An IDC White Paper, sponsored by Dell Technologies

IDC

YOUR INNOVATION ENGINE:

Fit-for-Purpose IT Infrastructure for Digitally Determined Organizations

RESEARCH BY:

Ashish Nadkarni
Group Vice President, Infrastructure Systems, Platforms and Technologies Group, IDC

Heather West, Ph.D.
Senior Research Analyst, Infrastructure Systems, Platforms and Technologies Group, IDC

May 2021 | IDC Doc. #US47628621

Navigating this White Paper

Click on titles or page numbers to navigate to each section.

IDC Opinion 3

Situation Overview 4

The Digitally Determined Invest in Fit-for-Purpose Infrastructure 5

Server Platform Selection 7

Processor Platforms Are a Crucial Building Block for Infrastructure 8

Investing in a Trusted Vendor 8

Choosing Dell Servers with AMD Third-Generation EPYC Processors 9

Advanced Micro Devices (AMD) EPYC Processor Family 9

Dell Is a Trusted Vendor in Datacenter Infrastructure 11

Dell AMD Third-Generation EPYC-Based Servers 12

Dell Project Apex: Flexible/as-a-Service Option for Enterprises 13

Benefits of Adopting Dell AMD Third-Generation EPYC-Based Servers 14

Adoption Drivers 14

Service-Quality Improvements 16

Business Benefits 18

Financial Improvements 19

Hardware/Software Optimization 20

Essential Guidance for IT Buyers 22

Taking a Systematic Approach for Infrastructure 23

Step 1: Profiling key workloads 23

Step 2A: Matching workload requirements to infrastructure (the platform) 24

Step 2B: Matching key workload requirements to the platform (the processor) 27

A Shift to Operational Expenses: Flexible Consumption Options 28

All Servers and Server Vendors Are Not the Same 29

The Importance of Server Security 30

Challenges and Opportunities for Dell 31

Conclusion 33

Appendix 34

Glossary 34

About the Analysts 35

IDC Opinion

Digital transformation is upon us.

CIOs and IT leaders everywhere must guide their organizations through this change while enabling their businesses to thrive in the new digital economy—one that has complex channels and new business models that shape new revenue-generating products and services. This process of transformation thrusts the role of digital infrastructure forward, prioritizing it as a strategic investment.

This means that CIOs and IT leaders must make conscious and forward-looking choices that enable their organization's digital infrastructure to be:

- Scalable, elastic, and agile from a performance and capacity perspective, and secure and compliant from a data management perspective
- Capable of meeting stated objectives from a cost-of-ownership and return-on-investment perspective.
- Consumed as software-defined resources (e.g., infrastructure as code) supporting newer (e.g., cloud-native) workloads and newer development processes and operations methodologies (e.g., DevOps) and using newer pay-as-you-go consumption models.

Investment criteria in digital infrastructure assets up until now have focused on driving homogeneity in the datacenter. A key benefit of this is that it delivers people and process scaling and efficiency while reducing the technology drag (i.e., pseudo-automation). However, a key disadvantage is that homogeneity makes the infrastructure very inflexible: It cannot scale to meet the demands of newer workloads, nor can it be optimized to meet the requirements of newer development and operating paradigms. In fact, over half (56%) of organizations looking to procure new AMD servers are doing so to improve scalability for certain workloads (*Dell AMD Buyer Behavior Survey,* IDC, October 2020).

Digital transformation and the ensuing infrastructure investment strategy demands a different approach. Divergent workload requirements demand a shift away from homogeneity and a one-size-fits-all approach. This shift opens the door for IT organizations to examine newer platforms and systems for their compute, storage, and networking needs. The evaluation of processor subsystems, for example, is now table stakes when making the selection of new platforms and systems. In a study conducted by IDC in 2020, respondents indicated that they plan to increase their investments in heterogeneous infrastructure in the coming years, to be prepared for the onslaught of digital transformation initiatives. The same study found that most customers were willing to consider the processor platform when procuring new infrastructure *(Computing Platforms Trends Survey, 1Q20,* IDC #US46413220, May 2020).

The third-generation EPYC processors from Advanced Micro Devices (AMD), code-named "Milan," represent a unique opportunity for addressing the modern workload-optimized infrastructure world. Several OEM vendors are offering computing platforms (aka servers) with AMD EPYC processors. Identifying the server platform that takes advantage of this opportunity is not always straightforward. IT organizations must make a comprehensive evaluation of the options and ensure that the server platform they invest in can meet their future business requirements for flexibility, autonomy, and security.

Situation Overview

A digitally transformed business is a technology and data business.

IDC refers to businesses that have a vision and strategy for digital transformation and are farther along in implementing this transformation as "digitally determined" businesses. For such businesses, key decisions are made entirely on insights gained from timely analysis of diverse data sets, stitched together by various artificial intelligence as well as machine learning and analytics technologies. Considered a necessary element for business success, data is often referred to as "the new oil." Like oil, data needs to be refined in order to be transformed into a usable state or product. For example, data collected at various edge locations often needs to be refined, processed, and analyzed so it can be used to develop new insights for the business.

To protect the digital asset value, data and infrastructure security is the top IT infrastructure-related investment that organizations make, followed by datacenter IT compute and storage infrastructure (illustrated in **Figure 1**, next page). Other key IT infrastructure-related investments include mobility, applications, organizational skills, cloud, and data analytics. IDC confirms these findings to be consistent year over year, emphasizing the importance of IT infrastructure to digital transformation. It is no longer an afterthought but rather a strategic investment area.

FIGURE 1

Top IT-Related DX Investment Priorities

(% of respondents)

Q. Which of the following IT transformation initiative areas were the largest investments and/or changes made in response to the digital transformation (DX) project?

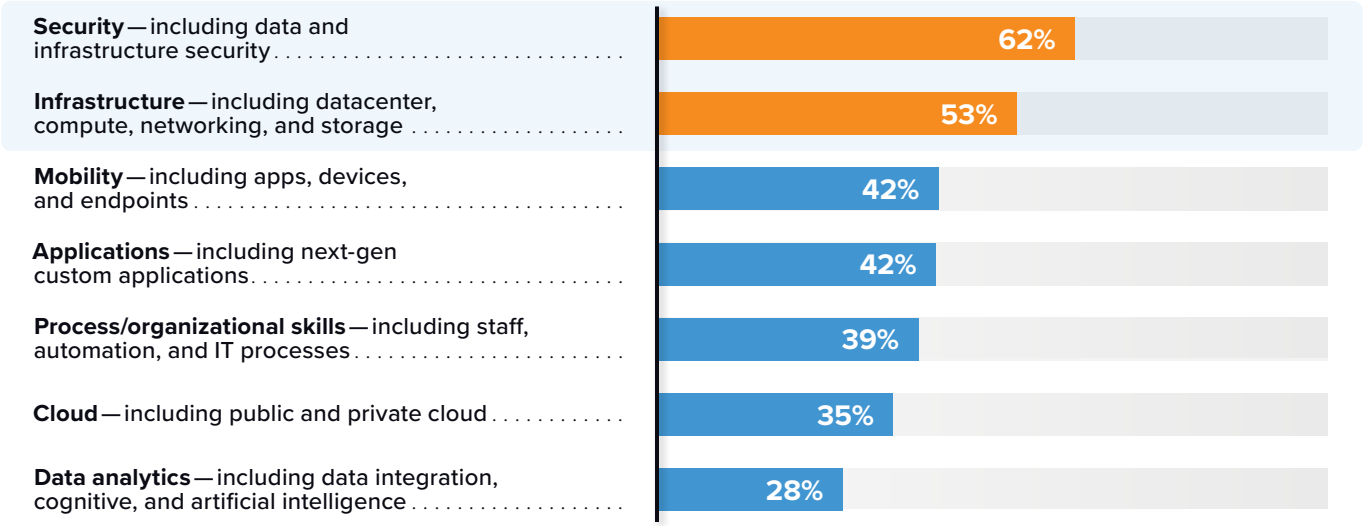

Note: Multiple dichotomous table—total will not sum to 100%
n = 304, Base = All respondents
Source: IDC *DX Datacenter Study—Web Survey 2017,* 2017

The Digitally Determined Invest in Fit-for-Purpose Infrastructure

Digitally determined businesses rely on a combination of current- and new-generation workloads (applications and their associated data sets). Each workload has its own unique infrastructure requirements. With most current-generation workloads being virtualized or run on bare metal, it is assumed that the infrastructure will handle operational characteristics such as resiliency, performance, and scalability. Comparatively, new-generation (or cloud-native) workloads often run inside containers, are function-driven, or are deployed in a microservices architecture. New-generation workloads account for operational characteristics in their own intrinsic design and can adapt well to features and optimizations provided by newer hardware platforms, including low-level performance and security features.

The IT infrastructure for a digitally determined business must support a varying combination of current- and new-generation workloads. Businesses increase investments in cloud-native workloads as they accelerate their transformation. In the process, businesses reduce their reliance on current-generation workloads.

IT organizations must scale their infrastructure accordingly, by investing in fit-for-purpose infrastructure while maintaining consistency and without creating too many islands or siloes. In other words, fit-for-purpose infrastructure lends itself well to divergent software requirements—which are commonplace in most digital businesses today.

A fit-for-purpose IT infrastructure provides a solid foundation on which firms can accelerate their digital journey. From the perspective of IT organizations, it allows them to shift to strategic enablers of their vision laid out by their businesses to differentiate themselves. The key benefits to CIOs and enterprise architects are illustrated in **Figure 2**.

FIGURE 2

Benefits of Fit-for-Purpose Infrastructure

For CIOs	For Enterprise Architects
Delivers consistent service quality and the backing of services provided by a trusted vendor.	Provides optimized scaling of performance and capacity independent of each other.
Maximizes their return on IT investments while reducing total cost of ownership of IT infrastructure assets.	Scales in terms of bandwidth to support mixed performance-intensive and business-critical workloads. Enables effortless systems management.
Assures security of data and intellectual property at the infrastructure hardware and software platform layer.	Delivers a secure infrastructure foundation with built-in security that starts at the processor level.

