

Information Technology Intelligence Consulting Corp.

Information Technology Intelligence Consulting



ITIC 2025 Global Server Hardware, Server OS Reliability Report

February 2026

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: 55% Estimate Losses Exceed \\$1M](#).....23
- [Security, Human Error and Complexity Top Causes of Downtime](#) 26
- [Security, Resiliency Reduce Downtime Costs](#)27
- [Enterprises that Prioritize Security, Best Practices Bolster Reliability](#)28
- [AI Projects, Supply Chain Issues and IT Skills Shortages Impact Reliability](#).....29
- [Conclusions](#) 36
- [Recommendations](#) 39
- [Survey Methodology](#) 41
- [Survey Demographics](#) 41
- [Appendices](#)..... 41

Executive Summary

The IBM Z and IBM Power continue to dominate, delivering the best server reliability, uptime, and security for the 17th straight year.

For the second consecutive year, Lenovo's ThinkSystem high-end mission critical servers are now tied with IBM Power Systems for uptime and reliability with the latest versions of the Lenovo ThinkSystem and IBM Power10 and Power11 delivering an average of "eight nines" of uptime. Lenovo's ThinkSystem hardware also continues to deliver the top reliability among all x86 servers for the 12th consecutive year.

Cisco Systems, Huawei KunLun, Hewlett-Packard Enterprise (HPE) Superdome mission critical servers also improved their uptime, registering an average of five, six and seven nines and higher reliability/availability to lessen the gap between their servers and the leaders. Cisco UCS improved its uptime statistics with robust network edge reliability and security of 1.10 minutes of unplanned per server downtime, annually.

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 servers likewise provide over 40x more uptime than the least reliable systems. Cisco UCS, HPE Superdome and Huawei KunLun delivered the best TCO among x86 platforms ranging from 15x to >35x more economical than less reliable rivals.

An 88% majority of corporate respondents cited security as the top cause of unplanned downtime followed by 75% saying human error causes unplanned outages. Some 67% of large enterprises indicated they are encountering delays relating to the complexity of their AI project. In some instances, this necessitates taking servers offline resulting in unavailability and downtime.

IBM, Lenovo, Cisco, HPE and Huawei maintained and improved reliability scores with large investments in AI, Generative AI and cloud offerings/services.

For the 17th consecutive year, IBM's Z mainframe reigned supreme delivering true fault tolerant "nine 9s" – 99.9999999% reliability.

Among mainstream server distributions, IBM's Power Systems along with the Lenovo ThinkSystem servers for the 12th straight year - dominated the reliability and security landscape delivering the most uptime. These three server distributions continue to deliver the highest levels of uptime, availability and best economies of scale among 18 server hardware and operating

system platforms. Those are the results of ITIC's 2025 Global Server Hardware, Server OS Reliability Survey which polled nearly two thousand organizations worldwide.

All versions of the [IBM Z](#) systems (z13, z14, z15, z16 and z17) averaged “eight nines” or better of uptime and availability. These were the best results among 18 different server platforms. This equals 315 milliseconds per server annual downtime. The ITIC survey found that the z16 release (introduced in April 2022) and the newest model z17 (shipping since June 2025) both attained a near perfect “nine nines” or 99.9999999% continuous fault tolerant reliability of 31.56 milliseconds of per server annual downtime (**See Exhibit 1**). 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.

IBM's Power10, [IBM Power Virtual Server \(VS\)](#) a family of configurable, multi-tenant, virtual IBM Power servers with access to IBM Cloud services and several high-end mission critical Lenovo [ThinkSystem servers](#) – including the SR950 V3 Mission Critical Server and the SD650-I V3 Supercomputing Server – achieved parity with IBM's Power10 and Power11 delivering a superior “eight nines” or 315 milliseconds of unplanned annual server downtime.

All Lenovo ThinkSystem servers recorded their best-ever reliability ratings for the 12th straight year, continuing to outperform all x86 servers. A 90%+ majority of **all versions** Lenovo ThinkSystem servers that were one, two and three years old, recorded “six and, seven nines” or better of uptime.

ITIC's 2025 Global Server Hardware, Server OS Reliability survey indicated that 96% of the legacy, mature IBM Power8 and Power9 hardware 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, a 97% majority of the newest IBM Power11 (released July 25, 2025) and its predecessor Power10 distributions (shipping since September 2022) 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. Therefore, Power11 and Power 10 corporate enterprises spend a scant \$0.7 cents per server in annual remediation costs that occur due to unplanned server outages.

ITIC calculated the IBM z16, z17 and Power10 and Power11 server-specific reliability data by analyzing the results of 215 enterprises utilizing the z16 (shipping since April/May 2022) and 38 organizations deploying the latest z17, released in June 2025. A 98% majority of z16 and z17 users reported their businesses achieved nine nines—99.9999999%—of reliability or 31.56


milliseconds of unplanned annual per server downtime (**See Table 1**). To put these statistics into perspective: the latest z16 and z17 enterprises spend virtually no operational monies on restoring the IBM mainframes because of unanticipated server outages.

The IBM Z and Power10 and Power11 server-specific uptime statistics were obtained by breaking out the results of more than 200 respondent organizations that deploy those distributions in production environments. A 96% majority of z16 and z17 customers 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 similarly posted their best reliability scores: 96% of the over 400 Lenovo survey participants said they achieved at least six nines or better of uptime. Lenovo enterprise respondents likewise indicated their organizations have minimal to This is the highest reliability score among all x86 hardware platforms for the 12th straight year.

Cisco UCS, the HPE Superdome and Huawei KunLun and Fusion hardware again rounded out the top five most reliable server platforms. These servers likewise improved their uptime and reliability and 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 2025 Global Server Hardware Server OS Reliability Survey

In the digital age, even few minutes of unplanned downtime can devastate daily operations, upend productivity and negatively impact revenue. Outages that occur during peak usage times, interrupt or abrogate critical business transactions (s) potentially can cost organizations thousands, tens of thousands or even millions for each minute servers, devices and applications are unavailable. The potential consequences of an unplanned outage will have a ripple or 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.

