

Google Apps Incident Report

Google Docs - March 21, 2013

Prepared for Google Apps customers

The following is the incident report for the Google Drive access disruption that occurred on March 21, 2013. We understand this service issue has impacted our valued customers and users, and we apologize to everyone who was affected.

Issue Summary

Beginning at 5:15 AM PDT on 21 March 2013 and lasting until 11:55 AM, approximately 3.6% of users experienced long load times, "Server Error 503", and "still trying" warnings when trying to access their Google Drive list. Applications using the Google Drive and Docs APIs also returned errors. The impact was most severe in the first few minutes of the event, and was sharply reduced once the engineering team responded. The issue was resolved for the vast majority of users at 11:55 AM. However, the Google Apps Status Dashboard continued to track the resolution over the next few hours for the small number of users who may have experienced intermittent slowness with Drive. Throughout the incident, users could continue to access, view, and edit individual Drive files by direct link or URL.

Actions and Root Cause Analysis

At 4:45 AM PDT, scheduled maintenance on a cluster of servers used to index files in Google Drive triggered a performance bug in the indexing system. This bug had not been triggered during the dry run conducted days earlier for this same maintenance. The Google Drive engineering team was alerted to increasing latency for directory operations at 4:55 AM, and at 5:15 AM determined that the latency increase was sufficient to affect users' experience with Drive. The engineering team first attempted to mitigate the issue by reducing load from nonessential operations, and then by reducing other Google services' use of the same indexing system which was operating slowly for Drive. After eight separate mitigation actions between 5:15 AM and 7:45 AM, the team determined that their efforts were insufficient to restore normal operation for Drive users, and at 7:51 AM they requested that the scheduled maintenance be reverted so that normal user operation could be assured. The server maintenance required several hours to completely unwind. User experience improved at 8:21 AM with another set of mitigation changes, and normal operation was restored at 11:55 AM.

The root cause of this outage was the failure of the pre-maintenance performance testing to trigger the performance bug described above.

Corrective and Preventative Measures

The Google Engineering team conducted an internal review and analysis of the March 21 event. They are taking the following actions to address the underlying causes of the issue and to help prevent recurrence.

- Fix the performance bug in the indexing system. This work has been completed.
- Change the test procedure used for performance assessments to more accurately simulate maintenance events.
- Modify the Drive software to more reliably serve user requests during short periods where overall server capacity is reduced.

- Improve the redundancy of Drive storage systems to assure resiliency and availability during network events that potentially affect performance, such as hardware upgrades.
- Review the preparation and reversion procedures for this type of server upgrade to speed the recovery process.
- Improve flexibility in managing and applying Drive system resources to best support user traffic and requests.

The following is an update on the preventative and corrective measures for the March 18 and 19 Drive issues. The root cause was different from the issue of March 21. However, these measures, which were completed before March 21, helped lessen the impact to users.

- Fixed the bug in the software that manages user connections and sessions with Google Drive, and changed internal structures and resources to make Drive more resilient to latency and errors.
- Accelerated the work in progress that ensures user traffic for Drive is properly prioritized during network events. The core work has been completed, and additional improvements are in progress.
- Increased the capacity of the systems that serve Drive requests well beyond peak demand estimates.
- Improved the Drive alert systems and expanded monitoring of Drive systems for faster detection of issues. New monitoring and alerts to be released shortly, with more rolling out over time.

Google is committed to continually and quickly improving our technology and operational processes to prevent service disruptions. We appreciate your patience and again apologize for the impact to your organization. We thank you for your business and continued support.

Sincerely,

The Google Apps Team