Identifying and Engaging with Stakeholders for an AI (AAA) System Project

Exec Summary

This document provides comprehensive guidance for identifying and engaging stakeholders in AI system projects, moving beyond traditional shareholder-focused approaches to embrace a holistic stakeholder framework that addresses the critical gap between technical AI development and broader societal impacts. The methodology establishes ForHumanity's definition of stakeholders as encompassing both direct stakeholders such as internal employees, customers, and regulatory bodies, and indirect stakeholders including society, non-profit organizations, and the environment. This comprehensive approach recognizes that AI systems create ripple effects far beyond their immediate operational context, tracing the evolution from 1980s shareholder primacy toward broader stakeholder accountability as corporations increasingly recognize they cannot operate in isolation from their environmental and social impacts.

The document introduces a structured "double diamond" methodology adapted for AI projects, featuring four sequential phases that move from initial stakeholder discovery through detailed analysis, systematic categorization, and final integration into AI system governance. This process emphasizes treating any internally created stakeholder list as inherently incomplete and "always open" to expansion, while gathering comprehensive information about each stakeholder's role, interests, influence, and stance to enable strategic prioritization. The

methodology employs three primary visual frameworks for stakeholder mapping: the Onion Framework organizing stakeholders in concentric rings from development teams to society-level impacts, Graph Network Diagrams providing computer-readable representations of complex relationships, and Stakeholder Matrices enabling analysis based on paired characteristics such as power-interest and knowledge-support dynamics.

The framework identifies five critical stakeholder categories ranging from Team Level personnel who shape day-to-day development decisions, through Organization Level executives and investors, to Ecosystem Level customers and regulators, Directly Affected Consumers who experience impacts but may lack design influence, and Society Level stakeholders including government agencies and marginalized communities who face broad societal consequences. Particular attention is dedicated to environmental stakeholders, recognizing both substantial negative impacts from AI energy consumption projected to reach 0.5% of global electricity generation by 2027, and potential positive contributions to conservation, climate change mitigation, and ecosystem monitoring through optimized resource management and environmental data analysis.

The methodology emphasizes three core implementation principles of inclusivity, people-centered design, and iterative refinement, while stressing the importance of statistical validity in stakeholder representation and continuous education for technical stakeholders in AI ethics and security. The framework reveals potential friction points between technically-oriented AI companies and non-technical stakeholders in government agencies and non-profit organizations, where disconnects in technical expertise can lead to inappropriate deployment or inadequate risk management. This analysis suggests that effective AI governance requires stronger board independence and a return to broader organizational accountability, moving away from narrow profit maximization toward comprehensive stakeholder consideration.

This stakeholder identification and engagement framework provides essential infrastructure for responsible AI development by systematically mapping stakeholder relationships, power dynamics, and interests to help organizations better anticipate and address the wide-ranging impacts of AI systems. The methodology supports the development of AI systems that are not only technically sound but also ethically grounded and socially beneficial, with practical applications ranging from AI startups to large enterprises and government agencies. As AI systems become increasingly integrated into society, this comprehensive approach to stakeholder engagement becomes not just advisable but essential for sustainable and responsible AI development that serves the interests of all affected parties while maintaining accountability for environmental and social impacts.

For a complete set of practical guidance materials, please visit our library.

Introduction - Who is a Stakeholder?

In this document, "Identifying and Engaging with Stakeholders for an AAA (AI) System Project," we explore the concept of a stakeholder in the context of AI project management. We will start

with a review of the concept of the stakeholder; the exact definition of the stakeholder according to ForHumanity is given in Section 1 below.

In public services, in Westminster systems, where PBMOK (Project Management Body of Knowledge, see Appendix V) derives from, the term is used routinely to describe any representative, category, group or natural person with a recognisable relevant interest in a decision or policy proposal. This neutral definition is one that should be relied on.

However, the term "stakeholder" is originally American and specifically derived from another word - "shareholder". - meaning its roots and more complex influence and origin story derive from private sector governance. Specifically, it arose within corporations and during the Reaganomics era of the 1980s. This term has historically been contentious, perhaps unnecessarily so, with two primary schools of thought shaping its understanding. One that views shareholders as a class or category of stakeholders and another which rejects that stakeholders are relevant for corporations to take into account. The latter approach has transferred to other business structures because of its ability to rationalise decisions in favor of maximizing profits, to the exclusion of all other factors.

This perspective views stakeholders as an "externality", considering the environment and other non-market transactional elements as outside the purview of corporate social responsibility and decision making.

It holds that markets are neither intended nor equipped to effectively manage such externalities. Indeed, this view goes beyond such a “stand-off” perspective, actively advocating that they should not even attempt to manage such externalities. Conversely, the alternative perspective of shareholders as a class of stakeholder advocates for a broader accountability of corporations. It argues that corporate operations, profit, and cash flow do not exist in isolation. Corporations should be accountable for the full cost of their operations, including their impacts on the commons or commonwealth. This approach emphasizes the need for corporations to avoid causing unnecessary harm for short-term profit, considering the potential for long-term liability, as well as personal culpability of those who declined to robustly consider the risks being negligently, recklessly or intentionally taken, and the opportunity cost of innovation due to narrow, short-term focus at the top management and oversight level.

While the detached view on externalities has been widely popular until recently, the second view is becoming increasingly prominent. For example, Angus Deaton, the British-American economist and Nobel Prize winner in Economic Sciences in 2015, has recently expressed a critical viewpoint towards mainstream economics in the International Monetary Fund's (IMF) Finance and Development magazine. Deaton argues that mainstream economics is in disarray because it overlooks the impact of power dynamics on economic outcomes, neglects questions of equity, and often results in policy recommendations that exacerbate wealth inequality. He highlights the profession's failure to address the real-world issues facing people, particularly in the United States, where economic growth no longer benefits the broader population as it once did. Deaton points to the rising "deaths of despair" and the growing income inequality as symptoms of these failures.

Deaton suggests that the economics profession needs to broaden its focus beyond monetary measures of welfare to consider other aspects of human well-being, such as job satisfaction,

family, and community. He proposes that economists pay more attention to pre-distribution—how income is distributed by the market before taxes and transfers—and to policies that could prevent economic distress. This shift would require economists to engage with more complex and politically sensitive areas like union promotion, immigration control, and industrial policy.

While this view is growing in prominence, there has been pushback from a number of sectors - not just other economists, but also corporate structures and organizations, such as the American Business Roundtable. The party most inhibited by the shift to shareholder primacy has been the board of directors. Whereas previously they acted as a balance and check on the operations, their independence has been actively eroded by the doctrine of shareholder (rather than stakeholder) primacy. This doctrine has resulted in the effective replacement of a broad duty at law "to act in the best interests of the organisation" with a narrow one that has redefined their role, taking it to be tantamount to exclusively focusing on the share price. However, the application of stakeholder theory extends beyond the private sector. It is more prevalent and clearly understood in not-for-profit and public sector organizations. Here, "stakeholder" aptly captures the idea of parties with legitimate and relevant interests, who should be consulted and considered in relation to policies, public benefits, or other matters. This consultation aims to enhance the robustness of analysis and decision-making.

Moreover, stakeholder analysis has practical applications, such as in ensuring statistical validity, and driving inclusive and accessible methodologies such as co-design. These methodologies may be implemented to align with indigenous cultural norms and practices, allowing for culturally sensitive inputs in not-for-profit and public decisions. It also supports the cultural norms and practices of widely varied groups, including ethnic, religious, national, gender/gender identity based, sexual orientation based, age based, class based, professional, and for those who have different types of disabilities and/or need for accommodations. Additionally, the advent of technology-enabled inputs presents a new frontier in stakeholder engagement, though this area requires more consistent and effective implementation, embracing a 'no wrong door' approach.

We have considered these varied perspectives and applications of stakeholder theory, particularly in the realm of AI systems, where such considerations are increasingly critical.

Board Independence and AI Governance

The shift toward shareholder primacy over the past four decades has fundamentally altered corporate governance structures in ways that create particular vulnerabilities for AI oversight and regulation. Historically, boards of directors served as independent checks on executive management, with fiduciary duties broadly defined as acting "in the best interests of the organisation." This broad mandate traditionally required directors to consider the long-term sustainability of the enterprise, including its relationships with employees, customers, communities, and other stakeholders whose interests might diverge from short-term profit maximization.

However, the doctrine of shareholder primacy has systematically eroded this independence by redefining the board's role in narrow terms focused exclusively on share price performance. This redefinition has created structural conflicts of interest where board members, often selected by and beholden to major shareholders or executive management, find their oversight responsibilities compromised by pressure to prioritize immediate financial returns over comprehensive risk assessment. In practice, this has led to boards that function more as advisors to management than as independent oversight bodies, undermining their capacity to challenge strategic decisions or demand thorough consideration of stakeholder impacts.

For AI systems, this erosion of board independence creates particularly acute governance challenges. AI technologies present complex, often poorly understood risks that extend far beyond traditional business considerations to encompass algorithmic bias, privacy violations, safety hazards, environmental impacts, and broader societal consequences. Effective oversight of these risks requires board members who can independently evaluate technical assessments, challenge management assumptions, and insist on comprehensive stakeholder analysis even when such scrutiny might delay product launches or increase development costs. Boards compromised by shareholder primacy are structurally ill-equipped to provide this level of independent oversight.

The technical complexity of AI systems exacerbates these governance challenges by creating information asymmetries between management and boards. When board members lack technical expertise and are primarily focused on financial performance metrics, they become heavily dependent on management representations about AI system capabilities, limitations, and risks. This dependency is particularly problematic given that management teams, under pressure to deliver rapid growth and market advantages, may have institutional incentives to downplay risks or oversimplify complex technical trade-offs. Independent boards with diverse expertise and genuine authority to challenge management decisions are essential for breaking through these information asymmetries and ensuring thorough risk assessment.

Furthermore, the stakeholder-inclusive approach to AI governance outlined in this document requires board-level commitment to considering diverse perspectives and long-term consequences that may not align with short-term shareholder interests. Boards constrained by shareholder primacy doctrine lack both the mandate and practical independence necessary to enforce comprehensive stakeholder engagement processes or to make decisions that prioritize societal benefit over immediate profit maximization. For AI regulations and governance frameworks to be fully operative, corporate governance structures must evolve to restore genuine board independence, with fiduciary duties that explicitly encompass stakeholder considerations and long-term organizational sustainability rather than narrow focus on share price performance.

This restoration of board independence is not merely a matter of good corporate governance; it is a prerequisite for effective AI regulation and responsible innovation. Regulatory frameworks that rely on corporate self-governance and compliance will inevitably fail if the internal governance structures they depend upon are compromised by conflicts of interest and narrow mandates. As AI systems become increasingly consequential for society, the alignment of corporate governance structures with broader stakeholder interests becomes not just advisable

but essential for maintaining public trust and ensuring that technological advancement serves human flourishing rather than merely financial extraction.

Section 1: ForHumanity Definition of a Stakeholder

> DEFINE STAKEHOLDERS - Ensuring that stakeholders are considered from a holistic perspective of impacted groups, including both direct stakeholders that may be internal (e.g., employees, customers, shareholders, AI ethics officers, legal compliance officers, data protection officers and other personnel in charge of compliance) or external (e.g., human users, communities, regulatory bodies) and indirect stakeholders (e.g., society, non-profit organizations and the environment)

> *"Simple ideas create complex questions, and we proceed as follows. In the next section we examine why the dominant story or model of business that is deeply embedded in our culture is… resistant to change, not consistent with the law, and for the most part, simply ignores matters of ethics. Each of these flaws is fatal in the business world of the twenty-first (and current) century"* (Freeman, R. Edward. Stakeholder Theory: The State of the Art. Cambridge University Press, 2010, p. 5)

> *"Every business creates, and sometimes destroys, value for customers, suppliers, employees, communities and financiers. The idea that business is about maximizing profits for shareholders is outdated and doesn't work very well, as the recent global financial crisis has taught us. The 21st Century is one of "Managing for Stakeholders." The task of executives is to create as much value as possible for stakeholders without resorting to tradeoffs. Great companies endure because they manage to get stakeholder interests aligned in the same direction."* (Freeman, R. Edward. "Managing for Stakeholders." Effective Executive, vol. 12, no. 5, 2009, pp. 18–23)

> Put simply, "...*using "stakeholder" as a basic unit of analysis makes it more difficult to ignore matters of ethics"* in top management and oversight deliberations about what constitutes the best interests of the organisation (quote Ibid, page 56).

