NTT Smart Solutions

Tech Brief

NTT

Table of contents

"Open **data platforms** are the last crucial ingredient of the technology base. **Smart technologies run on data**… yet this data only becomes useful once it is made available to actors who can build smart applications out of it. **Open data platforms do exactly that.**"

McKinsey Global Institute.
Smart Cities: Digital Solutions for a more livable future.

1. Introduction

This tech brief is an overview of the NTT Smart Platform, the technological engine for the NTT Smart Solutions. It provides an overview of how data is ingested, transformed and analyzed. It also describes the main segments, units and components that compose the NTT Smart Platform.

The INTRODUCTION describes how the platform is the foundation for all of the NTT Smart Solutions.

The FUNCTIONAL ARCHITECTURE section focuses on the performance of the platform, describing its main divisions and how they interact with each other. It is divided into two parts: a high-level approach where the three main functional blocks are introduced, and a deeper look to its main components.

The TECHNOLOGICAL ARCHITECTURE section presents the key principles used to define the design and implementation of the platform, and an overview of the technological structure.

1.1 Smart landscape: the future

The modern organization lives in a data-driven environment. With the advent of the Internet of Things (IoT), increasingly more sensors and data devices are connected to the internet and each other, generating an unprecedented amount of data. Knowing how to handle this data and how to take advantage of it is the key to creating actionable information and predictive analytics and organizational success.

To remain competitive and deliver high-value services in this environment, organizations are challenged with navigating through the noise and volume created by this massive explosion of data. These challenges start at the Edge, where each sensor or device creates data from a source and format-specific perspective. The absence of common data standards results in a disconnect in the structure, meaning and topology across many disparate data sources.This in turn creates a tremendous amount of effort for companies who do not always have the skills and competencies to develop the common data structures required for analysis. Furthermore, because the Edge is in a constant state of change, keeping these data sources and definitions in-synch, requires constant monitoring and support. Finally, consuming, transporting and storing all the data produced by the Edge can be costly and in most cases is not necessary. Finding the right balance of data analytics and storage at the Edge versus data transport, storage and analytics in a central location (cloud) is critical to harvest the most value from any data-driven solution.

Once data is synchronized and connected through a common data architecture, the real value can be realized and the power of data-driven AI, ML and Data Science can be applied to produce valuable results by enabling better, more insightful decision-making for the organization. These insights must be delivered in a form that is easily consumable and appropriate for the industry, role, or activity. For example, city transportation officials might require the data to be presented on a virtual map of the city with KPI's shown as alerts and notifications at intersections where attention is required. A plant manager might want to see a digital-twin representation of the machinery on the shop floor with indicators showing operational efficiency, preventative maintenance, and other important productivity measures.

To survive the increasing flood of data being generated by IoT sensors and devices, companies must make every effort to understand the data lifecycle, its structure, and relative value to the organization. NTT Smart Solutions can support the ingestion, normalization, analysis, and overall delivery of data-driven insights and help to accelerate the time-to-value for better business decision-making.

1.2 NTT Smart Solutions

NTT Smart Solutions are the best way to transform data into actionable intelligence. The system quickly calibrates and curate's data to reveal trends, discover powerful insights and generate predictive analytics that can drive better business decisions and planning. The platform helps to improve the quality of decisions and makes an efficient resource allocation that saves money.

NTT Smart Solutions are designed with flexibility in mind and can deliver intelligent insights across multiple industries. The innovative NTT Smart Platform technology drives the solutions and enables organizations to start with a small and low-risk pilot program or a phased deployment and then scale fast and efficiently. Moreover, NTT's position as a global technology leader means unparalleled support for its partners and clients across the full lifecycle of activities from planning and strategic design, to implementation and deployment. It is also especially important to emphasize that all the NTT Smart Solutions are always committed to complying with each customers' data policies.

What defines this innovative approach from a technological point of view is:

• **The platform adapts to every specific need.** The nature of this solution set is 100% configurable, allowing the building of different feature configurations according to the customer and their needs.

