

# The DETER Testbed: Overview

25 August 2004

## 1. INTRODUCTION

The DETER (Cyber Defense Technology Experimental Research testbed) is a computer facility to support experiments in a broad range of cyber-security research projects, including those experiments that involve "risky" code.

*Note: The name "DETER" is a play on words; please, always stress the second syllable!*

The DETER testbed is designed to provide an experimental environment in which government, academic, and industry cyber-security researchers can safely analyze and measure attacks and develop attack mitigation and confinement strategies. In addition, the DETER project will provide tools and resources to enable repeatable experiment methodologies, allowing different researchers to duplicate and analyze the same experiments.

This note provides a brief overview of the capabilities, architecture, and usage of the DETER testbed.

## 2. BACKGROUND

Development and operation of the DETER testbed is being performed by the DETER project funded by the National Science Foundation (NSF) and the Department of Homeland Security (DHS). Participating organizations are USC/ISI, UC Berkeley, and McAfee Research (formerly NAI).

DETER is constructed using the *cluster testbed* technology developed by the University of Utah and known as "Emulab" – see <http://www.emulab.net/>. Much of the online documentation for DETER is taken from Emulab, since much of the control and administrative software is the same. However, there are some differences between DETER and Emulab, primarily to assure greater safety for malevolent code in DETER. For example, a DETER experiment does not have a direct IP path to the Internet, unlike an Emulab experiment.

There will be no charge for the use of the DETER testbed. Acceptable use policies are approved by the sponsoring agencies.

The DETER testbed is targeted, at least initially, at support for open and publishable research projects, typically academic research. It does not currently have the kinds of privacy protection that might be required for use in testing vendor products, for example. We believe that the testbed technology that is being developed in DETER could easily be cloned and extended to support a product testing laboratory.

An initial version of the DETER testbed has been in operation since March 2004. DETER has been used by three research teams under another NSF/DHS-sponsored project, EMIST (Evaluation Methods for Internet Security Technology). EMIST teams performed experiments on DDoS attacks, worm propagation, and BGP attacks using the initial testbed.

During the first 116 days of availability of the initial DETER cluster, a total of 28 users had run 3284 experimental sessions ("swap ins") on DETER, using an aggregate total of 26,062 experimental nodes over 723 experiment days.

We plan to gradually extend access to the testbed beyond the EMIST project teams to a broader community of academic, government, and industrial researchers.

