

MULTIFAMILY NEXT

AI Regulations and the Multifamily Industry

What You Need to Know in 2026

A Plain-Language Guide for Multifamily Leadership

AI regulation is changing fast. New federal policies, state laws, and legal settlements are reshaping how multifamily operators can use AI for pricing, screening, chatbots, and more. This guide cuts through the legal language and gives you what matters: what changed, what it means for your properties, and what to do about it.

Multifamily NEXT | Build AI Fluency. Develop Real Skills. Create Working Tools.
March 2026 | multifamilynext.com

Why This Guide Matters to You

If you run, manage, or invest in apartment communities, AI regulation is no longer a "someday" topic. It is a right-now topic. In the first quarter of 2026 alone, the federal government released a national AI policy framework, HUD proposed removing its disparate impact rules, multiple states enacted new AI laws, and the RealPage antitrust settlement reshaped how revenue management software works.

Every one of these changes affects how you screen residents, price units, use chatbots, and manage your teams. The problem is that most of this information lives in dense legal documents written for attorneys, not operators.

This guide translates the regulatory landscape into plain language so you can make informed decisions, ask the right questions of your vendors, and protect your portfolio.

You do not need to be a lawyer to lead responsibly with AI. You need to understand what the rules are, where they apply, and what action to take. That is exactly what this guide delivers.

The Big Picture: What Is Happening in 2026

Two big forces are pulling in opposite directions right now. Understanding this tension is the key to understanding everything else in this guide.

Force 1: The Federal Government Wants One Set of Rules

On March 20, 2026, the White House released the National Policy Framework for Artificial Intelligence. The core message: AI is an interstate technology, and having 50 different state laws creates confusion and slows innovation. The federal government wants to replace the patchwork of state AI laws with a single national standard.

The framework also created an AI Litigation Task Force within the Department of Justice. This group is specifically looking at ways to challenge state laws that conflict with the federal approach.

Force 2: States Are Not Waiting

While Washington talks about one national rule, states like Texas, Illinois, Colorado, California, New York, and Oregon have already passed their own AI laws. Several took effect on January 1, 2026. These laws have real teeth, including fines up to \$200,000 per violation in Texas and potential class-action exposure in Oregon.

The bottom line: Federal preemption may come, but it is not here yet. The state laws are enforceable today. You need to comply with both tracks simultaneously.

What This Means in Practice

If your portfolio spans multiple states, you are operating under multiple sets of AI rules at the same time. Some states care about the intent behind your AI decisions (did you mean to discriminate?). Other states care about the effect of your AI decisions (did the outcome disproportionately affect a protected class, regardless of intent?). Your compliance approach needs to cover both.

HUD's Big Shift: Disparate Impact Rules

On January 14, 2026, HUD proposed removing its "discriminatory effects" regulations. This is one of the most significant changes in Fair Housing enforcement in over a decade. Here is what it means in plain terms.

What Changed

Before (Disparate Impact Standard)	Proposed New Standard (Intent-Based)
You could be held liable if a neutral policy had a disproportionate negative effect on a protected class, even if you did not intend to discriminate.	HUD will now prioritize investigations where there is evidence of intentional discrimination.
HUD provided a specific three-step test for evaluating disparate impact claims.	No regulatory test provided. Courts will decide on a case-by-case basis.
HUD actively interpreted how the Fair Housing Act applied to algorithms and AI.	HUD will no longer provide interpretive guidance on algorithmic effects.

What It Does NOT Mean

This does not eliminate the Fair Housing Act. Private individuals and state attorneys general can still bring disparate impact claims through the courts. The Fair Housing Act itself has not changed. What changed is HUD's role in providing a clear standard and actively investigating effect-based claims.

The practical takeaway: Do not treat this as a green light to relax your Fair Housing compliance. State regulators in California, Illinois, and New York are filling the enforcement gap that HUD is leaving. If anything, this makes state-level compliance more important, not less.

State-by-State: What You Need to Know

Not every state law affects you the same way. Here is a quick reference, followed by the details that matter most for apartment operations.

