

From pilot to production

The IBM AI infrastructure story



Table of contents

03	Foreword
04	Why building AI infrastructure from the ground up matters
06	The state of AI infrastructure
11	Challenges in AI infrastructure implementation
13	Why IBM Infrastructure for AI
19	Recommendations: How to go from pilot to production
22	Next steps



Foreword

Any sufficiently advanced technology is indistinguishable from magic.

—Arthur C. Clarke

Sixty years since it was formulated, with AI omnipresent in all aspects of our lives, Arthur C. Clarke’s 3rd law has never been more relevant.

AI may often feel like magic, especially to the people outside its development sphere. We are creating code, content and even art seemingly out of nothing. We’re accomplishing complex tasks with a few simple prompts, without much thought to the mechanisms making it all possible.

But, of course, AI is not magic. It’s engineering. And it has forever transformed the way people live and work every day. Behind the curtain, a complex infrastructure made up of specialized hardware, software and data is working hard to deliver exceptional computational power so we can achieve these “magical” feats.

Today, enterprise AI adoption is at an all-time high: **92%** of companies plan to invest in generative AI (gen AI) over the next three years.¹ And though there have been reports of AI pilots failing, businesses still believe in the promise of AI—because now that we can see past the hype and hope of AI, it’s clear that AI can deliver value. Gartner predicts that by 2028,

more than **95%** of enterprises will have used gen AI APIs or models or deployed gen AI-enabled applications in production environments, or done both.² It also reports that by 2028, **33%** of enterprise software applications will include agentic AI, and at least **15%** of daily work decisions will be made autonomously through agentic AI.³

The conversation, therefore, has shifted from whether AI can deliver value to how we can successfully take AI pilots to production and spread the value across the enterprise.

Having an AI infrastructure foundation is a key step toward success. The right AI infrastructure for business must be built with reliability and security as a core foundational layer—these qualities are just as critical as raw speed or performance—and not merely an afterthought. According to Gartner, **80%** of gen AI spending is going to hardware systems,⁴ indicating just how important infrastructure is to the success and the future of AI deployments.


Let’s delve deeper to understand why a robust foundation is critical to AI success and explore the true wizardry that lies in the infrastructure.



Why building AI infrastructure from the ground up matters

It's easy to focus on AI chips to the exclusion of other elements when talking about AI infrastructure. GPUs are trending, TPUs are exciting. But as AI research advances and we are introduced to newer AI models and newly minted agentic systems, one thing's becoming clear: the need for sophisticated AI infrastructure goes far beyond these high-performance computing (HPC) technologies.

Whether it's training, tuning or inferencing, AI models need an infrastructure foundation of hardware, software and networking systems to work well. The performance, reliability and scalability of the models are directly determined by the robustness of that foundation. Businesses banking on AI, therefore, need a unified AI infrastructure that can manage data, optimize workloads and ensure security and governance at every step—even when AI models evolve and their compute demands escalate.



New chips on the block

AI chips are specialized processors, such as GPUs (graphics processing units), TPUs (Tensor Processing Units), and ASICs (application-specific integrated circuits), designed to accelerate AI tasks. They excel at parallel processing and deliver the massive computational power needed for training and running AI models.



Show me the data! Challenges in data accessibility, management and security

AI is complicated. Without the right data to train and fine-tune models, AI journeys can end before they even begin. In fact, more than **60%** of AI projects that aren't supported by an AI-ready data practice are fated to fail.⁵ Yet, most organizations are still figuring out how to access, manage and secure the data needed for AI workloads. The data is all over the place—siloed away in on-premises data centers or scattered across multiple clouds—and managed by different business units that rarely communicate. And for many organizations, the data may be too voluminous and unstructured for their existing data infrastructure. According to IDC, **90%** of enterprise data is unstructured,⁶ and handling unstructured information can significantly strain older platforms designed for structured data. An overwhelmed data platform is an underperforming data platform—unable to process data efficiently and thus begetting poor AI performance and insights.


AI models often synthesize and process sensitive information, so data security is another critical concern. Can the existing infrastructure provide robust data storage and management? Can it implement mature governance practices and stringent security measures? Without these safeguards, organizations risk data breaches, user privacy violations and regulatory non-compliance.



