

INFORMATION TECHNOLOGY  
INTELLIGENCE CONSULTING

Information Technology Intelligence Consulting



**ITIC 2023 Global Server Hardware,  
Server OS Reliability Report**

**August/September 2023**

Table of Contents

- Executive Summary .....3
- Introduction .....6
- High Reliability and Operational Expenditure Costs .....7
- Survey Highlights.....8
- The Nines of Reliability and the Cost of Downtime .....9
- Data Analysis ..... 12
- Overview: Top Survey Findings ..... 13
- Downtime Comparison Costs by Server Platform..... 16
- Hourly Downtime Costs: 48% Estimate Losses Exceed \$1M.....23
- Security Hacks, Data Breaches are Top Cause of Downtime .....26
- Security, Resiliency Reduce Downtime Costs .....27
  - Enterprises that Prioritize Security, Best Practices Bolster Reliability.....28
- Business and Technology Trends Impacting Reliability .....29
- Conclusions ..... 31
- Recommendations ..... 33
- Survey Methodology ..... 35
- Survey Demographics .....36
- Appendices.....36

## Executive Summary

*The IBM Z and IBM Power continue to dominate, delivering the best server reliability, uptime and security for the 15<sup>th</sup> straight year.*

*Lenovo's ThinkSystem servers provide the best reliability among all x86 servers for the ninth consecutive year and best server security for the last five years.*

*Cisco Systems, Huawei KunLun, Hewlett-Packard Enterprise (HPE) Superdome mission critical servers also register high reliability and security rankings challenging the leaders. Cisco, in particular, continues to up its game with robust network edge reliability and security.*

*IBM Z and IBM Power deliver over 40x more uptime than least efficient competing platforms and up to 60x lower Total Cost of Ownership (TCO depending on age, usage, and configuration). The Lenovo ThinkSystem, Cisco UCS, Huawei KunLun and HPE Superdome delivered the best TCO among x86 platforms ranging from 17x to 40x more economical than less reliable rivals.*

*An 84% majority of corporate respondents cite security as the top cause of unplanned downtime and 67% say human error causes unplanned outages.*

---

For the fifteenth year in a row, IBM Z and Power Systems servers dominated the reliability and security landscape, delivering the highest levels of uptime and security among 18 server hardware and operating system platforms.

The IBM Z mainframe servers (z13, z14 and z15) achieved an average score of “eight nines” of uptime – or 315 milliseconds of per server annual downtime (**See Exhibit 1**). However, the latest z16 server, (introduced in April 2022), delivers nine nines—99.9999999%—of near perfect uptime and reliability. This is 31.56 milliseconds of per server annual downtime, according to the results of the ITIC 2023 Global Server Hardware, Server OS Reliability Survey. From a monetary standpoint, IBM z16 businesses spend literally next to nothing on per server, per annum operational expenditures performing IT remediation due to unplanned server downtime.

Those are the results of the ITIC 2023 Global Server Hardware, Server OS Reliability independent Web-based survey. It polled 1,900 C-level executives, IT, and security administrators in corporations across 37 vertical market segments worldwide. It compares the reliability, performance, and security of 18 leading mainstream on-premises and cloud-based servers and operating systems from March through mid-July 2023. The corporate enterprise participants were selected based on the length of usage of the various server hardware and server OS distributions. For customers' enterprise reliability responses to be counted, ITIC specified that all server hardware distributions and specific releases were deployed in production


environments for at least one year. The survey also examines the monetary and business costs associated with outages. To maintain objectivity, ITIC accepted no vendor sponsorship.

ITIC's 2023 Global Server Hardware, Server OS Reliability survey also found that overall, 91%, of IBM Power8 and Power9 customers averaged six nines (99.9999%) – the equivalent of 31.5 seconds of unplanned annual per server downtime and seven nines (99.99999%) which equals 3.15 seconds of unplanned per server yearly downtime. Separately, an 88% majority of the newest IBM Power10 server (shipping since September 2021) enterprises achieved eight nines—99.999999%—of uptime. This is 315 milliseconds of unplanned, per server, per annum outage time due to underlying system flaws or component failures. So, Power10 corporate enterprises might spend \$0.7 cents per server/per year performing remediation due to unplanned server outages.

The IBM z16 and Power10 server-specific uptime statistics were obtained by breaking out the results of more than 200 respondent organizations that deployed the z16 since it began shipping in April/May 2022. A 96% majority of z16 enterprises say their businesses achieved nine nines—99.9999999%—of reliability or 31.56 milliseconds of unplanned annual per server downtime (**See Table 1**).

The Lenovo ThinkSystem servers followed closely; 78% of customers said they posted six nines of uptime which is the highest reliability score among all x86 hardware distributions for the ninth consecutive year. Overall, a 94% majority of Lenovo ThinkSystem server enterprises attained five, six and seven nines of reliability. Cisco UCS, Huawei KunLun and Fusion servers, the HPE Superdome hardware, rounded out the top five most reliable server platforms. They averaged a robust five nines of reliability equal to 5.26 minutes of unplanned annual per server downtime.

**Exhibit 1.** Unplanned Annual per Server Downtime in Minutes by Vendor Platform



**Source:** ITIC 2023 Global Server Hardware, Server OS Reliability Survey

Time is money.

A few minutes of downtime can halt business operations and immediately impact revenues. This is particularly true if an outage occurs during peak usage times or interrupts a crucial business transaction (s). The potential consequences of an unplanned outage will have a negative, domino effect on customers, business partners and suppliers.

Reliability or lack thereof, raises the risk of litigation and non-compliance with industry and government regulations. This leaves the business open to civil and criminal fines and penalties. Downtime associated with unreliable systems can also potentially damage a company’s reputation and result in lost business.

Corporate enterprises owe it to themselves and their businesses to deploy the most reliable, robust, and secure servers and server OS software.

The IBM Z, IBM Power, Lenovo ThinkSystem, Cisco UCS, HPE Superdome and Huawei KunLun servers (in that order) are the most consistently reliable, secure and robust servers.

These six platforms offer the lowest Total Cost of Ownership (TCO) and fastest Return on Investment (ROI) among 18 different server hardware and server OS distributions.

Highly reliable servers, operating systems, and applications – that register a minimum of four and preferably five nines – 99.99% and 99.999% - are required to support uninterrupted, continuous data transactions and conduct business. Midsize and large enterprises in the top verticals: banking/finance, education, government, healthcare, manufacturing, retail, transportation, and utilities, increasingly require six, seven and eight nines of reliability. Organizations that fail to protect the uptime/availability of the core network infrastructure and vital applications will almost certainly suffer the monetary and business consequences.

## Introduction

ITIC's 2023 Global Server Hardware, Server OS Reliability Survey examines the internal and external issues that affect the reliability and security of server hardware, server OS and mission critical applications. It also details the myriad ways in which the inherent reliability or instability of the server hardware impacts enterprises' finances and operations at on premises datacenters, the network edge, hybrid cloud computing, remote and hybrid work environments and IoT ecosystems.

Server hardware, server operating systems and the business-critical applications running on them are the foundation upon which the lifeblood of business operations rely.

With each passing year, organizations grow most cost conscious and risk averse. They depend on highly reliable, robust, and secure servers, server OS and application software to conduct uninterrupted their daily business operations across their entire interconnected network ecosystem(s). This includes on-premises datacenters, hybrid cloud environments, the network edge and remote hybrid work environments.

The *inherent* reliability of servers, server OSES and component parts continue to demonstrate steady year over year (YoY) performance, reliability, and security gains.

These reliability gains are achieved via advances in the underlying semiconductor technology (e.g., processors, memory, hard drives, storage etc.); improvements in source code and network connectivity. There are however many other factors that can precipitate unplanned outages. They include but are not limited to:

- Security
- Human error
- Complexity in configuring and provisioning new applications

- Incompatibilities among hardware and applications
- Old or inadequate server hardware
- Challenges arising from managing remote and hybrid work environments
- Understaffed or inadequately trained IT, security, and software administrators
- Catastrophic man-made or natural disasters (e.g., severe weather events)

Organizations require the most reliable, robust, secure and feature rich server technology to fortify their infrastructure and mitigate risk.

