DESIGNING DECISION SYSTEMS FOR SMART CITIES

FROM SIGNALS TO DECISIONS: A CTO FRAMEWORK FOR OPERATIONAL AI

Sponsored by HPE

Case Study: Town of Vail

Author Keith Townsend, The CTO Advisor

www.thectoadvisor.com/blog/hpesmartcities

a00156405enw

Case Study: Town of Vail

INDEX

Executive Summary

Smart city initiatives have stalled not because of a lack of technology, but because cities have struggled to translate infrastructure investments into trusted, repeatable decision-making systems.

Over the past decade, municipalities have deployed sensors, cameras, connectivity, and cloud platforms at scale. Yet the operational impact has been uneven. Many cities now find themselves in a familiar position: sensors everywhere, decisions nowhere. Data is abundant, dashboards are sophisticated, and pilots demonstrate promise—but real authority remains gated by manual review, fragmented workflows, and institutional risk aversion.

> *The result is not autonomous cities, but decisive cities where AI shortens the time between observation and action without removing accountability from human leaders.*

This whitepaper introduces a CTO-grade framework for smart cities grounded in the 4 + 1 AI Infrastructure Model, designed to help public-sector leaders move from infrastructure enablement to operational intelligence.

Using the Town of Vail as a real-world case study, the paper examines how cities can architect AI systems that support faster, more consistent decisions while preserving human oversight, public accountability, and trust. Rather than treating AI as a standalone analytics capability, the town is approaching AI as part of a control system—embedded into existing operations, constrained by governance, and aligned with how city decisions are actually made.

At the center of this approach is the HPE Unleash AI ecosystem, not as a product portfolio, but as an operational mechanism. Unleash AI enables cities to bring AI capabilities closer to where signals originate and decisions occur, while maintaining centralized governance, economic sustainability, and architectural coherence across departments.

The result is not autonomous cities, but decisive cities—where AI shortens the time between observation and action without removing accountability from human leaders.

To support this transition, this paper introduces a diagnostic lens for identifying where smart city initiatives stall—not due to ambition or funding, but because architectural design constrains how decisions move across systems and departments. By examining progress through the lens of decision flow rather than deployment maturity, city leaders can focus investment where it restores authority, trust, and operational coherence.

This paper is written for city CTOs, CIOs, CISOs, and executive leaders who are no longer asking whether AI belongs in municipal operations, but how to deploy it responsibly, at scale, and with authority.

1. The Smart City Reality Check

Why Infrastructure Alone Has Not Delivered Intelligence

The first generation of smart city initiatives focused on visibility. Cities invested in connectivity, sensors, and data platforms with the expectation that insight would naturally translate into action.

In practice, most initiatives plateaued.

Common symptoms include:

- Analytics that inform but do not decide
- Pilots that succeed in isolation but fail to scale
- AI systems that remain advisory-only
- Decisions that still require manual coordination across departments

The problem is not data quality or model accuracy. It is decision architecture.

Cities are complex socio-technical systems. Authority is distributed, risk is political, and accountability extends beyond IT into public trust. Any AI system that ignores these realities will remain constrained, regardless of technical sophistication.

2. Reframing Smart Cities as Control Systems

From Insight Generation to Decision Acceleration

To move forward, cities must stop treating AI as an analytics upgrade and start treating it as part of an operational control system.

A control system does four things:

- Observes signals
- Interprets context
- Recommends or initiates action
- Preserves accountability

Most smart city platforms stop at step two.

The goal is not to remove humans from the loop, but to design the loop.

Common questions include:

- What decisions can be accelerated?
- Which decisions require human confirmation?
- Where does trust break down—and why?

Answering these questions requires an architectural model that aligns technology with governance. That model is the 4 + 1 AI Infrastructure Framework.

3. The Decision-Centered AI Engagement Method (Core Contribution)

How the Town of Vail Is Approaching AI Differently

To launch its approach, the Town of Vail surfaced a consistent set of decision-design principles, allowing the town to move faster than typical municipal AI programs without triggering organizational or political resistance. These principles were not applied as a formal methodology at the outset; rather, they emerged and continue to evolve through practice into what is now described as the Decision-Centered AI Engagement Method.

