

Fujitsu Private GPT use cases for real-world impact

Fujitsu Private GPT use cases for real-world impact

Imagine a world where you can chat with your own data – instantly receiving relevant, reliable answers that help you streamline operations, enhance customer experiences, and accelerate business growth. With Fujitsu Private GPT, this isn't a future dream – it's your business advantage today.

Designed to operate securely on-premise, Fujitsu Private GPT lets you use the full potential of generative AI (GenAI) while maintaining control

over your most valuable asset: your data. Whether your focus is on protecting intellectual property, ensuring data compliance, or simply speeding up workflows, Private GPT brings the power of GenAI to your fingertips without compromising on security or privacy.

Fujitsu Private GPT is transforming businesses across industries by addressing unique challenges. In this use case study, you'll find examples of how it's being applied today.

CONTENTS

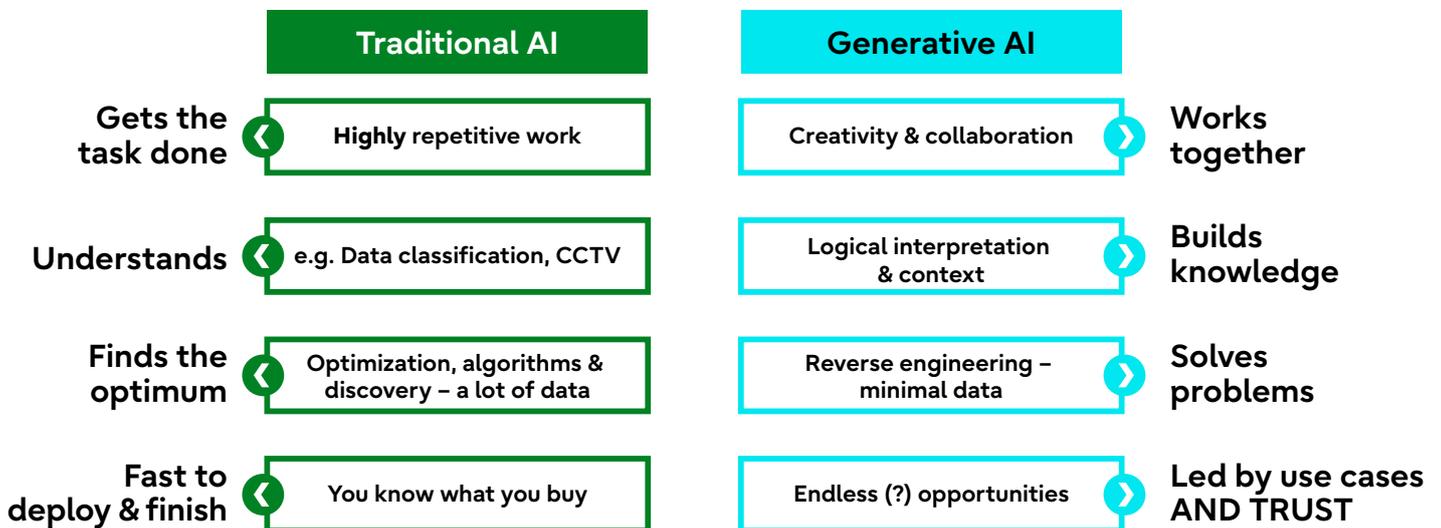
Know the difference between traditional AI and generative AI	3
What does a good use case look like?	4
Key ingredients for a successful use case	5
Case studies	
German SME transforms operations with Fujitsu Private GPT	6
Improving knowledge management and efficiency	8
Legal sector – managing complex data and contracts	10
Solving overwhelming tasks with Private GPT	12

Know the difference between traditional AI and generative AI

Generative AI represents a significant advancement beyond traditional AI technologies. While traditional AI automates well-defined, repetitive tasks like natural language processing or computer vision, generative AI introduces a new level of flexibility and creativity. It's not just an improvement – it's a fundamentally different approach. Generative AI enables businesses to

solve problems in innovative ways and generate new insights that traditional AI cannot provide.

