

The Ultimate Guide to AI Readiness

Embracing the cloud to
unlock the full potential of AI





Contents

Meet the Experts	03
Chapter 1 Getting AI-Ready	04
Chapter 2 Bridging the Data Gap	06
Chapter 3 Implementing AI Securely and Responsibly	09
Chapter 4 Preparing for AI	13
Chapter 5 Moving to an AI-Enabled Future	20
Chapter 6 Industry Considerations	25
Chapter 7 Why Linux on Azure	33
Chapter 8 Why Neudesic	35
Next Steps	37
References	37

Meet the Experts

We'd like to thank the following people for contributing their time and expertise to The Ultimate Guide to AI Readiness. All thought leaders in AI innovation, their perspectives have been invaluable in shaping this guide.



Additional Experts

Charles Feddersen

Partner Director of Product Management –
Postgres and MySQL on Azure
Microsoft

Sandi Besen

Artificial Intelligence, Applied Research
Neudesic

David Bess

Global Vice President, Industry Solutions –
Energy and Communications
Neudesic

Raj Biswas

Global Vice President, Industry Solutions
Neudesic

Pedro Franco

Global Vice President, Industry Solutions
Neudesic

Sai Nageshwaran

Directory of Technology, Data, and AI
Neudesic

Tula Masterman

Principal Data Scientist, Applied AI Practice Lead
Neudesic

Erin Sanders

Global Director, Business Strategy and
Transformation, Responsible AI Leader
Neudesic

Chad Little

Director of Client Solutions
Neudesic

Steve Tetrault

Vice President, Innovation and Growth –
Financial Services and Insurance
Neudesic

Shannon Wu-Lebron

Vice President, Industry Strategy and Growth
Neudesic

Chapter 1

Getting AI-Ready

Getting AI-Ready

AI is changing everything, and early adopters are already at work putting their own AI environments into play. Indeed, more than 80% of enterprises will have used or deployed AI technologies by 2026.¹

Through purpose-built infrastructure with AI-first capabilities, businesses can streamline their operations, improve efficiencies, and maintain the flexibility needed to thrive in today's rapidly evolving landscape. When we sampled enterprise users who leverage AI in their workflows, they reported a 30% increase in productivity.²

When you incorporate AI into your day-to-day processes and technology stacks, you'll see some more impressive benefits, including:



Cost savings: AI provides new cost efficiencies by automating common expenses and redundancies, enabling digital transformation within budget.



Agility and innovation: AI enables enhanced central management and security, rapid application development, and data-driven insights, propelling new ideas forward and going to market faster.



Data optimization: AI is incredibly data intensive. Co-locating data on purpose-built infrastructure reduces siloes, improves responsiveness, and cuts latency, enabling faster development and deployment of generative AI solutions.

For every dollar invested in AI, the average return is \$3.50, with some organizations seeing returns as high as \$8.

Source: Neudesic. "Risk to Reward: The Business Case for Responsible AI." February 27, 2024.



Chapter 2

Bridging the Data Gap

Bridging the Data Gap

At the heart of AI is data, but **data preparedness is the most common blocker** for enterprises looking to develop and deploy AI solutions. The challenges—including latency/concurrency issues, data silos, and the high costs of data processing and on-premises infrastructure—can be daunting.

But the challenges don't need to slow momentum. By embracing the latest technologies, businesses can overcome these obstacles and succeed in their AI initiatives, achieving greater efficiency and scalability.

Let's start at the beginning, with some common considerations that will help you bridge gaps in data preparedness:

Data governance

Effective data governance ensures the integrity, security, and usability of data used in AI systems. AI solutions, especially those handling sensitive data, need data governance and compliance measures from the start to prevent risks and costly adjustments later.

Data security

Prioritizing data security ensures AI systems operate safely and responsibly. Businesses that adopt a security-first mindset will gain ready compliance with data regulations and a reputation for trustworthiness. Trusted data protection measures, including encryption, access controls, and advanced threat detection, help to safeguard private and confidential information and resist attempts to corrupt AI systems.

Upskilling

Successful AI adoption requires employees who have the right knowledge and skills, and executives who provide strong support for AI initiatives. While the pace of innovation and lack of standardization make it tough for even motivated employees to become AI experts, dedicated training and continuous learning programs go a long way. Remember, too, that incentive drives behavior, so companies that embrace and reward AI-first skills will see faster progress toward AI readiness.

Tool selection

Creating successful enterprise AI applications often requires multiple tools. Enterprises must carefully identify and use the right tools for the right job, while avoiding tool overload and ensuring efficient AI operations. By doing so, they can streamline processes, reduce complexity, and enhance the overall effectiveness of their AI initiatives.

Trust and reliability

As enterprises explore AI, trustworthiness and integration with broader IT systems become essential concerns. Ensuring that AI systems are transparent and free from biases is the formula for building trust. Transparency involves making AI decision processes clear to users, explaining how models arrive at conclusions, and showing these processes to be unbiased. Explainability provides some insight into an AI model's inner workings. This includes understanding the data sources, factors, and features that the model considers when making decisions; the importance it assigns to different inputs; and the logic behind its predictions. Knowing how the model arrived at a certain output is especially crucial in sensitive sectors like healthcare and finance, where this understanding can help ensure fairness and trustworthiness.



Bridging the Data Gap



Change management

This won't surprise you: implementing AI requires a robust change management plan. To address concerns such as fear of job displacement and resistance to change, companies should focus on culture, organizational structure, and leadership. Fostering a culture of continuous learning, open communication, and adaptability encourages employees to embrace AI technologies and provide feedback along the way.

Cost management

There's no denying that developing and deploying AI solutions can be costly, driven by factors like data quality, model complexity, infrastructure demands, and integration requirements. But companies can manage these costs effectively, starting with a repeatable process framework that considers a clearly defined AI problem, infrastructure needs, and long-term maintenance. Investing in data readiness and quality is also crucial, as clean data significantly reduces development time and minimizes costs.

Because AI processing is notoriously intensive, having access to hybrid or cloud solutions (like Microsoft Azure) provides much-needed scalability and flexibility, allowing organizations to adjust resources based on their needs. Likewise, through seamless integration, AI solutions can work in harmony with existing systems, which saves costs on multiple fronts, including minimizing disruptions and reducing the need for extensive modifications.

Chapter 3

Implementing AI Securely and Responsibly

Implementing AI Securely and Responsibly

When an enterprise uses AI in a secure and responsible manner, it not only ethically harnesses the technology's benefits but also minimizes risk to the business and its customers. By following AI best practices, frameworks, and principles, your company can become AI-ready more quickly and confidently.

Considerations for AI safety

As AI continues to become more common in our lives, AI safety grows increasingly critical. Advanced AI models have the potential to cause harm, either intentionally or unintentionally. By implementing strong AI safety practices, we can ensure these systems deliver meaningful benefits while protecting users from unintended outcomes.


Today, two major concerns in generative AI are content safety—effectively identifying harmful user- or AI-generated content—and prompt injection, which safeguards against malicious attempts to bypass system protections. It's essential to have guardrails in place to ensure AI systems operate as intended, protecting users from risks like misinformation that could undermine trust and cause harm.

Taking the lead on advancements in generative AI, Neudesic's solution is notable for both its security and efficiency, requiring less compute and training time. This simplifies customization to meet specific AI needs while keeping user protection at the forefront.³

Importance of responsible AI

AI is transforming industries by driving better decision making, improving customer experiences, increasing productivity, and generating higher revenue. According to PwC, AI could boost global GDP by 14%—or \$15.7 trillion—by 2030.⁴ However, challenges like privacy, bias, data quality, and regulation must be addressed to ensure AI's responsible use.

