

BIG DATA EDITION

Unleashing IT

VOLUME 5 / ISSUE 3

BIGGER DATA

In expanding its data resources, Australia's Quantum is realizing the true potential of consumer analytics. Page 8

PAGE 5

IT in reverse?

PAGE 6

Faster data prep

PAGE 10

Treasure trove



Unleashing IT

VOLUME 5 / ISSUE 3

BEYOND THE LAB

For many companies, the time has come. After building new data repositories and testing new analytics software, they are eager to push their big data systems from the IT lab into the hands of the business.


This edition of *Unleashing IT* highlights the maturation of big data technologies and strategies, and how companies are transitioning from one-off pilot projects to powerful enterprise capabilities that can be leveraged on a daily basis.

For communications titans like CenturyLink, this transition has required integration between new data lakes and preexisting enterprise systems (page 10). For analytics companies like Quantum, making a “giant leap forward” involved an entirely new set of data resources (page 8). And for others, it demands new tools (pages 6 and 7) and fundamentally different IT processes (pages 5 and 14).

Regardless of the data, tools, and processes being utilized, you need a rock solid infrastructure to take big data from the lab to an enterprise capability.

For more information, follow the links inside or contact Cisco at 1-866-428-9596. We welcome your feedback on the articles in this publication at www.UnleashingIT.com.

Sincerely,


DD Dasgupta
Vice President
Cisco Systems, Inc.


Ron Kasabian
Vice President
Intel Corporation

STRATEGIES & SOLUTIONS

INFLECTION POINT	3
IT IN REVERSE?	5
PUTTING ANALYTICS INTO THE HANDS OF BUSINESS USERS	6
BIG DATA TOOLS FOR THE LONG HAUL	7
GOING DEEP WITH CONSUMER ANALYTICS	8
MONETIZING A TREASURE TROVE OF DATA	10
A NEW RESOURCE FOR IT TROUBLESHOOTING	12
THE HOSPITALITY INDUSTRY REGAINS CONTROL	14

Unleashing IT is published by Cisco Systems, Inc. To receive future editions of *Unleashing IT*, visit UnleashingIT.com.

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco, the Cisco logo, Cisco ACI, Cisco Unified Computing System, and Cisco UCS are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, visit: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1607)

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

I INFLECTION POINT



How companies are turning their big data science projects into enterprise advantages.

Countless companies have taken preliminary steps to extract more value from their data. Some have adopted the latest open source repositories to reduce the cost of storage and free up precious space in their enterprise data warehouse. Others have brought together a limited number of data sets to conduct trial analyses and see what insights can be attained.

Today, many of these companies are reaching an inflection point. One where cost reduction and experimentation must give way to production capabilities and revenue generation.

“Most companies are over the hump of gathering and managing high volumes of data,” says Ron Kasabian, vice president and general manager of big data solutions at Intel.

“And now they are looking for more return on their big data investments. They’re ready to make the leap.”

As with any advanced technology, this leap requires the right underpinnings—an infrastructure that delivers world-class performance, management, and scalability.

“As they move beyond pilot projects and targeted use cases, companies must learn how to deal with the increased complexity and scale,” says DD Dasgupta, vice president of enterprise product marketing at Cisco. “It requires an infrastructure that can bring together Hadoop platforms, analytics software, and other enterprise systems with immense processing power and automated, policy-driven governance.”



PREPARING TO SCALE AND EVOLVE

Scalability is often the greatest inhibitor of a company's big data evolution. Pulling a few data sources into an isolated data lake is one thing. Correlating hundreds of disparate data streams across thousands of nodes and dozens of enterprise systems is quite another.

"Scalability and management are critical, especially as IoT [Internet of Things] becomes more mainstream," says Dasgupta. "These new workloads are so different from traditional enterprise data, which is more structured and static. Companies need an infrastructure that can align new and legacy systems without negatively impacting the security and governance policies that are already in place."

The infrastructure must also be able to support an evolving set of software tools.

