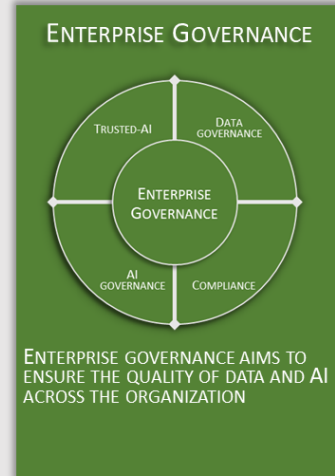


ENTERPRISE GOVERNANCE

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 TRAI DA 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 TRAIIDA 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 TRAIIDA 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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 have multiple versions with different training prompts and documents.
Document Management	<ul style="list-style-type: none"> • 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 "no 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.

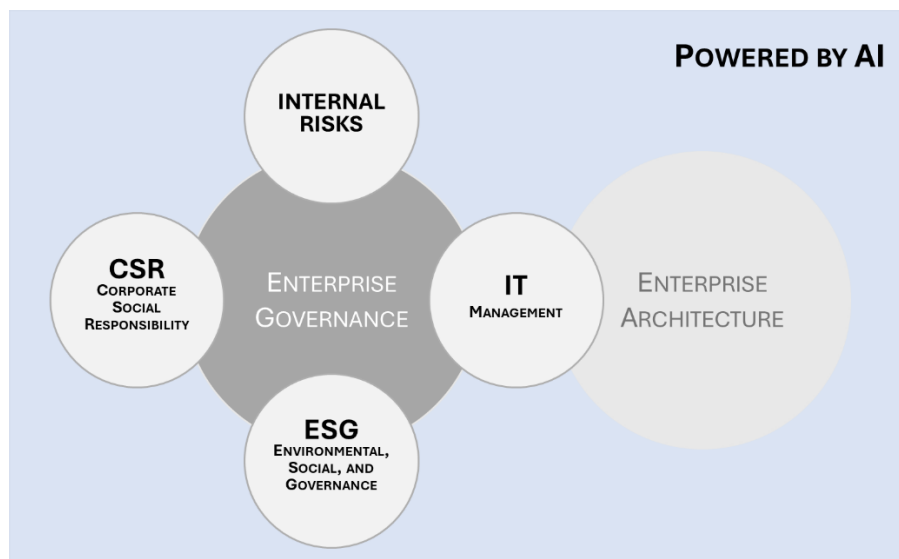
With the introduction of AI at the enterprise level, the entire behavior of the enterprise system becomes more intelligent, along with the risks of losing control of it for the two aforementioned reasons.

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