City Operating System and City Intelligence Layer

July 27, 2023

Revision History

Date	Version	Changes
14 June 2023	1.0	Initial document
27 July 2023	1.1	First release

Table of Contents

1. Introduction

It is well known that the lack of availability of data within cities presents a major challenge. As a result, the concept of a data exchange has emerged and such exchanges are increasingly common within cities as a fundamental enabler for data sharing. However, a data exchange is not enough. In this paper, we argue that a data exchange is only one aspect of a "City Operating System" (COS) that should become the basic building block within the IT infrastructure of cities. All ICCC and other IT systems in cities must be built upon a COS to facilitate rapid development and portability of applications and services. We discuss what comprises a full COS and present the requirements of each component. One key component is a City Intelligence Layer (CIL) which makes derived insights about the city available through standardised API's in addition to the raw data. This can substantially reduce the time-to-market for new applications and services.

The core idea behind COS is to provide standardised access (APIs) of raw and cleansed data from a variety of sources, as well as their various "views," including processed data. This will enable a baseline of curated city operational services to be created for administrators and make it easy for application programmers, dashboard builders and analytics-as-a-service providers to build applications, visualisations and other services that benefit city administrators (e.g., for monitoring and decision making), and citizens (e.g., for planning and utilisation of services). The applications will be able to invoke and combine the APIs seamlessly in different ways to produce useful visualisations and advanced solutions. The availability of cleansed and processed data makes it fast and simple for application developers to build diverse analytics functions. This is because the cleansed and processed data saves the application developers the time and cost of accessing, downloading, and processing huge amounts of raw data.

Various aspects of city management and urban planning are interlinked and cannot be managed effectively in silos. The COS, along with its City Intelligence Layer (CIL), provides standardised interfaces to access data and the associated metadata and will facilitate the use of data across silos. It will enable seamless and efficient management of various aspects of urban life for the stakeholders, such as city officials, delivery companies, transportation aggregators, law enforcement authorities, disaster management agencies, and ordinary citizens. More generally, the proposed effort will play a crucial role not only in "urban informatics," but also in data-driven solutions in many other domains such as agriculture, energy, and water management.

As shown in Figure 1, the COS is an encapsulating layer which wraps the critical Data Exchange entity as well as the City Intelligence Layer (CIL). The CIL provides domain-specific analytic data for various domains and for various classes of use cases. We envision the COS as an intermediary infrastructure between the Data Layer, which may consist of middleware platforms (e.g., for IoT) and the Service Delivery Layer, which directly interfaces with Dashboards and Command and Control Centre displays. The existence of the Data Exchange Layer acts as a data standardisation interface for domain-specific analytics of the City Intelligence Layer. The COS becomes the interaction interface over which dashboards can be built. In addition, applications and business use cases can be built using the standardised APIs of the City Intelligence Layer and the Data Exchange Layer. The COS will provide data access, under appropriate access controls and authorizations, to third party application developers. Thus, it will play a crucial role in seeding an application ecosystem for solving complex urban problems grounded on digital innovation.

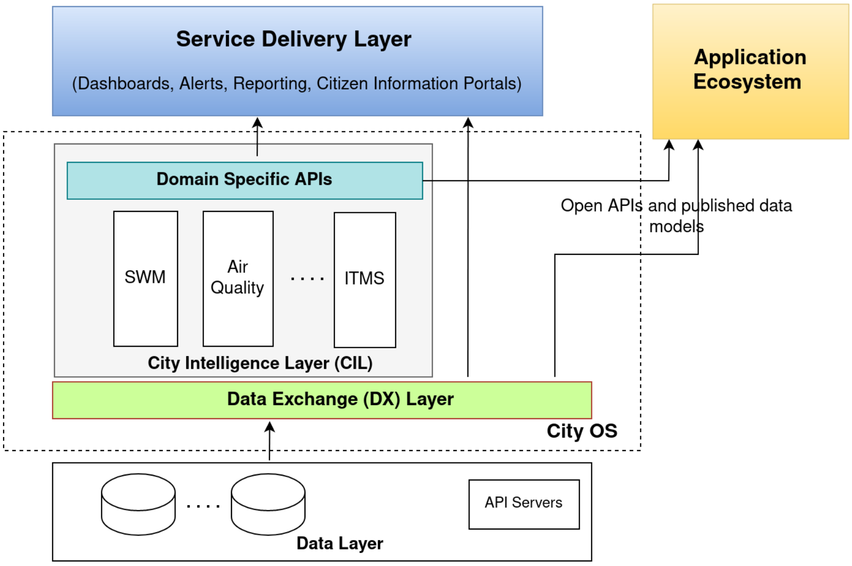

Figure 1: Architecture Overview of the COS in City Deployments

The COS will be deployed in a federated manner as shown in Figure 2, possibly with each smart city containing a COS. The Data Exchange Layer of the COS acts as a data backplane between multiple COS's and provides for aggregations at state levels and national levels. Various data collecting systems such as IoT systems, citizen grievance aggregators, etc., can be directly interfaced with the Data Exchange Layer of the COSs. In other words, the nature of data sources may be different in each city, and the use cases and dashboards may be different in each city, but all the COS's will all have common APIs at the output level, so that various types of aggregations are possible. Centralised access control based on data sharing policies may also be promulgated at the national level.

