

# Remedy for Generative AI-Related Harms

Guide 8 of the Responsible AI Practitioner Guides for Taking a Human Rights-Based Approach to Generative AI

February 2025



BSR®

# Summary

This paper provides guidance about how to take a human rights-based approach to remedy for the harms linked to the development, deployment, and use of generative AI (genAI). It includes the following sections:

- 1 The Human Rights Foundations for Remedy:** A high-level overview of the theoretical basis for remedy for human rights impacts, covering the extent of the duty to provide remedy, the five elements of remedy, and eight components of an effective grievance mechanism.
- 2 The Landscape of Remedy for GenAI-Related Harms:** An application of remedy's theoretical principles with real-world concerns, covering the challenges with providing remedy for genAI products and key observations on remedy across the genAI value chain.
- 3 Examples of Remedy and Cooperation in the Remedy Ecosystem:** Illustrative examples of remedy provision by value chain entities, as well as explanations of how value chain entities may cooperate with the broader remedy ecosystem.
- 4 Key Resources**

## Key Points

- Remedy for genAI's harms will often require action by multiple actors, including businesses in the genAI value chain, nonprofit organizations, and public entities. This is known as the "remedy ecosystem." Direct remediation by businesses is important, but not the only tool in the human rights-based remedy toolkit.
- Where direct remediation by businesses is required, and that remedy requires collaboration across the genAI value chain, the duty to coordinate that remedy should lie with a single point of contact. This single point of contact should coordinate different elements of remedy across the value chain and communicate those actions to impacted rightsholders.

### 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](#)

- The single point of contact will usually be the value chain entity that is directly interfacing with the impacted rightsholder. Examples include the developers of direct-to-consumer genAI apps, the deployers of business-to-business genAI solutions, or deployers that integrate genAI tools developed by other companies into their own systems or products.
- Given the speed and scale of potential harm from genAI systems, many harms will be cumulative or societal in nature, such as the harms related to disinformation or job loss. Effectively addressing these will necessitate an ecosystem approach that relies on cooperation between state and non-state actors, including courts, governmental bodies, and private or nonprofit actors.
- Entities from the broader remedy ecosystem—specifically law enforcement and the courts— must sometimes play a role in effective remedy. In these situations, value chain entities should cooperate with those separate processes as required.

# 1. The Human Rights Foundations for Remedy

Principle 29 of the [UN Guiding Principles on Business and Human Rights \(UNGPs\)](#) states that businesses must provide effective remedy for individuals whose rights have been impacted. In simple terms, this means that businesses who harm people must try to return those affected to the same or equivalent position before the adverse impact occurred.

Under the UNGPs, the extent of the connection between the human rights impact and the company's actions determines the extent of their obligation to provide or cooperate in remedy.<sup>1</sup> Businesses must directly provide or cooperate in remedy when they **cause** or **contribute** to a harm. Businesses are not required to directly provide for or cooperate in remedy if they are **directly linked or not linked** to a harm, although they may choose to do so. For more information on how to determine the extent of connection to an impact, see [Guide 4: A Human Rights-Based Approach to Risk Mitigation](#).

Businesses should provide remedy for human rights harms not just because it is the right thing to do; it is often also an effective strategy for mitigating risks to the business. People who are unable to settle their grievances directly may resort to alternatives that create business risk, including public shaming campaigns or lawsuits. The cumulative price of negative press or sizable damages may outweigh what it would have cost the business to proactively manage rights violations. Furthermore, some remedy mechanisms, such as user reporting channels, are a useful source of business intelligence on how their genAI products are impacting people in the real world.

## The Five Elements of Remedy and Effectiveness Criteria

[International standards](#) provide for five elements of remedy. Remedies may cover one or more of these elements (the "examples of remedy" section below illustrates what this might look like in practice across the genAI value chain).

<sup>1</sup> The technical term for this concept under the UNGPs is "attribution."

