White Paper

HPE's Superdome X: The Mission-Critical Scale-Up x86 Platform for SAP, Oracle, and SQL Server

Sponsored by: Hewlett Packard Enterprise and Intel

Peter Rutten
June 2016

EXECUTIVE SUMMARY

Customers of high-end, mission-critical systems are increasingly moving their large applications to x86 infrastructures. However, years of x86 "server sprawl" have caused datacenter complexity, high software license costs, and growing maintenance and power and cooling costs — not to mention that this sprawl has pushed the limits in terms of datacenter footprint. Furthermore, x86 servers have proven to be unable to deliver the kind of performance; reliability, availability, and serviceability (RAS); and scalability that mission-critical workloads in the 3rd Platform era — mobile, social, Big Data, and cloud — require.

As customers transform their datacenters to address 3rd Platform dynamics, there has been increasing demand for "Unix on RISC"-like performance, RAS, scalability, and security on x86. This, in turn, has led to increased interest in x86-based scale-up systems, with a larger number of sockets and cores to achieve a better cost-performance ratio.

This white paper investigates the current environment as these various dynamics play out and takes a closer look at the HPE Integrity Superdome X, a mission-critical platform based on Intel x86 processors, supporting Linux, Windows, and VMware, that was launched in December 2014 and significantly updated throughout 2015 and 2016, including a new product launch in June 2016. The Superdome X seems to have been designed to directly address the above-mentioned confluence of customer needs and, as such, may be the first of its kind.

The Superdome X is aimed at a specific set of mission-critical workloads: business processing, such as ERP and CRM, and real-time analytics on top of in-memory database platforms such as SAP HANA, Microsoft SQL Server, or Oracle Database. These are the kinds of 3rd Platform workloads that benefit greatly from near-linear scalability in a scale-up system with large socket counts, high RAS, a large memory footprint for in-memory processing, and raw performance. The system was initially launched as Linux only (Red Hat and SUSE), but Superdome X has since been certified with Windows and VMware.

HPE has increased the Superdome X memory from 12TB to 24TB; made the system available with the latest Intel Xeon processor E7-8800 v4 family, with core counts ranging 4–24 cores; and improved the crossbar fabric to increase bandwidth by 33%. Furthermore, HPE has introduced a nonuniform memory access (NUMA) awareness technology called Application Tuner Express (ATX) to improve linear scalability, brought in load-balancing software called System Resource Orchestrator (SRO) to ensure that mission-critical workloads always get the resources they need, and increased the number of partition sizes in its popular hard-partitioning technology. This provides customers with a flexible and easily customizable platform for running multiple workloads (such as Oracle and SAP HANA) on various operating systems (OSs) in the same enclosure.

SITUATION OVERVIEW

3rd Platform Dynamics

We are in the midst of what IDC calls the 3rd Platform of information technology (IT), which is dominated by four major trends: Big Data, cloud, mobile, and social computing. Datacenters and IT have evolved continuously to keep up with the demands of the business and support top-line and market share growth as these dynamics play out, all the while trying to maintain a balance between:

- **Aligning with the business:** Accelerating innovation, delivering real-time analytics for decision making, supporting mission-critical workloads, and enabling mobile and social
- **Managing risk:** Guaranteeing high availability (HA), disaster recovery, security, governance, and compliance
- **Reducing costs:** Working toward consolidation, virtualization, automation, integrated systems, and cloud sourcing

Reducing costs has become increasingly critical because, over the past several years, operational expenses (opex) have been growing much faster than capital expenses (capex) in enterprise datacenters, primarily because of skyrocketing maintenance costs. This has resulted in significantly less IT budget for delivering what's most important: innovation that can drive business growth in the 3rd Platform world.

The Irreversible Standardization on x86

One of the main reasons for the growth in maintenance costs has been the industry move toward distributed environments as low-cost x86 systems take on an increased share of the workloads. Before this shift occurred in the late 1990s, big, high-end Unix solutions had been the preferred choice for business-oriented workloads. However, the Unix offerings from multiple competing vendors were too different, and supporting the various Unix flavors was complicated and costly for ISVs. The Unix solutions in the market failed to unify into a larger, standardized ecosystem and diverged even further until customers began to prefer the standardized, less expensive solutions based on x86 offered by multiple vendors. The Unix market has been in slow decline ever since.

Figure 1 shows worldwide x86 versus non-x86 server shipments for 2011-2020.