This suggestion aligns with evolving views in stakeholder theory, where the generational focus on shareholder primacy is gradually shifting towards broader sustainability goals, including ESG and the United Nations Sustainable Development Goals (UN SDGs). This shift is influenced by growing concerns about climate change and international regulations, as well as the need for compliance with emerging regulations on AI and corporate responsibility. There is a developing

consensus in stakeholder literature regarding the inclusion of entities without agency, such as the natural and social/urban environment. Recognizing these entities as stakeholders reflects a more comprehensive approach to understanding the impact of corporate actions, balancing profit maximization with broader societal and environmental responsibilities.

2. Finding and Defining Stakeholders: A Systematic Process

2.2 Visual Models for Stakeholder Mapping

Visual diagrams are helpful tools for stakeholder mapping. They transform complex relational data into intuitive spatial representations. Diagrams such as onion models, network graphs, and matrices enable project teams to visualize stakeholder positions, relationships, and relative influence at a glance, revealing patterns and power dynamics that might remain obscured in text-based descriptions.

These visual frameworks facilitate more strategic stakeholder engagement by making clear which groups require different levels of attention based on their proximity to the project, degree of influence, or level of interest. Moreover, well-designed stakeholder diagrams create a shared mental model among team members, ensuring alignment on who matters, why they matter, and how they should be engaged throughout an AI project's lifecycle. When regularly updated, these visual maps also help track shifting stakeholder positions and evolving relationships, providing a dynamic rather than static understanding of the stakeholder landscape.

2.2.1 The Onion Framework: Concentric Layers of Influence

When describing stakeholder roles and power dynamics to humans, an onion type diagram may be used as shown below:

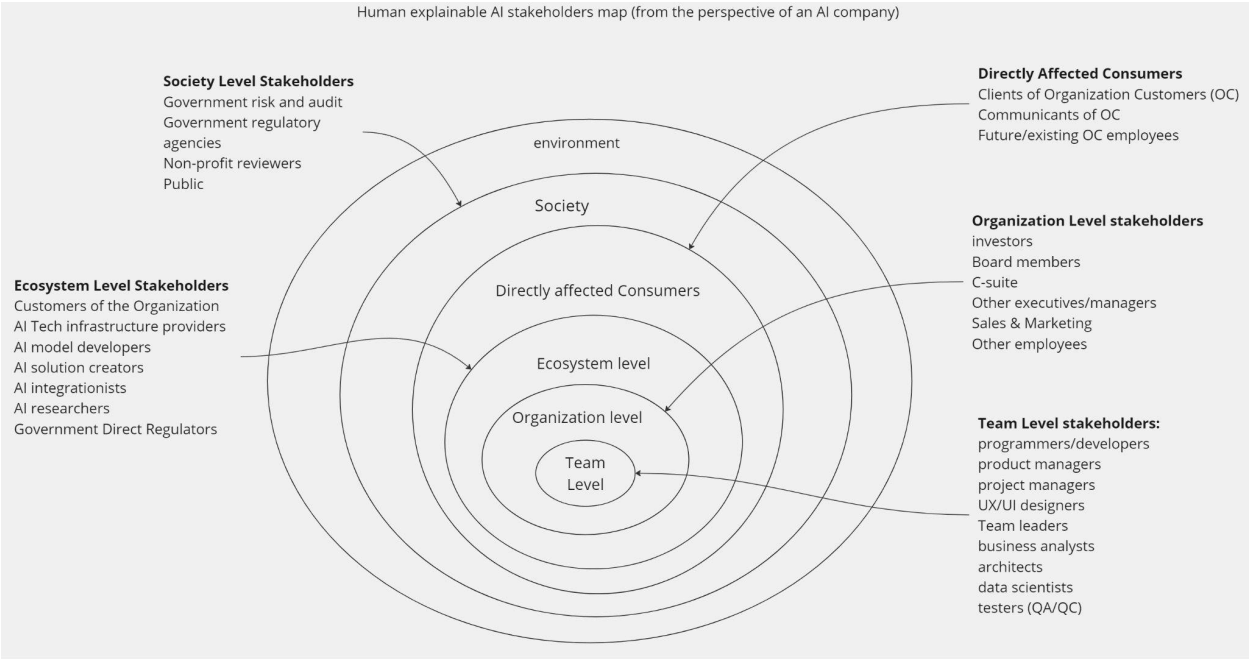

The diagram presents a layered, concentric view of an AI company's stakeholders, with the company at the center and successive rings representing different categories of stakeholders with varying degrees of proximity and influence. This type of diagram, also referred to as an onion diagram, shows how stakeholders who are still affected by the AI, may be much farther from the center of influencing the design and construction of the AI.

The innermost ring, the "Team Level", includes those closest to the AI development process, such as programmers/developers, product managers, project managers, UX/UI designers, team leaders, business analysts, architects, data scientists, and testers (QA/QC). Team members are expected to have the most direct, day-to-day influence on how the AI is designed and implemented.

The next ring, the "Organization Level", encompasses a wider set of internal stakeholders like C-suite executives, other managers, sales & marketing teams, and other employees. These stakeholders may set direction and vision, and/or give feedback, but are not as involved in the day-to-day activities of AI construction.

Moving outward, the "Ecosystem Level" includes stakeholders who are more external but still closely linked to the company, such as the organization's customers, AI tech infrastructure providers, AI model developers, AI solution creators, AI integrationists, AI researchers, and government direct regulators. Ecosystem stakeholders may have hard power (such as government direct regulators of AI and the organization's customers), or more indirect soft power (such as AI researchers and tech infrastructure providers).

The "Directly Affected Consumers" ring represents clients of the organization's customers (OC), communicants of OC, and future/existing OC employees. These stakeholders are directly impacted by the AI system but are one step removed from the company's ecosystem – and are two steps removed from the company itself.

Finally, the outermost "Society Level Stakeholders" ring includes entities with a broader interest in the AI's implications, such as government risk and audit agencies, non-profit reviewers, and the general public. It also includes the environment, recognizing the AI's potential ecological impact.

This diagram effectively illustrates how the level of direct involvement and influence on the AI company diminishes as we move from the central team to the wider societal context. However, it also acknowledges that the AI's effects extend far beyond the company's immediate vicinity, ultimately impacting a diverse array of stakeholders and the environment we all share. The radial structure emphasizes the ripple effects of AI development and the importance of considering the full scope of its influence.

Section 2: The Process of Finding and Defining Stakeholders

1. **Understanding Stakeholder Roles**: Stakeholders in an AI project can include a wide range of individuals, groups, or organizations that can affect or are affected by the project. This includes internal stakeholders like employees, managers, and owners, and external stakeholders like investors, suppliers, policymakers, and customers.

The diagrams shown below illustrate various types of stakeholders and relationships between them. Stakeholders are important to AI projects for many reasons; as separate entities/individuals, but also according to their power dynamics and overall workflows. It is important to continuously reassess such diagrams throughout the AI lifecycle, as the lifecycle may result in alterations, such that the stakeholders might change.

Graph Network Diagrams

A computer needs a computer-readable diagram, such as a graph network diagram, a basic example of which is shown below:

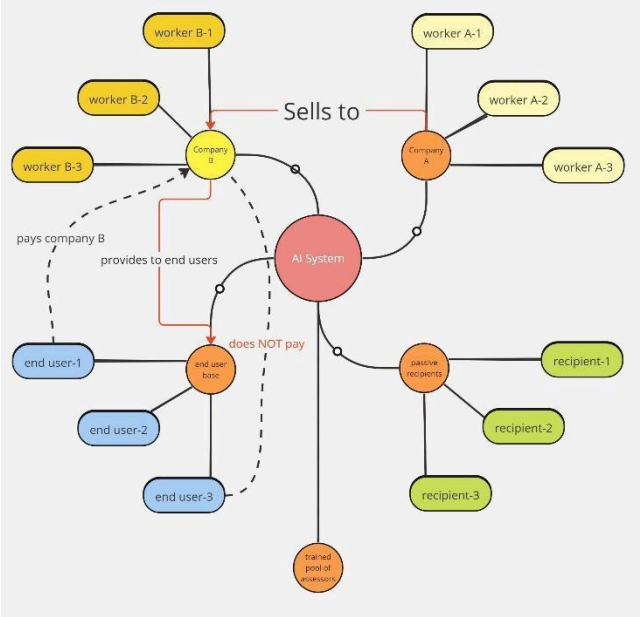

The diagram is a sociogram representing the stakeholders and their relationships in an AI company's ecosystem. The sociogram shows the interactions between workers, the company, the AI system, end users, and recipients. It also allows for including different types of open-source software and customizing the stakeholder relationships and directions.

In this example, the company can create AI directly or through contractors and/or other resellers. For example, there's a bidirectional relationship where Company A creates AI and sells it to Company B, a government contractor, which sets up the AI for government use.

The company sells the AI system to clients. End users 1-3 may purchase the AI directly from the company, or indirectly from another company that has purchased the AI from the original providing company.

The diagram distinguishes between end users who directly interact with the AI and passive recipients who are affected by it. Recipients can be affected by industry (secondary market) or by the energy and other indirect effects of the AI system.

Additional details that would be required include the environmental effects caused by the industry around the AI system, as well as other stakeholders (direct and indirect; not shown - see below for a detailed description).

Overall, the sociogram is an example of the complex web of stakeholders and their relationships in an AI company's ecosystem, considering not just direct users but also passive recipients and environmental impacts (see below for such a diagram).

Stakeholder Matrices

Stakeholder matrices are diagrams used to visualize and map stakeholders based on specific characteristics. Stakeholder matrix types are binary, so only two characteristics are considered and compared. The matrix typically has four quadrants, with the X-axis and Y-axis labeled with different characteristics (e.g., level of interest, impact, influence, support, etc.). The main benefits of using a stakeholder matrix include categorizing stakeholders for tailored communication and engagement, identifying potential risks, resistance, or opportunities, setting priorities for managing or engaging with stakeholders, communicating strategies to colleagues and leadership, and tracking progress over time.

There are three common stakeholder matrix frameworks: The Power-Interest Grid, The Power-Predictability Matrix, and The Stakeholder Knowledge Base Diagram. Each framework maps stakeholders based on two specific characteristics, and the four quadrants in each matrix represent different combinations of those characteristics. For example, in the Power-Interest Grid, stakeholders are mapped based on their level of power/influence and level of interest, with the four quadrants being High Power/High Interest, High Power/Low Interest, Low Power/High Interest, and Low Power/Low Interest.

