

Google Troubles - A Case Study in Cloud Computing

October 2009

It's tough when you're the big guy on the block. Everybody's watching, and your missteps are often reported with glee. Such is the case with Google. Its services are so broad and so popular that it is bound to stumble once in a while, and when it does it seems to be the target of a feeding frenzy in the press.

At the risk of joining this feeding frenzy, we look at some of Google's latest outages over the last year. Gmail seems to be the Achilles heel, but Google Search, Google News, Google Talk, and other Google services are not immune. As good as Google is, it is understandably not perfect, and this imperfection can send a message to us all. Redundancy.

Tuesday, September 1, 2009 – Gmail Down

During the afternoon of September 1, 2009, Gmail went down for over two hours for its million and a half users. The failure chain began when Google upgraded its routers with the intent of improving reliability. Unfortunately, this upgrade appeared to negatively impact the capacity of the routers – an impact not immediately realized by Google.

But they soon found out, to their chagrin. On that afternoon, Google staff took down several Gmail routers for maintenance – a common practice done hundreds of times in the past without incident. What the staff had underestimated was the additional load that this would put on the remaining routers as they routed Gmail traffic.

The overloaded routers rejected traffic that they could not carry, and this traffic was rerouted through other routers that then became overloaded. This created a classic cascade of increasing workload on fewer and fewer routers. Within minutes, all of the routers in the Gmail network were overloaded, and Gmail crashed.

It took two hours for Google's staff to reroute Gmail traffic across the rest of its immense network that normally carried search traffic in order that Gmail could be restored.

Google has since announced that it is working to make sure that its routers have enough capacity to handle the loss of several Gmail routers due to upgrade activity and that, should such a problem occur again, it can be isolated to just a sector of the network.

Thursday, May 14, 2009 – Google Network Slows to a Crawl

Google Search and Google News slowed to a crawl on May 14, 2009 when 14% of Google's network capacity was lost. The outage that took Google offline for millions of users for over an

hour was caused simply by a routing error which affected most of Google's services including Gmail, YouTube (owned by Google), and even Google's web site.

Google explained that the failure was a result of misdirected traffic. An error in one of its systems caused its network to direct some of its web traffic erroneously through Asia. Many routes became overloaded, causing service requests to repeatedly time out.

Users across the United States, China, Australia, France, and other countries all reported problems.

Tuesday, March 10, 2009 – Google Apps Down for 22 Hours

Subscribers for Google Apps were taken down for several hours on Tuesday, March 10, 2009. Service for many was restored within an hour, but an unlucky subset of users were offline for up to 22 hours.¹

The problem appeared to be a failover fault. A Google Apps server failed, and its backup in a remote data center failed to take over. Matthew Glotzback, at the time the Director of Product Management for Google Enterprise, said in an interview with IDG News Service² that “[when] we can't fail [a] user over for whatever reason – there's an error with the account, or the master and the slave are out of sync ... we can't restore that user's access to the service until we restore that physical location. This is an area where we are constantly getting better ...”

Just a week earlier, a Google glitch allowed some Google Apps files to be accessible to unauthorized users.

Tuesday, February 24, 2009 – Gmail Becomes “Gfail”

Gmail was down around the world for two and a half hours on Tuesday, February 24, 2009. This outage earned Gmail the infamous nickname of “Gfail.”³

Google later explained that the cause was a new feature that they had installed to keep email geographically close its owner. In preparation for a routine maintenance event at one of their European data centers, users were routed to another nearby data center. This inadvertently overloaded that data center, which caused a cascading effect from one data center to another.

Google added actions to its apologies. The very next day, it launched the Google Apps Status Dashboard. In addition, though its Google Apps Premier customers were due an additional three days of free service under Google's SLA, Google extended 15 days of additional service to these users.

Tuesday, February 24, 2009 – Google News Out for 14 Hours

On the same day as the “Gfail” outage, Google News went down for about fourteen hours. Google has not said what caused this outage, and it is not known whether or not it is related to the Gmail problems which also occurred on that day.

¹ “Gmail Becomes Gfail Once Again,” *More Never Agains III*, *Availability Digest*; July 2009.
http://www.availabilitydigest.com/public_articles/0407/more_never_agains_3.pdf

² Gmail Down: Outage Could Last 36 Hours for Some People, *IDG News Service*; March 10, 2009.

³ Has Gmail Become Gfail?, *Availability Digest*; March 2009.
http://www.availabilitydigest.com/public_articles/0403/gmail.pdf

January 31, 2009 – Google Declares Entire Internet as Malware

Google works with StopBadware.com, a non-profit project headed by Harvard and Oxford, to identify suspicious sites that export malware such as viruses and Trojans. Whenever Google lists such a site in response to a search, it cautions the user with the message, “This site may harm your computer,” and links the user to StopBadware.com for further information.

