BUSINESS DOMAIN OVERVIEW

General introduction to TRAIDA cards in the business domain. No matter how powerful a new technology is, its use is unlikely to be profitable if it doesn't sufficiently take into account the requirements of the business. This is especially true for AI, whose use cases are limitless and which raises questions about human employability.

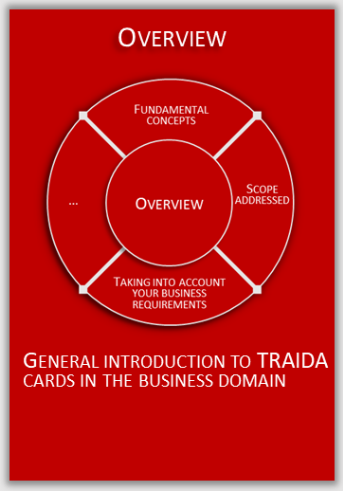

1. CONDITIONS OF SUCCESS

The TRAIDA framework (Transformative AI and Data Solutions) is based on three domains:

1. Technical (blue cards).
2. Governance (green cards).
3. Business (red cards).

The business domain is based on these three fundamental objectives that support the profitability of AI:

1. **Achieving productivity gains**. These gains address business inefficiencies by eliminating hidden costs. The productivity card of the business domain is the first to be considered for enterprise-wide AI deployment. It is used during the "Boost" phase of the AI transformation plan (see TRAIDA Treasury & Assurance card).

2. **Transforming business models**. This transformation is more secure when productivity gains are already significant. The creativity card of the business domain comes into play following the productivity card. It is used to modify business models during the "Institutionalize" phase of the AI transformation plan (see TRAIDA Treasury & Assurance card).

3. **Building human trust in AI**. Without this trust, it is difficult to scale AI within the organization, as users may harbor doubts and resistance. AI's reliability must be regularly demonstrated and monitored. The TRAIDA Trustworthiness business card addresses this issue, viewing AI as a new stakeholder to be integrated into the organization.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

Regardless of how powerful a new technology may be, if its use does not sufficiently consider business requirements, it is unlikely to be profitable. This is even more true with AI, whose use cases are limitless and which raises questions about human employability. In other words, without serious business management, AI will at best be a failure with no vital consequences for the company and at worst a black hole that will eventually destroy it. **In this drastic context, if you are discovering the impacts of AI, you would be well advised to first consult the business domain cards and the Human Resources card from the governance domain**.

FUNDAMENTAL CONCEPTS

Depending on your context, you should start with the TRAIDA card that is most relevant, for example:

- If you want to enhance your organization's capacity to accumulate knowledge: the technical card "Enterprise Knowledge Graph (EKG)" helps you start with the right technology to achieve this.
- If you need a metadata catalog that describes your existing applications and databases: the technical card "Core system data" shows you how to combine AI with knowledge graph-oriented databases to implement this catalog.
- If you are planning to deploy a regulatory repository for managing your clients: the "Treasury & Assurance" card helps formalize the return on investment by highlighting productivity and creativity gains. You would also use the technical knowledge graph cards to implement the repository and the AI governance card to support its deployment.

By selecting the most useful cards for your project, it is possible to ensure alignment between your business objectives (red cards), your technical choices (blue cards), and your governance (green cards). For example, here is a representation of the most significant alignment points in a project to implement a regulatory repository for client management:

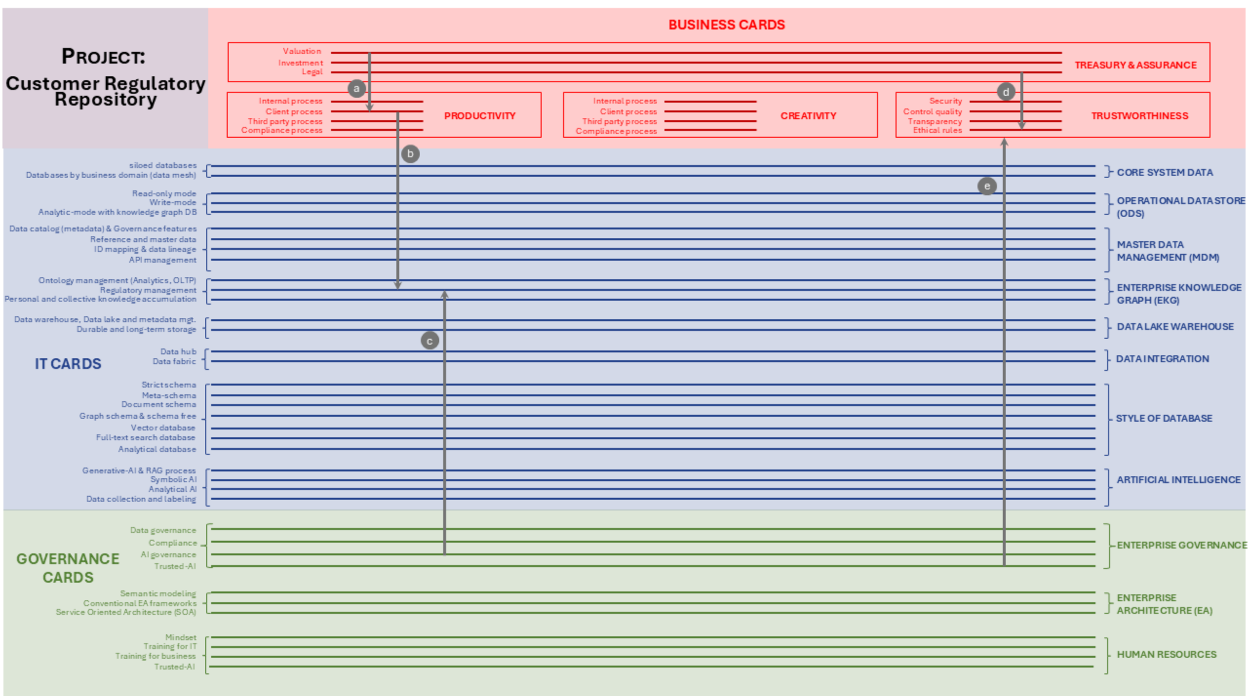

Here is the meaning of the dependency arrows between the cards:

a) The value of the project is based on productivity gains related to client management.

b) This client management benefits from a regulatory repository powered by a knowledge graph-oriented database technology, namely the EKG (Enterprise Knowledge Graph) repository.

c) AI governance is used to strengthen the control of the new repository that has been implemented.

d) A legal requirement is identified within the scope of business challenges, which refers to the ethical rules adopted by the company.

e) The contribution of trusted AI is leveraged to enhance compliance with the aforementioned ethical rules.

This schematic representation does not attempt to capture all interdependencies between the cards required for the completion of a project, but rather focuses on the most significant subset. It thus allows confirmation that a minimum alignment between business needs, technical capabilities, and governance is achieved. During project implementation, this alignment can be better documented to capitalize on knowledge and avoid divergences from the initial objectives.

TAKING INTO ACCOUNT YOUR BUSINESS REQUIREMENTS

By default, TRAIDA contains four business domain cards that are useful in all business contexts. Regardless of your organization, the integration of AI requires addressing productivity and creativity gains, managing human trust in AI, and ensuring its profitability.

However, unlike the TRAIDA cards from the technical and governance domains, the business domain cards are not always sufficient to express all the company's requirements. Indeed, **depending on your context, it may be beneficial to create your own cards to better describe your business needs**, for example, for marketing, supply chain, finance, etc. Similarly, it is possible to use the same business domain cards multiple times to account for the varying requirements of legal entities, such as a headquarters, subsidiaries, or joint ventures. For instance, the productivity card can be drafted to meet the needs of the headquarters of an international company, and then a variant can be created to address the needs of a subsidiary.

SCOPE ADDRESSED

The cards in the business domain are listed in the table below. There is no preferred reading order to follow. From an academic perspective, that is, for discovering the cards with the aim of learning general culture, the order of the cards in the table is the most advisable to follow.

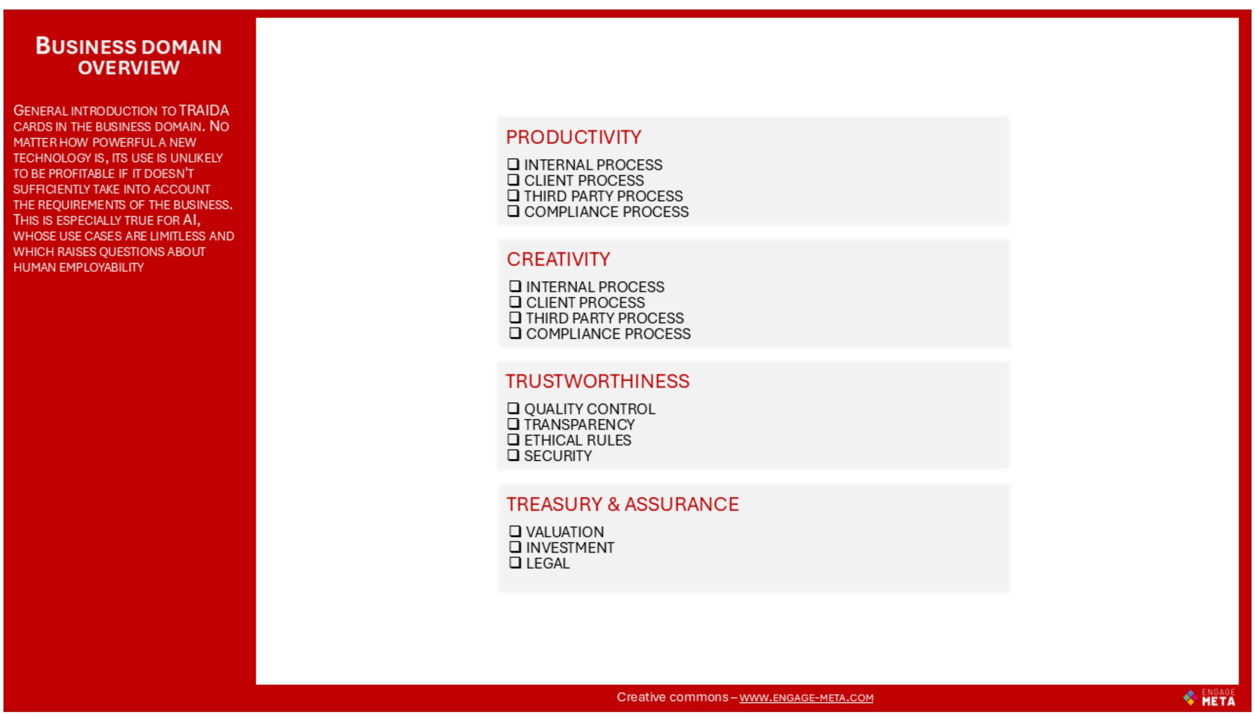

3. YOUR SITUATION & OBJECTIVES

PRODUCTIVITY

Improving productivity across all company processes is a key objective of AI. In the TRAIDA approach, achieving productivity gains is the primary objective to reach an initial return on investment from AI at the enterprise level. This is achieved through an analysis of hidden costs.

1. CONDITIONS OF SUCCESS

In the TRAIDA approach, achieving productivity gains is the primary objective to reach an initial return on investment from AI at the enterprise level. In other words, AI is first deployed to improve existing processes before being used for business model transformation. This is an important step aimed at securing initial successes and gaining experience, allowing for more creative action later on.

According to the consulting firm McKinsey (2024), 70% of tasks performed by each employee can be automated by 50% thanks to AI. This represents a significant source of productivity that does not require disrupting business models. By leveraging this productivity potential, the benefits for managing the transformation with AI are as follows:

- It does not require prior consideration of changing business models.
- In the event of failure, it does not disrupt the company's operations.
- It offers the opportunity to achieve financial gains through incremental deployments, without tunnel effects or big-bang scenarios.

These productivity gains must cover the cost of the minimal architecture necessary for AI deployment at the enterprise level (see the TRAIDA technical cards, particularly ODS, MDM, and EKG). To recall, the goal is to set up a semantic platform from the deployment of the first AI use case. Since the cost of this architecture is added to that of the initial use cases, it is important for it to become profitable quickly.

Let’s take the example of a company starting its transformation in this way:

- An impact study shows that AI will save two workdays per employee. With 10 employees, each with an average monthly salary of 5,000 euros, the total payroll is 600,000 euros per year. The estimated productivity gain is 60,000 euros per year, or 240,000 euros over four years. This amount is allocated for implementing the first version of the minimal viable architecture for AI.
- The workload saved by this AI exceeds 200 days per year. This productivity gain will enable team reorganization and increase value creation (see the TRAIDA business card for Creativity).
- Once in place, the semantic platform serves as a springboard to quickly deploy additional AI and data governance mechanisms, thus adding other use cases that will target both productivity gains and creativity in business models.
- Before committing this 240,000-euro budget, a decision-making dossier demonstrates the reality of the expected gains and proposes a roadmap with intermediate results. An initial release of 20% of the financial resources is used to develop an AI prototype. Thus, the initial commitment of 48,000 euros represents the maximum financial risk to confirm that the business and technical

constraints are well understood by stakeholders. Once the prototype succeeds, the remaining budget of 192,000 euros is released to continue the implementation.

The transformation with AI is therefore initiated by pursuing productivity gains before even starting value creation projects. This approach is important to avoid the trap of unprofitable AI projects that could lead the company into an AI winter.

The socio-economic approach

To benefit from productivity gains, hidden costs in work processes are reduced or eliminated through AI. To identify them, it is useful to rely on the socio-economic approach of the ISEOR school (*), which provides a classification:

- Quality-related extra costs: reduction of errors; production defects.
- Non-productivity extra costs: poor resource utilization; time loss.
- Absenteeism-related extra costs: unplanned absences; difficulty in replacing and reorganizing.
- Turnover-related extra costs: loss of knowledge; loss of motivation.
- Workplace accident-related extra costs: lack of employee information; poor practices.
- Social climate-related extra costs: conflicts; lack of communication.
- Etc.

(*) Socio-Economic Approach to Management (SEAM) - ISEOR has created a set of processes and tools to help organizations turn dysfunctions into productivity - https://recherche.iseor.com.

The AI transformation plan

Each actor in the organization reviews their work processes to document their hidden costs and those related to other stakeholders. Think tanks are set up to encourage the sharing of results. This exercise takes place over two or three weeks at most, following a TRAIDA master class to raise awareness about AI. The McKinsey study (2024), mentioned earlier, helps define quantified profitability objectives.

The results of the hidden cost analysis feed into the AI transformation plan. The first step of this plan is to implement an initial version of the minimal architecture to scale AI within the company. Therefore, AI use cases that sufficiently reduce hidden costs should be selected to make this semantic platform profitable. As mentioned earlier, TRAIDA recommends starting with the productivity card to create financial flexibility, ensuring a concrete return on investment from AI. Only after the initial productivity successes is the creativity card used.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The ease of access to AI tools allows everyone to use them, regardless of their level of training and professional experience. For example, using ChatGPT is simpler than using an Excel spreadsheet. However, the power of AI surpasses that of office tools, which is akin to putting potentially dangerous technology in everyone's hands. The goal is not to prohibit the free use of AI to increase knowledge, but it should certainly not be allowed on company processes and data without strict governance. If decision-makers fail to take this point of caution into account, the risk of AI usage failure is likely for the following reasons:

1. Augmented one-off tasks using AI do not guarantee an overall and long-term gain for the organization. Worse, in the absence of minimal governance, these free implementations backfire on their creators and cause dysfunctions within the company. In other words, AI is too powerful a technology to be deployed in successive patches without a governance and security architecture.

2. Misuse of AI can create a negative atmosphere and spread false ideas about its impacts, which are then difficult to correct.

Thus, the success of AI at the enterprise level relies on a minimal technical architecture to accommodate use cases in a profitable and secure manner. TRAIDA provides the technical and governance frameworks for its implementation in the form of a semantic platform.

This architectural effort is a *sine qua non* condition for sustainable AI profitability and a well-managed transformation of business models. However, its startup cost must be justified; otherwise, teams may develop AI applications without an architectural framework, leading to financial losses and poor quality, risks we have already highlighted.

The first AI use cases aimed at seeking productivity gains are selected in the following activity domains: internal organizational processes, those related to clients, and more broadly, all external stakeholders of the company, as well as compliance support with regulations.

The next part of this TRAIDA card presents some of the most common and easy-to-deploy use cases. They are limited to seeking productivity gains without any specific innovation effort.

INTERNAL PROCESS

Certain time-consuming administrative tasks, such as writing meeting minutes, drafting summary notes, or translation, represent significant sources of productivity. AI assistants are capable of automating 50% of these tasks.

The recruitment field also benefits from AI in the pursuit of productivity gains, with automatic analysis of applications, responding to candidates, and training new employees through assistants that act as virtual mentors. For example, a company can assign a specific AI assistant to each individual, playing the role of their digital twin. It then helps the employee with their daily tasks by accumulating knowledge on their behalf.

Still within the realm of human resource management, it is also beneficial to use AI to anticipate job dissatisfaction and correct it early enough to improve team retention.

AI is also an effective tool for optimizing the time spent on data entry, complementing the traditional operation of management applications: suggesting default values, verifying the relevance of information, and generally supporting users. Each IT application is thus enhanced by AI to optimize data entry, opening the door to productivity gains.

CLIENT PROCESS

The first use case is customer support, with the implementation of chatbot-type assistants to improve support availability and accumulate knowledge.

For better commercial management, predictive AI is used to detect risks of customer loss through weak signals that are sometimes difficult to interpret by humans.

Another area focuses on optimizing sales cycles to detect overly complex commercial negotiations with a low probability of leading to profitable sales. For complex offers, the customer's needs are analyzed by an AI assistant to determine whether it is profitable to respond to them or not.

Finally, AI is also a powerful tool for anticipating market needs to better forecast resource allocations, merchandise purchases, production means assignments, etc.

THIRD PARTY PROCESS

In the pursuit of productivity gains in external processes (after client process) of the company, the first area of interest is optimizing the supply chain through better coordination of stakeholders. For example, an AI assistant takes in the delivery conditions of several suppliers that need to be synchronized in an overall plan. It identifies possible optimizations in collaboration with the supply chain manager.

Still within the supply chain domain, reviewing supplier contracts based on general conditions, return policies, or negotiated pricing is time-consuming. The use of AI significantly optimizes these reviews and reduces errors.

Another area for seeking productivity gains is supplier review. Monitoring their financial stability, the quality of their services, and publicly available information is time-consuming and can be partially automated with AI. Finally, for complex organizations exposed to the risk of errors in invoice payments, an AI assistant detects overpayments.

COMPLIANCE PROCESS

It is possible to train an AI assistant using regulatory texts. It then acts as a legal advisor for teams that have questions about compliance requirements. It also operates in a practical manner by analyzing data flows generated by business activities to monitor regulatory compliance.

Another significant source of productivity arises when a new version of the regulation is introduced, requiring an analysis of its impact on existing processes. AI instantly conducts a gap analysis between the different versions of the regulation to propose an action plan to address the impacts.

Finally, a last example involves data protection regulations. Rather than relying solely on human inspections or developing complex verification software, AI is trained with the applicable rules and then analyzes the exchanged data flows to detect non-compliance. For example, this AI is useful for verifying that communications with clients comply with data protection and ethical standards.

3. BLUEPRINT

4. YOUR SITUATION & OBJECTIVES

CREATIVITY

Enhancing the creativity of certain company processes is an AI objective that complements the goal of improving productivity. The way decision-makers perceive the impact of AI on their own role also influences the relevance of the choices they will make for their organization's transformation. Indeed, AI is also competing with the intelligence of executives at all levels of the hierarchy.

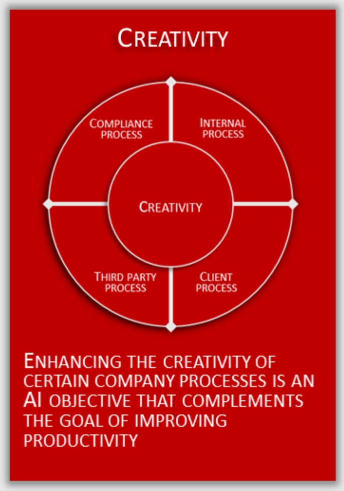

1. CONDITIONS OF SUCCESS

To ensure the large-scale integration of AI into the company, TRAIDA proposes a three-phase transformation plan:

1. The TRAIDA productivity business card is used to improve work processes through AI. The goal is to achieve concrete results based on the existing situation, while postponing a deeper transformation of the organization and business models (see the TRAIDA Productivity card).
2. Subsequently, an initial version of the minimum viable architecture to scale AI is implemented. This leads to the semantic platform recommended by TRAIDA with ODS, MDM, and EKG repositories (see respective TRAIDA technical cards). The profitability of this platform is achieved through the productivity gains generated during the previous phase.
3. Finally, thanks to the experience gained from implementing AI for productivity gains and the availability of the semantic platform, the TRAIDA creativity card is activated to transform the organization and business models with better risk control.

To maximize the profitability of AI-driven creativity and ensure stakeholder support, the company's ambition for its medium- and long-term transformation must be clearly defined. **Since AI raises concerns about the employability of individuals responsible for the company’s activities, total transparency regarding the transformation strategy is essential** and is based on the following observations:

- AI improves people’s daily lives, especially in health and education. In these areas, AI assistants will increase the availability of services with a quality superior to that offered by humans without AI. They will be accessible remotely by isolated individuals and poor countries. Thus, humanity should benefit from AI to better meet basic needs, including agriculture, transport, construction, etc. The more citizens become happy users of AI, the more its use will be facilitated in companies with the support of employees. In other words, the more a company trains its employees in using AI in their daily lives, the more it prepares for its positive integration into its own organization.
- AI will alter business models in all industries due to intelligence superior to that of humans. They will have to learn to collaborate with it.
- AI will have multiple forms: replacement AI to fully substitute humans; collaborative AI when it enhances human capabilities; and autonomous AI when it performs new tasks that humans have never undertaken.
- AI is multi-channel, meaning it can absorb written knowledge, as well as audio, visual, tactile, and perhaps even olfactory inputs. In this context, the fusion of AI and robotics opens up possibilities for versatile and human-free warehouses and factories.

- AI is also the driving force behind transhumanism, for the fusion of humans and machines, for example, through electronic chips implanted in the brain. In a less intrusive way, this also involves the use of 3D headsets and digital glasses for the metaverse.
- In the military field, AI poses a threat to humanity, with drones and other more or less autonomous destruction devices.
- Finally, AI is a technology with no limit for improvement. There is no known physical law that would set a ceiling to its evolution. Furthermore, it is likely that AI will become its own master architect for developing next-generation AIs. Consequently, no one can claim that this autonomous progress loop will ever reach an asymptote.

Given the strength of these observations, questions about the future of civilization arise in the following terms:

- What is the residual value of human labor when it is mostly performed by AIs with superior quality?
- What will be the useful jobs for training and controlling AIs?
- What will the workforce look like to maintain company operations when AIs are deployed on a large scale? For example, is it conceivable that technical inspections of AI-based vehicles could be carried out in smart workshops without mechanics? Could an insurance company replace its experts with AI assistants? Will research and development become more productive with AIs that think faster than researchers? Will humanity still need radiologists, dentists, factory workers, drivers, or teachers when AI assistants become increasingly intelligent?

These questions create anxiety among individuals and shake up every industry. TRAIDA does not claim to provide answers to these societal questions, but rather to help companies establish their survival plans in a context where collaboration between humans and AI is inevitable. Thus, when AI is merely seen as a technology that will eventually find its natural place in the organization, it is the company's competitiveness that is at stake, and perhaps even its survival.

The way decision-makers perceive the impact of AI on their own role also influences the relevance of the choices they will make for their organization's transformation. Indeed, AI also competes with the intelligence of executives at all levels of the hierarchy. In other words, if they do not make the effort to use AI in their daily work, the risk of them misunderstanding its use is considerable. This is why TRAIDA's master class emphasizes the importance of each member of the organization, including top-level decision-makers, creating their own AI assistant with their own knowledge base that they must build. A decision-maker who does not make this effort should not lead their company's transformation with AI.

Ultimately, everything we have just described raises a philosophical question about the meaning of life for humans with AI. Since the common ground between them is intelligence, it is naturally in this field that competition is open. To remain active, the individual must then demonstrate a level of creativity superior to AI. Thus, a human less creative than AI will be replaced by it, and a human more creative than AI will be enhanced by it. This race for intelligence is not lost for humanity, but it will be difficult to sustain with AIs improving autonomously and with exponential gains. To increase their chances of survival, individuals must then take these new principles into account:

- Manual work without creativity will be replaced by AI.
- Intellectual work with an insufficient level of creativity will be replaced by AI.
- Intermediate management tasks that involve project monitoring, resource management, or reporting will be replaced by AI. Consequently, organizational hierarchies will be reduced to give more freedom to autonomous teams. These teams will be managed by AI systems whose availability, speed, and efficiency will surpass those of human managers.
- The ability to formalize knowledge in writing to train and collaborate with AI assistants will become essential in all professions and at all levels of qualification. Even if a person's professional expertise is of a high level, a human who cannot interact in writing with AI will be replaced by it. In other

words, AI is always marginally more intelligent than humans in its ability to manage information. To remain active, humans will need to continuously provide new knowledge, which requires strong writing and analytical skills with intelligence.

At this stage of describing this card, it is clear that creativity with AI is a vast subject that raises the question of the company's survival, beyond even its transformation with new technology.

Before describing the topics of this card in the following sections, it is important to remind that TRAIDA advises starting your AI strategy with a focus on productivity gains without immediately changing the organization or business models. Your teams need time to understand AI's impacts, and it is better to do so without risking a big bang in your way of working. Only after achieving productivity gains with AI will your survival plan become clearer with the help of the TRAIDA creativity card.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

Value creation with AI begins with a complete overhaul of internal management within the company, then extends externally to customer relations, followed by other stakeholders, and finally toward the legislator to influence regulations.

INTERNAL PROCESS

The significant contribution of AI within the company is the reduction of its administrative functions in favor of management AIs. This is not only about targeting productivity gains but also about increasing the speed of collective work through the simplification of decision-making layers. The new organization with AI is then based on these principles:

- Elimination of intermediate managers, whose roles are replaced by AI.
- Faster and more efficient coordination between AIs compared to existing coordination between human managers.
- Increased work efficiency of teams, who act faster with AIs by removing intermediate human layers.

For this mode of working with management AIs to function properly, teams collaborate with them in the following ways:

- Ability to formalize knowledge about work processes in order to train management AIs and gradually replace managers.
- These managers are redeployed to operational teams, some moving into AI governance roles, while others become super-managers who will remain necessary to consolidate the work of management AIs. Thus, depending on the company's context, this redeployment takes different forms and considers the simplification brought by AI.
- Ability to collaborate with management AIs to achieve the expected benefits in terms of project monitoring, resource allocation, result analysis, proposing actions to address malfunctions or difficulties, etc. For this collaboration to be effective, teams must demonstrate critical thinking about the results from management AIs to help them improve continuously.

Thanks to this intelligent management approach, the creative capacity of teams increases. They invent better solutions to improve work by proposing new approaches that do not necessarily rely on AI. **In other words, creativity with AI does not necessarily mean using AI for creative use cases but rather using AI to free up creative thinking**.

CLIENT PROCESS

The more information a company has about how its offerings are used, the better it understands how to adapt them to increase profits. TRAIDA already addresses this aspect for customer support to improve its productivity through AI (see TRAIDA's Productivity card). In a more comprehensive way, the Creativity card opens up other avenues for increasing knowledge about the use of offerings, as follows:

1. Streamline the transmission of data from the customer by avoiding taking up too much of their time. Thus, rather than relying solely on manual satisfaction form entries, it is more efficient to implement multimedia communication using sound, images, and video. In other words, the customer communicates with the company through voice, photos, and video recordings, which are automatically processed by AI. For physical products, using a QR code creates a bridge to the company's website to capture customer feedback. Marketing intelligence then comes into play to motivate customers to provide more information about their use of the offerings.
2. Provide each customer with a digital twin of the offering in the form of an AI assistant. Beyond the obvious role of customer support, this assistant primarily accumulates knowledge about the customer to enhance its effectiveness and better understand how to improve offerings to better serve the market. This assistant is even more useful to the customer when it addresses a domain that goes beyond just the company's offerings. For example, it could be interested in other products in collaboration with partners to expand the value proposition of the assistant. The more useful the AI assistant is, the more incentive the customer has to stay loyal to it and provide more information. This accumulation of data is strategic for ensuring that the offerings evolve in the right direction and increase profitability.

THIRD PARTY PROCESS

Management AIs optimize purchase requests from teams more quickly and efficiently than human managers. In other words, the less internal work processes are slowed down by intermediate human management, the better the relationships with external stakeholders improve. Thus, the process of value creation with AI starts with a more streamlined management style, which then creates value through smarter and faster external processes.

AI also provides an opportunity to rethink data management in order to better capitalize on knowledge about stakeholders, particularly suppliers, to increase profitability. Finally, supplier sourcing also benefits from AI with better comparative analysis of market offerings and strengthened competition.

COMPLIANCE PROCESS

TRAIDA explains how AI helps companies better comply with regulations, particularly in governance with the Enterprise Governance card and implementation with the technical EKG (Enterprise Knowledge Graph) card. When considering AI as a tool for creating value in the regulatory space, the possibilities are limited. Indeed, companies are not responsible for creating new standards but for complying with those set by legislators.

However, large companies and those in disruptive sectors play an influential role (lobbying) with legislators. This is a legal activity based on trustful human relationships and more or less targeted communication actions, for example, through think tanks. AI plays a role in this influence strategy on two levels:

- It allows the company to gather information faster and on a larger scale, processing it intelligently to enhance its capacity for dialogue with legislators. This can include broad technological monitoring, i.e., analyzing existing and potential competitors and their own ability to influence regulations.
- It also enables the company to create AI assistants trained to convince legislators to evolve regulations for the company's benefit. The ultimate goal of such a system is for legislators to use these assistants themselves to better draft their laws. These assistants can be embodied through think tanks.

3. BLUEPRINT

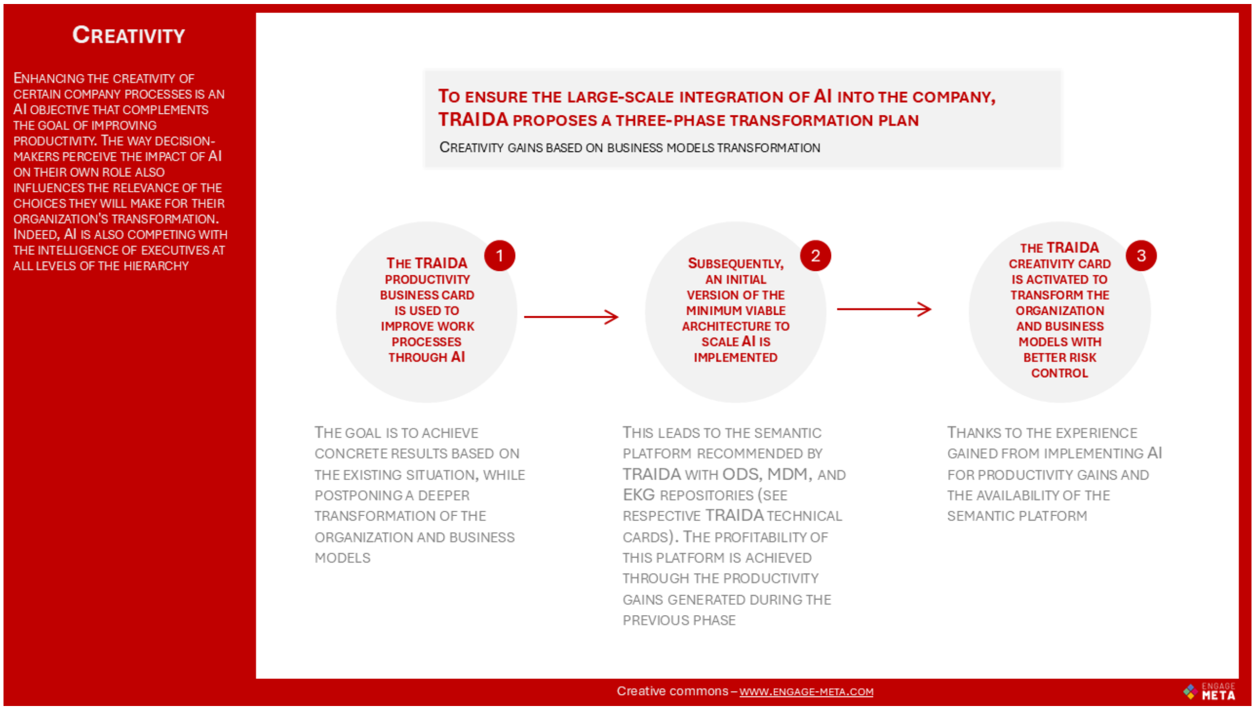

4. YOUR SITUATION & OBJECTIVES

TRUSTWORTHINESS

Trust in data and AI must be objectively assessed to successfully implement AI throughout the enterprise. The coupling of humans and AI enhances the intelligence of the organization, provided they complement each other to ensure reliable management. To achieve this, the user's trust in AI must be strong and can be improved by promoting AI that upholds the following qualities: reliability, honesty, competence, and integrity.

1. CONDITIONS OF SUCCESS

With generative, symbolic, or analytical AI, the dialogue between humans and computers is not limited to the deterministic scope of traditional software. Indeed, AI adapts to management situations by considering unforeseen events and incomplete information. Thus, the user no longer merely manages data to execute a predefined process but engages in a constructive dialogue with the AI to obtain responses tailored to their work situation.

For example, when a doctor classifies domestic accidents according to administrative criteria, they select values in the management application's interface: the time slot of the accident, location, object involved, height of the fall, water level, type of fire, etc. This data is used for statistical studies. The more precise the classification, the more time this administrative task consumes for the doctor. With AI, it is no longer necessary to predefine possible classifications in advance. The practitioner simply expresses the accident's context in natural language, and the AI handles its classification. By using voice input processed by the AI, the doctor further reduces the time spent on classifying each accident. The old application, at least its user interface, becomes obsolete. As such, the scope of digitization through AI is broader than that of traditional software.

TRAIDA advises first leveraging this strength to enhance productivity without changing existing applications and processes, and then focusing on creativity to deeply transform the organization and applications (see TRAIDA's business cards on productivity and creativity).

In other words, AI invites the user to contribute knowledge, clearly articulate their requests, analyze the responses, and ask for clarifications or additional information when needed. Thanks to this more intelligent dialogue between humans and machines, new task automations become possible. This setup is especially powerful for logics not fixed in algorithms, benefiting from the collaboration between the user and the machine.