### **3. OVERVIEW OF THE TESTBED FACILITY**

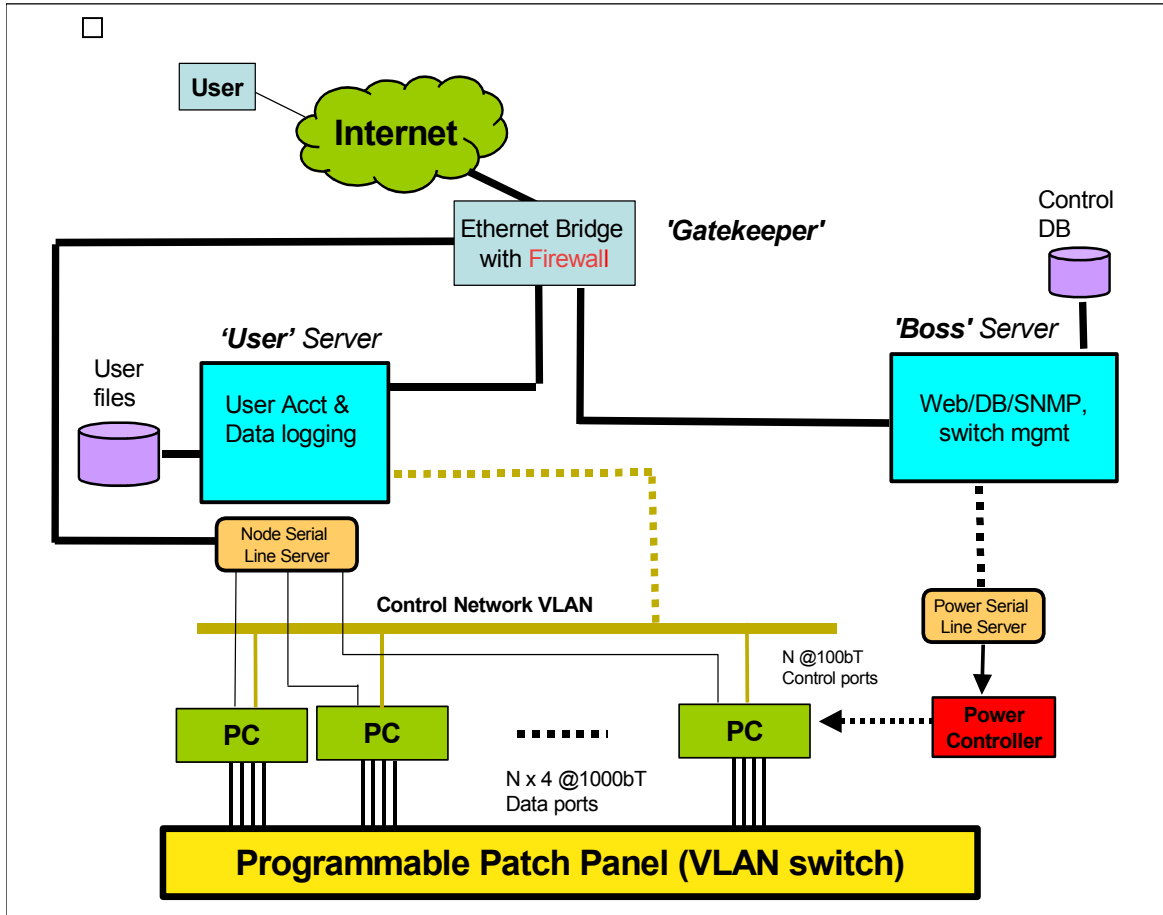
The DETER testbed is constructed using the cluster testbed technology developed by the University of Utah and known as "Emulab" (see <http://www.emulab.net/>). A *cluster* is a set of experimental nodes in a single room, interconnected through a programmable "backplane" or "patch panel", which can dynamically establish a distinct network topology for each experiment. In the DETER testbed, each node is a PC machine with significant disk storage, at least 2MB of memory, and four 10/100/1000bT Ethernet interfaces to the programmable patch panel.

Under the control of an experimental script, the Emulab control software automatically allocates a set of free nodes and sets up the interconnections to allow a particular experiment to proceed. The control software automatically loads kernels and other software into the nodes and configures the VLAN switch, firewalls, NFS mount points, and other system parameters to isolate it from any other experiments that may be running simultaneously. Control of the PCs is then turned over to the experimenter.

The Emulab control software used by DETER supports space-sharing on a node-level granularity (i.e., an individual node is the smallest unit of allocation), and experiments can be swapped out for large-grain time-sharing (i.e., after an experiment completes, it can be swapped out and back in at a later time).

During the experiment, the user can employ the control plane to monitor the nodes, to reload crashed nodes, or to swap out the entire experiment, using a web GUI on the Boss server (see Figure 1). When all else fails, the user can deliberately power-cycle a node. The user has access to nodes through their serial consoles, and through their control ports

via the User server. On the other hand, a running experiment in DETER will not have direct IP connectivity to the Internet outside their testbed. It is also important to note that experimenters are not able to change the configuration of the VLAN switch while the experiment runs.



**Figure 1: Schematic of DETER Testbed**

Figure 1 shows a schematic of the DETER testbed. This diagram omits the detailed configuration of the DETER control plane, since this should not be relevant to most users. The functional paths are shown as dotted lines.

The DETER testbed includes two clusters, one at USC ISI (ISI) and one at UC Berkeley (UCB). Within each cluster, there are one or more sets of identical nodes. The script for an experiment may ask for allocation of nodes from any single set or from a selection of these sets of nodes. The ISI/UCB link is high speed and should be largely, but not entirely, free of interference from other experiments or other Internet traffic. On the other hand, nodes from within a given set are always locally connected to the same

programmable backplane switch, whose very high aggregate bandwidth should guarantee no interference.

The N nodes and the VLAN switch that are shown in Fig. 1 are physically split between the ISI and the UCB clusters. Except for the performance variation imposed by the ISI/UCB link, the different physical location should be transparent to DETER users. ISI is the primary entry point for users in every case. The VLAN switch and the Control Network VLAN are bridged between the two sites, using IPsec for security.

### 3.1 Experimental Nodes

At present the experimental nodes in the DETER testbed are all PC machines, with a high degree of homogeneity. There are plans to add a few commercial routers to the testbed during the coming year.

