

Community technology advisory commission

The community technology advisory commission is meeting in person at St. Louis Park City Hall, 5005 Minnetonka Blvd., in the Westwood Room on the third floor. Members of the public are welcome to attend in person.

Visit www.stlouispark.org/government/boards-commissions to view the agenda and reports.

Agenda

1. Call to order – roll call (6 p.m.)
2. Approval of minutes – July 15, 2026
3. Presentations - none
4. Business
 - a. Review submitted signal items (10 minutes)
 - i. Next steps
 - ii. September signal assignments
 - b. Review AI whitepaper (20 minutes)
 - c. Prep for Sept. 14 annual meeting with city council (10 minutes)
5. Communications and announcements (5 minutes)
 - a. Staff
 - b. Commission members
6. Adjournment

Future meeting/event dates:

- Sept. 14, 6 p.m., CTAC annual meeting with council
- Sept. 16, 6 p.m., regular meeting
- Oct. 21, 6 p.m., regular meeting
- Nov. 18, 6 p.m., regular meeting
- Dec. 16, 6 p.m., regular meeting

CTAC Purpose Statement: *The Community Technology Advisory Commission advises the city council and other city bodies on how technology services can improve resident services and enhance quality of life. By researching emerging technologies, assessing their potential impact, and recommending solutions grounded in clearly defined problems, the commission helps ensure that technology decisions support the city's long-term strategic goals.*

Community technology advisory commission

Minutes

Members present: Elsa Anderson, Reid Anderson, Edward Bulliner, Rudyard Dyer, Lindsay Keogh, Konnor Slaats, Kayla Stautz
Members absent: Ben Straus, Tom Marble
Staff liaison: Jacque Smith, communications and technology director

1. Call to order – roll call

The meeting was called to order at 6:04 p.m.

Members introduced themselves. Chair Slaats noted he has served on the commission since 2021 and works in product operations. Commission Member Stautz said she has been on the commission for close to a year and works as general manager of the consumer business for a technology company. Commission Member Reid Anderson noted he works as a software engineer and has been on the commission for just under a year. Vice Chair Dyer noted he has lived in St. Louis Park for over a decade, has served on the commission since 2022, and works in due diligence at an alternative investment management company.

Commission Member Bulliner introduced himself as new to the commission and to St. Louis Park, having moved here approximately a year and a half prior. He noted he works in geographic information systems software, has a background in statistics and data science, and previously worked in government. Commission Member Keogh introduced herself as new to the commission and noted she works as a product manager, focused on search and digital experiences. Commission Member Elsa Anderson introduced herself as a 16-year-old high school junior.

2. Approval of minutes – April 15, 2026

It was moved by Commissioner Slaats, seconded by Commissioner Anderson, to approve the April 15, 2026, minutes. The motion passed.

3. Presentations – none

4. Business

Review and approve purpose statement

Ms. Smith displayed the current purpose statement alongside a proposed draft for the commission's review. The draft statement introduced language around city-provided public-facing technology.

Chair Slaats expressed a preference for broader language, suggesting that a focus solely on public-facing technology could exclude relevant matters such as technology by outside entities which the city may have to respond to, for example, data centers. He proposed language referencing how technology can improve resident services and enhance quality of life.

Commission Member Keogh agreed that framing the purpose around residents and quality of life would appropriately capture the commission's advisory scope, noting that the draft reflected

the city's strategic direction and the increasingly interwoven role of technology across city functions.

Vice Chair Dyer suggested the language could also acknowledge technology that could negatively affect quality of life. He proposed morphing the original and draft statements to reference how city-provided technology can affect residents.

Commission Member Bulliner noted the importance of a holistic perspective and acknowledged that back-end systems often affect public-facing outcomes.

Commission Member Reid Anderson suggested simplified language referencing city-provided technology without the public-facing qualifier.

After discussion the commission agreed to modify the draft to reference how technology can improve resident services and enhance quality of life, removing the public-facing qualifier while retaining the general structure of the draft statement. The revised statement was approved, as follows:

“The Community Technology Advisory Commission advises the city council and other city bodies on how technology services can improve resident services and enhance quality of life. By researching emerging technologies, assessing their potential impact, and recommending solutions grounded in clearly defined problems, the commission helps ensure that technology decisions support the city’s long-term strategic goals.”

Work plan items: background, workgroups, progress

Chair Slaats provided background on the commission's work over the past year. He described the commission's role as an advisory body to city council and noted that a meeting with council in April 2026 helped clarify expectations, with a follow-up council meeting scheduled for September 2026. He explained that the commission had split into two workgroups: one focused on artificial intelligence and one focused on evaluating the city's digital presence, including the city website and the My St. Louis Park app.

Chair Slaats then presented a proposed framework for how the commission might structure its future work. He described a multi-stage process involving three phases: signals, exploration and active work.

In the signals phase, commission members would bring forward brief summaries of technology trends, peer city activities, policy developments or resident-relevant issues. These would be submitted to Ms. Smith in advance of meetings for distribution with meeting materials and discussed at the start of each meeting, with the commission deciding collectively whether to watch, explore further or set aside each signal. Chair Slaats suggested that two members might bring signals per meeting but that anyone could contribute.

