Dell Technologies intel

USING DISTRIBUTED ANALYTICS TO MAKE THE INTERNET OF THINGS SMARTER

Machine Learning and Artificial Intelligence drive transformative value from the flood of data generated by IoT devices

The digital economy has arrived, and businesses are challenged with how to transform to remain competitive. Successful businesses thrive on the ability to convert data into insights. In fact, the capacity to continuously analyze data and deploy insights for improving customer experience or driving new operational efficiencies is now a core business process for delivering competitive advantage.

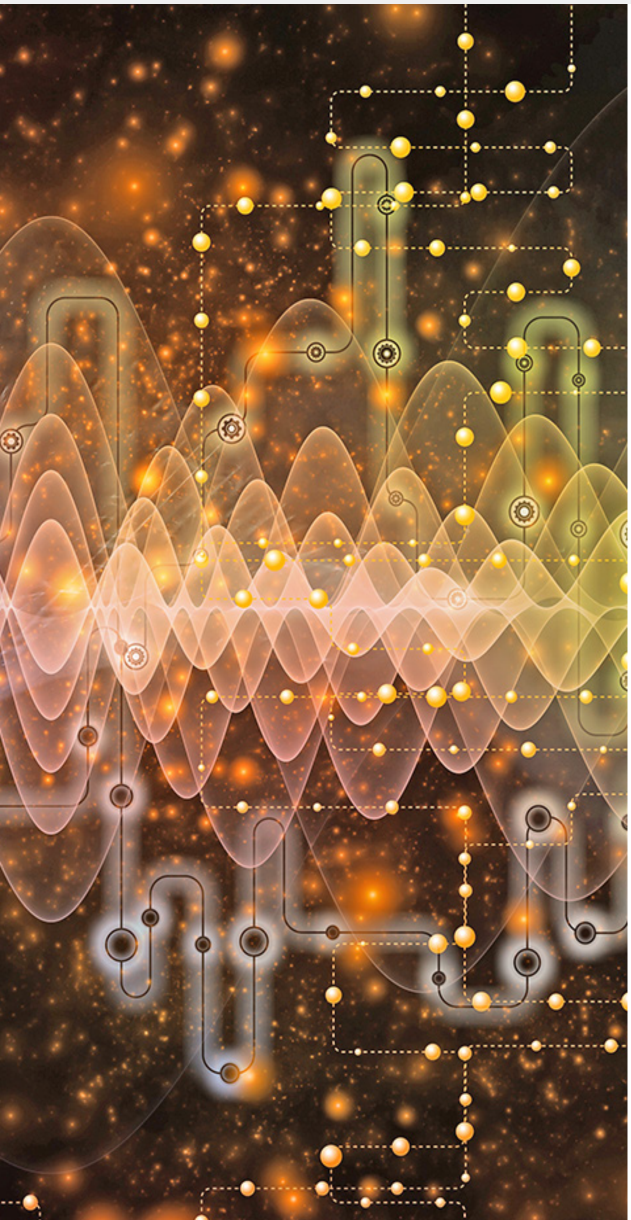

As the Internet of Things (IoT) continues to take shape and brings more and more data sources online, one of the biggest challenges is to understand and act on—both in real time and over the long term—the data generated by anything that can be connected. To do this will require layers of intelligence across the IoT continuum, from the edge, where things first get connected, through to the core, where core data center technology enables real-time machine learning, and to the cloud, where the most sophisticated data centers perform deep learning. In fact, the true potential of IoT can only be unlocked when it comes together with artificial intelligence (AI) and is seen as one complete, interdependent distributed ecosystem. This ecosystem can be thought of as IQT, which stands for the IQ of Things.

As devices at the edge become smarter, organizations can leverage the rising "IQ of Things" to transform businesses and industry. Driving this evolutionary step forward in computing is the relentless power of Moore's law, which is pushing the cost of silicone to near-zero, and enabling compute everywhere, coupled with advances in applying data analytics and machine learning to data at the required speed, scale, and cost to achieve real business value.