A core analytical tool used to explain and repeat these principles is what CTO Advisor refers to as the Decision Authority Placement Model (DAPM), a framework to systematically determine where decision authority should reside across human and AI systems. In this paper, DAPM is applied as an explanatory and forward-looking model—illustrated by the Town of Vail's engagement—to make the method repeatable for other cities and enterprises.

This section codifies those principles so they can be evaluated, repeated, and demanded elsewhere.

CTO Advisor Observation:
Why Most AI Programs Start in the Wrong Place

In nearly every enterprise and public-sector AI program I review, teams start by asking "What models should we deploy?" or "What data do we have?"

That framing is backwards. The programs that stall are the ones that never explicitly answer a more basic question: "Which decisions are we trying to accelerate, and who owns them today?"

The Town of Vail moved faster because it treated decision friction—not data availability or model selection—as the primary design input.

3.1 Start with Decision Friction, Not Use Cases

DAPM implication: Begin by explicitly mapping decision authority: where a decision is made today, how often it occurs, and what happens when it is delayed or inconsistent.

Town of Vail application: In the town, the engagement started by identifying cross-department decisions that repeatedly stalled due to unclear ownership or serial escalation. This allows the team to focus AI support on reducing coordination friction rather than creating new tools.

3.2 Treat Decision Flow as the Primary Design Surface

DAPM implication: DAPM distinguishes between decision support (AI informs), decision acceleration (AI reduces latency), and decision handoff (AI routes or structures escalation while preserving authority).

Town of Vail application: Rather than embedding AI inside departmental systems, the town is designing around decision handoff points between departments. AI is used to accelerate context sharing and routing, while authority and accountability remain clearly human-owned.

3.3 Design AI as a Shared City Capability from Day One

DAPM implication: Decisions that affect multiple departments require shared authority context and consistent governance.

Town of Vail application: Using DAPM, the town identifies decisions that span public safety, transportation, and facilities. AI capabilities are therefore designed as a shared city resource, coordinated through HPE Unleash AI member Kamiwaza's orchestration layer, rather than duplicated across departments. Kamiwaza is a Colorado-based company that specializes in secure AI Orchestration.

3.4 Target Mundane, Cross-Department Decisions First

DAPM implication: Low-risk, high-frequency decisions offer the best opportunity to place AI assistance without threatening authority or triggering political resistance.

Town of Vail application: The engagement deliberately focuses on repetitive coordination and routing decisions—areas where reduced latency and improved context have immediate value. This sequencing allows trust to accrue incrementally before expanding AI involvement into more sensitive domains.

3.5 Orchestrate an Ecosystem, Not a Product Stack

DAPM implication: Placing decision authority correctly often requires multiple specialized capabilities—governance, data locality, reasoning, and orchestration—working together.

Town of Vail application: HPE's Unleash AI program provides the ecosystem framework, while Kamiwaza supplies agentic orchestration—defined here as the coordination of multiple specialized AI agents and human workflows to execute multi-step decisions under explicit authority, governance, and audit constraints.

