

Incident Jul. 26, 2021: Delay on Intelligent Search Indexing Process

Availability: **Stores operational**

% of clients affected: **All customers using Intelligent Search**

Duration of incident: 20 hours and 22 minutes

Symptoms

From July 26 at 23h30 to July 27 at 19h52 UTC, customers that tried to index new products in Intelligent Search could experience slow indexing.

Summary

On July 27 at 12h34 UTC, we identified that products were taking too long to commit to the search index. We discovered that the system responsible for the commits was degraded and the backlog of unprocessed messages was increasing. At 13h58 UTC, our team mitigated the problem, and the backlog of unprocessed messages started to decrease. To speed the recovery process the team also created a new environment that was better scaled to deal with the backlog.

From the time the issue started happening until it's full resolution, customers navigating through accounts that used Intelligent Search may have seen outdated product information on search pages. Price and availability information are always real-time, so they are excluded from this scenario. During the incident, we've also used an emergency queue to accelerate the indexing process of more critical scenarios.

To be the trusted partner to your success, our team is working on follow-up actions to make sure 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

[2021-07-26 23:30 UTC] Our monitoring systems started notifying about the unprocessed messages backlog increase, but this did not trigger any alert.

[2021-07-27 12:34 UTC] The team noticed the backlog increase on our monitoring stack.

[2021-07-27 13:58 UTC] The indexing queue started committing processed products again.

[2021-07-27 17:00 UTC] We deployed the service to the new environment.

[2021-07-27 19:52 UTC] Normal operations were fully reestablished.

Mitigation strategy

We created and migrated the committing service to a new dedicated environment, better scaled to deal with the backlog of unprocessed messages that was generated during the incident.

Follow-up actions: preventing future failures

As follow-up to this incident, we will deploy a new architecture for some of our services that relies less on shared resources, making the monitoring process for those resources more efficient. We'll improve our alerts of unprocessed messages for message processing services, thus ensuring that we'll act faster on these scenarios avoiding an increasing backlog.