

INFORMATION TECHNOLOGY INTELLIGENCE CONSULTING

Information Technology Intelligence Consulting



ITIC 2016 - 2017 Global Server Hardware, Server OS Reliability Report

October 2016

Table of Contents

Executive Summary.....3

Introduction.....5

Data & Analysis.....6

 Survey Highlights.....9

 Human Error, Security Undermine Reliability9

 Key Reliability Trends.....13

 Minimum Reliability Requirements Increase Year over Year14

 Processor Technology17

 Real World Reliability Scenarios19

 IBM Rock Solid Reliability20

 Lenovo System x Servers Best in Class Reliability.....24

 Cisco UCS Comes On Strong, HP Reliability Rebounds.....29

 Server Virtualization Reliability.....29

 Cost of Hourly Downtime Soars30

Conclusions and Recommendations.....32

Recommendations.....33

 Survey Methodology35

 Survey Demographics35

Appendices.....35

Executive Summary

IBM Power Systems Servers Most Reliable for Ninth Straight Year; Lenovo x86 Servers Deliver Highest Uptime/Availability among all Intel x86-based Systems; Cisco UCS Stays Strong; HPE Integrity Superdome Garners Strong Reliability Ratings; Intel Xeon Processor E7 v3, v4 chips incorporate advanced analytics; significantly boost reliability of x86-based servers.

Human Error and Security are biggest issues negatively impacting server hardware/server operating system reliability.

The reliability of core server hardware and server operating systems has never been more crucial. Systems, networks and applications in the 21st Century Digital Age are always-on.

Corporations conduct business around the clock and around the globe irrespective of time zones or the location of their increasingly mobile and remote workforces. Workloads, which include larger, more complex and sophisticated applications have more functionality and greater intelligence. But those added features are increasingly compute-intensive, placing more demands on the underlying server hardware. Organizations' increasing use of cloud computing and virtualization, coupled with the use of emerging new technologies like the Internet of Things (IoT), Big Data Analytics, cognitive computing and mobility make high reliability imperative for organizations' core on-premises datacenters, collocations and cloud based deployments.

The latest ITIC 2016 – 2017 Global Server Hardware and Server OS Reliability survey finds that 72% of organizations now require a minimum of “four nines” or 99.99% reliability and uptime. That is the equivalent of 52 minutes of **unplanned** per server/per annum downtime or 4.33 minutes per month. Additionally, 24% of businesses now demand even greater – 99.999% availability -- which equates to 5.25 minutes of **unplanned** annual downtime or a scant “blink and you miss it” 43.7 seconds per month for their mission critical servers and main line of business applications.

Overall, the inherent reliability of the majority of server hardware platforms, server operating systems and the underlying processor technology continues to improve year over year. However, external threats including security and human error are proliferating and can significantly undermine the overall health and stability of the corporate infrastructure, to the detriment of the entire enterprise ecosystem reliability.

Among the top survey findings:

- **IBM, Lenovo and Cisco UCS** scored highest reliability.
- **IBM's z Enterprise** mainframe class server customers reported the least amount of unplanned downtime and the highest percentage of “five nines” 99.999% uptime of any server hardware platform.
- **HPE (formerly Hewlett-Packard)** reliability rebounded from prior years, thanks to the HPE Superdome X. However, the ProLiant server reliability still lags.
- **Cisco, Dell, IBM and Lenovo** rated the highest in customer satisfaction.
- **Red Hat Enterprise Linux (RHEL) and Ubuntu** continue to post the highest reliability scores among all the Linux OS distributions for the third straight year.
- **Red Hat Enterprise Linux and SuSE on IBM Power VM and Lenovo System x** followed by RHEL and Ubuntu on Cisco UCS and RHEL on HP Integrity offer the highest availability among the major server/OS platforms.
- **Cisco, IBM, Lenovo & Dell w/Hyper-V** remain the most reliable virtualization offerings.

Those are the results of the ITIC 2016 – 2017 Global Server Hardware and Server OS Reliability survey which polled 750 organizations worldwide during August through October 2016.

The survey also showed that the three **technology** issues of most concern according to this year's survey once again: Security, Backup and Business Continuity and Integration/Interoperability.

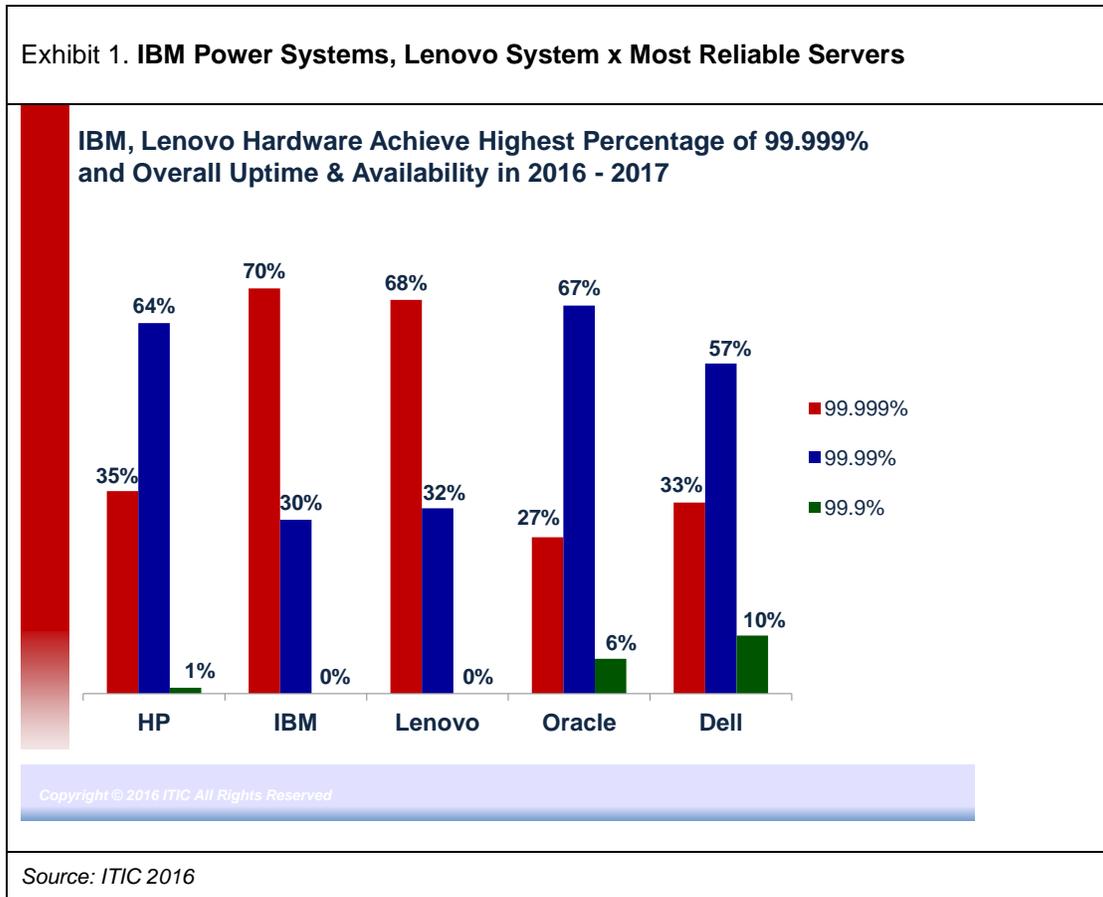
The inherent reliability and robustness of server hardware and the server operating systems are the singularly most critical factors that influence, impact and ultimately determine the uptime and availability of mission critical line of business applications, virtual machines (VMs) that run on top of them and the connectivity devices that access them.

For the ninth year in a row, corporate enterprise users said IBM server hardware delivered the highest levels of reliability/uptime among 14 server hardware and 11 different server hardware virtualization platforms. A 79% majority of IBM Power servers and 68% Lenovo System x servers achieved “five nines” or 99.999% availability – the equivalent of 5.25 minutes of unplanned per server downtime compared to 35% of HPE servers; 33% of Dell hardware and 27% of Oracle server equipment (**See Exhibit 1**).

ITIC's ninth annual Global Server Hardware and Server OS Reliability poll indicates that the inherent reliability of server hardware and server operating system software continues to improve. It must. Today's businesses will not tolerate downtime and businesses of all sizes – from SMBs to global enterprises – are extremely risk-averse. The survey results indicate that external issues - most notably human error and security breaches - can undermine system and network reliability if the organization's core infrastructure is not sufficiently robust or reliable.

The overall health of network operations, applications, management and security functions all depend on the core foundational elements: server hardware, server operating systems and virtualization to deliver high availability, robust management and solid security. The reliability of the server, server OS and virtualization platforms form the foundation of the entire network

infrastructure. The individual and collective reliability of these platforms have a direct, immediate and long lasting impact on daily operations and business results.



As Exhibit 1 above also illustrates, none – 0% of the IBM and Lenovo survey respondents said their servers had a low 99.9% and just 1% of HP servers recorded three nines of uptime – equivalent to 8.76 hours of per server/per annum downtime. By contrast, 6% of Oracle servers and 10% of Dell servers notched 99.9% downtime, according to the latest poll.

Introduction

This Report compares the reliability of 14 major server platforms, 18 server operating system distributions and 11 server hardware virtualization layers. It also provides an in-depth look at the internal issues that have the potential to positively or negatively influence the inherent reliability

of servers and operating systems. It also delves into the overarching industry dynamics and the impact that new and emerging technologies like cloud computing, the Internet of Things (IoT) and Big Analytics can have on planned and unplanned system downtime.

This report quantifies and qualifies the inherent reliability based on key metrics including:

- Automated and manual patch management
- Percentage of Tier 1, Tier 2 and Tier 3 Help Desk calls and length of outages
- Inherent server and server OS reliability
- System unavailability due to planned outages for routine system maintenance, upgrades and the application of patches
- The impact of security issues, including the inherent how quickly vendors are able to respond to security flaws/vulnerabilities with effective patches
- The impact of improperly configuring or right-sizing the server to accommodate virtualization and more compute-intensive workloads
- Server virtualization reliability
- Vendor technical service and support and availability of documentation
- Human error
- Overworked, understaffed IT departments
- The impact of aging server hardware on reliability
- Integration and interoperability issues

This report also utilizes information gathered from previous ITIC surveys to compare and contrast the reliability of the various server hardware, server OS and virtualization platforms and track uptime trends. The survey findings provide crucial reliability metrics to assist organizations in making informed purchasing, management and upgrade decisions for their specific business and budgetary needs.

Data & Analysis

In the mid-1990s, two nines or 99% uptime, the equivalent of nearly 88 hours of per server downtime, was the acceptable norm. In the Digital Era of “always connected” networks, 99% or even “three nines” 99.9% -- 8.76 hours of per server/per annum downtime is unacceptable and unthinkable. Again, the latest ITIC 2016 – 2017 Reliability poll indicates that 72% of respondents now consider 99.99 % to be the minimum acceptable level of reliability for their main line of business (LOB) servers. That’s an increase of 22 % compared to the 49% of respondents in the 2014 survey that required a minimum of “four nines” availability.

