Intelligence Fabric: How to weave data and business strategies for AI transformation success

Content

Introduction

Artificial intelligence (AI) technologies aren't new, but recent advances – especially the release of OpenAI's ChatGPT in late 2022 – have set off a massive wave of investment and adoption by businesses of all kinds. ChatGPT gained more than 100 million users within just two months, putting consumer-facing generative AI technologies under the spotlight on a global stage for the first time. AI's potential to transform every part of an organisation became undeniable. As a result, any company today without a well-thought-out AI strategy is at risk of falling behind.

But that doesn't mean your organisation should implement an AI programme as quickly as possible – you need a proper data foundation and AI strategy first. Establishing a foundation and strategy requires time and effort, although it's certainly possible to be bold and work fast. What's important is working methodically and following a proven framework for AI adoption. With the right approach, you can use AI to transform your business in significant and beneficial ways over both the short and long term.

This methodical approach starts with a focus on your data. Because accurate, verifiable, clean, complete and well-managed data is the cornerstone for any AI adoption journey. But achieving this is a challenge. In a survey of AI decision makers and project leaders, S&P Global found that the most frequently cited technological inhibitors of AI and machine learning (ML) programmes are:

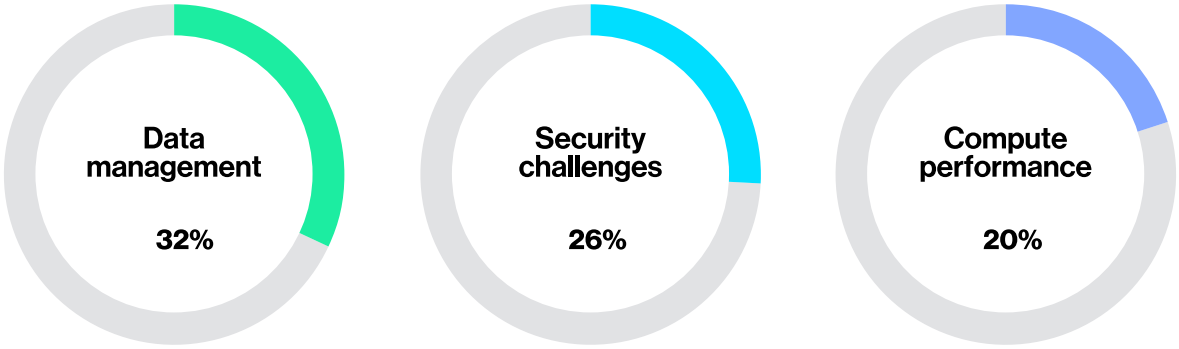

Without good data management, you can't become a data-driven organisation. And research shows that many businesses are struggling with this: only 24% of chief data officers (CDOs) say their organisation is data driven, according to Wavestone/NewVantage Partners. S&P Global's survey found that just 28% of businesses have achieved AI deployments at enterprise scale, even though 69% have one or more AI/ML projects in production.

This paper is designed to help enterprises understand how to weave together the essential data strategies and other necessary business elements for AI in a way that minimises risks and uses repeatable industry patterns for success. It outlines a methodology built on proven practices and assets for growing data and AI maturity.

AI for enterprises: Vast potential for transformation

Despite the challenges around adoption, AI is rapidly becoming a must-have capability. That's because it's more than a type of technology – it's an enabler for many sources of value:

- Insights
- Efficiency
- Innovation
- Improved customer experience, and more

Well-implemented AI and data intelligence tools can transform business operations and performance by helping businesses to make better, more data-driven decisions. Most organisations hold vast amounts of data – far too much for people to analyse and understand without help. With the right data foundation and AI deployment, though, those stores of information can yield valuable insights for understanding customers, targeted marketing, sales forecasts, process automation, long-term planning and more.

AI's ability to quickly analyse large volumes of customer data enables businesses to segment audiences and personalise their products and services, which can drive revenue growth and market expansion. And applying AI analytics to internal data can help to improve the employee experience too, making it easier for organisations to streamline workflows, support skills development, deliver automated support through chatbots and increase job satisfaction and retention.

AI can also be used to automate time-consuming and manual processes, freeing employees to focus on more important and creative tasks, as well as reducing human error, improving efficiency and helping to generate greater value for the business. And manufacturers are increasingly using ML and AI to not only automate industrial processes but to harness predictive analytics to identify potential issues and prevent downtime in production.

These are just a few of the use cases for AI, which can help with so many critical business processes. Generative AI promises an even wider subset of use cases.

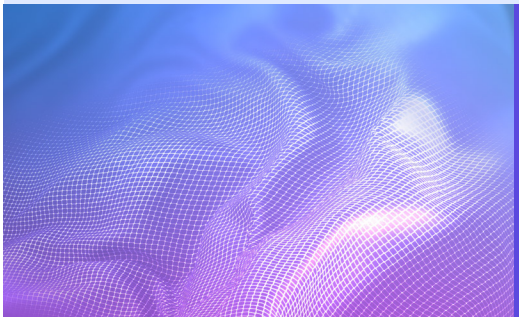

Gaining any of these benefits, however, isn't possible without the right foundation and data for AI. Neglect the fundamentals, and you risk wasting large amounts of money, time and effort, and seeing your business fall behind in an increasingly AI-driven world.

