Pilot-to-production GenAI guide

Lessons from the front lines of enterprise AI adoption

Real-world insights from TELUS's enterprise AI transformation

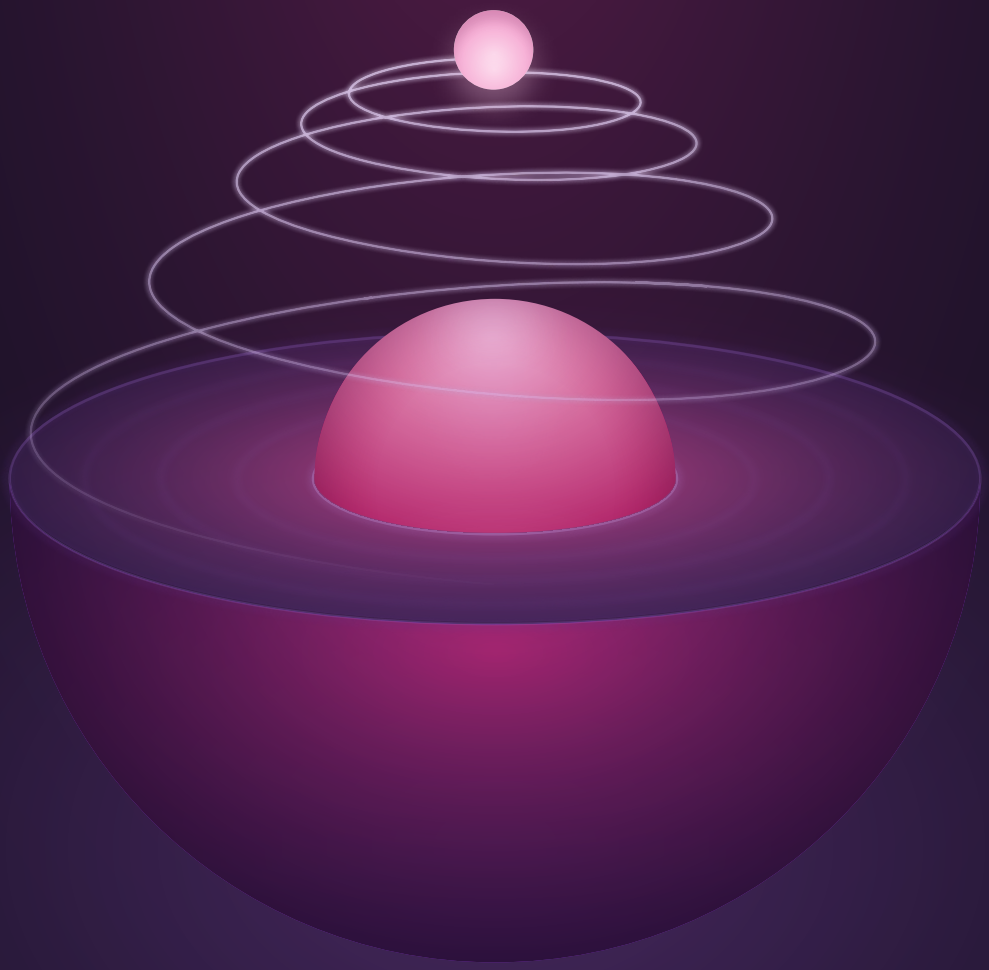

Fuel iX

TELUS Digital

Table of Contents

Introduction

Generative AI (GenAI) is no longer a novel experiment—it's rapidly becoming a fundamental layer of enterprise operations. From large language models (LLMs) powering intelligent chat interfaces, to multimodal systems capable of reasoning across many different content types, GenAI is changing how businesses ideate, operate and compete. However, the shift is as much about people and processes as it is about technology—ushering in an era where AI becomes a true copilot in decision-making, service delivery and innovation.

However, the shift is as much about people and processes as it is about technology—ushering in an era where AI becomes a true copilot in decision-making, service delivery and innovation.

While early hype sparked almost instant and widespread curiosity, most enterprises have now moved beyond the novelty phase. According to Gartner, GenAI has passed the peak of its hype cycle. Today, the focus is shifting from experimentation to execution—operationalizing GenAI in a way that's safe, strategic and scalable. As Hesham Fahmy, CIO of TELUS, aptly puts it; "We had three choices—we could either block it entirely, we could ignore it or we could lean in and understand what it can and can't do." TELUS chose the third path, marking the beginning of their GenAI adoption journey.

"We had three choices—we could either block it entirely, we could ignore it or we could lean in and understand what it can and can't do."