• **Leverage Existing Technology Investments.** Because the IoT market is constantly evolving, it is critical that Smart Solutions not only adapt to new technology, but also embrace current and legacy environments as well. The NTT Smart Platform is built to handle this broad spectrum and does the heavy lifting to synchronize technology assets to minimize disruption and increase value.

• **Faster analysis and better value from data.** The platform implements an innovative data architecture that assists the data scientist in reducing the time-consuming activity of data preparation so that they can focus on performing higher value data analysis and prediction. Furthermore, the system can ingest data from virtually any sensor and a wide variety of other data sources (cameras, audio, social media and operational systems…) creating a rich data analysis environment.

2. Functional Architecture

The platform is at the core of the solutions, it is the centerpiece that keeps everything connected and on track. It is like a set of springs and gears that processes the data from the edge through to when it is translated to numbers, graphs and statistics at the other end.

2.1 NTT Smart Platform

It is a platform focused on the analysis of the data, that works with data but never retains it, always being compliant with customer data privacy and retention policies. A platform built by modules that communicate between themselves through standard APIs, ready to be deployed whether on-premise or in the cloud and able to integrate with third-party software along the way. It is quick to deploy and ready for scale. The platform is built for a purpose and can be synthesized in the following diagram:

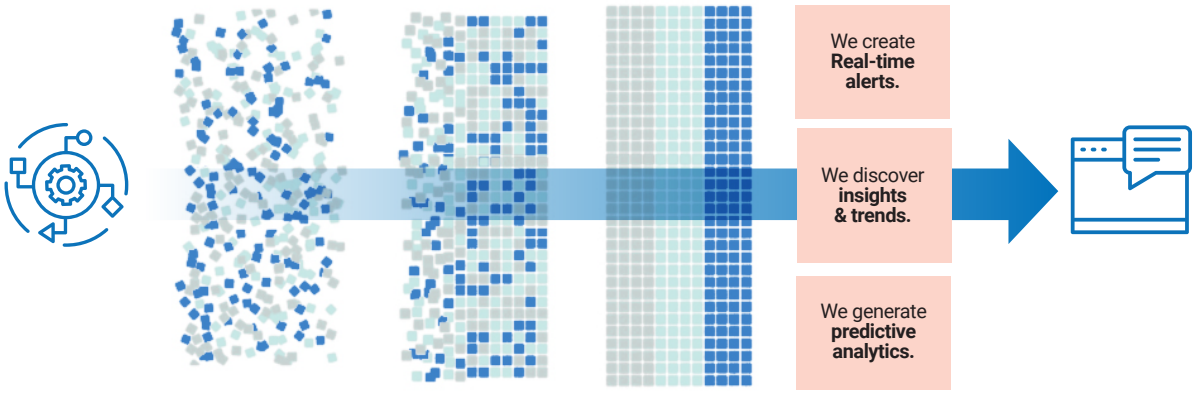

Figure 1. Main purpose of the platform.

The data is collected and aggregated from multiple sources, in different formats and topologies, and all of it must be filtered, organized and prioritized to be used by the Analytics Modules, to derive insights and predictions and display conclusions for the Business Specific Solution.

To achieve this purpose, the platform operates in a distributed environment that is composed by different parts that can be summarized in three big segments:

1.The first one is at the edge, where the data is captured from to the different sources to the Micro Data Centers (MDC).

2.The second one is the Customer Core, where all the normalizations and analytic algorithms are applied to the data.

3.The third one is the Global Core, that orchestrates deployments and keeps the system updated.

