

Future Defined: AI Innovation through Smart Data Insights

How the collaborative power of Dell Technologies and AMD is helping organizations harness AI and deep data insights, paving the way to a stronger future.

Jay Limbasiya
Global Strategist, AI | Data Science | Data Management

Introduction

We stand at the edge of a new era. Today, we cannot fully imagine the novel possibilities ahead. However, one thing is certain: life as we know it is about to change – in a profound, amazing and irreversible way – thanks to artificial intelligence (AI) and its burgeoning potential.

Activating AI is all about the data. Estimates show that data practitioners spend 80% of their time just finding and organizing data source. That makes deriving value from data challenging. AI cannot be effectively deployed without a proper data strategy and data management approach. I am passionate about helping companies adopt and implement a data-first strategy – it unlocks the power of technology.

Humans and the role of technology in our lives are becoming intertwined. As they do, AI technology itself is becoming more “human” in its capabilities, poised to extend and enhance our lives as never before. From predicting and preventing personal chronic diseases years before they occur to curing cancer, uncovering mysteries of the universe, skyrocketing operational efficiency and elevating our business potential, AI promises a new world filled with incredible possibilities and innovations.

Dell and AMD are collaborating to drive AI innovation forward. Our shared goal is to help organizations embrace this new and exciting data-driven, AI-enabled era – for today, in the near term future and beyond.



80% of a data practitioner's time is spent finding the right data. Solving this problem can greatly accelerate time to value.

When our customers succeed, we succeed.

Helping our customers win is why AMD and Dell are committed and passionate about innovating for the future, both individually and together. We know that innovation never rests. And collaborations – like AMD and Dell working together – further accelerate and elevate innovation to produce breakthrough opportunities, especially in AI.

Artificial intelligence (AI) is everywhere, yet just scratching the surface of its tremendous potential. AI is dependent on one important enabler: Data. Today, many organizations are wrestling with the right data strategy and approach to set up successful AI use cases and deployments at scale. AMD and Dell are helping organizations do just that and unleash AI's full potential.

We hope this eBook is a resource to organizations on a data journey to modernize data centers for AI and analytics, catalyzing new possibilities. Ultimately, those organizations will be tomorrow's winners. We are here to help them on that data journey and ultimately lead to a stronger, brighter future for all.

Jay Limbasiya

Global Strategist, AI | Data Science | Data Management



Contents

Chapter 1: The AI Revolution	5
Chapter 2: Fundamentals of a Data Strategy	9
Chapter 3: Exploring Data Lakehouses.....	13
Chapter 4: Modernizing Data Centers and Embracing Tech Refreshes.....	17
Chapter 5: Case Studies and Success Stories	19
Chapter 6: Future Trends and Predictions	23
Shaping a Stronger Future with AMD and Dell Technologies	26



Chapter 1: The AI Revolution

Understanding AI and Its Impact on Industries

AI is a revolution, now capable of doing things that have previously been beyond the reach of humankind and technology. While organizations understand the potential, they struggle to deploy AI at scale. Consider these predictions:

- 90% of organizations support AI for gaining a competitive edge over rivals.
- 75% of CEOs surveyed believe the organization with the most advanced generative AI wins
- 72% of business leaders believe AI will improve employee productivity

The power of AI is unmatched in speed, scale and accuracy...and somewhat incomprehensible.

—

“Some computers have now crossed the exascale threshold, meaning they can perform as many calculations in a single second as an individual could in 31,688,765,000 years.”



AI's impact on industries:

Healthcare

- Predicting and preventing chronic diseases a decade before they impact lives
- 90% of \$4.1 trillion annual healthcare expenditures are for people with chronic conditions
- Enabling personalized medicine
- Catalyzing new drug discovery

Finance and banking:

- Improving customer interactions (chatbots and financial assistants)
- Detecting and preventing fraud
- Tracking and reporting regulatory compliance

Manufacturing

- Mapping and tagging IoT data with factory schematics and diagrams
- Automating processes
- Proactively managing supply chains

Energy

- Bolstering resilience against natural events (fires, blizzards, etc.)
- Complying with regulatory and environmental requirements
- Improving customer experiences

Education

- Assisting teachers and tutors
- Supporting students
- Customizing educational plans



Key Technological Capabilities Needed for AI

- Data is the foundation of AI
- Tremendous compute power is required to process AI workloads – at AI speed and scale
- Data management is essential to handle volumes of unstructured data – at AI speed and scale
- Modern data storage is needed to provide distributed storage, data compression and efficient data indexing – at AI speed and scale

[Link to learn more about Dell UDS](#)



Dell & AMD’s Role in Powering AI Innovations

Dell and AMD are collaborating to power data-based AI innovations. Each company brings state-of-the-art technology forward and together, integrates that technology to achieve next-level AI capabilities.