FIGURE 1

Worldwide x86 Versus Non-x86 Server Shipments, 2011-2020

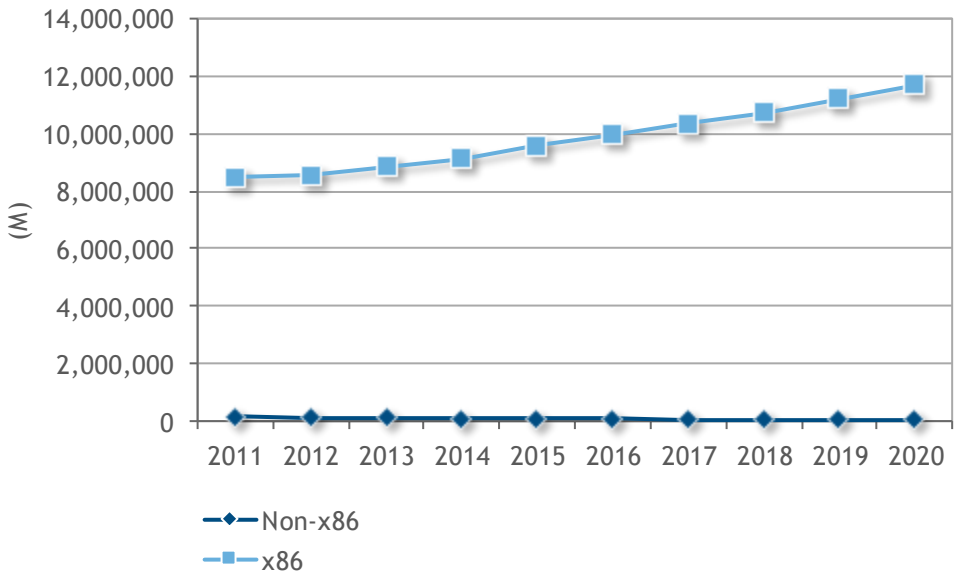

Source: IDC, 2016

The shift to x86 allowed customers to control capital expenditure with lower-cost servers and add capacity incrementally. At the same time, Linux became increasingly viable as a commercial platform, delivering the portability and flexibility that Unix could not provide. Linux has served well as a Unix alternative for customers that want a Unix-like experience on x86 systems. Over time, the evolution of high-availability features within Windows Server (load balancing, failover clustering, and live migration capabilities), recently further enhanced with Windows Server 2016, has enabled IT to start deploying mission-critical workloads on Windows as well.

The Opex Effects of Scale Out

The shift to low-capex scale-out x86 systems also had undesirable effects. The increased opex of managing the expanding x86-installed base has been taxing IT budgets, making it difficult to fund new initiatives that might have delivered revenue-generating innovations.

Figure 2 shows the worldwide annual server opex-to-capex ratio for 1996–2018.

FIGURE 2

Worldwide Annual Server Opex-to-Capex Ratio, 1996–2018

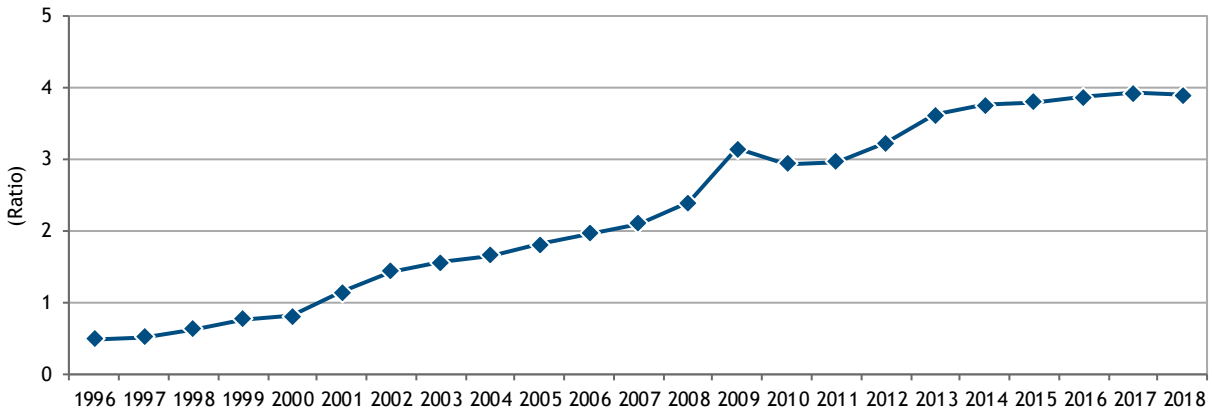

Source: IDC, 2016

This phenomenon, sometimes referred to as "x86 sprawl," has been somewhat muted by ongoing virtualization in terms of hardware proliferation, but virtualization itself has also had an impact on opex because every operating system instance in every virtual machine (VM) needs to be provisioned, life cycled, and managed; backed up; and protected with a disaster recovery strategy. IDC estimates that since the start of x86 server virtualization in 2004, VM deployments have increased almost 100-fold, necessitating a doubling of the number of server administrators worldwide over the same period. Server sprawl has also led to high power and cooling costs and caused the datacenter footprint to expand to nearly unsustainable levels (see Figure 3).

