

CENIC: High-speed Connectivity for the DETER Testbed

Terry V. Benzel

Bob Braden

**Information Sciences Institute
University of Southern California**

Anthony Joseph

University of California at Berkeley



UC DAVIS



Cyber Defense Technology Experimental Research (DETER) & Evaluation Methods for Internet Security Technology (EMIST)

- Inadequate wide scale deployment of security technologies
 - Despite 10+ years investment in network security research
- Lack of experimental infrastructure
 - Testing and validation in small to medium-scale private research labs
 - Missing objective test data, traffic and metrics

DETER + EMIST

- DETER & EMIST: companion projects funded by NSF and DHS HSARPA
 - 1) Design and construct a **testbed** for network security experiments
 - 2) Do research on **experimental methodology** for network security
 - 3) Do research on **network security**

Examples of Experiments

- DDoS attacks and defense
 - Understand dynamics, try various defense strategies
- Worms
 - Understand worm dynamics
 - Try various defense strategies
- Routing security
 - Understand BGP dynamics and threats
 - Build and test countermeasures
- Advance intrusion detection
- (and many more)

DETER “Requirements” (1)

- **Versatility:** Support wide range of security scenarios.
- **Repeatability:** Complete control of environment, for repeatable experiments without artifacts.
- **Containment:** Confine dangerous code.
- **Realism:** Real router and end-system behavior.
- **Fidelity:** Represent topology and traffic mix of Internet.
- **Programmability:** Add new algorithms to routers

DETER “Requirements” (2)

- **Accessibility:** Remote control over Internet.
- **Efficiency:** Testbed partitionable among simultaneous independent experiments
- **Functionality:** Rich set of traffic and topology generators and experimental profiles.
- **Economy:** Accomplish all this with very limited \$\$

Basic Design Choice

- **Cluster testbed**
 - Many nodes in one laboratory
 - Dedicated local inter-node links => Perfectly controlled
 - Prime example: Univ of Utah's Emulab
- **Distributed TB**
 - Nodes scattered across Internet
 - Prime example: Planetlab
 - Links subject to "normal" Internet interference

Basic Design Choice

- Two reasons to choose clusters for DETERlab
 1. Security & containment ...
 - would be impossible in distributed testbed.
 2. Experimental repeatability
- There is no perfect solution ...
 - Use Utah's Emulab software
- Objective: biggest scientific bang for the bucks

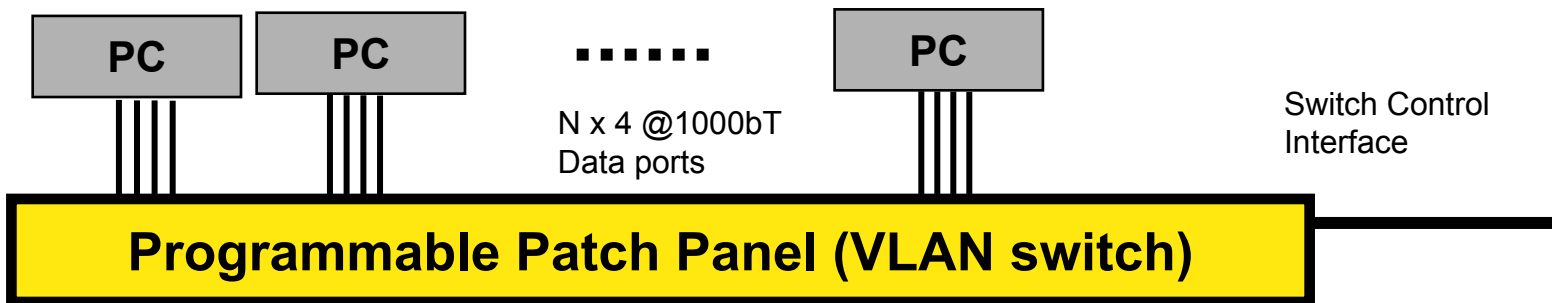
Deterlab Experimental Plane

- Basic experimental node in cluster:
high-end PC
- Each node may be configured to:
 - Emulate an end node or a router,
 - Generate traffic,
 - Emulate link characteristics, or
 - Make measurements
- Nodes identical in large groups
 - e.g., 32 or 64

