PATHFINDER REPORT

451 Research® | Advisory

The Role of Operational Analytics and Interoperability in the Era of IoT

NOVEMBER 2016

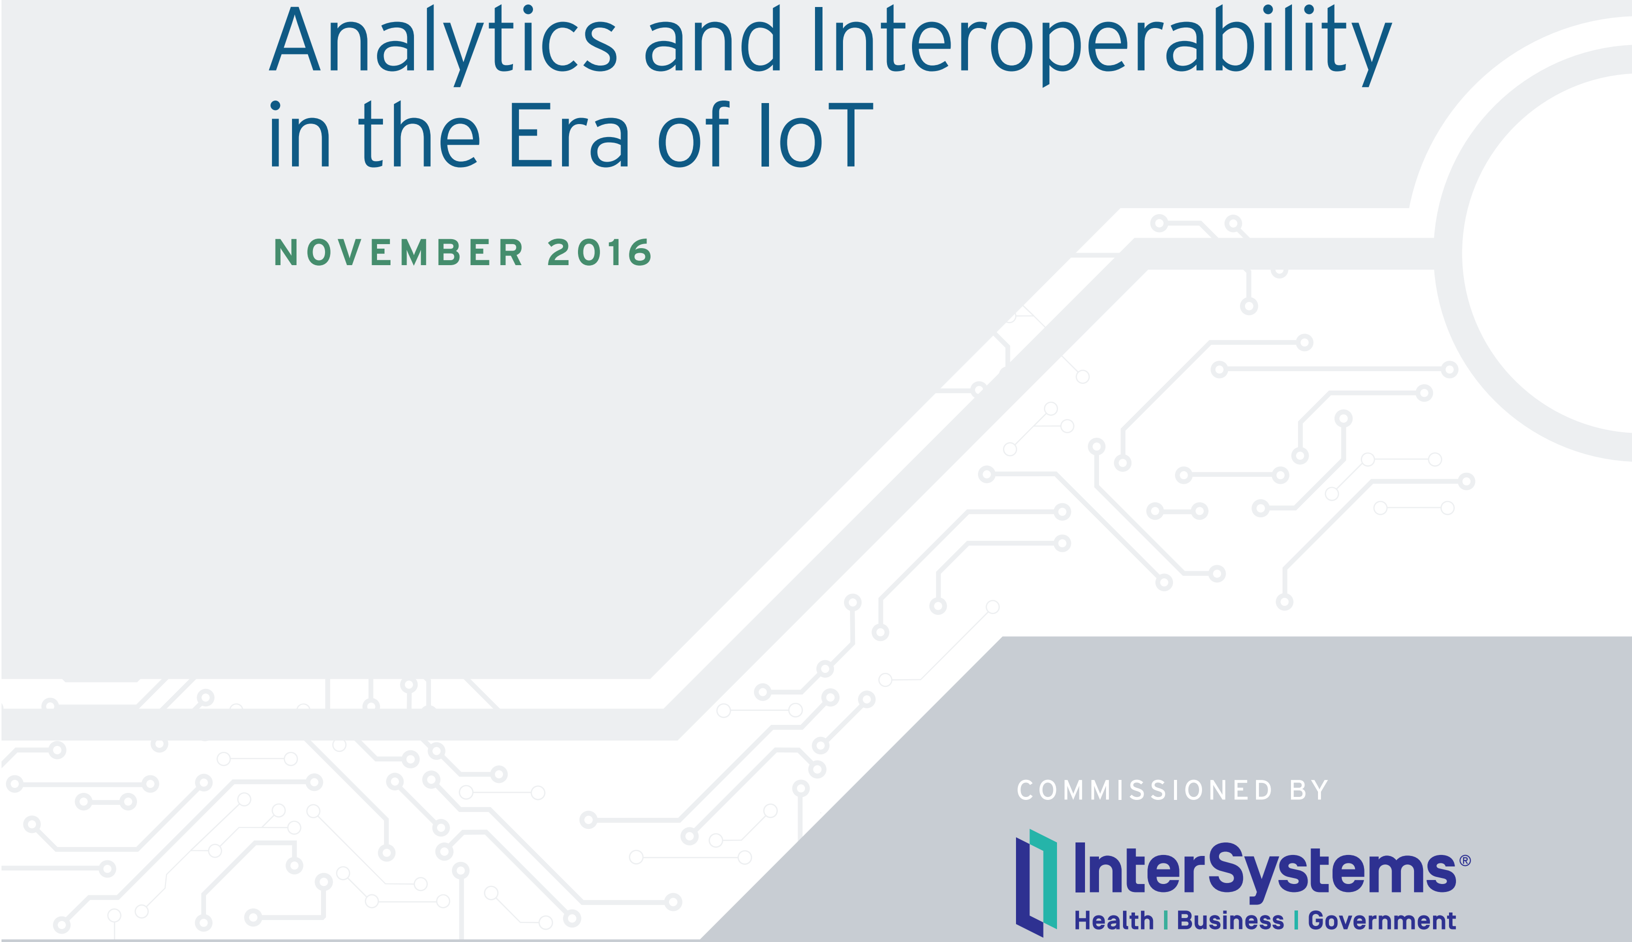

COMMISSIONED BY

InterSystems®
Health | Business | Government

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About 451 Research

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
20 West 37th Street
New York, NY 10018
+1 212 505 3030

SAN FRANCISCO
140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON
Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 207 426 1050

BOSTON
One Liberty Square
Boston, MA 02109
+1 617 598 7200

Abstract

The diversity and scale of the Internet of Things (IoT) will require many organizations to re-architect their data estates to enable support for multiple devices, data formats and new approaches to analytics.

There are distinct challenges to accomplishing this – in most instances this requires data to be ingested, curated and analyzed from millions of endpoints, with multiple data formats and multiple data models, at a high rate of speed. Meanwhile, there has been a strong drive to handle both operational and analytical needs in the same platform to reduce complexity – something we call 'combined operational and analytic processing' – in order to support operational analytics workloads driven by 'multi-modal' databases.

There are also new types of data platforms in the modern enterprise – SQL, NoSQL, Hadoop and so on – and a growing trend toward storing and processing data in multiple data formats in a single data platform, which we call 'multi-model'.

These multi-model databases enable organizations to store and process data in multiple formats, enabling interoperability and providing the flexibility to evolve to keep up with the rapid pace of innovation involved in IoT, in terms of both changing data formats and the ability to apply the result of analytics processing to operational applications.

Current relational database technologies are challenged with delivering these benefits due to the inflexibility of fixed schema. However, SQL continues to be the default approach to analyzing data, so organizations deploying IoT need to rethink how they can optimize analytics for these deployments and consider multi-model, extremely scalable data platforms that combine support for SQL-based analysis with the flexibility to process multiple data formats and the agility to evolve to meet changing requirements.

