



THE AI DATA COMPANY

Artificial Intelligence Success Guide

Solving 5 of the
Most Common
AI Infrastructure
Challenges





You've planned your AI strategy to the best of your organization's abilities, but somehow something doesn't feel right.

Maybe your applications seem to be running fine, but since when has "fine" been good enough? Instead of celebrating the new business value of AI, you find yourself resetting expectations.

DOES THIS SOUND LIKE YOUR AI IMPLEMENTATION EXPERIENCE?

Unfortunately, despite its potential, your AI strategy may have yet to deliver on its business goals. Your KPIs could be set, and your ROI calculated. Yet, neither has been accomplished. If this sounds even a little like your implementation, or exactly like what you're trying to avoid, rest assured, you're not alone.

THE QUESTION IS, WHY ARE SO MANY COMPANIES SLOW TO IMPLEMENT AI?

What missteps have others made that you could learn from? In reality, there are several. Some of them involve planning, and others, technology. Still, others touch on company culture. Below are five common challenges companies experience with their AI implementations and how you can meet them head-on.

5

KEY CHALLENGES

For AI Infrastructure

01

YOUR AI PROJECT TAKES TOO LONG TO REACH PRODUCTION

02

YOUR SYSTEMS ARE OVERWHELMED BY THE VOLUME OF DATA

03

YOUR SYSTEMS ARE NOT OPTIMIZED FOR AI

04

YOUR AI SYSTEMS STRUGGLE TO SCALE UP TO PRODUCTION

05

SHADOW AI PROJECTS ARE AN IT HEADACHE

AI Challenge 01

Your AI Project Takes Too Long to Reach Production

When an organization sets out to leverage AI, it's rarely without forethought. Artificial Intelligence is not a strategy a company typically assigns to entry-level IT personnel. Often the organization's best and brightest are tapped to work on the project.

Still, the challenge is that specialized AI workloads are hard to integrate and optimize at scale, deploying an AI workflow on an existing enterprise storage infrastructure seems like a sensible approach, however this is often the first misstep. Problems start to appear as more users take advantage of the system, and applications slow to a crawl. Maybe you add capacity, but jobs continue to run slowly, and then there are intermittent failures, and network, storage, applications issues - the list goes on.

As a result, the timescales used in planning are slipping, and your project has been unintentionally set up to miss its deadlines. Still, the good news is that there are industry experts that can help align your needs and expectations. If you bring the right skills to the challenge, you can set your AI strategy up for success.

AI Challenge 02

Your Systems are Overwhelmed by the Volume of Data

AI needs to be fed with huge volumes of data - video, imaging, language processing - not only to build initial deep learning models, but also to apply those models in production, and evolve through reinforcement learning models and MLOps techniques.

Faced with this surge in data volumes, systems can slow, applications can underproduce, and the return on AI investments can decrease accordingly. However, it's also important to note that solving the problem is not a simple matter of increased throughput or processing speed or more storage.

AI APPLICATIONS AND WORKFLOWS HAVE SPECIFIC REQUIREMENTS WHICH NEED A SPECIALIZED INFRASTRUCTURE THAT IS OPTIMIZED TO ACCOMMODATE AND MAXIMIZE BUSINESS VALUE.

When AI systems are unable to drive sufficient data throughput, organizations sometimes attempt to reduce the amount of data for learning, or reduce the precision and accuracy of their AI models - which compromises the depth of your AI system's insights. It's like expecting a college student to ace an exam without a sufficient amount of studying.

SO HOW CAN YOU BUILD AN INFRASTRUCTURE THAT WILL HANDLE THE NEEDS OF AI WORKLOADS AND EFFICIENTLY MANAGE THE DATA THAT IS REQUIRED TO FEED HUNGRY AI APPLICATIONS?

The answer is to adopt a data-first strategy. Consider the system's data needs from the beginning - in other words, in the design phase of the project. Also, you should address any data privacy, data attribution and intellectual property issues.

Only then, once all of these needs are taken into account, can you design your AI-optimized infrastructure on a reference architecture that accounts for the optimal processing, storage, and network needed and get the most business value from your investment.

AI Challenge 03

Your Systems Are Not Optimized for AI

Although building any computing environment brings unique challenges, AI workloads are particularly challenging, with extreme performance demands.

It's fairly easy to build small systems with very fast data access and low latency, but more difficult to support sustained high-bandwidth data throughput needed by AI systems with massively parallel GPUs.

However, at production scale, AI and Deep Learning architectures need to take it to another level: with very large numbers of small files, and the need to manage vast peta-scale datasets for machine learning, real-time processing and archiving. It's no wonder that regular enterprise storage systems can't handle the data at the scales needed for enterprise-class AI.

Even with the latest in solid-state memory and high-performance networking, regular enterprise storage has to make compromises. AI applications are hungry for data, and enterprise storage infrastructure simply can't feed the data to keep your AI systems working efficiently.

AI Challenge 04

Your AI Systems Struggle to Scale Up to Production

As if implementing AI wasn't tricky enough, even your greatest wins can be bittersweet when you finally switch into production.

