

INFORMATION TECHNOLOGY
INTELLIGENCE CONSULTING

Information Technology Intelligence Consulting



**ITIC 2019 Global Server Hardware,
Server OS Reliability Report**

March 2019

Table of Contents

Executive Summary.....3

Introduction.....7

Reliability and Uptime by the Numbers9

 Reliability versus Availability 10

Reliability Dollars and Sense: The Actual Cost of Downtime 15

 Data & Analysis 16

 IBM and Lenovo Reliability Success: Innovation, High Performance, Security and
 Top Technical Support Deliver High Reliability 17

 HPE’s Reliability Ratchets Up 18

Other Notable Survey Findings.....20

Server Hardware Platform Overview22

 IBM Power Systems and IBM Z22

 Lenovo System x and ThinkSystem23

 ThinkSystem SE35023

 HPE Integrity24

 Huawei KunLun and FusionServer24

2019 Reliability Trends25

 Human Error and Security Threats Top the List25

 Hourly Cost of Downtime Continues to Rise.....26

 Minimum Reliability Requirements Rise27

 Higher Server Workloads Negatively Impact Reliability28

 Hardware Continues to Evolve, and So Does Risk29

 Commodity Hardware Can Have Expensive Consequences29

 Robust, Reliable Servers Pay Off.....30

Conclusions31

Recommendations.....34

Appendix37

 Survey Methodology37

 Survey Demographics37

 Links to Supplemental Information37

Executive Summary

IBM delivered the highest reliability for the 11th year in a row, while Lenovo had the best uptime for the sixth straight year.

HPE Integrity, Huawei KunLun, and newer Dell Servers post reliability gains.

IBM System Z and Power Systems, Lenovo x86, and Huawei KunLun hardware deliver highest availability and shortest duration outages.

IBM Power Systems, Lenovo x86 System x and ThinkSystem, HPE Integrity, and Huawei KunLun deliver up to 24 times better reliability than the least-efficient white-box platforms.

Human error, security, and IT departments' failure to monitor BYOD device security are top external threats and causes of downtime.

High reliability, uptime, and availability are imperative for today's "always on" digital networks.

For the 11th straight year, IBM Z mainframe and Power Systems achieved the highest server reliability rankings in ITIC's 2019 Global Server Hardware and Server OS Reliability survey. Lenovo System x and ThinkSystem delivered the best uptime among Intel x86 servers for the sixth year in a row. Hewlett Packard Enterprise (HPE) Integrity and Huawei KunLun mission-critical servers continued with impressive gains, and moved to the head of the reliability class.

ITIC's independent web-based survey polled over 1,000 businesses worldwide from October 2018 through January 2019. It compared the reliability and availability of 18 mainstream server platforms and one dozen operating system (OS) distributions. To obtain the most accurate and unbiased results, ITIC accepted no vendor sponsorship.

IBM System Z server is in a class of its own and maintained its best-in-class rating among all server platforms. An 83% majority of IBM respondent organizations said their firms achieved five and six nines — 99.999% and 99.9999% — or greater uptime. Nine in 10 IBM System Z customers reported that the mainframe recorded a mere 0.74 seconds of unplanned per server downtime each month, and 8.87 seconds of downtime annually due to inherent flaws in server hardware or its component parts. Less than 0.5% of IBM System Z respondents said the mainframe experienced unplanned outages exceeding four hours of annual downtime.

Among mainstream servers, IBM POWER8 and POWER9, along with Lenovo x86, HPE Integrity Superdome X, and Huawei mission-critical KunLun servers delivered the highest levels of reliability/uptime among 18 platforms. Each server platform averaged just under two minutes of *unplanned* per annum/per server downtime due to inherent flaws or quality issues involving server hardware, or components such as hard disk drives or memory (**see Exhibit 1**).

Depending on individual implementation, configuration and usage scenarios, the IBM Power Systems, Lenovo x86, HPE Integrity, and Huawei servers deliver up to 24 times better reliability than the least efficient rivals. For the second year, unbranded white box servers recorded highest per server/per annum downtime — reaching 49 minutes of downtime in the 2019 poll — due to system flaws or component issues.

The economic annual downtime cost comparisons among the top performing and the least reliable server hardware platforms is staggering.

A single hour of downtime with an estimated cost of \$300,000 equates to \$4,998 per server/per minute.

Organizations using the most reliable servers — IBM POWER8 and POWER9, Lenovo x86 System x and ThinkSystem, HPE Integrity, or Huawei KunLun — that experienced just under or just over two minutes downtime would spend approximately \$9,996 in annual per server downtime costs due to inherent flaws in server hardware or component parts (**see Table 2**).

By contrast, corporations using Dell PowerEdge servers, which experienced an overall average of 25 minutes of per server/per annum downtime, racked up yearly outage costs of \$124,950 at the same \$300,000 per hourly downtime rate for a single server. This exceeded yearly remediation costs by more than 14 times for organizations deploying IBM Power Systems, Lenovo System x and ThinkSystem, HPE Integrity, Huawei KunLun, and FusionServer.

Outage costs are steepest for unbranded white box servers.

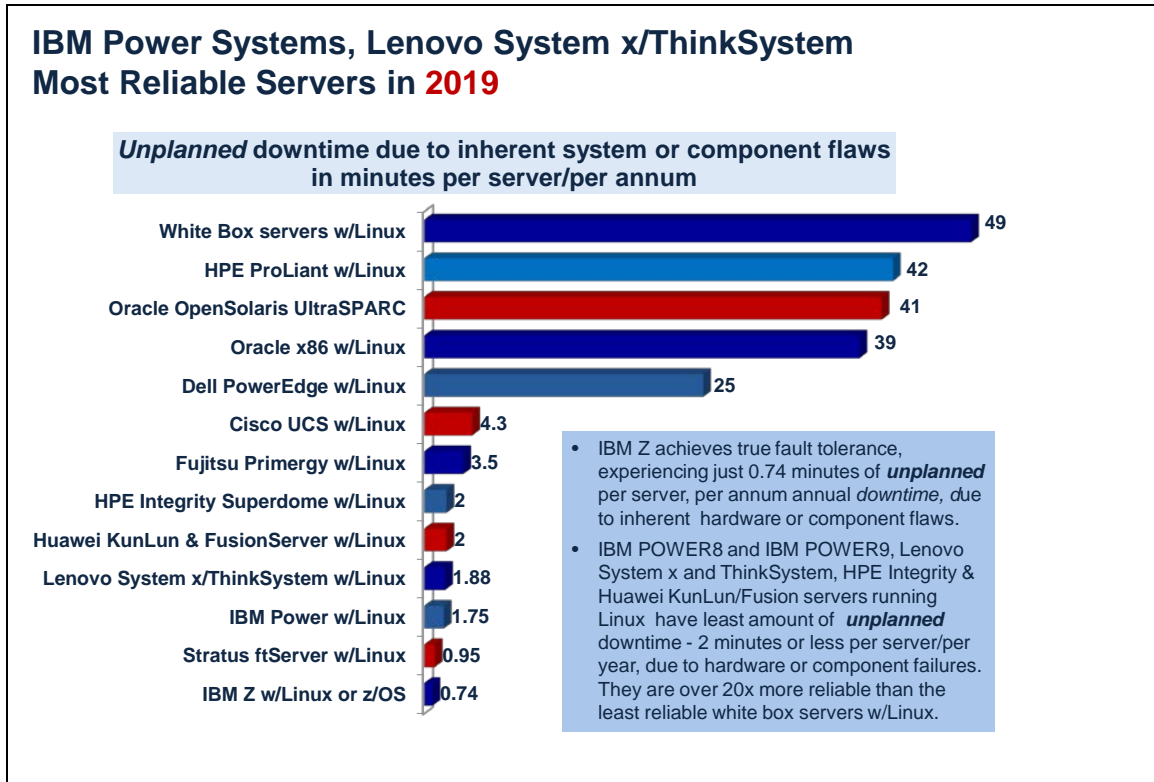
Corporations deploying the least reliable unbranded white box servers — that experience 49 minutes of per server/per annum downtime — potentially could incur downtime losses of \$244,902, specifically related to server hardware flaws and bugs in the OS and applications. This is not surprising, since white box servers frequently incorporate inexpensive components and some businesses recklessly run unsupported or pirated versions of OSes and applications. The aforementioned hourly downtime examples are for **just one server**. Downtime costs mount quickly; costs can reach into the millions for corporations with dozens or hundreds of highly unreliable servers.