Additionally, the scale of IoT projects will push computational and analytical capabilities closer to the network edge, where IoT devices live, in order to reduce the potential for the huge amount of data streaming from 'things' to saturate datacenter networks, storage and processing capabilities. Edge processing does not exist in isolation, and needs to be constantly updated based on analysis of the combination of real-time and historical data.

Introduction

The term Internet of Things (IoT) describes an emerging technology paradigm in which sensors, machines and other smart devices create and then disseminate data. All sorts of 'things' that previously had little or no intelligence built into them – including energy meters in the home, equipment in a factory or traffic management systems in an urban area – can be made 'smarter' to deliver new benefits to consumers, companies, cities and society.

However, the nature of the data that such smart devices are likely to create, and the speed with which it needs to be analyzed, is placing new demands on data storage, networking, processing and, perhaps most importantly, analytics.

In order to gain insight and value from data generated by IoT, enterprises need to:

- Capture and process data coming from sensors and other devices.
- Ensure interoperability and integration based on data coming from multiple sensors with multiple data formats and multiple protocols.
- Analyze that data in real time to compare it with historical trends.
- Ensure that appropriate responses to the results of that analysis are built in to operational application workflows and business processes.

This paper examines why the Internet of Things is ushering in new data types and new data characteristics, and therefore why new approaches to application development, data processing, data interoperability and operational analytics are required. It also examines why other tried-and-tested techniques for data storage and application development, such as object-relational databases, have the potential to see a resurgence in the era of IoT.

The paper will explain why more traditional data platforms and analytics that were conceived before IoT became a reality are unlikely to cope with the speed, types and sheer amount of data that we're not just likely to see in years to come, but are already seeing today. However, the purpose of this paper is not to focus on the 'things' themselves – the smart devices, machines, sensors and so on – but just how the data that they will create will be stored, processed, analyzed and acted upon.

