

# Upgrading the Internet for the mobile age

August 1 2012, by John Sullivan

---



A team of Princeton researchers, including Erik Nordström (left), a research scholar, and Michael Freedman, an assistant professor of computer science, is developing a system to address the problems involved with the Internet's growing pains. Called Serval, the system makes a small change to the way programs download and manage data that could have a big impact on the future development of the Web. (Photo by Frank Wojciechowski)

(Phys.org) -- Some strands of the World Wide Web are getting a little tattered.

When it comes to delivering data to users, the Web still works brilliantly. But for other functions such as allowing users to move between [wireless networks](#) or companies to shift traffic among [servers](#), engineers are forced to implement increasingly cumbersome [tweaks](#).

In part, that is a function of the Internet's growth. What began as an

elegant method to transfer data among a small number of computers has evolved into a massive makeshift system in which [designers](#) have added workarounds on top of earlier layers of fixes.

Now, a team of Princeton University researchers has released a plan to cut through that tangle and provide a simple solution to many of the problems involved with the Internet's growing pains. Called Serval, the system makes a small change to the way programs download and manage data that could have a big impact on the future development of the Web.

"Smart engineers are great at mastering complexity and hacking around it. But with the current approach, everybody needs to reinvent the wheel," said Michael Freedman, an assistant professor of [computer science](#) who heads the Serval development team. "Serval is something that will make things easier to build and easier to manage without relying on these hacks."

The team has put together a small Princeton network to demonstrate Serval's function and is moving into testing the system in larger and more complex deployments. Serval fixes problems by augmenting the role of one of the foundations of the Internet — the IP address. The IP, or Internet Protocol, address is the fundamental identifier for computers and other devices on the Internet. It is critical both for transferring information and for programs that use the data at its ultimate destination.

One way to think about this is to imagine the IP address is a street address such as 1 Main St. When packets of information need to find their way to your computer, they use the IP address. When the information arrives, the IP address of the transmitting computer also plays a role in allowing your computer to recognize information needed for specific applications such as Firefox or QuickTime (there are a number of other programs involved, but the IP address is one step in

allowing your computer to say 'Oh, yes, this is the data I need to run this YouTube video.')

That worked well in the early days of the Internet when users wanted to connect to a specific computer. But now, people are not interested in connecting to a single computer, they want to connect to services such as Google or Wikipedia. Unfortunately, because the IP address is still linked to specific machines, there is no easy way to do that.

The problem is that the IP address cannot change without messing everything up. Computer programs that rely on the IP address have no easy way to adjust if the IP address changes, a possibility that is increasingly common as mobile devices proliferate. Without a clever hack, the program just stops.

"A lot of applications can be rewritten to hide the fact that you break your connection and re-establish it," said Jennifer Rexford, the Gordon Y.S. Wu Professor in Engineering who is also working on Serval. "But it makes things unnecessarily complex and more expensive."

Serval does not replace the IP address, which continues to function as usual. Instead, Serval works by adding another layer of information, called the service access layer, to certain packets contained in a data stream. Unlike the IP address, which identifies a specific computer, the service access layer identifies a service such as Google or Facebook. So if the IP address changes, the program does not stop because it receives the functional information it needs from the service access layer.

This makes it easier for users to watch a movie while their tablet computer switches between a Wi-Fi stream and a cellular one. It also allows a mobile device to use Wi-Fi and cellular service at the same time.

"The Internet is increasingly a platform for connecting mobile users with the cloud and the Internet was not designed to do that," Rexford said.

"All the work that goes into shoehorning that into the old network architecture is very difficult."

Serval also makes it much easier for Internet companies to shift traffic among different servers. While not as immediately important to individual users, this function could allow for drastic increases in efficiency for Internet companies, and it could play an important role in the future development of the Internet in general.

Remember the 1 Main St. analogy? Big companies such as Google have expanded far beyond an initial address such as 1 Main St. Now, the company has thousands of servers, branch offices if you will, and Serval will allow company managers to switch users among offices seamlessly or allow subdivisions within individual offices to serve different users. To do that now, the companies have to rely on expensive and hard-to-manage workarounds. With Serval, it would be straightforward.

Gary Berger, a lead engineer at Cisco Systems, said bringing greater flexibility to the network is a goal that many organizations would welcome. Although not affiliated with the development of Serval, Berger is reviewing the project to see whether Cisco could perform some in-house testing of the Serval system. Internet companies are looking at a number of options, but Berger said that Serval offers a possible solution to basic problems that are limiting the growth of the Internet — many of which can't be easily solved by further patches.

"You can't solve all the problems by incrementally adding to the architecture because it is fundamentally broken at a basic level," Berger said. "Serval helps us to think differently about the problem, which will benefit everyone."

One advantage of Serval is that the system can be deployed incrementally and does not require massive changes to current networks. For users who do not use Serval, the additional data layer in packets will be invisible and the IP address will remain as the identifier. In cases where Serval is used, the researchers estimate that it will increase the amount of time data takes to transit the Web by less than 1 percent.

Berger said the Serval team has already addressed the issue of making applications work with the service access layer rather than the IP address alone. He said that was the most technically [complex](#) and the most important part of Serval. But Berger noted there are still questions that have to be answered regarding how Serval will function in the broader context of the Internet, and what potential problems could crop up if it were widely deployed.

Erik Nordström, a research scholar in the computer science department and the lead author of the team's conference [paper](#) on Serval, said the group has been meeting with Internet companies and service providers to inform them of the system's advantages. He said it can be difficult to convince large organizations to make fundamental changes, particularly when they already invested considerable sums into the software that runs their systems.

"We'd like to understand how we can make Serval more attractive for others to adopt, especially when it comes to managing large-scale Web services," Nordström said.

Serval was developed with support from the National Science Foundation, the Defense Advanced Research Projects Agency, the Office of Naval Research and Cisco Systems. The authors of the Serval paper were: Nordström, David Shue, Prem Gopalan, Robert Kiefer, Matvey Arye, Steven Y. Ko, Rexford and Freedman. All are with the Princeton School of Engineering and Applied Science's Department of

Computer Science.

Nordström said the Serval team would like to follow up with a wider deployment of the system to understand engineering challenges at a larger scale and in a more realistic environment.

He said the next step will allow the team "to better understand practical limitations and benefits of Serval, and also provide us with a complete deployment strategy that we can present to interested parties."

Provided by Princeton University

Citation: Upgrading the Internet for the mobile age (2012, August 1) retrieved 27 April 2024 from <https://phys.org/news/2012-08-internet-mobile-age.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--