Hesham Fahmy
CHIEF INFORMATION OFFICER

TELUS

But that journey is far from simple. There's a long way to go between establishing a culture of curiosity and calculated risk to turning GenAI into a force multiplier throughout the business. Because of the many challenges along the way, it's easy to get stuck in the 'pilot trap'; launching small-scale proofs of concept without a viable roadmap to scale. Others face concerns around governance, model reliability and the soaring threat of shadow AI. As TELUS learned, achieving maximum business impact requires not only choosing the right platform, but also the right strategy, people and safeguards.

This guide provides a framework for addressing that complexity. Learning from TELUS's real-world experience, we'll share guidance on building cross-functional AI teams, implementing governance guardrails and defining meaningful success metrics. Whether you're just starting out or scaling across business units, this guide will help you move from pilot to production—and capture the true value of GenAI.

Whether you're just starting out or scaling across business units, this guide will help you move from pilot to production—and capture the true value of GenAI.

Beyond the hype: GenAI use cases that matter

In the rush to adopt GenAI, many organizations begin with an overly narrow mindset, viewing it solely as a tool to save time and cut costs. While these are, of course, important benefits, they represent a mere fraction of what GenAI is capable of. The real opportunity lies not in replacing human creativity or expertise, but in enhancing their capabilities for increased agility and improved decision-making. It's about reimagining how work gets done across the enterprise.

That said, narrowness in scope can be a strength, especially during the early stages of your GenAI adoption journey. Starting small with targeted, high-value use cases—like using GenAI to automate policy document reviews—allows organizations to move fast and prove impact, without breaking things. The first thing to consider is whether your proposed use cases are both valuable and viable. These are areas where GenAI can solve real business problems without needing months of development and widespread organizational overhaul. **These 'quick wins' help build confidence, generate internal momentum and validate your investments.**