On the morning of Saturday, January 31st, Google did a periodic update of its malware site list. Unfortunately, this update included the URL “/” - a wild card meaning all URLs. Thereafter, for an hour, every site included in a search result was given the warning message.⁴ Google dutifully even labeled itself a malware site. The load on StopBadware.com was so great that it went down under a denial-of-service attack.

October 17, 2008 – Update Takes Down Google Apps

Google Apps is an SaaS (Software as a Service) offering aimed at easing the movement of small businesses to the cloud. A “Start” portal page provides a single sign-on for remote access to a company’s email, instant messaging, voice, video, calendar, storage, document, intranet, and collaborative services provided by Google.

On Friday, October 17, 2008, Google decided unilaterally and without prior warning to update its portal pages to look more like its iGoogle personalized home pages following a major upgrade to iGoogle. Suddenly, links were broken, buttons were misconfigured, and strange “gadgets” caused confusion, preventing access to many Google Apps services. It took days for Google to correct the problems.

Other Outages

The above stories are just some of the recent problems that have created Google headlines in the last several months. Others outages reported by the press include:

August 6, 2008: A small number of Google Apps Premier users suffered a lengthy outage of up to 15 hours.

August 11, 2008: Google Apps was down for two hours. Premier customers were affected.

August 15, 2008: Some Google Apps Premier customers were denied service for up to 80 hours.

Mid-October, 2008: Some Gmail users were locked out of their accounts for about 30 hours.

December, 2008: Google Talk, Google’s chat service, refused to deliver text messages for several hours.

May 18, 2009: Google News was down for about two hours.

September 22, 2009: Google News had an outage lasting two hours.

September 24, 2009: Gmail was down for almost two hours.

October 6, 2009: Gmail went down for one and a half hours.

⁴ “Google Declares Itself to be Malware,” *More Never Agains II*, *Availability Digest*, February 2009
http://www.availabilitydigest.com/public_articles/0402/more_never_agains_2.pdf

Google's New Policy for Transparency

Enterprises using Gmail and other Google services expect communication and transparency regarding outages. After several of its earlier failures and many complaints by users that Google was not communicating with them concerning outages, Google intensified its communication efforts.

Google decided to be very proactive in keeping its users informed about the status of problems after the "Gfail" outage. The day after that event, on February 25, 2009, Google launched its Google Apps Status Dashboard, which can be found at <http://www.google.com/appsstatus#hl=en>.

Though the dashboard covers only Google Apps components (Gmail, Calendar, Talk, Sites, and others), it does not provide information on other services such as Google News. However, for Google Apps, it notes the daily status of each component and provides information concerning the status of outages and other problems.

Google is now also posting outage information on two Twitter accounts.

Google also committed to one-on-one post-mortem discussions with its larger customers following a problem. Google has been much more open with the press about the reasons for outages, as is evidenced by some of the stories above.

The Premier Apps SLA

Google Apps is a free service. However, support can be obtained by paying for a Google Apps Premier subscription. With this subscription comes a Service Level Agreement (SLA) that guarantees 99.9% availability.

Three 9s availability equates to 43 minutes per month. For each month that Google Apps outages exceed 43 minutes a month, Google will provide all customers with three extra days of service at the end of their contract. Should Google Apps not achieve 99% availability within a month (about 7 hours of downtime), Google will give an additional 15 days service to its customers at the end of their contracts.

What About Redundancy?

In this day of high availability and continuous availability, redundancy is king. With all of Google's problems, one might question the extent of redundancy incorporated into Google's systems and networks.

Google admits that it is not fully redundant. In its most recent quarterly 10Q regulatory report filed with the U.S. Securities and Exchange Commission, it states that some of its important data centers are not fully redundant, and that it faces risks from network issues:

"The availability of our products and services depends on the continuing operation of our information technology and communications systems. Our systems are vulnerable to damage or interruption from earthquakes, terrorist attacks, floods, fires, power loss, telecommunications failures, computer viruses, computer denial of service attacks, or other attempts to harm our systems. ... Some of our systems are not fully redundant, and our disaster recovery planning cannot account for all eventualities.

"The occurrence of a natural disaster, a decision to close a facility we are using without adequate notice for financial reasons, or other unanticipated problems at our data centers

could result in lengthy interruptions in our service. In addition, our products and services are highly technical and complex and may contain errors or vulnerabilities.”

