

WINNING THE DATA-CENTRIC DIGITAL BUSINESS IN THIS DECADE

DELL'S ADAPTIVE, SECURE, AND RESILIENT PORTFOLIO FOR THE DIGITAL BUSINESS

INTRODUCTION

The enterprise datacenter has evolved from one of a location (or locations) to one that has no physical construct. Today's modern datacenter emphasizes data – and spans the core to the cloud and the edge.

It is critical for the modern IT executive to understand this, as it informs the strategy and approach to building out the next generation of infrastructure that drives the modern, data-centric business. Data collected across the business must be aggregated, shaped, and transformed into intelligence that runs the digital business and drives rapid business decisions based on near real-time data.

The data-centric business requires data-centric infrastructure – i.e., infrastructure designed and optimized to enable the fastest time to value, time to action, and time to outcome. Infrastructure that delivers this business outcome with the most robust security and availability profile while simultaneously reducing costs.

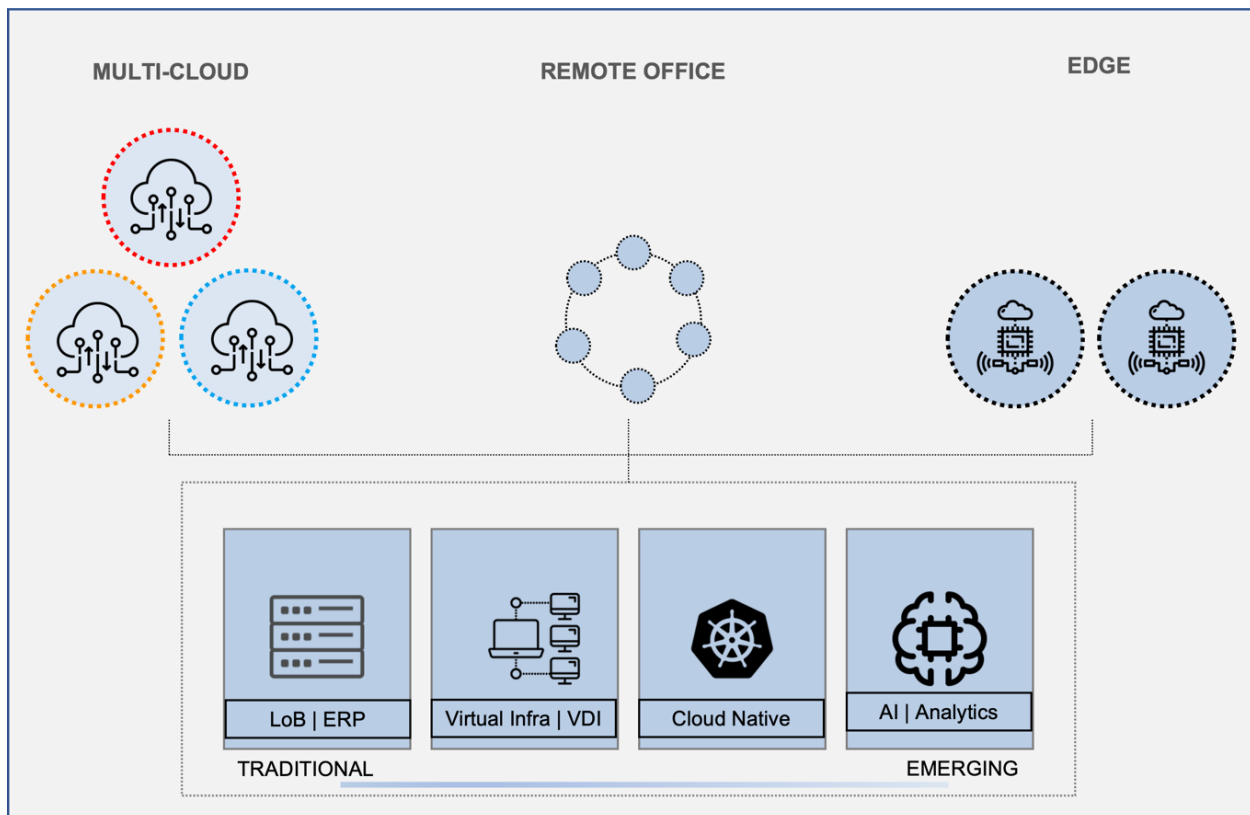
This brief will explore how the Dell PowerEdge 16G server portfolio, powered by 4th Gen Intel Xeon Scalable processors can support the range of workloads and applications critical to the modern business – from distributed cloud-native applications to big data analytics to artificial intelligence and machine learning (AI/ML).