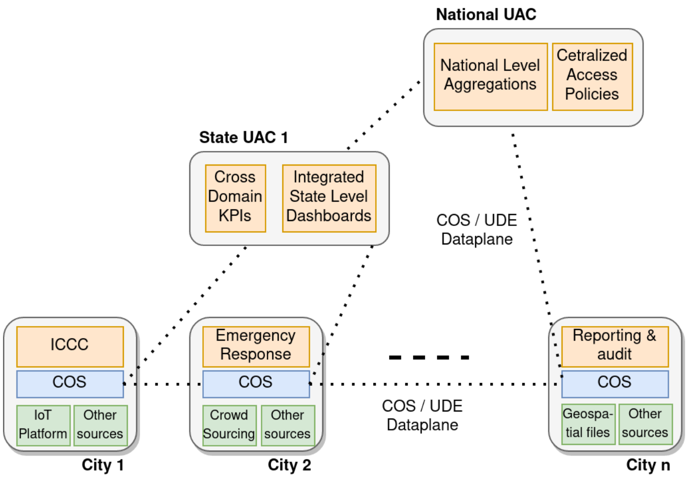

Figure 2: Federated Deployment of COS

2. Usage and benefits

COS will significantly improve and streamline the process of building solutions for data-driven planning and decision making in many domains. For example, as shown in Figure 3, a typical ICCC RFP Lifecycle for City Administrators will involve specifying (i) the requirements (i.e., identifying the data sources and the end use cases), (ii) the solution process (i.e., the problems to be solved, the methods and algorithms to be used, and the solution providers they may entail), (iii) the visualisation/delivery aspects (namely, the parameters to be displayed, the look and feel of the interface or dashboard, and the mode of delivery to the end user who may be the general public or a line officer), followed by (iv) raising tenders for final consumer application development . The execution of the entire lifecycle of an ICCC is greatly protracted by the existence of these four phases, each with its own procedural delays and considerable cost overhead as different solution providers may be required for the execution of each of the life cycles.

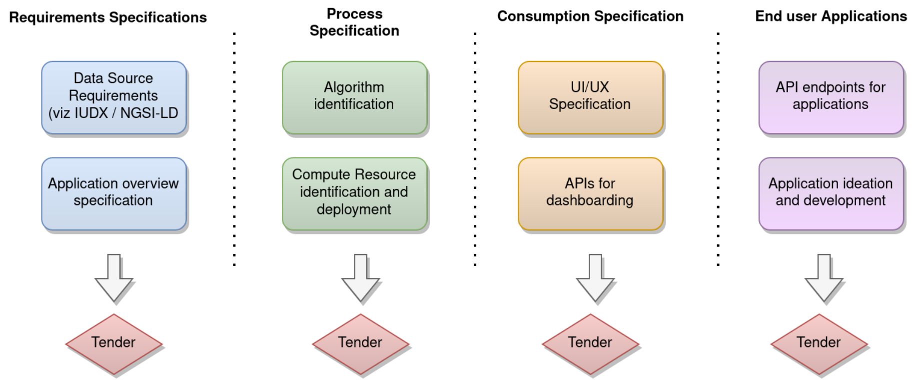

Figure 3: City RFP Lifecycle Pre-COS

COS will amalgamate the four lifecycles into two lifecycles as shown in Figure 4, namely, (i) deploying the City Operating System, and (ii) development of end user applications.

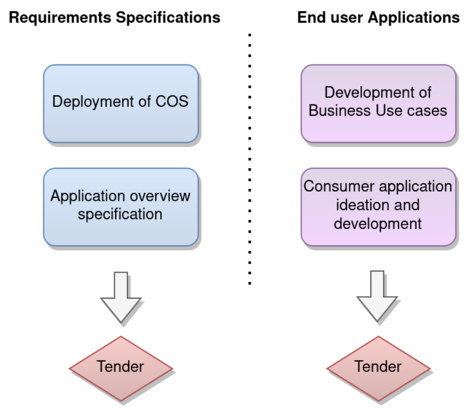

Figure 4: City RFP Lifecycle with deployment of COS

Owing to the reduction of turnover time for turnkey domain specific intelligence deployments, and the availability of key analytics building blocks from the COS directly, city officials can now shift their attention and resources to more impactful business use cases such as improvement of quality of services and optimization of revenue and expenditures, as shown in Figure 5.

