

ITIC 2018 Global Reliability Survey Mid-Year Update

August 3, 2018

IBM Z, IBM Power Systems, Lenovo System x, Huawei and Stratus Servers Top ITIC Reliability Poll

ITIC Position

Rapid adoption of new technologies coupled with complexity, increased workloads, the proliferation of endpoint connections and pervasive security threats all make core infrastructure reliability imperative.

Demands on servers and operating systems are soaring. Applications are bigger and more complex. Escalating workloads consume more precious system resources. Internet of Things (IoT) and Network Edge deployments are ramping up as are the use of Bring Your Own Device (BYOD) computers and mobile devices. The result: more endpoints and connections to oversee. Ever more sophisticated security threats also threaten the reliability of the core infrastructure. If reliability is compromised and servers and applications are unavailable, business ceases. Businesses are risk averse; there is no tolerance for downtime.

IBM Z, IBM POWER, Lenovo System x, HPE Integrity and Huawei KunLun Servers Deliver Highest Uptime/Reliability

ITIC Global Server Hardware and Server OS Reliability Survey 2018 Mid-Year update reveals that the reliability gap among high end and inexpensive commodity based servers is widening.

For the 10th year in a row, corporate enterprise users said IBM's Z Systems Enterprise mainframe class server achieved near flawless uptime, recording just nine seconds of unplanned monthly server downtime (**See Exhibit 1**). The IBM Z was closely followed by veteran niche market vendor Stratus Technologies' ftServer platform with just over nine seconds of unplanned downtime per month. Like the IBM Z, Stratus' ftServer is known for delivering mainframe-like fault tolerant reliability and availability.

Among mainstream server platforms, IBM POWER, the Lenovo System x platform, Hewlett Packard Enterprise (HPE) Integrity Superdome X and Huawei's mission critical KunLun and x-86-based FusionServers demonstrated the highest levels of reliability/uptime among 18 server hardware platforms.

© 2018 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.

Information Technology Intelligence Consulting



Audience

Corporate enterprises, CEO, CIO, CTO, VP/Manager of IT, Security Administrators, Compliance Professionals, Vendors and Third party systems integrators, consultants.

Relevance

Reliability and continuous uninterrupted access to servers, systems, applications and data assets are essential for business productivity, revenue and the corporation's reputation.

Laura DiDio
Principal Analyst
ldidio@itic-corp.com

Page 1 of 13
Phone: (508) 887-9814
Mobile: (508) 740-1513
Twitter: @lauradidio
Skype: laura.didio

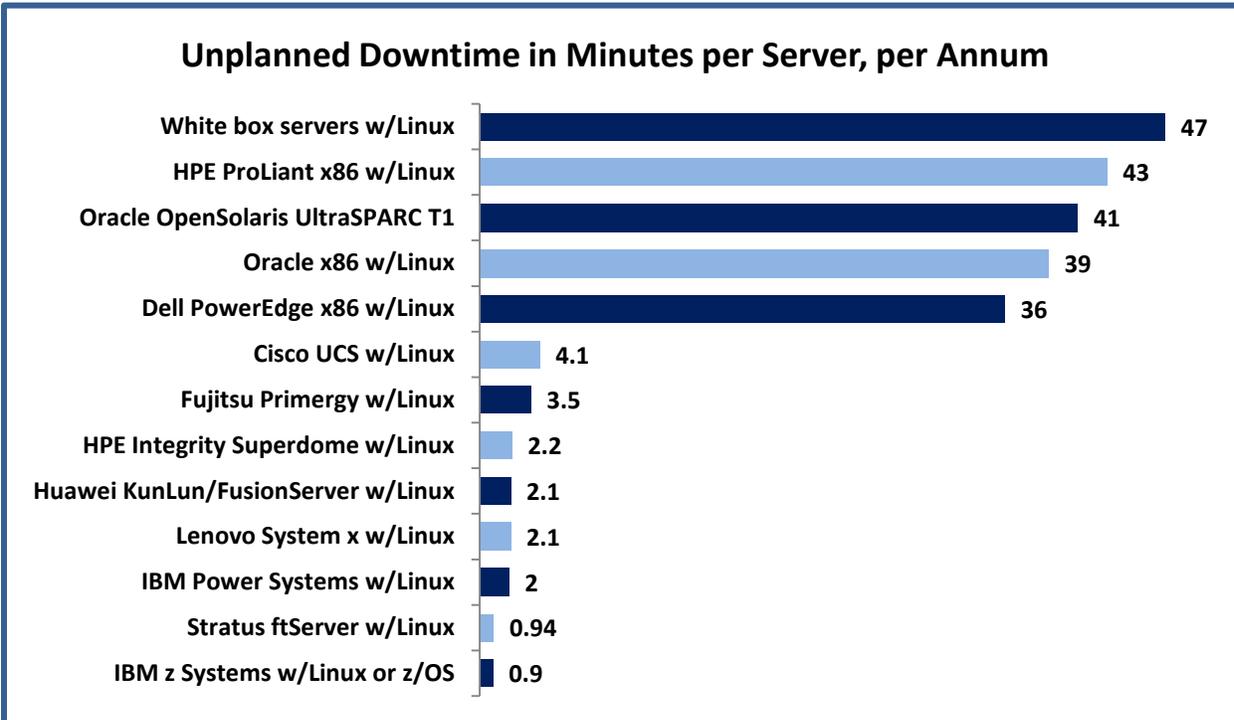
ITIC 2018 Global Server Hardware and Server OS Reliability Mid-Year update survey polled over 800 organizations worldwide from April through mid-June 2018. To obtain the most accurate and unbiased results, ITIC accepted no vendor sponsorship.

