

INFORMATION TECHNOLOGY INTELLIGENCE CONSULTING

Information Technology Intelligence Consulting



ITIC 2015 - 2016 Global Server Hardware, Server OS Reliability Report

July 2015

Executive Summary

IBM Power Systems Servers Most Reliable for Seventh Straight Year; Lenovo x86 Servers Deliver Highest Uptime/Availability among all Intel x86-based Systems; Cisco UCS Stays Strong; Dell Reliability Ratchets Up; Intel Xeon Processor E7 v3 chips incorporate advanced analytics; significantly boost reliability of x86-based servers

In 2015, infrastructure reliability is more essential than ever.

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 are the cornerstones 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.

For the seventh 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 61% majority of IBM Power servers and Lenovo System x servers achieved “five nines” or 99.999% availability – the equivalent of 5.25 minutes of unplanned per server downtime compared to 46% of Hewlett-Packard servers and 40% of Oracle server hardware (See Exhibit 1).

Those are the results of the latest independent ITIC 2015 - 2016 Global Server Hardware and Server OS Reliability Survey which polled C-level executives and IT managers at over 600 organizations worldwide during June and July 2015. ITIC’s seventh 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. However, the results also reveal that external issues, most notably human error, security breaches and Bring Your Own Device (BYOD) are playing a more pivotal and direct role in solidifying or undermining reliability.

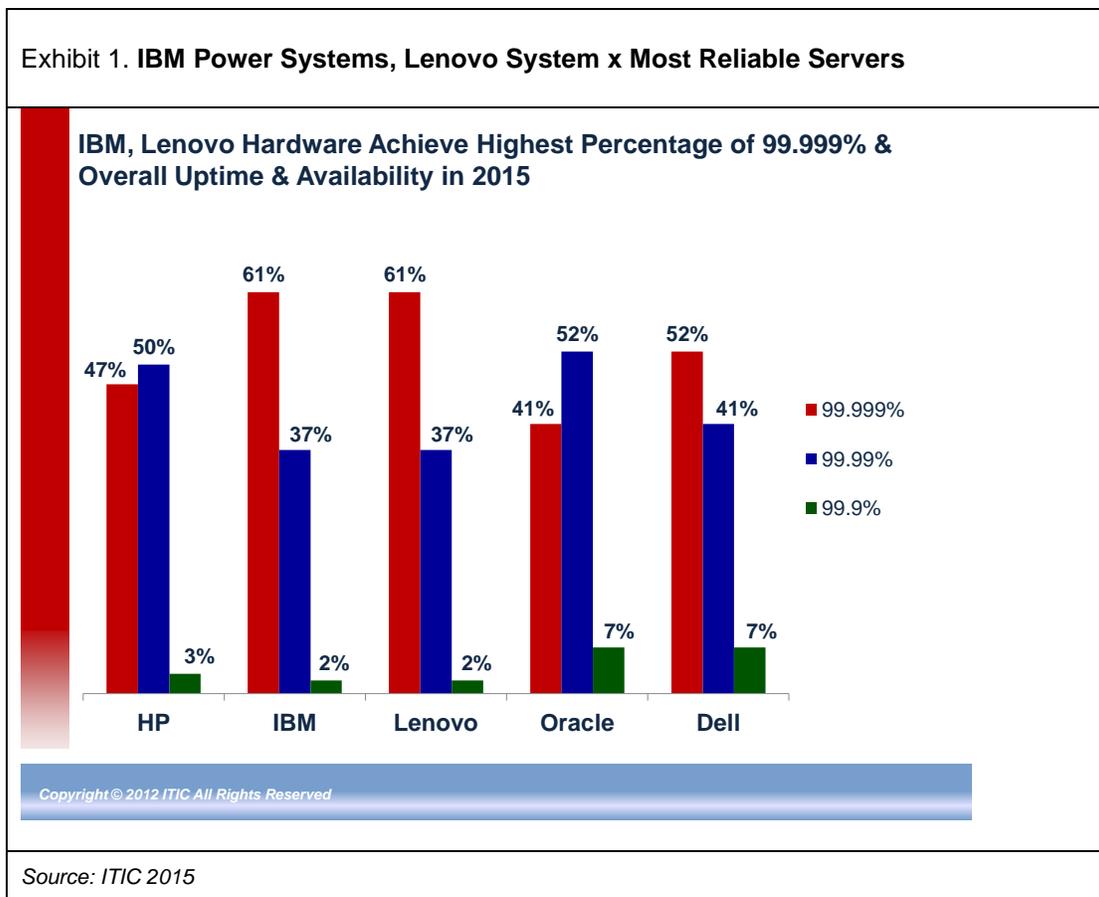
Overall, the inherent reliability of the majority of server hardware platforms, server operating systems and the underlying processor technology continues to improve year over year.

Among the top survey findings:

- **IBM, Lenovo, Cisco UCS, Dell and Fujitsu** score highest reliability
- **Cisco, Dell, IBM, Lenovo** rated the highest in customer satisfaction.

- **Dell hardware exhibited** biggest year over year (YoY) improvements; reliability increased seven percent from 2014.
- **Red Hat Enterprise Linux (RHEL) & Ubuntu posted** the highest reliability scores among the Linux OS distributions for the second 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** are the most reliable virtualization offerings.

