Overcoming AI Paralysis: Practical Solutions for Public Entities

Moving from Analysis to Action

John Huisman, COO of Blue Cypress

POOL/PACT Conference • October 14, 2024

Today's Journey: From Understanding to Action

Hour 1: Foundation & Strategy (60 minutes)

- Why AI matters for public entities
- Understanding what AI actually is
- Your implementation options
- Technology requirements for executives

Hour 2: Interactive Workshop (60 minutes)

- Live demonstrations with your use cases
- Individual hands-on exercise
- Group presentations & discussion

This session emphasizes **practical, hands-on learning** where you'll work directly with your real challenges. The second hour will transform today's concepts into actionable plans for your organization.

Why AI Matters for Public Entities

Setting the foundation for moving from paralysis to action

The AI Acceleration That Changes Everything

1

Early AI Development

15+ Years to reach human performance

Traditional machine learning required extensive data preparation and manual feature engineering

2

Recent AI Advances

2-3 Years to reach human performance

Deep learning accelerated development but still required significant technical expertise

3

Latest AI Breakthrough

0.5-1 Years to reach superhuman performance

Generative AI democratizes access with natural language interfaces

- AI capabilities have accelerated dramatically, creating unprecedented opportunities.
- Public entities can now "leapfrog" traditional technology adoption, accessing advanced solutions previously limited to large corporations.

Public Sector Challenges vs. AI Opportunities

Traditional Challenges

- Budget constraints and lengthy approval cycles
- Risk aversion and strict compliance requirements
- Technical debt from legacy systems
- Limited technical expertise and resources
- Mission delivery prioritized over infrastructure

These constraints are real and won't disappear overnight. However, AI offers a path to work *with* these limitations rather than against them.

AI Opportunities

- Leapfrog expensive infrastructure investments
- Improve service delivery with existing resources
- Turn data into actionable insights
- Automate routine tasks for strategic focus
- Enhance decision-making capabilities

AI can amplify your existing capabilities without requiring complete system overhauls or massive budget increases.

The Cost of Staying on the Sidelines

Analysis Paralysis

- Endless research and committee meetings
- Missed opportunities while "studying the problem"
- Members served with outdated approaches

Competitive Disadvantage

- Other pools gaining AI-powered insights
- Falling behind in service capabilities
- Higher operational costs over time

Increasing Complexity

- AI landscape becomes more complex to navigate
- Legacy system integration becomes harder
- Staff expertise gap widens

The Alternative

Action-Oriented Approach: Start with measured pilots, learn through doing, build confidence with manageable wins

