



**Hitachi Vantara**

# Thrive in an AI-driven future

Three steps to build the solid foundation, capabilities, and resiliency you need to enable AI today, and more confidently lead the AI innovations of tomorrow.





# Executive summary

**As AI continues to evolve and mature, so must our application, management, and governance of it.**

It's a changing landscape like no other; the risks are undeniable, but the opportunity for greater insights, efficiencies and cost savings are enormous.

As we move beyond simple chatbots and internal systems, accountability and explainability become more critical. High-quality data ensures AI outcomes are explainable and compliant with regulations—building trust with users and customers and ensuring unquestionable transparency. This is especially critical as the systems take on a more autonomous role in our daily decision-making and enterprise operations.

In this eBook, we'll address the big challenges and concerns that are holding back innovation and progress and provide the actionable steps you need to form a solid foundation for AI, increasing its value now and tomorrow.

## Overview

<i>AI adoption today</i>	3
<i>Considerations when deploying AI</i>	4
<i>Data quality builds ethical AI</i>	5
<b>Three steps to thrive in an AI-driven future</b>	6
<i>Turn AI into ROI with an AI solution tailored for business growth</i>	9
<i>Hitachi iQ is</i>	10
<b>Industry-focused AI-powered use cases</b>	11
<i>Optimize and streamline the development and deployment of AI and GenAI models</i>	13
<b>A nod to the future</b>	14

## AI adoption today

**76%** of organizations are rolling out AI for widespread or critical functions.

**37%** AI is already critical to a third of their businesses.

The democratization of AI has opened the floodgates to a generative AI solution for everything at every turn—but with the promise of opportunity comes the risks. From platform lock-in and questions around sustainability and ethics, to confusing pricing models and shadow AI. Those grey areas might cause you to slow, or perhaps even halt AI adoption. And that's not to mention the added challenge AI presents to maintaining resilience and compliance.

But IDC has projected that the global AI market will reach \$631bn by 2028. To put that into context, it stands at nearly \$235bn today! The growth of AI is exponential, and that's because, despite the question marks, the opportunities are enormous. **Speeding processes, freeing manual resource by automating repetitive tasks, and delivering deep insights that accelerate innovation and enable more personalized services.**

### The AI opportunity



*Increased performance and optimized operations*



*Increased reliability and resilience*




*Streamlined processes and joined-up visibility*



*Deeper insights and more accurate decisions*



*Improved efficiency*



*Faster outcomes and up to 75% improved ROI*

# Considerations when deploying AI

## Before deploying AI, organizations must consider:

💡 Scalability and reliable performance

💡 Ethical considerations and compliance

💡 Cost and ROI

💡 Sustainability

## AI has the potential to touch every part of an organization's IT ecosystem, and this can make it difficult to know where to start.

### Scalability and reliable performance:

In one study, when asked to rate the most important characteristics in enterprise data and storage infrastructure, 28% said flexibility and adaptability to meet changing data demands, 28% said the ability to track data and its usage/performance, and 28% also said scalability to accommodate growing data volume.<sup>2</sup> This need will only rise with the success of AI and GenAI, both of which are hugely dependent on the free flow of data and the scalability of systems.

### Cost and ROI:

Cost is a conflicting issue; on the one hand, it's holding organizations back from pursuing AI ambitions, with 27% most concerned about the cost of adopting GenAI. There is also the investment in new skillsets and talent to consider. On the other hand, it's seen as a significant source of value, with 37% stating the business objective for GenAI is to reduce operational costs, and 26% say the priority is to implement it to create new revenue streams.<sup>3</sup>

### Ethics and compliance:

Only 38% of organizations enhance training data quality to explain their AI outputs, and almost a quarter (24%) don't review the datasets they use to train AI for quality at all.<sup>2</sup> Poor quality data and/or training could lead to AI 'hallucinations' (causing inaccurate outputs) or bias. With many AI-focused regulations coming down the line globally, accuracy and high-quality data are key to avoiding brand damage and monetary fines.

