Segment

Security Essentials Overview

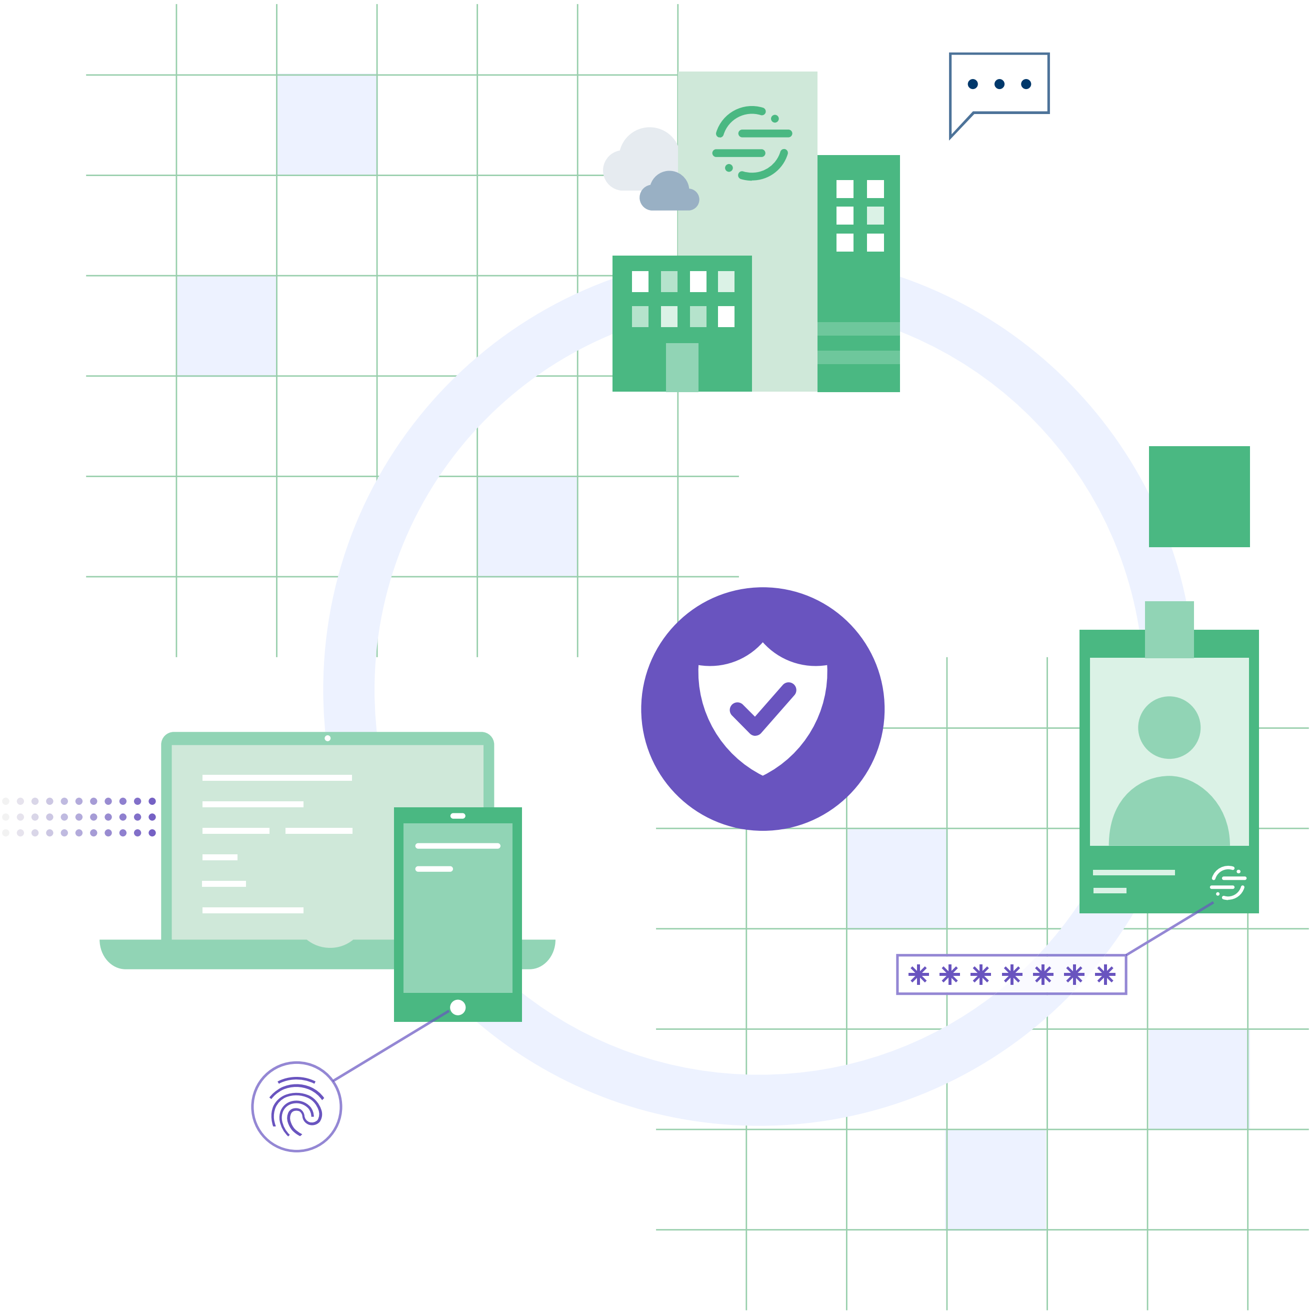

Overview

At Segment, security is a top priority. Segment customers trust us with a significant amount of data, and we do not take this responsibility lightly.

To earn our customers' trust, we handle data with the utmost care and integrity. Whether it's encrypting your data over the Internet and at rest, or our commitment to open-sourcing core pieces of our infrastructure, we want you to have confidence in how your customer data is being collected, transported, and stored.

In addition to our public resources, we frequently get questions that fall into a handful of major groupings surrounding your data and who has access to it. These are the same questions we ask of third-party services handling our