



Business

**Artificial Intelligence.
Real Wisdom.**

Sophie's COO doesn't know where to start. But Sophie has a clear action plan for taking her PoC enterprise-wide. Find out here what she and other Orange Business customers know.

Your route to operationalizing AI:

Transforming PoCs into production-grade GenAI services





Business

| | |
|--|-----------|
| Introduction | 3 |
| <hr/> | |
| Tackling data quality | 4 |
| <hr/> | |
| Using language models properly | 5 |
| <hr/> | |
| Ensuring security | 7 |
| <hr/> | |
| Designing an infrastructure for AI | 9 |
| <hr/> | |
| Value creation | 11 |
| <hr/> | |
| Your AI-driven future is harder to realize than you think | 12 |
| <hr/> | |
| How Orange helps | 12 |

Introduction

89% of business leaders think AI will transform their organizationsⁱ. Yet 90% of CEOs are waiting for GenAI to move past the hype or Proof-of-Concept (PoC) stage. And 66% of leaders are ambivalent or dissatisfied with the progress their companies have made with AIⁱⁱ. So, what's going wrong?

We understand that the urgency attached to operationalizing GenAI services can sometimes make for a stressful working environment. But, to borrow the words of Rudyard Kipling, I hope that the information in this document will help you keep your head when all about may be losing theirs.

Kristof Symons, CEO International, Orange Business

Today, the pressure is on to transform AI-based Proof of Concepts (PoCs) into scaled AI services – a process we describe as operationalization. However, those responsible for this task are waking up to the scale of the challenges they are facing.

It is relatively trivial to throw together a PoC (which is, of course, why many companies have already done so). With a PoC, there is only a very limited need to: coordinate teams; overcome silos in data repositories and fix outputs that may have low quality and many 'hallucinations' (false findings); ensure a robust security posture; or upskill large teams.

However, as you operationalize that PoC, each of these areas assumes a huge new significance. As a result, a yawning gap is beginning to appear between the demands of business leaders for production-grade services that can be rolled out enterprise-wide and the ability of technical leaders to deliver them. In fact, a recent Economist Impact report found that only 37%

of executives believe their GenAI applications are production-ready, a figure that falls to just 29% among practitionersⁱⁱⁱ. And (as recently observed in an AI survey by Orange Business and GlobalData) cloud costs that could be managed as part of existing budgets suddenly start to 'skyrocket'.

Pulling all these threads together to create scalable AI services that deliver real value takes more than mere technical know-how (although it does need plenty of that) – it requires the experience and deep expertise that informs good judgment to make the right decisions. It requires wisdom.

In this case, it's a wisdom born from a deep familiarity with the foundations of (Gen) AI – on which businesses are betting their futures. This includes (but is not limited to) an ability to overcome the silos that exist between cloud, cybersecurity, and infrastructure teams and technology. And the intelligence not only to set and monitor governance and security policies but to be able to adjust them as the risk and AI landscape evolves.

At Orange Business, we have this wisdom in abundance. It's demonstrated in how we guide our customers to operationalize and get value from AI, and it is reflected in the way that AI infuses our products and services. In this document, we surface some of this wisdom, describing the key considerations for operationalizing AI services, covering five core areas:

- 1. Tackling data quality**
- 2. Using LLMs properly**
- 3. Security**
- 4. Infrastructure**
- 5. Value creation**

Tackling data quality

For all its potential use cases, there is one fundamental rule that underpins all AI deployments: rubbish in, rubbish out.

What determines whether rubbish 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 who wants access can have it.

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 issues, security flaws, broken data entry processes – if you have an underlying issue with your data, GenAI will amplify it. It's also vital that these issues are addressed from the outset. It is far more expensive to remediate problems with data quality once the model is built than it is to ensure clean data at the start of the project. And, in the future, we may well see AIs trained on data produced by other AIs so any quality issues will become deeply embedded.