DATA DRIVES THE BUSINESS – SPEED DEFINES THE WINNERS

Only 52 of the Fortune 500 from 1955 are on the list as of [2021](#). Digging a little further, between 1990 and 2021, [71% of the Fortune 100 companies dropped out of the list](#). And between 2010 and 2021, the 23 new entrants to the Fortune 100 list represented disruptors across a variety of major industries. As interesting as this is, some believe most of today's Fortune 500 companies won't exist in 30 years, due to digital disruption. As the rate and pace of innovation increases, those organizations that embrace and leverage this innovation will continue to surge ahead, while those that resist fall behind.

Across industries and around the globe, both accuracy and timeliness are becoming increasingly important and interconnected. Smart factory floors that can manage assembly lines for maximum productivity can save a company millions of dollars. And an e-commerce site that can respond to a customer inquiry measurably faster – with relevance – will beat out the competition. These aren't simple examples designed to make a point; they are real-world dynamics leading to volatility.

FIGURE 1: THE MODERN DATA CENTER ENVIRONMENT



Today's business requires an "all of the above" strategy to support the wide range of workloads
Source: Moor Insights & Strategy

THE MODERN BUSINESS REQUIRES THE MODERN DATACENTER

As enterprise IT organizations strategically plan to support this new business requirement - where data is collected and used everywhere – accounting for accessibility at the edge and the cloud have become table stakes. And while many organizations have sourced or migrated functions to the cloud, recent rationalization efforts have found that most enterprise organizations have a balance between on-premises and public cloud services – a hybrid, multi-cloud strategy. Moor Insights &

Strategy (MI&S) finds strong anecdotal evidence to support this conclusion through our conversations with IT organizations that span industries and sizes.

Enterprise organizations require a cloud operating model – consumption-based and agile. The business needs development environments, data management, and analytics in an "as-a-Service" fashion. To support this requirement, IT requires highly performant server platforms capable of supporting the most demanding workloads. These platforms must not only be performant, but secure, cost-effective, and easily managed. Server infrastructure is the critical cornerstone of the modern datacenter – and critical to the IT executive measured on responsiveness to (and success of) the business.

The workloads and applications running in the modern datacenter are diverse. Resource-hungry workloads and applications get more benefit from servers that are more richly configured, whereas lightweight, highly distributed workloads such as cloud-native containerized applications can run in a "scale-out" kind of fashion. These resources include CPUs (sockets) to support processing more data faster, memory capacity for large data footprints, accelerators to offload processing needs, storage for data locality, or some combination of them all.

EXPLORING THE BUSINESS-CRITICAL APPS IN THE ENTERPRISE

As a business continues to evolve and modernize to stay competitive getting to outcomes faster is highly dependent on its ability to generate and act on intelligence. And this responsiveness is, in turn, tied directly to the underlying server infrastructure.

The most effective server platforms for managing enterprise data can house large amounts of data in close physical proximity to compute. An architecture that can balance many cores to large memory footprints, closely connected to high-performing local storage, is ideal for business-critical workloads.

Enterprise use cases and deployment models where we believe a richness of server infrastructure is required to drive successful business outcomes:

- 1. Enterprise resource management and planning**
- 2. Data analytics**
- 3. High-density virtualization**
- 4. Artificial Intelligence and machine learning (AI/ML)**
- 5. High-performance computing (HPC)**
- 6. Enterprise virtual desktop infrastructure (VDI)**

MI&S identifies several enterprise use cases and deployment models where we believe a richness of server infrastructure is required to drive successful business outcomes:

1. **Enterprise resource management and planning:** As data becomes more vital to an organization, the underlying infrastructure used to collect, synthesize, transform, and analyze this precious commodity grows in importance.

The amount of time it takes to move data from storage to memory and compute (latency) increases as datasets stored in databases such as Oracle, Microsoft SQL Server, and SAP HANA grow. Because of this, investments in converged and operational databases that have higher (and better performing) storage, combined with high memory capacity and more compute nodes, will significantly benefit the analysis of bigger datasets.

