DATA TRANSFER INITIATIVE

A third-party trust model for direct personal data transfers

Data Transfer Initiative

The notion of data portability is in the midst of a paradigm shift. The original concept of data portability is a data protection centered, individual user right. It says that users should be able to download personal data they have contributed to services along with data reflecting their activity within services. That use case remains intact. However, with the rise of competing services in the market and new opportunities for businesses to deliver user value through personal data, increasingly data portability also encompasses the direct transfer of personal data associated with a user, at the user's request, from one service to another. This offers many technical advantages, including avoiding any limitations that might arise from poor connectivity or limited local device storage, and shifts the burden of ensuring the data is useful on the receiving service, towards a shared responsibility, as the Data Transfer Project has demonstrated.

The Digital Markets Act in the European Union solidifies this paradigm shift further by requiring designated gatekeepers to make certain personal data available at the request of a user to be transferred directly to third parties. Proposed laws in other countries and at least one U.S. state would arguably impose comparable requirements. Between the improvements in usability and regulatory pressure, direct third-party transfers represent the future for data portability.

Yet, there is no universal method for establishing trust between the parties involved in the direct transfer process. Can the receiving service be confident that the data will not introduce security or other risks? Can the sending service be confident that the data will be secured following the transfer? Can both parties be confident that the transfer reflects the user's understanding and intention? Do users understand the potential safety and privacy implications of authorizing the transfer? Most importantly: How can the necessary trust be established through a process that does not undermine the user making the request, nor compromise the potential for data portability to deliver powerful competitive benefits for the market as a whole?

This report proposes a series of processes and questions suitable to address this gap. While inspired by the near-term challenge of facilitating effective DMA compliance, its purpose and intended impact is broader: to help grow trust in direct transfers regardless of the nature of the parties or their location.

The Data Transfer Initiative began consultations with a range of stakeholders in the fall of 2023 and began circulating concepts and language to organizations from industry large and small and civil society in early 2024. DTI's Executive Director, Chris Riley, led a working session in Brussels in January to further present and seek input on this effort.

The output, as presented in the appendix to this report, includes questions that may be asked by sending and receiving services of the other party. These questions allow the asking party to assess risk associated with the transaction and consequently to make a determination whether or not to allow the transfer. They touch on the legal jurisdictions of the parties (and thus relevant legal frameworks), security and privacy practices, and how the parties establish user understanding and intent. The model also provides guidance on mechanisms for expiration, revocation, and appeal. DTI expects that rejections of transfer requests under this model will be less frequent than approvals; however, there are indeed circumstances, as this report will elaborate, where rejection is appropriate. Even in those cases, the individual's right to data portability remains intact, as the user can download

their data and upload it to any receiving service, without the original service having any opportunity to limit that subsequent use (or even to be aware of it).

DTI is uniquely positioned to shepherd this exercise, both through the experience of building tools – as part of the open source Data Transfer Project – for precisely such kinds of transfers, and through the organization's broad network and reach. Furthermore, there is inherent value in conducting this exercise in improving commonality and consistency among the models implemented by each major data service, which in turn improves the efficiency and experience for receiving services in gaining access to different data resources. Lowering the barriers to entry for both sending and receiving services in taking part in direct transfers works to advance DTI's organizational mission: to empower people by building a vibrant ecosystem for simple and secure data transfers.

This report is intended to be a static document. The model itself, in the appendix, is expected to change and improve over time, as stakeholders implement trust mechanisms in varying ways and learn from those experiences.

1. Introduction

Data portability is a powerful tool to allow individuals to exercise more agency over personal data. While its origins lie in data protection and fundamental rights contexts, and that context remains salient, portability's nexus to competition gives it new relevance. Portability helps consumers have more choices of services, and it helps businesses both to compete over similar service offerings and to innovate new features and functions. To take full advantage of these benefits, portability is evolving from a tool principally exercised by individual users requesting and receiving their data from a service provider via download, to implementation through direct transfers, where a user requests a service to transfer personal data directly to a third party.

The question that arises as a consequence of this shift is, how to build trust in these direct transfers. Exchanges of personal data are not without risk. DTI counts privacy and security among its core principles: "Platforms should not be required to substantially threaten the privacy and security of individual users or the platforms. Platforms may require baseline privacy and security standards, including shared certification or accreditation processes, as a condition of enabling data transfer." This report is an exercise in coordinating the development of such standards.

