INFORMATION TECHNOLOGY
INTELLIGENCE CONSULTING

ITIC 2022 Global Server Hardware, Server OS Reliability Report

August 2022

Table of Contents

Executive Summary

The IBM Z and IBM Power Systems continue to dominate, delivering the best server reliability, uptime and security for the 14th straight year.

Lenovo's ThinkSystem servers provide the top reliability and security among all x86 server distributions for the ninth straight year.

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

IBM Z and IBM Power Systems deliver over 40x more uptime than least efficient "White box" platforms and 60x lower Total Cost of Ownership (TCO). The Lenovo ThinkSystem, Huawei KunLun and HPE Superdome (in that order) delivered the highest reliability among x86 platforms.

Over three-quarters of businesses – 78% – cite security as the top cause of unplanned downtime and 64% said human error causes unplanned outages.

Mission critical server and server OS distributions from IBM, Lenovo, Hewlett-Packard Enterprise (HPE), Huawei and Cisco continue to deliver the highest levels of inherent reliability and availability among 18 different server platforms despite a continuing spike in security hacks, increasing ecosystem complexities and ongoing supply chain challenges.

For the 14th consecutive year, the IBM Z, the LinuxONE III and the IBM Power Systems remained the preeminent server platforms posting the best across-the-board reliability ratings among 18 mainstream distributions. Some 96% of IBM Z mainframes and LinuxONE III server customers recorded seven nines (99.99999%) of true fault tolerant reliability and availability. The IBM Z, and LinuxONE III recorded a near-imperceptible 0.0043 minutes of per server unplanned monthly outages or just 3.15 seconds of unplanned per server downtime **annually (See Table 1).** This was followed by 93% of IBM Power Systems clients said the IBM systems achieved five and six nines of system reliability and availability **(See Exhibit 1)**. The IBM Power8, Power9 and Power10 servers posted just one (1) minute each of unplanned per server monthly downtime.

The Lenovo ThinkSystem servers followed closely and posted the highest levels of reliability among all x86 hardware distributions for the eighth consecutive year. A 92% majority of Lenovo servers attained five and six nines of reliability, posting just over one minute – 1.10 - of unplanned per server monthly downtime. The Huawei KunLun and Fusion servers, the HPE Superdome and the Cisco UCS hardware (in that order), rounded out the top five most reliable server platforms.

Those are the results of the ITIC 2022 Global Server Hardware, Server OS Reliability independent Web-based survey. It polled 1,550 corporations across 30 vertical market segments worldwide on the reliability, performance and security of the leading mainstream on-premises and cloud-based servers from January through July 2022. In order to maintain objectivity, ITIC accepted no vendor sponsorship.

Exhibit 1. Unplanned Monthly Per Server Downtime in Minutes

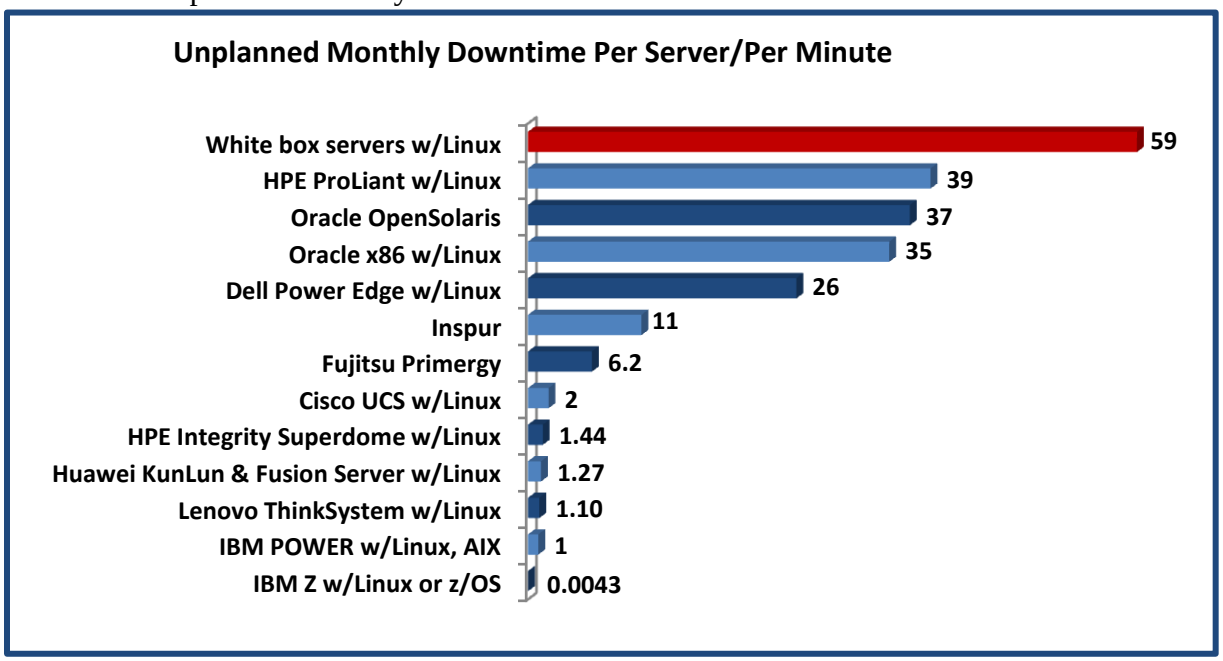

Source: ITIC 2022 Global Server Hardware, Server OS Reliability Survey

The increased server and server operating system uptime and availability enabled the IBM, Lenovo, Huawei, HPE and Cisco servers (in that order) to deliver, the most economical Total Cost of Ownership (TCO) among all mainstream distributions in datacenters, at the network edge and in hybrid cloud environments.

The Lenovo ThinkSystem servers likewise improved their uptime and availability recording the best reliability among all x86 servers – a scant 1.10 minutes of per server unplanned monthly outages. The Huawei KunLun and Fusion platforms also improved uptimes with 1.27 minutes each of unplanned per server outage, along with the HPE Superdome platform which averaged 1.44 minutes of unanticipated per server downtime. Cisco's UCS servers also hung tough. Cisco servers frequently are installed at the network edge/perimeter, which is often the first line of attack. The Cisco UCS servers registered two (2) minutes of monthly unplanned per server downtime.

The top server reliability vendors – led by IBM, Lenovo, HPE and Huawei – also delivered the strongest server security, experiencing the fewest number of successful data breaches and the least amount of downtime due to security-related incidents.

Introduction

ITIC's 2022 Global Server Hardware, Server OS Reliability Survey examines the profound and lingering affects the global pandemic continues to exact on the reliability and security of server hardware, server OS and mission critical applications and its impact on customers. It also delves into new issues impacting the reliability of critical server infrastructure and applications. These include cloud computing and the ongoing data deluge and data sprawl.

The reliability and security of server hardware, server operating systems and mission critical applications are critical elements of the core datacenter, network edge and cloud infrastructure.

Server hardware reliability is essential in the Digital Age of "always on" and interconnected networks. Downtime of even a few minutes impedes productivity; interrupts daily operations and negatively impacts revenue. Over 90% of corporations now require a minimum of "four nines" (99.99%) reliability and nearly 40% of organizations now strive for "five nines" (99.999%) uptime or higher.

There is a significant difference between four and five nines of uptime. Four nines of reliability are equal to 52.56 minutes of unplanned **annual** per server downtime. In contrast, five nines of uptime is the equivalent of just 5.26 minutes of unplanned, **annual** per server downtime. Organizations are risk averse. They rely on servers, applications, devices and networks to transact business 24 x 7 irrespective of worldwide location. There is no tolerance for downtime. Corporate enterprises need the most reliable, robust servers to ensure continuous, uninterrupted secure data access; maintain regulatory compliance and mitigate risk to an acceptable level.

In the digital era of interconnected intelligent systems and networks, unplanned downtime of even a few minutes is expensive and disruptive and can reverberate across the entire ecosystem. This includes datacenters; virtualized public, private and hybrid clouds; remote work and learning environments and the intelligent network edge.

ITIC's 2022 Hourly Cost of Downtime survey indicates a single hour of server downtime can result in potential losses of $300,000 or more for 91% percent of mid-sized enterprises (SMEs) and large enterprises during a single hour. And among that 91% majority, nearly half or 44% – of corporations said hourly outage costs exceed one million ($1M) to over five million ($5M).

To reiterate, the most reliable server distributions: IBM, Lenovo, Huawei, HPE and Cisco successfully maintained the best ever uptime scores they posted in ITIC's prior reliability studies

over the past 20 months. This is despite the business and technology challenges precipitated by the pandemic. They include security hacks and data breaches; increased IT management challenges associated with teleworking and remote learning and supply chain disruptions.

Other mainstream server distributions, like Dell, Inspur, HPE ProLiant and unbranded white box server customers, either maintained the same unplanned outage levels from ITIC's 2021 Global Reliability survey or recorded slight, one-to-two minute per server monthly increases in additional server downtime during the past 12 to 18 months. These seemingly slight increases in unplanned downtime can have a significant and negative impact on the corporate bottom line. Any additional downtime can disrupt daily operations and they can be costly, particularly when outages involve multiple servers in on-premises datacenters as well as those running at the network edge and in virtualized cloud environments.

Since 2008, the ITIC Global Server Hardware, Server OS Reliability Report has compared the reliability of the top mainstream server platforms. It also examines the internal and external issues that positively impact or undermine core server hardware, server operating system and application reliability, based on key metrics and business practices.