Among the other top survey findings:

- **Reliability:** IBM Power Systems, Lenovo System x, and ThinkSystem hardware and Linux OS distributions were first or second in every reliability category, including server, virtualization, and security.
- **Availability:** IBM Z mainframe, Power Systems, Lenovo System x and ThinkSystem, HPE Integrity, and Huawei KunLun all provided the highest levels of server, application, and service availability. When servers did experience an outage due to an inherent system flaw, they were of the shortest duration, typically one to five minutes.
- **Technical support:** Businesses gave high marks to IBM, Lenovo, HPE, Huawei, and Dell tech support. Only 1% of IBM and Lenovo customers and 3% of HPE and Huawei users gave those vendors “poor” or “unsatisfactory” customer support ratings.
- **Hard drive failures most common technical server flaw:** Faulty hard drives are the chief culprits in inherent server reliability/quality issues (58%) followed by motherboard issues (43%), and processor problems (38%).
- **IBM, Lenovo, and Huawei KunLun Servers had fewest hard drive failures:** IBM, Lenovo, and Huawei KunLun platforms saw the fewest hard drive quality or failure issues among all server distributions within the first one, two, and three years of service. Fewer than 1% (0.4%) of IBM Z mainframes, for example, experienced technical problems with hard drives in the first year, followed by IBM Power Systems, Lenovo System x, and ThinkSystem with 1% each during the first 12 months.
- **External issues negatively affect reliability:** Human error (59%) and security (51%) are the top external causes of downtime and unanticipated reasons for taking servers offline.
- **Minimum reliability requirements increase:** An 85% majority of corporations require a minimum of four nines (99.99%) of uptime for mission-critical hardware, OSes and line of business (LOB) applications. This is up 4% from ITIC’s 2017-2018 Reliability Survey.
- **Patch time increases:** Seven in 10 businesses now devote from one hour to over four hours applying patches. This is primarily due to a spike in wide ranging security issues such as email phishing scams, ransom ware, CEO fraud, and malware and viruses.
- **Increased server workloads cause reliability declines:** Reliability declined in 64% of servers over 3½ years old, when corporations failed to retrofit or upgrade hardware to accommodate increased workloads and larger, more compute-intensive applications. This is up 19% from the 45% of businesses that said uptime declined due to higher workloads in the ITIC 2017-2018 Reliability Survey.
- **Hourly downtime costs rise:** A 98% majority of firms say hourly downtime costs exceed \$150,000 and **86%** estimate hourly downtime costs over \$300,000. Some 34% of companies say the cost of a single hour of downtime tops \$1 million.

Server hardware, OS, and by extension, virtualization reliability, uptime, and availability are core elements of the health of an organization’s Digital Age ecosystem — and the lifeblood of daily business operations.

Exhibit 1. Server Reliability by Hardware Platform Running Linux OS



Source: ITIC 2019

Introduction

Since 2008, the ITIC Global Server Hardware, Server OS Reliability Report has compared the reliability of 18 mainstream server platforms, one dozen server OS distributions (Linux, Unix, Ubuntu, Debian, Z/OS, and Microsoft Windows) and one dozen server hardware virtualization layers. It also delves into the internal issues that strengthen or undermine core server hardware and OS reliability.

The report quantifies and qualifies the overarching reliability of mainstream server hardware, based on key metrics and corporate policies including:

- Automated and manual patch management
- Percentage of Tier 1, Tier 2, and Tier 3 help desk calls
- Inherent server and server OS reliability
- Inherent server availability and the duration of outages
- System unavailability due to planned outages for routine system maintenance, upgrades, and application of patches
- Security issues, e.g., how quickly companies can identify and thwart hacks or malicious code, and how quickly vendors respond and deliver fixes and patches
- Human error such as an IT manager improperly configuring the server or the company's failure to right-size or upgrade servers to accommodate virtualization and more compute-intensive workloads
- Server virtualization reliability
- Vendor technical service, support, and documentation
- Budget/cost constraints
- Length of upgrade cycles
- Overworked, understaffed IT departments
- The impact of aging server hardware on reliability
- Integration and interoperability issues

ITIC expanded 2019 Global Server Hardware, Server OS Reliability survey questions to delve into specific technical issues that cause server performance problems and failure. The latest poll also incorporates new questions that address availability metrics or the length of unplanned outages of each server and OS distribution. For the first time, ITIC provides a detailed breakdown of monetary costs associated with hourly downtime ranging from \$100,000 to \$10 million in various server quantities ranging from a single server to 1,000 servers according to the amount of per server/per annum downtime by minutes.

New questions include:

- Specific causes of server and component reliability/quality issues, including processor and hard disk problems, memory failure, fans/cooling, audio, and video.
- How servers and components failed, e.g., power surge, exceeded capacity, overheating, and power surge.
- How frequently various server distributions experience component parts fail.
- The percentage of servers that experienced quality issues by specific vendor platform.

As always, ITIC's Global Server Hardware and Server OS 2019 Reliability Report utilizes information gathered from prior ITIC surveys to compare and contrast the reliability of the various platforms and to track industry trends. These objective findings, based on direct, independent user data, helps organizations make informed purchasing/upgrade decisions that align to their specific business and budget needs.

Reliability and Uptime by the Numbers

When it comes to reliability, there's only one direction that matters: up. Technical and business issues associated with infrastructure reliability have been steadily and inexorably increasing, based on wide ranging trends and use factors. These include:

- **Demand for reliability continues to rise:** Since 2013, demand for four nines (99.99%) uptime has increased by 46%. At that time, 39% of ITIC 2013 Reliability Survey respondents required 99.99% uptime. ITIC's 2019 Reliability poll finds 85% of respondents consider 99.99 % (52.56 minutes unplanned per server/per annum downtime) the minimum acceptable level of reliability for mission-critical servers and applications.
- **Applications are getting larger and more compute-intensive** and consume more system resources.
- **Server workloads increased by an average of 17% to 34%** during the last three years, depending on specific vertical markets and enterprise configurations and deployments.
- **Average hourly downtime costs continue to rise with no end in sight.** Currently, 86% of ITIC 2019 respondents said a single hour of downtime costs in excess of \$300,000. Over one-third (34%) estimate 60 minutes of downtime costs companies \$1 million or more. The price tag for a single minute of downtime vary widely depending on use configurations, timing, duration, and severity. The cost of one minute of downtime ranges from \$1,666 for an hourly outage of \$100,000, to \$16,667 per minute for an hourly downtime rate estimated at \$1 million (**see Table 2**). These figures do not include costs associated with potential legal fees, fines, or civil or criminal penalties that may arise.
- **Increase in connected devices and networks.** An estimated 39 billion devices will connect via the Internet by 2020. This is triple the number of Internet-connected devices in 2015. Connected devices and networks facilitate higher productivity, faster transaction speeds, and a greater number of transactions. When a server fails and is unavailable due to reliability issues, it can have a domino effect resulting in greater collateral damage.
- **Security and data breaches are on the increase.** Weak, porous server hardware and operating security, and lax corporate security, undermine reliability. Hamburg, Germany based research firm, [Statista](#), which tracks data breaches and records exposed, reports hacks increased exponentially in the last 19 years. In 2005, Statista ¹ reported 157 cyber security incidents in the U.S. that exposed 66.9 million records. By 2014, this increased 500%, with 783 data breaches that exposed 85.61 million records. Its most recent 2018 study found the U.S. experienced 1,244 data hacks, exposing over 446.5 million records.

¹ Statista, "Annual number of data breaches and exposed records in the United States from 2005 to 2018 (in millions)." URL: <https://www.statista.com/statistics/273550/data-breaches-recorded-in-the-united-states-by-number-of-breaches-and-records-exposed/>

All of these issues directly impact reliability positively or negatively. Based on industry trends, corporate enterprises will continue to require very high reliability beginning with the inherent reliability of their core server hardware and server OS infrastructure.

What does *reliability* actually mean?

A generation ago in the 1990s, two nines (99% uptime), or nearly 88 hours of per server downtime, was considered acceptable. Today, 99% or even three nines (99.9%), 8.76 hours of per server/per annum downtime, is unthinkable and unworkable.

Metrics of three, four, and five nines — 99.9%, 99.99%, and 99.999%, — equate to 8.76 hours, 52.56 minutes, and 5.26 minutes of per server/per annum downtime, respectively. Technology advances in virtualization, cloud computing, and IoT ecosystems built on disparate interconnected devices demand near-flawless, uninterrupted reliability and availability.

Reliability versus Availability

The terms *reliability* and *availability* are often used interchangeably, but are two distinct metrics.

Reliability tracks the overarching dependability of server hardware and OS and the amount or percentage of time it performs the intended function (e.g. LOB application, web server).


Reliability metrics also measure the mean time between failures (MTBF).

Availability measures the percentage of time server hardware and OS are operable and data and services are accessible. Availability chronicles duration of downtime: The amount of time the system, data, and services are unavailable per incident. A business with 99.99% server uptime may experience a Tier 2 outage lasting 15 minutes to one hour. This can cause more significant productivity disruption, data loss, and monetary loss than several shorter Tier 1 outages.

It is equally important organizations and IT departments monitor **both** reliability (number of incidents and amount of overall server downtime) and availability (length of each outage) to accurately determine the productivity and monetary cost to the company.

Exhibit 2 details the cost of one minute of hourly downtime of \$300,000 for a single server at \$4,998 per minute. Some 86% of firms say one hour of downtime costs their firms \$300,000+.

