

Data Transfer Project Overview and Fundamentals

July 20, 2018

Table of Contents

[Abstract](#)

[Background](#)

[Principles](#)

[The Future: Service-to-Service Portability](#)

[The Data Transfer Project Overview](#)

[Use Cases](#)

[The project participants' role with respect to data](#)

[Architecture](#)

[System Components](#)

[Data Models](#)

[Adapters](#)

[Data Adapters](#)

[Authentication Adapters](#)

[Task Management](#)

[Worker](#)

[Securely Storing Data](#)

[System Access](#)

[Deployment](#)

[Deploying a Host Platform](#)

[Acquire API keys](#)

[Deploy](#)

[Create a User Interface \(UI\)](#)

[Deployment Models](#)

[Distributed](#)

[Centralized](#)

[Self-managed](#)

[Security & Privacy](#)

[Data minimization](#)

[User notification](#)

[Rate Limiting](#)

[Token Revocation](#)

[Minimal Scopes for Auth Tokens](#)

[Data retention](#)

[Abuse](#)

[Ecosystem Issues](#)

[Project Governance](#)

[Inconsistent API Landscape](#)

[Reciprocity](#)

[In source code](#)

[Transparency](#)

[Automated fidelity test](#)

[Data Portability Provider Pledge](#)

[Personal Information Management Systems \(PIMS\)](#)

[Implementation](#)

[Ways to Participate](#)

[Add a Provider](#)

[Contribute to the Core Framework](#)

[Conclusion](#)

[Resources](#)

[Mailing list: dtp-discuss@googlegroups.com](mailto:ntp-discuss@googlegroups.com)

[Glossary](#)

Abstract

The Data Transfer Project (DTP) extends data portability beyond a user's ability to download a copy of their data from their service provider¹ ("provider"), to providing the user the ability to initiate a direct transfer of their data into and out of any participating provider.

The Data Transfer Project is an open source initiative to encourage participation of as many providers as possible. The DTP will enhance the data portability ecosystem by reducing the infrastructure burden on both providers and users, which should in turn increase the number of services offering portability. The protocols and methodology of the DTP enable direct, service-to-service data transfer with streamlined engineering work.

Background

Data portability and interoperability are central to innovation. The DTP Partners believe that people should use products because they provide unique value and features. If a user wants to switch to another product or service because they think it is better, they should be able to do so as easily as possible. This concept of allowing users to choose products and services based on choice, rather than being locked in, helps drive innovation and facilitates competition.

Data portability can also provide security benefits for users. Practical tools that let users backup or archive important information, organize information within multiple accounts, recover from account hijacking, and retrieve data from deprecated services all work to improve user security.

Perhaps most importantly, data portability provides peace of mind to users.

Principles

We believe the following principles around interoperability and portability of data promote user choice and encourage responsible product development, maximizing the benefits to users and mitigating the potential drawbacks.

- **Build for users:** Data portability tools should be easy to find, intuitive to use, and readily available for users. They should also be open and interoperable with standard industry formats, where applicable, so that users can easily transfer data between services or download it for their own purposes.
- **Privacy and security:** Providers on each side of the portability transaction should have strong privacy and security measures—such as encryption in transit—to guard against unauthorized access, diversion of data, or other types of fraud. It is important to apply privacy principles such as data minimization and transparency when transferring data between providers. When users initiate a transfer they should be told in a clear and concise manner about the types and scope of data being transferred as well as how the data will be used at the destination service. Users should also be advised about the privacy and security

¹ See ISO/IEC 17788:2014 section 3.2.15

practices of the destination service. These measures will help to educate users about the data being transferred and how the data will be used at the destination service. More details are in the [Security & Privacy](#) section below.

- **Reciprocity:** While portability offers more choice and flexibility for users, it will be important to guard against incentives that are misaligned with user interests. A user's decision to move data to another service should not result in any loss of transparency or control over that data. Individuals should have assurance that data imported to a provider can likewise be exported again, if they so choose. Ultimately, users should be able to make informed choices about where to store their data. We believe that providing transparency around portability will lead to users preferring providers that are committed to reciprocal data portability, over those that are not.
- **Focus on user's data:** Portability efforts should emphasize data and use cases that support the individual user. Focusing on content a user creates, imports, approves for collection, or has control over reduces the friction for users who want to switch among products or services or use their data in novel ways, because the data they export is meaningful to them. Portability should not extend to data that may negatively impact the privacy of other users, or data collected to improve a service, including data generated to improve system performance or train models that may be commercially sensitive or proprietary. This approach encourages companies to continue to support data portability, knowing that their proprietary technologies are not threatened by data portability requirements. For a detailed taxonomy of such data, see [ISO/IEC 19944:2017](#).
- **Respect Everyone:** We live in a collaborative world: people connect and share on social media, they edit docs together, and they comment on videos, pictures and more. Data portability tools should focus only on providing data that is directly tied to the person requesting the transfer. We think this strikes the right balance between portability, privacy, and benefits of trying a new service.

The Future: Service-to-Service Portability

Given our understanding of the current ecosystem, particularly the infrastructure constraints faced by people around the globe, it is clear that the future of portability will need to be more inclusive, flexible, and open. We believe users should be able to seamlessly and securely transfer their data directly from one provider to another. We also believe the approach outlined below aligns closely with the principles of other thought leaders in the space².

Our vision for this project is that it will enable a connection between any two public-facing product interfaces for importing and exporting data directly. This is especially important for users in emerging markets, or on slow or metered connections, as our project does not require a user to upload and download the data over what may be low bandwidth connections and at potentially significant personal expense.

² http://mesinfos.fing.org/wp-content/uploads/2018/03/PrezDataccess_EN_V1.21.pdf

Data Transfer Project Overview

The DTP was developed to test concepts and feasibility of transferring specific types of user data between online services. The recently published international standard on Cloud Computing—Interoperability and Portability ([ISO/IEC 19941:2017](#))—notes that “there is no single way of handling portability issues within cloud computing.” Each portability scenario needs separate analysis of the systems, data and human aspects (involving syntactic, semantic and policy facets), both for the sending cloud service as well as the receiving one.

Moreover, for each portability scenario, issues of identity, security and privacy must be addressed by both the initial cloud service and the receiving one. With this community-driven effort, the DTP Contributors expect to learn more about the development of tools and techniques to facilitate portability of specific, desirable user data between providers. Establishing an open source ecosystem for the creation of these tools is an important step for understanding and enabling data portability for our users, within a community-driven, inclusive, multi-stakeholder process. This will also help prioritize those data portability scenarios that the community and users deem in-demand, in a pragmatic, real-world setting.

