



White Paper

Mission-Critical Infrastructure for the Data-Driven Enterprise

Sponsored by: Hewlett Packard Enterprise and Intel Corporation

Ashish Nadkarni
November 2017

Peter Rutten

EXECUTIVE SUMMARY

The role played by modern analytics environments at data-driven enterprises is clear. By leveraging Big Data and analytics (BDA), lines of business can analyze massive and growing amounts of data from the digital core to the intelligent edge. BDA helps modern enterprises compete more effectively in the digital economy via the analysis of data coming from core enterprise applications and emerging initiatives such as the internet of things, robotics, next-generation security, and next-generation supply chain automation.

IDC predicts that IT spend on Big Data and analytics will reach \$59 billion by 2020. Much of this spend is going toward modernizing the application landscape and, therefore, the IT infrastructure that hosts these applications. From a data perspective, a move to a future-ready real-time enterprise includes the converging of business-centric transaction processing and data-centric analytics systems. To speed up analytics, businesses are deploying in-memory and memory-centric databases, infusing data analytics platforms with high-performance technologies, and using a highly available and secure conduit for data movement between the various application tiers.

The new HPE Superdome Flex server is designed to host data-intensive and memory-centric enterprise workloads with evolving demands. Specific examples include in-memory database platforms such as SAP HANA, Microsoft SQL Server with in-memory capabilities, and Oracle Database In-Memory, as well as in-memory high-performance computing (HPC) applications. When deployed in mission-critical environments, these applications benefit greatly from the near-linear compute, memory, and I/O scalability; extreme availability; and simplified management capabilities of Superdome Flex.

IDC believes that Superdome Flex sets a high standard for mission-critical servers for data-driven enterprises and is worthy of consideration by firms embarking on a journey to modernize their applications and infrastructure and, crucially, to unlock the value of their data in a timely manner.

SITUATION OVERVIEW

Modern analytics environments provide the crucial underpinning for firms transforming themselves into data-driven enterprises. Big Data and analytics enable modern enterprises to compete more effectively in the digital economy via analysis of data coming from core enterprise applications and emerging initiatives such as the internet of things, robotics, next-generation security, and next-generation supply chain automation. As firms seek to create and deliver digital offerings and experiences, insights obtained from data are paramount in making key business decisions.

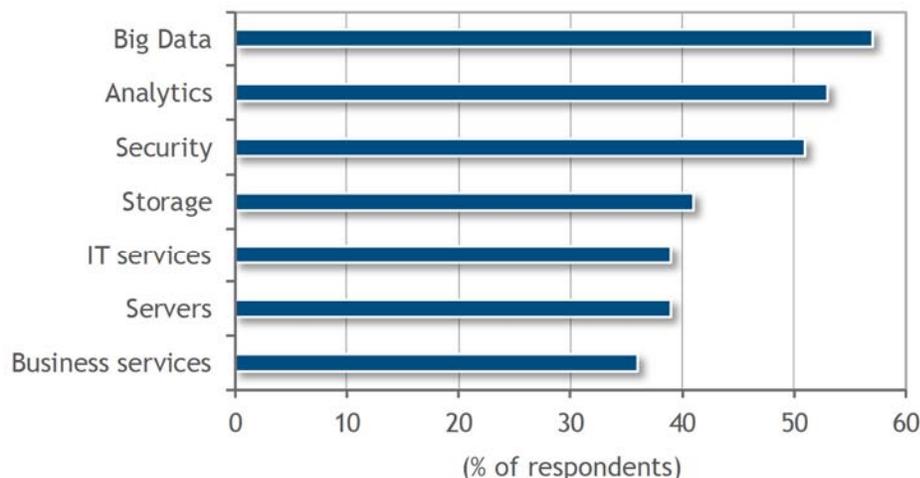
Data is the new basis of competitive advantage. Big Data and analytics have made their way to the top of the technology investments list and are areas of increased spend (see Figures 1-3). IDC finds that businesses are realizing that accelerating analytics and, ultimately, unlocking the value of data in real time (i.e., real-time Big Data and analytics) are crucial to their ability to lead in the digital economy (see *IDC Survey Spotlight: Big Data and Analytics Go Mainstream*, IDC #US42388517, March 2017). Furthermore:

- As data continues to pervade organizations at all levels at an ever-increasing pace, organizations are challenged to handle the volume, velocity, and veracity of data as leadership strives to derive value from the data and drive business impact in a real-time fashion.
- Generating data intelligence requires the analysis of vast quantities of diverse data, either structured or unstructured, generated by humans or by machines, to uncover patterns and pursue breakthrough ideas.

FIGURE 1

Big Data and Analytics Software Most Likely Areas for Increased Spending

Q. For each of the Big Data and analytics technologies listed, do you expect your organization's 2017 spending to increase, decrease, or remain the same compared with 2016?



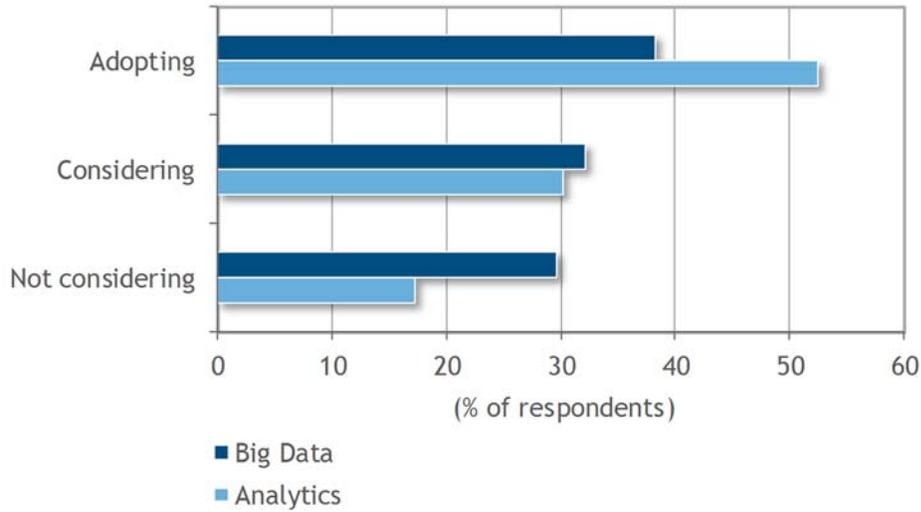
n = 3,602

Source: IDC's *Industry IT and Communications Survey*, April 2017

FIGURE 2

Big Data and Analytics Adoption

Q. At what stage is your organization today in the deployment of each of the following?



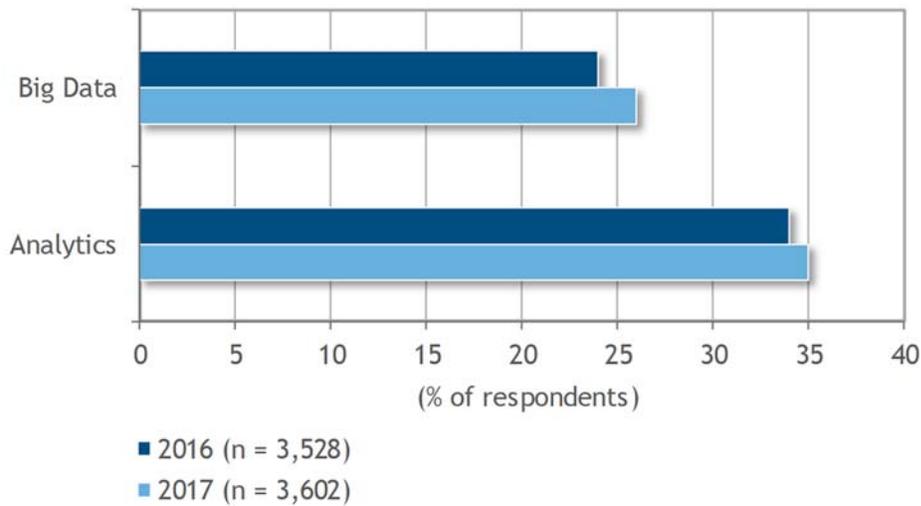
n = 3,602

Source: IDC's *Industry IT and Communications Survey*, April 2017

FIGURE 3

Big Data and Analytics Are Increasing in Priority

Q. Thinking about your organization's top priorities in the coming year, please rank the following initiatives in order of importance from 1 to 6, where 1 is most important and 6 is least important to your organization. (Big Data, analytics)



Base = percentage of respondents who ranked Big Data and analytics 1 or 2

Source: IDC's *Industry IT and Communications Survey*, 2016 and April 2017

Unlocking the intelligence from data in real time requires a modern application and data management environment. The IT infrastructure that hosts these applications and data management platforms serves as a critical foundation layer. The move to a real-time enterprise includes:

- Converging business-centric transaction processing and data-centric analytics systems to increase the quality and timeliness of insight (i.e., systems of record, engagement, and insight)
- Deploying in-memory databases for low-latency response times as part of the application environment
- Infusing data analytics platforms with high-performance technologies to optimize application performance for large data sets
- Using a highly available and secure conduit for data movement between the various application tiers
- Implementing an appropriate data persistence tier that can support the storing, securing, and fast access of rapidly changing data sets

