

INFORMATION TECHNOLOGY
INTELLIGENCE CONSULTING

Information Technology Intelligence Consulting



**ITIC 2020 Global Server Hardware,
Server OS Reliability Report**

April 2020

Table of Contents

Executive Summary.....3

 Introduction6

 Reliability and Uptime by the Numbers.....7

 Reliability versus Availability9

 ITIC 2020 Reliability Results9

 Reliability Dollars and Sense: The Actual Cost of Downtime.....15

 Analysis17

 IBM and Lenovo Reliability Success: Innovation, High Performance, Security and
 Top Technical Support Deliver High Reliability18

 Other Notable Survey Findings23

 Server Hardware Platform Overview24

 IBM Z24

 IBM Power Systems.....26

 Lenovo ThinkSystem27

 HPE Integrity29

 Huawei KunLun and Fusion Servers.....29

 2020 Reliability Trends: Top Notch Security is Crucial30

 Hourly Cost of Downtime Continues to Rise.....31

 Minimum Reliability Requirements Rise31

 Conclusions.....34

 Recommendations36

 Survey Methodology39

 Survey Demographics39

 Appendices.....40

Executive Summary

IBM, Lenovo maintain top server Reliability ranking for 12th and seventh straight year, respectively

Cisco, Hewlett-Packard Enterprise (HPE) and Huawei close gap, challenge leaders with strong reliability showings

IBM Z, Power Systems, Lenovo x86 and Huawei KunLun hardware deliver highest Availability; faster resiliency with outages of shortest durations; service interruptions

IBM Z, IBM Power Systems, Lenovo ThinkSystem, HPE Integrity and Huawei KunLun deliver up to 28xs better reliability than least efficient rival “White box” platforms and 34xs better economies of scale

Security/Data Breaches; Human Error; Software bugs are top Reliability threats

High reliability, uptime and availability are imperative in today’s “always on” Digital networks.

For the 12th straight year, IBM’s Z and Power Systems, achieved the highest server; server operating system reliability and server application availability rankings, along with Lenovo’s ThinkSystem servers which delivered the best uptime among all Intel x 86 servers for the last seven consecutive years. Those are the results of ITIC’s latest Global Server Hardware and Server OS Reliability survey, which polled 1,200 global enterprises from January through March 2020.

ITIC’s latest survey data finds that the elite IBM z14 and z15 servers are in a class of their own, averaging just 0.62 seconds of monthly/7.44 seconds of annual unplanned per server downtime due to bugs in the underlying hardware. This is consistent with very high availability and fault tolerance of “six nines” – 99.9999% and greater. The IBM Z family – the “Z” stands for zero downtime - consistently outperforms **all** competitors in every reliability category and delivers the lowest total cost of ownership (TCO) and fastest return on investment (ROI).

Additionally, IBM’s Power Systems, Lenovo ThinkSystem, Hewlett-Packard Enterprise (HPE) and Huawei KunLun, respectively, also scored high reliability marks, recording up to 28xs more uptime and availability than the least dependable unbranded “White box” servers. The superior uptime of the above top ranked mission critical hardware platforms makes them up to 34xs more economical and cost effective than the least stable White box servers.

The mission critical IBM Power Systems and Lenovo x86-based ThinkSystem both notched improvements in their 2019 reliability scores; each registered under two (2) minutes of per server unplanned downtime due to inherent flaws in the underlying hardware or component parts. Cisco, Hewlett-Packard Enterprise (HPE) and Huawei server platforms were close behind: each recorded approximately two minutes or a few seconds more downtime attributable to inherent issues with the hardware (**See Exhibit 1**).

The least consistent hardware - unbranded white box servers - averaged 53 minutes of unplanned per server downtime due to problems with the server or its components (e.g. hard drive, memory, cooling systems etc.). This represents an increase of four (4) minutes of downtime compared with ITIC's 2019 Global Server Hardware, Server OS Mid-Year Update survey.

ITIC's independent Web-based survey, (which also included an essay question and first person customer interviews), compares and analyzes the reliability and availability of over one dozen mainstream server platforms and one dozen operating system (OS) distributions. To obtain the most accurate and unbiased results, ITIC accepts no vendor sponsorship.

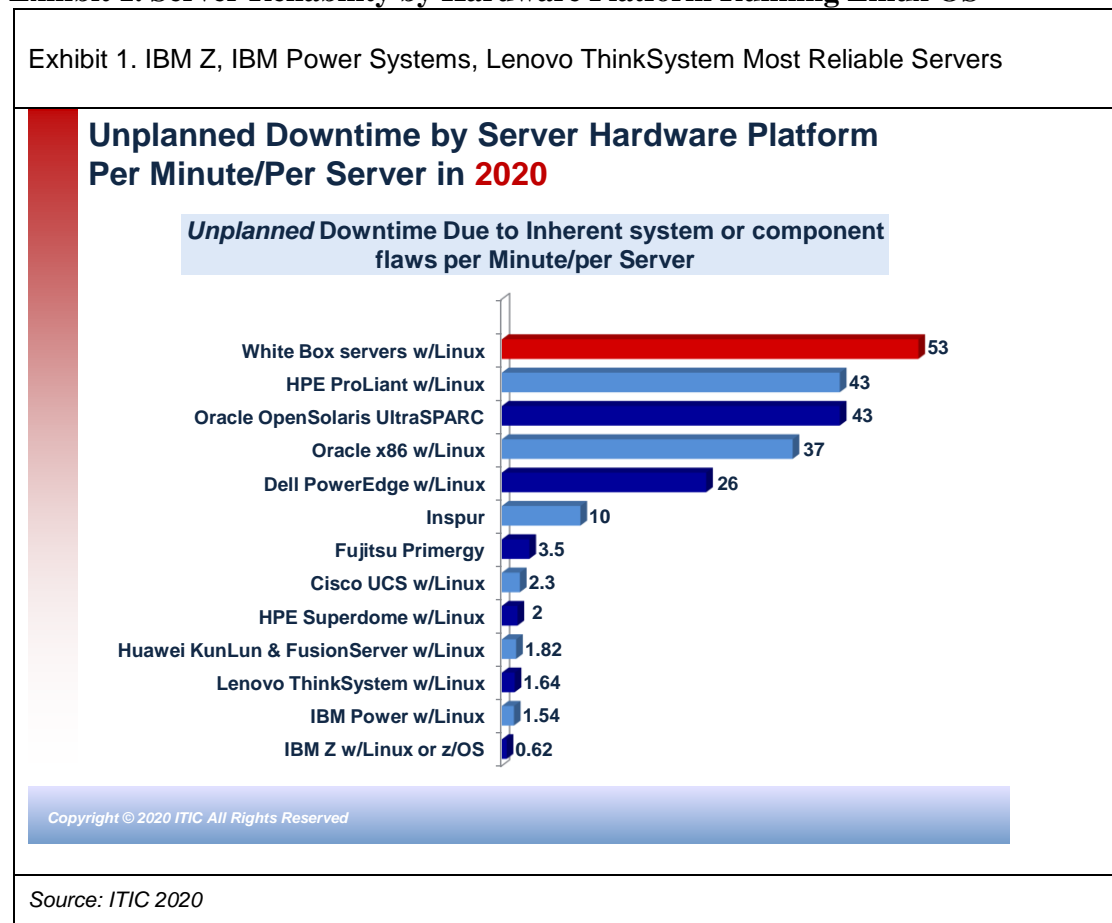
Among the other top survey findings:

- **Reliability:** IBM Z, IBM Power Systems and Lenovo ThinkSystem hardware and the Linux operating system distributions were once again either first or second in every reliability category, including server, virtualization and security.
- **Availability:** IBM Z, Power Systems, Lenovo ThinkSystem, HPE Integrity and Huawei KunLun all provided the highest levels of server, applications and service availability. That is, when the 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 one percent of IBM and Lenovo customers and two percent 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 remain 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's KunLun platforms experienced the fewest hard drive quality or failure issues among all of the server distributions within the first one, two and three years of service. Less than one percent – 0.4% - of IBM Z servers experienced technical problems with their hard drives in the first year of usage, followed by the IBM Power Systems and Lenovo ThinkSystem with one percent (1%) each during the first year.
- **Security is Top External Issue Negatively Impacting Reliability:** Security and data breaches now have the dubious distinction of being the top cause of downtime cited by 64% of survey participants, followed by “user error” at 60%.
- **Minimum Reliability Requirements Increase:** An 88% majority of corporations now require a minimum of “four nines” of uptime - 99.99% for mission critical hardware,

operating systems and main line of business (LOB) applications. This in an increase of five (5) percentage points from ITIC’s 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, Ransomware, CEO fraud as well as malware and viruses.
- **Increased Server Workloads Cause Reliability Declines:** The survey found that reliability declined in 67% of servers over four (4) years old when corporations failed to retrofit or upgrade the hardware to accommodate increased workloads.
- **Hourly Downtime Costs Rise:** A 98% majority of firms say hourly downtime costs exceed \$150,000 and 88% of respondents estimate hourly downtime expenses exceed \$300,000. Four-in-10 ITIC survey respondents - 40% - estimate the cost of a single hour of downtime is over one million (\$1,000,000).

Exhibit 1. Server Reliability by Hardware Platform Running Linux OS



Source: ITIC 2020 Global Server Hardware Server OS Reliability Survey

The server hardware, server OS and application infrastructure is the bedrock upon which the reliability and uptime of the entire organizational operation rests. Very high reliability and uptime is imperative to ensure uninterrupted and continuous daily business operations.

©Copyright 2020 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.

