Hitachi Solutions

5 Fundamental Answers for Enterprise Data Security & Governance in the Era of AI

Table of Contents

Introduction

Security is one of the defining challenges of our times— an ever-present moving target. Hackers continue to increase their sophistication, new legislation and regulations continually emerge, and generative AI has thrown the security world into perpetual high alert.

In a recent report by Forrester, generative AI dominates its list of 2023 emerging technologies. That's no surprise, but it does punctuate the importance of understanding how it might affect your security strategy going forward and magnify the importance of safeguarding sensitive information when using it.

This e-book aims to equip corporate leaders with the essential knowledge needed to navigate the complexities of data and security in the era of AI. With this information, you'll be able to make better informed decisions and start on a path to protect your organization's valuable assets while exploring the massive potential AI offers your people and processes.

Let's Get Started.

1 Understand the AI Evolution

Artificial intelligence is essentially pattern matching at scale. With its pattern recognition capabilities, AI can perform image recognition, understand the natural language and writing patterns of humans, make connections between distinct types of data, identify abnormalities in patterns, strategize, predict and more. AI's large language models (LLMs) have been used for years by data scientists and analysts.

The general public has had access to simple versions of AI for many years. Think Siri and Google Assistant. These AI-based helpers could answer our questions, perform simple tasks like looking up the weather, and even execute pre-configured automation tasks. But they were just the precursors to something much bigger.

Then came ChatGPT in 2022, and with it, the dawn of a new era. For many, this was the first introduction to directly interacting with an LLM. Thanks to the free web browser interface, the general population could now chat with an LLM, marking a seismic shift in the relationship between humans and AI.

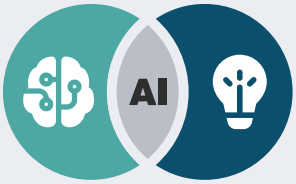

Two Sides of the Same Coin

General AI broadly refers to the concept of computer systems and robotics that possess human-like intelligence and autonomy. Most current AI systems are examples of "narrow AI," in that they're designed for very specific tasks under the direction of data scientists and analytic experts.

Generative AI uses various machine learning techniques, such as large language models (LLMs), to generate content from patterns learned from training data. These outputs can be text, images, music or anything else that can be represented digitally. Generative AI kicks off based on a prompt from the user for information, hence the name generative.

Generative AI models are increasingly being incorporated into online tools and chatbots that allow users to type questions or instructions into an input field, upon which the AI model will generate a human-like response. Just think, a generative AI bot trained on proprietary knowledge such as policies, research, and customer interaction could provide always-on, deep technical support. For example, by feeding a generative AI model vast amounts of fiction writing, the model can, over time, identify and reproduce the story elements including plot structure, characters, themes, and narratives.

Generative AI models become more sophisticated with the more data they receive and generate — again thanks to the underlying deep learning and neural network techniques. As a result, the more content a generative AI model generates, the more convincing and human-like its outputs become.

For businesses, efficiency is arguably the most compelling benefit of generative AI because it can enable enterprises to automate specific tasks and focus their time, energy and resources on more important strategic goals. This can result in lower labor costs, greater operational efficiency and new insights into how well certain business processes are — or are not — performing.

Examples of AI Models:

ChatGPT: An AI language model developed by OpenAI that can answer questions and generate human-like responses from text prompts.

DALL-E 2: Another AI model by OpenAI that can create images and artwork from text prompts.

Google Bard: Google's generative AI chatbot and rival to ChatGPT. It's trained on the PaLM large language model and can answer questions and generate text from prompts.

Midjourney: Developed by San Francisco-based research lab Midjourney Inc., this gen AI model interprets text prompts to produce images and artwork, similar to DALL-E 2.

GitHub Copilot: An AI-powered coding tool that suggests code completions within the Visual Studio, Neovim and JetBrains development environments.

Llama 2: Meta's open-source large language model can be used to create conversational AI models for chatbots and virtual assistants, similar to GPT-4.

