

AI Technologies and The Water Utility Sustainability Gap

*Helping Rural Utilities Build Capacity
Without Deepening
Cyber-Operational Risk*

WRITTEN BY
Jason M. Nichols, PhD

2026



**INSTITUTE FOR
HOMELAND SECURITY**
SAM HOUSTON STATE UNIVERSITY



AI Technologies and the Water Utility Sustainability Gap

Helping Rural Utilities Build Capacity
Without Deepening Cyber-Operational Risk

Jason M. Nichols, PhD
HydrospherIQ Consulting, LLC
jason@hydrospheriq.com

ABSTRACT

This research paper examines how adoption of generative and agentic artificial intelligence (AI) may affect the operational sustainability gap between larger water utilities and many small or rural water and wastewater systems. This gap is the uneven capacity of these utilities to operate, maintain, and secure the systems needed to provide safe and reliable service over time. The gap is shaped by workforce constraints, compliance and infrastructure burdens, vendor dependence, and cybersecurity limitations. Well-governed AI adoption could help narrow that gap by reducing administrative, compliance, funding, and operational-planning burdens. However, the same tools could widen the gap if they are adopted before utilities can meet the cybersecurity, data-governance, vendor-management, and operational-fallback requirements for safe use. The problem is not rurality itself. It is the interaction of capacity constraints, operational technology (OT) and data exposure, vendor and access-control weaknesses, and limited recovery capability that undermines safe use. The brief recommends an incremental adoption approach. The approach starts with low risk uses that build capacity, protect sensitive and operational information, and preserve human accountability. It delays higher risk uses until the organization can establish and verify AI governance and safeguards.

Keywords

Water utilities; wastewater systems; rural infrastructure; cybersecurity; operational technology; generative AI; agentic AI; critical infrastructure resilience; vendor risk.

Table of Contents

INTRODUCTION	1
ANALYSIS AND FINDINGS	3
The Sustainability Gap Is Structural, Widespread, and Diffuse	3
Foreseeable AI Risks in Water Utility Operations	5
External exploitation of OT and cyber weaknesses	5
External exploitation of people and trust relationships	5
Internal governance failures in AI-enabled utility workflows	6
Foreseeable AI Benefit Pathways in Water Utility Operations	8
Administrative and Compliance Capacity	8
Operational Knowledge and Workforce Support	8
Maintenance and Operational Planning Documentation	9
Cybersecurity and Resilience Documentation	9
AI and the Sustainability Gap: Incremental Adoption or Unmanaged Drift	11
Policy Implications for Safe AI Adoption	12
A Texas Perspective: Rural Water Infrastructure, AI Governance, And Regional Resilience	13
PRACTICAL APPLICABILITY / RECOMMENDATIONS	15
CONCLUSION	18
DISCLOSURE STATEMENT	19
FUNDING / GRANT INFORMATION	19
AUTHOR BIOGRAPHY	19
REFERENCES	20

INTRODUCTION

Since the public release of ChatGPT in 2022, generative AI has moved rapidly into public, commercial, and workplace use, with more agentic systems following.

Generative AI produces text, code, images, summaries, recommendations, and other outputs in response to prompts or data inputs. For water utilities, these tools can serve as writing assistants, search-and-synthesis tools, coding aides, training aids, and decision-support interfaces for technical, regulatory, operational, and cybersecurity information.¹ Agentic AI goes further by pursuing goals through multiple steps, using external tools or data sources, making intermediate decisions, and recommending or initiating actions through connected software. In a water utility setting, agentic tools may support compliance reporting, work-order triage, supply-chain coordination, maintenance planning, vendor-access review, and incident documentation.

Many small and rural water and wastewater systems entered this new era with an operational sustainability gap that predates AI. Operational sustainability refers to a utility's continuing ability to operate, maintain, and secure the systems needed to provide safe and reliable water and wastewater service. The gap refers to the uneven capacity of larger urban utilities and many small or rural utilities to sustain those functions over time. The difference is not simply geography. Larger utilities are more likely to have dedicated engineering, legal, procurement, IT, cybersecurity, and compliance functions, while many small systems rely on limited staff, part-time managers, contract operators, vendors, or regional assistance to meet similar needs. The U.S. Environmental Protection Agency (U.S. EPA) identifies small-system sustainability as a persistent challenge because small utilities often spread fixed operating, compliance, capital, and cybersecurity costs across a limited customer base.² Workforce pressures compound these constraints as experienced operators age out of the sector without adequate replacements.³

Cybersecurity is a significant contributor to this capacity divide. As utilities adopt more connected digital technologies, they also take on more recurring technical work. Remote monitoring, cloud platforms, digital maintenance systems, and connected operational technology can improve visibility, documentation, efficiency, and planning. They also create more accounts, vendor connections, logs, backups, and remote-access paths to manage. Cybersecurity is therefore not a one-time purchase. It requires recurring work

¹ Autio et al., 2024.

² U.S. EPA, n.d.

³ Kane & Tomer, 2018; U.S. EPA, 2024a.

across access management, asset monitoring, incident preparation, staff training, and vendor oversight.⁴

In earlier waves of digital technology adoption, larger urban utilities were often better positioned to realize the benefits because they could absorb new cybersecurity demands through dedicated staff, broader rate bases, and formal governance processes. Many small or rural utilities faced a harder choice. They could forgo sustainability-enhancing technologies or manage the same expanding cybersecurity demands with fewer personnel, older systems, and greater vendor dependence. Generative and agentic AI may sharpen that divide. Better-resourced utilities are more likely to gain additional capacity from AI while managing the required safeguards. Smaller utilities may need the benefits more urgently, but be less able to use AI safely. Without practical support, AI adoption could widen the operational divide.

