
JLab Computing Tips and Tricks

Brad Sawatzky

First up:
A Couple Quick Tricks
to make your
Computing Work Suck Less

How to find information

- JLab's web search sucks and no one cares...

→ *Still* working on improving this...

- » Baby steps: [ServiceNow SciComp Portal](#) “Knowledge Base”
- » [Getting Started](#) and [Experimental Physics User’s Guide](#) pages are being updated
 - Info still not widely searchable, sigh...

→ **Search trick: do this in Firefox:**

- » Go to www.google.com and search for 'site:jlab.org foo'
- » Right click on the bookmark and choose 'Properties'
 - Give it a good name
 - Give it a short 'keyword' like 'jj'
 - Clean up the Location as shown, replace 'foo' with %s

→ **Now type 'jj jget' in URL bar**

- » %s in 'Location' string is replaced with text following Keyword

Name:	[jj] JLab Search
Location:	http://www.google.com/search?hl=en&q=site:jlab.org%20%s&btnG=Search
Tags:	Separate tags with commas
Keyword:	jj
Description:	

- » 'site:jlab.org' is google-fu to restrict search to jlab.org domain

How to find information

- Trick works great for many things
 - **JLab staff page** (<https://misportal.jlab.org/mis/staff/staff.cfm>)
 - » Keyword: 'page'
 - » Location (can extract from search on 'smith' above):
 - » `https://misportal.jlab.org/mis/staff/staff.cfm?field=all&name=%s&Search.x=36&Search.y=11&Search=Search&field=all`
 - **ROOT / G4**
 - » Keyword: 'gr'
 - » Location:
 - `https://www.google.com/search?hl=en&btnG=Search&q=site:cern.ch%20%`
 - **Stackoverflow.com**
 - **JLab Logbook (a little trickier, but you can work it out)**
 - ...

How to work from Offsite

- How to work from offsite without tearing your eyes out because, holy hell, the graphics and menus are just so slow...
- VNC + ssh tunnel to the rescue
 - VNC: Virtual Network Computing
 - ssh used to securely move VNC traffic through jlab firewall



- Computer Center How-to
 - <https://cc.jlab.org/accessingvnc>
- Old 'howto' I wrote for my collaboration
 - adapt to machine you use
 - Search: 'jj vnc session'
 - » Pick: [Using a VNC Server/Client](#)

How to work from Offsite

- How to work from offsite without tearing your eyes out because, holy hell, the graphics and menus are just so slow...
- Virtual Desktop Environment (VDI)
 - <https://vdi.jlab.org>
 - Fewer “hoops” than VNC, but...
 - » limited number of ‘slots’ available
 - » logins are not as persistent



- Computer Center How-to
 - <https://cc.jlab.org/remotearchive>

Offline Analysis Farm Usage / General JLab Computing

Nuts to the Farm, I analyze on my Desktop

- Simple tasks, some analysis OK on the desktop, BUT!!
 - Thou shalt backup your code!
 - Thou shalt backup your results!
 - Who among us has done
 - `% rm -rf stuff/`
 - » Followed by !@#\$?
- Don't keep only copies on your laptop
- Don't keep only copies on your desktop's hard drive
- Do use git for all code and scripts!
 - Commit early, commit often
 - 'git push' often too!
 - » It's a backup!
- Hard drives die and the data are gone.
 - Drives are large and cheap
 - But reliability on consumer drives is worse that it used to be!
 - SSDs are (weirdly) no better!
- IF your hard drive died today, how long would it take to recover?
 - » a day, a week,
 - » a month???

JLab Systems can help!

- **/home, /group** are automatically backed up
 - They are snapshotted hourly!

```
% cd .snapshot/  
% ls -lrt
```
 - Longer term backups are on tape
- **/work, /volatile** are on heavily redundant filesystems
 - NOT backed up
 - » Use tape
 - More on this later...
- **NOTE:** Your JLab RHEL system can mount these directories if needed
 - Talk to me if this would help

The JLab Farm • Power at your Fingertips

- Farm has many pieces
 - ~30000 compute cores
 - ~6 PB Lustre
 - ~5 PB NFS/XRootD (ZFS)
 - ~100+ PB of Tape
 - Consumes ~400kW of power!
- Growth is \$\$\$ and based on projections from Halls
 - Expenditures often switch between storage + CPU every other year




The JLab Farm • Batch Computing

- The Farm: Batch Computing
 - No direct access to these machines
 - » Use “Interactive” farm nodes for testing
 - ie. ifarm1802
 - DB and other network access (git, http, etc) generally constrained
 - Jobs controlled by automated system called “slurm”
 - You submit a job via slurm or swif and slurm schedules it to run
- All about trade offs:
 - “Latency” can be high (hours+ from submission to job execution)
 - » BUT!
 - Throughput is enormous
 - » 100s (1000s) of jobs can run simultaneously
 - » High bandwidth access to fast storage
 - A full replay (100s of runs) can be completed in the time it would take 2–3 runs to complete in series on your desktop.