"Hadoop has been around for a decade, so it has matured quite a bit. But analytics software is still relatively nascent and continues to change at a rapid clip," says Kasabian, who recommends infrastructure consistency with software flexibility. "You need to be open minded and willing to change course as new tools and capabilities come to the fore. At the same time, you don't want to overhaul the infrastructure every time you adopt a new piece of software."

A POWERFUL FOUNDATION

Cisco and Intel work closely with big data vendors to make sure Hadoop distributions and analytics software perform best on the Intel® Xeon® processor-based Cisco Unified Computing System™ (Cisco UCS®).

"Cisco UCS Integrated Infrastructure for Big Data provides a secure and scalable infrastructure to support enterprise requirements," Forrester stated in a recent report, listing Cisco as a "Leader" in Big Data Hadoop-optimized systems¹. "Cisco's UCS solution comes pretested and prevalidated for Cloudera, Hortonworks, IBM, and MapR, providing a lower cost and scalable storage platform to support Hadoop deployments."

It's a powerful foundation that includes centralized management and policy automation, which are critical as a big data environment grows.

"Going from less than ten nodes to hundreds or even thousands of nodes is a huge difference," says Dave Kloempken, director of global data center solution sales at Cisco. "You need to spin them all up, load them with software, and keep the firmware up to date. Without centralized management and automation, it can be extremely difficult to expand an environment and maintain that level of scale."

"Management tools such as Cisco UCS Manager and Cisco UCS Director allow for simple configuration of big data Hadoop clusters that can adapt dynamically to changing workloads," the Forrester report noted. "Cisco's key differentiators lie in its ability to offer a wide range of configurations, its strong focus on Internet of Things (IoT) use cases, and its broad partner ecosystem."

¹ Forrester, *The Forrester Wave™: Big Data Hadoop-Optimized Systems, Q2 2016*, May 2016.

GET ANALYST INSIGHT

To download the IDC white paper, "The Business Value of Cisco UCS Integrated Infrastructure for Big Data," visit UnleashingIT.com/bigdata.

IT IN REVERSE?

Turning traditional IT methods on their head to maximize the value of big data.

IT processes, especially those involving data management and analytics, have historically been premeditated and methodical. As business needs arise, IT specialists define the requirements, identify the data sources, and establish the schema and Extract, Transform, and Load (ETL) processes—all in painstaking fashion. Only after those steps have been completed will most IT teams build an analytics system to support the business need.

“Traditional data warehouses were so rigid and expensive, they forced everyone to be extremely deliberate,” says Nik Rouda, a senior analyst at ESG focused on big data. “But the economics and possibilities have shifted with Hadoop.”

The traditional “Schema On Write” approach of data management requires a lot of forethought and upfront IT involvement, he explains, whereas Hadoop’s “Schema on Read” approach empowers users to quickly store data in any format and apply structure in a more flexible and agile way.

“To maximize the potential and value of analytics, companies need to rethink how they capture, collect, and curate their data,” says Scott Gnau, CTO of Hortonworks, a leading provider of enterprise-ready data platforms and applications. “The traditional waterfall approach is too slow.”

For data analytics purposes, Gnau recommends flipping standard IT processes on their head.

“With big data, you don’t always know what you’re looking for or what you’ll find, so you can’t meticulously define everything upfront,” he explains. “You need to think about data and data-related projects in reverse order, where requirements are determined at the very end—after the data has been landed and analyzed. It’s not hard, but it’s very different.”



EMPOWERMENT WITH CONTROL

Most legacy technologies can’t operate without schema, but Hadoop-based systems allow any type of data—structured and unstructured, at rest and in motion—to be stored and explored. According to Gnau, ad hoc analyses often lead to the greatest insights.

“Invest in data scientists and arm them with data,” he recommends. “The key is providing access to the data and facilitating experimentation in ways that don’t hinder security or introduce risk.”

This often requires new governance processes, and a big data platform that complements legacy systems and data sources.

