

Incident Oct. 12, 2022: Intelligent Search outage

Availability: **Product list pages partially unavailable**

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

Duration of incident: **40min for 80% accounts, 4h48min for the other 20%**

Symptoms

From 14h02 UTC to 18h50 UTC, the Intelligent Search service faced availability issues that resulted in no products being shown for some users on search and navigation pages.

Summary

At 13h20 UTC on Oct. 12, 2022, the team started a preventive maintenance on our search infrastructure. While this maintenance was taking place, at 14h02 UTC, our monitoring system started reporting failures on multiple environments. Promptly, the team began to work on disaster recovery actions for this incident.

While working on stabilizing the original infrastructure, the team also started to spin up a new fallback environment at 14h20 UTC, a maneuver we knew would take a long time to conclude but would guarantee availability if we were unable to stabilize the original environment.

At 14h47 UTC, the original infrastructure started showing signs of recovery, with all accounts now starting to see search results. The situation reverted to a degraded state a few minutes later. At this point, at 15h00 UTC, 80% of accounts had fully recovered to a normal state. The other half of the environment was only returning stale results, so we focused our mitigation strategies on those degraded accounts.

Once it became clear that part of the original environment was not responding to any of our disaster recovery strategies, we started to roll over accounts to a secondary service that had some of the Intelligent Search features turned off, but was able to return results on full-text search pages. The fallback cluster initiative was completed at 17h00 UTC, and we started rolling over the affected accounts to this new environment. At 18h50, the Intelligent Search service was stable for all accounts.

Timeline

[2022-10-12 13:20 UTC] We started a preventive maintenance on our search infrastructure.

[2022-10-12 14:02 UTC] Our alert system reported failure in multiple environments, and the team started working on disaster recovery procedures.

[2022-10-12 14:20 UTC] We started spinning up a fallback environment.

[2022-10-12 14:47 UTC] The original environment showed signs of recovery, and all accounts began seeing results in the searches.

[2022-10-12 14:50 UTC] 80% of the accounts were back to a normal state.

[2022-10-12 15:00 UTC] The original environment went back to an unstable state.

[2022-10-12 17:00 UTC] The fallback environment was stable, and we started migrating affected stores to it.

[2022-10-12 18:50 UTC] All the accounts were back to a normal, stable state.

Mitigation strategy

We returned Intelligent Search to normal operations by moving traffic out of the degraded infrastructure.

Follow-up actions: preventing future failures

As a follow-up to this incident, we will improve our preventive maintenance and disaster recovery strategies. We are reviewing the Intelligent Search infrastructure to allow for faster fallback environment rollout. We're also working on a new architecture where infrastructure issues become more localized, affecting fewer accounts at a time, and providing an even quicker mitigation strategy.