IDC

DELL AMD SERVERS:

Selecting a Fit-for-Purpose Server Platform for Datacenter Infrastructure

RESEARCH BY:

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

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

December 2021 | IDC Doc. #US48482721

Table of Contents

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

Datacenter Outcomes Are a Journey . . . 3

If the Path Is Different for Each Company, Then the Journey Is What Matters . . . 4

Let Fit-for-Purpose Design Guide You Through Technical, Business, and Financial Decisions for Your Datacenter . . . 5

Workload Drives the Selection of a Fit-for-Purpose Infrastructure Platform . . . 6

...And Then Shifts to Taking a Closer Look at the Profile of Key Business-Critical Workloads . . . 7

AMD EPYC Processors with Up to 64 Cores Provide a Critical Building Block for an Optimized, Scalable, and Secure IT Infrastructure . . . 8

What Businesses Look for in AMD EPYC Server Platforms . . . 9

A Trusted Vendor Delivers a Best-in-Class Fit-for-Purpose AMD EPYC Platform . . . 10

How to Benefit from Investing in Fit-for-Purpose Infrastructure . . . 11

Breaking the "Good Enough" Barrier: Investing in Infrastructure Is a Business Decision . . . 11

Considering Dell Technologies for Its AMD EPYC Server Portfolio . . . 12

IDC Methodology . . . 13

For Dell Partners . . . 14

About the Analysts . . . 15

Datacenter Outcomes Are a Journey

The destination is the same across industry and company size.

FIGURE 1

Business Benefits and Service Quality Improvements Being Realized by Dell Adopters

(% of benefits)

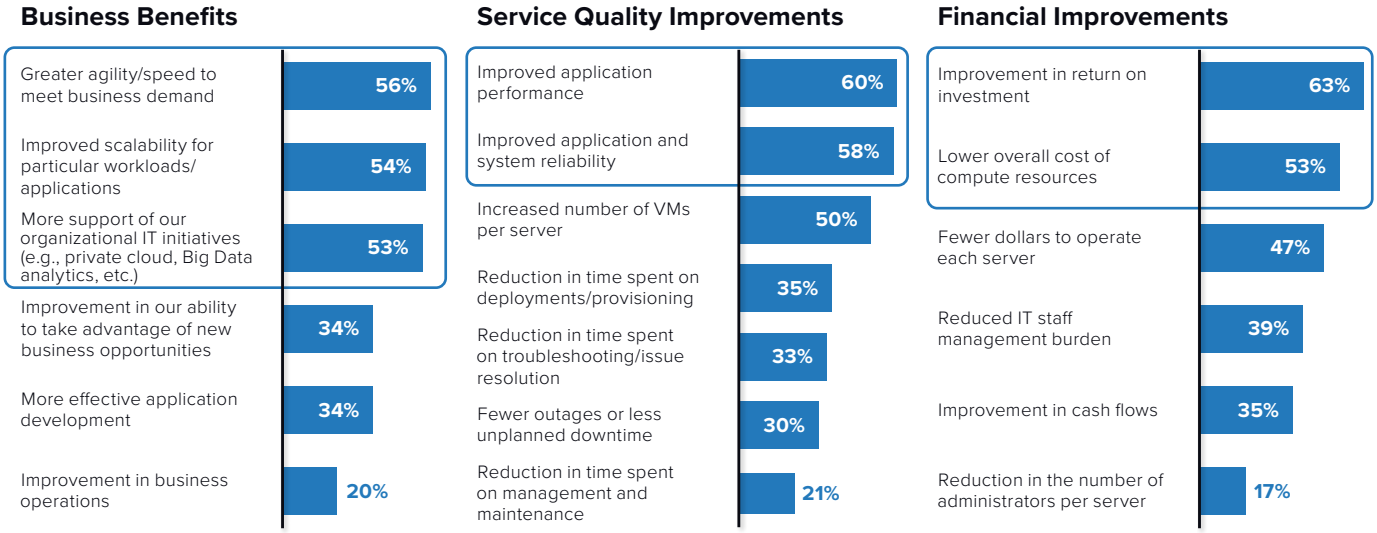

n = 703 (Dell adopters); Source: IDC *Dell AMD Buyer Behavior Survey*, October 2020; Note: Use caution when interpreting small sample sizes.

