A Brief Trip into Spining Disk

Brock Palen

Center for Advanced Computing

Cluster Admins Meeting
February 23, 2007

Working Storage Options

Options

- Local Disk
- Shared Local Disk
- Shared Clustered Disk
- Ram Disk

Local Disk (ATA, SAS, SCSI, etc.)

Points Good/Bad

- User Isolation
- User Can Stream
- Cheap (Free with every node)
- Small
- Hard for Users
- Hard to Program

Local Disk (ATA, SAS, SCSI, etc.)

Points Good/Bad

- User Isolation
- User Can Stream
- Cheap (Free with every node)
- Small
- Hard for Users
- Hard to Program

Shared Local (NFS, Large SAS)

Points Good/Bad

- Most Common form of Storage
- Cheap
- Parallel Access (Simple to Program)
- Does not Grow
- Poor Performance

Shared Local (NFS, Large SAS)

Points Good/Bad

- Most Common form of Storage
- Cheap
- Parallel Access (Simple to Program)
- Does not Grow
- Poor Performance

Shared Cluster (Luster, PVFS2, Proprietary)

Points Good/Bad

- Fast
- Scales
- MPI-2 IO Support
- Cost
- Complex
- Small File IO Problems

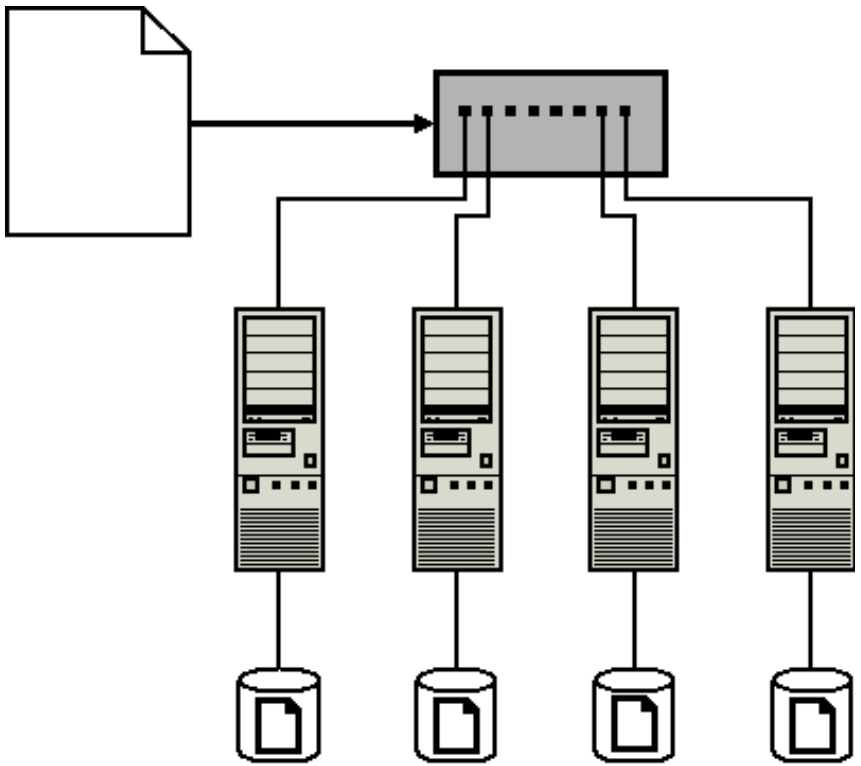

Shared Cluster (Luster, PVFS2, Proprietary)