AI services also provide access to a much richer and more diverse set of information sources, such as text documents, audio and video recordings, email messages, scanned documents, and sensor readings. How to assess the quality of this unstructured data is a question that we are helping lots of customers to answer.

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 – this 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.

This approach should also be applied to unstructured data – documents such as emails and presentations – which will be critical to fine-tuning the model and essential to the successful implementation of a RAG strategy (see section below). Therefore, you should apply governance that sets out criteria for 'document quality'.

2.

Scale incrementally

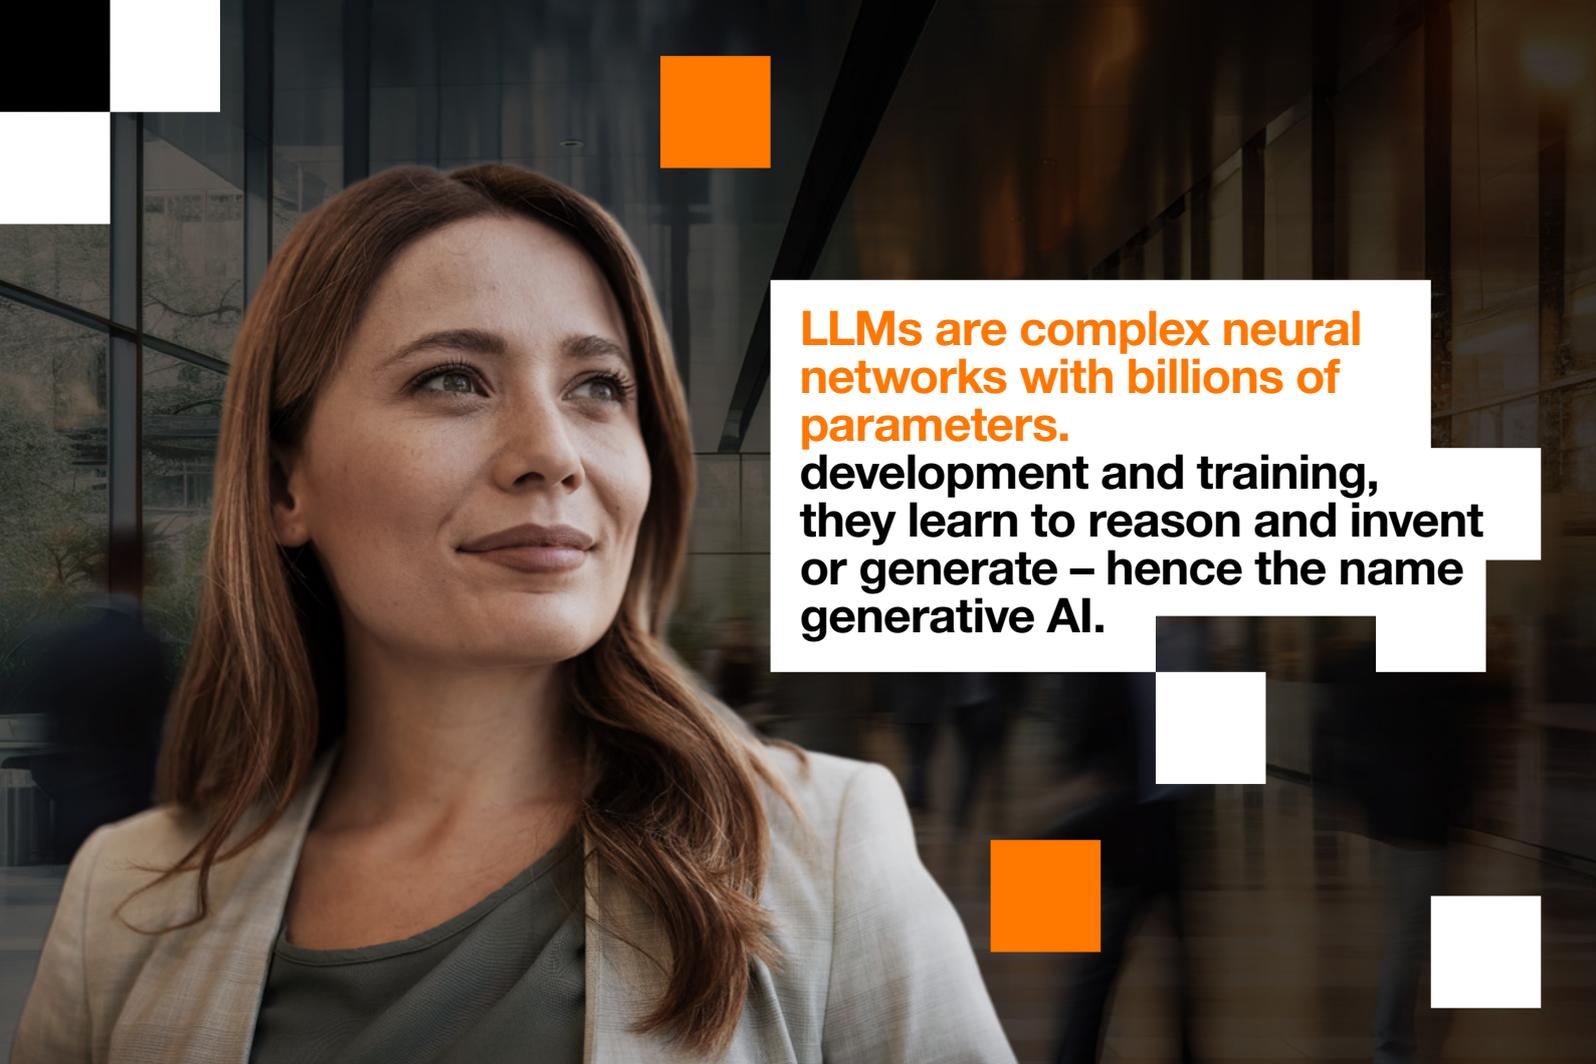
By taking a value-driven approach, you scale incrementally. With each step, you ensure that the data involved meets your requirements. This helps you build a body of use cases with demonstrable results and allows you to remain in control of all required governance. Rather than trying to comply with every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