The need for trust

This human-AI coupling increases the organization's intelligence, provided they complement each other to ensure reliable management. To achieve this, the user's trust in the AI must be strong, built on the following qualities:

- **Reliability**: Working professionally within a domain of expertise.
- **Honesty**: Telling the truth and not seeking to deceive or manipulate others.
- **Competence**: Possessing the necessary skills to be reliable in the domain of expertise.
- **Integrity**: Acting according to ethical standards and the company's values.

These qualities form the foundation of trust required to successfully deploy AI at an enterprise scale.

AI and consciousness

At the time of writing this card (October 2024), AI does not possess consciousness. It does not, therefore, conceptualize its relationship with humans. However, the data used to train AI essentially injects a cognitive heritage into it, embodying a form of artificial consciousness. Indeed, this heritage contains deliberate biases and others introduced by error. Thus, depending on its learning process, an AI is more or less confident in the role for which it was trained, such as pattern recognition, data analysis in a domain of expertise, solving specific problems, etc. **To avoid usage errors, it is important that AI refuses to respond to requests for which its training is insufficient**. If this boundary does not exist, AI responds outside its field of competence with hallucinations and approximations, which deteriorate the user's trust. Conversely, when this boundary is in place, a form of artificial consciousness emerges.

Thus, the decision of AI to respond or not to a user's request is a first level of singularity. It shifts AI into a realm beyond just new technology. With generative AI, the degree of tolerated hallucination is a parameter that adjusts the expected level of trust in the responses. The more the user expects a creative AI, the more hallucinations are encouraged; conversely, for a scientific AI that relies on facts, hallucinations will be minimized. Between these two behaviors, the degree of hallucination varies to allow AI to innovate from real facts. Thus, a subtle relationship takes place between the AI user (a), the AI itself (b), and indirectly the system that trained it (c). This relationship becomes systemic when several AIs are deployed within an organization for different users and with distinct training modes. This triplet (a, b, c) multiplies and generates exponential complexity in the interactions, with varying degrees of hallucinations desired or endured.

For example, an AI assistant (a1) manages logistical flows and interacts with another AI assistant (a2) specializing in comparative supplier analysis. A request is made by a1 for the urgent selection of a service provider to solve a delivery issue. However, a2 has not been trained to guarantee reliable contracting with a service provider, but only to establish a comparative list of potential suppliers. AI a2 then responds to AI a1 that it is not capable of producing the requested work. It justifies this decision by reminding that its training scope covers a different need. At this point, AI a1 has two possibilities: either it no longer solicits AI a2 because the risk of contracting in real-time with a service provider is too high; or it forces AI a2 to provide a limited selection so that it can assign an urgent order to a service provider. This use case shows that regulating exchanges between AIs and with humans is not trivial. These exchanges fall outside the usual scope of traditional software programming. Since it is impossible to anticipate all scenarios, it is necessary to build barriers to ensure that AIs do not exceed their area of competence and responsibility.

These barriers rely on establishing a sufficient level of trust and elevate AI to the rank of a stakeholder within the organization.

The level of trust is quickly established when it comes to integrating new technology into a company: either it brings a benefit, and trust is built, or it fails, and the company can move on without it. With AI, it is impossible to accept a situation where its usage fails and for the company to simply discard it without facing significant consequences. Indeed, it is human behavior toward AI that can lead to failure, not the technology itself. **In other words, the level of intelligence brought by AI is already too important for it to be disqualified (as of October 2024)**. The company then has no other option but to succeed in its transformation with AI, ensuring a sufficient level of trust in its use.

In TRAIDA's business card on productivity, we explained that AI is more than just a simple technology; it is a new stakeholder that possesses superior intelligence to humans in many use cases. Therefore, an efficient and harmonious relationship between humans and AI requires defining a stable and clear framework for work methods. It is not about claiming perfection in all operational rules from the first deployments of AI but ensuring that this framework is built with respect to the interests of all parties, in a transparent and committed manner. In other words, with an affirmed, durable, and proven level of trust.

This approach calls for resources that are not economic but human. This is a fundamental point for the success of AI deployment at the company-wide level. Indeed, although it is always possible to force an organization to adopt a new technology by imposing work processes, this does not work with AI. It is not enough to implement processes that include new technology; it is necessary to continuously reinvent the relationship between humans and AI. **This is a perpetual re-engineering that requires a critical mindset from all actors in how they collaborate with AI**. Without trust, this perpetual re-engineering is devastating for the organization, and AI risks entirely replacing humans.

However, trust cannot be decreed. It is a quality that is built progressively and is never definitively acquired. The company must then establish an organization that nurtures this trust. TRAIDA advises setting up two independent bodies that address the following objectives:

1. **AI Compliance**: Define the rules for AI transparency, ethics, and security.

2. **AI Quality Control**: Ensure that these rules are applied in accordance with expectations.

Each body operates with its own resources to guarantee its autonomy and independence. They are under the responsibility of the enterprise governance (see TRAIDA's card on this topic).

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

In the first part of this card, we emphasized the importance of trust in AI to ensure its successful large-scale deployment within the company. Building this trust is based on the consideration of three cardinal values: transparency, ethics, and security. To ensure these values are upheld over time, an AI quality control process is added.

TRANSPARENCY

Transparency is the responsibility of the AI compliance body (see the first part of this card). This involves documenting practices and making them known. The areas of application are varied, such as:

- The uses and impact on employability.
- The data used for training both internally and with stakeholders.
- Ethical and security rules.
- Expected and actual results.
- Investments.
- Risks and opportunities.
- Training and career plans with AI.
- Traceability and auditability of results.
- Detection of deviations and fraud in the use of AI.
- ../..

ETHICAL RULES

Best practices for AI ethics can be found in the public domain, particularly in government regulations. They address at least the following topics:

- Respect for the company's values.
- Respect for HR policies.
- Compliance with regulations.
- Democratization of AI usage.
- Reduction of the carbon footprint.

The definition of ethical rules is the responsibility of the AI compliance body (see the first part of this card).

SECURITY

A security breach involving AI inevitably leads to a lack of user trust in the technology. It is therefore important to address this by reviewing all processes that incorporate AI, such as:

- Data protection.
- Licensing for commercial and open-source AIs.
- Onboarding a new employee (AI usage rights, access to training data, etc.).
- Offboarding an employee.
- Work with contractors.
- Mergers and acquisitions.
- Rollback in case of malfunction.
- Backup and archiving.
- Etc.

In general, the stronger the security rules, the fewer innovative uses of the technology are possible. To avoid ossifying practices, it is useful to set up a free-use mode for AI in a technical environment and with data that does not pose security risks. This acts as a kind of sandbox for AI, where users can install new tools, test innovative AI behaviors with fictitious data, and more.

The definition of security rules is the responsibility of the AI compliance body (see the first part of this card).

CONTROL QUALITY

A dedicated AI quality control body is planned within the organization, alongside the body responsible for defining transparency, ethics, and security rules.

This control is applied at all levels of AI involvement: design, budgetary decisions, implementation, training, evaluation of results, etc.

3. BLUEPRINT

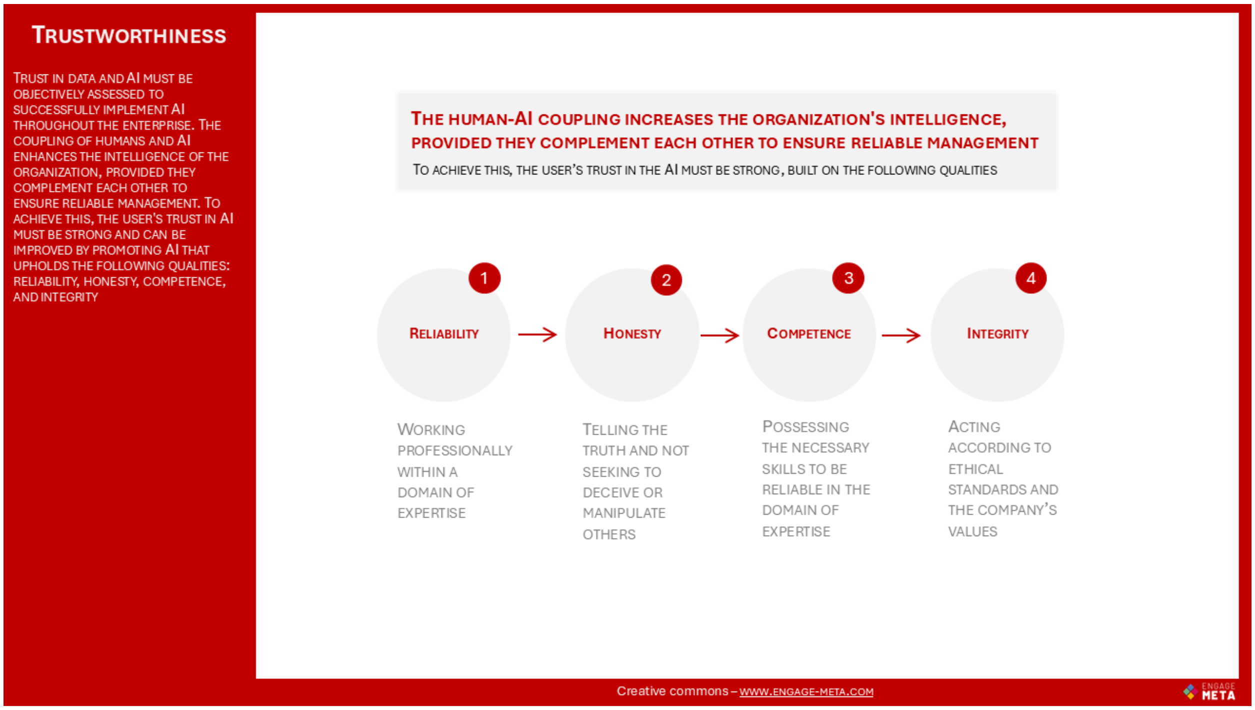

4. YOUR SITUATION & OBJECTIVES

TREASURY & ASSURANCE

Properly managing budgets and mastering value analysis are essential for successfully scaling AI. TRAIDA plans to deploy AI in three phases to manage financial commitments and economic risks: Boost (Phase 1), Expand (Phase 2), and Institutionalize (Phase 3).

1. CONDITIONS OF SUCCESS

The financial approach to large-scale AI integration is specific to each company's context. CAPEX (Capital Expenditure) and OPEX (Operating Expense) are not based on universal data. However, each company can follow an AI deployment plan to gradually gather the necessary information to control AI investments and optimize return on investment. To achieve this, TRAIDA proposes a three-phase deployment:

- **Boost** (Phase #1): Implementation of a minimal viable architecture (semantic platform) to deploy AI at scale, focusing on productivity gains (see TRAIDA's technical domain cards and business card on productivity).
- **Expand** (Phase #2): Enhancement of the minimal architecture to target initial creativity gains (see TRAIDA's card on this topic).
- **Institutionalize** (Phase #3): Full-scale exploitation of the architecture to leverage AI for transforming business models.

During each phase, the company increases its mastery of AI, cost structures, profitability criteria, and regulatory requirements. Thus, investment budgets, expected gains, and legal constraints are documented for each phase.

This gradual approach increases the likelihood of successfully integrating AI while avoiding the risks of deep usage too early in the process. Nevertheless, it advocates for the immediate deployment of a minimal viable architecture that facilitates the subsequent scaling of AI across the company. The following table outlines the concerns to address in each of the three phases.

CONCERNS	BOOST (PHASE #1) IMPLEMENTATION OF A MINIMAL VIABLE ARCHITECTURE TO SCALE AI, FOCUSING SOLELY ON PRODUCTIVITY GAINS	EXPAND (PHASE #2) ENHANCEMENT OF THE MINIMAL ARCHITECTURE TO TARGET INITIAL CREATIVITY GAINS	INSTITUTIONALIZE (PHASE #3) FULL-SCALE USE OF THE ARCHITECTURE TO LEVERAGE AI FOR TRANSFORMING BUSINESS MODELS
IMPLEMENTATION OF THE MINIMAL VIABLE ARCHITECTURE (SEMANTIC PLATFORM)	Version Boost Minimal viable architecture	Version Expand Improved evolution	Version Institutionalize Major evolution

PRIMARY TARGETED GAIN	Productivity	Creativity	Transformation
RISK LEVEL	Low	Meduim	High
ESTIMATED INVESTMENT COSTS FOR THE SEMANTIC PLATFORM (CAPEX) SIMULATION BASED ON THE COMPANY'S REVENUE (LINES BELOW)	Straightforward	Challenging	Highly complex
• SMALL : CA = $1M	$10.000	$20.000 - $50.000	Context dependent
• MEDIUM : CA = $10M	$60.000	$120.000 - $250.000	
• LARGE : CA > $500M	$200.000	$400.000 and above	
ESTIMATION OF OPERATING COSTS (OPEX)	Challenging	Challenging	Highly complex
ESTIMATION OF RETURN ON INVESTMENT	Straightforward	Straightforward	Straightforward
LEGAL IMPACTS	Straightforward	Challenging	Highly complex
AVERAGE DURATION	6 months maximum	12 months maximum	Depends on the context

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

Although there is plenty of financial data available to explain the costs and profitability of AI, each company's context is unique and requires adaptation. Additionally, the cost of AI software and its infrastructure frequently changes, necessitating constant adjustments in financial analysis. In this context, TRAIDA recommends managing AI integration in three phases (Boost, Expand, Institutionalize). These phases allow the gradual accumulation of financial knowledge within the specific context of each deployment.

INVESTMENT

The table in the first part of this document presents examples of CAPEX for the semantic platform according to the three phases: Boost, Expand, and Institutionalize.

Boost (phase #1)

During the Boost phase, the semantic platform is established to create a minimum viable architecture for scaling AI. The implementation of ODS, MDM, and EKG data repositories is a priority (see the respective TRAIDA cards in the technical domain).

Initial AI use cases are developed to target productivity gains.

Expand (phase #2)

During the Expand phase, the semantic platform is enhanced to enable the deployment of creative AI use cases. This goes beyond seeking the productivity gains targeted in the Boost phase.

Institutionalize (phase #3)

In the Institutionalize phase, the company decides to invest in a deep transformation of its business models. This phase is not mandatory and depends on the competitive context of each company. The deeper the transformation, the more the semantic platform strengthens its processing power to replace humans, including with AI-augmented robotics.

VALUATION

The table in the first part of this document indicates the levels of difficulty in estimating OPEX across the three phases: Boost, Expand, and Institutionalize. OPEX is added to CAPEX to calculate the profitability thresholds of AI solutions, starting with productivity gains (Boost phase), followed by creativity gains (Expand phase), and finally through business model transformation (Institutionalize phase).

Boost (phase #1)

During the Boost phase, CAPEX estimation for the semantic platform is feasible, whereas OPEX estimation is more delicate. However, since the goal is to deploy initial AI use cases focused solely on productivity gains, operational costs can be easily observed and stabilized without the need for significant initial financial resources. The economic approach is oriented toward usage-based billing for AI, enabling near real-time management of return on investment. Depending on the AI solutions used, the cost of user queries, tokens consumed in user-AI interactions, and AI training fees vary. It is impossible to know these costs precisely in advance, and they often change. To mitigate these uncertainties, the Boost phase provides an opportunity to better understand and manage AI OPEX, reducing risks before moving on to the more challenging Expand phase.

Expand (phase #2)

In the Expand phase, AI OPEX becomes more significant as the technology is used to create new use cases. As a result, it is quite difficult to predict AI usage frequency and return on investment. However, the experience gained during the Boost phase helps to better control the economic equation. Additionally, by maintaining a strategy of on-demand AI billing, the risk of financial overrun is eliminated. It then becomes possible to implement financial control measures to ensure that each dollar invested in AI use contributes sufficiently to creativity and productivity gains.

The key point here is not to begin the Expand phase without having sufficient control over the previous Boost phase. It is also important to account for the costs of training and supporting teams, which are at the intersection of CAPEX and OPEX.

Institutionalize (phase #3)

This phase is not mandatory and depends on the company’s strategy for using AI to deeply transform its business models. Similar to CAPEX, OPEX estimation depends on the specific context of each company. With TRAIDA, this phase is considered feasible only if the preceding Boost and Expand phases have been sustainably successful.

LEGAL

The table in the first part of this document indicates the levels of difficulty for legal efforts across the three phases: Boost, Expand, and Institutionalize. For the first two phases, the following issues should be taken into account:

- Legal protection of the data used by AI.
- Legal protection of AI-generated outputs.
- Updating employment and subcontracting contracts to reflect AI usage rights and obligations.
- Understanding and tracking the licenses of AI software used.
- Considering the impacts of AI on insurance contracts, particularly in cases where decisions are delegated to AI.

During the Institutionalize phase, the deep transformation of business models may lead to large-scale layoffs, which will require corresponding legal support.

3. BLUEPRINT

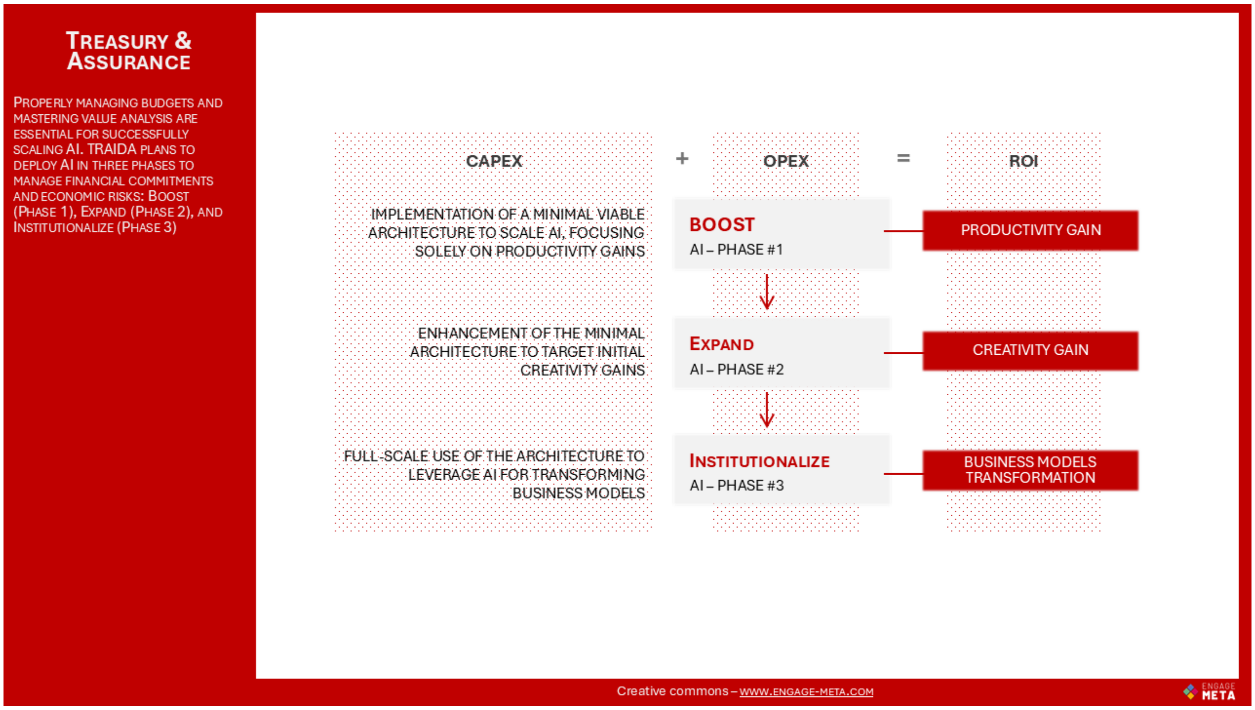

4. YOUR SITUATION & OBJECTIVES

GOVERNANCE DOMAIN OVERVIEW

General introduction to TRAIDA cards in the governance domain. The cards in this domain are universal and apply to all business contexts. You select the practices that correspond to your needs and complete them to manage a roadmap for implementing your minimum architecture to scale AI and data management solutions in your company.

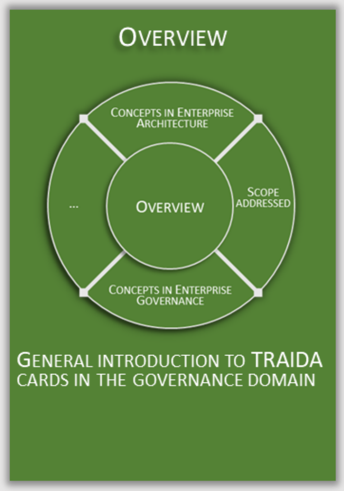

1. CONDITIONS OF SUCCESS

The TRAIDA framework (Transformative AI and Data Solutions) is based on three domains:

1. Technical (blue cards).
2. Governance (green cards).
3. Business (red cards).

To scale AI profitably across the enterprise, these three domains must be aligned.

The field of governance is based on a foundational principle: **AI is not just a new technology, but a stakeholder to be integrated into the company**. In other words, it is a kind of super collaborator that can intervene in all processes. It optimizes the way people work, helps humans be more productive, and makes decisions with a level of autonomy that depends on its configuration. This is a revolution that is transforming the world.

The benefits of AI are already visible, but this is only the beginning. Innovation in this field is dynamic. As of the writing of this TRAIDA card (September 2024), competition among players in the field is primarily focused on the IT infrastructure necessary for AI training. However, the next step is already in sight, with the idea that the benefits of these massive trainings on billions of parameters are approaching an asymptote in the creation of intelligence.

Moreover, after absorbing the entire Internet, sources of information are not infinite, which poses a structural limit to the large-scale training of AI models. **It is, therefore, time to open a new chapter to improve generative AI with an additional intelligence called deductive**, meaning it is capable of conducting complex reasoning based on a chain of thought.

Generative AI would then be able to question itself about the user's request, and then about the results it proposes to improve the relevance of its final answer. During this reflection, it can detect issues in the initial request, inconsistencies in the data, and gaps in information that it will seek to fill either on its own or with the support of the user. This system reduces hallucinations and refines the quality of the final answer.

With innovations like this, and others sure to follow, it is likely that artificial general intelligence (AGI) will emerge by 2030. It is not a certainty, but it signals at least that much more powerful AIs will be available in the coming years. AI will be able to address any problem with a level of intelligence superior to the best human experts in the relevant field.

To be convinced, one must ask how innovation emerges in the thought process. For example, in my personal case(*), my engineering background and experience lead me to seek innovation by practicing these principles:

a) Creating mental representations of my knowledge in the form of graphs, and regularly accumulating new knowledge in my field of expertise and in other areas to build knowledge.
b) Formulating a response to a problem by recycling my knowledge and adding ideas outside the context of study to create something new; this is the innovative effect.
c) Critically analyzing the result of my analysis to identify points of inconsistency, improvement, and clarification, and looping back to the previous step (b) as much as necessary and possible.
d) Sharing my work with others to benefit from intellectual impulses that contribute to value creation.

(*) Pierre Bonnet, main author of TRAIDA.

This magic sauce for innovation is not unique to me. Most people work according to these principles without even questioning their way of thinking and acting. The question, then, is how AI uses the same sauce; let's revisit the four principles from the perspective of AI practice:

a) AI absorbs large amounts of information to generate usable internal representations. It applies this first principle more efficiently than humans.
b) AI recycles its vast knowledge to combine it into a response. Like humans, a degree of hallucination is introduced, varying in intensity depending on its configuration.
c) Classic generative AI is less efficient than humans at critiquing its own responses to enrich them. Without additional systems, the hallucinations produced earlier (b) are not identified and corrected. This is where the addition of deductive intelligence significantly enhances the power of generative AI. It can now critique its own responses to detect reasoning errors and loop back to the previous principle (b).
d) AI can share its results with humans, who can then contribute to improving the responses. It can also autonomously interact with other AIs, especially in the context of the reasoning process inherent to deductive AI.

If you are convinced that the integration of this general AI is essential on a large scale in your organization, it is important to prepare and implement its proper governance. With TRAIDA, this is addressed through enterprise governance and enterprise architecture. Additionally, human resource management plays a transversal role.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

This card is an introduction to the governance domain of the TRAIDA framework. It helps you become familiar with the other cards in this domain. The following provides some additional information to facilitate your reading and the necessary reflection for your own context.

CONCEPTS IN ENTERPRISE GOVERNANCE

Generally speaking, enterprise governance focuses on risk management and compliance with both internal and external regulations affecting the company. It is a broad area of application that varies from one company to another.

However, the trend is an increase in requirements in this field, as the social, economic, political, financial, and technical worlds become more regulated. For instance, in managing its information for the public, a company must control all its communication channels to avoid disseminating false data, information outside of its ethics, or data prohibited by regulations. Given the speed of exchanges on social media, it is challenging to verify every message unless dedicated teams are mobilized, which may still prove insufficient.

To ease this heavy regulatory burden, AI assists in automating controls by implementing monitoring systems and reducing the need for human intervention. In other words, the principle of "code is law" is supported by AI. Enterprise governance also covers data governance, which is essential for the reliable and profitable large-scale use of AI. Finally, it also addresses AI governance, which must align with the governance applied to data. New-generation software, sometimes referred to as data fabric or more broadly as semantic platforms, offer solutions in this area.

CONCEPTS IN ENTERPRISE ARCHITECTURE

Enterprise architecture (EA) is often seen as a theoretical discipline, removed from the concrete concerns of IT projects. Yet, its goal of documenting the information system to better transform it is legitimate. Without this knowledge, it is difficult to deploy IT solutions in a coordinated manner across the enterprise.

The obstacles to the deployment of enterprise architecture lie in the difficulty of keeping documentation up to date and in the cumbersome nature of using best practices to transform the information system. Caught between a lack of alignment with the formalization of an ever-changing reality and technological variations that make cross-cutting decisions difficult, enterprise architecture struggles to justify its profitability.

In this context, AI is both a revealer of the importance of enterprise architecture and an accelerator for its more profitable implementation. Indeed, AI integrates into numerous processes and requires high-quality data. Without precise documentation of the information system, it is challenging to properly manage its transformation with AI. However, traditional practices must be adapted to give greater importance to semantic modeling (ontology), data governance, and knowledge management.

Finally, AI helps improve the automatic management of EA documentation to ensure better tracking of necessary updates as systems evolve. The combination of knowledge graph databases with generative AI enables the creation of highly efficient document repositories (see the TRAIDA Enterprise Knowledge Graph - EKG card).

SCOPE ADDRESSED

The cards in the governance domain are listed in the table below. There is no preferred reading order to follow. From an academic perspective, that is, for discovering the cards with the aim of learning general technical culture, the order of the cards in the table is the most advisable to follow.

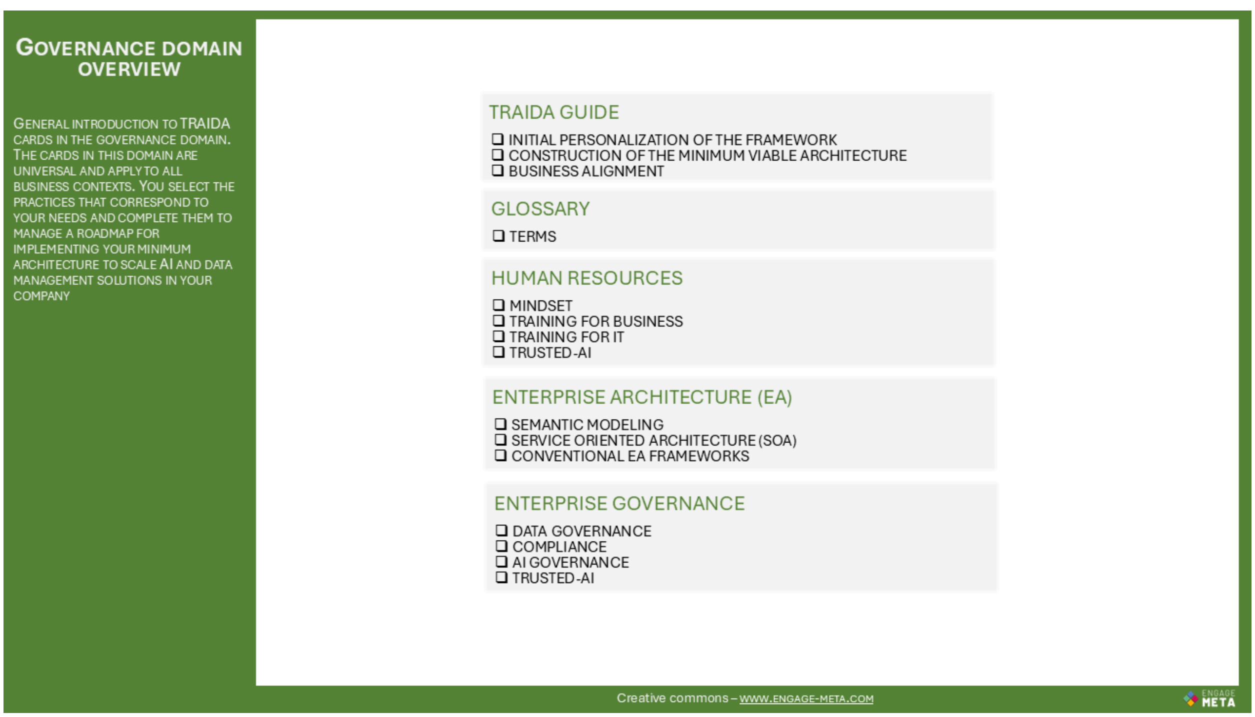

3. YOUR SITUATION & OBJECTIVES

TRAIDA GUIDE

TRAIDA is a knowledge repository with an educational purpose on AI and data solutions. Its primary use is therefore the culture development of your teams on the architectural consequences of AI and data solutions on your information system. Once your teams are sufficiently aware of the architectural impacts of AI and associated data, TRAIDA is used as an operational tool to assist in the gradual transformation of your information system towards large-scale AI implementation. It relies on three stages: Initial personalization of the framework (1); construction of the minimum viable architecture (2); business alignment (3).

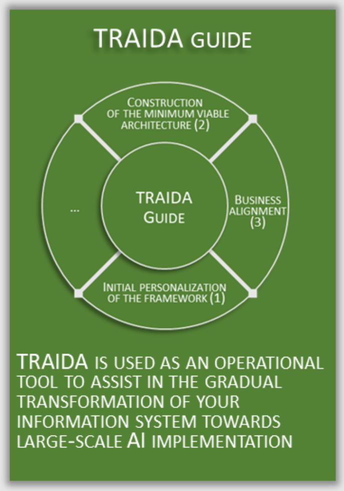

1. CONDITIONS OF SUCCESS

Thanks to its ready-to-use knowledge base, the TRAIDA framework helps you spread a uniform culture of AI and data solutions among your teams. It's an essential step before utilizing the framework for the transformation of your information system with AI.

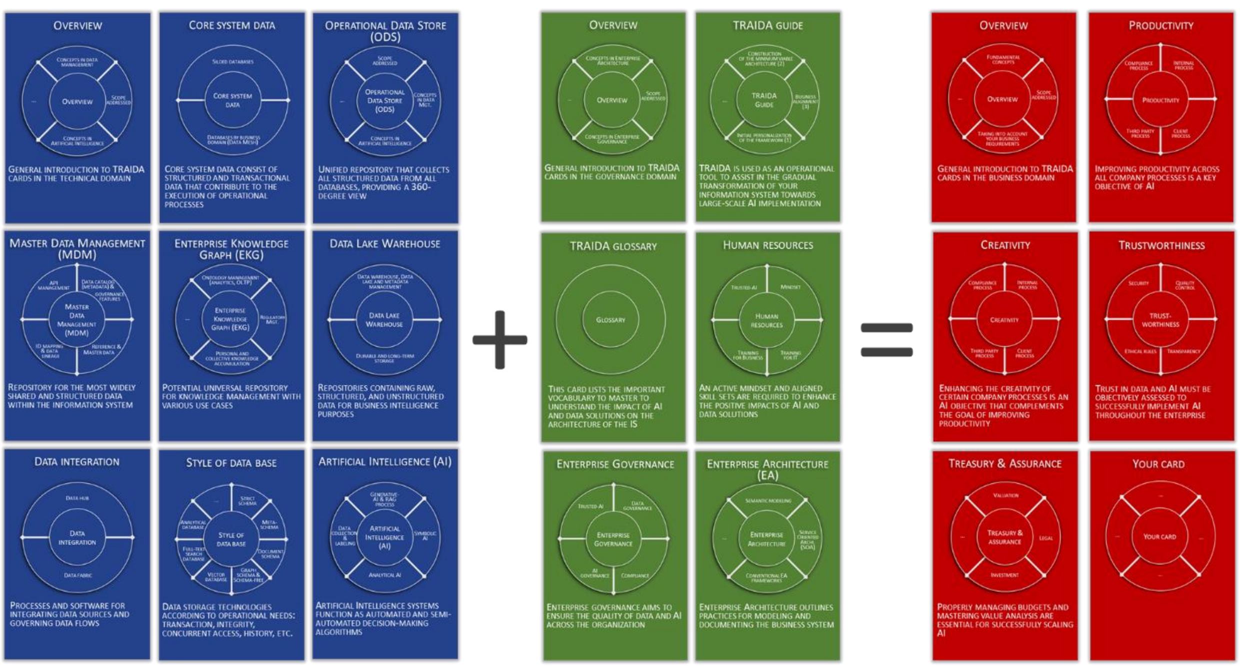

TRAIDA consists of technical cards (blue), governance cards (green), and business cards (red). Each card is described in writing and revolves around a set of a few key topics that the company must consider.

This sharing of knowledge fosters the commitment of stakeholders to the success of projects and the quality of their results over the long term. Even if you already have significant AI expertise and a good understanding of the impacts on data management, it remains costly to formalize a wide-reaching knowledge framework like that proposed by TRAIDA. To save time and optimize your costs, the framework is a catalyst for drafting the essential knowledge to support your educational approach.

The knowledge formalized in TRAIDA is useful for training your teams, your service providers, but also for implementing quality control processes such as the selection of AI and data management software, or for increasing the relevance of the governance of your information system.

During the educational phase of spreading general AI culture, it is preferable not to alter the content of the framework. Only limited adaptations to a few fundamental terms of the TRAIDA vocabulary should suffice. Indeed, we advise not to modify the other cards, including those of the business. Your goal should be to rapidly spread a general culture of AI and data management without it being fundamental to detail use cases specific to your company. Nevertheless, if these exist and can be formalized quickly, they will always be useful. Conversely, if their drafting imposes a foundational work while the general AI culture is not yet forged, it may create unnecessary confusion in your teams. Worse, if these projects are not properly understood by the stakeholders, they can be counterexamples and hinder the rapid sharing of a common culture.

The objective is to proceed to a first phase of education in less than two months.