Exhibit 2. IBM Power Systems, Lenovo System x/ThinkSystem Achieve Lowest Downtime Costs



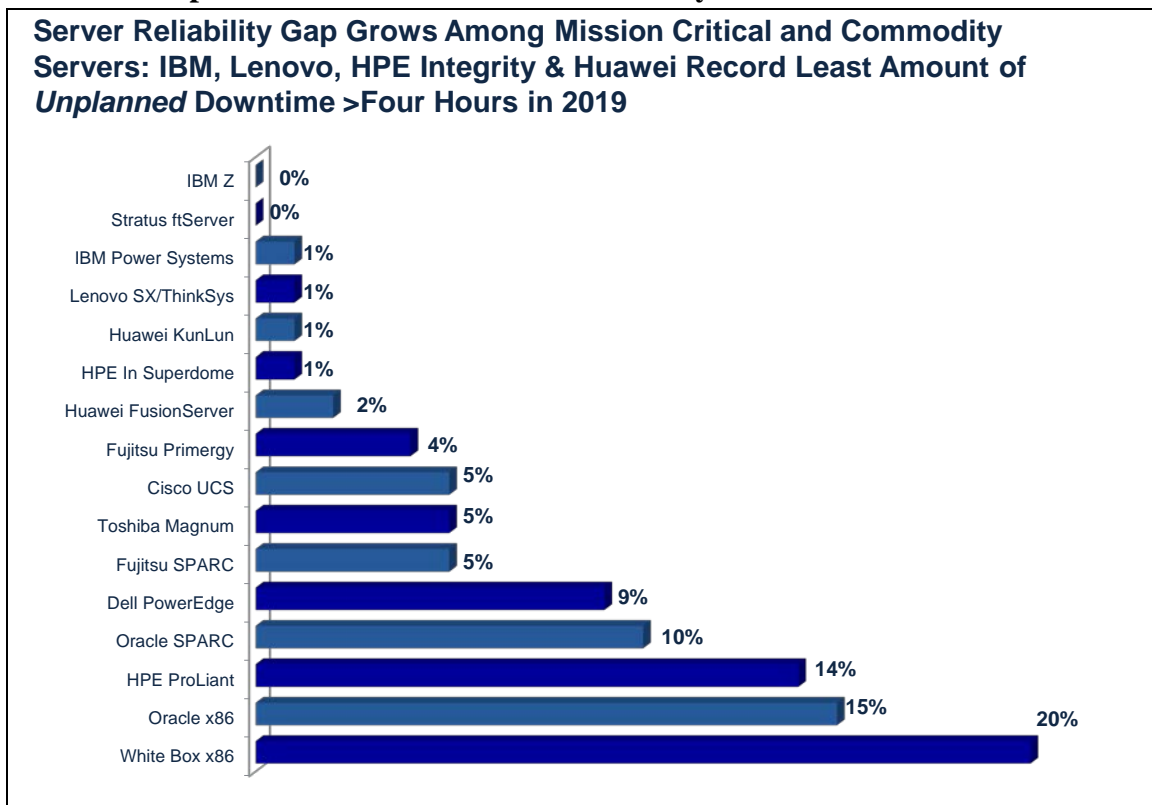
Source: ITIC 2019

Using the \$300,000 metric, ITIC calculated the estimated cost for one server based on the number of per server/per annum minutes of unplanned downtime experienced by each vendor platform due to inherent problems with the server hardware or components. As **Exhibit 2** illustrates, Lenovo x86 System x and ThinkSystem distributions were the most economical servers among all x86-based platforms. Specifically unplanned downtime related to hardware failures or problems cost IBM Power Systems customers who experienced 1.75 minutes of downtime an estimated \$8,746 per server/per annum. Lenovo x86 System x and ThinkSystem users, whose servers experienced 1.88 minutes of per server/per annum downtime, spent \$9,396. Corporations that deployed unbranded white box servers, which had 49 minutes of per server/per annum unplanned downtime due to server hardware flaws, racked up charges of \$244,902. This is a vivid example of how the least expensive servers can actually cost companies more than higher end distributions, which have more built-in features and robust capabilities.

Exhibit 2 also depicts how downtime disparities of even a few minutes can quickly add up and take a huge toll on an organization’s ongoing operational budget.

As **Exhibit 3** illustrates, ITIC’s 2019 Global Server Hardware, Server OS Reliability study found only a niche 1% of IBM Power Systems, Lenovo System x and ThinkSystem, HPE Integrity, and Huawei KunLun servers experienced over four hours of unplanned annual downtime due to problems with server hardware or component parts. These are the best scores among mainstream server platforms surveyed. This level of high reliability ensures excellent availability to servers, applications, data, and system resources. Corporate enterprises deploying IBM Z and Power Systems, Lenovo System x and ThinkSystem, HPE Integrity, and Huawei KunLun servers all achieved uninterrupted productivity along with high performance even during peak usage. This lowers total cost of ownership (TCO), accelerates return on investment (ROI), and mitigates risk.

Exhibit 3. Unplanned Downtime of >Four Hours by Server Platform



Source: ITIC 2019

Organizations whose server hardware, OS, and virtualization components fail to deliver a minimum of four nines and preferably, five nines of uptime, will have a very different usage experience. These corporations will immediately experience negative effects including:

- Productivity disruptions
- 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
- Heightened security risks
- Damage to the company's reputation
- Potential loss of existing customers and new business

Time is money.

Connected devices, applications, and people have increased by orders of magnitude. Consequently, *the potential* for collateral damage has increased commensurately. An outage on a virtual server running multiple instances of a crucial main 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 occurring during peak usage hours or interrupting a crucial business transaction can prove catastrophic and expensive. It can cost tens of thousands to millions per hour or millions per minute depending on specific circumstances and use cases. Firms must also factor in the cost of remediation efforts — time, labor, and expense involved to restore systems and services to full operational status and recovery of damaged or lost data.

Table1 below depicts the availability percentages and the equivalent number of annual, monthly, and weekly hours and minutes of per server/per annum downtime. It is a useful reference to calculate downtime and measure business and monetary impact on operations.

Table 2 illustrates the monetary costs associated with specific hourly downtimes ranging from \$100,000 to \$5 million based on per-minute costs from a single server to 1,000 servers. Taken together, the two tables paint a clear and dramatic picture of how quickly downtime costs can escalate to the detriment of a corporation's productivity and bottom line.

Table 1: Reliability/Uptime by the Numbers

Reliability %	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 2019

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

- Tier 1:** Typically minor, common, annoying occurrences. Network administrators can usually resolve such incidents in one to fewer than 30 minutes for dependent users. Tier 1 incidents can usually be fixed by rebooting the server (locally and remotely) and rarely involve any data loss. An example of Tier 1 outages is someone accidentally unplugging the server. A Tier 1 outage, even a minor one, can be costly *if* it occurs during peak usage hours or interrupts a crucial transaction.
- Tier 2:** Moderate-to-serious issues in which the server, OS, and key applications may be offline or unavailable 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 at least a portion of 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 prolonged, lasting four hours or longer. During this period, applications and services are unavailable to the corporation's associated dependent users: The IT department and external customers, partners, and suppliers. Tier 3 outages almost always require multiple network administrators to resolve 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. Examples include a backhoe severing a power line, a natural disaster such as a hurricane, flood or tornado, security breaches, and integration issues.

Reliability Dollars and Sense: The Actual Cost of Downtime

There is a marked economic disparity in the hourly cost of hourly downtime between the most reliable high-end server distributions and inexpensive platforms.

No server, OS, or application is immune to bugs or flaws. Even the most expensive, high-end hardware will occasionally experience a hard drive failure, memory leak, or OS bug.

However, IBM Power Systems, Lenovo x86 System x and ThinkSystem have delivered demonstrably superior reliability and availability and security over the past 11 years (IBM) and six years (Lenovo), than commodity-based x86 platforms that do not incorporate the same advanced functionality. Over the past several years, the same is true for higher end, mission-critical systems from HPE, Huawei, Fujitsu, and niche market vendor, Stratus.

Additionally, corporate enterprises that deploy higher-end systems are not “penny wise and pound foolish.” They purchase robust server configurations that are adequate or more than adequate to suit present and future needs. They adhere to regular refresh cycles and typically have more experienced IT and security administrators who can identify issues, isolate problems and perform remediation quickly. These factors all combine to limit and minimize downtime and the associated costs when technical issues do arise.

Table 2 below is a detailed illustration of the hourly cost of downtime ranging from \$100,000 to \$10 million per hour for a single server and in configurations of 10, 100, and 1,000 servers.

Table 2. Monetary Cost of Hourly Server Downtime: Per Minute/Per Server(s)

Hourly Cost of Downtime	Per Minute, Per Server	Per Minute, 10 Servers	Per Minute, 100 Servers	Per Minute, 1,000 Servers
\$100,000	\$1,667	\$16,670	\$166,667	\$1,666,670
\$300,000	\$4,998	\$49,980	\$499,800	\$4,999,800
\$400,000	\$6,666	\$66,660	\$666,600	\$6,666,670
\$500,000	\$8,333	\$83,330	\$833,300	\$8,333,300
\$1,000,000	\$16,667	\$166,670	\$1,666,700	\$16,667,000
\$2,000,000	\$33,333	\$333,330	\$3,333,300	\$33,333,000
\$3,000,000	\$49,998	\$499,980	\$4,999,800	\$49,998,000
\$5,000,000	\$83,333	\$833,330	\$8,333,300	\$83,333,000
\$10,000,000	\$166,667	\$1,666,670	\$16,666,700	\$166,667,000

Source: ITIC 2019

©Copyright 2019 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.