Data Lakehouse



To make your data actionable for AI

A data lakehouse addresses the significant AI challenges of data accessibility and usability – anywhere and anytime. Without a data lakehouse, organizations are stalled and cannot make data actionable and extract AI-enabled insights. With a data lakehouse, organizations gain access to all their data and can analyze it in place, regardless of where the data is located.

AI-optimized infrastructure



To meet the IT demands of AI

AI requires massive compute capabilities for training, retraining and optimizing AI models. Simultaneously, these AI workloads put greater demands on storage infrastructure, mandating scalable, secure and cost-effective storage, designed to handle AI demands. Finally, AI-enabled devices, systems and components should be connected throughout an organization.

Deployment



To scale AI and improve business outcomes

As organizations deploy and scale AI, access to actionable data accelerates insights and improves outcomes. In addition, an open ecosystem and partnerships of leading data technologies, ensure an agile, apt, actionable AI-enabled and scalable approach.



Chapter 2: Fundamentals of a Data Strategy


A data-driven, AI-enabled organization is a future-forward one. And at-scale data and AI capabilities are no longer a nice-to-have option; they are table stakes. AI holds massive potential. A Forrester study and financial analysis found that AI-enabled organizations can gain an ROI of 183% over three years.

A foundational and holistic data strategy is required to unleash the full potential of data – and enable AI at scale. In somewhat of a plot twist, generative AI is accelerating data strategies and AI deployments in organizations due to its quick time to value.

Challenges in Modern Data Strategy / Management

Data is significantly underutilized – across organizational sizes, industries and geographies. That’s because data is complex, siloed, and massive in scale, preventing timely data usage and access when and where it is needed. Often, organizations create stores of “hidden” data that could provide rich insights and inform decision making when uncovered. Of the three data types – unstructured, semi-structured, and structured – unstructured data has the most tremendous and untapped potential. Estimates indicate that 90% of data is unstructured and semi-structured; 10-20% of data is structured.

Further, nearly all growth is in unstructured data.