Key Findings

THE INTERNET OF THINGS WILL PLACE UNPRECEDENTED DEMANDS ON DATA STORAGE, PROCESSING, ANALYTICS AND WORKFLOW. Existing relational databases built with fewer, longer-running transactions in mind are already starting to struggle. That's because we're seeing a shift from systems designed to handle a few hundred events per minute toward technologies that are able to cope with trillions of events per minute in extreme cases.

THREE OF THE KEY DATA TYPES WITH IOT ARE MEASURES AND METRICS, TRANSACTIONS, AND DIAGNOSTICS. Each has different characteristics, and may need to be treated differently from a data processing and analytics point of view – unless they can all be harmonized, analyzed and made actionable in a single platform that is able to provide interoperability between multiple data formats and protocols.

THE CHANGING TYPES OF DATA THAT WE'RE STARTING TO SEE IN THE ERA OF IOT MEAN THAT, IN MANY CASES, RELATIONAL DATABASE TECHNOLOGIES WITH A RELIANCE ON FIXED SCHEMA WILL STRUGGLE TO KEEP UP. While there are a number of newer technologies that can help, object-relational databases that have been maturing for some years and are inherently designed to deliver agility could get a new lease on life.

Introducing the Internet of Things

The term 'Internet of Things' was coined by Kevin Ashton, a British technology pioneer who cofounded the Auto-ID Center at the Massachusetts Institute of Technology (MIT), which created a global standard system for RFID and other sensors.

The Internet of Things is essentially the interconnected set of sensors, machines and smart devices designed specifically to enable the virtualization of the physical world. The devices can create new value, services and perspectives from the sharing and analysis of their observations of the world around them, and are evolving to also be able to 'act' in response to the results of analysis of the data they generate.

IoT encompasses simple tags that convey identity and location; devices with sensors and the computational capacity to provide and communicate results; the gateways and communications facilities to share and distribute that data; and the digital infrastructure to store, analyze and present it to consumers. Nearly every type of business that exists today can employ IoT systems to create cost savings and efficiencies or support new business models. IoT will form the core value proposition for many new businesses.

Given the far-reaching potential, it is no surprise that IoT has been identified as a strategic area of investment for hundreds and thousands of IT suppliers, distributors, systems integrators, ISVs, cloud service providers, network operators, device makers, crowd funders, financiers and application providers. Governments around the globe have made IoT a central component of strategic initiatives that focus on competitiveness, energy and eco-efficiency policies, public safety, quality of life, and citizen services.

Our research also indicates growing interest among enterprises. More than two-thirds of enterprises surveyed by 451 Research's Voice of the Enterprise are already using or piloting IoT.

Figure 1: IoT Adoption

Q. Do you Have an IoT

Source: 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects, 2016

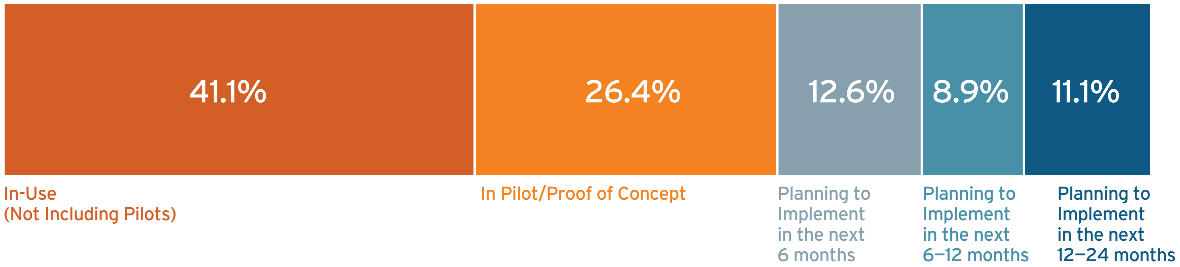

Taken together, IoT is expected to revolutionize how we work; entertain ourselves; drive; monitor our health; shop; interact with friends, children and pets; manage our homes; and move around. Today, much of this potential has yet to be realized due to numerous challenges, including security risks, a lack of IoT standards in some areas, unclear business and monetization models, and high start-up costs. Despite these challenges, the potential of IoT has led to a frenzy of activity as companies look to identify potential IoT opportunities.

Essentially, IoT is expected to generate vast amounts of data from diverse locations that needs to be collected, analyzed and acted upon in real time, thereby increasing the need to better index, store and process such data. The trouble is, some of the more traditional data platforms and analytics were simply not designed with IoT in mind, and are therefore unlikely to be able to cope with the latest demands.

