Virtual Machines for ROC: Initial Impressions

Pete Broadwell
pbwell@cs.berkeley.edu

Talk Outline

1. Virtual Machines & ROC: Common Paths
2. Quick Review of VMware Terminology
3. Case Study: Using VMware for Fault Insertion
4. Future Directions

Background

- Virtual machine: an efficient, isolated duplicate of a real machine – Popek & Goldberg
- VMware: an x86-based virtual machine environment
  - Runs on PCs, workstations, servers
  - Supports Linux and Windows
  - Began as a research project at Stanford

ROC & Virtual Machines: A Perfect Match?

Recovery-Oriented Features of VMs

- VM "sandboxing" provides effective **isolation**.
- Multiple VMs on one machine yields **redundancy**.
- Suspend/resume capability means fast failover and **restartability**.
- Support for checkpointing, **undo**able sessions
- Significant support for **monitoring** and **diagnostics**
- **Online verification** of recovery mechanisms?

Type I VM: Stand-Alone

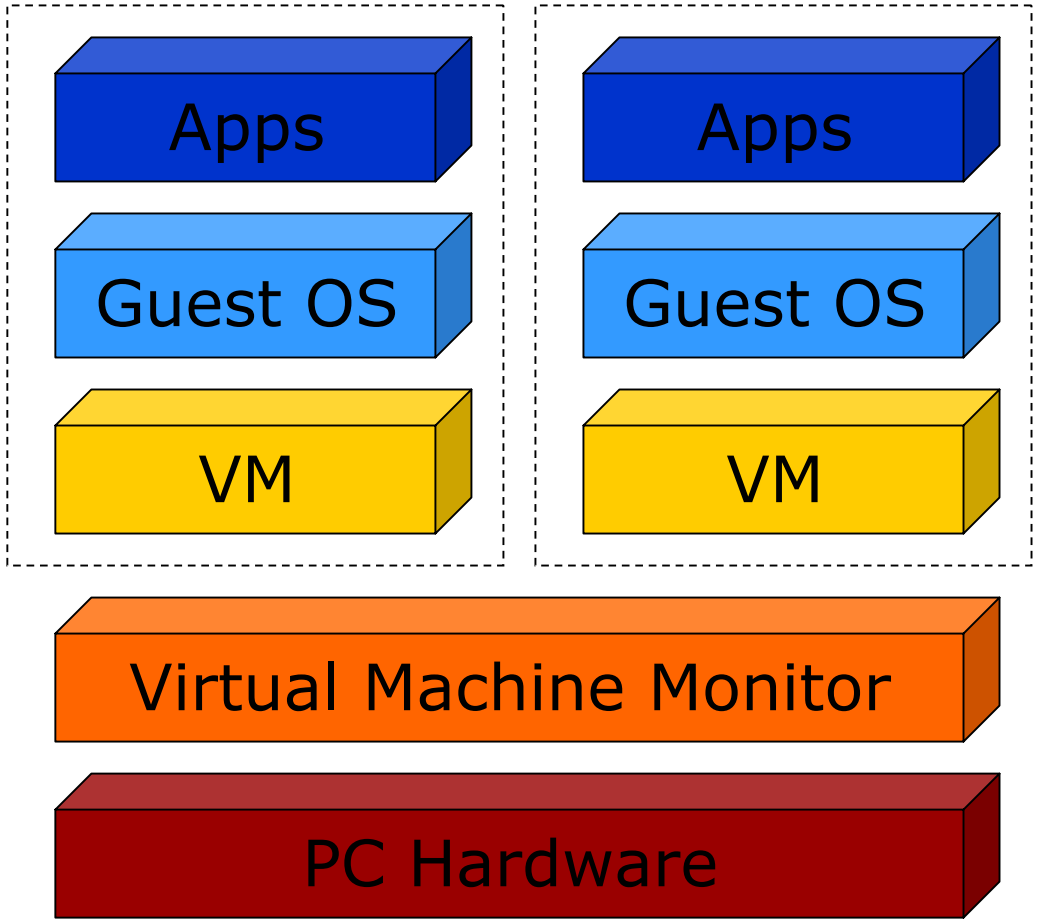

- Virtual machine monitor runs on bare hardware, supports multiple virtual machines.
- Examples: VMware ESX Server, IBM z/VM