Commission Member Bulliner expressed support for receiving advance materials so members could formulate informed opinions prior to discussion.

Commission Member Reid Anderson noted the framework helpfully defines what substantive commission work can look like for members who wish to contribute more.

Commission Member Stautz noted that technology evolves quickly and that the framework would allow the commission to be more current and relevant. She also confirmed her understanding that all three stages would operate simultaneously, with work moving through the pipeline.

Commission Member Keogh noted the framework allows work to be taken in manageable increments and raised the question of how resident voices might be incorporated into the process, acknowledging that was likely a future consideration.

In the exploration phase, one or two members would draft a one-page framing of a potential topic, working toward a clear problem statement. Chair Slaats offered an example question related to the digital channels workgroup: what principles should the city use to determine if a resident-facing digital service provides sufficient public value.

Vice Chair Dyer indicated he was open to reframing or restarting the digital channels workgroup given the extended period of limited progress, much of which was due to waiting for the city's website redesign. He noted the commission had previously reviewed website analytics and encountered questions about the overlap and distinction between the city website and the city's app. Vice Chair Dyer suggested that a one-page exploration document could be a useful way to reframe that work with new members.

In the active work phase, a lead commission member would outline the project and identify key decisions, with two or three contributors drafting sections. The full commission would then critique, debate and approve a final version for submission to council.

Commission Member Stautz asked whether it would be worth attempting the signals component before the next meeting as a trial. Chair Slaats agreed it would. The commission decided to ask members to submit signal summaries to Ms. Smith by Aug. 12, 2026, ahead of the Aug. 19, 2026 meeting.

Ms. Smith noted the timing is well suited to the upcoming September council meeting and suggested that the next meeting agenda include discussion of how to present the commission's new direction to council.

Chair Slaats noted the draft AI paper compiled by Commission Member Stautz would be distributed with the next meeting packet for review and discussion, along with the framework document he had emailed the prior evening.

Regarding the AI workgroup, Vice Chair Dyer noted progress had been set back when a former member departed. Chair Slaats indicated the workgroup was approaching completion and that distributing the draft ahead of the next meeting would allow the full commission to review and provide feedback.

Ms. Smith confirmed she would distribute all relevant materials to members in advance of the Aug. 19, 2026 meeting. She also noted she would resend calendar invites for upcoming meetings to all members.

Ms. Smith provided an update on the city's website redesign, noting a soft rollout is scheduled for Aug. 17, 2026. She described improvements including enhanced accessibility, improved search functionality and better mobile responsiveness.

5. Communications and announcements
There were no communications or announcements

6. Adjournment
It was moved by Chair Slaats, seconded by Commission Member Dyer, to adjourn. The motion passed. The meeting adjourned at 6:47 p.m.

CTAC signals

A technology trend I have observed in St. Louis Park is the rapid growth of AI use, specifically in schools. As a student, I spend most of my time surrounded by other students and in the past year irresponsible use of AI emerged as a big discussion in schools. Many students utilized the increasing availability and quality of AI for school work and as a result many assignments had to be changed and even some curriculum. Many assignments had to be completed on paper and only in class which I, along with many other students found reduced the quality of the work. I believe that this will become an even bigger issue in the upcoming school year as AI continues to improve and the school board and city council will have to find a solution to ensure students' quality of education remains strong. CTAC could be useful in providing resources for teachers on how to avoid the misuse of AI while keeping the quality of education strong.

The growing importance of AI- and data-driven technology in city technology means that responsible data stewardship is becoming core to good, transparent governance. City government plays a critical role in re-establishing trust in government overall. If residents do not trust how their data is handled by the city, it leads to more distrust in government and overall dissatisfaction.

Common solutions used by cities, especially in the surveillance/policing space, gather and share more data than ever with outside parties. Once data is collected by these sources, how the data is used & accessed can be a black hole or pose a risk to the safety of the community, depending on who gets access to the data.

However, some cities take a community-centered, transparent approach to solutions that gather & process resident data, which both mitigates risk of data misuse, unwarranted surveillance, and improves community trust.

Trends

- Privacy/security concerns
 - Flock: West St Paul, Burnsville, Duluth, Isanti county all have ended or chose to not enter into an agreement with Flock, citing privacy concerns & spuriousness about Flock's behavior as a vendor partner
 - Unwarranted/rampant usage of data by ICE during the federal invasion to intimidate observers and community members - sourced from tools such as Accurint
 - Backlash against Skydio drones in minneapolis
- High-profile problems lead to distrust of AI-driven systems, wasting community resources and underscoring the need to build trust through other mechanisms
 - 1.4 billion bad notifications a month lead to unnecessary traffic stops or effort
 - Flock: LA saw massive accuracy issues, leading them to end their pilot
 - Plymouth man pulled over by four cops due to a plate read error
 - Tennessee woman wrongfully jailed for six months based on AI facial recognition

Potential policy question: How should St Louis Park evolve its data usage & sourcing practices to achieve its strategic vision, particularly around good governance & creating a welcoming, safe community?