The Impact of IoT on Data Platforms and Analytics

From a data point of view, the Internet of Things is a real game-changer. Traditionally, most transactional systems used to handle commerce were designed with the need to be able to cope with the fact that a consumer might make one or two transactions every few minutes – at the most. In many cases, such as the purchase of a television or an airplane ticket, it's unlikely any one consumer would transact with a company's back-end systems more than once or twice in a week.

It's a totally different scenario in the era of IoT: a sensor or smart device that is monitoring temperature, humidity, vibration, acceleration or numerous other variables could potentially generate data that needs to be handled by back-end systems in some way every millisecond.

Indeed, this is already happening in numerous industries. For example, a typical Formula One car already carries 150-300 sensors, and design teams are adding more all the time. Today, those couple hundred sensors already capture data in milliseconds. Racecars generate 100-200KB of data per second, amounting to several terabytes in a race season.

So while each individual reading from a sensor might translate into a relatively small amount of data – a temperature reading along with a time stamp, for example – there are hundreds or thousands of them being generated each second. For back-end data platforms, this is very different from more traditional transactional systems, which were designed for an entirely different use case.

For them, transactional consistency is paramount, and they need to follow a certain chain of logic to complete each transaction securely. For example, when someone makes a purchase through a website, the systems need to check if the item is in stock, take a credit card payment, update current stock levels, mark the item to be dispatched to the correct address and confirm all of this with the customer.

The transaction is longer-running than simply capturing time-series data coming from a Formula One car, and it can afford to take a few seconds or even minutes to run – after all, the consumer takes a certain amount of time to confirm their details anyway. In the Formula One case, missing a few seconds of sensor data, or being unable to analyze it efficiently and rapidly, could mean the difference between winning or losing a race (which in turn has a knock-on effect of millions of dollars in earned or lost revenue for the team).

In fact, there are three major types of data that come together to make up the Internet of Things:

1: METRICS AND MEASURES (METADATA AND STATE)

This type of data is what most people think of when they consider IoT. It consists of the data that comes from the 'things' themselves – measures from sensors such as temperature, humidity, acceleration, vibration, speed, video feeds, medical devices and biometric data, and so on. In many cases, readings (or small packets of data) are created every few seconds or even milliseconds, so there is a high frequency of data creation. While the packets of data created are small, there are so many readings taken by such a large number of sensors that we have heard of a company in the oil and gas space that is already dealing with over 100TB of such data per day.

2: TRANSACTIONS (COMMANDS)

Rather than simple metrics and measures, 'transactions' – or commands – are slightly more complex, and could include an interaction between two machines, or between a system and a human being. The commands could come in the form of an adjustment to the parameters of a machine or system, such as an alteration to a generator or air conditioning unit. The sort of data that transactions give rise to is likely to be less frequent than metrics and measures, but potentially somewhat more complex.

3: DIAGNOSTICS (TELEMETRY)

This is the type of data that gives insight into the overall health of a machine, system or process. Diagnostic data might show the overall health of a system, but also shows whether the monitoring of that system is working effectively. Diagnostic data is also likely to be somewhat more complex than metrics and measures, and potentially an alert that a system is no longer functioning within normal parameters and might need further analysis and investigation to determine the root cause. This could also include predictive analytics in order to support preventive maintenance of the sensor itself. To some extent, the goal of diagnostic data is to be able to filter out all of the 'normal' data in order to focus on the data that suggests there is a problem that needs further investigation. Equally, too many false positives reduce the efficacy of the monitoring system.

As we will discuss later, one of the challenges for IoT is how these different types of data can be brought together to add value in new ways, especially as new sensors are added over time, delivering data in different formats and using different protocols. It's a challenge that will impact the way that companies ingest data into a data platform and subsequently analyze it. Indeed, it will also impact the choice of data platform that is best suited to providing interoperability, given the changing data speeds and characteristics that we are witnessing in the era of IoT.

Introducing Total Data

Total Data is a term that was introduced by 451 Research, offering our perspective on the trends shaping the data platform and analytics sectors.

As Figure 2 illustrates, Total Data goes beyond issues related to the size and complexity of data, and the rate at which it is being produced, to also take into account the evolving ways users want to engage with their data – with more frequent and more exploratory analysis – as well as the implications from ongoing dependency on existing data processing technologies.

Additionally, while Total Data is focused on emerging and new data-processing and -analysis technologies, it also takes into account the importance of nontechnological factors, such as domain knowledge and the underlying business case, which provide the overall context to ensure that the right data is being captured and analyzed.

