Which door do you expect?

Across the whole development process, from land to lease-up

Faster, cheaper, fairer development

Quietly more exclusionary

Two ways to pick a site

CONVENTIONAL DEVELOPMENT

You underwrite a market

WHO YOU HAVE TO SATISFY

Tenants and buyers

HOW YOU KNOW YOU ARE RIGHT

Rents, absorption, land prices

THE SITE QUESTION

Will it lease, and will it exit?

TAX-CREDIT HOUSING (LIHTC)

You underwrite a rubric

WHO YOU HAVE TO SATISFY

The allocating agency (NCHFA)

HOW YOU KNOW YOU ARE RIGHT

A published point scale, one competitive round a year

THE SITE QUESTION

Will it score, and will it close?

A published scoring rule is exactly what AI optimizes. That is the promise, and the risk: whatever the rubric leaves out becomes invisible to the optimizer.

A market study still gates every NC application (capture rate, absorption, vacancy), so it is a rubric with a market check.

How a 9% deal works, in plain terms

What is being won?

Federal tax credits. Investors buy them, and that money becomes the project's equity. In return, rents stay restricted for 30 years in NC.

Who decides?

The state housing agency, NCHFA, scores every application against its published rulebook, the Qualified Allocation Plan (QAP).

How often?

The 9% credit: one competitive round a year. The 4% credit is a separate path tied to tax-exempt bonds, filed May through October.

THE 2026 CALENDAR (NC, 9% CREDIT)

Jan 23

Preliminary application. Site control (option, contract or deed), a plot plan, and a $6,200 nonrefundable fee.

Mar 20

Final site scores announced.

May 15

Full application. Required zoning must be in place.

August

Tax credit awards.

You must control the land before you know your score. That is why the site hunt is where AI is being sold.

68 site points, mostly geometry

QAP = Qualified Allocation Plan: the state housing agency's published scoring rules.

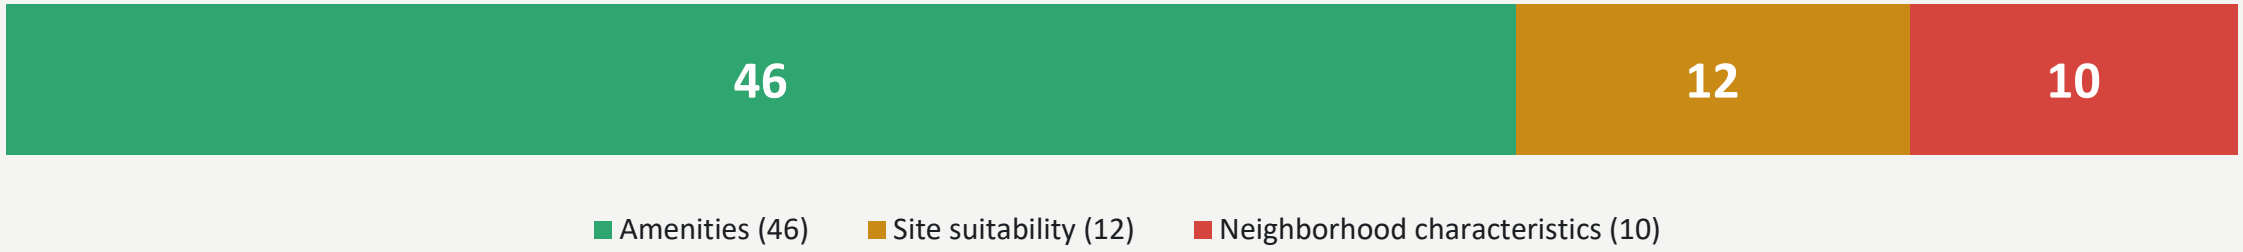

MEASURABLE

Amenities: 46

Google Maps driving distance to a named grocer (12 points within 2 miles), shopping and pharmacy, plus services, parks, schools, and a sheltered transit stop within a quarter mile.

MOSTLY CHECKABLE

Site suitability: 12

Incompatible uses nearby, flood and slope, safe traffic access, visibility from active buildings. GIS layers plus some judgment.

JUDGMENT

Neighborhood: 10

Are structures within a half mile "well maintained," visibly deteriorated, or blighted? A reviewer's eye, or a model's.

