




Modern Data Center Consolidation, Migration, and Build Considerations in the Context of Artificial Intelligence Demands

In the era of exponential data growth and burgeoning artificial intelligence (AI) applications, the role of data centers has evolved dramatically. With the increasing demands placed on infrastructure by AI workloads, organizations are reevaluating their data center strategies. This whitepaper explores the nuances of data center consolidation, migration, and build considerations in the context of AI workloads, addressing the challenges, best practices, and emerging trends shaping the modern data center landscape.



Introduction

The proliferation of AI applications and hardware has refocused industries across the globe, driving unprecedented levels of data generation and processing requirements. As organizations harness the power of AI to derive actionable insights and automate increasing numbers of tasks to gain a competitive edge, the demand for robust, scalable, and efficient data center infrastructure has surged. This section provides an overview of the intersection between AI and data center operations, setting the stage for subsequent discussions.

The need for data center consolidation still exists, now more than ever

Inefficiencies stemming from disparate data center infrastructures can impede agility, scalability, and cost-effectiveness. Consolidating data centers offers a strategic approach to streamline operations, optimize resource utilization, and enhance overall performance. The drivers for migration and consolidation have largely not changed in the context of AI and range from cost containment to risk management, operational efficiencies, performance management, and more. Many organizations are reviewing infrastructure strategies to optimize their investment in cloud hyper-scalers and on-premises deployments to meet the needs of each workload, including AI training and production workloads.

Migration strategies for AI workloads: Unique, but meet the new boss, same as the old boss

The migration of AI workloads poses unique challenges, necessitating careful planning, execution, and risk mitigation strategies. Whether transitioning to cloud-based architectures or adopting hybrid models, organizations must prioritize factors such as data integrity, latency, and regulatory compliance. The right migration approaches, including lift-and-shift, replatforming, and rearchitecting, must be considered, since many organizations cannot afford downtime while executing these changes. These events are the right time to explore the 7 R's of workload alignment and baseline your application lifecycle position for the entire estate. The challenges for AI workloads are significant, but experienced infrastructure strategists have a toolbox of time-tested methods to bring to bear.

Building future-ready data centers for AI: Workload alignment meets well-architected frameworks

As AI continues to reshape business landscapes, data center design and construction must adapt to accommodate evolving computational and storage demands. From high-performance computing (HPC) clusters to specialized AI hardware accelerators, cutting-edge technologies must be integrated by architects and engineers to support AI workloads efficiently. Are we building in data protection and recoverability to combat disasters and ransomware attacks? Key design principles, infrastructure considerations, and emerging trends shaping the development of AI-ready data centers represent the practical implementation of technology to support the 7 R framework.






Optimization and performance management, where observability meets its purpose

Ensuring optimal performance and resource utilization is paramount in AI-centric data centers. Advanced monitoring tools, predictive analytics, and automation play a crucial role in preemptively identifying bottlenecks, optimizing workload distribution, and mitigating performance degradation. Observability has had its journey through the hype cycle, and strategies for real-time monitoring, workload optimization, and capacity planning to maximize efficiency and ROI become not just desirable, but required, for AI workloads.

Security and compliance considerations, more important than ever before

The proliferation of sensitive data and AI-driven applications underscores the importance of robust security and compliance measures. In addition, the concerns of individuals and governments with respect to how intellectual property and end-user data are used to train AI models must be addressed as a risk management problem. Data encryption, access controls, and threat detection mechanisms are essential safeguards against cyberthreats and regulatory violations. On the risk management front, we need to protect our client positions and ensure our language models do not violate the law or the privacy rights of individuals in the use of the data we train our models against. Security best practices, risk management, and compliance frameworks tailored to the unique challenges posed by AI workloads in data center environments bleed over into business compliance and regulatory concerns.



Conclusion

As organizations navigate the complex landscape of AI-driven innovation, the role of data centers as the backbone of digital transformation cannot be overstated. By embracing consolidation, migration, and forward-looking design principles, enterprises can build agile, resilient, and future-ready infrastructure capable of supporting the relentless demands of AI workloads.

How we can help: Our team is built on the experience of 30-year veterans who construct their own acceleration tools and refine a reliable methodology every day on the largest consolidation and migration engagements. Our baseline and discovery phase will rapidly uncover technical data and organizational knowledge crucial to the overall success of the program. As we reimagine the estate, this knowledge is applied to optimization efforts and application lifecycle planning BEFORE assets and workloads are moved, leading to a much better outcome in the execution and success planning phase.

Where will your ambition take you? Let's talk.

Visit www.tetoncloudconsulting.com

