

# Aligning Transparency and Disclosure Practices with Human Rights Responsibilities

Guide 7 of the Responsible AI Practitioner Guides for Taking  
a Human Rights-Based Approach to Generative AI

February 2025



# Summary

This paper provides guidance about how entities across the generative AI (genAI) value chain can ground their disclosures and transparency efforts in best practices for human rights and sustainability reporting. It includes the following sections:

- 1 The Human Rights Disclosure Landscape:** A high-level overview of human rights disclosure, including important definitions, key voluntary and mandatory standards, and the benefits of disclosure.
- 2 How GenAI Value Chain Entities Can Apply Disclosure Best Practices for Human Rights Impacts:** This section begins with some key observations to bridge the gap between disclosure best practices and the genAI value chain context. It then provides examples of disclosure for each value chain entity, including suggestions on format, content, and indicators.
- 3 Key Resources**

## Key Points

- Disclosures are a key human rights risk mitigation measure in the genAI value chain. Negative human rights impacts associated with actions from upstream actors may be passed downstream, and remedies for those impacts will often require coordination. Proactive information sharing by all value chain entities is thus core to a cross-value chain strategy for mitigating potential human rights impacts.
- Disclosures may occur at different levels, with content that varies according to the information needs of the reader:
  - Corporate-level reporting, usually annually and following formal reporting standards, for use by investors, civil society organizations, and other company analysts.
  - Model / product / system-level disclosures. These are targeted at audiences with specialist AI knowledge and may be based on the principles of interpretability and explainability.

## ACCOMPANYING RESOURCES

- [A HRA of the GenerativeAI Value Chain](#)
- [Overview of the Practitioner Guide](#)
- [Guide 1: Human Rights Fundamentals](#)
- [Guide 2: Governance and Management](#)
- [Guide 3: Impact Assessment](#)
- [Guide 4: Risk Mitigation](#)
- [Guide 5: Stakeholder Engagement](#)
- [Guide 6: Policies and Enforcement](#)
- [Guide 7: Transparency and Disclosures](#)
- [Guide 8: Remedy for GenAI Related Harms](#)

- Disclosures for users, such as model cards or dataset datasheets. These are aimed at the general public, including information about basic capabilities and safe usage, and should use accessible language.
- Currently, foundation model developers have the most advanced disclosure practices, which have included as standard practice documents such as system cards, model cards, datasheets for datasets, or blogposts setting out training methods or research direction. Best practices on disclosure have yet to emerge for suppliers, application developers, and system deployers.
- The capacity to provide updated disclosure will likely significantly vary among genAI value chain entities, a trend that is also observable in digital companies more broadly.
- The format and depth of disclosure for gen AI value chain entities may vary, in part due to capacity limitations and transparency trade-offs mentioned above.
- Disclosures can feel risky due to fears over reputational damage and the exposure of trade secrets.
- Disclosures can provide too much information, thereby enabling bad actors or overwhelming a nontechnical audience; striking a balance is a key consideration.
- Disclosures should focus on risks to people that are the most salient, in accordance with the UN Guiding Principles on Business and Human Rights (UNGPs) criteria of severity and likelihood.
- It is important to carefully select which metrics are used as indicators for disclosure. Some are more informative on human rights impacts than others.

# 1. The Human Rights Disclosure Landscape

## Key Definitions

- A **disclosure** is information that a company or other entity reveals about itself that would not otherwise be available or easily discoverable. Disclosures differ based on whether they occur at the organization-wide level or at the level of the AI system.
  - Company-level disclosures may cover governance, strategy, risks and impacts, indicators, and other data that is decision-useful for readers.
  - Model- or product-level disclosures include model and/or system cards and dataset datasheets.

To be useful, disclosures should provide sufficient detail to enable internal and external evaluation of how a company or other entity is posing risks to, and/or enabling the realization of human rights. Some publications may also use the word “transparency” or “reporting” to refer to disclosures.

