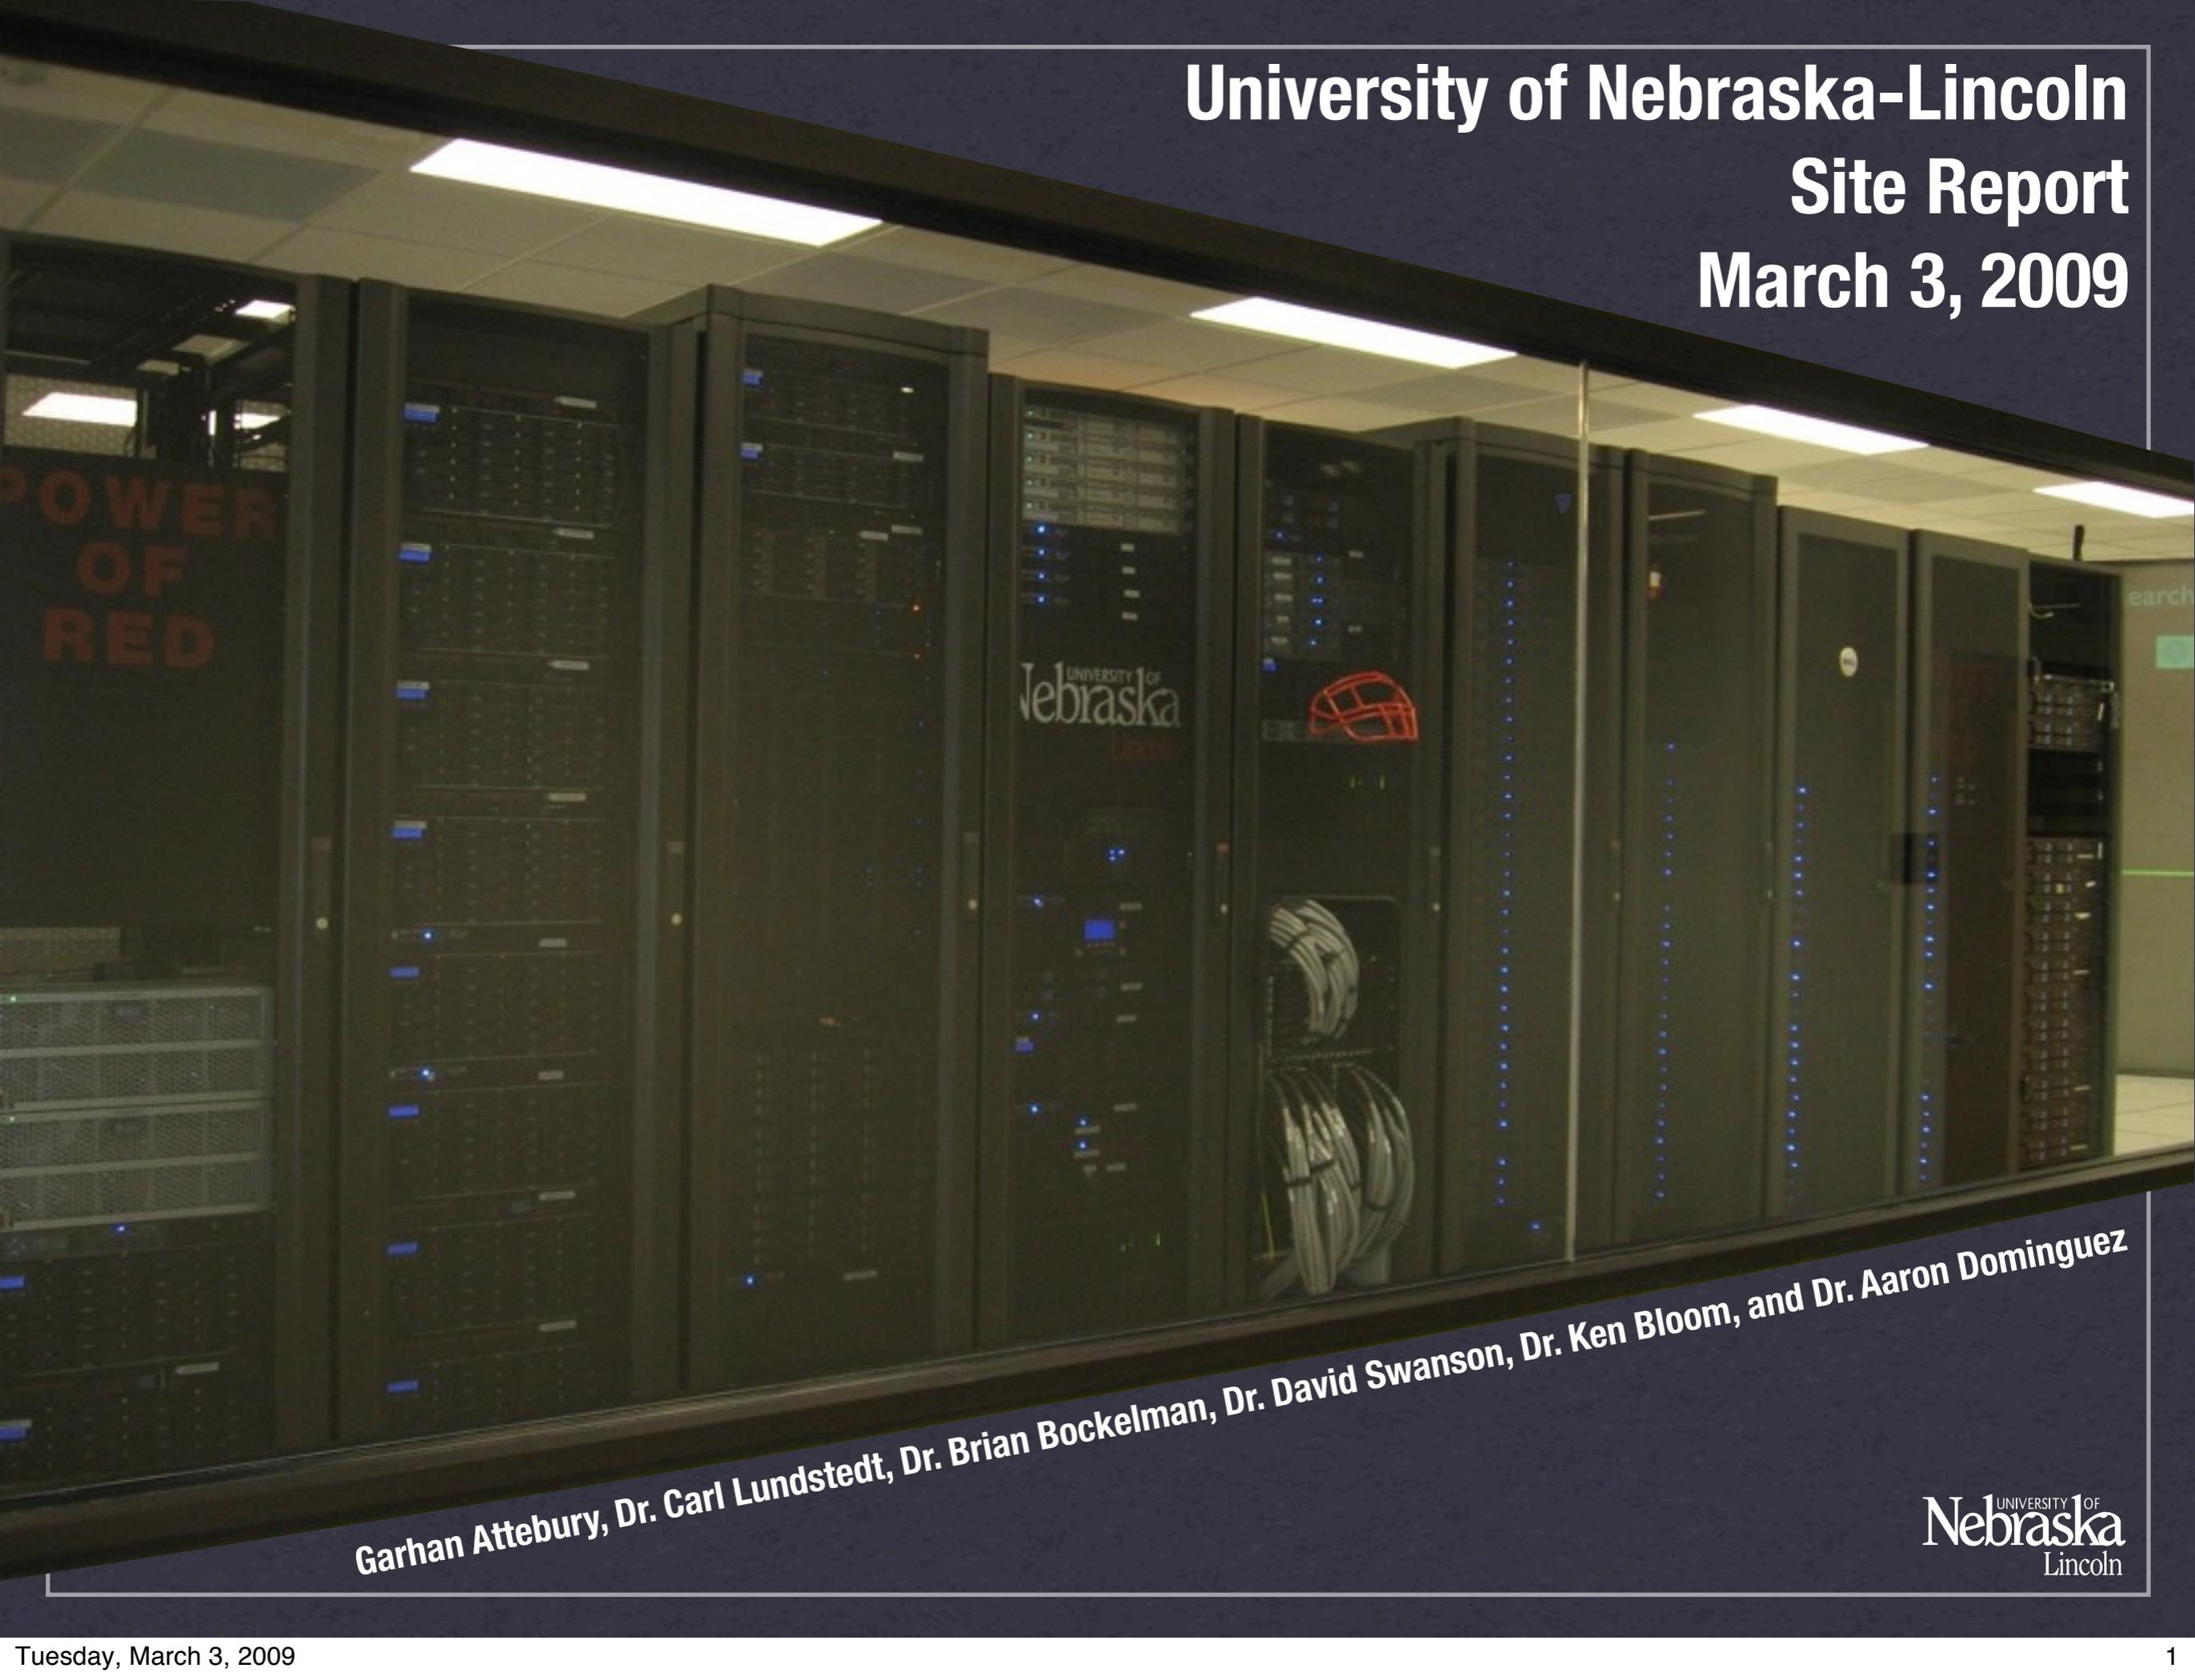


# University of Nebraska-Lincoln Site Report March 3, 2009



Garhan Attebury, Dr. Carl Lundstedt, Dr. Brian Bockelman, Dr. David Swanson, Dr. Ken Bloom, and Dr. Aaron Dominguez