The ITIC 2022 Global Server Hardware, Server OS Reliability survey calculated specific per minute, per server unplanned monetary downtime costs ranging from $10,000 to $10 million in various server quantities from a single server to one thousand (1,000) servers **(See Table 2)**.

As always, ITIC's 2022 Global Server Hardware and Server OS Reliability Report utilizes information gathered from previous ITIC surveys to compare, contrast and analyze the reliability of the various platforms and to track current and future trends. ITIC's 2022 poll incorporated additional security questions into the main survey. ITIC also fielded a standalone separate survey: the ITIC 2022 Global Server Hardware, Server OS Security Report.

The global pandemic and the ensuing supply chain disruptions, the sharp spike in data breaches have all combined to precipitate shifts in the ways organizations deploy their foundational IT architectures across the network ecosystem. These changes include: a surge in demand for daily Ebusiness transactions; a hybrid work environment with increased teleworking; remote learning; interactive virtual collaboration and audio and video streaming. Business is conducted virtually; in datacenters; at the network edge or perimeter; in clouds and in homes on BYOD and mobile devices.

These trends are the "new normal." The infrastructure expansion and shift towards public, private and hybrid clouds and intelligent edge computing will continue unabated. Organizations will continue to pursue hybrid office and remote working and remote learning models.

Enterprises must be able to ascertain the server reliability and calculate the associated costs.

Reliability/Uptime by the Numbers

Organizations measure server and application reliability percentages in "nines." There is an order of magnitude difference of server and application reliability and uptime between each additional "nine." Four nines – 99.99% – reliability equals 52.56 minutes of unplanned per server/per annum downtime or 4.32 minutes of per server monthly unplanned downtime **(See Table 1)**. By contrast, five nines – 99.999% – is the equivalent of 5.26 minutes of unplanned per server/per annum and just 25.9 seconds of monthly unplanned system downtime. The highly sought after continuous uptime and availability levels of six nines equals a near-imperceptible 2.59 seconds of per server unplanned monthly downtime and just 31.5 seconds of yearly system downtime.

Table1 below depicts the availability percentages and the equivalent number of annual, monthly and weekly hours and minutes of per server/per annum downtime. It illustrates the business and monetary impact on operations. ITIC publishes this table in every one of its Global Server Hardware, Server OS Reliability reports. It serves as a useful reference guide to enable organizations to calculate downtime and determine their levels of server uptime.

Table 2 details the monetary costs associated with specific hourly downtimes ranging from $100,000 to $10 million (USD) based on per minute costs from a single server to one thousand servers. Taken together, the two tables paint a clear and compelling overview of server performance and how quickly and significantly monetary downtime costs can escalate.

Table 1: Reliability/Uptime by the Numbers

Reliability %	Downtime per year	Downtime per month	Downtime per week
90% (one nine)	36.5 days	72 hours	16.8 hours
95%	18.25 days	36 hours	8.4 hours
97%	10.96 days	21.6 hours	5.04 hours
98%	7.30 days	14.4 hours	3.36 hours
99% (two nines)	3.65 days	7.20 hours	1.68 hours
99.5%	1.83 days	3.60 hours	50.4 minutes
99.8%	17.52 hours	86.23 minutes	20.16 minutes
99.9% (three nines)	8.76 hours	43.8 minutes	10.1 minutes
99.95%	4.38 hours	21.56 minutes	5.04 minutes
99.99% (four nines)	52.56 minutes	4.32 minutes	1.01 minutes
99.999% (five nines)	5.26 minutes	25.9 seconds	6.05 seconds
99.9999% (six nines)	31.5 seconds	2.59 seconds	0.605 seconds
99.99999% (seven nines)	3.15 seconds	0.259 seconds	0.0605 seconds

Source: ITIC 2022 Global Server Hardware, Server OS Reliability Survey

The aforementioned metrics clearly underscore that the IBM z14 and z15; along with the LinuxONE III platform continue to maintain continuous levels of reliability, with just 0.0043 minutes of unplanned monthly per server downtime; this equates to just 3.15 seconds of unplanned per server annual downtime. They were followed closely by the IBM Power8, Power9 and Power10 with one (1) minute of per server unplanned monthly downtime and the Lenovo x86-based ThinkSystem with 1.10 minutes of per server unplanned downtime each month. In practical terms, this means there is minimal or imperceptible impact on daily business operations, end user productivity and corporate revenue.

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 2022 Global Server Hardware, Server OS Reliability Survey

In 2022 a price tag of $100,000 (USD) for one hour of downtime for a single server is extremely conservative for all but the smallest micro SMBs with one to 25 employees. It equates to $1,670 per minute/per server. Hourly cost of downtime calculated at $300,000 equals about $5,000 per server/per minute. The cost of a more severe or protracted hourly outage that a business estimated at $1 million (USD) is the equivalent of $16,700 per server/per minute.

ITIC's 2022 Global Server Hardware and Server OS Reliability Survey found that 91% 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, 44% of those polled indicated that hourly downtime costs now exceed $1 million. Since 2021, only one (1%) percent of respondents said a single hour of downtime costs them $100,000 or less. Nine percent (9%) of respondents valued hourly downtime at $101,000 to $300,000.

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.

The aforementioned ITIC Hourly Downtime monetary figures represent only the costs associated with remediating the actual technical issues and business problems that caused the server or OS to fail. They 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).

Overview: Top Survey Findings

The IBM Z; IBM LinuxONE; IBM Power; Lenovo ThinkSystem; Huawei KunLun and Fusion and HPE Superdome server distributions (in that order) took the top spots 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).

Among the most notable survey findings:

- **Server Reliability:** IBM z14 and z15 and the IBM LinuxONE III outpaced all rivals; both posted their best ever reliability results: just 0.0043 minutes of per server monthly unplanned downtime. The reliability of the IBM Power servers also posted notable uptime improvements, experiencing just one (1) minute of unplanned per server monthly downtime. The Lenovo ThinkSystem hardware averaged 1.10 minutes of per server unplanned monthly outages, followed by the Huawei KunLun hardware with 1.27 minutes of unplanned per server outages and the HPE Superdome distributions experiencing 1.44 minutes of per server monthly downtime. Inspur was in the middle of the pack with 11 minutes of unplanned per server downtime, while the Dell PowerEdge servers again recorded 26 minutes of unplanned outages. Unbranded White box servers again were the least reliable servers with 59 minutes of unplanned per server downtime; this is up two (2) minutes from ITIC's previous 2021 poll.
- **Server Availability:** The IBM Z servers are in a class by themselves: a 96% majority of IBM Z customers said their businesses achieved unparalleled fault tolerant levels of seven nines or better - 99.99999%+ reliability and continuous availability. IBM's Power Systems are close behind with 93% of customers reporting that the Power9 and latest Power10 models deliver a minimum of five and six nines availability/uptime. Meanwhile, 92% of Lenovo ThinkSystem; 90% of Huawei KunLun and HPE Superdome enterprises, respectively said their businesses achieved a minimum of five and six nines server availability.
- **Cost Effectiveness/Total Cost of Ownership:** The most reliable IBM z14 and z15; the IBM LinuxONE III and the Power8, Power9 and Power10 servers deliver the best TCO and near immediate Return on Investment (ROI). A single minute of per server unplanned monthly downtime on IBM z14 and z15 servers and the LinuxONE III solution, calculated at a rate of $100,000, costs enterprise customers just $7.19 for one minute of unplanned per server monthly downtime. One minute of unplanned downtime on a single IBM Power8, Power9 and Power10 calculated at $100,000 hourly, can cost $1,670. The latest Power10, which began shipping in September 2021, has delivered tangible reliability improvements in just under one year of mainstream production deployment; in addition to improving reliability, it also delivers approximately 30% better performance, manageability and security, lowers TCO and accelerates ROI.
- **The Lenovo ThinkSystem servers** provide the highest uptime and availability among more than one dozen x86 server distributions. The Lenovo ThinkSystem servers averaged 1.10 minutes of unplanned per server monthly downtime. Assuming hourly downtime losses of $100,000, one minute of per server/per minute downtime on a Lenovo ThinkSystem server could cost companies $1,837 per minute.

- **The Huawei KunLun and Fusion** mission critical servers clocked a per server monthly average downtime rate of 1.27 minutes of unplanned per server outages. That equates to per server/per minute downtime charges of $2,120.
- **HPE's high end Superdome** servers also registered high reliability: with customers reporting 1.44 minutes of per server unplanned monthly downtime. Enterprises that estimate a single of downtime calculated at $100,000 would incur potential monetary per server/per minute losses of $2,404.
- **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 an increase of two (2) minutes per server compared to ITIC's 2021 Reliability survey. That has the potential to cost corporations $98,530 when hourly downtime losses are calculated at $100,000.

As **Exhibit 2** illustrates the most reliable servers: the IBM Z, IBM LinuxONE III; IBM Power; Lenovo ThinkSystem, Huawei KunLun and HPE Superdome (in that order) experienced the lowest percentages – 0.1% and 1% each respectively of the most severe outages of four hours.