- **Sustainability and human rights** are related but distinct fields. Sustainability is a well-developed field with formal reporting standards (these are covered in the next section). Human rights is a component of sustainability, and the UNGPs have been deliberately integrated into the Global Reporting Initiative (GRI) standards and [European Sustainability Reporting Standards \(ESRS\)](#) (see below), which are also unambiguously human rights disclosure standards. However, practical guidance for how to apply these standards in the context of AI, including genAI, is lacking. Much of the Responsible AI field, and responsible tech more broadly, has leaned toward creating its own bespoke approach to disclosure, rather than using well established sustainability and human rights reporting standards.

## Key Voluntary and Mandatory Disclosure Standards

Voluntary standards for corporate disclosure of human rights impacts have existed for decades. While not legally binding, these standards have become best practice and are core to stakeholder expectations. Four key voluntary standards are relevant:

1

The [UN Guiding Principles on Business and Human Rights](#) establish public reporting as the fourth of four elements in a company's human rights due diligence process—assessment, action, tracking, and public communications. Principle 21 of the UNGPs states that disclosure should:

- Include analysis of salient human rights impacts based on severity and likelihood
- Be of a form and frequency that reflect an enterprise's human rights impacts and that are accessible to its intended audiences
- Provide information that is sufficient to evaluate the adequacy of an enterprise's response to the particular human rights impact involved
- Not pose risks to affected stakeholders, personnel, or legitimate requirements of commercial confidentiality

### KEY ELEMENTS OF ENTITY-LEVEL HUMAN RIGHTS DISCLOSURES

#### Strategy elements

- Explanation of how the human rights commitment is embedded across the business and value chain, including, for example, issue-specific policies, trainings, contractual clauses
- Explanation of challenges, examples of how these are addressed, and lessons learned
- Forward focused with clear strategy, targets, and goals that are measurable
- Risk management elements
- Salient human rights issues, including impacts that occur in reporting period
- Description of methodology to identify and prioritize impacts for action and disclosure
- Measures taken to prevent and mitigate salient human rights issues, including, for example, efforts to address systemic challenges such as participation in multi-stakeholder initiatives and engagement in public policy

- Process to address grievances and examples of grievances received and remedy outcomes
- Affected stakeholder engagement, including how this informs actions to address salient issues

#### Governance elements

- Description of human rights governance demonstrating top-level oversight and well-integrated cross-functional structures

#### Indicators, metrics, and targets

- Key performance indicators and other quantitative and qualitative metrics to support the narrative

2

**The Global Reporting Initiative (GRI)** provides guidance for sustainability reporting for all organizations of all types, sizes, and locations, and across all sustainability topics. The GRI is by far the most commonly used sustainability reporting standard—it is used by over 80% of the world’s 250 largest companies. The [GRI standards](#) have been [updated](#) to align with the UNGPs, including recommendations to disclose:

- The salience of human rights impacts based on severity and likelihood, which should be assessed by external stakeholders and senior leadership within the corporate entity.
- Relevant quantitative and qualitative indicators for those human rights impacts.

3

**The International Financial Reporting Standards (IFRS) Foundation’s ISSB S1 on General Requirements for Disclosure of Sustainability-related Financial Information.** The ISSB S1 is a formal requirement that companies disclose useful information about their sustainability-related risks and opportunities. The ISSB S1 also prescribes how these reports should be prepared and their content.

- Additionally, the [Sustainability Accounting Standards Board \(SASB\) Standards](#) provide industry-specific metrics for companies to disclose material and decision-useful sustainability information to investors. SASB includes metrics for specific topics, such as privacy and supply chain labor standards, but it does not contain metrics on the overall human rights due diligence process.

4

**The [OECD Guidelines for Due Diligence for Responsible Business Conduct](#)** provide practical support to enterprises on the implementation of the [OECD Guidelines for Multinational Enterprises](#), an authoritative global standard on responsible business conduct. The guidelines provide plain language explanations of due diligence recommendations and associated provisions. One of the core steps of due diligence is

communicating how human rights impacts have been identified and addressed. The guidelines provide instruction on:

- The appropriate forms of communicating publicly and to impacted stakeholders
- Communicating relevant information that is also commercially sensitive

The Corporate Sustainability Reporting Directive (CSRD) has transformed the corporate reporting landscape by making these voluntary expectations mandatory. The CSRD, which is implemented by the European Sustainability Reporting Standards (ESRS), is bringing best practices into the mainstream and assuring higher standards of disclosure quality. There is increasing alignment and interoperability between these mandatory reporting requirements and the above-mentioned voluntary reporting standards.