As **Exhibit 1** below illustrates, the IBM Z mainframe and Power Systems and Lenovo System x servers successfully maintained their decade-long status as best in class for reliability and performance. In the latest poll, they were joined by HPE's Integrity Superdome X, Huawei's KunLun high end server and the FusionServer x86 platforms and the Stratus Technologies' ftServers – all of which achieved equivalent reliability scores. Each of these distributions delivered solid “four nines,” 99.99% and “five nines,” 99.999% of *inherent* hardware reliability. This means that the systems and infrastructure experienced minimal amounts of unplanned downtime that occurred due to a temporary or permanent failure in the server hardware or any of its component parts. The poorest performing distributions were “White box” unbranded servers with 47 minutes and HPE's ProLiant x86 with 43 minutes of unplanned downtime. They are 17 to 18 times less reliable than IBM Power Systems (POWER8 and POWER9), Lenovo Systems X, HPE Integrity Superdome and the Huawei KunLun and FusionServer platforms.

The disparity in reliability amongst the most and least reliable server hardware platforms is attributable to several factors. They include:

- **Aging server hardware.** The reliability declines are most pronounced in aged servers over 3 ½ years old that have not been upgraded or retrofitted to accommodate higher workloads. Older hardware is more common among the less expensive, commodity platforms (e.g. White box unbranded servers, Dell, HPE ProLiant and Oracle x86) where approximately six-in-10 businesses retain the hardware for >3 ½ years and fail to right-size the server. Aged servers (four, five and six years old) are also prone to slower performance and response times and hard drive failures
- **The 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 high end features/functions like RAS 2.0 and hot swap capabilities.
- **Vendor Service and Support.** After market technical service and support as exemplified by IBM and Lenovo – both of which have been consistent, stable and quick to respond, also plays big role. The more quickly a vendor responds to its customer base with answers and fixes, there is more of likelihood that downtime will be minimized. Corporations that choose White box unbranded servers may find themselves left entirely to their own devices if and when problems arise and outages occur.

Exhibit 1. IBM, Lenovo, Huawei & HPE Servers Achieve Four and Five Nines Reliability



Source: ITIC 2018 Global Server Hardware/Server OS Reliability Survey Mid-Year Update

Metrics of three, four and five nines of uptime – 99.9%, 99.99% and 99.999%, – equate to 8.76 hours; 4.38 hours, 52.56 and 5.26 minutes of per server/per annum downtime, respectively.

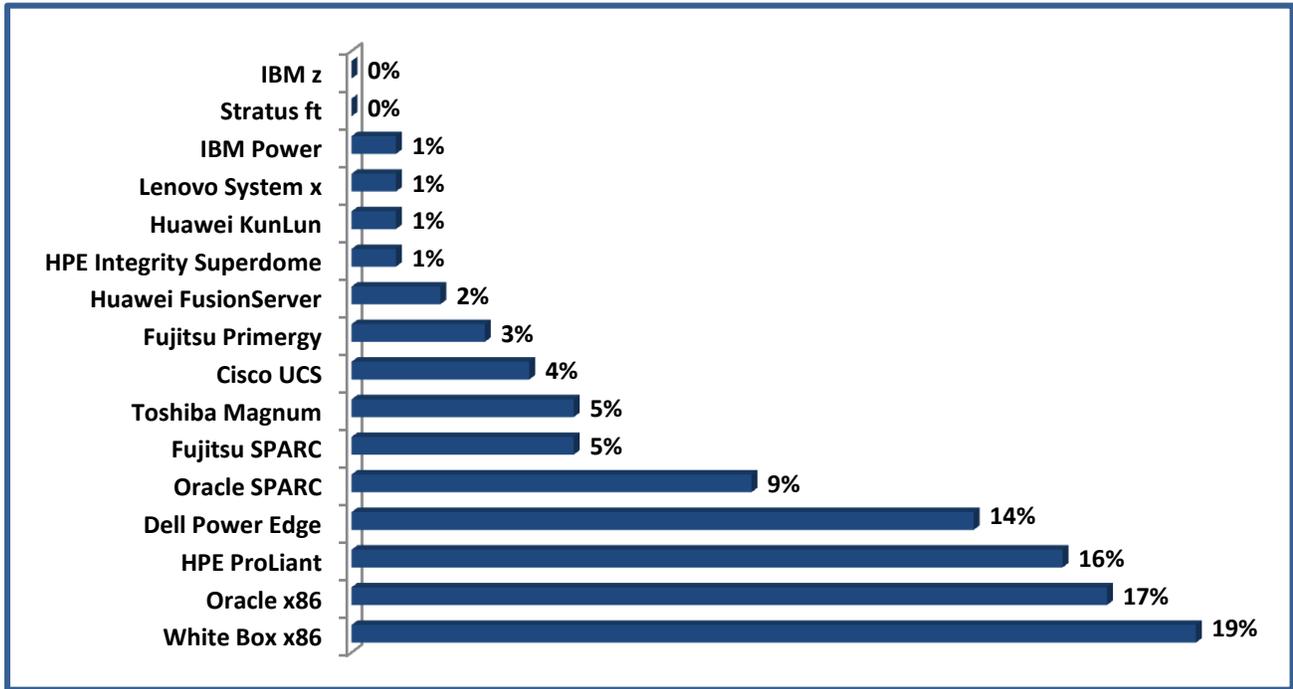
Unplanned downtime is a critical barometer corporations use to calculate and measure lost productivity, data losses, and time to recover and perform remediation activities following an outage. This in turn, enables them to estimate hourly, weekly, monthly and yearly data and monetary losses and to determine Total Cost of Ownership (TCO) and Return on Investment (ROI).

