LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY

LIGO Laboratory / LIGO Scientific Collaboration

LIGO-T080243-00-C *LIGO* November 20, 2008

Advanced LIGO CDS DAQS Design Review Report

CDS DAQS Review Committee, edited by V. Sandberg (chair)

Distribution of this document:
LIGO Science Collaboration

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project – MS 18-34
1200 E. California Blvd.
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@LIGO.caltech.edu

Massachusetts Institute of Technology
LIGO Project – NW17-161
175 Albany St
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@LIGO.mit.edu

LIGO Hanford Observatory
P.O. Box 159
Richland WA 99352-0159
Phone 509-372-8106
Fax 509-372-8137

LIGO Livingston Observatory
P.O. Box 940
Livingston, LA 70754
Phone 225-686-3100
Fax 225-686-7189

http://www.LIGO.caltech.edu/

Document numbering and revision history

LIGO-T080243-00-C	original draft	2008 November 17
LIGO-T080243-00rv2-C	committee’s additional comments added	2008 November 20
LIGO-T080243-00rv3-C	final edit completed and checked in to DCC	2009 January 8

Advanced LIGO CDS DAQS Design Review Report

LIGO-T080243-00-C

Review held 2008 September 18 by teleconference.

Presentation materials posted on the advLIGO "wiki" page at the URL:

http://lhocds.LIGO-wa.caltech.edu:8000/advLIGO/CdsDesignDocuments

Review Committee Membership

Stuart Anderson, Matthew Evans, Michael Landry, Brian O'Reilly, Keith Riles, Daniel Sigg, Joshua R. Smith, John Zweizig, Kent Blackburn (cybersecurity), William Tyler (safety), and, Vern Sandberg (chair)

Introduction

This report presents our comments, opinions, and recommendations on the proposed advLIGO CDS Data Acquisition System. The scope of work is as defined in the advLIGO WBS. It is the responsibility of the CDS group to specify the data conversion hardware (ADCs and DACs), computer interfaces, and computing platforms for the other advLIGO subsystems (e.g., SUS, ISI, and ISC). It is the responsibility of the CDS group to provide and maintain the data acquisition computing systems that acquires data from the subsystems, assembles the data into frames, and provides the infrastructure to distribute the data to the DCS archive group. CDS is also charged with providing the infrastructure services, both hardware and software, such as clock signals and clock signal distribution, GPS referenced timestamps, machine controls in the form of EPICS (slow controls), NDS (realtime data flowing in the DAQ system and a local repository of "trend data"), and computing platforms for the servo systems that control the interferometers.

Much of the design has been based on the experience gained with initial LIGO. A new part is the development of a software toolkit to engage "users" in the development of DAQ code specific to their subsystems. New "tools" are being developed to supplement the venerable control room tools data viewer and dtt. New software also is coming from outside the core group in CDS.

The presentation reflected the vigorous state of development and in general *the view by the review committee is very positive*.

A Re-Interpretation of this Review from PDR to FDR

The CDS DAQ system is very dependent upon current and projected computer technology. Designs necessarily make assumptions on technology development to define capabilities and costs. It was proposed that this review be interpreted as a partial "final design", with changes and additional pre-procurement reviews being made as technology settles. From this point of view the question to ask as a reviewer is: Is the proposed system (CDS DAQS) defined and mature to the degree that the proposers can be given the approval to proceed and move to the procurement stage. There will be "Procurement Readiness Reviews" on a regular basis, so the approval does not "open the flood gates" to spending. The Readiness Review(s) can be used to communicate choices and selections, and provide a mechanism for feedback and comment.

The review committee is not comfortable with this policy. Members expressed a strong belief that more detailed descriptions of key subsystems needs to be in place before such an overall approval can be given. For example, the DAQ software architecture needs to be documented in more detail. A top down presentation (e.g., good set of flow charts or similar diagrams from a functional level) is needed and a review program instituted and followed as the system takes shape. Generic hardware and software (e.g., network communication systems and computers) purchases should be made in a prudent and timely manner as overseen by senior Laboratory management.

General Comments

A good, clean, elegant design was presented, but not enough top level detail was available to fully judge it. Much of the software design was inherited from initial LIGO, which provides a stable starting point. As work proceeds to create the system outlined, now is a good time to produce documentation as components are created. All designs, both circuits and software, must be documented and reviewed (at some level).