We are noticing a marked shift in the digital economy in terms of how businesses rely on infrastructure. Business priorities are driving workload transformation and placement, which in turn are driving infrastructure transformation. Investing in "fit-for-purpose" workload infrastructure is one of the crucial pillars of this transformational journey. AMD EPYC processors, with up to 64 cores, represent a unique opportunity for addressing the modern workload-optimized infrastructure world.

However, identifying the server platform that takes advantage of this opportunity is not always straightforward. Some server platforms may support the new processor, but don't enable the new processor features. Businesses must check the processor specs against the vendor's specs and ensure the server platform is fully optimized. This eBook will help guide you in the process and help you confirm long-term return on investment and reduced total cost of ownership (TCO).

To support superior business outcomes, CIOs must make strategic investments in IT infrastructure to support current and next-generation workloads. This infrastructure must:

- → Be highly available and scalable
- → Support end-to-end security
- → Support high-performance compute
- → Enable fast data access

If the Path Is Different for Each Company, Then the Journey Is What Matters

FIGURE 2

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

(% of respondents)

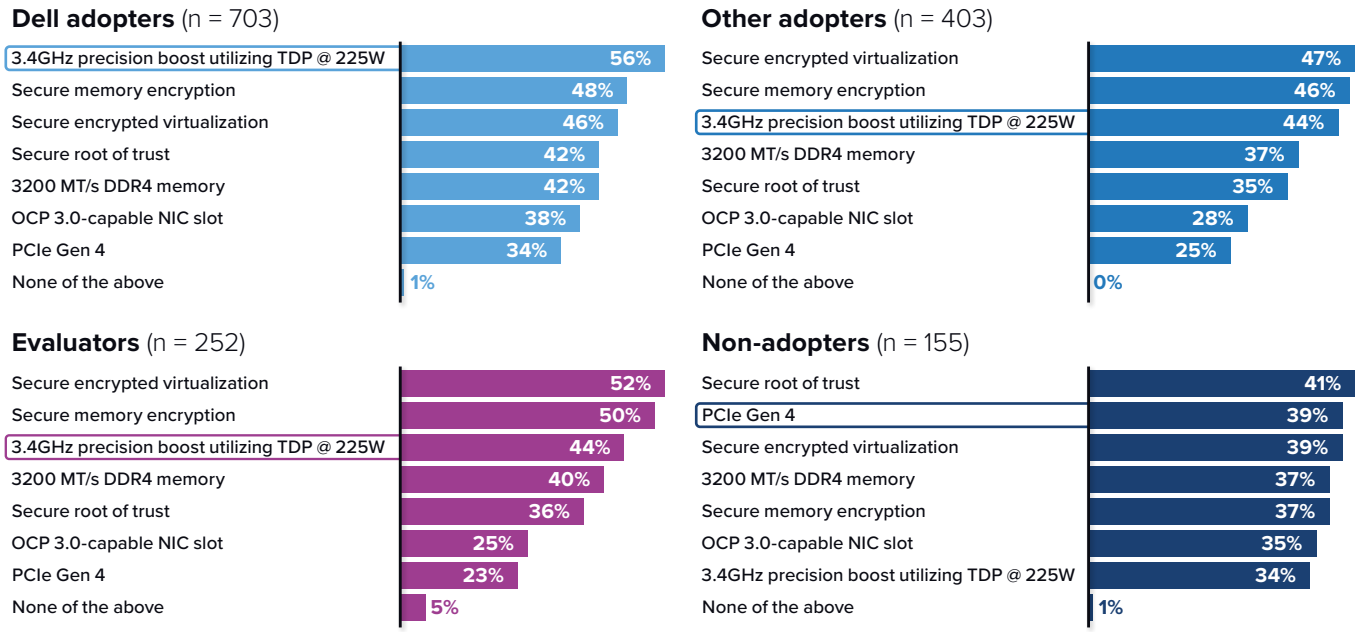

Source: IDC *Dell AMD Buyer Behavior Survey*, October 2020; Note: Use caution when interpreting small sample sizes.

