

# Incident Mar. 17, 2022: Degraded messaging on VTEX IO

Availability: **Degraded developer experience**

% of clients affected: **100% of VTEX IO clients**

Duration of incident: **2 hours and 45 minutes**

## Symptoms

From 17h10 to 19h55 UTC, store operators that tried to link VTEX IO apps or had any other processes that relied on the VTEX IO messaging system experienced intermittent higher latencies and increased 5xx errors.

Some operations that rely on asynchronous messaging may also have been affected.

## Summary

At 17h10 UTC, we saw an increase in timeouts and errors in the app build system and toolbelt commands. Upon further investigation, we found out that a misbehaving app was using most of the resources of our messaging engine. Once we disabled the app, the resource usage and all systems went back to normal.

We also identified another, more prolonged incident related to the app build system in VTEX IO, which also degraded the development experience. We are working on a report about that as well.

## Timeline

**[2022-03-17 17:10 UTC]** An app started to consume most of the resources available for the VTEX IO's messaging system.

**[2022-03-17 17:15 UTC]** Our monitoring systems showed an increase of 5xx errors in the development cluster.

**[2022-03-17 17:25 UTC]** The team tried to mitigate the problem by swapping dev clusters. This strategy had no effect.

**[2022-03-17 17:28 UTC]** Our monitoring systems reported abnormal behavior in the pub-sub system in admin and dev clusters. We also received reports of event server errors in VTEX CLI. After that, we focused the investigation on the messaging system.

**[2022-03-17 19:39 UTC]** We identified that a VTEX IO app has been consuming most of the available resources for the messaging system.

**[2022-03-17 19:48 UTC]** We disabled the app on the platform. Errors gradually started returning to normal levels.

**[2022-03-17 19:55 UTC]** Normal operations were fully reestablished.

## Mitigation strategy

We first tried migrating traffic from one cluster to another. When that didn't give the expected results, we continued searching for the root cause. Upon discovering that a certain app was responsible for degrading our services, we disabled it, making errors go back to normal levels.

## Follow-up actions: preventing future failures

As a follow-up to this incident, we will work on improving our visibility of shared resources throughout VTEX IO so that we can detect the degradation of all shared systems more quickly. We'll also review our policies of shared resources so that no single misbehaving app can degrade the experience for everyone.

More follow-up actions will be specified in the report on other development instabilities throughout the week.