Your route to AI success

A practical guide to getting ahead of the game

orange™ Business

Introduction

Artificial intelligence is dominating business. Or at least, it dominates business conversations. Everyone wants to understand what AI is, what it means to them, and how they can deploy it to maintain a competitive edge. Every time we talk to customers, AI comes up. Even when it isn't the main focus, it is the elephant in the room.

Of course, AI has been around for years. Why is everyone focused on it now? Because the planets have aligned; we have the massive data sets and compute power to use it effectively, and the emergence of Gen AI tools has brought the barriers to entry crashing down.

So, what's missing? Specific use cases. The more specific the business need, the greater the likelihood of success. Saying you will deploy AI to improve productivity is too general: the objective needs to be hyper-focused.

Once you have your use cases, you must be clear on designing and building your approach. Businesses need to be aware of certain areas, put aside their assumptions, and invest the time and energy to get the implementations right.

And they do need to be successful. While AI's potential impact is undoubted, there are still many concerns surrounding its use; poorly managed AI applications will hinder the technology's acceptance within the business.

What, then, do enterprises need to consider when planning AI rollouts? There are five core areas:

1. **The business case**
2. **Tackling data quality**
3. **Using LLMs properly**
4. **Security**
5. **Network**

To help, we've gathered insights from across Orange Business to break down these areas, identify what businesses need to be aware of, and highlight how they can incorporate these learnings into their AI rollouts.

Kristof Symons
CEO International, Orange Business

Building the business case

AI presents businesses with significant opportunities to change how quickly they can work and how they operate as a company. It can be a foundational building block for growth and continued success, from new revenue streams to partnerships.

Enterprises recognize this. In generative AI (GenAI) alone, Gartner predicts that by 2026, more than 80% of enterprises will have used AI application programming interfaces (APIs) or models and/or deployed AI-enabled applications in production environments, up from less than 5% in 2023[1].

In other words, whether businesses use AI is no longer a question. Every company that wants to survive never mind thrive, must use AI. Therefore, it must understand how to deploy it successfully within its organization.

And it is easier than ever to do so. The barriers to entry have fallen, with more tools and services than ever before, from a vast array of providers, from Big Tech to open source. One trap many fall into is thinking that, as prices fall, they should wait for the costs to keep coming down. Yet, when the ability to train models and extract value is tied so closely to deploying the technology, it is more important to use it quickly than save a bit on capital expenditure.

Building the business case

It is also important to know that AI has lost its edge as a competitive differentiator as it becomes more readily available. Using AI is not enough when everyone has access to the same tools. How you use it becomes the critical enabler of value.

To do that, there are several areas that you need to think about as you build your business case. They include:

1 How will your deployment scale?

AI in the lab or in pilot differs from AI in the wild, so you must be clear on how to scale sustainably. This becomes harder when tools are over-engineered, particularly if it's to be unique. Ask yourself whether there is more value to your business in having a truly unique AI deployment or one that's easy to scale.

2 How close do you want your tooling to your clients?

There are already examples of AI engaging with customers directly – chatbots being the most obvious example – but that should only happen in very specific, guarded deployments. For instance, one bad experience, such as a hallucination, could damage your customer relationships and brand reputation. Using chatbots for triaging customer queries makes sense; more complex, open-ended escalations should be left to human agents.

3 Have you checked your tools' bias?

We all have biases. Everything we develop, including the AI tools you deploy, will have biases. To minimize the potential negative impact, you need to be selective in your AIs—they are not all the same—and use the ones most relevant to your use case.

4 How will it integrate with your other technology?

Depending on your existing tech stack, adding AI could be the equivalent of having rocket science alongside rocks and sticks. So, as part of your deployment plan, you need to know how it will integrate; otherwise, any savings, efficiencies, or new revenue opportunities will disappear.

5 What are the privacy and security implications?

You must know the legal implications of everything you do. Where is the data coming from that's training your AI? Is it infringing copyright laws? Is it exposing your systems to cyberattacks? Are you breaching privacy regulations by using certain data types to train and inform the AI? You have to have answers to all these questions before you begin.

Tackling data quality

For all its potential use cases, there is one fundamental rule that underpins all AI deployments: if you put crap in, you'll get crap out.

What determines whether crap gets in? Data, and more specifically, the quality of data. Without data, there can be no AI; the latter needs the former to learn, whether it's informing the training of new models or helping existing deployments evolve and adapt in the real world. Conversely, data needs AI to be valuable and to process, analyze, and extract insights from the vast amounts of information being created today.