Fast-rising pressure to implement AI

Businesses of all kinds are increasingly recognising the potential value of AI. IDC predicts that global spending on AI will pass $500 billion by 2027, driven in part by the recent groundswell of interest in new generative AI tools. In fact, IDC expects that, by 2025, organisations will be directing more than 40% of their core IT budgets to AI programmes. This rapid shift in focus will lead to IT infrastructure turbulence and an intense focus on data as a critical business asset.

Navigating such rapid and dramatic change requires a smart approach that can mitigate risks while optimising outcomes. But many organisations aren't prepared for the challenge. According to a Gartner survey of CIOs and technology leaders, only 9% of businesses are currently guided by an AI vision statement. Beyond establishing such a vision, organisations will also need to ensure they have AI-ready principles, data and security, Gartner says – and this is something that businesses will need to prioritise over the next 12–24 months.

All of this underscores why it's critical for your organisation to have a solid foundation in place for AI adoption. You need the right approach to data, strategy and a host of other business considerations. The good news is that developing this approach does more than prepare your business for future AI initiatives – it also benefits you in more fundamental ways in the short term, ensuring that everything you do is based on complete and reliable data, best practices and a strong network of business and technology partners.

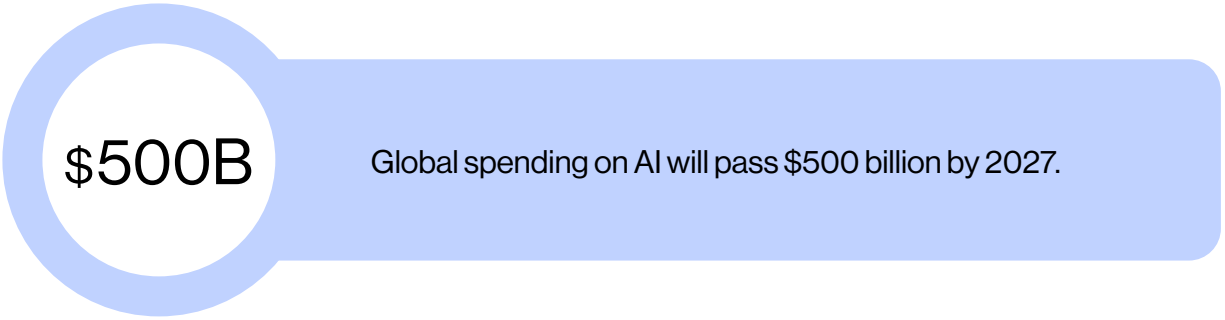

The essential threads for any AI programme

Making AI work effectively for your business isn't a quick or simple proposition. It requires you to build a foundation and strategy that supports every part of your organisation, ensuring that you have accurate and verifiable data, mature data practices and ways to support your users' needs. And every part of your foundation must include strong security protections and robust governance.

Nor is an AI programme something you can implement without support. Very few organisations have all the varied capabilities and resources needed to bring an AI project from conception to production on their own, and doing so isn't necessary anyway. Proven tools and technologies are already available through leading hyperscalers such as Amazon Web Services (AWS), Google Cloud Platform and Microsoft Azure, as well as through many reliable technology partners. Successful AI deployments depend on smart business networks, strong collaboration and trusted partnerships.

Creating an effective AI framework requires several essential threads. Weave these together properly, and you'll have a sturdy fabric on which to model your AI ambitions. Leave out any of these essential threads, though, and your ambitions are at risk of unravelling.

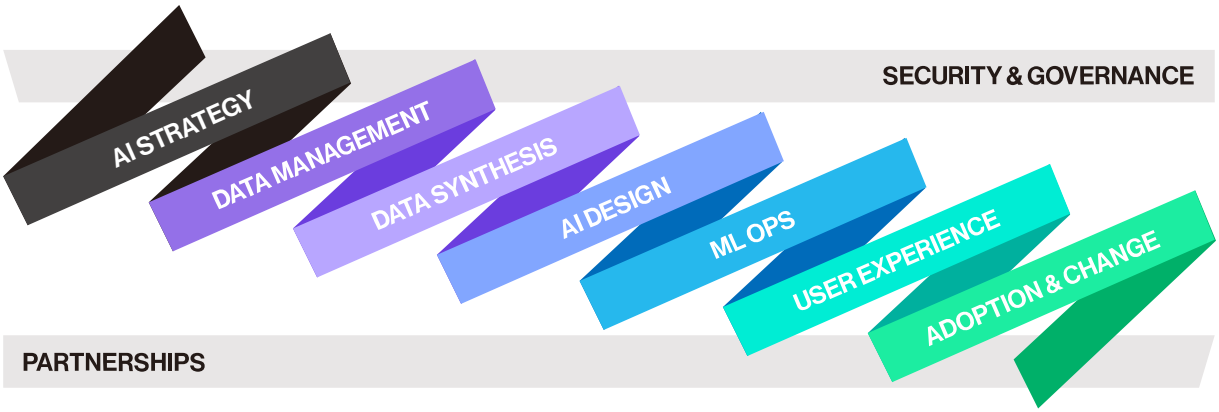

Business outcomes and AI strategy

Before you embark on an AI programme, it's wise to seek out a knowledgeable and experienced partner that can assess your organisation's readiness for AI adoption and can provide advisory support and strategy guidance during the design process. Because any successful journey requires knowing not only where you want to go, but where you're starting from. This means understanding your current data and analytics capabilities and then identifying how to grow and improve these capabilities to achieve your AI goals. Doing this can be difficult if your current AI foundation is immature.