1. **Satisfaction** involves ceasing the violation, acknowledging the harm, disclosing the truth, providing an apology, and sanctioning those responsible. One example is publishing a blog post detailing recent trends of product misuse and what the entity has done about it.
2. **Restitution** entails restoring, to the extent possible, whatever has been lost and returning the rightsholder to the state before the harm occurred. One example is the restoration of an account disabled because the deployer mistakenly believed the user breached their acceptable use policy.
3. **Guarantees of non-repetition** entail changes to policies and procedures to prevent future harms, or the taking of disciplinary action. One example is the adjustment to a privacy policy to cover a new vector for data breaches.
4. **Rehabilitation** may include medical, psychological, legal, and social services to restore the victim. Company-based grievance mechanisms may be less equipped to provide rehabilitation than other entities in the remedy ecosystem, such as nonprofits.
5. **Compensation** includes money or other benefits, where damage can be financially assessed. One example may be cooperation with a legal ruling for damages.

Remediation processes must be effective. Principle 31 of the UNGPs sets out eight criteria for an effective company-based grievance mechanism:

1. **Legitimate**—enabling trust from the stakeholder groups for whose use they are intended, and being accountable for the fair conduct of grievance processes.
2. **Accessible**—being known to all stakeholder groups for whose use they are intended, and providing adequate assistance for those who may face particular barriers to access.
3. **Predictable**—providing a clear and known procedure with an indicative time frame for each stage, and clarity on the types of process and outcome available and means of monitoring implementation.
4. **Equitable**—seeking to ensure that aggrieved parties have reasonable access to sources of information, advice, and expertise necessary to engage in a grievance process on fair, informed, and respectful terms.
5. **Transparent**—keeping parties to a grievance informed about its progress, and providing sufficient information about the mechanism’s performance to build confidence in its effectiveness and meet any public interest at stake.
6. **Rights-compatible**—ensuring that outcomes and remedies accord with internationally recognized human rights.
7. **A source of continuous learning**—drawing on relevant measures to identify lessons for improving the mechanism and preventing future grievances and harms.
8. **Based on engagement and dialogue**—consulting the stakeholder groups for whose use they are intended on their design and performance, and focusing on dialogue as the means to address and resolve grievances.

# 2. Remedy For Generative AI-Related Harms

## Challenges

While remedy can be challenging for companies in all sectors, there are unique difficulties to remedying harm for related genAI products. Some of these challenges are applicable to remedy provision in the technology sector more broadly, including some lessons from providing remedy for social media companies.

- › **The speed and scale at which harm from genAI systems may spread. ChatGPT alone reported 200 million monthly active users in August 2024.**

If only 1% of these users is impacted and requires a remedy, that is still millions of remedy-seekers for one genAI product. This scale makes it challenging to identify which individuals have been harmed, make remedy accessible to all of them, and effectively engage with them to improve remedies.

- › **Given the speed and scale of potential harm from genAI systems, many harms will be cumulative or societal in nature, such as harms related to disinformation or job loss.**

These may be difficult to address. For example, it is challenging to remedy individuals for impacts to their right to vote related to AI-generated disinformation because it is difficult to clearly articulate the harm and to establish the connection between that harm and the disinformation. Effectively addressing these will necessitate an ecosystem approach that relies on cooperation between state and non-state actors, including courts, governmental bodies, and private or nonprofit actors.

- › **The complexity of genAI systems makes it challenging to identify the connection between value chain actors and the eventual harm.**

One example of this is AI's "interpretability" problem, which means AI systems are currently too complex for anyone to fully understand the reasons behind their decision-making process. This can create ambiguity over which value chain actors have caused, contributed to, or are directly linked to a harm. This ambiguity may be used by companies to "point the

finger” at other value chain actors or to justify holding no actors responsible. For example, a downstream developer might claim that a harmful output was due to a defect at the model level and is therefore the fault of the foundation model developer, while the foundation model developer might blame the downstream developer’s poor safeguards in application design. This complexity also makes it more challenging for users to understand the sources of harm.