The DMA, Article 6(9), requires designated gatekeepers to facilitate the direct porting of personal data to third parties at the direction of the user. Similarly, Article 6(10) requires gatekeepers to provide business users with access to data generated by them or their end users (which may include personal data, subject to their consent). But nothing in the law specifies how harm is to be mitigated as a result of these increases in access. The DMA does not include any details or requirements as to how third parties could or should be authorized from a security or privacy perspective.

In the absence of any harm mitigation, direct transfer obligations can pose significant risk of harm to users and to the platforms themselves, and can facilitate the commission of violations of the General Data Protection Regulation and other privacy laws. The notion of "continuous and real time" portability, as required under the DMA, creates the potential to compound that risk further in various ways. While there may be room for debate as to the extent to which this risk could result in liability for the service that is effectively providing access to the data, this report's exercise to build trust is not contingent on such a determination. Building trust is good in general, including to the extent that users may blame the services sending their data for actions taken by the services receiving the data, regardless of how legal liability is determined. And certainly, the DMA, in Recital 12 and Article 8, indicates that compliance with portability obligations must not undermine the GDPR.

There remains the risk that an overly conservative review could re-introduce unwanted behavior, or that an overly permissive review could empower bad actors and cause users to lose faith and trust in data portability – in either case, undermining the objectives of the DMA. Such risk can be mitigated through thoughtful, shared design of a model approach for building trust. Alignment on approach will greatly further the DMA's core values of fairness and contestability.

2. Goal and intended use of a shared trust model

To protect the interests of third parties, end users, and data services, this report will describe and provide a security and safety trust model, intended for use by both services sending and receiving data as they mutually authorize each other for a direct transfer. In the context of the DMA, much of the focus is on sending services approving third party receivers; however, its intended utility is broader in both the parties and the geographies involved. The report describes its output as a "model" rather than a framework or standard, however in many senses, these alternate descriptions could be used. The objective is to help guide, structure, and align independent implementations of trust processes in

a way that provides consistency, including for companies of all sizes and countries operating as both senders and receivers of data. Use of a shared model in this broader context is preferable to companies independently developing their own mechanisms, particularly in that it facilitates greater neutrality and consistent application of review, as well as stakeholder buy-in to and acceptance of the framework.

In this way, the proposed trust model mirrors the Data Transfer Project, at DTI's core of technology; both exercises seek to reduce what would otherwise be an "n-squared" (n^2) scaling challenge in connecting parties to each other for data transfers. Just as a shared data model reduces the cost of connecting to new providers of comparable technology by reducing the number of necessary adapters, so too will a shared trust model make it easier for a service provider to prove its trustworthiness to multiple data senders by reusing answers and evidence or simply relying on another party's signal of trust.

The model is intended to be implemented individually by service providers, as senders and receivers, with variations specific to individual contexts. Because of the varying sensitivity of underlying personal data, the same questions with the same answers may lead to different assessments and outcomes depending on the context. In some cases, and in some ways, these differences may converge over time as better risk assessments become feasible; in others, though, the context at hand, including the nature of the personal data and of its use, will support different assessments and outcomes in perpetuity.

While the future use of data portability is impossible to predict, it seems likely that the vast majority of third parties and direct data transfer requests should be capable of establishing trust. However, in those circumstances where trust cannot be adequately established, direct transfer can and should be barred. The user requesting the transfer would, and should always, retain the "manual option" to download their data and subsequently upload it anywhere, without any need for a trust relationship between the source and destination (who would not necessarily be aware even of the other party's existence).

Implementing this model will require new, or reused, internal company processes. These parallel existing processes that are in place today in many companies, notably those used to review requests for API keys, sometimes characterized as developer access to systems. Many companies have well-documented internal review processes for such access. To the extent that trust model implementations align with these existing mechanisms for system access, the reuse of interfaces, validity durations, and other elements may help reduce the costs of implementing this model, and in general increase the clarity and alignment of service provider interactions.