“There needs to be a balance between the two worlds, a balance between business empowerment and IT control,” Rouda claims. “Big data doesn’t get a pass because it’s new. It still has to meet enterprise requirements for security and governance.”

Fortunately, front-end tools are getting better, he notes. Especially those that span new data lakes and legacy business intelligence systems, allowing data to be queried without requirements or schema. And without turning an enterprise IT environment into the Wild West.

“At first blush, doing IT in reverse seems like a horrible idea,” Rouda chuckles. “But the more you think about it, the more it makes good sense.”

To learn how Cisco and Hortonworks can help you take advantage of big data opportunities, download the solution brochure at UnleashingIT.com/bigdata.

PUTTING ANALYTICS INTO THE HANDS OF BUSINESS USERS

How self-service data preparation tools are helping business teams unlock the value of big data.

As Bob Eve sees it, big data should be a team sport. One where the IT department deploys and maintains foundational tools and business teams use those tools to perform ad hoc analyses and unearth new insights. But despite the emergence of powerful data management platforms and sophisticated analytics software, most teams are stuck on the practice field.

“Eighty percent of every analytics project is devoted to data preparation,” says Eve, a director of data and analytics products at Cisco. “It needs to be easier and faster to find, combine, and normalize data in advance of analysis. Not just for IT specialists and data scientists, but for everyone.”

To get the most out of big data, companies need to put the power of analytics into the hands of business users. After all, they are the ones with the questions and the ability to act on the answers. Access to data is the key, but it must be given in a way that doesn't create chaos in the IT environment.



CONFRONTING A BOTTLENECK

Many companies have built data lakes to store large volumes of information and advance their analytics capabilities. But the results haven't always matched the hype.

"Companies are getting frustrated because they are pouring tons of money into their big data systems and not seeing a comparable return," says Michele Goetz, principal analyst at Forrester. "But you can't just dump IoT data into a Hadoop environment, for example, and expect business insights to magically appear. Sensor data is just a log. It doesn't mean anything without some context or correlation."

Today, aggregating diverse data sets from diverse systems and applying contextual relevance is largely a manual process, the complexity of which places a tremendous strain on IT groups and slows down the business.

"Big data environments and analytics appliances aren't the problem," says Goetz. "It's the bottleneck that is created when every request and query is reliant on the IT organization."

What's needed is a self-service, front-end data portal, she suggests. One that is platform agnostic and helps business users find, pull, and prepare data from a number of repositories. One that allows data manipulation and experimentation without adversely affecting the underlying systems or overarching governance policies.

SELF-SERVICE DATA PREPARATION

Cisco® Data Preparation—which runs on the Intel® Xeon® processor-based Cisco Unified Computing System™—allows raw data to be quickly and easily gathered, combined, and enriched. The self-service application puts the power of analytics into the hands of nontechnical business users.

"Data from multiple sources can be integrated, cleansed, and explored without coding or scripting," says Eve. "It works a bit like Excel, where columns can be added and things can be moved around. And you don't have to be an IT whiz."

What may seem like a risky relinquishing of control can greatly benefit an IT organization, says Goetz.

"IT groups can focus on foundational capabilities instead of one-off projects," she says. "They can maintain security and access parameters that protect underlying systems and the data within. And they can learn which data sources and queries are most valuable to the business, and use that knowledge to continually enhance their analytics capabilities."

Self-service data preparation not only brings IT and business teams closer together. It helps them get off the big data practice field and into the game.

WATCH THE VIDEO

To see an introductory video on Cisco Data Preparation, visit cisco.com/go/datapreparation.

BIG DATA TOOLS FOR THE LONG HAUL

Where a company starts with big data isn't necessarily where it will end up. As with any advanced technology, the learning and capability curve can take years.

"Every organization is at a different stage of big data maturity," says Pandit Prasad, manager of Hadoop and open analytics systems at IBM.

The first stage, he explains, is typically focused on cost reduction, supplementing or replacing expensive enterprise data warehouse (EDW) systems with commodity hardware. Many companies will leverage their new hardware environment to create a data lake—the second stage—that supports data governance, cleansing, and matching.

