

Google Apps Incident Report

Google Services - March 17, 2014

Prepared for Google Apps customers

The following is the incident report for the Google services disruption that occurred on March 17, 2014. We understand this issue has impacted our valued customers and users, and we apologize to everyone who was affected.

Issue Summary

From 8:35 AM to 12:10 PM PT, Google Talk, Google Hangouts (including Chat and Video), Google Voice, and the Google App Engine XMPP and Channel APIs, were unavailable for the majority of users. Certain features of other services were affected, including the multi-player features of Google Play Store games. From 9:45 AM to 10:50 AM, some Google Sheets users also experienced very slow responses or 502 errors. The root cause of this disruption was a miscalculation in available capacity during a hardware maintenance event.

Actions and Root Cause Analysis

Background: Google Engineering regularly performs scheduled maintenance on data center systems. Some procedures involve upgrading groups of servers and redirecting the traffic to other available servers. Normally, these maintenance procedures occur in the background with no impact to users.

At 8:25 AM, Google Engineering began a maintenance procedure on a group of backend servers that support Google Hangouts, Google Sheets, and other services, and they redirected the processing load to a new set of backend servers. Due to a miscalculation of memory usage, the new set of backend servers lacked sufficient capacity to process the redirected traffic. These backend servers could not process the volume of incoming requests and returned errors. This led to some Google services retrying requests which further compounded the high levels of user traffic, and backend servers rapidly became overloaded. At 8:35 AM, as incoming requests were dropped, the affected services were no longer accessible to users.

When Google Engineering was alerted at 8:42 AM, they diagnosed the problem and identified the root cause. To mitigate the issue, Google Engineering brought in additional backend server capacity, stopped the maintenance procedure, and began to bring the original backend servers back online.

As the backend servers became overloaded, they began to drop users' connections to their devices; these connections are critical to services such as Chat and Hangouts. The recovery was prolonged because Google Engineering had to reestablish users' connection in stages to avoid overwhelming the backend servers with sudden high volumes of traffic. The initial traffic overload also impacted the capacity of the messaging routers, and this affected the recovery efforts as the backend servers came back online.

Google Engineering continued to add both router and backend server capacity, and some users began to regain access to the affected services beginning at 9:45 AM. Due to the nature of this issue, the Hangouts service was restored for all users at once rather than ramping up as the affected systems came online.

Corrective and Preventative Measures

The Google Engineering team conducted an internal review and analysis of the March 17 event. They are taking the following actions to address the underlying causes of the issue and to help prevent recurrence:

- Review memory requirements and increase the memory capacity for the affected backend servers to meet peak load needs.
- Implement better monitoring for memory utilization and usage tracking to ensure that servers have sufficient capacity available.
- Lower the alert threshold for errors with the Hangouts service to improve Engineering response time.
- Review internal procedures for bringing up emergency capacity to speed mitigation efforts.
- Continue work in progress to improve the resilience of Hangouts service during high load conditions.

Google is committed to continually and quickly improving our technology and operations to prevent service disruptions. We appreciate your patience and again apologize for the impact to your organization.

Sincerely,

The Google Apps Team