The DTP is powered by an ecosystem of adapters (Adapters) that convert a range of proprietary formats into a small number of canonical formats (Data Models) useful for transferring data. This allows data transfer between any two providers using the provider’s existing authorization mechanism, and allows each provider to maintain control over the security of their service. This also adds to the sustainability of the ecosystem, since companies can attract new customers, or build a user base for new products, by supporting and maintaining the ability to easily import and export a user’s data.

Transferring data using canonical formats will not necessarily mitigate problems such as formatting limitations or inconsistent feature support. However, our approach illustrates that a substantial degree of industry-wide data portability can be achieved without dramatic changes to existing products or authorization mechanisms, while still providing a flexible enough platform to adapt and expand to support new formats and use cases brought by future innovation. Additionally, the Data Transfer Project has been developed to increase participation by motivating providers to build both export and import functionality into their services.

This paper describes the technical foundation of the Data Transfer Project. Importantly, this work is available in an open-source format to be universally accessible and promote transparency.

Use Cases

Individuals have many reasons to transfer data, but we want to highlight a few examples that demonstrate the additional value of service-to-service portability.

- A user discovers a new photo printing service offering beautiful and innovative photo book formats, but their photos are stored in their social media account. With the Data Transfer

Project, they could visit a website or app offered by the photo printing service and initiate a transfer directly from their social media platform to the photo book service.

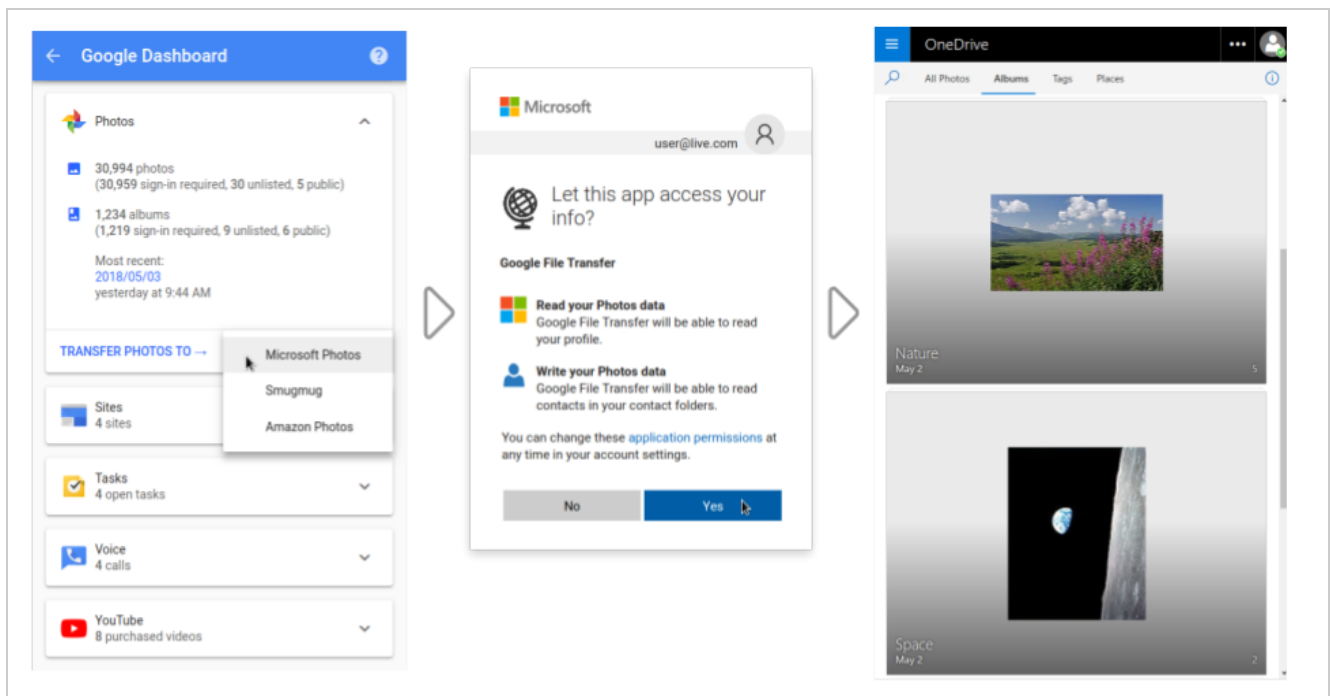
- A user doesn't agree with the privacy policy of their music service. They want to stop using it immediately, but don't want to lose the playlists they have created. Using this open-source software, they could use the export functionality of the original provider to save a copy of their playlists to the cloud. This enables them to import the playlists to a new provider, or multiple providers, once they decide on a new service.
- A company is getting requests from customers who would like to import data from a legacy provider that is going out of business. The legacy provider has limited options for letting customers move their data. The company writes an Adapter for the legacy provider's Application Program Interfaces (APIs) that permits users to transfer data to the company's service, also benefiting other providers that handle the same data type.
- A user in a low bandwidth area has been working with an architect on drawings and graphics for a new house. At the end of the project, they both want to transfer all the files from a shared storage system to the user's cloud storage drive. They go to the cloud storage Data Transfer Project User Interface (UI) and move hundreds of large files directly, without straining their bandwidth.
- An industry association for supermarkets wants to allow customers to transfer their purchase history from one member grocer to another, so they can get coupons based on buying habits between stores. The Association would do this by hosting an industry-specific Host Platform of the DTP.

The innovation in each of these examples lies behind the scenes: the Data Transfer Project will make it easy for providers to meet their users' expectation that they will be able to migrate data with minimal engineering effort. In most cases, the direct data transfer experience will be branded and managed by the receiving provider, and the customer will authenticate their existing account with both parties. They won't need to see the DTP branding or infrastructure at all as part of the process.

It is worth noting that the Data Transfer Project doesn't include any automated deletion architecture. Once a user has verified that the desired data is migrated, they would have to delete their data from their original service using that service's deletion tool if they wanted the data deleted.

Generally, any two providers participating in the DTP ecosystem can enable users to transfer data directly between them and participation is self-determined. There is no reason participation would be more attractive to large providers; in fact small providers would reap relatively larger benefits from participating, as the DTP reduces the amount of engineering work.