"Most companies are in these first two stages," Prasad claims. "It's a matter of building foundational capabilities and bringing information assets together. But to get the most out of big data, those capabilities must be turned over to the business."

This handoff is the third and arguably hardest stage of a company's big data evolution. And it must be carefully considered in the very beginning.

"We always recommend starting small," says Prasad, "but you have to have a long-term vision and the tools to support it."

There are different tools for different stages of big data maturity. Companies focused on cost reduction and EDW offloading, for example, may have little need for sophisticated data science and analytics tools—but they might eventually.

"As your big data capabilities and objectives evolve, you don't want to find yourself with incongruent tools," Prasad says. "It's a huge amount of work to integrate new components that don't align with your pre-existing tools. And you certainly don't want to redesign your entire stack just to take advantage of something like Spark 2.0 or new machine learning languages and libraries."

This underscores the need for enterprise-grade platforms that come with an entire suite of integrated tools. IBM's Open Platform and BigInsights, for example, leverage Hadoop and Spark to support each stage of a company's big data evolution. The platform is currently being integrated with the Intel® Xeon® processor-based Cisco Unified Computing System™ for world-class performance and manageability.

"Our philosophy is to provide a stable and reliable Hadoop and Spark data platform while supplementing that platform with complementary systems and premium components to provide the enterprise-class functionality clients expect for their mission-critical data," says Prasad.

In doing so, IBM is helping its clients focus on their data—not their tools.

Get a Cisco and IBM big data solution overview at UnleashingIT.com/bigdata.

GOING DEEP WITH CONSUMER ANALYTICS

With a depth and breadth of consumer information and a powerful platform to process it, Quantum is realizing the true potential of big data.



Quantum has been processing data and conducting in-depth analytics since its formation in 2002, helping companies better understand their market and their customers. But until recently, most of that data came directly from its clients.

“We saw an opportunity to take a giant leap forward,” says Alex Shaw, head of technology operations at Quantum, Australia’s leading data analytics company, “by bringing our own data assets to the table.”

Retailers, insurers, healthcare providers, and financial services companies, to name a few, are eager for more knowledge about their customers. And any company with consumer-oriented products or pricing complexity can benefit from the analyses that Quantum provides. But the resulting insights are only as deep as the data being scoured.

“You can only go so far with in-store transactions,” explains Shaw, using the retail industry as an example. “But when you marry that data with online purchases and media consumption, you get a much more complete view of the customer. Those insights can be used for more targeted advertising, promotions, loyalty programs, inventory, product placement, you name it.”

To bolster its data resources, Quantum formed partnerships with media, banking, supermarket, and cable television providers in Australia. The data it receives from those partners

is extremely detailed, but anonymous to ensure full compliance with all privacy requirements. The information provides deep insights into the spending habits, product preferences, and media consumption—both past and present—of Australians.

“We have a breadth and depth of consumer data that is unmatched in the region,” says Victor Bajanov, a member of Quantum’s senior leadership team. “Compliance and privacy are very important to us, so we don’t have individual names, addresses, or account information. Essentially, we hold no personal information but can analyze every credit card transaction and every line of a grocery store receipt. And when you combine and correlate data from multiple sources, you get a very granular picture of modern consumers.”

REDUCING QUERY TIME BY 92 PERCENT

Quantum’s pre-existing analytics engine couldn’t support its desire to process more data and conduct faster analyses, so the company decided to build a world-class big data platform from the ground up. After a rigorous selection process, Quantum chose the MapR Converged Data Platform running on the Intel® Xeon® processor-based Cisco Unified Computing System™ (Cisco UCS®).



“Our banking partner has more than two million customers who generate 14 million transactions per week. That’s a lot of data to crunch, and it’s just one source,” says Shaw. “We chose MapR because it had the best performance and Cisco UCS because it had the best management capabilities and lowest power consumption.”

The new platform has reduced query times by 92 percent, allowing Quantum data scientists and engineers to design complex queries that run against multi-terabyte data sets—and get results in minutes.

