



From AI Prompts to Policy

▶ IOWA LEAGUE OF CITIES | SEPTEMBER 24, 2026 | BETTENDORF

Jason Peters

AI & Data Solutions Lead, VC3

jason.peters@vc3.com

Who's in the room



CURRENT CUSTOMERS

Hannah Pauly

Strategic Advisor

hannah.pauly@vc3.com




NEW CUSTOMERS

Shane Glenn

Account Executive

shane.glenn@vc3.com



YOUR HOST

Jason Peters

AI & Data Solutions Lead

jason.peters@vc3.com

Where we're headed today

01 Icebreaker: AI vs. The Room

Beat the clock against AI

03 How to prompt AI

Better prompts, better results

05 The risks are real

What goes wrong, and how to prevent it

07 Your path forward with VC3

Assess, govern, enable, and Q&A

02 Intro to AI

What it is, and where it stands in 2026

04 AI assistants

Grounded in your data, working 24/7

06 From prompts to policy

AI governance that keeps you safe



AI vs. The Room



ICEBREAKER | BEAT THE CLOCK, THEN SEE WHAT AI DOES

How it works

3

ROUNDS

Real municipal scenarios, one at a time

90s

ON THE CLOCK

Your table drafts first, 60 to 90 seconds

1

LIVE REVEAL

Then AI gets the same brief, live on screen



THE POINT

Spot exactly where AI earns its keep

ROUND 1

The Water Main Break

THE BRIEF

It's **7:15 a.m.** A 12-inch water main just broke on Main Street. Public Works has shut off the section, and about **1,800 homes and businesses** on the north side have low or no water pressure. State rules require a **precautionary boil advisory**: residents should boil tap water for one minute before drinking or cooking, or use bottled water. Crews expect repairs by **6 p.m. today**; the advisory lifts after two clean water samples, about 48 hours.

 **YOUR TABLE (90 sec):** Write the text / SMS alert, under 160 characters, that goes out right now.

ROUND 2

Decode the Grant Award

THE BRIEF

Riverton just won a **\$600,000 federal grant** to rebuild a park pavilion and add lighting. Now the fine print. The award requires compliance with **NEPA**, **Davis-Bacon prevailing wages**, **Build America, Buy America (BABA)**, and reporting under **2 CFR 200 (Uniform Guidance)**.

 **YOUR TABLE (90 sec):** In plain English, what does each of the four mean, and name one thing each makes you do differently on this project?

ROUND 3

Translate the Notice

THE BRIEF

Council just voted to move trash and recycling pickup from **Fridays to Mondays**, starting the week of **October 6**. Carts must be at the curb by **6 a.m.**, and missed-pickup calls go to Public Works at **867-5309**. About 12% of households speak **Spanish** at home, and a growing number of families speak **Vietnamese**.

 **YOUR TABLE (90 sec):** Write the resident notice, ready for everyone in town: English, Spanish, and Vietnamese, at a 6th-grade reading level.

PROMPT ROULETTE

You write it. AI builds it.

A **[ADJECTIVE]** city council meeting where **[NOUN]** are the council members, set in **[CITY]**, in the style of **[ARTIST / TV SHOW]**

Shout out your answers, and we'll generate it live.

THE BOTTOM LINE

AI doesn't replace your judgment.
It gives you more time to use it.

Where AI stands for cities in 2026

44%

of local governments are
experimenting with generative AI
today

+50%

growth in public-sector AI use
from 2024 to 2026

63%

say limited staff and expertise is
their #1 barrier

33%

have an AI policy in place

The pattern: tight budgets, staffing gaps, and rising resident expectations are pushing cities toward AI faster than the guardrails are being built.

THE TRAJECTORY

The best is yet to come

“AI today is at the AOL stage of the internet, and it will improve exponentially faster than anything we've ever seen.”

Which means:

The tools you evaluate today are the worst they will ever be. Start learning now, on your terms, with guardrails.

What is Artificial Intelligence?

Defined

A branch of computer science focused on creating machines or software capable of tasks normally associated with human intelligence.

