

SegMetrics Quantifies Touchpoint Values throughout Customer Journeys with Time Series Analytics from SingleStore on AWS

1B

Tag apps parsed day

50GB

of daily UPSERTS

1M Rows

ingested per second

5,000

concurrent queries

200ms

average query speed

4TB 12Hrs

Only 12 hours to migrate 4TB to SingleStore via Arcion Cloud

“On the old system, we couldn’t bring in too much data. We had a large enterprise account sign up as a client, bringing a million web views a day. We had to add smart filtering to ensure it wouldn’t make our database explode. With SingleStore, I don’t have that concern anymore.”



Keith Perhac,
Founder and
CEO, SegMetrics

SegMetrics offers a martech application that bridges gaps between the top, middle, and bottom of the funnel to give digital marketers an unrivaled view into how each step in their marketing strategy is performing: who is engaging with their messaging and who is not. The platform gives marketers the insights they need to optimize campaigns and customer lifetime value (CLV), progress leads down funnel, and improve their bottom line.

The SegMetrics platform works by pulling in huge volumes of tagged data from disparate sources, like Stripe, Shopify, Facebook and CRM systems into a singular customer journey for each of its clients’ end customers.

Challenges/Goals

Most marketers spend a lot of time mastering website analytics, learning all the different variables, terms, and metrics. With SegMetrics, marketers gain the ability to look at how customers progress throughout the entire marketing journey, not just things like website interactions. SegMetrics works with thousands of clients every year, and what it found challenging was ingesting terabytes of data at any one time and delivering real-time results to each client.

SegMetrics is similar to companies like Fathom Analytics in that it pulls a huge volume of data from web tracking and CRM systems like HubSpot and Salesforce, as well as purchase and refund data, into its database daily, via APIs enabling nearly 100 integrations. The data volumes are significant. For example, when one of its clients tags a contact in its CRM based on a behavior or action, that contact can end up having up to 200 tags. Multiply that by the hundreds of thousands or millions of people in a CRM system, and suddenly SegMetrics is importing and parsing over a billion tag applications per day across its client base.

“Our challenge has been the amount of data we’re processing. SegMetrics pulls in many different data types: click data, CRM data, purchase data, refund data. Essentially we pull in data from every source our clients have that touches and tags their customers, so SegMetrics imports and parses about a billion tag applications per day. The issue we ran into with our previous database was the inability to do this in a highly performant way.”

“SegMetrics has API connections to a lot of our clients’ data systems. We pull in and normalize the data into what we need, parse the data into SQL files, and then run them against the database. We can run 200,000 rows at a time against the database to UPSERT them in. So for example, if someone in a client’s CRM system updates the time they want to meet with their customer or updates the industry, SegMetrics catches all that and then pushes the data into our system,” Perhac said.

In its previous system, a managed DigitalOcean MySQL database, SegMetrics had a limitation on the amount of data that it could ingest, and parsing speeds were very slow. Getting data into DigitalOcean was fast because it was a non-distributed, single source insert — but getting the data out as the database grew meant double, triple, or 4X the cost. “We were at, I think 80 or 90% of the database size, and our next step up would have raised our pricing to about \$10,000 a month,” said Perhac. “We were a big customer for DigitalOcean, which concerned me because at the time we were only at five terabytes.”

“We pull in data from every source our clients have that touches and tags their customers, so SegMetrics imports and parses about a billion tag applications per day. The issue we ran into with our previous database was the inability to do this in a highly performant way.”

— Keith Perhac, Founder and CEO, Segmetrics

SegMetrics found itself out of data. “We were at a specific size, and then we had to upgrade to the new server in MySQL, in DigitalOcean, and it took 72 hours to upgrade. Essentially going from an S2 to an S4 took three days for DigitalOcean, which was not great for us,” Perhac said.

Technology Requirements

Two compelling events prompted SegMetrics to look for another database. One was the slowness of the MySQL on DigitalOcean solution, as a single customer could slow it down the reporting servers considerably if they were running multiple heavy queries. The second was that whenever SegMetrics was low on database space or various issues popped up, it could take days to provision a new instance to get the support they needed.

“We were at a specific size, and then we had to upgrade to the new server in MySQL, in DigitalOcean, and it took 72 hours to upgrade. Essentially going from an S2 to an S4 took three days for DigitalOcean, which was not great for us.”

— Keith Perhac, Founder and CEO, Segmetrics

SegMetrics wanted a solution that would enable it to increase the amount of data stored while maintaining compute. Pricing and performance were also of top concern, as was finding a partner that could facilitate the migration from MySQL on DigitalOcean to a new solution.