Location, location, location: Scaling AI across a distributed infrastructure

Why are so many organizations struggling with identifying the best location for AI workloads? With factors such as data gravity, latency and data management regulations at play, workloads can no longer be confined to single data centers. To realize AI's full potential for their specific needs, organizations must optimize application and data placement across a distributed infrastructure. The strategic move is to place them close to where transactions happen and where insights are derived.

But this distributed placement can mean that some data might reside on premises, some in the public cloud, and still some in private clouds. In such a situation, AI solutions that can operate seamlessly across hybrid cloud environments are not just a nice-to-have; they are a must-have. And to enable these solutions, organizations need a robust, scalable and reliable infrastructure that can support workload mobility and data integrity while upholding stringent security measures across the entire hybrid IT estate.



The state of AI infrastructure

Though early adopters have deployed AI in the enterprise in some form since 2017, the watershed moment for modern AI came in late 2022.¹ That's when gen AI, particularly large language models (LLMs), broke containment and entered the mainstream. And it became apparent that AI, to quote renowned AI expert Kai-Fu Lee, could help “humans to think and create in ways we've never imagined.” Things, we know, have not been the same since.


AI technology has evolved at a dizzying speed. In under just three years, we've seen a dramatic acceleration in model capabilities, a definite shift toward multimodality and the inexorable rise of agentic AI. The need for more powerful, more efficient and more specialized infrastructure has never been more acute.

All models large and small

The transformer architecture was introduced to the world by Vaswani et al, through the Attention Is All You Need research paper in 2017.⁷ In the eight years since then, LLMs have undergone a rapid evolution. The inner workings of LLMs continue to remain mysterious and inscrutable, but their importance in the enterprise has grown exponentially, putting a spotlight on the infrastructure behind the technology.

Unlike traditional machine learning models, LLMs require high-performance, real-time processing to generate content on demand. LLM's breakthrough for gen AI created a massive demand for specialized computing power, shifting the focus from general-purpose computing to hardware and software optimized for parallel processing. As LLMs change and accelerate further, the demand will grow. The more powerful and complex these models become, the greater the computational load they'll place on the infrastructure.

With the emergence of lower-cost, small language and optimized models such as DeepSeek's R1, we're seeing a change in the prevailing wisdom about AI models and their infrastructure needs. DeepSeek turned the focus on to smaller, open-source models and highlighted how they can be more efficient than massive, general-purpose models, validating IBM's long-standing AI strategy.⁸ It also showed us that building a performance and cost-optimized model with limited compute resources is possible. IBM's fit-for-purpose models have already led to reductions in AI inference costs, making deployment and customization more accessible and cost-effective for businesses.⁸ These smaller models can run on powerful CPUs or specialized AI chips and be deployed on premises or at the edge, making them ideal for GPU-poor organizations.




When small models score big on ROI

IBM's success with domain-specific AI models subverts the common assumption that big, general-purpose models are inherently better. It proves that smaller, fit-for-purpose models can deliver more value than those massive systems—especially when they are open source and tailored to specific needs.

“Smaller, open-source models are how that future will be built ... For too long, AI has been seen as a game of scale—where bigger models meant better outcomes. But the real breakthrough is as much about size as it is about efficiency. In our work at IBM, we've seen that fit-for-purpose models have already led to up to 30-fold reductions in AI inference costs, making training more efficient and accessible.”⁸

Arvind Krishna
IBM Chairman and CEO




Hot chips, cool AI

In response to the escalating compute demand from AI models, there has been a flurry of innovation to accelerate and scale AI infrastructure solutions. As a result, there are now powerful chips, higher-performance networking solutions and energy-efficient platforms entering in the market.

NVIDIA continues to dominate the AI chips market, with a 92% market share as of 2024.⁹ But now many AI companies are building their own custom chips to reap the benefits. When tailored to specific language models for selected use cases, these chips can help cut costs, improve latency and accelerate inferencing.

There are also new entrants in the market, designing specialized hardware—such as AI accelerators—specifically for inferencing, training and in-process, transactional AI. These chips are engineered to deliver high performance while minimizing power consumption and are ideal for companies that are exploring energy-efficient hardware solutions to power their AI operations.