The most critical ingredient to AI is your data

175 ZB

The total global datasphere will
hit more than **175 ZB** by 2025



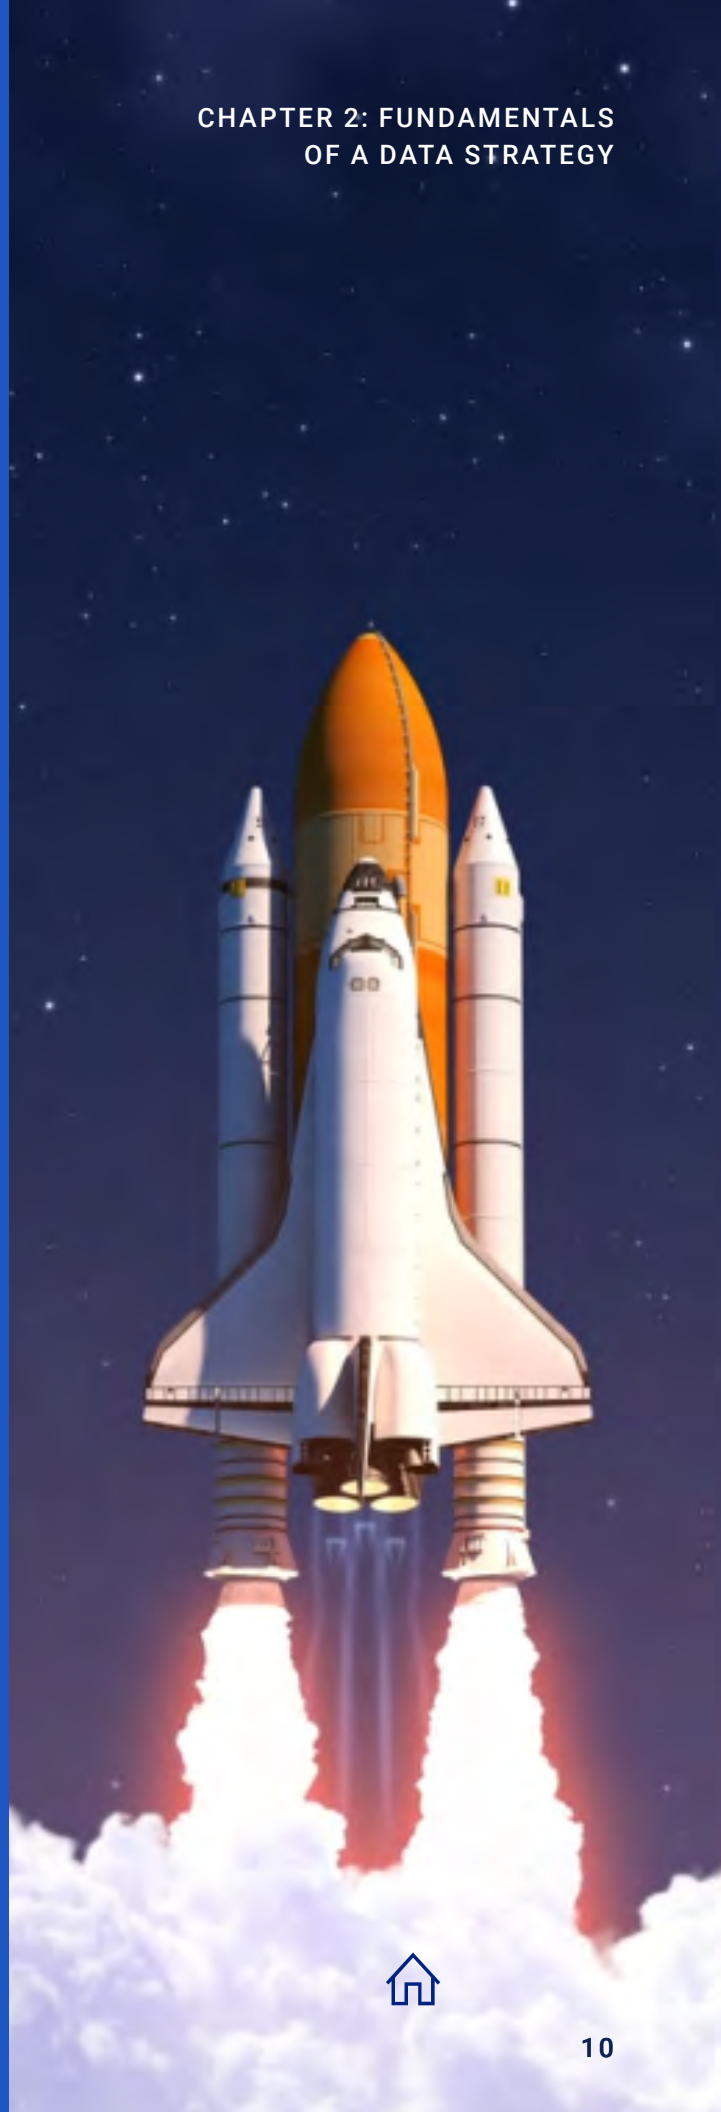
90% of the global datasphere is
comprised of unstructured data

IMAGES • VIDEOS • DOCUMENTS • AUDIO • SENSOR DATA



10% of the global datasphere is
valuable structured data stuck in silos

Source



Implementing a Data Strategy

To overcome these data obstacles, organizations need a holistic data strategy. The goal is straightforward: "How can more people securely extract value from your data?" Achieving this goal is imperative yet far from simple. A structured, intentional approach is paramount, using three high-level steps.

STEP 1.

Locate and classify data across data estate.

Move beyond traditional, highly structured data approaches that fail to account for unstructured data.

STEP 2.

Prioritize and manage your full suite of data for use.

- Identify your most important data (for use) and least important data (for cold storage or deletion).
- Adopt a modern, flexible IT architecture to store, protect, prepare and govern data.
- Include hybrid clouds on-premises, off-premises and the edge.
- Evolve for future scalability to create a single, shared data platform.
- Employ modern storage and other tools to provide data accessibility and processability within sub-second response times.
- Legacy infrastructure often impedes this step.

STEP 3.

Develop data and AI tools.

