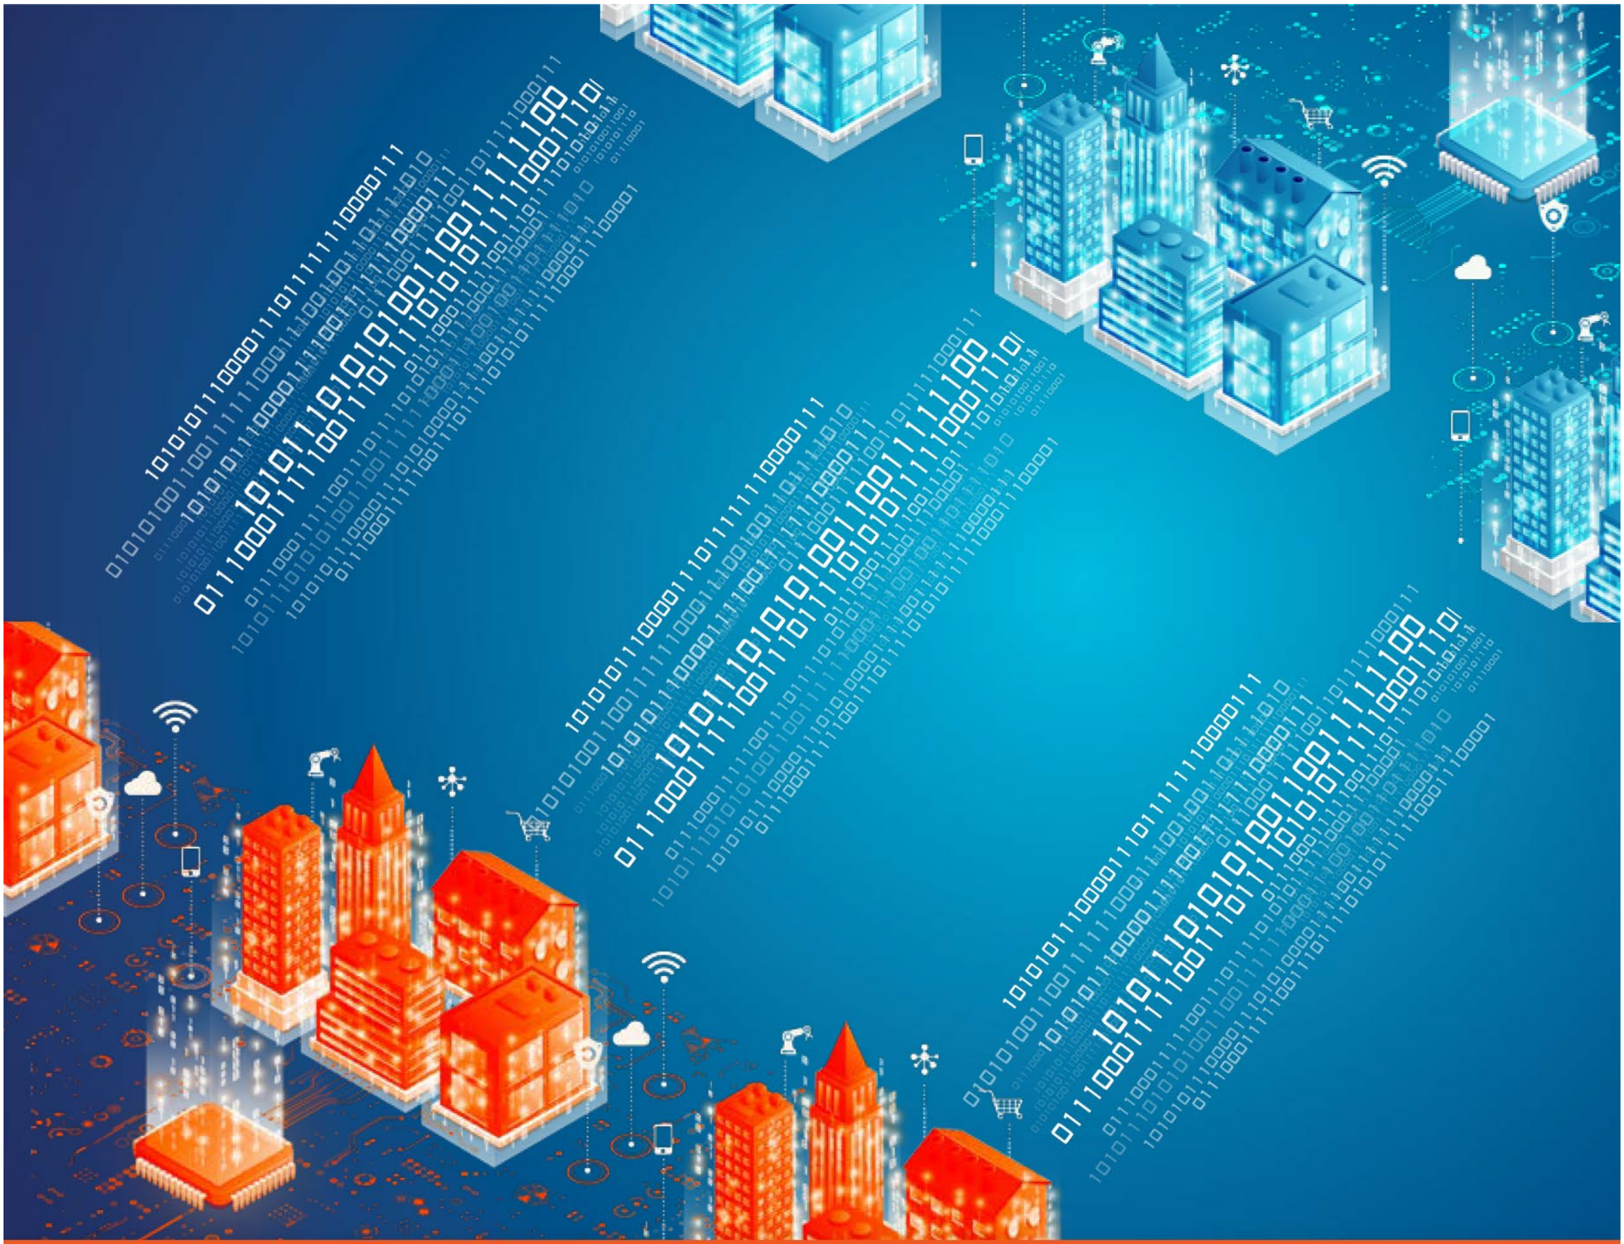

Smart Cities Data Catalog Specification

ATIS-I-0000080 | August 2020

Abstract

As smart city projects continue to expand and evolve, data sharing sits at the intersection of business opportunities and technology developments. Cities can certainly benefit from data sharing across smart city applications and sectors. However, sharing among cities, as well as the development of data exchanges and marketplaces, will signal that smart cities are moving to the next level of value creation for citizens and local governments.

This report assesses data sharing alternatives for smart cities. It also proposes a blueprint for a common framework, a set of critical components and an evolutionary path from data collection to value added capabilities (e.g., economic development, data monetization, third-party relationships).

Foreword

As a leading technology and solutions development organization, the Alliance for Telecommunications Industry Solutions (ATIS) brings together the top global ICT companies to advance the industry's business priorities. ATIS' 150 member companies are currently working to address 5G, cybersecurity, robocall mitigation, IoT, artificial intelligence-enabled networks, the all-IP transition, network functions virtualization, smart cities, emergency services, network evolution, quality of service, billing support, operations, and much more. These priorities follow a fast-track development lifecycle – from design and innovation through standards, specifications, requirements, business use cases, software toolkits, open source solutions, and interoperability testing.

ATIS is accredited by the American National Standards Institute (ANSI). ATIS is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a founding Partner of the oneM2M global initiative, a member of the International Telecommunication Union (ITU), and a member of the Inter-American Telecommunication Commission (CITEL). For more information, visit www.atis.org.

This document was developed as part of a partnership between ATIS and US Ignite.

US Ignite is a non-profit organization accelerating the smart community movement by guiding communities into the connected future, creating a path for private sector growth, and advancing technology research that is at the heart of smarter development. For more information, visit www.us-ignite.org.

Notice of Disclaimer and Limitation of Liability

The information provided in this document is directed solely to professionals who have the appropriate degree of experience to understand and interpret its contents in accordance with generally accepted engineering or other professional standards and applicable regulations. No recommendation as to products or vendors is made or should be implied.