FIGURE 3

Worldwide New Server, Power and Cooling, and Management and Administration Spending, 1996-2018

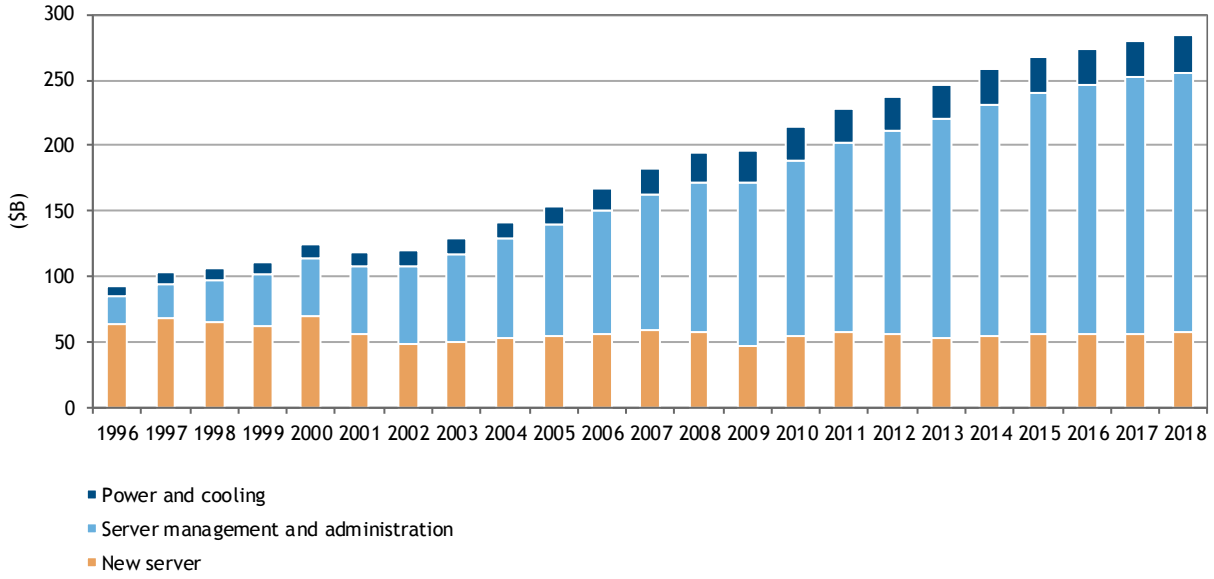

Source: IDC, 2016

Scale-out x86 systems have furthermore presented customers with performance and RAS gaps that relate directly to the migration away from powerful, scalable, and highly reliable Unix-based systems, while 3rd Platform requirements are becoming more demanding with regard to all these system attributes. As a result, these deficiencies in terms of performance and RAS have translated into an enterprise-level customer need for cost-efficient ways to manage very data-intensive, mission-critical workloads with guaranteed high RAS and scalability on x86. To resolve these issues, customers have begun considering a range of increasingly potent scale-up solutions that vendors are offering for mission-critical workloads.

IDC MARKET ANALYSIS

The Trend Toward Scale Up

IDC is seeing customers embrace the virtues of a scale-up approach for the most demanding enterprise applications after more than a decade of scale-out strategies. Most vendors today offer scale-up products in the form of relatively cost-efficient, high-end, multicore, and multiprocessor servers with four or even eight sockets, very high throughput, and high RAS. Such systems deliver excellent performance per system watt per dollar for large, mission-critical applications. They can perform various critical workloads, such as online transaction processing, database, data analytics, business processing, and ERP. The HPE ProLiant DL580 line, which was among the first to include systems built around Intel's scale-up-focused Xeon E7 chips, falls into this category.

Scale-up servers tackle server sprawl quite efficiently. Thanks to their multiple processors and cores, scale-up servers can consolidate onto a single physical host many workloads and virtual machines that used to be hosted across multiple physical servers. Scale up also has significant application performance and resource utilization benefits as well as improved RAS attributes with the right configuration and software.