Source: IDC, 2021

A fit-for-purpose infrastructure means using best-of-breed components to build a heterogeneous infrastructure. Starting from the components that make up servers, it delivers on three key elements:

- **Adaptive compute:**
  Address evolving compute demands with a platform engineered to optimize the latest technology advances across processor, memory, networking, storage, and acceleration. It delivers infrastructure optimized for performance-intensive computing (AI/ML, high-performance computing, and Big Data/analytics) initiatives.
- **Autonomous compute infrastructure:**
  Respond rapidly to business opportunities with intelligent systems that work together and independently, delivering to the parameters that are set by organizations. APIs and integration tools provide a seamless setup for autonomous compute while defining automation levels and strategic workstreams to improve datacenter operations.

- **Proactive resilience:**
  Embed trust into digital transformation with an infrastructure designed for secure interactions and the capability to predict potential threats. Organizations can operate workloads on a cyber-resilient platform using a cryptographically trusted booting cycle and immutable silicon root of trust, starting at the vendor's factory.

Such an approach differs from the de facto one-size-fits-all stance many organizations take today. The latter can potentially slow down the pace of business innovation due to limitations in performance and scalability. It is inefficient when it comes to provisioning and re-provisioning hardware. A one-size-fits-all approach cannot fully accommodate divergent software requirements.

Finally, an aging infrastructure can quickly become a liability for an organization. Infrastructure operating costs spiral upward if businesses do not upgrade in a timely manner. Further, aging infrastructure platforms are vulnerable to data breaches. Recent disclosures have revealed that vulnerabilities whose root lies in the processor design are difficult to patch without incurring heavy penalties.

IT organizations enable their business units to produce consistent outcomes by making timely investments in workload-optimized and innovation-packed infrastructure. They rely on a trusted vendor as a partner to deliver a fully optimized and secure platform.

Server Platform Selection

A modern server infrastructure enables organizations to host strategic and/or mission-critical workloads where they are needed most: on-premises (part of the hybrid cloud infrastructure). For others it is a matter of meeting stringent regulatory mandates by choosing between the ability to anchor their infrastructure on premises or control their consumption of public cloud resources using an on-premises control plane. Organizations seek a cost-effective, efficient, and scalable in-house modernized infrastructure. They are keen to refresh their infrastructure using a flexible/as-a-service pricing model to align expenditures more closely with actual usage and/or consumption.

Server technologies are foundational. They influence the outcomes of other infrastructure investments. IDC found that organizations consider cybersecurity capabilities, cloud-based data protection capabilities for cloud-based workloads, and accelerated compute and solid-state storage as the top technology investments to leverage in a modernized infrastructure. These criteria require a robust compute infrastructure that delivers consistent service quality to the workloads running on it.

Examples of technical specifications of servers for optimal workload placement include:

- Form factors and rack density
- Memory speed and capacity
- CPU performance, core density, accelerator density
- Local storage capacity and performance (I/O latency and bandwidth)

Server form factors determine the performance and availability characteristics of that server and hence are crucial to workload optimization. Organizations prefer servers with built-in failover capabilities for their strategic workloads—workloads they consider to be essential for their business mandate. Some organizations deploy standalone blade or multinode servers for their strategic workloads, while others deploy them on single node servers. IDC's research finds that single node and rack optimized form factors are best suited for scale-out workloads that use cloud-native application approaches.

Processor Platforms Are a Crucial Building Block for Infrastructure

Over the last 15 years, datacenter homogenization eliminated the make of the processor as a key performance criterion. While organizations acknowledged the processor as a crucial building block for an infrastructure platform, the options were limited to "fast, faster, and fastest"—causing the make and model never to be in play. This situation is changing rapidly, and organizations now have a choice to make: Stick with the status quo or make processor selection a part of the platform investment process.

Now organizations inclined to change can make a more informed decision and simultaneously avert impact on business outcomes due to processor shortcomings. Businesses acknowledge that switching out processors in an existing server platform is arduous. The selection must happen during initial investments.

Investing in a Trusted Vendor

Firms consider their IT vendors as partners in their digital transformation journey. The level of trust that firms place on an infrastructure vendor is directly related to the vendor's ability to build and deliver a secure server platform.

The National Institute of Standards and Technology defines a vendor's cyber-resilient architecture using four key attributes:

- **Built-in versus bolted-on security:**
  Security should not be an afterthought, but rather an integral part of the overall server design. The vendor builds deep protection into the server architecture, providing end-to-end security for the datacenter from server deployment through retirement.
- **Security by design:**
  Server security should be designed to protect, detect, and recover data across your entire system. The servers include powerful features like a cryptographic silicon root of trust that authenticates BIOS and firmware during the server boot process.

- **Server integrity from silicon to system:**
  Whether caused by malware, unauthorized changes, or failed patches, issues that threaten system integrity must be detected and managed quickly. The vendor continuously monitors the system to detect, log, and send alerts about changes in configuration parameters across the firmware.
- **Rapid recovery:**
  If an attack does change a component outside of the operating system environment, the right tools for rapid recovery are essential. The vendor includes a built-in automatic recovery feature that restores the server to a trusted state, even if the attack is made against the BIOS, firmware, signed drivers, or any signed components.

Choosing Dell Servers with AMD Third-Generation EPYC Processors

Advanced Micro Devices (AMD) EPYC Processor Family

Organizations often worry that changing a processor type will be very disruptive to their infrastructure. That is a fair point if they are in fact opting for processors with a different instruction set architecture (ISA). An ISA is a computation model implemented in a central processing unit. IT ISA defines the supported data types, the registers, the hardware support for managing main memory, fundamental features such as the memory consistency, addressing modes, virtual memory, and the input/output model of a family of implementations of the ISA. (AMD EPYC-family processors use an ISA known as x86 that is commonly used in datacenter infrastructure.)

The launch of the EPYC processor family in mid-2017 signaled a renewed focus from AMD on the primary datacenter server market, which IDC believes crossed $100 billion in size in 2020. EPYC is based on the x86-based Zen microarchitecture, which is designed to optimize for core performance and power consumption. EPYC processors are available in -8 to 64+ core SKUs, with support for single- and dual-socket configurations.

Key value propositions and differentiators that EPYC introduced into the public cloud services market were:

- **Flexibility to configure socket count:**
  With their high core count and full support for single- and dual-socket motherboards, EPYC processors give operators the flexibility to optimize their server built to the specific workload or customer need. The scalable I/O design and parallel PCIe bus lanes for each core allow the single-socket configuration to meet compute requirements for several workloads that typically use two sockets to meet scale-out needs.
- **Performance-intensive computing applicability:**
  EPYC's high core density and I/O bus design also make it optimal for certain performance-intensive workloads—in particular, implementations with highly parallel execution with low interaction between the tasks. These are applications for which (in the current compute landscape) accelerator-enabled options are an overkill or require a lot of refactoring, while a normal compute VM or bare metal environment is not cost-optimal.
- **Embedded security capability:**
  The growing enterprise reliance on public cloud has increased the focus on security. The growing investments in Trusted Platform Module (TPM) and recent security capability announcements from Google and Microsoft highlight the importance of hardware-enabled security. EPYC's built-in security processor allows full memory encryption with no application-level changes and allows providers to deliver enhanced security capabilities to end customers.
- **Low cost per core for providers and end customers:**
  The high core density and power efficiency of EPYC allows providers to build and deliver virtual compute-as-a-service offerings at lower price points per core than other similar offerings in the market. Organizations can leverage the same capabilities to build out dedicated cloud infrastructure.

For the first time in several years, the AMD EPYC is considered more than an alternative commercial processor for the datacenter server market and a first choice for fit-to-purpose infrastructure. It improves single-socket and dual-socket performance with reduced TCO. With features such as higher core counts, more PCIe lanes, support for larger amounts of RAM, and larger cache memory, the EPYC processor platform directly enables the capabilities of the key workload requirements illustrated in **Figure 12** with several advantages.

The 3rd Gen AMD EPYC Processor cements its position as a critical building block for an optimized, scalable, and secure IT infrastructure.

The following are noteworthy capabilities of the processor platform.

- Single I/O/memory die removes processor subsystem bottleneck and provides 2x performance over the second-generation EPYC family.
- First x86 platform to support faster PCIe Gen 4 interface. Improved I/O performance in virtualized environments, especially for high-bandwidth PCIe Gen 4 endpoint attach use cases.

- 19% IPC (instructions per cycle) enhancements with the Zen 3 architecture over the previous generation.
- L3 cache shared across entire CCD (more available cache/core and better intercore latency).
- Reduced memory latency helps improve memory performance. Reduction in latency for fabric clock coupled with DRAM MEMCLK at DDR3200 (versus fabric clock decoupled DDR3200 DRAM in 7002).
- Improved IOMMU performance for attached peripherals and virtualized environments (distributed workloads, networking, NVMe). Improvement to achieve better line rate with 200Gbits.
- Boost frequency enhancements and uniform speed across cores, and high-bandwidth and low-latency network connectivity.
- Separate I/O die from Zen chiplets allows flexible core configurations. (Chiplets are multiple smaller pieces of silicon that make up a processor versus the older process of carving new processors from silicon as single chips.)
- Dedicated encryption offload sub-processor provides full stack encryption without any overhead.
- Dedicated I/O plus memory die provides improvements for NUMA workloads while also reducing latency for non-NUMA workloads.

The hardware security features of AMD EPYC processors are also used in confidential computing environments to protect data in use in addition to data in flight and data at rest. AMD EPYC features that are part of its Infinity Guard portfolio—such as Secure Encrypted Virtualization (SEV), Secure Encrypted Virtualization-Encrypted State (SEV-ES), and Secure Memory Encryption (SME)—enable organizations to encrypt data in use, data in flight, and data at rest.

Dell Is a Trusted Vendor in Datacenter Infrastructure

Organizations must consider the role of IT vendors as partners when selecting a suitable datacenter computing infrastructure.

Customers of Dell place a level of trust in Dell's ability to build and deliver the platform, including:

- **Maintaining a secure supply chain:**
  This includes verifying the authenticity of components or parts, procuring components or parts from trusted suppliers, and physically securing the build environment, the system build process, and the process of shipping the system to the customer.

- **Building security into every design step:**
  This includes incorporating security features in the hardware to prevent malicious attacks as well as developing, testing, and verifying the integrity of the firmware prior to installing it.
- **Providing proactive updates and system recovery:**
  This includes proactively patching firmware bugs or vulnerabilities and ensuring that updates are delivered and applied in a timely fashion. This also includes providing a restore to a known good state when a system recovery is necessary.

