From Amarillo to Carmel

A practical path to AI concierge services in city hall

A whitepaper by The Pennar LLC

- 21 May 2026*

Executive summary

In December 2024, the City of Amarillo, Texas (population 200,000+) deployed Emma, an AI-powered digital human, on its city website. Within ten months, Emma was answering 16,000 resident questions per month — half the volume the city's call center was previously handling by phone — and had effectively saved the city more than $1 million while deferring a $1.8 million planned investment in new call center staff. Resident satisfaction stands at 98%.

Amarillo's success is no longer an experiment. It is a precedent.

Carmel, Indiana — a city Mayor Sue Finkam has publicly committed to making "agile, data-driven, innovation-minded, and customer-focused" — is well positioned to be the next city to lead. In December 2025, Mayor Finkam stated that AI could turn a two-to-four week permit wait into a "near spontaneous" request, and signaled that the city is actively considering AI for public records and planning workflows.

This whitepaper is for city leaders, council members, and CIOs who want to understand what is genuinely possible today, what it costs, what governance it requires, and how a mid-sized city can move from idea to working pilot in 90 days. We focus specifically on the form of conversational AI that has shown the strongest results in municipal settings: the AI digital human, or "concierge" — a humanlike conversational interface that lives on a website, in a city hall lobby, or both.

The conclusion in brief: cities that move now — before the technology becomes a commodity sold by the Tier-1 enterprise vendors — will define how their residents experience local government for the next decade. The cost of leading is modest. The cost of waiting is higher than it appears.

---

1. The Amarillo precedent

Amarillo is not a particularly large or wealthy city. It is the largest city in the Texas panhandle, with a population just above 200,000 — comparable in scale to Indianapolis's suburbs combined, and smaller than the City of Indianapolis itself. Nearly a quarter of Amarillo's residents do not speak English at home, and Amarillo public schools serve students who collectively speak 62 languages and dialects.

In other words, Amarillo had a hard problem and a limited budget.

In August 2023, the Amarillo City Council unanimously approved a contract with Dell Technologies for $582,948.89 to design, build, and implement an AI digital human for the city. The contract was renewed in December 2024 at an annual operating cost of $553,052 for hosting and licensing. The technology stack:

- **Dell Technologies** — prime contractor, infrastructure, large language model integration
- **UneeQ** — the digital human avatar and conversational interface layer
- **Pryon** — knowledge fabric, designed to prevent the system from "hallucinating" answers it does not know

Emma was launched publicly on October 1, 2024, and upgraded to Emma 2.0 in August 2025 with a more natural conversational style and explicit "digital concierge" routing — directing residents to the correct city web pages rather than simply answering text questions.

The measurable results, ten months in:

- Approximately 16,000 of Amarillo's monthly 32,000 call center questions are now answered by Emma rather than by staff.
- The city has deferred a planned $1.8 million investment in new call center agents.
- Total reported savings exceed $1 million.
- Resident satisfaction with Emma's responses is 98%.
- Service is offered in English and Spanish today; five additional languages are being added based on input from refugee-serving nonprofits in the community.

The roadmap, as stated by Dell and the City of Amarillo:

Emma is now being extended beyond the website. The next phase places her in libraries, airports, and physical kiosks throughout the city. Her capabilities are being expanded to include emergency messaging, public library lookups, building and permitting information, and — eventually — a "24/7 insight engine" that asks residents qualitative questions rather than only answering them.

The trajectory is clear: from chatbot, to website avatar, to physical concierge. Amarillo is moving in that order. The next wave of cities can move differently.

2. What Amarillo got right — and what comes next

Amarillo's success rests on four decisions that any city replicating this work should understand explicitly.

First, they treated this as a citizen service problem, not an IT project. The work was led from the City Manager's office — Assistant City Manager and CIO Rich Gagnon, reporting directly to leadership — and approved by the full council with public input. It was not delegated to the IT department to figure out alone.

Second, they wrote the governance framework before they wrote the technology. Before Emma launched, Amarillo passed a Digital Dignity, Rights and Privacy Resolution, then later worked toward making it a binding ordinance. The city collaborated with a group called Digital Rights House, an advisory board of residents and legal experts, and more than 100 resident beta testers. The principle they adopted: *AI can enhance human dignity and must respect fundamental rights; privacy must be protected by design and by default.* This governance work is what allowed Emma to launch in a politically sensitive environment without backlash.