Responsible AI serves as a governance framework, guiding the development and deployment of AI technologies to maximize value while minimizing risks. It incorporates key principles like accountability, fairness, transparency, and privacy, ensuring that AI aligns with ethical standards and regulatory requirements. In high-stakes use cases, such as mortgage approvals or medical chatbots, unreliable AI can erode trust, harm users, and damage reputations. Implementing responsible AI requires accountability at all levels, from executives to engineers, ensuring everyone is accountable for their actions and decisions.⁵



Implementing AI Securely and Responsibly

Guiding principles of responsible AI

In general, the guiding principles of responsible AI include accountability, reliability, inclusiveness, fairness, transparency, security and privacy, sustainability, and governance. Each enterprise will have its own strategies for supporting and managing these principles, as what works for one may not work for another. It's important to keep this in mind when creating your organizational plan for responsible AI.

Here's a breakdown of each principle:

—> ACCOUNTABILITY

Ensures human oversight over AI systems, with mechanisms to hold teams and technology responsible throughout the project lifecycle.

—> RELIABILITY

Requires rigorous testing, validation, and monitoring to maintain consistent performance. It's important to note that AI systems can lose reliability over time if training data becomes outdated or if model assumptions no longer align with real-world conditions, which can erode trust and increase risks.

—> INCLUSIVENESS

Must be prioritized to guarantee that autonomous systems deliver optimal outcomes for a diverse range of users and stakeholders.

—> FAIRNESS

Ensures unbiased decisions and equitable treatment of all users, with continuous assessment to mitigate biases.

—> TRANSPARENCY

Is key to building trust in AI, requiring clear explanations of decision-making processes and associated risks. Organizations should document transparency-related questions for clarity.

—> PRIVACY AND SECURITY

Are essential, ensuring sensitive data is protected and access is restricted.

—> SUSTAINABILITY

Should be prioritized by minimizing the environmental impact of AI systems. Architectural decisions, such as selecting a smaller model, can have positive environmental, cost, and social impacts.

—> GOVERNANCE

Is necessary to align AI with ethical principles and with policies guiding development and monitoring for compliance.

Responsible AI principles must be paired with actionable strategies, and as AI evolves, organizations must continue to embrace safe and ethical practices to truly maximize the potential of AI.^{6,7}



Implementing AI Securely and Responsibly

Security and responsible AI

AI models require large quantities of high-quality data, including private data, for complex tasks. One approach to responsible AI involves securing data, model, and system infrastructure to enable the security and privacy necessary to protect sensitive data. Confidential computing technology ensures data remains encrypted even while being processed, providing an additional layer of security. Integrating security measures such as Trusted Execution Environments (TEEs) into a confidential computing strategy is one of the best ways to ensure your company's use of AI technologies is both effective and trustworthy. (A TEE is a secure area within a main processor that protects the confidentiality and integrity of the data and code loaded inside.)

Leading corporations worldwide are developing principles and frameworks for responsibly using AI technologies. Partnerships like Microsoft and AMD are front and center, deeply collaborating on confidential computing and AI initiatives.

AMD's broad compute portfolio delivers industry-leading performance, helping to save power, cost, and time to value—while assuring unparalleled security through support for confidential computing:

- AMD EPYC™ 9005 processor-based servers and cloud instances enable fast, efficient AI inference close to enterprise data, driving transformative business performance.
- AMD's high-end CPUs on Azure cloud, like 4th Gen AMD EPYC™ processor-based cloud virtual machines (VMs), enable quick, efficient AI with the fastest x86 CPU performance in Azure, helping transform businesses (source: [Azure Compute Blog](#)).
- Users of modern, general purpose AMD VMs can save 17% on cloud operational expenses on average, enabling additional investment into Azure AI services (source: [AMD Data Center Insights](#)).

Moreover, Azure leverages AMD Secure Encrypted Virtualization and Secure Nested Paging (SEV-SNP) at all three layers of the technology stack:

- **Container layer:** Confidential VMs enable the lift-and-shift of existing VM workloads to a TEE in Azure.
- **Infrastructure layer:** Confidential VM node pools on Azure Kubernetes Service (AKS) are available, along with specific container group isolation offerings.
- **Infrastructure layer:** Azure Confidential Clean Rooms provide a platform for multiple parties to share sensitive data, ensuring that only trusted software running inside the TEE can access that data.

NOTE: Confidential computing is only enabled on certain AMD processors on Azure. Please inquire for more information.

Chapter 4

Preparing for AI

Preparing for AI

While the thought of having to make infrastructure changes may seem overwhelming, inertia is even more so. Without moving forward, without innovating, organizations risk becoming obsolete.

A well-planned initial strategy and phased approach will keep things moving forward steadily and with purpose. To that end, we've compiled the following checklists of common considerations to help you determine your company's AI readiness. These checklists outline key areas to assess and address as you prepare the business to use AI effectively.

Strategy

- ☐ Define what you want to achieve with AI. Set clear objectives (measurable steps) and goals (long-term outcomes).
- ☐ Identify some specific use cases for AI within your company. For each use case, evaluate the problems being addressed and the value that resolution will bring to the business.
- ☐ Set key performance indicators (KPIs) and measures to evaluate the success and impact of your AI use cases.
- ☐ Assess whether AI is the right solution for the problem in each use case. Note, too, if there are any other tools or technologies that might be needed in conjunction with AI.
- ☐ Create a strategy for how data will be exposed to AI and other applications. Options include using custom APIs, following protocols like MCP, or enabling an AI-native platform and providing data based on natural language queries.

Data quality & readiness

- ☐ Start small: focus on data related to specific AI use cases you've identified.
- ☐ Plan to clean only the data you need. Ask yourself: "Do we need all this data or only a subset of it?"
- ☐ Identify the goals of exposing data and which applications or systems will consume it.
- ☐ Choose the methods to expose the data (custom APIs, standard protocols, or AI-native applications that support natural language queries).
- ☐ Map out required potential data sources and ensure access to them.
- ☐ Analyze characteristics like data types, distribution, and potential biases.
- ☐ Establish policies and procedures for data access, usage, security, and ethical considerations.
- ☐ Establish clear and structured processes for data cleansing, categorization, and labeling.
- ☐ Implement proper authentication and authorization mechanisms to control data access and ensure compliance with regulatory requirements.

Compute resources

- ☐ Analyze existing compute power, storage capacity, network bandwidth, and potential bottlenecks for AI workloads.
- ☐ Assess the computational power needed for your AI projects, whether it's for training or inference.
- ☐ Research and choose a cloud provider with robust AI services, scalability options, and suitable pricing models.
- ☐ Evaluate the need for specialized hardware like graphics processing units (GPUs) for computationally intensive AI tasks.
- ☐ Evaluate the pricing and performance models for resource usage to manage costs effectively.

Data infrastructure

- ☐ Transition the data plan from fragmented data systems (silos) to centralized solutions like data lakes and data warehouses.
- ☐ Evaluate and ensure your infrastructure is designed to handle both structured and unstructured data.
- ☐ Plan to migrate data to an accessible environment (on-premises or cloud-based).
- ☐ Implement tools for data integration, storage, and retrieval.
- ☐ Regularly audit and update data infrastructure to meet AI needs.

