

# Incident Report

## High Error Rate in Stores

|                   |   |
|-------------------|---|
| Status Page URL   | <a href="https://status.vtex.com/incidents/01KM6G2MFWSDN97ANMBX10GGFM">https://status.vtex.com/incidents/01KM6G2MFWSDN97ANMBX10GGFM</a> |
| Impacted accounts | Stores using StoreFramework or Faststore  |
| Impact            | Partial storefront degradation caused latency, errors, and page access issues.  |
| Duration          | Sales flow impact: 3 hours and 13 minutes<br>Search Degradation: 8 hours and 43 minutes   |

## Summary

On Mar 20, 2026, from 19:56 UTC to 05:49 UTC, some VTEX customers experienced degraded storefront performance, including increased latency, intermittent errors, page loading failures, and occasional difficulty accessing specific pages.

This incident was triggered by a combination of infrastructure instability and an exceptionally high volume of search requests, which together overloaded our system responsible for processing searches and loading pages. We mitigated the impact by increasing system capacity and stabilizing the affected components, allowing requests to be processed normally again and restoring the experience.

Our sales flow was partially affected for approximately 3 hours and 13 minutes during this incident. After this period, some products for a subset of clients continued to experience intermittent loading issues due to residual instability in the search infrastructure.

All issues were subsequently mitigated, and full platform stability was restored. We sincerely apologize for any inconvenience this incident may have caused.

## Symptoms

A subset of accounts using Store Framework and FastStore experienced intermittent disruptions in the sales flow. Users encountered 504 errors and missing content while browsing the storefront. In Store Framework stores, search functionality was impacted, with products not being returned in search results. In FastStore stores, search results failed to load, and PDP (Product Detail Page) rendering issues led to "product not found" errors. The issue was caused by instability in the platform's request processing layer, resulting in inconsistent delivery of pages and resources.

## Timeline

Impact window: 8 hours and 41 minutes

|                             |   |
|-----------------------------|---|
| <b>2026-03-20 19:38 UTC</b> | Impact on sales began. Initial infrastructure degradation observed, including increased latency and errors in the request handling layer. |
| <b>2026-03-20 19:40 UTC</b> | Customer-facing impact became evident, with elevated error rates and partial failures in page and asset loading.                          |
| <b>2026-03-20 19:48 UTC</b> | Initial signals from the search layer indicated abnormal behavior, including increased error rates in downstream services.                |
| <b>2026-03-20 20:02 UTC</b> | Search infrastructure degradation became more visible, with increased latency and processing delays affecting search-related operations.  |
| <b>2026-03-20 20:07 UTC</b> | Internal alerts confirmed elevated error rates  |
| <b>2026-03-20 21:20 UTC</b> | Engineering teams engaged AWS support to investigate infrastructure behavior and elevated error rates.                                    |
| <b>2026-03-20 21:34 UTC</b> | Teams analyzed errors and shared evidence to support the investigation.   |


|                             |   |
|-----------------------------|---|
| <b>2026-03-20 21:40 UTC</b> | Mitigation attempts focused on isolating potential causes within the search and indexing layers.                                  |
| <b>2026-03-20 22:13 UTC</b> | The team confirmed that increased latency in the request handling layer was a key driver of the incident.                         |
| <b>2026-03-20 22:30 UTC</b> | Initial mitigation actions were applied by recycling affected cache components to restore system stability.                       |
| <b>2026-03-20 22:36 UTC</b> | Platform metrics showed early signs of recovery.  |
| <b>2026-03-20 22:51 UTC</b> | Sales impact was mitigated, with platform performance returning to expected levels for core transaction flows.                    |
| <b>2026-03-20 23:00 UTC</b> | Although the main sales flow recovered, intermittent errors continued due to instability in the search infrastructure.            |
| <b>2026-03-21 00:15 UTC</b> | Targeted mitigation actions were applied, including workload redistribution across search clusters, but instability persisted.    |
| <b>2026-03-21 00:40 UTC</b> | The team initiated rollback procedures for recent changes in the search service.  |
| <b>2026-03-21 01:50 UTC</b> | Analysis of search system logs identified an unusually high volume of heavy queries, significantly impacting cluster performance. |
| <b>2026-03-21 02:23 UTC</b> | The decision was made to isolate a workload from the infrastructure.  |
| <b>2026-03-21 03:10 UTC</b> | Additional capacity was added to the search infrastructure to stabilize performance under load.                                   |
| <b>2026-03-21 04:31 UTC</b> | The incident was fully mitigated.   |

## Mitigation strategy

We reestablished normal operations of the platform by investigating infrastructure, identifying and recycling unhealthy cache instances, and stabilizing request processing and traffic levels. In parallel, we isolated and mitigated search-related pressure by temporarily stopping indexing, redistributing workload across environments, rolling back recent changes, and scaling infrastructure capacity.

## Follow-up actions

We are committed to being your trusted partner for success, and uphold the stability you expect from our platform. To prevent incident recurrence, we are taking these actions:

-  **(DONE) Improved infrastructure observability:** Implementing new alerts and dashboards to detect unhealthy nodes, memory exhaustion, and latency degradation earlier.
- **Stronger protection against traffic spikes:** Improved our ability to detect and isolate high-load scenarios, reducing the risk of cascading impact across shared environments.
- **Enhanced search resilience and performance:** Reviewing timeouts, query behavior, and cluster configuration to better handle load and avoid latency spikes.
- **Safer platform operations:** Strengthening processes for workload migration, rollback, and infrastructure scaling to ensure more predictable recovery during incidents.
- **End-to-end storefront visibility:** Expanding monitoring across all storefront layers to provide clearer visibility into customer experience and quickly identify where degradation occurs.