- **The CSRD establishes EU-wide rules on corporate sustainability disclosure, which includes negative and positive impacts on affected stakeholders.** Disclosed information must be published in a company's sustainability statements of their management report, and those entities must seek external assurance on the full sustainability statements. The CSRD has a global scope, covering EU companies of all sizes and large companies operating in the EU.

Notably, the CSRD requires companies to report on impacts **up and down the value chain**. GenAI companies that are in scope must report on the impacts not only of their own operations but also those associated with the actions of other value chain entities. For instance, large EU-based foundation model developers, as well as large non-EU based foundation model developers doing business in the EU, would have to report on the upstream impacts of their data supplier choices, as well as the downstream impacts of both the application developers they sell their models to and the deployers of those systems.

- **The European Sustainability Reporting Standards (ESRS) outline the detailed rules for disclosure that companies subject to the CSRD are mandated to follow.** For example, draft guidance for section 4 of the ESRS notes that disclosures should include:
  - How actual and potential human rights impacts on consumers and end users arise from a company's strategy and business model, including its value proposition (e.g., providing online platforms with potential for online and offline harm, providing products that harm when overused, misused, or used as intended).
  - The disclosure may also explain whether these initiatives or processes also play a role in mitigating negative impacts disclosed (e.g., where a program that aims to support heightened digital awareness of the risk of online scams has led to a reduction in the number of cases of end users experiencing breaches of data privacy).

The ESRS also cross-references two other international standards described in this section. Section 3.4 of the ESRS states that the assessment of impacts that are the most important for the company and people—known as “material” impacts—should be informed by the “due diligence process defined in the international instruments of the UN Guiding Principles on Business and Human Rights and the OECD Guidelines for Multinational Enterprises.”



## The Benefits of Disclosure

Companies within scope of these regulations should not view these obligations as merely a compliance exercise, nor should companies that are not within scope feel that they should avoid disclosure. Companies have been making disclosures in alignment with voluntary standards like SASB and the GRI for decades due to the varied business benefits of doing so:

### › Meeting disclosure standards have required companies to focus their efforts and resources on producing reports.

This can improve various workflows within a company, including more effective risk management and tracking over time. For instance, the need to produce an externally scrutinized report may impose discipline on a company to gather and organize information in a manner that helps them prioritize key issues to address, establish baselines, craft mitigations, track progress, and guide resource allocation.

### › Creating disclosure workflows for genAI systems helps companies to prepare for external audits or evaluations.

This may include obligations under laws such as California's [AB2013](#), which requires AI developers to disclose a range of information about training datasets (e.g., the purpose of the datasets, whether they contain personal data, the number of datapoints in the datasets). It may also include voluntary, collaborative audits by third-party organizations, such as [Model Evaluation and Threat Research \(METR\)](#). By proactively establishing systems to track genAI's impacts, companies can be prepared with information ahead of regulatory mandates, make informed decisions, and efficiently respond to data requests. Adherence to voluntary reporting standards now can signal that a company is a leader in the space, while helping them establish robust systems to be as efficient as possible.

### › Disclosures can address investor and other key stakeholder concerns.

Investors and other stakeholders are increasingly expecting companies in the genAI value chain to address their human rights impacts. Disclosure can help companies to reassure investors and stakeholders that they are tracking the key issues and taking steps to protect human rights.

### › Disclosure about impacts to people helps to spur progress on disclosure across the AI industry.

Reporting makes more robust data available in the public domain. More disclosures offer other companies clear examples of how to disclose impacts, and can contribute to a more nuanced public understanding of what it means for genAI to impact people and society. Transparency about human rights by genAI companies can also engage concerned consumers, legislators and authoritative bodies in conversations about human rights that result in better coordination and more informed policymaking.








