INTEGRATED SOLUTIONS INSIGHTS

HELPING YOU BRING YOUR IP TO MARKET FASTER ON WORLD-CLASS TECHNOLOGIES

MODERNIZE FOR AI. OPTIMIZE FOR AGILITY. DEPLOY GLOBALLY WITH CERTAINTY.

The product design decisions you make today will determine your organization's ability to compete in the AI-driven world of tomorrow.

According to the Gartner® Planning Guide for Cloud, Data Center and Edge, "Organizations must unify and optimize their cloud, data center and edge infrastructures to flexibly distribute AI workloads based on factors like latency, privacy, cost, and compliance. They must simultaneously address the increased complexity, specialized skills and governance challenges that come with managing such interconnected environments."

The five key trends for cloud, data center and edge infrastructure in 2026 are

- Organizations will risk-manage their vendor dependencies.
- I&O will deploy GenAI in production across cloud, data center and edge.
- GenAI will disrupt I&O skill sets and require well-defined scopes for use.
- Modern infrastructure will require a platform-first mindset.
- Edge solutions will evolve to integrate AI-driven analytics and enhanced security.

In the Planning Guide for Cloud, Data Center and Edge, Gartner® advises Infrastructure & Operations (I&O) leaders to focus on five trends. These are: vendor risk, production GenAI across hybrid estates, skills shifts, platform first delivery, and AI driven edge security and analytics.

In order to get to market faster, maximize resources, and get maximum return on investment, solution builders should align offerings accordingly and de-risk execution by using an integration partner.

FIND OUT HOW TO WIN

Build on a proven AI-ready foundation: Dell PowerEdge servers powered by AMD EPYC™ processors deliver the performance, efficiency, and automation required to modernize for AI, enabling consolidation, lowering power consumption, and supporting platform-first architectures.

Execute at global scale with confidence:
Avnet Integrated Solutions turns that infrastructure foundation into production-ready platforms, integrated, validated, and deployed worldwide - so OEMs can move faster, reduce risk, and scale consistently across markets.

The result: Faster time to market. Lower operational risk. Predictable global outcomes.

BUILD FASTER, RUN LEANER, SCALE SMARTER

AI WORKLOAD MIX IS SHIFTING:

According to Gartner®, "By 2029, 70% of large organizations will utilize AI infrastructure platforms to deliver applications based on frontier models, compared with less than 10% in 2025."

PLATFORM FIRST IS MANDATORY:

Modern workloads demand automated, governance rich platforms, not siloed stacks.

ENERGY IS A CONSTRAINT:

Efficiency, cooling innovation, and performance per watt are key in determining feasibility and economics.

CONSOLIDATION ECONOMICS:

Replace up to 7 five year old servers with one modern system giving up to 65% CPU power reduction and up to 86% fewer racks to match legacy performance.*

FORWARD THINKING

By 2029, AI inference workloads in public clouds will surpass AI training workloads, increasing from less than 20% in 2024 to over 60% of total AI workloads in the public cloud.*

*Gartner, 2026 Planning Guide for Cloud, Data Center and Edge Infrastructure, October 2025

Source: https://www.dell.com/en-us/blog/rapidly-evolve-business-operations-with-poweredge-rack-servers

WHAT DO THESE TRENDS MEAN FOR OEMS BUILDING FOR DATA CENTERS?

Recommendations from Gartner® to manage the critical trends shaping infrastructure and operations in 2026.

- **Risk-manage your vendor dependencies** by refreshing multicloud strategies to access best-of-breed AI capabilities and for regulatory compliance. Seek "just sovereign enough" solutions to meet data residency and compliance needs, considering the increased costs and complexity.
- **Deploy AI where it will deliver the greatest impact** and best meet business requirements, whether that be cloud, data center or edge. Utilize modern approaches, such as Kubernetes and AI-aware scheduling, combined with good governance, to effectively manage the range of environments.
- **Utilize GenAI for individual and team skill development** both to build AI infrastructure and to quickly develop prototypes. However, take care not to supplant well-developed and differentiated skills with merely "average" GenAI solutions and results.
- **Adopt a platform-first mindset** by building robust, automated platforms to support modern workloads and drive AI adoption. This approach will enable innovation and operational excellence by delivering agility, scalability, and resilience, while simplifying infrastructure and accelerating value from AI investments.
- **Revamp your edge strategy** by implementing unified data management and cyber-physical systems (CPS) protection. Expand your edge technology and platform into new, adaptive forms of connectivity. Utilize micro data centers to standardize and scale edge solutions.

HOW THIS AFFECTS YOUR BUILD STRATEGY:

These trends converge on a single requirement: purpose-built infrastructure for AI workloads. Modernization isn't optional - it's the future for data centers and the solutions built to operate in them.