These days, a price tag of \$100,000 for one hour of downtime for a single server is considered “cheap,” but it still equates to \$1,667 per minute/per server. The cost of a more severe outage, estimated at \$1 million, equals \$16,667 per server/per minute.

Downtime costs increase exponentially based on the number of affected servers and the corporation’s estimate on the hourly cost of downtime. ITIC’s 2019 Global Server Hardware and Server OS Reliability Survey found 86% of respondents calculate that one hour of downtime costs the firm \$301,000 or more. Of that number, 34% indicate hourly downtime costs exceed \$1 million. Overall, only 2% of firms said one hour of downtime costs them \$100,000 or less, and 12% of respondents placed the valuation of a single hour of downtime at \$101,000 to \$300,000.

Except in rare instances, when an outage does occur it will involve multiple servers. Businesses must also consider whether or not the affected hardware was a virtual server running two, three, or four instances in a single machine.

Table 2 depicts the best and worst case financial impact of one hour of downtime, from the least expensive incident involving a single server due to inherent problems with the underlying hardware or component parts (e.g., hard drive failures). In this instance, it costs \$1,667 per minute. The much more extreme hourly downtime scenario presented in Table 2 affects 1,000 servers at an organization that values an hour of downtime at \$10 million. In the latter case, a very large enterprise could conceivably sustain crippling losses of \$166,667,000 per minute!

Keep in mind the above referenced ITIC Hourly Downtime monetary figures represent only the costs associated with remediating technical problems that caused the server or OS to fail. They do **not** include legal costs, criminal or civil penalties the company may incur or any “good will gestures” the corporation may make to customers (discounted or free equipment or services).

Data & Analysis

Reliability and uptime are not static metrics. Technology advances and business drivers evolve. No server hardware platform, OS, or software application is immune to failure. Even the most robust and reliable hardware and software distribution can fail. Any vendor or corporate enterprise may experience declines or increases in downtime related to bugs or flaws in its respective server platforms, OSes, and applications, year over year.

Corporate reliability and availability requirements are by necessity, markedly different in 2019 than they were 10, 5, or even 2 years ago. In the Digital Age, business is conducted 24 x 7, 365 days a year. Corporations frequently allow customers, suppliers, and business partners to access data on their systems and networks. Workforces are increasingly mobile. Fulltime and contract employees often work remotely, travel, and use their own devices to access the corporate network. Business is conducted irrespective of geographic location and time zones.

Today’s organizations are highly risk averse and intolerant of downtime.

Server failures create a domino effect. They adversely impact the OSes, applications, processes, user productivity, and business transactions in quick succession. Daily business operations and end-user productivity depend on the reliability and availability of server hardware. Network ecosystems now span on-premises datacenters, private, public, and hybrid clouds to geographically dispersed deployments that may be located hundreds or thousands of miles away at the network edge.

High availability ensures uninterrupted productivity, supports the bottom line, strengthens security and compliance, and mitigates risk.

The IBM Z and Power Systems and Lenovo x86System x and ThinkSystem platforms have consistently delivered the highest levels of uptime availability based on every metric and measure of reliability over the past 11 years and six years, respectively.

So how have the IBM (since 2008) and Lenovo (since 2014) servers achieved and maintained best in class reliability status for over a decade?

IBM and Lenovo Reliability Success: Innovation, High Performance, Security and Top Technical Support Deliver High Reliability

Both IBM and Lenovo have mapped out cogent, consistent strategies and successfully executed against it.

This includes regular, planned product releases, an emphasis on innovation such as the ability to support more compute-intensive workloads, and advanced functionality to support emerging technologies like AI and analytics. IBM for example, typically ships a new Power processor approximately every three to four years.

At its Think 2019 Conference in San Francisco in February, IBM unveiled the newest POWER9 servers that are purpose-built to handle some of the world's most demanding workloads, from AI model building and training to data analytics, SAP HANA, and Oracle applications. For example, the IBM Power Systems AC922 capitalizes on POWER9 processors and NVIDIA GPUs connected through NVLink technology to help deliver the robust performance required for model development. The AC922 is the same server used as a building block for the Oak Ridge National Laboratory's Summit supercomputer — the world's fastest.¹

IBM and Lenovo's server strategies continually accentuate security for both the server hardware and associated OSes. Security constitutes one of the most daunting threats to server and OS stability and reliability. The IBM Z, Power Systems and LinuxONE, for example, provide pervasive encryption and SSC workload isolation for constant protection against internal and external threats.

Furthermore, both vendors received high marks for technical service and support from their installed base of customers, and responded quickly when issues arose. Both vendors' respective support organizations are established and stable.

HPE's Reliability Ratchets Up

IBM and Lenovo managed to avoid the management pitfalls that plagued rival HPE (formerly HP) prior to being split into two separate entities in 2016. Previously, HP spent the better part of a decade undergoing tumultuous management shifts, board infighting and shakeups, costly and disastrous acquisitions, and multimillion dollar payouts to fired or exiting company executives.

In 2016, at the time of the split, HP shed 30,000 jobs — many of them in product management, marketing, sales, and technical support. Following the split, HP Enterprise split again; it [spun off its IT services business](#) to save an estimated \$1 billion in operating costs. HP originally acquired the IT services business in 2008 with its \$14 billion acquisition of EDS Corp. HPE merged the former EDS services unit with Computer Sciences Corp (CSC) and created a new corporation called SpinCo. HPE took a 50% stake in SpinCo.

ITIC interviewed scores of HP and HPE IT managers over the past several years unhappy with being forced to cope with diminished technical support. IT departments reported longer lead times for HPE replacement parts and long wait times for technical documentation and phone support. This caused a decline in the reliability of HPE servers. ProLiant was particularly hard hit: It landed near the bottom of reliability scores in the last five ITIC Global Server Hardware Server OS Reliability surveys. **Exhibit 1** shows HPE ProLiant servers averaged 42 minutes of per server/per annum downtime in the ITIC 2019 Reliability poll. Only white box servers fared worse with 49 minutes of unanticipated downtime due to server problems or failures.

IBM technical support has been a bastion of stability. ITIC interviews found IT managers at IBM enterprises pleased with the quality and speed with which IBM tech support responded to issues.

Lenovo too, delivers excellent service and support, according to IT managers interviewed by ITIC. IT managers gave Lenovo high marks for technical support and responsiveness over the last five years following IBM's 2014 sale of its Intel Corp. x86 server businesses. And with good reason: When IBM and Lenovo inked the deal, the two vendors stipulated IBM would continue to service the Lenovo System x installed base for five years after the purchase date. This prescient move ensured a smooth, trouble-free transition, and no post-sale service hiccoughs.

HPE is back on track post-split and its mission-critical HPE Integrity line is demonstrating top notch reliability and uptime, rivaling IBM Power Systems and Lenovo System x and ThinkSystem. The HPE ProLiant platform however, continues to languish near the bottom of the server offerings in terms of uptime, due to protracted user upgrade cycles.


Besides the obvious technical merits of IBM Z, Power Systems and Lenovo System x and ThinkSystem hardware, server reliability is reinforced by the expertise of IT administrators in corporations that use IBM and Lenovo equipment. ITIC's reliability survey data shows IBM and Lenovo IT administrators typically have 10 or more years of experience. By contrast, firms (with exceptions) that use Dell, HPE ProLiant, Oracle, and unbranded white box server hardware are more likely to hire IT managers with one to five years' experience.

IBM and Lenovo server shops are less price sensitive than organizations that have a high percentage of commodity servers. ITIC survey data consistently finds organizations that purchase IBM, Lenovo, and other high-end mission-critical systems, like Huawei and niche market vendor Stratus Technology are willing to pay more for advanced features/functions.

A higher percentage of IBM and Lenovo organizations use a regular three-year upgrade cycle and retrofit IBM, Lenovo, Huawei, and HPE Integrity servers. This is crucial, since applications like AI, analytics, blockchain, IoT and virtual reality (VR) are compute and resource-intensive.

Exhibit 4 shows a majority (66%) of respondents find larger workloads negatively impact reliability; up 6% since 2018. Only 16% said server reliability was not hurt by higher workloads.

Exhibit 4. Two-thirds of Companies Say Higher Workloads Lessen Reliability



Source: ITIC 2019

Among 66% of users whose server reliability dropped, about 74% use commodity servers: white box, older Dell, HPE ProLiant, and Oracle hardware more than 3½ years old that have not been retrofitted/updated.


The message is clear: Commodity server users should not increase server workloads without retrofitting hardware. Do not delay upgrades or retain servers well beyond the three-year upgrade cycle. Over 60% of Dell PowerEdge and white box users keep servers for over three-and-a-half, four or even five years while increasing application workload. This is just asking for trouble. The exceptions to the rule: Very small businesses whose application environment remains static.