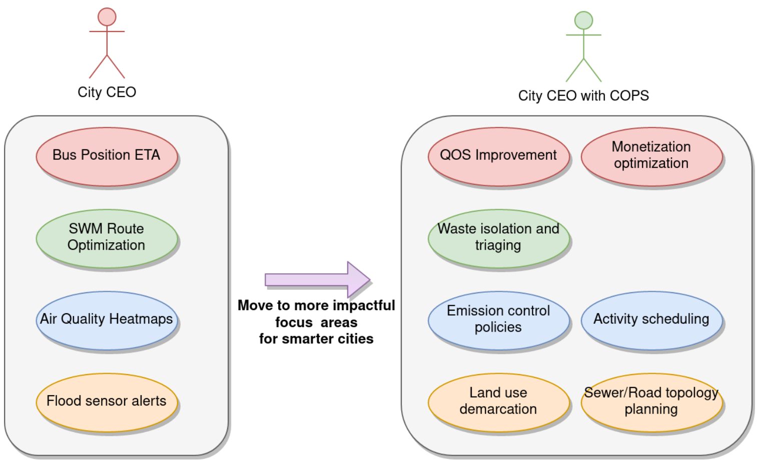

Figure 5: Use case identification for City CEO with COS

More generally, the COS along with its inbuilt City Intelligence Layer (CIL) caters to a wide array of dashboarding and application development requirements. Previously as shown in Figure 6, application developers would have to develop a plethora of intelligence algorithms for developing end user applications, whereas with the COS/CIL the application developers can choose from a variety of (processed) input sources, combine them in various ways by using the APIs, and feed them into any of the standard algorithms or methods available (again as APIs), which will in turn provide a variety of outputs through APIs. The application developer can also choose a suitable set of outputs to funnel into desired user interfaces. For example, in the context of a transit management system of a city, without COS/CIL, application developers would have to build many components, as shown in Figure 6. However, with COS/CIL, many APIs for “Estimated Time of Arrival,” etc. can be directly used by application developers to build a multi-modal transit application very easily (see Figure 7). As a result, the application developer can now focus on the parts of the application that will differentiate the app from the competition (such as personalised optimization and offer suggestions for the end user). Such efforts have been piloted in the UK [5] and have resulted in similar successes.

