CLAS12 Software Readiness Review

Jerry Gilfoyle (for the CLAS12 Software Group)

Physics Department, University of Richmond, Virginia

Outline:
1. Introduction
2. Software Framework
3. Management
4. Requirements
5. Manpower
6. Conclusions

Introduction

- CLAS12 - Large acceptance spectrometer based on CLAS6.
- Luminosity increases by a factor of ten over CLAS6 ($L = 10^{35}\ cm^{-2}s^{-1}$).
- Software Goal:

  Ready to analyze data at turn on (October, 2014).

- Software development is far along.
- Planning has been ongoing.

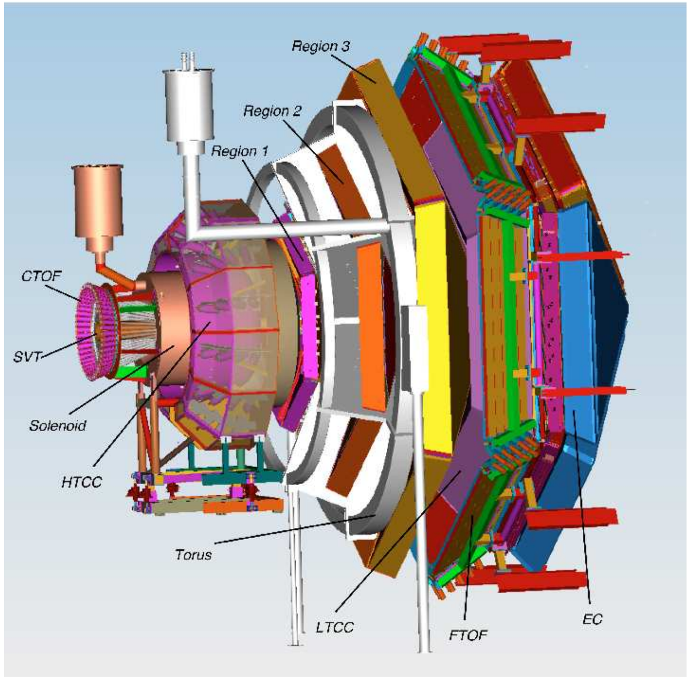

	Forward	Central
θ_{track}	$5^\circ - 40^\circ$	$35^\circ - 135^\circ$
θ_{photon}	$2.5^\circ - 40^\circ$	
$\Delta p/p$	< 0.01	< 0.05
$\Delta\theta$	$< 1\ mr$	$< 10 - 20\ mr$
$\Delta\phi$	$< 3\ mr$	$< 3 - 5\ mr$

Software Framework

- CLAS12 Reconstruction and Analysis Framework (CLARA)
  - Service Oriented Architecture - build/maintain complex, distributed software system.
  - Example: CERN Technical Infrastructure Monitoring.
- SCons
  - Open Source software construction tool.
  - Improved, cross-platform substitute for Make.
- SVN
  - Open source software versioning and revision control.
  - Successor to CVS.
- Already adopted SCons and SVN for CLAS6.

CLARA

- Service Oriented Architecture (SOA) for physics data processing - multi-threaded, distributed environment.
- The fundamental unit is the ‘Service’ - physically independent software programs with a common interface.
- Services are loosely coupled, and may participate in a variety of algorithms.
- Interface.
  - Specifies a set of methods an object can perform but not the implementation of those methods.
  - Promote flexibility and reusability in code by connecting objects in terms of what they can do rather than how they do it.

CLARA/SOA Example - 1

- Service Integration Unit - allows user applications to be presented as CLARA services.
- PCEP layer - Physics Complex Event Processing.
- Services can originate on different nodes.