The alternative – multiple server platforms that require analytics to be more distributed – results in a longer time to outcome. Likewise, cloud-based data management services can result in performance (latency) degradation and unexpected costs as datasets grow. Latency and costs are key factors that turn a leader into a laggard.

In short, performance is so critical that companies such as SAP provide standard benchmarks to assist IT organizations select optimal solutions.

2. **Data analytics:** Tightly coupled with enterprise data management, data analytics is more than simply running a report against data residing in a SQL database. This is deep analysis on the transactional, historical, structured, and unstructured data collected across the business – in the datacenter, on the edge, and on devices. Being able to glean accurate insights from all this data is what separates winners from losers.

One of the more popular use cases MI&S has seen emerging is around real-time analysis of transactional data – otherwise known as hybrid transaction/analytical processing (HTAP). By performing analytics on transactional data in-memory, the requirement for extract, transform, and load (ETL) into a data warehouse or data mart is removed, resulting in much lower latency. For HTAP, a server platform must have a large memory footprint and rich compute capabilities.

3. **High-density virtualization:** Virtualization took hold in the enterprise to drive down the total ownership cost through higher server utilization. Higher virtual machine (VM) density levels translate into more significant savings. For enterprise organizations looking for a reduction in management complexity and

costs, virtualization on server platforms with the richest compute resources is a must.

4. **Artificial Intelligence and machine learning (AI/ML):** AI/ML efficiency is highly dependent on the ability to train and analyze *more data faster*. Data locality is also critical to the speed and accuracy of training models and inference.

Equally important to maximizing the compute-memory-storage footprint per rack unit (RU) is the ability to provide acceleration for AI/ML workloads. Graphics processing units (GPUs) are the most widely deployed AI/ML accelerators and require additional PCIe support. It is important for IT organizations to map the needs of their AI environment in terms of local storage and memory and the underlying PCIe lanes to support the required acceleration.

A real-world use case where AI/ML has proven invaluable comes from the financial services industry. ML is used in support of risk/return scenarios – a good example being the processing of consumer loans. A decision tree algorithm supported by GPU offload enables banks to quickly assess risk and deliver an explained decision.

5. **High-performance computing (HPC):** HPC is a very broad workload term. And workload requirements change, based on the specific application types. How critical is data locality to application performance? Are high message passing and low latency important? Some workloads, such as high-frequency trading, require the lowest latency a system can provide, which only comes through large, localized storage and memory.
6. **Enterprise virtual desktop infrastructure (VDI):** VDI deployments have seen significant growth since the COVID pandemic and have been invaluable in building secure collaboration for distributed workforces. Cost-effective deployments of VDI support the range of users that comprise large organizations.

While highly performant VDI implementations require adequate compute, performance constraints tend to be found in memory allocation and GPU sharing for rich user experiences. Servers rich in CPU resources that can pack a lot of memory and support rich GPU configurations are best suited for enterprise VDI.

ADDRESSING THE DATA DELUGE

The theme of this brief is that data drives the modern business. A business's success or failure is tied to the ability to generate good intelligence based on the seemingly infinite sources of data it collects – and acting on that intelligence faster than the competition.

That said, data management in the enterprise has evolved considerably. The days of a SQL data warehouse ingesting data from SQL-based database instances across various departments are long gone. Data now comes in all forms and sizes: structured, unstructured, SQL, NoSQL, graph, document, data lakes, data lakehouses.

It is not uncommon for the modern enterprise to utilize a range of these database types and deployments to support its digital transformation. And deploying is not a function of IT bending to the whims of the business; instead, it is because each database type serves a particular purpose that helps drive to that modernized end state.

Further, there is a tight coupling of these platforms and the cloud-native applications that generate and utilize diverse data types.