- Data scientists build interrelated algorithms for use across processes and use-case applications.
- They also gain the option to build data models or leverage existing frameworks models.

[Learn more about a modern data strategy here.](#)



Chapter 3: Exploring Data Lakehouses

When it comes to AI-enabled data architectures, organizations face a case of “what got you here, won’t get you there.” Data challenges are steep. However, to move AI forward, organizations must establish an optimized, future-proofed, holistic data architecture that can deal with escalating data complexity. A data lakehouse does just that.

What is a Data Lakehouse?

A data lakehouse handles all data types via a single repository, eliminating the need for separate systems. It seamlessly integrates infrastructure, software and services into a holistic architecture to mitigate complexity. A data lakehouse unifies access to data, removes multiple areas of ingress / egress and simplifies security and management, achieving both value extraction and security. Data lakehouses support AI and real-time data applications with streamlined, fast and effective access to data.

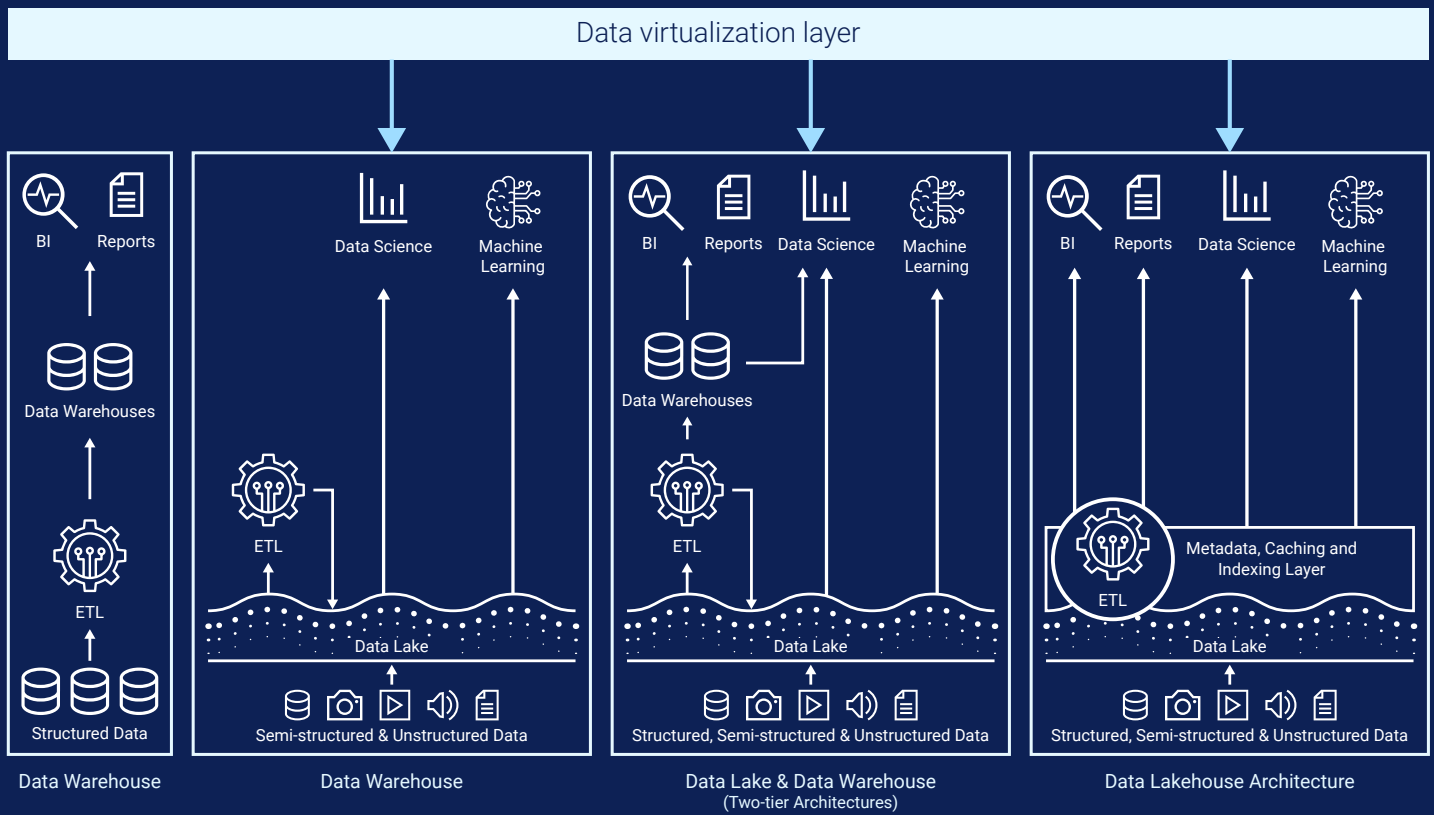
A holistic, integrated approach is essential because a data lakehouse without dedicated storage solutions struggles with scalability, performance, and data management challenges, leading to inefficiencies and potential data silos. In contrast, the Dell Data Lakehouse is purpose-built to work optimally with Dell object storage, ensuring scalability, security, and efficient storage to handling diverse data types and workloads. This results in streamlined data operations, enhanced performance for analytics and data science tasks, and better overall cost management.