Source: NCHFA 2026 QAP, Sec. IV(A)(1). Minimum total site score is 50; up to 3 negative points for sites deemed unsuitable.

THE OPTIMIST'S CASE

Score every parcel before January.

Screen every parcel in a county

Run distances, incompatible-use buffers, flood and slope checks for thousands of parcels, not the dozen you already know.

Assemble the evidence as you go

The Agency wants route maps in satellite view and a photo of each amenity. A tool can build that file while it scores.

Choose which sites go forward

A principal can be in at most five applications and win at most two, so ranking a pipeline is a real optimization problem.

ALREADY HERE

Parcell

Markets AI site scoring for LIHTC-eligible parcels.

AURA (Feb 2026 preprint)

Frames siting as constrained optimization under QCT, DDA and LIHTC rules.

Syndicators

Using AI for data extraction and analysis (Tax Credit Advisor, July 2026).

Where AI helps, and where it can hurt

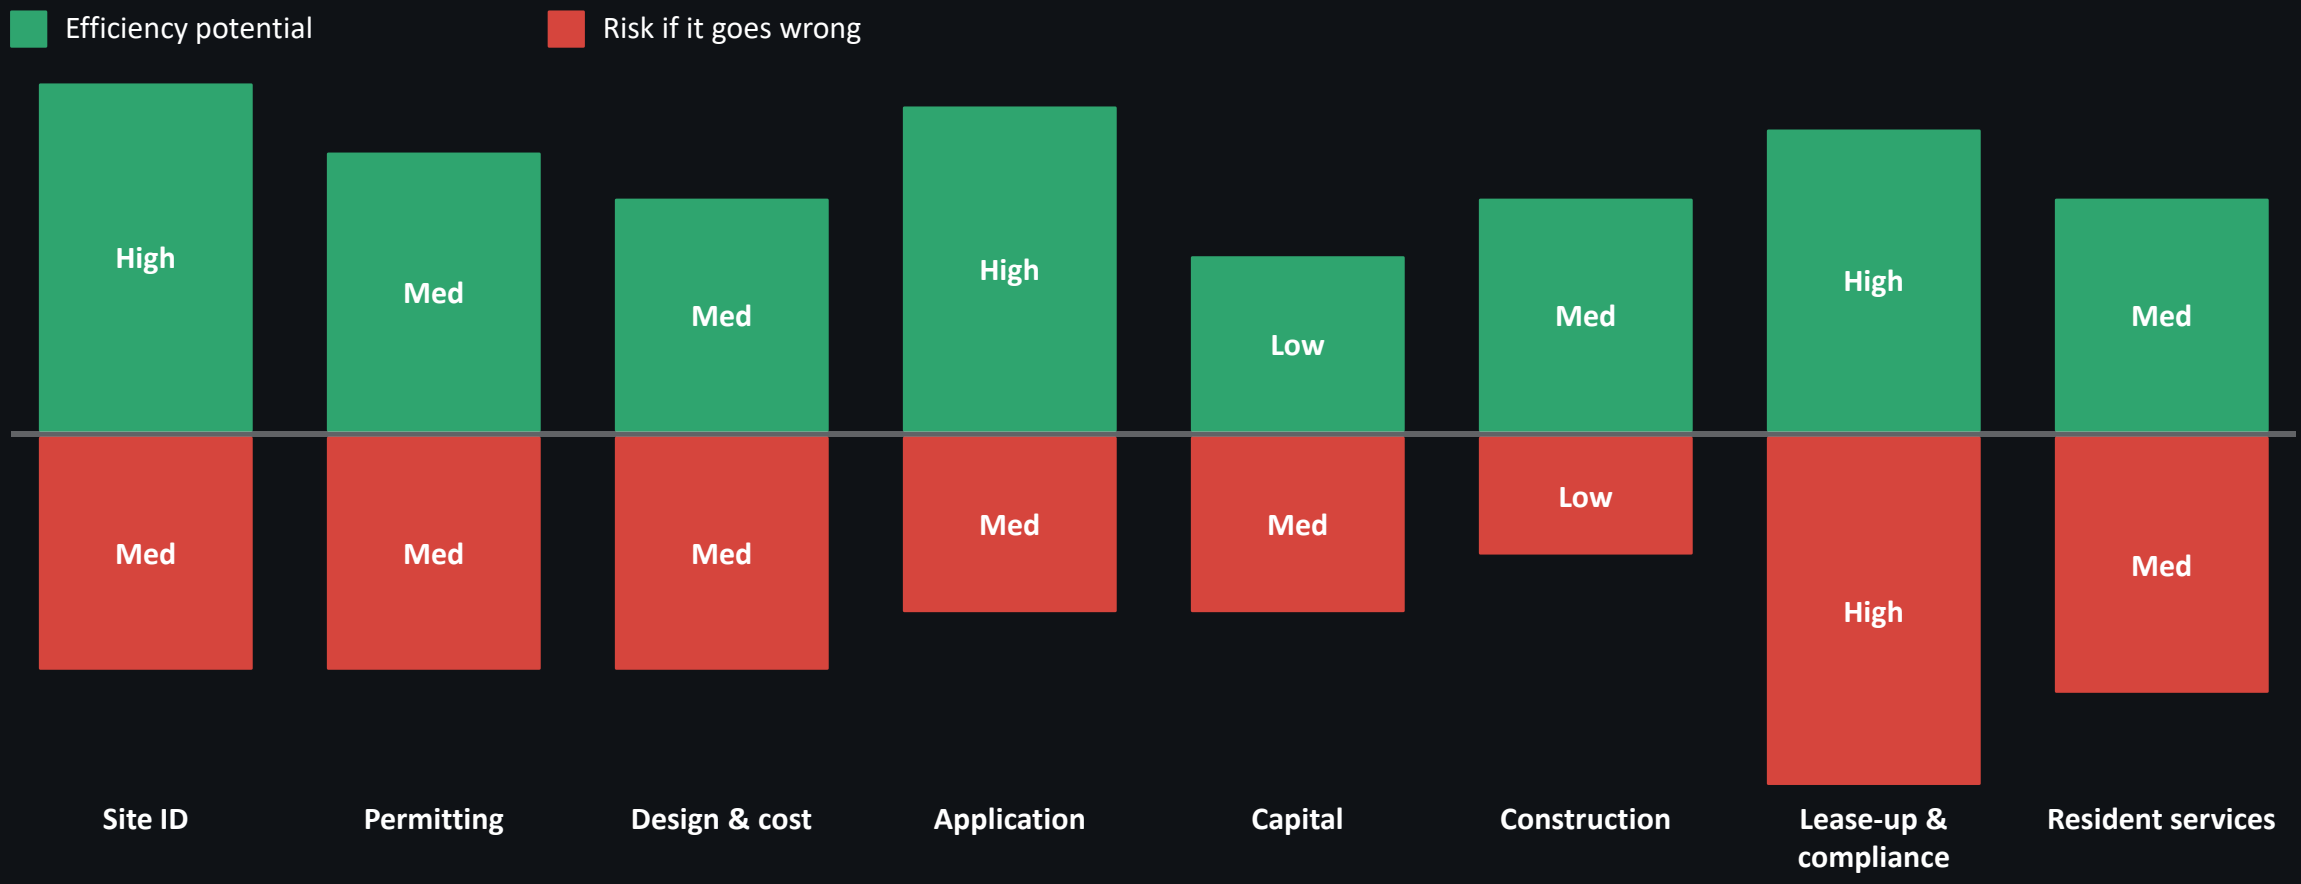

A perfect 68/68 site score still needs site control, zoning, utilities and neighbors who say yes. Efficiency clusters in paperwork; risk clusters where decisions touch people.

Illustrative: the author's judgment, not measured data.

AI on the other side of the desk

CITY SIDE: PERMITTING

3.4 → 1.4

review cycles per application after AI pre-screening (Honolulu, single- and two-family homes)

45–60 days

expected cut to initial review (Jacksonville, AI-assisted review; human reviewers keep final authority)

HELPS

- Flags incomplete applications before a reviewer opens them
- Cities from Denver to Jacksonville are adopting; HUD is funding pilots

RISKS

- Accuracy on complex zoning is questioned by some officials and local architects
- Reporting says applicants cause at least half the delay; evidence is for simple residential

When applicants and agencies both use AI, the rubric is optimized by one machine and enforced by another. **Who audits whom?**

Sources: Stateline (Aug 2026), HousingWire, local news reporting. Reporting is not independent evaluation.