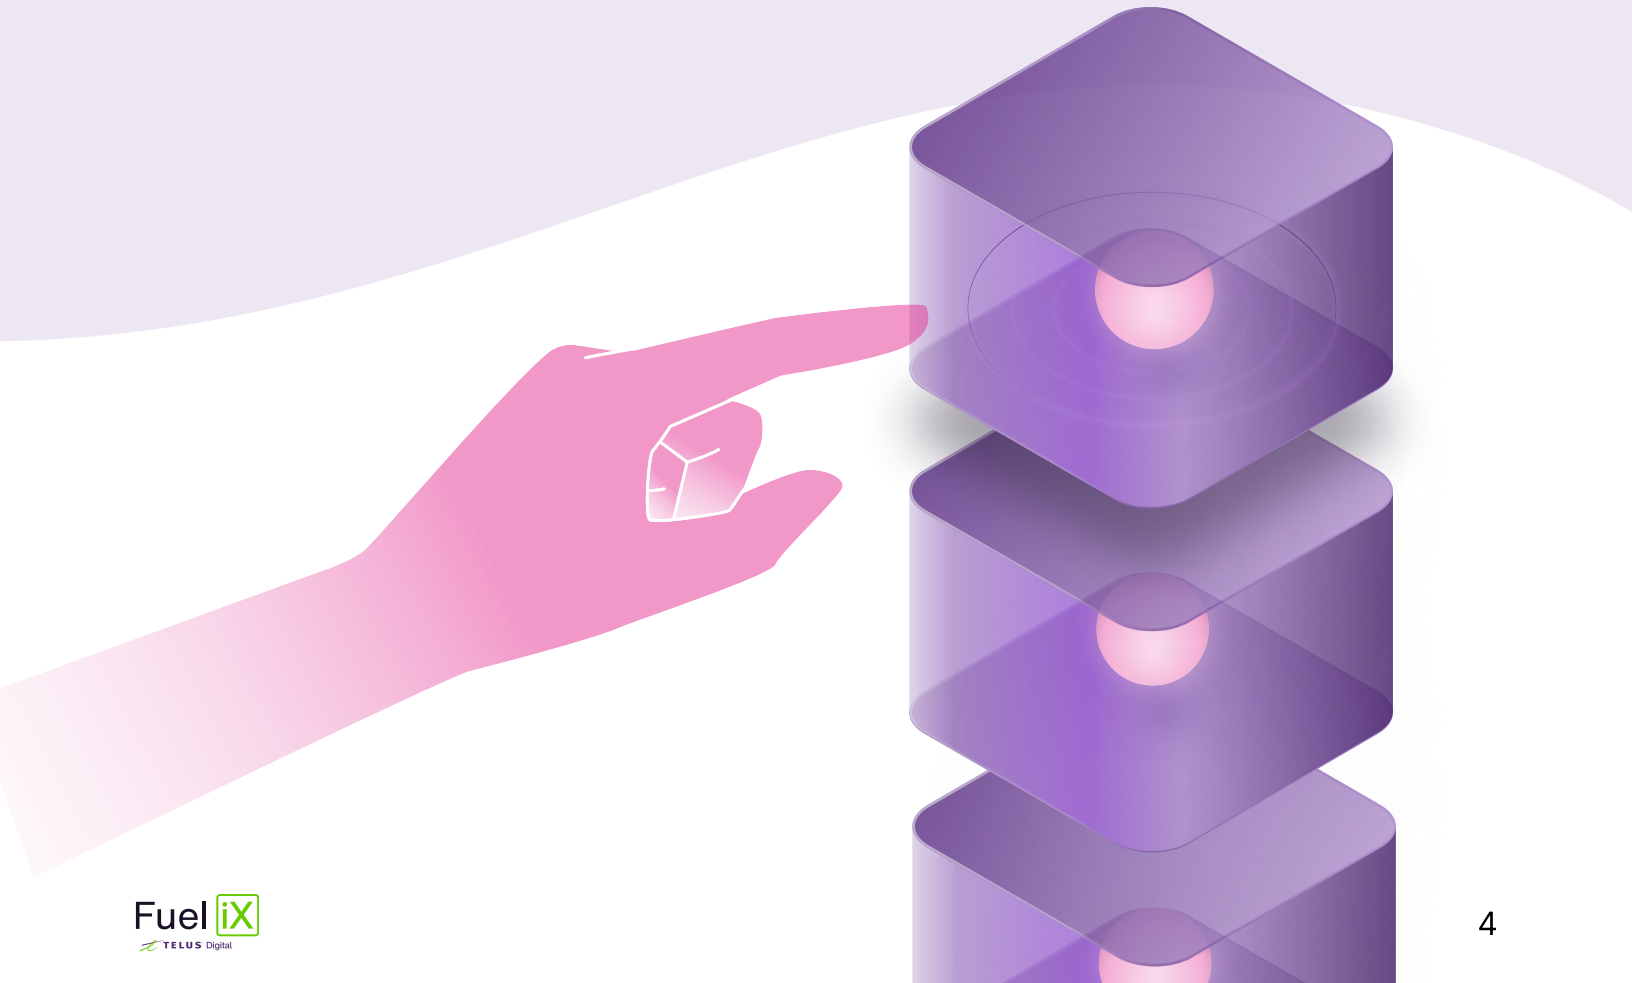

In practice, this means resisting the urge to deploy AI everywhere all at once. That doesn't necessarily mean limiting access to a small group—organizations that enable broader employee access to copilots often see faster, more effective adoption. On the other hand, you'll also want to avoid moving too slowly, lest the unsanctioned use of AI assistants (shadow AI) become a factor. Starting with a narrower scope—by addressing specific, clearly defined business problems—can still yield impressive returns. **It's about a 1% improvement each day, for example, leading to a transformational strategy over time.**

One such example is how many organizations started out with internal copilots as knowledge-retrieval assistants—surfacing hard-to-find institutional insights buried in documents scattered across teams. As Fahmy puts it: "There's so much tribal knowledge and so much that's undocumented. A lot of early use cases were about leveling up every employee to have the same amount of knowledge, because they had this assistant that they could talk to, to tap into that information."

As Fahmy puts it:

"There's so much tribal knowledge and so much that's undocumented. A lot of early use cases were about leveling up every employee to have the same amount of knowledge, because they had this assistant that they could talk to, to tap into that information."

Building on these initial use cases, other organizations have deployed copilots to assist with things like IT support, drastically reducing ticket resolution time by automating routine things like password reset requests. Field service teams are now using GenAI tools to diagnose unfamiliar hardware, while HR or legal departments are using copilots to review policies, draft documents and summarize regulatory content.

These aren't just proofs of concept or experiments. They're focused, practical applications that embed GenAI directly into everyday workflows to the point they can drive measurable business value from day one. By starting with clear, constrained problems, organizations can identify the viable, high-impact use cases that really matter and scale up with what works—without getting stuck in a cycle of endless experimentation.

