451 Research | PATHFINDER REPORT

Data Analytics at the Edge

COMMISSIONED BY

SAS

DELL Technologies

©COPYRIGHT 2019 451 RESEARCH. ALL RIGHTS RESERVED.

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 THE AUTHOR

MATT ASLETT

RESEARCH VICE PRESIDENT

Matt Aslett is a Research Vice President with responsibility for 451 Research's Data, AI and Analytics Channel – including operational and analytic databases, Hadoop, grid/cache, stream processing, data integration, data governance, and data management, as well as data science and analytics, machine learning and AI. Matt's own primary area of focus currently includes distributed data management, data catalogs, business intelligence and analytics, data science management, and enterprise knowledge graphs.

Table of Contents

Executive Summary

The value of data in driving business decisions has long been understood by industrial companies. Traditionally, however, business intelligence has almost exclusively involved humans analyzing data generated by enterprise applications. To remain competitive, this time-consuming, arduous process must be transformed, especially given the innovations emerging with Industry 4.0, as well as growing expectations of real-time customer engagement and responsiveness.

We have moved from the transactional era, through the interaction era to the engagement era, in which industrial enterprises have recognized that they must store, process and analyze as much, if not all, data that is available to them in order to survive and thrive in the digital economy. This includes data produced by the myriad of sensors, embedded computers, industrial controllers and connected devices such as vehicles, wearable computing devices, robots and drones that make up the emerging Internet of Things (IoT).

Figure 1: Increasing data volumes and the engagement era

Source: 451 Research

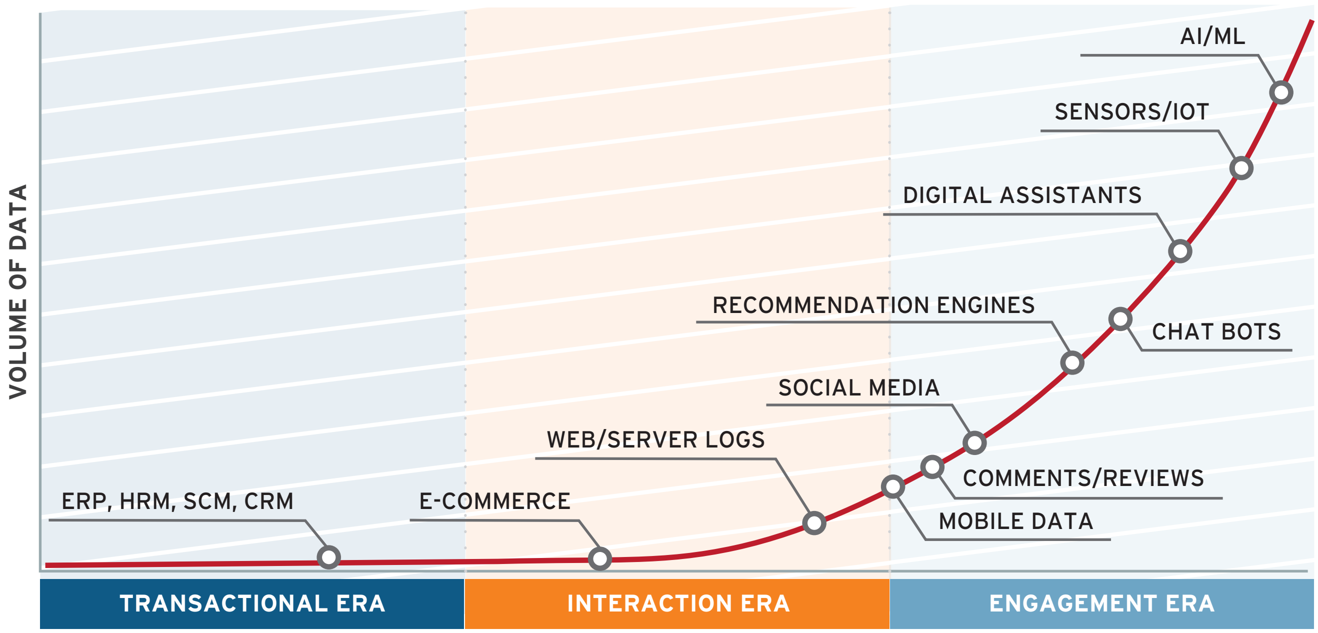

Processing and analyzing this data at the edge, close to its point of generation, has the potential to enable improved processes and operational efficiencies, as well as a better return on investment and competitive advantage. Edge data processing is a critical element in the ability to deliver the ultra-low latency required for real-time applications while minimizing the inefficient movement of unnecessary amounts of data to the cloud or remote on-premises datacenters for processing.

Indeed, given the volumes of data being produced by the Internet of Things, and the rate at which it is being generated, the case could be made that it is no longer a matter of enterprises deciding what data to process at the edge, but that processing data at the edge is a necessary step in deciding what data to store and analyze in centralized infrastructure.

With data processing at the edge, enterprises are enabling more cost-effective data storage and gaining insight from their IoT data earlier. This has the potential to be translated into millions of dollars saved via the reduction of unplanned down and the realization of operational efficiencies, as well as providing the opportunity for the development of new revenue streams and differentiated customer experiences.

Despite the potential, many organizations have yet to unlock the advantages of processing and analyzing data at the edge, and much of the data being produced at the edge goes untapped or underutilized. That is beginning to change. Data from 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018 indicates that processing data at the edge is on the rise, with 62% of respondents processing data on the data-generating device or nearby IT infrastructure, compared with 44% in 2017.