One important caveat in interpreting and understanding this model is that there is a distinction between provider-level trust and transactional review. Even between two providers with maximum trust, there will often be necessary steps taken to process individual transactions, such as ensuring that the user is intending to make the transfer in that instance. While data portability will become more trusted and more straightforward, it cannot inherently be instantaneous, considering the scale and potential sensitivity of the data being transferred.

A second caveat worth noting is that any trust model, any trust exercise, will have some inherent limitations. Service providers are not regulators, and lack both the authority and the tools to enforce data protection and security laws. A trust model can help reduce the possibility and frequency of harm, but cannot eliminate it; there will remain a critical role for government enforcers to monitor and penalize those in violation of the law.

3. Potential harm of unprotected transfers

Because data transfers reduce friction for portability, they also reduce friction for potential harm. As one example, where a sending service lists another service provider in a drop-down menu for potential direct data transfers, that inclusion signals to the user that the sender has somehow reviewed and approved the potential destination for safety. It becomes easier, in this instance, for a malicious third party service provider to trick a user into transferring their data.

But the scale of potential harm in data portability is broader than this. DTI's writings on threat modeling for data portability draw on existing threat model frameworks to articulate the potential threats as including:

- Disclosure of the data being transferred (i.e. to other parties not authorized by the user);
- Indirect loss of privacy, e.g. linking, identifying, detecting;
- Tampering;
- Non-Repudiation;
- Denial of Service;
- Elevation of Privilege;
- Non-Compliance;
- Harmful Content;
- Spoofing and related bad actor activity; and
- Permission and access control challenges.

Notably, these threats can, collectively, cause harm to all parties involved in direct data transfers: users, data senders, and data receivers, although users are at greatest risk. Personal data reflects personal lives. Video viewing histories, searches within search engines or mapping applications, browsing or shopping activity, and other kinds of data typically made available through portability interfaces can provide indications of medical information, political or religious beliefs, racial or ethnic origin, or sexual orientation, as well as home and work addresses.

Trust between the parties helps mitigate the risk of abuse. And without any process or review before allowing direct data access and transfer, it could become impossible to prevent any actors, even those known to be in violation of established data protection laws, from posing harm. At the same time, service providers would be unwise simply to accept and process any data that a user transfers in without some level of malware detection. Since receiving, storing, and processing data can introduce legal as well as security liabilities, and further can bypass normal mechanisms for content moderation and permission-checking, receivers of data benefit in establishing trust with data senders too.

A few scenarios can help illustrate the potential for harm and the opportunity for trust between service providers to prevent harm:

1. **Horizontal**: Two similar companies offer comparable social media services, and a user seeks to move their past media from one service, where it was available only to a select group of users, to a new service, which by default allows media to be world-viewable including to on-lookers who do not have accounts in the system. This widely-understood example illustrates the potential for harm to the user if this change in data visibility is not known or able to be controlled in some way; a trust model can catch such misalignment in expectations.
2. **New use**: A service receiving data plans future uses of the data that are not consistent with the original sending service's use of the data or with what users might expect, and also is not

made clear to and consented to by the user. Examples could include the unclear or undisclosed resale of portions of the data for advertising, or mining of data to identify individuals who might be particularly susceptible to scam attempts. Furthermore, where that receiving service is outside the European Union or a jurisdiction of effective data protection, such undisclosed and potentially harmful future uses may not be in violation of any laws or rules. However, if the sender and receiver properly communicate and align regarding post-transfer use and ensure user understanding and intent, these harms can be largely mitigated.

3. **Poor protection**: In a range of ways, poor data security practices by either the sender or receiver may result in harm. Where the user is not properly authenticated, the transaction may not be intended or wanted; where the data is not stored and processed securely, private data may be inadvertently revealed, or harmful data or software code could be introduced into a system. Core to the foundation of trust between the parties is a shared understanding of sufficiency of data security practices.

This is not by any means offered as an exhaustive list. At the same time, it is also not meant to convey that data portability as implemented today is an inherently or particularly risky endeavor – provided that suitable measures are taken. Building a foundation of trust sufficient to greatly mitigate harm in portability is feasible, and essential.

4. Key criteria for a trust process

This section will describe what it looks like for a service provider to put a trust process into practice. This requires defining what questions could be asked of another party; determining how to handle internal review; and setting up mechanisms to communicate the outcomes of the process to an inquiring party, and particularly notifications and appeals for denials. Finally, the service provider should make some information on its process public for transparency purposes. While this is framed typically in a one-directional perspective – one service provider, whether sender or receiver, evaluating the other party for trust – in some circumstances it may be a mutual and parallel exercise where both parties are evaluating the other at the same time.

