

Tecton: A Rapidly Scaling Startup Outgrows Prometheus

Chronosphere provides a scalable, reliable and customizable monitoring-as-a-service purpose built for cloud-native applications



The Challenge

Tecton is a startup that provides an enterprise-grade feature store for machine learning applications. When the product first launched, the team relied on an out-of-the-box open-source Prometheus setup to monitor the fleet of Tecton instances, some of which are deployed directly in the customer account.

This approach worked at first, when Tecton had relatively few customers, meaning the number of Tecton instances was small. But as the company grew, the monitoring system started breaking down.

“We’re not in the business of building monitoring stacks,” explained Ravi Trivedi, software engineer at Tecton. “We have the engineering know-how to do it, but we don’t want to be wasting our time with monitoring.”

But even at a relatively small, startup scale the monitoring system was a source of frustration.



Each instance was completely siloed

With the Prometheus setup, it was impossible to get a global view of the entire fleet of Tecton instances. This was the underlying cause of a number of frustrations.

First of all, during incidents the on-call engineer would have to toggle between isolated Prometheus deployments, increasing the amount of time it took them to uncover the metrics needed to address the issue. In addition, any change to Prometheus had to be made to each deployment. When the Tecton team wanted to silence a noisy alert, they would have to go through every Prometheus instance and manually make the change.



There was no long-term storage

Monitoring data was stored on persistent volumes, which generally stored data for about a week. Metrics would constantly be dropping out of memory, and if a customer brought up a minor incident that happened 8 days prior, there was little Tecton could do to investigate it. Just as importantly, the persistent volumes weren't reliable, and Trivedi would often have to restart the volume, wiping away any historical monitoring data.



Looking for an alternative

The Tecton team had three criteria for a new monitoring solution.

They were:

- A global view across regions and customers
- Control over data retention and the ability to keep historical data
- High availability



The system kept breaking

Among the biggest problems with the open source Prometheus setup was that it was buggy. When it inevitably broke, Trivedi had to drop everything and focus on fixing the problem. “It’s not acceptable to have your monitoring broken,” he said. “If something happens, it preempts everything else. I could easily burn the better part of a day figuring out monitoring.”

These frustrations with the monitoring system were a slow burn, but when the team started building a custom way to get multi-region support for Prometheus, they knew that their approach had to change. “We said, this is not the right way to do it,” Trivedi said. “We’re basically building what we know is technical debt.” Not only was the team investing significant engineering time building custom features to make Prometheus meet everyone’s needs, they were also spending a non-trivial amount of time firefighting monitoring issues.

“I actually worked on Google’s monitoring stack,” Trivedi said.

“I knew that we don’t have the human resources at Tecton to build a good monitoring stack.”

Success with Chronosphere

The biggest benefit of moving to Chronosphere is having an out-of-the-box monitoring solution that doesn't break all the time and has a full feature set that doesn't need custom engineering work.

"We basically don't think about monitoring anymore as we spin Tecton deployments up and down," Trivedi said. The days of spending hours urgently debugging the monitoring system are gone, leaving Trivedi and the rest of the team with more time to focus on Tecton's core product.



Less burdensome on-call rotations

Using Chronosphere has made on-call rotations less of a burden in a couple ways. First of all, there is a single point of entry into the monitoring system, which reduces the amount of time it takes to get visibility during an incident and figure out what is happening. Second, the ability to apply silences across the entire fleet of instances has made it easier to tune the alerting system and reduce the amount of noisy alerts, generally reducing alert fatigue.



Useful data retention

Now that Tecton has more control over data retention — as well as the ability to mix data from multiple places — Tecton has been able to create custom internal dashboards that pull data from the Chronosphere API and track business intelligence metrics. That wouldn't be possible to do with Prometheus.



"It's way better than it was before."

Ravi Trivedi, Software Engineer

Engineers are happier, better equipped to handle their on-call rotations and don't have to waste time building custom monitoring features just to get the visibility they need.