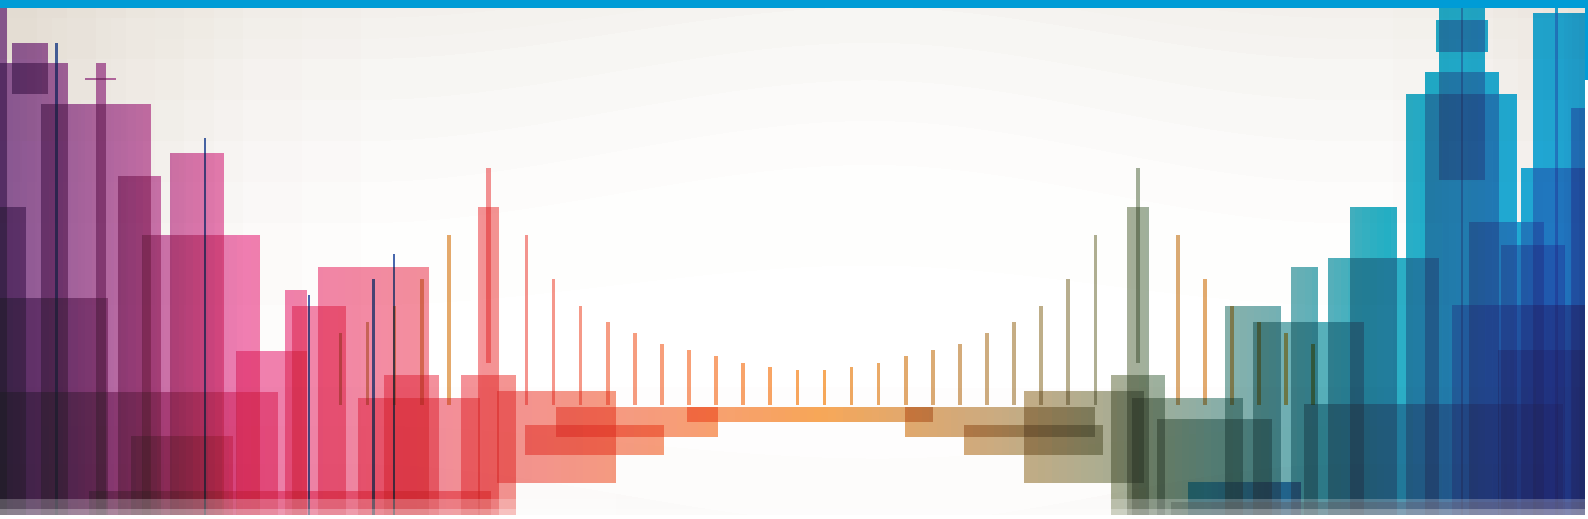
An urban federated agentic AI framework

Enabling responsible federated AI agents for future-ready cities



Convention on Biological Diversity

United Nations ECLAC

Food and Agriculture Organization of the United Nations

ITU

UN-HABITAT

UNDP

United Nations Economic Commission for Africa

UNECE

UNIDO

unesco

UN environment programme

UNEP FINANCE INITIATIVE

United Nations Framework Convention on Climate Change

unitar United Nations Institute for Training and Research

UNFPA

United Nations | Department of Economic and Social Affairs

UNJSPF United Nations Joint Staff Pension Fund

UN OFFICE 4 PARTNERSHIPS

UNU EGOV

UN Tourism

WORLD METEOROLOGICAL ORGANIZATION



An urban federated agentic AI framework

Enabling responsible federated AI agents for future-ready cities

Foreword

This publication was developed within the framework of the United for Smart Sustainable Cities (U4SSC) initiative.

Acknowledgements

The development of this deliverable was led and coordinated by Okan Geray (Digital Dubai, United Arab Emirates).

This publication was authored by Okan Geray (Digital Dubai, United Arab Emirates).

The author would like to thank the following contributing individuals: Adam Obeid (Independent Researcher), Anna Badimo (Sigmateq), Eduardo Alencar (Brazilian Chamber of Deputies), Grace Jimenez (Researcher), Jose Manuel Pérez Marzabal (Member of the Barcelona Bar Association), Kwi-Gon Kim (Seoul National University), Paul Wilson (Smart Cities World), Rajagopal Tampi (AIPathfinder), Shalini S Gopalkrishnan (San Francisco Bay University), Silvia Blanco (SB Analytics), Val Wise (Citiverse Council).

The author wishes to thank the U4SSC management team: Okan Geray (U4SSC Chair), Ramy Ahmed Fathy, Giampiero Bambagioni, Paolo Gemma, Wendy Teresa Goico Campagna, Tania Marcos and Emily Royall (U4SSC Vice-Chairs) for their assistance and contributions.

The author also extend their gratitude to the contributing organizations along with their representatives: Oliver Hillel from the Convention on Biological Diversity (CBD), Lucy Winchester and Vera Kiss from the Economic Commission for Latin America and the Caribbean (ECLAC), Nevena Alexandrova from the Food and Agriculture Organization (FAO), Cristina Bueti from the International Telecommunication Union (ITU), Deniz Susar from United Nations Department of Economic and Social Affairs (UNDESA), Iryna Usava from the United Nations Development Programme (UNDP), James Murombedzi from the United Nations Economic Commission for Africa (UNECA), Dario Liguti and Tea Aulavuo from the United Nations Economic Commission for Europe (UNECE), Guilherme Canela from the Regional Bureau for Sciences in Latin America and the Caribbean of the United Nations Educational, Scientific and Cultural Organization (UNESCO), Gulnara Roll and Mrityunjay Sharma from United Nations Environment Programme (UNEP), Matthew Ulterino from the United Nations Environment Programme Finance Initiative (UNEP-FI), Motsomi Maletjane from the United Nations Framework Convention for Climate Change (UNFCCC), Soren Thomassen from United Nations Population Fund (UNFPA), Aline Machado da Matta and Isabel Wetzel from the United Nations Human Settlements Programme (UN-Habitat), Katarina Barunica Spoljaric and Nicholas Dehod from the United Nations Industrial Development Organization (UNIDO), Ebru Canan-Sokullu from the United Nations Institute for Training and Research (UNITAR), Dino Cataldo Dell'Accio from United Nations Joint Staff Pension Fund (UNJSPF), William Kennedy from the United Nations Office for Partnerships (UNOP), Soumaya Ben Dhaou from the United Nations University – Operating Unit

on Policy-Driven Electronic Governance (UNU-EGOV), Roberta Boscolo from World Meteorological Organization (WMO) and Sandra Carvao from the World Tourism Organization (UN Tourism).

Disclaimer

The opinions expressed in this publication are those of the authors and do not necessarily represent the views of their respective organizations or U4SSC members. In line with the U4SSC principles, this report does not promote the adoption and use of Smart City technology. It advocates for policies encouraging responsible use of information and communications technologies (ICTs) that contribute to the economic, environmental and social sustainability as well as the advancement of the 2030 Agenda for Sustainable Development and the Pact for the Future and its Global Digital Compact.

