

RESPONSIBLE AI

Manifesto

NOVEMBER, 2026

RESPONSIBLE AI
VISION



ISABELLE QUAINON
Chief Human Resources Officer

As AI develops quickly, Veolia is choosing a responsible approach. This means understanding both its benefits and its risks, and setting clear rules for its use.

AI is about serving our core purpose: delivering essential water, waste, and energy services daily. To keep our high level of quality, technology must always support what matters most: our expertise, collaboration, and commitment of our teams on the ground.

If used incorrectly, AI can weaken people's skills. Used well, it can strengthen human expertise.
AI must remain a support tool that helps us spot problems faster. People must always stay in control and use their critical thinking.

As AI transforms our work, we purposefully deploy it with and for our users, empowering everyone to evolve with confidence and ensuring that AI creates progress for all.



CHRISTOPHE MAQUET
Chief Business Performance Officer

At Veolia, we view AI as a tremendous opportunity, but we refuse to adopt it blindly.

We know that these digital tools require significant water, energy, and physical resources. That is why we only deploy AI where we create value—using it as a lever to protect resources across our operations. And we also provide circular solutions to reduce the footprint of the tech industry's own data centers.

Our approach is built on two core principles. First, efficiency and utility: we select only secure AI solutions tailored to our needs to prevent digital waste. Second, security: with these tools, all our data remains strictly within Veolia, safe from leaks, and under our control.

This is how we leverage technology to drive environmental security, with absolute responsibility to all our stakeholders.

Table of Contents



Responsible AI
MANIFESTO

Executive Summary **From Dialogue to Action: Veolia's Approach to RESPONSIBLE AI**	**P.3**
Commitments, Actions by theme **Responsible AI for PEOPLE**	**P.5**
Responsible AI for ENVIRONMENTAL SECURITY	**P.10**
Responsible AI for RESILIENCE & GOVERNANCE	**P.14**
AI in writing this Manifesto	**P.18**
Glossary	**P.19**

Words with an asterisk (*) are defined in the glossary on page 19

From Dialogue to Action: Veolia's Approach to RESPONSIBLE AI

At Veolia, AI must be responsible, or not be used at all. We believe artificial intelligence can create lasting value only when it strengthens human expertise, protects resources, and supports the resilience of the essential services we provide in water, energy, and waste management.

This Manifesto defines Veolia's convictions and commitments for the responsible use of AI. It is not simply a statement of principles: it is a practical framework for guiding our decisions, investments, and uses of AI across the Group. It reflects our ambition to ensure that technology remains at the service of our purpose, ecological transformation, reconciling human progress with environmental security.

From experience to shared commitments







The Manifesto is the result of a three-year process. In September 2023, Veolia rolled out VeoliaSecureGPT—an internal platform of secure AI tools—to all its employees. This first milestone enabled us to identify the operational potential of AI while assessing its impacts on our employees, our operations, our clients, our stakeholders, and the environment.

Drawing on this feedback, we established an ongoing dialogue among governing bodies, experts, social partners, and other stakeholders. This collaborative consultation connects frontline reality with a clear-eyed analysis of the risks involved. It shapes our approach: establishing a shared diagnosis ("What we know") and making concrete commitments ("We act").

This Manifesto is the culmination of that collective journey. It translates ground-level insights and our shared convictions into tangible commitments, while acknowledging that responsible AI is not a one-time achievement, but an ongoing responsibility.

Words with an asterisk (*) are defined in the glossary on page 19

Responsible AI approach, the 5 guiding principles

The Manifesto defines Veolia's stance on the responsible use of AI: an AI that enhances human expertise without replacing it; that creates more value than the resources it consumes; and that operates within resilient systems where humans retain control over decision-making.

Our approach is built on five principles that define the conditions for responsible AI at Veolia:



From principles to actionable commitments around 3 themes

Our responsible AI approach is built around three core themes: People, Environmental Safety, and Resilience & Governance.For each of these themes, we have established six operational commitments, summarized in the overview table on the following page.
To explore these assessments further and discover the details of our concrete actions, we invite you to consult the dedicated thematic chapters (pages 6 to 18).