Other Notable Survey Findings

Among the other key survey highlights:

- **Unplanned server outages due to bugs in the OS:** This is a continuation of a trend seen in the prior three ITIC Reliability surveys. The 88% of IBM Power Systems and 87% of Lenovo System x and ThinkSystem users running Red Hat Enterprise Linux (RHEL), SUSE, or Ubuntu Linux experience fewer than one *unplanned* outage per server/per year due to any flaws/bugs in the OS.
- **IBM Z mainframe class servers running Z/OS or Linux rank “best in class” for reliability, accessibility, performance, and security among all server platforms.** The Z mainframes had highest reliability/uptime ratings across the board in terms of actual minutes of unplanned per server/per annum downtime. An 83% majority of IBM Z mainframe users achieved at least five nines (99.999%) reliability. IBM Z servers exhibit true fault tolerance, experiencing just 0.74 — less than one minute of *unplanned* per server, per annum annual downtime due to inherent problems with the server or component parts. That equates to just over eight seconds per month or “blink and you miss it,” two seconds of weekly downtime.
- **IBM Power Systems, Lenovo System x and ThinkSystem, HPE Integrity, and Huawei KunLun and FusionServer platforms running Linux recorded the least *unplanned* downtime** of less than two minutes to just over two minutes per server/per year respectively due to inherent flaws in server hardware or component parts among all mainstream Linux server platforms. HPE’s Integrity Superdome and the Huawei KunLun mission-critical server each also recorded 2 minutes of *unplanned* downtime related to inherent issues with the underlying hardware or its components.
- **Dell PowerEdge server reliability improves:** Dell’s PowerEdge, in particular, new to two-year old models, showed a marked uptick in reliability. Dell PowerEdge servers recorded 25 minutes of unplanned per server/per annum downtime due to inherent system and component flaws. Dell server uptime increased by an average of four minutes in this year’s poll compared to ITIC’s 2018 survey. Survey data and anecdotal evidence based on first-person interviews indicate the reliability and availability of Dell PowerEdge servers heavily depend on server age, configuration, and experience of IT administrators.

Exhibit 5. IBM Power, Lenovo x86 and Huawei KunLun Running Linux Tops in Enterprise Server OS Reliability in 2019



Source: ITIC 2019

Server Hardware Platform Overview

IBM Power Systems and IBM Z

IBM's POWER9 scale-out systems began shipping in March 2018. The line includes the S914, S922, S924, H922, H924, and L922. They support in-memory databases, advanced analytics, and cloud environments. POWER9 servers are also cloud-ready and include built-in PowerVM virtualization capabilities. POWER9 scale-out servers for IBM i, AIX, and Linux integrate into organizations' cloud and AI strategies to ensure high performance and RAS capabilities required to support mission-critical workloads like IBM's Db2 and Oracle databases and SAP HANA. IBM POWER9 servers incorporate the latest I/O technology; this includes 25 GB/sec high-speed interconnect for Coherent Accelerator Processor Interface (CAPI) and OpenCAPI along with embedded PCI-Express 4.0 connectivity. IBM claims this doubles the I/O bandwidth versus PCI-Express 3.0. POWER9 2-socket systems provide up to 4TB of memory, which IBM asserts is 33% greater than comparable Intel x86 Xeon systems, delivering additional benefit to in-memory databases such as SAP HANA. POWER9 continues in the tradition of previous generations by delivering improved per core performance capabilities compared to its predecessor.

IBM says, similar to the prior generation POWER8 systems, POWER9 delivers 1.25 to 1.5 times per core the performance capability, enabling clients to further reduce software spending on stacks that are licensed per core. This focus on core architecture and capabilities delivers over twice the per-core performance over compared x86 systems across a wide range of benchmarks. IBM says improvements to POWER core capabilities results in lower cost of acquisition and lower cost of ownership over the entire server product lifecycle. POWER9 architecture is open for licensing and modification by [OpenPOWER Foundation](#) members. The next version, POWER10, which will be based on 10 nm technology, is due out sometime in 2020. IBM aims to deliver a very high core count and high performance I/O.

The IBM Z mainframe is in a class of its own in terms of advanced features/functions, reliability and availability, and security. It has near-flawless reliability of 0.74 seconds per month and just less than nine seconds per server/per annum. IBM's Z portfolio of servers focus on pervasive encryption, cloud computing, and real-time analytics. IBM Z boasts virtually no downtime as a result of inherent system or component flaws. That alone makes it worth its weight in gold. In spring 2018, IBM unveiled the z14 ZR1 and LinuxONE Rockhopper II, single-frame systems. IBM continues to advance Z capabilities with more power in a smaller footprint. IBM z14 ZR1 can process over 850 million fully encrypted transactions per day in the space of two floor tiles. The IBM z14 portfolio delivers end-to-end encryption at rest and in "flight" or transmission.

Lenovo System x and ThinkSystem

Lenovo Data Center Group's (DCG) newest strategy focuses on leveraging ongoing partnerships with best-in-class next-generation technology partners like SAP, Nutanix, and others. More recently, Lenovo DCG embraced "a people-first" approach to support customers and channel partners at the local and global level to drive technology adoption and deployment. In the five years since Lenovo purchased IBM's x86 line of servers, it maintains a high level of performance, reliability and after-market technical service and support. Lenovo's System x and ThinkSystem averaged just under two minutes of unplanned per server/per annum downtime due to inherent flaws in server or component parts. Lenovo responded to the intensely competitive nature of the server market by doubling down on its commitment to core reliability, availability, and performance in System x servers. This continues with the Lenovo ThinkSystem server portfolio announced in mid-2017. DCG [recently announced](#) its fifth consecutive quarter of profit growth on revenue of \$1.6 billion. Fueled by success, Lenovo seeks to aggressively expand its presence in the battle for server market share. At the 2019 Mobile World Congress in Barcelona in February, Lenovo showcased its new line of IoT-related products and services targeting the red-hot network edge/perimeter market, where many IoT deployments are taking hold.

ThinkSystem SE350

[ThinkSystem SE350](#) is Lenovo's first purpose-built edge server, targeted at the network edge to deliver optimal bandwidth, bolster security, and reduce downtime. It's a small-footprint server measuring 1.75 inches high, 8.1 inches wide and 14.9 inches deep that can be mounted on a wall, stacked on a shelf or installed in a rack. ThinkSystem SE350 is a high-performance server based on Intel's [Xeon-D](#) processor with 256GB of RAM and 16TB of internal solid-state storage. ThinkSystem SE350 has enhanced physical security features like a locking bezel, intrusion detection, tamper detection, and encrypted storage. It boasts zero-touch deployment software.

Lenovo says it is combining innovation with reliable, flexible, and secure data center systems. It emphasizes its open server, storage, networking, and system management platforms seamlessly integrate with existing and legacy environments. This is critical, since ITIC's latest survey data indicates human error, flaws in server OS software, and complexity in provisioning new applications are pivotal issues negatively impacting reliability. In first-person interviews with ITIC analysts, Lenovo customers lauded ease of deployment and integration and backwards compatibility as contributing to reliability and stability of System x and ThinkSystem platforms. Lenovo users also had high praise for after-market service and support.

Lenovo customers interviewed by ITIC for the 2019 Reliability poll and 2018 Mid-Year Update praised Lenovo's "agile system design" and hot swap capabilities. IT can easily upgrade for scalability on-demand via front and rear access, and add components without removing the server from the rack. Lenovo's server design is resilient. It maximizes application uptime and easy integration in virtual environments. This is a win-win for customers and Lenovo.

HPE Integrity

HPE's Integrity Superdome line of servers also exhibit exceptional reliability and are in a virtual dead heat with IBM, Lenovo and Huawei KunLun distributions for least amount of per server/per annum downtime, averaging around two minutes. They benefit from the renewed stability of the HPE brand and business nearly two years after Hewlett-Packard split into two separate entities. This enabled HPE to renew its focus on innovation and after-market technical service and support, crucial in the current digital computing environment. HPE is a legacy server brand entrenched in corporate enterprises, and built a solid reputation over six decades. Now that the corporate tumult of last 12 years is behind it, HPE can concentrate on innovation, which it is doing with a vengeance, evidenced by reliability ratings of HPE Integrity Superdome servers.

Huawei KunLun and FusionServer

Huawei is among the newest entrants to challenge established players in the high-end server market with its three-year-old KunLun family of 16- and 32-socket mission-critical servers. Huawei is a top five vendor in global server hardware shipments. In three years, Huawei, based in Shenzhen, China, emerged as a top five server hardware vendor worldwide, with high-end KunLun mission-critical server and general purpose FusionServer x86-based servers.

Huawei's rise in the server market —except for North America — is due to a multipronged series of bold tactical moves and its ability to formulate and execute a compelling strategy for long-term success. Huawei servers support a wide range of needs, from general-purpose rack and blade servers to mission-critical hardware to address high-performance computing (HPC). Advanced capabilities support emerging compute-intensive applications like AI, big data analytics, deep learning, and machine learning. According to the 2018 EU Industrial R&D Investment Scoreboard report,² Huawei was fifth in global R&D investment among 2,500 companies surveyed. In 2018, Huawei's R&D hit \$13.1 billion; 15% of annual sales. Huawei's R&D spending bested both Intel with \$12.6 billion and Apple with \$11.1 billion, with about one billion going to server technology. In July 2018, executives announced Huawei's R&D budget for 2019 and beyond³ would reach \$15 to \$20 billion annually.