TABLE 1

IT Organizations Focus on Different Technical Outcomes with Their Infrastructure Investments

AMD EPYC Platforms	Design Capabilities That Enable Optimization
Are built with a high-performance processor subsystem, reducing latency and delivering fast response times	• Support for full processor stack — up to 240W TDP • 20% faster memory speed at 3200 MT/s
Enable faster data access, faster transfer speeds for storage workloads, and faster connectivity for networking and storage workloads	• 25% more PCIe lanes and 2x PCIe speed — up to 160 PCIe Gen4 • 60% faster interconnect fabric with xGMI-2 at 16 GT/s (up to 24 GT/s with a supported peripheral)
Are designed to make full use of processor capabilities for integrated end-to-end security	• Secure data at rest with AMD Secure Memory Encryption (SME) • Workload VM to hypervisor isolation with AMD Secure Encrypted Virtualization (SEV) — 509 unique keys per hypervisor

Let Fit-for-Purpose Design Guide You Through Technical, Business, and Financial Decisions for Your Datacenter

A fit-for-purpose infrastructure platform incorporates technology innovations that address critical business priorities:

- Dynamic workload-optimized scaling
- Effortless management
- High-bandwidth capacity
- Low-latency data access
- Faster network connectivity
- Robust end-to-end security

Why fit-for-purpose infrastructure should set your priorities:

FOR THE CXO

- Maximized return on investment
- Reduced total cost of ownership
- Security of data and intellectual property
- Backed by the capabilities of a trusted vendor

FOR THE ENTERPRISE ARCHITECT

- Optimized scaling
- Faster hardware
- Effortless systems management
- Larger bandwidth capacity
- Robust built-in security that starts at the processor

FIGURE 3

Dell Adopters: Hardware Acceleration Technologies

(% of respondents)

Q. Does your organization use any hardware acceleration technologies in your AMD EPYC CPUs?

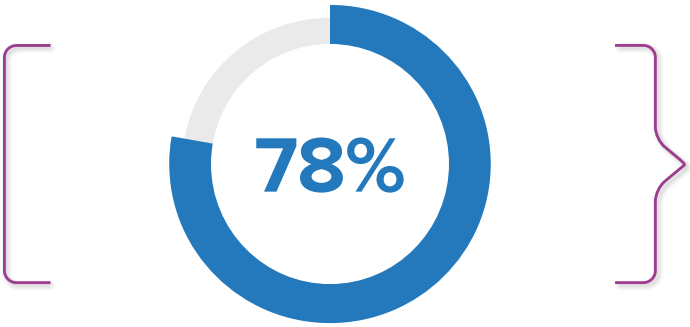

Q. What benefits do your organization's servers with AMD EPYC CPUs provide (or are expected to provide) with regard to accelerators?

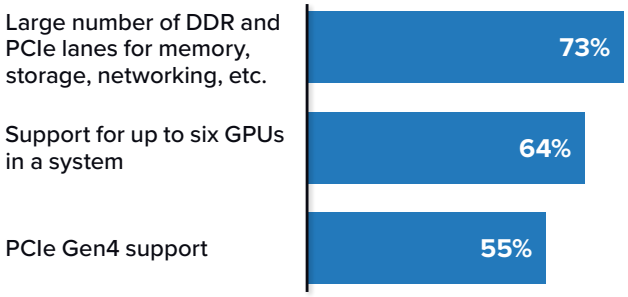

n = 703 (Dell adopter); n = 547 (Using acceleration hardware); Source: IDC *Dell AMD Buyer Behavior Survey*, October 2020; Note: Use caution when interpreting small sample sizes.