The JLab Farm • Scheduling

- The Farm is a Lab-wide shared resource
 - Each Hall's budget includes \$\$\$ to support their usage
 - *Rough* allocation:
 - » A: 9%, C: 9%
 - » B: 34%, D: 34%
 - » EIC: 14%
- Ruled by Slurm workflow manager
 - Allocations *not* written in stone and are adjusted based on needs
- The balance is trickier to manage than you may think...
 - Jobs take time to run (system doesn't know how long beforehand)
 - Upcoming job load is hard to predict
 - System balances allocations over a few days, not hours
- More documentation here:
 - <https://scicomp.jlab.org/>
 - <https://data.jlab.org/>

Farm Cluster Daily Usage by Account



Slurm Fairshare Setting/Usage Info

Farm Utilization

Farm Usage

GPU Usage

Monthly Report

Summary Report


Fairshare

Summary Chart


Select date Range

4/16/2023 - 5/17/2023

Slurm Fairshare Setting



Slurm Accounts Usage (CPU Hours)



Do use the Farm!

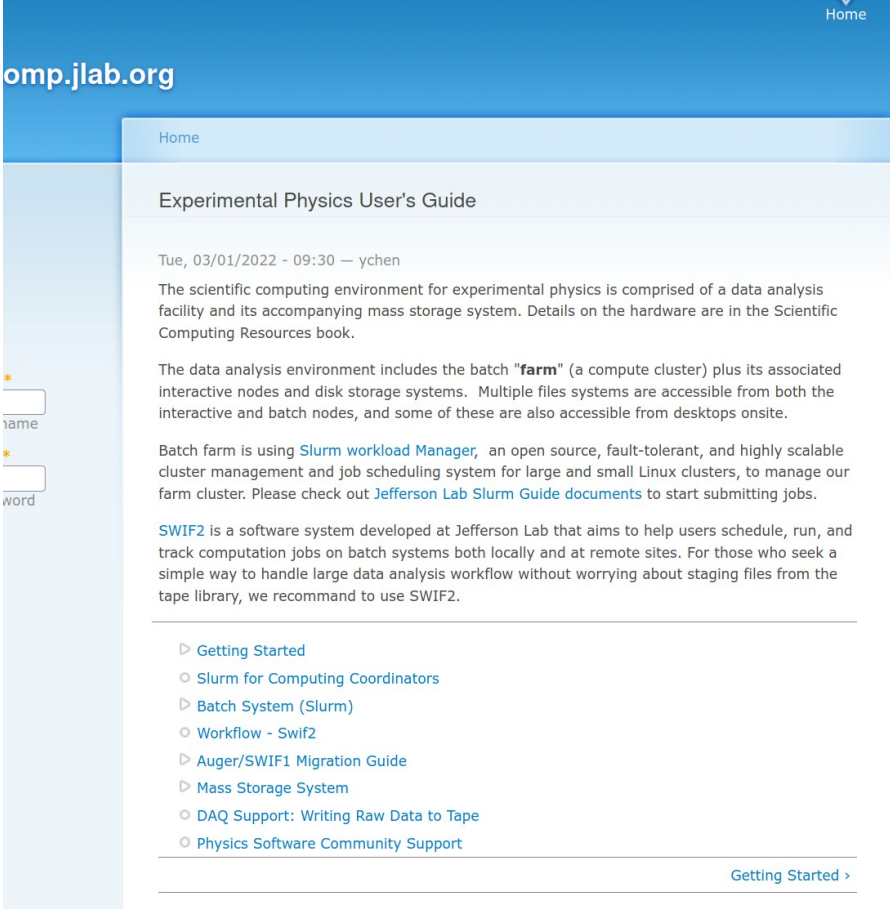
- The Farm is not your desktop
 - Best to plan, test, and fire off groups of jobs
- Test your job first!
 - Can it run reliably?
 - » If it doesn't run on ifarm180x, it won't run on the farm!
 - Is the output what you want?
 - » Check before firing off 100 jobs
- Simple tasks, some types of analysis can be done on small systems, BUT!!
 - Thou shalt back up your code!
 - Thou shalt back up your results!
 - IF your hard drive died today, how long would it take to recover?
- Don't keep only copies on your laptop
- Don't keep only copies on your desktop's hard drive



What's a “Job”?

- A 'Job' often maps to a shell script
 - It can do multiple things, but usually it executes a single instance of your software
 - » Analyze one run, or
 - » Simulate “1M” events,
 - » *etc...*
- **NOTE:** Output that would normally go to a terminal goes to special file system:
 - `/farm_out/$USER/job_id.out`
 - `/farm_out/$USER/job_id.err`

<https://scicomp.jlab.org/docs/FarmUsersGuide>



The screenshot shows the 'Experimental Physics User's Guide' page on the scicomp.jlab.org website. The page title is 'Experimental Physics User's Guide' and it is dated 'Tue, 03/01/2022 - 09:30' by 'ychen'. The content describes the scientific computing environment, including the 'farm' cluster and the 'Slurm workload Manager'. It also mentions 'SWIF2' as a software system for scheduling and tracking jobs. A table of contents is visible at the bottom of the page, listing sections like 'Getting Started', 'Slurm for Computing Coordinators', 'Batch System (Slurm)', 'Workflow - Swif2', 'Auger/SWIF1 Migration Guide', 'Mass Storage System', 'DAQ Support: Writing Raw Data to Tape', and 'Physics Software Community Support'. The page has a blue header with 'scicomp.jlab.org' and a 'Home' link. There are also search input fields labeled 'name' and 'word' on the left side.