In 2016 – 2017, corporate virtualization, cloud computing and IoT ecosystems which are built on interconnected devices demand near-flawless, uninterrupted availability. In virtualized and IoT environments, *the potential* for collateral damage has increased by orders of magnitudes. An outage on virtual server running multiple instances of a crucial main line of business (LOB) application will have a greater impact on productivity, operations and the corporate bottom line compared to a server running a single instance of an application. Similarly, in IoT environments where devices, applications, people and processes are all interconnected, reliability is paramount to the corporation’s ability to conduct business in an uninterrupted fashion. A few minutes of downtime can prove catastrophic, disrupting productivity and cost tens of thousands to millions an hour or event minutes. Firms must also factor in the cost of remediation efforts – time, manpower and expense involved to restore systems and networks to full operational status and /recovering lost data.

As ITIC has done every year since 2008, we publish the specific Table depicting the availability percentages and the equivalent number of annual, monthly and weekly hours and minutes of per server/per annum downtime. **Table 1** provides a useful reference to let organizations or IT departments calculate downtime to measure the business and monetary impact on the firm. Metrics of three or four nines of uptime – 99.9% and 99.99%, – equate to 8.76 hours; 4.38 hours and 52.56 minutes of per server/per annum downtime, respectively.

TABLE 1: Reliability/Uptime by the Numbers

Availability %	Downtime per year	Downtime per month*	Downtime per week
90% (one nine)	36.5 days	72 hours	16.8 hours
95%	18.25 days	36 hours	8.4 hours
97%	10.96 days	21.6 hours	5.04 hours
98%	7.30 days	14.4 hours	3.36 hours
99% (two nines)	3.65 days	7.20 hours	1.68 hours
99.5%	1.83 days	3.60 hours	50.4 minutes
99.8%	17.52 hours	86.23 minutes	20.16 minutes
99.9% (three nines)	8.76 hours	43.8 minutes	10.1 minutes
99.95%	4.38 hours	21.56 minutes	5.04 minutes
99.99% (four nines)	52.56 minutes	4.32 minutes	1.01 minutes
99.999% (five nines)	5.26 minutes	25.9 seconds	6.05 seconds
99.9999% (six nines)	31.5 seconds	2.59 seconds	0.605 seconds
99.99999% (seven nines)	3.15 seconds	0.259 seconds	0.0605 seconds

Source: ITIC 2016

ITIC defines Tier 1, Tier 2 and Tier 3 server outages as follows:

- **Tier 1:** These are the typically **minor** common, albeit annoying occurrences. Network administrators can usually resolve such incidents in 1 to less than 30 minutes for dependent users. Tier 1 incidents can usually be resolved by rebooting the server (locally and remotely) and rarely involve any data loss. Tier 1 outages range from something as innocuous as accidentally unplugging the server to applying a quick update. While outages of five, 10 or 15 minutes were acceptable in the 1990s or early 2000s, that's not the case today.
- **Tier 2:** These are **moderate issues** in which the server may be offline from one hour to four hours. Tier 2 problems may require the intervention of more than one network administrator to troubleshoot. They frequently disrupt network operations for the company's end users and potentially impact business partners, customers and suppliers attempting to access data on an affected corporate extranet. Data loss is possible and some remediation is required.
- **Tier 3:** This is the most **severe** incident. Tier 3 outages are of longer than four hours duration in terms of service unavailability, the corporation's associated dependent users and the remediation efforts of IT. Tier 3 outages almost always require multiple network administrators to resolve issues and there is a greater probability of data loss or damage to systems. Another real threat associated with a protracted Tier 3 outage is potential lost business and damage to the company's reputation. Tier 3 outages can also be man-made such as when e.g., a backhoe cuts a power line; a prolonged power outage; or a natural disaster, such as a hurricane, flood or tornado occurs. Other causes of Tier 3 outages include an external security breach/hack, integration/interoperability problems, or when the IT department cannot obtain the necessary technical support from their vendors, or if no fix is available for an otherwise a minor issue.

Any downtime is an anathema to business operations. Unplanned service outages of even a two to three minutes can wreak havoc, bringing network operations and end user productivity to a standstill. When the servers, operating systems and applications stop, so does the business. Tier 1 and Tier 2 outage of several minutes, one-to-four hours or a severe Tier 3 outage of four hours or more, can result in significant monetary losses, disrupt productivity, damage the company's reputation and raise the risk of litigation. Enterprises in vertical markets such as banking and finance, stock exchanges, communications/media, insurance, healthcare, manufacturing, retail and transportation, whose businesses are based on intensive data transactions, can lose millions if service is interrupted for two, five, 10 or 30 minutes. The consequences can be catastrophic if the outage occurs during peak usage or during a crucial transaction. Small and midsize businesses (SMBs) and midsize enterprises (SMEs) are also vulnerable and just as risk-averse as and even more so than their enterprise counterparts. SMBs and SMEs typically lack the manpower, resources and financial means to withstand the impact of a moderate Tier 2 or severe Tier 3 outage. Worst case scenario: SMBs and SMEs could be out of business.

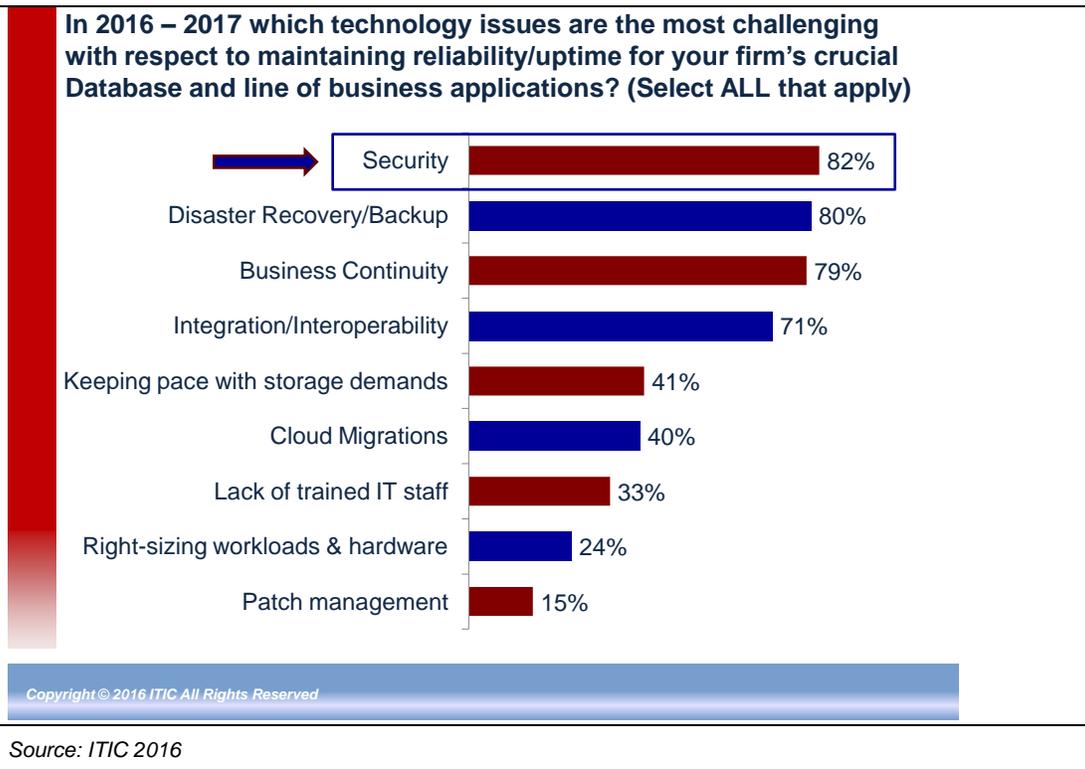
Top Survey Highlights

- IBM Power Systems servers recorded the least amount of downtime just one percent (1%) -- of > 4 hours of per server/per annum downtime of any major mainstream hardware platform.
- IBM Power Systems and Lenovo x86 server hardware running the Red Hat Enterprise Linux (RHEL) and Ubuntu open source operating systems were either first or second in every reliability category, including virtualization and security.
- Cisco, Dell, IBM and Lenovo server hardware received the highest marks for customer satisfaction and overall technical support.
- Oracle x86 and HP ProLiant servers experienced the highest percentages -- 10%, each -- of >4 hours of per server annual downtime. Much of the higher outage time on the Oracle, HP and Dell platforms is attributable to a high proportion of users retaining the hardware for four, five and even six+ years without refreshing/retrofitting the servers.
- Some 51% of survey respondents reported that aged or inadequate hardware (3 ½+ years) has had a negative impact on server uptime and reliability vs. 31% that said it has not impacted reliability/uptime.
- Some 15% of Oracle customers continue to rate service & support as Poor or Unsatisfactory. Dissatisfaction with Oracle licensing and pricing policies remains high.
- Only 1% of Cisco, 2% of Dell, 1% of IBM and Lenovo, 3% of HP, 4% of Fujitsu and 6% of Toshiba users gave those vendors “Poor” or “Unsatisfactory” customer support ratings.
- About 30% of businesses spend more than one hour applying patches manually; this percentage is unchanged from the prior year’s polls.

Human Error, Security Undermine Reliability

To reiterate, ITIC’s survey results indicate that the *inherent* reliability and uptime among most of the 14 major server hardware and 18 server operating system distributions generally continues to improve year-over-year. At the same time, ITIC’s 2016 - 2017 Reliability research reveals that a variety of external factors have more of a direct impact on system downtime and overall availability. As **Exhibit 2** below illustrates, survey respondents indicated that their top technical challenges with respect to maintaining reliability are: security; disaster recovery and backup, business continuity and integration and interoperability. Additionally, organizations indicated they are also challenged by the rapid mainstream adoption of complex new technologies such as virtualization and increasing cloud computing deployments, the proliferation of Bring Your Own Device (BYOD) and mobile technologies.

Exhibit 2 Security is Most Challenging Technical Issue Impacting Reliability



As Exhibit 2 shows, pragmatic daily operational issues: Security, Disaster Recovery/Backup and Business Continuity remain the top three most challenging technology issues that concern corporations with respect to maintaining uptime and reliability. ITIC believes these issues will continue to be among the most crucial going forward with little change in the foreseeable future. Likewise, as more systems become interconnected and the overarching infrastructure, application and network environments become more complex, integration/interoperability issues also plague businesses.

2016 – 2017 Reliability Trends

In the context of its Reliability Surveys, ITIC broadly defines human error to encompass both the technology *and* business mistakes organizations make with respect to their network equipment and strategies.

Human error as it relates to technology includes but is not limited to:

- Configuration, deployment and management mistakes

- Failure to upgrade or right size servers to accommodate more data and compute intensive workloads.
- Failure to migrate and upgrade outmoded applications that are no longer supported by the vendor.
- Failing to keep up to date on patches and security.