Survey Highlights

To reiterate, as **Exhibit 2** shows, the reliability gap is widening among the high end server distributions and the less expensive commodity based platforms. Again, there are several reasons for this. High end server platforms incorporate many more embedded features such as RAS, in-memory databases, advanced analytics, support for virtualization and cloud computing, which deliver superior reliability and performance compared to less expensive commodity-based servers. Additionally, corporate enterprises that purchase high end servers are up to three times more likely to retrofit and upgrade on a regular two-to-three year schedule. This enables them to avoid “server fatigue” the natural decline in performance and reliability that occurs in

aging hardware when components begin to slow or fail, rendering the servers ill-equipped to accommodate more compute intensive workloads.

Exhibit 2. Reliability Gap of >4 Hours Widens Among Mission Critical & Commodity Servers



Source: ITIC 2018 Global Server Hardware, Server OS Reliability Survey Mid-Year Update

Among the other Survey Highlights:

- IBM z Systems Enterprise** mainframe class servers once again had the distinction of being ranked “best in class” for reliability, accessibility, performance, and security among all server platforms. The z Systems servers had highest reliability/uptime ratings across the board in terms of actual minutes of unplanned per server/per annum downtime. Specifically, IBM z Systems mainframe class servers exhibit true fault tolerance experiencing just 0.90 - less than one minute of *unplanned* per server, per annum annual downtime due to inherent problems with the server or its component parts. That equates to just eight (8) seconds per month or “blink and you miss it,” 2 seconds of weekly downtime.
- The IBM Z family of servers also continues to experience the lowest incidence – **0%** — of **> 4 hours of per server/per annum downtime due to inherent flaws in the server hardware, components or architecture** of any hardware platform.
- Veteran niche market vendor Stratus Technologies ftServer, long notable for its fault tolerance and high availability capabilities – notched just 0.94 seconds – less than one minute of unplanned per server/per annum downtime attributable to any hardware flaws.

- **IBM Power Systems, Lenovo System x and newcomer Huawei KunLun and FusionServer platforms running Linux** all recorded the least amounts of *unplanned* downtime of 2.0 and 2.1 minutes per server/per year respectively due to any inherent flaws in the server hardware or its component parts among all mainstream Linux server platforms. The **HPE Integrity Superdome** platform closely followed notching just 2.2 minutes of *unplanned* downtime related to inherent issues with the underlying hardware or its components. The IBM and Lenovo results are noteworthy for the continuing excellent reliability ratings and the consistency of the IBM Power Systems and Lenovo Systems x platforms over the past 10 years since ITIC began conducting the Reliability poll.
- **IBM Power Systems:** IBM's POWER9 scale-out systems began shipping in March 2018. The line includes the S914, S922, S924, H922, H924 and L922. They support in-memory databases; advanced analytics and cloud environments. POWER9 servers are also cloud-ready and include built-in PowerVM virtualization capabilities. The POWER9 scale-out servers for IBM i, AIX and Linux integrate into organizations' cloud and AI strategies to ensure high performance and RAS capabilities required to support mission-critical workloads like IBM's Db2 and Oracle databases as well as SAP HANA. The IBM POWER9 servers incorporate the latest I/O technology; this includes 25 GB/sec high-speed interconnect for CAPI and OpenCAPI along with embedded PCI-Express 4.0 connectivity. IBM claims this doubles the I/O bandwidth versus PCI-Express 3.0. POWER9 2-socket systems provide up to 4TB of memory which IBM asserts is 33% greater than comparable Intel x86 Xeon systems delivering additional benefit to in-memory databases such as SAP HANA. POWER9 continues in the tradition of previous generations by delivering improved per core performance capabilities compared to its predecessor. Additionally, IBM says that compared to similar to the prior generation POWER8 systems, POWER9 delivers 1.25-1.5X per core the performance capability, enabling clients to further reduce software spending on stacks that are licensed per core. This focus on core architecture and capabilities delivers over 2X the per-core performance over compared x86 systems across a wide range of benchmarks. IBM says that the improvements to the POWER core capabilities also results in lower cost of acquisition and lower cost of ownership over the entire server product lifecycle.
- **Lenovo System x:** In the nearly four years since it purchased IBM's x86 line of servers, Lenovo has maintained a high level of performance, reliability and after-market technical service and support. Lenovo's System x averaged just 2.1 minutes of unplanned per server/per annum downtime due to any inherent flaws in the server or its component parts. Lenovo has responded to the intensely competitive nature of the server market by doubling down on its commitment to embedding its systems with capabilities that bolster the core reliability, availability and performance in its mission critical System x servers. And that should continue with the new Lenovo ThinkSystem server portfolio announced in mid-2017, which will be included in the future ITIC reliability studies. Lenovo says it is combining innovation with reliable, flexible and secure data center systems. It also emphasizes that its open server, storage, networking and system management platforms seamlessly integrate with existing and legacy environments. This last point is critical since ITIC's latest survey data indicates that Human Error, flaws in the server operating system software and the complexity involved in provisioning new applications are all pivotal issues negatively impacting reliability. In