Dell PowerEdge servers have featured robust security for several generations, including the innovation of using silicon-based data security. Dell 14G PowerEdge servers extended silicon-based security to authenticate BIOS and firmware with a cryptographic root of trust during the server boot process.

Dell's product team considered several key requirements during the design of the 14th and 15th generation of PowerEdge servers in response to security threats faced in modern IT environments:

- **Protect** the server during every aspect of the life cycle, including BIOS, firmware, data, and physical hardware.
- **Detect** malicious cyberattacks and unapproved changes; engage IT administrators proactively.
- **Recover** BIOS, firmware, and OS to a known good state; securely retire or repurpose servers.

Dell PowerEdge servers conform to key industry standards on cryptography and security and perform ongoing tracking and management of new vulnerabilities. Dell EMC has implemented the Security Development Lifecycle process with security as a key element in every aspect of development, procurement, manufacturing, shipping, and support, resulting in a cyber-resilient architecture.

Trusted and reputable vendors often take their time to design, build, and deliver a quality product. They prioritize getting the product right versus superficial accolades. It is worth investigating what such vendors have planned for release when investing in new infrastructure. It is crucial for organizations to partner with vendors like Dell that are investing in design and engineering capabilities to produce such a platform at scale.

Dell AMD Third-Generation EPYC-Based Servers

Dell Technologies rose to this challenge with its AMD Rome–based, second-generation servers.

Overall, customers are increasingly satisfied with Dell's ability to deliver:

- Extra and high core configuration for running multi/mixed workloads
- PCIe 4 lanes for I/O intensive workloads
- Reduction in licensing costs with single-socket configurations

Dell's family of server platforms has been designed from the ground up to take full advantage of capabilities offered by the third-generation AMD EPYC processors. Such server platforms offer investment protection as their capabilities are future-ready.

Configuration enabled by Dell's design and engineering include:

- Higher sustained performance per processor
- Configuration flexibility for mixed workloads using optimized risers
- Flexible use of PCIe lanes for network, storage, and peripheral access
- Socket direct networking for balanced network I/O
- Direct NVMe configuration eliminates switches or bridges
- Liquid cooling options via partners
- Reduced deployment time through automation capabilities and system management solutions
- Improvement in datacenter cooling and power utilization efficiency
- Easy BIOS tuning with workload-optimized server configuration profiles
- Verifiable hardware secure root of trust that is etched at the factory

Dell supports the Infinity Guard features of AMD EPYC processors, enabling organizations to deploy Dell servers in hosted or collated environments where physical security could be compromised.

Dell Project Apex: Flexible/as-a-Service Option for Enterprises

Dell Technologies Cloud builds a common interface, all-way connections between edge, public, and private cloud, and a single management console to maintain these environments. Dell offers Flex on Demand as a pay-per-use consumption model for technology solutions across its entire ISG portfolio. Dell servers based on AMD EPYC-based CPUs provide end users with hardware and software optimization alternatives that enhance compute and memory capabilities, increase scalability, improve reliability and performance, provide more support of organizational IT initiatives for organizations looking to be technical innovators, and, finally, reduce overall compute costs.

Benefits of Adopting Dell AMD Third-Generation EPYC-Based Servers

Prior to the release of the third-generation AMD EPYC processors, Dell Technologies commissioned IDC to study the business, service quality, and financial improvements that end users realized with Dell AMD EPYC-based servers (that used first- and second-generation AMD EPYC processors). For its qualitative analysis, IDC relied on empirical data obtained via 17 in-depth interviews split between AMD EPYC server adopters and evaluators that did not adopt.

Additional quantitative analysis was gathered via a web survey of 1,513 IT practitioners and decision makers in organizations that were current:

- Adopters of Dell AMD EPYC-based servers (703 respondents)
- Adopters of AMD EPYC-based servers from other vendors (403 respondents)
- Evaluators considering AMD EPYC-based servers (252 respondents)
- Non-adopters that considered but ultimately did not adopt AMD EPYC-based servers (155 respondents)

Additionally, IDC's observations, insight, and recommendations are based on over six decades of research and intelligence on IT infrastructure industry and markets. All monetary values are in USD.

Adoption Drivers

Dell AMD EPYC-based servers prove to be the fit-for-purpose infrastructure needed to achieve service quality, business, and financial benefits among Dell's customers. By prioritizing features that improve performance (3.4GHz precision boost utilizing TDP@225W), memory (3200 MT/s DDR4 memory), and security (secure encrypted virtualization and secure memory encryption), Dell affords its customers the technological innovation, technical support, and configuration flexibility needed to support data-intensive, business-critical workloads (illustrated in **Figures 3**, **4**, and **5**, next pages).

FIGURE 3

Preferred Features of Dell Servers with AMD EPYC CPUs

(% of respondents)

Q. Which of the following AMD features is your organization taking or planning to take advantage of or evaluated?

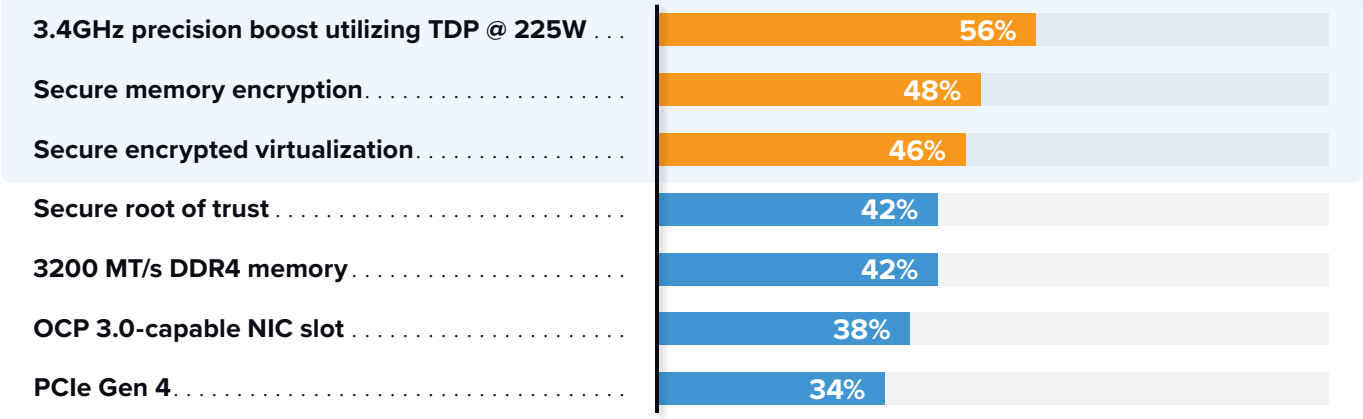

Note: Multiple dichotomous table—total will not sum to 100%
n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

FIGURE 4

Top Technical Criteria Considered when Selecting Dell Servers with AMD EPYC CPUs

(% of respondents)

Q. What are the top technical criteria for selecting Dell servers with AMD EPYC CPUs in your organization?

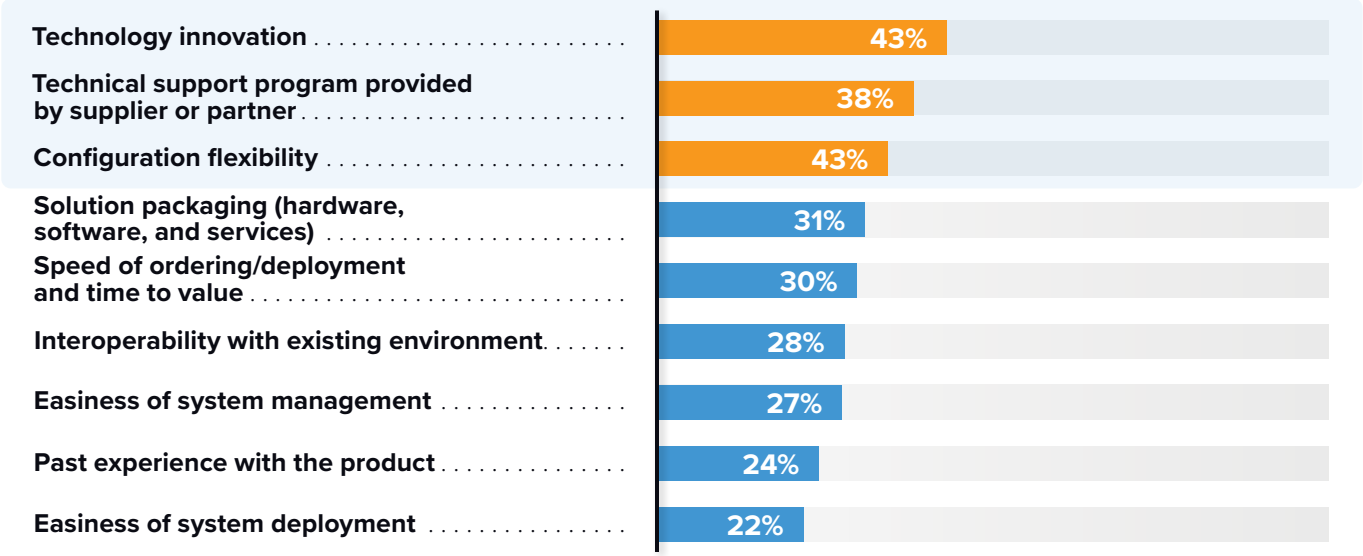

Note: Multiple dichotomous table—total will not sum to 100%
n = 533, Base = Path A (Have Dell with AMD) respondents indicated technical or business and technical reasons primarily led organization to buy Dell servers with AMD EPYC CPUs
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

FIGURE 5

Workloads Run on Dell Servers with AMD EPYC CPUs

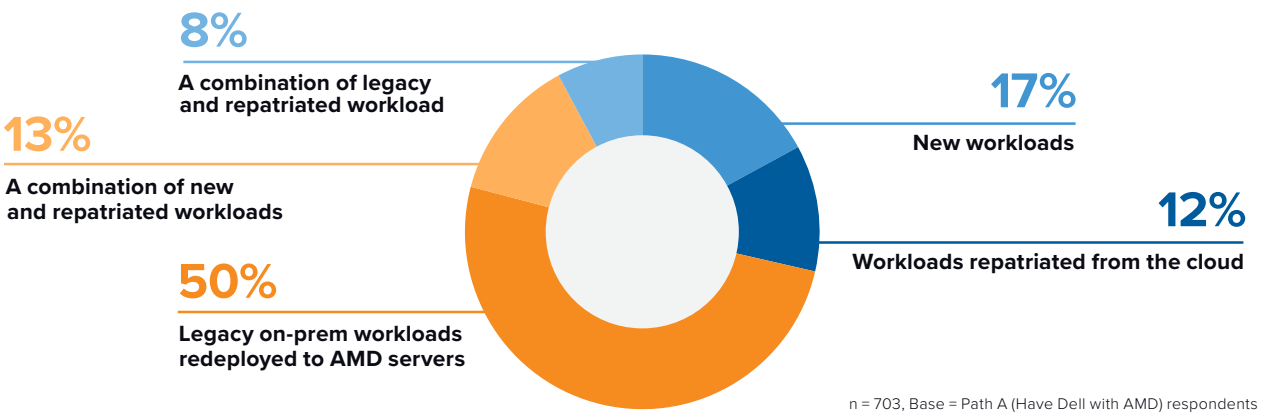

n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey*, October 2020

Service-Quality Improvements

Service-quality factors ultimately show up in how well the business can rely on its infrastructure. The better the service quality, the more scalable the business is in terms of meeting its objectives. Adopters of Dell AMD EPYC-based servers reported significant service-quality gains following the adoption of the Dell servers with second-generation AMD EPYC processors, specifically with regard to compute density, reliability, and performance (illustrated in **Figure 6,** next page).