State	Law	Effective	Key Focus	Penalties
Texas	TRAIGA	Jan 1, 2026	Intent-based. Must prove purposeful discrimination.	Up to \$200,000 per violation
Illinois	HB 3773	Jan 1, 2026	Effect-based. Bans ZIP code proxies. AI disclosure for employment.	Civil rights enforcement via IL Dept. of Human Rights
Colorado	SB 24-205 (revised)	June 30, 2026	Near-total rewrite. Shifted to transparency and consumer redress.	AG enforcement with 90-day cure period
California	FEHA ADS rules	Oct 1, 2025	Automated Decision Systems subject to anti-discrimination law.	Litigation risk: AI scores may be treated as consumer reports
NYC	Local Law 144	July 2023	Annual bias audits for AI in hiring. Enforcement audit revealed major gaps.	Heightened enforcement expected after Comptroller audit
Oregon	SB 1546	March 2026	Toughest chatbot law. Mandatory disclosure, crisis protocols.	Private lawsuits: \$1,000 per violation starting Jan 2027

Texas: The Intent Standard

Texas passed the Responsible AI Governance Act (TRAIGA), which applies to any company serving Texas residents, regardless of where the company is headquartered. The law prohibits using AI with the specific intent to discriminate against a protected class.

The critical distinction: TRAIGA explicitly states that disparate impact alone is not enough to prove intent to discriminate. This is a higher bar for enforcement compared to states like Illinois or California.

Safe harbor: If you can demonstrate compliance with the NIST AI Risk Management Framework, you have a legal defense. This is not optional best-practice guidance. It is a statutory shield written into the law.

Enforcement: Only the Texas Attorney General can enforce the law. No private lawsuits. But violations carry fines up to \$200,000 each, and you get a 60-day window to fix the problem after being notified.

Illinois: ZIP Codes Are Now a Liability

Illinois integrated AI rules into the Illinois Human Rights Act, making this a civil rights issue, not just a technology issue. The law focuses on AI in employment decisions, but the principles apply broadly to how you manage on-site teams.

The ZIP code ban: Illinois explicitly prohibits using ZIP codes as a proxy for protected characteristics in AI-driven decisions. If your HR screening tools or applicant tracking systems use geographic data to filter candidates, that is a civil rights violation. ZIP codes frequently correlate with race and income, and the law treats that correlation as discrimination.

Mandatory notification: You must tell candidates when AI is involved in any part of your hiring process. Records must be kept for four years. And importantly, you cannot pass the blame to your vendor. Illinois views the employer as owning the outcome of the software they use.

Colorado: The Rewrite

Colorado's original AI law (SB 24-205) was one of the most aggressive in the country, requiring annual impact assessments and mandatory risk management programs. Then the backlash came. Governor Polis criticized the law's complexity, and in March 2026 a working group proposed a near-total rewrite.

What got removed: Annual impact assessments, mandatory risk management programs, and the requirement to report bias findings to the Attorney General within 90 days. All gone.

What stayed or expanded: Consumer protections. Colorado residents must receive notice before AI is used in a housing or credit decision. If denied, they have the right to correct inaccurate data and request a human review. Records must be kept for three years. And the definition expanded: any "computational process" that influences decisions now falls under the law, not just systems that would technically qualify as AI.

California: The Most Aggressive State

California has taken the position that AI tools used in housing and employment are subject to existing anti-discrimination law under the Fair Employment and Housing Act. But California is pushing further with the concept of "architectural authority," which means building fairness into the system at the engineering level rather than auditing for bias after the fact.

The litigation risk: In January 2026, a proposed class action argued that AI-generated applicant rankings are "consumer reports" under federal and state law. If courts agree, every AI tool that provides a score or ranking to a landlord would be subject to the same disclosure and dispute rights as a traditional credit report. That is a significant increase in administrative burden for property managers.

New York City: Enforcement Is Tightening

NYC's Local Law 144 has required annual bias audits for AI hiring tools since July 2023. In early 2026, the Comptroller's office released an audit that found enforcement was "ineffective." 75% of test calls about AI bias were misdirected or never reached the right agency. Out of 32 reviewed companies, the enforcement agency found only one violation. The Comptroller found at least 17.