first person interviews with ITIC analysts, Lenovo customers cited the ease of deployment and ease of integration and backwards compatibility as contributing to the underlying reliability and stability of the System x platform. Lenovo users also had high praise for the company's after-market service and support. The Lenovo System x3950 X6 server, an eight-socket 8U rack-mounted server and the next generation eight-socket, 4U ThinkSystem SR950 for example, are optimized for resiliency, reliability and extremely fast application performance. Lenovo's agile system design supports mission-critical databases, enterprise applications, big data analytics, 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.

- **HPE: Integrity Superdome:** HPE's Integrity Superdome line of servers also exhibit exceptional reliability. And they benefit from the renewed stability of the HPE brand and business nearly two years after Hewlett-Packard split into two separate entities. This has enabled HPE to once again renew its focus on innovation and after-market technical service and support which is crucial in the current complex, digital computing environment. HPE is a legacy server brand that is entrenched in corporate enterprises and has built a reputation as a solid vendor over the last six decades. Now that the corporate tumult of last 12 years is behind it, HPE can concentrate on innovation, which it is doing with a vengeance as evidenced by the reliability ratings of the various HPE Integrity Superdome servers.
- **Huawei:** Huawei is a newcomer to ITIC's Reliability Server. Although Huawei has been in the server market for years, until recently, it was better known for its mobile phones. That's changing fast – particularly in the Asia Pacific region. In the last two years Huawei, headquartered in Shenzhen, China, has emerged as one of the top five server hardware vendors worldwide with its high end KunLun mission critical server and its general purpose FusionServer x 86-based servers. Huawei's ascension in the server market –with the exception of North America - is due to a multi-pronged series of bold tactical moves and its ability to formulate and execute a cogent, compelling strategy to ensure long-term success. To successfully compete with rivals Fujitsu, HPE, IBM, Lenovo and others, Huawei's servers support a wide range of customer needs for everything from general purpose rack and blade servers to mission critical hardware to address high performance computing (HPC). Huawei has also imbued its servers with advanced capabilities to support emerging compute intensive applications like Artificial Intelligence (AI), Big Data Analytics, Deep Learning and Machine Learning. [In 2017, Huawei spent a record \\$13.8 billion \(€11.2 billion\) on R&D](#); this is the equivalent of 15% of its annual sales; with approximately one billion invested in server technology. Huawei now ranks third, just behind Alphabet, Google's parent company in total R& D global spending. Huawei estimates its R&D budget for 2018 and beyond, may reach \$20 billion annually.

