AODOCS

Enterprise chatbots: The keys to success

Table of contents

1. Generative AI: rapid adoption, sometimes fraught with pitfalls

In recent years, the use of Artificial Intelligence (AI) has become very popular, both within companies and among individuals. Recent developments in LLMs (Large Language Models) have demonstrated surprising capabilities in natural language processing (NLP) and text generation. They have enabled the creation of intelligent conversational robots capable of providing relevant answers to a user's questions. A growing number of such conversational agents have emerged, with LLM publishers competing in ingenuity to constantly improve the relevance of their tools.

Every day, we are all confronted with the promises of AI and LLM: increasingly sophisticated dialogues with chatbots integrated into our smartphones' operating systems, significant improvements in information search results on the web, contextualized proposals to support document drafting in office applications, and so more. The adoption of these man-machine dialogues in natural language is proceeding at dizzying speed.

The promises of digital intelligence have rapidly created new expectations within companies, particularly when it comes to the quality of the user experience and an application's ability to "understand" employees' requests. According to a survey conducted by Forbes in April 2023, 64% of companies are convinced that AI will help boost their productivity.

Many companies have reacted, setting up prototypes of solutions and applications based on generative AI, in response to a variety of relatively fuzzy challenges. AI experts and business application architects jointly examined the relevance of such prototypes. Their research led to the development of the RAG (Retrieval-Augmented Generation) technique, which aims to enable LLM to adapt to a company's specific context by giving it access to a precise, dedicated knowledge base (KB). Rather than relying solely on its initial training, the LLM can then exploit this KB to enrich its answers with relevant facts, in a lexical field in line with the company's activities.

Building up the knowledge base is therefore a key challenge, and one that will determine the quality of the answers provided by LLM. Although the triplet (LLM, RAG, KB) looks promising, an objective analysis of the initial feedback is essential. After a phase of euphoria, most companies have decided to postpone the deployment of their prototype

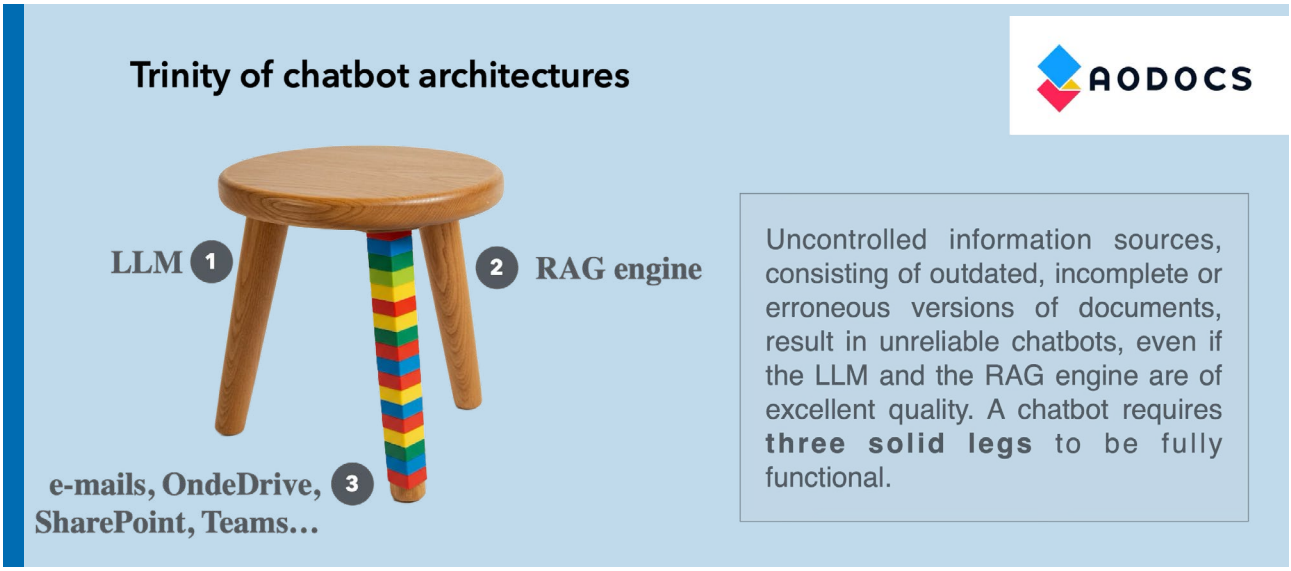

Chatbot : LLM + RAG + Knowledge Base

applications, while others have scaled back their initial ambitions.

LLMs can generate fluent, coherent text based on billions of parameters learned during trai–ning. But they have a major limitation: their knowledge is frozen at the date of their last update and can sometimes contain errors or biases. RAG is an innovative approach that improves response quality by feeding LLMs with external information that the LLM can use before generating a response. RAG is thus an informational integration architec-ture, enabling the LLM to query an external knowledge base (e.g., a document base, API or search engine) for specialized, up-to-date and relevant information, often in the lexical field of the user company.

Following the recommendations of enthu-siastic experts or the reassuring advice of suppliers of platforms exploiting generative AI, sometimes under pressure from manage-ment anxious not to “let the train pass”, some organizations have taken the plunge and put applications exploiting these new approa-ches into production. Analysis of their results prompts caution: chatbots producing weird responses, accidental and undesirable expo-sure of critical and confidential documents, errors in the conduct of key processes, etc.

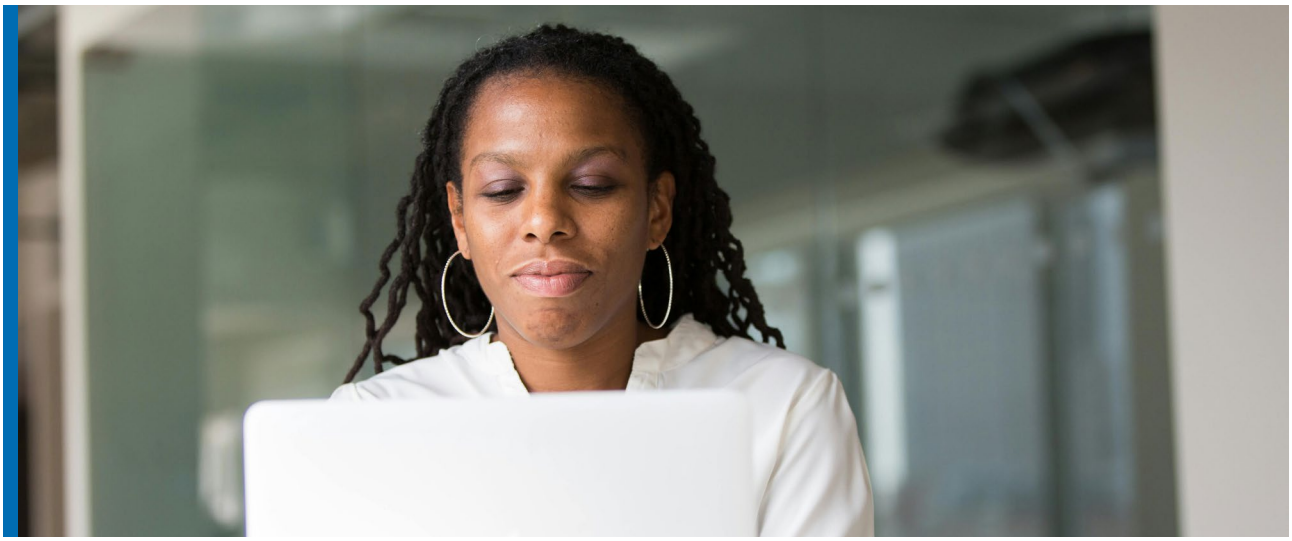

Here are some real-life situations

Costs resulting from a chatbot producing approximate or erroneous answers can be prohibitive, as the following examples from real-life situations demonstrate.

Price list error costs service company several hundred thousand dollars

A business engineer checks an enterprise chatbot for help in creating a quote for a prospect. Unfortunately, he is unaware that the generative AI he is using has been fed with an outdated pricing grid. As a result, the sales assistant produces a quotation containing prices 15% lower than current rates. The prospect accepts the offer, resulting in a dead loss for the company.

A legal file violates the intellectual property practices

An IT project manager at a healthcare company is preparing the certification of a digital solution as a medical device (FDA). He interrogates an AI to prepare his work plan. An incorrect setting of confidentiality rules for documents accessible to the AI leads to the use of information concerning patents filed by a competitor. The whole certification process is called into question.

Quality control error at aircraft manufacturer brings assembly plant to a standstill

An inexperienced employee oversees the quality control of the thousands of holes that pierce the wings of an aircraft. Convinced of the relevance of the documents used to feed the AI, he asks the latter to generate the quality control rules to be implemented. The Technical Department has recently adopted modifications to the drilling plan, but the employee is unaware of this. Not having access to the latest reference documents, the AI provides an erroneous result, leading the employee to establish an inept procedure which causes unjustified rejections.

