

Revitalize Your Aging Datacenter – the Real Value of Datacenter Modernization

Commissioned by



Introduction

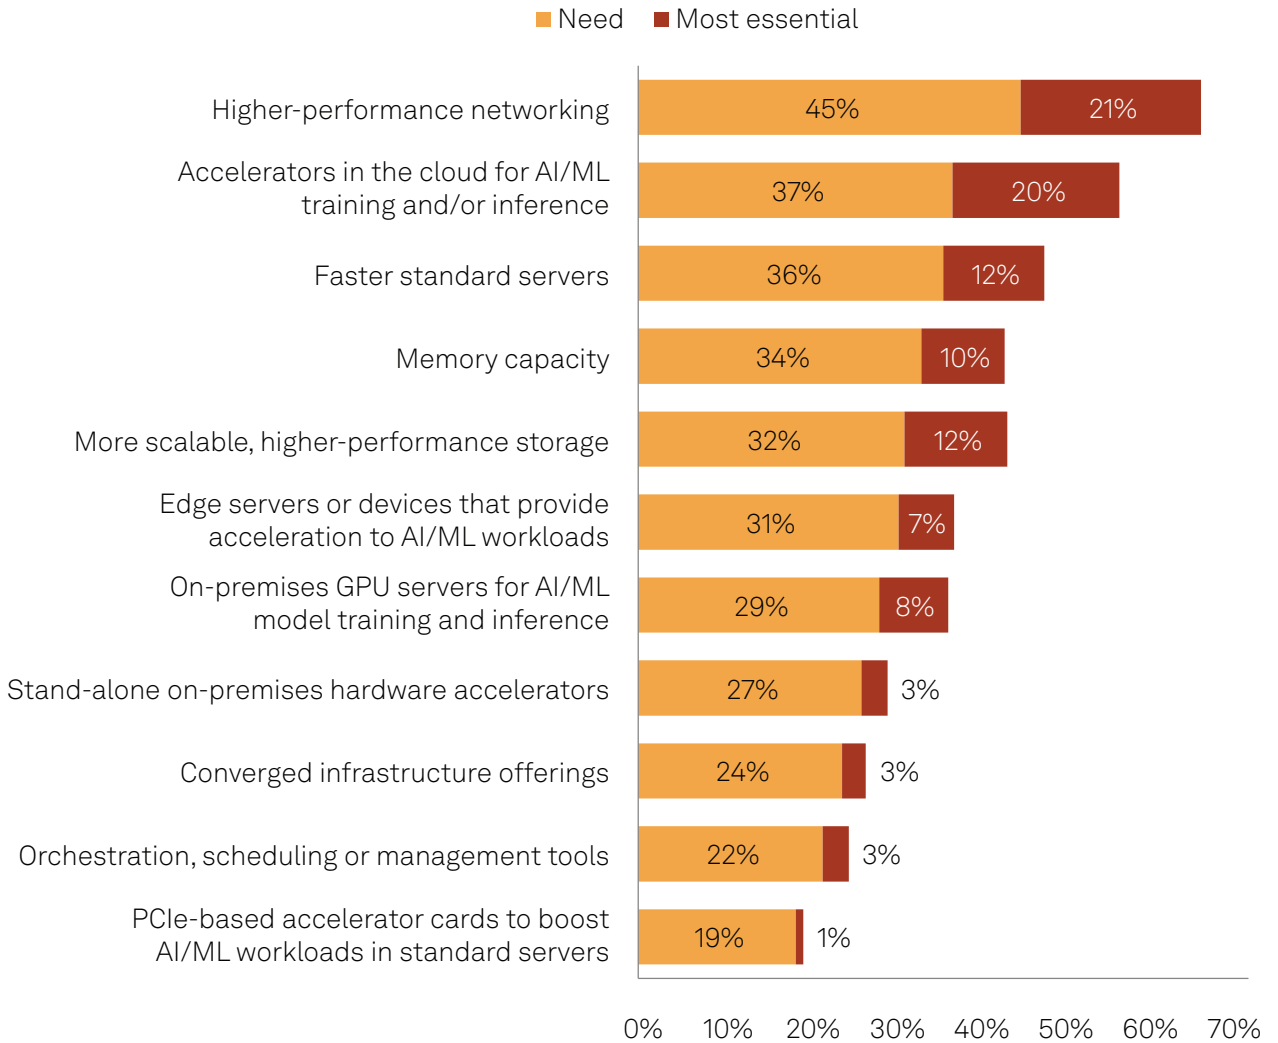
The exponential increase in application support demands for new technologies such as data visualization, augmented reality/virtual reality (AR/VR) tools and generative AI (GenAI) is driving today's datacenter modernization requirements. Simply maintaining the datacenter with software updates will not be adequate to support the new compute, network and latency requirements for these types of applications. Postponing modernization could have a significant business cost by inhibiting the enterprise's ability to rapidly deploy and scale applications and pursue new product and service innovations in response to business needs. Revitalizing an aging datacenter via IT infrastructure modernization not only can provide substantial business benefits to enable new applications, but also cost savings through reducing operational power and cooling costs along with the physical footprint.