Workload Drives the Selection of a Fit-for-Purpose Infrastructure Platform

FIGURE 4

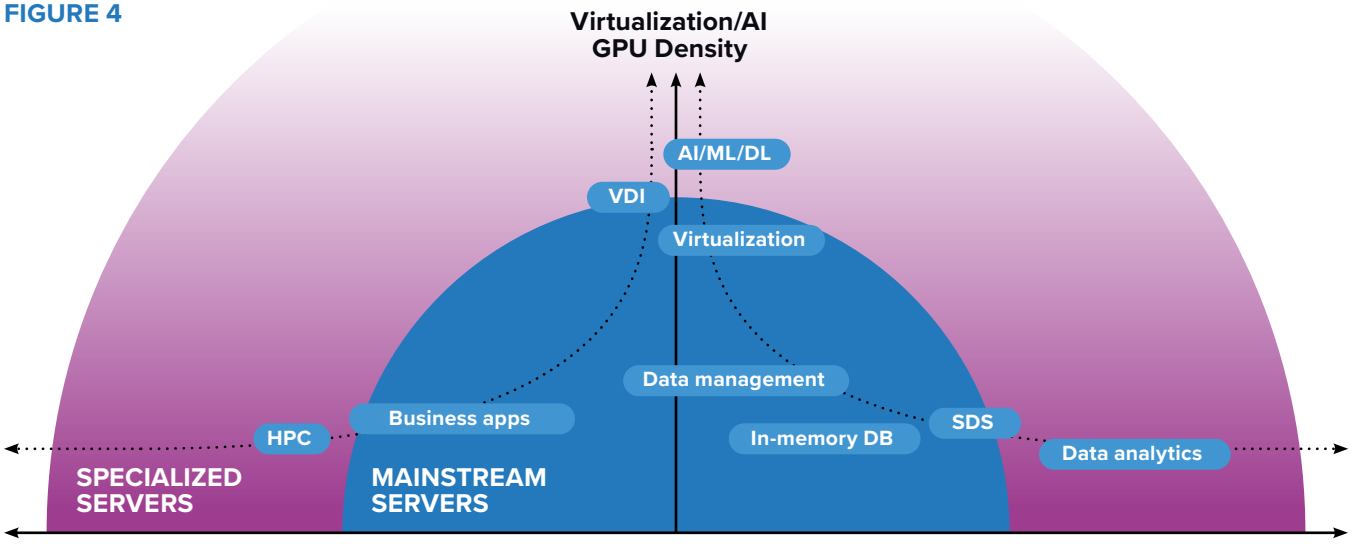

A modern business is a technology and data business. It has multiple functions that need to be deployed efficiently and executed in a timely manner. These may be simple functions such as file and print, standard customer relationship management applications, or complex applications that use artificial intelligence for predictive outcomes.

Each business has unique workload requirements — performance, agility, or investment returns — which, in turn, direct the mix of current and next-generation workloads. Multicloud environments (which are becoming the new normal) place unique demands on IT infrastructure.

The IT infrastructure for each business must support any of these operational or strategic needs. The deployment of an infrastructure platform (servers) therefore must be executed in a workload-centric manner — along five dimensions:

1. **Performance:** How fast can my platform run? How can I accelerate the performance?
2. **Compute scalability:** How well can my server systems scale for compute? An easy way to think about this is on sockets/U-height basis.
3. **Acquisition cost:** What kind of choices are available on a cost/node basis?
4. **Storage capacity:** What kind of storage capacity is available on the system? This is a critical factor for several data-intensive, low-latency applications, expressed in TB/node.
5. **Accelerator density:** Does your application performance increase with GPUs?

...And Then Shifts to Taking a Closer Look at the Profile of Key Business-Critical Workloads

TABLE 2

Compute and Storage Workload Profile	Data Analytics and AI/ML/DL	Software-Defined Storage (SDS)	Modeling and Simulation (HPC)	Virtualization and VDI	Network Virtualization
	Accelerator optimized (with support for GPUs/accelerators)	Storage and I/O optimized	Compute optimized	Accelerator and I/O optimized	Compute, I/O, and connectivity optimized
	Specialized servers	**Mainstream servers**	**Specialized servers**	**Mainstream servers**	**Mainstream servers**
	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
	• Hadoop/Spark • Tensor Flow	• Hyperconverged and non-hyperconverged file, block, and object SDS	• High-performance applications	• VMware Horizon • Citrix Xen Desk/ Xen App • Autodesk	• OpenStack • NSX • AHCI
	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

The approach outlined above must be the first step in a workload-optimized infrastructure platform selection process. It is imperative for enterprise systems architects to examine each workload type for its requirements on compute, I/O, connectivity, and storage optimization. For example:

- **Data analytics** workloads require matching compute with storage capacity and speed.
- **High-performance computing (HPC)** workloads need bandwidth and memory.
- **AI/ML/DL** workloads present IT architects with a key challenge in choosing a platform with the right performance characteristics to tackle complex compute-intensive algorithms.
- **Desktop virtualization** workloads demand reliable performance and bandwidth to match dynamic user demand.
- **In-memory database workloads** require that the platform support the bandwidth and throughput to manage high-volume transactions.