The concept dates to a 1956 Dartmouth workshop, where a team first framed AI as a research discipline.

We've used it for decades

- ▶ Narrow AI has been around for years: spam filters, fraud detection, routing logic, rules-based systems, each very good at one task
- ▶ What changed: large foundational models (LLMs) that write, summarize, translate, analyze, and code without being custom-built for each job
- ▶ That flexibility is why AI is suddenly useful across so many roles

What is a Large Language Model?

The ultimate predictor

An LLM is a complex function that simply predicts the most probable next word in a sequence. It generates text by calculating probabilities.

Massive scale

The “Large” comes from billions of internal parameters (weights). The model learns by constantly adjusting these internal “dials.”

Learns context

Using an operation called “attention,” the model figures out a word's meaning from its context (e.g., river bank vs. financial bank).

Trained by reading and humans

It learns by reading vast amounts of text, then human feedback refines its behavior to make it more helpful and safe.



How to Prompt AI



PROMPT ENGINEERING | COMMUNICATING CLEARLY WITH AI

Prompt Engineering

1

Assign a Role

Clearly define the AI's role to provide context.

2

Provide Context

Give specific details and requirements.

3

Step-by-Step & Verification

Request AI to outline its reasoning and verify its own results.

4

Divide and Conquer

Break down complex tasks into smaller, manageable subtasks.

5

Reference & External Tools

Provide relevant data and leverage specialized tools when appropriate.

6

Input = Output

The quality of your prompt directly impacts the quality of AI's response.

Practical ways to use AI in your city

Grant applications

Analyze the solicitation, research regulations and community impact, then draft the proposal section by section.

Resident communications

Newsletters, notices, FAQs, and plain-language translations in seconds.

Meeting minutes

Turn a recording or transcript into clean minutes and a clear action list.

Staff reports & policy

First drafts of memos, staff reports, and policy language you then refine.

Research & lookups

“What does this ordinance mean?” Summarize code, contracts, and precedent.

Data & spreadsheets

Clean messy data, summarize trends, and explain the numbers in words.

One rule: *do this in an approved, secure tool, never paste resident or protected data into a public chatbot.*

Artificial Intelligence Assistants

Customizable AI Assistant

An assistant is a tailored AI that can be designed to perform specific tasks, answer questions, or automate workflows.

Grounded in Your Data

Assistants use your organization's files, emails, calendars, and systems to provide accurate, context-aware responses.

Automates Repetitive Work

Assistants can handle routine requests, generate documents, analyze data, and trigger actions.

Secure and Governed

Assistants respect your organization's security, permissions, and compliance policies, so sensitive information stays protected.

Assistants built for local government

Meeting Minutes Assistant

Transcribes meeting recordings, summarizes content, and provides key insights and action items.

Website Assistant

Tier 1 customer service available 24/7/365. Reduces call volumes while increasing accessibility of information and transparency of governance.

Policy Interpretation Assistant

Trains on your unique organizational policies and provides consistent, accurate interpretation for your teams, ensuring a unified message across the organization.

HR Assistant

Guides employees through HR questions with 24/7 support, freeing up HR time for the personal instruction and guidance where it has the highest impact.

The Website Assistant in action

The setup

-  Small city hall team fielding every call
-  No after-hours coverage
-  A rich website and decades of minutes, hard to search
-  Now: answers instantly 24/7, or routes to the right staff

What residents can ask, any hour

“When is trash pickup on my street?”

“What happened in the November 1987 commission meeting?”

“What rec programs are open this summer?”

“When does the pool open?”



The risks are real



WHY GUARDRAILS AREN'T OPTIONAL

Artificial Intelligence Risks

Unauthorized Data Access

Without proper security and file permissions, AI may surface sensitive or confidential documents to users who shouldn't have access.

Loss of Credibility

Untrained users may rely on AI-generated answers that are inaccurate, outdated, or “hallucinated” (fabricated by the model).

Compliance & Regulatory Violations

Integrated AI may inadvertently expose data subject to legal, regulatory, or privacy requirements (e.g., CJIS, FOIA, PCI) if not properly configured.