The IDC research mentioned previously shows that Big Data and analytics initiatives will continue to become more mainstream across many different applications. Using BDA to optimize business processes was the most popular choice among respondents, followed closely by helping improve security. Customer-facing activities and product life-cycle activities were the next most popular uses for BDA within organizations. 84% of respondents have some BDA initiatives implemented within their organizations. These findings support the notion that BDA has hit mainstream adoption and will continue to grow strongly.

IDC believes that to support business agility needs as well as to manage the resources and skills shortage, more and more businesses will choose to deploy enterprise infrastructure platforms that are designed specifically for Big Data and analytics.

The Role of Mission-Critical Platforms for Modern Data-Centric Applications

Firms are increasingly deploying x86-based mission-critical platforms that scale up to optimize their BDA-centric application and IT infrastructure transformation. They are favored over commodity rack-optimized servers that cannot deliver performance at scale, especially for analytics workloads that require a high in-memory footprint and low-latency multicore performance. Furthermore, commodity servers have contributed to "server sprawl" in the datacenter over the years, leading to complexity, high software license costs, and growing maintenance and power and cooling costs.

In addition, some standard x86 servers have been evolving to provide high availability and meet the service levels demanded by mission-critical workloads for the data-driven enterprise by delivering extreme reliability, availability, and serviceability (RAS), even reaching the highest levels in IDC's High Availability framework.

IDC classifies servers in four levels of availability, with the highest level being availability level 4 (AL4). At this level, the combination of multiple hardware and software components allows a near-instantaneous failover to alternate hardware/software resources so that business processing continues as before without interruption. In 2016, IDC noted that AL4 servers were becoming increasingly suitable for the modern datacenter with capabilities such as mobile apps, cloud, APIs, open source software, and next-generation application development. In 2017, IDC noted that platforms with these capabilities for digital transformation have joined the ranks of the fault-tolerant category. This trend is in sync with the increasing desire for zero downtime in today's always-on world.

Next-generation mission-critical platforms such as HPE Integrity Superdome X are designed to provide "Unix on RISC"-like RAS at a system level, which can be augmented further by using clustering technologies; HPE is a key player in the AL4 market. Superdome X is now included in the AL4 market, and as HPE Superdome Flex inherits the Superdome X RAS framework, we can expect it to be classified at this level.