Furthermore, IDC research has determined that IT can achieve more than 35% in savings from the economies of scale of a scale-up approach, with the savings contributors being (ranked in order) increased IT staff productivity, increased resource utilization (CPU, memory, and network), reduced software licensing costs, reduced energy consumption, reduced IT infrastructure costs, reduced end-user productivity losses, and improved performance. Workloads that specifically benefit from a scale-up approach are SAP, Oracle, SQL Server, and custom application stacks.

The Reversal of the High-Availability Downgrade on x86

As mentioned previously, a direct effect of the collective jump from Unix onto x86 has been the resulting downgrade in terms of built-in high-availability features. Unix servers were traditionally providing much higher levels of availability than x86, albeit at a price premium. As datacenters switched to x86, IT lived dangerously for a while, especially if running earlier versions of Windows and Linux, but then these operating systems began to ship with built-in HA extensions and third-party software for workload balancing, and clustering began to fill the availability gap. Virtualization, too, has had an effect on availability because switching to another virtual machine is easier than firing up a physical server.

Availability demands are becoming more stringent and urgent, however. They are increasingly mandated by governments or are a condition for participating in a supply chain, including severe monetary penalties for downtime. The business impact of downtime and/or data loss has also become higher, and few firms today think they can afford their customer-facing, plant-managing, or transaction processing system to be down for even half a minute — let alone longer. The same holds true for such mission-critical workloads as database and data analysis.

The resulting need to combine scale up, high performance, and high availability within a standardized operating environment is being addressed by multiple vendors and from multiple angles: x86 scale-up servers, for example, especially those based on the Xeon E7 chip, are being designed with improved availability features, and traditional high-end, mainframe-like fault-tolerant servers are being reinvented to run Linux on x86, including HPE with its Integrity Superdome systems and Fujitsu and Unisys with their high-end systems. In addition, other vendors, such as IBM with POWER8 and z Systems as well as Oracle, are innovating toward Linux on non-x86 hardware.

In other words, with downtime becoming universally unacceptable in the always-on business world, the mission criticality of workloads dramatically increasing, scalability becoming a top priority to provide the datacenter with fluid flexibility, and security climbing to be among the top 3 concerns, high RAS, scalability, and x86 are in the process of — finally — being tightly joined. Some vendors are going as far as imbuing their high-end x86 products with a variety of RAS attributes that start at the microchip level, are embedded throughout the hardware, and spread into the firmware, the OS, and all the way through the application stack — essentially establishing a pyramid of availability. HPE's Superdome X may be the prime example of this coming together.

Extreme Scalable Performance and In-Memory

As to scalability, vendors are attempting to grow the performance of their high-end systems in as much a linear fashion as possible, with every doubling of sockets. While 8 sockets have become the norm, HPE's Superdome X scales to 16 sockets and, as such, HPE claims to hold several performance world records for x86 servers with this platform.

Further affecting system requirements in this new world of Big Data and real-time data analysis is the need to bring the data that is being processed and stored as close as possible to the data that is being analyzed. Various vendors have boosted their systems' in-memory capabilities for this purpose or have added flash storage close to the RAM to facilitate a very fast interplay between stored data and in-memory data. IBM's POWER8, Oracle's M-Series, and HPE's Superdome X, all deliver extensive capabilities in this area.

This white paper takes a deep dive into HPE's scale-up system for tier 1 mission-critical workloads, the HPE Superdome X Server, an x86-based platform that represents an evolutionary branch of HPE's big and powerful HP-UX-based Superdome. While IDC does not expect that x86-based scale-up solutions will, as of yet, be able to fully match the largest RISC- and EPIC-based systems in terms of scalability and availability, HPE is coming close with Superdome X, which scales to 16 sockets and has a slew of additional RAS features built in at every level — from silicon to firmware to application stack to services.

HPE SUPERDOME X

To address the market trend toward standardizing on x86 architectures, even for mission-critical workloads, HP, now HPE, introduced the Superdome X server in December 2014. To develop Superdome X, HPE engineered a large number of Unix server capabilities around the Intel Xeon E7 processor architecture. In addition, HPE worked closely with Intel to fine-tune the chips for improved reliability as well as with the operating system vendors and the Linux community to enable Linux to take advantage of the platform's RAS features. HPE optimized the firmware; brought forward the unique hard partitioning technology, HPE nPartitions (HPE nPars), to Superdome X to make it flexible and reliable; and included support with its high-availability clustering offering, HPE Serviceguard for Linux.