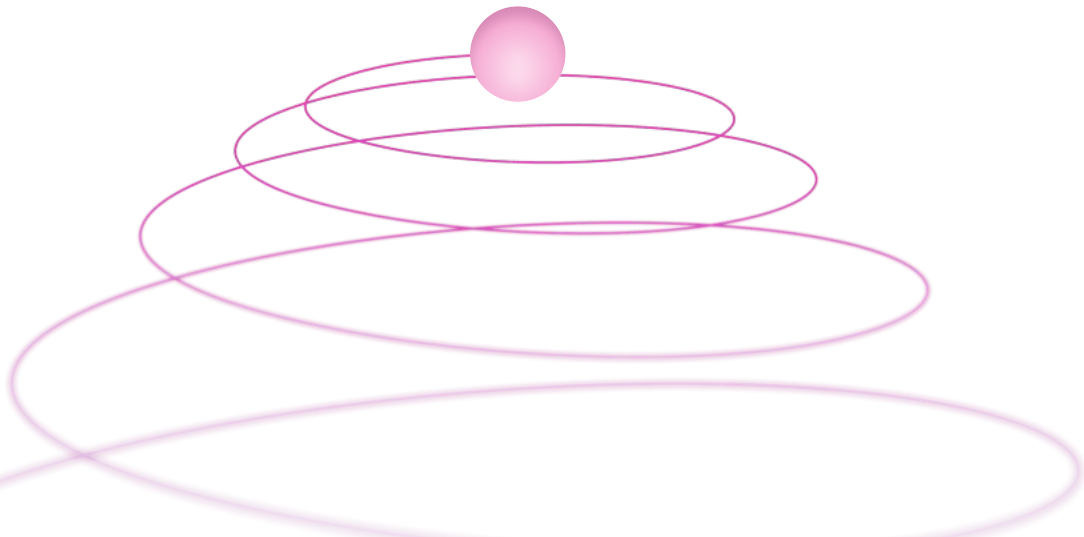

Strategic foundations for GenAI success

As generative AI adoption matures across virtually all industries, many enterprises are ready to move beyond the relatively narrow use of AI assistants and copilots. As organizations start seeing tangible results from their initial use cases, they must lay the strategic and operational groundwork needed to scale those efforts. However, few are fully prepared to scale in a way that drives innovation without adding risk. Successful adoption isn't just about technology, but also gaining a clear-eyed assessment of your organizational readiness. **This means evaluating GenAI maturity across people, processes, platforms and—most importantly—culture.**

Success starts with having an aligned strategy across the enterprise. If you're not clear on the value you're trying to create, you'll end up with different answers and competing ideas across different departments. That's why it's essential to bring your leadership team together to develop an enterprise-wide vision. As Fahmy puts it: "You can run workshops with your leadership team to get them familiar with how you actually use GenAI to generate new ideas." This sort of collaborative experimentation is essential for informing what your long-term strategy should look like.

Achieving enterprise-wide alignment has to start at the top, where line-of-business decision makers from across departments collaborate to establish the key performance indicators (KPIs) that matter. These KPIs will be vital in measuring progress—beyond productivity alone. For instance, these might include reduced cycle times for specific workloads, improved quality in decision-making or increased customer satisfaction or employee engagement.

As Fahmy puts it:

"You can run workshops with your leadership team to get them familiar with how you actually use GenAI to generate new ideas."

This sort of collaborative experimentation is essential for informing what your long-term strategy should look like.

A phased, departmentalized—but clearly aligned—roadmap can help ensure orderly deployment and avoid organizational chaos. While this approach may not eliminate shadow AI entirely, it can reduce risk by offering sanctioned, supported alternatives early on. Centralized teams should provide guidance, tooling and governance, while individual business units have a degree of freedom to experiment and execute use cases relevant to their domains. That's why the most important part of establishing the foundations for the broader adoption of GenAI requires a careful balance between control and creativity.

"You have to have a culture that's curious by nature and willing to experiment," says Fahmy; "But this could also be a very big risk, so you need to break it into smaller chunks so that any negative impacts are very small." To strike an optimal balance, many organizations pilot new GenAI projects in a secure sandbox environment, testing them thoroughly before putting them into production—albeit under strict guardrails.

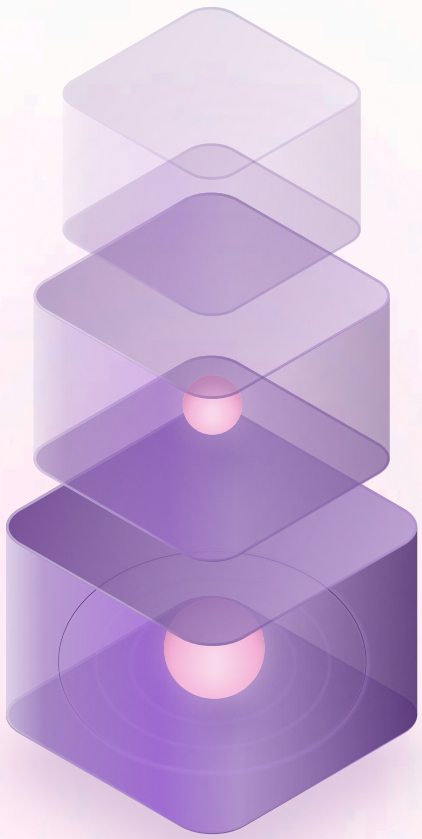

"You have to have a culture that's curious by nature and willing to experiment." says Fahmy; "But this could also be a very big risk, so you need to break it into smaller chunks so that any negative impacts are very small."