Data quality is at the heart of AI success and always has been. What's changed is the scale of adoption is evolving. Go back a couple of years, and only the most well-resourced enterprises could deploy AI tools; now, with the advent of GenAI, everyone has access.

That means the tools you have and the algorithms you're deploying are no longer differentiators, but your data can be.

You need to ensure you have quality data to minimize risks. Privacy, security, bias, rubbish—if you have an underlying issue with your data, GenAI will amplify it. It is also vital that these issues are addressed now: in the future, we may well see AIs trained on data produced by other AIs so that any quality issues will become deeply embedded.

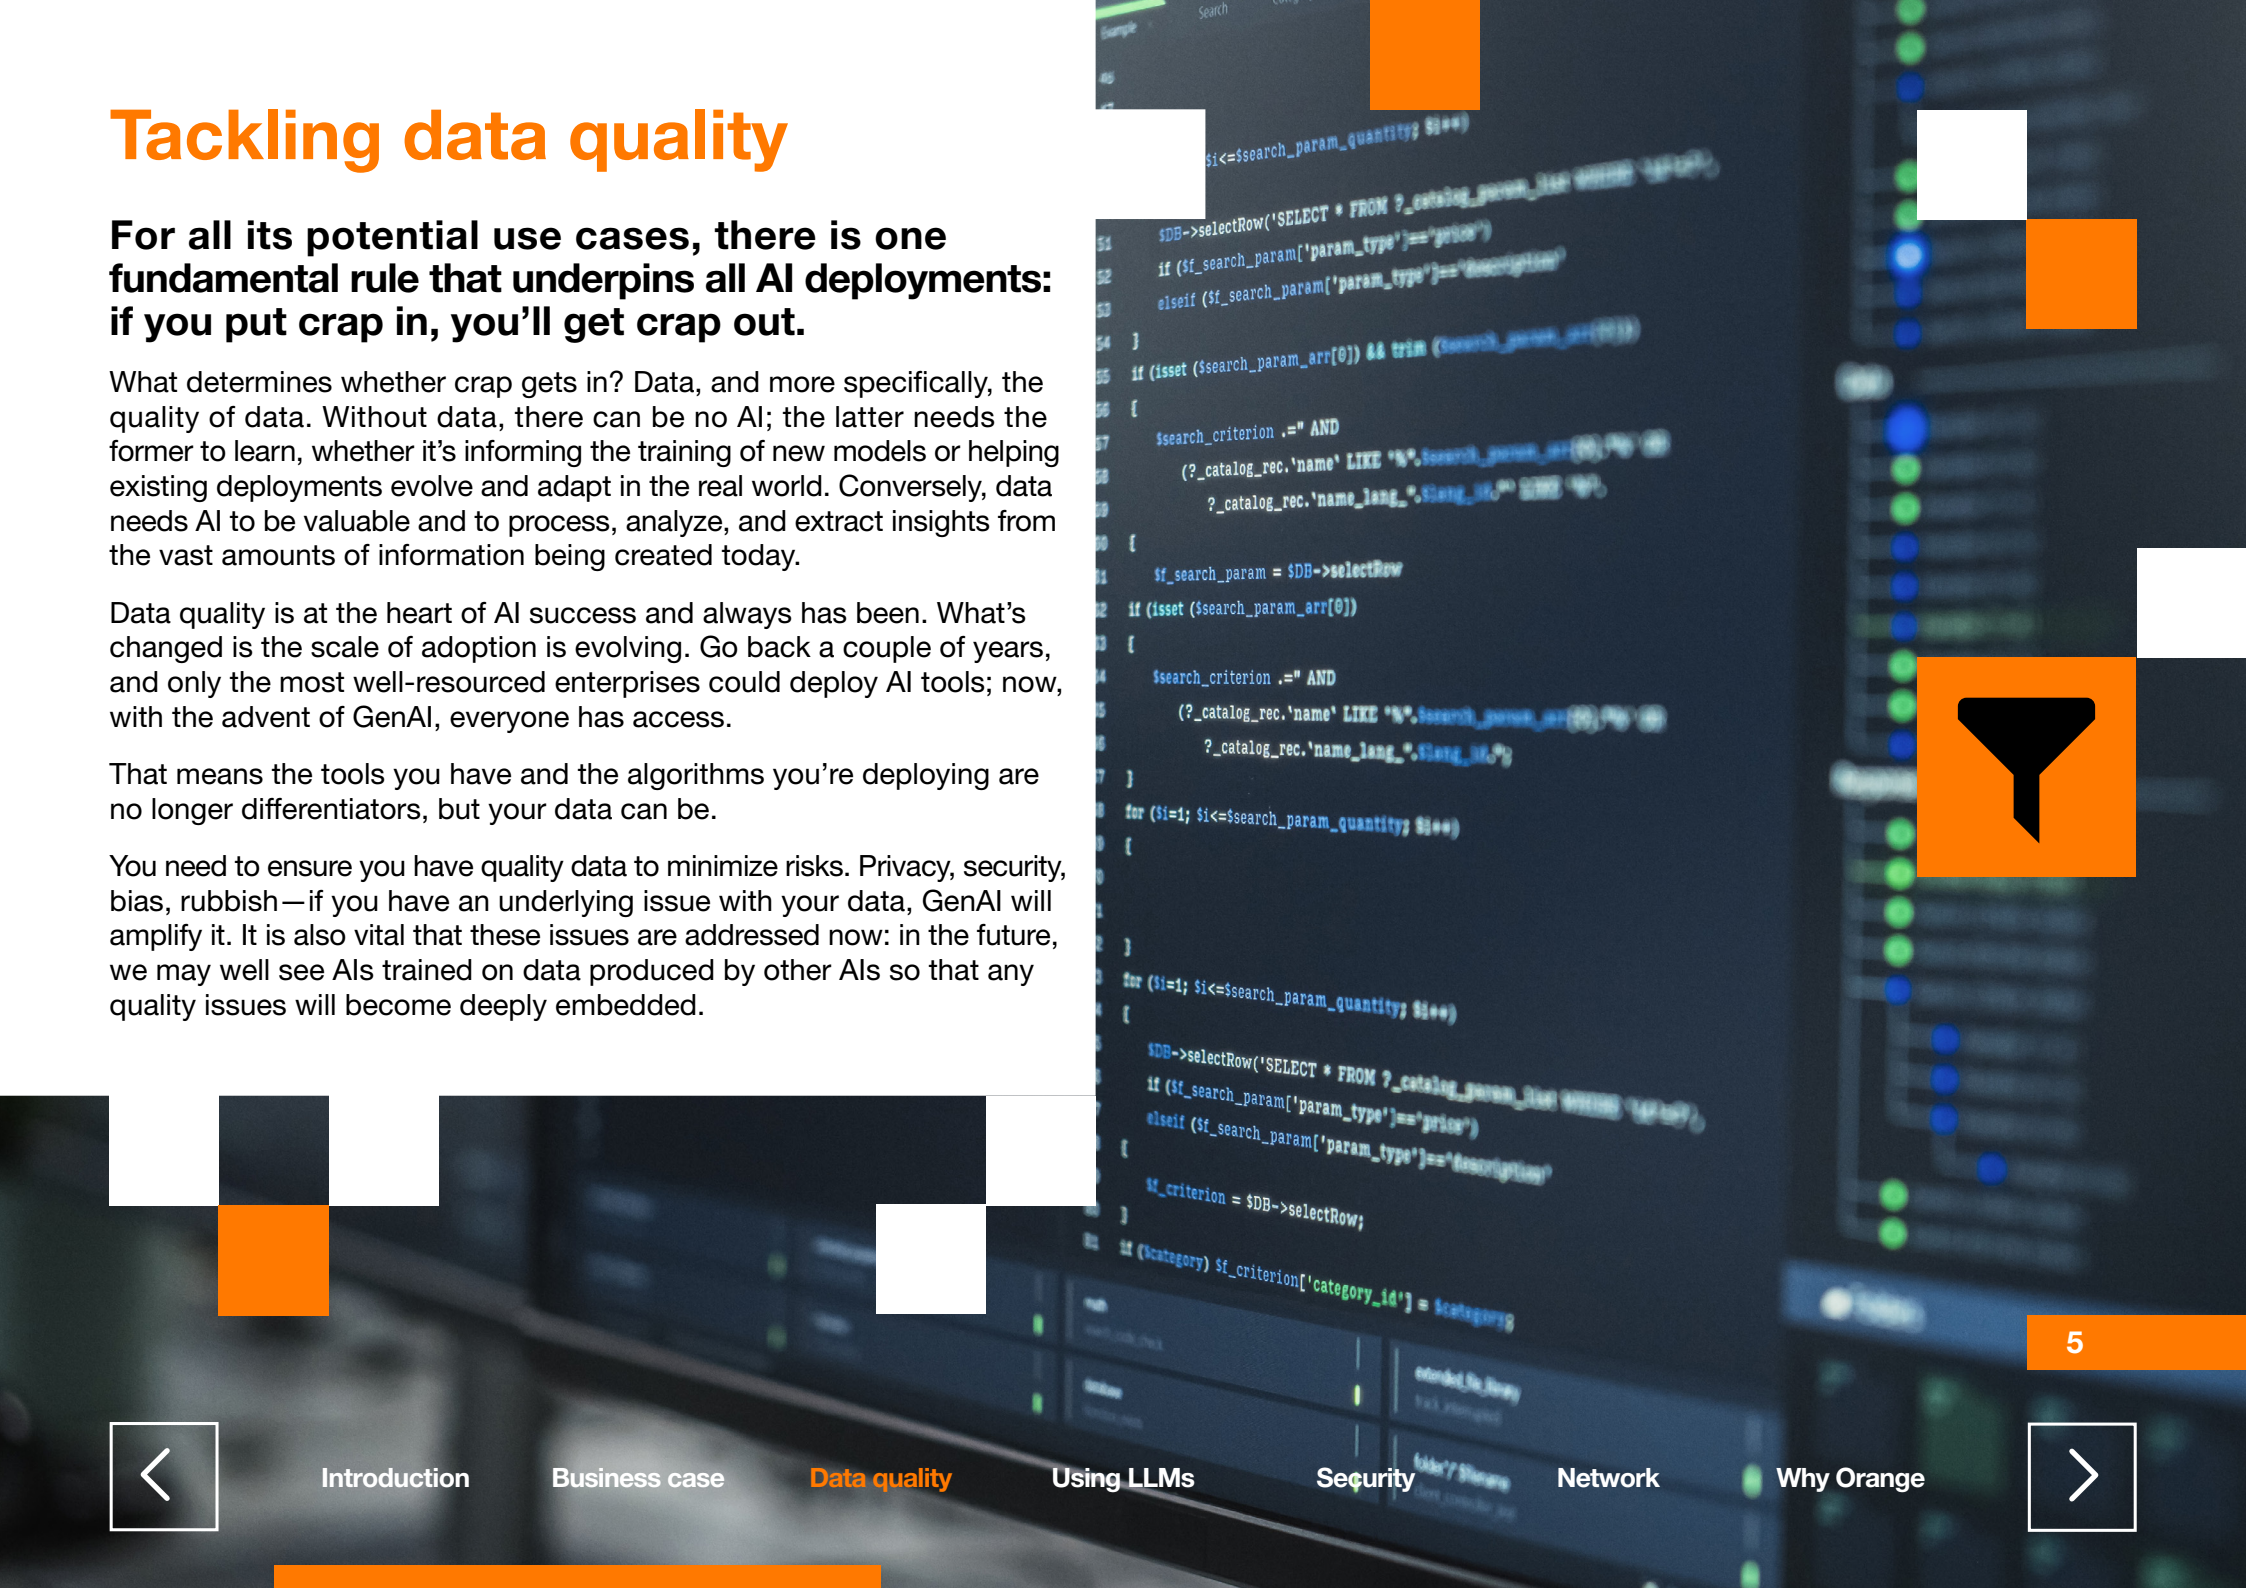

Tackling data quality

Yet how do enterprises ensure their data has the necessary quality when there is so much of it? By doing three things:

1 Take a value-driven approach

Attempting to improve the quality of all your data is an invitation to paralysis; nothing will get done. Instead, start by identifying where AI will bring value. These will help you assess the use cases that will most likely be a success and allow you to establish best practice data foundations as you go. If you have a specific use case in mind, you can identify the data required and focus on ensuring that it meets the necessary quality standards. That way, you can trust the data going in and, therefore, the outputs at the other end.

2 Scale incrementally

By taking a value-driven approach, you scale incrementally. With each step, you ensure that the data involved meets your requirements. This helps you build a body of use cases with demonstrable results and allows you to remain in control of all required governance. Rather than trying to comply with every possible regulation, you will only spend time where specific laws apply to your AI.

3 Learn and iterate

In a field moving as quickly as AI, success depends on your ability to learn and iterate at the same speed. As use cases prove successful, their processes can inform future deployments, particularly the refining and improvement of data sets. Learnings can be incorporated into the next stages while architecture, governance, security, and strategy remain controlled and driven by use cases.