KEY TRENDS: DATA CENTER MODERNIZATION FOR 2026 AND BEYOND

Avnet believes that Gartner® 2026 signals a structural disconnect between legacy infrastructure and AI-era demands. Cloud, data center, and edge strategies must evolve to support modern AI workloads.

THE EXTREME-SCALE PARALLELISM CHALLENGE

As GenAI workloads grow in complexity and computational output, they require extreme-scale parallelism, utilizing a unified cloud-native platform. Organizations that continue to rely on legacy systems may struggle to keep pace with these demands, putting their competitive advantage at risk. Older systems typically lack:

- Sufficient CPU/GPU core density for parallel processing
- Necessary bandwidth to feed AI accelerators and prevent bottlenecks
- Gen5 connectivity for high-speed GPU communication
- Power efficiency to manage dramatically increased compute density.

THE PLATFORM ENGINEERING REQUIREMENT

Gartner® emphasizes "Modern infrastructure requires a shift from traditional and siloed methods to a platform-first mindset that delivers the agility, scalability and resilience needed for today's workloads." This translates into moving beyond isolated servers to integrated infrastructure platforms that support:

- Kubernetes and AI-aware scheduling
- Consistent governance across cloud, data center and edge
- Self-service capabilities with built-in guardrails
- Unified management and observability

The bottom line: Servers more than five years old will likely lack the architecture needed for AI workloads. They consume more power than modern servers, occupy valuable rack space and require expensive maintenance, while delivering a fraction of modern performance.

The Balance of AI Workloads Is Moving Toward Inference

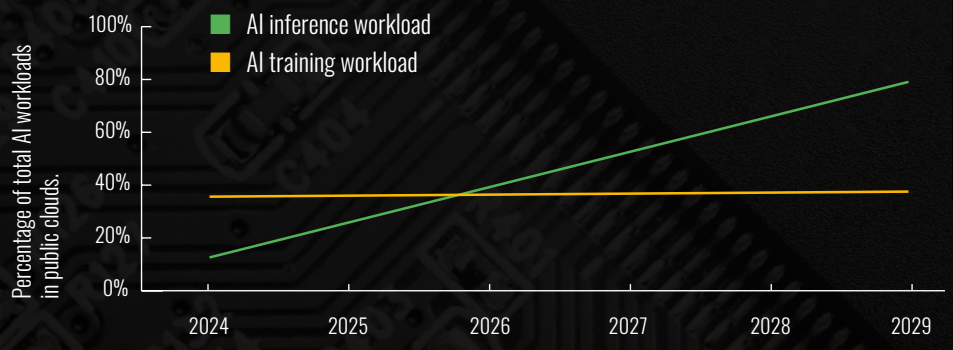

Inference workloads will exceed 60%

60%

of public cloud AI activity by 2029

Why it matters: Inference workloads demand sustained performance, higher availability, and greater operational efficiency than training-centric environments.

AVNET'S PARTNERSHIP WITH DELL AND AMD

The 2026 AI infrastructure trends outlined in the Gartner® Planning Guide are realized through Dell PowerEdge servers powered by AMD EPYC™ 9004 Series processors, with Avnet Integrated Solutions transforming that foundation into production-ready platforms deployed at global scale.

DELL'S EDGE COMPUTING SOLUTIONS OFFER:

Flexible architecture

- Open standards: Avoid proprietary lock-in with industry-standard architectures
- Hybrid deployment: Seamlessly deploy across on-premises, colocation and cloud environments

Proven platforms for AI

- Dell OpenManage enterprise integration simplifies AI infrastructure management
- Validated designs reduce complexity and accelerate deployment
- Extensive partner ecosystem provides training and best practices

Extend AI to the Edge

- Compact form factors for edge deployment
- Integrated cooling for diverse environments
- Remote management capabilities for distributed operations

AMD EPYC™ PROCESSORS CHIPSETS

AMD EPYC™ processor-based servers offer leadership performance and efficiency to enable material workload consolidation, allowing more space and energy to support new AI workloads.

Breakthrough performance

Accelerate productivity, make more informed decisions, and speed time to market with a common platform that delivers the performance you need.

Modernizing infrastructure

When upgrading data center infrastructure, what's required is the most innovative design, high density, excellent energy efficiency, and low total cost of ownership.

Compute with confidence

In order to move forward, it is essential to be able to confidently navigate today's risks, complexities and requirement.