Understanding AI

What It Actually Is

Demystifying the technology behind the headlines

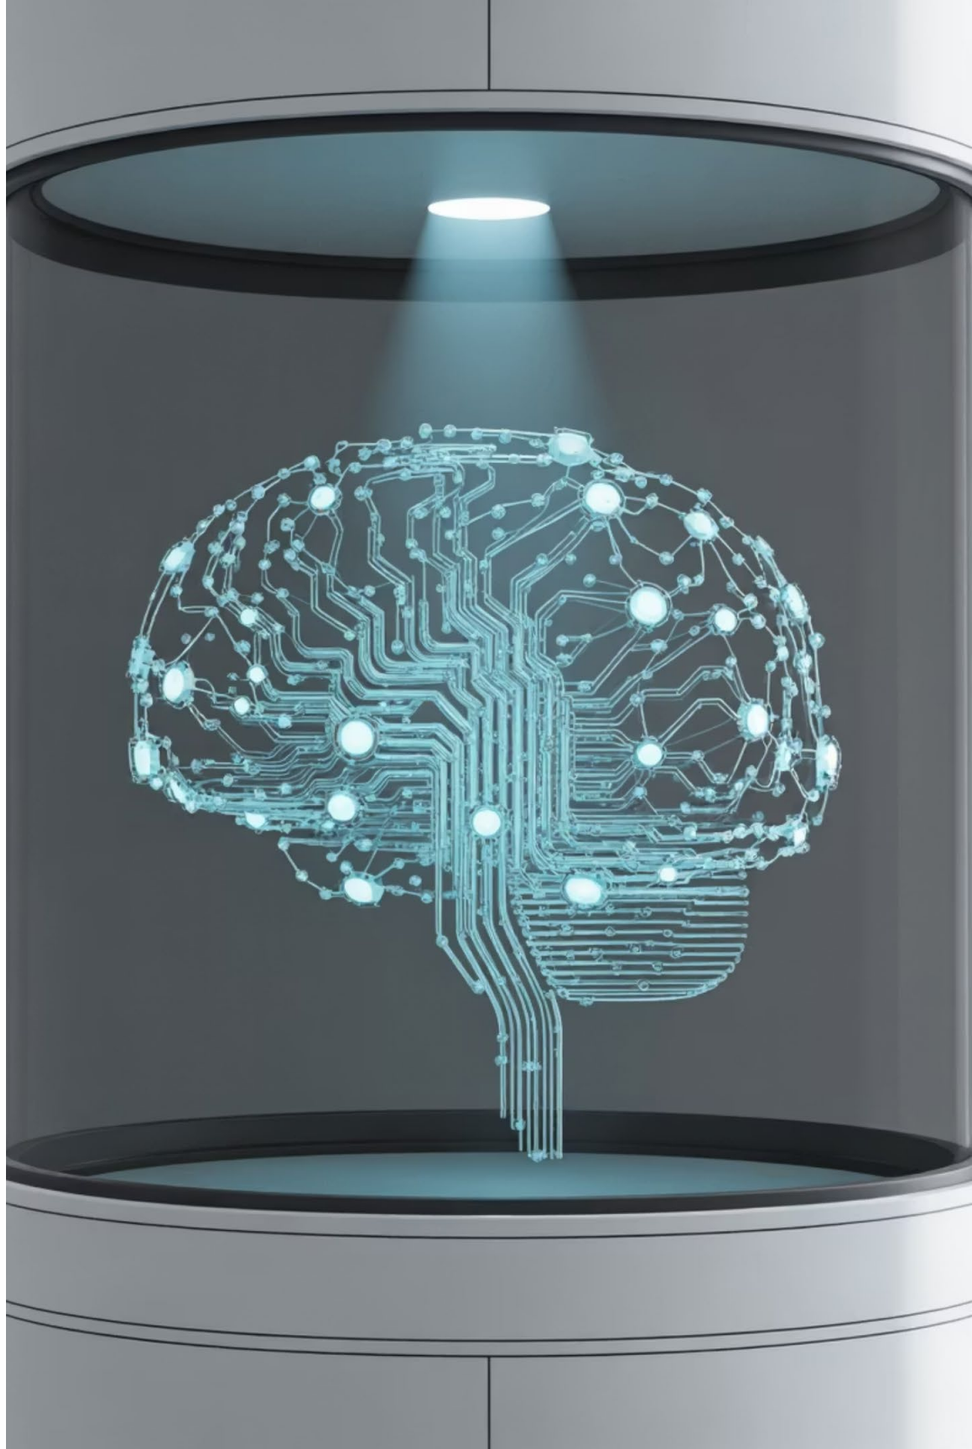

AI is a Toolkit, Not a Single Tool

Predictive Analytics

Forecasting claim costs and frequency across your member base. Uses historical patterns to anticipate future trends and identify potential risk areas before they become costly problems.

Computer Vision

Automated damage assessment from photos submitted in claims. Can analyze images to estimate repair costs, identify damage patterns, and flag potential fraud indicators.

Natural Language Processing

Extracting insights from claim narratives and loss descriptions. Processes written reports to identify trends, categorize incidents, and surface important details automatically.

Pattern Recognition

Looking for unusual claim patterns, problem scenarios, or emerging member issues. Identifies anomalies that human reviewers might miss in large datasets.

Generative AI

"Chat with your data" interfaces without any technical skills required. Ask questions in plain English and get intelligent responses from your organization's information.

Generative AI gets the most attention in media coverage, but it's just one powerful tool in a comprehensive toolkit. The real magic happens when these tools work together.

Generative AI: The Game Changer for Public Entities

What it is:

AI trained on billions of parameters that understands and generates human language, allowing natural conversation with your data systems.

