

How Mesosphere Measures and Optimizes its Support Teams

Mesosphere is an enterprise software company that builds the Datacenter Operating System (DCOS). The software behind the Mesosphere DCOS (including Apache Mesos) runs some of the largest data centers in the world, along with all of Apple, Twitter, and big portions of Verizon and others.

In this Q&A, Graham covers:

- The top analyses support teams should be running
- How to prove your support team is contributing to, not just eating up revenue
- Tips for making sure product and engineering teams listen to your customers
- How Mesosphere is using Segment Sources to pull together Zendesk, Salesforce, and Intercom data in Amazon Redshift to uncover these insights



“With Segment Sources, we’re calculating cost per ticket and prioritizing feedback for the product team.”

Graham Murphy

Head of Customer Operations

 MESOSPHERE



by **Diana Smith**
Product Marketing
Segment

DIANA: Thanks for taking the time with us today, Graham. So, tell us about Mesosphere and your role.

GRAHAM: Sure! I'm the head of customer operations for Mesosphere. In short, we're building an operating system for datacenters, the DCOS. Mesosphere organizes your entire datacenter and cloud infrastructure as if it were a single computer. We help businesses that run distributed systems build, deploy, and manage them much more efficiently.

My role here is to make sure we are providing exceptional service to our customers. Basically everything customer-facing falls into my world: support, escalating tickets to engineering, prioritizing product feedback, and resolving customer problems

DIANA: So what does the support team look like at Mesosphere? How important is it to the business?

GRAHAM: Well, we are an enterprise software company, and we have very large customers who rely on DCOS to power their datacenters. We need to be ultra-responsive, answer customer questions, and squash issues as quickly as possible. Each contract has strict support SLAs (service-level agreements for response time, etc), so closing tickets quickly and efficiently is critical.

As far as our team goes, we're made up of highly skilled engineers. You can't support Mesosphere if you're not an engineer, so.... our team is not cheap. We need to make sure we're maximizing customer happiness while also being fiscally responsible.

DIANA: Seems like your team is under a lot of pressure to perform quickly. How are you measuring performance and working to improve it?

GRAHAM: Right now, I'm focused on a few things to make sure our team is really contributing to the overall organization:

- Better understanding the scope and costs of support
- Tying support interactions to potential and existing revenue

- Prioritizing customer feedback for our product and engineering teams

These deliverables made me really excited about Segment Sources because they each require support and sales data in one place. We're already fairly heavy users of Segment's other products, but Sources is extremely valuable to me because I used to have to tie together my data from Zendesk, Intercom, and Salesforce by hand.

DIANA: What do you mean "by hand"?

GRAHAM: In the past I had to pull data from the tools' API's with generally hacky scripts that would push the data into Redshift or another SQL database. I could build the initial script myself or put a few engineers on it for a few weeks, but they are already at capacity handling support. And, thinking about other companies, I doubt most support teams would be technical enough to do this themselves.

And then, even when you get a hacky cron job put together (or chronos job, in our case), you still have the headache of really maintaining it. For example, what happens when APIs and schemas change? Simple things like adding an additional custom field to your Zendesk script become a complete pain when you consider the fact that they will also require adding another column to your Zendesk table in Redshift.

We learned the hard way that there are inevitably going to be 10 to 20 different idiosyncrasies between various systems, data structures, and APIs that we're not going to have thought of the first time around. This means we are going to have to go and fix them and, frankly, that is not our specialty. That's your focus. I'd rather have you guys do it for us.

DIANA: Glad to hear it! So, what Segment Sources are you pulling data from right now?

GRAHAM: Zendesk is a big one — it's the support channel that is under SLA for us. That's usually where meatier customer questions come in. We also offer support through Intercom in certain parts of our product, so we need that too. We're syncing Slack data from our customer channel over now with

“We learned the hard way that there won't be 10-20 idiosyncrasies of APIs that we won't find the first time around building our own ETL pipelines.”

the `.set()` API for custom object loading, but I'd be more than happy to take a few Segment engineers out for happy hour if I could fast track that source.

To tie our support data to cash, we also use Salesforce, which is our source of truth for revenue. Having worked at Salesforce for a while, I'm a big advocate of the software. Easy access to this data — accounts and opportunities, specifically — in Redshift is very useful for me because it allows me to slice and dice the revenue data and tie it directly to our support and customer data.