# 2. How GenAI Value Chain Entities Can Apply Disclosure Best Practices For Human Rights Impacts




## BSR's 10 Principles of Good Reporting

The following 10 Principles of Good Reporting are based on BSR's long experience working on sustainability reporting with companies. The principles also integrate features of several prominent reporting standards and frameworks, including the GRI and the SASB. While they are drawn from sustainability reporting, they are broadly applicable to disclosure on the human rights impacts of AI. For more information on the principles, please see the longer form report: [Five Steps to Good Sustainability Reporting](#).

### Report Content

 <p><b>PRINCIPLE 1</b> MATERIALITY AND CONCISENESS</p>	 <p><b>PRINCIPLE 2</b> STRATEGIC AND FORWARD LOOKING</p>	 <p><b>PRINCIPLE 3</b> SUSTAINABILITY CONTEXT</p>	 <p><b>PRINCIPLE 4</b> KEY PERFORMANCES INDICATORS AND NARRATIVE</p>	 <p><b>PRINCIPLE 5</b> COMPLETENESS</p>
---	---	--	---	--

### Report Quality

 <p><b>PRINCIPLE 6</b> STAKEHOLDER ENGAGEMENT</p>	 <p><b>PRINCIPLE 7</b> BALANCE</p>	 <p><b>PRINCIPLE 8</b> ASSURANCE</p>	 <p><b>PRINCIPLE 9</b> CONSISTENCY AND COMPARABILITY</p>	 <p><b>PRINCIPLE 10</b> CONNECTIVITY OF INFORMATION</p>
--	---	---	---	--

## Key observations on the application of disclosure best practices to AI companies

- Disclosures are a key human rights risk mitigation measure in the genAI value chain.** As described in the [Human Rights Assessment of the Generative AI Value Chain](#), value chain entities operate in relationships of enablement with each other. This means that they can “inherit” potential human rights impacts, and remedies for such impacts may require cross-coordination. Proactive disclosure can notify other value chain entities about inherited risks, thereby enhancing their ability to address them. For instance, a warning from a downstream developer about an increased risk of hallucination in certain domains can help system deployers avoid those high-risk use cases. Greater information flow across the value chain also enhances collaboration to mitigate risks, since it helps define the human rights risk areas that value chain entities may pool resources to address. For more information on remedy, please see the [Guide 8: Remedy for Generative AI Related Harms](#).
- Currently, foundation model developers have the most advanced disclosure practices, although these tend to be oriented toward a technical audience and rarely disclose impacts to people.** Foundation model developers have created standard practice disclosure formats, such as [system cards](#), [model cards](#), [datasheets for datasets](#), or [blogposts setting out training methods or research direction](#). However, these tend to be written for a technical audience and may be difficult for the average user to understand. Furthermore, disclosure of risks to people (e.g., human rights impacts) are rare, with frontier AI companies leading the way. See examples from [Microsoft](#), [OpenAI](#), and [Google](#). Best practices on disclosure have yet to emerge for suppliers, downstream developers, and deployers.
- The capacity to provide updated disclosure will likely significantly vary among genAI value chain entities, a trend that is also observable in digital companies more broadly.** Small dataset suppliers and application developers may feel they do not have the capacity to voluntarily disclose impacts. Given the diversity of company/entity types across the value chain, not all will be accustomed to disclosure. Some entities may also feel confused about what they should disclose. Disclosure about genAI systems may also have to be regularly updated, given the emerging landscape of risks, model adjustments, and rapid innovation. Disclosure will be a resource-intensive task that some value chain entities may struggle to fulfill. Automation may be a partial solution to the problem of updated disclosure.
- The format and depth of disclosure may vary, in part due to capacity limitations and transparency trade-offs mentioned above.** A well-resourced AI company that straddles multiple segments of the genAI value chain may choose to disclose human rights impacts and mitigations in long-form, in-depth Responsible AI reporting, as [Microsoft](#) and [Google](#) have chosen. These reports include in-depth analysis at the dataset, model, and application levels. Disclosure about genAI risks may also be present in corporate human rights reports (see, for example, [Meta’s report](#)). Foundation model developers may focus on specific model-level reporting, as [OpenAI](#), [Anthropic](#), and [Google DeepMind](#) have done with their [system cards](#). However, many value chain entities will not be so well resourced. Small application developers may release shorter disclosures about human rights impacts,

such as blog posts or technical specs on their website. Dataset suppliers may rely on data documentation templates, such as those that have been published by [Microsoft Research](#).