The survey also showed that the three **technology** issues of most concern according to this year's survey are: Security, Disaster Recovery and Backup and Business Continuity.



Introduction

This Report examines and compares the reliability of 14 major server platforms, 18 server operating system distributions and 11 server hardware virtualization layers. It also explores the underlying issues that positively or negatively influence the inherent reliability of servers and operating systems. And it delves into the external elements that directly impact 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 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% availability is unacceptable and unthinkable.

Corporations are increasingly deploying virtualization, cloud computing and move to an Internet of Things (IoT) environment that is built on interconnected devices embedded with sensors to deliver business intelligence and Big Data Analytics. In a world where everything is interconnected, the old adage: “A chain is only as strong as its weakest link,” has never rung truer. In virtualized and IoT environments, *the potential* for collateral damage has increased by orders of magnitudes. A virtual server running multiple instances of a crucial main line of business (LOB) application will have a greater impact on employee productivity, operations and the corporate bottom line in the event of a service interruption than a server running a single 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.

Corporations, their end users, customers, business partners and suppliers expect 24 x 7 access to data resources irrespective of time or location. Business is conducted round-the-clock in today’s Digital Age and global economy. Even a few minutes of downtime is an anathema that can disrupt productivity and cost businesses tens of thousands to millions in monetary and business losses.

The latest ITIC 2015 - 2016 Reliability survey data found that over three-quarters of respondents –71% – now consider 99.99 % to be the minimum acceptable levels 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 found businesses require a minimum of “four nines” availability.

As ITIC has done every year since 2008, we have published the specific Table that depicts the availability percentages and the equivalent number of annual, monthly and weekly hours and minutes of per server/per annum. **Table 1** below provides a quick and eye opening reference for any organization or IT department to calculate downtime and reliability and the business and monetary impact on their corporation.

The accepted metrics of three or four nines of uptime – 99.9%, 99.95% and 99.99%, – equate to 8.76 hours; 4.38 hours and 52.56 minutes of per server, per annum downtime, respectively, (**See Table 1**). Additionally, the most recent ITIC poll shows that 25% of the over 600 respondent organizations require a very high 99.999% or better degree of reliability; 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 in the last 12 months. In the ITIC 2013 Reliability survey only 11% of participants said their firms required “five nines” reliability. The 2015 – 2015 poll also found that five percent of corporations require *in excess* of 99.999% reliability – mere seconds of downtime.

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

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 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.
- **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. It frequently disrupts network operations for the company's end users and possibly 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 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 few minutes can quickly disrupt network operations, end user productivity and adversely impact business operations. The result of a Tier 2 outage of one-to-four hours or a severe Tier 3 outage of four hours or more, can cause significant business losses, damage the company's reputation and raise the risk of litigation. Enterprises in vertical markets such as banking and finance, stock exchanges, insurance, healthcare and legal, whose businesses are based on intensive data transactions, can lose millions if service is interrupted for even five, 10 or 15 minutes. The consequences can be catastrophic if the outage occurs during peak usage or during a crucial transaction. Small and mid-sized businesses (SMBs) and mid-sized enterprises (SMEs) are also vulnerable and equally as risk-averse as 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.

Survey Highlights

- IBM System z Enterprise had the lowest incidence – 0% -- of > 4 hours of per server/per annum downtime of any major mainstream hardware platform.
- IBM and Lenovo hardware running the Red Hat Enterprise Linux (RHEL), 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. Dell PowerEdge Server reliability notched a measurable increase with only six percent of the Dell servers experiencing in excess of 4 hours of downtime versus 13% in the 2014 survey. 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.
- Reliability of the HP ProLiant and Oracle SPARC & x86 hardware and Solaris OS remains spotty mainly due to aging and protracted product refresh cycles. Some Oracle hardware customers are delaying upgrades or switching to other platforms.
- 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 2014 poll.

