

ACCELERATING INNOVATION BEGINS WITH DATA

AND THE ENTERPRISE DATA ESTATE SPANS THE DATACENTER, CLOUD AND EDGE

SITUATION ANALYSIS

While many say we live in the era of the data-driven enterprise, the reality is that we've always been driven by data. Three factors make today unique though: how much data we are generating, where we are generating this data, and what we can do with it. Where responsiveness to the needs of the market and customers is critical, organizations that can best utilize data are poised to succeed.

Undoubtedly, data is the fuel that powers the enterprise. However, utilizing data-gleaned insights and intelligence is increasingly difficult in a world where the amount of data generated and collected has exploded.

While many organizations understand the importance of - and challenges posed by - the data management lifecycle, the required resources to succeed are insufficient, from both a human and technology perspective. As a result, the all-important time-to-value (TTV) metric that drives the business is not being met.

At the heart of this challenge is a tension that exists between business users who want every analytics service and capability at the speed of the cloud and an IT organization that is increasingly operating beyond its full capacity.

This research brief will dig deeper into these challenges and explore how implementing a hybrid multicloud environment built on Nutanix and Intel can help deliver the performance that enables the business and the centralized control that IT requires.

THE CHANGING ENTERPRISE APP AND DATA LANDSCAPE

The enterprise technology landscape has dramatically changed in recent years. While digital transformation has been an evergreen process over decades, this latest wave, driven by data, data analytics, and artificial intelligence (AI), has led to perhaps the most profound technology inflection point.

An enterprise IT landscape that a few short years ago consisted of virtualized and bare metal servers supporting traditional line-of-business applications and functions is gone.

The conventional datacenter has given way to data estates that span the cloud(s), the edge(s), and wherever else data is created and housed.

The era of the three-tiered application architecture developed, deployed, and maintained by a centralized IT staff is also ending. Cloud-native, distributed applications designed in newer methodologies such as SecDevOps by architects and developers embedded in the business units are being deployed in their place. Further, while some applications drive the business longer term, others are being created and used with a fixed duration and limited scope to support campaigns and other business unit initiatives.


Seemingly, every application populating the data estate is a snowflake with compute, security, and latency requirements that are unique and can't be compromised—all while running across multiple clouds, the edge, and traditional data environments. Further, much of the data generated serves two purposes. First, some data – log files, performance data, and the like – is monitored and analyzed in real time for security, performance, and other purposes. For example, consider chatbots that must parse conversation data to understand a customer's needs and sentiment or security log files that are constantly scanned for indications of breaches or other malicious behavior. In these scenarios, faster-performing servers with more memory bandwidth that can feed a CPU or accelerator will drive the best performance.

On the other end of the spectrum, data is collected across multiple endpoints, integrated with other data from various sources across an enterprise, and housed in a repository where deeper analysis takes place by a team of data scientists. This could be climate scientists measuring weather patterns around the globe, epidemiologists tracking a pandemic or seasonal flu outbreak, or financial analysts studying historical data to predict a stock's performance. More powerful servers that can pack more computational resources (CPU cores), memory capacity, and storage would be ideal for these deployments.

It is important to note that while application modernization and the use of analytics and AI to drive transformation and modernization projects are prevalent, Moor Insights & Strategy (MI&S) sees many of the traditional large-scale database deployments as continuing to drive modern business. Oracle, SAP, PostgreSQL, and many traditional database players used in data warehousing efforts are vital to the digitally transformed environment. The data from these platforms feed the foundational and large language models (LLMs) that power organization-specific generative AI (GAI) implementations.

And the data from these warehouses deliver the historical context that can aid in strategic planning.

FIGURE 1: DATA LIVES EVERYWHERE IN THE MODERN ENTERPRISE



*Data is generated, used, and stored everywhere in today's data estate
Source: Moor Insights & Strategy*

The above examples show the value of data. More specifically, they show how that data helps organizations of all types and sizes drive better outcomes. Data feeds applications that enable organizations to plan, act, and respond faster and with greater precision.

Finally, the above demonstrates the need for an environment where data can be easily consumed by applications and workloads as fast as possible and with absolute security.

