

AI IN THE MODERN ENTERPRISE

THE AI-POWERED HYBRID CLOUD STARTS WITH INFRASTRUCTURE

SUMMARY

We're in perhaps the most dynamic era of enterprise IT. Modernization initiatives have been rescoped and accelerated to support generative AI projects, which have captured the attention of every executive with good reason. Gen AI promises to transform businesses in ways we haven't witnessed.

However, as the need to accelerate and alter modernization efforts to support this new wave increases, IT budgets are only rising incrementally at best. Sustainability is another variable in the equation. While AI initiatives require more compute, storage, and other resources, CIOs are tasked with lowering power footprints to drive sustainability goals.

How can enterprise IT organizations simultaneously achieve modernization, AI, and sustainability goals, which seem to directly contradict one another? Moor Insights & Strategy (MI&S) sees the solution as rooted in infrastructure.

Outdated operating stacks powered by outdated hardware and processors unable to deliver the required performance, agility, security, and targeted acceleration are, in some cases, used as the building blocks for the AI-driven workloads running the modern business. This is a recipe for failure.

This research brief will explore enterprise IT organizations' technical and operational challenges and how technology vendors are responding with hybrid cloud environments powered by modern AI-ready infrastructure. Further, it will evaluate how Nutanix, Dell, and Intel have partnered to deliver the Dell XC Plus running the Nutanix Cloud Platform (NCP) and GPT-in-a-Box powered by AI-accelerated Intel Xeon CPUs.

MODERNIZATION AND AI — COMPLEMENTARY YET COMPETING

For purposes of this brief, IT modernization encompasses the infrastructure and the operating and application environments necessary to support the workloads that power today's business. It also considers the operational and organizational restructuring that enables these technology updates to maximize the value delivered to the business. Similarly, AI explicitly refers to gen AI and the emergence of agentic AI, whereby

business functions and processes achieve efficiency through embedded AI agents that operate autonomously.

Data is the foundation of AI and the analytics tools that empower organizations, but it's the underlying architectures and infrastructure that unlock the potential of that data.

Success in the data-driven world is influenced by speed and agility. "Time to value" essentially measures how quickly and ably an organization can recognize and respond to market needs. "Time to first token" is a newer term that reflects how quickly organizations can go from concept to operational in AI projects. While often referenced in different contexts, the two terms are inextricably related, as gen AI can be an accelerant that enables organizations to serve the market more effectively.

It makes sense that some IT organizations view modernization and AI projects as competitors for budget, resources, and prioritization. In some cases, AI is being introduced after many modernization projects are already in flight. Additionally, there's pressure on IT organizations from stakeholders who don't necessarily understand the connective tissue between the two. Because of this dynamic, it's critical for IT leaders to educate stakeholders and solicit their support for this integrated approach.

CAN IT MODERNIZATION AND AI OPERATIONALIZATION OCCUR SIMULTANEOUSLY?

While this notion of a wholly integrated project is theoretically the ideal approach, prioritization becomes a question. Is it better to deploy new infrastructure and a cloud operating model to alleviate existing IT stressors? Or should this be deprioritized in favor of stakeholder demands to focus on deploying AI? Can both projects be managed at the same time?

The realities of budget and resource allocation often prevent IT modernization and the operationalization of AI simultaneously. Simply put, there isn't enough money or people to achieve both goals simultaneously using traditional infrastructure deployment methods. Furthermore, the employed IT professionals, while capable, simply haven't been trained to stand up and support these new environments.

WHERE DO WE GET ENOUGH POWER?

The electricity consumed by the platforms and accelerators running AI exceeds many datacenters' space and power budgets. While lowering power footprints is ideal for environmental purposes, it appears to be in direct conflict with operationalizing AI in the enterprise. This final element of the challenges outlined previously hits on a difficult question for organizations. Considering the common mandates to drive down power

consumption, how can an IT organization embark on modernization and AI projects simultaneously?

MI&S regularly speaks with IT leaders who say that these budget, staffing, and power challenges make this balance impossible to achieve. However, there is a solution to address these issues: bringing the cloud on-premises. This starts with Dell's XC Plus.

THE OPTIMAL AI FOUNDATION BEGINS WITH THE CLOUD

While many enterprise IT organizations encounter significantly higher costs in the cloud than originally anticipated, adoption continues to accelerate. Elasticity and simplicity are two big reasons for this. As resource needs change, the cloud can scale up or down. Provisioning resources is as simple as navigating a console. This is true for virtually any workload, including AI.