Research from Boston Consulting Group found that among more than 50% of data leaders, architectural complexity is a major pain point – forcing organizations into significant complexity, generating avoidable costs and losing critical time to value.

Benefits of Data Lakehouses over Traditional Data Warehouses and Lakes

Data architectures are evolving to meet expanding data and AI needs while mitigating inherent complexities.

[Learn more about solutions leveraging Dell Data Lakehouse with AMD compute.](#)





Data warehouses

Historically, data warehouses have stored curated, structured data to support analytics and business intelligence, with fast, easy access to data. Data warehouses, however, were not designed to support the demands of AI or semi-structured and unstructured data sources. Consequently, traditional data warehouses confine data within proprietary formats, hindering universal access.



Data lakes

Data lakes emerged to help solve complex data organizational challenges and store data in its natural format. Used in tandem with data warehouses, data lakes, while helpful, simultaneously create more data silos and increase cost. As a result, data lakes lack reliability and governance and don't perform well.



Data lakehouses

Data lakehouses combine the best of data warehouses and data lakes, without their inherent limitations. The benefits of a data lakehouse address complexity, value and security:

- Create more value quickly and efficiently from all data sources
- Simplify the data landscape via carefully engineered design features
- Secure data and ensure data availability at the right time for the right requirements

The Dell Data Lakehouse can do this in several ways.



Eliminate data silos

Enhance data exploration, accelerating time to insights by up to 90%¹ and revealing usage patterns that enable smarter data centralization into the data lakehouse.



Unleash performance at scale

With a distributed, massively parallelized engine running on tailor-made infrastructure that separates compute and storage, achieve unparalleled performance that scales as your needs grow.



Take control of your data

100% open format driven and future-ready with modern industry standards such as file formats like Parquet, Avro, ORC and table formats like Iceberg and Delta Lake. Built-in data governance helps you remain in control of your data and empowers you to navigate evolving landscapes with confidence and clarity.



Democratize insights

Give your data team self-service access so they can create high-quality data products, fostering a culture of collaboration and exploration to move your business forward. Integrate with a wide ecosystem of tools such as BI, AI and ML tools, enabling a wider reach for innovation across the organization.



One simplified platform

Designed to streamline deployment, lifecycle management and support services, this turnkey solution encompassing compute, software and storage components delivers a cost-effective and predictable outlay versus cloud-based options. Dell Data Analytics Engine enables 3x faster time to insight at half the cost of other comparable technologies.² Dell ECS storage can save up to 76% in total cost of ownership versus public cloud offers.³ Finally, Dell Lakehouse System Software built on AMD EPYC can deliver significant operational savings by reducing manual effort across the lifecycle.



Chapter 4: Modernizing Data Centers and Embracing Tech Refreshes

Data lakehouses are an important foundational step to enabling AI and a key but partial component to a modernized data center that can enable digital transformation. As AI is deployed at scale and business needs fluctuate and evolve, it's important to fully modernize data centers and embrace ongoing IT management and tech refreshes. That's how IT keeps pace with the accelerating speed of business changes. Evergreen modernization advances agility, digital transformation and future-proofed business relevance.



The Imperative for Datacenter Modernization and Regular Tech Refreshes

According to a recent survey, organizations are transforming digitally to prepare for a stronger future in [three primary ways](#).

