

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

Data Management System

Operator's Manual

**CERES AuTomAteD job Loading sYSTem
(CATALYST) Server**

Version 2

Primary Authors

T. Nelson Hillyer, 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: January 2017
Document Date: January 2017

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	1065	<ul style="list-style-type: none"> • Separated document from original combined CATALYST operator's manual. • Initial version of document. (09/06/2015) 	All All
09/28/2016	V2	1176	<ul style="list-style-type: none"> • Added new configuration settings for CATALYST v2.1. 	Sec. 1.3

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 CATALYST software, and outlines installation, execution, and monitoring procedures. In addition, all CATALYST error messages and subsequent actions required by the ASDC operations staff are included.

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

TABLE OF CONTENTS

<u>Section</u>	<u>Page</u>
Document Revision Record	ii
Preface.....	iii
1.0 CATALYST Server	3
1.1 Server Details	3
1.1.1 Responsible Persons.....	3
1.2 Operating Environment	3
1.2.1 Runtime Parameters	3
1.2.2 Environment Script Requirements	3
1.2.3 Execution Frequency	3
1.2.4 Memory/Disk Space/Time Requirements.....	3
1.3 Server Configuration Parameters	4
1.4 Operating Procedures	6
1.4.1 Installing the CATALYST Server Software.....	6
1.4.2 PGE Handler Management	6
1.4.2.1 Installing PGE Handlers.....	7
1.4.2.2 Removing PGE Handlers	7
1.4.3 Purpose of CATALYST Components	7
1.4.4 How to Start the CATALYST Server.....	8
1.4.5 How to Verify the CATALYST Server is Running.....	9
1.4.6 Editing the CATALYST Server's Access Control List.....	9
1.4.7 Stopping the CATALYST Server.....	10
1.5 Expected PR Instantiation Time.....	10
1.6 Execution Evaluation	11
1.6.1 Exit Codes.....	11

TABLE OF CONTENTS

<u>Section</u>	<u>Page</u>
1.7 Operational Log Files/Directories.....	11
1.8 Operational Files/Databases.....	12
1.9 CATALYST Monitoring Utility	13
1.9.1 Installation of the Monitoring Utility.....	13
1.10 Best Practices	15
Appendix A - Acronyms and Abbreviations	A-1
Appendix B - Product Availability in CATALYST Database.....	B-1
B.1 CATALYST 1.0.4 Installation Package + Logging Database Patch (SCCR 996)	B-1

LIST OF TABLES

<u>Table</u>	<u>Page</u>
Table 1-1. CATALYST Software Developer Contacts	3
Table 1-2. Memory/Disk Space/Time Requirements for CATALYST Server	3
Table 1-3. Configuration Parameters	4
Table 1-4. CATALYST Components	7
Table 1-5. CATALYST Privilege Categories.....	10
Table 1-6. Exit Codes for CATALYST Server	11
Table 1-7. Log File Listing for CATALYST Server	12
Table 1-8. Operational File Listing for CATALYST Server.....	13

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 Manual](#) is part of the CERES CATALYST delivery package provided to the Langley Atmospheric Science Data Center (ASDC). It provides a description of the CERES CATALYST software and explains the procedures for executing the software. A description of acronyms and abbreviations is provided in [Appendix A](#), and a Product Availability in CATALYST Database list is provided in [Appendix B](#).

This document is organized as follows:

[Introduction](#)

[Document Overview](#)

[CATALYST System Overview](#)

Section 1.0 CATALYST Server

[Appendix A](#) - Acronyms and Abbreviations

[Appendix B](#) - Error Messages for CATALYST

CATALYST System Overview

CATALYST is a semi-automated workflow manager designed and implemented by the CERES team to streamline production processing. It consists of a server and Operator's Console. The server is responsible for ingesting a PR sent from the PR Web Application, generating specific instances of the PGE within the PR, determining if all data inputs are available through the logging database, handling the submission and monitoring of PGE instances, and updating the logging database with status. The Operator's Console allows tracking of job instances and provides the ability for an operator to start and stop PGE instances. For user authentication to CATALYST and the PRDB, the *AMI* LDAP is used.

