

INFORMATION TECHNOLOGY
INTELLIGENCE CONSULTING

Information Technology Intelligence Consulting



ITIC 2017 – 2018 Global Server Hardware,
Server OS Reliability Report

December 2017

Table of Contents

Executive Summary.....3

 Introduction5

 Data & Analysis5

 Survey Highlights.....7

 Reliability and Uptime by the Numbers.....8

 Minimum Reliability Requirements Rise10

 2017 – 2018 Reliability Trends.....12

 IBM Rock Solid Reliability14

 Lenovo System x Servers Best in Class Reliability.....17

 Huawei KunLun, Fujitsu PRIMEQUEST E Series Mission Critical Platforms Notch Strong Reliability Ratings20

 Huawei KunLun Supports SUSE Linux and SAP HANA22

 Processor Technology22

 Cisco UCS Comes On Strong, HP Reliability Rebounds.....26

 Cost of Hourly Downtime Soars26

 Conclusions and Recommendations28

 Recommendations29

 Survey Methodology31

 Survey Demographics31

 Appendices.....32

Executive Summary

IBM, Lenovo servers record highest Reliability,

IBM Power Systems, Lenovo x86 X6 Servers Deliver up to 18x more uptime and reliability than some rivals

Cisco UCS, Fujitsu PRIMEQUEST, HPE Integrity and Huawei Kun Lun also record strong reliability numbers

Reliability declines on Dell, HPE and Oracle x86 servers > 3 years old

Human Error, Security and Complexity are Top Three Causes of Downtime

For the tenth straight year, IBM and Lenovo servers again achieved top rankings in ITIC's 2017 – 2018 Global Server Hardware and Server OS Reliability survey.

IBM's Z Systems Enterprise server is in a class of its own. The IBM mainframe continues to exhibit peerless reliability besting all competitors. The Z recorded less than 10 seconds of unplanned per server downtime each month. Additionally less than one-half of one percent of all IBM Z customers reported unplanned outages that totaled greater than four (4) hours of system downtime in a single year.

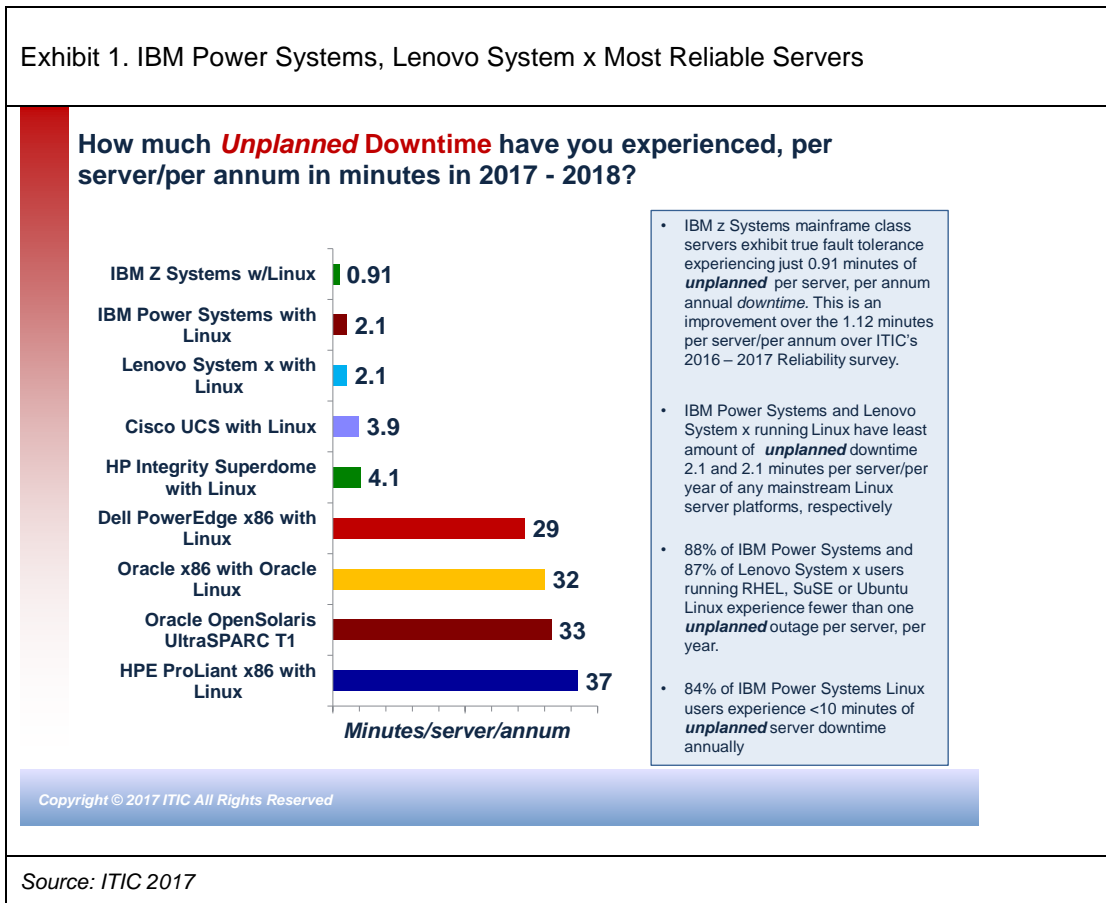
Among mainstream servers, IBM Power Systems 7 and 8 and the Lenovo x86 X6 mission critical hardware consistently deliver the highest levels of reliability/uptime among 14 server hardware and 11 different mainstream server hardware virtualization platforms. Each platform averaged just 2.1 minutes of *unplanned* per annum/per server downtime (See Exhibit 1).

That makes the IBM Power Systems and Lenovo x 86 servers approximately 17 to 18 times more stable and available, than the least reliable distributions – the rival Oracle and HPE ProLiant servers.

Additionally, the latest ITIC survey results indicate just one percent of IBM Power Systems and Lenovo System x servers experienced over four (4) hours of unplanned annual downtime. This is the best showing among the 14 different server platforms surveyed (**See Exhibit 2**).

ITIC's 10th annual independent ITIC 2017 - 2018 Global Server Hardware and Server OS Reliability survey polled 800 organizations worldwide from August through December 2017. In order to obtain the most accurate and unbiased results, ITIC accepted no vendor sponsorship.