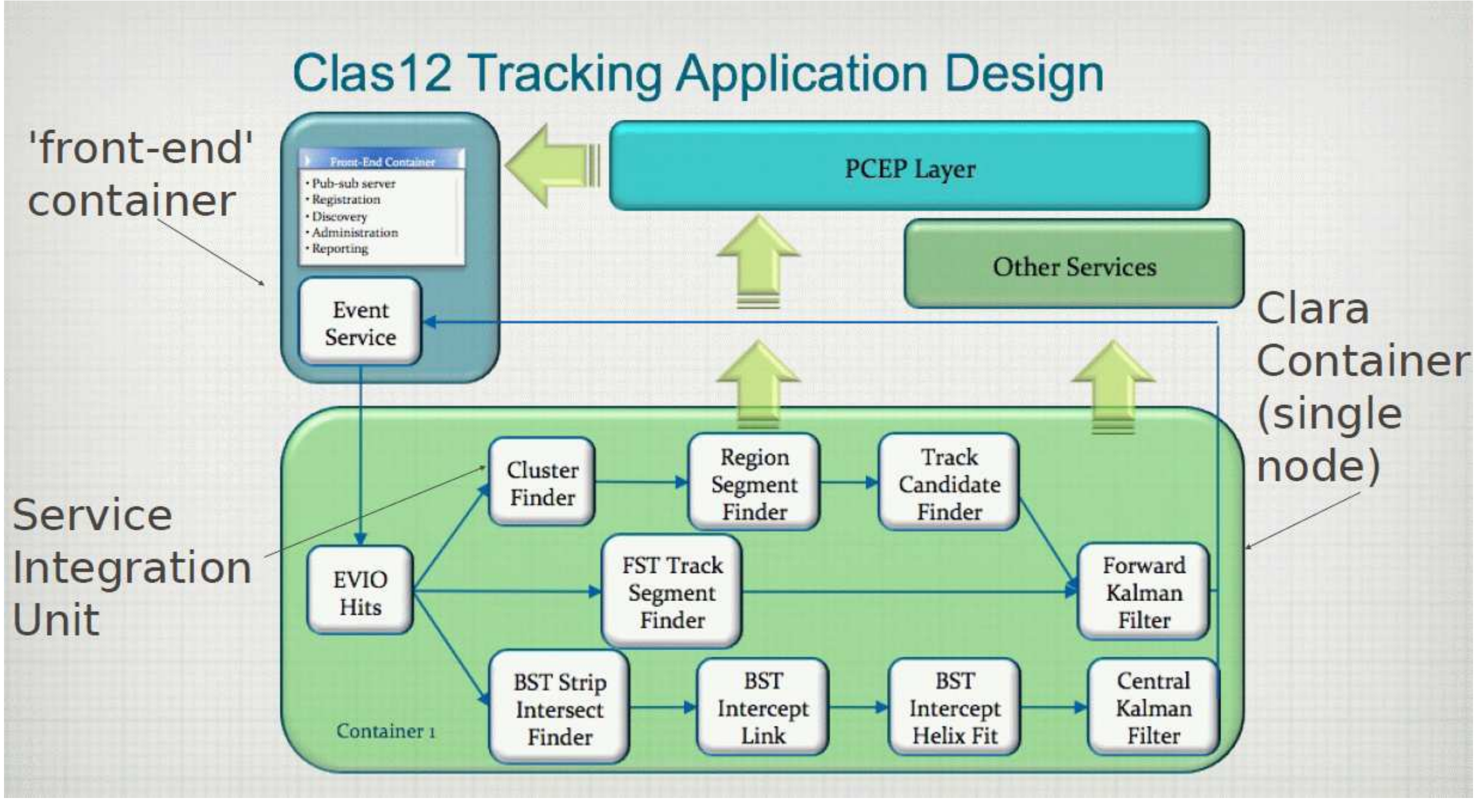

CLARA/SOA Example - 2

- Testing the CLAS12 tracking service.
- Tested on Spiderwulf - University of Richmond Nuclear and Astro-Physics Cluster: 17 nodes, Xeon, 2×6 Westmere nodes.
- Electron events generated from CLAS12 simulation *gemc*.