Reliability Factors

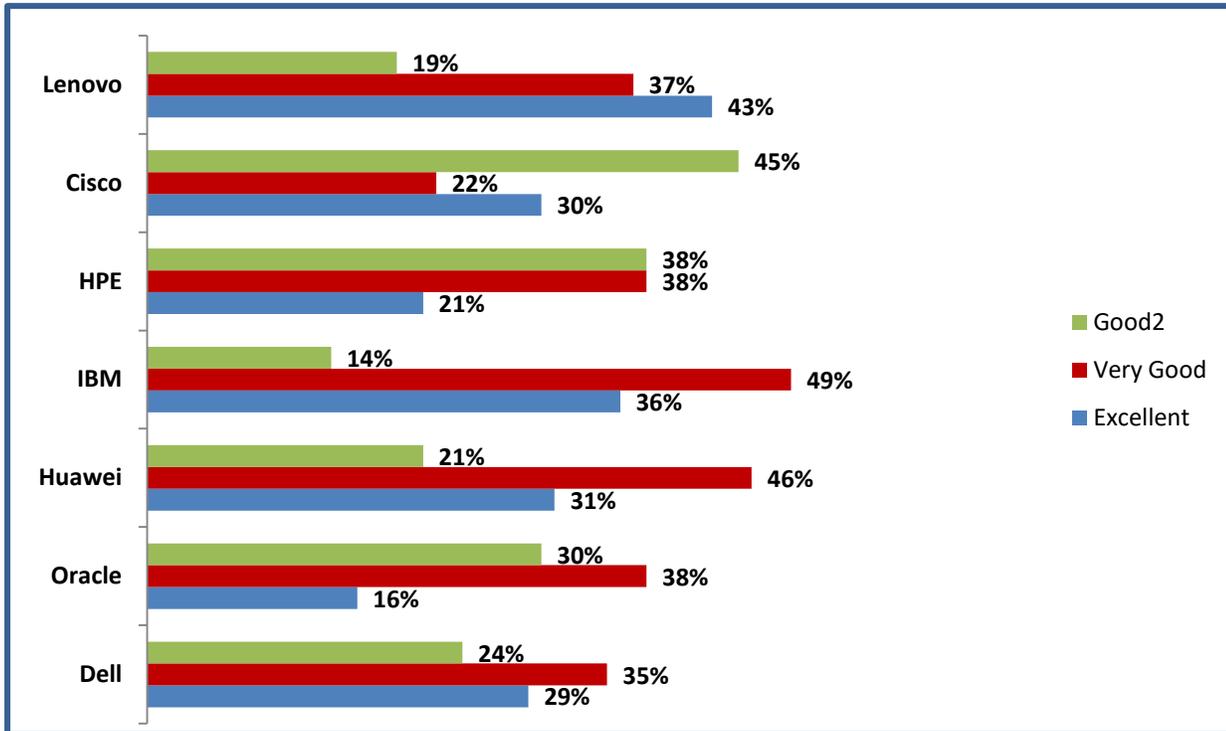
Multiple issues contribute to the high reliability ratings among the various server hardware distributions.

- **Research and Development (R&D)** The leading server vendors like Cisco, HPE, Huawei, IBM and Lenovo have made an ongoing commitment to research and development (R&D) and continually refresh/update their solutions.
- **RAS 2.0.**The higher end servers incorporate the latest Reliability, Accessibility and Serviceability (RAS) 2.0 features/functions and are fine-tuned for manageability and security.
- **Price is not the top consideration.** Businesses that purchase higher end mission critical and x86 systems like Fujitsu's Primergy, HPE's Integrity, Huawei's KunLun, IBM Z and Power Systems and Lenovo System x want a best in class product offering, first and foremost. These corporations in verticals like Banking/Finance, Government, Healthcare, Manufacturing, Retail and Utilities are more motivated with the historical ability of the vendor to act as a true responsive "partner" delivering a highly robust, leading edge hardware. They also want top-notch after market technical service and support, quick response to problems and fast, efficient access to patches and fixes.
- **More experienced IT Managers.** In general, IT Managers, application developers, systems engineers and security professionals who purchase higher end servers from IBM, HPE, Lenovo, Huawei tend to have more experience – approximately 12 to 13 years compared with an average of six years experience among IT managers and systems engineers at companies that purchase less expensive commodity based servers. The more experienced IT managers are more likely to spot problems before they become a major issue and lead to downtime and in the event of an outage, they are more likely to get the servers and applications up and running faster than less experienced peers.

Customer Satisfaction

IBM's Z mainframe and POWER servers and Lenovo Systems X platforms once again garnered high marks for customer satisfaction in the latest ITIC Reliability survey, while HPE's ratings rebounded thanks to its mission critical Integrity Superdome line of servers. Newcomer Huawei's KunLun and FusionServers, making just their second appearance in the ITIC Reliability poll also received high customer approval ratings for product performance, reliability and after-market technical service and support (**See Exhibit 3**).

Exhibit 3. IBM, Lenovo and Huawei Score High Customer Satisfaction Ratings



Source: ITIC 2018 Global Server Hardware/Server OS Reliability Survey Mid-Year Update

Reliability Trends

ITIC’s latest Reliability mid-year update also chronicled the strain that external issues placed on organizations and their IT departments to ensure that the servers and operating systems deliver a high degree of reliability and availability. Among the factors that can potentially undercut reliability:

- **Heavier, more data intensive workloads.** The latest ITIC survey data finds that workloads have increased by 14% to 39% over the past 18 months.
- **Provisioning complex new applications** that must integrate and interoperate with legacy systems and applications. Some 40% of survey respondents rate application deployment and provisioning as among their biggest challenges and one that can negatively impact reliability.
- **Security.** Security and “Human Error” continue to rank as the top issues adversely impacting reliability. Some 62% of businesses indicated that their Security and IT administrators are grappling with a near constant deluge of more pervasive and pernicious security threats. If the availability, reliability and access to servers, operating systems and mission critical main LOB applications is compromised or denied, end user productivity and business operations suffer immediately.

