

Incident Dec. 16, 2021:

Products search may not work as expected

Availability: **Some searches and product pages were down**

% of stores affected: **3.7%**

Duration of incident: **65 minutes, plus 40 minutes of the inactive indexing process**

Symptoms

From 17h20 to 18h25 UTC, customers who tried to use VTEX stores might have had issues searching for products. For these clients, products that were not cached in our systems wouldn't be displayed on searches, which led to missing components on pages, resulting in a "Something Went Wrong" error message in some cases. Additionally, VTEX clients were unable to notice their updates on products for the entire incident duration.

Summary

At 17h20 UTC, we identified an increase in the latency of one of our catalog search engine services. At 17h51 UTC, we discovered that one of our search catalogs had elevated errors for a while. We then started the failover cluster and redirected traffic to it. Finally, we reactivated the indexing process after the degraded cluster was fully recovered.

Timeline

[2021-12-16 17:20] Our monitoring systems notified us of high latency on one of the catalog search engine services.

[2021-12-16 17:25] We stopped the indexing process to avoid higher loads.

[2021-12-16 17:51] We identified elevated errors in this catalog.

[2021-12-16 17:55] We provisioned the failover cluster so it could handle the load.

[2021-12-16 18:09] We started migrating traffic from the affected cluster to the failover one.

[2021-12-16 18:19] We finished traffic migration to the failover cluster. The search was now fully operational, but the indexing process was still turned off.

[2021-12-16 18:54] The catalog search engine affected cluster was fixed, and the indexing process was reactivated.

[2021-12-16 19:23] Normal operations were fully reestablished.

Mitigation strategy

We started our failover cluster for the affected index. Then, after the failover cluster was ready, we gradually migrated the traffic from the affected cluster to the new one. That was the solution to make the search operational again. Lastly, we had to recover the affected cluster before enabling indexing again.

Follow-up actions: preventing future failures

As follow-up actions to this incident, we will work towards improving our monitoring alerts to speed up the identification process for this module. Additionally, we will improve catalog search engine configurations to reduce the pressure on the cluster.