There are three separate operating groups within the LIGO Laboratory that deal with information coming directly from the interferometers: CDS, DCS, and GDS. CDS has direct responsibility for digitizing and collecting the data and providing infrastructure to control the interferometers. CDS also collects data from ancillary systems such as the facilities management equipment (HVAC, room temperatures, etc.) and the PEM system. CDS provides access of the collected data to DCS and GDS. CDS systems receive stimulus signals from GDS and the broader LSC community in the form of excitations and simulated waveforms (for calibrations, test injections of waveforms, and blind injections). CDS has a responsibility to protect the interferometers. Injected signals must be carefully controlled so as to create the expected effects (actuation on a test mass) and so as to not damage the instruments or corrupt data. As LIGO has grown from the original organizational structure, the subdivisions have changed, blurring their responsibilities. A clear identification of responsibilities, ownership, and methods of communicating between CDS, DCS, GDS, Calibration, PEM, ... is needed. This is not the responsibility of CDS alone, but extends across CDS, DCS, and GDS and should be overseen by LIGO management.

Action Items

1. One set of questions arose during the discussion of the Frame Writer and the interface between CDS and DCS.

> i. What is the maximum data rate? Total data rate? Data rate per IFO? Data Rate for sub-frames, such as IFO channels, PEM channels, EPICS data, filter bank and other parameter data?
>
> ii. Questions about modification of the "frame": How many frames? How to break up the frame(s)? Where to package the EPICS data? The filter parameter data? The IFO channels? The question about the frame specification and the inclusion of filter coefficients, parameters, other EPICS data, etc. came up here. The use of the "static frame" in the frame definition was raised. All of this flows back to the first item listed above on the maximum data rate and the size of the frame(s).

These questions should be resolved in consultation with DCS and GDS.

2. A second set of questions arose during the discussion of what contribution the LSC could make to CDS activities. The creation of the RCG told to develop sub-system

specific DAQ software raised the issue of the setting up of a working group to define the detailed design and to provide resources for the necessary code reviews. This effort would draw from LIGO staff outside of CDS and the LSC. The action item would be to define and setup the working group.

3. LSC Engagement

The LSC is providing mDV and LIGODV, which are Matlab-based tools for getting and processing data through the NDS server. mDV is a "data getter and analysis" library for use with MatLab or similar analysis packages that has the potential to augment the present control room tools, dtt and data viewer. The development of Matlab-based tools is being carried on outside of CDS oversight and control, unlike RCG. An MOU or similar mechanism needs to be set up to define contracted development work so we can be assured of a well defined deliverable.

Detailed Comments

Network traffic:

EPICS and BURT were suspected of producing "glitches" at LLO during S5. The loading reported in section 8.2 from local frame writing is a reminder that the FE processors are touchy. Moving from hourly backup to quasi-realtime is attractive during commissioning - the faster the better, but a problem for science-mode running. A throttled-back BURT could be used during a science run. The capability of adapting the archiving and restore capabilities to the running mode would be useful.

Is there a planned upgrade of PEM monitoring for AdvLIGO? PEM supplies most of the channels of interest to our outreach customers who are not involved in gravitational wave searches. These channels are of interest to a larger scientific community as well.

Data Monitoring and Quality

Consider adding additional information on data integrity checks, e.g., explain why 2 FrameWriters, and if used for data validation how that will be done. It might also be useful to track the life of a data sample from the ADC all the way to a frame written to disk for archiving, describing all the data integrity checks made along the way? For example, do we still have ADC overflow bits being recorded? Are all buffers time tagged so we never have to worry about out of order transfer due to bugs in network device drivers, for example? Are all buffers initialized with easily identified patterns so that a failure to properly fill a new buffer does not send old data through with the appearance of normality?

DAQS Network concerns:

Packet collisions on the DAQS network, especially when broadcasting data will make it necessary to keep up with the DCB. With the GDS broadcast network we found that any extra traffic would screw up the data flow. We also had to move to 9kB packets (jumbo frames) to minimize the frame overhead. With faster computers, the problems are decreased considerably, but our experiences should be considered in increasing the error margin in case of transient load increases.

Two memory buffers seems to be a bare minimum. Wouldn't increasing the number of memory buffers to 3-4 and decoupling the writers via a queue mechanism be safer?

