Securing Data in the Agentic AI Era

Summary

Moor Insights & Strategy (MI&S) sees AI as transitioning from a nascent, emerging technology to achieving widespread adoption in the enterprise. Proof-of-concept projects have given way to cross-organization strategic planning and controlled pilots. But while AI is an asset to the modern enterprise, it is also a liability and a threat.

Data is more important in the AI era than ever since it fuels the automation that comes from agentic AI. However, data has also become increasingly vulnerable, as it is generated at unprecedented rates across the enterprise data estate, from the core datacenter to the edge and the public cloud. The criticality of this function makes protecting data the top priority for enterprise IT organizations.

This research brief examines how AI is impacting the modern business and how IT organizations face challenges in protecting the modern, distributed data estate as AI workloads accelerate from nascent to ubiquitous. Further, it explores how enterprise IT organizations manage this heightened security profile through platform security measures, such as those found in HPE ProLiant Compute Gen12 servers, protected by an HPE Integrated Lights-Out (iLO) 7 baseboard management controller (BMC).

The Agentic AI Enterprise

It is not hyperbolic to say that the avalanche of data accompanying the AI revolution has impacted the modern datacenter unlike any other technology inflection point. Data is being generated like never before across every endpoint. New and historical data feeds large language models, which in turn, feed the agents and assistants that enable enterprise-wide automation.

While discrete AI is familiar to many organizations, generative and agentic AI show the greatest promise for business transformation. Whereas discrete AI is used to solve a single task by performing a single function repeatedly, solutions built on these AI technologies shift the orientation of AI from vertical to horizontal without sacrificing depth. In fact, it is fair to say that the opposite is true — that the potential for generative AI reaches much further. And agentic AI will usher in the digital worker revolution, augmenting employees in ways that will increase productivity and operational efficiency almost exponentially.

Consider the example of a consumer electronics company trying to optimize its supply chain in a global environment where demand, costs, and availability are erratic. To deal with this, the company deploys an agentic AI environment that includes:

- A demand forecasting engine that analyzes sales data, market trends, and consumer sentiment via social media platforms
- A supplier risk monitor that tracks geopolitical events (including tariffs) and supplier health in the face of these issues
- An inventory balancer that looks at real-time SKU-level adjustments across factories and warehouses
- An autonomous logistics controller to enable route and shipment optimization, partly enabled by IoT devices

This environment leverages predictive analytics and multi-agent negotiation protocols with vendors and partners to optimize inventory levels and ensure resilient supply chains. It also employs reinforcement learning for dynamic adaptation.

FIGURE 1: ENTERPRISE AGENTIC AI

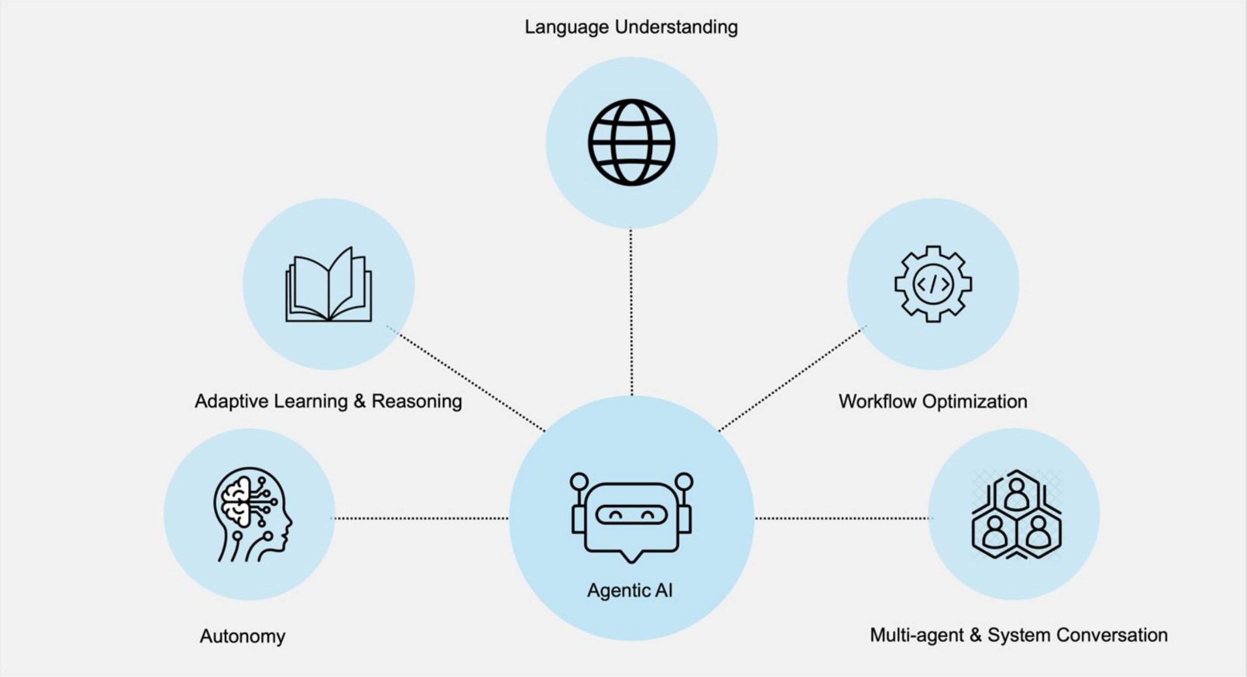
Language Understanding
Adaptive Learning & Reasoning
Workflow Optimization
Autonomy
Agentic AI
Multi-agent & System Conversation