Optimizing to the cliffs

Site ID | Permitting | **Design & cost** | Application | Capital | Construction | Lease-up & compliance | Resident services

Illustrative: where applications might land if everyone optimizes against one published line

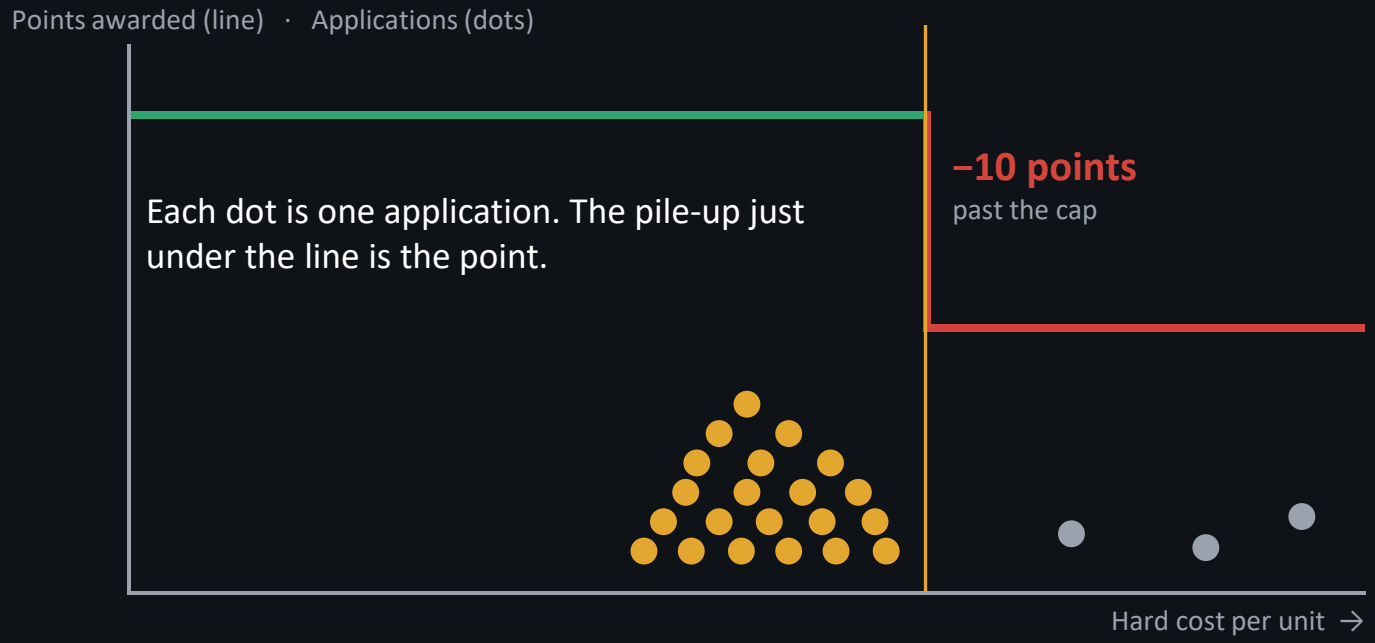

$135,000 per unit
Chart A hard-cost threshold; $150,000 for some project types

WHAT THE QAP COUNTS

- **Design:** up to 25 points for countable features: roof lines, siding and brick variety, porches, railings
- **Cost:** −10 points if hard costs top $135,000 per unit ($150,000 for some project types), per QAP IV(C)(1)
- **Cap:** $350,000 per unit total, excluding land, unless waived

Anything that is counted can be optimized, including what is easy to count. **Is the best building the one that scores best?**

Illustrative chart, not data. Thresholds from NCHFA 2026 QAP, Sec. IV(C)(1) (Chart A) and IV(G)(2), VI(B)(6)(e), III(C)(2). No documented AI design tool in LIHTC found; this is a hypothesis.