Introduction

For the past 12 years, the ITIC Global Server Hardware, Server OS Reliability Report has compared the reliability of up to 18 mainstream server platforms, and over one dozen server operating system 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 improve 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 the 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 the 2020 Global Server Hardware, Server OS Reliability survey questions to delve into crucial external issues like Security and Data breaches that can wreak havoc and cause server and application downtime and unavailability. The latest poll also incorporates queries that address the “Availability” metrics or the length of unplanned outages of each server and OS distribution. And continuing the question first introduced in the 2019 Reliability poll, ITIC’s 2020 Global Reliability study provides a detailed breakdown of the monetary costs associated with Hourly Downtime. ITIC calculated per minute/per server downtime costs ranging from \$10,000 to \$10 Million in various server quantities encompassing a single server to one thousand (1,000) servers.

As always, ITIC's 2020 Global Server Hardware and Server OS Reliability Report utilizes information gathered from previous ITIC surveys to compare, contrast and analyze the reliability of the various platforms and to track current and future trends. These findings help organizations make informed purchasing/upgrade decisions aligned to their specific business and budget needs.

Reliability and Uptime by the Numbers

There is no such thing as too much reliability and availability.

The technical and business metrics associated with infrastructure reliability have all steadily and inexorably increased based on wide ranging trends and use factors. These include:

- **Demand for Continuous Operations:** Since 2013 the demand for high reliability of “four nines” or 99.99% uptime has increased by 46% percentage points. At that time, 39% of ITIC 2013 Reliability survey respondents required 99.99% uptime. ITIC's 2020 Reliability poll finds that 87% of respondents consider 99.99% - which equals 52.56 minutes of unplanned per server/per annum downtime -to be the minimum acceptable level of reliability for mission critical servers and applications.
- **Larger Applications:** Additionally, applications are getting larger and more compute-intensive and consuming more system resources.
- **More Compute-Intensive Server Workloads:** Mission critical business workloads increased an average of 15% to 36% over the last three years, depending on specific vertical markets and individual enterprise configurations. This trend will continue.
- **Average Hourly downtime costs:** Continue to rise with no end in sight. Currently, 88% of ITIC 2020 survey respondents said that single hour of downtime costs their firms in excess of \$300,000. And four-in-10, or 40% - of respondents estimate that 60 minutes of downtime costs their companies one million (\$1,000,000) or higher. The price tag for a single minute of downtime will vary widely depending on specific use configurations, as well as the timing, duration and severity of a particular outage. The cost of one minute of per server downtime ranges from \$1,670 for an hourly outage of \$100,000 to \$16,700 per server/per minute for an hourly outage cost estimated at one million (\$1,000,000) (See **Table 2**). Even a small business that gauges the cost of downtime at a “modest” \$10,000 an hour will rack up outage fees of \$167 per minute for a single server! These figures are **exclusive** of the costs associated with potential legal fees, fines or civil or criminal penalties that may arise as the result of the outage.
- **Increase in connected devices and networks.** An estimated 39 billion devices will be connected via the Internet by the end of 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. But connectivity does present its own set of risks. Server failures, particularly in cloud and virtualized environments have a domino effect because they impact greater numbers of servers and business critical applications. This potentially results in higher collateral

damage due to productivity and revenue losses. It also places organizations at heightened risk for litigation with customers and regulatory agencies over non-compliance.

- **Security and data breaches.** Security issues – including targeted data breaches, Ransomware, Business Email Compromises (BEC), CEO fraud, Phishing attacks and malware are now the top external causes that undermine reliability. The research firm, [Statista](#), headquartered in Hamburg, Germany which tracks the number of data breaches and records exposed, reported that hacks rose exponentially in the last 19 years. In 2005 for example, Statista¹ reported 157 cyber security incidents in the U.S. that exposed 66.9 million records. By 2014, that number increased 500%: there were 783 data breaches reported that exposed 85.61 million total records. The most recent 2019 Statista study found that the U.S. experienced 1,244 data hacks exposing over 446.5 million records. Additionally, the latest [Federal Bureau of Investigation \(FBI\) statistics](#) indicates that Email BEC and Phishing scams caused \$26 billion in losses from 2013 to 2019².
- Over 160,000 data-breach notifications have been made to authorities in the two years since Europe's [General Data Protection Regulation](#) (GDPR) new digital privacy regulation took effect in May 2018. Since, the GDPR became law, the number of breaches and other security incidents rose approximately 12%, according to research by law firm DLA Piper. [European authorities now receive 278 breach notifications daily.](#)

All of the aforementioned issues have a direct negative or positive effect on reliability. And in the Digital age of “always on” and interconnected systems and networks, the business and financial consequences are felt immediately. These facts underscore the continuing need for the highest levels of reliability. This starts with the foundational elements of core server hardware and server operating system infrastructure.

Thirty years ago, in 1990, two nines or 99% uptime, equaling nearly 88 hours of per server downtime, was considered the “gold standard” and quite acceptable. In 2020, “two nines” - 99% or even “three nines,” - 99.9% - equaling 8.76 hours of per server/per annum downtime is unthinkable.

¹ Statista, “Annual number of data breaches and exposed records in the United States from 2005 to 2019 (in millions).” URL: <https://www.statista.com/statistics/273550/data-breaches-recorded-in-the-united-states-by-number-of-breaches-and-records-exposed/>

² FBI, Public Service Announcement, September 10, 2019, Alert No. I-091019-PSA “Business Email Compromise, the \$26 Billion Scam.” URL: <https://www.ic3.gov/media/2019/190910.aspx>

Metrics of three, four and five nines – 99.9%, 99.99% and 99.999%, – equate to 8.76 hours; 4.38 hours, 52.56 minutes and 5.26 minutes of per server/per annum downtime, respectively. Technology advances in virtualization, cloud computing and Internet of Things (IoT) ecosystems, built on disparate interconnected devices, now demand near-flawless, uninterrupted reliability and availability.

Reliability versus Availability

Reliability and availability frequently are used interchangeably but they are two distinct metrics. Reliability tracks the overarching dependability of the server hardware and server OS and the amount or percentage of time it performs the intended function (e.g. line of business application, Web server). Reliability metrics also measure the mean time between failures (MTBF).

Availability measures the percentage of time the server hardware and server OS are operable and the data and services are accessible. Availability tracks the percentage of time that the organization’s crucial business processes are available in the event of unplanned downtime due to hardware failures, network or application issues or manmade or natural disasters. In other words what percentage of the time will your data and services remain available in the face of an A business that achieves 99.99% or “four nines” of uptime, which is considered the current minimum acceptable reliability standard, will still record 52.56 minutes of server downtime. Availability also takes into consideration the duration of an outage. A company might experience a Tier 2 outage lasting 15 minutes to half an hour or an hour. This can still cause significant productivity disruptions, data losses and more expensive monetary losses than several Tier 1 outages of shorter duration. It’s equally important that organizations and IT departments monitor **both** the reliability (number of incidents and amount of overall server downtime) and the availability (length of outage and amount of time applications were unavailable/offline) to accurately determine the productivity and monetary cost to the company.

ITIC 2020 Reliability Results

The evidence is in the results.

To reiterate: the IBM System Z server (mainframe) is in a class of its own. It has maintained its best in class rating among all server platforms since 2008 when ITIC first began conducting its annual Global Server Hardware Server OS Reliability poll. An overall 83% majority of IBM (both Z and Power Systems) respondent organizations said their firms achieved five and six nines – 99.999% and 99.9999% - or greater uptime. The IBM Z platform registered the highest reliability, uptime and availability numbers: 94% of enterprise organizations said the mainframe consistently achieved the vaunted six and seven nines of reliability. IBM Z customers reported

that it recorded just 0.620 seconds of unplanned per server downtime each month and 7.44 seconds annually due to any bugs in the server hardware or components. Additionally, less than one-half of one percent, of IBM Z respondents said the mainframe experienced unplanned outages exceeding four (4) hours of annual downtime.

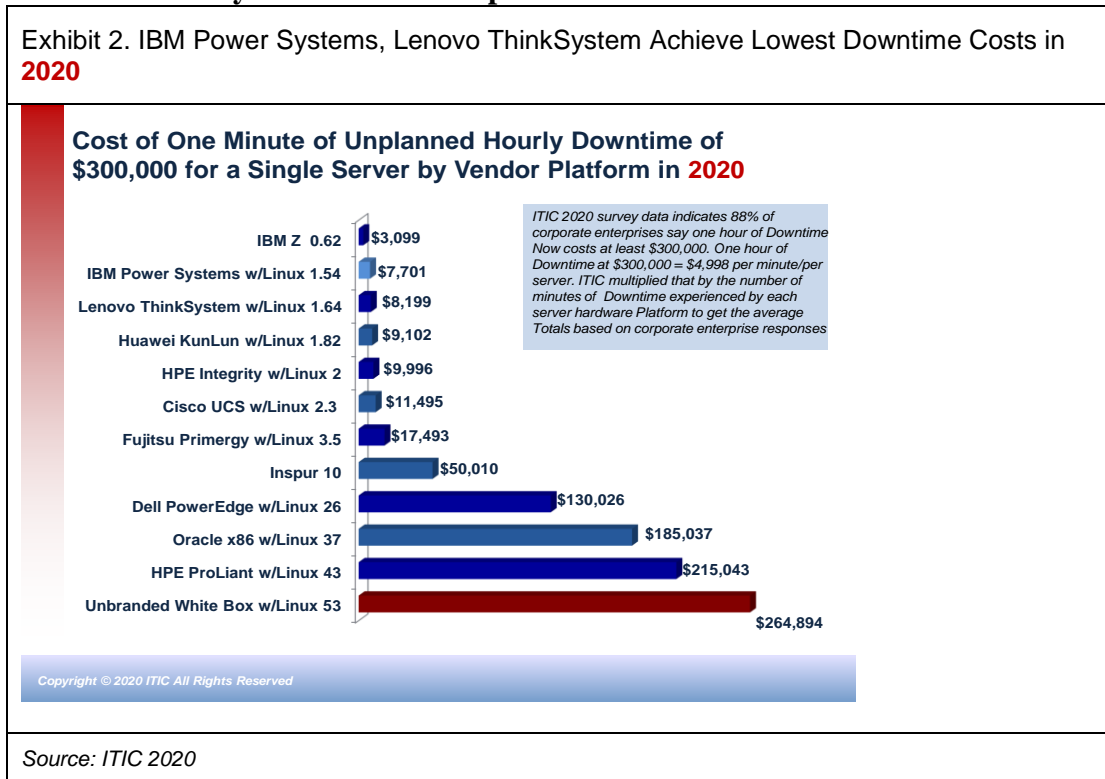
The economic annual downtime cost comparisons among the top performing and the least reliable server hardware platforms is staggering.