Preparing for AI

Security & compliance

- ☐ Implement robust data security and compliance measures (for example, encryption and access controls).
- ☐ Protect AI systems from adversarial attacks, threats, and data breaches.
- ☐ Ensure compliance with legal frameworks such as the General Data Protection Regulation (GDPR), Health Insurance Portability and Accountability Act (HIPAA), and standards issued by the International Organization for Standardization (ISO).
- ☐ Conduct regular security audits and risk assessments.
- ☐ Stay updated on industry-specific regulations and standards.

Governance

- ☐ Develop a governance framework for responsible and transparent AI use.
- ☐ Define ethical guidelines for AI development and deployment.
- ☐ Ensure AI models are explainable and auditable.
- ☐ Align AI initiatives with business objectives and regulatory requirements.
- ☐ Monitor AI systems for bias, fairness, and accountability.

Implementation & monitoring

- ☐ Start with pilot projects to test AI solutions before scaling.
- ☐ Define clear KPIs to measure the success of AI initiatives.
- ☐ Continuously monitor AI performance and make iterative improvements.
- ☐ Gather feedback from stakeholders to refine AI applications.
- ☐ Document lessons learned and best practices for future projects.

Skills & expertise

- ☐ Build cross-functional teams with data scientists, engineers, and AI specialists.
- ☐ Invest in AI literacy training for employees across the organization.
- ☐ Enhance the skills of your current workforce to effectively use AI tools and technologies.
- ☐ Partner with external experts or consultants if in-house expertise is limited.
- ☐ Foster a culture of continuous learning to keep pace with AI advancements.



Preparing for AI

Considerations for AI adoption

So, what is an “ideal” AI Infrastructure? There’s no one-size-fits-all model, but there are some common factors.


Like **cloud nativity**. By building on cloud-native platforms that offer scalability, flexibility, and high performance, future-proofing is achieved. Continuously improved cloud services provide the necessary computational resources and storage capabilities to support AI workloads.

Data management is another area to consider. A good place to start is by centralizing data through data lakes or data warehouses, but there’s more to the puzzle, as it can be a challenge to navigate data lakes and identify what data is important for a specific use case. In addition, making that data ready and available for AI to consume is a crucial step that can become complex without a logical approach to discovery and management. To ensure you have the right data, implement a metadata protocol and establish data governance procedures. From there, automate Extract, Transform, Load (ETL) and Extract, Load, Transform (ELT) pipelines, build quality checks into those pipelines, and optimize data wrangling processes to ensure that data is cleaned, validated, and prepared for AI use.

Moving to **model selection**, different AI models excel at different tasks, so it’s important to choose an existing model that aligns with the problem you’re trying to solve. For example, basic large language models (LLMs) are good for text generation, while models that power natural language services like ChatGPT are effective for tasks involving text processing, customer interactions, and content creation. If required, there are options for fine-tuning or customizing a pretrained model, and for full generative AI model training, high-performance computing resources like GPUs (such as those offered by AMD) ensure low-latency, real-time performance for AI-driven applications.

And then there’s a unique aspect of AI infrastructure—the concept of a **governance framework**. These frameworks ensure ethical AI use by promoting transparency, reducing biases, and aligning AI initiatives with regulatory requirements. As AI models are trained and deployed, governance tools help monitor their behavior, ensuring they remain effective and compliant. Security features must also be embedded to safeguard sensitive data and ensure AI systems remain resilient against threats.

Regulated businesses, especially, should establish **comprehensive governance and security policies** early on to ensure responsible AI adoption and deployment.



Preparing for AI

Comparing AI infrastructures

With these factors in mind, consider the following diagrams of a traditional versus modern infrastructure for AI and machine learning.

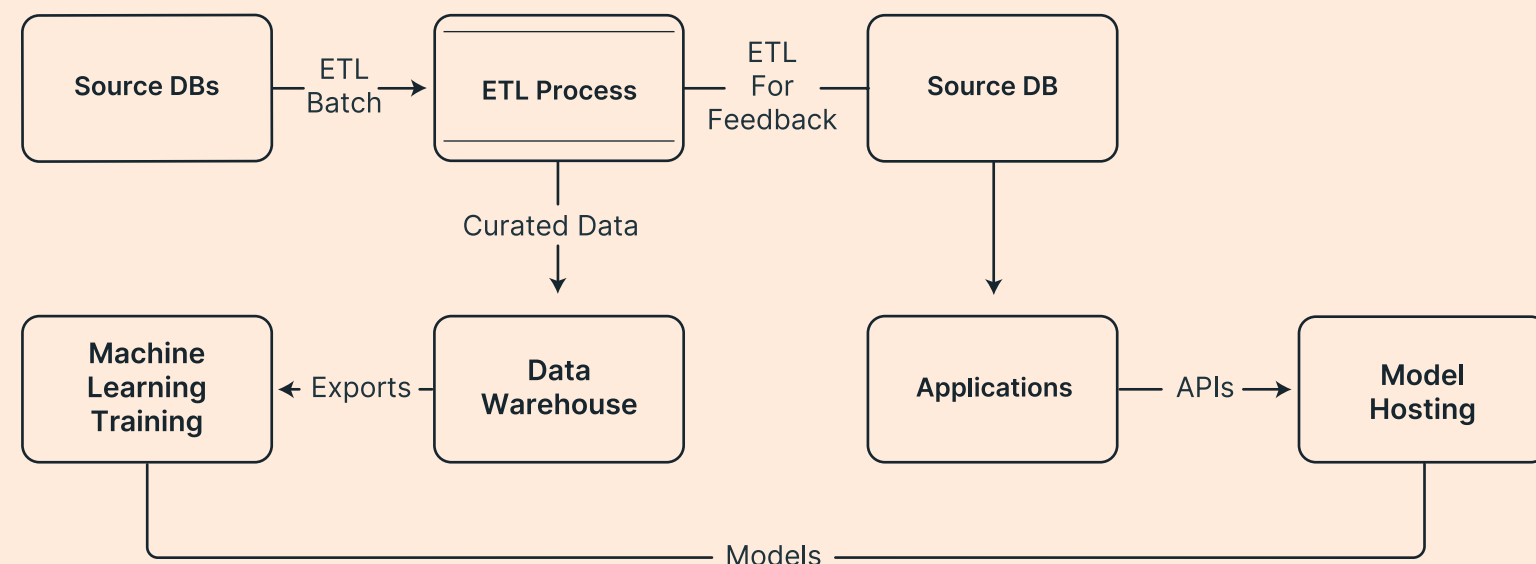
Traditional infrastructure typically follows a structured and sequential approach to data processing and requires manual model training:

- Data is sourced from various databases and applications, and then undergoes an ETL process. This is often done in batches, leading to delays in data availability.
- Curated data is stored in a data warehouse optimized for SQL-like queries, which may limit flexibility and scalability.
- Machine learning models are trained offline using data from the warehouse, with no pretrained models.
- The trained models are then manually deployed for use in applications, resulting in limited scalability and minimal real-time interaction.
- Updates require manual and full retraining, with feedback mechanisms that are slow or non-existent.