- **Disclosures can feel risky.** Companies or other entities that disclose more information, including about impacts on people, may suffer reputational damage, even though the increased transparency may reflect the effort the company/entity is putting into risk mitigation. Following several high-profile [lawsuits](#) against tech companies, companies may also feel that disclosing impacts to people may further expose them to litigation risk. The CSRD may reverse this dynamic somewhat by creating legal risk for companies that under-disclose. Furthermore, businesses have legitimate concerns that disclosing too much information may expose trade secrets and harm their competitiveness. AI companies may thus feel disincentivized from clearly and candidly disclosing risks.
- **Disclosures can provide too much information, thereby enabling bad actors or overwhelming a nontechnical audience; providing disclosures that are targeted to relevant stakeholders is therefore key.** Too much transparency about how a model works, such as details of classifiers or input filters, can enable bad actors to circumvent safety protections. Providing too much information can also overwhelm an audience. The average user is more likely to read a two-pager highlighting key risks of using a genAI system than to sift through an 80-page system card that contains dense technical information. That information, however, is likely to be useful for engineers or other technical stakeholders. While technical disclosures are useful for technical audiences, genAI disclosures in general should [include more nontechnical disclosures regarding risks to people](#) that are understandable by the average user to improve genAI safety literacy at scale.
- **It is important to focus on disclosing information that helps report readers understand how the company/entity is addressing their most salient (those that the company/entity has assessed to be the most severe and likely) impacts on people.** GenAI is a disruptive general purpose technology, which means that it is likely to impact all human rights. However, AI companies have limited space for disclosure, and should focus on reporting those impacts that are the most salient (see [Guide 3: A Human Rights-Based Approach to Impact Assessment](#)). The process of identifying salient impacts can vary widely in depth, including deep-dive impact assessments by well-resourced frontier labs, or more rapid analysis by smaller application developers based on internal metrics and user feedback. Regardless of the length or resource intensiveness of the assessment, a description of that process and the salient human rights impacts are key components of disclosure.
- **It is important to carefully select and build out “numbers and narrative” in disclosures.** The numbers refer to metrics and the narrative refers to explanations of why the numbers are important; some are more informative on human rights impacts than others. This combination of numbers and narrative (Principle 4 of the 10 Good Reporting Principles above) should clearly point to progress on human rights. One example of this is disclosure about the hallucination rate of a model and an explanation of what it means. From the quantitative standpoint, the hallucination rate is useful because hallucinations can be associated with adverse human rights impacts for any use case where accuracy is

important (e.g., legal advice). Less useful indicators include those that do not provide clear directionality on human rights impact. For example, a higher proportion of model outputs caught by a genAI application's output filters (as measured by toxicity score with output filters) may be because the model is outputting more toxic content, which may be negative for human rights. However, it could also be because the application's filters are becoming more accurate, which may be positive for human rights.

- **Disclosure will also be important for entities that are not a part of the genAI value chain but are adjacent to it, such as academic researchers and policymakers.** Disclosures address an imbalance in expertise between the private and public sectors on genAI. This can empower policymakers to create laws that are more responsive to the realities of the technology, thereby reducing the costs of regulatory compliance for businesses. Academics may use disclosures as a foundation for future research, which may help them contribute technical or safety breakthroughs in genAI.

## Disclosure Examples for Value Chain Entities

### Suppliers

The audiences for supplier disclosures will include investors who may wish to understand portfolio risk, as well as other value chain entities who use their datasets to train or finetune models. To properly understand the risk profile of using these datasets, developers and deployers require information about the composition, purpose, and potential risks of the datasets they are using to train their models. This will assist those value chain entities to anticipate downstream harms, such as compromised model performance due to lack of data relevance / inappropriate applications of data, or representational harms.