Debugging a job

- Generally want a single script that does everything!

→ Set up full environment

→ Use full paths

» /group/myExp/myscript.sh

» ~~./myscript.sh~~

- Testing your script:

→ 1st: Run on ifarm180x

→ 2nd: Submit job to Farm

- Test with the 'debug' Farm track

→ Max priority, fast sched.

→ Limited 4 hour runtime

→ Limited jobs/user

- Test on ifarm180x

```
% ssh you@ifarm1802
```

```
% /group/myExp/myscript.sh
```

→ Make sure it worked!

» check histos, report files

- Quick Test on Farm

```
% swif2 add-job -create \  
-track 'debug' \  
<other options> ... \  
/group/myExp/myscript.sh
```

→ Make sure it worked!

» check histos, files

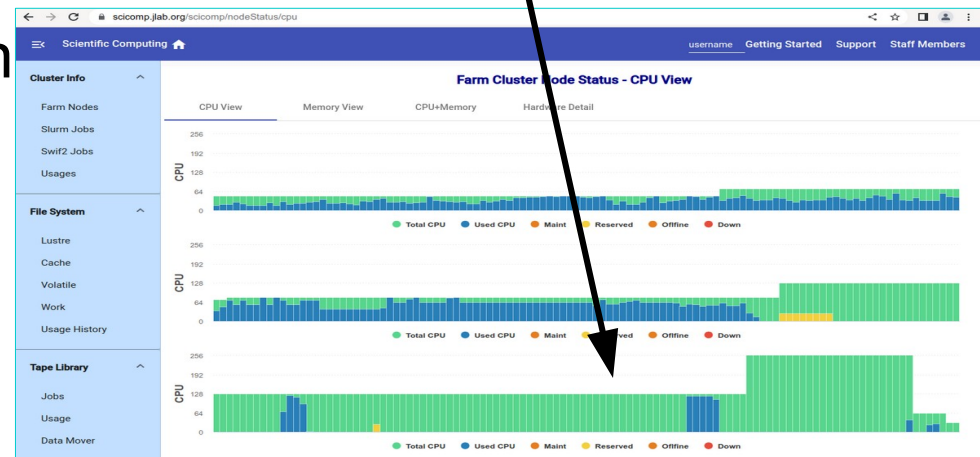
» check /farm_out/\$USER/

- Then submit full set!

→ SWIF2!

Make your jobs run faster!

- Scheduling jobs takes many things into account
 - File availability from tape
 - Memory request
 - CPU/core request
 - » >1 is useless for podd/hcana
 - 'Fairshare' metric
 - » Average Hall utilization
 - » Hall Usage can be subdivided further
- Details
 - [Fairshare Web Page](#)
- If a Hall / Project is not using 'their' fraction, then those Farm resources are available to anyone on a first-come, first-serve, basis!
 - If the Farm is idle, you can take advantage!
 - » Like now!



Make your jobs run faster!

- Common Bottlenecks/ Mistakes
 - CPU count
 - » use 1 core only (unless you know the job can multi-thread!)
 - Memory allocation
 - » < 2GB is best!
 - » Smaller → Faster scheduling!
 - Insufficient debugging/ cross checks
 - » Fire off 100s of jobs with bad config, buggy code



Check Job Status

← → ↻ 🏠 🔒 🔍 https://scicomp.jlab.org/scicomp/swif/active

Scientific Computing 🏠 username Getting Started Support Staff Members

Cluster Info ^

- Farm Nodes
- Slurm Jobs
- Swif2 Jobs
- Usages

File System ^

- Lustre

Active Workflows

Active Workflows Dormant Workflows Workflow Summary File Queue Globus Status

Filter

Id	Site Name	Workflow Name
54180	jlab/enp	offmon_2023-01_ver03_post
54179	jlab/enp	analysis_2017-01_ver64_batch01_merge
54178	jlab/enp	rgc-dra-dst_sqlite3-16327
54177	ilah/enn	rgc-dra-dst_sqlite 2-16327

Workflow Summary

Active Workflows Dormant Workflows Workflow Summary File Queue Globus Status

Choose a user: **jjagle** Choose a workflow: **wf-RunPeriod-2019-01-target-nobfield-primex-eta-full-ver...**

Succeeded Job Usage Distribution

Mean waittime: 0 (hrs)

Mean cputime: 2 (hrs)

Mean memory used: 2.275 (MB) (4,000 MB requested)

VMemory Used (MB)