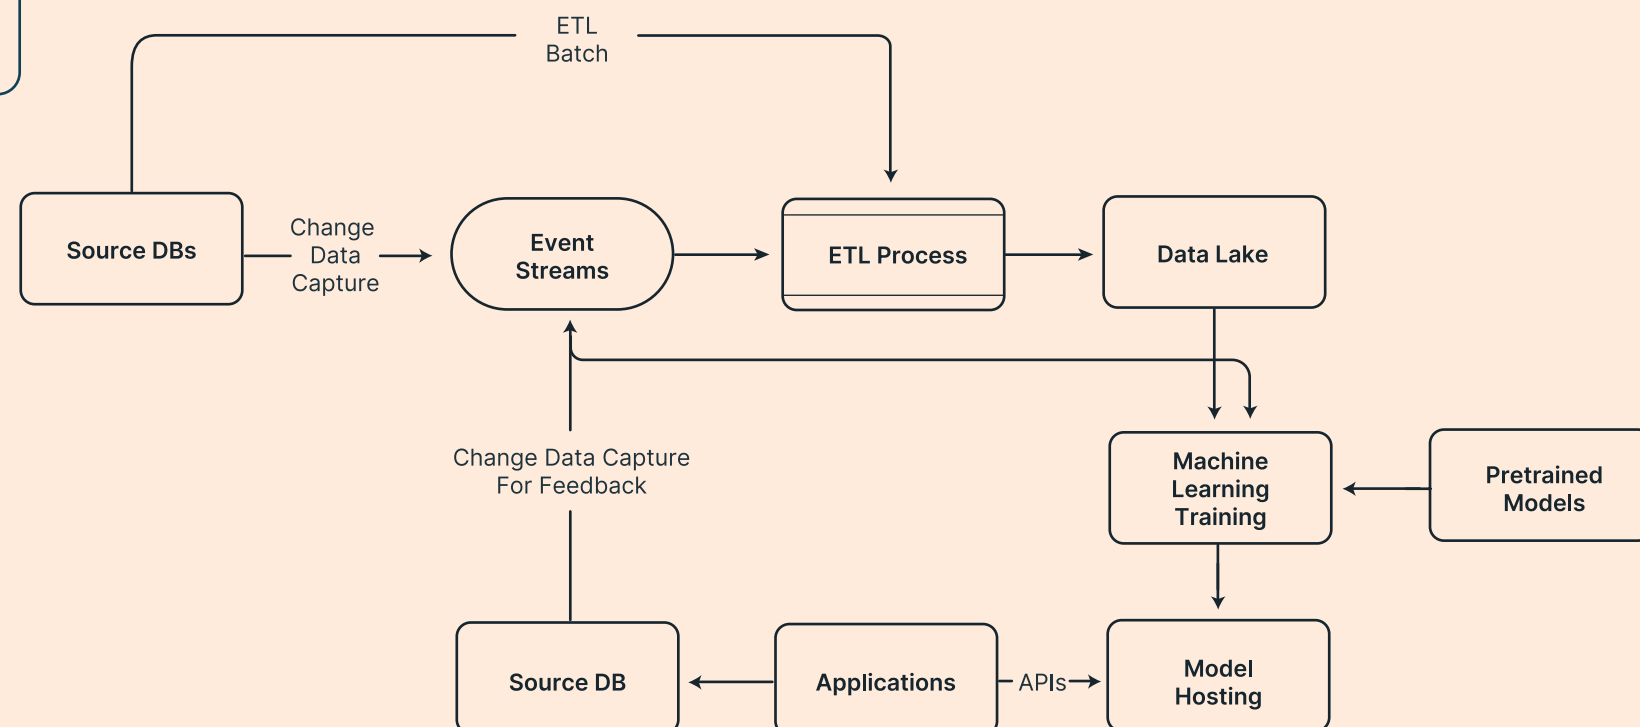
This traditional architecture results in limitations to performance and scale:

- Latency in data processing due to batch processing, requiring hours or even days.
- Rigid data schema design that struggles with unstructured data.
- Scalability bottlenecks in storage.
- Offline machine learning training, always from scratch.
- Manual and monolithic deployment with slow updates.
- Minimal adaptability to changing data patterns.
- Manual retraining feedback mechanisms.

Sample Traditional Architecture



Sample Modern Architecture



Preparing for AI

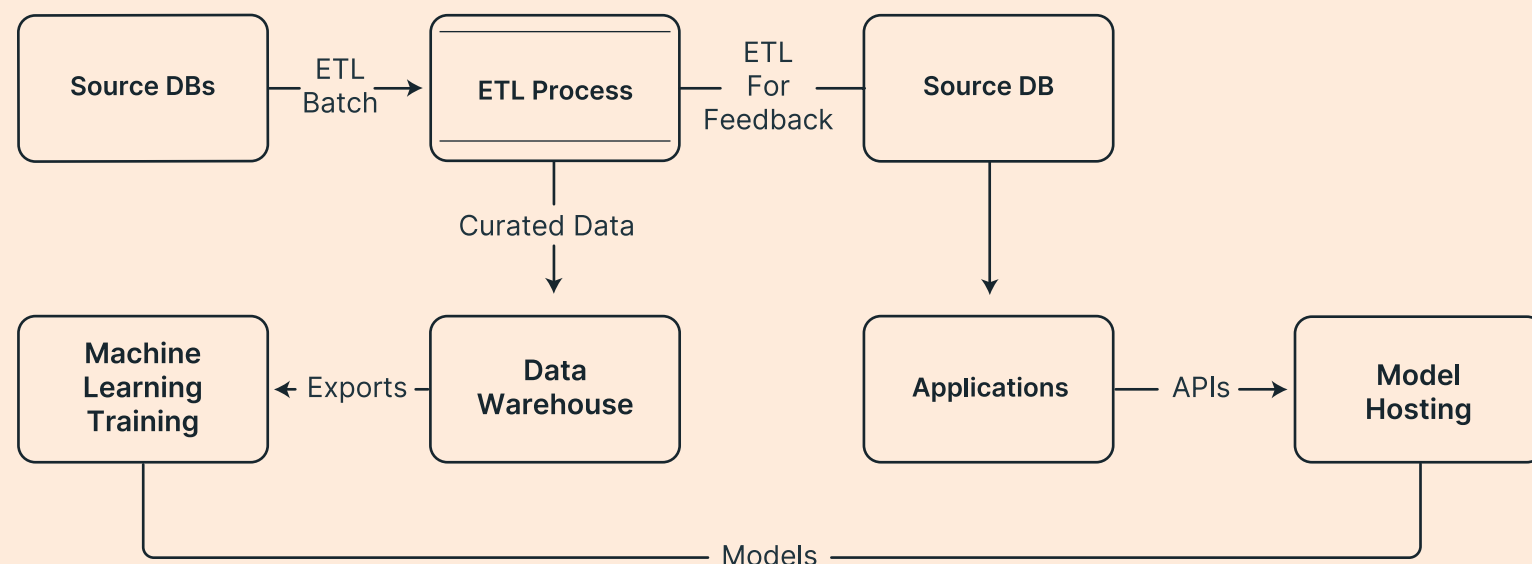
In contrast, modern infrastructure boosts performance and scale through real-time data ingestion and processing coupled with decentralized, API-driven, and iterative learning:

- Data sourced from databases and applications is ingested in near real-time through Change Data Capture (CDC) and streaming services.
- Data is stored and batch processed in a data lake, which can store raw structured/unstructured data. This enables scalable and flexible data storage.
- Machine learning training leverages pretrained models and supports online learning with continuous updates from event streams.
- Trained models are deployed, hosted, and accessible through API-first serving and scalable microservices, increasing flexibility and efficiency.
- Applications consume real-time predictions through APIs, while feedback loops capture new data and retain models for continuous improvement and retraining.