Exhibit 1. Server Reliability by Hardware Platform



The reliability of the core server hardware and server operating system (OS) platforms is the bedrock upon which the business applications and daily operations rest. The reliability of these key infrastructure components are of pivotal importance as corporations scale the size, scope and complexity of their networks and applications. This is especially crucial as corporate enterprises move more operations to virtualized and cloud environments. Corporations are also ever-more dependent for their servers and OSes to deliver top notch reliability and uptime. The servers and operating systems now support more compute-intensive workloads and applications like Big Data Analytics, cognitive computing and mobility. The network environment now spans on-premises datacenters, private, public and hybrid clouds to geographically dispersed deployments located thousands of miles away at the network edge or perimeter.

Daily business operations now demand that server hardware and operating systems deliver near-perfect reliability. High availability ensures uninterrupted productivity; supports the business' bottom line; strengthens security and compliance and mitigates risk.

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 21st Century Digital Age are always-on. Businesses require their mission critical applications to deliver reliable, uninterrupted access to data in real-time, 24 hours a day, seven

days a week. Corporations conduct business globally irrespective of time zones or the location of their increasingly mobile and remote workforces.

Organizations whose server hardware, operating system and virtualization components fail to deliver a minimum of “four nines” of availability, will feel the negative impact immediately on their bottom line and daily business operations in a variety of ways. These include:

- Productivity disruptions to internal knowledge workers and IT administrators
- Potential for lost, damaged or destroyed data
- Revenue losses
- Failure to complete key transactions within a specified time period
- Failure to meet Regulatory Compliance and legal requirements
- Failure to meet Service Level Agreements (SLAs)
- Damage to the corporate brand
- Increased security risks

All of these issues make it imperative that organizations have server hardware, operating systems and virtualization solutions that can deliver rock solid reliability, uptime and uninterrupted access to mission critical applications and data resources.

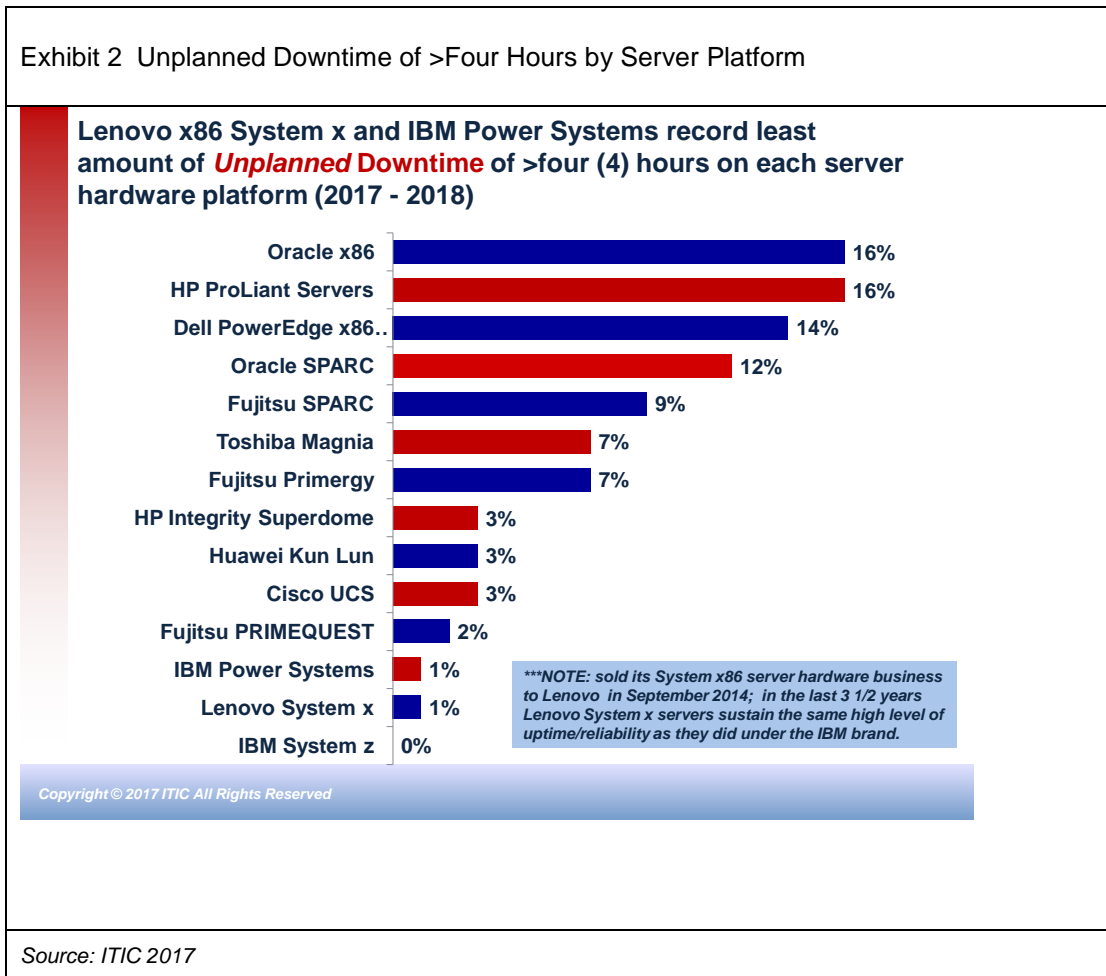
As Exhibit 1 above also illustrates the widening availability chasm among the most reliable mainstream server platforms – IBM and Lenovo – which averaged 2.1 minutes each per server/per annum and the least reliable – HPE’s ProLiant x 86 hardware which experienced approximately 37 minutes of unplanned per server downtime annually.

IBM’s Z mainframe platform continues to best all platforms, recording a near-perfect reliability score of less than 10 seconds of unplanned downtime each month.

These results mean that 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.

As Exhibit 2 shows, the IBM Z (0%), IBM Power Systems and Lenovo x86 platforms (1% each) also had the lowest incidents of unplanned downtime of greater than four hours. By contrast, some 16% of HPE ProLiant and Dell PowerEdge servers recorded over four hours each of unanticipated per server/per annum downtime.

