The India Urban Data Exchange

An Overview of the Rationale, Architecture and Methodology

Abstract: The Smart City program in India can benefit greatly from the creation and deployment of a new software platform called the India Urban Data Exchange (IUDX). This paper describes the rationale for IUDX, provides a high-level view of the architecture and deployment models, and describes a private/public consortium to enable its creation and maintenance.

1. Smart Cities in India

The Indian Ministry of Housing and Urban Affairs has launched an ambitious program for the development of 100 Smart Cities[1]. The total investment in these cities over a period of 5 years is expected to be around US$30 billion. The first wave of projects is well underway and implementation efforts are proceeding rapidly. There has been considerable learning from these initial efforts and it is worthwhile to examine how to tune the process to improve efficiency and maximize benefit.

The SmartCity projects have been chosen through a comprehensive process of citizen engagement and public/private collaboration. The resulting outcomes are therefore quite different for each city, with each city investing in applications, infrastructure, and services that meet hopes and aspirations of their citizens and stakeholders. Each city chooses systemintegrators (SIs), hardware and software vendors (OEMs), and application developers that best meet its needs.

An illustration of various urban services and Integrated Command and Control Centre (ICCC)in a smart city is shown below:

Figure 1

[1]See http://smartcities.gov.in/content/

While the first implementations have just begun and are starting to show positive results, they are largely independent, created in vertical silos, with no standardization of software components, or of their interfaces, or of the underlying data models. Data created by a specific application is usually available only to that application and cannot be leveraged more broadly.

This limits the ability to gain broader insights or value from the enormous amount of data being generated. Cities are unable to benefit from data gathered or lessons learned in other cities, and often across different departments in thesame city. The use cases described in a latersection demonstrate this through specific examples. This is a clear case where each city may optimize around its immediate objectives, but may fail to achieve its goals of maximizing citizen benefits with their investments.

2. The City as a Data Exchange

Cities around the world have learned that they possess a new valuable asset, namely the data generated by their various departments and agencies. Some of the data consists of streams of IOT data from installed sensors (e.g. Air Quality, Traffic, etc), some of the data is demographic or geographical, some may be from municipal tax or property records, some from legal documents or registrations, and some may be historical data from archival sources. Each set of data has its own security and privacy consideration, as well as commercial, monetary or subscription aspects which must be observed. Cities such as Copenhagen[2], Columbus[3] and Manchester[4] have taken ownership of their data assets by creatingdata exchanges, which are software platforms that allow controlled sharing of data by providing common ways of accessing and representing the data.

An importantidea behind a data exchange is that data silos are actually not a bad thing, as each silo often represents a domain-optimized service that performs that function very well. Instead of breaking silos or moving data en masse into a central repository, this approach chooses to interconnect the disparate and distributed entities through a common data exchange. This provides a way for accessing data in a unified, common format, allowing for sharing of data between different departments in a city, as well as opening up data for third party developers to create innovative new applications and citizen services. In addition, there is an opportunity for third party providers of data, or third party providers of data analytics or data annotation, to participate in what becomes a data marketplace.

Data Exchanges are starting to give cities around the world an opportunity to step up and take control of their vast data resources, enabling improved function of their own departments as well as creating new sources of revenue.This paper describes how Smart Cities can be in the forefront of this wave, through the creation of the

[2]Seehttps://cphsolutionslab.dk/en/news/learnings-from-the-city-data-exchange

[3]See https://datasmart.ash.harvard.edu/news/article/how-columbus-accomplishing-its-smart-city-vision

[4]See https://www.manchester.gov.uk/smartercity

India Urban Data Exchange (IUDX). As described in this paper, IUDX will be completely open source, based on an underlying framework of open APIs, data models, and the security, privacy and accounting mechanisms that will facilitate, easy and efficient exchange of data among disparate urban data silos. IUDX draws on ideas and, where feasible, code, from best-of-breed global projects such as CityVerve[5] and Fiware[6]. However, we are heavily skewed towards the Indian ecosystem, cultural norms, city nomenclature, payment and identification systems, etc. and therefore our solution is particularly suitable for Indian Smart Cities.

The paper outlines the objectives, structure and deployment model of the IUDX platform. It also suggests a collaborative & multi-institutional approach that will create and maintain the specifications, models and a software implementation as an open source project, with participation from all stakeholders in the smart cities ecosystem in India. This will be referred to as the *Open Smart Cities of India (*OSCI) consortium.