A living framework for continuous improvement

This Manifesto is not a static declaration. It is both a statement of intent and a roadmap for action. It is a living framework that will evolve alongside our practices, our feedback, and technological advances. It marks the transition from dialogue to action, with a clear ambition: to make AI a driver of performance, collective progress, and ecological transformation, within a framework of trust and responsibility.

Responsible AI thematic summary

Theme Our Guiding Principles	We know Our Diagnostic	We act Our Commitments
Responsible AI for PEOPLE  ➜ AI augments expertise, it does not replace it. ➜ AI demands continuous dialogue, awareness, and training.	**We know** that without training, both general and tailored to the hardest-hit sectors, AI risks creating a deep digital divide within teams. An overreliance on automated tools can erode essential hands-on expertise, teamwork, and critical thinking. If left unchecked, adoption sparks legitimate concerns regarding job security, day-to-day working conditions, and the intergenerational transfer of knowledge.	1. We invest massively in AI literacy and training. 2. We identify and protect critical skills and expertise & We use AI as a personalized coach at scale to develop our employees. 3. We equip managers to make AI a shared resource for their teams. 4. We see AI as a lever to enhance our employees, not to replace them. 5. We maintain dialogue on AI usage with all stakeholders, starting with employee representatives. 6. We promote responsible AI use and commit publicly with other companies. *Please refer to pages 5–9 to discover our detailed actions.*
Responsible AI for ENVIRONMENTAL SECURITY  ➜ AI creates value only if it saves more than it consumes.	**We know** that operating complex AI models requires immense amounts of electricity and water. Without strict discipline, digital deployment leads straight to energy waste. Meanwhile, the rapid global expansion of data centers strains local ecosystems, making circular resource solutions urgent.	1. We only deploy AI where we create value, using it as a lever to protect resources. 2. We apply "frugal by design" principles to minimize AI resource consumption. 3. We make AI's ecological impact visible to ensure responsible usage. 4. We tackle digital waste by promoting sober and useful AI. 5. We put our industrial ecology solutions at the service of data center sustainability and performance (DataCenter Resource 360). 6. We commit publicly alongside coalitions, stakeholders, and especially future generations. *Please refer to pages 10–14 to discover our detailed actions.*
Responsible AI for RESILIENCE & GOVERNANCE  ➜ Humans remain in control of all AI decisions. ➜ AI requires resilient IT systems and strict data control.	**We know** that managing drinking water, clean power, and waste plants leaves zero margin for machine error on critical infrastructure. Unauthorized tools expose company data, depending on a single vendor creates systemic vulnerabilities, and unmonitored systems quickly trigger cost inflation.	1. We coordinate AI deployment through clear governance, from strategy to execution. 2. We protect company data through robust governance, quality management, and security controls. 3. We ensure our digital autonomy through modular and interchangeable solutions. 4. We prevent unauthorized AI (Shadow AI*) through secure internal solutions and clear guardrails. 5. We maintain strict control over our AI expenditures. 6. We secure our critical operations without ever delegating the final decision to AI. *Please refer to pages 14–18 to discover our detailed actions.*

Commitments & Actions by theme



Responsible AI for PEOPLE

Guiding Principles

- ➜ AI augments expertise, it does not replace it.
- ➜ AI demands continuous dialogue, awareness, and training.

We want to give artificial intelligence its full place, not as individual tools, but as shared resources serving the collective. We choose to deploy AI with and for its users, with a clear intention, consistent with our values and our purpose, so that it is a source of progress for all.