With experienced advisory support, you're better able to pinpoint technology and skills gaps, choose the right solutions and build momentum towards greater data and analytics maturity. This will enable you to establish a vision for your data-driven and AI-powered future and then begin aligning your business outcomes to those goals.

Data management

There's no AI without getting your data in order. So your first step is to assess the status of your data and analytics capabilities and ensure that your organisation's data is trustworthy. If you want to become a modern, data-driven business, your decisions and actions will only be as good as the data you've woven them from. The most advanced and intelligent IT systems won't produce good results if you feed them inaccurate or unreliable data.

Data management means taking steps to ensure your data foundation is as strong as possible. It means ensuring that your data is based on reliable, verifiable measurements and calculations. That it is checked for error, cleansed and accurately recorded. That it is well-stored, secure and properly integrated into the systems you use. And, most of all, that it is in the cloud – because the cloud is essential for making AI work.

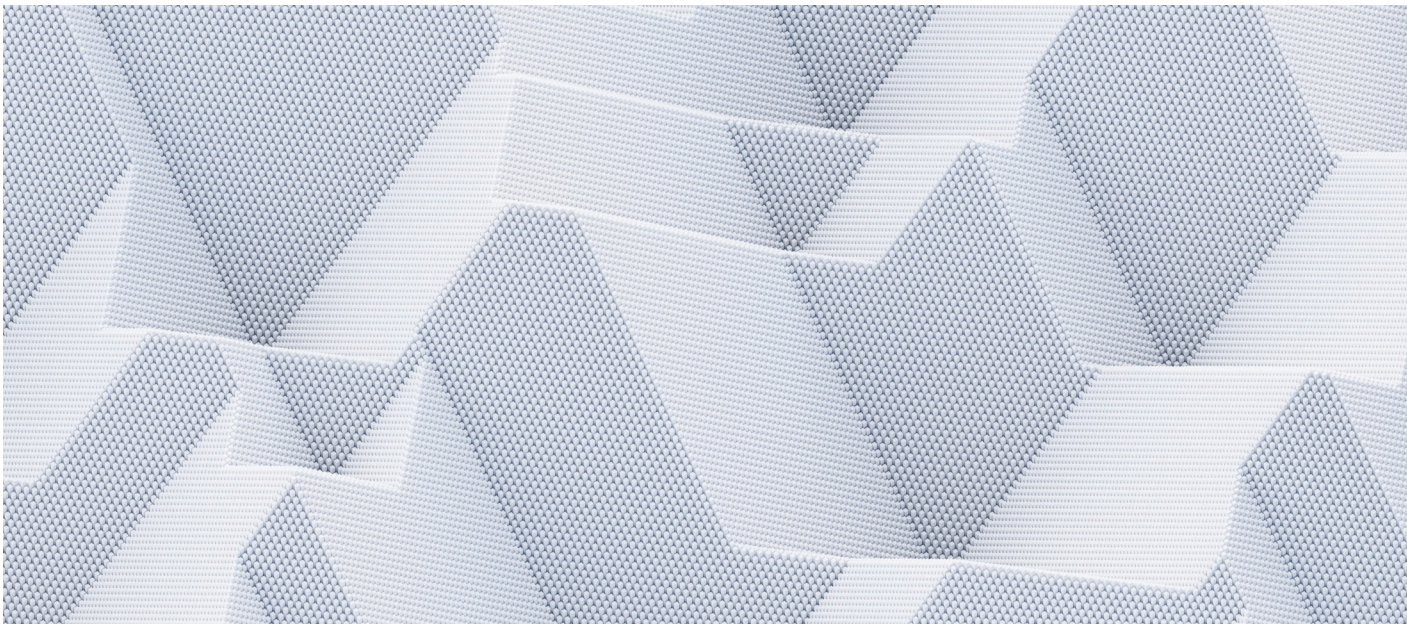

Data synthesis

Clean, accurate, secure and cloud-based data alone is not enough to build an AI implementation on. You also need to make sure that all of your data can be managed and analysed from a central location, and that information from various sources is recorded in compatible units and formats. For instance, this might require you to transform some types of information to meet common standards. And you'll need to eliminate data silos across your organisation by unifying disparate stores of data – systems for customer relationship management, enterprise resource planning, project management, human resources, accounting, sales and marketing and so on.

Data synthesis also involves examining all of your available data to determine if there are any gaps. These might include areas where you will need to add new information from new data sources. The need for feature engineering is another possibility – this involves manipulating or extracting information from existing raw data to create new kinds of variables or datasets that will be digestible by the ML and AI models you plan to use. Your data synthesis requirements will depend on what you hope to achieve using AI.