AI can help reduce the gap if utilities adopt it safely. Generative AI can help draft administrative, training, funding, maintenance, public-facing, and compliance materials. Agentic AI may eventually support multi-step work, such as gathering records, flagging missing information, routing drafts for review, and documenting incident actions. These uses could give small utilities support that previously required more staff time or outside expertise. The same value proposition creates risk. A tool that helps a small staff do more could bypass governance, expose sensitive information, increase vendor dependence, or influence decisions before the utility can verify the output. AI features may also appear through routine software updates, vendor-managed platforms, business and operational software, or informal staff use of public chatbots. In those settings, a tool can become operationally important before the utility has approved its use or reviewed its safeguards.

The central argument for this brief is not that rural water utilities are inherently insecure and therefore must avoid emerging AI technologies. It is that AI can help narrow the urban-rural sustainability gap only if small utilities and their support partners build the capacity to govern, secure, review, and recover from AI-enabled workflows. Without those supports, AI could deepen the same divide it promises to reduce. The policy challenge is to create structures that support incremental, safe AI adoption in ways that strengthen rather than undermine utility sustainability.

⁴ Cybersecurity and Infrastructure Security Agency, Environmental Protection Agency, & Federal Bureau of Investigation, 2024e.

ANALYSIS AND FINDINGS

The Sustainability Gap Is Structural, Widespread, and Diffuse

The operational sustainability gap is not limited to small, unusually fragile utilities. It is a structural condition spread across a large and diverse population of systems. In the context of AI technology, the gap is more about operational capacity than infrastructure. It is about the recurring ability to manage people, technology, vendors, cybersecurity, compliance, and recovery. EPA reports that 95 percent of the 145,720 active public water systems serve 10,000 or fewer people, and approximately 79 percent serve 500 or fewer people.⁵ The reality is that the water sector includes a very large number of systems that must operate with limited staff and uneven access to technical support.⁶ While this does not mean most Americans receive water from small or insecure systems, it does mean that the sustainability gap is widespread.

The issues are not uniformly a rural or small system phenomenon. EPA's capacity-development framework defines sustainable performance through technical, managerial, and financial capacity, which includes the ability to maintain compliance and provide safe drinking water over time.⁷ In 2022, a large system in Jackson, Mississippi caught in the sustainability gap suffered a service disruption involving around 150,000 residents. The crisis left residents without drinkable water and under a boil-water notice for 40 days. Beyond infrastructure deficiencies, the crisis was reported to be driven by staffing shortages, operations and maintenance issues, and long-term regulatory violations.⁸ For AI technology, cybersecurity gaps are critical. GAO found that systems across the water sector struggle with workforce skills gaps, older technologies that are difficult to update, and limited cybersecurity investment.⁹ The American Water Works Association's 2025 State of the Water Industry summary also identified public trust, workforce, and cybersecurity among the sector's top challenges. It reported that 26 percent of small systems had low to no ability to implement cybersecurity measures.¹⁰

The challenge posed by the sustainability gap is difficult to solve because the condition exists in diffuse operational environments. Water utilities operate under multiple types of jurisdictions, ownership models, vendor relationships, rate bases, and regulatory

⁵ U.S. EPA, 2025.

⁶ U.S. EPA, n.d.; U.S. EPA, 2024a; Government Accountability Office, 2024.

⁷ U.S. EPA, 2025.

⁸ Associated Press, 2023.

⁹ Government Accountability Office, 2024

¹⁰ American Water Works Association, 2025

environments.¹¹ This fragmentation is a long-standing feature of U.S. drinking-water governance. The National Research Council described the small-systems problem as long recognized but unresolved. It traced the fragmented U.S. water-supply network to local development patterns in which towns and cities built their own systems under local, state, and federal oversight.¹² EPA's water-system partnership guidance similarly treats partnerships as a means by which to build technical, managerial, and financial capacity across systems that face similar challenges but may differ in ownership, geography, operations, and governance structure.¹³ This fragmentation makes sector-wide improvement difficult because no single governance structure, procurement channel, or technical-assistance model reaches all systems equally.

Fragmentation may localize some operational risks, but modern AI can also scale risk across many organizations. A failure confined to one small utility may have local consequences. By contrast, a shared AI service, vulnerability, or attack method could affect many utilities in similar ways. The same public chatbot, vendor platform, compliance assistant, or automated workflow may be used across many water and wastewater systems. As with cloud services and software supply chains, this shared dependence can create common points of failure. NIST describes one form of this risk as an algorithmic monoculture, where many users rely on the same model or algorithm and may therefore experience correlated failures.¹⁴

AI systems also introduce some less familiar risks. "Prompt injections" occur when malicious or misleading instructions cause an AI system to ignore its intended rules or misuse connected data and tools. "Excessive agency" occurs when an AI system is given more authority to act than is necessary or adequately controlled. Other risks include disclosure of sensitive information, insecure third-party components, unsafe or inaccurate outputs, and overreliance on AI-generated recommendations¹⁵. Recent cyber-threat assessments warn that AI can increase the effectiveness, scale, and efficiency of malicious activity. Examples include phishing, impersonation, reconnaissance, and support for exploiting vulnerable systems.¹⁶

Given the water sector's uneven operational capacity, AI adoption should be understood as both a partial remedy for the sustainability gap and a source of new risk. Many small

¹¹ U.S. EPA, n.d.; Government Accountability Office, 2024

¹² National Research Council, 1997

¹³ U.S. EPA, 2025

¹⁴ Autio et al., 2024

¹⁵ Autio et al., 2024; Open Worldwide Application Security Project, 2025.