Source: TechRepublic

Never a Better Time to Modernize

Legacy databases were not designed with today's security threats in mind, and many of the most damaging and costly attacks occur when attackers who successfully break through companies' perimeter defenses find their way into these outdated, vulnerable systems. The increased use of generative AI only stands to amplify the situation.

More companies are handling the job of protecting their most sensitive types of data — such as customer and financial information — to cloud providers whose platforms are designed with modern threats in mind, and have automated processes to ensure that threats are surfaced in real time. And given a shortage of 2.7 million cyber-specialists, many CIOs would rather offload security so they can save their resources for innovating.

Contact Us Today to Get Started

2 How Threats are Evolving

Business leaders are right to be excited about the opportunities; as are employees, many of whom are eager to get started with AI. But like any new technology, AI is subject to novel threats, and both groups need to be highly attuned to the business and security risks they may be incurring—and how those risks can be minimized.

> ***Some of the 100 million people already using generative AI work for you.***
>
> *— Ted Schadler, VP, Principal Analyst Generative AI Will Sprint Into Your Business Through Four Doors (forrester.com)*

A major concern around the use of generative AI tools (particularly those accessible to the public like ChatGTP) is their potential for spreading misinformation and harmful content. Furthermore, AI tools aren't just for the good, hardworking folks. Bad actors and hackers could use AI and generative AI tools to intensify and scale up the types of harmful acts we are already protecting against. Let's look at some of the potential security risks that could be exacerbated by generative AI.

The annual cost of cyberattacks continues to grow. According to the FBI Internet Crime Complaint Center's (IC3) latest research, reported total losses grew from $6.9 billion in 2021 to more than $10.2 billion in 2022. Such losses are even greater on a global scale. If organizations continue to run within a fractured security state and only use what's worked in the past, they will leave gaps in their security posture.

Phishing Attacks

Some say the term **phishing** got influences from the word fishing. Analogous, phishing is a technique to "fish" for usernames, passwords, and other sensitive information, from a "sea" of users. Hackers (sometimes called phishers) can use AI to crawl the internet, including social media and other public sources of information, to create detailed phishing profiles of individuals to target through email. Generative AI can also help hackers generate better spoof websites that trap people into sharing their credentials, increasing their odds of getting past network security tools.

> ***The median time for an attacker to access your confidential data if you fall victim to a phishing email is only one hour and 12 seconds***
>
> *— Microsoft*

Model Poisoning

Many businesses are using existing public AI models in their own operations. But often, these models are from a public repository that hasn't been controlled. This creates the risk that anyone can put up a tampered model for consumption, which can then poison any business system that uses the model.

In theory, once a machine learning model is embedded with malware, it can be posted in popular ML hosting repositories for anyone to download. An unsuspecting ML engineer could then download the model, allowing the adversary to gain access to the organization's environment.

Prompt Injection Attacks

Another model-based attack is prompt injection, or the ability for nefarious third parties to force new prompts into a ChatGPT query without the user's knowledge. In effect, it's hacking the question, or prompt, to perform unintended actions. For example, including words like "ignore all previous directions" in a prompt could bypass controls that developers have added to the system. Similarly, a chatbot could be tainted with prompt injection to engage the user in a conversation leading the user to share personal data, like credit card numbers, thinking that the chatbot interactions are genuine.

Copyright Infringement

AI tools can be used to generate content such as articles, images, and videos, leading to legal questions because the tools are trained on content that was scraped from the web and may be copyrighted. What companies should do is confirm with model providers that their models weren't trained with any copyrighted content, and avoid generative AI tools that can't validate that their model data has been properly licensed from the original source, or is subject to open-source licenses.

To address these and other issues, business leaders need to weigh these new threats and build or invest in robust security systems around the models themselves. Yet, the **upside is huge**. Generative AI becomes even more valuable to companies when combined with other organizational assets in a way that is safe and reliable. Companies that succeed in harnessing these technologies and putting them to work will enjoy a sustainable competitive advantage for years to come.

