Sponsored by: **Cisco**

Authors:
Peter Rutten
Dan Vesset
Matthew Marden

February 2016

The Business Value of Cisco UCS Integrated Infrastructure for Big Data

EXECUTIVE SUMMARY

The hype surrounding technology capabilities and inflated promises of outcomes from Big Data projects has finally subsided. As evidence mounts for the competitive edge demonstrated by organizations embracing data-driven decision making, Big Data and analytics (BDA) has become a top agenda item for a growing number of executives.

Today, a new, more pragmatic view is taking hold, dominated not by discussions about the volume, velocity, and variety of data but by the value of data — the ability to generate actionable insights and deliver them to all the relevant stakeholders inside and outside the organization. With the opportunity to unlock the value of Big Data to accelerate innovation, drive optimization, and improve governance comes the need to navigate expanding technology alternatives, recreate business processes, and ensure the availability of appropriately skilled staff.

Gone are the days of endless arguments about the superiority of either relational or non-relational technologies. The reality is that the emerging Big Data architectures encompass the best of both deployed on optimized infrastructure for a broad set of data exploration and discovery, operational intelligence, and performance management use cases.

To better understand the value component of Big Data, IDC conducted interviews with organizations using the Cisco UCS Integrated Infrastructure solution as a platform for their Big Data operations (Cisco UCS for Big Data). These organizations reported creating substantial value through Big Data operations, with the Cisco UCS platform supporting this value generation with scalability, performance, and cost and time efficiencies. IDC projects that on average, these organizations will achieve a three-year return on investment (ROI) of 366%

Business Value Highlights

Three-year ROI
366%

Breakeven in just under
5 months

Faster Big Data deployment*
56%

Faster server deployment*
62%

Lower operational cost*
32%

** than other solutions considered by the interviewed organizations*

through their deployment of Big Data solutions on Cisco UCS, with Cisco UCS contributing to this value by:

- **Reducing time to value from Big Data projects.** The ease of deploying Cisco UCS as a platform for Big Data means that organizations capture its benefits earlier.
- **Making Big Data solutions scalable.** Supporting expanding Big Data operations with servers and other datacenter resources takes less time for these organizations with Cisco UCS, enabling them to better support their businesses and users and to ensure sufficient capacity for their Big Data operations.
- **Ensuring optimized Big Data solution performance.** Interviewed organizations noted the strong performance of their Big Data operations on their Cisco UCS platform.
- **Having a lower cost of operating Big Data solutions.** Cisco UCS costs less and needs less staff time to manage and support than other solutions these organizations considered for running Big Data operations.

Situation Overview

The Big Data and Analytics Imperative

Many factors have a direct or an indirect impact on the decisions made by IT and business leaders as they work on BDA strategy and plan for specific projects. Some of the key drivers are:

- **Digitization of everything.** The availability of personal data, data from the Internet of Things, and enterprise data from digital transformation efforts is driving demand for solutions to take advantage of all this data to derive insights and actions. IDC research shows that in the past 12–24 months, 77% of large and midsize organizations expanded the number of data types and sources they analyze.
- **Promise of predictability.** Organizations across industries are seeking predictability across all fronts, such as customer interaction, cash management, service delivery, physical asset management, and product development and support. One manifestation of this trend is the demand for a range of applications with predictive functionality.
- **Demand for self-service.** Historically, the state of the art in business intelligence has involved the delivery of reports or dashboards that answer predefined questions to decision makers. More recently, self-service analytics tools have moved analytic capabilities into the hands of business users. This is fundamentally changing how organizations interact with data, how they develop new hypotheses and scenarios, and how they react to changes in the market. This

driver is beginning to affect not only the "last mile" data visualization and exploration software but also the data acquisition and preparation steps of the full analytics life cycle.

- **High demand for and scarce supply of Big Data business/IT skills.** The ability to acquire BDA competency is constrained by a talent pool that is growing too slowly to meet business demand. While much of the focus in the BDA market has been on the scarcity of data scientists, the same issue applies to other data and infrastructure management professionals. In the absence of a short-term solution to labor availability, technology takes on added importance in fulfilling the demand-supply gap.