› **There may be resource and capacity imbalances between value chain actors.**

For instance, a large foundation model developer likely has significantly more money and human resources than vendors supplying data or the third-party downstream developers who use their models. This may create obstacles to coordinating remedies that require action across the value chain. For instance, a foundation model developer might require third-party dataset vendors to implement rigorous auditing processes to ensure data quality and fairness, but these smaller vendors may lack the technical expertise, financial resources, or manpower to conduct such audits effectively and at scale.

## Key observations on remedy across the genAI value chain

› **Remedy for genAI-related harms often requires ecosystem-wide action, but the duty to coordinate remedy for harms that require business action should generally lie with a single point of contact.<sup>2</sup>**

A single point of contact ensures that remedies are accessible and adequate for impacted rightsholders and prevents “finger-pointing” by value chain entities that may frustrate the effective resolution of grievances. This single point of contact should coordinate the different elements of remedy that are needed across the value chain. For instance, remedy for AI-generated hate speech may require updating the application’s safety filters, new fine-tuning of the foundation model, and the cleaning of datasets of harmful content and associations.

› **The “single point of contact” approach aligns with the principle that remedy should be grounded in the needs of affected stakeholders.**

Expecting users to understand the complexity of the technical mitigations described in the previous paragraph or approach value chain entities individually would be onerous and unrealistic, frustrating the requirements that remedies be accessible and affordable for users. The point of contact should communicate its coordinative actions to impacted rightsholders. For more information about designing remedy to meet the needs of affected stakeholders, see the [UN B-Tech Project’s guidance on company-based grievance mechanisms](#).

<sup>2</sup> Known in human rights terminology as “operational-level grievance mechanisms.” For more on this, please see: <https://www.ohchr.org/sites/default/files/Documents/Issues/Business/B-Tech/access-to-remedy-company-based-grievance-mechanisms.pdf>.

› **The single point of contact will usually be the value chain entity that is directly interfacing with the affected stakeholder.**

Examples include the developers of direct-to-consumer genAI apps or the deployers of business-to-business genAI solutions. The existing interface between the value chain entity and impacted rightsholder provides a useful surface for providing remedy, including through in-app features such as reporting channels, existing lines of communication, and terms of service. There may be exceptions to the directly interfacing rule, such as when the impacts pertain to the dataset (e.g., intellectual property violations or data collection without consent). In such cases, the point of contact should be the owner of those impacts, which—in the case of the dataset—would be the supplier.

› **There will be cases, such as those that involve criminal law, where entities from the broader remedy ecosystem, such as law enforcement or social services organizations, must play a role in effective remedy.**

In these situations, value chain entities should cooperate with or facilitate those separate processes as required. This may include the provision of information or the removal of hosted or dataset content, such as copyrighted content, where applicable. This cooperation should be subject to industry standards regarding law enforcement requests and respect for human rights. Cooperation may also include proactive partnerships, such as with organizations such as the National Center for Missing and Exploited Children (NCMEC) that specialize in identifying victims of child sexual abuse and exploitation. In other cases, value chain entities may simply direct users to the appropriate organizations, such as providing contact details for social services providers. More detail on the remedy ecosystem is provided in the last section of this brief.

› **Remedies for societal impacts are more complex than for individual harms and may require proactive action from multiple actors.**

It may be unrealistic for an entity such as a three-person app developer to coordinate remedy for societal harms, such as those related to election or climate-change misinformation. Corporate entities may also be less likely than NGOs, academic centers, or public bodies to have the expertise or networks to address societal-level issues. In such cases, those entities should take the initiative, with genAI value chain entities cooperating to provide assistance as needed, including information disclosure, participation in stakeholder engagement, or product and policy changes.