### Sustainability:

According to a study by **Goldman Sachs**, AI will drive a 160% rise in data center power demand; unsurprising when it's also claimed that a single ChatGPT search can use ten times more electricity to process than a Google search. Depending on where you look, the figures vary. Still, the trends all point in one direction: the rise of AI-driven processes, particularly in data-intensive industries, will amplify the pressure on computer processing requirements and on the grid. It's why many large technology leaders are already turning to options such as nuclear power in a bid to find cleaner power sources.

## Build, borrow or buy?

It is possible to scale AI initiatives without stretching internal resources or needing a big upfront investment. Read the guide to learn about the different options.

[Read the guide](#) →

# Data quality builds ethical AI

As new AI regulations emerge and many are in development around the globe, it's crucial to get ahead by making sure that AI models are only trained with the most accurate data. This ensures fully transparent and explainable AI outputs, reducing risk in a space that's ripe with public attention. For this reason, also building resiliency and security into the data and systems that feed AI and GenAI models is paramount.

## Ten dimensions for data quality

Leave no stone unturned: the questions to ask when you're addressing data quality and resilience:





1. **Accurate:** does the data reflect reality?
2. **Complete:** is any crucial data missing?
3. **Consistent:** is the data uniform in formats, units, and naming conventions?
4. **Unique:** is there duplicate data?
5. **Reliable:** is data based on typical scenarios?
6. **Accessible:** is data readily accessible as needed?
7. **Timely:** is the data current and relevant?
8. **Traceable:** can data origins and processing be tracked?
9. **Tagged:** is metadata available to add context and meaning?
10. **Clean:** does our data contain personal, sensitive, or proprietary information?



# Three steps to thrive in an AI-driven future


## Step


To ensure the reliable performance of AI and GenAI solutions, today and tomorrow, you must **start at your core IT infrastructure**. Here are the most important questions to address...


-  **Do you need more computing power, more storage, or faster networking to process the massive volumes of data that will be created and moved around?**  
The reality is, legacy was never designed with this kind of demand in mind.
-  **Where will you deploy AI... on-premises, in the cloud or using a more flexible hybrid cloud model? Will it be managed in-house, or using third-party expertise?**  
The type of data, volume of data and sensitivity of the data will all need to be considered here.
-  **Do you have the capabilities to move safely from testing to deployment?**  
It's important to have a process in place for the entire AI lifecycle so that everything is optimized and fine-tuned at every stage.
-  **Are you ready with a data management platform to overcome data silos and unlock AI-based analytics?**  
It will need the ability to scale up quickly as its use becomes more widely adopted.





## Drive operational efficiency and unlock ROI with improved visibility and optimization.

 **Data is a huge contributory factor to the cost of running AI and GenAI solutions.** As such, scalable, intelligently automated storage that moves data where it's needed, when it's needed will deliver real value back to the business (low-cost long-term storage for archive data, only keeping high-value operational data in high-cost locations).

 **Next, it's important to look at data storage pricing models:** they should be transparent, and with options for consumption-based pricing to help scale costs with usage.

 **As you address the infrastructure, look to fast-performance, low-latency systems.** You want data that can move seamlessly to maximize the value of AI deployments.


 **There's value to be had in AI-based data analytics;** done correctly, it can unlock significant insights and ROI for business through opening new revenue streams and delivering greater efficiencies.

 **Crucially, all stakeholders should have awareness and buy-in of AI and GenAI projects to ensure a smooth road to value.** Cross-functional collaboration is also key here to drive success, while training at every level will further optimize investments and ensure the correct use of AI.

Done correctly, operational efficiencies driven by AI and GenAI solutions will speak volumes through the ROI back to the business.

**Step**

**2**



## To protect sustainability and ensure ethical AI you must address your data quality.

**How do you get data in the ‘right’ state?** To start, by optimizing data models through classification and data deduplication to ensure everything is being stored in the right place, and that everything being stored is necessary. Some data can be moved off-premises and into the cloud.