This involves delivering a TRAIDA master class and introductory workshops on AI concepts in a spirit of exchange and listening to participants. It's an opportunity to address potential obstacles and answer some questions. To act quickly, you must avoid the trap of personalizing the framework that would seek to prematurely take into account a complex existing IT environment. The goal is to rally as many of your stakeholders as possible to the TRAIDA framework in its initial version. You will explain that it will undergo a specific adaptation to the company in upcoming work.

After a successful educational approach, the TRAIDA framework is used as a tool to aid the progressive transformation of your information system with AI. It proposes three fundamental steps which are detailed in the continuation of this card:

1. Initial personalization of the framework.
2. Construction of the minimum viable architecture.
3. Business alignment.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The success of deploying AI across your company primarily depends on two fundamental elements. On one hand, the rallying of your teams to a common culture surrounding AI and the management of associated data. We discussed this at the beginning of this card. TRAIDA is your ideal educational tool for spreading this culture.

On the other hand, the specification of a business system architecture and more basically of an information system, which allows you to deploy your first AI projects while ensuring a gradual scaling. This is about creating a minimum viable architecture for scaling. Since you cannot put everything in place at once, this minimum architecture will help you manage the different stages of your transformation with AI. The TRAIDA cycle described in this card helps you converge towards this minimum viable architecture.

INITIAL PERSONALIZATION OF THE FRAMEWORK (1)

The technical and governance cards of the TRAIDA framework are universal and do not need to be customized to your context. However, the vocabulary listed in the "TRAIDA glossary" card can be adapted to your organization. These changes will then necessitate adjustments in the texts of the cards. The stability of this vocabulary and its adoption by your teams is a key element of success for scaling AI. It reduces misunderstandings and misconceptions that prove detrimental in any transformation project.

In the business domain, the default TRAIDA cards are universal. They offer a general perspective on the impact of AI in the company, covering productivity, creativity, trust, and finally finance and legal aspects. They enable your teams to start their reflection by avoiding the fear of a blank slate or, conversely, a premature confrontation with a too-long list of micro-needs that do not help in building a solid and lasting vision for the information system.

With a well-defined corpus of terms and an initial set of sufficiently broad business requirements to support a global reflection, your teams are well-positioned to start in-depth work on the architecture. It is important not to block your teams at the start or let them get lost in details that would be premature to analyze.

This initial customization is not final since the framework undergoes regular changes during the iterations in the subsequent stages. At this stage, however, it is important to establish the initial pillars of the business on which the AI and data management strategy must rest.

Construction of the minimum viable architecture (2)

This stage involves a comparative analysis between the technical requirements formulated in TRAIDA and the solutions provided by the architecture of your information system. In the first iteration, you have access to the default TRAIDA business cards and those specific to your context described during the initial customization. As iterations progress, the business cards will express new needs that must then be taken into account in the evolution of the architecture.

The comparative analysis is conducted according to two complementary scenarios. The first is independent of business requirements. It involves reviewing all technical and governance topics without considering business priorities. The second scenario is business-dependent and focuses the analytical effort on only the technical and governance cards needed to meet the requirements.

Initially, we advise conducting an analysis independent of business requirements to review the entire architecture of your existing information system, followed by a second, medium-term analysis. In subsequent iterations, you will work from business needs. This will allow you to better understand the gap between the architecture of the information system resulting from business needs and the theoretical architecture that should be deployed. From these analytical elements, you can construct your minimal architecture that does not permanently stray from evolving towards the theoretical target.

Since there is no universal architecture for AI and data solutions, the work carried out in this stage is not conducted with the mindset of a maturity study. The goal is to clarify a minimal information system architecture that is acceptable in your context and facilitates widespread use of AI and data management solutions. It should enable the gradual deployment of AI and the accompanying data solutions. The company cannot deploy all the technologies, methods, and practices for AI at once and across the entire scope of the information system. Therefore, you must build a framework of thought that embodies a powerful and global conceptual vision to better determine the path to follow to meet your needs in a pragmatic and sustainable manner.

To construct the minimal architecture necessary for scaling AI, you will need to deeply assimilate each of the TRAIDA cards to objectively compare them with your existing setup and then with your business objectives. These will evolve over time and are formalized in the next stage of alignment.

Business alignment (3)

This stage is devoted to the analysis and adaptation of business cards that serve to question the architecture developed in the previous step (2).

The formalization of business requirements revolves around two categories. First, there are generic or cross-cutting needs that are not directly related to a business project. The default cards provided in TRAIDA fall into this category. Next, there are needs that arise during a transformation project, such as the implementation of new software or a database.

The new business cards are subject to the formalization of requirements with a level of drafting identical to that of the other cards in the TRAIDA framework. They are used at two levels. First, to analyze the alignment between business needs and the capabilities of the architecture. The reference point considered may be the existing architecture of the information system or a medium- to long-term target. Then, as requirements to be taken into account to feed a new iteration with the previous step (2) in order to question the architecture again and evolve it.

3. BLUEPRINT

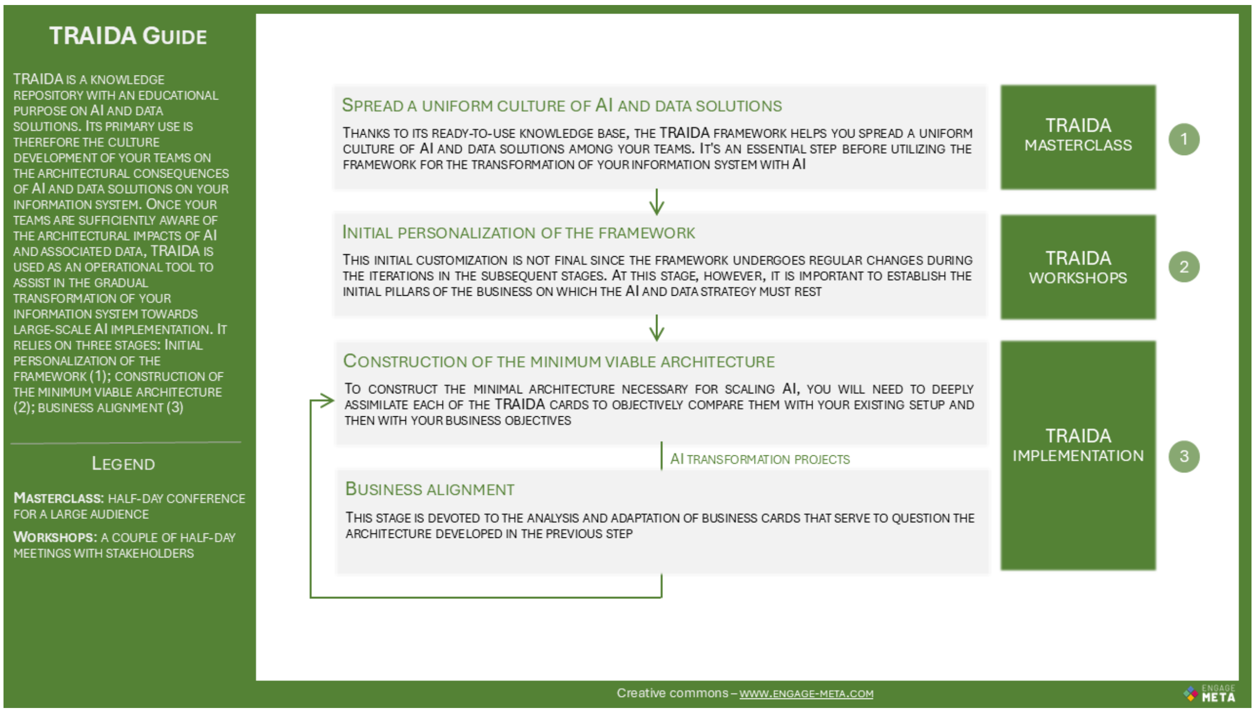

4. YOUR SITUATION & OBJECTIVES

TRAIDA GLOSSARY

To increase your speed of spreading a culture of AI and data management that is understandable by all of your technical and business teams, it is essential to establish and share a glossary of AI and data solutions terms. Although popular, some of these terms do not always have a definition commonly recognized by the market. You will therefore need to decide on your vocabulary choices. This card gives you the starting point for this semantic work, which is fundamental to building and managing your transformation with AI and data management.

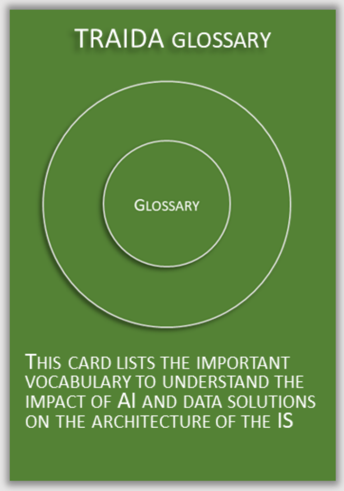

The definitions are customized for the TRAIDA framework. They are not intended to conform to the marketing presentations of software vendors or IT analysis firms. Based on these definitions, you can create your own company glossary and update the various cards of the TRAIDA framework according to your context. **However, the more you maintain definitions that are neutral in relation to marketing trends, the more comprehensible your AI and data solutions strategy will be to your stakeholders, and the more robust it will remain over time**. The worst scenario would be to introduce terms and definitions that change too frequently and are challenged by the marketing and sales rhetoric of solution providers, whether they are technology companies or consultants. By relying on the most neutral definitions possible, TRAIDA helps you build a stable communication strategy for AI and data solutions in your context.

D

Data fabric, data hub and data mesh (overview)

Data fabric and **data hub** are complex to define precisely, as different software vendors encompass various concepts within these terms. At TRAIDA, we prioritize identifying the needs of the three fundamental repositories regardless of the chosen data fabric and data hub solutions: Master Data Management (MDM), Operational Data Store (ODS), and Enterprise Knowledge Graph (EKG). No single technology can universally manage these three repositories on the same platform. **To choose the best data fabric and data hub tools for your context, it is important first to have a clear understanding of your needs in MDM, ODS, and EKG** (refer to the respective TRAIDA cards). It is based on these needs that scaling AI and data solutions will be properly managed. Otherwise, you risk selecting technological solutions that are suitable for an initial deployment but not appropriate for scaling AI and data management solutions.

The term **data mesh** is relatively straightforward to define, as it refers to a data architecture that organizes data by business concepts to reduce silos (micro databases).

Data fabric

A data fabric is a comprehensive set of technologies designed to streamline data integration processes, including referencing data sources (data sets), data cleaning procedures, and unifying data structures through semantic modeling. It relies on robust metadata management systems and often uses graph knowledge database technology.

Modern data fabric supports the configuration and testing of AI decision-making algorithms (such as machine learning, AI training, and rule-based systems), as well as the deployment and monitoring of AI processes and data in production environments.

While a data fabric can assume some roles of a data hub (data integration), its primary focus is to enhance data and AI governance at scale. Rather than replacing MDM (Master Data Management), ODS (Operational Data Store), and EKG (Enterprise Knowledge Graph) repositories, it should coordinate them. **However, the overlap between a data fabric and core repositories like MDM, ODS, and EKG must be carefully evaluated before deciding on large-scale deployment**.

In a data mesh context, a data fabric can also offer additional features for controlling micro databases, such as data caching, inter-database transactions, workflow management, and support for long transactions.

Data hub

A data hub primarily functions as a data flow integration bus, incorporating technologies like EAI (Enterprise Application Integration), ETL (Extract - Transform - Load), and ESB (Enterprise Service Bus).

Depending on the solution, a data hub can manage metadata (mainly at the flow level), map IDs across silos, visualize unified data, and store certain operational data akin to an ODS (Operational Data Store).

Coupled with a data mesh approach, it can also handle data caching and long transaction management.

While some vendors market data hubs as universal data management platforms, they often fall short of fully implementing MDM, ODS, and EKG systems. It's typically more effective to use data hubs for integrating data flows and supplement them with dedicated solutions for MDM, ODS, and EKG.

More generally, the concept of a data hub is gradually being absorbed by the broader concept of a data fabric. We can therefore say that a data fabric solution either natively includes or integrates with a data hub solution. Open-source offerings facilitate this kind of integration, which should be carefully considered when selecting tools.

Data mesh

Data Mesh is a data architecture approach that organizes data by business domains or concepts, rather than by functional or organizational silos. It uses semantic modeling and a technical infrastructure to manage transactions between business concepts spread across different micro databases.

Data Mesh enhances data governance and reduces data duplication. It is a set of architectural principles rather than a specific technology. Implementing a Data Mesh requires leveraging data fabric and data hub technologies, tailored to the specific context of each company.

E

Enterprise Knowledge Graph (EKG)

The Enterprise Knowledge Graph (EKG) is a repository specialized in knowledge accumulation. It manages both structured and unstructured data, with the capability to receive data sources without requiring prior modeling. It is based on the technology of knowledge graph-oriented databases.

Unlike MDM, it does not have as advanced governance processes; and unlike the ODS, it does not offer as powerful transactional management (OLTP).

Depending on the context of each company, it is necessary to find the best combination between the needs of the MDM, ODS, and EKG. However, it is important that all these repositories share the same ontologies to avoid the negative effects of siloing. For a small or medium-sized enterprise, it is feasible to manage everything within a knowledge graph-oriented database, that is, within the EKG.

On the other hand, for a larger information system supported by a rationalization policy, it may be necessary to opt for three different technologies for the MDM, ODS, and EKG. The worst approach would be to implement as many EKG repositories as there are functional domains without considering the cross-functional needs of the MDM and ODS. This would lead to siloed EKGs with associated quality issues. In such a case, large-scale AI integration within the enterprise would be compromised.

Each of these repositories—MDM, ODS, and EKG—is covered by a dedicated TRAIDA card.

M

Master Data Management (MDM)

Master Data Management (MDM) is a data repository specialized in managing reference and master data. These are the most shared data between applications. Their lifecycle is less rapid than that of transactional data.

The strength of MDM lies in its agility to accommodate changes in the structures of reference and master data, and in the richness of its data governance processes: quality, security, traceability, data entry UI, reporting, version and variant management, workflow, etc. MDM is also the preferred repository for creating a metadata catalog, which benefits from the full power of its governance. In this context, the MDM includes descriptions of the ontologies managed within the company, which form the core of the semantic platform recommended by TRAIDA.

MDM works in collaboration with the Operational Data Store (ODS) and the Enterprise Knowledge Graph (EKG). Each of these repositories—MDM, ODS, and EKG—is covered by a dedicated TRAIDA card.

O

Operational Data Store (ODS)

The Operational Data Store (ODS) is a data repository specialized in the unified management of operational data. It provides a unified access point to data from multiple sources, meaning data located in heterogeneous databases (silos). Unlike MDM, the ODS deals with transactional data, which has a rapid lifecycle. It is therefore specialized in transaction management and does not offer governance processes like those for reference and master data.

A vertical implementation of the ODS for a specific functional domain leads to the concept of a data hub, such as with Customer Data Integration (CDI) or Product Information Management (PIM). This siloing results in unnecessary data duplication, increasing the risk of poor data quality.

In TRAIDA, the embodiment of ontologies in the semantic platform does not rely on such verticalization. Instead, it is advisable to build a solution that establishes a single ODS. This ODS then works in close collaboration with the MDM, which provides a central access point to reference data, master data, and metadata. Each of these repositories—MDM, ODS, and EKG—is covered by a dedicated TRAIDA card.

Ontology

An ontology is a structured representation of a domain of knowledge. It is based on these four fundamental properties:

1. Exhaustive: The entire semantics of the domain is expressed in the form of concepts and relationships.
2. Unified: There is no redundancy.
3. Explicit: There is no ambiguity.
4. Universal: It is independent of information processing technologies.

Maintaining these four properties throughout the lifecycle of a domain is challenging. Thus, an ontology is a living representation that evolves to improve and accommodate changes. It is necessary, therefore, to plan for version management, variants, and the impact of ontology deployment, which means adopting appropriate governance.

To achieve such a powerful representation, an ontology requires these components:

- Glossary: Unambiguous definitions of concepts.
- Thesaurus: An extension of the glossary with synonymous terms and expression equivalences according to the contexts in which the concepts are used, including multilingual contexts.
- Taxonomy: Hierarchies among the concepts.
- State Machine: The lifecycle of each concept and synchronization between concepts.
- Identifiers: Format and semantics of the identifiers for each concept.

The combination of these components, along with the procedure used for their construction, forms semantic modeling. Semantic modeling is thus the discipline that allows the construction of ontologies.

From the perspective of the tooled representation of ontology, standards such as RDF and OWL are used. It is also possible to opt for a representation with UML, which is also used for semantic modeling. To illustrate these choices of technical representations, here are two possible use cases:

- Using RDF for the ontology, then transcribing it into UML to obtain a semantic model. This model is then derived into a logical model in a database.
- Using UML for the ontology and deriving it to the logical level for implementation.

With TRAIDA, ontologies and semantic modeling form the foundation of the semantic platform that enables the construction of a digital twin of the information system. It is from this digital twin that the integration of AI systems is ensured with associated data management solutions.

S

Semantic modeling

Semantic modeling brings together the design processes for the following components: glossary, thesaurus, taxonomy, ontology, state machine, and identifiers. All of these are necessary to formalize the knowledge of a domain, such as an organization, a business, an activity, or an area of expertise. This formalization is carried out independently of any specific technological implementations.

In TRAIDA, semantic modeling is used to build the semantic platform that powers the ontologies, from which AI systems and associated data management solutions are integrated (digital twin). This approach avoids integrating AI at a lower level of abstraction, that of the physical flows of data and application systems. In most cases, these physical layers do not have the required quality to ensure reliable AI execution and the agility needed to adapt quickly enough to business requirements.

HUMAN RESOURCES

An active mindset and aligned skill sets are required to enhance the positive impacts of AI and data solutions. Reducing AI to just another technology does not reflect reality. Indeed, it brings a level of intelligence that gives it a special role. Therefore, a traditional approach to change management is insufficient.

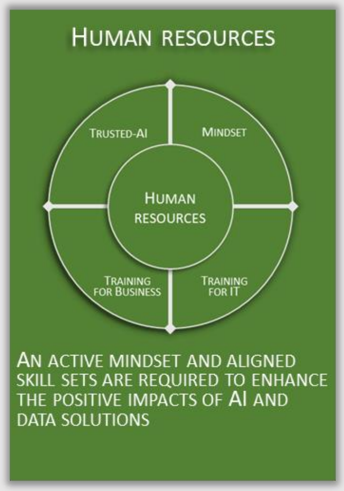

1. CONDITIONS OF SUCCESS

The integration of new technologies is generally accompanied by change management involving training and process reengineering. When AI is perceived as just an additional technology, these practices are reused.

However, reducing AI to just another technology does not reflect reality. Indeed, it brings a level of intelligence that gives it a special role. Therefore, a traditional approach to change management is insufficient.

In fact, AI is a new stakeholder that needs to be integrated into the organization. In other words, it involves welcoming a new actor who will impact all work processes. It is therefore natural that human resource management takes an interest in it. To be convinced of this, the following fundamental characteristics of AI should be considered:

- It is the only technology that explains to the user how it can help in their activity or, more generally, in their life. In other words, generative AI relies on a dialogue with its user that is not pre-written. This conversational aspect, personalized to each usage context, is revolutionary. It fosters a mutual enrichment between humans and AI. This embodiment justifies its role as a stakeholder in the organization.
- With improvements in generative AI, this conversation becomes increasingly intelligent. For example, at the time of writing this TRAIDA document, the ChatGPT o1 version offers a new deductive working mode that improves use cases for research and planning (see the following paragraph). Conversations between the user and this AI resemble a dialogue between humans.
- Its access is immediate and does not require prior investment in a technical infrastructure. On-demand service platforms democratize the use of AI. Its power is within everyone's reach, at least for common usage. Only massive AI training requires significant computing power and is handled by major tech operators.
- For the first time in human history, a competition of intelligence between humans and machines emerges: a human who works with AI is more productive than a human working alone. The most intelligent AIs will outperform even humans augmented with AI. From a systemic perspective, the collective intelligence of an organization interacts with another intelligence that emerges through interactions with AI assistants. A clarification of the operating rules between these two intelligences is necessary, leading to the concept of trusted AI (see the rest of this document).

If your company considers AI to be just another technology, you may not be convinced by the aforementioned characteristics. Conversely, if you adopt a more impactful scenario of AI, you will pay close attention to the considerations outlined in this TRAIDA document. It does not redefine a traditional change management strategy but explains the specific challenges of AI for human resource management: mindset, training for business, technical team training, and trusted AI.

Example of the enhancement of conversational intelligence with ChatGPT o1

Here is the statement of a basic exercise that is first submitted to ChatGPT 4o, then to version o1, which includes a deductive working mode (chain of thought): *"I have a mathematical problem. A boy has 3 apples, and a girl has 1 apple. Suddenly, an event occurs that makes 20 more apples available. After some discussion between the boy and the girl, they decide to share the apples. How many apples does each of them have after the sharing?*". This test is conducted on September 14, 2024. The response with ChatGPT 4o is limited to a single scenario, without additional analysis. The dialogue with the user is thus restricted.

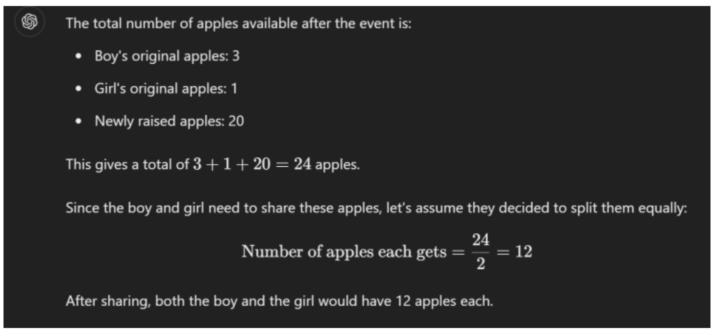

The total number of apples available after the event is:

- Boy's original apples: 3
- Girl's original apples: 1
- Newly raised apples: 20

This gives a total of $3 + 1 + 20 = 24$ apples.

Since the boy and girl need to share these apples, let's assume they decided to split them equally:

$$\text{Number of apples each gets} = \frac{24}{2} = 12$$

After sharing, both the boy and the girl would have 12 apples each.

The response with ChatGPT o1 is more comprehensive, with a description of multiple scenarios that allows for a much richer conversation between the human and the AI.

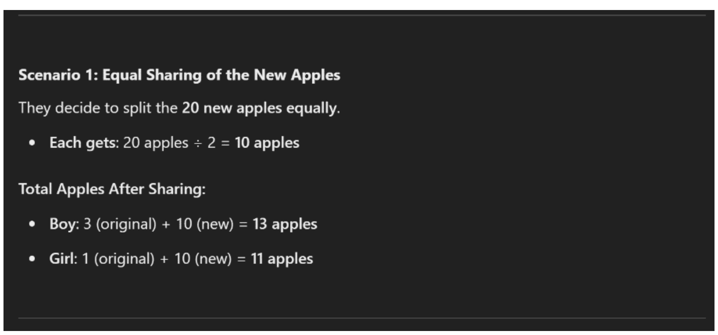

Scenario 1: Equal Sharing of the New Apples

They decide to split the **20 new apples equally**.

- **Each gets**: 20 apples ÷ 2 = **10 apples**

Total Apples After Sharing:

- **Boy**: 3 (original) + 10 (new) = **13 apples**
- **Girl**: 1 (original) + 10 (new) = **11 apples**

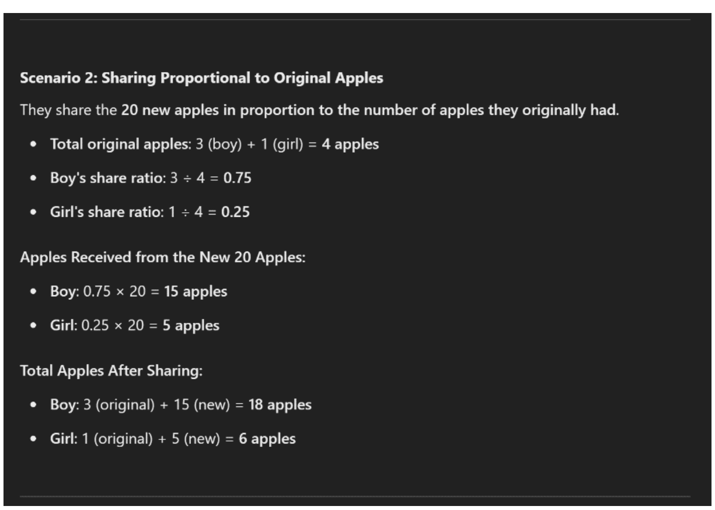

Scenario 2: Sharing Proportional to Original Apples

They share the **20 new apples in proportion to the number of apples they originally had.**

- **Total original apples**: 3 (boy) + 1 (girl) = **4 apples**
- **Boy's share ratio**: 3 ÷ 4 = **0.75**
- **Girl's share ratio**: 1 ÷ 4 = **0.25**

Apples Received from the New 20 Apples:

- **Boy**: 0.75 × 20 = **15 apples**
- **Girl**: 0.25 × 20 = **5 apples**

Total Apples After Sharing:

- **Boy**: 3 (original) + 15 (new) = **18 apples**
- **Girl**: 1 (original) + 5 (new) = **6 apples**

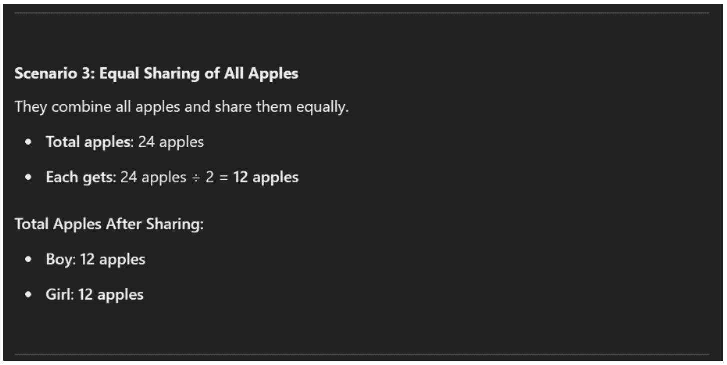

Scenario 3: Equal Sharing of All Apples

They combine all apples and share them equally.

- **Total apples**: 24 apples
- **Each gets**: 24 apples ÷ 2 = **12 apples**

Total Apples After Sharing:

- **Boy: 12 apples**
- **Girl: 12 apples**

The resolution of this problem illustrates the power of the deductive mode in the new version. Its application to complex cases in research & development, mathematics, and planning offers a power equivalent to, or even greater than, human intelligence alone.

In this sense, the introduction of AI in an organization cannot be reduced to the usual technological change management. As mentioned earlier, generative AI like ChatGPT presents itself more as a new stakeholder finding its place in the organization, like a new collaborator with superpowers.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The contribution of this TRAIDA card to your AI strategy depends on the answer to this question: ***'Do you consider AI as just another technology or as the embodiment of a stakeholder to be integrated into your organization?***'. When this question applies to Internet technologies, mobile telephony, or blockchain, the idea of them being stakeholders in your organization does not come to mind. These technologies are tools serving human actors.

With AI, particularly in its generative form, the conversational aspect and the other characteristics mentioned above invite us to consider it as a stakeholder. In this case, it is no longer just a tool serving human actors but a new collaborator to be integrated into the organization. More precisely, an unlimited number of new collaborators trained to intervene in processes in the form of AI assistants. They possess superpowers that humans do not have, such as the ability to instantly assimilate large amounts of knowledge, work continuously, and multiply at a low cost. They also present disadvantages, which the TRAIDA approach controls, such as the lack of reliability due to poor-quality data.

The deeper this AI penetrates the organization, the more human actors must collaborate with it to train, improve, and monitor it. This collaborative working mode reinforces the embodiment of AI as a full-fledged stakeholder that requires a new kind of change management.

MINDSET

In the business world, the purpose of actors is to create value. In other words, an individual whose contribution is not sufficient has no future. Regardless of the level of responsibility and expertise, everyone must participate in the wealth produced by the organization. For this value to be sustainable and growing, it relies on a collective effort. Indeed, no one is skilled in everything, and the complexity of organizations requires a division of tasks with overall coordination. The quality of interactions between colleagues, managers, subordinates, clients, and partners determines value creation. Individuals who work like free agents are rarely long-term creators of wealth. **To succeed, the mindset must therefore be oriented towards the collective**.

Thus, although AI is useful for improving individual productivity, its more strategic and profitable contribution emerges when it acts on interactions between the stakeholders of the organization. To operate at this systemic level, each individual must learn to use AI with a collective mindset. In other words, the more AI penetrates the organization, the more actors must excel in how they communicate, exchange, and collaborate with each other and with AI. This is not about technical skills but general aptitudes in human relations.

Immersed in an AI-augmented company, an individual with limited skills in writing, analysis, sharing, and innovation will find it difficult to fit in. They will not be able to properly train their AI assistants, analyze the responses obtained to enhance the training and improve results, or share them with colleagues and other AI assistants. They will become a hindrance to the organization's velocity, and their work could be called into question. Conversely, an actor with relational skills will interact better with AI and with stakeholders who also use AI. They will be a positive contributor in this new environment of more dynamic, complex, and intelligent interaction.

To embody these relational skills, TRAIDA uses the WASI approach, an acronym for the following skills: 'Write, Analysis, Share, and Innovate.' Integrating AI without WASI skills means using it merely as a tool to improve individual productivity. The accumulation of these gains does not guarantee the triggering of sufficient overall benefit for the organization. Yet, as we have already mentioned, the profound profitability of AI lies at the level of interactions between actors. Knowledge is then accumulated and formalized to share, enhance, secure, and project it into AI, thereby generating reinforced gains. In this context, the explanation of the skills highlighted by the WASI approach is as follows:

1. **Write**: Writing down knowledge strengthens mastery and enables improvement. An actor who cannot put their expertise into writing is less effective than someone who can. Regardless of the level of expertise and the field, every individual must be able to document their work regularly to improve.
2. **Analyze**: Analysis is the prerequisite for good knowledge writing. A well-written text does not rely solely on correct syntax and grammar. It is also essential to step back from the knowledge, dedicate time to observation and listening, and synthesize a clear and relevant thought.
3. **Share**: Knowledge is easier to share when it is formalized in writing. Although videos and podcasts are widely used for knowledge dissemination, their quality depends on the clarity of their authors. If they have not put in the effort to write down their knowledge to deepen it, their multimedia content is often mediocre. Moreover, the learner engages their intellect more effectively with a written document than with a video or podcast.
4. **Innovate**: Updating knowledge is faster and richer when it is formalized in writing. By sharing texts, a confrontation of knowledge begins, contributing to innovation. Conversely, it is more difficult to innovate from a series of videos or podcasts, which do not facilitate the mental construction of mapping to grasp the complexity of knowledge. In other words, writing is the most appropriate format for innovation."

WASI skills are useful regardless of AI use, but they become essential with it. Particularly with generative AI, it is important to write knowledge in a rich, clear, and relevant manner to train AI engines. It is also important for actors to have a critical mindset to analyze AI results and engage in discussions with AI to improve responses. Finally, as we have already mentioned, the profitability of AI is much stronger when it operates at the organizational level and not just as an individual productivity tool. To achieve this, sharing and innovation are key skills for actors who successfully operate in the AI universe. Conversely, actors who are not supported to enhance their performance in WASI will not be able to create the expected gains with AI.

TRAINING FOR BUSINESS

Operational training in the use of AI tools should be planned. However, their effectiveness depends on general skills necessary for working with AI, covered in the following training areas:

- **Aptitude for formalizing individual and collective knowledge in writing**. This involves transforming tacit know-how into a wealth of explicit knowledge. The application of this training is based on defining a new role within the organization, with users responsible for overseeing the accumulation of knowledge: the Knowledge Accumulation Leader (ACL). Rather than a central team imposing a common mode of operation from the start of AI implementation, it is beneficial to appoint an ACL in each department of the company. This way, teams can organize autonomously according to their skills, work habits, and availability. Coordination among ACLs encourages the sharing of certain practices and supports the formalization of knowledge at the collective level. Written knowledge is then utilized through the following use cases:
  a) AI training. This includes knowledge governance to ensure the maintenance of training, the security of the information used, and the ability to audit AI responses.
  b) The creation of a knowledge graph-oriented database that is paired with generative AI to automatically load documents. Tacit knowledge then becomes explicitly usable to comply with regulations, train new actors, conduct organizational optimization studies, explore opportunities, or perform benchmarking. This reference system also supports governance functions (a) and corresponds to the Enterprise Knowledge Graph (EKG) of TRAIDA (see EKG card).
- **Aptitude for identifying tasks that benefit from partial or total automation with AI**. According to McKinsey (2024), on average, 70% of each actor's activity can be automated by 50%. These ratios are useful for setting individual productivity goals, then extending them to each team and the organization as a whole.
- **Aptitude for supporting personal development so that actors engage positively in their work with AI**. It is important for each individual to understand why formalizing, sharing, and enhancing knowledge is strategic for producing more efficient AI. In this regard, the consideration of AI in career planning is outlined. For example, an individual who does not mention an AI assistant on their resume has a lower level of employability than candidates with the same profile who are proficient in AI.

Finally, raising awareness of the systemic aspect of AI and data management is necessary, supported by the TRAIDA masterclass planned as part of the technical team training (see the following section).

TRAINING FOR IT

For technical teams, that is, IT professionals, the following training areas are priorities for scaling AI in the company:

- **Semantic modeling**: This is the essential discipline for creating ontologies that form the core of the semantic platform for AI recommended in TRAIDA (see the documents on ODS, MDM, and EKG data repositories).
- **Enterprise architecture and enterprise governance**: These are the two pillars for managing complexity and the governance of AI associated with data management solutions (see the respective TRAIDA documents).
- **Transformative AI and Data solutions**: This refers to the TRAIDA framework with all its technical, governance, and business cards so that each IT professional is aware of the systemic aspect of AI impacts and data management. This is a quick one-day awareness session in the form of a TRAIDA masterclass that also includes business actors.

Operational training on AI tools, AI automatic code generation techniques, the use of AI for testing, data science, etc., should be planned according to the IT professionals' profiles.

Finally, the general skills we mentioned for business teams are also useful for technical teams.

Based on these training areas, each company adapts its own support programs for business and technical teams, depending on the existing and missing skills, as well as the transformation projects to be managed.

TRUSTED AI

We mentioned earlier that the introduction of AI into the organization is similar to welcoming a new stakeholder. It has exceptional capabilities for machine learning on large volumes of information to replicate and enhance human work, automate decision-making, and perform tasks, including those in the physical world using robots.

Each AI assistant becomes the companion of an actor, a team, or a decision-maker and deeply integrates into management processes.