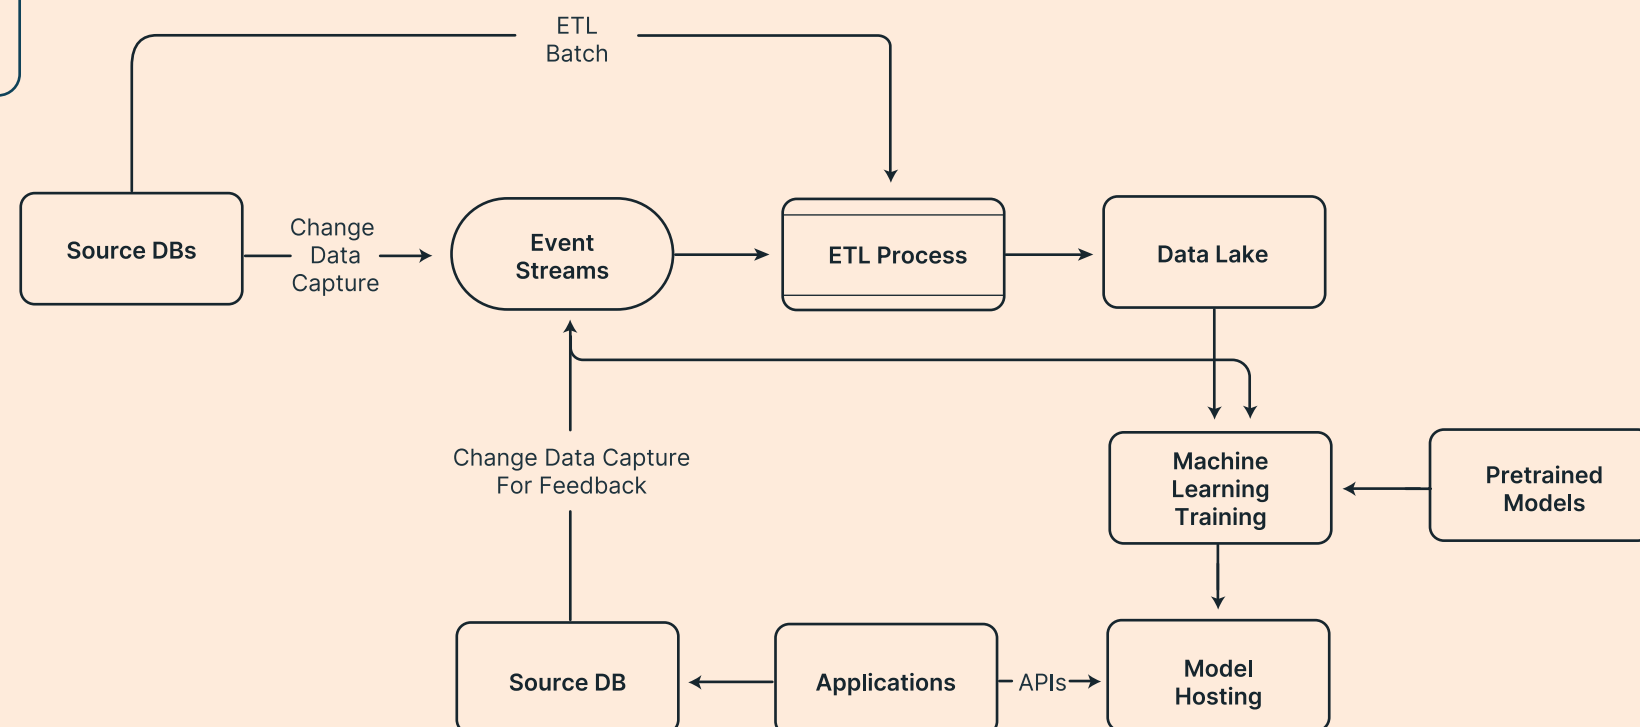
This modern architecture offers greater flexibility, scalability, and responsiveness, making it better suited for the evolving needs of AI applications. Key advantages include:

- Real-time data processing with low latency.
- Flexible and scalable data storage for raw and unstructured data.
- Agility through pretrained models, which enables rapid experimentation with new models and data sources.
- Continuous machine learning with auto-retraining through feedback loops.

Sample Traditional Architecture



Sample Modern Architecture



Preparing for AI



Building an ideal AI infrastructure

While the ideal infrastructure represents a fully mature AI ecosystem, most organizations start with more basic setups. Many begin with on-premises or hybrid cloud environments, which may limit scalability and flexibility as AI workloads grow. Data may initially be fragmented across disparate sources, making it difficult to integrate and centralize. This can hinder the full potential of AI, as AI thrives on consolidated and accessible data.

For initial AI adoption, companies often rely on simple, off-the-shelf AI solutions like preconfigured chatbots or productivity tools. These applications offer quick wins but may not deliver significant business impact without further customization. Many companies also focus on low-complexity use cases, such as automating customer service interactions or providing basic insights, before moving on to more complex, high-value AI projects.

Note, too, that AI governance and security may not be fully integrated at the outset. While essential, these aspects can be overlooked during early-stage adoption when organizations are primarily focused on getting AI systems up and running. As AI implementations scale, the need for robust governance and security policies becomes more pressing to ensure that AI is used responsibly and securely.

In terms of compute, most organizations start with basic cloud services and may not yet use specialized GPU resources. Early-stage AI models often rely on less intensive computational resources, focusing on predictive analytics or simple automation rather than training deep learning models. Similarly, data infrastructure may begin with data warehouses or databases that do not yet have the full capabilities of a data lake.

Looking ahead

The infrastructure required to adopt AI effectively varies from the initial stages to full deployment. While starting with basic systems and tools is common, moving toward a cloud-native, AI-driven infrastructure ensures the scalability, flexibility, and power needed to achieve higher-value business outcomes.

As AI technologies evolve, organizations must prioritize data readiness, computational power, governance, and security to ensure their AI initiatives deliver lasting value while remaining compliant and responsible.

Chapter 5

Moving to an AI-Enabled Future

Moving to an AI-Enabled Future

When moving to an AI-enabled environment, the cloud comes into play as the forge where enterprises can quickly access and leverage AI capabilities. There, you're free to experiment, build, and test, while taking advantage of built-in, guaranteed performance, resilience, and security. The cloud will move you quickly from AI as a concept to AI as a practical, production-ready solution—all enabled with the right mix of tools, services, and infrastructure to develop, deploy, optimize, scale, and otherwise support your workloads.

Sounds great, right? But the fact remains that most enterprises aren't currently set up to implement AI effectively.

Fortunately, there are ways to fast-track the process.

Companies seeking to accelerate AI adoption can engage an external AI consultant, like Neudesic, that has experience scoping and implementing AI across different enterprise environments. (Of course, we're biased, as we wrote this guide.)

Another way to "start fast" is to engage third parties that have already built the necessary infrastructure, data hosting capabilities, and AI services you need to adopt AI on your own timelines. In this case, of course, we mean Microsoft (incidentally, another sponsor of this guide).

So if the cloud is the way forward for AI, how do you get there? You start by getting your data ready, then you make the move—in your own way, at your own pace—and that gains you access to AI capabilities that can supercharge innovation and growth. The overall process looks like this:



Unlocking AI with Linux on Azure

When we think about enterprise data, there is no **single** database, application, or program that needs to be prepared for AI. There is always an infinite array of possibilities, which creates a frustrating amount of complexity.

To instill clarity here, we're going to focus on the specific use case of enabling enterprise Linux databases in the cloud—and unlocking the potential of AI. By bringing Linux and open-source databases to the cloud, we carry critical workloads forward into the future.

If your company already uses Linux-based systems, moving to Linux on Azure is seamless, with existing familiarity allowing your development team to optimize Azure's capabilities using their current skills.