Modern software is already AI-infused. As its prevalence expands, and more and more applications—from analytics to operations—demand instant, model-driven intelligence, the need for high-performance infrastructure will rise even higher.



When you optimize the architecture of the hardware to work with the architecture of the software, that makes magic. It decreases the total cost and latency while increasing the throughput.”

Shobhit Varshney
Global Head of AI
Citi

Agentic AI is expected to drive nearly USD

6 trillion

in economic value by 2028.¹⁰

Here comes the agentic AI

Reshaping the AI infrastructure landscape further is the rise of agentic AI. Agentic AI refers to systems or programs that operate autonomously or semi-autonomously to perform tasks on behalf of a user or another system. Unlike AI assistants and chatbots that only respond to queries, agentic systems have “agency.” They are programmed to make decisions, take actions and interact with external environments beyond the data on which their machine learning models were trained—all to solve complex problems.

Agentic AI is expected to drive nearly USD 6 trillion in economic value by 2028,¹⁰ accelerating its role in enterprise decision-making, operations and automation. So, it’s not surprising to see agentic AI adoption soaring among leading organizations. According to the May 2025 AI Agent Survey from PwC, 88% of 300 surveyed senior executives say they plan to increase AI-related budgets in the next 12 months due to agentic AI, and 79% percent say AI agents are already being adopted in their companies.¹¹

AI agents are expected to make decisions and take actions in real time, fast. To deliver low latency and lag-free responses, AI agents need high-performance compute resources, reliable runtime environments, such as cloud services, containers and on-premises servers, and smart orchestration tools such as Kubernetes. They also need the right memory and storage tools to maintain the data and scale as needed.


Shifting gears from static inference to dynamic and autonomous orchestration puts even more focus on the underlying AI infrastructure. It’s the infrastructure that provides the foundational systems and environments that allow these autonomous, goal-driven agents to function effectively, reliably and safely.

88%

of 300 surveyed senior executives say they plan to increase AI-related budgets in the next 12 months due to agentic AI.

79%

say AI agents are already being adopted in their companies.¹¹



Challenges in AI infrastructure implementation

Despite the obvious benefits of having a robust infrastructure foundation for AI, organizations are still struggling to figure out their infrastructure strategy. Let's find out why.



Talent and expertise

AI requires a high level of expertise. But organizations are often unable to close the significant gap that exists between the demand for and the supply of skilled AI professionals.



Data storage and management

For AI applications to work well, vast amounts of data must be processed, stored and retrieved quickly. Without a data infrastructure designed to handle large-scale, unstructured data, eliminate persistent data silos, and address inconsistent formats and poor data governance, AI initiatives may fail before they even begin.



Security and governance

AI models often process sensitive data. But organizations' existing infrastructure may lack the necessary security and governance frameworks to protect both the data and the models from breaches, adversarial attacks and misuse.



Regulations

Governments and regulatory bodies are introducing new rules to address the ethical, security and privacy risks associated with AI. Organizations must ensure compliance management is built into their AI infrastructure strategy to meet current and future AI regulations.



Cost and complexity

As AI is increasingly woven into the fabric of business applications—from the core to consumer-facing services to the edge—cost and complexity can climb. Reducing infrastructure cost, maintaining deployment flexibility and managing implementations are necessary for avoiding project failures.




Compute and power requirements

AI consumes significant amounts of power, especially when training large neural networks. Traditional CPUs are often insufficient, and scaling up AI deployments requires high-performance GPUs, TPUs or specialized hardware. Not only can it be costly, it can also place enormous strain on the power grid and cause profound environmental changes.



Integration with legacy systems

Many organizations still rely on legacy IT infrastructure that is not designed to support modern AI needs. Integrating new AI workloads and infrastructure with these existing systems without disruptions can be a major technical challenge.



Why IBM Infrastructure for AI

Addressing today's complex infrastructure challenges requires more than just off-the-shelf solutions; it demands a comprehensive, integrated approach. Here's where it's important to highlight IBM's long history of AI research and enterprise computing as well as our commitment to hybrid cloud and open technology.