It is likely that the first action in your workday will be to ask your AI assistant about the tasks to perform and how to approach them, including automatic responses to emails, reports to review, summaries to draft, recruitment suggestions, expense and savings proposals, activation of robots in a workshop, etc.

With hundreds of AI assistants spreading throughout the organization, it is essential to implement trustworthy AI to maintain control.

Thus, in addition to the best TRAIDA practices that enhance data quality and governance, an independent artificial intelligence is trained to monitor the functioning of the information system. This particular AI contributes to achieving trusted AI within the company. For example, it observes the behavior of AI assistants to detect anomalies that may violate predefined compliance rules, especially in terms of security and ethics.

This supervisory AI, also called the second brain or nerve center, must be considered in the human resource management approach for the following reasons:

- It requires the establishment of a trustworthy AI manager whose role is to collect all the documents and rules the information system must follow to train the supervisory AI. They collaborate with the enterprise architecture and governance managers, who formalize these rules and manage their maintenance.
- The very existence of this supervisory AI indicates that new AI stakeholders spreading throughout the company, notably in the form of assistants, must be monitored. It is thus relevant for human resource management to contribute to the proper governance of these new stakeholders, for example, by setting objectives and defining usage rules that comply with the company’s HR and ethical policies.

From a technical perspective, the trusted or supervisory AI relies on a knowledge graph-oriented database augmented with generative AI to build the reference framework of rules to follow (see the TRAIDA EKG – Enterprise Knowledge Graph card). This reference framework is fed with all documents describing the expected behavior of the information system and regulatory texts. During its execution, this AI compares the outputs of the information system (data flows, calculation results, activation of AI assistants...) with the rule reference to detect any abnormal operating cases.

3. BLUEPRINT

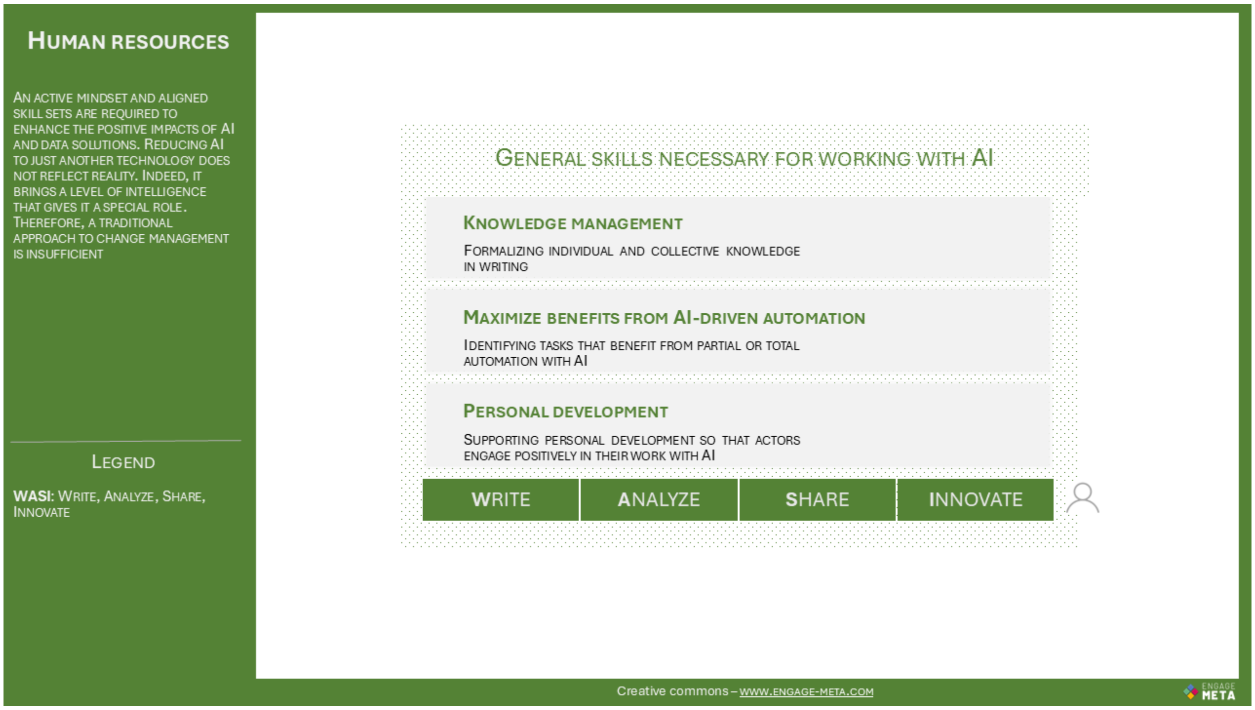

4. YOUR SITUATION & OBJECTIVES

ENTERPRISE ARCHITECTURE

Enterprise Architecture (EA) outlines practices for modeling and documenting the business system. It enables the preparation and support for large-scale deployment of AI by promoting the consideration of semantic modeling (ontology) and service-oriented architecture (SOA).

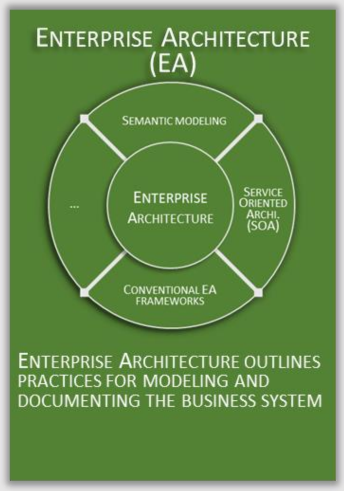

1. CONDITIONS OF SUCCESS

The profitability of AI relies on the use of best practices described in TRAIDA, particularly those concerning data quality, ontology modeling, and knowledge management. Their implementation is closely linked with the company's information system, which includes the processes, rules, and data that support the execution of operations.

With TRAIDA, the goal is not to create a new AI-based system from scratch that would operate parallel to the existing one, but rather to promote a symbiosis between AI and the information system. To extend the metaphor, it's similar to the relationship between a clownfish and an anemone. Both derive benefits: the fish is immune to the stinging tentacles of the anemone, allowing it to hide there, and the anemone feeds on the fish's waste. The coupling is the same for AI and the information system. One cannot survive sustainably without the other, especially when it comes to maintaining the company's competitiveness through new information management technologies.

This coupling revolves around the value chain of the information system, which starts with the organizational processes (a) operated by the company's actors. These processes trigger rules (b) that exploit data (c). Like any chain (a-b-c), its strength depends on its weakest element. A defect in one of these intangible assets—processes (a), rules (b), or data (c)—impairs the execution of the whole. For example, an information system built around rigid silos is prone to data quality defects, which hampers the proper execution of rules. In other words, the interdependence between processes, rules, and data leaves no room for errors in any of the assets. A defect in any one of them, even minor, can have negative consequences for all the others.

Given the importance of this value chain, the integration of AI must preserve its quality. Moreover, it should contribute to greater efficiency while respecting the integrity of the three intangible assets. For example, when AI calculates the assignment of a task to an actor within an organization, the reasons behind this decision must be auditable according to the elements of the value chain:

- Processes (a): Optimize the actors' time within a general planning framework.
- Rules (b): Determine whether a treatment should be automatic, manual, or mixed, depending on the nature of the case and the regulatory context.
- Data (c): Analyze the case requiring the task to determine its nature within a predefined classification, then verify compatibility with regulatory clauses that must be adhered to.

An AI-based system that opaquely mixes several of these levels would make decision audibility and overall system maintenance difficult. In other words, each level or type of intangible asset in the information system has its own AI system dedicated to its specific concern. Of course, these levels accumulate according to automation:

- Processes: Decision-making about the steps the organization must follow to meet a need, respond to an unforeseen event, comply with a regulatory requirement (organizational level), request an actor within a team, etc.
- Rules: Automation of calculations, decision support, deduction, etc.
- Data: Analysis, compliance, consolidation, aggregation, pattern recognition, content generation, preparation of data sets for rule execution, regulatory verification (business level), etc.

Without quality data, and processes and rules that are formalized and executed with precision, it is difficult to implement intelligent and reliable algorithms whose responsibilities are clearly defined at each level. To achieve this, enterprise architecture helps improve the quality of these intangible assets within the information system:

- Instead of implementing fragmented databases with disorganized processes and rules, the goal is to streamline the quality of these assets to leverage digitization and the benefits of AI with a sufficient level of reliability and security. Good mastery of enterprise architecture helps achieve this objective.

Basic principles of enterprise architecture

Enterprise Architecture (EA) is a set of practices for classifying, modeling, and improving the quality of the intangible assets of an information system, namely the processes, rules, and data that a company uses to conduct its activities. The desired contributions are multiple, aimed at enhancing the company's operations with transparency, auditability, agility, security, regulatory compliance, reliability, and efficiency.

More concretely, EA is structured using four layers of abstraction:

1. **Business architecture**: this layer describes the processes used by the company to conduct its operations. These processes rely on steps that are automated by IT tools, as well as others that are manual. The scope of analysis for business architecture is therefore global to the company, extending beyond the IT perimeter.
2. **Data architecture**: this layer describes the architecture of the data, primarily focusing on databases. The scope of analysis is thus limited to the IT domain.
3. **Application architecture**: this layer describes the IT applications, distinguishing between different software solutions, such as custom developments and packaged software.
4. **Technical architecture**: this layer describes the technical infrastructure, including hardware, networks, and security devices.

Integration of enterprise architecture with AI

To take into account the ontologies and knowledge management required for optimal AI deployment, enterprise architecture is strengthened at two levels:

1. **The business architecture layer**: it is extended to include the consideration of ontologies to integrate data modeling at the semantic level in a unified manner, independent of databases. Thus, business architecture is no longer limited to merely analyzing processes.
2. **The data architecture layer**: in addition to the data stored in traditional databases, this layer deals with knowledge using the Enterprise Knowledge Graph (EKG) repository as described in TRAIDA. Therefore, data architecture is no longer confined to the IT perimeter; it opens up to the formalization of knowledge, independent of its level of computerization.

By adding these two components, enterprise architecture becomes an interesting framework to support large-scale AI deployment within the company, for the following three reasons:

1. **Modeling of intangible assets of the information system (processes, rules, and data)**: this modeling clarifies the value chain. It is from this chain that the best anchor points for AI-based algorithms are identified in a more efficient, reliable, and traceable manner. Without a layered structure (processes, rules, and data), the use of AI systems tends to mix organizational, business, and data processing levels, making overall governance more difficult (auditability, maintainability, impact analysis, etc.).
2. **Ontology modeling at the business architecture level**: this modeling reinforces the importance of this practice to successfully implement AI on a large scale. It is the core of the semantic platform recommended by the TRAIDA architecture.
3. **Knowledge management consideration at the same level as data architecture**: this consideration contributes to the proper integration of the Enterprise Knowledge Graph (EKG) repository.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

Enterprise architecture is primarily implemented with the objective of documenting the information system. Given the complexity of databases, workflows, and applications, the resulting descriptions are complex. They focus on a high level of abstraction that does not allow for concrete action to transform systems.

Worse still, since the information system is constantly evolving to meet business needs and regulatory requirements, the documentation produced by EA is rarely up to date.

Although this observation is unfavorable, the documentation effort supported by EA, even if imperfect, must be maintained; otherwise, the overall knowledge of the information system will be lost.

A more positive use of EA is nonetheless possible, especially by taking into account the specific needs for the deployment of AI systems. Indeed, adding ontologies at the business architecture layer solidifies the documentation. It makes it more sustainable than simply modeling around processes. In other words, the velocity of changes at the ontology level is lower than that observed on processes:

By clearly distinguishing the two documentation spaces (ontologies, processes), knowledge capture is more robust, easier to update, and actionable.

At the data architecture layer, knowledge management paves the way for a more powerful enterprise architecture than simply documenting databases.

This renovation of enterprise architecture is important based on your experience:

- You already have EA practice in your company, with mixed results or even a perceived failure. Rather than abandoning this documentation effort, you will take advantage of integrating AI to renovate your enterprise architecture practices. Although you consider EA as an abstract approach, distant from project needs, costly... it remains essential to better manage complexity and, therefore, contribute to better AI integration.
- You have no experience in EA and identify the need for an architectural framework to embody the information system's value chain based on processes, rules, and data. This is necessary to better control the integration points of AI-based systems in your information system.

More generally, the power of generative AI allows for the optimization of information system documentation through the following use cases:

- Automatic detection of discrepancies between documentation in the EA repository and a database or applications. Generative AI can absorb all documents, database schemas, specifications, and user guides for applications to check the overall consistency of the documentation.

- Automatic generation of documentation from descriptions of existing systems, whether they are databases, applications, data flows, regulations, security directives, user support tickets, etc. This is an opportunity to accelerate re-engineering.
- Storing EA documentation in a knowledge graph-oriented database to obtain visual representations (graphs), conduct analyses, etc. This repository functions like an Enterprise Knowledge Graph (EKG) as described in TRAIDA, but applied to knowledge about the information system.

This EKG repository, provided there is a version that describes the expected behavior of the information system, can then be used to train an AI aimed at observing the actual behavior of the information system to detect deviations from expectations. This AI is referred to as a second brain, nerve center, or trusted-AI.

CONVENTIONAL EA FRAMEWORKS

Due to their complexity, conventional enterprise architecture frameworks, such as TOGAF or Zachman, are difficult to make profitable. They are mainly used by experts, in isolation from operational teams like application designers and software developers. These frameworks are presented as repositories of best practices that cover all aspects of formalizing and governing information systems. This encyclopedic positioning does not facilitate their readability or usability.

This criticism could also apply to TRAIDA, as it is a type of enterprise architecture framework. However, since it is limited to the domain of transformative AI and data management solutions, the number of best practices remains small, which makes it easier to read and allows it to be actionable for operational teams. Conversely, because it does not cover all concerns related to the formalization and governance of information systems, it does not replace conventional frameworks.

Thus, the encyclopedic nature of EA frameworks makes them too general to provide concrete guidance, and they lack specificity to be used as a structured methodology. Although they are useful for initiating an enterprise architecture culture, it is necessary to customize them for each company's context, especially considering AI.

A major flaw of conventional enterprise architecture lies in its process-oriented approach. Indeed, the first layer of enterprise architecture emphasizes processes alone to document the information system. The next layer focuses on data; however, it is limited to logical and physical models. Consequently, the conceptual level of data is not a priority in conventional enterprise architecture frameworks.

There is a lack of conceptual data analysis, which should occur at the business architecture level, in alignment with process documentation. Without considering data early enough in the analysis, there is a risk of documenting processes without generating a positive impact on data architecture. As a result, data remains confined in silos centered around processes, leading to duplication and quality issues. This lack of data governance hinders the effective use of digitalization and AI.

A positive aspect of conventional enterprise architecture is its business focus, which aims to document processes beyond the scope of IT tools. However, as mentioned above, it is problematic that this business level is not also seen as an opportunity to model data conceptually, that is, ontologies independently of their implementation in IT databases. Worse, the next level of the data architecture layer is reduced to the IT domain, essentially databases. Yet, a wealth of information exists beyond IT systems that companies must exploit for training AI systems.

Contributions to improve EA coupled with AI

Based on these observations, conventional enterprise architecture is enhanced to serve as a facilitator for large-scale AI integration. It is then necessary to consider the four contributions presented below.

1. **Key strategic goals for enterprise architecture : t**he first contribution ensures that the enterprise architecture approach is understood and adopted by all stakeholders. It is necessary to define its objectives, including those for integrating artificial intelligence. A document titled "Key Strategic Goals for Enterprise Architecture" is prepared, containing the following four chapters:

- MOTION to clarify the objectives of the EA+AI approach.
- ENGAGEMENT to identify the tasks that need to be prioritized to achieve tangible results with EA+AI, typically within a timeframe of less than six months to avoid a tunnel effect.
- TREASURY to allocate the necessary financial resources for EA+AI operations. This chapter also describes the rules for calculating return on investment.
- ASSURANCE to outline the major governance rules for EA-AI.

To be useful to the organization, this document should be drafted in the shortest possible time, in an initial version. Ideally, within two to three weeks, stakeholders should agree on key objectives accompanied by concrete results to be delivered. It is then updated based on the progress of EA+AI in the company, usually on an annual basis.

2. **Data governance**: the second contribution is data governance with semantic modeling (ontologies). Its integration starts at the business architecture level. This step is crucial for improving data quality, a prerequisite for the effective integration of AI at the enterprise scale. This holistic approach to data analysis promotes a comprehensive understanding of the flow of information within the organization and its interactions with processes. Data governance is thus established, ensuring data quality, integrity, and accessibility. Moreover, this holistic approach ensures synergy between digitalization, AI, and business processes. It enables organizations to leverage data-driven analytics for smarter decision-making, better process automation, and ultimately greater efficiency and competitiveness.
3. **Knowledge governance**: the third contribution concerns knowledge governance. It begins with ontology modeling (as mentioned in the previous point). It now extends to the formalization of knowledge at the data architecture level. Instead of limiting the analysis to the IT scope alone, knowledge governance focuses on data that is not yet digitized, whether structured, unstructured, already formalized in writing, or tacit as individual and collective knowledge. In TRAIDA, this involves building the Enterprise Knowledge Graph (EKG), which is essential for training AI systems.
4. **AI-assisted automated governance**: the final contribution is AI-assisted automated governance. This plays a crucial role in controlling the entire information system. This mechanism is based on implementing an intelligence layer above the information system to supervise decision-making algorithms. It is known by various names, such as second brain, nerve center, or trusted AI. This AI is continuously fed with software specifications, application documentation, data structures, regulations, KPIs, etc. It observes the behavior of processes and the information system as a whole to alert on executions that do not meet expectations.

SEMANTIC MODELING

The deployment of AI systems in companies and on a large scale need to use a lot of data from the company's databases, both during their training and during prompts to enrich requests (RAG: Retrieval Augmented Generation). Since these databases and other sources such as files, archives, etc., are often heterogeneous and of varying quality, it is dangerous to connect the AIs directly to these storage areas. It is smarter to build a unified vision of all the company's data using a powerful business model that sits in front of the heterogeneous storage areas (digital twin). The AIs can then draw their data from a clean source, accompanied by security rules.

Software platforms exist for setting up this kind of system, either with a graph-oriented database approach or with the NoCode database. But regardless of the technology used, an effort of modeling is required to achieve this unified vision of the data. It also needs to be done in a way that allows for its evolution to keep up with business changes that occur regularly. Therefore, the model must be both very clear and strict in quality management, but also well-constructed enough to accept extensions without questioning everything.

This modeling involves expertise in ontology construction, also known as the art of semantic modeling. Ontology is the art of documenting the business concepts of the company and defining their relationships as well as the rules for controlling their quality.

First of all, a business concept is a key management entity for the company, such as a Client, Supplier, Invoice, Production Unit, etc. A startup has about fifteen of these, an SME more than twenty, and a large company even more. Each business concept is defined to constitute a glossary shared by the entire company. It is accompanied by a thesaurus to standardize term equivalences.

Next, the business concepts are organized into a hierarchy that describes the parent-child structures that exist between them. For example, a Client is specialized by B2B, Retail markets, etc.

Once the glossary, taxonomy, and hierarchy are in place, it is time to model the attributes of the business concepts and specify the relationships they have with each other. The semantic power of the data model greatly depends on the quality of the modeling of the relationships between business concepts. The first time you do semantic modeling, be accompanied by an expert in this discipline, at least to verify that your model is solid. You can also use an AI assistant for data modeling, but you will need to train it well before it can help you effectively.

At the end of this semantic modeling, you will have built your ontologies. At this stage, it is still a static vision of unified data. A final modeling step is needed to add a more dynamic dimension. Its purpose is to control the quality of the data contained in the business concepts. These are axioms that are added to the ontology. Here, focus on universal control rules that do not depend on organizational choices. A powerful way to formalize these business axioms is to use state machines. For example, a Product business concept could have this list of possible states: R&D, Offer Catalog, Maintenance, Out of Sale… Depending on the state of a product (instance of the Product business concept), update, delete, and usage actions are possible or not.

List of key advantages of having well-constructed ontologies

They allow the implementation of a unified data layer in front of your heterogeneous databases, or if you are starting from scratch, to have a very clean database that will follow your business evolution without creating chaos for data storage. This approach creates a digital twin of your IT on which you can plug your AIs both for their training and for prompt augmentation (RAG) by fetching real-time data in vectorized ontology instances.

They provide the necessary classification for organizing knowledge, beyond data from databases. To better train your AIs, you will need to formalize your organization's tacit knowledge, i.e., what your teams know but is not yet documented or well explained. All explicit knowledge is then loaded into the ontologies to complement structured data, thus increasing the knowledge base used by the AIs.

During prompt execution, real-time access to ontologies allows on-the-fly enrichment of the request context, enabling the AIs to work better. This is the principle of RAG.

Conversely, during the reception of the AI-generated response, access to ontologies will allow verifying the quality of the result, for example, by checking the data sources used. This significantly reduces the negative effects of hallucinations when the AI is not used in a creative context but rather for deterministic analysis.

If you are starting a business, a NoCode database with ontologies is the right way to go. If you already have an existing setup, you still need ontologies, but perhaps with a technological choice oriented towards NoCode and graph-oriented databases. Depending on the scope of your existing IT, you will need to consider the best data architecture. The TRAIDA cards will help you to decide the best choice in your context.

SERVICE ORIENTED ARCHITECTURE (SOA)

Service-Oriented Architecture (SOA) emerged in the early 1990s with the advent of Client/Server solutions. Today, it remains a valuable approach for structuring the information system around reusable services. It is reinforced by cloud platforms' microservices and DevOps engineering. It also benefits from improved implementation through databases organized by business object domains, known as data mesh architectures.

SOA is addressed at the data architecture level in the EA approach, then extends into application architecture and technical architecture. It facilitates AI implementation by clarifying the levels of responsibility in software execution (service providers and consumers) and enabling their reuse. The more sustainable the software architecture, the better AI integration benefits from good governance. Drawing a parallel with data quality, SOA acts as a powerful tool to enhance software quality, thereby improving the quality of integrating new digitalization technologies, especially AI.

To better understand this benefit, it is important to recall the essential properties that services bring with SOA architecture. A service is a process that adheres to the five properties detailed below: loose coupling, remote and interoperable activation, asynchronous operation, exposes a usage contract (interface), and complies with the SOA architecture pattern.

Property #1: Loose coupling

- A service cannot directly call another service. It delegates this function to a process specialized in chaining (orchestration).
- A service can be activated independently of its technology. To do this, activation is performed by sending (and receiving) an XML message. Therefore, it is not a binary call.
- A service can be activated in an asynchronous mode. In this case, the service subscribes to an event via an orchestration function.

Property #2: Remote and interoperable activation

- A service exposes a usage interface that is consistent regardless of its network location. The service call works regardless of the consumer's language and operating system. To promote interoperability, XML is preferred. When the service is backed by a data mesh database, the technical architecture is termed microservices. In this case, the service is reusable in a plug-and-play mode since it operates with its dedicated database for its functional scope. In practice, this type of micro-database handles a set of services around a business concept derived from ontologies.

Property #3: Asynchronous operation

- A service operates asynchronously, meaning it does not block the consumer while it executes. This principle is useful for reducing bottlenecks (performance, robustness). This type of architecture is known as Event Driven Architecture (EDA) and is combined with SOA.

Property #4: Exposes a usage contract

- A service exposes a usage contract described in two parts. The abstract part declares the input and response messages of the provided service. The concrete part describes the technical standards and protocols used for service activation. Depending on the implementation and deployment choices, there can be multiple concrete parts for the same abstract part. The usage contract is also referred to as a service interface expression.

Property #5: Complies with the SOA architecture pattern

- The SOA architecture pattern involves creating an application architecture that breaks down processes into services attached to class packages. These packages form categories (business objects, concepts, or business subjects), each with an access facade that contains all the services it exposes (also referred to as a port).

Types of services

- **Business service** : this is the highest-level service in the SOA architecture. It is directly understandable by users, meaning the service providers and consumers.

- **Service exposed** by a business concept : this service is situated at the application architecture level. It represents the preferred unit for managing and reusing services. Using a "data mesh" approach allows for autonomous coupling of these services with the underlying databases that handle their business concepts.
- **Internal service to a business concept** : these are the services that implement those exposed by business concepts. These services operate at the level of detailed software engineering, particularly with components. They are not visible to users.

3. BLUEPRINT

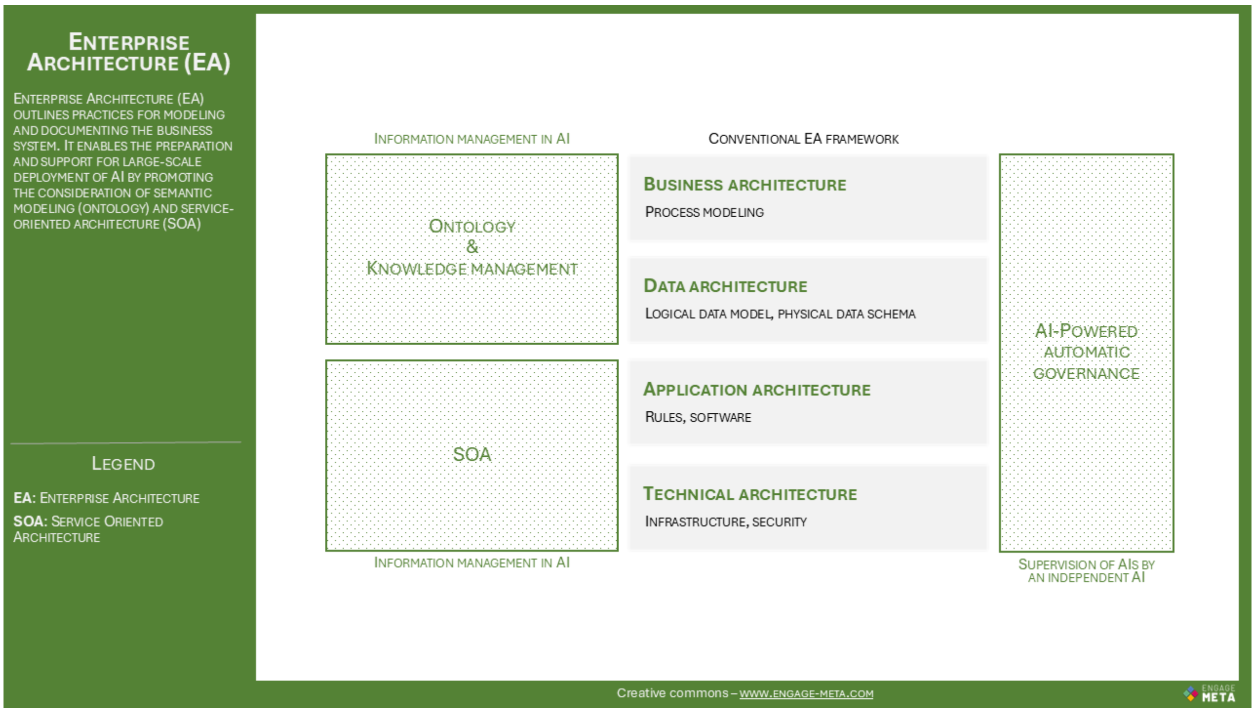

4. YOUR SITUATION & OBJECTIVES

ENTERPRISE GOVERNANCE

Enterprise governance aims to ensure the quality of data and AI across the organization. It revolves around risk management and regulatory compliance, the application of ESG (Environmental, Social, and Governance) and CSR (Corporate Social Responsibility) principles, as well as ensuring the reliability of the IT system.

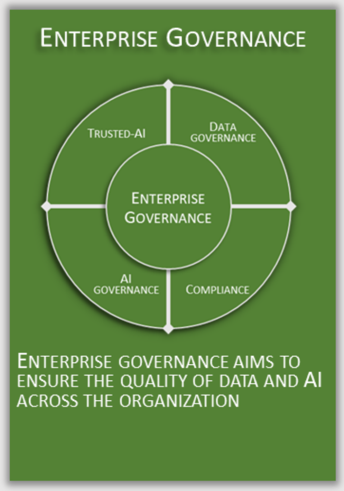

1. CONDITIONS OF SUCCESS

Enterprise governance ensures that decision-making and management processes are conducted in compliance with the company's internal rules and regulations. Given the complexity of the organization, it often mobilizes significant human and technical resources. These resources focus on two major areas: risk control and regulatory compliance. The following key domains of governance are then considered:

- The management of internal risks and compliance with industry-specific regulations.
- The application of ESG (Environmental, Social, and Governance) principles for non-financial performance and their translation into regulations.
- The application of CSR (Corporate Social Responsibility) principles and their translation into regulations.

IT management is delegated to the governance of the information system, which uses frameworks such as COBIT and ITIL, as well as enterprise architecture with TOGAF (see TRAIDA card on Enterprise Architecture).

Governance quality

The quality of enterprise governance increases with its level of automation.

In other words, the less human intervention is required to execute processes, the more governance is embedded in the software code. For example, the control of an expense commitment amount, based on a matrix that cross-references actors and needs, is integrated into the order processing software. However, if this matrix or software has flaws, it challenges governance as seriously as a human error would. A balance between controlled automation and human intervention is a goal to be clarified, especially since AI enhances this potential for automation and shifts the usual balance.

It introduces new use cases depending on the context of each company. Here are some examples for illustration:

- A human resources management AI is integrated into the employee promotion process to automate certain decision-making steps that were previously exclusively human. Enterprise governance ensures that this AI's training aligns with HR policy and complies with regulations, such as CSR criteria.
- The organization finds that increasing the use of AI for decision-making correlates with a decrease in informal communication between actors. Enterprise governance examines the risk of

deteriorating social relations and its consequences on the organization's non-financial performance according to ESG criteria.

- AI-augmented humanoid robots have the potential to replace workstations in a factory. Enterprise governance studies this opportunity and demonstrates that wealth creation increases while the wage bill decreases. The social impact of this change must not be left unanswered. It is accompanied by a retraining plan for factory employees and a new distribution of financial value (CSR).
- A company deploys AI tools alongside traditional IT solutions. This dual operation creates inconsistencies between AI decisions and application systems. A reassessment of enterprise architecture is necessary. In this context, enterprise governance is responsible for better controlling transformative AI and associated data solutions.
- A knowledge graph database augmented with AI loads regulatory text to automatically create an interactive model. Each clause of the regulation is then linked to elements of the IT system, such as actor types, roles, applications, databases, information flows, etc. When the regulation is updated, a new model is loaded to compare it with the previous one and identify differences. This automatic impact analysis is crucial for monitoring regulatory compliance. In other words, AI allows the computerization of regulatory tracking.
- A company implements a workplace wellness AI to supplement the support offered by a psychologist, whose time is limited. The confidential and empathetic nature of the relationship between the employee and this AI facilitates the request for psychological assistance at work. The ESG benefit is significant, provided the AI has undergone relevant and respectful training in line with the company's values.
- An AI is implemented by a works council to enhance employee assistance and training on workplace hygiene and safety conditions. Each employee has a personalized assistant at their workstation, helping them improve their tasks. The CSR score improves as the rate of illness and accidents at work decreases thanks to this AI assistant.
- To better detect IT system execution failures, an AI is trained with the company's application, process, and database specifications. This AI captures data flows exchanged between applications, execution reports, database contents... to identify deviations between the IT system's execution and its specifications.

In conclusion, enterprise governance plays a critical role in ensuring that decision-making processes are aligned with internal policies and regulatory frameworks, while balancing human oversight and automation.

As AI increasingly integrates into various aspects of governance, it offers both opportunities for improved efficiency and challenges related to maintaining social responsibility and regulatory compliance. Achieving the right balance between automated systems and human intervention will be key to enhancing governance quality in an ever-evolving technological landscape.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The large-scale deployment of AI is accompanied by enterprise governance that focuses on the following topics: data governance, AI governance, compliance, and Trusted-AI.

DATA GOVERNANCE

Data governance contributes to achieving the required level of data quality so that AI can operate at an enterprise scale in a reliable and cost-effective manner.

It integrates the three data repositories: MDM, ODS, and EKG, which structure the core of the semantic platform (see respective TRAIDA cards). Ideally, this governance is shared among the three repositories. However, depending on the technologies chosen by the enterprise, more or less complex integration efforts are necessary. For example, creating a dataset that combines master data (MDM), transactional data (ODS), and data from a knowledge graph (EKG) requires the provision of a common governance function, according to the following scenarios:

- Master and transactional data involved in the dataset are copied into a knowledge graph loading area. In this case, the EKG repository is chosen as the central point for the dataset.
- The three repositories are built on the same technology, ensuring unified dataset creation functionality.
- An independent database is used to implement a data loading function from MDM, ODS, and EKG. This could involve a dedicated storage space integrated into a data fabric solution.

Other technical scenarios are possible and depend both on the enterprise context and the governance functions to be implemented. These functions are numerous, such as:

- Datasets and data spaces.
- Versions and variants with comparison and merging solutions, both on datasets and data spaces, as well as on data models (ontologies, dictionaries, metadata, etc.).
- Processes for adding, archiving, and removing data.
- UI for administration, day-to-day management, and reporting.
- Data cleaning and deduplication.
- Security.
- History, traceability of operations on the data, and archiving.
- Integration between repositories and with applications.
- And more.

To control the roadmap and costs of these functions, it is necessary to establish a vision that considers the following elements:

- The quality level of the governance functions natively offered by each of the technologies used for the MDM, ODS, and EKG repositories. Consequently, their choice considers governance needs to avoid hidden costs during implementation. It is unfortunate to start with a repository that seems technically attractive but has governance functions too weak for deployment at the enterprise level.
- Clarification of the end-to-end integration principles of the three repositories (MDM, ODS, and EKG) to avoid heavy technological barriers that generate insurmountable technical costs for large-scale deployment. Consideration should be given to the implementation of an independent data hub or one integrated into a data fabric (see TRAIDA Data Integration card).

Data governance also defines the organizational processes that provide these functions to the stakeholders involved in data management, such as:

- **Data owner**: responsible for the data concerning its structure (model, metadata, dictionaries), values, and uses.
- **Data modeler, data analyst, data architect**: stakeholders in charge of data modeling and organizing it in an enterprise architecture model that defines data domains and integrity rules

between their boundaries. With a data mesh approach (see TRAIDA Core system data card), modeling is done by business domain to elevate the data to the level of a managed product, rather than through application silos, which lead to duplication and quality issues.

- **Data steward**: operational user of the data, according to usage rights (read, write, copy, delete, etc.).
- **Data scientist**: stakeholder responsible for using the data for analysis purposes (see TRAIDA Data lake warehouse and Artificial intelligence cards).