The illustration below demonstrates what portability enabled by the DTP might look like. In this case, the customer wants to join a new provider (Microsoft) and is requesting their files from their existing provider (Google):



In this hypothetical example, a Google Photos user wants to move their photos from Google to Microsoft OneDrive. They go to Google’s file transfer interface, choose the destination, and hit ‘send.’ They then must authorize the transfer using both services’ chosen methods, in this case OAuth. The selected files are automatically copied and routed to the destination, without using either bandwidth or hardware of the user.

The project participants’ role with respect to data

The focus of the DTP is for consumers to be able to have more control over their data. It should be noted that this solution can support data subjects’ individual rights and applies to services for which the user has a direct relationship with a DTP service, for example Gmail.com or Outlook.com (free services for consumers) or Office 365 Home (subscription-based service for home users). In enterprise cases, the customers of those services, for example their domain administrators, should have the choice to use or not use the DTP based on their individual situation. Each enterprise or commercial customer and their admins need to make the choice of how to implement portability of their users’ data based on their individual scenario³ (note that such users are often their employees). It should also be noted that scenarios involving commercial services and their professional users are often addressed by their respective contracts or existing tools.

Architecture

The DTP solution was designed so that it is easy for a provider to adopt and enable the standards, and took into consideration the following constraints:

³ Privacy experts may find it easier to think of the DTP as a solution that can support data subjects’ individual rights and is aimed at enabling data controllers to offer portability solutions. This solution, however, cannot be directly used for the commercial customers when they are the data controllers, and the services integrated with the DTP are just data processors. In such cases, the data controllers (who may be using the services integrated with the DTP) should make the choice of using the framework to enable portability. In either case, it is the data controller that decides the data portability requirements.

- **Use existing standards; don't create new ones**
By supporting existing standards where possible (like OAuth and REST), we aim to minimize the foundational work required before the DTP can be built and put into action. Widespread adoption and understanding of existing standards makes this possible. As new standards are developed and adopted, they will be reviewed and, where applicable, incorporated into the DTP.
- **Minimize the work required to add a service**
We designed the DTP to allow providers to participate without impacting their own core infrastructure. Service providers can build Adapters and enable import and export functionality that works with their existing APIs and authorization mechanisms.
- **Support standard deployment infrastructure**
The DTP infrastructure was designed with the flexibility to be hosted by anyone, anywhere. See [Deployment](#) for more details.

System Components

As noted above, the system comprises three main components:

- **Data Models** are the canonical formats that establish a common understanding of how to transfer data.
- **Adapters** provide a method for converting each provider's proprietary data and authentication formats into a form that is usable by the system.
- **Task Management Library** provides the plumbing to power the system.

Data Models

The Data Model for transferring files consists of two parts: a file type and the additional metadata needed by the receiving provider to accurately import the data. For example with photos, the file type might be a standard format such as JPEG, and the metadata would include information such as title, description, album, and so forth.

Data Models are clustered together, typically by industry grouping, to form Verticals. A provider could have data in one or more Verticals. Verticals could be photos, email, contacts, or music. Each Vertical has its own set of Data Models that enable seamless transfer of the relevant file types. For example, the Music Vertical could have Data Models for music, playlists and videos.

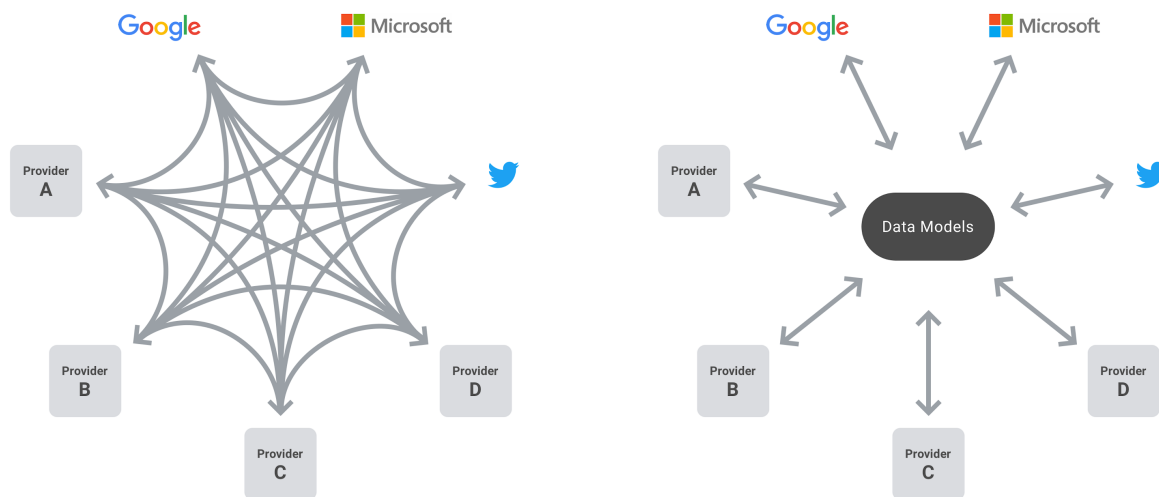
Ideally, a Vertical will have a small number of well-defined and widely-adopted Data Models. In such a situation, the generally accepted standard will be used as the Data Model for that Vertical across companies. This is not currently the case for most Verticals because Data Models have emerged organically in a largely disconnected ecosystem.

One goal of the DTP is to encourage organizations to use common Data Models in their systems, which will happen if organizations take importing and exporting data into consideration when

initially designing their systems or providing updates. Using a common Data Model will significantly reduce the need for companies to maintain and update proprietary APIs.

In the case where there is no standard Data Model for a Vertical, companies will want to collaborate and agree upon standardized Data Models, either during the DTP development or in collaboration with external standards bodies. Without collaboration, each provider could have their own Data Model, and would have to create adapters that would have to support the same number of Data Models as there are companies in the Vertical, which would reduce the usefulness of the DTP.

Even where standard Data Models do exist, collaboration will be an ongoing and mutually beneficial shared commitment as APIs will need to be maintained to handle new features, evolving standards, or innovative new formats.