- **54%** Become more operationally efficient
- **47%** Develop more data-centric products and services
- **45%** Deliver better customer experiences

AI is a core ingredient in enabling these capabilities, making data center modernization essential. Traditional IT infrastructures are limited in handling the high workload requirements of AI and large data sets with flexibility and speed. Mounting IT complexity and skyrocketing data growth introduce challenges to achieving data center modernization.

Key Components of a Modern Datacenter

Dell and AMD have partnered to jointly design and develop solutions that deliver accelerated performance, automation, efficiency, security and sustainability. Dell PowerEdge servers with AMD EPYC™ processors serve as the backbone of modern data centers and offer scalability to handle the changing, high workload demands of AI and large data sets.

Leveraging AMD and Dell Technologies for a Future-Ready Data Center


[Dell PowerEdge servers with AMD EPYC™ processors power the AI and the modern data center.](#) The solutions deliver industry-leading performance and efficiency for AI workloads and are time-tested. Results include:

- **Up to 121% higher performance** comparing 96-core top-of-stack 4th Gen AMD EPYC processor with 64-core 3rd Gen AMD EPYC processor using SPEC CPU.¹
- **Up to 64% savings** in CPU power.²
- **Up to 34% improved performance per core for virtualization** comparing the 64-core 3rd Gen AMD EPYC processor with the 96-core 4th Gen AMD EPYC processor using VMmark®.³



Chapter 5: Case Studies and Success Stories

Every story is unique. These customer stories share one powerful commonality: they are leading innovation forward to push the boundaries of what is possible. AI and high-performance compute deployments are facilitated by cutting-edge infrastructure. IT solutions help answer challenges, improve efficiencies, navigate roadblocks and help deliver next-level business outcomes.



Real-world Applications of AI and Data Management Solutions

“With PowerEdge, we have more control over our data. We’re now able to run our AI models much faster without any surprises.”

– Marcel van Verk, Vice President
Crop Data Science, KeyGene

KeyGene is an international research and development organization focused on delivering crop improvements for sustainable agriculture all over the world. Based in the Netherlands, the company is driving technological innovations in data science, proteomics, metabolomics, cell and tissue technology, genome editing and genomic breeding.

Business needs:



- KeyGene needed powerful compute to process large data sets and overcome data bottlenecks when it comes to artificial intelligence (AI), machine learning (ML) and deep learning (DL) workloads
- KeyGene also wanted to reduce rack space, energy consumption and its overall data center footprint

Business results:



- Supports large and demanding data sets without compromise
- Enables AI models to run 24/7 on-premises with lower cost of ownership
- Accelerates insights and analytics with AI, ML and DL across all data science teams
- Reduces rack space and data center footprint by 50% to help meet the company’s sustainability goals

[KeyGene full case study here.](#)



“It used to take about 10 years to fully sequence a seawater sample. Now it takes about less than a week to analyze and sequence all of the DNA in a sample. That’s how technology has progressed science.”

– Christopher Sullivan, Assistant Director of Biocomputing, OSU Center for Genome Research and Biocomputing

Built on a rich, 150-year heritage, OSU is an international public research university focused on creating a smarter, healthier and more prosperous future. Through its Center for Genome Research and Biocomputing (CGRB), the university runs about 20,000 jobs a day to process massive data sets and enable more sophisticated scientific breakthroughs that help humankind.

Business needs:


- The university needed to upgrade its high-performance compute (HPC) technology to achieve more processing capacity, increased memory and lower operating costs.
- They required tailored configurations to accommodate massive amounts of data for repeated use.

Business results:

- OSU increased productivity and processing capabilities while also reducing operating costs.
- The university accelerated the speed of scientific research – analysis that used to take 10 years now takes about one week.
- They increased the sophistication of scientific breakthroughs to accelerate smarter scientific discoveries and outcomes.

[Oregon State University full case study here.](#)

Want to learn more about how businesses leverage Dell UDS storage and AMD EPYC Servers in their data lakehouses? [Read more here.](#)



[Redacted text block]

[Redacted text block]

[Redacted text block]

[Redacted text block]



Chapter 6: Future Trends and Predictions

Emerging Trends in AI, Data Management, and Datacenter Technologies

While it's tricky to predict the future, one thing is certain: changes in AI, data management and datacenter technologies are happening at breakneck speed. Staying agile and aware is the key to a future-focused approach.

Emerging Trends in Artificial Intelligence

AI democratization is underway and no industry is immune to its impact. Widespread implementation is being catalyzed by tools and platforms that make AI accessible to everyone, moving beyond tech-savvy individuals to include experts, non-experts and those in between.

