

Incident Report

Mar 12, 2025 : Increased errors in VTEX IO

Status Page URL	https://status.vtex.com/incidents/01JP4W4E56AAQPMBPD5SZFM41Q
Impacted accounts	All accounts
Impact	Most merchants were unable to login to their Admin Some shoppers experienced errors navigating or placing orders
Duration	Symptom 1: 2 hours and 50 minutes (from 09:13 to 12:03 UTC) Symptom 2: 2 hours and 5 minutes (from 09:49 to 11:54 UTC) Symptom 3: 2 hours and 37 minutes (from 09:13 to 11:50 UTC)

Summary

On Mar 12, 2025, from 09:13 to 12:03 UTC, most merchants were unable to login to their Admin and some shoppers experienced errors navigating or placing orders. This was triggered by a downstream dependency failure in the VTEX IO platform.

Our global sales flow was unaffected during this incident. Some accounts may have experienced sales below forecasted levels due to customized Checkout flows and payment connectors depending on VTEX IO infrastructure.

Symptoms

Symptom 1: “Something went wrong” errors on Admin login page

An error page appeared when merchants attempted to access their Admin login page, due to increased errors in the VTEX IO render server used by that page, causing most merchants not to be able to use their Admin for 2 hours and 50 minutes (from 09:13 to 12:03 UTC).

Symptom 2: “The circuit is now open” errors / infinite loading storefront components
Increased errors in VTEX IO triggered circuit breakers (infrastructure protection mechanisms) to open, preventing requests associated with these errors from reaching other parts of our

infrastructure and leading some shoppers to experience errors navigating or placing orders for 2 hours and 5 minutes (from 09:49 to 11:54 UTC).

Symptom 3: CHK0223 error code / payment authorization denied

Payment connectors running on VTEX IO presented intermittent communication failures for 2 hours and 37 minutes (from 09:13 to 11:50 UTC), causing AuthorizationDenied errors in some transactions and a generic error to be displayed to shoppers. Failed order attempts were stuck in the “Processing” status as incomplete orders.

Timeline

[2025-03-12 09:13 UTC]	Posterior analysis of our platform health logs indicate increased errors in some VTEX IO clusters.
[2025-03-12 09:17 UTC]	Our incident response team was notified of Symptom 1 : merchants experiencing difficulties logging in to their Admin. We started investigating.
[2025-03-12 09:35 UTC]	<p>Preliminary investigation indicated that the VTEX IO request routing service stopped sending metrics to our observability services. An increase in 5xx errors was also identified in the VTEX IO authorization (authZ) service. At that moment, no correlation was found between these pieces of evidence.</p> <p>We started scaling impacted services as the investigation to find the trigger of the reported symptoms continued.</p>
[2025-03-12 09:49 UTC]	Symptom 2 started impacting stores after the VTEX IO request routing service opened a circuit breaker to protect the authorization service, after the increase in 5xx errors.
[2025-03-12 10:30 UTC]	Symptom 3 was reported and linked to this incident.
[2025-03-12 10:39 UTC]	Our incident response team reverted all recent changes to the request routing service. We identified that pods for the authZ service were stuck in a crash loop.
[2025-03-12 11:02 UTC]	We identified that the authZ service also presented an error sending metrics to our observability services. The error message indicated no healthy upstream was available.
[2025-03-12 11:16 UTC]	We identified that the VTEX IO package manager service also presented an error sending metrics to our observability services. We started investigating potential correlations.

[2025-03-12 11:35 UTC]	We observed a gradual service degradation following recent changes to our observability services. The update affected the metrics collector for the authZ service, impacting both the request routing and package manager services.
[2025-03-12 11:36 UTC]	Our incident response team started reverting recent changes to our observability services. We start observing a gradual improvement in service health metrics.
[2025-03-12 11:47 UTC]	We completed the rollback of the recent changes in observability services and further scaled VTEX IO services to stabilize the platform after a surge of requests.
[2025-03-12 11:53 UTC]	Multiple reports confirmed that Symptom 1 was ceasing, as most merchants were able to load the Admin login page.
[2025-03-12 11:55 UTC]	VTEX IO application metrics were back online and visible in the dashboards of our observability platform.
[2025-03-12 11:59 UTC]	VTEX IO error rates were back to normal operating levels.
[2025-03-12 12:03 UTC]	The incident was fully mitigated.

Mitigation strategy

In the course of our investigation, our incident response team identified that the reported symptoms were triggered by increased errors in the VTEX IO authorization service, following the deployment of changes to our observability services. This deployment was performed several hours prior to the first occurrence of the reported symptoms, but led to a gradual service degradation of the metrics collector used by the VTEX IO platform.

We reestablished normal operations of the platform by reverting recent changes to our observability services and scaling multiple services to stabilize the platform.

Follow-up actions

To prevent similar incidents in the future, our team will work on the following:

- Perform controlled blackouts to identify critical dependencies on the o11y stack
- Perform a technical deep dive on the failure mode observed between the VTEX IO authorization service and the metrics collector used by the VTEX IO platform, in order to completely remove that critical dependency

- Create new alarms to identify unhealthy metrics collectors
- Create automations to more closely monitor deployments to observability services
- Improve monitoring, alerts and tracing for VTEX IO infra services

We are committed to being your trusted partner for success.