Type II VM: Hosted

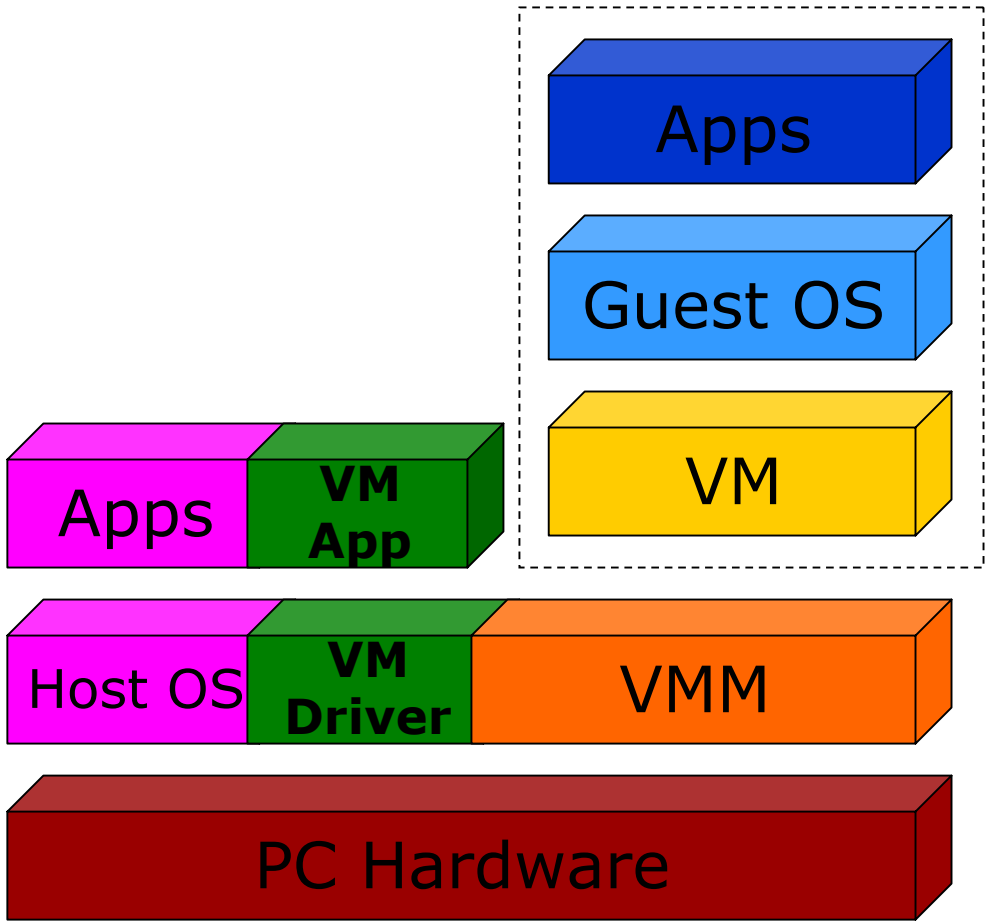

- VM app uses driver to load VMM at privileged level. VMM uses host OS I/O services through VM app.
- Examples: VMware Workstation, VMware GSX Server, Connectix Virtual PC, Plex86