Why it matters:

Eliminates the need for technical expertise to access insights. You can ask questions of your data in plain English and receive intelligent, contextual responses.

Example queries you could ask:

Public Works Director: Show me all road repair work orders over $50K and identify common delay factors

Human Resources Director/City Manager: What are the top three reasons for employee turnover in our department over the past year, and what interventions have worked in similar-sized counties?

Sheriff/Chief: Analyze property crime patterns in our jurisdiction over the past 12 months, identify the top three hotspots, and suggest operational adjustments for patrol schedules

Real-World Applications for Your Organization

Service Request Processing

Automated categorization and consistency checks for faster processing. Reduce manual review time while maintaining accuracy and identifying potential issues early in the process.

- Automatic claim categorization
- Duplicate detection
- Missing information alerts

Community Issue Pattern Identification

Spot trends across service requests before they become costly problems. Early warning systems that help prevent infrastructure failures or service gaps rather than just manage them after they occur.

- Emerging issue detection
- Neighborhood vulnerability scoring
- Seasonal pattern analysis

Predictive Maintenance

Schedule infrastructure maintenance before failures occur. Optimize maintenance budgets while reducing unexpected downtime and emergency repairs.

- Equipment lifecycle tracking
- Failure prediction models
- Maintenance optimization

Resident Communication

Personalized service and proactive outreach programs. Deliver the right information to the right residents at the right time for maximum impact.

- Need-based messaging
- Service recommendations
- Policy updates targeting

Service Improvement Targeting

Data-driven programs that prevent problems rather than just manage them. Focus limited resources on the highest-impact interventions for maximum community benefit.

- High-need area identification
- Intervention effectiveness tracking
- Resource allocation optimization

Your Implementation Options

Security vs. Speed

Finding the right balance for your organization's needs and risk tolerance

Three Approaches to Implementing Generative AI Based on Data Sensitivity and Risk Tolerance

Public/Consumer AI

Leverage off-the-shelf AI models and public cloud services. Best for non-sensitive data and rapid prototyping

Using public AI to analyze general industry trends or optimize your website content
(e.g., ChatGPT for drafting member communications)

Private/Walled Access

Utilize third-party hosted AI solutions in controlled, secure cloud environments (e.g., Azure, AWS). Balances capability with data privacy

Vendor-hosted AI analyzing your claims data in a secure, dedicated environment (e.g., Microsoft Azure AI services with private cloud instances)

Custom/Closed Environment

Develop bespoke AI models within your own on-premise infrastructure. Offers maximum security and tailored capabilities

AI running entirely on your own servers, processing ultra-sensitive member data (e.g., on-premise AI servers in your own data center)

This is not an all-or-nothing choice. Mix-and-match is common: public tools for non-sensitive tasks, walled/closed environments for anything touching member or claims data

Option 1: Public/Consumer AI Solutions

What it is:

Using publicly available AI tools that, unless an enterprise opt-out is offered, reserve the right to use your data for training or model improvement. No formal data protection addendum (DPA) or SOC 2 guarantees in the consumer tiers.

Examples for Public Entities

- Analyzing community survey responses
- Generating resident newsletters
- Summarizing industry news and trends
- Brainstorming public meeting agendas

Real Gains

Low cost, fast implementation, proven technology

Real Risks

Data exposure, no control over model training, does not meet regulatory requirements for protected information

Best For

- Strictly non-PII tasks (marketing copy, trend summaries)
- Never for resident data, employee records, or proprietary entity data

Option 2: Private/Walled Access Solutions

What it is:

Enterprise AI services with contractual data protection, SOC 2 compliance, and dedicated instances. Your data is encrypted at rest/in transit, not used for training, and protected by enterprise controls.

Examples for Public Entities

- Service request categorization and trend analysis
- Budget variance analysis and pattern identification
- Conversational interfaces for resident inquiries

Real Gains

Balance of security and capability, managed infrastructure, protection for sensitive data

Real Risks

Vendor dependency (feature drift, model updates, data-residency jurisdiction), ongoing subscription (OpEx) costs, limited deep customization

Best For

- Recommended for most public entities
- Handles resident data with enterprise-grade protection
- Includes data processing agreements, audit trails, and liability provisions
- Meets common state cyber-insurance requirements

Option 3: Custom Solutions in Closed Environment

What it is:

AI infrastructure within your complete control, often on-premises or in dedicated private cloud. Security and compliance are self-managed; no shared responsibility with a vendor.

