Cyberinfrastructure for Distributed Rapid Response to National Emergencies

Henry Neeman, Director
Horst Severini, Associate Director
OU Supercomputing Center for Education & Research
University of Oklahoma
Condor Week 2006, University of Wisconsin

it@ou

SUPERCOMPUTING CENTER for EDUCATION & RESEARCH
OSCER

Disasters

http://abyss.ecs.umass.edu/tornado/may-3-99.html

Congressional Anthrax 2001

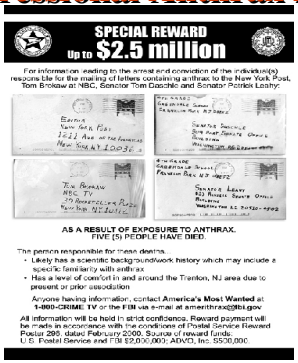

http://en.wikipedia.org/wiki/2001_anthrax_attacks

http://en.wikipedia.org/wiki/2004_Indian_Ocean_earthquake

OKC Wildfires Jan 2006

http://www.usatoday.com/weather/news/2006-01-01-grass-fires_x.htm?POE=WEAISVA

The Problem and the Solution

- **The Problem**: Problems will happen.
- The problem is that we don't know the problem.
- The solution is to be able to respond to unknown problems with unknown solutions.
- Unknown problems that have unknown solutions may require lots of resources.
- But, we don't want to buy resources just for the unknown solutions to the unknown problems – which might not even happen.
- **The Solution**: Be able to use existing resources for emergencies.

Who Knew?