It is expected that *all* stakeholders in the Indian Smart City ecosystem will gain from the initiative. Within each of the Cities, citizens and the community will benefit through the availability of better, more innovative, and cheaper applications and services. The Cities themselves will benefit by the reduced development cost and faster development times enabled by a standard platform, together with the ability to choose vendors freely and avoid vendor lock-in. They will see new source of revenue through the unlocking of data assets, and will unleash innovation from entrepreneurs and community, without any cost to themselves. Industry will benefit enormously through the improved ability to find skills and rapidly ramp up projects. They will also see reduced development expense enabled by a standardized and open-source platform, and be able to focus on innovation and differentiated value rather than design basic platform software. Start-ups, in particular, will benefit from the decrease in heterogeneity IUDX provides. Third party sources of data (such as private apartment complexes) will have an new opportunity to share and monetize their assets. Academic institutions and research labs will be able to conduct more meaningful research by having direct access to a wide variety of data.

3. Illustrativeuse cases

In order to set the perspective for the IUDX and OSCI, it will be useful to examine a few illustrative use cases to motivate their need.

[5] See https://cityverve.org.uk/
[6] See https://www.fiware.org/

Case 1relates to connecting Fire, Traffic Police and Hospitals for rapid and effective emergency response. It highlights importance of data sharing to allow collaboration between departments withina city.

Use Case 1

The city's traffic police department has just completed an expensive installation of cameras, Public Address (PA) systems and smart signals at traffic junctions. Separately, the fire department has installed smoke detectors, panic buttons and provided citizen apps which can provide alerts and early warnings of dangerous fire situations. Hospitals have installed a digital inventory/patient management system. The city administrator would like to enable these entities to work with each other seamlessly to provide the best emergency response during a fire event. Ideally we would like the smoke alarms to trigger a chain of events leading to dispatch of fire trucks, ambulances, creation of green corridors and readiness of emergency rooms in a seamless fashion. This can only happen if there is a **trusted** fire emergency response application, which can work seamlessly across the three entities. This application will ***receive real time data from the fire department*** as well as the hospitals, and will be able to ***send alerts/requests for green corridors*** to the traffic police. This requires a data and application level ***connectivity between*** the emergency response application and the three Entities. Furthermore, having a ***standardized way*** to connect via published APIs and data models, will enable the ***same emergency response application to be deployed across all cities*** – regardless of the specific vendors used by each city/agency.

Case 2 deals withoptimizing streetlight power consumption. It shows how an application designed for a particular purpose can benefit by using data generated for a completely different purpose.

Use Case 2

A large city can have anywhere between 400k to one million street lights. The electricity costs for street lighting could range from 55 Cr. to 100 Cr/year. With smart streetlights throughout the entire city, these can be monitored remotely and their brightness controlled proactively.The city manager could optimize energy usage by reducing brightness during peak activity times as there will be sufficient lighting from street shops as well as vehicular traffic. In addition, she could reduce the brightness when there are no people on the street, but ensure that as soon as some person is detected, the lights are on at full brightness. She would like to use existing traffic cameras to provide this data. An application which can derive the traffic and people density ***information from traffic video feeds***, could ***provide this real-time data to the streetlight management*** system. Having published APIs and data models could allow this new application to interact with both the systems. ***In addition, with standardized interfaces, the same application could be deployed countrywide***. In the case of large cities, even a 10% reduction will save the city 5-10 Cr. annually, which could then be redeployed to other projects.

Case 3 deals with re-using a successful citizen application. It shows the average citizen benefit of rapid sharing of ideas and success stories between cities.

Use Case 3

The city has created a highly successful application for citizens to interact with its sanitation department. Citizens can use an app on their smart phones to photograph unsanitary conditions or improperly disposed waste matter and pinpoint the location through the GPS. Through some back-end analytics, the issue is then analyzed and passed on automatically to the appropriate sanitation crew that may already be in the vicinity, leading to rapid resolution of issues, sometimes within minutes of the report. A sister city would like to adopt this application (with the city's full permission). ***By having standard interfaces to solid waste department's system of the city, such countrywide portability will be possible***.

Case 4 is based on the high implementation costs experienced by an innovative start-up company doing parking management. It shows how standardization of API's and data formats can significantly help start-ups.

Use Case 4

An innovative start-up has created a citizen app that finds open parking spaces and directs the citizens to them. The app manages both street and garage parking. It interacts with a wide variety of sensors, cameras, drop-gates, and other parking infrastructure. Unfortunately, the heterogeneity of the infrastructures has made the cost of implementation prohibitive, causing great strain on the finances of a cash-strapped start-up. ***Having standard API's and data models across vendors, will make this implementation far less costly***.

