxy variable

A harmless looking input that carries protected information: address, name, school, employer, phone type. The reason “we do not collect race” is not a defense.

impact assessment

A written pre launch review of what a system does, what could go wrong, and for whom. Required outright in some states for consequential decisions, and the first document anyone will ask for. Do it before launch, not after.

meaningful human review

Not a person clicking approve. A person with the information, the authority, and the time to reach a different conclusion. Regulators test for the difference.

protected class

A characteristic the law forbids you to decide on. Seven are federal; many states and cities add more, and source of income is the one that most often catches screening models.

adverse action

A decline, or worse terms, based on a consumer report. Triggers a duty to notify and to give the applicant a real route to dispute.

algorithmic collusion

Competing firms reaching a coordinated outcome through a shared system rather than a conversation. Increasingly treated as an agreement regardless.

06.

What are the AI rules for property managers?

There is no single AI law waiting to tell you what to do.

There is a patchwork, and most of what already binds you was written before any of this existed.

The most common regulatory mistake we see is a leadership team waiting for AI legislation before setting a policy. The wait is unnecessary. Housing discrimination law, consumer reporting law, data protection law, and competition law all already apply to what your systems are doing today, and they applied last year too.

There is no comprehensive federal AI statute, and the federal posture has swung toward encouraging development rather than constraining it. That pushed the action to the states, which now differ materially. The practical consequence for a multi state portfolio: your AI compliance position is a map of markets, not one company wide answer.

What binds you regardless of state

The Fair Housing Act and its disparate impact doctrine. The Fair Credit Reporting Act, for anything screening related. And the Federal Trade Commission's authority over deceptive claims, which it has used against AI marketing that overstated what a product could do. That last one cuts both ways, because it also applies to what your own leasing chat promises a prospect.

Where the states have moved

Colorado has enacted a broad framework aimed at consequential decisions in areas including housing, with implementation repeatedly adjusted. Texas, California, and Illinois have each brought in AI specific obligations, largely centered on transparency, disclosure, and, in California's case, what a model was trained on. New York City requires bias auditing for automated hiring tools, which matters if you use them for your own site teams. Several states have also pulled automated decision making inside their general privacy statutes, which gives residents rights over decisions made about them whether or not an AI specific law applies in that state.

The one to watch in this industry

Municipal and state restrictions on algorithmic rent setting. Several major cities have already legislated and the list is growing. This is

the item most likely to have changed between the writing of this guide and your next pricing review, and it genuinely varies market by market.

Does the EU AI Act apply to you?

The EU AI Act is the most prescriptive regime in force anywhere, and it applies on the basis of where the effect lands rather than where your company is registered. It classifies systems by risk, prohibits a small set outright, and imposes substantial documentation, testing, oversight, and transparency duties on high risk uses. Access to essential services, a category that has been read to include housing in some analyses, sits in the high risk tier. Its obligations phase in over several years.

For most US operators the EU Act matters less as a compliance obligation than as a template, because it is where the large vendors are building their controls, and because the specific things it demands, documentation of the system, human oversight, testing for bias, and records of what happened, are the same things a state regulator or an opposing attorney will ask you for without needing an AI statute to point at.

What should be in your AI policy?

You do not need a forty page document. You need eight decisions, written down, approved at board or executive level, and shared with everyone. Most organizations can draft this in a single workshop.

- **Sanctioned tools.** Which AI systems staff may use for work, and the explicit statement that using anything else with resident, applicant, or commercially sensitive information is not permitted.
- **Prohibited inputs.** A short, concrete list. Applicant financial details, health information, safeguarding records, anything about a child, anything under legal privilege.
- **The human decision line.** The decisions that may never be made without a person: declining an applicant, serving notice, closing a safety related report, anything affecting someone's tenancy.
- **Disclosure.** When you tell a resident, applicant, or client that they are dealing with an AI system, or that one contributed to a decision about them. Default to telling them.
- **The register.** A living list of every AI system in use, what it does, who owns it, and when it was last reviewed. This is the artifact a regulator will ask for first.
- **Outcome testing.** Who checks for discriminatory outcomes, against what, and how often. Annual is a floor, not a target.

- **Vendor terms.** The minimum contractual position: your material is not used for training, retention is bounded and specified, you get audit rights, and liability is not entirely disclaimed.
- **The escalation route.** How a member of staff reports an AI output that looks wrong or unfair, and who has to act on it. Make it easier than staying quiet.

The above is general orientation, not legal advice. Effective dates and municipal rules here change frequently; verify anything you intend to act on. See the disclaimer at the end of this guide.

WHO REGULATES WHAT

HUD & DOJ

Fair housing enforcement, including guidance on screening and advertising. DOJ also brings the antitrust cases on algorithmic pricing.

State AG offices

Increasingly the most active enforcers on both algorithmic pricing and AI transparency. Do not plan only for federal.

Privacy regulators

State level, and newly relevant. Several state privacy statutes now carry rights over automated decisions regardless of any AI specific law.

CFPB & FTC

Consumer reporting accuracy and adverse action; deceptive claims about what a product can do.