Azure supports various Linux distributions and integrates well with existing systems—and tools like Azure Arc enable management of Linux servers across on-premises, edge, and multicloud environments from a single control plane.

Moving to an AI-Enabled Future

1 Preparing your data

If AI is the engine, then data is the fuel. And the cleaner your fuel, the better your engine runs. It's essential to refine and otherwise get your data estate in order before moving it to the cloud. Also, remember to clean only the data you need. In today's world, nearly everything involves data, so it's necessary to accept that with massive volumes of data, not all of it will be quality—and that's expected and perfectly okay. Just focus on the AI project at hand and the subset of relevant data.

Preparing your data and moving it to the cloud requires a structured approach to ensure the data is cleaned, organized, and ready to be consumed by AI models.

Tools and services within a cloud-based platform like Microsoft Azure can help at each of the following high-level steps:

- 1 Discover data:** Identify Linux-hosted data sources, including databases like PostgreSQL or MySQL; log files in /var/log; and APIs running on Linux servers. Also, determine the nature of the data: structured like SQL, semi-structured like JSON/CSV, or unstructured like syslog or Apache logs.
- 2 Collect data:** Collect data from identified Linux systems into a central repository.
- 3 Clean data:** Remove or repair incomplete, inaccurate, duplicate, or improperly formatted data.
- 4 Transform data:** Apply data transformation techniques like data format conversion, aggregation, or filtering. An example would be converting syslog formats to structured JSON.
- 5 Store and manage data:** Store cleaned and transformed data in a centralized, scalable, and secure location like Azure Storage Services to ensure it's accessible for AI model training and inference.
- 6 Establish data governance:** Implement policies for data privacy, security, and compliance. Monitor data usage and ensure proper access controls.

2 Moving your workloads

Now then, let's start with the most basic of questions: how does one get their Linux workloads onto Azure?

The answer is Azure Migrate, a free tool providing guidance on the best migration strategy, whether it's a lift-and-shift approach (rehosting) or a more complex modernization effort (refactoring or rearchitecting).

As you move through four steps—Discover, Assess, Migrate, and Activate—you'll be able to track the on-premises-to-cloud process with a centralized dashboard. Throughout the process, you'll gain valuable technical and business insights into your Linux infrastructure, databases, and applications, empowering you to plan your company's migration with confidence. Plus, with tools like Azure Migrate, Azure helps simplify the migration of Linux workloads.

3

Discover

In the Discover step, you'll evaluate your existing Linux environment to understand workloads, dependencies, and potential issues. This involves:

- **Discover:** Locating the Linux VMs running on Hyper-V, VMware, or physical servers, and identifying the applications, services, and dependencies on these servers.
- **Analyze:** Performing dependency analysis to understand the dependencies across the Linux servers you want to assess for migration to Azure.

Moving to an AI-Enabled Future

Migrate

5

It's finally time to make the official move, after settling on one of these two approaches:

- **Rehost (lift-and-shift):** For many Linux workloads, a lift-and-shift approach is the fastest way to migrate. This involves moving or copying existing VMs to Azure. Integrated tools like Azure Site Recovery (ASR) support rehosting by simply replicating Linux VMs to Azure.
- **Refactor/rearchitect (modernization):** Some older Linux workloads may need to be modernized. This is an opportunity to refactor or rearchitect applications with Azure platform-as-a-service (PaaS) offerings. For example, you can containerize applications using AKS for improved scalability, performance, and cost savings.

And what about your open-source databases like MySQL or PostgreSQL? You can migrate them to Azure fully managed database services like Azure Database for MySQL or Azure Database for PostgreSQL. These fully managed database solutions automate tasks like backups, patching, and monitoring. Plus, this move ensures your Linux workloads and database components are optimized and maintained on Azure, leading to reduced operational costs and improved business agility.

Take, for example, Azure Database for PostgreSQL, which offers a fully managed, AI-ready open-source relational database that can deliver substantial cost savings and performance improvements over an on-premises PostgreSQL database. Azure Database for PostgreSQL takes care of routine database management tasks so developers can focus on high-value AI app development. The offering also supports the latest PostgreSQL versions and a range of PostgreSQL extensions, with new AI features to improve the accuracy of retrieval augmented generation (RAG)-based AI applications.

Assess

4

Once discovery is complete, you'll have an assessment report in hand, which tells you more about key details:

- **Identify workloads:** Identifies which Linux workloads are ready for migration to Azure and highlights any compatibility issues.
- **Estimate costs:** Provides a detailed one-year cash flow analysis. This includes compute, storage, and network costs and will detail the financial benefits of moving to the cloud, validating that your company's migration strategy aligns with its financial and operational goals.
- **Recommend resources:** Analyzes the performance data of Linux workloads to recommend appropriate Azure VM sizes and services. This helps to ensure resources in the cloud are not over-or under-provisioned.

Economic benefits of moving to Azure Database for PostgreSQL:

- Up to 58% lower TCO migrating from on-premises PostgreSQL databases
- \$770K savings over three years in avoided downtime costs
- 65% improvement in database performance

Source: [Nathan McAfee, Enterprise Security Group. "Economic Validation: The Economic Benefits of Migrating On-Premises Databases to Azure Database for PostgreSQL." March 2025.](#)

Moving to an AI-Enabled Future

Activate

6

After your Linux workloads and databases are migrated, you're ready to gain a competitive advantage by harnessing the power of [Azure AI Services](#). This is where it gets good, as you now have a complete, authoritative set of services to build your AI applications. Azure AI Services deliver advanced capabilities like natural language processing, computer vision, intelligent search, image recognition, automation, translation, and more. **Linux Workloads on Azure empower your teams to:**

- > Build AI applications powered by foundation models, including those from Microsoft, OpenAI, Meta, Cohere, and more. Rely on a wide range of predefined tools, frameworks, and templates designed for development, deployment, and management.
- > Develop dedicated AI applications with prebuilt and customizable models, and use cognitive services to enhance automation, insights, and experiences.
- > Build your own agents with cutting-edge language and vision models. Use prebuilt AI-powered features and machine learning to create conversational interfaces, read text, and analyze images.
- > Enhance the safety of generative AI applications with advanced content guardrails, including custom AI filters and groundedness detection. Identify either protected or harmful user- and AI-generated content, including text, images, and multimodal APIs.
- > Deliver a better search experience with a sophisticated retrieval strategy that returns the most relevant data using keyword, vector, and hybrid search.
- > Automate tasks, like HR or IT support requests, to enhance user experiences and improve efficiency.
- > Quickly infuse generative AI into production workloads using studios, SDKs, and APIs.
- > Continuously innovate while staying up to date and secure—with new features and services tailored for Linux, including Azure Confidential Computing for secure data processing.
- > Capitalize on intrinsic, Azure-based AI and machine learning capabilities. Azure's AI and machine learning services, such as Azure Machine Learning and Cognitive Services, are fully compatible with Linux environments. You can build, train, and deploy AI models using Linux-based frameworks like TensorFlow and PyTorch.

Canonical Ubuntu on Azure has resulted in an average **306% ROI** with an 11-month payback on investment, as well as a **35% lower three-year cost of operations**.