LLMs are complex neural networks with billions of parameters.

development and training, they learn to reason and invent or generate – hence the name generative AI.

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 Copilot. The number of new LLMs released worldwide in 2023 doubled over the previous year^{iv}.

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 time and resources in training models, understanding how to engineer appropriate prompts, and being clear on the guardrails necessary 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 safeguards. 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?

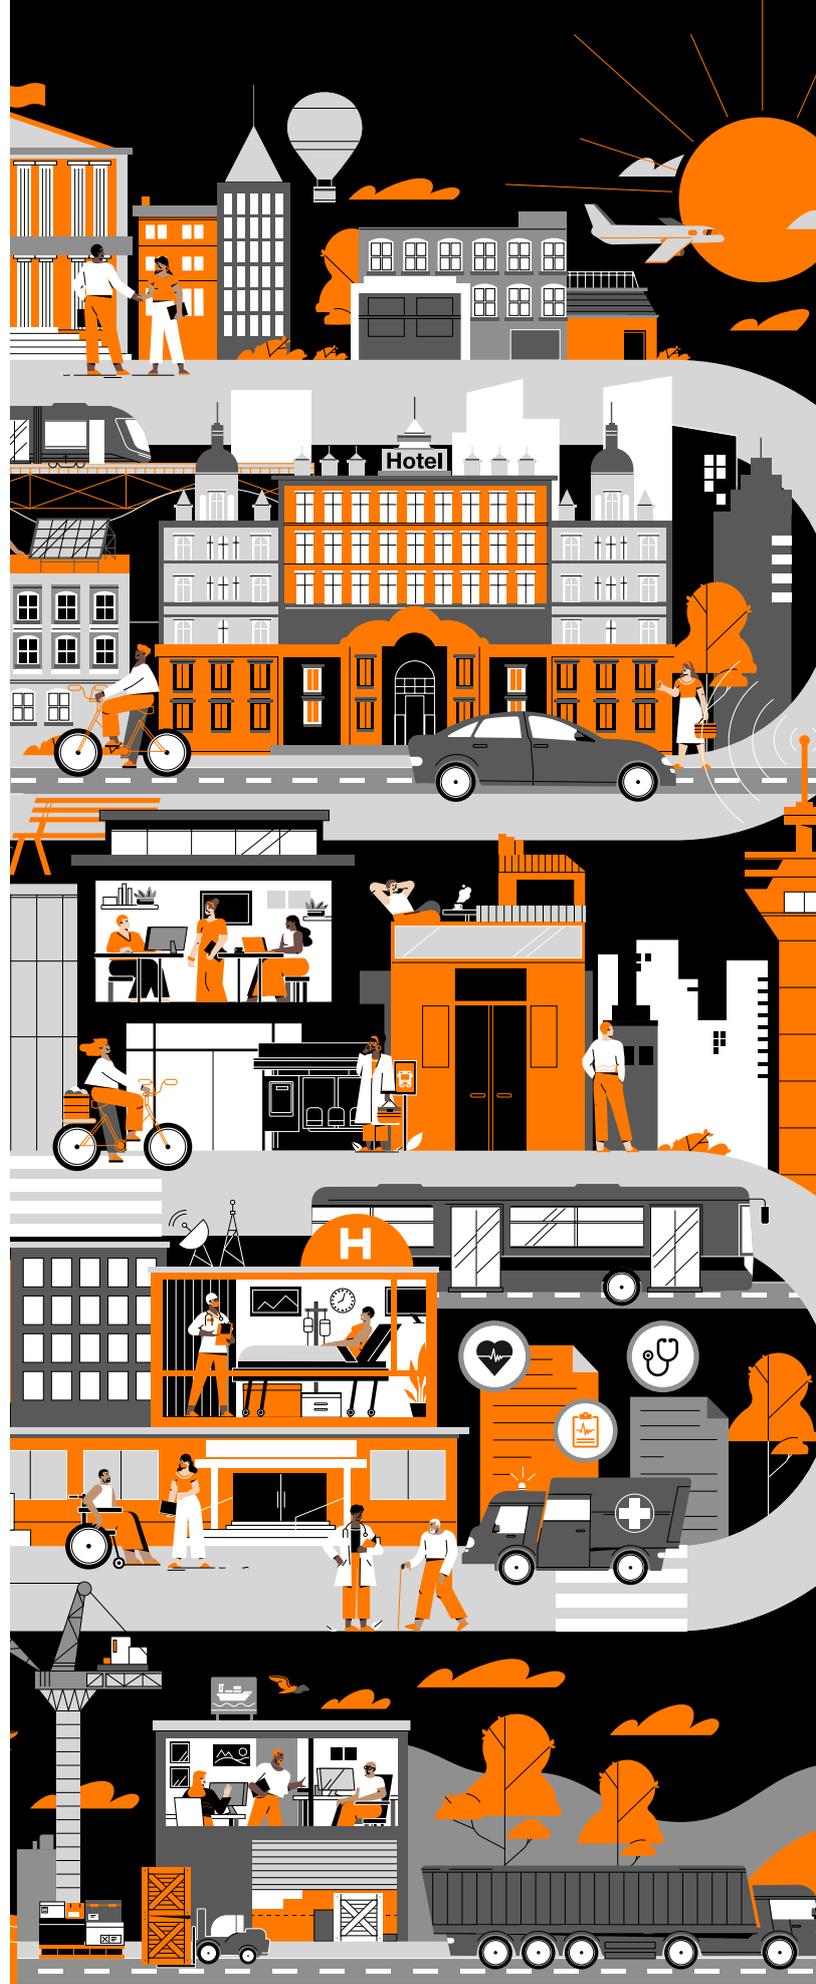
LLMs consume a significant amount of compute power. With new models offering increased prompt size and answers, the energy needed will also grow. Be clear on whether you need the full capacity and understand the implications of detailed prompts asking for large answers. Plus, it is worth considering 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 of doing so, 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.