But, of course, as more computing moves to the Edge (and for other mission-critical data like ERP), demand for data stored on-premises to ensure resiliency and security will rise. **This is where modernization comes in**, with new, vastly more efficient servers and storage devices (that offer advanced file compression, for example) available on the market to reduce the overall energy needs of your data center.

### **It’s important to stay abreast of developments in processing capabilities too.**

Take the ARM processors as an example. By combining the ARM Neoverse-based Grace CPU with their Blackwell GPU architecture, NVIDIA has reduced energy consumption by 25x and boosted performance by 30x per GPU compared with NVIDIA H100 GPUs.

**Finally, by looking at your data center layout and the cooling systems you have in place, you can further drive down energy consumption and support rising sustainability targets.**

Step

## Unlocking optimal, sustainable AI

**Data optimization** to ensure efficient storage without compromising access

**Data center modernization**, reducing the physical footprint, and swapping in efficient solutions

**Consider energy sources**, switch to renewable energy and choose more sustainable, environmentally conscious suppliers and partners

# Turn AI into ROI with an AI solution tailored for business growth

## Value starts with understanding

A clear goal and visibility of the IT and data landscape are key to deploying AI solutions that deliver. But it's not always straightforward to pinpoint the most high-value AI use cases, evaluate data readiness, determine ROI and develop a strategic roadmap for successful AI implementation. It's why Hitachi Vantara has developed the *AI Discovery Service*. In just three weeks it gives customers:

### Week 01



**Opportunity analysis** of the organization's current landscape, from company priorities to inventory of internal data sources, risk identification, and opportunity prioritization.

### Week 02



**Use case selection** to optimize the business and priority candidates for pilots, based on value creation opportunity, risk, and data availability.

### Week 03



**Planning** with a proof-of-concept vision, approach and timeline aligned to value creation goals, hypotheses to validate success criteria—and a roadmap for next actions.

## → AI Discovery Service Results

Tangible deliverables from the AI Discovery process includes:

1. Research report on user insights
2. Prioritized AI use case roadmap
3. AI concept documentation and recommendations for proof-of-concept

**Book your AI  
Discovery Service**



## Hitachi iQ is...

**a comprehensive suite of services and integrated approach to propel the success of AI deployments**

Once you have a clear roadmap of where to start, you can speed success with Hitachi iQ. An industry-optimized AI solution suite, Hitachi iQ equips IT leaders with end-to-end capabilities as they seek new ways to automate processes, speed time-to-market and uncover new insights that inspire innovation and accelerate growth. The powerful, bespoke solution suite allows organizations to automate and speed progress with intelligent, performant, scalable and flexible GenAI solutions and services in a hybrid cloud environment.

With Hitachi iQ, AI deployments are empowered with unified access to data irrespective of where it resides, while ensuring explainability, lineage, data accuracy, security, and traceability at any given point for mission-critical solutions. And unlike other approaches, Hitachi iQ goes beyond basic integration and testing by layering industry-specific capabilities on top of the AI solution stack, so outcomes can be more specific and relevant to an organization's business.

## Hitachi iQ delivers



65%

*storage cost savings*



75%

*improved ROI*

# Industry-focused AI-powered use cases




By combining the deep data and storage infrastructure expertise of Hitachi Vantara, with the heritage and deep, cross-industry insight of the Hitachi Group, we have created AI use cases that are proven, reliable and demonstrable from the moment of deployment. These use cases are flexible, and can be applied across multiple, diverse industries where the underlying need is the same.

Here are a few examples to show what's possible...



## Up to 40% forecasting accuracy in logistics

Businesses face significant challenges in predicting demand for parts, products and services, with inaccurate forecasting leading to overstocking, stockouts, increased costs and lost revenue. While traditional forecasting methods struggle to adapt to market fluctuations and evolving preferences, with our AI-powered demand forecasting solution, machine learning models analyze historical data, market trends, customer behavior and external factors. The result is:




-  **20-40%** improved forecast accuracy
-  **up to 30%** reduced inventory costs
-  **up to 20%** boosted revenue



## Turning data into actionable insights for the energy sector

Mounting challenges face the energy sector, such as environmental impact, regulatory pressures, operational complexity from demand patterns, and cost and reputational risks.