Points Good/Bad

- Fast
- Scales
- MPI-2 IO Support
- Cost
- Complex
- Small File IO Problems

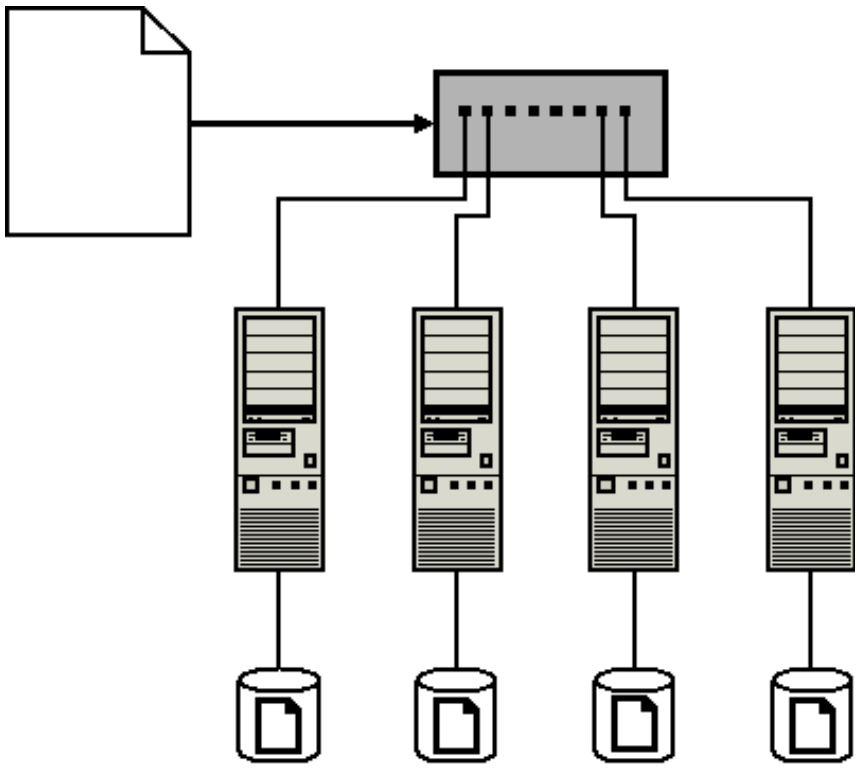

Ramdisk (OS Dependent)

Points Good/Bad

- `/dev/shm` on Linux
- Very Fast
- Expensive
- Volatile
- Swap

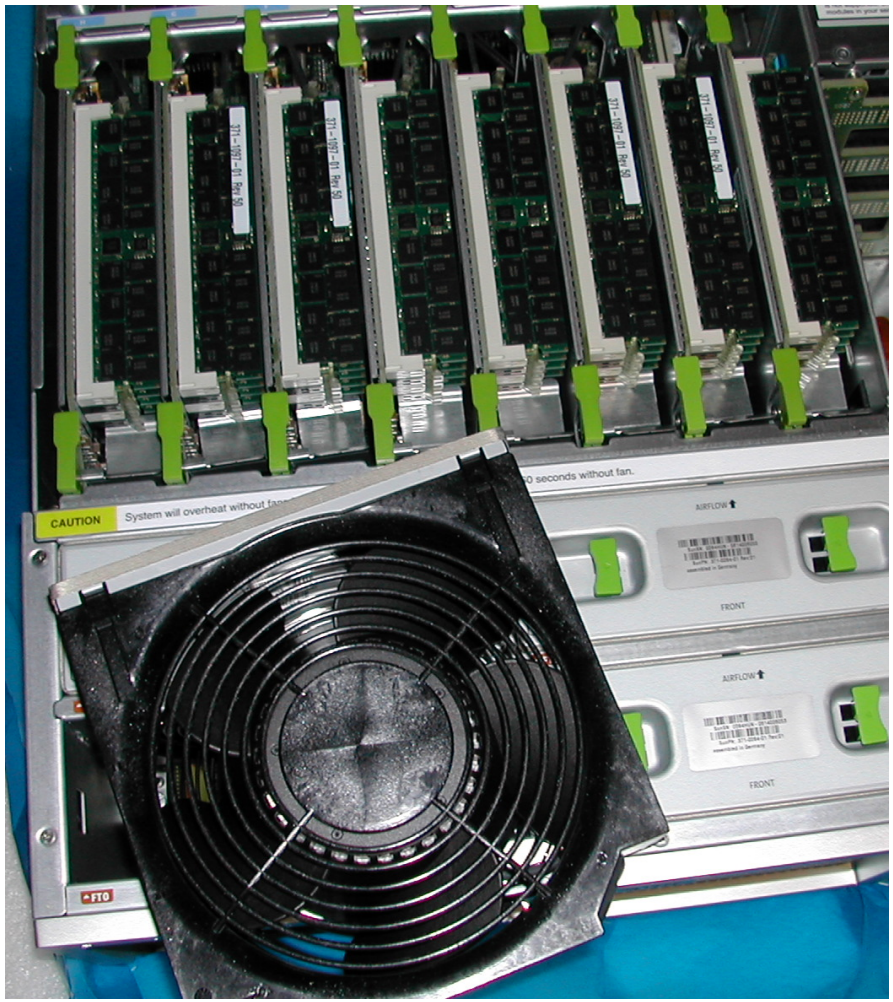

Ramdisk (OS Dependent)

