CASE STUDY

SaaS Case Study

HRDD Across the Technology Ecosystem

SEPTEMBER 2022

This case study on a hypothetical Software as a Service (SaaS) company demonstrates how the technology ecosystem mapping tool can be used to understand how exploring connections across the ecosystem can enhance human rights due diligence.

1. CONTEXT

BACKGROUND

XRLY is a SaaS Service Provider that delivers enterprise solutions, such as customer relationship management (CRM) and human capital management (HCM) services. We sell our solutions through distributors and resellers who customize the services before selling them to customers. One of our primary customers is a retail company–we provide CRM solutions for their sales and marketing teams, and HCM solutions for the Human Resources team. Accordingly, we consider the end user to be employees at the retail company. The rightsholders (and potentially impacted parties) include employees as well as the customers and communities of the retail company.

XRLY was prompted to examine our human rights impacts after customer inquiries into our business practices. We seek to identify:

- The key human rights issues we need to focus on
- What other tech ecosystem actors are involved in our impacts and how we can work with them
- The activities we need to undertake to meet our responsibility to respect human rights

This exercise is being led by Jane Doe, Director of the Sustainability Team. The Sustainability Team is collaborating with staff from other teams, specifically sales, security, compliance, and engineering, because we understand that different teams have different insights into risks and use cases.

2. HUMAN RIGHTS DUE DILIGENCE QUESTIONS

This section explores the key questions and human rights issues for enterprise software developers to consider as prompted by the HRDD questions on the Summary slide.

WHAT IS THE SERVICE DESIGNED TO ACHIEVE?

What are the use cases? We provide CRM solutions for retail company sales and marketing teams. Our services help employees on these teams store contact information of partners and prospects, identify sales opportunities, and manage marketing campaigns.

We also provide HCM solutions for the HR team. These software solutions help the Human Resources team with administrative processes and workforce engagement programs like hiring, tracking, evaluation, and retention.

What are the actual and potential human rights impacts of the use cases? We know that SaaS services can be (mis)used in ways that adversely impact human rights when:

- We have no control or insight into how the end-user is customizing and deploying our software.
- User data collection is extensive and not properly secured, or collected without the informed consent of users, and/or shared with third parties for our financial gain. This data may tell very personal stories about individuals that could be utilized for government surveillance, harassment, denial of services and opporuntities, and discrimination.
- Decisions are made using artificial intelligence algorithmic models that are opaque or not explainable or contestable.
- User-generated content is moderated through algorithmic systems that may erroneously over- or under-moderate content.

The use of XRLY's SaaS services in these ways could be connected with adverse human rights impacts to freedom of expression, freedom of movement, freedom of association and privacy. We sense the likelihood of our retail customers using our services in this way is low but nonetheless acknowledge there is some potential for harm.

HOW IS THE SERVICE USED?

What insight do we have into how and where our software is being used? Reports from our resellers provide high-level insights into how our software is being used. The sales and marketing teams use our services to support their engagement and revenue goals. The human resources team uses our service to collect and organize sensitive employee information and records. Our reseller partners learn about intended uses during their pre-sale conversations as they inquire about the capabilities the prospective client needs. Once our software is operational in their system, we can see the types of features and functionality they are using most, but for finer grained insights into the exact ways they are using features and data we would need their permission. If we were successful in getting these deeper insights, we would also need to develop a protocol for responding to evidence of AUP violation or software misuse.

How could our software be misused, abused, or used for purposes other than intended? Risks to privacy, freedom of expression, and freedom of movement created by unintended software use primarily pertain to the personal data the retail company collects and stores on their customers and employees. For example:

- If there were a security breach and those records were obtained by an unauthorized party, customers and employees would be vulnerable to surveillance, identity theft, financial fraud, and harassment.

- If our software were used to generate large consumer databases on race, ethnicity, socioeconomic class, or other protected categories, that were then used in marketing efforts, or sold to data brokers without the explicit consent of customers and employees, they would become vulnerable to discriminatory profiling, invasive advertising, and privacy violations.

- Law enforcement agencies may demand that the retailer shares personal data with them, and this data may be much more useful to law enforcement owing to the use of XRLY products.