These drivers and other drivers are influencing Big Data technology purchasing and deployment across industries and regions. IDC estimates that in 2015 organizations spent over $20 billion on Big Data technology and services. This spending, which includes functionality for data integration and preparation; data management using massively parallel processing (MPP) relational databases, various NoSQL databases, and Hadoop; advanced and predictive analytics and data visualization; and text and rich media analysis, was distributed across a wide range of use cases. Examples include:

- **Retail.** Cross-channel customer behavior analysis, physical store layout optimization, and product cross-selling
- **Banking.** Transaction fraud detection and prevention and risk management
- **Utilities and energy.** Predictive and preventative asset maintenances
- **Healthcare.** Genetic research and diagnosing, drug effectiveness measuring, and patient outcomes analysis
- **Logistics and transportation.** Route optimization and fuel efficiency management
- **Media and entertainment.** Content distribution optimization, audience behavior analysis, personalized ad targeting, cross-industry security analytics, and IT operational analytics

Infrastructure Implications

Increasing Big Data initiatives have infrastructure implications; enterprises need to make deliberate infrastructure decisions that take into account specific workload requirements. These require an understanding of analytic and decision-making processes and the needs of specific internal and external users (including data scientists, business analysts, managers, and front-line employees as well as suppliers and customers). Data volume, variety, and velocity; query complexity; user concurrency; skills availability; and other variables influence infrastructure decisions and in turn determine the value derived from data.

The required infrastructure for a Big Data environment cannot simply consist of an extension of the enterprise's data warehouse because this approach lacks scalability and can cause an undesirable increase in capex and opex due to server sprawl. Growing small clusters of standard servers into large clusters is not efficient in terms of hardware costs (especially switches) and from a cluster management perspective. IDC research has shown that the server sprawl from growing standard servers into large clusters can cause opex to become much greater than capex.

A Big Data infrastructure needs to be high performing and highly available, must be able to handle large variations in demand, and has to scale well (ideally linearly) and easily. The infrastructure cannot be overly complicated to manage, taking up large amounts of staff time and leading to staff productivity issues, and it must be capable of rapidly deploying resources and redeploying them. It also needs to deliver a high utilization rate, and it has to be secure, including an ability to isolate sensitive data in a shared environment.

The following section of this white paper investigates Cisco UCS for Big Data infrastructure, which Cisco has designed as a "programmable infrastructure" to address these infrastructure requirements.

Cisco UCS Integrated Architecture for Big Data

Cisco UCS for Big Data is a comprehensive, automated architecture to simplify the operational complexity of managing Big Data clusters. Compute in the system consists of up to 16 Cisco UCS C240 M4 servers per rack. The servers run on dual Intel Xeon E5-2600 v3 processors, which Intel designed for high-performance infrastructures. The processors support up to 768GB of main memory.

Storage in the C240 M4 is provided with two approaches in mind: to support performance or to support capacity. The performance-optimized approach provides 24 small (in terms of form factor) disk drives; the capacity-optimized approach uses 12 large disk drives. The system facilitates data tiering, and for less frequently accessed data, a Cisco UCS C3260 Rack Server can be included for up to 360TB of archival RAID storage.

The network consists of UCS 6296 Server Fabric Interconnects, which are installed in pairs for redundant active-active connectivity. The system's performance and scalability advantages are directly related to the network design as the network acts as a single point of high-bandwidth, low-latency connectivity that provides IT with unified management of the entire system through UCS Manager. UCS Manager allows for very fast server configurations using service profiles, automated maintenance, and cluster health monitoring.

In addition, Cisco's Application Centric Infrastructure (ACI) provides high-level policies to dynamically improve cluster performance, automatically reposition parts of a cluster for other types of workloads (for easy redeployment of resources), and ensure secure isolation of users as well as data streams in this shared environment.

Cisco UCS for Big Data has been designed to achieve various requirements for a Big Data cluster environment, with the most critical being:

- **Scalability:** A single rack consists of 16 servers, and a single domain contains up to 80 servers (or up to 160 servers if outfitted with Cisco Nexus 2232PP Fabric Extenders). Scaling beyond 160 servers is achieved by interconnecting multiple UCS domains with Cisco Nexus 9000 switches. This way, a cluster can consist of thousands of servers and petabytes of storage, all managed with a single pane of view on the entire infrastructure using Cisco UCS Central, whether it is housed in a single datacenter or in multiple datacenters around the world.
- **Low TCO:** This IDC Business Value study demonstrates that UCS for Big Data delivers an advantageous compute price/performance for Big Data. Some of the savings are achieved with the system's switching and cabling technology SingleConnect, which requires fewer cable and switches and services up to 160 servers without the need to add cables or switches. Another benefit from SingleConnect is that only one connection type is needed for all the system's protocols whether LAN, SAN, or management. SingleConnect further provides direct SAN access as well as automated I/O bandwidth allocation.
- **High performance:** The processors and storage components provide the required processing-intensive performance for Hadoop 1.0 workloads. However, with Hadoop 2.0, clusters are getting larger, the data variety is increasing as is the mix of workloads, and the sheer number of processes is growing rapidly, challenging the network more than the processors. In response, Cisco has designed the system for greater network performance via the aforementioned network capabilities.
- **Easy, comprehensive management and maintenance:** Cisco UCS for Big Data is managed with UCS Manager, which oversees the entire infrastructure, enables easy configuration of Hadoop nodes, and provides an opex benefit with its ability to quickly adapt to changing workloads. With single-pane management, UCS Manager gives IT complete control of up to 6,000 servers. In addition, UCS Director provides central, integrated, and policy-based management of the Hadoop applications running on the cluster, essentially single-pane control of the hardware and the software. UCS Director lets IT provision Hadoop clusters on-demand and provides the flexibility for such clusters to coexist with enterprise applications.