¹⁶ Federal Bureau of Investigation, Internet Crime Complaint Center, 2024; National Cyber Security Centre, 2024.

and rural systems have the strongest practical need for AI-enabled support because they must manage operations, compliance, modernization, public communication, cybersecurity, and incident recovery with limited staff and uneven access to technical assistance. Yet those same capacity limits can make AI adoption harder to govern safely. This creates a central tension. The systems that may benefit most from AI-enabled capacity support may be least able to manage the associated risks.

Foreseeable AI Risks in Water Utility Operations

The water-sector evidence base for actual generative or agentic AI incidents remains limited, and there is not yet a large public dataset of AI-driven water utility disruptions. At the same time, the absence of widespread reported AI incidents should not be mistaken for an absence of foreseeable risk. The available evidence points to three broad pathways by which AI could affect water and wastewater utilities:

1. External exploitation of OT and cyber weaknesses
2. External exploitation of people and trust relationships
3. Internal governance failures in AI-enabled utility workflows

External exploitation of OT and cyber weaknesses

Direct water-sector incidents show that attackers already exploit exposed human-machine interfaces (HMIs), virtual network computing (VNC) remote access, default passwords, weak credentials, missing multifactor authentication, and fragile fallback capacity. CISA and partner agencies reported that pro-Russia hackers manipulated water and wastewater HMIs, changed settings, disabled alarms, changed administrative passwords, and forced some operators to revert to manual control.¹⁷ IRGC-affiliated actors similarly compromised Unitronics programmable logic controllers (PLCs) used in water and wastewater facilities and other sectors. The devices were often connected to the internet and protected by default passwords or none at all.¹⁸ Generative and agentic AI could make this pathway faster and easier. Adversaries could use it to identify exposed systems, generate phishing or exploit-support material, summarize public utility information, or scale attacks across similarly configured small systems.¹⁹

External exploitation of people and trust relationships

Water utilities depend on trusted communications among operators, managers, board members, vendors, engineers, regulators, emergency-management partners, and

¹⁷ Cybersecurity and Infrastructure Security Agency et al., 2024a.

¹⁸ Cybersecurity and Infrastructure Security Agency et al., 2024c.

¹⁹ National Cyber Security Centre, 2024.

customers. Generative AI can make malicious messages more persuasive by producing realistic emails, voice messages, meeting summaries, instructions, public notices, or vendor communications. Federal reporting warns that criminals are already using generative AI to facilitate fraud. The United Kingdom’s National Cyber Security Centre also assesses that AI will likely increase the effectiveness and scale of phishing, reconnaissance, social engineering, and related cyber activity.²⁰ For water utilities, this matters because a false instruction, fake vendor request, fraudulent payment approval, or misleading emergency message can affect access decisions, public communications, incident response, or operational coordination.

Internal governance failures in AI-enabled utility workflows

These risks arise when a utility, vendor, or staff member adopts AI without clear limits on data, permissions, review, accountability, and fallback procedures. NIST and OWASP identify generative and agentic AI risks such as hallucination, prompt injection, sensitive-information disclosure, unsafe output, supply-chain weakness, excessive agency, and overreliance.²¹ OT security guidance also emphasizes that decisions affecting connected operational environments should be evaluated through safety, business knowledge, OT data protection, segmentation, supply-chain security, and people-centered operations.²² In water-utility practice, these risks can arise when staff enter customer data, credentials, system diagrams, supervisory control and data acquisition (SCADA) screenshots, cybersecurity details, or incident records into poorly governed tools. They can also arise when a vendor adds AI features without review or a chatbot produces confident but incorrect compliance or process guidance. The risk becomes more direct when an AI agent routes work orders, drafts notices, updates records, or recommends operational actions before a qualified person verifies the result.

Table 1 translates these three pathways into practical AI risk classifications for water utilities. The classifications do not claim that each risk has already produced a documented water-sector AI incident. Rather, they identify foreseeable risk categories supported by direct water-sector cyber incidents, broader critical-infrastructure threat reporting, AI security guidance, and cross-sector AI analogies. This framing supports a cautious but actionable conclusion. Realized water-sector AI incidents remain few, but the risk landscape is foreseeable enough to inform adoption policy now.

²⁰ Federal Bureau of Investigation, Internet Crime Complaint Center, 2024; National Cyber Security Centre, 2024.

²¹ Autio et al., 2024; Open Worldwide Application Security Project, 2025.

²² Australian Signals Directorate’s Australian Cyber Security Centre et al., 2024.

Table 1. AI risk classifications and water-sector relevance.

<i>Risk Classification</i>	<i>Water-Sector Example</i>
AI-enabled reconnaissance and exploitation	An attacker uses generative AI to identify exposed HMIs, vulnerable remote-access tools, default credential patterns, public utility documents, vendor relationships, or likely phishing targets, then scales attacks against small utilities with limited monitoring capacity.
Impersonation and authorization fraud	AI-generated emails, voice messages, meeting summaries, or video/audio impersonations appear to come from a utility manager, board member, regulator, vendor, engineer, or emergency-management partner and are used to approve payments, change access, redirect staff, or create false urgency.
Malicious input and prompt-injection attacks	A malicious email, PDF, calendar invite, vendor attachment, maintenance note, or fake advisory contains instructions that manipulate an AI assistant when the tool summarizes documents, drafts work orders, reviews tickets, or retrieves records.
Unsafe output and decision-support failure	An AI tool produces confident but incorrect guidance about compliance language, public notices, cybersecurity triage, maintenance priorities, alarms, chemical dosing context, pump operations, or incident response, and staff rely on the output without sufficient review.
Sensitive data and OT information exposure	Staff or vendors enter customer data, credentials, emergency plans, cybersecurity details, network diagrams, SCADA screenshots, asset inventories, incident records or nonpublic operational information into public or poorly governed AI tools.
Data, model, vendor, and supply-chain compromise	A vendor adds an AI feature to a compliance, billing, asset-management, dashboard, customer-service, or maintenance platform without utility review; alternatively, compromised source data, retrieval systems, AI-generated scripts, or model updates distort outputs or create insecure behavior.
Overreliance and excessive agency	A utility allows an AI agent to send emails, create work orders, update records, route vendor tickets, summarize cybersecurity alerts, or recommend operational actions before a qualified person verifies the result, creating a pathway for errors to become actions.