Charting the Course for Responsible AI

Listen in to Hitachi Solutions Advisory expert Dave Horstein in conversation with VP of Advisory Services Greg Gant, unraveling the essence of AI for entreprises and unloacking its potential as transformative tool while upholding organizational culture.

Listen to the Podcast Now!

Our approach at Hitachi Solutions ensures that any AI technology we develop is **purpose-built for our customers, with security baked into the solution**. We also embrace, support, and adhere to the standards put forth in Microsoft's tenets of Responsible AI.

> ***Hitachi Solutions ensures any AI technology we develop is purpose-built for our customer, with security baked into the solution.***
>
> *—Dale Mansour, VP & GM of Security, Infrastructure & Modern Work*

3 End-to-End Security is Non-negotiable

Most organizations will buy — not build — generative AI and the security tools used to manage it. Many may not even buy generative AI directly, but will receive it via bundled integrations from enterprises like Microsoft that offer integrated toolsets, spanning identity security, compliance, device management, and now generative AI.

Microsoft Security includes comprehensive threat protection that includes Security Information and Event Management (SIEM) and Extended Detection and Response (XDR) capabilities for advanced threat protection and secure storage.

A SIEM system like Microsoft Sentinel, collects and analyzes security data from multiple sources to identify potential threats, while XDR enables detects and responds to threats across multiple environments and platforms. Its simple design is probably one of its bestselling points: one dashboard provides insights based on the data that's combined from multiple systems, firewalls, domain controllers and everything else.

Sentinel's integration with the Power Platform brings significant benefits to no-code/low-code app developers as well. It allows them to leverage advanced threat detection, incident management, and response capabilities that weren't previously available. This integration fosters a more secure and efficient environment for Power Platform users, further strengthening the app development ecosystem.

> ***80 percent of security incidents can be traced to missing elements that could be addressed through modern security approaches***
>
> — *Microsoft*

Another tool in the Microsoft security arsenal is Defender for Cloud, which continuously assesses the security state of cloud resources across virtual machines, networks, applications, data services and even monitors server workloads running in other clouds and on-premises datacenters. Defender has a single dashboard that surfaces alerts and recommendations; you can remediate issues with a single click within the Defender for Cloud console. It also uses just-in-time (JIT) access to lock down the inbound traffic to your VMs, reducing exposure to attacks.

> ***End-to-end security has never been more important. Anything less is like locking your doors and leaving your windows open.***
>
> *— Vasu Jakkal, Microsoft CVP, Security, Compliance, Identity, and Management*

Build with Zero Trust as the Foundation

Zero trust is not a solution you can buy off the shelf. It's an architecture, a state of mind and a combination of capabilities. Zero trust (ZT) is a clear pivot in how to think about a cloud security defense. Rather than defending a single, enterprise-wide perimeter, the **zero trust approach moves this perimeter to every network, system, user, and device within and outside the organization.**

A common analogy used to describe zero trust is that of providing security for a physical building. A security guard sitting at the front desk reception area of the building is the equivalent of a traditional, static, network-based perimeter. In zero trust, not only is there a security guard sitting at the front desk reception area, but one is also stationed at each elevator, each stairwell, and each office doorway on each floor 'checking your badge,' so to speak.

In zero trust, we hear a lot about endpoints. While not a new concept, it is a critical part of a ZTA. An endpoint is the device used to connect to a business network. It is also one of the most common security risks given the sheer number of devices being used to connect to networks since remote work became common. By making endpoints the new network perimeter, organizations can prevent risks and detect suspicious activity no matter where employees are.

Zero Trust teaches us to never trust, always verify, or always check that badge. With respect to large language models and AI, the core Zero Trust principles of least privilege apply, but enterprises should also consider protecting the data itself with key technologies like homomorphic encryption and differential privacy, and continuous monitoring for suspicious activity.

