

# Google and IPv6: Another Y2K?

Conference papers make my heart go pitter-patter. You may find the tiny type, academic impedimenta, and wishy washy conclusions the next best thing to Ambitropin.

For the USENIX LISA 2011 Conference in Boston, December 4 to 9, 2011, several Googlers provided a glimpse of the transition from today's Internet to the brave new world of IPv6.

If you are not familiar with IPv6, it is the new version of the Internet Protocol, which is the language that computers use to communicate with each other via the Internet. The addresses used by the current version of the Internet Protocol, IPv4, are exhausted. While technologies such as Network Address Translation (NAT) offer a partial fix, workarounds can cause certain functions to behave in an unexpected or undesirable manner. IPv4 allows 32 bits for an Internet Protocol address. Thus, IPv4 can support 4,294,967,296 addresses. However, IPv6 uses 128-bit addresses. The IPv6 address space supports about  $3.4 \times 10^{38}$  addresses. Your watch and your running shoes may have an IP address in the near future. (For a more thorough explanation of IPv6, visit <http://ipv6.com/>, which is a commercial site offering a range of links and other IPv6 information. The Wikipedia article provides basic information as well. See <http://en.wikipedia.org/wiki/IPv6>.)

The main premise of the paper was a technical road trip. The Googlers piled into a new Internet Protocol and headed down the information highway. Now most organizations will not have a corporate network on a par with Google's. Nevertheless, the Googlers' experiences provide some insight into the shift from a world running out of Internet address space to the promised land of an Internet address for every thing. I use the word "thing" in the sense that refrigerators, children's toys, and remote controlled drones like the Beast of Kandahar will have an Internet address. There will be enough address space to give each intern at the X Factor his or her own Internet address so questions about who did what does not rely on the full time staff's collective memory.

Google has a sprawling footprint for its more than 55,000 full time equivalents, the eight airplanes which belong to senior management, and the thousands upon thousands of servers which contribute to MOMA, the company's Intranet. The shift to IPv6 is not the search network which handles more than three billion queries per day, serves up millions of advertising messages, and delivers the myriad public-facing Google services.

The shift affected Google's internal corporate network that involves desktops, offices and so on. The Google enterprise network, according to *Deploying IPv6 in the Google Enterprise Network: Lessons Learned*,

*consists of heterogeneous vendors, equipment, devices, and hundreds of in-house developed applications and setups; not only different OSes like Linux, Mac OS X, and Microsoft Windows, but also different networking vendors and device models including Cisco, Juniper, Aruba, and Silverpeak. These devices are deployed globally in hundreds of offices, corporate data centers and other locations around the world. They support tens of thousands of employees, using a variety of network topologies and access mechanisms to provide connectivity.*

The authors of the paper are Haythum Babiker, Irena Nikolova, and Kiran Kumar Chittimaneni.

IPv6 is in everyone's future. The reason is that the world has run out of Internet addresses. There have been stop gaps and work arounds, but as the authors assert, it is "the only strategy that makes sense for the long term since only IPv6 can assure the continuous growth of the Internet, improved openness, and the simplicity and innovation that comes [*sic*] with end-to-end connectivity."

IPv6 is important to Google. Without address space, the growth of the Internet and concomitantly Google will be affected. In short, not good.

How did Google shift from the outdated corset of IPv4 to the stretchable waistband of IPv6? The company set up a team and developed a plan that supported an iterative approach. The idea was that "launching small pieces rather than try to complete everything at once." The Google method focused on reliability; otherwise, Googlers would not use the IPv6 system. Finally, Google had to avoid downtime. The deployment of IPv6, therefore, fit into the normal upgrade cycles.

The core of the method used by Google was to follow the guidelines in the Internet Engineering Task Force's RFC 5375. According to Ars Technica's "USENIX: Google Deploys IPv6 for Internal Network":

*Each campus or office got a /48 address block, which meant that it was allotted 280 addresses. In turn, each building got a /56 block of those addresses (or about 272 addresses) and each VLAN (Virtual Local Area Network) received a /64 block, or about 264 addresses. To assign numbers to specific devices, the engineers used the Stateless Address Auto-Configuration capability (SLAAC), which allows the devices to assign numbers to themselves. This approach eliminated the need to manually assign numbers to each device. It was also necessary in that at least some operating systems do not yet support DHCPv6, the version of Dynamic Host Configuration Protocol server-based addressing mechanism for IPv6 networks. (Source: ITWorld at <http://www.itworld.com/networking/231929/usenix-google-deploys-ipv6-internal-network>)*

What were the hurdles the Google team encountered?

Google learned that vendors were not yet fully IPv6 ready. Among the problems Google encountered were IPv6 devices which lacked IPv6 features. Some vendors shipped IPv6 devices with software which caused unacceptably high CPU usage, thus degrading performance. A wireless equipment vendor did not support IPv6 access control lists. There were issues with virtual local area network pooling. Wide area network devices did not support encrypting of Web Cache Control Protocol or WCCP. The Google team said: "We also faced some big challenges when working with various service providers."

