Weizmann Institute Chemistry Faculty HPC

Mark Vilensky
Weizmann Institute of Science
Rehovot, Israel

About us

מכון ויצמן למדע
WEIZMANN INSTITUTE OF SCIENCE

- Six departments
- Ten heavy usage groups
- Long HPC history (clusters, Altix)
- Faculty cluster was planned in 2009 and built in 2010. It is continuously growing

Software types

- Commercial and semi-commercial
  - DFT (Vasp, QChem) - MPI
  - Ab initio (Molpro, Gaussian, Orca) - OpenMP
  - Matlab – serial, parallel toolbox
- Open source
  - Molecular dynamics (Gromacs, Namd, Amber) - MPI
  - Flexible Modeling System (FMS) - MPI
  - MIT GSM (Global Sky Model) - MPI
- Home made
  - Trajectories - serial
  - Monte Carlo – serial, OpenMP, MPI

Cluster topology

- 362 compute nodes
  - 354 dual socket
  - 8 quad socket
  - Memory ranges between 4 to 16 GB/core
  - Total 8,484 cores
- InfiniBand interconnect 5:1 blocking
  - Going to grow with non-blocking area
- GPFS/DDN
- CentOS 7.4
- PBS Professional
- Management servers on VMmware
- Staff = 2!

Storage and file systems

- DDN 7700, Infiniband
- 4 single socket NSD servers
- 2 file systems
  - Apps - /usr/local
  - Data - /gpfs/work – 100M files; /gpfs/home – 20M files. Metadata on 8 SSDs. 170 data drives. Block size – 4M, with 40M files smaller than a subblock (128kB)
- cNFS

Points of taking a decision

- At the beginning, in 2009
  - We knew nothing
  - Many players
- Storage replacement
  - Let's save some money

Parallel FS Comparison

	GPFS	Lustre	Bee GFS
HA	√	-	only mirror
Kernel module	√	own kernel	√
Own file system	√	-	-
Backup	mmbackup	external	external
Multirail	√	?	one port per service
Data in metadata	√	-	-

Parallel FS Comparison

	GPFS	Lustre	Bee GFS
Autotiering	√	-	-
Policies	√	-	-
Monitoring	ok	-	better
Management	central	local	local
Snapshots	√	~	~
Filesets	√	-	-
Call backs	√	-	-

Summary

- GPFS is a product
- Reasonable performance for many workloads
- Enterprise features
- Price?
- Node expels!
- Monitoring