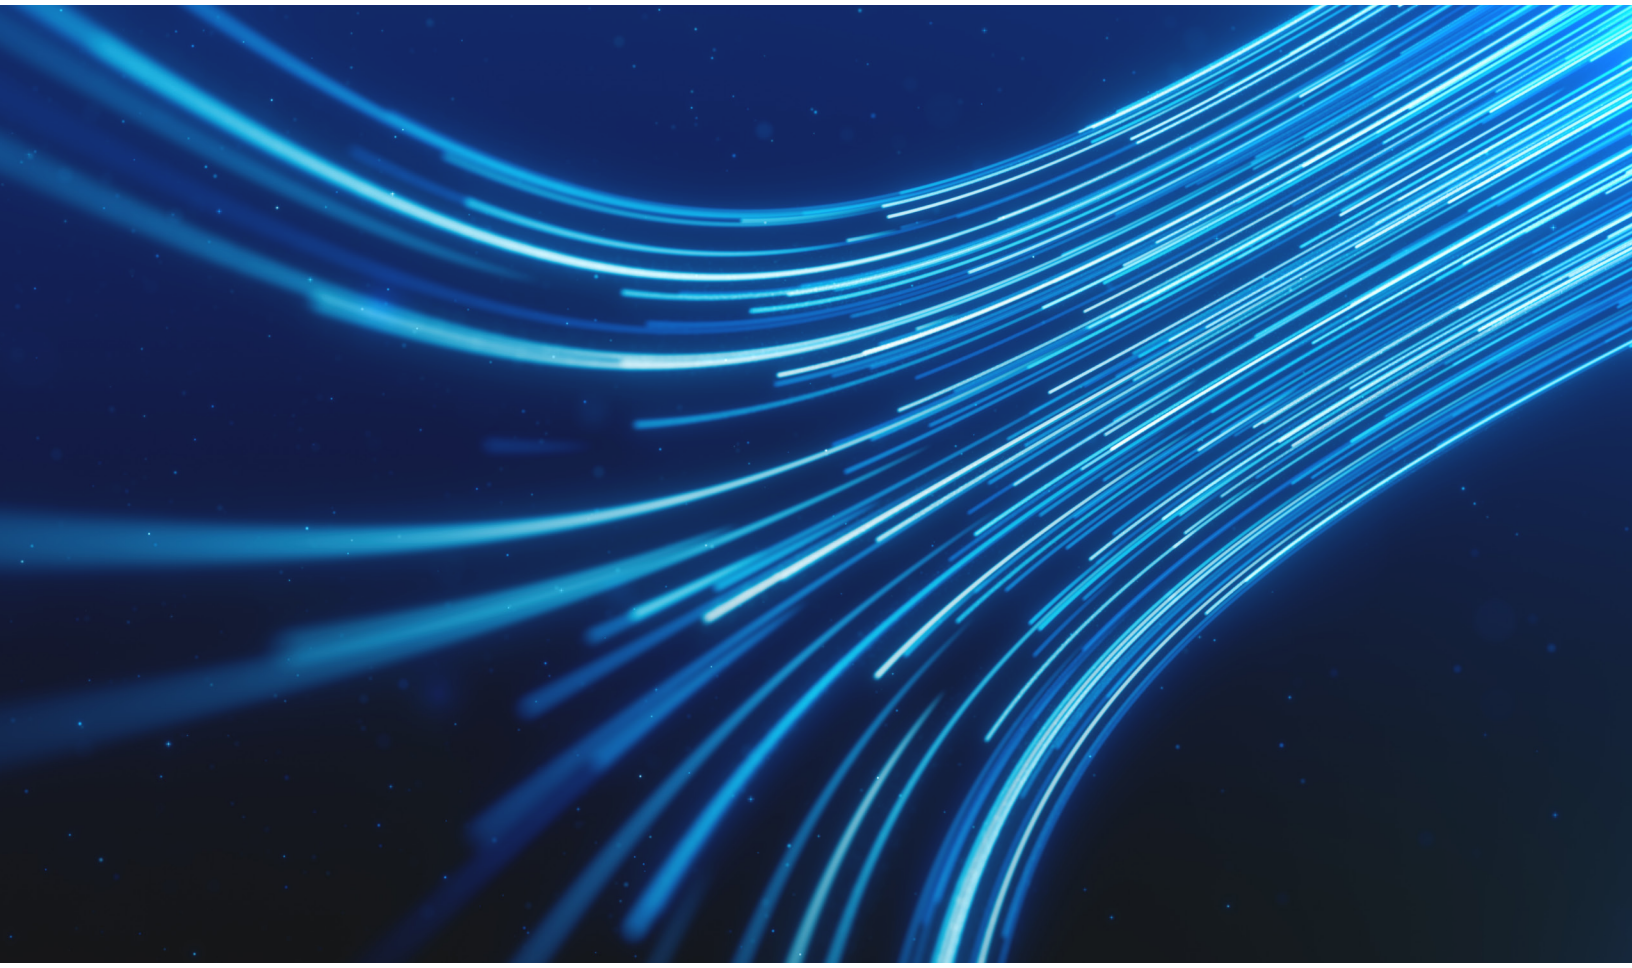

4. Case Study: The Town of Vail

Validating the Decision-Centered Method in Practice

Rather than attempting to "boil the ocean," the Town of Vail engagement deliberately focuses on what participants describe as "boiling mud puddle by mud puddle"—mundane, high-frequency decisions that created real operational drag across departments.

Section 508 / Title II Accessibility Compliance

With the Accessibility Remediation Intelligence Agent (ARIA)—by Kamiwaza, the town is implementing AI-assisted decision flows to identify documents requiring remediation, generate accessibility guidance, and route outputs through human review—accelerating compliance without removing accountability.

Public Safety and Security

Multiple Unleash AI ISV partners are providing solutions to quickly detect hazards, including fire, using existing camera installations.

Housing Deeds Audit and Verification

AI is applied to reduce friction in document ingestion and routing while preserving authority over validity decisions.

CTO Advisor Observation:
Why Starting Mundane Is Strategic

Most smart city initiatives fail because they try to demonstrate value with highly visible, politically sensitive use cases first.

The Town of Vail's approach—starting with mundane, cross-department decisions—reduces operational drag without threatening authority. This is how trust scales.

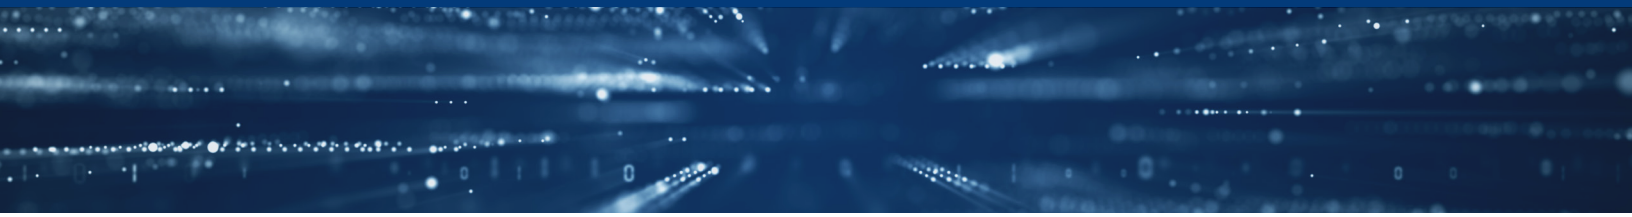

Timeline as Proof of Method

The Town of Vail moved from concept to first phase deployment of the initial use cases in approximately three months—a timeline that is highly atypical for municipal AI initiatives.