By migrating to more modern servers, datacenters can decrease the number they need to service workloads, enabling consolidation into a smaller physical footprint. While some organizations may find it challenging to justify the investment needed to modernize their IT estates, inaction can have its own monetary repercussions in the form of sustainability and business opportunity costs.

The Take

Application performance demands from new AI and machine learning (ML) workloads and the need for increased efficiency and sustainability to support all workloads are driving datacenter modernization. The benefits of server modernization span processor performance gains, improved data access speeds with faster memory and bus architectures, and network throughput — all of which are necessary to support new power-hungry applications. Almost half (48%) of respondents to 451 Research's Voice of the Enterprise: AI & Machine Learning, Infrastructure 2023 survey say they need faster standard servers — or that they are the most essential infrastructure resource — to improve AI/ML production workload performance (see Figure 1).

Figure 1: Infrastructure requirements for AI/ML production performance improvement



Q. Which of the following infrastructure resources does your organization need to improve the performance of its AI/ML production workloads? Please select all that apply.


Base: All respondents, abbreviated fielding (n=434).

Q. And which of the following infrastructure resources is most essential to improve the performance of your organization's AI/ML production workloads?

Base: Organization needs infrastructure resources to improve the performance of its AI/ML production workloads, abbreviated fielding (n=425).

Source: 451 Research's Voice of the Enterprise: AI & Machine Learning, Infrastructure 2023.

Modern servers typically consume less power per unit of work and have a better thermal design with reduced cooling requirements compared to prior generations. They also support more robust management and automated administration and have enhanced security features and greater resiliency and reliability. Modern servers can handle more demanding workloads, enabling datacenters to consolidate applications and services onto fewer physical devices, thus leading to improved scalability, more efficient facility utilization and decreased operating costs.



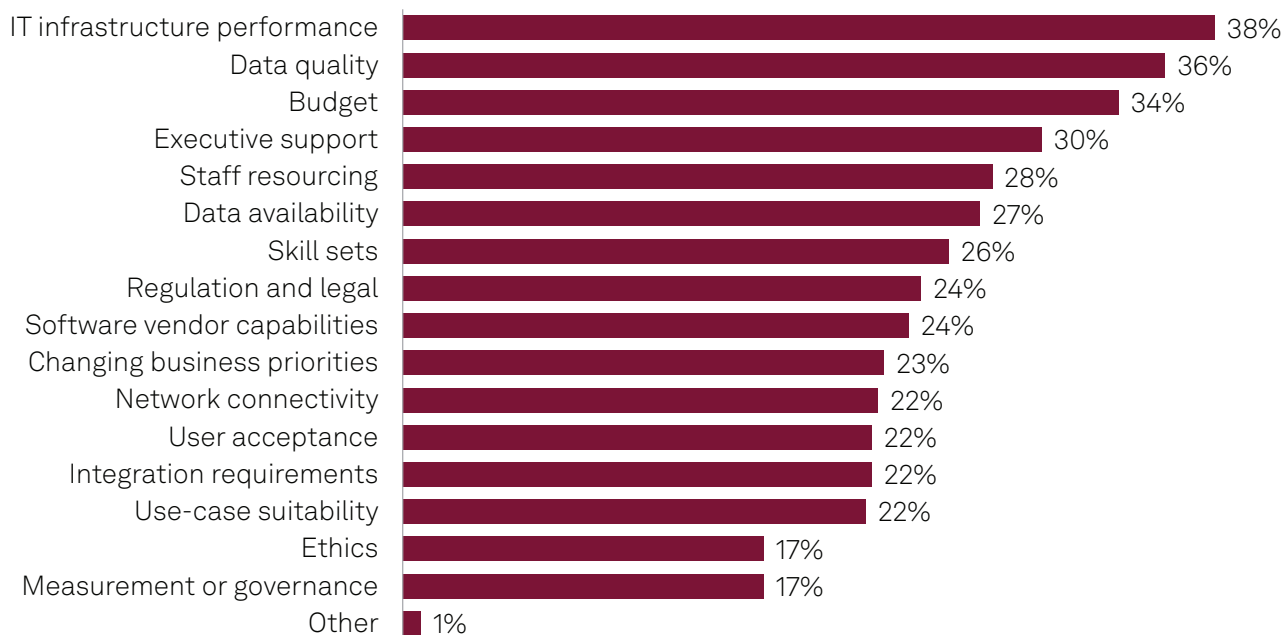
While many applications are deployed across hybrid private/public clouds, some cannot be moved off-premises due to regulatory compliance requirements, data transfer size and volatility, and minimum latency demands. To take advantage of new server architectures, these legacy applications will have to be moved to modern processors within the datacenter. Modernizing to a new processor architecture could entail platform migration, especially in virtualized environments; however, migration tools help lessen the costs and effort to migrate to modern servers. The up-front investment to modernize servers can be significant, but the long-term benefits — improved performance, efficiency and sustainability gains, and reduced maintenance and operating costs — will likely justify the expense of modernizing.