The term "traditional AI" can be misleading, as AI itself is still evolving. Technologies like natural language processing and computer vision remain highly valuable, offering significant capabilities even within the framework of traditional AI.



Traditional AI is designed for well-defined, repetitive tasks, such as analyzing large sets of data or processing camera footage to identify patterns – like counting electric cars on a highway. These tasks are highly structured and can't be efficiently performed by humans, which is where AI shines.

However, generative AI offers something different: creativity and collaboration. It's not just about processing data but solving problems in ways that traditional AI cannot. Generative AI can build knowledge, interpret context, and generate new solutions with minimal input

data. It works together with traditional AI, complementing its strengths while enabling more flexible and innovative outcomes.

The power of generative AI lies in its ability to address a wide range of use cases, but it must be guided properly. By condensing data down to what's truly relevant, generative AI can create endless possibilities. As we continue to explore its potential, it's clear that its success comes from working closely with businesses to solve specific challenges, grounded in trust and use-case relevance.

What does a good use case look like?

A successful use case should deliver measurable benefits, such as time savings or significant improvements in quality at scale. The key trigger for identifying a strong use case often begins with the question: **How often do we need to perform this task, and what is the potential for optimization?** Whether it's automating processes or enhancing decision-making, the goal is always to achieve efficiency gains. In many cases, saving time directly translates into saving money, and improvements in quality can elevate a mid or low performer to a top performer, driving scalability.

The discussion around use cases is not just about automating tasks that were previously impossible (as with traditional AI) but also about enhancing the tasks you do today. You're already using data and knowledge in your processes, and for AI to be effective, this information needs to be made accessible to it. One of the challenges is that much of this knowledge is not formally documented but exists in people's minds. The first step in a successful use case is to capture this expertise so the AI can begin to assist.

Another important factor is localization. Use cases are often highly localized, not just in terms of language but in how people think and approach problems. For example, even within the same country, regions may have different ways of interpreting processes. This is why

customizing use cases for specific customers is so critical. Working with partners who understand the local context, often better than the customers themselves, helps ensure that solutions are tailored to specific needs. While the technology might be globally consistent, its application is deeply local.

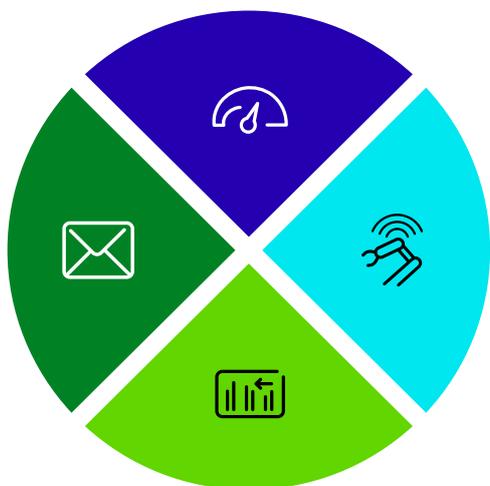
Finally, there's the role of IT architecture. A well-validated IT infrastructure is key to supporting AI systems. The process of implementation involves two essential components: first, making sound investment decisions in IT and AI infrastructure; and second, focusing on the adoption of these technologies through strategic partnerships and business processes. Balancing ambition with cost-effectiveness is crucial to ensure that the solution is not only innovative but also scalable and sustainable for long-term growth.



Key ingredients for a successful use case

The key advantage of Fujitsu Private GPT lies in keeping your data sovereign and protected. For many businesses, the primary reason for choosing a private GPT system is the need to retain full control over their sensitive data, such as intellectual property, company knowledge, and proprietary processes. Losing access to this data could be as damaging as losing your top employees, whose knowledge is vital to your company's success. Ensuring data sovereignty is not just a priority – it's essential, often mandated by strict data governance and compliance requirements.