The compliance tail

Site ID | Permitting | Design & cost | Application | Capital | Construction | **Lease-up & compliance** | Resident services

A single clerical error can climb a ladder. Each rung is in the QAP or in industry reporting.

1	**A certification error**	Industry-sponsored coverage: credits at risk, "millions" in the worst cases	REPORTED
2	**A fine, per instance**	Up to $2,000 each time a written requirement is not met [III(B)(7)]	IN THE QAP
3	**Uncorrected for over 3 months**	Principals may be disqualified from future awards [IV(D)(3)(e)]	IN THE QAP
4	**An adverse fair housing finding**	Within 10 years, may disqualify the owner, principal or management agent [IV(D)(3)(b)]	IN THE QAP

The AI may not be the thing that breaks. A screening or chatbot error can still reach your next application. **Who signs the certification?**

Sources: NCHFA 2026 QAP; Multi-Housing News (vendor-sponsored survey). The QAP says "may"; this is a stakes ladder, not a documented AI incident.

Closing the gap, or moving it?

Where the rules already level the field

Nonprofit set-aside

10% of the state credit ceiling goes to nonprofit-sponsored projects

CHDO set-aside

15% of the Agency's HOME funds are reserved for CHDOs

Caps on the big players

At most two new-construction awards and $2.6M in credits per principal

Where AI could tilt it, or can't help

Speed to site control

Whoever acts on model output first ties up the land, with option money at risk

Application capacity

A well-resourced shop can screen and document five applications at once

The experience threshold

A principal needs one NC 9% project, or six totaling 200+ units. AI does not fix that

Hypothesis: the most exposed may be the mid-sized for-profit developer, with neither a set-aside nor scale.

CHDO: a nonprofit certified for federal HOME funds. Principal: the owner and developer parties named on an application.