² The European Commission, 2018 EU Industrial R&D Investment Scoreboard. URL: <http://iri.jrc.europa.eu/scoreboard18.html>

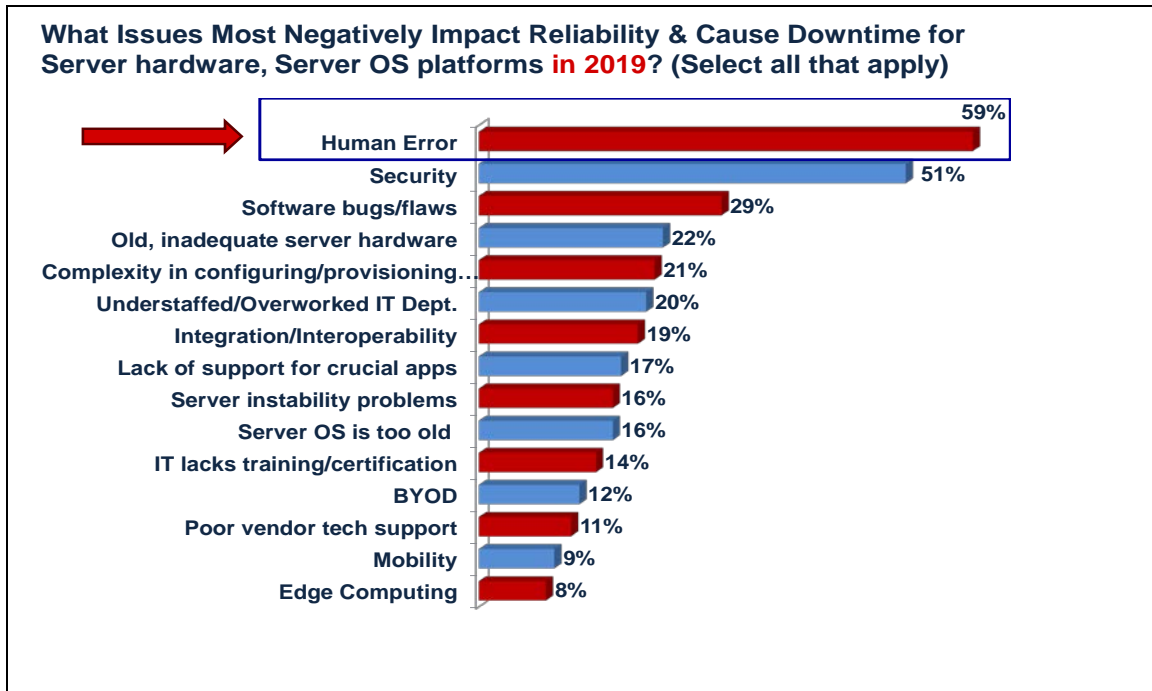
³ Reuters, "Huawei to raise minimum annual R&D spending to at least US\$15 billion," July 26, 2018. URL: <https://scmp.com/tech/social-gadgets/article/2157024/huawei-raise-minimum-annual-rd-spending-least-us15-billion>

2019 Reliability Trends

Human Error and Security Threats Top the List

As **Exhibit 6** illustrates, 2019 survey results indicate human error and security top the list as issues that most negatively impact reliability, followed by 29% who cited software bugs and 22% of users who said old, inadequate server hardware precipitated downtime.

Exhibit 6. Human Error and Security Top Causes of Downtime



Source: ITIC 2019

Human error and security are inextricably linked. It's no surprise these two issues are the top two causes of downtime. ITIC expects that trend continue. Happily, one trend that has not been a big issue thus far in 2019, since the appearance of the infamous [Meltdown and Spectre](#) security attacks, which appeared back in January 2018.

These security attack methods, which rely on vulnerabilities roughly two decades old, allow malicious applications to bypass memory isolation mechanisms and access passwords, photos, documents, emails, and other sensitive data. The bulk of the publicity (surrounding Meltdown and Spectre) focused on Intel, which commands 95% of the microprocessor market share. However, corporate users would do well to remember that nothing is "hack proof"; every processor is vulnerable. Meltdown and Spectre security attacks can be launched against systems using Advanced Micro Devices (AMD), ARM, and others.

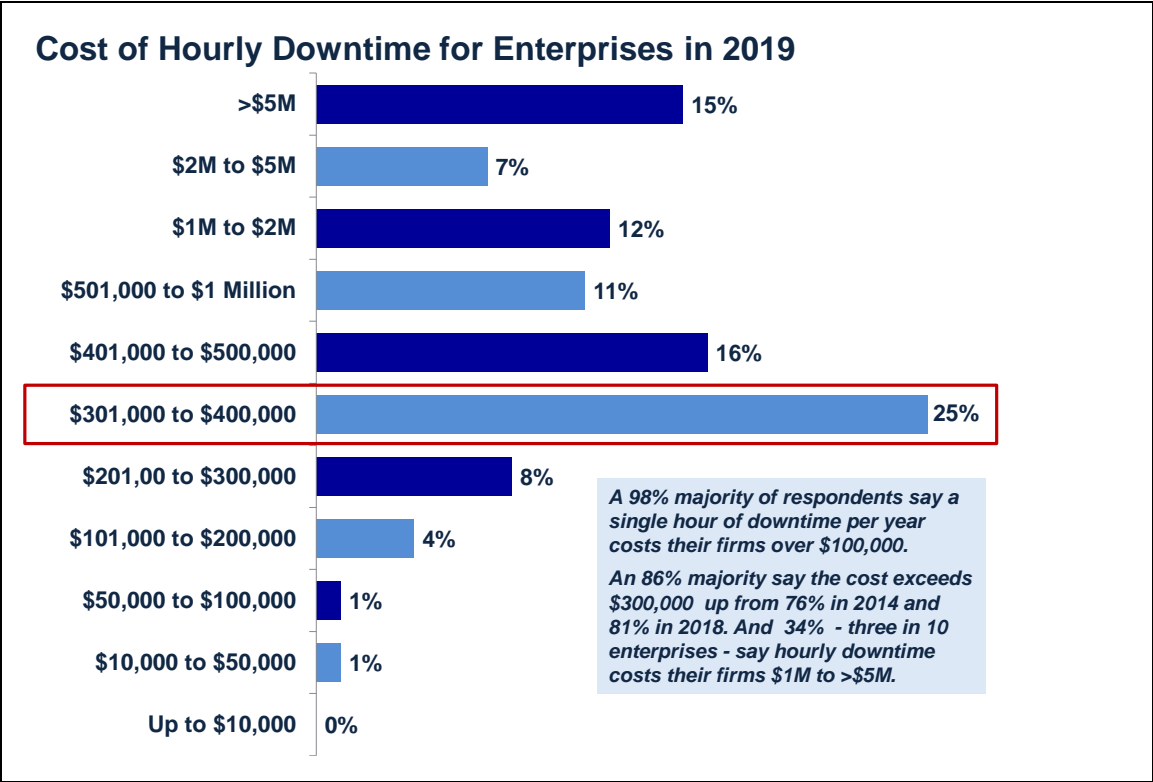
Security attacks are commonplace. The Meltdown and Spectre attacks are a reminder that old technology can lay dormant for years and still rise up to wreak havoc. It is crucial that businesses closely communicate with vendors and stay up to date on news of latest patches and fixes to find the most effective mechanism for dealing with hardware and software-based security flaws.

Hourly Cost of Downtime Continues to Rise

Hourly downtime costs continue to increase for all businesses irrespective of size or vertical market. This trend has been evident over the last five years. Given organizations’ near-total reliance on systems, networks, and applications to conduct business 24x7, it is safe to say the cost of downtime will continue to increase for the foreseeable future.

Although large enterprises with over 1,000 employees may experience the largest actual monetary losses, downtime can be equally devastating to small and mid-sized businesses that typically lack the financial resources of larger firms. An hour or two of downtime, or even 30 minutes of downtime during peak usage hours, disrupts productivity and business transactions, can deal SMB firms a crippling monetary blow that could potentially put them out of business.