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

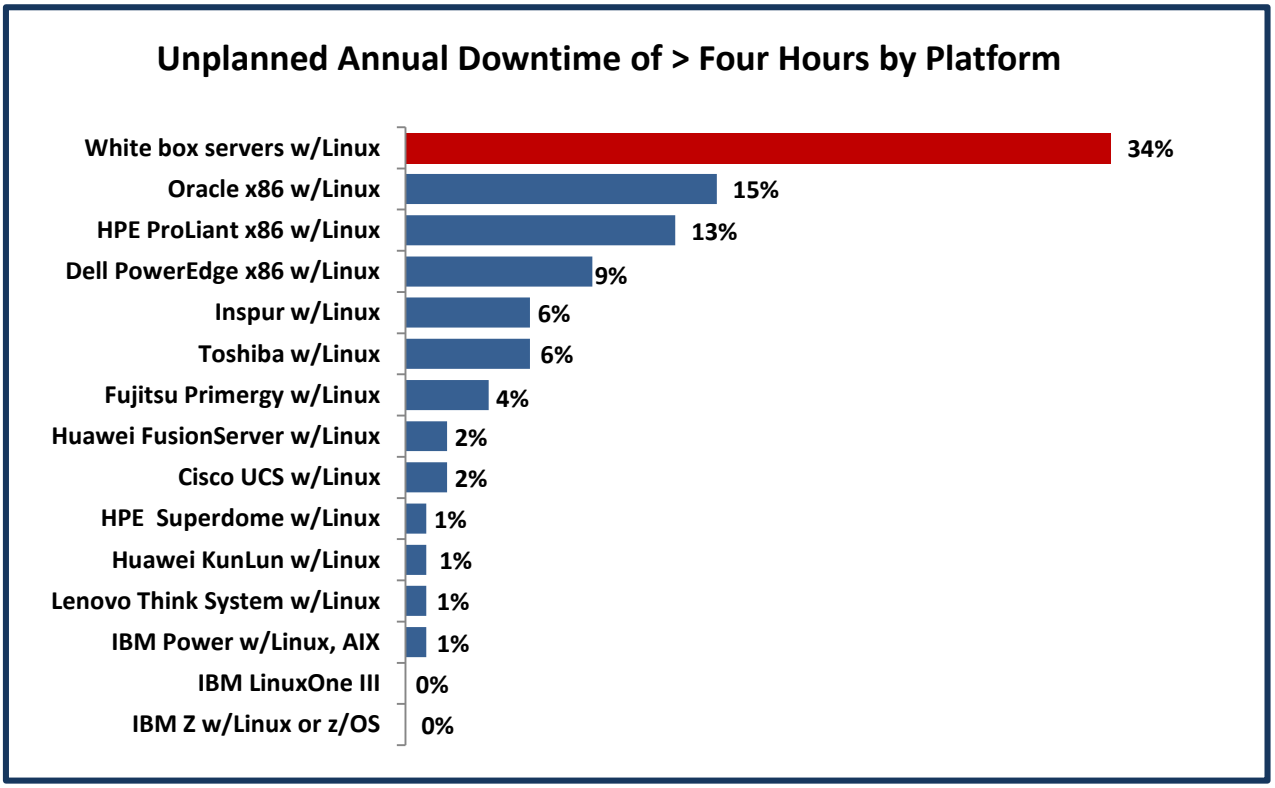

Source: ITIC 2022 Global Server Hardware/Server OS Reliability Survey

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

Server Hardware Vendor Platform Overview

There is no way to overstate the correlation between the server reliability gains of the top vendors like IBM, Lenovo, Huawei and HPE and the release of new, higher performing servers.

IBM for example released the Power10 server in September 2021. The general availability of the Power10 contributed significantly to the additional 49 seconds of per server monthly uptime by the platform's users over the last year.

In February 2022, the Lenovo Infrastructure Solutions Group (ISG) likewise unveiled several new offerings, including new server options and new features for its Lenovo ThinkSystem portfolio.

The latest generations of the top performing server distributions utilize advanced 7nm technology provides shrink down transistors, improvements in silicon on-chip utilization, and improvements in power efficiency and overall greater economies of scale.

IBM Reliability and Security Survey Highlights

- **IBM Z** servers continue to achieve top grades for overall reliability, accessibility, performance, and security among all server platforms. The IBM Z family – the "Z" stands for zero downtime – consistently outperforms **all** competitors in every reliability category and delivers the lowest total cost of ownership (TCO) and fastest return on investment (ROI). The z14 and z15 servers scored the best reliability/uptime, application availability ratings and security across the board in terms of actual minutes of unplanned per server/per annum downtime. The IBM Z mainframe and the IBM LinuxONE distributions both exhibit true fault tolerance and deliver the vaunted seven nines – 99.99999% or greater uptime and availability. The IBM z14, z15 and LinuxONE III offerings now average 0.0043 minutes of ***unplanned*** per server monthly downtime due to server flaws or just 3.15 seconds of per server **annual** unplanned downtime. This is a 58% reduction in unplanned per server monthly downtime compared to the IBM Z and IBM LinuxONE III reliability scores in ITIC's 2019 survey. The improved monthly uptime in turn, lowers the TCO of the IBM Z and LinuxONE from $12.32 per server/per minute in 2019 to $7.18 per server/per minute in the latest ITIC 2022 Global Server Hardware Security study. Overall, the IBM Z registers just seconds of near imperceptible monthly downtime. Equally important, given the ongoing surge in security hacks and data breaches, is the IBM Z server's superlative security. The Z continues to register the lowest percentage – less than one percent – of successful data breaches from January through August 2022. Additionally, IBM Z and LinuxONE III survey respondents also

reported the quickest Mean Time to Detection (MTTD) with 96% of ITIC enterprise respondents stating their security and IT administrators were able to detect and shut down hacks on these platforms. The new IBM z16 industry-first quantum-safe system is the latest in a long history of security innovations on IBM Z. Singularly and collectively, these results underscore the success of the Z and LinuxONE III offerings. The platforms have also been bolstered by IBM's 2019 acquisition of Red Hat for $34 billion. Industry pundits no longer question the rationale or efficacy of the Red Hat acquisition. It has produced a significant increase of Linux workloads on the Z and LinuxONE III platforms. IBM executives publicly stated the company has seen a 55% increase in Linux MIPS. IBM also said that 92 of IBM's top 100 Z clients run Linux workloads.

- **IBM's LinuxONE III** is based on the IBM Z platform. The LinuxONE server is provisioned entirely for Linux workloads. The LinuxONE, in combination with robust secure service containers, is a key component of IBM's blockchain offerings, IBM Cloud's blockchain services and IBM private cloud implementations. The LinuxONE III addresses hybrid cloud environments and utilizes the IBM Z's pervasive encryption. Like the IBM z14 and z15, the LinuxONE III platform also incorporates the IBM Hyper Protect Data Controller, which delivers transparent, end-to-end, data-level protection and privacy. The IBM Hyper Protect Data Controller enables corporations to encrypt or mask eligible data, that is, data accessible through Java or REST APIS; grant and revoke access to it and maintain control of it – even if it moves off the system of record**.** The result: IBM LinuxONE III shared the highest security and reliability rankings in ITIC's 2022 poll. The LinuxONE III experienced 0.60 minutes of unplanned per server downtime. On the security front, less than one percent (0.1%) of enterprise users said they suffered a successful hack. And 96% of LinuxONE III enterprises said they detected and shut down data breaches "Immediately or within the first 10 minutes" of the attack.

- **IBM Power8, Power9 and Power10 Servers** recorded just one (1) minute of per server/per month of unplanned downtime. These results are the best evidence that IBM continues to "power up" its Power Systems servers. In ITIC's latest 2022 Reliability poll, a 93% majority of IBM Power customers reported their IT and security administrators were able to detect and thwart attacks "immediately or within the first 10 minutes" of a breach. IBM's Power9 scale-out systems have been out three years. IBM continually refreshes and updates the line – placing particular emphasis on performance, support for mission critical workloads, support for advanced analytics, in-memory databases and embedded security. The latest generation Power10 server – the IBM Power E1080 – shipped in September 2021. It is the first scale up system based on the Power10 processor. IBM Power servers have security built in at all layers in the stack – processor, systems, firmware, OS and Hypervisor. With accelerated encryption built into the chip, data is protected in motion and at rest. IBM claims that its PowerVM hypervisor has no reported security vulnerabilities. Power9 and Power10 servers are cloud-ready and include built-in PowerVM virtualization capabilities. The Power9 and Power10 scale-out servers are designed for integration into organizations' cloud and AI strategies. This provides the high performance and RAS capabilities required to support mission-critical workloads like IBM's Db2 and Oracle databases as well as SAP HANA. The Power10 processor is designed for energy efficiency and performance in a 7nm form factor. IBM

claims Power E1080, the first in a generation of servers based on Power10 processor, delivers 50% more performance compared to Power E980. It also lowers energy consumption by 33% for the same workload. The Power10 also features hardware-enabled security capabilities like transparent memory encryption for end-to-end security. The IBM Power10 processor is engineered to achieve significantly faster encryption performance with quadruple the number of AES encryption engines per core compared to the IBM Power9 for today's most demanding standards and anticipated future cryptographic standards like quantum-safe cryptography and fully homomorphic encryption. It also brings new enhancements to container security.

