

OSG storage at Minnesota

Graham Allan
allan@physics.umn.edu

CMS at Minnesota

- Current CMS infrastructure
 - File server: NFSv3 (plus SRM copy)
 - Approx 10TB storage (RAID6), using ext3
 - Approx 28 compute cores now
 - 1st gen servers: each 2xDC Opteron, 6GB.
- Need ability to scale to ~100TB storage and ~150 compute cores, later this year.
 - Dcache or hadoop+ bestman

How it's worked so far

- Linux systems, storage, condor managed by dept IT staff
- CMS software stack by faculty (Jeremiah Mans)
- The two may need to be integrated more as the software becomes more complex.

Minnesota Physics systems

- CMS systems are part of overall MN Physics linux systems
- Approx 80 server nodes, 250 workstations
- CMS, MINOS, CDMS, CLEO, BES-III, LIGO (etc)
- Scientific Linux 4 and 5
- Fully automated installation and configuration
- kickstart+yum+cfengine+svn
- Condor: single instance with multiple CondorGroup ClassAds

Minnesota Physics network

- In-house network infrastructure
 - 1Gbps router/firewalls to campus (failover pair)
 - HP Procurve switches (mix of 3500yl, 2910, 2848)
 - 1Gbps to all nodes/servers; 2/4Gbps LACP trunks between switches.
 - All systems on public address space
- Monitoring
 - Nagios, cricket

Storage hardware

- RAID6 storage on SAN
 - Some of this is a multi-group shared storage facility
 - AC&NC Jetstors with 750G and 1TB SATA drives (current); 2TB drives (in pipeline).
 - 4Gbps Brocade 4100 switches (two fabrics)
- Many local SATA drives

OSG storage at Minnesota

- Software – what's important
 - Dcache or BeStMAN/hadoop
 - Low management overhead for small staff
 - Installation and configuration are automated/repeatable
 - Natively packaged software (RPMs) vs pacman?
- Hardware
 - Arranging RAID6 for performance (eg drives/set)
 - Balance use of expensive RAID vs single drives
 - Linux SAN multipath support