› **Remedies can rarely guarantee non-repetition of genAI-related harms.**

GenAI outputs are based on probabilistic algorithms, which means that even with strong controls, a model may still randomly output harmful content. Occasional mistakes do not indicate that a remedy has been ineffective, similar to how a content moderation system with a 0.1% failure rate across all categories of content would be considered highly effective. Value chain entities should focus on demonstrating that they have made best efforts to address risks by enacting mitigations to ensure that harms are rare, and on providing effective remedies when harms do occur.



› **Companies should engage with affected stakeholders to improve the effectiveness of remedies.**

There are a variety of mechanisms that can facilitate stakeholder engagement on remedy. For instance, engagement on a user-reporting mechanism might include in-app features for user reviews of the mechanism, manually following up with users who have used the mechanisms, or proactively soliciting their views during mechanism design. The information should then be integrated into the design of the reporting mechanism. Ultimately, the users of remedy systems are the best judges of whether it is effective and what it should look like.

# 3. Examples of Remedy and Cooperation In The Remedy Ecosystem

## Examples of Remedy for Value Chain Entities

The following tables provide concrete examples of how genAI value chain entities could provide effective remedy for genAI-related harms. The tables begin with an overview of how each value chain actor could participate in remedy. Next, the “remedy examples” cells describe how each entity could take actions that cover different categories of remedy. Not every category of remedy needs to be covered by every value chain entity. For instance, not many corporate entities are equipped to handle rehabilitative remedies, which will generally be better handled by other remedy ecosystem entities, such as nonprofits or government bodies. The examples are also calibrated to fulfill the eight effectiveness criteria listed in [Principle 31 of the UNGPs](#).

### Suppliers

Supplier mechanisms for providing remedy will revolve around due diligence and mitigations for their key value creation workflows, such as dataset cleaning or adjustments to annotation instructions. Supplier remedies may also involve clear communication and collaboration with downstream actors, such as foundation model developers who purchase their datasets.

#### REMEDY EXAMPLES

- **Satisfaction:** A supplier includes simple, easy-to-use channels to flag and remove personal data from the datasets they use. For instance, suppliers can create an online form on their site for users or foundation model / downstream developers to submit complaints alleging that personal data has been included.
- **Restitution:** Upon receipt of a data removal request, the supplier complies and notifies the impacted stakeholder.

- **Compensation:** A supplier provides direct compensation, or cooperates in class action settlements, for the inclusion of copyrighted material in a dataset.
- **Rehabilitation:** N/A
- **Guarantees of non-repetition:** Following repeated complaints that their datasets include personal data, a vendor creates and publicizes a new method for procuring and cleaning data.

## Foundation model developers

---

Model-level interventions, such as safety fine-tuning or dataset filtering, will be key mechanisms for providing remedy available to foundation model developers. Given the probabilistic nature of genAI algorithms, such mechanisms may reduce the likelihood of, but cannot guarantee, further human rights harm. Another important element of effective remedy for developers is directly engaging with users, such as through public communications or cooperation with entities in the remedy ecosystem (e.g., courts).

### REMEDY EXAMPLES

- **Satisfaction:** A developer publishes a blog post detailing how it identified an error in its image-generation model, which was disproportionately generating oversexualized images of women. The blog post acknowledges the error, attributes its discovery to user feedback, and sets out a plan to address it. The plan includes engagement with suppliers to source datasets that are less likely to create oversexualization risk, such as more diverse female representation or refraining from comingling sexualized images with those of women in the same dataset.
- **Restitution:** A model was generating incorrect information about an upcoming election's candidates. The developer fine-tunes the model with reinforcement learning by human feedback until it reliably outputs correct information and publishes a correction blog post.
- **Compensation:** Developers negotiate a settlement in a lawsuit alleging the use of copyrighted material to train their models.
- **Rehabilitation:** In addition to existing remedies for harms, users are redirected to organizations that may provide services such as counseling or legal aid.
- **Guarantees of non-repetition:** A developer researches and publicizes a new safety fine-tuning method that reduces hallucinations by 20%. While this cannot guarantee that there will be no more hallucinations, it makes recurrence significantly less likely. Part of the method involves new data annotation instructions, which the developer publicizes on its website for upstream actors to use.

