

Incident Jul. 01, 2022: Degraded cluster on VTEX IO impacting sessions

Availability: **All accounts were available but probably faced timeouts for part of their requests**

% of clients affected: **Still being investigated**

Duration of incident: **1h43**

Symptoms

From 18h00 to 19h43 UTC, most customers that tried to use VTEX stores might face some timeouts while navigating the stores.

Summary

At 18h00 UTC, we identified an increase in the overall latency on our platform, leading to timeouts. We discovered that one of the authentication/authorization system modules had an increase in latency due to a problem with the infrastructure it was set in. We started mitigating the problem by migrating traffic to another cluster in our infrastructure.

To be the trusted partner to your success, our team is working on follow-up actions to ensure that this incident does not happen again and that we identify and recover from future incidents faster. We are committed to improving our systems to guarantee a reliable and trusted experience.

Timeline

[2022-07-01 3:00 PM UTC -3] Errors concentrated in one of the VTEX IO clusters triggered alarms.

[2022-07-01 3:01 PM UTC -3] We started investigating the errors and identified that these were concentrated in one of our clusters.

[2022-07-01 3:05 PM UTC -3] Mitigation strategies were taken by moving part of the traffic to another cluster while we investigated the root cause of the problem.

[2022-07-01 3:06 PM UTC -3] We identified an overall latency increase on the platform.

[2022-07-01 3:13 PM UTC -3] We also identified that one of our platform modules was facing timeouts from another system deployed on VTEX IO.

[2022-07-01 3:20 PM UTC -3] We began provisioning infrastructure to replace the infrastructure that faced a degraded performance.

[2022-07-01 4:12 PM UTC -3] We started migrating traffic to the new infrastructure.

[2022-07-01 4:43 PM UTC -3] Normal operations were fully reestablished.

Mitigation strategy

We have provisioned a new cluster on VTEX's IO infrastructure to move traffic out of the degraded infrastructure.

Follow-up actions: preventing future failures

As a follow-up to this incident, we will work on improving our alarms to detect this kind of cluster degradation faster. We'll also work on enhancing the tools we have to provide infrastructure for a faster traffic migration.