An example Power-Interest stakeholder matrix for an AI project is shown below:

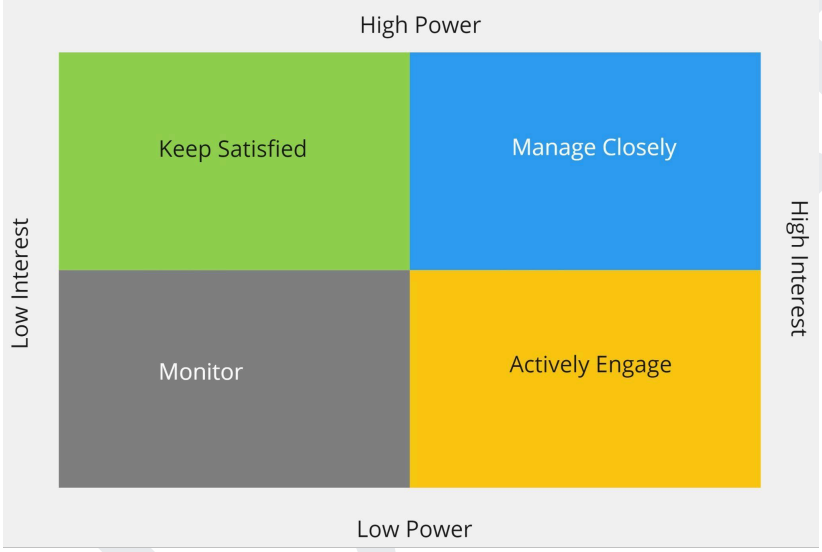

Quadrants of this diagram

1. High Power/High Interest (Top Right Quadrant; Blue; “Manage Closely”)

Description: Stakeholders in this quadrant have both significant influence over the project and a high level of interest in its outcomes. They are crucial for the project's success and should be closely managed and engaged regularly.

Example Stakeholders: Senior Executives, Key Investors, Regulatory Bodies (if the AI affects compliance issues)

2. High Power/Low Interest (Top Left Quadrant; Green; “Keep Satisfied”)

Description: These stakeholders have significant power to impact the project but have low interest or less concern about the project’s outcomes. It is important to keep these stakeholders satisfied to avoid potential obstacles.

Example Stakeholders: High-Level Directors not directly involved, Government bodies not actively regulating

3. Low Power/High Interest (Bottom Right Quadrant; Orange; “Actively Engage”)

Description: Stakeholders here have less power or influence but a high interest in the project. They can be important as supporters or blockers based on how their interests are managed. Traditional “shareholder oriented” design and implementation processes have often neglected this group, as in such processes, lack of power is seen as making them unimportant. Typical stakeholder matrices may refer to such stakeholders as a group to keep informed only. Significant clashes and friction are more likely to occur between Group 1 and this group.

Example Stakeholders: End Users, Academic Communities, Non-Governmental Organizations (NGOs)

4. Low Power/Low Interest (Bottom Left Quadrant; Gray; “Monitor”)

Description: These stakeholders have minimal power and interest relative to the project. They require less active management but should be monitored for any changes in their position.

Example Stakeholders: General Public, Media (if not focused on the tech industry)

Some additional matrix frameworks include: Power-Support Stakeholder Analysis Matrix, Support-Importance Stakeholder Matrix and Influence-Interest Matrix.

To use a stakeholder matrix effectively, one should determine relevant characteristics of the stakeholders, determine characteristic levels, and plot stakeholders on the matrix according to their characteristics. As only one pair of characteristics can be considered at once, to maintain the quadrant format, multiple such matrices may be needed to fully capture stakeholder characteristics.

It is possible to create a stakeholder matrix according to the following process:

1. Define the purpose of the stakeholder matrix. Understand what information you want to capture and communicate. In particular, consider all of the stakeholder characteristics that you may wish to include (note above examples).

2. Brainstorm and identify all potential stakeholders, both internal (e.g. employees, managers) and external (e.g. customers, suppliers, regulators). Be comprehensive in your list.

3. Identify each stakeholder's key interests, goals, and concerns. Understand their motivations and potential conflicts between stakeholders. Turn this information into stakeholder characteristics.

4. Determine the level of each characteristic for each stakeholder, which may include knowledge, power, interest, involvement, influence, support, predictability, and so forth. Categorize them as having high, medium or low of each of these characteristics. This can be visualized using a grid or network diagram.

5. Create one or more stakeholder matrices according to each selected pair of characteristics.

6. Develop an engagement plan for how you will communicate with and involve each stakeholder group according to these matrices. Use a variety of methods like surveys, focus groups, workshops, and online platforms.

7. Continuously monitor and update the stakeholder matrices as the project or organization evolves. Stakeholder relationships and priorities can change over time.

The goal is to create a comprehensive, visual representation of the stakeholder landscape that informs your strategy for effectively engaging and managing stakeholders throughout a project or initiative, using stakeholder matrices as one tool.

For more information, see Appendix VI - stakeholder matrices.

For all of these diagrams, the exact lists are less important than the conceptual roles and their relationship to each other, including with regard to power dynamics. The exact lists are considered below.

2. **Creating Comprehensive Stakeholder Lists**: We need to create lists of all potential stakeholders, and then sort them into the various categories described in the visual models.

It is fair to say that every organisational project will have different stakeholders and that organisations who are unused to having to identify and engage stakeholders in their work would benefit from looking at how the practice has evolved in the public and not for profit space.

Statistical validity is important to at least ensure no underrepresentation of a group (for example, if a specific group is 1% of the stakeholders, then its members should be 1% of the identified stakeholders plus 1% of the diverse group that engages with the AI project). However, processes like co-design, that focus on equity (and not just inclusion), the percentage of members in a particular group may need to be higher to support the process.

Overall, the process is expected to feature a “double diamond” type process, adapted to the requirements of these AI guidelines. This process is a central tenet of human centred design. It therefore is already connected to private sector product development and focus group practices that are common in industry development, and so could be leveraged by groups who are new to these AI guidelines.

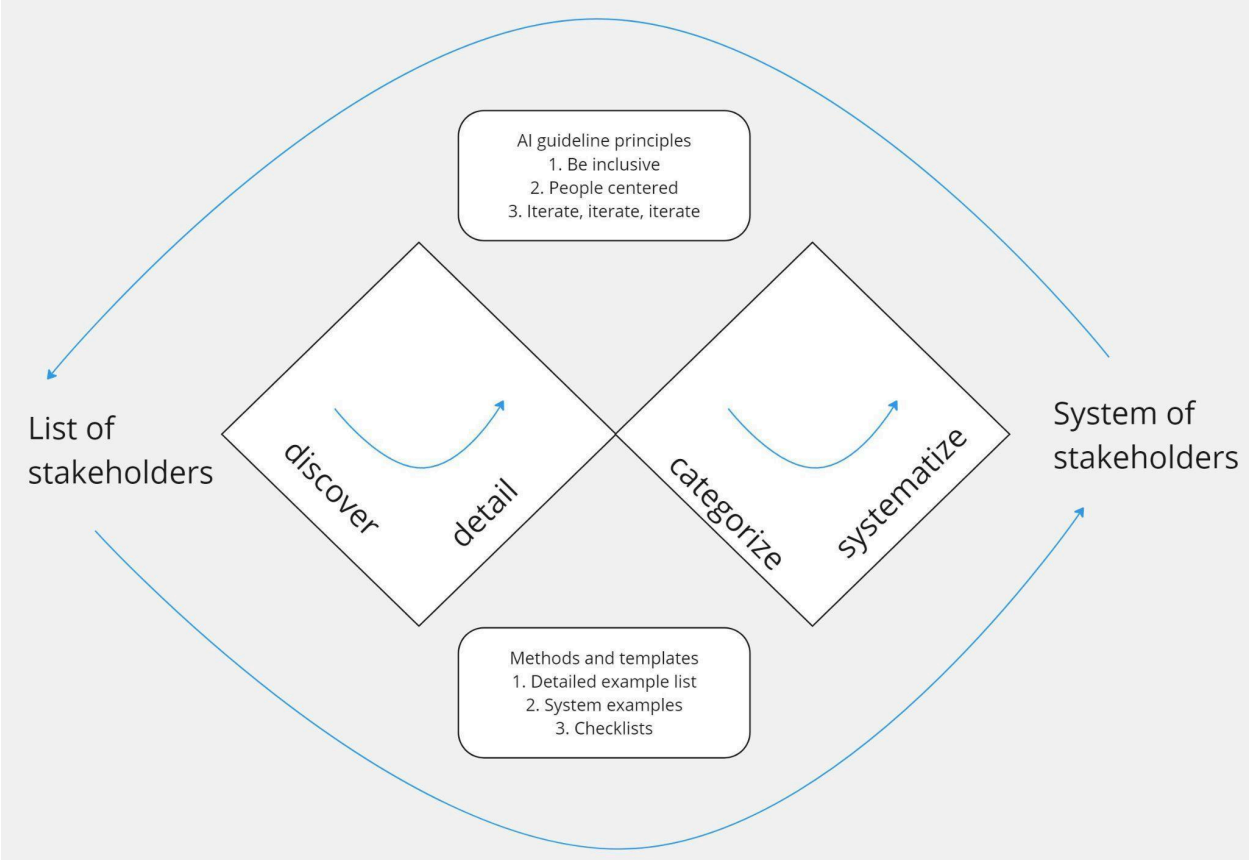

As the above diagram shows, we begin by creating lists of stakeholders, by various main groups.

An example list is given in Appendix III below.

Double diamond thinking is a useful approach to answering the questions:

1. Who has a relevant, recognised stake?
2. What is that stake?
3. In what category do they fall?

Because a stake can be self-identified, and our process is always biased (toward what Donald Rumsfeld termed ‘known knowns’), it’s important to the process and its purpose that any internally created list be viewed as incomplete. Alternatively, any such list may be treated as ‘always open’ to ensure that the point of stakeholder management - which is engagement and robust consensus between diverse perspectives as to the best way to proceed toward an objective.

The double diamond design process is a framework used in design thinking to structure the creative process from initial idea to finished product. It consists of four main stages: Discover, Define, Develop, and Deliver. This process can be adapted and applied effectively to stakeholder engagement for AI guideline principles.

In the context of stakeholder engagement, the first "Discover" stage involves identifying and understanding the relevant stakeholders. This includes listing out all potential stakeholders, such as AI developers, policymakers, industry leaders, civil society organizations, and affected communities. Methods like stakeholder mapping can help provide a comprehensive view.

The second "Detail" (rather than “Define”) stage focuses on narrowing down and prioritizing the key stakeholders. This involves gathering detailed information on each stakeholder's role, interests, influence and stance on the AI guidelines. Categorizing stakeholders based on these factors helps define who is the most essential to engage. Ideally, all stakeholders would be equally engaged in the process at the same time. However, given the realities of limited time and resources, the Detail stage can help with stakeholder engagement prioritization.

Transitioning into the “Categorize” (rather than "Develop”) stage, the engagement process becomes more systematic and targeted. With a prioritized list of key stakeholders, tailored engagement strategies can be developed. This could include methods like one-on-one interviews, focus groups, surveys, or participatory workshops depending on the stakeholder. Categorization can also help with further stakeholder engagement prioritization – as well as the potential need to obtain more resources, for effective engagement of all necessary stakeholders.

Finally, the “Systematize” (rather than "Deliver") stage is about synthesizing the insights gathered from stakeholder engagement and integrating them into the development and implementation of the AI guideline principles. This ensures the guidelines are inclusive, people-centered, and responsive to stakeholder needs and concerns.

Importantly, the double diamond process is not purely linear but allows for iteration and feedback loops. Engaging stakeholders should be a continuous process, with insights from the

"Categorize" and "Systematize" stages looping back to inform and refine the understanding of stakeholders and engagement strategies. The arrows form a cyclical flow, showing that this is an iterative process of identifying stakeholders, providing details on them, categorizing them, and integrating them into a stakeholder system or analysis.

The "AI guideline principles" includes three basic principles: "1. Be inclusive, 2. People centered, 3. Iterate, iterate, iterate". These principles provide a clear and concise set of values to guide the entire stakeholder engagement process. The first principle, "Be inclusive," emphasizes the importance of considering and involving a diverse range of stakeholders. This includes not just those directly involved in AI development, but also those who may be impacted by the technology, such as end-users, communities, and underrepresented groups. Inclusivity ensures that the AI guidelines reflect a broad range of perspectives and needs.

