

# Data and Computing Resources at JLab

Mark Ito  
Jefferson Lab  
GlueX Analysis Workshop  
Friday, May 13, 2016


# Outline

- CPU
- Disk
- Tape
- Build Systems
- Getting Help



# CPU: General Purpose Computing

- Web browsing, document writing, email reading...
- Linux:
  - [jlab13.jlab.org](http://jlab13.jlab.org) (RedHat Enterprise 6)
  - [jlab15.jlab.org](http://jlab15.jlab.org) (RedHat Enterprise 7)
- Windows:
  - public computers on the mezzanine of CEBAF Center
  - Windows Terminal Server Cluster (see [cc.jlab.org](http://cc.jlab.org))
- Mac:
  - Not that I know of




# CPU: Batch Farm

- Platform for large-scale scientific computing, batch mode
- Linux only, CentOS 6.5
- [farmyynn.jlab.org](http://farmyynn.jlab.org) (yy = year, nn = serial number)
- 3,800 cores on 220 nodes
- No direct login



www.shutterstock.com · 393906964

# Batch Farm



# CPU: Interactive Farm

- Same OS/configuration as batch farm machines
- ifarm.jlab.org (front end for ifarm1102, ifarm1401,...)
- Direct login to interactive shell
- Develop farm jobs scripts
- Access to scientific computing disks (work, cache, volatile...)
- Batch Farm job submission commands
  - `jsub`: submit a single job
  - SWIF: manage multiple, closely-related jobs
  - See Sean's talk and JLab SciComp page



# Disk: general purpose

- /home, new users start at 2 GB, backed up
  - Your home area: /home/username
- /group, 700 GB for Hall D, backed up
  - For shared files: /group/halld
  - Webpages, software builds, documents
- /apps, read-only
  - Non-standard versions, applications not supplied by distributions:  
/apps/bin
  - gcc, python,...
- /site, read-only
  - JLab-centric applications (/site/bin, e. g. staff command)
- /scratch, 4.6 TB, not backed up
  - Files unaccessed for two weeks are deleted





# Disk: SciComp (farm/ifarm only)


- /work
  - Large files, personal builds
  - /work/halld, 115 TB, Lustre
  - /work/halld2, 3 TB, traditional RAID (small files, SQLite database files)
- /mss
  - Front-end representation of tape library
- /cache, Hall D: 35 TB of pinned files
  - Actual file realization of /mss
  - Use jcache command
- /volatile, Hall D: 20 TB
  - Temporary storage, old files deleted to make space for new
- Local /scratch (farm and ifarm only)





# Tape

- IBM TS3500 Tape Library, 10 PB
  - 9909 slots in 9 frames
  - 10 LTO-5 drives, 140 MB/sec, 1.5 TB (raw) per LTO-5 cartridge
  - 6 LTO-6 drives, 160 MB/sec, 2.5 TB (raw) per LTO-6 cartridge
- Tape files presented as a disk file directory
  - /mss
  - One-to-one correspondence between files in directory tree and files on tapes
  - Structure in tree only loosely related to structure on tape
- Volume Sets
  - Nodes in directory associated with a named volume set
  - All files in and below that node get written to



# Tape Library



# Tape Library Robot in Action



# Build Systems

- Basic problem: many packages, many versions, many possible directory structures
  - Package list: sim-recon, JANA, HDDS, Xerces-C++, CERNLIB, ROOT, Geant4,...
  - Only certain combinations work
- A build system provides:
  - Directory structure
  - Help building each required package
  - Version management system
- GlueX has two:
  - build\_scripts (Mark Ito, [marki@jlab.org](mailto:marki@jlab.org))
  - Hall D Package Manager (HDPM) (Nathan Sparks, [nsparks@jlab.org](mailto:nsparks@jlab.org))

# GlueX Build Systems

- `build_scripts`
  - Fixed directory structure (alternate structures OK as well)
  - One makefile per package, version number an input
  - Version number for builds communicated via environment, XML configuration file drives environment definitions
    - Same XML file sets up user environment
- Hall D Package Manager (HDPM)
  - Specify a list of paths to sim-recon and its dependencies via text files
  - Based on Julia
  - Additional commands are provided for downloads, updates, and clean builds.
  - Can be used to install precompiled packages on Linux and Mac OS X