A personal accident occurs in a factory, a complaint is made.

In the design office of an industrial company, a project manager asks an AI agent for help in preparing a document describing assembly procedures for a new product. The agent does not have the latest information on assembly recommendations for certain parts produced by foreign suppliers. She produces a document tainted with errors. An accident ensues in the factory, endangering the life of an employee on the assembly line.

The risks mentioned above are not new. They are inherent to the exploitation (digital or otherwise) of a documentary heritage, potentially marred by errors. But the use of AI adds a difficulty that it is essential not to overlook SAA: Self-confidence, Accuracy and Assertiveness. These characteristics are common to all chatbots on the market. Our AIs are very assertive in the way they answer our questions (prompts). More so as we tend to consider that these digital solutions can "understand" our queries. As a result, we believe them to be infallible.

A baffling Self-confidence

When a chatbot is asked a question, it responds with unwavering confidence, whatever the question. This self-confidence, which results from the way these models are trained to produce fluent, well-constructed sentences, gives users the impression of interacting with an entity with complete mastery of its subject. Yet this assurance is deceptive: LLMs possess neither genuine understanding nor awareness of the veracity of their assertions.

Accurate answers that misleadingly reinforce credibility

The answers provided by RAG-powered generative AIs are not only confident, they're also highly accurate. They include detailed facts, structured references and a technical vocabulary that reinforces their apparent reliability. Once again, this phenomenon is the result of a purely statistical mechanism: the models implemented produce sentences based on the probability of word placement in relation to each other, and not on objective validation of the information provided. AI does not master semantics or ontologies, on which it seems to rely. This apparent precision is not synonymous with accuracy. Even when an answer refers to information extracted from the critical documents used to feed it, the rendering of this information must be treated with caution.

A misleading assertiveness

LLMs never express the slightest doubt. Unlike humans, who can qualify their statements or make explicit a degree of uncertainty, these models systematically present their answers as established truths. The absence of any sign of doubt or contextualization of the answer's limits contributes to the illusion of infallible knowledge.

The use of documentary heritage by AI does not represent a new risk in itself: any database or documentary system can contain errors. However, the use of LLMs adds an additional difficulty: the biased reformulation and artificial increase in credibility of erroneous information.

The experience feedback presented above demonstrates the need to put in place a rigorous framework for managing critical documents, to enable AI to exploit reliable and consistent data. More than ever, the quality and relevance of documents are strategic levers in the service of business competitiveness.

While the SAA phenomenon described here is a risk, it can also be exploited constructively. The self-confidence demonstrated by LLMs not only masks possible errors, but it also has a revealing effect. When generative AI is used within a structured framework, it can highlight inconsistencies, contradictions and errors present in source documents, offering companies a new approach to auditing and improving the quality of their document repositories.

Organizations operating in highly regulated sectors (finance, industry, healthcare, legal…) rely on voluminous, constantly changing bodies of documentation. A policy, procedure or regulation may be updated without all related documents being modified accordingly. Internal discrepancies may emerge, making it difficult to maintain overall consistency.

It's here that LLM+RAG models bring unprecedented added value. Their modus operandi is based on retrieving information from a defined corpus (RAG) and reformulating it in natural language. This ability to link disparate elements and synthesize information enables the rapid identification of contradictions or flaws that humans, overwhelmed by the sheer volume of documents, might not immediately notice.

AI can highlight errors and inconsistencies to test the robustness of a document corpus. Rather than simply being a tool for generating answers, generative AI can be used as an anomaly detector and document verification assistant.

Of course, AI does not replace human expertise in auditing critical documents. It is a powerful pre-analysis tool, capable of speeding up the detection of errors and inconsistencies. By integrating AI into a supervised process, experts can concentrate on validating the alerts issued, rather than on the tedious rereading of thousands of pages.

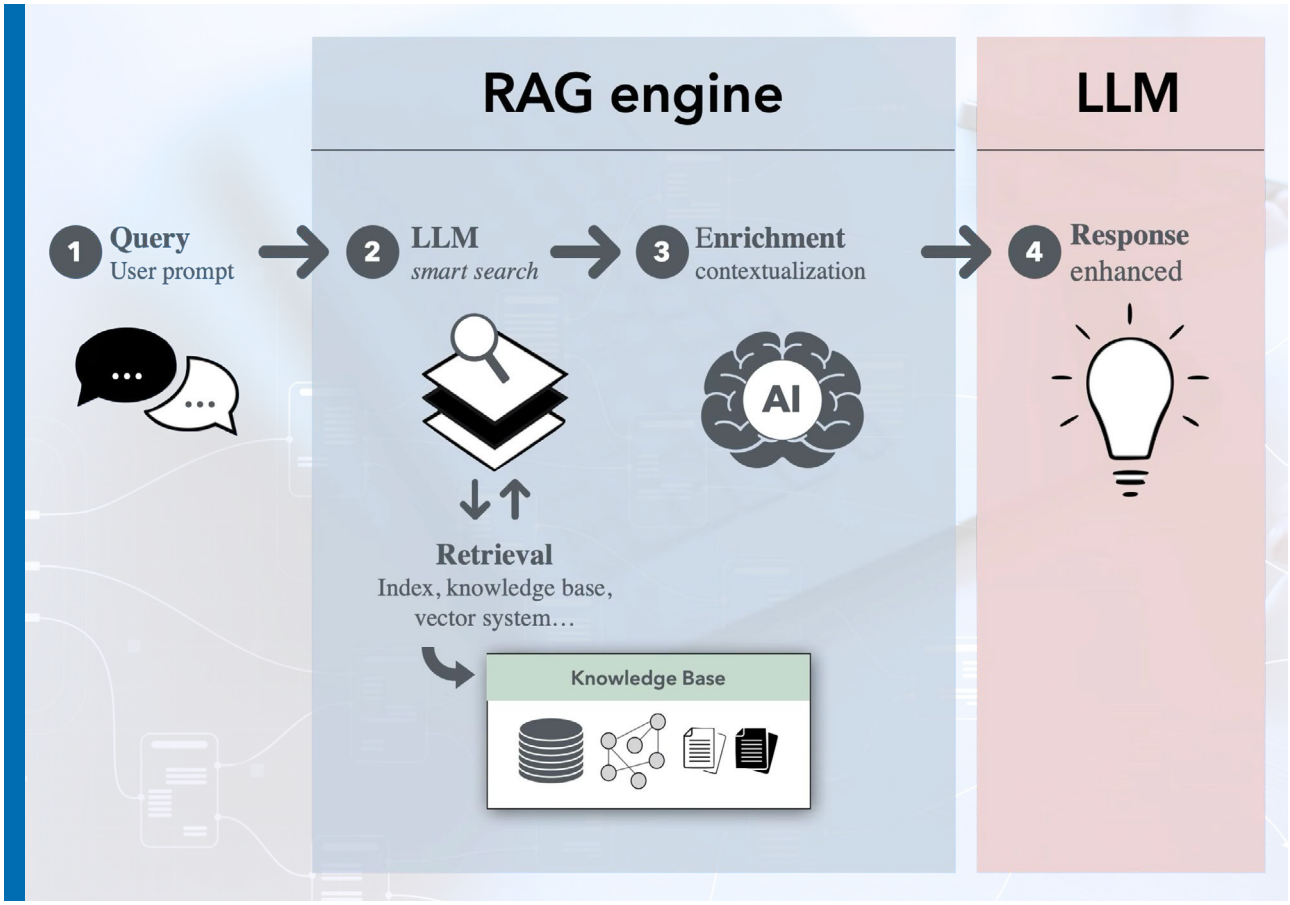

Overview of user request processing within an LLM-RAG architecture

1 + 1 = 3: Generative AI and critical document management platform in the cloud

Experts agree that, without robust document management, AI remains a risky bet. This is where AODocs comes in. Since its creation, this cloud-based document management platform has enabled companies to structure, secure and intelligently exploit their critical document assets.

Analysts make no mistake when they recognize AODocs as the most advanced solution on the market for facilitating the efficient and secure integration of AI.

> *The AODocs document management platform, natively designed for the cloud, is the ideal solution for deploying applications based on generative AI that are relevant, scalable and reliable.*

More specifically, AODocs:

✓ Is build on a Cloud native architecture, allowing customers to keep their documents in their own cloud storage and integrate with the LLM of their choice

✓ Contributes to the automation of certain processes, thanks to the use of AI combined with powerful document workflow functions

✓ Guarantees systematic access to the right document, at the right time, by the right person

✓ Guarantees respect for the principles of confidentiality linked to the company's document assets

✓ Allows the company to choose its own document storage solution in the cloud

With AODocs, generative AI leverages a reliable, relevant and secure document base, enabling it to reveal its full potential. The risk of errors is considerably reduced, and the relevance of answers is maximized, in the user's business context.