Testing and Monitoring:

The DAQ system, in the advent of real GW detections, has to be free of artifacts. Since the same software runs on all three interferometers one can easily imagine a coherent glitch which can be confused with a real signal. In particular, I think there needs to be special effort in testing and development to guarantee that there are no undetected transmission errors, that all data is processed in time and arrives at its destination without exception, that no hidden glitches can be introduced by the controls and excitation processes, that hardware errors do not go undetected and finally that software algorithms are correct and stay correct as the system is commissioned and improved.

There need to be plans and support for testing after the Final Design Phase, i.e., during installation and commissioning and beyond. We recommend establishing a permanent dedicated test system. The recent roll out of the new DAQS system at LHO would have benefited from a test system at Caltech.

CDS should develop a set of documents and reviews which addresses the lessons learned in initial LIGO. We don't want to see the same problems again. Most problems in initial LIGO seem to stem from scaling up a simple system. The fact that there is a term "bootfest" is itself a problem. But, we also had problems with the EPICS system not being able to handle the requested size, with links going up and down and bringing other systems with them, with discrepancies between what we thought the state of a controls variable is and what it actually was, with software updates which couldn't be tested with anything but the full system, with data corruption caused by unrelated modules, with hardware modules not performing to specs or getting into unknown states, and with maintaining obsolete software and hardware.

Code Reviews

We would like to see an equivalent of a "code walk-through" for the major components of the control systems; e.g., the differential arm length data path. In particular, we would like to know more about the software which accesses the hardware, the logic behind time stamping, the algorithms which are used to process the data, and the network layer which transports the data to the other end. I would like to see a worst case analysis of the computer and network timing. And finally, we would like to see what kind of testing capabilities will be built-in and what kind of diagnostics will be available to guarantee performance during operations. The time frame for this task would be the procurement reviews for the sub-systems.

It is a general impression that there are no more than one or two people that are truly knowledgeable about these systems. We believe this exposes LIGO to an unacceptable risk. We recommend that the circle of people who develop software for the DAQ system, are able to install the software, or perform testing is significantly enlarged.

We also have a different paradigm of software development this time around, as it will be mostly scientists and engineers using the CDS code generator to build applications and the CDS core group developing the underlying ’generic’ code. This leads to two levels of code review and test. We have built, and continue to build, test code for the underlying software. The hope is to automate this as much as possible, where as code changes are made, the latest rev is run against automatic test systems to make sure all the pieces still work. The primary focus here will be for the DAQ and networking software, as they are perhaps the simplest to test in an automated fashion. Testing all the possible combinations of applications built with the code generator will be a more daunting task and needs further thought.

Script Management

The premise, "Various scripts, based on EPICS channel access, have been developed for LIGO to automate operator tasks" is not quite correct. Many of the scripts also rely on access to quasi-real-time data from the NDS. This represents a non-negligible network and processor load as the scripts perform filtering and demodulation for 16kHz data and should be kept in mind when choosing a script server (or servers).

There are at least 2 different kinds of scripts. Some run continuously and perform sub-system automation (e.g., the MC auto-locker) by responding to requested and unrequested changes in state (e.g., lock and unlock). Other scripts are task based, generally short lived, and are run as needed (e.g., angle to length decoupling). We may benefit from distinguishing between these two types.

For instance, we suggest a scripting system in which most (if not all) sub-systems have a "guardian" script which runs continuously and is automatically restarted in case of trouble. Each guardian would have a set of known system states and sub-scripts for performing requested state changes, and responding to unrequested ones. Guardians would also have a set of known tasks that can be performed in each state (these could be treated as sub-states which last as long as the task requires to be completed).

The system described is similar to the one that Virgo has used successfully for a few years. The primary difference is that Virgo developed its own scripting language (Alp) to run on its script servers. The motivation for this was, more or less, to provide specialized data access commands, but we think we have already achieved that capability in a more elegant way with our use of standard scripting languages and small binaries for data access.

All scripts must be in a version control system.

LSC Developed Software and Control Room Tools

We were pleased to see LIGODV specifically mentioned as an operator tool for Advanced LIGO. Will the control room workstations be standardized, say to CentOS 5? Will a similar "standardization" apply to other software packages, such as defining a default version of Matlab? CDS should create a management plan for a "default" Matlab release for operator stations that could change (say updating once a year - using half-year old releases previously tested to be suitable for the applications) and a Matlab license plan that would ensure that there are sufficient licenses to allow operators to run several Matlab instances.

The definitions of channel names needs to be defined or re-defined and firmly enforced. Facilities channels have to get into the frames.