At the heart of it all is recognition that the biggest driver of change in the data processing industry is not functional or technological, but economic: the cost of storing, processing and analyzing data, as well as acting on the data by adapting operational workflows as required.

It's worth noting that many of the different technologies that can be considered part of Total Data end up producing data with different characteristics. As a result, there are more types of data in the typical enterprise, and it's often accessible at different speeds. And as everyone is only too aware, the total volume of data in the enterprise is only going one way – up – a trend that is set to continue in the era of IoT. All in all, it's something of a perfect storm in terms of data processing, analytics and interoperability.

Figure 2: Total Data

Source: 451 Research

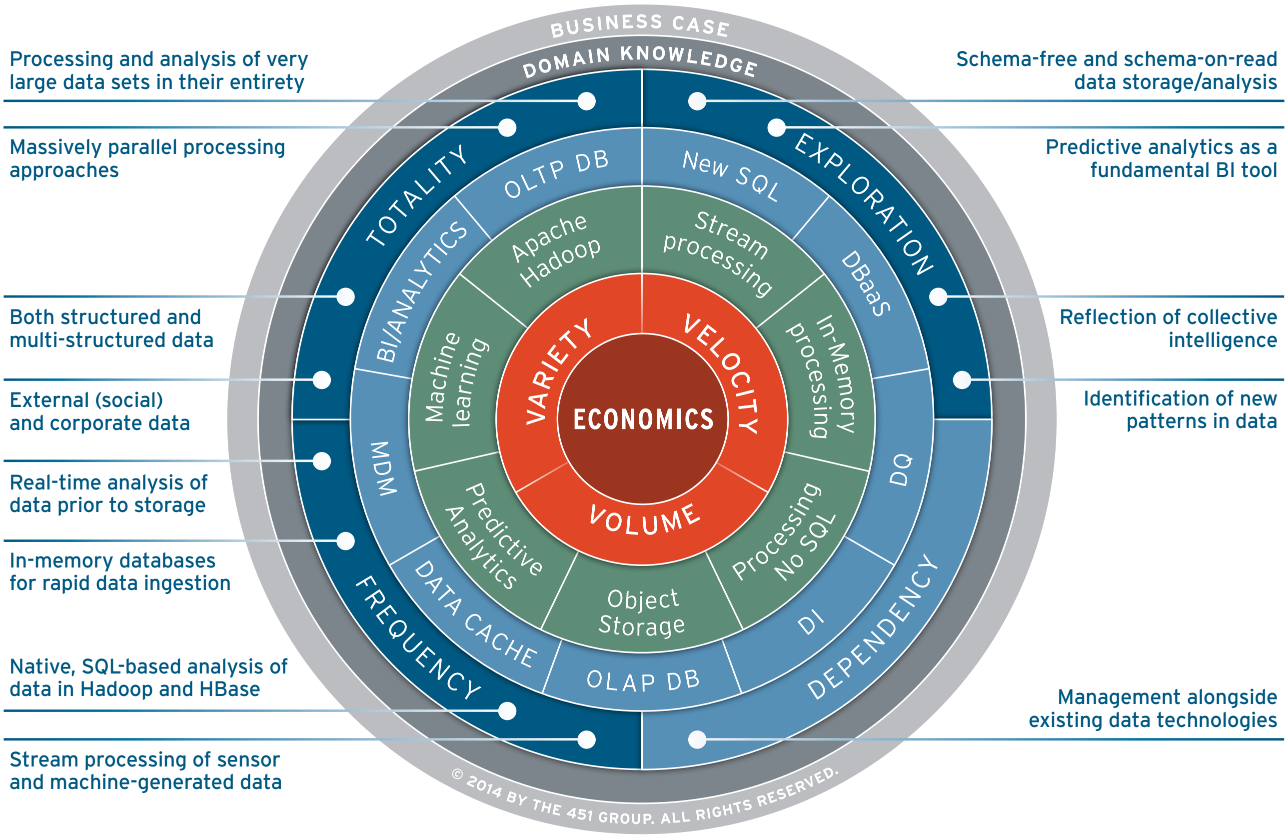

Multi-Modal Databases: Back to the Future?

Database products have traditionally fallen into two categories: Online Transaction Processing (OLTP) and Online Analytical Processing (OLAP). General-purpose databases can be tuned to work for both, although not necessarily at the same time.