AMD EPYC Processors with Up to 64 Cores Provide a Critical Building Block for an Optimized, Scalable, and Secure IT Infrastructure

Advanced Micro Devices (AMD) introduced its AMD EPYC brand of x86-64 microprocessors in 2017. Based on the company's Zen microarchitecture, the server variant of these processors is specifically targeted for computationally and I/O-intensive datacenter and cloud workloads.

With features such as higher core counts, more PCIe lanes, support for larger amounts of RAM, and larger cache memory, these processors support single- and dual-socket system configurations.

AMD EPYC x86 server processors offer an integrated architecture with higher performance, improved bandwidth, balanced energy consumption, and enhanced security.

TABLE 3

Workload Attribute	AMD EPYC Capabilities	Naples	Rome	Milan +
High-compute parallelism	Higher core count with configurable TDP	• 32 cores • 64 threads • 2 x L3 cache/core (16MB per 4 cores) • 200W TDP	• 64 cores • 128 threads • 2 x L3 cache/core (16MB per 4 cores) • Up to 240W	• 64 cores • 128 threads • 2 x L3 cache/core (32MB per core) • Up to 280W
Faster in-memory workloads	RAM capacity and speed	• 2666 MT/s • Max. system memory 18TB	• 3200 MT/s • 4TB per socket	• 3200 MT/s • 4TB per socket • Option for 6 channel memory interleaving
Massive I/O for data access and connectivity	More PCI lanes	• Up to 128 PCI Gen 3 lanes • xGMI at 10 GT/s	• Up to 160 PCIe Gen4 lanes • xGMI-2 at 16 GT/s	• Up to 162 lane option (2P config) • xGMI-2 at 16 GT/x
Cryptographic isolation between hypervisor and VM	Secure Encryption Virtualization (SEV)	Up to 16 SEV keys per hypervisor	Up to 511 SEV keys per hypervisor	• Secure Nested Paging • Up to 511 keys
Security for business-critical data	Secure Memory Encryption (SME)	Not supported	Configurable	Configurable

- ✓ Single I/O/memory die removes processor subsystem bottleneck, and provides 2x performance over Naples.
- ✓ Uniform speed across cores, and high-bandwidth and low-latency network connectivity.
- ✓ Separate I/O die from Zen2 chiplets* allows flexible core configurations.
- ✓ Dedicated encryption offload subprocessor 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.
- ✓ Improved single-socket and reimagined dual-socket performance reduces TCO.

* 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.
** Non-uniform memory access (NUMA) is 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 non-local 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.

What Businesses Look for in AMD EPYC Server Platforms

FIGURE 5

Key Features of AMD EPYC Servers

(% of respondents)

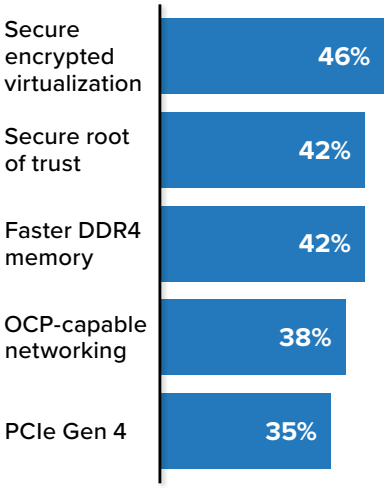

Additional configuration enabled by vendor's design and engineering:

- Design that makes use of higher sustained TDP of 240W per processor
- Risers optimized to balance workloads (flexible configuration)
- Flexible use of PCIe lanes for network and management access
- Socket direct networking for balanced network I/O
- Direct NVMe configuration that eliminates switches or bridges
- Extended speed mode ready for faster peripherals operating at 24 GT/s when available
- BIOS-supported AMD Secure Memory Encryption
- Liquid cooling options via partners
- Reduction in deployment time through automation capabilities and system management solutions
- Improvement in datacenter cooling power utilization efficiency (cooling PUE)
- Easy BIOS tuning with workload-optimized server configuration profiles
- Verifiable secure root of trust that is etched at the factory

n = 703 (Dell adopter); n = 174 (Rome-based adopters); Source: IDC *Dell AMD Buyer Behavior Survey*, October 2020; Note: Use caution when interpreting small sample sizes.