- Computing density is a combination of the number of servers consolidated with each server cycle and the (increased) number of virtual machines per server. Most Dell customers (60%) reported improvement in virtual machine density.
- Improved application and system reliability is also a top indicator of how the business benefits from the adoption of Dell servers with second-generation AMD EPYC processors, with 58% of respondents reporting this improvement.

Platform selection

"The high-speed memory, again, some of our workloads are fairly memory-bound. So, it is important to be able to get lots of data through the memory. The high clock rate is important because not all portions of the code scale well to many cores. So, there are small bits of the code where there are only a few cores working and now we want to get through those as quickly as possible. It is good to be able to reduce the clock rate."

—Senior research associate, physics, at a large university research facility

- Faster application performance is a leading indicator of how the business benefits from the adoption of Dell servers with second-generation AMD EPYC processors. About half (50%) of Dell's customers cited this improvement.
- Fewer hours spent on tactical activities means there are more hours available for strategic priorities. Redeploying IT staff to focus on unplanned outages, troubleshooting or issue resolution, provisioning, and maintenance reduces the number of resources available for tasks that improve the business and create revenue. About a third of respondents cited decreases in unplanned downtime (35%) as well as time spent on troubleshooting and issue resolution (33%) and deployments and provisioning (30%). About a fifth of respondents experienced a reduction in time spent on management and maintenance (21%).

FIGURE 6

Service-Quality Improvements Being Realized by Dell AMD Rome-Based Adopters (% of respondents)

Q. What service-quality improvements has your organization experienced from the adoption of Dell servers with AMD EPYC CPUs?

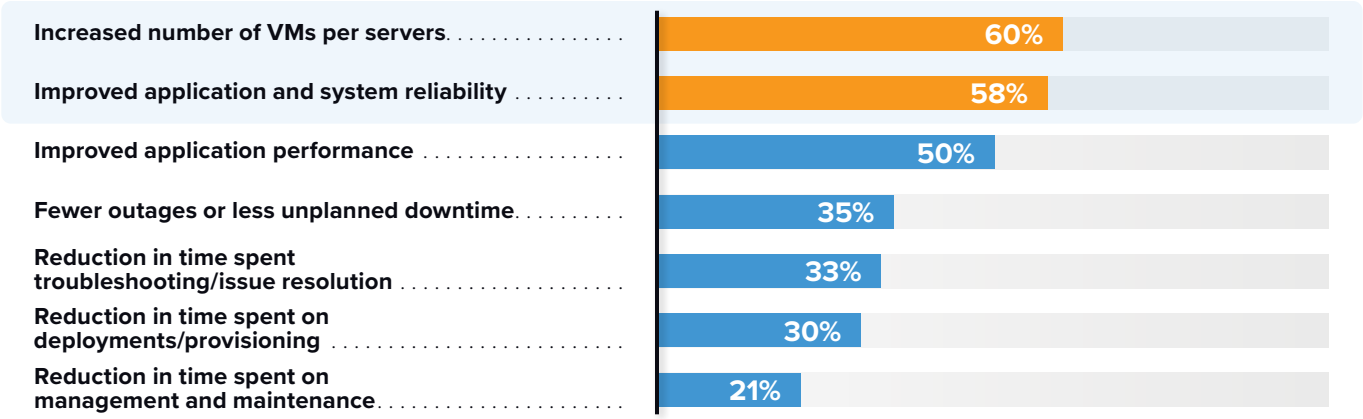

Note: Multiple dichotomous table—total will not sum to 100%
n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

Virtual machine density

"We are still very much VM-based. We are supposed to do the containerization, Docker, or Kubernetes, we are supposed to move to that stack, but we have not got time. We have to spend all the money to keep the lights on. We do not have much time to do the innovation. So, we are still stuck on the VM because VM is sort of a bit legacy at this point. We are supposed to move to Docker and Kubernetes, but we do not have the time and resources."

—VP of IT at a major global bank and financial institution

Business Benefits

Business benefits indicate the level of success an organization has obtained when fulfilling its internal and external objectives and how server infrastructure impacts these objectives. Adopters of Dell AMD EPYC-based servers reported significant business gains following the adoption of the Dell servers with second-generation AMD EPYC processors, specifically among revenue-centric and internally focused factors (illustrated **Figure 7**).

- Improved revenue-centric factors include direct impact to the business itself. Most Dell customers cited a positive impact on business agility (56%), scalability for certain workloads and applications (54%), and more effective application development (34%).
- Improved internally focused factors include how well the business is taking care of itself. Here, respondents reported impressive numbers, which include 53% of Dell's customers reporting a positive impact on support for organizational IT initiatives. Additionally, respondents cited improvements in business operations (20%) as well as their ability to take advantage of new business opportunities (34%).

FIGURE 7

Business Benefits Being Realized by Dell AMD Rome-Based Adopters

(% of respondents)

Q. What are the top business benefits that your organization is realizing from the adoption of Dell servers with AMD EPYC CPUs?

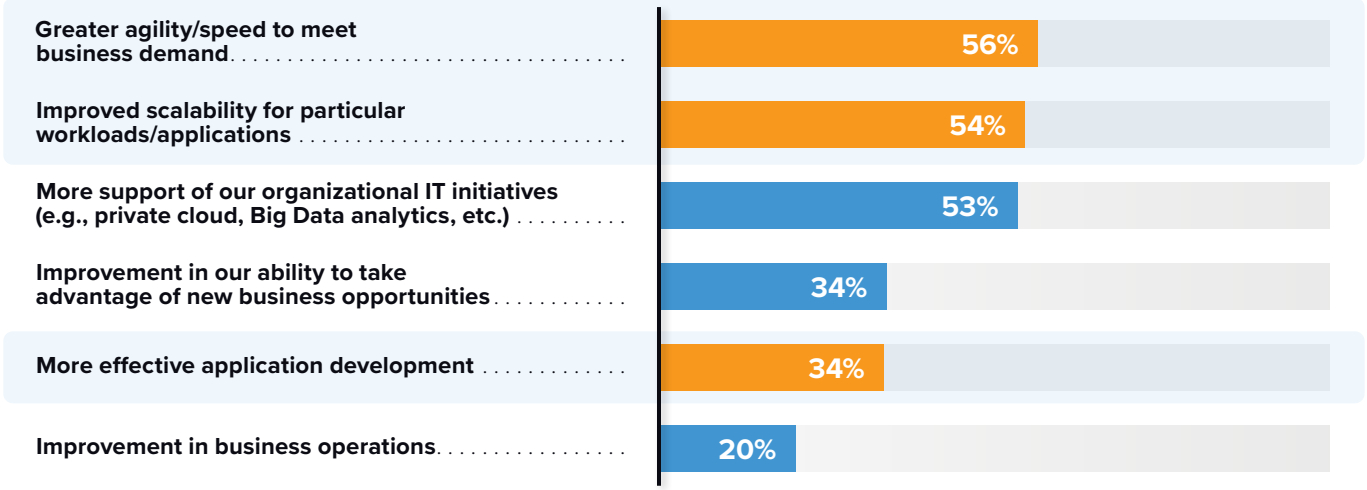

Note: Multiple dichotomous table—total will not sum to 100%
n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

Financial Improvements

Financial factors indicate the level of financial success an organization has obtained. Ideally, this entails spending less than the organization is making in revenue. Adopters of Dell AMD-based servers with AMD EPYC-based servers reported significant cost savings following the adoption of Dell servers with second-generation AMD EPYC processors, specifically with regard to IT staff burden and overall cost savings (illustrated in **Figure 8**).

- Improvements to staff burden enable the organization to better allocate IT staff resources to tasks that improve the business. Dell's customers reported reductions in IT staff burden (63%) and in the number of administrators per server (17%).
- Improvements to cost savings increase an organization's net growth. Spending more on IT reduces the amount of revenue realized. Most respondents reported improvements to cash flows (53%) and return on investment (47%). Overall cost of compute resources (39%) and dollars spent to operate each server (35%) also decreased among respondents.

