ORACLE CLOUD

Driving Real-Time Insight

The Convergence of Big Data and the Internet of Things

ORACLE WHITE PAPER | JULY 2016

ORACLE

Executive Summary

Big data and the Internet of Things (IoT) are two of the most exciting developments in the business world, but forward-looking organizations need new types of technologies to fully realize their potential. Both are important in their own right. Together they offer even greater potential. IoT is an important thread in a broader conversation that is being articulated with big data analytics. The value is in the data, but much of that value is buried. The challenge comes with finding precious data elements and uncovering unique insights, and then using those learnings to impact enterprise applications and processes. When properly configured, big data and IoT reinforce each other, with the whole becoming greater than the sum of the parts.

The value of a connected enterprise goes beyond its remote systems. It's what those remote machines and sensors connect to that adds value. Every large business depends on a growing set of data sources and entrenched enterprise applications. The key tenets of IoT and big data projects are to ingest the data and react to it in real time, to analyze sensor data alongside existing enterprise data for deeper insights, and to use those insights to enhance and reinforce the processes within your applications.

For example, a courier service could use IoT data to monitor its drivers and vehicles, including routes followed, fuel used, and even the temperature in the cargo bins when sensitive materials are transported. Tracking these data points would allow the company to improve efficiency and performance, as well as to ensure optimum conditions for each shipment. A sporting goods manufacturer could put sensors in its rackets and bats to allow athletes to analyze ball speed, spin, and impact location. A retail chain could use IoT data to track customer locations, tailor unique offers, and boost participation in its loyalty program. A heavy equipment manufacturer could install sensors in its machines to alert technicians when critical processes are out of tolerance, and then analyze the data stream to predict upcoming maintenance.

All of these examples illustrate the potential of IoT in the context of big data analytics. However, organizations can't realize the full value of IoT and big data initiatives until they succeed not only in collecting information, but also integrating it into their existing business processes. In most cases, they need to be able to merge IoT data with other types of enterprise data and incorporate it into the applications and information systems that run the enterprise, from enterprise resource planning (ERP) systems to business intelligence apps. Most organizations implementing IoT projects today are not making that connection and are thus not taking full advantage of the true value of these initiatives.

Oracle offers a scalable, secure, platform that empowers organizations to acquire and manage data, integrate it throughout their businesses, analyze it to discern patterns, and take action to improve critical processes, services, and operations. Oracle's IoT and big data platforms work together and complement each other as part of a cohesive platform as a service (PaaS) strategy. Read on to learn how you can:

- Securely connect and acquire data from any device
- Analyze in-flight data and drive insights from IoT analytics
- Integrate IoT data into your enterprise applications
- Create a cost-effective, scalable, open architecture for big data analytics in the cloud
- Automate lifecycle management to run a wide variety of big data workloads
- Analyze data across Hadoop, NoSQL and Oracle Database Service - Exadata Edition
- Enable in-depth analytics using Oracle Advanced Analytics and Oracle R Advanced Analytics for Hadoop

Defining Terms

Internet of things (IoT) refers to automated networks of computers, devices, and sensors that can process their own data instead of depending solely on people for input. These internet-connected systems gather data and communicate with external processes via onboard sensors, impacting businesses in just about every industry. Gartner predicts that the number of "connected things" will reach 25 billion by 2020.

There are four primary types of IoT projects:

- *Monitoring products* – Embedding sensors, software and other technologies into the offerings that a company brings to market
- *Monitoring customers* – Tracking digital devices that customers carry or wear
- *Monitoring supply chains* – Putting sensors, digital cameras, and other digital devices in production and distribution operations
- *Monitoring premises* – Putting sensors, digital cameras, and other devices in places where companies do business with their customers

Of course, monitoring products, customers, supply chains, and facilities is important, but they are merely the first steps in a complete IoT scenario. By loading a real-time IoT stream into a big data reservoir you can also conduct *after the fact analytics* to pull more value from that data as you gradually move from reactive to proactive operations. This more comprehensive approach goes beyond merely ingesting data to perform analytics on the real time data stream. Once something breaks you can react by fixing it—and you can also study the event data so you can foresee and prevent similar events from happening in the future.

For example, sensors on an assembly line can detect when a machine is out of tolerance, and react by sending an alert or shutting down the machine to minimize unwanted defects. Later, an analyst can examine this data stream to determine what caused the problem. This is after-the-fact analysis. Combining these two types of inquiry helps organizations address critical issues, and paves the way for predictive analytics to circumvent problems in the future.

Why Cloud?