The second principle, "People-centered," puts human well-being and interests at the forefront. It recognizes that AI should be developed and used in service of people, not the other way around. This principle guides stakeholder engagement to prioritize understanding and addressing the needs, concerns, and aspirations of the people who will be affected by AI systems.

The third principle, "Iterate, iterate, iterate," acknowledges that developing effective AI guidelines is an ongoing process. It requires continuous learning, refinement, and adaptation as new insights, challenges, and stakeholders emerge. This principle encourages a mindset of experimentation, feedback, and improvement throughout the stakeholder engagement process.

By keeping these principles central, the framework ensures that stakeholder engagement is not just an exercise in collecting input, but a values-driven process aimed at creating AI guidelines that are truly inclusive, people-centered, and responsive to evolving needs and contexts. The principles serve as a constant reminder of the ultimate goals and help align the various stages and methods of stakeholder engagement towards these ends.

Moreover, having explicit AI guideline principles helps communicate the purpose and values behind the stakeholder engagement process to both internal teams and external stakeholders. It provides a clear rationale for why stakeholder engagement is critical and how it contributes to the responsible development and governance of AI.

Overall, the "AI guideline principles" are the compass that orients the entire stakeholder engagement journey. They ensure that the process stays true to its inclusive, people-centered, and iterative values, ultimately leading to AI guidelines that are more ethical, beneficial, and widely accepted.

The "Methods and templates" box in the diagram provides practical resources to support each stage of the stakeholder engagement process. The "Detailed example list" offers concrete illustrations of how to identify, categorize, and prioritize stakeholders. This could include case studies or best practices from organizations that have successfully engaged stakeholders in developing AI guidelines.

The "System examples" go a step further, demonstrating how to systematically analyze and map stakeholder relationships, influence, and potential impact on the AI guidelines. These examples might showcase tools like stakeholder matrices, power-interest grids, or network maps to visualize the stakeholder landscape.

Lastly, the "Checklists" serve as quick reference guides to ensure thoroughness and consistency in the engagement process. These could include checklists for stakeholder identification criteria, information to gather about each stakeholder, engagement methods to consider, or steps to integrate stakeholder feedback.

These practical resources are provided as part of these guidelines, thereby enabling the framework to be more actionable and easier to implement. The methods and templates can be adapted to fit the specific context and needs of the organization developing the AI guidelines. They also promote a more systematic and comprehensive approach to stakeholder engagement, reducing the risk of overlooking important stakeholders or engagement opportunities.

Moreover, having a library of methods and templates builds internal capacity and knowledge sharing around stakeholder engagement. As teams iterate through the double diamond process across different AI guideline initiatives, they can also continually refine and expand these resources based on their learnings and successes.

Thus, the "Methods and templates" component of the framework offers valuable tools to operationalize the stakeholder engagement process. It equips teams with concrete examples, systematic approaches, and practical checklists to effectively identify, analyze, and engage stakeholders in the development of AI guidelines that are inclusive, people-centered, and responsive to diverse needs and perspectives.

Overall, leveraging the double diamond process provides a structured yet flexible approach to stakeholder engagement. It enables a deep understanding of stakeholders, targeted engagement strategies, and the ability to iterate based on feedback. This results in AI guideline principles that are more robust, inclusive, and people-centered, ultimately leading to more responsible and beneficial AI development.

3. **Stakeholder Analysis**: We need to identify each stakeholder's goals, motivations, interests, influence, and interest level and establish relevance.

We also need to place each stakeholder within the above visual models, by creating a system view.

By describing each stakeholder in these terms, and then locating them within the stakeholder models, we will be able to create both a manual template and computer software for AI auditing of stakeholders.

When analyzing stakeholders, we need to ask the following questions. First, who has an interest internally or externally? Who definitely has such an interest - as opposed to someone who might have such an interest? Who could have such an interest - and who should have such an interest?

We also need to consider how to handle whose interests should definitely be recognized. For example: a rival firm has a relevant recognised interest in what you do, but are you open to collaboration or is this not in the firm's best interest? Coopetition is a better model in healthcare than competition and collaboration might work to produce economies of scale. However, in many countries, such as the US, this type of operation is not the norm - and might even be illegal, due to anti-cartel laws.

Below is an example of a system view of stakeholders from the perspective of a company creating an AI platform. This system view does not feature power relationships.

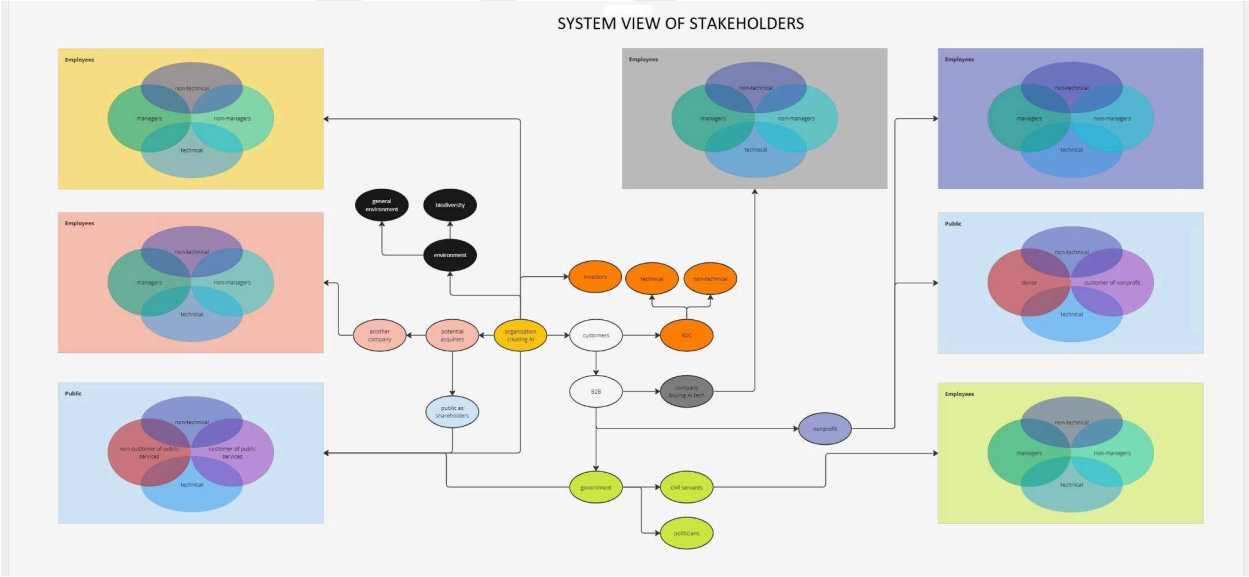

The organization creating the AI platform (the AI company) is shown in pale orange. The public is shown in light blue. The government is shown in bright green. The enterprise client is shown in gray. The non-profit organization is shown in purple. The environment is shown in black. A potential acquiring company (who could acquire the AI company) is shown in salmon.

We can further add hierarchies to the above diagram. In the below diagram, these hierarchies have been selected to account for a potentially concerning situation. In this situation, the company creating the AI platform (the AI company) is mainly driven by its technical employees, both managers and non-managers. However, all other groups of employees are shown as being within organizations - an acquiring company, a non-profit organization, the government as purchaser or an enterprise customer - in which the organization is mainly driven by its non-technical employees. We realize that this may not be the case, but such a situation is most likely to lead to problems, simply because of differences in expectations, language and communication between the AI company and all other organizations in this system. Therefore, we have attempted to mitigate such a situation in this document.

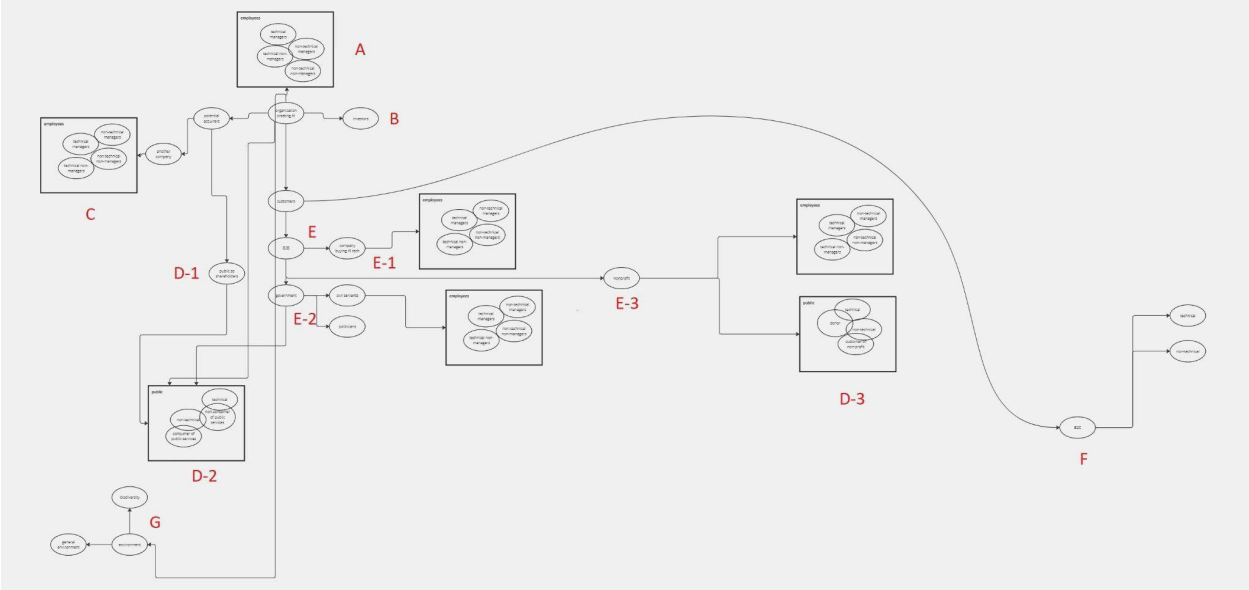

We will provide the above diagrams in a more readable manner, that can be zoomed into/out of. We are also working on manual templates that include checklists.

However, for now, we will focus on each part of the drawing. First, we will consider the company creating the AI software and its employees (labeled “A”; also referred to as the “AI company”); its investors (“B”); and potential acquirers of the company itself (“C”). For a startup or even a larger organization, typically such acquirers are other companies as shown; however, as described in greater detail below, acquisition through an IPO (Initial Public Offering) on the stock market is also possible.