One of the earliest companies to invest in AI research, IBM has a storied AI history. IBM's AI infrastructure history, though not as long or well-known, is pretty remarkable too. Researchers at IBM began explorations into making AI-specific hardware more than a decade ago,¹² around the same time the chips arms race in Silicon Valley started. Their work led to the opening of IBM's [AI Hardware Center](#) and the adoption of a holistic approach to building AI systems from the ground up—spanning materials, chips, devices, architecture and the entire software stack. IBM Telum[®] microprocessor chips, which brought AI inferencing to IBM Z[®] and IBM LinuxONE, and [IBM Spyre™ Accelerator chip](#), which will bring gen AI to IBM Z, LinuxONE and IBM[®] Power[®] platforms, were developed at the AI Hardware Center.




IBM Infrastructure is bringing to bear that experience and working with organizations around the world to help them make the right infrastructure decisions.




data

 **Agentic AI**

 **AI platform**
 IBM watsonx.ai® + IBM watsonx.data® +
 IBM watsonx.governance® + IBM watsonx Orchestrate®

 **Hybrid cloud platform**
 Red Hat OpenShift

 **IBM Infrastructure**
 IBM Cloud, IBM Z, IBM Power, IBM Storage,
 IBM Technology Lifecycle Services



0101



The AI inferencing market is projected to reach USD

254.98 billion

in 2027, driven by the need for real-time, low-latency insights across industries.¹³

A comprehensive full-stack portfolio

That holistic approach to building AI systems is reflected in the IBM Infrastructure portfolio, which runs the gamut from cloud to core. Unlike other AI companies, IBM offers a full-stack AI solution, covering data storage, processing, model training, deployment and management. With end-to-end infrastructure solutions, IBM can help build a secured, scalable and robust environment designed specifically for AI workloads. Our full-stack approach, combined with access to IBM's software offerings, makes IBM the ideal AI infrastructure partner for businesses looking for an integrated approach to AI.

This cloud to core strategy is further bolstered by IBM's red-hot Red Hat advantage. Red Hat® AI solutions available on IBM Cloud® deliver open-source innovation to help organizations operationalize their AI strategies, leveraging the flexibility and security of hybrid cloud environments.

Tuning and inferencing support

In line with our strategy for fit-for-purpose models, IBM has infrastructure solutions to support the tuning market, primarily for organizations that are looking to turn generic AI intelligence into business-specific intelligence. These organizations want a base model that has been trained on general-purpose data and can be tailored to their specific use case, but they also want to retain their own data and proprietary knowledge. IBM can help the organizations train their models on our platforms, helping them build the model they need while also making sure their data is secured through our infrastructure solutions.

Tuning is a one-time—or infrequent—event, but inferencing happens continuously and at massive scale. The market for AI inferencing is expected to grow exponentially and reach USD 254.98 billion by 2030,¹³ driven by the need for real-time, low-latency insights across all industries. To support this growing market and the need for optimized AI software and hardware, IBM has been focused on getting a large set of LLMs attached to our workloads. On IBM Z and LinuxONE, we leverage the Telum II chip and the Spyre™ Accelerator card, which is available on Power as well. We also offer support on the cloud, independent of our own platforms.



Training support

Organizations involved in large-scale training are always looking for solutions that can support their growing data management needs. In the training market, IBM is working with these organizations—partners in many instances—to support their needs using IBM Storage solutions. Additionally, IBM offers GPU resources on IBM Cloud for model training and full deployment as needed. Red Hat AI InstructLab™ and Red Hat OpenShift® AI on IBM Cloud are other offerings IBM provides for training and tuning support.

Data storage support

AI's data needs are wide, varied and never-ending. No surprise then that the AI-powered storage market is expected to touch **USD 118.38** billion by 2030.¹⁴ To support this flourishing market, IBM has high-performance file, block and object storage solutions that are designed to meet the unique needs of AI workloads. IBM's storage portfolio for AI includes IBM Storage Scale, a scale-out file and object storage platform; IBM Fusion, a modern application data platform; and IBM Storage Ceph®, an open-source solution unifying block, file and object data.



Bringing AI deployments from pilot to production

AI adoption has never been higher. Yet, gen AI pilots at **95%** of organizations have failed to yield any returns so far.¹⁵ According to the GenAI Divide: State of AI in Business 2025 report from MIT, only 5% of integrated AI pilots are extracting millions in value.¹⁵

Why do many organizations get stuck in pilot purgatory, unable to scale their successful AI prototypes into production, their projects stalling forever?