Human error with respect to business issues includes:

- Failure to allocate the appropriate Capital Expenditure and Operational Expenditure funds for equipment purchases and ongoing management and maintenance functions
- Failure to devise, implement and upgrade the necessary computer and network to address issues like Cloud computing, Mobility, Remote Access, and Bring Your Own Device (BYOD).
- Failure to construct and enforce strong computer and network security policies.
- Ignorance of Total Cost of Ownership (TCO), Return on Investment (ROI).
- Failure to track hourly downtime costs.
- Failure to track and assess the impact of Service Level Agreements and regulatory compliance issues like Sarbanes-Oxley (SOX), Health Insurance Portability and Accountability Act (HIPAA).

All of the aforementioned human errors can have an immediate, tangible and far reaching impact on daily, monthly and annual system, application and network reliability and availability.

- **Security:** Not surprisingly, given the near constant reports of data breaches – 82% of survey participants said security was their biggest concern and constitutes the biggest ongoing threat to network reliability. Security problems negatively impact overall system and network reliability according to 46% of respondents versus 32% who said security does not affect uptime.
- **Cost of Hourly Downtime Increases: 98%** of firms say hourly downtime costs exceed \$100K and 37% of respondents estimate hourly downtime costs their companies \$300 to \$500K. Nearly one-third of survey respondents indicated that **on average** a single hour of downtime can cost their firms from \$1 million to over \$5 million.
- **Human Error:** Nearly half of all survey respondents indicated that human error is one of the top reliability concerns. ITIC defines “human error” broadly. It is the result of IT administrators who misconfigure systems and that fail to apply the latest patches and security fixes. Human error also refers to the failure of upper

management/C-level executives to approve the necessary Capex budgets upgrade or retrofit key servers and applications on a regular three-year upgrade cycle or to provide the necessary Opex funds to retrain and certify existing IT staffers with the latest certifications.

- **Understaffed IT departments or inadequately trained administrators** were cited by one third of survey respondents as negatively impacting reliability.
- **Flaws in the server operating system** were also referenced by 33% of users.
- **Hardware problems** – (e.g. aging hardware; servers that are outmoded or not robust enough to carry today’s more demanding workloads and trouble getting replacement parts) was cited by 26% or one-in-four businesses as the root cause of their network reliability issues. *Inherent flaws in server hardware* have steadily declined as one of the primary causes of reliability issues over the last 10 years. Nonetheless, organizations are well-advised to regularly upgrade their hardware and right-size their server platforms to accommodate the workloads.

All of the aforementioned human errors will result in immediate, tangible and far reaching consequences on daily operations as well as monthly and annual system, application and network reliability and availability.

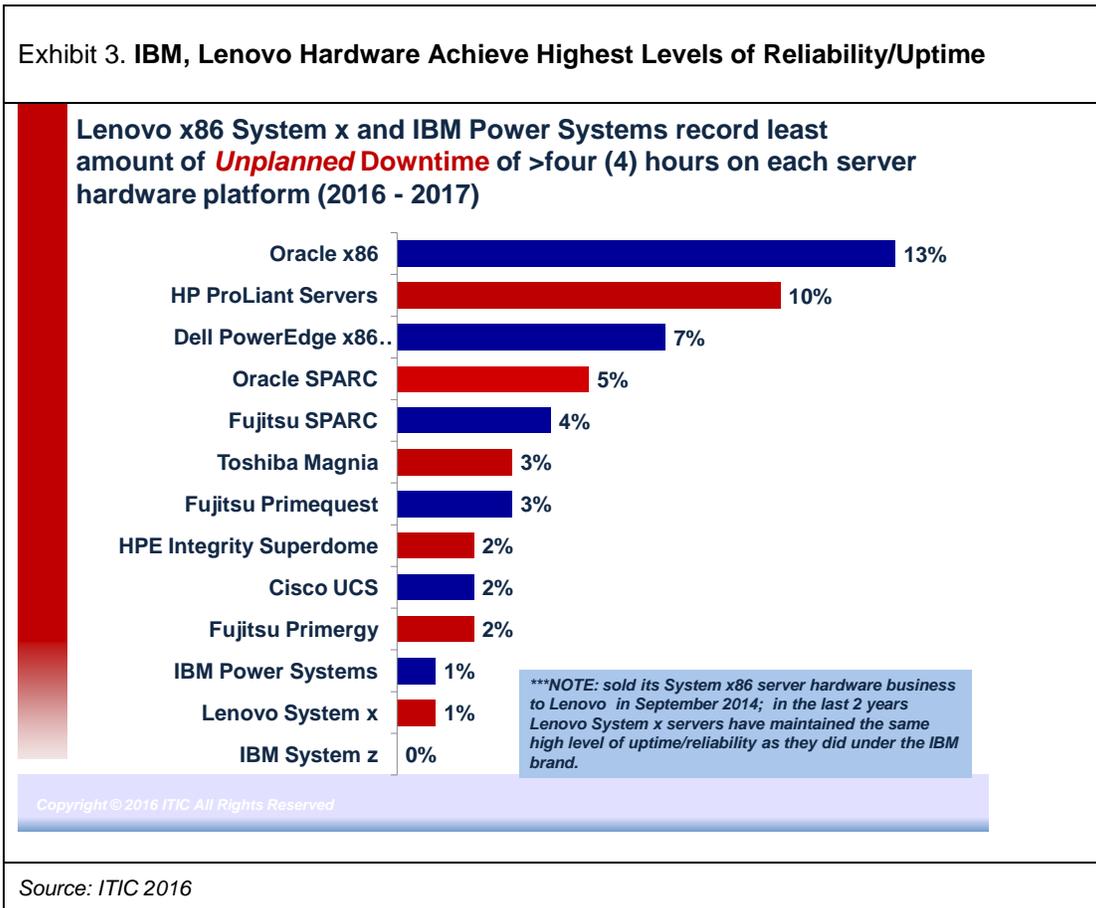
IBM, Lenovo Server Platforms Achieve Highest Reliability

IBM and Lenovo servers recorded the least amount of Unplanned *Downtime of over four hours* – just 1% each, followed by Cisco UCS, Fujitsu Primergy and HPE’s Superdome with only 2% of those platforms recording over four hours of unplanned downtime (**See Exhibit 3**). And as we said above, the IBM and Lenovo servers achieved the highest levels of “five nines” or 99.999% reliability among all mainstream server hardware platforms.

IBM’s System z Enterprise high end mainframe class server was the only major mainstream hardware platform that had **no – 0%** - unplanned server outages of over four (4) hours duration.

The results continue to be upbeat for Cisco Systems, Inc.’s Cisco Unified Computing System (UCS) servers. Appearing for the third year in ITIC Reliability poll, Cisco UCS once again made an impressive showing, posting uptime equal to or better than HP (in certain categories) and bested only by IBM and Lenovo server reliability. Half – 50% - the survey participants said their Cisco UCS servers achieved 99.999% of per server/per annum availability. Cisco, Dell, IBM, Apple and HP (in that order) achieved the highest customer satisfaction ratings for product functionality, reliability, service and support. The Cisco survey respondents also gave the firm high marks for quick problem resolution and the technical savvy of both its telephone and onsite

tech specialists. It is apparent that when corporations retain their servers for protracted periods of four, five and even six years or more without upgrading, they will experience significant increases in unplanned downtime. Vendors' technical service and support; the ready availability of documentation and fixes/patches for known issues and security vulnerabilities also play a crucial role in helping to mitigate or exacerbate and extend outages.



Key Reliability Trends

Among the other top survey highlights:

- Corporate Enterprise Minimum Required Levels of Reliability/Uptime Have Increased Dramatically from 2008 to 2016.** Over the last eight (8) years the minimum reliability requirements for businesses of all sizes and across all vertical markets has increased significantly.

- **Some 30% of businesses** spend more than one hour applying patches manually; up from 26% in 2013 but nowhere near the peak of 40% in the ITIC 2011 Reliability survey
- **Over three-quarters – 77% - of IBM AIX users and 69% HP UX users** reported they “never” or “rarely” reboot the server OS, including planned reboots to add or reconfigure system resources. In contrast, less than half -- 49% of Oracle Solaris and 45% of Windows Server survey respondents indicated they “rarely” or “never” had to reboot their servers for planned maintenance.
- **Nearly one-third or 31% of businesses** don’t provide for hardware failover and redundancy and 12% of companies don’t bother to track hardware failure rates.
- **Some 45% of respondents rely on the built-in redundant hardware capabilities** of their servers to provide high availability and failover protection.
- A **51% majority of respondents** said their main line of business (LOB) servers were two-to-four years old. Of that number 30% revealed their LOB servers were two-to-three years old; 21% said they were three-to-four years old. Another 21% indicated their servers were four, five or greater than five years old.
- **Customer satisfaction:** Dell (75%), IBM (73%), Lenovo (73% and Cisco (70%) achieved the highest customer satisfaction ratings for products, service and support. Oracle’s customer satisfaction ratings were again the lowest in the survey. Only 47% ranked its technical service and support as “excellent” or “very good – although that was an increase of two percent over the ITIC 2013 Reliability survey numbers.” However, for the third year in a row, Oracle’s customers were the most dissatisfied in the survey. Some 14% of respondents rated Oracle’s technical service and support “poor” or “unsatisfactory.”

Minimum Reliability Requirements Increase Year over Year

In 2016 – 2017 time is measured by money, productivity, return on investment (ROI) and managing the corporation’s level of risk and exposure to potential litigation as a consequence of unplanned downtime. This is evidenced by corporations’ reliability requirements which have increased every year since 2008 when ITIC began polling organizations on these metrics.

