Clouds and the Earth's Radiant Energy System (CERES)

Data Management System

CERES AuTomAted job Loading sYSTem (CATALYST) Operator's Manual

Operator's Console

Version 1

Primary Author

Joshua C. Wilkins

Science Systems and Applications Incorporated (SSAI)
One Enterprise Parkway, Suite 200
Hampton, VA 23666

NASA Langley Research Center
Climate Science Branch
Science Directorate
21 Langley Boulevard
Hampton, VA 23681-2199

SW Delivered to CM: September 2015
Document Date: September 2015

Document Revision Record

The Document Revision Record contains information pertaining to approved document changes. The table lists the date the Software Configuration Change Request (SCCR) was approved, the Version Number, the SCCR number, a short description of the revision, and the revised sections. The document authors are listed on the cover. The Head of the CERES Data Management Team approves or disapproves the requested changes based on recommendations of the Configuration Control Board.

Document Revision Record

SCCR Approval Date	Version Number	SCCR Number	Description of Revision	Section(s) Affected
04/02/2015	V1	1066	• Initial version of document.	All

Preface

The Clouds and the Earth's Radiant Energy System (CERES) Data Management System (DMS) supports the data processing needs of the CERES Science Team research to increase understanding of the Earth's climate and radiant environment. The CERES Data Management Team works with the CERES Science Team to develop the software necessary to support the science algorithms. This software, being developed to operate at the Langley Atmospheric Science Data Center (ASDC), produces an extensive set of science data products.

The DMS consists of 12 subsystems; each subsystem contains one or more Product Generation Executables (PGEs). Each subsystem executes when all of its required input data sets are available and produces one or more archival science products.

This Operator's Manual is written for the data processing operations staff at the Langley ASDC by the Data Management Team Systems group who are responsible for the CATALYST system. This document describes the Operator's Console software, and outlines installation and execution procedures.

Acknowledgment is given to the CERES Documentation Team for their support in preparing this document.

TABLE OF CONTENTS

Section Page

Document Revision Record ii

Preface iii

1.0 Operator's Console 3

1.1 Operator's Console Details 3

1.1.1 Responsible Persons 3

1.2 Operator's Console Dependencies 3

1.3 Operating Environment 3

1.4 Obtaining the Operator's Console Software 3

1.5 Running the Operator's Console Software 3

1.6 Logging in to the CATALYST Server 4

2.0 Using the Operator's Console 5

2.1 Layout and Features 5

2.1.1 Operation Permissions 5

2.1.2 Menu Items 6

2.1.2.1 File Menu 6

2.1.2.1.1 Preferences Option 6

2.1.2.1.2 Exit Option 7

2.1.2.2 CATALYST Server Menu 7

2.1.2.2.1 Connect Option 7

2.1.2.2.2 Disconnect Option 7

2.1.2.2.3 Server Configuration Option 8

2.1.2.2.4 Server Environment Option 8

2.1.2.2.5 Server Status Option 9

2.1.2.2.6 Pending Archive Ingests Option 11

2.1.2.2.7 Blade List Option 11

2.1.2.2.8 Current Epilog Job Option 13

TABLE OF CONTENTS

Section Page

2.1.2.2.9 PGE Settings Option 13
2.1.2.2.10 Reload ACL Option 14
2.1.2.2.11 Start Processing (Ready for Processing -> Not Ready for Processing) Option 14
2.1.2.3 View Menu 15
2.1.2.3.1 Hide Log Console Option 15
2.1.2.4 Tools Menu 15
2.1.2.4.1 Global Job Search (by id) Option 15
2.1.2.5 Help Menu 16
2.1.2.5.1 About Option 16
2.1.3 PR View 17
2.1.3.1 PR Sorting 18
2.1.3.2 Viewing a PR's CATALYST Jobs 18
2.1.3.3 Viewing a PR's Chunk Details 18
2.1.3.3.1 Viewing Different Chunk Numbers 19
2.1.3.3.2 Viewing a PR's Initialization Information 19
2.1.3.4 Locking a PR 20
2.1.3.5 Unlocking a PR 20
2.1.3.6 Closing a PR 20
2.1.3.7 Deleting a PR 21
2.1.3.8 Determining a PR's Overall Status 21
2.1.4 Job View 21
2.1.4.1 Navigating the Job View 21
2.1.4.2 Determining the Status of CATALYST Jobs At a Glance 22
2.1.4.2.1 Date Granularity 22
2.1.4.2.2 CATALYST Jobs Status 23

TABLE OF CONTENTS

Section Page

2.1.4.2.3 Total 23

2.1.4.2.4 Science Completed 23

2.1.4.2.5 Science Failed 23

2.1.4.2.6 Epilogs Completed 23

2.1.4.2.7 Epilogs Failed 23

2.1.4.2.8 Instance 24

2.1.4.2.9 CATALYST Job Status 24

2.1.4.2.10 Science Status 26

2.1.4.2.11 Epilog Status 26

2.1.4.3 Quickly Navigating to Failed CATALYST Jobs 27

2.1.4.4 Searching for CATALYST Jobs by Datadate and Status Code 27

2.1.4.4.1 Searching by datadate 27

2.1.4.4.2 Searching by status code 28

2.1.4.5 Performing Operations on CATALYST Jobs 29

2.1.4.5.1 Rescan ( ) 30

2.1.4.5.2 Pause ( ) 30

2.1.4.5.3 Resume ( ) 30

2.1.4.5.4 Force Start ( ) 30

2.1.4.5.5 Flag as Won't Run ( ) 30

2.1.4.5.6 Rerun Science ( ) 31

2.1.4.5.7 Rerun Epilog ( ) 31

2.1.4.5.8 Stop ( ) 31

2.1.4.5.9 Skip Epilog ( ) 31

2.1.4.6 Viewing a CATALYST Job's Details 31

2.1.4.6.1 PCF 33

2.1.4.6.2 PCF Log 34

TABLE OF CONTENTS

Section Page

2.1.4.6.3 SGE Log 34

2.1.4.6.4 Log Report 35

2.1.4.6.5 Log User 35

2.1.4.6.6 Log Status 36

2.1.4.6.7 Epilog Log 36

2.1.5 Log Console 37

2.2 Recovering from an Operator's Console Failure 37

Appendix A - Acronyms and Abbreviations A-1

Appendix B - Messages for Operator's Console B-1

LIST OF FIGURES

Figure Page

Figure 1-1. Login Window 4
Figure 2-1. Operator's Console Layout 5
Figure 2-2. Different Console Views Based on User Privileges 6
Figure 2-3. File Menu 6
Figure 2-4. Operator's Console Preferences 6
Figure 2-5. CATALYST Server Menu 7
Figure 2-6. Server Configuration Window 8
Figure 2-7. CATALYST Server Environment Window 9
Figure 2-8. Server Status Window 10
Figure 2-9. Kernel Log File 11
Figure 2-10. Blade List 12
Figure 2-11. Blade Details 13
Figure 2-12. PGE Settings 14
Figure 2-13. View Menu 15
Figure 2-14. Tools Menu 15
Figure 2-15. Global Job Search 16
Figure 2-16. Help Menu 16
Figure 2-17. About Window 17
Figure 2-18. PR View 18
Figure 2-19. PR Details 19
Figure 2-20. Viewing PR Chunks 19
Figure 2-21. PR Initialization Results 20
Figure 2-22. Job View 21
Figure 2-23. Job View 22
Figure 2-24. Epilog Status 24

LIST OF FIGURES

Figure Page

Figure 2-25. Job Status 24
Figure 2-26. CATALYST Job State Flow 25
Figure 2-27. Science Status 26
Figure 2-28. Search by Datadate 28
Figure 2-29. Search by Status Code 28
Figure 2-30. Job Operations 29
Figure 2-31. Pause Action Submission Status Bar 29
Figure 2-32. Pause Actions Summary 30
Figure 2-33. Job Details Window Part 1 31
Figure 2-34. Job Details Window Part 2 32
Figure 2-35. Links to Predecessor Jobs 32
Figure 2-36. PCF Part 1 33
Figure 2-37. PCF Part 2 33
Figure 2-38. PCF Log 34
Figure 2-39. SGE Log 34
Figure 2-40. Log Report 35
Figure 2-41. Log User 35
Figure 2-42. Log Status 36
Figure 2-43. Epilog Log 36

LIST OF TABLES

Table Page

Table 1-1. CATALYST Software Developer Contacts 3

Table 1-2. Dependencies 3

Table B-1. Operator's Console Messages B-1

Introduction

CERES is a key component of EOS and NPP. The first CERES instrument (PFM) flew on TRMM, four instruments are currently operating on the EOS Terra (FM1 and FM2) and Aqua (FM3 and FM4) platforms, and NPP (FM5) platform. CERES measures radiances in three broadband channels: a shortwave channel (0.3 - 5 μm), a total channel (0.3 - 200 μm), and an infrared window channel (8 - 12 μm). The last data processed from the PFM instrument aboard TRMM was March 2000; no additional data are expected. Until June 2005, one instrument on each EOS platform operated in a fixed azimuth scanning mode and the other operated in a rotating azimuth scanning mode; now all are typically operating in the fixed azimuth scanning mode. The NPP platform carries the FM5 instrument, which operates in the fixed azimuth scanning mode though it has the capability to operate in a rotating azimuth scanning mode.

CERES climate data records involve an unprecedented level of data fusion: CERES measurements are combined with imager data (e.g., MODIS on Terra and Aqua, VIIRS on NPP), 4-D weather assimilation data, microwave sea-ice observations, and measurements from five geostationary satellites to produce climate-quality radiative fluxes at the top-of-atmosphere, within the atmosphere and at the surface, together with the associated cloud and aerosol properties.

The CERES project management and implementation responsibility is at NASA Langley. The CERES Science Team is responsible for the instrument design and the derivation and validation of the scientific algorithms used to produce the data products distributed to the atmospheric sciences community. The CERES DMT is responsible for the development and maintenance of the software that implements the science team's algorithms in the production environment to produce CERES data products. The Langley ASDC is responsible for the production environment, data ingest, and the processing, archival, and distribution of the CERES data products.

Document Overview

This document, CERES CATALYST Operator's Console Operator's Manual is part of the CERES CATALYST Operator's Console delivery packages provided to the Langley Atmospheric Science Data Center (ASDC). It provides a description of the CERES CATALYST Operator's Console software and explains the procedures for executing the software. A description of acronyms and abbreviations is provided in Appendix A, and a comprehensive list of messages that can be generated during the execution of the Operator's Console are contained in Appendix B.

This document is organized as follows:

Introduction
Document Overview
CATALYST Operator's Console
Section 1.0 - Operator's Console
Section 2.0 - Using the Operator's Console

Appendix A - Acronyms and Abbreviations
Appendix B - Error Messages for Operator's Console

CATALYST Operator's Console

The Operator's Console is a Java client side application that displays information from the CATALYST Server (Refer to the CATALYST Server Operator's Manual for more information about the CATALYST Server) and provides the ability to control and monitor the execution of individual PRs. The Operator's Console will show which PRs are active within CATALYST, PGE instances and their status, and *AMI* cluster node availability. Depending on the assigned privileges a user can pause, resume, rerun, and flag PGE instances not to run, check PGE computer resource use, and monitor, as well as start and stop multiple server processes. Each operator will run the console on their local workstation.

1.0 Operator's Console

1.1 Operator's Console Details

1.1.1 Responsible Persons

Table 1-1. CATALYST Software Developer Contacts

Item	Primary	Alternate
Contact Name	Joshua C. Wilkins	T. Nelson Hillyer
Organization	SSAI	SSAI
Address	1 Enterprise Parkway	1 Enterprise Parkway
City	Hampton	Hampton
State	VA 23666	VA 23666
Phone	757-951-1618	757-951-1951
Fax	757-951-1900	757-951-1900
LaRC email	joshua.c.wilkins@nasa.gov	thomas.n.hillyer@nasa.gov

1.2 Operator's Console Dependencies

Table 1-2. Dependencies

Dependency	Version Supported
JAVA Runtime	Version 7 or Newer

1.3 Operating Environment

The Operator's Console is run from the user's local desktop computer.

1.4 Obtaining the Operator's Console Software

The latest version of the Operator's Console can be downloaded from the CATALYST Home page at the following URL:

http://ceres.larc.nasa.gov/Internal/intern_docs_catalyst.php

Download either the zip or tar file of the latest Operator's Console (choose the format you are most comfortable with) and extract it.

1.5 Running the Operator's Console Software

To run the Operator's Console choose from the following based on your operating system:

1. Windows → double click start_oc.bat or double click OperatorsConsole_YYYY-MM-DD.jar

2. Mac → double click OperatorsConsole_YYYY-MM-DD.jar or run ./start_oc.sh from the command line

3. Linux → run ./start_oc.sh from the command line (double clicking the OperatorsConsole_YYYY-MM-DD.jar file may also work depending on your linux distribution)

1.6 Logging in to the CATALYST Server

Figure 1-1. Login Window

Once the console starts up, you are prompted to enter in the server details for CATALYST. The remote hostname will typically be catalyst, but if you are not sure ask a member of the CATALYST team. While the remote port will default to the production server port, listed below are the remote port numbers for all possible environments:

- Production – use remote port 4020
- PPE – use remote port 4021
- CM Testing – use remote port 8020

To change the default remote port, see File Menu. The username and password will be the one you use to login to the *AMI* systems. The Operator's Console generates a log file which is displayed at the bottom of the console window. If you notice any unexpected errors, please send the log file (client.log) and a description of what you were doing when the error occurred to the CATALYST development team. The log file is located in the directory from which the Operator's Console was opened.

2.0 Using the Operator's Console

2.1 Layout and Features

The figure below shows the overall layout of the Operator's Console – PR View, Job View, and Log Console. Further layout details and features are explained more thoroughly in later sections.

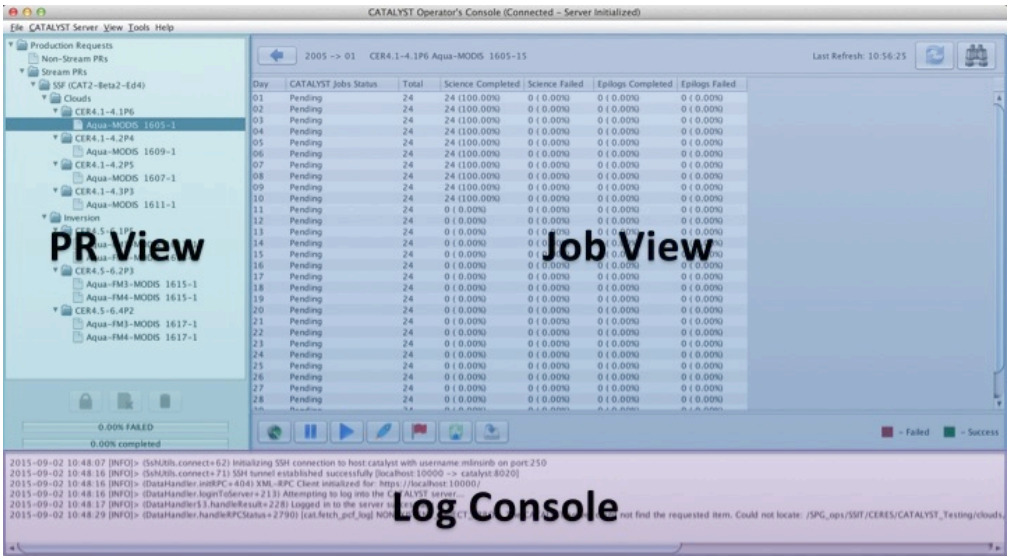

Figure 2-1. Operator's Console Layout

2.1.1 Operation Permissions

Once connected to the CATALYST Server, the Operator's Console periodically confirms the user's privileges as set in the ACL. Options that the CATALYST Server does not support or that the user does not have permission to access are grayed out in the Operator's Console (see the CATALYST Server Operator's Manual for more on user permissions).