Without the DTP

Each provider has to build and maintain Adapters for every other provider's proprietary APIs and, potentially, data formats

With the DTP

Each provider only has to build and maintain an API that supports the DTP Data Models, which are based on standard formats where available

Adapters

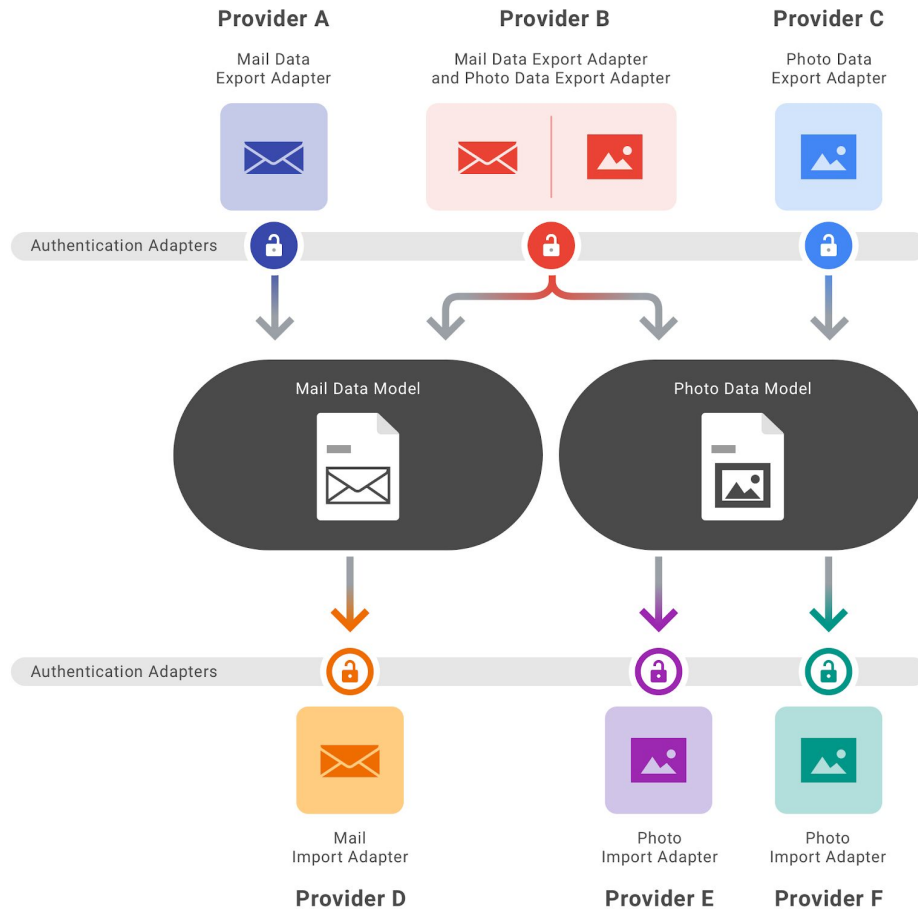
There are two main kinds of Adapters: Data Adapters and Authentication Adapters. These Adapters exist outside of a provider's core infrastructure and can be written either by the provider itself, or by third parties that would like to enable data transfer to and from a provider.

Data Adapters

Data Adapters are pieces of code that translate a given provider's APIs into Data Models used by the DTP. Data Adapters come in pairs: an exporter that translates from the provider's API into the Data Model, and an importer that translates from the Data Model into the provider's API.

Authentication Adapters

Authentication Adapters are pieces of code that allow users to authenticate their accounts before transferring data out of or into another provider. OAuth is likely to be the choice for most providers, however the DTP is agnostic to the type of authentication.



Task Management

The rest is just plumbing.

The Task Management Libraries handle background tasks, such as calls between the two relevant Adapters, secure data storage, retry logic, rate limiting, pagination management, failure handling, and individual notifications.

The DTP has developed a collection of Task Management Libraries as a reference implementation for how to utilize the Adapters to transfer data between two providers. If preferred, providers can choose to write their own implementation of the Task Management Libraries that utilize the Data Models and Adapters of the DTP.

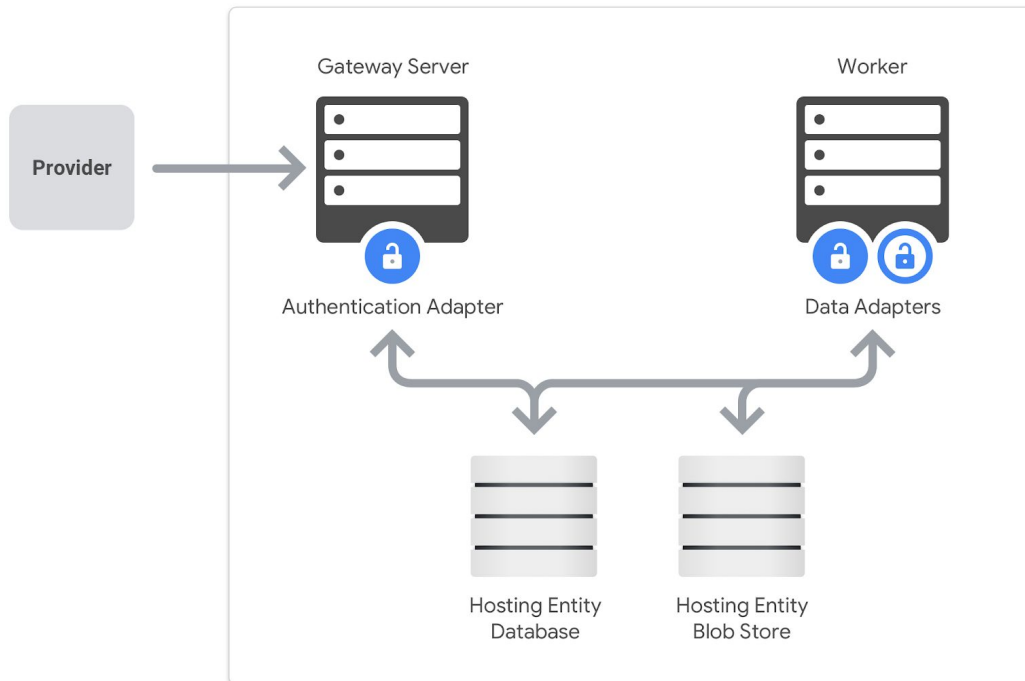


Figure 2: This graphic provides an overview of the interaction between the components of the DTP system. The gateway server facilitates the user's authorization for data export and import via their Authentication Adapter and stores encrypted credential and metadata for the transfer request in the database. A worker process is assigned to a particular transfer request and invokes the Task Management Library to coordinate and execute export and import tasks, optionally, tentatively storing data in encrypted form in a blob store between export and import.

