


November 22-29, 2023

Root Cause Analysis & Mitigation Plan



Agenda

Incident summary

Incident impact

Root cause analysis

Incident timeline

Piano solutions

Summary

Incident date

November 22-29, 2023

Affected software modules

Insight, DMP, CCE

Affected environment or location

All regions

Incident scope

partial outage

Incident duration

8 days

Impact

Incident severity

Major

Impact description

Customers were not constantly impacted throughout the 8 days incident period. The problem came in waves and lasted for short periods of time at the beginning, but as time went on, these periods lasted longer.

During the problem periods, the latency of some queries increased significantly, and traffic ingestion was much slower than normal, with some cubes experiencing 1-2 hours delay in receiving traffic.

No traffic lost.

Root cause

We installed two new machines in our main European data center. Due to an undiscovered long standing bug in the automated build and deploy process, which was very difficult to trigger, both of the machines were deployed to production with incorrect network configuration. This incorrect config resulted in services setup to communicate with the new machines, periodically incorrectly switching between them. This created a significant increase in query lookup retries, and a huge increase in machine CPU usage.

Incident timeline

| Date / Time (UTC) | Description |
|----------------------|---|
| November 21st, 17:00 | New machines installed and brought online |
| November 21st, 18:10 | The system experienced a relatively insignificant small wave of higher CPU, memory usage and System CPU |
| November 22nd, 14:59 | First wave of noticeable high load |
| November 28th, 11:40 | Slow rollout of the new GC setting in the first 8 servers |
| November 29th, 11:05 | End of the last recorded wave |
| November 29th, 11:32 | Network configuration fixed. Incident is resolved. |

Piano solutions and corrective measures

Mitigation: new GC settings. The new settings allow the JVM to return the memory faster to the OS, keeping the System CPU under control, reducing significantly the impact of the load waves and preventing them from being self-sustaining. We can estimate the new settings to reduce the impact of the waves by about 80%.

Status

Done

Estimated delivery date

November 28, 2023

Better information: we plan to track retries in two systems, to be able to better understand if we are experiencing many retries. In addition, a planned new ingestion dashboard should be make easier to understand the load on the cluster and if the problem is in ingestion or in queries execution.

Status

Planned

Estimated delivery date

December 31, 2023

Piano solutions and corrective measures


Fixed the bug in the automated machine build and deploy process

Status

Done

Estimated delivery date

November 28, 2023



piano

ANALYTICS + ACTIVATION

We help businesses drive revenue by accelerating their ability to understand and influence customer behavior.

Learn more by reaching out to hello@piano.io