CTAC could:

- Propose a framework for gathering resident input to study the topic further
- Investigate strategies, principles, and risks for current & future resident-facing technology and data collection

Sources:

<https://www.govtech.com/transportation/for-cities-collecting-visual-data-brings-privacy-concerns>

Continued - August 2026 Signal-Lindsay Keogh

<https://www.brookings.edu/articles/how-citizens-and-local-governments-can-advance-data-driven-policy-making-together/>

https://mcity.umich.edu/wp-content/uploads/2023/03/Privacy-Frameworks-for-Smart-Cities_White-Paper_2023.pdf

August 2026 Signal

<https://www.lmc.org/news-publications/magazine/july-august-2026/letter-of-the-law-july-2026/>

What's happening?

The League of Minnesota Cities' July–August 2026 guidance says information entered into AI systems, AI-generated outputs and associated system data may qualify as government data under the Minnesota Government Data Practices Act.

It warns that putting nonpublic information into third-party AI tools could constitute a data breach and that prompts or outputs may need to be retained, searched and disclosed.

Why might this matter to council?

SLP already has a GenAI use policy that prohibits use of medium-to-high risk data in GenAI tools ([March 18 packet](#)). SLP also already urges compliance with the Minnesota Government Data Practices Act, which the League article is referencing.

Is the League guidance going further by saying all inputs/outputs may be government data and need to be tracked? Is that something SLP and other cities are tracking today, or are they just tracking final outputs?

What policy or strategic question could this eventually create for Council?

Does the city's current approach account for prompts, input histories, intermediate outputs, and other AI system data, or primarily final documents?

If the city's current policy covers these, this signal may not require further action.

If a gap exists, Council could eventually consider whether the city needs clearer standards for retaining, retrieving, and classifying AI-related records. Any policy change would likely be narrow.

Why might CTAC be useful?

CTAC could help if staff identifies an unresolved policy question. Potential work could include comparing how peer cities treat prompts, intermediate outputs, and AI metadata under public-records requirements.

Otherwise, the League guidance could validate St. Louis Park's existing approach rather than justify additional CTAC work.

Future of AI for Cities

A Policy White Paper

Community Technology Advisory Commission | City of St. Louis Park, Minnesota

2026

Executive Summary

Artificial intelligence is no longer an emerging technology. Across the United States and internationally, cities of every size are deploying AI tools to detect potholes, process permit

applications, answer resident questions, translate public meetings, and manage aging infrastructure. The question for St. Louis Park is not whether AI will affect city government, but how the city will govern, evaluate, and integrate it responsibly.

This white paper, prepared by the Community Technology Advisory Commission, explores how cities can adapt their policies, processes, and governance frameworks for an AI-enabled future. It draws on examples from peer cities, national guidance from the National League of Cities (NLC), and the city's own existing technology posture.

The paper is organized around five topic areas: (1) the cross-cutting risks and governance principles that apply to all AI use; (2) internal city operations, including back-office tools and workflow acceleration; (3) public-facing services such as resident chatbots and meeting transparency tools; (4) public works and infrastructure operations; (5) planning, zoning, and permitting; and a forward-looking look at agentic AI and what may come next.

Across all areas, the Commission reaches a common conclusion: the value of AI depends on the quality of underlying data, the clarity of governance structures, the preservation of human accountability, and a sustained commitment to equity. St. Louis Park is better positioned than many peer cities because it has already adopted internal generative AI guidelines and has existing resident-service channels like MySLP. The next step is to extend that framework thoughtfully as new use cases emerge.

Section 1: Risks, Governance, and Guiding Principles

Before evaluating individual AI applications, the Commission believes it is important to establish a shared governance vocabulary and identify the risks that apply across all municipal AI use cases. This section frames the core issues that should inform every subsequent decision the city makes about AI.

Cross-Cutting Risks

AI Slop and Hallucination. Generative AI systems can produce confident, fluent, and entirely incorrect answers. A prominent municipal example: in 2024, New York City's MyCity chatbot advised business owners that they could legally reject Section 8 housing vouchers and that serving cheese bitten by a rat was permissible. Deploying AI tools that interact directly with the public without rigorous accuracy review creates real reputational and legal exposure.

Resident Data and Privacy. Public AI models may retain user inputs for training. Staff using unapproved tools can inadvertently expose resident financial records, health information, or other protected data. In early 2026, the acting director of the U.S. Cybersecurity and Infrastructure Security Agency reportedly triggered internal security warnings after uploading sensitive government documents into a public ChatGPT instance. Third-party contractors may also process city data through unauthorized AI systems without the city's knowledge.

Surveillance Creep. Camera networks, mounted sensors, and detection tools deployed for narrow purposes such as pothole detection can be tempting to repurpose for enforcement, public safety monitoring, or behavioral surveillance. Governance frameworks should establish clear use-case boundaries and require council review before any expansion.

Equity in Service Delivery. AI tools trained on complaint data may prioritize areas where residents are more likely to report problems, systematically underserving neighborhoods with lower reporting rates, language barriers, or lower institutional trust. Cities must actively monitor whether AI improves or exacerbates service equity.