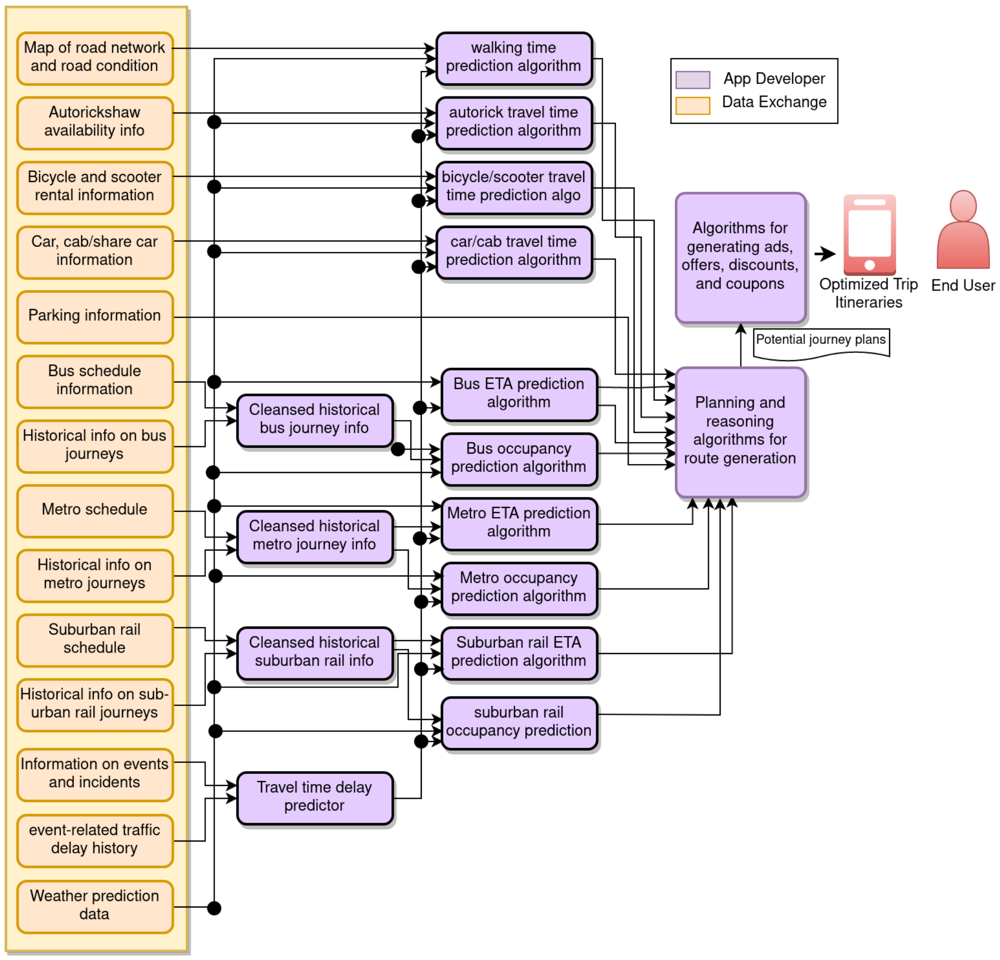

Figure 6: Application development considerations without COS/CIL

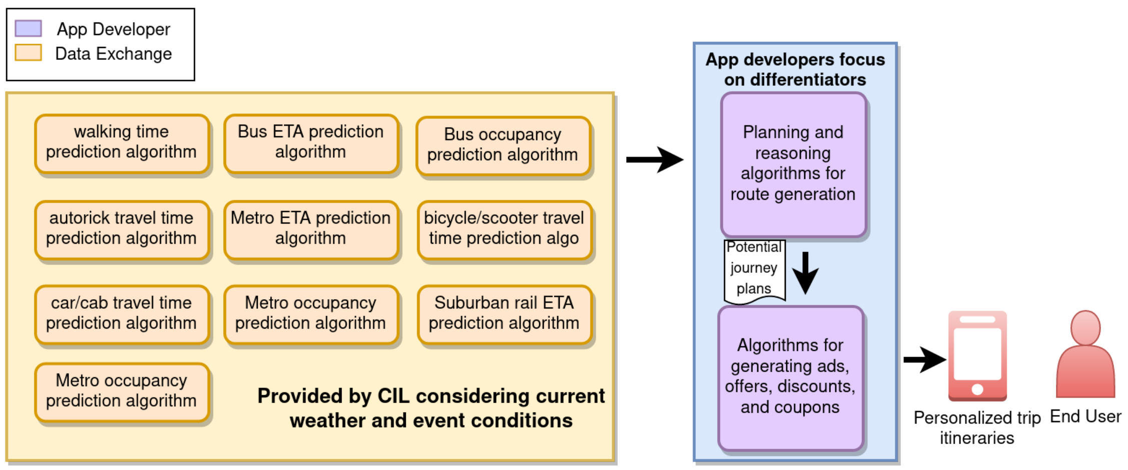

Figure 7: Application development considerations with COS/CIL

There is also another important aspect of the APIs that provide cleansed and processed data from an application developer's perspective. Without COS/CIL, the application developer would have to download huge amounts of raw data and spend computational resources to pre-process the raw data to generate inputs (such as expected time of arrival or occupancy) that go into the generation of the final recommendations to the end user. On the other hand, with COS/CIL, the developer would save on data downloading/storing cost, computational cost, as well as algorithm development cost of the pre-processing components.

Intelligence provided by the CIL is developed through a consolidated effort by stakeholder organisations participating in the collaborative development of the intelligence layer. The Intelligence adheres to commonly agreed upon terminologies and data formats for the data sources and the output data. Figure 12 shows an example of the agreed up semantics for data sources, analytics processes and data consumption formats.

The CIL will be developed bearing in mind the interoperability and reusability of intelligence across cities and across similar domains. (See also the discussion on the federated model in Section 1.) This will also foster the development of cross-domain aggregate analytics and lead to advanced business use cases as shown in Figure 8. In this example related to transit management systems, key performance indicators such as the regularity (is the bus on time?) of transit fleets can be compared across two cities. Also, on a more advanced cross-domain level, the financial performance of buses and metros can also be compared since the underlying domain connotations are the same.