Third, they used a digital human, not a chatbot. This was a deliberate choice. Text chatbots had been available for years and were already deployed in cities like Denver, Boston, Midland, and Palo Alto. Amarillo's leadership concluded that for residents who don't speak English well, who lack digital literacy, or who have low trust in government, a humanlike face delivering information builds connection in a way that a text box does not. The data has supported this: an 89% user preference for digital humans over chatbots is consistent across enterprise deployments. Emma's 98% satisfaction score is an extension of that finding into local government.

Fourth, they kept humans in the loop. Emma is trained only on Amarillo's public website content, refreshed every 24 hours, and is not connected to internal city systems. Anything sensitive — anything that could harm a resident if answered incorrectly — is escalated to a human. The system's confidence is calibrated to say "I don't know, here is the number to call" rather than invent an answer.

What Amarillo did not do — and where the next wave will differ:

Amarillo built a website-first digital human, custom-engineered through a Tier-1 enterprise vendor at a six-figure first-year cost and a comparable annual recurring cost. That model worked for a city of 200,000 with federal pandemic-recovery funding to deploy. It is not the right model for every city.

The next wave of municipal AI concierge deployments will, we believe, differ in three respects:

1. **Physical-first, not web-first.** The city hall lobby is where residents who *most* need help actually show up. A kiosk-based concierge meets them where they are.

2. **Open architecture, not vendor lock-in.** A modular stack built on widely available components — open-source avatar rendering, configurable voice and language models, standards-based infrastructure — gives cities the ability to evolve their system without renegotiating with a single prime contractor every renewal cycle.
3. **Right-sized for the city.** A city of 106,000 should not pay a contract designed for an enterprise of 200,000 plus refugee multilingual support. The cost structure should match the deployment scope.

This is not a critique of Amarillo's choices. It is a forecast of what becomes possible once the precedent has been set.

---

3. Five questions every city leader should ask

City leaders evaluating conversational AI for municipal services are often handed a vendor demo before they have a framework for judging it. The following five questions, drawn from interviews with practitioners in Amarillo, Edmonton, Palo Alto, and other deploying cities, are the ones we believe matter most.

Question 1: What problem are we actually solving — and for whom?

Cities deploy conversational AI for very different reasons. Some are trying to extend service hours when 311 call centers are over capacity. Some are trying to reach non-English speakers or residents with low digital literacy. Some are trying to compress permit and records-request turnaround times. The right architecture for each goal is different.

Practical use cases that have demonstrated value in municipal settings include:

- *Citizen services* — handling the roughly 85% of inbound questions that are repetitive: property tax inquiries, recreation center schedules, trash pickup, hours of operation, where to apply for what.
- *Legislative support* — summarizing large document stacks for council members ahead of meetings, and assisting in drafting or updating bylaws and briefing notes.
- *Internal knowledge management* — letting staff query building codes, safety policies, or accessibility requirements in seconds instead of sifting through PDFs.
- *Administrative automation* — converting recorded meeting audio into structured minutes and agendas, a process that traditionally takes staff hours.

Carmel's specific stated interest in AI for permitting and public records aligns with categories one and three above. That is a tractable starting point.

Question 2: Chatbot, voice assistant, or digital human?

This is the most consequential technical choice and the one most commonly skipped. A text chatbot is the cheapest to deploy and the lowest-engagement format; it works for digitally fluent residents seeking quick answers. A voice assistant adds accessibility for residents who struggle with reading or typing. A digital human — a visible, expressive, conversational face — is the most engaging format and the one with the highest measured citizen satisfaction in municipal deployments. It is also the format that translates into a physical kiosk experience.

For a city deploying *only* on a website, a chatbot may be sufficient. For a city that wants service in the city hall lobby — where many residents who most need help actually come in person — a digital human on a kiosk is the right unit of design.

Question 3: Website-only, kiosk, or both?

Most early municipal deployments have been website-only because that is what the established vendors (Dell + UneeQ, Citibot, others) currently productize. Yet residents who lack home broadband, who have low digital literacy, or who are most likely to be intimidated by a website are also the residents most likely to physically visit city hall. A kiosk-based concierge in the lobby is not a luxury feature; it is the most equitable form factor. We expect every serious municipal deployment over the next 36 months to include a physical kiosk component.

Question 4: On-premise, cloud, or hybrid — and who owns the data?

The single most contested question in municipal AI procurement is data residency. Cities have legitimate concerns about resident data leaving their custody, particularly when the vendor's infrastructure is operated by a national enterprise. The practical framework used by experienced practitioners is the **green data / red data** distinction:

- *Green data* — publicly available information (city website content, posted ordinances, published schedules, public meeting recordings). This data can safely sit in cloud infrastructure.
- *Red data* — private resident information (birth dates, income data, addresses tied to identity, anything covered by FOIP, HIPAA, or local privacy ordinance). This data should never be exposed to the AI layer at all.