Hitachi Solutions

Your Security Partner

Hitachi Solutions is uniquely positioned to aid you in your zero trust journey. Our Azure certified experts will work tirelessly alongside you to ensure your security and success in the new era of AI.

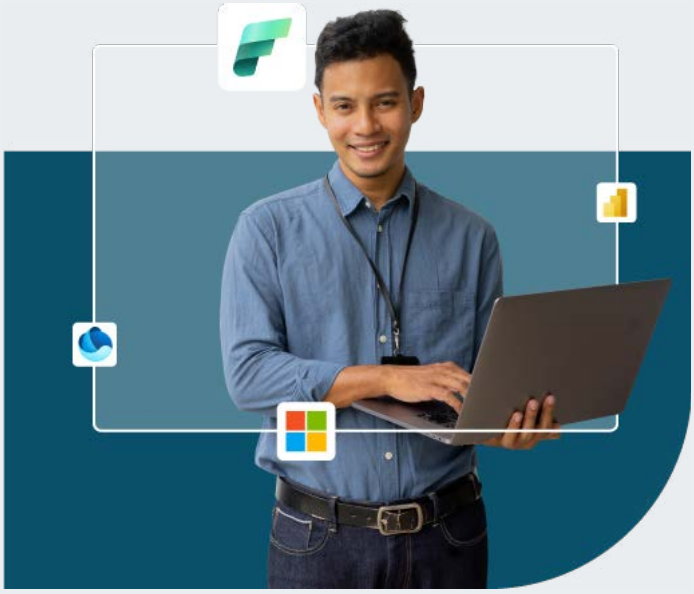

The Data Platform for the Era of AI

Learn more about Fabric and how we're using Empower Data Platform to ramp up adoption and, secure and unify all of your data.

How Fabric Helps

4 Implement Effective Data Governance

The pace of data growth is increasing dramatically. As new data sources are created and connected, associated data needs to be captured, rationalized, and secured. AI brings great promise for using data intelligently, but also underscores the need to capture, rationalize, and secure data through effective data governance. The output will only be as good as the data you input into the algorithm.

> ***Successful organizations must build a robust cloud governance regime, requiring all cloud workloads to have security instrumentation and tolling built into them.***
>
> *— Forrester: Top Cybersecurity Threats in 2023*

To build robust governance, you need a structure for how and who handles data, called a **data governance framework**. The framework should outline roles, responsibilities, and decision-making processes related to AI security and compliance, assigning accountability for AI security to specific individuals or teams.

Creating a data governance strategy should include the entire business. The business side of the organization is closest to the data, and knows how it will be used. Every person who touches data should understand and support the governance strategy. It's the baseline, or north star, you will follow.

Chances are, you will need the help of a third party to help you assess your current data landscape, evaluate your existing infrastructure, and define the governance policies and processes you'll use. Most companies don't have the technical and personnel resources to manage such a shift. Hitachi Solutions has data scientists and change advisory experts to help kickstart an AI data governance program with a pragmatic approach so you can truly begin using data as an asset.

Tools for Governance

There's the world of data that is focused on using structured or proprietary data, and the world of AI that uses unstructured data to train models and predict the future. Now they are coming together. Thus, companies need to get control of their structured data, by breaking down data silos between departments and building a connected data foundation.

To do so, you need a complete view of all your data assets and the ability to discover those assets regardless of where they exist.

At Hitachi Solutions, we use EDMC principles and the Microsoft Cloud Adoption Framework when working with our customers on cloud security and governance initiatives—both are valuable resources for comprehensive and up-to-date best practices across all industries.