The Task Management Libraries are built on top of a generic cloud interface so that the Host Platform can be run locally, on a corporate production environment, or on a cloud platform. The cloud interface aims to use only high level abstractions so that it can be implemented on any vendor's cloud platform.

Worker

The Worker utilizes the Task Management Library to execute the Adapters. The Worker is an isolated virtual machine that is created when a data transfer is initiated, and destroyed when that transfer is completed. The Worker generates an ephemeral key when it is created, and that key is destroyed when the Worker is destroyed.

Securely Storing Data

The security of the data passing through the Data Transfer Project is critical. An important objective of the DTP design is that Hosting Entities do not have access to a user's data either in transit or at rest. The DTP security follows industry best practices, and includes requiring transport layer security and storing all data encrypted with the ephemeral key generated by the Worker described above. The DTP system design ensures that administrators do not have access to the encryption key, which protects a user's data.

Details on security measures can be found in the core framework developer documentation on the Data Transfer Project GitHub (see [Resources](#)).

System Access

Each provider will maintain full control over determining who has access to the data stored on their systems. When a Hosting Entity runs a Host Platform of the DTP, the Hosting Entity will need to request keys from each provider they would like to be able to transfer data to and from. The DTP will not mediate data access rights between providers.

This ensures that API quotas continue to be managed by the provider, thereby helping to mitigate traffic spikes and negative impacts across providers.

While anyone can contribute Adapters to the DTP, each Hosting Entity decides which providers they will interact with. Each Hosting Entity determines which providers to request API keys from, and each provider chooses which Hosting Entity they grant API keys to. When making these choices, the Hosting Entity should consider the privacy practices of the provider, as well as its reputation, and benefit to the user to ensure the data will be used appropriately. It is up to all actors in the portability ecosystem (providers, Hosting Entities, Contributors, and Users) to be diligent in ensuring user data is handled safely.

Deployment

Deploying a Host Platform

Deploying an instance of the DTP involves several steps that allow a Hosting Entity to customize the behavior of the DTP to meet its needs.

Acquire API keys

The Hosting Entity must acquire API keys for all the providers it wants to be able to transfer data to or from. This allows the Hosting Entity to decide which providers it wishes to interact with. This step also allows each provider to have control, if they wish, over which other providers or Hosting Entities they give keys to. This step may also require the Hosting Entity to agree to the terms and conditions required by providers to use their APIs.

Deploy

The DTP is built to be run using docker images. The DTP Open Source Code Repository repository contains script to deploy the DTP to a number of cloud providers or locally (see [Resources](#)).

Create a User Interface (UI)

When deployed, the DTP exposes a Restful API that allows for HTTPs requests to be sent to it to start and manage transfer jobs. Hosting Entities will likely create a UI on top of the DTP API to allow their users to easily interact with the DTP.

Deployment Models

There are three models for successfully deploying a Host Platform of the DTP: distributed, centralized, and self-managed. Each has advantages and disadvantages and no one solution is right for all use cases.

Distributed

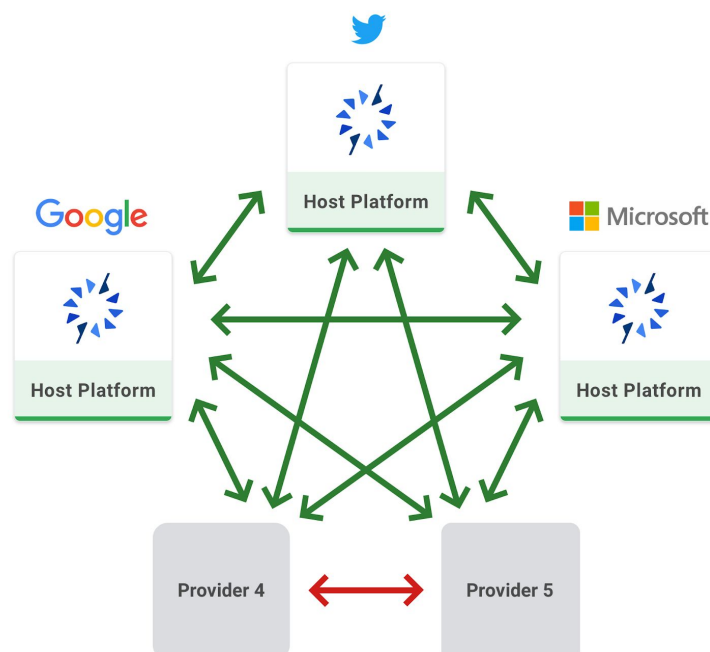
In a distributed hosting model, the provider becomes the Hosting Entity and runs a Host Platform. When a user wants to transfer data to or from a provider, the user initiates the transfer using the Host Platform of the provider they are initiating the transfer to or from.

Example: Google and Microsoft are both running Host Platforms. If a user goes to the Microsoft UI and initiates an import from Google, they would use the Host Platform run by Microsoft.

The advantage of distributed hosting is that data is never handled by a third-party. Only the source provider and the receiving provider have access to the data itself. The disadvantage is that it limits transfers to providers that are running Host Platforms.

Additionally, every provider having to maintain a separate Host Platform could create excessive Host Platform overhead costs.

Distributed DTP Environment



In this example of a distributed model, Microsoft, Twitter, and Google are each a Hosting Entity, and can transfer data from each other using the DTP Adapters and Data Models. providers 4 and 5 have the DTP Adapters but have chosen not to run a Host Platform. Users can transfer data from Google Microsoft, and/or Twitter to providers 4 and/or 5. However users can not transfer data directly between providers 4 and 5 because neither provider is running a Host Platform of the DTP.

Centralized

In a centralized system, an independent third-party is the Hosting Entity, and sets up and configures a Host Platform to handle requests to and from any provider (or a subset of them).

Example: An NGO interested in data portability agrees to run a Host Platform. Users can go to the NGO's portal and transfer data between any two providers that have Adapters in the DTP.

The advantages of this system are that many small companies will want to be part of the DTP, yet they don't have the resources or expertise to run a Host Platform themselves. With the centralized system, they will only have to write and maintain an Adapter.