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.

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.

2. Tailor your security

Historically, corporate cybersecurity 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 – and they all have 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 cybersecurity is about creating ecosystems of partners that understand today's realities and your specific scenarios. That might be through 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 an infrastructure 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^v. 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 into dedicated or public cloud infrastructure in the next 18 months^{vi}. 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 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. You must also factor in the sovereignty implications of where you store your data, 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.



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.

Value creation

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 GenAI alone, Gartner predicts that, 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.

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

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.

And don't forget the need to implement the right key performance indicators (KPIs). One Orange Business customer believes that getting to 60 or 70% data quality is a good enough threshold to deliver value from their AI services – but most organizations haven't put in place an equivalent metric for their projects. This is a classic case of, 'if you fail to prepare, you should prepare to fail': if you haven't defined what good looks like, how will you know if your AI project is delivering the value you need from it?

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.

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



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.



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.



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.



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.

Your AI-driven future is harder to realize than you think

You will hit a wall – the critical success factor is how you deal with that

AI isn't perfect. Anyone 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 if – and only if – they can successfully operationalize their AI deployments.

Here, we've provided practical guidance that we hope will help you realize your AI-driven future. If we could leave you with one more message, it would be this: don't underestimate the challenges involved in scaling a PoC to a production-grade AI

service. The GlobalData survey found that projects are taking longer and require more resources than originally anticipated because the complexities of operationalizing AI services are far greater than most people assume.

However, as Henry Ford once observed, "Obstacles are those frightful things you see when you take your eyes off your goal." So, in support of your objectives, ensure your network can support the use cases that make sense for your business, be borderline obsessed with data quality, ensure your governance and security policies are constantly updated, and most importantly, engage the right assistance as you tackle the headwinds you will inevitably encounter...

How Orange Business helps

We are a leading network and digital integrator that has been trusted to guard our customers' mission-critical data for the last several decades. Orange Cyberdefense is also one of the world's leading security services companies. And, we have been deploying AI in our own business for years: our AI tooling touches everything on our network; our employees use GenAI to generate content and manage projects; and AI-powered

threat intelligence is at the heart of our risk-based approach to cybersecurity. The accumulated wisdom acquired through this unique combination of expertise and experience – together with an ecosystem of best-of-breed vendors – means we can partner with our customers to co-create AI-driven services that drive lasting value for their businesses.

We have the skills and experience to help you:

- ensure you have the high-quality data you need to inform your AI deployments
- harness the power of LLMs
- develop and deliver a security posture fit for the AI era – today and in the future
- deliver a high-speed, low-latency network able to support your AI service in any location
- optimize the value of your AI service by understanding the KPIs and risks

Get in touch today to successfully operationalize AI in your organization.

Sources

- i. <https://www.dataiq.global/articles/2025-ai-and-data-leadership/#:~:text=Most%20organizations%20believe%20that%20AI,2024%20to%2089%25%20in%202025.>
- ii. <https://www.bcg.com/press/12january2024-ceos-genai-hype-or-experimenting#:~:text=According%20to%20a%20new%20report,and%20GenAI%20roadmap%20and%20investment>
- iii. https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_2024_AI-Index-Report.pdf
- iv. <https://blogs.idc.com/2022/12/09/idc-futurescape-worldwide-future-of-digital-infrastructure-2023-predictions/>
- v. <https://www.idc.com/getdoc.jsp?containerId=US51313423>
- vi. <https://www.gartner.com/en/newsroom/press-releases/2023-10-11-gartner-says-more-than-80-percent-ofenterprises-will-have-used-generative-ai-apis-or-deployed-generative-ai-enabled-applications-by-2026>