The excitation/stimulus awg system, test points, and the dtt tools need to be stabilized into a standard release. Logs of excitations need to be recorded. Filter files, both IIR and FIR, need a standardized filter format definition and suitable software tools need to be either upgraded or developed to manage these critical files. All files must be archived in a suitable repository and be traceable as to date and time of use.

NDS and NDS2 Support

CDS should lead and define the use and inclusion of EPICS slow channels in the NDS and NDS2 (GDS sponsored) network data servers. (NDS2 is not in CDS's scope and properly belongs within the responsibility of GDS and DCS. CDS DAQS should make every effort to accommodate NDS2 in supplying data to the degree it can do so safely without compromising the interferometers.) NDS2 is very important for opening up LIGODV, mDV, noisebudget and other software to be used by collaboration members.

Verification of running codes and filters, control data states, etc.:

It must be possible to completely establish what code (front-end, filters, parameters etc) was running during a given science mode epoch. How will this be done? By tagged SVN versions? What about Astrowatch segments, i.e. AdL prior to, or in between, science runs?

Cyber security

While technically out of scope for this review, we encourage CDS to build it into CDS DAQS from the beginning.

Risk Reduction

We risk lots of lost time during installation and commissioning if we don't get our real-time CDS controls and DAQ systems operating robustly. Because of the critical role the CDS systems play to enable operation and testing of other LIGO systems a serious effort at controlling the releases of system upgrades and attention to maintenance activities is of critical importance. A more managed release plan for hardware and software would help in this area.

Calendar and Schedule

Selection and charge to review committee	2008 July 30
Comments and questions back from review committee	2008 August 13
Presentation teleconference	2008 September 18

Documents of relevance to this review

Design documents for the preliminary design review

1. LIGO-G080446-00-C, “AdvLIGO DAQ Design Review,” R. Bork
2. LIGO-T080135-00-C, “AdvLIGO CDS Realtime Code Generator (RCG) Application Developer’s Guide,” R. Bork
3. LIGO-T080136-00-C, “AdvLIGO CDS Realtime Software SysAdmin Guide,” R. Bork
4. LIGO-T080182-00-C, “AdvLIGO Control and Data Acquisition System Preliminary Design,” R. Bork and CDS Staff

Documents from the conceptual design review held May 15, 2007

5. LIGO-T070056-00-C, “Advanced LIGO Control and Data System Infrastructure Requirements,” R. Bork and D. Sigg
6. LIGO-T070059-00-C, “Advanced LIGO Control and Data System Conceptual Design,” R. Bork
7. LIGO-T070057-00-C, “AdvLIGO CDS Subsystem Block Diagrams,” R. Bork
8. LIGO-G070291-00-C, “AdvLIGO Control and Data System Infrastructure Conceptual Design Review,” R. Bork

Acronyms

Acronyms, definitions, and naming conventions are being collected in the document T080179, “advLIGO CDS Software Development Plan Guidelines,” which is a “living document” and will be kept updated and stored on the advLIGO wiki, as well as periodic updates to be checked into the LIGO DCC. The existing naming conventions are confusing and should be revised. (This is a broader problem across LIGO.) Some examples: It is unfortunate that “IOC” is used for “input/output chassis” under the ADC/DAC converter electronics and “input/output controller” under standard EPICS usage. The following acronyms are all synonymous: "DC", "Data Broadcaster", "DC/B", and "DCB". In the iLIGO and eLIGO systems the names frame builder (fb) and network data server (NDS) are used. As the LIGO DAQ system grows, the inevitable change takes place and “growth happens”. A consistent naming convention should be selected and used. For example, “fbc” for frame builder collector or concentrator and “fbw” for frame writer, or "DCB" for Data Concentrator/Broadcaster" as the canonical name. There is confusion over the system to serve data: Is it “NDS” or “NDS2”. Note: NDS is a CDS product, while NDS2 is a GDS product. We need to work collectively to try to resolve these ambiguities.

Acronyms (reproduced from LIGO-T080179-01-C)

1. Definitions and acronyms

Lists of abbreviations and acronyms used in the DAQ system documentation are provided in LIGO-E950111-A, LIGO-T960004-A, and LIGO-T960135-00-C. (Many of these are used in other contexts and often repeated. The reader should be cautious of possible ambiguities.) A brief list of some more commonly used acronyms are given below.