Cisco UCS for Big Data supports Hadoop solutions from Actian, Cloudera, DataStax, Elastic Search, Hortonworks, IBM, Intel, MapR, Microsoft, MongoDB, Pivotal, Platfora, SAP, SAS, and Splunk.

The system can manage the data growth stemming from streaming and interactive workloads that Hadoop 2.0 processes and the new applications that are evolving from this environment. It supports NoSQL solutions including Oracle DB, DataStax, and MarkLogic as well as MPP relational databases Greenplum DB (Pivotal Labs) and ParAccel (Actian).

Application Centric Infrastructure

ACI deserves some additional mention as this solution is strongly benefiting the overall efficiency of the UCS for Big Data system. ACI improves cluster performance by lowering network latency through telemetry that is embedded in the Nexus 9000 leaf switch. It monitors the network and makes adjustments to reduce network congestion. These adjustments may include load balancing data across racks to alternate paths, prioritizing data flows (smaller first), and estimating when a job will be completed so as to prioritize accordingly. These actions are important in a Hadoop environment as Hadoop copies data across the cluster and distributes processes across racks. Mixed workloads with real-time and/or interactive data put further demands on the system.

ACI also helps improve resource utilization. It uses application policy profiles to reposition the infrastructure (the same network, compute, and storage policies used by such tools as UCS Director), with a different profile for every specific need. Cluster resources can be redeployed quickly, enabling a highly efficient utilization scenario with, for example, enterprise applications such as CRM or ERP running during the day and Hadoop jobs at night.

In addition, ACI provides data isolation and security in a heavily shared environment. ACI isolates a data path by using programmable fabric based on predefined policies, while the data maintains vault access policies in-flight across the infrastructure's physical and virtual endpoints. This is important as Hadoop processes may be fed by numerous data sources and received by just as many clients. Data becomes vulnerable during job transitions or redeployments to other applications. ACI helps avoid the typical situation in which multiple databases and users and compliance mandates lead to a tough-to-manage-and-scale security patchwork that is vulnerable to human error and that demands significant staff time.

The Business Value Of Cisco Ucs Integrated Infrastructure For Big Data

Study Demographics

IDC interviewed seven organizations in North America and EMEA using Cisco UCS Integrated Infrastructure as a platform for Big Data solutions. IDC designed the interviews to understand the qualitative and quantitative impact for these organizations of deploying Big Data solutions on

Cisco UCS, with an emphasis on how the Cisco UCS platform supports their Big Data operations. By size, organizations ranged from 800 employees to 157,000 employees, with an average employee base of 60,000 (see Table 1). On average, these organizations had deployed 141 servers in their Cisco UCS for Big Data environments as of the time of their interviews and were supporting an average of three business applications requiring 219TB of data with Cisco UCS.

TABLE 1

Demographics of Seven Interviewed Organizations Using Cisco UCS for Big Data

	Average	Range
Number of employees	60,000	800 to 157,000
Number of Cisco UCS servers*	141	18 to 750
Number of applications on Cisco UCS*	3	0 to 10
Number of terabytes on Cisco UCS*	219	1 to 750
Number of users on Cisco UCS*	7,500	15 to 20,000
Countries	United States, Canada, Italy	
Industries	Professional services, financial services, government, automotive, healthcare, cloud services	

n = 7

** within Cisco UCS environment supporting Big Data operations*

Source: IDC, 2015

Overview of Cisco UCS for Big Data Use by Interviewed Organizations

For the most part, interviewed organizations have used Cisco UCS as a platform for deploying new Big Data solutions rather than migrating existing Big Data workloads. As a result, these organizations have seen an impact from running Big Data operations on Cisco UCS mostly on their operations and business, despite cost and staff efficiencies that they attribute to the Cisco UCS platform.