- **Deploying new technologies** like Artificial Intelligence (AI), Big Data Analytics which require special expertise by IT managers and application developers as well as a high degree of compatibility and interoperability.
- **A rise in Internet of Things (IoT) and edge computing deployments** which in turn, increase the number of connections that organizations and their IT departments must oversee and manage.

Minimum Reliability Requirements: 99.99% is Industry Standard

The aforementioned issues are converging at a time when organizations' demand for uninterrupted uptime has reached a peak.

In 2018 and beyond time is not only measured by money. It also equates to productivity and the efficiency, continuity of ongoing, *uninterrupted access* to data assets to maintain the flow of daily operations. If any of these activities are compromised by outages for any reason – business issues that force the company to shut the networks down or operational failure that render the systems, applications or networks unavailable, or that make it impossible for the IT Department and the organization's knowledge workers to do their jobs, business grinds to a halt. As **Exhibit 4** shows, 84% of organizations now require a minimum of “four nines” or 99.99% reliability and uptime. That's an increase of four percentage points from the ITIC 2017 Reliability survey. Metrics of three, four and five nines of uptime – 99.9%, 99.99% and 99.999%, – equate to 8.76 hours; 4.38 hours, 52.56 and 5.26 minutes of per server/per annum downtime, respectively.

In 2008, the first year that ITIC surveyed enterprises on their Reliability requirements, 27% of businesses said their firms could get by on a mere 99% uptime – that equates to nearly 88 hours of unplanned downtime per server/per annum. That much downtime would be shocking and unacceptable in today's “always on” business environment. And a decade ago, four-in-10 corporations – 40% required 99.9% or “three nines” of availability. In that same 2008 survey, only 23% of firms indicated they required a minimum of “four nines” or 99.99% uptime for their servers, operating systems and virtualized environments, while a seven percent (7%) minority demanded the highest levels of “five nines” – 99.999% or greater availability.

Five years ago in ITIC's 2013 Hourly Cost of Downtime poll, 67% of businesses needed a minimum “four nines” or 99.99% or greater reliability/uptime; up 34% from 2008. Currently, 84% of ITIC survey respondents consider 99.99% availability to be the accepted minimum or norm for mission-critical activities. ITIC expects those percentages to continue to rise.

In the latest 2018 survey – *none – 0% of survey respondents* indicated their organizations could live with just “two nines” – 99% uptime or nearly 88 hours of annual unplanned per server downtime! Only a miniscule one percent minority said their organizations required “three nines” or 99.9% reliability which equates to 8.76 hours of per server/per annum unplanned downtime.

Time is money.

ITIC's 2018 Hourly cost of Downtime Survey found that a 98% majority of 800 respondents say that, on average, a single hour of downtime per year costs their company over \$100,000. Additionally, an 81%

© 2018 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.

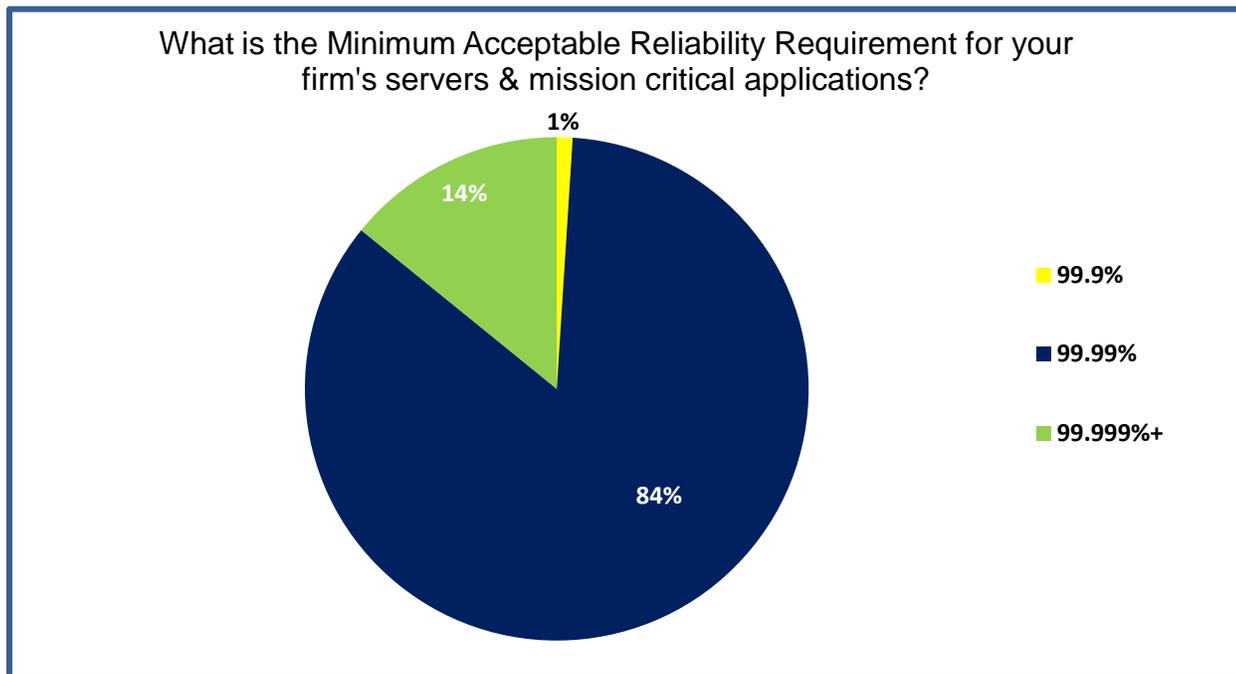
Laura DiDio
Principal Analyst
ldidio@itic-corp.com