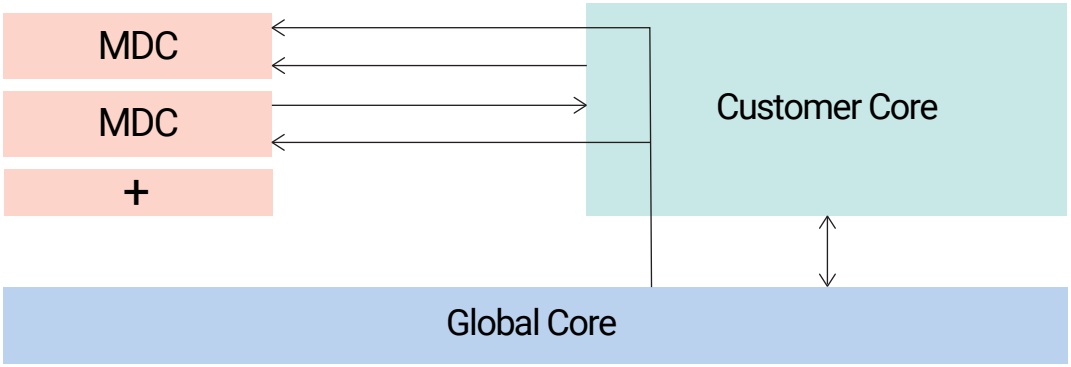

Figure 2. Main segments.

The fuel of the platform is the data. It is the main resource not only to make it work, but to understand the context, discover trends and generate insights. The MDC's at the edge ingest data from any kind of source, physical or virtual, as it can be video from a camera, sound from a microphone, ticket information from a transportation system, telemetry from machinery or an automobile, air quality from an environmental sensor, social media data, weather data, etc.

The MDCs are always located close to the source, so they can extract the valuable data right there, analyze it and send to the Customer Core only the data required for further insights. For example, if a system analyzes the crowd that passes through a street throughout the day, it does not need people's facial, nor its distribution by gender or age, much less the raw video; the only information that will be saved is the number of people on that street throughout the day. This way, the sources are the receptors of the raw data and the MDCs are the senses of the system, which evaluate and sort the information before it gets to the core for deeper analysis and decision-making.

The Customer Core is the central nervous system and includes a Data Storage module for the client to guard it and the platform to work with it. In the Customer Core all the algorithms and predictive analytics are applied to the data to get results and transform it into information that is useful for decision making. It is also the heart of the platform, for all the edge data streams end up here, and from here they are displayed as insights in dashboards, KPI's, metrics or fed to downstream custom applications.

Whether the Customer Core is deployed on-premise or in a public/private cloud, the procedure is always the same: the platform implements its analysis in this part of the system. The Customer Core derives intelligent insights and trends from the complexity of IoT systems and the massive amounts of data they produce.

Finally, the third block, the Global Core, represents the brain, the command center that supports and maintains the platform performance from the edge to the final screen. It is located in the NTT facilities, and it does not touch any customer data. It keeps the systems oiled and spinning, always synchronized and updated.

2.2 Components and data flow

Taking a closer look at the Figure 2, the whole platform looks like this:

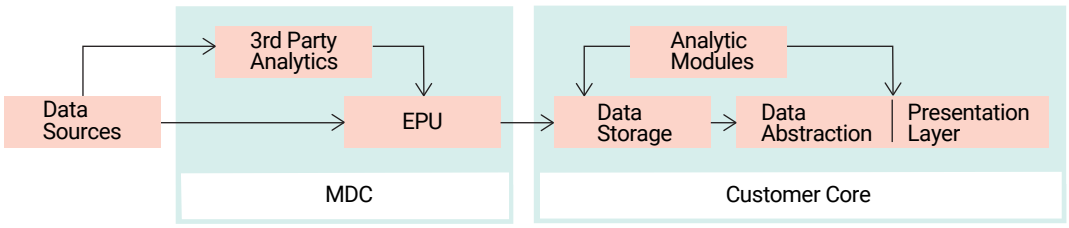

Figure 3. Extended flow chart of the first two segments.

It all starts in the Edge, where the MDC's are located to analyze the incoming data. Using standard API's, the raw data is ingested from any kind of source, physical or virtual. The big amount of data that is ingested needs to be analyzed and filtered at the edge to send only relevant and normalized data to the Customer Core in subsequent steps.This process dramatically reduces the payload of data flowing from the MDC's to the Customer Core, saving both time and money.