Unpredictable Costs. Operational costs for AI can fluctuate by orders of magnitude based on small changes in how systems are prompted or how often they self-correct. Without hard spending limits and performance boundaries, AI tools can quietly become fiscal liabilities.

Vendor Dependence and Data Lock-In. Many AI tools are embedded in proprietary platforms. Cities should evaluate whether they can export their data, audit model performance, switch vendors, and integrate tools with core systems before signing contracts.

Records and Legal Obligations. Images, videos, work orders, and resident reports may become government records. Under the Minnesota Government Data Practices Act, data

created with AI assistance may need to be retained and produced in response to public data requests.

A Governance Framework for Council Questions

The National Institute of Standards and Technology (NIST) AI Risk Management Framework provides a useful foundation. The Commission has distilled its core functions into four questions that council members and staff can apply to any proposed AI deployment:

NIST Function	Council Question	What to Look For
Govern	Who is ultimately accountable if this fails and responsible for associated costs?	A specific human owner or position, not just an IT vendor.
Map	What specific problem does this solve, what data does it touch, and how is it paid for?	A clear use case with strict boundaries on what resident data is involved. Control mechanisms to prevent runaway costs.
Measure	How will we prove this system is unbiased, accurate, and cost-constrained?	Pre-deployment testing plan, ongoing performance audits, and consistent cost measurement.
Manage	What is the human fallback plan when the AI makes a mistake or is unavailable?	A documented mitigation strategy and a clear off-switch if things go wrong.

These four questions do not resolve every governance challenge, but they ensure that each AI deployment has a named human owner, a defined problem scope, a measurement plan, and a documented fallback. The Commission recommends that St. Louis Park require documented answers to these questions before approving any new AI system.

Principles for Responsible AI Use

Drawing on NLC guidance and peer city experience, the Commission recommends the following principles for all AI activity in St. Louis Park:

- AI should support human decisions, not replace them. Staff, commissioners, and council should retain accountability for all interpretations, recommendations, and approvals.
- AI should be attached to defined operational problems, not deployed speculatively.
- All AI tools should go through IT and legal review before deployment, including free and pilot products.
- AI-generated content in public-facing channels should be disclosed and cite authoritative sources.
- Performance should be measured against equity metrics, not just efficiency metrics.

- The city should maintain a public inventory of AI systems in use, consistent with best practices from cities like San Jose and Lebanon, NH.

St. Louis Park's existing internal generative AI guidelines already reflect many of these principles. The task ahead is extending that framework to cover operational, perceptive, and predictive AI applications as they emerge.

Section 2: Internal City Operations

City employees across departments increasingly interact with AI tools in their daily work, whether the city has sanctioned those tools or not. A proactive strategy for internal AI use reduces risk, improves consistency, and creates opportunities to demonstrate AI's value to staff and elected officials before expanding into higher-stakes resident-facing applications.

The Current Landscape

Cities are deploying AI for internal operations in three broad categories: productivity tools for individual staff, workflow acceleration across departments, and data analysis and decision support.

San Francisco has led among peer cities, adopting Microsoft Copilot for writing city reports, analyzing data, and summarizing documents across departments. San Jose issued an RFP in 2025 for a city-wide generative AI platform for employee use. Seattle and Boston have each published internal generative AI policies governing staff use of tools like ChatGPT, Gemini, and Copilot.

According to a 2025 global study of 250 cities, 56 percent are actively piloting or using AI to upgrade government operations and services, and 83 percent plan to do so within three years. The trend is clear: cities that establish governance frameworks now will be better positioned to adopt tools responsibly as they mature.

St. Louis Park has already taken a meaningful step by publishing internal generative AI guidelines. The next phase is building on that foundation by clarifying which tools are approved, how staff are trained, and how the city will evaluate and expand internal AI use over time.

Key Internal Use Cases

Staff Productivity and Writing Assistance

AI copilot tools integrated with productivity suites (Microsoft 365, Google Workspace) can help staff draft reports, summarize meeting notes, respond to routine correspondence, and search internal knowledge bases. Early research suggests these tools are especially helpful for newer or less-experienced employees. The risk is over-reliance on AI-generated content that has not been fact-checked or verified against authoritative sources.

Peer example — Chattanooga, TN: The city created a “Prompt Library” of pre-written prompts to guide generative AI tools for specific staff tasks, standardizing outputs and reducing errors across departments.

Grant Writing and Administrative Support

AI can assist staff with grant research, identifying eligible programs, drafting application narratives, and checking compliance with grant requirements. This has particular value for smaller cities where grant-writing capacity is limited.

Permitting and Workflow Acceleration

Several cities have integrated AI into permitting workflows to reduce review cycles, catch incomplete applications early, and improve consistency across staff reviewers.

Peer example — Hamilton, ON: The city used AI to scan first-stage building permit applications for compliance with city rules, building codes, and zoning requirements, achieving a 60 percent decrease in permit processing time and a 70 percent decrease in overall review time from application receipt to conditional approval.