All of a sudden, bottlenecks can appear and you're not sure what is causing them.

As a result, the overall system starts to slow, applications aren't performing, inference workloads are not coping, and timescales start to slide. Sadly, these bottlenecks may only become apparent when your AI systems start to take on the stress of a real-world, production-sized workload.

PLANNING FOR SCALE NEEDS END-TO-END SYSTEM DESIGN, AND A STREAMLINED DATA WORKFLOW.

The greater the complexity of your AI environment, the more likely it is to have problems scaling. Think of your AI infrastructure as an office building you want to grow to an indefinite height; if the foundation isn't optimized and strong, its growth will be limited.

To meet the challenge of scaling, focus your design and capacity planning using an AI reference architecture. From day one, scalability has to be a priority in the design of the environment. It's also important to plan for both the system itself and the operations surrounding it, like backup and recovery. Again, when your environment is optimized for AI workloads in terms of current and future needs, everybody wins and you are the hero.

AI Challenge 05

Shadow AI Projects are an IT Headache

Unfortunately, disappointment resulting from any of the above factors (implementation, data management, optimized platforms, and scalability) can motivate other teams within your company to branch off and pursue their own AI strategies.

This split causes multiple AI projects to spring up within the same organization. Whereas an organization may have started with one AI project, all of a sudden, they have several.

While a “do it yourself” mentality may be commendable in other business areas, it can be costly when it comes to AI projects. The organization as a whole can pay the price either literally for investing in redundant AI tools or internally for the extra work hours devoted to implementing and supporting multiple systems. In this scenario, the company loses economies of scale and the benefits of standardization.

JUST AS SHADOW IT CAN BE A DANGER TO THE ENTERPRISE, SO CAN SHADOW AI.

In reality, most organizations need only one AI-optimized infrastructure strategy. A fundamental way to prevent multiple approaches is by establishing a scalable, centralized AI infrastructure or a center of excellence.

If a company’s first AI investment is properly designed and built, other teams won’t feel the need to create their own and the initial setup can be leveraged to meet everyone’s growing needs and plans.

5 STEPS

To AI Success

01

ARTIFICIAL INTELLIGENCE
AND CORPORATE CULTURE

02

THE HIDDEN VALUE OF
BUSINESS DATA

03

AI INNOVATIONS IN DRUG RESEARCH
AND PHARMACEUTICALS

04

AI INNOVATIONS IN HPC
AND SUPERCOMPUTING

05

AI INNOVATIONS IN MEDICAL
RESEARCH AND GENOMICS

Steps to AI Success



AI and Its Effect on Corporate Culture

Another benefit of a centralized AI function is the management of expectations.

With all the excitement in recent years over AI's potential, it helps to have a central, grounded group of experts to differentiate hype from reality.

Your AI Center of Excellence can be a focus for both technology strategy, as well as for business metrics, setting KPIs and measuring outcomes.

AS AI EVOLVES, SO MIGHT THE WAY WE LEVERAGE IT AND, PERHAPS, EVEN THE WAY WE LEAD OUR ORGANIZATIONS.

Steps to AI Success

Discovering the Hidden Value of Your Business Data

Organizations such as Amazon, Google, Apple and Facebook harvest vast amounts of data every day, for targeted advertising and intent marketing right through to natural language processing and facial recognition.


Still, companies in a variety of industries are finding they have plenty of their own data. They may have access to operational data, contextual data, and ambient data from multiple sources, which can be used to augment or enhance their existing business data. Implement AI successfully, and your organization is in a position to drive market disruption, by turning data innovation into competitive advantage.

Steps to AI Success



Another organization who is putting AI infrastructure at the forefront of their objectives is the University of Florida with their AI Supercomputer, HiPerGator.

The University's mission is to arm every one of their graduates with an education in AI and understanding of how AI applies to their discipline.



CLICK ME TO PLAY

Steps to AI Success

ST. JUDE CHILDREN'S RESEARCH HOSPITAL



As more enterprise organizations apply AI innovation, it's nice for those who have forged the path to share lessons learned.

In this video, Wei Guo, Computational Engineer, from St Jude Children's Research Hospital, talks about considerations in building an AI infrastructure and recommendations to others.



CLICK ME TO PLAY

Setting Your AI Project Up for Success

As you can see, for organizations to attain business value from data and get a true ROI from their AI strategy, an investment must be made in planning.

Thus, an AI-optimized infrastructure must be powered with a unique processing and storage architecture that can deliver success.

When considering your AI infrastructure needs, your first logical step is to connect with an expert that can best understand your goals and requirements. Only then can you, together, build a solution that truly fits your organization's needs.

When you partner with DDN, you can start with a turnkey, easy-to-implement, scalable storage architecture that unleashes the power of AI infrastructure. With it, you can drive maximum innovation and gain competitive advantage.

Contact one of our storage experts today to put your AI project on the path to success!

[GET IN TOUCH TODAY!](#)