In-depth knowledge on generative AI validated use cases:



- Ecosystem partnering**
 - Working in-depth with the right AI partners
 - Ready to scale & partnerships
 - Act local, think global
- User / Business process / Industry**
 - Time saving potential + quality gains at scale
 - Data/information: make the right data available
- IT & architecture**
 - Pre-validated IT infrastructure stacks
 - Speed is everything: fast innovation cycles
 - Scaling options by default + efficiency
- Generative AI engine**
 - Define the required LLM skill
 - Focused & optimized open-source models
 - We believe in mixed models / experts

By securing your data, you enable the foundation for trusted AI. A trusted AI solution gives you complete control:

- You decide how the AI works.
- No third parties have access to your data.
- Your AI is customized to your needs, learning from your unique data through training or inference.

This trust is further strengthened by focusing on cost and energy efficiency. In the IT world, cost efficiency is directly linked to energy efficiency. While there are upfront costs associated with acquiring high-performance systems, there are also ongoing operational costs, particularly with power-intensive technologies like GPUs. Fujitsu Private GPT ensures that both the initial and running costs are optimized for performance without unnecessary expenditure.

Finally, simplicity and performance are essential. For end users, the technology must be easy to use and seamlessly integrate into their daily work, improving processes and making their tasks more efficient. Whether it's simplifying workflows or improving productivity, Fujitsu Private GPT is designed to make your life easier, using your own knowledge and data to deliver meaningful, long-term improvements.

For example, Private GPT enhances the debugging process by making it faster and more efficient, but it doesn't replace the actual work. This is a common misconception about generative AI – it doesn't substitute human effort but rather supports people in doing their jobs more effectively.

A validated IT infrastructure is essential for making generative AI work effectively. The bottom half of the process represents the investment decisions related to IT and AI, while the top half involves adoption decisions based on partnerships and business processes. It's crucial to balance ambition with cost-effectiveness to ensure a scalable solution.

German SME transforms operations with Fujitsu Private GPT

One of our “hidden champions” in the German SME sector has been testing Fujitsu Private GPT over the past few months. This company specializes in digital transformation for medium-sized, factory-based enterprises, focusing on automation, digitalization, and electronics. Their expertise lies in areas like coding, integrating IT and operational technology (OT), data processing, engineering, and robotics.

They identified several ways that Private GPT could help them enhance their operations:

- **Knowledge access and sharing:**
Private GPT provides workers with access to essential knowledge, even when it’s not immediately available in their minds. It serves as a valuable resource for sharing knowledge across teams, enabling smoother workflows.
- **Communication and content creation:**
Private GPT assists in crafting and managing emails and documents, ensuring they’re grammatically correct, well-styled, and professional. Private GPT assists in creating and maintaining accurate documentation, and keeping it up-to-date. This frees up valuable time for the team to focus on higher-level tasks.

- **Software development:**
It is not about the actual coding, but about supporting the coding work by offering code snippets and debugging assistance. Private GPT aids in identifying and resolving code errors, facilitating code debugging, testing, and querying. It interprets error messages and offers step-by-step guidance in simple, human-readable language. The tool doesn’t automate code creation but streamlining the coding process without replacing the human developer’s role.



Why Fujitsu Private GPT?

Fujitsu Private GPT was chosen by this customer because of its robust data security and control features:

- **Data sovereignty and protection:** since the system runs on-premise on their own servers, the customer maintains full control over their data, ensuring compliance with stringent governance rules. It is completely secure since the customer maintains control over the data, within the framework of any strict data governance rules they may have. The data never leaves their premises - it's not even on the internet. It's only available on site. This level of security is a key advantage over public cloud-based GPT solutions, where data privacy may be at risk.
- **Open-source models and expertise:** Fujitsu's open-source solutions offer flexibility and customization. The customer has significant control over the system's capabilities and can fine-tune it to their specific needs. Furthermore, the Retrieval-Augmented Generation (RAG) technology enables efficient documentation integration, improving the quality and accuracy of AI outputs.
- **Cost and energy efficiency:** Fujitsu Private GPT provides technical excellence and ensures cost and energy efficiency, which is crucial for SMEs. The system is designed to be scalable, allowing customers to benefit from AI without incurring excessive operational costs.

A commitment to long-term partnership

Fujitsu's Private GPT represents a long-term partnership in the world of AI. This AI platform is continually evolving, and our roadmap ensures that it will remain aligned with the customer's growing needs. Future updates will include new models, expanded data types, and enhanced capabilities – delivering even more value as the technology continues to evolve.