This has been changing though, with increasingly more databases that claim to be able to handle both operational and transactional workload needs, often through the copious use of in-memory techniques, but sometimes by adding a columnar store that is optimized for analytics, alongside the more traditional row store that handles OLTP.

In the era of IoT, we expect this 'multi-modal' data platform trend to increase as combined operational and analytical processing becomes a requirement to support operational analytics use cases and enable organizations to react rapidly to insight generated by IoT data.

For example, traditional approaches to analytics – whereby data from operational applications is transformed and loaded into a spare data warehouse – might help an auto parts manufacturer spot trends in reports of faulty parts that may require a recall to protect the brand image. However, real-time operational analysis of sensor data could be used to detect quality-control variance in the manufacturing process in order to warn that tolerance thresholds are about to be threatened, thereby preventing the faulty parts from being manufactured in the first place.

Some of the latest approaches to handling transactions and analytics – with different types of data coming into a data platform with different speeds and latencies – have included in-memory databases, NoSQL approaches, in-memory data grids, data streaming and various work around the Hadoop open source data processing framework.

But it's not simply a case of 'out with the old, in with the new.' That's because databases that are able to handle numerous complex data types aren't all as modern as you might think. It turns out that object-relational databases, originally conceived as an alternative to relational databases to support applications that needed to store additional data types, such as images, videos and spatial data, may well see something of a renaissance in the era of IoT.

Object-oriented databases are inherently more flexible than relational databases that rely on fixed schema, since object inheritance provides interoperability of data in multiple formats or protocols. As noted below, this interoperability is likely to be critical in enabling enterprises to keep up with the rapid pace of innovation, the evolution of IoT protocols, and the obsolescence of sensors and processors.

Additionally, many object databases offer mature support for SQL-based transactions and analytics, enabling enterprises to bring their existing applications, tools and skills to data from IoT.

In addition to requiring support for multi-modal databases that deliver combined operational and analytical processing, the range of data formats generated by IoT is also expected to drive requirements for multi-model databases that can store and process data in multiple models – relational, key value, document and graph models, for example.

Having multiple data models within a database is not necessarily a new concept, but we have started to see an uptick in database vendors adopting additional models. The definition of a multi-model database is somewhat subjective, and the idea – at least in principle – has been around for some time. If we adhere to a strict definition, then one could make an argument that relational databases, which provide extensible architecture, the ability to store data as XML documents or objects, and varied index types, could technically qualify as providing multi-model capability, which was present long before NoSQL databases started to gain traction. More recently, the term multi-model has been used to describe nonrelational-database vendors offering support for storing and processing data as key value pairs, documents, objects and graphs – even via SQL.

One of the oft-cited advantages of multi-model databases is that they address the polyglot persistence challenge (i.e., that multiple databases support a single application because of the need to store data in multiple formats). That fundamentally makes sense, but in practicality, it can present some real challenges if there are multiple databases driving a single application. Incorporating multiple databases based on separate data models makes for a highly complex environment, let alone the time and resources required to maintain them.

Unlike some rival approaches, object-relational databases have the advantage of being able to keep those different data types in the same place, easing not only analytics across that variable dataset, but also application development. After all, the alternative is often to have different data platforms (SQL, NoSQL, Hadoop, etc.) that all have different programming paradigms, to support multiple data models, with the associated requirements for maintaining multiple data platforms and dealing with data integration challenges.

Another key requirement in IoT is the ability to analyze not only the very latest data that is coming in from sensors and smart devices, but also the ability to compare that data with historical data and trends. Indeed, it could be argued that IoT projects generate not only more real-time data, but also the requirements for larger volumes of historical data, which is used for trend analysis and generating the business logic that drives real-time decision-making. One way to do this might be to keep the freshest data in a NoSQL or in-memory database, and persist older, larger datasets to another platform such as Hadoop.

This is not without its challenges when it comes to analyzing the data – neither in terms of the management nor development overhead of those different data silos.

While object-relational databases are able to cope with the numerous different data types that are found in the era of IoT, object programming is generally considered simpler and more efficient than rival approaches. There is more chance of code reuse; it's easier to isolate the user interface from the underlying application code, so changes are easier to make without rewriting; and it's easier for developers to improve one area of an application without potentially breaking another area – because objects essentially act like black boxes (also known as encapsulation).

IoT: Driving Analytics to the Edge?

Another trend that we are seeing is the growing adoption of edge analytics, sometimes also known as fog computing. This is all about doing some filtering of the data at the edge of the network – closer to the source of the data – in order to reduce the chance of that data saturating networks and datacenters.

