

Incident May 06, 2022: Promotions' database increased errors

Availability: **100% of orders were not finished during the incident**

% of clients affected: **100% of clients affected during the incident**

Duration of incident: **31 minutes**

Symptoms

From 19h51 to 20h22 UTC, customers that tried to use VTEX stores were not able to make purchases.

Summary

At 19h51 UTC, our automated monitoring system identified a quick and substantial increase in processing in one of the databases used by the Promotions system. We discovered that there was a query responsible for it. We proceeded to roll back changes made prior to that day, working to mitigate the problem and prevent it from happening again in the future.

Timeline

[2022-05-06 19:51 UTC] Our monitoring systems warned of anomalous behavior in the Promotions system.

[2022-05-06 19:57 UTC] We identified anomalous spikes on Promotion's database connections and processing.

[2022-05-06 20:05 UTC] We identified the main contributing factor of our Promotions service layer failures and started mitigation actions.

[2022-05-06 20:22 UTC] Normal operations were fully reestablished.

Mitigation strategy

We scaled the servers from the old version prior to the day and moved the Promotions traffic to it. After that, we had to guarantee the queries impacting our database negatively were not being overused. Gradually, our services started returning fewer navigation and transaction errors. For some services, manual intervention was necessary to accelerate their recovery.

Follow-up actions: preventing future failures

As a follow-up to this incident, we will work on improving our internal architecture to mitigate errors that affect the shopper experience in VTEX stores. We will also create tools to allow us to block API routes and queries that are degrading our databases. Besides, we will create a failover that can still work in case of degradation of our databases.