Improving knowledge management and efficiency

A second use case comes from a software company with a similar objective to the first example: improving access to internal knowledge across their workforce. The company identified a need to streamline information retrieval for teams that were in transition, making their daily tasks easier and more efficient. The main focus was on enabling end users, such as consultants, project managers, and even customers, to access critical information quickly and easily.

Goal

To enhance workforce efficiency, we aim to make the internal company knowledge more accessible by employing a Large Language Model (LLM).

Situation

Despite the heavy reliance on task management and knowledge base tools by all users, including consultants, QA staff, documentation specialists, and pre-sales colleagues, the sheer volume of information and complexity of search results made it challenging to locate relevant data efficiently.

Specialized Large Language Models (LLMs), akin to those used in public cloud GPTs, have already been reviewed as a potential solution to make knowledge management more accessible.

Key problems

- **Time wasted on searching for information:** employees spent significant amounts of time, sometimes up to two hours a day, searching for data, depending on the project stage.
- **Redundant queries:** colleagues would often ask basic questions that could be answered by existing resources if the structure were more accessible.
- **Redundant content creation:** employees suspected that many “how-to” guides had already been created by others but could not be easily found.
- **Information overload:** the sheer volume of similar search results made it difficult to find relevant information.
- **Conflicting methodologies:** the company’s waterfall project structure clashed with the more agile, unstructured information retrieval methods used by certain teams.
- **Reliance on colleagues for guidance:** employees often had to improve their search queries by consulting colleagues, a time-consuming process that could be better handled by AI.

Solution

Fujitsu Private GPT, using an LLM, provided a local solution leveraging the company's existing internal knowledge while keeping all output internal and secure. A user-friendly, chatbot-like interface allowed employees to ask questions and be quickly directed to relevant sources. Private GPT also allowed for the activation or deactivation of documents based on context, improving search relevance and output.

Benefits and impact

- **Increased efficiency:** preliminary results indicated that consultants saved 1-4 hours per week, leading to more productive billing hours.
- **Faster onboarding:** new employees became operational more quickly, with easier access to essential knowledge.
- **Potential for monetization:** the company explored the possibility of upselling its LLM-based query service to partners already using its knowledge database.

This case highlights the importance of defining tasks clearly, optimizing internal knowledge sharing, and ensuring the efficient use of data sources. Fujitsu Private GPT's ability to structure data, group relevant information, and provide targeted results at the click of a button helped improve productivity and reduce redundant workloads across teams. By supporting end users like consultants in their daily work, Private GPT allowed the company to focus on core tasks, leading to higher efficiency and better overall performance.



Legal sector – managing complex data and contracts

In the legal sector, the challenge is often handling vast amounts of sensitive and highly confidential data that cannot be shared externally. This data includes critical documents like contracts, case files, and due diligence reports. Protecting this information is essential, especially from a client-confidentiality perspective, which is why many legal firms are exploring generative AI solutions like Fujitsu Private GPT to handle complex workloads while ensuring data privacy.

Together with experts in the legal field, we identified 26 use cases that could benefit from generative AI and evaluated each based on the feasibility of the technology, the potential benefits, ease of use, and the frequency of the task. Both CIOs and legal professionals contributed to this assessment, ranking the highest-value use cases for practical implementation.

Key use cases

1.

Contract analysis:

Private GPT helps legal professionals analyze contracts by identifying risks, unusual clauses, or a lack of detail. The AI can also act as a sparring partner by evaluating specific terms and suggesting ways to phrase them more clearly. This helps increase the efficiency of analyzing contracts, boost consistency and minimizes the risk of missing key terms.

2.

Knowledge management:

By utilizing the RAG technology, Private GPT helps you easily manage large amounts of knowledge. You can use the system to research past cases and find relevant information for your current case. Consider queries like “Summarize 5 cases we won in the area of copyright infringement within the last year and highlight the key success factors”. Additionally, when preparing for a presentation, Private GPT can suggest slides from other meetings to save valuable time and ensure you’re well-prepared.

3.

