



Artlogic Database

How we deliver our services

Introduction

If you are interested in how Artlogic provides its services and where the servers are located for auditing purposes this document tells you what you need to know about:

- the services we use
- the location of these services for audit purposes



Data storage

Artlogic use a variety of world-class cloud data service providers to deliver the Artlogic Database service (and its related assets such as images and documents) to its clients. We have made this choice in order to provide the best possible security, connection speed, resilience and functionality.

Google Cloud Platform

Our application and database servers are located on a cluster of 'Google Compute Engine' servers ('GCE') behind a load balancer and some web-serving host computers. These are high-powered servers that are *only* used by Artlogic Media Limited, and to which nobody else has access. In addition to these servers, we also store data in Google's 'Google Cloud' infrastructure, via 'App Engine' apps, utilising data storage in their 'Blobstore', 'Memcache', 'Big Query' and 'Search API' (the latter should not be confused with Google's Search Engine – 'Search API' is a private, closed system protected by strong security). All Google's storage infrastructure is 'namespaced', in other words, data belonging to one client is stored completely separately from all other clients.

Amazon S3

Some functionality within our 'PrivateViews' service stores data on Amazon S3, part of Amazon Web Services ('AWS'). Like Google Cloud Store Storage, this is a closed system protected by rigorous security.

Backups with Google and Amazon Web Services

To protect clients against data loss, we take a backup of your database at least once per day and store each backup for a minimum of 90 days. This means we can recover data you have deleted or overwritten in error if required. We run an incremental backup once a day of all your uploaded files (images, documents) in Google Coldline storage. We copy an incremental, non-destructive backup of these backup files and to a resource on **Amazon Web Services** (AWS). In this location, we could still recover a file you have deleted up to 3 months after you deleted it.

For our own server operation purposes, we take a complete copy of the hard disks every four hours using Google Snapshot to mitigate against total hard disk failure. Please note that we cannot recover individual files from these disk images without considerable effort and expense.

Google Compliance

As is expected of the world's largest web-based company, Google has rigorous security measures in place. They have independent auditors assess the security of their server centres to ensure they meet the highest internationally certified standards of data security.

For more information about compliance visit <https://cloud.google.com/security/compliance>

For an overview of Google's security visit <https://cloud.google.com/security/>

Amazon Web Services (AWS)

We also then take a further backup of our backups at Amazon Web Services for disaster recovery. Amazon Web services also conform to internationally recognised compliance schemes. For more information, visit <https://aws.amazon.com/compliance/>

Other Services

In addition to Google Cloud and AWS, we also use the following third-party providers:

- Cloudinary – for storing, scaling and serving images and for creating and serving PDF documents ‘on demand’. Cloudinary use the Akamai CDN infrastructure to distribute assets globally so that the performance of your Artlogic Online instance is optimised to your location.
- Pusher (<https://pusher.com/>) – for real-time ‘web sockets’ based notifications, used for feedback in image uploading, and for the ‘Guest list’ feature in Marketing & Events. Please note that we never send your private data via pusher – we simply use it to trigger interactive events on the server.
- SendGrid – for sending email campaigns. Email activity is logged at SendGrid and on Google ‘Search API’ as mentioned above.
- CloudFlare is a service that caches public web page data and speeds up asset delivery by using Content Delivery Networks. Clients mostly pay CloudFlare directly to improve the performance of websites, especially when they may be subject to bursts of popularity. Artlogic Database material used for public presentations as Online Viewing Rooms may be stored on CloudFlare for a short time.

Locations

The Google services we use are located in the European Union (GCE servers) and U.S. Central (Blobstore and Search API). Within the E.U., the GCE servers are located in Belgium. As they open new datacentres, we may relocate some or all of these services to the UK. Some of the secondary services may be more widely distributed, for resilience and performance.

The Amazon Web Services storage are located in their datacentres in Ireland (E.U.) or in the USA. Amazon are opening datacentres in the UK and we may relocate some services here.

Cloudinary are officially located in Israel but use the global Akamai service so their services should be considered global.

Similarly, CloudFlare are headquartered in San Francisco but use servers all over the world.

Pusher are based in the U.S.. As mentioned above, we do not send private data via Pusher.

SendGrid are based in the U.S. but will themselves use cloud services from Google or Amazon and of course emails sent from your system should be considered ‘out there’ as they will be distributed to email servers and clients globally.

Artlogic offices rarely ever have copies of client data on site. This may happen if a client has sent files on a hard drive for evaluation or to work on import scripts. Nearly all work developing new features or debugging a reported problem that uses data and much of the data migration activity takes place on a cloud-based developer environment. Our development and client teams only access client data from within our network (using two factor authentication) for the purposes of responding to queries or resolving issues that might affect your users.