

Tweet 920

1.1k

Share 10

Going With the Flow: Google's Secret Switch to the Next Wave of Networking

By [Steven Levy](#)

04.17.12

11:45 AM

Follow @StevenLevy



In early 1999, an associate computer science professor at UC Santa Barbara climbed the steps to the second floor headquarters of a small startup in Palo Alto, and wound up surprising himself by accepting a job offer. Even so, [Urs Hölzle](#) hedged his bet by not resigning from his university post, but taking a year-long leave.

He would never return. Hölzle became a fixture in the company — called Google. As its czar of infrastructure, Hölzle oversaw the growth of its network operations from a few cages in a San Jose co-location center to a massive internet power; a 2010 study by Arbor Networks concluded that if Google was an ISP it would be the second largest in the world (the largest is [Level 3](#), which services over 2,700 major corporations in 450 markets over 100,000 fiber miles.)

‘You have all those multiple devices on a network but you’re not really interested in the devices — you’re interested in the *fabric*, and the functions the network performs for you,’ Hölzle says.

Google treats its infrastructure like a state secret, so Hölzle rarely speaks about it in public. Today is one of those rare days: at the [Open Networking Summit](#) in Santa Clara, California, Hölzle is announcing that Google essentially has remade a major part of its massive internal network, providing the company a bonanza in savings and efficiency. Google has done this by brashly adopting a new and radical open-source technology called [OpenFlow](#).

Hölzle says that the idea behind this advance is the most significant change in networking in the entire lifetime of Google.

In the course of his presentation Hölzle will also confirm for the first time that Google — already famous for making its own servers — has been designing and manufacturing much of its own networking equipment as well.

“It’s not hard to build networking hardware,” says Hölzle, in an advance briefing provided exclusively to Wired. “What’s hard is to build the software itself as well.”

In this case, Google has used its software expertise to overturn the current networking paradigm.

If any company has potential to change the networking game, it is Google. The company has essentially two huge networks: the one that connects users to Google services (Search, Gmail, YouTube, etc.) and another that connects Google data centers to each other. It makes sense to bifurcate the information that way because the data flow in each case has different characteristics and demand. The user network has a smooth flow, generally adopting a [diurnal pattern](#) as users in a geographic region work and sleep. The performance of the user network also has higher standards, as users will get impatient (or leave!) if services are slow. In the user-facing network you also need every packet to arrive intact — customers would be pretty unhappy if a key sentence in a document or e-mail was dropped.



Urs Hölzle. Photo provided by Google.

The internal backbone, in contrast, has wild swings in demand — it is “bursty” rather than steady. Google is in control of scheduling internal traffic, but it faces difficulties in traffic engineering. Often Google has to move many petabytes of data (indexes of the entire web, millions of backup copies of user Gmail) from one place to another. When Google updates or creates a new service, it wants it available worldwide in a timely fashion — and it wants to be able to predict accurately how quickly the process will take.

“There’s a lot of data center to data center traffic that has different business priorities,” says Stephen Stuart, a Google distinguished engineer who specializes in infrastructure. “Figuring out

the right thing to move out of the way so that more important traffic could go through was a challenge.”

But Google found an answer in OpenFlow, an open source system jointly devised by scientists at Stanford and the University of California at Berkeley. Adopting an approach known as Software Defined Networking (SDN), OpenFlow gives network operators a dramatically increased level of control by separating the two functions of networking equipment: packet switching and management. OpenFlow moves the control functions to servers, allowing for more complexity, efficiency and flexibility.

“We were already going down that path, working on an inferior way of doing software-defined networking,” says Hölzle. “But once we looked at OpenFlow, it was clear that this was the way to go. Why invent your own if you don’t have to?”

Google became one of several organizations to sign on to the [Open Networking Foundation](#), which is devoted to promoting OpenFlow. (Other members include Yahoo, Microsoft, Facebook, Verizon and Deutsche Telekom, and an innovative startup called [Nicira](#).) But none of the partners so far have announced any implementation as extensive as Google’s.

Why is OpenFlow so advantageous to a company like Google? In the traditional model you can think of routers as akin to taxicabs getting passengers from one place to another. If a street is blocked, the taxi driver takes another route — but the detour may be time-consuming. If the weather is lousy, the taxi driver has to go slower. In short, the taxi driver will get you there, but you don’t want to bet the house on your exact arrival time.