As AI becomes pervasive across industries and applications, AI models are simultaneously advancing in sophistication and capabilities to encompass deep learning and neural networks. Generative AI applications are evolving rapidly. Quantum computing is supercharging AI algorithms with faster operation and more efficient processing of large datasets.

AI and edge make a powerful partnership. Today, edge applications have only scratched the surface of what is possible. Emerging AI and edge developments will integrate AI with edge devices in more industries and improve real-time processing.

Emerging Trends in Data Management

As AI moves more mainstream, big data will transition to fast data due to advanced AI analytics and real-time processing. AI will also be instrumental in automating data management tasks. In turn, these new capabilities will improve business decisions, enabling smarter choices in real time and help manage data governance and quality.

The need for data integration is spurring the development of new tools and platforms that provide data interoperability in multi-cloud environments.

AI advancements are escalating the importance of data privacy and security. New strategies for ensuring compliance and managing breaches are emerging, alongside advanced encryption and security technologies and data privacy regulations.

Emerging Trends in Datacenter Technologies

With the adoption of AI, datacenters are changing. The growth of hyperconverged infrastructure (HCI) will accelerate and help streamline datacenter management. In addition, edge datacenters that support edge computing are on the rise.

The design of datacenters is also shifting. Most notable is the emerging trend toward greener, more sustainable data centers. Tomorrow's datacenters will use renewable energy sources, employ energy-efficient designs and use advanced cooling technologies. Regulatory requirements will continue to drive new trends and technologies.



Preparing for the Next Wave of Technological Advancements

Emerging trends are the foundation that is shaping the decade ahead. For example, the synergies between AI and data management will supercharge integration efforts and drive autonomous data center management. In automation, the rise of AI, robotics and IoT will create next-level operational and business benefits, as well as challenges. And, innovations in cloud interoperability will forge hybrid and multi-cloud strategies that enable seamless data and workload migration across diverse environments.



Shaping a Stronger Future with AMD and Dell Technologies

An AI-enabled future is filled with promise, potential and challenges. No one company can go it alone. Collaborations and partnerships are the new mandates for business viability. AMD and Dell Technologies are investing and innovating to help bring AI and better business outcomes to industries and organizations worldwide. Our collective efforts can help make your data actionable for AI, optimize your infrastructure to meet the IT demands of AI and scale AI deployments for improved results.

Learn more by contacting us or using the links below. Together with you, we look forward to defining the future.

- [Empowering Enterprises with AI](#)
- [Accelerate AI Applications](#)
- [AMD and Dell Partnership](#)
- [Dell Data Management and Data Lakehouse](#)
- [Dell AI-Ready Storage](#)

Jay Limbasiya

Global Strategist, AI | Data Science | Data Management



- [1] Based on Dell analysis of the submitted SPECfPRate score of 1410 achieved on a Dell PowerEdge R7625 with AMD Epyc 9654s compared to the previous high score of 636 on a Dell PowerEdge R7525 with AMD Epyc 7763 processors as of 11/3/2022. Actual performance will vary.
- [2] Based on Dell testing using the VMware VAMT tool to migrate 380 VMs from 5x 2U servers with Intel Xeon® 8180 processors (TDP: 205 W) to 1 Dell PowerEdge R7625 server with the AMD EPYC 4th Gen 9654 (TDP: 360 W) processors on 11/5/2022. Total CPU TDP of 10x Xeon 8180 processors is 2050 W. Total CPU TDP of 2x EPYC 9654 processors is 720W. The difference is 64% savings in power. VMware vSphere licensing is based on multiples of 32 cores per socket. 5x servers with 2x Xeon 8180 processors would need 10 licenses. 1x server with 2x EPYC 9654 processors would need six licenses, which is a 40% savings. Actual results will vary depending on actual product configuration, usage, operating conditions, power management settings and other factors.
- [3] Based on Dell analysis, in September 2023, of VMmark 3.x server performance score of 40.51 achieved on a PowerEdge R7625 2-node cluster with AMD EPYC 4th Gen 9654 processors with 96 cores each for a total of 384 cores, compared to a prior score of 20.04 achieved on an R7525 2-node cluster with AMD EPYC 3rd Gen 7763 processors with 64 cores each for a total of 256 cores. Actual performance may vary.

