

Incident Nov. 19, 2021: Catalog Search Engine Partial Outage

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

% of stores affected: **7%**

Duration of incident: **35 minutes, plus 134 minutes of inactive indexing process**

Symptoms

From 20h41 to 21h16 UTC, customers who tried to use VTEX stores might have had issues searching for products. For these clients, products not cached in our cache systems wouldn't be displayed on searches, which may have also led to missing components on pages. Additionally, VTEX clients were unable to update their products for the entire duration of the incident.

Summary

At 20h41 UTC, we identified a massive increase in the load of one of our catalog search engine services. We discovered that one of our catalog search indexes stopped working as expected, and started returning these errors. We started the failover clusters and redirected traffic to it. Finally, after the affected cluster was fully recovered, we reactivated the indexing process.

Timeline

[2021-11-19 20:42] Our monitoring systems notified us of a high load on one of the catalog search engine services.

[2021-11-19 20:46] We stop the indexing process to avoid higher loads.

[2021-11-19 20:50] We provision the failover cluster so it could handle the load.

[2021-11-19 21:01] We start migrating traffic from the affected cluster to the failover one.

[2021-11-19 21:16] We finish traffic migration to the failover cluster. Search is fully operational but the indexing process is still turned off.

[2021-11-19 21:50] The catalog search engine affected cluster is healed and the indexing process is reactivated.

[2021-11-19 23:30 GMT-3] Normal operations were fully reestablished.

Mitigation strategy

We started our failover clusters related to the affected index. After the failover clusters were ready, we gradually migrated traffic from the affected clusters to the failover ones. That was enough to make the search fully operational again. Then, in order to enable indexing again, we recovered the affected cluster.

Follow-up actions: preventing future failures

As follow-up actions to this incident, we will work towards improving our catalog search engine configurations to reduce the pressure on the cluster, and additionally, increase observability overall.