AI design

The next step is to actually create and train an AI model that's tailored to your business goals. With so many models now readily available from the main hyperscalers, you're no longer required to start this process from scratch. However, you do need to carefully assess the available models to make sure they're right for your needs. This includes making sure they're appropriate for the types of data your organisation has, the platforms you currently use, the types of challenges you are applying AI to and your delivery requirements. For example, how quickly do you need to train your AI model, and how much accuracy/reliability are you willing to compromise on to achieve results quickly?

Your use cases can vary considerably. For example, an AI model designed to help IT integrate a fragmented data system will look different from one that's designed to help the chief financial officer of a global corporation speed up and automate processes for financial forecasting, or from one that's aimed at updating marketing campaign processes to improve ROI. Your way forward will depend on what are currently your biggest pain points and most important business objectives. Is operational efficiency the key goal or is the top priority to increase revenues and decrease costs? You want to ensure that your AI solution is tailor-made to deliver maximum value and relevance.

MLOps

MLOps applies DevOps principles to machine learning, looking for ways to standardise practices and processes – and automating when possible – to monitor, manage, verify and accelerate AI model performance. This is important because, even if you fine-tune an AI model to produce the results you're looking for, conditions will change. If you don't adapt your model to evolving conditions, you risk deteriorating outcomes that could lead to bad business decisions.

During the COVID-19 pandemic, for example, many advanced sales forecasting models failed because market conditions changed drastically overnight. Without rapid adjustments, these models no longer matched reality.

So it's important to ensure that your organisation has an automated system that can monitor AI model performance, verify results and alert you to changing conditions that might require retraining with new data. Not long ago, doing this required weeks or months of painstaking manual effort. Today, all of the major hyperscalers offer tools that can help with this – so finding the right partner to support you will be critical.

Another area to focus on is the emerging discipline of LLMOps (large language model operations). This applies DevOps principles to applications involving generative AI, helping to improve the efficiency of GenAI development and training processes, reducing risk and supporting scalability.

User experience

Building and deploying an accurate AI model isn't enough to bring value to your organisation, though. You also need to ensure that the model's output is usable and useful in real life. This involves thinking about the kind of apps that you will build on top of AI, and how effectively they integrate AI capabilities into a user-centred experience – for example, can results be delivered in user-friendly and digestible formats such as Power BI or Tableau? Without such accessibility, the raw numbers produced by an AI model will be difficult, if not impossible, for employees to draw insights from.

Consider what Microsoft has done with Microsoft Copilot, which was designed to work seamlessly across the entire Microsoft 365 (M365) stack, as well as in Windows 11 and web browsers. Although Microsoft Copilot can work as an app on its own, Microsoft says it will also "reveal itself" with a single click in any M365 application when users need it.

Whatever AI model you build on, it's also important that users understand how the model works so they can have confidence in the output and know how to seek additional insights when needed for effective decision making.

Adoption and change management

This leads to the last – and possibly most important – element needed for a successful AI deployment: adoption and change management. Even with the most complete data inputs and best-designed AI models, your organisation won't see the benefits if your people don't adopt and use the technology as it was designed to be used. To see a return on your considerable investment of money, time and effort, it's essential to win buy-in from employees.

How do you do this? By establishing and following an adoption framework that clearly spells out roles and processes for everyone involved. This means answering employees' questions, demystifying the technology and addressing their concerns about AI – for example, assuring them that it won't put them out of a job, but will free them from tedious, repetitive and time-consuming manual tasks so they can focus on more interesting work. And it means providing people with the skills and knowledge they need to get the most out of your AI deployment and build a truly data-driven work culture.

Effective adoption and change management also means encouraging people to actively participate in your AI strategy. This should include, for instance, establishing a cross-functional AI centre of excellence to drive adoption across every area of the business. It could also include identifying people who can act as promoters and champions for AI to build enthusiasm about the technology's potential. And people across the organisation should all be encouraged to look for and suggest new use cases that could help with their tasks and responsibilities.

One more note: if you're lacking capabilities in any of the areas above, it's wise to seek out a knowledgeable and experienced partner that can assess your organisation's readiness for AI adoption and can provide advisory support during the design process. Because any successful journey requires knowing not only where you want to go, but where you're starting from. This means understanding your current data and analytics capabilities and then identifying how to grow and improve these to achieve your AI goals. Doing this can be difficult if your current AI foundation is immature.

With experienced advisory support, you're better able to pinpoint technology and skills gaps, choose the right solutions and build momentum towards greater data and analytics maturity. This will enable you to establish a vision for your data-driven and AI-powered future and then begin aligning your business outcomes to those goals.