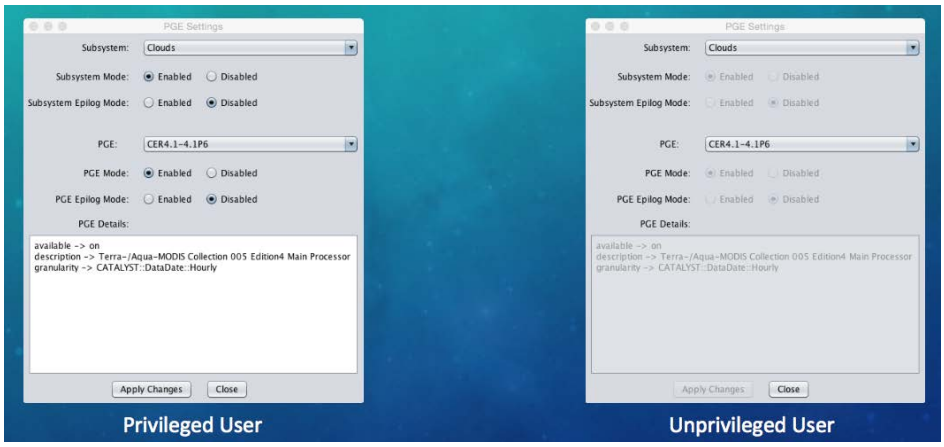

Figure 2-2. Different Console Views Based on User Privileges

2.1.2 Menu Items

2.1.2.1 File Menu

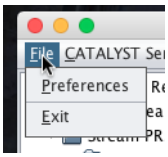

Figure 2-3. File Menu

2.1.2.1.1 Preferences Option

Opens the Operator's Console preferences window, which allows the user to change the default values for several variables used by the Operator's Console.

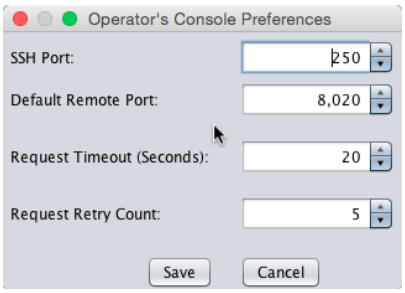

Figure 2-4. Operator's Console Preferences

a) SSH Port – The SSH port number that the Operator's Console will use when logging into the CATALYST Server
b) Default Remote Port – The port number of the CATALYST Server that the Operator's Console will use by default when logging into CATALYST. See Logging in to the CATALYST Server for more information.
c) Request Timeout (Seconds) – The time in seconds before the Operator's Console gives up waiting for a response from CATALYST server for a particular request. It is recommended to contact the Operator's Console development team for help with the Request Timeout value before changing it.
d) Request Retry Count – The number of retries the Operator's Console will attempt for a particular request sent to the CATALYST Server before giving up completely. After the maximum amount of retries has failed, the Operator's Console will assume that the Server connection has been lost and will disconnect. See Appendix B for console related error messages. It is recommended to contact the Operator's Console development team for help with the Request Retry Count value before changing it.

2.1.2.1.2 Exit Option

Closes the Operator's Console and disconnects from the CATALYST Server.

2.1.2.2 CATALYST Server Menu

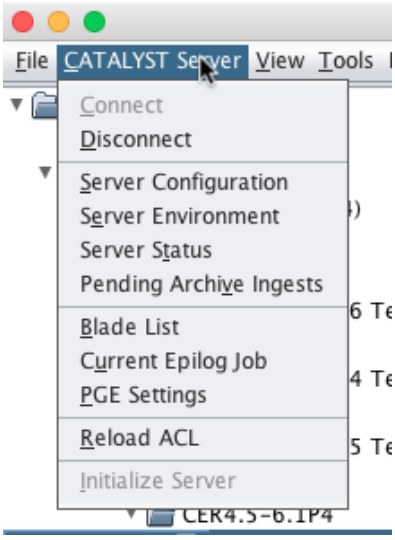

Figure 2-5. CATALYST Server Menu

2.1.2.2.1 Connect Option

Brings up the login window. See Logging in to the CATALYST Server for more information.

2.1.2.2.2 Disconnect Option

Disconnects from the CATALYST Server.

2.1.2.2.3 Server Configuration Option

Displays the CATALYST Server configuration properties found in its configuration file (usually called catalyst.conf). See the CATALYST Server operator's manual for more information.

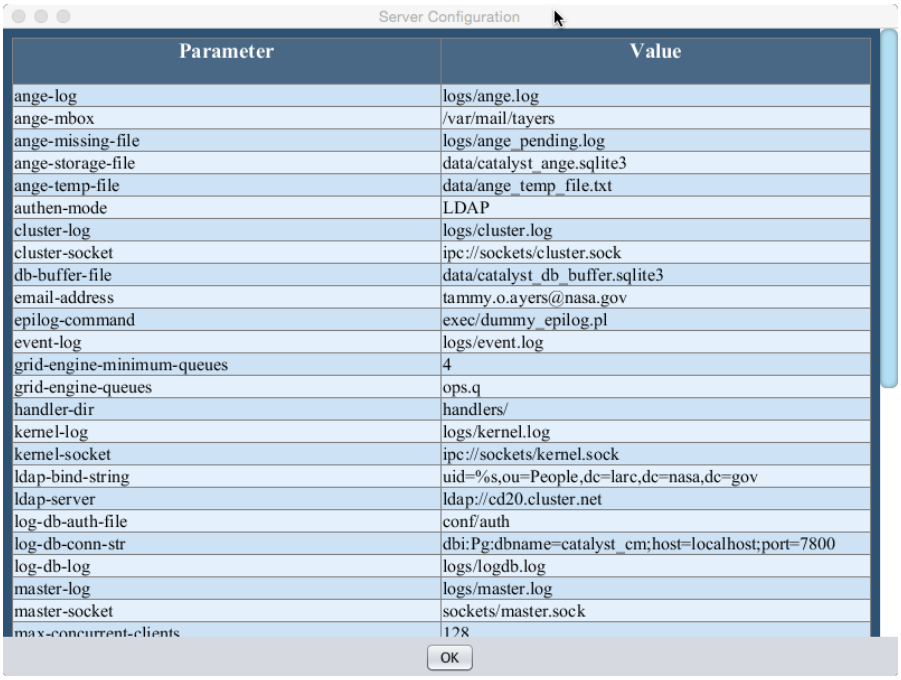

Parameter	Value
ange-log	logs/ange.log
ange-mbox	/var/mail/tayers
ange-missing-file	logs/ange_pending.log
ange-storage-file	data/catalyst_ange.sqlite3
ange-temp-file	data/ange_temp_file.txt
authen-mode	LDAP
cluster-log	logs/cluster.log
cluster-socket	ipc://sockets/cluster.sock
db-buffer-file	data/catalyst_db_buffer.sqlite3
email-address	tammy.o.ayers@nasa.gov
epilog-command	exec/dummy_epilog.pl
event-log	logs/event.log
grid-engine-minimum-queues	4
grid-engine-queues	ops.q
handler-dir	handlers/
kernel-log	logs/kernel.log
kernel-socket	ipc://sockets/kernel.sock
ldap-bind-string	uid=%s,ou=People,dc=larc,dc=nasa,dc=gov
ldap-server	ldap://cd20.cluster.net
log-db-auth-file	conf/auth
log-db-conn-str	dbi:Pg:dbname=catalyst_cm;host=localhost;port=7800
log-db-log	logs/logdb.log
master-log	logs/master.log
master-socket	sockets/master.sock
max-concurrent-clients	128

Figure 2-6. Server Configuration Window

2.1.2.2.4 Server Environment Option

Brings up a window listing the environment variables with which the CATALYST server was launched.

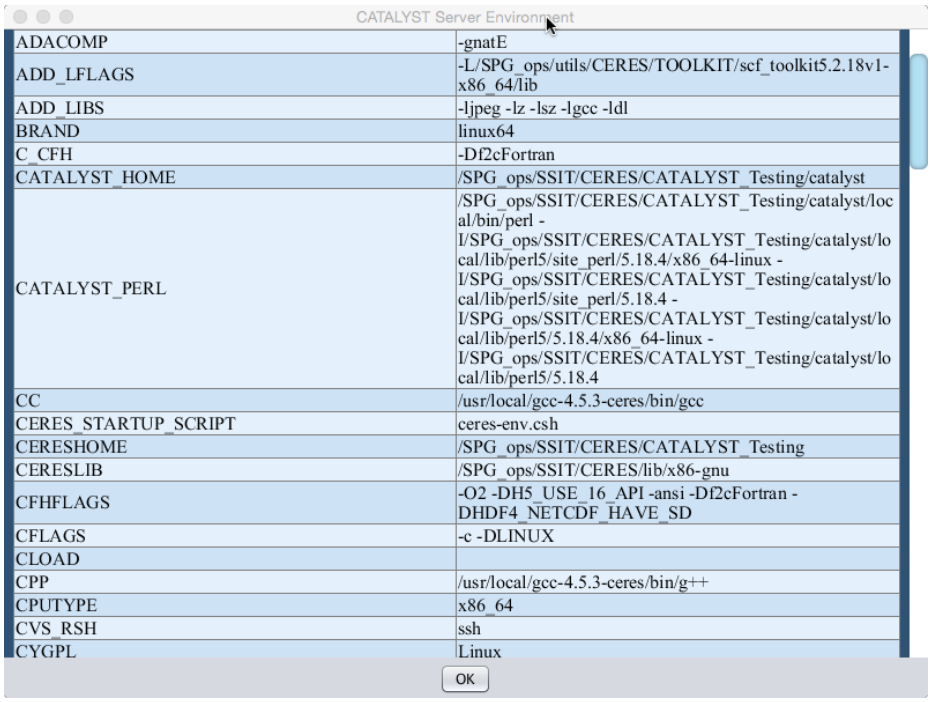

CATALYST Server Environment

ADACOMP	-gnatE
ADD_LFLAGS	-L/SPG_ops/utils/CERES/TOOLKIT/scf_toolkit5.2.18v1-x86_64/lib
ADD_LIBS	-ljpeg -lz -lsz -lgcc -ldl
BRAND	linux64
C_CFH	-Df2cFortran
CATALYST_HOME	/SPG_ops/SSIT/CERES/CATALYST_Testing/catalyst
CATALYST_PERL	/SPG_ops/SSIT/CERES/CATALYST_Testing/catalyst/local/bin/perl -I/SPG_ops/SSIT/CERES/CATALYST_Testing/catalyst/local/lib/perl5/site_perl/5.18.4/x86_64-linux -I/SPG_ops/SSIT/CERES/CATALYST_Testing/catalyst/local/lib/perl5/site_perl/5.18.4 -I/SPG_ops/SSIT/CERES/CATALYST_Testing/catalyst/local/lib/perl5/5.18.4/x86_64-linux -I/SPG_ops/SSIT/CERES/CATALYST_Testing/catalyst/local/lib/perl5/5.18.4
CC	/usr/local/gcc-4.5.3-ceres/bin/gcc
CERES_STARTUP_SCRIPT	ceres-env.csh
CERESHOME	/SPG_ops/SSIT/CERES/CATALYST_Testing
CERESLIB	/SPG_ops/SSIT/CERES/lib/x86-gnu
CFHFLAGS	-O2 -DH5_USE_16_API -ansi -Df2cFortran -DHDF4_NETCDF_HAVE_SD
CFLAGS	-c -DLINUX
CLOAD	
CPP	/usr/local/gcc-4.5.3-ceres/bin/g++
CPUTYPE	x86_64
CVS_RSH	ssh
CYGPL	Linux

OK

Figure 2-7. CATALYST Server Environment Window

2.1.2.2.5 Server Status Option

Brings up a window which shows the processes running under the CATALYST Server and how much memory resources each is using on the system. Users with the right privileges can start and stop several of the CATALYST Server processes as needed. See the CATALYST Server Operator's Manual for more details on each server process.

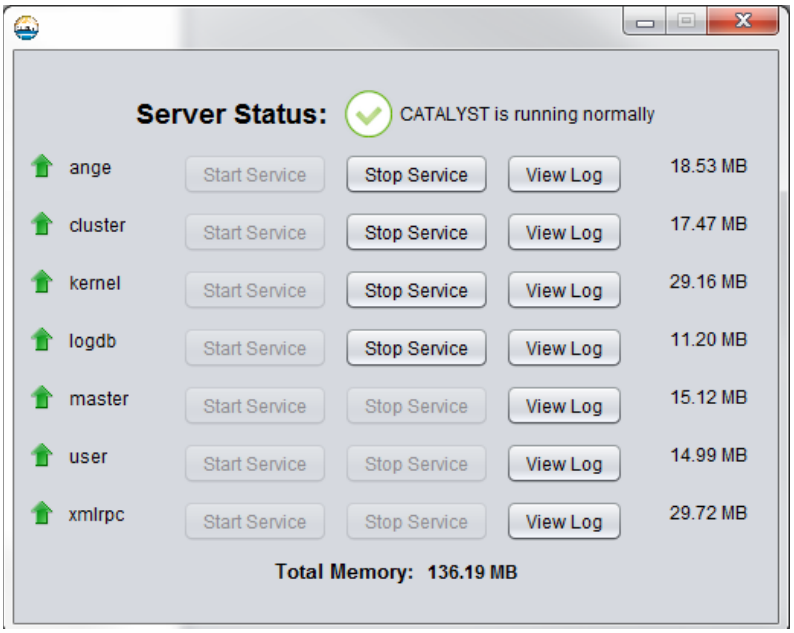

Figure 2-8. Server Status Window

At the top of the window is an overall Server Status. If one or more of the CATALYST Server processes are down, the text: "CATALYST is running in a degraded state" is shown.

Clicking on any of the "View Log" buttons will open up the server log for that particular process. This log window formats and displays the last 100 lines of the log file maintained by the CATALYST Server for that process. An example is shown below.

