Snapbar Studio is a tool that allows clients to capture and distribute employee or member headshot photos. Studios are created and customized on a per-client basis for their specific needs and distributed to their users either by an internal invitation email or publicly available URL.

Types of user data we collect

We refer to a user submitting photos to a studio as a “Contact”. We only store necessary data based on the client’s needs. Most contact and photo data persists until a client closes their account or ends a subscription to the service.

Standard fields for contacts

For the purposes of properly tracking and organizing photos we collect the following fields by default for a contact:

- First name
- Last name
- Email
- Phone (optional)

Custom fields for contacts

In addition to these basic fields, we allow each client to create custom contact properties that can be configured on each account. Any information may be collected, but should not include special personal data (race, ethnic origin, religious beliefs, etc). It is the responsibility of the client to only collect the data required for their use case.

Contact photos

As a part of the core offering within the app, we offer the ability for a contact to submit photos taken on their personal device. These photos most commonly are a detailed image of the contact’s face and computer-edited after capturing. A user always has the chance to review a photo before it is submitted.

User and device information

We may also record a user’s IP address, browser details, interaction behavior when they submit photos through a studio. We use this information strictly for internally debugging and improving the product experience. These types of logs are most commonly retained for a maximum of 60 days (but often less).

User authentication

We never store passwords or other authentication details for contacts. However, we do require a contact to verify their email address before they may proceed with a photo submission. Email verification happens either through a private invitation link or a self-service verification email.

For group coordinators and admins, authentication is provided by Snapbar Hub. Snapbar Studio never directly stores an admin's password, but rather uses an OAuth2 process of obtaining a token for API access.

System infrastructure

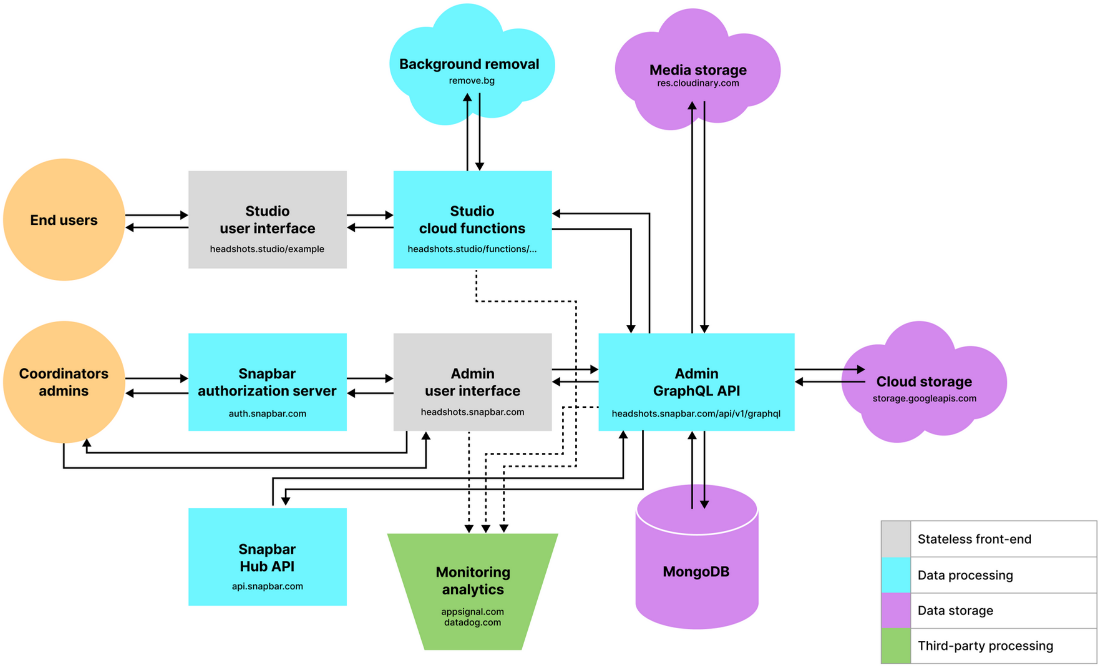

We prefer to use managed Cloud resources whenever possible for the services we provide. This means we don't implement or maintain local servers or networks, which allows us to focus primarily on the security and experience of our own applications without a major concern for network privacy and security.

Google Cloud Platform

We use Google Cloud as our primary Cloud provider. They provide services such as VMs, networking, cloud functions, and bucket storage. All data stored through GCP is by default encrypted at rest and transported with an encrypted TLS connection.

Netlify

All of the contact-facing Studio experiences are hosted on Netlify, which provides the interface by which a user may submit photos. We use cloud functions to securely communicate with our backend API to fetch Studio details, send an image for background removal, and to deliver a photo submission.

Cloud66

We use Cloud66 to manage our DevOps for our backend infrastructure. This service allows us to deploy our application code to certain Cloud servers with their Maestro service, which under the hood uses Kubernetes to provision servers with a variety of different services.