Shadow IT & Lack of Oversight

Staff may use unsafe AI tools or create and share agents without IT or leadership approval, leading to inconsistent practices and loss of control.

The risks are real: real-life examples

New York City

NYC's AI chatbot told businesses to break the law, telling landlords they can reject tenants with housing vouchers.

Fresno, CA

A school official resigned after using quotes fabricated by ChatGPT. She moved too fast under pressure and didn't fact-check. The school had no AI use policy or guidance.

Detroit, MI

Detroit police wrongfully arrested a man due to incorrect facial recognition, leading them to change their policy for its use.

Seattle, WA

AI-drafted police reports (no AI use policy) were found to undermine credibility in court. One test report included an officer who wasn't there.

Artificial Intelligence Risk Mitigation

1

AI Readiness Assessment

Identify gaps in security and compliance, evaluate preparedness, and get tailored safeguards aligned to NIST and OECD.

2

Implement AI Governance

Stand up an AI Use Policy and governance so acceptable use, transparency, and accountability are defined.

3

Staff Training

Prepare staff for their AI role, how to use it safely, verify output, and protect data.

4

Implement Secure Chat

Give staff a governed, secure alternative so they aren't reaching for risky public tools.



From prompts to policy



AI GOVERNANCE | BEYOND THE POLICY DOCUMENT

AI Governance: beyond policy

The AI Use Policy is the anchor. A working program adds the pieces that make it live:



AI Governance Council



AI Use Policy



AI Governance Handbook



Incident Response Procedure



**AI Tool Inventory & Risk
Assessment**



Define risk tolerance



AI Training Roadmap



Implementation Guide

AI Governance examined: the AI Use Policy

01 Purpose & Scope

02 Guiding Principles

03 AI Governance & Roles

04 AI Risk Management Process

05 Acceptable Use

06 Prohibited Uses

07 Data Governance

08 Procurement & Third-Party
Management

09 Enforcement, Monitoring &
Review

AI policy resources and best practices

ISO/IEC 42001:2023

LEADING STANDARD

The newest AI policy standard. Sets requirements to establish, run, and continually improve responsible AI governance across your whole organization.

OECD AI Policy Principles

Promote AI that is innovative, trustworthy, and respects human rights and democratic values. Adopted 2019 and used across the US, EU, Council of Europe, and UN.

NIST AI Risk Management Framework

A voluntary, structured approach to identifying, assessing, and managing AI risk, with a companion Playbook that turns the principles into actionable steps.

San Jose GovAI Policy Coalition

A coalition of local governments focused on responsible AI, sharing policy guidance, templates, and collaboration. You don't have to start from a blank page.



Your path forward



HOW VC3 HELPS YOU GET THERE SAFELY

Three steps to AI you can trust

1

ASSESS

AI Readiness Assessment

We scan for shadow AI and the data-exposure risks AI would amplify, before you roll anything out. Available on its own or as part of Managed AI.

2

GOVERN

Governance, done for you

AI Use Policy, Governance Handbook, and Incident Response Procedure, tailored to your city and built on NIST and OECD.

3

ENABLE

Managed AI Assistants

VC3 Secure Chat plus the assistants, meeting minutes, website, policy, HR, grounded in your data, with foundational training.

When you want a partner, not just a tool

Your virtual Chief AI Officer

Structured like a vCISO engagement: a strategic AI roadmap plus ongoing monthly and quarterly reviews, so AI keeps paying off long after the first project.

What you get

-  **A prioritized AI roadmap tied to your goals**
-  Secure access to your own data through MCP infrastructure
-  Bespoke agents that optimize whole workflows, not one-off tasks
-  Data segregation done right, so the wrong people never see the wrong records
-  Monthly and quarterly reviews to keep it on track

Thank you.



Leave ready to act: pick one use case, write one page of policy, and choose one secure tool.

Jason Peters

AI & Data Solutions Lead, VC3

jason.peters@vc3.com

Managed AI Assistants • vCAIO Advisory • AI Governance for local government