- Moderate freedom of expression risks exist with our services that enable retail customers to leave public reviews for products. We enable this kind of user-generated content to enhance the shopping experience for other prospective customers. We require that customers be a verified purchaser before being permitted to leave a review and allow customers to decide whether their name appears along with their review commentary. Other customizations and content moderation features are available to our software users, though our reseller partners are the ones with insight into which features they activate and what policies they adhere to.

Can we implement technical limitations to restrict the functionality of the service or the addition of features/customizations that have the potential for adverse impacts? Theoretically we could, but we have not yet developed criteria for what would trigger such an intervention, nor have we developed a process for doing so. We anticipate that such interventions would need to be thoroughly discussed with our reseller partners, so they can notify prospective clients about this authority.

Do we have an Acceptable Use Policy? How do we enforce this AUP (e.g., termination of service or denial of software updates)? We have an AUP that customers must agree to before partnership deals are finalized and SaaS installation begins. We have not yet had to enforce the AUP or terminate an account but working with our compliance team to define that process is a priority for next quarter.

Are any clauses in the AUP addressing product misuse or human rights risks? Not really. Our AUP includes clauses prohibiting the use of our products in violation of the law, which could catch some misuse risks. However, there are no clauses relating to human rights risks specifically—for example, a prohibition on the use of our products to violate the rights of others, or a list of prohibited use cases.

How do we share and educate our customers about best practices for harm mitigations? Could we do more? We prohibit some practices in ourAUP but don't get too deep into it elsewhere. There is some resistance internally and from our resellers to "doing more" because the sales team don't want to slow the process or behave in a way that compromises the deal. Perhaps we can address this in our annual trainings and reseller retreats.

WHAT DO WE KNOW ABOUT OUR CUSTOMERS?

Do we conduct due diligence on prospective customers or prospective uses of our software? The nature of our business partnership with distributors and resellers means that the onus for end-use due diligence falls primarily on the distributors and resellers, but at present this is limited to compliance with export controls, sanctions laws, and corruption risk. To support distributors and resellers, we could develop a list of the characteristics of high-risk industries or use-cases that are likely to be connected to harm. Further, we could expect distributors and resellers to ask prospective customers questions about corporate policies, practices and how they intend to use the service as part of our sales process. However, we don't do any of this today as it might slow the sales process down or adversely impact trusted relationships.

This due diligence could include assessments of the prospective customers' public human rights commitments, company human rights rating and rankings, nature of company ownership, countries of concern, high-risk application areas, and nature of the intended use of the SaaS product or service. If this process uncovers a proposed use-case that we know can contribute to human rights harms, the account could be flagged with members of our Sustainability team to start a thorough human rights review process. If that process determines this potential customer to be a high-risk user, we could advise the distributor or reseller to diplomatically discontinue the sales discussions. However, we are concerned that this volume of activity might not be scalable.

Who oversees our ongoing due diligence efforts? On an individual account basis, the sales rep from the distributor or reseller oversees ongoing due diligence because they are in regular communication with our customers and are expected to track end-use as part of account management. However, it can be challenging to effectively track end-use in practice because customers' privacy and data security needs can prevent us from having full insight into how they are using our products. Distributors and resellers also may not know what flags to look for that might indicate misuse. We don't currently require distributors and resellers to share instances of potential or confirmed misuse and worry that they would fear punitive action if they disclosed these to us.

What are the salient human rights risks we may be involved in because of 1) how our software is designed and developed, 2) how is it used by the customer? See above on the risks to privacy, free expression, freedom of movement, and discrimination.

3. STRATEGIC HUMAN RIGHTS ISSUES

Our key strategic human rights issues are:

- Acceptable Use Policies (AUPs)
- Downstream impacts
- Freedom from Discrimination
- "Know your customer"
- Privacy

4. CONNECTIONS WITH OTHER PARTS OF THE ECOSYSTEM

This section evaluates how XRLY may be involved with human rights impacts through our relationships with other parts of the ecosystem by following the prompts on the "Connections | Enterprise Software Provider" slide. It assesses the relevance of each part of the ecosystem for XRLY's HRDD as either low, medium, or high.