Cloud technology is ideal for IoT due to its inherent flexibility, low relative cost, and innate scalability—all of which are important as IoT projects gain momentum. You can start small and expand quickly, just like many IoT projects do, with minimal upfront costs. The cloud offers elastic scalability and the flexibility to support a wide range of IoT and big data projects. It's easy to get started with cloud initiatives because you can start small and scale up quickly as data volume grows and more resources are needed. When combined with big data technologies, you can analyze massive amounts of IoT data as quickly as it is streamed into the cloud.

In addition, a complete solution should allow you to take action when an event requires a specific response based on predefined rules. Big data and IoT technologies should work together to help you realize the three primary objectives of IoT:

- Generate insights into processes, products, and quality
- Predict failures before they occur to reduce downtime
- Reduce costs by optimizing equipment and maximizing resources

These three objectives are relevant to almost every industry and vertical market segment. Let's look at some examples.

Industry Examples

Manufacturing

Consider a pharmaceutical company that depends on high-precision pumps to control the active ingredients added to pills during the pharmaceutical production process. The company that manufactures these pumps can differentiate itself by offering monitoring services for this equipment to sense and track operational variables. Instead of just selling a pump with a warranty, they can also provide complete operational reports about each pump's behavior.

Whether it is pumps or valves or filters or any other type of precision equipment, this type of third party monitoring allows the manufacturer to recommend maintenance based on actual usage and day-to-day conditions. In essence, the pharmaceutical company not only purchases equipment. They also have a contract with the manufacturer to keep the equipment running optimally. The manufacturer's competitive edge comes not only from

supplying precision equipment, but also from being able to monitor its usage by tracking and analyzing sensor data from an active data stream.

Clearly, most pharmaceutical companies already have a high degree of automation on their factory floors. However, there is an opportunity for companies that manufacture pumps, valves, filters, and other precision equipment to add additional value with IoT analytics. These OEM manufacturers are in the best position to predict failure or degradation of behavior—assuming they can acquire and analyze the IoT data. They can monitor the data stream to assess the performance of systems and equipment and react to problems as they occur. Real time data can be loaded into a big data cloud for additional analysis. Sensors on the equipment can detect when a pump is out of tolerance. After-the-fact analytics can determine what caused a particular problem. It's a huge market that is expected to reach $140 billion by 2020.[1]

A complete platform utilizes this type of two-pronged approach to monitoring and improving conditions on the manufacturing floor. It can monitor a real-time data stream for immediate problems. Furthermore, businesses can periodically analyze the historical data. These secondary analyses enable them to determine whether an isolated problem is a one-time anomaly or indicative of a pattern. Insight derived from IoT data helps these organizations close gaps between operational systems on the factory floor and the enterprise systems in the back office as managers strive to reduce costs, improve quality, and maximize available resources.

An IoT and big data cloud service can bring in the data and evaluate the stream in real time. It can react appropriately based on predefined rules. In addition, it can store and manage all the data for in-depth analysis. Having these two types of analysis working together helps uncover hidden problems and enables the manufacturer to modify the rules that are applied to the IoT stream.

Transportation and Logistics

IoT data helps logistics companies with load and dispatch operations by enabling location-based routing of a fleet of vehicles. Sensors can monitor, capture, analyze, and report updates about each vehicle's whereabouts and operating status. For example, one trucking company is combining sensor data from the vehicles with scheduling data from the ERP system and weather information from a third party service to set the optimal tire pressure for each vehicle. The data can be loaded into a Hadoop database for after-the-fact analysis, helping these companies maximize fuel efficiency, optimize routes, and locate goods in transit.

Utilities

This same type of bi-directional communication and control pertains to many other industries, markets, and use cases as well. Equipment sensors are prevalent in heavy machinery, automobiles, assembly lines, electric grids, computer equipment, and many other domains. For example, utility companies commonly gather usage data from smart meters to help customers cut costs and conserve energy. Monitoring this real-time IoT stream allows the utility to detect potential overloads developing throughout the electric grid. In extreme cases they can selectively switch off high consumption equipment to prevent brownouts. In addition, the historical data can be analyzed to determine the cause of chronic problems, to recommend cost saving measures, and to forecast long-term power needs.

Retail

Retailers gain insight by analyzing consumer transaction histories and web-behavior. Data is commonly gathered from social media feeds, purchase histories, and location information on mobile devices. The goal is to understand each customer's behavior and preferences so the retailer can present personalized offers to them. For example, location data allows marketers to issue targeted offers as customers get close to stores. The real-time IoT data stream triggers real-time offers based on in-depth analytics of each customer's profile. After-the-fact analysis can combine the location information with footfall analysis to determine the optimum location of new stores.