Figure 2: Increasing adoption of IoT edge processing

Source: 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

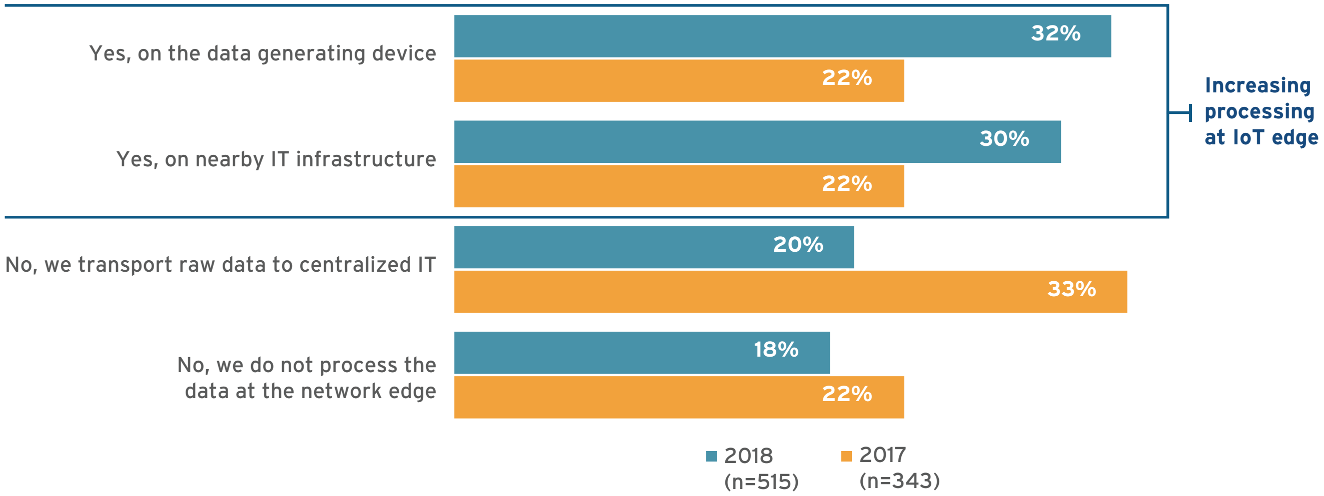

Even those companies that understand the potential benefits of processing data generated at the edge can struggle to generate value from it, however. Given the multitude of different devices that can be considered 'the edge' – as well their respective performance constraints – identifying where, and why, to perform data processing is not a simple task. Additionally, deciding when, and how, to move IoT data from the edge is also essential to getting the most value from edge compute architecture.

The Rise of Edge Computing

Reasons for Rising Adoption of Edge Computing

There are a multitude of reasons for processing and analyzing data at the edge. However, the industry has coalesced around seven primary reasons for edge computing: the need for low/ultra-low latency; the cost of bandwidth; the lack of high-speed connectivity; security; data sovereignty; reliability; and interoperability with legacy systems.

Figure 3: Seven primary reasons for edge computing

Source: 451 Research

Defining the Edge

One of the challenges in considering when, where, why and how to process and analyze data is that enterprises and vendors alike have different definitions of 'the edge' depending on where data processing is currently carried out and (for vendors in particular) the relative data processing capabilities of their existing products and services. For industrial enterprises that have traditionally processed and analyzed the majority of their data in centralized datacenters (either on-premises or in the cloud), anything that is not a central datacenter could legitimately be considered 'the edge.'

Rather like early astronomers who studied Saturn with relatively weak telescopes and observed that it was encircled by a ring, their perspective of edge is defined by their distance from it. Later astronomers with more powerful telescopes effectively gained a closer vantage point from which they could identify that Saturn was, in fact, encircled by multiple rings. Similarly, those with a vantage point closer to the edge recognize that it cannot be considered a single distinct entity, but a spectrum of multiple edge devices.

Figure 4: The spectrum of edge computing devices

Source: 451 Research

The definition of the edge, therefore, is relative to the position of the enterprise or vendor perceiving it:

- Mobile network operators (MNOs) perceive the edge as the end of their radio access network (RAN), and the potentially large opportunities for multi-access edge computing (MEC).
- Datacenter operators may view the edge as infrastructure deliberately deployed at key locations to minimize Internet latency.
- High-performance computing and edge server vendors view devices at remote sites, either within datacenters or sitting in cages and closets, as their edge.
- Network equipment providers (NEPs) focus on the edge of the network with robust gateway devices connecting non-networked equipment and performing edge analytics.
- Understandably, vendors of sophisticated sensors or embedded computing, or even complex devices with integrated analytics, view their offerings as 'the edge.'

The distinction between these categories of edge devices is important – not just in ensuring a common vocabulary – but also in understanding the relative data storage and processing capabilities of each.

As Figure 5 illustrates, data processing requirements range from ultra-low latency and real-time analysis at the edge, through medium latency and local data processing at the 'near edge,' to high latency and high-capacity storage and networking in centralized datacenters. Enterprises need to take these latency capabilities into account when considering where to execute data processing workloads.