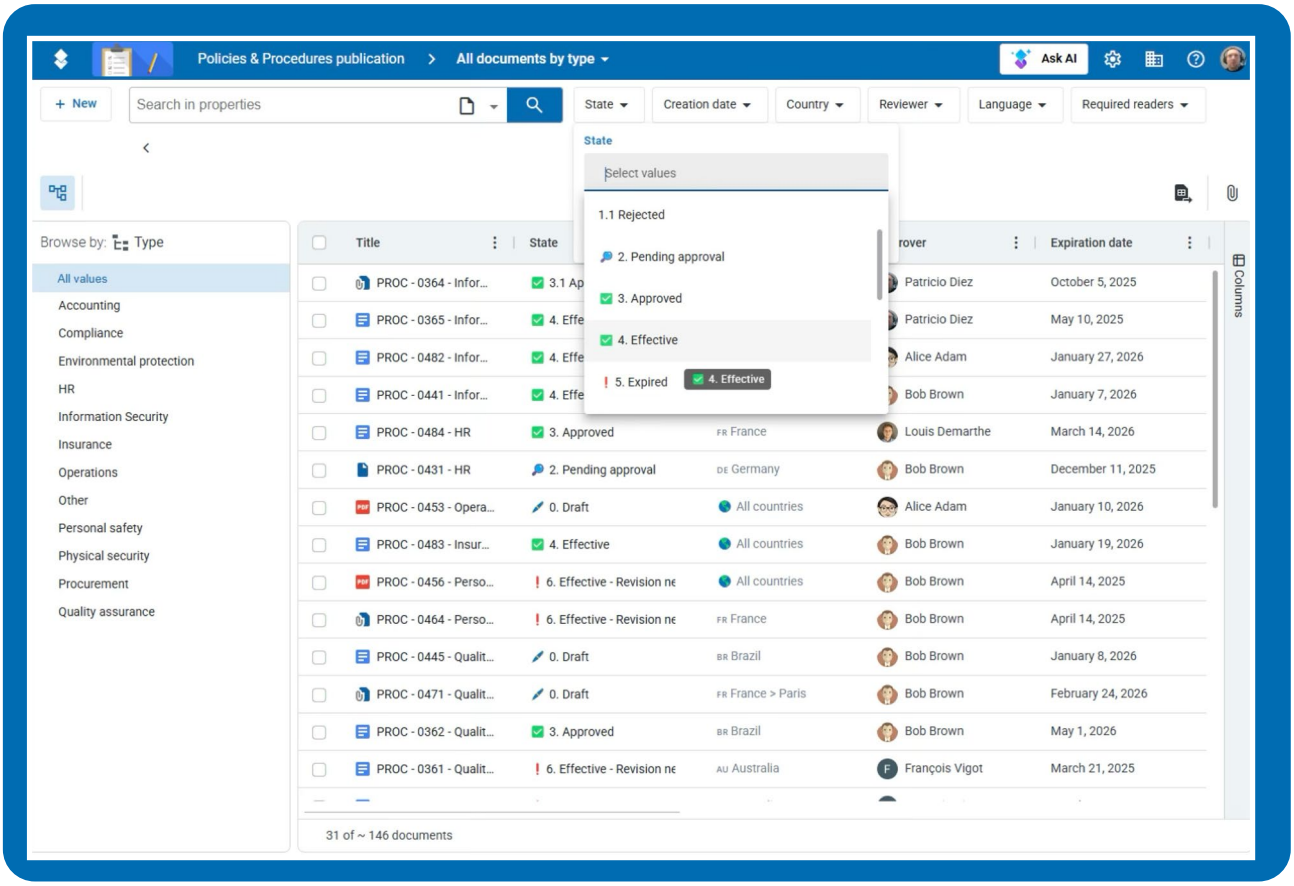

Thanks to its cloud-native architecture and remarkable flexibility, AODocs is the ideal solution for creating, rapidly deploying and operating your enterprise chatbots.

2. Critical documents: invaluable assets

Documents are at the heart of any company's business operations. Let's start with a quick overview of the concepts that characterize it.

2.1 What is a critical document?

During its activities, a company manages a wide range of documents: customer invoices, supplier quotations, calls for tender and replies to calls for tender, e-mails from customers, assembly instructions, sales presentations, study results, press articles, annual reports, safety regulations, quality management procedures, etc.

We usually classify these documents into categories, according to criteria that depend on the company's activities and organizational principles:

- Accounting, legal or human resources documents,
- Documents relating to prospects, customers, marketing and sales,
- Documents relating to production, procurement and quality control,
- Documents useful for company management,
- Technical documents (specifications, maintenance sheets, user manuals, etc.),
- Etc.

Not all documents produced or used by a company have the same value or strategic importance. Some are merely informative, while others are essential to the smooth running of the organization. These levels of criticality vary according to the company's structure, processes and legal or contractual obligations.

Critical documents are those that contain information vital to the execution of a company's key activities, particularly those linked to employee safety in the workplace, production, sales, accounting, regulatory compliance and customer service. They are the pillars of internal governance and external relations, and their accuracy is a sine qua non for controlling operational and financial risks.

The criticality of a document is assessed in terms of the commitments it formalizes and the potential consequences of an error in its content:

- An incorrect financial document (invoice, contract) can lead to disputes with a customer or supplier, affecting the company's cash flow or reputation,
- A poorly drafted safety procedure can expose the organization to major legal and financial risks in the event of an accident in the workplace, up to and including criminal sanctions,
- Imprecise or erroneous technical documentation can lead to costly product recalls, damaging the reliability perceived by customers and partners.

Every error can have profound repercussions on the company, its compliance, the quality of its relations with its stakeholders, and its

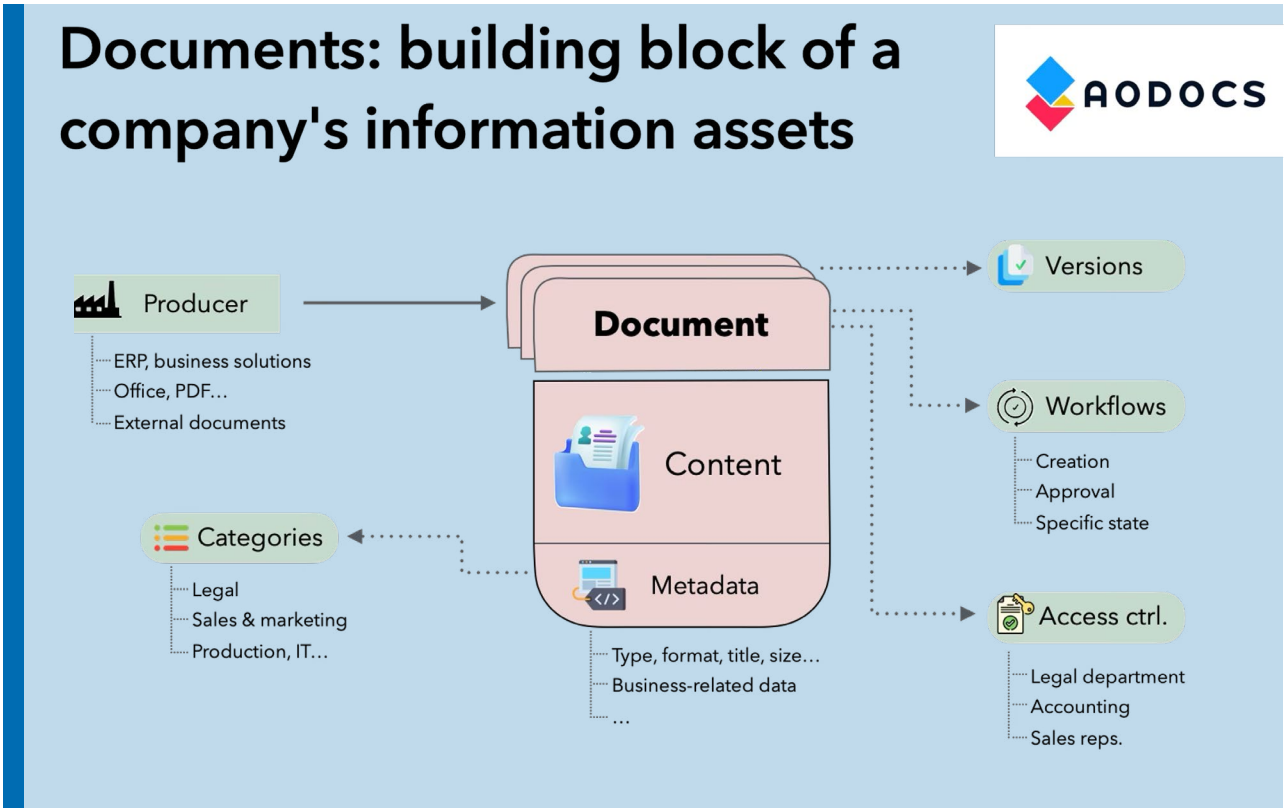

Main characteristics of a document

overall performance. That's why the rigorous management of critical documents is a strategic challenge, requiring appropriate solutions to guarantee their reliability, traceability and accessibility in all circumstances.