Making the data-driven enterprise work smarter

Generating and collecting significant amounts of data creates an extensive pool of information that organizations can leverage to derive business insights for key strategic decisions, recognize patterns from across vast fields of IoT devices, and develop algorithms for controlling devices at the edge. For less time-sensitive information, many organizations are turning to cloud-based analytics and the benefits of deep learning enabled by layers of neural networks using large and complex datasets.

IDG INTERNATIONAL DATA GROUP
WHITE PAPER

Yet key decisions for today's heavily automated systems that control everything from autonomous vehicles to complicated manufacturing processes can't afford the latency while data is sent to the cloud for analysis and back to the site for execution. Instead, there is a growing need for analytical capabilities at the edge as well as in the "core" layer that can more rapidly interpret and act on data while it matters.

The true potential of IoT can only be unlocked with the benefit of a complete, interdependent ecosystem, in which connected things, artificial intelligence and machine learning all come together to make things smarter. This is about using data and the insights from analytics to build on each layer of a solution to generate a virtuous circle of ever increasing intelligence along the IoT continuum.

Organizations with the ability to distribute analytics across the edge, the core, and the cloud stand to gain the most from IoT. Distributed intelligence is part of a larger trend, with the swing of the pendulum back to decentralized computing; just when we're getting used to the cloud, we're seeing another tectonic shift! Driving this new dynamic is the need for computing closer to the edge to provide real time response while alleviating upstream networks from the onslaught of billions of connected things.

But edge isn't the answer for every IoT challenge. Some solutions require a progression from edge devices to central computing capable of more sophisticated learning and understanding of crucial patterns, and then creating and updating algorithms that will control the devices. Consider, for instance, that today's automotive giants produce far more than automobiles and trucks. Those vehicles are true IoT devices and highly-sophisticated data collectors. Business intelligence and data science professionals constantly leverage the resulting information to improve end-product performance, safety, and the customer experience, such as determining the extent that weather, wear, or other factors are impacting tire performance.

Success here requires looking at data beyond a single device, or vehicle, to recognize patterns across a whole population. It's only after studying the whole vehicle population that these automakers are able to create new algorithms that can be transmitted to individual vehicles, so that the embedded analytics engine can make real-time decisions based on, for example, the information generated by sensors in the tire, suspension and braking systems correlated with weather conditions.

Consider a collision avoidance system where edge analytics performed within the vehicle tell us about the operating environment or enable us to assess the health of the system. The edge handles the intelligence layer that must take immediate action with 100% reliability and predictability, such as the rapid autonomous braking response to stop the car. It is where split second decisions are made that require analyzing and correlating many sources of data. The Cloud, however, is the knowledge center where the manufacturer analyzes events from all the cars, conditions, and diverse data sets, and then uses deep learning to train the cars to perform faster, smarter and safer.

There will always be a practical limit for edge compute based on cost, space and power constraints. However, there will also always be a need for even more powerful compute close to the edge for reasons such as responsiveness, security, and governance. On a manufacturing line specialized Programmable Logic Controls (PLCs) are used at the very edge to control things like robots and actuators in deterministic real-time, a parallel to the edge computing for immediate response within an autonomous car. Edge gateways can be used to aggregate data from these PLCs and other field devices, and depending on capability perform a range of functions from filtering to basic machine learning. In turn each of these gateways can feed steams of data to on-premise server-class processing in the core for more advanced machine learning. While this core compute cannot be relied on to respond as immediately as a PLC or even proximal edge gateway, it still enables more rapid and reliable response than the cloud because it is not dependent on the potential bandwidth, reliability and latency issues of a wide-area network connection.