# Tier2 Computing @ NE

- Current CE/SE Status
- Site experiences
- Future Plans



# Existing Hardware



- Compute Element
  - red.unl.edu + gpn-husker.unl.edu
    - 60x 4 core THPC, Opteron 275 w/4GB RAM
    - 53x 4 core Dells, Opteron 2216 w/4GB RAM + HDFS
    - 78x 8 core Suns, Opteron 2354 w/16GB RAM + HDFS
    - CPU = ~1076 worker cores / 2.32M si2k
    - Disk = ~235TB usable (170TB dCache + 300TB HDFS raw)

# Existing Hardware



- Storage Element
  - ~15 12-24 disk storage vaults (dCache)
  - 2x 'Thumpers' now dedicated to HDFS
  - 131x worker nodes with disks for HDFS
  - Disk = ~235TB usable (170TB dCache + 300TB HDFS raw)

# Existing Hardware

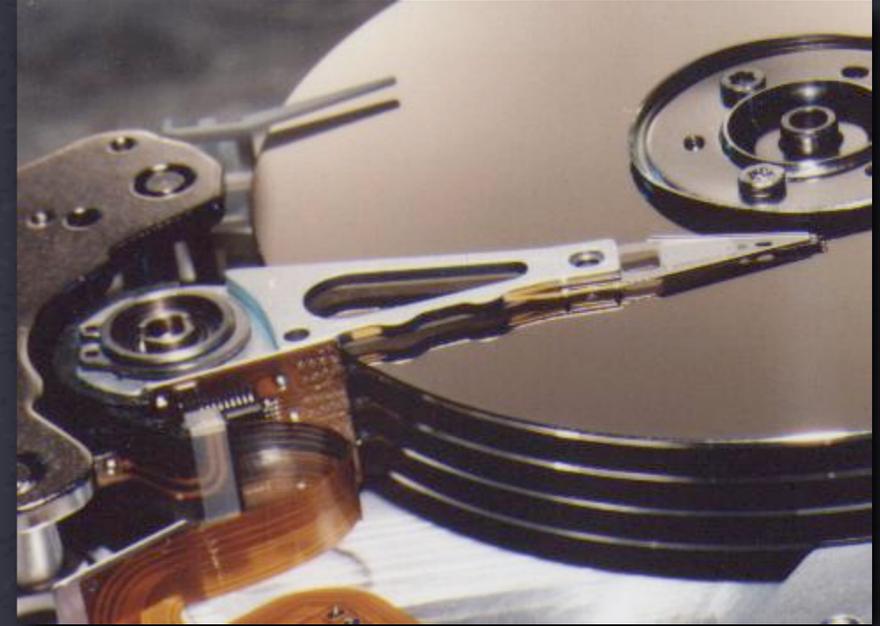


- WAN connection via 10Gb I2 + 10Gb DCN
- t3.unl.edu
  - 4x single socket dual core Opteron 175 w/2GB RAM
  - 4x dual socket dual core Opteron 2216 w/4GB RAM
  - login and submit machine for local HEP users