At one point, we considered doing all of this analysis in Salesforce by pulling more support information in there but, honestly, I find it lacking a lot of the integrations I'd want. With Segment Sources, I can supplement in our data in Salesforce with queries in Redshift across support and product data, which means it's much more flexible.

Unifying the data sets through Segment makes my performance analysis easier. Now that I can run a single query against a few Redshift tables, I've saved dozens of hours of customization work. It's the unification of the data sources that's really powerful for me.

DIANA: What types of analyses are you running now that you have Segment Sources? How are you tying your team's contributions to revenue? What's your main goal?

GRAHAM: That is a multifaceted question. The short answer is that the ultimate goal of customer support in any organization is to prove you're providing value. After all, it's very easy to view customer support as a cost center. The good news is that you can demonstrate customer support is not a cost center, but actually protects and grows revenue, if you have the right data.

Here's an example: Say you had a hundred customers that encountered a bug. Fifty came and contacted support, and 50 didn't. With your product usage data, your support data, and your Salesforce data, you can look at

“Unifying Zendesk, Salesforce, Intercom and in-app data sets through Segment saved me dozens of hours of customization work. Now, I just run a query against a few RedShift tables.”

the differences between the two groups and speculate as to whether the support interaction influenced whether or not they stayed on as a customer. Note that this is a simplification of a principle that John Goodman covers in detail in “Strategic Customer Service.”

DIANA: Nice! What else are you doing with your support data?

GRAHAM: The second way we’re using the data is to help prioritize fixes for our product and engineering teams and make sure our customers’ voices are heard. When we say X number of customers, equaling Y amount of revenue are experiencing this bug, that’s a lot more compelling when we ask for a fix. And, I can also tie in things like actual impact on revenue, potentially impacted revenue, and future revenue. This makes the support team’s product requests hold a lot more weight.

Besides feature feedback, we’re also working on building individual timelines for each customer across all channels. With my Zendesk, Salesforce, and Intercom data together in one database, it’s really easy to figure out exactly when our last contact with a certain organization happened. I want to do that in a single query, and use it for all sorts of stuff. It can be a dashboard. It can be a report for the exec team. And, it can be a tool for support and sales agents reaching out to high-value customers. There are a lot of possibilities.

DIANA: That’s such a great idea! We might steal that one. I want to bring back cost per ticket for a second, which you mentioned earlier. What are you doing with that?

GRAHAM: What I’m really trying to do is look across all channels (SLA-governed or not) and all customers (community customers that aren’t paying us anything, plus customers that are paying us a ton) and get a holistic view of what my team is doing. I want to measure every output.

Once I get all of the top level data (replies, notes, tickets), it’s easy to run some pretty critical projects:

- Calculating cost per ticket
- Extrapolating cost to support specific and future customers

“We’re using Segment Sources to identify the root cause of email bounce behavior from SendGrid.”

- Building a hiring model

The way that you generally calculate the cost of a ticket is to take the target earnings of all the people on your team. You then add your overhead costs associated with the team, tooling or anything else, and the variable costs. This is your “all in” costs. You then compare that to the number of tickets that were created in that time period. That would be industry standard way to calculate cost per ticket.

With this baseline you can get into more creative metrics, like support cost over customer buckets—such as paying customers versus community customers. And, what I look at now is how much it costs for each interaction, because maybe one ticket had 20 interactions.

DIANA: Those seem like very useful analyses!

GRAHAM: Yes, especially because the support we provide is very costly. Right now our enterprise ticket costs are pretty high. Therefore, it’s important for us to actually tie the cost per ticket numbers back to the contracts on specific deals. I’m making up numbers, but let’s say one customer pays us \$100 a month, but it costs \$500 dollars to support them. Obviously, that’s not sustainable.

So in terms of comparing the cost of goods sold to internal operational costs, it’s very important for us to be able to have that holistic view of every ticket and conversation. And again, for me, it’s really important because if I were to look at all the Sources, I could make a generalization that an Intercom reply costs \$100 dollars, while a support ticket costs us \$2,000. Then applying that to requests per specific customer, I have a much more realistic view of the impact support is having on our organization. I also know how many agents we need to make sure our customers have a great experience.

These types of analyses would really annoying to do without all of my data in Redshift. Segment Sources makes it easy.

“These types of analyses would really annoying to do without all of my data in Redshift. Segment Sources makes it easy.”