We Know **Our Diagnostic**	**We Act** **Our Commitments**
We know that without training, both general and specific to the sectors most affected, AI risks creating a deep digital divide across teams. An over-reliance on automated tools can erode essential field expertise, team work and critical thinking. Left unmanaged, adoption raises legitimate concerns about job security, daily working conditions, and knowledge transfer across generations.	1. We invest massively in AI literacy and training. 2. We identify and protect critical skills and expertise & We use AI as a personalized coach at scale to develop our employees. 3. We equip managers to make AI a shared resource for their teams. 4. We see AI as a lever to enhance our employees, not to replace them. 5. We maintain dialogue on AI usage with all stakeholders, starting with employee representatives. 6. We promote responsible AI use and commit publicly with other companies. *Please refer to the following pages to discover our detailed actions by commitment.*

1.We invest massively in AI literacy and training

Our strategic objective centers on equipping every employee with essential AI literacy and training, ensuring seamless alignment with the EU AI Act requirements while embedding ethical AI standards throughout our global operations. We are committed to extending this comprehensive initiative to our entire workforce by the close of 2027.

We know

The generative AI literacy gap is a major societal challenge. Global trends show a strong adoption of generative AI among students and managerial staff, contrasted with significantly lower usage among

seniors and non-managerial staff. Without universal training, there is a direct risk of reproducing a digital divide across Veolia's workforce.

We act

We have:

- Trained over 60,000 employees in AI by October 2026, primarily through a 45-minute e-learning course covering foundational concepts, environmental and societal impacts including biases, and responsible usage practices, followed by specialized prompt* engineering modules.
- Built and deployed a global network of over 1,600 "AI Champions" equipped through a 19-hour learning pathway (refreshed quarterly) to drive AI adoption across all regions and business lines.
- Provided a comprehensive catalog of training programs across workshop and e-learning formats, spanning prompt* engineering, responsible AI principles, data governance, and role-specific use cases.
- Created and mobilized “LoveAI”, an internal community of more than 3,300 members, to foster peer-to-peer knowledge sharing, real-time support, and best-practice dissemination. The name “LoveAI” is an anagram of “Veolia”.
- Organized interactive team events ("AI Cafés", informal team sessions where colleagues discuss practical ways of using AI) to facilitate open discussion and exchange of practical usage techniques among colleagues.
- Proposed operational field discussions ("AI field talk", short discussions on AI held directly with frontline operational teams) to drive awareness and structured dialogue directly within frontline operational teams.

2. We identify and protect critical skills and expertise & We use AI as a personalized coach at scale to develop our employees

When generative AI is placed directly in the hands of operators, used the wrong way it can erode expertise, but used the right way it builds expertise. Generative AI can be used as a "personalized coach at scale", complementing initial training.

We know

Over-reliance on automated systems carries a direct threat to domain know-how. As algorithms take over operational steps, critical thinking and foundational field expertise risk gradual atrophy. More generally, the development of AI in technical scopes will lead to an even stronger need to develop the "soft skills" of our employees (critical thinking, empathy, etc.)

We act

We have:

- Deployed AI-driven interactive training to prepare technical and commercial teams for client dialogues.
- Designed systems in a "co-reflection" mode (a mode in which the AI asks questions and supports the user in finding the appropriate solution) in order to preserve the expertise of our employees. This ensures that the AI is used as a source of reflection rather than a replacement for it. This “co-reflection” methodology was developed as part of the AI program at Water France for the

Words with an asterisk (*) are defined in the glossary on page 19

optimization of wastewater treatment plant operations.

- Built on and shared the co-reflection methodology with AI Champions and AI Coordinators as well as in the context of a global challenge called “Veolia Agentic Challenge”, launched in July 2026 to scale agents* built in our proprietary platform VeoliaSecureGPT. This methodology is also accessible through a special e-learning.
- Defined a methodology of analysis of AI-impacted tasks to safeguard human expertise, collective work, engagement and foster critical thinking in order to implement adapted solutions:
  ➜ Limit the scope entrusted to AI to uses that do not put critical skills at risk;
  ➜ Design a "co-reflection" AI that maintains engagement, or even allows learning;
  ➜ Acculturate the right reflexes to have when using AI, to encourage critical thinking.