Security and Governance: Beyond RBAC

The Town of Vail engagement also surfaced a governance challenge common in agentic systems: traditional role-based access control (RBAC) breaks down when autonomous agents operate across systems and cascade permissions.

To address this, the team is implementing a relationship-based access control (REBAC) approach. Unlike RBAC, which grants static permissions to roles, REBAC evaluates access dynamically based on relationships between actors, data, decision context, and scope of action. This ensures permissions are precisely bounded to the decision at hand and cannot be escalated beyond authorized authority boundaries.

Importantly, this is being deployed as an architectural governance pattern, not as a standalone security product—reinforcing the paper's core thesis that governance must be embedded into decision systems rather than enforced after execution.

Ecosystem-Orchestrated Delivery

The HPE Agentic Smart City Solution, as illustrated through the Town of Vail's ongoing deployment, is explicitly ecosystem-driven.

The solution runs on **HPE Private Cloud AI**, co-developed with NVIDIA, bringing together hardware, software, and services to simplify AI deployment and management. In the Town of Vail's case, this includes deployment on **HPE ProLiant Compute DL380a systems**, which provide a standardized, enterprise-supported platform for running AI workloads on-premises without requiring the town to assemble or operate a bespoke hardware stack. The HPE Unleash AI team delivers expertise on HPE solutions and curated ISV collaboration.