2.1.1 Digitized documents

In a digital environment, a document refers both to its ***content*** and to all the information that characterizes its existence within the company, known as ***metadata***: name, categories, author(s), creation date, current status, size, life-cycle history and other information that the company wishes to associate with it.

All the critical documents managed by a company bear witness to its identity, values, culture, strategy, operational activities and, more generally, the way it conducts its business.

As a company grows and develops in its markets, so does the volume of documents it receives, manages and issues. This increase in document volume invariably translates into an increase in the number of human errors committed in the manual processing of documents, leading to an explosion in induced costs and requiring the company to equip itself with a solution to manage its critical documents, such as AODocs. This is even more critical when a company's operations are based on manual document processing, or processing that is only partially automated.

2.1.2 Numerous document producers

The introduction of business applications, notably ERP (such as Salesforce, SAP, Oracle Applications, Autodesk...), the use of office suites by company employees and the deployment of collaborative or shared workspaces (Microsoft Office 365, Google Workspace, etc.) are leading to an explosion in the volume of documents managed by the company. Some are produced and hosted by

business solutions, while others are stored in the cloud, within collaborative workspaces. The diversity of document sources - often referred to as producers - combined with the heterogeneity of document formats (spreadsheets, text files, images and diagrams, PDF documents, etc.), as well as the multiplicity of physical storage environments for critical documents, make finding up-to-date, reliable and relevant information a complex task.

2.1.3 Document access permissions and confidentiality rules

Each company defines its own operating and organizational rules for conducting and managing its activities. It assigns responsibilities and roles to certain individuals, whether inside (employees) or outside (service providers, consultants, auditors, etc.) the company. These organizational principles are systematically translated into a set of access permissions (to critical documents) that the company assigns to certain individuals.

Of course, depending on the tasks they carry out - and therefore, by extension, on their role and responsibilities - individuals producing, using or modifying documents managed by the company must have access permissions compatible with their prerogatives. This rule may seem obvious. We shall see, however, that it is more difficult to implement once we decide to use generative AIs.

2.2 Document life cycle: creation, enrichment, validation, transfer and sharing, storage, archiving, destruction

Some critical documents are produced by the company itself, such as responses to invitations to tender, contract proposals, or notices describing the workflow of a production unit.

Some are intended for transmission to an external entity, such as a government department, a customer or another organization. Others have been produced outside the company and need to be brought in for processing. These may include invoices from suppliers, invitations to tender, requests from customers, regulatory changes, and so on.

Finally, some documents are intended to facilitate cooperation between several teams, with certain stakeholders even located outside the company: suppliers, consultants, controllers, service providers, etc.

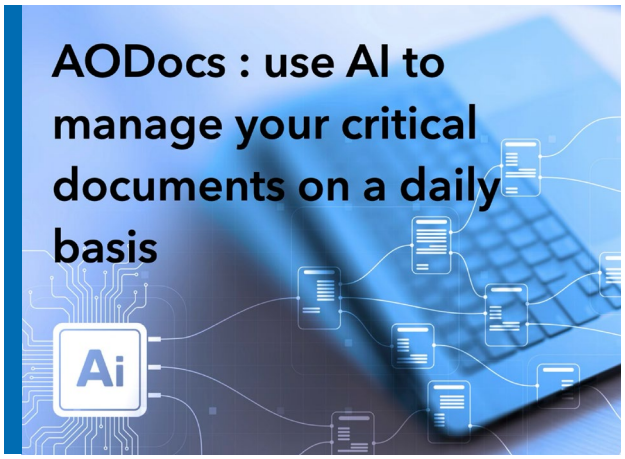

Exploiting the tremendous operational potential of documents

Just as a cell develops at the heart of a living organism, a document is characterized by its own ***life cycle***. This is represented in the form of a directed graph, the nodes of which correspond to the states that the document in question can assume. Experts use the term ***workflow*** to describe:

- The life cycle of a document, i.e. the list of states it can accept,
- The conditions that drive any change of state,
- The actions to be taken when such a change occurs: sending an e-mail, generating an alert, reaching a deadline, generating an electronic signature, or notifying a company department.

2.2.1 State-to-state transitions

Consider a critical document, *Dc*. Its current state is designated E_1. Several types of state transitions are likely to characterize *Dc*.

A. Human intervention (manual)

This is the most trivial transition, relying on human intelligence, informed by the company's business processes and governance rules.

Clearly, a duly empowered individual manually modifies the state of *Dc*. The AODocs platform presents him with the various possible states, resulting from the document's initial state (E_1) and the access permissions assigned to the individual concerned in relation to *Dc*.

B. Time-based trigger

When document *Dc* assumes state E_1, it is assigned a date stamp *T* and a duration *d* (5 minutes, 2 hours, 3 days, 15 years, etc.). As soon as the duration *d* has elapsed (*T* + *d*), *Dc* automatically changes state.

This approach makes it possible, for example, to automate the management of processing priorities for certain documents, or to temporize processing constraints for a given document. It offers the advantage of automating process dynamics, reducing the risk of human error and making the company more reactive.

C. Metadata value conditions

Rules can be defined based on the metadata values associated with document *Dc*. They may relate to document characteristics, such as a naming convention, or to metadata automatically extracted by the AODocs platform from document content, such as the amount of a customer order.

The advantage of this approach is that several individuals can act jointly on the metadata associated with Dc, without being explicitly informed of the state transition conditions applicable to the latter.

2.2.2 Manage multiple versions of a document

The existence of several versions of the same document adds to the complexity described

above. Determining, at any given moment, which version of a document is the most recent or the most complete, is often a nightmare, especially with the explosion of collaborative work environments; employees are quick to duplicate a document without maintaining a link between its different versions.

The uncontrolled duplication of critical documents is a real challenge for most companies. It makes version control difficult and considerably increases the risk of errors. Each copy in circulation becomes a potential source of discrepancy, liable to be modified, stored or used independently, without synchronized updating with the reference version.

Critical documents are particularly exposed to this phenomenon, as they are frequently shared and used by several teams. This dispersal of use also leads to duplication of the same resource on several digital systems.

To avoid the risks associated with the proliferation of versions of a critical document, it is essential to implement best practices and appropriate tools. Of course, the use of a single repository, capable of centralizing documents in a secure document management system, is essential. This is key as the DMS can ensure a single, automatic update of the versions of the documents it references, while restricting downloads when not strictly necessary.

Educating employees to prefer sharing by link rather than attachment is also very effective, particularly when sending e-mails. Reduce the risks of “shadow IT” by providing appropriate solutions, including official and secure collaboration tools, to prevent users from resorting to uncontrolled platforms. Finally, it is often essential to apply document lifecycle management rules to automatically delete or archive obsolete versions.

Mechanisms favoring the proliferation of copies of critical documents:

E-mail • Sending e-mails containing critical documents as attachments is one of the first sources of massive duplication. Each time an e-mail is sent, an independent copy of the file is created in the recipient's inbox. If several people reply to or forward this e-mail, new versions are created, sometimes modified locally, the changes made rarely being reflected in the original document.

Local printing and saving • A document printed or saved on the local hard disk associated with an employee's workstation is not traceable. It may be considered as a reference, even though a more recent version exists elsewhere. Such "dead copies" significantly increase the risk of using obsolete information.

Digital cohabitation • In many companies, a critical document may transit between several platforms. An employee downloads a file from a document management system with the aim of modifying it. They import it into Salesforce or another CRM tool to fill in missing data. He then saves this updated version of the document on his workstation, before reloading it into an internal tool to share it with colleagues. Each of these transitions creates a new instance of the document, multiplying versions and complicating the traceability of modifications.

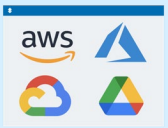

Shadow IT • When the company's IT department doesn't offer official solutions tailored to users' needs, they tend to resort to unapproved tools such as Dropbox, Google Drive, WeTransfer and so on. They even resort to simple USB sticks to store and share document resources. These practices get out of hand and encourage the dispersal of critical documents outside the company's secure systems, creating greater risks in terms of confidentiality, compliance and updating of data.

3. AI and critical documents: a key challenge

AODocs guarantees the success of enterprise chatbot projects

Artificial intelligence, in particular generative AI, is creating remarkable opportunities to boost competitiveness and operational efficiency.

Understanding the capabilities and characteristics of AI is essential for identifying the areas of progress it can bring about within a company.

When it comes to deploying chatbots in production, 2 essential characteristics of gene–rative AI need to be highlighted:

- Reduced information overload,
- Lack of discernment in the choice of documents used to feed the AI.

3.1 AI reduces document overload and transcends information dissemination

The multiplication of documents and their dispersal across various systems (internal servers, cloud, business solutions, collaborative spaces, etc.) is a major challenge for many companies. Finding relevant information can become a complex task, with uncertain results. This reality is exacerbated by the existence of duplicates, obsolete versions and poorly identified or unidentified document silos.