With digital transformation and IT modernization efforts taking over in recent years, the cloud operating model became the target end state for IT organizations struggling to meet business needs. For many organizations, that goal hasn't been realized for a number of reasons, one of which is the AI rush.

In Dell XC Plus, powered by Intel Xeon processors, MI&S sees a platform ideally suited to meet the needs of today's enterprise, including traditional virtualized, cloud-native, and even gen AI workloads. Dell XC Plus is plug-and-play cloud infrastructure combining Dell PowerEdge servers with the Nutanix Cloud Platform. This co-engineered solution delivers a hybrid cloud environment with a single management plane for compute, storage, and networking.


NUTANIX CLOUD PLATFORM — SIMPLICITY THROUGH ABSTRACTION

NCP is the operating stack that delivers the cloud operating model and hybrid cloud connectivity. NCP comprises several modules that bring this functionality to bear:

- **Nutanix Cloud Infrastructure (NCI)** is the heart of NCP. It's the underlying infrastructure that unifies storage, compute, and networking into a single logical pool, enabling fast and automated provisioning, scaling, and management of critical resources.
- At the heart of NCI is AHV, or the **Acropolis Hypervisor**. AHV, the Acropolis operating system (AOS), and the Distributed Storage Fabric (DSF) work in concert to deliver the essence of the hybrid cloud platform: virtualized and containerized environments that can reside on-prem or in the cloud.

- **Nutanix Central** delivers control, governance, and consistency across the entire data estate: on-prem, in the cloud, and multi-cloud environments.
- A key element of Nutanix is **Nutanix Cloud Manager** (NCM), which is where MI&S sees Nutanix delivering particular value and market differentiation. NCM's Intelligent Operations module enables ITOps, AIOps, and FinOps through several machine learning algorithms. This use of multiple models is essential as it accounts for various operating environments.
- One last element of the Nutanix Cloud Platform worth exploring is the **Nutanix Security Central** (NSC) module. NSC isn't a security platform but an aggregation point whereby IT managers can look across their public-private environments to better understand their security posture and quickly identify vulnerabilities. Further, NSC can automatically respond to incidents to mitigate the impact of cyber incidents.

FIGURE 1: NUTANIX CLOUD PLATFORM AND GPT-IN-A-BOX



Nutanix Cloud Platform supports virtualized, containerized, and enterprise AI applications.

Source: Moor Insights & Strategy

While NCP delivers the hybrid cloud platform and operating model to enterprises, GPT-in-a-Box is how those enterprises can operationalize AI quickly. This full-stack AI platform exposes the full capabilities of AI-accelerated Intel Xeon CPUs with Dell as well as an opinionated AI stack by working with Dell, Intel, and AI software ecosystem partners such as Kubeflow, PyTorch, and TouchServe. GPT-in-a-Box also supports deploying many popular large language models. In addition to out-of-the-box support for

Llama 2, MosaicML, and Falcon GPT, the platform provides integrated access to the Hugging Face library of models.

Nutanix has evolved from an HCI startup to the foundation of the hybrid cloud (and hybrid AI) environment. During this evolution, the company has smartly adhered to its core tenets of abstracting complexity and embracing openness, two critical requirements for the enterprise IT organization tasked with simultaneously driving modernization and AI efforts.

DELL XC PLUS — PERFORMANCE AND SECURITY

Like any operating environment, Nutanix's richness and power depend strongly on its underlying infrastructure, which must be highly performant and finely tuned for the environments and workloads that are running. Dell XC Plus, a family of server platforms tuned to support a variety of deployment and usage scenarios, is exactly this. These servers include:

- **XC660xs** — This server is tailored for smaller operating environments, such as remote office/branch office (ROBO), where the basic HCI is run to quickly deliver IT services to a location or office.
- **XC660** — This dual-socket 1U server combines compute, memory, and I/O for organizations requiring more horsepower to support enterprise applications such as VDI and data analytics.
- **XC4000** — The edge often requires unique form factors due to operating environment constraints around space, power, and the like. However, the workloads supported are often performance- and latency-sensitive. Point-of-sale, manufacturing, and defense are good examples of environments that require hardened, smaller server form factors. The XC4000, powered by Nutanix, is designed specifically for these edge environments. It's a highly performant server yet physically designed to withstand non-datacenter environments. Further, NCP and its integration with Dell's OpenManage platform allow for the remote deployment, provisioning, and lifecycle management of these servers.
- **XC760** — Database and data-intensive workloads have high storage requirements. Servers for these workloads require fast processors, as much high-performance storage as possible, and potentially acceleration. The XC760 is a storage-optimized platform designed specifically for these environments by supporting up to 368TB of highly performant flash storage with Intel Xeon CPUs and support for accelerators.