Organizations reported that they are running a number of Big Data solutions on Cisco UCS. Some organizations are running a version of Hadoop on Cisco UCS, and others are supporting an MPP database solution (e.g., Teradata and Vertica) on Cisco UCS. Most interviewed organizations are using Big Data solutions on Cisco UCS for both data management (e.g., data warehousing

and data lakes) and analytic applications or business intelligence. In terms of how they are using Big Data solutions running on Cisco UCS, many of these organizations are supporting business intelligence and exploration/discovery operations, while fewer of them are carrying out online transaction processing or operational intelligence (see Figure 1).

FIGURE 1

Use of Cisco UCS for Big Data

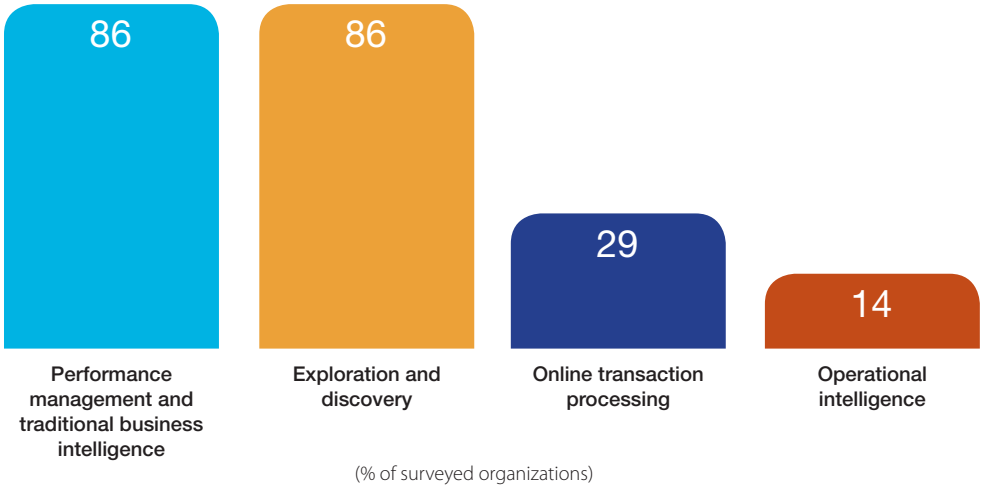

Source: IDC, 2015

Business Value Analysis

Organizations interviewed for this study reported enabling their operations and businesses by running Big Data solutions on Cisco UCS. IDC projects that on average these organizations will realize benefits worth $4.76 million per organization per year over three years ($33,638 per Cisco UCS server) as a result of increased employee productivity, higher revenue, and reduced costs achieved through their use of Big Data solutions running on Cisco UCS (see Figure 2).

These organizations attributed this value to both the Big Data solutions they are using and the Cisco UCS platform. They reported that Cisco UCS creates value for their Big Data operations in the following areas:

» **Business productivity benefits.** IDC projects that these organizations will achieve operational and business benefits worth an average of $4.19 million per year over three years ($29,654 per Cisco UCS server) attributable to their Big Data operations running on Cisco UCS. These benefits include increased productivity for employees such as data scientists and data analysts and higher revenue through services and products supported by Big Data. According to interviewed organizations, Cisco UCS contributes to this value

by enabling earlier deployment, scalability, and higher performance levels of their Big Data solutions.

- **IT staff productivity gains.** IDC calculates that on average these organizations will achieve IT staff productivity benefits — including application development teams — worth $550,000 per year over three years ($3,861 per Cisco UCS server). The Cisco UCS platform supports higher productivity for application development teams by reducing the time to deploy server and other resources, and through higher application performance, the Cisco UCS platform makes their efforts more beneficial.

- **IT infrastructure cost reductions.** IDC calculates that these organizations will achieve IT infrastructure cost reductions worth an average of $17,000 per year over three years ($123 per Cisco UCS server). However, the cost-effective and operationally efficient nature of the Cisco UCS played a substantial role in these organizations' decisions to use Cisco UCS as a Big Data platform. IDC calculates that on average, Cisco UCS will have a three-year cost of operations that is 32% less than that of alternative platforms these organizations considered. In addition, Cisco UCS serves as a highly reliable platform for Big Data operations, with each user losing just over one-half hour of productive time per year due to unplanned downtime.