Generative AI, powered by a coherent subset of critical documents, introduces a radical paradigm shift. Unlike humans, who must manually navigate complex document trees, or juggle multiple search tools, AI frees itself from the physical and organizational constraints of digital data storage. What's more, unlike the vast majority of humans, AI is capable of processing large volumes of information with no marginal increase in error rates.

It synthesizes information in real time, extracts the essence from dispersed documents and provides immediate answers, without the user having to worry about the location or format of the sources used.

Generative AI is not concerned with the rele–vance of the information it processes, or even with the relevance of the answers it produces. Caricaturally, and using a crude anthropomorphism, we should bear in mind

that generative AI has no free will: it ignores what is true, what is admissible, what is acceptable... It simply digests the information it is provided with (RAG) and renders it in natural language (LLM), understandable and usable by a human user. It has no way of determining whether information is obsolete, true or ethical.

While we still think in terms of files, directories, folders and compartmentalized databa–ses, AI operates in a fluid, highly meshed and interconnected universe, where only relevant data counts.

If the sources exploited by AI are rigorously selected and kept up to date, this new paradigm opens the way to unprecedented efficiency for company employees.

3.2 The AI uses all available documents

When it must process a user request in a given context, generative AI makes no distinction between the documents it is supposed to exploit and those it should not consider. Unlike a human, who is naturally aware of the confidentiality rules and sensitivity levels of the information it consults, AI acts indiscrimi–nately on the whole range of documentary resources made available to it.

It is imperative for a company to put in place a strict framework for managing access rights to its critical documents, otherwise it will have to assume the risk of a chatbot using confidential, obsolete or irrelevant information in inappropriate contexts, leading to unintentional and potentially serious breaches of security and compliance rules.

Similarly, providing generative AI with the “right” documents is essential to guarantee the quality and relevance of responses. The proliferation of versions of the same document is a major issue for most organizations.

Consider the caricatural example of a conversational assistant based on generative AI with unrestricted access to all internal company documents. A user could then, voluntarily or not, obtain and exploit sensitive data, such as patents, strategic contracts, accounting data, employees’ personal data and so on.

In short, AI feeds on the documentary resources to which it has access, without any discernment or strategy of preference. It does not spontaneously detect risks linked to possible information leaks or to the quality of the information extracted from the documents made available to it: it produces answers without conscience or filter.
The operational risks resulting from this situation require companies to implement strict governance of access to critical document resource:

- Definition and rigorous application of access permissions to critical documents,
- Segmentation of document sources according to their level of sensitivity,
- Regular auditing and control of access granted to AI to avoid any drift or unintentional leakage of information,
- Centralized management of document versions,
- Assiduous control of reference versions of critical documents.

Without these precautions, AI potentially becomes an uncontrolled vector for the disclosure of sensitive or erroneous data, jeopardizing a company’s regulatory compliance, penalizing its production system and damaging the quality of its customer relations.

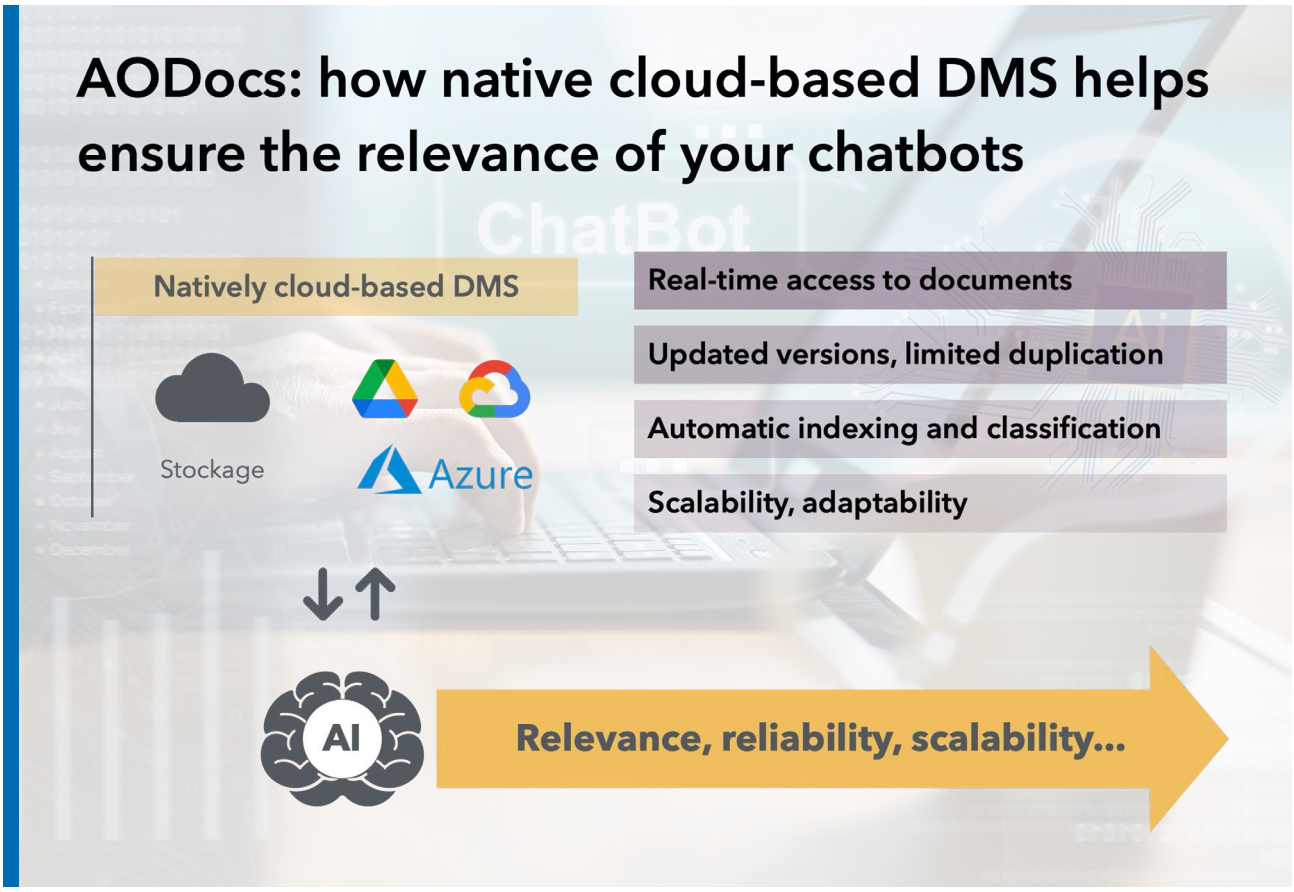

Cloud storage brings flexibility and guaranteed relevance to chatbot projects

3.3 Flexibility and scalability of a DMS natively compatible with cloud-based document storaged

It is important to use a document management system that is natively compatible with cloud-based document storage.

The effectiveness of a RAG model coupled with an LLM depends on the quality and accessibility of the source documents. To guarantee accurate and relevant answers, AI must be able to interrogate reliable, up-to-date and contextualized data. In their environment, most companies use the cloud to store and share large volumes of documents. Using a DMS natively connected to the main cloud platforms, such as Google Drive, Google Cloud Storage or Azure Blob Storage, offers a host of advantages.

A cloud-native DMS guarantees successful implementation of a RAG architecture:

Always up to date

AI queries always up-to-date data, without duplication or latency. Synchronization is instantaneous, and up-to-date versions of documents are perfectly identified.

Optimized search

The DMS designed for the cloud is capable of classifying and enriching documents with metadata that will improve the relevance and performance of information retrieval.

Easier scaling

The cloud enables companies to rely on a scalable, high-performance infrastructure to manage large volumes of documents, while guaranteeing a high level of availability.

A DMS natively compatible with a document architecture deployed in the cloud, like AODocs, avoids the need for costly and insecure workarounds, such as temporary local copies of documents, non-optimized APIs or manual data import/export processes. These approaches undermine the solution's performance, increase the risk of bias and significantly complicate document management.

