

# CIO Essential Guidance

AI for the Enterprise



vmware®

 NVIDIA.



## Table of Contents

Introduction: The Importance of AI to Transform your Business

The Power of Data To Unlock Business Insights

Optimizing Your Strategy for AI to Overcome Obstacles

Avoiding the Downside of Infrastructure Silos with Modern AI Applications

Creating a Unified Data Center

Simplifying Proof of Concepts

Democratizing AI

Bringing AI to Data

Maintaining Privacy and Security with AI

The CIO Takeaway

Solutions

“The companies that will survive are the ones that actually will do this well.”

Paul Turner  
Vice President Product  
Management vSphere  
VMware

---

## Introduction: The Importance of AI to Transform Your Business

Artificial Intelligence is transforming industries today and is at an important inflection point, as adoption transitions from pioneers in the tech sector to mainstream enterprise. From CIOs whose organizations are just getting started in AI to those needing robust infrastructure, everyone is wrestling with the challenge of platform. Whether you are asking, "What do I need if I'm just trying to get a project off the ground," all the way to "How do I build an AI Center of Excellence," every organization needs the right architecture, tools and strategy to enable faster insights from data.

As enterprises integrate AI into their workflows, all aspects of their businesses are becoming infused with AI. Human resources teams are using AI to scan resumes and filter profiles to find workers who might be a good fit for an open role. Marketing and sales groups are using AI to determine pricing and create virtual showrooms and experiences. IT and Information Security use AI to catch cyber security threats and fraud in real time before the threat proliferates. Operations teams use AI for predictive maintenance so they can service machines before they break, and perform demand modeling to forecast and stock supply. Customer service leaders are already successfully deploying chatbots and virtual assistants to shorten wait times. Product development groups are using AI to do digital prototyping and design to quickly and cost-effectively bring products to market.

When used thoughtfully and efficiently, AI solutions are delivering real business and cultural value by harnessing the power of data to develop intelligent solutions across every industry, from advanced diagnostics in healthcare and smart factories for manufacturing to fraud detection in financial services.

“If you have the data that everyone wants, but no computational capability you're still out of luck.”

Tom Hite  
Senior Director, Accelerated  
Co-Innovation Engineering  
VMware


---

## The Power of Data to Unlock Business Insights

Companies are collecting more and more data that has the potential to transform their business. But as organizations accumulate increasingly large data sets, a high-quality problem arises: How do they store, evaluate, understand, communicate, and act on that data in a strategic, cost-effective way?

Tasking data scientists with combing through data is costly, time-consuming and leaves little opportunity for real-time decision-making. This is where leveraging state of the art data science tools powered by AI is essential.

Business leaders across the board are recognizing the benefits of incorporating AI into their day-to-day processes and the impact of AI will only grow in the coming years. In a 2020 IDC survey, 70 percent of CEOs said their organizations needed to be more data-driven, while 87 percent of CXOs said their top priority for the next five years was to be a more intelligent enterprise.<sup>1</sup>



of executives are incorporating Artificial Intelligence and Machine Learning into their 2021 digital transformation projects.<sup>2</sup>

---

1. IDC. "Scaling Artificial Intelligence and Machine Learning Workloads." 2021

2. VMware. "Q2 Executive Pulse." June 2021.

“The first reason [to invest in AI] is that [you need] to better serve your customers. And the second reason is to improve your own internal operations, so that you spend less and do more.”

Manuvir Das  
Head of Enterprise Computing  
NVIDIA


Those who don't take advantage of AI technology may find it increasingly difficult to keep up. VMware Senior Vice President and General Manager, Cloud Infrastructure Business Group Krish Prasad says, "Whether they like it or not, to compete in the market and gain a business advantage they have to utilize the AI techniques as part of their applications to get a leg up on their competition."

Ultimately, AI could be the key to survival. VMware Vice President Product Management, Cloud Platform Business Unit Paul Turner says that no matter the industry or product, every company needs to become a data company: "The heart of IT transformation is data understanding."

### Why Leaders Say They Are Investing in AI\*



The economics of AI: Larger revenue increase or great cost reduction



The dilemma of fast followers: Delay makes us vulnerable to competitors and start-ups



Board approval is quicker with AI tied to solution

\*Gartner. "P-19019 AI in Organisations." 2020.

“Analysts say that up to 50% of AI projects actually fail. They don't see daylight because [they] are not integrated with the mainstream.”

Krish Prasad  
Senior Vice President and  
General Manager, Cloud  
Infrastructure Business Group  
VMware

---

## Optimizing Your AI Strategy to Overcome Obstacles

AI is transforming every industry, whether it's by improving customer relationships in financial services, streamlining manufacturer supply chains, or helping doctors deliver better outcomes for patients. While most organizations know they need to invest in AI to secure their future, they struggle with finding the strategy and platform that can enable success.