3. We equip managers to make AI a shared resource for their teams

We know

The deployment of AI confronts managers with new challenges: overseeing its use, preventing risks (including mental health impacts), and maintaining a collaborative dynamic and a healthy working climate. The manager's role lies, and will increasingly lie, in mastering what surrounds AI: ensuring it does not become a source of discrimination, keeping humans at the center of decision-making, upholding ethics, and preserving social bonds. When properly framed, however, AI can also play a key role in supporting diversity, equity, and inclusion across collaborative workflows and team interactions.
At the same time, AI disrupts the traditional relationship between knowledge and expertise, notably by challenging the transmission of knowledge between senior and junior profiles. If left unmanaged, these new usages could accentuate generational fractures, compelling us to rethink how we onboard and train junior profiles.

We act

We have:

- Developed workshops dedicated to Executive Committees and Management Committees.
- Created dedicated training for managers and leaders to navigate workforce dynamics in the AI era.
- Implemented a systematic approach to question the impact of AI use cases on collaboration and team engagement.

4. We see AI as a lever to enhance our employees, not to replace them

We know

AI boosts productivity and frees up time from repetitive tasks, it also opens new ways of serving our clients. Therefore, AI inevitably transforms job roles, skills requirements, and operational workflows.

We act

We have:

- Conducted an initial study in 2025 based on interviews with business coordinators to understand the

Words with an asterisk (*) are defined in the glossary on page 19

real impact of AI on tasks and jobs.
- Launched an AI Observatory early 2026 to monitor the real and localized impacts of AI on daily tasks, skills, teamwork, and engagement across 10 use cases and more upcoming field observations.

5. We maintain dialogue on AI usage with all stakeholders, starting with employee representatives

We know

Digital transformation raises legitimate workforce concerns regarding job security, tool intrusiveness, and quality of working life.

We act

We have:

- Initiated a transparent dialogue with French and European social partners since 2024 to identify risks and prioritize actions.
- Pursued this dialogue on a regular basis, sharing our observations, progress, experiments, throughout presentations, workshops and learning expeditions.

6. We promote responsible AI use and commit publicly with other companies

We know

Responsible AI adoption requires collective action. Without shared standards and cross-industry collaboration, companies risk developing fragmented approaches that fail to address systemic risks. Public commitments and coalitions accelerate the adoption of responsible practices across the entire ecosystem.

We act

We have:

- Organized a week on responsible AI inside the company to promote our commitments and widely shared communication toolkit across the Group.
- Co-signed the Manifesto for Collaborative AI and Human Responsibility alongside 16 other major enterprise partners (Projet SENS collective).

Commitments & Actions by theme



Responsible AI for ENVIRONMENTAL SECURITY

Guiding Principles

➜ AI creates value only if it saves more than it consumes.

As a global leader in ecological transformation, Veolia has a responsibility to lead by example. Artificial Intelligence should contribute to addressing the climate crisis while its environmental footprint is measured and actively reduced. Driven by its core mission to "Resourcing the World," Veolia firmly believes in AI as a powerful catalyst to enhance environmental security.

We Know **Our Diagnostic**	**We Act** **Our Commitments**
We know that operating complex AI models requires immense amounts of electricity and water. Without strict discipline, digital deployment leads straight to energy waste. Meanwhile, the rapid global expansion of data centers strains local ecosystems, making circular resource solutions urgent.	1. We only deploy AI where we create value, using it as a lever to protect resources. 2. We apply "frugal by design" principles to minimize AI resource consumption. 3. We make AI's ecological impact visible to ensure responsible usage. 4. We tackle digital waste by promoting sober and useful AI. 5. We put our industrial ecology solutions at the service of data center sustainability and performance (DataCenter Resource 360). 6. We commit publicly alongside coalitions, stakeholders, and especially future generations. *Please refer to the following pages to discover our detailed actions by commitment.*

1. We only deploy AI where we create value, using it as a lever to protect resources

We know

The large models often used in Generative and Agentic* AI require a lot of computing power, whereas standard automated tools such as Machine Learning* are much simpler. Using advanced AI without making sure it brings real value risks wasting energy and water without delivering meaningful environmental benefits.

