

# **A new way to protect computer networks from Internet worms**

June 4 2008

---

Scientists may have found a new way to combat the most dangerous form of computer virus. The method automatically detects within minutes when an Internet worm has infected a computer network. Network administrators can then isolate infected machines and hold them in quarantine for repairs.

Ness Shroff, Ohio Eminent Scholar in Networking and Communications at Ohio State University, and his colleagues describe their strategy in the current issue of IEEE Transactions on Dependable and Secure Computing.

They discovered how to contain the most virulent kind of worm: the kind that scans the Internet randomly, looking for vulnerable hosts to infect.

"These worms spread very quickly," Shroff said. "They flood the Net with junk traffic, and at their most benign, they overload computer networks and shut them down."

Code Red was a random scanning worm, and it caused \$2.6 billion in lost productivity to businesses worldwide in 2001. Even worse, Shroff said, the worm blocked network traffic to important physical facilities such as subway stations and 911 call centers.

"Code Red infected more than 350,000 machines in less than 14 hours. We wanted to find a way to catch infections in their earliest stages,

before they get that far," Shroff said.

The key, they found, is for software to monitor the number of scans that machines on a network send out. When a machine starts sending out too many scans -- a sign that it has been infected -- administrators should take it off line and check it for viruses.

The strategy sounds straightforward enough. A scan is just a search for Internet addresses -- what we do every time we use search engines such as Google. The difference is, a virus sends out many scans to many different destinations in a very short period of time, as it searches for machines to infect.

"The difficulty was figuring out how many scans were too many," Shroff said. "How many could you allow before an infection would spread wildly? You want to make sure the number is small to contain the infection. But if you make it too small, you'll interfere with normal network traffic."

"It turns out that you can allow quite a large number of scans, and you'll still catch the worm."

Shroff was working at Purdue University in 2006 when doctoral student Sarah Sellke suggested making a mathematical model of the early stages of worm growth. With Saurabh Bagchi, assistant professor of electrical and computer engineering at Purdue, they developed a model that calculated the probability that a virus would spread, depending on the maximum number of scans allowed before a machine was taken off line.

In simulations, they pitted their model against the Code Red worm, as well as the SQL Slammer worm of 2003. They simulated how far the virus would spread, depending on how many networks on the Internet were using the same containment strategy: quarantine any machine that

sends out more than 10,000 scans.

They chose 10,000 because it is well above the number of scans that a typical computer network would send out in a month.

"An infected machine would reach this value very quickly, while a regular machine would not," Shroff explained. "A worm has to hit so many IP addresses so quickly in order to survive."

In the simulations pitted against the Code Red worm, they were able to prevent the spread of the infection to less than 150 hosts on the whole Internet, 95 percent of the time.

A variant of Code Red worm (Code Red II) scans the local network more efficiently, and finds vulnerable targets much faster. Their method was effective in containing such worms. In the simulations, they were able to trap the worm in its original network -- the one that would have started the outbreak -- 77 percent of the time.

Anywhere from 10 to 20 percent of the time, it spread to one other network, but no further. The remaining 3 to 13 percent of the time, it escaped to more networks, but the infection was slowed.

In all cases, there was a dramatic decrease in the spread of the worm within the first hour.

To use this strategy, network administrators would have to install software to monitor the number of scans on their networks, and would have to allow for some downtime among computers when they initiate a quarantine.

According to Shroff, that wouldn't be a problem for most organizations. Very small businesses -- ones with only a few servers -- may have more

difficulty taking their machines off line.

"Unfortunately there is no complete foolproof solution," Shroff said.  
"You just keep trying to come up with techniques that limit a virus's ability to do harm."

He and his colleagues are working on adapting their strategy to stop targeted Internet worms -- ones that have been designed specifically to attack certain vulnerable IP addresses.

Source: Ohio State University

Citation: A new way to protect computer networks from Internet worms (2008, June 4) retrieved 8 May 2024 from <https://phys.org/news/2008-06-networks-internet-worms.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>
--