Source: [IDC Business Value White Paper sponsored by Microsoft. "The Business Value of Ubuntu on Microsoft Azure." January 2025.](#)

Red Hat Enterprise Linux on Azure reduces costs* across the business:

- > \$3.9M business continuity savings
- > \$4.4M datacenter cost savings
- > \$2.7M reallocated FTE savings
- > \$917.9K legacy solution consolidation savings

*Three-year, risk-adjusted

Source: [Forrester Total Economic Impact™ Study Commissioned By Red Hat And Microsoft. "The Total Economic Impact™ Of Red Hat Enterprise Linux On Microsoft Azure." January 2024.](#)

Chapter 6

Industry Considerations

Industry Considerations



AI readiness looks different based the distinct wants and needs of any industry. From business goals and customer expectations to technical requirements and regulatory compliance, becoming AI-ready is as unique as your own business. That said, there are two constants that span every industry and every enterprise: first, being AI-ready is now an imperative. Second, AI should be viewed as a strategic investment, with a focus on long-term feasibility and value to drive sustained growth and resilience.

Understanding these constants, Neudesic is dedicated to helping businesses across industries transform themselves with breakthrough technologies. In this section, we take a look at some key industries where AI is disrupting existing boundaries and creating new opportunities—as well as how Neudesic is helping to identify and capitalize on untapped value. For additional industry stories, check out [Neudesic's Industries](#) page.

AI for utilities

The utility industry is rapidly transforming, with data, security, and infrastructure requirements in an ever-changing state of flux. Utility companies must focus on the reliability, speed, and security of their services, while also streamlining their field employees' workflows and protecting their safety. In this environment, the question quickly becomes, "How do we seamlessly connect production, delivery, and everything in between?"

The answer: AI and emerging technologies designed for the specific needs of utilities. With these solutions, it's possible to unlock new insights in the areas of planning, investment, generation, storage, transmission, delivery, and maintenance. As a result, companies can drive forward with selective modernization of infrastructure, services based on customer usage patterns, predictive versus reactive maintenance, and improved productivity in the field.^{8,9}

At Neudesic, we're committed to achieving remarkable outcomes by modernizing utility technologies. By connecting data, leveraging AI, and enhancing stakeholder experiences, we combine strategy and IT to help our clients succeed.

[Discover how we support the largest utility customers in navigating their most disruptive challenges.](#)

Industry Considerations



AI for retail


Perhaps more than any other, the retail industry relies on making people happy. To grow brand loyalty, customers must feel seen, heard, and satisfied, and to retain top talent, employees must stay energized and engaged. With AI-powered solutions, retailers can deliver exceptional personalized experiences for customers. For employees, they can provide streamlined onboarding processes and enhanced engagement programs.

AI-powered solutions also alleviate the operational inefficiencies that come with managing vast amounts of data across multiple disconnected systems. By streamlining operations, these solutions can simplify workflows, improve communication, and boost overall efficiency.¹⁰

At Neudesic, we're ready to help you use cutting-edge technologies tailored for the retail industry.

Learn how we're dedicated to empowering visionaries and innovators to shape the future of manufacturing.

Industry Considerations



AI for manufacturing

As the manufacturing industry continues to be redefined at intense speed, many companies find themselves uncertain about how to invest in the future. Factories must stay in production, but operational inefficiencies and data siloes drive higher costs, lower productivity, and wasted resources. Moreover, as environmental regulations tighten and consumer expectations shift toward ecofriendly products, companies must find innovative ways to reduce energy consumption, minimize waste, and lower carbon emissions, all while maintaining profitability.

AI can play a pivotal role in helping these companies remain competitive by enabling the factory of the future. Key to this approach are the connected systems and real-time insights that allow information to flow across departments, empowering manufacturers to make faster decisions and seize opportunities for improved efficiency.¹¹

Equipped with AI-powered solutions, manufacturers are positioned to revolutionize traditional methods of production, quality control, predictive maintenance, and process optimization. At Neudesic, we have the experience and expertise required to assist in the revolution.

[Learn how we're dedicated to empowering visionaries and innovators to shape the future of manufacturing.](#)

Industry Considerations



AI for financial services

Today's financial services industry faces increasingly complex economic and regulatory challenges that produce ripple effects across banking, insurance, and capital market enterprises. These organizations are being pressed to accelerate time to value and deliver data-driven customer experiences, all while minimizing risk and protecting personal information.

With AI-powered solutions, companies can leverage data at scale, surfacing intelligent insights that help them design product and service offerings based on customers' preferences and recent actions.

By capitalizing on the latest technology innovations, including machine learning and analytics, these companies are poised to ride the wave of disruption, optimizing risk and becoming smarter, more agile, and more secure.¹²

The financial transformation consultants at Neudesic drive positive change by implementing data and AI technologies tailored for the world's most innovative financial institutions.

[Discover how we can help transform your financial services data ecosystem with AI.](#)

Industry Considerations

AI for healthcare and life sciences

The fast-evolving healthcare and life sciences landscape has led to an explosion of data and rising consumer expectations. For this industry, AI readiness is all about getting—and staying—ahead of the curve when it comes to ensuring the security of sensitive patient data and delivering highly personalized health services.

This means AI systems must be designed and deployed strategically to integrate smoothly into existing infrastructure while also safeguarding against unauthorized access. When AI is implemented in this way, it can drive long-term value by improving decision making; enhancing customer, patient, and provider experiences; boosting productivity; and generating higher revenue, all while protecting users and ensuring system integrity.^{13,14,15}

At Neudesic, we help healthcare clients deliver the access, experiences, and outcomes necessary for success, both today and in the future.

[Learn more about how we manage data at the heart of healthcare.](#)



Industry Considerations



AI for entertainment and media


With user expectations increasing almost daily, the entertainment and media industry runs on speed and personalization. Not only do consumers demand instant access to content, but they also expect it to be immediately relevant, immersive, and engaging. Moreover, from an infrastructure perspective, media workloads are expensive, as they are compute and egress intensive and traditionally rely on on-premises equipment.

For many enterprises, the trick is determining how to move to the cloud for current and future growth, while still realizing value from existing content and ways of consuming media. A purpose-built infrastructure with enhanced AI capabilities enables expansive scaling and reliability—on any screen or device, all at the speed of a click. Likewise, advanced analytics and interactive technologies set the stage for precisely personalized and captivating experiences that extend interaction.

For an up-close look at AI in action in entertainment and media, check out how [MediaKind, a global video technology company, partnered with Microsoft](#) to deliver video streaming at broadcast quality and scale while also reducing downtime, enhancing security, and improving costs.

NOTE: The MediaKind story is one of several Microsoft customer success stories showing how AI readiness can be achieved. Neudesic partners with similar Fortune 500 companies to build customized, AI-enabled capabilities.

Industry Considerations



AI for education

Across the globe, people look to their systems of education, especially higher education, to provide opportunities for personal and cultural growth. A diverse body of traditional and nontraditional students drives exciting possibilities for inclusion, engagement, and innovation—as well as more complex requirements for learning, including hybrid and remote options for instruction. Charged with supporting students across physical and virtual locations without sacrificing personal attention, today's universities urgently need technologies that can keep pace.

To unlock access for all students, regardless of location or situation, universities need to harness their datacenter environments and take advantage of the benefits of the cloud. Moving to a cloud-based infrastructure and AI-powered services enables intelligent tutoring systems and adaptive learning platforms that meet students where they are. For administration, AI-powered tools increase the efficiency of services like enrollment, scheduling, and billing, helping universities better respond to the business demands of education.

For an up-close look at AI in action in education, read how [Western Sydney University partnered with Microsoft](#) to migrate its infrastructure in three months, increasing elasticity and efficiency across locations and saving significant costs along the way.