We act

We focus on use cases where we can measure environmental and operational benefits, and we set our deployment priorities accordingly. Below are concrete examples of our impact in action.

- The progressive implementation of Hubgrade, Veolia's suite of digital and AI solutions for operating water, energy and waste infrastructure, has contributed to the following results:
  - **Wastewater:**
    A 20% reduction in aeration time, according to a study by the Collège de France led by Prof. P. Aghion; more than 10% energy savings across 300 sites, with AI-related consumption representing less than 1% of site energy use. Average savings of 30% (energy/chemicals) across 100 plants (EU, US, Japan, Australia);
  - **Energy:**
    Over 25% production/distribution through dispatch optimization and flexibility enhancements;
  - **Waste collection:**
    50 AI-equipped trucks (33 sites in France) reduce pollutants in recyclables from 20% to close to zero, improve compliance, and prevent fire risks;
  - **Other promising use cases with ongoing pilots:**
    Cooling water chemistry, AI-powered substations, membrane maintenance optimization, leak detection, embedded robotics.

- **The Gen AI and Agents* "Hubgrade AI Companion"** service suite will enable us to go even further in large-scale use of real-time industrial and operational data from our plants to unlock additional savings and deliver a positive environmental impact.
  Early results from selected pilots indicate the following potential benefits:
  - In desalination plants: up to 10% energy reduction and 5% chemical dosage reduction;
  - In waste-to-energy combustion efficiency: up to 3% steam production increase;
  - In district heating network decarbonization: 1% to 5% CO_2 emissions reduction;
  - In water leak detection: faster localization reducing unnecessary site visits.

2. We apply "frugal by design" principles to minimize AI resource consumption

We know

Developing large-scale AI systems from scratch, such as Large Language Models (LLMs*), requires substantial energy and water resources. For instance, training GPT-4 consumed an estimated 50-70 GWh of energy and approximately 600 million liters of water. By deploying these tools with clear and targeted objectives, through optimized architectures, appropriately scaled models, and efficient prompt* iterations, it is possible to minimize 'token drift*,' thereby reducing resource waste and financial overhead while enhancing overall performance and sustainability.

We act

We have:

- Refrained from training proprietary foundation models from scratch, thereby avoiding the additional energy and water consumption associated with undertaking such large-scale training internally.

- Implemented intelligent model routing in VeoliaSecureGPT that automatically matches each request with the most appropriate model, whether Small Language Models (SLMs*) for simple tasks or larger models for complex queries, ensuring we use only the computing power necessary for each use case.
- Built prompt* optimization capabilities to compress multi-turn back-and-forth interactions into a single optimized request, reducing token* consumption.

3. We make AI's ecological impact visible to ensure responsible usage

We know

Employee AI usage carries a real footprint that must be made visible to drive accountability. However, as AI interactions become increasingly complex and embedded within business tools, accurately measuring their environmental impact becomes more challenging. Unlike standalone queries, agentic* workflows, automated processes, and integrated AI features generate distributed computational loads that are difficult to isolate and track. Cloud providers often lack granular reporting on the carbon footprint of specific AI operations, making it nearly impossible to obtain precise, real-time environmental data at the individual use case level.

We act

We have:

- Required our cloud and AI suppliers to provide environmental impact data and carbon footprint reporting, pushing for greater industry transparency.
- Integrated environmental impact estimation from the open-source EcoLogits tool into VeoliaSecureGPT, displaying carbon equivalents to raise user awareness, while acknowledging current measurement limitations.

4. We tackle digital waste by promoting sober and useful AI

We know

Unrestricted AI usage can trigger unnecessary computational processes that consume energy without delivering proportional value. Features like automatic content generation, uncontrolled query iterations, and data duplication across multiple systems amplify resource consumption. Without clear guidelines and data architecture discipline, AI deployment risks creating digital waste at scale.