Figure 5: Latency requirements from edge to datacenter

Source: 451 Research

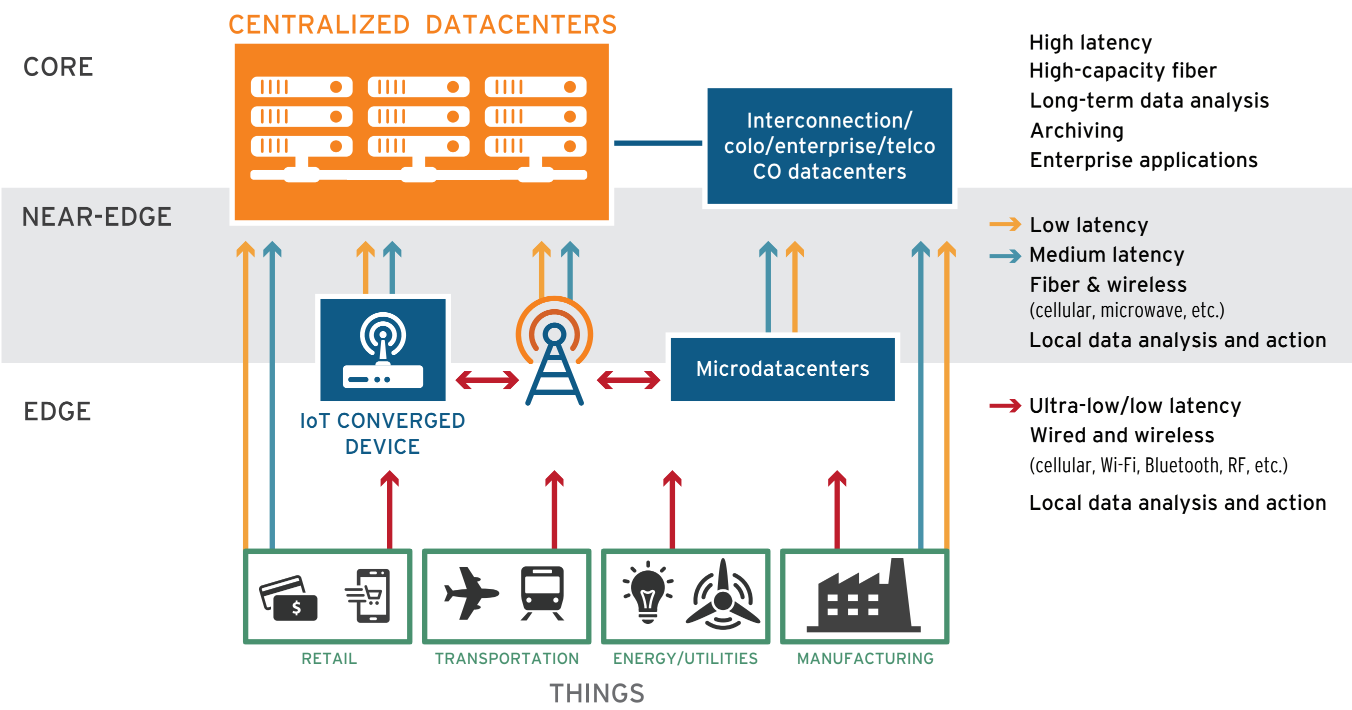

The Edge as a Complement to Centralized Infrastructure

As Figure 5 suggests, the cloud and the edge are complementary in relation to analytic processing workloads. Indeed, there is a good argument to suggest that the increasing trend of capturing data at the edge is elevating demand for cloud-based analytics in order to process at least some of that data, as well as for advanced workloads such as the development and testing of machine learning models.

451 Research data[1] shows that 63% of IoT users initially store and analyze IoT data on centralized infrastructure, including public cloud environments; on-premises datacenters; and both owned and third-party off-premises datacenters. The remaining 37% of respondents are initially storing and analyzing IoT data on what might be considered the edge, including on or near the device, intelligent gateways, micro-datacenters, and network operator infrastructure. Additionally, more than half of respondents report that having performed this initial analysis at the edge, they also subsequently perform additional analysis on the data, the majority of which is performed on centralized infrastructure.

Figure 6: Initial locations for storing and analyzing IoT data

Source: 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

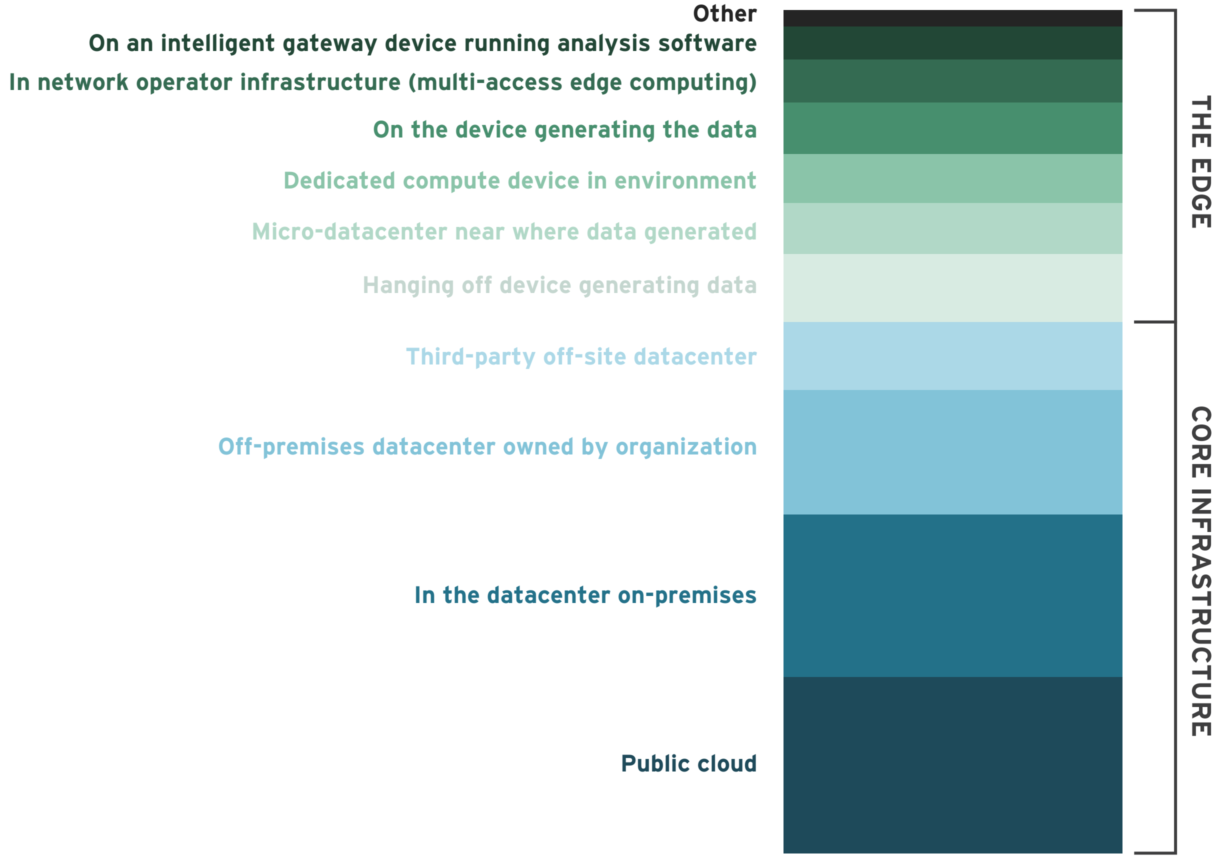

1. 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

In terms of the nature of the data processing being carried out at the edge, 451 Research data shows that analytics at the network edge/perimeter is most popular, followed by the aggregation of data from multiple sources. Meanwhile a growing number of respondents are performing processing designed to reduce data volumes at the edge, including data filtering, data synchronization and data modeling.