Superdome X combines scale to handle the largest workloads with 16 sockets and — since its latest update — 24TB of memory, including extensive in-memory capabilities and the RAS needed for mission-critical operations, thanks to its hardware partitions. It runs critical business processing and decision support workloads, such as Oracle, SAP, and SQL Server stacks. As such, the Superdome X, which can be referred to as a "system of record" (in essence, the server where the data resides), has been brought very close to the "systems of engagement" (where the apps live and the transactions are happening) as well as to the "systems of insight" that analyze the data for patterns to provide new business insights. IDC believes that this coming together of systems of record, engagement, and insight is a powerful trend across high-end enterprise datacenters.

New HPE Focus Areas Since the Superdome X Launch

The Superdome X became available in late 2014 as a general-purpose x86 platform for the most demanding enterprise workloads. Over time, HPE has begun to concentrate on two mission-critical workloads for enterprise customers: business processing (ERP, CRM, etc.) and real-time analytics on top of three mainstream database solutions: Oracle Database, Microsoft SQL Server, and SAP HANA.

Oracle, SQL Server, and SAP Customers

HPE is focused on helping Oracle Database customers reduce complexity and licensing costs with various attributes of the Superdome X, including the fact that multicore processors are priced as (number of cores)*(multicore factor) processors, where the multicore factor is 0.5 for x86 processors versus 1.00 for modern processors commonly used in Unix systems. Also, the cost of maintenance and support tends to be lower on a modern scale-up x86 system than on a legacy Unix system. HPE is also trying to migrate customers on Oracle RAC, a popular but costly scale-out clustering solution that delivers high availability, to a scale-up approach on the Superdome X, which — HPE believes — would deliver similar levels of availability but at lower cost.

HPE is focused on providing Microsoft SQL Server customers the scalability that they can achieve now that Windows 2016 will be supporting 12TB, a threefold increase compared with the 4TB support with Windows 2014. HPE expects to continue engaging with Microsoft on multiple levels to coordinate the optimization of the Superdome X/SQL Server combination, especially on two fronts: consolidation of multiple SQL Server databases, which is an important strategy for Microsoft and for which HPE has issued a new reference architecture, and the in-memory capability of SQL Server.

HPE is concentrating on helping SAP customers on their journey to HANA because SAP has indicated that its support for non-HANA databases will end in 2025. HPE believes that the Superdome X is a very suitable platform for an enterprise's journey to HANA, not just because of its performance, scalability, and RAS features, but especially because its HPE nPars allow for a customized step-by-step migration to the HANA platform while maintaining the existing database for as long as desirable. HPE says it aims to leverage its long-standing partnership with SAP to help SAP customers pave the way to HANA on Superdome X.

Typical Customers for Superdome X

The abovementioned focus areas have evolved from a crystallization of Superdome X customers into four distinct segments:

- Customers that are performing a migration from Unix to Linux or Windows on Superdome X (e.g., Iowa-based doors and windows manufacturer the Pella Corporation says that it has completed an Oracle Database migration from HP-UX to Superdome X)
- Customers that are deploying a mission-critical database they wish to run on scale-up x86 (e.g., Missouri-based health information technologies provider Cerner runs the database tier of its electronic health records application Millennium on a Superdome X)
- Customers of ERP solutions from various vendors, such as SAP and Infor (e.g., California-based Mag Instrument, manufacturer of the MagLite flashlight, runs Infor on a Superdome X)
- Customers that are moving to SAP HANA on Superdome X, either as an appliance in the form of the CS900, which is built around the Superdome X, or as a Tailored Data Center Integration approach (e.g., one of Singapore's largest grocery chains NTUC FairPrice runs SAP HANA on a CS900)

Superdome X System Details

The Superdome X is – in its largest configuration – built up with 8 BL920s Gen9 dual-socket server blades in an 18U enclosure and supports up to 16 processors. The system allows for efficient sizing in the range of 2–16 sockets and provides a range from 8 cores (16 threads) to 384 cores (768 threads). The scalability factor, the percentage of resources that remain usable as the system scales up, is near linear. For example, from 8 to 16 sockets, HPE assesses its scalability factor as 1.77, meaning that a system with 16 sockets will perform 1.77 times the transactions as a system with 8 sockets. With HPE's new NUMA awareness offering ATX (which is discussed in the New Product Introduction for Superdome X section), applications can now truly realize this hardware scalability factor.

Superdome X supports up to 384 DDR4 memory DIMMs (48 DIMMs per blade), with a combined capacity of 24TB (3TB per blade). I/O is delivered through 16 dual-port FlexLOMs and up to 24 dual-port mezzanine I/O cards.

