

FY 2007 Minos Analysis Computing Resource Expansion

In response to the Young Minos requests in DocDB 3254-v1, we used the remaining FY07 Minos computing budget to purchase computing capacity for analysis batch jobs, and enough NFS served disk to support that additional computing.

The Young Minos request was to triple/quadruple the dedicated 12 core/36 GHz being delivered by 6 Minos Cluster nodes at that time, a net of up to 150 GHz

<i>Computing Resources</i>	<i>Capacity in GHertz</i>
FNALU Batch	137
Minos Cluster (in LSF)	75
Minos Cluster (notLSF)	75
Minos Grid (new nodes)	170

<i>Storage Resources</i>	<i>Capacity in TBytes</i>
Old Minos Cluster /local/scratch	5.7
AFS MINOS_DATA	10.5
/minos/scratch	10
/minos/data	20

We have asked that condor queues be established on the entire Minos Cluster. This is being actively pursued, we will work with the support people daily.

We have purchased eight new hosts for Analysis computing. They are Dell PowerEdge 1950 systems, with 2 x Quad-Core Intel Xeon 2.66 Ghz, 16GB RAM, and 500GB SATAu HD. Deliver is due by 8 October.

The new Analysis nodes will be deployed as part of the FermiGrid compute pool, so that other groups can use it when we do not. We will do what is necessary to make these nodes useful for Minos analysis, including making /afs/fnl.gov available on these systems.

Jobs will be submitted via the Condor system already used by the farm and by other experiments. Steven Cavanaugh has started testing this for Minos, and has run Loon jobs. The addition of AFS will greatly simplify this process.

Our queues are guaranteed to get as much capacity as we have purchased. More capacity beyond our purchase is available via the general FermiGrid queues. Use of those nodes may require the more general grid tools being tested by Steve and pioneered by Nick West.

The Minos farm is allocated about 175 execution slots.
But we do not count this toward our peak analysis capacity.

The new disk is a 32 TB SATA Beast array served by a BlueArc NFS server.
Usage allocations can be adjusted without interrupting service.

We plan to deploy the disk as follows :

10 TB – to replace the /local/scratch disks on the Cluster,
allowing us to move more Cluster nodes to analysis,
and to use some existing Cluster nodes for special servers.

20 TB – to supplement the existing 5 TB AFS physics work space

N.B. – CPU Terminology

A host computer can have multiple CPU's (usually 2).
A CPU can have multiple cores.
And the batch system may schedule more than one execution slot per core.

For the sake of simplifying the discussion above,
I use the term slot and core interchangeably.
In operation, we will need to verify that our slot allocation is adjusted
to deliver the cores we have purchased.