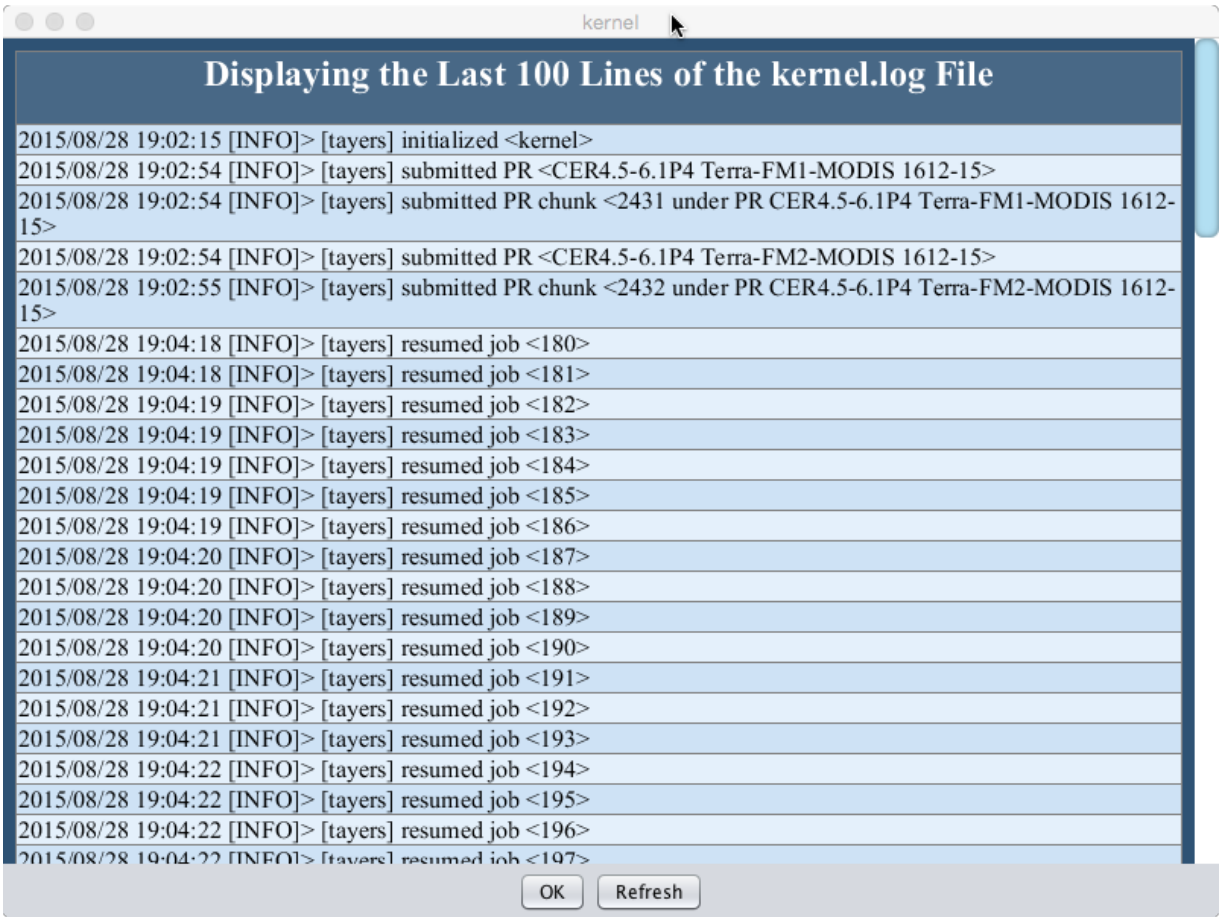

Figure 2-9. Kernel Log File

2.1.2.2.6 Pending Archive Ingests Option

Retrieves and displays the list of archive ingests that have been pending being written to the DPO for a time period of greater than a day. Usually this list will be empty but on rare occasions the list can potentially be large and take a while to load. Items on the list show the ANGe archived time, and the size of the file provided in the ANGe ingest email. Files that CATALYST finds on the DPO that do not match the expected file size provided by ANGe will also appear on this list and can be investigated from the command line. Please contact the CATALYST team with questions about any items that show on this list, as typically they will need further investigation in order to provide the best course of action for the particular item.

2.1.2.2.7 Blade List Option

Displays the list of blades known to the CATALYST Server. The server regularly updates the status of by syncing with UGE and disabling blades if errors are encountered. When hovering the mouse over a particular blade name in this window, more details gleaned from UGE by the CATALYST Server will be displayed. Note that any status set by CATALYST for a particular blade is overridden by the UGE status of "unavailable" for a blade – that means that a blade that

SGE/UGE has determined to be down and unavailable for use will be down in CATALYST and unavailable regardless of whether it is enabled or disabled by the CATALYST Server.

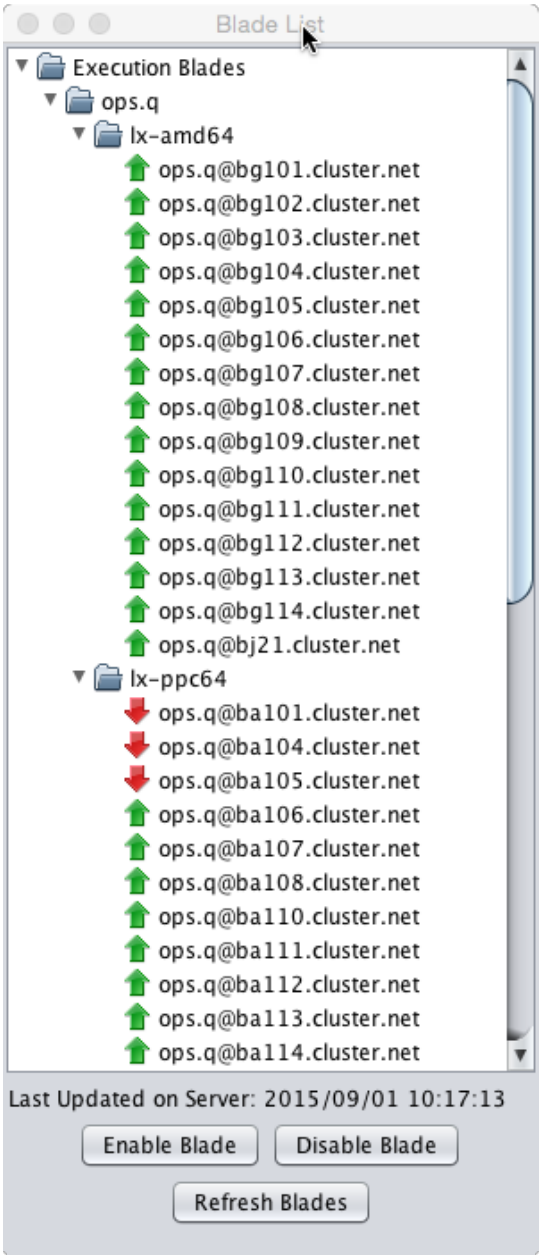

Figure 2-10. Blade List

a) Last Updated On Server – shows the time the CATALYST Server last updated its list of blades based on its periodic queries to SGE or user changes to enabled blades.
b) Enable Blade – Requests that the given blade be set to enabled in the CATALYST Server. Note, a blade that is unavailable in SGE cannot be enabled and a message will be sent back to the Operator's Console explaining why the blade could not be enabled.

c) Disable Blade – Requests that the given blade be set to disabled in the CATALYST Server. CAUTION: Blades that have been disabled, either manually or automatically from a BLD_ERROR science status code, are not automatically enabled again by CATALYST when the issues are fixed. A user needs to re-enable blades that have gone offline or were disabled previously in order for CATALYST to use them again for processing.
d) Refresh Blades – Gets the most up-to-date list of blades from the CATALYST Server.

The hover text previously mentioned showing more blade details:

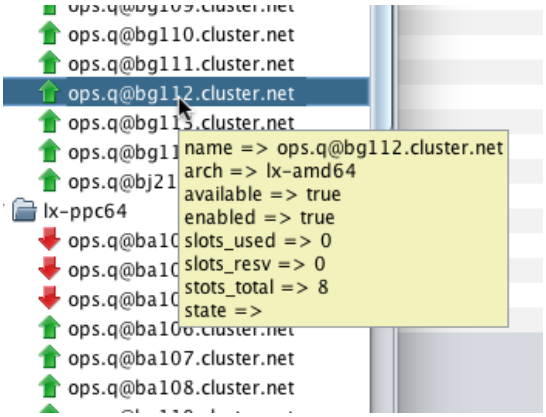

Figure 2-11. Blade Details

2.1.2.2.8 Current Epilog Job Option

Opens up the Global Job Search window and displays the CATALYST job associated with the currently running epilog process, if any. If there are no completed jobs and therefore no epilogs to run at any given time the message, "Epilog work queue is empty" will be displayed. If the most recently executed epilog did not complete successfully its associated job will be displayed. Epilogs are processed in a first-in, first-out order so that when an epilog fails to run, all of the follow on epilogs are blocked until that error is fixed. This prevents epilogs from running out of order and cleaning up science data earlier than desired, especially when dealing with some follow-on PGE epilogs that clean up data from a predecessor PGE's output.

2.1.2.2.9 PGE Settings Option

Allows users with elevated privileges to disable subsystems, individual PGEs, or epilogs in the CATALYST Server. Disabling a subsystem, for instance, stops all PRs for the PGEs under that subsystem from progressing any further. When a PGE is disabled, or its subsystem is disabled, new PRs using the PGE will be rejected when submitted from the PR tool. Once you have changed either a subsystem's setting or a PGE's setting, click "Apply Changes" to make sure the CATALYST server gets the updated settings before closing the window, or selecting a different subsystem or PGE. "Apply Changes" only sends the changes for the subsystem or PGE you are currently viewing. **Please note, if you don't click "Apply Changes" before switching to a different PGE or subsystem, the CATALYST server will not get the changes.** After re-enabling a previously disabled PGE, you may need to unlock any PR(s) for that PGE (see

Unlocking a PR for more information on how to unlock a PR). Users without the proper privileges can view the settings in the PGE Settings but cannot change them. Please see the CATALYST Server Operator's Manual for more information on managing PGE handlers.

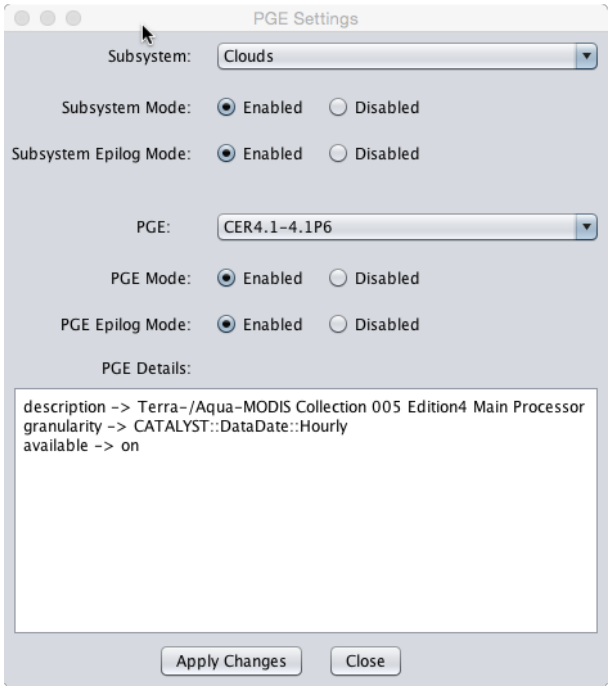

Figure 2-12. PGE Settings

2.1.2.2.10 Reload ACL Option

Sends a request to the CATALYST Server to refresh the Access Control List. See the CATALYST Server Operator's Manual for more information about the ACL.

2.1.2.2.11 Start Processing (Ready for Processing -> Not Ready for Processing) Option

Sends a request to the CATALYST Server to start processing. The CATALYST Server must be set to the "Ready for Processing" mode before it can run jobs, accept new PRs, or apply actions on CATALYST Jobs such as pause, etc. Restarting the kernel process, or when CATALYST is initially launched, will set the CATALYST Server into the mode "Not Ready for Processing".

2.1.2.3 View Menu

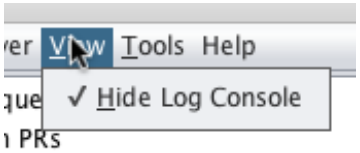

Figure 2-13. View Menu

2.1.2.3.1 Hide Log Console Option

Toggles showing or hiding the log console near the bottom of the Operator's Console. Since all errors and warnings are also written to the log file for the console, users can safely hide the log console if they desire more space to be given to the rest of the console.

2.1.2.4 Tools Menu

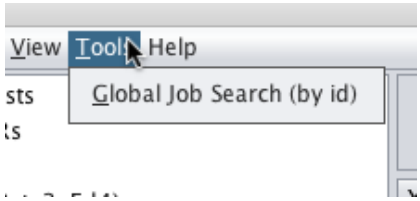

Figure 2-14. Tools Menu

2.1.2.4.1 Global Job Search (by id) Option

Opens up a job search window that allows the user to search for a specific CATALYST Job by its CATALYST ID number. Since all the CATALYST server logs list jobs by their CATALYST ID the Global Job Search gives users a way to search for any ID listed in those log files.

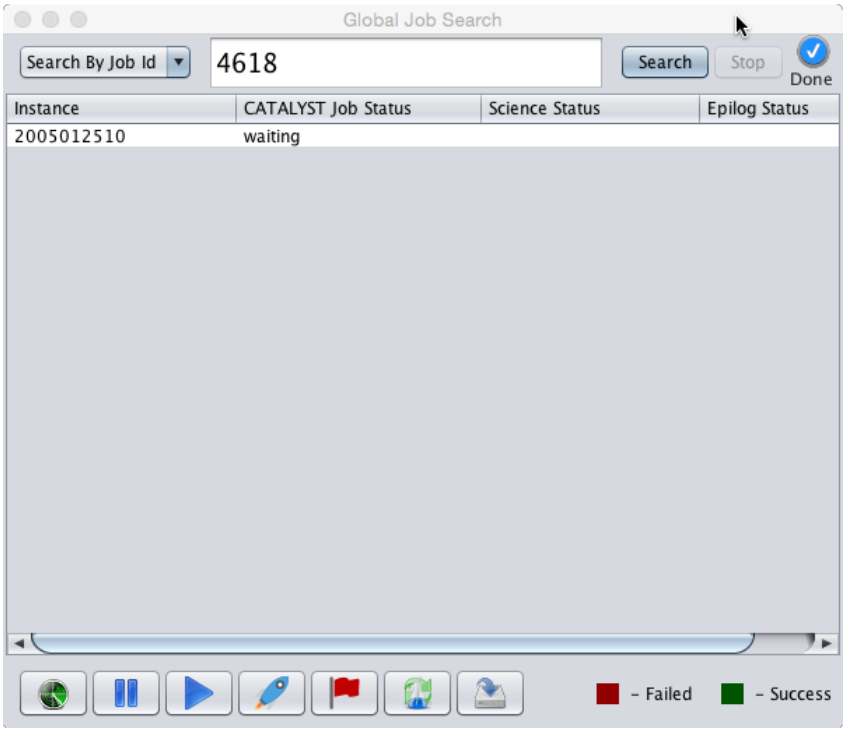

Figure 2-15. Global Job Search

a) Search Bar – Enter the CATALYST Job ID in the textbox and then click search to start a search for that job ID.
b) Job View – Shows the jobs, if any are found, in a table. Performing Operations on CATALYST Jobs
c) Job Commands – See Performing Operations on CATALYST Jobs in a PR for more information.