## High Reliability and Operational Expenditure Costs

The most reliable servers that demonstrate five, six, seven, eight and even nine nines of reliability will deliver the lowest TCO and greatest economies of scale. However, this does **not** mean there are **no costs** associated with owning, operating, managing, and securing the core server and server OS infrastructure.

Every organization has operating expenses (OpEx). This is the cost organizations incur as a result of doing business every day. OpEx includes line items like salaries, rent and leasing, equipment operation and maintenance, utilities, marketing, insurance, accounting, legal fees, research and development (R&D) funds and inventory costs.

Companies will still pay operating costs regardless of how reliable their servers are. However, having a highly reliable server infrastructure ensures that a company can stay within its OpEx budget. An unstable, unreliable server environment will increase OpEx costs and has the potential to impact the company's capital expenditure (CapEx) budget as well.

ITIC launched its inaugural Global Server Hardware, Server OS Reliability survey in 2008. Every year since, ITIC updates and expands the survey questions to reflect technology advances and current business trends. Technology and the way enterprises deploy technology and transact business have undergone substantial changes in the last decade. Indeed, even the last three-to-five years have seen the rapid emergence of new technologies like: Artificial Intelligence (AI), analytics, cloud computing (particularly hybrid cloud), cognitive computing, Internet of Things (IoT) Quantum safe computing (from IBM) and the expansion of the network edge.

These technology advances center on scalability, performance, sustainability, and security and are increasingly optimized for burgeoning hybrid cloud and network edge environments. New hardware technology has yielded positive gains on server hardware reliability.

ITIC's 2023 Global Server Hardware, Server OS Reliability Survey segmented the reliability of the latest versions of mainstream mission critical servers to quantify these gains. Platforms including the IBM z16 with its quantum-safe computing security capabilities and AI on-chip inferencing performance capabilities; the IBM Power E1080 with transparent in-memory encryption for hybrid cloud security; the Lenovo ThinkSystem SR950, which engineered for "always-on" reliability with multiple levels of resiliency to protect data and the Dell PowerEdge R740XD, to name a few. This enabled ITIC to calculate the tangible business and economic improvements associated with the latest server versions.

ITIC's latest 2023 Reliability poll found that 95% of organizations cut downtime by 7% to 23% when their firms upgraded to the newest, more robust and feature-rich brand name servers from Cisco, Dell, Fujitsu, HPE, Huawei, IBM, Inspur, Lenovo, Oracle and Toshiba. Less downtime also improved employee productivity and enabled firms to better meet their revenue goals.

But new technological advances are sometimes offset by external market dynamics.

The COVID-19 pandemic (2020 through 2022) caused cataclysmic upheavals in business and economic trends that persist to the present day. During the pandemic lockdown, businesses and academic institutions transitioned to remote and hybrid workforce and education environments. Every market sector experienced disruption in their respective supply chains. This resulted in longer lead times to purchase or replace products, components, and services. This precipitated price increases and inflation.

The last three years have also seen an explosion of targeted security hacks and data breaches like ransomware and CEO fraud, which continue unabated in 2023. The number, severity and monetary cost of security incidents continue to soar. The [Ponemon Institute IBM 2023 Cost of a Data Breach Report](#) found that the global average cost of a data breach in 2023 was \$4.45 million (USD); this is a 15% increase over the last three years. And in a worrying trend, the survey revealed that 82% of data breaches involved data stored in the cloud.

All these issues make the underlying reliability of the core network infrastructure imperative.

## Survey Highlights

To reiterate, the most reliable server distributions: IBM, Lenovo, Cisco, HPE and Huawei successfully achieved their individual personal best reliability and availability scores compared to the results posted in ITIC's prior reliability studies.

- The IBM Z, IBM Power servers and the LinuxONE III and LinuxONE Emperor 4 solution ranked first or were tied for first in every category including reliability, security, performance, leading edge features and support.



- Lenovo ThinkSystem servers scored top reliability marks among all x86-based hardware for the ninth straight year averaging “six and seven nines” uptime depending on age and configuration of servers. Cisco UCS, the HPE Superdome and Huawei KunLun (in that order) mission critical servers rounded out the top five most reliable servers.
- Cisco UCS also came on strong as the company doubled down on its security initiatives, particularly at the network edge to maximize server and application availability. An 83% majority of Cisco UCS customers achieved a solid “five nines” (5.26 minutes of per server annual downtime) to “six nines” (31.5 seconds of downtime) depending on age, configuration, and workload. Cisco UCS servers cut unplanned per server annual downtime to an average of 1.27 minutes – just behind the IBM Z, IBM Power and Lenovo ThinkSystem distributions. High end mission critical Huawei KunLun and Fusion servers and HPE Superdome rounded out the top five most reliable platforms.
- Dell PowerEdge servers also improved their reliability, registering 24 minutes of unplanned per server annual outage time; this is a reduction of two (2) minutes compared to the 26 minutes of downtime Dell servers registered in ITIC’s 2022 Global Server Hardware, Server OS Reliability Survey. A 71% majority of the 22% of survey respondents who upgraded to Dell high end PowerEdge servers (e.g., the PowerEdge R6525 Server and the PowerEdge MX750c Server) in the last 12 to 18 months scored five nines – 99.999% - uptime equaling 5.26 minutes of per server annual downtime.
- Some 84% of businesses cited security and data breaches as the number one cause of unplanned downtime. This was followed by 67% of respondents who said Human error was responsible for unanticipated outages and 51% who blamed complexity in configuring/provisioning new applications.
- The IBM Z, IBM Power, IBM LinuxONE III and LinuxONE Emperor 4, Lenovo ThinkSystem and Cisco UCS (in that order) also delivered the most consistent and highest levels of security with the lowest percentage of downtime due to on-premises and cloud security hacks and data breaches.
- Finally, Cisco, Dell, IBM, and Lenovo were in a virtual four-way tie for the highest customer satisfaction and service and support. Nine-in-10 ITIC survey respondents gave them “excellent” or “very good” ratings for quality, speed, and responsiveness.

## The Nines of Reliability and the Cost of Downtime

As always, ITIC’s Global Server Hardware, Server OS Reliability Report utilizes information gathered from ITIC’s prior surveys (2008 through 2022) to compare the reliability of the various server hardware and server OS platforms. The study also tracks reliability, security, cloud and

other pertinent technology and business trends, like AI, analytics, IoT and virtualization. The findings provide crucial metrics to assist organizations in making informed purchasing decisions on which hardware and software offerings best suits their organizations’ specific business and budgetary needs.

In 2023 a 93% majority of midsized and large enterprises (with over 500 employees) now require a minimum of “four nines” (99.99%) reliability. However, 49% of organizations now strive for “five nines” (99.999%) uptime or higher – up from 25% in ITIC’s 2018 survey.

As **Table 1** illustrates, four nines of reliability equals 52.56 minutes of unplanned annual per server downtime or 4.38 minutes of per server monthly downtime. By comparison, five nines of uptime is the equivalent of just 5.26 minutes of unplanned, annual per server downtime. By contrast, six nines of uptime equals 31.5 seconds of unplanned per server yearly downtime!

## The Nines of Reliability and the Cost of Downtime

**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     |
| 99.999999% (eight nines) | 315 milliseconds   | 26.25 milliseconds | 6.05 milliseconds  |
| 99.9999999% (nine nines) | 31.56 milliseconds | 2.63 milliseconds  | 606.9 microseconds |

**Source:** ITIC 2023 Global Server Hardware, Server OS Reliability Survey

There is an **order of magnitude** of difference between the “nines” of reliability. The difference between each “nine” is usually 10; each order is either 10xs greater or 10xs smaller than the next amount.

©Copyright 2023 **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 example, the order of magnitude of difference between a millisecond and a microsecond is 1/1,000,000th of a second. A millisecond is one-thousandth of a second while a microsecond is one-millionth of a second. The ratio between the two is 1,000 to 1. Therefore, a millisecond is 1,000 times longer than a microsecond. And while a nanosecond is not yet depicted on **Table 1 above**, it is defined as one-billionth of a second.

The emergence of advanced capabilities like the Telum processor and on-chip AI inferencing make it possible for high end mission-critical servers like the IBM z16 to achieve eight nines and nine nines - 99.999999% and 99.9999999% or true fault tolerance.