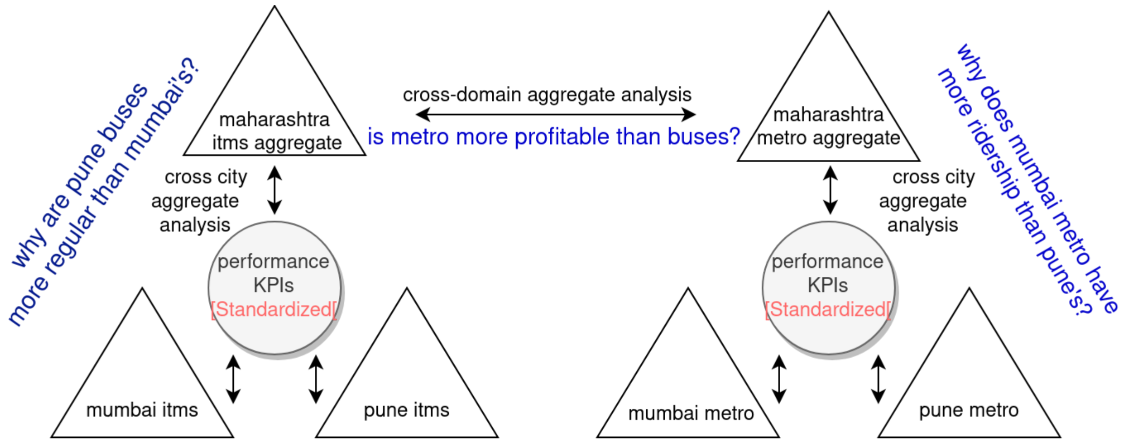

Figure 8: Advanced Cross-Domain / Cross-City Aggregations enabled by COS Standardization

3. Target domains and API Examples

Few of the initial target domains and domain specific intelligence that the CIL will be catering to are listed below:

1. Air Quality Management Systems: spatio-temporal interpolations, alerts, and hotspot identification.
2. Intelligent Transit Management Systems: Estimated times of arrival, traffic hotspots, fleet performance and status, and deviations from scheduled routes.
3. Solid Waste Management Systems: Daily dumping estimations, dumping zone identification, dumping trend anomalies, daily routes and schedules based on trends, anomalous surge handling, and dumping redistribution recommendations.
4. Flood management systems: Flood alerts, vulnerability analysis, storm water drain flow, and level status updates and alerts.
5. Weather: Spatio-temporal interpolations, weather forecasts, wind maps, heat-island maps
6. Citizen Grievance: Ward-wise top grievances, and top grievances.

CIL APIs will be "OpenAPI" compliant and will only include ontology-specific request and response bodies.

Figure 9 shows an example API exposed by the City Intelligence Layer that provides spatially interpolated air quality data. Various smart cities currently only have point-based visualisations of air quality sensor data which is not effective in identifying problem hotspots and is not intuitive to operators in responding to alarming levels of pollution.