CATALYST Server

The CATALYST server is an application that provides a service to coordinate the execution of CERES PGEs in a production environment. The server runs as a collection of Perl daemon processes that provides a front-end XML-RPC API with which external applications (both Perl and non-Perl) can connect and query/modify the state of the CATALYST server. The CATALYST server ingests PRs sent from the CERES PR Tool. The system can accept all PRs for a given processing stream of interdependent PGEs. The system builds job instance information from ingested PRs and implements PGE specific logic to continuously poll for availability of input file preconditions through the logging database. CATALYST uses a logging database to store job instance information. It can also submit the epilogue scripts to archive output files.

A handler is needed for each PGE that CATALYST manages. The handler contains specific logic for creating job instances and determining its ability to run. Please refer to the CATALYST PGE Handler Document on the CATALYST Home Page for a listing of CATALYST supported PGEs.

CATALYST Operator's Console

For more information on the CATALYST Operator's Console, please refer to the CATALYST Operator's Console Operator's Manual available on the CERES Online Documentation website.

1.0 CATALYST Server

1.1 Server Details

1.1.1 Responsible Persons

Table 1-1. CATALYST Software Developer Contacts

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

1.2 Operating Environment

1.2.1 Runtime Parameters

The CATALYST command line tools are argument-less.

1.2.2 Environment Script Requirements

CATALYST only requires that the CERES environment script and the ASDC epilogue environment script be sourced prior to starting.

1.2.3 Execution Frequency

Single Instance – This service runs in the background as a daemon; only a single instance of this program is intended to run in the production environment at any moment of time.

1.2.4 Memory/Disk Space/Time Requirements

Table 1-2. Memory/Disk Space/Time Requirements for CATALYST Server

System	System Name	Hardware	Total Run Time (HR:MN:SEC)	Memory (MB)	Disk Storage (MB)
	CATALYST	x86_64	Indefinite	Dependent on PR Load	Dependent on PR Load

1.3 Server Configuration Parameters

The server has several parameters that can be adjusted in “\$CATALYST_HOME/conf/catalyst.conf”. A summary of them, along with possible values are listed in the table below; *italicized text* denotes items highly suggested to remain there as delivered defaults (unless directed by the CATALYST development team):

Table 1-3. Configuration Parameters

Parameter	Purpose	Possible Values
server-description	This is used to specify clearly in the status emails which instance of CATALYST is described.	<i>User-defined text data</i>
<i>timezone</i>	<i>The timezone the CATALYST server is running in.</i>	<i>timezone (such as US/Eastern)</i>
<i>master-log</i>	<i>The file name of the master log file.</i>	<i>path to log file</i>
<i>master-socket</i>	<i>The socket file name of the master process.</i>	<i>path to socket file</i>
minimum-free-space	The minimum number of gigabytes of free disk space. Exceeding this will result in warnings in the status email.	Integer (gigabytes)
memory-warn-threshold	The upper limit in megabytes to the maximum combined size of CATALYST processes in memory. Exceeding this value will result in warnings in the status email.	Integer (megabytes)
<i>email-address</i>	<i>The email address to appear in the from field of status emails.</i>	<i>single email address</i>
status-email-recipients	The list of recipients for status emails.	space separated list of email addresses. (please leave ceres-dmt-catalyst@lists.nasa.gov on the status email list so the development team can monitor the server effectively)
minimum-processes	The minimum processes that should be active to not result in a degraded status email	One or more of the following: master ange logdb kernel cluster xmlrpc user
<i>authen-mode</i>	<i>The mode of authenticating users.</i>	<i>LDAP</i>
ldap-server	The URL of the LDAP server to be used.	LDAP URL
<i>ldap-bind-string</i>	<i>The binding string to be used for LDAP authentication.</i>	<i>String</i>
<i>user-socket</i>	<i>The socket file to be used by the CATALYST user process.</i>	<i>path to socket file</i>
<i>user-log</i>	<i>The log file to be used for capturing messages emitted by the CATALYST user process.</i>	<i>path to log file</i>
<i>user-list</i>	<i>The file containing the list of privileged users in CATALYST.</i>	<i>path to user list file</i>