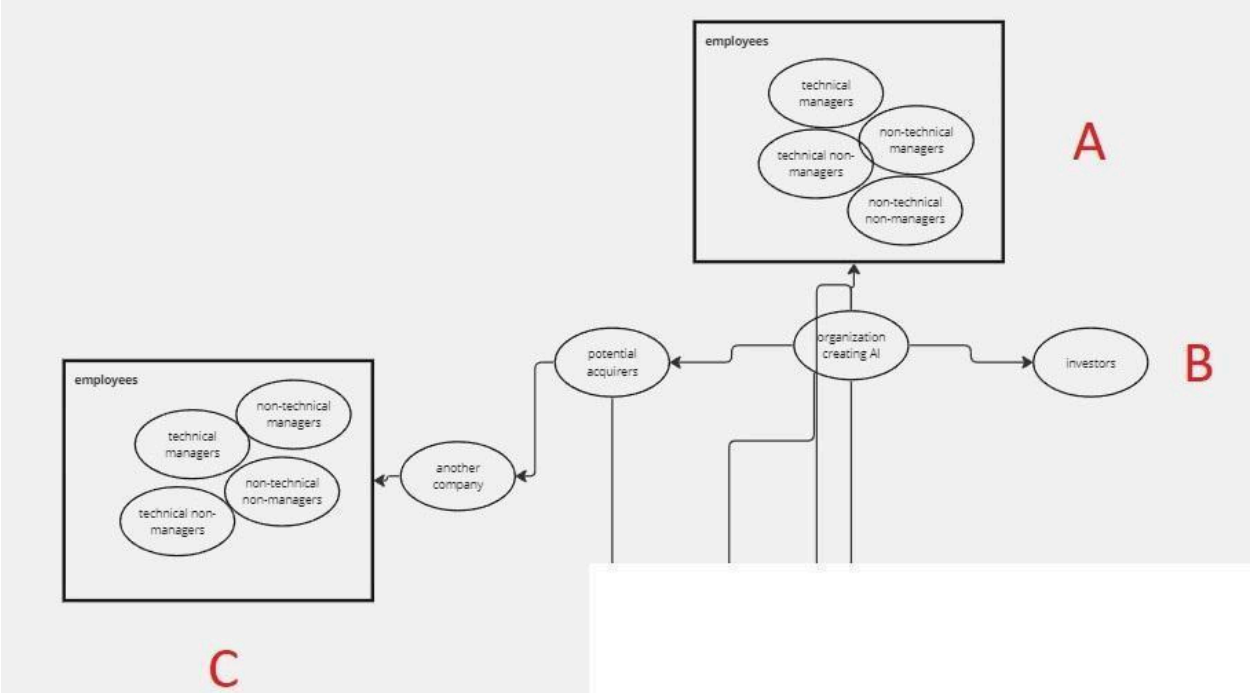

Section A depicts a typical hierarchy at AI companies for their employees. Usually, the top positions are held by technical managers. Technical non-managers may have more soft power than non-technical managers, even though they may be lower down on the org chart. Non-technical non-managers typically have the least soft power in such an organization.

Section B relates to investors, who are very important to young AI companies (startups) in particular – but are also important to older AI companies that are publicly traded on the stock market. Unless a company is privately held by a single person or by family members, or has some other situation that results in less pressure from investors, all AI companies do need to pay attention to their investors.

Section C relates to a potential acquirer. A privately held company – including a startup – can be acquired by another company as shown. This type of acquisition is currently most common for AI companies. Such a company also has employees. In this scenario, we have reversed the order of top management and those with soft power, so that non-technical managers and non-managers have the most power and/or soft power within this organization. We are not stating that this situation is typical; however, it may represent the most dangerous situation for an acquired AI company, as the acquiring company may have much less technical understanding of how the AI works and the potential dangers that it may represent, including with regard to stakeholders.

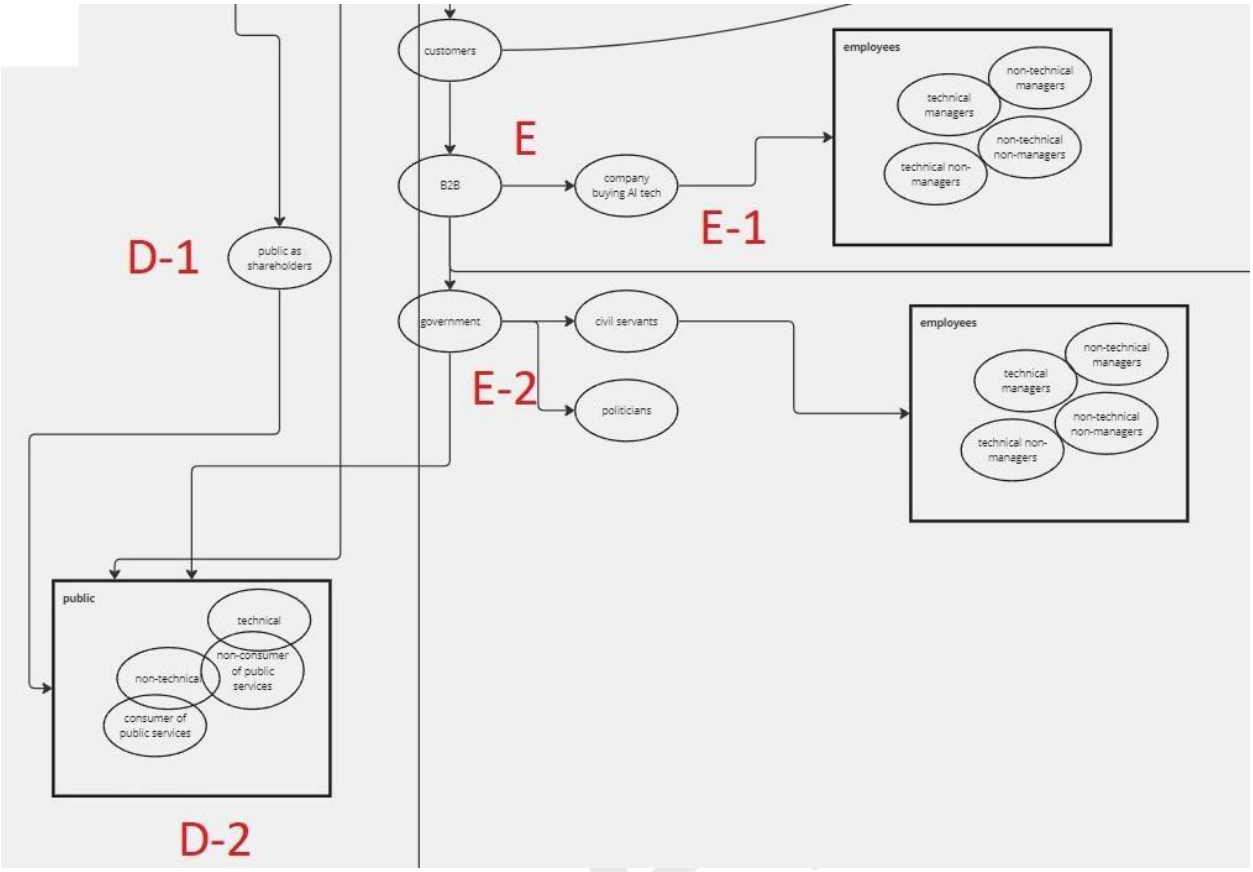

In the next diagram, we focus on some of the potential customers of the AI company, in particular some of the B2B customers (“E”). These B2B customers may include a company buying AI technology (“E-1”), which, like the potential acquiring company, may be led more by the non-technical employees, rather than the technical employees - with the same potential risks. Such customers may not know the right questions to ask the AI company for example, and may have no way to model the downstream risks of using the AI within their products, services or processes.

The government (“E-2”) as a customer adds another wrinkle, in terms of its own customers - who are the public (“D-2”). The “public” plays many roles within this ecosystem. It is important to define exactly which aspects of the public are being considered, the particular members being affected and their role. Even within this portion of the overall diagram, the “public” could be consumers of public services or non-consumers - or even acquirers, if they buy stock in an AI company when it goes public (“D-1”).

Although this is not necessarily correct, we have assumed that the consumers of public services are more likely to be non-technical, in comparison to the non-consumers of such services. This assumption relates to a potentially more dangerous situation than the reverse, or a situation in which both technical and non-technical members of the public are consumers of the public service in question. Those who are non-technical may be more vulnerable to deleterious effects of the AI. Furthermore, if most consumers of a public service are non-technical, the public as a

whole may not realize or consider the potential dangers of applying AI to, or involving AI in, the public service.

The government itself may not be a source of help. In general, politicians are not technical, and so must rely on civil servant advisers. However, most civil servants - especially those higher up in the hierarchy - are either non-technical or have been out of direct contact with technology for such a long time that they may not be able to advise politicians, and/or to make the best purchasing decisions in regard to AI.

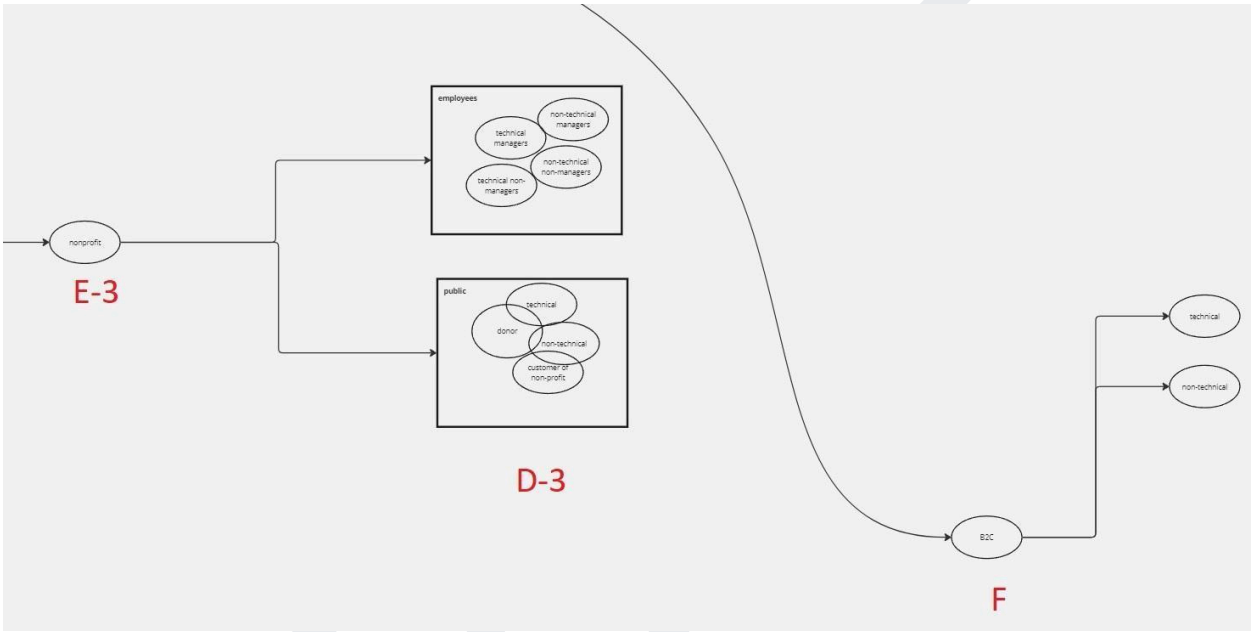

This diagram shows non-profit organisations, which are the third type of B2B customer (“E-3”). Their employees may also be more non-technical than technical, which may cause the previously described problems when it comes to purchasing and integrating AI to their services. They also service the public (“D-3”), which in this case includes donors - who are more likely to have technical knowledge or background than the customers of the non-profits. Again, this situation can cause a disconnect and/or problems between the donors, the non-profit organization and the customers that are served.

The last type of customer are B2C customers (“F”), who may be technical or non-technical. When non-technical B2C customers use an AI company's products, several potential issues could arise – especially since B2C customers may rely on the customer support capabilities of the AI company, rather than having co-workers to assist them (as for a B2B customer). For example, non-technical customers may struggle to comprehend how the AI product works, its capabilities, and limitations. This can lead to unrealistic expectations or frustration when the product doesn't perform as anticipated.

Traditional tech problems for consumers may be amplified by the more complex nature of AI. Non-technical customers might find it challenging to navigate the process of setup or configuration, leading to improper usage of the product altogether – with additional risks to

them. When issues arise, non-technical customers may struggle to diagnose and resolve problems independently. Non-technical customers might misinterpret the outputs provided by the AI product, leading to incorrect decisions or actions based on the information.

Non-technical customers may also fail to sufficiently analyze the terms, conditions and disclosures of the AI company, which may lead to excessive trust and/or a failure to demand transparency. Non-technical customers may not have sufficient knowledge to fully evaluate the ethical implications of AI, such as algorithmic bias, job losses, or the potential for AI systems to be used maliciously.