FIGURE 8

Financial Improvements Being Realized by Dell AMD Rome-Based Adopters

(% of respondents)

Q. What financial improvements has your organization experienced from the adoption of Dell servers with AMD EPYC CPUs?

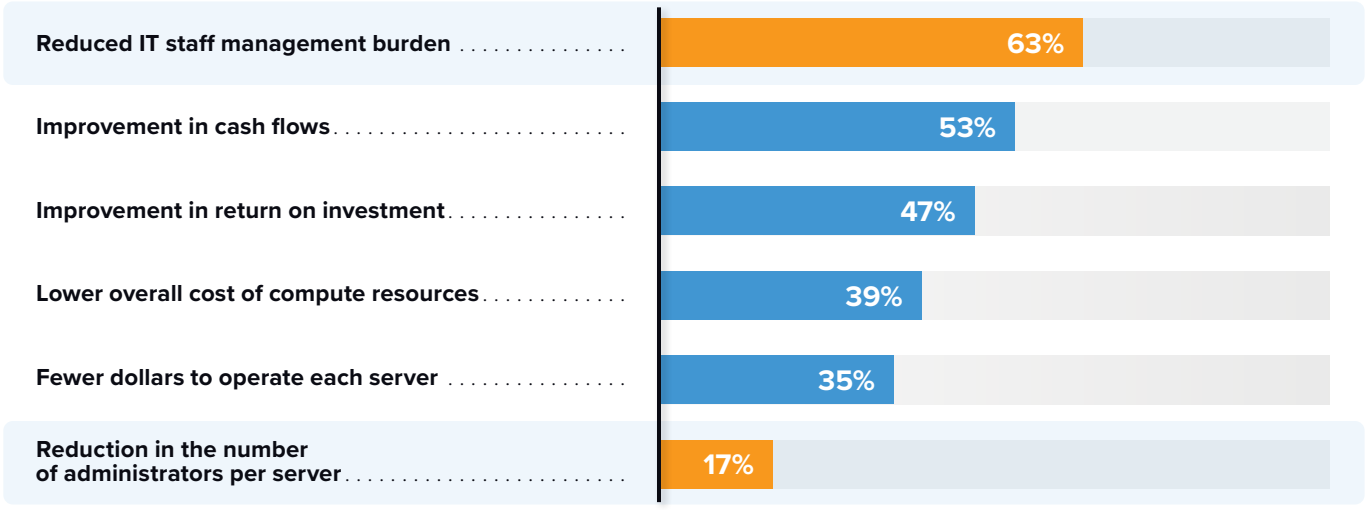

Note: Multiple dichotomous table—total will not sum to 100%
n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

Hardware/Software Optimization

All Dell customers reported improvements in service-quality, business, and financial factors following the adoption of second-generation AMD EPYC processors. However, customers that supplemented this technology with hardware (1U versus 2U, one socket versus two sockets, hardware accelerators) and software optimization (AMD-specific customizations on the software stack) options expect a greater magnitude of improvements. For example, the percentage of Dell's customers that recognized improvement in virtual machine density increased to 61% for those that adopted 2U and AMD-specific customizations on the software stack. The percentage of Dell's customers that experienced greater business agility increased from 56% to 58% with the adoption of AMD-specific customizations on the software stack, and to 59% with the adoption of hardware accelerators. Among financial factors, the percentage of Dell's customers citing reductions in IT staff management burden increased by 2% when using 2U form factors or AMD-specific customizations on the software stack. These findings may lead organizations to assume that adopting more hardware and software optimizations will result in increased benefits. This is not necessarily the case.

The largest percentage increase of Dell's customers citing service, business, or financial improvements occurred among those reporting reductions in time spent on management and maintenance when using hardware accelerators (+24%) and 1U form factors (+9%). In other instances, the adoption of these optimization technologies resulted in negative or small outcomes. For example, the percentage of Dell's customers reporting a reduction in the number of virtual machines when using 1U form factors decreased by 5%. The percentage of Dell's customers reporting the same improvements saw no change when using the hardware accelerators.

The overarching theme of the results supports the notion that when used as a fit-for-purpose infrastructure, Dell AMD EPYC-based servers, in combination with hardware and software optimizations, produce service-quality, business, and financial benefits among Dell's customers.

Staff reallocation

"Our developers do not bother to do any code optimization anymore because they want to delegate the dirty work to the server. That is the only way to improve performance, partly because any time you change your business logic, we must do very tedious regression testing and have someone do UAT testing, then have someone sign off on the change. The only way to improve application performance is from the hardware. That is why we constantly are looking for significant hardware improvements, in this case more cores."

—VP of IT at a major global bank and financial institution

FIGURE 9

Directional Change in the Percentage of Dell's Customers Experiencing Service Quality Improvements After Adopting AMD EPYC-Based Servers and Hardware/Software Optimizers

Q. What service-quality improvements has your organization experienced from the adoption of Dell servers with AMD EPYC CPUs?

Service-Quality Factor	All Dell Adopters (n = 703)	Financial Factor				Hardware Accelerators (n = 547)	AMD-Specific Customizations on the Software Stack (n = 519)
		1 Socket (n = 245)	2 Socket (n = 368)	1U (n = 245)	2U (n = 368)		
Increased number of VMs per servers	60%	–	–	~	–	+	+
Improved application and system reliability	58%	–	+	+	–	+	~
Improved application performance	50%	+	+	+	+	+	+
Fewer outages or less unplanned downtime	35%	+	–	~	–	+	+
Reduction in time spent troubleshooting/issue resolution	33%	~	+	–	~	+	~
Reduction in time spent on deployments/provisioning	30%	–	~	+	+	–	–
Reduction in time spent on management and maintenance	21%	–	+	+	+	–	–

Note: Multiple dichotomous table—total will not sum to 100%
n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

FIGURE 10

Directional Change in the Percentage of Dell's Customers Experiencing Business Improvements After Adopting AMD EPYC-Based Servers and Hardware/Software Optimizers

Q. What are the top business benefits that your organization is realizing from the adoption of Dell servers with AMD EPYC CPUs?

Business Factor	All Dell Adopters (n = 703)	Financial Factor				Hardware Accelerators (n = 547)	AMD-Specific Customizations on the Software Stack (n = 519)
		1 Socket (n = 245)	2 Socket (n = 368)	1U (n = 245)	2U (n = 368)		
Greater agility/speed to meet business demand	56%	+	–	+	–	+	+
Improved scalability for particular workloads/applications	54%	–	+	–	+	~	~
More support of our organizational IT initiatives	53%	–	+	+	+	+	+
Improvement in our ability to take advantage of new business opportunities	34%	–	+	–	+	~	–
More effective application development	34%	+	~	+	–	–	–
Improvement in business operations	20%	+	–	+	–	–	+

Note: Multiple dichotomous table—total will not sum to 100%
n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

FIGURE 11

Directional Change in the Percentage of Dell's Customers Experiencing Financial Improvements After Adopting AMD EPYC-Based Servers and Hardware/Software Optimizers

Q. What financial improvements has your organization experienced from the adoption of Dell servers with AMD EPYC CPUs?

Financial Factor	All Dell Adopters (n = 703)	Financial Factor				Hardware Accelerators (n = 547)	AMD-Specific Customizations on the Software Stack (n = 519)
		1 Socket (n = 245)	2 Socket (n = 368)	1U (n = 245)	2U (n = 368)		
Reduced IT staff management burden	63%	+	–	+	–	+	+
Improvement in cash flows	53%	+	–	~	–	+	+
Improvement in return on investment	47%	+	–	+	–	+	~
Lower overall cost of compute resources	39%	–	+	–	+	–	~
Fewer dollars to operate each server	35%	+	+	+	+	+	~
Reduction in the number of administrators per server	17%	~	~	–	+	–	~

Note: Multiple dichotomous table—total will not sum to 100%
n = 703, Base = Path A (Have Dell with AMD) respondents
Source: IDC *Dell AMD Buyer Behavior Survey,* October 2020

Essential Guidance for IT Buyers

Infrastructure is going to play an increasingly important role in business outcomes as firms undergo digital transformation. As the stewards of digital currency (i.e., data), IT organizations have a fiduciary responsibility to secure the firm's IT infrastructure. Partnering with an IT infrastructure vendor to procure a fit-for-purpose infrastructure, using products and solutions that can deliver performance and security, is of paramount importance. The vendor must back up its offerings with a services organization which can in turn assist in in protecting, detecting, and recovering from security threats.

Taking a Systematic Approach for Infrastructure

IDC believes that IT organizations benefit from developing and implementing a fit-for-purpose infrastructure strategy. Organizations often take the systematic approach outlined below (and illustrated in **Figure 12**) when selecting infrastructure computing platforms (servers).

FIGURE 12

A Systematic Approach for Infrastructure Selection

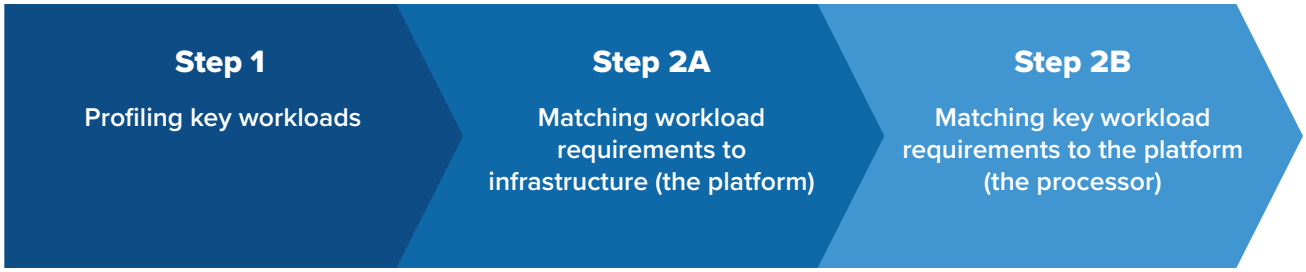

Source: IDC, 2021

Step 1: Profiling key workloads

First, IT organizations profile each workload type for its compute, storage, and networking (connectivity) requirements (illustrated in **Figure 13**, next page). For example:

- Performance-intensive computing workloads require varying degrees of compute.
  - Data analytics workloads call for matching compute.
  - Modeling and simulation workloads need bandwidth and memory.
  - Artificial intelligence and machine learning workloads rely on accelerated compute to parallelize the execution of complex compute-intensive algorithms.
- Desktop virtualization workloads demand reliable performance and bandwidth to match dynamic user demand.
- In-memory database workloads that require bandwidth and throughput for high volume transactions.

FIGURE 13

Compute and Storage Workload Profiles

Data Analytics and AI/ML/DL	Software-Defined Storage	Massively Parallel Computing	Virtualization and Virtual Desktop Infrastructure (VDI)	Network Virtualization
High parallelism for complex analysis	Direct high-performance and capacity-optimized storage support	Highly parallelized processor optimized for accelerators and real-time data streams	Single-socket performance that reduces TCO without compromising availability	Accelerated provisioning/agility
Massive I/O bandwidth for faster data loads	Low latency and high I/O parallelism for data access and persistence	Massive I/O for cluster connectivity	Higher core count to enable dense user base	High-bandwidth network connectivity
High integrity for floating-point capacity	More memory for larger cache	Memory capacity for larger data sets	Larger memory capacity for VM density	High integer and floating-point capacity
Security for business-critical data	Memory encryption for data security	Massive I/O bandwidth	Cryptographic isolation between hypervisor and VM	Security for business-critical data
Hadoop/Spark; TensorFlow	Hyperconverged and non-hyperconverged file, block and object SDS	High-performance applications	VMware Horizon; Citrix Xen Desk/Xen App; Autodesk	OpenStack; NSX; AHCI
Compute optimized (with support for GPUs/accelerators)	Storage and I/O optimized	Compute optimized	Compute and I/O optimized	Compute, I/O, and connectivity optimized

Source: IDC, 2021

Step 2A: Matching workload requirements to infrastructure (the platform)

The second step is to evaluate the platform and match workload requirements to the capability of the platform itself. Answers to questions asked by IT decision makers provide the weighting necessary to make a proper workload-centric selection.