Based on its Google Apps SLA, Google has architected its systems to achieve a 99.9% availability goal, and it is struggling to achieve this level. It has violated this goal for some subset of users at a frequency of over once per month (sixteen outages in fifteen months, as reported above).

An availability goal of four or five 9s (as little as five minutes per year of downtime) would be warmly welcomed by Google’s subscriber base, But the increase in subscription costs and the potential loss of free services would not be greeted with such enthusiasm.

There is no question that availability costs money – and perhaps lots of it. Google is striving to balance cost and availability to maximize its market share, and by all measures it seems to be achieving that balance. It has shown that there is a mass market in which three 9s, though perhaps painful, is perfectly acceptable.

Amazon’s Approach

Amazon is a cloud provider that has taken high availability to heart. Its cloud computing and storage services, EC2 and S3, are similar technologically to Google’s systems. Amazon also says in its SEC filings that its systems are not fully redundant.

However, after some early disastrous outages, including one that took down its own Amazon.com store for several hours in June of 2008,⁵ it embraced the redundancy mantra. Amazon now provides redundancy service via its *Availability Zones*.⁶

Amazon divides the world into geographic regions, each containing several Availability Zones. A customer can select an Availability Zone to launch an instance of his application. He can also launch a backup instance in another Availability Zone in the same region. The database in the backup instance is kept synchronized with the primary data database via data replication. Following a primary failure, the backup application instance will assume the IP addresses used to access the application; and the application will be back up and running immediately.

Lessons Learned

Many see cloud computing, in which computing, storage, and application services are available anywhere, anytime, and at a low cost, as the future of computing. This may be so, but we have a long way to go.

Today, by and large, the cloud is useful for applications that are not mission-critical. If your application simply cannot be down, it doesn’t belong in the cloud. But as seen by the commercial success of cloud providers, there are a myriad of applications that benefit from cloud computing.

However, when something goes wrong in your cloud vendor’s data center and the performance or availability of your critical application is compromised, there is little that you can do but wait for the problem to be solved while your users and customers complain and demand information.

Cloud computing is like any other computing system. If you are to achieve high availability, you must have redundancy; and failover has to be rapid, reliable, and automatic. For extremely

⁵ [How Many 9s In Amazon?](http://www.availabilitydigest.com/public_articles/0307/amazon.pdf), *Availability Digest*, July 2008.

http://www.availabilitydigest.com/public_articles/0307/amazon.pdf

⁶ [Can You Trust the Compute Cloud?](http://www.availabilitydigest.com/public_articles/0308/amazon.pdf), *Availability Digest*, August 2008.

http://www.availabilitydigest.com/public_articles/0308/amazon.pdf

important applications, you must have some control over your fate. When something goes wrong with the hardware or software in your cloud vendors' data centers and the performance and availability of the cloud's services are affected, you don't want to be in the position that all you can do is wait for the problem to be solved.

Cloud computing simply provides space, hardware, software, and IT management in someone else's data center, but this is an environment over which you have little control. Maybe it will function flawlessly, and maybe it won't. When it doesn't, what do you do? For critical applications, you need to continue on even if your cloud vendor can't. You need to have planned your own backup contingency plans that are independent of the cloud vendor.

At the very least, you should back up your data at an offsite facility. Perhaps you can have your applications ready to go in another cloud. Perhaps your application data can be replicated from your primary cloud to your backup cloud, as in Amazon's Availability Zones. Perhaps you can manage manually for hours or days if you lose your application.⁷ Whatever your contingency plan is, you must have one.

No matter how reliable a service is, the question is not will it fail? The question is when will it fail?

Acknowledgements

In addition to the references cited above, material for this article was taken from the following sources:

Google's Cloud "Not Fully Redundant," Company Admits, *IT Channel News*; September 25, 2009.
Outages Plague Google, *Searchnewz*; September 25, 2009.

Google Outages Damage Cloud Credibility, *PC World*; September 24, 2009.

Google News hit by downtime, *Periscope IT*; September 23, 2009.

Google News back up after outage, *Computerworld*; September 22, 2009.

Businesses Respond To Gmail Outage, *Information Week*; September 3, 2009.

Gmail outage blamed on capacity miscalculation, CNET; September 1, 2009.

Update: Google's Gmail hit with outage – again, *Computerworld*; September 1, 2009.

Google Apps Hit by Prolonged Gmail Access Problems, *PC World*; August 7, 2008.

⁷ Google has recently made a small step in the direction of manual backup. It has made available an offline version of Gmail. With this service, you can download your email to your own PC, read it, and create replies to be sent when you are next online. At least this provides a manual backup so that some email activities can take place when Gmail is down.