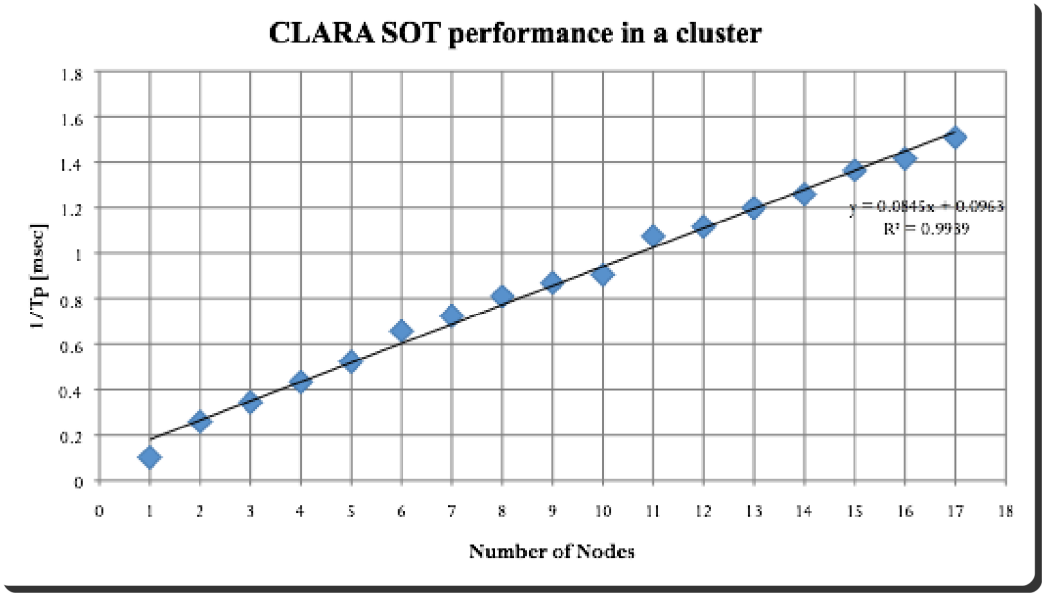

CLARA/SOA Example - 2

- Testing the CLAS12 tracking service.
- Tested on Spiderwulf - University of Richmond Nuclear and Astro-Physics Cluster: 17 nodes, Xeon, 2×6 Westmere nodes.
- Electron events generated from CLAS12 simulation *gemc*.

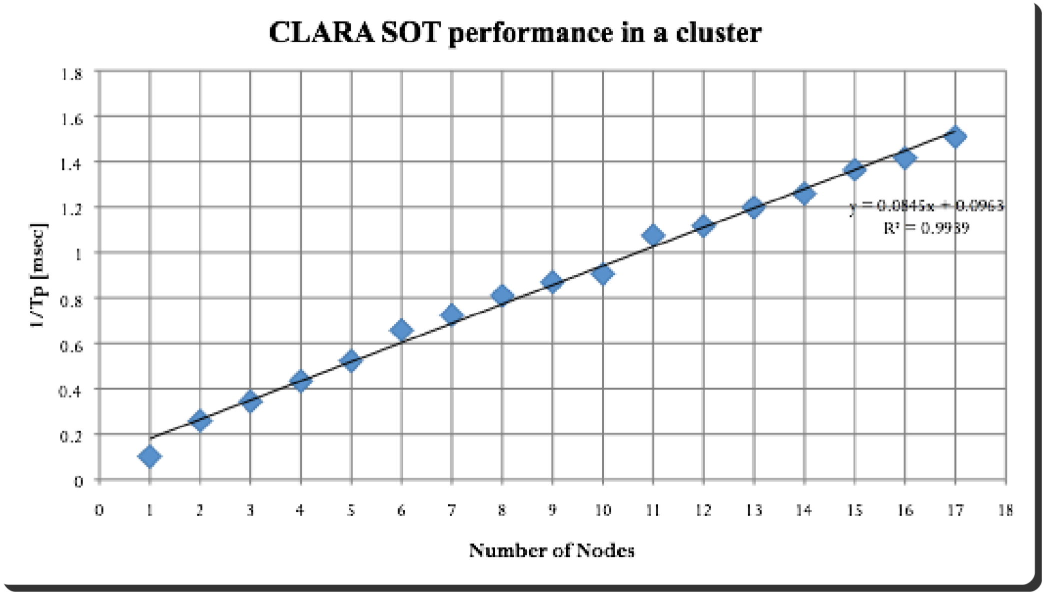

ONGOING ISSUE: CLARA access at JLab blocked by security barriers.