THE MODERN DATA ESTATE REQUIRES MODERN PLATFORMS

As detailed, the workloads and data powering the modern enterprise differ in architecture and utility across the organization. Cloud-native application architectures benefit from scale-out compute infrastructure, whereas data-centric workloads such as data warehouses and ERP systems can benefit from more of a scale-up deployment. While MI&S sees single-socket platforms gaining popularity across the enterprise, data warehouses and large-scale database environments still consume two-socket, four-socket, and even eight-socket platforms with maximum memory to enable performant database operations.

However, what is consistent across these different classes of workloads is a requirement for modern platforms designed from the silicon to the application stack with performance, mobility, and security as priority vectors. This is a critical consideration for application and infrastructure architects. The requirements of modernized applications require modernized platforms. From cores to cache to memory and I/O, these applications have very low latency requirements not served well by servers that don't support DDR-5 RAM or PCIe 5.0, for example.

IS MODERNIZATION TENABLE?

Any experienced IT professional reading the above should note two issues that would likely be considered showstoppers – cost and complexity. From a cost perspective, modernizing the datacenter requires the potential replacement of infrastructure with bespoke platforms supporting bespoke workloads, not simply in terms of scale-out versus scale-up but also in terms of performance requirements around memory, I/O, acceleration, and the like.

From a complexity perspective, supporting such an environment traditionally requires adding more resources. This includes more IT professionals with perhaps specialized skills to deploy, provision, and manage the infrastructure and application environment that powers this digitally transformed, modernized business.

While these concerns are certainly valid, the solution to both challenges may already be deployed in the enterprise datacenter – hyperconverged infrastructure (HCI).

HCI IS THE FOUNDATION OF THE MODERN PLATFORM

HCI, a technology pioneered by Nutanix, is a model designed to bring the concept of hyperscale-level deployment and provisioning to the masses. This means rack it, stack

it, and get a server up and running with a few clicks. The technology quickly found favor with IT organizations wanting to deploy infrastructure at scale for some of the more traditional workloads in the datacenter. However, in its early releases, HCI had latency challenges that prevented it from supporting some of the most demanding business-critical workloads.

During the COVID pandemic, the HCI market accelerated due to supporting a remote workforce, with workloads like virtual desktop infrastructure (VDI) driving much of that expansion.

Interestingly, COVID was an opportunity for IT organizations to rediscover the benefits of HCI. First was the simplified control and cloud experience that HCI can deliver for IT administrators and staff. With HCI, deploying a fully functioning VDI environment goes from days or perhaps even weeks to minutes.

The second benefit was how capable some HCI distributions had become at supporting performance and latency-sensitive workloads. While the HCI solution stack of some vendors matured since its introduction, many optimizations were made around the scalability of resources and performance, supported by advancements in silicon and associated technologies from companies such as Intel. As many IT organizations struggled to deploy development, database, and application environments during COVID, some HCI platforms delivered the simplicity and performance necessary as transformation and modernization efforts regained momentum.

MI&S can anecdotally validate this dynamic through IT executive engagements. In particular, MI&S has found that organizations undertaking transformation projects that span the cloud, the edge, and the traditional datacenter have found Nutanix to deliver the tools, control, and performance that ease the IT/business tension around responsiveness, agility, performance, security, and cost.

NUTANIX – WHERE BUSINESS CRITICAL AND CLOUD INTERSECT

HCI from Nutanix has evolved from simplifying the deployment of IT infrastructure to enabling the full cloud experience. A critical element of this is enabling the support of any workload anywhere – meaning on-prem, in the cloud, on the edge, or even all these environments simultaneously. The distributed nature of cloud-native makes this "run anywhere and run everywhere" scenario not just possible but plausible.

Technologies such as Intel Virtualization Technology, which is still enhanced with each generation, make it easier to place workloads and run them anywhere across different clouds and different generations of hardware. So, delivering a real cloud experience means enabling this deployment without compromising performance, latency, and security.

As many IT organizations look to “cloudify” operations and deliver the self-service, billable consumption model to the business, MI&S sees HCI from Nutanix as a compelling enabling foundational cornerstone. Further, the architecture of the Nutanix Cloud Platform delivers cloud compute, storage, and networking services to the enterprise with minimal latency. This means developers embedded in a business unit can quickly deploy infrastructure to support their specific needs with point-and-click simplicity without sacrificing performance, consistency of performance, or security.