Figure 3: Edge Computing Architecture and Components

Source: 451 Research

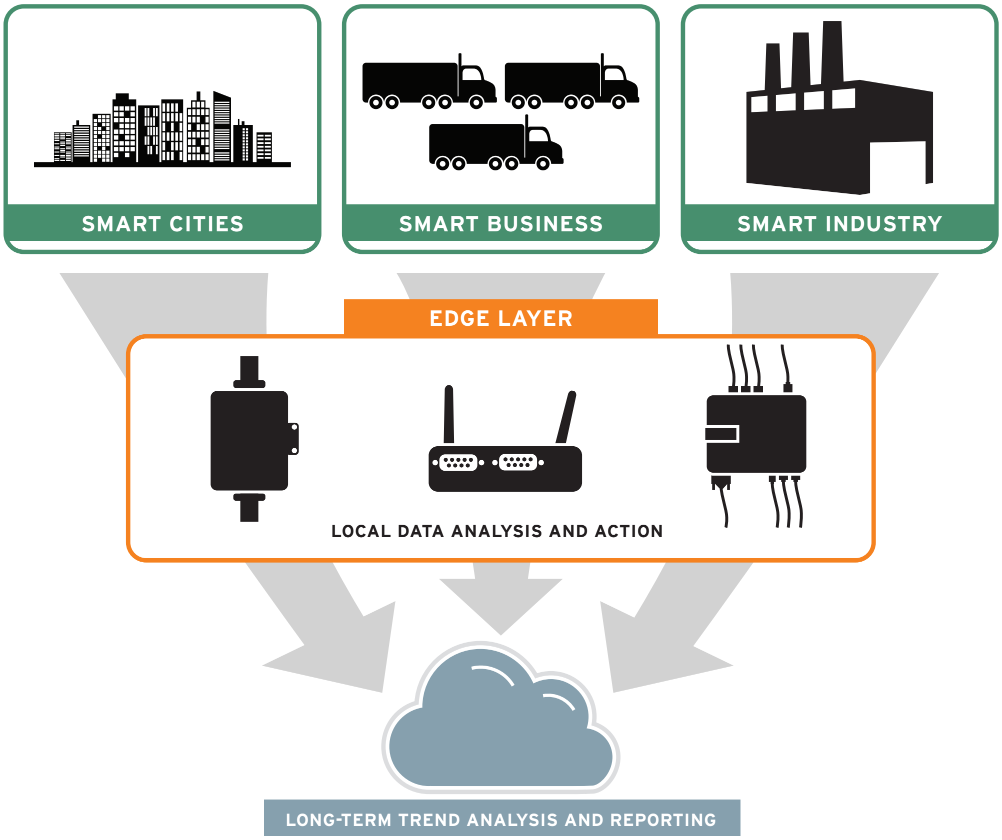

As can be seen in Figure 3, the edge computing layer is most likely to be used to filter 'fresh' data – data that is streaming in from devices, sensors and so on. Using edge analytics, it's possible to remove any data that fits within a normal range, and instead only send data that is an anomaly, or varies from the norm in some way, for further analysis.

However, as can also be seen in Figure 3, there's a need to compare the latest data with historical data – without the historical context it's far harder to glean any real value from the latest information coming from sensors and smart devices. This kind of analysis – blending the latest data with historical data – is less likely to happen at the edge of the network, and is more likely to occur in a datacenter or in the cloud.

It is important to note that the analytics of processing at the edge does not occur in isolation. The definitions of a 'normal range' and 'anomaly,' as referenced above, may change over time, and the edge layer needs to be constantly updated with new business logic and workflows as requirements evolve based on long-term analysis of real-time and historical data.

This also highlights another element in data processing and analytics that should not be overlooked: the need for data interoperability to enable organizations to keep up with the pace of innovation as new sensors, using new protocols and formats, are added to IoT.

Data Interoperability: A Key Enabler for IoT

As noted above, IoT is expected to generate vast amounts of data from diverse locations that must be aggregated very quickly. As such, we expect the Internet of Things to place additional demands on requirements for data interoperability and data integration.

Data integration uses tools to deal with the synchronization, standardization, transformation, mapping, quality and transport of data between different applications and systems. This is especially necessary for IoT because it will play a role in assuring data quality from a variety of IoT applications and devices for optimal analytics capabilities. With data integration, the data is abstracted and 'mapped' between different data sources and targets. The structure and format of the data are important at this stage. Data validation checks to see if the source and target formats are compatible, and if not, the data is rejected. Transformation is then used to standardize or 'normalize' the data structure, to make it useable to each application or system.