The architecture is also built for multitenancy and scalability. Quantum can segment and safeguard client data while simultaneously giving clients access through a secure portal. And the clustered environment will be easy to expand as the business grows.

FLEXING AN “INNOVATION CENTER”

Beyond speed and scalability, the new platform is fostering innovation, helping Quantum develop new products and enter new markets.

“We have a lot of smart people who have been hamstrung by technology and its inability to implement their ideas,” says Bajanov. “Now they have a powerful platform at their disposal to test new scenarios quickly and cost effectively. It’s an innovation center that allows us to create new solutions and services for our clients.”

Quantum is already flexing the muscle of its new analytics engine for Woolworths, the largest supermarket chain in Australia. The platform is crunching the data from nine million loyalty program members, 25,000 products, and 52 weeks of purchasing behavior to automate the selection of merchandise that is featured in weekly promotional emails to members.

“It’s a massive data set, and we are processing it weekly,” says Shaw. “If every person on the planet were doing the same calculation, we would all need to produce six decisions a second to keep up with our MapR and Cisco cluster.”

This is the true potential of big data. And it’s being realized Down Under.

To see how Quantum is using Cisco UCS and MapR Converged Data Platform to revolutionize consumer analytics, watch the webcast at UnleashingIT.com/MapR.

MONETIZING A TREASURE TROVE OF DATA

CenturyLink has built a new big data platform that will fuel a variety of industry-specific services.

As one of the largest communications, hosting, cloud, and IT service providers on the planet, CenturyLink has a treasure trove of data. Roughly 20 percent of the world's Internet traffic flows through its fiber. And 98 percent of Fortune 500 companies utilize its services, which are delivered from 60 data centers spread around the globe.

Today, CenturyLink is working to monetize its data riches.

"All telecommunications companies are going through a transition," says Shivkumar Thiagarajan, vice president of business development for CenturyLink Cognilytics, the company's new big data and predictive analytics division. "We want to leverage our core assets to reach new markets and create higher value services for our customers."

From its network, data center, and cloud footprints to its core IT and application services, CenturyLink has a wealth of data at its disposal. But to maximize the value within, it needed to find ways to combine and analyze those sources collectively.

BEYOND PILOT PROJECTS

Most big data projects—at least in the early stages—are simplified and limited by design, pulling together a few workloads and seeing what can be learned. To get the most out of information assets, however, companies must move beyond discrete data lakes and finite pilot projects.

"When everything is the same or carefully curated, it's easy to automate and scale," says Richard Probst, vice president of architecture at SAP. "But at some point, you need to integrate big data capabilities and pilot projects with enterprise systems."

With the help of SAP, Cisco, and Cloudera, CenturyLink did just that. The company has created a large-scale Hadoop data lake using Cloudera Enterprise. The data lake integrates with SAP HANA® and SAP HANA Vora, and runs on the Intel® Xeon® processor-based Cisco Unified Computing System™ (Cisco UCS®).

"It's a platform that delivers the scalability of Hadoop, the intelligence of SAP HANA, and the power of Cisco UCS," says Thiagarajan. "It's very unique. No other architecture can do what it does."

The platform can accommodate a variety of data sets from any number of systems. And it delivers policy-driven automation, scalability, and control using Vnomic software and Cisco Application Centric Infrastructure (Cisco ACI™).

"Policy-based governance is key," says Probst, "transforming big data from a science project to an enterprise capability."

PUTTING THE PLATFORM TO THE TEST

CenturyLink is already using the platform for its own purposes. The company is combining and analyzing finance and supply chain data to enhance operational performance and reporting.



It is scouring the entirety of its global fiber network to improve signal and anomaly detection. And it will soon collect data from its 1600-vehicle fleet in North America to improve preventive maintenance and reduce downtime and costs.

"We wanted to put the platform to the test, and it has performed exceptionally well," says Thiagarajan. "All of our predetermined KPIs have been validated."