HPE Superdome Flex – Combining the Best of Superdome X and MC990 X

HPE introduced Superdome X to address the market trend toward standardizing on x86 architectures for mission-critical workloads, delivering a level of reliability previously unseen on standard platforms. Superdome X runs critical business processing and decision support workloads, such as Oracle, SAP (including SAP HANA), and SQL Server stacks on Linux and Windows.

With the SGI acquisition, HPE gained access to a portfolio of high-performance computing systems such as HPE MC990 X (renaming of SGI UV 300). MC990 X is based on a modular architecture that is optimized for high-performance low-latency workloads such as in-memory databases. MC990 X supports 4-32 sockets, up to 48TB shared memory, and up to 96 PCIe-based I/O cards in a single system.

Superdome Flex is born out of HPE's efforts to combine the best of Superdome X and MC990 X to deliver a modular, standards-based mission-critical system with maximum flexibility, performance, and reliability.

Noteworthy Characteristics of HPE Superdome Flex

Superdome Flex features a next-generation multisoocket, multicore x86 architecture and is built with Memory-Driven Computing principles derived from HPE's The Machine research project.

Unmatched Scale and Flexibility

Designed to address the needs of environments of all sizes, Superdome Flex has a unique modular design that enables firms to start small and scale up seamlessly as their needs grow. Utilizing a 5U 4-socket chassis "building block," the platform scales from 4 sockets to 32 sockets and from 768GB to 48TB of shared memory in a single system. The chassis is connected via a high-bandwidth, ultra-low-latency fabric via cabling.

Unbound I/O

When fully configured, Superdome Flex supports 128 Gen3 PCIe card slots that can be used for external storage connectivity, hardware accelerators, and other peripherals. Along with the compute capabilities, a highly scalable I/O subsystem enables the deployment of HPC software (which often requires high IOPS and low-latency bandwidth access to storage or accelerator cards).

Extreme Availability

Superdome Flex features many RAS capabilities not present in other x86 servers – a key differentiator for the platform. Some of those capabilities are:

- **Firmware First:** This approach ensures error containment at the firmware level, including memory errors, before any interruption can occur at the operating system (OS) layer. Firmware First covers correctable errors and uncorrectable errors and gives firmware the ability to collect error data and diagnose faults even when the system processors have limited functionality.
- **Self-healing capabilities:** When faults do occur, Superdome Flex provides several mechanisms to avoid unplanned downtime, including disabling failed or failing components during boot and attempting recovery on failed or failing components during runtime.

- **Processor RAS:** Superdome Flex servers use the latest-generation Intel Xeon Scalable processors. These processors include extensive capabilities for detecting, correcting, and reporting hard and soft errors. Because these RAS capabilities require firmware support from the platform, they are often not supported in other industry-standard servers. Superdome Flex implements the RAS functionality provided in Xeon Scalable series processors, including corrupt data containment, PCIe live error containment, poison error containment, processor interconnect fault resiliency, and advanced MCA recovery.
- **Memory RAS:** Superdome Flex servers use several technologies for enhancing the reliability of memory, including proactive memory scrubbing, Advanced Double Device Data Correction (ADDDC), or Single Device Data Correction (SDDC). Although ADDDC is based upon an Intel Xeon E7 processor feature, HPE enhanced it with specific firmware and hardware algorithms to substantially reduce memory outage rates over other x86 offerings.
- **Platform RAS:** Superdome Flex uses a fabric interconnect scheme featuring adaptive routing capabilities. The system routes traffic down the optimal latency path for performance and provides the ability to route traffic around failing or failed links in the fabric and while the system is running.
- **Application-level RAS:** Superdome Flex supports Serviceguard for Linux to enable software failover.

Simplified User Experience

Superdome Flex provides a simplified management experience by supporting HPE-specific tools such as HPE OneView, Insight Remote Support, and Proactive Care, as well as the open source Redfish API and OpenStack.