We act

We have:

- Maintained a curated library of business-specific prompts* and agents* to reduce unnecessary query iterations.
- Provided controlled access to compute-intensive features, such as image and video generation, through usage quotas and environmental guidance. This approach is intended to reduce uncontrolled use while preserving access to legitimate business applications.
- Incorporated digital footprint awareness into our employee AI training programs, with specific focus on resource-intensive features.

- Organized Digital Cleanup campaigns to raise awareness and lead employees to delete unnecessary files, emails, and data, reducing storage and energy consumption.
- Worked toward data centralization in data domains and data lakes, aiming to store data once and access it for multiple AI use cases rather than replicating it across systems, which will further reduce storage and processing overhead.

5. We put our industrial ecology solutions at the service of data center sustainability and performance (DataCenter Resource 360)

We know

Global data center energy consumption is set to double by 2030, with cooling drawing heavily on municipal freshwater supplies and waste heat* being vented into the atmosphere. As AI-driven digital infrastructure accelerates worldwide, the primary challenge is no longer just technological, it is territorial. The "social license to operate" depends on the industry's ability to prove data centers are long-term contributors to their host communities rather than resource-intensive islands. Without industrial-scale environmental solutions that transform data centers into integrated, circular assets, the expansion of AI infrastructure risks triggering resource conflicts and community resistance, ultimately threatening the sustainable growth of the global digital economy.

We act

We have:

- Launched our “DataCenter Resource 360” Offer in April 2026, enabling Veolia to offer an integrated approach across the data center value chain, including circular water management, waste heat recovery, energy supply, CO_2 capture and territorial integration.
- Engaged with policymakers globally to advocate for regulatory frameworks that enforce operational resource efficiency, accelerate territorial decarbonization, and safeguard long-term water and energy accessibility for local communities.

6. We commit publicly alongside coalitions, stakeholders, and especially future generations

We know

The environmental impact of AI is a global, systemic challenge that no single company can solve in isolation. Driving sustainable AI across industrial operations requires collective advocacy, transparent knowledge-sharing, and joint action - not only between industrial leaders and technology providers, but also with future generations who will inherit the long-term consequences of today's technological choices and expect companies to act responsibly.

We act

We have:

- Signed the AI Pledge and are a founding member of the Coalition for Sustainable AI launched at the

Words with an asterisk (*) are defined in the glossary on page 19

Paris AI Action Summit (Feb 2025).

- Co-published a white paper with Institut Veolia and Microsoft on AI energy efficiency in industry.
- Engaged directly with young leaders through Veolia's Future Generations Council, co-creating actionable proposals such as embedding environmental stewardship and community dialogue from the outset of AI/data infrastructure projects, and leveraging AI and climate data to quantify the true cost of environmental inaction.

Commitments & Actions by theme



Responsible AI for RESILIENCE & GOVERNANCE

Guiding Principles

➜ Humans remain in control of all AI decisions.

➜ AI requires resilient IT systems and strict data control.

As an operator of critical infrastructure serving millions of people across more than 60 countries, Veolia bears a unique responsibility: any AI failure could directly impact public health, environmental safety, and essential services. This operational reality, combined with the Group's decentralized structure where no single team can control every use case from the center, demands that reliability, security, and accountability be built into every AI system from the start. These safeguards are the foundation of trust with clients, employees, and regulators. That is why Veolia scales AI technology once it has proven to be secure, reliable, and worthy of trust.

We Know **Our Diagnostic**	**We Act** **Our Commitments**
We know that managing drinking water, clean power, and waste plants leaves zero margin for machine error on critical infrastructure. Unauthorized tools expose company data, depending on a single vendor creates systemic vulnerabilities, and unmonitored systems quickly trigger cost inflation.	1. We coordinate AI deployment through clear governance, from strategy to execution. 2. We protect company data through robust governance, quality management, and security controls. 3. We ensure our digital autonomy through modular and interchangeable solutions. 4. We prevent unauthorized AI (Shadow AI*) through secure internal solutions and clear guardrails. 5. We maintain strict control over our AI expenditures. 6. We secure our critical operations without ever delegating the final decision to AI. *Please refer to the following pages to discover our detailed actions by commitment.*

1. We coordinate AI deployment through clear governance, from strategy to execution

We know

Operating across more than 60 countries requires a structured governance framework that balances Group-wide strategic alignment with local operational agility and strict regulatory compliance, including the EU AI Act and emerging AI regulations worldwide.