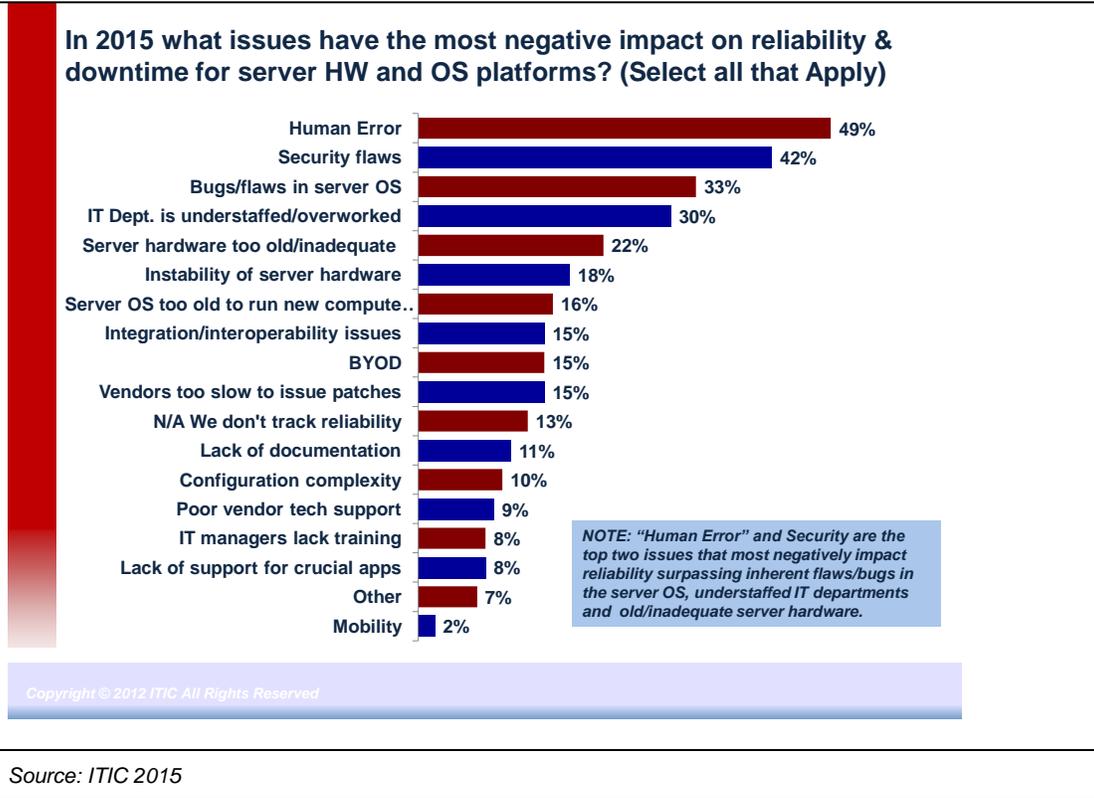
Human Error, Security Undermine Reliability

To reiterate, ITIC’s survey results indicate that the *inherent* reliability and uptime of nearly all 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 2015 Reliability research reveals that a variety of external factors are having more of a direct impact on system downtime and overall availability. As **Exhibit 2** below illustrates, these include human error; overworked and understaffed IT departments; the rapid mainstream adoption of complex new technologies such as virtualization and increasing cloud computing deployments and an explosion of Bring Your Own Device (BYOD) technology.

2015 Top Reliability Trends:

- A 71% **majority** of corporations now require a minimum of 99.99% uptime for mission critical hardware, operating systems and main line of business (LOB) applications, up from 49% in 2014-2015 survey and an increase of 32% from the 17% of respondents who indicated their firms required “four nines” minimum reliability in ITIC’s 2012-2013 survey.
- **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.

Exhibit 2 Five out of 10 Enterprises Rank Human Error as the Top Cause of Downtime, Security Issues is the Second Biggest Cause of Outages



Top Issues Negatively impacting network reliability:

- Human Error – **49% vs. 44% in the 2014 poll**
- Flaws in the server operating system – **33%**
- Understaffed IT departments or inadequately trained administrators – **33%**
- Hardware problems – (e.g. aging hardware; trouble getting replacement parts) was cited by **26%**

Most significantly, nearly half - 49% - of survey respondents cited “human error” as the Number One issue negatively impacting reliability (See **Exhibit 2**). Other top trouble sources include: bugs and flaws in the server OS; understaffed and overworked IT departments and aging hardware that is too old or inadequate to support current compute-intensive physical and virtual workloads. Inherent instability of server hardware was cited by 18% as undercutting reliability.

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.

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.

On the opposite end of the spectrum, while Hewlett-Packard received high marks for customer satisfaction, some 10% HP ProLiant model servers recorded unplanned downtime exceeding four (4) hours. The results were noticeably better for HP's Integrity line; only three percent recorded more than four (4) hours of unplanned per server/per annum downtime.

Additionally, HP users' anecdotal essay comments cited growing concerns/questions surrounding HP's imminent split into two organizations. On November 1, 2015, HP will officially become two separate entities. HP Enterprise will be organized into five separate segments for financial reporting purposes: the Enterprise Group, Enterprise Services, Software, Financial Services, and Corporate Investments. The second organization, HP Inc. will focus more on the consumer segment and sell printers and personal computers. While survey respondents are still generally satisfied with the overall reliability of HP servers, they are understandably nervous on how the split into two companies will impact HP Enterprise's ability to deliver innovation and the top notch technical service and support they've grown accustomed to over the last six decades. Additionally, as HP's commodity businesses have come under increasing pressure, there has been significant tumult and changes at the executive management level. The company has also continued to shave its headcount. In the last 10 years, HP has cut an estimated 100,000 workers -- more than any other high tech firm as a means to maintain profitability. Since taking over as chief executive in 2011, Meg Whitman has overseen a dramatic downsizing that has eliminated about 55,000 jobs thus far. And more layoffs are imminent. In September, HP announced it will chop another 25,000 to 30,000 employees; this equates to approximately 10% of its total workforce. HP said that the layoffs will save the company \$2.7 billion annually. Ironically however, one of the hardest hit groups is HP's Enterprise Services team; Ms. Whitman claims that services are still a strategic priority for the struggling vendor. In this context, the inherent reliability of HP's systems will be more crucial than ever. And although HP servers and services have long been a staple in corporate enterprises, survey respondents indicated they are nervous about HP's future direction in light of the continuing job cuts and the company's plans to split into two separate entities in November 2015. Specifically, in anecdotal essay comments and first person interviews, IT managers noted that HP's formerly stellar technical service and support is slipping. IT administrators disclosed they are spending longer time on hold waiting to reach a live technician and oftentimes when they finally do speak to tech support they are compelled to escalate the call. HP users also complained of extended lead times to obtain replacement parts.