#### DISCLOSURE EXAMPLES

**Dataset documentation.** There are a variety of data documentation formats that serve different purposes while accommodating varying levels of capacity, resources, and expertise:

- Data statements provide essential information about the characteristics of datasets, including but not limited to the data collection methodology, curation rationale, and data sources.
- Datasheets for datasets provide a structured questionnaire that suppliers may fill in with details about model creation, composition, intended uses, maintenance, and other properties such as human rights impacts.
- Data nutrition labels are brief documents that utilize punchy visuals such as badges or quantitative metrics such as racial bias indices.
- Dataset cards provide concise, structured summaries of key dataset characteristics such as data collection methods, ethical considerations, potential biases, intended use cases, and limitations.

- Dedicated research papers are longer-form works that dive into different elements of datasets. These are more resource-intensive.

## Foundation Model Developers

---

Disclosures by foundational model developers will have a broad audience:

- Suppliers may need to understand what and how their datasets or annotated data will be used.
- Application developers will need to know sufficient information about the models to appropriately design applications around them.
- System deployers need to know the capabilities and risks of the models that power the systems they deploy.
- There is a broad public interest in foundation model transparency, given that foundation models are an advanced technology that may have significant human rights impacts. Users will want to know how the underlying technology works when interacting with it, or for their own general curiosity and GenAI literacy.
- Value-chain adjacent entities such as academics and policymakers. Regulators may rely on GenAI disclosures to ensure that their policymaking is informed by technical realities and a comprehensive understanding of risk. Academics may use the disclosures as a foundation for further research.

### DISCLOSURE EXAMPLES

#### Model and system cards:

- A model card contains core information about a model, including model details, training processes, evaluation metrics, and safety / human rights considerations.
- A system card is similar, but applies to an entire AI system, which includes the model and other components, including the dataset and the application that is built on top of it. System cards are generally produced by businesses that are both foundation model and application developers.
- While the type of information that is important to disclose may vary by the model and use case, all model or system cards should include an overview of salient human rights impacts. Frontier AI labs such as OpenAI, Google DeepMind, and Anthropic have developed many examples of this type of disclosure.

**Responsible use guides:** Responsible use guides by foundation model developers are a combination of important information about the model and tailored instructions to developers on how to utilize it safely. Meta's LLaMa Responsible Use guide provides a useful template for the types of information that should be disclosed by foundation model developers.

**Disclosures for a nontechnical audience:** Provided that these are communicated in clear and accessible language that is easily understood by nonspecialist audiences, they can adopt a variety of formats, such as:

- Explanatory blog posts that accompany model release (see example from [Anthropic](#)).
- Nontechnical “executive summaries” of complex technical language. For instance, a model developer may release a system card that uses technical safety terminology, and an executive summary version that has been translated into easily understandable language.
- Model cards written in simple language (see an example for Google’s [family of foundation models](#) fine-tuned for medical use cases).

## Downstream Developers

---

Downstream developer will also have a broad audience for their disclosures:

- Suppliers of datasets for fine-tuning will need to know the application has been fine-tuned.
- Foundation model developers will want to know how their model has been used for a range of reasons, including whether their terms of service have been met or to provide feedback on application design.
- Users will want to know the potential impacts of using the application across different contexts, including to them and to other users.

The format and content of disclosure should differ for each audience category; users might want less technical information, for instance.

### DISCLOSURE EXAMPLES

#### Nontechnical disclosures for users

- A core element of disclosure for downstream developers will be intended use cases, misuse cases, and limitations for their application. For instance, developers may disclose to users that an image generation application is intended for creative and personal use, but that they should not use it to attempt to breach copyright.
- Disclosures should also include a summary of salient human rights impacts, as well as a description of how the developer identified and mitigated them. This does not have to be phrased in human rights language, although developers may rely on international human rights standards to ensure their list of salient impacts is appropriately comprehensive.
- Quantitative indicators for developers may include rates of hallucination. These indicators can also be use case-specific. For example, a genAI image recognition tool for identifying anomalies on security footage may indicate how often anomalies were accurately iden-

tified during testing and post-deployment. Developers can explore using automation to refresh some quantitative indicators to reduce the cost of providing updated metrics. An example of a qualitative indicator may be the progress of alignment of internal AI governance systems with regulatory requirements, such as under the EU AI Act.