Organizations owe it to themselves to deploy the most reliable, robust, and secure servers and server OS software to deliver continuous business operations; ensure regulatory compliance, mitigate risk and protect the integrity and security of sensitive data and intellectual property (IP).

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.

©Copyright 2026 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.

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 ensure the uptime/availability of the core network infrastructure and vital applications in their on-premises datacenters, the network edge and the cloud will almost certainly suffer monetary and business consequences. Unreliable servers, operating systems and applications will heighten the risk of regulatory non-compliance and leave companies vulnerable to litigation as well as civil, criminal penalties and fines.

Introduction

The ITIC Global Server Hardware, Server OS Reliability is an independent Web-based survey conducted annually since 2009. The study compares the reliability, performance, and security of 18 leading mainstream on-premises and cloud-based servers and operating systems. The most recent ITIC 2025 Global Server Hardware, Server OS Reliability survey polled 2,000 organizations worldwide from February through December 2025.

Respondents included small and mid-sized businesses as well as multinational enterprises with over 100,000 employees across 39 vertical markets. 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 ITIC annual Global Server Hardware, Server OS Reliability survey and the separate, but associated yearly ITIC Global Server Security survey both examine the top technical and business challenges, which can positively or negatively affect the performance, availability and reliability of organizations' infrastructures. The surveys also polled customers on the monetary and business costs of outages and the commensurate risks. To maintain objectivity, ITIC accepted no vendor sponsorship. No participants received remuneration.

ITIC's 2025 Global Server Hardware, Server OS Reliability Survey examines the internal and external issues that positively and negatively impact the reliability and security of server hardware, server OS and mission critical applications. It also details the various ways in which the inherent reliability or instability of the server hardware affects 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 bedrock and lifeblood of business operations.

With each passing year, organizations grow more risk averse. They depend on highly reliable, robust, and secure servers, server OS and application software to conduct uninterrupted daily business operations across their entire interconnected network ecosystem(s). This includes on-premises datacenters, public, private and 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 technologies and applications, most recently Agentic AI, IoT and network edge deployments.
- Incompatibility among server hardware and applications.
- Outdate 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 business has operating expenses (OpEx). These are the costs incurred in daily routine business operations. 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 pay operating costs regardless of how reliable their servers are. However, a highly reliable server infrastructure ensures businesses can stay within its OpEx budget. An unstable, unreliable server environment will increase OpEx costs and has the potential to negatively influence the company's capital expenditure (CapEx) budget as well.

The five years have seen the rapid emergence and adoption of advanced technologies. These include but are not limited to: AI, AI on-chip inferencing; advanced prescriptive analytics; AI powered surveillance systems; Biometrics and touchless access; cloud computing (particularly hybrid cloud); cognitive computing; Edge AI and Environmental Internet of Things (IoT) sensors; Quantum safe computing (from IBM) and the expansion of the network edge.

These technological advances revolve around scalability, performance, sustainability and both cyber-hardened hardware infrastructure security as well as increased physical security, i.e., such as [UL 752 Level 3 certified ballistic-resistant sliding doors](#), designed for high-risk environments. These cutting-edge technologies are specifically optimized for existing on-premises datacenters, burgeoning hybrid cloud, network edge and remote environments. They have resulted in positive gains on server hardware uptime and by association, server OS, application and network reliability.

ITIC's 2025 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 and z17 with 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 is engineered for "always-on" reliability with multiple levels of resiliency to protect data and the Dell PowerEdge R740XD, to name a few. This allows ITIC to calculate tangible business and economic improvements associated with the latest server distributions as well as determine ROI.

ITIC's latest 2025 Reliability poll found that the 92% of organizations that regularly refreshed or retrofitted their servers every three years – or as necessary - to accommodate increased compute-intensive workloads reduced downtime by 5% to 25% depending on individual server configuration, workload and capacity. Less downtime improved daily operational efficiency, employee productivity and enabled firms to meet their revenue targets.

But even the most impressive new technological advances are *potentially undermined* by both internal corporate issues and external market dynamics.

The last five years have also seen an explosion of targeted security hacks and data breaches like ransomware and CEO fraud, which continue unabated in 2026. The number, severity and

monetary cost of security incidents continue to soar. The [IBM Cost of a Data Breach Report 2025](#) found that the global average cost of a data breach in global average breach cost declined to USD 4.44 million from USD 4.88 million in 2024. This is a 9% decrease and a return to 2023 cost levels. On average, 13% of organizations reported breaches that involved their AI models or applications. However, among those that did, almost all (97%) lacked proper AI access controls.

Security teams using AI and automation extensively shortened their breach times by 80 days and lowered their average breach costs by USD 1.9 million compared to organizations that didn't use these solutions. Nearly a third of organizations said they used these tools extensively across the security lifecycle—in prevention, detection, investigation and response.

The IBM Cost of a Data Breach Report 2025 further reported that a “...majority of breached organizations (63%) either don't have an AI governance policy or are still developing one. Even when they have a policy, less than half have an approval process for AI deployments, and 61% lack AI governance technologies.” Among organizations that have governance policies in place, only a minority (34%) perform regular audits for unsanctioned AI, according to the IBM study. It shows AI remains largely unchecked as adoption outpaces both security and governance. On a more positive note, the latest IBM Cost of a Data Breach 2025 Report found that organizations that utilize security AI and automation extensively to prevent hacks, realized an average cost savings of \$1.9 million (USD).

All these issues are persistent. This underscores the need for the foundational network infrastructure to have robust reliability and strong security.