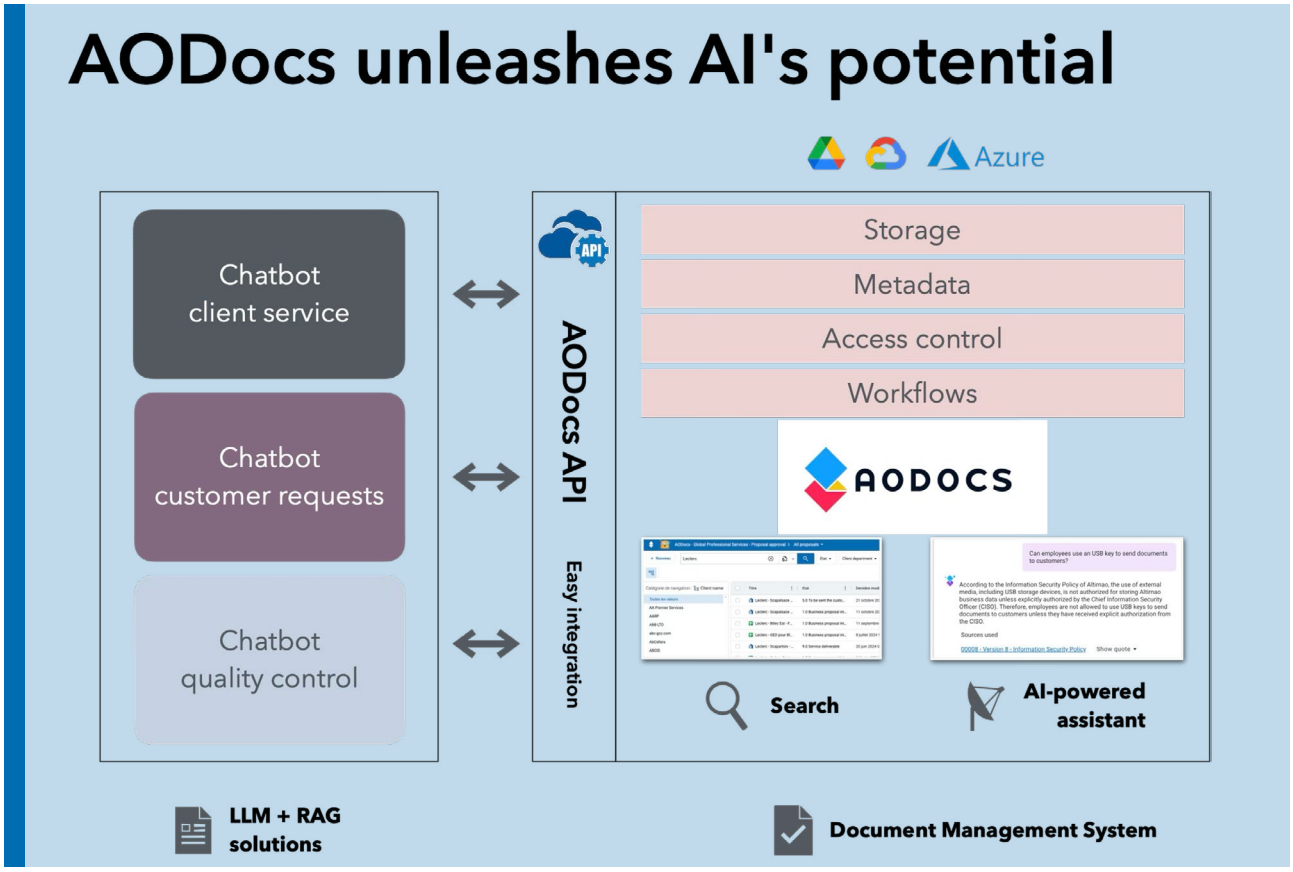

AODocs unleashes the potential of intelligent chatbots

4. Case studies: operational risks

Here are a few real-life examples, anonymized for obvious reasons of confidentiality.

4.1 Users of a legal chatbot mistakenly send sensitive information to an external service provider

Context

A major company in the financial sector implemented a chatbot designed to facilitate the drafting and validation of contracts, based on a rich documentary history.

AI deployment objective

The company wanted to reduce the time spent by legal teams checking contractual clauses and speed up the signing of new agreements.

Implementation

Mistakes were made when building the file containing the documents to be fed to AI in an LLM+RAG architecture:

- Contracts not finalized, under negotiation or not yet signed,
- Some internal documents with no contractual value, such as informal exchanges between lawyers and legal experts.

Assessment

A consultant working within the company, wishing to call on the services of an external company, asked the legal chatbot at his disposal whether a confidentiality agreement existed between his company and the external company. The chatbot replied (with self-confidence): “Yes, I did find an NDA on behalf of this company”. This reassured the consultant, who had no legal expertise, and he began sharing confidential information with the external company’s employees.

Consequences

The contract found by the chatbot had not yet been signed: the clauses contained in the confidentiality agreement were still being negotiated. After a few weeks, it was decided not to sign an agreement with this external service provider. A few months later, however, the company’s legal department became aware of leaks of sensitive information. Discussions took place with the legal department of the external company in question, which asserted that it was not the source of the leaks.

Conclusion

Without rigorous control of the documentary sources used to feed the RAG engine, legal AI can become a risk rather than an asset for the company running it. Strong document governance and human supervision are more essential than ever.

4.2 An AI used by a customer service department provides incorrect answers

Context

A player in the mass retail sector observed a significant increase in the volume of requests from its customers, processed manually by its employees.

AI deployment objective

Automate the management of customer requests: product returns, refund requests, warranty information, etc.

Implementation

To feed this AI, several thousand internal documents were used, including:

- Internal FAQs, sometimes obsolete or incomplete,
- Informal exchanges between customer advisors,
- Technical manuals, some corresponding to earlier versions of the products marketed.

Assessment

The AI began to provide incorrect information on the return and warranty conditions of certain products. It referred to products no longer marketed by the company, and provided inconsistent answers to over 30% of customer queries.

Consequences

The company had to react quickly and stop the project, otherwise customer satisfaction would plummet, and support costs would soar, overloading the teams and generating significant additional costs.

The direct financial impact was assessed, in terms of refund claims or unjustified financial compensation.

Conclusion

AI must not be fed with unverified or obsolete data. The quality of the documents used in its training is crucial to guaranteeing relevant and reliable responses.

A system of continuous validation backed up using a modern platform for managing critical documents, and regular auditing of the AI's responses help to avoid costly errors.

5. IA and Governance of critical documents: challenges and best practices

When we ask an employee who uses a solution based on generative AI what precisely he or she expects from it, the first response we get is: "That the solution provides me with precise information that I can rely on, or that is simple to control and saves me time". This response highlights a legitimate expectation of **precision**, **reliability**, **traceability** and **speed**. Interviews with other employee profiles, individuals working in a variety of roles at different levels of responsibility, reveal a few more key words: **confidentiality**, **adaptation**, **compliance**.

5.1 AI-based solution evaluation criteria

Reliability

In a professional setting where contracts, production or sales data, complex regulations, international standards, market research, etc. are handled, an AI must not only produce plausible answers, but also verified, usable and correct information. Ideally, AI should be able to associate its answers with a self-assessment of a level of confidence. This "confidence index" would enable an application or a user to decide whether or not to carry out verifications, e.g. manually, of the proposed response elements, before considering them to be truly exploitable.

Relevance

The relevance of a solution based on generative AI refers to its ability to provide useful answers, i.e., adapted to the user's needs in a given context. There's nothing trivial about this definition, despite the fact that, like a set of Russian dolls, it is based on other keywords whose definitions raise just as many questions.

Taking a simplified approach, the usefulness of the solution could be defined as a Boolean value (true or false), and the user would simply be asked to evaluate whether or not the answer obtained is useful to him or her. A simple couple of thumbs up or down would suffice. The usefulness objective of the AI-based solution would be reached as soon as the percentage of affirmative responses exceeded a certain threshold.

Contextualization and adaptation

The solution needs to understand the specific nature of the user's field of interest, to adjust its answers to the user's business needs: legal, commercial, technical, medical, etc. Each answer must also respect the specificity of the user's field of interest. Each response must also respect the level of expertise expected (deduced) or literally expressed by the user.