FIGURE 2

Average Annual Benefits per Cisco UCS Server

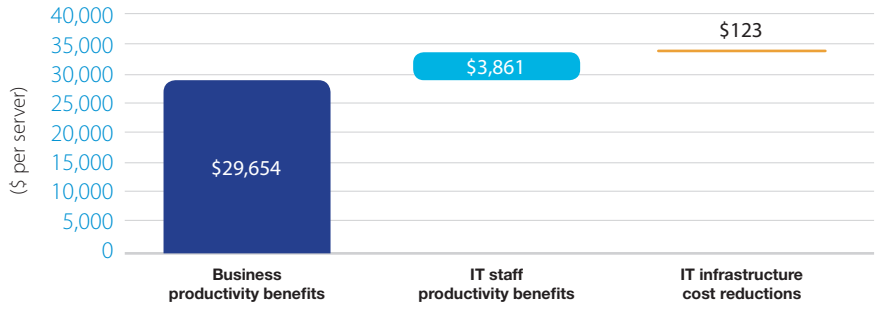

Average annual benefits per server: $33,638

Source: IDC, 2015

Business Productivity Benefits of Big Data on Cisco UCS

Organizations running Big Data solutions on Cisco UCS are achieving increased revenue and higher employee productivity, reflecting the value of Big Data installed on Cisco UCS by supporting the creation and improvement of data-driven, customer-facing services and

products and applications that support internal users. Table 2 reflects the extent to which these organizations are leveraging their use of Big Data solutions running on the Cisco UCS platform to create value for their operations and businesses.

TABLE 2

Business and User Impact of Cisco UCS for Big Data

	Per Organization
Revenue impact	
Average additional revenue per year	$6.6 million
Assumed operating margin	15%
Average net operating margin gain	$1 million
Productivity benefits	
Average number of users* impacted	138
Average productivity gain	8.8%
Average FTE impact	12

** does not include data scientists/analysts or application development teams*
Source: IDC, 2015

Impact of Cisco UCS on Supporting Business and Users

The value depicted in Table 2 represents how interviewed organizations are benefiting from Big Data solutions they are running on the Cisco UCS platform, but they reported that Cisco UCS helps them optimize the business and user benefits they achieve from Big Data operations in two significant ways:

- **Time to market and scalability.** Interviewed organizations reported being able to deploy and scale their Big Data solutions in less time with Cisco UCS than with alternative platforms.
- **Performance.** Interviewed organizations said that the Cisco UCS platform ensures strong Big Data performance.

For interviewed organizations, these benefits they attribute to the Cisco UCS platform enable them to capture more value by reducing the time to value of Big Data solutions, ensuring the scalability and continuity of their Big Data operations, and enhancing the value of Big Data through strong performance.

"Cisco UCS helps us reduce the time to respond to new capacity requests by typically two months. That's significant, because with UCS, you can just drop a rack in there and you're done. The impact is felt throughout the entire organization, because they're waiting for the [Big Data analytics] results."

Time to Market and Scalability with Cisco UCS for Big Data

Interviewed organizations reported that the speed with which they are able to deploy and extend their Big Data operations with Cisco UCS differentiated it from other platforms they considered. With Cisco UCS, they minimized the time required to realize value from and to scale their Big Data operations to meet evolving business demand. These efficiencies stem from the pre-configured and integrated nature of Cisco UCS as well as UCS management software that leverages service profiles and policy to enable fast deployment of server resources. Interviewed organizations reported to IDC that, compared with alternative Big Data platforms they considered, they needed 56% less time to deploy their Big Data solutions on Cisco UCS and can deploy new servers to support expanding Big Data operations in 62% less time (see Figure 3).

Organizations provided examples of how Cisco UCS is helping them achieving value with Big Data:

- **Faster deployment of Big Data.** *"Based on our experience with commodity blades, and the complexity of this deployment with Cisco UCS, I'd say that it would take more than two times longer to do everything with a non-Cisco UCS approach."*
- **Ease of scaling.** *"Cisco UCS helps us reduce the time to respond to new capacity requests by typically two months. That's significant, because with UCS, you can just drop a rack in there and you're done. The impact is felt throughout the entire organization, because they're waiting for the [Big Data analytics] results."*

