

IoT Analytics at the Edge

Vertica Brings Analytics to Where the “Things” Are

As the flow of data generated from the Internet of Things (IoT) swells in size, organizations increasingly recognize the advantages of leaving certain data right where it is captured—at the edge—and analyzing it there. This shift away from a centralized cloud model enables organizations to avoid the downsides of transmitting large data sets to distant data centers while gaining greater control over connected things.

Table of Contents

page

IoT Is Driving New Business Opportunities.....	1
Why Push Analytics to the Edge?	2
Why Shift Compute to the Edge? Consider These Seven Reasons.....	3
Vertica Advanced Analytics Platform at the Edge	4
Closing the Loop on IoT Analytics.....	5
Use Cases for Analytics at the Edge	5
Key Takeaways.....	6

Are You Ready for This?

The projected growth in the use of IoT devices and applications is enormous. Gartner estimates endpoints of IoT will grow at a 31.7% compound annual growth rate from 2013 through 2020, reaching an installed base of 20.8 billion units.²

IoT Is Driving New Business Opportunities

For organizations that want to move certain SQL analytics functions to the edge, Vertica is an ideal partner. Vertica enables the shift to analytics at the edge by bringing the power of the [Vertica Advanced Analytics Platform](#) directly to the Internet of Things. These solutions help organizations gain greater value from the data they generate remotely by delivering analytics and machine learning at the edge—where the “things” are.

For today's enterprises, the proliferation of connected devices and the rapid rise of the Internet of Things create unique opportunities to grow the business. IoT opens the door to new ways to engage customers, new business models, new insights into existing processes, and new ways to monitor and maintain machinery and equipment.

“Why are so many companies moving ahead with IoT in light of the complexity that it represents?” asks Enterprise Strategy Group (ESG). “They see immediate business value in the form of operational efficiencies (45% of respondents), better and differentiated customer service (39%), creation of new products and services (38%), and development of new business models (26%).”¹

Viewed from a broader level, IoT also opens the door to market disruption. Organizations that are poised to capitalize on massive amounts of data from connected devices have the opportunity to change the ground rules for their current markets and to segue quickly into new markets.

Those who succeed in this new era are likely to be the organizations that understand the potential of IoT and put the correct systems in place to capture, process, and analyze terabytes or even petabytes of data from connected things. For these digital disruptors, IoT will equate to real-time control, faster decision-making, and a competitive advantage.

While there are many unknowns about IoT and how it will impact current business and operating models, this much is clear: Now is the time to lay the groundwork for your organization to capitalize on the growing volume of machine and sensor data.

¹ *ESG Solutions Showcase, “HPE and IoT Compute at the Edge,”* March 2016, citing survey findings from the ESG Research Report “*2Q16 IT Spending Intentions Survey*,” February 2016.

² *Forecast Alert: Internet of Things—Endpoints and Associated Services, Worldwide, 2015,* Gartner, 2015.



With conventional approaches, massive amounts of data captured at remote locations must be transmitted to a distant corporate or cloud data center for processing and analysis. When you're talking about terabytes of data, this data transfer process can be a slow, risky, costly, and inefficient undertaking.

Why Push Analytics to the Edge?

Data from the Internet of Things originates remotely. Depending on the industry, the edge might be on a plant floor, in a field, on an oil rig, or in a warehouse. The data from connected things at these remote sites has the potential to generate valuable business, engineering, and scientific insights.

Oil and gas, energy, manufacturing, telecommunications, and consumer tech are among the industries that have big IoT opportunities. If they have the correct systems in place, they can now harness data from diverse environments—from windmills in the desert to smart energy grids, from pipelines to manufacturing plants—to enable predictive maintenance and timely process optimization.

With conventional approaches, massive amounts of data captured at remote locations must be transmitted to a distant corporate or cloud data center for processing and analysis. When you're talking about terabytes of data, this data transfer process can be a slow, risky, costly, and inefficient undertaking. What's more, in many industrial settings, latencies in data transmission and analysis can have serious consequences in the form of failed equipment and operational downtime.

In light of these IoT challenges, organizations increasingly recognize the advantages of leaving certain data right where it is captured—at the edge—and analyzing it there. The result is the ability to take instant action and affect immediate control over the connected things.

Data from the Internet of Things originates remotely. Depending on the industry, the edge might be on a plant floor, in a field, on an oil rig, or in a warehouse.