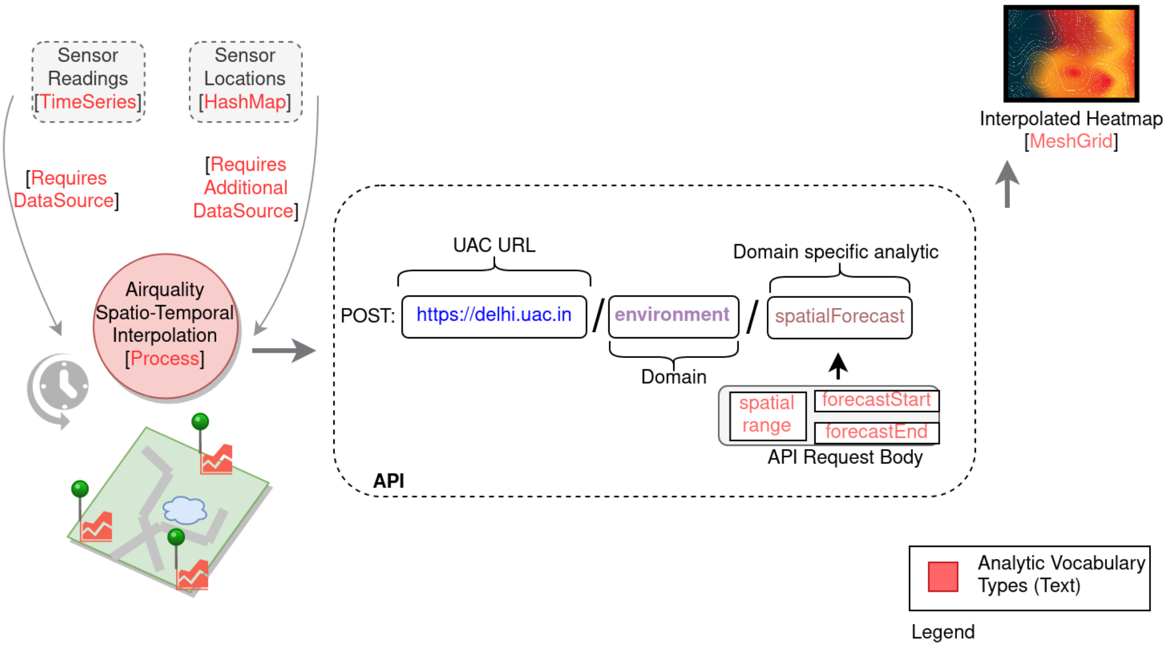

Figure 9: CIL API Example for Air Quality Domain

Figure 10 shows an example API exposed by the City Intelligence Layer that provides Estimated Time of Arrival for a bus belonging to an Intelligent Transit Management System. The Estimated Time of Arrival will be useful for application developers building multimodal City Transit Applications for use by commuting citizens to experience a more holistic transit experience.

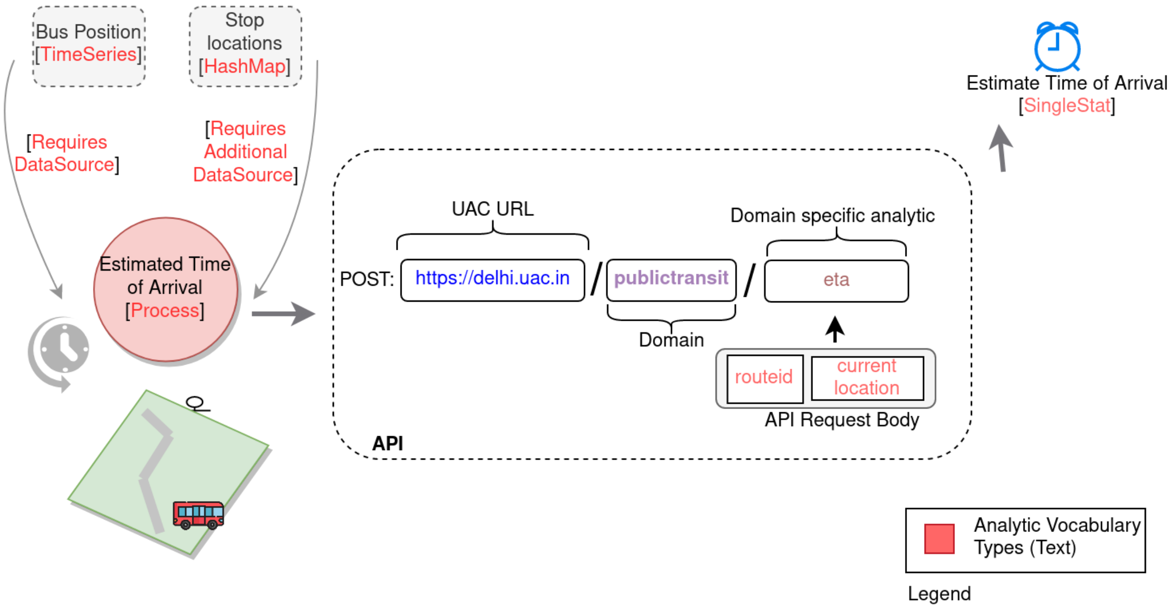

Figure 10: CIL API Example for Intelligent Transit Domain

4. Design principles

The design of the COS has been made bearing in mind the interoperability and interpretability aspects of homogenised data as promulgated by the Urban Data Exchange Specifications, which was published as a data standard in the Bureau of Indian Standards [1]. The Data Exchange acts as the heart of the COS, and provides standard data access interfaces that serve semantically homogenised data for consumption by the CIL to provide intelligence. In addition to the standard APIs of the Data Exchange, domain-specific APIs serving outputs of domain intelligence algorithms are also provided to facilitate the development of even more advanced domain specific business use cases (such as multi-model applications, cost optimizations, quality of service reporting, etc). Such domain specific APIs have shown utility in advanced use cases of Smart City Operations as described in [2] and have seen wide adoption in several exemplary Smart Cities in countries such as Singapore [3] and Finland [4].

Figure 1 shows the high level architecture of the deployment model of the COS. The COS essentially comprises (i) the Data Exchange Layer which serves as a data and API standardisation mechanism and (ii) the City Intelligence Layer (CIL) which consumes this standardised data and performs analytics operations to provide for domain specific APIs. Connecting the two layers are various types of general (standardised) as well as domain-specific analytics engines, schedulers, and OLAP access. The analytics engines include those that produce cleansed or processed data. Figure 11.