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

As Exhibit 2 illustrates, the IBM Z, the IBM Power Systems and Lenovo ThinkSystem distributions, continue to be the most reliable and the most economical mainstream servers.

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.

Exhibit 2. Hourly Downtime Costs per Server/Per Minute



Source: ITIC 2020 Global Server Hardware, Server OS Reliability Survey

IBM Z provides a true fault tolerant server solution that delivers the best economies of scale among all server platforms. A corporation that deploys an IBM Z and estimates one hour of

downtime cost \$300,000 would spend just \$3,099 per server/per minute, based on the platform's miniscule 0.620 seconds of unplanned monthly downtime.

IBM Power Systems, Lenovo ThinkSystem, Huawei KunLun and HPE Integrity Superdome also deliver very high and continuous availability of four and five nines.

Unplanned downtime related to hardware failures or problems cost IBM Power Systems customers who experienced 1.54 minutes an estimated cost of \$7,701 per server/per minute for a 60 minute outage calculated at \$300,000. Lenovo ThinkSystem users, whose servers experienced 1.64 minutes of per server/per minute downtime, for an outage estimated at \$300,000 would see costs of \$8,199.

Huawei KunLun servers with just fewer than two (2) minutes of per server/per minute downtime would see costs of \$9,102 and HPE Integrity at two minutes of unplanned annual per server downtime would average costs of \$9,996 due to inherent flaws in server hardware or component parts (**See Table 2**).

Corporations using Dell PowerEdge servers which experienced 26 minutes of per server/per minute downtime at the same \$300,000 per hourly downtime rate potentially would rack up outage costs of \$130,026 for a single server.

Meanwhile, businesses deploying the least reliable unbranded White box servers that registered 53 minutes of per server, per minute downtime in the latest ITIC 2020 Global Reliability survey can expect to incur downtime losses of \$264,894 specifically related to server hardware flaws and bugs in the OS and applications. The four additional minutes of downtime – from 49 minutes per server in 2019 to 53 minutes of per server outage time due to hardware flaws in 2020 represents a cost increase of \$19,992 compared with the White box server 2019 per server, per minute downtime price tag of \$244,902.

The wide ranging reliability disparities among server platforms are a compelling example of how the least expensive servers can actually cost companies more than higher end distributions which are specifically architected to support reliability with more robust capabilities.

Time is money.

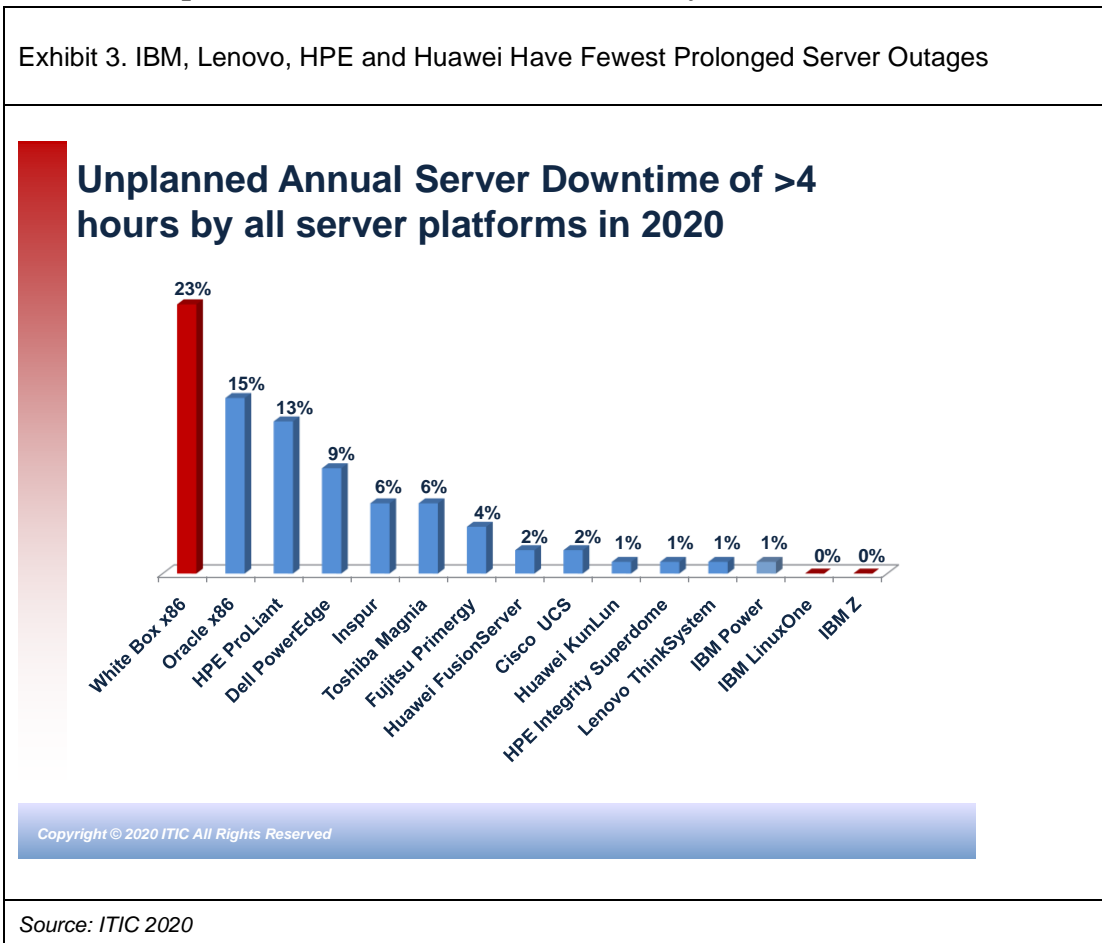
The higher monetary costs associated with unbranded White box servers are not surprising. The unbranded White box servers frequently incorporate inexpensive components. And some businesses recklessly run unsupported or pirated versions of operating systems and applications. The aforementioned hourly downtime examples are for **just one server**. Downtime costs can mount quickly and reach into the millions for corporations with dozens or hundreds of highly unreliable servers.

Keep in mind that organizations may still rack up unplanned server downtime expenses due to other external issues that may be entirely out of their control. These include: targeted security

hacks and data breaches; catastrophic weather events; catastrophic manmade events (e.g. terror attacks); human error (misconfiguration or provisioning issues).

As **Exhibit 3** illustrates, ITIC’s 2020 Global Server Hardware, Server OS Reliability study found the IBM Z achieved the best scores among all server platforms: 0% percent or none of IBM Z shops, experienced over four (4) hours of unplanned annual downtime due to problems with the server hardware or component parts. And only a niche one percent of IBM Power Systems, Lenovo ThinkSystem, HPE Integrity and Huawei KunLun servers registered over four (4) hours of unplanned annual downtime due to server flaws. This very level of high reliability also ensures excellent availability to servers, applications, and data and system resources.

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



Source: ITIC 2020 Global Server Hardware Server OS Reliability Survey

Organizations whose server hardware, operating system 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 firms will experience such negative effects as:

©Copyright 2020 **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.

- Productivity disruptions
- Inability to access mission critical applications/data
- Potential for lost, stolen, damaged, destroyed or changed 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

Connected devices, applications and people have increased by orders of magnitude. Consequently, *the potential* for collateral damage increases commensurately the minute servers and applications are offline and data is inaccessible. An outage on a virtual server running multiple instances of a crucial main line of business (LOB) application will have a greater impact on productivity, operations and revenue 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, manpower and expense involved to restore systems and services to full operational status and recovery of damaged or lost data.

Table 1 below depicts the availability percentages and the equivalent number of annual, monthly and weekly hours and minutes of per server/per annum downtime. ITIC publishes this table every year. It is a useful reference that lets organizations calculate downtime and measure the business and monetary impact on operations.

Table 2 illustrates the monetary costs associated with specific hourly downtimes ranging from \$100,000 to \$10 million based on per minute costs from a single server to one thousand servers. Taken together, the two tables paint a clear, concise and compelling overview of how quickly downtime costs can escalate to the detriment of a corporation's productivity and bottom line.

Table 1 below illustrates the wide discrepancies among the various “nines” of reliability.