The bombshell in the paper was, "We learned a lot of valuable lessons during the deployment process. Unfortunately, the majority of the problems we've faced were unexpected." The authors hit the main points of their information highway road trip:

1. Vendors and providers are not yet ready for IPv6
2. Performance issues abound, putting additional loads on key devices such as routers

3. Software is often not able to support IPv6
4. Applications cause problems. These range from no WCCP to no voice over Internet protocol support for IPv6.

The authors conclude:

*In summary, when it comes to technical problems we can confirm that there is a lot of new, unproven and therefore buggy code, and getting our vendors aligned so that everything supports IPv6 has been a challenge.*

The paper reports that by December 2011, about “95 percent of the engineers accessing our corporate network have IPv6 access on their desks and are whitelisted for accessing Google public services (search, Gmail, YouTube, etc.) over IPv6.” The paper did not provide details about the time and cost of the shift from IPv4 to IPv6. Google held an IPv6 conference in June 2010, so I concluded that the Google change over has taken at least four years. The agenda for the conference and the presentations are available at <http://goo.gl/f7V2>.

Almost one year after the June 2010 conference, Google held an IPv6 test on “World IPv6 Day.” The IPv6 test was available when I submitted this column in mid December 2011. Navigate to <http://ipv6test.google.com/>. If you want to access Google services over IPv6, the company has prepared an information page at <http://www.google.com/intl/en/ipv6/>. Despite the brevity of the document, there are links to a more substantial FAQ at <http://www.google.com/intl/en/ipv6/faq.html>.

First, a tip of the hat to Google for collecting its experiences in the USENIX presentation. Google has provided enterprise systems professionals with useful, real-life information about making the shift from IPv4 to IPv6. Like it or not, most organizations will have to bite the bullet despite the cost and the technical hurdles many will face.

Second, Google’s reputation for engineering excellence is well deserved. The challenges reported in the Google paper are interesting and suggest that a seamless change over may not be possible.

Third, with the change over consuming about four years, organizations may want to do a reality check on their IPv6 status.

Resources which you may find useful include the IPv6 Deployment Aggregated Status pager for 50 top level domains. The table includes clickable links so that IPv6 data for individual countries can be viewed. You can find this resource at <http://www.vyncke.org/ipv6status/>. An interesting and somewhat surprising run down of IPv6 status information appears at IPv6 Status Survey ([http://www.mrp.net/IPv6\\_Survey.html](http://www.mrp.net/IPv6_Survey.html)).

Is IPv6 the 2012 equivalent of the Y2K issue? I am of two minds. For most organizations, the IPv6 shift will be handled as part of normal upgrades. If and when a problem arises, most information technology departments will deal with the problem. On the other hand, as the cloud becomes a larger part of the enterprise applications approach, specialists will take care of the types of problems which Google and other firms have reported.

Stepping back from the Google paper, several thoughts struck me. Google is making clear that it jumped in early, encountered quite a few problems, and--not surprisingly--solved them. The culprits in the trouble were the "vendors". Hardware and software developers were just not in the racing shell quickly enough nor up to the task of IPv6. Google, therefore, suffered so that others could learn.

Another angle I considered was that Google revealed its travels down the IPv6 *autobahn*. Amazon, Microsoft, Rackspace, among others did not. Presumably these other companies made the trip or are making the trip. But the "other guys" did not do the technical equivalent of a tell all in a London tabloid.

Finally, the white paper is a reminder that Google is going to be the technical choice for next-generation services. If I were pitching Google's enterprise applications, I would emphasize that the company is IPv6 ready. By implication, Google's message would be: "Rely on us for technical leadership." My view is that quite a few organizations may respond to this message. Google is serious about maintaining the perception that it is a technical leader. Others may point to Amazon as the technical leader in cloud services, but for the Google faithful, Amazon is a retailer. Google is an engineering company. Personally I am not sure I am comfortable with this argument. Amazon just maintains a lower engineering profile. Google promotes its engineering excellence more enthusiastically.

The adage "If it ain't broke, don't fix it" may not apply to some organizations. Sure, IPv4 is broken. However, the shift to IPv6 is taking place so most services are not disrupted. Google's experiences are instructive, and the company deserves a pat on the back for making the information about its journey available. I think the return to time sharing--today the preferred buzzword is "cloud computing"--is one way to skip the grueling hours of travel and get to the beach quickly. IPv6 is not Y2K yet. Consultants, grab your oars. The race is about to begin.

Stephen E Arnold, December 12, 2011

Mr. Arnold is a consultant. More information about his practice is available at [www.arnoldit.com](http://www.arnoldit.com) and in his Web log at [www.arnoldit.com/wordpress](http://www.arnoldit.com/wordpress). His most recent monograph is *The New Landscape of Search*, which is available from [www.pandia.com](http://www.pandia.com).