Page 9 of 13

Phone: (508) 887-9814
Mobile: (508) 740-1513
Twitter: @lauradidio
Skype: laura.didio

majority indicates that the average hourly cost of a 60 minute outage exceeds \$300,000. And 33% of respondents - three in 10 corporate enterprises - say hourly downtime costs their firms from \$1million to over \$5 million (USD). These statistics are exclusive of any civil or criminal penalties that may ensue in the wake of the outage. These figures also do not take into account any additional monies a company might spend on “goodwill” gestures to customers or business partners whose business transactions may have suffered because of the outage.

Exhibit 4. Eighty-four percent of Businesses Require Minimum 99.99% Uptime

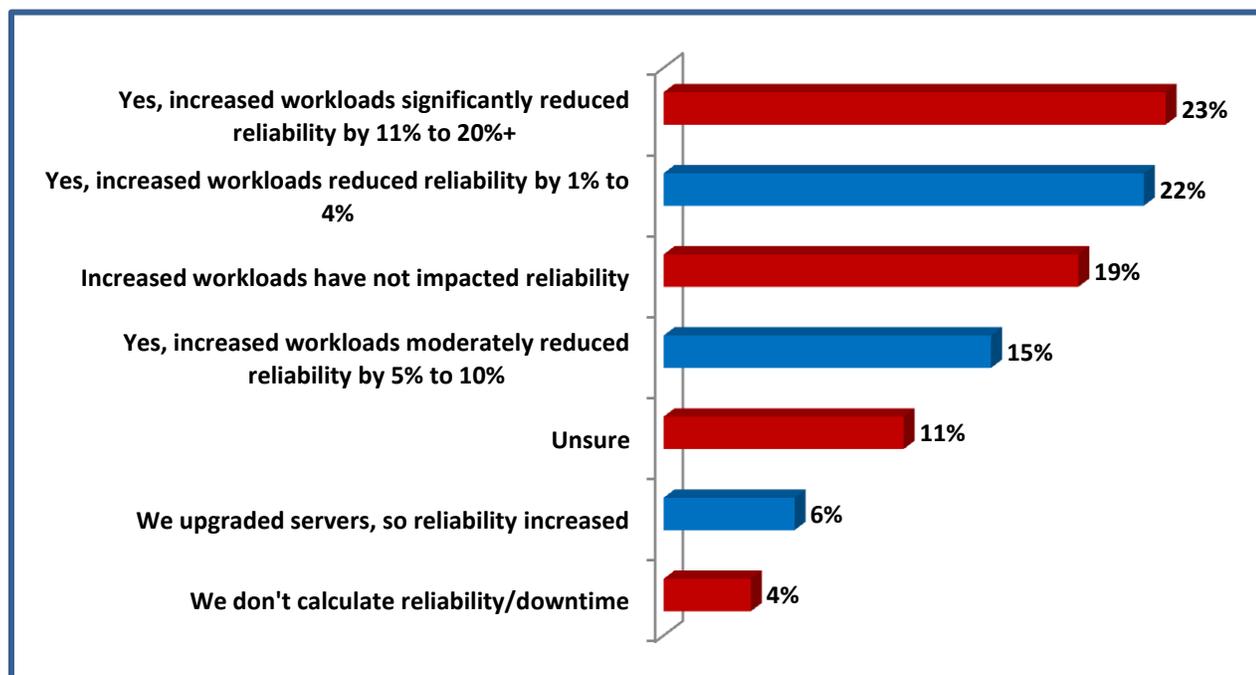


Source: ITIC 2018 Global Server Hardware/Server OS Reliability Survey Mid-Year Update

Six-in-10 Firms Say Workload Increases Negatively Impact Reliability

Another intriguing trend that has both direct and indirect implications for corporate enterprises is the increase in server workloads. Workloads across the network ecosystem from on-premises datacenters, to the cloud, to remote offices and at the network edge have risen by an average of 27% over the last 18 months, according to ITIC’s most recent survey data. The increases range from 14% to as much as 39% among large corporate enterprises that have added compute-intensive applications like AI and Data Analytics or beefed up their storage and/or launched IoT projects. As **Exhibit 5** illustrates, a 60% majority of respondents indicated that increased workloads negatively impact reliability. This is a 15% increase from the 45% of respondents polled in ITIC’s 2017 Global Server Hardware and Server OS Reliability Survey who said reliability suffered when workloads increased. By contrast 19% of respondents said there’s been no impact on reliability and six percent of respondents that said their firms upgraded their servers to accommodate the increased demand so that reliability actually increased.