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 OverLinuxONE Emperor 4 and LinuxONE Emperor 5 solutions ranked first or tied for first in every category including reliability, security, performance, leading edge features and support.
- Lenovo high end enterprise ThinkSystem servers achieved parity with IBM Power10 and Power11 servers attaining “eight nines” or 315 milliseconds of uptime and reliability for the second straight year. All Lenovo ThinkSystem server distributions achieved top reliability marks among 18 different x86-based hardware for the 12th straight year averaging “six and seven nines” uptime depending on age and configuration of 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 87% majority of Cisco UCS customers achieved a solid “five nines” (5.26 minutes of per server annual downtime) and “six nines” (31.5 seconds of downtime) depending on age, configuration, and workload. The latest Cisco UCS servers cut unplanned per server annual downtime to an average of 1.10 minutes, down from 1.20 minutes.
- High end mission critical HPE Superdome and the Huawei KunLun and Fusion servers and rounded out the top five most reliable platforms also scored between “four, five and six nines” of reliability averaging 1.25 minutes and 1.34 minutes, respectively of unplanned per server annual downtime.
- Dell PowerEdge servers maintained their reliability, registering 20 minutes of unplanned per server annual outage time for newer models; this is a reduction of four (4) minutes compared to the 24 minutes of downtime Dell servers registered in ITIC’s 2023 Global Server Hardware, Server OS Reliability Survey. However, Dell servers, older than three years, saw an uptick in downtime of approximately two to three minutes per server for an average of 22 to 23 minutes per server, per annum downtime. 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 89% - or nine-in-10 businesses - cited security and data breaches as the number one cause of unplanned downtime. This was followed by 77% of respondents who attributed human error to unanticipated outages and 66% who blamed complexity in configuring/provisioning new applications.
- Over two-thirds of survey respondents – 67% - reported delays in rolling out new AI projects in production environments due to complexity and backwards compatibility issues. These delays often cause them to take servers offline, resulting in unexpected downtime, which in turn, is reflected in higher OpEx costs.
- The Lenovo ThinkSystem and IBM Power10 and Power11 hardware followed by Cisco UCS, HPE Superdome and Huawei KunLun (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, HPE and Lenovo all received high marks for customer satisfaction, 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 2025 Global Server Hardware, Server OS Reliability Report utilizes information gathered from ITIC's prior surveys (2008 through 2024) 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 guide organizations to make informed purchasing decisions on which hardware and software offerings align with the organization's specific business, technology and budgetary needs.

In 2025-2026, a 96% majority of midsized and large enterprises (with over 500 employees) now require a **minimum** of "four nines" (99.99%) reliability. However, "five nines" of reliability is quickly becoming the gold standard; 58% of organizations say their organizations strive for 99.999% uptime or higher. This is an increase of 29% percentage points compared to ITIC's 2020 survey, just five years ago.

As **Table 1** illustrates, four nines of reliability equal 52.56 minutes of unplanned annual per server downtime or 4.38 minutes of per server monthly downtime. By comparison, five nines of uptime are 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 2025 Global Server Hardware, Server OS Reliability Survey

There is an **order of magnitude** of difference between each of 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.

The order of magnitude 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 and z17 to achieve eight nines and nine nines - 99.999999% and 99.9999999% or true fault tolerance.

©Copyright 2026 **Information Technology Intelligence Consulting Corp. (ITIC)** All rights reserved. Other products and companies referred to herein are trademarks or registered trademarks of their respective companies or mark holders.

These superior levels of uptime and availability would have been considered unattainable a decade ago. And true, mainframe fault tolerance continues to elude beyond what most servers can consistently deliver.

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

ITIC’s 2025 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 2025 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 multiple factors, including:

- **Advanced technology and management capabilities.**
- **Strong Security.**
- **The number of employees** who have physical access to the server as well as the number employees who can make changes to source code key applications.
- **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 mid-size 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 55% to 60% 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 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 2026 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 2025 Global Server Hardware, Server OS Reliability segmented the standalone reliability for the IBM z16 and z17 mainframes. The results indicate that the z16 (introduced in April 2022) and the z17 (shipping since June 2025) both deliver nine nines—99.9999999%—of uptime and reliability. This is 31.56 milliseconds of per server annual downtime, according to the results of the ITIC 2024 Global Server Hardware, Server OS Reliability Survey.

From a monetary standpoint, IBM z16 and z17 businesses spend a negligible amount 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 2025 Global Server Hardware, Server OS Reliability survey also found that an 91% majority of IBM Power10 and Power11 server (shipping since mid-year 2025) 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. Power10 corporate enterprises spend mere pennies per server/per year performing remediation due to unplanned server outages.

Once again, the IBM z16, z17 and Power10 and Power11 server-specific reliability statistics were obtained by breaking out the results of more than 200 respondent organizations that deployed these distributions since their deployments in production environments. A 96% majority of these z16 and z17 customers 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 (from the z13 to the newest z17); along with the LinuxONE Emperor 4 and LinuxONE 5 (introduced in April 2025) platforms record continuous, fault tolerant levels of eight nines of reliability; this is a nearly imperceptible

315 milliseconds of per server annual downtime. The older IBM Power8 and Power9 distributions averaged close to seven nines or 3.15 seconds of annual per server downtime. This equates to \$5.26 per server yearly downtime, assuming an hourly downtime loss of \$100,000.