## Downstream developers

---

Downstream developers may build **business-to-business** genAI products on top of foundation models, such as workflow automation software or data analytics tools, for the purposes of

sale to deployers. Developers may effect remedy through model fine-tuning, product feature design, or coordination with the deployer. This remedy may be informed by feedback from the deployer or from affected stakeholders.

Downstream developers may also build **direct-to-consumer** products on top of foundation models, such as genAI chatbots or video editing apps. In these cases, the downstream developer will also be the deployer. These entities are covered in the next section.

The examples below cover remedy for downstream developers of **business-to-business** products, with direct-to-consumer examples listed in the next section.

## REMEDY EXAMPLES

- **Satisfaction:** Satisfaction: Users who work at the deployer receive emails notifying them that developers updated the app’s content policy. This update is a direct response to user feedback collected by the deployer about policy gaps that made the chatbot vulnerable to jailbreak prompts.
- **Restitution:** A law firm deploys a genAI tool that helps users organize and retrieve their files more efficiently. Following feedback over privacy concerns, the developers of the tool ship a new feature that allows users to easily delete all conversation history with the tool. This feature also prevents the input and output data from being used to train the model further.
- **Compensation:** Developers of a genAI video-generator pay part of a damages settlement alongside deployers when they utilize the app to imitate copyrighted works.
- **Rehabilitation:** A telehealth provider deploys a genAI chatbot to assist users with preliminary medical inquiries. Due to flaws in the chatbot’s responses, some users receive incorrect medical advice that leads to adverse health effects. The deployer offers affected users free access to certified healthcare professionals for proper diagnosis and treatment, and the developers create in-app notifications to communicate this remedy to users.
- **Guarantees of non-repetition:** Developers of an informational chatbot deployed in the public sector create reporting channels for users to directly report issues to the developer. Developers use that feedback to inform updates to their prompt / output filters that prevent further harmful content from being generated. The downstream developer may also swap their choice of foundation model to power their product if the initial foundation model developer is incapable or unwilling to address human rights impacts.

## Deployers

Deployers are often best placed to provide remedy because they are the closest to the application context and can thus fashion appropriate remedies. Furthermore, they have a direct interface with affected stakeholders for implementing remedy, such as through existing user or customer relationships, or in the case of public entities deploying genAI, through their role as the provider of public services. In the case of deployment to users who are employees, the deployer has a direct interface through their role as employer.

**Developers of direct-to-consumer products are also often the deployers**, and therefore have direct interface with users. App interfaces and other channels of communication, such as user accounts that may be linked to e-mail addresses, create multiple opportunities to provide remedy. Some examples include in-app reporting channels and dashboards or automated email updates regarding complaints.

If the downstream developer is also the foundation model developer and deployer, such as OpenAI, Anthropic, and Google DeepMind and their respective chatbots, they will be best placed to effectively provide remedy. These entities will have some insight into how decisions at the model and application levels may be connected to downstream harms, as well as the ability to make corresponding changes at both levels. Furthermore, entities that both develop models and applications are likely to be highly resourced, increasing their ability to design and maintain remedies that are accessible, adequate, and timely for users. Note that within these organizations, further coordination may be needed between the teams that build the foundation model and those that develop applications to provide effective remedy.

The examples below cover scenarios of the deployment of **direct-to-consumer** and **business-to-business** products.