Insurance

Insurance companies gather weather data from third parties to prepare for natural disasters and the consequent influx of claims. These real-time weather feeds can be analyzed to develop risk profiles. Actuaries can also use this historical data to forecast revenue and liabilities. Other uses of IoT data in the insurance industry include the use of sensor devices to share details on vehicles and driving behavior. Gathering this information enables insurance companies to offer usage-based policies that reduce costs and lower rates for good drivers.

[1] *BI Intelligence*, "The Internet of Everything," 2014.

Automotive

Connected vehicles gather reams of data about engine performance, fuel efficiency, and other mechanical conditions to convey insights to occupants, service centers, and other vehicles. Sensors throughout the vehicle monitor fuel levels, tire pressure, and dozens of engine control functions—amassing as much as 25GB per hour of operation. In addition to flashing a warning light on the dashboard when a sensor detects an adverse condition, these sensors can send fault codes directly to the manufacturer for troubleshooting and quality control monitoring. Auto dealers can use this information to schedule service, estimate repairs cost, and order parts. Additional capabilities can potentially be available through embedded applications such as navigational map updates, maintenance reminders, and occupant entertainment options. This use case illustrates an important shift from preventive maintenance to predictive maintenance. Collecting and analyzing this connected-vehicle data also enables new types of pay-per-use models.

Municipal Services

Today's "smart cities" monitor the status of streetlights, parking facilities, and field equipment to improve roads, reduce traffic congestion, and enhance public transportation options. For example, each streetlight has a unique identity on the network and is constantly emitting data. If a streetlight at a busy intersection is acting up it will try to reboot itself, and if that doesn't work it will send an alert so that a repair technician can be dispatched immediately. The City of San Francisco provides real-time information about available parking spaces by sending updates to a parking application on drivers' mobile phones. The city continually analyzes this real-time data stream to adjust the price of parking in relation to demand. Data is collected from parking meters, automatic garage gates, and on-street sensors.

Agriculture

The farming industry uses IoT data to forge tighter relationships between agriculture suppliers and farmers. Sensors are prevalent on farm equipment, shipping containers, and delivery vehicles to communicate valuable information for tracking assets. Large industrial farms can analyze the data to accelerate deliveries, improve forecasting of yields, and optimize a huge variety of production processes.

The Oracle Offering: Ready for the Enterprise

Oracle is developing integrated PaaS solutions that empower organizations to acquire and manage IoT data streams, extend data throughout their businesses, analyze data to discern patterns, and take action to improve critical processes, services, and operations. This mature solution is anchored by three primary cloud services:

- *Oracle Internet of Things Cloud Service* lets you build and deploy applications that can capture and analyze IoT data. It can accommodate data from many types of connected devices and react to events in a real time data stream. Oracle IoT Cloud Service can securely connect any device to bi-directionally transact data. It can also perform real-time analytics to enrich that data. Oracle simplifies connectivity using open interfaces and integrations to ensure that the right data is available for the right application at the right time. Oracle IoT Cloud Service can easily integrate with your existing enterprise applications. For example, when a fault is detected, you can automatically issue a work order.

- *Oracle Big Data Cloud Service* is an automated, high-powered cloud environment that has been designed to advance existing analytical capabilities. With automated lifecycle management and one-click security, Big Data Cloud Service can securely run a wide variety of big data workloads. When used with *Oracle Big Data SQL Cloud Service*, you can analyze data across *Hadoop*, *NoSQL* and *Oracle Database Service - Exadata Edition*, while leveraging your existing SQL skills. In an IoT setting, Oracle Big Data Cloud Service combines data from multiple sources to enable in-depth analytics using Oracle Advanced Analytics and Oracle R Advanced Analytics for Hadoop, which are used to build the predictive models.

- *Oracle Big Data Discovery Cloud Service,* running on Big Data Cloud Service, lets you securely run a wide range of big data workloads. When you work with these new data sets the first step is often to figure out the realm of possibilities for exploration and discovery. Oracle Big Data Discovery Cloud Service

helps you determine what questions you can potentially answer from your data. This is essential information for the data scientists who build the predictive models.

So, to summarize: You can store your data in Oracle Big Data Cloud Service. You analyze it and discover patterns with Big Data Discovery Cloud Service. You build predictive models using R. And you can communicate those results using Oracle Big Data SQL Cloud Service. Oracle's comprehensive big data solution includes Oracle Data Integration, Oracle R, and Oracle Spatial and Graph software to analyze the data in these reservoirs.