In every MDC there are one or more Edge Processing Units (EPU), which process data from a specific source and provide the common data structure, naming and semantics for it so that it can be analyzed with a common vernacular in the Customer Core. Also, the EPU performs edge analysis on the source data to support certain real-time alerting, trending or other active business rules. In other words, in the EPU the data is categorized, analyzed and prioritized from a "source dependent" definition to one that is "source independent" (see **Figure 4**) in a process called Common Data Object (CDO) Normalization. CDO is a proprietary data format, and it unifies the data structure for all components enhancing serviceability and re-usability of key components like the Analytics Modules.

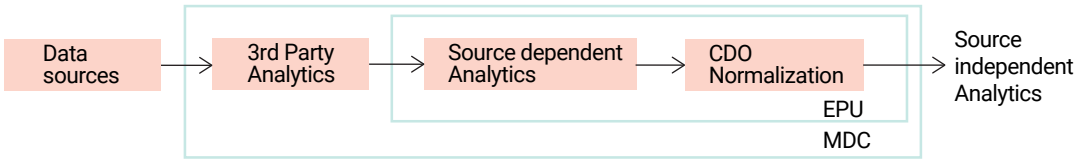

Figure 4. Data processing in the edge.

Depending on the business case, the data will typically be passed to the EPU in one of three potential methods:

• From 3rd party processors such as a Video Management System (VMS) or a Manufacturing Execution System (MES) which collects certain IoT data for non-analytic operational purposes.

• From non-intelligent sensors that provide raw data describing their operational activity.

• From intelligent sensors and other sources with some embedded analytic capabilities.

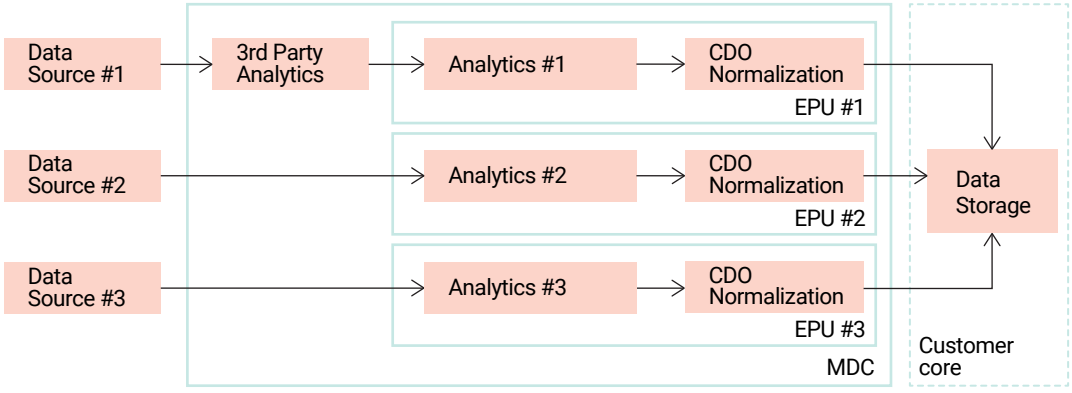

Figure 5. Data from different sources is filtered and normalized.

The MDC's only send to the Customer Core the necessary information for further analysis, hence most of the information will remain at the Edge, and only for the time-cycle required by the client's Privacy and Data Retention Policies.
In consequence, most of the data gathered at the edge will not even get to be normalized, and no personal information will be shared from the MDC's. Once the process of assimilation and abstraction of the data has been completed and all of it is in a standardize common form, it is sent to Customer Core to have analytics applied to it.

As data comes into the Customer Core it is loaded into a structured data store. Traditionally, data scientists would have spent most of their time at this point, *fishing* for the right data at data lakes that can reveal insights and potentially provide good predictions. However, this time is dramatically reduced thanks to the prior CDO normalization and abstraction process. The Analytic Modules implement specific analytic functions using Artificial Intelligence (AI) and Machine Learning (ML), and they operate on the normalized data instead of an unstructured data lake. In other words, the extra work that data scientists normally had to do is removed, and it allows them to focus on the high value of analyzing.

And, in the end, the data will be adapted and rendered in the Presentation Layer through role-based specific UIs, application and dashboards across delivery channels (Mobile, Web, Text, etc.), focused on the outcome needed from the user.