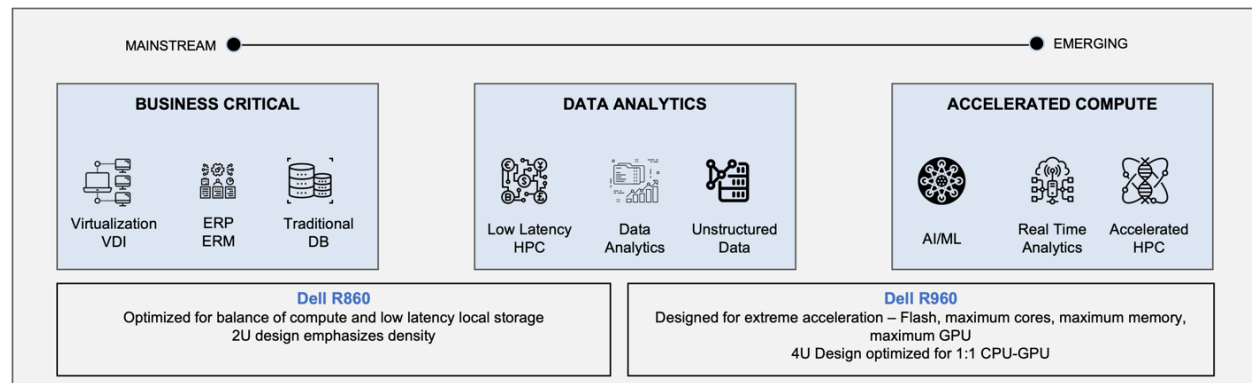
EXPLORING & ASSESSING DELL'S PORTFOLIO – AND STRATEGY

Dell is a leader in the IT solutions space, [claiming the top spot](#) for storage and server units shipped. This success is grounded in a pragmatism the company has toward entering markets. Dell has a long history of delivering products to the market at just the right time as a result of good product planning and strong supply chain management.

In conversations with Dell executives about its portfolio, two themes emerged: tailoring designs to outcomes and enabling the intelligent, data-centric enterprise. It is clear that Dell understands that the core of the modern business is the data executives and business managers use to derive value. And it is this understanding that has led to design principles the company employs to ensure innovation is fueling the market's needs – today and tomorrow.

In support of this strategy, the Dell PowerEdge server portfolio meets the computational demands for the range of workload requirements we have discussed in this paper. On one end of the portfolio are the single and dual-socket servers that support the highly distributed, scale-out needs of lighter-weight applications.

FIGURE 2: DELL FOUR-SOCKET SERVER PORTFOLIO ALIGNMENT BY DATA TYPE



Dell's four socket portfolio addresses the compute hungry needs of modern workloads

Source: Moor Insights & Strategy

For the data-centric workloads that power the business, Dell employs a rich 4-socket design that aligns the diverse needs of these applications, such as richer compute, more memory, high-performance localized storage, and/or performance enhancement through GPUs or other accelerators. Below is a categorization of workload types, and the optimal Dell server platform:

- **Business-critical** – The applications that power business-critical functions are diverse and require a platform that is performant and extremely flexible. The Dell PowerEdge R960 is well positioned for this performant flexibility, targeting workloads that require a combination of large memory and storage footprints.
 - Enterprise resource planning and management (ERP/ERM), dense virtualization and VDI, and line of business (LoB) applications requiring data analytics are good candidates for the R960
- **Accelerated Compute** – AI/ML, database acceleration for real-time analytics, HPC, and next-gen sequencing are examples of workloads that require a GPU richness to enable best performance. And for this class of workload, Dell designed the PowerEdge R960. Workloads that require rich acceleration will find this server to be ideal as it enables a 1:1 ratio of GPU to CPU for extreme application performance.
 - In particular, MI&S sees the R960 as a robust platform for ML as the extreme GPU support can significantly drive down ML training time across virtually any training algorithm. The R960 is a prime example of how infrastructure drives faster outcomes.

- **Data Analytics** – Data analytics platforms are optimized when there is enough data residing in memory to feed the many cores in a CPU complex. The emphasis in data analytics is richness of memory. And this is what drove the design of Dell’s PowerEdge R860 server. The R860 is a 2U server seemingly designed for low latency, packing up to 240 Intel Xeon cores and up to 24 NVMe drives.
 - Target workloads for the R860 include HPC workloads such as high-frequency trading, data analytics, and dense server virtualization.

Data comes in different forms and sizes, and database types vary greatly, impacting compute requirements. This is where Dell shows how its innovation aligns with real business needs. The graph above shows how the Dell 4-socket server portfolio aligns with the range of data and database types that fuel the modern business.

4TH GEN INTEL XEON SCALABLE PROCESSORS – DESIGNED FOR ACCELERATION

Innovation in infrastructure is also only innovation if downstream silicon partners can design and deliver compute platforms that Dell, its software ecosystem, and IT organizations can easily adopt. With 4th Gen Intel Xeon Scalable processors, it seems the company has indeed delivered on that innovation.