Table 1-3. Configuration Parameters

Parameter	Purpose	Possible Values
<i>user-rpc-list</i>	<i>The file containing the list of CATALYST RPCs and the group allowed to call the RPC</i>	<i>path to user RPC list file</i>
<i>user-token-expiry</i>	<i>The number of seconds a user's login token is allowed to be idle.</i>	<i>Integer (seconds)</i>
<i>kernel-socket</i>	<i>The socket file to be used by the CATALYST kernel process.</i>	<i>path to socket file</i>
<i>kernel-log</i>	<i>The log file to capture messages emitted by the CATALYST kernel process.</i>	<i>path to log file</i>
<i>event-log</i>	<i>Not used in this release. Will be removed in next release.</i>	<i>path to log file</i>
<i>storage-file</i>	<i>File containing PRs, jobs, and PGE settings for this CATALYST installation.</i>	<i>path to CATALYST storage file</i>
<i>db-buffer-file</i>	<i>Not used in this release. Will be removed in next release.</i>	<i>path to CATALYST LogDB buffer file</i>
<i>handler-dir</i>	<i>Directory containing PGE handlers.</i>	<i>path to CATALYST PGE handlers</i>
<i>prod2pge-mapfile</i>	<i>File containing to CATALYST Product-PGE mappings</i>	<i>path to PGE mapping file</i>
<i>epilog-command</i>	<i>Path to the Epilogue-ing program to be run after the PGE is complete</i>	<i>path to epilogue program</i>
<i>log-db-conn-str</i>	<i>Perl DBI connection string for LogDB</i>	<i>DBI compatible connection string</i>
<i>log-db-auth-file</i>	<i>File containing DBI authorization</i>	<i>path to DBI authorization</i>
<i>log-db-log</i>	<i>Log file containing messages emitted from the CATALYST logdb process.</i>	<i>path to log file</i>
<i>grid-engine-queues</i>	<i>Space separated list of acceptable queues to submit AJSS jobs to</i>	<i>space separated list of SGE queues</i>
<i>grid-engine-minimum-queues</i>	<i>Number of queues at a minimum for the cluster process to remain active.</i>	<i>Integer (number of queues)</i>
<i>max-grid-engine-running</i>	<i>Maximum number of SGE jobs to have active at any one time</i>	<i>Integer (maximum number of jobs)</i>
<i>grid-engine-qacct-retry-delay</i>	<i>The waiting period between subsequent attempts to query SGE's qacct</i>	<i>Integer (seconds)</i>
<i>grid-engine-qacct-retry-count</i>	<i>The number of reattempts to query a job's status in SGE's qacct</i>	<i>Integer (retry count)</i>
<i>cluster-log</i>	<i>Log file containing messages emitted by the CATALYST cluster process.</i>	<i>path to log file</i>
<i>cluster-socket</i>	<i>Socket file to be used by the CATALYST cluster process.</i>	<i>path to socket file</i>
<i>xmlrpc-port</i>	<i>Port number to be used for handling XMLRPC requests from</i>	<i>Integer (greater than 1024)</i>

Table 1-3. Configuration Parameters

Parameter	Purpose	Possible Values
	<i>the Operator's Console or PR Tool.</i>	
<i>xmlrpc-proto</i>	<i>Protocol for handling XMLRPC. HTTP or HTTPS.</i>	<i>http or https</i>
<i>xmlrpc-ssl-key</i>	<i>Key file for XMLRPC SSL encryption</i>	<i>path to SSL key file</i>
<i>xmlrpc-ssl-crt</i>	<i>Certificate file for XMLRPC SSL encryption</i>	<i>path to SSL cert file</i>
<i>xmlrpc-log</i>	<i>Log file for capturing messages emitted by the CATALYST xmlrpc process</i>	<i>path to log file</i>
<i>ange-storage-file</i>	<i>File containing ANGe email processing activity</i>	<i>path to ANGe storage file</i>
<i>ange-mbox</i>	<i>Mbox file where ANGe emails will be delivered.</i>	<i>path to mbox file</i>
<i>ange-temp-file</i>	<i>Temporary file to be used by ANGe-CATALYST email processing</i>	<i>path the ANGe temporary file</i>
<i>ange-missing-file</i>	<i>Log file containing listing of files reported as archived by ANGe but have yet to appear in AMI DPO</i>	<i>path the ANGe missing file log</i>
<i>ange-log</i>	<i>Log file containing messages emitted by the CATALYST ange process.</i>	<i>path to log file</i>

1.4 Operating Procedures

1.4.1 Installing the CATALYST Server Software

Refer to Section 2.1 in the CATALYST Test Plan document for installation instructions.