AI GOVERNANCE

AI governance relies on functionalities, some of which are already covered by data governance. This includes, in particular, the consolidation of datasets with their archiving and version management, as well as the functions of cleaning, enriching, and analyzing the quality of datasets.

Specific functions of AI governance include, for example:

- Unified interface for accessing various AI engines.
- Data labeling.
- Workflow for integrating datasets with AI and training requirements.
- Archiving of results.
- Auditing of results.
- Etc.

When implementing an AI governance software, it is necessary to clarify its integration with the existing data governance system. Depending on the enterprise context, different scenarios may arise, such as:

1. There is no data governance in place, and the introduction of AI governance provides an opportunity to build on a common toolset.
2. Conversely, data governance is already in place. In this case, two approaches are possible:
   a) Data governance is the preferred system for preparing datasets, which are then provided to the AI governance tool. The latter limits the use of its data management functions to focus on AI governance functions.
   b) AI governance is allowed to deploy its data governance functions independently of the existing system. In this case, the minimum coordination rules between the two governance tools must be specified to avoid inconsistencies in the datasets used and to optimize implementation and maintenance costs.

Whichever scenario is chosen, it is necessary to have a data integration tool such as a data hub or data fabric (see TRAIDA Data Integration card) to provide unified access to the semantic platform's MDM (Master Data Management), ODS (Operational Data Store), and EKG (Enterprise Knowledge Graph) repositories (see respective TRAIDA cards).

AI governance tools

At the time of writing this TRAIDA card, the available AI governance tools on the market are categorized according to the nature of the data they manage:

a) Logical or physical data flows, without a semantic layer.
b) Business concepts modeled in ontologies that hide the physical implementations of the data.
c) Prompts for generative AIs.

Approach (a) does not align with the vision of the semantic platform recommended by TRAIDA. Indeed, it is preferable to handle data at a business abstraction level, meaning ontologies with approach (b). The last approach (c) focuses on managing prompts with AI assistants and should be combined with approach (b).

AI prompt manager

The governance of prompts for generative AIs is an important topic for scaling AI. It is not enough to manage catalogs of prompts; they must also be associated with assistants, which are themselves connected to information sources. The system aims to manage prompts used to interact with GPT-like AI assistants and the resulting outputs. It also manages different versions of AI assistants, considering the data used for their training (stemming from the EKG repository). The system facilitates collaboration among team members, sharing of prompts, and tracking prompt quality improvement over time. Here are the key business concepts:

Prompt Management	• **Prompts**: Contains basic information about prompts, including the creation date and the author. For example, a prompt could be "Writing a LinkedIn post." This table provides a brief description of the general objective of the prompt. • **Prompt Text**: Stores the specific texts of the prompts used to interact with the AI. Each prompt text is associated with a prompt from the Prompts Table. This relationship includes an attribute specifying the prompt text's objective (e.g., "for the month of June"). An additional attribute indicates whether the prompt text is usable or deprecated. The author of the prompt text is also recorded.
Execution Management	• **Executions**: Records the results of prompt text executions, including the execution date, the documents used as input, and the execution result. The author of the execution and an analysis of the prompt quality (strengths and weaknesses) are also stored.
User Management	• **Users**: References authors and users present in other tables. This table contains basic information such as name and email address. A user role (administrator, editor, reader) can be added to manage permissions. • **Prompt Sharing**: Users can share prompts with other team members. A sharing attribute can indicate which users or groups the prompt is shared with.
AI Assistant Management	• **Assistants**: Contains basic information about AI assistants, including the creation date and the author. For example, an assistant might have the objective "Writing social media content." This table provides a brief description of the assistant's general objective. • **Assistant Versions**: Stores different versions of AI assistants. Each version is associated with an assistant from the Assistants Table and includes the list of prompt texts used for training as well as the list of documents (in the form of URLs) used for training. The same assistant can h²ave multiple versions with different training prompts and documents.
Document Management	• **Documents**: Stores documents used for training assistants and executing prompts in the form of URLs to storage locations and/or stemming from the EKG repository. Documents can be shared among different users and assistant versions.

COMPLIANCE

The more organizational and decision-making processes are automated, the more it becomes possible to implement the approach known as **governance by design**, which integrates governance rules directly into software. Budgets allocated for AI automation and data management solutions thus benefit from an increased return on investment due to better application of enterprise governance.

For example, controlling the quality of financial reports is a strategic aspect of enterprise governance. In immature or complex organizations, these reports may contain errors, omissions, sometimes deliberate oversights, or even fraud, which is crucial to detect as quickly as possible. Manually verifying them involves controllers whose availability is limited, and they can also make mistakes. AI then supports enterprise governance on two levels:

a) An AI is implemented independently of the applications that generate financial reports. It is trained with governance rules, past detected error cases, fraud risks, the expertise of controllers, etc. In parallel with the applications, this AI receives datasets and the generated financial reports to detect anomalies that require human oversight.

b) Anchor points to a finance-specialized AI are added to the applications generating financial reports to increase the level of control over their operations. Instead of coding all reporting rules into hard-to-audit algorithms, the most strategic parts are outsourced. This AI is independent of the AI mentioned in (a) to ensure total separation of responsibilities between them.

Depending on the use case, the types of AI invoked are selected to best support enterprise governance. Thus, the control AI (a) is likely generative, while the one connected to the applications responsible for producing financial reports is more likely symbolic (b).

Beyond this example, all processes managed by enterprise governance benefit from control AIs embedded in applications. By reducing manual interventions in favor of AI training, the risks of non-compliance with rules executed daily within the company are minimized.

Application of AI for regulatory impact monitoring

Another use case for AI in enterprise governance concerns the monitoring of regulatory impacts.

Let's take the example of a banking regulation presented in textual form, such as a PDF file with several dozen pages. This regulation lists rules to be integrated into various processes and during the management of certain strategic data. The company must then create a mapping between this regulation and its applications and databases. This involves inventory work, mapping, architecture, and is carried out collaboratively between IT and a compliance team within the company. Once this mapping is established, the IT team develops the rules and control mechanisms, prioritizing the use of AI as previously mentioned.

Beyond the initial implementation, the maintenance cycle proves to be complex:

- Applications and databases evolve, requiring updates to the rule integration. For example, when deploying a new application, it is necessary to ensure that the rules are properly integrated and executed with the correct data, among other considerations.
- The regulation itself also evolves. When a new version of the regulatory text needs to be taken into account, both the intrinsic changes must be identified, as well as the impacts on applications and databases. This requires moving rules, removing some, modifying others, etc.

An exclusively human management of this maintenance is cumbersome and not very responsive. It is also prone to errors. AI can then assist in this work, following this use case:

- The regulatory text is loaded into a knowledge graph database augmented with a generative AI. Its training allows it to detect business concepts, rules, and other fundamental entities in the regulation in order to produce a graph enabling computerized management. This results in a knowledge base. With TRAIDA, this is a use case for the Enterprise Knowledge Graph (see TRAIDA EKG card).
- Another knowledge base, created with AI, describes the company's applications and databases.
- The compliance team, in collaboration with IT, then establishes links between the regulation graph and the system graph to document the implementation of the regulation within the company.

By coupling AI with a knowledge graph database, the company has a repository containing the formalization of regulatory implementation.

When a new version of the regulation is released, a new knowledge graph is generated using the same process described earlier. A difference calculation between the current and new version instantly and unambiguously provides the exhaustive list of changes and their impacts on the IT system. Enterprise governance is thus improved in terms of responsiveness, cost efficiency, and relevance. This is a significant asset for the smooth execution of operations.

Finally, it also ensures the sustainability of the knowledge within the compliance team and IT, as their expertise is formalized and used to train the AIs. A team member who changes roles or leaves the company no longer results in a total loss of knowledge, as the AI they worked with remains in place.

TRUSTED AI

The more intelligent a system becomes, the harder it is to control internally for two reasons:

a) It is accompanied by increasing complexity and a growing body of knowledge. It becomes incomprehensible for a single human intelligence. The observation "n*o one is an expert in everything*" progresses to "***the whole escapes the understanding of the collective***."

b) It adapts to new situations through self-learning, which increases the risk of internal opacity.

With the introduction of AI at the enterprise level, the entire behavior of the enterprise system becomes more intelligent, along with the risks of losing control of it for the two aforementioned reasons.

TRAIDA proposes the implementation of a semantic platform to enhance the quality of AI and the associated data solutions. This platform helps improve internal control of the system by positioning itself as an active component in the execution of processes. However, this is not enough to guarantee enterprise governance of such a system in the long term. As we have said, the more intelligent a system becomes, the harder it is to control from the inside.

To control it externally, an autonomous intelligence of at least an equivalent level must be installed. This AI is designed to observe the system’s behavior and detect atypical operations, potential errors, possible fraud, and also to propose improvements, optimizations, preventive maintenance actions, etc.

This superintelligence acts as a second brain or central nervous system for supervision. It is trained based on regulations, documentation, and specifications, key objectives, known error and fraud cases, and a list of actors with their responsibilities to fulfill its role of overseeing proper enterprise governance. By incorporating the company's values, ethics, and social and environmental responsibility objectives, this super AI plays a global role in trust, acting as a Trusted-AI.

3. BLUEPRINT

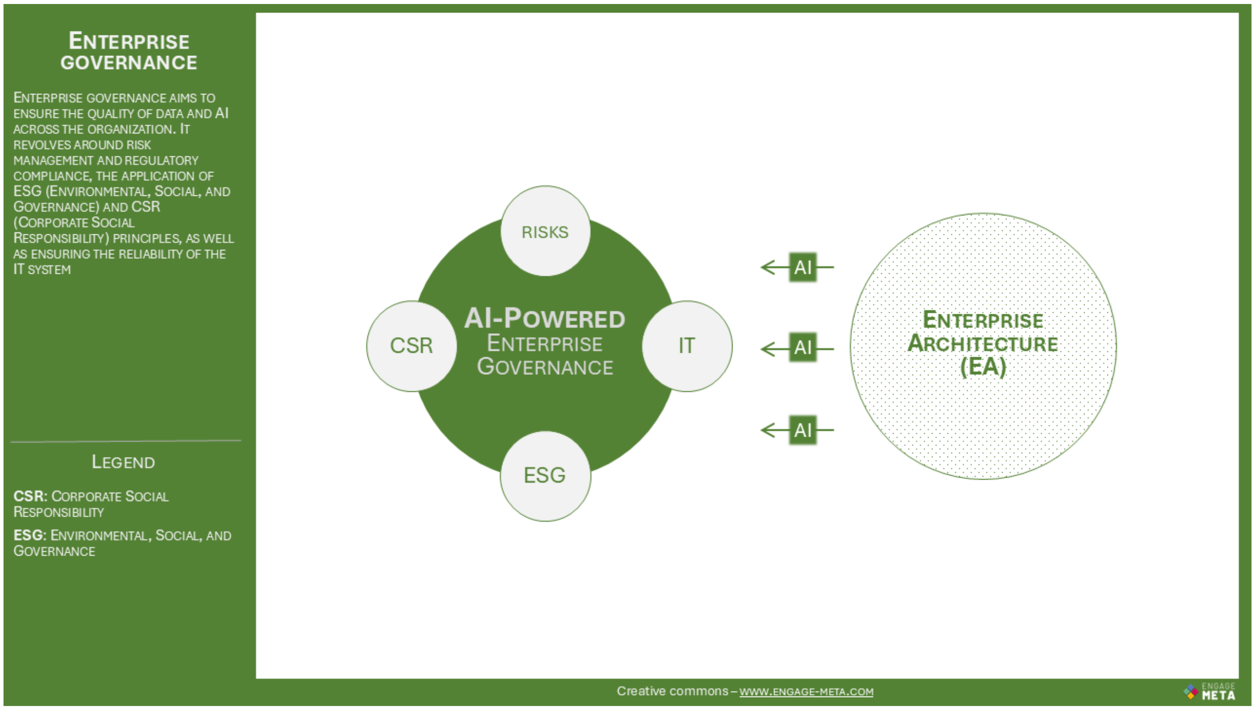

4. YOUR SITUATION & OBJECTIVES

IT DOMAIN OVERVIEW

General introduction to TRAIDA cards in the technical domain. The cards in this domain are universal and apply to all business contexts. You select the practices that correspond to your needs and complete them to manage a roadmap for implementing your minimum architecture to scale AI and data management solutions in your company.

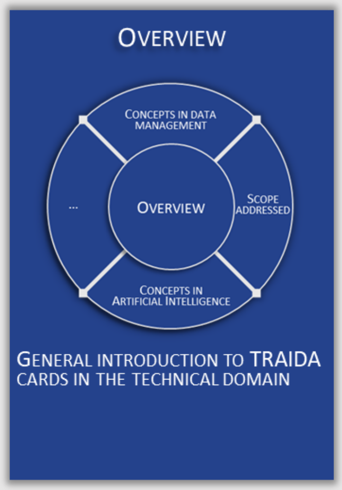

1. CONDITIONS OF SUCCESS

The TRAIDA framework (Transformative AI and Data Solutions) is based on three domains:

1. Technical (blue cards).
2. Governance (green cards).
3. Business (red cards).

To scale AI profitably across the enterprise, these three domains must be aligned. The technical domain is based on a foundational principle that serves as the cornerstone of the entire TRAIDA approach: "***The idea of integrating AI with existing databases is rejected***." The reasons for this recommendation are as follows:

- A strong coupling between AI and the databases of the existing information system creates point-to-point connections that are fragile (difficult to maintain) and poorly auditable (lack of central governance). From a software engineering perspective, this coupling creates technical debt and must be replaced by loose coupling. This allows AI systems to be independent of the physical access layers to production databases.
- A new data repository is necessary to store the tacit knowledge required for AI training. This type of knowledge, also known as informal knowledge, exists in the minds of human actors and is increasingly necessary to enhance AI's capabilities. This new repository is disconnected from production databases and aligns with the objective of loose coupling.

To ensure this separation of concerns between AI and production systems, a semantic platform is implemented. It relies on three repositories that create a digital twin of the existing databases (see the respective TRAIDA cards): MDM (Master Data Management), ODS (Operational Data Store), EKG (Enterprise Knowledge Graph). The semantic platform also integrates processes for data quality control and integration with production systems.

AI systems can then draw training data from this digital twin. The repositories are modeled using ontologies shared at the global enterprise level to ensure a unified view of the data.

Success criteria for AI

To successfully scale AI, TRAIDA highlights the following points:

1. Have a semantic platform with MDM, ODS, and EKG repositories.

2. Clarify the strategic contributions of AI targeted by the company. TRAIDA identifies two universal contributions to consider as a foundation for drafting your strategic approach with AI (see the TRAIDA Artificial Intelligence card):

   a. Process automation.

   b. Knowledge accumulation.

3. Define a progressive roadmap for the implementation of the semantic platform, considering a minimum viable architecture to scale AI. You will find practices in TRAIDA that correspond to your needs. You will need to adapt them to build your roadmap for the implementation of your digital twin. The goal is to establish a minimum architecture within a timeframe that limits AI deployments outside the architectural framework. Without this, it is likely that heterogeneous implementations will lead to AI malfunctions. This situation would negatively impact user adoption and the motivation of decision-makers to support AI investment.

Coordination with the TRAIDA governance domain

The priority coordination points between the technical domain and the governance domain are as follows: Enterprise Architecture card: semantic modeling; Enterprise Governance card: data governance and AI governance. The other cards and topics in the governance domain should be used according to your needs.

Coordination with the TRAIDA business domain

When implementing the first version of your minimum architecture to scale AI, the strategic framing of AI around the two universal contributions proposed by TRAIDA may suffice (see the technical card on Artificial Intelligence). However, you can further detail them with the following cards from the business domain:

- **Productivity card**: all topics on this card should be studied from the perspective of process automation (the first universal contribution of AI), including: Internal process, Client process, Third party process and Compliance process.
- **Creativity card**: all topics on this card can also be analyzed from the perspective of knowledge accumulation to enhance AI capabilities and foster collaboration with humans and robots (the second universal contribution of AI). The topics are identical to those on the previous card regarding productivity.

The other two cards in the business domain, namely Trustworthiness and Treasury & Assurance, are more suitable for use when studying the cards in the governance domain.

Coordination during the implementation of a project with AI

Once the minimum viable architecture to scale AI is established, the coordination of TRAIDA cards revolves around the successive deployment projects for AI use cases, which involves:

1. Updating the architecture according to a predetermined roadmap and considering the needs of business projects as AI and data management solutions are deployed within the company.
2. Ensuring the alignment of the technical and governance domains with business needs. Each project is then analyzed using a set of topics from the business cards (red), a set of governance topics involved in business implementation (green), and the technologies involved (blue).

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

This card is an introduction to the technical domain of the TRAIDA framework. It helps you become familiar with the other cards in this domain. The following provides some additional information to facilitate your reading and the necessary reflection for your own context.

Concepts in Data Management

If you are a newcomer in the field of data management, it is advisable to read the glossary card of the TRAIDA framework, located in the governance domain (green). The concepts that require a deeper level of expertise in data management are as follows:

- **Digital twin or semantic platform**: This involves creating a data repository that cleanly unifies all existing databases within applications. It does not replace them but acts as a clearinghouse for all data. As mentioned earlier, AI systems are then connected to this digital twin rather than directly to heterogeneous databases, which are often organized in silos. However, even in the case of a more streamlined data architecture, such as one based on a data mesh approach (service-oriented architecture), it remains important to establish the digital twin to decouple AI from production systems. This does not mean that AI is not utilized in operational processes running in production, but rather that the data it uses for training and execution comes from the digital twin, which ensures its quality and security.
- **Ontology**: This refers to the embodiment of a semantic data model that is conducted at the enterprise level to overcome imperfections generated by silos. With databases verticalized on application and organizational domains, duplications of information are inevitable, and sometimes semantic ambiguities arise, reducing the reliability of consolidations. The repositories that form the digital twin, namely MDM, ODS, EKG (see the respective TRAIDA technical cards), share the same unified data model, i.e., the same ontologies. Such a data model requires a specific, phased modeling process led by experienced experts. It is a significant investment; the return of which is the profitable deployment of AI at the enterprise scale.

Concepts in Artificial Intelligence

The internal functioning of AI relies on complex and rapidly evolving technologies. For a company that uses AI systems, it is not necessary to have expertise in mastering them. The concepts to understand and disseminate within your teams are as follows:

- Distinguish between different types of AI, such as generative, symbolic, and analytical (see the respective TRAIDA technical cards).
- Understand the paradigm shift from conventional software development, which is moving from coding algorithms to training AI. The contribution of NoCode solutions combined with AI further accelerates this movement.
- The corollary of this new paradigm on skill and career management is significant. In particular, developers and analysts will no longer have a monopoly on specifying databases and use cases. Users, provided they are trained to document their needs in a formalism compatible with AI training, will become *de facto* super-analysts and application producers.
- The approximate results of AI in certain situations should not be perceived as definitive malfunctions for two reasons:
    a) AI collaborates with human actors to improve results. AI use cases are constructed with the possibility of human intervention depending on the execution context: see the TRAIDA technical card on Artificial Intelligence with the principles of preconditions and postconditions for AI-driven steps. This working method requires that users involved in the processes maintain a critical mindset toward AI-produced results, with the ability to propose the addition of knowledge to improve AI training.
    b) The level of hallucination generated by AI is regulated by settings. Depending on the desired level of creativity in executing a use case, these settings should be carefully adjusted. For example, in generating text for a marketing context, the level of hallucination (and therefore creativity) can be high to produce original results. Conversely, in the context of legal analysis of a contract, the level of hallucination should be minimized, and AI training should be reinforced to prevent the creation of new ideas.

SCOPE ADDRESSED

The cards in the technical domain are listed in the table below. There is no preferred reading order to follow. From an academic perspective, that is, for discovering the cards with the aim of learning general technical culture, the order of the cards in the table is the most advisable to follow.

IT DOMAIN OVERVIEW

GENERAL INTRODUCTION TO TRAIDA CARDS IN THE TECHNICAL DOMAIN. THE CARDS IN THIS DOMAIN ARE UNIVERSAL AND APPLY TO ALL BUSINESS CONTEXTS. YOU SELECT THE PRACTICES THAT CORRESPOND TO YOUR NEEDS AND COMPLETE THEM TO MANAGE A ROADMAP FOR IMPLEMENTING YOUR MINIMUM ARCHITECTURE TO SCALE AI AND DATA MANAGEMENT SOLUTIONS IN YOUR COMPANY

CORE SYSTEM DATA
- ❑ SILOED DATABASES
- ❑ DATABASES BY BUSINESS DOMAIN (DATA MESH)

OPERATIONAL DATA STORE (ODS)
- ❑ READ-ONLY MODE
- ❑ WRITE-MODE
- ❑ ANALYTIC-MODE WITH KNOWLEDGE GRAPH

MASTER DATA MANAGEMENT (MDM)
- ❑ DATA CATALOG (METADATA) & GV. FEATURES
- ❑ REFERENCE & MASTER DATA
- ❑ ID MAPPING & DATA LINEAGE
- ❑ API MANAGEMENT

ENTERPRISE KNOWLEDGE GRAPH (EKG)
- ❑ ONTOLOGY MANAGEMENT (ANALYTICS, OLTP)
- ❑ REGULATORY MANAGEMENT
- ❑ PERSONAL & COLLECTIVE KNOWLEDGE ACC.

DATA LAKE WAREHOUSE
- ❑ DATA WAREHOUSE, DATA LAKE & METADATA MGT.
- ❑ DURABLE AND LONG-TERM STORAGE

DATA INTEGRATION
- ❑ DATA HUB
- ❑ DATA FABRIC

STYLE OF DATABASE
- ❑ STRICT-SCHEMA
- ❑ META-SCHEMA
- ❑ DOCUMENT SCHEMA
- ❑ GRAPH SCHEMA & SCHEMA-FREE
- ❑ VECTOR DATABASE
- ❑ FULL-TEXT SEARCH DATABASE
- ❑ ANALYTICAL DATABASE

ARTIFICIAL INTELLIGENCE (AI)
- ❑ GENERATIVE AI & RAG PROCESS
- ❑ SYMBOLIC AI
- ❑ ANALYTICAL AI
- ❑ DATA COLLECTION & LABELING

3. BLUEPRINT

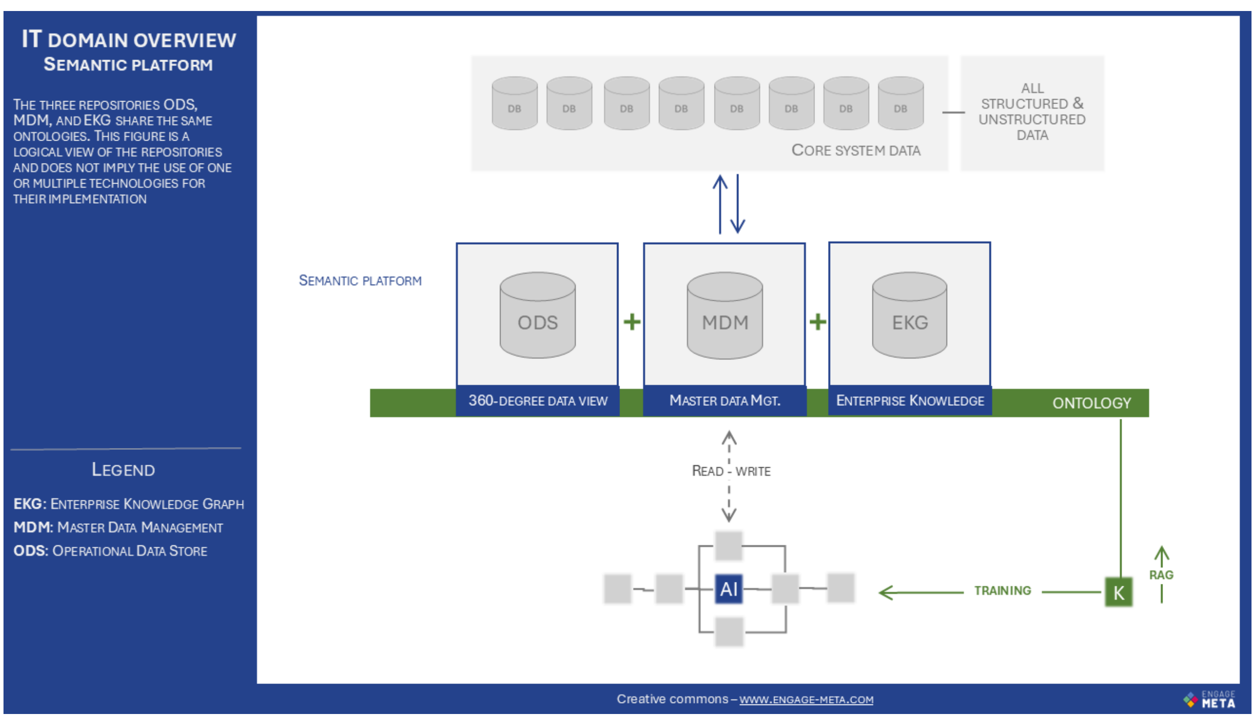

4. YOUR SITUATION & OBJECTIVES

CORE SYSTEM DATA

Core system data consists of your structured and transactional data that contribute to the execution of operational processes, as well as links to unstructured and multimedia data structures. These data elements have predetermined usage objectives. This does not refer to decision-making system data (business intelligence, data analytics…). Core system data relies on OLTP technologies capable of handling high-frequency multi-user and multi-system concurrent access.

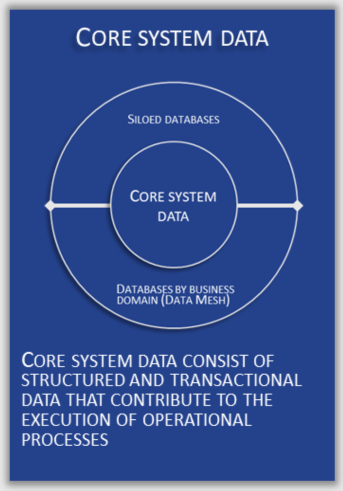

1. CONDITIONS OF SUCCESS

Implement a metadata catalog.

If you do not have unified and up-to-date knowledge of your core-system data structures such as dataset names, table names, field names, relation names, you need to build or strengthen your metadata repository while avoiding extensive semantic modeling that could be lengthy and costly. This repository isn't meant to handle the data values but to help you understand the metadata managed in your core-system databases. It facilitates the creation of a business terms glossary that must be synchronized across all your operational systems.

This is a sort of data catalog, but it is limited to the work of capitalizing on the knowledge applied to core system databases. It does not replace a complete data catalog repository, which is usually managed through a Master Data Management (MDM) (see the related card).

The knowledge accumulated within the metadata repository highlighted in this card is essential to support and enhance your efforts in semantic modeling. This will provide the initial versions of the ontologies needed to increasingly scale your AI systems. More broadly, it will help you regain control of your data quality.

To achieve this goal, utilize graph-oriented database technology, which offers a schema-free approach for loading existing core system data along with their documentation and automatically computes an initial version of your metadata portfolio. This computation is driven by a generative AI (LLM) at the entry-point of the data injection. By combining agile graph technology with generative AI, you will quickly enhance your understanding of core-system data structures. You will apply a prompt similar to this one:

“Develop an ontology from the provided data repository, utilizing the initial list of business concepts, which you may further enrich. Ensure the removal of any duplicate concepts and clearly articulate the relationships between business concepts and existing elements, including applications, tables, fields, and relationships. The ontology should document all metadata, such as application names, table names, field names, and relationship names, to form a comprehensive knowledge graph”.

The result of this prompt is then used to generate the graph. AI prompting must be guided by your business terms glossary to create triples from every metadata item stemming from your core system to your official business terms. A triple consists of (1) a unified business concept, (2) a relation (linked to), and (3) an existing concept in your Information System, such as application and dataset names, table names, and field names.

The AI approach will enable a faster engineering process and avoid cumbersome modeling procedures that are inefficient due to rapidly changing data structures and complexity in your operational systems.

This knowledge enables you to determine actions to correct defects in existing silos, clean up data (meaning, completeness, accuracy, deduplication), optimize your data flow integration APIs (pivot format enhancement) and align better with regulations.

Additionally, you will accumulate the necessary knowledge to progressively redesign your siloed databases to organize them by business concepts (e.g. data mesh) and set up operational repositories such as your Operational Data Store (ODS, storage of data values) and Master Data Management (MDM, metadata repository at scale with data governance features applied to master and reference data).

Unstructured data

Your core system data have links to unstructured data sources, such as document management repositories or big data. The metadata catalog presented here should also take this into account. These are specific types of metadata that are considered "unstructured." Since this is not about storing the value of the data, the issue of storing these multimedia contents is not addressed here.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The data from your core-system are necessary to train your AIs.

Static data (which does not change or changes infrequently) is used during mass training sessions (batch) and the more volatile data is used to refine queries (prompt injections, RAG technology, etc.) in real time (on the fly). To facilitate the implementation of unified data repositories necessary for AI systems, such as Operational Data Store (ODS), Master Data Management (MDM), or Enterprise Knowledge Graph (EKG), you must have good documentation of your core-system data structures and meanings. To achieve this goal, it is beneficial to set up a metadata catalog that describes them in synchronization with the MDM system.

EXISTING SILOED DB

Depending on the quality of your existing core system databases, you will be more or less well-prepared for the deployment of an enterprise-wide data governance policy with the right foundation for your AI systems.

If you encounter obstacles when deploying new transactional data and when developing processes that extend across multiple databases, then consider the success conditions described in this card. It is unlikely that you can significantly improve your situation through a single data governance initiative or data mesh project alone. Indeed, having up-to-date and detailed knowledge of your data structures (metadata) is a prerequisite for any improvement actions.

As indicated in this card, you need to build this knowledge quickly, cost-effectively, and so that it can be easily updated as your core-system evolve. This objective is addressed by the use of graph-oriented database technology in schema-free mode coupled with generative AI.

BUSINESS DOMAIN DB WITH DATA MESH

Data mesh is a data architecture approach. Its objective is to organize databases around business concepts, as opposed to organizational and functional silos. It's an "object-oriented" approach at a systemic data scale.

With a database for each business concept, the issues of data duplication across silos are eliminated. Operational processes then source their data through standardized access in each database of the involved business concepts.

To implement this approach, you must act cautiously to avoid malfunctions at these levels:

- management of links between data located in different databases in terms of integrity, transaction management, and data flows,
- normalization of metadata that must be common to all databases,
- service level agreement on response times and concurrent accesses,
- unified data preservation in histories and archives,
- engineering of software development that exploits these new databases and security.

To succeed in a data mesh program, first, ensure sufficient and up-to-date knowledge of the existing core-system databases structures (metadata), which we have already discussed in the success conditions of this card.

3. BLUEPRINT

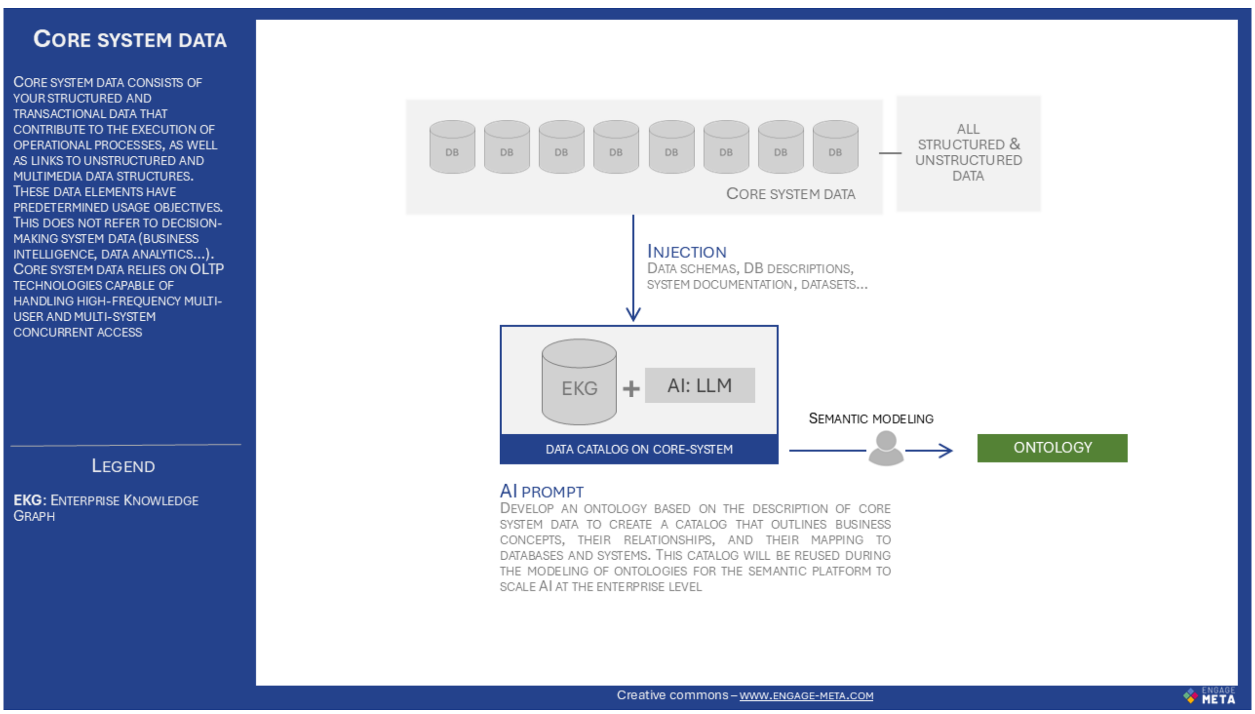

4. YOUR SITUATION & OBJECTIVES

OPERATIONAL DATA STORE

The Operational Data Store (ODS) is a unified repository that collects all structured data from all databases, providing a 360-degree view. In practice, a read-only ODS can cover just one functional or business domain of the enterprise to build a unified view of data within this limited scope.

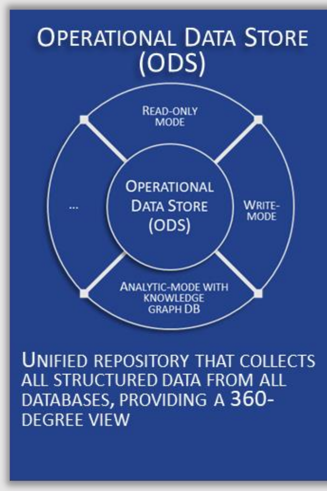

1. CONDITIONS OF SUCCESS

History

Since the beginning, information systems have gradually structured around multiple data sources. These systems generate information quality issues due to duplications and complex relations between objects stored in these different sources.