Exhibit 2. Unplanned Downtime of >Four Hours by Server Platform



Survey Highlights

Among the other key survey highlights:

- IBM Z Systems Enterprise mainframe class servers topped all platforms for reliability, accessibility, performance, and security among. The Z servers deliver true mainframe fault tolerance experiencing just 0.91 - less than one minute of *unplanned* per server, per annum annual downtime. That equates to a barely perceptible 7.6 seconds per month or “blink and you miss it,” 1.8 seconds of weekly unplanned downtime.
- The z Systems family of servers also the lowest incidence – 0% — of > 4 hours of per server/per annum downtime of any hardware platform.

- To reiterate, **IBM Power Systems and Lenovo System x** servers once again consistently demonstrated the least amount of *unplanned* downtime 2.1 minutes per server/per year among all mainstream server offerings. Both the IBM and Lenovo offerings shaved several seconds off their previous results of 2.5 and 2.8 minutes per server/per year, respectively in ITIC’s 2017 Mid-year Reliability poll in June. These latest results underscore the stability and consistency of the IBM and Lenovo distributions running Linux as well as top notch service and support. Additionally, both vendors are aggressively updating their platforms to support Data Analytics, Artificial Intelligence (AI), Machine Learning and stronger embedded security and management capabilities.
- And 88% of IBM Power Systems and 87% of Lenovo System x users running RHEL, SUSE or Ubuntu Linux experience fewer than one *unplanned* outage per server/per year.
- Only one percent of IBM and Lenovo servers recorded over four (4) hours of unplanned per server/per annum downtime; followed by six percent of HPE servers; eight percent of Dell servers and 10% of Oracle servers (**See Exhibit 2**).
- IBM and Lenovo hardware and the Linux operating system distributions were either first or second in every reliability category, including virtualization and security.
- **Fujitsu PRIMEQUEST and newcomer Huawei’s Kun Lun** server platforms both scored high reliability rankings. It’s noteworthy that both the Fujitsu and Huawei server distributions are mainly deployed in Asia/Pacific and Europe. To date, neither has seen significant deployment in the North American region.
- Cisco UCS and HPE Integrity Superdome servers are getting more competitive with each subsequent ITIC Global Server Hardware and Server OS Reliability survey. In this latest 2017 Mid-Year Update poll, only two percent of Fujitsu PRIMEQUEST and three percent of Cisco, HPE Integrity and Huawei Kun Lun servers recorded in excess of four (4) hours of unplanned downtime.
- **Lenovo x86 X6 mission critical servers** scored the highest reliability ratings among all x86 platforms.
- **Commodity Aging Hardware:** Some 69% of survey respondents said aged hardware (3 ½+ years old) had a negative impact on server uptime and reliability vs. 31% that said it has not impacted reliability/uptime.
- **Reliability continues to decline** for the sixth year in a row on the HPE ProLiant and Oracle’s SPARC & x86 hardware and Solaris OS – each of which experienced 37 minutes and 33 minutes of unplanned outages per server/per annum. Reliability on the Oracle platforms declined slightly mainly due to aging. Many Oracle hardware customers continue to eschew or prolong upgrades and instead migrate to rival platforms.

Reliability and Uptime by the Numbers

A generation ago, in the 1990s, two nines or 99% uptime, equaling nearly 88 hours of per server downtime, was considered acceptable. In the current 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.

ITIC’s 2017 – 2018 Reliability poll indicates that 80% 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 eight percent in the last two years.

And it’s easy to see why. Technology advances in virtualization, cloud computing and IoT ecosystems which are built on myriad interconnected devices all demand near-flawless, uninterrupted availability.

At the same time as connected devices, applications and people have increased by orders of magnitude; *the potential* for collateral damage has increased commensurately. 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. 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 again publish the Table depicting the actual 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 enable organizations and their IT departments to calculate downtime in order to measure the business and monetary impact on the firm. Metrics of three, four and five nines of uptime – 99.9%, 99.99% and 99.999%, – equate to 8.76 hours; 4.38 hours, 52.56 and 5.26 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 2017

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

© Copyright 2017 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.

- **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.

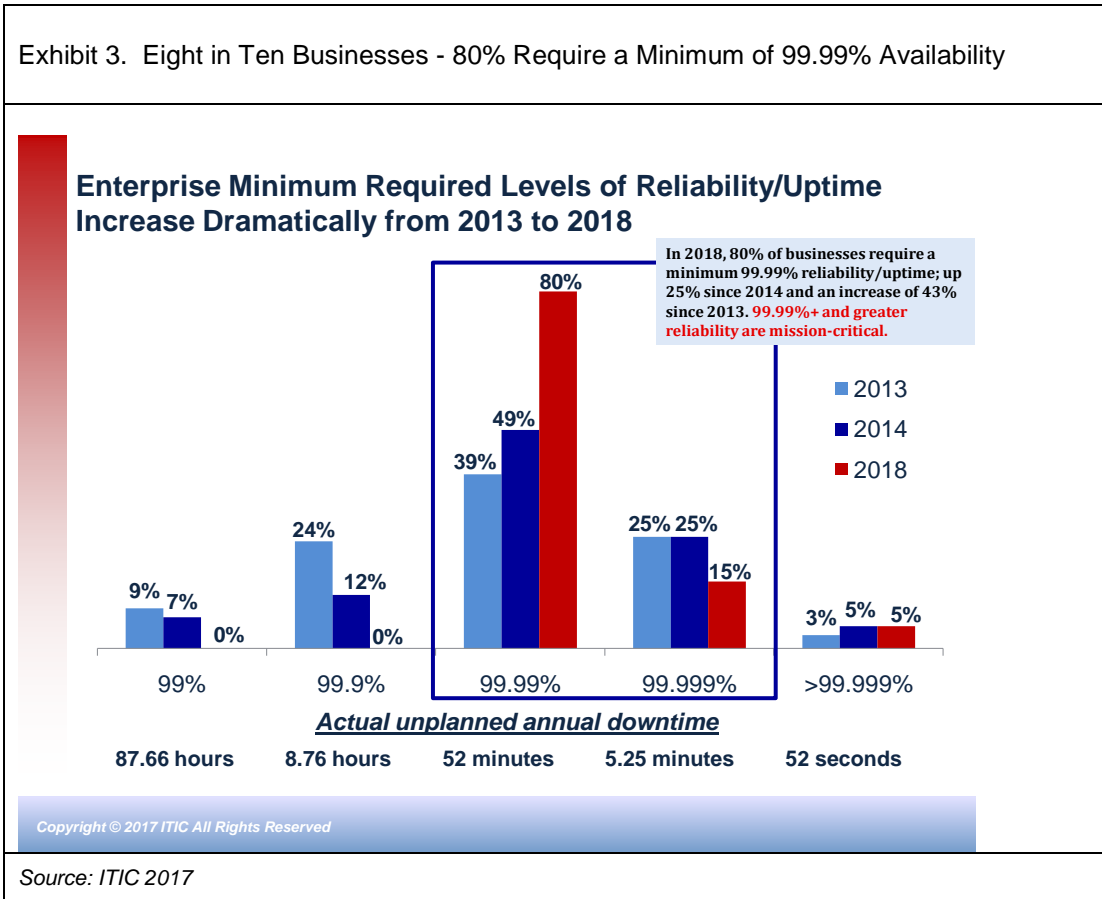
Minimum Reliability Requirements Rise

Organizations are acutely aware of the consequences of downtime.