Additionally, the most recent 2016 – 2017 ITIC poll shows that 27% of the over 750 respondent organizations require a very high 99.999% or better degree of reliability (**See Exhibit 4**). This is a scant 5.25 minutes of unplanned per server, per annum downtime. The percentage of respondents that indicated their businesses need “five nines” or better availability more than doubled over the last several years. Consider the following: in 2008, fewer than three-in-10

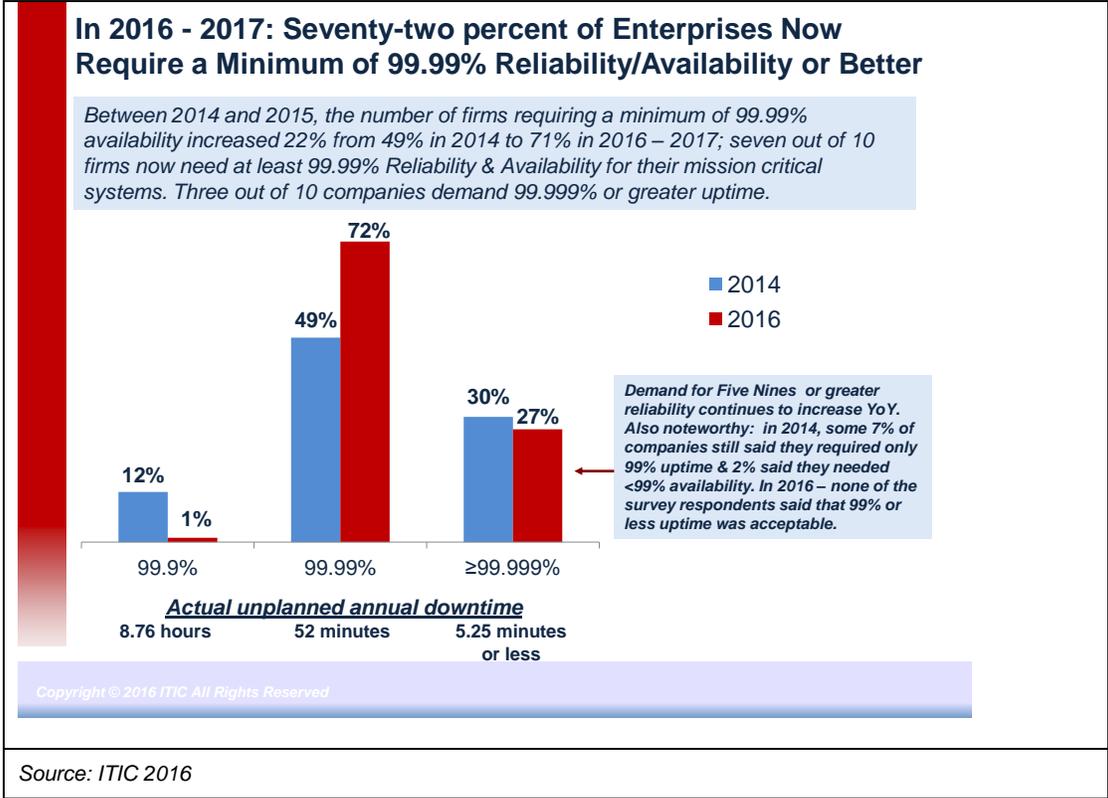
businesses or 27% of survey respondents required just 99% uptime. That same year (2008) four-in-10 corporations – 40% required 99.9% availability, while 23% of firms required a minimum of “four nines” or 99.99% uptime for their servers, operating systems and virtualized environments. And only a seven percent (7%) minority demanded the highest levels of “five nines” – 99.999% or greater availability. In the ITIC 2013 Reliability survey, for example, only 11% of participants said their firms required “five nines” reliability.

By 2013, 67% of businesses cited a requirement for 99.99% or greater reliability/uptime; up 34% from 2008. 99.99%+ and greater reliability are mission-critical. In the 2015 – 2016 and in our latest 2016 – 2017 ITIC Reliability surveys – *none – 0% of survey respondents* indicated their organizations could live with just “two nines” – 99% uptime or nearly 88 hours of annual unplanned per server downtime! And only a miniscule one percent – 1% - minority said their firms required just 99.9% reliability.

Four nines or 99.99% uptime and availability is the new normal or average minimum requirement for 72% of organizations.

Additionally, 24% of those polled indicated their businesses require “five nines” or 99.999% server and operating system uptime which equates to 5.26 minutes of per server/per annum unplanned downtime. And three percent (3%) of leading edge businesses need “six nines” 99.9999% near-flawless mainframe class fault tolerant server availability of 31.5 **seconds** per server/per month.

Exhibit 4. Nearly Three-Quarters – 72% of Enterprises Now Require 99.99% Reliability



Server Hardware Vendor Landscape

It's no secret that the server hardware vendors are locked in a battle for market share and dominance.

HPE, Dell, IBM, Lenovo and Cisco in that order, are the Top Five vendors in 2016 to date, in both market share and revenue. Together these five vendors account for roughly 65% of server revenue and market share, with all other vendors representing the remaining 35%.

It's also no secret that worldwide server hardware vendors are facing pressure on revenue and slowing shipment growth.

Competition is cutthroat among all vendors.

In order to retain and up sell to existing customers and add new customers, all of the hardware vendors are differentiating their offerings and solutions.

Dell and HPE have traditionally competed on low pricing and high volume for their workhorse server lines (e.g. Dell PowerEdge, HP ProLiant). While IBM, Lenovo and more recently Cisco are increasingly focused on embedded features and enhanced functionality – and of course, reliability and superior service and support.

All of the server vendors are showcasing performance, reliability, scalability, security and management capabilities as a means of promoting their server solutions.

Processor Technology

The servers are only as good as their component parts. The memory, CPU and other functionality like management and security must stay abreast of the accelerated rate of technology innovation and change. Any firm that expects to achieve optimal server uptime and reliability must invest in the underlying processor technology that best aligns with their applications and workloads. Market leader Intel, which has close to 90% of the microprocessor market continually, updates its processors.

The Intel Xeon Processor E7 v2, E7 v3 and the latest E7 v4 released in the 2016 second calendar quarter are specifically geared towards reliability. All of the E7 v2, E7 v3 and E7 v4 accommodate and accelerate in-memory analytics and the scalability to manage large data sets and the ability to boost I/O performance. In a definitive nod to reliability, the Intel Xeon Processor E7 v2, v3 and v4 families incorporate *Intel's Run Sure Technology*. It integrates processor, firmware, and software layers to diagnose fatal errors, contain faults and automatically recover to keep the server operating. The Xeon Processor E7 v2, v3 and v4 chips also feature Intel's *Resilient Memory Technologies*, to solidify and ensure data integrity within the memory subsystem.

Intel claims that the RAS enhancements in the E7 v3 and v4 families of chips ensure up to “five nines” or 99.999% uptime.

Real-time business intelligence is a top priority for small and midsize businesses (SMBs) to the largest enterprises and spanning all vertical industries. In addition to high performance and reliability, the E7 v3 and v4 processors assist customers in parsing and analyzing the increased amounts of data – including unstructured data to extract actionable insights. The Intel Xeon processor E7 v3 and v4 families help organizations to securely process and analyze massive data sets in system memory for faster decision-making and improved operational efficiency, giving companies a competitive edge. According to Intel, the E7 v3 processor family delivers 6x to 8x improvements in business processing application performance for in-memory transactional workloads and is optimized with the new Intel® Transactional Synchronization Extensions (Intel® TSX). The Intel Xeon Processor E7 v2, v3 and v4 offerings are built around the philosophy of continuous self-monitoring and self-healing. Self-healing features enable the server to proactively and reactively repair known errors and minimize future errors. This in turn, bolsters reliability and uptime.

The synergies between the server hardware and the underlying Intel processors enable organizations to realize *average* performance gains of 20% to 30% (depending on individual workloads and configurations) and reduced power consumption of 15% to 25% compared with prior models. A dual core processor may deliver sufficient performance, speed and response time for a small server that services a department or remote branch office. However, it could prove entirely inadequate for a main LOB server-based application that features digital audio, video and large, very dense file formats. Compute-intensive, business critical workloads require stronger, more advanced four- and eight-socket processors.

Intel processors are installed in a majority of today's servers, desktops, notebooks and tablets. For the third year in a row the ITIC Global Reliability survey incorporated specific questions on corporations' experiences with the performance and reliability of the latest Intel Xeon E7 v2 Processors. Among the key findings:

- A 67% majority of corporations continue to experience 35% to 60% improvement in reliability and performance (the variation was determined by the age of the server hardware and the workload) with the Intel Xeon E7 v2, v3 and v4 processor across-the-board on x86 based servers including Cisco UCS Dell, HP, IBM and Oracle.
- Eight out of 10 servers equipped with the Intel Xeon E7 v3 and v4 achieved “four nines” or 99.99% of per server/per annum uptime/availability. That equates to 52 minutes of per server/per annum downtime or 4.33 minutes per month.
- Customer satisfaction with Intel performance, reliability, service and support is extremely high. Some 48% of respondents rated it “Excellent,” 42% rated it “Very Good” and 10% rated it “Good.” **None** of the respondents gave the Intel processors or Intel's technical support a “Poor” or “Unsatisfactory” rating.
- Virtual servers equipped with the Intel Xeon E7 v3 and v4 Processors received similarly high performance and reliability grades. An 80% majority achieved a minimum of 99.99%; 44% achieved 99.999% and four percent attained 99.9999%.

The *inherent* reliability of the Intel x86-based and RISC processor platforms is extremely close. ITIC's survey results and subsequent first person conversations with IT managers indicates that the improvements in the Xeon E7 v2, v3 and the newest E7 v4 families can deliver comparable native performance, reliability and uptime when installed on robust servers that are new or up to 3 years old. Achieving these high levels of reliability demands that corporate enterprises upgrade and right-size the server configurations every two-to-three years to ensure they are adequate for current and future compute intensive physical, virtual and cloud workloads.

The breakdown in x86 server reliability **generally** (but not always) occurs not because of any **inherent flaws** in the underlying server hardware because x86 customers unwisely “push their luck” and retain their server hardware for 4 ½ to sometimes 6+ years without retro-fitting,

upgrading. A business that overloading old, outmoded servers or misconfigures a server will consequently experience availability problems. This is particularly true of organizations that purchase entry level or inexpensive commodity servers.

ITIC's first person customer interviews yielded invaluable anecdotal information regarding the net positive performance and reliability gains specifically related to the newer processor technology.

Businesses are well advised to refresh or upgrade their servers with the latest processors in order to accommodate higher workloads and optimize reliability.

Real World Reliability Scenarios

Once again, as with the prior year's survey, the latest 2016 – 2017 ITIC survey results clearly demonstrate that the human element is having a greater influence on reliability – and not always for the better – than the underlying hardware technology. Once again the survey results showed that human errors related to configuration and management; overworked/understaffed and less skilled IT managers along with utilizing too-old or inadequate server hardware were the primary causes of unreliable systems.

The fact that 10% of HP ProLiant and Oracle systems racked up the highest percentages of downtime exceeding four (4) hours of unplanned downtime is an eye opener. That is, until an in-depth analysis of the results indicate that the prolonged unplanned downtime of over four hours has less to do with the inherent reliability or instability of the HP and Oracle servers and is more indicative of end user behavioral patterns. Some 53% of HP ProLiant users, for example, indicated they retained their servers for four, five or even six years or longer without upgrading/retrofitting or right-sizing the servers to accommodate today's more compute intensive workloads.

By contrast, only 20% Lenovo System x, 18% IBM Power Systems users and 21% of Cisco UCS users retained their servers for four or more years without upgrading or retrofitting them.

The high end enterprise mainframe IBM System z Enterprise is a different breed of server and in a class of its own. The System z is more hardened and robust than the overwhelming majority of mainstream LOB servers. And the System z is typically managed and maintained by highly experienced/guru administrators.

Organizations that opt to retain server hardware for four+ years, exceeding capacity limitations, top notch technical service and support from their vendors assumes much greater importance. The vendors' ability to respond quickly; the ready availability of replacement parts and

immediate access to documentation and fixes/patches for known issues and security vulnerabilities will help to mitigate or exacerbate and extend corporate outages.