The system is also available as part of a converged and fully packaged HPE appliance, the ConvergedSystem 900, in both scale-up and scale-out versions. These solutions include an HPE 3PAR StoreServ 8400 all-flash array, Linux, a higher software stack, HPE Serviceguard for Linux, and services for SAP HANA.

The Superdome X runs on standard versions of SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL) developed in collaboration with SUSE and Red Hat, respectively. It is also certified with Windows and VMware. HPE has made a point of *not* developing a specialized Linux operating environment for the system because it believes customers are asking for standard Linux distributions. However, HPE also says that it has worked closely with the x86 operating environment community and distributors, resulting in thousands of code changes that take advantage of the Superdome X architecture and RAS features.

Performance and RAS

An important component of the Superdome X that helps it achieve the extreme performance and very high RAS is its crossbar fabric, which connects the 16 sockets, 384 DIMM slots, and the I/O, providing a per-blade I/O bandwidth of about 100GBps and supplying the Superdome X with enough throughput for very large workloads.

The crossbar fabric also functions as an important reliability component for mission-critical workloads in that it facilitates link failover and end-to-end retries of an instruction set through multiple paths to ensure that a transaction is always completed. Furthermore, the crossbar fabric separates the physical partitions (called HPE nPartitions), which makes the Superdome X more scalable in terms of performance and more reliable in that partitions cannot affect each other in any way. The HPE nPars technology is essential for the system and therefore warrants a brief explanation.

HPE nPartitions are a hard partition technology that, unlike virtual partitioning, is based on electrical isolation. HPE nPars allow IT to configure each Superdome X system as a single server or as many smaller, independent servers. Each HPE nPartition has its own CPU, memory, and I/O that can be removed and added to another HPE nPartition via management software – there is no physical manipulation involved, except perhaps if a blade needs to be added.

The way this works is that a single Superdome enclosure could be set up with, for example, six HPE nPars, each with its own dedicated CPU, memory, operating environment, and I/O. The first three HPE nPars might be dedicated to mission-critical workloads, possibly a development, test, and production environment. The fourth HPE nPar could contain a number of resource-managed partitions with stacked workloads, for example, databases that are running on a single native OS. The fifth HPE nPar might run KVM, supporting a number of VMs that each runs its own OS and applications on shared resources. And the sixth HPE nPar might run a hypervisor that supports multiple VMs with different OSs.

The HPE nPars technology has various advantages. From a TCO perspective, capacity is simply expanded by adding another blade, but ISV license costs are incurred only for cores that are included in a partition. In terms of resource utilization, the ability to set up and manipulate multiple electrically isolated environments in a single enclosure provides IT with a large amount of freedom to get the most out of the system's capabilities. And from a business continuity point of view, scheduled downtime can be nearly eliminated because partitions can be serviced separately without affecting each other.

HPE has found that the HPE nPars technology is developing into a key differentiating factor for the Superdome X. A public sector company in Europe, for example, says that it has an Oracle Database on bare metal in one HPE nPar and a VMware guest in another. A bank in Africa runs Oracle in one HPE nPar, with a low core count to reduce licensing costs, and Temenos' core banking application in another. Another HPE customer, a pharmaceutical distributor, says that it is deploying SAP HANA for its wholesale applications in one HPE nPar and for the pharmacies in another. Finally, a European SAP customer that has begun the journey to HANA but has not yet made the switch to HANA applications is running SAP ERP on DB2 in one HPE nPar and SAP HANA in another.

The Superdome X's RAS features — some of which have been mentioned previously in the context of the crossbar fabric and the HPE nPars technology — unequivocally show the system's heritage from the Superdome and Integrity product lines. The system's availability — which HPE claims is at 99.999% uptime — is best described as a pyramid, starting at the bottom with the infrastructure where system components, firmware, fault-tolerant fabric, and HPE nPars have been designed for reliability and can be optimized and, if necessary, repaired online.

The blades and the enclosure have redundancy and failover of key components built in. As for the processor, HPE worked closely with Intel to get the most out of the Xeon processor E7's RAS features, including fine-tuning the error recovery process between the processor, the OS, and the applications. HPE has also developed the firmware in-house, maintaining full control over the code and the way it controls the hardware to ensure that the system handles any errors before they can affect the next layer, the operating system. This firmware has decades of mission-critical heritage built into it, bringing the full experience to x86 for the first time.

As mentioned previously, HPE has collaborated closely with OS vendors to make sure the Superdome X can take full advantage of the operating systems' availability features. Above the OS sits HPE's high-availability solution, HPE Serviceguard for Linux, as well as a disaster recovery solution, remote support, and what HPE terms *mission-critical services.* These mission-critical services consist of a range of services to ensure the system's workload availability, such as support for platform migration, workload management, partitioning, virtualization, and datacenter operations.