With the software-defined network Google has implemented, the taxi situation no longer resembles the decentralized model of drivers making their own decisions. Instead you have a system like the one envisioned when all cars are autonomous, and can report their whereabouts and plans to some central repository which also knows of weather conditions and aggregate traffic information. Such a system doesn’t need independent taxi drivers, because the system knows where the quickest routes are and what streets are blocked, and can set an ideal route from the outset. The system knows all the conditions and can institute a more sophisticated set of rules that determines how the taxis proceed, and even figure whether some taxis should stay in their garages while fire trucks pass.

Therefore, operators can slate trips with confidence that everyone will get to their destinations in the shortest times, and precisely on schedule.

Continue reading ‘Going With The Flow: Google’s Secret Switch To The Next Wave Of Networking’ ...

Making Google’s entire internal network work with SDN thus provides all sorts of advantages. In planning big data moves, Google can simulate everything offline with pinpoint accuracy, without having to access a single networking switch. Products can be rolled out more quickly. And since “the control plane” is the element in routers that most often needs updating, networking equipment is simpler and enduring, requiring less labor to service.

Most important, the move makes network management much easier.

By early this year, all of Google’s internal network was running on OpenFlow. ‘Soon we will be able to get very close to 100 percent utilization of our network,’ Hölzle says.

“You have all those multiple devices on a network but you’re not really interested in the devices — you’re interested in the *fabric*, and the functions the network performs for you,” says Hölzle. “Now we don’t have to worry about those devices — we manage the network as an overall thing. The network just sort of understands.”

The routers Google built to accommodate OpenFlow on what it is calling “the G-Scale Network” probably did not mark not the company’s first effort in making such devices. ([One former Google employee](#) has told *Wired’s* Cade Metz that the company was designing its own

equipment as early as 2005. Google hasn't confirmed this, but its job postings in the field over the past few years have provided plenty of evidence of such activities.) With SDN, though, Google absolutely had to go its own way in that regard.

"In 2010, when we were seriously starting the project, you could not buy any piece of equipment that was even remotely suitable for this task," says Hölzle. "It was not an option."

The process was conducted, naturally, with stealth — even the academics who were Google's closest collaborators in hammering out the OpenFlow standards weren't briefed on the extent of the implementation. In early 2010, Google established its first SDN links, among its triangle of data centers in North Carolina, South Carolina and Georgia. Then it began replacing the old internal network with G-Scale machines and software — a tricky process since everything had to be done without disrupting normal business operations.

As Hölzle explains in his speech, the method was to pre-deploy the equipment at a site, take down half the site's networking machines, and hook them up to the new system. After testing to see if the upgrade worked, Google's engineers would then repeat the process for the remaining 50 percent of the networking in the site. The process went briskly in Google's data centers around the world. By early this year, all of Google's internal network was running on OpenFlow. Though Google says it's too soon to get a measurement of the benefits, Hölzle does confirm that they are considerable. "Soon we will be able to get very close to 100 percent utilization of our network," he says. In other words, all the lanes in Google's humongous internal data highway can be occupied, with information moving at top speed. The industry considers thirty or forty percent utilization a reasonable payload — so this implementation is like boosting network capacity two or three times. (This doesn't apply to the user-facing network, of course.)

Though Google has made a considerable investment in the transformation — hundreds of engineers were involved, and the equipment itself (when design and engineering expenses are considered) may cost more than buying vendor equipment — Hölzle clearly thinks it's worth it. Hölzle doesn't want people to make too big a deal of the confirmation that Google is making its own networking switches — and he emphatically says that it would be wrong to conclude that because of this announcement Google intends to compete with Cisco and Juniper. "Our general philosophy is that we'll only build something ourselves if there's an advantage to do it — which means that we're getting something we can't get elsewhere."


To Hölzle, this news is all about the new paradigm. He does acknowledge that challenges still remain in the shift to SDN, but thinks they are all surmountable. If SDN is widely adopted across the industry, that's great for Google, because virtually anything that happens to make the internet run more efficiently is a boon for the company.

As for Cisco and Juniper, he hopes that as more big operations seek to adopt OpenFlow, those networking manufacturers will design equipment that supports it. If so, Hölzle says, Google will probably be a customer.