In the early 1990s, the need for a unified repository to consolidate these sources into a single point emerged. At that time, it was about preparing data downstream from business intelligence repositories like data warehouses. In this context, the term Operational Data Store (ODS) became widespread. It didn't introduce new storage technologies since the use of relational databases was the norm. It was used as a new data source exclusively for consultation in business intelligence. Although its data model needed to be properly constructed, it was not yet a semantic modeling. It was just necessary to ensure an organized structure of data for their use in decision-making systems, in a context where data warehouses presented significant constraints for the volumes of data managed.

A few decades later, the emergence of massive data storage technologies with big data made the use of ODS less useful: why spend money on this repository when it was possible to dump all data sources into big data? Unfortunately, experience showed that the lack of data structuring in big data harms the quality of analyses.

Today, many companies are dissatisfied with their big data projects partly due to the absence of an ODS upstream of decision-making systems. This results in a lack of semantics in big data that prevents leveraging the deep richness of data.

In parallel with the deployment of big data, the ODS survived outside the needs of decision-making systems, under different names and in a manner limited to certain business or functional domains. The most common are CDI (Customer Data Integration), PIM/PLM (Product Information Management / Product Lifecycle Management), and to some extent MDM (Master Data Management).

The return of the ODS

In this context of losing data meaning in decision-making systems, generative AI seems to offer a miraculous solution to regain meaning in data repositories, whether structured or not. Unfortunately, two new problems arise:

1. The use of AI on decision-making data sources (big data) is not sufficient since the company generally wants to leverage operational data in all its extent to train AIs, with the most accurate freshness level and sometimes in real-time for certain use cases.

2. AI needs information about the meaning of data in their usage contexts to reduce biases, analysis errors, or hallucinations. The more AI systems spread into the company's operational processes, the more these data interpretation flaws become unacceptable and can lead to significant value losses.

To counter these two problems, setting up an operational data repository upstream of AI systems usage becomes a necessity. In other words, AIs draw their data from this repository, whose quality, depth of details, and semantics are sufficient to build systems that make AIs more reliable in all usage contexts. Thus, it marks the strong return of the ODS.

Implementing the ODS

Once the interest in the ODS is understood, choosing its implementation can be delicate due to existing IT systems that already use repositories like CDI, PIM/PLM, or MDM. Some characteristics of the Operational Data Store might appear redundant with these repositories. These architectural decisions depend on each company's context, but you should follow these initial principles:

- The construction of the ODS must aim for coverage of the entire enterprise scope, meaning all business concepts like products, customers, organizations, manufacturing processes, etc. It's often necessary to start a first version of the ODS on a limited scope corresponding to a business or functional domain. However, its future extension should be planned from its foundations to ensure the establishment of a minimum viable architecture to scale. Since the ODS serves to train and feed your AI systems on all your operational processes, it must cover all the company's business concepts.
- Be careful not to create ODSs in silos by favoring short-term agility for a project over a global technical solution associated with enterprise-level modeling work. Indeed, heterogeneous ODSs encourage data duplications and semantic divergences with associated problems.
- The ODS should rely on the metadata catalog of core system data (see the TRAIDA card that covers this area). This ensures that the ODS's semantics are sufficiently solid even if you have a modeling know-how deficit. The ODS's projection on the entire information system scope requires relevant and robust data modeling.
- Depending on whether the ODS usage is read-only or write-mode, information storage technologies can vary between relational databases and knowledge graph-oriented solutions. We will revisit the issue of unstructured data storage later in this document.

With these fairly simple initial principles, you should study the synchronization rules with existing repositories mentioned earlier. It is not relevant to integrate an ODS into your architecture if you haven't clarified the integration rules with your CDI, PIM/PLM, and MDM.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

If you do not have a good understanding of your data, it is difficult to train AI systems correctly and connect them to your real-time information sources (RAG integration). Therefore, you need to have a data catalog with up-to-date and reliable contents. Two data repositories need to be made available to the AI:

1. Structured Data: Mostly from your operational applications, primarily your backends, which TRAIDA refers to as core system data.
2. Unstructured Data: This includes multimedia data, encompassing various formats such as images, sounds, videos, documentation, emails, etc.

The first repository is handled using a metadata catalog on core system data (see the TRAIDA core system data card) coupled with the implementation of an ODS (the focus of this document). The second repository requires the use of a knowledge management repository, addressed in the form of an Enterprise Knowledge Graph (see the relevant TRAIDA card).

The use of the ODS can be considered in three modes, which we will describe in the following sections.

READ-ONLY MODE

The read-only ODS feeds AI systems and does not accept update flows in return.

More specifically, ODS data is exclusively updated by injection flows from source systems. Applications and AI do not directly update it. The feedback from the ODS to the source systems is limited to data cleaning and quality improvement decisions. ID deduplication and data normalization mechanisms should rely on specific technologies outside the ODS, typically within the scope of MDM (Master Data Management).

Since this ODS operates in read-only mode, using knowledge graph-oriented database technology is pertinent. This approach allows for reusing the metadata catalog related to core system data (see the corresponding TRAIDA card). This "schema-free" system offers great flexibility in implementing the ODS data model. Additionally, as it is not intended to support updates, the level of semantic modeling can remain quite basic. For example, there should be no complex integrity control rules to model.

Given that the repository is read-only, it is feasible to develop multiple ODSs for isolated business or functional domains without the risk of introducing update silos with associated data duplications and quality issues. However, if the future goal is to implement a write-mode ODS, careful consideration is needed regarding the use of this flexibility; isolated ODS repositories would need to be deconstructed and integrated into the write-mode ODS.

Unstructured data can be incorporated into this knowledge graph-oriented repository. Alternatively, integration with a specialized knowledge management repository, itself based on graph database technology, can be considered (see the TRAIDA Enterprise Knowledge Graph card).

Governance processes for data injection flows are necessary to manage versions and data landing zones.

Finally, since this ODS is limited to providing data without direct updates, a simple user interface for consulting data sets by business concepts is generally sufficient. For data analysis needs, see the next topic, "Analytic-mode with knowledge graph DB."

WRITE-MODE

The write-mode ODS follows the recommendations already described for the read-only ODS, but introduces differences that we will now list.

First, the use of knowledge graph technology is no longer as obvious as for the read-only ODS. This ODS accepts data updates directly from application systems and AIs, giving it greater responsibility within the overall information system but also increasing the requirements for data quality control, especially regarding integrity constraints. Additionally, the frequency of transactional update flows may necessitate specific database technologies. The best solution between relational database technology and knowledge graphs must be decided. Although relational databases are often the best choice for repositories with intense updates on structured data (OLTP), two disadvantages should be considered:

1. Since this is not a "schema-free" approach, creating relatively rigid data schemas is required, which comes with stronger modeling quality demands. This point turns into an advantage in the long term by ensuring the durability of data structures and avoiding the trap of poor agile iterations in a "schema-free" environment, which sometimes leads to chaotic solutions.
2. The difficulty of handling unstructured data. The solution here is to integrate directly with a dedicated knowledge management technology for multimedia data (see the TRAIDA Enterprise Knowledge Graph card).

The modeling effort for a write-mode ODS is more significant than for read-mode alone. It requires modeling complete ontologies based on a glossary and taxonomy of business concepts shared across the enterprise. Although these semantic mechanisms are also necessary for Master Data Management (MDM), the ODS does not replace MDM for at least two reasons:

1. Data governance features remain simple with the ODS, as its goal is to provide data to AIs rather than managing the data itself.
2. Unlike MDM, which offers rich governance features for business but whose scope is limited to the most shared data in the enterprise (reference and master data, see TRAIDA card on MDM), the ODS specializes in managing operational data.

This point about MDM encourages studying the integration with write-mode ODS: the ODS specializes in managing operational data, and the MDM governs the most widely shared data in the enterprise. Both repositories then share responsibility for certain integrity rules, which should be adjusted in your context.

It should also be noted that deploying a write-mode ODS across multiple isolated functional or business domains is not possible. This would result in data duplication risks and a reinforced silo effect, detrimental to overall information quality. Therefore, a systemic approach to modeling with a global enterprise scope is required, referring to our previous point on the need for complete ontologies. Remember, the read-mode ODS does not introduce such risks to data quality and can be deployed by contexts.

Governance processes around the write-mode ODS are more complex than those necessary for the read-only ODS. It is no longer just about governing data injection flows into the ODS from source systems, but also the upward flows from the ODS to them and to the AIs for cascading updates. It is essential to precisely address the needs for managing histories, versions, security, archives, traceability of flows, and rollback if necessary. With such needs, considering a data fabric solution might be relevant. A particular point to carefully study is the governance of ontologies. You must consider variants and versions of your data model over time and its synchronization with all integration points regarding data sources and update targets. This is a complex subject requiring appropriate expertise and technologies, not just marketing intentions.

Finally, the user interfaces of the write-mode ODS are richer than those of the read-only ODS. It is no longer just about displaying data sets by business concepts but also allowing their update, including the links that express their relationships in more or less complex data hierarchies according to operational processes, and with the required level of security.

The management of unstructured data follows the same principles already indicated for the read-only ODS. To reiterate, the recommendation is to use a dedicated knowledge management repository (see the TRAIDA Enterprise Knowledge Graph card).

ANALYTIC-MODE WITH KNOWLEDGE GRAPH DB

An analytical ODS aims to analyze operational data, particularly through AI. It is no longer about providing ODS data to AI systems but about applying AI directly on the ODS. Since knowledge graph technologies enable powerful visualization and the use of inference rules, their use can be favored. Additionally, the "schema-free" approach facilitates the implementation of the analytical ODS by reducing the modeling effort required. Knowledge graph technology works well with the read-only ODS that reuses it, but less so with the write-mode ODS when it relies on relational database technology.

Ultimately, the technological choice for implementing the ODS involves considering the combined needs of read-only, write-mode, and analytical ODS. The universal choice of a graph-oriented database is not always possible, especially when the write-mode ODS is highly transactional with intense and complex injection and restitution flows (OLTP). It is then conceivable to implement two ODSs simultaneously and non-competitively: an analytical ODS using a "schema-free" approach for local AI use, and a write-mode ODS on relational database technology for feeding AI systems and synchronizing data updates to application systems.

3. BLUEPRINT

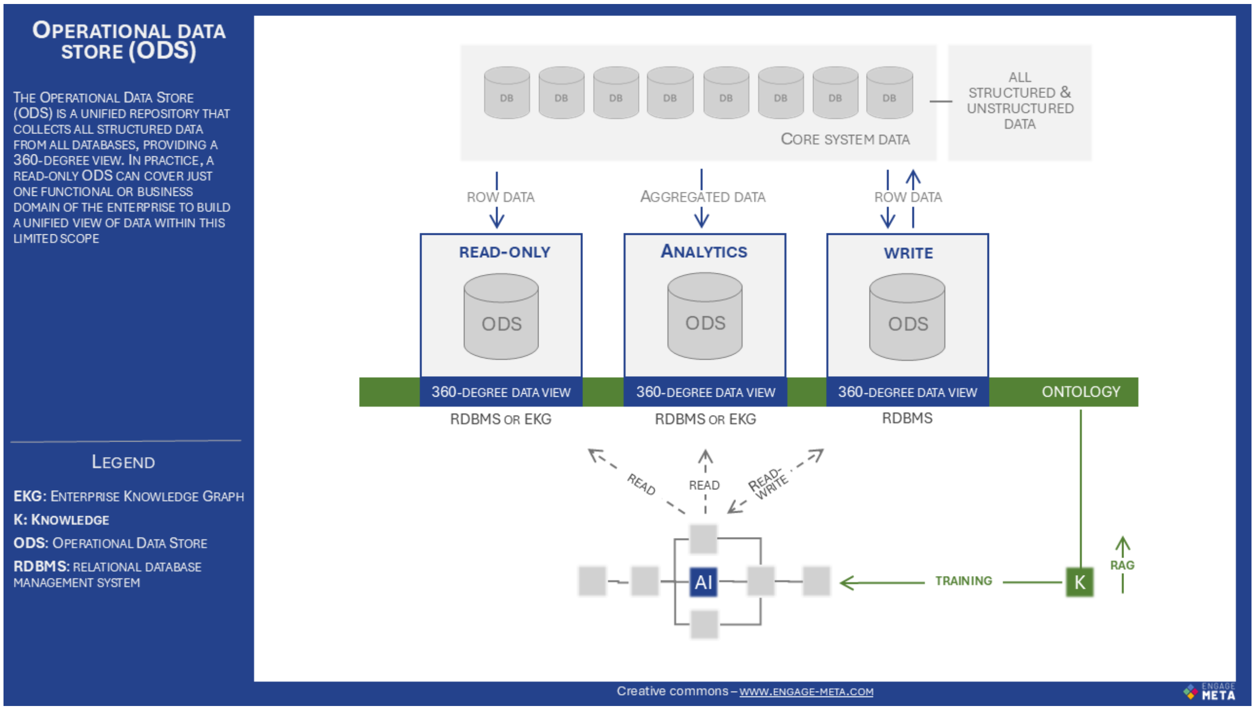

4. YOUR SITUATION & OBJECTIVES

MASTER DATA MANAGEMENT

Master Data Management (MDM) serves as a repository for the most widely shared and structured data across the information system. It is particularly important for AI at scale, as it plays a crucial role in creating ontologies in conjunction with the Operational Data Store (ODS).

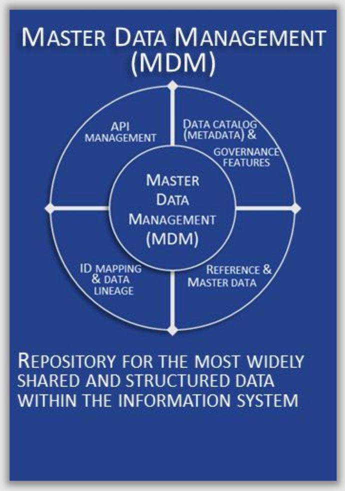

1. CONDITIONS OF SUCCESS

Master Data Management (MDM) offers advanced data governance features such as version and variant management, temporal management (historical), version comparison and merging, data deduplication, data cleaning, data authoring UI, etc. The richer this governance is, the less feasible it is to apply it to data that is frequently and massively (OLTP) modified. Therefore, master and reference data are primarily concerned with MDM.

For instance, the stock of a product in a company's offer catalog evolves in real-time with the flow of orders. However, the physical locations of these stocks in warehouses remain stable over a predetermined period, such as a day, week, or longer. MDM does not manage stock values for each order but handles data concerning their warehouse locations. This is a meta-knowledge applied to the concept of stock. Specifically, MDM manages the metadata of the business concept of "stock" (name, format, nature, application linkage, etc.) without knowing the successive stock values of products. Conversely, for product storage locations, MDM manages both the metadata of associated business concepts (warehouses, geographic location) and the values with warehouse instances and their physical addresses.

The previous example highlights two principles essential for establishing a minimum architecture to scale Artificial Intelligence:

- Metadata is indispensable for describing business concepts used by the company in a unified manner without semantic ambiguities, regardless of their formats, nature, and life cycles: Format: integer, character string, video, sound, multimedia; Nature: operational, decision-making, governance; Life cycle: update frequency.
- The richer the data governance features, the more their usage is limited to long-life cycle data. This mainly concerns the most shared data in the company, namely reference, master, and metadata. This limitation results from technical constraints and the commitment of data management teams (data stewards) whose role is to work on the most shared data within the company. Most of the time, it is the MDM that provides these rich governance features.

In other words, MDM enhances the quality of the most shared data in the information system, which: Carries the core business referential integrity rules; Is used for data consolidation at the reporting level; Is deeply integrated into operational processes.

These data, and thus the underlying business concepts they embody, cannot be managed in silos without risking semantic discrepancies that compromise quality.

The goal is not just to consolidate the most shared data to create a single access point, like an ODS. The objective is also to manage their updates in compliance with globally applicable governance rules within the company. These updates are then reflected in the consuming applications.

To better understand the importance of MDM, here is a list of metadata that passes through its governance:

1. Identifiers of business concepts (sometimes called business objects or strategic data elements) and the relationships between them, describing the mesh between business concepts in the form of taxonomies and a semantic model.
2. The nomenclature (or identity card) of each business concept: name, description, creation date, modification date, and other widely shared data between applications.
3. The life cycle of each business concept. For example, the business concept Customer can follow this life cycle: prospect, new customer, active customer, passive customer, former customer, closed customer. Depending on the state of the business concept, integrity rules are declared to frame updates (governance).
4. The company's business glossary.

This set of metadata embodies a unified data model independent of the specific data structures of existing applications. Therefore, it is in the MDM that the necessary ontologies for anchoring with AI, as recommended in the TRAIDA framework, are found.

This MDM should serve as a launchpad for creating your ontologies, loading them from your applications, updating them by users in charge of governance (data stewards), and resynchronizing them with your application systems.

Here, you need to develop your roadmap to clarify the collaboration between MDM, ODS, and the metadata catalog on core system data (see respective TRAIDA sheets).

In the scenario where MDM and ODS share the same technological solution, they can then be integrated or merged into the same tool. Here are two possible cases:

- ODS in a knowledge graph database with a no-schema MDM in relational database: integration needs to be planned between the two, and a pivot repository for ontology management should be chosen.
- ODS and MDM in a same knowledge graph database: fusion is possible.

The main obstacle to fusion is the lack of governance functions in the ODS to fulfill the role of an MDM. This governance must be exercised not only at the metadata and data levels but also at the underlying data model level, i.e., ontologies. Deploying ontologies without properly managing their versions over time is not recommended. As applications evolve independently of ODS and MDM repositories, change management is required to ensure proper synchronization of the systems in place.

Finally, attention should be drawn to the opportunity of using knowledge graph technology for implementing MDM. The advantage of this technology is the possibility of automatically obtaining an initial ontology version from data sources, with AI support for better interpreting the flows. This capability to avoid modeling work is an attractive aspect of the "free-schema" brought by graph databases (see also TRAIDA card on core system data). However, the semantic accuracy and long-term solidity of automatically generated ontologies are not the best. Indeed, semantic accuracy does not come from existing data flows since they are altered by accumulated poor designs over time (technical debt) and accompanying semantic ambiguities. Instead of starting too quickly with a "free-schema" approach, it is better to go through specific semantic modeling work and consider automatically created ontologies as drafts of a target version to be built. Therefore, it is uncertain that knowledge graph technology is the best option for MDM, especially since this repository requires flawless OLTP management. The alternative is to use a "no-schema" database technology backed by a relational DBMS that ensures better referential integrity and OLTP management than graph technology.

Why is MDM important for AI?

Ontologies form the semantic backbone from which your AI training and enrichment (RAG) processes must be built. You should not integrate your AI directly with heterogeneous data sources from applications. There would be a significant risk of letting quality defects in the data flows, which would cause errors in AI systems without even being able to trace the source of these defects (biases, hallucinations, bugs). The use of ontologies imposes an effort to clarify and clean your data and promotes decoupling between your data and AI.

For example, consider an AI system that needs access to all knowledge about your customers. Without MDM/ODS systems and ontologies, the AI would have to query a series of heterogeneous applications and databases to find customer data, with risks of semantic ambiguities and poor quality due to technical debts accumulated over time in these systems. Conversely, with MDM/ODS, the AI system directly accesses a single, reliable point that provides, through ontologies, a high-quality data source.

Integrating a data mesh strategy

Data mesh is a data architecture that aims to break down silos to organize databases by business domain (see TRAIDA core system data card). Therefore, there should no longer be data duplication and poor data quality. In practice, the roadmap to transition the entire information system to a data mesh can take several years, providing valuable time for MDM. Moreover, unless the different data mesh databases can offer cross-functional and shared governance functions, a central mechanism for managing this governance will still be necessary. The longevity of MDM is even more evident when different database technologies are used to deploy the data mesh. In this case, it is unlikely that these databases can share governance without an MDM acting as a pivot.

Considering multimedia data

MDM handles structured data and integrates multimedia data through links to specific storage areas, ideally a graph database repository (see TRAIDA card about Enterprise Knowledge Graph), big data, or a cloud storage server like Google Drive.

Complement on ODS and MDM integration

When using MDM as a pivot repository for managing metadata, reference, and master data, it is possible to use complementary data access virtualization technology to retrieve operational data located in the ODS for reading purposes.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

AI systems are initially trained on large volumes of data, where semantic structuring is not fundamental. However, the subsequent fine-tuning processes involve smaller volumes of data that require increased semantic mastery to achieve relevant results. At the most detailed level of this training, it may involve real-time access to specific business concept information located in a database (RAG: Retrieval Augmented Generation). At this access level, it is crucial to have metadata describing this business concept.

Thus, without effective metadata management across the entire information system, it will be impossible to fine-tune AI systems, leading to disappointing results. This is a fundamental reason for the shift from big data, which lacks semantic management, to ODS and MDM systems that rely on powerful semantic management.

If this observation seems relevant to your context, you should establish a roadmap for systemic metadata (ontology) management. It is not just about creating a metadata catalog to understand your existing data (see TRAIDA core system data card) but about building your minimum semantic model to gradually scale your AI strategy. You have two possible approaches:

1. Operational Data Store (see TRAIDA ODS card): This goes beyond metadata management by handling all operational data. However, its lack of governance functions may prevent its use as a pivotal metadata repository. It is crucial to manage the lifecycle of ontologies to avoid failures in scaling due to misalignments between the ODS and applications.
2. Master Data Management: Specialized in managing metadata, reference, and master data, MDM's governance functions are more powerful than those of the ODS, giving it an advantage as the company's central metadata catalog. It must be capable of managing the lifecycle of ontologies to ensure synchronization with applications.

Therefore, MDM is a pivotal element in building the semantic management platform necessary for scaling your AI systems. Depending on your context, you will need to determine your own roadmap to synchronize the ODS, MDM, and Enterprise Knowledge Graph (EKG).

REFERENCE AND MASTER DATA

Reference data includes the codifications used in your applications. Some are standardized, such as country codes, others are industry standards, and some are specific to your context. They often follow a simple structure like code and label.

Master data refers to the identity cards of your business concepts. First, you need to list these concepts to build a catalog and then a glossary. Each identity card consists of the most stable and shared data among applications. For example, for the business concept Customer, the master data would include: identifier, name, surname, address, email, date of first purchase, status (payment in progress, payment OK, pending validation, no longer active, archived...), and relations to other business concepts (Product, Sales, Billing, etc.).

ID MAPPING AND DATA LINEAGE

Each business concept referenced in the MDM is fed by source applications and synchronized with consumer applications, which can be the same or different. Typically, a primary (or master) source application is designated as responsible for the main ID in the ID mapping that needs to be constructed to reference all source and consumer applications. The materialization of this data mapping in the MDM (ID mapping) can take various forms, often involving metadata.

Accumulating the identifier mappings of different business concepts allows us to identify chains in their usage across applications. For example, a Product Management application sends an update of product descriptions to the Sales and Marketing applications. In this case, there is a chaining for the Product business concept that starts with the Product Management application and links to the other two applications, Sales and Marketing. This formalization of knowledge enhances the semantic scope in the MDM and improves governance. For instance, it could be decided that the MDM is responsible for directly propagating product description changes to the target applications.

DATA CATALOG (METADATA) & GOVERNANCE FEATURES

As mentioned earlier and reiterated several times, metadata management is crucial for gradually scaling your AI systems. This is an essential part of establishing a semantic management platform as recommended by TRAIDA. The data catalog, and more specifically the metadata catalog, is better managed in the MDM due to its powerful governance features. Here is a non-exhaustive list of these functions: Data model versions with data migration functions between models; Data spaces and data sets by variants and versions; Data update screens automatically available from data models and customizable as needed; Data hierarchy manipulation; Data update and validation workflows; Security; History; Archiving; And more.

API MANAGEMENT

API management involves documenting and configuring service contracts used for interactions between applications. A service contract is a business concept whose identity card consists of the description of the service's purpose, its input and response parameters, configuration possibilities, etc. Instead of leaving this information solely in technical documentation (e.g., Javadoc), it is beneficial to elevate it as metadata within the MDM to manage their life cycles.

For example, consider an API for retrieving a product sheet described in a Javadoc. This description is brought into the MDM to reference the consumers of this API: the Sales application uses it with a filter that only retrieves pricing data, while the R&D application uses it with a filter that only provides technical data.

3. BLUEPRINT

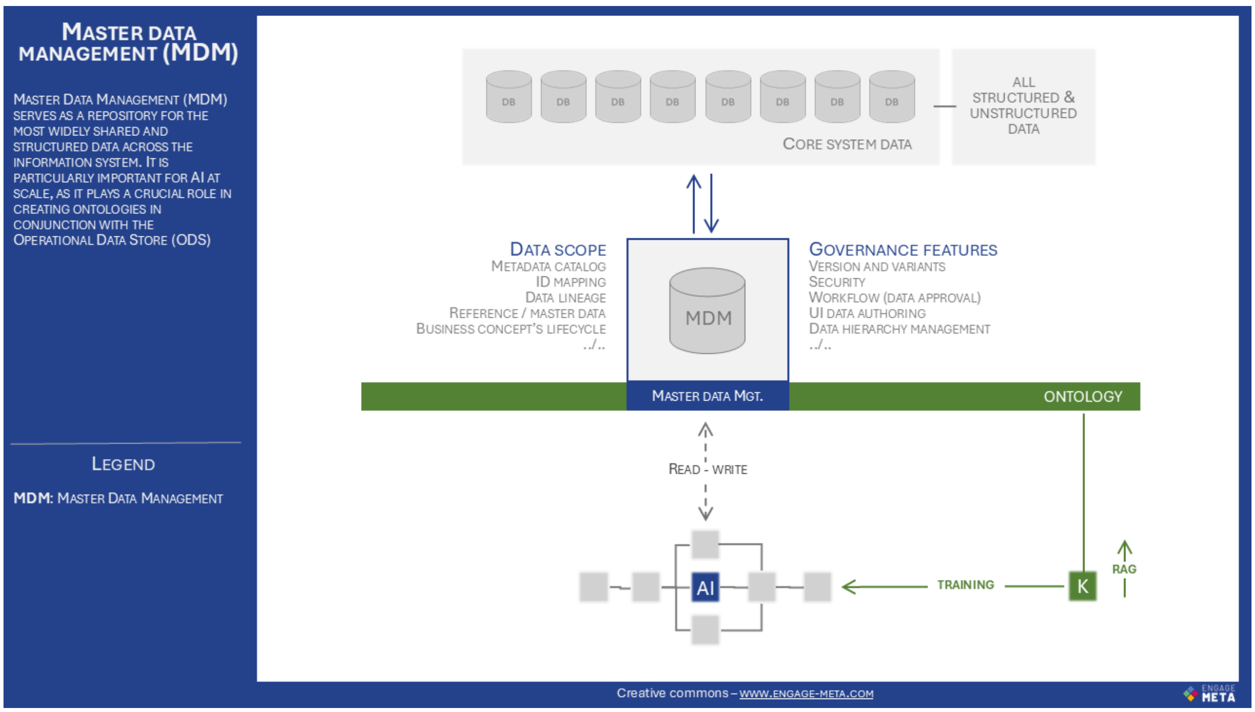

4. YOUR SITUATION & OBJECTIVES

ENTERPRISE KNOWLEDGE GRAPH

The Enterprise Knowledge Graph (EKG) is a potential universal repository for knowledge management with various use cases. It is the cornerstone of the semantic platform promoted by the TRAIDA platform. It serves as the unique point of contact for all AI systems within the company (digital twin).

1. CONDITIONS OF SUCCESS

To properly train AI systems, it is necessary to gather the maximum amount of knowledge according to three levels:

1. Data available on the Internet, paying attention to usage rights. Large AI models like ChatGPT or Llama are trained on these data. As a user of these LLMs, you benefit from the training already done on large amounts of information. However, keep in mind that with an open-source LLM like Llama (Meta), you will still need to find a solution to run it on sufficiently powerful infrastructure, likely in the cloud.
2. Your company's data that already exists in your databases, office files, and physical documents (paper). This data is essential to enhance the training of LLMs in order to personalize their behaviors to your company. This is a fine-tuning task.
3. Your company's data known by your teams but not yet formalized in databases, files, or even in writing. This wealth of data is a reservoir of tacit knowledge that represents a significant percentage of the total knowledge the company possesses, around 60% to 80%. This includes the know-how of operators, how they adapt work procedures to the realities on the ground, information exchanged between actors and stakeholders to meet objectives, etc. This informal knowledge must be transformed into formal knowledge to enrich AI systems and improve their profitability.

This data is of all kinds, both structured and multimedia. It evolves with the company and requires version management. For example, a set of data used to train an AI system in an initial version may become obsolete later and will then need to be removed from the AI system's training. In other words, for each AI system training, it is necessary to keep the sources of data used and ensure that rights and security are respected.

This management is particularly delicate because the structures of the collected data are very diverse. Indeed, the training scope of AI systems encompasses the entire company. For example, starting from an internet-based LLM like ChatGPT, the company will proceed to a first level of global fine-tuning to its activity before carrying out finer settings for its different activities, such as its marketing, manufacturing, human resources departments, etc. As these activities coordinate through cross-functional processes, other knowledge will enrich AI systems to optimize operations at the boundaries of departments. Thus, it is a bidirectional movement of AI system training that operates from global to local and vice versa.

From this description, a need for data and knowledge governance (multimedia, tacit data) emerges. It must be flexible enough to instantly adapt to new data structures whose prior modeling would be too cumbersome, lengthy, and even impossible in some situations. However, it must also be able to work with metadata to classify, document, organize, manage versions, rights, security, traceability, and more for all the knowledge used to train AI systems.

Finally, a last criterion to consider is the dual mode of knowledge exploitation:

1. First, in asynchronous mode for the massive training of AI systems. Thus, the company trains its AI from the company's data as OpenAI or Meta does from internet data. This training mobilizes the maximum knowledge to personalize large LLMs.
2. Then, in synchronous mode at the time of prompt execution, through real-time enrichment of information injected into AI systems. This principle relies on RAG (Retrieval Augmented Generation) technology. It is no longer about mobilizing the maximum knowledge but the minimum necessary to help the AI better respond in the specific context of a query. For example, it may involve adding updated data from a customer relationship management database to a prompt analyzing a customer file.

Need for a Knowledge Management Repository

To address the needs described earlier, it is pertinent to set up a knowledge management repository that offers the following characteristics:

- Data injection in compliance with a predefined data model (schema-oriented) that relies on ontology modeling.
- Data injection with free loading (schema-free) and the ability to automatically generate the data schema reflecting the source data (automatic ontologies).
- Management of rules for quality and security controls.
- Management of all types of data, both structured and multimedia.
- Version management at the data schema level.
- Version management at the data level.
- Connectors for data transformation during import and export.
- Management and visualization UI for data.

This repository is both schema-oriented and schema-free. An interesting implementation can be found with knowledge graph-oriented database technology, which is the subject of this TRAIDA card. The EKG (Enterprise Knowledge Graph) repository then becomes the pivot of the semantic management platform necessary for the governance and execution of AI systems. It is also identified as the digital twin of the information system, providing a single and secure access point to all AI systems within the company.

How to Integrate the EKG with Other Data Repositories

The EKG repository should be the unique integration point with AI systems, thereby centralizing governance. This enables tracing the versions of data used to train AI systems as well as the data sources requested during real-time prompt enrichment (RAG).

However, upstream of the EKG, it is essential to consider the use of two other repositories addressed in specific TRAIDA cards:

1. Master Data Management (MDM) for reference and master data. It is the natural source for feeding the EKG with this data. It also provides the data models for ontologies. In other words, the MDM is the pivot repository for ontologies. The EKG can have its own ontologies during free data injection (schema-free) but must rely on ontologies from the MDM for schema-controlled data injections (schema-oriented).
2. The Operational Data Store (ODS) also relies on MDM ontologies. It contains operational data that are injected into the EKG as needed, both during massive AI system training processes and for RAG processes.

Data integration must be carefully designed. For example, it starts with an information modification in a production database that triggers an update in the ODS, then propagates it to the EKG to make it available to AI systems. This integration preferably relies on an event-based architecture to avoid tight coupling between subsystems. In other words, the ODS listens to a data injection stream from the application to initiate its update. Similarly, the EKG listens to a data injection stream from the ODS to trigger its own update. If the MDM repository is also involved in propagating reference and master data, the ODS will listen to a data stream from the MDM.

Is it Possible to Merge MDM, ODS, and EKG Repositories?

It is not straightforward to merge the three repositories into one, for the following reasons:

- MDM requires a transaction-oriented database with strong data typing, meaning a technology that relies on a formal data schema. This may involve an internal meta-schema to ensure sufficient flexibility when updating data structures and their relationships. Moreover, MDM is also an application system with specific governance functions for business teams that are not found in ODS and EKG solutions.
- ODS is an operational data repository that unifies structured data from application systems. It requires schema-oriented database technology to guarantee integrity and transactional management for large volumes and high-frequency access (multi-user in parallel).
- EKG is a more flexible repository that benefits from a schema-free approach to absorb multiple types of data beyond structured data. As indicated in this TRAIDA card, it provides a relevant solution for accumulating the knowledge necessary for AI system execution.

Depending on the technological quality of the software products used, the following integration scenarios are possible:

- For a company with data volumes and transactional requirements compatible with the processing power of the MDM repository, an additional ODS is unnecessary. The MDM also serves as the ODS. In this scenario, the company has two repositories: the MDM, which also handles ODS, and the EKG.
- Conversely, a company with high data volume and transactional demands and a simple governance requirement for its reference and master data can use the ODS as an MDM repository. In this scenario, the company has two repositories: the ODS, which also handles MDM, and the EKG.
- If the EKG technology is robust enough to handle the company's structured data volumes and transactional requirements, it is conceivable to use it as the single repository that also addresses governance needs for reference and master data (MDM) and operational data unification from application systems (ODS). As of the writing of this TRAIDA card and to our knowledge, such technology does not exist. Indeed, graph-oriented databases are not well-suited for multi-user transactional management on large volumes (ODS) and lack business governance functions for reference and master data (MDM).

The TRAIDA MDM and ODS cards provide a more detailed description of these two repositories.

For small businesses (start-ups, SMEs), it is also conceivable to implement a NoCode database (such as Knack or Airtable) to handle the three repositories MDM, ODS, and EKG within the same technology. If the company is starting its information system, it is even possible to consider the ODS as the core system data. In other words, operational processes are directly built around the ODS without needing to install application systems like ERP.

How Does a Knowledge Graph Database Work?

The technology behind knowledge graph databases implements a meta-schema for data storage based on the description of triplets: subject, predicate, object. For example, a customer (subject) buys (predicate) a product (object). Therefore, to inject data into the database, it is not necessary to model data structures beforehand. This meta-schema leads to a schema-free mode of operation, which is very agile for data manipulation. Conversely, in the absence of a data model, there are few rules for data quality verification. However, these rules can be added in addition to the triplet meta-schema. Thus, the knowledge graph database reconciles the flexibility of schema-free operation with the power of quality controls configured according to the use cases implemented in the company.