Superior technical support explains why Dell consistently gets high marks for customer satisfaction even when reliability is spotty. Dell has also benefitted greatly from its decision to move its corporate technical service and support organization back onshore. Many Dell and Oracle customers told ITIC in their anecdotal comments and acknowledged first person interviews in this latest survey and over the past several years, that they have only themselves to blame when the servers crash because they are insufficiently robust enough to carry heavy workloads. Vendors' ability to deliver superlative technical service and support figures prominently in reliability and increased system availability. Technical service and support – good and bad –distinguishes and differentiates vendors from rivals. How promptly, efficiently and effectively vendors respond to corporate customers when issues arise has a definite impact on customer retention and the company's willingness to upgrade and purchase new equipment and software, expand usage of specific products and renew service contracts.

IBM Rock Solid Reliability

Corporate enterprises have given IBM hardware the highest reliability ratings every year since 2008 when ITIC began conducting the Global Server Hardware and Server OS Reliability poll.

This is no accident.

IBM's high reliability ratings over the past decade speak to the technical excellence and robustness of the hardware platform. The rock solid reliability also reflects and underscores the consistency of IBM's technical service, support, security and customer responsiveness over the last six years and the stability of Big Blue overall. IBM like many vendors has suffered contraction in its hardware sales. In 2014 IBM notably sold its commodity x86 server business (which included approximately seven thousand employees) to Lenovo for just over \$2 billion. This move benefits both IBM and Lenovo. IBM is now free to concentrate on high end servers like its Power Systems and mainframe class System z Enterprise servers, while Lenovo intends to make a success of the x86 platforms.

IBM under chief executive and President Virginia Rometty is a bastion of stability. This stability extends to Big Blue's continuous investment in improving the core RAS and performance capabilities across its server lines and working closely with Linux OS and open source vendors like Red Hat, SUSE and Canonical which makes the highly regarded Ubuntu.

Thanks to an ongoing commitment to R&D and continual refresh of its embedded performance, reliability, security and management functions, IBM hardware retains its status as best in class in terms of reliability, stability and performance and customer satisfaction. IBM consistently has demonstrated its ability to articulate and craft a comprehensive product roadmap and strategy and execute against it. IBM (along with Cisco Systems, Intel and Microsoft) is perennially on the Top 10 list of companies that spend the most of research and development (R&D), according to Standard & P Capital IQ. For the past several years – from 2014 through 2016- IBM has spent nearly \$6 billion on R&D (six percent of annual revenue). This is approximately twice as much as rival HP which spent just over \$3 billion on R&D and 19% more than Oracle which allocated \$5 billion to R&D. IBM also continues to rank very high in customer satisfaction. IBM servers recorded the lowest incidences of the more significant Tier 2 and Tier 3 server outages lasting from one-to-four hours or more.

IBM's POWER8-based processor systems provide several key feature/function advantages that advance reliability and enable customers to lower Total Cost of Ownership (TCO) and achieve near-immediate ROI.

IBM has made several key improvements to the processor to turbo-charge performance and throughput to accommodate today's more compute-intensive workloads. The POWER8 systems have 4x the threads and 4x the memory bandwidth per core, making them capable of processing more data than many rival processors. Additionally, the POWER8 servers contain built-in virtualization and OpenStack-based management to make handling private and public cloud deployments easier. Organizations can take advantage of all the newest systems software irrespective of whether or not the servers run Linux, AIX or IBM i operating systems.

While performance and throughput will vary according to the individual enterprise's specific workloads and implementations, IBM internal benchmark tests indicate that companies running prior generation POWER6 and POWER7 technology will realize a 2.5 and 2x per-core performance boost, respectively when they upgrade to POWER8. In additional response times will accelerate will accelerate with up to 3.5x greater throughput and as much as 85% faster response time depending on the underlying server hardware platform and configuration. The responses times and throughput achieved on POWER8 technology will also vary according to workload types e.g. Open Source databases.

As a founding member of the OpenPOWER Foundation, IBM and the other members are aiming for performance improvements of an order or two orders of magnitude – that is, 10x and 20x performance gains.

IBM architects each new generation of POWER technology to deliver more robust per-core performance and faster overall system throughput. This is crucial because many corporations pay for their software according to the number of cores. Boosting the per-core performance can help enterprises lower the overall infrastructure and software licensing costs.

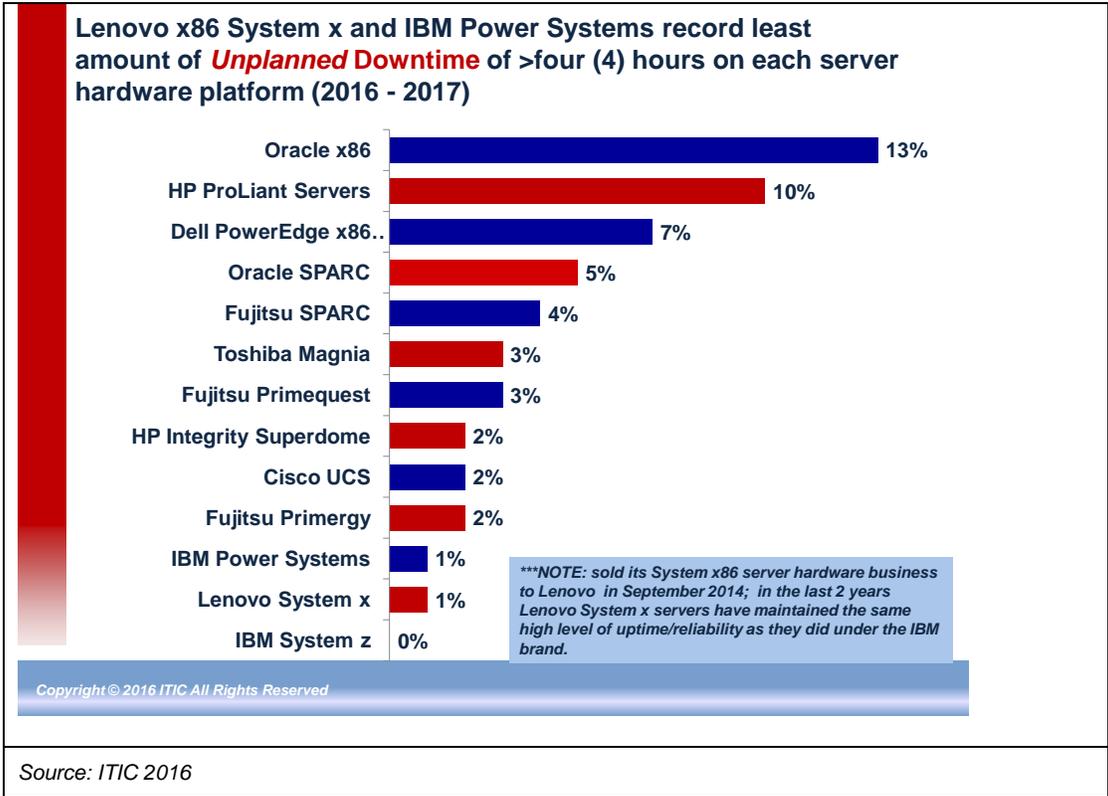
IBM is also emphasizing that a move to POWER8-based solutions can leverage cloud capabilities. This is due to the BIM Power Virtualization Center (PowerVC), which is IBM's OpenStack-based cloud management offering, deployed in combination with POWER* technology. Together, these two technologies allow organizations to move and manage cloud-based workloads quickly and efficiently. Another key feature of the POWER8 technology is its scalability and elastic capacity. Corporations have the flexibility to move or add capacity around on their datacenters as needed to accommodate workloads when and where they're needed most during peak usage hours.

IBM POWER8 technology also includes such features as:

- Live partition mobility capability, which enables organizations to move virtual machines from POWER7 to POWER8 processors, while they're running and without taking the system offline. This provides businesses with improved productivity and system availability by eliminating/reducing even planned downtime.
- Backwards compatibility with older versions of POWER operating systems, which obviates the need for businesses to update their OSes.

As **Exhibit 5** illustrates, these improvements have paid off for IBM POWER Systems, which along with Lenovo, once again achieved the highest reliability ratings in the ITIC 2016 – 2017 Global Server Hardware and Server OS survey.

Exhibit 5. Lenovo System x, IBM Power Systems Servers Deliver Highest Reliability



Additionally, IBM Power Systems and Lenovo System x servers averaged the lowest percentage (4%) of annual server downtime of one to over four (>4) hours compared to an average of 6% for Dell servers and 8% of all HP servers and 8% of all Oracle server hardware. IBM’s System z13 Enterprise mainframe – which is in a class by itself -- was the only mainstream server offering that had no – 0% - of unplanned system downtime due to any inherent flaws in the hardware. Of the mainstream server hardware platforms, IBM Power Systems and Lenovo System x topped the poll besting all other server distributions.

In the semiconductor sector, Intel continues to outspend all of its rivals on R&D as it pursues an aggressive course of innovation to keep pace with rivals like IBM. In 2015, Intel’s R&D spending hit \$12.1 billion, up from its previous record-high of \$10.6 billion in 2014. And the world’s number one chipmaker remains on track with its 2016 R&D spending as well. Intel’s R&D expenditures were triple that of its closest competitors Qualcomm which spent \$3.7 billion on R&D and Samsung which spent \$3.1 billion on R&D. To put this in perspective, Intel’s 2015 R&D spending accounted for 22% of **all** R&D expenditures in the semiconductor sector.

© Copyright 2016 Information Technology Intelligence Consulting Corp. (ITIC) All rights reserved. Other products and companies referred to herein are trademarks or registered trademarks of their respective companies or mark holders.

A significant portion of Intel's R&D investments are specifically aimed at ameliorating the reliability of its latest processor family the Intel Xeon Processor E7 v3 and E7 v4 which are widely used in Lenovo System x, Dell PowerEdge, HP ProLiant and Oracle systems. To reiterate, The E7 v3 and E7 v4 chips accommodate and accelerate in-memory analytics and the scalability to manage large data sets and the ability to boost I/O performance. In a definitive nod to reliability, the latest Intel Xeon Processor E7 v3 family incorporates the *Intel Run Sure Technology* that integrates processor, firmware, and software layers to help diagnose fatal errors, contain faults, and automatically recover to keep the server operational.