These levels of uptime and availability would have been considered a pipe dream a decade ago in 2013. In fact, this level of reliability is still far beyond what the majority of servers can reliably deliver at present.

The monetary cost of downtime is another crucial barometer that measures the impact reliability and availability has on corporations' bottom line.

ITIC's 2023 Hourly Cost of Downtime survey indicates a single hour of server downtime can result in potential losses of \$300,000 or more for 93% percent of mid-sized (SMEs) and large enterprises during a single hour. And among that 93% majority, over half or 46% – of firms said hourly outage costs now exceed one million (\$1M) to over five million (\$5M).

**Table 2** illustrates the Per Minute Cost of Downtime ranging from \$100,000 to \$10 million per hour for a single server in configurations of one, 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 2023 Global Server Hardware, Server OS Reliability Survey

## Data Analysis

The disparity in reliability amongst the most and least reliable server hardware platforms is attributable to several factors, including:

- **Advanced technology and management capabilities.**
- **Strong Security.**
- **Aging server hardware.** Unsurprisingly, when ITIC sliced the survey results, it found respondents reported the most pronounced reliability declines in aged servers over 3 ½ years old that they failed to upgrade or retrofit to accommodate mission-critical workloads. Retaining servers for four, five and even six years is more common among SMBs and smaller midsize enterprises that utilize less expensive, commodity platforms (e.g., White box unbranded servers, Dell, HPE ProLiant and Oracle x86). In this group, ITIC’s survey found that approximately 53% to 64% of businesses retain the hardware for over three-and-a-half, to four, five and six years. Some firms upgrade “only as needed” or when the server performance significantly degrades, causing the server to freeze up or crash. Many customers likewise do not sufficiently right-size their servers to carry the more memory and compute-intensive application workloads. Aged servers are also much more prone to hard drive and component failures. This can prove especially

problematic if customers encounter supply chain shortages or are unable to replace the parts at all.

- **Server configuration.** Many organizations that purchase less expensive commodity servers skimp on configuration (e.g., less memory, storage, slower CPU). Additionally, the lower end servers frequently lack true enterprise capabilities that are designed for 24 x 7 operations. Commodity servers for instance may not be equipped with hot swap capabilities like memory boards, disk drives, cooling units and power supplies. They may also lack High Availability (HA) and clustering capabilities and may not support the latest and most rigorous AES encryption and security standards.
- **Vendor Service and Support.** After market technical service and support as exemplified by IBM, Lenovo, Cisco, Dell and HPE – all of whom are consistent and responsive - plays a pivotal role in customer satisfaction and retention. The more quickly a vendor responds to its customer with answers and fixes, the more likelihood that downtime will be minimized. Corporations that choose White box unbranded servers often find their businesses left entirely to their own devices if and when problems arise.
- **Regular server hardware, server OS and application upgrade and refresh cycles.**
- **Experienced IT and security administrators.**
- **Corporations' ability to adhere to the best computing practices and establish governance.**

In 2023 and beyond, \$100,000 (USD) for one hour of downtime for a single mission-critical enterprise server is extremely conservative for all but the smallest micro SMBs with one to 25 employees (**See Exhibit 7**). It equates to \$1,670 per minute/per server. The hourly cost of downtime calculated at \$300,000 equals \$4,998 per server/per minute. The cost of significant, extended outage incident that a business estimates at \$1 million (USD) per hour equates to \$16,700 per server/per minute.

## Overview: Top Survey Findings

To reiterate, ITIC's 2023 Global Server Hardware, Server OS Reliability also segmented the standalone reliability for the IBM z16 server. The results showed that the z16 (introduced in April 2022) delivers nine nines—99.999999%— of uptime and reliability. This is a scant **31.56 milliseconds** of per server annual downtime, according to the results of the ITIC 2023 Global Server Hardware, Server OS Reliability Survey.

From a monetary standpoint, IBM z16 businesses spend mere pennies on per server yearly operational expenditures related to managing the platform and performing IT remediation activities resulting from unplanned server and server OS outages.

ITIC's 2023 Global Server Hardware, Server OS Reliability survey also found that an 88% majority of the newest IBM Power 10 servers (shipping since September 2021) customers

reported their organizations achieved eight nines—99.999999%—of uptime. This is **315 milliseconds** of unplanned, per server, per annum outage time due to underlying system flaws or component failures. So, Power10 corporate enterprises spend mere pennies per server/per year performing remediation due to unplanned server outages.

Once again, the IBM z16 and Power 10 server-specific uptime statistics were obtained by breaking out the results of more than 200 respondent organizations that deployed the z16 since it began shipping in April/May 2022. A 96% majority of these z16 enterprises say their businesses achieved nine nines—99.9999999%—of server uptime. This is the equivalent of a near-imperceptible 31.56 milliseconds of per server annual downtime due to any inherent flaws in the server hardware and its various components (**See Table 1**).

The above metrics illustrate that all versions of the IBM Z (z13, z14 z15 and z16); along with the LinuxONE III and LinuxONE Emperor 4 platforms record continuous, fault tolerant levels of eight nines of reliability or, a near imperceptible 315 milliseconds of per server annual downtime. They were followed by the IBM Power8 and Power9 distributions which averaged close to seven nines or 3.15 seconds of annual per server downtime. This equates to \$5.26 of per server yearly downtime assuming an hourly downtime loss of \$100,000.

**The Lenovo ThinkSystem** servers provide the highest uptime and availability among more than one dozen x86 server distributions. The Lenovo ThinkSystem servers averaged six nines – 31.5 seconds of unplanned per server annual downtime, which is their best ever reliability results. This amounts to an annual per server, per minute potential cost of \$52.60 assuming hourly downtime losses of \$100,000. In the ITIC 2022 Reliability study, Lenovo ThinkSystem servers averaged 1.10 minutes per server, per minute downtime with a potential associated cost of \$1,837 per server, per minute.

**The Huawei KunLun and Fusion** mission critical servers clocked a per server annual average downtime rate of 1.27 minutes of unplanned per server outages. That equates to per server/per minute downtime charges of \$2,120 for a single hour of downtime calculated at \$100,000.

**HPE's high end Superdome** servers also registered high reliability: with customers reporting 1.38 minutes of per server, per minute unplanned yearly downtime. Enterprises that estimate a single hour of downtime calculated at \$100,000 would incur potential monetary per server/per minute losses of \$2,321 down from the \$2,404 based on 1.44 minutes of unanticipated per server, per minute outage costs.

**Unbranded White box servers** were the least economical and least reliable as they continued to experience the highest rate of unplanned per server monthly downtime – an average of 59 minutes of unavailability. This is the same statistic reported in ITIC's 2022 Reliability survey.

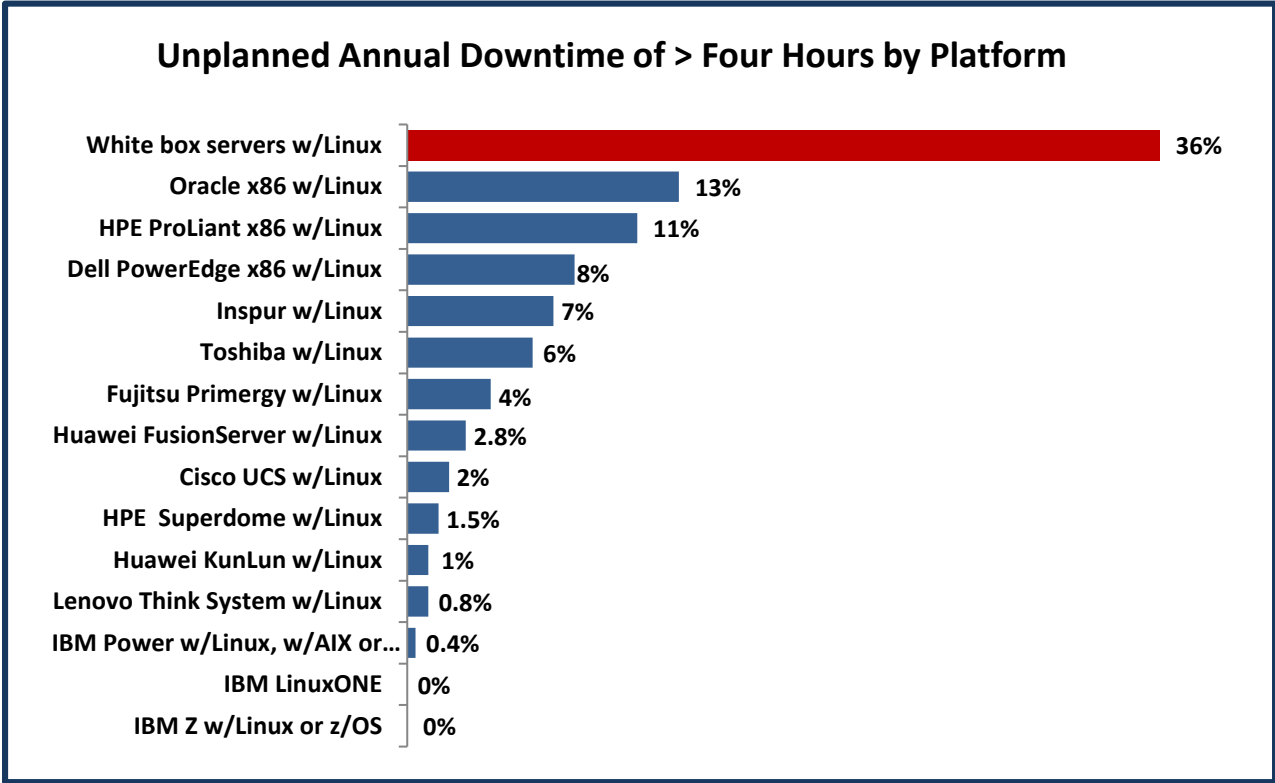
That has the potential to cost corporations \$98,530 when hourly downtime losses are calculated at \$100,000.