New application requirements call for a datacenter revitalization strategy

As new enterprise application demands collide with existing IT infrastructure limitations, many organizations will have to address infrastructure deficiencies to service these new resource-intensive requirements. The rapid rise of GenAI and other forms of AI enablement will place greater pressure on IT architectures that are already struggling to service existing workloads. The scale and diversity requirements of GenAI are outstripping the ability for many organizations to service these projects, and organizations will not be able to meet the challenge by simply accessing additional graphics processing units (GPUs). Companies wanting to meaningfully expand AI initiatives will have to construct a comprehensive strategy to address infrastructure processing power, storage and networking bottlenecks.

The majority of respondents (61%) to the same Voice of the Enterprise AI/ML survey mentioned above say that infrastructure limitations are preventing, or will prevent, more frequent retraining of models in production. Only a third of respondents (32%) believe their organization's IT environment is always able to meet existing AI/ML workload demands. On average, respondent organizations abandon more than a third of their AI projects, and IT infrastructure is most frequently cited as the reason (see Figure 2). Most organizations (82%) predict that AI workload demands will increase over the next two years, and just 33% of respondents believe their organization's existing IT environment will be able to meet that demand without upgrades. These workloads are creating challenges because they require faster servers capable of supporting higher-performance storage, greater memory speed/capacity and higher-bandwidth networking. As datacenters move to support large language models (LLMs) in production, the training datasets will become extremely large, often exceeding a petabyte of data.

Figure 2: IT infrastructure is the top driver of AI project abandonment



Q. In which of these areas did challenges lead to project abandonment at your organization within the past 12 months? Please select as many as apply.

Base: All respondents (n=686).

Source: 451 Research's Voice of the Enterprise: AI & Machine Learning, Infrastructure 2023.

Modern servers are also better suited for virtualized workloads and cloud integration, enabling datacenters to take greater advantage of virtualization for more efficient resource allocation, operational flexibility and server management. Modernizing to a new processor architecture may entail platform-migration efforts, especially in virtualized environments. The cost and effort to migrate to modern servers is lessened by migration tools, which enable moving from one processor architecture to another. Architecture-agnostic and open-source migration tools are available, allowing anyone to leverage the technology.

Consolidation advantages

Server modernization provides the advantages of increasing compute capacity in a given physical footprint while decreasing the power and cooling requirement per compute cycle, thereby decreasing overall physical space and operational costs for current datacenter workloads. This increase in compute capacity and reduction in power and cooling usage also enables consolidation of new and existing workloads within current datacenter facilities. Consolidation addresses the challenge of finding colocation and/or lease space, which is difficult in several geographic regions. As datacenters pivot to support new applications such as GenAI, they can add capacity without increasing facility footprint.

Another benefit of moving to a modern server architecture is the ability to standardize on up-to-date security, device management and DevOps models. Datacenters must often deal with multiple generations of incompatible technologies, requiring significant effort to custom code device support and address legacy security holes. This may also complicate new application deployment and management if older platforms are inadequate to support modern application requirements.

The cost of inaction — why modernization may cost less than doing nothing

While server modernization does entail an initial investment, there are also costs — beyond operational expenses — to remaining on legacy hardware deployments. When weighing whether to invest in server modernization, organizations should consider the full spectrum of new technology advantages. The reduction in physical footprint for equivalent or even greater processing capacity creates obvious facility savings, and this is accompanied by more efficient cooling and a reduction in power per compute cycle as well. Organizations could think of the recurring cost savings as a budget increase and use it to enhance existing products, services and technology innovation projects or internal system improvements.