Raleigh, today

Median household income by census tract

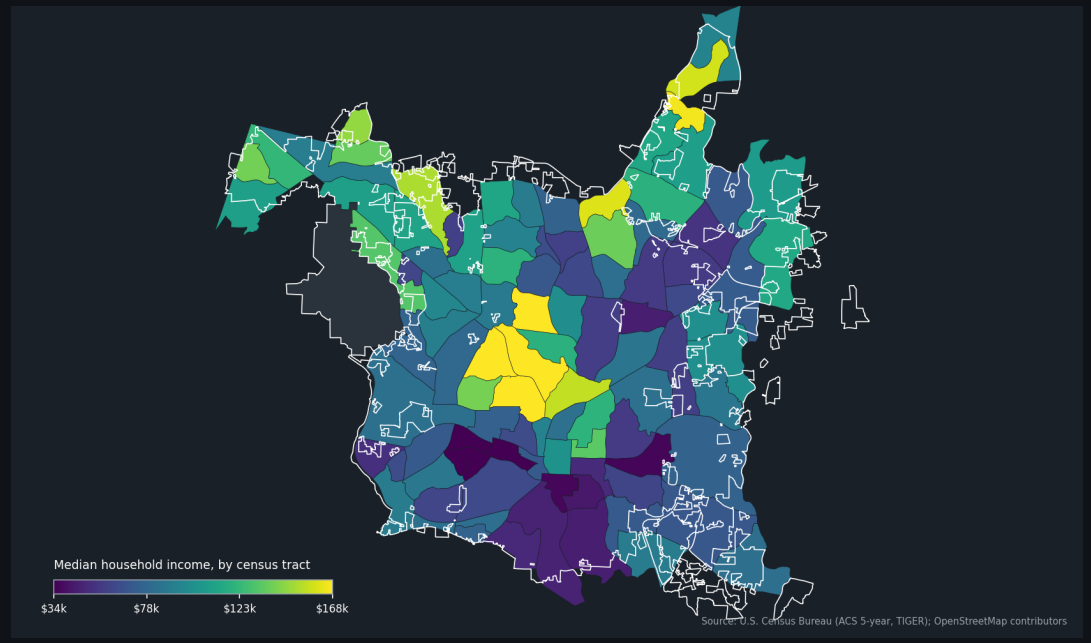

Within about 2 miles (straight line) of a named grocer

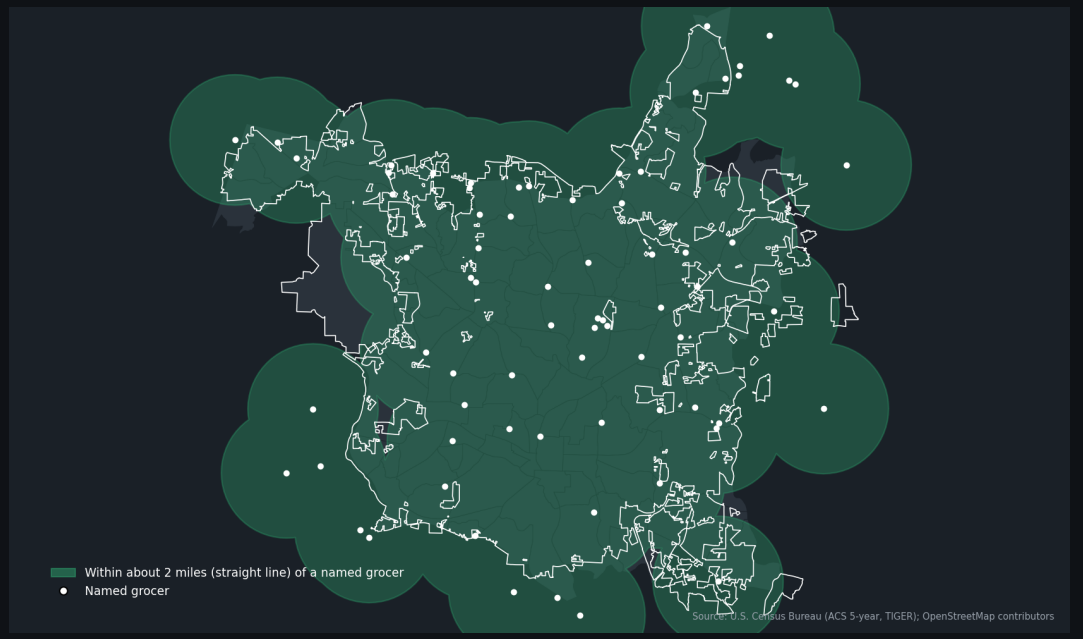

When nearly every site passes, the test can't tell them apart.

By straight-line distance, 92% to 100% of residents in every income quartile live in tracts within about 2 miles of a named grocer. When nearly everyone scores the same, the points that decide are elsewhere.

Sources: U.S. Census Bureau (ACS 5-year, TIGER); OpenStreetMap contributors. The 2-mile radius is a straight-line approximation of the QAP's driving-distance test.

The 10 points a model would grade by eye

10 — **Well maintained**

Structures within a half mile look well kept, or the site is a Redevelopment Project.

5 — **Visible deterioration**

Structures within a half mile are not well maintained.

0 — **Blighted or fortified**

Blight, or security modifications such as bars on windows or barbed wire.

The optimist's answer

Consistent, scalable, and auditable. You can test a model for disparate impact in ways you cannot test a reviewer's eye.

The skeptic's answer

Only if someone runs the audit. And "well maintained" tracks income as much as need, so a 0 can lock out the blocks that need housing most.

Contested policy, hard-coded: the QAP bars sites in areas of minority and low-income concentration, with a revitalization exception. A tool will enforce it as a filter.

Three futures

1 Filter

AI as a first-pass screen; humans and the community still decide

True if: human review of judgment criteria; cheap, widely available tools

2 Gatekeeper

Scores get baked into agency and lender processes

True if: scores are proprietary and unexplained; reviewers defer to the number

3 Split market

Advantage concentrates with the best-resourced

True if: land goes to whoever moves first; lenders trust institutional analysis

IF YOU DEVELOP

1. **Screen with AI, select with your feet.**
2. **Build the evidence file as you score.**
3. **Log what the model rejected, and spot-check it.**
4. **Price the real bottleneck: land, entitlements, gap.**

Investors

Are you pricing in model risk, or just the speed?

Lenders

Could you underwrite a deal you can't explain if the model is wrong?

Policymakers

Which QAP criteria should stay human-judged?

Vote again.

Whichever door you picked, the work is the same: keep humans in the loop, and make the model show its work.