Threading the AI needle: The benefits of getting it right

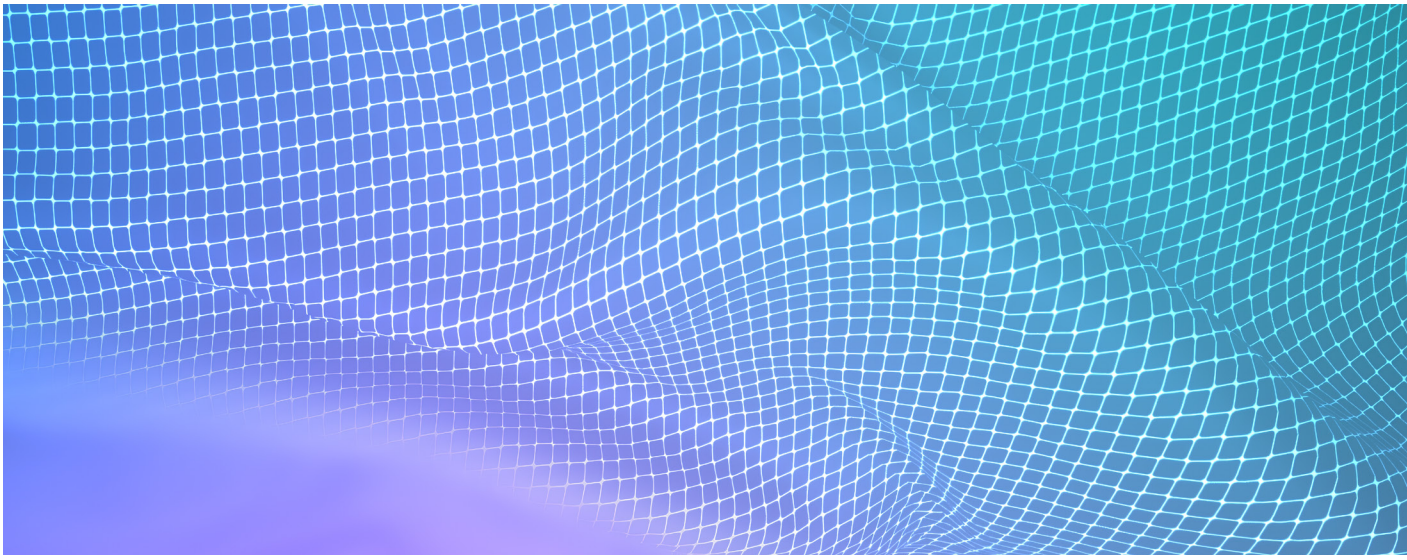

Although you need all of the essential threads described in this paper to make the most of AI's capabilities, improving each one individually brings rewards of its own. Every step you take to progress on the AI maturity curve helps to strengthen and modernise your business in many other ways.

For example, strong data management practices ensure that the information you use to guide business decisions is reliable and verifiable. And synthesising that data so that it's standardised, complete and accessible from a centralised interface makes it easier to keep your data infrastructure up to date, secure and ready for modernisation in the cloud. Similarly, paying attention to AI design, MLOps, user experience and change management in line with your larger business goals helps you to identify, clarify and focus on your strategic objectives.

Consider one company that spent almost seven years building the foundation necessary for an ambitious AI project. Although that organisation is only just now getting ready to go into production, the updates that it has made to its data, processes and practices have helped it to become more data-driven, generate better insights from the data it has and even implement self-service tools that are helping employees to do their jobs better and more efficiently.

There's also the possibility that, after laying all of the above groundwork for an AI project, you conclude that you've overestimated the use case or potential return on investment. If you decide at that point to stop pursuing the project, you've still modernised your overall approach to data – and you have saved your company the financial (and potential reputational) expense of a failed or underwhelming AI deployment.

With a modern approach to data, you can improve the efficiency of your internal operations, make better business decisions, save costs and reduce errors. And all of these things can contribute to better service for your organisation's customers – whether that's in the form of faster service, more accurate answers, new self-service capabilities, greater personalisation or lower costs.

Finally, weaving together the essential threads for AI adoption also helps to strengthen your organisation's security and governance. You're better able to defend decisions and actions when they're well-grounded in reliable, accurate and up-to-date data, and better able to prove compliance with high-stakes regulations such as the EU's General Data Protection Regulation (GDPR).

To sum up, then: by following a methodical, multi-step roadmap to AI adoption that improves your data, processes and practices, your business can gain many benefits. These include greater operational efficiency, stronger support for innovation, more business agility, cost savings through standardisation and automation, improved customer service and the potential for increased revenues through new or more modern products and services.

The wrong approach: Big risks and a cascade of impacts

Without putting in the time and effort needed to build the right foundation and strategy for AI, organisations can lose millions through failed projects with no prospects for a return on investment. Such failures can also bring many knock-on effects: lost business opportunities, damaged reputations and reduced competitiveness, to name a few.

AI projects that manage to go live despite not having the proper building blocks in place can be just as dangerous to an organisation. That's because a poor implementation with disappointing results can create a lack of trust in AI technology. It can also discourage leaders from pursuing future projects, which can result in the business falling even further behind competitors that have greater AI capabilities. When this happens, organisations also find it harder to attract – and retain – skilled and experienced people, leaving the business at an even greater disadvantage over time.