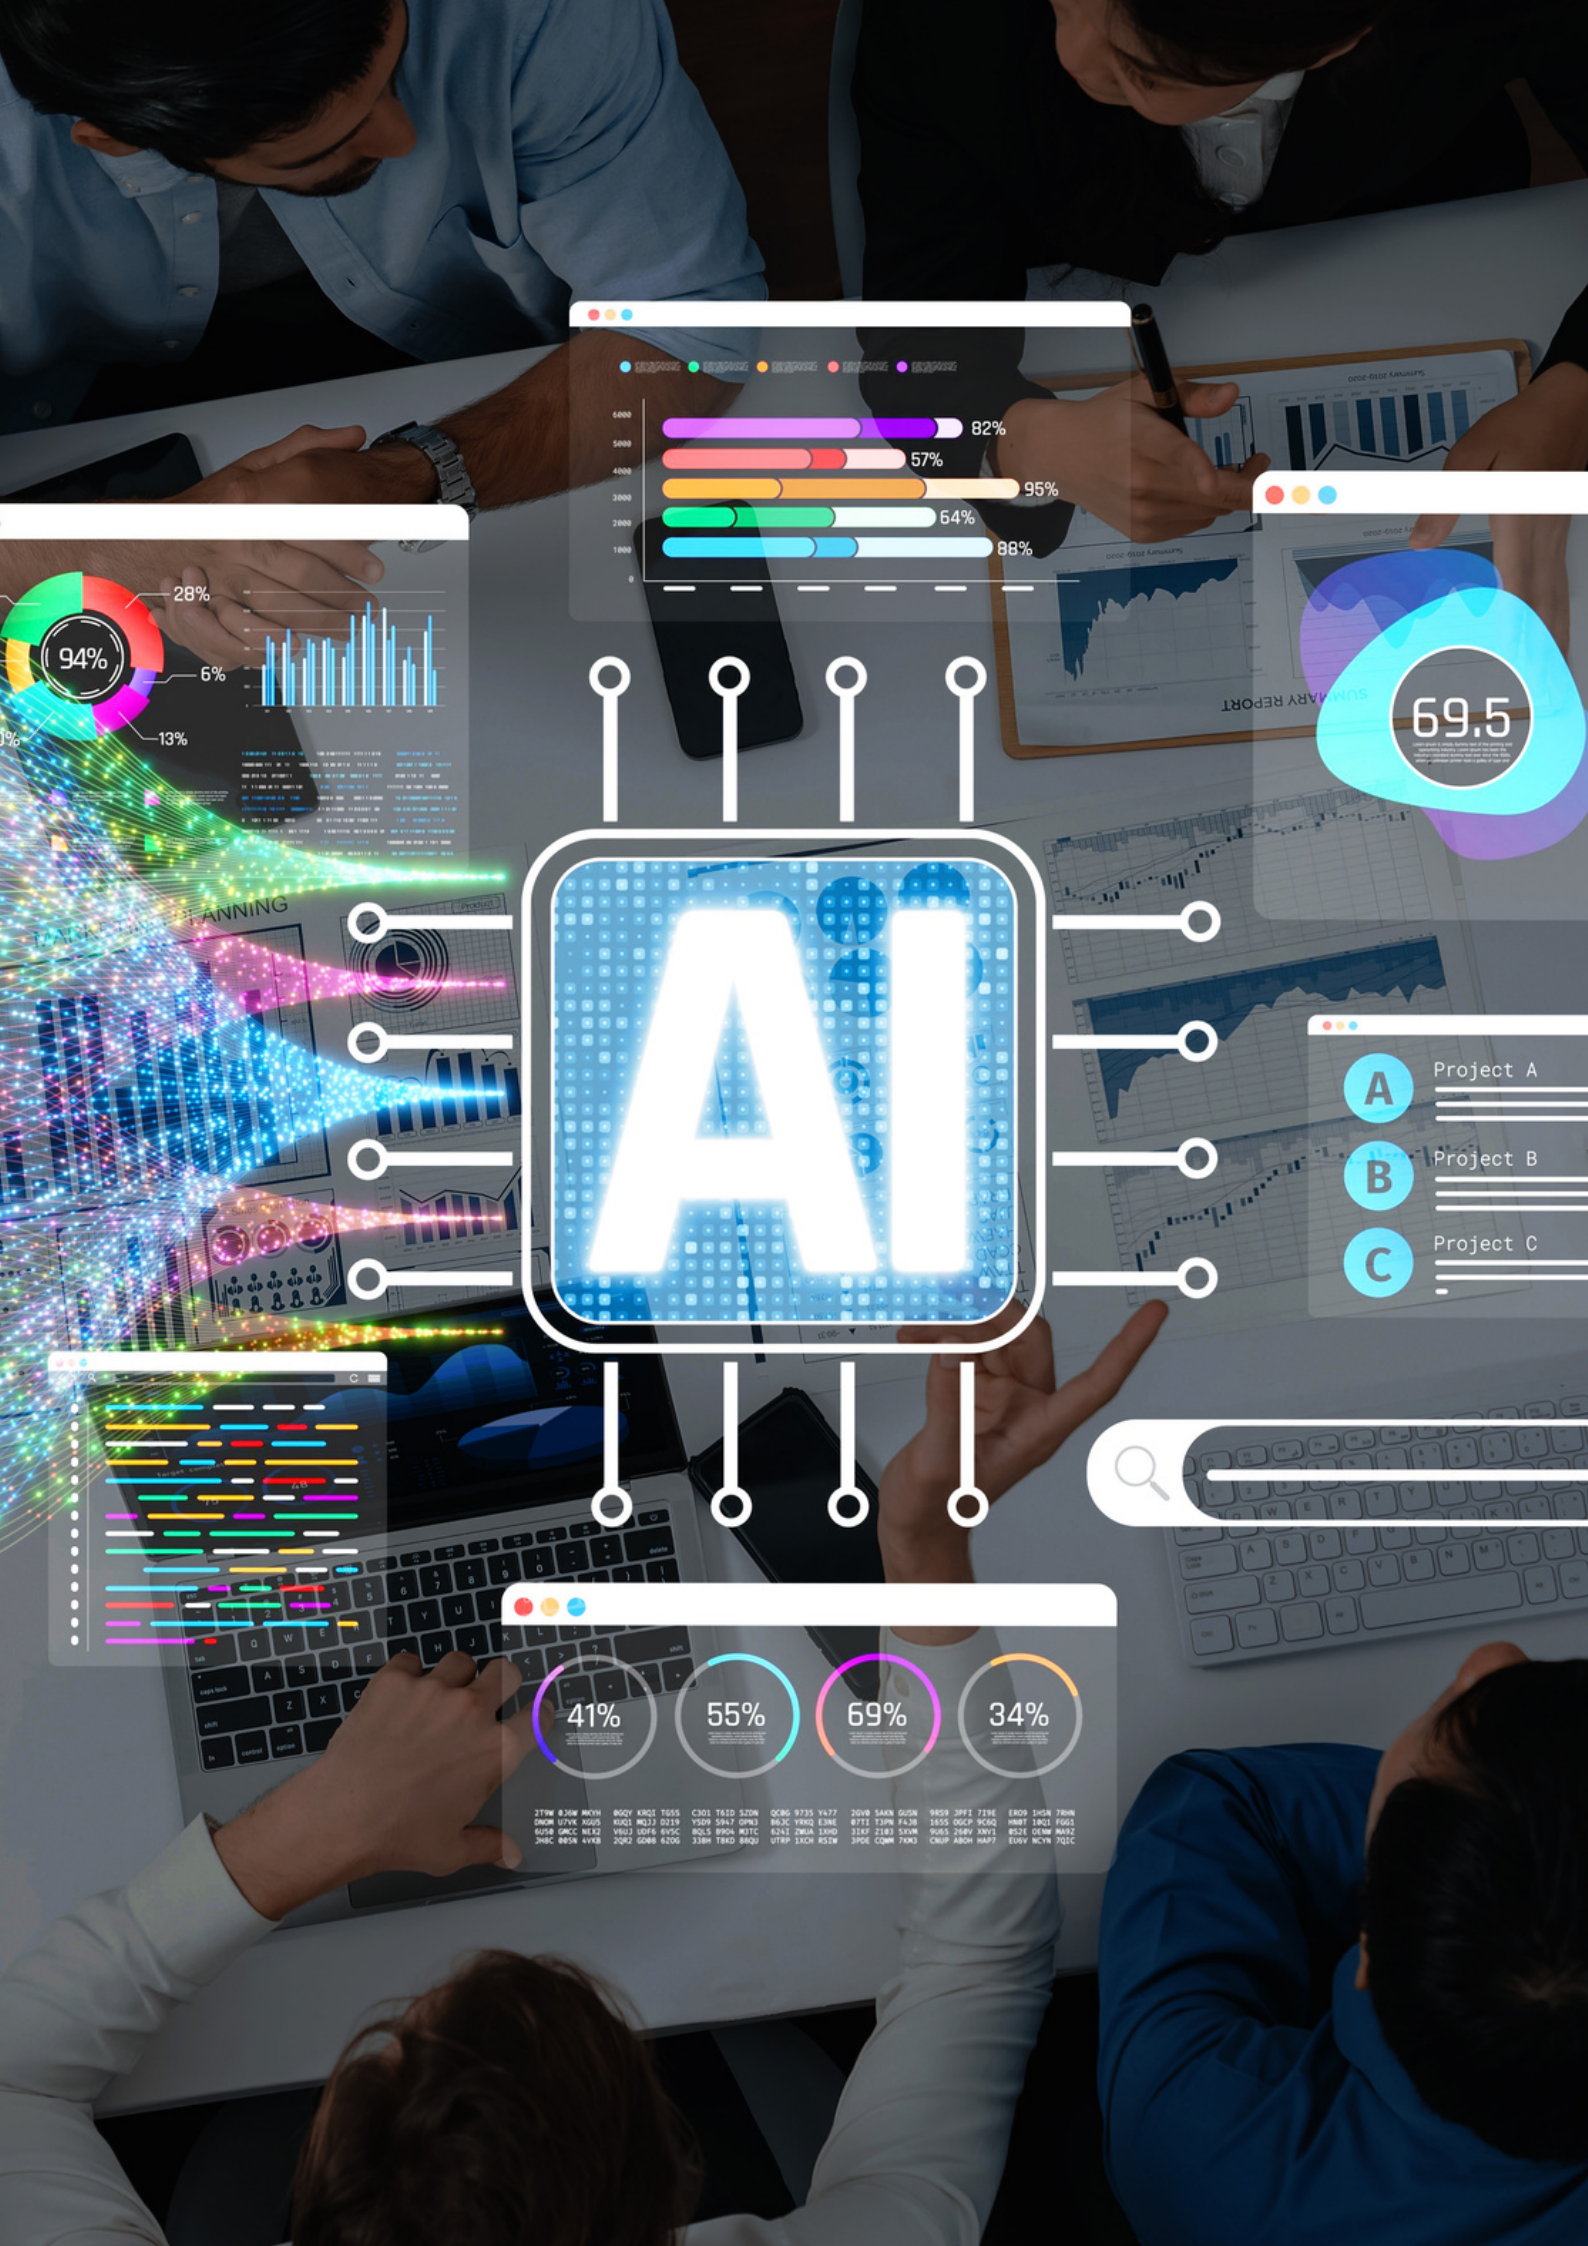
ISBN

978-92-61-43851-7 (electronic version)
978-92-61-43861-6 (EPUB version)

This work is licensed to the public through a Creative Commons Attribution-Non-Commercial-Share Alike 3.0 IGO license (CC BY-NC-SA 3.0 IGO).

For more information, please visit https://creativecommons.org/licenses/by-nc-sa/3.0/igo/

© CBD, ECLAC, FAO, ITU, UNDESA, UNDP, UNECA, UNECE, UNESCO, UNEP, UNEP-FI, UNFCCC, UNFPA, UN-Habitat, UNIDO, UNITAR, UNJSPF, UNOP, UNU-EGOV, WMO and UN Tourism.


82%
57%
95%
64%
88%
28%
94%
6%
13%
69.5
AI
Project A
Project B
Project C
41%
55%
69%
34%

Contents

Abbreviations and acronyms vii

Executive Summary viii

1 Purpose and Overview of the Framework 1

2 The Rise of Agentic AI in Cities 1

3 Why Cities Need AI Agents Now 2

4 Key Benefits of AI Agents to Residents, Businesses, and Government 4

4.1 Benefits for residents 4

4.2 Benefits for businesses 4

4.3 Benefits for government 5

5 What Federated Agentic AI Means for City Governance 5

6 Benefits of Federated Agentic AI vs Standalone Agentic AI in Cities 6

7 A Model Architecture for Urban Federated Agentic AI 9

7.1 Core Design Principles for Federated Agentic AI 9

7.2 Architecture overview 10

7.3 User experience & interaction layer (The Federation Gateway) 12

7.4 Orchestration & choreography layer (The Collective Intelligence) 14

7.5 Messaging & transport layer (The Communication Backbone) 14

7.6 Interoperability & knowledge (ontology) layer 15

7.7 Data & privacy layer 15

7.8 Security & Identity Layer (Trust Infrastructure) 16

7.9 Governance, Observability & Auditing Layer 16

8 Call to Action for Cities on Urban Federated Agentic AI 20

Appendix 1: A Brief Note on Practical Standards for Urban Federated Agentic AI in 2026 23

List of figures

Figures

Figure 1: Federated Agentic AI for Cities ix
Figure 2: Why Cities Need AI Agents Now 3
Figure 3: Key Benefits for Residents, Businesses and Government 4
Figure 4: Standalone AI vs Federated Agentic AI for Cities: From Siloed Systems to Connected, Interoperable and Coordinated City Services 8
Figure 5: A Model Architecture for Urban Federated Agentic AI 12
Figure 6: Implementation Roadmap for Cities 20

Abbreviations and acronyms

Abbreviation	Full form
A2A	Agent-to-Agent
ACL	Agent Communication Language
AI	Artificial intelligence
API	Application programming interface
DCAT	Data Catalog Vocabulary
DID	Decentralized identifier
EMS	Emergency Management System
EPUB	Electronic publication
ETSI	European Telecommunications Standards Institute
FIPA	Foundation for Intelligent Physical Agents
GPS	Global Positioning System
gRPC	gRPC Remote Procedure Calls
HVAC	Heating, ventilation and air conditioning