Underneath the two segments previously described, the Global Core implements the orchestration components for deployment and maintenance. It is responsible for managing the control of the entire environment, ensuring that all parts are synchronized and updated. Global Core is unique across the MDCs and the Customer Core (or Cores, if a client's environment requires more than one), and it provides the ability to push upgrades and new functionalities, as well as the capability to monitor how the whole system is working.

Completing **Figure 3**, the Global Core supports the system as follows:

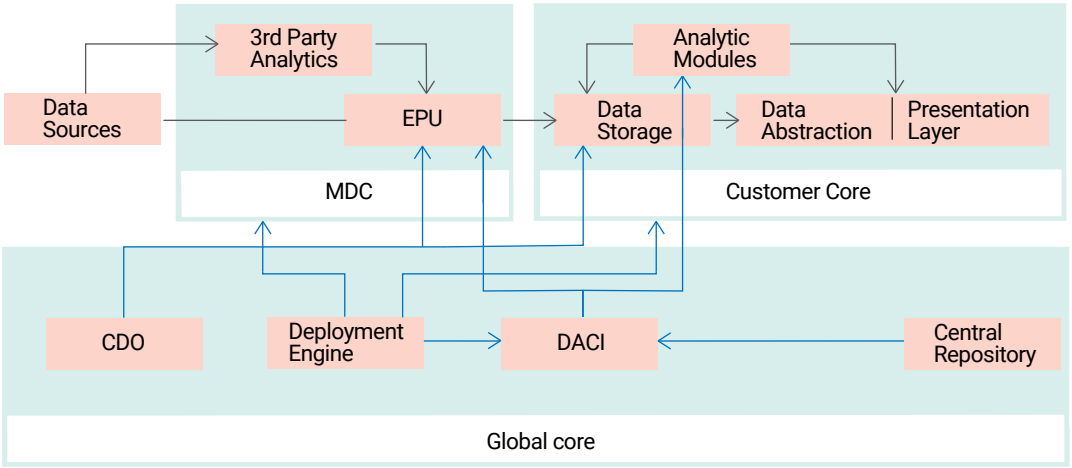

Figure 6. Extended flow chart.

As the Global Core is agnostic to the data stored in the Customer Core, it only keeps configuration information, allocated in the Repository. Repository and Data Architecture Continuous Integrator (DACI) works together in order to understand the configuration of each customer and apply the necessary changes or updates.

Specifically, the DACI provides a mechanism for Continuous Integration (CI) and Continuous Deployment (CD), as a standard CI/CD framework but specifically oriented to data architecture, enabling not only a rapid deployment, but also rapid upgrades. This way, when some changes are required from the configuration of a particular client, DACI makes sure that those changes are implemented and are in sync with the rest of the deployment.

In conclusion, DACI is the core management of the platform: it keeps everything in sync and allows any changes to be propagated throughout the system, enabling scalability and operational efficiency.

3. Technological Architecture

The previous section described how the platform works and how the data is processed throughout the different segments, units and components. This section aims to show how the components are built and how they all configure the layers of the platform.

3.1 Key principles

The NTT Smart Platform is conceived as the engine that drives NTT Smart Solutions, and as any engine, it was designed under a set of principles. These are the key principles on which it is built:

1. Open Architecture: The digital technology market is always changing and growing, looking for new ways to process data to create value. This is why NTT Smart Solutions are not tied to a particular technology platform. With an open architecture, it can be adapted quickly, allowing for access to a variety of deployment choices, from on-prem to public or private cloud.

2. Interoperable: The system can connect with different platforms or third-party software's through standard API's. In fact, it can ingest data from virtually any source. It allows the platform not only to be addressing rapid technology change generation, but also to work with a strong operational ecosystem. It is flexible enough to work in any environment.

3. Modular Design: Each of the components of the Smart Architecture Platform are separated from the others but work together through a common interface. Every installation is fully adaptable, and it can scale broadly.

As a result of the intersection of these three fundamentals, the platform achieves two significant outcomes:

1.Rapid Deployment: It is a containerized/micro-services deployment architecture that isolates the software from its environment to ensure uniformity across operational environments for rapid deployment and CI/CD changes.