# How to Get Help

- Analysis Wiki Page
  - [https://halldweb.jlab.org/wiki-private/index.php/Gluex\\_Physics\\_Analyses](https://halldweb.jlab.org/wiki-private/index.php/Gluex_Physics_Analyses)
- Offline Software Wiki Page
  - Package documentation
  - HOWTO pages
  - Meeting minutes
  - [https://halldweb.jlab.org/wiki/index.php/Gluex\\_Offline\\_Software](https://halldweb.jlab.org/wiki/index.php/Gluex_Offline_Software)
- JLab Scientific Computing
  - <https://scicomp.jlab.org/docs/>
- Email Lists
  - [halld-offline@jlab.org](mailto:halld-offline@jlab.org), <https://mailman.jlab.org/mailman/listinfo/halld-offline>
  - [halld-physics@jlab.org](mailto:halld-physics@jlab.org), <https://mailman.jlab.org/mailman/listinfo/halld-physics>
  - [gluex-software@googlegroups.com](mailto:gluex-software@googlegroups.com), <https://groups.google.com/forum/#!forum/gluex-software>
  - Simple lists
    - notifications from automated tasks
    - [https://halldweb.jlab.org/wiki/index.php/Email\\_Lists#Simple\\_Email\\_Lists](https://halldweb.jlab.org/wiki/index.php/Email_Lists#Simple_Email_Lists)





Main page

Recent changes

Random page

Help

Tools

What links here

Related changes

Upload file

Special pages

Printable version

Permanent link

Page information

Cite this page

Page Discussion

Read

Edit

View history



More ▾

Search



# GlueX Physics Analyses

## Contents [hide]

- 1 Information (Links)
- 2 How to Setup and Do an Analysis
- 3 Projects/Studies
- 4 Analyses - References
  - 4.1 Analysis Workshops
- 5 Analyses - Commissioning Data
  - 5.1 Useful Links
  - 5.2 Non-Strange Meson Channels
  - 5.3 Strange Meson Channels
  - 5.4 Hyperon Channels
- 6 Analyses - Simulation Studies
  - 6.1 Non-Strange Meson Channels
  - 6.2 Strange Meson Channels
  - 6.3 Hyperon Channels
- 7 Analyses - Long-Term
  - 7.1 Non-Strange Meson Channels
  - 7.2 Strange Meson Channels
  - 7.3 Hyperon Channels
  - 7.4 Charm Channels

My Drive - Google Dri... x

Data and Computing ... x

Halld-physics Info Page x

GlueX Offline Softwa... x

https://halldweb.jlab.org/wiki/index.php/GlueX\_Offline\_Software

Search

Google vHall D vJLab vMisc. vNews vStandard Sites vAdd

Marki

talk

preferences

watchlist

contributions

log out

page

discussion

edit


history

delete

move

protect

unwatch



navigation

Main page

Recent changes

Random page

Help

search

Search

Go

Search

tools

What links here

Related changes

Upload file

Special pages

Printable version

Permanent link

Page information

Cite this page

# GlueX Offline Software

## News:

- March 8, 2016: [sim-recon 1.10.0 released](#).
- March 2, 2016: [Minutes of the March 2nd meeting](#) are available.
- February 22, 2016: [2016 GlueX Physics and Analysis Workshop](#) announced.
- February 9, 2016: [Kinematic Fitter Overhaul](#) done.
- February 3, 2016: [Minutes of the February 3rd meeting](#) are available.
- [Previous news items...](#)

The legacy version of this page is [here](#).

Contents [hide]

1 General Information

1.1 Shell Environment Set-Up

1.2 Building GlueX Software

2 Software Documentation

2.1 GlueX Software

2.1.1 Software Packages

2.1.2 Special Topics

2.2 Documentation for External Software Packages

3 Data Sets

4 Offline Data Monitoring

5 Computing Facilities

5.1 JLab

5.2 Grid

6

# JLab Scientific Computing

Search this site:

## Navigation

- [Books](#)

## User login

Username: \*

Password: \*

- [Request new password](#)

## Scientific Computing at Jefferson Lab

**Scientific Computing** has two missions: **Experimental Physics Analysis and Simulation**, and **Lattice QCD Computing**. Activities include deploying and operating computing resources, and research and developments in **High Performance Computing**.

Many resources (file servers, offline storage, wide area networking and sometimes compute nodes) are shared between the two missions with appropriate allocations. Details are found at the following links:

- [Scientific Computing Resources](#) -- details of hardware
- [LQCD User's Guide](#) -- how to use the LQCD / HPC computing systems
- [Farm User's Guide](#) -- how to use the Experimental Physics "Farm" and associated systems

User Portals:

- [LQCD Computing User Portal](#)
- [Experimental Physics User Portal](#)

## Halld-offline --

### About Halld-offline

English (USA)

To see the collection of prior postings to the list, visit the [Halld-offline Archives](#).

