

The Compute & Storage Planning Guide

Navigating DDR4 Phase-Down, DDR5
Adoption, and NVMe Demand



The Compute & Storage Planning Guide

Executive Summary

Enterprise infrastructure is entering a period where **planning matters more than procurement**. Memory standards are shifting, storage demands are rising, and platform compatibility is becoming more restrictive. DDR4 is being phased down across server platforms, DDR5 is now the standard for new deployments, and NVMe adoption is accelerating as **workloads demand higher performance and density**. These changes are happening at the same time, compressing timelines and increasing pressure on IT teams.

For many organizations, the risk is not choosing the wrong technology. The risk is waiting too long to plan and being forced into decisions by availability, compatibility, or urgency.

This guidebook is designed for enterprise IT teams planning compute, memory, and storage decisions over the next 12–36 months. It is intentionally educational. The goal is not to sell faster, but to help teams plan earlier and avoid forced decisions later.

Inside, we outline what is changing in the compute and storage landscape, why those changes matter, and how they affect real-world platform decisions. We explore the implications of the DDR4 phase-down, the planning requirements of DDR5 adoption, and the growing role of NVMe in modern architectures. We also highlight the cost of delayed planning and the benefits of taking a deliberate, roadmap-driven approach.

This guide is meant to support informed conversations, not rush outcomes. Teams that plan early retain flexibility, control costs, and align infrastructure decisions with business priorities. Teams that wait often lose options and are forced to react under pressure.

**The infrastructure shift is already underway.
The advantage belongs to the teams that prepare for it.**

Why This Guide Exists

Enterprise infrastructure is entering a period where **planning matters more than procurement**. For years, teams could rely on incremental upgrades and predictable refresh cycles. That model is breaking down as memory standards change, storage demands grow, and platform compatibility becomes more restrictive.

Today, decisions around compute, memory, and storage are tightly linked. A constraint in one area often forces trade-offs in the others. This guide helps teams understand **what is changing, why it matters, and how to plan ahead** - rather than react under pressure.

The Market Transition Underway

The enterprise compute and storage market is going through a major shift. **DDR4 memory, once the standard in servers, is being phased out. DDR5 is now the standard for new deployments.** At the same time, demand for faster storage -- especially NVMe SSDs -- is rising as workloads require more speed and capacity.

Because these changes are happening simultaneously, many IT teams are feeling pressure sooner than expected.

DDR4 Is Being Phased Down


DDR4 will remain deployed in production environments for years, but it's becoming **less common in new platforms**. OEMs are reducing DDR4 production, availability is tightening, and prices are less stable. Many next-generation CPUs no longer support DDR4 at all.

For IT teams, this makes extending existing platforms riskier. What used to be a safe way to delay upgrades now comes with more uncertainty around cost and supply.

DDR5 Is Now the Standard

DDR5 is not an emerging technology. **It is the standard memory architecture for new enterprise servers.** It offers higher performance and efficiency, but it also requires more careful planning.

DDR5 systems must be designed with the right CPU, motherboard, power, cooling, and memory configuration. It isn't a simple swap for DDR4. **Without clear planning, teams can end up overspending or not getting the performance they expect.**



NVMe Demand Is Accelerating

Storage needs are increasing as workloads become more demanding. AI workloads, higher virtualization ratios, and I/O-intensive applications are driving broader NVMe adoption.

NVMe is moving from specialized use cases into mainstream enterprise architectures. As a result, storage decisions now play a bigger role in overall system design instead of remaining an afterthought.

Common Platforms and Components Impacted

The shift to DDR5 and broader NVMe adoption affects widely deployed server platforms across enterprise environments.

Most organizations are not simply upgrading components. They are transitioning platforms.

Category	Current Environment (DDR4-era platforms)	Transition Direction (DDR5/NVMe era)	Planning Impact
Dell PowerEdge	R640, R740	R650, R750, R660, R760	Newer platforms require DDR5 and updated CPUs, power, and cooling. Most upgrades require full platform refresh rather than memory swap.
HPE ProLiant	DL360/DL380 Gen10 or Gen10 Plus	DL360/DL380 Gen11	Gen11 platforms move to DDR5 and newer CPU architectures. Timing matters to avoid rushed refresh decisions.
Cisco UCS	C220/C240 M5 or M6	UCS X-Series and newer M-series	Many environments still DDR4-based. Transition typically requires coordinated server and memory planning.
Supermicro	X11, X12 platforms	X13 and newer	DDR5 support tied to CPU generation and motherboard compatibility. Most transitions require full platform replacement.
Memory	DDR4 RDIMMs (16GB, 32GB, 64GB; higher capacities tightening in supply)	DDR5 RDIMMs (32GB, 64GB, 128GB common)	DDR5 requires platform alignment and stricter population rules. Availability and pricing variability increasing for DDR4.
Storage	SAS/SATA SSDs and partial NVMe adoption	Broader NVMe (U.2, PCIe Gen4/Gen5)	Storage decisions increasingly influence server refresh timing. NVMe adoption often coordinated with platform upgrades.

Older systems are built around DDR4. New platforms are designed for DDR5 and typically require a full platform.

Planning memory and storage together prevents rushed, high-cost decisions later.

The Cost of Waiting

Most infrastructure challenges do not come from poor technology choices. They come from delayed planning.

Teams that wait too long often face limited component availability, forced platform refreshes instead of phased upgrades, and higher costs driven by urgency. Architectures built under pressure tend to solve immediate problems but create other issues in the long-term.

What Good Planning Looks Like

Effective infrastructure planning focuses on workload requirements not just choosing parts. **It aligns platform lifecycles, spreads investments over time, and keeps options open.**

Good plans leave room to adapt as availability, budgets, and business priorities change. Bad plans lock teams into narrow paths with few alternatives.

Lifecycle Strategy and Transition Planning

Infrastructure upgrades rarely happen all at once. **Smart teams plan for a period of overlap between old and new platforms.**


This often includes extending support where appropriate, phasing in new systems deliberately, and recovering value from retired equipment rather than writing it off. **A clear lifecycle strategy reduces pressure and keeps decisions aligned with long-term goals.**

How OSI Global Supports Planning-First Teams

DDR4 production has shifted to DDR5 across the major manufacturers. While DDR4 will remain in market for some time, availability will depend on remaining supply rather than new production. A move to DDR5 is inevitable, but timing still matters.

OSI helps preserve flexibility during this transition by sourcing DDR4 components when OEM channels are constrained or pricing becomes prohibitive, extending current platforms where appropriate, and providing Sustain maintenance as bridge coverage while next-generation plans are finalized. We also offer ITAD and buyback options to recover value from retired equipment.

OSI continues to support DDR4 environments while delivering full DDR5-ready systems across Dell, HPE, Cisco UCS, and Supermicro platforms when the timing is right.



Checklist

Use the checklist below to identify planning gaps and areas that require attention.

Compute and Memory:

- ☐ Identify platforms approaching DDR4 or CPU support limits
- ☐ Determine which workloads benefit from DDR5 today
- ☐ Define a clear memory capacity and density strategy

Storage:

- ☐ Identify I/O-bound workloads
- ☐ Evaluate where NVMe provides measurable benefit
- ☐ Align storage planning with compute refresh timelines

Lifecycle and Timing:

- ☐ Assess exposure to component availability risk
- ☐ Identify alternative configurations
- ☐ Extend support where appropriate and evaluate buyback options

Final Thought

The infrastructure shift is already underway. The risk is not change, but waiting until the market removes your options.

Teams that plan early retain control over cost, performance, and timing.