One of the highlights of the CIL is the “Pluggable Analytics”. The CIL acts as an execution and data delivery platform for “Analytics Providers” to host and execute their analytics on the data coming from the Data Exchange System as well as transformed versions of the data, thereby democratising the development and deployment of analytics. The platform will then take the responsibility of periodic execution, generation of alerts based on rules, and most importantly “Data Serving” through domain-specific APIs. Applications can then consume advanced analytics data through these APIs to build consumer-facing applications. In addition, the CIL could also support specific business intelligence use cases and even “Natural Language Querying" capabilities to facilitate easier user interaction. City officials can plug in (or use) dashboards to visualise the results in the ICCC’s and also receive monthly reports for gauging system performance.

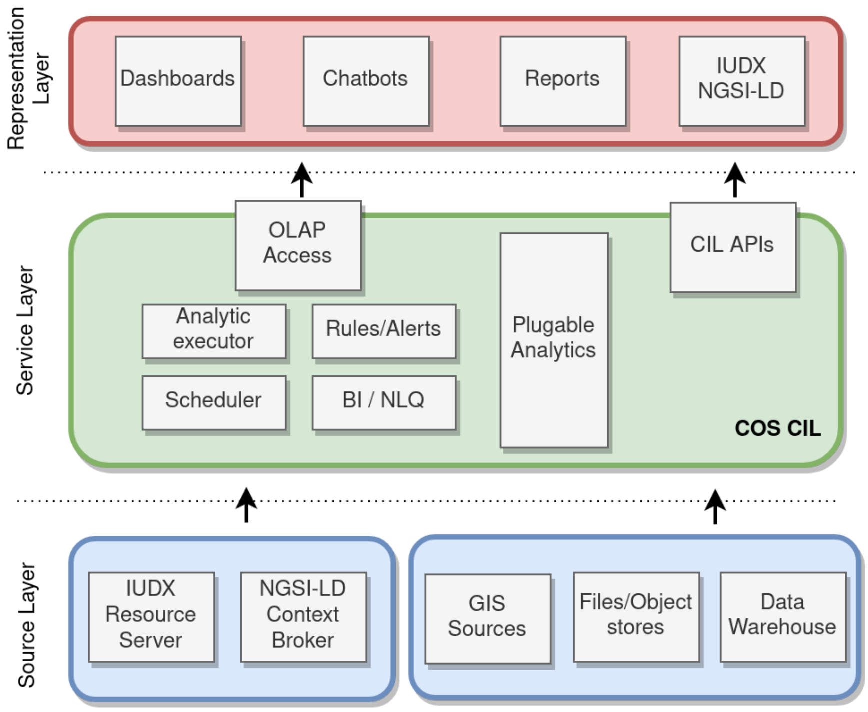

Figure 11: Architectural overview of CIL showing its components

Another highlight of the City Intelligence Layer is the standardisation of analytics terminologies for machine interpretability and application interoperability. All analytics processes are described through a commonly agreed upon analytics ontology as shown in Figure 12.