The IBM Z; IBM LinuxONE; IBM Power; Lenovo ThinkSystem; Cisco Systems UCS; Huawei KunLun and Fusion and HPE Superdome server distributions (in that order) achieved the highest levels of reliability and scored high in every category; including:

- The least amount of per server/per minute unplanned downtime due to server flaws.
- The least amount of unplanned per server downtime over four (4) hours.
- The fewest number of successful security hacks resulting in server outages.
- The least amount of unplanned per server downtime due to security and data breaches.
- The least amount of security-related data losses, data theft.
- The lowest amount of monetary losses.
- The fastest Mean Time to Detection (MTTD) and Meantime to Recovery (MTTR).
- The lowest Total Cost of Ownership (TCO) and fastest Return on Investment (ROI).

As **Exhibit 2** illustrates the most reliable servers: the IBM Z, IBM LinuxONE, IBM Power; Lenovo ThinkSystem, Huawei KunLun, HPE Superdome and Cisco UCS (in that order) again registered the lowest percentages of the most protracted four hours or greater per server.

**Exhibit 2.** IBM, Lenovo, Huawei and HPE Register Least Amount of Extended Outages



**Source:** ITIC 2023 Global Server Hardware/Server OS Reliability Survey

©Copyright 2023 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 IBM Power, Lenovo ThinkSystem, Huawei KunLun, HPE Integrity Superdome and Cisco UCS were close behind. Only a niche, one percent minority of each platform experienced over four hours annual downtime due to server or component flaws.

## **Downtime Comparison Costs by Server Platform**

There is no way to overstate the symbiotic correlation of the server reliability and performance gains of the top vendors like IBM, Lenovo, Cisco, HPE and Huawei and their ability to double down on security at time when security hacks and data breaches are soaring. Additionally, the most reliable server hardware distributions continually advance their respective solutions by supporting the latest chip technology and incorporating new technologies like AI, analytics and cognitive computing into their hardware.

Every increase or decline in the amount of server reliability – however slight or prolonged – will result in a commensurate positive or negative monetary cost. The reliability of the core server hardware, server OS and business-critical application infrastructure directly impacts customers’ ongoing daily business transactions and operations; employee productivity; security and intellectual property (IP); the business reputation and ultimately, the revenue stream. There are immediate monetary consequences associated with server outages of even short durations. The reliability/uptime of each server platform yields substantial financial economies of scale in terms of lower or higher TCO. Additional per server outage time of even a minute or two, can cause daily operational costs to skyrocket and raise the corporation’s risk of litigation and potential penalties associated with non-compliance or failure to meet the terms and conditions of Service Level Agreements (SLAs) with customers and business partners.

As **Exhibits 3, 4, 5 and 6** below illustrate, there is a substantial gap in the annual unplanned downtime cost comparisons among the top performers and the least stable hardware.


**A single hour of downtime calculated at \$100,000 = \$1,670 per server/per minute.**

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

**A single hour of downtime estimated at \$1,000,000 = \$16,670 per server/per minute.**



**Exhibit 3. Unplanned Annual Downtime Per Server/Per Minute Assuming Cost of \$100K**




**Source:** ITIC 2023 Global Server Hardware, Server OS Reliability Survey

**Exhibit 4** depicts the downtime YoY cost comparison statistics between ITIC’s 2022 and 2023 Global Server Hardware, Server OS Reliability surveys.

All versions of the IBM Z, the IBM LinuxONE, IBM Power, Lenovo ThinkSystem, Cisco UCS Huawei KunLun and Fusion and HPE Superdome bested their prior year reliability statistics despite challenges posed by the ongoing spikes in security hacks, supply chain disruptions, interoperability challenges and increasingly complex deployments.

**Exhibit 4. Unplanned Downtime Per Server/Per Minute Assuming Hourly Cost of \$100K**



**Source:** ITIC 2023 Global Server Hardware, Server OS Reliability Survey

The pivotal role server reliability plays in cost effectiveness is clearly evident among the various server distributions in **Exhibit 4** above. IBM’s Power8 and Power9 servers averaged close to seven nines or 3.15 seconds of annual per server downtime; this potentially costs enterprise customers \$5.26 in per server yearly unplanned outages. This is a significant cost saving compared to the \$1,670 in annual outage remediation costs due to one (1) minute of per server downtime customers spent in 2022.

In the latest 2023 ITIC Reliability poll, Lenovo ThinkSystem Server respondents improved to an average of six nines of uptime; this equates to 31.5 seconds of unplanned per server annual downtime. Assuming an hourly cost of downtime of \$100,000, a business would potentially incur remediation costs of just \$52.60 per server yearly. By contrast, in 2022 Lenovo’s ThinkSystem servers registered 1.10 minutes of yearly downtime at a potential cost of \$1,837 for one minute of per server downtime, assuming hourly downtime losses of \$100,000.

Inspur, hardware recorded 10 minutes of per server/per annum downtime – the monetary equivalent of \$16,700. Inspur’s current yearly uptime improved by one minute compared to ITIC’s 2022 Reliability survey results. This enables Inspur corporate users to realize potential cost savings of \$1,670 per server. Inspur reliability costs are roughly 5x less than unbranded

White box servers. In ITIC's latest 2023 Reliability Study, unbranded White box again recorded per server downtime of 59 minutes; calculated at an hourly downtime rate of \$100,000 which could cost corporations \$98,530 per server.

And as **Exhibit 4** also shows, Dell's popular line of PowerEdge servers averaged 24 minutes of unplanned downtime, shaving off a full two minutes off the 26 minutes of per server, per minute yearly downtime versus ITIC's 2022 Reliability survey results. In 2023, Dell PowerEdge customers have the potential to spend \$40,080 compared to the \$43,420 they may have had to spend for 26 minutes of per server unplanned downtime a year ago. That's a potential cost savings of \$3,340 per server annually.

The economies of scale between the most reliable server distributions and the less stable servers are immediately evident.

For instance, IBM Power8 and Power9 servers which registered 3.15 seconds or seven nines of unplanned yearly per server downtime costing \$5.26 can potentially save corporate customers nearly \$16,694 per server annually versus an Inspur machine that experienced an average of 10 minutes of per server unplanned annual downtime at a cost of \$16,700 assuming hourly outage loss of \$100,000.

This is not to say that the more commodity servers are bad; they are not. Dell, Fujitsu, Inspur, Oracle and Toshiba hardware customer respondents to ITIC's 2023 Reliability all achieved an average of four nines – 99.99% and five nines – 99.999% per server reliability, which is the current and acceptable industry standard.