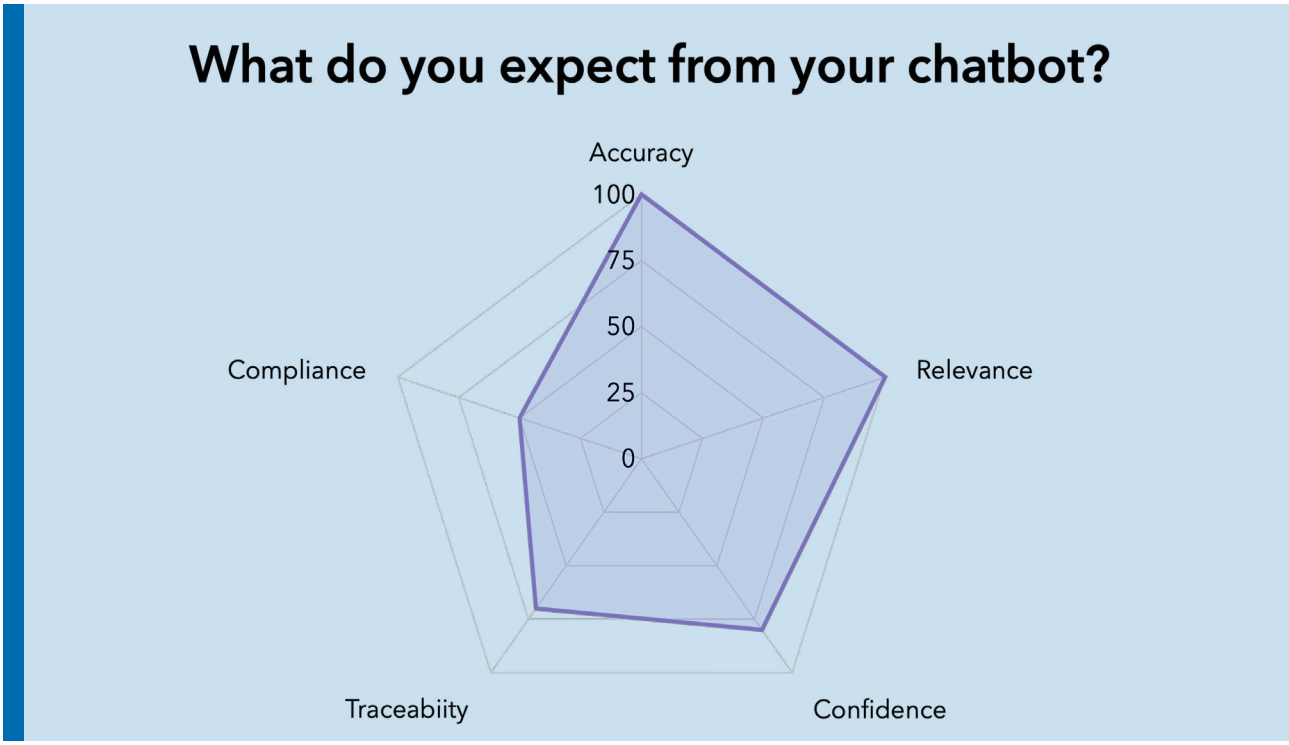

Defining the relevance objectives of a solution based on generative AI

Traceability and auditability

It is absolutely essential to be able, at any time, to justify the decisions made by the AI, to meet regulatory and compliance requirements, but also to be able to carry out rapid checks on the relevance of the proposed answers.

Security and confidentiality

Critical documents often contain sensitive data (intellectual property, strategic information, key figures). We need to ensure that AI complies with access and data protection rules. While it's not up to the AI to block any risk of sensitive information leaking out, the information architecture implemented upstream of the solution running the AI must guarantee the confidentiality of the data it exposes to the LLM.

Compliance

It is often desirable for the AI-powered solution to guarantee the compliance of its responses with regulations such as GDPR, a QMS, a specific standard, etc. The company must clearly formulate its requirements in this area, and compliance with these must be an integral part of the evaluation process for a prototype solution, before any production launch is considered.

Speed and efficiency

The solution must provide sufficiently rapid responses. The right balance between processing time and response quality is essential to optimize the user experience.

5.2 Evaluating the performance of a chatbot prototype

All the experts agree that AI is a formidable tool, an ideal medium for building new solutions, especially powerful chatbots, capable of responding quickly and appropriately to requests expressed in plain language by users.

But AI is not a magic solution, an off-the-shelf product that you simply plug in to benefit from its contributions. The feedback we have today, covering all sectors of activity, highlights the "potential" power of LLM, combined with other technological foundations as required. To unlock this potential and reduce the risks involved, a few recommendations should be borne in mind.

5.2.1 Define and document expectations

A company wishing to implement a chatbot using generative AI needs to decide on all the criteria mentioned above: relevance, confidentiality, compliance, etc. It is imperative to define its objectives in terms of the relevance of the solution envisaged, before considering any production.

It is imperative to define its objectives in terms of the suitability of the envisaged solution, before considering any production launch. This will avoid many implementation errors, while maximizing the chances of project success. Documenting expectations in a shared document, setting detailed quantitative and qualitative objectives, and presenting these to the teams concerned as well as to management bodies, all help to avoid many pitfalls.

5.2.2 Rely on the right documents

As we have seen, the cost to the company of implementing an AI trained on incomplete, erroneous or obsolete documents, or on documents with an inappropriate level of confidentiality, can be exorbitant.

5.2.3 Adopting the right production architecture

The elements presented so far are not, of course, limited to the context of a prototype project. They are part of a global dynamic of monitoring and operation.

As soon as a chatbot is put into production, a few precautions must be taken to guarantee its lasting usefulness from the point of view of the teams involved:

- Ensure compliance with the qualitative and quantitative objectives defined at the outset, and revised on a regular basis in the light of user feedback,
- Scrupulously maintain the relevance of the documents used to feed the RAG engine,
- Carry out regular assessments to consider corrective measures, particularly with regard to the documentary heritage used for AI training.

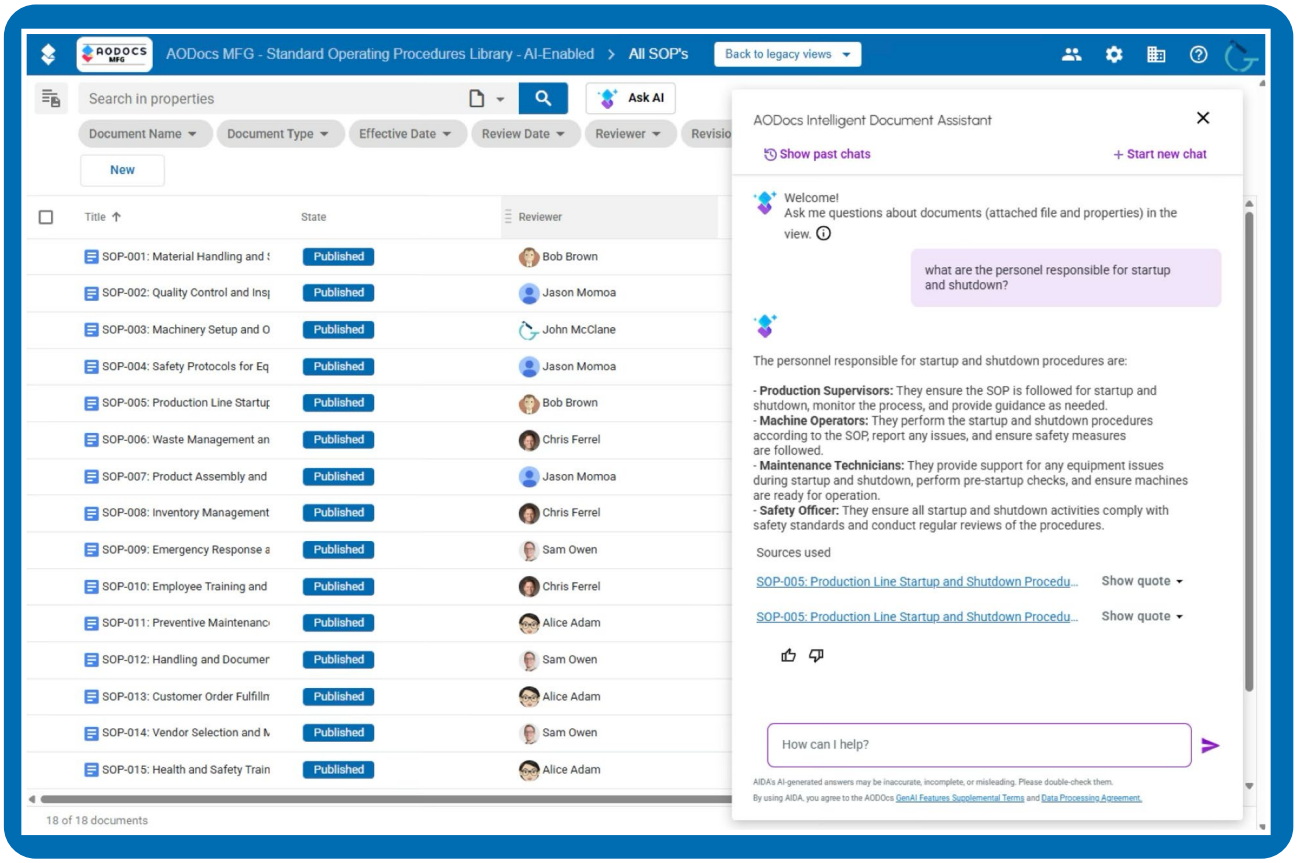

AODocs leverages AI to ensure that every new document is processed appropriately: intelligent chunking, metadata extraction, automatic summarization, classification... key functions that help improve the quality of documents exposed to the chatbot.

AODocs guarantees the relevance of your knowledge bases.

The documents proposed to the AI are the ontological foundation of the future solution.

The AODocs platform offers the advantage of guaranteeing the quality of the documentary heritage to be mobilized to feed the AI. In addition to the innumerable advantages that result from the implementation of a business-critical document management solution, AODocs is the best guarantee of success for any truly useful chatbot project.

For more information: www.aodocs.com

AODocs: the most AI-enabled DMS platform

The era of generative artificial intelligence is opening fascinating prospects for businesses. Chatbots and other virtual assistants, enriched by advanced language models (LLMs) and RAG engines, are now much more than simple automation tools. They are becoming true partners, capable of analyzing, synthesizing and rendering information with unprecedented relevance.