FIGURE 6

(% of respondents)

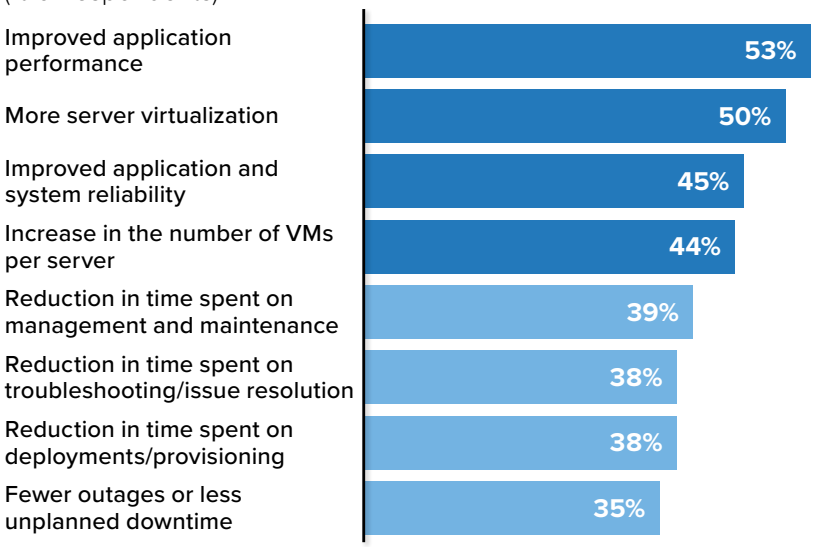

- Extra cores for multi/mixed workloads
- Single-socket option, high core counts, and PCIe 4 lanes
- Reduction in virtualization and licensing costs (32+ core counts in specific environments being the exception)
- Volume of virtual machines supported per server
- Lower cost of CPU supported flat cost for more highly configured systems
- Communication of road map, even if it's "in vague terms and happy feelings"

n = 703; Source: IDC *Dell AMD Buyer Behavior Survey*, October 2020; Note: Use caution when interpreting small sample sizes.

A Trusted Vendor Delivers a Best-in-Class Fit-for-Purpose AMD EPYC Platform

Organizations must consider the role of IT vendors as partners when selecting a suitable AMD EPYC–based server platform. It is crucial for organizations to partner with vendors that are investing in design and engineering capabilities to produce such a platform at scale.

Equally imperative is the level of trust that firms must place on the vendor'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 system recovery is necessary.

FIGURE 7

Key Capabilities of Trusted Vendors

Q: You indicated 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?

(% of respondents)

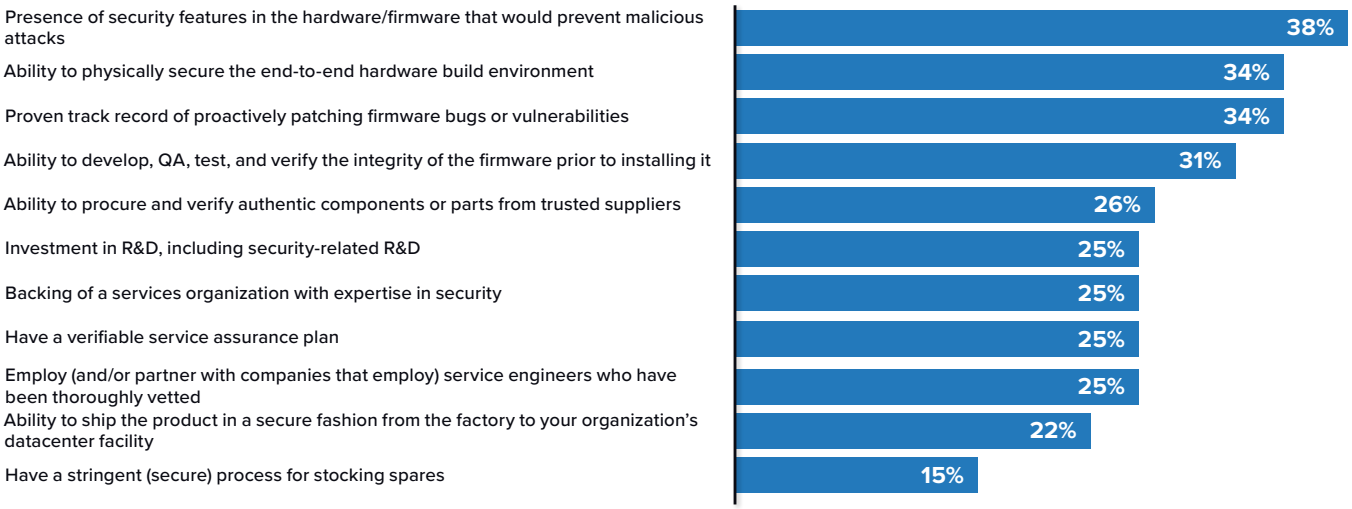

n = 273; Base = respondents who indicated that their organization maintains a roster of trusted vendors for providing infrastructure; Note: multiple responses were allowed; Source: IDC's *Value of Secure Server Infrastructure Web Survey*, February 2018