Hosted VM I/O Virtualization

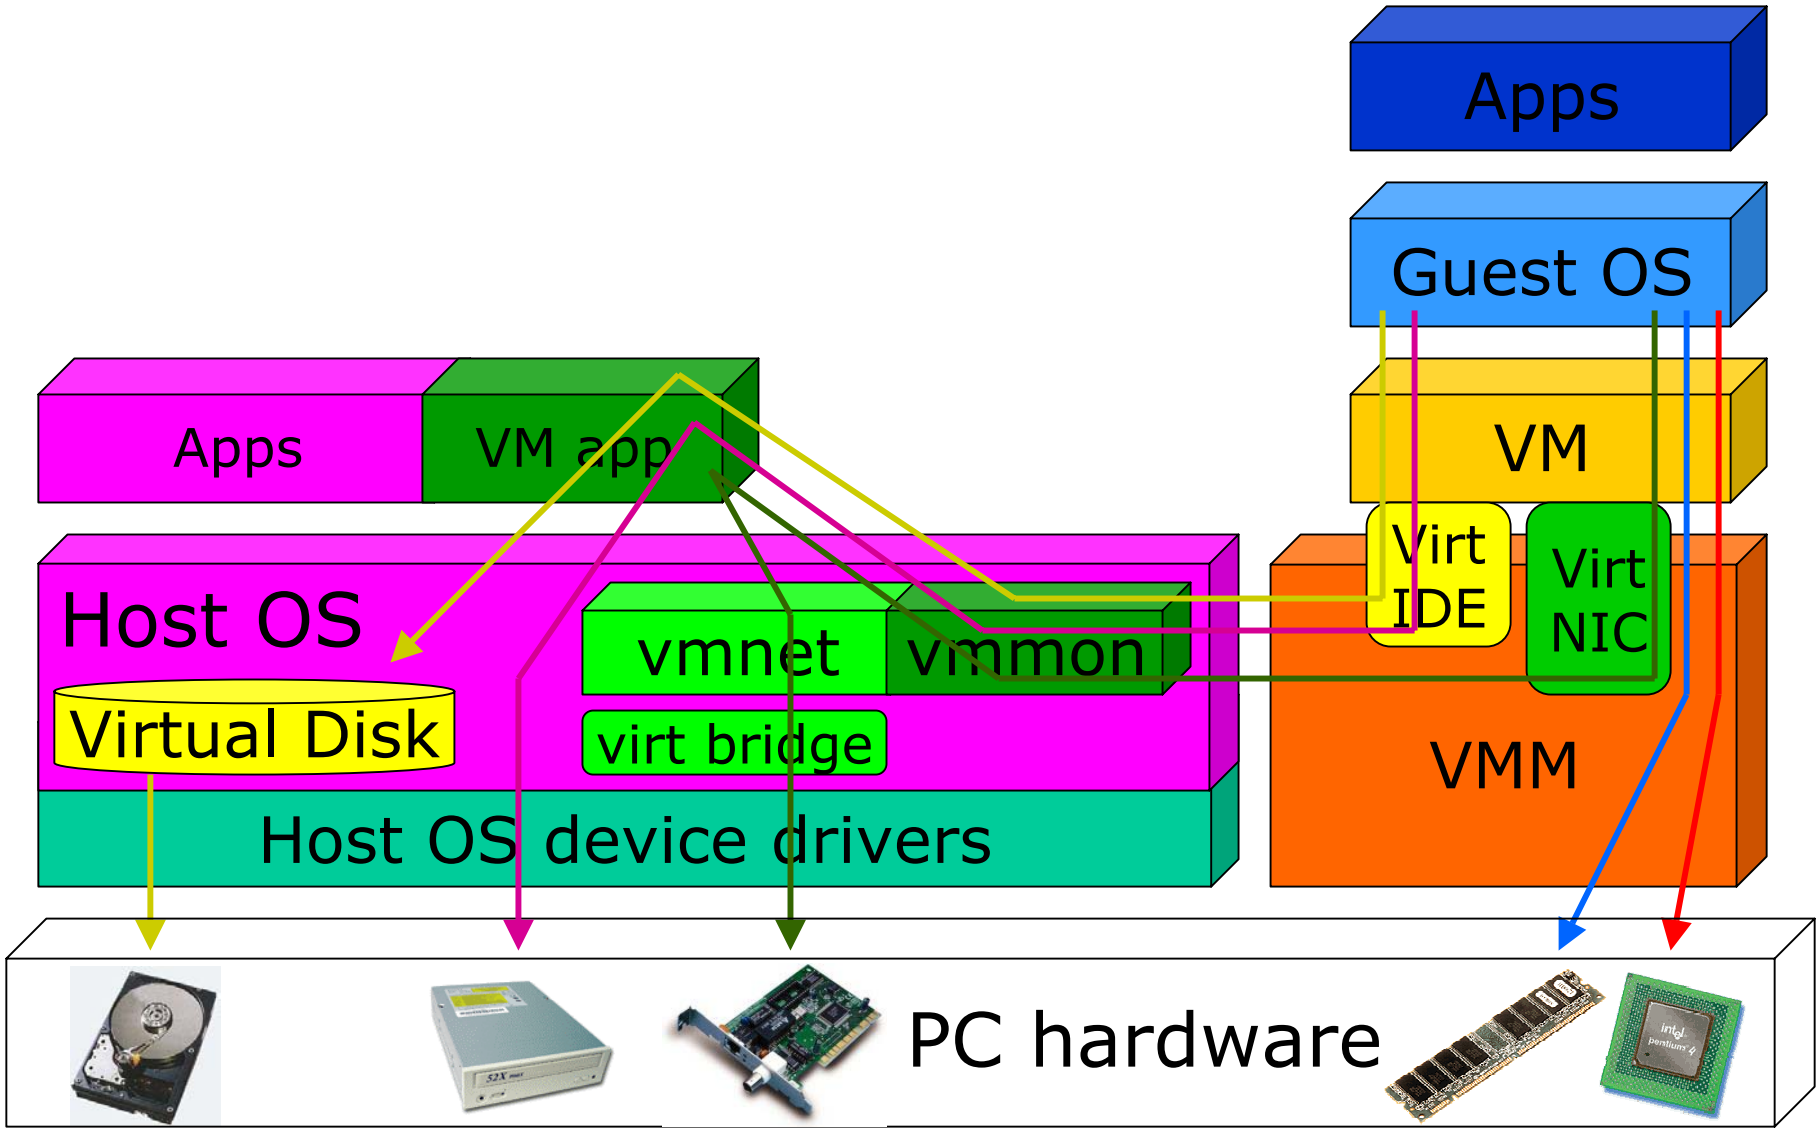

Case Study: Opportunities for Online Fault Injection in VMware GSX Server

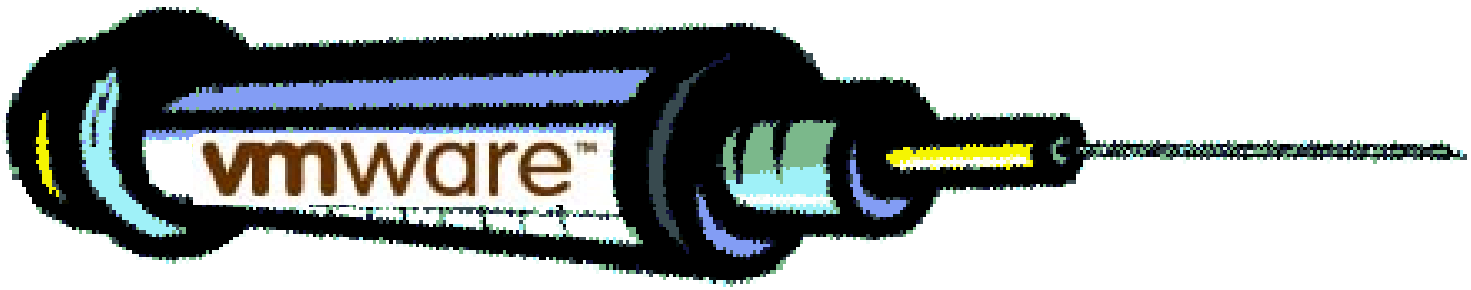

Why VMs for Fault Injection?

Fault injection is old news!

- ROC goals for fault injection:
  - Integrated with operating environment
  - Capable of injecting multiple types
  - Low overhead, high configurability
  - Able to expose latent errors in production systems

Which Faults are Important to Inject?

- Consider errors that have been observed on x86 PCs.

- Of these errors,
  - Which can be inserted using the existing capabilities of VMware?
  - Which require that VMware source code must be modified?
  - Which can’t be injected at all?

VMware does checking of its own!