"That's actually part of the reason for giving the talk and being open," he says. "To encourage the industry — hardware, software and ISP's — to look down this path and say, 'I can benefit from this.'"

For proof, big players in networking can now look to Google. The search giant claims that it's already reaping benefits from its bet on the new revolution in networking. Big time.

Pages: 1 2 View All



Steven Levy's deep dive into Google, [In The Plex: How Google Thinks, Works And Shapes Our Lives](#), was published in April, 2011. Steven also blogs at [StevenLevy.com](#).

Check out [Steve's Google+ Profile](#) .

[Read more by Steven Levy](#)

Follow [@StevenLevy](#) on Twitter.

Tags: [Google](#)

[Post Comment](#) | [61 Comments](#) | [Permalink](#)

[Back to top](#)

[Tweet](#) 920

1.1k

[Reddit](#) [Digg](#) [Stumble Upon](#) [Email](#)

61 comments

5 Stars ▼

Comments for this thread are now closed. ✕

Discussion ▼

Community |



[Adam Williamson](#) • 7 months ago

"In the user-facing network you also need every packet to arrive intact — customers would be pretty unhappy if a key sentence in a document or e-mail was dropped."

facepalm

You...do realize that's not how packet loss works, right? Right? Please, please, tell us you realize that. This is supposed to be a vaguely technical publication. It's not fricking Reader's Digest.

30 ^ | | v • [Share](#) >



[Impk](#) → [Adam Williamson](#) • 7 months ago

It's a popular science text, not even in any tech section, but in Enterprise. It would make the text less readable if it actually included a lecture on CSc to explain a very secondary sentence. In other words: those who know what packet loss is don't need that explanation at all, it's for those who don't, for whom it's an easily understandable explanation, sufficient for understanding the text, even if indeed not very accurate.



Philip Whitehouse → Impk • 7 months ago

No it's extremely important. This why the US has senators on committees overseeing the internet that think it's a bunch of tubes. Wired Enterprise is a respectable technical-focussed look on Enterprise systems.

6 ^ | v • Share ›



horsey_horse → Philip Whitehouse • 7 months ago

Hence the budget item for greased ferrets to keep those tubes working at top speed. At least that's what the congressman said they were for!

8 ^ | v • Share ›



intertubes → Impk • 7 months ago

It's the opposite of accurate; TCP relies on dropped packets to manage its sending window.

2 ^ | v • Share ›



William Raney → intertubes • 7 months ago

I'd make a UDP joke, but I don't know if you'd actually get it...

2 ^ | v • Share ›



windump → intertubes • 7 months ago

no wonder

Hölzle rarely speaks in public - he gets flamed for every little thing. hmm, google infrastructure czar doesn't know what constitutes a dropped packet, /ok/

2 ^ | v • Share ›



uberfu → Impk • 7 months ago

You defeat your own argument by saying it's not very accurate. So this article is simply just made up fluff to entertain the general masses with no real bearing on reality. Granted it did not have to be very technical for those folks who only know insider terms and detailed specifics - BUT if a Writer at a Tech Site that sways influence over a large mass of people is going to report about something as significant as Google shifting it's methods and making them public - you'd think a more accurate analogy could be used for the illustration. If someone is going to write on a subject - be accurate - even if they are watering down the technical for the general public. Don't pass on misinformation and dismiss it as "okay" because those readers wouldn't know the difference. Once upon a time Wired was geared for more technically minded folks and we are still around. And we will call you on the flaws.

1 ^ | v • Share ›



andrew → Adam Williamson • 7 months ago



Maybe he incorrectly assumes email and documents are really just a collection of UDP packets independently assembled?

...nm, I have no good excuses for that ridiculousness.

A UDP packet walks into a

8 ^ | | v • Share ›



paco cornholio → andrew • 7 months ago

man, "Byte her?! I didn't even

1 ^ | | v • Share ›



quiop → Adam Williamson • 7 months ago

"

You... do realize that's not how packet loss works, right? Right?"

It is a Wired article so probably not.

1 ^ | | v • Share ›



David Khanaferov → Adam Williamson • 7 months ago

Dude when I read that I was like ??????????

1 ^ | | v • Share ›



Jeff Hancock • 7 months ago

This looks like the typical hyper-article with all the ingredients tossed in to sell -- 1 cup of tech titan, a dash of bearded engineer with a weird name, a teaspoon of "open" and voila, tasty!