DELL Technologies

AMD

Up to

93%

faster throughput ranking and classification using XGBoost compared to competitor CPUs

Up to

79%

faster Facebook AI Similarity Search (FAISS) compared to competitor CPUs

Up to

3.8X

the throughput for end-to-end AI compared to competitor CPUs

Source: https://www.amd.com/en/products/processors/server/epyc/9005-series/demos.html

THE BUSINESS CASE: PERFORMANCE, EFFICIENCY, CONSOLIDATION

What modern AI-ready infrastructure delivers in practice

AI PERFORMANCE LEADERSHIP

Up to 3.7X performance
on end-to-end AI workloads (TPCx-AI derivative)

1.9X throughput on Meta Llama 3.1-8B inference vs. competition

Up to 37% higher IPC
(instructions per clock) in AI and HPC workloads

DATA CENTER MODERNIZATION

Replace up to 7 five-year-old servers
with a single modern system

Up to 65% reduction in CPU power
consumption

Up to 86% fewer racks
needed to match legacy integer performance

COST OPTIMIZATION

Reduce expensive maintenance
contracts on aging hardware

Lower power and cooling costs
with modern efficiency

Consolidate software licenses
across fewer systems

Source: https://www.amd.com/en/products/processors/server/epyc/ai.html

AVNET INTEGRATED SOLUTIONS ADVANTAGE

From Design to Deployment at Scale

Avnet converts plans into production with global design, integration, supply chain orchestration, and lifecycle services, so that OEMs can focus on IP and growth.

DESIGN & ENGINEERING SUPPORT

- Custom server configurations optimized for specific AI workloads
- Thermal and power analysis to ensure infrastructure works in data centers for your customers
- Integration planning for accelerators, networking and storage
- Validation testing before production deployment

DEPLOYMENT & LOGISTICS

- Rack integration and pre-staging to accelerate deployment
- Global logistics coordination for distributed deployments
- On-site deployment support and knowledge transfer
- Asset tagging and inventory management

GLOBAL MANUFACTURING & INTEGRATION

- Build-to-order capabilities across six global integration centers
- Regional manufacturing to meet local compliance requirements
- Quality assurance and testing at scale
- Secure supply chain management

LIFECYCLE SERVICES

- Warranty management and support coordination
- Break-fix services and spare parts logistics
- Technology refresh planning and migration services
- Ongoing optimization and performance tuning

Avnet Integrated Solutions brings together Dell PowerEdge platforms, AMD EPYC™ performance, and Avnet's global integration and supply chain expertise to help solution providers design, build, deploy, and scale AI-ready infrastructure worldwide—faster, with less risk, and with predictable outcomes.

WEB HOSTING PROVIDER RADICALLY IMPROVES INFRASTRUCTURE WITH AMD & DELL

As a global web hosting company experiencing exceptional customer growth and rapid expansion into new offerings, Hostinger needed to modernize its infrastructure at speed. This modernization enabled Hostinger to scale faster, operate more efficiently, and deliver measurable performance gains across its global footprint.

SOLUTION

Hostinger deployed 800 Dell PowerEdge servers with AMD EPYC™ processors and Dell Technologies support to ensure responsive, reliable customer-facing services and fast server deployments. IT teams use built-in capabilities to deliver Zero Trust environments, automated operational processes and scalable resources to meet demand.

IMPACT

With its latest Dell PowerEdge server deployment, Hostinger has improved its ability to compete, grow and innovate. Customers across 150 countries now enjoy 20% faster services. Additionally, because the new servers deliver up to 50% more cores in a 1U footprint, Hostinger improved the price/performance of its servers by 20%.

HOSTINGER

50%
year on year growth

2 million
customers globally

20%
faster services

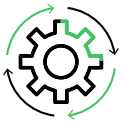

50%
more cores in a 1U footprint

20%
Improved price/performance

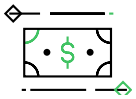

"Given that we operate in a highly competitive market, improving service performance by 20% with Dell PowerEdge servers powered by AMD ... helps us maintain our edge."

Skirmantas Juraška,
Head of Infrastructure, Hostinger

DELL Technologies

START WITH SERVERS THAT ARE OPTIMIZED FOR DATA CENTER PERFORMANCE

Intelligent and cyber-resilient servers built for hybrid and multicloud environments that support a diverse range of applications, including those that require extreme acceleration.