The agentic AI enterprise workflow
Source: Moor Insights & Strategy

This example is a single process in a single business function. When considering this capability across the entire business, the number of agents augmenting business workflows will reach into the tens or perhaps hundreds of thousands, especially if we include personal assistants and agents that help workers with schedules, travel, and other routine tasks.

Jensen Huang, CEO of AI giant NVIDIA, envisions a future where his company could have 50,000 employees supported by upwards of 100 million agents or digital workers. While some may look at NVIDIA as an outlier, in an important sense it is simply an enterprise that designs, manufactures, and sells products. The same principle applies to a pharmaceutical company or an auto manufacturer.

Agentic AI Benefits Are Also Agentic AI Risks

Our consumer electronics example describes the incredible value agentic AI can deliver to an organization. However, when considered critically, the risks introduced into an enterprise become clear.

The ability to dynamically look across the organization, as well as across multiple suppliers and logistics companies, requires coordination and communication among hundreds of agents belonging to the company and its partners and suppliers. If manipulated, these agents could have an incredibly far-reaching, cascading impact where one agent infects another agent and so on. In short order, an exploit embedded in a supplier agent has propagated across an enterprise.

In addition to the increased attack surfaces and vectors, agentic AI can introduce risk through tools such as Model Context Protocol (MCP) servers, which connect agents with data sources. Like agents, MCP servers have the potential to propagate malware across the enterprise — and to partners — rapidly. The infection can happen via agents or the poisoning of data sources. Consider a supply chain scenario in which a compromised MCP server can potentially infect an organization's most critical repositories, halting manufacturing processes and other operations.

AI as a Tool and a Weapon

As enterprise AI adoption transitions from proofs of concept to deployments, MI&S has observed a trend where IT acts as customer zero. This model is helpful, as the organization can benefit from the automation that AIOps drives while learning how to roll out more broadly.

We have also seen AI infused into many of the cybersecurity platforms and solutions that protect our data estates. This can also happen through copilots and assistants that enable natural-language interfaces for IT professionals.

However, what can be used for good can also be used for nefarious purposes. Part of what drives the use of AI in security platforms is that hackers and nation-states utilize this technology to breach organizations at the lowest levels of the operating stack in a self-learning, amorphous way. This means that malware can change its form and properties in response to its environment to evade detection. When combining this amorphous behavior with the highly interactive nature of agents, the magnitude of the agentic AI risk is apparent. It is fair to say that agentic AI enterprises will be rich targets for hackers and bad actors.

Another wrinkle to this equation is the rise of post-quantum cryptography (PQC), a future threat that has to be addressed today. As most security professionals are aware, using quantum computing with Shor's algorithm can easily decrypt today's encrypted data in minutes. Shor's algorithm was developed to factor very large numbers quickly, something classical computers fail to achieve. This algorithm is fundamental to bad actors being able to break RSA encryption in hours (compared to roughly *300 trillion years* for a classical computer).

Although PQC has a horizon of five years or more, its anticipated emergence has opened a new battlefront — harvest now, decrypt later (HNDL) attacks. HNDL is a movement to capture encrypted data today so that when quantum resources are available, this data can be decrypted and used. MI&S believes bad-acting nation-states will be at the forefront of this movement, as they will have the earliest access to quantum resources.