Peer example — Los Angeles Planning: The Planning Department uses AI to analyze land use regulations and zoning data to proactively identify issues before formal permit application review begins, reducing back-and-forth between applicants and staff.

Peer example — Honolulu: CivCheck, an AI-guided permit preparation tool, helps applicants check for code compliance and required documents before submission, reducing review time by catching problems upstream.

Policy Implications for St. Louis Park

For internal operations, the Commission recommends the following:

- Extend the existing generative AI guidelines to cover productivity tools, copilots, and any AI embedded in core software platforms such as the city's permitting or work-order systems.
- Establish an approved tools list. Following Seattle's model, require that all AI software, including free and pilot products, go through IT and legal review before staff use.
- Invest in staff training. Partner with the League of Minnesota Cities, MNIT, or local educational institutions to develop AI literacy programs for staff at all levels.
- Create a feedback channel for staff to surface AI use cases, identify problems, and share what is working across departments.
- Set spending controls. Any AI tool with usage-based pricing should have a hard spending cap and monitoring process in place from day one.

Section 3: Public-Facing AI and Resident Engagement

Public-facing AI applications carry higher reputational and governance stakes than internal tools because errors affect residents directly and can undermine trust in city government. At the same time, well-designed resident-facing tools can meaningfully improve access to services, reduce barriers for non-English speakers, and extend the city’s reach outside business hours.

Chatbots and Self-Service Assistance

AI-powered chat interfaces are increasingly common in city government. When well-governed, they can handle routine informational inquiries, help residents navigate city services, submit service requests, and identify the right staff contact for complex issues. When poorly governed, they produce incorrect guidance with high confidence, as New York City’s MyCity chatbot demonstrated in 2024.

The key design principles for safe resident-facing chatbots are: ground responses in authoritative city documents and ordinances; include clear disclaimers; provide a path to human review for any consequential decision; and disclose that the user is interacting with an AI.

Ann Arbor, MI	"Ask Ann" provides 24/7 chat assistance in 71 languages, helping residents find city services, submit requests, and contact staff. Unresolved questions are escalated to staff during business hours.
Dearborn, MI	AI translation and virtual chatbot tools serve a community where more than half of 110,000 residents speak a language other than English at home.
Atlanta, GA	ATL311 chatbot handles routine service requests and common questions, reducing call-center volume and improving after-hours responsiveness.
Midland, TX	"Ask Jacky" is integrated with SeeClickFix to allow residents to report issues and track status via an AI-powered chat interface.
Kelowna, BC	Building permit chatbots had over 5,000 weekend conversations and 6,500 after-hours conversations in the first six months of 2024, reducing routine staff inquiries.

Meeting Transparency and Accessibility

AI tools for meeting accessibility address two related challenges: making government proceedings understandable to residents who cannot attend in person, and serving residents who do not read or speak English fluently.

Meeting agenda explainers: Saratoga, CA deployed an AI tool called Hamlet that summarizes council agendas, supporting documents, and meeting recordings in plain language. Government Technology reports that several cities are using similar tools to improve transparency and resident engagement.

AI translation and captioning: Sunnyvale, CA provides real-time translation and captioning at city council meetings. Smart Cities Dive reports growing use of AI captioning tools like Wordly for public hearings to serve non-English speaking residents. The NLC notes that Dearborn's AI translation services have significantly improved information access for its large Arabic-speaking population.

Open data and civic intelligence: Washington, D.C. piloted DC Compass, a generative AI assistant integrated with the city's open data portal, allowing residents to ask questions about datasets, generate visualizations, and understand neighborhood-level statistics without specialized technical knowledge.

Equity and Participation

AI-assisted public engagement tools have genuine potential to improve equity by reaching residents who cannot attend meetings, do not speak English fluently, or are unfamiliar with government processes. But they can also replicate or amplify existing inequities if deployment is not thoughtful.

Digital access remains a prerequisite for many of these tools. Residents without smartphones, reliable internet, or digital literacy may be excluded from AI-mediated engagement channels. The city should monitor whether AI tools improve participation among historically underrepresented groups or consolidate engagement among those already most engaged.

Policy Implications for St. Louis Park

- Any resident-facing chatbot should be grounded in authoritative city documents, updated regularly, and subject to accuracy review before launch and on an ongoing basis.
- AI-generated responses in public channels must disclose AI involvement and provide a clear path to human review.
- Meeting accessibility tools should be evaluated for translation quality across the languages most spoken by St. Louis Park residents.
- The city should monitor whether public-facing AI tools improve participation and access for underserved residents, not only convenient metrics like total interactions.
- All resident-facing AI applications should be included in a public AI inventory.

Section 4: Public Works and Infrastructure Operations

Public works may be the most operationally mature area for municipal AI because the work is data-rich, repetitive, and highly visible to residents. Streets, sidewalks, snow removal, sewer systems, waste collection, traffic signals, and resident-reported issues all generate information that can be used to detect problems earlier, prioritize work more effectively, and improve service delivery consistency.