Vertica is enabling this shift to analytics at the edge using the same core Vertica Analytics Platform that is currently deployed on premises, in the clouds, or on Apache Hadoop. This helps organizations harness the power of the data they generate remotely by delivering advanced analytics and in-database machine learning at both the point of data collection, and at data aggregation points in the clouds or on-premises.

Shifting analytics to the left in the end-to-end IoT solution

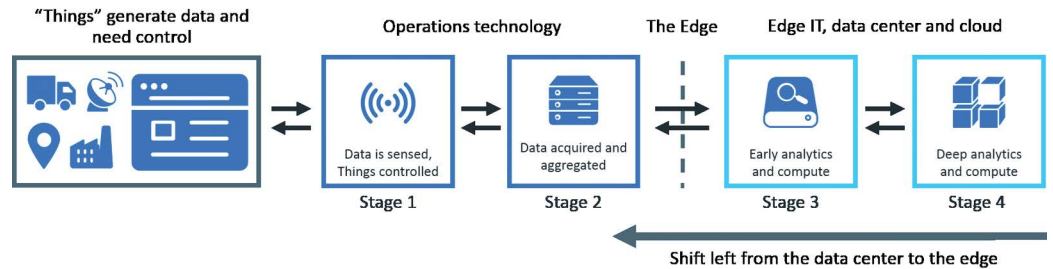


Figure 1. Moving analytics to the edge can be characterized as a shift to the left, or to an earlier stage of the end-to-end IoT solution. This shift moves certain processing and analytics functions out of the data center and closer to the point of data capture.

Why Shift Compute to the Edge? Consider These Seven Reasons

Data from the Internet of Things originates remotely. Depending on the industry, the edge might be on a plant floor, in a field, on an oil rig, or in a warehouse. The data from connected things at these remote sites has the potential to generate valuable business, engineering, and scientific insights. Although many IoT use cases depend on performing advanced analytics on large, aggregated datasets, in some cases, shifting analytics to the edge can create an advantage.

- 1. Latency:** What good is an Internet-connected car if there's a lag between when a pedestrian appears in front of the car and when the system actually tells the car to stop? For some mission-critical functions, this sort of latency is intolerable, so compute on the edge is a must. Latency in data transfer reduces time to insight, which slows time to action for business and responses to the data.
- 2. Bandwidth:** Sending data from edge devices to the cloud or a data center can use a tremendous amount of bandwidth. Using available but limited network bandwidth prevents other business-critical uses of that network bandwidth. Many companies simply can't handle the bandwidth needs of IoT without compromising the performance of mission-critical systems.
- 3. Cost:** Sending large amounts of data from remote sensors to a distant data center over wireless networks can be costly. Extra wireless bandwidth and extra security precautions will inevitably add to the costs. Processing data at the edge reduces the network-related costs.

4. **Threats:** Transferring data over long distances and remote networks inevitably exposes data to threats of attacks and other security breaches. These risks strengthen the case for securing and analyzing data at the edge.
5. **Duplication:** If you are going to collect remote data and send it to a distant data center there will inevitably be some duplication. In addition, the complexity and cost of storage and other assets must be duplicated to accommodate the data that is sent to a corporate or cloud data center.
6. **Corruption:** Even without any nefarious activity from hackers, data can be corrupted on its own. Data transmission, especially large amounts across long distances, can incur drops and delays associated with correction and recovery, while missed connections can plague edge-to-data-center communications.
7. **Compliance:** Regional and country compliance regulations can restrict or complicate data transfer across borders and over long distances. Companies that embrace IoT solutions often run up against such compliance issues, which are mitigated with edge analytics.

Vertica Analytics Platform has a wide range of data processing functions that enable analytics at the edge.


Vertica Advanced Analytics Platform at the Edge

Vertica Analytics Platform runs at the edge on the HPE Edgeline EL4000 and other systems to deliver historical and predictive analytic insights from in-database machine-learning functions across a broad range of IoT analytics use cases. It enables organizations to manage and analyze massive volumes of structured and semi-structured data quickly and reliably.

Vertica Analytics Platform has a wide range of data processing functions that enable analytics at the edge. For example, the platform offers:

- Blazing fast query performance
- Integration with Apache Kafka and Spark
- Geospatial, time-series and log text analytics
- In-database machine learning functions for regression, clustering and more
- Edge, cloud, Hadoop, and on-premises deployment

With integrated machine-learning capabilities and a closed-loop approach—spanning from the data center to the edge—Vertica Analytics Platform enables organizations to train predictive analytics models on large sets of data in the data center and then run those models at the edge.