While the benefits of using AI systems are clear, implementation can be challenging for a variety of reasons, from startup cost and infrastructure to privacy and silos.

## Avoiding the Downside of Infrastructure Silos with Modern AI Applications

AI-accelerated applications can have intensive resource demands. Many AI models, while in the training process with large data sets, may require a vast number of compute cycles, and AI inference often demands real-time response. Data analytics often involves the transfer of data across multiple systems. Traditional server designs with only CPUs and commodity networking are not able to provide the necessary compute power to run these applications.

These AI applications often have a cloud native architecture, with a large number of microservices organized into complex workflows. This increase in size and scale generates a large amount of network traffic in the data center, which is harder to manage and secure.

Companies often end up deploying these AI applications into one-off, single-purpose clusters, or in the cloud. This results in operational overhead, and often these silos do not adhere to IT standards for visibility, security and governance. This lack of integration can create real barriers to managing data effectively, and can lead to security vulnerabilities.

“When AI workloads are set up this way, both the data and the talent are cut off from the mainstream and there is a higher risk of failure for a multitude of reasons,” Prasad says.



### Challenges and Barriers to AI\*



**53%**

The average number of AI projects that make it from pilot to production.



**30%**

Percentage of respondents who cite complexity of AI solution integration with existing infrastructure as a top barrier.



**30%**

Percentage of respondents who cite security or privacy concerns as a top barrier.

\* Gartner. "P-19019 AI in Organisations." 2020.

“Virtualization is helping healthcare systems to deliver services to clinicians and patients at scale, across radiology departments and facilities. It has the potential to significantly increase the adoption of GPU-based AI applications. This allows for better utilization of technology infrastructure and minimizes the need for dedicated GPU systems for each project, which means AI can be applied more broadly to improve patient services.”

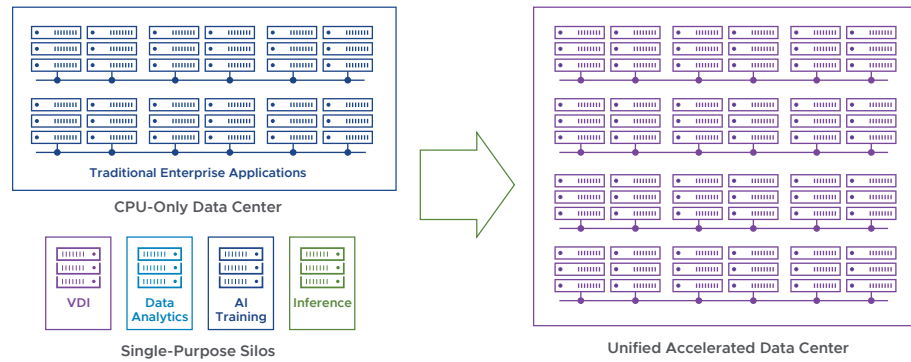
—Tom Schultz  
 Director of Information Systems,  
 Enterprise Medical Imaging and  
 Clinical Data Science  
 Mass General Brigham

## Creating a Unified Data Center

To address the challenges of implementing AI solutions, the data center must evolve to an integrated solution which can streamline both the development and deployment of AI workloads. This solution should include all of the AI tools and frameworks required to simplify the process, and it should include management and orchestration tools to ease the burden on the IT department.

### Path to a Simplified Accelerated Data Center

Prepare for the future while driving down data center costs.



Ideally, the solution would run on accelerated computing platforms, eliminating the need for expensive specialty hardware purchases and the creation of separate infrastructure silos for running AI.





## Simplifying Proof of Concepts

AI infrastructure is often perceived to require a large investment in terms of both time and finances as well as highly specialized IT staff. This may be accurate for companies that are building their own AI platform. However, for many enterprises, it's possible to take a more "off the shelf" approach and leverage pre-tested and integrated solutions.

Ideally teams should build a production environment that creates the essential platform for the evolution of enterprise data centers, delivering infrastructure that can handle a diverse range of accelerated workloads. With systems for accelerated computing from leading partners, organizations can use mainstream servers to cost-effectively perform proof of concepts (POCs) while providing performance, manageability, scalability and security. The most successfully scalable AI solutions will be those that can be folded into existing enterprise infrastructure.

Getting started with AI is low risk with re-integrated solutions. Organizations can get started with minimal initial investment, and purchase a server that can be used for AI and other enterprise workloads like ERP.

## Democratizing AI

While there are certainly barriers organizations must overcome in order to implement AI, the good news is that new technologies are bringing down the cost as well as speeding up the process, giving more organizations the power to harness their own data.

When we think of models, it's important to realize that they don't always have to be created from scratch. There are models available that teams can implement and modify for their own organization's needs. That front-end work has already been done.

VMware Chief Technologist for High Performance Computing Josh Simons says, "For example, Google has built huge models based on vast amounts of data for doing language models. And what you can do is you can take these, and you can customize them by taking some amount of your own data and applying it, continuing the training."