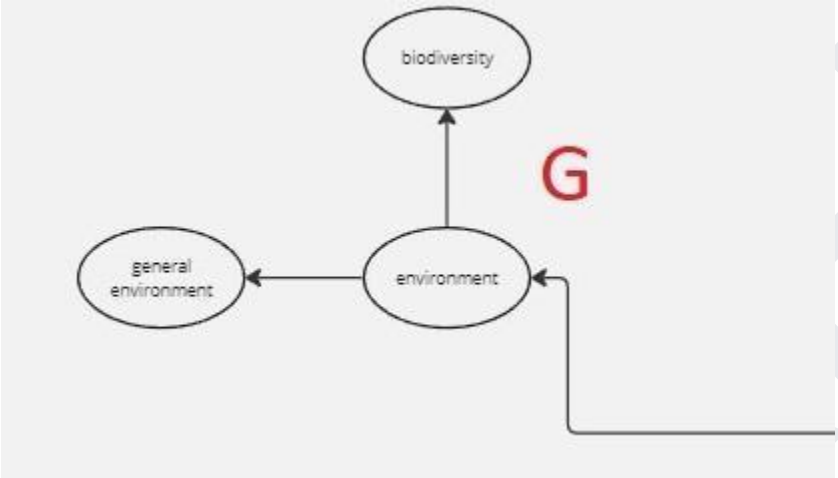

Last - and sadly for many companies, by far the least concern - is the environment (“G”). The environment is divided into biodiversity and the general environment - with the latter including inanimate but important aspects of the environment, such as bodies of water, land areas and regions, the air, and so forth. Energy usage can be a significant issue for AI companies, as their models require a great deal of computing power and massive amounts of stored data (see Appendix IV below).

To assist auditors and others, we have included a package of materials, including the above diagrams. These materials include a manual template for a complete list of stakeholders.

Example manual template row (from the complete manual template):

	A	B	C	D	E	F	G	H
1	Interest	Stakeholder category	Stakeholder group	Goals, motivations and interests	Influence	Interest	Action	Win/win strategies
2	affected by project	Community	directly or indirectly affected by AI, without being able to control its design or use	At a minimum, maintenance of privacy with the AI, and understanding of how and why the AI reached a decision (transparency)	No	High	Meet their needs and earn their trust	Use focus groups and analysts to determine potential vulnerabilities, and how to help these downstream customers.
3	affected by project	Community	other members of the public	At a minimum, maintenance of privacy with the AI, and understanding of how and why the AI reached a decision (transparency)	No	Varied	Earn their trust	Determine public opinion through surveys and focus groups.
4	Regulatory	Government	Regulators - may be direct or indirect; also policymakers and politicians	Government is both a regulator of AI and AI companies, whether direct or indirect. Policymakers wish to influence the direction of AI technological and business development. Politicians may favor AI companies as economic drivers, or else seek to diminish their influence. Some government agencies may purchase AI technology from companies that are regulated by other government agencies. Different parts of the government may have different objectives - and may even work at cross-purposes.	Varied	High	Understand the different stakeholders within government and their roles	1. Maintain transparency and open communication channels with all relevant government entities. 2. Demonstrate a commitment to responsible and ethical AI development, addressing concerns around privacy, security, and fairness. 3. Collaborate with government stakeholders to develop industry standards, guidelines, and best practices for AI governance. 4. Contribute to public discourse and education efforts to build trust and understanding around AI technologies. 5. Adapt engagement strategies to the specific needs and priorities of each government stakeholder while seeking opportunities for collaboration and alignment.
	affected by project	Environment	Environment includes both biodiversity and the general	As detailed in the Appendix on AI and the environment, AI can both help improve the	low	High	Review environmental	1. Prioritize energy efficiency and renewable energy 2. Analyze hardware and energy usage of an AI project before

4. **Diverse Perspectives**:

Ensure a diverse and statistically valid range of stakeholders are included to bring different lived experience relevant perspectives and adjacent disciplines, especially in areas like AI system design, to promote fairness and inclusivity. We also need to educate stakeholders who will be participating in discussions and/or reviews. For example, members at the Team and Organization levels need to be trained in AI security and ethics, as well as best data and integration practices for ethical reasons and to avoid bias.

One aspect of ensuring diverse perspectives is for the company to reach out beyond its immediate ecosystem and tech community. Companies should engage with stakeholders through various channels, such as surveys, focus groups, workshops, and advisory boards, to gather their input, concerns, and expectations regarding the AI product. They should also ensure that the stakeholder engagement process is inclusive and accessible, considering factors such as language, cultural sensitivity, and accessibility needs.

Furthermore, companies should strive for diverse representation among stakeholders, including individuals from different socioeconomic backgrounds, ethnicities, age groups, genders, and abilities. They should also ensure that the perspectives of vulnerable and underrepresented groups are actively sought and incorporated into the AI product development process. They should foster an inclusive environment that encourages open dialogue, respect for diverse opinions, and equal opportunity for participation.

To help stakeholders better understand the specific AI technology, and provide their diverse perspectives more effectively, at least some education and capacity building would be desirable. For example, the company could provide educational resources and training programs to stakeholders, particularly those at the Team and Organization levels, to enhance their understanding of AI security, ethics, and best data practices. The company could conduct workshops and seminars to raise awareness about potential biases, fairness considerations, and the importance of responsible AI development. The company could encourage cross-functional collaboration and knowledge-sharing among stakeholders to promote a holistic understanding of AI's impact on various domains.

Going deeper, Inclusive Design and Development is a great way to be certain that these diverse perspectives are considered from the beginning of the design process - not when development

is complete and the AI product is about to be launched (or has been launched). Some examples of how companies could engage with their stakeholders during this process include:

- Adopting inclusive design principles throughout the AI product development lifecycle, considering the needs and preferences of diverse user groups.
- Conducting thorough testing and evaluation of the AI system to identify and mitigate any biases or disparate impacts on different stakeholder groups.
- Gather feedback from stakeholders, and iterate on the AI product based on their insights and experiences.

Companies should maintain transparency about the AI product's purpose, functionality, and decision-making processes, ensuring that stakeholders have a clear understanding of how the system operates, during all stages of the product development lifecycle. Companies should also establish accountability mechanisms to address any concerns, complaints, or unintended consequences raised by stakeholders. The company should regularly communicate progress, challenges, and mitigation strategies to stakeholders, fostering trust and confidence in the AI product development process.

Going even deeper, companies may wish to engage in a process of ongoing Stakeholder Engagement and Feedback. For example, companies could establish ongoing channels for stakeholder engagement and feedback throughout the AI product lifecycle, from ideation to deployment and beyond. They should continuously monitor and assess the AI product's impact on different stakeholder groups, making necessary adjustments and improvements based on their feedback and evolving needs. They could also foster a culture of continuous learning and improvement, acknowledging that stakeholder needs and expectations may change over time.

Appendix I - Background material on who is a stakeholder

In this document, we consider the concept of a stakeholder in the context of AI project management. In public services, particularly within Westminster systems where PBMOK (Project Management Body of Knowledge) derives from, the term is routinely used to describe any representative, category, group, or natural person with a recognizable, relevant interest in a decision or policy proposal. This neutral definition should be relied upon for clarity and inclusiveness.

However, the term "stakeholder" has an American origin, derived specifically from "shareholder," and its roots are embedded in private sector governance, particularly during the Reaganomics era of the 1980s. Historically, this term has been contentious, shaped by two primary schools of thought. One school views shareholders as a class of stakeholders, advocating that the primary responsibility of corporations is to maximize shareholder profits. The other school rejects the relevance of stakeholders in corporate decision-making, focusing exclusively on profit maximization.

The Shareholder Primacy Perspective

This perspective, rooted in Milton Friedman's doctrine, views stakeholders as "externalities"—non-market transactional elements outside the purview of corporate social responsibility and decision-making. Proponents argue that markets are neither intended nor equipped to manage such externalities effectively. This perspective holds that corporations should focus solely on maximizing shareholder value, dismissing broader societal and environmental considerations.

The Broader Accountability Perspective

Conversely, the second perspective advocates for a broader accountability of corporations. It argues that corporate operations, profit, and cash flow do not exist in isolation. Corporations should be accountable for the full cost of their operations, including their impacts on the commons or commonwealth. This approach emphasizes the need for corporations to avoid causing unnecessary harm for short-term profit, considering the potential for long-term liability, as well as personal culpability of those who declined to robustly consider the risks being negligently, recklessly, or intentionally taken, and the opportunity cost of innovation due to a narrow, short-term focus at the top management and oversight level.

Changing Perspectives in Economics

While the first, “stay away” view on externalities has been the primary view until now, the second view is becoming increasingly prominent. For example, Angus Deaton, the British-American economist and Nobel Prize winner in Economic Sciences in 2015, has recently expressed a critical viewpoint towards mainstream economics in the International Monetary Fund's (IMF) Finance and Development magazine. Deaton argues that mainstream economics is in disarray because it overlooks the impact of power dynamics on economic outcomes, neglects questions of equity, and often results in policy recommendations that exacerbate wealth inequality. He

highlights the profession's failure to address the real-world issues facing people, particularly in the United States, where economic growth no longer benefits the broader population as it once did. Deaton points to the rising "deaths of despair" and the growing income inequality as symptoms of these failures.

Deaton suggests that the economics profession needs to broaden its focus beyond monetary measures of welfare to consider other aspects of human well-being, such as job satisfaction, family, and community. He proposes that economists pay more attention to pre-distribution—how income is distributed by the market before taxes and transfers—and to policies that could prevent economic distress. This shift would require economists to engage with more complex and politically sensitive areas like union promotion, immigration control, and industrial policy.

Pushback and Corporate Structures

While this broader accountability view is growing in prominence, there has been pushback from various sectors, including other economists and corporate structures, and even corporate organizations. For example, the American Business Roundtable released an updated “Statement on the Purpose of a Corporation” in 2019, which updated this Purpose to include benefitting all stakeholders. The party most inhibited by the shift to stakeholder primacy has been the board of directors. Whereas previously they acted as a balance and check on operations, their independence has been actively eroded by the doctrine of shareholder primacy. This doctrine has resulted in the effective replacement of a broad duty at law "to act in the best interests of the organization" with a narrow one that redefines their role as exclusively focusing on share price. For AI regulations and governance structures to be fully operative, a return to stronger board independence is required.

Stakeholder Theory in Public and Not-for-Profit Sectors

The application of stakeholder theory extends beyond the private sector. It is more prevalent and clearly understood in not-for-profit and public sector organizations. Here, "stakeholder" aptly captures the idea of parties with legitimate and relevant interests, who should be consulted and considered in relation to policies, public benefits, or other matters. This consultation aims to enhance the robustness of analysis and decision-making.

Practical Applications of Stakeholder Analysis

Stakeholder analysis has practical applications, such as ensuring statistical validity and co-design methodologies. Co-design particularly aligns with indigenous cultural norms and practices, allowing for culturally sensitive inputs in not-for-profit and public decisions. It also supports the cultural norms and practices of various groups, including ethnic, religious, national, gender/gender identity-based, sexual orientation-based, age-based, class-based, professional, and those who have different types of disabilities and/or need accommodations.

The advent of technology-enabled inputs presents a new frontier in stakeholder engagement, though this area requires more consistent and effective implementation, embracing a 'no wrong door' approach.

Appendix II – summaries of material used for this document

The document "Australian AI ethics principles: Flamingo AI looks at small vendor challenges" from the Department of Industry, Science and Resources discusses Flamingo AI, a small Australian AI start-up. It details Flamingo AI's participation in the Australian AI Ethics Principles pilot, highlighting their approach to ethical AI development. Key points include the company's commitment to ethical culture, staff training in security and ethics, data analysis for bias prevention, and ongoing client engagement in ethical AI practices. Challenges and opportunities for small businesses in adopting AI ethics principles are also explored, with emphasis on the need for continuous monitoring, effective training, public demand for ethical AI, and the importance of transparency and accountability in AI products. The document underscores the growing importance of ethical AI principles in business practices.