Examples for Public Entities

- Comprehensive resident service profiles with PII
- Sensitive personnel or financial analysis
- Multi-system integration with custom security

Real Gains

Maximum security for sensitive data, full customization, regulatory compliance with state/local government requirements

Real Risks

High CapEx & staffing burden; patching, model updates, and insider-threat controls are entirely your responsibility

Best For

- Complete data isolation; air-gapped or single-tenant setups with no external dependencies
- Full control over security protocols, audit trails, and access
- Meets the most stringent regulatory requirements

Common Security Questions About Generative AI Options

How can we know our data is safe? I heard data can leak from ChatGPT.

Public ChatGPT (Option 1) does have sharing risks - that's why it’s only recommended for non-sensitive tasks. Option 2 uses enterprise AI with isolated instances, encryption, and contractual protections - your own private AI that others cannot access.

What about data deletion rights and regulatory compliance?

Enterprise agreements (Option 2) include data ownership rights, deletion capabilities, and "no training" clauses. These are different from consumer services and include SOC 2 compliance, BAAs, and liability provisions. Vendors provide financially backed SLAs and 72-hour incident-notification clauses.

Our data is our most valuable asset. How can we trust a third party?

Start with less sensitive use cases to build confidence. Option 2 providers offer the same security standards used by major banks and health systems, with full audit trails and cyber insurance requirements.

What would a data breach look like in Option 2?

Unlike public AI, an Option 2 breach requires defeating enterprise-grade isolation, encryption, and access controls; similar to breaching a major bank’s cloud environment. Risk is similar to any enterprise cloud service, with better controls than email or file sharing systems.

Remember: You can start with Option 1 for public content while using Option 2 for member data - it's not all or nothing