Accumulated Done Job

Number of attempts for each job

Problem Type Distribution

swif Job Id	Slurm Job Id	Attempt Id	Job Name	Problem Code	Node	Resolution	Complete Time
16255661	64052684	22256631	wf-RunPeriod-2019-01-target-nobfield-primex-eta-full-ver-17052023-skim-eta2g-skim_061378	SLURM_FAILED	farm140211		May 17, 2023 7:32:30 PM
16255663	64052688	22256633	wf-RunPeriod-2019-01-target-nobfield-primex-eta-full-ver-17052023-skim-eta2g-skim_061391	SLURM_FAILED	farm140209		May 17, 2023 7:27:32 PM
16255665	64052686	22256635	wf-RunPeriod-2019-01-target-nobfield-primex-eta-full-ver-17052023-skim-eta2g-skim_061435	SLURM_FAILED	farm140122		May 17, 2023 7:14:59 PM
16255667	64052692	22256637	wf-RunPeriod-2019-01-target-nobfield-primex-eta-full-ver-17052023-skim-eta2g-skim_061437	SLURM_FAILED	farm140127		May 17, 2023 7:18:26 PM

- https://scicomp.jlab.org/scicomp/swif/active
- Workflow Summary tab can help you find information how jobs ran (or didn't run...)
 - ie. Memory usage!
 - See also: /farm_out/\$USER/*

Jefferson Lab

Thomas Jefferson Lab Software & Cloud

Small I/O Problems

- Small read/write operations are very inefficient
 - Old/legacy code defaults can be very small (~4kB)
 - Should be closer to 4MB chunks for decent performance
 - Buffered IO can bridge the gap if needed
 - » Common errors:
 - 'Debugging' output
 - » `stderr << "got here" << endl;`
 - » `fprintf(stderr, "event %d\n", eventNum);`
 - Opening/closing files very frequently
 - **Frequent** random I/O
 - » ie. searching through a file for a parameter every event
- Workflows / procedures that may work on desktops or older systems do not scale well on modern systems (100s or 1000s of simultaneous jobs)
 - **Can take down / degrade system-wide filesystems**
 - Always be mindful you are on a large-scale shared system, not a personal desktop

File Systems: Where do I put my stuff?

- SciComp/IT provides
 - /group - a space for groups to put software and some files, backup up by CST
 - /home - your home directory, backed up by CST
 - Cache - write through cache
 - Volatile - acts as a scratch space
 - Work – unmanaged outside of quotas / reservations

File Systems: Where do I put my stuff?

- SciComp/IT provides
 - /group - a space for groups to put software and some files, backup up by CST
 - /home - your home directory, backed up by CST
 - Cache - write through cache
 - Volatile - acts as a scratch space
 - Work – unmanaged outside of quotas / reservations
- No really, that doesn't help. Where do I put my stuff so Brad won't hassle me and I can get my work done in peace!

Where do I put my JLab stuff?

- `/home/<you>/`
 - hourly snapshots
 - » `cd .snapshot/`
 - personal, non-analysis files
 - » papers, notes, thesis, etc...
 - analysis scripts: ~OK
 - » use git!
 - source code: ~OK
 - » /work better
 - NEVER store ROOT files or CODA files in /home
- Your laptop / desktop
 - Should **really** be just a front-end for working on JLab systems
 - Everybody plans to do backups, but almost no one actually does backups until **after** they've lost data...



Where do I put my stuff?

- /group

- Think “/home” for work groups
 - » papers, thesis, etc
- hourly snapshots
 - » `cd .snapshot/`
- analysis scripts: YES
 - » use git!
- source code: ~OK
 - » /work is better
- papers, thesis, etc in user subdirs is great

- /work

- Tuned for speed, small files
 - » ie. source, binaries, etc.
- NOT backed up
 - » but is resilient
 - » snapshots under `.zfs/snapshot/` for *some* directories
 - » Do **NOT** count on this
- Source code: YES
 - » use git!
- ROOT output: ~ick (don't)
- CODA data: No
- **YOU must backup to tape**
 - » `tar + jput` (*more on this soon*)

Where do I put my stuff?

- /group

- Think “/home” for work groups
 - » papers, thesis, etc
- hourly snapshots
 - » `cd .snapshot/`
- analysis scripts: YES
 - » use git!
- source code: ~OK
 - » /work is better
- papers, thesis, etc in user subdirs is great

- /work

- Tuned for speed, small files
 - » ie. source, binaries, etc.
- NOT backed up
 - » but is resilient
 - » snapshots **may** be available under `.zfs/snapshot/`

PSA: /work snapshots can be a pain because they count towards the quota for that space! (But you can't see them.)

- Generate big files, fill quota, whoops!
 - `rm -rf <all the big files>`
- quota still full!?!
 - Talk to helpdesk... (nothing you can do)

Where do I put my stuff?