2.1.2.5 Help Menu

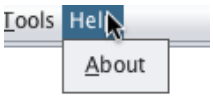

Figure 2-16. Help Menu

2.1.2.5.1 About Option

Opens the About window for the Operator's Console which lists the build date, developer contact information, and the current user's Java environment.

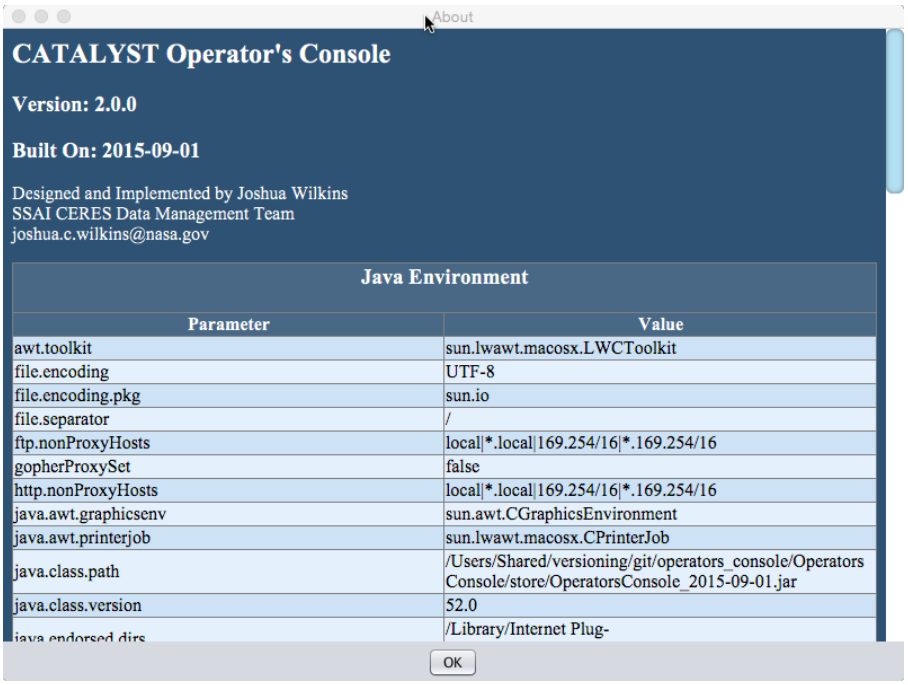

Figure 2-17. About Window

2.1.3 PR View

The PR View allows the user to monitor and perform several operations on the active PRs in the CATALYST Server.

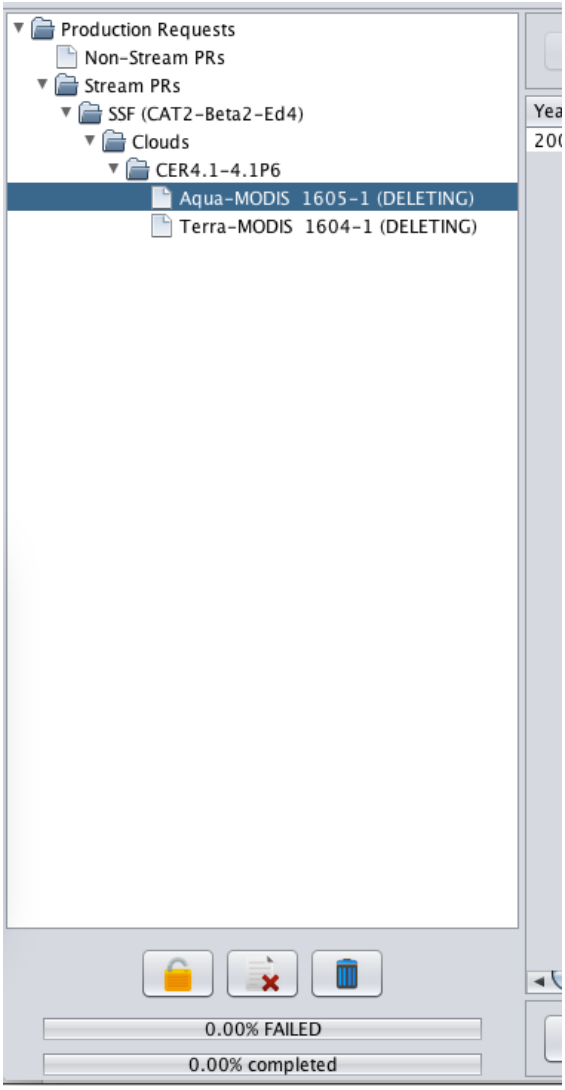

Figure 2-18. PR View

2.1.3.1 PR Sorting

PRs are organized in the following way: Stream name (if available) → subsystem name → PGE name → PR name. To hide a PGE, subsystem, or even stream that you do not wish to view in the list simply click on the small arrows beside the folder icons to the left of the category you wish to collapse and hide.

2.1.3.2 Viewing a PR's CATALYST Jobs

Single click on the PR name you wish to view the jobs of, and the jobs will show up in the Job View, located in the center of the Operator's Console.

2.1.3.3 Viewing a PR's Chunk Details

To view a PR's details by chunk, double click on the PR name and the PR Details Window will appear.

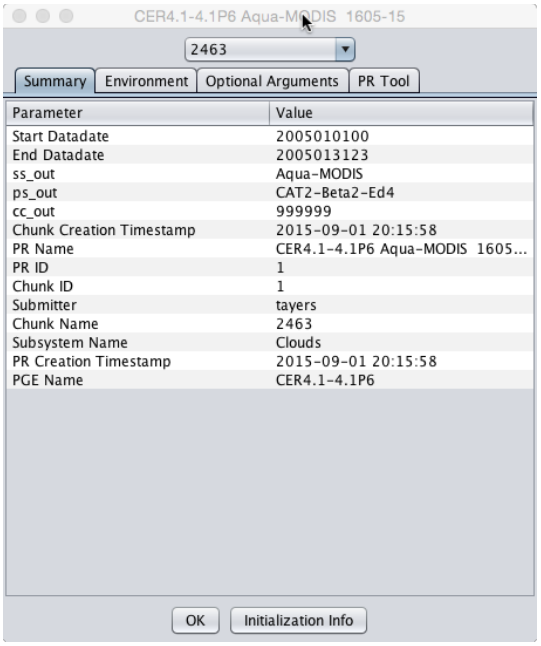

Figure 2-19. PR Details

2.1.3.3.1 Viewing Different Chunk Numbers

To view a different chunk's details, click the dropdown and select the desired chunk (the chunk name or number is provided by the PR Web Tool when the chunk is submitted to CATALYST).

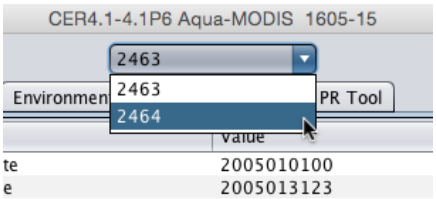

Figure 2-20. Viewing PR Chunks

2.1.3.3.2 Viewing a PR's Initialization Information

Some PR's have specific initialization steps that are performed when the PR is submitted to CATALYST, such as Clouds PGE CER4.1-4.1P6's CopyECS routine (see the Clouds Operator's Manual for more details on this). Clicking on "Initialization Info" will open a window containing the output generated by any initialization steps CATALYST executed when the PR was submitted.

Figure 2-21. PR Initialization Results

2.1.3.4 Locking a PR

A privileged user can lock a PR by either right clicking on the PR name and selecting "lock PR" or by selecting the PR and clicking on the lock icon button near the bottom of the PR view ( ). Locking puts the PR into essentially a frozen state. This means no more jobs will run in SGE and no actions can be performed on the PR. Once locked a privileged user can either unlock, close, or delete the PR. Locked PR's will show that it has been locked beside the PR's name for all users. Other users trying to apply commands such as pause or re-running science code for a job will receive a message saying the PR is locked.

2.1.3.5 Unlocking a PR

A privileged user can unlock a PR by either right clicking on the PR name and selecting "unlock PR" or by selecting the PR and clicking on the unlock icon button near the bottom of the PR view ( ). This puts the PR back into a normal state, allowing jobs to run and job actions to be applied.

2.1.3.6 Closing a PR

Once a PR has been locked, a privileged user can close a PR by right clicking on the PR and selecting "Close PR" or by clicking on the close PR button ( ). Closing a PR removes it from the CATALYST Server without calling any cleanup on the science data on disk. It is recommended to close a PR either before any jobs have run (no data has been created) or after it is 100% completed which means the epilog wrappers have run and have cleaned up the files. Once a PR is removed from the CATALYST Server all Operator's Console users will see a window appear listing the PR that has been removed.

2.1.3.7 Deleting a PR

Once a PR has been locked, a privileged user can delete a PR by right clicking on the PR and selecting "Delete PR" or by clicking on the delete PR button ( ). Deleting a PR will change its state to "DELETING" which is shown to the right of the PR name. PRs in the process of deleting will clean up any data associated with each job. Users can still view the PR's jobs that have not yet been deleted while it's being deleted in the Job View. This process can take some time depending on the PGE's cleanup steps needed. Once the deletion steps are completed the PR is removed from the CATALYST Server. All Operator's Console users will see a window appear showing which PR has been removed from the CATALYST Server.

2.1.3.8 Determining a PR's Overall Status

The two progress bars at the bottom of the PR View show the total percentage of CATALYST jobs complete (both science and epilog) as well as the percentage of CATALYST jobs that resulted in a failure of some sort. This provides an at-a-glance look at how far a PR has progressed and if there are any problems that need to be investigated further in the Job View.

2.1.4 Job View

The Job View allows the user to monitor and perform operations on the individual CATALYST Jobs for a given PR. Selecting a PR is detailed in Viewing a PR's CATALYST Jobs.

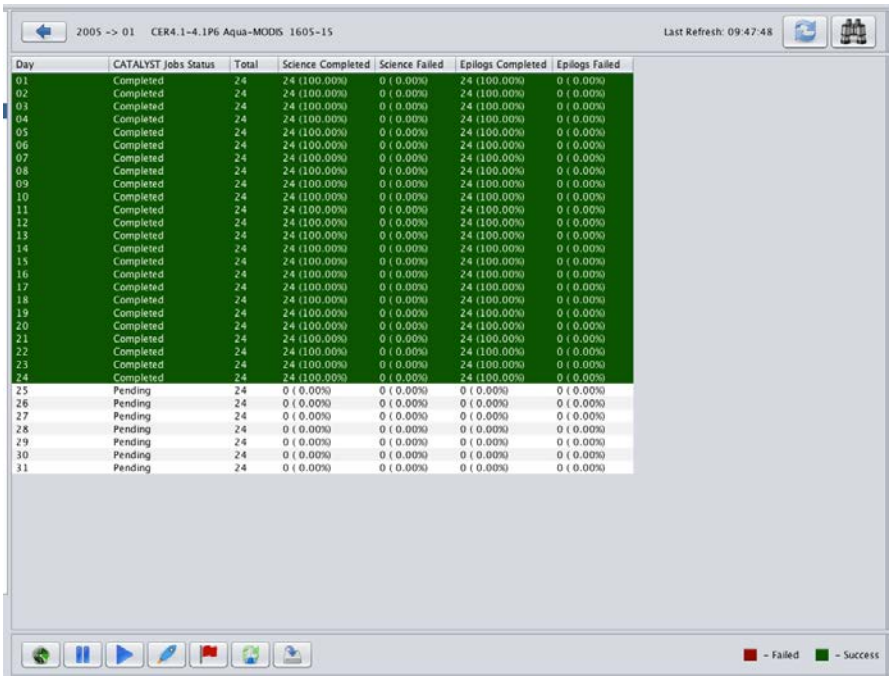

Day	CATALYST Jobs Status	Total	Science Completed	Science Failed	Epilogs Completed	Epilogs Failed
01	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
02	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
03	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
04	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
05	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
06	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
07	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
08	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
09	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
10	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
11	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
12	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
13	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
14	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
15	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
16	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
17	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
18	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
19	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
20	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
21	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
22	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
23	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
24	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
25	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
26	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
27	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
28	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
29	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
30	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
31	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)

Figure 2-22. Job View

2.1.4.1 Navigating the Job View

The Job View uses a drill-down system organized by date (or zones for some PGEs), represented essentially as folders containing either CATALYST Jobs or other folders, following the pattern

year -> month -> day -> hour. To navigate to a particular CATALYST job for the selected PR (see Viewing a PR's CATALYST Jobs for how to select a PR in CATALYST), double click a row in the Job View that would have the desired datadate (for example: looking for the hourly job 2005010115, double click on year 2005, then month 01, followed by day 01, and listed should be the CATALYST jobs by their full datadate including hour 15). To go back to a higher date (such as from day back to month) click the back arrow button ( ). To force a refresh of the current list, click the refresh button near the top right of the Job View ( ), this will retrieve the most current jobs at that date level from the CATALYST Server for the selected PR in the PR View.

2.1.4.2 Determining the Status of CATALYST Jobs At a Glance

When looking at CATALYST Jobs in the Job View after selecting a PR in the PR View, the columns in the Job View show the overall status for CATALYST Jobs that fall under that date for each row (year,month,etc). At the higher date granularities, each row represents a folder in the drill down scheme detailed in Navigating the Job View. Folders have a different set of columns from the lowest granularity, which contains the actual CATALYST Jobs after drilling down to the lowest level and are detailed further down. The folder columns are detailed below:

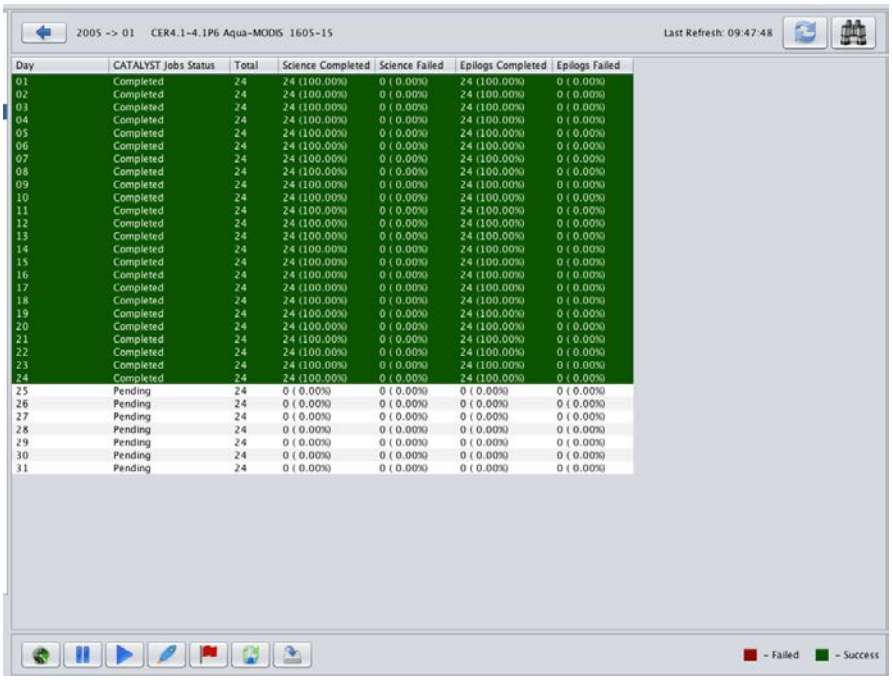

Day	CATALYST Jobs Status	Total	Science Completed	Science Failed	Epilogs Completed	Epilogs Failed
01	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
02	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
03	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
04	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
05	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
06	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
07	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
08	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
09	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
10	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
11	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
12	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
13	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
14	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
15	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
16	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
17	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
18	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
19	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
20	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
21	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
22	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
23	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
24	Completed	24	24 (100.00%)	0 (0.00%)	24 (100.00%)	0 (0.00%)
25	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
26	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
27	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
28	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
29	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
30	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)
31	Pending	24	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)

Figure 2-23. Job View

2.1.4.2.1 Date Granularity

Shows which date granularity (In the screenshot above the granularity is Day) applies to the rows of folders containing Jobs.