However, this flexibility generally comes with weaknesses in response times with large volumes, particularly during queries that traverse multiple nodes (equivalent to table joins in relational databases), and in cases of massive concurrent and transactional access (difficulty in guaranteeing ACID criteria: Atomicity, Consistency, Isolation, and Durability).

As mentioned earlier in this TRAIDA card, the state-of-the-art in knowledge graph databases offers an advantage for integrating AI systems, but they do not yet have the sufficient capabilities to serve as ODS and MDM.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

Establishing a unified knowledge repository exposed to AI systems is an indispensable condition for mastering AI and managing data at the enterprise level. Without this repository, point-to-point calls between AI systems and core-system databases (applications, software) would be necessary, as well as all other information sources like files and other documentation. The quality control of data used by AI and their traceability would be compromised. To avoid this point-to-point mode, it is essential to build an EKG repository. It forms a central element of the semantic platform adopted by the TRAIDA framework, in conjunction with the MDM and ODS repositories.

This unified repository thus enables the implementation of security rules, traceability, version management, and more, across all the knowledge used by AI systems.

PERSONAL AND COLLECTIVE KNOWLEDGE ACCUMULATION

An important aspect of the successful and profitable deployment of AI in the enterprise is the ability of stakeholders to document their tacit knowledge and elevate it to a collective level. According to most studies, this knowledge represents about 60% to 80% of the information used by the enterprise. Therefore, a comprehensive program must be initiated to support stakeholders in formalizing their know-how in writing. In the Engage-Meta community, this work relies on the WASI process for Write, Analyze, Share, and Innovate (see the community website for more information).

Of course, all these writings must be stored, versioned, quality-controlled, shared, and used to train AI systems. Therefore, a knowledge storage repository is necessary, and naturally, the EKG stands out as the solution.

This knowledge is stored according to two classifications: either in an ontology specifically created for the classification of knowledge, for example, following the main functions of the company such as marketing, sales, production, etc., or in the business ontologies already in place to operate the organization's processes. In this case, it involves attaching the knowledge directly to the business concepts that form the ontologies.

This new and decisive approach to formalizing tacit knowledge is an ongoing, daily activity that the organization must adopt. Strategies for classifying new knowledge in the EKG may vary depending on use cases and cultural practices within the company, such as in terms of information sharing. However, in all cases, the goal is to have the best possible and up-to-date knowledge to train AI systems and thereby increase their intelligence and value-creation potential.

ONTOLOGY MANAGEMENT (ANALYTICS, OLTP)

The basic principle to remember is the unification of ontologies across the three repositories: MDM, ODS, and EKG. In other words, the list of business concepts, their hierarchies and relationships, as well as their life cycles (business states) are shared by the three repositories. As mentioned earlier, it will be necessary to decide on the pivot repository for the ontologies, the one capable of managing their versions. This is generally the MDM. The EKG, which powers these ontologies shared within the company, resembles the OLTP usage mode.

Next, given the flexibility of the EKG, it can be used to generate tactical or temporary data analysis ontologies. Indeed, when dealing with a data set whose underlying data structure is unknown or partially compatible with the shared ontologies, it is useful to load them in schema-free mode into the knowledge graph database. This allows for data discovery by navigating through the information triplets, and even automatically generating an ontology based on the injected data. This process helps in better understanding the data and even calculating the differences between the default ontology obtained and those officially shared within the company. The EKG that powers these automatic and non-shared ontologies within the company resembles the Analytics usage mode.

How to Build Shared Ontologies?

In the TRAIDA core-system data card, we describe the process of creating a draft of shared ontologies through AI analysis of existing databases. Further construction work is then initiated to build and maintain the shared ontologies within the company, taking into account the needs for optimization, automation, and value creation. This work requires expertise in semantic data modeling, which can be provided by AI systems specialized in this field. In the TRAIDA Initial Engagement offer, each step of data and process modeling is accompanied by an AI assistant under ChatGPT, helping teams produce the ontologies. Here is an excerpt from the work process, with the full description available on the Engage-Meta community website:

- Process
- Business Concepts
- Business Concepts Control
- Ontology
- Data Modeling
- Identifiers Design
- Business Concepts States

- Process Modeling Refinement
- Integration
- Database Implementation
- Process Implementation
- Security Policies
- Governance
- Review

REGULATORY MANAGEMENT

The EKG is very powerful for implementing regulatory compliance monitoring based on these basic principles:

1. A regulation describes rules that the company must apply to different business concepts and processes.
2. The company implements these rules and references them in documentation.

When regulations are extensive, they can include dozens of rules whose implementations in the company's applications can result in hundreds of impact points. Beyond the initial system compliance, the company must find a way to monitor changes in both the regulations and the application systems, whose updates may render the impact points obsolete.

To establish this governance, the EKG repository is used as follows:

- The regulatory text is analyzed by an AI (LLM) to inject it in the form of triplets, for example (Rule, Objective, Business Concept), into the repository.
- A regulatory manager verifies and enhances the relevance of the triplets and complements them with references to the impact points in the company's application systems and processes.

Once this repository is in place, it is used for regulatory monitoring at two levels:

1. The documentation of the application systems and processes, as well as their IT descriptions (source code, data flows, etc.), are processed by an AI (LLM) and injected into the EKG repository to calculate discrepancies between the intended impact points and the reality of the systems. This is an audit operation to ensure regulatory implementation compliance.
2. When a regulatory update arrives, a new repository is built using the same process as described previously. An analysis of the discrepancies between the two versions of the regulations then allows for automatically identifying the impact points to be considered.

Without the power and flexibility of knowledge graph databases and LLMs, it would not be possible to automate the governance described above. Only manual management would be feasible, with the associated costs and risks of non-compliance errors.

3. BLUEPRINT

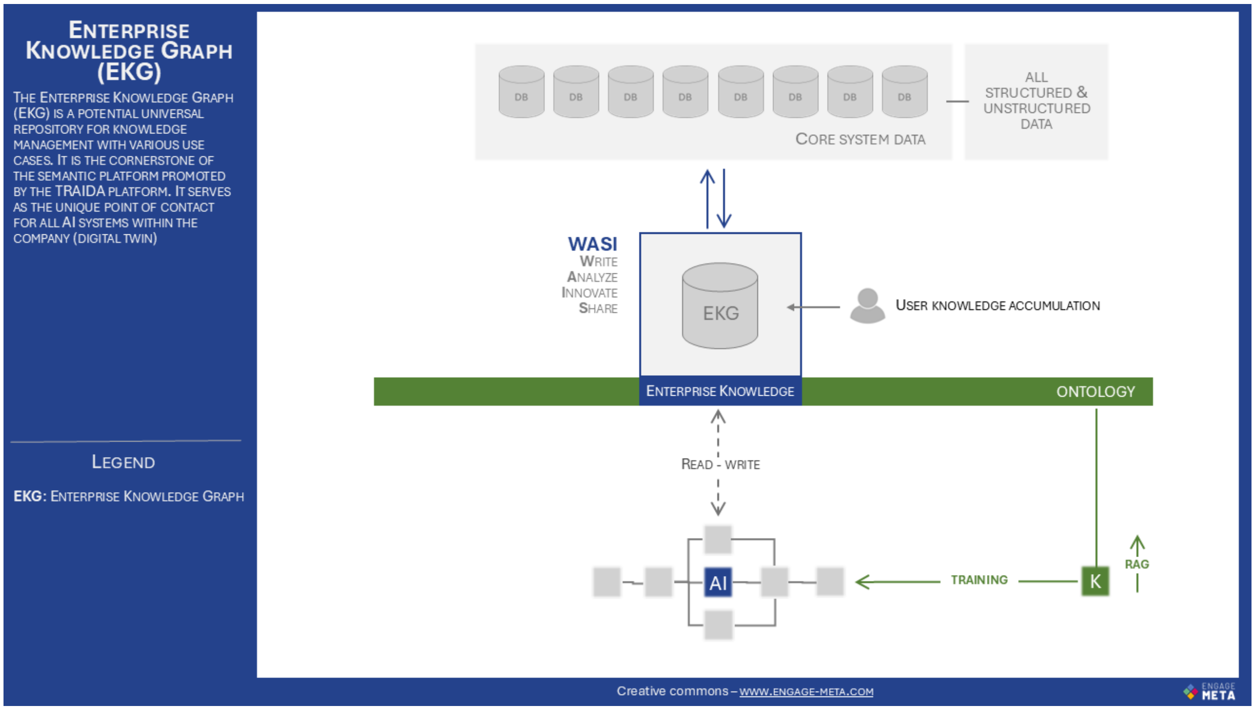

4. YOUR SITUATION & OBJECTIVES

DATA LAKE WAREHOUSE

Repositories contain raw, structured, and unstructured data for business intelligence and data analytics purposes. In TRAIDA, the term 'Data lake warehouse' encompasses data warehouse, data lake, and data lakehouse. The term 'Business intelligence' includes data reporting and OLAP. The term 'data analytics' refers to data science.

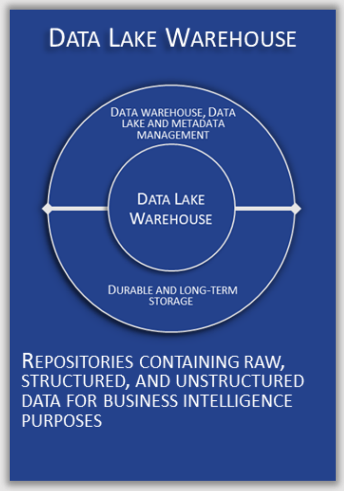

1. CONDITIONS OF SUCCESS

When "big data" solutions do not fully meet expectations, most decision-makers believe that AI and knowledge graphs are the solution to better address data analysis needs. However, successfully integrating transformative AI at the decision-making system level requires clarifying the architecture. With TRAIDA, the effort made at the semantic platform level and with shared ontologies facilitates this integration. We will explain how in this TRAIDA card, but first, we need to clarify the meaning of the term "big data" by reducing it to the identification of multimedia databases. Since this term does not impose specific technologies or use cases, it becomes a commodity that is not structurally important for architectural choices.

We need to move beyond the term big data and return to the company's objectives in these two classic realms of decision-making IT, which we group under the generic term "Data Lake Warehouse":

- **Business Intelligence**: Focuses on reporting needs and structured data analysis. These data are described using metadata that provide their structures, definitions, and quality control rules. The technologies used are SQL-type databases and OLAP (Online Analytical Processing), including meta-schema and NoCode approaches. They are grouped under the generic term data warehouse.
- **Data Analytics**: Refers to the domain of data science, which works on more or less extensive multimedia data sets, with or without metadata. The goal is trend calculation, data discovery, detection of atypical cases, general classification, etc. The technologies used are NoSQL and schema-free. They are grouped under the generic term data lake.

AI's power is expressed in each of these two realms separately. However, it brings more potential when applied to a data repository that unifies the data warehouse and the data lake. This is the promise of new data lakehouse solutions. At the time of writing this TRAIDA card, the feedback from such solutions is still recent, making it difficult to assess their maturity. Nevertheless, it is certain that the convergence of data warehouse and data lake will be realized through such mechanisms:

- The ability to extend OLAP technologies to include multimedia data.
- Adding metadata management in the data lake to enhance query power and quality controls. These metadata must be shared with the OLAP part of the unified solution.
- Standardizing mass data storage solutions for both structured (enriched with their OLAP dimensions) and unstructured (multimedia) data inherent to the data lake.
- Unifying data manipulation languages between the data warehouse and the data lake necessary for injections, cleaning, aggregations, etc.

- Sharing a universal data access layer (OLAP, SQL, data lake) usable by data visualization tools.
- The ability to export vectorized databases from all data sources, both OLAP and data lake. This vectorization is necessary for enriching AI conversations with RAG (Retrieval Augmented Generation) technique and training AI assistants.
- Advances in transaction management (ACID) for both the data warehouse and the data lake will allow the implementation of data update processes directly at the decision-making IT architecture level. It will be possible to build integrated business intelligence and data analytics solutions in operational application systems. For example, a data lakehouse could modify a customer's data as part of an end-to-end process with a CRM application. Transactional management will then encompass both decision-making and operational systems.

To prepare for this kind of evolution, a good practice is to build the necessary ontologies for your MDM, ODS, and EKG repositories (see the respective TRAIDA cards). Indeed, it is from these shared ontologies that metadata are developed and implemented. They bring to life the semantic platform highlighted by the TRAIDA framework. This metadata is necessary to enhance the power of your data lake and to configure the OLAP dimensions in your data warehouse. They will be used for the unification of the two solutions through the evolving data lakehouse technology on the market. All of these MDM, ODS, and particularly EKG repositories are also the necessary data sources to train your AIs and enrich them on demand (RAG).

In parallel with this technological landscape, it is also possible to use the EKG repository as a data analysis solution (see TRAIDA EKG card). Indeed, the technology of knowledge graph-oriented databases offers specific benefits that OLAP and data lake solutions do not:

- The OLAP approach stores data to aggregate them according to multiple axes of analysis.
- The data lake stores data in a raw manner without imposing strong structuring.
- The knowledge graph stores data in the form of a meta-structure in triplets (starting object, relationship, ending object).

Therefore, the EKG is a complementary opportunity for value creation in the field of decision-making IT. It is unlikely that graph technology will replace OLAP and data lake solutions in a reasonable timeframe. However, it can handle certain data analysis cases that avoid the use of OLAP or data lake, and offer others that are impossible to implement without graph management, such as inference algorithms.

Finally, a powerful way to combine the EKG with OLAP and the data lake is to consider knowledge graphs as a layer above the data lakehouse to manage metadata and enrich data analysis systems. In choosing a technical solution for the data lakehouse, it is important to understand the mode of mass data storage (OLAP, SQL, multimedia) but also to evaluate the availability of a knowledge graph-oriented database used as a supervisory layer. If the barycenter of the IT system is the data lakehouse, EKG technology can then come with the data analysis solution. However, as mentioned earlier, it is still too early to validate the technical maturity of such an apparatus. A more reasonable approach is to choose EKG technology that does not depend on the future choice of a data lakehouse. If the latter has a semantic layer based on a knowledge graph, it will need to be integrated with your EKG.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The TRAIDA "Data Lake Warehouse" card is not necessary for successfully deploying AI and its associated data at scale. Indeed, the semantic platform recommended by TRAIDA relies on the MDM, ODS, and EKG repositories described in specific cards.

However, companies need to conduct data analyses, and the repositories of the semantic platform are insufficient when it comes to multidimensional analysis (OLAP) or the exploitation of large amounts of minimally or unstructured information (data lake).

The contributions of TRAIDA to decision-making IT are at two levels:

1. First, the ontologies modeled in the semantic platform can be reused to improve data warehouse and data lake solutions. These ontologies provide the metadata necessary for enriching analyses and enhancing data quality.
2. Second, the MDM, ODS, and EKG repositories are key data sources for the data warehouse and data lake.

DATA WAREHOUSE, DATA LAKE AND METADATA MANAGEMENT

The following architectural principles are proposed:

- The ODS is the preferred repository for feeding data warehouses with operational data. The MDM is the source of reference and master data, which are used to build data hierarchies and analysis dimensions (OLAP).
- The EKG is the preferred repository for feeding data lakes with already accumulated knowledge for AI needs. Additional sources of multimedia data can be added to complement the content from the EKG.
- In all cases, the shared ontologies at the semantic platform level serve as a reference for the metadata handled in data warehouses and data lakes.
- In the case of a data lakehouse equipped with a semantic layer based on a knowledge graph, integration with the EKG is even more natural. The knowledge graph at the EKG level should then be considered the master repository.

According to needs, AI can be used at all levels of the data warehouse and data lake. For example, an AI can be trained at the EKG level to be used on data from the data lake, or directly trained at the data lake level without attempting to reuse it at the EKG level. It is not possible to define generic governance rules that would apply to all companies. It is preferable to adopt a pragmatic approach and frame the training and use of AIs according to their operational or decision-making scope. In other words, a distinction should be made between AIs that act on operational systems (MDM, ODS, and EKG) and AIs that operate at the data analysis level (data warehouse, data lake).

DURABLE AND LONG-TERM STORAGE

In the field of repositories for data analysis, the volumes handled are generally large. It is therefore necessary to specify the means used to ensure their storage by distinguishing between two conservation horizons:

- **Durable Storage**: This corresponds to the ability to provide data on demand, including large volumes, with redundancy mechanisms, multiple backups, real-time access security guarantees, and durability of access APIs. The physical storage media (disks, memories) must be able to change with technological advancements without altering the access mechanisms. Some well-known solutions include: Amazon S3, Google Cloud Storage, Azure Blob Storage.
- **Long-term Storage**: This corresponds to the archiving and recycling of data over long time horizons (several years, decades) without allowing real-time access to the data. When an archive needs to be loaded, a process lasting several hours or days is initiated in accordance with a service contract. The physical storage media must be maintained or even transferred from an old technology to a more recent one transparently for the archive user. Over a long period of several years, it is essential to ensure that the physical storage media does not degrade and remains readable with the most recent technologies. Therefore, regular maintenance is required. Some well-known solutions include: Amazon Glacier, Google Cloud Archive Storage, Azure Archive Storage.

AI solutions contribute to better management of these storage systems by monitoring access to detect potential fraud, anticipating failures, optimizing the transition between durable and long-term storage (hybrid storage), eliminating redundant data sets, and more.

3. BLUEPRINT

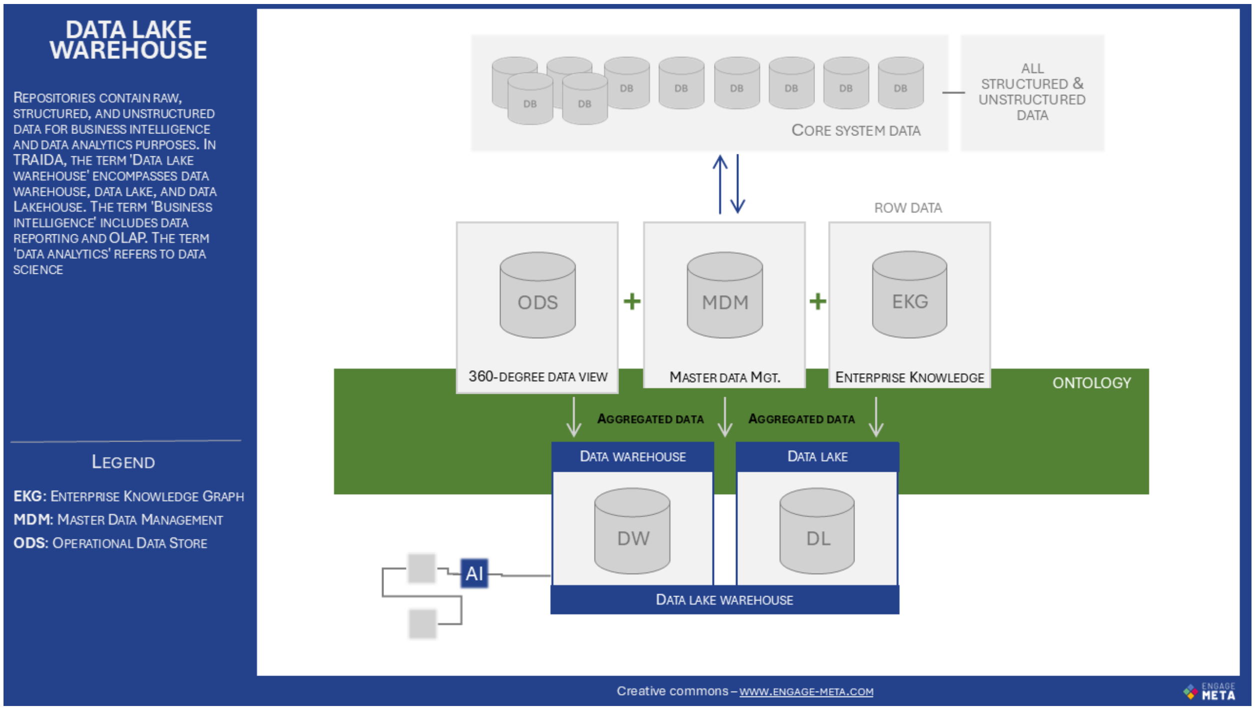

4. YOUR SITUATION & OBJECTIVES

DATA INTEGRATION

Processes and software for integrating data sources and governing data flows. The data hub might compete with the ODS (Operational Data Store) of the semantic platform; and the data fabric might compete with the EKG (Enterprise Knowledge Graph). Therefore, a choice must be made to either use the data fabric as a component of the semantic platform or integrate it with more transversal MDM (Master Data Management), ODS, and EKG.

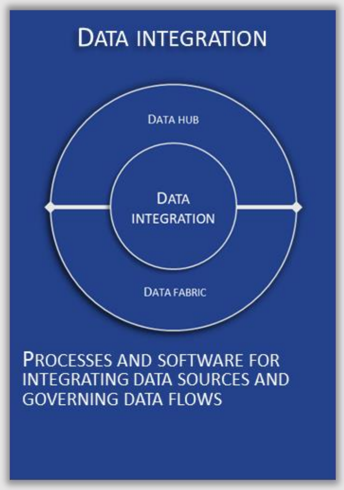

1. CONDITIONS OF SUCCESS

Data integration synchronizes and transforms multiple sources of information to provide a standardized data flow to consumers. These consumers can be repositories like MDM (Master Data Management), ODS (Operational Data Store), EKG (Enterprise Knowledge Graph), data warehouses, data lakes or application systems and AI systems for training.

Historically, this need has been covered by ETL (Extract, Transform, Load) and EAI (Enterprise Application Integration). However, to handle the complexity of integration processes, specific developments are often necessary to adapt them. These implementations become a significant technical debt and create a high rigidity in data flow integration. This rigidity is incompatible with agile governance. For instance, a simple change in data type requiring several days of maintenance would be unacceptable in a business emergency.

To address this rigidity of ETL-EAI, data hub and data fabric solutions have emerged.

Although the boundaries of these solutions vary depending on software vendors, their value proposition is based on greater agility in data flow integration. To achieve this, they use metadata and repositories for information storage that contribute to flow management. Consequently, they not only integrate data flows but also manage repositories. As vendors of these solutions ride technological and marketing waves, defining a solid architectural framework is not straightforward.

In this difficult-to-decipher marketing context, TRAIDA approaches the choice of data hub and data fabric by considering that unified data repositories like MDM, ODS, and EKG (see respective TRAIDA cards) must be preserved. They form the foundation of the semantic platform for AI.

Therefore, when considering a data hub or data fabric solution, it is essential to evaluate its ability to provide robust MDM, ODS, EKG repositories or to integrate with those of the semantic platform. For example, if the data hub establishes a metadata catalog, its integration with the shared ontologies in the semantic platform must be carefully examined. Neglecting this issue would result in managing two metadata catalogs: one at the global level housed in the semantic platform and the other accompanying data flow integration in the data hub. These two catalogs should share the same ontologies to avoid creating silos, which could lead to poor data quality and high maintenance costs.

To help you build the best data flow integration solution with the semantic platform for AI, here are the criteria to consider; for the data hub :

- The origin of the data hub is primarily technical and resembles an ETL-EAI solution enhanced with metadata management. This enhancement promotes better governance of data flows and quality control rules. The functions of data transformation, mapping, and flow propagation then rely on metadata. Consequently, the need for specific software development is reduced, making room for configuration and a NoCode approach.
- Some data hubs integrate data repository management, such as CDI (Customer Data Integration). In this case, the data hub is no longer limited to just data integration; it also provides a new repository. It then resembles a specialized ODS focused on a functional domain (here, customer management), competing with the generalist ODS of the semantic platform. Over time, some companies find themselves with a CDI data hub coexisting with other data hubs like HR, Supplier, Marketing, etc. Unfortunately, the creation of these siloed ODS repositories does not support unified data governance and generates maintenance difficulties. In the TRAIDA vision, it is preferable to build a unified ODS independent of the data hub. This would limit the data hub's scope to only managing data flow integration.

For the data fabric :

- A data fabric is a multi-technology framework that offers unified data management, with AI governance functions (management of data spaces for AI training, vectorized databases, prompt management, a unified user interface over different LLMs, etc.). It integrates data preparation and transformation functions identical to those of the data hub, based on metadata management.
- Most data fabrics incorporate a knowledge graph-oriented database technology, raising the question of integration with the EKG of the semantic platform. Unlike the data hub, which offers a solution that might compete with the ODS, the data fabric extends into the EKG repository.

How to make the right choice?

To avoid creating technological silos, the choice of a data hub or a data fabric must align with the unified repositories of the semantic platform, namely the ODS and EKG, followed by the MDM for metadata management. These repositories have a scope of action that transcends silos. They should not create new silos during data integration, as this would harm the information system's governance and data quality. These issues can arise if a data hub or data fabric imposes its own ODS and EKG repositories without sufficient integration capability with the semantic platform. To avoid this risk, follow these three recommendations:

1. **Specify your ODS and EKG needs independently**: It is necessary to specify your ODS and EKG needs independently of the study of data hub and data fabric solutions. The TRAIDA framework offers a sufficiently generic knowledge to achieve this (see the MDM, ODS, and EKG cards).

2. **Metadata management synchronization**: Data hub and data fabric solutions rely on metadata management. They impose a sort of verticalized MDM on the metadata that should be synchronized with the MDM of the semantic platform. Without such synchronization, different ontologies would be maintained between the operational management level (semantic platform) and flow integration (data hub or data fabric). These divergences degrade data quality and increase maintenance costs.

3. **Event-driven architecture**: To avoid point-to-point exchanges between systems, data flow integration relies on an asynchronous architecture, meaning event-driven management. For example, the ODS listens to a data channel for updates without being directly connected to the source application providing the flow. Modern data hubs offer this type of automation. The decoupling between repositories (MDM, ODS, EKG) and the systems that provide and consume data is crucial for architectural robustness. The alternative solution of exposing each system's access to all others resembles point-to-point exchanges. The negative consequences of this architecture are well-known and often illustrated by the metaphor of "spaghetti architecture."

By following these recommendations, you can ensure a cohesive and well-governed data integration strategy that leverages the strengths of your semantic platform while maintaining flexibility and data quality.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The domain of data integration is strategic for successfully scaling AI. The MDM, ODS, and EKG repositories of the semantic platform must be synchronized with the upstream systems that provide them with data. They are also synchronized with the downstream systems that utilize them, including AI systems for training and real-time conversation enrichment (Retrieval Augmented Generation). In the realm of business intelligence, data warehouse, data lake, and data lakehouse repositories must also be fed from the semantic platform's repositories. These repositories then expose their data flows to information visualization tools and other AI systems.

Governance and quality control of data flows

There is a need for governance of data flows, including a description of data producers, consumers, transformation and normalization rules, as well as quality control, enrichment, security, traceability, version management rules, and more.

Choosing a data hub

When choosing a data hub, the decision to use its ODS component is relatively straightforward. It depends on the transactional quality of the database. For an SME (Small and Medium-sized Enterprise) and provided that the data hub's ODS is not verticalized on a specific functional domain, it is possible to make it a transversal solution for the entire company, thus becoming a component of the semantic platform. Regarding metadata management, the architectural principles presented for selecting a data fabric should also be applied (see below).

Choosing a data fabric

Selecting a data fabric is more complex than the choice of a datahub. Indeed, it must be coordinated with the EKG repository and the MDM. Depending on the size of the company and the complexity of its technical infrastructure, several scenarios are possible; the following two are the most significant:

- For an SME: A solution centered on a data fabric with a knowledge graph-oriented data repository should also be usable as the EKG repository of the semantic platform. Metadata management could also be handled within the data fabric, replacing a more general MDM, at least during the initial deployment phase and while waiting for a more robust semantic platform. It's important to note that the MDM repository provides business governance functions similar to a complete application system, which typically do not exist in the data fabric. The ODS needs to be addressed with a transactional framework (OLTP), which is not always scalable with knowledge graph technology.

- For a Large Enterprise: Implementing MDM, ODS, and EKG repositories independently of the data fabric seems essential. The repositories in the data fabric should be considered tactical and synchronized with the master repositories in the semantic platform. If this integration is too burdensome, reconsidering the choice of data fabric is advisable; otherwise, there is a significant risk of creating silos at the level of strategic repositories. In a multi-year deployment program, it is also possible to consider some data fabric repositories as master during a limited time, awaiting reversibility into the semantic platform.

In all scenarios, ensuring the coherence of decisions hinges on sharing ontologies. No technical solution should impose specific ontologies that diverge from those shared at the enterprise level.

DATA HUB

As previously mentioned in this document, the data hub provides two main functions:

1. Integration of data flows.
2. A data repository resembling a verticalized ODS, such as a Customer Data Integration (CDI).

Deploying verticalized ODS within a data hub exacerbates the problems caused by silos, which unnecessarily duplicate data. Therefore, it is better to limit the data hub's scope to the integration of data flows, essentially an ETL-EAI augmented with a metadata catalog. This approach reduces the need for specific development in favor of flow configuration. It is essential to study the synchronization of this catalog with the MDM of the semantic platform. Finally, the data hub must offer asynchronous mechanisms for flow integration, meaning an event-driven architecture.

Ideal architecture

In summary, the ideal architecture relies on a data hub that handles data flow integration by integrating with the MDM of the semantic platform to unify metadata management (structure of flows, list of data-producing and consuming systems, service contracts, etc.). The ODS needs are not addressed at the data hub level and remain the responsibility of the semantic platform.

Contribution to AI scaling

The data hub does not directly contribute to scaling AI. However, it is essential for industrializing data flow integration both upstream and downstream of the semantic platform.

DATA FABRIC

The data fabric is a technological assembly based on data flow integration, similar to a data hub. Beyond this integration, the innovative aspect of the data fabric lies in the provision of a knowledge graph-oriented data repository akin to the EKG of the semantic platform. The technical quality of this repository and the functions offered by the data fabric determine its proximity to the semantic platform. This technical quality can be assessed through the following main criteria:

- Ability to handle increased data volume.
- Compliance with transactional management (ACID) even under intense multi-user demands.
- Ability to synchronize the metadata catalog and ontologies with third-party tools (unified MDM).
- Availability of governance functions for version management, including comparison and branch merging.

Key functions of the data fabric

The main functions offered by the data fabric include:

- Metadata manager: Some solutions provide a semantic modeling tool for building ontologies.
- Metadata and ontology use: For data flow integration, reducing specific development in favor of configuration.
- Integration with AI systems: Configuring data training flows (datasets), managing fine-tuning processes, and prompt catalogs.
- Version management: Managing versions of metadata, ontologies, and integration processes to control changes over time.
- User and access management: Security based on profiles.
- Interfaces for data visualization Tools Integration.

Data Fabric deployment scenarios

For an SME or as a first tactical deployment in a large enterprise, a data fabric with powerful knowledge graph technology can serve as a component of the semantic platform as described in TRAIDA. After the initial implementation effort, the sustainability of the technology within the semantic platform should be evaluated. This deployment mode quickly provides AI governance (training flow configuration, prompt catalog, version management, etc.).

If a tactical project is not desired for deciding the target solution, consider integrating the EKG with the knowledge graph-oriented database of the data fabric. The benefits of this choice include:

- Reduced dependency: Maintaining an EKG repository in addition to the data fabric technology reduces dependency on the latter. Changing the data fabric can leverage the data capitalized in the EKG.
- Selective data storage: The EKG can contain information that doesn't need to be copied into the data fabric. Without it, there's a risk of treating the data fabric's repository as a universal storage solution, increasing technological dependency. For example, a knowledge graph that imports regulatory documentation can be managed at the EKG level without using the data fabric's storage repository.

Key considerations for AI governance

Regardless of the path chosen for analyzing and deploying your data fabric solution, it's crucial to specify your AI governance needs. Refer to the governance cards in the TRAIDA framework for more information (green cards). The advantage of the data fabric lies in its ability to manage AI system training data flows (datasets). However, it's not just about creating prompts and uploading information into AI assistants. It's also about connecting these elements with the business concepts (ontologies) the company uses in its operations and managing versions. For example, consider how to ensure an AI forgets certain data when necessary.

Therefore, the choice of a data fabric also depends on your AI governance needs. If no existing solution meets your expectations, consider developing specific enhancements around the EKG of the semantic platform.

3. BLUEPRINT

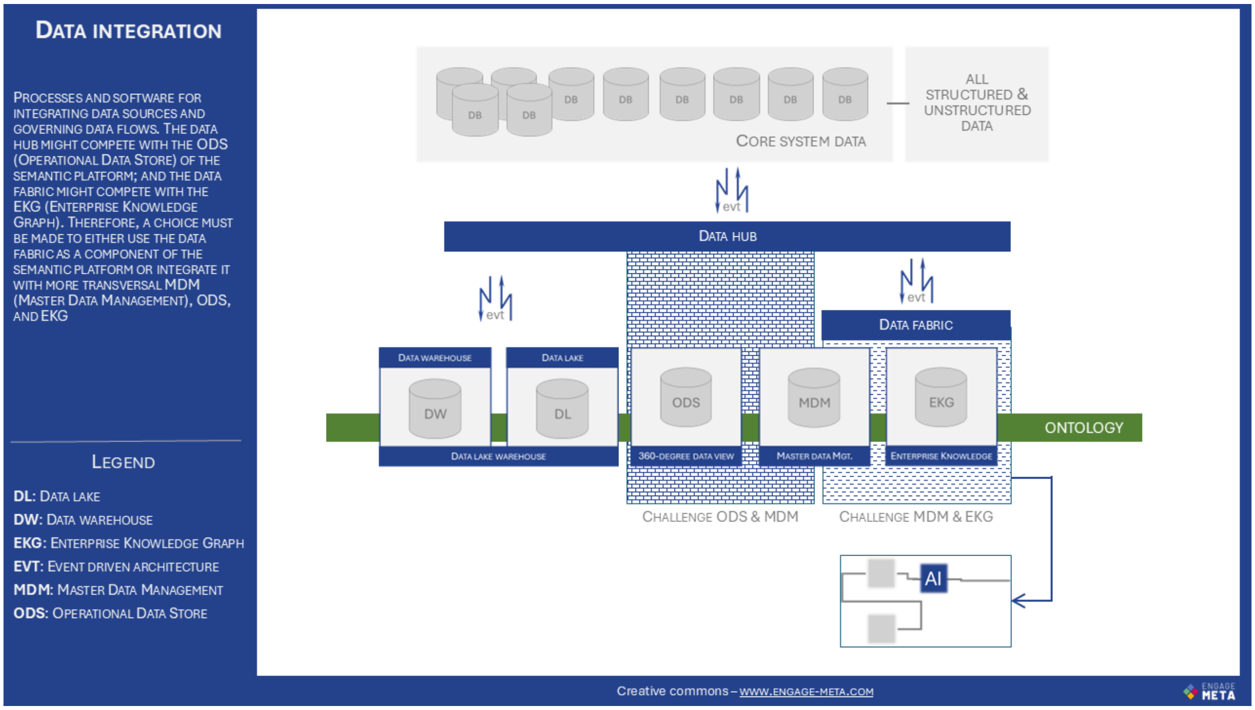

4. YOUR SITUATION & OBJECTIVES

STYLE OF DATABASE

Data storage technologies according to operational needs: transaction, integrity, concurrent access, history, data natures; volume, governance, etc. The choice of these technologies is important for deciding the architecture of the semantic platform and more specifically the MDM, ODS, and EKG repositories.