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 2020 Global Server Hardware, Server OS Reliability Survey

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

- **Tier 1:** These are typically minor common, annoying occurrences. Network administrators can usually resolve such incidents in one (1) to less 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:** These are moderate-to-serious issues in which the server, operating system 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 the services, applications and services are unavailable to the corporation’s associated dependent users; the IT Department as well as external customers, partners and suppliers. 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. Examples include a backhoe severing a power line; or a natural disaster, such as a hurricane, flood or tornado, security breaches and integration issues.

Reliability Dollars and Sense: The Actual Cost of Downtime

To reiterate, 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, operating system 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 at some point in time.

However, the IBM Z, the IBM Power Systems, Lenovo x86 ThinkSystem have delivered demonstrably superior reliability and availability and security over the past 12 years (IBM) and seven years (Lenovo), respectively than commodity-based x86 platforms that do not incorporate the same level of advanced functionality. Over the past several years, the same is true for higher end, mission critical systems from HPE, Huawei, Fujitsu and more recently, Cisco Systems UCS line of servers have exhibited much improved reliability.

Table 2 below illustrates the Hourly Cost of Downtime ranging from \$100,000 to \$10 million per hour for a single server, 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
\$10,000	\$167	\$1,670	\$16,700	\$167,000
\$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 2020 Global Server Hardware, Server OS Reliability Survey

A price tag of \$100,000 for one hour of downtime for a single server appears relatively inexpensive. However, that equates to \$1,670 per minute/per server. The cost of a more severe outage estimated at \$1 million is the equivalent of \$16,700 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 2020 Global Server Hardware and Server OS Reliability Survey found that 88% of respondents calculate one hour of downtime costs the firm \$301,000 or more. Of that number, 40% of those polled indicated that hourly downtime costs now exceed \$1 million. Overall, only two percent 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. And businesses must also consider whether or not the affected hardware was a virtual server running two, three or four instances in a single machine.

©Copyright 2020 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.

The above **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 server/per minute!

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

Analysis

Reliability and uptime are in constant flux based on the performance, reliability and uptime of the servers and the main LOB applications that run on them. Technology advances and business drivers evolve. Any vendor or corporate enterprise may experience declines or increases in downtime related to bugs or flaws in its respective server platforms, operating systems and applications, year over year.

Corporate reliability and availability requirements are by necessity, markedly different in 2020 than they were 10, five or even two years ago. Business is conducted 24 x 7, 365 days a year irrespective of time zones. Corporations frequently allow their 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 operating systems, 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 business’ bottom line; strengthens security and compliance and mitigates risk.

The IBM Z followed by the IBM Power Systems and Lenovo ThinkSystem platforms have consistently delivered the highest levels of uptime availability based on every metric and measure of reliability over the past 12 years and seven years, respectively.

How have the IBM (since 2008) and Lenovo (since 2014) servers achieved and maintained best in class reliability status for over a decade and half a decade, respectively?

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

Both IBM and Lenovo have cogent, consistent strategies for high reliability, availability, recovery and all-important security in the Digital Age. And both vendors continue to successfully execute against them.

The IBM Z platform has historically been purposefully engineered for zero downtime since its initial design in 1952. The IBM Z servers, for example, are built with spare components capable of hot failovers to redundant standby processors in order to ensure continuous operations. And in a very noteworthy achievement, the IBM servers have been engineered for backwards application and system software compatibility for the last six decades.

Furthermore, both companies have invested heavily in security technology for servers and in Lenovo's case, its workstation and server platforms. The ongoing attention and focus on security has yielded tangible results.

In the latest 2020 ITIC Global Server Hardware Server OS Reliability poll, the IBM Z (z14 and z15) reduced monthly downtime to 0.62 seconds per server from 0.74 seconds per server in the 2019 survey. This improved the TCO of the z14 and z15 offerings by \$599 per server/per minute in a one hour outage where downtime is estimated at \$300,000. In the 2019 ITIC study, the same one minute of z14 or z15 per server/per minute downtime cost \$3,698 compared with the price of \$3,099 in the current survey.

Similarly, in the latest survey IBM Power Systems and Lenovo ThinkSystem servers were able to shave off approximately 20 seconds each, off their 2019 per server downtime results, respectively. This lowered the cost of hourly downtime by over one thousand dollars for both IBM Power Systems and Lenovo ThinkSystem servers.

The cost of one minute of per server/per minute downtime for an IBM Power Systems server declined from \$8,747 (based on 1.75 minutes of per server downtime) in the ITIC 2019 Global Server Hardware, Server OS Reliability Survey to \$7,701 (based on 1.54 minutes of per server downtime) in the 2020 poll, a cost savings of \$1,046 per server/per minute in an hourly outage costing \$300,000.

Lenovo ThinkSystem x86 customers also reaped hefty year over year (YoY) savings of \$1,197. Lenovo ThinkSystem server outage costs decreased from \$9,396 (based on 1.88 minutes of per

server/per minute downtime) in 2019 to \$8,199 (based on 1.64 minutes of per server/per minute downtime) in the 2020 study.

These reliability improvements are the result of IBM and Lenovo's continuing commitment to 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. However in 2019, IBM did push back the release of its POWER10 systems until sometime in 2021.

IBM's POWER9 servers 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

Additionally IBM and Lenovo's respective server strategies continually accentuate security for both the server hardware and associated operating systems. Security constitutes one of the most daunting threats to server and operating system stability and reliability.

The IBM Z and LinuxONE for example, provide pervasive encryption, [Data Privacy Passports](#) and workload isolation through partitioning; container isolation through Hyper Protect; Virtual servers and granular, virtual machine isolation through Secure Execution, to deliver constant/continuous protection against internal and external threats.

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

IBM and Lenovo also managed to avoid the management pitfalls that plagued rival HPE (formerly HP) prior to that vendor 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 multi-million 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 functions. Following the split, HP Enterprise then split again; it [spun off its IT services business](#) in order to save an estimated \$1 billion in operating costs. HP originally acquired the IT services business in 2008 as a result of its \$14 billion acquisition of EDS Corp. HPE subsequently merged the former EDS services unit with Computer Sciences Corp (CSC) and created a new corporation called SpinCo. HPE took a 50% stake or half-ownership of SpinCo.

ITIC interviewed dozens of HP and HPE IT managers over the past several years that were unhappy at being forced to cope with diminished technical support. IT departments also reported longer lead times to receive HPE replacement parts and long wait times for technical documentation and phone support. This in turn, caused a decline in the reliability of the commodity-priced HPE servers. The ProLiant line was particularly hard hit: it's landed near the bottom of the reliability scores in the last five ITIC Global Server Hardware Server OS Reliability surveys. As **Exhibit 1** depicts, the HPE ProLiant servers averaged 43 minutes of per server/per annum downtime in the latest ITIC 2020 Reliability poll. Only unbranded White box servers fared worse with 53 minutes of unanticipated downtime due to server problems or failures. On the plus side, HPE's services and high end systems have rebounded considerably and are once again scoring top marks for reliability and HPE is once again innovating.

IBM technical support continues to be robust and timely. ITIC anecdotal customer interviews found that IT managers at IBM corporate enterprises were pleased with both the quality and speed with which IBM technical support responded to issues.

Lenovo too, also delivers excellent service and support, according to IT managers interviewed by ITIC for this Report. The IT managers gave Lenovo high marks for technical support and responsiveness following IBM's 2014 sale of its Intel Corp. x 86 server businesses. And with good reason: when IBM and Lenovo inked the deal, the two vendors stipulated that IBM would continue to service the Lenovo installed base through 2019. Since then Lenovo's after-market support has continued to be extremely robust.

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 the Lenovo ThinkSystem. The HPE ProLiant platform however, is still mired near the bottom of the server offerings in terms of uptime, due to protracted user upgrade cycles.

Besides the obvious technical merits of the IBM Z, Power Systems and Lenovo ThinkSystem hardware, the server reliability is reinforced by the greater expertise of the IT administrators in corporations that use IBM and Lenovo equipment.

