STRATEGIC ADVISORY

# AI Opportunity & Governance Assessment

Where AI genuinely helps your business, where it doesn't, and how it runs without putting your license, your brand, or your clients at risk.

**Prepared for: Meridian Realty Group — a 14-agent residential real estate team, Nassau County, NY (fictional illustrative client)**

**SAMPLE DELIVERABLE**

Meridian Realty Group is a fictional business. This document shows the exact format, depth, and honesty of the assessment Fortera delivers to real clients.

Fortera Labs LLC · One World Trade Center, Floor 85, New York, NY 10007 · forteralabs.com

# 1. Executive Summary

Meridian's biggest AI opportunity is not marketing content or chatbots — it is the gap between when a lead reaches you and when a human responds. Meridian currently responds to after-hours portal inquiries the next business morning. In residential real estate, response time is the single strongest controllable factor in lead conversion, and roughly half of Meridian's inbound inquiries arrive outside business hours.

Three findings drive this assessment:

- **Speed-to-lead is the highest-value fix.** Automated acknowledgment and qualification within minutes, with a human taking over every conversation — no bot ever negotiates, advises, or discusses terms.
- **Transaction coordination is drowning in manual status work.** Deadline tracking and status updates across ~9 active transactions consume roughly a day of admin time per week that automation can absorb.
- **Two commonly-pitched AI uses are wrong for Meridian right now:** AI-written listing descriptions (brand risk exceeds time saved at your volume) and automated buyer-matching outreach (fair-housing exposure with no compliant guardrail we would sign off on). We recommend against both, and we say so up front because that is the point of governance.

**The 90-day recommendation: implement after-hours lead response and CRM hygiene automation first, transaction-deadline automation second, and revisit content automation only after the first two are stable. Every automated action that a client could ever see is gated behind human approval.**

# 2. Current-State Snapshot

Compiled from owner and admin interviews and a walkthrough of Meridian's actual working systems (CRM, portal accounts, shared inbox, transaction checklists).

| Function | How it works today | Where it strains |
|---|---|---|
| Lead intake | Portal leads (Zillow, Realtor.com), sign calls, and website form land in a shared inbox; agents claim manually. | No response outside business hours; no ownership rule — fastest inbox checker wins; some leads never get a first touch. |
| Follow-up | Each agent follows up from personal notes; CRM updated “when there's time.” | CRM is ~60% current; re-engagement of 30–90 day-old leads effectively doesn't happen. |
| Transaction coordination | One coordinator tracks deadlines for ~9 active deals in spreadsheets and calendar reminders. | Manual status chasing (attorneys, lenders, inspectors) ≈ a full admin day weekly; deadline risk concentrated in one person. |
| Listing operations | Agents write descriptions; photos via vendor; posting is manual per channel. | Slow but low-risk; not a bottleneck at current volume (~6 new listings/month). |
| Marketing | Sporadic social posts; no consistent nurture to past clients or sphere. | Referral engine underused — past-client contact is ad hoc. |

## 3. Opportunity Map

Each opportunity is scored for impact and effort, and — because this is a licensed, regulated business — for risk. Anything touching protected-class considerations, advice, or negotiation is either human-gated or off the table.

| # | Opportunity | Impact | Effort | Risk | Verdict |
|---|---|---|---|---|---|
| 1 | After-hours lead acknowledgment & qualification — instant reply captures intent, timeline, and preferred contact; human takes over next morning (or is paged for hot leads). | High | Low | Low — scripted, non-advisory, disclosed as automated | Do first |
| 2 | CRM hygiene automation — every inquiry auto-logged, tagged, and assigned by rule; stale-lead flags surface weekly. | High | Low | Low — internal only | Do first |
| 3 | Transaction deadline engine — dates extracted at contract, reminder cadence to all parties, human approves every outbound message. | High | Med | Low–Med — outbound gated | Do second |
| 4 | Past-client & sphere nurture — quarterly touch drafts (home anniversary, tax season, market note) queued for agent approval. | Med | Low | Low — nothing sends unapproved | Do second |
| 5 | Weekly ops digest — lead volume, response times, pipeline by stage, deadline exposure in one owner email. | Med | Low | None | Do second |
| 6 | AI-drafted listing descriptions | Low | Low | Med — brand voice + accuracy liability on MLS | Not now — revisit at 2× listing volume |
| 7 | Automated buyer-match outreach (“this home fits buyers like you”) | Med | Med | High — fair-housing exposure: any targeting proxy for protected classes is untenable | Recommend against |
| 8 | Public-facing AI chat on website answering market/valuation questions | Low | Med | High — unlicensed-advice territory | Recommend against |

Note what is *not* here: nothing that negotiates, advises on price, or communicates with a client without a human having approved the message. That is a permanent design rule, not a phase-one limitation.

## 4. How It Stays Governed

Fortera's delivery methodology (FIPOS) comes from a decade of shipping automation inside highly regulated financial services, where nothing goes live without proving it can be supervised, audited, and switched off. Translated to plain terms, three rules govern everything we run for Meridian:

### Rule 1 — Humans approve anything a client can see.

Automated systems draft, queue, and remind. A named person at Meridian approves every outbound message beyond the initial scripted acknowledgment — one tap, from a phone. Approval queues are designed to take minutes per day, not hours.

### Rule 2 — Compliance lines are engineered in, not policied in.

Fair-housing safety is structural: automations never segment, score, or target by any attribute that could proxy for a protected class — they operate on behavior only (inquired, viewed, went stale). Text/SMS automation follows opt-in consent rules from day one. Where a use case can't be made structurally safe (opportunities 7 and 8), we don't build it.

### Rule 3 — Everything is visible and reversible.

Every automated action is logged in plain English — what ran, when, why, approved by whom. You can read the log, and any automation can be paused instantly without breaking the rest. If you ever end the engagement, the CRM, the data, and the documentation are yours; there is no lock-in by design.

# 5. The First 90 Days

| Phase | What happens | What you should see |
|---|---|---|
| Days 1–30 Foundation | CRM cleanup and ownership rules; after-hours acknowledgment live on portal + website leads; response scripts approved by you before anything activates. | Every lead gets a first touch within 5 minutes, 24/7. Zero leads unowned. You approve every word the system can say. |
| Days 31–60 Coordination | Transaction deadline engine live on active deals; weekly ops digest starts; nurture drafts begin queuing for approval. | Deadline reminders go out without chasing; the Monday digest replaces “where are we on…” meetings. |
| Days 61–90 Steady state | Tune from real usage; retire what isn't earning its keep; quarterly advisory rhythm set. | A measured baseline: response times, pipeline movement, admin hours recovered — the numbers that decide what, if anything, we build next. |

**By day 90 you have data instead of impressions. The decision to expand, hold, or cut is made from your numbers — and we will recommend cutting anything that isn't working.**

# 6. What Fortera Runs vs. What You Own

| Fortera runs (managed service) | Meridian owns |
|---|---|
| Build, monitor, and maintain every automation; approval queues and escalation paths; monthly tuning; the audit log; a quarterly advisory session with the owner. | All client relationships and every client-facing decision; final approval on all outbound content; the CRM, the data, and all documentation — portable at any time. |

This is the standing structure of a Fortera managed-service engagement: we run the machine; you run the business; nothing reaches your clients without your sign-off.

**About this sample:** Meridian Realty Group is fictional, but the assessment structure, the scoring discipline, the recommendations against certain AI uses, and the governance rules are exactly what a real Fortera engagement delivers — grounded in your systems, your numbers, and your market.