Kamiwaza provides agentic orchestration and decision routing.

ProHawk AI contributes video restoration capabilities designed to preserve evidentiary integrity rather than generate synthetic enhancements.

Vaidio supplies vision AI analytics.

Gambit delivers a personal concierge.

For the Town of Vail, HPE solution provider **SHI** is providing integration and deployment services.

5. Unleash AI as the Mechanism, Not the Message

Enabling Operational AI Without Fragmentation

For the Town of Vail, HPE's Unleash AI program serves as an enabling mechanism for aligning edge, data, AI, governance, and partner capabilities into a coherent operational system. The value is not in individual components, but in coherence—supporting local action while preserving centralized oversight and economic sustainability.

A critical factor in the town's ability to move quickly is the use of an integrated appliance-style delivery model for the AI system as a whole. While the solution incorporates multiple specialized partners, the town does not have to manage these relationships independently. From an operational standpoint, the system is supported and evolving as a single integrated capability.

For CIOs and CTOs, this distinction matters. Ecosystem-based solutions often fail not because the components are inadequate, but because responsibility is fragmented across vendors. By packaging the ecosystem into an appliance-like operating model—with unified deployment, lifecycle management, and support—HPE's Unleash AI program reduces coordination overhead while preserving architectural flexibility.

This approach gives town leaders a familiar operating construct: one accountable provider, predictable operations, and a clear escalation path, even as the underlying system remains modular and extensible.

6. From Pilot to Policy

Scaling Smart City Intelligence

Scaling AI requires moving from isolated success to repeatable patterns: shared platforms, architectural standards, embedded governance, and sustainable operating models.

7. Where Architecture Becomes the Constraint

Diagnosing Failure Modes Through the 4 + 1 AI Infrastructure Model

CTO Advisor Observation:
Why More GPUs Don't Fix Broken Automation

When AI-driven automation produces bad outcomes, the issue is rarely the model or the hardware. It's that the decision was placed in the wrong part of the system.

Teams automate end-to-end workflows only to discover AI is approving exceptions, escalating incidents incorrectly, or triggering downstream actions without the right context. Adding more GPUs just makes the wrong automation run faster.

Layer 0: Physical & Digital Signal Sources
Failure occurs when cities treat signal collection as the objective rather than a means to support decisions.

Layer 1: Data Foundation
Layer 1 becomes the constraint when data cannot move with context across departments.

Layer 2: AI & Orchestration
This is the most common point of failure when reasoning and orchestration are implicit, opaque, or hard-coded applications.

Layer 3: Applications & Workflows
Failure occurs when AI outputs cannot be absorbed into real operational workflows.

+1: Governance, Trust, and Accountability
Without architectural governance, AI initiatives retreat back to pilots.

8. Conclusion

The Town of Vail demonstrates that smart cities advance not by collecting more data, but by designing better decision systems. When authority placement, governance, and architecture align, AI can scale responsibly.

Keith Townsend

Founder & Executive Strategist
The CTO Advisor

Keith Townsend operates at the intersection of enterprise IT strategy and business outcomes. While often described as a CTO, his work extends well beyond traditional organizational roles, encompassing strategy, advisory, and practitioner-led research.

Over the course of his career, Townsend has built and successfully exited a research and media firm, advised Fortune 500 CIOs and CTOs, and led large-scale enterprise infrastructure and cloud transformations across multiple industries. His background spans hands-on systems architecture, cloud and AI strategy, and executive-level communication—enabling him to translate complex technical decisions into clear business impact.

Today, Townsend leads executive programs designed to bridge the gap between enterprise IT decision-makers and the technology vendors seeking to engage them. His work combines practitioner-grounded research, narrative development, and strategic advisory, helping organizations align product strategy, go-to-market execution, and customer outcomes. His focus remains consistent: turning technical capability into durable business value.