Notes:

- Report any and all installation warnings/errors to the CATALYST development team (see [Table 1-1](#)) when encountered.

1.4.2 PGE Handler Management

To provide greater degrees of flexibility and extensibility, CATALYST makes a clear separation of server code and PGE handler code. The PGE handler code is loaded dynamically at the starting of the CATALYST kernel process (described in Purpose of CATALYST Components). Changes to the contents of the PGE handlers directory (`$CERESHOME/catalyst/handlers`) requires a restart of the CATALYST kernel process.

1.4.2.1 Installing PGE Handlers

Installing a CATALYST PGE handler, which is a Perl module following the name format of “**CERX_XPX.pm**” only requires placing the file in the handlers directory (**\$CERESHOME/catalyst/handlers**). Once placing the file in the correct directory, a restart of the kernel process is required.

NOTE: When a PGE has been added and after the kernel process has restarted, the PGE (and the PGE's subsystem) must be enabled. Enabling a PGE makes it ready to accept and process PRs. You can enable the PGE (and the PGE's subsystem) using the CATALYST Operator's Console (see “PGE Settings Option” in the Operator's Manual for the CATALYST Operator's Console for details on how to do this).

1.4.2.2 Removing PGE Handlers

Removing a PGE Handler from CATALYST requires first removing the associated Perl module file (in the form of “**CERX_XPX.pm**”) in the handlers directory (**\$CERESHOME/catalyst/handlers**) and restarting the CATALYST kernel process for the change to come into effect. It is **HIGHLY** recommended that all PRs associated with the PGE being removed to first be closed or deleted prior to removing the PGE handler.

1.4.3 Purpose of CATALYST Components

For purposes of improving stability of the CATALYST system, CATALYST's functionality has been separated into several individual processes. The following table lists and describes those processes as well as describes what can be done with those processes online.

Table 1-4. CATALYST Components

Process Name	Description	Required for PR Tool Logins?	Required for Submitting PRs to CATALYST?	Required for Processing PRs in CATALYST?
master	Parent process that owns all other CATALYST processes. (everything below)	yes	yes	yes
ange	Process that receives ANGe emails, extracts ingested products listed in the emails and searches the DPO for their availability. These products will be placed into the CATALYST LogDB via the kernel and logdb processes.	no	no	no
cluster	Process that submits jobs to SGE.	no	no	yes
kernel	Process that validates PRs, builds jobs from PRs, and verifies all dependencies have been met for jobs. This process interfaces with the PGE handler code.	no	yes	yes

Table 1-4. CATALYST Components

Process Name	Description	Required for PR Tool Logins?	Required for Submitting PRs to CATALYST?	Required for Processing PRs in CATALYST?
logdb	Process that interfaces with the long term record of jobs run in CATALYST.	no	no	yes
user	Process that authenticates users and provides users with login tokens to be used with session activity.	yes	yes	no
xmlrpc	Process that manages the front-end to the CATALYST system. All communication from the PR Tool and Operator's Console flows through this.	yes	yes	no

1.4.4 How to Start the CATALYST Server

The CATALYST server first requires the CATALYST master process to be started. This master process is the coordinator/parent for all of the specialized CATALYST processes. The master is started using these commands:

```
> cd $CERESHOME/catalyst
> bin/start_catalyst.sh
```

Information regarding the startup of the CATALYST master process will appear on screen. Take note of error messages that are displayed and report them to the CATALYST development team, if necessary. Once the CATALYST master process has started successfully, the remaining CATALYST components can be started using the management utility on the command line.

```
> cd $CERESHOME/catalyst
> bin/catalyst_manager.sh
```

