Incident Aug. 01, 2022: Increased latency and errors in the order listing operations

Availability: **A percentage of our clients had intermittent issues while listing orders using the Admin UI or the Listing API in the OMS**

% of clients affected: **16%**

Duration of incident: **2 hours and 16 minutes**

Symptoms

From 13:57 UTC to 16:13 UTC, one of our order indexing clusters failed some queries due to high CPU usage. Some of the accounts that use this cluster failed some queries. This might have affected our clients using the Admin Order UI while searching for an order or clients with integrations using the Listing API in the OMS/Orders. This didn't impact integrations using Feed/Hook or purchases in general.

Summary

At 13:57 UTC on Aug 4, 2022, our alarms fired, notifying us of high CPU usage in one of our clusters. This was due to an increase in regular use. We started monitoring the situation immediately to check if the cluster would heal itself, but it didn't. At 14:20 UTC, we started scaling up our systems to accommodate the current traffic better. At 14:39 UTC, our

monitors showed us that the maneuver had good results. The CPU usage was progressively lowering, and our systems were returning to normal. We continued to monitor the situation up to 16:13 UTC when we declared the problem solved.

Our clients could still navigate their Orders Admin and make queries during this time, though some were unsuccessful. No impact was detected in purchases nor the order flow.

Timeline

[2022-08-02 13:57 UTC] Our alarms fired, notifying us about high CPU usage in one of our clusters.
[2022-08-02 14:20 UTC] Scaling up procedure.
[2022-08-02 14:39 UTC] First signs that the cluster was recovering.
[2022-08-02 14:39 UTC] The team monitored the whole situation until it returned to normal.
[2022-08-02 16:13 UTC] Incident closed.

Mitigation strategy

A scale-up procedure solved the problem.

Follow-up actions: preventing future failures

As a follow-up to this incident, we will work on splitting the data of our shards to better help with the CPU load where there is high demand.