For these reasons, security embedded at the platform level that can protect against the threats of today and tomorrow is not just nice-to-have but critical for the generative and agentic AI environment.

Full-Stack Security Is Rooted in Silicon

It is obvious that the AI datacenter requires full-stack protection; however, the meaning of full-stack protection depends on the organization. MI&S considers full-stack security to be protecting the platform from racking to recycling.

Frankly, full-stack security begins even before a server is manufactured. It starts with a design process that considers the real security challenges that organizations face in the present and future. And it requires turning those designs into protections that anchor upstream/up-stack security solutions by establishing an immutably clean operating environment on the servers that put data to work for generative and agentic AI-driven purposes.

During the full-stack security process, the silicon root-of-trust, embedded in the BMC, examines millions of lines of pre-boot code to ensure that the server operating environment is pristine at the lowest levels. After booting, the server operating environment and hardware components are verified as authentic and untampered. At that time, a broader security framework is employed to enforce ongoing security policies.

This pre-boot, hardware-based anchor ensures that the enterprise operating environment is safe. Otherwise, malware embedded in the OS's most privileged environment — ring 0 — can run undetected for months. This kernel mode is where device drivers run with full system access.

A BMC that includes silicon root-of-trust is the key to security. However, nearly all server vendors rely on third-party BMCs to deliver this functionality. In an era when geopolitics and other factors have caused supply chain angst, most security professionals are hesitant to simply trust a third party.

HPE — Platform Security for Agentic AI

HPE designs and manufactures its own BMC — HPE with iLO 7. While HPE has been designing and building iLO for generations, HPE iLO 7 enhances server security posture by incorporating a secure enclave to store and protect encryption keys and perform cryptographic functions. Effectively, iLO 7 acts as a local hardware security module (HSM) and connects to enterprise HSMs such as Thales, Utimaco, Entrust, and Atos.

The iLO 7 chip sits directly on the HPE ProLiant Compute Gen12 motherboard, with its own networking and storage designed to deliver this initial security and ongoing out-of-band monitoring and management. Having iLO 7 embedded on the motherboard enables administrators to monitor, manage, and secure the server regardless of power and CPU state.