Case 5 deals with splitting a contract to maximize flexibility and bargaining power.

It shows how a common framework can assist in improving the RFP process.

Use Case 5

The city is going to tender on a major Smart City project for transit, traffic and public safety. The city manager wants a single Integrated Command and Control System (ICCC) system with all these three different aspects on a single dashboard. However, she also knows that system integrator A has superior skills and expertise on transit, and System Integrator B is better for traffic and public safety. She would like to split the tender and award one piece to A and another piece to B, while being assured that everything would work together. ***A standard data exchange framework would enable this easily.***

4. Overview of India Urban Data Exchange platform (IUDX)

India Urban Data Exchange (IUDX) will be an open source software platform that will facilitate secure, authenticated and managedexchange of data amongst various data platforms, 3rd party authenticated & authorized applicationsand other data sources, data producers and consumers, both within a city to begin with and scaled up across cities eventually at a national level, in a uniform &seamless way.

The platform will provide full control to the data owners as to what data to expose and to whom. Built-in accounting mechanisms will enable connect with payment gatewayswhich will form the foundations for a data marketplace. The whole platform will be developer friendly, via definitions of open APIs (application program interfaces) and data schema templates (formats for interpreting data), so that a whole new application ecosystemgets created.

The Initial focus will be to enable data exchange between various city departments, between governments &citizens and governments & private sector within a city. Going forward, the initiative will scale up to data sharing between various cities &their stakeholders on a national level data sharing platform. It will directly address the issues that inhibit sharing & extraction of maximum value from the City's data.

The initiative would essentially create a unified single-point data market place for various smart cities ecosystem stakeholders. This will help cities with new revenue sources and create a fertile environment for innovation.

Figure 4 below illustrates the architecture of this platform:

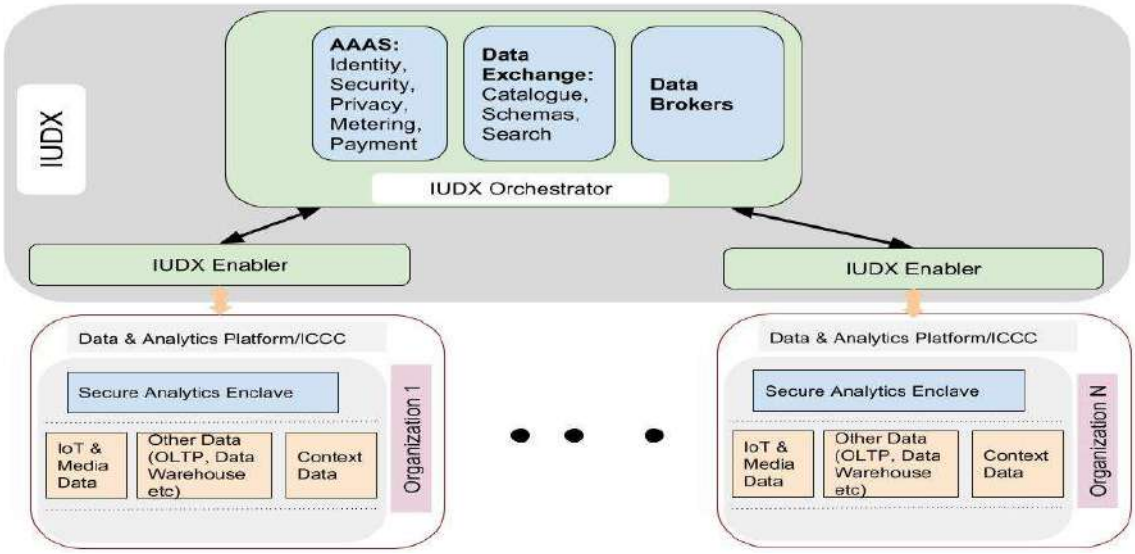

Figure 4: Proposed Architecture of IUDX platform.

The architecture consists of two primary building blocks, the IUDX Orchestrator and the IUDX Enablers.

The IUDX Orchestrator will provide a holistic view of all the features and services offered by IUDX and via IUDX Enablers, the IUDX will tie together diverse data platforms – each serving their own organization or department,

The two main enablement services working as IUDX orchestrator are:

a) Authentication, Authorization, Accounting & Security (AAAS) Enablement Services: These services would provide for identity, security, privacy, authorization, metering, payment support and other management related activities.

b) Data Exchange Enablement Services: These will host the meta data catalog (data about the data) like schemas, information models, ontologies etc., and will be searchable.