CenturyLink is now in the process of developing industry-focused solutions that take advantage of the platform and can be delivered as a service. Manufacturing, healthcare, retail, and other data-driven services are all in the works.

"The infrastructure, plumbing, and governance used to be the hard parts, but we've got those covered," says Thiagarajan. "Now it's just a matter of identifying the business problems, the desired outcomes, and the data that can help. The platform is ready to go."

LEARN MORE

The combination of Cisco UCS, Cisco ACI, and SAP HANA Vora helps unlock the intelligence in your data and interpret it with a new dimension of context and insight.

To learn more, download the solution brief at UnleashingIT.com/bigdata.

A NEW RESOURCE FOR IT TROUBLESHOOTING

Why machine data analysis is becoming increasingly critical to the health and performance of an application environment.

BOOST IT OPERATIONS

IT operations have never been more complex—or more critical to the businesses they serve. Increasing application and infrastructure performance requires real-time, end-to-end visibility across applications and physical, virtual, and cloud environments.

Cisco and Splunk provide a comprehensive, highly efficient IT operations analytics solution. To learn more, download the At-a-Glance brochure at UnleashingIT.com/bigdata.

Application monitoring and troubleshooting used to be far easier. If something went wrong, there were only so many places where the error could have occurred. Only a few potential culprits.

In today's world of hybrid IT, however, where systems and applications are spread among physical, virtual, and cloud environments, finding the root cause of a problem can be daunting. What used to be a defined trail of breadcrumbs has been scattered across the floor, in a variety of rooms.

"Modern application environments are so distributed and complex," says Jon Rooney, senior director of solutions marketing for Splunk, a leading provider of operational intelligence and data analytics solutions. "If a problem occurs, was it the application server, the web server, the database tier, the API gateway, something else? Without full visibility across all of them, you just don't know."

Application Performance Management (APM) tools can help, he adds, but they are inherently limited. While they can detect hiccups with availability and performance, they can't show where or why the problem is happening unless it is related to the application code itself.

That's why companies are beginning to combine APM information with other kinds of data—like machine data. Created by every application, operating system, network device, and IT service, machine data comes in a variety of formats, from logs and wire data to mobile application and API endpoint data.

"Log data is becoming increasingly valuable," says Tim Grieser, program vice president for enterprise system management software at IDC. "But you need to combine and correlate the logs from multiple systems to effectively triage, troubleshoot, and remediate problems."

REAL-TIME MONITORING AND PATTERN DETECTION

Many companies disregard their machine data entirely. Some have developed what Rooney calls "duct tape and super glue solutions" for monitoring and managing machine data formats like log files. Others have attempted conventional methods to process these unconventional files.

"You can't monitor and analyze log files with a relational database. That's like placing a fire hose in a Dixie Cup," Rooney quips. "Batch processing isn't viable either, because you need to analyze log files in real time and respond immediately to anomalies or performance dips."

Splunk provides industry-leading software to aggregate, integrate, and analyze every type of machine-generated data—including logs—in both structured and unstructured formats. Pretested and prevalidated on the Intel® Xeon® processor-based Cisco Unified Computing System™, the software helps identify, resolve, and prevent operational, security, and business issues through real-time monitoring and pattern detection.

"Companies have been doing this in bits and pieces, but they have an opportunity to go deeper," says Grieser. "They can correlate logs from many systems to see across the layers of an environment, from the server and virtualization layers to the network and web layers to mobile and device layers."

AN ADDED DIMENSION

With a depth and breadth of visibility, organizations don't just improve their ability to troubleshoot problems. They also increase their understanding of the systems and circumstances that affect application behavior. And they can work to proactively prevent problems altogether.

"Troubleshooting is a big part of the equation, but effective log monitoring can lead to predictive and preventive IT maintenance," says Grieser. "Organizations can apply machine learning to better understand application usage patterns and contingencies. And they can set dynamic, automated thresholds and alerts that help circumvent potential problems before they occur."

In doing so, they can better manage the health and performance of their application environment, which is vital for the health and performance of the business.