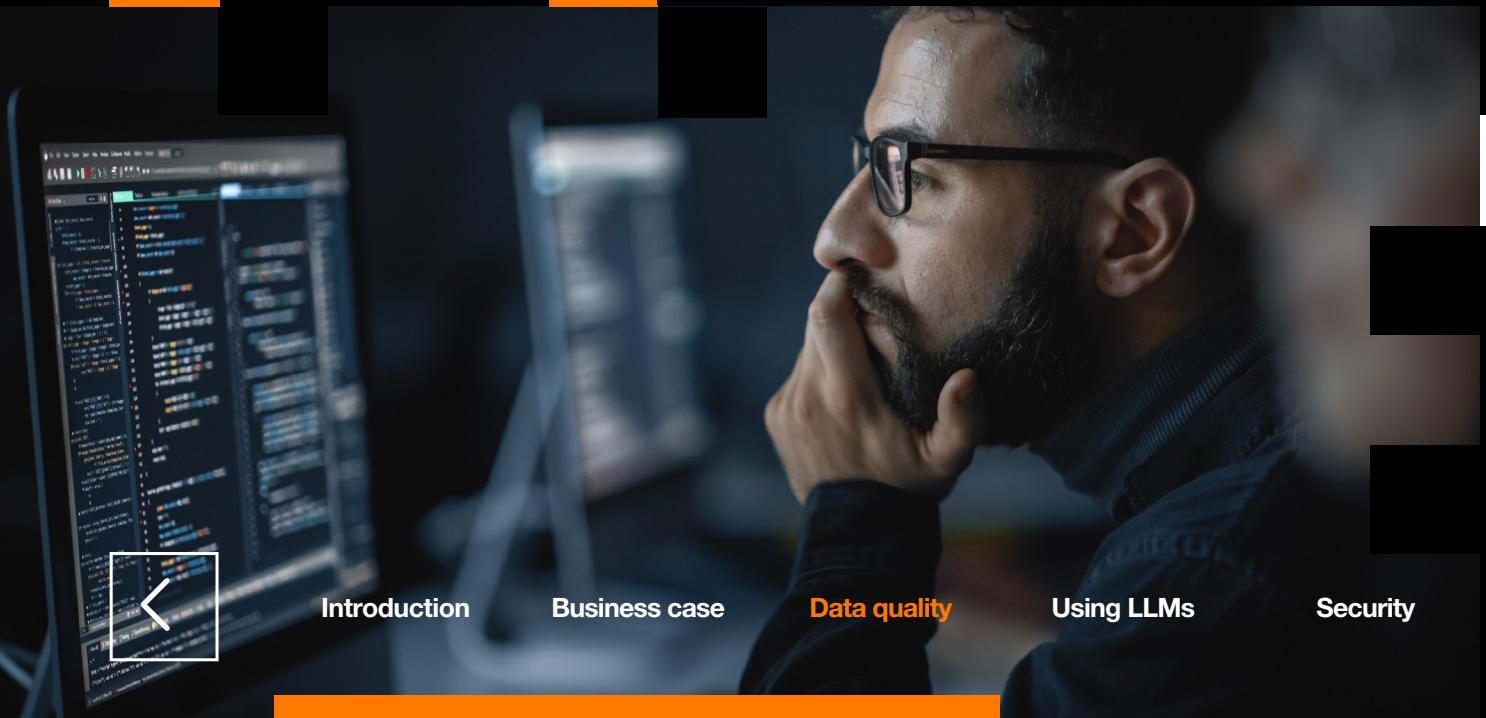

Using language models properly

Large language models (LLMs) have become one of AI's most visible uses, thanks to the widespread availability of services such as Bard, ChatGPT, and Co-Pilot. The number of new LLMs released worldwide in 2023 doubled over the previous year[2].

Perhaps partly because of that visibility, businesses often think that one LLM, run through an application like Bard or ChatGPT, can solve all their problems.

This misconception overlooks that, by being so large, LLMs are by their very nature generic. They are not specifically designed for individual business use cases. As mentioned earlier, when AI success is predicated on aligning with use cases, it is clear that LLMs must be tailored and trained to meet organizational goals.

What does that mean? Investing the time and resources in training models, understanding how to engineer appropriate prompts, and being clear on the guardrails that are needed to ensure LLM outputs are aligned with business objectives.