These services will provide holistic view across all the data platforms. IUDX enabled platforms can seamlessly **inter-work not only across the various city departments, but also across cities at a national level**.

IUDX Enabler function is elaborated on in Figure 3 below:

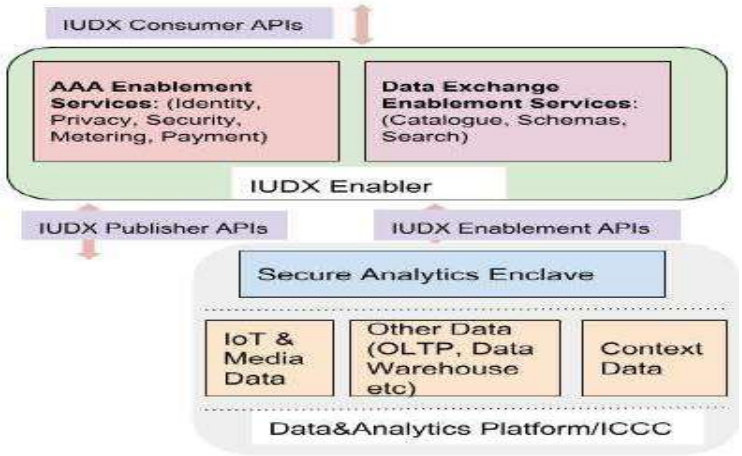

Figure 5: IUDX enabler function

As shown in the Figure 5, any data platform can be IUDX enabled, via a set of enablement services, to facilitate the access to the data within the platform. An IUDX enabled platform will have Open APIs and standardized schema mechanisms for data consumption and publications.

We envision these services to seamlessly inter-work not only across the various city departments, but also across cities at a national level.

IUDX compliant applications will be able to use consumer APIs to pull data from any of the underlying data platforms and using the publisher APIs to push data to any of the applications behind the individual platforms. Standardized APIs and data schema templates, will enable an IUDX compliant application to work in a city without needing any modification. Additionally, the standardized publisher APIs along with common data schemas, will enable vendor neutrality for IoT devices.

In summary, the IUDX platform will consist of:

a) Definitions of the APIs for the consumer, producer and enablement interfaces.
b) Definitions of the schema templates (where needed) for both IoT data as well as context data[7].
c) Open source libraries that implement the APIs and other exchange related functions to help the application and device developer community to easily work with IUDX.
d) Open source implementation of the enablement and orchestration services.

In IUDX, There shall be clear definition of **data ownership and sharing mechanism**, under the control of the data owner. The platform itself is only a vehicle for the data owners to more effectively manage and share the data as per their discretion/policies.

[7] Context data digitally represents the existing physical, organizational and other facets of a city. Examples include GIS information, utility grids, road networks, organizational structure, building and other structures etc.

Use Case 1 revisited with the use of IUDX: Emergency Response (ER)