4 ^ | | v • Share ›



Kevin • 7 months ago

What a novel idea, separating network management from the switching function! Why that's only been around for a few decades in the telecommunications world since at least the advent of Signaling System 7 (SS7).

4 ^ | | v • Share ›



D B → Kevin • 7 months ago

Oh my god, you're right! How could every enterprise and carrier network have been missing that for all these years! The way they accomplish this currently is to.....wait a minute....

1 ^ | | v • Share ›



Kevin → D B • 7 months ago

Did I say that OpenFlow is a bad idea? No.... I said it's not a novel idea as in not new, been there done that in an analogous industry and technology. What took them so long and the hype of it's coming as a concept is being way overdone. Obviously from

practical, implementation, management and operations perspectives the hype of the actual existence of functional product may be well deserved

0 ^ | v • Share ›



Alex Murphy → Kevin • 7 months ago

"What took them so long" - Well, since Google is the first major internet company to implement OpenFlow, your real question should be: what's taking every other company in the world so much LONGER?

3 ^ | v • Share ›



rapier1 → Alex Murphy • 7 months ago

Depends on what you mean by 'major' company. OpenFlow has been used in the Internet2 and NLR networks for a while. More on a testbed basis but that's the point in experimental research backbones.

2 ^ | v • Share ›



rapier1 → Kevin • 7 months ago

Novel? No. It is based in part on XORP which goes back to 2000. So OpenFlow is clearly an evolutionary advancement over previous ideas. Anyway, the article doesn't do a good job on explaining what OpenFlow is - it does separate the functionality *but* more importantly it is an opensource extensible framework to build and implement your own routing protocols without having to wait for vendors to roll them out as a product. There is a lot of neat stuff you can do with it that goes beyond what they discussed here.

Note: I'm not saying that this article was bad but it was written for the popular science wired crowd who likes to hear more about gee whiz stuff rather than drilling into the details.

1 ^ | v • Share ›



itsaustralian • 7 months ago

Lol, the only reason they are making this announcement is because their shipping company delivered the switch to the building next door and they started to post pics about this new device. Its basically it's a massive fuck up at google right now and they decided to save face and just go public with the new tech.

1 ^ | v • Share ›



ccie5000 → itsaustralian • 7 months ago

Or maybe Google doesn't want to be in the switch manufacturing business, and would like major vendors to adopt OpenFlow.

0 ^ | v • Share ›



Kai → itsaustralian • 7 months ago

W... ..

What? Note that Urs Hoelzle gave the talk at the Open Networking Summit. I o
organise the summit, the organisers probably have already been looking for speakers for
months in advance, and then announced the programme for the summit, and then sold tickets
to the summit. It's not as though it is a sudden announcement.

0 ^ | | v • Share ›



jnemesh • 7 months ago

It seems like the comments so far are extremely negative. Why? This tech seems to
revolutionize the structure of the internet, making it VASTLY more efficient and much more
manageable on the large scale! I think many of the people commenting are forgetting that this tech
is for the "big iron" on the net...the stuff running on the backbone, NOT the systems connected to
the average user. Anything they can do to keep costs down and improve efficiency on the 'Net
should be welcomed with open arms!

1 ^ | | v • Share ›



rapier1 → jnemesh • 7 months ago

It's not revolutionary. OpenFlow is more evolutionary and comes from a basic
understanding of the problems facing backbones. The real advantage of OpenFlow, which I
don't think they really discussed enough here, is that you don't have to rely on a vendor to
supply your management methodologies. You really can write your own and experiment with
different ideas.

2 ^ | | v • Share ›



Eve Moynihan → jnemesh • 7 months ago

Google isn't building anything that runs on any of the Internet backbone networks.
They're doing this for their own internal backbone.

The Internet backbone networks are orders of magnitude larger than Google's internal
networks, and those networks are (largely) built and maintained by the carriers, using purpose-
built technology which really wouldn't make sense for individual companies to use.

0 ^ | | v • Share ›



jnemesh → Eve Moynihan • 7 months ago

Interesting. I got that they weren't deploying this for connections to end users,
but I guess I just assumed that their internal networks were on the internet backbone. It
didn't occur to me that they have a different WAN in place.