FIGURE 2: SECURE BOOT WITH iLO 7

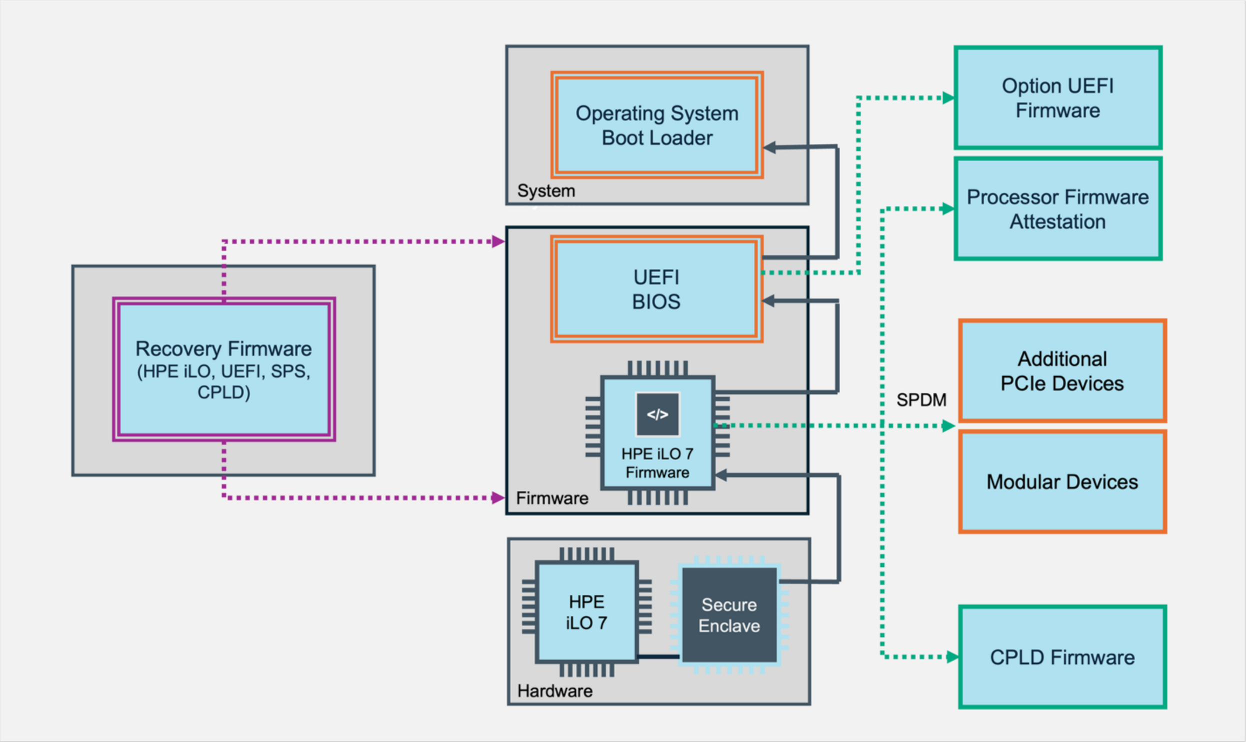

HPE iLO 7 ensures a pristine operating environment.
Source: Moor Insights & Strategy

iLO 7 is also tied to PQC, protecting against future threats. In addition to its secure enclave, iLO 7 employs the Leighton-Micali Signature (LMS), a hash-based quantum-resistant signature algorithm, to securely manage firmware updates, authentication, and remote management. The LMS-based chain of trust begins during boot, when silicon root-of-trust validates firmware components. By using LMS, users are assured that this highly vulnerable ring 0 state is fully protected from quantum-based compromise. LMS also secures against remote management functions performed with iLO 7, such as configuration changes and powering servers off and on. Again, LMS prevents quantum-based attacks from controlling these functions and injecting malware into a server's lowest levels.

In addition to ensuring that HPE ProLiant Compute Gen12 servers are booted into a pristine state, iLO 7 provides continuous firmware monitoring to prevent servers from being tampered with or hijacked. If these servers are found to have encountered firmware drift due to tampering, iLO 7 can automatically restore the server with authentic firmware and alert administrators. This process, known as runtime firmware verification, is part of why MI&S considers HPE ProLiant Compute servers worthy of serious consideration for those concerned with data privacy and security.

HPE ProLiant Compute Gen12 servers are the only commercial servers MI&S can identify as having achieved FIPS 140-3 Level 3 certification. FIPS, the Federal Information Processing Standards, was created to ensure that all U.S. government systems adhere to specific standards for security and interoperability.

In addition to FIPS, HPE ProLiant Compute Gen12 servers also conform to NIST (National Institute of Standards and Technology) and CNSA (Commercial National Security Algorithm) 2.0. Whereas CNSA is a set of guidelines applicable to both the public and private sectors, the stricter NIST deals with systems supporting U.S. national security. Interestingly, HPE does not disclose its full use of quantum-resistant algorithms. As the table below shows, the security frameworks converge on ML-KEM, ML-DSA, and SLH-DSA, the approved quantum-resistant algorithms.