Figure 7: IoT edge data processing

Source: 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018
Q: What data process(es) do you undertake at the network edge/perimeter for IoT? Please select all that apply.

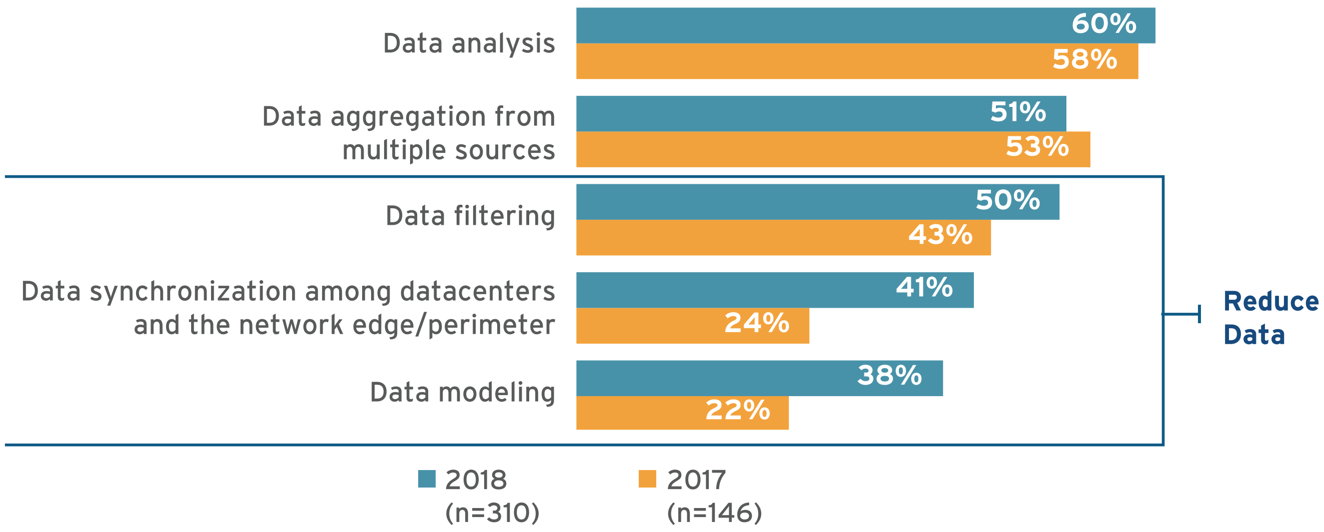

Successfully taking advantage of edge data processing requires a combination of technologies that fulfill all aspects of the data processing lifecycle: data ingestion and inventory (creating a catalog of the available data); data preparation and data delivery; and data discovery and data visualization. Although some automated data analysis is being performed on or near edge devices, data delivery is a critical component in ensuring that the results of this analysis, aggregation and filtering are also made available in centralized infrastructure for data analysis and visualization by data scientists and data analysts.

One of the technologies that is increasingly used to support data processing and analytics, from edge to core, is stream processing, which enables the processing and analysis of events or streams of events in real-time. Indeed, stream processing has a dual role to play in edge data processing: it can be used to develop applications that perform low-latency data ingestion, aggregation and processing of data from streaming IoT sources at the edge, as well as streaming the resulting data to centralized infrastructure for long-term storage and further analysis.

AI at the Edge

While much of the analytics performed at the edge involves diagnostic analysis (filtering and monitoring) as well as real-time analysis of metrics and measurement, data from 451 Research[2] indicates that the IoT analytics workloads being performed at the edge are becoming increasingly sophisticated, with a shift away from rules-based actions and toward artificial intelligence and machine learning.

While the vast volumes of data and compute power needed to train machine learning models is not an ideal fit for edge computing, inference – or running the trained models on new data – is a perfect fit to be performed on data as it is being generated. In particular, anomaly detection models can be pushed down and executed directly on edge devices to drive alerts that prompt immediate action, while the data causing those alerts can also be transmitted to central data processing locations for further analysis.

Figure 8: Streaming analytics and machine learning across the analytics lifecycle