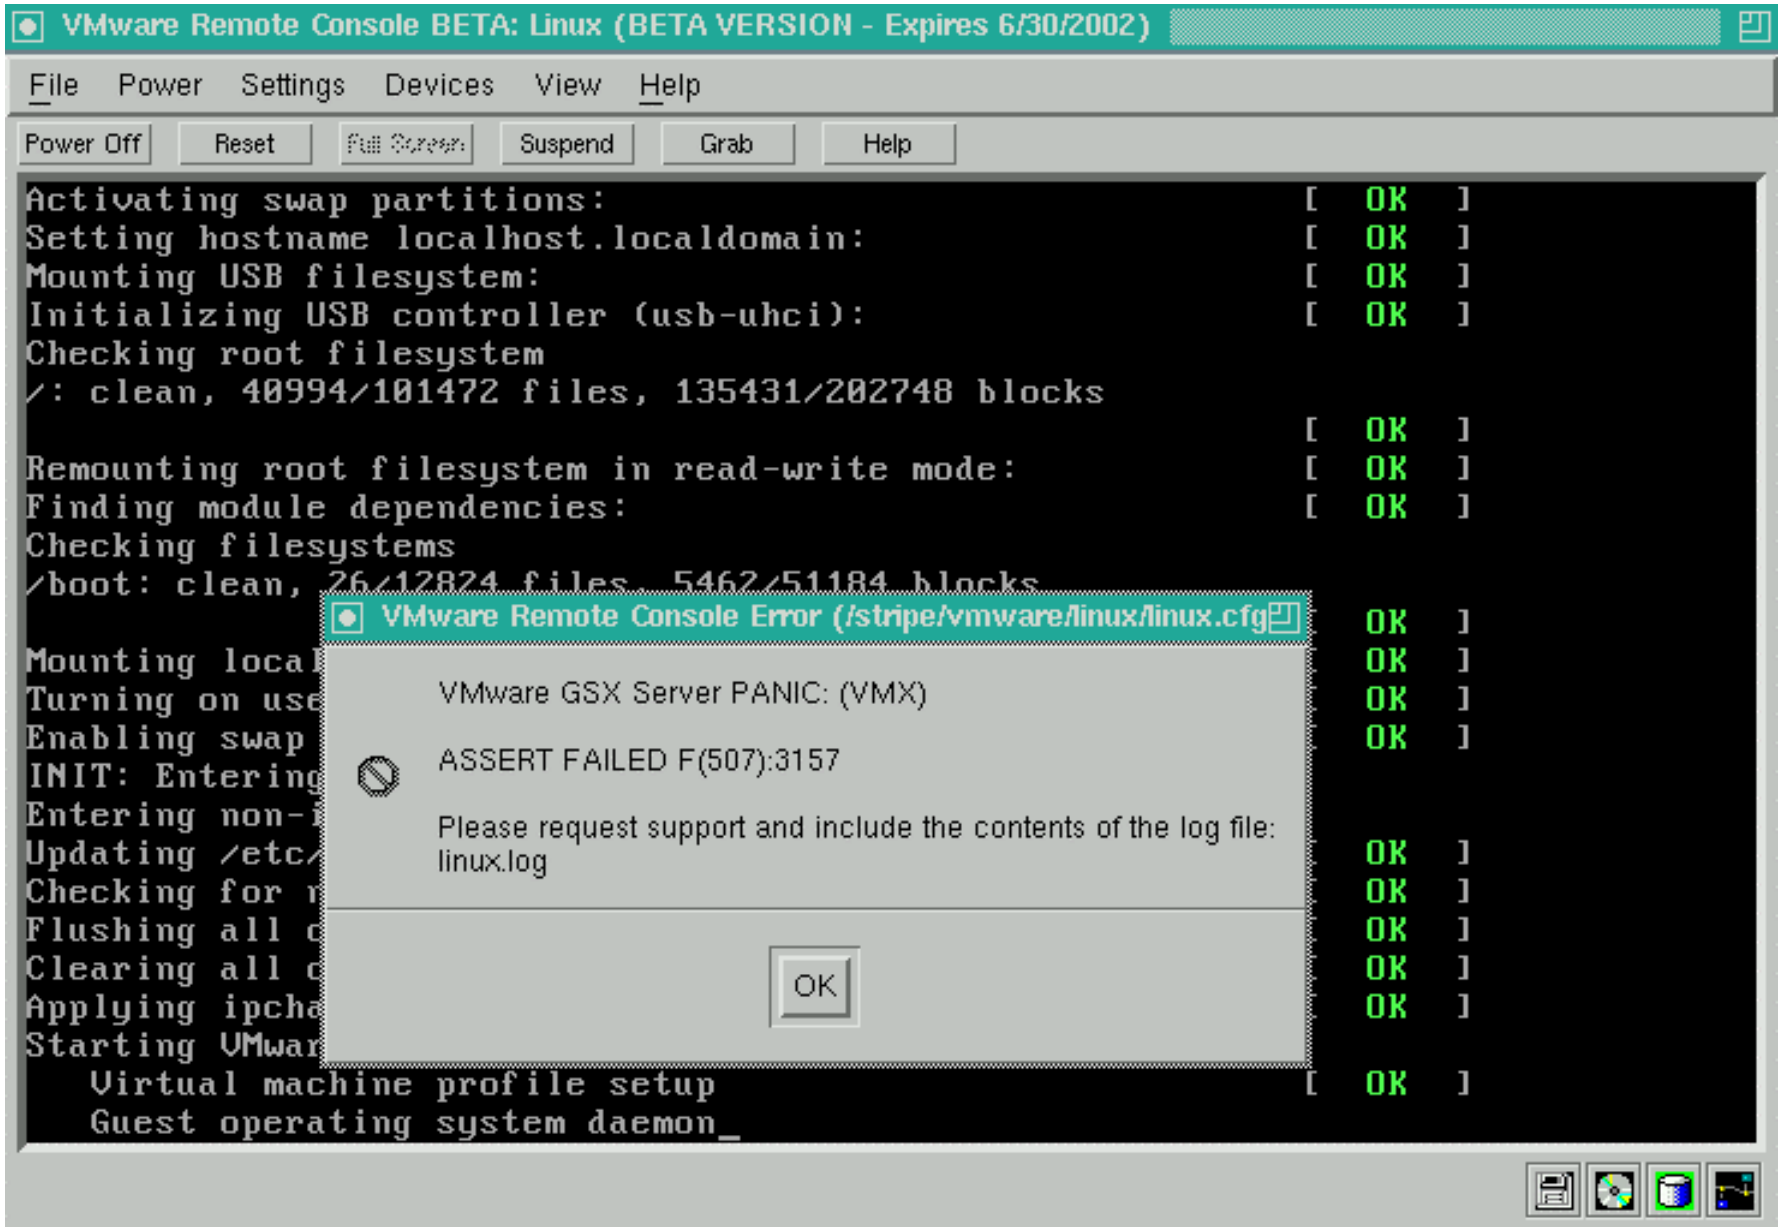

Memory/Processor Errors

- Want to simulate processor faults, memory ECC errors.
- Problem: in VMware, processor ops & memory accesses execute directly on hardware (not simulated).
- Need to allow VM to return "machine check" exception to guest OS.

Not difficult to guess what will happen:
kernel panic or blue screen.

Memory Corruption

- VMs use file system as backing for pinned memory pages – point for inserting corruption errors.
- VM driver (open source) interposes upon memory requests between VMs & host OS – can insert memory errors here.

Easy to do, but not very interesting or realistic.

Disk Fault Injection

- By default, a VM's virtual disk image is a flat file.
- Failures: catch read/write calls to the file, return errors indicating bad blocks, device failures to OS.
- Transient failures: overwrite random portions of disk image.

Should be relatively straightforward.