The nodes fall into three homogeneous pools of identical PC hardware configurations. The machines used in a particular experiment can be allocated from a particular pool or from a combination of pools, as specified by the experimenter.

Each PC node in an experiment may run FreeBSD, Linux, or a (nearly-) arbitrary OS of the experimenter's choice. As in Emulab, the experimenter may have root access to each assigned node. PCs in a configuration can be configured as routers, as end systems, as traffic generators, or as traffic policers (e.g., using the 'dummynet' mechanism) to emulate arbitrary link characteristics.

The DETER testbed includes two clusters of nodes. There is an operational cluster of 72 nodes at ISI. A second cluster of 32 nodes at UC Berkeley is expected to be operational August.

There are currently three pools of nodes in DETER:

- 64 IBM Netfinity 4500R (Dual Pentium III 733MHz CPUs with 1GB RAM and 17GB SCSI disk storage), at USC/ISI.
- 8 Sun Microsystems Sun Fire V65X nodes (dual Pentium 2.8 Ghz Xeon CPUs with 2GB RAM and 216GB disk storage) nodes, at USC/ISI.
- 32 Sun Microsystems Sun Fire v60x (dual Pentium 3.06 Ghz CPUs with 2GB RAM and 72GB disk storage) nodes, at UC Berkeley (to become operational in October 2004).

### 3.2 Inter-node Links

Network connections between experimental nodes are created by a "programmable backplane" consisting of a large, high-speed Ethernet switch. Each experimental node

has four 10/100/1000Mbps Ethernet ports connected to the local high-speed Ethernet switch using VLANs. VLANs are used to create the desired experimental topology for each concurrent experiment.

In addition to the four backplane ports, each node has a fifth port, of at least 100Mbps, for downloading and controlling the experiment. The experimenter is also given access to remote power-cycling control and to the serial console interface of each node. However, remote access to a running experiment may be limited or even disallowed if the experiment poses a risk of spreading outside the designated experimental nodes.

Different Ethernet switch hardware is used at each of the ISI and UCB sites. Specifically, the ISI cluster uses a Cisco 6509 switch, while the UCB cluster uses a Foundry Fast Iron 1500 switch. Locally, each switch is sufficiently over-provisioned to avoid any interference between different inter-node logical links. Thus, the nodes within a single site are interconnected over high-speed LAN connections whose performance is constant and repeatable. On the other hand, the ISI and Berkeley pools are interconnected over high-speed Internet links whose performance is necessarily somewhat variable. An experiment can deliberately incorporate, allow but ignore, or avoid this inter-node variability by choosing which pool(s) to use for allocating nodes.

### 3.3 Future Plans

There are plans to extend the DETER facility significantly during the funding year beginning in October 2004. Current funding will support the addition of roughly 60 experimental nodes, in the near future. Other improvements are expected to involve (1) addition of (a few) commercial routers, (2) control software to support experiments using risky code, and (3) monitors to detect security or containment breaches.

## 4. Experiment Access

To define and set up an experiment on the DETER testbed, a user (i.e., an experimenter) defines a configuration using a scripting language that is derived from the "ns" simulation language. In accordance with this script, the testbed control plane allocates nodes, loads the specified disk images, and starts the experiment.

Each user is given an account on a DETER server machine called "Users" (or "Ops"). This account may be used for staging input to an experiment and output from it, as well as experiment-related processing. Each experiment has its own file system that can be mounted from the user's experimental nodes.

The Boss server (Figure 1) runs the web GUI for [deterlab.net](http://deterlab.net), used for registering users and projects and for initiating and controlling experiments. Each registered user has an account of the Users machine, and the file system under that account can be NFS-mounted from experimental nodes.

A user does not have direct IP connectivity to an experimental node, but the user will be able to SSH from his/her account on the "Users" machine into any of the allocated experimental nodes. The user can also communicate with the experiment through shared files on the user's file system on the User node. If the node OS crashes, the user can reload the OS using the web interface. OS reload includes power cycling the node. The user can also employ serial console access to each node when severe containment requirements preclude even SSH or NFS access, so that the experiment must be run in an isolated manner.

Experimental GUIs specifically designed for defining, controlling, and analyzing major classes of security experiments on DETER are under development by the EMIST project.