Importantly, this requires that a centralized third-party be both trustworthy and highly skilled in data protection. The third-party could alter the DTP code to gain access to a user's data and thus a provider could be a target of attackers. Additionally, the third party would have to find a means of financially supporting this initiative, since hardware and bandwidth do have a cost.

Self-managed

In a self-managed environment, a user can download and run a copy of the DTP either locally on their machines or in their private cloud instance.

The advantage of self-managed is that it allows users to fully control the transfer experience. For example, users could transfer their data between providers with end-to-end encryption and not have to upload or share their private keys. Having this option also ensures that if a provider stops running a Host Platform, users still have the ability to transfer data to and from that provider. The disadvantage is that running a Host Platform is more complex and resource intensive than most users will be able to take on.

Security & Privacy

The security and privacy of user data is a foundational principle of the Data Transfer Project. Because there are multiple parties involved in the data transfer (the user, Hosting Entity, providers, and Contributors) no one person or entity can fully ensure the security and privacy of the entire system. Instead, responsibility is shared among all the participants. Here are some of the responsibilities and leading practices that contribute to the security and privacy of the DTP.

Data minimization⁴

When transferring data between providers, data minimization should be practiced. Practically this means that the receiving provider should only process and retain the minimum set of data for the individual that is needed to provide their service. The sending provider should provide all needed information, but no more.⁵

User notification

The Hosting Entity should configure their Host Platform to notify the user that a data transfer has been initiated by the user. Ideally, the user of the source and destination account are the same. However, user notification is designed to help protect against situations where that is not the case, and so notifications alerting the user of the transfer request should be sent to both the source account and the destination account. Depending on the sensitivity of the data being transferred, the Hosting Entity should consider delaying the start of the transfer so that the user has the opportunity to cancel the transfer after receiving the notification.

Rate Limiting

Hosting Entities, as well as providers, should consider rate limiting the number and frequency of transfers for a given user. This approach can help limit the impact of an account compromise. The trade off between ease of use and security with this method means there is not a one size fits all answer as to what rate limit should be set. Instead, providers and Hosting Entities should evaluate the sensitivity of the data, as well as known and possible attacks, when determining the appropriate rate limiting.

Token Revocation

When a transfer is completed, the DTP will attempt to revoke the authorization tokens used for the transfer. Providers should ensure their API supports token revocation. This approach ensures that if one of the security mechanisms is compromised, a second layer is in place to provide protection (defense in depth) to ensure that if a token is leaked, its effectiveness will be limited to the duration of the transfer.

Minimal Scopes for Auth Tokens

Providers should offer granular scopes for their authentication tokens. This provides two benefits: first, providing transparency into exactly what data will be moved; second, as a defense in depth mechanism so that if tokens are somehow leaked they have the minimal possible privilege. At a minimum there should be read-only versions of all the scopes so no write/modify/delete access is granted on the sending provider.

⁴ De Hert, Paul, Malgieri G. 'User-provided personal content' in the EU: digital currency between data protection and intellectual property. *International Review of Law, Computers and Technology*, Volume 32, 2018. <https://doi.org/10.1016/j.clsr.2017.10.003>

⁵ The DTP won't delete data from the sending provider as part of the transfer. However, participating providers should allow users to delete their data after a successful transfer has been verified.

Data retention

The DTP stores data only for the duration of the transfer job. Also, all data passing through the system is encrypted both at rest and in transit. Specifically, all data stored at rest is encrypted with a per-user session key that is created and stored in memory of the ephemeral virtual machine that is used to process a single user's job. The Hosting Entity and provider are responsible for ensuring that any stored aggregated statistics maintain user privacy.

Abuse

Providers and the Hosting Entity (if separate from the provider) should have strong abuse protections built into their APIs. Due to the fact that the DTP retains no user data beyond the life of a single transfer, and that there might be multiple Hosting Entities utilized in an attack, the providers have the best tools to be able to detect and respond to abusive behavior. Providers should carefully control which Hosting Entities are able to obtain API keys. Providers are also encouraged to have strong security around granting auth tokens. Examples of this include requiring a reauthentication or asking security challenge questions before granting access.

Shared Responsibilities Table: Security and Privacy

Task	User	Provider-exporter	Provider-Importer	Hosting Entity	DTP System
Data Minimization	Selects data to transfer	Provides granular controls of what data to export	Discards any data not needed for their service	Configure only appropriate transfer partners	N/A
Rate Limiting	N/A	Implements	N/A	Sets reasonable limits to prevent abuse	Supports provider-specific rate limiting
User Notification	Receives and reviews notification of transfer	N/A	N/A	Configure mail sender and delay policy	Send notification, optionally with delay to allow for cancellation
Token Revocation	May need to manually revoke tokens if provider doesn't support automated revocation	Support Token Revocation	Support Token Revocation	N/A	Revoke Auth tokens after use (if supported by providers)
Minimal Scopes for Auth Tokens	Verify Appropriate Scopes requested	Implements granular scopes	Implements granular scopes	N/A	Requests minimal scopes for each transfer
Data Retention	Transfer of data is not deletion, user should delete source	Store only data needed to prevent fraud and abuse	Only retain imported data in compliance with privacy policies;	Configures system to not retain any identifiable	Retains no data after transfer completed

	data if desired		Store metadata needed to prevent fraud and abuse	information	
Abuse	Protect account credentials (strong passwords, two-factor authentication, etc.)	Implement appropriate fraud and abuse protections on APIs	Implement appropriate fraud and abuse protections on APIs	Implement appropriate fraud and abuse protections on UI	Encrypts data in transit and at rest using ephemeral key; Uses isolated/dedicated VMs per transfer

Ecosystem Issues

Project Governance

As the DTP matures it may benefit from the formation of a neutral governance body. The goals of the governance body should include:

- a) Advocacy for the program
- b) Publishing and maintaining a registry website that helps discover participating providers
- c) Managing the open source repository
- d) Proposing and reporting on security and privacy best practices and enforcement mechanisms.

Inconsistent API Landscape

Despite the project’s emphasis on using open web and standard technologies, there still exist technical and public policy issues. One obstacle the DTP will face is that the API landscape is inconsistent, particularly with respect to the flexibility afforded by a service. The problem is twofold: some companies lack open APIs and others don’t support or maintain them sufficiently to enable service-to-service portability at scale. Even when companies offer different types of open APIs, they may have restrictions in their Terms of Service prohibiting use cases like the DTP from using them. By restricting how users can relocate their data, these rules discourage consumers from trying new services.