Source: 451 Research

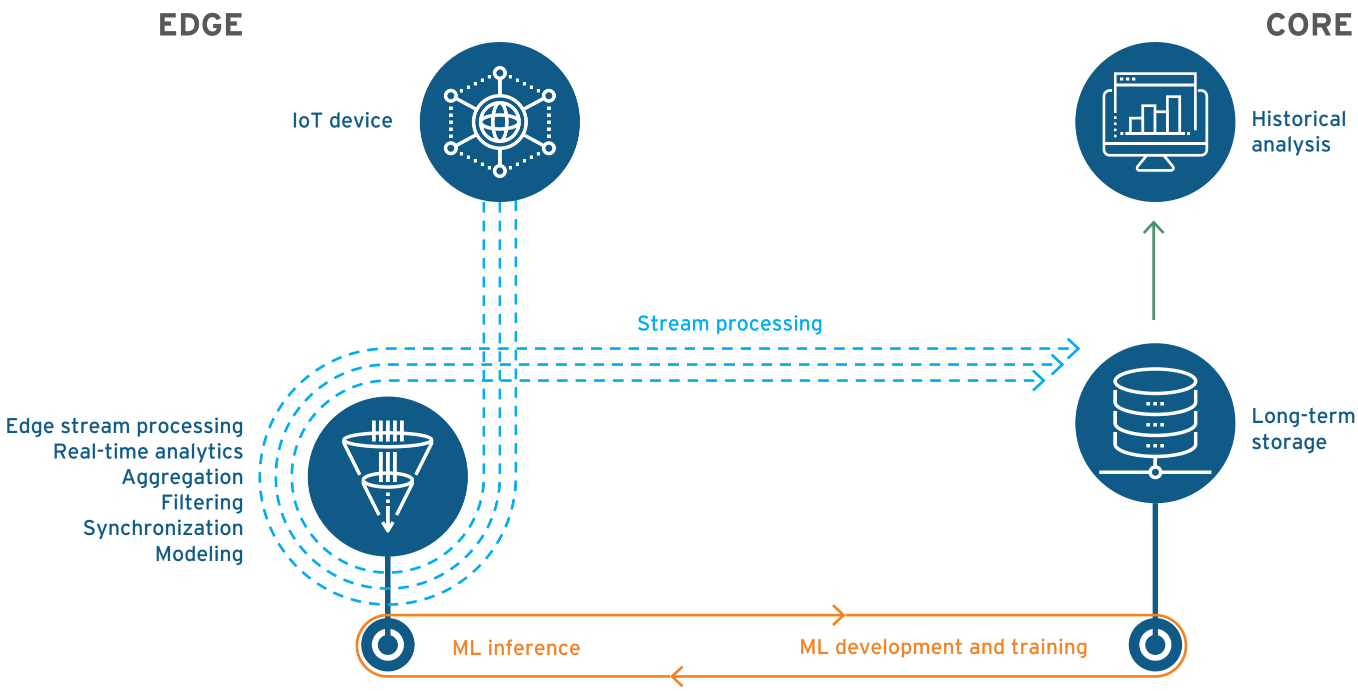

The benefits of this combination of IoT datasets and edge analytics include avoiding downtime and generating completely new revenue streams – for example through preventative maintenance subscription packages. Enterprises at the forefront of the overlap of IoT and machine learning have the potential to create meaningful differentiation in terms of increasing operational efficiency and competitive advantage through improved customer engagement and care.

2. 451 Research: Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

DRIVERS OF EDGE ANALYTICS IN IOT

The success or failure of any IoT project comes down to the ability for a business to transform IoT data into timely, accurate and actionable insights for the business. Data analytics is the critical enabler of faster insight from IoT data – reducing latency, speeding time to insight, streamlining storage and bandwidth costs, and addressing privacy and data-sovereignty concerns – by enabling the analysis of data at the edge, without all data needing to be shipped back to centralized IT locations.

Data from 451 Research[3] indicates that the number one reason for edge computing is to improve security (cited by 49% of respondents in 2018). This can potentially be achieved thanks to the decentralized architecture, which reduces the potential for a compromise taking down a whole network and reduces the volume of data that is transmitted to centralized infrastructure. However, security is followed by multiple reasons that involve faster data processing, including the speed of data analysis and the velocity of data transmission. Edge analytics performed on IoT data in real-time has the potential to unlock the value of IoT systems investment for many vertical industries.

In addition to early adopter IoT industries such as manufacturing, transportation, connected home and wearables, adoption has also accelerated in healthcare, agriculture, smart cities and retail. In healthcare, for example, edge-based analysis of patient data can detect anomalies such as an irregular and rapid heart rate that suggest critical situations, and alert patients and doctors to take preventative action.

Each of these industries makes use of IoT applications where real-time responsiveness or localized data processing can drive positive business outcomes. For example, localized analytics conducted on industrial machine data can reduce the time between identifying a potential issue (such as manufacturing device failure) and fixing it to proactively avoid unplanned downtime (which might involve halting the entire production line).

While localized analytics can drive faster insights, it also reduces the amount of aggregate data traffic that needs to be transported to centralized IT locations. In that context, only signals that 'matter' are transported for action or long-term storage. This is particularly relevant when very-high-bandwidth data streams are involved, such as video analytics. For instance, rather than transporting the full-bandwidth video files, edge-based video analytics can automate surveillance monitoring of a hazardous environment (such as an oil rig) and provide alerts when a potential security or safety issue is detected.

In manufacturing, visual analysis of components on the production line and comparison to models trained on historical defects can be used for quality control. Local processing of the images on the factory floor enables the identification of defective components without centralized processing of the entire video stream.

Reducing the volume of data transmissions contributes to faster data processing but could also potentially mean lower storage costs, one of three cost-related IoT drivers alongside lowering network connectivity and hardware compute costs. As figure 9 indicates, in the last year, there has been a significantly increased focus on edge drivers that lower costs, which perhaps points to an increased realism about the cost of cloud processing and the most appropriate location for data processing.