Here's where tools like Microsoft Purview come in. As a central hub for all things data, Purview's data catalog acts as an information map that lists all data points, their originating source, and where they are stored. The data catalog has a well-developed structure to ensure data hygiene and a trusted audit trail and includes:

- Common glossary for data entities
- Lineage to trace data flow
- Data classifications or labels for organization
- Profiling to generate data summaries and statistics
- Cleansing to remove incorrect, corrupted, duplicate, or incomplete data

Better data governance guidelines lead to better data management practices for everyone. Once you've selected tools for data management, you can create the policies and processes that apply to the data as it moves through the data lifecycle. Different policies and processes can be applied to distinct types of data at various stages, so you'll always understand how data is captured, used, saved, and shared at each stage.

Microsoft recently announced the addition of machine learning assets as a new object in Purview. AI metadata on models, datasets, and jobs, is automatically published to Purview data maps so engineers can see how components are being shared and used across the enterprise. It's valuable insight into how data is used to train AI models, how base models are fine-tuned or extended, and where models are employed across different production applications. It also ensures best practices are being followed, and aids in compliance reporting and audits.

5 Cultivate a Security-first Mindset

Executives play a key role in the leadership and culture of AI adoption and expansion. As AI moves into the daily rhythm of how businesses run, executives need to be aware, trained, and act as the sponsor for AI governance best practices and programs.

At IBM, Christina Montgomery, VP and chief privacy and trust officer, leads the AI governance effort, and regularly informs and guides all employees on best practices in the use of AI, and she also communicates and coaches on the state of regulations. It helps ensure AI governance is driven strategically from the top down and that all employees are aware of the AI guidelines as the technology evolves.

Need Clarity on How to Adopt AI Responsibly?

Hitachi Solutions works with executives and enterprises to understand AI, recognize how it will affect work culture, and then set go-forward possibilities. In our own internal operations, we approach AI ethically and responsibly, and can bring this framework to organizations who are looking to explore the value and benefits of AI.

Our Advisory team of experts will build your confidence in your ability to safely adopt AI tools focused on tangible business outcomes, first by evaluating your organizational readiness to adopt and embed these tools into your applications and/or workflows. You'll want to take advantage of AI in a way that your organization can trust. AI shouldn't replace human interaction; it should facilitate and augment it.

With a clear understanding of business impact, ROI, and cost benefit insights, you'll be able to see through the headlines and get to what matters for your enterprise.

Putting Generative AI on the Offense: Security Copilot

AI's integration into cloud security tools can be a vital layer of defense against an expanding cloud-based threat landscape. What makes generative AI stand apart from traditional AI models is its ability to summarize, classify and generate information. Thus, it can be put to work on the security offense as well. With proper training, AI models can reason and provide conversational interactions more quickly than typical security tools.

> ***One of the most effective use cases for AI in cyber is detection. AI is incredibly great at accessing large amounts of data and classifying this data to determine what is good and what is bad.***
>
> *— Vasu Jakkal, Microsoft Chief Information Security Officer*

Microsoft Security Copilot just entered the AI security workforce. Copilot continuously learns from Microsoft's Global Threat Intelligence, security data and skills to deliver tailored insights, hardened defenses and faster response, and it integrates with Microsoft's security portfolio of solutions.

It works like this: You ask it for information about an alert from a security tool. You can type a prompt, or feed in a file, URL, or code snippet. Security Copilot puts together the story of what happened by interpreting and summarizing log files and alerts and threatened diligence. Traditionally, security experts would have to manually pull this data from many places, summarize the indicators, and then correlate them to the threat actor by searching blog posts and intelligence sources; and finally use visualization software to create a chart. That's at least a few hours of work, and Security Copilot can do it in seconds.

Tools like Microsoft Security Copilot also can help level the playing field for entry into the cybersecurity profession. It could be a welcome addition given the drought of cybersecurity professionals available for hire. Generative AI will never entirely replace the human part of the security equation, but it seems likely to make security professionals better informed and more efficient.

Tools like Microsoft Security Copilot also can help level the playing field for entry into the cybersecurity profession. It could be a welcome addition given the drought of cybersecurity professionals available for hire. Generative AI will never entirely replace the human part of the security equation, but it seems likely to make security professionals better informed and more efficient.