Lenovo System x Servers Best in Class Reliability

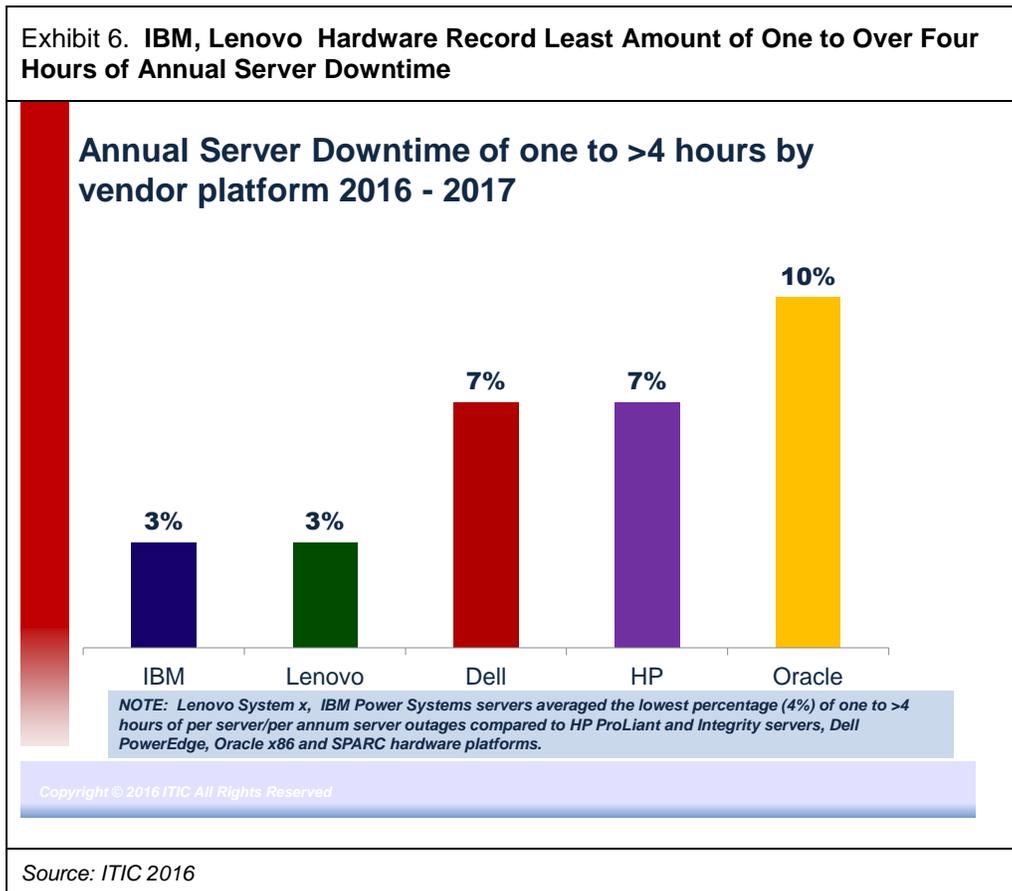
It has been a full year since IBM completed the sale of its x86 server business to Lenovo in a deal that was valued at \$2.3 billion.

Over the last two years the Lenovo System x servers have maintained the same high levels of reliability they had previously achieved while under the IBM brand. In public statements, said the IBM deal represented a new "growth engine" for his company. On May 26, 2016 Lenovo announced the results for its fourth fiscal quarter and full-year ended March 31, 2016. In the **Enterprise Business Group (EBG)**, which incorporates servers, storage, software and services sold under both the Lenovo ThinkServer brand and the System x Business Unit, sales were \$4.6 billion (US dollars), up 73 percent driven by hyperscale wins in China in the full year, while quarterly revenue fell 8 percent to US\$1 billion, primarily as a result of a sales force model that was not fully aligned to maximize opportunities for the company. For the sixth full quarter since the System x acquisition, EBG delivered positive operational pre-tax income with 1.7 percent operational margin.

Going forward, Lenovo Chief Executive Yang Yuanqing said in a public statement that the strategy for EBG which has been renamed as the Data Center Group will be to attack top line growth in mature markets with a better aligned sales model, end to end business structure and dedicated leadership focus. It will fully leverage its new partnership approach with best-in-class next generation technology partners like SAP, Nutanix, Juniper and Red Hat to capture high growth segments of the full United States-based \$87 billion data center market.

Lenovo also markets its line of entry level ThinkServers but it is clearly counting on the System x servers to compete in the enterprise. These higher-end machines have the proven performance, reliability and advanced features to accommodate more compute-intensive and complex workloads and incorporate complex analytics and database-related functions.

This is excellent news for Lenovo customers and Lenovo itself. As **Exhibit 6** shows, Lenovo System x servers retain the same high levels of reliability they delivered while still under the IBM brand. Additionally, IBM and Lenovo server platforms recorded the least amount of unplanned annual server downtime of one to over four (4) hours of any of the mainstream server hardware distributions.

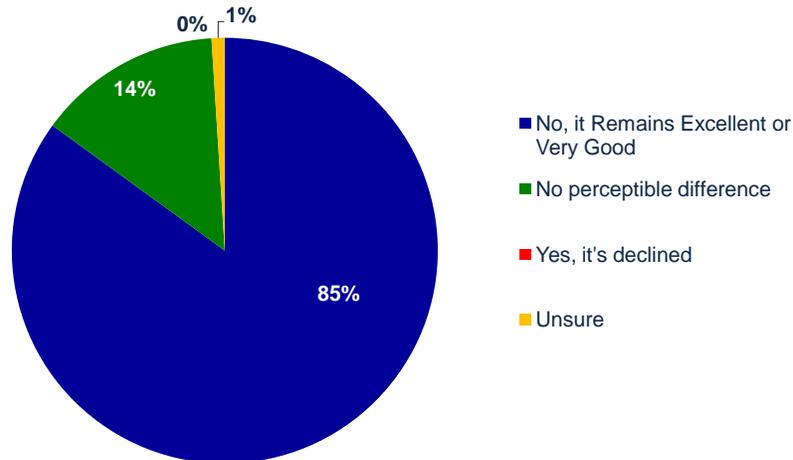


Additionally IBM Power Systems and Lenovo System x servers recorded the least amount of downtime of one to 40 minutes and 41 minutes to up to four (4) hours of per annum/per system downtime of any of the over one dozen major server hardware platforms.

ITIC's 2016 - 2017 Reliability Survey again polled organizations on how the System x servers were faring in terms of product quality and the responsiveness and quality of after-market technical service and support. As **Exhibit 7** illustrates, an overwhelming 99% of respondents rated service and support, product quality as excellent and had experience no perceptible differences.

Exhibit 7. Lenovo System x Servers Maintain High Reliability, Excellent Technical Service and Support Record

Has the Quality of the Products, Technical Service & Support Changed Since Lenovo Acquired IBM's x86 Server Business in 2014?



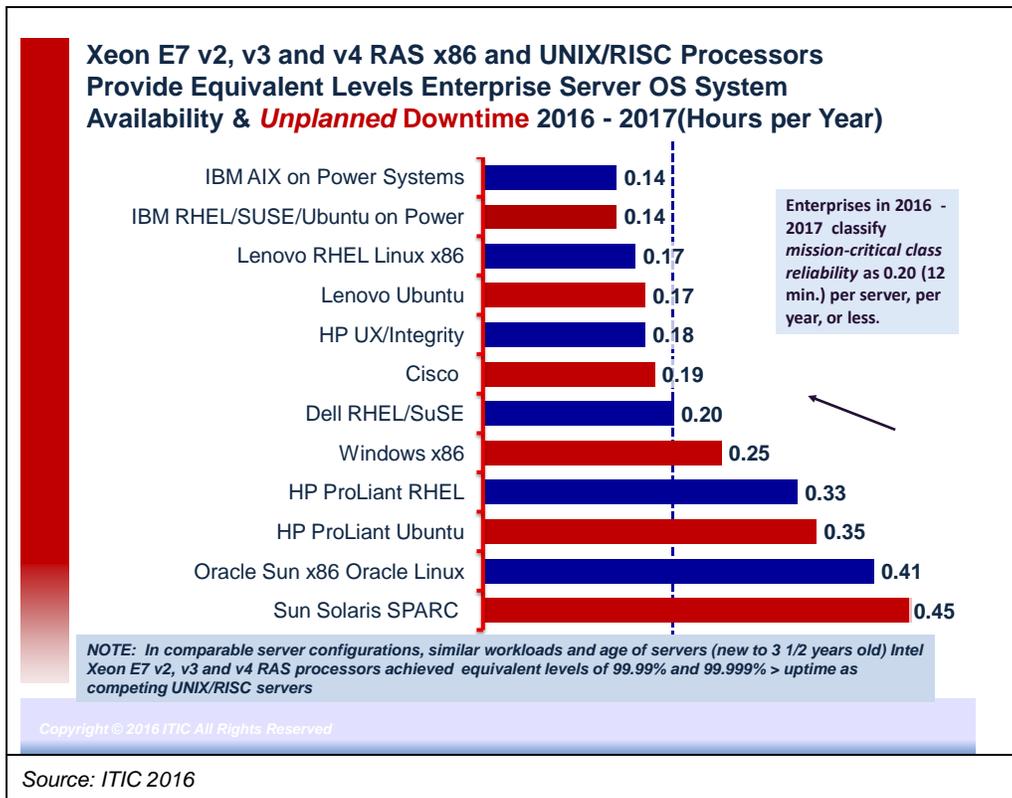
NOTE: NONE – 0% - of survey respondents said that the quality of products or technical service & support has declined since Lenovo purchased IBM's x86 Server Business in 2014

Copyright © 2012 ITIC All Rights Reserved

Source: ITIC 2016

As **Exhibit 8** indicates, the Intel Xeon E7 v3 and E7 v4 RAS processors when properly deployed in newer systems can help organizations achieve equivalent levels of high reliability in comparison to competing UNIX and RISC-based servers in comparable configurations and workloads. Enterprises are demanding higher levels of uptime/reliability from their server platforms with each passing year. The ITIC 2016 – 2017 Reliability poll found that respondent organizations now classify mission critical class reliability as 0.20 or 12 minutes of ***unplanned server downtime per server/per annum***. That equates to just one minute per month of unplanned downtime due to any type of outage.