IBM Infrastructure solutions are built to solve this problem. IBM understands that having the right infrastructure is critical to bringing AI deployments from pilot to production. Here is how IBM's infrastructure offerings help organizations successfully move from experimentation to full operationalization.

AI's infrastructure challenge

IBM's infrastructure advantage

AI is power hungry.
Deep learning—also inferencing and tuning—requires significant compute and energy resources.

IBM offers everything from high-performance GPUs, TPUs and specialized AI hardware to deliver the compute power required to make AI successful.

Our high-performance servers can meet low-latency, real-time processing demands and help scale AI across a global network—all while delivering best-in-class energy efficiency.

Our AI accelerators such as the IBM Spyre Accelerator chip are made for in-transaction inferencing and can be more efficient for the organization.

AI needs good data, fast.
Proper data storage, access and management are key to AI system health.

IBM's storage offerings and cloud solutions help address all the data concerns associated with AI production, whether it's retrieving data quickly, reducing silos, eliminating sprawl or improving governance.

IBM Storage Scale, notably, provides the high-performance, scalable data platform that AI workloads need. It's designed for storing, managing and accessing large amounts of unstructured data. This means rapid parallel access to massive datasets across hybrid cloud environments, faster model training and inferencing, and better data consistency and availability for distributed AI pipelines.

AI's infrastructure challenge

IBM's infrastructure advantage

AI has trust issues.

With the threat of misuse, breaches and attacks always present, minimizing the impact on trust and data security is critical.

IBM infrastructure platforms offer quantum-safe security and best-in-class reliability with up to eight nines of uptime—because your AI is only as secured and resilient as the infrastructure it runs on.

AI will break rules.

Ethics and safety of AI will always be a concern, so implementing governance processes is crucial to avoid any potential financial, legal or regulatory risks.

IBM solutions help automate and scale end-to-end AI governance—from selecting the right use cases to developing, deploying, monitoring and replacing models. We simplify compliance management, helping organizations reduce the risks and challenges associated with AI use.

Our high-performance servers such as IBM Z, LinuxONE and Power systems are designed to the highest standards to help manage even the strictest regulatory requirements.

AI craves integration.

AI projects often require integration with various tools, frameworks and platforms.

IBM takes an open and flexible approach to AI, fostering a strong open-source and partner ecosystem to help our clients prepare for their current and future AI needs.

AI has varied infrastructure needs.

A one-size-fits-all approach to AI seldom works. Success means embracing multi-environment deployment to meet AI's diverse requirements.

IBM provides flexible and high-performance AI infrastructure and supports different frameworks and tools that enable organizations to deploy AI models not just in the cloud, but also in distributed systems, on-premises environments, GPUs and AI accelerators.



Recommendations: How to go from pilot to production

When IT leaders evaluate infrastructure solutions to drive AI adoption and acceleration, they must look beyond a simple comparison of price and performance. They should consider solutions and best practices that support performance, scale, governance and security while aligning with the organization's current and future goals.



Assess, modernize and secure the core infrastructure

First things first: make sure your foundation is solid. Without a solid infrastructure foundation, your AI estate can crumble. Building that foundation is about more than just buying new hardware; it's about rethinking your entire IT backbone.

- Prioritize hybrid or multicloud environments with containerization and serverless options to handle dynamic AI workloads on the fly, scaling up or down as needed.
- Ensure ready access to HPC resources such as GPUs, TPUs and AI accelerators.
- Implement a modern data storage and management platform for seamless access, movement and governance of structured and unstructured data.



Enable a robust data strategy

Your data is your most valuable asset; so a smart data strategy is non-negotiable for AI.

- Establish controls for access, management and governance to ensure your data is AI-ready.
- Use gen AI to create synthetic data that mimics your real data, augmenting limited or sensitive datasets wherever appropriate.



Invest in scalable compute and storage

Once your foundation is in place, you need to think about growth. A pilot is one thing, but a production-ready enterprise AI solution requires more.