- **Performance:**
  How fast can my platform run? How can I accelerate the performance? What are some of the intrinsic scaling capabilities of the platform?
- **Compute scaling:**
  How well can my server systems scale for compute? What kind of compute density (i.e., cores per unit height) do they offer? Support for peripherals is often an important criterion as well (e.g., accelerator and network interface support, the types of accelerators supported, the number of peripherals supported).

- **Storage scaling:**
  What kind of storage density is available in the platform in terms of terabytes per node? What kind support does it offer for concurrent data-intensive, low-latency applications?
- **Acquisition costs:**
  What kinds of choices are available on a cost-per-node basis? Can I procure these platforms on a buy, lease, or flexible-consumption (pay-as-you-go) basis?

A server platform provides a solid foundation on which firms launch key digital initiatives. The security of the server platform is crucial to maintaining the integrity of the data that lives on it. Therefore, as demonstrated in a recent IDC study, server security is a crucial starting point when making any hardware purchasing decisions. Study results indicate that 41% of respondents ranked security of hardware (including BIOS, BMC, and other firmware) in their top selection criteria when procuring server infrastructure, surpassing total cost of operations (34%) (illustrated in **Figure 14**). It must be core to the design of the platform, not "bolted on" as an afterthought. Specifically, one fifth of survey respondents considered only server infrastructure with system lockdown and secure customizable alerts for health, configuration changes, and updates to security hardware and firmware (illustrated in **Figure 15**, next page).

FIGURE 14

Selection Criteria During the Purchase of Server Infrastructure

(% of respondents)

Q. When it comes to evaluating the purchase of server infrastructure, what are the three most important selection criteria for your organization?

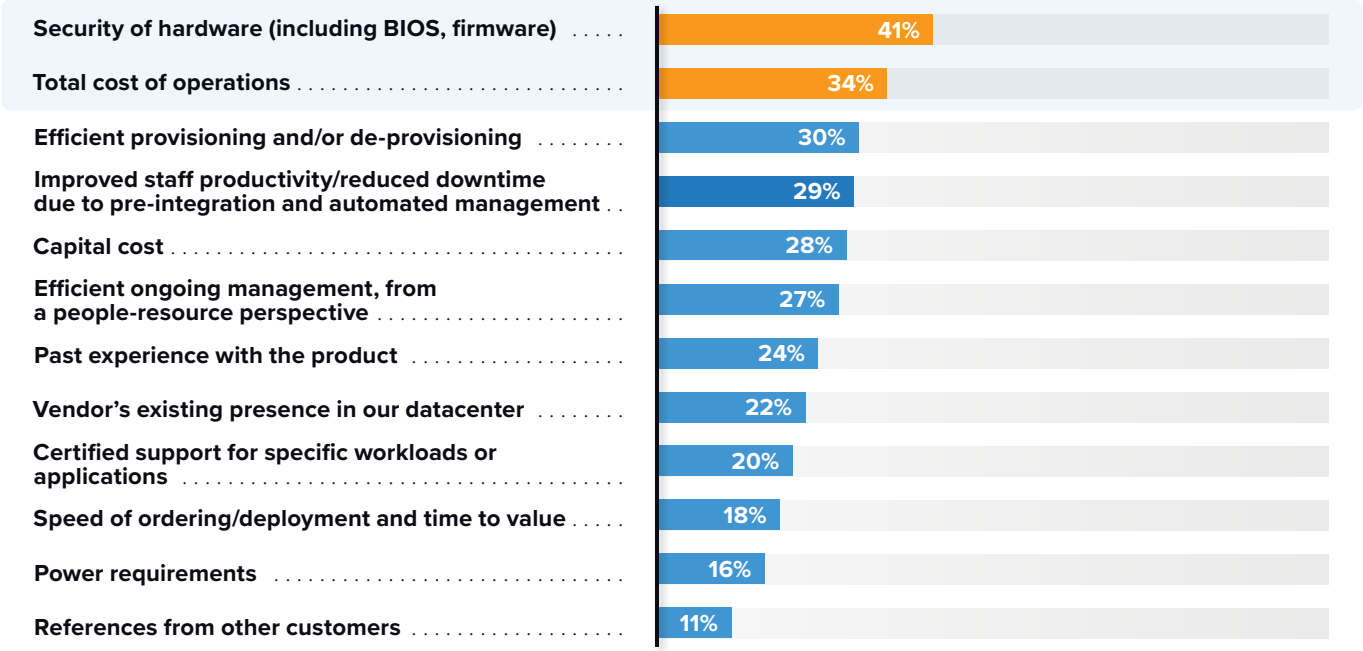

Note: Multiple dichotomous table—total will not sum to 100%
n = 301, Base = All respondents
Source: IDC *Value of Secure Server Infrastructure—Web Survey 2018*, February 2018

FIGURE 15

Hardware/Firmware Security Features, Ranked

(% of respondents)

Q. Which hardware/firmware security features would you consider mandatory for you to select a vendor's server hardware?

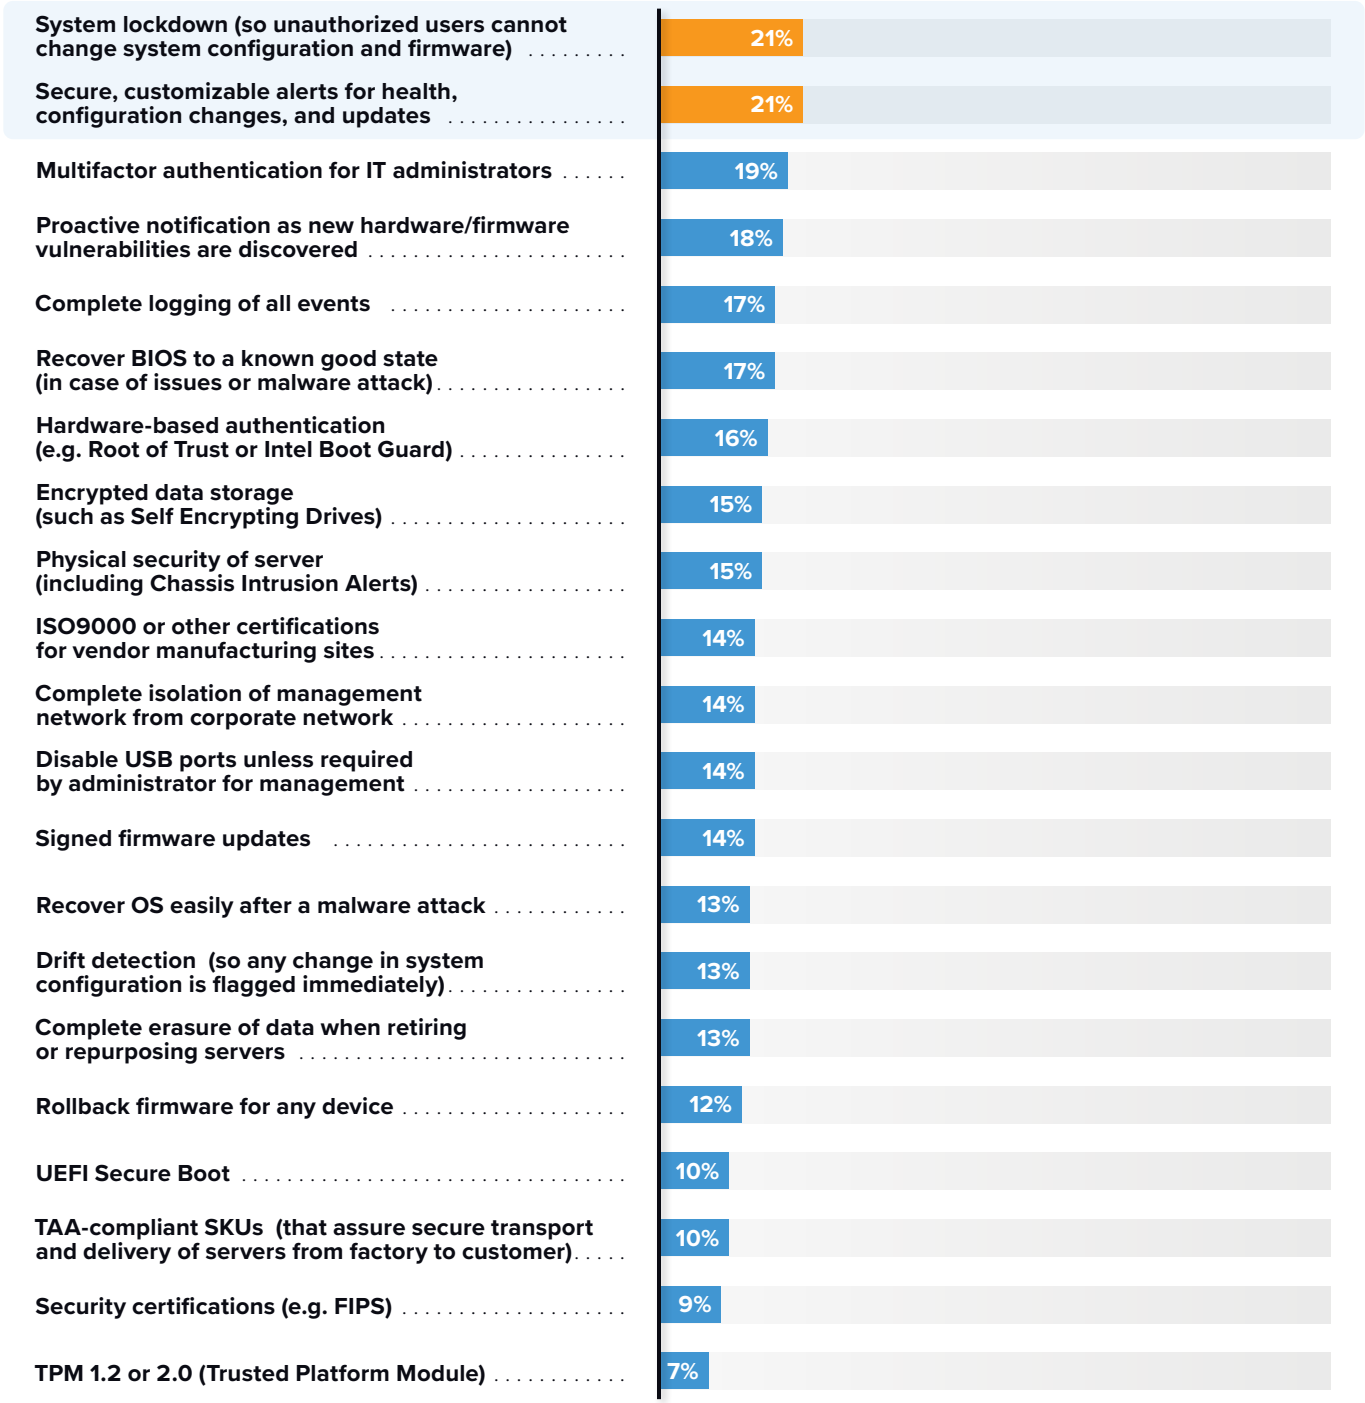

Note: Multiple dichotomous table—total will not sum to 100%
n = 301, Base = All respondents
Source: IDC *Value of Secure Server Infrastructure—Web Survey 2018*, February 2018