The document "Australia's AI Ethics Principles" outlines eight principles to ensure AI is safe, secure, and reliable. These principles focus on human, societal, and environmental well-being, human-centered values, fairness, privacy, security, reliability, transparency, contestability, and accountability. They aim to achieve fairer outcomes for all Australians, reduce negative AI impacts, and maintain high ethical standards. The framework is voluntary, designed to complement existing AI regulations, and emphasizes building public trust, consumer loyalty, and positive AI outcomes. It also underlines the importance of diversity in AI system design and the need for accountability throughout the AI lifecycle.

Problematically, despite the existence of this AI ethics principles and a further, older guideline on achieving ethical and accountable automated decision making in government organisations, Australia has simultaneously experienced a catastrophic failure of the public interest, fairness, reliability, contestability and human centred values tests with an income averaging algorithm that has earned the nickname 'Robodebt' that seems to have been engineered to ensure that latent distrust in government use of technology and data is perpetuated - in ways a cynic might think was done deliberately in order to manufacture conditions under which existing culture and practices in the Australian Public Service that are already unaccountable and not transparent could be continued with public support.

The article "Getting all the Stakeholders on the Same Page" by Emotive Brand emphasizes the importance of engaging stakeholders in a project. It describes stakeholders as anyone who can affect or is affected by a corporation's actions. The article divides stakeholders into internal (employees, managers, owners) and external (funders, investors, shareholders, advisors, banks, suppliers, policymakers, legislators, social media influencers, and customers). It advises early engagement of stakeholders, ideally during the planning stage, to establish a common understanding of project scope, timing, budget, and resources. The article also suggests strategies for handling difficult stakeholders, such as understanding their motivations and goals, maintaining open communication, and relying on data to support project direction.

The document "IT Stakeholder List - 112 stakeholders" provides an extensive list of roles, job titles, departments, or groups who might be stakeholders in an IT project. This comprehensive list serves as a starting point for stakeholder analysis or as a checklist to ensure no roles are overlooked. While focused on IT projects, the list also includes some generic roles, pointing to another document for a more comprehensive generic stakeholder list. The variety in this list reflects the wide range of individuals and groups that can have an interest in or be affected by IT projects.

The "Example Mapping Template - IT Project" document is a detailed template for stakeholder analysis in IT projects. It categorizes stakeholders into groups and outlines their goals, motivations, interests, influence, and interest level. The document also suggests actions and win/win strategies to engage each stakeholder effectively. This template is a practical tool for identifying and managing stakeholders in IT projects, ensuring their needs and influences are considered throughout the project lifecycle.

The "Basic Stakeholder Analysis Template" document provides a structured format for analyzing stakeholders in a project. It prompts the user to list stakeholders, assess their expectations, and categorize their influence and concerns using a Red-Amber-Green (RAG) status. This tool is designed to assist project managers in identifying key stakeholders, understanding their needs and expectations, and developing strategies to engage them effectively throughout the project lifecycle.

Appendix III - Example Lists of Stakeholders

We suggest including at least the following stakeholders:

1. **Team Level Stakeholders:**
    a. programmers/developers
    b. product managers
    c. project managers
    d. UX/UI designers
    e. Team leaders
    f. business analysts
    g. architects
    h. data scientists
    i. testers (QA/QC)
2. **Organization Level stakeholders**
    a. investors
    b. Board members
    c. C-suite
    d. Other executives/managers
    e. Sales & Marketing
    f. Other employees
3. **Ecosystem Level Stakeholders**
    a. Customers of the Organization (OC, or Organization Customers)
    b. AI Tech infrastructure providers
    c. AI model developers
    d. AI solution creators
    e. AI integrationists
    f. AI researchers
    g. Government Direct Regulators (who may be subject to regulatory capture)
4. **Directly Affected Consumers and Others**
    a. Clients of OC
    b. Communicants of OC
    c. Future/existing OC employees
    d. Target of the system (e.g. biometric surveillance)
5. **Society**
    a. Government risk and audit
    b. Government regulatory agencies
    c. Non-profit reviewers
    d. Public
    e. Marginalized Communities

6. Environment

a. Biodiversity
b. Emissions
c. Water Consumption
d. General environment

Appendix IV - the environment as stakeholder

AI companies could significantly impact the environment in terms of energy consumption, particularly by 2030. For example, AI servers alone may use 0.5% of the world's electrical generation even as early as 2027, on top of the current 1% of global electrical generation used by data centers. Amazon is just one of the many companies that is investing heavily in data centers, committing to spending $148 billion over the next 15 years to build and operate data centers worldwide. Much of this expansion is geared towards meeting the rising demand for AI services. One study estimates that data centers' electricity use in the European Union could increase by 28% by 2030. This is driven by the increased computational power and storage requirements of AI models.

The need for data centers to have reliable power at all times could lead them to rely more on fossil fuel sources, as renewable energy may not always be able to meet the fluctuating demand. Data centers, and the need for more computing power more generally, may lead to increased e-waste. The environmental impact of hardware production and disposal, including rare earth mineral extraction and e-waste generation should also therefore be considered.

Higher energy demands for AI training and inference are also increasingly an important driver of AI energy usage, adding strain on an existing energy grid that is challenged with demand/supply optimization. Training large AI models like GPT-4o can consume massive amounts of electricity, equivalent to the emissions of driving 112 gasoline-powered cars for a year. Additionally, the energy required for AI inference (making predictions) may be even higher than the training phase, with estimates suggesting 60% of AI energy use goes towards inference.

For example, interacting with a Language Learning Model (LLM) could cost 10 times more than a standard keyword search, with a single LLM interaction consuming approximately 3 Wh of electricity, according to Google. If every Google search for a year used AI, it would consume the equivalent amount of electricity used to power a small country like Ireland.

While AI companies are working on making their tools more efficient, the increased accessibility and widespread adoption of AI could lead to a significant surge in energy consumption, contributing to excess carbon emissions, excess heat dispensation and overarching climate change concerns. As AI becomes more prevalent, it is crucial for companies to be mindful of their energy usage and to invest in sustainable and renewable energy sources to mitigate the environmental impact of this rapidly growing technology.

Even with improvements in hardware and software efficiency, the growing computational intensity of AI may not be entirely offset, leading to a "rebound effect" where the extra resources fuel more demand for AI applications.

The increased energy consumption of AI is expected to significantly contribute to global greenhouse gas emissions. Estimates suggest AI could account for up to 5% of global emissions by 2030 if left unchecked. AI-related electricity consumption could increase by 85 to 134 TWh annually by 2027, comparable to the annual electricity needs of countries like the Netherlands, Argentina, and Sweden.

However, as Professor Roberto Verdecchia of the University of Florence noted in “Generative AI’s Energy Problem Today Is Foundational” (IEEE, link below), “In the race to produce faster and more-accurate AI models, environmental sustainability is often regarded as a second-class citizen.”

Still, AI can also have positive impacts on the environment. Thus, the environment, encompassing both biodiversity and the general environment (including bodies of water, regions of land, air, etc.), has a complex relationship with AI development and deployment. While the environment itself does not have explicit goals or motivations, there are inherent interests that should be considered when assessing the impact of AI on the environment. Some of these interests relate to the negative impact, particularly of energy usage. However, there are potential positive impacts as well. Here's a summary of some potential positive impacts - and how AI can further the interests of the environment:

1. Conservation and Protection:

- The primary interest of the environment is its conservation and protection from harmful impacts.
- AI technologies can be leveraged to monitor and analyze environmental data, enabling better decision-making for conservation efforts.
- AI can support initiatives such as wildlife tracking, habitat monitoring, and ecological forecasting to safeguard biodiversity.

2. Sustainable Resource Management:

- The environment has an inherent interest in sustainable resource management to ensure long-term ecological balance.
- AI can optimize resource utilization, such as water management, precision agriculture, and sustainable forestry practices.
- Intelligent systems can help monitor and mitigate the environmental impact of resource extraction and industrial activities.

3. Climate Change Mitigation:

- The environment is directly affected by climate change, and there is a pressing need to mitigate its impacts.
- AI can contribute to climate change mitigation by optimizing energy systems, improving energy efficiency, and supporting the transition to renewable energy sources.
- AI-powered tools can help model and predict climate patterns, enabling better adaptation strategies and risk assessment.

4. Pollution Reduction:

- The environment has an interest in reducing various forms of pollution, including air, water, and land pollution.
- AI can be applied to monitor and control pollution levels, optimize waste management systems, and support the development of cleaner technologies.
- Intelligent systems can help identify pollution sources, predict dispersion patterns, and inform targeted interventions.

5. Ecosystem Balance and Resilience:

- The environment thrives on maintaining a delicate balance between various ecosystems and their components.
- AI can assist in modeling and understanding complex ecological interactions, predicting ecosystem responses to disturbances, and informing conservation strategies.
- AI-driven insights can help build resilience in ecosystems by identifying vulnerabilities and guiding restoration efforts.

6. Energy Efficiency and Renewables:

- The environment benefits from the adoption of energy-efficient practices and the transition to renewable energy sources.
- AI can optimize energy systems, improve grid management, and facilitate the integration of renewable energy into existing infrastructure.
- Intelligent algorithms can help predict energy demand, optimize energy storage, and enhance the performance of renewable energy technologies.

Finally, we are left with the “unknown unknowns”, such as indirect effects on the environment due to changes in human behavior and decision-making influenced by AI, which also need to be considered.

To address these concerns, AI companies (and their auditors) need to consider mitigation strategies, such as:

- Developing energy-efficient AI algorithms and hardware solutions.
- Adopting sustainable practices in AI development and deployment, such as using renewable energy sources and responsible disposal of hardware.
- Conducting thorough environmental impact assessments and implementing mitigation strategies.

Collaborating with environmental experts and organizations to ensure AI aligns with conservation and sustainability goals.

Citations:

https://www.euronews.com/next/2023/10/10/demand-for-ai-could-mean-technology-consumes-same-energy-as-a-country-analysis-shows

https://www.axios.com/2023/12/04/cop28-artificial-intelligence-energy

https://www.thestreet.com/technology/ai-electricity-requirements

https://spectrum.ieee.org/ai-energy-consumption

https://www.edwardconard.com/macro-roundup/ai-servers-could-use-0-5-of-the-worlds-electrical-generation-by-2027-for-context-data-centers-currently-use-around-1-of-global-electrical-generation/?view=detail

https://www.theatlantic.com/technology/archive/2023/08/ai-carbon-emissions-data-centers/675094/

https://www.bairesdev.com/blog/ai-energy-consumption/

https://news.climate.columbia.edu/2023/06/09/ais-growing-carbon-footprint/

https://www.theguardian.com/technology/2024/mar/07/ai-climate-change-energy-disinformation-report

https://agsiw.org/the-rising-energy-demand-of-ai-and-clouds-unraveling-the-environmental-conundrum/

Appendix V - PBMOK (Project Management Body of Knowledge)

The PMBOK (Project Management Body of Knowledge) guide provides guidance on how to effectively manage project stakeholders as part of the overall project management process:

1. Identify Stakeholders: The PMBOK states that identifying all the stakeholders, both internal and external, is a critical first step in the Initiating process group. This involves creating a stakeholder register that lists all the key individuals, groups, or organizations that can affect or be affected by the project.

2. Plan Stakeholder Engagement: The PMBOK recommends creating a Stakeholder Engagement Plan as part of the overall Project Management Plan. This plan outlines the strategy for interacting with each stakeholder based on their level of interest and influence in the project.

3. Manage Stakeholder Engagement: The PMBOK provides guidance on actively managing stakeholder engagement throughout the project lifecycle. This includes communicating project information, addressing concerns, and maintaining stakeholder satisfaction.