This also means that IT organizations that have historically left business-critical applications outside of the HCI environment due to application latency issues can simplify management and control without sacrificing performance. Business users and customers see no measurable difference in application responsiveness, while IT significantly reduces the time (and cost) associated with managing these applications.

The Nutanix Cloud Platform is powered by Nutanix Cloud Infrastructure (NCI). NCI is the underlying HCI platform that delivers on the promise of performance, security, and availability and bridges an on-prem to off-prem environment.

Working with NCI to deliver the full cloud operating model is Nutanix Cloud Manager (NCM). Enterprise IT organizations hoping to achieve the automation and cost savings associated with the cloud operating model require automation through AI and intelligent operations (AIOps), self-service capabilities for users, security, regulatory compliance, and cost control for hybrid multi-cloud environments. And this is what NCM delivers.

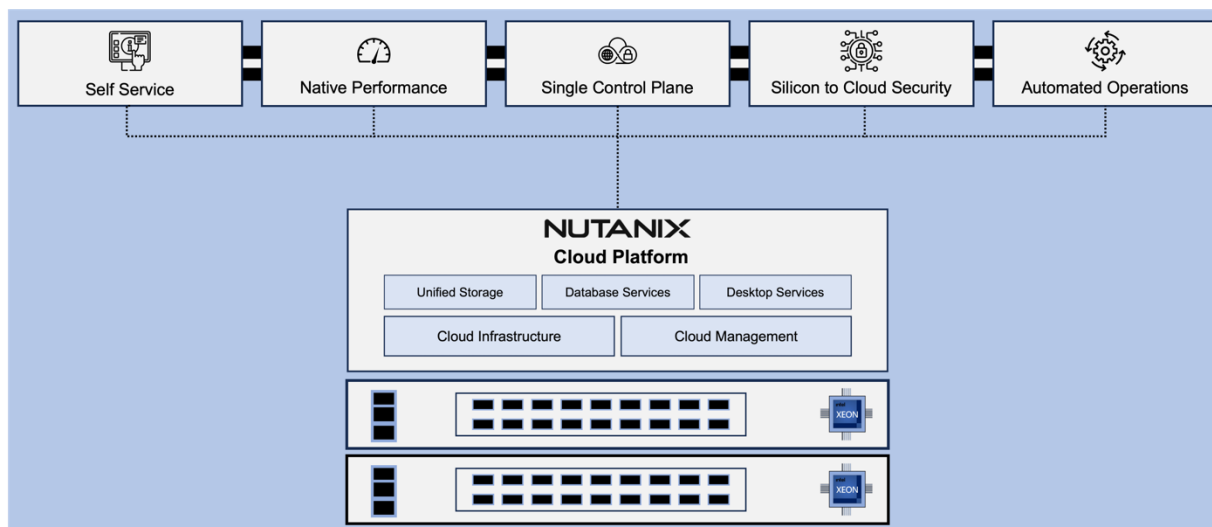
While MI&S finds a number of NCM capabilities compelling, its ability to analyze cost spend across the data estate and suggest optimizations based on cloud usage is especially attractive. One of the constant challenges MI&S hears from IT and finance executives is the challenge of controlling cloud spend. NCM can enable this governance.

The Nutanix Database Service is another key element of the Nutanix Cloud Platform and vital to enabling business-critical applications. Through the Nutanix Database Service, data base administrators (DBAs) and database professionals can automate

many operational and security tasks that burden IT organizations. This service enables consistency of database operations and delivers that "few clicks to support" for application developers. For a deeper analysis of Nutanix Database Service visit the MI&S research brief [here](#).

Finally, MI&S sees Nutanix Unified Storage as a critical enabler of the overall cloud experience. Storage management is costly, complex, and highly siloed. File, object, and block storage traditionally require three environments, each complete with separate consoles for administering, protecting, and managing. With Unified Storage, Nutanix removes these complexities from IT and, as a result, delivers incredible efficiency gains and cost savings.