# Operational Experiences

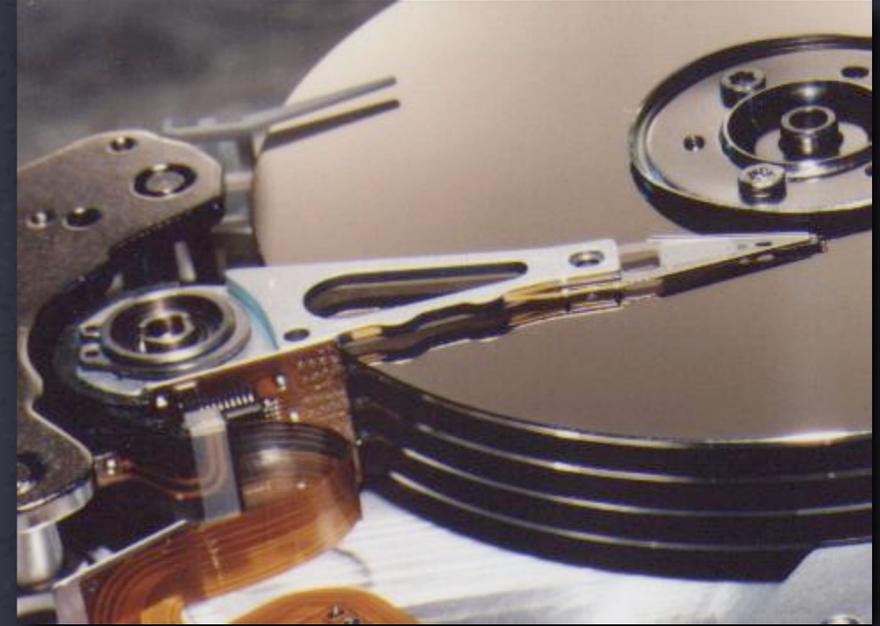
- Hadoop
  - Dual SEs can be a pain
    - TFC, anything but
  - GridFTP, SRM, fuse, etc...
  - Best practices, what works and doesn't, performance, etc ...
- Gratia
  - Local non-OSG cluster
  - HDFS transfers

# Future Plans



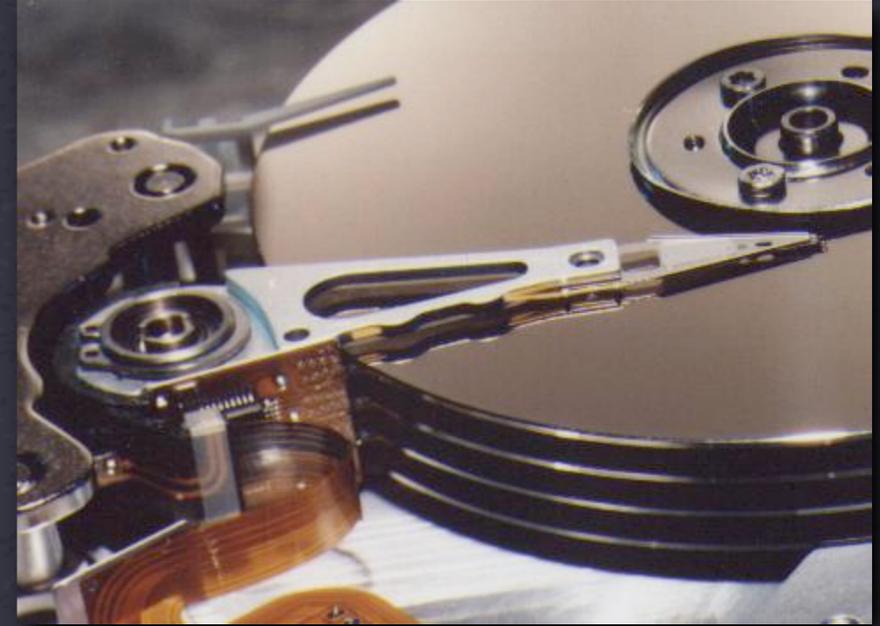
- Storage storage storage!
  - We have enough CPU (2.32M si2k) already
  - There never seems to be enough space... need ~200TB more
  - HDFS ideally suited to worker nodes with large disks
    - Cost effective
    - More CPUs
    - Perhaps harder to manage...

# Future Plans



- Need ~165TB minimum, round to 200TB
- Dell 2970 w/6x drives = ~8TB per 2U node, ~\$700 per usable TB
- 4 bay Supermicro/Asus 1Us, ~\$600 per usable TB
- Thumpers? \$31,000 per 41.5TB = ~\$1500 per usable TB
- Could meet 400TB usable with ~\$29,000 'absolute minimum'
- Better numbers? Can't quite say yet ...

# Future Plans



- HDFS -- documentation and findings
- Infrastructure upgrades (web services, virtualization, HA)
  - 10g GridFTP nodes?
- New resources -- local campus and Omaha -- grid-ify!