The strongest AI applications in public works are not AI replacing staff. They are tools that help staff convert messy, high-volume real-world inputs—photos, sensor readings, service requests, inspection notes, and asset data—into structured information that supports better decisions. The city still needs public works professionals to validate conditions, authorize work, understand field constraints, and communicate tradeoffs.

Examples from Peer Cities

San José, CA	Vehicle-mounted cameras identify potholes (97% accuracy) and trash/debris (88% accuracy) in real time. The city is expanding the pilot to bike lanes and sidewalks, and also captures parking violations and lived-in vehicles—raising surveillance creep concerns.
Memphis, TN	AI analyzes video from city vehicles to identify potholes and property blight, combining footage with existing data sources to predict areas at risk of urban decay before residents report them.
Eagan, MN	A second-ring Twin Cities suburb with 244 centerline miles of streets uses LiDAR and AI for pavement condition scoring, producing dashboard-ready condition maps and data downloads for capital planning.
Cranberry Twp., PA	A pavement sensor and forecasting system providing localized 72-hour pavement temperature forecasts saved \$160,000 in winter maintenance costs by reducing unnecessary road salting.
Montreal, Canada	AI analyzes dump truck images to verify snow levels and prevent contractor fraud without capturing license plates or other identifying information.
Seattle, WA	Participates in Google's Project Green Light, using AI and Google Maps data to identify inefficient signal timings. Early results show potential to reduce stops by 30% and CO2 emissions by 10%.

Potential Applications for St. Louis Park

For St. Louis Park, the most relevant public works applications would likely be incremental and workflow-driven rather than large “smart city” transformations. The city’s existing MySLP/Accela platform, with 17,288 subscribers and over 6,000 service requests submitted since 2014, provides a foundation that AI can build on rather than replace.

- Converting resident service requests, photos, and notes into standardized categories, locations, urgency levels, and routing queues.
- Using camera or inspection data to identify potholes, pavement markings, sidewalk obstructions, illegal dumping, waste overflow, or snow-related conditions.
- Supporting pavement, sewer, and asset-management planning by combining inspection data, condition ratings, historical repairs, and GIS.
- Helping field staff retrieve internal knowledge, prior work orders, maintenance history, and standard operating procedures while in the field.
- Summarizing public works trends for staff and council, such as recurring hotspots, seasonal patterns, response times, and service-level performance.

Risks and Watch-Outs

Public works AI can appear operationally low-risk but still raises serious governance questions. Surveillance creep is a particular concern: the same camera network used for pothole detection can be repurposed for monitoring people, encampments, or behavior. Equity in service prioritization is another: AI trained on complaint data may systematically underserve neighborhoods with lower reporting rates. And AI outputs should always be treated as signals for human review, not final determinations—false negatives in public safety contexts can miss genuine hazards.

Policy Implications for St. Louis Park

- Start with defined operational problems. Attach AI to specific service goals—faster pothole identification, improved snow response, better asset planning—rather than deploying speculatively.
- Prioritize integration before automation. If resident reports, work-order systems, and GIS are disconnected, AI will add complexity rather than reduce it.
- Require human review for any action. AI can flag, sort, summarize, or recommend; it should not independently authorize repairs, enforcement, or billing decisions.
- Draw clear surveillance boundaries. Establish policy on what camera networks can and cannot be used for, and require council approval before expanding use cases.
- Measure service equity. Evaluate AI tools by neighborhood-level response times, complaint resolution rates, missed issues, and whether benefits are distributed fairly.

Section 5: Planning, Zoning, and Permitting

Planning, zoning, and permitting are among the most promising municipal AI use cases because the work is complex, document-heavy, rules-based, and highly consequential. Residents, builders, developers, board members, and elected officials must navigate zoning ordinances, building codes, application requirements, site plans, variances, setbacks, design standards, and public hearing processes. These systems are difficult for non-experts to understand, and delays can affect housing production, economic development, and resident trust.

AI can help by searching large bodies of rules, identifying missing documents, summarizing applications, comparing submissions against requirements, and providing clearer guidance to applicants and staff. The strongest use case is not having AI approve development; it is using AI to reduce avoidable back-and-forth, improve completeness at submission, and help staff focus on judgment-heavy issues.

Examples from Peer Cities

Hamilton, ON	AI-assisted site plan review for non-profit development projects produced a 60% decrease in review days and 70% decrease in overall processing time.
Kelowna, BC	Building permit chatbots launched in 2023 guide homeowners and builders through permit applications, reducing routine staff questions and extending service to evenings and weekends.
Austin, TX	Beta-testing an AI Pre-Check tool for residential zoning review, allowing applicants to identify zoning compliance issues before formal submission.
Honolulu, HI	CivCheck helps applicants verify code compliance, permitting requirements, and required documents before submission, reducing review cycles.
Sacramento, CA	AI roadmap identifies electronic plan check as a near-term priority to automate code compliance and design validation.
Los Angeles, CA	Planning Department uses AI to analyze land use and zoning data to proactively identify issues during the permit review process.