FIGURE 3: COMPARING FIPS, CNSA, AND NIST

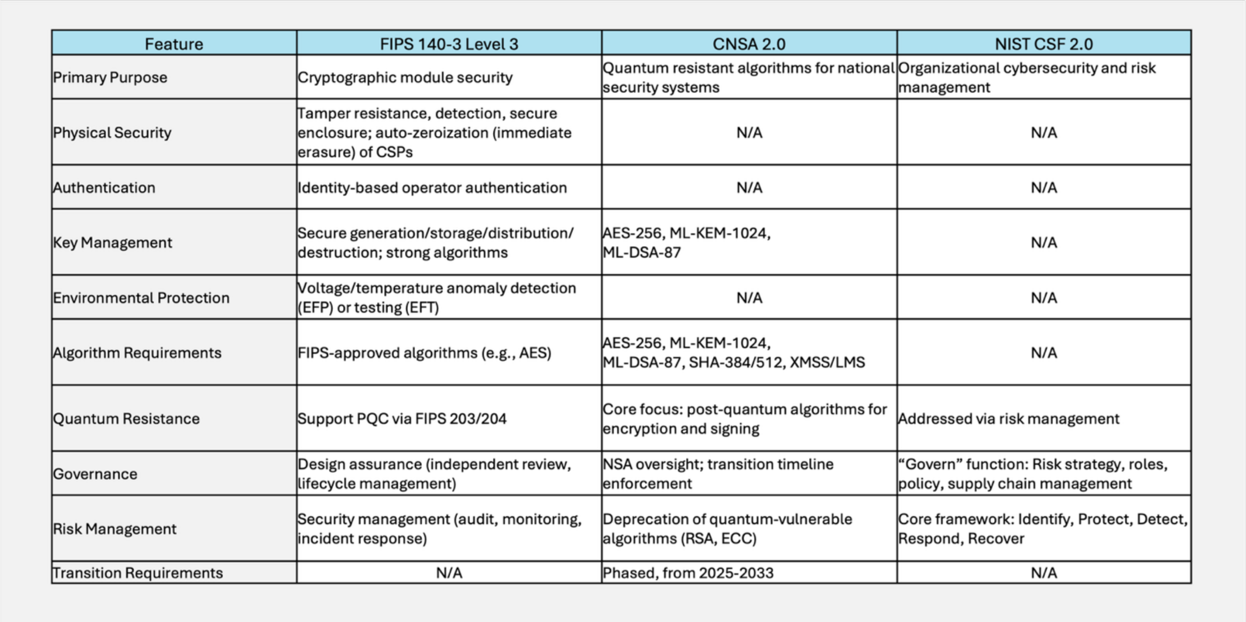

Feature	FIPS 140-3 Level 3	CNSA 2.0	NIST CSF 2.0
Primary Purpose	Cryptographic module security	Quantum resistant algorithms for national security systems	Organizational cybersecurity and risk management
Physical Security	Tamper resistance, detection, secure enclosure; auto-zeroization (immediate erasure) of CSPs	N/A	N/A
Authentication	Identity-based operator authentication	N/A	N/A
Key Management	Secure generation/storage/distribution/ destruction; strong algorithms	AES-256, ML-KEM-1024, ML-DSA-87	N/A
Environmental Protection	Voltage/temperature anomaly detection (EFP) or testing (EFT)	N/A	N/A
Algorithm Requirements	FIPS-approved algorithms (e.g., AES)	AES-256, ML-KEM-1024, ML-DSA-87, SHA-384/512, XMSS/LMS	N/A
Quantum Resistance	Support PQC via FIPS 203/204	Core focus: post-quantum algorithms for encryption and signing	Addressed via risk management
Governance	Design assurance (independent review, lifecycle management)	NSA oversight; transition timeline enforcement	"Govern" function: Risk strategy, roles, policy, supply chain management
Risk Management	Security management (audit, monitoring, incident response)	Deprecation of quantum-vulnerable algorithms (RSA, ECC)	Core framework: Identify, Protect, Detect, Respond, Recover
Transition Requirements	N/A	Phased, from 2025-2033	N/A

The U.S. Government details requirements for PQC.
Source: Moor Insights & Strategy

These algorithms are built on a mathematical structure known as a lattice. This is a grid of points formed by combining vectors in various ways. In functions such as key generation, public and private keys are derived from lattice structures. While the public key is associated with a lattice that is very difficult to reverse, a private key holds the secret to solving for it efficiently. Lattice-based cryptography is used for both encryption and digital signatures requiring quantum resistance.

Is iLO 7 an HSM?

This research brief previously discussed iLO 7 as a local HSM. This means it delivers much of the functionality of an HSM, and its secure enclave protects cryptographic keys and operations much like an HSM would. Additionally, iLO 7 has many low-level management capabilities. But iLO 7 is not an enterprise HSM like Luna by Thales, which is designed for enterprise-wide cryptographic key management and protection.

However, for enterprise organizations using enterprise HSM platforms such as Thales or Utimaco's SecurityServer, iLO 7 integrates by utilizing such servers to back up and store keys.

Measuring Security Readiness

Protecting the enterprise can seem to require both technical skill and a little bit of magic. With AI as a workload, tool, and weapon, the emphasis on magic may be understated. However, the first step to assuring readiness is a baseline measurement that an organization can build from and continually measure against.

The U.S. Department of Defense (DoD) has developed a Cybersecurity Maturity Model Certification (CMMC) to assess the readiness of contractors and organizations. In its 1.0 release, the DoD CMMC marked five levels of maturity, from basic cyber hygiene (ad hoc, no formal documentation) to advanced/progressive (continuously optimizing). They measure how an organization protects its information across 17 domains — from access controls to situational awareness to recovery.