Essential Guidance

How to Benefit from Investing in Fit-for-Purpose Infrastructure

- Fit-for-purpose IT infrastructure provides a solid foundation on which firms can accelerate their digital journey. IT staff can enable their business to produce consistent outcomes by making timely investments in a workload-optimized and innovation-packed infrastructure platform from a trusted vendor.
- Additionally, IDC studies consistently find that infrastructure operating costs continue to spiral upward if businesses do not upgrade in a timely manner.
- One-size-fits-all approaches to infrastructure platforms can potentially slow down the pace of business innovation due to performance and scalability. At the same time, an aging infrastructure can quickly become a liability for an organization.
- Older or generic infrastructure platforms are vulnerable to data breaches. Recent disclosures have revealed that vulnerabilities whose roots lie in the processor design are difficult to patch without incurring heavy penalties.
- Organizations must therefore discard the one-size-fits-all dictum and replace it with a workload-optimized approach to making infrastructure decisions – which is based on technology and business economics. They must rely on a trusted vendor as a partner to deliver a fully optimized and secure platform. Only then can they enable business outcomes at the pace of innovation.

Breaking the "Good Enough" Barrier: Investing in Infrastructure Is a Business Decision

Server operating costs in years 4–6 of deployed life are more than 10 times higher than the initial acquisition cost of the server.

Compared with continuing to operate installed servers, refreshed servers deliver:

- 59% lower three-year cost of operations
- $4.66 million additional revenue per year per organization ($123,400 per server)

Potential benefits of a regular, faster refresh cycle (two three-year life-cycles versus one six-year life-cycle):

- 33% lower net cash flow
- $14.6 million cash flow over six years per 300 servers
- Less than one year return on investment

Considering Dell Technologies for Its AMD EPYC Server Portfolio

FIGURE 8
Dell's Cyber Resilient Architecture provides best-in-class capabilities for AMD EPYC servers.

Dell Technologies commands a high

satisfaction rate

among its AMD EPYC server customers.

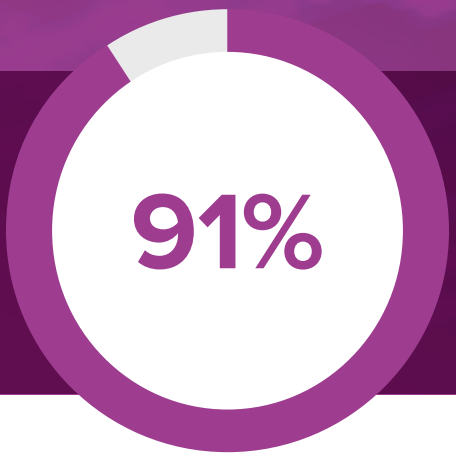

Being a trusted vendor is more than supplying a good product at a fair price. It requires partnering to identify goals and plot the best course of action. Fit-for-purpose infrastructure is an extension of this principle, as is providing quality customer support, services, and financing options. When businesses design their datacenters with this in mind, they deliver superior business, financial, and service quality outcomes. Vendors that help their customers reach these goals achieve a high satisfaction rate.

IDC Methodology

This IDC eBook provides a summary of an extensive validation process performed by IDC in collaboration with the vendor's teams. IDC relied on data from the vendor, as well as IDC's own independent research, to make the statements in this document.

This document is meant to provide a quick set of inferences and insights for IT professionals and business decision makers seeking to perform further due diligence on the capabilities of the product and/or services that have been evaluated in this eBook. However, the goal of this eBook is not to supply detailed, hands-on test plans and validation jobs. It is not meant to replace the evaluation process that most businesses will conduct before making any decision to purchase the product and/or services. It is for this reason that this eBook is designed not to be an all-inclusive document on all the capabilities of the product, but rather as a concise briefing that highlights features/functions of the product, their relative performance with respect to a traditional environment, and the value these features bring to businesses looking to solve certain problems for Hadoop workloads.