NO REPRESENTATION OR WARRANTY IS MADE THAT THE INFORMATION IS TECHNICALLY ACCURATE OR SUFFICIENT OR CONFORMS TO ANY STATUTE, GOVERNMENTAL RULE OR REGULATION, AND FURTHER, NO REPRESENTATION OR WARRANTY IS MADE OFMERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR AGAINST INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. ATIS SHALL NOT BE LIABLE, BEYOND THE AMOUNT OF ANY SUM RECEIVED IN PAYMENT BY ATIS FOR THIS DOCUMENT, AND IN NO EVENT SHALL ATIS BE LIABLE FOR LOST PROFITS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES. ATIS EXPRESSLY ADVISES THAT ANY AND ALL USE OF OR RELIANCE UPON THE INFORMATION PROVIDED IN THIS DOCUMENT IS AT THE RISK OF THE USER.

NOTE - The user's attention is called to the possibility that compliance with this standard may require use of an invention covered by patent rights. By publication of this standard, no position is taken with respect to whether use of an invention covered by patent rights will be required, and if any such use is required no position is taken regarding the validity of this claim or any patent rights in connection therewith. Please refer to www.atis.org/01_legal/patent-policy/ to determine if any statement has been filed by a patent holder indicating a willingness to grant a license either without compensation or on reasonable and non-discriminatory terms and conditions to applicants desiring to obtain a license.

Copyright Information

ATIS-I-0000080

Alliance for Telecommunications Industry Solutions
1200 G Street, NW, Suite 500
Washington, DC 20005

Table of Contents

1. Introduction

1.1 Purpose

This document was developed as part of a partnership between cities and industry to advance the interoperability of smart city *data catalog* solutions. It was created as part of a larger collaboration between ATIS and US Ignite that is focused on developing a set of smart city *data exchange* specifications and bringing together a core of data-centric cities and leading industry solution providers.

The technical and business requirements contained in this document are intended to create a data catalog blueprint, which can be used to exchange data between cities and city data partners. The data catalog is a key part of a smart city data exchange, providing a linkage between data producers and data consumers.

Data catalogs will play a pivotal role in smart cities, given the growing demand for smart city data and the opportunities for data sharing. It is expected that this specification will be used by municipal chief data officer and chief technology officer/chief information officer organizations to plan, procure, implement and operate smart city data exchanges.

1.2 Description

The use of data catalogs has increased in the past few years in big data and other data management applications because they enable data users to discover and search datasets and data sources, without the cumbersome need to know the exact data connection string, syntax or path. Data catalogs can leverage open APIs to create a simplified mechanism to identify data that is relevant to a data consumer’s application or need.

In the context of a smart city, data catalogs allow city operations and affiliated third parties to register metadata from any number of data sources to a common repository that can be discovered, searched and filtered by a data consumer partner. Potential data consumers using the catalog include other smart cities, government agencies (i.e., local, state, federal), trusted partners, application developers, citizens and businesses. It is important to note that data catalogs act as a repository and visualization of metadata and thus do not store the actual data itself.

At the core of a smart city data catalog is the fundamental concept of data producers and data consumers. By using data catalogs, *data producers* can register, profile and, in some cases, preview their metadata. This metadata may be associated with data located in data warehouses, data lakes, databases, cloud-stored data, third-party data sources and crowdsourced data. Similarly, *data consumers* are presented with a comprehensive and centralized view of available data sources within a city’s data governance domain, thus creating a powerful data acquisition tool. In the future, an additional role will be added to the smart city data ecosystem: *data enrichers*, who will utilize capabilities such as advanced analytics, machine learning (ML), artificial intelligence (AI) and other value-added features to enhance the data prior to consumers requesting the data.

1.3 Benefits

The value of a data catalog must be considered within the larger context of a data exchange. Smart city data exchanges are not intended to operate within the data management boundaries of city operations, where cities already manage data across municipal silos and departments. Data exchanges exist at a layer on top of city data management platforms and provide the means to exchange data between cities and their data partners. Although cities could develop unique requirements and apply them on a pairwise basis with other cities or data entities, a common approach to data exchanges will provide an interoperable and extensible framework that can operate across a broad range of partners and applications, well into the future.

The smart city data catalog functions as a fundamental building block of the data exchange. As discussed later in this document, data catalogs can be deployed in several different scenarios. However, the common thread is the ability to publish metadata in a consistent manner, allowing data partners to search, request and consume data through a simplified set of open APIs. Given that smart city data has an acquisition cost, data catalogs provide cities with a means to create significantly greater visibility and usability of their data assets. In addition, they create new opportunities for cities to work with their data partners to monetize and create value on top of data collection.

Perhaps the best means to illustrate the benefits of a data catalog is to consider some of the most likely use case classes, as discussed below:

Economic Development – Private companies, city developers, government agencies and others interested in business development and investment want to access a broad range of city data across a number of municipalities. This data enables them to explore snapshots of development potential on a property-by-property basis based on zoning, development standards, infrastructure requirements and funding options.

Transportation – Commuters, ride-sharing services, taxi services and private transportation companies would like to assess real-time traffic flow, road conditions, accident reporting, parking availability, public transportation scheduling and more across a large metropolitan area or a regional set of municipalities. The information can be used to generate arrival times, parking reservations, route planning and other services. The data sources for this information include city Internet of Things (IoT) sensors, third-party data and crowdsourced data. A common cataloging of data across an entire transportation region, combining real-time and time-series data and integrating public/private data reporting would deliver an end-to-end set of transportation applications for citizens and business.

Environmental – Public and private IoT sensors can collect a broad range of environmental data, including temperature, air quality, wind, water levels and seismic activity. These data sources provide valuable information and can be made available on a municipality-by-municipality basis. However, collecting and analyzing this information comprehensively across a larger mesh of data catalogs could provide city agencies and private entities with a much greater visualization of impending weather conditions, health warnings, flood control and fire advisories.

Public Safety – Emergency response is one of the most critical functions of local, state and federal government. Emergency communications and inter-agency databases provide valuable information to assist first responders. However, in many cases, an emergency response agency must leverage city data on a real-time basis across a broad base of municipalities. A common data catalog approach that provides first responders with restricted but rapid access to mission-critical data (combined with relevant third-party or crowdsourced data) could vastly expand the situational awareness of on-the-ground emergency response teams.

Tourism – Similar to the economic development use case, tourism is a valuable component of a city's economic base. But tourism is not always limited to a specific municipality. Instead, it often involves applications beyond city information, such as public transportation alternatives, parking availability, sports events, optimum visit times for attractions, museums, etc. Enabling visitors to search and filter related services across a region of municipalities would deliver a vastly improved tourism experience and help to move tourists through a city or metropolitan area in the most efficient manner.

Designers need to consider a broad array of datasets that a data catalog would handle. This includes relatively static datasets that are used in infrastructure planning and economic development such as:

- City pedestrian and traffic sensor data
- City building or housing permits
- Industry ride-sharing data
- City transit routes
- Public and affordable housing stock
- Locations of major employers and small business clusters
- Environmental data such as air quality, noise levels and greenhouse gas emissions
- Population and census data

- 911/311/crime data
- Schools/education data
- Budget data
- Drone data

There is also a need to manage streaming/time-dependent data from sources such as traffic sensors, CCTVs, environmental sensors and pedestrian/vehicle counters.

1.4 Deployment Scenarios

Smart city data catalogs will be deployed across a broad range of geographical, urban/suburban/rural, demographic and application-oriented environments. This section describes the two most likely architectures, with the understanding that a data catalog configuration may evolve into another configuration, based on municipal and regional considerations.

Given the practical considerations related to city autonomy of services, budgets, jurisdiction and other factors, it is expected that *federated (local) data catalogs* will have a significant presence in the smart city development cycle. A federated data catalog acts as a local catalog but is part of a larger federation of data catalogs from other cities and municipalities, based on a common approach to registration and discovery of data across various catalogs.

Although cities will determine which data is allowed into a local catalog and how that data is used, the actual management and governance of the data catalog could be provided by city personnel or by an authorized third-party provider. Figure 1.1 provides an architectural view of local data catalogs.