These states of readiness, built on a cumulative structure, are designed to help organizations thoughtfully and thoroughly increase their cybersecurity posture. While many different guides can be used, MI&S has found the DoD CMMC to be especially helpful and thorough, as it applies perhaps the most stringent standards around cybersecurity.

However, it is important to note that security is a three-legged stool comprising people, processes, and technology. Even organizations with the tightest processes, the best-trained IT staff, and a vigilant workforce will be at risk for compromise without underlying technologies that protect from the lowest levels throughout the operating stack and network.

MI&S views platform security from HPE and its 12th generation of ProLiant servers as a key building block.

GETTING TO AN OPTIMIZED SECURITY POSTURE

It is natural and easy to approach enterprise AI security based on previous deployments of discrete AI, where machine learning was used to perform a specific function, such as visually inspecting products on an assembly line. However, agentic AI differs in terms of the number of tools required and its potential for widespread adoption. Not only will AI agents be deployed at incredibly high rates, but enterprise software applications will also include agentic AI.

Because of this, enterprise IT organizations must step back and think about cybersecurity through a different lens. This is additive to current practices, not a replacement for them. Below are five top-of-mind considerations for enterprise IT security organizations:

1. **Harden access controls and authentication**
   - Enforce strong authentication by requiring (phishing-resistant) multifactor authentication for access to agentic AI.
   - Apply the least privilege principle by using role- and attribute-based access to make sure agents have only the permissions required to perform a function.
   - Separate users from administrator roles and use privileged access workstations for sensitive tasks.

2. **Secure data and communications**
   - Encrypt all data, whether at rest or in transit.
   - Employ HSM to manage keys and cryptographic functions.
   - Protect against future threats today by deploying quantum-resistant protections.

3. **Continuously monitor and validate**
   - Perform red-team and penetration testing continually.
   - Continuously check data flowing in and out of agents to prevent injection attacks and other exploitations.
   - Regularly review logs and conduct audits to detect breaches and ensure compliance.

4. **Enforce zero trust**
   - While this may seem like an impossible task in agentic AI, every user, device, and agent must be continuously verified.
   - To mitigate exposure, agentic AI systems should be isolated and in secure subnets and containers.
   - Define and enforce strict constraints on which data agents can access.

5. **Secure your platforms**
   - Server and storage platforms are foundational to AI. Assure that the latest platforms are deployed and that embedded security is activated.
   - Enable runtime scanning of low-level drivers and firmware to assure that compromises are immediately recognized and mitigated.
   - Continuously test your organizational readiness in responding to (and recovering from) exploits.

Some of these actions may seem quite obvious, while others may not. However, in the attempt to deliver comprehensive solutions to complex challenges, a mindset that focuses on the operational execution of fundamental disciplines is critical for success. Mapping this to the DoD CMMC, or any good maturity model for that matter, is about executing the basics and building on a strong foundation.

CALL TO ACTION

The era of agentic AI is upon us. While we are still in the early stages of its enterprise adoption, MI&S sees a flywheel-like effect, where momentum makes broader adoption easier, faster, and more impactful over time. We believe the enterprise employee will eventually be augmented by dozens, if not hundreds, of digital assistants and workers.

With the promise of agentic AI, it is important to note the obvious: Its activation has made securing data exponentially more difficult. Hundreds of thousands of agents connecting with and sharing/extracting data from other agents across the enterprise and external partners is a security professional’s worst nightmare come true.

Data security strategies must be rethought to consider where data is generated and used. Protections at the lowest levels of the platform, where malware can reside undetected for months at a time while exfiltrating data, must be considered. Ensuring that platforms are booted to a pristine state and that integrity is continuously monitored and managed are foundational steps in a zero-trust environment.

HPE ProLiant Gen12 servers with iLO 7 deliver levels of security that can protect the modern data environment and AI. These servers are the anchor — the cornerstone for many organizations on the leading edge of AI adoption. With ProLiant Gen12 platforms, enterprise IT organizations can maintain the highest levels of integrity and protect against the threats of today and tomorrow. Because of this, MI&S suggests that IT organizations consider HPE and its ProLiant Gen12 platform as the first step in their agentic AI journey.

For more information, visit: https://www.hpe.com/us/en/hpe-integrated-lights-out-ilo.html

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 cited 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

Disclosures

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