FIGURE 2: THE CLOUD OPERATING MODEL DELIVERED BY NUTANIX AND INTEL



The Nutanix Cloud Platform delivers control for IT and agility for the business

Source: Moor Insights & Strategy

When deployed, these technologies are the beginning of the modern platform that supports modern business. But, as with any software, Nutanix is limited or empowered by the underlying hardware architecture, which begins with silicon.

4TH GENERATION INTEL XEON SCALABLE PROCESSORS – WHERE INNOVATION AND ACCELERATION BEGIN

When Intel introduced its 4th Generation Xeon Scalable Processor (codenamed Sapphire Rapids) in 2023, it significantly focused on enabling and accelerating real-world application performance. These modern workloads with unique performance

requirements need more than just many cores. Architectural optimizations and silicon-based acceleration are needed for a business to hit that vital time-to-value metric. Whether that value measures the responsiveness to a customer inquiry, the time it takes to transact a stock trade, or the development lifecycle of a breakthrough medicine, TTV is a real metric with real-world impact.

In each of the above scenarios, more cores can help – but only so much. The modern workload requires marrying the right balance of cores with memory, I/O, and networking resources. Depending on workload, the supporting CPU is better enabled if it can offload some of these compute tasks to accelerators that drive better application-level performance. This is Intel's 4th Generation Xeon Scalable Processor.

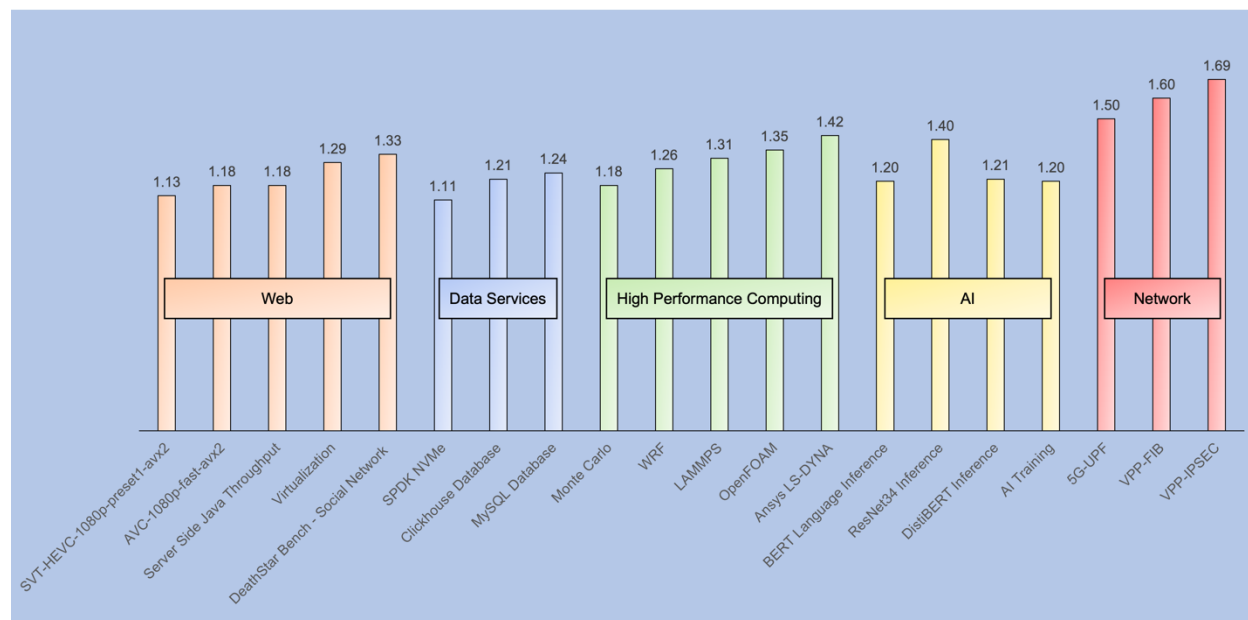
The concept of workload acceleration is nothing new to Intel, as Xeon processors have more accelerators than any CPU on the market. However, the addition of several new acceleration engines in the 4th Generation Xeon Scalable Processors is especially compelling:

- **Intel Advanced Matrix Extensions (Intel AMX)** offloads matrix operations such as matrix multiplication and convolution to a dedicated accelerator. This leads to better performance of AI functions like deep learning, training, and inferencing. As a result, workloads such as natural language processing (NLP), image recognition, and recommendation systems perform dramatically better and show that a discrete GPU may not be required.
- **Intel QuickAssist Technology (Intel QAT)** is an accelerator engine that speeds up computationally intensive functions such as encryption and compression. Additionally, Intel QAT frees up Xeon cores to focus on application functions.
- **Intel Dynamic Load Balancer (Intel DLB)** balances traffic across multiple CPUs and CPU cores. Workloads that generate a lot of traffic – HPC and cloud computing included – can run faster and more efficiently with Intel DLB routing traffic to free resources in real time.
- **Intel Data Streaming Accelerator (Intel DSA)** is precisely as its name implies. This accelerator engine offloads dedicated silicon's most common data streaming, movement, and transformation functions. This improves networking, storage, and data-intensive workloads that drive the digitally transformed business.
- **Intel In-Memory Analytics Accelerator (Intel IAA)** accelerates the performance of database and data analytics.

The partnership between Nutanix and Intel is significant, as the Nutanix Cloud Platform enables workloads to benefit from the acceleration capabilities of the 4th Generation Xeon Scalable platform.

With Intel’s 5th Generation Xeon Scalable Processor (codenamed Emerald Rapids), MI&S believes enterprise organizations will see an even greater workload performance benefit as new architectural capabilities are delivered. One such capability is a design point around reducing the number of compute tiles in each Xeon CPU. By reducing the number of tiles from four to two, Intel can further reduce the latency experienced in application performance.

FIGURE 3: 5TH GEN XEON DELIVERS SIGNIFICANT PERFORMANCE GAINS



*Performance relative to 4th Gen Scalable Xeon
Source: Intel*

An even stronger security posture is also enabled in the 5th Generation Xeon Scalable Processor. With Intel Trust Domain Extensions (TDX), Xeon delivers VM-level isolation and security. This allows applications to run with predictable, secure performance. Perhaps equally important is the negligible impact on VM and application performance.

The ability to lock down virtualized environments while enabling near-native performance is unique.

Finally, Intel made several improvements to power management, making this 5th Generation CPU highly power efficient. In fact, while the CPU runs in the same power envelope as its predecessor, Intel claims to achieve a 1.34x performance per watt advantage. Additionally, Xeon's Optimized Power Mode can save up to 100W per system in a two socket server as a server sits idle. When considering both of these capabilities across data estate, an IT organization can realize considerable progress toward sustainability goals as well as real TCO savings.

ACHIEVING CLOUD ECONOMICS WITH GREENLAKE

Delivering the cloud experience means delivering the promise of cloud economics. Public cloud deployments have seen a rationalization over the past few years as costs around data and utilization proved to be prohibitive to long-term adoption. Stated more plainly: The reality of cloud economics failed to meet the promise of cloud economics.

However, enterprise organizations have embraced the cloud operating model. As repatriation efforts took place, IT organizations were less prepared to reabsorb those operational and maintenance tasks that were shed.

With HPE GreenLake as the delivery mechanism for the Nutanix Cloud Platform, organizations can finally realize the promise of cloud economics without having to reallocate resources to those critical yet mundane management tasks.

Deploying modernized workloads on Intel-based HPE ProLiant DX servers running Nutanix Cloud Platform should enable the best of both worlds – a digitally transformed business that can spin up and run virtually any workload at any time in any location. An IT organization can deliver this environment without adding significant resources and additional staff.

PLATFORM MODERNIZATION CHECKLIST – EIGHT QUESTIONS TO ASK

Many options are available to IT organizations tasked with supporting the modern workloads required for the digitally transformed business. Before evaluating solutions, MI&S believes the following questions should be explored and answered in the affirmative before a solution is considered:

1. Can the platform deliver an authentic, hybrid, multi-cloud environment?

Put simply, does the solution under consideration enable the secure and performant deployment and management of applications across the entire data estate – multiple clouds, the edge, and the traditional datacenter? This is a critical consideration for organizations that require real-time analytics of data where it is generated.

2. Can the platform support even the most performance-intensive workloads?

The right solution has no limits - from VDI sessions to online transaction processing to advanced analytics and AI. The right solution delivers the control IT craves with the business's required performance of its applications.