What this means for you: Superficial compliance is no longer a viable strategy in NYC. The Comptroller's report has put pressure on the enforcement agency to get serious. If your organization uses AI in hiring for NYC properties, ensure your independent auditor has no financial interest in the tool being audited and is using the tool's actual historical data.

Oregon: Chatbots Under the Microscope

Oregon passed the nation's toughest chatbot law in March 2026. If you use an AI chatbot at any of your Oregon communities, these rules apply to you.

Mandatory disclosure: You must clearly tell people they are talking to an AI, not a person.

Crisis protocols: Your chatbot must be able to detect signs of suicidal ideation or self-harm and immediately interrupt the conversation to provide crisis resources.

Private lawsuits: Starting January 1, 2027, users can sue for \$1,000 per violation. In a class action, where every individual message could be counted as a separate violation, the financial exposure for a property management firm with multiple communities could reach tens of millions of dollars.

The RealPage Settlement: Revenue Management Redefined

The Department of Justice antitrust case against RealPage is the most direct regulatory impact on multifamily revenue management in 2026. The settlement fundamentally changes how pricing software works in this industry.

What the DOJ Alleged

The DOJ argued that RealPage's software facilitated a "pricing cartel" by sharing nonpublic, competitively sensitive information between landlords who should have been setting prices independently. The software made it easy to accept pricing recommendations in bulk but forced managers to override them one at a time, with mandatory commentary that was escalated to regional leadership. The DOJ described this as a system designed to discourage managers from lowering rents.

What Changed

Rule	What It Means
Data must be at least 12 months old	Revenue management AI can no longer train on real-time competitor data. Pricing must be based on your own property's performance and historical trends.
No sub-market granularity	Pricing information cannot be reported more narrowly than at a state-wide level. No more benchmarking against the property across the street.
No identical recommendations	The software cannot make the same pricing recommendation to different owners in the same market.
No coercive override design	Managers must be free to lower rents without facing escalation, mandatory commentary, or monitored rejections.
Court-appointed monitors	Compliance will be overseen by independent monitors for up to seven years.
Vendor certification required	Landlords must obtain a signed certification from their software vendors confirming the tool meets these restrictions.

If your revenue management vendor has not proactively reached out to explain how they are complying with the settlement terms, that is a question you need to ask immediately.

Chatbots and Screening: Legal Precedents You Need to Know

Two landmark settlements in the PropTech space have established clear legal precedent: the property manager shares responsibility with the software developer for AI outcomes. "The algorithm did it" is not a defense.

SafeRent Solutions (\$2.3 Million Settlement)

SafeRent's tenant screening algorithm was found to discriminate against housing voucher holders because it failed to account for the financial benefit of the voucher in its scoring model. The court ruled that national tenant screening providers are subject to the Fair Housing Act. SafeRent is now prohibited from issuing "approve" or "decline" recommendations for voucher holders unless verified by civil rights experts.

Harbor Group / PERQ (Chatbot Discrimination)

A settlement was reached after a lawsuit alleged that an AI chatbot automatically rejected Section 8 applicants. The settlement requires the software developer to audit its AI tools for Fair Housing compliance twice per year. This case confirmed that deployers, meaning property managers, share legal responsibility with developers for the outcomes their AI tools produce.

These cases establish a clear standard: you own the outcomes of the AI tools you use at your properties. Vendor diligence is not optional. It is a legal obligation.

Your Compliance Playbook: NIST and ISO Frameworks

With no single federal law governing AI in housing, multifamily operators are converging on two primary frameworks to establish "reasonable care" and satisfy state-level safe harbors. You do not need to master these frameworks in detail, but you need to understand what they are and why they matter.

NIST AI Risk Management Framework (AI RMF 1.1)

The National Institute of Standards and Technology created a practical, step-by-step approach to AI governance. The framework has four core functions.

