

Google Apps Incident Report

Google Docs - March 18, 2013

Prepared for Google Apps customers

The following is the incident report for the Google Drive access disruption that occurred on March 18, 2013. We understand this service issue has impacted our valued customers and users, and we apologize to everyone who was affected.

Issue Summary

From 6:15 AM to 9:10 AM PT, some users experienced “Server Error 503” messages, long load times, or timeouts when trying to access Google Drive. Applications using the Google Drive and Docs APIs also returned errors. The issue affected up to 33% of all user requests to Google Drive during this period. Users could continue to access individual Drive files by direct link or URL. The root cause of this service disruption was an issue in the network control software.

Actions and Root Cause Analysis

At 6:09 AM PT, a portion of Google’s network capacity went offline due to a novel bug in its control software that surfaced during a scheduled network event. As designed, the Google network routing and load balancing systems responded within seconds, redirecting user traffic to other network connections and servers. The servers receiving the additional user traffic continued to function normally, and as expected, their latency increased with this load. However, the latency increase triggered a second bug in the software that manages user connections and sessions with Google Drive. This resulted in errors and timeouts for some users who were attempting to access Google Drive through the affected servers.

By 6:15 AM, Google Engineering was taking steps to first mitigate the impact to user traffic and then to restore the network capacity. The Google Drive team also took actions targeted at restoring user access as well as beginning the analysis of the bug affecting the Drive interface servers. The network was restored to normal operation by 7:50 AM, after Google Engineering identified the control system bug and implemented a workaround. The recovery was prolonged in part because of the nature of the outage and the care required to avoid expanding the scope of the outage in the process of mitigating or fixing it. Access to Google Drive began to return for users at 8:10 AM, and the issue was resolved for all users by 9:10 AM.

Corrective and Preventative Measures

The Google Engineering team conducted an internal review and analysis of the March 18 event. They are taking the following preliminary actions to address the underlying causes of the issue and to help prevent recurrence:

- Fix the issue in the controls systems which led to the initial event.
Completed. The Google Networking team issued a fix to the affected systems on the evening of March 18, and will release the fix to all systems within days.
- Implement a load balancing policy change to provide greater isolation between different services on the network in the event of a partial failure. This change will reduce the complexity of expected behaviors during partial failure, making software service testing easier and more deterministic.

- Fix the bug within Drive and change internal structures and resources to make Drive far more resilient to latency and errors.
- Improve the Drive alert systems to address the delay between the initial incidence (6:15 AM) and the start of work (6:40 AM), and expand monitoring of Drive systems.
- Accelerate the work in progress that ensures user traffic for Drive is properly prioritized during network events.

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