MongoDB Atlas

For database storage, we use MongoDB Atlas, which provides permanent data storage for our application. Servers are provisioned within Google Cloud's centers and communicate with our VMs with a VPC peering connection, which prevents non-authorized networks from connecting to the database. All data stored through MongoDB Atlas is by default encrypted at rest and transported with a private encrypted TLS connection.

Cloudinary

Media (such as photos, animations, and videos) are stored in Cloudinary and served from their CDN. We also instruct Cloudinary to transform images into different sizes, formats, colors, etc. Images have publicly-accessible URLs, but are signed with a unique identifier within the URL that prevents tampering or common access. Cloudinary uses TLS for connections and stores media encrypted at rest.

Monitoring / Analytics

We utilize analytics services AppSignal and Datadog for application error and performance tracking.

Exporting data

When a client requests photos or contact data to be exported, we will process the requested data and store within a secure Google Cloud bucket. As a logged in user with proper permissions, the admin may access a page where a signed URL is generated and provided to the client. The URL is intentionally short-lived (~1 hour) so as to avoid accidental public exposure. The object is auto-deleted from the bucket after 10 days.

DNS

DNS is hosted via Cloudflare.

Provider Selection Process

We take care to ensure that any third party provider that we use meets our obligations as a data controller/processor. We look for GDPR compliance in each and review their security practices as well.

We are working towards executing a data processing agreement (DPA) with each and a way for us to notify our clients of changes to our subprocessors.

Transport Layer Security

Every connection between services in the diagram above is secured via TLS. In most cases, we are relying on the TLS best practices of each cloud provider.

- Netlify
  - Documentation: https://docs.netlify.com/domains-https/https-ssl/
  - Issuer: Let’s Encrypt
  - SSL Labs Report (Grade A)
- Google Cloud
  - Documentation: https://cloud.google.com/load-balancing/docs/ssl-certificates
  - Issuer: Google
  - SSL Labs Report (Grade B, supports TLS 1.0/1.1)
- Cloudinary
  - Issuer: Go Daddy
  - SSL Labs Report (Grade B, supports TLS 1.0/1.1)

Software Development Process

Change Management

Changes to the product are tracked in GitLab issues. The Product team follows a conventional Scrum process with two week sprints.

Source Code Management

All source code is hosted in GitLab. We enforce access controls over who can read/write to the repositories and also protect certain branches to enforce code review of changes.

Code Review

Code Review is a critical part of our software development process. All changes are reviewed by another engineer who is familiar with that product. Engineers are expected to review every single line, look for proper testing, question code that is unnecessarily complicated or obscure, suggest performance improvements, and validate any third-party dependencies that are added for licenses, project health, relevance, etc. Continuous integration is run during code review, looking for a full suite of passing tests, style compliance, and language-specific linting. Any failures block the pull request from being merged.

QA

In addition to automated testing, we have a QA team that is responsible for performing exploratory testing with new features before their release. Feedback from this determines whether completed features are ready for release.

Packaging

We build Docker-compatible images of the API and authorization services that we deploy. These are carefully crafted to only contain the necessary libraries and tools to run in production. Cloud66 is responsible for building these images upon CI/CD pipeline completion. We have build logs to review the steps that go into these built images. Images are stored by Cloud66 with BuildGrid.

Dependency Updates

We aim to keep dependencies patched on a quarterly basis or sooner. The same goes for rebuilding from updated Docker base images.

Data retention

A contact has a right to the data they store and may at any time request its access or deletion. A client is responsible for reviewing these requests and processing the request through the Snapbar Studio admin portal. In certain cases, it may take up to 30 days to fully eradicate the data requested for deletion.

When a client's account is closed, we'll provide up to 30 days before we permanently delete all data associated with the account. This includes studio preferences, contacts, submissions, and other settings.

Netlify

No personal data is retained here.

Application server and database

Data stored within the database are retained until a client either deletes a resource or closes their account. Backups are kept for the MongoDB server for 21 days and we have the ability to recover to a point in time. Access logs are retained for 30 days.

Cloudinary

Media is retained for as long as the Studio is active and until a Snapbar staff member removes the collection associated with the Studio. Backups are currently not made of these.

remove.bg

remove.bg does not retain images that we send to them after performing background removal.

AppSignal and Datadog

We're working on removing any personal information from these logging platforms. Datadog logs are retained for a maximum of 15 days and AppSignal a maximum of 60 days.

Information Security Practices

Regular Training

Regular sessions are held where all employees are trained on security best practices. Topics include password usage and sharing, phishing, computer security, etc.

Password Policy

We have a well-defined password policy that we expect all employees and contractors to follow.

Full Disk Encryption

We have a policy that all employees configure full disk encryption of their issued computers so that the risk of data loss due to theft is minimized.

Non-Disclosure Agreements

All employees and contractors sign a non-disclosure agreement.