Running this command will present the user with a prompt, as in the following screenshot:

```

ncc1701:bin nhillyer$ ./catalyst_manager.sh
Checking "/Users/nhillyer/development/catalyst/sockets" => okay!
CATALYST Manager

Available Commands:
  help           => Prints this list of available commands
  exit           => Exits this program. (CTRL+D also works)
  start|stop all => Starts/stops all of the CATALYST processes
                  (except the master process)
  start|stop xmlrpc => Start/stop the XMLRPC interface for CATALYST

  start|stop user   => Start/stop the user login manager for CATALYST
  start|stop kernel => Start/stop the kernel process for CATALYST
  start|stop ange   => Start/stop the ANGe process for CATALYST
  start|stop cluster => Start/stop the cluster process for CATALYST
  start|stop logdb  => Start/stop the CATALYST log database interface
  shutdown master  => Stops the CATALYST master process (and all
                  other CATALYST processes)
  status           => Lists the status of the CATALYST processes

command (? for help)> 

```

This program can be used to start and stop the individual CATALYST processes. For example, typing “start user” will start the user process and “stop user” will stop the user process. For convenience, a “start all” and “stop all” will start and stop all of the CATALYST components (except the master process itself).

Typing “shutdown master” will stop all of the CATALYST components and stop the CATALYST master process itself. The master process would then have to be restarted using the “start_catalyst.sh” script mentioned at the beginning of this section.

Typing “status” will present you with listing of the current status of the CATALYST components.

Typing “help” or “?” will bring up the list of available commands.

Typing “exit” or (CTRL+D) will exit the CATALYST management program.

1.4.5 How to Verify the CATALYST Server is Running

To verify that the CATALYST server is running, the Operator’s Console can be used. Use the CATALYST Operator’s Console to connect to the appropriate CATALYST Server installation. If you can successfully connect to the CATALYST server, that signifies that the server is running and can accept requests. Please refer to the CATALYST Operator’s Console Operator’s Manual for more information regarding host and port selections.

1.4.6 Editing the CATALYST Server’s Access Control List

To restrict the abilities of various users, CATALYST utilizes an Access Control List (ACL). This ACL describes what privileges a particular user has in CATALYST: admin, manager, operator, or a combination of the three. [Table 1-5](#) describes the rights granted by the three security categories:

Table 1-5. CATALYST Privilege Categories

Name	Abilities
admin	Gives the user the ability to start/stop/initialize server components.
manager	Gives the user the ability to delete PRs, mark PRs complete, and enable/disable subsystems and PGEs.
operator	Gives the user the ability to submit PRs and pause, resume, delete, force start, and resubmit jobs and/or epilogs.

The ACL file is located in “**\$CERESHOME/catalyst/conf/catalyst_users.conf**”. The file is a plain text file with one user (represented by their *AMI* username) per line. Following the user's name is a list of the privilege categories to which they belong. For example:

```
thillye1 operator
```

grants thillye1 only operator privileges, while:

```
jcwilki1 operator manager
```

grants jcwilki1 both manager and operator privileges.

In order for changes in this file to come into effect, this requires either “CATALYST Server->Reload ACL” from the Operator's Console or restarting the CATALYST user process.

NOTE: Users not listed in the CATALYST ACL are given read-only privileges to the CATALYST server; they will not be able to modify the state of the server or jobs managed by the server. Only users with AMI-p login privileges will be able to interact with the CATALYST server.

1.4.7 Stopping the CATALYST Server

In the event the CATALYST server must be shutdown completely, a CATALYST administrator can initiate the stop_catalyst script on the command line:

At the command-line (>) type:

```
> cd $CERESHOME/catalyst
> bin/stop_catalyst.sh
```

The server and all of its components should stop momentarily.

1.5 Expected PR Instantiation Time

The amount of time required for CATALYST to generate jobs from a PR is dependent on the following factors:

- The CPU and I/O load of the host and/or disk storage system (this is especially true in an environment where resources are shared amongst CATALYST and non-CATALYST operations)
- The size of the PR that is to be loaded.
- The amount of other PRs which have been or are being loaded.