Potential Applications for St. Louis Park

The most relevant near-term applications fall into four categories:

1. Applicant self-service. AI can help residents, small businesses, and property owners understand basic rules such as setbacks, fences, accessory dwelling units, signage, and parking—and help them determine whether a permit is likely needed—outside of business hours. Tools must cite authoritative ordinances and provide links to current sources.

2. Submission completeness checks. AI can review application packages for missing documents, incomplete fields, conflicting information, or obvious inconsistencies before staff begins formal review.

3. Staff copilots. AI can help staff summarize application materials, prior permits, property history, relevant ordinances, variances, and public comments, reducing administrative time while preserving staff judgment.

4. Policy and scenario analysis. AI can help staff, boards, and council summarize how peer cities handle zoning questions, identify common ordinance language, or model high-level implications of proposed zoning changes.

Risks and Watch-Outs

Planning and zoning AI carries higher governance risk than public works triage because land-use decisions affect property rights, housing, and neighborhood change. Due process and transparency require that residents and applicants understand how AI was used in any decision affecting them. Incorrect legal interpretation by a chatbot or copilot can cause residents or applicants to spend money, delay projects, or misunderstand their rights. And tools trained on historical permitting data may perpetuate patterns of unequal access or enforcement.

A specific technical risk worth noting: applicants could embed hidden or misleading text in PDFs or plan documents to attempt to manipulate AI summaries or reviews. This “prompt injection” risk is a real procurement and security consideration for any document-review AI tool.

Policy Implications for St. Louis Park

- AI should not make binding land-use decisions. Staff, boards, commissions, and council remain accountable for interpretations, recommendations, and approvals.
- AI-generated guidance should cite current ordinances, city webpages, application forms, and relevant policies. Residents must be able to verify the source of any guidance they receive.
- Residents and applicants should always have a path to human review. AI should make the process easier, not create a barrier that residents cannot understand or challenge.
- The city should preserve records of AI use in planning and permitting decisions consistent with public records requirements.
- Procurement should require vendors to explain what data the tool uses, how it interprets city rules, how often it is updated, and how the city can test for errors.
- Performance metrics should include application completeness, number of correction cycles, total review time, applicant satisfaction, and equity of access across applicant types.

Section 6: Looking Ahead — Agentic AI and the Evolving Landscape

The AI use cases described in this white paper are largely assistive: AI that flags, summarizes, translates, or classifies information for human review. The next wave of AI technology—already emerging in enterprise and government contexts—is agentic AI, in which multiple AI agents work together to automate and coordinate multi-step tasks with minimal human intervention at each step.

What Agentic AI Means for Cities

The National League of Cities describes agentic AI this way: “Agentic AI can coordinate the actions of several business processes, from checking applications for errors or missing information to validating plans and generating compliance reports and then notifying applicants of approvals and collecting applicable fees.”

In practical terms, agentic AI could eventually handle workflows that today require multiple staff handoffs: receiving a permit application, checking it for completeness, routing it to the appropriate reviewer, generating a compliance checklist, sending the applicant a status update, and processing payment—all without human intervention at each step.

This is not science fiction. Agentic AI capabilities are being built into enterprise platforms today, and cities that have strong data foundations and well-defined workflows will be first in line to benefit. But the governance implications are significantly greater than for assistive AI, because agentic systems can take consequential actions on behalf of the city without a human reviewing each one.

Implications for St. Louis Park

The Commission does not recommend that St. Louis Park deploy agentic AI in the near term. However, the city should be aware of this trajectory and begin laying the governance and data foundations that will make responsible agentic adoption possible when the time comes.

Specifically, this means:

- Establishing clear workflow documentation so that any AI system—agentic or otherwise—can be configured against defined, auditable processes rather than informal practices.
- Ensuring that AI governance frameworks explicitly address autonomous action: under what circumstances can an AI system take an action without human review, and who bears accountability when it errs?

- Maintaining a human-in-the-loop requirement for all decisions that affect resident rights, benefits, or obligations until the city has sufficient experience, audit capability, and public trust to consider exceptions.
- Watching peer cities that are piloting agentic workflows in lower-stakes contexts—scheduling, document routing, status notifications—to learn from their experience before adopting similar tools.

External Resources and Learning Networks

St. Louis Park does not need to navigate this landscape alone. The following networks and resources provide valuable guidance, peer learning, and policy models:

National League of Cities (NLC). NLC’s AI in Cities Report and Toolkit, City AI Governance Dashboard, and AI Policy Dashboard provide case studies, readiness assessment tools, and examples of city AI policies from across the country.

League of Minnesota Cities (LMC). LMC has published guidance on AI use consistent with the Minnesota Government Data Practices Act, and its frameworks for data classification and risk management are already referenced in St. Louis Park’s internal guidelines.

MNIT (Minnesota IT Services). MNIT guidance on AI risk classification and data governance provides a state-level framework that St. Louis Park can align with.

Bloomberg Philanthropies / City AI Connect. A global learning community for cities exploring generative AI in public services, providing peer exchange, expert access, and a shared resource hub.

AI4Cities (European). A collaborative project involving Helsinki, Amsterdam, Copenhagen, Paris, Stavanger, and Tallinn using pre-commercial procurement to develop AI solutions for mobility and energy, offering useful models for collaborative AI procurement.