More specifically, to successfully harness LLMs, enterprises need to consider:

1 Recognizing and tackling hallucinations

LLMs are complex neural networks with billions of parameters. Through their development and training, they learn to reason and invent or generate – hence the name generative AI. Yet, while it learns from data, an LLM's primary objective is to create content, even if it doesn't know the answer. In these instances, it will generate what's known as a hallucination, where it invents answers to prompts. If that happens in a business situation, the impact could be significant.

Enterprises, therefore, need to be alert to the potential for hallucinations and put in place the necessary guardrails. That includes defining policies and procedures that ensure users know how to deploy prompts correctly, directing the model with the appropriate context, and providing background information to support it with the right answer.

2 Whether you fine-tune or RAG

It's worth reiterating that LLMs, for all their sophistication, are generic. To apply them to specific use cases, decide whether you will fine-tune your LLM or implement Retrieval Augmented Generation (RAG).

Fine-tuning involves teaching models using your private data. This has the benefit of only needing to do it once and the potential to be highly relevant to your use case and your business as a whole. The downsides are that it is costly, time-consuming, and, if a prompt comes in that the LLM doesn't know, prone to hallucinations.

RAG, conversely, is a way of engineering prompts using your specific knowledge to provide context. It can be done on several levels: simple, where you provide everything the LLM needs to know in a prompt, or more advanced approaches, in which you will split up your prompt and have the model focus on specific sections. RAG doesn't require you to train the LLM on private data; subject matter experts must work on the prompts and review output quality.

3 How sustainable is your LLM use?

LLMs consume a significant amount of compute power. With new models offering increased prompt size and answers, the energy needed will also grow. Be clear on whether you need the full capacity, and understand the implications of detailed prompts asking for large answers. Plus, it is worth considering how repetitive the prompts will become – will your organization ask the model different variations of the same question repeatedly? Storing and making answers accessible can reduce the number of times large prompts are deployed and cut down the compute demand.

4 The data privacy implications

If you tailor the LLM to your use case, you will use business data. As such, you'll need to be clear on the regulatory implications, as the models are third parties. Depending on the use case, you may need to consider anonymizing your data—for instance, using a model to analyze contracts. However, this could compromise the effectiveness of the output, so it's important to be aware of this potential drawback.

5 Evaluation

Underpinning all of the above requires an effective, rigorous feedback loop. Outputs need to be assessed, prompts reviewed and refined, to ensure that models perform as expected. Establishing an evaluation framework right at the outset would help track answer quality and how often a model hallucinates, insights that could help shape further training and optimization.

Ensuring security

Cybersecurity remains a huge problem for organizations. As they embrace digital innovation and harness powerful technologies, including AI, they also introduce new vulnerabilities, which increases the risk of being attacked. Additionally, as AI is as available to attackers as to companies using it for good, exploiting the already old vulnerabilities can now be done with superpowers.

This is the global digitalization paradox: the technologies that will unlock new levels of growth also have the potential to accelerate the explosion of toxic assets. These are the people, processes, and tech that burden resources, generate employee dissatisfaction, and propagate weak spots in corporate systems.

According to industry analysts and security companies, over 80% of companies have vulnerabilities in their IT landscape. They also suggest that it equates to one vulnerability per application. Now, consider how many applications the average digitally-enabled business has, and we start understanding the issue.

Historically, this would have been deeply troubling. The saving grace was that cyber attackers needed specialist knowledge to exploit these vulnerabilities. Not anymore. AI's increasing maturity is having a major impact on cybersecurity in two ways:

First, it's lowering the barrier to entry. AI is turning everyone into a superhero, able to complete so much more, so much faster. That's great when it's accelerating employee productivity and terrible when it's doing the same for cyber attackers. Now, anyone can break into a system, network, or application. Where that once used to take weeks or months per attack, now it's down to minutes, with multiple simultaneous attacks. If you've got several hundred applications with vulnerabilities, scale is no longer a form of protection; even a relative novice can exploit them all simultaneously.

Second, AI's ability to remove repetitive work makes it attractive to various functions and departments. New tools are being deployed in HR, IT, development, marketing, and legal. How do you know if they've got vulnerabilities before they're being deployed? Every new application is a potential security hole.

How do you avoid this? By focusing on four areas:

1 Get rid of toxic assets

With so many vulnerabilities, most enterprises don't have the resources to patch everything. So, you need to identify where the biggest vulnerabilities are and get rid of those assets.