Lenovo Reliability and Security Survey Highlights

- **Lenovo ThinkSystem** servers had the best reliability rankings among all Intel x86 platforms. The Lenovo ThinkSystem hardware registered just 1.10 minutes of unplanned per server downtime; a reduction of 33% since 2019. This makes the Lenovo ThinkSystem the most reliable and cost efficient x86-based server. The Lenovo ThinkSystem security is also highly rated: it achieved the best MTTD rates among all Intel x86 servers. A 95% majority of survey respondents stating their IT and security administrators detected and shut down attempted hacks and data breaches immediately or within the first 10 minutes of the penetration. In the eight years since Lenovo purchased IBM's x86-based server business and the 13 years since it acquired IBM's line of PCs and notebooks, Lenovo prioritized security with a vengeance. Consequently, Lenovo servers and desktops have gone from strength to strength as the company continually enhances and fortifies the performance, reliability and security of the servers and its desktop PCs and laptops. In March 2022 Lenovo unveiled a suite of edge-to-cloud IT infrastructure solutions optimized to address remote and hybrid work situations. The latest offerings include the single-socket ThinkSystem V2 Servers, which are flexible, energy-efficient, and low-noise solutions designed to allow customers to manage constrained spaces inside. They are also well suited for network edge deployments. At its Lenovo Infrastructure Solutions Group (ISG) Summit in mid-June 2022, Kirk Skaugen, executive vice president of ISG, noted that Lenovo differentiates itself from server competitors via its "full edge-to-cloud portfolio," and by providing more choices and solutions for customers who want greater simplicity. Skaugen also emphasized Lenovo's ability to "bundle and kit" its solutions to best align the technology with the business requirements. Lenovo's technical service and support is also first-rate. Like IBM, Lenovo has constructed and executed an excellent and effective tactical and strategic security strategy. In 2018, Lenovo unveiled the ThinkShield end-to-end security technology for its PCs and laptops. While the advanced ThinkShield technology is primarily focused on desktops and mobile devices, it has had a notable and noticeable impact in helping to secure Lenovo ThinkSystem servers as well. This was particularly true over the last three years following the onset of the global pandemic when security and data breaches soared. Many servers were successful penetrated by attached desktop and mobile devices as many organizations shifted to a remote workforce for workers and students alike. At its Lenovo Infrastructure Solutions Group (ISG) Summit in mid-June 2022, Skaugen, noted

that Lenovo differentiates itself from server competitors via its "full edge-to-cloud portfolio," and by providing more choices and solutions for customers who want greater simplicity. Skaugen also emphasized Lenovo's ability to "bundle and kit" its solutions to best align the technology with the business requirements. In March 2022 Lenovo unveiled a suite of edge-to-cloud IT infrastructure solutions optimized to address remote and hybrid work situations. The latest offerings include the single-socket ThinkSystem V2 Servers, which are flexible, energy-efficient, and low-noise solutions designed to allow customers to manage constrained spaces inside. They are also well suited for network edge deployments. The single-socket ThinkSystem ST50 V2, ST250 V2 and the SR250 V2 servers offer companies simple solutions that are easily tailored for running their business, including support for business-critical applications in retail, manufacturing, and financial services. The line also includes the ThinkSystem SE450, Lenovo's latest AI server for the network edge. It delivers optimal bandwidth, bolsters security and reduces downtime. The SE450 has 10 to 36 server cores with up to 1TB of memory; it supports broad-wired and wireless connectivity and it can continuously operate at temperatures between 5°C and 45°C. Lenovo is also targeting midrange and high end enterprises with its SR630 V2 Rack server which supports data analytics, hybrid cloud, hyperconverged infrastructure, video surveillance and high performance computing. Lenovo's overarching ThinkSystem strategy melds innovation with reliable, flexible, and secure data center systems. This is a savvy move that also has far reaching ramifications for Lenovo's servers, networks and ultimately its corporate customers. Human error is by far the biggest cause of server downtime. Lenovo enforces rigorous security standards, policies and procedures at its manufacturing facilities and global supply chain. Lenovo's Quality Engineers retain the right to audit the company's Trusted Suppliers at any time, giving the company even further control and insight into the security of its devices' components. ThinkShield also delivers design level security. This includes secure BIOS and firmware, as well as privacy screens and laptop camera shutters into its devices to help minimize "visual hacking" when mobile users are in public places. ThinkShield is designed to protect users' identities and credentials, offering FIDO-certified authenticators and integration with Intel Authenticate (offering up to 7 authentication factors). ThinkShield also features BIOS-based Smart USB protection, which functions by configuring USB ports to only respond to keyboards and pointing devices. Lenovo also emphasizes that its open server, storage, networking and system management platforms seamlessly integrate with existing and legacy environments. In first person interviews with ITIC analysts, Lenovo customers cited the ease of deployment and ease of integration and backwards compatibility as contributing to the underlying reliability and stability of the ThinkSystem platform. Lenovo users also lauded the vendor's after-market service and support. Lenovo's system design supports mission-critical databases, enterprise applications, big data analytics, and cloud and virtualized environments. Both these systems incorporate numerous fault-tolerant and high-availability features into a high-density, rack-optimized lid-less package that minimizes the space needed to support "massive network computing operations" and simplify servicing, as the system never needs to be removed from the rack.

Cisco UCS Reliability and Security Survey Highlights

- **Cisco's Unified Computing System (UCS)** continues to score well registering, two (2) minutes of per server monthly unplanned downtime. Since 2019, Cisco UCS servers have shaved monthly unplanned downtime per server, per minute downtime costs by an impressive 65%. This is a noteworthy and welcome accomplishment since many Cisco UCS servers are positioned at the network edge which is on the front line of security attacks. Despite the onslaught, 87% of Cisco UCS survey respondents said they were able to detect, isolate and shut down security hacks immediately or within the first 10 minutes. Cisco UCS survey respondents also reported that the servers experienced five successful security hacks each over the last 12 months. In response to the increase in data breaches, Cisco began publishing the Cisco UCS Hardening Guide. The document is available for free download. It contains detailed information to help users secure Cisco UCS platform devices to improve network security. Structured around the three planes by which the functions of a network device are categorized, this document provides an overview of each Cisco UCS Software feature and references related documentation. Additionally, Cisco introduced a number of management and performance upgrades aimed at improving TCO and accelerating installation and deployment. Cisco claims its UCS will allow an 86% reduction in cabling and allow provisioning in a matter of minutes (rather than days or weeks), while reducing capital expenses by more than 40%. Manufacturers assure users of 100% compatibility between and among components. And load balancing is a non-issue.

Huawei Reliability and Security Survey Highlights

Huawei KunLun and Fusion Servers. Over the last five years Huawei, headquartered in Shenzhen, China has emerged as one of the top five server hardware vendors worldwide with its high end KunLun mission critical server and its general purpose FusionServer x 86-based servers. Based on ITIC's 2022 Global Server Hardware, Server OS Reliability Survey and the ITIC 2022 Global Server Hardware Security Survey, the Huawei KunLun and Fusion Servers are also among the top three most reliable and secure hardware platforms. The Huawei KunLun high end servers experienced just 1.27 minutes of per server unplanned downtime. Nine-in-10 Huawei customers reported they achieved five or six nines uptime. Additionally, a 91% majority of Huawei survey respondents noted their IT and security administrators detected and shut down attempted breaches "Immediately or in under 10 minutes." Huawei survey respondents indicated that the KunLun and Fusion servers each experienced only 1.5 hacks during the last 20 months. Since 2015, Huawei fortified the advanced features, inherent security and overall performance of its servers. To successfully compete with rivals including Cisco, Fujitsu, HPE, IBM, Inspur, Lenovo and others, Huawei's server family includes general purpose rack and blade servers to mission critical hardware to address high performance computing (HPC). Huawei has also imbued its servers with advanced capabilities to support emerging compute intensive applications like AI, Big Data Analytics, Deep Learning and Machine Learning. To successfully compete with rivals like Dell, HPE, IBM, Inspur, Lenovo and others, Huawei's servers address a wide range of customer needs for everything from general purpose rack and blade servers to mission critical hardware and high performance computing (HPC). Huawei has also imbued its servers with advanced capabilities to support emerging compute intensive applications like

Artificial Intelligence (AI), Big Data Analytics, Deep Learning and Machine Learning. In 2021, Huawei spent $22.1 billion or 22.4% of its FY 2021 sales on R&D. As Bloomberg News noted, Huawei, China's largest technology vendor nearly doubled its R&D budget over the past half-decade. This is more than any company in the world outside America. Although, Huawei is best known for its telecommunications and phone equipment, it has made great strides worldwide (except in the U.S. due to government sanctions) in expanding its footprint in the server market. Huawei invests well over $1 billion of its R&D budget in server technology. Huawei is emphasizing security via best practices documents on "How to Build a Proactive Defense System?" via its HiSec solution which enables more intelligent threat detection, threat response, security operations and maintenance. Huawei says HiSec improves the threat prevention capabilities of enterprise networks and the telecom infrastructure, thus increasing security O&M efficiency and reducing O&M costs. Huawei also offers a number of new security offerings for its various server solutions in the data center, the cloud and the network.