Step 2B: Matching key workload requirements to the platform (the processor)

Using the infrastructure selection process, organizations perform a critical examination of the processor subsystem. During this examination, they examine processor-centric optimizations (e.g., memory, I/O, connectivity, and storage) and methods for enhancing workload performance. For example, some end users indicate a preference for enhanced memory capability for targeted workloads. For others, the choice of a suitable processor in fact enhances existing workload performance (illustrated in **Figure 16**).

FIGURE 16

Processor Attributes

Attribute Category	What to Look For
Processor subsystem performance	▸ High core and thread count per socket ▸ High L2 and L3 cache per core ▸ Higher TDP wattage ▸ Memory performance and capacity
I/O bandwidth and latency	▸ PCIe capabilities for high-bandwidth/low-latency access to accelerators, connectivity interfaces, and storage media ▸ Number of PCIe lanes—flexibility in configuration for optimized performance
Storage performance and latency	▸ NVMe direct performance ▸ PCIe lanes for nondirect I/O
Processor subsystem security	▸ Processor security to prevent hostile code execution attacks ▸ Memory encryption ▸ Workload and VM-aware isolation ▸ Key management ▸ Operating system and hypervisor support
Operating and ownership costs	▸ x86-based ISA eliminates the need to rewrite applications ▸ Single-socket performance for reduced TCO in socket-based licensing ▸ Dual-socket performance for improved ROI in mixed or specialized workload configuration ▸ Based on industry-standard x86 ISA requiring no migration costs ▸ Integration with existing management solution

Source: IDC, 2021

A Shift to Operational Expenses: Flexible Consumption Options

infrastructure and financial services portfolio, specifically supporting usage-based financial payments. These offerings ensure that an organization's technology and financial needs are being met consistently and in a timely manner. With Dell Technologies Cloud as an option for hybrid environments, Dell Technologies can offer a customized infrastructure-as-a-service option to clients seeking capex-friendly solutions.

Infrastructure updates on a flexible/pay-as-you-go procurement and consumption model are becoming more difficult.

Organizations are driven by many reasons to shift to operational expenses — for example, the ability to:

- Try out new technologies in the public cloud without making upfront capital-intensive investments
- Exit facilities management and instead go all in on a hosted private cloud, either self-managed or as a service
- Place newer workloads on best-of-breed server solutions without creating an imbalance in existing infrastructure
- Preserve cash during stressful times and/or during times when the business is seeking to redefine itself and generate new routes to market

Public cloud infrastructure services are not the only avenue to achieving these objectives. Organizations must examine options provided by OEM vendors like Dell.

Such options are quite compelling, as they:

- Are designed for organizations that want to consume infrastructure either traditionally or as a hybrid cloud anchored on premises and want to be in control of their own security apparatus
- Meet the needs of data-intensive, business-critical, high-performing workloads and offer new technologies in support of such workloads (for example, GPUs)
- Offer customization options for clients seeking opex-friendly solutions that are specific to their tax and accounting practices
- Offer a cloud-based control and management plane enabling a true multicloud experience

Companies are no doubt changing their mindset and strategy for their cloud operations.The cloud operating model has become popular with its ability to:

- Scale IT infrastructure without the need for more human capital or fixed assets — i.e., servers and datacenter space

- Help organizations pay for technology only when it is needed, creating a more predictable and consistent budget plan for IT spending
- Enable the business to take advantage of elastic capacity—critical in environments where growth is unpredictable and capacity demands are in flux, eliminating the need to purchase excess capacity upfront
- Ensure that IT budgeats are more predictable and consistent using vendor-provided tools (like Dell CloudIQ)

As this consumption business model evolves, IDC expects vendors to develop newer offerings to push the market toward flexible consumption as a de facto procurement model. Organizations must be prepared for it.

FIGURE 17

Differences Between Capex and Opex Options

Capital Expenses	Operational Expenses
Enables customers to procure and place assets on their books.	Gives customers the option to either procure assets and place them on their books or have the vendor carry them on their books.
Customers can choose to redeploy assets once they are no longer on the maintenance cycle.	Customers can move assets out of their datacenter once the commitment period is over.
Customers get better TCO and ROI.	Customers are able to preserve capital.

Source: IDC, 2021

All Servers and Server Vendors Are Not the Same

Firms must be wary when purchasing equipment from white box server manufacturers. Most of these vendors have limited resources and do not invest heavily in developing and managing firmware-level security. They instead rely on their partners or ODM clients to provide such functionality on their own.

Many such vendors also lack the setup to meet the requirements of enterprise IT, such as the ability to meet stringent service-level agreements for feature updates and bug fixes. They lack the systems or resources to prevent or respond to threats on their own. This means their ability to aid with post-incident recovery is limited at best.

Buying servers from such vendors can expose firms and make their investments in servers more expensive in the long run.

The Importance of Server Security

Implementing and managing server security is a complex task and requires the IT organization to hold itself accountable for ensuring that the firm's key data assets are always secure.

Using a simple Protect, Detect, Recover framework (based on NIST recommendations), IT organizations can build a secure server infrastructure and maintain its integrity by incorporating four key areas of server security into the "people, process, and technology" fabric of the IT organization:

- **Physical server security:**
  This involves protecting the physical servers' enclosure, power supply, components, and fabric interconnection from being intentionally damaged or altered.
- **Firmware and software security:**
  This includes the ability to efficiently carry out complex encryption and cryptographic algorithms to protect both server data and application data in storage devices attached to the server, and the ability to monitor changes made to firmware and software.
- **Attestation trust features:**
  These capabilities are built into the server using special-purpose processor features and cryptographic on-board elements. These provide a root of trust (preferably based on an immutable authentication element) that enables checking at start-up and other points in the server life cycle whether the server has been altered from its expected configuration. Administrators can set policies to take the server offline and send a notification of the error to IT staff. This ensures that components of the system software stack (hypervisor, OS, applications) are aware the underlying server can be trusted when the server is operational. This layer establishes the foundation for a chain of trust between servers and creates a trusted and secure distributed server platform.
- **Secure server management:**
  This is carried out using an embedded baseboard management controller (such as the iDRAC, a proprietary remote server management technology from Dell) to handle server firmware and software updates and other operations for server security across the datacenter.

An effective way to deal with security-related organizational and infrastructure challenges is to move the organization toward IT automation. This makes large-scale asset management and auditing easier; automates routine activities, reducing the risk of human errors; and tags and identifies potential vulnerabilities, like out-of-date patches, before they become issues. Any automation solution should include comprehensive APIs and the ability to simplify complex tasks through software functionality.

Challenges and Opportunities for Dell

As Dell knows very well, organizations consider their IT vendors as partners in their digital transformation journey.

The level of trust that firms place in an infrastructure vendor is directly related to the vendor's ability to build and deliver a performant and secure server platform.

For a vendor, delivering a quality product (see **Figure 18**, next page) means having the ability to:

- **Maintain a secure supply chain:**
  Verifying the authenticity of components or parts, procuring them from trusted suppliers, and physically securing the build environment, the system build process, and the process of shipping the system to the customer. Dell's cyber-resilient design uses a cryptographically trusted booting cycle and immutable silicon root of trust, starting at Dell's factory.
- **Build security into every design step:**
  Incorporating security features in the hardware to prevent malicious attacks as well as developing, testing, and verifying the integrity of the firmware prior to installing it. Dell's servers enable IT organizations to wipe all data securely and quickly from storage media including hard drives, SSDs, and system memory with System Erase.
- **Provide proactive updates and system recovery:**
  Proactively patching firmware bugs or vulnerabilities and ensuring that updates are delivered and applied in a timely fashion; when a system recovery is necessary, providing a restore to a known good state. Dell Repository Manager and Update Manager plug-ins automate driver, BIOS, firmware, and software updates.

IT organizations are generally open to working with a vendor that expends the extra effort to make key influencers and buyers aware of the build integrity, and strongly affirm the need for individuals and teams involved in making or influencing purchasing decisions to be fully aware of the vendor's build process. Furthermore, firms hold vendors to high standards when it comes to the vendor's ability to protect (prevent), detect, and recover from security breaches. This is where the vendor's ability to stand behind its product with an equally competent services organization comes into play.

Maintaining a trusted roster of infrastructure vendors is a key practice in most modern IT organizations. Firms can add a vendor to this trusted list in return for the vendor's demonstrating consistent product and service quality. Firms can disqualify or change server vendors should the vendor fail to deliver on its commitments (see **Figure 19**, next page).

Examples that can lead firms to switch vendors include:

- Dissatisfaction with the quality of services and support, as well as inadequate or lacking security capabilities of the server platform
- Discovery of bugs in the security of the platform, a lack of or poor security features in the product, a lack of guarantees from the vendor on the integrity of its end-to-end build environment, and its inability to support security-related situations
- Security certification: In some industries, it may be a critical factor that could be the sole cause of vendor disqualification.