data, so we appreciate that you're asking. We're hopeful that our answers in this document will give you a better view into the exact mechanics we use to secure your data within our internal systems.

Securing data in transit

Segment acts as a routing and storage layer for your data. That means data goes through three distinct phases:

1. Data is **COLLECTED** via our libraries and Sources
2. Data is **ROUTED AND TRANSFORMED** inside the Segment infrastructure
3. Data is then **SENT** to your data warehouses or your enabled integrations

Let's go through each portion in turn.

Data collection

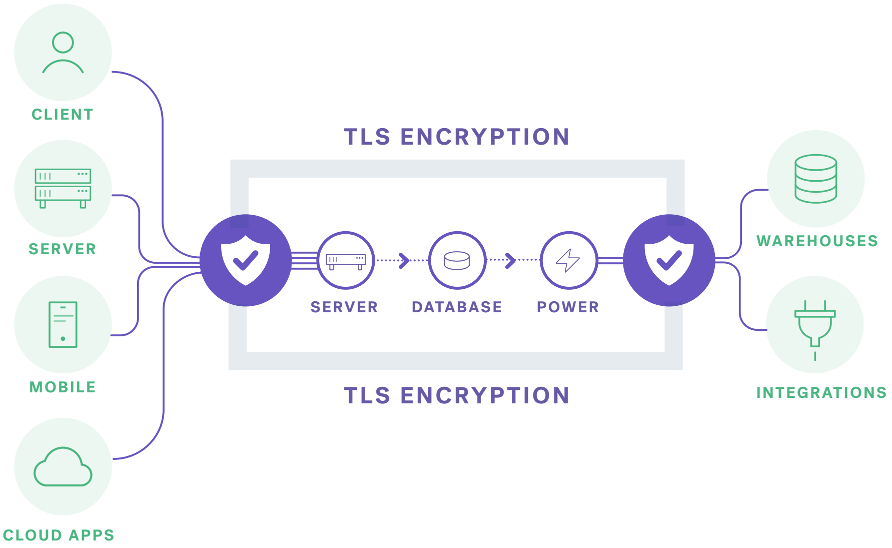

Segment collects your data from two places: 1) **your application code** and 2) **your Cloud apps**, like Zendesk, Stripe, Intercom, and more.