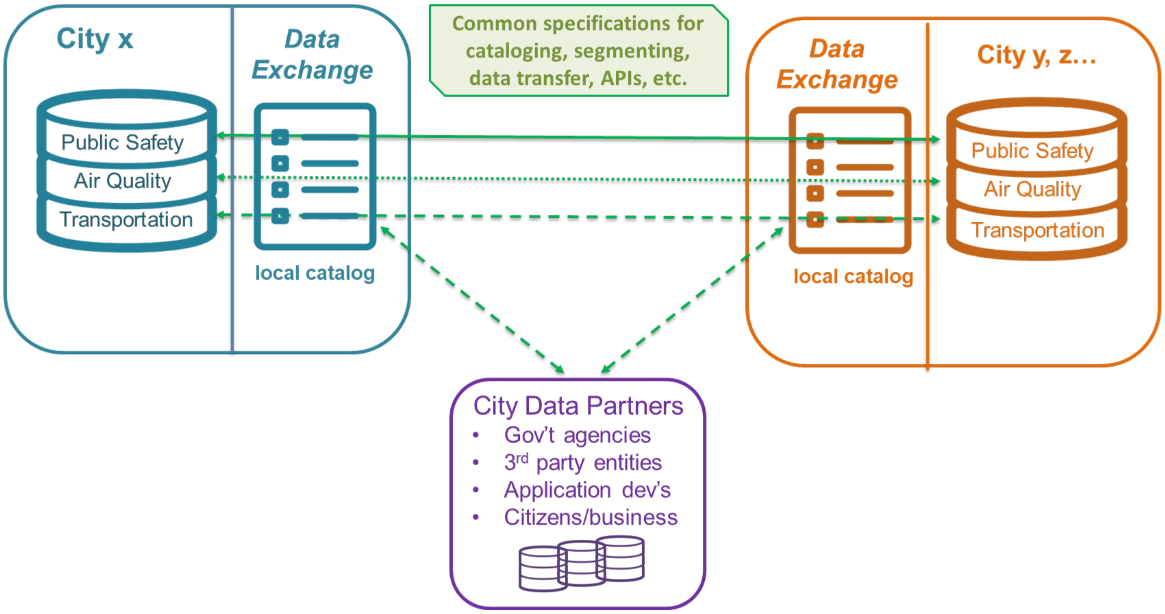

Figure 1.1 Federated Local Catalog Example

The following is a list of federated (local) data catalog characteristics:

- Cities maintain local catalogs of shareable data
- Data is segmented by field of use and restriction levels
- Federated partners evolve over time based on overlapping dataset needs
- Cities control which data can be federated and which is kept local
- In some cases, federations might be a requirement of a state or federal body
- Federated catalogs could outwardly present themselves as a single catalog offering coordinated data access
- Cities create metadata repositories in local catalogs based on data partner agreements and local data-sharing policies
- Security and privacy control are enforced at city level

The second type of catalog configuration covered in this document is a *centralized (regional) data catalog*. In this arrangement, a single catalog (as part of a shared data exchange) supports multiple municipalities. The most common application space for a centralized data catalog is a group of municipalities that are geographically related and choose to pool data as part of a centralized catalog. However, a centralized catalog is not necessarily limited to adjoining local governments and may act as the preferred option for any group of cities that share a common set of needs.

Similar to federated local catalogs, centralized catalogs can offer a variety of data ownership and management scenarios, but there are more drivers for third-party stewardship of centralized catalogs because the data is generated and registered from a number of city governments.

Figure 1.2 shows the basic architecture for a centralized catalog, illustrating the flow of data for a transportation application.

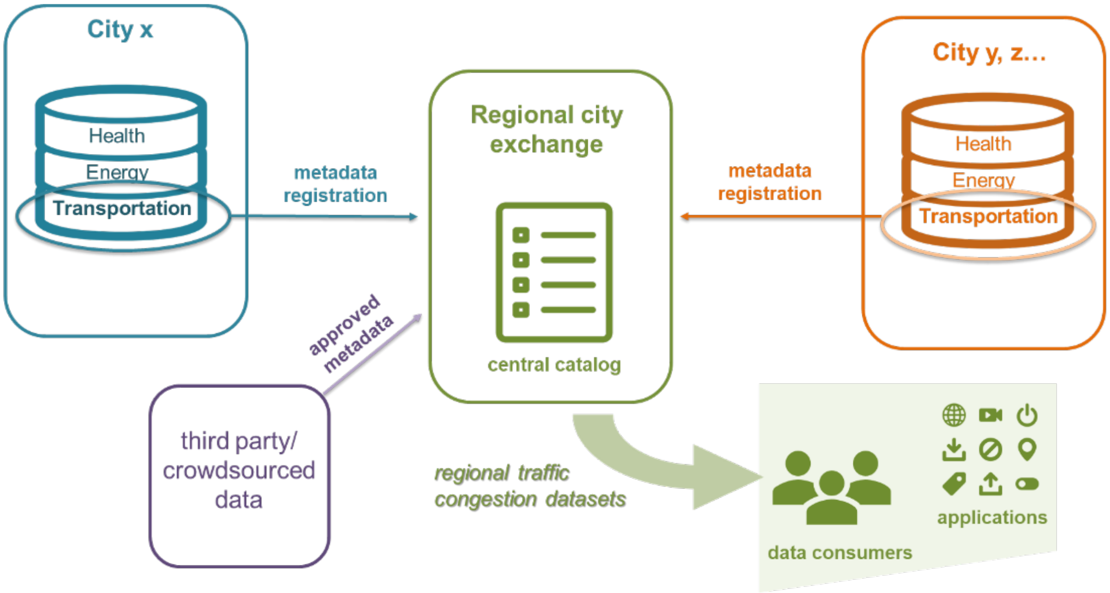

Figure 1.2 Centralized Catalog Example

The following is a list of basic characteristics of a centralized data catalog:

- Acts as a regional catalog for a group of related municipalities (geographical or other common relationship)
- Builds upon the needs of a set of common regional applications, so there is a mutual benefit in pooling data
- Some entity in the ecosystem host is selected to host the centralized exchange (local government or third-party partner)
- From a user perspective, it operates as a central catalog for discovery and access to data
- Cities independently approve all cataloged data
- Consumers of data are subscribed and segmented by access level, which must be administered across multiple municipalities

While this section has provided an overview of different catalog configurations, it is important to note that the basic requirements surrounding registration and discovery of metadata are generally consistent across the various deployment scenarios.

2. Definitions, Acronyms and Abbreviations

For a list of common communications terms and definitions, visit the ATIS Telecom Glossary at https://glossary.atis.org.

2.1 Definitions

Centralized Data Catalog: Collection of data catalog resources from multiple cities managed on a regional basis by a city or affiliated organization.

Crowdsourced Data: Data sourced from individuals, devices or organizations that is relevant to a specific application, service or need.

Data Lake: A repository that stores data from many sources, typically in multiple raw or native formats, which may include structured or unstructured data.

Data Warehouse: A centralized repository that stores and integrates data from many different sources to facilitate analytics and decision-making.

Discover: Process used by a user, machine or system to become aware of data that exists within the producer or enricher domain.

Federated Data Catalog: Locally managed metadata resource that is managed within the jurisdiction of a single city or local government.

Filter: A data-refinement process for extracting specific data from a large pool of data resources that is specifically relevant to a need or purpose.

Metadata: Information that provides descriptive, relational or locational characteristics related to the data for the purpose of creating a linkage between producers and consumers.

Open API: Publicly accessible application programming interface, where the owner of a resource, service or application gives universal access to developers or consumers.

Policy Management: A mechanized process that manages the use, licensing, privacy, consent and sharing of city-owned data assets, third-party data and crowdsourced information.

Register: Process of creating or extracting metadata from a data source and publishing approved metadata to a cataloging resource.

Smart City Data Exchange (SCDE): Architecture that supports a common approach to sharing city data assets and approved third-party data between cities and their data partners.

SCDE Authentication Services: Manages the identity verification of a user, device or system to a SCDE.

SCDE Authorization Services: Manages the roles, privileges and access to specific SCDE based on access control levels and policy management requirements.

SCDE Data Broker: City or third-party business entity that collects or enriches data and provides a value-added service to SCDE partners or consumers in accordance with local regulations and policy management requirements.

SCDE Data Consumer: A user interface, system or machine that searches and discovers relevant metadata from the data catalog.

