

# Incident Mar. 8, 2022: Elevated Errors on deployments of new versions with VTEX IO services

Availability: **The development environment from VTEX IO was impacted up to 25%. No sales were degraded.**

% of clients affected: **Developers using *vtex-cli* to test and deploy apps improvements experienced issues during the incident.**

Duration of incident: **3 hours and 46 minutes**

## Symptoms

From 14h56 to 18h42 UTC, developers were experiencing errors when linking apps within their workspaces. There were also reported issues communicating with the storage service in the development environment.

## Summary

At 14h56 UTC, we were notified about errors in our development environment. Our engineering team investigated immediately and determined there was disk pressure causing issues on services related to linking, building, and deploying applications. The storage service was also affected.

We tried several mitigation strategies between 14h56 and 18h42 UTC, like removing instances with disk pressure. Meanwhile, we were provisioning new infrastructure for our development environment, and we ultimately moved all traffic to the newly equipped infrastructure.

At 18h42 UTC, the disk pressure was sorted out, and the environment was fully operational.

## Timeline

**[2022-03-08 14:56 UTC]** Engineering team was notified of random errors in the development environment.

**[2022-03-08 15:10 UTC]** Our on-call engineering began the investigation.

**[2022-03-08 15:14 UTC]** We replaced an instance which was experiencing issues and causing abnormal behavior for all services on that instance.

**[2022-03-08 15:37 UTC]** We identified additional issues with services being affected by disk pressure on additional instances.

**[2022-03-08 17:04 UTC]** We began resizing root volumes on affected instances.

**[2022-03-08 17:13 UTC]** Volume increases had no effect; existing instances did not see volume size changes.

**[2022-03-08 17:29 UTC]** We started spinning up replacement infrastructure.

**[2022-03-08 18:00 UTC]** We attempted to roll out new instances in the existing environment with updated volume sizes.

**[2022-03-08 18:10 UTC]** We successfully provisioned replacement infrastructure. In addition, we continued monitoring rollout of instances in the existing environment.

**[2022-03-08 18:20 UTC]** We started moving traffic to the replacement infrastructure.

**[2022-03-08 18:42 UTC]** Operations were completely reestablished after this maneuver.

## Mitigation Strategy

At 14h14 UTC, the engineering team replaced an instance that was no longer communicating properly with the environment. At 14h37 UTC, the engineering on-call team attempted to increase volume sizes for the affected environment using several methods, unsuccessfully. At 18h10 UTC, we started the process of replacing the infrastructure. At 18h42

UTC, all traffic was moved to the replacement infrastructure, and normal operations were restored.

## Follow-up actions: preventing future failures

As follow-up actions to this incident, we will update our infrastructure to prevent disks from filling up. We will also increase monitoring to promptly identify issues related to disk pressure.