At the same time, mission critical distributions of Lenovo ThinkSystem servers are now in a virtual dead heat with the IBM Power10 and Power11 for reliability and uptime. Additionally, **all** versions of the Lenovo ThinkSystem hardware continue to provide the highest uptime and availability among more than one dozen x86 server distributions for the 12th consecutive year. In the ITIC 2025 Reliability study, Lenovo ThinkSystem servers averaged seven nines or 3.15 seconds of annual per server downtime. This equates to \$5.26 per server yearly downtime, assuming an hourly downtime loss of \$100,000. This is a big decline from the over one minute per server, per minute downtime that Lenovo ThinkSystem servers recorded in ITIC's 2023 Global Server Hardware, Server OS Reliability survey. In 2023, the Lenovo servers' one minute of per server/per annum downtime carried a potential cost of \$1,837 per server, per minute based on an estimated Hourly Downtime cost of \$100,000. In the latest ITIC 2025 Reliability poll, Lenovo ThinkSystem servers experienced 315 milliseconds of per server, per annum downtime. This effectively reduced the yearly cost of ThinkSystem hardware downtime to \$5 per server down from \$1,837. This is a substantial \$1,832 decrease in annual outage costs per server. Therefore, in 2025 and 2026, a business that has 50 Lenovo ThinkSystem servers would potentially only accrue \$250 in downtime costs compared with \$91,850 related to outage costs two years ago in 2024 (See Exhibit 4).

HPE's high-end Superdome servers also registered high reliability: with customers reporting 1.35 minutes of per server, per minute unplanned yearly downtime in 2025. 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 compared with just two years ago.

Unbranded White box servers were the least economical and least dependable as they continued to experience the highest rate of unplanned per server monthly downtime – an average of 60 to 63 minutes of unavailability. This is an increase of one minute of additional annual per server unplanned downtime reported in ITIC's 2023 Reliability survey. That has the potential to cost corporations \$100,200 when hourly downtime losses are calculated at \$100,000.

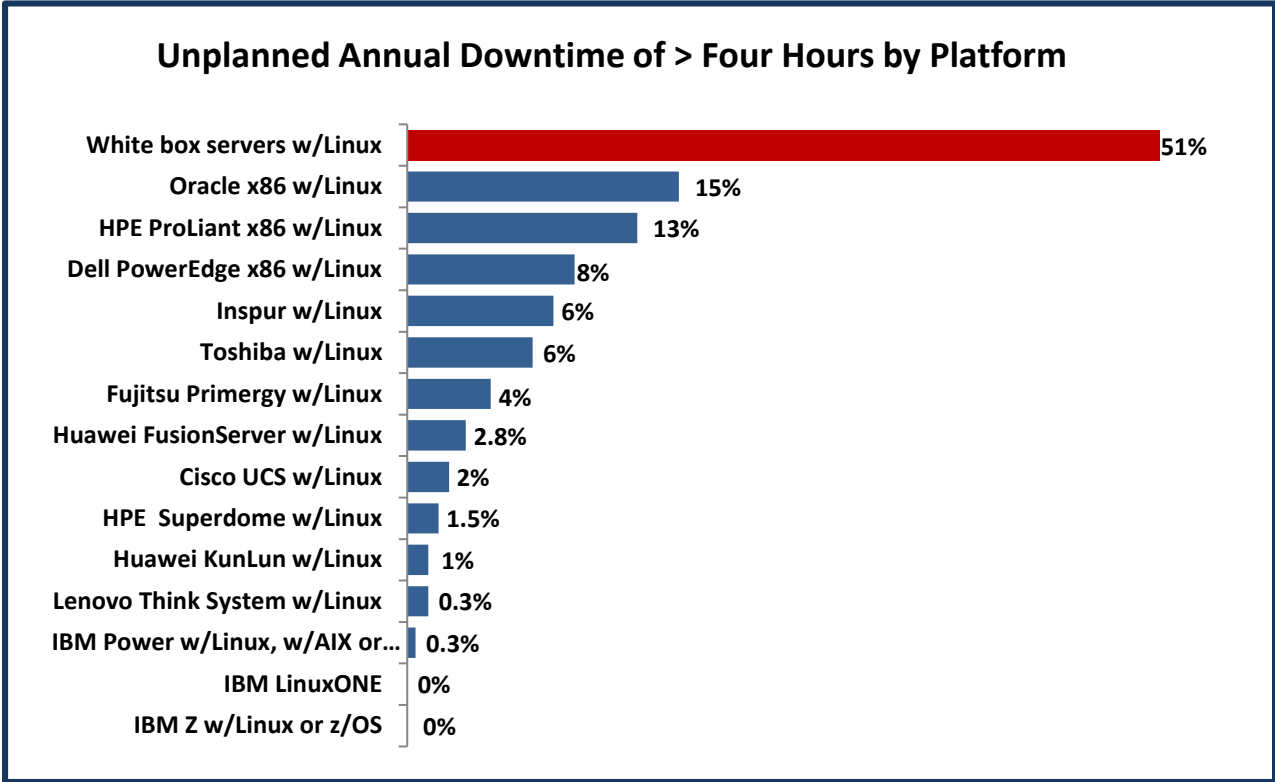
The IBM Z; IBM LinuxONE; IBM Power; Lenovo ThinkSystem; Cisco Systems UCS; Huawei KunLun and Fusion and HPE Superdome server distributions (in that order) achieved the highest levels of reliability and scored high in every category; including:

- 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 number 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 2025 Global Server Hardware/Server OS Reliability Survey

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 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 innovative 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’

©Copyright 2026 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.

ongoing daily business transactions and operations; employee productivity; security and intellectual property (IP); the business reputation and 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. An 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.