The results were more upbeat for Cisco Systems, Inc.'s Cisco Unified Computing System (UCS) servers. Appearing for the second 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.

On the server hardware side 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:

- **Some 30% of businesses** spend more than 1 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.”

Processor Technology

The servers are only as good as their parts. The memory and CPU must keep pace. To reiterate, any company that expects to achieve optimal server uptime and availability should 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 and Advanced Micro Devices (AMD) continually update their processors.

The Intel Xeon Processor E7 v2 and the latest E7 v3 for example, are specifically geared towards reliability. Both the E7 v2 chips and the E7 v3 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 and v3 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 and v3 chips also feature Intel's *Resilient Memory Technologies*, to solidify and ensure data integrity within the memory subsystem.

The Intel processor E7-8800/4800 v3 product families debuted in May 2015. They provide accelerated business insight through real-time analytics and enhanced performance and reliability for mission-critical computing. Intel claims that the RAS enhancements in the E7 v3 family 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 processors assist customers to sift through the data deluge and extract actionable insights from these massive volumes of data. The Intel Xeon processor E7 v3 family helps customers 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 up to 6 x improvements in business processing application performance for in-memory transactional workloads and is optimized with the new Intel® Transactional Synchronization Extensions (Intel® TSX). The E7 v3 family also supports up to 18 cores; this is a 20% increase in cores compared to the prior generation. It also has up to 45 megabytes of last-level cache, which enables it to deliver up to 70% percent more decision support analytic sessions per hour, according to Intel. The Intel Xeon Processor E7 v2 and v3 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 second year in a row the 2015 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 saw a 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 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 v2 and v3 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 46% of respondents rated it "Excellent," 37% rated it "Very Good" and 17% 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 v2 and v3 Processors received similarly high performance and reliability grades. An 82% majority achieved a minimum of 99.99%; 48% achieved 99.999% and six 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 and v3 families can deliver comparable native performance, reliability and uptime when installed on robust servers that are new or up to 3 years old and when corporation right-size the server configurations 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 2015 – 2016 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.

On a positive note the reliability of Dell PowerEdge servers increased perceptibly in this year's survey – with only six percent of survey respondents stating their Dell servers recorded per server/per annum outages in excess of four hours. This is less than half of the 13% of Dell users who reported protracted server downtime exceeding four hours in the 2014 Reliability poll. Not surprisingly, over one-third of the Dell server respondents in this year's poll said they had upgraded or retrofitted their servers within the past 12 months.

By contrast, only 20% Lenovo System x or IBM Power Systems users and 24% of Cisco UCS users retained their servers for four or more years without upgrading or retrofitting them. It is a testament to the technical elegance of the former Sun Microsystems platforms – now the Oracle x86 and Oracle SPARC servers -- that “only” 10% percent and eight percent of those platforms respectively, experienced over four hours of per server/per annum unplanned downtime. This is despite the fact that 55% of those users reported keeping the Oracle servers for four, five or six years without upgrading.

The high end enterprise mainframe IBM System z Enterprise is a different breed of server and in a class of its. The System z is more hardened and robust and it is typically managed and maintained by highly experienced administrators.

For organizations that opt to retain their server hardware for over four years and test the limits of its capacity, 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.

Top notch 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 first person interviews 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 zEnterprise servers, while Lenovo intends to make a success of the x86 platform.

Unlike its chief rivals, which have been battered by management upheavals (HP) or in transition due to mergers and acquisitions (Oracle's purchase of Sun Microsystems), IBM under chief executive and president Virginia Rometty is a bastion of stability. This stability extends to Big Blues 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 all of the above, 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. In 2014, IBM 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.

In addition, IBM Power Systems and Lenovo 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 .

In the semiconductor sector, Intel outspent all of its rivals on R&D. In 2014, Intel's R&D spending reached a record-high \$10.6 billion and the world's number one chipmaker remains on track with its 2015 R&D spending as well. Intel's R&D expenditures were triple that of its closest competitors Qualcomm which spent nearly \$3.4 billion on R&D and Samsung which spent \$2.8 billion on R&D.

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 which are widely used in Lenovo System x, Dell PowerEdge, HP ProLiant and Oracle systems. To reiterate, The E7 v3 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