The crux of the issue is that enterprises are extremely risk averse. Downtime is expensive, disruptive to daily operations and it leaves corporate enterprises vulnerable to security hacks.

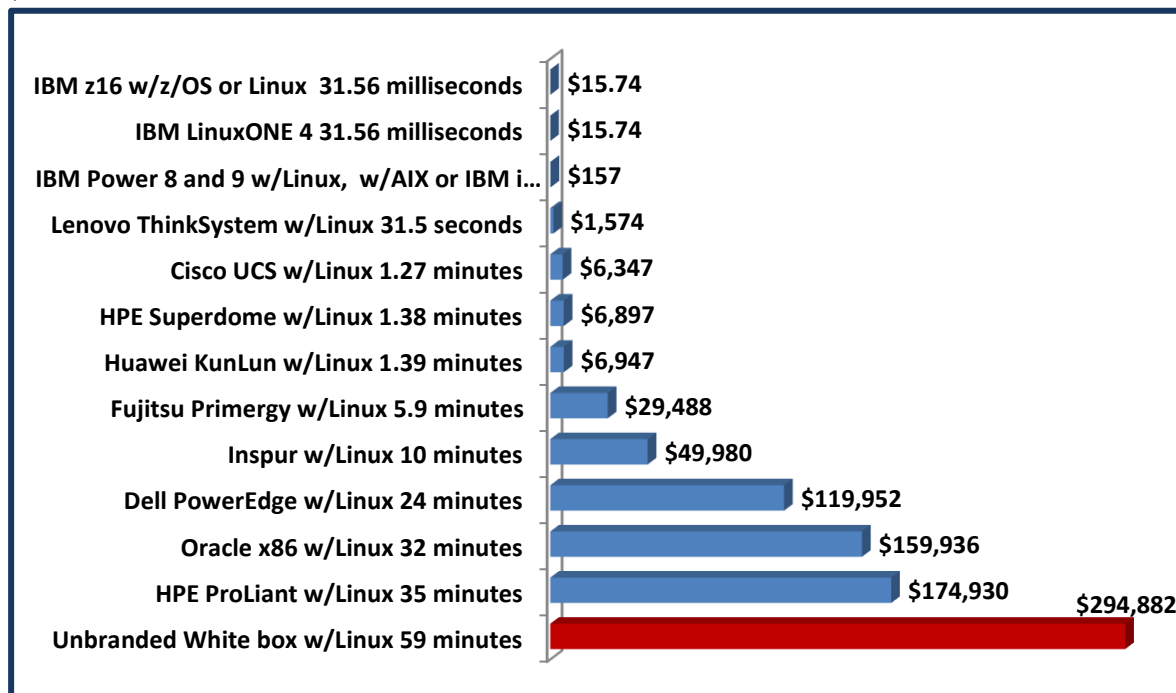
Costs quickly add up when businesses factor in the total number of affected servers across the entire network ecosystem – on premises datacenters, virtualized servers, cloud (hybrid, public and private), the network edge and remote offices. Many corporations today have virtualized server farms in their on-premises data centers; essentially all cloud computing environments are virtualized. A corporation that experienced one minute of downtime involving a single server running three or four mission critical applications would incur outage costs of \$6,680 per minute. Similarly, a company that experienced a single minute of downtime impacting 10 corporate servers, at an estimated hourly downtime rate of \$100,000 would register \$16,700 in outage-related revenue and productivity losses. These statistics are exclusive of any litigation or civil and/or criminal penalties or fines arising from the downtime. The figures also do not include the cost of any “good will” gestures, in terms of refunds or credits, a firm might make to customers, business partners or suppliers whose operations were affected by any outages.

One quarter – 25% of survey respondents said hourly downtime costs their organizations from \$301,000 to \$400,000. Overall, 93% of SME and large enterprises estimate that a single hour of downtime ranges from \$300,000 to over \$5 million (USD).

As **Exhibit 5** illustrates, higher hourly downtime cost estimates of \$300,000 for a single hour; this essentially triples per server/per minute downtime costs. Once again, the IBM, Lenovo, Huawei, HPE and Cisco hardware deliver the greatest economies of scale based on their high reliability and availability.

The IBM z14 and z15 as well as the IBM LinuxONE III deliver several orders of magnitude greater cost savings. For instance, IBM Z and IBM LinuxONE customers reported eight nines or 99.999999% reliability which is just 315 milliseconds of unplanned per server/per minute annual downtime. This equates to \$15.74 per server/per minute outage cost calculated at an hourly downtime cost of \$300,000.

**Exhibit 5.** Unplanned Annual Downtime Per Server/Per Minute Assuming Hourly Cost of \$300K



**Source:** ITIC 2023 Global Server Hardware, Server OS Reliability Survey

The IBM Power8 and Power9 per server, per minute reliability rate of approximately seven nines – **3.15 seconds** of annual per server unplanned downtime potentially cost businesses \$157 when assuming an hourly outage rate of \$300,000.

In contrast, Dell PowerEdge server survey respondents reported an average of 24 minutes of per server annual unplanned downtime. This is between four and five nines of uptime – 99.99% and 99.999%, respectively. This is the current industry minimum accepted average uptime rate. Yet Dell enterprise shops could potentially incur per server losses of cost of \$119,952 assuming an hourly outage cost of \$300,000.

This means the IBM Power servers can deliver a nearly better 100% economies of scale cost savings of \$119,795 a single server compared to the Dell PowerEdge.

The Lenovo ThinkSystem servers likewise deliver impressive cost savings. Lenovo ThinkSystem servers recorded an average of six nines or 31.5 seconds of unplanned yearly server downtime which equals \$1,574 assuming a cost of \$300,000 of hourly downtime. The Lenovo ThinkSystem servers potentially deliver an annual cost savings of \$118,378 per server, per minute compared to the Dell PowerEdge servers.

The above statistics are fluid not static. Individual corporate enterprise reliability results and costs will fluctuate depending on the amount of time in minutes and hours that each server distribution experiences annually, monthly, and weekly.

There are many other issues that also factor into final costs such as: whether or not the unplanned outage occurred during peak usage time; whether or not data was lost, stolen or damaged, compromised or destroyed; the amount of time and the number of affected employees who experienced lost productivity; the costs associated with interrupted business transactions as well as the number of IT administrators and the length of time it took to return the servers and applications to full operations.

Corporations must also tally any costs associated with litigation: was the firm sued or was it required to pay civil or criminal penalties due to regulatory non-compliance?

The uptime and availability of the least reliable hardware platforms can be improved considerably when corporate enterprises adhere to best practices. This includes right-sizing and configuring servers to accommodate current and future compute intensive applications and workloads. Organizations that elect to purchase inexpensive servers to cut capital expenditure costs should also review their upgrade cycles and not push servers beyond their acceptable limits. While a three-and-a-half or four-year refresh cycle may be adequate for a server that is not running a business-critical application, it's not advisable for a hardware platform running mission critical applications containing sensitive data or IP that is directly tied to the company's revenue stream. Strong security and getting the appropriate training and certification for IT staff and security professionals are also crucial to improving reliability. A reduction of even a few minutes of unplanned downtime can save enterprises substantial sums and mitigate risk.

Large enterprises with over 1,000 employees comprised 54% of ITIC's 2022 Global Server Hardware, Server OS survey respondents. From a monetary perspective, large enterprises

typically suffer the largest amounts per server/per minute losses. ITIC's latest survey data revealed that 44% of those polled, estimated that the average price tag for one hour of unexpected downtime exceeds \$1 million (USD) and 18% said their firms' losses surpass \$5 million (**See Exhibit 7**).

This makes the inherent reliability and security features/functions of the server all the more important. Server hardware, server operating systems and the business-critical applications they run are the foundational elements of the entire connected network ecosystem.


The superior economics of the most reliable versus least reliable servers is even more apparent for businesses that estimate or calculate their hourly downtime losses to be \$300,000; \$500,000 or \$1,000,000 or higher as depicted in **Exhibits 3, 4, 5 and 6**.

**Exhibit 6** depicts the cost of one minute of per server hourly downtime calculated at \$1 million (USD) associated with each server hardware platform based on their annual per server, per minute unplanned downtime amounts. Once again, the IBM Z and IBM LinuxONE delivered the best economies of scale: the lowest TCO and immediate ROI.