Added Functionalities Since the Superdome X Launch

Since the Superdome X's initial launch, HPE has introduced the following updates:

- **March 2015:**
  - Gen8 blades certified for Windows and VMware
- **September 2015:**
  - Availability of Intel Xeon processor E7 v3 family on Gen9 blades, first for Linux then also for Windows and VMware
  - Maximum memory increase from 12TB to 24TB
  - An added HPE nPar size; originally, available HPE nPar sizes were 1, 2, 4, or 8 blades; from September 2015, 3-blade nPARs were supported as well, providing additional flexibility
  - New, 18-core processor so that a 2-socket blade can range from 8 to 36 cores (based on a 4-core to 18-core processor); the lower core count is offered to enable customers to reduce software licensing costs
  - Investment protection, allowing customers to mix Gen8 and Gen9 blades in separate hard partitions of the same enclosure
- **March 2016:**
  - Gen9 blades certified for Windows and VMware

New Product Introduction for Superdome X

In June 2016, HPE announced the following new Superdome X features and functionalities:

- **Intel Xeon processor E7-8800 v4 family added to the existing hardware.** The HPE Superdome BL920s Gen9 Server Blade supports 5 new processors with 4, 10, 14, 22, or 24 cores, essentially going from the previous E7 v3 version for each processor to the E7 v4 version, except for the 14-core processor (E7-8855 v4), which is a new offering from Intel and which HPE uses to replace the previous generation's E7-4850 v3.
- **Crossbar fabric enhancements to provide more bandwidth to the Superdome X infrastructure, including layout changes and additional internal connections to the current crossbar fabric.** These changes will increase the bandwidth to the processors and to the memory from 75GBps to 100GBps per blade – a 33% increase. The new crossbar fabric will be forward- and backward compatible with existing blades and next-generation blades. Furthermore, the current crossbar fabric will be compatible with the future blades. In other words, it doesn't matter which crossbar fabric module a customer selects – it is supported both with future and past generations of blades. Customers that have a fully populated enclosure (8 blades) will enjoy the greatest benefits, especially when using the higher-performance Xeon processor E7-8800 v4 family in those blades.
- **Yet another HPE nPar size addition, namely, for 6 blades (or 12 sockets).** The current offering is a 1-blade, a 2-blade, a 3-blade (which was added in 2015), a 4-blade, and an 8-blade partition. The new 6-blade HPE nPar addition fills the space between the 4- and the 8-blade partitions, providing customers with an additional incremental step.
- **Introduction of the System Resource Orchestrator (SRO) for Superdome X, a software solution for consolidating mission-critical workloads on a single server, specifically for large, multicore environments.** SRO guarantees that certain applications will always have the system resources (cores and memory, with I/O added at a later stage) they need based on time of day, current load of the server, and other customizable policies. SRO operates at the OS level, in the context of a

single OS image, unlike a hypervisor that balances resources across multiple VMs; SRO is ideal for running multiple workloads simultaneously – for example, compilers and a database – while guaranteeing that the required resources are available at the right time.

- **Introduction of Application Tuner Express, a tool that provides NUMA awareness without requiring changes to the application.** NUMA groups microprocessors in a multiprocessing system in such a way that they optimally share nearby memory. This avoids overloading the interconnect between the processors of a large system and helps achieve near-linear performance increases as processors are added. However, for applications to benefit from NUMA, they need to be "NUMA aware," which is what ATX is designed to do. ATX helps an application avoid getting bogged down in a large system such as the Superdome X.

  When running ATX with a NUMA-aware database like Oracle, ATX will automatically make decisions to keep the memory close to the CPUs that are executing the tasks. Recent HPE tests have shown significant OLTP multitenant performance gains with ATX. With 8 processors, the gain was 24%, going from non-NUMA aware to NUMA aware with ATX with a HammerDB workload. For 16 processors, performance improved by 58% by making it NUMA aware. The more processors involved, the greater the performance impact of keeping the compute close to the memory.

  An interesting result with these tests is that an 8-processor configuration with ATX has almost the same performance as 16 processors without ATX. This means that customers can add ATX and as such almost double the workload the server can handle while reducing the core count (within a partition) and licensing costs. ATX is currently available for Linux only. HPE says that it is looking into making ATX and SRO available for Windows and VMware later in 2016, if there is customer demand.