FIGURE 3

Time-to-Market Metrics with Cisco UCS

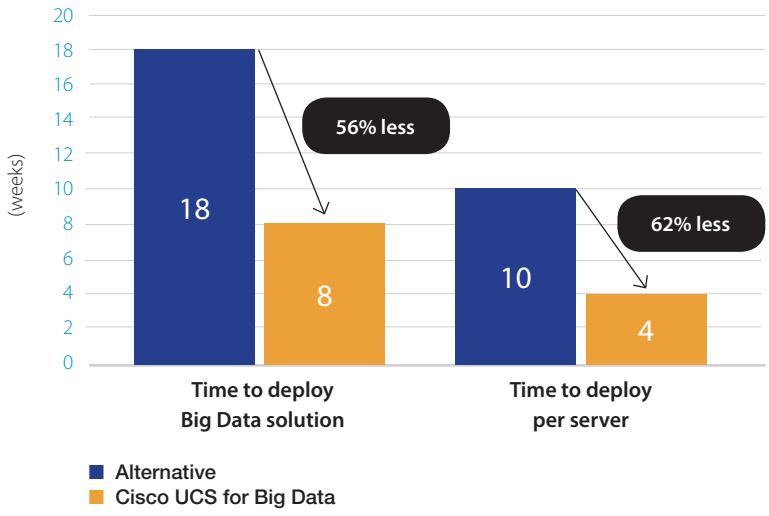

Source: IDC, 2015

The speed of deployment and scalability of Cisco UCS can translate to improved operational and business results; in other words, Cisco UCS can amplify the benefits that these organizations otherwise achieve with Big Data. Getting Big Data operations up and running in less than half the time it would take with an alternative hardware platform gives these organizations two additional months (18 weeks compared with 8 weeks) to benefit from enhanced productivity, operational benefits, and higher revenue they attribute to Big Data. Meanwhile, faster server provisioning enables their IT teams to move at the speed their businesses require. In addition, faster server deployment means that these organizations are continually generating value through Big Data faster with Cisco UCS than they otherwise could. Several organizations described these benefits:

- **Impact on business and operations.** *"Getting servers provisioned faster with Cisco UCS means that we can better respond to performance needs and support the onboarding of additional clients more rapidly . . . . My business counterparts are out selling, and they tell people they can get them going sooner . . . the scaling with Cisco UCS is about serving, getting clients onboard."*
- **Quality of services.** *"The main driver for us in using Cisco UCS for Big Data was operational efficiency to be gained around time to provision the resources . . . . We're taking advantage of features of Cisco UCS to define standard configurations, so that we are confident in our initial configurations so that when something gets deployed, it's 100% accurate the first time it goes out."*

"We have 300 data scientists who are 50% more productive since we deployed Cisco UCS for Big Data. I'd say about half of this is due to UCS, because if we scale faster and its performance is high, then it improves their productivity."

Performance with Cisco UCS for Big Data

Interviewed organizations concluded that Cisco UCS offered strong performance for their Big Data workloads based on their experience and tests conducted in association with their choice of Cisco UCS. These organizations benefit from the compute power of Cisco UCS and its integration of network and compute resources, as well as their ability to manage workloads efficiently through policy and automation.

With Cisco UCS supporting Big Data operations, organizations told IDC that they are benefiting from the timeliness and quality of analytics deliverables. In particular, teams such as data scientists and application developers are more productive when they have consistent, relevant, timely, and

high-quality data. On average, interviewed organizations reported that their data scientist and data analyst teams are 26% more productive since deploying Cisco UCS for Big Data and that their application development teams are 15% more productive. Much of this productivity gain is attributable to the Big Data solutions running on Cisco UCS, but interviewees also attributed value to the performance, reliability, and scalability of their Cisco UCS for Big Data platforms (see Figure 4).

Interviewed organizations provided a number of examples of the impact the performance of Cisco UCS has on these individuals:

- *We have 300 data scientists who are 50% more productive since we deployed Cisco UCS for Big Data. I'd say about half of this is due to UCS, because if we scale faster and its performance is high, then it improves their productivity."*
- *"80 data scientists save time with Cisco UCS by getting stuff faster, at least a number of days faster. They save a couple of hours per week each."*

FIGURE 4

Productivity Gains Among Data Scientists and Data Analysts

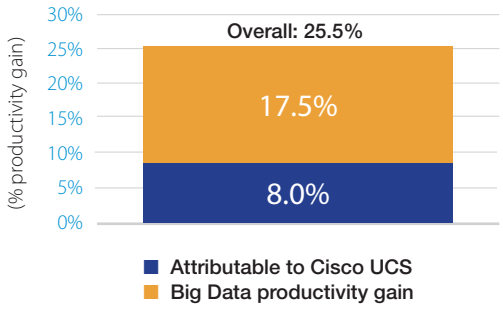

Source: IDC, 2015

Cost of Operations and Reliability

Interviewed organizations reported that in addition to serving as a high-performing and scalable platform for Big Data solutions, Cisco UCS is a cost-effective and efficient platform. In particular, organizations need less time to manage their Cisco UCS environments than alternative solutions they considered, thanks to UCS management software, while the ability to push through software updates and patches automatically also saves time. They also benefit from efficiencies that result from service profiles that UCS supports for Big Data workloads and time savings from minimizing the staff time required to deploy and scale their Big Data environments. One organization said: *"The UCS management piece of Cisco UCS, where we can manage an entire stack as opposed to having a network guy take care of network switches and a server guy who handles servers, is really important. If we were going to do the network separately, we'd probably need to hire 2–3 extra people."*

