STEVEN LEVY BUSINESS APR 17, 2012 11:45 AM

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

Google treats its infrastructure like a state secret, so Google czar of infrastructure Urs Hölzle rarely ventures out into the public to speak about it. 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.

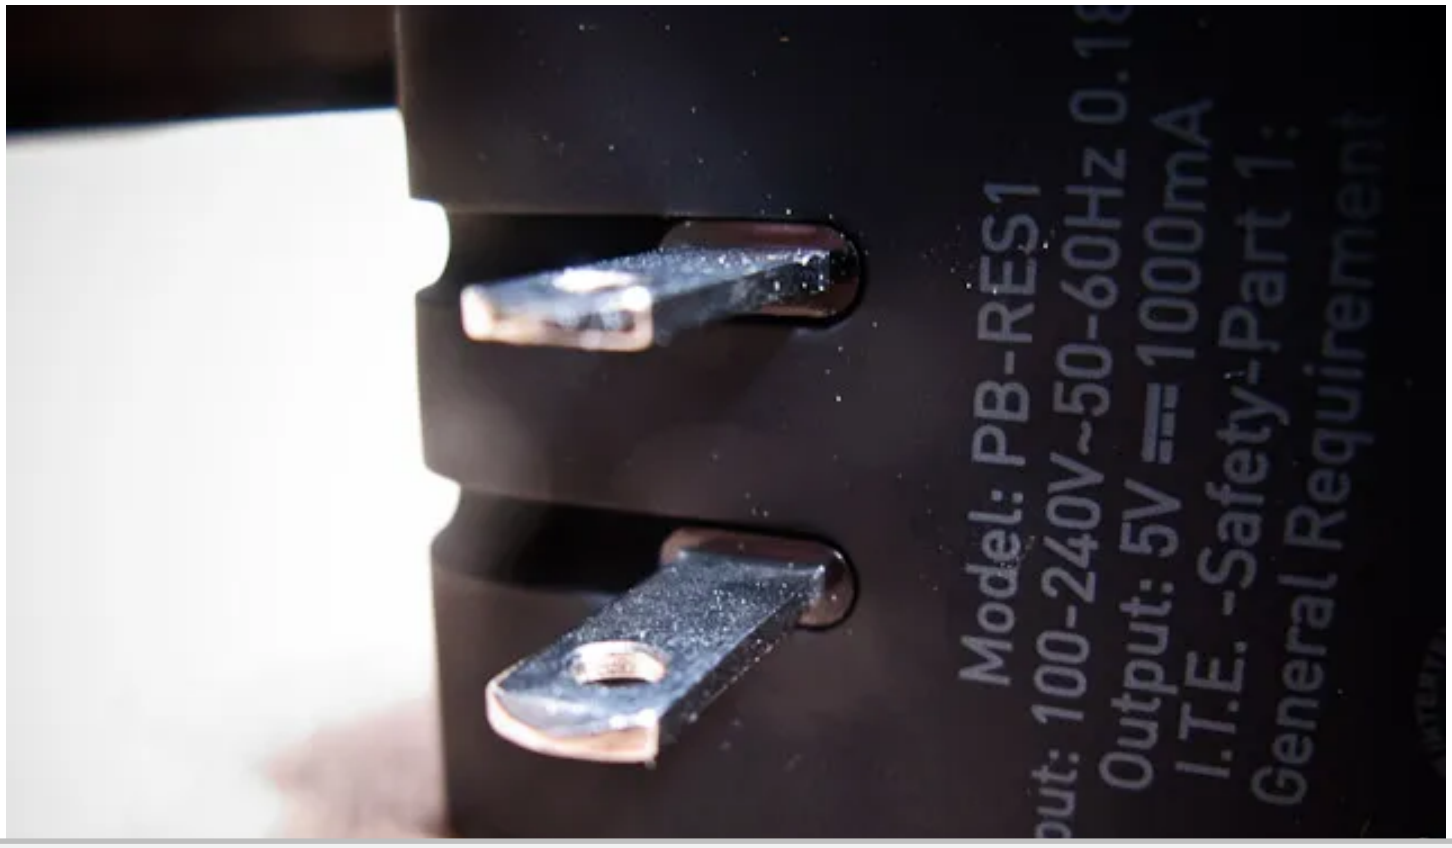

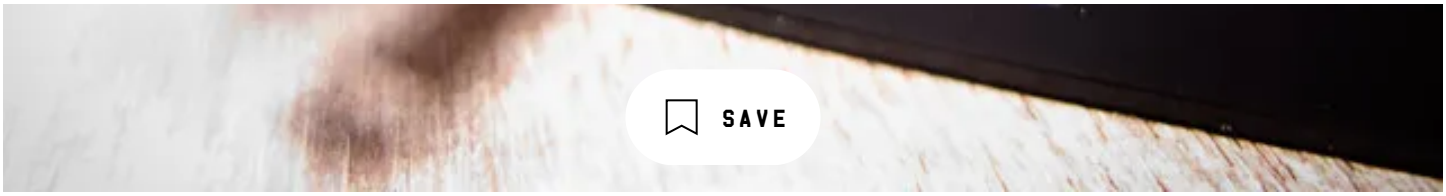

CHARLIE SORREL

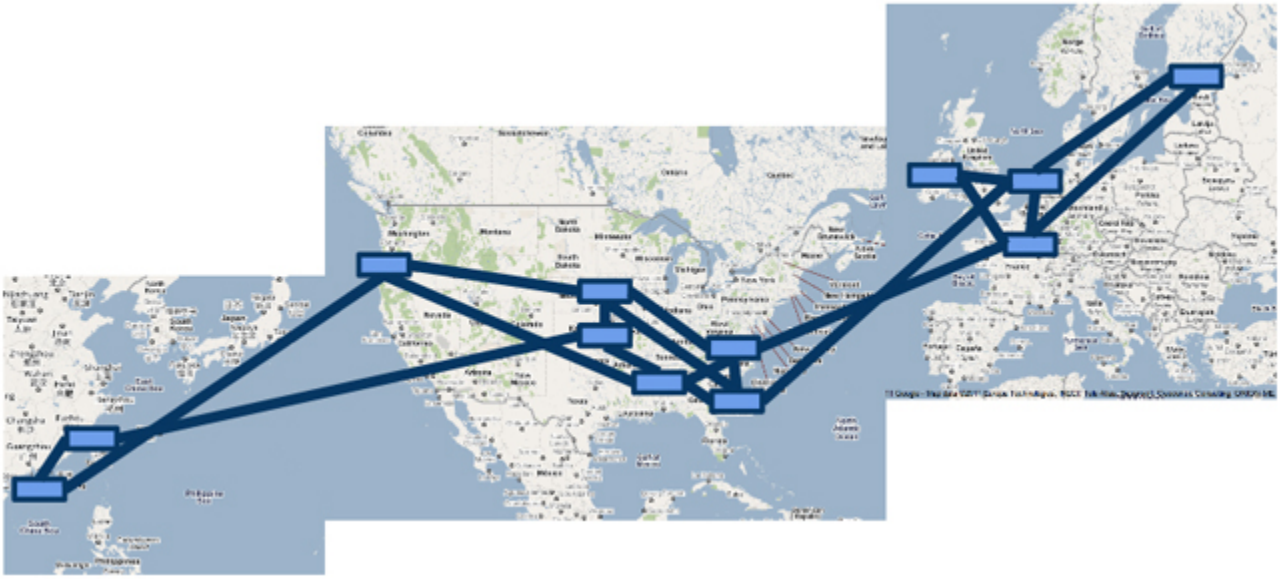

In early 1999, an associate computer science professor at UC Santa Barbara climbed the steps to the second floor headquarters of a small startup in Palo Alto, and wound up surprising himself by accepting a job offer. Even so, Urs Hölzle hedged his bet by not resigning from his university post, but taking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

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 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.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.