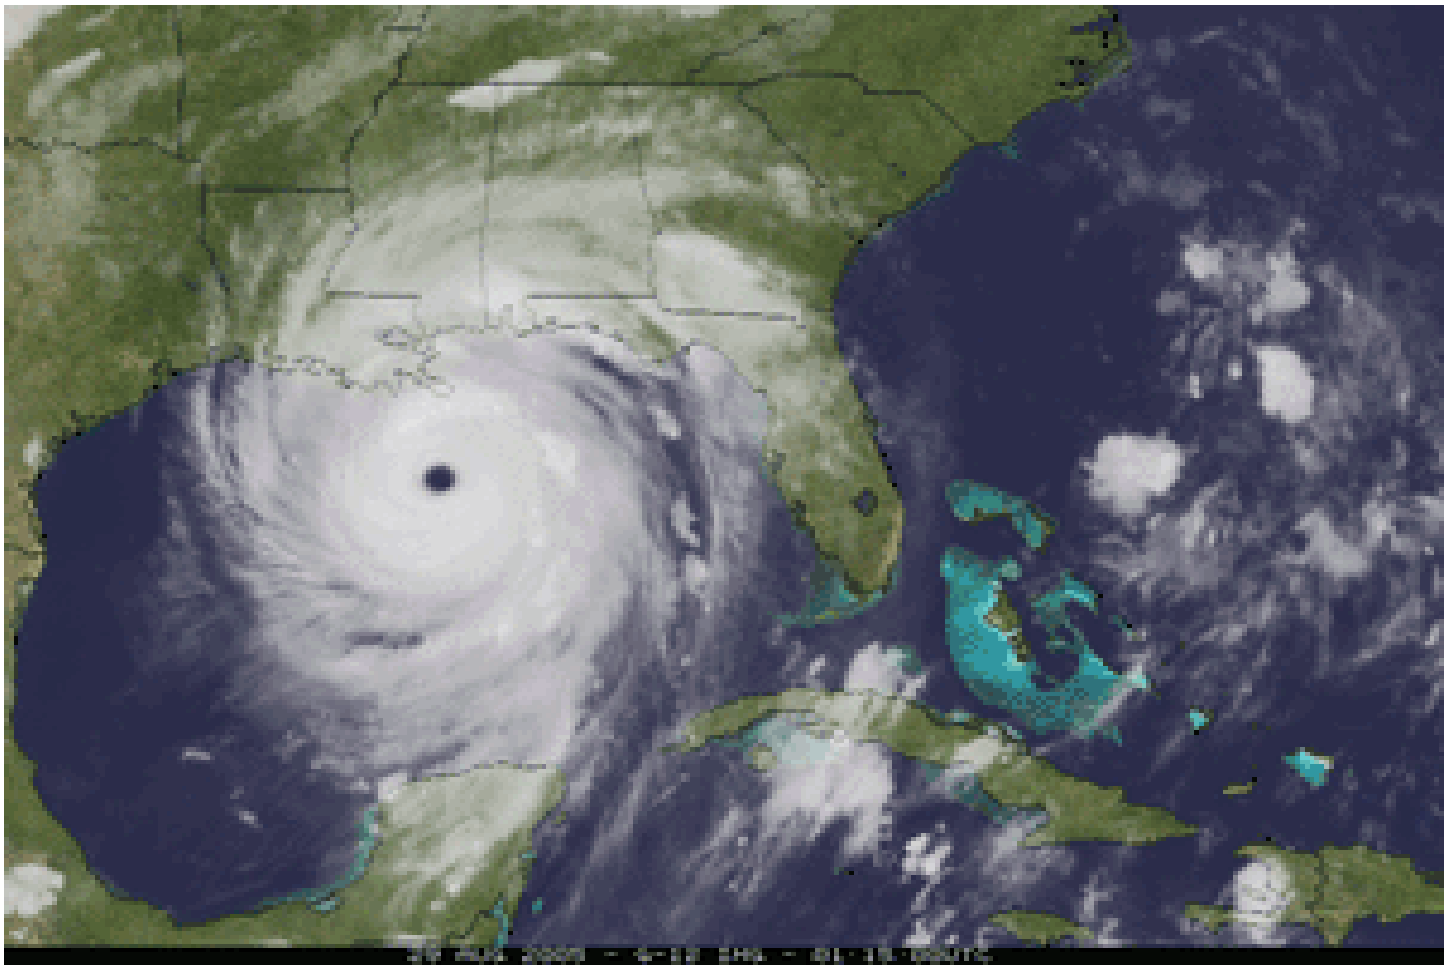

http://www.ncdc.noaa.gov/oa/climate/research/2005/katrina.html

National Emergencies

- **Natural**
  - Severe storms (e.g., hurricanes, tornadoes, floods)
  - Wildfires
  - Tsunamis
  - Earthquakes
  - Plagues (e.g., bird flu)
- **Intentional**
  - Dirty bombs
  - Bioweapons (e.g., anthrax in the mail)
  - Poisoning the water supply
  - (See Bruce Willis/Harrison Ford movies for more ideas.)

How to Handle a Disaster?

- **Prediction**
  - Forecast phenomenon's behavior, path, etc.
- **Amelioration**
  - Genetic analysis of biological agent (find cure)
  - Forecasting of contaminant spread (evacuate whom?)

OSCER's Project

NSF Small Grant for Exploratory Research (SGER)

- **Configure** machines for rapid switch to Condor
- **Maintain** resources in state of readiness
- **Train** operational personnel: maintain, react, analyze
- Fire **drills**
- Generate, conduct and analyze **scenarios** of possible incidents

@ OU: Available for Emergencies

- 512 node Xeon64 cluster (6.5 TFLOPs peak)
- 135 node Xeon32 cluster (1.08 TFLOPs peak)
- 32 node Itanium2 cluster (256 GFLOPs peak)
- Desktop Condor pool – growing to 750 Pentium4 PCs (4.5 TFLOPs peak)

TOTAL: 12.4 TFLOPs

Dell Xeon64 Cluster

1,024 Pentium4 Xeon64 CPUs
2,180 GB RAM
14 TB disk (SAN+IBRIX)
Infiniband & Gigabit Ethernet
Red Hat Linux Enterprise
Peak speed: 6.5 TFLOPs
Usual scheduler: LSF
Emergency Scheduler:
Condor

topdawg.oscer.ou.ed

DEBUTED AT #54 WORLDWIDE, #9 AMONG US UNIVS, #4 EXCLUDING BIG 3 NSF CENTERS

www.top500.org

Aspen Systems Xeon32 Cluster

270 Xeon32 CPUs
270 GB RAM
~10 TB disk
Myrinet2000
Red Hat Linux
Peak speed: 1.08 TFLOPs
Scheduler: **Condor**

Will be owned by High Energy Physics group

DEBUTED at #197 on the Top500 list in Nov 2002

www.top500.org

boomer.oscer.ou.edu

Aspen Systems Itanium2 Cluster

64 Itanium2 1.0 GHz CPUs
128 GB RAM
5.7 TB disk
Infiniband & Gigabit Ethernet
Red Hat Linux Enterprise 3
Peak speed: 256 GFLOPs
Usual scheduler: LSF
Emergency scheduler: **Condor**

Aspen Systems Inc.

