



The Compliance Playbook for AI-Powered Leasing

Note: This content is for general informational purposes only, is not legal advice. You should consult their own counsel regarding your specific legal and compliance obligations.



WELCOME

Table of Contents

3

Introduction

4

I: AI is moving faster than laws can keep up

5

II: Building frameworks to mitigate risk

7

III: Why having a “human-in-the-loop” is necessary

8

V: Building a compliance-first AI strategy

10

VI: How Entrata Helps

11

Conclusion



Introduction

AI is becoming part of every step of the leasing process. Implemented thoughtfully, it takes over the repetitive, manual work that used to fill the day: answering leasing questions, scheduling tours, verifying identities and income, detecting fraud, and executing leases.

But AI also adds new complexity. Chief among it is compliance, particularly with the federal Fair Housing Act, FCRA, and similar state laws. And federal regulation is only the beginning. Many states and municipalities are considering their own laws to regulate AI products and protect renters.

Regulators are working through the same questions you are: how AI behaves in real-world situations, and how it affects the people it touches. One concern is AI rejecting applicants based on historical data that's biased, intentionally or not, against certain demographics. To keep the process fair, many new regulations require every leasing decision to be backed by data, documented, and shared with the prospect on request when they're denied.

AI shouldn't solve one problem and create another. The right guardrails keep you compliant as you adopt it. In this e-book, we'll cover how AI is changing the leasing process, how it's shaping legislation, the risks it introduces and how to address them, and why the human/AI partnership still matters in leasing decisions.



SECTION I

AI is moving faster than we can keep up

New AI companies and capabilities seem to launch monthly, weekly, even daily. First came generative AI, handling content creation, emails, and website copy. Then came agentic AI, which automated workflows and took over many repetitive leasing tasks.

To protect consumers, many states and municipalities are considering legislation on how AI can be used in leasing, particularly for vetting prospects. So far, only Colorado has enacted a law.

Leasing decisions sit at the intersection of two federal laws: the Fair Housing Act (FHA), which prohibits discrimination based on race, color, national origin, religion, sex, familial status, and disability, and the Fair Credit Reporting Act (FCRA), which requires you to provide notice and documentation whenever you make an adverse decision.

One way to mitigate this risk: skip general AI tools and use ones purpose-built for multifamily. General AI models typically can't explain what's driving their decisions. If a rejected applicant asks why, "the algorithm said so" isn't good enough for FCRA compliance.

AI models trained on historical data can inherit the biases embedded in that data. General tools that aren't built for multifamily, and that offer less flexibility, may factor in information outside your policies. That can skew decisions and put you out of compliance with Fair Housing and FCRA.

It doesn't stop there. General tools may not account for jurisdiction-specific rules. As noted, cities and states aren't waiting on federal legislation. They're moving on their own to protect the people most likely to be harmed by AI-driven decision-making.

If you operate in multiple jurisdictions, your AI needs to enforce each one's specific restrictions, or you risk non-compliance and costly litigation. Before deploying AI across your portfolio, make sure you can answer these questions:

- What information is the AI using to make its decision?
- Do you know exactly what decisions the AI will influence?
- Is every recommendation based on documented company policy?
- Can you explain every recommendation if questioned by an applicant or regulatory agency?
- Which recommendations require human review?
- Are your screening criteria applied consistently for every applicant across all your communities?
- How are you monitoring AI performance over time?
- Do onsite teams understand their role with AI?

SECTION II

Building frameworks to mitigate risk

Every leasing decision lives at the intersection of multiple laws, and new legislation is trending toward regulating automated decision-making in housing. To stay ahead, properties should adopt policies and tools flexible enough to change as fast as the legal landscape does.

Instead of expecting onsite teams to memorize a legal textbook, build decision-making frameworks around FHA, FCRA, and other laws, such as clearly defined, written screening criteria. These frameworks create guardrails that simplify approvals and keep you compliant. Here's what that might look like.

The FHA was passed to protect residents from discriminatory leasing practices. Build your FHA framework around these questions:

- Are approval criteria applied consistently?
- Could this create a disparate impact?
- Can applicants receive documentation about their screening?



The FCRA regulates consumer reporting agencies, protects personal privacy, and ensures accurate, fair consumer report data. It:

- Guarantees the right to view consumer files.
- Gives individuals the right to dispute errors and provide updated information to agencies.
- Gives applicants the right to know why they were denied, with documentation to support it.
- Requires consumer reporting agencies to share this information only with businesses that have a legal, valid purpose.

When it comes to AI and compliance, focus on the decisions AI produces, not the technology itself. That's why a person, using all available data from AI and other sources, should make the final leasing decision. Other principles for your framework:

- Technology should follow your policies, not the other way around.
- Transparency above everything: if you can't explain it, don't automate it.
- Use the same criteria for every applicant at every property.
- Someone owns every decision, always.
- Good documentation is your best defense against litigation.