SCDE Data Enricher: An entity that improves or creates value from raw data through analytics or automated means.

SCDE Data Producer: A user interface, device or system that creates or collects data, enabling the corresponding metadata to be registered and published to a data catalog.

Search: A structured process for identifying specific data from a broad set of data resources.

Smart City Data Catalog: Component of the SCDE that manages metadata for use by data producers, enrichers and consumers.

Swagger File: A file, typically in JSON or YAML format, built on Swagger open source software implementation allowing developers to act upon REST APIs.

Third-Party Data: Data produced outside of city-owned data assets by data partners and affiliated agencies.

2.2 Acronyms & Abbreviations

AI	Artificial Intelligence
API	Application Programing Interface
ATIS	Alliance for Telecommunications Industry Solutions
AutoCAD DWG	Automated Computer Aided Design Drawing
CSV	Comma Separated Values
ESRI Geodatabase	JavaScript Object Notation
GIS	Geographic Information System
Google KML	Google Keyhole Markup Language
JSON	JavaScript Object Notation
ML	Machine Learning

REST	Representational State Transfer
XLSX	Microsoft Excel Spreadsheet
XML	Extensible Markup Language
YAML	a recursive acronym for "YAML Ain't Markup Language

3. Functions

3.1 General Features

The functions of a smart city data catalog include producer-level attributes related to the publishing of a city's data assets and consumer-level capabilities associated with the discovery, searching and retrieval of relevant datasets. One of the value-added capabilities of modern data catalogs is data enrichment. Leveraging the growing availability of city data assets, data enrichment is the analyzing of datasets (single data source or multiple data sources from data producers) to draw conclusions about the information they contain and deliver that insight to data consumers. Therefore, data enrichers act upon the collected data to increase the effectiveness and usability of the data by consumers. In figure 3.1, the producer domain is shown on the left and the consumer domain is shown on the right:

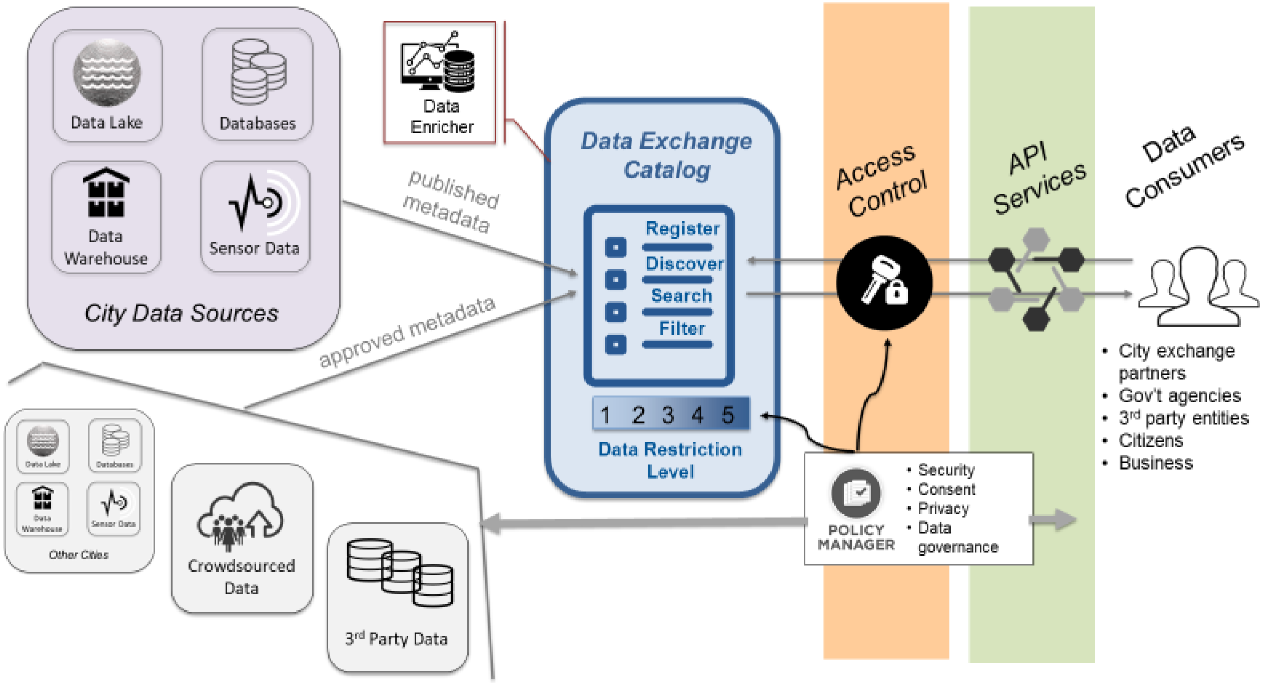

Figure 3.1 Data Catalog Architecture

As shown in this diagram, the smart city data catalog exists between the data source assets (city-owned data or third party) and the data consumers. Data enrichers may utilize analytics tools, ML and AI to enrich metadata through profiling, organizing, auto tagging and other capabilities.

To access the catalog, users are authenticated and authorized via access control in coordination with the policy manager, which enforces security, consent and privacy procedures. An API services layer provides a simplified view of the data catalog and facilitates the eventual data request processes.

The key features of a smart city data catalog include:

- Allows users/partners to discover and consume city data assets
- Integrates metadata and serves as a searchable metadata repository
- Includes semi-static metadata, which is likely to form the parameters by which search and discovery requests are made – i.e., geography (non-moving entities), entity type, key semi-invariant properties (e.g., opening times, capacity, unit measured)
- Manages/gates access to the source of the actual dynamic data
- Creates a connection path to city-stored data and approved third-party or crowdsourced data
- Presents a view of authorized data from a city's data warehouses or data lakes
- Applies a city's security, privacy and policy management controls (e.g., randomization, non-identifying data, privacy levels)
- Relates a library of available datasets to open APIs for simplified access

Data catalogs continue to evolve to open data registries that may incorporate cloud-based search and discovery capabilities. SCDE data catalogs can be supported with on-premises or cloud-based servers. Public cloud approaches to a SCDE data catalog may reduce smart city infrastructure requirements in some cases. Generally, the decision about how and where to host the metadata and associated datasets is part of a city's broader data management plan, which includes considerations about data security, privacy and resiliency, and third-party integration needs.

3.2 Registration and Data Publishing

This section describes the basic registration function for publishing metadata to a smart city data catalog. Through the authorized registration of data sources to the catalog, metadata becomes discoverable and searchable.

The registration process associated with publishing metadata to a catalog will include both manual and automated entry processes. During the initial establishment of a data catalog, it is understood that datasets will be identified and manually registered in many cases. As data catalogs become more heavily populated, both manual entry and automated movement of new and updated metadata into the catalog will be made possible through techniques such as machine learning and advanced analytics.

As the first step in registering data sources to a catalog, the user must be authenticated through access control functions and authorized to enter metadata resources at specific levels of access (e.g., open data, privileged, restricted, highly restricted).

As part of the registration process, object-level or structural metadata is extracted. This includes the data source location, file size, file type and certain data governance and data lineage information. In addition, descriptive information about the data asset is published to the catalog, such as keywords or identifiers that will relate the relevant data to a data consumer.

The process by which metadata is imported to a catalog will include the use of a registration open API, which may be utilized as part of a single registered data asset or can be implemented through an automated bulk registration application that uploads metadata in a standard format specified by the catalog.

An important set of additional data catalog features is the ability to preview and profile datasets, as defined by the publisher of the data asset. A data catalog preview may contain a section of a table, an example, a short stream of live content or any number of other data attributes that will demonstrate the dataset's value and relevancy to the user. A data catalog profile is additional information added to the metadata that more specially describes the use factors surrounding the data source. Common examples might include the file size, number of rows/columns in a table, number of entries, video format and when the information was last updated. Data profiles are commonly added to the search and filter fields to provide a higher quality user experience.

It is understood that datasets registered to a data catalog will include various types of temporal, time-series and event-driven data. Thus, it is critical that the metadata published to a data catalog and the related APIs provide a field for date or time value.

3.3 Using a Data Catalog

The primary objective for a data consumer, as a user of a data catalog, is the discovery of relevant data assets through searching, filtering and other API-assisted functions. Data catalogs give users optimized access to data assets from city and third-party data sources that are deemed to be sharable in a smart city ecosystem.