ENABLING HARDWARE

Overlap with our service: the availability of hardware may impact access to or successful delivery of services (e.g., bandwidth / latency)

Relevance: **Low**

How will the local hardware context impact the provision of enterprise software services or the ability of rightsholders to use software or access data? Computers and an internet connection are needed to access our services and data. If computers are not available, users' ability to use software and access data will be stalled. XRLY has no control over this, though we do encourage customers to equip themselves with proper hardware and make our systems and engineering team available to customers at the beginning of all contracts.

How could hardware vulnerabilities / limitations be connected to human rights harms? Security breaches that enable unauthorized parties to access employee devices present privacy risks, though data is stored on the cloud with numerous security gates and hardware devices are password protected and have multiple layers of authorization requirements. We consider hardware vulnerabilities a low risk for adverse human rights impacts, especially compared to software vulnerabilities.

INTERNET STACK

Overlap with our service: Decisions taken across the internet stack (e.g., CDNs, Telcos, Cloud Services Providers) may impact access, availability, and functionality of our services.

Relevance: **High**

What laws, regulations, or other pressure may lead to other companies (e.g., ISPs, CDNs, hosting companies, cloud services providers) blocking or restricting access to our services? This is an area of some uncertainty for us. Perhaps this would happen if XRLY services are found to intentionally and consistently violate applicable privacy and security laws, like the GDPR. For example, our services could be customized by customers to violate GDPR by collecting and storing more data than is necessary for the purported purpose, not informing data subjects of how and which kinds of their data is being collected, or providing them opportunity to give consent, rectify records or have them erased entirely, and/or by changing our default privacy settings to essentially disable our 'privacy by design' features.

One way to avoid this from becoming a possibility is to disallow changes or customizations to the features of our software that ensure GDPR compliance. We'll need to discuss the implications of these technical gates with our engineering team and check with the sales team to ensure that such changes wouldn't create other customization barriers that devalue our product. While unlikely, our user generated content platform XRLYchatterbox could result in blocks or restrictions if illegal or harmful content becomes prevalent or threatening to local authorities.

How could vulnerabilities / service interruptions impact the use / functionality of enterprise software in ways that lead to harm? Consistent internet service is key to our business and to preventing some adverse human rights impacts associated with the loss of access to our products due to connectivity issues. Since enterprise customers use our services to manage their human resources programs, service interruptions that make data inaccessible or damaged could impair an employee's ability to be promoted or take earned leave time. Similarly, sales and marketing personnel could lose out on deals (and thus commission payments) if service interruptions are significant enough to bar them from accessing records, email, or other account details.

Are there entities in the internet stack that we should not work with given their lack of human rights considerations? So far, we have not identified any industry segments that we would not work with in their entirety. We recognize that many adverse human rights impacts can arise from content moderation practices of UGC; as we venture into servicing more clients whose primary service is content hosting and sharing, we will need to develop some criteria to evaluate those customers and their moderation policies.

What responsibility does the software provider have in ensuring fair content moderation practices on UGC? We have to configure our services to allow for some amount of content moderation in product reviews, as a matter of safety. For example, product reviews can contain harmful content like hate speech, threats of violence, or inappropriate imagery that violate our content policy and negatively impact user experience and trust in the platform. However, the exact ways our customers apply their content moderation policies are beyond our purview. Perhaps we will add reporting on active content moderation practices to our due diligence checklist and quarterly reporting template.

CONNECTED SECTORS

Overlap with our service: Availability and access to enterprise software may impact the functionality or efficiency of connected sector services.

Relevance: High

What service terms / acceptable use policies should apply when providing service to connected sectors? Connected sector entities are our primary customer base, though the group is not monolithic, and some segments present much greater human rights risks than others. Our AUP is a purpose-oriented document that describes the intended use of our SaaS service; prohibited technical and behavioral uses, content, or activity; and our enforcement practices (terms of service, termination clauses and reporting channels). Since our business model is primarily closed (we serve as backend infrastructure) our AUP focuses more on data governance and user privacy than types of acceptable content and moderation practices; however, the sudden popularity of XRLYchatterbox may require us to fill that gap. We may also need to create service or sector specific terms to address risks that are unique to each sector, such as the use of XRLY in the criminal legal system, logistics, or financial services sectors. While we primarily target the retail industry, XRLY tools are gaining momentum in other sectors too.