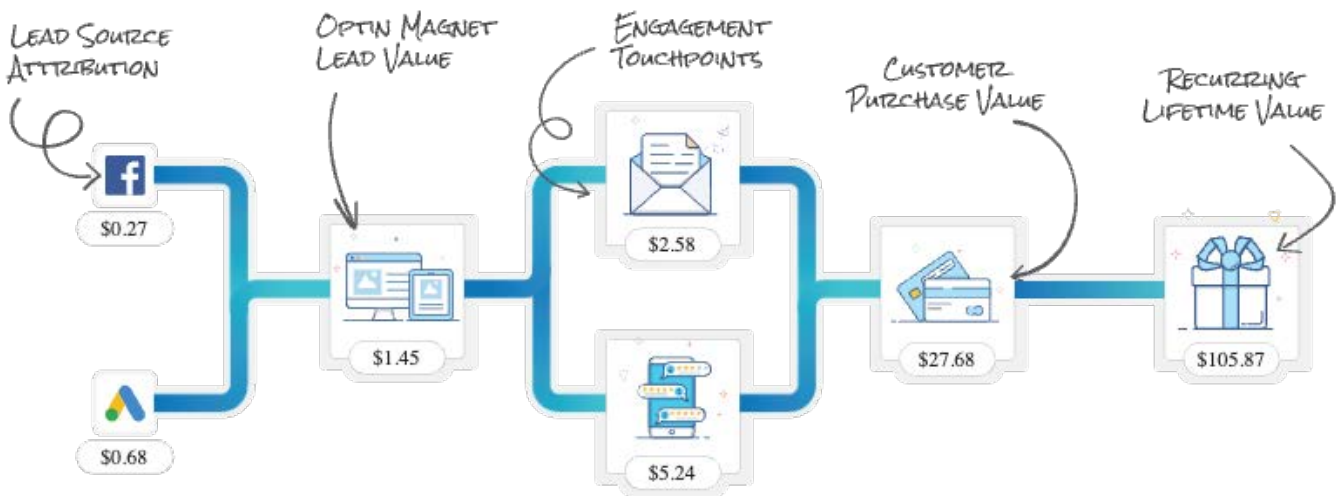


IMAGE | SegMetrics helps companies optimize for revenue at each touchpoint in their marketing funnel



Technology Requirements (con't)

The new database needed to:

- Enable ultra-fast ingest for millions of events per second and provide instant availability
- Deliver sub-second latencies with immediate consistency
- Provide universal storage for transactions and analytics
- Facilitate unlimited storage with separation of storage and compute in an operational - analytic database
- Be compatible with MySQL
- Work well with database replication partner solution Arcion to facilitate database migration

Why SingleStore

Perhac and team evaluated AWS Aurora, AWS Redshift, PlanetScale — and SingleStore, thanks to Perhac's thanks to Perhac's colleague and friend Aaron Francis, Co-Founder of Hammerstone Dev. Aaron also introduced Keith to [Jack Ellis](#), Co-Founder of [Fathom Analytics](#) and a SingleStore customer. "I was talking with Jack Ellis about our need to migrate from MySQL on DigitalOcean. [He is a big SingleStore proponent](#)," said Perhac. "I couldn't believe how much more data Fathom is processing compared to SegMetrics while maintaining great performance.

Compared to the Amazon solutions, SingleStore had the best pricing model for what SegMetrics needed, as well as top-notch support, and would enable the ease of migration so it didn't have to change its entire language. "On paper, SingleStore looked like the best and easiest solution to switch to, and in practice it was even better than it would have been going to something like Amazon Aurora," said Perhac. As for PlanetScale, its pricing structure eliminated it immediately from consideration, as it charges per row. "I couldn't even price out PlanetScale, but back of the napkin calculations made it seem incredibly expensive for our use case," he added.

SingleStoreDB offers the fastest real-time distributed SQL database designed to power modern applications and real-time analytics. It scales real-time applications and analytics to any amount of data and users while maintaining millisecond (ms) response times. SingleStoreDB is built on a distributed scale-out architecture that scales linearly based on customer business needs, and it is built from the ground up to drive complex analytics on operational data as it lands. SegMetrics chose SingleStore for those reasons and more:

- Real-time analytics: SingleStore reduces time-to-insights with a database optimized for fast ingest and low-latency query performance with high concurrency.
- The ability to eliminate database sprawl: This modern, cloud-native, relational database delivers unlimited storage and scalability in a system that is durable, available, secure and resilient to failure, across cloud providers and on-premises. SingleStore is multi-model, supporting all major data types and unifying transactional and analytical workloads.
- MySQL compliance: The SegMetrics team wouldn't have to learn a new query language to work with SingleStore.
- Cost-efficiency: SingleStore's cost was significantly lower than PlanetScale for comparable performance levels.
- Its partnership with [Arcion Cloud](#), as discussed in part [here](#).