Together, these classifications show that safe adoption policy should focus on how proposed AI use cases might fail in a water-utility setting. Before adoption, utilities should assess the data involved, the system’s permissions and authority, and its proximity to OT. They should also evaluate vendor dependence, human review requirements, and their ability to continue operating if the tool fails.

Foreseeable AI Benefit Pathways in Water Utility Operations

AI adoption remains relatively low in the water sector, so evidence for generative and agentic AI benefits is limited. Recent work on generative AI in water utility operations and federal AI risk guidance suggests that bounded and governed AI tools can support several tasks. These include drafting, data summarization, information retrieval, training, documentation management, and human-reviewed decision support.²³ Near-term benefits are most likely found in documentation, procurement, maintenance coordination, compliance, and administrative workflows. Realizing these benefits still requires investment. Utilities may need to prepare data, train staff, review vendors, strengthen cybersecurity controls, and maintain ongoing oversight.

Administrative and Compliance Capacity

The first general benefit pathway is administrative and compliance capacity. Many small utilities spend scarce staff time drafting public notices, Consumer Confidence Reports, board materials, grant narratives, permit summaries, meeting notes, and routine customer communications. Generative AI can reduce the time required to create first drafts, translate technical material into plain language, organize source records, and prepare documents for review by a responsible operator, manager, or compliance lead. Commercial tools are already emerging in this area. For example, “1water.ai” describes an AI-supported Consumer Confidence Report platform for U.S. community water systems, and Klir’s “Boots” is marketed as an AI assistant that can answer utility-data questions from centralized records.²⁴

Operational Knowledge and Workforce Support

The second benefit pathway is operational knowledge and workforce support. Small utilities often depend on a small number of experienced operators, contract personnel, vendors, or regional technical-assistance providers. Generative and agentic AI tools can help preserve and organize institutional knowledge. They can summarize standard operating procedures, turn maintenance notes into searchable records, draft training materials, produce shift handoff summaries, and help newer staff understand historical

²³ Autio et al., 2024; Sela et al., 2025

²⁴ 1water.ai, 2026; Klir, n.d.

maintenance or compliance records. These uses align with water-sector research emphasizing the relationship among knowledge management, human resources, and operational capacity²⁵. The outputs are safest if they remain advisory, are checked against source records, and do not replace operator judgment.

Maintenance and Operational Planning Documentation

The third benefit pathway is maintenance and operational-planning documentation. Generative and agentic tools can organize approved maintenance notes, inspection reports, operator logs, service tickets, vendor manuals, and past work orders into summaries, checklists, and questions for staff, engineers, or technical-assistance providers. These tools should not be presented as predictive-maintenance engines or digital twins. For small utilities, their near-term value is lower-risk workflow support. They can help staff find records, prepare draft plans, and document decisions while qualified personnel remain responsible for analysis and action.

Cybersecurity and Resilience Documentation

The fourth benefit pathway is cybersecurity and resilience documentation. Generative and agentic tools may help small utilities summarize logs for qualified review, organize incident records, draft tabletop-exercise materials, identify missing asset-inventory fields, summarize vulnerability information, prepare after-action reports, and convert incident lessons learned into updated procedures. Properly bounded AI may help small utilities perform recurring cybersecurity work more consistently. Final security decisions should remain with responsible staff or managed service providers.²⁶

Table 2 translates these pathways into practical AI benefit classifications for water utilities. The classifications do not claim that every small utility should adopt each type of tool. Rather, they identify foreseeable benefit categories that can be evaluated against the utility's staffing, data, cybersecurity, vendor-management, and fallback capacity.

²⁵ Feliciano et al., 2021.

²⁶ Cybersecurity and Infrastructure Security Agency, Environmental Protection Agency, & Federal Bureau of Investigation, 2024e.

Table 2. AI benefit classifications and water-sector relevance.

<i>Benefit Classification</i>	<i>Water-Sector Example</i>
Administrative and compliance drafting	AI drafts a Consumer Confidence Report explanation, grant narrative, board memo, permit summary, or public notice from approved non-sensitive source records for human review. Commercial examples include 1water.ai for CCR support and Klir's Boots for utility-data questions.
Plain-language customer and public communication	AI converts technical information into plain-language FAQs, customer notices, translation drafts, meeting summaries, or emergency-message templates that are reviewed by utility staff before release.
Workforce training and institutional knowledge	AI summarizes standard operating procedures, maintenance histories, training materials, shift notes, and lessons learned so newer staff can more easily understand routine utility practices.
Utility records search and retrieval	AI helps staff ask natural-language questions about approved records, inspections, work orders, compliance data, asset records, or maintenance notes without manually searching across disconnected files.
Maintenance and operational-planning documentation	AI summarizes approved maintenance notes, inspection reports, operator logs, service tickets, vendor manuals, and past work orders to help staff prepare checklists, agenda items, draft plans, or questions for human review.
Cybersecurity and resilience documentation	AI helps summarize logs for qualified review, draft incident documentation, organize asset inventories, prepare tabletop exercises, identify missing recovery information, and convert after-action findings into updated procedures.