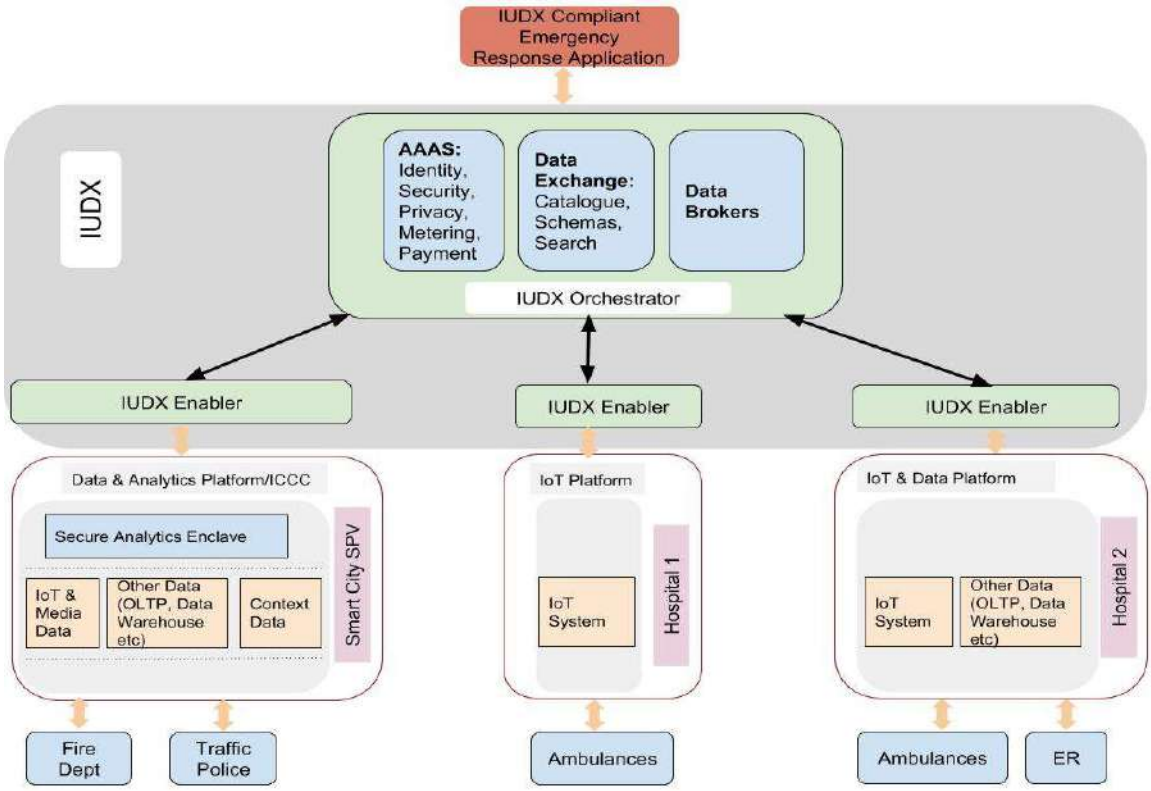

Figure 6: Illustration of how IUDX supports use case 1.

Consider the use case of a new emergency response applicationwhich coordinates across three systems: Fire department, traffic police and a bunch of hospitals. Typically, the hospital systems will be separate from a smart city deployment. However, with IUDX enablement of each platform, it will become very easy to coordinate across these different systems. When a new emergency response application is registered with IUDX, it gets permission to access the fire alarm data from the fire department, the hospital ER as well as ambulance data from the hospital system and be allowed to post requests for green corridor to the traffic police. Once the application is live, it interacts with the city data platform and the hospitals' data platforms to thus effectively coordinate the emergency response. With IUDX the same application **can now run at all IUDX enabled cities**.

Use Case 2 Revisitedwith the use of IUDX: Streetlight Management

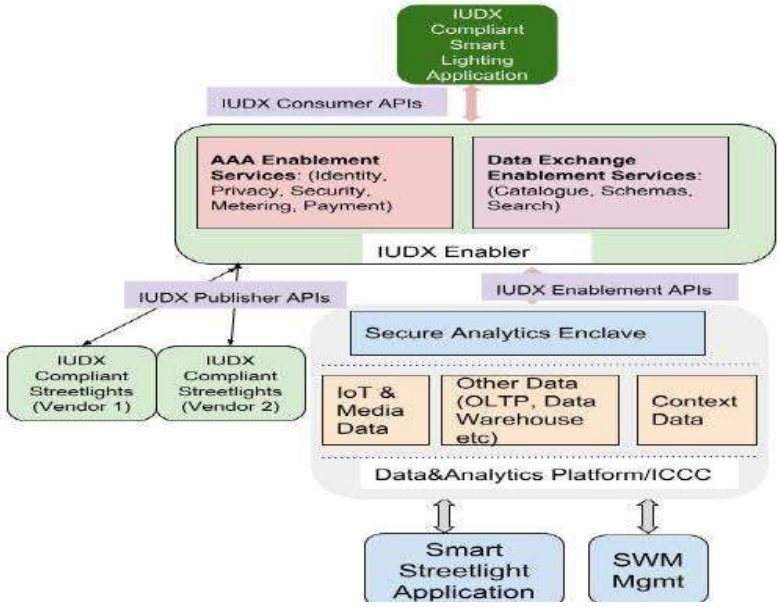

Figure 7: Enabling vendor agnostic deployment of smart streetlight solutions

Consider a smart city platform that hosts a smart streetlight application. With IUDX enablement of the platform, vendors can agree on common schemas for exchanging data between the smart streetlights and the city's platform. This will allow new streetlight devices from multiple different vendors, that are IUDX compliant, to work with the city's platform. Similarly, a new streetlight management application that is IUDX compliant, can manage all the streetlights (including the non-IUDX compliant ones), in a more efficient way, for example, by combining analytics from traffic junction cameras to detect pedestrian activities. The traffic analytics can run in the secure analytics enclave, to ensure the privacy of video data. **More importantly, all these devices and applications will be able to run seamlessly across all IUDX compliant cities**.