Integrating all the data related to an IoT project into one place for analysis is therefore a considerable challenge; however, a more fundamental challenge is that of providing data interoperability to ensure that data from multiple sensors can be synthesized despite the use of multiple protocols, and to allow for sensors to be added, altered and removed without an impact on existing application workflow.

Consider the example of an airline that is taking sensor readings from its aircraft in order to (among other things) enable predictive maintenance, as well as improve operational efficiency and safety.

Given the range of aircraft manufacturers, designs and ages, it is inevitable that different sensors will be utilized across an airline's fleet. British Airways, to pick an example, currently has 289 different individual aircraft in its fleet, including 14 different aircraft designs, from three different manufacturers. Those sensors will inevitably utilize different data formats and protocols, even when the data being measured – for example, engine temperature – is the same.

The challenge is therefore not just to integrate the data from multiple sensors and bring it all into one place for analysis, but to provide the interoperability that can normalize the data from multiple sensors delivered in multiple formats using multiple protocols.

This challenge will be exacerbated as older aircraft are updated and new aircraft are put into service, with the requirements that the underlying data platform can continue to integrate and synthesize data to allow for new formats and protocols without needing to change business workflow or rewrite associated applications.

In the world of IoT, several protocols have been developed for specific purposes – for example, IoT devices or machines must communicate with each other (D2D or M2M); device data must be collected and sent to servers (D2S); and servers have to share device data (S2S), possibly providing it back to devices, analysis programs or people. The protocols that enable these exchanges include:

MESSAGE QUEUE TELEMETRY TRANSPORT (MQTT)
A protocol for collecting device data and communicating it to servers (D2S).

EXTENSIBLE MESSAGING AND PRESENCE PROTOCOL (XMPP)
A protocol for connecting devices to people (adaptation of D2S – i.e., people are connected to servers).

DATA DISTRIBUTION SERVICE (DDS)
A fast bus for integrating intelligent machines (D2D).

ADVANCED MESSAGE QUEUING PROTOCOL (AMQP)
A queuing system designed to connect servers (S2S).

CONSTRAINED APPLICATION PROTOCOL (COAP)
A web transfer protocol used with constrained nodes and networks, designed for M2M applications (e.g., smart energy and building automation).

Companies will struggle to handle this array of protocols by writing interoperability scripts of their own – these will quickly prove brittle and a barrier to making changes to applications and back-end systems.

Summary

In this paper, we have explained the origin and implications of the Internet of Things from a data, data-platform and data-analytics perspective. In particular, we have seen how the three different types of data that are likely to pervade IoT – metrics and measures, transactions, and diagnostics – have different characteristics that may not all be ideally ingested, stored and analyzed in more traditional relational databases.

We then went on to look at what 451 Research calls Total Data – our way of categorizing the different latencies, varieties and volumes of data that we increasingly see in the era of the Internet of Things.

Finally, we looked at some of the options for coping with the changing nature of data in the era of IoT, including in-memory databases and data grids, NoSQL databases and Hadoop, and object-relational database technology.

We explained why we think object-relational technology, although not new, will be given a shot-in-the-arm thanks to IoT due to its ability to handle multiple data types, and thanks to its simplified programming model that makes application development faster and more efficient. We have also espoused the benefits of being able to do all analytics – not just on different types of data, but also on data of different ages – on one data platform.

Finally, we addressed some of the issues around analyzing data at the edge of the network in order to avoid network or datacenter saturation, and the continued importance of data interoperability, especially to maintain consistency given the numerous different protocols that are being adopted in the era of IoT. These will become difficult to handle if companies aren't using some kind of data management technology capable of abstracting away at least some of the complexity.

As with many things in data processing and analytics, there will always be more than one way to crack a nut, and it's nearly impossible to say definitively which approach is the 'best' way to handle analytics and interoperability in the era of IoT.

However, as we see IoT projects rolled out to support industries such as financial markets, manufacturing, healthcare and transportation, it is clear that the applications enabled by IoT are quickly becoming mission-critical. As such, enterprises are advised to be cognizant of the need for reliable and mature data platforms that enable them to not just store, process and analyze data from sensors based on multiple data formats and protocols, but also to act on it, by providing a platform for operational analytics applications – driven by the results of that analysis – that will enable the business to act upon it.