Fair housing agencies

State and local. They take complaints, they run tests, and their protected classes are frequently broader than the federal list.

07.

What to buy, and what to build

Everything so far would be true whoever published it. This is the part where we tell you what we build and why. Short version: you should buy and build, and most of this industry is drawing the line in the wrong place.

Here is the pattern we watch play out. A housing enterprise decides to take AI seriously. Within eighteen months it owns nine AI features, from nine vendors, in nine parts of the business, and not one of them knows what the others know.

The leasing tool has never seen a maintenance history. The pricing engine has no view of retention risk. The screening score cannot see the resident relationship that came before it. Each product demonstrated well in isolation. Together they produce a business that is more expensive to run than it was before, and no more intelligent.

To be clear about what went wrong there, because it is easily misread: the mistake was not buying. The mistake was buying nine of everything. Nine infrastructures, nine permissions models, nine roadmaps. We think that happens for three reasons.

Cost, context, and trust

Cost. Every point solution ships its own infrastructure. Its own messaging, its own telephony, its own integrations, its own permissions model. You pay for that stack nine times, and then you pay again to make the nine talk to each other. The endless demand for another API and another webhook is not an implementation phase. It is the architecture.

Context. Each system sits in its own orbit and shares nothing. Which means the thing that would actually make AI useful, everything your business already knows, stays split nine ways. Your data team becomes the bottleneck for every question worth asking, and the promise stays siloed.

Trust. You are dependent on nine roadmaps. When the business changes, and in this industry it changes every quarter, you raise a ticket and wait. And the intelligence being built on top of your own business becomes the vendor's asset rather than yours.

The shift that actually matters

Software has always been bought department by department, and AI is being sold the same way: a smarter feature bolted onto one team's tool. But that misreads what changed. Reasoning models can think *across* a business, and that only works with enterprise wide context. So the opportunity is not a smarter app for leasing. It is a business where intelligence, decisions, and the work itself are joined up.

That is what **AI-native** means, and it is worth defining precisely because the phrase is used loosely. An AI-native enterprise is built backwards from four things:

- **What the business knows is the foundation**, not an afterthought bolted on at the end of an implementation.
- **Work is carried out by agents that reason**, not by rigid rules that break the first time reality is unusual.
- **Intelligence, decisions, and action are joined up**, so the enterprise behaves as one rather than as nine departments with nine tools.
- **The capability is yours**, so you open up one use case after another on your own timetable, and what you build becomes your advantage rather than your vendor's.

Two things follow that leadership teams tell us they have never had. Real transparency, because when intelligence is joined up and always on you can finally see what is actually happening across every property without being on the ground. And flexibility, because when you build the reports, profiles, scores, and agentic workflows yourself, you decide what matters and change it when the business changes.

So: buy the platform, and build what only you can

This is our position, and it is the whole reason Travtus exists. It is not an argument against buying. You should buy a great deal of this, and building your own AI infrastructure from the ground up would be a two to three year mistake for almost every housing business. The useful question is not whether to buy. It is **where you draw the line between what you buy and what stays yours**.

Buy the platform. Nobody in this industry should be hand building governance, role based security, property based permissioning, model routing, or integration plumbing. That is the engineering that comes in the box. The Everyday AI™ Platform is model agnostic, so you are not tied to one AI provider as the models keep improving, it integrates without disruption, and the housing context is already built in rather than configured from scratch.

Buy the ready-made, wherever ready-made does the job. A lot works on day one without you building anything. You connect what your business already generates, and in **Gateway** anyone can explore it and simply ask, in plain language, across conversations, records, and numbers alike. Pre-built **profiles** and **scores** are there to use immediately. Building something that already exists is not a virtue, it is a delay.

Build the part that is actually you. Then, in **Studio**, your own people build what no vendor can ship: the **reports** only your business needs, described in a sentence, **profiles** cut the way each role needs to see a resident, community, or asset, **scores** that encode what *your* business thinks is worth watching, and **agents and agentic workflows** shaped around how you genuinely run, carrying out real jobs consistently and flagging the exceptions for a person.

That is also the order we would recommend you do it in. Start by exploring and using what is ready. Build your own with your team once you know what you want. Bring us in to build alongside you when the ambition outgrows the appetite to do it yourselves.

The part that matters most: when you do build, your business users do the building. Not a tech team, not a professional services engagement, not a wait behind someone else's roadmap.

Build your own AI, in a sentence

That line is our product claim and we mean it literally. In Studio you describe what you want in plain language and it gets built: a report, a score, a profile, an agent that handles a real job end to end. Not a configuration screen, not a services engagement, not a ticket in someone else's backlog. A sentence.

That is what makes the buy and build split work in practice. Buying the platform is only a good trade if building on it is genuinely within reach of the people who understand the work. If building requires engineers, you have not bought a platform, you have bought a dependency.

So the honest test to apply to us, and to anyone else pitching you a platform: **can a regional manager build something useful on it, unaided, in an afternoon?** Ask to watch someone do it in the demo, with your own material. That single question separates a platform from a product with an API.