- /volatile
 - Largest file system
 - » Petabyte scale
 - High performance for large files
 - » ie. ROOT output
 - NOT backed up
 - Files auto-cleaned based on quota/ reservation/ and filesystem pressure
 - » https://scicomp.jlab.org/docs/volatile_disk_pool
 - Analysis output goes here!
 - » Check, then push to tape if good!
- Tape System
 - Even bigger
 - » 100+ PB and growing
 - /mss/hallX/...
 - » Stubs: shows what is in the tape system!
 - » not the actual files
 - /cache/hallX/...
 - » actual files
 - » auto-clean up in play
 - next slide

Accessing files from Tape

- Retrieving files from tape
 - `jcache get /mss/.../foo.dat`
 - » Manual pull from tape to `/cache/.../foo.dat`
 - » **Never** call this (or `jget`) in a farm script!
 - Let SWIF2 do it!
 - » List needed files as `<Input>` tag(s)
 - » Backend will prestage them for you in advance
 - `jget /mss/.../foo.dat $PWD/`
 - » pull file from tape to any filesystem
 - » generally **not** the right tool

File duration in /cache

Scientific Computing [username](#) [Getting Started](#) [Support](#) [Staff Members](#)

Write-through Cache System 3100 (TB)

Project Usage | jcache Requests | jcache Query | File Pin Info | Usage By User | Small File Usage | File Distribution

Filter

Name	High Quota (GB)	Guarantee (GB)	Pin Quota (GB)	Cached (GB)	NeedTape (GB)	SmallFileCount*	Pinned (GB)
halld	1,550,000	800,000	800,000	1,590,648	14,552	24,855,716	305,610
clas12	1,050,000	500,000	500,000	1,051,307	0	3,028	162,189
halla	400,000	200,000	200,000	321,113	5	60,471	0
hallb	140,000	70,000	60,000	140,170	308	152,732	15,024
halld	130,000	70,000	70,000	102,083	0	769,410	2,818
clas	70,000	35,000	20,000	38,854	0	2,326	0
ceba24gev	5,000	2,000	2,000	0	0	0	0
eic	4,000	2,000	200	1,670	0	2	0
home	3,000	1,000	1,000	1,146	0	136,308	0
accel	2,000	1,000	800	1	0	1,191	0
Sum:	3,354,000	1,681,000	1,654,000	3,246,992	14,865	25,981,184	485,641

- Files auto-cleaned based on quota and system pressure on /cache
 - Clean up least-recently-used files first
 - Can 'pin' files to keep them stable
 - » Shared resource, don't abuse!

Copying files to Tape

- Storing files on tape

- `jput file /mss/.../`

- » 'jput -h'

- » [Online Docs](#)

- 'write-through cache' ([Online Docs](#))

- » write large file output directly to /cache/hallX/...

- no 'staging' on /volatile

- » automagically backed up to tape after a few days

- guaranteed to be safe on tape **before** /cache auto-removal kicks in

- » **Gotchas:**

- small files (<1MB) not backed up to tape


- avoid pathname collisions with files already on tape

- » ie. 'overwriting' files with same name, etc

Now ask Questions!

Hall Computing while Running / Online Analysis

Hall Computing Do's and Don'ts!



Rules to live by so your colleagues
don't curse your name

Donny Don't!

- **Don't** copy or write large files to any “home” directory. This includes
 - CODA files
 - ROOT files
 - tar or zip archives
- When these shared filesystems fill, *many things break*
 - processes fail with corrupt output (not just yours!)
 - machines may require rebooting (disruptive!)
 - DAQ can die!
- **Do use** destinations like:
 - /chafs/work*, /chafs2/work*
 - /net/cdaq/*
 - Use symlinks in local dirs
- **Do use** the tape silo:
 - /mss/cache/hallX/...



Donny Don't!

- **Don't** change global config files in common accounts (*adaq, cdaq, a-onl, cvxwrks, coda, etc*)
 - » ie .bashrc, .cshrc
- **Don't** change environment variables (ie. 'setenv ...')
- **Don't** edit or save files in the global replay directories
- **Don't** run analysis on the DAQ machines
 - All can have unexpected, and difficult to debug impacts on processes running behind the scenes



- **Do use** the appropriate machines and accounts
- **Do use** the JLab Farm:
 - » See next talk and/or
 - » Lunch seminar Wed

Donny Don't!

- **Don't** use buggy/untested code in production environments

→ If your code is:

- » generating enormous 'log' files
- » generating core.NNNN dumps
- » littering hv.1.NNNN.tmp files
- » pegging a core at 100% without good reason
- » spewing warnings / debug info

