



Root Cause Analysis

Event: Google Apps Webmail Outage, September 1, 2009

Summary: On September 1, 2009, Google Apps webmail users were unable to access webmail interface for approximately 2 hours. IMAP and POP access continued to function normally.

Detail:

Synopsis:

Between 11:45 am and 1:30 pm on Tuesday, September 1, 2009, Google Apps Gmail users were unable to access their accounts through the Gmail web interface. Users could continue to access their accounts via IMAP and POP. No data was lost during this time; messages were received and delivered, but could not be displayed.

Cause:

On Tuesday, September 1, a small portion of Gmail's web capacity was taken offline during a routine upgrade and service update. This is normal operating procedure as the Gmail web interface runs in multiple locations, and Gmail's request routing automatically directs users' requests to available servers.

However, Google underestimated the increased load that some of the new updates placed on request routing. As a result, at approximately 12:30 PDT, a few request routers became overloaded and responded by refusing all incoming requests. This response transferred the load to the other request routers, and as the effect rippled through the system, almost all of the request routers became overloaded. As a result, users could not access Gmail through the web interface since their requests could not be routed to a Gmail server. Gmail processing and access through the IMAP/POP interfaces continued as usual because these processes use different request systems.

Actions Taken:

Upon receiving the error alerts, the Gmail Engineering team immediately began analyzing the issue and initiated a series of actions to help alleviate the symptoms. After determining the root cause to be insufficient available capacity, the Engineering team deployed a large-scale addition of request routers through Google's flexible capacity server systems. As they distributed incoming traffic across the expanded pool of request routers, access to the Gmail web interface returned to normal. During the incident, Google published ongoing reports to the Google Apps dashboard, Gmail Help Center, the Enterprise and Gmail blogs, and the GoogleAtWork and Google Twitter feeds, to help provide customers with the latest status and available workarounds.

University of Alaska Support Center monitored the status throughout and kept key user groups informed.

Future Prevention:

The Gmail Engineering team conducted an internal review and analysis, and determined the following actions to address the underlying causes of the issue and help prevent recurrence:

- Increasing request router capacity well beyond peak demand estimates. This action was completed immediately following the incident, and helps prevent recurrence under similar conditions.
- Isolating failure of request routers so that issues are limited to the specific datacenter, and do not affect servers in other datacenters.
- Addressing request router behavior under load: if multiple routers are simultaneously overloaded, they should continue to perform at a reduced rate rather than refusing connections and attempting to defer their load

Google's RCA can be found at:

http://72.14.213.132/external_content/untrusted/www.google.com/appsstatus/ir/buuqdn6fcervea.pdf