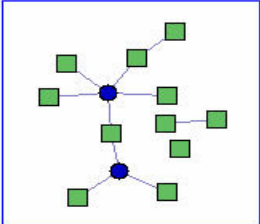
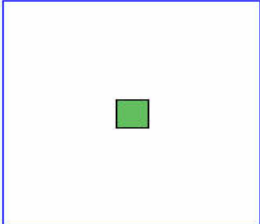
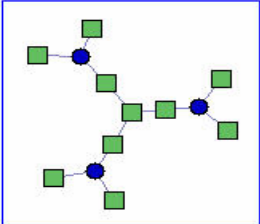
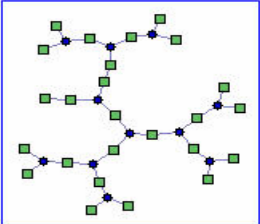
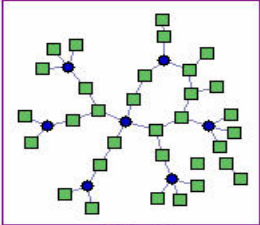
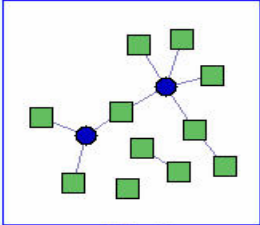
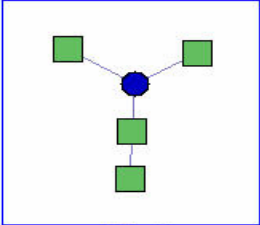
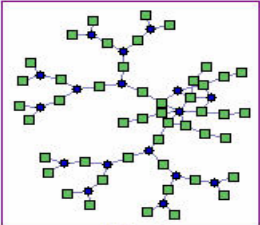
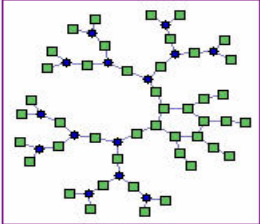
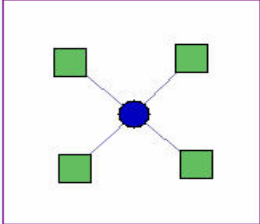
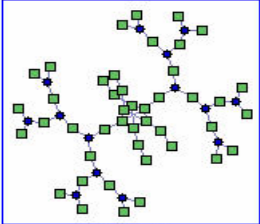
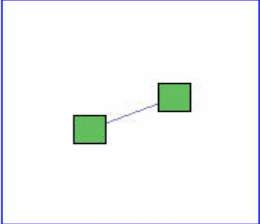
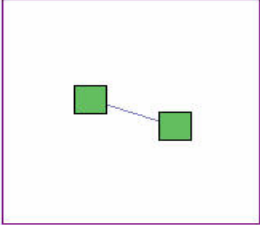
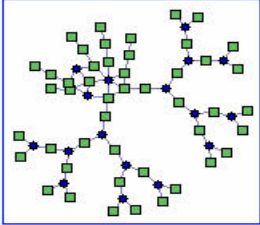
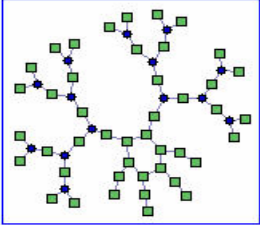
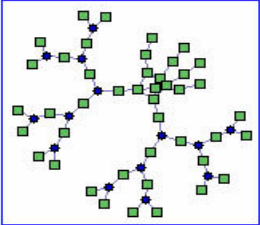
DETER Experimental Network

Cluster of N nearly identical experimental nodes, interconnected dynamically into arbitrary topologies using VLAN switch.