In some cases output from edge and core compute will be streamed to the manufacturer's cloud to be combined with data from multiple sites for longer term optimization through deep learning, or even to its parts suppliers for near real-time visibility into quality

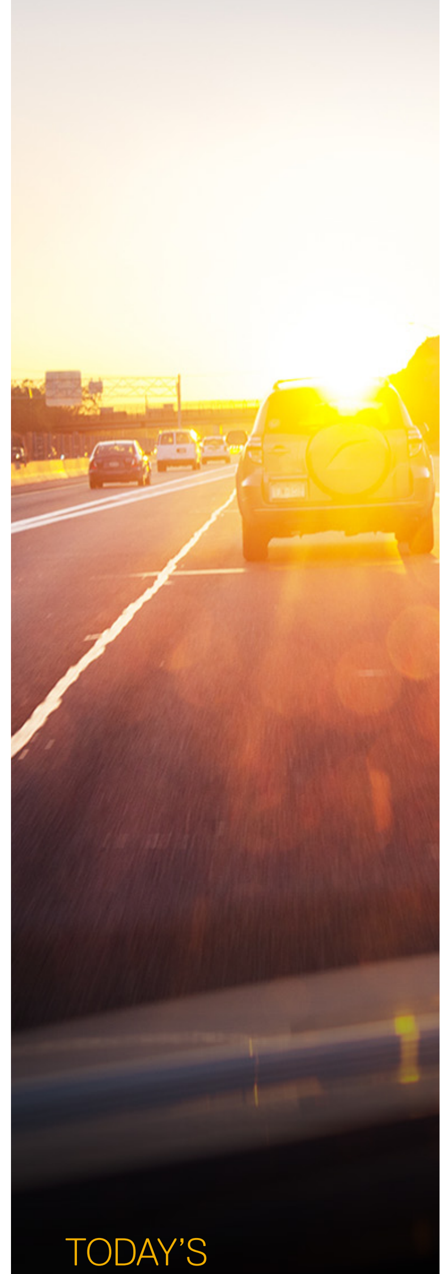

TODAY'S VEHICLES ARE TRUE IoT DEVICES AND HIGHLY SOPHISTICATED DATA COLLECTORS.

BI AND DATA SCIENCE PROFESSIONALS LEVERAGE THE INFORMATION THEY COLLECT TO IMPROVE END-PRODUCT PERFORMANCE, SAFETY, AND THE CUSTOMER EXPERIENCE.