ICT	Information and communications technology
IGO	Intergovernmental organization
IoT	Internet of Things
ISBN	International Standard Book Number
ISO	International Organization for Standardization
JSON-LD	JavaScript Object Notation for Linked Data
MCP	Model Context Protocol
NGSI-LD	Next Generation Service Interfaces-Linked Data
OAuth	Open Authorization
OGC	Open Geospatial Consortium
OWL	Web Ontology Language
Pub-sub	Publish-subscribe
RDF	Resource Description Framework
SLA	Service-level agreement
SOSA/SSN	Sensor, Observation, Sample and Actuator / Semantic Sensor Network
TLS	Transport Layer Security
U4SSC	United for Smart Sustainable Cities
UUID	Universally unique identifier
UX	User experience
VC	Verifiable credential
W3C	World Wide Web Consortium
ZKP	Zero-knowledge proof

Executive Summary

Cities do not need one giant artificial intelligence (AI) system to become smarter. They need public services that work together. From flood response and emergency transport to business permits and heat-wave protection, residents expect city institutions to coordinate quickly, safely and transparently, even when responsibilities and information are spread across different organizations.

This report presents an Urban Federated Agentic AI Framework: a practical approach for connecting AI agents operated by city departments, utilities, hospitals, transport authorities and other partners. An AI agent is software that can pursue an assigned goal, use approved tools and support or carry out tasks within defined limits. Federation allows these agents to exchange necessary information and coordinate actions under shared rules, while each participating organization retains control of its systems, data and legal responsibilities.

For mayors and policymakers, the opportunity is tangible: faster and more accessible services for residents; simpler administrative journeys for businesses; better coordination of roadworks, mobility and public infrastructure; and more timely responses to floods, heat waves and other emergencies. The framework illustrates these opportunities through cross-agency scenarios rather than claiming that they have already delivered measured results.

The essential policy choice is not how much government can automate, but where coordination creates demonstrable public value. Federated systems require clear mandates, interoperability, lawful and minimal data sharing, verified agent identities, security safeguards, auditable decisions and effective human oversight. Decisions affecting rights, safety and other high-stakes matters must remain subject to appropriate human authority and review. Federation is not a substitute for accountability.