Target Use Cases and Workloads

Superdome Flex is designed for mission-critical workloads, in-memory databases and real-time analytics, and in-memory high-performance computing. Like its predecessors, it continues to target SAP HANA, Oracle 12c, Microsoft SQL Server, in-memory high-performance computing, and Unix-to-Linux migration.

SAP HANA

SAP has made its in-memory database, SAP HANA, the foundation of the entire environment for combined analytical and transactional processing under SAP S/4HANA. Superdome Flex's abundant memory and modular architecture make it particularly optimized for SAP HANA environments. Customers can start with a 4-socket scale-up system and add 4-socket chassis as workload needs increase. Also, the systems can be readily repurposed between scale-up and scale-out configurations as workload requirements change across analytics, SAP BW/4HANA, or SAP S/4HANA, adding investment protection.

Oracle 12c

Oracle continues to innovate on its core data management product. With 12c, it has raised the bar higher by adding in-memory analytics, pluggable databases, in-database archiving, and support for multiple indexes on a single column. Oracle 12c can be configured as a scale-up database or scale-out using clustering via Oracle RAC. By deploying 12c as a scale-up database on Superdome Flex, businesses can increase their database performance, add in-memory options for real-time workloads, and reduce TCO by saving on licensing costs.

Microsoft SQL Server

Like SAP and Oracle, Microsoft continues to innovate, raising the bar on its venerable SQL Server database. Recent features include in-memory OLTP, always-on data encryption, and support for Hadoop File System, Stretch Databases, and most notably, Linux. Superdome Flex is ideal for midsize to large SQL Server on bare-metal or virtualized server deployments. It is also suitable for database consolidation and migration initiatives, where the target database is SQL Server, and for cases where customers need reliability levels that they can't achieve with other industry-standard servers for their critical SQL workloads.

In-Memory High-Performance Computing

This use case is notably one that Superdome Flex inherits from the MC990 X and in which SGI had strong leadership. Superdome Flex equips scientific, engineering, and other technical computing environments with the ability to solve complex, data-intensive problems holistically at extreme scalability with "single-system simplicity." These types of problems are often challenging to distribute across multiple nodes in an HPC cluster and benefit from "fat" nodes (more processors and memory). They include CAE, genomics, fraud detection and prevention, and large data visualization, among others.

Unix-to-Linux Migration

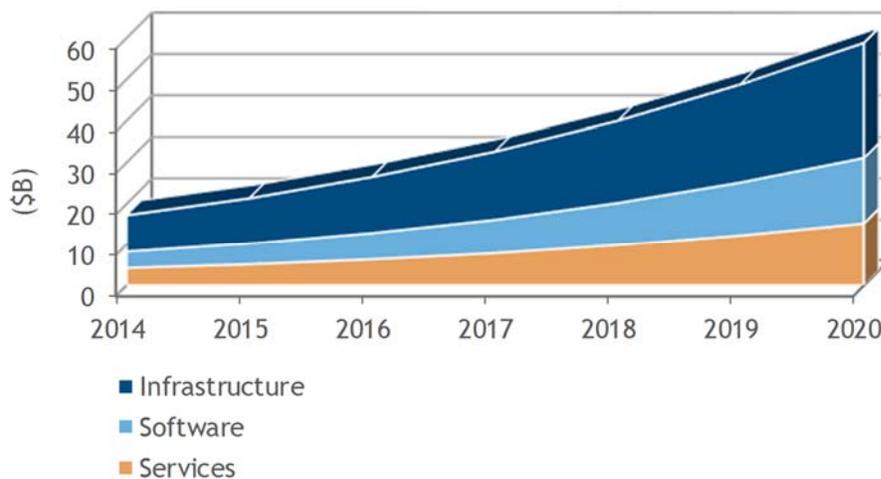
Like its predecessor, Superdome Flex is ideal for firms that want to standardize on x86-based compute infrastructure but do not want to compromise on performance, or RAS. With support for standard operating environments and virtualization technologies, firms get a wide set of options for migrating their mission-critical databases and workloads from Unix systems.