3. 451 Research: Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

Figure 9: Drivers for processing IoT data at the edge

Source: 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

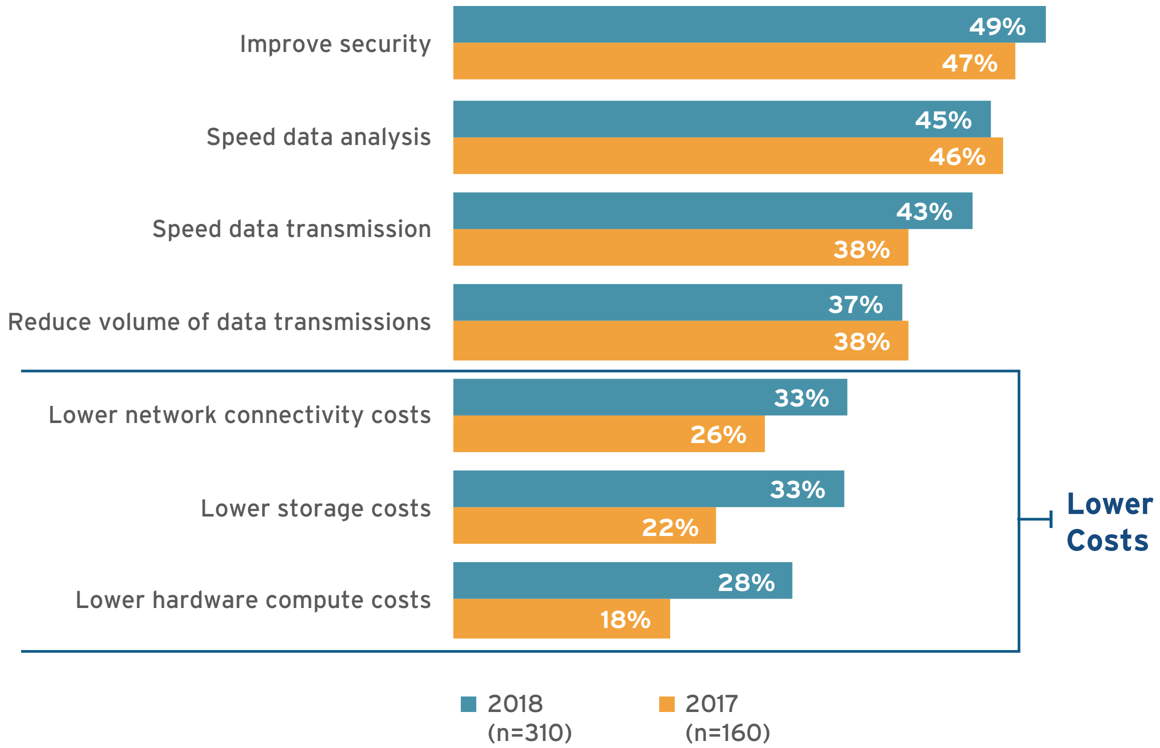

Edge analytics enables analysis of IoT data without having to move it to a centralized location. That does not eradicate the value in also moving some, if not all, of the data to a datacenter for long-term storage and analysis, however. This highlights the symbiotic role that exists between edge and centralized computing in industrial IoT initiatives. It is also one reason why established data management and analytics vendors have a good story to tell in relation to IoT: being well-practiced in terms of efficiently moving data and providing integration between the edge and central data analytics environments.

Given the potential value at stake, it's no surprise that it is not only the established data management vendors that are lining up to deliver edge analytics offerings focused on IoT. Each of the major cloud providers has released edge computing products, while major IT equipment vendors are also in the mix with edge hardware-based solutions, as well as telecom providers that see an opportunity to either directly support edge analytics via their own B2B/B2C solutions or by opening their infrastructure up to partners to accomplish the same.

For enterprises, there is a plethora of potential approaches and products available, each of which might have different advantages for the range of stakeholders – not least operational technology (OT) and information technology (IT) professionals, as well as data analysts and business users. It is important to understand the merits and risks of each vendor and its specific approach.

While the most appropriate location for analytics (from edge, to near-edge to centralized) will be driven primarily by latency and bandwidth requirements, enterprises will be wary of locking themselves in to specific data processing products and services. The utilization of platform-independent cloud-native architecture based on loosely coupled microservices from edge to cloud can provide some peace of mind by enabling workload portability as data volumes grow and data processing needs evolve.

Edge Use Cases and Workloads in Industrial Settings

Manufacturing

The primary benefit of greater intelligence in the manufacturing sector is automation. At the most basic level, this means monitoring the manufacturing and production process with a view to ensuring that the various machines in the production line are operating as desired. While traditionally that has meant monitoring to ascertain simply whether a machine was functioning or not, it increasingly involves monitoring more nuanced information such as equipment performance, which can be used to increase productivity levels, as well as production quality. Taking monitoring to the next level means monitoring not just the production process but utilizing monitoring to predict the risk of failure and deliver preemptive maintenance to drive improvements in productivity and operational efficiency.

Figure 10: IoT use cases implemented/planned for implementation in manufacturing

Source: 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

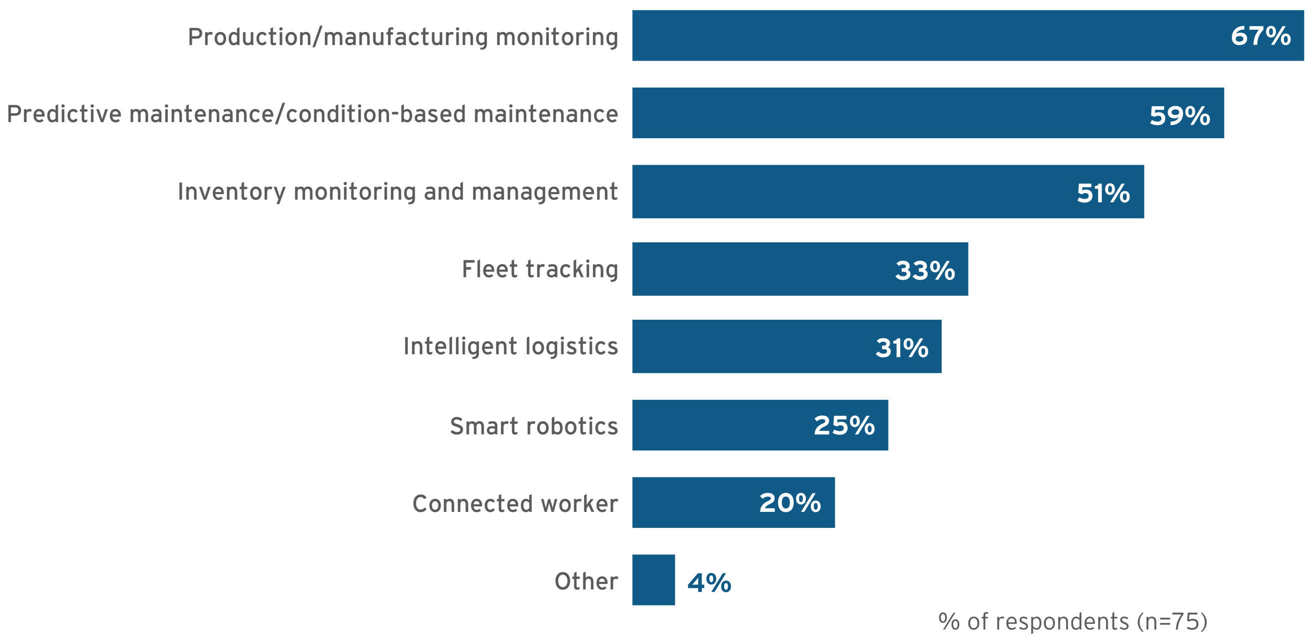

Manufacturing isn't just about the equipment that is used in the production process, but also the various components and raw materials that are involved. To ensure the smooth running of the production line, inventory monitoring and management are critical, as well as fleet tracking and intelligent logistics to ensure that the components and raw materials are available as and when required, according to just-in-time processes.