A badly designed AI project that neglects data safeguards can:

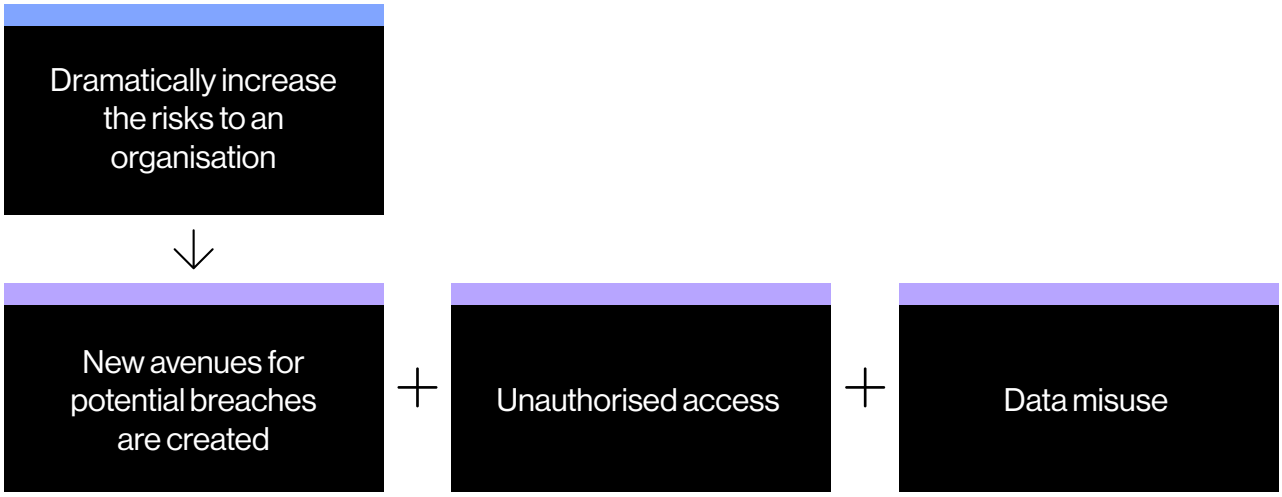

This is why knowledgeable guidance from a trusted adviser is such a critical part of any AI deployment. An experienced partner can provide vital support and help to steer an organisation away from short-term thinking, unrealistic expectations and rushed implementations. You want a partner that – after assessing your AI plans and current level of readiness – won't be afraid to tell you, "No, don't do this. Because it won't deliver the results and revenues you're hoping for."

What kinds of organisations can benefit?

Ultimately, any kind of business can benefit from a well-implemented AI project built on the right foundation. But because AI outcomes are powered by data, the greatest benefits will typically go to organisations with the largest volumes of data – especially data related to their customers or key business processes.

This means, for example, the financial services industry. With all of the data that banks, insurance companies and other businesses in this sector have about their customers and their investments, this sector has great potential for using AI to innovate new financial products and services, personalise existing offerings and provide deeper insights for financial health.

The retail sector is another one that could see outsized benefits from AI, which can be applied to everything from personalised marketing and automated customer service to secure payment processing, price optimisation and sales forecasting. Generative AI in particular is already producing great results for retailers by creating product descriptions for online listings and catalogues, which are time- and labour-intensive to develop manually. Meanwhile, the public sector – which is often understaffed and underfunded – could apply AI to many processes that are currently handled manually, helping to deliver answers and services to citizens faster and more efficiently.

In manufacturing, AI can be used to identify inefficient processes, automate manual tasks and monitor equipment for potential maintenance issues. Such applications can help to reduce repair costs and downtime, speed up production and reduce operating expenses. Manufacturers and other industrial businesses can also use generative AI to automatically update maintenance plans based on past operational data and experiences, or to help design new parts or products according to desired characteristics such as energy efficiency.