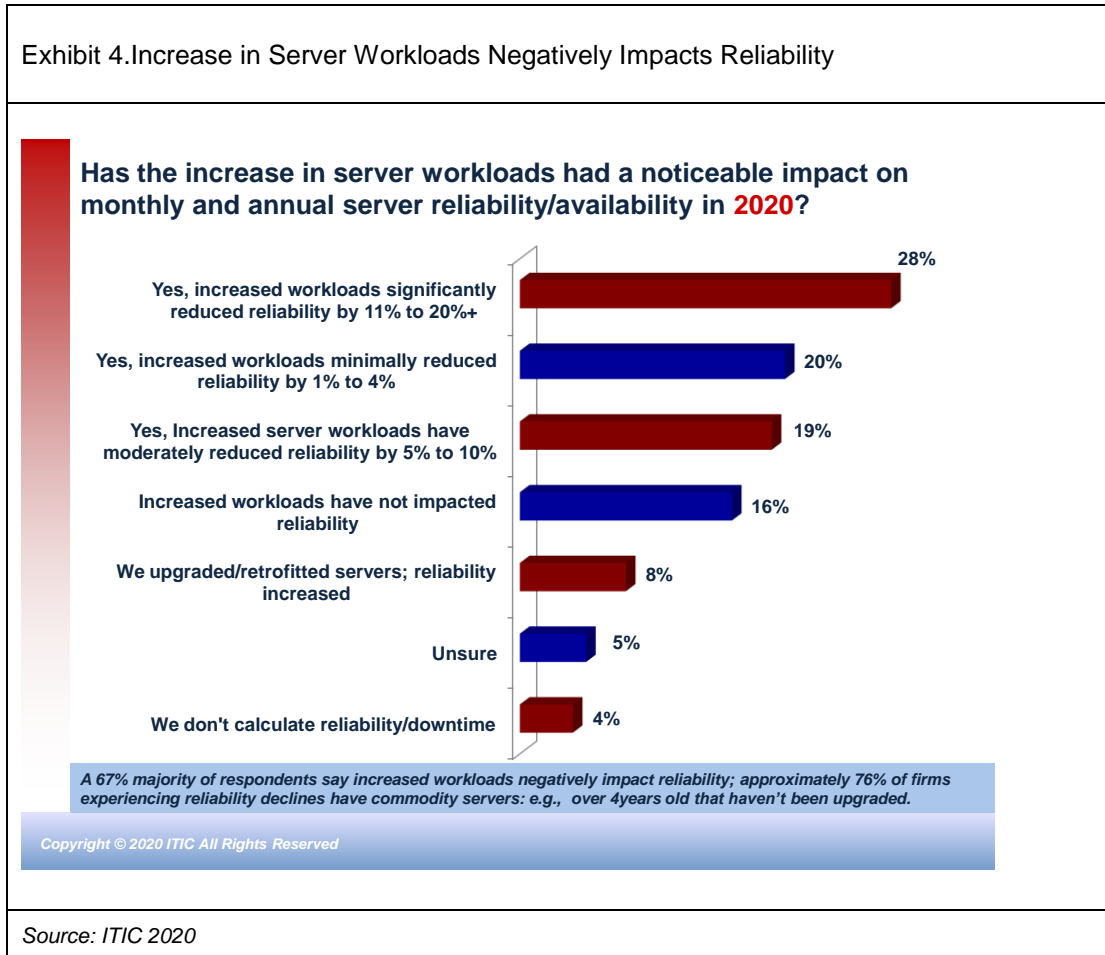
ITIC's reliability survey data shows that IBM and Lenovo IT administrators typically have 10 or more years experience. By contrast, firms (with some exceptions) that use Dell, HPE ProLiant, Oracle and unbranded White box server hardware are more likely to hire IT managers with one-to-five year's experience.

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

Another differentiator: a higher percentage of IBM and Lenovo organizations also adhere to a regular three year upgrade cycle and retrofit their IBM, Lenovo, Huawei and HPE Integrity servers as needed. This is crucial since applications like AI, Analytics, Blockchain, IoT and Virtual Reality (VR) are compute and resource-intensive.

Exhibit 4 shows a two-thirds majority of corporate respondents find larger workloads negatively impact reliability; up six percentage points since 2019. Only 16% of those polled indicated that server reliability was not negatively affected by higher workloads.

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



Source: ITIC 2020 Global Server Hardware, Server OS Reliability Survey

Among the 67% of users whose server reliability dropped, approximately 76% have commodity servers: e.g., White box; older Dell, HPE ProLiant & Oracle hardware over four (4) years old that haven't been retrofitted or upgraded.

Commodity server users should not defer upgrades or retain servers well beyond the recommended three year upgrade cycle. Over 60% of Dell PowerEdge and unbranded White box

users retain the servers for four, five or even six years while increasing the application workload. And this has been the case for the last four years. This is just asking for trouble. The exceptions to this rule: very small businesses whose application environment remains static.

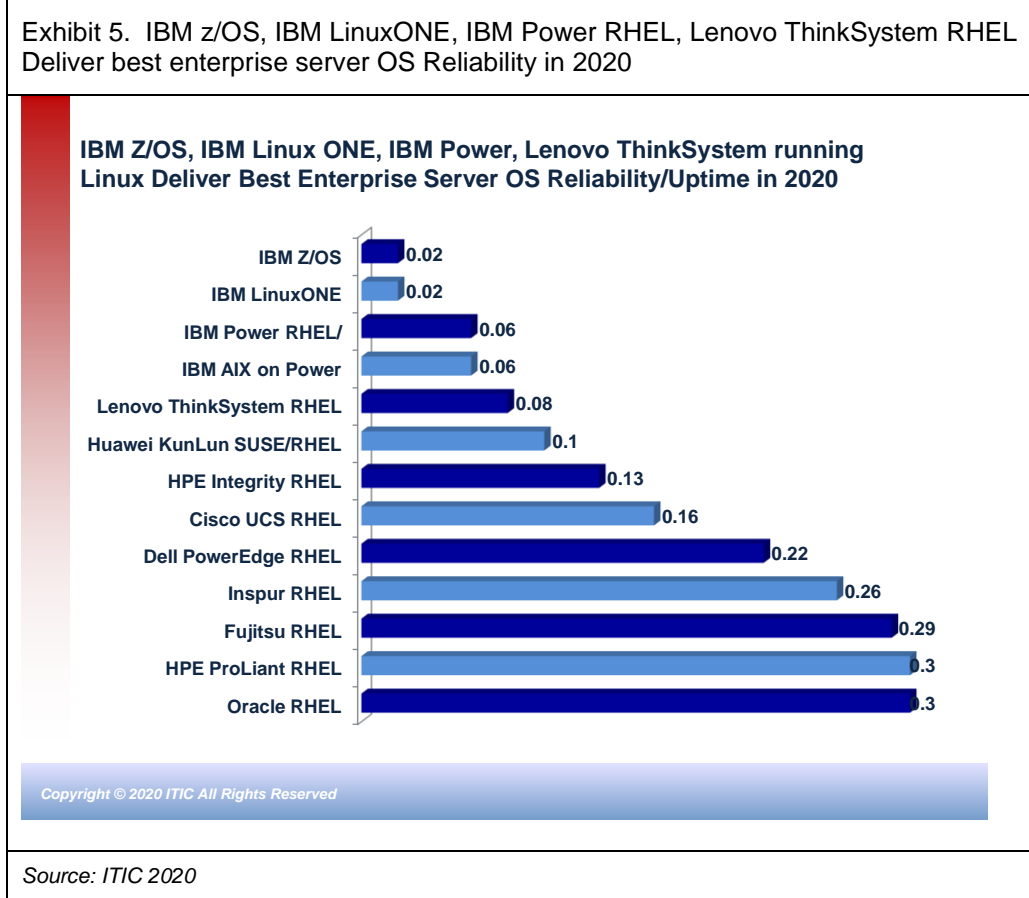
Other Notable Survey Findings

Among the other key survey highlights:

- **Unplanned Server Outages Due to Bugs in the Operating System:** In a continuation of a trend that ITIC saw in the prior three (3) Reliability surveys and 90% of IBM Power Systems and 89% of Lenovo ThinkSystem users running RHEL, SUSE or Ubuntu Linux experience fewer than one *unplanned* outage per server/per year due to any flaws/bugs in the operating system.
- **IBM Z** ranked “best in class” for reliability, accessibility, performance, and security among all server platforms. The IBM Z servers had highest reliability/uptime ratings across the board in terms of actual minutes of unplanned per server/per annum downtime. A 90% majority of IBM Z users achieved at least six nines or 99.9999% reliability, with some reporting seven nines – 99.99999% of uptime. Specifically, IBM Z servers exhibit true fault tolerance experiencing just 0.62 - less than one minute of *unplanned* per server, per annum annual downtime due to inherent problems with the server or its component parts. Recently introduced, System Recovery Boost is also available to minimize the impact of downtime by providing a boost to image processing capacity after an IPL
- **Cisco Systems’ Unified Computing System (UCS) servers** – which are frequently deployed at the network edge – showed a marked improvement in reliability beginning with ITIC’s 2019 Global Server Hardware, Server OS Mid-Year Update survey. Cisco servers have maintained that high level of reliability and uptime in the latest poll. The Cisco UCS servers reduced per server/per annum downtime by nearly 50% from the 4.1 minutes in ITIC’s prior first quarter reliability survey to 2.3 minutes in the 2020 survey.
- **Dell PowerEdge server reliability remains stable:** Dell’s PowerEdge – in particular, new to two-year old models - showed a marked uptick in reliability. Overall, Dell PowerEdge servers recorded 26 minutes of unplanned per server/per annum downtime due to inherent system and component flaws in the latest 2020 Reliability poll. Both the survey data and anecdotal evidence based on first person interviews indicate that the reliability and availability of the Dell PowerEdge servers are contingent upon the server’s age, configuration and the experience of the IT administrators.

As **Exhibit 5** details, the IBM Z server running z/OS and the IBM LinuxONE offered best in class uptime and availability owing to a very bug-free server operating system environment. IBM’s Power Systems servers running IBM’s Red Hat Linux (RHEL) operating system also made a very strong showing along with Lenovo’s x 86 ThinkSystem servers running RHEL and Huawei’s KunLun running RHEL and the SUSE operating system.

Exhibit 5. Enterprise Server OS Reliability/Uptime



Source: ITIC 2020 Global Server Hardware, Server OS Reliability Survey

Server Hardware Platform Overview

IBM Z

The IBM mainframe made its debut in 1952 – some 68 years ago. Its popularity soared in 1964 with the introduction of the IBM System 360.

©Copyright 2020 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.

For at least the last quarter century industry pundits have predicted its demise. In fact, the IBM Z server is enjoying a renaissance. That's attributable to a number of factors. In the late 1990s, IBM made a series of crucial – and at the time, controversial decisions, one of which was to open up the platform to the then, fledgling Linux and open source OS distributions. It also debuted the “zSeries” mainframes; a variety of new form factors; and the more economically priced midrange z13s and z14 ZR1 mainframes, IBM followed up these offerings with the LinuxONE server, which was provisioned entirely for Linux workloads. The LinuxONE, in combination with robust secure service containers, is a key component of IBM's Blockchain offerings, IBM Cloud's Blockchain services and IBM private cloud implementations.