Together, these classifications show why generative and agentic AI should be evaluated as capacity-building tools, not only as sources of risk. *The practical question is whether a proposed use helps the utility perform routine work more consistently without expanding access to sensitive records, obscuring accountability, or moving AI closer to operational control than the utility can safely govern.*

AI and the Sustainability Gap: Incremental Adoption or Unmanaged Drift

AI will not narrow or widen the water utility sustainability gap by itself. The effect depends on how adoption occurs. For small and rural utilities, the safest and most useful path is incremental adoption. It should treat generative and agentic AI as capacity-building tools, not replacements for operators, engineers, cybersecurity support, regulatory judgment, or accountable management. Utilities can begin with narrow, low risk uses which help staff complete work they already struggle to manage. Strategic selection of use cases can increase operational capacity and build the governance needed to manage more sophisticated technologies, including agentic AI that automates whole workflows.

This incremental approach can narrow the sustainability gap because it allows small utilities to build usable administrative and technical capacity without immediately exposing sensitive systems or creating new operational dependencies. A utility that lacks a full-time compliance specialist, planner, communications officer, data manager, or cybersecurity analyst may still use a well-controlled AI tool to prepare draft materials for those functions. The benefit comes from reducing friction in routine work, improving documentation, preserving institutional knowledge, and helping limited staff make better use of outside expertise. These uses are most appropriate when the tool does not receive customer records, credentials, system diagrams, cybersecurity details, or nonpublic operational information. It should also have no access to OT. The output should remain advisory, traceable, and checked against source records before it becomes part of official utility work.

AI widens the gap when adoption occurs through unmanaged drift rather than deliberate capacity building. Drift can occur when staff use public AI tools informally, a vendor adds AI features through a routine software update, or an AI support channel becomes part of process troubleshooting. It can also occur when a tool is adopted before the utility understands its requirements and risks. These problems are especially important for utilities with limited operational capacity because failure further constrains that capacity. For example, an ungoverned reporting tool could create errors that increase regulatory burden instead of building capacity.

The dividing line is not AI versus no AI. It is whether the adoption pathway fits the utility's capacity to govern the use. AI narrows the sustainability gap when adoption starts with low-risk, reviewable uses and gradually builds the policies, contracts, staff habits, records, and fallback procedures needed for more complex tools. AI widens the gap when a tool handles sensitive records, cybersecurity workflows, vendor-access

processes, public communications, or OT-adjacent decisions before the utility can manage the risk. Small utilities should not be pushed toward rapid AI adoption or left to manage it alone. They need practical support that builds capacity first and delays higher-risk uses until governance and safeguards are in place.

Policy Implications for Safe AI Adoption

The policy goal should be to make safe, incremental AI adoption easier than informal or unmanaged adoption. Small and rural utilities should not be pushed toward advanced automation before they can govern it safely. At the same time, they should not be left to encounter AI one public chatbot, vendor feature, software update, or support tool at a time. Policy should provide a deliberate adoption path that begins with bounded, reviewable uses and postpones higher-risk applications until appropriate governance, cybersecurity, vendor oversight, and fallback capacity are in place.

Policy guidance should direct utilities to evaluate proposed AI uses by function and risk, not by product label. They should first identify what data the system uses and how sensitive it is. They should then assess the system's permissions, authority to act, and proximity to OT. Finally, they should determine whether the utility can oversee the vendor, provide adequate human review, meet records requirements, and continue operating if the tool fails.

Low risk uses may be appropriate when they involve non-sensitive information, advisory outputs, and meaningful human review. Higher risk uses require stronger controls and more formal oversight. Applications that could affect OT control, chemical dosing, pump operations, alarm response, cybersecurity actions, or vendor remote access should be deferred unless they have been validated and governed. They also require logging and tested manual procedures.

Public support should initially emphasize capacity-building uses rather than advanced automation. Appropriate early applications include drafting and summarizing documents, training support, data organization, and organizational planning that do not require sensitive data, system access, or operational authority. These uses can provide practical value while helping utilities build the governance habits needed for more complex applications. They also reduce the likelihood that utilities learn AI governance through unmanaged experimentation.

State agencies, associations, technical-assistance providers, federal partners, and vendors should treat safe AI adoption as part of water-sector resilience policy. Their role

is to create shared support structures that capacity-limited utilities can use in practice. These structures should include plain-language guidance, template procedures, training, procurement support, technical assistance, cybersecurity exercises, and access to qualified review.

Vendor governance should be a central part of that support. Many utilities will encounter AI through billing systems, compliance platforms, asset-management tools, cybersecurity products, customer-service platforms, or routine software updates. Before those features become operationally important, vendors should explain what the AI does and what data it uses. They should also describe data retention, activity logging, and access controls. Vendors should disclose material model or feature changes, provide practical disablement options, and explain how the technology connects to other systems. They should identify whether failures, updates, or automated actions could affect related workflows, data, or operational functions.

A Texas Perspective: Rural Water Infrastructure, AI Governance, And Regional Resilience

Texas illustrates how geography, dispersed infrastructure, uneven capacity, drought pressure, local dependence on small utilities, and broad resilience needs can turn a local cyber or AI-governance failure into a regional concern. It also shows why the issue belongs in homeland-security and resilience discussions, not only technology governance. The state combines large geography, population growth, drought pressure, agricultural dependence, and many small or dispersed communities. The Texas Water Development Board adopted the 2022 State Water Plan after projecting substantial long-term growth and future water-supply needs. The Texas Comptroller similarly framed the plan as central to statewide economic and community resilience.²⁷

Rural water and wastewater resilience in Texas affects public health, emergency response, agriculture, local employers, environmental compliance, and community trust. The April 2024 South Plains incidents, including the Muleshoe water overflow event, show how even limited disruptions can become regional confidence and emergency-management concerns.²⁸

Texas also shows how the “rural-does-not-always-mean-insecure” qualification can guide support. Some communities have strong operators, mutual-aid ties, and

²⁷ Texas Comptroller of Public Accounts, 2022; Texas Water Development Board, 2021.