Based on the above factors, the exact amount of time for loading a PR into CATALYST can vary as its runtime is a non-deterministic process. To alleviate some of this, CATALYST does perform some task switching between PRs as it create jobs and perform predecessor background checks so as to maximize execution progress across the scope of active PRs. Note that it is **not** required to load a PR in its entirety before being able to process the workunits comprising that PR.

1.6 Execution Evaluation

1.6.1 Exit Codes

The CATALYST server can exit with the codes listed in [Table 1-6](#). Exit codes marked as failure should be reported to the responsible persons listed in [Table 1-1](#).

Table 1-6. Exit Codes for CATALYST Server

Exit Code	Definition	Action
0	Normal Exit	Server started normally
Other	Failure	Notify persons listed in Table 1-1

Exit codes from PGEs run inside of CATALYST should be compared against the CERES *Standard Exit Codes* document found on the CERES documentation website.

1.7 Operational Log Files/Directories.

CATALYST updates several log files during operation.

Table 1-7. Log File Listing for CATALYST Server

Directory	File Name	Purpose
\$CERESHOME/catalyst/logs	ange.log	Contains messages and errors encountered by ANGe email processor process.
\$CERESHOME/catalyst/logs	ange_pending.log	Contains a list of files that ANGe reports as archived, but have not been placed in the AMI DPO or who's DPO file size does not equal that of the ANGe email's listed file size.
\$CERESHOME/catalyst/logs	cluster.log	Contains information regarding SGE-relating information and errors. Messages regarding users enabling/disabling blades will appear in this file. All messages in this file are emitted by the CATALYST cluster process.
\$CERESHOME/catalyst/logs	kernel.log	Contains PR and job management activity managed by the CATALYST kernel process. Manual actions performed by the operator on PRs and Jobs will appear here.
\$CERESHOME/catalyst/logs	logdb.log	Contains log database connection information and/or error messages emitted by the CATALYST logdb process.
\$CERESHOME/catalyst/logs	master.log	Contains process management activity managed by the CATALYST master process.
\$CERESHOME/catalyst/logs	user.log	Contains messages regarding user logins/logouts which is managed by the CATALYST user process.
\$CERESHOME/catalyst/logs	xmlrpc.log	Contains messages and errors emitted by the CATALYST xmlrpc (front-end) process.
\$CERESHOME/catalyst/logs/epilogue/\$PGE/\$YYYY/...	\$CATALYST_JOB_NAME.epilogue	Contains the STDOUT/STDERR from the ASDC epilogue process that is called by CATALYST.

1.8 Operational Files/Databases.

The table below outlines files that CATALYST produces. These files should **never** be removed and are part of the normal operation of CATALYST.

Table 1-8. Operational File Listing for CATALYST Server

Location	Name
\$CERESHOME/catalyst/data	all files with the extension .sqlite3
\$CERESHOME/catalyst/data	all files with the extension .sqlite3-shm
\$CERESHOME/catalyst/data	all files with the extension .sqlite3-wal
\$CERESHOME/catalyst/data	catalyst.pid (contains CATALYST master process's PID)
\$CERESHOME/catalyst/sockets	all files with the extension .sock
dsrvr205.cluster.net	This server hosts the CATALYST Logging Database. Later builds of CATALYST will provide a database inspection API. ceres_logdb_ppe (<i>only when running in PPE</i>) ceres_logdb_prod (<i>only when running in Production</i>)

The CATALYST databases on “dsrvr205.cluster.net” are initially populated with static data included in the CATALYST installer. They contain records of the existence and approved non-existence of predecessor products used by the PGE Handlers.

1.9 CATALYST Monitoring Utility

The CATALYST monitoring utility is used to report on the status of the CATALYST server and its various components at regular intervals. The monitor checks and reports on:

1. is the CATALYST master process running?
2. is the free disk space available to \$CERESHOME equal to or greater than the threshold value (see “minimum-free-space” in the server configuration file)?
3. is the memory utilization for the CATALYST components equal to or greater than the threshold value (see “memory-warn-threshold” in the server configuration file)?
4. are the CATALYST processes specified in the configuration file (see “minimum-processes”) running?
5. have any of the CATALYST server logs recorded an error within the past hour?