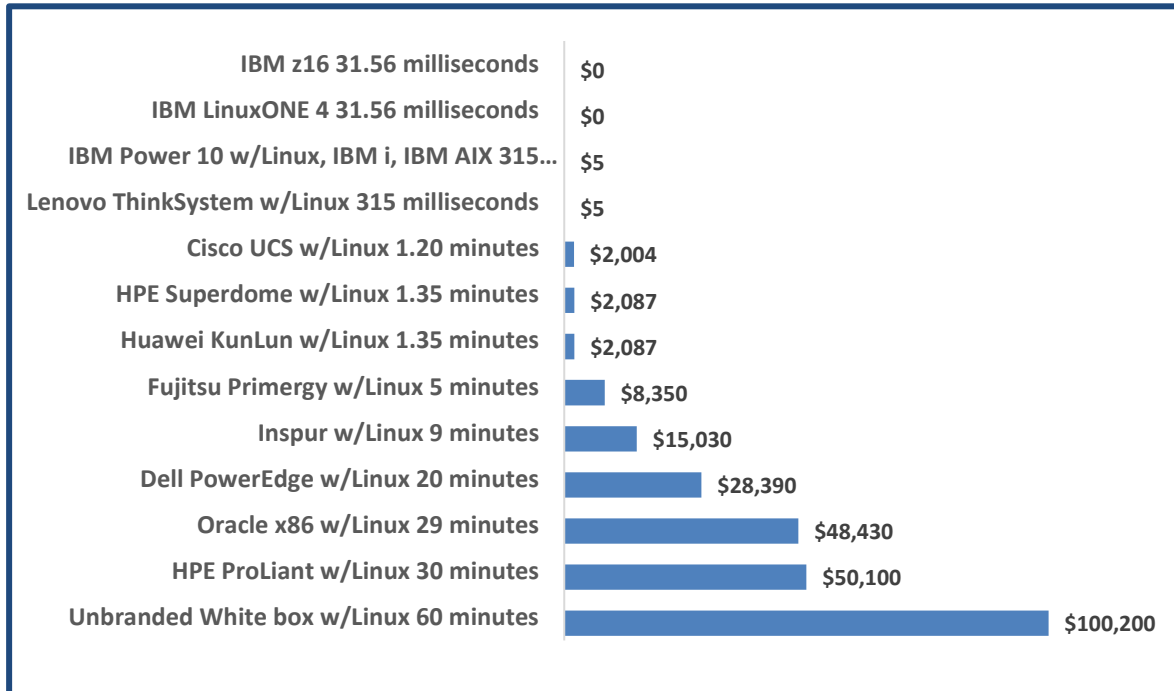
As **Exhibits 3, 4, 5 and 6** below illustrate, there is a substantial gap in the annual unplanned downtime cost comparisons among the top performers and the least stable hardware.

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 2025 Global Server Hardware, Server OS Reliability Survey

Exhibit 4 depicts the downtime YoY cost comparison statistics between ITIC’s 2024 and 2025 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. 2024 vs. 2025 Downtime Per Server/Per Minute Assuming Hourly Cost of \$100K



Source: ITIC 2025 Global Server Hardware, Server OS Reliability Survey

The pivotal role server reliability plays in cost effectiveness is evident in the varying costs associated 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 2023.

In the latest 2025 ITIC Reliability poll, *all* Lenovo ThinkSystem Server distributions improved to an average of between “six and seven nines” of uptime; this equates to 31.5 seconds of unplanned per server annual downtime.

However, the high-end mission critical Lenovo servers, achieved between “seven nines” and “eight nines” of reliability – a scant 315 milliseconds of per server annual downtime, thus drawing even with IBM’s Power10 and Power11 distributions. Assuming an hourly cost of downtime of \$100,000, a business would potentially incur remediation costs of just \$5.26 per server yearly. In 2023 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. The improvement in Lenovo’s uptime statistics from 2023 to 2025 tangibly lower Lenovo customers TCO and delivers near-immediate ROI.

In 2025, Inspur, hardware recorded 9 minutes of per server/per annum downtime – the monetary equivalent of \$16,700. In 2025 Inspur maintained its yearly 2024 one-minute improvement compared to ITIC’s 2023 Reliability survey results. This enables Inspur corporate users to realize potential cost savings of \$1,670 per server. Inspur reliability costs are 5x less than unbranded White box servers. In ITIC’s latest 2025 Reliability Study, unbranded White box recorded per server downtime of 60 to 63 minutes depending on the age of the servers. While that doesn’t sound like much; calculated at an hourly downtime rate of \$100,000, that extra minute of unplanned downtime raises the potential cost to White box server users to \$100,200 in 2024. That’s an increase of \$1,670 per server from the \$98,530 per server reported in 2023.

And as **Exhibit 4** illustrates, Dell’s popular line of PowerEdge servers averaged 20 minutes of unplanned downtime, shaving four (4) minutes off the 24 minutes of per server, per minute yearly downtime versus ITIC’s 2023 Reliability survey results. In 2024, Dell PowerEdge customers have the potential to spend \$28,390 versus \$40,080 they may have had to spend for 24 minutes of per server unplanned downtime a year ago. That four minutes of additional uptime and availability nets Dell customers a potential cost savings of \$11,690 per server annually.

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

For instance, even older versions like the 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 \$15,025 per server annually versus an Inspur machine. In 2024 Inspur servers experienced an average of nine (9) minutes of per server unplanned annual downtime at a cost of \$15,030 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 2024 Reliability achieved an average of “four nines and five nines” reliability – 99.99% and 99.999% per server reliability across the board. To reiterate, is the current and acceptable industry standard.

Enterprises continue to be extremely risk averse minimal patience or tolerance for downtime.

Heightened risk aversion is compounded by the expanding use of AI and Generative AI projects which are experiencing an uptick in mainstream adoption.

Security issues and data breaches are an ever-present threat. The hacks themselves are more pervasive, more pernicious, and targeted. That will not change.