IBM z15 family debuted in September 2019, while the smaller mid-range z15 offerings came out in spring 2020. The three key tenets of the z15 platform offerings are:

- **Encryption everywhere:** which builds upon IBM's existing pervasive encryption, model which is: secure by design, secure in use, secure in transmission and secure at rest (in storage.) IBM's z15 has now advanced security with [Data Privacy Passports](#) technology that enables businesses to control how data privacy is maintained when data is stored and shared. They now have the ability to protect and provision data and revoke access to that data at any time, not only within the z15 environment but across the enterprise. The z15 can also encrypt data everywhere including across hybrid multicloud environments. This ensures that data remains secure regardless of where it is transmitted.
- **Secure Execution:** IBM z15 and LinuxONE also now provide Secure Execution, a Trusted Execution Environment (TEE) for enhanced isolation and access control of assets which can perform at the enterprise scale required for mission critical applications.
- **Cloud-Native Development:** Can give clients a competitive advantage by evolving how they modernize apps in place, build new cloud-native apps and securely integrate their most important workloads across clouds. Clients are already using IBM Z for their mission-critical workloads to build, deploy and manage next-gen apps and protect data through advanced security.
- **System Recovery Boost:** This provides exceptional resiliency designed to minimize the cost and impact of both planned and unplanned downtime, so that organizations realize “instant recovery” from an outage. It allows enterprises to access additional system capacity for a period of time to accelerate shutdown and restart of IBM z15 and LinuxONE III images and provide a temporary capacity boost to rapidly recover service levels and process backlog. The so-called instant recovery capability with IBM z15 unlocks additional processing capacity during shutdown, restart and workload catch-up to help clients mitigate the impact of downtime to get back up and running to previous service level agreements (SLAs) 50% faster than previously possible, all without any increase in IBM software cost. This benefit comes from both the architecture based on the way the z15 is designed and delivered with compute headroom built in to the offering structure, along with IBM having an integrated software stack to provide a total solution

to clients. There are also benefits in improved parallelism in GDPS processing to help accelerate recovery.

To reiterate, the IBM Z offerings are in a class of their own in terms of advanced features/functions, reliability and availability and security. The platform delivers near-flawless reliability of 0.620 seconds per month and just 7.44 seconds per server of unplanned downtime due to server hardware problems. IBM's Z portfolio of servers focus on: encryption, cyber resiliency, cloud computing and real time analytics. The IBM Z boasts virtually no downtime as a result of inherent system or component flaws. That fact alone makes it worth its weight in gold. IBM continues to advance the z15 capabilities with more power in smaller footprints via the mid-range offerings.

The native cloud application development advancements; the instant recovery capabilities for **both** planned and unplanned downtime events and data security and privacy advancements, tick all the boxes that long-time and prospective IBM Z enterprise customers want to further performance, reliability and security across their entire ecosystem from datacenter to private, hybrid and public cloud to the network edge.

IBM Power Systems

IBM's POWER9 scale-out systems began shipping in March 2018 and they offer significant performance, security and reliability improvements over the previous POWER8 servers. Specifically, the POWER9 line delivers up to 1.5x speed and the performance of its POWER8 predecessor to support increasingly demanding compute-intensive workloads and applications.

Equally crucial are POWER9 servers' security capabilities. Just as it does with the Z, IBM builds security into every layer of the POWER9 stack – starting at the chip level. For customers, this means that POWER9 servers offer end-to-end security: they are secure by design, secure in use, secure in transmission and secure at rest (in storage). Security is an absolutely essential element to ensure reliability/uptime, availability and continuous data operations. POWER9 servers have double the number of crypto engines; they encrypt data twice as fast as the prior POWER8 servers. And the on-chip accelerators with new GZIP compression/decompression similarly perform much faster than software.

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. The 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 as well as SAP HANA. The IBM POWER9 servers incorporate the latest I/O technology; this includes 25 GB/sec high-speed interconnect for 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.

Additionally, IBM says compared to the prior generation POWER8 systems, POWER9 servers deliver 1.25-1.5X 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 2X the per-core performance over compared x86 systems across a wide range of benchmarks. IBM says that the improvements to the POWER core capabilities also results in lower cost of acquisition and lower cost of ownership over the entire server product lifecycle. Overall, these advancements all undoubtedly attributed to POWER9 servers' continuing reliability improvements. To reiterate: ITIC's 2020 Global Server Hardware Server OS Reliability survey statistics reveal the POWER platform recorded just 1.54 minutes of unplanned per server downtime due to any inherent hardware flaws, a decrease of 21 seconds from the 1.75 minutes it posted in the 2019 Reliability poll. This extra 21 seconds of reliability represents significant monetary cost savings depending on the length, severity and timing of a corporate enterprise's specific outage.

The POWER9 architecture is open for licensing and modification by [OpenPOWER Foundation](#) members. The next version – POWER10, which will be based on 7nm technology, is due out sometime in 2021. IBM aims to deliver a very high core count and high performance I/O.

Lenovo ThinkSystem

Lenovo's Data Center Infrastructure model offers customers edge to core to cloud solutions based on the latest technologies. It delivers leading edge performance, lower TCO, enhanced security and best in class reliability. ThinkSystem servers now average 1.64 minutes of unplanned per server/per annum downtime, in 11% less time ³than the next closest x86 competitors. Additionally, in ITIC's 2020 Global Reliability survey, the Lenovo ThinkSystem platform improved its reliability by 15% over its' previous year results. ⁴Lenovo's repeated x86

³Unplanned Downtime by Server Hardware Platform in Minutes Per Server/Per Annum: Lenovo ThinkSystem w/Linux = 1.64 minutes vs. Huawei KunLun and FusionServer w/Linux = 1.82 minutes

⁴Unplanned Downtime by Server Hardware Platform in Minutes Per Server/Per Annum: 2020 Lenovo ThinkSystem w/Linux = 1.64 vs. 2019 Lenovo ThinkSystem w/Linux =

leadership in reliability highlights its commitment to providing customer solutions that are focused on quality through a meticulous attention on the details within the product design.

To handle customers' evolving data-intensive workloads, Lenovo introduced their first AMD servers, ThinkSystem SR635 and ThinkSystem SR655, featuring the new AMD EPYC 7002 family of processors. Lenovo's ThinkSystem SR635 and SR655 single-socket server platforms harness the full capabilities of the AMD EPYC™ 7002 Series processor family. These systems are optimized to handle evolving, data-intensive workloads such as video security, software-defined storage and network intelligence, as well as support for virtualized and edge environments. They provide more throughput, lower latency; higher core density, enhanced security, and the largest NVMe drive capacity of any one-socket server on the market. The same focus around performance, security and reliability that have made the ThinkSystem brand of products one that customers have come to trust in their datacenters were leveraged in the development of these new AMD offerings.

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

Lenovo customers interviewed by ITIC for both the latest 2020 Reliability poll and the 2019 Mid-Year Update survey praised Lenovo's "agile system design" and hot swap capabilities. The Lenovo customers were also very encouraged by the vendor's well documented investment in the ThinkShield embedded security and management 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. This is a win-win for Lenovo customers and Lenovo.

1.88; 15% decrease in unplanned downtime (<https://en.resources.lenovo.com/analyst-reports/itic-2019-global-server-hardware-server-os-reliability-report>)

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's KunLun distributions for least amount of unplanned per server downtime averaging two minutes. And they benefit from the renewed stability of the HPE brand and business five years after Hewlett-Packard split into two separate entities. This has enabled HPE to once again renew its focus on innovation and after-market technical service and support which is crucial in the current complex, digital computing environment. HPE is a legacy server brand that is entrenched in corporate enterprises; it has built a reputation as a solid vendor over the last six decades. Now that the corporate tumult is over, HPE is focused on innovation and reliability. This is yielding tangible results as evidenced by the improved reliability uptime of the HPE Integrity Superdome servers.