2.1.4.2.2 CATALYST Jobs Status

Shows the overall status of any jobs that fall under the date listed for that row. The individual CATALYST Job states and their respective state transition flow are described in detail below:

a) Pending – Not all of CATALYST jobs under this date have finished running their science and epilog processes. No failures have occurred for the CATALYST jobs under this date that have completed their respective processes. The row color for this value will show as white.
b) Pending w/failures – Not all of the CATALYST jobs under this date have finished running their science and epilog processes. One or more CATALYST jobs under this date have resulted in a failure for either their science or epilog process. The row color for this value will show as red.
c) Completed – All of the CATALYST jobs under this date have finished running their science and epilog processes. No failures have occurred for the CATALYST jobs under this date. The row color with this value will show as green.
d) Completed w/failures – All of the CATALYST jobs under this date have finished running their science and epilog processes. One or more of the completed jobs resulted in a failure for either their science process or epilog process. The row color for this value will show as red.

2.1.4.2.3 Total

Shows the total number of CATALYST Jobs under this date.

2.1.4.2.4 Science Completed

Shows the total number and percentage of science processes that have completed under this date.

2.1.4.2.5 Science Failed

Shows the total number and percentage of science processes that have failed under this date.

2.1.4.2.6 Epilogs Completed

Shows the total number and percentage of epilog processes that have completed under this date.

2.1.4.2.7 Epilogs Failed

Shows the total number and percentage of epilog processes that have failed under this date.
After drilling down to the lowest date granularity (dependent on the PGE), the individual CATALYST jobs can be seen, which have a different set of columns and values. Those columns are detailed below:

Instance	CATALYST Job Status	Science Status	Epilog Status
2005012700	completed	SUCCESS	SUCCESS
2005012701	completed	SUCCESS	SUCCESS
2005012702	completed	SUCCESS	SUCCESS
2005012703	completed	SUCCESS	SUCCESS
2005012704	completed	SUCCESS	SUCCESS
2005012705	completed	SUCCESS	SUCCESS
2005012706	completed	SUCCESS	SUCCESS
2005012707	completed	SUCCESS	SUCCESS
2005012708	completed	SUCCESS	SUCCESS
2005012709	completed	SUCCESS	SUCCESS
2005012710	completed	SUCCESS	SUCCESS
2005012711	completed	SUCCESS	SUCCESS
2005012712	completed	SUCCESS	SUCCESS
2005012713	completed	SUCCESS	SUCCESS
2005012714	completed	SUCCESS	SUCCESS
2005012715	completed	SUCCESS	SUCCESS
2005012716	completed	SUCCESS	SUCCESS
2005012717	completed	SUCCESS	SUCCESS
2005012718	completed	SUCCESS	SUCCESS
2005012719	completed	SUCCESS	SUCCESS
2005012720	completed	SUCCESS	SUCCESS
2005012721	completed	SUCCESS	SUCCESS
2005012722	completed	SUCCESS	SUCCESS
2005012723	completed	SUCCESS	SUCCESS

Figure 2-24. Epilog Status

2.1.4.2.8 Instance

Shows the datadate for this CATALYST Job.

2.1.4.2.9 CATALYST Job Status

Shows the state of the CATALYST Job in the CATALYST Server. Hovering over this column with the mouse will show a description of the state in the Operator's Console.

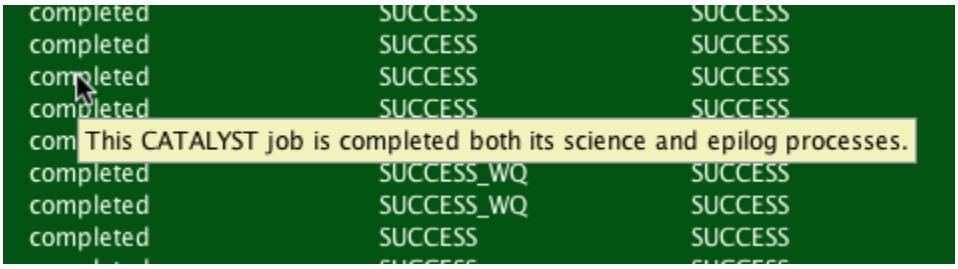

Figure 2-25. Job Status

A CATALYST Job goes through several states in its lifetime. State changes are initiated either by the CATALYST Server for normal job flow (checking inputs, running the job, etc.) or by a user performed action (see Performing Operations on CATALYST Jobs). The CATALYST Job state flow is detailed below:

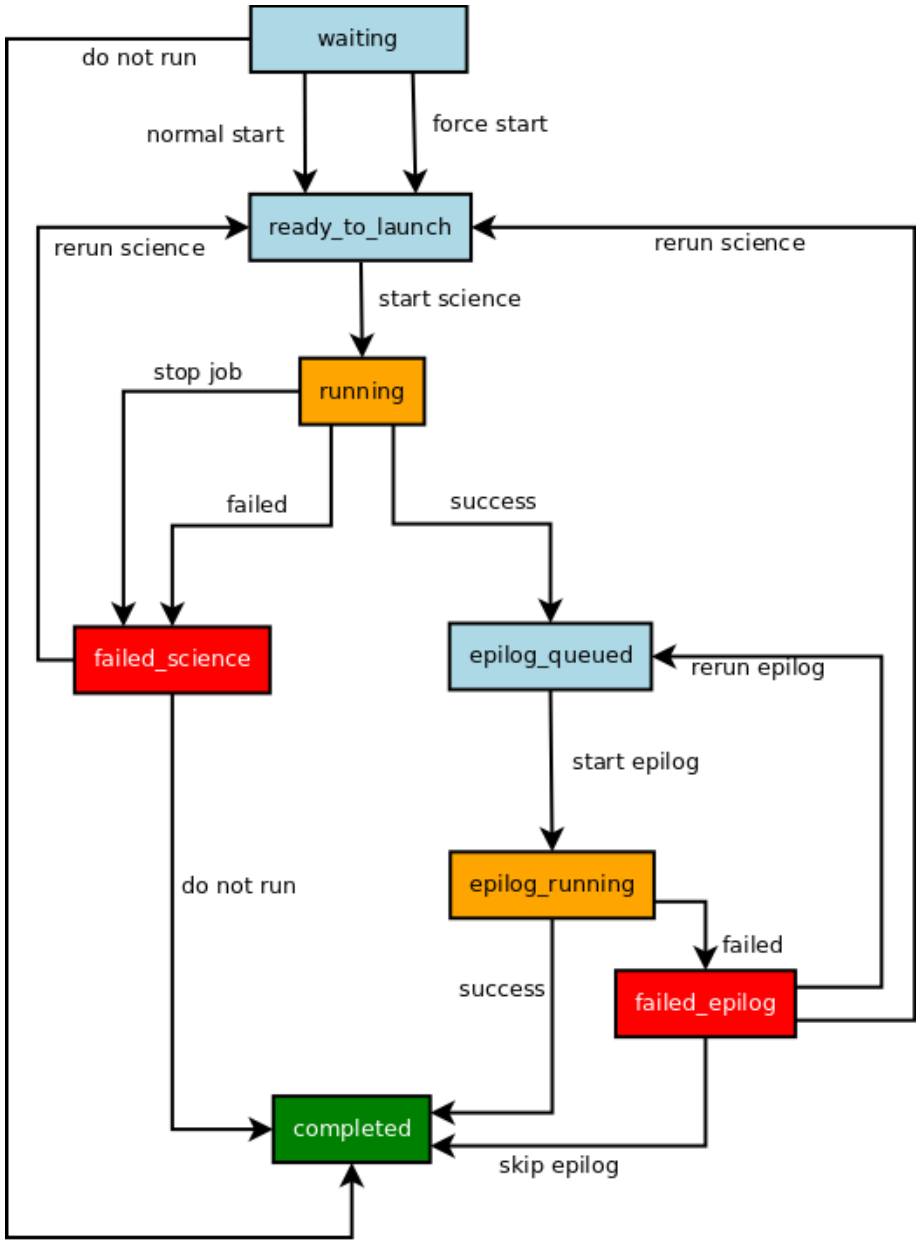

Figure 2-26. CATALYST Job State Flow

a) waiting – The CATALYST job is waiting for all of its inputs to be checked off as being available. CATALYST gleans whether an input is available from the CATALYST log database (See the CATALYST Server Operator's Manual for more details). A CATALYST job that is waiting can move to the ready_to_launch state or the completed state. When a waiting job moves directly to completed it is because CATALYST has marked its science process and epilog process statuses as CATALYST_WONT_RUN. This is due to either to known missing inputs from the log database (the required input is known to never be available for this job's datadate and the job will never run successfully).

b) ready_to_launch – The CATALYST Server has accounted for all of this job's inputs and is therefore now ready to run on the cluster. The next state for this job is that of running.
c) running – This CATALYST job is running its science process on the cluster. If successful the state moves to epilog_queued. If an error is encountered by the science process it moves to failed_science. A CATALYST Job running its science process on the cluster can also be stopped manually (see Performing Operations on CATALYST Jobs) in which case it also moves to the failed_science state.
d) failed_science – This CATALYST job has either been stopped while running on the cluster manually or its science process resulted in an error. It can only move from this state via manually re-running the science process, in which case it moves to ready_to_run, or by marking the job as CATALYST_WONT_RUN, in which case it moves to completed.
e) epilog_queued – This CATALYST job is ready to run its epilog process and has been added to the queue for epilog processes. This state moves to the epilog_running state.
f) epilog_running – This CATALYST job's epilog process is being run by the CATALYST Server. If successful the state moves to completed. Alternatively, if there was an error with the epilog process, the CATALYST job state moves to failed_epilog.
g) failed_epilog – This CATALYST job's epilog process encountered an error. At this point, its state can move to epilog_queued by manually re-running the epilog process, or to completed by marking the epilog as being skipped. See Performing Operations on CATALYST Jobs for details on these operations.
h) completed – This CATALYST job's science and epilog process have finished running or have been marked that they will not run.

2.1.4.2.10 Science Status

Shows the science process' exit status. The Operator's Console displays the science process' status code by its name instead of its code number. See the CERES Standard Exit Codes document for the full listing of exit codes. Hovering over this column with the mouse in the Operator's Console will bring up a description of the status.

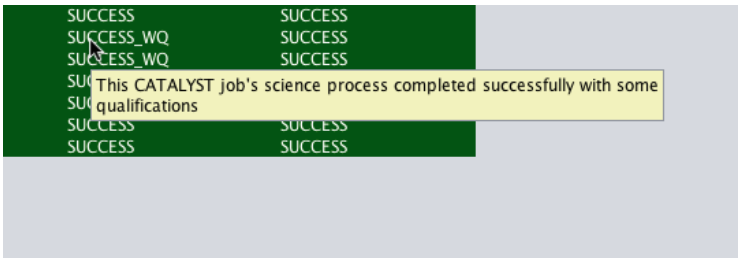

Figure 2-27. Science Status

2.1.4.2.11 Epilog Status

Shows the epilog process' exit status. Hovering over this column with the mouse in the Operator's Console will bring up a description of the status like with the science status. Possible statuses for epilog processes are listed below:

a) EPILOG_SKIPPED – The CATALYST Job's epilog has been manually marked as skipped. The CATALYST Server will skip over epilogs with this status.
b) CATALYST_WONT_RUN – CATALYST has automatically marked that this epilog will not be run. This is set when CATALYST knows this epilog cannot be run based on the science process' exit code or because of an error condition.
c) SUCCESS – The epilog process completed successfully.
d) VERIFY_ERROR – The epilog process resulted in a verification error.
e) PARSE_ERROR – The epilog process resulted in a parsing error.
f) ARCHIVE_ERROR – The epilog process resulted in an archiving error.
g) REMOVE_ERROR – The epilog process resulted in a removal error.
h) UREMOVE_ERROR – The epilog process resulted in an uremoval error.
i) REPORT_ERROR – The epilog process resulted in a report error.
j) OTHER_ERROR – The epilog process encountered a non-standard error.

For more details on the epilog statuses, please contact the development team responsible for epilog wrapper scripts.

2.1.4.3 Quickly Navigating to Failed CATALYST Jobs

In the section, Determining the Status of CATALYST Jobs At a Glance, row colors based on status and columns for total failures were described. You can quickly navigate to a failed job by following (drilling down - double clicking on the row) either the red row color or by following the total failure columns (for science or epilog) where there is greater than zero failures until the individual CATALYST jobs with datadates are listed. At the lowest level, you should see the individual failed job or jobs.

2.1.4.4 Searching for CATALYST Jobs by Datadate and Status Code

To search for CATALYST jobs in the selected PR (see Viewing a PR's CATALYST Jobs) by their datadate or status code, press the search button ( ) to open the search dialog.

2.1.4.4.1 Searching by datadate

To search for jobs that fall under a particular date, click the top left dropdown box and select "Search by Datadate", if it's not already selected. Next click on the search text box in the top middle and type in either a specific datadate (in the format of YYYYMMDD or YYYYMMZZ if it is a zonal datadate) or a general date you want to see the jobs under. For instance, if you wanted to see all the jobs for January 2005, you would type: 200501. You can also search for a particular status code under this date using the status code dropdown boxes. Once ready, click the search button to begin the search.

