

Municipal Housing Service Transformation

Omid Ganji | Independent BA portfolio

A fictional municipal housing-support service experiences incomplete applications, repeated handling and slow decisions. This case study connects the problem to process, requirements, data, reporting and a controlled pilot proposal.

Synthetic baseline

2,400 applications received in 2025; 1,393 first-pass complete (58.0%); 46-day median decision cycle among closed cases. Year-end backlog: 394 cases. Follow-up cutoff: 31 March 2026.

Contribution and evidence

Problem framing, simulated stakeholder discovery, AS-IS / TO-BE modelling, requirements and rule governance, SQL / ETL analysis, a native six-page Power BI report, a sample intake experience and validation evidence.

What exists vs what is proposed

The portfolio artifacts, report and sample website exist. A municipal rollout, real-user research, production integration and operational benefits have not occurred.

Frame the problem before selecting technology

Problem signals

42.0% incomplete at first review; 9.6 average manual touches; 29.8% of closed cases within fictional SLA thresholds. These signals motivate discovery but do not prove causes.

Stakeholder and governance work

11 stakeholder groups and eight RACI decisions. Discovery is simulated, not interviews with real municipal staff or residents. Applicants and frontline staff inform design; policy and control owners retain approval authority.

Recommendation and trade-off

Phased service orchestration: improve intake and case creation first, then status and reporting. Retain assisted channels and human eligibility decisions. Costs and benefits require real validation.

Inspect the work

Business case: </project/business-case>

Stakeholder map and downloadable RACI: </project/stakeholders>

Turn the problem into testable controls

AS-IS to TO-BE

The AS-IS shows document chasing and manual handoffs. The TO-BE checks fields and conditional evidence before case creation, standardizes routing and keeps consequential decisions with people.

Decisions and exceptions

Editable public-review BPMN models show completeness branches, resubmission, supervisor review and a non-interrupting review timer. These are non-executable analysis models, refined from the original drafts.

A concrete trace

PP-01 incomplete intake -> FR-001 dynamic checklist -> RULE-002 evidence conditions -> UAT-003 conditional-document tests -> KPI-02 first-pass completeness. The baseline contains 15 requirements and eight governed rules.

Inspect the work

BPMN and sample journey: </project/process-redesign>
Requirements and UAT coverage: </project/requirements>

Show how the numbers are produced

Reconciled synthetic data

2,400 applications; 10,341 document records; 15,617 status events; 8,674 communications; 852 survey responses. These counts replace the earlier draft counts.

SQL proof

The baseline query returns 2,400 cases, 0.5804 completeness and 9.61 average touches. Re-executed read-only on 5 September 2026. Keep application totals at case grain when working with child events.

Actual Power BI model and report

The verified semantic model uses Application Data as the case hub. The active date relationship uses received date; closure date is inactive. Six native report exports cover executive overview, flow, documents, equity, team operations and data quality.

Interpretation limits

The public report images are static. Zero currently missing documents does not mean intake was complete: 7,270 were submitted later. Survey experience describes respondents, not all applicants. Inspect: [/project/measurement](#)

Preserve failures, corrections and conditions

Historical simulated UAT

On 1 September 2026, 12 scenarios were executed with three seeded defects and retests. The 15 execution records retain failures rather than hiding them. This is case-study evidence, not production acceptance.

One traceable defect

UAT-006 first exposed an internal risk note in applicant-facing status (DEF-002, Critical). The model was corrected to allow-list public fields and passed retest. A formal privacy assessment remains required.

Real-pilot gates

Before a real pilot: validate policy authority, usability with real applicants and staff, accessibility, privacy, security, performance, reconciliation, operational support and rollback. Simulation cannot replace these checks.

Targets are hypotheses

80% first-pass completeness, 25-day median cycle and 85% SLA attainment are strategic hypotheses, not achieved results or automatic launch gates. Lower pilot continuation gates support controlled learning.

A portfolio you can inspect

Three ways to review

Read the seven phase pages, inspect the actual report, or try the sample-only intake journey. The demo does not submit an application, upload documents or send email.

Report and demonstration

Power BI:

<https://municipal-housing-ba-portfolio.omid-industry4.chatgpt.site/project/measurement#power-bi-report>

Intake demo: <https://municipal-housing-ba-portfolio.omid-industry4.chatgpt.site/application-demo>

Discuss the decisions

Why validate at entry? How are exceptions owned? Why preserve assisted access? How do SQL totals reconcile? What does simulated UAT prove, and what still needs real-world assurance?

Evidence boundary

Independent fictional case study with synthetic data and simulated discovery / testing. No City of Toronto or municipal affiliation. No claimed employment, client engagement or achieved operational improvement.