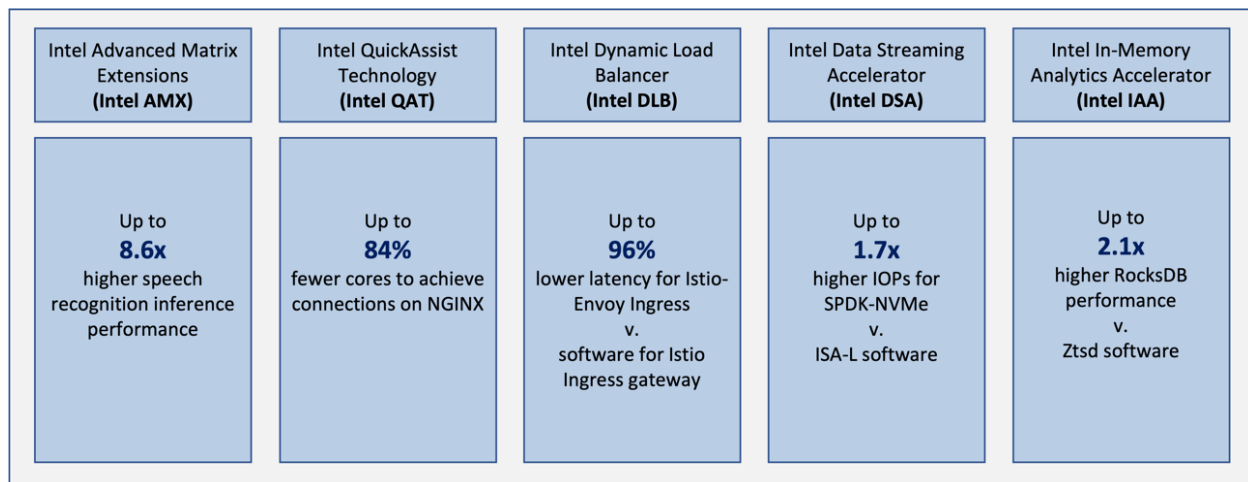
The design focus of this latest Intel Xeon processor seems to have been on seamlessly accelerating the performance of workloads that shine on Dell’s R860 and R960 server platforms. Yes, the requisite high core counts, memory channels, and PCIe Gen5 lanes are included in the 4th Gen Xeon processor packaging. What MI&S finds perhaps more compelling is the addition of built-in acceleration engines, on-chip components dedicated to offloading specific functions from the CPU.

This notion of workload acceleration is nothing new to Intel, as Xeon processors have more accelerators than any other CPU on the market. The addition of several new acceleration engines is especially intriguing:

- **Intel Advanced Matrix Extensions (Intel AMX)** offloads matrix operations such as matrix multiplication and convolution to a dedicated accelerator. This leads to better performance of AI functions like deep learning, training, and inferencing. As a result, workloads such as natural language processing (NLP), image recognition, and recommendation systems perform dramatically better.

- **Intel QuickAssist Technology (Intel QAT)** is an accelerator engine that speeds up computationally intensive functions such as encryption and compression. Additionally, Intel QAT frees up Xeon to focus on core functions.
- **Intel Dynamic Load Balancer (Intel DLB)** balances traffic across multiple CPUs and CPU cores. Workloads that generate a lot of traffic – HPC and cloud computing included – can run faster and more efficiently with Intel DLB routing traffic to free resources in real time.
- **Intel Data Streaming Accelerator (Intel DSA)** is exactly as its name implies. This accelerator engine offloads the most common data streaming, movement, and transformation functions to dedicated silicon. This improves networking, storage, and data-intensive workloads that drive the digitally transformed business.
- **Intel In-Memory Analytics Accelerator (Intel IAA)** accelerates the performance of database and data analytics.

FIGURE 3: REAL-WORLD WORKLOAD ACCELERATION¹



Intel’s accelerator engines drive substantial performance boost

Source: Intel

While these acceleration engines sound interesting in concept, the real-world impact makes them especially compelling. And the results in the above graphic do not lie.

¹ See [A26, W6, N18, D1, N15] at <https://edc.intel.com/content/www/us/en/products/performance/benchmarks/4th-generation-intel-xeon-scalable-processors/>. Results may vary

The workloads that drive the modern business have performance characteristics and requirements that vary, which makes this latest Xeon CPU so compelling. Marrying high-performance cores with rich memory configurations, PCIe v5, and these acceleration engines makes the 4th Gen Xeon Processor ideal for business-critical applications.