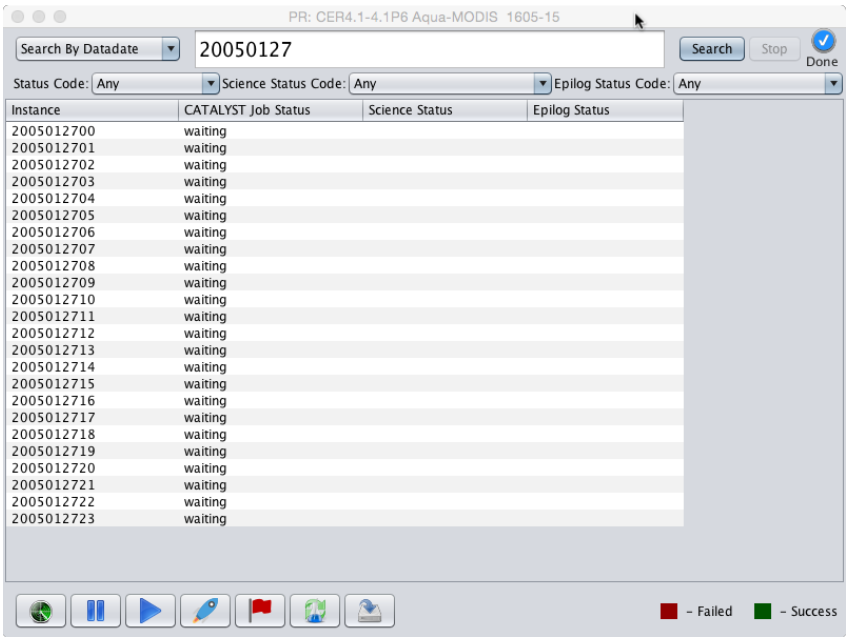

Figure 2-28. Search by Datadate

2.1.4.4.2 Searching by status code

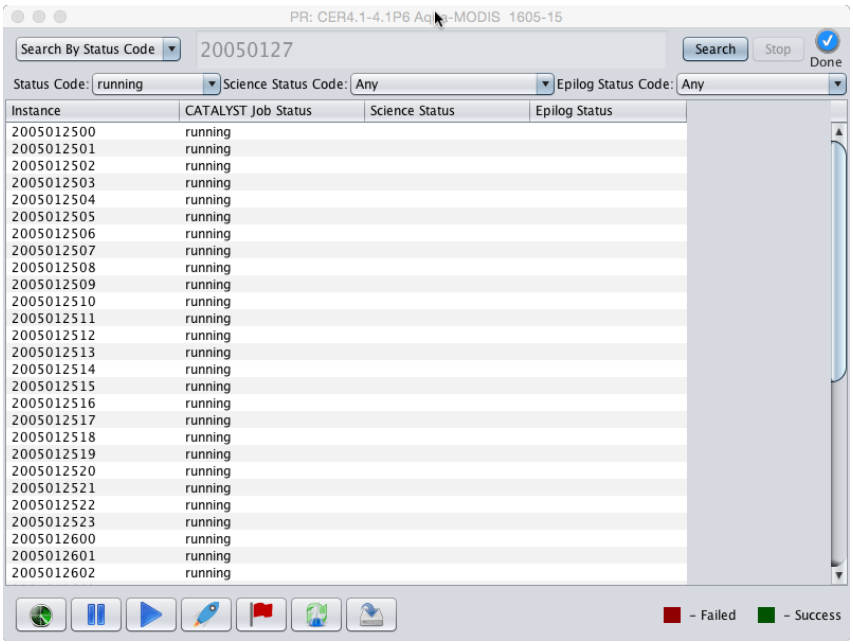

Figure 2-29. Search by Status Code

To search for a particular status code click the desired drop down, then click the status code you want to search for. You can search for any combination of three status dropdowns: CATALYST job status code, science status code, or epilog status code. Once ready, click the search button to begin the search.

2.1.4.5 Performing Operations on CATALYST Jobs

Operations can be performed on CATALYST jobs by using the buttons at the bottom of the Job View or by right clicking on a specific CATALYST job to open up the right click menu.

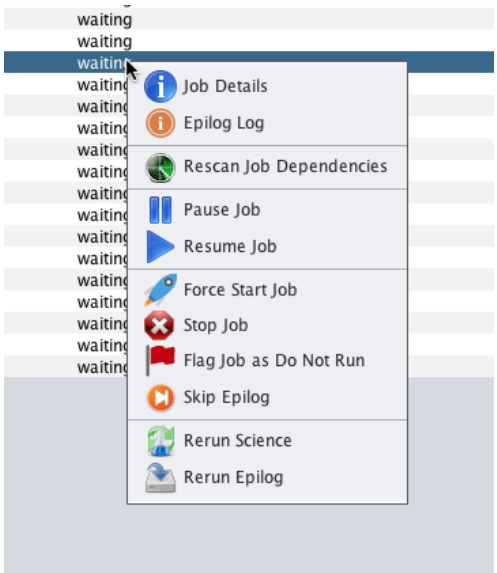

Figure 2-30. Job Operations

To select multiple jobs click and drag your mouse to highlight multiple rows of jobs, or alternatively before drilling down all the way to the individual jobs, select one or more of the rows (see Navigating the Job View) that contain multiple jobs (total given by the total column – see Determining the Status of CATALYST Jobs At a Glance). When performing an action on multiple jobs, a window containing a progress bar as the operations are submitted to the CATALYST server will appear.

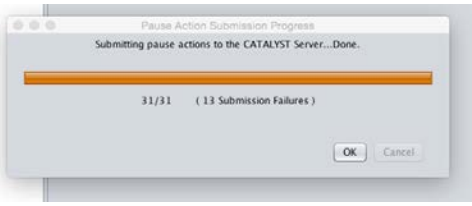

Figure 2-31. Pause Action Submission Status Bar

Clicking cancel will stop the actions where they are and show a summary of the ones that have completed. After finishing all the requests, clicking ok will also bring up the summary window.

The summary window lists each job, if any, where the operation failed and why it failed along with any additional message form the server if any.

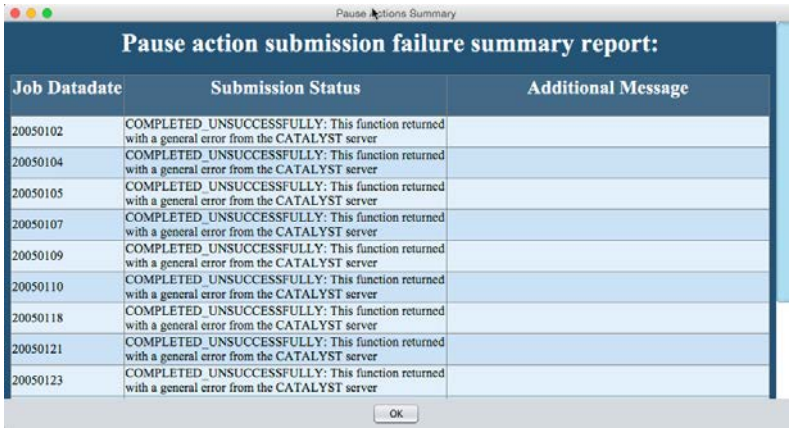

Job Datadate	Submission Status	Additional Message
20050102	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050104	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050105	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050107	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050109	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050110	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050118	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050121	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	
20050123	COMPLETED_UNSUCCESSFULLY: This function returned with a general error from the CATALYST server	

Figure 2-32. Pause Actions Summary

Below the operations that can be performed on CATALYST jobs are detailed. Note all job operation buttons have detailed hovertext in the Operator's Console when you hover the mouse over them.

2.1.4.5.1 Rescan ( )

Rescanning tells CATALYST to recheck the job's inputs to see if they are now available.

2.1.4.5.2 Pause ( )

Pauses this job in CATALYST, preventing it from proceeding to another state. Paused jobs show the current CATALYST job state with "(PAUSED)" beside it.

2.1.4.5.3 Resume ( )

Unpauses the job in CATALYST, allowing it to proceed as normal.

2.1.4.5.4 Force Start ( )

Forces this job to start, regardless of whether CATALYST has confirmed all of its inputs as being available. Note: the job submission script for the PGE associated with the job can still cause the job to fail if the submission script determines too many inputs are missing and in this case the science process will typically result with a JSS_ERROR.

2.1.4.5.5 Flag as Won't Run ( )

Flags the job as CATALYST_WONT_RUN for both the science and epilog process. It also notifies any other CATALYST jobs that were waiting for this job for input. You may want to flag a job as wont run if you know it will never have enough input to run successfully and it is holding up follow on jobs. If you are unsure contact the CATALYST development team.

2.1.4.5.6 Rerun Science ( )

Reruns the science process for this job. You may want to rerun the science process for a job that had a science failure if you have since fixed the issue that was causing it to fail, for instance.

2.1.4.5.7 Rerun Epilog ( )

Reruns the epilog process for this job. You may want to rerun a failed epilog process if the problem causing the failure has been cleared up, for instance.

2.1.4.5.8 Stop ( )

Stops the job if its science process running on the cluster currently and returns it to the ready_to_run state. This operation is only accessible from the right click menu.

2.1.4.5.9 Skip Epilog ( )

Notifies CATALYST that this job's epilog process can be skipped. Use caution when skipping an epilog process. It should only be skipped if the epilog will never be successful and will permanently prevent other epilogs in the queue from running. If you are unsure, contact the CATALYST development team. This operation is only found in the right click menu.

2.1.4.6 Viewing a CATALYST Job's Details

Once you have navigated to a list of specific CATALYST jobs under a given date (see Navigating the Job View) you can double click on the row containing the job to open up the job's details. A new window will open containing the jobs details.

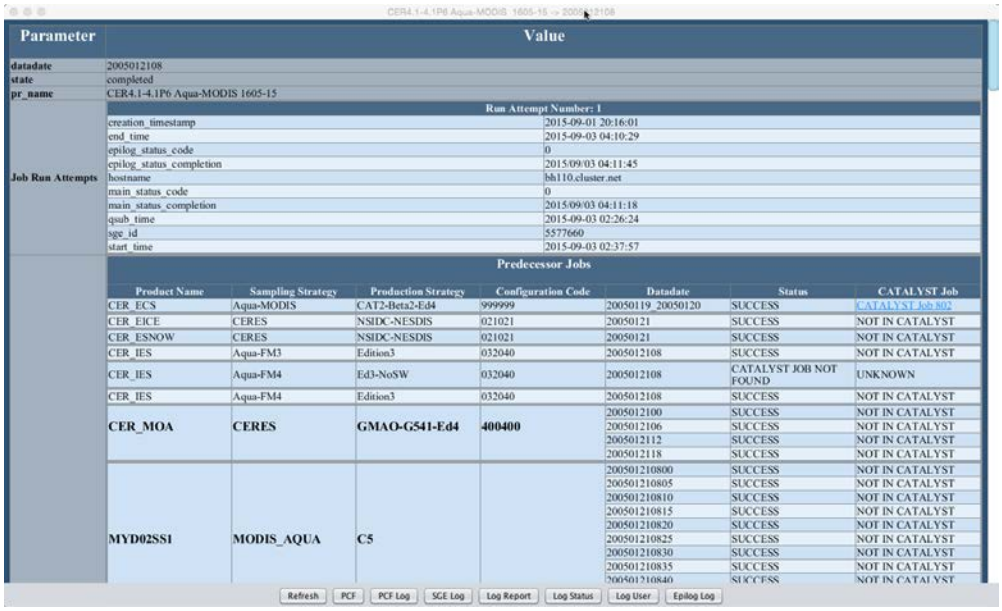

Parameter	Value
datadate	2005012108
state	completed
pr_name	CER4.1-4.1P6 Aqua-MODIS 1605-15
Job Run Attempts	Run Attempt Number: 1
	creation_timestamp: 2015-09-01 20:16:01
	end_time: 2015-09-03 04:10:29
	epilog_status_code: 0
	epilog_status_completion: 2015/09/03 04:11:45
	hostname: bh110.cluster.net
	main_status_code: 0
	main_status_completion: 2015/09/03 04:11:18
	qsub_time: 2015-09-03 02:26:24
	sge_id: 5577660
	start_time: 2015-09-03 02:37:57

Predecessor Jobs

Product Name	Sampling Strategy	Production Strategy	Configuration Code	Datadate	Status	CATALYST Job
CER_ECS	Aqua-MODIS	CAT2-Beta2-Ed4	999999	20050119_20050120	SUCCESS	CATALYST Job 802
CER_EICE	CERES	NSIDC-NESDIS	021021	20050121	SUCCESS	NOT IN CATALYST
CER_ESNOW	CERES	NSIDC-NESDIS	021021	20050121	SUCCESS	NOT IN CATALYST
CER_IES	Aqua-FM3	Edition3	032040	2005012108	SUCCESS	NOT IN CATALYST
CER_IES	Aqua-FM4	Ed3-NoSW	032040	2005012108	CATALYST JOB NOT FOUND	UNKNOWN
CER_IES	Aqua-FM4	Edition3	032040	2005012108	SUCCESS	NOT IN CATALYST
CER_MOA	CERES	GMAO-G541-Ed4	400400	2005012100	SUCCESS	NOT IN CATALYST
				2005012106	SUCCESS	NOT IN CATALYST
				2005012112	SUCCESS	NOT IN CATALYST
				2005012118	SUCCESS	NOT IN CATALYST
MYD02SS1	MODIS_AQUA	C5		200501210800	SUCCESS	NOT IN CATALYST
				200501210805	SUCCESS	NOT IN CATALYST
				200501210810	SUCCESS	NOT IN CATALYST
				200501210815	SUCCESS	NOT IN CATALYST
				200501210820	SUCCESS	NOT IN CATALYST
				200501210825	SUCCESS	NOT IN CATALYST
				200501210830	SUCCESS	NOT IN CATALYST
				200501210835	SUCCESS	NOT IN CATALYST
				200501210840	SUCCESS	NOT IN CATALYST

Refresh | PCF | PCF Log | SGE Log | Log Report | Log Status | Log User | Epilog Log

Figure 2-33. Job Details Window Part 1

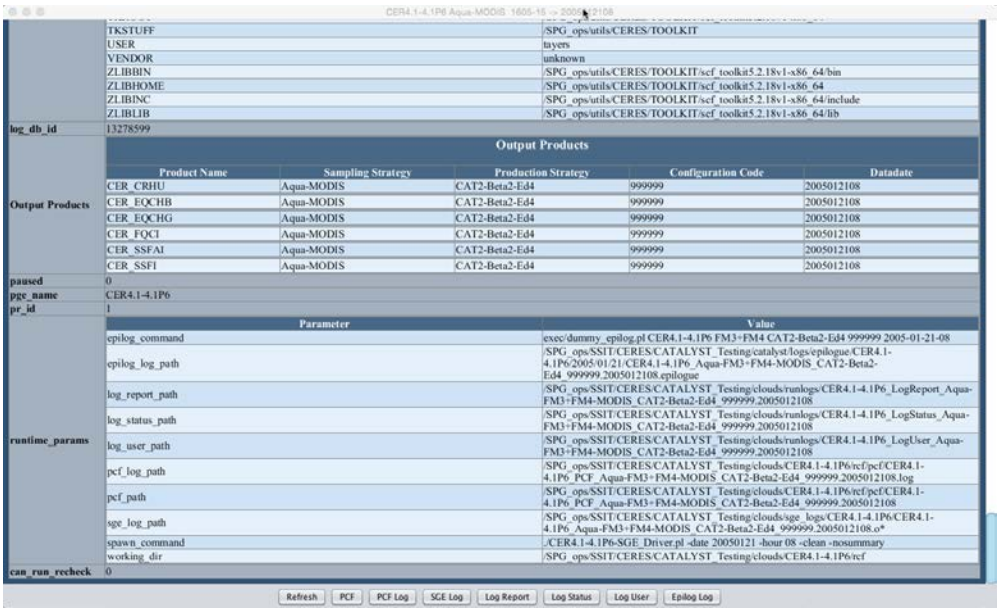

	TKSTUFF	/SPG_ops/utils/CERES/TOOLKIT
	USER	tayers
	VENDOR	unknown
	ZLIBBIN	/SPG_ops/utils/CERES/TOOLKIT/scf_toolkit5.2.18v1-x86_64/bin
	ZLIBHOME	/SPG_ops/utils/CERES/TOOLKIT/scf_toolkit5.2.18v1-x86_64
	ZLIBINC	/SPG_ops/utils/CERES/TOOLKIT/scf_toolkit5.2.18v1-x86_64/include
	ZLIBLIB	/SPG_ops/utils/CERES/TOOLKIT/scf_toolkit5.2.18v1-x86_64/lib
log_db_id	13278599	