Network Device Faults

- VMware's virtual network module is open-source.
- Modify module, introduce failure code at virtual bridges and hubs
  - Drop packets
  - Corrupt packets
  - Simulate slowdown
  - Simulate DOS attacks

Virtual Hub: No Faults

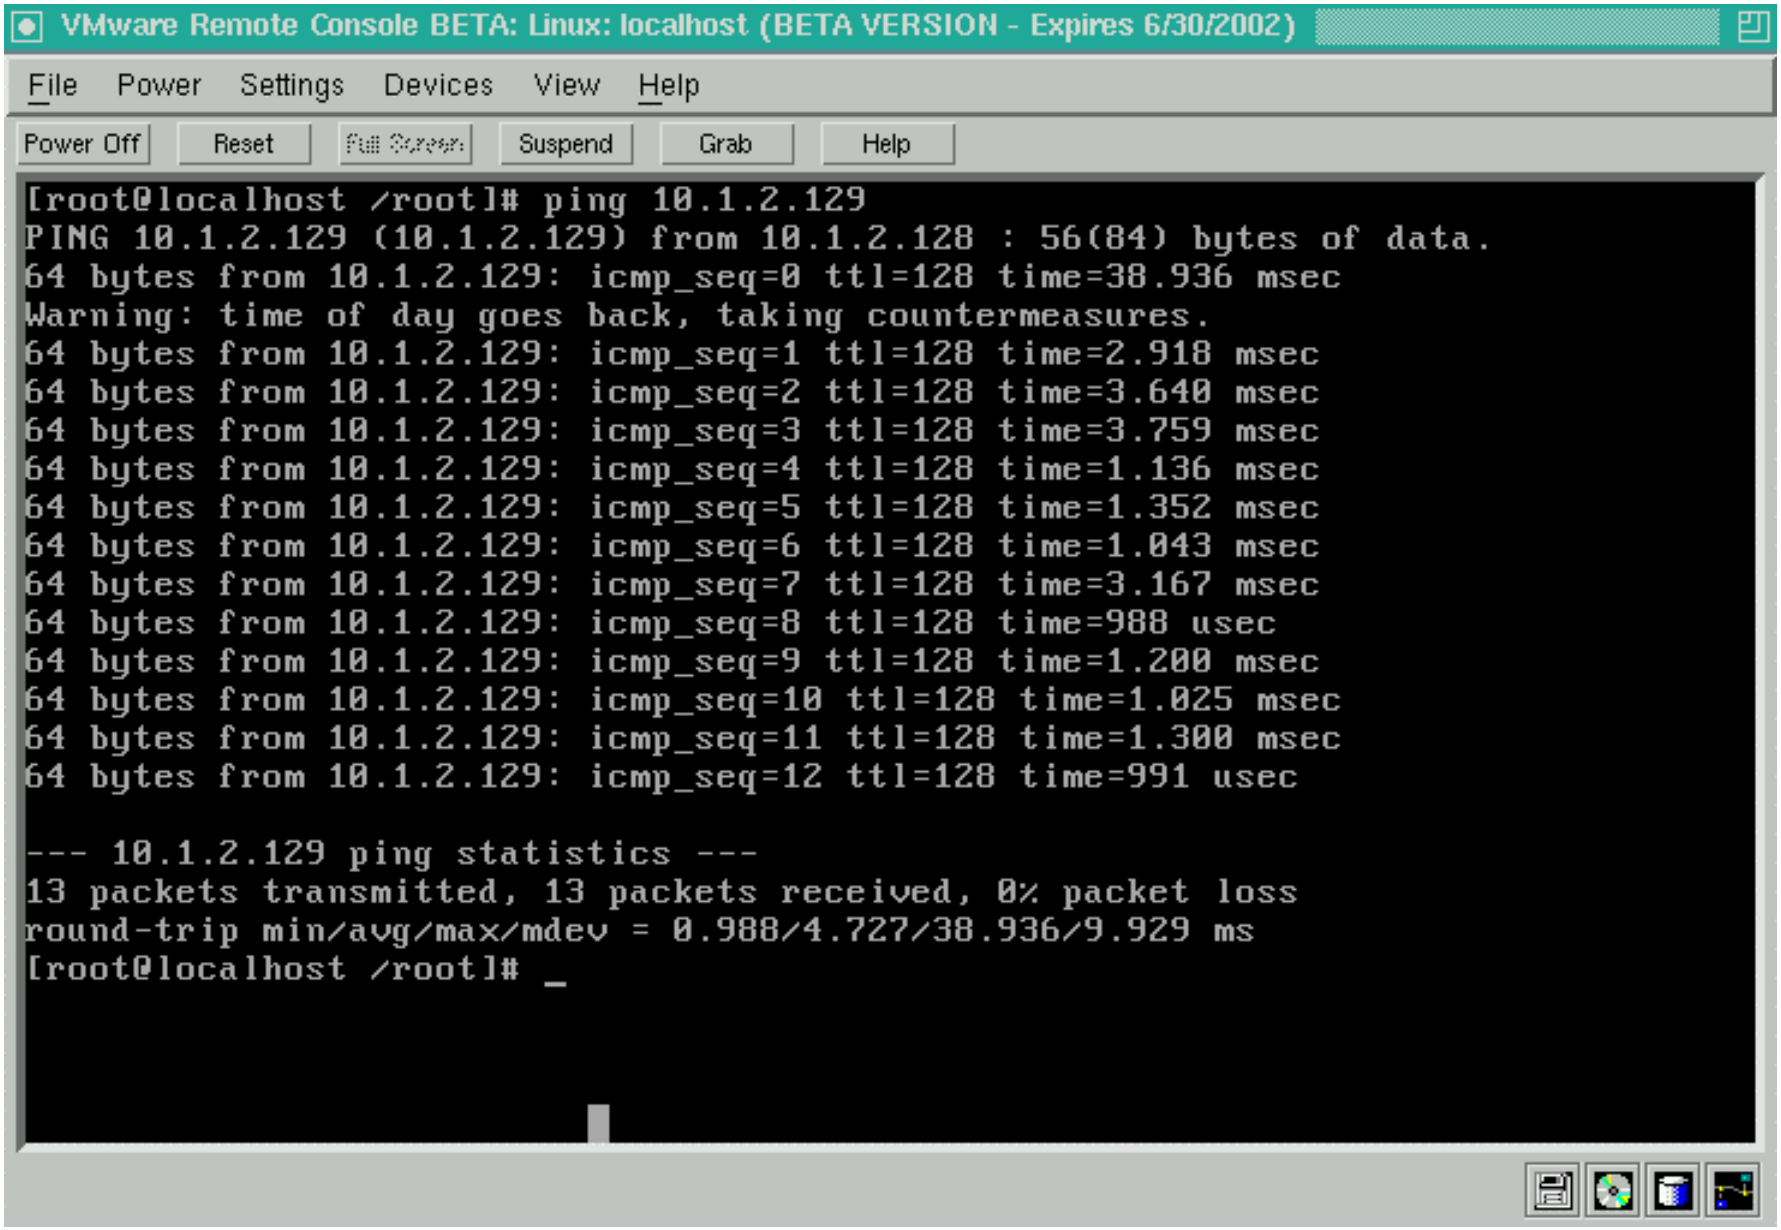

Virtual Hub: Injected Faults

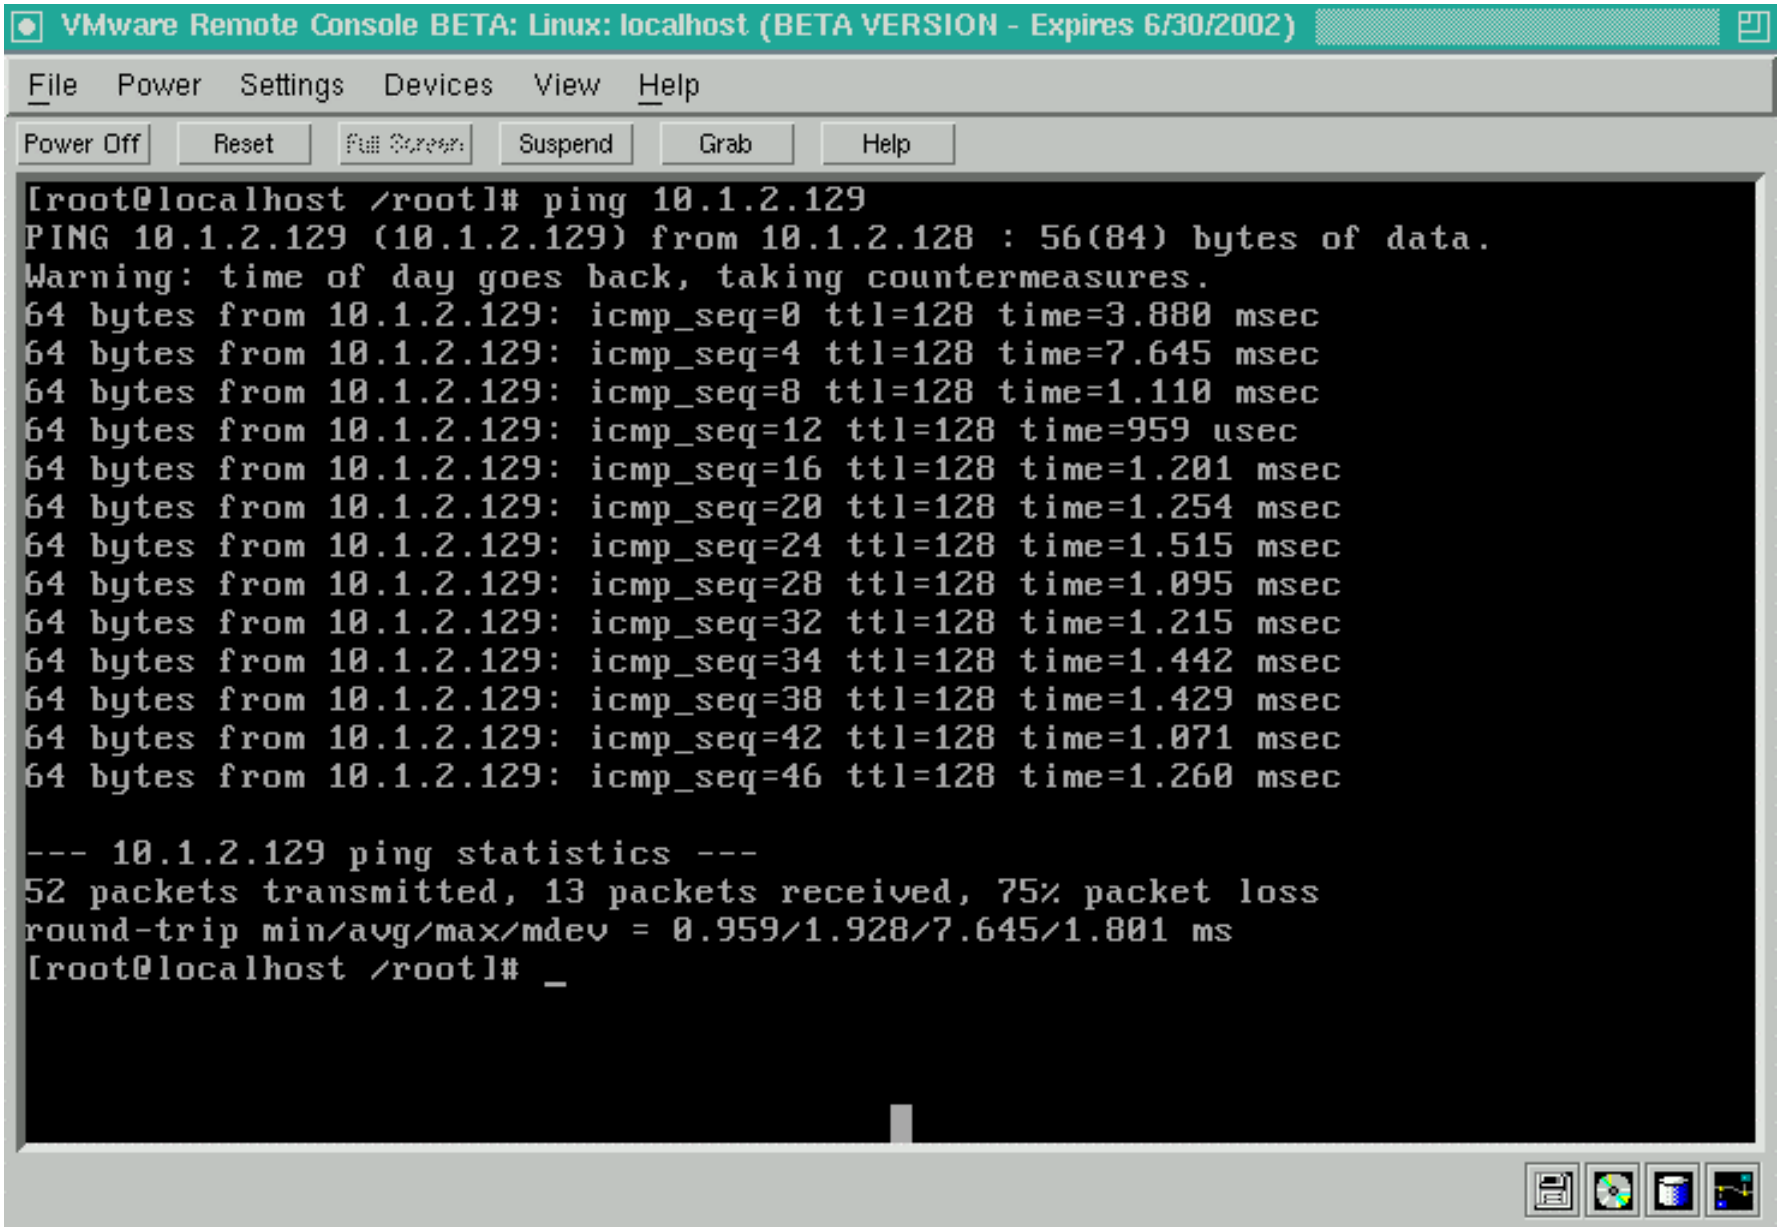

Cluster-Level Faults

- Use VMware's built-in remote management interface to hard-suspend nodes in a cluster, remove network bridges.
- Verify recovery/failover routines in cluster management software.
  - Dell Scalable Enterprise Computing
  - MS Cluster Server
  - NetWare Cluster Services
  - Microsoft SQL Server!

(Virtual) Cluster Management Interface

Analysis

- Levels of difficulty for different fault injection types:
  - CPU, cache, & memory (non-corruption) are hard to do.
  - Memory corruption, disk, NIC, peripherals may be medium.
  - Network, cluster level is easy.

The Big Picture

- Want to develop models for multiple correlated faults & implement them.
- Combine fault injection with introspection tools for anomaly detection & root-cause analysis.