Global Observatory of Urban Artificial Intelligence (GOUAI). Supported by UN-Habitat, GOUAI maintains the Atlas of Urban AI, a repository of ethical AI initiatives and projects from cities worldwide.

Conclusion: Toward an AI-Compatible City

The question this white paper set out to answer is how cities need to adapt to become AI-compatible—not just technically, but in terms of governance, culture, and policy. The Commission’s answer, drawn from peer city experience and national guidance, is that AI compatibility is less about adopting any particular technology and more about building the underlying conditions that make AI trustworthy and useful.

Those conditions include: clean, integrated data that AI systems can act on reliably; clear governance structures that assign human accountability for every AI system; staff who understand both the capabilities and the limitations of the tools they use; procurement practices that require vendors to be transparent and auditable; and a commitment to measuring whether AI improves service equity, not only efficiency.

St. Louis Park is better positioned than many peer cities because it has already established an internal AI use policy, operates an existing resident service platform in MySLP, and has a Community Technology Advisory Commission actively engaged in these questions. The work ahead is to extend and deepen that foundation—incrementally, measurably, and with residents at the center.

The Commission recommends that the city:

- Formalize a public AI inventory and governance process, requiring council review for any AI tool that affects resident rights or services.
- Conduct an AI readiness assessment covering data quality, workforce capability, system integration, and cybersecurity posture before adopting new AI tools in any department.
- Prioritize internal productivity tools and low-risk pilot applications in the near term, using those experiences to build institutional knowledge before moving to higher-stakes use cases.
- Establish equity measurement as a non-negotiable component of any AI evaluation, with regular reporting to council on whether AI tools are improving or worsening service equity across neighborhoods.
- Stay connected to the NLC, LMC, and peer city networks to monitor how the technology—and the governance approaches to it—continue to evolve.

AI will not transform city government on its own. But cities that invest now in the governance, data, and human capacity to use AI responsibly will be better positioned to deliver more efficient, equitable, and transparent services to their residents in the years ahead.

Appendix: Key Resources and References

National and State Guidance

- National League of Cities — AI in Cities: Report and Toolkit (2024)
- National League of Cities — City AI Governance Dashboard: nlc.org/resource/city-ai-governance-dashboard/
- National League of Cities — Use AI to Transform City Operations (2025): nlc.org/article/2025/07/31/use-ai-to-transform-city-operations/
- League of Minnesota Cities — AI guidance aligned with Minnesota Government Data Practices Act
- MNIT — Minnesota AI risk classification guidance
- NIST AI Risk Management Framework: nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

Peer City Policy Documents

- City of Seattle — Generative Artificial Intelligence Policy
- City of Boston — Interim Guidelines for Using Generative AI (2023)
- City of San Jose — AI Handbook / Digital Privacy Office (2024)
- City of Tempe — Ethical Artificial Intelligence Policy (2023)
- City of Lebanon, NH — AI Registry
- New York City — Artificial Intelligence Action Plan (2023)
- City of Chattanooga — AI Prompt Library

Selected Case Study Sources

- Sacramento 2025–2028 AI Strategy and Roadmap
- Hamilton, ON — Bloomberg Harvard City Program All4One Pilot
- APWA / Eagan, MN — AI for Pavement Assessment Case Study
- Google Research — Project Green Light: sites.research.google/greenlight/
- Bloomberg Philanthropies / Johns Hopkins — City AI Connect: cityaiconnect.jhu.edu
- Global Observatory of Urban Artificial Intelligence — Atlas of Urban AI: gouai.cidob.org/atlas/
- Eurocities — Algorithmic Transparency Standard: algorithmregister.org

This white paper was prepared by the St. Louis Park Community Technology Advisory Commission. It draws on publicly available municipal policy documents, the NLC's AI in Cities Report and Toolkit, the Sacramento and Hamilton AI strategy documents, and reporting from Government Technology, Smart Cities Dive, and other municipal technology publications.

Community technology advisory commission

Presented by: Konnor Slaats and Rudyard Dyer
Staff Liaison: Jacque Smith

Policy question

Does the council have any work direction for this commission?

Workplan summary

Workplan Item	Status Update	Completion Timeline
Assess the city's customer response management tool	In process. Statistics and information shared by staff with commission members for their assessment. Discussion continues. Some information was included in the AI whitepaper.	4Q 2026
Produce a whitepaper on the future of generative AI for cities	Completed and shared with staff liaison for use in communications and technology department	July 2026
Participate in city website redesign review process	CTAC members participated in user experience workshops and provided comments on user feedback, wireframe and design. Staff liaison has provided updates on redesign process.	Sept 2025
Support citywide Vision 4.0 process	Commission members received information from the community engagement coordinator throughout the Vision 4.0 process; including survey invites; a commission member hosted a community meeting; commission members reviewed Dec. 8, 2025, council report to inform CTAC workplan goals for 2026	December 2025
Stay informed about pending legislation affecting cable and technology	Ongoing	Ongoing