Huawei KunLun and Fusion Servers

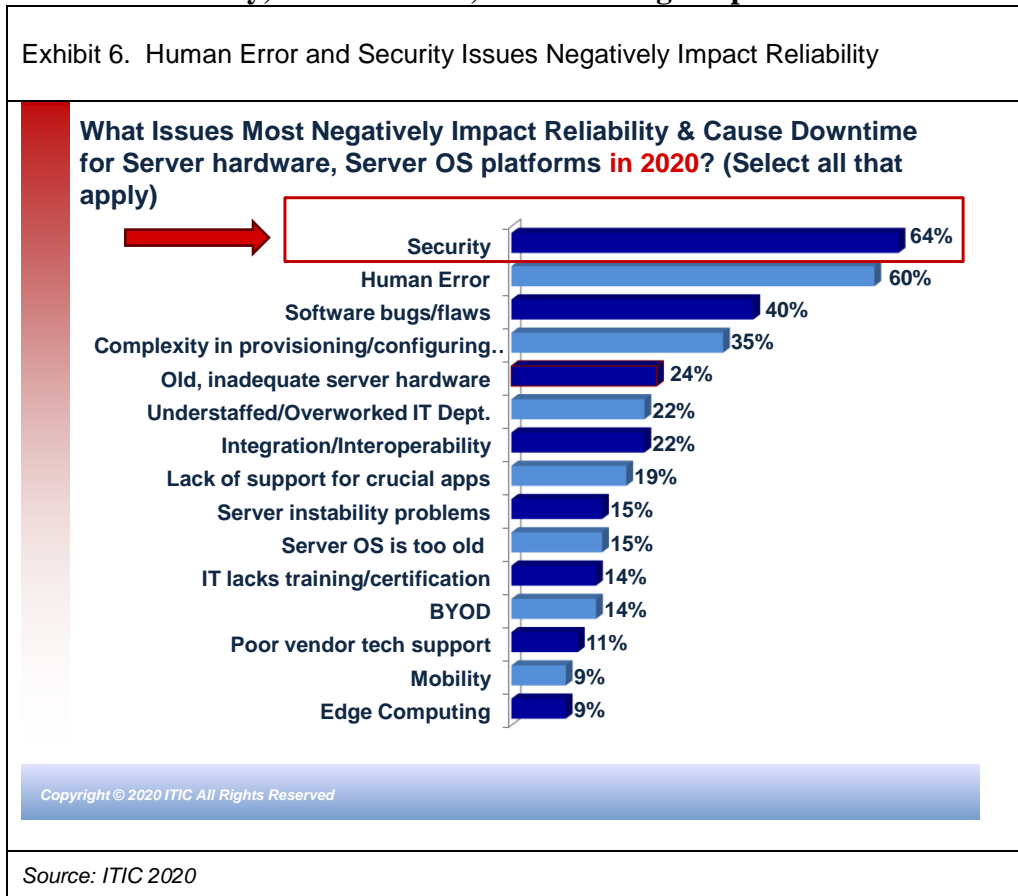
Huawei is among the newest entrants to challenge established players in the high end server market with its three year old KunLun family of 16- and 32-socket mission critical servers. Huawei is now among the Top Five vendors in terms of global server hardware shipments. In the last three years Huawei, headquartered in Shenzhen, China, has emerged as one of the top five server hardware vendors worldwide with its high end KunLun mission critical server and its general purpose FusionServer x 86-based servers. Huawei's ascension in the server market –with the exception of North America - is due to a multi-pronged series of bold tactical moves and its ability to formulate and execute a cogent, compelling strategy to ensure long-term success. To successfully compete with rivals Fujitsu, HPE, IBM, Lenovo and others, Huawei's servers support a wide range of customer needs for everything from general purpose rack and blade servers to mission critical hardware to address high performance computing (HPC). Huawei has also imbued its servers with advanced capabilities to support emerging compute intensive applications like Artificial Intelligence (AI), Big Data Analytics, Deep Learning and Machine Learning. According to the 2019 EU Industrial R&D Investment Scoreboard report,⁵ Huawei ranked fifth in global R&D investment among 2,500 companies the European Commission surveyed. In 2019 Huawei's R&D hit \$13.1 billion; this equals 15% of the Chinese firm's annual sales. Huawei's R&D spending bested both Intel with \$12.6 billion and Apple with \$11.1 billion. Approximately one billion of Huawei's R&D was invested in server technology. In July 2019, company executives said publicly that Huawei's R&D budget for 2020 and beyond, would reach \$15 to \$20 billion annually.

⁵ The European Commission, 2019 EU Industrial R&D Investment Scoreboard. URL: <http://iri.jrc.ec.europa.eu/scoreboard18.html>

2020 Reliability Trends: Top Notch Security is Crucial

As **Exhibit 6** illustrates, 64% of respondent say security is the top cause of reliability issues. , followed by 60% of respondents who blamed Human error.

Exhibit 6. Security, Human Error, Software Bugs Top Causes of Downtime



Source: ITIC 2020 Global Server Hardware, Server OS Reliability Survey

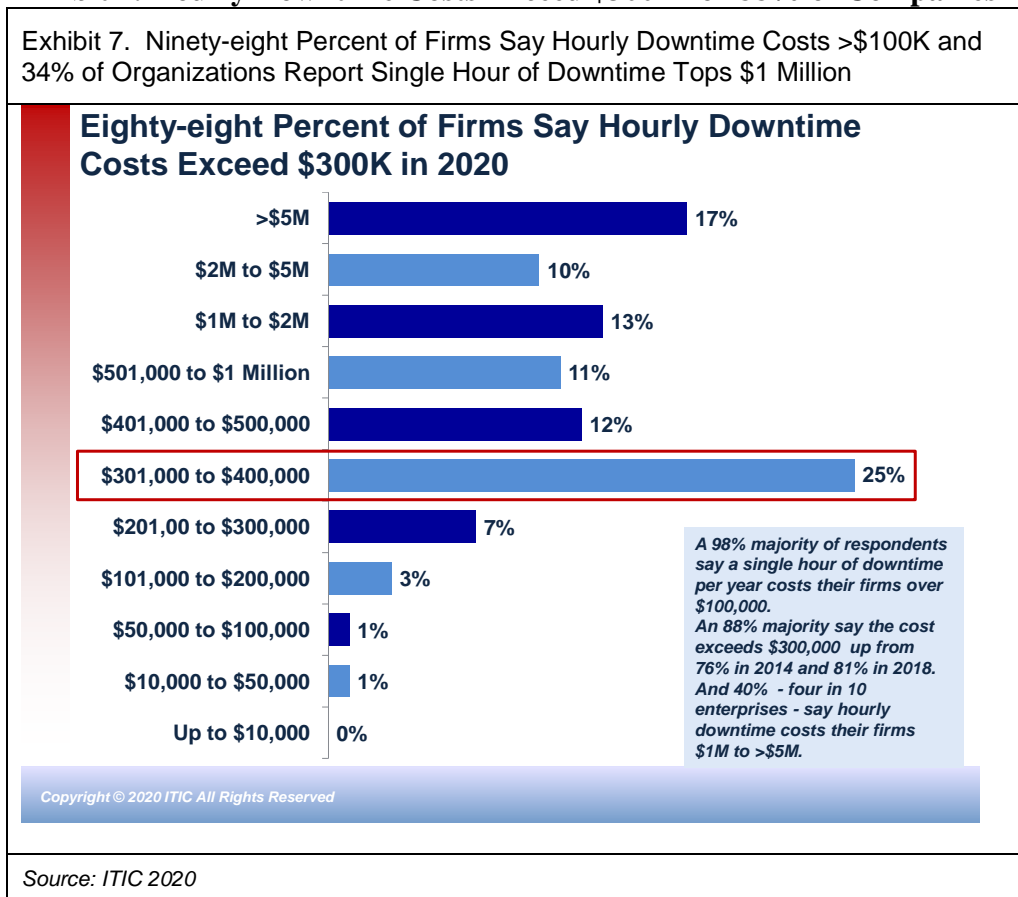
Security and Human error are inextricably interwoven. Unfortunately, ITIC expects no foreseeable end in sight for this trend as targeted security and data breach attacks like ransomware, BEC, Phishing scams, CEO fraud and even corporate espionage attacks are a reality of doing business. Targeted attacks are commonplace in the Digital age of connected systems and network ecosystems. No server, device, application or network is immune. It's crucial that businesses closely communicate with their vendors and stay up to date on news of latest patches

and fixes to find out the most effective mechanism for dealing with hardware and software-based security flaws.

Hourly Cost of Downtime Continues to Rise

As ITIC has referenced repeatedly in this report, 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 24 x 7, it's safe to say that the cost of downtime will continue to rise.

Exhibit 7. Hourly Downtime Costs Exceed \$300K for 88% of Companies



Source: ITIC 2020 Global Server Hardware, Server OS Reliability Survey

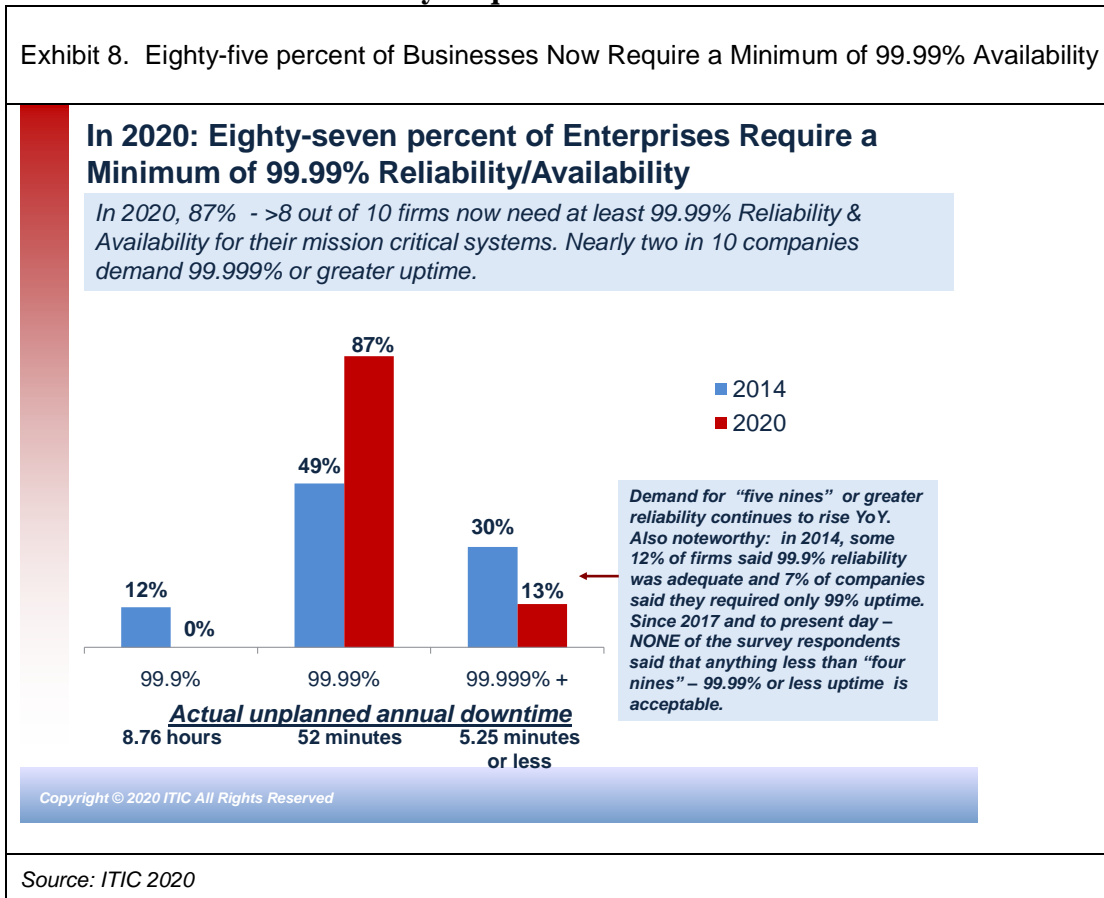
Minimum Reliability Requirements Rise

Similarly, based on all of the other trends, the latest ITIC 2020 Global Server Hardware and Server OS Reliability survey found that 87% of respondent corporations now require a minimum

“four nines” or 99.99% reliability/uptime. This is more than double the 39% of respondents that said their firms required 99.99% availability in ITIC’s 2013 poll, seven years ago.