Application data sent to Segment is transmitted via our client and server libraries. These libraries are all open-sourced and available on GitHub so you can investigate further if you'd like.

All data sent is opt-in, so we only collect the data you specifically instrument via our API. We don't automatically pick up fields like passwords or text inputs. Additionally, our Analytics.js library will automatically pick up page metadata, such as page title and URL. Our mobile libraries will capture common user lifecycle events, such as "Application Opened," **but only if you opt-in to this feature**.

Data from **Cloud apps** is collected via Source integrations that run on your behalf. In order to collect this data, you must OAuth with the Cloud app vendor and grant Segment-specific access to the tools.

In cases where Cloud apps do not support OAuth, we will ask you to grant Segment access via API keys.

Both application Sources and Cloud app Sources send data securely to our TLS endpoint.

The Segment event-ingestion API supports the use of the most up-to-date Amazon recommended cipher suite and supports **TLSv1, TLSv1.1,** and **TLSv1.2**. The Segment app supports only **TLSv1.2.**

Data routed within Segment's infrastructure

Data routed within our infrastructure happens inside an Amazon Virtual Private Cloud (VPC). It is transferred within a private subnet that is inaccessible from the public Internet.

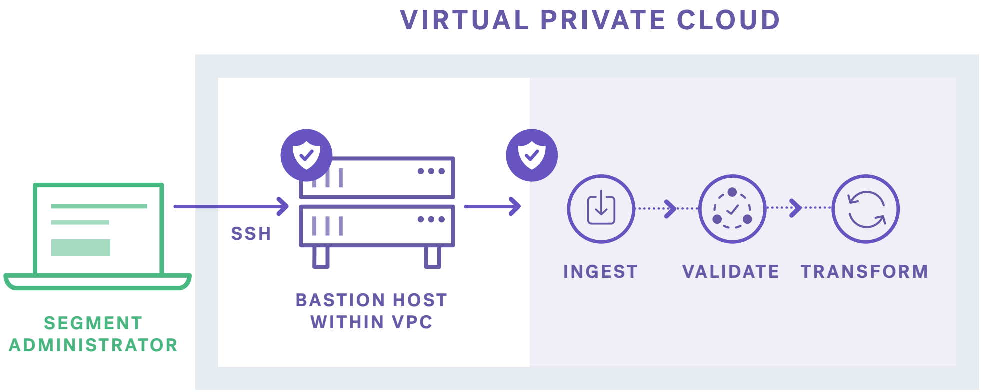

Aside from our bastion nodes, only instances inside the VPC are able to access one another. Public-facing servers as part of our API, CDN, or Web app all receive traffic through public load balancers managed by Amazon. Additionally, security groups are based upon CIDR ranges.

Data sent to warehouses and partner integrations from Segment

Your data is protected on its way to your integrations. It's sent over TLS as specified by the customer. We've open-sourced many of these integration points inside our segment-integrations Github organization, so you can take a look at them as well.

For piping data into your data warehouse, we will authenticate with access granted to a dedicated Segment user you create. For streaming integrations (tools for marketing, analytics, etc.), we use a shared API key provided by the integration.

Securing data at rest

There are two core types of data that Segment manages: 1) **metadata** related to API keys and authentication, and 2) **customer data**, which is collected via our API.

Metadata is stored within our Mongo database. Customer login passwords for the Segment application are salted and hashed using the industry standard bcrypt. All metadata is backed up hourly to avoid data loss and service disruption.

