

FermiGrid Bluearc Unmount Task Force Report

DRAFT

Version 0.8

2015/03/16

By the FermiGrid Bluearc Unmount Task Force

Gerard Altayo - DCSO

Dmitry Litvintsev - DMS

Marc Mengel - SDPS/DMA

Andy Romero - ESO/SVS

Marco Slyz - SDCS/USDC

Steve Timm - ECF/VFP

Matt Tamsett - NOvA

Arthur Kreymer (chair) - SDPS/DMA

6. Storage System summary

Bluearc is a low latency robust file server, accessed via NFS, CIFS or FTP. Aggregate throughput to the present data areas is over 1/2 GByte/second. It has a limited ability to handle large scale parallel access, which we have regulated with ifdhc locks.

DCache is a moderate latency system with built in regulation of parallel file access. It supports both tape backed and scratch storage, with file aggregation for archival of small files. It has a high throughput, 1GByte/second per pool, and multiple pools, with the network setting the practical limit. Procols include DCAP, NFS 4.1, WEBDAV and xrootd. File access is limited to file creations and reading, no appending or modification.

In February 2015 the experiment-specific FTP servers moved to 10 GBit hosts, which are a better match to the Bluearc servers than the former hosts, which had limited capacity and support.

7. Working Documents

The work of the task force is primarily contained in the Redmine Wiki,

<https://cdcvs.fnal.gov/redmine/projects/fife/wiki/FermiGridBlue>

The task force was focused mainly on preparing the Strawman plan, and providing pointers to supporting documentation.

Since there is a nearly complete overlap with FIFE membership, we have had short meetings immediately after FIFE weekly meetings as needed.

Redmine Issues track detailed technical discussions.

- Overall milestones: <https://cdcvs.fnal.gov/redmine/issues/7100>
- Alternate data mount points: <https://cdcvs.fnal.gov/redmine/issues/7123>
- Move /grid/data to Data head: <https://cdcvs.fnal.gov/redmine/issues/7121>
- Locks based on Bluearc disk: <https://cdcvs.fnal.gov/redmine/issues/7122>
- GridFTP scaling: <https://cdcvs.fnal.gov/redmine/issues/7124>

Performance Metrics

The present metrics are linked to from <https://cdcvs.fnal.gov/redmine/projects/fife/wiki/FGBMon>

The principal metric we will use to measure availability is the fraction of time under 5 MB/sec as summarized at <https://cdcvs.fnal.gov/redmine/projects/ifdhcpn/wiki/PERFORMANCE#section-2>

Historically, the APP failure rate has been well under 1/10000 , and we would like to see similar numbers for the DATA areas, presently slowing at around 5/10000 (.05%)