



Power Systems Servers Receive Top Reliability Ranking in Latest ITIC Survey

The platform recorded less than 2 minutes of downtime per server/per annum

When it comes to reliability, there's only one direction that matters: up.

For the 11th straight year, IBM Power Systems* servers achieved the highest server reliability rankings in the Information Technology Intelligence Consulting's (ITIC) 2019 Global Server Hardware and Server OS Reliability survey.

ITIC's independent web-based survey polled over 1,000 businesses worldwide from October 2018 through January 2019. It compared the reliability and availability of

18 different mainstream server platforms and a dozen different OS distributions. To obtain the most accurate and unbiased results, ITIC accepted no vendor sponsorship.

Depending on individual implementation, configuration and usage scenarios, the IBM Power Systems platform recorded less than two minutes (1.75 minutes) of per server/per annum downtime, delivering up to 24x better reliability than the least efficient competing hardware.

The reliability of the Power Systems platform trans-


lates into substantial cost savings. As Figure 1, page 31, illustrates, IBM Power Systems users accrued the lowest monetary costs due to inherent server flaws of any platform: \$2,917 versus \$81,683 for unbranded "white box servers," which experienced 49 minutes of per server/per annum downtime.

Other top survey findings include:

- Reliability: POWER8* and POWER9* servers running Linux* distributions were either first or second in every reliability category, including server, virtualization and security for the 11th straight year
- Availability: Power Systems hardware also provided the

✍ **Laura DiDio** is a principal analyst with Information Technology Intelligence Consulting. ITIC conducts primary research and independent surveys on a wide range of topics including server hardware and server OS reliability, security, hourly cost of downtime, Internet of Things, and gender issues in the workplace.

FIGURE 1: COST OF ONE MINUTE OF HOURLY DOWNTIME OF \$100,000 FOR A SINGLE SERVER BY SERVER VENDOR PLATFORM



highest levels of server, applications and service availability. That is, when the servers did experience an outage due to an inherent system flaw, the outages were of the shortest duration—typically one to five minutes.

- **Technical support:** Businesses also gave high marks to IBM technical support. More than 8-in-10 survey respondents rated IBM technical support “excellent,” “very good” or “good.”
- **Hard drive failures:** Faulty hard drives are the chief culprits in inherent server reliability/quality issues (58%) followed by motherboard issues (43%) and processor problems (38%). POWER8 and POWER9 processor-based servers likewise experienced the fewest hard drive quality or failure issues among all of the server distributions within the first three years of service. Only 1% of IBM Power Systems servers experienced technical problems with their hard drives in the first year of usage.
- **External issues:** End-user carelessness (74%), human error (59%) and security (51%) are the top external causes of downtime and unanticipated reasons for taking servers offline
- **Minimum reliability requirements increase:** 85% of corporations now require a minimum of “four nines” of uptime—99.99%—for mission-critical hardware, OSes and main line of business applications. This is a 4% increase from ITIC’s 2017-2018 reliability survey.
- **Patch time increases:** 7 in 10 businesses now devote one to four hours applying patches. This is primarily due to a spike in wide-ranging security issues such as email phishing scams, ransomware and CEO fraud, as well as malware and viruses.
- **Increased server workloads cause reliability declines:** The survey found that reliability declined in 64% of servers more than three and a half years old, when corporations failed to retrofit or upgrade the hardware to accommodate increased workloads and larger, more compute-intensive applications. This is up 19%

Automation and High Availability/Disaster Recovery Go Together

Every high availability (HA) environment is unique. We all have different applications with different requirements. But whether you’re testing a role swap or dealing with a disaster, it’s never wise to rely on manual processes.

Instead, plan for—and practice—an automated process from the beginning. This way, your role swaps are more likely to succeed, no matter whether you’re using PowerHA* high availability or logical replication.

Here’s what I recommend:

1. Use automation to verify that you’re ready for a role swap and continuously confirm that your source and target systems are in sync
2. Use automation to shut down your applications, check resources, end jobs, end subsystems, and then do the opposite on your target system. All HA solutions have an exit program where you can place your shutdown and startup process.

If you have staff executing more than two or three manual processes during a role swap, they’re adding minutes or hours to the swap, while you watch your recovery time objective sail on by.



Tom Huntington

Executive vice president of technical solutions, HelpSystems

Tom heads up the worldwide presales team at HelpSystems and is a four-year IBM Champion for Power Systems.