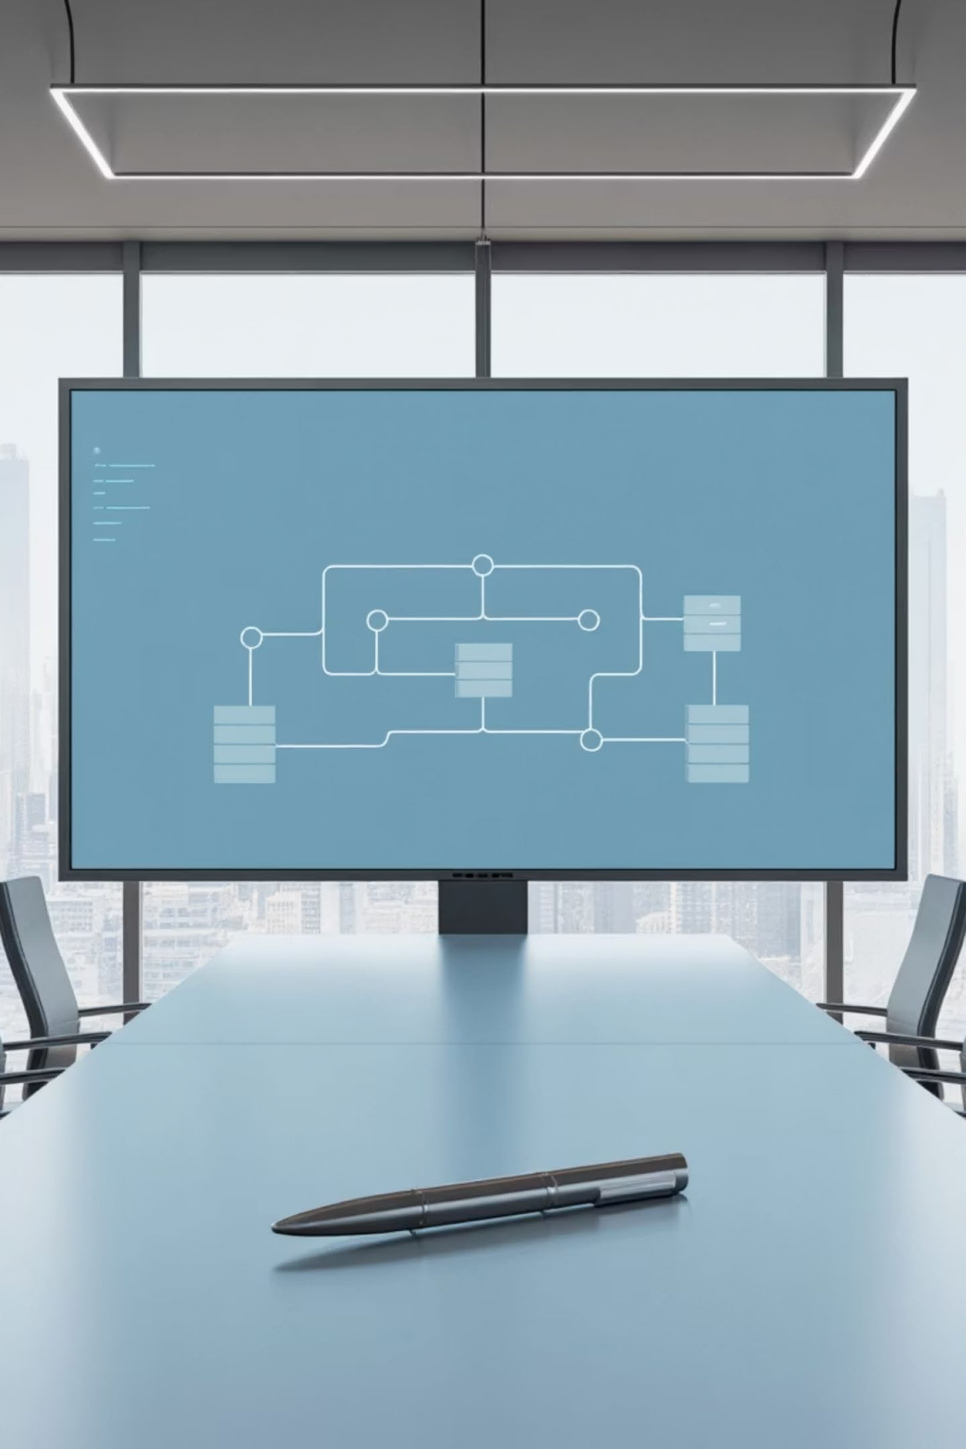

Technology Requirements

What Every Executive Should Know

Essential concepts without the technical complexity

Four Core Components Every Leader Should Know

Data Storage

Where your data lives and who can access it. The fundamental choice between cloud services (someone else's computers) and on-premise storage (your computers) affects everything else.

Key consideration: Control vs. convenience

Integration Systems

APIs and connectors that let your systems communicate without complete replacement. Think of these as translators between different software applications.

Key consideration: What systems need to talk to each other?

Security Layer

Role-based access controls plus encryption protecting data in transit and at rest. Determines who can see what information and how it's protected during processing.

Key consideration: Compliance requirements vs. usability

User Interface

How AI insights reach your team through dashboards, chat interfaces, alerts, or embedded features. The bridge between complex AI processing and practical daily use.

Key consideration: Staff adoption and ease of use

You don't need to become a technical expert, but understanding these building blocks helps you ask the right questions and make informed decisions about AI implementation in your organization.

Questions Every CEO Should Ask Their Tech Team

1 Data Control

"Who actually controls our data - us or our vendors? Can we access it directly for AI analysis without paying extra fees?"

2 System Integration

"Which of our systems can't currently share data? What would it take to connect them?" Note: AI can work with disconnected systems and actually help identify integration priorities

3 Security Requirements

"What are our member expectations for data protection vs. our actual liability risks?"

4 AI Readiness

"For each AI approach (public/private/closed), what specific technical obstacles do we face in the AI implementation?

5 Vendor Strategy

"If our current claims software vendor houses our data, what are the contractual and technical limitations on accessing it for AI analytics? Are there fees for data exports or API access?"