The framework sets out the benefits of federation compared with standalone agents, the principles for responsible cooperation and a layered architecture covering user access, orchestration, messaging, shared meaning, data protection, security and oversight. It is intended to be technology-neutral and adaptable to different levels of city digital maturity.

A practical agenda for city leaders:

- establish common governance before scaling;
- identify two or three cross-agency problems where coordination would make a measurable difference;
- agree participating institutions' roles, permissions and safeguards;
- require open international standards and meaningful public access to human assistance; and
- evaluate pilots against service quality, safety, inclusion, resilience and public trust.

The objective is not autonomous government. It is better-connected, more responsive government that remains answerable to people.

Figure 1: Federated Agentic AI for Cities



The headline: Don't build one giant city brain. Connect the agents your agencies already have.

A flood does not respect departmental boundaries, and neither should the city's response. This framework shows how the AI agents run by transport, utilities, hospitals and emergency services can work as one team in a crisis, without any of them handing its data or its authority to a central system. Each agency stays in charge of its own agent; the city sets the shared rules they all play by.

What mayors should know

- **Federation beats centralization**. Pooling all city data in one AI creates a single point of failure and a legal headache. In a federation, data stays where it is and agents share only the minimum needed.
- **The biggest wins are cross-agency**. Emergency routing, flooding, heat waves, public health, roadworks and major events all involve split data and split authority. No single agent can solve them alone.

- **Residents get one front door**. A “no wrong door” gateway tells people they are dealing with AI, who is involved, and how to reach a human. Residents should never have to work out which agent caused a problem.
- **Accountability is built in, not bolted on**. Every agent has a verified identity, a published mandate, the right to refuse a request, and a full audit trail usable in legal review. Humans decide high-stakes matters.

The mayor's move: Set the rules before the agents multiply. Establish citywide governance, pick two or three high-impact cross-agency use cases, insist on open standards, and bring residents into the design. **The soundbite:** Innovation must not outpace governance. The aim is not to automate government, but to help cities govern wisely and act decisively.

1 Purpose and Overview of the Framework

This framework gives cities a practical, high-level method for designing, deploying and governing Federated Agentic AI.

It has four aims:

1) **Explain clearly** what Federated Agentic AI is and why it matters to cities.
2) **Set out the governance** needed to deploy it safely and responsibly.
3) **Outline a simple architecture** for AI agents working across city platforms.
4) **Show the practical value** of AI agents through real use cases.

The framework is **technology-neutral and adaptable**. Cities at any level of digital maturity can use it.

2 The Rise of Agentic AI in Cities

AI has advanced quickly over the past decade. Early uses were narrow: predictive analytics and other fairly simple tools.

Generative AI and reasoning models have now made possible **agentic AI**: systems that pursue goals and act on their own.

Unlike earlier AI tools, these systems can:

- Plan multistep tasks.
- Interact with software tools and APIs.
- Coordinate with other AI agents.
- Learn and improve on their own or through feedback loops.

That makes AI agents well suited to the public sector, where work often cuts across several agencies.

Cities around the world are already running pilots such as:

- AI assistants for city service centres.
- AI-driven traffic optimization systems.
- Automated document analysis for regulatory compliance.
- Intelligent infrastructure monitoring systems.

Agentic AI could become a core part of how cities are run.

Most cities have already adopted digital transformation strategies aimed at improving public services and operational efficiency.

Agentic AI builds on those strategies by adding the ability to reason and act.

3 Why Cities Need AI Agents Now

Urban governance is becoming increasingly complex. Cities should simultaneously address, among other things:

- Population growth.
- Infrastructure pressures.
- Climate risks.
- Economic transformation.
- Rising citizen expectations for digital services.

Traditional ways of working struggle to keep up. Workflows are fragmented, decisions are slow and often manual, and a limited number of staff must process huge volumes of information.

Agentic AI can extend what city governments are able to do by:

- Processing large volumes of information.
- Automating routine administrative tasks.
- Supporting evidence-based policy decisions.
- Coordinating complex urban systems.

Figure 2: Why Cities Need AI Agents Now



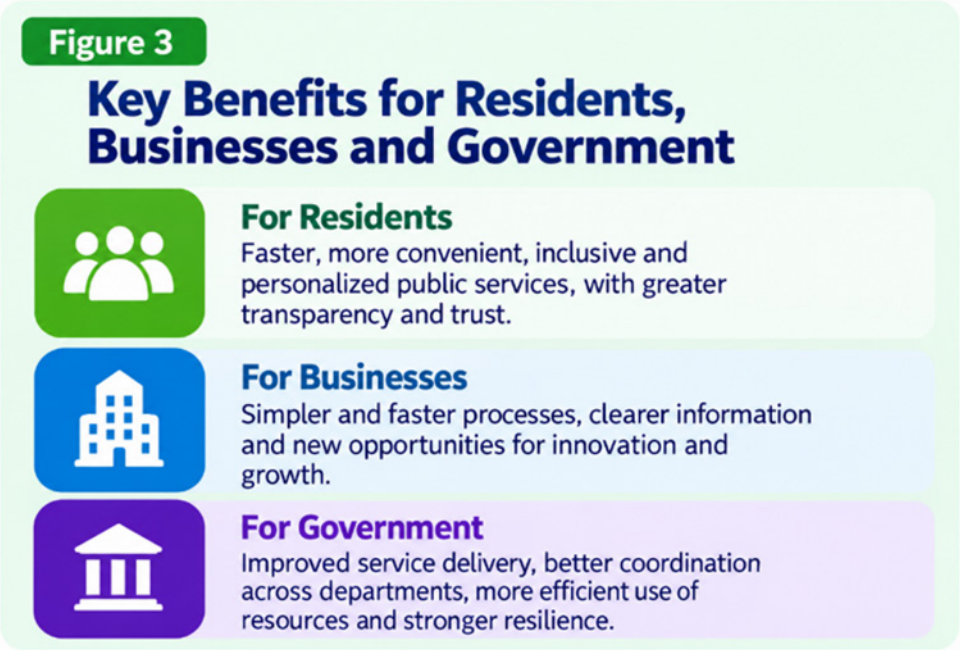
Large language models, reasoning systems and tools for coordinating multiple agents now make it possible to handle demanding tasks such as:

- Document analysis.
- Planning and reasoning.
- Policy simulation.
- Natural language interaction.
- Real-time system monitoring.

For cities, this is a major opportunity to govern better and serve residents better.

4 Key Benefits of AI Agents to Residents, Businesses, and Government

Figure 3: Key Benefits for Residents, Businesses and Government



4.1 Benefits for residents

AI agents can understand requests in everyday language and coordinate services across agencies. For residents, that means public services that are faster and more personal.

Benefits include:

- Unified digital assistants for city services.
- Faster response to service requests.
- Proactive notifications about permits, benefits and obligations.
- Improved accessibility for diverse populations.

4.2 Benefits for businesses

Businesses often have to deal with several agencies to get one thing done. AI agents can guide them through the rules and the paperwork.

Benefits include:

- Automated permit assistance.

- Regulatory compliance guidance.
- Faster licensing procedures.
- Improved transparency in government interactions.

4.3 Benefits for government

For city administrations, AI agents can raise productivity and efficiency.

Benefits include:

- Automated administrative workflows.
- Real-time analytics and decision support.
- Improved inter-agency coordination.
- Reduced operational costs.

Public servants can then spend less time on repetitive tasks and more on work that needs their judgment.

5 What Federated Agentic AI Means for City Governance

A city is not one organization but many: agencies, utilities and public institutions, each with its own job. No single central AI system can manage all of this well.

This framework therefore proposes a federated model instead of a centralized one.

In Federated Agentic AI, many independent AI agents work under one shared set of rules.

In practice:

- Each city agency can deploy its own AI agents tailored to its operations.
- AI agents operate within a shared governance framework.
- Common standards ensure interoperability and contribute to trust across agencies.
- Agents work together and get smarter as a group, without pooling sensitive or proprietary data in one place.

This approach balances three things:

Autonomy – each agency can innovate and build its own specialized agents.