2.Scalability: It is a distributed architecture that accommodates significant scale through the complete lifecycle of operations—from small pilots to large deployments.

3.2 Layers and components

The NTT Smart Platform is built on a modern container-based model that is universally deployable and secure. The diagram below is an overview of the architecture of its layers and main components:

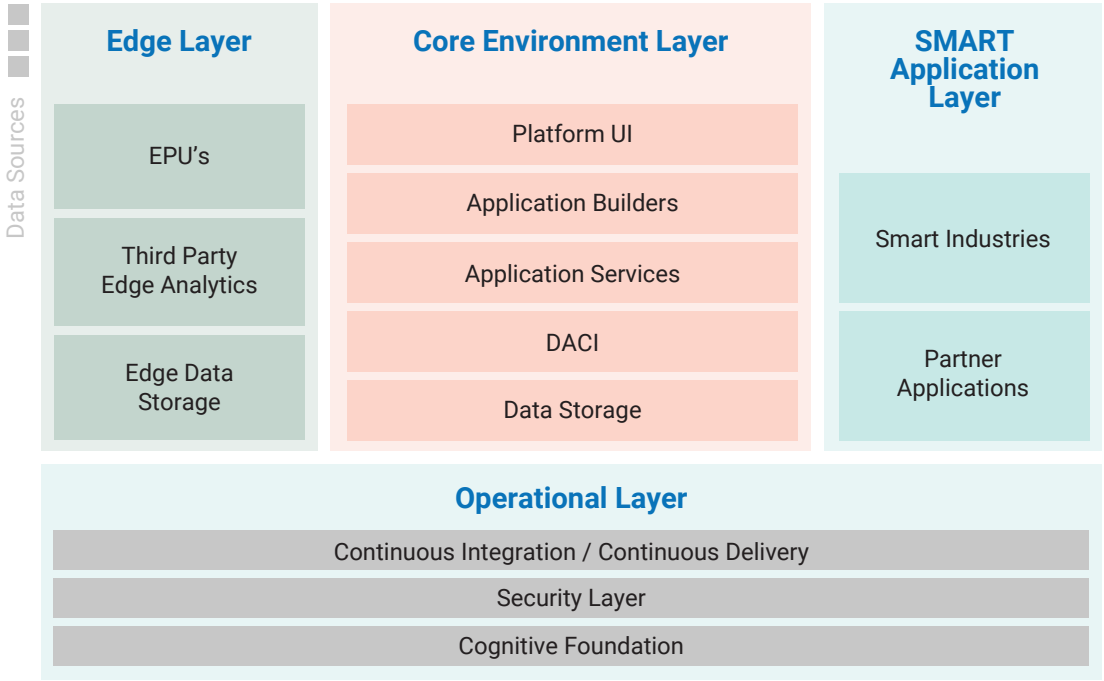

Figure 7. Architecture overview by layers.

3.2.1 Edge Layer

The Edge Layer of the platform contains both the sensors that provide information about the world, and the processing capabilities that are more efficiently implemented near this data.

One of the key capabilities of the layer is to pre-process source data to leave the bulk of it close to the edge (for example, video streams are never sent to the cloud). Also, this is the point in the pipeline that removes the coupling with the specific event data mapping to the common independent data architecture and so is key for functional scalability.

The Edge Layer is composed of:

• **EPU's:** The EPU's ingest source-specific formatted data from IoT sensors, a 3rd party source, or other data sources, analyze and transform it into a source-independent proprietary CDO format, through a process of normalization. Each EPU is sensor/function-specific, allowing a single repeatable use case to be bundled from multiple sources, and combined into a single package to ease distribution. EPUs enable the SMART Applications to scale horizontally to address small pilots or large multi-source, multi-use case scenarios.

• **Edge Data Storage:** Data Storage provides temporary event storage capabilities for the EPU's. To make an efficient and scalable solution, this storage is emptied periodically keeping only the key information needed for further analysis along with special rules based on privacy and retention policies from the customer.