5. IUDXdeployment scenarios

The IUDX platform described above is ultimately a piece of software and needs to be deployed and operated as part of the overall Smart City IT infrastructure in order to be of value. Since IUDX is a data exchange and not a data repository, it provides considerable deployment flexibility and can add value to a wide variety of scenarios. Some of these scenarios and considerations are mentioned below. Each city has the freedom to make the choice that best suits its requirements and project maturity.

Importantly, IUDX is not a *rip-and-replace* deployment. It is designed to add valuein existing or planned deployments (so, called "brownfield" deployments) *withminimal modification or disruption to existing platforms.*For these brownfield deployments, the IUDX enabler components along with the necessary API adapters, will allow existing data in their data platforms to be exchanged. Some

rework (mainly at the API adapter level) may be required.For new deployments (“greenfield” deployments), adapter rework will not be required. System Integrators can include IUDX components natively in the deployment.

IUDX can be deployed in a traditional on-premise model; or in a private or public cloud model; or in some form of hybrid deployment. The cloud service providers may setup a cloud in the city itself – thus ensuring that the data lakes/warehouses/marts are local. With high bandwidth video data, local analytics (at the fog or edge level) will become important.

In a fully deployed IUDX scenario, each city will have a unique instance of an IUDX system associated with it. As mentioned, oneCity’s instance may be physically running in a City owned data center while another City’s instance may be in a public or private cloud. IUDX instances can be federated together to create the virtual appearance of a larger IUDX system, even though the instances themselves operate as separate systems. So, it is possible to create an India IUDX, or a State IUDX, which will provide a virtually aggregated view of data at an all-India level, or a state or regional level (see Figure 8). This enables analysts and developers, to create applications or dashboards that operate at different geographicalor administrative levels.

Another important aspect is the operating model. In some cases, the cities may themselves own and operate their instance of IUDX (as an on-premises or cloud deployment). In this case, they (or their agents) make key implementation choices, choose the deployment model and infrastructure, and ensure 24x7 operation of the platform. It is also possible, the city may simply subscribe to a cloud provider offering an *IUDX-as-a-service* and obtain a turn-key instance of IUDX with issues such as availability, disaster-recovery, and security completely outsourced to the cloud provider. Hybrid models are possible where Cities maintain some operational control while still using IUDX-as-a-service for some functionality.