We act

We have:

- Established a multi-tier governance structure: overall strategy steered by a Steering Committee, cross-functional alignment driven by Group "AI Coordinators", and local adoption executed by Zone AI Strategic Leads.
- Collected ideas and existing AI projects in a global AI use case register to consolidate use cases and foster reuse and best practice sharing, and implemented a value-driven prioritization framework.
- Established a cross-functional AI Act Task Force and conducted a comprehensive inventory and classification of all AI systems across the Group, confirming compliance with EU AI Act prohibited use cases.
- Reinforced the existing mandatory validation process for all digital solutions used in the Group (Design & Legal Assessment) with AI-specific controls.

2. We protect company data through robust governance, quality management, and security controls

We know

Inputting proprietary operational or client data into public AI models risks sensitive data leaks, IP exposure, and algorithms being trained on company assets. Beyond security threats, deploying AI at scale requires robust data foundations: without proper organization and governance, master data management, quality standards, and harmonized data models across Business Units, AI systems cannot deliver reliable outputs or enable effective knowledge sharing across our Business Units. Data governance is both a security imperative and a performance enabler.

We act

We have:

- Deployed a comprehensive data transformation and awareness program and trained more than 3,900 employees through Group-led and Business Unit initiatives, including face-to-face sessions, while over 5,000 employees completed the "Data Passport" e-learning course. This data transformation program develops data literacy and understanding of the data lifecycle, from data collection to AI consumption, as well as awareness of data security risks, the proper handling of sensitive information, and the importance of data quality for AI effectiveness.
- Enforced mandatory Design & Legal Assessment validation for all AI applications to verify compliance with data standards and intellectual property policies.

- Architected an end-to-end time-series data pattern ensuring high quality, security, and reliability, and enabled more trustworthy AI applications, with reduced risks of hallucination* through rich contextual data enrichment.
- Introduced binding contractual GenAI clauses prohibiting third-party AI suppliers from using Veolia's or clients' data to train, test, or fine-tune public models without prior written consent.
- Implemented master data management frameworks to establish single sources of truth for critical operational data across our operations.
- Created data catalogs and metadata management systems to improve data discoverability and traceability.
- Compartmentalized core operational data within private cloud environments, AWS, Azure, and GCP, with strict encryption and access controls to prevent exposure to public AI models.
- Isolated industrial SCADA* data and real-time client operational inputs from all external AI systems to guarantee data confidentiality across our global operations.

3. We ensure our digital autonomy through modular and interchangeable solutions

We know

Relying exclusively on a single vendor for IT or AI services introduces severe lock-in risks, creating operational vulnerabilities to technical outages, vendor price inflation, or geopolitical shifts. Furthermore, deploying rigid, oversized models leads to unnecessary financial costs and deprives operations of the agility required to meet real-world needs.

We act

We have:

- Architected *VeoliaSecureGPT*, an enterprise-grade Generative AI platform focused on technological resilience, a multi-cloud solution with a model-agnostic gateway, enabling the business to easily transition between different AI providers while ensuring continuous, uninterrupted operations.
- Implemented dynamic query routing that automatically selects the most appropriate model (LLM* or SLM*) based on task complexity and resource efficiency, applying "right-sizing" methodology.

4. We prevent unauthorized AI (Shadow AI*) through secure internal solutions and clear guardrails

We know

Unapproved external AI tools used on personal devices or corporate accounts bypass IT security, with the risk of leaking corporate intellectual property and expanding the attack surface for credential theft. Without visibility and governance, Shadow AI proliferates uncontrolled across the organization.