### REMEDY EXAMPLES

- **Satisfaction:** An app developer creates and deploys a genAI tool that helps users determine their eligibility for social security benefits. Users from a certain demographic are systematically told by the app that they are ineligible because the genAI tool incorrectly hallucinates disqualifying factors. Realizing this, the developer sends apology letters directly to the affected users and commits to fixing the mistake. They promptly fine-tune the AI model to correct the error and keep the users informed about the progress and improvements made.
- **Restitution:** An affected stakeholder is erroneously denied a social security benefit due to an error made by a genAI tool deployed by the administering public body. The stakeholder has their payment, with applicable interest, reinstated by the body.
- **Compensation:** A grocery supplier's genAI planning tool commits errors that result in delayed inventory delivery for a caterer. The supplier compensates the caterer for lost revenue.
- **Rehabilitation:** A business deploys a genAI-powered tool that outputs discriminatory content to female employees. The company pays for therapists to provide free counseling to affected employees.
- **Guarantees of non-repetition:** An entity that has deployed an error-prone workflow genAI tool halts its use.

## Cooperation with the Remedy Ecosystem

Businesses are not the only entity that have responsibility for remedy—the UNGPs promote an **ecosystem-based** approach. The ecosystem includes the following categories of state and non-state actors:<sup>3</sup>

- **Judicial:** Judgments through lawsuits or decisions made by national, regional, or international courts. Remedies for serious human rights violations, such as child sexual exploitation, may require involvement of police and the courts. Examples of such pathways to remedy include:
  - Article 82 of the **EU General Data Protection Regulation (GDPR)** provides for the right to compensation for damage suffered, including potentially from data processing related to genAI.
  - **U.S. Copyright Law** provides a comprehensive legislative regime for violations of intellectual property rights. Numerous lawsuits have already been initiated against frontier AI companies for alleged breaches of these rights.
  - Affected stakeholders may also be able to access **sector-specific legislation** for the use of genAI by certain companies, such as applicable laws that govern the healthcare or legal services.
- **Governmental:** This is a broad category that includes mediation services by government bodies, decisions by public bodies such as consumer protection agencies, and legal aid to pursue judicial remedies. This may also include regulatory action under laws such as the EU AI Act, where developers and deployers may be fined for failing to comply with a range of obligations based on a predetermined “risk level.”
- **Nonprofit:** Nonprofits may collectively represent or advocate for stakeholder interests to value chain entities. They may also provide social, legal, or health services for impacted rightsholders (e.g., targets of hate speech).

When a value chain entity is involved with human rights harm but cannot fully remediate that harm alone, it must cooperate in the provision of remedy by actors in the broader remedy ecosystem. For instance, a foundation model developer may prevent a synthetic child sexual abuse image of a real person from being regenerated by their model. However, the image may continue to spread online, and the subject depicted in the image may require counseling or legal advice. These negative impacts cannot be addressed by businesses alone and will require remedy from entities such as courts or social services providers. Value chain entities should cooperate to ensure comprehensive and more efficient remedies.

While the UNGPs do not provide definitive guidance on all the forms that this cooperation should take, value chain entities may consider the following pathways:

<sup>3</sup> The following is a simplified categorization. The UNGPs divide remediation mechanisms into three types: “judicial mechanisms,” “state-based non-judicial mechanisms,” and “non-state-based grievance mechanisms.” For explanations of each type, please see: <https://www.ohchr.org/sites/default/files/Documents/Issues/Business/B-Tech/access-to-remedy-concepts-and-principles.pdf#page=4>.

- **Information provision and sharing:** Providing information to public entities, such as regulators, can help address an imbalance in expertise between the private and public sectors on genAI. There is also a medium-term business case for this collaboration, which may empower regulators to create legislation that is more responsive to the realities of the technology, thereby reducing the costs of regulatory compliance. Knowledge provision may take a variety of forms, including training, joint research projects, technical collaboration, and secondment programs. One model for this is the [collaborations between frontier AI labs and AI Safety Institutes](#).