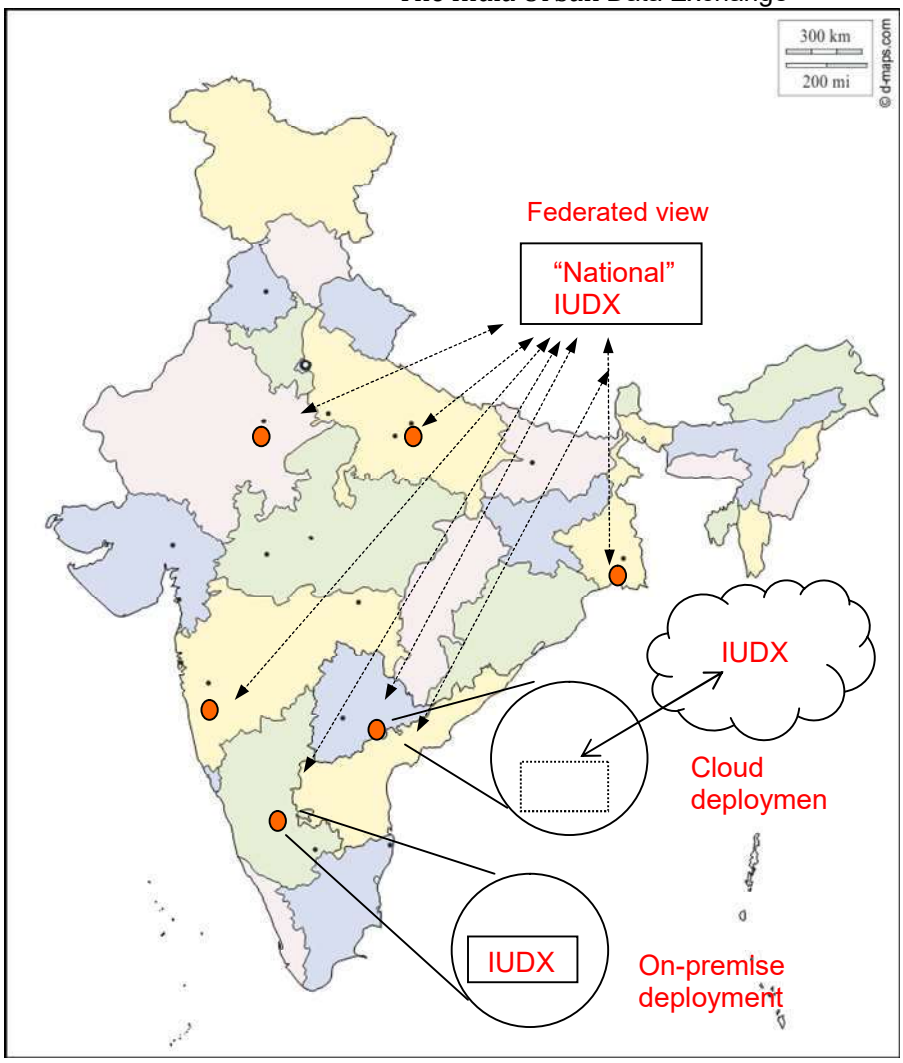

Figure 8: Examples of IUDX deployment scenarios

6. Open Smart Cities of India (OSCI): A collaboration initiative

In order to set up the IUDX platform, it is proposed to establish the Open Smart Cities of India (**OSCI)** consortium, that will develop and maintain IUDX as a collaborative project.

There are several key requirements, whichdrive the choice of OSCI organization structure:

1. There must be broad based participation in the design and evolution of a common platform such as IUDX. No single stakeholder or group of stakeholders should have an undue advantage.
2. It is vital that IUDX platform evolves and expands over time, benefiting from the latest and greatest in technology and software engineering.

3. The best minds from industry, academia and community must be engaged in creating and advancing IUDX platform.
4. The platform must be responsive to the evolving needs of Indian cities.

The OSCI structure and governance mechanism have been carefully chosen to achieve these goals. OSCI is proposed as a nonprofit company under Section 8 of the Companies Act 2013[8]. The organization, governance, and financing structure draws heavily from similar organizations of national scope[9], and from successful international open-source collaborative projects[10]. The stakeholders involved will be from Central and State government, Smart Cities, research and academic institutions, and relevant segments of the industry (software and hardware companies, system integrators, service providers, start-ups and individual entrepreneurs, etc.).

OSCIwill mostly operate as a fast-moving start-up company, albeit with a board and technical steering committee representing the Smart City ecosystem. It will hire top-notch core professional staff that will execute the day-day activities. But, because of its collaborative governance and financing, it will be possible to augment the core staff with technical experts on loan from industry, academic institutions and the government. OSCI would act as the bridge between the city and the civic-innovation community and hence, would enable bottom-up innovation within the cities.

The immediate priority of OSCI is to set up and scale the initial project, i.e. IUDX. This involves developing the API specifications, data schemas, and code for the IUDX core services. As OSCI evolves new projects will be created based on member interest and advocacy. It is likely that the future project will include reference implementations for most commonly deployed smart city applications.

OSCI would not only focus on creating state of the art middleware but would also catalyze creativity through the Smart City ecosystem through a multi-disciplinary approach that includes education programs, advocacy, popularization of best-practices and success stories, internships and training, industry and start-up outreach, hackathons, etc.

It is important to note that OSCI is not a standards body. It will work with BIS, TSDSI and other standards bodies with India and globally, which adopt and influence standards. In fact, as the interface standards for IUDX relate to programmatic APIs and data formats, It is largely accepted that paper specifications are inadequate in these areas and a reference software implementation is the only way to provide an adequate specification.

[8] See http://www.mca.gov.in/MinistryV2/companiesact2013.html

[9] See, for example, National Payments Corporation of India https://www.npci.org.in/ or Egovernments Foundation https://www.egovernments.org/

[10] See, for example, Linux Foundation https://www.linuxfoundation.org or OpenDaylight Foundation https://www.opendaylight.org/

For the initial years, OSCI will be funded by member contributions, grants from governmental agencies, sponsorships and support from philanthropic organizations. As OSCI matures, it will evaluate ways of making the organization self-sufficient and not in need of external funds.

Figure 9shows the operation of OSCI in a schematic representation. An important observation is that OSCI does not interfere in the relationship or the tender process between the Smart City Special Purpose Vehicles (SPVs) and their respective vendors.