`schooner.oscer.ou.edu`

Dell Desktop Condor Pool

OU IT is deploying a large Condor pool (750 desktop PCs) over the course of the 2006:

3 GHz Pentium4 (32 bit), 1 GB RAM, 100 Mbps network connection.

When deployed, it'll provide 4.5 TFLOPs (peak) of additional computing power – more than is currently available at most supercomputing centers.

Currently, the pool is 136 PCs in a few of the student labs.

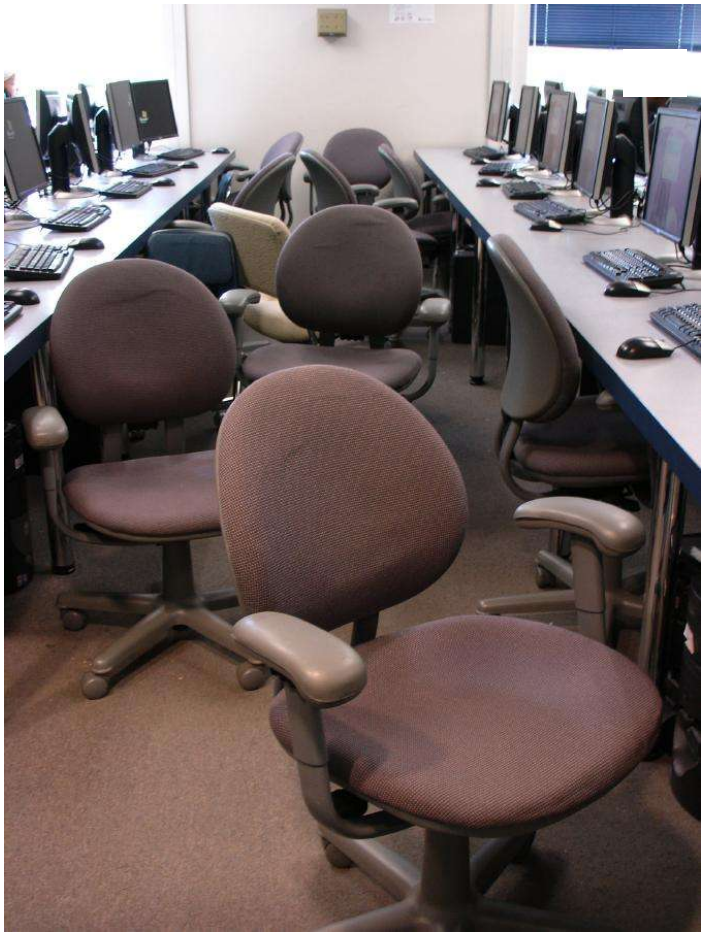

National Lambda Rail @ OU

Oklahoma has just gotten onto NLR; the pieces are all in place but we're still configuring.

