

Rudderstack








by VP of Product

Customer Data Platform

Details

Review Date	12/07/2023
Purchase Date	Q3'22
Implementation Time	< 1 week
Product Still in Use	Yes
Purchase Amount	\$12k / year for 5M events
Intent to Renew	0%
Review Source	Elion

Product Rating

Product Overall	 5.0
Use Case Fit	 5.0
Ease of Use	 4.5
API	 5.0
Integrations	 5.0
Support	 4.5
Value	 4.0

About the Reviewer

Purchasing Team Implementation Team
Product Oversight

Reviewer Organization

Virtual-First Provider Behavioral Health

Reviewer Tech Stack

SimplePractice CertifyOS BigQuery

Other Products Considered

Freshpaint Segment

Summary

- **Product Usage:** The reviewer uses Rudderstack as a customer data platform to centralize tracking data from various marketing campaigns.
- **Strengths:** Rudderstack's strengths include its extensive library of integrations, reliable operation, similar design to Segment, and an event-based pricing model that fits the reviewer's economic model.
- **Weaknesses:** The reviewer has identified Rudderstack's documentation as a weakness, mentioning that it lacks details on nuances and data model idiosyncrasies; they also note that their current pricing tier doesn't provide access to historical event logs.
- **Overall Judgment:** The reviewer believes that they made the right decision in choosing Rudderstack due to its cost-effectiveness, extensive features, and easy integration with other systems.

Review

So today we're chatting about Rudderstack and how it's used at your company. Before we jump into that, could you give a brief overview of the company and your role there?

We are a virtual behavioral health clinic, and I'm the VP of product.

What was the need that drove you to look for a customer data platform (CDP)?

We run multiple marketing campaigns to attract patients and needed a CDP to centralize tracking data from various channels. This data then feeds into our marketing tools, including email and product analytics, to evaluate channel performance in patient acquisition.

You were using Segment for CDP. What drove you to change from Segment to Rudderstack?

There were two reasons. First, Segment's pricing model is based on MTUs (monthly tracked users), and since we attract a high volume of MTUs to our website through free content for patient acquisition, this was cost-prohibitive. Rudderstack's event-based pricing was more aligned with our economic model. Second, since Segment's acquisition, we've noticed a slowdown in innovation. We preferred partnering with a company like Rudderstack that remains highly focused on its core product and business.

Which vendors did you evaluate as potential replacements for Segment, and how did they compare?

We wanted a CDP that offered comprehensive data integrations with the services we used. Rudderstack matched Segment with nearly 100% coverage, whereas Freshpaint lacked sufficient integration capabilities at the time. Additionally, we required a platform with flexible and well-documented APIs to facilitate back-end server API calls for data ingestion. Another key feature was the ability to filter out unwanted data from being sent to specific downstream providers, a functionality that Rudderstack provides through their transformations. We also considered PostHog, which has built-in analytics alongside its other features.

How did you find the sales process overall with Rudderstack?

The sales process was very reasonable. We had an initial discussion with a sales representative followed by a technical session with a specialist who addressed our questions about the platform's capabilities. The technical expert was knowledgeable, and their familiarity with organizations transitioning from Segment to Rudderstack was reassuring. They provided case studies of similar companies to ours that had switched without experiencing any downtime.

What was the onboarding and setup process like?

Building the initial proof of concept with Rudderstack was straightforward, as it replicated Segment's functionality closely. The onboarding and setup was primarily a self-service process, with support available through a dedicated Slack channel. The plug-and-play integrations, like connecting to Mixpanel, were simple, involving simply copying the project ID and secret key into Rudderstack's fields. API integrations pretty much matched Segment's format, so we just needed to update the URL endpoint.

What are the use cases you're using Rudderstack for?

Rudderstack helps us maintain HIPAA compliance. When data is ingested into Rudderstack, we have transformations set up to anonymize personal identifiable information (PII). This ensures that when we pass the data to downstream vendors we don't need a Business Associate Agreement (BAA) with them to stay compliant. These transformations, using Python or JavaScript, allow us to mask or hash sensitive data before sharing it with external parties.