FIGURE 18

Key Capabilities of Trusted Vendors

(% of respondents)

Q. You indicated that your organization maintains a roster of trusted vendors for providing your infrastructure. What are the key capabilities that a vendor must meet to make it to this roster?

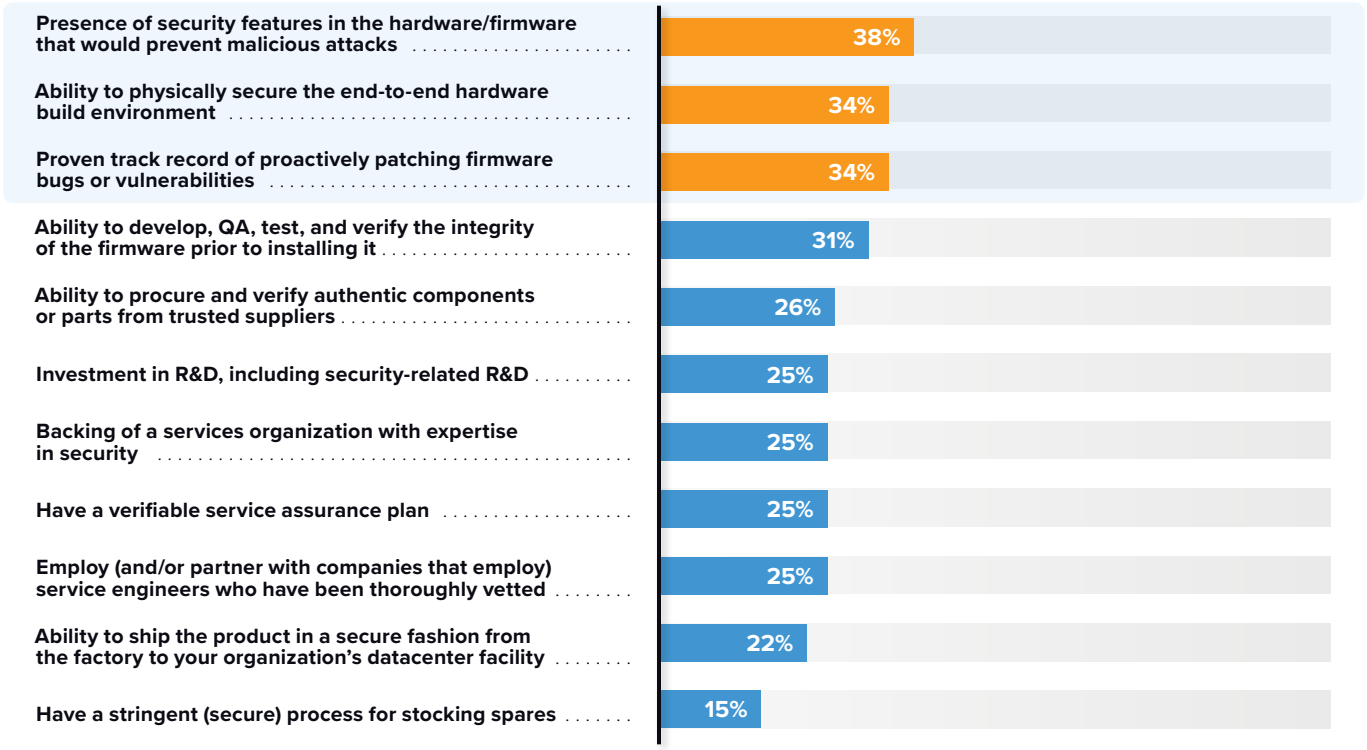

Note: Multiple dichotomous table—total will not sum to 100%
n = 273, Base = Respondents indicated their organization maintains a roster of trusted vendors for providing infrastructure
Source: IDC *Value of Secure Server Infrastructure—Web Survey 2018*, February 2018

FIGURE 19

Reasons for Switching Infrastructure Suppliers

(% of respondents)

Q. You indicated that you switched infrastructure suppliers in the past 12–18 months. What was the main reason for this switch?

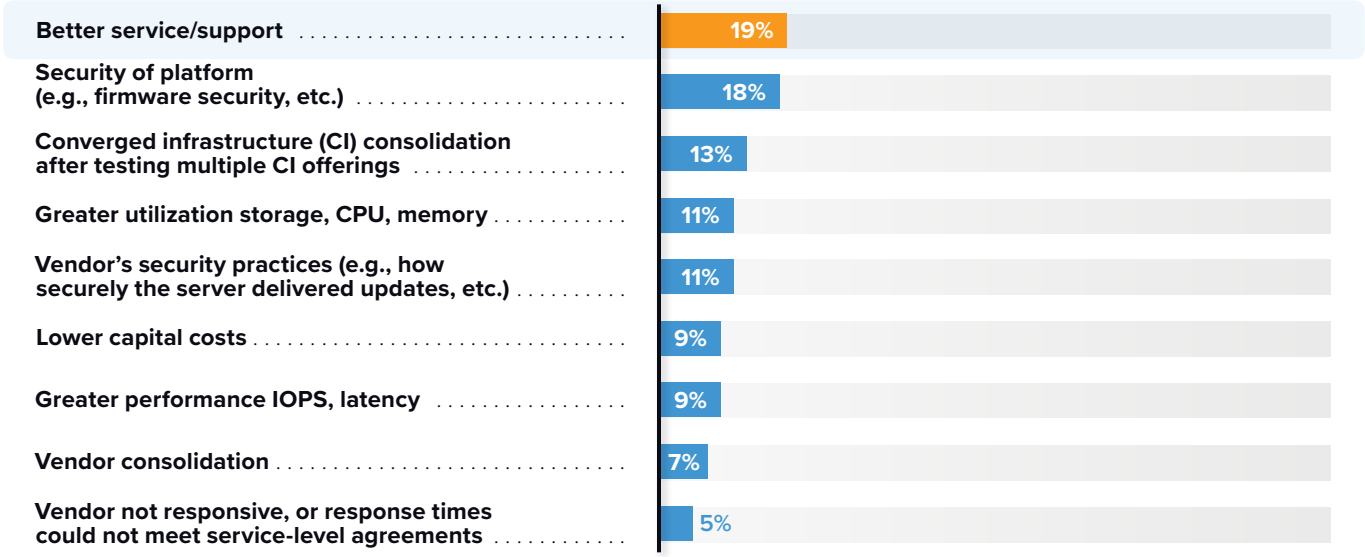

Note: Multiple dichotomous table—total will not sum to 100%
n = 150, Base = Respondents indicated they have switched infrastructure suppliers in the past 12-18 months
Source: IDC *Value of Secure Server Infrastructure—Web Survey 2018,* February 2018

Conclusion

Maintaining a modern server environment is becoming a critical success factor for enterprises in most industries. End-user expectations have risen to a point where IT is expected to be ubiquitous in the same way electricity is. Unexpected outages and downtime can have a direct impact on revenue and customer satisfaction; hence the need for a modern server environment. Until recently, IT decision makers did not have much of a choice when it came to the selection of fit-for-purpose infrastructure that would scale for performance-intensive and virtualized workloads while also lowering the total cost of environments. With AMD EPYC servers, and specifically AMD EPYC-based servers from Dell Technologies, they have that option. IDC encourages organizations to take a second look at this option: It enables IT staff to optimize the performance and security of their infrastructure while increasing their return on investment.

Appendix

Glossary

- **AMD Secure Encrypted Virtualization (SEV)**
  Uses one key per virtual machine to isolate guests and the hypervisor from one another. The keys are managed by the AMD Secure Processor. SEV requires enablement in the guest operating system and hypervisor. The guest changes allow the VM to indicate which pages in memory should be encrypted. The hypervisor changes use hardware virtualization instructions and communication with the AMD Secure Processor to manage the appropriate keys in the memory controller.
- **AMD Secure Encrypted Virtualization–Encrypted State (SEV-ES)**
  Encrypts all CPU register contents when a VM stops running. This prevents the leakage of information in CPU registers to components like the hypervisor and can even detect malicious modifications to a CPU register state.
- **AMD Secure Memory Encryption (SME)**
  Uses a single key to encrypt system memory. The key is generated by the AMD Secure Processor at boot. SME requires enablement in the system BIOS or operating system. When enabled in the BIOS, memory encryption is transparent and can be run with any operating system.
- **Non-uniform memory access (NUMA)**
  A computer memory design used in multiprocessing, where the memory access time depends on the memory location relative to the processor. Under NUMA, a processor can access its own local memory faster than nonlocal memory (memory local to another processor or memory shared between processors). The benefits of NUMA are limited to workloads where the data is often associated strongly with certain tasks or users.
- **Thermal design power (TDP)**
  Also called thermal design point, TDP is the maximum amount of heat generated by a computer processor or component that the cooling system in a computing platform (like a server) is designed to dissipate under any workload.

About the Analysts

Ashish Nadkarni
Group Vice President, Infrastructure Systems, Platforms and Technologies Group, IDC

Ashish Nadkarni is Group Vice President within IDC's Worldwide Infrastructure Practice. He leads a team of analysts who engage in delivering qualitative and quantitative research on computing, storage, and data management infrastructure platforms and technologies, via syndicated research programs (subscription services), data products (IDC Trackers), and custom engagements. Ashish's vision for his team is to take a holistic, forwarding-looking, and long-term view on emerging as well as established infrastructure-related areas in the datacenter, in the cloud, and at the edge. His core research starts with an objective assessment of heterogeneous, accelerated, fog, edge, and quantum computing architectures, silicon, memory, and data persistence technologies, composable and disaggregated systems, rackscale design, software-defined infrastructure, modern operating system environments, and physical, virtual, and cloud computing software. It is complemented by research on current and next-gen applications and workloads, vertical and industry-specific use cases, emerging storage and server form factors and deployment models, and upcoming IT vendors. Ashish also takes a keen interest in tracking the ongoing influence of open and open- source communities like OpenStack and Open Compute Project on infrastructure.

More about Ashish Nadkarni

Heather West, Ph.D.
Senior Research Analyst, Infrastructure Systems, Platforms and Technologies Group, IDC

Heather West is a Senior Research Analyst within IDC's Enterprise Infrastructure practice. In this role, Heather contributes to semi-annual Server and Storage Workloads Trackers, primary market research, and custom data modelling.

More about Heather West, Ph.D.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

IDC Custom Solutions

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

idc.com @idc