We act

We have:

- Published an official Group-wide policy directive to all employees, clarifying authorized AI tools, prohibited practices, and baseline security rules.
- Deployed secure internal alternatives (VeoliaSecureGPT and Gemini), giving all over 150,000 connected employees safe, monitored, and performant tools that eliminate the need for external AI services.
- Activated perimeter monitoring to flag unapproved external AI tool usage.
- Implemented automated blocking mechanisms for high-risk Shadow AI* platforms, preventing access to unauthorized generative AI services from corporate networks.
- Established guidance and redirection workflows: when employees attempt to access blocked external AI tools, they receive immediate in-browser notifications explaining risks and directing them to approved secure internal platforms.
- Integrated Shadow AI* awareness into our AI training programs, educating employees on security risks, policy compliance, and the availability of enterprise-grade alternatives.

5. We maintain strict control over our AI expenditures

We know

The AI market has shifted to consumption-based billing. Autonomous agentic* workflows consume 5 to 30 times more tokens* than standard queries, creating "token drift*" and silent exponential cost spikes. Without real-time visibility and controls, AI costs can spiral unpredictably.

We act

We have:

- Established a clear Group-wide rule to manage token usage, setting spending caps, user quotas, and technical guardrails to prevent AI assistants from running out of control and generating unexpected costs;
- Deployed built-in security features directly within our AI platforms (including VeoliaSecureGPT) to automatically enforce our financial rules without manual intervention;
- Implemented a global real-time tracking dashboard that provides full visibility into digital resource consumption, making it easy to spot anomalies and spending across all company levels;
- Set up an automated cost-allocation system to accurately assign AI expenses to specific projects and teams, ensuring full financial transparency;
- Configured automated alerts to notify teams as consumption approaches established limits, preventing unexpected charges.

6. We secure our critical operations without ever delegating the final decision to AI

We know

Cyberattacks, physical disruptions, or cloud outages on critical water, energy, and waste infrastructure can threaten public health and service continuity. Algorithmic hallucinations* in operational SCADA* systems could also pose severe physical risks.

We act

We have:

- Implemented dynamic multi-region cloud replication with proven failover resilience, enabling instant rerouting across geographies during infrastructure disruptions.
- Instituted mandatory human validation authority (Human-in-the-Loop*) for all consequential operational decisions, ensuring AI remains a decision-support tool rather than an autonomous decision-maker in critical activities.

AI in writing this Manifesto

This Manifesto is the result of collective work across Veolia's teams.

It is a living framework that will evolve with our practices, our feedback, and the development of technology.

Several generative AI tools were used in co-reflection mode to organize the document, challenge and rewrite some sections, check overall consistency, and correct spelling and typography.

All content was reviewed and validated by the teams involved: **humans stayed in control until the end.**

Words with an asterisk (*) are defined in the glossary on page 19

Lexique

Agent, Agentic AI	AI system capable of chaining actions and executing complex tasks autonomously.
Prompt / Prompting	Formulation of precise instructions, guidelines, or requests designed to guide and optimize the response produced by generative AI.
Machine Learning	AI method allowing a system to learn from data to perform a task or generate predictions.
Token / Token drift	Elementary unit of information processed by an AI. Drift refers to an uncontrolled overconsumption of these units during automated exchanges, causing energy waste and extra costs.
Hallucination	Error or factually incorrect response generated by a language model with an appearance of logical certainty.
LLM / SLM	Language models (Large or Small Language Models). LLMs are versatile but very compute-intensive, while SLMs are lighter, specialized, and resource-efficient.
Shadow AI	Unregulated use of external AI tools on the corporate network without IT security validation.
Agnostic architecture	System designed to work interchangeably with different tech vendors or AI models without exclusive technical dependency.
SCADA	Acronym for "Supervisory Control and Data Acquisition": industrial control computer systems providing real-time management of plants and critical infrastructure (water, energy, waste).
Waste heat	Thermal energy inevitably released by equipment (especially data center servers) that can be recovered to heat district networks.
Human-in-the-Loop	Principle ensuring that no sensitive industrial or operational decision is delegated to AI without human review and approval.