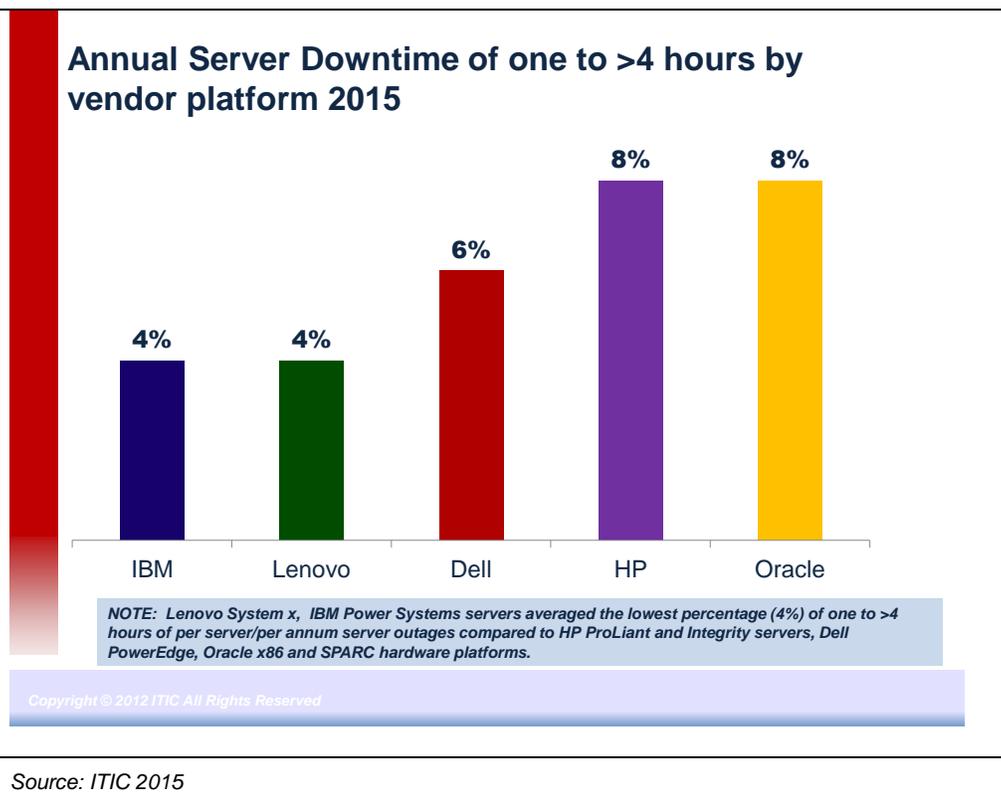
It has been a full year since IBM completed the sale of its x86 server business to Lenovo in a deal that was valued at \$2.3 billion.

The purchase is Lenovo's latest since it overtook HP as the world's top PC maker in 2014. In the ensuing 13 months, 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, Lenovo Chief Executive Yang Yuanqing said the IBM deal represented a new "growth engine" for his company. Yuanqing estimates that the Lenovo x86 Business Unit will net \$5 billion in revenue during 2015, its first full year of operations and deliver profit margins higher than the four percent of Lenovo's PC business.

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 3** shows, Lenovo System x servers retain the same high levels of reliability they delivered while still under the IBM brand. Additionally,

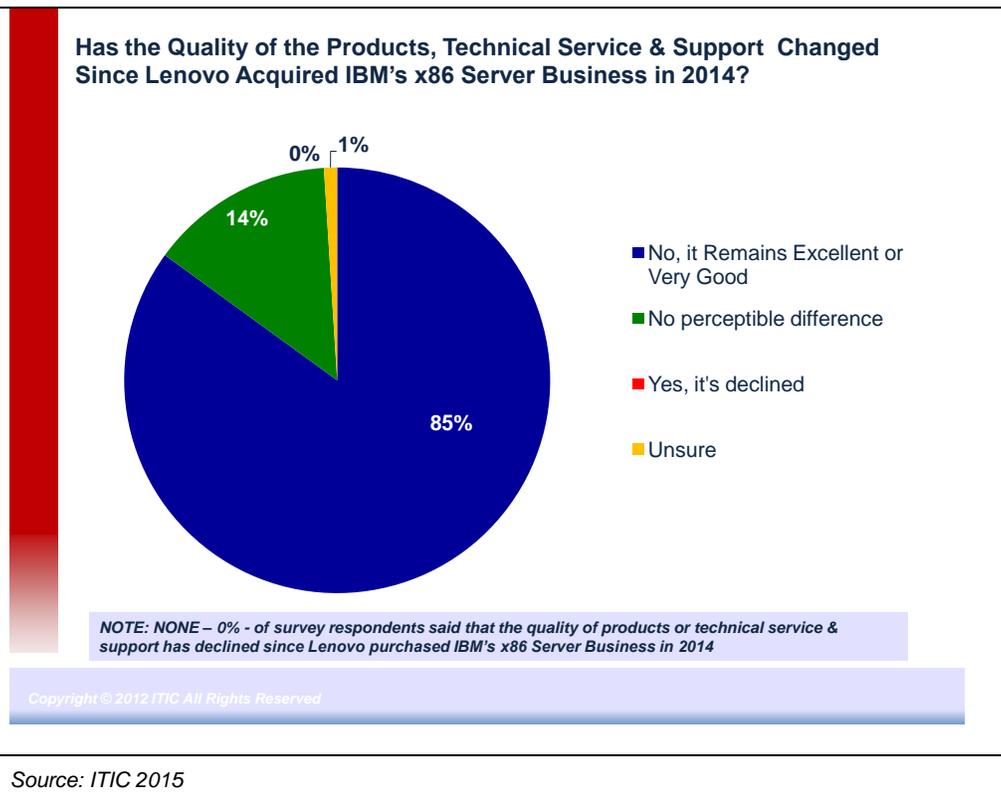
Exhibit 3. IBM, Lenovo Hardware Record Least Amount of One to Over Four Hours of Annual Server Downtime



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.