Solution

SegMetrics chose [SingleStoreDB Cloud](#) as the foundational database for its platform. Next, the team needed a way to migrate from MySQL on DigitalOcean to SingleStore on AWS while minimizing downtime so clients would experience only minimal disruption, allowing SegMetrics to adhere to committed SLAs.

The initial migration plan involved a one-time manual migration of a snapshot of data at a specific point in time, followed by a final migration of the latest data. Due to the millions of PSERT operations per hour and 50GB of daily upserted data, the downtime would have spanned more than 32 hours. SegMetrics connected with SingleStore and its migration partner, Arcion, for a different approach that resulted in a migration rate of 110 million rows in less than four minutes. The teams went through a three-week process to troubleshoot the migration strategy and fine-tune the configuration to get the data import delay under 10 minutes.

SegMetrics could switch between SingleStore and DigitalOcean for its data through a single line of code in the application's configuration file. The team used that capability to work through a feature-flag release. They were able to scale the SingleStore data queries while ensuring they could quickly address any issues during the migration without impacting the user experience.

Ultimately it took just 12 hours to migrate the TB SegMetrics database from MySQL on DigitalOcean to SingleStore on AWS. Real-time Change Data Capture from Arcion enabled SegMetrics to achieve ongoing synchronization between the old system and SingleStore, thus ensuring the data was up-to-date when SegMetrics cut over to the new system.

“I couldn't believe how much more data Fathom is processing compared to SegMetrics while maintaining great performance”

— Keith Perhac, Founder and CEO, Segmetrics

The system pulls in and normalizes time series data, then parses it into SQL files and runs against the SingleStore column store, tuned to accelerate query performance by aggregating on one or two columns. SegMetrics clients access data via a multi-tenant architecture. “Our clients can make highly specific queries,” said Perhac, “such as: ‘I want to see everyone who clicked this ad four days ago, who are named Bob, did not make a purchase, have a subscription to media and they like cats.’ A query such as this may join eight tables and runs filters on six of them, then aggregates the data down to the final results. Our queries are complicated, which is a blessing and a curse. SingleStore has made it so that these queries are processed in less than 50ms, rather than the 30 seconds it took previously.”

SegMetrics is using SingleStore Platform features, including:

- [SingleStore Pipelines](#), which support and streamline SegMetrics' data ingestion requirements by continuously loading data as it arrives from external sources. As a built-in component of SingleStore, Pipelines can extract, shape, modify, and load external data without third-party tools or middleware. Pipelines pull data from Amazon S3, and SegMetrics also uses raw SQL files for inserts.
- [SingleStore Workspaces](#), which enable customers to run multiple workloads on isolated compute deployments while providing ultra-low-latency access to shared data. This ensures applications are always operating on the latest data.

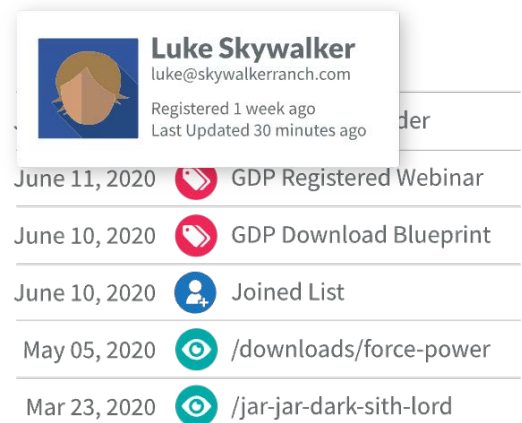


IMAGE | SegMetrics lets its customers see each contact, purchase, or event that drives its reports



Outcomes

By migrating from MySQL on DigitalOcean to SingleStore on AWS, SegMetrics expanded its real-time analytics capabilities, scaled its data capacity, and is growing its customer base and revenue.

Ability to Handle Twice as Much Data Enables Insights and Growth

"On the old system, we used to filter what data we brought in and couldn't bring in too much data. We had a large enterprise account sign up as a client, bringing a million web views a day. We had to add smart filtering to our processing to ensure it wouldn't make our database explode. With SingleStore, I don't have that concern anymore," said Perhac.

"One of the things I'm really happy about with SingleStore is that compute and storage are separate, which means that we can ingest 100TB of data and only pay the compute on that volume of data"

— Keith Perhac, Founder and CEO, Segmetrics

SegMetrics has doubled the data it works with, ingesting one million rows per second. "We're growing, revenue is growing, and new customers are coming in," said Perhac. "I'm very glad we moved over. We have a lot more confidence now to tell our clients we're bringing in a lot more data now, without sacrificing reporting speed."