Exhibit 7. Hourly Downtime Costs Exceed \$300K for 86% of Companies



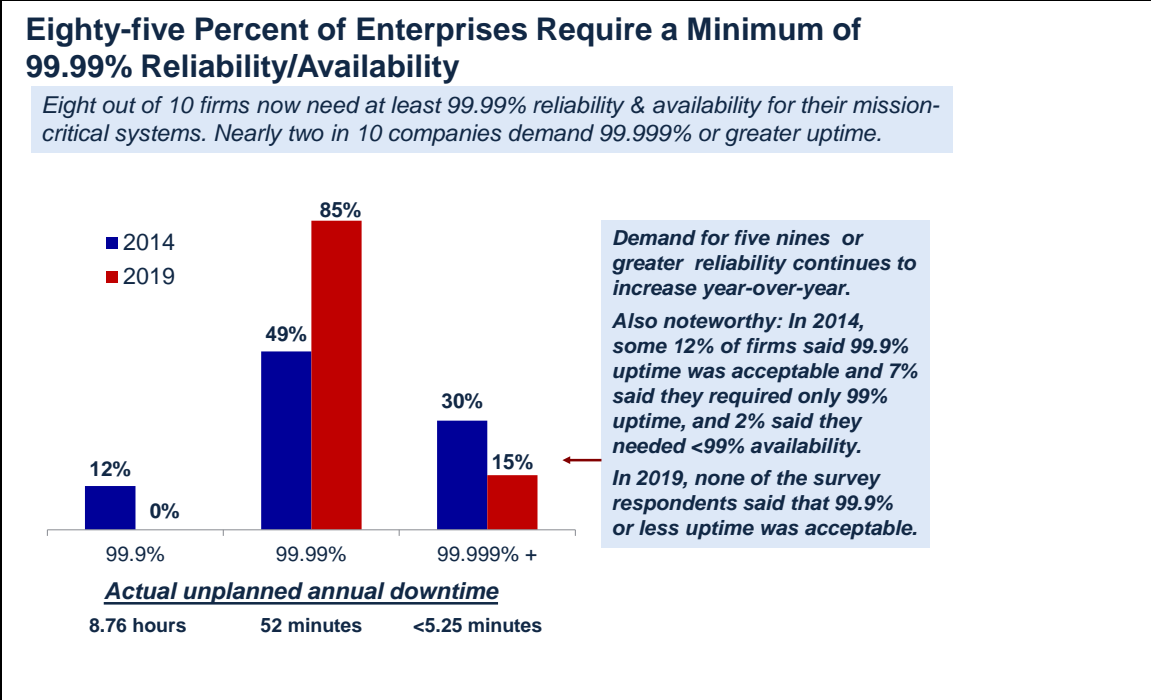
Source: ITIC 2019

Minimum Reliability Requirements Rise

Similarly, based on all of the other trends, the latest ITIC 2019 Global Server Hardware and Server OS Reliability survey showed more than eight out of 10 businesses — 85 % — now require a minimum four nines (99.99%) reliability and uptime. This is more than double the 39% that said their firms required 99.99% availability in ITIC’s 2013 poll, six years ago.

For the last several years four nines of uptime has been the minimum 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 mission-critical servers and main LOB applications (see Exhibit 3). The remaining 5% of enterprises indicated their firms need six nines (99.9999%) or better availability/uptime.

Exhibit 8. Minimum Reliability Requirements Rise



Source: ITIC 2019

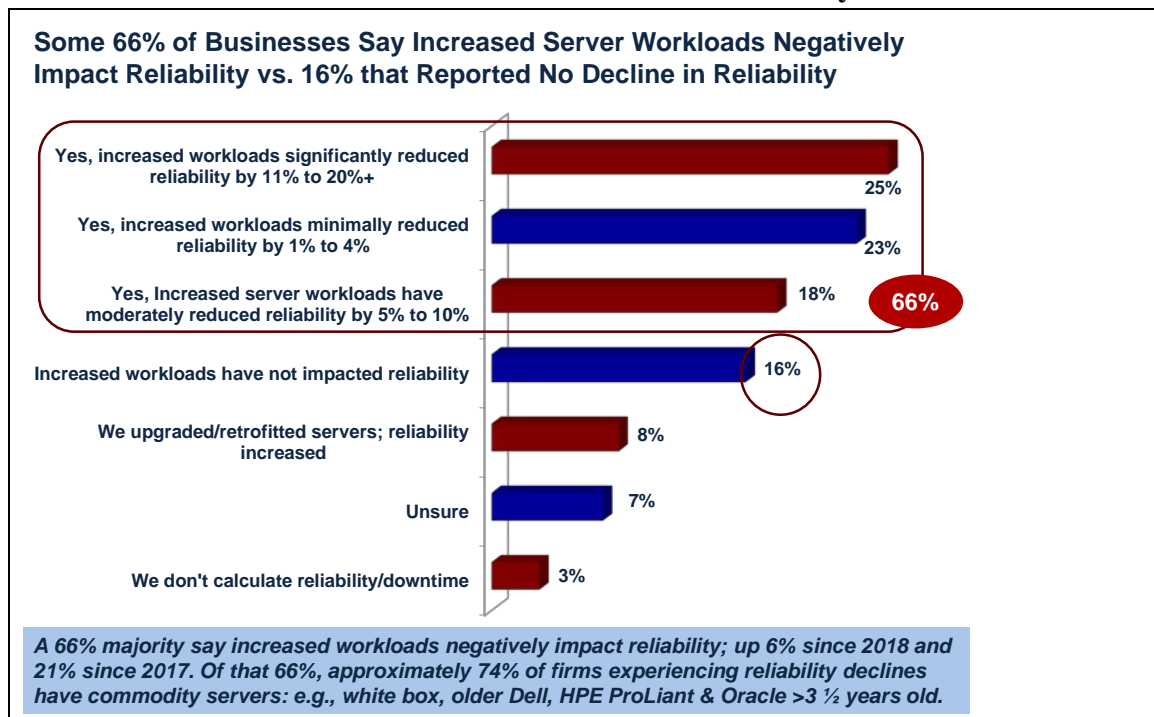
Once again, we can expect businesses demand for continuous uptime and availability will continue unabated.

Higher Server Workloads Negatively Impact Reliability

The most recent ITIC 2019 Reliability survey revealed server workloads have increased an average of 34% since 2014. Consequently, two-thirds — 66% — of organizations indicated that increased server workloads negatively impacted reliability and uptime. That is an increase of 21% from the ITIC 2017-2018 study. By contrast only 16% of those polled said increased workloads had no impact on server reliability As **Exhibit 8** shows, 25% of corporations indicated increased server workloads significantly cut reliability by 11% to 20%+; some 23% of respondents said higher server workloads minimally undercut reliability by 1% to 4%.

Only 16% of respondents said increased server workloads had **no impact** on reliability, a decline of 6% from the 2018 ITIC Reliability survey. Notably, 87% of respondents who reported no discernable impact on server uptime after upping workloads hail from high end, mission-critical server distributions. These include IBM, Lenovo, Stratus ftServer, HPE Integrity, Huawei's KunLun and Fujitsu Primequest.

Exhibit 9. Increased Server Workloads Undermine Reliability



Source: ITIC 2019

Hardware Continues to Evolve, and So Does Risk

Overall, the *inherent* reliability of the majority of server hardware platforms, server OSes 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. Technologies including AI, analytics, blockchain, cloud computing, IoT, security, virtualization, and VR are more compute intensive and place greater demands on core infrastructure components than technologies dating back to the 1990s, early 2000s or even those as recent as 2015. All of them can tax system resources to the breaking point *if* organizations fail to upgrade and right size hardware, OSes and applications to accommodate higher demands.

Breakdown in x86 server reliability generally (but not always) occurs *not* because of inherent flaws in underlying server hardware, but because x86 customers “push their luck” and retain server hardware for 4½ to sometimes 6+ years without retrofitting or upgrading. A business that overloads outmoded servers or misconfigures a server will experience availability problems. This is particularly true of organizations that purchase entry level or inexpensive commodity servers and fail to rightsize them to accommodate compute-intensive applications and workloads.

Commodity Hardware Can Have Expensive Consequences

Inexpensive and commodity hardware definitely has a place in organizations. However, the device must align with its usage within the business. Purchasing the least expensive device to keep capital expenditure costs down won't save the corporation money in the long run, if the inexpensive server configuration is not robust enough for the application and workload.

Shaving a few hundred or a few thousand off the upfront purchase price of a server will be of no avail if the server proves inadequate for the task and suffers increased downtime. Consider that in 2019 one hour of downtime costs an estimated \$300,000 or approximately \$5,000 for a one-minute outage and \$25,000 for five minutes of downtime. Those five minutes of downtime calculated at \$300,000 or more, are equivalent to the price of a fully loaded server. Keep in mind that many — albeit not all — outages involve multiple servers.

These examples are average estimates and will vary according to corporate implementations and use cases. They do not include worst-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 do not take into account legal fees, penalties or financial settlements that result from litigation.

Robust, Reliable Servers Pay Off

All the statistics presented in ITIC's 2019 Global Server Hardware and Server OS Reliability Survey and Report reinforce the necessity of purchasing and deploying the most reliable servers in the most robust configuration that the company's budget will allow. The underlying reliability and availability of the crucial infrastructure server hardware and server OS has a direct and immediate impact on the health of the entire network ecosystem and all associated data, services, daily business operations, transactions, and revenue.

Downtime is disruptive and expensive. It can also irreparably damage a company's reputation. In 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.

Conclusions

ITIC 2019 Global Server Hardware and Server OS Reliability Survey findings indicate that for the 11th straight year, the IBM Power Systems along with the Lenovo System x and ThinkSystem servers for the sixth straight year, are the most reliable hardware platforms. Each registered 99.99% and 99.999% for well over nine out of 10 corporate survey respondents. The IBM Power Systems and Lenovo System x platforms are up to 24x more reliable (each registering just under two minutes of downtime) than the worst performing unbranded white box servers.