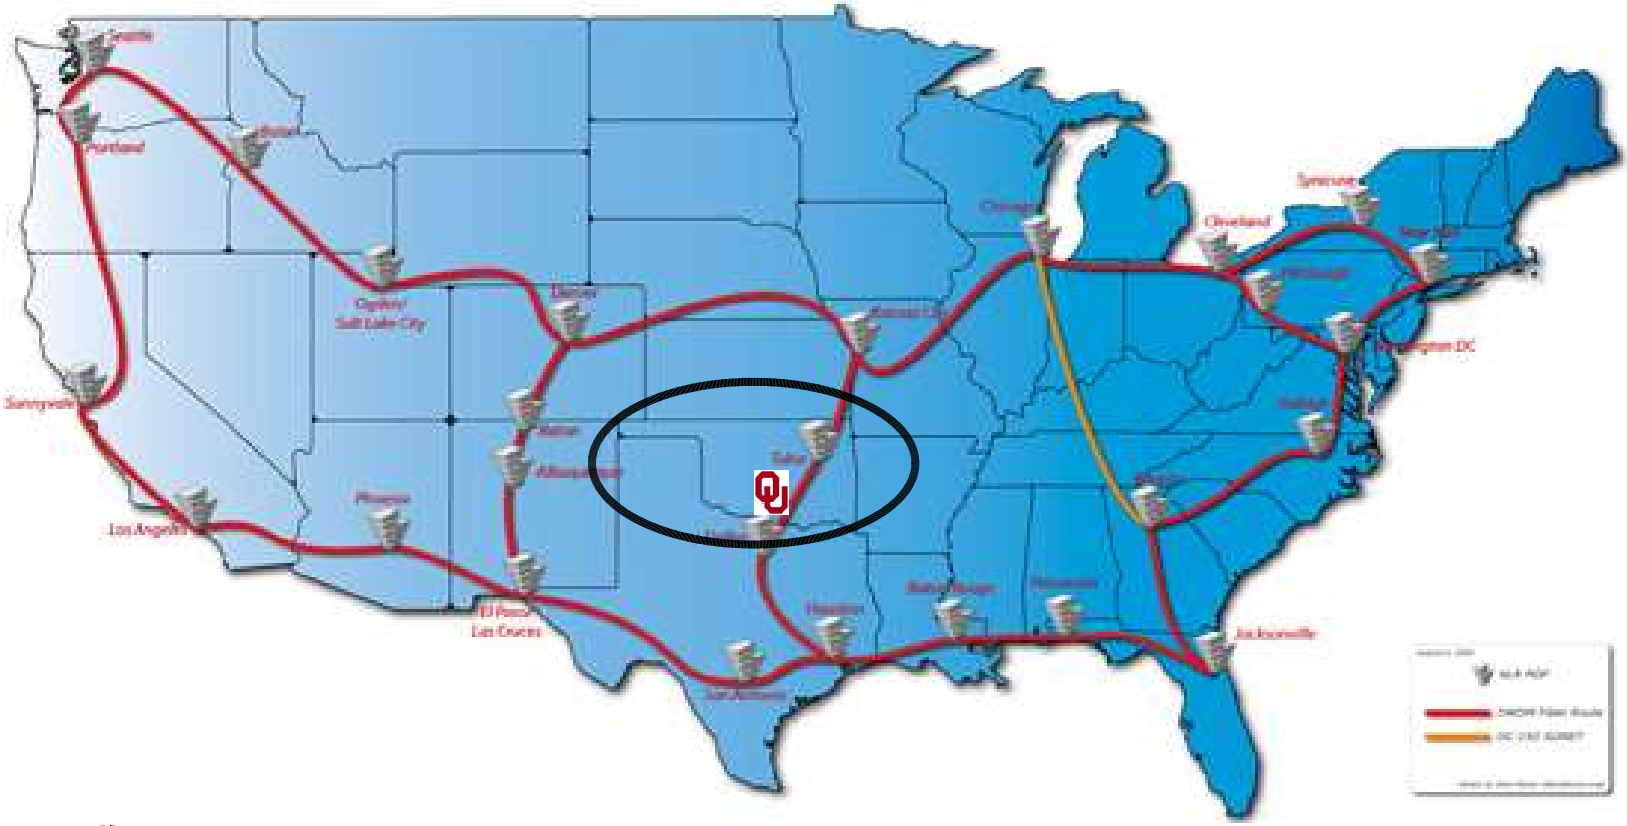

MPI Capability

- **Many kinds of national emergencies** – weather forecasting, floods, contaminant distribution, etc. – use fluid flow and related methods, which are tightly coupled and therefore **require MPI**.
- Condor provides the **MPI universe**.
- **Most of the available resources** – 7.9 TFLOPs out of 12.8 – **are clusters**, ranging from ¼ TFLOP to 6.5 TFLOPs.
- So, **providing MPI capability is straightforward**.

Fire Drills

- **Switchover** from production to emergency Condor:
  1. Shut down all user jobs on the production scheduler.
  2. Shut down the production scheduler (if not Condor; e.g., LSF).
  3. Start Condor (if necessary).
- **Condor jobs** for national emergency **discover** these resources and start themselves.
- We've done this several times at OU.
  - Only during scheduled downtimes!
  - Switchover times range from 9 minutes down to 2.5 min.
  - Pretty much we have this down to a science.

Thanks for your attention!

Questions?