According to Das, companies will be able to level the playing field with their competitors, "Your largest competitor had the money to go do all this on their own, but what we've now done is, we've packaged it all up so that you can just adopt it."



Pre-trained models help break down the barriers of adopting AI. VMware CTO Kit Colbert says pre-trained models “enable not only specially trained data scientists, but also a broader set of users in your organization.”

Leveraging out-of-the-box AI-infused software makes it easy and efficient to utilize your current talent to manage and leverage your AI processes.

In an ideal world, companies will run on software that's already been infused with AI. Das says, "You're just adopting the next version of software you already have, and it happens to be better because it's infused with AI. That's the true democratization, because that's the way that the maximum number of companies can adopt it without having to learn about AI."



## Bringing AI to Data

Rather than creating costly silos solely for the purpose of storing data and running an AI workload on an enormous data set, it's beneficial to integrate AI into your standard environments and practices.

Using an infrastructure of virtual machines that can scale up or down depending on need makes the best use of your distributed server, cloud and/or edge availability. Resources can be pooled and aggregated based on AI processing demands. It also allows you to bring AI to the data, rather than the other way around.

Optimizing for performance to run AI workloads at near bare-metal speeds enables data scientists to efficiently access their valuable data, no matter how large and complex their training and inferencing models are.



## AI Myths

**Myth:** AI Eliminates the Need for Data Scientists

**Reality:** Data Scientists are the key to training and monitoring any AI system.

**Myth:** AI Workloads Require A Bare-Metal Server

**Reality:** According to experts, AI systems can run efficiently in a virtualized environment.

## Maintaining Privacy and Security with AI

Companies that lack proper privacy and security policies will be challenged to implement a responsible AI strategy. Whether using customer data, operational data, or product performance data, a company needs to exercise caution and protect the privacy of their valued customers and their company's intellectual property.

A best practice, for example, is to anonymize customer data when it is collected, cleansing it of any identifying information. AI processes can also be useful in carrying out and ensuring customer privacy. Of course, keeping data secure and protected against cyber threats is also critical. In short, the team tasked with implementing AI should include experts from the company's IT security and privacy team, and their collaboration will be critical to the success of AI across the company.

## The CIO Takeaway

While AI systems have historically been expensive, time-consuming to launch and difficult to manage, new technology is making it less expensive, faster and easier to harness the power of your data.

Integrating your AI processes with your mainstream workloads across virtualized systems helps to better allocate resources, access all of your data faster, and better manage your AI applications.

Companies should not hesitate to implement AI processes that could transform their business. Das stresses, "You need to care about AI, because going forward the applications that you will need to use in your business will be built on AI."

CIOs should act now or risk their organization being left behind.



“NVIDIA AI Enterprise allowed us to expand our support for our researchers and students who utilize data analytics and AI deep learning and machine learning, while making these applications easier to deploy and manage. Our testing has shown that these latest collaborations between NVIDIA and VMware deliver the full potential of our GPU-accelerated virtualized infrastructure at near bare-metal speeds.”

—Maurizio Davini  
Chief Technology Officer  
University of Pisa

---

## Solutions:

VMware and NVIDIA have partnered to unlock the power of AI for every business by delivering an end-to-end enterprise platform optimized for AI workloads. This integrated platform delivers best-in-class AI software, the NVIDIA AI Enterprise suite, optimized and exclusively certified for the industry’s leading cloud platform, VMware vSphere®.

Running on mainstream NVIDIA-Certified Systems™, available from industry-leading server vendors, this platform accelerates the speed at which developers can build AI and high-performance data analytics, enables organizations to scale modern workloads on the same VMware vSphere infrastructure they have already invested in, and delivers enterprise-class manageability, security and availability. Furthermore, with VMware vSphere with Tanzu, enterprises can run containers side by side with their existing VMs.

Enterprises can confidently deploy an end-to-end AI solution by using AI solutions optimized and exclusively certified by NVIDIA to run on VMware vSphere. They can also prevent data center infrastructure silos and simplify management by leveraging virtualization to incorporate AI deployments into existing enterprise infrastructure with optimized software – accelerating AI adoption in the enterprise.

At a fireside chat, NVIDIA CTO Michael Kagan, said, “NVIDIA and VMware are providing a turnkey AI solution, allowing CIOs to deliver an AI platform across their organization.”



To learn more, visit:

[www.nvidia.com/en-us/data-center/vmware/](http://www.nvidia.com/en-us/data-center/vmware/)



Copyright © 2022 VMware, Inc. All rights reserved. VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001  
VMware and the VMware logo are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware products are covered by one or more patents listed at [vmware.com/go/patents](http://vmware.com/go/patents).  
Item No: 1080707aq-bro-cio-esntl-gdnc-ai-ent-uslet 1/22