Customer data, or event, object, and trait data sent through Segment, is securely stored in Amazon S3. This archive is used to power our warehouses integrations, give you access to data export, or to replay historical data into new tools. The data is logically partitioned by the ID of the data source, and access is controlled via IAM permissions. Data in S3 is encrypted at rest using Amazon's SSE offering and replicated to a separate bucket for redundancy. Customers can elect to have this data deleted at any time (note: this will limit the use of Replay).

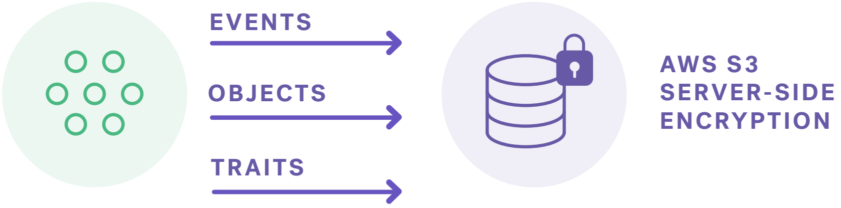

For customers using Segment to send data to a warehouse (Redshift, Snowflake, etc.), your data will be stored inside the data warehouse that you own, manage, and secure.

Securing data access

From a customer perspective, there are several ways of managing access to your data. While there is no means of getting raw data directly from Segment's infrastructure, we do give you the ability to receive raw events via our S3 integration, direct webhooks, our events debugger, and warehouses integrations. Access to this data is controlled via the Segment app, and we provide customers the ability to manage this access on a per-user level.

We designed our access management system with the principle of least privilege in mind. This means that your team members can access the information and resources they need in Segment, but nothing more.

If you have granted a user access to your Workspace as an owner, he or she may modify the downstream settings for any Source within that workspace. If a user has only been granted access as an admin for a single Source, that person can only make modifications to that Source, and to the integrations for the data coming from that Source.

For Business Tier customers, we also offer more granular roles that enable you to:

- Assign **admin** or **read-only** access to **specific Sources** or **all warehouses**
- Control who can make changes to or view **Protocols** and **Personas** within your team with **admin** and **read-only** roles

Read-only members can view assigned Source(s) or sub-resource(s) to see their layout, settings, and live data in the events debugger, but they cannot create or modify any Sources or integrations. Read-only members can also view Workspace settings, but cannot modify any settings.

On Segment's side, members of the engineering and success team have access to your data for debugging purposes only. The following safeguards are in place to ensure this access is strongly protected:

1. All production access is federated through SSH and IAM.
2. The only nodes open to the outside world are our bastion nodes which are secured with individual user auth and individual SSH keys.
3. Our employees have no IAM user keys and can only access AWS through our

identity provider, Okta. In order to perform administrative actions, users must use multi-factor authentication (MFA) to authenticate with our identity provider and obtain temporary credentials using the AWS AssumeRole API. This was important enough to us that we wrote about it and open sourced aws-okta.

Continuous monitoring and updating

Amazon Web Services (AWS) serves as the front-line of protection for all network-level attacks, like DDoS. AWS manages the software and selection of cipher suites on our load balancers and will automatically flag and block malicious behavior.

To ensure that Segment adheres to security best practices, we maintain relationships with our vendor-security consultancies, who perform regular application-security assessments on our highest-risk applications. To cover any additional vulnerabilities, we run a bug-bounty program, which attracts individual researchers and allows them to safely disclose potential vulnerabilities. Critical vulnerabilities are handled within 1 business day.

We also keep audit logs provided by AWS Cloudtrail on all user administrative actions.

Fortified infrastructure

Our infrastructure is entirely provisioned on AWS and Google Cloud Platform (GCP). Our Amazon infrastructure is contained within Segment managed VPCs (Virtual Private Clouds) and provides total isolation from other instances in the same datacenter. We use Google Cloud Platform for managing parts of our Personas infrastructure. We have a single GCP project for production that provides total isolation from other tenants on Google Cloud.