Coordination – all agents follow the same city-wide rules.

Resilience – a failure in one agency or agent is contained and does not disrupt the whole city.

Shared governance prevents fragmentation while letting the system grow and agencies collaborate.

A federation will often reach beyond city limits, to metropolitan transit authorities, regional utilities, national health systems or private operators based elsewhere. That raises questions about which laws apply and how data may be transferred. The federation agreement should settle these in advance: which law governs each type of exchange, which transfer mechanisms apply, and which authority supervises joint outcomes.

6 Benefits of Federated Agentic AI vs Standalone Agentic AI in Cities

A standalone AI agent works well for a self-contained task inside one agency. It is far less effective for the wider business of running a city, which is fragmented across tightly regulated jurisdictions, critical infrastructure and highly sensitive data about residents.

Compared with standalone agents, a federated approach offers four strategic benefits:

- **Legal and jurisdictional boundaries are respected**: A standalone system must gather data in one place to make decisions that cut across functions. Federated agents act locally and share only what is needed, so data stays under its owner's control and local rules are met.
- **No single point of failure**: Putting all data in one place creates a security liability. If a central system is breached, the whole city's data may be exposed. In a federation, data stay spread across the agencies that own it.
- **Services that fit local needs**: A standalone agent built on broad, averaged models often misses the particular needs of a sector or neighbourhood. Federated agents learn from local data and each agency's own policies, while still contributing to a smarter picture of the city as a whole.
- **City-wide collaboration**: Federation works best for challenges with three features: data held in many places, authority split between many bodies, and a need for agencies to coordinate. Emergency response, flooding, heat waves, resilience and public health all fit that pattern.

The table below shows city challenges where a federated approach works better than a standalone one.

City issue	Challenge with Standalone Agentic AI	Benefits of Federated Agentic AI
Emergency response and hospital routing	A single emergency-response AI agent may know emergency calls and ambulance locations, but not necessarily real-time hospital capacity, specialist availability, emergency-room saturation, private ambulance status, traffic-light control, or transit disruptions. Much of that data is sensitive or owned by different entities.	Emergency agents, hospital agents, traffic-signal agents, fire agents, police agents, and mobility agents coordinate. Hospitals share capacity information without exposing patient records. Traffic agents create green corridors. Emergency agents route patients based on live constraints. The federation solves the problem because no single AI agent needs to own all medical, mobility, and public-safety data.
Flood and stormwater management	Flood response depends on storm drains, pumps, river gauges, wastewater systems, road closures, transit tunnels, private basements, construction sites, and weather models. A central agent usually cannot directly control all pumps, gates, private retention tanks, or mobility infrastructure.	Water-utility agents, transportation agents, building agents, transit agents, weather agents, and emergency-management agents share forecasts, local sensor readings, and action constraints. Pump stations can act locally while coordinating with downstream drainage capacity. Road agents close underpasses before vehicles are trapped. The result is system-level flood mitigation without centralizing every sensor and actuator.
Heat-wave protection for vulnerable residents	Heat vulnerability is distributed across health records, utility shutoff data, housing conditions, age, disability status, social-service records, and building cooling capacity. A single agent cannot ethically or legally centralize all that personal data.	Health agents, housing agents, utility agents, community-organization agents, cooling-centre agents, and building agents compute local risk and share limited, privacy-preserving alerts. The federation can identify high-risk zones, open cooling centres, trigger wellness checks, adjust utility policies, and prioritize outreach without exposing individual medical or social-service records to a central AI system.
Multimodal traffic, transit, and event management	A city traffic AI agent can optimize signals, but congestion is shaped by buses, trains, ride-hail vehicles, delivery fleets, parking garages, event venues, emergency vehicles, construction crews, and pedestrians. Many of these actors are outside direct city control.	Traffic agents, transit agents, ride-hail agents, delivery-fleet agents, parking agents, curb-management agents, and venue agents exchange constraints and intentions. The system can coordinate signal timing, bus priority, dynamic curb allocation, freight delivery windows, pedestrian flows, and event egress. A single traffic agent cannot do this well because it lacks authority and visibility across all mobility actors.

City issue	Challenge with Standalone Agentic AI	Benefits of Federated Agentic AI
Construction, road excavation, and infrastructure permitting	Roadworks are often poorly coordinated because utilities, telecom firms, developers, mobility agencies, public works departments, and private contractors all plan separately. A single city permitting agent may see applications but not private construction timelines, utility maintenance plans, or contractor constraints.	Permit agents, utility agents, telecom agents, public-works agents, mobility agents, and contractor agents coordinate schedules and constraints. They can combine trenching, avoid repeated excavations, protect bus corridors, schedule night work, and minimize neighborhood disruption. The value comes from federated planning among entities that would not give full operational control to a single AI system.
Public-health early warning and targeted intervention	Outbreak detection or environmental-health monitoring may require signals from clinics, labs, pharmacies, schools, wastewater plants, employers, and community organizations. A single city health AI usually cannot ingest all identifiable health, school, and workplace data.	Clinic agents, lab agents, pharmacy agents, school agents, wastewater agents, and public-health agents detect local anomalies and share aggregate signals. The federation can identify emerging clusters, deploy testing, target sanitation, adjust school guidance, and alert hospitals while keeping sensitive records local.

Figure 4: Standalone AI vs Federated Agentic AI for Cities: From Siloed Systems to Connected, Interoperable and Coordinated City Services



7 A Model Architecture for Urban Federated Agentic AI

The Urban Federated Agentic AI Framework proposes a high-level **layered architecture** for Federated Agentic AI systems in city environments.

It assumes that many AI agents already exist in city agencies, and already have their own models, tools, memories, data, policies and operating procedures.

7.1 Core Design Principles for Federated Agentic AI

Sovereignty by design

Each AI agent stays under the control of its owner. Joining a federation does not hand authority, control or ownership to anyone else.

Capability contracts over technical interfaces

Agents should publish what they can do and what they are accountable for, not just expose technical connection points. These published "capability contracts" should also count as legal mandates: authority delegated in advance by the organization behind the agent. Commitments the agent makes then bind that organization, but only within the limits of the mandate.

Data minimization by default

Agents should share only the minimum information needed for the task at hand.

Policy precedence over action

Agents may not interact or act across agencies until the relevant policies have been checked and enforced.

Explainable coordination

The system should be able to explain not only what each agent did, but how the agents coordinated.

Commitments beyond communication

It must be clear which agents have formally committed to act. Exchanging messages is not enough.

Human oversight for high-stakes decisions

Humans must review and decide high-stakes matters, such as those affecting rights, safety or politically sensitive issues.

Mandatory semantic interoperability

Agents must understand each other's meaning, not just each other's message formats. Each agent inherits the data models, vocabulary and definitions of its own field, so the same word can mean different things in different organizations. A translation layer, built on agreed reference vocabularies (ontologies) and governed mappings, resolves these differences at the point of exchange. Agents keep their own models while working from a common understanding of intentions, entities and outcomes.

Right to refuse participation

Agents can decline requests that conflict with their own policies, privacy limits, authority or safety rules.

Auditability as a core capability

As far as infrastructure and data-retention rules allow, every action and decision should be traceable and reconstructable afterwards for review, oversight and accountability.

Lawful provenance and legal interoperability

To join the federation, each agent must be able to show that its models, data and knowledge sources were obtained lawfully, and legal roles and responsibilities must be agreed among the participants. Federation agreements should turn legal obligations into machine-readable policies. These obligations include data-protection roles, purpose limitation, records of processing, intellectual-property warranties, liability and duties of redress. Legal interoperability then goes hand in hand with technical and semantic interoperability.

7.2 Architecture overview

The architecture is the connective tissue that lets existing city AI agents find one another, exchange messages, share context, agree commitments and coordinate action safely. It also enforces governance, preserves each agent's autonomy and gives users a coherent experience.

It is designed specifically to connect many city AI agents.