Ensuring security

2 Tailor your security

Historically, corporate cyber security has been defined by security professionals and applied across the entire organization. But your security teams aren't HR specialists, they aren't sales executives, they aren't marketing managers, all with their own priorities and pressures. It creates situations where security might restrict and hinder, so employees look for workarounds and inadvertently create more weaknesses. So, security needs to accommodate the needs of the employee. That means defenses that protect but don't hinder, but it also means training personalized to that function based on scenarios that reflect how teams operate.

3 Revise your security strategy

Most cyber defenses are based on traditional practices and policies. They're not designed to cope with superpowered attacks. It's time to revisit, revise, and build a strategy that reflects the continuously evolving nature of today's cyber landscape and business needs. It's about adjusting and adapting, using standards and compliance as a minimum, not an ambition, and accepting that security is not a one-and-done situation.

4 Know your limitations

No one can do everything, whether a company trying to do it all themselves or a vendor promising to cover all eventualities. Modern cyber security is about creating ecosystems of partners that understand the realities of today and your specific scenarios. That might be using existing partners in new ways or engaging new support; whatever it looks like, it needs to be a step change away from existing principles and practices.

Designing a network for AI

Network services have constantly been influenced by broader technology demands. In the 1990s, we saw the adoption of global voice, while a decade ago, we had the proliferation of cloud services. Each necessitated an evolution in how networks were acquired, deployed, and used.

Most companies accept they need a digital infrastructure: 80% of decision-makers worldwide recognize that it is important or mission-critical to achieving business goals[3]. That includes networks.

Now, we have the AI explosion. We already know there will be huge demands on infrastructure, with 52% of GenAI investments going on dedicated or public cloud infrastructure in the next 18 months[4]. But what will that do to the network?

No one has a crystal ball, so we can't predict how enterprises will use AI and what they need from their networks to support those implementations.

What is clear is that the network will be required to adapt to enable AI use cases. It needs to have the capacity, bandwidth, and latency to manage the mass of data being created and processed by AI at the edge, thanks to the Internet of Things and other connected devices. It must also deliver similar capabilities to enable huge LLMs to operate effectively, whether getting the mammoth data sets needed to train them or ensuring that the apps and services they support perform as expected.

It is challenging for those tasked with delivering that network infrastructure to know what is required. Yet waiting and seeing won't work; the accelerated pace of AI adoption means enterprises need a foundation to cover all eventualities with a future-proofed, adaptable, and scalable strategy.

What does that look like? It's a strategy that encompasses:

1 Design for use

As we've already seen, the use case shapes how an AI is used. That, in turn, influences the network requirements, such as where the AI is located (whether at the edge or in the cloud) or how decentralized the business (and its data) is. Yet whatever the demands, there is likely to be a need for compute power, high-speed connectivity, and data on the move, all of which the network needs to support.

2 Privacy, security, and regulation

We've also spoken extensively about the privacy and security implications of using data in AI, but what about getting the data to the applications? With data storage increasingly decentralized, securing networks between data, compute, and applications is critical to maintaining privacy. Where you store your data can have sovereignty implications, which must also be factored in, along with inconsistencies between geographies that will require added elasticity. Plus, while there has not been much movement on regulation, more will come. That means any network deployment must be able to respond to legislation-driven changes.