→ *Then it needs fixing before the experiment*

- **Don't** ignore warning in your replay scripts

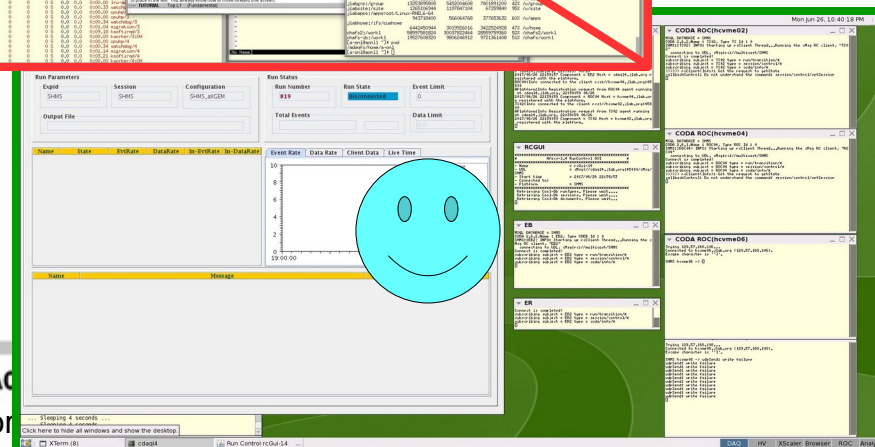
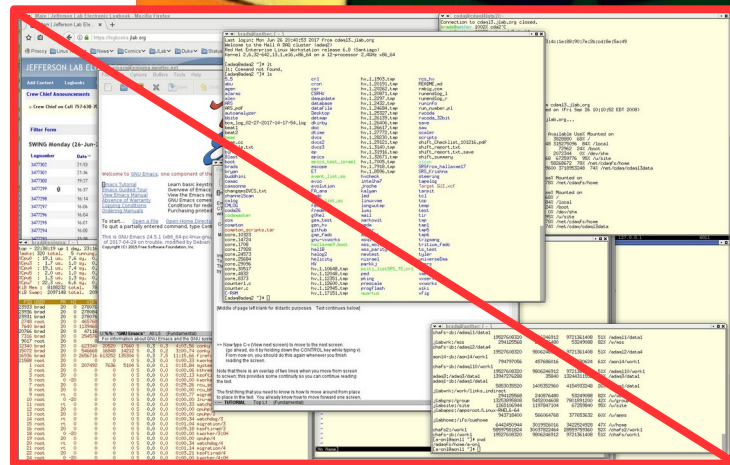
→ *NEVER* disable/hide warnings



- We've lost way too many beam hours (\$\$\$ + data) to such problems.
- Poor saps have to come in at crazy hours to fix things, and hassle folks to clean up.
→ Have mercy on us :-)

Donny Don't!

- Remember the Hall
Computers are shared
machines
 - **Don't** clutter CH screens
with 100s of windows
 - **Don't** shuffle windows
around on Shift Crew
Machines
 - **Clean up** after yourselves
 - » Copy to tape, and/or
remove obsolete files
 - » **NEVER** move or touch
CODA data files though
– talk to Hall expert



Online Help is Available

File Edit View History Bookmarks Tools Help

12 GeV DVCS/GMp Oc... x +

https://hallweb.jlab.org/wiki/index.php/12_GeV_DVCS/GMp_October_2014_r

BTA | HALOG | ELOG | Hall A | Hall A Wiki | Tech On-Call | Fall 2014 Run | DVCS - Hall A Wiki | Gmp - Hall A Wiki

page discussion view source history

12 GeV DVCS/GMp October 2014 running

Important information

General instructions for shift takers

Other useful links

- RC run plan
- How to HRS / DVCS (Detailed documentation and troubleshooting)
- Gmp wiki
- DVCS 12 GeV Wiki
- Main Hall A Wiki
- HALog
- December 2014 BPM calibration

navigation

- Main Page
- Community portal
- Current events
- Recent changes
- Random page
- Help

search

Go Search

toolbox

- What links here
- Related changes
- Upload file
- Special pages
- Printable version
- Permanent link

This page was last modified on 9 December 2014, at 21:56. This page has been accessed 3,799 times. Privacy Disclaimers

• Default web page on CH computers links to “How-tos” and pointers on Shift Crew Duties

Information for Shift Takers

Counting house: 5503, 6666 & 6000 | Run coordinator: 270-8916 | MCC: 7046 & 7047 & 7048 | Crew chief: 7045 | Guard house: 7045

Before the start of your shift, read and understand the Safety Documentation for this experiment.

- You must then sign in the Yellow Binder in the Hall C Counting House!

Detailed instructions are available on the How-to page. If you encounter problems, first look at those instructions for assistance.

Spring 2017 Commissioning Plans

- KPP Run Plan
- Commissioning Plan 2017
- RC daily meetings and daily runplan
- Experts on call

Your Responsibilities

Shift Leader

- Communicate clearly and effectively with shift crews and MCC (7047), and log any status information to the HCLLOG
- Maintain and update the **Shift Summary** detailing the events which occur while on shift
 - This is done by starting the summary as a log entry titled “X Shift Summary” where X = Day, Swing, Owl
 - Save & edit** the summary after **ANY** event of interest occurs
 - This provides the readers with an up-to-date play by play of the current shift
 - See example: <https://logbooks.jlab.org/entry/3402742>
- Keep track of beam time accounting
- Consult the daily run plans and communicate with Run Coordinator whenever problems happen that cannot be solved by shift workers.
- Maintain data taking quality and an efficient use of beam time.
- Follow the directives in the COO and associated Safety Documentation.
- Log the following in a shift summary in the Hall C Logbook:
 - run list (describing the goal of this run: eg production on LH2, BCM calibration...) and report main statistic numbers
 - any major events, including accesses

