Summary of Agile Pilot PrimAssist

Company name	**GESEIDL EURO FINANCE CONSULTING SRL**
Company location	**Ploiești, Romania**
Domain	**Digital services for public administration — AI-assisted citizen-to-municipality interaction**
Municipality	Comuna Măgurele, Prahova County, Romania (rural, ≈5,000 inhabitants; Mayor Vasilică Diaconu)
Project period	Preparation Nov 2025 - Feb 2026; contract in force 09.03.2026 (No. 2413/02.03.2026, EUR 12,000 + VAT); development & testing March-May 2026; reporting June 2026.
Solution	PrimAssist is an AI citizen assistant that answers questions on municipal procedures (fiscal, urban planning, civil status, social assistance, public relations) exclusively from indexed municipal documents — no fabricated information, every answer grounded in a cited source. The pipeline is multi-agent: planning agent → domain specialist → challenger agent (validates against the source) → synthesizer → GDPR guardrails. Citizens interact via tablet (hall kiosk), web browser or Telegram share; voice input via Whisper STT. Stack: FastAPI + Next.js + a cloud LLM with PageIndex tree-based retrieval; GDPR-compliant by design (no conversation persistence, no tracking); intranet-first deployment on a single server inside the municipality. This phase was a technical validation (proof-of-pilot) of feasibility and quality, not a production operation with citizens.
Stakeholders	Municipality of Comuna Măgurele (pilot site, document provider); UTCB — Technical University of Civil Engineering Bucharest (coordinator, academic advisor); Geseidl Euro Finance Consulting SRL (solution provider); Wallachia eHub / Danube Engineering Hub + Spiru Haret University (complementary Digital Europe co-financing, mentoring, networking). Targeted end users: the municipality's citizens and clerks.
Lessons learned	• Multi-agent debate works against hallucination — a separate challenger agent validates each answer against its source. • Anti-defensive prompting is non-trivial and required a dedicated sprint to avoid the "no information found" reflex. • Tablet kiosk deployment exposes hidden assumptions (PDF rendering, soft keyboard, microphone HTTPS). • Curation quality beats corpus size. • Cost discipline (caching + small corpora) is critical for small municipalities. • Overall, the agile pilot succeeded as a technical validation; measuring operational impact requires a follow-up phase with real users.

KPI 1 Adoption & usage	Targets: >85% substantial answers, >85% correct-source retrieval (top-K), <10% defensive redirects. Results (May 2026): 100% substantial answers, 87.5% correct-source retrieval, 0% defensive redirects. Two test rounds — a representative 10-question set across 4 domains (08.05.2026) and a targeted 4-question municipal set after a retriever fix (18.05.2026).
KPI 2 Time savings	Targets: <20 s average latency per question, 0 connection errors, 0 GDPR/accessibility issues. Results: average latency 11–14 s (best case ~8 s after a single-call optimisation), 0 connection errors, full GDPR conformity (no conversation persistence, PII sanitisation, WCAG AA). Note: the operational impact indicators committed in the application — ≥40% faster access to information, +25% requests processed per clerk, −30% citizen returns, ≥85% clerk satisfaction — require a production setting with real users and are deferred to the post-pilot operational phase.
KPI 3 User experience	KPI 3 measures the percentage reduction in citizens' returns for the same request, after the introduction of PrimAssist, compared to the reference situation (without assistant). It is expressed as: KPI 3 (%) = ((RR_reference - RR_PrimAssist) / RR_reference) x 100 Target: ≥ 30%
Evaluation of the business model focusing on its viability and potential for growth	Target customers: small and medium Romanian municipalities seeking affordable, GDPR-compliant AI for citizen information. Before the pilot, Geseidl's penetration of the public-administration market was zero. Commercial offering post-pilot with subscription tiers (Basic 500–1,000 RON/month; Standard 1,500–3,000; Premium 4,000–8,000; regional SaaS Pilot 200–500). Estimated operating cost ≈50–200 EUR/month cloud LLM per municipality. No vendor lock-in — the LLM provider is swappable, including a local Ollama option for fully offline deployments. Strengths: demonstrated multi-stack delivery and trusted access via UTCB and Wallachia eHub co-financing. Barriers: limited development capacity and no dedicated public-administration sales/marketing function yet. The pilot validated technical feasibility and a low total cost of ownership; profitability at scale depends on adding human resources and business development.
Impacts	For Comuna Măgurele: first practical experience with AI-assisted citizen interaction; exposure to document-indexing and GDPR-compliant deployment; a reference solution available for operational use after administrative approval. For Geseidl: validation of the multi-agent architecture in the Romanian administrative context; a replicable delivery process and reusable transfer artefacts (indexing recipes, KPI rubric, prompt patterns); entry into the public-administration market. For the partnership: a working academia–SME–local-authority model (UTCB–Geseidl–UAT Măgurele, associated with Wallachia eHub) transferable to other Danube cities. Operational impact on citizens and clerks (time saved, productivity, satisfaction) is expected but will be measured only in the operational phase.

Suggestions for future actions, especially focusing on sustainability and replication	Run an operational phase with real users to measure the committed impact indicators that a technical pilot cannot capture. Add a continuous indexing pipeline (scheduled crawl + re-indexing) and proactive website consultation to reduce manual curation. Extend the benchmark to ~30 questions with a periodic groundedness audit; resolve the partial "vehicle tax" retrieval case. Extend Whisper voice beyond Romanian and complete EU AI Act Article 50 disclosure for cross-border replication. Replication requires only localisation (document corpus + prompts); the architecture is language-agnostic and reusable across BG, HU, CZ, SK, SI, HR.
Next steps	Q3 2026 — replication in 2–3 Prahova municipalities (Măgurele as reference) plus an operational phase to measure impact. Q4 2026 — regional SaaS Pilot package via Wallachia eHub. H1 2027 — national rollout (20–30 municipalities) and integration with national platforms. H2 2027 — cross-border pilot in 1–2 Danube countries via an Interreg follow-up call.
Provider's Reflection	The agile piloting model proved decisive in turning a generic LLM into a reliable, grounded administrative assistant. The challenger-agent architecture is a transferable pattern. The technical pilot validated feasibility and quality; an operational phase with real users is the recommended next step to confirm the committed impact indicators. Constructive feedback: coordinator communication arrived concentrated and on short notice — monthly coordination with formal milestones, bilingual RO+EN templates and a communication budget would help future pilots.
Municipality's Reflection	Proposed text: Hosting PrimAssist gave our administration its first practical experience with AI-assisted citizen information. The assistant answers strictly from our own documents, which builds trust, and it runs without storing personal data. The pilot also showed where we need to strengthen internal capacity (document curation, digital skills). We are interested in continuing on a commercial basis once an operational phase confirms the benefits for citizens and clerks.
Expert's Reflection	The expert considers the pilot a sound technical validation of grounded, GDPR-compliant conversational AI for local administration. The KPI framework and the multi-agent, anti-hallucination design are methodologically solid and transferable. For scaling, we recommend an operational phase with real users to measure the committed impact indicators and an independent security review before large-scale deployment.