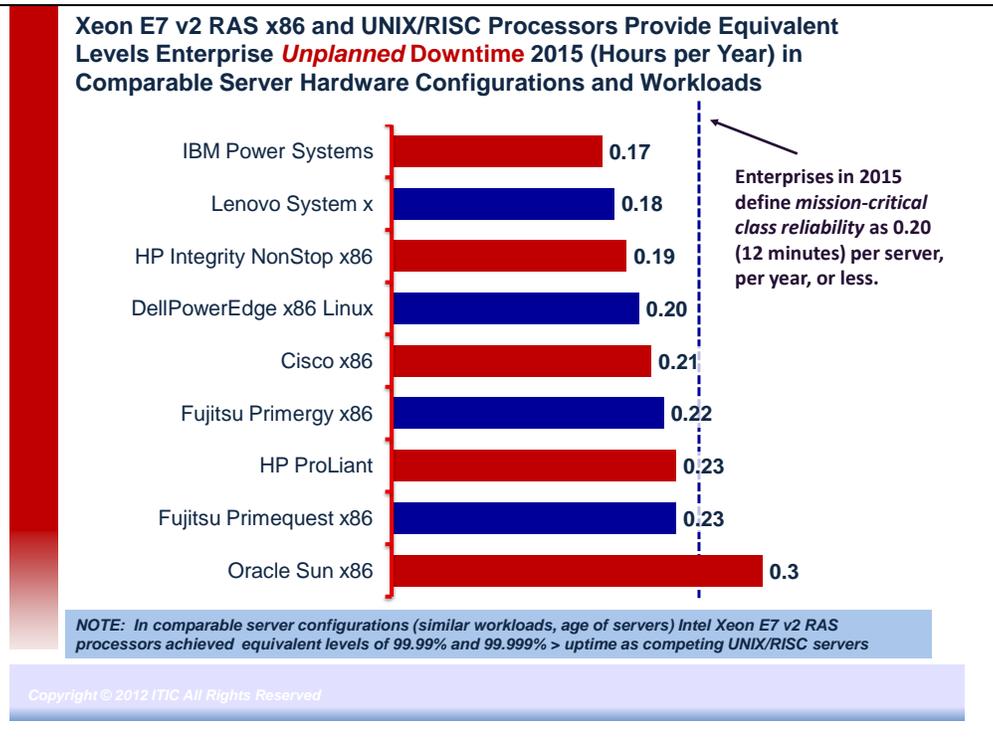
ITIC’s 2015 Reliability Survey also queried customers on how the System x servers were faring in terms of product quality and the responsiveness and quality of after-market technical service and support. The results show that an overwhelming 99% of respondents rated service and support, product quality as excellent and had experience no perceptible differences.

Exhibit 4. Lenovo System x Servers Maintain High Reliability, Excellent Technical Service and Support Record



As **Exhibit 5** illustrates, 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.

Exhibit 5. Intel Xeon E7 v2 RAS Processors Deliver Equivalent Level of 99.99% & 99.999% > Reliability as UNIX/RISC Servers in Comparable Configurations



Source: ITIC 2015

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.

Exhibit 6. IBM Power Systems, Lenovo System x Servers Experience Least Amount of Unplanned Annual Downtime