It assumes the agents already exist. They may belong to city agencies, utilities, hospitals, transport operators, telecom providers, emergency services, private companies, community organizations or regional authorities.

Each agent may already have its own:

- Models.
- Tools.
- APIs.
- Workflows.
- Local memory.
- Local data.
- Local users.
- Local governance.
- Local legal authority.
- Local operating constraints.
- Local risk controls.

The federation does not replace existing agents. It lets them work together under agreed rules.

In an urban setting, existing agents may include:

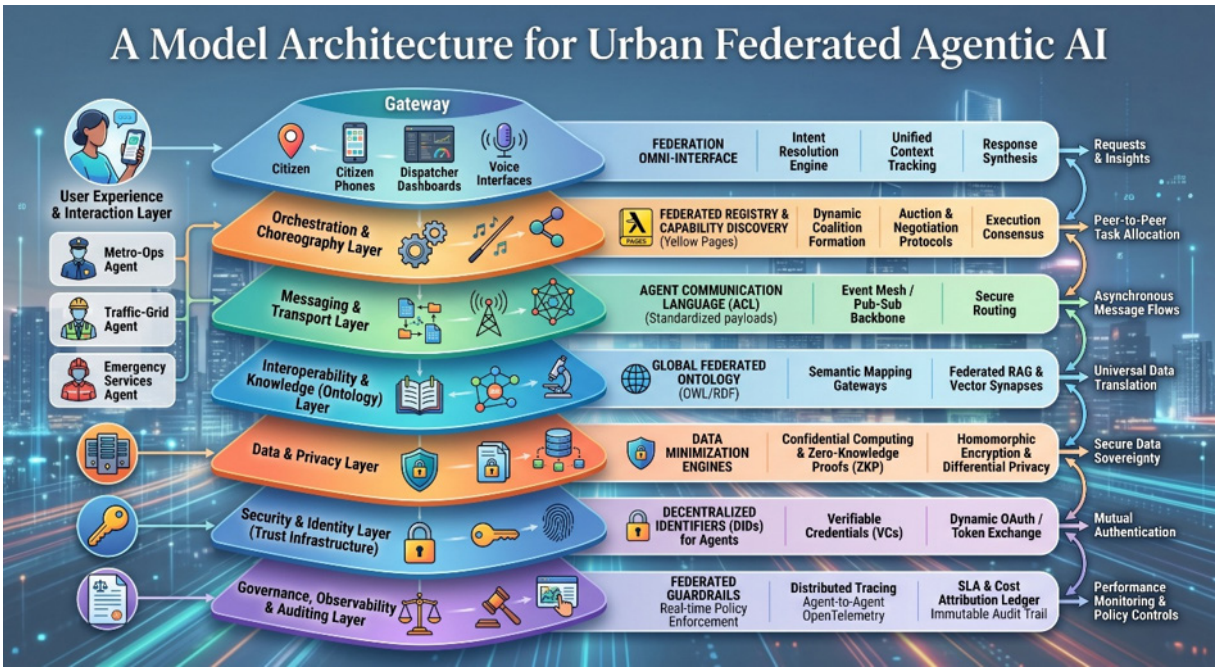
Existing agent	Local responsibility
Traffic signal agent	Optimizes signal timing and corridor flow.
Transit operations agent	Manages bus, rail, paratransit, and service disruptions.
Hospital capacity agent	Reports emergency department capacity and specialty availability.
Emergency Management System (EMS) dispatch agent	Allocates ambulances and recommends routes.
Water utility agent	Detects leaks, flooding, pump failures, and pressure anomalies.
Energy grid agent	Manages feeder load, outages, and distributed energy resources.
Housing services agent	Tracks shelter capacity and housing-service workflows.
Public-health agent	Detects outbreaks, heat-risk patterns, and health-resource needs.
Waste-management agent	Optimizes collection routes and detects illegal dumping.
Building-management agent	Controls heating, ventilation, and air conditioning (HVAC), occupancy, energy use, and safety systems.

Existing agent	Local responsibility
Environmental agent	Air quality, flood prediction, wildfire detection, climate monitoring, etc.

Each agent is autonomous within its own domain. The federation lets them collaborate while staying independent, with no merging into one central AI system.

Figure 5 shows a model architecture.

Figure 5: A Model Architecture for Urban Federated Agentic AI



Source: Author.

It is a simplified, layered design: a network of independent agents connected through a governed federation layer.

7.3 User experience & interaction layer (The Federation Gateway)

This layer is where people meet the federation, whether they are members of the public or system operators. However many agents are at work behind the scenes, users should see one seamless service.

It is the only layer that users and operators deal with directly.

User groups

This layer may serve:

- Residents, businesses and customers in general.
- Field workers.
- Dispatchers.
- City operators.
- Agency managers.
- Elected officials.
- Community organizations.
- Private operators.
- Regional partners.
- Auditors.

Interfaces

This layer may include:

- Conversational interfaces.
- Dashboards.
- Command-centre views.
- Mobile field-worker apps.
- Resident portals.
- Public alert channels.

Users should get one coherent experience, however many agents are involved. This layer hides the complexity so that the system works, and feels, like a single service.

- **Federation omni-interface**: One access point, such as a chat window or dashboard. It accepts any request, even a vague or multipart one, and the user does not need to understand the system behind it.
- **Intent resolution engine**: Works out whether a request can be handled by one agent or needs several working together.
- **Unified context management**: Keeps track of the request, its history and progress as it passes between agents, so users never have to repeat themselves.

- **Response synthesis**: Combines the answers from different agents into one clear, user-friendly reply.
- **Single point of notice and redress ("No Wrong Door")**: The gateway is where the federation meets its duties to the public. It tells users that they are dealing with AI agents, which organizations are involved and how to reach a human. It also receives complaints and requests for redress on behalf of the whole federation; responsibility is then assigned internally using the audit trail. Residents should never have to work out which agent caused an outcome.

7.4 Orchestration & choreography layer (The Collective Intelligence)

Traditional systems rely on one central controller directing every action (orchestration). A federation relies mainly on choreography: independent agents deal with each other directly, following shared protocols.

- **Federated AI agents registry & capability discovery**: A live directory, like a "yellow pages", where agents publish what they can do, their areas of expertise, their tools and how to connect to them. Agents use it to find and work with each other.
- **Dynamic coalition formation**: Assembles the right team of agents for each request. There is no fixed line-up; the system picks the best agents for the task.
- **Auction & negotiation protocols**: When several agents could do the same job, structured negotiation decides which one takes it, based on factors such as speed, cost, reliability or access to the right data.
- **Execution coordination and consensus**: Supervises tasks that involve several agents. It manages dependencies, detects failures and triggers fixes such as retrying or reversing a step that did not complete.

7.5 Messaging & transport layer (The Communication Backbone)

Agents can only cooperate if they have a reliable and secure way to communicate, one that does not require them all to respond at the same moment. This layer sets the rules for how agents exchange information and coordinate.

Agent communication language (ACL): A formal standard for how messages are structured, built on established frameworks (e.g., FIPA-ACL, JSON-LD, gRPC payloads). Every message identifies the sender, recipient, intent, content and conversation ID, so that exchanges are clear and traceable.

Event mesh/Pub-sub backbone: A decentralized messaging system (e.g., Kafka, NATS or distributed ledger-based systems) in which agents subscribe to the event streams relevant to them and respond to the tasks they can handle. This keeps coordination scalable and agents loosely linked.

Secure routing: Encrypts all agent-to-agent messages end to end and routes them safely across network boundaries, including firewalls. This is critical between organizations, where trust, confidentiality and integrity must be protected.

7.6 Interoperability & knowledge (ontology) layer

This layer lets very different AI agents communicate, even when they were built by different vendors, organizations or teams and do not share protocols, message formats, data models or security frameworks.

Existing agents often use different internal "languages" and data structures. This layer acts as a universal translator so that all agents can understand one another.

Global federated ontology: A shared vocabulary of terms, concepts and relationships for the whole federation, usually built as a semantic knowledge graph (e.g., OWL/RDF). It ensures that every agent interprets information the same way.