So it's no surprise that the latest ITIC 2017 – 2018 Global Server Hardware and Server OS Reliability survey showed that eight-in-10 businesses - 80% - now require a minimum of "four nines" or 99.99% reliability and uptime. This is double the 39% of respondents that said their firms required 99.99% availability in ITIC's 2013 poll.

Four nines availability is the new acceptable reliability standard. It equates to 52 minutes of *unplanned* per server/per annum downtime or 4.33 minutes per month. Additionally, 15% 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 (**See Exhibit 3**). The remaining five percent of enterprises indicated their firms need "six nines" – 99.9999% or better availability/uptime.

Exhibit 3. Minimum Reliability Requirements Rise



And that’s not all. Today’s workloads 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. The most recent ITIC 2017 – 2018 survey revealed server workloads have increased an average 34% since 2014.

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.

Any downtime is an anathema to business operations. Unplanned service outages of even 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 mid-sized businesses (SMBs) and mid-sized enterprises (SMEs) are also vulnerable and just as risk-averse as an 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.

2017 – 2018 Reliability Trends

The 2017 – 2018 survey results indicate that Human Error continues to be the biggest cause of unplanned reliability incidents,

However, from a pure technology standpoint the top three technology issues adversely impacting reliability are in order: Security, Backup and Business Continuity and Integration/Interoperability.

Reliability is not merely a percentage or a statistic. In order to fully comprehend the crucial nature of infrastructure reliability, one must consider it in the context of “cause and effect.” That is accomplished by measuring Reliability in terms of its impact on daily business operations; lost productivity; hourly downtime costs; the manpower time and costs associated with remediation efforts and calculating the cost of lost, damaged, stolen, destroyed or altered data. Additionally, corporate enterprises must also take into account the negative impact that downtime may have had on business partners, customers and suppliers. In a worst case scenario, the corporation may find itself the target of expensive litigation resulting from failure to meet existing Service Level Agreements (SLAs) or failure to meet Regulatory Compliance laws. This can put the organization at risk of civil and criminal penalties and fines and even jail time.

To reiterate, in cloud computing, virtualization, IoT, mobility and BYOD where a much greater percentage of servers, networks, people, devices and applications are interconnected, the risks are commensurately greater. “Four nines” or 99.99% uptime is now the required minimum level of reliability for 80% of survey respondents. This number will continue to rise.

Similarly the Hourly Cost of Downtime also continues to soar.

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 68% of respondents versus 32% who said security does not affect uptime.
- Human Error: Over 50% of survey respondents said human error is the top reliability concern. 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 Rock Solid Reliability

Corporate enterprises have given IBM hardware the highest reliability ratings every year for the past decade, 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 employed 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 servers deliver consistent performance, reliability, security, advanced features and 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.

IBM hardware has retained its status as best in class in terms of reliability, stability and performance and customer satisfaction throughout the nine years that ITIC has conducted its Global Server Hardware and Global Server OS Reliability survey. The excellent performance, reliability, stability and security of IBM servers are attributable to several key factors.

- **Patent Power:** For the last 24 years, IBM has been awarded more patents than any other vendor worldwide. IBM continues its dominance in patents to this day. According to IBM more than 2,700 of the patents it was awarded in 2017 involved Artificial Intelligence (AI), cognitive computing, cloud computing, cyber security, and data analytics. All of these technologies are crucial for next generation strategic initiatives. So IBM is extremely well-positioned to retain its market leadership position and boosting the bottom line revenue via a robust and diversified portfolio of server hardware infrastructure and key applications.
- **Research and Development (R&D):** IBM has maintained an unswerving commitment to R&D. and continual refresh of its embedded performance, reliability, security and management functions, IBM has demonstrated an unswerving ability to craft and articulate 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 high technology companies that spend the most of research and development (R&D), according to Standard & P Capital IQ. In 2016, IBM spent \$5.4 billion on R&D, which equates to approximately six percent of its annual revenue. Additionally, IBM spent more than half of its R&D monies in the United States.
- **Product Innovation:** The POWER8 and newly released POWER9 servers are optimized for Data Analytics, AI and security.
- **Customer Satisfaction:** 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.
- **Superior Service and Support:** IBM has achieved consistently high marks for the breadth and depth of its professional services and after-market technical service and support.

All of these aforementioned issues contributed to IBM hardware maintaining its status as best in class in terms of reliability, stability and performance and customer satisfaction since 2008 when ITIC first began conducting its Global Server Hardware and Server OS Reliability Surveys. IBM consistently has demonstrated its ability to articulate and craft a comprehensive and cogent product roadmap and strategy and execute against it. 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 Z mainframe is also boosting Big Blue's reliability standings. In June, 2017 IBM released the IBM Z, the 14th generation of IBM's industry-leading mainframe technology, advances the

already solid and robust security and reliability features inherent in the platform over the last decade. It also amps up the processing power to new levels. The IBM z14 has the ability to process 12 Billion encrypted transactions daily. It accomplishes this via the industry's fastest microprocessor and a new scalable system structure that delivers a 35 percent capacity increase for traditional workloads and a 50 percent capacity increase for Linux workloads compared to the previous generation IBM z13. The system can support:

- More than 12 billion encrypted transactions per day on a single system.
- The world's largest MongoDB instance with 2.5x faster Node JS performance than x86-based platforms.
- Two million Docker Containers.
- 1,000 concurrent NoSQL databases.