Reducing the Complexity and Cost of Real-Time Analytics

"SingleStore simplified a lot of our analytical queries because it has special features that are specific for time series data," said Perhac. With the previous system, SegMetrics had to do multiple subqueries and joins on the table to get a visitor and potential buyer's first click. Now, the system can get right to the first click and enable clients to take action. "I'm also really happy that SingleStore has a separation of storage and compute," he added. "I can throw in 100TB of data because I'm essentially paying Amazon S3 prices on storage, and the compute cost only applies to the data being worked with."

Supporting Customers by Importing and Parsing One Billion Tag Applications Daily

Each SegMetrics client may have 100 to 200 tags. This number is multiplied by the thousands or millions of people in the CRM. This adds up to over one billion applied tags imported and parsed daily. "SingleStore has helped us save a lot of developer time because it never goes down, and it has reduced our opportunity costs. We can now support clients with millions of web views per day with no concerns," said Perhac.

SegMetrics uses the analytics in SingleStore to parse that data to provide actionable insights for its clients. "At any point during that customer journey, we can look at one of those touch points and say, 'People who take this action are worth \$500. People who don't take that action are worth \$200. So doing this increases customer value by \$300,'" he explained.

"SingleStore simplified a lot of our analytical queries because it has special features that are specific for time series data."

— Keith Perhac, Founder and CEO, Segmetrics

SegMetrics' own internal reporting has also benefited from adopting SingleStore. The team has a report on the number of clicks they've seen across all clients over a time period such as 30 days. Before with MySQL, the page timed out after five minutes. With SingleStore, it loads 15 million clicks per day in just 12 seconds.

Achieving Ultra-fast Time-to-Insight Through 200ms Response on 5,000 Complex Queries Per Second

SegMetrics real-time query engines are highly complex, allowing marketers to query on any of their data. At any given moment, an average of 50 customers are actively running new queries. While that might sound like low concurrency, each customer runs between 10 and 100 queries at a time - which means processing around 500 to 5,000 new queries per minute. "We're seeing around 200ms — less than a second for any query, and can run a hundred of them simultaneously. We can query data a lot faster than we used to," said Perhac.



Outcomes (con't)

Achieving Ultra-fast Time-to-Insight Through 200ms Response on 5,000 Complex Queries Per Second

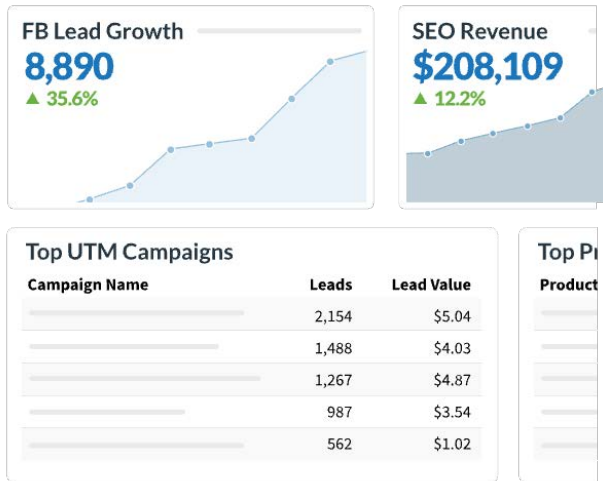


IMAGE | SegMetrics matches account info and ensures that each purchase is counted once and only once

SegMetrics has an alert that goes off if more than 200 queries are on the queue. Back in its MySQL days, it would have to delete those queries in the queue because the system couldn't catch up. "With SingleStore, even if we do get a backlog because we have so many people running queries at once, the queue clears in a second or two," said Perhac.



SingleStore is helping companies compete and win across every vertical

[Learn More >](#)

Eyeing a Future with AI Copilots, More Data, and Further Optimization

The SegMetrics team has its eyes on the future of its platform. One use case it wants to explore is using an AI copilot to parse the analytics results into information that people can understand better than numbers. "Even though we're an analytics tool, a lot of our users are not data people. We can get the results in English in a way that lets them understand and act on these insights better. It also explores what we can bring out of the data without specifically having to ask for it," said Perhac.

"We're seeing around 200ms — less than a second for any query. We can query data a lot faster than we used to."

— Keith Perhac, Founder and CEO, Segmetrics

SegMetrics also wants to ingest more data, since the platform running on SingleStoreDB Cloud is now able to reliably ingest and parse it. The team also wants to continue optimizing its SingleStore database through SingleStore-specific analytic queries to get better information and using features such as [SingleStore's Data API](#) combined with web performance and security leader [Cloudflare](#).

"We can ask, 'What's the first session someone had? What was the second page of their session, on the third Tuesday of the month?' We want to build systems on top of SingleStore to get those specific answers in a systematic way," he concluded.

Learn how SegMetrics successfully migrated from MySQL on DigitalOcean to SingleStore on AWS using Arcion Cloud, with zero downtime: [Avoiding SLA Breaches with Zero Downtime Database Migration using Arcion](#)