- **XC760xa** — HPC and AI are two workloads that require not only high performance and maximum capacity but maximum acceleration. When thinking about AI and training models, acceleration is critical for maximizing the time-to-value metric. The XC760xa is an Intel, Dell, and Nutanix platform that exposes each partner's highest performance by delivering acceleration that begins in the silicon and extends through the entire compute stack.

FIGURE 2: DELL XC PLUS PORTFOLIO WITH INTEL XEON

Dell XC660xs 1S or 2S 1U	ROBO	Light	Simple HCI
Dell XC660 2S 1U	Enterprise Datacenter	Balanced	Virtual Apps & VDI Test/Dev Private Cloud
Dell XC4000 1S 2U	Rugged Edge	Latency-Sensitive	Industrial Automation Telecommunication Military
Dell XC760 2S 2U	Enterprise Datacenter	Storage-Heavy	High Performance AI Database
Dell XC760xa 2S 2U	Enterprise Datacenter	GPU Density	AI & ML

The Dell XC Plus portfolio supports the broad range of deployment and computational needs.

Source: Moor Insights & Strategy

The XC Plus is built on the Dell PowerEdge server platform, which is highly secure and tightly managed. Intel Security Engines, combined with Dell's platform silicon root of trust and iDRAC, ensure a verified chain of trust across the platform stack that protects across the seven pillars of the [National Institute of Standards and Technology's \(NIST\) Zero Trust model](#). This secure operating environment is easily monitored and managed through Nutanix Security Central.

From a systems management perspective, Dell and Nutanix have partnered to tightly integrate iDRAC with Nutanix Lifecycle Manager (LCM). Through this integration, LCM can deliver updates that once happened in an out-of-band fashion (often requiring reboots in-band). This means commonly performed functions such as driver updates can be performed without forcing servers to reboot.

The tight integration of Dell PowerEdge with Intel Xeon CPUs and the Nutanix Cloud Platform can deliver a highly performant, significantly self-managed, and self-secured cloud environment. MI&S recommends that enterprise organizations grappling with the tensions of modernization projects competing with AI projects take a hard look at XC Plus.


INTEL XEON — MODERNIZATION STARTS IN SILICON

Just as operating environments and workloads depend highly on infrastructure, infrastructure relies on compute platforms that power these environments. In other words, infrastructure starts in silicon. The Dell XC Plus with Nutanix Cloud Platform is highly performant and highly secure because of Intel's Xeon processor. Intel's 5th Gen Xeon was designed for general and targeted performance (through acceleration engines), scale, and efficiency.

Intel claims that the [5th Gen Xeon processor family](#) can deliver up to 84% better general-purpose compute performance and an astounding 10 times gain in performance-per-watt efficiency compared to the current infrastructure in many datacenters. Perhaps the most compelling design element of the 5th Gen Xeon is the inclusion of acceleration engines, discrete silicon built into the SoC that offloads specific tasks from the CPU to enable greater workload performance and higher efficiency.

The x86 architecture is built on the principles of the Von Neumann architecture. While incredibly efficient, this architecture doesn't account for the many functions that a modern CPU must perform. Because of this, x86 CPUs perform some functions less efficiently (and performantly). This is where Intel's smart use of acceleration engines delivers performance improvements and efficiency relative to the market.

FIGURE 3: INTEL XEON ACCELERATION ENGINES



*Intel's acceleration engines offload computationally intensive functions to discrete silicon.
Source: Moor Insights & Strategy*

Intel's Advanced Matrix Extensions (Intel AMX) should be of particular interest to enterprise IT. AMX works as an acceleration engine to deliver better performance for AI workloads such as natural language processing, recommendation systems, and image recognition, three AI implementations that begin the AI journey for many enterprises.

Additionally, Intel has long been at the forefront of building instructions into its core CPU to enable greater support for complex tasks. One of those, AVX-512 (Advanced Vector Extensions, 512 bits), offloads and handles vector operations much more quickly. These operations are critical to AI and analytics workloads.