Management

- CLAS12 Software Group (leader: Dennis Weygand).
- Wiki

- Tutorials to set up services in C++ and Java.
- Collaborations with Hall D and DAQ group.
  - JANA
  - Database
  - Event display
  - EVIO
  - cMsg

Management

- Tools:
  - Interfaces of calibration database to JANA, C++, MySQL, ...
  - Simulation: *gemc*
  - Bug reporting - Mantis
- Policies in place:
  - Regular builds of CLAS6 and CLAS12.
  - test histograms.
  - data challenges.

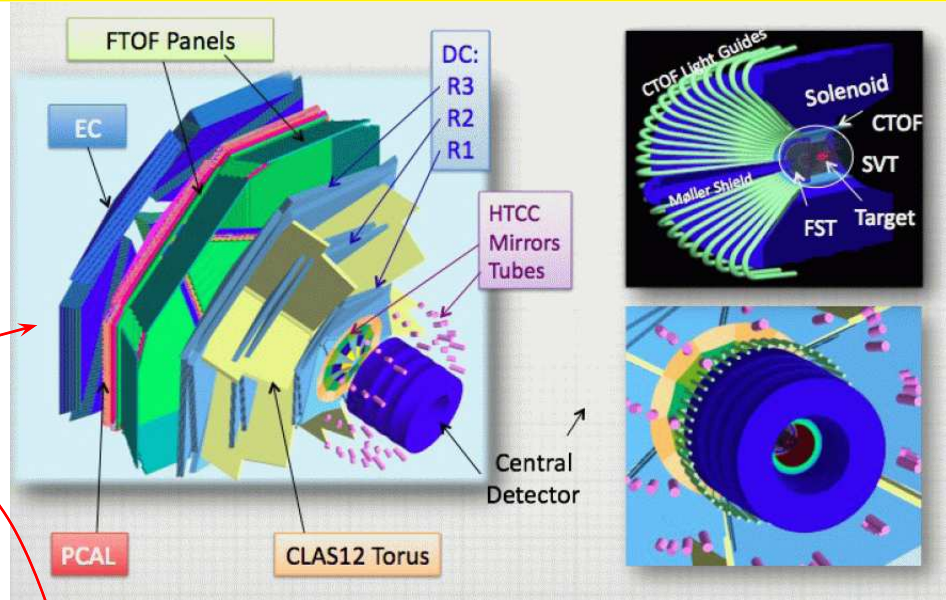

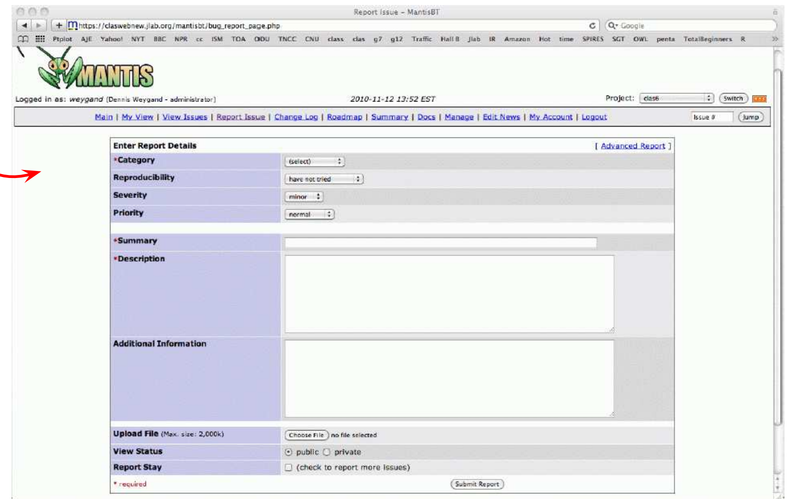

Requirements

- Calculations follow format from Graham Heyes.
- Assume an October, 2014 start date.
- Will present the major assumptions and results for:
  - data acquisition
  - calibration
  - simulation
  - reconstruction (formerly analysis in spreadsheet)
  - analysis

Ratio of Simulation:Reconstruction:Analysis

Process hours from Computer Center for the last year.