The strong reliability of the IBM Power Systems, Lenovo x86 System x and ThinkSystem, HPE Integrity, and Huawei KunLun hardware platforms yield demonstrably better uptime and availability resulting in lower TCO and accelerated ROI.

Both IBM and Lenovo continue to advance the core functionality of their hardware with inherent reliability, management, and security to support the demands of high transactional workloads and emerging technologies like AI, analytics, cloud computing, IoT, and VR.

The IBM Z mainframe is in a class by itself. No other server platform (excluding super computers and HA machines) even come close. An over 80% majority of IBM Z customers said their businesses achieved unparalleled fault tolerant levels of 99.999% and 99.9999% reliability and availability.

In another notable achievement, IBM and Lenovo were either first or second in every reliability and availability category or tied for first or second place in every uptime, security, or manageability metric in the survey. This includes:

- **Highest server availability ratings:** When IBM Power Systems and Lenovo System x and ThinkSystem servers did experience an outage due to inherent problems with server hardware or component parts, they were short. It typically took IT managers under 10 minutes and in most cases, approximately five minutes or less to bring them back online.
- **Least amount of server downtime of over four hours:** Only 1% of IBM Power Systems, Lenovo System x and ThinkSystem, along with the HPE Integrity and Huawei KunLun mission-critical server experienced downtime due to inherent system or component parts failures.
- **Most reliable enterprise server OS:** IBM Power Systems, Lenovo System x and ThinkSystem, and Huawei KunLun servers running Linux and open source distributions each delivered the least amount of unplanned downtime associated with system or OS technical problems, less than 12 minutes per server/per annum.
- **Least amount of unplanned server downtime:** IBM Power Systems, Lenovo System x and ThinkSystem, HPE Integrity, and Huawei KunLun servers all delivered the least amount of unplanned server downtime in 2019 (in hours per year) in comparable hardware configurations and workloads — all 12 minutes or less.

Mission-critical systems from HPE and Huawei continue to make impressive gains. HPE and Huawei distributions achieved near reliability parity with IBM and Lenovo, and enjoy a very high degree of customer satisfaction.

The reliability of newer Dell PowerEdge servers (new to 2½ years old) rebounded significantly, and Dell's overall reliability scores improved. However, Dell's higher levels of downtime are the direct result of corporate end users overtaxing server resources and retaining servers for 3½ years or more without upgrading or retrofitting to accommodate higher workloads.

Reliability is fluid, not static. No server, component, hard drive, memory or CPU, OS, application, device, or connectivity mechanism is immune from inherent problems or failure. Any system can be hacked and the reliability of any system can be compromised or undone by human error. The overarching ecosystem, with the server and OS infrastructure as its foundation, is becoming more complex. This presents challenges for vendors and corporate enterprises alike to easily deploy, provision and manage systems, and stay up to date.

This increased complexity comes at a time when organization requirements for unassailable reliability and near flawless uninterrupted availability is greater than ever.

Business applications, services, and data must be available round the clock, 365 days a year. Organizations are increasingly less tolerant and more risk averse to *any flaws* in their core infrastructure systems that disrupt productivity.

An 85% majority of ITIC 2019 Global Server Hardware and Server OS survey respondents across all vertical and horizontal markets, now require a minimum of at least four nines (99.99%) uptime. That equates to just 52.56 minutes of downtime per server/per annum or a scant 4.32 minutes of monthly downtime for each server.

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. The corporation must devote the necessary capital, operational expenditures, and labor resources to ensure peak levels of reliability. Achieving optimum uptime means upgrading refreshing server hardware as necessary 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 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 customers. The ability to meet SLAs hinges on server reliability, uptime, and manageability. These are key indicators that enable organizations to determine which server OS platform or combination thereof is most suitable.

Businesses must regularly replace, retrofit, and refresh server hardware and server Oses with the necessary patches, updates, and security fixes *as needed* to maintain system health. The onus is also on the server hardware and server OS 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 inform customers to the best of their ability about known incompatibility issues that may impact performance. Vendors should also be honest with customers in the event there is a problem or delay with delivering replacement parts.

Once again, time *is* money. In the Digital Age, even a few minutes of downtime can result in significant monetary and data losses and cause internal business operations to come to a standstill. Downtime can also impact adversely a company's relationship with customers, business suppliers, and partners. Reliability or lack thereof can potentially damage a company's reputation and result in lost business.

Recommendations

Corporate enterprises have every right to expect server vendors to continually improve hardware reliability and deliver a minimum of 99.99% uptime and availability and above. Vendors must also quickly address and resolve technical issues and security flaws when they arise and deliver advanced features/functions, and provide the necessary guidance and expert after-market technical service and support.

Organizations' ability to achieve four, five, and six nines of server and OS reliability is a two-way street, not a one-way footpath.

It is the responsibility of corporate enterprises to select the servers, OSES and applications in the correct configuration. Companies should also purchase servers that can accommodate not just current application workloads, but a robust platform that will flexibly scale to accommodate business needs over the next two to three years. Corporations must also employ IT and security administrators who have the skill to properly provision, upgrade and maintain a high degree of daily operational efficiency. Appropriate training and certification for IT and security administrators is also necessary.

ITIC strongly advises organizations to regularly measure the uptime and reliability of their main LOB server hardware, server OSES, 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 provides businesses with a mechanism to accurately assess the amount of downtime and its impact on business operations, the IT department, and end-user productivity. 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 enables companies to determine whether or not current server and server OS environment allows them to achieve optimal reliability.
- **Keep a comprehensive record of downtime.** IT departments should compile a detailed list of outages. Include facts like the cause of the outage (e.g., hard drive failure, human error, manmade disaster, etc.), the length/duration of downtime, and the severity of the event (e.g., lost, damaged, or stolen data, and interrupted transactions).

- **Closely monitor and track security issues and outages.** Security is one of the chief causes of downtime. Hacks are more pervasive and pernicious and hackers are more proficient and very well organized. An increasing number of connections and end points mean more potential vulnerabilities. Organizations should get training for their security administrators; stay up to date on security patches and fixes and work closely with their vendors to solicit guidance.
- **Don't delay updates.** Refresh and upgrade server hardware as needed to accommodate more data intensive and virtualized workloads. The server hardware (standalone, blade, cluster, etc.) and the server OS are inextricably linked. To achieve optimal performance from both components, corporations must ensure that 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 hardware is asking for trouble.
- **Calculate the cost of hourly downtime.** There is no "one size fits all." Hourly downtime costs will vary according to the length, severity, and duration, and whether data was lost, stolen, destroyed, or changed. A 15 or 20-minute outage that occurred in off-hours may have negligible consequences, while a server that goes down for three minutes disrupting a crucial transaction can potentially cost the business thousands or even millions.
- **Adopt formal SLAs.** SLAs 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 indicate nearly half of respondents (49%) do not calculate the hourly cost of downtime. This is a big mistake. Keep accurate records of outages and causes. Segment outages according to severity and length (e.g., Tier 1, Tier 2, and Tier 3). The appropriate IT managers should keep detailed logs of remediation efforts. Logs should include a full account of remediation activities, specifying how the problem was solved, how long it took, and what staff members participated. It should list costs as well as any material impact on the business, operations, and end users. This is an invaluable resource should the problem recur. It may help contain and minimize reliability-related incidents.

- **Calculate the cost of unplanned downtime.** Companies should determine the average cost of minor Tier 1 outages. They should also conduct detailed cost assessments of the extended and more severe unplanned Tier 2 and Tier 3 incidents. 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 and record whether litigation ensued or 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.
- **Construct 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 products they 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.

Appendix

Survey Methodology

ITIC's *2019 Global Server Hardware and Server OS Reliability Survey*, polled C-level executives and IT managers at over one thousand corporations worldwide from November 2018 through January 2019. The independent web-based survey included multiple-choice questions and one Essay question. To maintain 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 obtain valuable anecdotal data and gain deeper insights and contextual knowledge of the issues affecting corporate enterprises including C-suite executives, IT and security administrators and end users. ITIC employed authentication and tracking mechanisms to prevent tampering and to prohibit multiple responses by the same parties.

Survey Demographics

ITIC polled companies of all sizes and across 40 vertical markets for the survey. Respondents came from companies ranging from small and medium businesses (SMBs) with fewer than 50 workers, to multinational enterprises with over 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. 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.

Links to Supplemental Information

This section contains links to the various ITIC statistics and surveys cited in this report.

ITIC website and links to survey data and blog posts:

<https://itic-corp.com/blog/2018/11/1678/>

<https://itic-corp.com/blog/2018/08/itic-poll-human-error-and-security-are-top-issues-negatively-impacting-reliability/>

<https://itic-corp.com/blog/2018/08/itic-2018-server-reliability-mid-year-update-ibm-z-ibm-power-lenovo-system-x-hpe-integrity-superdome-huawei-kunlun-deliver-highest-uptime/>

<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/>