The aforementioned IBM POWER8-based processor systems and the latest POWER9 servers 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 z/OS 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 addition, response times 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 and POWER9 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. This is crucial, since companies whose server hardware does not expand capacity, will almost certainly see reliability decline without expensive retrofits and upgrades.

Lenovo System x Servers Best in Class Reliability

It has been three and a half years since IBM completed the sale of its x86 server business to Lenovo in a 2014 deal that was valued at \$2.3 billion.

And for the fourth straight year, Lenovo System x servers have maintained the same high levels of reliability they had previously achieved while under the IBM brand. Lenovo's Data Center Group's strategy is to leverage ongoing partnerships 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.

Survey respondents gave Lenovo high grades marks for availability based on the vendor's continuing product innovation and close attention to detail on ease of maintenance, technical support, documentation and responsiveness.

In particular, Lenovo's sixth generation System x3850 X6 mission critical server, which is based on Intel's Xeon E7-4800 v4 and E7 8800 v4 processors, delivered five nines – 99.999% reliability for close to three-quarters of the respondents to ITIC's latest poll. It incorporates a slew of fault-tolerant and high-availability features into a high-density, 4U rack-optimized lid-less package that reduces the space needed to support massive network computing operations and simplify servicing. The x3850 X6 supports up to four Intel Xeon E7 high performance processors and up to 12 TB of memory. The Lenovo System x3850 X6 server is optimized for enterprises who are running large mission critical databases; Big Data Analytics in virtualized, cloud computing and IoT environments. It supports up to four E7 v4 processors with a total of 96 cores and 192 threads to maximize the concurrent running of multi-threaded applications. This in turn, improves productivity by boosting system performance via processors

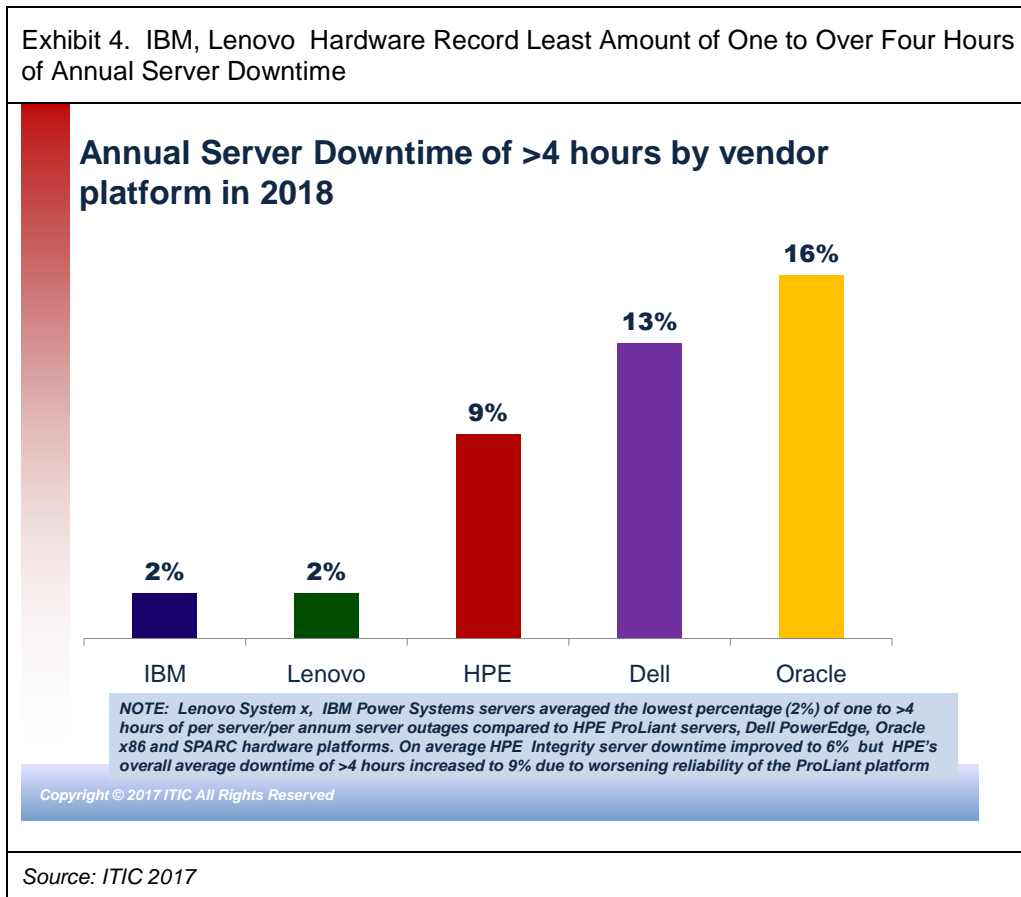
up to 24 cores, core frequencies up to 3.2 GHz core speeds, L3 cache of up to 60 MB, and three QPI interconnect links at up to 9.6 GTps.

Many Lenovo customer respondents interviewed by ITIC also extolled Lenovo's "agile system design" and hot swap capabilities. These features enable enterprise IT administrators to easily upgrade the system to provide scalability on-demand via front and rear access. IT managers can add components without removing the entire server from the rack. This reduces management time. Lenovo's server design is also resilient. It maximizes application uptime and provides easy integration in virtual environments. Integration and interoperability are among the key elements that can positively augment or negatively undermine overall system reliability. According to ITIC's latest 2017 -2018 survey results

Lenovo also markets its line of entry level ThinkServers, which likewise got high reliability ratings among the SMB respondents in ITIC's 2017 – 2018 Global Server Hardware and Server OS Reliability survey with over 86% of respondents reporting 99.99% uptime or better.

This is a win-win Lenovo customers and Lenovo itself. As **Exhibit 4** shows, Lenovo System x servers retain the same high levels of reliability they delivered while still under the IBM brand. Both the IBM and Lenovo server platforms recorded the least amount of unplanned annual server downtime – just two percent of respondents reported unplanned server downtime of over four (4) hours. These are the best results of any of the mainstream server hardware distributions.