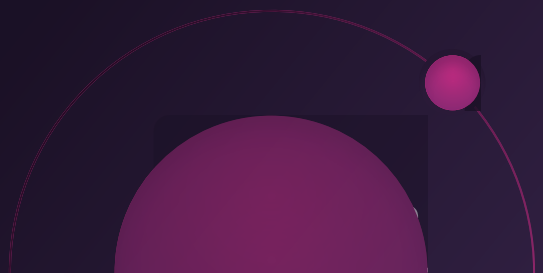

Another critical piece of the puzzle is achieving buy-in across the organization, from the C-Suite to the frontline employees who will incorporate GenAI into their daily work. Executives want to understand the risk profiles and ROI. Team leaders need implementation support. Employees need reinsurance and enablement.

Finally, clear governance must be baked in from the outset. Without clarity on who owns the platform, how the tools are vetted or how compliance is managed, shadow AI can become a real and serious problem. We'll explore the importance of governance more in Chapter 4—but its foundations must be laid early on.

Ultimately, scaling GenAI across the enterprise isn't all about deploying new tech—it's about aligning with business strategy, empowering your people and embedding accountability by design and default. Only that way, can innovation scale with confidence.

Building a team for GenAI enablement

Generative AI is redefining work processes and team structures. While roles continue to evolve, GenAI's purpose is to augment human capabilities, not replace people. “It isn’t about automating their jobs; it’s about empowering them in a different way,” says Fahmy.

For this to work, organizations must invest in AI literacy. Upskilling and reskilling initiatives are crucial, since they don’t just demystify the technology—they also reduce resistance and ease concerns around job disruption. Learning programs, internal workshops and peer-led communities can all contribute to building a more AI-confident workforce.

As GenAI capabilities expand, so too must collaboration between business units, IT, data science, and legal teams. Cross-functional collaboration is foundational to building a sustainable adoption strategy. As the boundaries between business and technology dissolve, organizations need to empower their teams to identify, develop and refine their own solutions. GenAI is the perfect tool for doing just that.

Central to this is democratization, where granting employees access to GenAI tools allows them to build and customize copilots suited to their daily workflows. These tools no longer require deep technical expertise—all they require is curiosity, context and an understanding of the problem to be solved. “Once you have a robust platform that allows anybody to build a high-value copilot, GenAI becomes as familiar as email or Excel spreadsheets,” says Fahmy; “Once you have the same kind of familiarity with these tools, you can empower teams by crowdsourcing ideation and allowing people to figure out how to make things function in their specific workflows.”

This level of enablement doesn’t just accelerate innovation—it also fosters a culture of ownership and accountability. However, democratization must be supported by structure. For instance, dedicated Centers of Excellence can help uphold best practices, avoid duplication of effort and support teams with templates, training and reusable assets. All the while, internal champions—early adopters who lead by example by combining business insight with a willingness to experiment—can help lead these changes from within.

Finally, while GenAI enablement should be democratized in execution, it should be centralized in the organization’s vision. The goal isn’t to let every employee choose their own tools, but to give them a safe, governed environment where creativity can flourish. Empowering teams with the right tools, training and trust is what ultimately turns GenAI from a pilot program into a powerful platform for innovation—one that scales from the bottom up, with alignment from the top down.

Managing risk with governance and guardrails

As with any major digital transformation initiative, the adoption of GenAI comes with significant risks. These span operational, legal, cybersecurity and reputational domains, among others. These risks are too serious to adopt the 'move fast and break things' approach—especially beyond the experimental phase.

While speed and experimentation are essential to innovation, they must be matched by governance and oversight from the outset. According to Monty Hamilton, Chief Product and Marketing Officer at TELUS Digital: "We want to make things easy, so you can set up an app and just have it running day and night. But we also need to have all the privacy and security guardrails in place, adhering to all the regulatory frameworks and responsible use of AI."

"We want to make things easy, so you can set up an app and just have it running day and night. But we also need to have all the privacy and security guardrails in place, adhering to all the regulatory frameworks and responsible use of AI."