For the last several years “four nines” of uptime has become the minimum acceptable reliability standard. It equates to 52.56 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 5.26 minutes of *unplanned* annual downtime or a scant “blink and you miss it” 25.9 seconds per month for their mission critical servers and main line of business applications (See Exhibit 8). The remaining five percent of enterprises indicated their firms need “six nines” – 99.9999% or better availability/uptime.

Exhibit 8. Minimum Reliability Requirements Rise



Source: ITIC 2020 Hourly Cost of Downtime Survey

Businesses’ demand for continuous uptime and availability will continue unabated.

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. Technologies including AI, Analytics, Blockchain, cloud computing, IoT, security, virtualization and VR, are all more compute intensive and place greater demands on the core infrastructure components than technologies dating back to the 1990s, early 2000s or even those as recent as 2015. All of them can place a greater toll and tax system resources to the breaking point *if* organizations fail to upgrade and right size their hardware, operating systems and applications to accommodate higher demands.

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.

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 2020, nearly nine-in-10 businesses estimate that a single hour of downtime costs \$300,000 or higher. That is approximately \$5,000 for a one minute outage and \$25,000 for five minutes of downtime. Those five minutes of downtime are worth the price of a fully loaded entry-level to mid-range server. And keep in mind that many – albeit not all – outages involve multiple servers.

The aforementioned examples are average estimates and will vary according to individual corporate implementations and use cases. 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.

All the statistics presented in ITIC’s 2020 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 operating system 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 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

In summary, the ITIC 2020 Global Server Hardware and Server OS Reliability Survey findings indicate that for the 12th straight year, the elite IBM Z (z14 and z15) is the clear winner outperforming all other platforms in every reliability category with over 90% of customers giving it five, six and seven nines of uptime. The IBM Z's performance was followed closely by the IBM Power Systems servers, which registered marked reliability improvements. The Lenovo ThinkSystem servers ranked as the most reliable x86-based hardware platform for the seventh straight year. Both the IBM Power and Lenovo ThinkSystem distributions achieved four and five nines of reliability - 99.99% and 99.999% - according to over nine-out of-10 corporate survey respondents. The IBM Power Systems and Lenovo ThinkSystem platforms are up to 28 xs more reliable than the worst performing unbranded "White Box" servers.

The strong reliability of the IBM Z, IBM Power Systems, Lenovo ThinkSystem, HPE Integrity and Huawei Kunlun hardware platforms yield demonstrably better uptime and availability resulting in lower TCO and accelerated ROI than the least reliable white box servers. .

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 server is in a class by itself. No other server platform (excluding super computers and HA machines) even come close. A 94% majority of IBM Z customers said their businesses achieved unparalleled fault tolerant levels of six and seven nines - 99.9999% and 99.99999% reliability and continuous 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 Z, IBM Power Systems and Lenovo ThinkSystem servers did experience an outage due to inherent problems with the server hardware or component parts – they were of short duration. It typically took IT managers under 10 minutes and in most cases, approximately five (5) minutes or less to bring them back online.
- Least amount of server downtime of over four (4) hours: zero percent of IBM Z and just one percent (1%) of IBM Power Systems, Lenovo ThinkSystem platforms, 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 operating system (OS): IBM Z was again the clear winner followed by IBM Power Systems, Lenovo 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.
- IBM Z, IBM Power Systems, Lenovo ThinkSystem, HPE Integrity and Huawei KunLun servers all delivered the least amount of *Unplanned Server Downtime* in 2020 (Hours per Year) in comparable hardware configurations and workloads.

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

The reliability of newer Dell PowerEdge servers (new to 2 ½ years old) rebounded and Dell's overall reliability scores remained unchanged. However, Dell's higher levels of downtime are the direct result of corporate end users pushing overtaxing the server resources and retaining servers for four (4) years or more without retrofitting the boxes to accommodate higher workloads.

Reliability is fluid, not static. No server, no component part – hard drive, memory or CPU; operating system; application, device or connectivity mechanism is immune from inherent problems or failure. Additionally 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 operating system 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 organizations' requirements for unassailable reliability and near flawless uninterrupted availability, is greater than ever.

The 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 87% majority of ITIC 2020 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. 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.

Businesses must regularly replace 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.

In the 21st century 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 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

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, provide the necessary guidance and top notch 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 foot path.

It is the responsibility of corporate enterprises to select the servers, operating systems 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 efficiencies. Appropriate training and certification for IT and Security administrators is also a must.

ITIC strongly advises organizations to conduct a regular and rigorous analysis that measures the uptime and reliability of their main LOB server hardware, server operating systems and applications. Identify the causes of the outage(s); determine whether any data was lost, stolen, damaged, destroyed or changed. Following the review, the IT department, C-level executives and the affected department heads should be able to ascertain the business and financial impact of the outage; how many employees were impacted and how many IT administrators and how much time they expended in remediation activities. Organizations should also verify whether or not any customers, business suppliers or partners were negatively affected by the outage.

Being cognizant of specific uptime and reliability statistics will enable the business and its IT department to identify and establish baseline metrics associated with all of their individual platforms. This will be useful for future reference. It will also provide companies with an accurate assessment of the inherent reliability and flaws in their hardware and software. Companies 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.

The ability to measure reliability also helps organizations gauge how downtime influences and possibly impinges on external business partners, customers and suppliers.

To optimize uptime and reliability, ITIC advises corporations to:

- **Be vigilant about Security.** Security is now the top cause of downtime. It will remain so for the foreseeable future. Hacks are more targeted and pervasive. Organizations of all sizes, across all vertical markets should construct a comprehensive security plan. Make sure your firm installs the latest security updates and patches. Conduct regular vulnerability testing to identify the potential weak spots and holes in your company's defenses – on premises and at the network edge. Make sure your Security and IT administrators and employees receive the proper training to enable them to recognize and thwart hacks.
- **Regularly analyze and review configurations, usage and performance levels.** This will enable companies to determine whether or not current server and server OS environment allows them to achieve optimal reliability.
- **Keep a comprehensive record of downtime and associated costs.** IT departments should compile a detailed list of outages and all pertinent remediation efforts. Include facts like the cause of the outage (e.g., hard drive failure, human error, manmade disaster etc.); the length/duration of downtime; the severity of the event (e.g. lost, damaged or stolen data; interrupted transactions). Also include the Mean Time to Detection and Mean

Time to Remediation and Recovery. All company stakeholders should compile a comprehensive list of the costs incurred by all affected departments (IT and employees) including the costs due to lost, damaged, destroyed or changed data. Companies should also keep detailed records of any litigation costs as well as civil, criminal or non-compliance penalties resulting from outages whatever the circumstances.

- **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 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.
- **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 of the outage and whether or not any 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 at peak usage time, disrupting a crucial transaction potentially can cost the business thousands or even millions.
- **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 indicate that nearly half of all respondents - 49% do not calculate the hourly cost of downtime. This is a mistake. To reiterate: maintain detailed and accurate records of outages and their causes. Classify outages according to their severity and length – e.g., Tier 1, Tier 2 and Tier 3. The appropriate IT and department 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 to restore full operations and what staff members participated in the event. This is an invaluable resource should the problem recur. It may also serve to contain and minimize reliability-related incidents.
- **Calculate the cost of unplanned and planned 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. It's also useful to log the amount of time spent on planned downtime to upgrade servers and applications and perform patch management. 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.

- **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 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 *2020 Global Server Hardware and Server OS Reliability Survey*, polled C-level executives and IT managers at over one thousand corporations worldwide from November 2019 through January 2020. 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 1,200 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 65% of respondents hailed from North America; 35% were international customers who hailed from more than 20 countries throughout Europe, Asia, Australia, New Zealand, South America and Africa.

Appendices

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/2019/11/ibm-lenovo-hpe-and-huawei-servers-maintain-top-reliability-rankings-cisco-makes-big-gains-ibm-lenovo-hardware-up-to-24x-more-reliable-28x-more-economical-vs-least-reliable-white-box-servers/>

<https://itic-corp.com/blog/2019/11/1678/>

<https://itic-corp.com/blog/2019/08/itic-poll-human-error-and-security-are-top-issues-negatively-impacting-reliability/>

<https://itic-corp.com/blog/2019/08/itic-2019-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/>