HPE Reliability and Security Survey Highlights

- **HPE's Superdome** line of servers (including the Integrity and Flex models) also exhibit high reliability of five and six nines for a 92% majority of its customers. HPE's Superdome had just 1.44 minutes of unplanned per server downtime. And 89% of HPE survey respondents said their firms discovered and shut down security breaches "Immediately or within the first 10 minutes." The ITIC survey data shows that HPE Superdome servers each experienced three (3) successful security hacks within the last 18 months. This puts the HPE hardware platforms in the top five most secure systems. The Superdome portfolio also benefits from the inherently strong stability of the HPE hardware. HPE has made security, feature/performance innovation and after-market technical service and support, its top priorities. All of this is critical in the increasingly insecure, complex and interconnected Digital Age. HPE is well entrenched in corporate enterprises from SMBs to the largest multinational businesses. The HPE Superdome Flex Server features RAS capabilities and end-to-end security to protect vital workloads. The HPE Superdome Flex Server, for example delivers scalability of up to 32 sockets. This is 2.3x the scalability of prior generation servers. It also features an In-memory design and memory capacity of 768GB - 48 TB in a single platform. HPE Superdome Flex Server has a modular design that scales flexibly from 4- to 32-sockets in 4-socket increments. HPE also says the Superdome Flex server has a more cost-efficient entry point for mission-critical workloads at 4 sockets; it delivers up to 45% lower acquisition cost compared to previous models. HPE also emphasizes reliability claiming that the Superdome Flex Server embedded RAS capabilities deliver five nines – 99.999% – of single-system availability. HPE also asserts that the Superdome Flex server reduces human error via its predictive fault handling Error Analysis Engine. Security and human error are two issues that are closely linked and undermine security and reliability. This engine predicts hardware faults and initiates self-repair with no need for human intervention or "operator assistance." It contains errors at the firmware level, including memory errors, before any interruption can occur at the Operating System layer with

HPE's "Firmware First" approach. HPE also provides continuity for Linux workloads with its HPE Serviceguard for Linux (SGLX) high availability and disaster recovery clustering solution. This enables businesses to safeguard their servers running Linux against a multitude of infrastructure and application faults across physical or virtual environments over any distance.

Downtime Comparison Costs by Server Platform

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 a few minutes. 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 significant disparity in the annual downtime cost comparisons among the top performing and the least reliable server 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.

Corporations that deploy the most highly reliable servers: the IBM z14 and z15; the IBM LinuxONE III; the IBM Power8, Power9 and IBM Power10; the Lenovo ThinkSystem; Huawei KunLun and HPE Superdome achieved the most economical TCO and immediate ROI.

Exhibit 3. Unplanned Monthly Downtime Per Server/Per Minute Assuming Cost of $100K

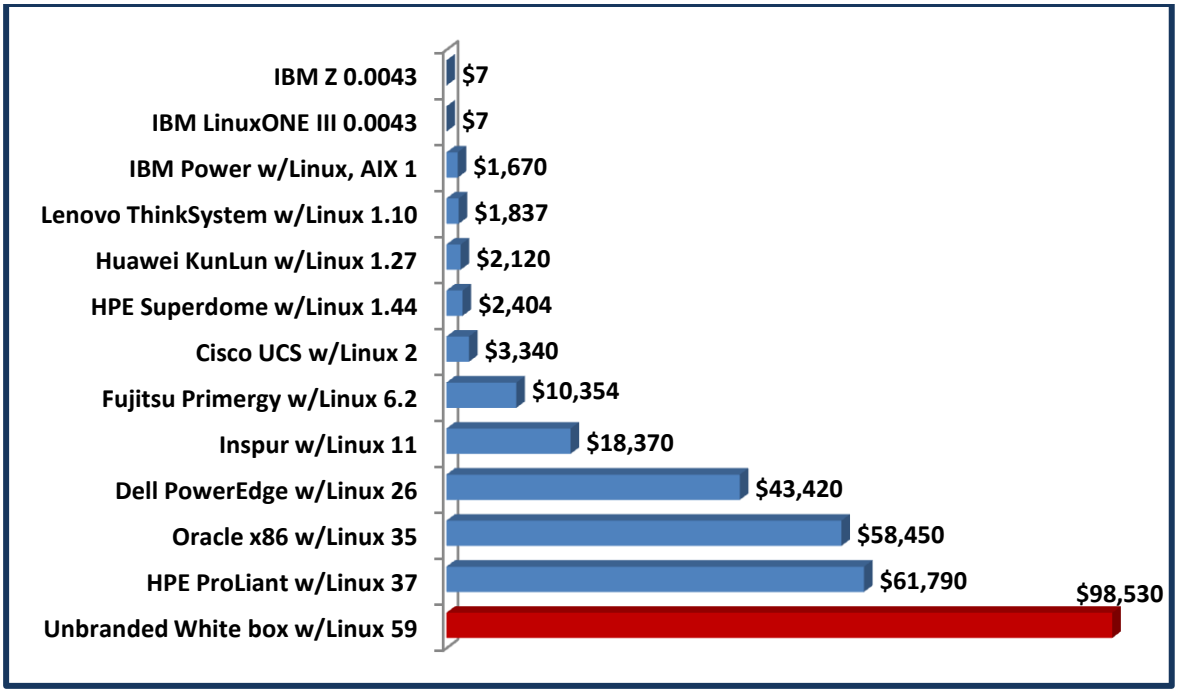

Source: ITIC 2022 Global Server Hardware, Server OS Reliability Survey

Exhibit 4 provides a cost comparison between ITIC's 2021 and 2022 Global Server Hardware, Server OS Reliability survey statistics. To reiterate, IBM Z, IBM LinuxONE III, IBM Power, Lenovo ThinkSystem, Huawei KunLun and Fusion and HPE Superdome maintained their previous high uptime rates despite the challenges posed by the sharp spike in security and data breaches, ongoing supply chain disruptions and increasingly complex deployments

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

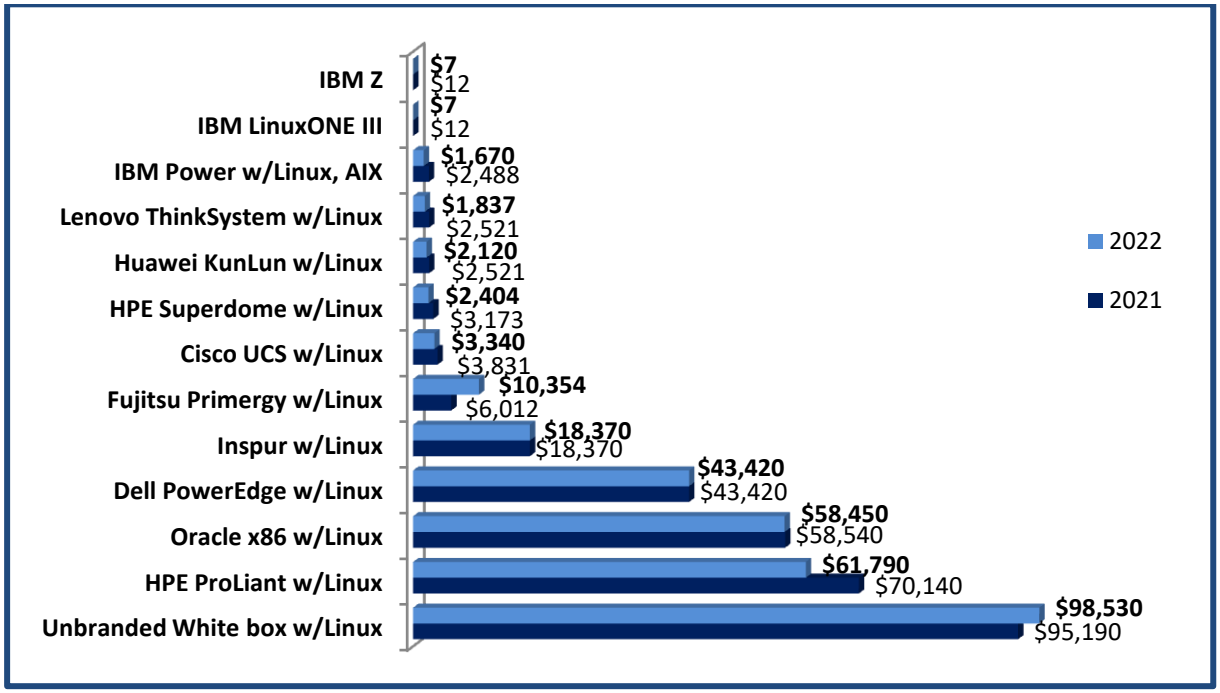

Source: ITIC 2022 Global Server Hardware, Server OS Reliability Survey

The pivotal role server reliability plays in cost effectiveness is clearly evident among the various server distributions in **Exhibit 4** above**.** IBM's Power8, Power9 and Power10 servers with just one (1) minute of monthly per server downtime potentially costs enterprise customers $1,670 in annual outages due to server flaws/outages while Lenovo's ThinkSystem servers with 1.10 minutes of yearly downtime due to inherent server flaws may potentially cost its customers about $1,837 for one minute of per server monthly downtime whose businesses calculate hourly downtime losses at $100,000 Inspur hardware scored in the middle of the pack with 11 minutes of per server/per annum down due to inherent hardware problems. By that calculation, Inspur corporate customers potentially could shell out $18,370 yearly in per server annual downtime costs, an increase of over nine percent from $16,670 posted in ITIC's previous 2021 Reliability poll. This is approximately 11x more than IBM Power Systems enterprises and 10x greater Lenovo ThinkSystem organizations spend. However, Inspur reliability costs are roughly 5x less than unbranded White box servers. In ITIC's 2022 Reliability Study, unbranded White box hardware recorded per server downtime of 59 minutes; calculated at an hourly downtime rate of $100,000, that could cost corporations $98,530 per server/per minute every month. That's up close to seven (7%) percent from 2020 when 55 minutes of White box servers cost $91,850 for each minute of unplanned per server monthly unavailability.