ITIC's 2023 Global Reliability report, all versions of the IBM Z (z13, z14, z15 and z16) and the IBM LinuxONE III and IBM LinuxONE 4 offerings recorded averaged eight nines of reliability or 315 milliseconds of annual unplanned per server downtime. The per server, per minute annual unplanned downtime cost is approximately \$53 dollars.

All versions of IBM Power servers (Power8, Power9 and Power 10) averaged seven nines or 3.15 seconds of unplanned downtime. Assuming hourly downtime losses of \$1 million, this equals \$526 per server, per minute downtime.

## Exhibit 6. Unplanned Downtime Per Server/Per Minute Assuming Hourly Cost of \$1Million



Source: ITIC 2023 Global Server Hardware, Server OS Reliability Survey

## Hourly Downtime Costs: 48% Estimate Losses Exceed \$1M

ITIC's 2023 Global Server Hardware and Server OS Reliability Survey found that 93% of respondents now estimate that one hour of downtime costs the firm \$301,000 or more; this is an increase of two (2) percentage points in less than two years (See Exhibit 5). Of that number, 46% of those polled indicated that hourly downtime costs now exceed \$1 million. For the third consecutive year, 99% of all sizes of small, mid, and large enterprises (SMEs) and large corporations' hourly downtime losses exceed \$100,000 (See Exhibit 7).

There are many cost variables. For instance, an issue that takes down a server(s) running a non-business essential application; or downtime that occurs in off-peak or non-usage hours, may have minimal to no impact on business operations and negligible financial consequences.

On the other end of the spectrum, cloud-based server outages involving a virtualized server running two, three or four instances of a business-critical application housed in a single physical machine have the potential to double, triple or quadruple business losses when daily business operations are interrupted and employees and business partners, suppliers and other stakeholders are denied access to critical data.

The most expensive hourly downtime scenario presented in Table 2 depicts per server/per minute outage expense impacting 1,000 servers at an organization that values an hour of downtime at \$10 million. In this example, a large enterprise could conceivably sustain crippling losses of \$166,667,000 per server/per minute.


Again, it's important to emphasize that the ITIC Hourly Downtime monetary figures represent **only the remediation costs** associated fixing the actual technical issues and business problems that caused the server or OS to fail and restoring the server and applications to full function. These statistics do **not** include legal fees, criminal or civil penalties the company may incur or any "goodwill gestures" that the firm may elect to pay customers (e.g., discounted, or free equipment or services). Enterprises must calculate any additional monies, fees and penalties separately and make a detailed accounting of each significant unplanned outage.

As **Exhibit 7** illustrates, the Hourly Cost of Downtime continues to increase. Nearly half of all midsize and large enterprise survey respondent organizations with 500 or more employees now say that a single hour of downtime results in \$1 million (USD) or more.

Thanks to inflation, in 2023 organizations of all sizes, spanning all vertical markets reported an uptick in their firms' hourly downtime costs compared to 2022. Some 29% of survey respondents say hourly downtime expenses total \$301,000 to \$400,000; this is an increase of four (4) percentage points versus 25% in 2022. Additionally, 48% of firms now say a single hour of downtime in 2023 costs their businesses one million (\$1 Million). Some 23% of survey respondents indicated their 2023 hourly downtime losses exceed five million (\$5M) an increase of three (3) percentage points from last year. Based on ITIC's Hourly Cost of Downtime research over the past decade as well as ongoing inflation impacting everything from consumer goods to transportation, fuel and utility costs, corporations should be prepared for hourly downtime costs associated with unplanned outages to increase. These costs are exclusive of any litigation, civil or criminal penalties, fines, or voluntary goodwill gestures that firms may incur.



**Exhibit 7. Hourly Cost of Server Downtime Tops \$1 Million for 48% of Enterprises**



**Source:** ITIC 2023 Global Server Hardware, Server OS Reliability Survey

Minimizing downtime means higher end user and IT and security administrative productivity; it ensures continuous daily operations; mitigates risk and improves satisfaction among customers, business partners and suppliers.

The escalating cost of computing/network outages is attributable to several factors, including:

- **Supply chain disruptions** are increasingly commonplace since COVID-19. Although the global pandemic is over, other issues stemming from climate change to volatile and unstable geo-political conflicts persist and negatively impact the global supply chain.
- **Inflationary pricing** impacting the cost of goods, services, transportation, and delivery.
- **An increase in the number of interconnected devices, systems and networks** via the Cloud and the Internet of Things (IoT) ecosystems.
- **The surge in security hacks and data breaches.** These include targeted security and ransomware attacks by organized hackers; email phishing scams; CEO fraud and a wide range of malware, viruses, and rogue code.
- **Human error.** The more people that “touch” systems and applications, the greater the likelihood of introducing errors. Everyone from CEOs, knowledge workers, IT and Security administrators down to part-time workers and company interns access corporate

servers, applications, and information. Users regularly access sensitive data assets and intellectual property (IP) via a wide array of devices and networks, many of which lack security. This creates more vulnerabilities and entry points into the network. All of which can contribute to increased downtime and higher costs.

- **Organizations’ increasing use of employee-owned and mobile devices creates opportunities and risk.** The proliferation of employee-owned personal devices: PCs, laptops, tablets and smart phones to access corporate servers, applications and sensitive data is a two-edged sword. On the plus side, the expanded access unfetters employees. Access to information is literally at their fingertips regardless of time or location. Organizations also save capital expenditure monies on the cost of purchasing desktops and laptops. Use of employee-owned devices is not without risk. They create new entry points and potential vulnerabilities into the network. Without the proper security mechanisms and corporate and IT security oversight, companies are at greater risk of being hacked if the employee-owned device is stolen or lacks sufficient security.

ITIC anticipates that all these trends – particularly security and data breaches; the ongoing hybrid work environment; the increase in cloud deployments as well as the data deluge and data sprawl will continue with no foreseeable.


Although large enterprises with over one thousand employees may experience the largest actual monetary loss totals, downtime can be equally devastating to SMBs. Smaller firms with one to 250 employees typically lack the financial resources of large corporations. Seemingly, short outages of five, 10 or 30 minutes during peak usage hours can deal SMBs a crippling monetary blow. Prolonged outages of 60 minutes or more, or a series of multiple outages of shorter durations could put SMBs at heightened risk of closure.

## **Security Hacks, Data Breaches are Top Cause of Downtime**

Security hacks and data breaches continue to surge, rising by 58% over the last 20 months, according to ITIC’s latest survey data. Those polled also said security issues constitute the most serious threat that can undermine the reliability and stability of servers throughout the entire corporate ecosystem – in datacenters, at the network edge and in public, private and hybrid clouds. Human error, which is closely linked to security outages, was cited by 67% of survey participants as a major cause of server downtime, followed by 54% of respondents who said remote working and remote learning precipitated unplanned outages.

As **Exhibit 8** illustrates, the IBM, Lenovo, Huawei, HPE and Cisco server platforms (in that order) delivered the highest levels of security and experienced the least amount of downtime related to a security hack.

## Exhibit 8. IBM, Lenovo Servers Most Secure, Toughest to Crack



Source: ITIC 2023 Global Server Hardware, Server OS Security Survey

## Security, Resiliency Reduce Downtime Costs

Presently, four nines or 99.99% uptime remains the minimum requirement for nine-in-10 companies. However, that is changing. To reiterate, ITIC's 2022 Global Server Hardware, Server OS Reliability poll found that 40% of respondents now say their firms strive for five nines - 99.999% - reliability or greater. To achieve that, robust security is imperative.

The most reliable servers all feature top-notch security, resiliency, and advanced system recovery capabilities. Consequently, their enterprise customers are well equipped to cope and quickly respond when an outage occurs.

Corporate enterprises also deserve a good deal of credit for IBM, Lenovo, Huawei, HPE and Cisco servers' high reliability and security scores. Typically, these enterprise customers are price conscious but not driven by the need to purchase the least expensive brands, delay upgrades and skimp on IT and security training for their administrators. IBM, Lenovo, Huawei, HPE and Cisco high end customers also retrofit and upgrade their server hardware and server operating systems on regular two- and three-year cycles or as needed. An 81% majority of IBM Z users and 72% of

IBM Power customers regularly retrofit or upgrade their hardware every three years compared with just 28% of Dell PowerEdge customers.