FIGURE 2: RELIABILITY AND 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 2019 GLOBAL SERVER HARDWARE, SERVER OS RELIABILITY SURVEY

from the 2018 survey. However, it's not a significant issue for IBM Power Systems clients: 76% indicate their businesses have a two-to-three-year upgrade cycle, or they retrofit the servers to accommodate higher workloads.

- Hourly downtime costs rise: 98% of firms say hourly downtime costs exceed \$150,000. Of that figure, 34% of survey respondents say the cost of a single hour of downtime now tops \$1 million.

No Tolerance for Downtime

ITIC's survey findings emphasize that digital age organizations have no tolerance for downtime. In the late 1980s and early 1990s, a metric of two nines or 99% uptime—which equates to nearly 88 hours of per server/per annum downtime—was adequate. Today, even “three nines”—8.76 hours of per server/per annum downtime—is unthinkable (see Figure 2, above).

In 2019, four and five nines—99.99% and 99.999%,

respectively—are the new gold standards of reliability. Downtime has a domino effect: Outages immediately impact the entire network. When servers fail and applications, data and services are unavailable, productivity halts and business ceases, costing thousands or even millions of dollars per minute if it occurs at a critical or peak usage time.

Innovation and High Performance

The recipe for the Power Systems platform's superior reliability and availability is deceptively simple and straightforward.

“IBM mapped out a cogent and consistent set of strategies for its Power Systems server portfolio and they've successfully executed against it,” says Andrew Baker, president of Brainwave Consulting, a technology consulting firm in Gassaway, West Virginia.

That includes regular, planned product releases; an emphasis on innovation such as the ability to support more compute-intensive workloads like analytics, Internet of Things (IoT) and virtualization; and advanced functionality to support emerging technologies like artificial intelligence (AI), blockchain and virtual reality.

“IBM shops realize that it makes more sense both from an economic and business standpoint to invest in robust, reliable servers with advanced functionality that offer the best uptime and scalability, than to shave a few dollars off list price.”

—Andrew Baker, president, Brainwave Consulting

IBM's technical support has been a bastion of stability. ITIC anecdotal client interviews found that IT managers were pleased with both the quality and speed with which IBM technical support responded to issues.

"IBM Power Systems work; they just work. I can't remember the last time we had a server crash due to system flaws, hard disk or memory issues. In the financial services industry, bullet-proof reliability is a necessity. Power Systems are robust and economical," says an IT manager at a large New York bank. "We rarely have to call IBM, tech support, but when we do, they're extremely helpful and responsive."

An IT architect at a large technology services provider says, "IBM Power Systems running the AIX* OS and PowerVM* deliver excellent performance, advanced features and high reliability. These are critical for our business." The Delhi, India, firm has more than 300 servers in its data center. "IBM Power Systems servers are leading edge. We want IBM to continue to deliver advanced capabilities, high reliability and security, so it retains its No. 1 spot as the server technology leader," the IT architect adds.

Finally, Baker notes Power Systems shops are less price sensitive than organizations that deploy inexpensive, commodity servers. "IBM clients aren't penny-wise and pound foolish," Baker says. "They want good deals, but they won't cut corners and sacrifice quality to get a lower list price. IBM shops realize that it makes more sense both from an economic and business standpoint to invest in robust, reliable servers with advanced functionality that offer the best uptime and scalability, than to shave a few dollars off list price."

Another plus: The majority of IBM clients adhere to a regular

three-year upgrade cycle and retrofit their Power Systems hardware as needed. This is crucial because complex AI, analytics, IoT and virtual reality applications are larger and more compute-intensive.

Power Systems Advantages

IBM Power Systems servers have consistently delivered and maintained the highest levels of uptime availability based on every metric and measure of reliability since

SPONSORED CONTENT



No-Code Modernization 101

Thursday, June 13 | 1 ET / 12 CT

From Green Screens to a New Digital UX That Drives Efficiency and Automation

It's no surprise that applications with poor usability hinder customer-facing employees and negatively impact your customer relationships. Many organizations find themselves struggling with applications on traditional platforms—mainframe or IBM i—which still successfully underpin business operations but lack modern user interfaces and workflows.

Now, advances in no-code technology and APIs are accelerating modernization and allowing organizations to turn older green-screen applications into new digital web and mobile user experiences that improve employee productivity, enable automation and increase operational efficiency.

FEATURING:



Charles Jones
Senior Sales Engineer
Rocket Software

SPECIAL GUEST:

Minesh Manilal
Head of Integration,
Platforms and DevOps
Prolifics, Inc.



Register Today: webcasts.com/ibmsystemsmag

ITIC's first reliability poll in 2008. And they continue to do so today.