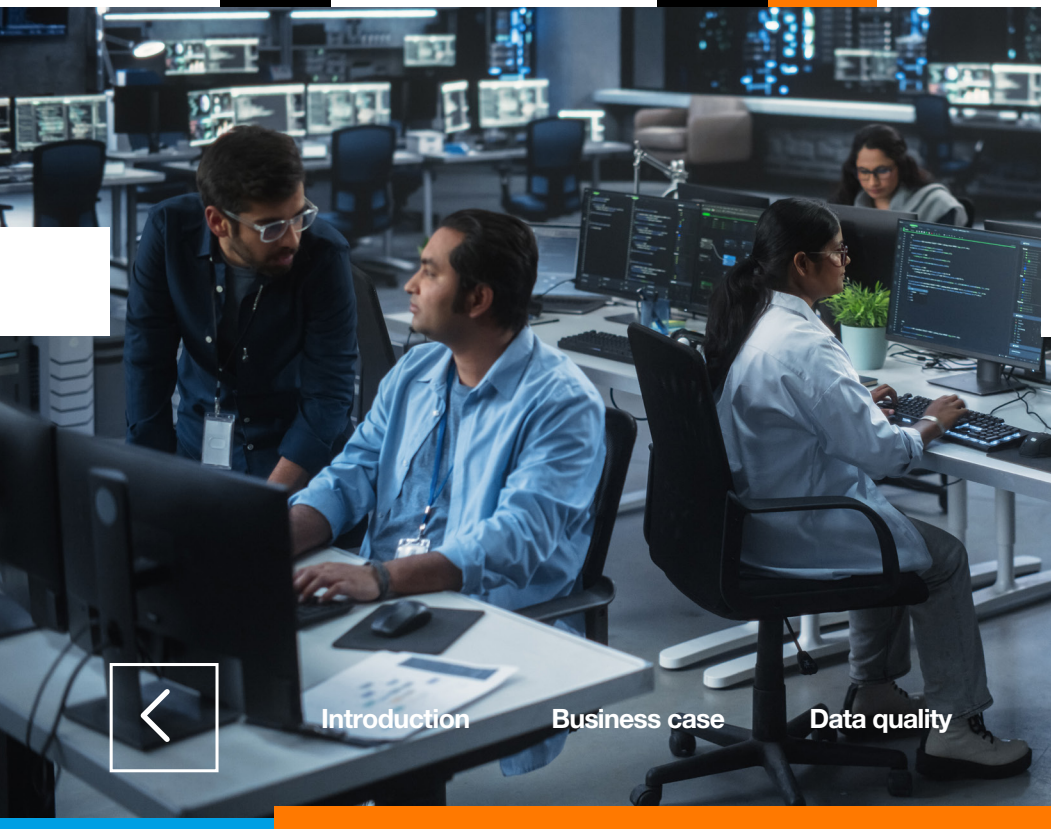

3 Managing congestion, traffic, and disruption

Traffic will grow; that's a given. With that comes the potential for congestion, impacting speeds and latency. Business applications of AI may, in the future, require access to fast, protected, and low-latency connectivity to function properly and deliver their full potential

At the same time, public and private networks are subject to disruption. We might exist in a cloud-based virtual world, but the connectivity that drives it all relies on very physical cables. As the news reminds us regularly, these can be at the mercy of environmental or geopolitical events.

So, any strategy must include contingencies that maintain service levels without degrading quality and ensure networking is part of the overall plan.

4 AI governance

The network also has an important role in AI governance. The importance of this cannot be overstated: poor governance can lead to implementation errors, breaches, and data exposure. So, knowing who manages data, determining whether it is trustworthy, and the visibility of the infrastructure it uses, including the network, is a key part of good governance. You can't trust data that is being delivered by an insecure network.

It’s time to plan for an AI future

Even if you aren’t ready to deploy AI now, start planning for it today.

AI isn’t perfect. No doubt anyone who’s been taking an interest in the technology will be aware of its current limitations. However, the fact is that it is already being used by businesses in a variety of industries, from manufacturing to healthcare. They might be experimenting, piloting, or testing concepts, but most importantly, they’re learning what works and what they can industrialize. These organizations will race ahead of the competition when they scale their deployments.

Here, we’ve covered how to build a business case, improve data quality, what to know about LLMs, security implications, and the need for AI-ready networks. If we could leave you with one more message, it would be this: even if you aren’t ready to deploy AI now, start planning for it today. Invest in the network, look at your data quality, update your security, and, most importantly, identify the use cases that will benefit from AI. The use case defines everything; if you shape your AI deployment to genuine business need, you will be much better placed for success.

How Orange helps

As one of the world’s leading network and digital integrators, we’ve worked on everything enabling AI for many years. From our industry-leading network and digital infrastructure to our security expertise and ecosystem of partners delivering the environments enterprise needs to fulfill their digital ambitions, we are laying the foundations for the AI era.

We have the skills and experience to help you:

- **Build the business case for AI by identifying the most valuable use cases for your organization**
- **Ensure you have the data you need to inform your AI deployments**
- **Harness the power of LLMs**
- **Develop and deliver a security posture fit for the AI era**
- **All delivered via a future-proof, agile, and scalable network infrastructure**

Get in touch today to deploy the optimal AI journey for your organization.

For more information visit:
www.orange-business.com/en/contact

Sources

1. https://www.gartner.com/en/newsroom/press-releases/2023-10-11-gartner-says-more-than-80-percent-of-enterprises-will-have-used-generative-ai-apis-or-deployed-generative-ai-enabled-applications-by-2026
2. https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_2024_AI-Index-Report.pdf
3. https://blogs.idc.com/2022/12/09/idc-futurescape-worldwide-future-of-digital-infrastructure-2023-predictions/
4. https://www.idc.com/getdoc.jsp?containerId=US51313423

orange™ Business