The architecture that follows from this distinction: the AI concierge is trained *only* on green data, and is structurally disconnected from any system that holds red data. Residents who need access to their own personal information are routed to a human or to an authenticated portal — not to the AI.

This is exactly how Amarillo structured Emma. It is the only approach that has survived political scrutiny in deploying cities to date.

Question 5: Who's accountable when it gives a wrong answer?

Every conversational AI system will occasionally answer a question incorrectly. The governance question is what happens when it does. Best practice from deploying cities:

- AI-generated content is reviewed by human staff before publication on city channels.
- The system is calibrated to say "I don't know" rather than invent, with explicit escalation paths to a human (a phone number, a specific department, an in-person desk).
- Daily or weekly logs of resident questions are reviewed by department heads, so the questions Emma cannot yet answer become the work tickets for content improvement.
- The city is transparent with residents that they are interacting with AI, and transparent with employees about how AI is being used and what it is and is not doing.

Together these five questions form a complete procurement framework. A vendor that cannot answer all five clearly is not yet ready to deploy in your city.

---

4. The broader landscape: who else is doing this

Amarillo is the most visible US municipal digital-human deployment, but it is not alone. The map below summarizes the state of play as of May 2026.

Cities that have deployed an AI digital human (avatar-based conversational AI):

- *Amarillo, TX* — Emma, web-based, expanding to kiosks. Production since October 2024.
- *Abu Dhabi (government)* — Fatima, holographic, bilingual English/Arabic. Showcase deployment by Capgemini and RAVATAR.

Cities with text or voice-based AI chatbots in production:

- *Palo Alto, CA* — CityAssist (Citibot platform), multilingual in 70+ languages, launched December 2025.
- *Denver, CO* — Sunny, expanded as 311 phone hours were cut in October 2025.
- *Boston, MA* — AI-native 311 platform on Creatio, deployed November 2025.
- *Midland, TX* — AskJacky chatbot, live since 2024.
- *Edmonton, AB* — Digital assistant on transit and waste service pages.
- *Windsor, ON* — Voice-enabled chatbot pilot for the Provincial Offences Administration office, approved December 2025.

Cities using AI for adjacent municipal workflows:

- *Pittsburgh, PA* — Real-time AI traffic data analysis to optimize signal timing.
- *Honolulu, HI* — AI permit prescreening; cut wait times from 6 months to 2–3 days.

- *Hamilton, ON* — AI building permit scanning, in partnership with the Bloomberg Center for Cities.
- *Quebec, Canada* — AI-driven predictive maintenance of municipal infrastructure.
- *Alberta, Canada* — AI-driven document analysis for building code queries.

Cities piloting or actively soliciting:

- *Austin, TX* — Citywide virtual assistant tested under "One City" initiative.
- *San Jose, CA* — GovAI Coalition (750+ government agencies), predictive AI in production.
- *Seattle, WA* — 39 AI pilots launched in the past year.
- *Oakland, CA* — "No-cost RFI" AI pilot program, November 2025.
- *Charlotte, NC* — Three-year $50M tech upgrade in phased AI rollout.

International leaders:

- *Seoul, South Korea* — Metaverse Seoul, an avatar-based public services platform for taxes, documents, and mayoral consultations.
- *Singapore* — Multiple digital human deployments across public-facing services.

The pattern is clear. Conversational AI for municipal services has moved from frontier experiment to mainstream procurement category in roughly 18 months. Carmel today is in approximately the position Amarillo was in 24 months ago: stated interest from the mayor, internal alignment forming, no production deployment yet, no other Indiana city ahead.

---

5. What it takes: the conditions for success

Across every deploying city, four conditions separate the projects that succeeded from those that stalled.

Leadership and political will. This work cannot be delegated to the IT department alone. It requires a mayor or city manager who has publicly committed to the outcome, council approval for funding tied to clear citizen-service goals, and a designated executive sponsor with authority to make tradeoffs. Amarillo's project succeeded because the assistant city manager owned it. Cities where the project lived only inside IT have generally not produced durable outcomes.

A governance framework adopted before the technology launches. Amarillo's Digital Dignity, Rights and Privacy work is the template. It does not need to be a full ordinance to start — a resolution with an annual review and a clear set of guiding principles is enough for v1. What matters is that the framework precedes the deployment, so when the inevitable questions arise (which they will, from council members, journalists, and engaged residents), the answers are already on record.