CHALLENGES/OPPORTUNITIES FOR HPE

IDC believes that firms are converging their systems of record, engagement, and insight as they advance on their journey to become data-driven enterprises. As a part of this journey, many firms are also standardizing on x86-based infrastructure, even for very data-intensive, mission-critical workloads. This is driving healthy growth for the Big Data and analytics infrastructure market (see Figure 4). Businesses are only going to spend more money on critical infrastructure platforms that enable them to accelerate their data-driven journey.

FIGURE 4

Worldwide Big Data Technology and Services Revenue by Segment, 2014-2020



Notes:

Infrastructure includes compute, storage, networking, and other infrastructure.

Revenue for compute, storage, and networking now includes cloud infrastructure and HPDA.

Infrastructure includes hardware and associated software such as operating environments, hypervisors, and DCIM.

Other infrastructure includes security and datacenter infrastructure software and hardware.

Source: IDC's *Worldwide Big Data Technology and Services Forecast, 2016-2020* (IDC #US40803116, December 2016)

With so much of a focus on scale-out architectures, it is easy to take for granted the importance of mission-critical scale-up platforms in data-driven enterprises. Similarly, it is also easy to overlook the fact that such platforms can be based on x86 architectures. IDC believes that a growing appetite for scale-up x86-based platforms such as Superdome Flex will continue to command traction in firms that require:

- Scale-up multisoocket design for high-performance scaling
- Security, availability, and reliability for mission-critical deployments
- Flexible and modular design for opex-friendly deployments
- Optimizations for in-memory databases and real-time analytics applications
- Support for an open standards-based and cloud-ready design for deploying hybrid IT

Superdome Flex is a powerful scale-up x86 platform that combines the best of HPE mission-critical reliability and SGI's scalable technology. It has been optimized for high-end performance at scale, in-memory databases, and a range of high-availability features throughout the platform – both hardware and software. It can handle the most demanding workloads quickly and without interruption. Because of its scale-up architecture, Superdome Flex also provides TCO efficiencies that, after a decade of x86 server sprawl and soaring opex in the datacenter, are in high demand.

For HPE, the opportunity lies in providing all the elements of a modern infrastructure environment in which fault tolerance truly matters. It is about positioning Superdome Flex as the platform that:

- Is flexible and powerful enough to handle the massive and growing amounts of data moving through a modern business
- Provides the ability to analyze data from the digital core to the intelligent edge in real time with an optimized in-memory design
- Is modular and cloud ready and is the right fit for any business of any size that is pursuing a traditional, private cloud or hybrid IT design

Enterprises are also embracing a world in which app developers would want a vibrant open source ecosystem on which they can develop complex, stateful apps that depend on the hardware to maintain their state, sometimes in multiple stages. Stateful apps expect the hardware to not fail, and stateful apps in many industries may have compliance requirements that mean they cannot fail.

Here, HPE should also enlist the developer community and ensure that Superdome Flex stays open and remains developer friendly. HPE should also make sure that its systems are capable of sustaining state without performance downgrades through superior compute, fabric, and storage components.

CONCLUSION

While traditionally mission-critical systems have represented a smaller part of the server market, they are well poised to grow in new areas as next-generation data analytics and in-memory databases and the expansion into the HPC space increase the market demand for this type of platform. With Superdome Flex, HPE has reaffirmed its commitment to delivering a high-end x86-based system for mission-critical workloads that run on standard operating environments. Superdome Flex differentiates itself from most of the competition by being a modular and flexible x86-based mission-critical platform. HPE must continue to convince prospective customers of the RAS features of Superdome Flex for hybrid IT deployments. HPE should now be the vendor that changes the dynamics of this market.

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.

Global Headquarters

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

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2017 IDC. Reproduction without written permission is completely forbidden.