From Theory to Practice

You now have the foundation - time to apply it to your real challenges

Next: I'll demonstrate the approach with real examples, then you'll apply it to your use cases

How This Workshop Works

What You'll Do

- Watch me work with AI on a real problem (15 min)
- Practice with your own challenge (30 min)
- Share what worked (15 min)

What You'll Learn

- How to ask questions, give context, and refine responses
- Techniques to get better results
- How to make AI do actual work, verify accuracy, and refine outputs confidently

💡 Think of AI as a really good junior analyst

Give it context, ask it to think through problems, review its work, request revisions

You already know how to do this with human teammates

The AI Partnership Approach

Four techniques to make AI conversations productive

Give Rich Context

Start with: 'I'm a [role] at [org type]. We're dealing with [situation]. Our constraints are [budget/timeline/resources]. I need help achieving [specific outcome].

You can say "I don't know' - let AI ask YOU questions to help focus the work

Ask AI to DO the Work

❌ "What should I do?" (with no follow-up)

✅ “Create a plan for how to deliver that. Then tell me which parts you can do and what you need to get started.”

✅ "Analyze this spreadsheet of 500 claims and identify the top 10 risk patterns by frequency and cost

✅ "Can you create a robust AI policy document for me?"

Use Meta-Prompting

"Rate the below prompt below on a scale of 1-10. After you share your rating and rationale, write me a prompt that would be a 10/10. [your original prompt]"

Iterate and Refine

Review AI's work → Request changes → Build on what works

Treat it as an ongoing dialogue, not one-and-done

Remember, treat it as a junior analyst – what would you do if a human gave you a draft that you didn’t think was ready?

Starting With 'I Don't Know' Is Actually Powerful

The Secret: Tell AI When You're Stuck

1. Give it your situation
2. Tell it you're not sure what to do
3. Ask it to ask YOU questions

Why this works: You don't need answers first—you need good questions to clarify your thinking.
AI is great at helping with this.

Example Opening Prompt

> I work for [a mid-sized county/a school district/a municipal government]. We need to improve citizen engagement but have limited staff time and budget. Different departments have different communication preferences and priorities.
>
> I'm not sure how to approach this effectively. Can you help me think through this by asking me questions to clarify the best approach?”

Live Demonstration: Effective AI Prompting in Action

We'll work through a real use case using these techniques:

Reminder: Our Starting Prompt:

> I work for [a mid-sized county/a school district/a municipal government]. We need to improve citizen engagement but have limited staff time and budget. Different departments have different communication preferences and priorities. I'm not sure how to approach this effectively. Can you help me think through this by asking me questions to clarify the best approach?”

Advanced Techniques We'll Use:

- Meta-prompting to improve our initial prompt
- Clear work division: what AI handles vs. what you do
- Asking AI to do actual work, not just suggest
- Getting step-by-step instructions for your parts

What to Watch For:

- How AI responses improve with better prompts
- How iterative refinement produces better results
- How AI shifts from suggestions to actual work

Your Turn: Apply AI to Your Use Case

Your 30-Minute Challenge:

- Open AI tool of your choice (e.g., ChatGPT, Claude)
- Work through your specific use case systematically
- Use the deck we've been reviewing as a practice guide (slides 22-23)
- Focus on YOUR challenge (this is an individual exercise)

Resources Available:

- Use the deck we've been reviewing as a practice guide (slides 22-23)
- Raise your hand and I'll come by to answer questions/brainstorm
- Talk amongst your table, but try to do the work yourself

Remember: The goal is learning AI partnership skills. Progress over perfection. Focus on effective conversation, not perfect solutions.

Get started now!

Share Your AI Conversation Experience

For Presenters (5 minutes each):

- What challenge did you work on?
- Describe your most effective prompt (original vs. AI-improved version)
- What actual work did AI do for you beyond just suggestions?
- How did you divide work between AI and yourself?
- What surprised you?

For Audience:

- Questions about prompting techniques that worked
- Suggestions for getting AI to do more actual work
- Ideas for better work delegation
- Similar challenges where this might work

Focus: We're learning AI partnership skills, not evaluating solutions. Share what you discovered about working with AI as a collaborator.