CHALLENGES/OPPORTUNITIES

Opportunities

IDC believes that enterprise customers are increasingly required to bring their systems of record, engagement, and insight closer together under the influence of 3rd Platform dynamics: the surge of mobile, social, Big Data, and the cloud. IDC research also shows that customers will continue to migrate off legacy non-x86 systems, even for very data-intensive, mission-critical workloads, toward x86 architectures. Furthermore, our research shows a growing appetite for scale-up x86-based systems and, tied in to this, a strong need for higher RAS features than standard x86 servers can typically provide.

The Superdome X is a powerful scale-up x86 system that has been equipped to deliver very high performance, in-memory capability, strong scalability (up to 16 sockets), and a range of high-availability features throughout the system – from silicon to HPE nPars, crossbar fabric, firmware, OS, apps, cluster software, and services. It can handle the most demanding workloads fast and without interruption. Because of its scale-up infrastructure, the Superdome X also provides TCO efficiencies that, after a decade of x86 server sprawl and soaring opex in the datacenter, are in high demand.

We believe that this system, therefore, represents an opportunity for high-end enterprise customers to start delivering on the many, sometimes contradictory, demands of the 3rd Platform era with a high degree of efficiency, reliability, and flexibility. We also believe that for HPE, the system represents a strong opportunity to continue developing a successful new branch of its business line of powerful systems for mission-critical workloads – one based on not only the performance and RAS philosophies of its traditional mission-critical systems but also x86.

Challenges

HPE has delivered a high-end x86-based system for mission-critical workloads that runs on standard Linux, Windows, and VMware and competes with high-end systems from IBM and Oracle. Both competitors are aggressively marketing enterprise-class solutions with powerful in-memory capabilities and high RAS features that can run Linux, albeit not within an x86 architecture. Several smaller vendors of mainframe-class systems are carefully converting their proprietary systems to x86, nurturing their existing customers in the process while trying to expand to new markets.

The Superdome X could have an advantage in this competitive environment by being x86 based, but HPE will have to convince prospective customers of the RAS features of Superdome X because x86 is traditionally not considered a high-RAS architecture. Over the long term, a challenge may emerge from cloud providers that are enabling mission-critical applications with high RAS.

The Superdome X may also have a cannibalizing effect on HPE's Unix-based Superdome 2, and HPE may find itself in the challenging position of continuing support for Superdome 2 customers while promoting the system's x86-based platform. HPE's strategy is to offer support for both sets of platforms, and HPE's ability to provide a logical continuum for each customer that is fully in line with that customer's long-term needs could also be an enrichment rather than a complication of HPE's sales strategies.

ESSENTIAL GUIDANCE

IDC believes that the following customer environments can benefit from the Superdome X:

- Customers with large, database-oriented environments (Oracle, SAP, and SQL Server)
- Customers with mission-critical OLTP environments
- Customers with Oracle or SAP stacks
- Customers running Unix-based systems for mission-critical workloads
- Customers that are converting from scale out to scale up
- Customers that are running or are considering running SAP HANA
- Customers that want to replace their high-end x86 systems with a larger, more flexible Linux-based system

CONCLUSION

It was more or less inevitable that one of the large server vendors would develop a system that addresses a range of unfulfilled customer needs as it navigates the many converging and contradictory dynamics of the 3rd Platform transition. Such a system would have to be x86 based because – for now – that is where most enterprises seem to be going, even with their mission-critical workloads. It would have to be very scalable and deliver outstanding reliability and availability, which have been sorely lacking since the migration off Unix. The system would need to provide extreme performance to tackle today's very data-intensive workloads as well as powerful in-memory capabilities to allow for real-time transactions and data analysis at the same time because Big Data will only get bigger. And it would have to provide a compelling TCO, countering the unsustainable rise in datacenter opex that customers will continue to face for the foreseeable future. With the Superdome X, HPE may be that vendor. It will be interesting to see how the Superdome X will compete with Oracle's and IBM's or even smaller vendors' latest attempts to fill the same gap with a slightly different philosophy.

RELATED RESEARCH

This white paper draws on the following IDC research:

- IDC's Worldwide Quarterly Server Tracker, 4Q15 (March 2016)
- IDC's Worldwide Quarterly Server Forecast, 4Q15 (March 2016)
- *Worldwide Server, Power and Cooling, and Management and Administration Spending 2014-2018 Forecast* (IDC #250082, August 2014)
- *The Business Value of Scale-Up x86 Servers: Platforms for Efficient Management and IT Simplification* (IDC white paper #244590, December 2013)

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-insights-community.com
www.idc.com