A user gains access to a data catalog through the access control domain, which authenticates a user's credentials and authorizes the appropriate level of access. Users will only be able to view metadata descriptions that are in their pre-assigned access control level. In most cases, users must be subscribed to a given data catalog to gain access to the relevant data.

As an initial step, users will search the data catalog using inputted search criteria (e.g., keywords), which will identify relevant data sources based on search parameters. Filter functions enable users to more finely identify data sources that meet their specific needs. Filtering attributes may include file size, file type, media type, data lineage and other factors. Open APIs will be utilized to establish common search and filter functions across smart city data catalogs that are compliant with this specification.

Additional capabilities may exist for a user to save search history and results, pin data sources for reuse or to profile a user's search criteria. A data catalog may also display related data assets that are associated with a selected metadata description.

3.4 Annotations, User Feedback and Ratings

One of the additional data producer benefits of a smart city data catalog, beyond registration, is the ability for those with knowledge of data sources to annotate additional descriptive information about the data source. This will normally include information about the data source or use of the dataset that is not captured as part of the metadata extraction or registration input. Examples may include tagging the data with the field of use or the metadata's context.

As is the case with social media and e-commerce applications, feedback, annotations and ratings have value. However, they should be managed and moderated by the smart city catalog producers to ensure information integrity is maintained across the ecosystem. Negative reviews are common, and while a majority of reviews are provided with honesty and accuracy, there are cases where such feedback is provided with alternative intent (e.g., sales, competition). Smart city data producers and enrichers must make sure that feedback, annotations and reviews are thoroughly assessed and applied in a manner that ensures consistency with a city's data governance policies.

From a user perspective, data catalogs can also act as a repository of user-initiated review and ratings on accessible datasets, including search patterns and history. Data catalog administrators can also benefit from collecting information about frequently searched or requested datasets. In some cases, data catalogs can allow users to rate the value of specific datasets or to insert review comments.

4. Technical Requirements

The SCDE data catalog as described in this document provides the foundation for the exchange of data resources between the data producer, data enricher and data consumer domains. This section addresses specific technical requirements for the functions that define the data cataloging resource within the SCDE. Data producers receive periodic reports, analytics and audits from data enrichers. There is a more frequent feedback loop between data enrichers and data consumers about demand for different datasets. Figure 4.1 illustrates the use and application of a SCDE data catalog across these respective domains.

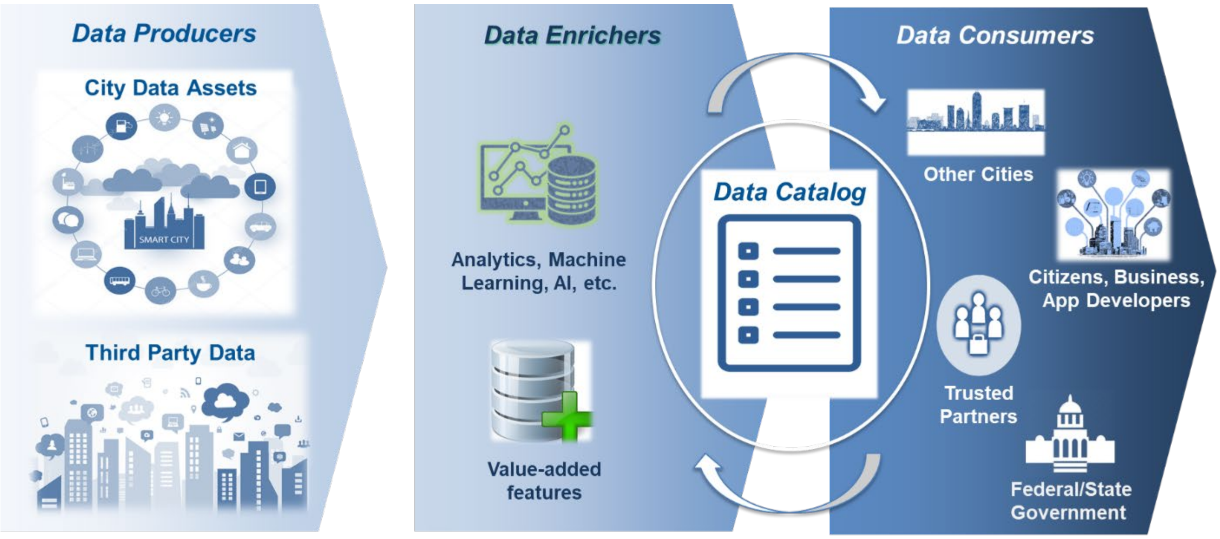

Figure 4.1 Data Catalog Domains

4.1 Data Producer Requirements

4.1.1 Supported Data Sources

City data assets will leverage a diverse set of data source types. These include data derived from sensors, databases, data lakes, data warehouses and data collected from approved third-party government agencies (federal, state, county and university) and negotiated arrangements with private data partners (e.g., financial, weather, transportation). Consequently, data stored as city-owned assets may exist in a broad set of data formats, including GIS shapefiles, CSV, XLSX, Google KML, ESRI Geodatabase, AutoCAD DWG, XML and REST endpoint.

The decision about which datasets are appropriate for data exchange will be ultimately left up to individual cities. However, the metadata published to data catalogs shall be consistent with the common data formats of JSON and XML, and may include other common data format types as listed above. In most cases, imported files (e.g., CSV) will be translated and formatted into the appropriate JSON file format for use in a data catalog.

Even though cities are conservative about sharing data to ensure privacy checks are in place, they have a wish and will to share as much data as possible to foster development. As the primary data producers, cities are also interested in exploring and collaborating with other third-party data sources (e.g., state, university, county) and negotiated third-party private sector data sources (e.g., credit card agency). This interest and collaboration highlights the need for standardized support for diverse data source types. The collaboration would eventually lead cities to become an active consumer of the data. Cities are also open to new data generation sources/pilots that can support futuristic technologies and can be integrated into their data portals and further enable sharing to the data exchange.

4.1.2 Policy for Data Sources

Most cities have their own regulated policy for data sharing, which may be governed by other regional authorities, such as county or statewide departments. City data management processes are at different levels of maturity, thereby resulting in a wide variety of data regulation policies. The diversity of policies supports the need for a standardized approach, including a set of data policy guidelines to ensure commonality between city data sharing regulations. Given that cities have adopted different data classification levels, it becomes more difficult for cities to create specific policies that relate to different categories of access. In most cases, cities revert to a policy that only allows sharing of data that has already been deemed public and been approved through their internal processes for publication. However, cities are increasingly working with third-party agencies to source city data. These projects are useful for understanding how cities will interface with data exchanges in the future, including how these third-party relationships will support data policy guideline development. This represents an important future area of work.

The positive aspect is that cities are open to emerging data policies. But they need clarity about the nature of classification of data interchanged with data exchanges. A data policy outlining the specifications would play a key role in this.

4.1.3 Registration of Data Sources

This section provides requirements for registering metadata into a data catalog from city data assets and third-party data sources, as well as the requirements for updating registration information.

A participating municipality will register each data source with the data exchange by providing the following properties (details in section 4.1.4):

- Owner::String
- OwnerId::String
- SourceType::String
- SourceSubtypes::String
- DataType::String
- Aggregation::String

- SourceId::String
- SourceLastUpdate::Datetime
- SourceUpdatePeriod::String
- Description::String
- ConnectionString::String
- ConnectionStringParameters::String
- *Owner* is a string of the name of the municipality and should be consistent across all the municipality's data sources.
- *OwnerID* is a string (or numeric) that is provided to the municipality upon registration with the data exchange.
- *ContactID* is a string that provides contact level information (department name or email) to understand source of data.

4.1.4 Object-Level Metadata

This section covers the metadata structure and object types for publishing and tagging data for use in a smart city data catalog, including indexing the metadata.