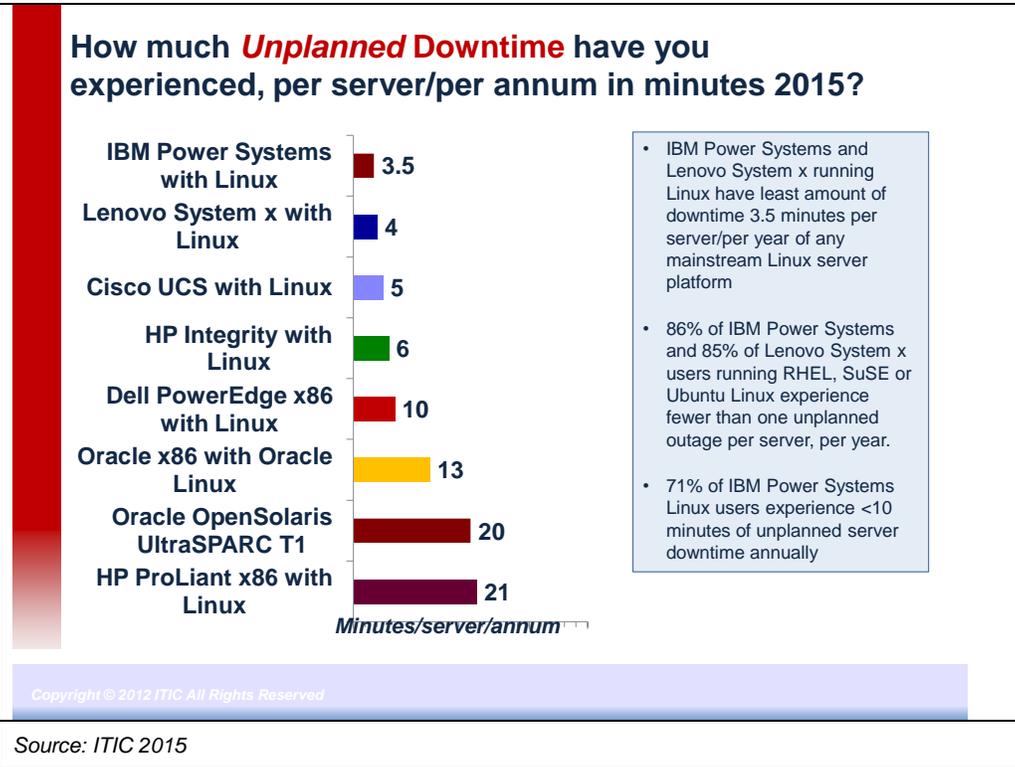
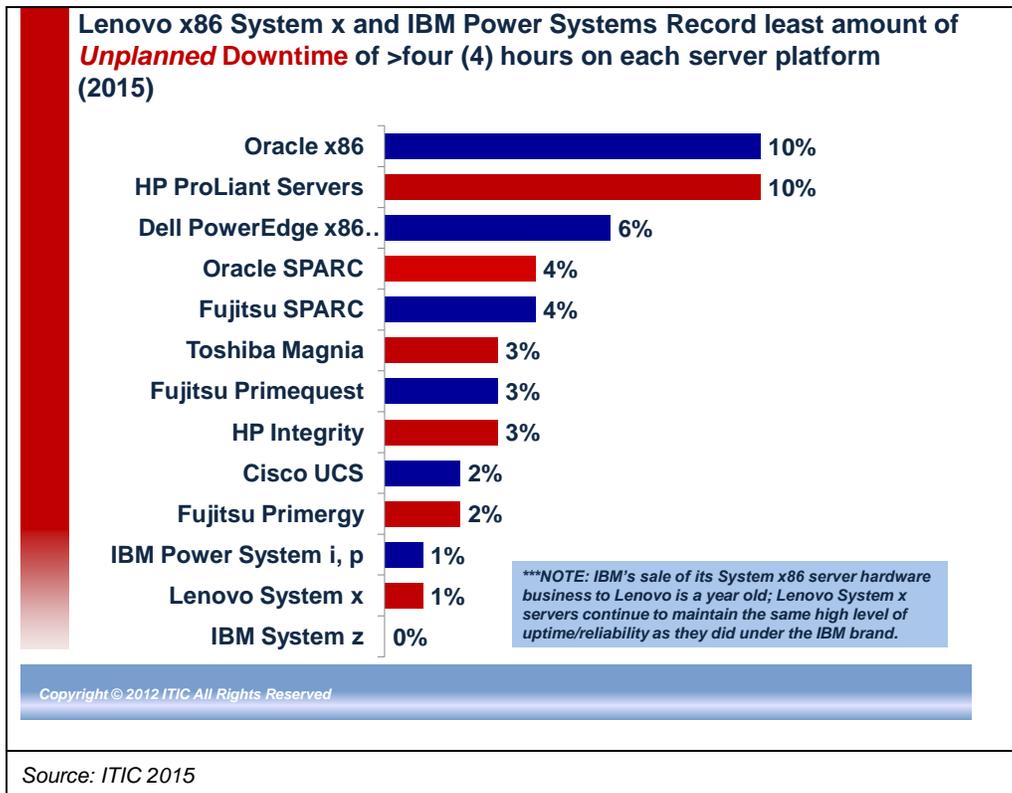


Exhibit 7. IBM System x Servers Register Least Amount of Prolonged Unplanned Annual Downtime Exceeding Four (>4) Hours of x86 Servers



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 gave Cisco UCS 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. Earlier this year when longtime CEO John Chambers stepped down (he’s still Executive Chairman), Chuck Robbins stepped in to fill his shoes. 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 Everything (IoE) as Cisco calls it. In April, 2014, Cisco Systems launched a \$150 Million initiative to fund IoE startups. 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.