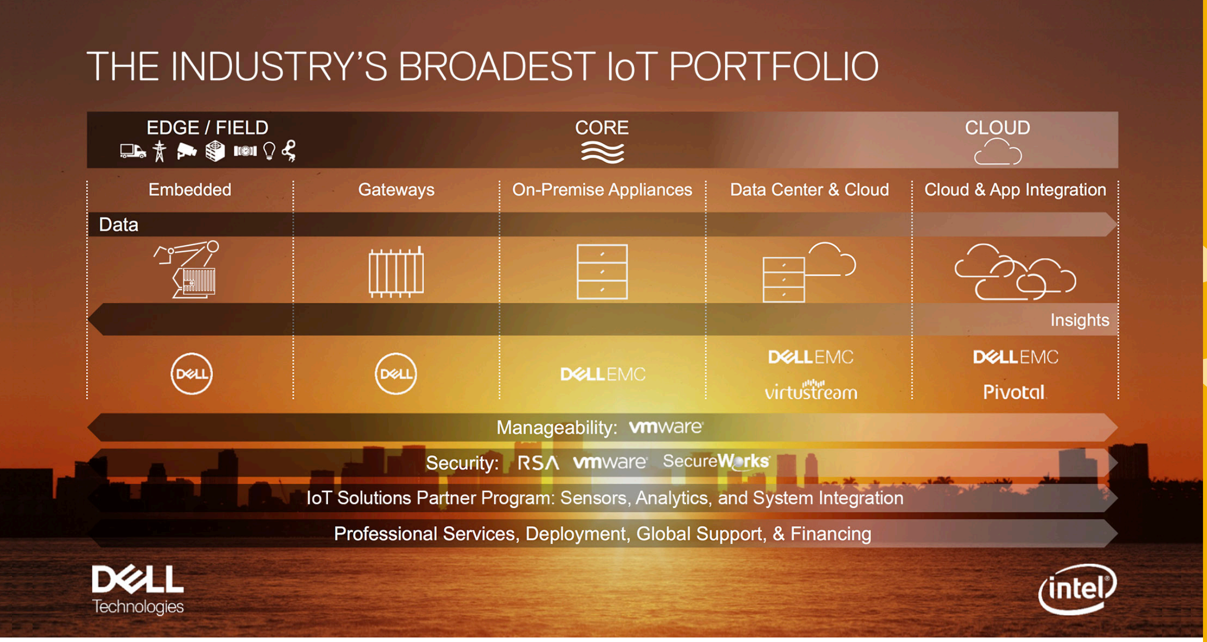

issues. In other cases, a manufacturer may only leverage the edge and core because they are not comfortable connecting to the cloud due to security, privacy and governance considerations. This comfortability depends on risk profile and company culture—a plant processing hazardous chemicals would likely be nervous connecting to the cloud whereas a manufacturer of clothing may not. Another benefit of utilizing core compute is to further reduce the cost for backhauling data to a cloud. For example, on a remote oil rig at sea PLCs would be used for controlling equipment at the very edge, with gateways and servers on the rig providing as much tiered data processing as possible to optimize operations and predict equipment failures prior to sending notifications of any issues over a satellite connection to central command.

Understanding the Layers of Intelligence

Simply having the many aspects of an IoT system in place is only part of the puzzle—the various components must be pieced together so the operational and information technologies are seamless. In coming years, machine intelligence will grow due to sustained deep learning in the cloud. And as compute efficiencies continue to increase machine leaning will continue to migrate to the edge, including being performed by the things themselves. Understanding the steps in creating and collecting meaningful data is crucial to realizing the value of IoT in a connected world.

The Edge. The edge is where things first get connected: for instance, robots in factories, self-driving vehicles, or surveillance cameras in a city. These things operate in the physical world of operations technology (OT) and often in harsh or dangerous environments. Sensors generate data about their operating environment and send it to compute systems—either embedded locally, to nearby on-premises IT infrastructure or to remote locations in the cloud—for processing and analyzing.

With deployments of hundreds or thousands of things, each outfitted with multiple sensors, rugged edge gateways alleviate strain on the network and perform essential, computing tasks. Examples include aggregating, filtering and running incoming data streams through simple rules engines or inference algorithms; grooming data to remove obvious errors; compressing repetitions; and packaging and sending datasets onward for further processing. But the value of data at the edge is typically time-sensitive, or perishable, such as data that alerts an operator to shut down a machine before it malfunctions and causes harm or damage.

The Core. In the core, compute and storage resources don't have the same cost, space and power constraints of edge devices and gateways. Thus, a higher order of intelligence and more complex decision making is possible, such as real-time quality control that takes into account machine, environmental, and quality data.

The core extends enterprise compute capabilities outward from the data center to increase analytical capabilities at intermediate points closer to the edge. Each layer in the core is an aggregation point where additional data sources can be integrated to enable greater insights with lower latency.

As the number of sensors grows exponentially and the sheer volume of data becomes too great to transport to the cloud, latency

POLE POSITION

As car racing has cruised to popularity, the level of competition and sophistication has intensified. JTG Daugherty Racing has garnered a great pole position by leveraging a collection of real-time data from sensors it places on its race cars, empowering crew-members to quickly analyze the data to optimize vehicle performance.

For instance, data downloaded from a car's electronic control unit provides real-time information on brake and throttle pressure, exhaust temperatures and speed after each lap of a race. A very small change in suspension can make a car handle much better.

Crew members in the pit can quickly identify whether they should make any changes and how each chassis modification impacts performance. The team is also saving money, because dynamic data-fueled 3D models of the car (also known as a digital twin) can more definitely pinpoint when tires and other crucial car parts are nearing end-of-life.

issues, as well as bandwidth costs, threaten to become insurmountable. The core lessens the amount of data that needs to be transmitted through wide area networks and to the cloud and has scalable intelligence to generate valuable insights and instruct machines to take action in real time.