Your AI Toolkit: From Conversation to Automation

Skills You Practiced Today:

01

Effective Prompting: Context setting and iterative refinement

02

Meta-Prompting: Getting AI to improve your prompts

03

Work Delegation: Having AI do actual work, not just make suggestions

04

Partnership Approach: Clear division of what AI handles vs. what you do

Next Level Tools (When You're Ready):

</>

Claude Code/Codex

AI that writes and executes code for you - enabling custom analysis and automation without programming skills

Custom AI Agents

AI assistants trained for your specific workflows and processes that have access to tools (e.g., web search, database access)

MCP/API Integrations

Connect AI directly to your existing systems (claims, finance, etc.) for automated workflows

Suggested Next Steps:

→ Use these techniques on a work challenge this week

→ Push AI to do more actual work in your daily tasks

→ Experiment with meta-prompting to improve your results

Taking This Forward

This Week:

- Apply these prompting techniques to 2-3 other work challenges (bring real data or scenarios to your AI conversations)
- Use meta-prompting on every significant prompt (ask AI to rate and improve your prompts it becomes natural)
- Push AI to do actual work – always consider asking 'Can you create this for me?' instead of 'How should I create this?

This Month:

- Share your best prompts with colleagues
- Start building a prompt library for common tasks
- Identify one workflow that could benefit from AI automation

When Ready:

- Explore advanced tools (Claude Code, custom agents)
- Use AI to help you craft a roadmap of where else you could use AI within your organization
- Share your AI wins with colleagues – teach others what you've learned to build organizational capability

Appendix

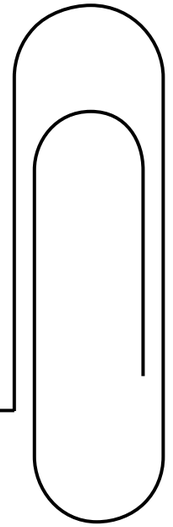

Key Questions When Evaluating AI Vendors

Technical Questions:

- ### AI Models and Capabilities

  "What specific AI models do you use, and how do you keep them updated?"

- ### Data Protection

  "How is our data protected and isolated from other customers' data?"

- ### Integration Options

  "What integration capabilities do you offer with our existing systems?"

- ### Model Management

  "How do you handle model updates and changes without disrupting our operations?"

Business Questions:

- ### Public Sector Experience

  "What's your experience with public sector clients and their unique requirements?"

- ### References and Case Studies

  "Can you provide references from similar organizations and specific success stories?"

- ### Pricing Structure

  "What's your pricing model and what's included versus additional costs?"

- ### Data Portability

  "What happens to our data if we terminate the contract? How is it returned or deleted?"

Start vendor conversations with these questions to quickly identify serious providers who understand public sector needs and can deliver on their promises.

Sample AI Policy Framework for Public Entities

Start simple. Adapt to your organization. Expand as you learn.

The Three Essentials:

Data Protection

What to include:

- Define data handling rules based on tool type:
- Public AI tools (ChatGPT free, Claude free): No PII, PHI, confidential member information, or personnel records
- Enterprise AI tools (paid versions): May be approved for sensitive data with proper agreements in place
- Rule of thumb: Match your data sensitivity to your tool's protection level

Accuracy Requirements

What to include:

- All AI outputs must be reviewed and verified by qualified staff
- Staff member remains accountable for final work product
- Document your verification process
- When in doubt, get a second review from a subject matter expert

Approved Use Cases

What to include:

- Start with low-risk applications (drafting, research, summarization)
- Require supervisor notification or approval for new use cases
- Expand approved uses as confidence grows
- Document what works and what doesn't for organizational learning

Starter Policy Language:

Employees may use AI tools for work purposes subject to the following requirements

- Public/Free AI Tools: May be used for drafting, research, and summarization. Confidential, personal, or sensitive information must never be input into free AI tools
- Enterprise AI Tools: May be used with sensitive data only when: (1) the organization has appropriate data protection agreements in place, (2) the specific use case has been approved by [designated authority], and (3) the tool has been vetted by IT/legal
- All AI-generated content requires human review and verification before use
- Employees remain fully accountable for accuracy and appropriateness of all final outputs
- Questions or concerns should be directed to [designated person/department]

This is a starting point, not legal advice. Consult your legal counsel and adapt to your organization's specific needs and risk tolerance.