Analyze court documents:

Private GPT can analyze incoming court documents and highlight crucial parts. It contextualizes the document into the correct case, helping you keep all information up to date. Additionally, the system can identify and even calculate the exact dates of deadlines mentioned in the documents, making it easy to manage them effectively. This reduces administrative workload as well as the risk of missing an important deadline.

The role of generative AI

Generative AI plays a critical role in simplifying the overwhelming complexity of legal tasks. Legal professionals often face situations where they are inundated with vast amounts of data – whether that's case files or regulatory documents. The challenge isn't necessarily the complexity of the legal work itself, but the difficulty of synthesizing all this information into clear, actionable insights.

Fujitsu Private GPT helps by processing large volumes of text, summarizing documents, and identifying key details, which significantly reduces the manual effort required. For example, when a lawyer is faced with hundreds of pages of case law, Private GPT can instantly summarize the relevant portions, saving time and reducing cognitive overload.

Benefits

- **Increased productivity:** by automating document review, contract drafting, and due diligence, Private GPT allows legal professionals to focus on higher-level tasks that require human judgment.
- **Enhanced accuracy:** the AI provides logical, context-based answers that help lawyers make more informed decisions, reducing the likelihood of errors in contract interpretation.
- **Time savings:** by analyzing and summarizing documents quickly, Private GPT can drastically reduce the time required to complete legal tasks, enabling legal teams to handle more cases efficiently.

In summary, Fujitsu Private GPT offers significant value in the legal domain by managing vast amounts of sensitive information while ensuring data protection. Its ability to automate complex processes, such as contract analysis and due diligence, transforms how legal teams work – allowing them to focus on what matters most.



Solving overwhelming tasks with Private GPT

We've all experienced moments of feeling overwhelmed, whether at work or during school. This use case highlights how Fujitsu Private GPT can simplify complex tasks, even in scenarios that don't directly involve AI.

A client's daughter had a daunting school assignment: to read, analyze, and present on a 90-page 17th-century diary. The task was time-consuming and overwhelming, as the text was dense, difficult to understand, and required thorough analysis.

Without any help, it would have taken her at least four days of intense work.

Instead, she turned to Fujitsu Private GPT. By uploading the diary text into the RAG (Retrieval-Augmented Generation) database, she was able to save a significant amount of time. The AI handled the complexity of the text, quickly analyzing the material and presenting the most relevant information. While she still needed to understand the text's content, Private GPT provided a foundation of insights, allowing her to focus on the key points and topics she wanted to discuss.

This case demonstrates the power of Private GPT in tackling overwhelming, repetitive tasks. The AI helped reduce the cognitive load by processing large amounts of complex data, making the assignment more manageable and allowing the student to concentrate on her presentation.

For businesses, this example highlights how Private GPT can help employees manage complex and time-consuming tasks more efficiently, saving time and ensuring accuracy in the process.



What makes Fujitsu Private GPT different?

Unlike public GPT solutions, Fujitsu Private GPT is entirely customizable to your business needs and operates within your infrastructure.

You get:

- **Tailored AI solutions:** models trained specifically on your data.
- **Complete data sovereignty:** full control over data security and compliance.
- **On-premise deployment:** keep your data private, with the flexibility of as-a-Service models.
- **Scalable infrastructure:** powered by Fujitsu PRIMERGY servers and Intel® technology.

Why choose Fujitsu Private GPT?

Fujitsu Private GPT isn't just another AI tool. It's a secure, scalable solution designed to help your business succeed. Here's why it stands out:

- **Security & privacy:** your data stays private, compliant, and under your control.
- **Flexibility:** available on-premise or as a pay-per-use model through Fujitsu uSCALE.
- **Custom solutions:** our expert team works with you to identify the most impactful AI use cases and ensure smooth implementation.





Start your AI journey with us

Ready to see what AI can do for your business?
Test your AI ideas with Fujitsu AI Test Drive – a free,
secure platform that lets you validate your business
case using the latest AI infrastructure.

More information about Private GPT at
www.fujitsu.com/emeia/private-gpt.

Discover more about the AI Test drive and sign up at
www.fujitsu.com/global/ai-test-drive.