4. Monitor Stakeholder Engagement: The PMBOK states that the project manager should continuously monitor the engagement and relationships with stakeholders, making adjustments to the stakeholder management approach as needed.

In summary, the PMBOK framework emphasizes the importance of comprehensive stakeholder management as a critical component of successful project delivery. It provides a structured approach to identifying, planning, managing, and monitoring stakeholder engagement.

Citations:
https://www.whizlabs.com/blog/pmp-project-stakeholders/
https://pmstudycircle.com/stakeholders-in-project-management-definition-and-types/
https://www.projectengineer.net/project-stakeholder-management-according-to-the-pmbok/
https://monday.com/blog/project-management/project-management-body-of-knowledge-pmbok/
https://openclassrooms.com/en/courses/4849646-project-management-101/5183906-identify-and-communicate-with-project-stakeholders

Appendix VI - stakeholder matrices

Stakeholder matrices are diagrams used to visualize and map stakeholders based on specific characteristics. The matrix typically has four quadrants, with the X-axis and Y-axis labeled with different characteristics (e.g., level of interest, impact, influence, support, etc.).

The main benefits of using a stakeholder matrix include:

1. Categorizing stakeholders for tailored communication and engagement

2. Identifying potential risks, resistance, or opportunities

3. Setting priorities for managing or engaging with stakeholders

4. Communicating strategies to colleagues and leadership

5. Tracking progress over time

Examples of Stakeholder Matrix Frameworks

There are three common stakeholder matrix frameworks, each of which is described in greater detail below:

1. The Power-Interest Grid: maps stakeholders on power to affect, vs interest in, a subject

2. The Power-Predictability Matrix: power vs likelihood to kick up a fuss (ie – be unpredictable)

3. The Stakeholder Knowledge Base Diagram: knowledge of, vs support for, a subject

1. The Power-Interest Grid:

The Power-Interest Grid maps stakeholders based on their level of power/influence (X-axis) and their level of interest (Y-axis). The four quadrants in this grid are:

1. High Power/High Interest: These stakeholders should be managed closely, as they have significant influence and are highly interested in the project's outcomes.

2. High Power/Low Interest: These stakeholders should have their needs met, as they have significant influence but may not be actively engaged in the project.

3. Low Power/High Interest: These stakeholders should be kept informed, as they are highly interested in the project but have limited influence.

4. Low Power/Low Interest: These stakeholders should be monitored, but with minimal effort, as they have limited influence and interest in the project.

2. The Power-Predictability Matrix:

The Power-Predictability Matrix maps stakeholders based on their level of power (X-axis) and their level of predictability (Y-axis). The four quadrants in this matrix are:

1. High Power/High Predictability (C): These stakeholders are powerful but predictable, so they are likely to cause few problems.

2. High Power/Low Predictability (D): These stakeholders represent the greatest danger or opportunity, as they are powerful and unpredictable.

3. Low Power/High Predictability (A): These stakeholders are likely to cause few problems, as they have low power and are predictable.

4. Low Power/Low Predictability (B): These stakeholders are somewhat unpredictable and may present some danger or opportunity, but their low power limits their impact.

3. The Stakeholder Knowledge Base Diagram:

The Stakeholder Knowledge Base Diagram maps stakeholders based on their level of knowledge (X-axis) and their degree of support (Y-axis). The four quadrants in this diagram are:

1. High Knowledge/High Support: These stakeholders are aware and supportive of the project and should be kept informed to maintain their support.

2. High Knowledge/Low Support: These stakeholders are aware but not supportive of the project. Efforts should be made to understand their concerns and address them.

3. Low Knowledge/High Support: These stakeholders are supportive but not well-informed about the project. They should be educated to ensure their continued support.

4. Low Knowledge/Low Support: These stakeholders are neither well-informed nor supportive of the project. Efforts should be made to understand their concerns, educate them, and build their support.

Other Examples of Stakeholder Matrices

- Power-Support Stakeholder Analysis Matrix
- Support-Importance Stakeholder Matrix
- Influence-Interest Matrix

The Power-Support Stakeholder Analysis Matrix maps stakeholders according to their level of power and the level of support they provide. This helps identify which stakeholders are critical to a project's success.

The matrix has four quadrants:

1. Promoters - Stakeholders with high power and high support. These are the champions who can be engaged to help drive the project forward.

2. Defenders - Stakeholders with high power but low support. These are the stakeholders that need to be carefully managed, as they have the ability to block or undermine the project if not properly engaged.

3. Supporters - Stakeholders with low power but high support. These are helpful allies who can provide assistance, but may not have the authority to directly influence the project.

4. Bystanders - Stakeholders with low power and low support. These are the least critical stakeholders who can be monitored with minimal effort.

The key benefit of this matrix is that it helps the project manager identify who the critical stakeholders are that need the most attention and engagement. If a completed matrix has few stakeholders in the "Promoters" quadrant, it may indicate the project lacks sufficient sponsorship and high-level support to be successful.

By understanding the power and support levels of each stakeholder, the project manager can develop targeted strategies to manage them effectively, such as increasing the commitment of "Defenders" or leveraging the support of "Supporters".

The Support-Importance Stakeholder Matrix, developed by Paul Nutt, is a grid for managing focus groups made up of stakeholders. Stakeholders are positioned on the matrix according to their expected level of support and importance to the issue.

The matrix has the following key elements:

Y-axis: Stakeholders' Issue Position - This ranges from 1 (strongly oppose) to 5 (strongly support) to identify the "intensity of the expected opposition or support" from each stakeholder.

X-axis: Importance - This ranges from 1 (least important) to 10 (most important) to measure the stakeholder's overall importance.

The four quadrants of the matrix are:

1. Antagonistic - Stakeholders who are likely to oppose the decision/action and are also highly important. These are critical stakeholders that need to be carefully managed.

2. Supportive - Stakeholders who are supportive of the decision/action and are also highly important. These are the key champions that should be leveraged.

3. Indifferent - Stakeholders who have low support and low importance. These can be monitored with minimal effort.

4. Concerned - Stakeholders who oppose the decision/action but have low importance. Their concerns should be acknowledged, but they do not require extensive engagement.

The key benefit of this matrix is that it helps identify the most critical stakeholders - those who are both highly important and highly opposed (the "Antagonistic" quadrant).

These are the stakeholders that require the most attention and engagement to address their concerns and secure their support.

By understanding the support and importance levels of each stakeholder, the project manager can develop targeted strategies to manage them effectively, such as increasing the commitment of "Antagonistic" stakeholders or leveraging the support of "Supportive" stakeholders.

The Influence-Interest Matrix is a tool used to categorize stakeholders based on their level of influence and interest in a project or initiative. It is similar to the Power-Interest Matrix. However “influence” is not necessarily power – or at least not classic “hard” power. Influence instead is more of a “soft power”. It helps project managers and organizations identify key stakeholders and develop appropriate strategies for managing them.

The matrix has the following key elements:

X-axis: Influence - This represents the stakeholder's ability to impact the project's outcomes, either positively or negatively. It considers factors like their authority, resources, and decision-making power.

Y-axis: Interest - This represents the stakeholder's level of concern or engagement with the project and its outcomes. It considers factors like how much the project affects them and their motivations.

The matrix is divided into four quadrants:

1. High Influence, High Interest (Manage Closely) - These are the most critical stakeholders who have both high influence and high interest in the project. They require the most attention and engagement to secure their support and address their concerns.

2. High Influence, Low Interest (Keep Satisfied) - These stakeholders have significant influence but low personal interest in the project. They need to be kept satisfied and engaged to maintain their support, but may not require as much direct involvement.

3. Low Influence, High Interest (Keep Informed) - These stakeholders have limited influence but a high personal stake in the project. They should be kept well-informed and their concerns addressed, even though they cannot directly impact the project.

4. Low Influence, Low Interest (Monitor) - These stakeholders have minimal influence and interest. They can be monitored with minimal effort, but it's still important to keep them informed of key developments.

By using the Influence-Interest Matrix, project managers can prioritize their stakeholder engagement efforts, allocate resources more effectively, and develop tailored communication and management strategies for each stakeholder group. This helps ensure critical stakeholders are actively engaged, potential opposition is managed, and overall stakeholder support is secured for the project's success.

Using Stakeholder Matrices

As a short summary, to use a stakeholder matrix effectively, one should:

1. Determine relevant characteristics

2. Determine characteristic levels

3. Plot stakeholders on the matrix

A more detailed version of the process of creating the stakeholder matrix is as follows:

1. Define the purpose of the stakeholder matrix. Understand what information you want to capture and communicate. In particular, consider all of the stakeholder characteristics that you may wish to include (note above examples).

2. Brainstorm and identify all potential stakeholders, both internal (e.g. employees, managers) and external (e.g. customers, suppliers, regulators). Be comprehensive in your list.

3. Identify each stakeholder's key interests, goals, and concerns. Understand their motivations and potential conflicts between stakeholders. Turn this information into stakeholder characteristics.

4. Determine the level of each characteristic for each stakeholder, which may include knowledge, power, interest, involvement, influence, support, predictability, and so forth. Categorize them as having high, medium or low of each of these characteristics. This can be visualized using a grid or network diagram.

5. Create one or more stakeholder matrices according to each selected pair of characteristics.

6. Develop an engagement plan for how you will communicate with and involve each stakeholder group according to these matrices. Use a variety of methods like surveys, focus groups, workshops, and online platforms.

7. Continuously monitor and update the stakeholder matrices as the project or organization evolves. Stakeholder relationships and priorities can change over time.

The goal is to create a comprehensive, visual representation of the stakeholder landscape that informs your strategy for effectively engaging and managing stakeholders throughout a project or initiative, using stakeholder matrices as one tool.

Practical Examples

Citation 1 below provides three practical examples of using stakeholder matrices in different contexts:

1. A healthcare organization using the Power-Interest Grid to determine stakeholder involvement, to be certain that individuals with high interest but low power (as well as lower degrees of interest) were included in the process.

2. A construction company using the Power-Predictability Matrix to assess stakeholders for a bridge-building project, to enable it to manage relationships also based on predictability, not just power – meaning that they had to engage with both local government (high power) and environmental groups (moderate power), due to their unpredictability.

3. A state government organization using the Stakeholder Knowledge Base Chart to determine engagement and communication needs for a new policy, in regard to support. Communication with groups with lower amounts of support for the new policy could be adjusted according to their level of knowledge.

Other stakeholder mapping methods and diagrams may be more suitable, as matrices have limitations due to their focus on only two characteristics.

Citations:

1. https://simplystakeholders.com/stakeholder-matrix/
2. https://www.mural.co/blog/stakeholder-mapping
3. https://fastercapital.com/keyword/stakeholder-input.html
4. https://fastercapital.com/keyword/stakeholder-feedback.html
5. https://implementationsciencecomms.biomedcentral.com/articles/10.1186/s43058-020-00023-7
6. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7427961/
7. https://www.linkedin.com/advice/1/how-do-you-use-powerinterest-matrix
8. https://pressbooks.ulib.csuohio.edu/project-management-navigating-the-complexity/chapter/5-2-stakeholder-analysis/
9. https://www.stakeholdermap.com/stakeholder-matrix.html
10. https://www.indeed.com/career-advice/career-development/types-of-stakeholder-analysis
11. https://www.solitaireconsulting.com/2020/07/stakeholder-management-using-the-power-interest-matrix/
12. https://sustainet.com/stakeholder-matrix/
13. https://arpedio.com/relationship-mapping/understanding-the-power-of-a-stakeholder-matrix-a-comprehensive-guide/
14. https://thinkinsights.net/strategy/stakeholder-analysis/
15. https://www.linkedin.com/pulse/stakeholder-influence-matrix-business-coaching-tool-chandan-patary
16. https://www.improvementservice.org.uk/business-analysis-framework/consider-perspectives/powerinterest-grid