Pool of N processors



Example DETER Topologies

			
ddos/ smh	ddos/ DT	ddos/ cossack-3networks	ddos/ topology-20
			
ddos/ Demo	ddos/ Flood	ddos/ trace	ddos/ as11537-5s-2t
			
ddos/ ring-4s-2t	ddos/ ring-4s-2t-min	ddos/ semiclique-4s-2t	ddos/ tgcontrol
			
ddos/ tgcontrol2800	FloodWatch/ as11537-5s-2t	FloodWatch/ ring-4s-2t	FloodWatch/ semiclique-4s-2t

Experimental Plane

- Internode links are 10/100/1Gbps
- Working on integration of a few commercial routers into the cluster, to provide realism and increase heterogeneity
- Will also add special-purpose hardware, e.g.:
 - high-speed synthetic traffic generators
 - hardware instrumentation devices

The Fidelity Issue

Would Ideally Like:

- Large and realistic topologies
- Diverse, realistic nodes and links
- **But:**
 - Fidelity is expensive
 - Large-scale fidelity may be unnecessary for (maybe even contrary to) good science
 - Plan to add limited *heterogeneity* and *realism* – e.g., a few vendor routers, network processors

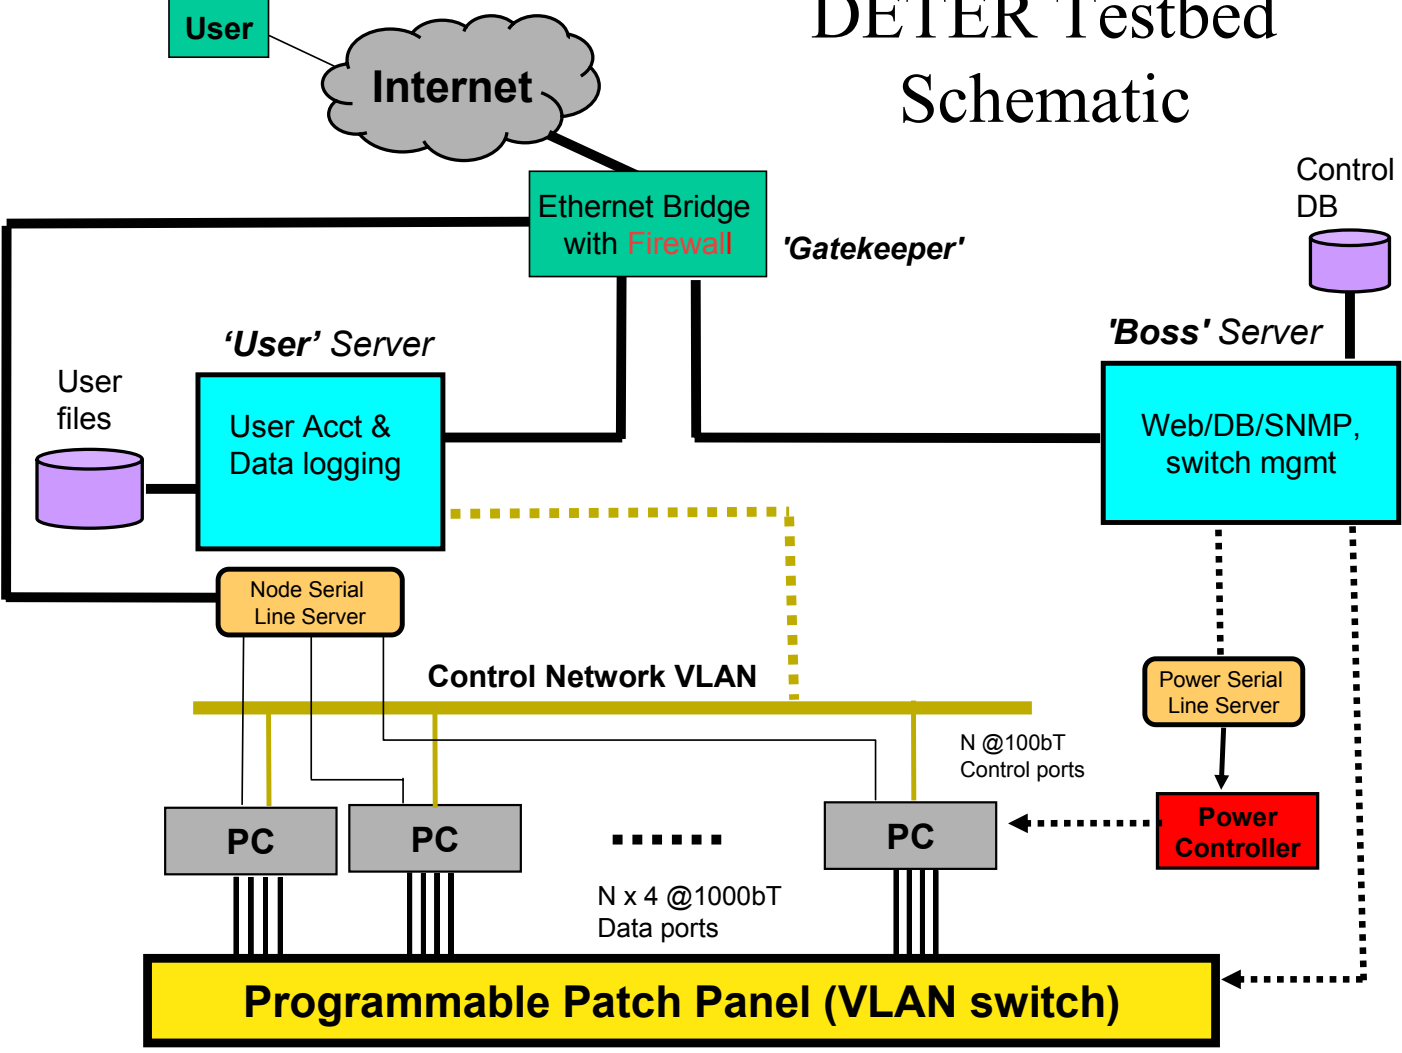
Experimental Backplane

- Switch hardware:
 - ISI: Cisco 6509, Nortel 5510 switches planned
 - UCB: Foundry 1500, Nortel 5510 switches planned
- Example:
 - ISI currently has 72 nodes, with 4 1000bT interfaces per node
 - All 10/100/1000 bT, VLAN'd
 - That is a LOT of wires...



Side view of the testbed.