"Log management helps blur the lines between IT and business performance, and can provide visibility on both sides," says Grieser. "It's an added dimension that is well worth exploring."



THE HOSPITALITY INDUSTRY REGAINS CONTROL

With Wi-Fi-based location tracking and analytics, hotels are rebuilding their relationship with guests.

In many ways, the hospitality industry has lost control. Rooms and rates are more commoditized. Online reviews are more influential. And powerful online brokers are usurping portions of the guest experience—and the profits.

“It wasn’t long ago that people would call a hotel directly to inquire about rates, check availability, and make a reservation,” says Natalie Osborn, hospitality lead for SAS, a leading provider of analytics and business intelligence software. “But that one-to-one relationship is gone.”

With intermediaries like Expedia, Travelocity, Orbitz, and Kayak providing comparison pricing, showcasing deals, and facilitating reservations, hotels have lost a significant measure of engagement and influence. And many are looking to build and reinforce their relationships with guests in other ways.

“Casinos have always done a great job with consumer analytics, leading to marketing investments and personalized services that keep guests coming back,” says Osborn. “Hotels are trying to do more of that.”

Their efforts go well beyond loyalty programs, she explains, which have become as ubiquitous and expected as airline miles. More than ever, hotels have an opportunity to

monitor and learn from multiple touchpoints—even those that are no longer under their control.

Before guests arrive, hotels can track how and when they booked a room, and whether it was at a discounted rate or part of a promotion. During their stay, hotels can analyze when they check in and out, how often they visit the gym, and whether they remain onsite for meals. And after their stay, hotels can monitor the feedback and reviews posted to social media.

“Hotels have a wealth of data at their disposal,” says Osborn. “They just need to tune in, listen, and learn.”

THE IMPORTANCE OF WI-FI

In addition to reservations, point of sale, and social media, Wi-Fi-based location tracking has become an increasingly valuable data source for the hospitality industry. It reveals the movement, patterns, and preferences of consumers when they are on the property—which remains firmly within a hotel’s sphere of control and presents the best opportunities to improve the guest experience.



SAS has been integrating its customer intelligence product suite with Cisco Connected Mobile Experiences (Cisco CMX) to deliver advanced location tracking and analytics. In stark contrast to other hospitality solutions, it doesn't require the use of a mobile app.

"It's tough to get people to download and use an app," says Osborn. "Our solution is passive and tracks location based purely on Wi-Fi connectivity."

This provides a wealth of opportunity to learn more about guests and respond in real time. But Osborn is quick to note that push notifications and unsolicited promotions can be a turnoff.

"You can't be creepy or intrusive," she insists. "Hotels need to stop thinking of marketing as sending an offer or a coupon to someone. Sometimes it's a friendly message or an in-person greeting. Sometimes it's nothing at all."

With real-time location tracking, hotels can be alerted when a guest enters the building and prompt their staff to welcome them—by name—before they reach the check-in counter. They can see when guests have been standing in line or waiting for a table, and quickly remediate such annoyances. They can personally

extend gym hours for guests who like to exercise in the evenings. And they can give unexpected upgrades and special amenities to loyal guests—even if they aren't a loyalty program member.

"Some hotels don't know a VIP was there until they leave. With better tracking and analytics, they can truly understand behaviors and preferences, they can anticipate needs, and they can proactively deliver personalized services," says Osborn. "And that goes for every guest, not just VIPs."

MAKE FASTER, BETTER DECISIONS

Succeeding in today's dynamic business environment requires faster, more informed decision making. These decisions often rely on complex analyses of ever-growing accumulations of data—in real time.

To learn how Cisco and SAS enable real-time analytics, get the solution brief at UnleashingIT.com/bigdata.



The center of endless possibilities

Insight is the key to disruption—changing the status quo. Cisco solutions for data centers and real-time analytics, built on Cisco UCS servers, help you turn your big data challenges into impactful business decisions.

cisco.com/go/disrupt

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.

Cisco UCS®
with Intel® Xeon®
processors