²⁸ Associated Press, 2024; KCBD, 2024.

responsive vendors. Others need regional assistance. The Texas Water/Wastewater Agency Response Network (TXWARN) already supports statewide preparedness, response, and mutual aid. State operator-support structures could extend that model to AI and cybersecurity training and tabletop exercises.²⁹

Table 3. Texas resilience impacts and practical support models.

<i>Resilience Concern</i>	<i>Potential Consequence</i>	<i>AI / Cyber Relevance</i>	<i>Practical Texas Support Model</i>
Public health	Unsafe water, service disruption, or public-notice confusion.	AI-generated notices or summaries may be inaccurate during incidents.	Templates, review procedures, and operator training.
Emergency response	Reduced fire suppression, hospital support, or shelter operations.	AI-assisted incident support may misread data or omit manual fallback needs.	Regional tabletop exercises and mutual-aid integration.
Agriculture and local economy	Livestock, processing, irrigation, and small industry disruptions.	Vendor or remote-access compromise can affect local continuity.	Shared managed security services and vetted vendor requirements.
Public trust	Rumor, misinformation, or loss of confidence after a cyber incident.	AI-generated impersonation or confusing messages can worsen trust.	Public-communication playbooks and callback verification.
Uneven capacity	Some utilities can govern AI; others need practical assistance.	One-size-fits-all AI policy misses local conditions.	Tiered AI policy templates and association-led training.

²⁹ Texas Water/Wastewater Agency Response Network, n.d.; Texas Commission on Environmental Quality, 2026.

PRACTICAL APPLICABILITY / RECOMMENDATIONS

For an individual utility, the first practical step is to create an inventory of use cases before formally or informally deploying a generative or agentic AI tool. The inventory should cover public tools, informal staff uses, vendor-provided features, AI support channels, and AI functions embedded in existing software. Its purpose is to identify where AI might safely build administrative, cybersecurity, and governance capacity before it affects sensitive records, cybersecurity workflows, public communications, vendor access, or operational technology.

The inventory should begin with a first-pass risk classification. The utility should ask simple questions before placing each potential use case into the framework in Table 4. What task would the tool support? What capacity problem would it help solve? What data would it need? Who would review the output? What could go wrong if the output was incomplete, misleading, unavailable, or compromised? All candidate uses should meet basic cybersecurity hygiene requirements. Public or poorly governed AI tools should not receive login credentials, customer records, security details, emergency plans, or nonpublic operational information. Early candidates should have clear action boundaries, human-reviewable outputs, and no access to critical systems, OT environments, security diagrams, or sensitive utility data.

Table 4. Use-case classification framework for generative and agentic AI.

<i>Risk Level</i>	<i>Example Uses</i>	<i>Required Controls</i>	<i>Stop Unless...</i>
Low	Drafting public notices; summarizing non-sensitive policies; training outlines; board materials.	Human review; no sensitive data entry; no reliance on legal, safety, or operational decisions.	The tool is not exposed to customer, security, credential, emergency, or OT data.
Moderate	Grant writing; compliance drafts; maintenance summaries; procurement drafts.	Human review; approved data sources; recordkeeping; staff training.	A person responsible can verify source material and final content.
High	Cybersecurity log summaries; vendor-access support; network documentation; incident communications.	Access limits; output logging; expert review; vendor controls; testing outside production.	Security staff or qualified support can review outputs and permissions.
Critical	Pump operations; chemical dosing; alarm response; PLC/HMI recommendations; automated OT actions.	Formal validation; human approval; OT isolation; change management; manual fallback.	The utility can prove safe operation, logging, disablement, and recovery.

After classifying a use case, the utility should decide whether it is ready for adoption. A simple readiness test is whether the proposed AI use can be explained in one page. The utility should be able to identify who is responsible for the tool, what data it uses, what task it supports, and who reviews the output. It should also state what records are kept and what fallback plan will be followed if the tool fails. If these questions cannot be answered, the use case is not ready for safe adoption or deployment.

As the utility collects and prioritizes its AI use cases, it will need a simple way to answer readiness questions consistently. That is the essence of an AI governance system. A small utility does not need a sophisticated AI governance program on day one. It needs a starter package that makes safe adoption repeatable. At minimum, that package should help the utility identify approved uses, define permitted data, assign human review, evaluate vendor conditions, preserve OT-control boundaries, and document fallback procedures. Table 5 identifies the core support roles and practical outputs that

can help small utilities build that starter package without having to design it from scratch.

Table 5. First steps for supporting safe AI governance in small utilities.

<i>Stakeholder</i>	<i>First Practical Action</i>	<i>Purpose</i>	<i>Output</i>
State agencies and water associations	Provide a starter AI governance and adoption package.	Give small utilities a practical way to identify, screen, and govern AI uses.	Model AI-use policy; use-case inventory template; risk worksheet; training module; tabletop scenario.
State agencies, associations, or technical-assistance providers	Provide model rules for safe AI use.	Help utilities set minimum boundaries before AI tools become routine.	Data-entry rule identifying prohibited inputs, including customer, credential, cybersecurity, OT, emergency, and vendor-access information.
State agencies, associations, or technical-assistance providers	Provide model review and accountability procedures.	Preserve human responsibility for safety, compliance, cybersecurity, public communication, and operations-adjacent decisions.	Human-review rule for public communications, compliance submissions, cybersecurity decisions, procurement, vendor access, and operations-adjacent work.
State agencies, procurement offices, and associations	Provide vendor-review language for AI-enabled tools.	Help utilities evaluate data use, retention, access, logging, feature changes, and disablement before purchase, renewal, or activation.	Vendor-review checklist; contract language; data-use summary; retention terms; logging description; disablement procedure.
Federal partners	Fund AI governance startup support through rural water cybersecurity and resilience programs.	Scale training, technical assistance, exercises, and expert review.	Grant language; technical-assistance guidance; exercise scenarios; shared support models.
Vendors	Provide plain-language AI deployment documentation before purchase, activation, renewal, or feature update.	Prevent hidden AI features, unclear data practices, and unsafe defaults.	AI feature description; data-use terms; retention terms; access-control documentation; logging description; disablement procedure.