Finally, even though this eBook is a sponsored document, it is not meant to be an IDC endorsement of the product, service, or sponsoring supplier. IDC's opinions are its own and not influenced by the production of this document.

For Dell Partners

→ Demonstrate strong partnership between Dell and AMD.
→ All PowerEdge servers get Dell support and services.

Q. What features would excite you most in future releases of servers with AMD EPYC CPUs?
(% of respondents)

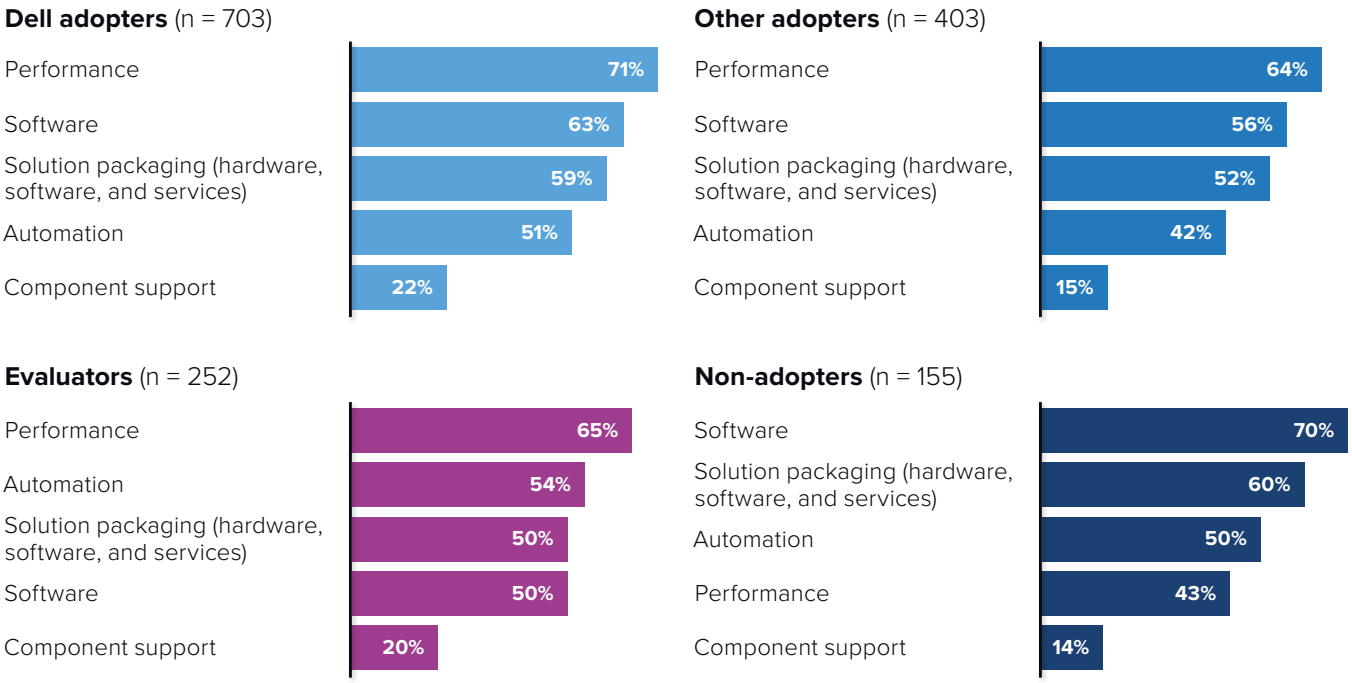

Source: IDC *Dell AMD Buyer Behavior Survey*, October 2020; Note: Use caution when interpreting small sample sizes.

Discover more about PowerEdge servers

Learn more about our PowerEdge servers

Learn more about our systems management solutions

Search our Resource Library

Follow PowerEdge servers on Twitter

Contact a Dell Technologies Expert for **sales or support**

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

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 semiannual Server and Storage Workloads Trackers, primary market research, and custom data modelling.

More about Heather West

IDC Custom Solutions

This publication was produced by IDC Custom Solutions. As a premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets, IDC's Custom Solutions group helps clients plan, market, sell, and succeed in the global marketplace. We create actionable market intelligence and influential content marketing programs that yield measurable results.