Value chain actors should also be open with users about the human rights implications of the product or service they provide. This can take the form of transparency reporting on the content, risks, and use cases of the components of AI systems. Examples include [datasheets for datasets](#), [model cards for model reporting](#), and responsible use guides (such as Meta’s guide for [developers using Llama](#)). This transparency will empower users to understand when and why their human rights have been impacted, which will also assist them to understand when and how to engage with the broader remedy ecosystem. Public or nonprofit actors may also use that public information to enhance their understanding of remedy. For example, a judge overseeing a copyright dispute may read a model card to understand the technology so they can render a more informed decision. For more information on transparency and disclosure, see [Guide 7: Aligning Transparency and Disclosure Practices with Human Rights Responsibilities](#).

GenAI value chain entities can also share information with each other and with governments to tackle coordinated harms. Many threats in the cybersecurity, disinformation, or fraud domains are orchestrated efforts that require a unified response. For example, [OpenAI partnered with Microsoft Threat Intelligence](#) to disrupt five state-affiliated actors that sought to use genAI services in support of malicious cyber activities. As frontier models become increasingly powerful, state actors may try to [steal them](#) through cyberattacks or espionage. This may necessitate greater cooperation between genAI value chain entities and governments to protect national security.

- **Multi-stakeholder collaboration:** GenAI value chain entities can also actively cooperate in the broader remedy ecosystem through multi-stakeholder participation. This may involve standard-setting that makes authoritative rules for providing remedy or the coordination of collaborative remedies. For example, a foundation model developer might become a member of the [Frontier Model Forum](#), which supports standards development on a range of remedy-relevant issues, such as [AI safety frameworks](#). Multi-stakeholder collaboration can also result in collaborative remedies. Certain elements of genAI, such as provenance interventions for synthetic content, often benefit from a critical mass of organizations agreeing to a set of standardized practices. One example of this is the [Partnership on AI’s Responsible Practices for Synthetic Media](#), which numerous AI companies have signed up to.
- **Redirection of users to entities that can provide rehabilitation:** Organizations outside of the value chain, such as nonprofits or healthcare providers, are better suited than most AI companies to address health or legal impacts. GenAI value chain entities can collaborate in the remedy ecosystem by redirecting users to organizations that provide relevant reha-

bilitation services. This would require due diligence to ensure that the organizations they are directing people to are qualified and have the appropriate license to dispense legal, health, or social services. GenAI value chain entities can employ a variety of methods for this, such as including model outputs directing users to a suicide hotline in response to prompts indicating a willingness to self-harm, or including information in its content policies or terms of use.

- **Funding entities that can provide rehabilitation:** GenAI value chain entities can also consider providing funding to ecosystem actors that can provide remedies that they cannot, such as NGOs or social services organizations. They may prioritize entities for funding to address risks that are especially salient or for which they cannot effectively mitigate alone. For instance, image or video model developers may choose to fund organizations that support targets of synthetic nonconsensual intimate imagery.



## 4. Key Resources

The following resources contain more detailed information about how to take a human rights-based approach providing remedy for genAI harms:

- **UN B-Tech Project's [paper on access to remedy in the technology sector](#)**—An authoritative resource on the human rights principles behind providing remedy in the tech industry. The key points in this paper formed part of the foundation for this guide.
- **UN B-Tech Project's [paper on the remedy ecosystem in the technology sector](#)**—An authoritative resource on the various actors that comprise the remedy ecosystem in the tech industry. The key points in this paper formed part of the foundation for this guide.
- **UN B-Tech Project's [paper on designing and implementing effective company-based grievance mechanisms](#)**—An authoritative resource on how value chain entities can design grievance mechanisms. The key points in this paper formed part of the foundation for this guide.
- **BSR's [Guide on Access to Remedy](#)**—An examination on the various components of remedy mentioned in this guide, generalized beyond the tech sector.
- **OECD [Due Diligence Guidance for Responsible AI](#)** (forthcoming)—A resource that builds on [OECD Due Diligence Guidance](#) to provide guidance for companies developing and using AI, including on remedy for AI-related harms.



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.