1. **Substance**: This is the principal subject of the appendix to this report, the information sought by the service provider of the third party seeking to establish trust to facilitate direct data transfers. Questions within this process can cover, for example, what the third party promises to do or not do to protect and secure data; authentication and authorization steps the third party will undertake to ensure the action reflects the user's intention and agency; and other matters related to the risk of direct transaction. These questions are not meant to allow the service provider to make judgments regarding the third party's legitimate business model, nor are they intended to deputize the service provider to enforce data protection law. Nevertheless, there is value in ensuring alignment and shared understanding, and as well a baseline to allow the data sender to revoke access should there be any deception later revealed (noting that the ability to revoke subsequently doesn't obviate the need for checking in advance).
2. **Process**: In some cases, automated means may be used to review answers provided by a third party; in others, or as a second step after initial automated review, human reviewers may be involved. The dual goals of expediency and efficacy can come into tension, and over time, service providers may learn from each other what balance works best to promote effective trust and data portability.

3. **Outcomes**: Either a positive or negative response to the process would be shared with the requesting party; if positive, some documentation or evidence may be provided, and if negative, the response should include some indications of the deficiency and any opportunities to appeal or to cure. This matter will also be presented in the appendix to this report. Additionally, a decision in favor of trust may under some circumstances need to be subsequently revoked; processes to consider and execute such revocation could be established.
4. **Transparency**: A service provider may wish to develop and make public some information on its approach to trust, even through a narrative such as a written policy, in the same way API access policies are typically documented to third parties.

While the details of how best to realize these goals will evolve over time, we expect that these four pieces of substance, process, clarity in outcomes, and transparency will always be necessary.

5. Next steps

This report is meant as a first step. Designated gatekeepers implementing the Digital Markets Act will put trust mechanisms in place, and the degree to which these approaches align with this model will be a subject for study – as will user behaviors and the analysis of regulators, in the EU and around the world. The model presented in the appendix to this report will be a living document, at least at first, and is expected to be adapted over time as more lessons are learned of how best to build trust in practice. DTI intends in producing this work to encourage alignment in trust mechanisms across stakeholders, and expects that increasing alignment will also build trust itself, in both the technical sense and in the general sense among services sending and receiving data, and regulators.

Many specific details around risk assessments based on the questions and issues raised in this model remain unresolved, including a number of thresholds that can be calibrated for individual contexts. But acknowledging shared principles and approaches through a shared model, and demonstrating a commitment to working out the details in a responsible way to help protect users without impeding legitimate business, will bring short and long term benefits for portability.

Appendix: Initial Trust Model

The proposed model establishes the key elements of the direct transfer trust-building process along with some suggested questions to identify specific data to be provided. While some of these elements and questions are symmetric, and the same whether asked by a sending service of a receiving service or vice-versa, most are asymmetric with either different or no applicability when asked in the reverse direction; as the risks involved differ between sender and receiver, the process for building trust must also vary, and the burdens of implementation will also fall differently, and often asymmetrically.

1. **Jurisdiction:** Determine the geographic location of a service partner and in particular, what if any data protection laws will apply. While not dispositive in itself, this information may be useful in an individual service's (either sender or receiver) risk assessment, particularly in how subsequent questions are evaluated for risk.
    a. Sample question: What legal jurisdiction governs the provider? Is data stored in a separate jurisdiction from the business? If so, which one? If not obvious from the stated jurisdiction, what data protection laws will apply to data? Does the business have a presence or representative in the European Union?
2. **User authentication**: Identify and authenticate users properly to both parties.
    a. Sample questions: What authentication method for end users is in use by the provider? Is a fresh authentication process initiated as part of a transfer request or is evidence of recent authentication such as a local device cookie sufficient?
        i. Note that the appropriate authentication method, and whether authentication is re-acquired or verified at time of transfer, will depend in practice on context and use cases, and are not necessarily symmetrical between source and destination.
3. **User authorization**: Ensure that users are properly indicating their desire for the transfer as scoped, with an understanding of what data is to be transferred, the frequency of transfer if appropriate, and how the data will be used. Additionally, if the user is not of legal age in the relevant jurisdiction, additional steps may be needed.
    a. Existing mechanisms for validating user download requests for personal data, such as to assess the risk of fraudulent activity, may be repurposed.
        i. Any mechanisms for seeking parental consent for minor users in downloads may be repurposed.
    b. Sample question: How is the scope of the data to be transferred made clear to the user, and is the user given specific controls to include or exclude portions of the data in/from the transfer?
    c. Sample question: How is the receiving service's intended use of the data shown to the user? Evidence may include screenshots and process statements to indicate user facing language such as a service saying "These photos will be added to your Drafts folder and left as private until you take action to post them to a feed." or instead "These photos will be added directly to your public feed."