The monitor will send status emails to the users designated in the configuration file (see “status-email-recipients”).

1.9.1 Installation of the Monitoring Utility

The monitor needs to be installed as a cron task on the CATALYST host (catalyst.larc.nasa.gov).

On the head node (*ab01.larc.nasa.gov*):

You will want to deactivate the CATALYST monitor from CATALYST v1.0.x before moving forward.

```
$> crontab -e
```

This will start your default editor. Delete the line(s) referencing the CATALYST v1.0.x monitoring utility. Save and close the editor.

On the CATALYST host (*catalyst.larc.nasa.gov*):

After the v1.0.x version of the monitor has been removed, the first step will be to open your crontab:

```
$> crontab -e
```

This will start your default editor. In it, you may see other cron jobs that you have already scheduled. Starting on an empty line, you will want to add the two lines from one of the groups below:

- CERES CM:

```
0 6,14 * * * /SSIT/CERES/CATALYST_Testing/catalyst/bin/send_email.sh /SSIT/CERES >& /dev/null
0 0-5,7-13,15-23 * * * /SSIT/CERES/CATALYST_Testing/catalyst/bin/send_email_if_issue.sh /SSIT/CERES
>& /dev/null
```

- ASDC PPE:

```
0 6,14 * * * /PPE/CERES/catalyst/bin/send_email.sh /PPE/CERES >& /dev/null
0 0-5,7-13,15-23 * * * /PPE/CERES/catalyst/bin/send_email_if_issue.sh /PPE/CERES >& /dev/null
```

- ASDC Production:

```
0 6,14 * * * /CERES/catalyst/bin/send_email.sh /CERES >& /dev/null
0 0-5,7-13,15-23 * * * /CERES/catalyst/bin/send_email_if_issue.sh /CERES >& /dev/null
```

The first of these two lines will always send a CATALYST monitor report email at 06:00 and 14:00. The second command will check for issues on the hour every hour other than 06:00 and 14:00 and only send an email if there is an issue. Saving and closing the editor will install the new crontab. You should see a message on your command prompt similar to the following: “crontab: installing new crontab”.

These times can be easily changed for the email notifications. Please let the CATALYST development team know if you need help in adjusting the cron job's times.

1.10 Best Practices

There are several best practices for efficiently using the CATALYST software for CERES data production:

Submit PR chunks as soon as possible to ensure CATALYST can keep busy during staff downtime. Not keeping CATALYST fed with enough workunits will cause starvation which has negative impacts on processing speed. This also allows CATALYST the opportunity to interleave a PR's PGE execution and background checks for the PR itself. The background checks specifically for CATALYST operation can take several minutes to complete, so performing this step as early as possible lessens the impact of this.

- **Avoid “block running” PR chunks.** For example, submitting and running an entire month of one PR before submitting a month of a successor PR. This may be convenient for executing PRs by hand, but it is a very inefficient strategy in the long-run because it results in large idle periods in both the predecessor and successor PRs. CATALYST knows how to run jobs in a PR in the CERES subsystem approved order.

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

Product Availability in CATALYST Database

B.1 CATALYST 1.0.4 Installation Package + Logging Database Patch (SCCR 996)

Based on the “as-delivered” static database contents listed below, CATALYST can successfully process the Edition4 Clouds/Inversion up to 2013-06-30. Attempting to process datadates beyond this will stall waiting for input; they will not run automatically producing non-optimal data. Static database patch deliveries will extend this date.