And as **Exhibit 4** also shows, Dell's popular line of PowerEdge servers recorded 26 minutes of unplanned downtime at a potential cost of $43,420 per server/per minute each month. The Dell reliability statistics held steady from the prior ITIC Global Reliability survey.

Another way of analyzing the data is to calculate the potential annual cost savings by server platform based on the average per server/per minute costs. For instance, IBM Power8, Power9 and Power10 servers which experienced one (1) minute of unplanned per server downtime costing $1,670 can potentially save corporate customers $41,750 annually versus a Dell PowerEdge server that averaged 26 minutes of per server monthly unplanned downtime at a cost of $43,420.

A single minute of per server downtime on a Dell PowerEdge server outage at an hourly downtime cost estimated at $100,000, potentially costs 26x more than comparable downtime involving IBM Power models and 24x more than the Lenovo ThinkSystem servers.

Costs quickly add up when businesses factor in the total number of affected servers across the entire ecosystem – datacenters, the cloud and the network edge. Many corporations today have virtualized server farms in their on-premises data centers; nearly 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 $16,670 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, an organization 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, 91% 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 III customers reported just 0.0043 minutes of unplanned per server/per minute monthly downtime. This equates to just under $22 per server/per minute outage cost calculated at an hourly downtime cost of $300,000.

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

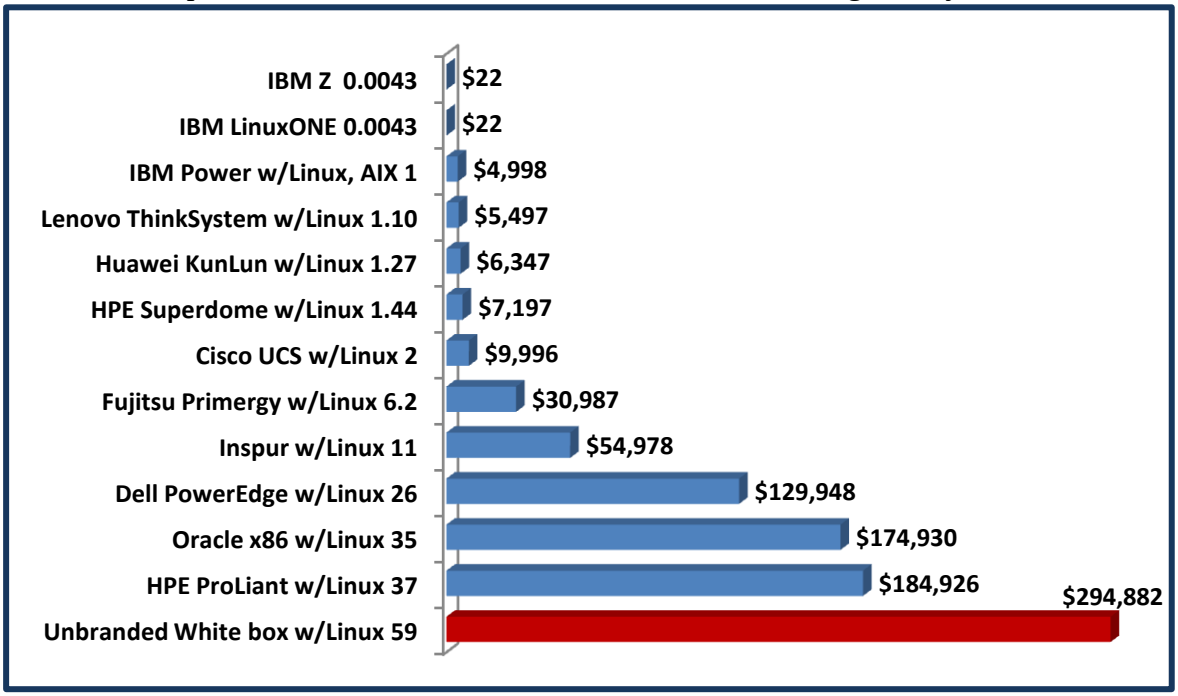

Source: ITIC 2022 Global Server Hardware, Server OS Reliability Survey

The IBM Power8, Power9 and Power10 per server, per minute of unplanned monthly downtime potentially cost businesses $4,998 when assuming an hourly outage rate of $300,000.

In contrast, Dell PowerEdge server survey respondents reported an average of 26 minutes of per server monthly unplanned outages which potentially costs of $129,948.

This makes the IBM Power8, Power9 and Power10 servers 26x more economical than the rival Dell PowerEdge servers. This means the IBM Power Systems servers can deliver a significant cost savings of $124,950 for a single server compared to the Dell PowerEdge.

The Lenovo ThinkSystem and Huawei KunLun servers likewise deliver impressive cost savings. Lenovo ThinkSystem recorded 1.10 minutes of unplanned monthly server downtime which equals $5,497 assuming a cost of $300,000 of hourly downtime. The Lenovo ThinkSystem servers deliver a cost savings of $124,451 per server, per minute compared to the Dell PowerEdge servers; that makes the Lenovo servers 23.5x more cost effective than the Dell hardware.

The Huawei KunLun hardware platform posted 1.27 minutes of per server monthly unplanned downtime; this totals $6,347 per server/per minute at an hourly downtime rate of $300,000. They likewise are over 20x more cost efficient than the competing Dell PowerEdge distributions.

It's noteworthy that the IBM, Lenovo, Huawei, HPE and Cisco servers have continually improved their YoY reliability scores and thus delivered measurably higher cost savings compared to competitors like Dell, Inspur and Oracle hardware whose per server/per minute unplanned monthly downtime statistics have remained relatively static for the last several years.

The aforementioned costs will fluctuate depending on the amount of time in minutes and hours that each server distribution experiences monthly and annually. Other factors such as whether or not the unplanned outage occurred during peak usage time; whether or not data was lost, stolen or damaged and whether or not the company was sued or had to pay civil or criminal penalties due to regulatory non-compliance, must also be calculated.

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 53% 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 individual unplanned downtime amounts. Once again, the IBM Z and IBM LinuxONE III delivered the best economies of scale and the lowest TCO results. To reiterate, in ITIC's most recent 2022 Global

Reliability report, the IBM Z and IBM LinuxONE III offerings recorded near imperceptible monthly unplanned downtime rates of just 0.0043 minutes per server. The per server, per minute monthly unplanned downtime cost is an extremely economical $72, which is a decline of 39% (YoY) from the $120 based on the per server unplanned monthly outage time recorded in ITIC's 2021 Global Server Hardware Server OS Reliability survey.

One minute of per server downtime on an IBM Power8, Power9 and Power10 server calculated at $1 million (USD) cost $29,165 in 2019. In ITIC's latest 2022 survey, that figure was reduced to $16,700, a cost savings of $12,465 over the last three years.

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

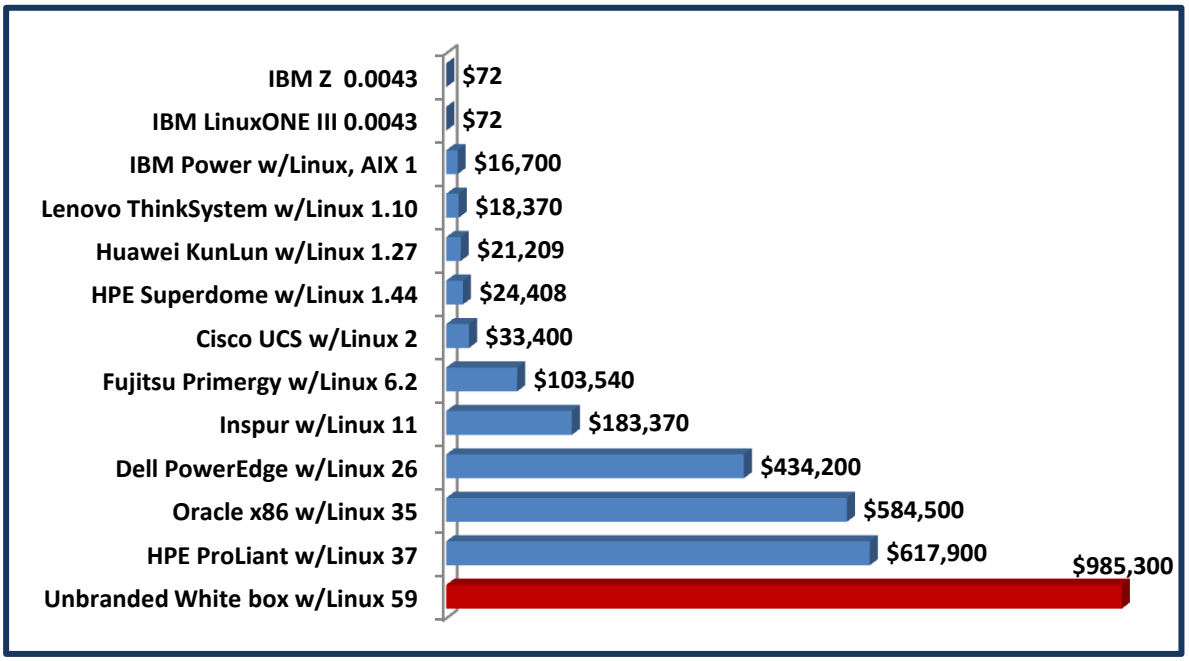

Source: ITIC 2022 Global Server Hardware, Server OS Reliability Survey

Hourly Cost of Downtime: 44% of Firms say Losses Top $1M

ITIC's latest 2022 Global Reliability poll now indicates that for the second year in a row, 99% of small and mid-sized enterprises (SMEs) and large corporations' hourly downtime losses exceed $100,000 **(See Exhibit 7)**. Some 25% of survey respondents say hourly downtime expenses total $301,000 to $400,000. And 44% of firms said a single hour of downtime in 2022 costs their businesses one million ($1 Million), while 20% of survey respondents indicated their hourly downtime losses exceed five million ($5M). Based on ITIC's Hourly Cost of Downtime

research over the past decade as well as rising costs, it's highly probable these percentages will continue to rise.