Interviewed organizations also described their Cisco UCS environments as cost effective compared with other Big Data platforms they considered. In particular, organizations reported benefiting from the fabric-based architecture of Cisco UCS that reduces the amount of

networking equipment required. One organization explained: *"The cost would have been greater with [the competing solution we considered] because of fabric interconnects with Cisco UCS. We would not have been able to connect servers directly to the fabric interconnect, and this would have meant buying additional switches for tens of thousands of dollars."*

IDC calculates that as a result of staff efficiencies in deploying, managing, and maintaining Cisco UCS environments for Big Data solutions, as well as lower infrastructure and ongoing operating expense costs, the three-year cost of operating Cisco UCS environments will be 32% lower than the three-year cost of operating the alternative hardware platforms considered (see Figure 5).

FIGURE 5

Three-Year Cost of Operations: Cisco UCS Versus Alternative

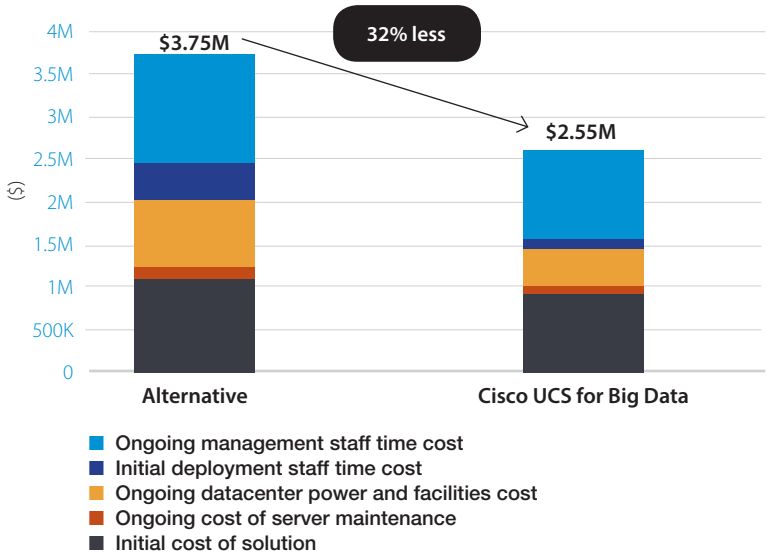

Source: IDC, 2015

Interviewed organizations also credited Cisco UCS with providing a reliable base for their Big Data operations. They reported they are losing just over one-half hour of productive time on a per-user basis to unplanned downtime with Cisco UCS, meaning that they are experiencing 65% less unplanned downtime with Cisco UCS for Big Data than they would expect running the same Big Data workloads with an alternative platform.

ROI Analysis

IDC interviewed multiple organizations running Big Data solutions on Cisco UCS Integrated Infrastructure and recorded their results to inform this study's analysis. IDC used the following three-step method for conducting the ROI analysis:

1. Gathered quantitative benefit information during the interviews using a before-and-after assessment. In this study, the benefits included staff productivity gains, increased revenue, and infrastructure-related cost reductions.

2. Created a complete investment (three-year total cost analysis) profile based on the interviews. Investments go beyond the initial and annual costs of using Cisco UCS as a platform for running Big Data solutions and can include additional costs, such as migrations, planning, consulting, configuration or maintenance, and staff or user training.

3. Calculated the ROI and payback period. IDC conducted a depreciated cash flow analysis of the benefits and investments for these organizations' use of Cisco UCS for Big Data over a three-year period. ROI is the ratio of the net present value (NPV) and the discounted investment. The payback period is the point at which cumulative benefits equal the initial investment.

Table 3 presents IDC's analysis of the average discounted benefits, discounted investment, and return on investment for the Cisco UCS for Big Data customers interviewed for this study. IDC projects that these organizations will spend a discounted average of $2.44 million ($17,244 per Cisco UCS server) and earn discounted business benefits worth $11.35 million over three years ($80,284 per Cisco UCS server). This means that they will achieve an average three-year ROI of 366% and break even on their investment in Cisco UCS for Big Data in an average of just under five months.