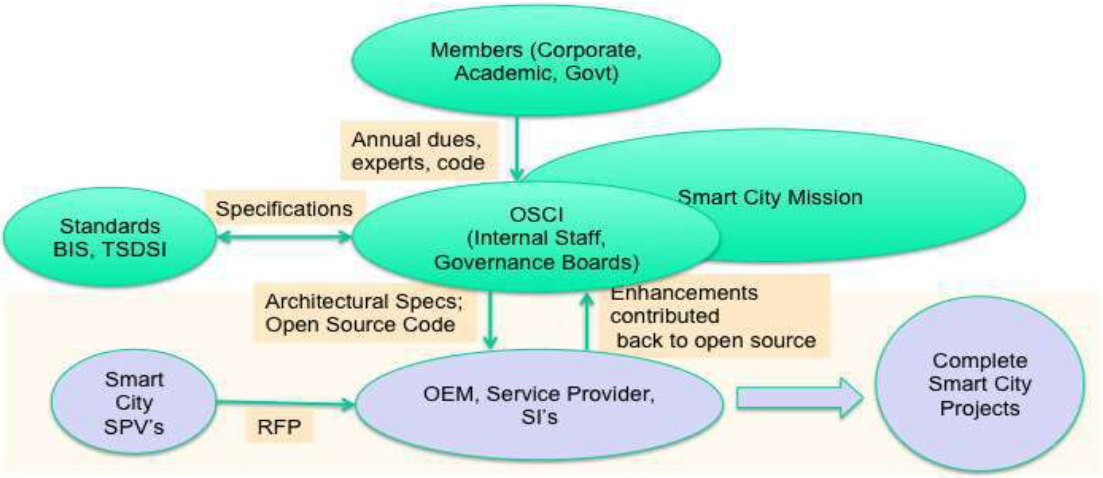

Figure 9: Proposed Operational framework of OSCI

7. Involvement of the Cities, Citizens and Industry

The board structure of OSCI ensures that the cities have a strong voice in the governance of OSCI. OSCI will define a process for requirements setting and prioritization for each release, and the cities will actively participate in it. In addition, to ensure grass roots engagement from citizens, OSCI will conduct outreach and focus groups to validate that the installation is actually addressing citizen needs.There shall be several checks in place to prevent a single vendor (or group of vendors) from unduly influencing the priorities or technical direction. Any board member will have only one vote on the board of OSCI. Most importantly, all artifactsshall be in open source (most likely under the Apache2[11] license) and the governance of OSCI will entirely be transparent.

When the IUDX code is mature enough to be used, the Smart Cities Mission will provide guidance to various Smart Cities. This will be in the form of an advisory, recommending that the implementations include IUDX. However, Smart Cities SPV's shall be free to make decisions in their best interest considering their city's requirements. It is expected that the Cities will create RFPs and manage the RFP process and OSCI will not interfere in that process.

[11]Apache2 is one of the most commonly used open source license. It is a very "permissive" license in that it provides considerable flexibility in the way the code can be used. See https://www.apache.org/licenses/LICENSE-2.0

Broad industry involvement is essential for the project to succeed. Industry partners will provide key skills and guidance for the project, and also be a source of financial support. It is also expected that industry members will contribute code to OSCI that they may possibly have developed initially for other purposes. This occurs extensively in many collaborative open source projects as companies find value in the reputational benefit and community accolades that result from having their code adopted in such a broad based platform. Nonetheless, the selection of a particular code to adopt, will be based on merits and not on any commercial or political interests.

While any organization can benefit from OSCI without becoming a member, it is believed that members benefit considerably by having an opportunity to steer andguide the organization. Moreover,supporting an open source organization that a company is benefitting from, should be considered a corporate social responsibility. Membership fee could be considered as part of corporate social responsibility donations.Members will gain a better understanding of trends and emerging requirements through the interaction and collaboration with other OSCI members. They willalso have the opportunities to participate in press opportunities and will be invited to meetings with Smart Cities, and be mentioned in press releases and other publicity from OSCI. Members will have early views of upcoming IUDX releases and possibly to some value added code released under an OSCI-specific open source license.

8. Next Steps

This paper describes *one* aspect of what is required for cities to leverage their data effectively and unlock the potential outlined[12]. In addition to creating their IUDX instances, there are many other aspects that relate to policy, governance, business model, etc that each city will need to address to derive useful benefit. For example, it is important for cities to create a set of data policies that govern the availability, usage, and monetization of their data. Some of these policies will come through a process of trial-and-error as these platforms are deployed in the Indian context. The OSCI consortium will have the responsibility to go beyond the technical aspects of IUDX and provide strong guidance and models for Smart Cities that address these critical non-technical issues.

[12]This paper is a work in progress. As the project evolves, the paper will be updated to reflect the current thinking and status of the effort. Comments, questions, expressions of interest in participation, etc are actively solicited and can be sent to iudx@rbccps.org