Points Good/Bad

- `/dev/shm` on Linux
- Very Fast
- Expensive
- Volatile
- Swap

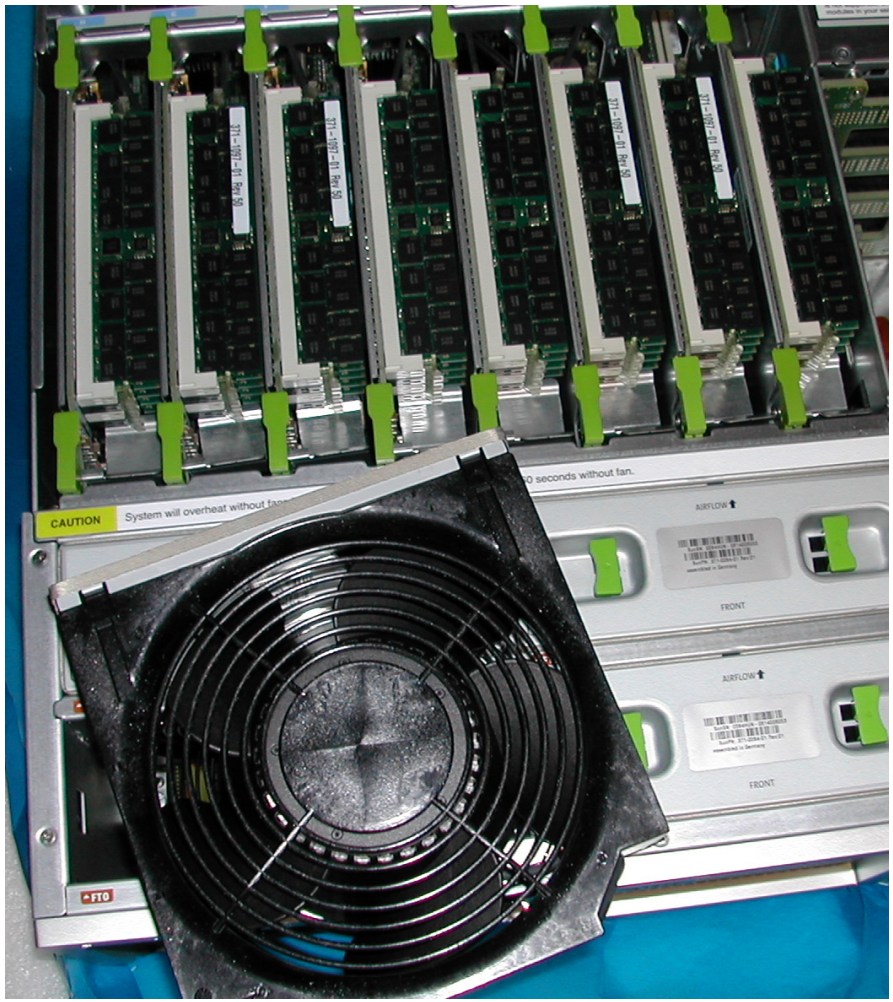

Abaqus/Standard

About Abaqus

A Direct Finite Element solver. Uses shared memory and threads for parallelism. Always writes a scratch file no matter ram size.

Example

- System: Sun X4600
- 8 3.0 GHz AMD Opteron 800 series cpus
- 64GB Ram
- single 10,000 RPM SAS drive 74GB

```
abaqus job=input user=user.f scratch=/tmp/$PBS_JOBID
cpus=8
```

http://cac.engin.umich.edu/resources/systems/nyxV2/bigmem.html

Abaqus Results

Results

Hardware	Walltime HH:MM:SS	CPU	Notes
x4600	4:35:29	Opt 856	4 CPUs Local Disk
x4600	3:23:30	Opt 856	4 CPUs `/dev/shm`
x4600	5:35:29	Opt 856	8 CPUs Local Disk
x4600	2:32:13	Opt 856	8 CPUs `/dev/shm`
x2100 [a]	7:19:13	Opt 175	2 CPUs Local Disk

[a] 72,000 RPM, 4GB DDR, 2.2GHz