Energy efficiency advantages of modern infrastructure

Nearly all (96%) organizations that own or operate datacenter facilities indicate that datacenter efficiency and sustainability is important or very important, according to 451 Research's Voice of the Enterprise: Datacenters, Sustainability 2023 survey. Modernizing datacenter infrastructure provides multiple energy-efficiency advantages. Many outdated enterprise datacenters use as much energy on cooling and operations as on powering their IT equipment. Modern servers are designed to leverage energy-efficient components and technologies, such as low-power processors, more effective processor cooling designs and advanced power management features. Coupled with intelligent processor load-based power management, this results in reduced power consumption and cooling requirements and, consequently, a reduced carbon footprint.

Modern servers are also forecast to have a longer life span than previous generations, and they are designed and constructed with more recyclable materials, resulting in less electronic waste, again helping to reduce environmental impact. These advantages go unrecognized as long as datacenters remain on older, less-efficient platforms.

Conclusion

Datacenter operators are continuously seeking the right technology and architecture to improve efficiency, reliability and sustainability, as well as reduce costs and the environmental impact of their facilities. The challenge for businesses is not just upgrading their infrastructure to address existing workloads more efficiently, but to ensure they design and deploy an architecture that will be able to handle the ever-expanding workload requirements of projects. Modernization brings immediate operating cost reductions, physical datacenter capacity extension through consolidation and significant energy-efficiency advantages. So, even though the initial investment in modernizing servers can be significant, these benefits, coupled with the overwhelming inbound demands of GenAI on IT infrastructure, suggest there is a compelling business case to invest in modernizing aging datacenters.



AMD addresses the needs of enterprises with a complete portfolio of innovative and proven data center products. A unique combination of security features, flexibility, and outstanding TCO supports your applications and platforms in the cloud, on-premises, and at the edge. An industry leader, AMD brings numerous strengths to help companies achieve their business goals: close engagement with customers to develop real-world solutions; a vibrant partner ecosystem offering; a well-defined roadmap; and stability and consistent growth to power advanced R&D. Discover more about [AMD EPYC™ processors](#) here.

About the authors



Dan Thompson

Principal Research Analyst, Datacenter Services & Infrastructure

Dan is a principal research analyst within the Technology, Media and Telecommunications group of S&P Global Market Intelligence leading the Datacenter Services and Infrastructure (DSI) team. The DSI team is charged with keeping tabs on the datacenter industry globally to better understand its trends and growth areas. His research includes analyses of datacenter providers, market size and supply/demand in key and emerging markets around the world. Dan provides coverage of datacenter providers offering services beyond colocation, such as managed and cloud-type services. Dan also provides research on the sustainability of the datacenter industry. Beyond just renewable energy purchasing and carbon offsetting, he has been investigating full life cycle emissions, including supply chain emissions, as well as efficiency gains and water usage.



Bobby Clay

Senior Consulting Analyst, Manager, TMT Consulting

Bobby is a senior consulting analyst/manager within the S&P Global Market Intelligence Technology, Media and Telecommunications Research group. He is focused on cloud, datacenter, generative AI and metaverse technologies and has a unique mix of technical and business skills with more than 30 years of industry experience. Bobby's recent areas of concentration include hyperscalers, generative AI cloud tools, metaverse virtual world and avatar construction technologies, and datacenter modernization and sustainability.

About S&P Global Market Intelligence

At S&P Global Market Intelligence, we understand the importance of accurate, deep and insightful information. Our team of experts delivers unrivaled insights and leading data and technology solutions, partnering with customers to expand their perspective, operate with confidence, and make decisions with conviction.

S&P Global Market Intelligence is a division of S&P Global (NYSE: SPGI). S&P Global is the world's foremost provider of credit ratings, benchmarks, analytics and workflow solutions in the global capital, commodity and automotive markets. With every one of our offerings, we help many of the world's leading organizations navigate the economic landscape so they can plan for tomorrow, today. For more information, visit www.spglobal.com/marketintelligence.

CONTACTS

Americas: +1 800 447 2273

Japan: +81 3 6262 1887

Asia Pacific: +60 4 291 3600

Europe, Middle East, Africa: +44 (0) 134 432 8300

www.spglobal.com/marketintelligence

www.spglobal.com/en/enterprise/about/contact-us.html

Copyright © 2023 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global keeps certain activities of its divisions separate from each other to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.