Finally, Intel has invested in the software ecosystem to make these accelerators application-aware and developer-friendly. This means a rich set of capabilities that simply work, and another example of processor capabilities that drive true differentiation.

ENTERPRISE MANAGED – ENTERPRISE SECURED

Mission- or business-critical is defined by two qualities – availability and performance. We have discussed performance in detail during this brief. However, the availability of data is where mission-critical starts.

One of the strengths of Dell's PowerEdge portfolio is the readiness of servers that comes by way of the company's security, manageability, and resilience intellectual property (IP).

From a manageability perspective, Dell's approach has been understated yet compelling. Through iDRAC and OpenManage Enterprise, IT administrators can manage their compute environments through an open and centralized console. This combination allows for levels of automation to simplify the infrastructure lifecycle.

And as autonomous IT operations transition from concept to reality, Dell has developed real solutions to further extend automated IT operations into autonomous operations through its CloudIQ platform.

Security is another area where Dell has both focused and delivered. Infrastructure security is a multi-vectored, multi-planed discipline that begins before a server is booted and ends well after a server is powered down. From silicon root of trust to environmental monitoring, Dell claims its servers can detect ransomware and other threats before they can take root – mitigating through removal and remediation.

The importance of Dell's server security capabilities has a direct tie to availability as it enables organizations to detect and respond to malware and ransomware attacks in real time, minimizing the downtime and potential hijacking of data.

CALL TO ACTION

Digital transformation is not a buzzword or trendy topic – rather, it is a business imperative for most businesses to stay competitive in the market. The driver for change is a new competitive landscape comprised of familiar players and new "cloud-born" companies that thrive on the responsiveness that comes from agility.

While an established enterprise organization may not yet have fully achieved the agility associated with cloud operating models, it does have something that these new market entrants don't have – historical data.

And using this data to inform and drive business strategies, along with a cloud operating model, can help

The 4-socket server is not only here to stay; MI&S sees it as a critical part of the enterprise datacenter .

established businesses not only fend off the digital upstart but deliver true separation in the market.

While many business users and IT organizations associate "cloud" with commoditized, scale-out server farms, the resources needed to transform seemingly infinite amounts of historical data are anything but commoditized. The ever-growing and ever-diverse needs of the digital business require a complementary compute platform. The 4-socket server platform can deliver the right resources of compute, memory, storage, and acceleration – and the balance between these resources to help the enterprise achieve outcomes faster.

Established and emerging workloads, from traditional databases and HPC to emerging workloads such as AI/ML and data analytics, benefit from the richness of resources and proximity of data to compute. This richness of resources for mission- and business-critical operations can only be delivered by a 4-socket server. And these workloads, in turn, drive the enterprise, service customers, and help decide the strategic direction a business should take.

Innovation in infrastructure is only innovation if it has real-world benefits. And real-world benefits are only realized when upstream partners and vendors across the software and services fully utilize and exploit such innovations. To this end, Dell has enjoyed broad software ecosystem support for its platforms. Its presence in the market requires ISVs and other IT solutions providers to optimize and support day 0 operations.

The 4-socket server is not only here to stay; MI&S sees it as a critical part of the enterprise datacenter. It is a building block of the digitally transformed business that is driven by data and designed for the future.

MI&S recommends IT executives and architects align infrastructure to the workloads guiding business outcomes. Is latency a factor? How much local storage is required? What about workload acceleration? These are considerations that should be mapped out and part of the evaluation criteria for server rightsizing decisions.

Dell is an established player in the datacenter, and its 4-socket portfolio is well-aligned to these varied (and sometimes competing) workload characteristics. It has smartly designed and mapped its portfolio to meet the needs of the workloads that power the datacenter. Further, MI&S sees an alignment to the emerging workloads of today that will be mainstream tomorrow.

Finally, Dell's management and security capabilities demonstrate an understanding of what IT operations teams manage daily. This frees IT organizations to automate many functions and free resources to better support the business.

For businesses looking to harness the power of their data, Dell should be given strong consideration.

For more information on the Dell 4-socket server portfolio – visit [here](#).

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR

Matt Kimball, Vice President and Principal Analyst at [Moor Insights & Strategy](#)

PUBLISHER

Patrick Moorhead, Founder, President, & Principal Analyst at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

This paper was commissioned by Dell Technologies, Inc. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2023 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.