Besides the obvious technical merits of the IBM Z, IBM LinuxONE and IBM Power mission critical hardware, the server reliability is further fortified by the greater expertise of the IT administrators in corporations that use IBM. ITIC's reliability survey data shows that large corporate enterprises with a strong contingent of mission critical IBM, Cisco, HPE and Lenovo servers, the IT administrators typically have 10 or more years' experience and are current on certifications and training – particularly with respect to security. By contrast, firms (with some notable exceptions) that use less expensive, commodity Dell, HPE ProLiant, Oracle and unbranded White box servers are more likely to hire IT managers who are entry level and have one-to-five years of experience.

Finally, the IBM, Lenovo, Huawei and HPE server organizations are price conscious, but not price driven in the same way that SMBs and midsize enterprises are. Another differentiator: a higher percentage of enterprise customers deploying high end mission critical IBM, Lenovo, HPE and Huawei servers also adhere to a regular three or three-and-one-half year upgrade cycle. They will also upgrade and right-size their server hardware to adequately support advanced, more compute intensive applications, as needed. This is crucial since applications like AI, Analytics, Blockchain, IoT and Virtual Reality (VR) are resource intensive. More recently, Cisco UCS shops (many of which are deployed at the network edge, often on the front line of security attack) are also upgrading their platforms with increasing frequency and regularity to bolster security.

Commodity server users should not defer upgrades or retain servers well beyond the recommended three-year upgrade cycle. Over 60% of businesses that deploy commodity servers and unbranded White box hardware 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.

## **Enterprises that Prioritize Security, Best Practices Bolster Reliability**

The most reliable server vendors also benefited from the proactive behavioral habits and expertise of their corporate customers. On average, 80% of IBM, Lenovo, Huawei, HPE and Cisco shops, regularly upgrade and refresh their servers every three years, or as needed. This enables the servers to accommodate mission-critical workloads and easily handle compute-intensive analytics, artificial intelligence (AI) and virtual reality applications, without taking a performance or reliability hit. Additionally, enterprises that deploy IBM Z, IBM LinuxONE, IBM Power, Lenovo ThinkSystem, Huawei KunLun, and Fusion, HPE Superdome and Cisco

UCS hardware are three to four times more likely to get the appropriate training and certification – especially security awareness training – for their IT departments and security professionals.

## Business and Technology Trends Impacting Reliability

The data deluge, data sprawl as well as the rapid shift to the cloud are factors cited by 60% of ITIC Reliability survey respondents as presenting challenges to overall server and network reliability. This is bolstered by multiple security and threat reports.

The [Netskope Cloud and Threat Report: Cloud Data Sprawl](#)<sup>1</sup> released in July 2022 found that organizations' cloud application usage has increased 35% since the beginning of 2022. The study, by Netskope, a security edge vendor based in Santa Clara, CA, found that an average company of 500 to 2,000 users uploads, creates, shares or stores data in 138 different apps. And they use an average of 1,558 distinct cloud apps each month. The report found that more than one in five (22%) of users upload, create, share or store data in personal apps and personal instances, with Gmail, WhatsApp, Google Drive, Facebook, WeTransfer, and LinkedIn ranking as the most popular personal apps and instances. This creates more potential areas of vulnerability and commensurately increases the organization's managerial burden.

And the [CloudStrike 2023 Global Threat Report](#)<sup>2</sup> by CloudStrike Holdings, Inc. an Austin, Texas cyber security technology firm reported adversaries are accelerating cloud exploitation and that hackers have created a new “state of the art” for vulnerability exploitation to sidestep patches.

The ITIC 2023 Global Server Hardware, Server OS Reliability Survey responses also highlighted several ongoing trends that have the potential to directly undermine system, application and network reliability.

Security is unsurprisingly, foremost among them. To put it simply: security is big business, so much so that it constitutes its own market segment.

- **Data Breaches Continue to Surge Post-Pandemic.** Some 61% of enterprises said targeted security issues including ransomware, Phishing, CEO fraud and assorted Email scams – increased during the last 16 months. This up 23 percentage points since 2020, further reinforcing the need for robust security in the foundational server infrastructure. Servers from IBM, Lenovo, Huawei, HPE and Cisco (in that order) were the most secure

---

<sup>1</sup> The Netskope Cloud and Threat Report: Cloud Data Sprawl, July 2022. URL: <https://www.netskope.com/netskope-threat-labs/cloud-threat-report>

<sup>2</sup> CrowdStrike 2023 Global Threat Report, URL: <https://go.crowdstrike.com/2023-global-threat-report>.

as those vendors make ongoing substantial investments/improvements in embedded security features.

- **Mixed Outlook for Semiconductors: Political instability, Supply Chain Issues and Inflation negatively impact Server Reliability.** Supply chain issues remain challenging and paint a nebulous picture for corporate enterprises. The situation is not as critical as it was during the COVID-19 pandemic. And certain sectors like Automotive are seeing supply chain woes easing somewhat, extended lead times in general are expected to persist at least through the end of 2023, according to investment and financial analyst firms. [Deloitte's 2023 Semiconductor Industry Outlook](#)<sup>3</sup> report says the war in Ukraine is disrupting supply chains, access to important raw materials, and energy prices worldwide and energy prices worldwide—and especially in Europe. Additionally, the Deloitte report notes that “...the US government’s steps in October 2022 to tighten the rules around the export of advanced semiconductor technologies to China will likely shape the entire industry (including the downstream customers) for 2023.” That is already happening as server vendors - including IBM, Lenovo, Huawei and others.
- **IBM, Lenovo, Huawei, HPE and Cisco servers deliver best security.** IBM, Lenovo, Cisco, Huawei, HPE and servers (in that order) also attained the highest levels of security for the third straight year recording the fewest number of successful hacks. Following on the results from the 2022 survey, ITIC’s latest 2023 Global Server Hardware Security and Reliability poll found that IBM, Lenovo, Huawei and HPE mission critical servers experienced the lowest percentages of downtime due to successful security hacks and data breaches. A miniscule 0.1% IBM Z, IBM LinuxONE III and LinuxONE Emperor 4 servers were hacked. And of that 0.1%, a 92% majority % of IBM survey respondents were able to detect, isolate and thwart the hack in five-to-10 minutes. Among the other server distributions: 74% of Lenovo; 71% of Huawei and HPE survey participants indicated the data breach was due to an unsecured attached employee-owned device (e.g., a PC, laptop, notebook, tablet, or smart phone) that enabled hackers to access the servers. Among mainstream hardware platforms, only four percent of IBM Power and Lenovo ThinkSystem users reported their systems were successfully hacked. And only five percent of Huawei KunLun and HPE Integrity Superdome servers suffered a security breach. Once again, unbranded White box servers had the highest percentage – 39% of successful security hacks and data breaches; consequently, 44% of White box servers also experienced downtime due to the attack.

---

<sup>3</sup> Deloitte 2023 Semiconductor Industry Outlook, March 2023. URL: <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/technology-media-telecommunications/us-tmt-semiconductor-industry-outlook.pdf>

- **IBM, Lenovo and Huawei KunLun Servers Lowest Percentage of Hardware Failures:** As in the prior ITIC 2022 Reliability Update Survey, the most recent 2023 survey statistics show that IBM, Lenovo and Huawei's KunLun platforms continue to experience the fewest hard drive quality or failure issues among all of the server distributions within the first year of usage. Less than one percent – 0.1% – of IBM z13, z14, z15 and z16 servers experienced technical problems with their hardware in the first year of usage, followed by the IBM Power with 0.5% and Lenovo ThinkSystem with 0.6% and Cisco UCS with 1.5% each during the first 12 months of deployment.
- **Increase in Server Workloads** causes reliability declines in 78% of servers over four (4) years old that haven't been retrofitted or upgraded to accommodate increased workloads.
- **Planned Downtime Increases:** Seven-in-10 enterprises now spend from two-to-eight hours monthly on planned downtime. The majority of planned downtime is attributable to applying security patches, resolving incompatibilities among various hardware and applications and testing and provisioning new applications.