### Using Halld-offline

To post a message to all the list members, send email to [halld-offline@jlab.org](mailto:halld-offline@jlab.org).

You can subscribe to the list, or change your existing subscription, in the sections below.

### Subscribing to Halld-offline

Subscribe to Halld-offline by filling out the following form. You will be sent email requesting confirmation, to prevent others from gratuitously subscribing you. This is a private list, which means that the list of members is not available to non-members.

Your email address:

Your name (optional):

You may enter a privacy password below. This provides only mild security, but should prevent others from messing with your subscription. **Do not use a valuable password** as it will occasionally be emailed back to you in cleartext.

If you choose not to enter a password, one will be automatically generated for you, and it will be sent to you once you've confirmed your subscription. You can always request a mail back of your

## Halld-physics -- Hall D physics discussion

### About Halld-physics

English (USA)

To see the collection of prior postings to the list, visit the [Halld-physics Archives](#).

### Using Halld-physics

To post a message to all the list members, send email to [halld-physics@jlab.org](mailto:halld-physics@jlab.org).

You can subscribe to the list, or change your existing subscription, in the sections below.

### Subscribing to Halld-physics

Subscribe to Halld-physics by filling out the following form. You will be sent email requesting confirmation, to prevent others from gratuitously subscribing you. This is a private list, which means that the list of members is not available to non-members.

Your email address:	<input type="text"/>
Your name (optional):	<input type="text"/>

You may enter a privacy password below. This provides only mild security, but should prevent others from messing with your subscription. **Do not use a valuable password** as it will occasionally be emailed back to you in cleartext.

If you choose not to enter a password, one will be automatically generated for you, and it will be



## Groups

NEW TOPIC



Mark all as read

Actions

Filters



### GlueX Software Help Shared publicly

22 of 22 topics (6 unread)

Manage · Members · About

Welcome to the GlueX Software Help Email List. Simple questions, tough questions, proposals, ideas, suggestions, discussion: all are welcome.

[Edit welcome message](#) [Clear welcome message](#)

- ☐ **Fwd: Hall D farm usage since Feb 1** (1)  
By Mark Ito - 1 post - 1 view May 11
- ☐ **2013 analysis workshop materials** (1)  
By Mark Ito - 5 posts - 22 views May 11
- ☐ **network outage at lab** (1)  
By me - 1 post - 5 views May 6
- ☐ **Delta\_T vs P**  
By bcannon@jlab.org - 8 posts - 7 views May 5
- ☐ **Fwd: possible problems with DChargedTrackHypothesis**  
By Paul Mattione - 8 posts - 29 views May 4
- ☐ **Today's farm fairshare status -- Fwd: Re: CCPR 116913 UPDATE (Re: farm priorities & threads)**  
By Mark Ito - 1 post - 3 views May 4
- ☐ **How many hypotheses per one track?**  
By semenov@jlab.org - 6 posts - 5 views May 4



naid-target [send email](#) [archive](#) [information page](#)

halld-tracking-hw [send email](#) [archive](#) [information page](#)

halld-trigger [send email](#) [archive](#) [information page](#)

## Simple Email Lists [\[edit\]](#)

See the [Simple Email Lists](#) wiki page for information on how to subscribe as well as other tasks.

List Name	Frequency	Description
nightly_build	daily	Errors and warnings from the <a href="#">nightly build of GlueX software</a>
b1pi	Mondays and Thursdays	Notification of completion of the bi-weekly <a href="#">b1pi reconstruction job</a>
single_track	Tuesdays and Fridays	Notification of completion of the bi-weekly single track reconstruction job
ccdb_changes	daily	A summary of changes to the CCDB during the last 24 hours.
git_update	daily	Sends a listing of files that have changed for all branches of all of the GlueX Git repositories.
monitoring_update	daily	Links to a minimal set of monitoring histograms, accumulated online, for runs taken during the previous twenty four hours.

These lists are sent to automatically by cron jobs running as user [gluex](#) on [jlabl1.jlab.org](#).

## Subversion Repository Check-In Email Notifications [\[edit\]](#)

Subversion has a feature where an arbitrary command can be issued on each check-in. We use this feature to send email notifications to interested parties. As such this acts just like an email list.

The scope of the messages can be customized; if you are interested in only particular branches of the tree, then notifications can be

# Summary

- JLab has CPU, Disk, and Tape
- Two systems for building software
  - build\_scripts
  - HDPM
- Omitted: how to actually use these resources
- See “How to Get Help” slide