AA	Anti-Aliasing filter
ADC	Analog to Digital Converter
AdvLIGO	Advanced LIGO
AI	Anti-Imaging filter
ASC	Alignment Sensing/Control
ASCII	American Standard Code for Information Interchange
AWG	Arbitrary Waveform Generator
BSC	Beam Splitter Chamber
BT	Beam Tube
CDS	Control and Data System
CIM	Computer Integrated Manufacturing
COC	Core Optics Component
CP	Chiller pad or Cold Plate
CPU	Central Processing Unit
CRC	Cyclic Redundant Check-sum

DAC	Digital to Analog Converter
DAQ or DAQS	Data Acquisition System
DTT	Diagnostic Test Tool, a suite of software tools to provide power spectra and related correlation tests on realtime data. Also provides stimulus signals to allow the measurement of transfer functions and system stimulus responses.
ELIGO	Enhanced LIGO
FCMS	Facility Control and Monitoring System
FB	Frame Builder
FB_C	Frame Concentrator or Frame Collector
FB_W	Frame Writer
FE	Front-End system, the computer, software, and electronic modules that make up the data acquisition system that is in immediate contact with the instrument hardware (e.g., sensors and actuators)
HAM	Horizontal Access Module
HVAC	Heating, Ventilation and Air Conditioning
IFO	Interferometer
IOC	Input Output Chassis or Input-Output Controller (EPICS context)
IOO	Input/Output Optics
ISC	Interferometer Sense and Control
ISI	In-vacuum Seismic Isolator
LA	Louisiana
LDAS (?))	LIGO Data Analysis System (a part of LIGO DCS, Data and Computer Services
LIGO	Laser Interferometer Gravitational-wave Observatory
LOS	Large Optics Suspension
LSC	Length Sensing/Control
LVEA	Laser Vacuum Equipment Area
MSR	Media Storage Room
NAT	Network Address Translation
NDS	Network Data Server
NIC	Network Interface Card
OMC	Output Mode Cleaner
OSB	Operations Support Building
PCIe	Peripheral Component Interconnect Express
PEM	Physical Environment Monitor, the collection of seismometers, magnetometers, accelerometers, anemometers, thermometers, barometers, etc. that monitor the environment around the interferometers.
PSL	Pre-stabilized Laser
RCG	Realtime Code Generator

RDMA	Remote Direct Memory Access
RFM	Reflective Memory
SEI	Seismic Isolation
SOS	Small Optics Suspension
SUS	Suspension
TCP/IP	Transmission Control Protocol/Internet Protocol
UDP	User Datagram Protocol
VE	Vacuum Equipment
VEA	Vacuum Equipment Area
VME	Versa Module European
VPN	Virtual Private Network
WA	Washington
CVS	Concurrent Version Control System
DRR	Design Requirements Review
EPICS	Experimental Physics and Industrial Control System
FDR	Final Design Review
IFO	Interferometer
PDR	Preliminary Design Review
PSL	Pre-Stabilized Laser
SCCS	Source Code Control System
SDL	Software Development Librarian
SMR	Software Maintenance Request
SNL	State Notation Language
SRS	Software Requirements Specification
STP	Software Test Plan
SVN	Subversion Code Repository Control System
TBD	To Be Determined
EPICS	Experimental Physics and Industrial Control System. The toolkit used as the LIGO control system.
MEDM	Motiff Editor/Display Manager
target	In EPICS jargon, a real-time processor which runs the real-time linux or VxWorks operating system. This includes all EPICS IOCs, and all other non-EPICS front end systems (DAQ, LSC, ASC, etc.) Originally it refered to a VME computer, but recent usage has expanded and generalized the term to cover any real-time processor in the EPICS system.
IOC	An EPICS target system (Input Output Controller) . A second use of this acronym has appeared as the input/output chassis used in the advLIGO DAQS converter boxes or blue boxes” that house the ADCs and DACs. This later is now being more commonly referred to as an I/O Chassis, so hopefully there will not be too much ambiguity.

PV	Process Variable. An EPICS variable.
XML	eXtensible Markup Language. A tag based database language. Used in the LIGO Light Weight data format.
h0	LHO systems which are not specific to the 2km or 4km interferometers
h1	LHO first 4km interferometer
h2	LHO initial 2km interferometer, advLIOG second 4km interferometer.
l0	LLO systems which are not specific to the 4km interferometer
l1	LLO 4km interferometer code