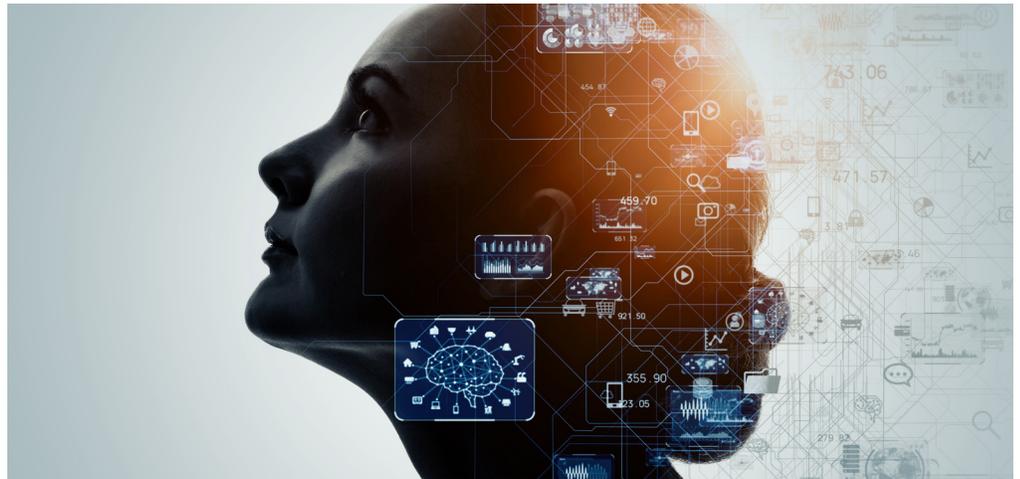




Don't Let Early Choices Hinder AI Project Success



Three Reasons the Right Storage Infrastructure is Key to Rapidly Moving AI POCs to Successful Production.

The last few years have seen a feeding frenzy around how artificial intelligence (AI) and Deep Learning is transforming all manners of businesses. The promise of lower costs, faster decisions, and entirely futuristic methods for achieving new research insights and paths revenue has lead to an intense focus on AI. Some fast movers have certainly achieved major breakthroughs. Yet, the landscape is littered with projects that never got off the ground—good ideas that couldn't make the transition from testing into production.

Organizations are finding the transition from a well-conceived Proof of Concept (POC) into at-scale production is an insurmountable hurdle. Just because a company considers itself data-driven, that doesn't mean it knows how to successfully manage and exploit its data for ultimate success. The more ambitious the project, the more important it is to ensure that data is well managed and available for analysis. There are enough non-technical hurdles to the successful operationalization of AI and Machine Learning; you shouldn't let early platform choices be an additional roadblock. Here are three reasons why choosing the right storage vendor is essential for the POC to production transition.

Reason #1: Starting AI POCs in the Cloud with the wrong vendor can make the transition to data center storage a non-starter.

There are a multitude of reasons to start AI and Machine Learning initiatives in the cloud. Getting access to high performance computing engines based on GPUs without a long-term capital commitment is a pretty good one. At the same time, choosing an architecture that does not mimic the local infrastructure could introduce complexity and unforeseen problems when trying to scale into production.

When choosing a cloud platform and storage provider on which to run I/O, make sure to choose one that can offer the same experience in the cloud and on premises. Poorly planned cloud to data center transitions can result in unpredictable operating expenses and failed data integration projects. Introducing new limitations or management parameters likely, at best, slow the rate at which a transition can be made, just when a competitive advantage could be realized.

Reason #2: Choosing the wrong data center supplier can mean a future of complexity and disappointing results.

Don't choose the infrastructure vendor that can meet your needs today, but plan for success and confirm that their solution can handle the anticipated future dataset requirements. Many storage vendors' "high performance" solutions will appear to have sufficient capabilities when datasets are a few hundred terabytes or a couple of petabytes, or initial analysis runs on only a few GPU nodes. Those same solutions will introduce incredible complexity and an inability to meet processing requirements as capacity and performance needs grow beyond those initial ranges. Point-to-point protocol overhead, file system limitations, and other storage management challenges can also introduce continuous barriers to success.

By ensuring your storage infrastructure vendor can maximize the performance of large-scale CPU and GPU systems and clusters of systems means you don't have to over-provision hardware, minimizing capital expenditures on expensive processing servers. Selecting a storage vendor that offers flexible access methods to AI data in tandem with scalable capacity ensures that management of multiple storage types and systems is kept to a minimum. Parallel client, NAS, and object access methods are all fundamental to AI data accessibility, and a system that supports all three makes management even simpler.

Reason #3: Relying on a vendor without proven expertise can compromise reliability and support on AI projects.

Plenty of infrastructure vendors claim customers with AI deployments, but how many can claim production-level support at the scale that your project will need? Make sure the vendor you select has experience in dealing with the data complexities that are directly relevant to your AI project. How many multi-petabyte implementations can they claim? How many engineers do they have focused on achieving rock solid availability and reliability for at-scale workloads? Can they quickly address new challenges that might emerge from new opportunities realized during an AI project?

DDN has a legacy of addressing complex data challenges.

DDN offers data center and Cloud solutions, including Hybrid Cloud capabilities, that serve as foundational elements to many successful production AI deployments. With over 20 years experience dealing with some of the most difficult data challenges on the planet, DDN is your trusted storage partner in AI and other data-intensive applications.

About DDN

DataDirect Networks (DDN) is the world's leading big data storage supplier to data-intensive, global organizations. DDN has designed, developed, deployed, and optimized systems, software, and solutions that enable enterprises, service providers, research facilities, and government agencies to generate more value and to accelerate time to insight from their data and information, on premise and in the cloud.

©DataDirect Networks. All Rights Reserved. The DDN logo and EXAScaler are trademarks of DataDirect Networks. Other Names and Brands May Be Claimed as the Property of Others.

v1 (7/20)