AI isn't perfect

AI can do a lot of good for your properties, but it's still early technology, and it can make mistakes or fail to document the reasoning behind a decision. Generalized AI screening tools, for example, can produce black-box decisions: a leasing recommendation with nothing to back it up. That's a serious problem for properties that need to stay compliant.

Anyone who's used AI knows it can hallucinate or make up data outright, and depending on what it was trained on, an LLM can carry historical bias. The goal is to eliminate that bias so every leasing decision rests on consistently applied fair housing policies.

AI can also experience scope creep, factoring in information outside your stated policies and producing inconsistent decisions. That inconsistency can put your property out of compliance and increase your exposure to litigation.



SECTION III

Why having a "human-in-the-loop" is necessary

Human-in-the-loop isn't just a buzzword. It's an essential part of a successful AI-driven leasing process. It ensures a qualified team member reviews the screening data behind every AI adverse recommendation to confirm the property's policies were followed.

Your screening policy should specify exactly what factors are considered, what thresholds apply, and what lookback periods you plan to use. AI can then pull the relevant information against those criteria and defer to a human for the final decision.

When making these decisions, avoid blanket rejection policies (e.g., denying anyone with a criminal record), which put you at risk of violating Fair Housing laws. Your screening policy may need more nuance, such as weighing the nature,

severity, and recency of an offense. A human reviewer makes that judgment call.

With the right policies and framework in place, human team members have everything they need to make the decision and, if an applicant requests it, provide supporting documentation. "Not knowing how AI came to the decision" isn't an excuse, and it can put you out of compliance.

AI also excels at applying your criteria consistently across applications and properties in your portfolio. But consistent recommendations are only part of the equation. Ultimately, a human decision-maker takes in the data AI provides and applies judgment to make the call.

AI does the bulk of the work; humans own the final decision.

AI responsibility

Fraud detection

Document analysis

Scheduling

Follow up

Data gathering

Human responsibility

Reasonable accommodations

Address Fair Housing concerns

Conduct in-person tours

Handle appeals process

Final accountability



SECTION V

Building a compliance-first AI strategy

Most properties won't dive into AI headfirst. They'll take a measured approach, letting adoption happen over time to build onsite team buy-in. Leasing teams typically start with an AI assistant, then add income and identity verification. Marketing begins using generative AI to draft blog posts and other resident-facing content. Eventually, you build AI-driven, automated workflows across the leasing process.

The challenge isn't the technology. It's making sure that technology helps, rather than hinders, compliance. Without clear governance, team members might unwittingly apply different standards, introduce inconsistent processes, or rely on AI in ways that create unnecessary risk.

Focusing on compliance first won't slow your leasing teams down. It puts the necessary guardrails in place so teams can use AI with confidence.

Governance vs. technology

One of the biggest mistakes properties make when deploying AI is asking, "What can I automate?" The better question is, "What decisions should AI be influencing?"

Not all workflows carry the same risk. Some AI interactions are negligible; others significantly increase a property's exposure (e.g., maintenance requests vs. leasing decisions). Categorizing AI use cases by risk helps you identify where additional human oversight is needed.

Low-risk activities

- OnDrafting emails
- Summarizing maintenance requests
- Scheduling appointments
- Creating marketing content
- Answering general leasing questions

Moderate-risk activities

- Fraud detection
- Income verification
- Lead scoring
- Maintenance request prioritization
- Renewal recommendations

High-risk activities

- Resident screening recommendations
- Adverse action decisions
- Reasonable accommodation requests
- Fair Housing-related interactions
- Workflows affecting someone's ability to obtain housing

The greater the potential impact on an applicant or resident, the greater the need for documented human oversight to mitigate risk.

Establish clear ownership

Ownership depends on the task at hand. AI touches operations, leasing, legal, IT, marketing, and executive leadership, so someone needs to keep those groups aligned. The goal isn't to approve every AI feature, but to set consistent standards

for how AI is evaluated, deployed, monitored, and improved over time. Just as organizations govern financial controls and cybersecurity with clear policy, AI needs the same.

Create written policies before deployment

Don't use AI to create your AI policies, and don't deploy AI before you've written them. When developing your policy, ask:

- What business problem is the AI solving?
- What data does AI have access to?
- Which decisions should AI make, and which shouldn't it?
- When is human review required?
- Who has the authority to override AI recommendations?
- How are decisions documented?
- What is your record retention policy?
- If required, what does the appeals process look like for an AI decision?