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

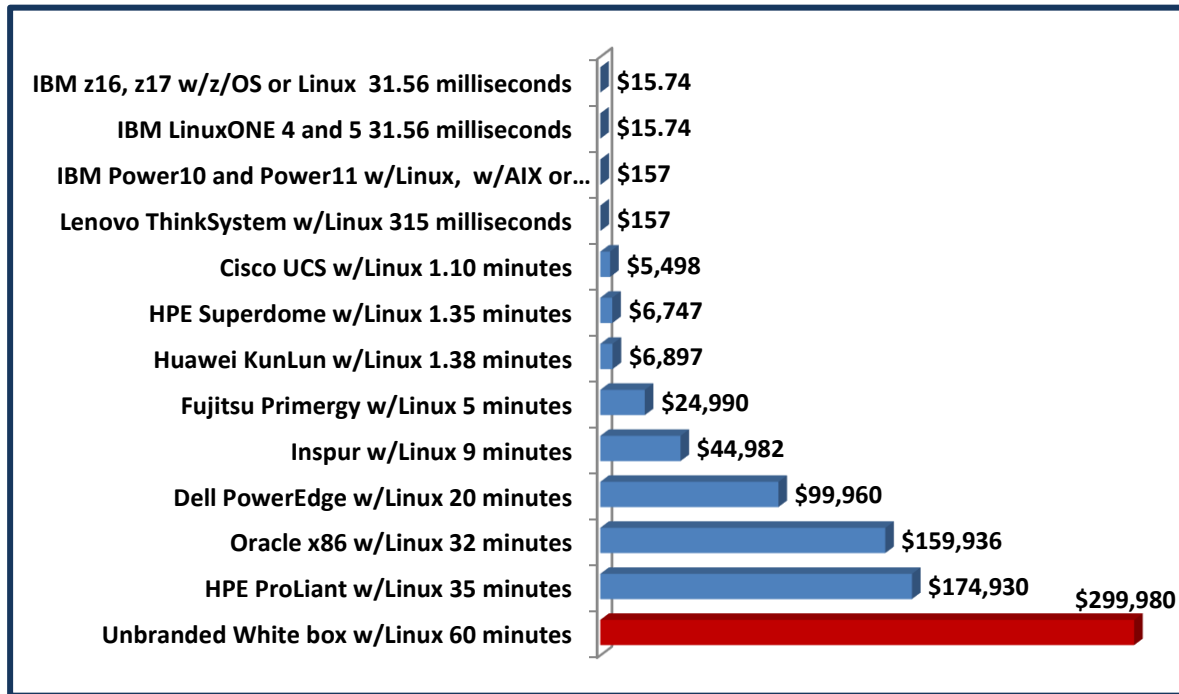
and private), the network edge and remote offices. Many corporations today have virtualized server farms in their on-premises data centers; all cloud computing environments are virtualized. A corporation that experienced one minute of downtime involving a single server running three or four mission critical 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, IBM, Lenovo, Cisco, HPE and Huawei hardware deliver the greatest economies of scale based on their high reliability and availability.

The IBM z14, z15, z16 and z17 along with the IBM LinuxONE distributions deliver several orders of magnitude greater cost savings. For instance, IBM Z and IBM LinuxONE 4 and the newest and LinuxONE 5 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 2025 Global Server Hardware, Server OS Reliability Survey

The legacy IBM Power8 and Power9 per server, per minute reliability rate of approximately “seven nines” – **3.15 seconds** of annual per server unplanned downtime while the IBM Power10 and Power11 distributions achieved “eight nines” of reliability. This could potentially cost businesses \$157 when assuming an hourly outage rate of \$300,000.

In contrast, Dell PowerEdge server survey respondents reported an average of 20 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. Dell enterprise shops improved uptime by four minutes per server in 2024 to 20 minutes of per server annual downtime compared with 24 minutes per machine in 2023. This lowered potential per server outage costs from \$119,952 in 2024 assuming an hourly outage cost of \$300,000 to \$99,960 in ITIC’s 2025 Global Server Hardware, Server OS Reliability study – a significant cost savings of \$19,992.

However, this is still far more expensive than the economies of scale that the IBM Z, IBM Power Systems and Lenovo ThinkSystem servers deliver.

The Lenovo ThinkSystem servers likewise continue to deliver the best cost savings and economies of scale among all x86 based server distributions for the 12th straight year. Lenovo’s high-end and fully loaded mission critical servers are now in a virtual tie with IBM Power10

distributions. 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. In 2025 the Lenovo ThinkSystem servers provided their enterprise customers with an average annual cost savings of \$99,803 per server, per minute unplanned yearly outage costs 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 several minutes of unplanned downtime can save enterprises substantial sums and mitigate risk.

Large enterprises with over 1,000 employees comprised 60% of ITIC’s 2024 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 dependable 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 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, 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 an average of “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, assuming an hourly downtime cost of \$1 million (USD).

IBM Power8 and Power9 averaged seven nines or 3.15 seconds of unplanned downtime while the Power10 and Power11 servers achieved “eight nines” or 315 milliseconds of per server annual downtime. Assuming hourly downtime losses of \$1 million, this equals \$5.26 per server, per minute downtime.

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



Source: ITIC 2025 Global Server Hardware, Server OS Reliability Survey

Hourly Downtime Costs: 51% Estimate Losses Exceed \$1M

ITIC's 2025 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 fourth 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 with fixing the actual technical issues and business problems that caused server failures as well as the time to restore 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 and ongoing supply chain shortages, in 2025 and continuing unabated into 2026, organizations of all sizes, spanning all vertical markets, reported an uptick in their firms' hourly downtime costs compared to 2024. Additionally, 51% of firms now say a single hour of downtime in 2025-2026 now costs their businesses one million (\$1 Million). Some 23% of survey respondents indicated their 2025 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 costs, civil or criminal penalties, fines, or voluntary goodwill gestures that corporations may incur.

Exhibit 7. Hourly Cost of Server Downtime Tops \$1 Million for 51% of Enterprises



Source: ITIC 2025 Cost of Hourly Downtime 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:

- **Global Trade Wars and Tariffs.**
- **Supply chain disruptions continue.** 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** impacts 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 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 end in sight.