How might government / law enforcement requests for data from actors within connected sectors impact rightsholders? What requests might enterprise software providers receive? Law enforcement agencies may demand that our customers share sensitive user data with them to help the agencies investigate crimes; while law enforcement agencies could approach XRLY directly, we'd pass them on to our customers since they own the data. However, the power of XRLY is such that data provided by our customers to law enforcement agencies could reveal more insight, patterns of behavior, or personal information than in the pre-XRLY era. These requests could violate individuals' rights to privacy, non-discrimiation and freedom of expression because the data collected by our service may be combined with other data enabling law enforcement to identify, track, or monitor individuals, which may result in a violation of individuals' rights.

How might the increase in data generated, collected, stored, and used by connected sectors create new privacy concerns/ impacts to human rights? Our XRLY tools increase the capability of the advertising and data broker industries--for example, in data collection, analysis, and pattern detection. Targeted advertising and data brokerage practices can reveal sensitive, private information about users that could be used for employment, medical, educational and/or social discrimination. Across retail, e-commerce and connected industries like healthcare, finance, and automotive, companies are increasingly shaping their business models to leverage associated user data for service delivery, advertising, and growth. As these practices proliferate, so do the risks of data breaches and discimination.

CONNECTED TECHNOLOGIES

Overlap with our service: Availability and access to enterprise software may impact functionality or efficiency of the connected technology products and platforms.

Relevance: High

What are the human rights impacts associated with the increased and combined use of connected technologies? How might enterprise software facilitate or enable these harms? Our SaaS services incorporate machine learning (ML) systems to analyze data and generate insights. For example, our customers may elect to use our ML models to facilitate individuals' access to a retail company credit card. If the model is trained using historical data, the inputs may lead to discriminatory outputs that suggest a given retail customer has bad credit or should not be made the offer.

When SaaS services include algorithmic decision-making and recommendation systems, what human rights-based approaches can be used to address adverse impacts that result from the AI models? Our proprietary recommender algorithms and decision models are core features of our sales and human resources software. We make our systems responsive to user feedback, so if they oppose the suggestion our AI presents to them (or their retail customers), they can give direct feedback about why. They can also elect to disable the predictive features, though we warn that this may diminish the efficacy of sales and marketing campaigns, and the completeness of employee performance reviews. We don't have an active practice of explaining our AI models but could develop a standard video or FAQ to make this part of our system more compliant with responsible AI principles on explainability and transparency.

What new policies are needed for successful collaboration between enterprise software providers and connected technologies? It would be useful if we developed and/or agreed upon a shared set of principles when using AI/ML that clearly outlined how we can best respect human rights and ensure fairness, transparency and explainability, human-centeredness, and privacy and security.

How might the increase in data generated, collected, stored, and used by connected technologies create new privacy concerns/ impacts to human rights (e.g., law enforcement relationships)? As mentioned above, all this data and associated predictions / insights could enable harmful advertising, discriminatory profiling, or more surveillance of individuals.

ENABLING SOFTWARE

Overlap with our service: Enterprise software providers may partner with other enabling software actors to provide holistic services (e.g., security software and app store backends).

Relevance: Medium

How is sensitive data protected against security breaches? Our Chief Technology Officers (CTOs) and Chief Security Officers (CSOs) ensure that our systems meet all legal requirements, industry standards, and security best practices. We require all internal staff and reseller/distributor partners to complete annual security training. Across our services, sensitive data is protected through technical gates like data encryption and masking in transit, at rest and at use; row level security; multi-level access authorizations and backup and recovery optimization.

What user-controlled data protections are in place? Retail company employees can request a report on the data our services capture about them. Sensitive and protected human resources files may require limited access or file visibility.

When establishing their online profiles, retail customers can toggle selections that control how much of their data we collect, who we are authorized to share it with, and in the event they want to cancel their account–the process for requesting that we delete all their data.

How can access be revoked to protect user and non-user safety? We can terminate licenses and refuse software updates that would render services unuseable.