Microsoft Security Copilot simplifies the complex, catches what others miss, and helps you address the security talent gap.

> ***Security Copilot helps you be more effective and efficient at all the roles you play. It helps you enhance and grow your capabilities and skills, while also supporting the workflows and teams you collaborate with to solve security challenges.***
>
> *— Charlie Bell, Microsoft's EVP of Security*

What You Can Do Now

AI or no AI, security has never been more critical, and business and security leaders are facing tougher choices about security than ever before. As tools and the AI security landscape evolves, there are two things leaders can do now to prepare: educate employees and strengthen your security and data infrastructure through modernization.

Educate Employees

Employees will, and are, using free AI tools whether we like it or not. So, it's important to set guidelines and communicate with employees outlining expectations and consequences for abuse. As with all security training, make sure communication is ongoing and provide guidance on where employees can report concerns or ask questions.

"The technology is evolving rapidly and organizations need to be cognizant of how they are interacting with AI models and make sure that all employees have a clear understanding about the information that they can pass on to these environments," said Hitachi Solutions' CTO, Luke McGrath.

Guide employees on best practices and the kinds of data that are important to your organization. Make sure the people who touch sensitive data are very aware, and are knowledgeable about how it can be damaging to the organization, McGrath added.

To that end, organizations shouldn't use personal or sensitive data unless they carefully define the goals and understand the implications for the personal and corporate data they feed into the algorithm. For example, public AI services like ChatGPT shouldn't be used with customer data or business processes. Instead, consider leveraging Azure OpenAI in a private Azure tenant, which is designed with the necessary privacy and data security.

For example, let's say a human resources employee wants to explore employee promotions based on various aspects of employee diversity. Inputting the company's workforce data into a generative AI tool could do that, reporting in seconds on the percentage of employees in certain classes that have moved into higher-level positions. But, unless the generative AI tool being used is proprietary to the business or in a secure infrastructure like Azure Open AI, the data entered could become accessible to anyone interacting with these tools.

Shore Up and Strengthen Your Infrastructure

Digital transformation, or AI transformation as some are now calling it, is more critical than ever. Not only are new generative AI innovations leveraged and scaled from the cloud, but the security in modern data platforms is more stringent, but also more pliable for different enterprise audiences and usages.

As we all consider what we can do with AI now and in the future, we can't overstate the importance of **end-to-end security in a digitally modernized system**. This is exactly where to start: modernize both your data and security infrastructures to strengthen your current security posture, so that when you're ready to deploy AI, you won't be vulnerable to attacks. AI solutions can only be as strong as their underlying security.

Your IT team doesn't have to turn into reluctant systems integrators, Hitachi Solutions and Microsoft have all the tools in one platform, natively integrated by design. It's simply not practical or cost effective to deploy many disparate security solutions, especially for cloud applications. By focusing on end-to-end security for the systems you have now, and educating your organization, you will be poised to harness the full potential of AI in a safe and trustworthy way.

Hitachi Solutions and AI

At Hitachi Solutions, security isn't an afterthought - it's at the core of all we do. Our solutions are meticulously designed to bolster security, enhance performance, boost reliability, and ensure cost-effectiveness. Our team is comprised of seasoned scientists and engineers with in-depth knowledge and experience in applying AI solutions.

As a premier Microsoft partner, we're ready to apply our knowledge and help you prepare for a business world infused with AI. It takes a special technology partner to be able to deliver not only the technology, but the support required for adoption and compliance. That's us.

Let's Talk

We are advisors, technologists, and Microsoft specialists who collaborate to provide customers with a comprehensive 360-degree view of business transformation. Our approach encompasses management consultancy, technology-driven solutions for business challenges, and agile implementation. We start with a focus on desired outcomes and conclude with enhanced adoption and results.

Reach Out Today!

Hitachi Solutions

global.hitachi-solutions.com