As ITIC has referenced repeatedly in this report, hourly downtime costs are on the upswing for all businesses irrespective of size or vertical market. ITIC's latest research indicates the average cost of a single hour of downtime now exceeds $300,000 for 91% of small and mid-size (SMEs) businesses and large enterprises, with 500+ employees. These costs are exclusive of any litigation, civil or criminal penalties or fines that an organization may incur.

As **Exhibit 7** illustrates, the IBM Power8, Power9 and Power10 systems also scored big reliability gains delivering an estimated 17% customer cost savings over the last three years.

Exhibit 7. Hourly Cost of Server Downtime Tops $1 Million for 44% of Enterprises

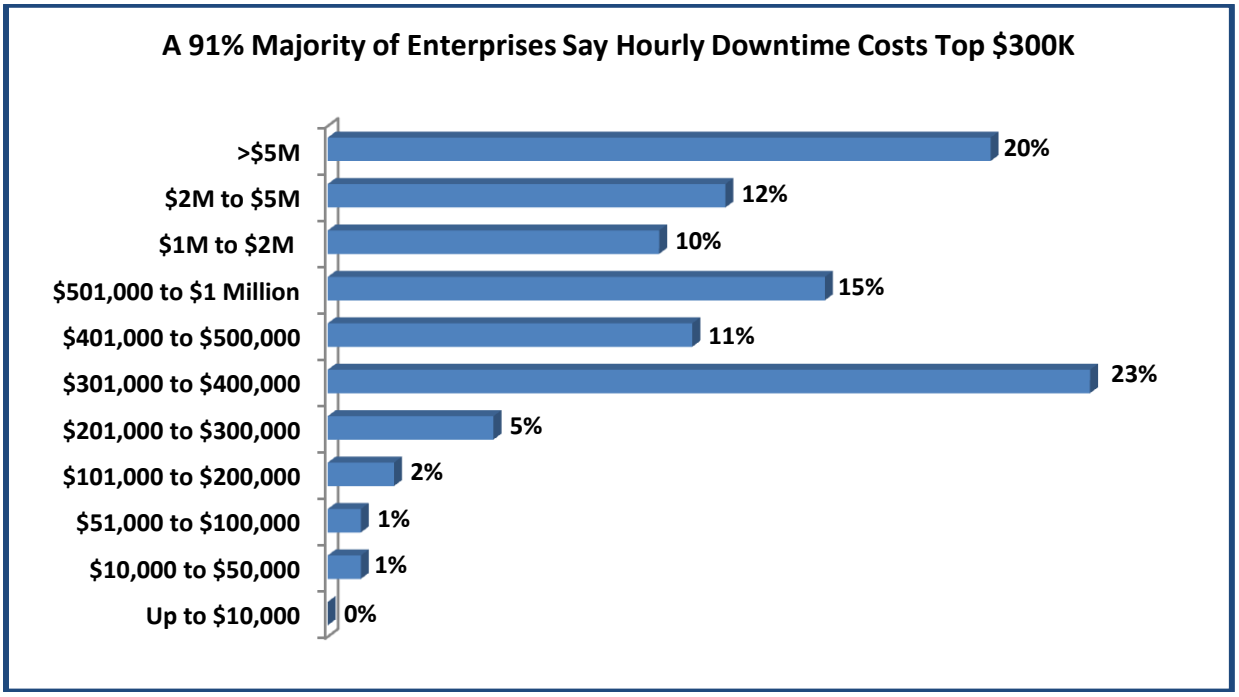

Source: ITIC 2022 Global Server Hardware, Server OS Reliability Survey

Anytime a company can reduce downtime, it solidifies and improves its bottom line.

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

- **The lingering effects of the global pandemic** which has created a slew of new management and technical issues for beleaguered IT and security administrators.
- **Supply chain disruptions.**
- **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.** 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' near-total reliance on computers and networks to conduct business.** Downtime of even a few minutes interrupts productivity and daily business operations. Downtime also has a domino effect, even if no data is lost, stolen, destroyed or hacked.

ITIC anticipates that all of 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 persist unabated. Consequently, businesses reasonably can expect hourly downtime costs to increase. It is imperative that organizations implement and enforce the necessary measures to ensure the reliability and security of their hardware, software applications and connectivity devices across the entire network ecosystem. Security and security awareness training are absolutely necessary to maintain the uptime and availability of devices and data assets. This will ensure continuous business operations and mitigate risk.

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. Even five, 10 or 30 minutes of downtime during peak usage hours can deal SMBs a crippling monetary blow. Prolonged outages or a series of multiple outages of shorter durations could put SMBs at heightened risk of closure.

The wide-ranging reliability disparities among the 18 mainstream server platforms present a cogent case for how the least expensive servers can actually cost companies more than higher end distributions are specifically architected to support reliability with robust capabilities.

Security Hacks, Data Breaches are Top Cause of Downtime

Security hacks, data breaches continue to surge, rising by 42% over the last 20 months, according to ITIC's latest survey data. Nearly three-quarters or 78% of survey respondents cited

security as the number one cause of unplanned downtime. 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 64% 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

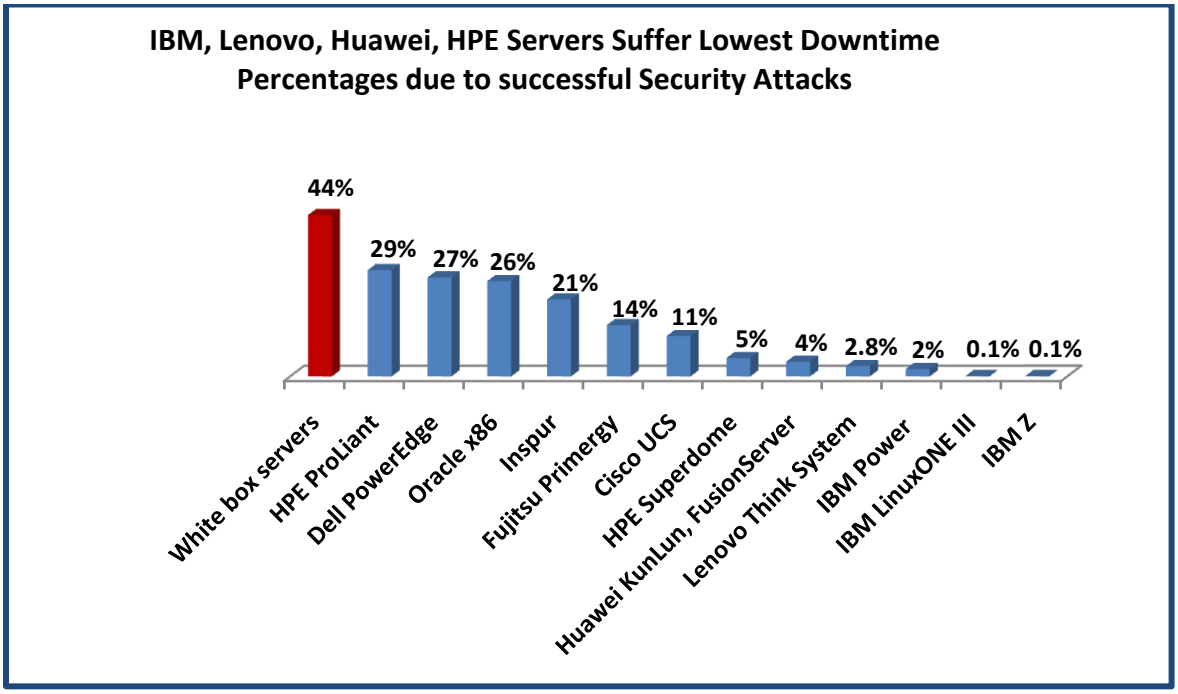

Source: ITIC 2022 Global Server Hardware, Server OS Security Survey

Data Analysis: Security and Resiliency Improve Uptime and 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 39% of respondents now aspire to "five nines" or 99.999% reliability or greater. In order 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. A 73% majority of IBM Z users and 69% of IBM Power customers regularly retrofit or upgrade their hardware every three years compared with 29% of Dell PowerEdge customers.

These reliability improvements are the result of IBM, Lenovo, Huawei, HPE and Cisco's continuing commitment to regular, planned product releases; an emphasis on innovation such as the ability to support more compute-intensive workloads and advanced functionality to support emerging technologies like AI and Analytics. IBM for example, typically ships a new Power processor approximately every three to four years. The new IBM Power10 servers are designed to securely and efficiently scale world's most business-critical workloads from ERPs and high-performance databases to AI inferencing and machine learning.