Exhibit 5. Sixty Percent of Firms Say Higher Workloads Negatively Impact Server Reliability



Source: ITIC 2018 Global Server Hardware/Server OS Reliability Survey Mid-Year Update

Other Survey Highlights and Trends

- Some **71%** of survey respondents said aged hardware (3 ½+ years old) had a negative impact on server uptime and reliability compared with just **16%** that said the older servers had not experienced any declines in reliability or availability. Still, this is an increase of five percentage points from the 66% of those polled who responded positively to that survey question in the ITIC 2017 Reliability Survey and it's a 27% increase from the **44%** who said outmoded hardware negatively impacted uptime in the ITIC 2014 Reliability poll.
- **Across all server distributions 39% of respondents** said the majority of their server hardware was on average two to, three years old; 11% said it was one to two years old; 21% said it was three to four years old and 10% said it was one to two years old. The remaining 20% said their server hardware was aged: four, five or six years old. However, over 80% of respondents who had servers that were over four years old, owned Dell, HPE ProLiant, and Oracle x86 or "White box" servers. For example, over 60% of Dell customers retain their servers for 3 ½ years or longer without upgrading or retrofitting. By contrast, only 17% of IBM Power Edge users and 19% of Lenovo Systems X customers decline to right size or retrofit servers that are 3 ½ years or older to accommodate increased workloads.

Conclusions

Reliability is and will continue to be among the most crucial metrics in the organization. Improvements or declines in reliability and availability will either 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.

The ITIC 2018 Global Server Hardware and Server OS Reliability survey finds that the performance/reliability chasm is widening among the high-end mission critical servers from HPE, Huawei, IBM, Lenovo, Stratus and their less expensive commodity counterparts. This is due to multiple factors. These include the vendors' ongoing investment in reliability, performance, manageability and after-market technical service and support. Additionally, reliability and availability are also determined by the ways in which the corporate enterprise customers deploy, utilize, manage, and secure their servers and the refresh cycles.

On the vendor side: High end mission critical servers: IBM Z, IBM POWER, Lenovo Systems X, HPE Integrity Superdome and Huawei KunLun and FusionServers exhibited the highest levels of server performance, reliability and manageability. And not surprisingly, these server distributions also garnered the highest after market customer satisfaction scores.

On the corporate customer side:

Corporate computing environments continue to increase in size, scope, complexity and connection points. Organizations are also expanding their use of complex technologies like virtualization and migrating key line of business applications to the cloud. Additionally, the scope of networks is expanding via Internet of Things (IoT) and the Network Edge/Perimeter deployments – all of which incorporate mobility and Bring Your Own Device (BYOD) solutions. Additionally, the increased use of more complex, advanced applications like Artificial Intelligence (AI), Data Analytics and Machine Learning, all place further demands on the server hardware and server OS resources.

The result: unplanned downtime is unacceptable and costly. It is an anathema to productivity and business operations. Organizations are more risk averse and have less tolerance for unreliable or unavailable systems and resources. Unplanned downtime for whatever reason will have a domino effect on corporate enterprises, customers, business partners and suppliers. Four nines – 99.99% uptime is now the minimum reliability required by 84% of organizations.

None of the server hardware vendors can rest on their laurels. Competition in the worldwide global server hardware market is cutthroat. It is, and will remain a buyer's market. While many companies, particularly SMBs, will make their purchasing decisions based on price, a significant portion of enterprises are opting to purchase more robust hardware, that come equipped with embedded security, advanced management and

even analytics functionality. In conducting anecdotal interviews for the ITIC 2018 Global Server Hardware, Server OS Reliability Mid-Year Update Survey, many customers volunteered that they are willing to pay a pricing premium for their hardware platforms and also purchase add-on technical support and service to ensure successful first time deployment, longer product lifecycles and continuing reliability and availability.

The survey data shows that corporate enterprises place an extremely high value on after-market vendor technical service and support. They need and expect their vendors to act as true partners who will provide prescriptive guidance. Server hardware and server operating system vendors should be up front and provide their customers with realistic recommendations and prescriptive guidance 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 and notify them of problems or delays in delivering replacement parts.

But partnerships are a two-way street. Organizations bear responsibility for maintaining the health and reliability of their network infrastructure. That means:

- “Right sizing” server hardware to accommodate increased workloads and larger applications.
- Businesses should also regularly replace, retrofit and refresh their server hardware and server operating systems with the necessary patches, updates and security fixes *as needed* to maintain system health.
- Whenever possible, never staying more than two or three revisions behind on server operating systems distributions.
- Installing the latest security patches and fixes.
- Ensure that your firm’s hardware and software vendors and cloud vendors meet or exceed the terms of their Service Level Agreements (SLAs) for agreed upon reliability levels.