Output Products

	Product Name	Sampling Strategy	Production Strategy	Configuration Code	Datadate
Output Products	CER_CRHU	Aqua-MODIS	CAT2-Beta2-Ed4	999999	2005012108
	CER_EQCHB	Aqua-MODIS	CAT2-Beta2-Ed4	999999	2005012108
	CER_EQCHG	Aqua-MODIS	CAT2-Beta2-Ed4	999999	2005012108
	CER_FQCI	Aqua-MODIS	CAT2-Beta2-Ed4	999999	2005012108
	CER_SSFAI	Aqua-MODIS	CAT2-Beta2-Ed4	999999	2005012108
	CER_SSFI	Aqua-MODIS	CAT2-Beta2-Ed4	999999	2005012108

paused	0	
pge_name	CER4.1-4.1P6	
pr_id	1	

	Parameter	Value
runtime_params	epilog_command	exec/dummy_epilog.pl CER4.1-4.1P6 FM3+FM4 CAT2-Beta2-Ed4 999999 2005-01-21-08
	epilog_log_path	/SPG_ops/SSIT/CERES/CATALYST_Testing/catalyst/logs/epilogue/CER4.1-4.1P6/2005/01/21/CER4.1-4.1P6_Aqua-FM3+FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108.epilogue
	log_report_path	/SPG_ops/SSIT/CERES/CATALYST_Testing/clouds/runlogs/CER4.1-4.1P6_LogReport_Aqua-FM3+FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108
	log_status_path	/SPG_ops/SSIT/CERES/CATALYST_Testing/clouds/runlogs/CER4.1-4.1P6_LogStatus_Aqua-FM3+FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108
	log_user_path	/SPG_ops/SSIT/CERES/CATALYST_Testing/clouds/runlogs/CER4.1-4.1P6_LogUser_Aqua-FM3+FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108
	pcf_log_path	/SPG_ops/SSIT/CERES/CATALYST_Testing/clouds/CER4.1-4.1P6/rcf/pcf/CER4.1-4.1P6_PCF_Aqua-FM3+FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108.log
	pcf_path	/SPG_ops/SSIT/CERES/CATALYST_Testing/clouds/CER4.1-4.1P6/rcf/pcf/CER4.1-4.1P6_PCF_Aqua-FM3+FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108
	sge_log_path	/SPG_ops/SSIT/CERES/CATALYST_Testing/clouds/sge_logs/CER4.1-4.1P6/CER4.1-4.1P6_Aqua-FM3+FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108.o*
	spawn_command	./CER4.1-4.1P6-SGE_Driver.pl -date 20050121 -hour 08 -clean -nosummary
	working_dir	/SPG_ops/SSIT/CERES/CATALYST_Testing/clouds/CER4.1-4.1P6/rcf
can_run_recheck	0	

Refresh | PCF | PCF Log | SGE Log | Log Report | Log Status | Log User | Epilog Log

Figure 2-34. Job Details Window Part 2

While most of the job details content is only informative, the predecessor jobs table has links to the details of those jobs if they happen to be CATALYST jobs currently in the CATALYST server. Click on their link to bring up the details for that job.

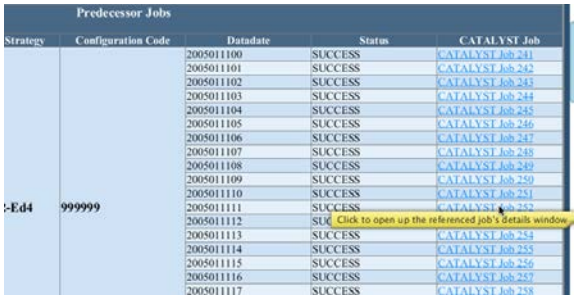

Predecessor Jobs

Strategy	Configuration Code	Datadate	Status	CATALYST Job
		2005011100	SUCCESS	CATALYST Job 241
		2005011101	SUCCESS	CATALYST Job 242
		2005011102	SUCCESS	CATALYST Job 243
		2005011103	SUCCESS	CATALYST Job 244
		2005011104	SUCCESS	CATALYST Job 245
		2005011105	SUCCESS	CATALYST Job 246
		2005011106	SUCCESS	CATALYST Job 247
		2005011107	SUCCESS	CATALYST Job 248
		2005011108	SUCCESS	CATALYST Job 249
		2005011109	SUCCESS	CATALYST Job 250
		2005011110	SUCCESS	CATALYST Job 251
-Ed4	999999	2005011111	SUCCESS	CATALYST Job 252
		2005011112	SUC	Click to open up the referenced job's details window
		2005011113	SUCCESS	CATALYST Job 254
		2005011114	SUCCESS	CATALYST Job 255
		2005011115	SUCCESS	CATALYST Job 256
		2005011116	SUCCESS	CATALYST Job 257
		2005011117	SUCCESS	CATALYST Job 258

Figure 2-35. Links to Predecessor Jobs

There are several log files for each CATALYST job. They can be viewed, if available, by clicking on the log buttons at the bottom of the job details window. Any underlined links in these logs can be clicked on to jump to different points in the log. The logs are detailed below:

2.1.4.6.1 PCF

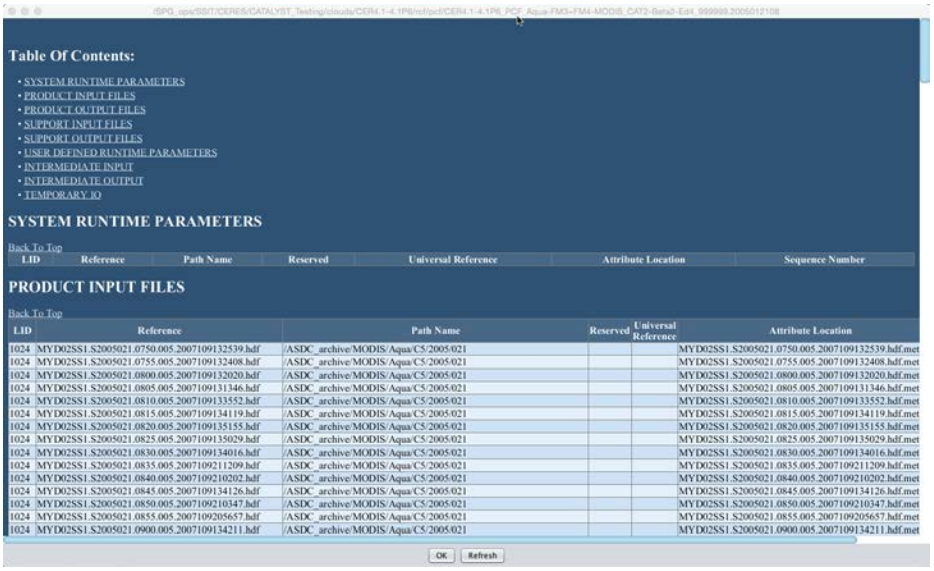

Table Of Contents:

- SYSTEM RUNTIME PARAMETERS
- PRODUCT INPUT FILES
- PRODUCT OUTPUT FILES
- SUPPORT INPUT FILES
- SUPPORT OUTPUT FILES
- USER DEFINED RUNTIME PARAMETERS
- INTERMEDIATE INPUT
- INTERMEDIATE OUTPUT
- TEMPORARY IO

SYSTEM RUNTIME PARAMETERS

Back To Top

LID	Reference	Path Name	Reserved	Universal Reference	Attribute Location	Sequence Number

PRODUCT INPUT FILES

Back To Top

LID	Reference	Path Name	Reserved	Universal Reference	Attribute Location
1024	MYD02SS1.S2005021.0750.005.2007109132539.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0750.005.2007109132539.hdf.met
1024	MYD02SS1.S2005021.0755.005.2007109132408.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0755.005.2007109132408.hdf.met
1024	MYD02SS1.S2005021.0800.005.2007109132020.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0800.005.2007109132020.hdf.met
1024	MYD02SS1.S2005021.0805.005.2007109131346.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0805.005.2007109131346.hdf.met
1024	MYD02SS1.S2005021.0810.005.2007109133552.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0810.005.2007109133552.hdf.met
1024	MYD02SS1.S2005021.0815.005.2007109134119.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0815.005.2007109134119.hdf.met
1024	MYD02SS1.S2005021.0820.005.2007109135155.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0820.005.2007109135155.hdf.met
1024	MYD02SS1.S2005021.0825.005.2007109135029.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0825.005.2007109135029.hdf.met
1024	MYD02SS1.S2005021.0830.005.2007109134016.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0830.005.2007109134016.hdf.met
1024	MYD02SS1.S2005021.0835.005.2007109211209.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0835.005.2007109211209.hdf.met
1024	MYD02SS1.S2005021.0840.005.2007109210202.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0840.005.2007109210202.hdf.met
1024	MYD02SS1.S2005021.0845.005.2007109134126.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0845.005.2007109134126.hdf.met
1024	MYD02SS1.S2005021.0850.005.2007109210347.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0850.005.2007109210347.hdf.met
1024	MYD02SS1.S2005021.0855.005.2007109205657.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0855.005.2007109205657.hdf.met
1024	MYD02SS1.S2005021.0900.005.2007109134211.hdf	/ASDC_archive/MODIS/Aqua/C5/2005/021			MYD02SS1.S2005021.0900.005.2007109134211.hdf.met

OK Refresh

Figure 2-36. PCF Part 1

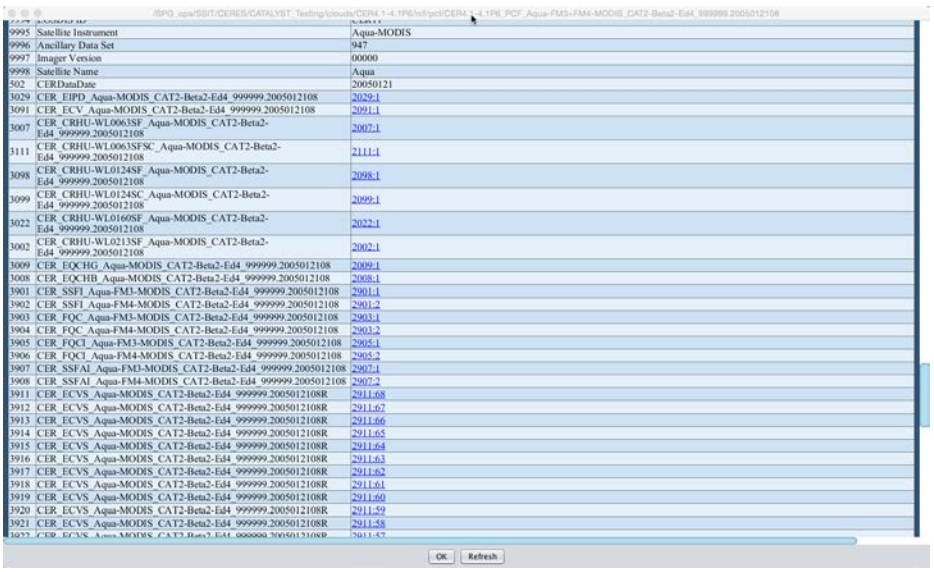

9995	Satellite Instrument	Aqua-MODIS
9996	Ancillary Data Set	947
9997	Imager Version	00000
9998	Satellite Name	Aqua
502	CERDataDate	20050121
3029	CER_EIPD_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2029:1
3091	CER_ECV_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2091:1
3007	CER_CRHU-WL0063SF_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2007:1
3111	CER_CRHU-WL0063SFSC_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2111:1
3098	CER_CRHU-WL0124SF_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2098:1
3099	CER_CRHU-WL0124SC_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2099:1
3022	CER_CRHU-WL0160SF_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2022:1
3002	CER_CRHU-WL0213SF_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2002:1
3009	CER_EQCHG_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2009:1
3008	CER_EQCHB_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2008:1
3901	CER_SSFI_Aqua-FM3-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2901:1
3902	CER_SSFI_Aqua-FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2901:2
3903	CER_FQC_Aqua-FM3-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2903:1
3904	CER_FQC_Aqua-FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2903:2
3905	CER_FQCI_Aqua-FM3-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2905:1
3906	CER_FQCI_Aqua-FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2905:2
3907	CER_SSFAI_Aqua-FM3-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2907:1
3908	CER_SSFAI_Aqua-FM4-MODIS_CAT2-Beta2-Ed4_999999.2005012108	2907:2
3911	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:68
3912	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:67
3913	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:66
3914	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:65
3915	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:64
3916	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:63
3917	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:62
3918	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:61
3919	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:60
3920	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:59
3921	CER_ECVS_Aqua-MODIS_CAT2-Beta2-Ed4_999999.2005012108R	2911:58