Besides the obvious technical merits of the IBM Z, IBM LinuxONE and IBM Power mission critical hardware, the server reliability is reinforced by the greater expertise of the IT administrators in corporations that use IBM. ITIC's reliability survey data shows that IBM IT administrators typically have 10 or more years experience. This is also true in many instances of large, high-end enterprises that deploy Lenovo, Huawei KunLun and HPE Superdome servers. 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 with one-to-five year's experience.

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

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-year upgrade cycle and retrofit their server hardware 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.

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

Commodity server users should not defer upgrades or retain servers well beyond the recommended three-year upgrade cycle. Over 60% of Dell PowerEdge and unbranded White box users retain the servers for four, five or even six years while increasing the application workload. And this has been the case for the last four years. This is just asking for trouble. The exceptions to this rule: very small businesses whose application environment remains static.

Proactive Users Bolster Reliability

The most reliable server vendors also benefited from the proactive behavioral habits and expertise of their corporate customers. On average, 82% 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 higher 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.

External Trends Impacting/Influencing Server Reliability

Finally, it comes as no surprise that security hacks, User error and remote working remain the top three causes of unplanned downtime. These issues will persist for the foreseeable future. A 78% majority of survey participants cited security as the number one cause of unplanned server downtime; 64% said human error caused unplanned server outages. Meanwhile, 64% of survey participants attributed increased downtime to management and security issues such as remote working. Post-pandemic, many businesses have transitioned to a hybrid work environment, requiring employees to work onsite two to three days a week. These factors combined with ongoing IT staff shortages and higher levels of security and data breaches, will continue to present significant challenges for IT and security administrators.

The data deluge, data sprawl as well as the rapid shift to the cloud are factors cited by 55% of ITIC Reliability survey respondents as presenting challenges to overall server and network

reliability. This is bolstered by the Netskope Cloud and Threat Report: Cloud Data Sprawl[1] released in July 2022. It 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.

Other Survey Highlights

Among the other significant survey highlights:

- **Data Breaches Continue to Surge Post-Pandemic.** Some 47% of enterprises said data breach attacks – most notably ransomware, phishing and email scams - rose during the January through the August 2022 time period. 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 as those vendors make ongoing substantial investments/improvements in embedded security features.
- **Semiconductor Supply Chain Issues Impact Server Reliability.** The well documented supply chain issues that emerged during the global COVID-19 pandemic continue to plague corporate enterprises. The ongoing chip shortage has the potential to impact corporate server hardware as well as a wide range of consumer devices. A January 25. 2022 New York Times article, "Commerce Dept. Survey Uncovers 'Alarming' Chip Shortages,[2]" said the United States is facing an "alarming" shortage of semiconductors, based on a government survey of more than 150 companies that make and buy chips. The

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

- [2]"Commerce Dept. Survey Uncovers 'Alarming' Chip Shortages, The New York Times, January 25, 2022. URL: https://www.nytimes.com/2022/01/25/business/economy/chips-semiconductors-shortage.html

situation is threatening American factory production and helping to fuel inflation, according to Commerce Secretary Gina M. Raimondo[3].

- **IBM, Lenovo, Huawei, HPE and Cisco servers deliver best security.** IBM, Lenovo, Huawei, HPE and Cisco servers (in that order) also attained the highest levels of security for the third straight year recording the fewest number of successful hacks. ITIC's latest 2022 Global 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**.** Less than one percent – 0.1% – of IBM Z and IBM LinuxONE III servers experienced a successful data breach. And of that 0.1%, approximately 82% of IBM survey respondents. 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.
- **IBM, Lenovo and Huawei KunLun Servers Lowest Percentage of Hardware Failures:** As in the prior ITIC 2021 Reliability Update Survey, the most recent 2022 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 – 0.3% – of IBM z13, z14 and z15 servers, for example, 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 Huawei KunLun 0.7% each during the first 12 months of deployment.
- **Increase in Server Workloads** causes reliability declines in 70% of servers over four (4) years old that haven't been retrofitted or upgraded to accommodate increased workloads.
- **Planned Downtime Increases:** Over two-thirds – 68% – of firms now spend from two-to-eight hours monthly on planned downtime. Much of this is attributable to applying security patches or provisioning new applications.

- [3] Results from Semiconductor Supply Chain Request for Information, US Department of Commerce, January 25, 2022. URL: https://www.commerce.gov/news/blog/2022/01/results-semiconductor-supply-chain-request-information

Conclusions

Downtime is disruptive and expensive. It can also irreparably damage a company's reputation. In some extreme cases, business and monetary losses as a result of unreliable servers can cause the company to go out of business due to sustained losses and possible litigation in the wake of a 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 and virtualization components fail to deliver a minimum of "four nines" – and preferably "five nines" of uptime – put their organizations at risk for increased downtime: higher costs and greater exposure to security hacks and data breaches. Unreliable systems can also have a domino effect that disrupts business operations and heightens security risks for business partners, suppliers and customers.

ITIC's 2022 Global Server Hardware and Server OS Reliability Survey findings indicate that for the 14th straight year, the elite IBM z14 and z15 and the LinuxONE III cemented their positions as the best-in-class hardware among all mainstream server distributions. The IBM Z mainframe and LinuxONE III again achieved its best results in every reliability, and security category. The IBM Z and the IBM LinuxONEIII are in a class of their own. A 96% majority of survey participants said the z14 and z15 and LinuxONE III delivered seven nines (99.99999%) or greater 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, Power9 and Power10 servers also met and exceeded their best ever reliability metrics – with 93% of the Power platforms averaging five and six nines of server uptime. The Power8, Power9 and Power10 servers posted just one (1) minute of unplanned per server monthly downtime.

The Lenovo ThinkSystem servers ranked as the most reliable x86-based hardware platform for the ninth straight year recording just 1.10 minutes of per server unplanned monthly downtime.

The Huawei KunLun and Fusion servers, registered 1.27 minutes respectively of unplanned per server downtime, followed closely by the HPE Superdome servers with 1.44 minutes of per server unplanned outages per month. They averaged four, five and six nines of reliability – according to over nine-out-of-10 corporate survey respondents.

The top performing and most reliable server distributions consistently achieve their top scores 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 and VR.

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 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.

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.

Businesses can however lower their exposure and mitigate risk by deploying the most reliable and efficient server hardware and server operating systems.

Organizations are increasingly risk averse with many enterprises adopting a zero tolerance for downtime, operational outages and productivity disruptions.

Reliability is among the most crucial metrics in the organization. Improvements or declines in reliability mitigate or increase technical and business risks to the organization's end users and its external customers. The ability to meet service-level agreements (SLAs) hinges on server reliability, uptime and manageability. These are key indicators that enable organizations to determine which server operating system platform or combination thereof is most suitable.

Businesses must regularly replace and refresh their server hardware and server operating systems with the necessary patches, updates and security fixes *as needed* to maintain system health. The onus is also on the server hardware and server operating system vendors to provide realistic recommendations for system configurations to achieve optimal performance. Vendors also bear the responsibility to deliver patches, fixes and updates in a timely manner and to inform customers to the best of their ability regarding any known incompatibility issues that may potentially impact performance. Vendors should also be honest with customers in the event there is a problem or delay with delivering replacement parts.

Recommendations

An organization's ability to achieve four, five, six and seven nines of server hardware, server OS and server application reliability/uptime and availability is a two-way street, not a one-way foot path. Corporate clients, their IT and security administrators must also invest in maintaining the reliability, security and management of their infrastructure.

Companies require robust servers that can accommodate current application workloads and flexibly scale to accommodate the technical demands of business operations for at least the next two-to-three years.

ITIC strongly 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 or not the company's reputation suffered as a result of a reliability incident; did any litigation ensue; were customers, business partners and suppliers impacted (and at what cost) and at least try and gauge whether or not the company lost business or potential business.
- **Construct a list of best practices.** Chief technology officers (CTOs), Chief Data Officers (CDOs), software developers, engineers, network administrators and managers should have extensive familiarity with the products they currently use and are considering. Check and adhere to your vendors' list of approved, compatible hardware, software and applications.
- **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. **I**nstall 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 or not 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 demands robust hardware. Organizations should purchase the beefiest server configuration their budgets will allow. Waiting four, five or six years to refresh servers while placing greater demands on the hardware, is asking for trouble.
- **Calculate the Cost of Hourly Downtime.** There is no "one size fits all." Hourly downtime costs will vary according to the length, severity and duration of the outage and whether or not any data was lost, stolen, destroyed or changed. 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 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 *2022 Global Server Hardware and Server OS Reliability Survey*, polled 1,550 C-level executives, IT and security administrators in corporations worldwide from January through August 2022. 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 impact positively and negatively impact and influence daily business operations and technology decisions. To ensure data integrity ITIC employed authentication and tracking mechanisms to prevent tampering and to prohibit multiple responses by the same parties.

Survey Demographics

Respondents came from companies ranging from small and medium businesses (SMBs) with fewer than 50 workers, to multinational enterprises with over 100,000 employees.

All market sectors were equally represented: SMBs with one-to-100 employees accounted for 28% of the respondents. Small and medium enterprises (SMEs) with 101-to-1,000 workers represented 27% of the participants. The remaining 45% of respondents came from large enterprises with 1,001 to over 100,000 employees. Survey respondents hailed from 49 different vertical markets. Approximately 65% of respondents hailed from North America; 35% were international customers who hailed from more than 30 countries throughout Europe, Asia, Australia, New Zealand, South America and Africa.

Appendices

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

ITIC Website and links to survey data and blog posts:

https://itic-corp.com/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-availiability/

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/