- The format and means of distribution of this information may vary, including blog posts, research papers with nontechnical executive summaries, or one-pager infographics with salient human rights indicators.

### **Disclosures for other value chain entities, policymakers, and auditors.**

- Upstream actors may benefit from quantitative indicators such as hallucination rates in deployment contexts, robustness against adversarial prompts, thresholds set by the company, or output filter performance. Qualitative indicators may include summaries of user complaints or feedback from external stakeholder engagement. These disclosures may be included in a separate section in the disclosures above, or they can be in a separate document that is selectively shared with upstream actors to collaboratively iterate on model safety and human rights impacts.

## **Deployers**

---

The primary audience for deployer disclosures will be the operators of the system, as well as those who will be impacted by the system (i.e., affected stakeholders). For instance, a genAI tool that is integrated into a government legal aid department should disclose information relevant for the government workers who use the tool, as well as for actual and potential users of legal aid services.

Upstream actors, especially downstream developers and foundation model developers, may also use the information disclosed by deployers to improve their models and applications.

### **DISCLOSURE EXAMPLES**

#### **Nontechnical disclosures for users**

- A deployer’s disclosure about a specific product, feature, or service integration should at minimum, provide:
  - » Key details about how the product, feature, or service works
  - » Disclosure that genAI has been integrated into the product, feature, or service (if relevant)
  - » A list of intended and prohibited use cases
  - » Identified impacts to people
  - » How the deployer is mitigating risks, including institutional governance processes such as advisory boards or the provision of training and usage guidelines
  - » Key qualitative and quantitative indicators



- Format choices may be flexible, provided the core elements listed above are included. For instance, the deployer may release a detailed PDF document or an interactive website that contains the above information.
- Disclosures by deployers should be pitched at a technical level appropriate to the audience. This may sometimes include technical language, such as when a genAI tool is deployed in a specialized enterprise context (e.g., as coding assistance for software engineers or drug development). In other cases, it should be phrased in simple language to be comprehensible by operators and affected stakeholders.

# 3. Key Resources

The following resources contain more detailed information on transparency and disclosure best practices on sustainability, human rights, and genAI.

- **The International Financial Reporting Standards (IFRS) Foundation's ISSB S1** sets out the general requirements for what and how companies should disclose sustainability-related financial information.
- **The European Sustainability Reporting Standards (ESRS)** outline the detailed rules for disclosure that companies subject to the Corporate Sustainability Reporting Directive (CSRD) are mandated to follow.
- **OECD Due Diligence Guidance for Responsible AI** (forthcoming) builds on **OECD Due Diligence Guidance** to provide guidance for companies developing and using AI, including on communicating how the impacts of AI are addressed.
- **BSR's Guide on the Five Steps to Good Sustainability Reporting** details five steps any company can take for the disclosure of sustainability impacts, supported by 10 principles on content and quality that characterize best-in-class sustainability reporting. The principles combine and draw on features of several prominent reporting standards and frameworks mentioned in this guide.
- **Microsoft Research's Data Documentation project** contains categories of information that dataset creators should disclose regarding their datasets that may be decision-useful for dataset consumers.
- **The Global Internet Forum to Combat Terrorism (GIFCT)'s paper on Transparency Reporting: Good Practices and Lessons from Global Assessment Frameworks** provides an overview of key themes, good practices, and lessons on transparency drawn from a variety of sustainability frameworks, along with academic, governmental, and civil society reporting.
- **BSR's blog post on A Coherent Approach to Sustainability Due Diligence and Reporting: Making Sense of CSDDD and CSRD** sets out how companies may take a coherent approach to complying with the transparency-related requirements of the CSDDD and CSRD, two seminal pieces of upcoming EU legislation.



BSR™ is an organization of sustainable business experts that works with its global network of the world's leading companies to build a just and sustainable world. With offices in Asia, Europe, and North America, BSR™ provides insight, advice, and collaborative initiatives to help you see a changing world more clearly, create long-term business value, and scale impact.

**[www.bsr.org](http://www.bsr.org)**

Copyright © 2025 by Business for Social Responsibility (BSR)

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.