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, Human Error & AI Deployments are Top Causes of Downtime

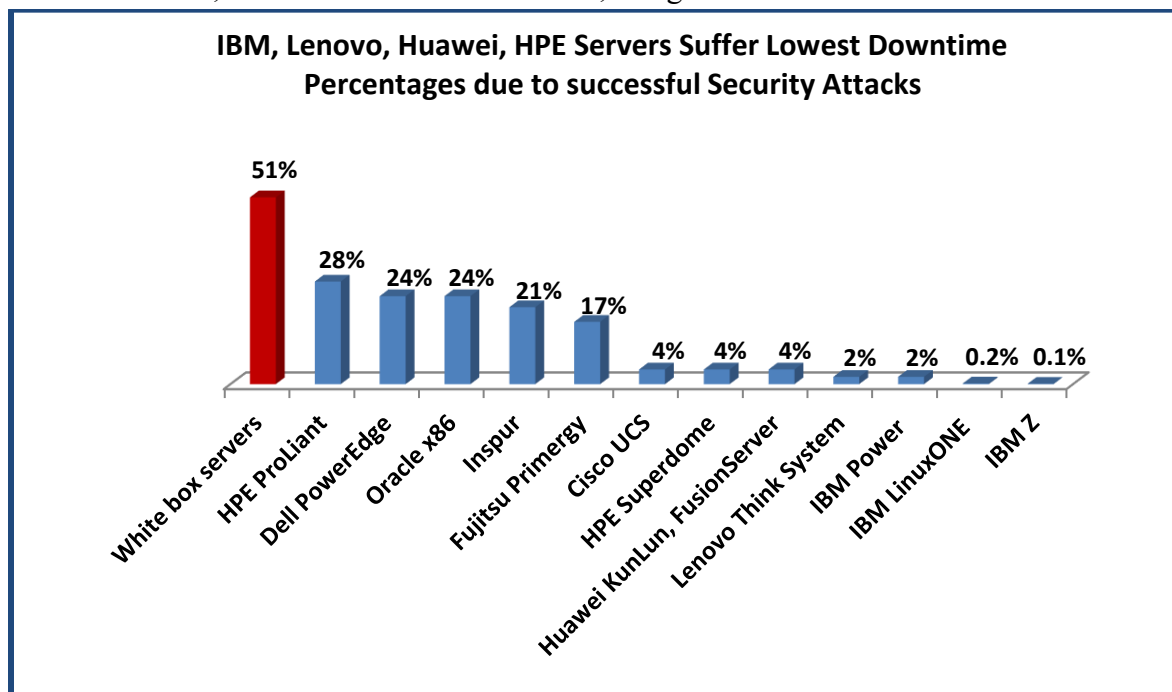
Security hacks and data breaches continue to surge, increasing by 58% over the last 20 months, according to ITIC's latest survey data. Those polled also said security issues constitute the most serious threat that can undermine the reliability and stability of servers throughout the entire corporate ecosystem – in datacenters, at the network edge and in public, private and hybrid clouds.

Complex AI deployments are also impacting uptime and reliability as software engineers and IT departments grapple with provisioning new applications and resolving incompatibility and integration issues. Some 44% of survey respondents noted that complexity and learning curve of their new AI projects were impacting uptime and lengthening the time from testing to full deployment in production environments.

Human error, which is closely linked to security outages, was cited by 76% of survey participants as a major cause of server downtime, followed by 66% 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 2025 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 quickly. ITIC’s latest 2025 Reliability Survey found that 61% of businesses say they strive for “five nines” or 99.999% uptime. This is an increase of 21 percentage points compared to ITIC’s 2024 Global Server Hardware, Server OS Reliability poll in which 40% of respondents said their firms require 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 most exceptionally reliable Lenovo ThinkSystem and 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 Lenovo and IBM machines. 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 that use less expensive, commodity Dell, HPE ProLiant, Oracle and unbranded White box servers are more likely to hire IT entry level IT managers with one-to-three years' experience.

The IBM, Lenovo, Huawei and HPE server organizations are price conscious, but not price driven in the same way that SMBs and midsized 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 Agentic 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: 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 manage 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 extreme challenges to overall server and network reliability. Hackers are extremely well organized and the data breaches and hacks themselves are ever more pernicious, sophisticated, and effective.

For example, the [CloudStrike 2025 Global Threat Report](#) by CloudStrike Holdings, Inc. an Austin, Texas cyber security technology firm revealed that “the breakout time — how long it takes for an adversary to start moving laterally across your network — reached an all-time low in the past year: The average fell to 48 minutes, and the fastest breakout time we observed dropped to a mere 51 seconds.”

The latest 2025 CrowdStrike report also highlighted the security threats related to Generative AI, noting, that “GenAI played a pivotal role in sophisticated cyberattack campaigns in 2024. It enabled FAMOUS CHOLLIMA to create highly convincing fake IT job candidates that infiltrated victim organizations, and it helped China-, Russia-, and Iran-affiliated threat actors conduct AI-driven disinformation and influence operations to disrupt elections.”

The ITIC 2025 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. Security is big business, so much so that it constitutes its own market segment. The hackers are sophisticated and organized.