The Cloud. The cloud is where the most sophisticated high-volume data storage, processing, and analytics occurs for realizing longer-term (e.g. minutes, days, weeks, months, quarters, years) business insights. It's also where organizations collect disparate data from multiple data sources, including public clouds, as well as information interpreted in historical context, leveraging data previously collected and analyzed over time.

The cloud is the most logical place for large-scale processing and analytics, utilizing deep learning to glean valuable and actionable insights that drive analytical models at the edge and augment machine learning in the core. As the cloud evolves, it is ushering in increasingly flexible platforms capable of changing the economics of machine intelligence, while lowering the cost of artificial intelligence. The result: Faster adoption and more practical implementations of AI, and machines that excel in both learning and inference to ultimately "train" the edge, as coaches train their athletes, to perform and execute better at each opportunity.

Deliberate Urgency

Staying competitive requires the ability to process real-time information and make data-based decisions that go well beyond the predictive, to the prescriptive. Prescriptive analytics not only predicts what will happen next, but also recommends next steps. The analytics involves combining historical data with an array of other available data sources, such as social media or highly localized weather data, to generate deeper insights.

Building an environment of layered intelligence by increasingly making the "things" smarter is an evolution. This only happens when businesses empower their organization to use IoT tools that provide the ability to optimize the distribution of analytics workloads across the edge the core and the cloud to securely and efficiently generate meaningful and measurable results.

Just as no single architecture exists for IoT, no single company alone can own every important aspect of IoT. Realizing the potential requires a vast, cooperative ecosystem capable of constant evolution. The seven strategically-aligned businesses that comprise Dell Technologies bring together many of the ingredients for building distributed IoT solutions for organizations around the world. These businesses provide the industry's broadest portfolio and when combined with curated partnerships for sensors, analytics and system integration, Dell Technologies can serve as a one-stop shop to deliver value across many industry verticals and use cases.

More specifically, Dell Technologies has built its IoT strategy on four key pillars to help customers make things smarter. This strategy begins with the industry's broadest IoT infrastructure portfolio to provide flexibility to deploy secure compute, analytics, and storage where you need it in the edge to cloud continuum based on your specific use case. Secondly, leveraging its expertise in both IT and OT, the different layers of the infrastructure portfolio are optimized for the needs of both parties. The third pillar is Dell Technologies' industry-leading partner program, which curates the best of breed for all layers of the IoT stack so you don't have to sort through the thousands of vendors. In our collaboration with Intel, we utilize the latest Intel technologies across our IoT portfolio to create a strong foundation and bring top performing solutions to market. Finally, the fourth pillar is bringing together proven, industry-specific engineered Ready Solutions and Blueprints to help achieve ROI faster.

These solutions harness the value of data analytics and the Internet of Things across the edge, the core, and the cloud to guide your digital transformation strategies for making things smarter, assisted by the power of Intel.

WE INVITE YOU TO LEARN MORE AT: DELLTECHNOLOGIES.COM/IOT

MAKING BUILDINGS SMARTER

Intel created an IoT enabled 10-story, 630K sq. ft. smart building, outfitted with around 9,000 sensors used to track and optimize temperature, lighting, energy consumption, and occupancy. The sensors, 70 percent in the ceiling, provide 24/7 real-time data. Analytics is run on the data gathered from sensors to generate actionable insights. Intel added advanced building analytics via an integrated energy management system, called iBEMS*, from L&T Technology Services. The 40 use cases include around 55 percent focused on energy conservation and operational efficiency, and the rest on employee satisfaction. The results forecast energy/water savings of $645K/year, payback in less than four years, and 30% capacity increase. Socially driven temperature control can increase worker satisfaction with workplace thermal comfort by 83 percent.