0 ^ | | v • Share ›



KC Berg • 7 months ago

Tier 3??? The company is Level 3 Communications and they also happen to be Google's
ISP.

1 ^ | | v • Share ›



KC Berg • 7 months ago



UB Berg • 7 months ago

Right, because Google only has 1 ISP.....

1 ^ | v • Share ›



Mark Simchock • 7 months ago

Geeky networking stuff aside.

"Our general philosophy is that we'll only build something ourselves if there's an advantage to do it — which means that we're getting something we can't get elsewhere."

Don't let Goggle's use of "open source" fool you. While most view open source as an ends, Google clearly treats it as a means. They take that means, add their own magic (i.e., have the premier resources) and that ends is a unique competitive uber-advantage. Of course Google is pro open source. It lets others do the R&D and much of the dirty work. Google then sweeps in, picks the free cherries and cooks'em up into a pretty sexy ROI pie. Meanwhile the proletariat waving the open source flag eat cold pizza and drink Red Bull.

Evil? You tell me.

1 ^ | 1 v • Share ›



David W • 7 months ago

Its always amazing to me that the greatest achievement that take companies like google to the next level of sophistication are barely noticed by the end users. I jumped over to this article from a post by Bill Slawski at SEObythesea.com am glad I did. It explains why google has been able to to create dramatic updates to its search functionality in such a short timespan. Bottleneck=gone! at least for now, until they add an entirely new level of data crunching sophistication.

0 ^ | v • Share ›



ccie5000 • 7 months ago

Juniper popularized the separation of packet switching and management years ago. (Juniper routers have separate Packet Forwarding Engines and Routing Engines.) I think what's new with OpenFlow is not separation, but centralization of management, as explained in the taxi analogy.

That's a big change. The Internet started as an ARPA project, and survivability was one of the design criteria. That had a huge influence on how networking evolved as separate network elements working together to make independent routing decisions.

Now, with OpenFlow, we see the emergence of centralized management. That has huge implications on network robustness, reliability and survivability. It will be very interesting to see how this develops.

One minor nit: it bugs me when people don't capitalize "Internet". An internet is any TCP/IP network. The singular and unique Internet (proper noun, thus capitalized) is the TCP/IP network to which billions of us are connected. :-)

Thank you,

Jerry J. Anderson, CCIE #5000

0 ^ | v • Share ›



MichaelADeBose • 7 months ago

"Hölzle doesn't want people to make too big a deal of the confirmation that Google is making its own networking switches" Indeed! Motorola Mobility makes home networking equipment as well as set top boxes for Verizon and others whos' set top boxes are becoming media WAS(customer mobile gear)/NAS(customer home bound gear) with increasingly sophisticated IP properties. Today most telephone guys from 10-15 years back who didn't recognize that VOIP was about the network guys slowly taking over what they do/did would probably warn all others that the digitization of content, any content, makes its routable and thus within the domain of packets and the network tasked to oversee their routing.

OpenFlow should really look a lot more concerning to the traditional guys if considered in conjunction with Intel's Integrated I/O and Data Direct I/O. Owning Motorola Mobility(MM) allows Google to take its expertise and bypass the competition because they already have a foothold with consumers with the gear they provide and with the cable companies and telcos because of the already enforce contracts they have with Motorola. The new box Verizon is moving with increased storage for FIOS TV and the other boxes they replace are built by MM. Even choosing to deal with another company would materially impact Verizon's bottom line. If you think Samsung and Android makers feel threatened by Google's MM purchase. I'm sure Verizon is feeling the heat because right now MM makes a lot of gear used by FIOS TV. It is rumored that Google may sell MM to Huawei, which does sell networking as well as Android gear. The media has never discussed Google's purchase of MM in any terms other than Android, but in the end the network which is the biggest elephant in a world's living room, manages to stay quite invisible. Indeed.

0 ^ | v • Share ›



Level380 • 7 months ago

Great read to see what Google is doing to get the most out of every dollar spent in the network. One would think that the higher costs of building your your networking gear would be offset by the savings of driving those links to nearly 100% usage. Fibre running across the ocean isn't cheap, so if you can save adding more there is huge cost savings to be had!

0 ^ | v • Share ›



Praseon Kumar • 7 months ago

Wow, just yesterday, I was reading how ex-Cisco are doing Nicira to make Cisco obsolete and here comes a big usecase of the same.