Monty Hamilton
CHIEF PRODUCT & MARKETING OFFICER

TELUS

Understanding the core risks of GenAI adoption

Before you can adopt an effective governance strategy, leaders—as well as everyone else in the organization—need to have a clear understanding of the following risk areas:

Data privacy and security

Frontier models might use your data to train future models, meaning that any confidential information shared with the tool could resurface later. Always choose solutions that incorporate privacy by design and default to ensure data protection.

Model vulnerabilities

GenAI systems may be susceptible to adversarial inputs designed to exploit model weaknesses—like jailbreaking—potentially leading to undesirable outputs. However, red-teaming efforts help identify and mitigate such vulnerabilities.

Lack of explainability

Generalist models like ChatGPT or Gemini may produce outputs that are hard to justify or explain...and even hallucinate. Explainable AI systems reduce the risk of hallucinations or bias to support safer and more ethical deployments.

Regulatory compliance

While regulations concerning the use of AI are still in their infancy, things are changing fast. For example, the EU AI Act—the first of its kind—sets a baseline for the safe and ethical use of AI, particularly in high-risk use cases.

Unauthorized AI usage

Shadow AI arises if people use GenAI tools without authorization—typically because they believe those tools help them do their jobs better or more efficiently, and no sanctioned alternative is available. To mitigate this risk, it's best to take a platform-based approach where you have a centralized control plane for managing and monitoring all GenAI tools in the enterprise.

Choosing the right tooling is only the first step. You also need a governance framework that sets clear rules around which tools and models are approved for use, which data can be used and how outputs are reviewed and validated. Governance should account for privacy, attestation, explainability and compliance—aligned with both the internal appetite for risk and external regulatory expectations.

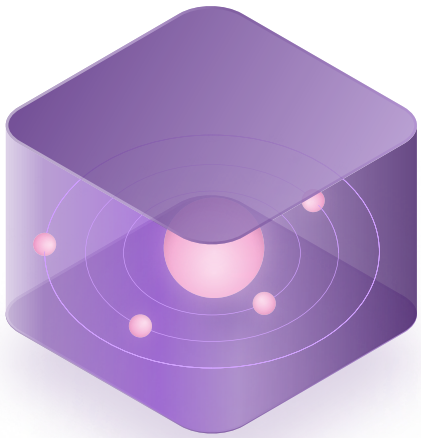

Keeping humans in the loop (HITL) helps monitor, enforce and continuously improve your governance initiatives. By contrast, fully autonomous systems can raise accountability concerns, especially if they don't feature end-to-end explainability.

However, GenAI-powered copilots are designed for human collaboration—they're not intended to replace human judgement, but to augment it. Combined with secure access controls, full audit trails and continuous monitoring, you can build a safe environment to scale without losing control.

