


# Computer scientists successfully boot one million Linux kernels as virtual machines

September 25 2009

---



Sandia National Laboratories computer scientists Ron Minnich (foreground) and Don Rudish (background) have successfully run more than a million Linux kernels as virtual machines, an achievement that will allow cybersecurity researchers to more effectively observe behavior found in malicious botnets. They utilized Sandia's powerful Thunderbird supercomputing cluster for the demonstration. (Photo by Randy Wong)

(PhysOrg.com) -- Computer scientists at Sandia National Laboratories in Livermore, Calif., have for the first time successfully demonstrated the ability to run more than a million Linux kernels as virtual machines.

The achievement will allow cyber security researchers to more effectively observe behavior found in malicious botnets, or networks of infected machines that can operate on the scale of a million nodes.

Botnets, said Sandia's Ron Minnich, are often difficult to analyze since they are geographically spread all over the world.

Sandia scientists used virtual machine (VM) technology and the power of its Thunderbird supercomputing cluster for the demonstration.

Running a high volume of VMs on one supercomputer — at a similar scale as a [botnet](#) — would allow cyber researchers to watch how botnets work and explore ways to stop them in their tracks. “We can get control at a level we never had before,” said Minnich.

Previously, Minnich said, researchers had only been able to run up to 20,000 kernels concurrently (a “kernel” is the central component of most computer operating systems). The more kernels that can be run at once, he said, the more effective [cyber security](#) professionals can be in combating the global botnet problem. “Eventually, we would like to be able to emulate the computer network of a small nation, or even one as large as the United States, in order to ‘virtualize’ and monitor a [cyber attack](#),” he said.

A related use for millions to tens of millions of operating systems, Sandia's researchers suggest, is to construct high-fidelity models of parts of the Internet.

“The sheer size of the Internet makes it very difficult to understand in even a limited way,” said Minnich. “Many phenomena occurring on the Internet are poorly understood, because we lack the ability to model it adequately. By running actual [operating system](#) instances to represent nodes on the Internet, we will be able not just to simulate the functioning of the Internet at the network level, but to emulate Internet functionality.”

A virtual machine, originally defined by researchers Gerald J. Popek and

Robert P. Goldberg as “an efficient, isolated duplicate of a real machine,” is essentially a set of software programs running on one computer that, collectively, acts like a separate, complete unit. “You fire it up and it looks like a full computer,” said Sandia’s Don Rudish. Within the virtual machine, one can then start up an operating system kernel, so “at some point you have this little world inside the [virtual machine](#) that looks just like a full machine, running a full operating system, browsers and other software, but it’s all contained within the real machine.”

The Sandia research, two years in the making, was funded by the Department of Energy’s Office of Science, the National Nuclear Security Administration’s (NNSA) Advanced Simulation and Computing (ASC) program and by internal Sandia funding.

To complete the project, Sandia utilized its Albuquerque-based 4,480-node Dell high-performance computer cluster, known as Thunderbird. To arrive at the one million [Linux](#) kernel figure, Sandia’s researchers ran one kernel in each of 250 VMs and coupled those with the 4,480 physical machines on Thunderbird. Dell and IBM both made key technical contributions to the experiments, as did a team at Sandia’s Albuquerque site that maintains Thunderbird and prepared it for the project.

The capability to run a high number of operating system instances inside of virtual machines on a high performance computing (HPC) cluster can also be used to model even larger HPC machines with millions to tens of millions of nodes that will be developed in the future, said Minnich. The successful Sandia demonstration, he asserts, means that development of operating systems, configuration and management tools, and even software for scientific computation can begin now before the hardware technology to build such machines is mature.

“Development of this software will take years, and the scientific

community cannot afford to wait to begin the process until the hardware is ready,” said Minnich. “Urgent problems such as modeling climate change, developing new medicines, and research into more efficient production of energy demand ever-increasing computational resources. Furthermore, virtualization will play an increasingly important role in the deployment of large-scale systems, enabling multiple operating systems on a single platform and application-specific operating systems.”

Sandia’s researchers plan to take their newfound capability to the next level.

“It has been estimated that we will need 100 million CPUs (central processing units) by 2018 in order to build a computer that will run at the speeds we want,” said Minnich. “This approach we’ve demonstrated is a good way to get us started on finding ways to program a machine with that many CPUs.” Continued research, he said, will help [computer scientists](#) to come up with ways to manage and control such vast quantities, “so that when we have a computer with 100 million CPUs we can actually use it.”

Provided by Sandia National Laboratories ([news](#) : [web](#))

Citation: Computer scientists successfully boot one million Linux kernels as virtual machines (2009, September 25) retrieved 20 March 2024 from <https://phys.org/news/2009-09-scientists-successfully-boot-million-linux.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>
--