

Incident Jul. 28, 2021: Orders Listing Partial Degradation

Availability: **Stores operational, but OMS Listing API affected**

% of clients affected: **20% of clients affected by the incident**

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

Symptoms

From 12h46 to 14h48 UTC, some customers did not receive updated data when listing orders using the OMS List API or the Admin UI. The reason is that the API was stalled for some time due to a database cluster failure. This issue did not affect our Feed and Hook integration.

Summary

At 12h46 UTC, we identified errors preventing some accounts from correctly indexing orders on our platform. We discovered that one of our clusters stopped working, and the accounts that were using it were affected.

We immediately started working on the recovery of that cluster and redirected the affected accounts to a healthy cluster. We were able to recover the failing cluster, and operations recovered correctly.

It's important to note that all accounts were receiving orders and that purchases were overall not affected.

Timeline

[2021-07-28 12:46 UTC] Our monitoring systems warned us of high latency in one of our clusters.

[2021-07-28 12:47 UTC] We started investigating the root cause of the problem and began recovery.

[2021-07-28 13:03 UTC] We identified that a single cluster and several stores were affected.

[2021-07-28 13:11 UTC] We started the scale-up of the unhealthy cluster.

[2021-07-28 13:39 UTC] We were unable to provision more instances and started the provisioning of the failover cluster.

[2021-07-28 14:00 UTC] We identified a problem with the instance launch of our clusters and started fixing it.

[2021-07-28 14:07 UTC] We fixed the cluster launch and began replacing the unhealthy instances.

[2021-07-28 14:10 UTC] We migrated the affected accounts to a healthy cluster.

[2021-07-28 14:20 UTC] Our system was able to restart the orders' indexing process.

[2021-07-28 14:48 UTC] Normal operations were fully reestablished.

Mitigation strategy

During the incident, we migrated affected accounts to a healthy backup cluster. This action fully restored the OMS List API operation while we worked on recovering the failing cluster.

Follow-up actions: preventing future failures

To be the trusted partner to your success, our team is working on doubling the efforts to do scheduled preventive maintenance on our clusters and implementing an anomaly detection system that would help us find out sooner if any degradation is in



progress. We are committed to improving our systems to guarantee a reliable and trusted experience.