1. CONDITIONS OF SUCCESS

AI systems need to be integrated with semantic data management; otherwise, the training processes weaken and profitability does not materialize. It is thanks to metadata and ontologies that AI better understands the meaning of information. Generally, the quality level of the data provided to AI conditions the level of intelligence obtained at the end of their training and execution.

In this context, the choice of database technologies to successfully implement AI is fundamental. It takes into account these four essential needs for obtaining high-performing AI systems:

1. **Data labeling**: AI learning processes rely on metadata that serves as labels describing their usage context. For example, the metadata of a bank credit file provides the history of its subscription, the calculation of its score, and the relationships to business concepts such as the client and the financed asset. The boundary between metadata and operational data is not always stable. In practice, metadata exists through ontologies, that is, unified data models to be implemented in the semantic platform as described by TRAIDA, with MDM, ODS, and EKG repositories. Therefore, their management must be intelligently integrated with production databases and shared ontologies at the enterprise level.
2. **Description of multimedia data**: Documents (file, image, video, text…) are enriched with metadata that helps AI systems interpret them. They also document the relationships that exist with the business concepts operated by the company. For example, a client email is classified according to the nature of the request and attached to the client file.
3. **Data grouping for AI system training**: The training process of an AI requires injecting datasets of different formats and origins. For example, an AI assistant for customer relationship support is trained with product descriptions, a user guide from the online order website, an ebook published by the company, the FAQ, etc. This set of files must be kept in an archive to retain the memory of the training carried out. It will be necessary to audit the functioning of the AI and for unlearning processes when certain outdated or erroneously loaded information needs to be removed from the AI.
4. **Data injection in AI conversations** (with the RAG - Retrieval Augmented Generation technique): This involves enriching the content of AI queries with access to databases. For example, submitting a ChatGPT prompt about a client file automatically generates a read in a database to retrieve the most up-to-date client information. Thus, the AI accesses information beyond the data already injected at the time of its training. This injection principle is also used to verify and complete the response formulated by the AI; it is then an interesting way to detect hallucinations and trigger alert and correction processes.

For these use cases to be successful, the company's data must be organized in high-quality repositories. In other words, using AI on a large scale while having poor data management is doomed to fail. To avoid this impasse, data preparation and governance rely on a semantic platform as described in TRAIDA. It implements three strategic repositories: MDM, ODS, and EKG. These are positioned downstream from operational applications and upstream from decision-making systems (data warehouse, data lake).

The data in these repositories must be of the highest possible quality. To achieve this, they share the same metadata and ontologies across the enterprise. This unification facilitates their interpretation for AI training. The semantic platform ensures the management of these ontologies.

In this context, choosing database technologies to meet the needs of AI and ensure integration into your information system that considers your specific constraints is essential. Therefore, depending on your requirements in transactional management, the nature of the data handled (structured data, multimedia), exploitation needs (unitary, collective, multidimensional, search...), volume and intensity of concurrent access, and maintenance agility, several technical approaches are available and complementary. These are described in the following sections of this card.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

Data management technologies form the foundation of the semantic platform for AI. It is necessary to reconcile both efficiency and agility criteria. In other words, it is not enough to choose a high-performance technology if, at the same time, it leads to system rigidity. Conversely, it is useless to opt for a flexible solution if it does not allow for scaling across the entire scope of the enterprise.

Depending on your context, you will likely need to choose multiple technologies for MDM, ODS, EKG repositories, as well as for data warehouses and data lakes. These strategic repositories can also be accompanied by more tactical ones included in data hub and data fabric solutions (see the TRAIDA "Data Integration" card). Depending on your needs for structured and unstructured data (multimedia) volumes, as well as your transactional management requirements and simultaneous multi-user access levels, the technological choices will differ.

The most fundamental way to categorize database technologies is based on the data schema management mode, that is, how the data model is described. We will now review these categories.

STRICT SCHEMA

With strict schema technology, the data description is formal, complete, and deterministic. For example, this involves describing tables, fields, and relationships for a relational database. A data description language is used, which also allows for expressing integrity and quality control rules. Since the data structure is explicitly described, this technology enables transactional management (ACID). Additionally, it allows for response time optimization through the creation of specific indexes.

The constraints of this approach are as follows:

- Little to no capacity to store and manage multimedia data whose structure is difficult to predict. Files are then referenced as simple links to third-party storage technology.
- Lack of agility for the process of modifying data structures. This maintenance requires technical intervention at the level of the data schema description, followed by redeployment of the database. This is a delicate IT expert task that does not allow for involving business teams.
- Since there is no predetermined data schema, database governance functions are limited to technical processes, such as backup and API exposure.

Examples of solutions: Oracle Database, Microsoft SQL Server, PostgreSQL.

META SCHEMA

This approach relies on a strict data schema that can describe any other data schema. In a minimalist approach, this meta-schema could be reduced to a single object (or table) with a reflexive association. For example, a Client object is then represented as a main record, and each of its fields (or columns) is represented by another secondary record (child of the main record). In reality, the meta-schema is more complex and richer than a simple object with a reflexive relationship. Nevertheless, its structure is ignored by technical and business teams who use the database through a data modeling tool. The internal system of the database then handles the mapping between business models and the meta-schema.

The agility of this approach is much better than that of the strict schema. Indeed, the creation or modification of a new data model is done declaratively in the modeling tool without the need to technically intervene at the level of the internal schema of the database. Moreover, since the data model is dynamically introspected by the database, it becomes possible to provide advanced governance functions whose behavior adapts to the semantics of the data.

However, the following weaknesses must be taken into account:

- Multimedia data that does not have predefined structures cannot be managed intrinsically. Only links to third-party storage systems are feasible.

- Data quality control rules are no longer integrated into the internal data schema of the database, as is the case with the strict schema approach. Therefore, these rules must be entrusted to an application layer whose reliability and performance depend on the quality of the developments.

- In the absence of a dedicated data schema for each data model, it is no longer possible to optimize it specifically. Only optimizations applied at the meta-schema level are ensured and thus escape the control of the database administrator. Depending on the data volumes managed and the specifics of certain access queries, performance issues may arise and hit an optimization barrier.

- Given the lack of optimization, the behavior of the database for transactional management (ACID) on large volumes and in the context of intense multi-user access must be carefully verified. The results are generally not as good as those obtained with the strict schema approach.

Examples of solutions: Apache Hive, Microsoft Azure Data Catalog, Informatica Metadata Manager.

DOCUMENT SCHEMA

This storage technology focuses on multimedia data that does not have predetermined structures. The term "document" aptly describes the organization mode of stored information, in the form of simple documents (JSON, BSON, or XML, and binary files). Therefore, it is not necessary to create a data schema since the generic concept of a document prevails. In this sense, it is also a meta-schema approach, but this time dedicated to multimedia data.

Depending on the database used, it is possible to attach descriptive metadata to the documents in the form of a list of keys and values. The advantage of this technology is the rapid and massive storage of multimedia data without the need for a modeling step. It is common in big data projects.

The main weaknesses to note are:

- Little or no querying capability that spans multiple documents. This is a corollary of the absence of relationship management. To implement querying between a parent document and child documents, it is necessary to denormalize the storage of related documents by duplicating them in each branch of the lineage.

- No introspection of the content of the documents. The database is limited to storing them in an atomic manner. Only descriptive metadata can accompany the document to describe it.

- Little or no propensity to store structured data, and thus no transactional management or management of relationships between documents.

Examples of solutions: MongoDB, CouchDB, Amazon DocumentDB.

GRAPH SCHEMA & SCHEMA FREE

With knowledge graph-oriented database technology, the storage mode is similar to that of the meta-schema. This relies on a definition of triplets consisting of a start node, a relationship, and an end node. This is known as the graph schema.

For example, modeling a customer who orders a product is expressed in the form of two nodes for the business concepts Customer and Product, connected by a relationship that expresses the act of purchase.

Unlike the meta-schema approach described earlier, it is not mandatory to model the data structure. More precisely, it is possible to activate a "schema-free" mode that generates triplets on the fly based on a data source whose structure is discovered at the time of loading. In this case, the quality of the obtained triplets depends on the quality of the injected data. This mode of operation amounts to creating an initial version of a data model from an information source. For example, with the help of a generative AI like ChatGPT, it is possible to extract business concepts and their relationships by introspecting the content of documentation, then automatically generate the knowledge graph without having to go through a preliminary modeling step.

Thus, the source document, for example, a PDF file of financial regulations containing several dozen pages, is visualized in the form of a knowledge graph. By injecting a catalog of metadata into the AI to help identify the business concepts of the company, this graph will correspond to a version close to the final result. The TRAIDA card that deals with "core system data" also addresses this use case, applied to the analysis of information system data. You can refer to it for additional information. The advantage of the graph schema and schema-free approach lies in its flexibility when transitioning from one to the other. Knowledge graph-oriented database technology thus offers the best of both worlds:

1. Explicit modeling with the graph schema that remains generic with meta-storage in the form of triplets.

2. Schema-free mode that does not require modeling and is similar to document storage. However, it is not about storing multimedia documents but structured triplets.

The disadvantages of the graph schema and schema-free approach are as follows:

- This technology is more suitable for managing structured data, although some solutions also claim to handle multimedia data. In practice, beyond marketing intentions, this essentially involves adding a third-party storage solution, such as a file system or cloud storage like Amazon S3, Google Cloud Storage, Azure Blob Storage, or directly a document-oriented database (see above).

- As with the meta-schema approach, quality control and referential integrity management rules are relegated to an application layer, which can suffer from malfunctions and performance issues.

- Similarly, transaction management (ACID) is an overlay to the storage technology, which can also experience malfunctions at the limits of volumes and simultaneous user access.

- Specific optimization of the data model is not possible, and it is necessary to rely on the quality of the technical solution used.

Examples of solutions: Neo4j, Amazon Neptune, ArangoDB.

VECTOR DATABASE

This is an additional storage area to the databases we have previously discussed. It is necessary for providing information to AI systems.

Vectorized database technology is not dependent on a specific type of data schema. In other words, there is no need for a preliminary data modeling step. It is sufficient to inject a data source into the vectorized database. From then on, it becomes available to load information for AI system training, as well as for on-the-fly access such as Retrieval Augmented Generation (RAG).

Generative AIs like ChatGPT ensure the storage of data in a vectorized form. They handle the process of loading data sources into their own vectorization technologies, eliminating the need for manual intervention.

Examples of solutions: Pinecone, Milvus, Weaviate.

FULL-TEXT SEARCH DATABASE

Full-text indexing technology complements database queries (such as SQL) with full-text and phonetic searches. The data from the database is then injected into an indexing technology dedicated to these types of searches. This allows for set-based searches that navigate through all the data and their relationships in a performant manner. However, it is necessary to establish a technical integration with the source databases to ensure that the indexes are updated at a frequency that meets the company's needs.

Examples of solutions: Elasticsearch, Apache Solr, Algolia.

ANALYTICAL DATABASE SCHEMA

In the realm of decision-support computing, there is a need for multidimensional data analysis. This involves navigating all the relationships between data from a point of interest without having to formalize a query that could be complex. Visual navigation is required, with progressive levels of data aggregation according to various dimensions. The implementation of this type of solution relies on OLAP (Online Analytical Processing) technology, which organizes the data in the form of cubes (stars) or snowflakes (fractals).

Examples of solutions: • Snowflake, Google BigQuery, Amazon Redshift.

3. BLUEPRINT

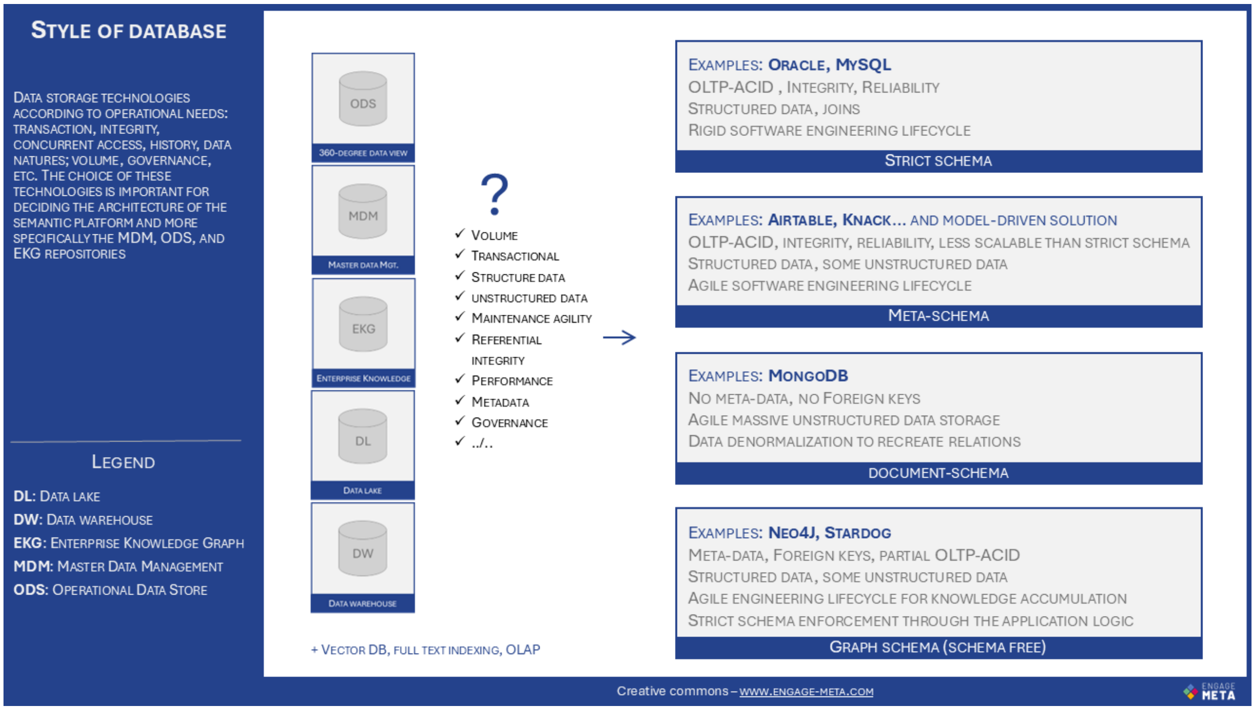

4. YOUR SITUATION & OBJECTIVES

ARTIFICIAL INTELLIGENCE

Artificial Intelligence systems function as automated and semi-automated decision-making algorithms. The different types of AI (generative, symbolic, analytical) share ontologies to facilitate their integration and use at the enterprise level.

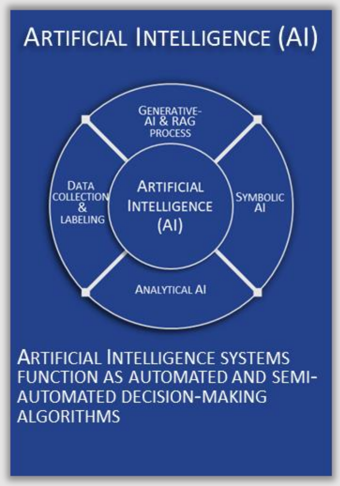

1. CONDITIONS OF SUCCESS

The interest in AI depends on the use cases of each company. Nevertheless, with broad application possibilities, significant gains are to be sought in all organizations. Indeed, AI covers a wide range of functionalities, such as:

- Creativity in communication and marketing, teaching, coaching, translation, text synthesis, report creation, financial optimization, customer tracking, trend calculations, pattern and video recognition, sound production, etc.

Beyond the specific case of a company, TRAIDA identifies two universal contributions of AI that do not depend on use cases. They form a strategic foundation so that stakeholders share certain fundamental objectives for the use of AI. Without this foundation, integrating AI into the organization encounters two riks:

- In the event of failure to implement AI in certain use cases, stakeholders may become demotivated. To counter this risk, it is important to have a framework that recalls the fundamental and shared objectives throughout the company.
- Poor implementation of AI leads to a misalignment with the company's fundamental objectives. Gains are then partially recognized by stakeholders. This context disrupts the organization and opens the debate towards questioning the profitability of AI. The strategic framework is necessary to counter this risk. It ensures that the contribution of AI for each use case aligns with the major objectives that bring stakeholders together.

To build this strategic foundation, the two universal contributions are as follows:

1. **Automate tasks**; that is, decision-making and the resulting actions.
2. **Accumulate and exploit knowledge**; in order to better control the organization.

These two contributions are identified by the majority of AI experts, but their formulation in the specific context of each company remains to be done. Indeed, automation is intimidating and requires an explanation to situate it within a framework of overall activity improvement. Similarly, knowledge management has been a recurring theme for decades, without much motivation. However, with AI, it becomes strategic and profitable.

By formalizing the two universal contributions of AI in terms that suit your company, you build your strategic AI framework. This is a document of a few pages, a sort of charter on the fundamental objectives of the company with AI.

To guide you in drafting this framework, the two universal contributions are detailed in the following section.

Automating tasks

Since its inception, computing has had the fundamental objective of automating tasks; AI is part of this trend. According to consulting firm McKinsey (2024), 70% of tasks could see a 50% productivity gain thanks to AI. Thus, it is not about 100% automation but about entrusting certain steps in work processes to AI.

This collaboration between humans and AI has consequences on the behavior of work processes. Four levels of analysis and execution must be distinguished: the design of the new process, the rules for triggering AI, its execution, and finally, its control:

1. **Design of the new process**: This involves determining the steps to entrust to AI in work processes. For each of these steps, a knowledge base is gathered to train the AI. Its execution must be consistent with the objectives of the step. This knowledge base consists of labeled or unlabeled data sets and documents already available in the company or newly created.
2. **Rules for triggering AI**: Before the AI-executed step is performed, the system checks that the conditions for its triggering are met. There is thus a context in which the AI must act; if this is not verified, the AI could be triggered incorrectly. In conventional programming logic, this preliminary check forms preconditions. They can be developed traditionally (algorithms) or rely on AI specialized in context management (symbolic AI, analytical AI, context learning, etc.).
3. **Execution**: Once the preconditions are verified, the AI executes the process step, which consists of two parts:

   a. Decision-making to build an action plan;

   b. Execution of this plan.

During process execution, points of collaboration with a human operator are possible, for example, in the form of a communication space in a workflow messenger. Execution under AI's responsibility can also act in the physical world with robot commands. The action plan constructed by AI then orchestrates its own executions, as well as those of human operators and robots.

4. **Control**: A fourth step ensures that the actions taken by AI meet the company's expectations. Two complementary intervention levels are provided:

   a. A verification that the AI's response conforms to an expected context at the process step's exit (guardrail). In conventional programming, these are postconditions. Just like preconditions, they can be developed in classical language or rely on AI. Barriers against hallucinations can also be provided by leveraging LLM's RAG (Retrieval Augmented Generation) technologies.

   b. Overall supervision of process behavior using AI dedicated to observing the information system. This AI is trained with the process specifications and all regulatory documentation. All process executions feed a knowledge base that serves as a source of information for analyzing this general AI (also called Trusted-AI, Second Brain, Nerve Center, etc.).

Accumulating knowledge

Companies manage an increasing amount of information. According to the observations of many experts[*], it is likely that only 30% to 40% of knowledge is recorded in databases. The rest, i.e., 60% to 70%, forms tacit or informal knowledge that is exclusively present in the minds of human actors.

(*) https://link.springer.com/article/10.1007/s12144-023-04994-3

Knowledge management is not a new topic for companies. It is an essential objective of computing. However, to go beyond databases and multimedia storage spaces, it is necessary to focus on transforming tacit knowledge into digital contents. Unfortunately, projects that encourage actors to document their practices in writing encounter this question:

- "*Why devote efforts to formalizing documentary repositories if they cannot be automatically exploited to improve processes?*".

In other words, once the know-how and practices are written down, it remains indispensable to employ consultants to analyze them in order to draw conclusions. This work is long, costly, often approximate, and not very compatible with the regular and necessary updates of the repositories.

Ultimately, many companies abandon overly ambitious projects for formalizing tacit knowledge. However, with AI, the potential for this formalization changes for two reasons:

1. The training of AIs is all the better when there is a formalization of tacit knowledge. It allows them to better understand the organization. In other words, AI systems that operate on 30% to 40% of the knowledge available in databases are much less efficient than those that also absorb tacit knowledge. The potential for process automation increases exponentially with each additional formalization of tacit knowledge.
2. AI enables the automatic exploitation of knowledge formalized in writing. It thus substitutes for consultants to accelerate the transition from writing to action. The barrier of the lack of profitability of documentation initiatives is then lifted.

In this context, actors must improve the quality of their writing to better formalize their work practices. For its part, the organization must explain the objectives and career plans with AI to maintain maximum trust in knowledge sharing.

The strategic foundation for profitability

The two universal contributions we have just described are fundamental for increasing the profitability of AI, particularly for its use at the enterprise level. Indeed, beyond a quick gain from carefully selected initial use cases, deploying AI across all processes of the organization naturally encounters difficulties and obstacles. In other words, you need to be precise about your fundamental objectives with AI:

- **Design work processes by entrusting certain steps to AI**: Implement a method for designing processes with preconditions, postconditions, training AI for action plan calculation, and then executing these plans collaboratively with humans and robots, and using AI for supervising the behavior of the information system.
- **Encourage teams to formalize their know-how and practices in writing**: Implement a method to improve writing quality and the automatic exploitation of this new knowledge by AI.

In addition to these two fundamental objectives, it should be noted that AI systems need high-quality data to function properly. TRAIDA is based on an architecture centered around the semantic platform with the MDM, ODS, and EKG repositories. They form the third fundamental objective in the technical field (see the respective TRAIDA cards).

In the continuation of this card's description, we detail the categories of AI to provide a better understanding of the solutions offered by the market. Although this general technical culture is useful, the management of the AI approach relies mainly on the strategic framework we have just described. Its implementation does not depend on technical mastery. To successfully scale AI in your company, you need to engage your stakeholders on this framework to build clear and widely shared commitment.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The use of AI at the enterprise level involves several types of technology. The boundaries between these types evolve with the state of the art. Given the strength of innovations, they could merge into a universal user solution. This is referred to as Artificial General Intelligence (AGI), which is predicted to arrive within a decade.

This potential for convergence reinforces the idea that companies should be interested in all types of AI, not just those highlighted by current technological trends. The effect of compounded interests should be sought by combining AIs together, as well as use cases that rely on different AIs.

To take this path toward unification, it is essential to rely on semantic data management that promotes the sharing of ontologies at the enterprise level; that is, by all use cases that rely on AI. It should be noted that an ontology is first a catalog of business concepts that the company uses to execute its activity. Then, these business concepts are organized in a hierarchy to describe their classifications and specializations in the sense of the object-oriented approach. Finally, the ontology itself emerges with the development of relationships between business concepts. This semantic modeling is the foundation for providing information to AIs and accumulating knowledge.

We now describe the different types of AI according to the current state of the art.

GENERATIVE AI & RAG PROCESS

This AI has become popular with the general public through ChatGPT by OpenAI. Since then, several similar solutions have emerged, such as Meta's LLAMA or Anthropic's Claude.

These are LLM (Large Language Model) AIs that generate text from prompts and knowledge bases. Contrary to the common belief that this AI merely predicts words to complete texts, deeper practice suggests that its functioning is more subtle. Indeed, beyond text translation and synthesis, it is possible to engage in a structured and intelligent conversation with the AI to receive help with goals that have nothing to do with text completion.

For teams within the company, it is important to experience this level of intelligence, but to achieve this, it is necessary to formulate prompts and enhance the AI's intelligence with your own knowledge bases. The AI then becomes a sort of piano, the use of which varies greatly depending on the musician's mastery. LLM is different from the application that use it as the core. For LLM it is just to do text completion - autoregressive text prediction. This kind of language model is enough to solve any natural language processing problem. The application like ChatGPT of course will do more things than just the language model (text completion), for example RLHF (reinforcement learning from human feedback), various NLP tasks, prompt analysis and guardrail.

At this point, the question of intelligence must be addressed. A lack of preparation to answer it leaves the field open to criticism from users resistant to change. Let's take the example of this negative assertion:

- "*No AI can replace human intelligence because it only recycles existing information of varying quality*."

Within the TRAIDA framework, we respond as follows:

- Intelligence consists of recycling existing knowledge to assemble and modify it so that new information emerges. The creation of knowledge ex nihilo is no longer merely intelligence but genius. At this stage, such a level of contribution is not expected within the company.
- The more powerful the learning, the more intelligence is unleashed. This power relies on two pillars for the company: the formalization of tacit knowledge and the high quality of data repositories. In other words, the perceived intelligence level of available AIs within the company reflects its own intelligence level in formalizing tacit knowledge and mastering data (ontologies). For example, using ChatGPT in its raw version does not know your company's data, so its intelligence level is close to zero for helping you in your management processes. You will need to provide it with knowledge about your activity to obtain relevant results.

This remark also applies to use cases that seem elementary, such as using ChatGPT for text translation. Indeed, two approaches are possible:

- Basic: a simple prompt with a copy-paste of the text to be translated.
- Intelligent: creation of an AI assistant to train it on a translation style, with a glossary of important terms, a directive requiring the AI not to add, change, or remove ideas, and finally, a request to add an explanatory note for the most delicate points of the translation. Moreover, this AI assistant would

be positioned as an expert in translation in the target language and an expert in the subject area of the texts to be translated (e.g., an expert in Greco-Roman history).

The translation quality then differs significantly. The first approach is even risky because the AI could engage in unnecessary hallucinations. The second approach guides the AI much better and improves the quality of the final result.

Thus, when choosing an LLM solution, it is important to study the creation of AI assistants. It is not enough to let your users create prompts; you need to organize a catalog of AI assistants so that each of them can enhance their skills in their respective fields.

Retrieval Augmented Generation (RAG)

AI Training is Conducted at Two Levels:

1. **General**: In batch mode, on large quantities of information through the loading of datasets and documents. The freshness level of the information used must be consistent with the batch mode. For example, if training is to be conducted every month, information with a lifespan of just one day should not be considered. Given the operational and financial cost of training AIs, batch mode is indispensable.
2. **Augmented**: In real-time mode, by retrieving the most up-to-date data possible from databases. This is where RAG (Retrieval-Augmented Generation) mode comes into play. It involves a technical decoupling to enrich the information flow sent to the AI (via the prompt or API activation) and control the information flow returned by the AI, for instance, to identify hallucinations. The operational and financial cost is indexed to the number of requests made.

Examples of solutions: Haystack (deepset), Replika.

Multimodal generation

Generative AI also works for formats other than text, such as images, photos, and more recently videos, as well as sounds and music. The application domains of multimodal AI are very broad, ranging from marketing content production to the creation of cinematic works, comics, music, etc.

The training process for these AIs is less within the control of companies compared to text generation. The large volumes of photos, videos, and sounds used for this training are managed by solution providers.

Nevertheless, a company with a significant repository of multimedia data can also train its own assistants to achieve AI personalization in its context. For example, a repository of users conversations (audio files) from a call center in a very specific language can be used to train a speech recognition AI.

Knowledge governance

The data used to train AIs must be archived for two reasons:

1. To keep a record of the information used in case of an audit.
2. To be able to duplicate this archive and update certain data in order to reset the AI assistant.

This second point is essential for unlearning outdated or incorrectly loaded data in the AI.

SYMBOLIC AI

Symbolic AI relies on formal rules for decision-making calculations. Unlike generative AI, the learning process is reduced to creating these rules, then organizing them into bundles, lineages, and inference logics. The rules receive input data (facts) and execute within the framework of a context of data shared across multiple executions.

Technical solutions for symbolic AI are known as expert systems or rule engines. They are useful when the knowledge domain can be formally described by rules. They are also used to increase the level of abstraction of certain programming logics, for example, for implementing preconditions and postconditions that govern the execution of services or steps in processes.

Unlike generative AI, the results of symbolic AI are auditable since it is sufficient to list the rules that were traversed to obtain an answer. Moreover, the deterministic functioning eliminates any risk of hallucination, though not of bugs if the rules are poorly organized.

The impact of this type of AI is more technical than operational. Indeed, they face a complex engineering barrier in rule modeling. Thus, it is a slower and more limited integration mode compared to what is possible with generative AI.

It is entirely possible to train a generative AI with rules described in natural language. However, with a large volume of rules (several hundred), maintenance difficulties are likely to arise with hallucination effects that are difficult to correct. When the number of rules is fewer than ten and the documentation format is clear and unambiguous, the LLM can act as a tactical rule manager.

Examples of solutions: Prolog, Drools.

ANALYTICAL AI

Analytical AI aims to predict outcomes from datasets. For example, in the medical field, a file that references diseases per patient along with associated symptoms is a dataset for training an analytical AI. Once training is successful, the symptoms of a new patient can be submitted to this AI to diagnose the potential disease. Similarly, in the healthcare domain, the principle of analyzing radiological or MRI images with AI is based on the same type of training, i.e., with a set of images that provide the disease result.

This type of AI uses machine learning technology, with numerous solutions whose performances vary depending on the targeted use cases. Their operation relies on statistical analysis, which remains more deterministic than that of LLMs. Thus, generative AI can also be used with the perspective of analytical AI but for limited datasets to avoid hallucinations. For example, based on a dozen criteria from a quote request, it is possible to determine the email response to send to prospects. It is sufficient to provide the generative AI with a few dozen examples of quote analysis along with the reference email used to achieve the desired automation. The complexity of cases here is much lower and less sensitive than for the medical domain. LLMs are fundamentally statistical in nature as well. The primary distinction lies in the sheer scale of LLMs compared to traditional statistical machine learning models.

Examples of solutions: Scikit-Learn, TensorFlow, RapidMiner.

DATA COLLECTION & LABELING

This final topic does not address a specific type of AI but rather the labeling of data necessary for learning processes. This labeling occurs at several levels:

- During the collection of tacit knowledge to attach it to ontologies. This information capture can be directly integrated into production applications. For example, an insurance operator analyzing a claim file has a button on the UI to enter a text explaining how they conduct their analysis. This knowledge is then linked to the type of claim being managed (ontology). A quality control process is initiated so that the AI manager for the claims department can assess the relevance of this new knowledge and accumulate it for future AI training enrichment.
- During the preparation of datasets, for example, to indicate the location of a problem on each photo of an electronic board with a manufacturing defect. A platform is needed to manage this enrichment process.
- During the evaluation of AI responses to improve future results. For example, the image generation solution MidJourney operates within Discord. It offers the user a series of four images to choose the best one. The user's response is given naturally in the discussion thread. By making this choice, the user helps improve the AI. This evaluation principle can be replicated in production applications. For instance, at the end of a process executed with AI, the system proposes a quiz to the user to

evaluate the positive and negative points of the work. This is similar to the performance evaluation of an employee but applied to AI within the execution of a process.

These data labeling mechanisms prepare the training and allow for the evaluation of results to continuously improve the AI (human in the loop). They should be integrated as closely as possible to applications and process execution.

Examples of solutions : Labelbox, Amazon SageMaker Ground Truth, Prodigy, SuperAnnotate.

3. BLUEPRINT

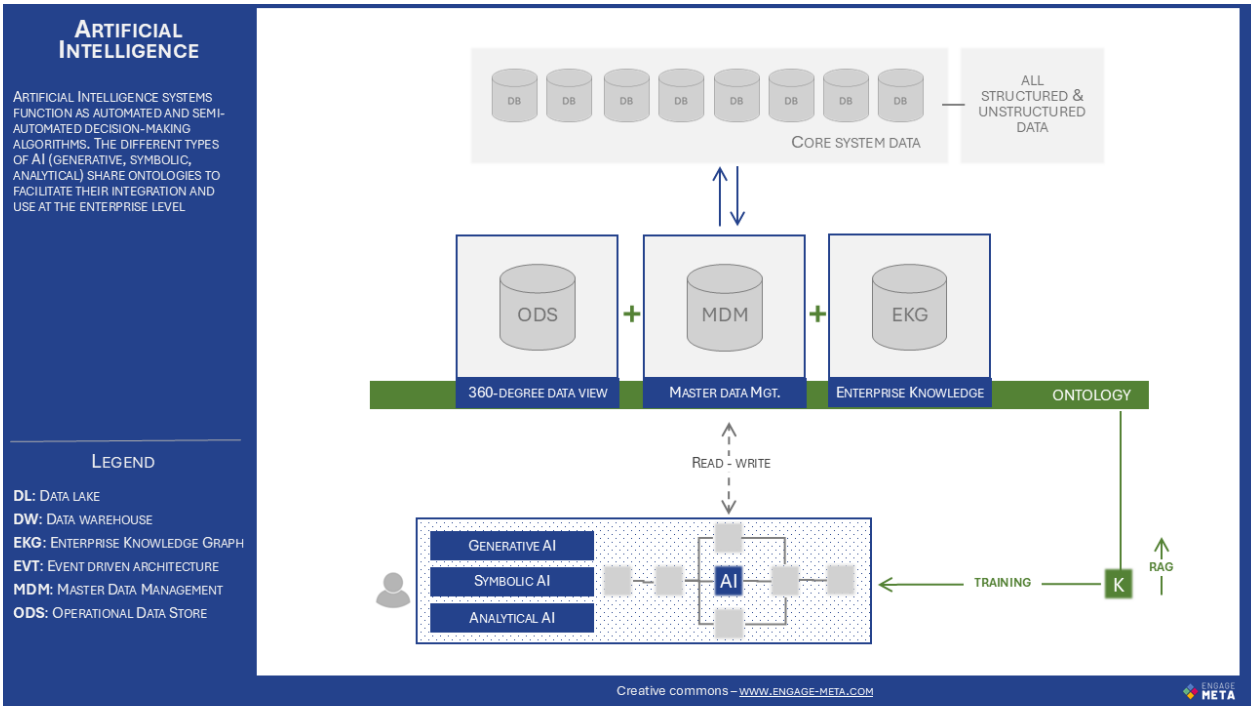

4. YOUR SITUATION & OBJECTIVES