3. Does the platform deliver resilience and security? And does this resilience and security come without a significant performance cost?

Resilience and security are the top priorities of every IT executive MI&S engages with. If a platform does not deliver on this, the solutions provider should be immediately removed from consideration. If these capabilities come at a performance cost, it indicates the vendor has not matured its solution appropriately.

4. Does the underlying hardware and CPU build in native acceleration for my most intensive workloads?

Virtually every workload benefits from acceleration. While discrete accelerators such as GPUs and ASICs can deliver the ideal platform, they are more costly, power-consuming, and sometimes less flexible. Further, the CPU powering the workloads, such as Intel Xeon, may already have accelerator engines designed into the silicon, pushing the need for discrete acceleration further to the right of the performance spectrum. These acceleration engines reduce cost, power, and management requirements.

5. Can my IT organization achieve the promise of cloud economics and infrastructure agility?

In a cloud environment, agility is critical. The expectation is that developers can consume infrastructure and platforms as a service, on-demand, and via self-service. The modern platform should be able to support this use case regardless of where workloads will run.

6. Does deploying this solution allow my IT organization to avoid acquiring workload-specific infrastructure?

If deploying modernized platforms in a cloud operating model still requires workload-specific infrastructure, the solution should be removed from consideration. Any application, anywhere, at any time is the very intent of deploying a cloud experience and cloud operating model.

7. Is my environment future-proof?

The next "killer app" that will enable a business to create differentiation is right around the corner. It may be some new variation on AI, a new analytics platform, or something entirely different. In any case, it's coming to market, and business users will demand it. Modernized platforms deployed today should be able to support the unknown through a richness of features, capabilities, and flexibility. Further, the providers behind such solutions should have a record of innovation.

8. Does this solution reduce my cost of ownership?

While investments are required in any technology deployed, such an investment should have a measurable return. Additionally, the right platform choice will reduce the total cost of ownership – not increase it.

While there are specific requirements every IT organization has around performance, cost, security, regulations, and the like, MI&S believes these eight questions can be asked universally. And if each is answered in the affirmative, the solution is worthy of deeper evaluation.

SUMMARY

Digital transformation is an evergreen trend. Businesses and IT organizations have been and will forever be transforming operations to better respond to the needs of customers, constituents, partners, and the market. This notion of perpetual change and its impact on technology has been evident with the hype of generative AI.

This cycle has always created a healthy tension between business users who require the latest and most performant applications and IT organizations tasked with deploying, provisioning, and managing such environments. This dynamic and the inability of IT organizations to keep pace with the business gave rise to shadow IT and the uncontrolled use of cloud instances to support specific needs of the business. And that

has resulted in exponential cost, complexity, and potential regulatory violations around data sovereignty and privacy.

While the most performance-sensitive business-critical applications have typically run on bespoke infrastructure deployed and managed separately, MI&S believes these days are past. Nutanix Cloud Platform, powered by Intel's Xeon Scalable Processors and delivered on HPE GreenLake, can provide the cloud experience for the user and the cloud operating model for IT. Further, MI&S believes this solution can deliver on the promise of cloud economics through a consistent, consumption-based charge-back model.

Each IT organization is unique in its mission, operations, and capabilities. Because of this, MI&S recommends each organization thoughtfully and strategically map out its requirements around platform modernization. This should be an exercise of asking difficult questions about required capabilities, performance requirements, and cost-benefit. Our eight-question checklist is a good starting point and will simplify the evaluation and purchase decision.

For more information on the Nutanix Cloud Platform and the Nutanix-Intel partnership, please visit www.nutanix.com/intel.

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR

[Matt Kimball](#), Vice President and Principal Analyst, Datacenter 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 mentioned in context, displaying the author's name, title, and "Moor Insights & Strategy." Non-press and non-analysts must receive prior written permission from 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 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 regarding such information's accuracy, completeness, or adequacy and shall have no liability for errors, omissions, or inadequacies. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The views expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators, not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on the future, they are subject to risks and uncertainties that could materially cause actual results to differ. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of this document's publication date. Please remember 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.

©2023 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of its respective owners.