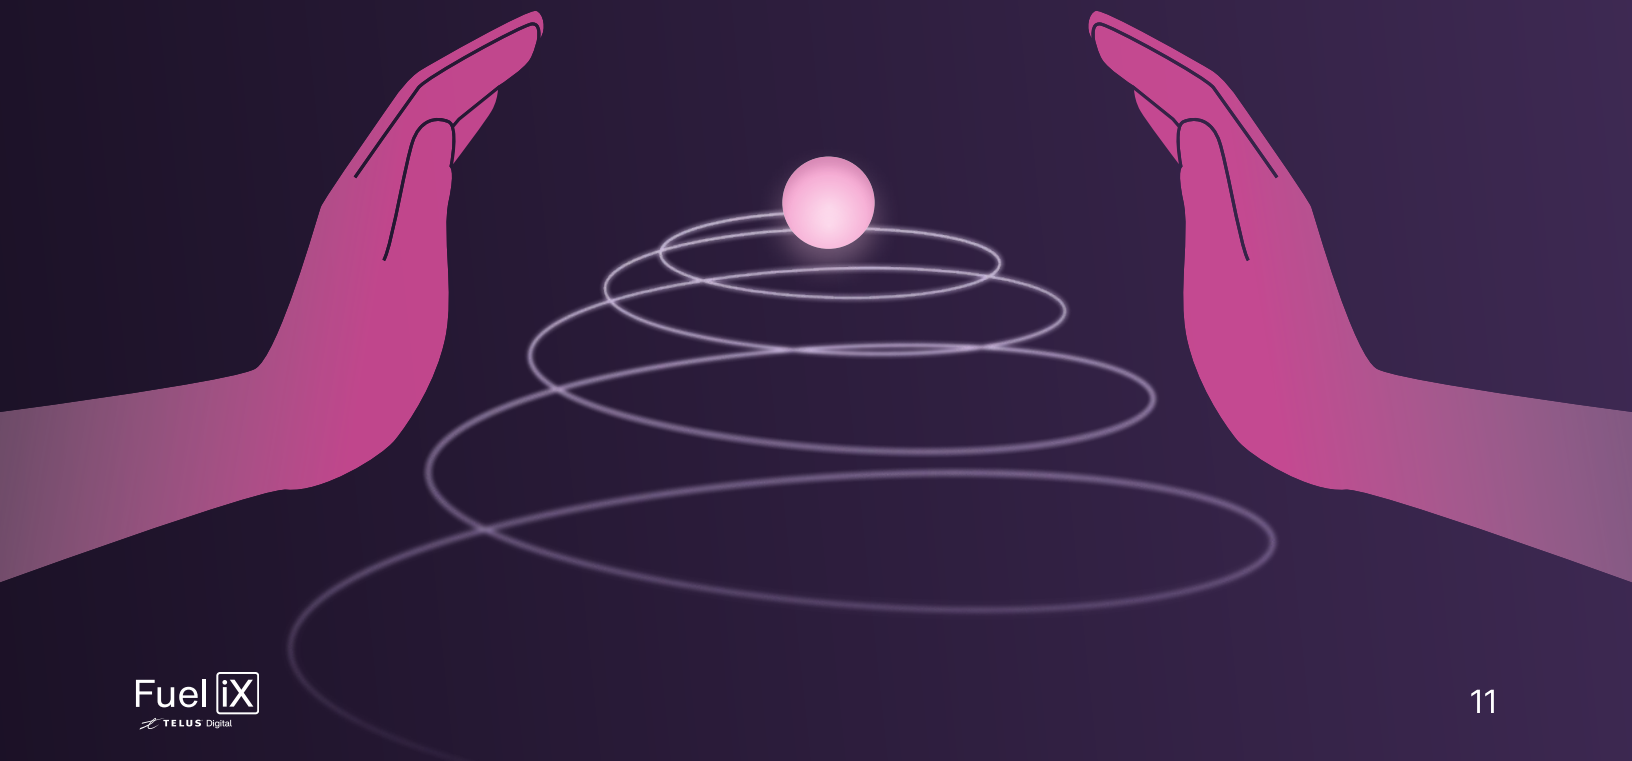

Tracking GenAI impact across the enterprise

When it comes to measuring the impact of generative AI copilots, we're often asked, "How can we demonstrate real value to stakeholders?" It's a crucial question, but the answer isn't always straightforward.

While our ultimate goal is to tie GenAI metrics to measurable business value, this can be a complex task. Instead of getting lost in the weeds of abstract potential or vanity metrics, **let's focus on three key areas that serve as leading indicators of business impact: productivity, quality and new capabilities.**

Think about it this way:

Does your GenAI tool enable users to complete tasks faster?

Does it improve consistency or output quality?

Or perhaps it unlocks new revenue streams or capabilities that were previously unfeasible with human labor alone?

These are the questions that can lead us to tangible outcomes like revenue growth, cost savings, operational improvements and enhanced job satisfaction.

Now, here's where things get interesting, and admittedly, a bit challenging. Measuring these indicators effectively requires a tailored approach for each copilot use case. It's not enough to say, "We want to improve productivity" or "enhance response quality." These goals are too vague and often lead to subjective opinions rather than concrete insights.

Let's look at a real-world example.

Say you've implemented a copilot to assist with IT or HR queries. In this case, you might want to track ticket deflection rate, time to resolution and employee satisfaction scores. You could even estimate cost savings based on the standard cost of L1 support per ticket. If your copilot is deflecting 25% of tickets, suddenly you have a clear picture of its value.

Or consider a project management copilot. Here, you might focus on time saved for specific tasks like summarizing notes or providing status updates. How often is it used? If it's saving an hour each time and it's used four times a week, you can start to quantify the labor costs that can now be redirected to more strategic activities.

Of course, not every organization has standard costs or task times readily available. This is where collaboration becomes key. Work with your team members to estimate baselines. It takes time, but it's a crucial step that's often overlooked. Without it, we tend to fall back on gut feelings, which can be misleading.

When it comes to measurement strategies, we've found that a mix of approaches works best.