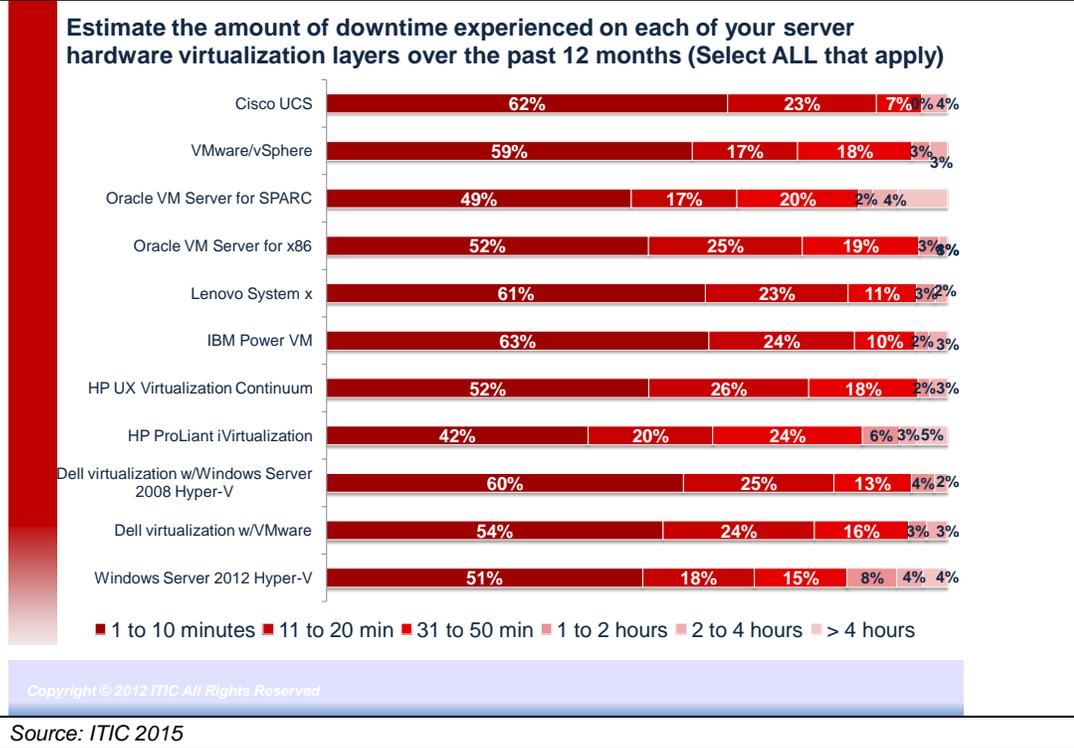
Server Virtualization Reliability

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 rightsize 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 2008 R2 Hyper-V were the top three most reliable hardware virtualization platforms (see **Exhibit 8**). 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.

Exhibit 8. IBM Power VM, Cisco UCS and Dell Score High Marks for Virtualization Reliability



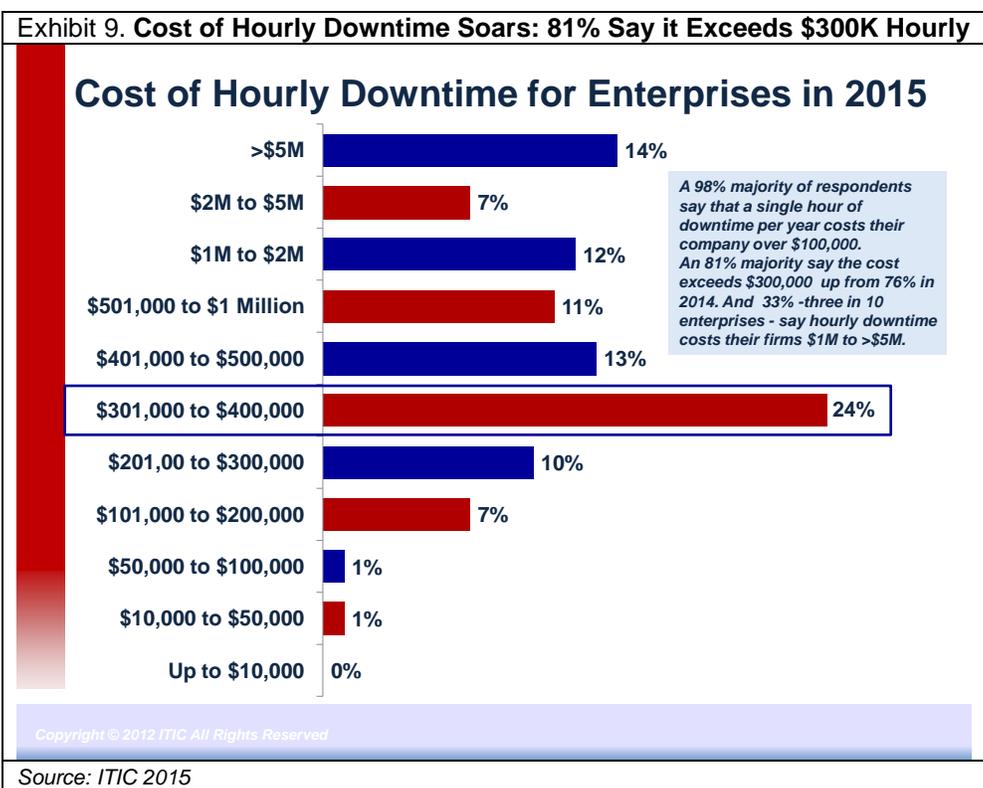
Cost of Hourly Downtime Soars

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 9** 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 2015 - 2015 Global Server Hardware and Server OS Reliability Survey findings indicate that all of the server operating system platforms have achieved a high degree of reliability. However, the IBM AIX 7.1 operating system, Ubuntu 12.04, HP UX 11i v3 and Red Hat Enterprise Linux v 6.5 are the top four most reliable server OS distributions. Server hardware reliability shows similar year-over-year improvements. 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 advances in technology suggested by the ITIC 2015 -2015 Global Server Hardware and OS Reliability Survey results are encouraging. But advanced technology is just one component, albeit a crucial factor in fortifying and maintaining reliability.

Vendors must deliver reliable products and top notch technical service and support. Corporations also bear responsibility for keeping their IT departments well staffed and providing the necessary training and certification to IT administrators. 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, 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 81% majority of corporations now require "four nines" or 99.99% minimum uptime. Businesses should 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.
- **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 *2015 - 2015 Global Server Hardware and Server OS Reliability Survey*, polled C-level executives and IT managers at over 550 corporations worldwide throughout June and July 2015. 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/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/>