Exhibit 4. Annual Downtime of >4 Hours by Server Vendor

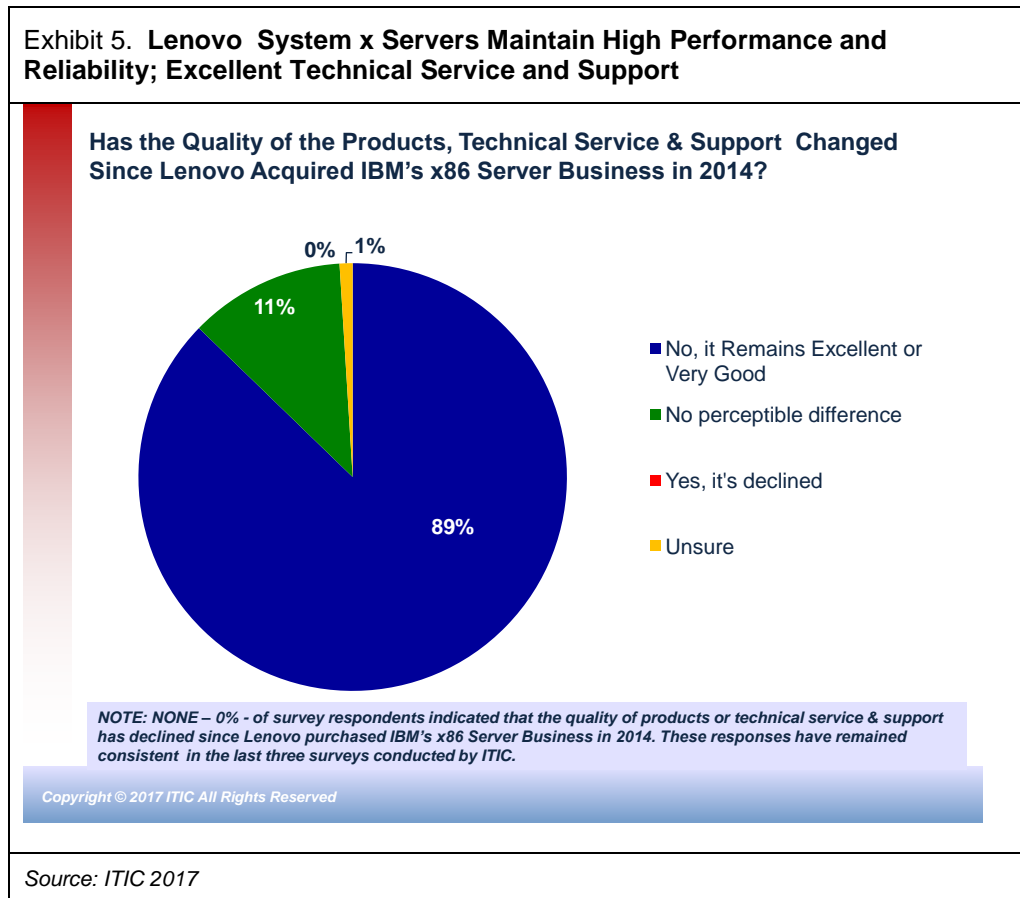


Source: ITIC 2017

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 2017 - 2018 Reliability Survey again polled organizations on how the Lenovo System x servers were faring in terms of product quality and the responsiveness and quality of after-market technical service and support in the three-and-a-half years since Lenovo acquired IBM's x86 server business unit in 2014. As **Exhibit 5** illustrates, an overwhelming 99% of respondents rated service and support, product quality as excellent and/or indicated their firms experienced no perceptible differences since the changeover from the IBM to the Lenovo brand.

Exhibit 5. Corporate Enterprises Say Lenovo System x Servers Maintain High Quality



Source: ITIC 2017

Huawei Kun Lun, Fujitsu PRIMEQUEST E Series Mission Critical Platforms Notch Strong Reliability Ratings

Mission critical servers from industry veteran hardware vendor Fujitsu and newcomer Huawei also posted high reliability statistics in the latest ITIC Global Reliability survey.

Fujitsu's PRIMEQUEST E Series and Huawei's Kun Lun servers both utilize the latest Intel Processor E7 v4 chips. The Fujitsu PRIMEQUEST supports hardware partitioning and targets enterprises that require "five nines" of reliability.

Huawei is among the newest entrants to challenge established players in the high end server market with its two year old Kun Lun family of 16- and 32-socket mission critical servers.

Huawei is now among the Top Five vendors in terms of global server hardware shipments. It is moving aggressively to target UNIX and RISC-based solutions from rivals like HPE, IBM and

© Copyright 2017 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.

others by promising turbo-charged performance with a significantly lower Total Cost of Ownership (TCO) than competitors.

Huawei's Kun Lun servers made their first appearance in ITIC's 2017 – 2018 Global Server Hardware, Server OS Reliability Survey and scored very well. Huawei's Kun Lun servers running Linux recorded just 3 minutes of unplanned per server/per annum downtime, which put it in the top three most reliable platforms, tied with HPE Integrity.

Huawei has fine tuned the Kun Lun server line to deliver the highest levels of four, five and six nines of reliability, The mission critical servers also deliver the scalability and fault tolerant capabilities necessary to support compute-intensive applications like high end databases, In-memory computing, cloud computing and virtualization, HPC fat node and IoT.

The Kun Lun servers leverage Huawei's Reliability, Availability and Serviceability (RAS 2.0) technologies, including core component hot swap, fault prediction, and partitioning. These advantages make Kun Lun the server of choice for building open, mission-critical platforms. Huawei has established an open platform alliance with mainstream operating system and database vendors, among them Microsoft, Oracle, IBM, Red Hat, and SUSE.

The Kun Lun server line (named after the Kun Lun mountain range is China's version of Mt. Olympus, the dwelling of the gods), is the industry's first 16- and 32-socket x86 mission-critical server series. They are based on the Intel Corp. open architecture x86 processors Intel Xeon E7 v3 and v4 chips. The Kun Lun servers support up to 32 CPUs, 768 cores and up to 32 TB DDR4 memory capacity. The latest Huawei mission critical servers incorporate the latest RAS 2.0 features and functionally to support "five nines" – 99.999% reliability/availability. Their performance, reliability, scalability and flexibility are ideally suited to a wide range of vertical markets such as Banking and Finance; Carriers, Government, Healthcare and Telecommunications.