Target Operator

- Watch the target, see Target Info
- Assist the shift leader and third person with their duties

Third person

- Start and stop the DAQ. Record the purpose for **every** single run in the “Run list” binders. Read *fixme* to create more blank forms.
- Online replay of **all** production runs (twice), as described in Analysis How-to: first replay the first 50 k events, check online plots (see next bullet), then run a full replay.
- Compare replay histograms with the sample ones and report to shift leader any unexplained differences. Hlog them.
- Fill shift checklist once per shift (for guidance, please see the checklist how-to). *fixme* go here to find blank copies of the shift check list and the shift check list how to
- The shift leader can help!

Online Help is Available

File Edit View History Bookmarks Tools Help

12 GeV DVCS/GMp Oc... x +

https://hallaweb.jlab.org/wiki/index.php/12_GeV_DVCS/GMp_October_2014_r

BTA | HALOG | ELOG | Hall A | Hall A Wiki | Tech On-Call | Fall 2014 Run | DVCS - Hall A Wiki | Gmp - Hall A Wiki

page discussion view source history

12 GeV DVCS/GMp October 2014 running

Important information

General instructions for shift takers

Other useful links

- RC run plan
- How to HRS / DVCS (Detailed documentation and troubleshooting)
- Gmp wiki
- DVCS 12 GeV Wiki
- Main Hall A Wiki
- HALog
- December 2014 BPM calibration

navigation

- Main Page
- Community portal
- Current events
- Recent changes
- Random page
- Help

search

Go Search

toolbox

- What links here
- Related changes
- Upload file
- Special pages
- Printable version
- Permanent link

• Default web page on CH computers links to “How-tos” and pointers on Shift Crew Duties

page discussion view source history

Information for Shift Takers

Counting house: 5503, 6666 & 6000 | Run coordinator: 270-8916 | MCC: 7046 & 7047 & 7048 | Crew chief: 7045 | Guard house: 7045

Before the start of your shift, read and understand the Safety Documentation for this experiment.

- You must then sign in the Yellow Binder in the Hall C Counting House!

Detailed instructions are available on the How-to page. If you encounter problems, first look at those instructions for assistance.

Spring 2017 Commissioning Plans

- KPP Run Plan
- Commissioning Plan 2017
- RC daily meetings and daily runplan
- Experts on call

Your Responsibilities

Shift Leader

- Communicate clearly and effectively with shift crews and MCC (7047), and log any status information to the HCLLOG
- Maintain and update the **Shift Summary** detailing the events which occur while on shift
 - This is done by starting the summary as a log entry titled “X Shift Summary” where X = Day, Swing, Owl
 - Save & edit** the summary after **ANY** event of interest occurs
 - This provides the readers with an up-to-date play by play of the current shift
 - See example: <https://logbooks.jlab.org/entry/3402742>
- Keep track of beam time accounting
- Consult the daily run plans and communicate with Run Coordinator whenever problems happen that cannot be solved by shift workers.
- Maintain data taking quality and an efficient use of beam time.
- Follow the directives in the COO and associated Safety Documentation
- Log the following in a shift summary in the Hall C Logbook:
 - run list (describing the goal of this run: eg production on LH2, BCM calibration...) and report main statistic numbers
 - any major events, including accesses

Target Operator

- Watch the target, see Target Info
- Assist the shift leader and third person with their duties

Third person

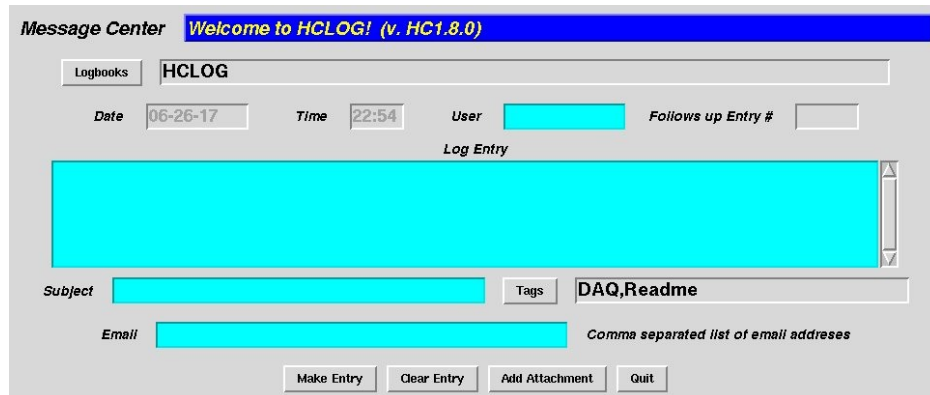
- Start and stop the DAQ. Record the purpose for every single run in the “Run list” binders. Read *fixme* to create more blank forms.
- Online replay of **all** production runs (twice), as described in Analysis How-to. first replay the first 50 k events, check online plots (see next bullet), then run a full replay.
- Compare replay histograms with the sample ones and report to shift leader any unexplained differences. Hlog them.
- Fill shift checklist once per shift (for guidance, please see the checklist how-to). *fixme* go here to find blank copies of the shift check list and the shift check list how to
- The shift leader can help!