How is data handled upon termination of a contract? We adhere to all National Institute of Standards and Technology (NIST) best practices for data governance and post-contract destruction. This includes deleting and clearing all PII from our cloud systems at termination, without the possibility of recovery.

How much data access and processing power should enterprise software providers have? This is an ongoing issue of debate for us. We need to make our systems interoperable with other platforms but also recognize that the more interoperability and data access can create more privacy risks. Our current approach is to only access the essential information needed for successful service delivery.

5. ROLE OF EXTERNAL STAKEHOLDERS

This section evaluates how XRLY can undertake meaningful engagement with the various stakeholders in the technology ecosystem by detailing the influence their work has on human rights in our sector, as well as the influence our sector has on their approach to human rights.

Governments and Regulators: Establish regulatory frameworks and policies that guide the market for compliance and data protection. While we primarily target the retail sector, we have also noted interest from governments and regulators to use XRLY products to both improve internal operations and assist with the efficient delivery of public services. When serving customers in countries with authoritarian governments and weak rule of law, this may create significant human rights risks. However, because we mainly sell to retail companies not government entities, this risk is less relevant to us.

Standards Setting Organizations: Interoperability standards are useful guides to our industry, but generally the work of standard setting organizations doesn't influence our approach to mitigating human rights risks outside of privacy and data security standards.

IGOs: The EU's GDPR has changed the ways we configure our systems to adhere to regulations for notice and consent and data governance.

Investors: Our investors are pushing for growth, which often incentivizes us to enhance service offerings rapidly, launch new products with minimal testing, and increase our data analytics capabilities in ways that get close to legal limits. Investors may be keen for us to expand from retail into higher risk sectors and may be skeptical of the need for end-use due diligence.

The Media: The media does not have much influence over how we address human rights, and we have not seen nuanced reporting on the human rights impacts of SaaS companies--indeed, the media's focus on social media and AI/ML companies have likely enabled us to get away with activities and customer relationships that should receive more scrutiny. However, we should prepare for increased scrutiny over time, especially as AI/ML in hiring software and workplace monitoring is receiving increased media coverage.

Civil Society: Civil society has a good sense of the human rights issues in retail, especially those stemming from extensive data collection, cookies, behavioral tracking, and targeted advertising--however, the role of XRLY in enabling these features is sometimes missed. Civil society is generally concerned with adverse impacts arising from misuse of AI-enabled analytics and data sharing in markets where individual identities, choices and preferences may be criminalized.

Academics: We do not consider this stakeholder group a key connection for us.

Users: Users' use habits shape our understanding of service relevancy and market demands. Their use habits can also implicate us as enablers of human rights harms.

Non-user Rightsholders: Non-user rightsholders can have their human rights adversely impacted if our cloud services are misused (e.g., for surveillance purposes), or our systems capture and/or share their personal data without consent. Further downstream, the location of the retail company stores we work with have economic impacts on the local community, such as impacting rental rates, employment markets, and infrastructure development.

6. Key Considerations to Address Adverse Human Rights Impacts

In our estimation, the greatest risks arise from data breaches (privacy harms); connected sector uses of our services (by governments, advertisers, or e-commerce) in ways that could adversely impact a range of rights, and UGC moderation practices (freedom of expression). There are a few key things we can do to address these risks. First is meaningfully engaging with sales partners (our resellers and distributors) and key customer segments to raise awareness around end-use human rights risks. Second is build out our AUPs with clauses that include both general and specific human rights protections. Third is to augment our sales due diligence systems to include an assessment of human rights risks associated with the use of our products and services, and a gating process that requires review and approval for high-risk sales. Fourth is to require our sales partners to implement appropriate HRDD as part of their own sales processes.

7. TO EXPLORE FURTHER

The following resources may be helpful for exploring the human rights risks associated with the SaaS sector and their interactions with other parts of the ecosystem in more depth:

- Human Rights Assessment of the Software-as-a-Service Sector | Reports | BSR
- Responsible Product Use in the SaaS Sector | Reports | BSR
- Taking Action to Address Human Rights Risks Related to End-Use | B-Tech
- Human Rights Due Diligence of Products and Services | Reports | BSR