VANDERBILT UNIVERSITY

ACCRE
Advanced Computing Center for Research & Education

Open Science Grid

Storage Hardware How to Choose and Experiences

Kevin L. Buterbaugh
klb@accre.vanderbilt.edu

March 9th, 2010

ACCRE
Advanced Computing Center for Research & Education

It's not a cartoon, but...

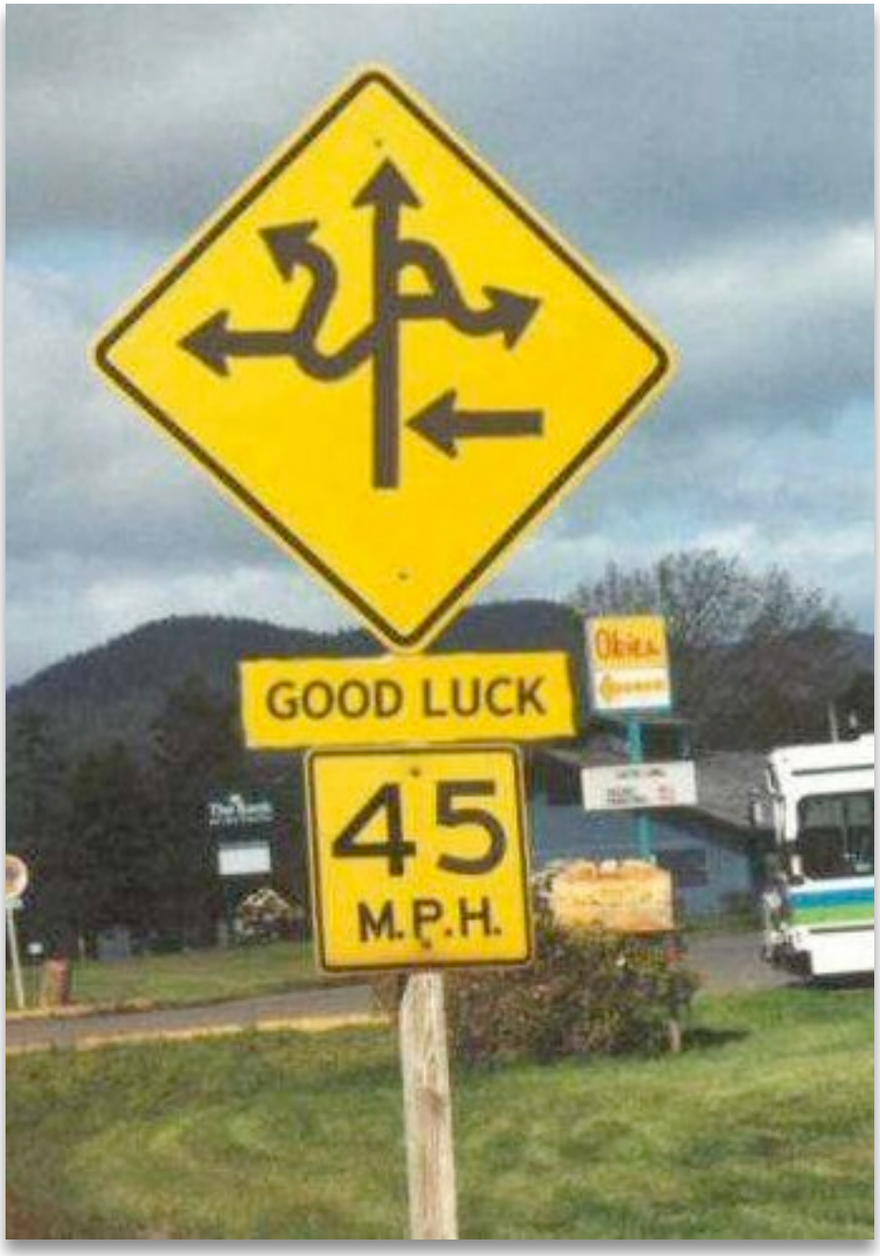

Three Common Options

- ☑ Storage Element with lots of local storage (possibly via a storage array). Shares storage to rest of cluster via NFS.
- ☑ Worker nodes with 4-6 local drives. Shares storage to rest of cluster via a distributed filesystem (Hadoop / REDDnet).
- ☑ Dedicated storage infrastructure. Typically a SAN with a parallel filesystem such as Lustre or GPFS.

Software on the SE

- There are numerous options for what software to run on your SE and what backend filesystem to use
- For the SE: BeStMan, dCache, xrootd, etc.
- For the filesystem: GPFS, Hadoop, Lustre, NFS, REDDnet, etc.

SE with local storage

- Who uses it? Maryland, probably most smaller Tier 3's
- Pro's...
- Easy to set up
- Uses standard utilities (i.e. NFS)
- Reasonable cost
- Con's...
- NFS, especially in conjunction with automounter, is notoriously flaky in Linux
- Doesn't scale well beyond ~150 clients

WN's with local storage

- Who uses it? Vanderbilt, Colorado, UC Davis?, others
- Pro's...
- Low cost
- Con's...
- Associated software typically has a learning curve for the SysAdmin
- Typically requires keeping multiple copies of all data to protect against the failure of a WN

Dedicated Storage

- ☑ Who uses it? Vanderbilt, Florida, others
- ☑ Pro's...
- ☑ Highest reliability and performance
- ☑ Con's...
- ☑ Cost (Dedicated I/O servers, storage arrays, SAN switches, storage software)