Reciprocity

A healthy data portability ecosystem has providers that allow equivalent import and export functionality. Providers that import data but don’t allow a similar level of export may pose a risk to users by trapping their data into a service. There are several possible ways to promote reciprocity in the DTP ecosystem. We have identified several methods, which we list below, and will work with Partners to further explore these and other options.

In source code

Contributions into the main Source Code Repository, hosted on [GitHub](#), are encouraged to contain an exporter coupled with each importer. This is to ensure at least an attempt at reciprocity.

Transparency

Each hosting provider is encouraged to provide aggregated statistics about problems users encountered when importing and exporting data to providers. This aims to ensure that providers are maintaining their exporting functionality to a similar level of reliability as their importing functionality.

Automated fidelity test

Hosting Entities can establish testing accounts with various providers and do periodic imports and exports of data to each service to ensure that data is exported with the appropriate fidelity. This information can again be provided in a transparent way to users at import time to ensure users can make an informed decision when choosing which providers to entrust their data with.

Data Portability Provider Pledge

Providers can work together to create a data portability pledge that requires them to follow best practices on portability. Host Platforms can seek to support providers that commit to the pledge, user interfaces can display providers that commit to the pledge to users, and reports can be published on the state of the ecosystem with regards to reciprocity.

Personal Information Management Systems (PIMS)

[Personal Information Management Systems](#) have some overlapping sets of features with this project. Both allow users to control their data and send it to different online providers.

For data minimization, we believe the DTP is a better alternative than PIMS since the DTP avoids storing the same data in multiple places. Typically, PIMS aggregate data from connected providers and create an extra copy of a user's data appended with data from multiple sources. Using the DTP, data resides in the originating service and destination service without a new copy being created at a third-party. The data only exists at the providers a user chooses.

PIMS are also subject to the risks and disadvantages outlined above in the description of the [Centralized Deployment Model](#).

The Data Transfer Project makes it easier for PIMS to interact with providers. Without the DTP, each PIMS would have to create Adapters for each data type and Vertical for each provider, thereby greatly increasing the number of Adapters they have to build and maintain.

Implementation

The success of the Data Transfer Project relies on generating an interested community of participants, including providers and users who want to transfer their data. The implementation described in this paper currently resides in prototype form on GitHub where we invite others to join the project: datatransferproject.dev. Links to specific repositories and guides are listed in the [Resources](#) section at the end of the paper.

The Open Source Code Repository for the Data Transfer project is also available on the Data Transfer Project GitHub site listed above. The repository contains a prototype implementation that supports data transfers for a handful of initial Verticals (photos, mail, contacts, etc.), Data Models, Authentication Adapters, and Data Adapters, as well as components needed to run the DTP infrastructure on two cloud Host Platforms (Google Cloud Platform and Microsoft's Azure).

The GitHub repository also contains developer documentation on how to add services to the DTP, how to integrate an existing Adapter, and how to deploy a Host Platform either locally or on a server.

Our goal is for the developer community to extend the DTP Platform to support more Verticals, providers, and Host Platform options.

Ways to Participate

Add a Provider

New providers can join the DTP using the set of interfaces described in the [EgdkYZgxiZ\gI'dc <j YZ](#) on GitHub. It requires writing a Data Adapter, Auth Adapter and, potentially, a new Data Model. In general the following steps would need to be taken

1. (If needed) Participating provider implements a new Data Model if an existing Data Model isn't sufficient. They are encouraged to use existing standards or norms (e.g. schema.org) where possible.
2. Participating provider implements Adapters to their system: Auth, Import, and Export.
3. On an ongoing basis the participating provider maintains their Adapter to ensure adequate reliability and fidelity as their service and API evolves.
4. On an ongoing basis providers should support the Data Transfer Project through actively participating in the evolution of the project.

See the [EgdkYZgxiZ\gI'dc <j YZ](#) on GitHub for more details (see [Resources](#)).

Contribute to the Core Framework

Core contributions are welcome, including support for additional providers, Host Platforms, Data Models, storage systems, etc. See the [9ZkZateZgh <j YZ](#) on GitHub (see Resources below).

Conclusion

The Data Transfer Project Contributors encourage the industry to adopt a bolder and broader view of the data portability ecosystem. We plan to work across Contributors to continue iterating on our design, fostering thought leadership in portability, and publishing updated information about our work and proposals.



Data Transfer Project

Resources

- **Data Transfer Project GitHub Site:**
 - **Open Source Code Repository** github.com/google/data-transfer-project
 - **Provider Integration Guide**
github.com/google/data-transfer-project/blob/master/Documentation/Integration.md
 - **Developers Guide**
github.com/google/data-transfer-project/blob/master/Documentation/Developer.md

Mailing list: ntp-discuss@googlegroups.com

Glossary

- **Adapters** - Adapters exist outside of a provider's core infrastructure and provide the translation between a provider's core infrastructure and the DTP environment.
 - Data Adapters translate a provider's APIs into Data Models, and vice versa
 - Authentication Adapters translate from a provider's Authentication to the DTP Authentication
- **Data Model** - Data Model is used to describe the file type and associated metadata to describe elements in a Vertical. For example, in the Photos Vertical, a Data Model would include a file type (.jpg, .tiff) and the associated metadata needed to fully describe the .jpg as a photo, such as title, description and album, if applicable.
- **DTP** - Data Transfer Project
- **Hosting Entity** - a Hosting Entity is the entity that runs a Host Platform of the DTP. In most cases it will be the provider sending or receiving the data, but could be a trusted third party that wants to enable data transfer among a specific group of organizations.
- **Host Platform** - a Host Platform is the technical environment where a DTP instance can be hosted. This can be a cloud environment, enterprise infrastructure, or local. As of March 2018, the supported cloud host platforms include Google Cloud Platform and Microsoft Azure.
- **User** - Users are any person who interacts with a provider. Users are interested in being able to manage and move the data they have stored in a provider's infrastructure.
- **Contributors** - Contributors are official members of the Data Transfer Project. They are most likely providers, but may also be organizations who are interested in enabling data transfer among their members. Contributors contribute in many ways to the DTP, including contributing code, tools, advice and insights.
- **Provider** - Providers are any company or entity that holds user data. Providers may or may not be Contributors. Provider is similar to Cloud Service Provider as defined in [ISO/IEC 17788:2014](#) section 3.2.15
- **Task Management Library** - The Task Management Library is a set of system components designed to coordinate and manage tasks that export and import data including retrieving data, pagination when applicable, handling exceptions utilizing retry strategies, re-creation of folder/container structures, and importing of data.
- **Vertical** - Verticals represent a collection of Data Models that make up a generic category of data. Some Verticals initially supported in the DTP include Photos, Calendar, Tasks, and so on.
- **Worker** - the Worker is an isolated virtual machine created when a data transfer is initiated, and destroyed when it is complete. The Worker holds the security key, and ensures that it does not exist longer than needed for the data transfer to complete.

