

# Postini Services Incident Report

## Mail Delivery - May 7, 2013

Prepared for Postini and Google Apps customers

The following is the incident report for the Postini services outage that occurred on May 7, 2013 (GMT). We understand this service issue has impacted our valued customers and users, and we apologize to everyone who was affected.

### Issue Summary

From 10:15 GMT May 7 to 3:52 GMT May 8, users on Postini System 200 (which comprises 12.7% of all Postini users) experienced severe delays in inbound and outbound mail delivery. The delays were most severe from 12:00 until 21:00, after which time delivery rates began to improve.

During this incident, inbound messages (messages sent to users) were deferred. Outbound messages (messages sent from users) were queued on customers' mail servers. Users who sent messages received a deferral notification with errors such as "421 Server busy, try again later - psmtip". Delivery of the deferred messages was retried based on the sending server's retry interval (which can range from minutes to hours).

A small portion of traffic continued to be processed and delivered throughout the incident. At no time were messages lost or deleted. The root cause of this service outage was a combination of load balancer failures in the primary data center and insufficient processing capacity in the continuation data center.

### Actions and Root Cause Analysis

Background: Postini services run in pairs of data centers, the primary and continuation. Messages are normally processed, filtered, and archived in the primary data center. If there is an issue affecting the primary data center, message traffic may be temporarily switched to the continuation data center.

At 10:15 GMT, mail processing performance began to degrade in the System 200 primary data center, and as designed, the automated monitoring systems directed message traffic to the continuation data center. Google Engineering diagnosed the issue, and at 11:30 GMT, they identified severe instability in the load balancer software, which is provided by a third party, as the core issue in the primary data center. The Engineering team escalated the issue to the third-party vendor and continued investigating the cause and restoration options.

As mail flowed through the continuation center, the message processing systems did not have the sufficient capacity for this sustained volume of traffic. As resources became consumed, this low rate of processing caused delivery delays, and the queued messages and retry attempts led to further processing latency.

At 15:48 GMT, the vendor reported that they had narrowed the source of the problem and were determining the root cause and solution. Throughout the day, Google Engineering continued to provide information to the third-party vendor and conduct their own investigation, and took actions to help reduce user impact.

Engineering detected an sub-optimal use of processing resources in the continuation data center and at 20:40 GMT, they implemented production configuration changes that increased delivery capacity and helped reduce deferrals. Additional performance tuning measures were implemented at 22:20 GMT and 23:20 GMT to provide incremental improvements to mail processing.

At 23:00 GMT, the vendor identified the root cause—a software defect in the load balancer that affected only certain operating system configurations—and began developing a fix. At 2:00 GMT, May 8, Google Engineering implemented the vendor-provided remediation and returned message traffic to the primary data center, and by 3:52 GMT, mail processing returned to normal. Customers' messages that were initially deferred were delivered according to the sending servers' retry interval.

### **Corrective and Preventative Measures**

We understand this was a severe service disruption that took a prolonged time to solve, which was frustrating for our users. The Google Engineering team conducted an internal review and analysis of the May 7 event. They are taking the following actions, a number of which are underway, to address the underlying causes of the issue and to help prevent recurrence:

- Implement fixes and recommendations provided by the vendor to the load balancer systems across all data centers.
- Assign additional storage capacity to the continuation data centers.
- Ensure consistency in performance tuning and configurations between the primary and continuation production systems to optimize performance in the continuation data center.
- Review the escalation response with the vendor to significantly improve the clarity and speed of resolution.
- Improve the Apps Status Dashboard to provide greater visibility and relevant detail about issues in progress.

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