DELL TECHNOLOGIES POWEREDGE RACK SERVERS POWERED BY AMD EPYC™ PROCESSORS

POWEREDGE R6625

1U Dual-Socket Performance Density

- Up to 2× AMD EPYC™ CPUs (up to 128 cores per CPU)
- Up to 6TB DDR5 memory (24 DIMMs, 4800 MT/s)
- PCIe Gen5, NVMe & E3.S storage options
- Up to 3 single-wide GPUs; optional liquid cooling

Best for HPC, Dense VDI and virtualization

POWEREDGE R6615

1U Single-Socket Value

- Single AMD EPYC™ CPU (up to 128 cores)
- Up to 3TB DDR5 memory (12 DIMMs)
- PCIe Gen5 with flexible storage configs
- Excellent TCO and power efficiency

Best for virtualization, Hyper-Converged Infrastructure (HCI) and NFV

POWEREDGE R7625

2U Dual-Socket Scale

- Up to 2× AMD EPYC™ CPUs
- Up to 6TB DDR5 memory (24 DIMMs)
- PCIe Gen5 with GPU and NVMe density
- Optional direct liquid cooling

Best for Data analytics, all-flash SDS, and VDI

POWEREDGE R7615

2U Single-Socket Expansion

- Single AMD EPYC™ CPU (up to 128 cores)
- Up to 3TB DDR5 memory
- Up to 8 PCIe slots with massive storage flexibility
- Up to 6 single-wide or 3 double-wide GPUs

Best for Software defined storage (SDS), virtualization, and data analytics

Source: https://www.dell.com/en-us/shop/data-center-servers/sf/poweredge-datacenter-servers#compare-module

POWER UP WITH AMD EPYC™

AI TECHNOLOGIES BRING BROAD INDUSTRY IMPACT

No matter the size or scale of your AI deployments, AMD EPYC™ server CPUs give you a high-performance, energy-efficient foundation for enterprise AI and general -purpose workloads.

The smart choice for purpose-built solutions
AMD EPYC™ processors lead the industry with higher core counts, higher frequency, and the best performance per watt of any server-class processor.

This means:

- Optimized workload performance: more cores and higher frequency deliver faster application response times.
- Efficiency gains: fewer servers are needed for the same workload, simplifying deployment and freeing up space for future growth.
- Energy savings: EPYC™ powers the most energy-efficient servers, lowering power and cooling requirements for end users.

AMD EPYC™ SERVER CPUS ARE THE BEST CPUS FOR ENTERPRISE AI

5th Gen AMD EPYC™ server CPUs stand up to the demands of AI with options that include high core counts or high frequency, plenty of memory and I/O bandwidth, and support for AVX-512 instructions.

For solution builders, **AMD EPYC™ specs offer up to 192 cores (EPYC™ 9965), boost up to 5.0 GHz (EPYC™ 9575F), 128 PCIe Gen5 lanes/CPU, DDR5 6400** - ideal for data hungry workloads.

CONSOLIDATE 7 TO 1

More efficient modern servers can help data centers consolidate, requiring fewer servers to deliver the same amount of compute power.

For example, it's possible to move from 2020-era Intel® "Cascade Lake" servers to 5th Gen AMD EPYC™ CPU-powered servers.

Fourteen AMD EPYC™ 9965 CPU-based servers can deliver the same integer performance as 100 old servers running Intel Xeon Platinum 8280 CPUs.

Up to
86%
Fewer Servers[1]

Up to
69%
Less Power[1]

Up to
41%
Lower 3-Year TCO[1]

Source: https://www.amd.com/en/products/processors/server/epyc/ai.html

AVNET®

LET'S TALK

Ready to help your customers modernize their data centers?
Avnet believes the Gartner® Planning Guide makes clear that 2026 will be a pivotal year for infrastructure modernization. Help your customers to act now to build AI-ready infrastructure and gain competitive advantages that compound over time.

Ask your customers:

- Are your inventory servers more than 5 years old?
- Do you have workloads that could benefit from AI acceleration?
- Do you know what your compliance and sovereignty requirements are and how they are likely to change in the future?

Avnet can help you:

- Use AMD's performance benchmarks to estimate your consolidation options
- Model power savings from replacing legacy infrastructure
- Quantify data center space reclaimation
- Build a TCO comparison over 3-5 years

GET STARTED TODAY:

Engage Avnet Integrated Solutions to quantify your AI infrastructure advantage before your next platform decision.

Speak to our team to get expert help. We can show you how modernizing your data centers can improve efficiency, ROI and set you up for future success.

Email: integrated@avnet.com
Web: www.avnet.com/integrated

AMD

DELL Technologies

About This Research
Gartner, 2026 Planning Guide for Cloud, Data Center and Edge Infrastructure, Simon Richard, Lucas Albuquerque, 13 October 2025

Sources:
Dell, Dell Technologies and PowerEdge are trademarks of Dell Inc.

By Dell https://www.dell.com/en-us/lp/dt/amd-servers

Dell PowerEdge Server https://www.dell.com/en-us/shop/data-center-servers/sf/poweredge-datacenter-servers

AMD, EPYC™ and combinations thereof are trademarks of Advanced Micro Devices, Inc.

By AMD https://www.amd.com/en/ecosystem/oem/dell.html

AMD EPYC™ Processors https://www.amd.com/en/products/processors/server/epyc/ai.html