Edge processing is particularly relevant where network access is intermittent or is limited by cellular or satellite bandwidth, such as fleet management. Other key use cases include smart robotics, with an example being real-time monitoring of precision engineering robots based on data streams indicating factors such as vibration, velocity and temperature, which can be analyzed to assess production quality as well as predictive maintenance.

Monitoring connected workers for health and safety reasons is also a primary use case. Preventing workplace injuries is a high priority, both in terms of employee well-being and avoiding costs associated with lost production.

Transportation

One of the most exciting and intriguing areas when it comes to IoT is transportation. Self-driving vehicles are at the forefront of the convergence of artificial intelligence and the Internet of Things, and there is growing interest in the potential for autonomous vehicles in industrial settings to improve the efficiency of the transportation of goods and raw materials, for example.

A survey of nearly 1,000 drivers in North America and Europe conducted by 451 Research in late 2018 showed that the majority of consumers believe that full self-driving (Level 5) cars would be available for purchase within the next five years or earlier, which is actually ahead of most automotive manufacturer expectations.

Figure 11: Consumer expectations for Level 5 autonomous car availability

Source: 451 Research

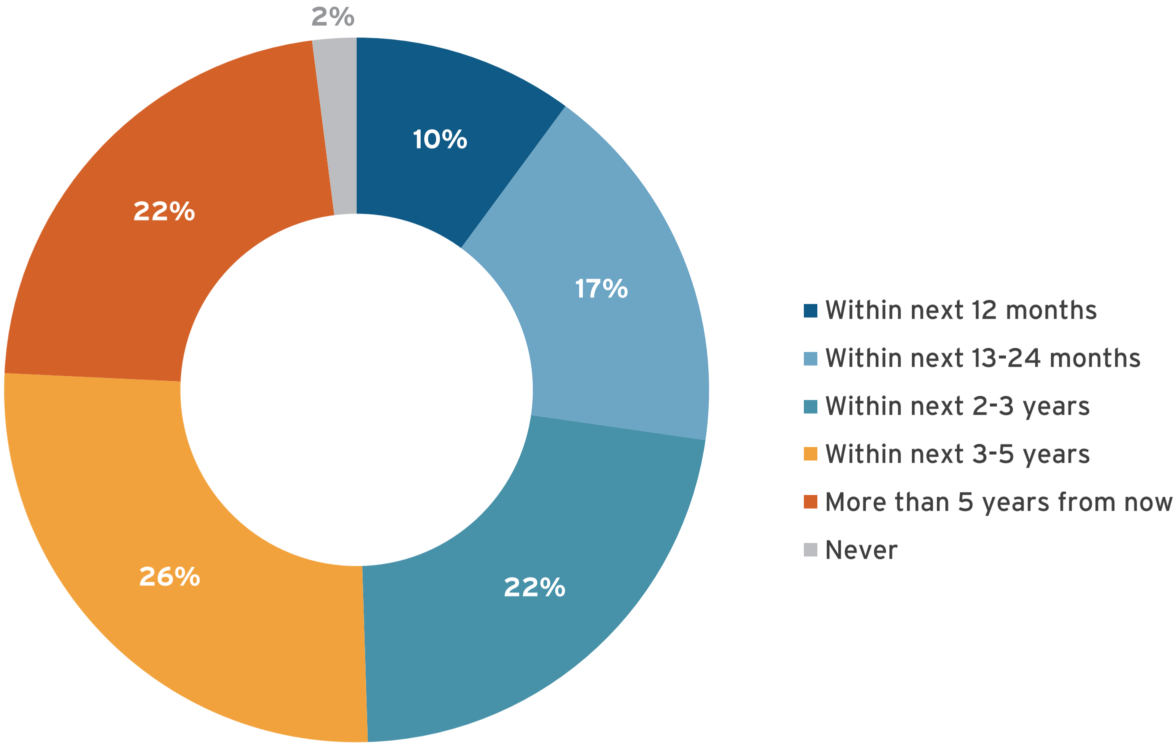

Self-driving vehicles are also a key area for edge data processing and analytics since automation requires in-vehicle data processing due to latency issues, as well as the need to avoid potential failures due to network connectivity issues.

Semi-autonomous advanced driver-assistance systems (ADAS) are a step in the direction of full autonomous vehicles, and ADAS is quickly becoming adopted as a standard feature that is increasingly legally mandated for new consumer automobiles. Just as seatbelts, airbags and now backup cameras have become mandatory equipment, automatic emergency braking, blind-spot detection, self-parking and adaptive cruise control are becoming commonplace in nearly all new automobiles.

However, data from 451 Research[4] indicates that the primary issue today in the transportation sector is not autonomous vehicles but keeping track of existing vehicles and their contents through fleet tracking and logistics, as well as supply chain logistics management and asset/shipment tracking. Other key use cases include vehicle-to-vehicle (V2V) communication, predictive maintenance and diagnostics, energy and fuel optimization, connected car services, and remote access to in-vehicle functions.

4. 451 Research: Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

Energy and Utilities

The energy and utilities sector is one of the most prominent adopters of IoT data analytics, driven by regulatory compliance, as well as the potential efficiency improvements that can come from monitoring and maintenance of the complex web of equipment and components that make up the smart grid and related initiatives, such as oil and gas exploration and monitoring.