Project View | Job Type View | Job Status View | Job Walltime View

Farm Job Stat (Job Type View)

From: 2010-04-19 (yyyy-mm-dd) To : 2011-05-19

Process Hour Summary (by org)

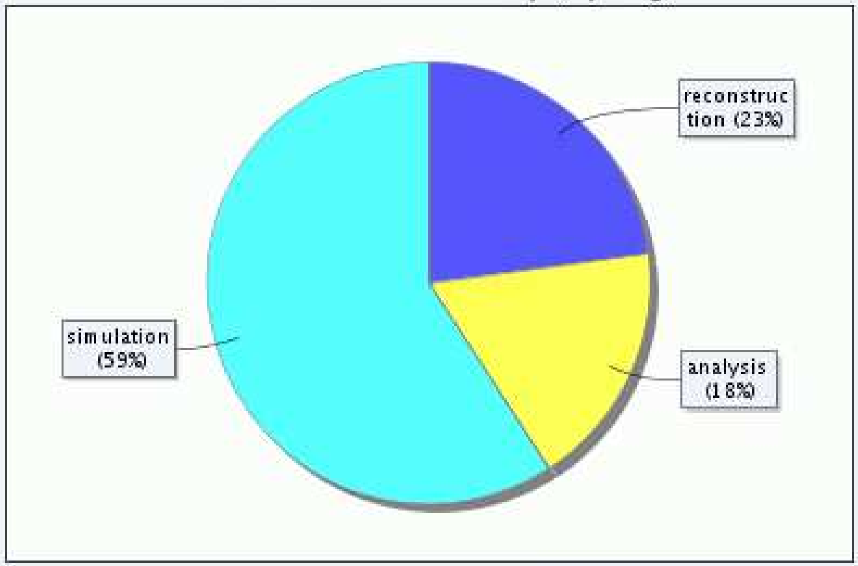

Job Type	Total Job Count	Used Process Hours
debug	36,432	5,511.1
reconstruction	1,182,092	1,557,314
one_pass	9,441	1,647.5
analysis	2,002,900	1,221,338.2
test	36,432	5,511.1
simulation	1,046,276	4,033,710.9
Total	4,313,573	6,825,032.7

Last updated: Thursday May 19, 11:14:03 EDT 2011

Data Acquisition

Event rate	10 kHz	Weeks running	30
Event size	10 kBytes	24 hour duty factor	60%

$$\text{Data Rate} = \text{Event Rate} \times \text{Event Size} = 100 \text{ MByte/s} \tag{1}$$

$$\text{Average 24-hour rate} = \text{Data Rate} \times \text{24 hour duty factor} = 60 \text{ MByte/s}$$

$$\text{Events/year} = \text{Event Rate} \times \text{Weeks Running} \times \text{24 hour duty factor} = 1.1 \times 10^{11} \text{ Events/yr} \tag{2}$$

$$\text{Data Volume/year} = \text{Events/year} \times \text{Event size} = 1100 \text{ TByte/yr} \tag{3}$$

Calibration - 1

CPU-time/event-core	67 ms	High-priority fraction	1%
Data passes	5	Desired processing time	20 min
Data fraction used	10%	Output size/input set size	5
Data set size	2 GB	Event size	10 kBytes

$$\begin{aligned}\text{Events/priority data set} &= \frac{\text{Data set size}}{\text{Event Size}} \\ &= 2 \times 10^5 \text{ events}\end{aligned} \tag{4}$$

$$\begin{aligned}\text{CPU time/priority data set} &= \text{Events/priority data set} \times \\ &\quad \text{CPU-time/event-core} \\ &= 1.3 \times 10^4 \text{ s}\end{aligned} \tag{5}$$

$$\begin{aligned}\text{Cores/data set for priority} &= \frac{\text{CPU time/priority data set}}{\text{Desired processing time}} \\ &= 11 \text{ cores}\end{aligned} \tag{6}$$