0 ^ | v • Share ›



Essex Blur • 7 months ago

While they've perhaps done something interesting, how does it deal with multiple fibre-cuts?

Not well you say, really ?

0 ^ | v • Share ›



amanfromMars • 7 months ago

"As for Cisco and Juniper, he hopes that as more big operations seek to adopt OpenFlow, those networking manufacturers will design equipment that supports it. If so, Hölzle says, Google will probably be a customer."

Err, is that not a bit like reinventing the wheel or trading with the enemy, and would not Google be missing out on a valuable revenue stream and market leading product placement in some pretty vital international services by not supplying their new super duper superior secretly developed switches to the likes of a Cisco and Juniper?

Or does the information business level they are operating in with such switches take the flow of available intelligence to higher and wider and even more sensitive [as in strictly need to know] levels, and thus would third party purchase be inappropriate and disadvantageous?

0 ^ | | v • Share ›



JK • 7 months ago

Google isn't first company to hack networking to meet is peculiar needs. In early 1990s, I worked at TRW Financial Systems in Oakland, CA, and their unix engineers had created "big packet" network drivers to move the large images of scanned documents between scanners and storage and display workstations.

0 ^ | | v • Share ›



Tex4n • 7 months ago

Sounds like OpenFlow means that the Internet would not be such a 'dumb pipe', which means more central control and less innovation on the edges. Network operators would have more control to restrict services and applications rather than just provide simple bandwidth and network speed.

0 ^ | | v • Share ›



rapier1 • 7 months ago

So it's cool that Google is OpenFlow to centralize it's flow management. It's a decent idea that really maximizes network utilization. Of course, I do wonder about how centralization might make it more prone to catastrophic failure in some circumstances.

0 ^ | | v • Share ›



Cai Cai • 7 months ago

So what does that mean for SNMPv3 and what level of security measures will be taken?

0 ^ | | v • Share ›



Guest • 7 months ago

I just *hate* it when an email arrives without some of its packets. Web pages, too!

0 ^ | | v • Share ›



AngryBuddha • 7 months ago

Would be nice for Google to have a positive cash flow from something other than search.

0 ^ | | v • Share ›



motytrah • 7 months ago

Start the countdown clock to a Cisco Patent dispute. I can see why Google is so cagey to talk about networking. It leaves them open to lawsuits from Patent Trolls.

0 ^ | | v • Share ›



rapier1 → motytrah • 7 months ago

I don't see Cisco doing anything like that. OpenFlow is hardly a secret and I know that Cisco is aware of it.

0 ^ | | v • Share ›



Trilogy Chiropractic Ballard • 7 months ago

Gmail is crashing big time this morning, millions of people are locked out of their accounts, cannot access their business docs, calendars, email, etc. Could this be why?

0 ^ | | v • Share ›



rapier1 → Trilogy Chiropractic Ballard • 7 months ago

No. This has been in place since 2010.

0 ^ | | v • Share ›



hemipw54 • 7 months ago

What, not a NASA idea, that disgusts me. ROFL

0 ^ | | v • Share ›



Guest • 7 months ago

Could this be responsible for the down service with Gmail now?

0 ^ | | v • Share ›



David Khanaferov → Guest • 7 months ago

if you were actually serious when you typed this comment you sir are obtuse.

0 ^ | | v • Share ›



ZangoHooo • 7 months ago

It has become pretty clear that Google will rule the world one day!

www.Total-Privacy dot US

0 ^ | | v • Share ›

Load more comments

Before the iPad, There Was The Honeywell Kitchen Computer

41 comments • 2 days ago



WayneofOZ — I think Apple should sue them for back to the future patent infringements. There has to be some in t...

The HP Way, and How It Completely Screwed HP

16 comments • 3 days ago



FrankOnSecurity — How is this a failure of the "HP Way". The author apparently does not know what the "HP Way" really ...

Video: Watch the World's Oldest Digital Computer Get a Reboot

30 comments • 2 days ago



mijikai — You have no feeling for human achievement. Possibly because you've never taken part in any.

Inside the James Bond Fantasy Data Center That Housed Wikileaks

18 comments • 7 days ago



WebDevery — Why can't Wired figure out how to use AJAX? Full page refreshes for a slide show of photos are obnox...