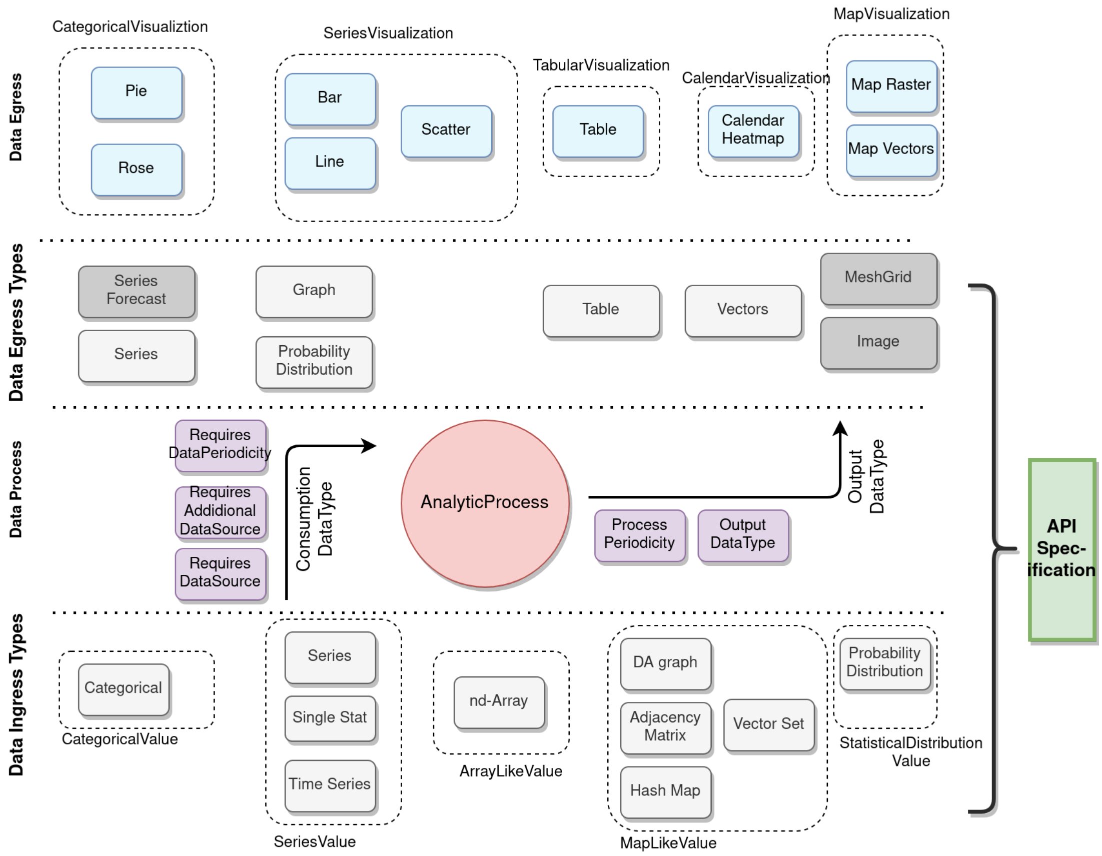

Figure 12: Intelligence Layer Ontology

All providers of analytics services need to adhere to the ontology and clearly specify the input datatypes, output datatypes, visualisation specifications and other details of analytics processes as shown in Figure 13. This will enable ease of use of complex processed data and allow for wider aggregations such as “state-level reports" and “sector-wise performance reports”.

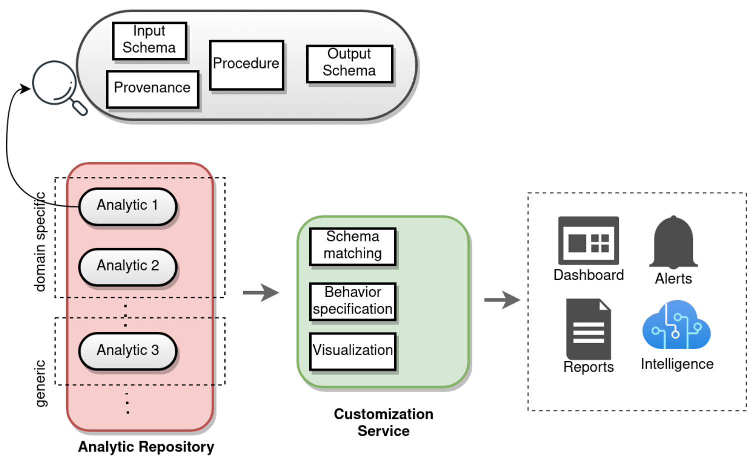

Figure 13: Analytic Process Specification

The COS and CIL will be developed using widely tested and adopted “Open Source” components and will be made available through publicly-accessible Git repositories to foster collaborative development.

5. Conclusions

In recent times, the concept of a data exchange has emerged as a result of the realisation that the lack of data availability is one of the main causes of poor governance in cities. However, a data exchange alone is not sufficient to address the many challenges related to rapid deployment of monitoring and decision-making solutions to city CEOs and development of applications and services for citizens. The City Operating System (COS), along with the City Intelligence Layer (CIL) addresses this gap. The benefit that COS/CIL provides to the Smart City is both the availability of Data Exchange capabilities along with City Intelligence infrastructure, which will result in a considerable reduction of development time for creating administrative solutions as well as for “Consumer Ready” applications. This will in turn lead to a significant reduction in the cost of state-of-the-art smart citydeployments and increase the efficiency and effectiveness of governance.

6. References

[1] Unified Data Exchange Architecture BIS Specification
https://iudx-catalogue-assets.s3.ap-south-1.amazonaws.com/documents/unified_document_architecture.pdf

[2] Raetzsch, C., Pereira, G., Vestergaard, L. S., & Brynskov, M. (2019). Weaving seams with data: Conceptualizing City APIs as elements of infrastructures. Big Data & Society, 6(1). https://doi.org/10.1177/2053951719827619

[3] Intelligence APIs from the Singapore Smart Cities
https://beta.data.gov.sg/datasets?formats=API

[4] Intelligence APIs from the Finland Smart Cities
https://www.hel.fi/hel2/tietokeskus/data/dokumentit/6Aika-esitteet/20171109_HarmonisedSmartCityAPIs_WEB.pdf

[5] Intelligent Transit Management Systems APIs from UK
https://developer.transportapi.com/