Exhibit 8. Intel Xeon E7 v3/E7 v4 RAS Processors Deliver Equivalent Level of 99.99% & 99.999% > Reliability as UNIX/RISC Servers in Comparable Configurations

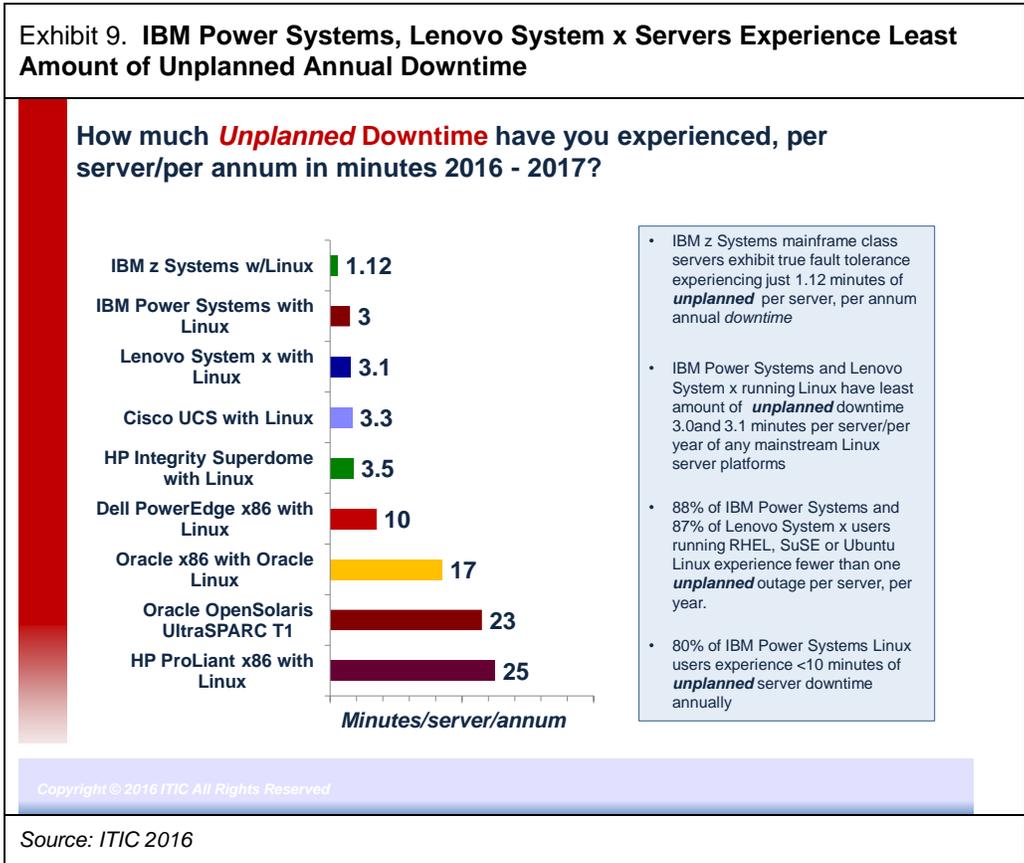


Survey respondents also said that IBM Power Systems and Lenovo System x servers recorded the least amount of per server/per annum of any of the x86-based hardware platforms, with just one percent (1%) of servers notching prolonged outages of over four (>4) hours duration. On the other end of the spectrum, corporate customers reported that 10% of the HP ProLiant and Oracle servers racked up downtime exceeding four hours per annum/per server downtime. According to customer interviews and essay comments, the high amount of Dell PowerEdge and HP ProLiant system unavailability is attributable mainly to the high proportion 60% of aged Dell PowerEdge and 53% HP ProLiant servers that are over four years old and have not been upgraded or replaced.

Upon delving more deeply into the research, ITIC found that in those instances where the HP ProLiant , Oracle and Dell PowerEdge servers were a more reasonable two-to-three and a half years old, the percentage of system downtime of over four (>4) hours dropped to a much more reliable three percent (3%) on both the Dell PowerEdge and HP ProLiant platforms.

As **Exhibit 9** indicates, IBM Power Systems and Lenovo System x servers again recorded the least amount of unplanned downtime among all mainstream server platforms – just three (3) minutes each respectively.

Once again, the IBM z Systems Enterprise mainframe class server had the best overall score: recording a scant 1.12 minutes of per server/per annum “blink and you miss it” outage.



Cisco UCS Comes On Strong, HP Reliability Rebounds

Cisco's UCS servers also achieved high reliability scores; 86% of respondents stated they experienced 40 minutes or less of annual downtime. Survey respondents also continue to give Cisco UCS as well as the company's service and support organization very high customer satisfaction marks; 70% of participants rated it "excellent" or "very good."

Cisco, like IBM also has strong, solid management and has ably managed the executive management transitions. Although Cisco has not been immune to the upheaval and contractions in the high technology market sector – since 2013, the company has restructured and cut approximately 10,000 workers in two separate rounds of layoffs, representing about 14% of its global workforce. On the plus side, Cisco continues to make targeted acquisitions, invest heavily in R&D and forge ahead into new markets most notably IoT – or the Internet of Things. In April, 2014, Cisco Systems launched a \$150 Million initiative to fund IoT startups. Since then, Cisco has aggressively moved to bolster and expand its presence in hardware, switches and data analytics software with a number of key acquisitions and by entering into several pivotal partnerships. Most notably, in the summer of 2016 Cisco and IBM inked a pact to embed IBM's Watson cognitive computing software on Cisco routers and switches at the edge of the network to further both firms' IoT ambitions. All of these efforts in turn, help both the visibility of the Cisco brand and further the functionality and performance of its core networking gear and UCS platforms. This is evidenced by the high marks survey respondents gave to Cisco UCS reliability and to its technical service and support organization.

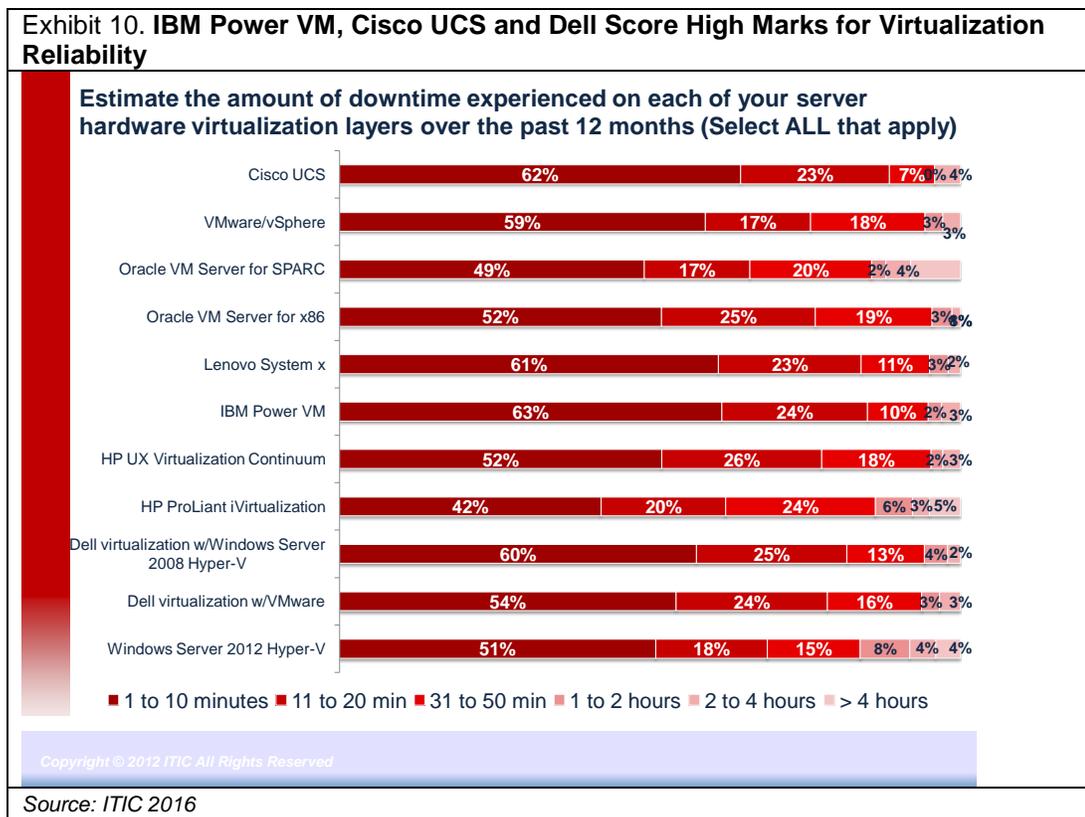
Server Virtualization Reliability

Virtualization is another key technology that has a pivotal impact on network reliability, particularly in the Digital Age. The benefits of virtualization are well known and well documented. Organizations are consolidating various workloads onto fewer servers to cut hardware costs; optimize facility space, conserve and cut manpower costs and administrative resources and lower power consumption.

However, this can be a two-edged sword. Businesses that fail to properly configure and right size workloads and adequately secure their environments run the risk of increased downtime and disrupting productivity. These are two very important considerations – particularly as companies interconnect more devices, applications and people and migrate to the cloud and the Internet of Things (IoT). Mobile workers and Bring Your Own Device (BYOD) can also potentially strain and undercut reliability if organizations fail to take the time and ensure end-to-end security and accountability as users, business partners and customers access the corporation's network and crucial data assets. Increased functionality and connectivity offers the potential for greater economies of scale and opportunity. But to reiterate without the proper controls, correct configuration, strong security and adequate server hardware, the virtualization environment can

turn into a giant Achilles Heel that can result in a higher degree of collateral damage because multiple instances of an application(s) are housed in a single server platform.

Over three-quarters of midsized and enterprise corporations utilize virtual servers for their business critical mainstream applications and workloads. IBM Power VM, Cisco UCS and Dell servers with Windows Server 2012 R2 Hyper-V were the top three most reliable hardware virtualization platforms (see Exhibit 10). Virtualization market leader VMware was in the middle of the pack with 59% of VMware vSphere customers reporting just one to 10 minutes of unplanned downtime followed by 52% of HP UX Virtualization Continuum businesses who had up to 10 minutes of unplanned per server/per annum downtime. One big surprise was that 60% of Dell Virtualization systems running Microsoft Hyper-V experienced one to 10 minutes of unplanned downtime compared to the 54% of Dell Virtualization systems running VMware.



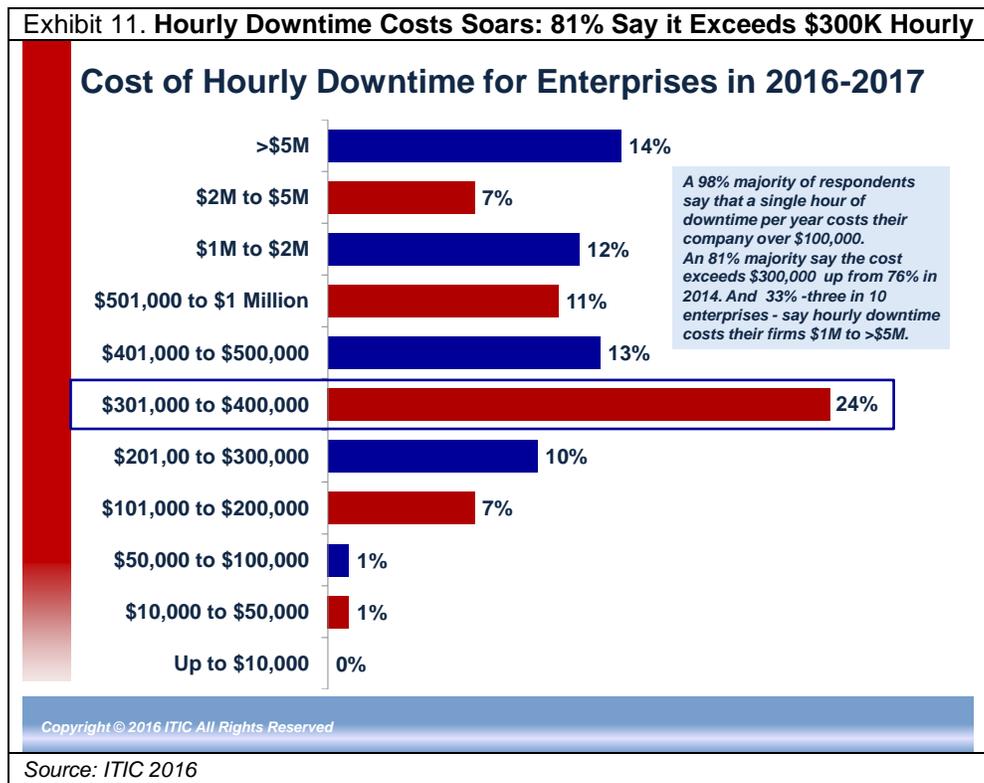
Cost of Hourly Downtime Soars

The only good downtime is **no downtime**.