- These are editable Wiki's
 - If you see an error, please update it!
 - If you don't, who will?

Logbook Tips

- Always use concise “Subject”
 - » No log entries titled “problem” please...
- Make clear entries that don't rely on being there to make sense
 - » We need to understand the entries months/years later
- Clean up the “Re: Follow-up: Follow-up: Re: Follow-up: ...” garbage
 - » it is just noise
- Manually link to relevant older entries
- Use Tags fields
 - » DAQ, Analysis, Techs, ...

- HALOG, HCLOG are web based:
 - <https://logbooks.jlab.org/book/halog>
 - <https://logbooks.jlab.org/book/hclog>
- On CH computers, can also make entries via standalone tool:
 - » halog / hclog
- **Simpler snapshots!**



The screenshot shows the HCLOG web interface. At the top, there is a blue header bar with the text "Message Center Welcome to HCLOG! (v. HC1.8.0)". Below this, there is a "Logbooks" dropdown menu set to "HCLOG". The form includes fields for "Date" (06-26-17), "Time" (22:54), "User" (a redacted field), and "Follows up Entry #" (an empty field). A large redacted area covers the "Log Entry" field. Below this, there is a "Subject" field (redacted), a "Tags" dropdown menu set to "DAQ,Readme", and an "Email" field (redacted) with the note "Comma separated list of email addresses". At the bottom, there are four buttons: "Make Entry", "Clear Entry", "Add Attachment", and "Quit".

Logbook Tips

- HALOG, HCLOG are web based:

→ <https://logbooks.jlab.org/book/halog>
→ <https://logbooks.jlab.org/book/hclog>

- Help is worth reading

- Useful 'tricks'

→ Use UPPER-CASE boolean logic in search

» ie. AND *not* and

→ Display Settings:
Hide Autologs

→ Useful Links

JEFFERSON LAB ELECTRONIC LOGBOOK

Logged in as brads (Logout)

Add content Logbooks Tags Useful Links Preferences Help/About

Crew Chief Announcements

- o Crew Chief on Call 757-630-7050

HCLOG

Filter From

DAY Friday (23-Jun-2017)

Lognumber	Date	Author	Title
3477165	14:47	ethanb	Follow-up Re: Chamber gas updates -- gas system back on

SWING Wednesday (21-Jun-2017)

Lognumber	Date	Author	Title
3476902	22:17	brads	Chamber gas updates -- gas system back on

SWING Tuesday (20-Jun-2017)

Lognumber	Date	Author	Title
3476684	15:03	beaufait	Follow-up Re: Follow-up Re: Chamber gas interlock problem

DAY Friday (16-Jun-2017)

Lognumber	Date	Author	Title
3476164	14:50	beaufait	Follow-up Re: Chamber gas interlock problem
3476105	09:02	pooser	Follow-up Re: GEM SRS DAQ standalone latency scan
3476104	08:56	pooser	Follow-up Re: Latency scan for GEM SRS DAQ in integrated Hall C system

OWL Friday (16-Jun-2017)

Lognumber	Date	Author	Title
3476094	01:05	latif	Latency scan for GEM SRS DAQ in integrated Hall C system
3476093	00:56	latif	GEM SRS DAQ standalone latency scan

SWING Thursday (15-Jun-2017)

Lognumber	Date	Author	Title
3476028	16:11	ethanb	HMS Quadrupoles

SWING Wednesday (14-Jun-2017)

Lognumber	Date	Author	Title
3475920	15:02	brads	Chamber gas interlock problem

DAY Wednesday (14-Jun-2017)

Date Picker

June 2017

Su	Mo	Tu	We	Th	Fr	Sa
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	

Date Range

Display Settings

Autorefresh

Useful Links

- o --> Hall C Shift Instructions <--
- o --> Hall C Wiki <--
- o Access Keys
- o ATLAS
- o CEBAF Status Information
- o Hall C EPICS archive
- o Hall C Run Entries -- End
- o Hall C Run Entries -- Start
- o Hall C Screens
- o Hall C Target [Hclog, Target tag]
- o Hall C Target [Target log, Hall C tag]
- o Hall C Tasklist (HCList)
- o Old Hall C Logbook (pre-2014)
- o RF Dashboard
- o Tech Notes
- o White Board (OPS)

And, Most Importantly

- If you see something you don't understand...
 - Ask someone
 - Make a log entry
 - Dig in and beat on the problem until it makes sense to *you*
 - » You're scientists – understanding weirdness is literally our business!



"Notice all the computations, theoretical scribbles, and lab equipment, Norm. ...
Yes, curiosity killed these cats."

Now ask More Questions!