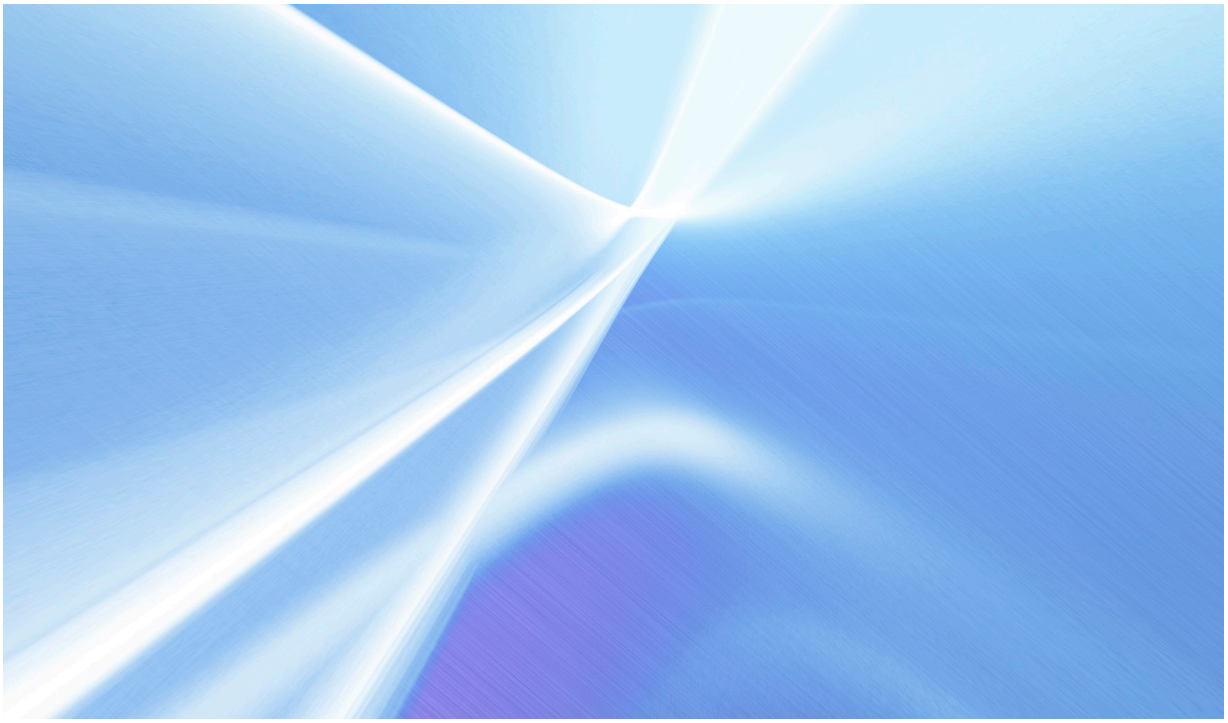

Consider a few of these real-life case studies, made possible by SoftwareOne:

Manufacturing

A fast-growing manufacturing business needed to consolidate its data so it could provide better services to customers and improve efficiency across multiple legacy systems. By migrating to the cloud and building a data lake for centralised data management, it was able to implement more data-driven sales and marketing activities to win more business. It's also working to use equipment data to implement predictive maintenance.

Pharma

A pharmaceutical manufacturer was looking to automate activities to reduce long turnaround times and reduce errors in its business forecasts. It consolidated its data in the cloud and trained an AI model, then made those outputs accessible to a variety of business applications. Using this solution, it is now able to automate its forecasting, better monitor product demand, gain new data insights and it has reduced forecasting errors to less than 5% across all product lines.

Logistics

An aviation logistics organisation wanted faster, more responsive analytics to improve the efficiency of the applications it used to manage fuel delivery to hundreds of flights per day. By consolidating its data in a cloud-based data warehouse, it has reduced the time spent on data analysis by 80% and doubled the accuracy of its predictions for asset purchases and resource management. The company can now also access on-demand statistics through real-time dashboards during conversations with prospects and customers, which has helped it to win new business.

Oil and gas

An oil and gas company lacked a centralised system for measuring the performance of its various business units and spent hours manually extracting data from emails and other sources. It wanted an automated way to extract information from files, so it set up a centralised data warehouse and an analytical business intelligence model. Using its new cloud-based systems and ML, the business can now view essential KPIs at a glance, receive automated alerts when KPIs drop below certain thresholds and automatically extract data for reporting as needed.

Hospitality

A restaurant business needed to manage hundreds of gigabytes of structured data every day, but lacked strong data management practices and the ability to personalise data for customer and marketing campaigns. By creating an enterprise data warehouse, it now has a single source of truth that users can access for customised reports and real-time sales data. The organisation is also using this data to develop AI solutions aimed at improving both internal operations and its overall restaurant business.

Closing thoughts

A word about technology

AI can't exist without technology, but technology alone doesn't enable a successful AI deployment. Nor is it the most demanding part of implementing AI. All of the other essential elements described earlier in this paper – data management, data synthesis, AI design, MLOps, user experience, adoption and change management and knowledgeable guidance – must be in place, with mature practices in each, if you want to optimise the outcomes of an AI project. Without this framework, even the most advanced and powerful AI tools will produce poor results and could damage your business.

And thanks to the many powerful and well-tested AI tools and platforms that have been developed by the major hyperscalers, you probably won't have to build your AI technology foundation from scratch. You'll undoubtedly have to do some customisation and train the AI models you select according to your specific business needs, but that work can be done only after you've first built the proper data foundation, clearly articulated your business goals and considered the needs of who will be using your AI solution.

Unfortunately, far too many organisations today still believe that AI is a plug-and-play technology – that all you need to do is choose the right AI tool, turn it on and then begin reaping the rewards.

That's simply not true.