Clear green-data / red-data separation. As discussed in Section 3, this distinction must be made architecturally, not procedurally. A system that *could* access private resident data and merely promises not to is not the same as a system that *cannot*. Cities that get this right structurally never have to defend a privacy breach.

Human-in-the-loop review. Every system needs a feedback path where staff see what residents asked, what the AI answered, and what it got wrong. The deploying cities that have improved fastest are the ones where department heads review the previous day's conversation logs each morning. This becomes a 15-minute habit that drives a virtuous cycle of content improvement.

These conditions do not require additional staff or significant new budget. They require executive attention and a few weeks of governance work before procurement begins.

---

6. A 90-day pilot blueprint

What does it actually take to get from "we should look at this" to "we have a working AI concierge serving residents"? Based on the deploying cities' published timelines, the answer is roughly 90 days for a meaningful pilot, and six to nine months for a full production launch.

Days 1–30: Foundation

- Designate an executive sponsor (deputy mayor, city manager, or CIO with authority).
- Pass a Digital Dignity-style resolution at council, or commit to one within 90 days.
- Identify the three highest-volume citizen question categories (often: permits, property tax, parks and recreation hours/registration). These become v1 scope.
- Inventory and tag city website content as green data; explicitly exclude red data systems from architectural reach.
- Select a vendor partner with a working pilot offer, not a six-figure custom build proposal.

Days 31–60: Build and train

- Ingest green data into the system. Most modern stacks complete this in under two weeks.
- Configure the avatar — appearance, voice, personality consistent with the city's identity.
- Define the escalation paths: what questions go to which human department, by what mechanism.
- Train a small beta group of residents (10–25, drawn from senior centers, the chamber, immigrant-serving nonprofits) on the working system.
- Daily review of beta interactions; fix content gaps.

Days 61–90: Public soft launch

- Deploy the concierge on a single high-traffic city web page, with clear AI disclosure.
- If kiosk-based, place a single kiosk in the city hall lobby alongside existing service desks.
- Establish a feedback channel and review cadence.
- Publish a 90-day report to council with usage data, satisfaction scores, content gaps identified, and a roadmap for v2.

The goal of the 90-day pilot is not a polished citywide deployment. It is a working, defensible, measurable v1 that builds the council's confidence in the next decision: whether and how to scale.

---

7. What this means for Carmel

Carmel is, by several measures, the Indiana city best positioned to lead on municipal conversational AI:

- A mayor who has publicly stated AI is part of her vision for "agile, data-driven, innovation-minded" government.
- Specific use cases already identified at the executive level: permitting and public records.
- A population (106,000+, newly ranked #1 Best Place to Live in the U.S. by U.S. News for 2026–2027) that is well-educated, multilingual, and digitally fluent — a community that will adopt the technology rather than resist it.
- A governance culture that has historically been willing to invest in long-horizon improvements to quality of life (the roundabouts, the Arts & Design District, the Palladium).
- No other Indiana city ahead of Carmel on this trajectory.

What Carmel would *not* do, in our view, is replicate the Amarillo path one-for-one. The Amarillo model was right for Amarillo — a custom build through a Tier-1 vendor at a six-figure annual cost, scaled for a refugee population speaking 62 languages. Carmel's needs are different: faster permit-related self-service, smarter records-request handling, and a city hall lobby experience that matches the rest of the city's quality bar. A right-sized, kiosk-inclusive, open-architecture deployment is the appropriate Carmel-scale approach.

We believe Carmel can be the first US city to deploy a physical AI concierge kiosk in its city hall lobby — not as a publicity initiative, but as a permanent improvement to how residents experience the city. Amarillo is moving toward kiosks but has not yet placed one. The window for Carmel to lead is open today and will close within roughly 18 months as the Dell-and-UneeQ productized version of this becomes a standard procurement option.

---

8. About Pennar

The Pennar LLC is a Carmel, Indiana company building Aria, a real-time AI Digital Concierge designed for the lobbies, front desks, and public-facing service settings where modern organizations meet the people they serve. Aria is voice-first, multilingual across more than 20 languages, and engineered for sub-500ms conversational latency — the threshold at which a conversation feels human rather than computational.

Pennar offers two production-ready form factors today. **Aria for Kiosks** is a 3D digital human deployed at physical locations, currently piloting in the hospitality sector in Carmel including a concept showcase at the Palladium. **Aria for Mobile** is a native mobile application delivering low-latency voice-to-voice AI conversation, designed for hospitality and service-experience use cases on any device.