Our AI-driven solution aggregates historical and real-time data from diverse sources, while advanced machine learning identifies anomalies and forecasts future CO<sub>2</sub> emissions, scenario simulation and more. Combined, these insights:


-  **accelerate** value
-  **guide** decision-making
-  **support** more sustainable performance in a sector under intense scrutiny



### 50% reduction in credit card fraud detection

Traditional fraud-detection methods often result in false positives, delayed response times and excessive manual effort by investigating teams. With our custom LLM-based fraud pattern recognition solution, financial organizations can achieve:


- 💡 **50%** reduction in investigation time
- 💡 **40%** fewer false positives



### Up to 30% cost savings in manufacturing

Unplanned downtime, equipment failures and suboptimal asset utilization all cause significant inefficiencies in manufacturing. To bring real-time visibility into asset health and predictive insights, we've designed an AI solution that integrates IoT and telematics sensor data with machine learning models to optimize asset performance. This continuous monitoring and automation mean manufacturers can expect enhanced operational efficiency and improved decision-making, together with:

- 💡 **25-40%** reduction in unplanned downtime
- 💡 **15-30%** cost savings in maintenance expenses
- 💡 **10-20%** increase in asset utilization



# Optimize and streamline the development and deployment of AI and GenAI models

The team at Hitachi iQ works with customers to implement an AI stack that delivers sustainable, ethical, reliable, secure, cost-effective and peak-value AI solutions.



## Personalized

Designed for the rigors of AI, engineered for industry outcomes, but customized to the specific customer needs.



## Faster insight

Accelerated architecture delivers massive performance starting at 600GB/s and 22M IOPS, resulting in as much as a 20x improvement in GPU resource utilization through improved processing.



## Scalable

Validated designs and solution blueprints provide the flexibility and scale to rapidly develop, test, and deploy modern AI solutions while adapting to fluctuating customer workload demands.





## Lower TCO

Lower cost, erasure coded, scale-out object storage safeguards data, catalogs for re-use and stores data long term while optimizing costs.



## Accuracy

Increase the quantity and quality of data to improve the reliability of results. Identify, classify, transform, move, consolidate and prepare data to get the most value out of AI initiatives.



# A nod to the future

As the technology continues to evolve, we're going to move from AI as a responsive tool to AI that operates independently, takes the initiative and self-optimizes—"Agentic AI". These 'mini experts', built on the data and systems driving our organizations, will increasingly rely on Small Language Models (SLMs), narrowing their focus. Complex processes will be handled by a network of smaller experts, dividing and conquering based on their specialized training.

As exciting as this is, the risks are undeniable. Organizations must prioritize accountability and responsibility, building robust frameworks to govern these systems and mitigate the potential for unintended consequences. Ensuring explainability will be key to maintaining trust and understanding the decisions made by these autonomous agents.

This future won't be without its challenges. Organizations must still balance innovation with governance to maintain trust as they scale. But for those willing to embrace this complex future, the opportunities are vast.

AI isn't just about technology; it's about reshaping how we think, interact with systems and create new ideas. And this next chapter of AI will belong to those who approach it with focus, intent, and the confidence necessary to evolve.

**Are there areas of your business where AI could unlock more efficiencies, insights, or power greater analytics? Good things are waiting.**

**To find out more and get a taste of what we could do together, click below.**

Turn AI into ROI →

## Sources

1. IDC Worldwide AI and Generative AI Spending - Industry Outlook
2. State of Data Infrastructure Global Report 2024: How AI is Shifting Data's Foundation
3. Analyst Report - Enterprise Infrastructure for Generative AI: A Foundation for Success

## Hitachi Vantara

Hitachi Vantara is transforming the way data fuels innovation. A wholly owned subsidiary of Hitachi, Ltd., we're the data foundation the world's leading innovators rely on. Through data storage, infrastructure systems, cloud management and digital expertise, we build the foundation for sustainable business growth.