Huawei's Kun Lun servers incorporate:

- RAS 2.0 intelligent fault management and the possibility of online key components maintenance, which make it easier and more efficient for network administrators to perform maintenance and upgrade tasks.
- A Proactive Failure Analysis Engine (PFAE), which delivers preventative maintenance capabilities to assist organizations in identifying and responding to performance issues in advance of any issues and resulting costly system downtime.
- Logical and physical partitioning to allow network administrators to set partitions containing from a single physical core or up to 32 CPUs.

Kun Lun 9032 is the flagship model with up to 32 CPUs (24 CPUs in non-full configuration), allowing for up to 576 cores when paired with 18-core members of the Intel's Xeon E7-8800 v3 and v4 processors. The Node Controller interconnection chips are all developed and

manufactured by Huawei to allow high speed interconnection for 32 and even 64, Intel Xeon E7 v3 and v4 CPUs. Huawei develops and manufactures its own Node Controller interconnection chips. This is a boon for customers and a competitive advantage for Huawei, because it is not dependent outside suppliers. Huawei can more quickly respond with updates, replacement parts and fixes in the event corporate customers experience issues and avoid time consuming delays that often occur in trying to get replacement parts from third party component manufacturers. Fast response time is a crucial factor in minimizing the impact of downtime on productivity and revenue.

Huawei Kun Lun Supports SUSE Linux and SAP HANA

Since the initial 2016 launch of the Kun Lun 9016 and 9032 servers, Huawei inked several key partnerships with vendors including SUSE, Oracle and SAP. Over the last 16 months Huawei also inked key alliances with SAP HANA for cloud computing and database veteran Oracle's DB 12C Rev 1.

Huawei also deepened its ties with SAP by certifying the Kun Lun servers to meet the requirement for 6/8/12/16 TB of In-Memory computing for OLTP scenarios. By early 2018, Huawei said the Kun Lun server family will be certified for 20/24 TB. Currently, only two vendors, including Huawei, in the world are capable of supporting 24 TB appliances.

Overall, Huawei's Kun Lun family of servers is well positioned to mount an aggressive and significant challenge to established mission critical high end server vendors.

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 near 90% of the microprocessor market continually updates its processors.

The Intel Xeon Processor E7 v3 and v4 chips are specifically geared towards reliability. All of the 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 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 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.

The caveat: the actual reliability results that organizations realize in production networks will vary according to individual firms' configuration, provisioning, applications and workloads and integration and interoperability within their specific environments.

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 6 x 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 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 v3 and v4 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 over the last four years and subsequent first person conversations with IT managers indicates that the improvements in the Xeon E7 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 overloads 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.

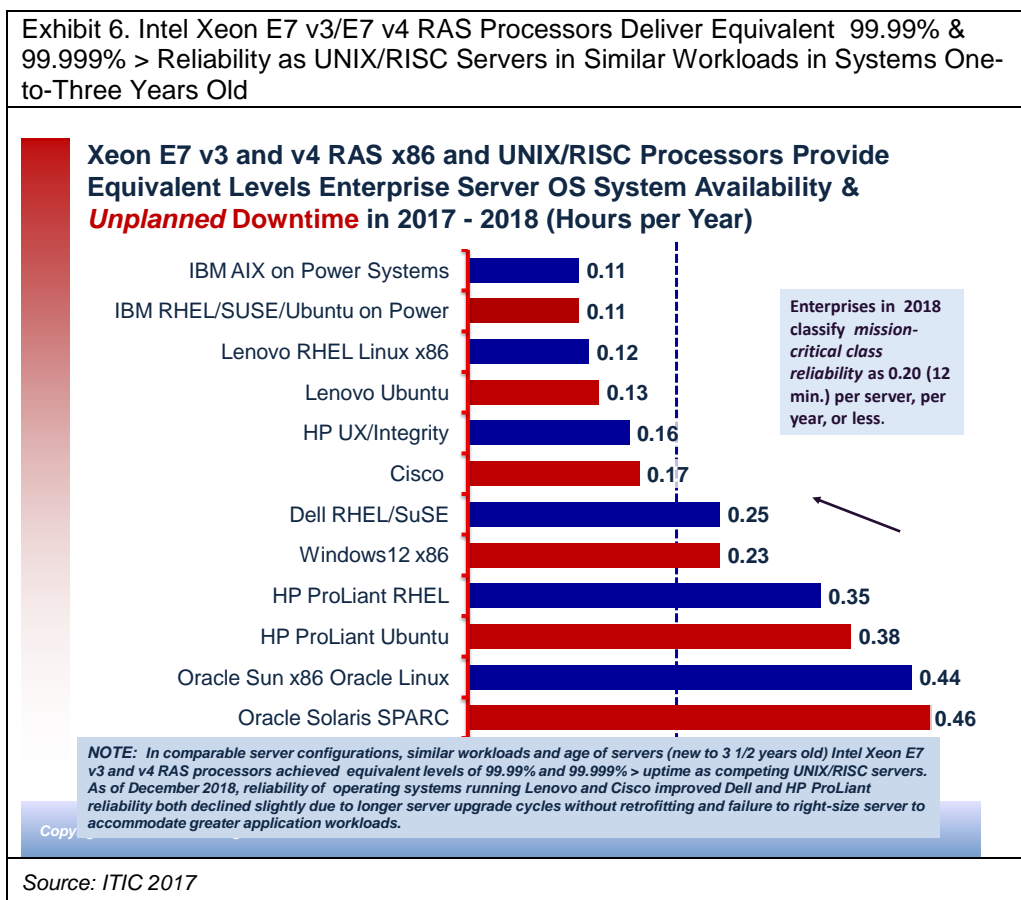
ITIC’s first person customer interviews yielded invaluable anecdotal information regarding the net positive performance and reliability gains specifically related to the newer processor technology.

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

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.

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 v4 family incorporates the *Intel Run Sure Technology* that integrates processor, firmware, and software layers to help diagnose fatal errors, contain faults, and automatic recovery. As **Exhibit 6** indicates, the Intel Xeon E7 v3 and E7 v4 RAS processors when properly deployed in newer systems can achieve equivalent levels of high reliability as rival UNIX and RISC-based servers in comparable configurations and workloads.

Exhibit 6. Comparison of Intel Xeon E7 v3 and v4 RAS Processors vs. UNIX/RISC



Source: ITIC 2017

According to customer interviews and essay comments, the spike in the percentage of Dell PowerEdge and HP ProLiant system unavailability is attributable mainly to the high proportion - about 62% of aged Dell PowerEdge and 58% HP ProLiant servers that are over three and a half years old and have not been upgraded or replaced.