Pennar's current production stack is deliberately open and modular: a real-time 3D avatar layer built on Three.js with GLB-format avatars and the TalkingHead lip-sync framework, Google Gemini Live for conversational intelligence, and Azure TTS for voice generation. Pennar is an **NVIDIA Inception Program partner** and an **ElevenLabs Grants Program member**, and the platform's evolution roadmap includes migration to NVIDIA GPU-accelerated rendering, Unity HDRP for higher-fidelity avatars, and ElevenLabs voice agents — upgrades that preserve the modular architecture while delivering increasing realism over time. The principle is consistent: every component is replaceable, and no single vendor controls the pipeline. The company's tagline is *Humanizing Conversational AI.*

About the author

Ram Samala is the Founder and CEO of The Pennar LLC. Before founding Pennar in 2023, Ram spent sixteen years at **Eli Lilly and Company** as Associate Director, where he led the design and implementation of a scalable cloud data warehouse for pharmaceutical commercial analytics, established enterprise data governance policies, implemented patient data tokenization for security and privacy compliance, and led CCPA compliance and information security initiatives across regulated patient datasets. That work — operating at the intersection of analytics, privacy regulation, and enterprise-scale data security — is the foundation on which Pennar's approach to municipal AI is built. Cities deploying conversational AI face the same fundamental challenges as a regulated enterprise: protecting sensitive data, building trustworthy AI systems, and earning the confidence of the people the data is about.

Ram holds a **Master of Technology in Industrial Engineering** from Sri Venkateswara University and completed the **Applied Data Science Program at MIT Professional Education**. He is based in Carmel, Indiana, where he is building Pennar with the explicit conviction that the future of generative AI in physical experiences can be built in Hamilton County rather than imported from a coastal tech hub.

Contact

Ram Samala — Founder & CEO, The Pennar LLC
Email: ram@the-pennar.com
Phone: 317.531.2933
Web: the-pennar.com
LinkedIn: linkedin.com/in/ramsamala

For inquiries about pilot deployments, governance frameworks, or briefings for city council and staff, contact Ram directly at the email or phone above.

About the author

Ram Samala is the Founder and CEO of The Pennar LLC. Before founding Pennar in 2023, Ram spent sixteen years at **Eli Lilly and Company** as Associate Director, where he led the design and implementation of a scalable cloud data warehouse for pharmaceutical commercial analytics, established enterprise data governance policies, implemented patient data tokenization for security and privacy compliance, and led CCPA compliance and information security initiatives across regulated patient datasets. That work — operating at the intersection of analytics, privacy regulation, and enterprise-scale data security — is the foundation on which Pennar's approach to municipal AI is built. Cities deploying conversational AI face the same fundamental challenges as a regulated enterprise: protecting sensitive data, building trustworthy AI systems, and earning the confidence of the people the data is about.

Ram holds a **Master of Technology in Industrial Engineering** from Sri Venkateswara University and completed the **Applied Data Science Program at MIT Professional Education**. He is based in Carmel, Indiana, where he is building Pennar with the explicit conviction that the future of generative AI in physical experiences can be built in Hamilton County rather than imported from a coastal tech hub.

Contact

Ram Samala — Founder & CEO, The Pennar LLC
Email: ram@the-pennar.com
Phone: 317.531.2933
Web: the-pennar.com
LinkedIn: linkedin.com/in/ramsamala

For inquiries about pilot deployments, governance frameworks, or briefings for city council and staff, contact Ram directly at the email or phone above.

Appendix: References and further reading

Amarillo and the Emma deployment:

- "Meet Emma, Amarillo's AI assistant and 'digital human'" — Route Fifty, May 2024
- "As AI Assistant Evolves, Officials Hope It Furthers Trust" — GovTech, October 2025
- City of Amarillo Emma overview — amarillo.gov/about-emma
- Dell Technologies City of Amarillo customer story — dell.com

Governance and digital rights:

- City of Amarillo Digital Dignity, Rights and Privacy Resolution and ordinance materials
- Digital Rights House — digitalrightshouse.org

Industry context:

- "Launch of UneeQ 2.0" — UneeQ press release, January 2025
- "Dell Technologies debuts On-Prem Digital Human Assistant" — IT Business Today, February 2026
- "Digital Cities 2025" — GovTech annual ranking

Indiana ecosystem:

- Accelerate Indiana Municipalities — aimindiana.org
- TechPoint Indiana AI Imperative — techpoint.org/indiana-ai-economy
- IU Public Policy Institute — policyinstitute.iu.edu

---

*This whitepaper is current as of May 2026. The municipal conversational AI landscape is evolving rapidly; an updated edition will be issued quarterly at the-pennar.com