Organizations can use Vertica on-premises, on Hadoop, or in the cloud to build and train predictive analytics models and apply machine-learning algorithms to incredibly large data sets without any unnecessary down-sampling.

Closing the Loop on IoT Analytics

When analytics shift to the left, they shouldn't exist in a world apart from that of traditional analytics environments. They should be part of a closed-loop system that also encompasses analytics run on-premises, on Hadoop, or in a cloud data center. This is the Vertica approach to advanced analytics.

Organizations can use Vertica on-premises, on Hadoop, or in the cloud to build and train predictive analytics models and apply machine-learning algorithms to incredibly large data sets without any unnecessary down-sampling. They can then deploy the models to the edge to enable targeted use cases, such as identifying anomalies, predicting maintenance issues, and analyzing product usage. Visualization of SQL analytics results can be done from either the edge or the data center.

With this closed-loop approach, organizations can gain the benefits of edge analytics as well as the benefits of more historical data analytics for typical IoT use cases that require larger volumes of data over longer time periods—such as warranty management and service history.



Use Cases for Analytics at the Edge

There are countless potential use cases for advanced analytics at the edge, including those that follow. Some of these are conventional IoT use cases that might benefit from the avoidance of the costs, risks, and latencies associated with transmitting large amounts of data over long-distance wireless networks.

Predictive Maintenance and Asset Management

In many industrial settings, edge analytics enables real-time monitoring, anomaly detection and alerts, failure prediction, and predictive servicing of critical equipment. These capabilities help organizations move from reactive to proactive maintenance. This shift to predictive maintenance can help reduce unplanned downtime, increase scheduled maintenance, improve safety, and pave the way for new operating models for industrial sites.

Real-Time Process Optimization

Factory operators capitalize on edge analytics and machine-learning algorithms to predict consumption of raw materials and optimize their supply chains in real time. These capabilities help operators maintain tighter control over costs while ensuring the availability of the materials that go into products.

Fleet Management

Equipment manufacturers integrate sensors and telemetry data analytics into fleets of vehicles to gain insights into driver behaviors and patterns and to help fleet operators improve the efficiency of both their drivers and their overall operations. For example, some companies use sensor data to enable intelligent



Without analytics, the massive amounts of data generated by the Internet of Things won't do anything for private enterprises or public entities.

route optimization to reduce fuel costs and to facilitate preventive maintenance before mechanical issues disrupt fleet operations.

Smart Agriculture

The agriculture industry analyzes data captured by sensors on farming equipment with publicly available data to enable smarter approaches to farming. Among other applications, IoT solutions help operators monitor field soil moisture levels to optimize the use of irrigation water, track soil conditions, and monitor the health of crops.

Building Management

Facilities personnel use IoT monitors to gather data on building systems, including heating, ventilation, and air conditioning (HVAC) equipment, and to enable alarm monitoring and active notifications. They also leverage sensor data to guide proactive maintenance and the optimization of systems to avoid equipment failures, improve energy efficiency, and keep tenants comfortable.

Key Takeaways

Without analytics, the massive amounts of data generated by the Internet of Things won't do anything for private enterprises or public entities. The real value lies in the use of advanced analytics to extract insights from the data. The only question is where to run the data analytics—in the data center or at the point of data capture or both?

Vertica enables the shift to advanced analytics at the edge by bringing the power of advanced analytics and in-database machine learning directly to the Internet of Things.

As the data from connected machinery, equipment, and other things swells in size, it often makes sense to leave certain data right where it is captured and analyze it there. This shift enables organizations to avoid the latencies, costs, and risks associated with transmitting large amounts of data over long distances. Instead, they can take instant action on the insights from the data and affect immediate control over the connected things.

For organizations that want to move certain analytics functions to the edge, Vertica is an ideal partner. Vertica enables the shift to advanced analytics at the edge by bringing the power of advanced analytics and in-database machine learning directly to the Internet of Things.

Learn More At

www.vertica.com/iot

www.vertica.com/HarborIoT

www.vertica.com/machinelearning



Additional contact information and office locations:
www.vertica.com

www.vertica.com