Prioritize Explainability

One surefire way to lose trust in AI is being unable to explain its recommendations. "Because AI says so" isn't an acceptable answer to give an applicant. You should be able to answer questions like:

- What information influenced this recommendation?
- Was the recommendation based on company policy?
- Did someone independently review the decision?
- Is there documentation to support the outcome?

Explainability isn't just about compliance. It builds trust with applicants, residents, and onsite teams.

Humans are responsible for consequential decisions

AI can process information faster, identify patterns humans might miss, and surface risk almost instantaneously. But humans, not AI, should make the final call on leasing decisions.


Humans should authorize screening exceptions, accommodate reasonable requests, review disputed results, verify conflicting data, assess unique resident circumstances, and address fair housing concerns. AI provides recommendations and context; humans provide judgment.

Conduct regular audits of AI

Just as you audit financial controls, accounting procedures, and operational performance, you should regularly review AI policies and procedures to confirm they're accurate, workflows aren't broken, and the data needed for decisions is available. Regular audits catch issues before they become regulatory or legal problems. Consider these questions when auditing your AI:

- Have the recommendations remained consistent?
- Have approval rates differed unexpectedly across protected classes?
- Are policies being applied consistently across communities?
- Are required review processes being bypassed by onsite teams?
- Do AI outputs still align with company policies and regulatory requirements?





SECTION VI

How Entrata helps

Responsible AI That Keeps People in Control

Artificial intelligence has the potential to make leasing teams faster, more efficient, and better equipped to detect fraud and process applications. But when it comes to decisions that can impact a person's ability to obtain housing, technology should support human judgment, not replace it.

When implementing AI, properties should ensure people remain accountable for consequential decisions. That's why Entrata, ResidentVerify®, and Entrata Layered Intelligence® (ELITM) are designed to give leasing teams better information, greater efficiency, and stronger fraud protection while preserving the human oversight necessary to support Fair Housing and FCRA compliance.

Human Oversight Built Into the Process

Entrata's approach uses machine learning to produce an RV Index Score that helps leasing teams make informed decisions.

Rather than issuing opaque approvals or denials, ResidentVerify gives properties comprehensive information to evaluate applicants against their organization's established policies. This allows onsite teams to review recommendations, consider unique circumstances when appropriate, and make informed decisions backed by documented screening results.

Keeping humans involved in consequential housing decisions helps operators apply policy consistently while maintaining accountability throughout the screening process.

Purpose-Built AI for Multifamily

Unlike general-purpose AI tools, Entrata Layered Intelligence is built specifically for multifamily operations.

Its AI capabilities are built around the workflows property managers use every day: resident screening, identity verification, fraud detection, and leasing operations. It's purpose-built for housing providers, designed to work within your established processes rather than replace them.

This focused approach helps reduce the risks of using consumer AI tools for business-critical decisions.

Greater Transparency for Better Decision-Making

One of the challenges with many AI systems is explainability. When applicants, regulators, or ownership groups question a housing decision, operators need to be able to explain and support it.

Entrata's screening solutions provide detailed screening information that helps leasing teams understand what factored into an application's evaluation. This transparency supports more consistent decision-making and helps teams fulfill compliance responsibilities, including individualized review when required.

Technology Designed to Support Responsible AI

Responsible AI isn't defined by how much decision-making technology can automate. It's defined by how effectively technology helps people make informed, consistent, and accountable decisions.

Entrata's approach reflects that philosophy: combining AI-powered efficiency with human expertise, transparent workflows, and purpose-built multifamily solutions. As regulations evolve and AI adoption accelerates, operators need technology partners who understand that innovation and compliance must work together, not compete against each other.





Conclusion

Artificial intelligence is reshaping the future of multifamily operations. From leasing and marketing to fraud detection and resident screening, AI has the potential to reduce manual work, improve operational efficiency, and help teams make faster, more informed decisions.

But with that opportunity comes responsibility.

When technology influences housing decisions, operators must ensure that efficiency never comes at the expense of fairness, transparency, or accountability. Fair Housing requirements, FCRA obligations, and emerging AI regulations all reinforce the same principle: organizations are responsible for their decisions, not the technology that helped make them.

The most successful operators won't be the ones who automate the most. They'll be the ones who implement AI thoughtfully, with clear governance, documented policies, consistent processes, and meaningful human oversight at every critical decision point.

Responsible AI isn't focused on efficiency at all costs. It's about making the right decision the first time. As the multifamily industry continues to evolve, the question is no longer whether AI belongs in your operations, but whether your AI strategy is built to support the people who depend on it.

With purpose-built technology, transparent workflows, and a commitment to keeping humans at the center of consequential decisions, operators can embrace AI with confidence, improving efficiency without sacrificing the fairness, consistency, and accountability that define exceptional resident experiences.