SAND2012-3802P

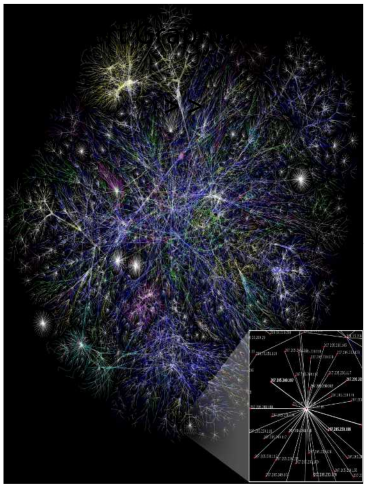

Large-Scale Network Emulation

Tom Kroeger
Sandia National Laboratories

Executive Summary: This technology uses lots of lightweight virtual machines (VMs) to create a virtual testbed to emulate a large network (such as a portion of the Internet). A computer cluster hosts the VMs, and up to 200 Windows or 1000 Linux or Android VMs can be run per physical processor, so that on a modest size computing cluster (with 1000+ processors) over 200,000 Windows computers can be emulated. This technology is useful for testing the effects of changes to security policies on enterprise networks, evaluating new security products in a contained environment, studying malware behavior and spread at large scale, or other network security experiments.

I. *Describe the problem that was the impetus for this work.* This technology addresses the issue of how to evaluate the behavior of large networks under stress or attack, as in Estonia.

   a. *Why is this problem important?* Understanding how networks will behave in the real world when they are attacked is important because computer networks are essential to the ability of government and critical infrastructure providers to fulfill their missions and provide essential services.

   b. *Who is affected by this problem?* This problem affects nearly every federal, state, and local government agencies, as well as critical infrastructure owners/operators.

II. *Describe the technology that you are researching.* This technology combines lightweight virtual machine technology with scalable cluster management software to enable emulation of large networks. A topology that represents a real network can be imposed on the virtual machines, which represent routers and hosts on the network, and experiments can be run on the virtual network.

   a. *How does your research address the problem?* Because the virtual machines run real operating system and application software, realistic modeling of attacks on a network can be done at scale. This allows countermeasures to be developed and tested for efficacy and safety before being deployed on real networks.

   b. *How does your approach differ from products currently on the market or other research?* The technology described here can achieve 100-1000 fold increases in size over existing network research testbeds such as Emulab or DETER.

   c. *Is this a standalone technology or is it a suite of solutions?* The entire system is standalone. However, there are components developed for the overall system that may have usefulness individually, and tools/software developed by others may be useful in conjunction with this system.

I. *Assess the research's technical readiness level.*

a. *Has the technology been thoroughly tested?* No

i. *If so, in an operational or research setting?* Has been tested in a research setting only.

b. *Is the technology currently being used?* The technology is being used to support other research, into file storage systems for example.

i. *If so, in an operational or research setting? By an internal or external organization?* Yes, being used to support other research by both internal and external research groups.

ii. *If not, when might it be ready for use, and how much funding over what timeframe would be needed if additional funds are required?* Development for use in doing network security experiments will require further maturation, depending on the features required. Areas that need more work are data collection/analysis and pause/restart. Approximately one more year and $1 million would enable maturation to where the technology would be useful for supporting network security experiments.

c. *What infrastructure must be in place to deploy the proposed technology?* A computing cluster using particular Intel or AMD processors and motherboards is required.

d. *What is the on-going overhead involved with operating and maintaining the technology?* At least one full-time system administrator is required.

i. *Will operators be required to learn new skill sets?* Yes.

ii. *Beyond operators, are there other support requirements?* Power, space, and cooling are required for the cluster.

iii. *Are these assessments based on operational experience with the technology, or are they estimates/expectations?* Based on operational experience.

II. *Has anyone expressed interest in transitioning this research to practice?* Yes, some other government agencies have expressed interest in adapting the system for particular needs.

a. *If so, who and when? What is the current status?* Various parts of DOD. We are engaged in discussions about potential work to mature the technology.

III. Intellectual Property (IP) Interests

a. *What is the funding history, and implications for government or private use?* Internal SNL funding (LDRD) has been used to develop the system.

b. *What are your organization's IP interests relative to this technology?* We are in the process of asserting copyright in the software. We would like to release it as open source software.

c. *Are there other collaborators?* Not yet.

i. *If so, who are they?* Prospective collaboration with UC Santa Cruz and Santa Clara University.

d. *Who owns the IP?* DOE

i. *Does anyone else have a claim to the IP?* No

e. *Are there any patents or patents pending for this technology?* No

i. *Does this approach require licensing any foundational IP?* No

f. *Are there processes in place for transitioning IP for commercialization?* No

g. *Are there any IP issues that would restrain open sourcing the technology?* No

U.S. DEPARTMENT OF ENERGY NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Sandia National Laboratories