Function	What It Means	Your Action
GOVERN	Build a culture of AI risk management.	Charter an AI governance committee. Assign a named owner for each AI tool you use.
MAP	Understand where AI risks exist in your operations.	Map data flows for screening tools, chatbots, and pricing software. Know what data goes in and what decisions come out.
MEASURE	Test and track AI outcomes with evidence.	Conduct bias testing. Look at demographic parity in screening outcomes. Document everything.
MANAGE	Act on what you find.	Build incident response playbooks. Establish "deployment gates" before new AI tools go live.

Why this matters legally: In Texas and Colorado, alignment with NIST AI RMF creates a statutory safe harbor. That means compliance with this framework gives you a legal defense if your AI use is challenged. This is not just a best practice. It is a shield.

ISO/IEC 42001: Certifiable AI Management

ISO 42001 is the world's first certifiable AI management system standard. While NIST provides the structure, ISO 42001 provides the certification that enterprise buyers, regulators, and insurers are increasingly requiring. Insurance carriers have begun introducing "AI Security Riders" that require documented evidence of risk assessments. ISO 42001 certification gives you that documentation.

What to Do Now: Your Action Items

This is not a "someday" list. These are the steps your leadership team should be taking in 2026 to protect your portfolio and position your organization to use AI responsibly and effectively.

1. Assign an AI Owner for Every Tool

Every AI system you use, whether for pricing, screening, chatbots, or employee management, needs a named person who is accountable for its compliance, performance, and outcomes. This is the single most important governance step you can take.

2. Audit Your Revenue Management Software

In light of the RealPage settlement, confirm with your vendor in writing that no real-time competitor data is being used. Require a signed certification that the tool meets the DOJ's settlement restrictions. If your vendor cannot provide that certification, treat it as a serious risk.

3. Review Your Chatbot Compliance

Does your chatbot clearly disclose that it is AI? Does it have crisis protocols? Does it screen voucher holders fairly? If your communities operate in Oregon, the stakes are especially high. Audit now, before the private right of action begins in January 2027.

4. Check Your Screening Tools for Voucher Compliance

Following the SafeRent settlement, ensure your screening provider accounts for the financial benefit of housing vouchers in its scoring model. If the tool issues blanket approve/decline recommendations for voucher holders without expert review, that is a liability.

5. Build Your NIST Alignment

Start with the GOVERN function: charter a governance committee and map every AI tool in your portfolio. You do not need to complete the entire framework in 90 days. But you need to start. In Texas and Colorado, this creates a legal safe harbor.

6. Adopt the "Strictest Common Denominator"

If your portfolio spans multiple states, build your enterprise AI standards to the most restrictive law that applies to you. This is simpler than maintaining different compliance programs for each jurisdiction, and it protects you everywhere.

7. Update Employee Handbooks

If you use AI in any employment-related decisions (hiring, promotions, performance reviews), add AI disclosure language to your 2026 handbooks and onboarding materials. Illinois requires this, and it is a best practice everywhere.

8. Stop Relying on Vendor Promises

"The vendor told us it was compliant" is not a legal defense. In Illinois, the law explicitly rejects it. In every other jurisdiction, the property manager is viewed as owning the outcome of the software they use. Require documentation, not reassurances.

ABOUT MULTIFAMILY NEXT

Multifamily NEXT is the multifamily industry's first hands-on AI implementation event series. Founded by Tami Siewruk, a 49-year industry veteran, Multifamily NEXT is built on one principle: Build, don't listen.

Our 2026 National Tour visits nine cities: Atlanta, Phoenix, Charlotte, Orlando, Chicago, Baltimore, Los Angeles, Denver, and Houston. Each event delivers 2.5 days of hands-on AI workshops where attendees build working tools, establish governance frameworks, and leave with skills they can put to work immediately.

Every workshop comes with our "Until You Win" guarantee.

Build AI Fluency. Develop Real Skills. Create Working Tools.

Learn more and register at multifamilynext.com

This guide is for informational and educational purposes only and does not constitute legal advice. Consult qualified legal counsel for guidance specific to your organization and jurisdiction.