    i. Note that any efforts by a service to track user comprehension of relevant disclosures, such as through user surveys, could add value and further reduce risk; but in most cases no such efforts should be required.

4. **Proper use**: Check that the use of transferred data by the receiving service is consistent with user expectations in context. As a component of this, asking for the other party's privacy policy is standard.
    a. Sample question: What is the nature of the service provided for users and what features or benefits will the transferred data enable for those users?
    b. Sample question: Will the provider process or use the data for any other purpose other than to directly deliver its services? If so, how do these uses relate to the service, and how have these uses been made clear to users?
    c. Sample question: Are all of the planned uses of this service legal within the relevant jurisdictions?
    d. Sample question: Have relevant legal authorities taken any actions directed to the provider related to actual or potential violations of relevant data protection law? If so, describe the nature of the action and its outcome, if resolved.
5. **Responsible data practices:** In reviewing receiving services in a transfer process, the sender may wish to determine the receiver's level of protection for the data after receipt, including securing the data from loss or unlawful access. Some amount of proportionality is appropriate in this context depending on the nature of the underlying data, e.g., distinguishing between data made publicly available and not.
    a. Sample question: How will the data be secured in internal systems storage?
    b. Sample question: What limits are in place on the service provider's employees and partners to control access to the data?
    c. Sample question: Are there internal controls in place to ensure the data is not made available for subsequent unauthorized secondary uses?
    d. Standards compliance, such as ISO 27001 certification or a SOC2 audit, may be designated as sufficient at the discretion of the implementer of the framework. Such compliance should not be mandatory, particularly as these methods are intended for cybersecurity rather than data protection; however, their use may accelerate or simplify trust processes by reducing the scale of or eliminating some portions of review.
6. **Security of source**: Receiving services may choose not to allow by default the transfer of data from all potential sending services, and may seek to ensure a sender is not transmitting data which could pose harm to the service.
    a. Sample question: Is the data secured and free from malware or other harm, to the best of the provider's knowledge and determination?
    b. Standards compliance, such as ISO 27001 certification or a SOC2 audit, may be designated as sufficient and may be useful to accelerate or simplify trust processes; however, such compliance should not be mandatory.
7. **Duration, accountability and appeal:** Establish transparent mechanisms to periodically review trust approvals, and to red-flag providers who are determined to have given false or unreliable information or otherwise developed circumstances of concern, as well as mechanisms for service providers to appeal such an action to another party.
    a. Existing mechanisms used for similar purposes for developer access to APIs may be repurposed for direct data portability.

b. Approvals of trust under this model should expire after a predetermined period of time, such as six (6) or twelve (12) months, although the appropriate period may vary depending on context including the nature of the data involved, and evidence of harm may trigger early review. Renewals may follow the original trust process or another process of re-approval, depending on sensitivity and context.
c. Approvals of trust should be revocable by the original party who authorized them or, potentially, a mutually trusted intermediary. While such actions should not be taken lightly without fair warning or notice, as revoking access to data can undermine business models and investment theses, evidence of harm such as a determination that a provider has violated data protection law by a relevant legal authority or that any information provided in an initial trust process was invalid may well justify revocation depending on the relevant context and severity.
d. Denials or revocations based on this model should include clear information on the reasons for action and what actions or new information would be sufficient to cure, along with mechanisms to appeal or restart a trust process based on such changed actions or new information. Such appeals should not be denied through a purely automatic process but should incorporate manual review, and ideally the opportunity for direct communication between providers.