- Ensure that your infrastructure can scale to support a fully deployed AI solution as you pivot from concept to implementation.





Prioritize security, compliance and responsible AI

AI without trust is not only worthless but also dangerous. Building trust in AI means getting proactive about security and ethics.

- Enable model and data auditing to ensure explainability and track decisions of gen AI outputs.
- Stay ahead of regulatory changes by ensuring compliance readiness. Maintain control over where and how data is stored and processed to uphold data sovereignty.



Choose the right AI stack and ecosystem

In a crowded field of AI solutions, your success lies in making the right choices for your business.


- Weigh the tradeoffs between flexibility and performance. Open-source options, such as Hugging Face and Mistral, offer flexibility while proprietary ones, such as OpenAI and Anthropic, may offer better performance and support.
- Build a solid infrastructure for API management that can scale and secure gen AI APIs.



Collaborate cross-functionally and develop skills

AI isn't an IT-only project. AI pilots that are successful and sustain are the result of collaboration across the entire organization.

- Develop governance structures that align infrastructure with enterprise goals and risk tolerance through cross-functional leadership collaboration—from the CIO and the CDO to the CISO and the legal team.
- Don't wait to upskill your team. Invest now in a skills development roadmap to get your people up to speed on AI use within your enterprise.





Next steps

Build that robust foundation with IBM Infrastructure solutions and move your AI from experimentation to production with greater speed, efficiency and confidence.

[Learn more →](#)

ENDNOTES

1. [Global AI Adoption Statistics: A Review from 2017 to 2025](#), G2, 28 May 2025.
2. [The 2025 Hype Cycle for GenAI Highlights Critical Innovations](#), Gartner Research, 29 July 2025.
3. [Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027](#), Gartner Research, 25 June 2025.
4. [Gartner Forecasts Worldwide GenAI Spending to Reach \\$644 Billion in 2025](#), Gartner Research, 31 March, 2025.
5. [A Journey Guide to Delivering AI Success Through 'AI-Ready' Data](#), Gartner Research, 18 October 2024.
6. [UNTAPPED VALUE: What Every Executive Needs to Know About Unstructured Data](#), IDC, August 2023.
7. [Attention Is All You Need](#), Vaswani et al, 12 June 2017.
8. [IBM CEO: DeepSeek proved us right—AI is not about big, proprietary systems](#), Fortune, 4 February 2025.
9. [Prediction: This Will Be Nvidia's Stock Price 5 Years From Now](#), Nasdaq, 7 September 2025.
10. [The Rise of Agentic AI: The Leading Solutions Transforming Enterprise Workflows in 2025](#), Futurum Group, 7 March 2025.
11. [PwC's AI Agent Survey](#), 17 May 2025.
12. [How the IBM Research AI Hardware Center is building tomorrow's processors](#), IBM Research, 12 August 2025.
13. [AI Inference Market Size, Share & Trends, 2025 To 2030](#), MarketsandMarkets Research, February 2025.
14. [AI Powered Storage Market \(2025 - 2030\)](#), Grand View Research, 2025.
15. [The GenAI Divide: State of AI in Business 2025](#), MIT NANDA, July 2025.

© Copyright IBM Corporation 2025

IBM, the IBM logo, IBM Cloud, IBM Spyre, IBM Telum, IBM Z, Power, Spyre, watsonx.ai, watsonx.data, watsonx.governance, and watsonx Orchestrate are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/legal/copytrade.

Red Hat, OpenShift, Ceph and InstructLab are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

This document is current as of the initial date of publication and may be changed by IBM at any time.

Not all offerings are available in every country in which IBM operates.

Examples presented as illustrative only. Actual results will vary based on client configurations and conditions and, therefore, generally expected results cannot be provided.

It is the user's responsibility to verify the operation of any non-IBM products or programs with IBM products and programs. IBM is not responsible for non-IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

No IT system or product should be considered completely secure, and no single product, service or security measure can be completely effective in preventing improper use or access. IBM does not warrant that any systems, products or services are immune from, or will make your enterprise immune from, the malicious or illegal conduct of any party.

The client is responsible for ensuring compliance with all applicable laws and regulations. IBM does not provide legal advice nor represent or warrant that its services or products will ensure that the client is compliant with any law or regulation.