OK Refresh

Figure 2-37. PCF Part 2

2.1.4.6.2 PCF Log

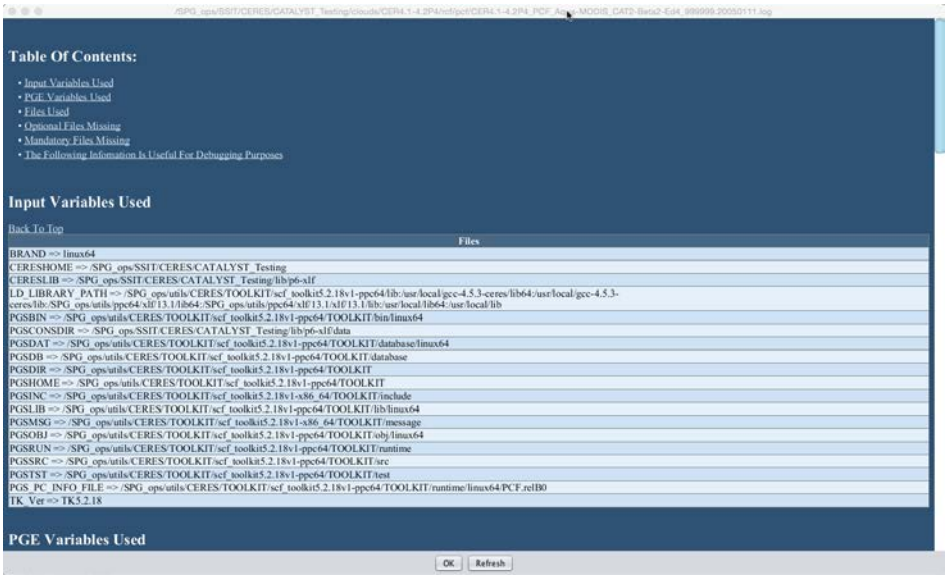

Figure 2-38. PCF Log

2.1.4.6.3 SGE Log

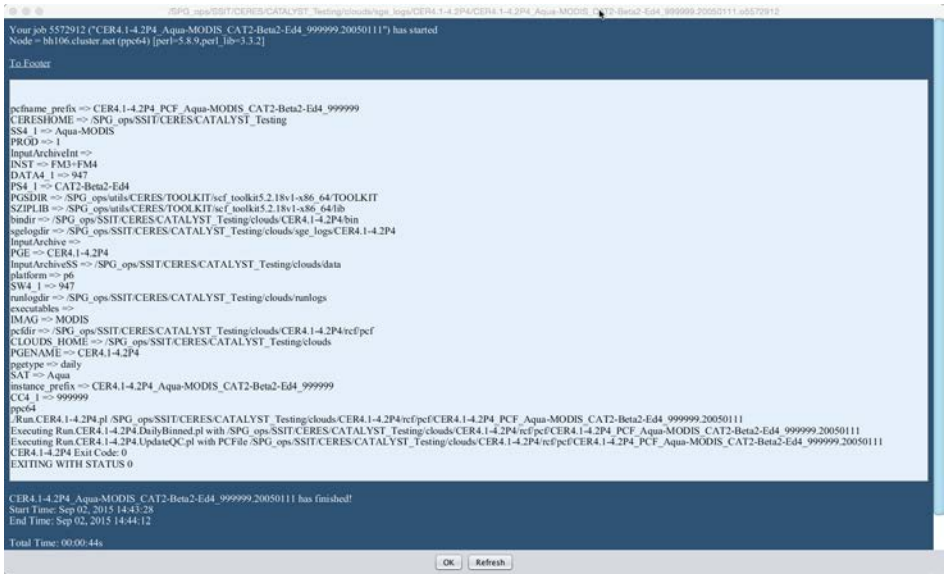

Figure 2-39. SGE Log

2.1.4.6.4 Log Report

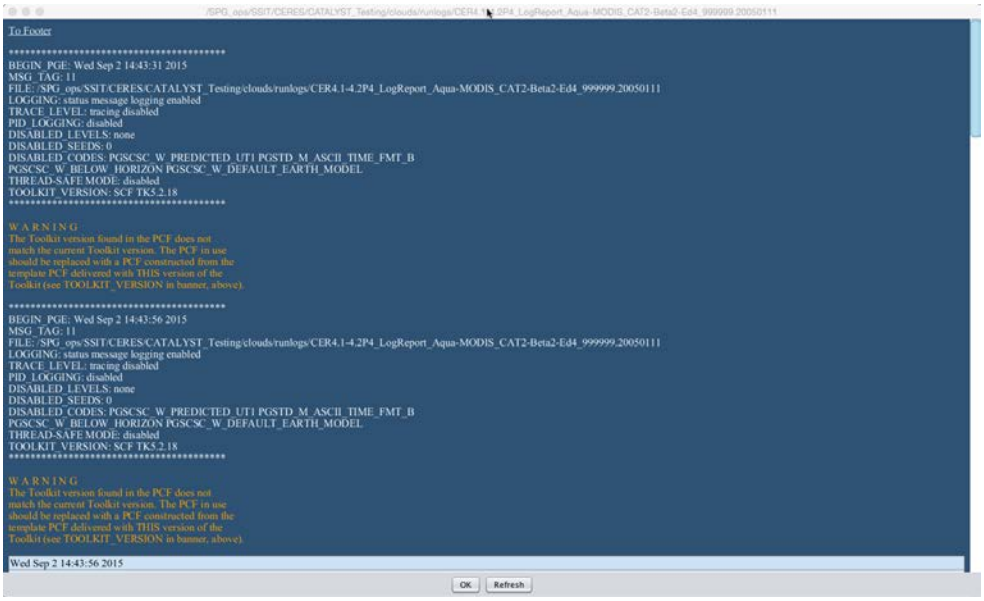

Figure 2-40. Log Report

2.1.4.6.5 Log User

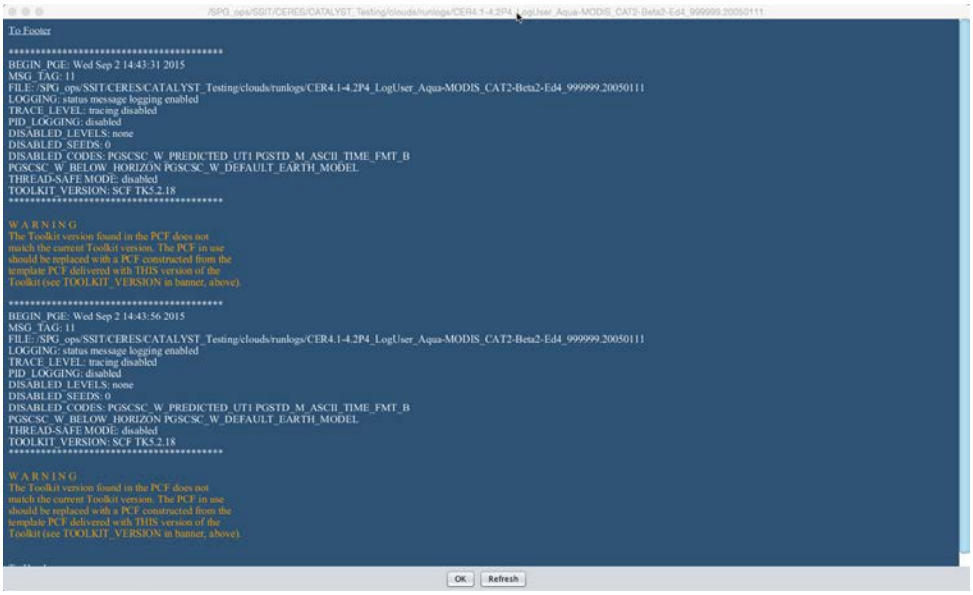

Figure 2-41. Log User

2.1.4.6.6 Log Status

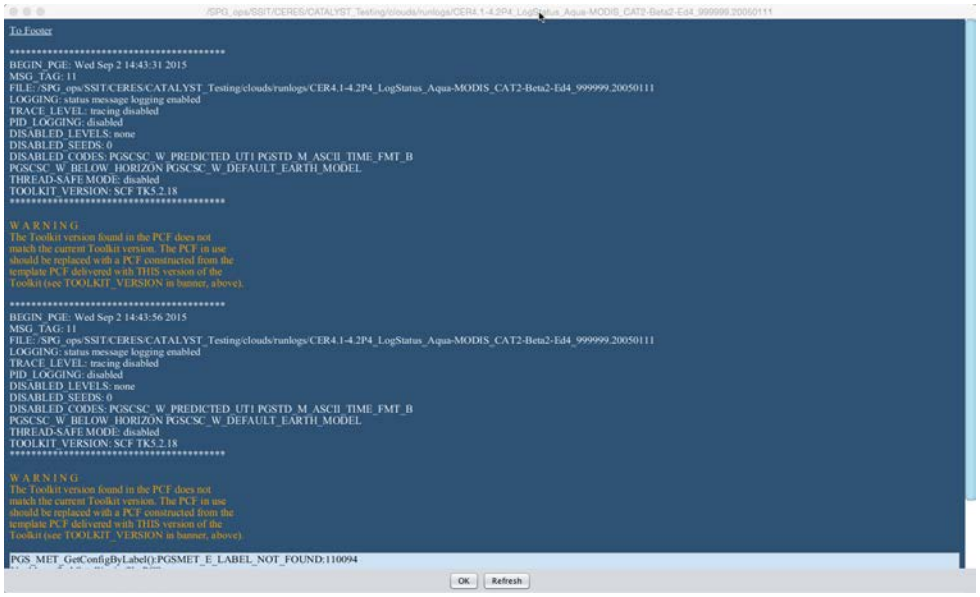

Figure 2-42. Log Status

2.1.4.6.7 Epilog Log

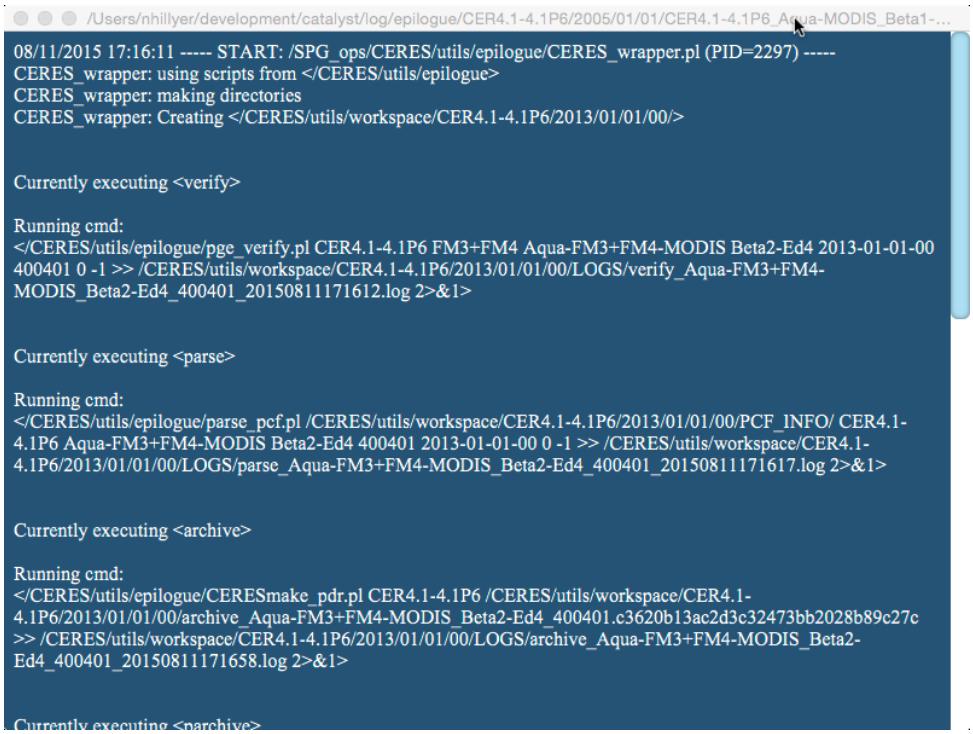

Figure 2-43. Epilog Log

2.1.5 Log Console

This area of the Operator's Console contains any log messages the program creates. Since the Operator's Console also writes an errors or warnings to the client.log file, the Log Console can be hidden via the menu options. Messages containing [INFO] are simply informative log messages and can safely be ignored. Hiding the Log Console provides more vertical space for the Job View and the PR View.

2.2 Recovering from an Operator's Console Failure

Refer to Appendix B for general error messages a user can receive in the Operator's Console. If the Operator's Console becomes unresponsive or has a failure that results in it being inoperable, there are several steps to take:

1. Click "CATALYST Server → Disconnect" followed by trying to reconnect to the server "CATALYST Server → Connect" and re-enter the corresponding login information.
2. Close and restart the Operator's Console. This should clear up any unresponsiveness.
3. Check that the login and server information is correct and that the server is running.
4. If any error messages appear in the console at the bottom or restarting the client does not work to solve your issue, please email the most recent client.log file generated by the console to joshua.c.wilkins@nasa.gov or your designated contact for CATALYST related bugs with a detailed description of what happened.

Appendix A
Acronyms and Abbreviations

ACL	Access Control List
API	Application Programming Interface
ASDC	Atmospheric Science Data Center
CATALYST	CERES AuTomAted job Loading sYSTem
CERES	Clouds and the Earth's Radiant Energy System
CM	Configuration Management
COTS	Commercial Off The Shelf
LaRC	Langley Research Center
LDAP	Lightweight Directory Access Protocol
NASA	National Aeronautics and Space Administration
PR	Processing Request
SSAI	Science Systems and Applications, Inc.
XML-RPC	Extensible Markup Language – Remote Procedure Call

Appendix B
Messages for Operator's Console

This table details the many messages the user may receive either in the log or as popup messages. These messages can be informational, warnings, or errors originating from the Operator's Console or the CATALYST server.

Table B-1. Operator's Console Messages

Message Keywords	Type	Diagnosis
SSLHandshakeException	Connection error	Check that the server is running (the master, xmlrpc and user CATALYST Server processes are running) and that your network is connected properly (or VPN if remote). Then retry to connect.
INVALID_UUID_ERROR – the submitted token was invalid	Authentication timeout error	The session with the CATALYST Server has timed out. Login to the server again.
NullPointerException	Coding error	Send the client.log file to the Operator's Console development team.
NON_EXISTENT_METHOD - Method lookup error	Server Request error	This can be caused when the server and console versions are out of sync. Check that your operator's console version matches the newest and that you are connected to the correct CATALYST server.
XMLRpcClientException – failure writing request	error attempting to send a request to the CATALYST server	Send the client.log file to the Operator's Console development team.
Lost connection to server, please reconnect.	Connection error	The server may have gone offline or network connection may have been lost (or VPN connection if remote). Check that the server is online and relogin into the server.
Code 4	Invalid Arguments error	Send the client.log to the Operator's Console development team.
Code 5	Authentication timeout error	The session with the CATALYST Server has timed out. Login to the server again.

Table B-1. Operator's Console Messages

Message Keywords	Type	Diagnosis
Code 7	Invalid Privileges	The user does not have the privileges needed to perform the request. Update the Access Control List for this user (see the CATALYST Server Operator's Manual)
Code 8	The Cluster is not ready to process jobs	Check the blade list for the status of the cluster. If there is not enough blades online this can happen. Enable blades as they become available after going down to prevent this scenario.
Code 9	Completed Unsuccessfully – general server error	If the message returned from the server is not apparent enough, contact the CATALYST development team for more information.
Code 12	PR no longer in the CATALYST server	This can happen if a PR was closed or deleted in CATALYST and removed by another user. The PR list should update momentarily with the correct list of PRs.
Code 15	CATALYST Server internal timeout	The server took too long for one of its processes to respond. Check that the expected server processes are online in the Server Status window. CATALYST Server -> Server Status in the menu. Contact the CATALYST development team for any questions.
Code 17	Unknown Error	Contact the CATALYST development team with details about the error.
Code 18	Nonexistent object error	The item you were requesting no longer exists in the CATALYST server. This can happen when another user removes a PR or if the job log files do not exist yet or have been removed. No action necessary.

Table B-1. Operator's Console Messages

Message Keywords	Type	Diagnosis
Code 21	PR Locked	This can happen when a user tries to perform an operation on a job while the PR is locked which prevents operations from being done on jobs for that PR. Another user may have locked the PR that has the proper privileges. No action necessary.
Code 22	Server not initialized	The server is not yet ready to process jobs and is in read only mode. Have a privileged user tell the server to start processing (CATALYST -> Start Processing). This can also occur when the kernel process has been restarted or gone offline due to an error. The CATALYST server defaults to "not ready to process" when the kernel process first starts.