As companies grow ever more reliant on their interconnected networks and applications to conduct business, the cost of downtime is rising commensurately. To put it bluntly: if the servers, applications and networks are unavailable *for any reason* business and productivity slow down or come to a complete halt.

For the fourth straight year, the survey data indicates that the cost of hourly downtime increased.

As **Exhibit 11** shows below, 98% of organizations say that a single hour of downtime costs over \$100,000; 81% of respondents indicated that 60 minutes of downtime costs their business over \$300,000 and a record one-third or 33% of enterprises report that one hour of downtime costs their firms \$1 million to over \$5 million.



These statistics reinforce what everyone knows: infrastructure matters. And server hardware, server OS and application reliability or instability, can have a direct and far reaching impact on the corporate bottom line, ongoing business operations and irreparably damage companies' reputation. In some extreme cases, business and monetary losses as a result of unreliable servers can cause the company to go out of business – as a consequence of sustained losses and possible litigation brought on by the outage.

Conclusions and Recommendations

In summary the ITIC 2016 - 2017 Global Server Hardware and Server OS Reliability Survey findings indicate that the reliability and uptime of the core infrastructure: server hardware, server operating systems and server virtualization are all more crucial than ever to the health of the overall on-premises datacenter and cloud-based environments.

In the 21st Century Digital Age systems, applications and networks must be available 24 x 7. Reliability and availability are crucial to uninterrupted business operations and productivity.

Among the “work horse” mainstream systems, the IBM Power Systems, Lenovo System x and Cisco UCS have the highest levels of inherent reliability.

Organizations must have confidence in the reliability and stability of their server hardware and server OS platforms. The underlying advances in semiconductor, software and management technology bolsters the reliability of the servers. However, these technology gains can be undone by human error and the ever-present threats posed by security issues, BYOD and the Internet of Things.

It is incumbent upon server hardware vendors to deliver top notch reliability and superior technical service and support. However, organizations and their IT departments ultimately bear responsibility for keeping their core infrastructure up-to-date and configured to accommodate the demands of increasingly compute-intensive applications and network operations. To accomplish this, the corporation must devote the necessary capital and operational expenditures and manpower resources to ensure peak levels of reliability. Achieving optimum uptime means upgrading refreshing server hardware as necessary in order to support more data intensive workloads and physical, virtual and cloud environments. Close attention must be paid to system integration and interoperability, security fixes, patch management and documentation. Business performance will almost certainly suffer if server configurations are inadequate for current tasks and requirements.

Companies should monitor their service level agreements (SLAs) to ensure that they meet the desired reliability levels. If they do not, corporations should ascertain the cause and make the necessary improvements.

Reliability is among the most crucial metrics in the organization. Improvements or declines in reliability mitigate or increase technical and business risks to the organization’s end users and its external customers. The ability to meet service-level agreements (SLAs) hinges on server reliability, uptime and manageability. These are key indicators that enable organizations to determine which server operating system platform or combination thereof is most suitable.

To ensure business continuity and increase end user productivity, it is imperative that businesses maximize the reliability and uptime of their server hardware and server operating systems. An

72% majority of corporations now require “four nines” or 99.99% minimum uptime. Businesses must regularly replace, retrofit and refresh their server hardware and server operating systems with the necessary patches, updates and security fixes *as needed* to maintain system health. The onus is also on the server hardware and server operating system vendors to provide realistic recommendations for system configurations to achieve optimal performance. Vendors also bear the responsibility to deliver patches, fixes and updates in a timely manner and to inform customers to the best of their ability regarding any known incompatibility issues that may potentially impact performance. Vendors should also be honest with customers in the event there is a problem or delay with delivering replacement parts.

Time *is* money. Even a few minutes of downtime can result in significant costs and cause internal business operations to come to a standstill. Downtime can also impact adversely a company’s relationship with its customers, business suppliers and partners. Reliability or lack thereof can potentially damage a company’s reputation and result in lost business.

Recommendations

ITIC strongly advises organizations to regularly measure the uptime and reliability of their main LOB server hardware, server operating systems and applications. Being cognizant of specific uptime and reliability statistics will enable the business and its IT department to identify baseline metrics associated with all of their individual platforms. It will also provide companies with an accurate assessment of the inherent reliability and flaws in their hardware and software. They can then compare and contrast that with downtime resulting from other issues such as: integration and interoperability; lack of readily available patches or fixes; problems with ISPs and carriers and unpredictable or unavoidable outages due to natural or manmade disasters.

This in turn, provides businesses with a mechanism to accurately assess the amount of downtime and its subsequent impact on business operations, the IT department and the productivity of its end users. The ability to measure reliability also helps organizations gauge how downtime affects external business partners, customers and suppliers.

To optimize uptime and reliability, ITIC advises corporations to:

- **Regularly analyze and review configurations, usage and performance levels.** This will enable companies to determine whether or not their current server and server OS environment allows them to achieve optimal reliability.
- **Don’t Wait to Update; refresh and upgrade the Server Hardware as needed to accommodate more data intensive and virtualized workloads.** The server hardware (standalone, blade, cluster, etc.) and the server operating system are inextricably linked. To achieve optimal performance from both components, corporations must ensure that

the server hardware is robust enough to carry both the current and anticipated workloads. Applications are getting larger. The number and percentage of virtualized servers continues to increase. Virtual servers hosting multiple instances of mainstream LOB business critical applications demands robust hardware. Organizations should purchase the beefiest server configuration their budgets will allow. Waiting four, five or six years to refresh servers while placing greater demands on the hardware, is asking for trouble.

- **Adopt formal SLAs.** Service level agreements enable organizations to define acceptable performance metrics. Companies should meet with their vendors and customers on at least an annual basis to ensure the terms are met.
- **Define measure and monitor reliability and performance metrics.** It is imperative that companies measure component, system, server hardware, server OS and desktop and server OS, security, network infrastructure, storage and application performance. Maintain records on the amount of planned and unplanned downtime.
- **Regularly track server and server OS reliability and downtime.** Keep accurate records of outages and their causes. Segment the outages according to their severity and length – e.g., Tier 1, Tier 2 and Tier 3. The appropriate IT managers should also keep detailed logs of remediation efforts in the event of the outage. These logs should include a full account of remediation activities, specifying how the problem was solved, how long it took and what staff members participated in the event. It should also list the monetary costs as well as any material impact on the business, its operations and its end users. This will prove invaluable resource should the problem recur. It may also make the difference in containing or curtailing the reliability-related incident, saving precious time for the IT department, the end users and corporate customers.
- **Calculate the cost of unplanned downtime.** Companies should determine the average cost of minor Tier 1 outages. They should also keep more detailed cost assessments of the more serious unplanned Tier 2 and Tier 3 incidents. It's essential for businesses to know the monetary amount of each outage – including IT and end user salaries due to troubleshooting and any lost productivity – as well as the impact on the business. C-level executives and IT managers should also pay close attention to whether or not the company's reputation suffered as a result of a reliability incident; did any litigation ensue; were customers, business partners and suppliers impacted (and at what cost) and at least try and gauge whether or not the company lost business or potential business.
- **Compile a list of best practices.** Chief technology officers (CTOs), Chief Data Officers (CDOs), software developers, engineers, network administrators and managers should have extensive familiarity with the products they currently use and are considering. Check and adhere to your vendors' list of approved, compatible hardware, software and applications.
- **Train and certify IT administrators.** Don't be penny wise and pound foolish. If your company can't afford the expense or time to certify the entire IT department, designate the most experienced or appropriate member of the team to get trained and let that person train the other IT staffers. Also enlist the aid of your server hardware and server OS

vendors and engage the services of third party systems integrators and consultants when appropriate.

- **Perform regular asset management testing.** Schedule asset management reviews on a yearly, bi-annual or quarterly basis, as needed. This will assist your company in remaining current on hardware and software and help you to adhere to the terms and conditions of licensing contracts. All of these issues influence network reliability.

Survey Methodology

ITIC's 2016 - 2017 Global Server Hardware and Server OS Reliability Survey, polled C-level executives and IT managers at over 750 corporations worldwide from August through October 2016. The independent Web-based survey included multiple choice questions and one Essay question. In order to ensure objectivity, ITIC accepted no vendor sponsorship and none of the participants received any remuneration. ITIC analysts also conducted two dozen first person customer interviews to validate or repudiate the Web survey responses and obtain anecdotal data. The anecdotal data provides broad and in-depth insight into the business and technology challenges confronting corporations in both the immediate and long term. ITIC employed authentication and tracking mechanisms to prevent tampering and to prohibit multiple responses by the same parties.

Survey Demographics

Companies of all sizes and all vertical markets were represented in the survey. Respondents came from companies ranging from small and medium businesses (SMBs) with fewer than 50 workers, to large enterprises with more than 100,000 employees.

All market sectors were equally represented: SMBs with one to 100 employees accounted for 33% of the respondents. Small and medium enterprises (SMEs) with 101 to 1,000 workers represented 30% of the participants and the remaining 37% of respondents came from large enterprises with 1,001 to over 100,000 employees. Survey respondents hailed from 49 different vertical markets. Approximately 84% of respondents hailed from North America; 16% were international customers who hailed from more than 20 countries throughout Europe, Asia, Australia, New Zealand, South America and Africa.

Appendices

This section contains a list and links to the various ITIC statistics and surveys cited in this Report.

ITIC Website and links to survey data and blog posts:

<http://itic-corp.com/blog/2016/08/cost-of-hourly-downtime-soars-81-of-enterprises-say-it-exceeds-300k-on-average/>

<http://itic-corp.com/blog/2016/02/ibm-z13s-delivers-power-performance-fault-tolerant-reliability-and-security-for-hybrid-clouds/>

<http://itic-corp.com/blog/2016/01/ibm-lenovo-top-itic-2016-reliability-poll-cisco-comes-on-strong/>

<http://itic-corp.com/blog/2015/02/iticknowbe4-security-survey-56-of-corporations-have-no-proactive-response-plan-to-deal-with-byod-security-hacks/>

<http://itic-corp.com/blog/2013/07/one-hour-of-downtime-costs-100k-for-95-of-enterprises/>

<http://itic-corp.com/blog/2011/04/itic-2011-reliability-survey-users-give-ibm-aix-v7-windows-server-2008-r2-highest-security-marks/>

<http://itic-corp.com/blog/2011/02/itic-reliability-survey-oracle-users-anxiousangry-over-service-support-slippage/>