Semantic mapping gateways: Translators attached to each agent that map its local data to the shared vocabulary. For example, one agent's internal field "client_id" can be mapped to a standard concept such as "Universal Customer UUID".

Federated retrieval and distributed knowledge access: Agents can securely query each other's knowledge sources such as vector databases or knowledge graphs, instead of pooling data centrally. They exchange insights and relevant results, not raw data.

7.7 Data & privacy layer

In a federation, data stays at its source and is not gathered into a central store. Agencies collaborate while each keeps managing and protecting its own data.

Data minimization mechanisms: When one agent asks another for help, only the minimum information needed is shared. This limits exposure of sensitive data and keeps sharing proportionate to the task.

Confidential computing & zero-knowledge proofs (ZKP): These techniques let an agent prove that it holds certain data, or has checked a condition, without revealing the data itself. Trust is built without giving up confidentiality.

Privacy-preserving computation techniques: Methods such as homomorphic encryption and differential privacy allow data to be processed securely. Calculations can run on encrypted data, or controlled statistical noise can be added to results, so that sensitive information cannot be reconstructed or misused.

Legal roles and federation data contracts: For each type of exchange between agents, data-protection roles (e.g., controller, joint controller, processor) should be assessed and written into federation data contracts. Data must stay tied to the purpose for which it was collected; it must not drift into unrelated scoring or enforcement uses. The layer should also automatically keep records of processing between agents.

7.8 Security & Identity Layer (Trust Infrastructure)

This layer provides a **trust-based security framework**. Every interaction between agents is verified, controlled and secured, on the principle that no participant is trusted automatically (zero trust).

Decentralized identifiers (DIDs) for agents: Each agent has a unique, cryptographically secure identity, usually based on key pairs. Agents can verify each other, and impostors ("spoofing") are kept out of the network.

Verifiable credentials (VCs): Agents can hold digitally signed credentials that certify what they are able to do, what they are permitted to do and whether they comply with the rules. For example, a credential could confirm that an agent is authorized to process sensitive health-care data.

Dynamic authorization and token exchange: When an agent needs to use another agent's resource or API, it receives an access token limited to a set time and purpose (e.g., OAuth-based mechanisms). Access is granted only when needed, which reduces the security risk.

7.9 Governance, Observability & Auditing Layer

Federated agents act with a high degree of autonomy. Strong safeguards and thorough audit trails are, therefore, essential to prevent unintended behaviour, errors that spread and breaches of regulation.

Federated guardrails: Policy engines check agents' messages and actions against the rules in real time. If an agent tries to share restricted information or take an unauthorized action, the system blocks or modifies the request.

Distributed tracing across agents: Tracks how a request moves from agent to agent (e.g., from AI agent A to AI agent B to AI agent C). This shows how decisions were made and helps answer the question, "Why was this result produced?"

Service-level accountability and cost attribution: Transparent logs record resource use, performance and transactions between agents, possibly on tamper-proof ledgers or blockchain-based systems. They show whether service-level agreements (SLAs) are met and allow costs such as computing resources or tokenized services to be tracked and settled between participants.

Regulation-aware guardrails (Policy-as-Code): Guardrails should encode legal triggers, as well as municipal rules. For example, they should flag when a coordinated action would amount to a legally significant automated decision that needs human authority, or when an exchange would move sensitive data beyond its original purpose.

Evidentiary-grade tracing: When several agents produce a decision together, responsibility is shared. Trace logs are the main way to establish afterwards who is liable. The federation agreement should, therefore, set rules for how long traces are kept, how their integrity is protected and who can access them. Traces should be fit for use in administrative review and legal proceedings.

7.9.1 What belongs inside the existing AI agents versus inside the federation

Belongs inside each sovereign AI agent	Belongs in the federation architecture
Local model selection	Agent identity and certification
Local tool execution	Agent discovery and capability registry
Local memory	Shared case context
Local planning	Cross-agent orchestration
Local data governance	Federated data contracts
Local human workflow	Cross-agent human approval
Local API design	API and protocol mediation
Local optimization	Multi-agent negotiation
Local policies	Federation policy-as-code
Local logging	End-to-end audit trail
Local authority	Commitment and execution assurance
Local risk controls	Federation-wide risk governance
Local UX	Unified federation experience

Appendix 1 gives a short note on practical standards for urban federated agentic AI in 2026.

7.9.2 Urban Federated agentic AI architecture applied to a real-life city example

Consider a major flash flood that disrupts public transport, road traffic and emergency services. In most cities these functions run on separate AI systems, built by different vendors, on different infrastructure, managed by different agencies. They cannot simply be merged into one platform. They need to work together as a federation.

The steps below show how the layered architecture coordinates these independent systems in a high-stakes, time-critical crisis.

The scenario

The request: A city dispatcher or an automated sensor detects the flood and issues an alert: "Flash flooding at the Central Metro Station intersection. Optimize safety, reroute traffic, and notify citizens."

Step-by-step lifecycle through the federated layers

1) **Request intake and intent analysis (user experience and interaction layer)**

- The alert arrives at the Federation Gateway, the single entry point for all requests.
- The **Intent Resolution Engine** recognizes that no single department or system can handle the situation. It breaks the request into coordinated tasks:
    - Assessing infrastructure damage.
    - Rerouting public transit services.
    - Redirecting surface traffic.
    - Issuing public safety notifications.
- To ensure consistent tracking and coordination, the system generates a unique **Global Context Identifier** (e.g., *Session_2026_Flood_091*), which is used to maintain continuity and traceability of the incident across all participating agents and systems.

2) **Discovery & team assembly (orchestration & choreography layer)**

- The orchestration broker searches the Federated AI Agents Registry, the "yellow pages", for relevant and available agents. It assembles a temporary Crisis Response Coalition:
    - Metro Operations AI Agent: Responsible for managing subway and rail systems.
    - Traffic Grid AI Agent: Responsible for traffic signals and smart road infrastructure.
    - Emergency Services AI Agent: Responsible for dispatching police, fire, and medical services.
- Where agents' capabilities overlap, the system decides who leads. For example, the Traffic Grid and Metro Operations agents can both access street-level camera feeds. Through an auction and negotiation protocol, the Traffic Grid AI Agent takes the lead on monitoring the intersection because it responds faster, while the Metro Operations AI Agent adds data from its underground sensors.

3) **Secure handshake & identity verification (security & identity layer)**

- Before sharing sensitive infrastructure data, the agents must establish that they can trust one another.
- The Traffic Grid AI Agent presents its Decentralized Identifier (DID) and a Verifiable Credential (VC) confirming that it is a city entity authorized to manage traffic systems.
- The agents receive short-lived, encrypted access tokens, so they can communicate securely with each other for the duration of the incident only.
- Verifiable credentials can also certify legal limits: for example, that the hospital-capacity agent may share only aggregate capacity figures that identify no one.

4) **Universal translation (interoperability & knowledge layer)**

- The Metro Operations AI Agent operates on a legacy data schema (e.g., "Station_ID: 402, Platform B"), while the Traffic Grid AI Agent uses geospatial references (e.g., GPS coordinates or street intersections such as "Main & 4th Street").
- Semantic Mapping Gateways translate both formats into the shared vocabulary of the city's Global Federated Ontology.
- Consequently, when Metro Operations reports a structural risk at Platform B, the Traffic Grid AI Agent knows at once where that is at street level and can anticipate risks such as the road above collapsing.

5) **Distributed data exchange & execution (messaging & transport layer)**

- Data is not gathered in one place. The agents exchange messages over a secure, encrypted messaging mesh, without waiting on each other.
- The Metro Operations AI Agent publishes an event such as: [Metro station flooded: Station closure initiated].
- The Traffic Grid AI Agent picks up this event, works out alternative routes, updates digital road signs ahead of the affected zone and retimes traffic signals to prioritize evacuation.
- At the same time, the Emergency Services AI Agent picks up the same event and routes emergency vehicles away from the flooded intersection.