<i>Stakeholder</i>	<i>First Practical Action</i>	<i>Purpose</i>	<i>Output</i>
Cybersecurity partners	Review AI tools that touch logs, network documentation, remote access, incident response, or other security workflows.	Prevent new security dependencies that exceed utility capacity.	Security review notes; access recommendations; logging requirements; escalation process.
Utility managers, operators, and support partners	Document OT-control and fallback boundaries for approved use cases.	Keep AI away from direct OT control unless rigorously validated and preserve service continuity if the tool fails.	OT-control policy; manual-fallback plan; one-page deployment record for each approved use case.

CONCLUSION

Generative and agentic AI will enter a water sector already marked by uneven operational capacity. Larger utilities are more likely to have the staff, expertise, cybersecurity maturity, and vendor-management capacity to adopt new tools safely. Many small and rural utilities face similar operational, compliance, communication, cybersecurity, and recovery demands with fewer resources. AI can therefore multiply both capacity and risk. Its effect on the sustainability gap will depend on whether adoption reduces existing constraints or adds new ones.

AI can help narrow the gap when adoption is incremental, bounded, and governed. Low risk uses can reduce staff burden without giving AI control over essential services or operational decisions. They can support drafting, records organization, grant applications, compliance summaries, training materials, and routine documentation. These uses can also build the governance practices needed for safer adoption. Each successful use case can strengthen the utility's ability to evaluate and manage the next one.

AI can widen the gap when adoption occurs through unmanaged drift. Public tools, vendor updates, cybersecurity products, asset-management systems, or operator-facing decision aids may become important before the utility has reviewed their risks. The utility may not know what data the tool receives, what decisions it influences, who reviews its outputs, or how work continues if it fails. The dividing line is deliberate adoption versus unmanaged adoption.

Small utilities can begin without a sophisticated AI governance program. They can identify current and proposed uses, classify risk, and prioritize the safest high-value candidates. For each selected use case, they should limit sensitive data and operational exposure, preserve human accountability, review vendor terms, and maintain fallback capacity. Supporting organizations should provide templates, procurement language, training, cybersecurity exercises, and technical assistance suited to small and rural systems.

For Texas and similar states, safe AI adoption is also a resilience issue. Small water-system disruptions can affect public health, emergency response, agriculture, local economies, environmental compliance, and community confidence. AI supports resilience when it helps utilities plan, communicate, comply, maintain systems, manage vendors, and recover from incidents. Policy should make safe adoption easier than unmanaged drift. With practical support, AI can narrow the sustainability gap and strengthen resilience across dispersed water systems. Without it, AI may amplify the existing capacity divide.

DISCLOSURE STATEMENT

The author reports no competing interests to declare.

FUNDING / GRANT INFORMATION

This work was funded by the Institute for Homeland Security at Sam Houston State University.

AUTHOR BIOGRAPHY

Jason M. Nichols, PhD, is CEO and Managing Partner of HydrospherIQ Consulting, where he advises water utilities and technology companies on AI product strategy, evaluation, and responsible deployment. He has more than 15 years of experience developing AI-enabled water-treatment solutions. As Director of Product at Pani Energy, he led optimization and expert-support software for reverse-osmosis and biological treatment plants; the platform received the 2023 Global Water Summit Breakthrough Technology of the Year award. Previously, at SUEZ Water Technologies & Solutions and GE Global Research, he led teams developing water analytics, expert systems, digital wastewater-treatment platforms, and model-predictive-control applications. He holds a PhD in organic chemistry from the University of Maryland, College Park.

REFERENCES

- 1water.ai. (2026). CCR Agent: Consumer Confidence Report software for water utilities. <https://1water.ai/product/ccr-agent>
- American Water Works Association. (2025). State of the water industry: Concerns over infrastructure financing, public trust, and cybersecurity. *Journal AWWA*, 117(6), 66-69. <https://doi.org/10.1002/awwa.2455>
- Associated Press. (2023, October 17). What's changed (and what hasn't) a year after Mississippi capital's water crisis? <https://apnews.com/article/jackson-water-crisis-federal-intervention-173ae30f9e087ace8fb45c7eac3d3924>
- Associated Press. (2024, April 18). Rural Texas towns report cyberattacks that caused one water system to overflow. <https://apnews.com/article/texas-muleshoe-water-systems-cyberattacks-russia-5f388bf0d581fc8eb94b1190a7f29c3a>
- Australian Signals Directorate's Australian Cyber Security Centre, Cybersecurity and Infrastructure Security Agency, National Security Agency, Federal Bureau of Investigation, Multi-State Information Sharing and Analysis Center, United Kingdom National Cyber Security Centre, Canadian Centre for Cyber Security, New Zealand National Cyber Security Centre, Bundesamt für Sicherheit in der Informationstechnik, Netherlands National Cyber Security Centre, Japan National Center of Incident Readiness and Strategy for Cybersecurity, Japan National Police Agency, Republic of Korea National Intelligence Service, & Republic of Korea National Cyber Security Center. (2024). Principles of operational technology cyber security. <https://www.cyber.gov.au/resources-business-and-government/governance-and-user-education/governance/principles-operational-technology-cyber-security>
- Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.600-1>
- Cybersecurity and Infrastructure Security Agency, Environmental Protection Agency, & Federal Bureau of Investigation. (2024e, February 21). Top cyber actions for securing water systems. <https://www.cisa.gov/resources-tools/resources/top-cyber-actions-securing-water-systems>

Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation, National Security Agency, Environmental Protection Agency, Department of Energy, Department of Agriculture, Food and Drug Administration, Multi-State Information Sharing and Analysis Center, Canadian Centre for Cyber Security, & United Kingdom National Cyber Security Centre. (2024a, May 1). Defending OT operations against ongoing pro-Russia hacktivist activity. <https://www.cisa.gov/resources-tools/resources/defending-ot-operations-against-ongoing-pro-russia-hacktivist-activity>

Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation, National Security Agency, Environmental Protection Agency, Israel National Cyber Directorate, Canadian Centre for Cyber Security, & United Kingdom National Cyber Security Centre. (2024c, December 18). IRGC-affiliated cyber actors exploit PLCs in multiple sectors, including U.S. water and wastewater systems facilities. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-335a>

Federal Bureau of Investigation, Internet Crime Complaint Center. (2024, December 3). Criminals use generative artificial intelligence to facilitate financial fraud (Alert No. I-120324-PSA). <https://www.ic3.gov/PSA/2024/PSA241203>

Feliciano, J. F., Arsénio, A. M., Cassidy, J., Santos, A. R., & Ganhão, A. (2021). Knowledge management and operational capacity in water utilities, a balance between human resources and digital maturity: The case of AGS. *Water*, 13(22), 3159. <https://doi.org/10.3390/w13223159>

Government Accountability Office. (2024). Critical infrastructure protection: EPA urgently needs a strategy to address cybersecurity risks to water and wastewater systems (GAO-24-106744). <https://www.gao.gov/products/gao-24-106744>

Kane, J. W., & Tomer, A. (2018). *Renewing the water workforce: Improving water infrastructure and creating a pipeline to opportunity*. Brookings Institution. <https://www.brookings.edu/articles/renewing-the-water-workforce-improving-water-infrastructure-and-creating-a-pipeline-to-opportunity/>

KCBD. (2024, April 18). Lubbock software company explains Russian South Plains cyberattack. <https://www.kcbd.com/2024/04/18/lubbock-software-company-explains-russian-south-plains-cyberattack/>

Klir. (n.d.). Boots: AI-powered insights for water utility data. <https://www.klir.com/boots>

- National Cyber Security Centre. (2024, January 24). The near-term impact of AI on the cyber threat. <https://www.ncsc.gov.uk/report/impact-of-ai-on-cyber-threat>
- National Research Council. (1997). Safe water from every tap: Improving water service to small communities. National Academies Press. <https://doi.org/10.17226/5291>
- Open Worldwide Application Security Project. (2025). OWASP Top 10 for large language model applications 2025. <https://owasp.org/www-project-top-10-for-large-language-model-applications/>
- Sela, L., Sowby, R. B., Salomons, E., & Housh, M. (2025). Making waves: The potential of generative AI in water utility operations. *Water Research*, 272, 122935. <https://doi.org/10.1016/j.watres.2024.122935>
- Texas Commission on Environmental Quality. (2026). Chapter 2: Agency activities. In Biennial report to the 89th Legislature. <https://www.tceq.texas.gov/publications/sfr/tceq-biennial-report/biennial-report-to-the-89th-legislature/agency-activities>
- Texas Comptroller of Public Accounts. (2022, July). The 2022 State Water Plan. Fiscal Notes. <https://comptroller.texas.gov/economy/fiscal-notes/archive/2022/jul/water-plan.php>
- Texas Water Development Board. (2021, July 7). Texas Water Development Board adopts 2022 State Water Plan. https://texaswaternewsroom.org/pressreleases/2021-07-07_swp_adoption.html
- Texas Water/Wastewater Agency Response Network. (n.d.). Texas Water/Wastewater Agency Response Network. <https://www.txwarn.org/>
- U.S. Environmental Protection Agency. (2024a). Interagency Water Workforce Working Group report to Congress. <https://www.epa.gov/waterutilityresponse/interagency-water-workforce-working-group-report-congress>
- U.S. Environmental Protection Agency. (2025). 21st annual EPA drinking water workshop: Small system challenges and solutions. <https://www.epa.gov/water-research/21st-annual-epa-drinking-water-workshop-small-system-challenges-and-solutions>
- U.S. Environmental Protection Agency. (n.d.). Small and rural wastewater systems. <https://www.epa.gov/small-and-rural-wastewater-systems>

The Institute for Homeland Security at Sam Houston State University is focused on building strategic partnerships between public and private organizations through education and applied research ventures in the critical infrastructure sectors of Transportation, Energy, Chemical, Water/Wastewater, Healthcare, and Public Health.

The Institute is a center for strategic thought with the goal of contributing to the security, resilience, and business continuity of these sectors from a Texas Homeland Security perspective. This is accomplished by facilitating collaboration activities, offering education programs, and conducting research to enhance the skills of practitioners specific to natural and human caused Homeland Security events.

[Institute for Homeland Security](#)

[Sam Houston State University](#)

© 2026 The Sam Houston State University Institute for Homeland Security.

Nichols, J. (2026). AI Technologies and the Water Utility Sustainability Gap Helping Rural Utilities Build Capacity Without Deepening Cyber-Operational Risk. (Report No. 2026 -1037). The Sam Houston State University Institute for Homeland Security. <https://doi.org/10.17605/OSF.IO/RQ4ZY>



**INSTITUTE FOR
HOMELAND SECURITY**
SAM HOUSTON STATE UNIVERSITY