NOTE: The Western Sydney University story is one of several Microsoft customer success stories showing how AI readiness can be achieved. Neudesic partners with similar Fortune 500 companies to build customized, AI-enabled capabilities.

Chapter 7

Why Linux on Azure

Why Linux on Azure

Enterprises running their Linux workloads on Azure will realize the optimal flexibility, performance, and cost efficiency of a purpose-built infrastructure boosted by AI capabilities. Sounds impressive, right? To back it up, here are six things you might not know about [Linux on Azure](#):

- 1 More than 60% of Microsoft Azure cores used by third parties run Linux-based workloads.
- 2 Hundreds of Azure and other Microsoft services—including Xbox, Microsoft 365, and Azure Kubernetes Service (AKS)—run on Linux.
- 3 Azure customers can access any Linux distributions they prefer. Azure closely collaborates with Canonical, Red Hat, SUSE, Oracle, and the broader Linux communities.
- 4 Azure offers Azure Linux, Microsoft's commercially backed distribution of CBL-Mariner, supported as a container host operating system for AKS and acting as a foundation for many Microsoft services.
- 5 Red Hat Enterprise Linux workloads run optimally on Azure—with access to the latest versions.
- 6 New AI workloads running on Azure, including OpenAI and ChatGPT, primarily run on Linux.

In terms of increasing efficiency, Azure's cloud-based global infrastructure ensures low latency and high availability for Linux applications. It also reduces the burden of maintenance by handling hardware updates, security patches, and other routine tasks. Auto-scaling capabilities ensure Linux applications run varying workloads effectively, without under- or over-provisioning resources. Flexible pricing models, including pay-as-you-go and reserved instances, also optimize costs.

Finally, no modern infrastructure is complete without a strong stance on security. With leading security features like Azure Security Center and Microsoft Defender for Cloud, Azure's approach is uncompromising. Azure delivers strong threat protection and vulnerability management for Linux workloads, while also meeting compliance requirements for highly regulated industries, with certifications such as the GDPR, HIPAA, and ISO standards.



Chapter 8

Why Neudesic

Why neudesic

an IBM Company

At Neudesic, we understand that the future of business lies in harnessing the power of cutting-edge technology, and we're here to help you navigate that journey with confidence. As an IBM company and a premier Microsoft partner, Neudesic is more than just a service provider; we're your dedicated partner in future-proofing your business.

Our ability to deliver rapid results is grounded in our fusion of extraordinary people, deep industry knowledge, and expertise in AI, cloud, data, and more. While many companies talk about the potential of AI, we've already implemented transformative, real-world solutions at some of the top companies in the world.

A culture of passion, discipline, and innovation

Our success begins with our people. Passion, discipline, and innovation are woven into the fabric of Neudesic. We love what we do, and that energy drives us to push the boundaries of what's possible.

With a relentless pursuit of excellence, we deliver consistency and accountability in everything we do. Our approach is rooted in disciplined execution, ensuring that we meet the challenges of today while preparing you for the uncertainties of tomorrow. At the same time, we embrace innovation by approaching every challenge with a fresh perspective, allowing us to create groundbreaking solutions that redefine industries and markets.

A collaborative partnership you can trust

At Neudesic, teamwork isn't just a value; it's how we operate. We trust one another, embrace diverse perspectives, and commit to achieving results together. This collaborative spirit ensures that we can move quickly, adapt to change, and always keep your business at the forefront of innovation. Our integrity is at the core of everything we do, from our transparent communication to our unwavering commitment to the highest standards of ethics.

In a rapidly evolving world, Neudesic is the trusted partner you need to not just keep pace with change, but also lead it. We bring a unique blend of experience, technology, and results-driven commitment to every project, helping your business unlock its full potential.

Let's talk
about how
we can drive
your business
forward,
together.

Next Steps

Take the next step toward AI readiness.

Contact us

Schedule a meeting



References

- ^{1,2}Neudesic. (2025). Technology Vision 2025. Neudesic. <https://www.neudesic.com/2025-tech-vision/>.
- ³Neudesic. (2025, February 13). Outperforming Best in Class AI Safety Approaches with LEC [Video]. LinkedIn. https://www.linkedin.com/posts/neudesic-an-ibm-company_outperforming-best-in-class-ai-safety-approaches-activity-7295471752887902208-4u6n/?utm_source=share&utm_medium=member_desktop&rcm=ACoAAAgLm8IBxj-gVFLUe8FKG76AhtsgJKtVQ1k.
- ⁴PwC. (2017). Sizing the Prize: What's the Real Value of AI for Your Business and How Can You Capitalise? PwC. <https://www.pwc.com/gx/en/issues/artificial-intelligence/publications/artificial-intelligence-study.html>.
- ⁵Sanders, E. (2024, February 8). What is Responsible AI: Understanding Its Importance. Neudesic Blog. <https://www.neudesic.com/blog/understand-responsible-ai/>.
- ⁶Sanders, E. (2024, March 13). Responsible AI in Practice: Accountability, Reliability, Inclusivity, and Fairness. Neudesic Blog. <https://www.neudesic.com/blog/practicing-responsible-ai/>.
- ⁷Sanders, E. (2024, March 8). Responsible AI in Practice: Transparency, Sustainability, Governance, Security, and Privacy. Neudesic Blog. <https://www.neudesic.com/blog/practicing-responsible-ai-2/>.
- ⁸Dvorak, C. (2023, August 18). Utilities and Generative AI: Examples from the Real World. Neudesic Blog. <https://www.neudesic.com/blog/utilities-generative-ai-real-world-examples/>.
- ⁹Dvorak, C. (2023, August 3). Generative AI and the Utilities Industry: A Giant Step Toward Optimization and Customer Satisfaction. Neudesic Blog. <https://www.neudesic.com/blog/generative-ai-utilities-industry-giant-step-toward-optimization-customer-satisfaction/>.
- ¹⁰Wu-Lebron, S. (2024, November 21). Giving Thanks to Improved Labor and Operations Management for Retail with Data and AI. Neudesic Blog. <https://www.neudesic.com/blog/giving-thanks-improved-labor-operations-management-retail-data-ai/>.
- ¹¹Biswas, R. (2025, February 25). How Does Agentic AI and Industrial Internet of Things (IIoT) Impact Manufacturing? Neudesic Blog. <https://www.neudesic.com/blog/agentic-ai-iiot-manufacturing/>.
- ¹²Neudesic. Scale AI for Smarter, Faster Decision-Making in Financial Services and Insurance [On-demand Webinar]. Neudesic. <https://www.neudesic.com/webinar/scale-ai-smarter-faster-decision-making-financial-services-insurance/>.
- ¹³Tetrault, S. (2025, February 27). Health Insurance Claims Automation Prioritizes People First. Neudesic Blog. <https://www.neudesic.com/blog/ai-health-insurance-claims-automation/>.
- ¹⁴Tetrault, S. (2025, February 13) How AI Closes the Gap for Health Insurance Claims Processing Challenges and Patient Care. Neudesic Blog. <https://www.neudesic.com/blog/ai-claims-processing-patient-care-impact/>.
- ¹⁵Thomas, C. (2024, April 11). What are Intelligent Applications? Neudesic Blog. <https://www.neudesic.com/blog/intelligent-applications-definition/>.
- ¹⁶Ravi, K. (2025, March 7). Microsoft PostgreSQL OSS Engine Team: Reflecting on 2024. Azure Database for PostgreSQL Blog. <https://techcommunity.microsoft.com/blog/adforpostgresql/microsoft-postgresql-oss-engine-team-reflecting-on-2024/4388700>.