Cisco UCS Remains Highly Reliable

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. Like every server vendor, Cisco has been impacted by the upheaval and contractions in the high technology market sector. Since 2013 and continuing into 2018, the company has restructured and cut well over 10,000 workers in several rounds of layoffs, representing about 16% 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.

Cost of Hourly Downtime Soars

The only good downtime is no downtime.

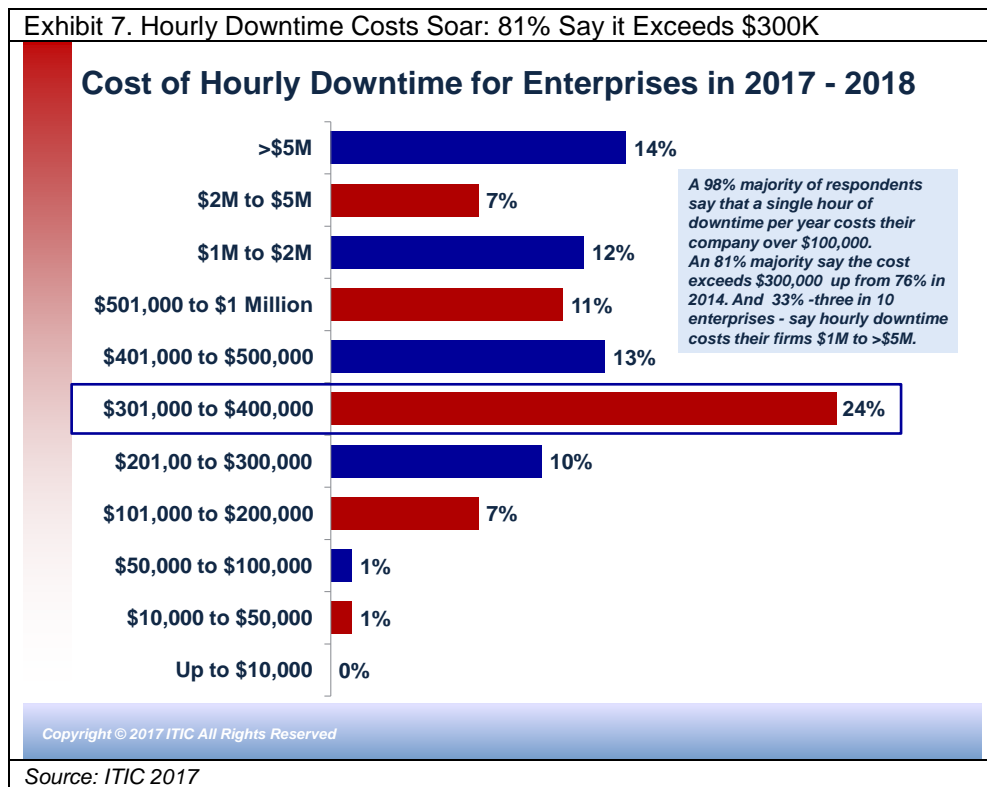
As companies grow ever more reliant on their interconnected networks and applications to conduct business, the cost of downtime is rising commensurately. To put it bluntly: if the servers, applications and networks are unavailable *for any reason* business and productivity slow down or cease completely until server and application access is restored.

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

As **Exhibit 7** illustrates, 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.

Consider this: for a company one hour of downtime estimated at \$300,000, a one minute outage costs \$5,000, while five minutes of downtime can cost \$25,000. These are average conservative estimates. They do not include worse case scenarios of outages that occur in peak usage times. Similarly, these figures do not cover the cost to recover lost, damaged or stolen data and they do not take into account legal fees, penalties or financial settlements that result from litigation.

Exhibit 7. Hourly Downtime Costs Top \$300K for 80% of Enterprises



Source: ITIC 2017

These statistics reinforce that infrastructure matters.

Server hardware, server OS and application reliability or instability have direct and far reaching impact on the corporate bottom line, ongoing business operations. They can also irreparably damage a company's reputation. In some extreme cases, business and monetary losses as a result of unreliable servers can cause the company to go out of business due to sustained losses and possible litigation in the wake of an outage.

Conclusions and Recommendations

In summary the ITIC 2017 - 2018 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.

Server hardware, operating systems, mission critical applications and networks must be available 24 x 7. Reliability and availability are crucial to uninterrupted business operations and productivity. IBM and Lenovo once again continued to deliver the highest levels of reliability. IBM's z Systems Enterprise mainframe is in a class by itself and is the clear winner in terms of unparalleled fault tolerant levels of 99.999% or greater reliability.

Among the "work horse" mainstream systems, the IBM Power Systems, Lenovo System x are the most reliable and offer the highest levels of technical service, support and customer satisfaction.

Cisco UCS reliability is likewise very strong and consistent. The most recent ITIC survey also showed impressive gains by veteran vendor Fujitsu's PRIMEQUEST E Series and newcomer Huawei's Kun Lun mission critical server. To date, though these firms' server offerings do not have a substantial presence in the North American region, though they are strong in Asia Pacific. Their strong showing in reliability makes them worth watching.

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 80% majority of corporations now require "four nines" or 99.99% minimum uptime, while one-in-five organizations – 20% require "five or six nines" – 99.999% and higher. 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 Put Off Updates.** 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.** The latest ITIC survey statistics show that 12% of respondents fail to calculate their server and OS reliability and the cost of downtime. This is a big mistake. 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.
- **Regularly conduct security awareness training and asset management testing.** Security training is essential and should be part of every organization's routine irrespective of company size or vertical market. 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 *2017 - 2018 Global Server Hardware and Server OS Reliability Survey*, polled C-level executives and IT managers at over 750 corporations worldwide from August through December 2017. 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 34% of the respondents. Small and medium enterprises (SMEs) with 101 to 1,000 workers represented 23% of the participants and the remaining 43% of respondents came from large enterprises with 1,001 to over 100,000 employees. Survey respondents hailed from 49 different vertical markets. Approximately 74% of respondents hailed from North America; 26% 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/2017/07/ibm-z14-mainframe-advances-security-reliability-processing-power/>

<http://itic-corp.com/blog/2017/06/ibm-lenovo-servers-deliver-top-reliability-cisco-ucs-hpe-integrity-gain/>

<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/>