Finally, the ability for enterprises to more effectively achieve sustainability goals by using 5th Gen Xeon bears repeating. Simply put, IT organizations can dramatically reduce their server footprints (hence energy footprints) through the higher levels of consolidation and higher performance of workloads delivered by these processors.

MANAGING THE MODERNIZATION-PLUS-AI JOURNEY

Driving a modernization strategy that includes operationalizing AI across the enterprise is a complex initiative that could easily fall into the "impossible" category. However, with a technology stack like the Dell XC Plus, which includes the Nutanix Cloud Platform (NCP) and supports integrated GPT-in-a-Box, this effort is not only feasible but can deliver incredible results. However, as with any pan-organization

IT project, intentionality is crucial when executing this strategy. MI&S sees the following considerations as critical to any enterprise:

- **Think big, plan small** — Consider the ideal end state of a global modernization effort that includes AI and agentic AI. Break this end state into smaller, manageable workstreams. When reaching phase gates, be sure to assess and apply learning progressively. MI&S believes a considerable amount of iteration will drive better outcomes.
- **Design from the bottom up** — This research paper underscores that infrastructure is critical to both modernization and AI. Without the right building blocks, the operating environment will be suboptimal at best. Therefore, it's critical to choose the right silicon, server, and operating stack for your environment. MI&S sees the combination of Intel, Dell, and Nutanix as a compelling foundation.
- **To accelerate or not accelerate** — The buzz around GPUs is part of the AI hype cycle. This buzz is often but not always warranted. As part of the planning process, understand where and when to deploy GPUs for acceleration. In some cases, a CPU such as Xeon with AMX will suffice. Other cases may require a targeted accelerator like Intel's Gaudi 3. Upfront planning can significantly impact power and resource allocation requirements.
- **Consider your data strategy** — Organizations have a lot of data. The older the organization, the more data. While data is gold in this AI era, some data has no bearing or value for training and inference purposes. Just like in database normalization, an organization can go too far in which data is deemed relevant and used for training. This can lead to accuracy issues and higher costs.
- **Mind the skills gap** — IT organizations are filled with bright, curious professionals. However, being bright and curious doesn't mean being able to support an initiative. This organizational challenge can be critical. MI&S strongly suggests that IT leaders do an upfront assessment that aligns the organization's skills with modernization and AI projects. Account for gaps and workloads. Consider where training can bridge gaps and where trusted partners are needed.
- **Find a trusted partner** — MI&S strongly recommends that IT organizations work with channel partners that have undergone the modernization and operationalization of AI and established best practices. Although this process has a cost, MI&S believes the cost will be worth it to avoid the unforeseen expenses that accompany the inefficiencies of attempting to go it alone.
- **Process is important** — Procedures and processes are critical to modernization fused with AI deployment. This includes everything from data governance to

where agentic AI can and should be used. Spend the time and effort to consider these things carefully and find an outside set of eyes to validate or challenge organizational positions.

CALL TO ACTION

The hype around AI is real, and a modernized environment is critical to leveraging the power of gen AI across the enterprise. Many enterprise organizations are woefully underprepared for this AI journey, partly because of outdated infrastructure and operating environments and partly because of organizational challenges. Additionally, the mandates around budget and sustainability seemingly contradict this notion of IT and business transformation, where more processing power is the answer.

MI&S regards a well-considered strategy and plan as the key first step in the AI-modernization journey. This strategy should include modernizing infrastructure that starts at the silicon and moves to the operating stack.

The Dell XC Plus, powered by Intel Xeon, is well-positioned to deliver the performance, security, agility, and cost savings that enable enterprise IT to meet current and future needs. Taking maximum advantage of Nutanix Cloud Platform, this combination can be the answer to that first key step for many enterprise organizations.

Because of the combined capacities outlined in this research brief, MI&S suggests that enterprises beginning their AI-modernization journeys consider the Dell XC Plus powered by Intel Xeon and Nutanix as their fundamental building block.

For more information, visit: nutanix.com/dell

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR(S)

[Matt Kimball](#), Vice President and Principal Analyst, Compute and Storage

PUBLISHER

[Patrick Moorhead](#), CEO, Founder, and Chief Analyst at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy." Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

Nutanix commissioned this paper. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

© 2025 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.