1 User Surveys

- Gather insights on time saved
- Measure perceived productivity gains
- Collect qualitative feedback on user experience

2 Usage Analytics

- Track engagement rates
- Monitor frequency of use
- Analyze patterns in tool utilization

3 Quality Ratings

- Implement simple feedback mechanisms (e.g., thumbs up/down)
- Use caution when interpreting results due to potential biases
- Consider more detailed rating systems for specific use cases

4 Job Satisfaction Monitoring

- Assess changes in employee satisfaction related to AI tool usage
- Measure team engagement levels
- Conduct regular check-ins or pulse surveys

5 Time-to-Deployment Metrics

- Track how quickly new copilots or GenAI functions are developed
- Measure implementation time for AI solutions
- Use as an indicator of your organization's AI maturity

6 Human Element Evaluation

- Observe changes in team dynamics and collaboration
- Assess impact on work-life balance
- Gather feedback on skill development and learning opportunities

The goal is to use a combination of these strategies to get a comprehensive view of your GenAI impact. Regularly review and adjust your measurement approach to ensure it aligns with your evolving business priorities and AI implementation.

Remember, tracking GenAI impact isn't just about justifying investments. It's about learning and driving ongoing enhancements. That's why creating feedback loops is so crucial. Regularly gather user feedback, analyze usage patterns and use this data to refine your prompts and retrain your models. Be prepared to adjust your KPIs as your business priorities evolve.

In the end, the key to successfully tracking the impact of GenAI across your enterprise lies in focusing on specific, measurable outcomes and creating robust feedback mechanisms. By doing so, you'll ensure that your GenAI deployments remain relevant, valuable and aligned with your organization's goals. It's an ongoing process, but one that can yield significant rewards in terms of productivity, quality and innovation.

Case study: The TELUS story

TELUS, one of Canada's largest telecommunications companies, has revolutionized enterprise AI adoption through Fuel iX, its innovative GenAI platform. What began as an internal initiative has grown into a remarkable success story, with over 50,000 users using AI to transform their daily work.

The platform's success lies in its unique approach to AI democratization. Rather than restricting AI access to technical teams, TELUS empowered employees across all departments—from legal to HR, product teams to frontline staff—to create and customize AI copilots within secure guardrails.

What made TELUS's approach distinctive from many other enterprise GenAI adoption strategies was establishing an optimal balance between structure and flexibility. With centralized governance to ensure effective risk management, consistency, and standardization, individual business units could retain their freedom to build and iterate on copilots tailored to their specific use cases.

The TELUS story proves that GenAI doesn't need to be driven solely by a small team of IT specialists. With the right platform—in this case Fuel iX—and cultural mindset, GenAI can become an enterprise-wide capability. If you're ready to move from pilots to platforms, now is the time to take the next step.

Key Outcomes through early 2025:

- Created over 18,000 custom GenAI employee copilots across more than 50,000 global employees.
- Saved employees over 500,000 hours, reducing time spent per task by over 40 minutes.
- Enhanced productivity and innovation while maintaining privacy and security guardrails.

How TELUS built a flexible enterprise AI solution that empowers 50,000+ employees

READ MORE ON THE BLOG

Driving AI adoption: Fuel iX Week of Learning

TELUS employs a multi-faceted approach to AI adoption, leveraging various strategies to ensure widespread organizational engagement and implementation. One standout example of their innovative tactics is the Fuel iX Week of Learning.

Recognizing that successful implementation goes beyond technology, TELUS launched this initiative in February 2025, driving over 14,700 hours of investment in AI learning across the organization. It included live and interactive demonstrations, Q&A sessions with Fuel iX leadership, regional watch parties, recorded sessions to watch later, daily summary emails, a dedicated event website and even a Copilots contest that awarded a Galaxy Tab to 30 winners.

Fuel iX Week of Learning watch parties across the globe

Fuel iX Week of Learning results

This program delivered impressive results:

Innovation Surge

96% increase in Fuel iX Copilots created

User Growth

31% increase in Fuel iX users

Active Learning

7,800 hours of interactive training sessions

Professional impact

Of the participants who responded to a post-event survey, 97.9% rated the program as either 4/5 (valuable) or 5/5 (very valuable) for their professional development.

Fuel iX is an enterprise AI platform developed through the collaborative efforts of TELUS and TELUS Digital. The platform seamlessly integrates company infrastructure with an extensive library of LLM models and generative AI applications, providing complete observability and control throughout the process. Designed to address a critical gap in the GenAI landscape, Fuel iX enables enterprise-scale management of AI applications, foundation models and data sources within a framework built to deliver safe, responsible and accurate AI-powered experiences.

The platform's model-agnostic approach serves over 50,000 users enterprise-wide, with its excellence recognized through the Responsible AI Institute's Outstanding Organization 2023 prize. TELUS's GenAI Customer Support tool, powered by Fuel iX, became the world's first to achieve ISO 31700-1 Privacy by Design Certification.

VISIT US AT FUELIX.AI