• **Third Party Analytics:** Third-party solutions can be integrated to implement specialized analytics on sensor data or any kind of general information source at the edge, for initial pre-processing that might be required in some use cases. Business events generated by these tools would be fed into the processing pipeline like any other data element.

3.2.2 Core Environment

The **Core Environment** provides both the core of event processing and analytics based on the common information architecture, and all the management capabilities for the users of the platform.The main objective of this layer is to be both efficient in the sense of being able to cope with complex and ever-evolving algorithms on many kinds of events and modular in the sense of being configurable for the needs of every implementation or customer.

Components of the Core Environment Layer are:

• **Platform UI:**It includes applications for administrators, operation teams and application developers. Platform UI provides end-user applications for platform administrators, platform operators and application developers, as well as a set of pre-configured working applications that are easily extended and customized via the application Builders.

• **Application Builders:** These provide an easy low code approach to extend the functionality of the UI. These Builders include the ability to develop custom alert & notifications, reports & dashboards through a standard RESTful interface. Application Builders sit on top of the Metadata Library which simplifies the extension without requiring extensive custom coding.

• **Application Services:** These implement the business logic that runs on the platform. It is composed of a set of services that include a rules engine, an analytics engine, a Data Science Workbench, and IoT Device Registry and a library of shared and extendable algorithms.

• **Data Architecture Continuous Integrator (DACI):** It implements the core configuration management for the data pipelines in the platform, from ingestion to processing, and end-user analytics. DACI implements a Data Abstraction Layer to agnostically interact with the underlying repositories, using the CDO as the common information model, and providing SQL, REST and XML as APIs. It allows data processing to be done both in batch, on-demand and implementing streaming models, on top of the Data Storage.

• The underlying **Data Storage**: It is designed to be heterogeneous, including both relational RDBMS, NoSQL DBMS and Big Data repositories, or data lakes, such as Hadoop. The previously explained CDO normalizing enables the efficient and fast data processing, providing an alternative approach to the current trend of creating large unstructured data lakes which are flexible but lack in efficiency and performance.

3.2.3 Smart Application Layer

The *vertical* business applications are implemented leveraging the core Cloud Layer APIs, to provide value-added services. NTT has implemented solutions for the different industries such as Smart Mobility, Smart Cities or Smart Manufacturing, that can be further customized for specific implementations.

The long-term model for the NTT Smart Solutions is to provide a purpose-built application development services layer where any NTT customer or partner can build their applications, and that could even be published in marketplaces or application stores.

3.2.4 Operational Layer

This set of capabilities ensures that cross-solutions concerns are implemented the right way at all the levels of the platform, and that the usage of the infrastructure is as efficient as possible. The units that compose this layer are:

• **Continuous Integration / Continuous Deployment (CI/CD):** All the development pipeline from the source code versioning to the build, release and deployment cycle is automated at all levels of the platform. This is critical to ensure that everything is kept in sync avoiding problems due to manual intervention, and change is managed adequately.

• **Security:** Proliferation of IOT devices and other kinds of sensors introduces exponential security vulnerabilities. The platform provides industry-standard security protocols and optional services for network threat detection. Also, Authentication and Access Control are managed uniformly along with all the platform layers.

• **Cognitive Foundation:** TThe runtime infrastructure for the NTT Smart platform has been designed to be platform agnostic so there is no coupling or lock-in to a specific vendor (such as Microsoft, Amazon, or Google). Container technology (docker) and container orchestration (Kubernetes and NTT Multi-Orchestrator) are used widely to achieve this uncoupling of the runtime components to the concrete runtime technology.

Empowering data to make a better world

NTT believes in resolving social issues through our business operations by applying technology for good. We help clients accelerate growth and innovate for current and new business models.

Our services include digital business consulting, technology and managed services for cybersecurity, applications, workplace, cloud, data center and networks – all supported by our deep industry expertise and innovation.

As a top 5 global technology and business solutions provider, our diverse teams operate in 80+ countries and regions and deliver services to over 190 of them. We serve over 80% of Fortune Global 100 companies and thousands of other clients and communities around the world.

http://www.global.ntt/NTT-accelerates-smart/