Calibration - 2

CPU-time/event-core	67 ms	High-priority fraction	1%
Data passes	5	Desired processing time	20 min
Data fraction used	10%	Output size/input set size	5
Data set size	2 GB	Event size	10 kBytes

$$\begin{aligned}\text{Output data set} &= \text{Data set size} \times \text{Output size/input set size} \\ &= 10 \text{ GByte}\end{aligned} \tag{7}$$

$$\begin{aligned}\text{Non-priority CPU time/year} &= \text{Events/year} \times \text{CPU-time/event-core} \times \\ &\qquad \text{Data fraction used} \times \text{Data passes} \\ &= 3.6 \times 10^{9}\text{s}\end{aligned} \tag{8}$$

$$\begin{aligned}\text{Cores for non priority} &= \frac{\text{Non-priority CPU time/year}}{\text{one year in seconds}} \\ &= 116 \text{ cores}\end{aligned} \tag{9}$$

Simulation - 1

CPU-sim-time/event-core	400 ms	Fraction to disk	2%
Sim-events/year	7×10^{10}	Fraction to tape	10%
Output event size	50 kBytes		

$$\begin{aligned}\text{CPU-time/year} &= \text{CPU-time/event-core} \times \text{Sim-events/year} \\ &= 3 \times 10^{10}\ \text{s}\end{aligned} \tag{10}$$

$$\begin{aligned}\text{Dedicated farm cores} &= \frac{\text{CPU-time/year}}{\text{one year in seconds}} \\ &= 828\ \text{cores}\end{aligned} \tag{11}$$

$$\begin{aligned}\text{Work disk} &= \text{Sim-events/year} \times \text{Output event size} \times \\ &\quad\ \text{Fraction to disk} \\ &= 65\ \text{TBytes}\end{aligned} \tag{12}$$

Simulation - 2

CPU-sim-time/event-core	400 ms	Fraction to disk	2%
Sim-events/year	7×10^{10}	Fraction to tape	10%
Output event size	50 kBytes		

$$\begin{aligned} \text{Tape storage} &= \text{Events/year} \times \text{Output event size} \times \\ &\quad \text{Fraction to tape} \\ &= 326 \text{ TBytes/year} \end{aligned} \tag{13}$$

$$\begin{aligned} \text{Average bandwidth} &= \frac{\text{Output event size} \times \text{Dedicated farm cores}}{\text{CPU-sim-time/event-core}} \\ &= 104 \text{ MByte/s} \end{aligned} \tag{14}$$

Reconstruction - 1

CPU-data-time/event-core	67 ms	Output size/input size	2
Data passes	1.7	Output fraction on work disk	10%
Event Size	10 kBytes	Events/year	1.1×10^{11}
Data volume/year	1.1 PBytes/yr		

$$\begin{aligned}\text{CPU time per year} &= \text{Events/year} \times \text{CPU-data-time/event-core} \\ &\quad \times \text{Data passes} \\ &= 1.2 \times 10^{10}\ \text{s}\end{aligned} \tag{15}$$

$$\begin{aligned}\text{Dedicated farm cores} &= \frac{\text{CPU time per year}}{\text{one year in seconds}} \\ &= 393\ \text{cores}\end{aligned} \tag{16}$$

$$\begin{aligned}\text{Cooked data to tape} &= \text{Data Volume/year} \times \text{Data passes} \\ &\quad \times \text{Output size/input size} \\ &= 3700\ \text{TByte/yr}\end{aligned} \tag{17}$$

Reconstruction - 2

CPU-data-time/event-core	67 ms	Output size/input size	2
Data passes	1.7	Output fraction on work disk	10%
Event Size	10 kBytes	Events/year	1.1×10^{11}
Data volume/year	1.1 PBytes/yr		

$$\begin{aligned}\text{Disk storage} &= \frac{\text{Cooked data to tape}}{10} \\ &= 370 \text{ TByte}\end{aligned} \quad (18)$$

$$\begin{aligned}\text{Average bandwidth} &= \text{Event size} \times (1 + \text{Output size/input size}) \times \\ &\quad \frac{\text{Dedicated farm cores}}{\text{CPU-data-time/event-core}} \\ &= 176 \text{ MBytes/s}\end{aligned} \quad (19)$$