With this extensive family of integrated solutions, you can explore the data to see what is there. You can build and test predictive models. And you run those models in the real world to predict significant events.

Impacting Enterprise Processes

Every business in the world runs on enterprise applications. The key is not just to analyze data but also to use the insights to modify business process as directed by those applications. Oracle's solution is like a "learning machine" that can detect problems, track events, and make corrections. Creating this type of real-time IoT feedback loop allows you to glean critical insights as you interactively develop, refine, and enhance your products and services.

Oracle IoT Cloud Service and Oracle Big Data Cloud Service are key elements of Oracle's broad cloud strategy. They play critical roles by collecting data, connecting it with devices, and conducting real-time analyses on that data to figure out where the problems are occurring—and send updates to enterprise applications. Oracle Big Data Cloud Service delivers Hadoop as a secure, automated, elastic service that can be easily integrated with enterprise data in Oracle Database. It also extends the mature security capabilities of Oracle Database to Hadoop and NoSQL. Oracle's extensive middleware family simplifies the connections and integrations.

Oracle leverages the cloud to ramp up big data projects quickly and help you get the most out of your big data investments. Rather than procuring and maintaining database infrastructure on premises you can pay a flat monthly fee for the cloud service. This approach enables you to enjoy the power and reach of big data without retraining employees or overhauling applications. Because the service is secure and scalable, you can access and analyze data when you need it. Analytical results can be seamlessly integrated with critical business processes using familiar SQL skills.

Oracle's comprehensive solution supports bi-directional communication. You can acquire data from sensors as you monitor and control IoT devices, and also send instructions to those devices—such as to reset them or switch them off—securely and reliably. By evaluating the IoT stream, Oracle IoT Cloud Service can determine if an event requires a specific response or action based on predefined rules. Furthermore, you can integrate IoT data with data from other sources to discern patterns, discover problems, and make changes to the surrounding infrastructure.

In today's growing IoT deployments, all devices must work together and be integrated seamlessly with all other devices, as well as be connected to the surrounding systems and infrastructure. To facilitate these connections, the Oracle big data and IoT solutions are backed by a comprehensive application portfolio that can help you automate everything from financial analysis to transportation management. Oracle Fusion Middleware provides security, integration, and business process management. Oracle's mature software infrastructure makes it easy to integrate data from Oracle applications, third party applications, and IoT streams, as well as from other independent data sources.

Of course, setting up the basic infrastructure is just the first step. Once you have the devices in place and the sensors activated you will start collecting IoT data—lots of it. Unless you can store it easily and integrate it with your existing applications and business processes, you won't realize any value. Oracle enables you to integrate IoT data with all types of applications, enterprise databases, third-party sources, and data feeds. All of this data can be brought together for discovery and statistical analysis. You can analyze the real-time stream and analyze big data sources. Combining these two methods helps you address critical issues.

Oracle has spent decades helping customers capture data, extract value from it, and make sure that the right data gets to the right application. In the world of big data and IoT, Oracle gets the right info to the right apps at the right time.

Conclusion: Transform Your Business with Big Data and IoT

IT leaders and business leaders in nearly every industry are beginning to roll out IoT and big data initiatives. Oracle enables you to acquire IoT data, integrate it throughout your business, analyze it to detect patterns, and use the insight to improve important processes, services, and operations. Big data analytics help you take the right action, whether it involves sending data back to a device, requesting a field service engineer, or updating a particular module within the ERP system. The Oracle solution can help you:

- Reliably and securely collect data from all types of devices
- Standardize the integration of devices within the enterprise
- Perform real-time, big data, and predictive analytics on IoT streams and events
- Seamlessly extend enterprise applications and processes with IoT data
- Allow enterprise and mobile applications to control devices for bi-directional communication and control

Oracle IoT Cloud Service and Oracle Big Data Cloud Service are ready to go to work in your organization to securely connect, analyze, and integrate real-time IoT data, at scale, as well as to transmit data between connected devices and enterprise applications. Let Oracle help you manage the complexities of the Internet of Things, from devices to applications, big data to analytics, security to scalability, all in the cloud.

For more information, please visit the following sites:

- oracle.com/iot
- cloud.oracle.com/iot
- oracle.com/bigdata
- cloud.oracle.com/bigdata

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

ORACLE

CONNECT WITH US


blogs.oracle.com/bigdata


facebook.com/OracleBigData


twitter.com/OracleBigData


oracle.com/bigdata


linkedin.com/company/oracle-big-data


youtube.com/user/OracleBigData

Hardware and Software, Engineered to Work Together