What other Rudderstack features are you using?

Rudderstack primarily serves as a data routing platform without built-in analytics features. It allows you to collect data from various sources, transform it, and then direct it to the desired destinations. A key aspect of Rudderstack is its warehouse-first CDP approach, which enables us to stream data into our data warehouse, Google BigQuery, in near real-time. This provides us with a well-structured single source of truth for all customer interactions.

What do you see as the relative strengths and weaknesses of Rudderstack?

Rudderstack's strengths include its extensive library of integrations that we've found compatible with every tool we've needed. The platform's design is nearly identical to Segment, enabling a smooth migration. Additionally, Rudderstack has demonstrated excellent reliability, with no noticeable data loss despite our high volume of event throughput.

As far as weaknesses, Rudderstack's documentation is less comprehensive than Segment's, offering basic guidance but lacking details on nuances and data model idiosyncrasies. Furthermore, our current pricing tier doesn't provide access to historical event logs, which limits our ability to troubleshoot historical issues, though we can activate a debugging mode for real-time event monitoring.

You mentioned that the platform is reliable and stable. Were there any bugs you encountered?

I personally didn't encounter any bugs within the platform.

How was the developer experience?

The developer experience with their API was pretty good, mainly due to its logical structure. Much of the code implementation involved simple drag-and-drop actions, requiring only modifications to the downstream URL and the authorization token. Access to Rudderstack via Slack was great for direct inquiries about API integrations, with their technical team providing competent responses.

Do you have any tactical advice for organizations that are implementing Rudderstack on how to do it more easily?

From a data model perspective, it's important to map out the essential business data you need to collect and determine which downstream vendors will receive each data piece. Establish a master definition list and a comprehensive model for all customer data components. For HIPAA-compliant companies, this approach allows for a streamlined data egress process, enabling a single transformation for all PII, which can then be distributed to all vendors, rather than crafting individual transformations for each one. And organizing your data from the outset simplifies future attribution tracking, such as multi-touch attribution, as opposed to trying to make sense of disorganized raw data formats in your data warehouse later on.

You're moving toward self-hosting Rudderstack (and not renewing the contract as a result). How did you come to that decision, and what is that process like?

Rudderstack's cloud-hosted product requires a minimum annual expenditure of \$25,000 to sign a BAA, which was not cost-effective for us. It also provided excess event capacity we didn't need. However, opting to self-host Rudderstack means we retain complete data ownership and remain HIPAA compliant, since our IT infrastructure already meets those standards.

Rudderstack is open source and offers documentation on self-hosting. We utilized their guidelines to set up a Kubernetes environment on our existing servers. Prior to this, we connected with several other companies who had successfully transitioned to self-hosting and were handling over 5 million events a month. This gave us the confidence to proceed with self-hosting Rudderstack.

The user interface remains unchanged with the self-hosted version, and we still manage it through Rudderstack's control plane. The primary limitation is that self-hosting only includes five transformations for data processing, whereas the paid version offers substantially more.

How has the account management and support with Rudderstack been?

We primarily interact through the Slack channel and have used it a few times. Each time, the support provided has been of high quality and very technical. Regarding general account management, at our pricing tier, we need to initiate contact for assistance, as proactive check-ins are infrequent.

Looking back, do you feel like you made the correct assessment in going with Rudderstack?

I do. From a cost perspective, certainly, as it's essentially zero cost after self-hosting. Product-wise, it fulfilled all our needs and integrated well with the required systems. Comparing it with other vendors, considering costs, feature completeness, and ease of integration, I would still rank Rudderstack at the top across most of those dimensions.

Do you have any recommendations for the Rudderstack team?

First off, if they had agreed to sign a BAA, we would have maintained our paid subscription, as we were satisfied with the product. We're now taking on some technical risk by self-hosting; we'd prefer to pay for their cloud service. And it would be helpful if they enhanced their documentation with more detail and nuance.