The concept of the smart grid is in itself a holistic use case for IoT edge analytics – involving maintenance, upgrade, monitoring alarms and control of the various components across the energy supply network – that enables utility companies to detect changes in usage and supply. There are also specific data processing and analytics requirements related to particular elements of the grid. These include metering to ensure accurate bills and customer satisfaction and engagement, as well as remote equipment diagnostics and performance monitoring that enable predictive maintenance and reduce outages.

Customer-focused energy companies are also able to combine these two elements to generate new revenue streams through subscription offerings that combine predictive servicing of equipment with smart metering that can provide customers with information to help them manage energy usage for improved energy efficiency. This in turn can also help the energy companies predict demand and respond accordingly via demand optimization.

Given the volumes of data, and the geographic range across which they are being generated, the energy and utilities sector is also a prime focus for edge computing. Examples include monitoring industrial machinery in remote areas such as mines or oil and gas fields where network access is intermittent or is limited by cellular or satellite bandwidth, while worker optimization and safety at these remote sites is also a key use case. Additionally, while the volumes of data involved with oil and gas exploration might seem ill-suited to the edge, localized edge processing of real-time data can complement predictive models created centrally on larger data volumes for faster decision-making about when or where to drill.

An unusual example of edge data processing in the energy sector is vegetation management: the need to stop vegetation growth impacting the delivery of electricity. This might sound trivial, but vegetation removal is the largest line item in operations and maintenance costs for the electrical grid. California Independent System Operator spends more than $250m per year on vegetation management on high-voltage distribution lines alone.

Traditionally, vegetation management was handled via manually intensive inspection and removal, but it is increasingly being automated via the use of drones as well as edge data processing and streaming analytics (on the drone itself) to identify types of vegetation and process predictive models of growth patterns to identify risk, maintain electricity supply, and enable more efficient decision-making to reduce operations and maintenance costs, enabling funds to be repurposed for digital transformation.

Retail

Improved customer experience is the number one benefit of data and analytics projects, according to the results of 451 Research's Voice of the Enterprise, Data and Analytics 2H 2018. As such, it is no surprise to see that customer behavior data is also the most popular data source used to generate business intelligence insight.

Alongside other existing business data sources (such as sales data and marketing data) enterprises are increasingly taking advantage of data that is generated by the various technologies that make up the Internet of Things, including device data (from smartphones, wearables and other connected devices); equipment/facilities data; video data, environmental data; and location data.

Many of these data sources are being processed at the edge to improve the speed of business decision-making. The two most prominent use cases, according to 451 Research data[5], involve ensuring that retailers can meet and drive demand. Specifically, demand-driven data warehousing and fulfilment enables retailers to reduce delivery times, improve customer satisfaction, and lower storage and logistics costs, while customer flow management and merchandising enables retailers to enhance customer service through improved responsiveness.

Another key use case includes fraud detection and reduction, which increasingly relies on real-time processing of machine learning models at the point of payment to ensure fraud is detected before it is committed. To achieve this, the processing needs to be carried out at or near the payment location, even if the model is created and trained centrally. Other key retail use cases that benefit from edge analytics include customer footfall and traffic analysis, which can be used to make rapid product positioning decisions, as well as smart digital signage, and just-in-time promotions and couponing.

5. 451 Research's Voice of the Enterprise: Internet of Things, Workloads and Key Projects 2018

Recommendations

- Enterprises in all industry sectors, but particularly manufacturing, transportation, energy and utilities, and retail, should investigate the potential business benefits of edge data processing, including faster time to insight, reduced bandwidth costs, and improved security.
- Given the rate of data generation and growth from IoT projects, enterprises should be wary of thinking about edge data processing as the exception but should see it as a necessary step in filtering and preprocessing data for centralized storage and analysis and competitive advantage.
- Enterprises should ensure they understand the latency and capacity characteristics of their various edge devices in order to understand their suitability when considering when and where to execute data processing and analytics workloads.
- Enterprises should investigate the use of stream processing to deliver ultra-low-latency data processing of events or streams of events at the edge in real-time, as well as the transmission of the resulting data to centralized infrastructure.
- Enterprises should be wary of locking themselves in to specific data processing products and services and explore the use of platform-independent cloud-native architecture from edge to cloud that enables workload portability as data volumes grow and data processing needs evolve.

Working Together to Solve Edge-Level Business Challenges

Combining Dell Technologies' infrastructure strength in the data center and out to the edge, Intel's latest generation of processors, and world-class SAS analytics capabilities isn't just a great idea in theory, it's one we've put into practice for customers.

Here are some examples:

- At the Edge: SAS Event Stream Processing (ESP) has been validated on the Intel powered Dell 3000 and 5000 series IoT gateways, enabling powerful edge analytics at the moment data is being generated
- In the Data Center or Cloud: The Dell EMC Solution for SAS Analytics uses the power of Intel Xeon technology and can leverage your SAS 9.X and SAS Grid Computing investments while providing a simple, yet performant, scale out and scale up building block roadmap for deploying and integrating SAS® Viya®. The Dell EMC Solution for SAS Analytics also enables SAS Viya in containers for data scientists and data analysts looking to streamline and accelerate their DevOps life cycle with open source tools, while delivering a self-service platform.

Learn more at:

sas.com/iotsolutions

delltechnologies.com/iot

intel.com/iot

CONTENT PROVIDED BY

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.

© 2019 451 Research, LLC and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your Service Agreement with 451 Research and/or its Affiliates. The information contained herein has been obtained from sources believed to be reliable. 451 Research disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although 451 Research may discuss legal issues related to the information technology business, 451 Research does not provide legal advice or services and their research should not be construed or used as such.

451 Research shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.

NEW YORK
1411 Broadway
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) 203 929 5700

BOSTON
75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

451 Research®

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. 110405 _G102731.0419