Companies that know how to exploit these technologies methodically will benefit from a major competitive advantage: improving certain processes, boosting operational efficiency, supporting a transformation project, reducing repetitive and time-consuming tasks, leveraging information assets, etc.

As with any technological (r)evolution, generative AI must be deployed with care. The relevance of an enterprise chatbot, however effective, reflects the quality and relevance of the documents on which it is based. AI fed with obsolete, biased or uncontrolled content risks introducing costly errors, compromising both the reliability of the information and the credibility of the company.

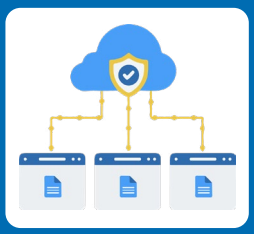

It is essential to adopt a controlled and secure approach, based on robust document governance.

AODocs plays an essential role in this equation, guaranteeing:

- Rigorous data structuring,
- Traceability and version control,
- Securing sensitive content,
- Compliance with regulatory requirements.

Successful enterprise chatbot projects

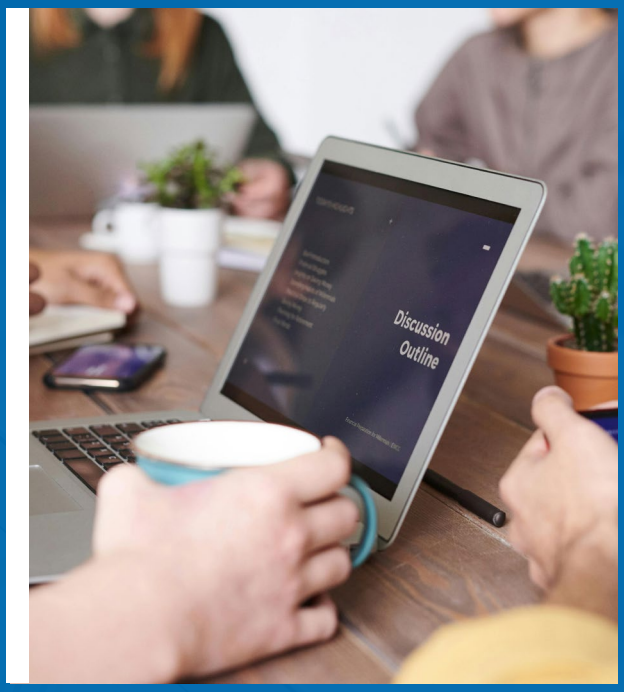

AODocs supports the major players in this revolution daily. Its agility, flexibility, cloud-based architecture and remarkable integration capabilities guarantee the success of enterprise chatbot projects.

For more information, visit www.aodocs.com

The future belongs to companies capable of integrating generative AI and document management in a complementary way. Rather than choosing between automation and control, we need to build a hybrid ecosystem, combining the power of LLM with document rigor.

By taking the right decisions today to frame and structure the use of AI, organizations can not only avoid the pitfalls resulting from chaotic, unsuitable or outdated document management, while unleashing the full potential of these technologies in the service of their operational excellence. AI is a formidable transformation accelerator. It's up to each company or organization to use it intelligently to make it a lasting asset.

AODocs goes beyond legacy document management with a cloud-native architecture built for seamless integration with LLMs and cloud-based AI. Its flexible, configurable design adapts easily to any enterprise environment, making it ideal for powering next-gen, AI-driven operations. What's more, AODocs offers an unrivalled level of flexibility and ease of configuration. This ultra-modern DMS adapts instantly to all types of configuration and operating architectures. Its ability to adapt to the evolutions characterizing the explosion in AI-based solutions makes AODocs unique on the market.

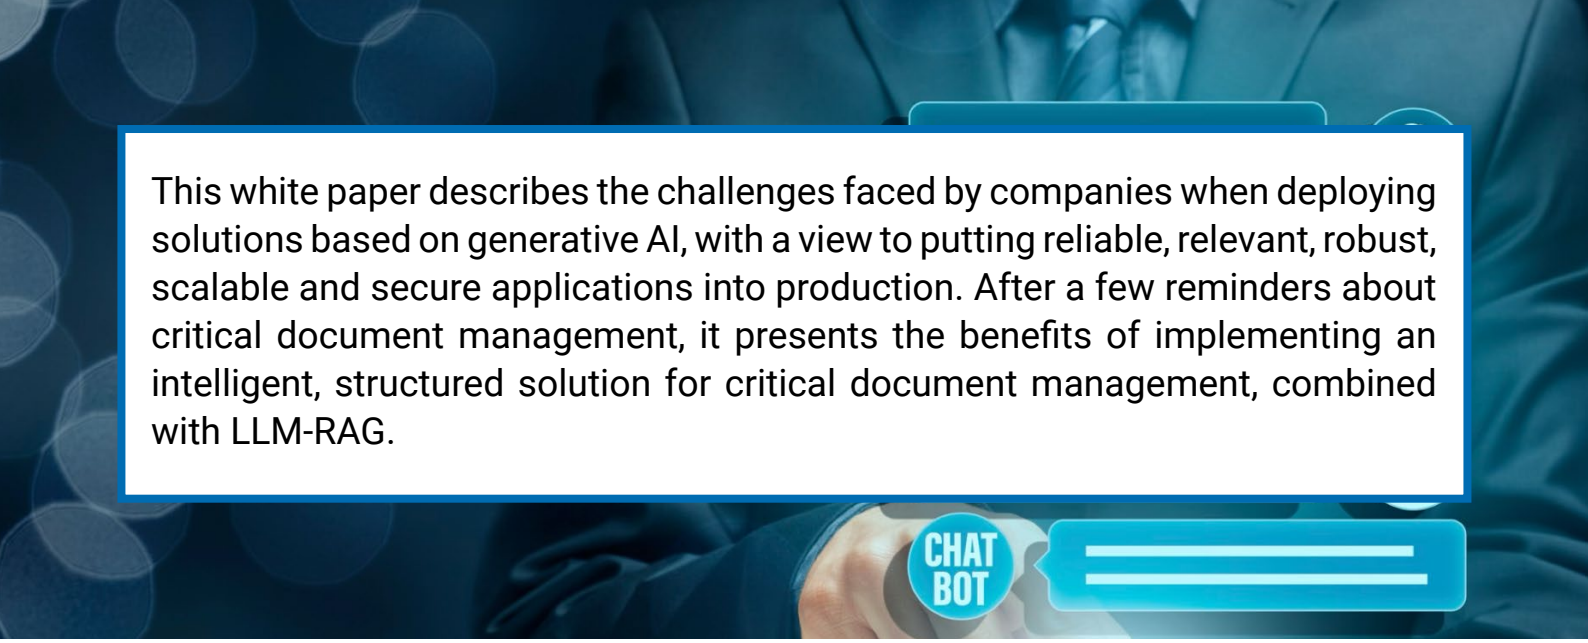

This white paper describes the challenges faced by companies when deploying solutions based on generative AI, with a view to putting reliable, relevant, robust, scalable and secure applications into production. After a few reminders about critical document management, it presents the benefits of implementing an intelligent, structured solution for critical document management, combined with LLM-RAG.

AODocs unleashes the power of AI-based solutions

Generative Artificial Intelligence (AI), driven by **Large Language Models *(LLMs)***, is radically transforming the way we interact with information. The rise of the **RAG *(Retrieval-Augmented Generation)*** approach now enables companies to exploit **their own knowledge base** to enrich the answers produced by AI, **contextualize** them and **adapt** them **to their operational needs**.

But these innovative technologies are not without risk. Analysis of the countless **chatbot prototypes** developed in recent years has highlighted **two major challenges**:

✓ **The quality and relevance of the source documents used** to guarantee reliable, consistent responses.

✓ **Compliance with confidentiality rules**, which are essential to avoid the inadvertent exposure of sensitive documents.

Poorly fed AI inevitably generates erroneous and inconsistent results, with **potentially serious consequences**: financial losses, strategic errors, even jeopardizing critical processes.

This white paper explores in depth **the opportunities and risks** involved in implementing enterprise LLMs and chatbots, detailing the challenges encountered and the solutions to be adopted. Based on real-life experience, it highlights **three key issues**:

✓ **The SAA (Self-Confidence, Accuracy, Assertiveness) phenomenon:** how LLMs artificially reinforce the credibility of sometimes erroneous answers.

✓ **Best practices** for building a reliable, up-to-date, evolving and relevant document base.

✓ **The quality guarantee of the LLM-RAG-KB triptych**, the keystone of a high-performance, reliable chatbot.

📌 **How to take advantage of generative AI while controlling its risks?**

This document offers concrete recommendations for building a robust documentation strategy, guaranteeing the success and scalability of AI-based projects.

AODOCS

AODocs
DOCUMENT MANAGEMENT
NATIVELY DESIGNED FOR THE CLOUD

www.aodocs.com
United States, France, Italy