DETER Testbed Schematic

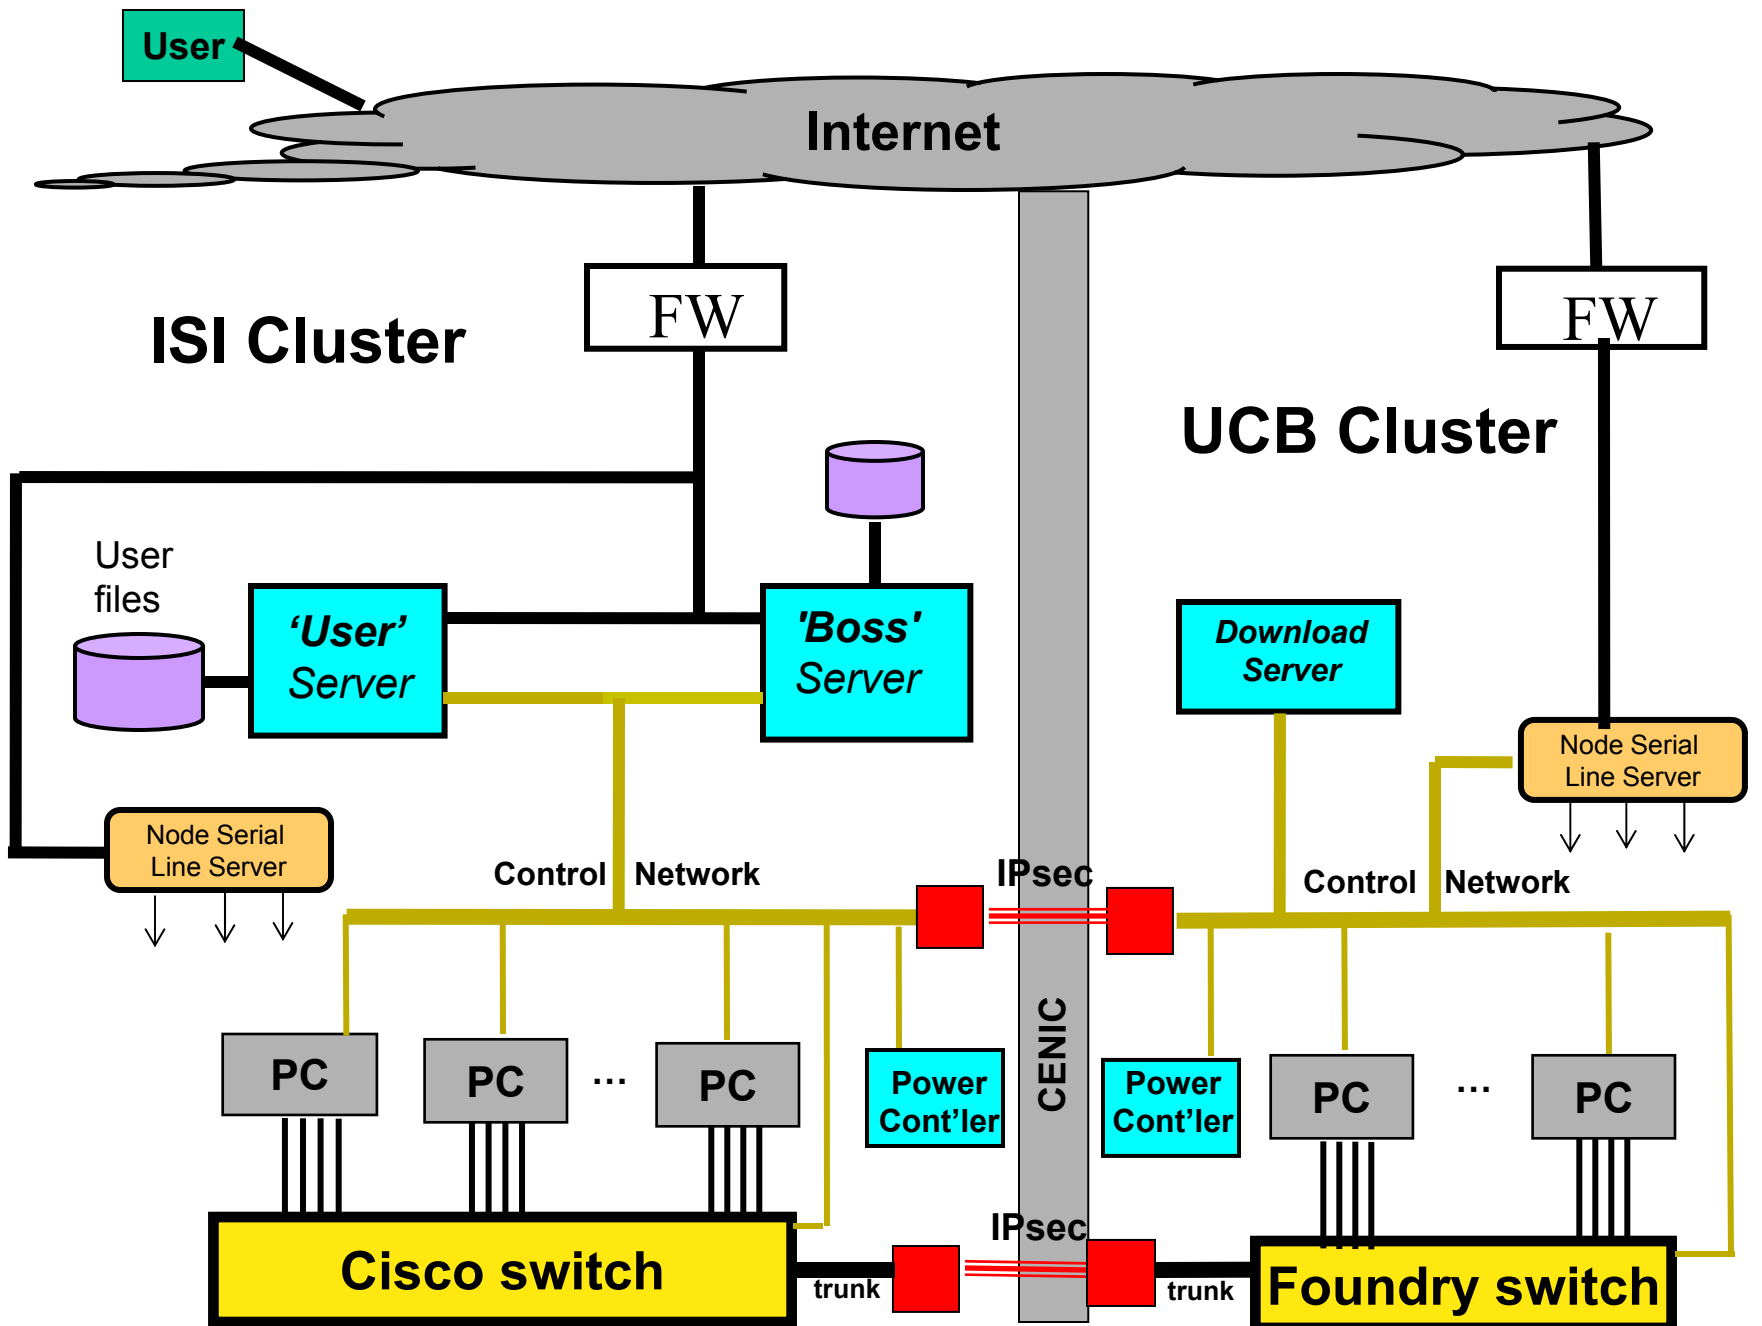


DETERlab Architecture

- Divide logical DETERlab cluster into two physical clusters
 - At USC/ISI and at UC Berkeley
 - One control plane, and one entry point (ISI)
 - "Centrally-controlled federation"
- Nodes from different clusters can be combined in one experiment when user chooses
 - When Internet introduces variability that will be desirable or at least tolerable

Interconnecting Clusters

- One control site (ISI)
 - One user entry point, accounts, control
- VLAN switches interconnected using IPsec tunnels
 - Distinct pools of nodes to be allocated
 - User can control whether span multiple pools
- IPsec tunnels should preserve security of link



ISI/UCB Links

- Two logical links:
 - Control plane link (layer-3 connection)
 - Experimental plane link (layer 2 connection – trunking ports between switches)
- For the experimental plane, this scheme demands a very high-speed link between Marina del Rey and Berkeley – at least 1 Gbps
- **CENIC fills that need**

Using CENIC

- Straightened out confusion about routing, so now have HPR routing
- No surprise: the last mile is the hardest part
 - Network people don't actually LIE to us, but...