- **Data Breaches Continue to Surge unabated.** An 89% of enterprises said targeted security issues including ransomware, Phishing, CEO fraud and assorted Email scams – increased by 10 percentage points during the last 12 to 18 months. This is an increase of 25 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 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. However, newer, and even more significant remain and persist owing ongoing geo-political tensions. For example, [Deloitte’s 2023 Semiconductor Industry Outlook](#)¹ report said the war in Ukraine disrupted and continues to disrupt supply chains, access to important raw materials, and energy prices worldwide and energy prices worldwide—and especially in Europe. The war in Ukraine is now entering its fourth year and these issues persist with no end in sight. Likewise, [Deloitte’s 2025 Semiconductor Industry Outlook](#)² predicts that countries and chip vendors face numerous, ongoing and severe challenges in building “resilient supply chains amid geopolitical tensions.” These include export restrictions on advanced chips; the fact that over 100 new entities - primarily Chinese firms were added to the restricted entity list. Most recently the Trump administration’s decision to impose tariffs on a wide array of goods from China, Mexico, Canada, and the European Union further complicate the situation with most of those countries responding with retaliatory tariffs of their own against U.S. goods. Deloitte’s 2025 Semiconductor Industry Outlook noted, “...Given the global nature of most semi supply chains, the proposed new AI-related chip-export controls (by the outgoing administration) and the planned higher tariffs...could make supply chains more complex to administer, shifting profits, costs, and more. And the impact could be felt across the supply chain, including R&D and manufacturing—as well as affecting how industry policies are shaped across countries and regions.”
- IBM, Lenovo, Huawei, HPE and Cisco servers deliver the 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 2024 survey, ITIC’s latest 2025 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 Emperor 4 and LinuxONE 5 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

¹ 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>

² Deloitte 2025 Semiconductor Industry Outlook, February 2025. URL: <https://www2.deloitte.com/us/en/insights/industry/technology/technology-media-telecom-outlooks/semiconductor-industry-outlook.html>

minutes. Among the other server distributions: 82% of Lenovo; 79% 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 two percent of IBM Power and Lenovo ThinkSystem users reported their systems were successfully penetrated and compromised. And only four percent of Huawei KunLun and HPE Integrity Superdome servers suffered a security breach. Once again, unbranded White box servers had the highest percentage – 51% of successful security hacks and data breaches; consequently, 44% of White box servers also experienced downtime due to the attack.

- **IBM, Lenovo and Huawei KunLun Servers Lowest Percentage of Hardware Failures:** As in the prior ITIC 2023 and 2024 Reliability Surveys, the latest 2025 polling statistics reaffirm 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, z16 and newest z17***³ 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 80% of servers over four (4) years old that have not been retrofitted or upgraded to accommodate higher workloads.
- **Planned Downtime Increases:** Some 74% of enterprises now spend from two to eight hours monthly on planned downtime. Most of the planned downtime is attributable to applying security patches, resolving incompatibilities among various hardware and applications and testing, and provisioning new applications.

Conclusions

ITIC’s 2025 Global Server Hardware and Server OS Reliability Survey found that for the 17th straight year, the IBM Z mainframe remains the undisputed and unsurpassed reliability server leader. The IBM z16 and newest z17 maintained their best-in-class status in every category across the board: reliability, security, scalability, performance, sustainability, and manageability.

³ NOTE: At the time of publication, the IBM z17 had only been shipping for six months. The full year statistics will not be available until July 2026.

The IBM Z servers remain in a class of their own. A 96% majority of survey participants said z13, z14, and z15 eight nines -99.999999% - of continuous fault tolerant reliability. IBM z16, the newest z17 and LinuxONE 5 enterprises, said they experienced no discernible downtime and the associated outage costs were likewise almost undetectable to the corporate bottom line.

All IBM Power distributions met and exceeded their best reliability metrics – with 95% of the Power platforms averaging six and seven nines of server uptime. The more recent IBM Power10 and Power11 hardware achieved eight nines – 315 milliseconds of unplanned per server annual downtime.

For the second consecutive year, mission critical Lenovo ThinkSystem servers are on par with their IBM Power rivals. Over 90% The high-end Lenovo servers recorded, “eight nines” – 99.999999% of uptime. Overall, every version of Linux ThinkSystem servers demonstrated superlative reliability by exceeding their year-over-year uptime scores.

The Lenovo ThinkSystem servers ranked as the most reliable x86-based hardware platform for the 12th straight year averaging between six nines or 31.5 seconds of per server unplanned annual downtime, seven nines or 3.15 seconds of unanticipated per server yearly downtime for all ThinkSystem distributions and eight nines of uptime for the latest mission critical systems.

The Cisco UCS servers registered 1.12 minutes of unplanned per server downtime, followed closely by the HPE Superdome and Huawei KunLun servers with 1.35 minutes and 1.38 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.

IBM and Lenovo both 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, IBM Power Systems, 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 will remain cost conscious and extremely risk averse for the foreseeable future. They depend on highly dependable, robust, and secure servers, server OS and application software to conduct uninterrupted 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 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. NO system is 100% secure. The reliability of any system can be undone by human error.

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

Unreliable systems raise 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 dependable, robust, and secure servers and server OS software.

Recommendations

An organization's business and fiscal success depend 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 collaborate and cooperate to attain these continuous reliability goals. Server hardware and server OS, AI, virtualization, security and cloud vendors and their clients bear equal responsibility to attain and maintain 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:

- **Take Inventory. 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 employee 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 pay close attention to whether the company's reputation suffered because 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 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 comprehensive downtime record and calculate 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 the 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 continue 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 21st 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 components, 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 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 resolved and the time to restore full operations.

Survey Methodology

ITIC's *2025 Global Server Hardware and Server OS Reliability Survey* polled 2,000 C-level executives, IT, and security administrators in corporations worldwide from February through November 2025. The independent Web-based survey included multiple choice questions and one Essay question. To maintain objectivity, ITIC accepted no vendor sponsorship. No participants received any remuneration. ITIC analysts also conducted two dozen first person customer interviews to obtain actual anecdotal data based on actual enterprise production data. This provides deeper context regarding the internal and external issues that directly impact the economics and efficiency of daily business operations and influence deployment 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 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 came from 41 vertical markets. Approximately 58% of respondents hailed from North America and 42% were international customers from more than 30 countries throughout Asia Pacific, Africa, Australia, New Zealand, Europe, Latin and South America.

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-2024-hourly-cost-of-downtime-report/>

<https://itic-corp.com/itic-2023-reliability-survey-ibm-z-results/>

<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/>