- *SourceType* is a string that specifies the primary source type (e.g., transportation, energy, safety). Primary source types are determined by the data exchange and are listed in section 4.1.3.
- *SourceSubtypes* is a delimited list of strings that distinguish subtypes of the SourceType. For instance, "train,stations" might be a SourceSubtype for the SourceType "Transportation." SouceSubtypes are freeform and can be determined by the municipality, though each municipality is encouraged to conform to standards set in the data exchange schema wiki.
- *DataType* is a string that specifies the type of data in the DataSource. Data types supported include string, numeric, datetime and point (lat/long pair).
- *Aggregation* is a string that specifies where in a chain of data aggregation the DataSource lies. Values can be:
    - "none" for no aggregation
    - "aggregates:<list of SourceIds>" to indicate which child SourceIds are aggregated
    - "aggregated:<SourceId>" to indicate the parent SourceId
- *SourceID* is a unique string (or numeric) provided to the DataSource upon registration.
- *SourceLastUpdate* is a datetime that is updated to the current datetime upon each update of the DataSource.
- *SourceUpdatePeriod* is a freeform string provided by the SourceOwner that indicates the frequency with which the DataSource is updated. Example include "hourly," "daily," and "weekly."
- *Description* is a freeform string provided by the SourceOwner that provides additional detail about the DataSource. Description is optional.

- *ConnectionString* is a string containing the URL to the API to query the DataSource.
- *ConnectionStringParameters* is a delimited string of filter parameters supported by the ConnectionString for the DataSource. Examples include “startdate,” “enddate,” “maxvalue,” and “minvalue.”

4.1.5 Descriptive Metadata

This section contains the metadata description fields that relate a data source to a catalog service, including attributes, keywords, captions and other information.

Metadata description fields relate a data source to a catalog service. In day-to-day usage, this metadata enables the data owner to provide contextual information for a data source and how to use it. It also can be searched and/or browsed by data exchange users for data discoverability and verification.

Descriptive metadata fields will include:

- Owner::String
- DataType::String
- SourceLastUpdate::DateTime
- Description::String
- Keywords::String
- ExtendedDescription::String
- ExampleQuery::String

Owner, DataType, SourceLastUpdate and Description are drawn automatically from the data source properties (see section 4.1.4) that are provided upon registration of the data source (see section 4.1.3). The exception is SourceLastUpdate, which is time stamped with the most recent update of the data source.

Keywords is a free-form text string of delimited keywords pertaining to the data source and/or the catalog service. ExtendedDescription is also free-form text and is made available so that data source owners can provide additional detail about how their data source relates to its catalog service. This may include details like contact information for the catalog service, relationships to other datasets and services, example queries and any other details, such as update frequency and the service’s version and life cycle information.

4.2 Data Enricher Requirements

4.2.1 Profiles and Previews

The data enricher provides the ability to:

- Create and onboard as an enricher partner

- Partner to create software offerings
- Allow the partner to limit or restrict who may access the software offerings
- Enable a partner to establish business terms for consuming offered insights and/or software
- Ingest data from partner data sources through a well-defined interface to that data source. The data could come from a cloud resource, a product resource or a device resource
- Use the offering interfaces to the exposed partner data to create a visualization dashboard
- Allow the offering interfaces to the exposed partner insight and/or software to be consumed by third-party applications to produce insights and additional data streams
- Throttle and managed the consumption of data and/or software via the offering interfaces
- Use an automatic reminder system that summarizes live datasets and reminders to upload new data when it becomes available

4.2.2 Usability and Tracking

Each de facto standard would then be further developed by different standardization bodies into proper standards that everyone can follow. Historically this has not been easy because even minor adaptations are costly to update, but the API-led connectivity streamlines this work. In the process, it enables a faster spread of the system as more and more organizations follow the same path.

Data can be consumed from any platform, with adapters for commonly used interfaces for IoT. Insight Creators are provided access within a community and are encouraged (or enforced, if required) to deploy code as containerized microservices to avoid the privacy risk of extracting data from the community. Finally, a marketplace exposes both raw data and insights to Insight consumers, with sharing as a portion of all transactions on the exchange.

Data enrichment provides an important level of value creation beyond basic data sharing. However, it is also important to have established policies and processes to recognize, manage and correct misrepresentations or erroneous contextualization of municipal data (either aggregated or processed on its own) that is outside the city’s control. Ideally, any enrichment of municipality data should be moderated or have the ability to be moderated by the data providers. Consequently, data producers and data enrichers must have clear and well-developed policies for reviewing, representing and correcting municipal data that is associated with a city’s data exchange.

The capabilities to be exposed and supported should include:

- Description of the data offering interfacing with IoT, robots and third-party systems for capturing updates about context information and translating required actuations.

- API Input method including API Host Name, API protocols, API usage policy and API support contact.
- API Details including Swagger file.
- Business terms of using the data - Context Data/API management, publication and monetization, bringing support to usage control and the opportunity to publish and monetize part of managed context data.
- Throttling Policy.
- Access control datasets matching right-time context data, the assignment of access terms and policies to those datasets, and the assignment of pricing and pay-per-use schemas to datasets.
- Test environment for the API Processing, analysis and visualization of context information implementing the expected smart behavior of applications and/or assisting end users in making smart decisions.
- Publishing possibilities for enriching the data based on single or multiple data sources.

4.2.3 Automated Populating and Tagging of Data Catalog

Populating metadata into the SCDE data catalog typically is a manual process, particularly in the early stages of implementing the cataloging function. As the volume and diversity of data sources increase, automated populating of metadata can vastly improve the velocity and usability of data for the data consumer domain.

For data that is derived from third-party or crowdsourced domains, automated populating of data will require an initial review and ongoing auditing to ensure that the data meets established data governance practices. Third-party data that is acquired as real-time data and approved for automated population will require ongoing validation of its adherence to established data sharing and governance policies.

Metadata that exists in the SCDE data catalog often includes native (or descriptive) metadata that is built at the data's source or storage level. When implemented as part of a manual function, tags are selected from a glossary of tagging options. Alternatively, tagging of metadata can occur as part of the data enrichment function. As data catalogs mature, automated tagging can offer significant benefits, both in terms of rapidly moving metadata to the catalog and creating a consistent view for the data consumer. These tags can also leverage feedback and ratings from the consumer domain to adjust the automated tagging to meet consumer needs and profiles.

ML- and AI-based automated tagging relies on a set of prescribed rules and policies for discovering relevant data, determining the level of access control (open, limited, restricted) and moving the metadata to the data catalog. In some cases, local policies may allow only data defined as "open data" to be tagged and published to the catalog. The benefit of using AI-supported data enrichment as part of the automated tagging function is the ability to learn and execute the appropriate match between types of data (e.g., images, names, addresses, numerical fields, etc.) and the appropriate metadata tag, without ongoing manual intervention.

Automated populating and tagging of metadata within the SCDE data enrichment function includes the following features:

- Accurate recognition of the data source and data asset
- Verification of the data lineage
- Rules-based application of tagging template
- Pre-determined glossary of tagging definitions
- Enforced conformance to local policy management rules and procedures
- Determination of appropriate sensitivity classification (e.g., open, limited, restricted)
- Discovery of relationships to other data assets that may offer value to the data consumer
- Optimization of descriptive metadata based on user feedback and ratings

4.3 Data Consumer Requirements

4.3.1 Discovery Requirements

This section describes the mechanisms and processes for discovering metadata that has been tagged and categorized. Discovery will be supported through the use of open APIs.

A metadata discovery API (e.g., http://scde.com/api/discovery) will return all metadata categories, including their data type:

- String
- Date
- Numeric

And the way in which they can be searched/filtered:

- Strings will be an exact match or a partial match
- Dates will be equals, before or after
- Numeric will be equals, less than or greater than

As a rule, each data source and object level metadata type (see sections 4.1.1 and 4.1.4) will be returned by the discovery API. The metadata discovery API will return all results in JSON format.

Each metadata category JSON object will also contain a list of values available. For instance, the SourceID metadata object will contain a list of strings of all Source IDs, the SourceType metadata object will contain a list of strings of all Source Types and so on.

4.3.2 Search and Filter Requirements

This section covers the mechanism to effectively search datasets and streaming data sources, and filter the most relevant data. This includes semi-static metadata, which is likely to form the parameters by which search and discovery requests are made: specifically, geography (non-moving entities), entity type and key semi-invariant properties (e.g., opening times, capacity, unit measured).

In common practice, metadata will be searched and filtered to narrow down a desired result set. For instance, an end user could issue an API query specifying string matches as follows: http://scde.com/api/metadata?SourceID=Portland&SourceType=Transportation. Such a query would, of course, return transportation-related metadata from the City of Portland, including the connection strings so that the user could then issue data queries.