6) **Governance, compliance, and safety enforcement (governance & guardrails layer)**

- Throughout, every action is checked against policy in real time.
- For example, if the Traffic Grid AI Agent tries to open a restricted bus lane to ease congestion, the Federated Guardrail Engine checks the action against municipal policy. Here it confirms that the measure is allowed during a flood and authorizes it.

- The guardrail engine also checks that the public emergency notice and any instructions on movement are within the issuing authority's legal powers and disclose no personal data.
- Every interaction, decision and system change is logged through distributed tracing and tamper-proof audit records (e.g., OpenTelemetry and secure ledgers), giving full transparency and accountability after the incident.

7) **Synthesis & egress (back-to-user experience layer)**

- The system gathers the responses from all agents:
  - Metro operations: Train routes adjusted; station bypassed.
  - Traffic grid: Road signage updated; traffic signals reconfigured.
  - Emergency services: Emergency units rerouted.
- These are combined into one clear operational update on the city's command dashboard.
- In parallel, the system sends a localized emergency notice to affected residents, so public communication is timely and coordinated.

8 Call to Action for Cities on Urban Federated Agentic AI

Figure 6: Implementation Roadmap for Cities



Cities are at a defining moment. Rapid advances in AI offer a historic opportunity to strengthen city administration, improve services and respond better to complex, fast-changing public needs. Urban Federated Agentic AI is a particularly powerful model. Intelligent systems work across departments, jurisdictions and services while staying anchored in public purpose, local accountability and prudent oversight.

City governments should move beyond experiments and start building the institutional, technical and ethical foundations for responsible deployment at scale. This is not just a technology upgrade. It is part of the strategic digital transformation of the public sector. Properly governed, these systems help cities coordinate services, spot risks earlier, allocate resources more intelligently and give residents more responsive, inclusive and personal support.

To deliver on this promise, cities should:

- Set up robust governance, with clear accountability, oversight and data stewardship rules, and mechanisms for audit, transparency and redress.
- Pick priority use cases that meet high-impact public needs, such as permitting, mobility, emergency response, social services, infrastructure maintenance and other public services.
- Build federated AI architectures that let agencies collaborate securely while keeping their autonomy, data protection and policy control.
- Invest in digital public infrastructure, interoperable data systems, cybersecurity and the technical capacity to deploy and maintain these systems responsibly.
- Prepare the workforce through targeted training and reskilling, and build multidisciplinary teams that combine policy, technical, legal and operational expertise.
- Involve the public. Engage residents, civil society, academia and the private sector in setting priorities, safeguards and performance expectations.
- Work with other cities and share knowledge internationally to learn faster, avoid duplication and agree shared standards and best practices.
- Prefer open international standards and, where appropriate and viable, open-source components for federation infrastructure. National federated data-exchange layers and open urban-platform ecosystems have proven this approach at scale. It keeps the coordination fabric of the city auditable, portable and under public control.
- Monitor and evaluate continuously, measuring public value, equity impacts, system reliability and alignment with human rights and the public interest.

Innovation must not outpace governance. Urban Federated Agentic AI should be deployed only in ways that are lawful, explainable, secure, inclusive and consistent with fairness, transparency and human-centred public service. The aim is not to automate government for its own sake. It is to help cities govern wisely, act decisively and serve all residents better.

City leaders should act with urgency, discipline and ambition. Investing now in responsible Urban Federated Agentic AI lays the foundation for city governance that is more adaptive, resilient and trusted, and able to meet the challenges of the twenty-first century.

Appendix 1: A Brief Note on Practical Standards for Urban Federated Agentic AI in 2026

In a federated urban environment, many independent organizations each run AI agents on their own platforms and keep ownership of their data, systems and decision rights. They include transport and mobility operators, utilities, health and emergency services, municipal departments and private providers. No single authority controls the whole, so interoperability cannot be achieved by centralizing. It has to be built from open standards that work at different layers and are designed to fit together. There are four layers: the agent interaction layer (how agents reach systems and each other); the knowledge or semantic layer (whether they mean the same thing); the trust and identity layer (who may act, and on whose authority); and the governance layer (what makes the whole accountable). Agents can collaborate autonomously when each layer is standardized separately and the boundaries between layers are explicit.

At the interaction layer, the Model Context Protocol (MCP) and the Agent-to-Agent (A2A) protocol have become the *de facto* foundations. Both are currently stewarded under the Linux Foundation's Agentic AI Foundation. MCP standardizes the vertical link between an agent and the tools, data sources and services of its own domain. A traffic-management agent can call signal-control systems, or a utilities agent can query grid telemetry, through one uniform, self-describing interface. A2A standardizes the horizontal link between agents in different organizations. Capabilities are published as Agent Cards, tasks are delegated with a defined lifecycle, and results come back as structured artefacts. An agent built on one vendor's framework can, therefore, work with a peer built on another's. Beneath these sits the vocabulary of communicative acts inherited from FIPA ACL, the Agent Communication Language: request, propose, accept, refuse, inform. It gives exchanges between agents an explicit structure, so that collaboration becomes negotiation and tracked commitment, not just message passing. Ahead of a large gathering, for example, an event-management agent can issue a call for proposals to mobility, parking and public-safety agents. Each replies with a binding acceptance, a counter-proposal or a reasoned refusal.

Agreeing on protocols does not guarantee shared meaning. City domains have long-established but different information models: a "zone", an "asset" or an "incident" means one thing to a transport authority, another to a utility and another to an emergency service. The semantic layer addresses this with recognized urban data standards: ETSI NGSI-LD and its Smart Data Models for context information, the W3C SOSA/SSN ontology for sensors and observations, the OGC SensorThings API for IoT feeds, JSON-LD and RDF for machine-readable representation, and DCAT for cataloguing data assets. These are combined with semantic mediation, which maps each organization's local vocabulary to shared reference ontologies at the point of exchange. Data carried over MCP and A2A therefore arrives with its context attached. The receiving agent can interpret, validate and reason over it, and no participant has to give up its own models. This is what turns a set of connected agents into a system that shares understanding.

In a federation, trust cannot be assumed. It must be established at every boundary. Mature identity and authorization standards are applied consistently across the stack: OAuth 2.0 and OpenID Connect for delegated access, mutual TLS for transport security, and W3C decentralized identifiers and verifiable credentials where organizations need portable, verifiable proof of which agent is acting and on whose behalf. At the agent layer, A2A introduced signed Agent Cards for verifiable identity, while MCP's authorization model governs which systems an agent may reach. Observability standards such as OpenTelemetry make delegated tasks traceable end to end across organizations. Management-system standards such as ISO/IEC 42001 provide the organizational frame in which policy, risk appetite and human oversight are set. Together, these layers let agents collaborate autonomously across the city while every action remains attributable, authorized and auditable. One caveat: the stack is still settling, and the protocols at each layer may be revised or replaced as the field matures.

International standards made specifically for cities are also on the way. ITU-T Study Group 20, the ITU standardization group for the Internet of Things, digital twins and smart sustainable cities and communities, is developing international standards on agentic AI for cities. To prepare the ground, in 2026 it set up the ITU Focus Group on AI for Smart Sustainable Cities and Communities (FG-AI4SSC), an open platform where cities, governments, industry, academia, other standards bodies and UN entities work together. The Focus Group held its first meeting in Tampere, Finland, in September 2026. Its pre-standardization work covers the requirements, interfaces and coordination frameworks that AI-enabled urban systems need, with agentic, physical and spatial AI among its key topics, and its results will feed into the standards of Study Group 20. Cities deploying federated AI agents should follow this work and take part in it, so that the resulting standards reflect what cities actually need.


Ai Search
AI
artificial intelligence
AI Prompt
Generate
AD
Coding



For more information,
please contact: u4ssc@itu.int
Website: https://u4ssc.itu.int/

ISBN 978-92-61-43851-7

9 789261 438517

Published in Switzerland
Geneva, 2026