THE SHORT VERSION

Buy and build, and be deliberate about the line between them.

Buy the platform and the enterprise engineering, because hand building that is a waste of two years. Use the ready-made features wherever ready-made does the job. Then build the reports, scores, and agentic workflows that are specific to how you run, because that is the part no vendor can ship you and the part that stays yours. What you should not do is buy nine of everything.

WHAT TO COMPARE, HONESTLY

a feature

Solves one job for one team. Fast to buy, easy to demonstrate, and it will never see the rest of your business. Fine for a genuinely isolated problem.

a platform

Infrastructure that ships useful features on day one **and** lets your own people build the rest. Harder to grasp in a demo, because the demo is only half the product; what you build on it is the other half. This is the category we are in, and the honest trade off is that it asks something of you.

a suite

Several features from one vendor. Better than nine vendors on integration cost. Still their data model, their roadmap, their definition of what matters.

build from scratch

Building the infrastructure itself, not building on top of it. The genuine alternative, and right for a small handful of very large enterprises. Budget two to three years, a platform team, and the ongoing cost of keeping pace with a field that moves monthly. For nearly everyone else this is the thing to buy.

Seven questions property teams ask about AI

The short answers in one place, if you want something to forward to a colleague. Each is worked through properly in the part named beside it.

What counts as AI in property management software?

Five quite different technologies share the label: written rules, machine learning, deep learning, generative AI, and reasoning models with agents. Most software sold as AI to this industry is rules or machine learning described in the language of the newer two, so the useful question in any vendor meeting is which of the five is actually under the hood. [Part 1 →](#)

Can an AI system discriminate without being given race or another protected characteristic?

Yes. Address, name, school, employer, benefit income, banking history and phone type all carry demographic signal, so a model can reproduce a discriminatory pattern with no protected field anywhere in it. Under the Fair Housing Act, liability can follow from the outcome alone with no intent required, and many states and cities protect classes the federal statute does not. [Part 5 →](#)

What can AI not do?

It cannot reliably tell you when it does not know something, distinguish cause from correlation, see anything that was never recorded, or carry accountability for a decision. It will also produce confident, fluent, entirely invented answers, which follows from how the technology works rather than being a defect awaiting a fix. [Part 4 →](#)

What is an AI agent in property management?

An agent is an AI system given a goal, access to your systems, and permission to carry out several steps on its own: look up the record, check the policy, draft the notice, log the outcome, and flag the exception for a person. That is the difference from a chatbot, which answers, and from a rule, which fires. An agent works a job. It is also why supervision matters more as these get more capable, because an agent that can act can act wrongly at speed. [Part 1 →](#)

Where is AI already used in property management?

In applicant screening, rent and renewal pricing, leasing chat, call handling, maintenance triage, marketing, collections, invoice coding, and automated valuation. Most portfolios are already running AI in six or more of these places, bought from different vendors at different times, which is why the first useful step is an audit rather than a purchase. [Part 2 →](#)

Is algorithmic rent pricing legal?

It depends where you operate and what the model is fed. Recommending rents from a system informed by competitors' non-public pricing and occupancy is being treated by enforcers as a potential coordination problem rather than a technology question. Federal and state antitrust action has advanced on that basis, and a growing list of cities and states restrict algorithmic rent setting outright. [Part 5 →](#)

Do property managers need an AI policy?

Yes, and it does not need to be long. Eight decisions cover it: which tools are sanctioned, which inputs are prohibited, which decisions a person must always make, when you disclose AI involvement, a register of every system in use, who tests for discriminatory outcomes, minimum vendor contract terms, and how staff escalate an output that looks wrong. [Part 6 →](#)

The housing businesses that get value from this are not the ones running the largest programs. They are the ones where a lot of people quietly stopped doing a lot of small painful tasks, on infrastructure they did not have to build themselves.

Buy the platform, the enterprise engineering and the housing context. Then have your own people build the reports, profiles, scores and agents that are specific to how you run.

95%

AUTOMATION RATE

25%

RETENTION LIFT

15%

PRODUCTIVITY GAIN

Build your own AI in a sentence.

TRAVTUS · THE EVERYDAY AI™ PLATFORM FOR HOUSING

Interested in AI platforms for the industry?

Visit www.travtus.com or email contact@travtus.com

ABOUT THIS GUIDE

Published by Trvtus. Trvtus™ and Everyday AI™ are trademarks of Trvtus. Written for property teams in the United States, covering multifamily, single family and build to rent, student, and affordable housing.

This guide is general information for a business audience. It is not legal advice and it has not been reviewed by counsel. The regulatory position in Parts 5 and 6 is as we understand it at the date of publication, in an area where statutes, effective dates and municipal rules change frequently and vary by market. Nothing here creates a professional relationship or an entitlement to rely; confirm anything you intend to act on with qualified counsel licensed in the relevant jurisdiction.

Performance figures reflect results achieved by specific Trvtus customers in their own environments and are not a projection of what any other organization will achieve.

All content is original to Trvtus. © 2026 Trvtus. You are welcome to share this guide whole and unaltered with attribution. Please do not republish extracts as your own.