To facilitate such API access, the data exchange must support relational-database-style querying in accordance with the matching rules specified in the discovery API: exact or partial match for strings; equals, before, after for dates; and equals, less than or greater than for numeric.

4.3.3 Ratings and Review Requirements

This section provides the requirements for data consumers to review and rate the accuracy and usability of datasets and streaming sources of data, including the metadata descriptions.

To facilitate consumer reviews and ratings of data exchange data, each metadata source and metadata object will have fields for reviews and ratings. These fields will contain lists of strings and numeric values, ordered chronologically, containing the reviews and ratings respectively, along with the ID of the consumer who posted the review or rating and date when submitted:

- *SourceReview* is an array of freeform strings, ordered chronologically, with each new review appended to the end. In the JSON that is returned to show these reviews, the reviews will be an array of objects. Each object is a pairing of the review text, ID of the consumer submitting the review and datetime stamp of review submission.
- *SourceRating* is an array of numeric rating scores, ordered chronologically with each new review appended to the end. In the JSON that is returned to show these ratings, the ratings will be an array of objects. Each object is a pairing of the numeric ratings, the ID of the consumer submitting the rating and datetime stamp of rating submission.
- *ObjectReview* is the same as *SourceReview* but for objects.
- *ObjectRating* is the same as *SourceRating* but for objects.

The data exchange will support two mechanisms for adding reviews and ratings:

- A web form interface containing a freeform text entry box for review and dropdown of numeric values for ratings. This page must exist for each source and object.
- An API call for programmatically adding reviews and ratings.

4.3.4 Data Request Requirements

This section integrates the data catalog discovery, search and filter requirements with the data request function, in cases where the data asset is requested via the data cataloging service. This includes managing the access to the source of dynamic data.

Access control: Data requests issued to the data catalog will adhere to one of three access control levels:

- Public: Datasets that are publicly accessible by any user of the data catalog
- Trusted: Datasets for which the user is a trusted partner
- Onetime: Datasets to which a user has been granted onetime access by the dataset owner

Access control levels will be established by the dataset owner upon publishing the dataset to the data catalog and may be edited at any point. Trusted partners are also specified by the dataset owner at any time, including upon publishing of the dataset to the data catalog. Requests can be made by any user for permission to access non-public datasets, either as trusted partners or as onetime users. Requests are sent to the dataset owner for evaluation, at which point the requesting user can be made a trusted partner or be given onetime access or be denied access.

Access rights for trusted partners and onetime users will be maintained by the data exchange service. For trusted partners, the service must check the credentials of the logged in user against a list of trusted partners for the dataset being requested. For onetime users, for each dataset the service must maintain a store of users who have been granted access but have not used their onetime credentials. Once the user has used onetime access, their credentials are removed from the list of onetime users with current access permission to the dataset.

In everyday usage, non-public datasets will be marked as Restricted. At that point, the user can specify that they are either a trusted partner or onetime user of the dataset, or they can initiate the request for access. Users who are not logged in will be asked to do so in order to access restricted data and/or request access.

4.4 Access Related Requirements

4.4.1 General Requirements

This section covers the access control requirements associated with the data catalog where metadata is presented and segmented in accordance with prescribed levels of access.

Authentication and authorization requirements associated with the use of a data catalog include the registration of metadata to the cataloging function, the application of data enrichment and the access to metadata by data consumers.

Within the scope of this document, authentication is used to validate the user's identity and credentials and must adhere to the policies established by the local municipality.

Authorization is the process of allowing access to specific resources based on the roles and privileges of the authenticated party. Authorization requirements are governed and enforced by the local municipality's policy management practices.

Within the producer domain, the means by which cities and local governments authenticate users within their own IT domain is considered out of scope for this document. Authentication requirements associated with data producers and data enrichers who operate as third-party partners are discussed later in this document.

Authentication requirements associated with the consumer domain depend on the classification of metadata as either open data or restricted data and are treated below. Cities will choose to categorize data access types per local data policies and appropriate open disclosure obligations, but the SCDE requirements contained in this document do support multiple accessibility options. In some cases, cities may only choose to exchange *open data* types as part of metadata published to a data catalog. Open data is defined as publicly accessible data with no SLA and no personally identifiable information (PII).

In other cases, cities may support additional levels of access control based on an appropriate level of restriction that is defined by local policy management controls. Additional classifications may include *limited to approved third-party arrangements* and *restricted to official government use*. In these cases, additional access control restrictions are required and should be calibrated to the level of SLA applied and to the degree of PII that may exist in specific circumstances.

In addition to the accessibility requirements stated above, cities may choose to apply access control to data based on duration aspects. Examples include continually available or limited to a specific timeframe of availability, such as datasets replaced on weekly basis.

A third delineation of access control may include a classification of data as regulated or non-regulated. This classification will depend on local policies and laws as well as state or regional obligations.

4.4.2 Authentication Requirements

This section addresses the basic set of identification and authentication requirements to gain access to a data catalog.

Authentication services provide verifiable proof of identity associated with a client or a server. Authentication does not provide access to a prescribed set of permissioned tasks. It is generally recommended that X.509 certificate-based processes or token-based solutions be used for mutual authentication.

In the first case, the certificate is signed with a secret private key and validated with a known public key. Certificate-based approaches provide significant benefits over “shared secret” authentication mechanisms.

Alternatively, token-based authentication processes would typically use standards such as JSON web tokens. In this case, the same key is used for the client and the server. Generally, token-based processes would use the following flow:

- User requests access to resources with a username/password.
- The application validates the credentials.
- The application provides a signed token back to the user’s client.
- The client has the ability to store the token and associates it with every request.
- The server verifies the token and returns the requested resource to the client.

In the producer and enricher domains, where third-party entities require access to the data catalog to register or enhance metadata, cities will require authentication of the client or server before posting, changing or deleting any metadata resources from the catalog. This guarantees that any data residing on a data catalog has been approved by the city or local government as part of the data governance.

In the consumer domain, authentication is required to access any stored metadata that is not deemed “open data” by the entity managing the data catalog. Metadata that exists at any restricted level of access should follow appropriate authentication processes. Upon verification of the digital certificate for authentication services, the authentication server shall share the verification status with the authorization services, so the permissioned view of the data catalog correlates with the level of authorized privileges.

4.4.3 Authorization Requirements

This section covers the administration of role-based security and privileges for a subscribed user or partner to gain access to specific levels of restricted and segmented metadata.

The authorization server operates in accordance with local access control and policy management and enforcement requirements and grants a set of permissioned actions to the provider, enricher and consumer domains.

In the case of a producer or enricher, authorization must be granted to register, modify, update or remove data to the data catalog. Authorization privileges must be consistent with the granted level of access given to any user or entity.

In the consumer domain, authorization levels to discover and utilize metadata are granted by the authorization server and conditioned on the classification of metadata as open or restricted level(s). These privileges allow a consumer user or entity to view, search, filter or preview metadata units. It is recommended to utilize OAuth, JSON Web Token or an equivalent open authorization protocol to authorize all authorization flows between the consumer and the operator of the data catalog.

It is recommended that the provider and enricher utilize a consent request and acknowledgement action to the consumer entity in accordance with local consent and licensing policies.

5. Business Requirements

5.1 Data Governance

Data governance is designed to maximize the usability of the data with the need to protect the integrity and appropriate accessibility of data throughout its lifecycle.

The basic components of data governance include:

- Understand the data
- Cleanse and curate the data
- Profile the data
- Apply appropriate controls to its usage
- Catalog and enrich the data
- Audit the quality and usefulness of the data

Data lineage allows data owners to track the usage of city data assets and third-party-generated data, from the source to its usage, including the various steps of data enrichment. In the smart city, data governance must balance the need to maintain the integrity, security and privacy of the data sources with the usability and accessibility needs of the consumers of the data.

The following is a list of data governance attributes that must be developed in support of a data catalog:

- Ownership of the data throughout its lifecycle, including third-party and crowdsourced data. While ownership is subject to legal, regulatory and ethical requirements within a given municipality, ownership and control of data must be well understood and documented within a data governance plan.
- Administration and operation of the data catalog and the metadata that is published to the catalog platform. Catalogs may be administered and managed within the city's domain or a third-party partner that may oversee its operation on behalf of one or more cities. It either case, a clear data governance strategy is critical.
- Data security, integrity and privacy controls must be developed across all of a city's data partners. They also must be integrated with data access and policy management requirements, covering the entire lifecycle of the data, including data enrichment.
- Specific governance requirements for administering real-time data, originating from city-owned assets (e.g., sensors, cameras) or published via third-party data assets. Dynamic data promotes the need for appropriate randomization and anonymization of streaming data with respect to local law and regulations.