Segment maintains many AWS accounts for better security isolation and reliability guarantees. Segment maintains a separate development and staging AWS account for our engineers to test their changes before deploying to production.

The gateways to our infrastructure are our bastion nodes. These nodes expose only the Secure Shell service (SSH) and are the only nodes we expose directly to the outside world. SSH is used by our employees to access and administer our infrastructure.

Developers can log in to our infrastructure using their own SSH keys. Keys are generated and stored locally and federated through Foxpass.

The only other publicly facing listeners are our load balancers, which are managed by AWS.

Once traffic is inside our infrastructure, there are two places it can be routed: the **public** subnet and the **private** subnet.

Nodes in the **public** subnet are assigned public IP addresses. This is where our **integration** and **warehouses** instances live. They are restricted from exposing any ports through AWS security groups and are located on the public subnet so that they can use direct networking with each of our partners.

All other instances live in the **private** subnet and are not publicly addressable. Traffic here goes through NAT instances before making it to the public Internet.

All instances live within our VPC and are only accessible from the IP ranges listed within that VPC.

Internal tools are secured via OAuth and SSO using our identity provider, Okta, as the authentication backend.

Commitment to security compliance

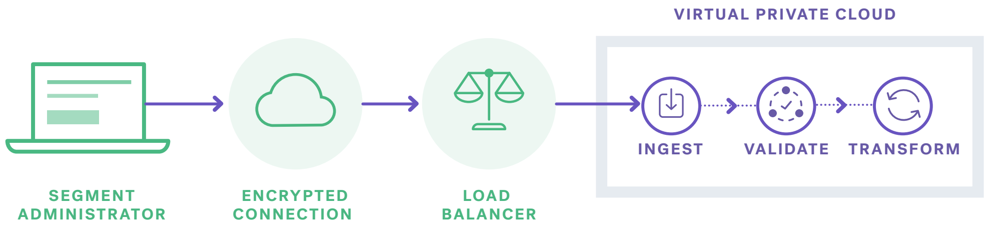

In alignment with Segment's commitment to the privacy and protection of customer and corporate data, we have developed a comprehensive **Information Security and Privacy Program (ISPP)**. The Segment ISPP is structured in alignment with ISO 27001 and 27018 guidance and is continually enhanced to align with new and evolving regulatory requirements such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA).

As a demonstration of our commitment to privacy and security, Segment is certified under the EU-US and Swiss-US Privacy Shield Frameworks and ISO/IEC 27001:2013. We have also completed the Cloud Security Alliance (CSA) Self Assessment CAIQ. Finally, Segment complies with ISO/IEC 27017:2015 and ISO/IEC 27018:2014. ISO 27017 has enhanced security controls for cloud services. ISO 27018 has privacy-specific, international guidelines to protect Personally Identifiable Information (PII) in the cloud.

When it comes to security, our focus at Segment never wavers, and our teams devote themselves to its practices to ensure the privacy of your data.

Certifications

ISO 27001 International Organization for Standardization

ISO 27017 International Organization for Standardization

ISO 27018 International Organization for Standardization

Privacy Shield Framework

CSA cloud security alliance®

For additional questions, contact us at: https://segment.com/contact
Or, visit our Security landing page at: https://segment.com/security

ABOUT SEGMENT

Segment provides the customer data infrastructure that businesses use to put their customers first. With Segment, companies can collect, unify, and connect their first-party data to over 250 marketing, analytics, and data warehousing tools. Today, thousands of companies across 71 countries use Segment, from fast-growing businesses such as Atlassian, Bonobos and Instacart to some of the world's largest organizations like Levi's, Intuit and Meredith. Segment enables these companies to achieve a common understanding of their users and make customer-centric decisions.

Levi's · intuit · 21ST CENTURY FOX · New Relic. · REUTERS