 - It took awhile to track down which organization had a 100Mbps link in the path.
 - "Sure, you can have a 1Gbps link", but when we try to USE it, alarms go off.
- Traceroute shows 8 hops, half in CENIC

Using CENIC

- 1Gbps CENIC path is actually much smaller than aggregate possible inter-cluster traffic
 - But an experiment that pushes this limit is probably a bad experiment.
 - The Emulab control plane understands about the relatively limited inter-switch connectivity of 1 Gbps.
- Measured performance ISI \leftrightarrow UCB (w/o IPsec!):
 - 930 Mbps UDP
 - (TCP measurement in progress)

IPsec Performance -- Hard

- Not easy to get desired throughput using IPsec.
- First attempt: 100 Mbps, currently 200 Mbps.
- Currently 200 Mbps using crypto boards
 - So far, not living up to specs
- Trying many variations of hardware, software
- May have to settle for 250, use 4-times striping

Testbed Software

- Utah Emulab software in control plane.
- Experimental node OS:
 - Current standard OS: RedHat Linux 7.3 or FreeBSD 4.7
 - Moving soon to Red Hat 9.0, FreeBSD 4.10.
 - Users can load arbitrary code, in fact
- Users have *root* access to all allocated nodes!

Security and Containment

- Threats
 - To other experiments from inside – isolation failure
 - To DETER from inside – *intrusion*.
 - To DETER from outside
 - To outside from inside – ***extrusion!!***
 - Accidents...

Security is Critical

- Defenses employed by the DETER test-bed must balance the requirements of containment, isolation, and confidentiality with the need for remote management of experiments
- DETER experiments are categorized according to the consequences of loss of containment, and procedures applied according to that categorization

Achieving Security

- Operational
 - Procedures for proposing and reviewing experiments
 - Guidelines for categorizing safety of experiments
 - Vetting of investigators and experiments
 - Procedures used by investigators
- Technical
 - Firewall, routing, intrusion detection and network isolation techniques
 - Data protection, system protection, and state destruction techniques

Conclusions

- The DETER testbed provides a vital experimental environment for network security research, using Utah's Emulab software
- CENIC allows us to combine testbed facilities at ISI and UCB into one logical cluster testbed
 - User convenience: one access point, administration
 - Larger pool of experimental resources
- IPsec is required for security, but performance is a hard problem