TABLE 3

Three-Year ROI Analysis

	Per Organization	Per Cisco UCS Server
Benefit (discounted)	$11.35 million	$80,284
Investment (discounted)	$2.44 million	$17,244
Net present value (NPV)	$8.91 million	$63,040
Return on investment (ROI)	366%	366%
Payback period	4.9 months	4.9 months
Discount rate	12%	12%

Source: IDC, 2015

Challenges and Opportunities

Leaders in many organizations believe they have what it takes to harness the power of Big Data for improving data-driven decision making; yet they are missing the competency or maturity to address the range of technology, staffing, process, and data requirements involved. With the opportunity to unlock the value of Big Data comes the need to navigate expanding technology alternatives.

Cisco has an opportunity to play a continued critical role in enabling deployment of an ever-increasing number of Big Data solutions. As the hype in the Big Data market has subsided, real-world issues about Big Data solution performance, scalability, security, and support have surfaced for organizations that were previously in the early stages of experimenting with some of the new Big Data management and analytics technologies. Cisco brings to the market a robust solution that can address a new set of customer needs and requirements.

At the same time, Cisco faces competition in the Big Data market and must continue to demonstrate the value of its technology and customer support processes.

As Big Data workloads start demanding more and more from customers' infrastructures, Cisco is in a unique position — thanks to its historical networking expertise — to shift the focus from ever more and faster processing power to intelligently designed and managed networks. Significant benefits in speed, efficiency, utilization, and TCO can be achieved with this approach. Cisco has demonstrated itself to be ahead of the game with its response to Hadoop 2.0 in this respect.

This also puts a burden on the company to continue staying abreast of new software developments in the Big Data market. Many of these are open source solutions, however, giving Cisco good visibility into what is being developed for what purpose. Indeed, the company might consider joining the Apache Software Foundation to start contributing to this open source community with its unique network perspective.

Summary And Conclusion

The competitive advantages and value of Big Data operations are becoming more widely acknowledged, shifting the focus at many organizations from whether and how to do Big Data to maximizing its value. With Big Data workloads requiring more from IT infrastructures and lines of business demanding higher quality insights in less time, choosing the right infrastructure platform for Big Data operations represents a core component of maximizing value. This IDC study considered the experiences of organizations using the Cisco UCS

Integrated Infrastructure solution as a platform for their Big Data operations and found that their Cisco UCS infrastructures contributed to the strong value they are achieving with Big Data operations through scalability, performance, time to market, and cost-effectiveness. As a result, these organizations attributed business benefits to Cisco UCS for Big Data such as earlier recognition of revenue and employee productivity gains while also maintaining efficient and cost-effective Big Data operations.

Appendix

IDC's standard ROI methodology was utilized for this project. This methodology is based on gathering data from current users of Cisco UCS as a platform for Big Data solutions as the foundation for the model. Based on these interviews, IDC performs a three-step process to calculate the ROI and payback period:

- Measure the savings from reduced IT costs (staff, hardware, software, maintenance, and IT support), increased user productivity, and improved revenue over the term of the deployment.
- Ascertain the investment made in deploying the solution and the associated migration, training, and support costs.
- Project the costs and savings over a three-year period and calculate the ROI and payback period for the deployed solution.

IDC bases the payback period and ROI calculations on a number of assumptions, which are summarized as follows:

- Time values are multiplied by burdened salary (salary + 28% for benefits and overhead) to quantify efficiency and manager productivity savings.
- Downtime values are a product of the number of hours of downtime multiplied by the number of users affected.
- The impact of unplanned downtime is quantified in terms of impaired end-user productivity and lost revenue.
- Lost productivity is a product of downtime multiplied by burdened salary.
- Lost revenue is a product of downtime multiplied by the average revenue generated per hour.

- The net present value of the three-year savings is calculated by subtracting the amount that would have been realized by investing the original sum in an instrument yielding a 12% return to allow for the missed opportunity cost. This accounts for both the assumed cost of money and the assumed rate of return.

Because every hour of downtime does not equate to a lost hour of productivity or revenue generation, IDC attributes only a fraction of the result to savings. As part of our assessment, we asked each company what fraction of downtime hours to use in calculating productivity savings and the reduction in lost revenue. IDC then taxes the revenue at that rate.

Further, because IT solutions require a deployment period, the full benefits of the solution are not available during deployment. To capture this reality, IDC prorates the benefits on a monthly basis and then subtracts the deployment time from the first-year savings.

Note: All numbers in this document may not be exact due to rounding.

IDC Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

IDC Analyze the Future