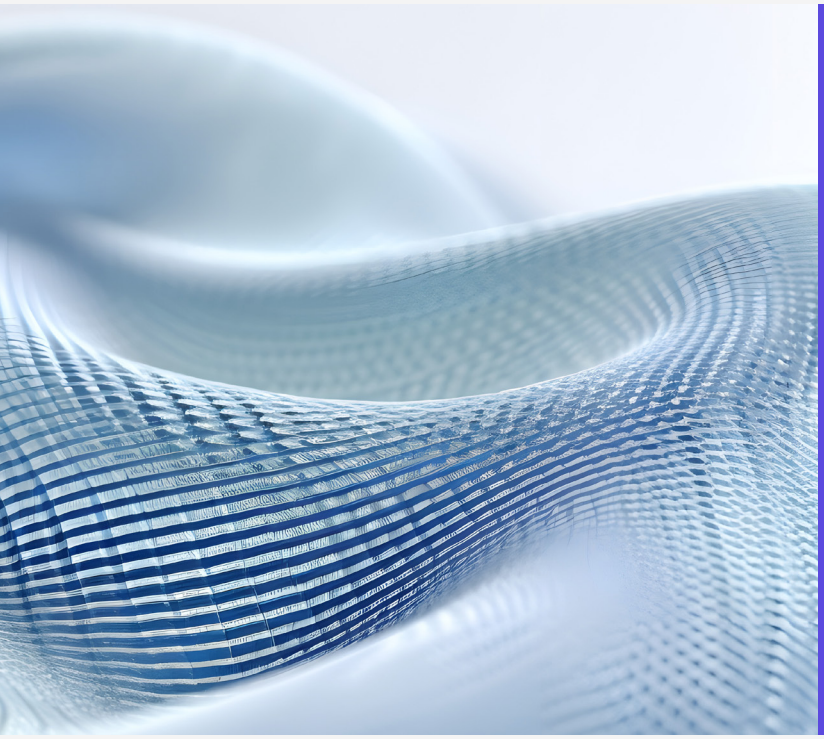

AI technology itself is enterprise ready. And it has proven its capabilities across many different use cases and industries for some time now. It also continues to evolve quickly, which means you could see many additional benefits in the years ahead – if you follow the path for deployment we've outlined here. The foundational elements in this framework must all be present to make AI work. Omit any one of them, or take shortcuts, and even the best-planned AI initiative could unravel.

Looking to the future

AI will touch everything, from automated customer experiences to personalised marketing to supply chain management to fraud detection and more. So it's critical to build the right foundation first and be prepared for continuous evolution. That's why having a solid framework with built-in security and governance is important. As the technology continues to change and new applications keep emerging, the businesses most likely to remain competitive and thrive will be those with firmly established data practices, well-supported workforces and knowledgeable, reliable partners.

The AI era is here, and businesses can't afford to ignore growing adoption pressures. Organisations that don't embrace AI will find themselves struggling to keep up with competitors that will harness the technology for better business insights and decision making. By continuing to rely on manual processes rather than optimising and automating them, these businesses will experience higher costs and inefficiencies. And they'll be unable to derive value out of the large and growing amounts of data they must manage, putting them at a disadvantage in an era of big data analytics.

But businesses also face many risks if they pursue AI but build their AI programmes on a poor foundation. Without strong data practices and a comprehensive framework for adoption, they could find themselves investing considerable time, effort and money into an AI project that doesn't deliver value. Worse still, a badly designed AI programme could increase reputational and security risks, raise legal and compliance concerns, amplify costs and inefficiencies and produce unreliable outcomes that compromise business decisions.

That's why the methodology outlined in this paper is so important.

SoftwareOne Intelligence Fabric

With our extensive experience across so many technologies and services that are fundamental to AI deployment, SoftwareOne has developed a comprehensive framework to help you get the most out of your AI initiatives: Intelligence Fabric.

It's a framework that provides security, governance and partnerships, along with all of the essential factors we described earlier for AI deployment: advisory services and support for improved data management, data synthesis, AI design, MLOps, user experience, and adoption and change management.

Intelligence Fabric is a suite of accelerators, practices and assets built on two advisory pillars: one, data foundations and, two, analytics and AI. This two-pronged approach is aimed at helping you to optimise every practice and process required for AI project success. It's a comprehensive, compliant, secure, efficient, scalable and future-proof methodology that uses proven, repeatable industry patterns and is designed to reduce risk. And it's built on our years of experience as a vendor-neutral partner to all of the major hyperscalers: AWS, Google Cloud Platform and Microsoft Azure.

As a software licensing expert, we're a one-stop partner that delivers full lifecycle support across a broad range of services, from cloud and software licensing to infrastructure services, application services, FinOps and managed services. **We have significant capabilities, with 250 data and AI experts and more than 230 projects delivered since 2021** across seven delivery hubs. We also have partnerships with more than 30 leading data and AI technology providers, including early access programmes to the latest technological advancements. Plus, our strategic partnerships with the three major hyperscalers can provide you with additional support, including a wide range of funding programmes.

Our Data Foundations advisory guides you through the Envision, Discovery and Solution phases to help you navigate the complexities of transforming into a data-driven business. We help you to understand the potential of your data, align it with your business goals and develop a practical roadmap towards transformation. This enables you to enhance decision making, improve operational efficiency, optimise costs, and mitigate risk and compliance issues, paving the way to a well-designed AI strategy.

On the Analytics and AI advisory side, we begin by working with you to explore the potential for AI in your organisation and to identify which areas could benefit most from the technology. We then help you to tailor a roadmap that aligns with your business goals, delivers an agile and scalable solution and mitigates risk through robust governance, security and compliance. The outcomes of this include improved operational efficiency, new revenue opportunities, cost savings, greater innovation, streamlined processes and faster times to market.

We developed SoftwareOne Intelligence Fabric to help you overcome the challenges of building a data-driven business that's ready to make the most of AI's opportunities and potential.

Begin weaving your AI strategy

Ready to pull together the essential data and analytics threads needed for business transformation in the AI era?

Reach out to our experts to schedule an AI strategy session for you and your team.

CONTACT US TODAY

Find out more at
www.softwareone.com/ai

SoftwareOne AG | Headquarters
T. +41 44 832 41 69
E. info@softwareone.com