5.2 Third-Party Data

This section covers the common requirements associated with publishing and utilizing third-party data across the functions of a data catalog. It is understood that cities may have specific requirements for third-party data in addition to a common set of business requirements.

Data catalogs enable the consolidation of approved third-party data with city data assets. Cities gain value from integration of third-party data because relevant data can be offered to smart city data consumers without the need for the municipality to collect, process and store data that is derived from affiliated data partners. One of the key benefits of a standardized data catalog approach is that all incoming third-party data can be cataloged and indexed according to a common glossary and taxonomy.

In assessing future business models for integrating third-party and city data assets, a corollary approach to city-operated open data platforms may exist where data is licensed from cities to third-party data marketplaces and blended with other data sources. Although those arrangements do not replace city open data platforms, they could promote monetization of data, creating new revenue opportunities for cities. While data could be simply replicated on third-party data marketplaces, some cities may choose to register their metadata more formally in third-party data catalogs, following common procedures outlined in this document.

The integration of city-owned and third-party data highlights the future need for data stewards. These would ensure that data collected via any city asset or city-owned infrastructure is shared and utilized in accordance with a municipality’s policies for transparency and open data obligations, balanced with appropriate ethical policies such as privacy. This requires a data licensing framework to ensure that city-owned data assets used in conjunction with third-party data assets meet an appropriate set of legal and ethical standards within a municipality.

5.3 Crowdsourced Data

Similar to third-party data, this section covers the requirements associated with conveying crowdsourced-based metadata to a data catalog.

Crowdsourced data is generally defined as creating a dataset through the collective resources of a group of people. The benefits of crowdsourced data for smart cities are generally twofold:

- For specific types of data, municipal organizations can collect data from crowd resources at a much lower acquisition cost than city-owned assets.
- The data received from crowdsourcing provides incremental value to data that is collected by city-operated sources because of specific attributes, such as

geographically dispersed data or real-time data generated by a large population of sources.

While the benefits of crowdsourced data can supplement a city’s open data, there are additional challenges of leveraging resources that are collected from a large volume of external contributors. The data must be validated, which generally requires that an appropriate sample size exists and that the crowdsourced resources contribute data on a fair and ethical basis. These factors contribute to the need for cities to periodically sample, test and validate the integrity of crowdsourced data that is published to a data catalog.

While the use of crowdsourced data in areas such as research activities is well accepted, publishing data from real-time contributors as part of a city’s open data framework is more challenging and will likely require the use of ML and AI to validate the data and its contributor resources. These capabilities can be incorporated into the data enrichment aspects of operating a data catalog.

In the context of crowdsourcing for smart cities and the publishing of crowdsourced metadata to data catalogs, two options have evolved and are generally viewed as being extensible for the future. The use of crowdsourced apps as a means of gathering data from contributors is directly relevant to a city’s open data needs and plans. These apps either provide permissive access to a crowdsourced data or rely on a user to provide a response to a specific subject matter.

In the first case, data personnel within city government will generally need to extract data from app contributors (e.g., traffic flow, road conditions, event-driven data) and provide a collective view as part of a dataset published and updated via a data catalog. In this case, cities will need to ensure that the data is sufficiently randomized and reliable.

Alternatively, crowdsourced platforms could be employed by either cities or third-party partners to gather, analyze and provide relevant datasets for publishing to a smart city data catalog. One of the benefits of leveraging crowdsourced platforms outside of city operations is that the platform investment is shared between cities and other parties who will find value in the contributed information.

5.4 Data Brokers

Data enrichment is the process of analyzing datasets (a single data source or multiple data sources from data producers) to draw conclusions about the information they contain and then deliver that insight to data/insight consumers. Not every data collection entity can analyze data on their own, but the use data analytics and AI to make more-informed business decisions and increase the effectiveness of their efforts is one solution for data/insight consumers to get actionable insights.

Businesses look for answers more than they look for data. This requirement has led to the formation of many new businesses – called data enrichers – whose sole purpose is providing data analytics and consultancy as a service. It is difficult for any entity whose primary focus isn't data collection, to store such ever-increasing data. This has led to the formation of another business model which focuses on providing data collection and handling tools analysis and visualization, i.e., data brokers.

In the general sense, data brokers have most often been associated with the collection of personal information, behaviors or preferences, and the reselling of this information to other parties. In the case of smart city data exchanges, data brokers may operate within a constrained space, where aggregated and fully anonymized data can be collected through data partnership agreements with cities and possibly combined with other third-party data to offer value-added services. This requires that all data provided to a data broker is approved by the municipality and that the data broker fully discloses the intended use of the data. While it is possible that data brokers could collect some level of data through a smart city's open data portal, there may be additional levels of aggregated data where a city and a data broker partner can find mutual benefit and citizens can realize additional value from the formalized licensing of data to the smart city data broker.

5.5 Privacy and Consent Policy

This section addresses the privacy control and enforcement policy requirements such as anonymization and randomization of metadata and data sources across the data partners.

Across any business, enterprise or city, an increase in the data's collection, storage and application correlates with increasing need to maintain privacy and informed consent processes. Understandably, smart cities have acknowledged the need to undertake privacy and consent policies that protect personal data and adequately address privacy needs with the exponentially increasing level of data that is collected by sensors, cameras and other devices.

Public spaces present a different privacy and consent environment than smartphone apps or online experiences. Data that is collected by street-level sensors or cameras does not provide an opportunity for opt-in consent as in the case of an ecommerce transaction. In some respects, data that is anonymized or randomized at the data source presents less privacy violation potential than data that is stored on a centralized basis. Nonetheless, privacy at any point in the ecosystem is equally critical. Generally, the approach has been to remove any PII as close to the source as possible and to present only anonymized information.

The application of smart city data catalogs must adhere to the overall privacy and consent policies of the local government. This requires that all metadata registered and published as part of the cataloging function must meet such policies and that any data that is identified, previewed or profiled on the data catalog must also meet all applicable privacy and consent policies. Examples of such smart city applications may include surveillance cameras, license or transponder readers, location services and applications that monitor social media feeds.

With the advent of GDPR, state privacy laws and the potential for federal privacy legislation in the near future, data catalogs must be adaptable to the applicable laws and regulations. This will require a strong commitment to tracking the data lineage and enforcing restricted access to data and strict adherence to authentication and authorization levels.

5 .6 Policy Management

Policy management associated with a data catalog requires strong coordination and alignment with the overall data governance structure of the smart city. Policy management becomes the intersection point for all policies that relate to security, authentication, authorization, data privacy, consent and licensing of data within the smart city. Significant value can be gained from policies that are clear and extensible for the future. Data catalogs are the data consumer’s window to a city’s data assets, so policy management and enforcement at the data catalog becomes critical in managing the smart city’s data resources.

As discussed earlier in this document, data exchanges may be deployed as a series of local exchanges that form a federation, or as a regional exchange that covers multiple municipalities. The application of policy management becomes more complex in a regional data exchange where multiple municipalities in a geographically connected region share a common data exchange and publish metadata to a merged data catalog. Whether managed by a regional partner or a third party, the data catalog must operate under an agreed upon set of policies, covering all aspects of publishing, enriching and consuming data.

As data catalogs within an organization move from manual registration of datasets by data stewards or data scientists to automated populating and tagging, policy management must also evolve to support processes that rely on ML, AI and other techniques. Consequently, policy management must be integrated into the systems that manage the automated movement and characterization of metadata into the city’s data catalog.

Acknowledgements

AT&T
Bell Canada
C-Spire
CenturyLink
Cisco
City of Austin
City of Chattanooga
City of Colorado Springs
City of Denver
City of Eugene
City of Independence
City of Kansas City
City of Las Vegas
City of Portland
City of San Diego
City of St. Petersburg
City of Vaughan
City of Virginia Beach
City of Washington, D.C.
Comcast
Ericsson
Fujitsu
GE
iconectiv
Intel Corporation
InterDigital
Microsoft
Oracle
Qualcomm
Sprint
TELUS
Verizon