## Conclusions

ITIC's 2023 Global Server Hardware and Server OS Reliability Survey findings indicate that for the 15<sup>th</sup> consecutive year, the IBM Z remains the undisputed and unsurpassed reliability server leader, along with the LinuxONE III and LinuxONE Emperor 4 solutions. The IBM Z and LinuxONE maintained their best-in-class status in every category across the board: reliability, security, scalability, performance, sustainability and manageability. The IBM Z and the IBM LinuxONE III and LinuxONE Emperor 4 are in a class of their own. A 96% majority of survey participants said the z13, z14 and z15 and LinuxONE III delivered eight nines -99.999999% - of continuous fault tolerant reliability. IBM Z and LinuxONE III enterprises said they essentially experienced no discernible downtime and the associated outage costs were likewise almost undetectable to the corporate bottom line.

The IBM Power8 and Power9 servers also met and exceeded their best ever reliability metrics – with 95% of the Power platforms averaging six and seven nines of server uptime, while the IBM Power10 achieved eight nines – 315 milliseconds of unplanned per server annual downtime.

The Lenovo ThinkSystem servers ranked as the most reliable x86-based hardware platform for the ninth straight year averaging between six nines or 31.5 seconds of per server unplanned annual downtime and seven nines or 3.15 seconds of unanticipated per server yearly downtime.

The Cisco UCS servers registered 1.27 minutes of unplanned per server downtime, followed closely by the HPE Superdome and Huawei KunLun servers with 1.38 minutes and 1.39 minutes respectively of per server, annual unplanned outages. They averaged between five and six nines of reliability – according to eight-in-10 corporate survey respondents.

The top performing and most reliable server distributions consistently achieve high reliability year after year because they 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, security and sustainability.

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. And when the IBM Z, IBM LinuxONE III, IBM Power, Lenovo ThinkSystem, Huawei KunLun and HPE Superdome 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, three-to-five (5) minutes to restore to full operation.

Organizations are cost conscious and extremely risk averse. They depend on highly reliable, robust, and secure servers, server OS and application software to conduct uninterrupted their daily business operations across their entire interconnected network ecosystem(s). This includes on-premises datacenters, hybrid cloud environments, the network edge and remote hybrid work environments.

It is the responsibility of the vendors to deliver reliable products and top-notch technical service and support. Corporations also bear responsibility for keeping their IT departments well-staffed and providing the necessary training and certification to IT administrators. Achieving optimum uptime means upgrading refreshing server hardware as necessary to support more data intensive workloads and physical, virtual and cloud environments. Close attention must be paid to system integration and interoperability, patch management and documentation. Business performance will almost certainly suffer if server configurations are inadequate for current tasks and requirements.

Unanticipated downtime is disruptive and expensive. It can also irreparably damage a company's reputation. In extreme cases, the business and monetary losses due to unstable servers and applications can cause companies to go out of business. particularly severe or prolonged outage from unreliable hardware, natural disasters or a targeted security incident like a ransomware attack or phishing scam.

Organizations whose server hardware, operating system, virtualization components fail to deliver the minimum acceptable standard of four, five and increasingly, six nines - of per server/per annum downtime are playing Russian roulette with the health of their foundational infrastructure and network ecosystem. Unreliable servers place organizations at heightened risk for prolonged downtime. Likewise, insecure servers will undermine the reliability of the entire corporate infrastructure leaving businesses vulnerable to security hacks and data breaches.

An unreliable server is an insecure server. An insecure server is an unreliable server.



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. No hardware, software or device can deliver 100% security. Any system can be hacked, and the reliability of any system can be compromised or undone by human error.

Server hardware, server operating system and the business-critical applications running on them are the bedrock and foundation upon which the lifeblood of business operations rely.

Reliability or lack thereof raises the risk of litigation and non-compliance with industry and government regulations which can leave the business open to civil and criminal penalties. The downtime associated with unreliable systems can also potentially damage a company's reputation and result in lost business.

Corporate enterprises owe to themselves and their businesses to test and comparison shop and deploy the most reliable, robust, and secure servers and server OS software.

## Recommendations

An organization's business and fiscal success depends on its ability to achieve four, five, six, seven and even the vaunted eight or nine nines of server hardware, server OS and server application reliability/uptime and availability.

Server vendors and corporate enterprise customers must partner to attain these continuous reliability goals. Server and server operating system, virtualization, security and cloud vendors and their clients bear equal responsibility to achieve high reliability and security from the core network infrastructure to the network edge, to the cloud and across all interconnected IoT systems.

ITIC advises organizations to:

- **Know what's on your network.** Conduct regular and thorough reviews of the current infrastructure. Analyze and measure the uptime and reliability of mission critical servers, server operating systems and applications.
- **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 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.
- **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. Compile a detailed list of what IT and security staff participated in the remediation and what actions were taken. This is an invaluable resource should the problem recur. It may also serve to contain and minimize reliability-related incidents.
- **Be vigilant about security.** Construct a comprehensive security plan and regularly review and update it annually or as needed. The hackers constantly hone their skills. Businesses must keep pace with cyber criminals. Organizations of all sizes and in all verticals should conduct vulnerability testing and regularly review and upgrade security policies, procedures and products. Install the latest security updates. Regular vulnerability testing will expose potential entry points and holes in your company's defenses – on premises and at the network edge. Make sure your security 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 current server and server OS environment allows them to achieve optimal reliability.
- **Maintain Regulatory Compliance.**
- **Don't Defer Upgrades.** 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 demand 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. In the 21<sup>st</sup> century digital era of 24 x 7 operations, there is also no “good time” for downtime. But there are worse case scenarios. For example, a 15- or 20-minute outage that occurred in off-hours may have negligible consequences, while a server that goes down for three minutes and disrupts 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 to conduct formal reviews on at least an annual basis to ensure all parties are fulfilling the terms and conditions of the SLA agreements.
- **Define measure and monitor reliability and performance metrics.** Always 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.

## Survey Methodology

ITIC’s *2023 Global Server Hardware and Server OS Reliability Survey*, polled 1,900 C-level executives, IT and security administrators in corporations worldwide from March through July 2023. The independent Web-based survey included multiple choice questions and one Essay question. To maintain objectivity, ITIC accepted no vendor sponsorship. None of the participants received any remuneration. ITIC analysts also conducted two dozen first person customer interviews to obtain anecdotal data and gain deeper, broader insights and contextual knowledge of the trends that can bolster and/or undermine the economics and efficiency of daily business operations and technology purchasing decisions. ITIC employed authentication and tracking mechanisms to ensure data integrity, prevent tampering and prohibit multiple responses by the same parties.

## Survey Demographics

Survey respondents were culled from a wide range of small and medium businesses (SMBs) with fewer than 50 workers, to the largest, global, multinational enterprises with over 100,000 employees.

All market sectors were well represented: SMBs with one-to-100 employees accounted for 24% of the respondents. Small and medium enterprises (SMEs) with 101-to-1,000 workers represented 27% of the participants. The remaining 49% of respondents came from large enterprises with 1,001 to over 100,000 employees. Survey respondents hailed from 49 different vertical markets. Approximately 60% of respondents hailed from North America and 40% were international customers from more than 30 countries throughout Europe, Asia, Australia, New Zealand, Central and 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/itic-2022-global-server-reliability-results/>

<https://itic-corp.com/ibm-z-ibm-power-systems-lenovo-thinksystem-servers-most-secure-toughest-to-crack/>

<https://itic-corp.com/server-and-application-by-the-numbers-understanding-the-nines/>

<https://itic-corp.com/security-data-breaches-top-cause-of-downtime-in-2022/>

<https://itic-corp.com/ibm-lenovo-hpe-and-huawei-servers-remain-reliable-and-secure-as-security-hacks-data-breaches-surge/>

<https://itic-corp.com/44-of-enterprises-say-hourly-downtime-costs-top-1-million-with-covid-19-security-hacks-and-remote-working-as-driving-factors/>

<https://itic-corp.com/blog/2021/05/ibm-lenovo-and-huawei-servers-most-secure-suffer-fewest-hacks-as-covid-19-data-breaches-surge/>

<https://itic-corp.com/blog/2020/06/forty-percent-of-enterprises-say-hourly-downtime-costs-top-1million/>

<https://itic-corp.com/blog/2020/05/itic-2020-reliability-poll-ibm-lenovo-hpe-huawei-mission-critical-servers-deliver-highest-uptime-availability/>

<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/>