Appendix A: Common Questions and Answers:

What type of providers are best matches for involvement in this project?

Any online service provider with consumer-facing services that contain Customer Content (as defined by [ISO/IEC 19944:2017](#)) that could conceivably be useful to that very user when using other online services.

What type of services are intended to benefit from the Data Transfer Project?

This solution is best suited to serve the portability needs of consumer-oriented services (i.e. those services having consumers as their end-users). There is no technical restriction for using this service for commercial services with professional users, but such services often have data portability requirements and associated solutions already addressed as part of the commercial agreements.

What does success look like for this project?

- A. A successful lifecycle for the DTP is one where many providers, made up of small, medium and large organizations worldwide, have come together to collaboratively build numerous data portability scenarios for end-user data, and have successfully worked together to develop Adapters to enable the direct transfer of user data between pairs of providers.
- B. Success also requires a robust technical collaboration of the Contributors to build a robust, secure, adaptable and reliable portability infrastructure that is easy to code for programmers and easy to use by end users.
- C. Success for this project should also be measured by how much it has been instrumental in raising the awareness of stakeholders in industry and user advocacy groups to better understand the challenges of direct transfer of end-user data – in the broader context of data portability – and what the best solutions and practices are to address them. We hope that such collaborative, open source development, would prove useful.

Do founding Contributors or early participants in this project have any privileges or advantages?

No, they do not, and should not. The only real advantage is that they may be slightly ahead of others in the process of jointly developing a generic and open solution framework for portability of end-user data. But the nature of the open source project is such that any progress already made is readily available for others to take full advantage of.

What if rogue actors take advantage of the technology in this open source project to trick the end users into revealing their credentials, and therefore gain unauthorized access to their data?

- A. The DTP utilizes the latest and most advanced security techniques available in the industry today to guard against unwanted and abusive use of the technology. For example, the use of API keys to control access to the utilized service APIs would inhibit such rogue efforts. The open source nature of this project would also allow the community stakeholders to review the code and identify any potential security and privacy breaches.
- B. Given that access to APIs is controlled via API keys, the Contributors would have the option of either not issuing the keys to rogue actors or revoking them.

- C. It is also worth noting that the DTP only utilizes APIs and interfaces already available, and therefore does not require the expansion of the existing attack surface of providers' publicly available interfaces.

What type of technical requirements are needed for a provider to successfully participate in this project?

A provider would need to have services that support existing APIs or other comparable mechanism to let their users access data that the users have created or uploaded to the service. The new Adapters needed to download and/or upload user data will be using these existing APIs or other comparable mechanisms offered by the participating providers. Therefore, part of the technical feasibility depends on the availability of such data extraction mechanisms.

Will this project support third-party providers who don't have their own end-user facing services, but may want to use the technology in this project to transfer user data from other providers' services in order to provide, for example, data aggregation, translation and management services?

- A. No, the technology can only be used between two providers who have agreed to send and receive their users' data. Any future, innovative use of this technology would have to be developed according to the principles of the DTP.
- B. It is also worth noting that the providers have total control of who is accessing and invoking their public interfaces to access user data (see use of API keys, for example)

Some academics and entrepreneurs predict that there will be a new era for unlocking personal data and aggregating and managing all of that up-to-date, personal data from all sources, in a manner of their choosing. Will this project help with such vision?

- A. No. This project is only about the portability of end-user data between two participating providers, at any given time. There is no provision for enabling data management and aggregation scenarios.

Which portability model is preferred, pull or push model?

- A. For a variety of reasons, including security and cost, it would be ideal for the system to work as a pull model. In other words, if a user is trying to move her data from provider A to B, the transfer would be best initiated from provider B.
- B. There may be a case where data is moving from LargeCorp to TinyStartup, and the latter does not have the infrastructure to initiate the pull, in which case LargeCorp may push the data. This should be examined case by case.

Why are the data portability requirements of professional users of commercial products and services addressed differently, and not by this project?

The DTP is only meant to support consumer-oriented services. The data portability needs of commercial services with professional users (where the service is only acting as the data processor) are addressed by their respective feature sets and service level agreements included in their commercial contracts.

What is the anticipated level of resource requirements by the new providers in the project?

The project is designed to minimize the amount of code a new provider would have to write in order to initiate new portability scenario. There is existing code that can be reused to develop new Adapters, while using the bulk of the existing necessary infrastructure and run-time components. Most components are designed to be replaceable in order to provide maximum flexibility and design choice to support as wide a variety of scenarios as possible. Therefore, new providers may end up doing as little work as they want, or as much work as they desire, depending on their requirements and scenarios.

What type of cloud infrastructure is needed for a Hosting Entity to run an instance of the DTP? Can a Hosting Entity choose their favorite cloud provider to run an instance of the DTP?

The DTP backend service is designed to run on any mainstream cloud infrastructure as long as basic features such as containers and Java VMs are supported. See the technical project documentation on the DTP GitHub site for more detailed on the technical cloud platform requirements.

Is the project code itself portable across various cloud platforms?

- A. Yes. The DTP code is written in Java and deployed via Docker containers. In addition, the modules are designed to be pluggable and replaceable and can be written in any other modern programming language, such as Python, C#/.NET, etc.
- B. The DTP code is being currently tested to run at least Google gcloud, Microsoft Azure and Amazon AWS, in addition to the already supported Google Cloud Platform.

How can a prospective Contributor join the project?

- A. Simply sign up with the open and freely available GitHub project:
github.com/google/data-transfer-project
- B. Join the DTP mailing list: ntp-discuss@googlegroups.com
- C. Introduce yourself to the other Contributors and start communicating with them
- D. Define your data portability scenarios and find another Provider who would want your common users to send and receive data between your service and theirs
- E. Start writing a pair of Adapters and publish them as part of the GitHub project



Data Transfer Project