And 88% of IBM Power Systems enterprises running Red Hat Enterprise Linux, SUSE Linux or the Ubuntu open-source distribution experience fewer than one unplanned outage per server/per year due to bugs or flaws in the OS.

This is a boon for IBM enterprise clients. High availability ensures uninterrupted productivity; supports the business' bottom line; strengthens security and compliance; and mitigates risk.

Optimize Uptime and Availability

Organizations have every right to expect server vendors to continually improve hardware reliability and deliver a minimum of 99.99% uptime. Vendors must also quickly address and resolve technical issues and security flaws when they arise to deliver advanced features/functions, and provide the necessary guidance and top support.

Organizations' ability to achieve 99.99-99.9999% uptime is a two-way street, not a one-way foot path. Enterprises bear the responsibility of selecting specific servers in the configuration that best serves their particular business and budgetary needs. They must also employ skilled IT and security administrators who can properly provision, upgrade and maintain a high degree of daily operational efficiency.

To optimize uptime and reliability, ITIC advises organizations to:

- Regularly analyze and review configurations, usage and performance levels. This will enable companies to determine whether the current server and server OS environment allows them to achieve optimal reliability. Being cognizant of specific uptime and reliability statistics will

DOWNTIME CAUSES

The top external causes of **downtime** and unanticipated reasons for taking servers offline

End-user Carelessness **74%**

Human Error

59%

Security

51%

enable the business and its IT department to identify baseline metrics associated with all of their individual platforms. It will also provide companies with an accurate assessment of the inherent reliability and flaws in their hardware and software. They can then compare and contrast that with downtime resulting from other issues such as integration and interoperability; lack of readily available patches or fixes; problems with ISPs and carriers, and unpredictable

or unavoidable outages due to natural or manmade disasters.

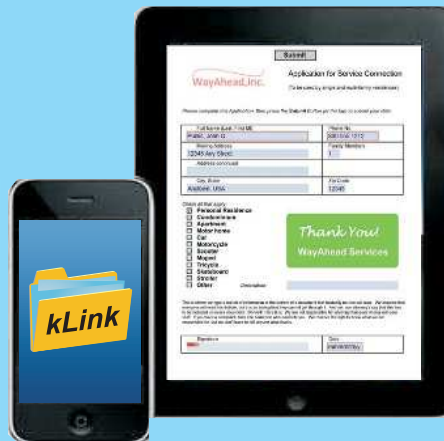
- Maintain detailed records on unplanned and planned downtime. IT departments should compile a detailed list of outages. Include facts like the cause of the outage (e.g., hard drive failure, human error, manmade disaster, etc.); the duration of downtime; the severity of the event (e.g. lost, damaged or stolen data; interrupted transactions). Be sure to log the remediation efforts including how many

users were affected and how many IT managers were involved in restoring the servers. Also, chronicle the speed and satisfaction of your vendor's response, if you called for technical support (phone or on-site).

- Keep tabs on security. Hacks are more pervasive and the hackers are more proficient and very well organized. Increasing connections and more end points, mean more potential vulnerabilities. Install the latest security patches and fixes and work closely with vendors to solicit guidance and inform them when issues arise.
- Don't delay updates. Refresh and upgrade hardware as needed to accommodate more data intensive and virtualized workloads. The server and OS are inextricably linked. To achieve optimal performance, corporations must ensure that the server hardware is robust enough to carry both the current and anticipated workloads. Applications are getting larger. 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. According to the survey, 49% of organizations don't take this step. A business that fails to calculate hourly downtime has no ability to accurately measure the server's reliability and cost effectiveness.

The continuing high reliability of the IBM Power Systems platform yields demonstrably better uptime, providing clients with improved performance and economies of scale, and lower TCO than competing hardware platforms. 🐼


Modern Software for IBM i



- Online Server
- Interactive PDF's
- Secure File Transfer
- Full Color Overlays
- Burst Spooled Files
- Email/Fax/Archive

KeyesMail • KeyesFax • KeyesOverlay
KeyesPDF • KeyesArchive • kLink

Computer Keyes
Free 30 Day Trial
Free Technical Support



www.ckeyes.com
info@ckeyes.com
USA: 800 356 0203
Tel: 425 776 6443



OCEAN TechCon19
July 25-27, 2019
Costa Mesa, California

World-Class Training by World-Class Speakers

Premier Vendor Solutions Expo

Exceptional Value

Over 50 Educational Sessions

oceanusergroup.org

#OCEANTechCon19