Post-Reconstruction Analysis

CPU-data-time/event-core	67 ms	Fraction of desired events	20
Data passes	10	Work disk space	370 TBytes
Tape storage	370		

$$\begin{aligned} \text{CPU time per year} &= \text{Fraction desired} \times \text{Events/year} \times \\ &\quad \text{CPU-data-time/event-core} \times \\ &\quad \text{Data passes} \\ &= 1.5 \times 10^{10} \text{ s} \end{aligned} \quad (20)$$

$$\begin{aligned} \text{Dedicated farm cores} &= \frac{\text{CPU time per year}}{\text{one year in seconds}} \\ &= 463 \text{ cores} \end{aligned} \quad (21)$$

Requirements Summary

	Cores	Disk (TByte)	Tape (TByte/yr)
DAQ	-	-	1100
Calibration	127	-	-
Simulation	828	65	327
Reconstruction	393	370	3700
Analysis	463	370	370
Sum	1811	805	5497

CLAS/CLAS12 Software Manpower (Preliminary)

	Function	Name
1	Management and Framework (CLAS)	Weygand, Gyurjyan, Heddle
2	Management and Framework (others)	Wolin, Lawrence, Abbott, Timmer, Lee
3	Core Developers (CLAS)	Ungaro, Gilfoyle, Wood, Procureur, Goetz
4	Developers (undergraduates)	Paul, Carbonneau, Frasier, Moog, Musalo, ...
5	Users	≈ 10 FTEs listed in SoS statements

Names listed in rows 1-3 provide ≈ 5 FTEs focused on CLAS12 software.

Conclusions

- Software framework is being developed; considerable progress in last year - CLARA, svn, SCons.
- Management tools are in place and a core group exists - exploiting overlaps with DAQ and Hall D.
- DAQ - $\approx 10^{11}$ events/year $\rightarrow 1$ petabyte/yr.
- Calibration - about 130 cores required.
- Simulation - 276 cores required with 109 TBytes/year of tape storage.
- Reconstruction - about 400 cores required with 3.7 PByte/yr of cooked data to tape.
- Post-reconstruction analysis - about 460 cores required.
- Manpower is still limited to a small group of core developers.

Conclusions and Questions

- Software framework is being developed; considerable progress in last year - CLARA, svn, SCons.
- Management tools are in place and a core group exists - exploiting overlaps with DAQ and Hall D.
- DAQ - $\approx 10^{11}$ events/year $\rightarrow 1$ petabyte/yr.
- Calibration - about 130 cores required.
- Simulation - 828 cores required with 91 TBytes/year of tape storage.
- Reconstruction - about 400 cores required with 3.7 PByte/yr of cooked data to tape.
- Post-reconstruction analysis - about 460 cores required.
- Manpower is still limited to a small group of core developers.

- Ratio of simulated events to data collected?
- Speed of simulation?
- Effect of user computing resources?

Additional Slides

Simulation - 1b: Sim events = data events

CPU-sim-time/event-core	400 ms	Fraction to disk	2%
Sim-events/year	10^{11}	Fraction to tape	10%
Output event size	50 kBytes		

$$\begin{aligned}\text{CPU-time/year} &= \text{CPU-time/event-core} \times \text{Sim-events/year} \\ &= 4.4 \times 10^{10}\ \text{s}\end{aligned} \tag{22}$$

$$\begin{aligned}\text{Dedicated farm cores} &= \frac{\text{CPU-time/year}}{\text{one year in seconds}} \\ &= 1,381\ \text{cores}\end{aligned} \tag{23}$$

$$\begin{aligned}\text{Work disk} &= \text{Sim-events/year} \times \text{Output event size} \times \\ &\quad \text{Fraction to disk} \\ &= 109\ \text{TBytes}\end{aligned} \tag{24}$$

Simulation - 2b: Sim events = data events

CPU-sim-time/event-core	400 ms	Fraction to disk	2%
Sim-events/year	2×10^{10}	Fraction to tape	10%
Output event size	50 kBytes		

$$\begin{aligned}\text{Tape storage} &= \text{Events/year} \times \text{Output event size} \times \\ &\quad \text{Fraction to tape} \\ &= 544 \text{ TBytes/year}\end{aligned} \tag{25}$$

$$\begin{aligned}\text{Average bandwidth} &= \frac{\text{Output event size} \times \text{Dedicated farm cores}}{\text{CPU-sim-time/event-core}} \\ &= 173 \text{ MByte/s}\end{aligned} \tag{26}$$