Product	Sampling Strategy	Production Strategy	Configuration Code	Starting Date	Ending Date
ESNOW/EICE	CERES	NSIDC-AFWA-Mesh16th	200200	20121230	20140223
ESNOW/EICE	CERES	NSIDC-CLASS	200200	20121211	20140223
ESNOW/EICE	CERES	NSIDC-NESDIS	020019	20000201	20021231
ESNOW/EICE	CERES	NSIDC-NESDIS	020018	20000301	20020930
ESNOW/EICE	CERES	NSIDC-NESDIS	021021	20030101	20050202
ESNOW/EICE	CERES	NSIDC-NESDIS	021022	20050201	20070131
ESNOW/EICE	CERES	NSIDC-NESDIS	021023	20070129	20090910
ESNOW/EICE	CERES	NSIDC-NESDIS	021024	20090831	20101031
ESNOW/EICE	CERES	NSIDC-NESDIS	022025	20101101	20130307
ESNOW/EICE	CERES	NSIDC-NESDIS-Mesh16th	022025	20081007	20130310
IES	Aqua-FM3	Edition3	032040	2002061800	2011010123
IES	Aqua-FM3	Edition3	300300	2010123100	2012100123
IES	Aqua-FM3	Edition3	300303	2012093000	2013063023
IES	Aqua-FM3	Edition3	300301	2012100100	2013022823
IES	Aqua-FM3	Edition4	400403	2002061800	2009123123
IES	Aqua-FM4	Ed3-NoSW	032040	2005033000	2011010123
IES	Aqua-FM4	Ed3-NoSW	300300	2010123100	2012100123
IES	Aqua-FM4	Ed3-NoSW	300303	2012093000	2013063023
IES	Aqua-FM4	Ed3-NoSW	300301	2012100100	2013022823
IES	Aqua-FM4	Ed4-NoSW	400403	2005033000	2009123123
IES	Aqua-FM4	Edition3	032040	2002061800	2005032923
IES	Aqua-FM4	Edition4	400403	2002061800	2005032923
IES	Terra-FM1	Edition3	032040	2000022900	2011010123
IES	Terra-FM1	Edition3	300300	2010123100	2012100123
IES	Terra-FM1	Edition3	300303	2012093000	2013063023
IES	Terra-FM1	Edition3	300301	2012100100	2013022823
IES	Terra-FM1	Edition4	400403	2000022900	2009123123
IES	Terra-FM2	Edition3	032040	2000022900	2011010123
IES	Terra-FM2	Edition3	300300	2010123100	2012100123
IES	Terra-FM2	Edition3	300303	2012093000	2013063023
IES	Terra-FM2	Edition3	300301	2012100100	2013022823
IES	Terra-FM2	Edition4	400403	2000022900	2009123123
MOA	CERES	DAO-G5-CERES	020032	2007100100	2010110100
MOA	CERES	DAO-G5-CERES	020033	2010110100	2012070118

Product	Sampling Strategy	Production Strategy	Configuration Code	Starting Date	Ending Date
MOA	CERES	DAO-G5-CERES	300300	2012063000	2014010118
MOA	CERES	DAO-G5re-CERES	019031	2007100100	2007103118
MOA	CERES	DAO-GEOS4	016024	1997121500	2004033018
MOA	CERES	DAO-GEOS4	016025	1998120100	2004120118
MOA	CERES	DAO-GEOS4	016023	2000022400	2004021518
MOA	CERES	DAO-GEOS4	016026	2004113000	2005010118
MOA	CERES	DAO-GEOS4	017027	2004123100	2005122918
MOA	CERES	DAO-GEOS4	017028	2005113000	2006020106
MOA	CERES	DAO-GEOS4	018029	2006020100	2007030118
MOA	CERES	DAO-GEOS4	018030	2007020100	2008010118
MOA	CERES	ECMWF-GEOS2	013016	1998010100	1998090118
MOA	CERES	ECMWF-GEOS3	015019	2000010100	2002080118
MOA	CERES	ECMWF-GEOS3	013016	2000022500	2000103118
MOA	CERES	ECMWF-GEOS3	014018	2000110100	2000111918
MOA	CERES	ECMWF-GEOS3	016020	2002080100	2002100112
MOA	CERES	ECMWF-GEOS4	016020	2002100100	2003043018
MOA	CERES	ECMWF-GEOS4	016021	2003050100	2003073118
MOA	CERES	ECMWF-GEOS4	016022	2003080100	2003093018
MOA	CERES	GMAO-G5-Beta2Ed3	300300	2004010100	2006013118
MOA	CERES	GMAO-G5-Beta2Ed3	300301	2006070100	2006073118
MOA	CERES	GMAO-G541-Ed4	400400	2000022800	2013070118
MOA	ValR2OS	DAO-GEOS4	018030	2006093000	2006110118
MOD02SS1	MODIS_TERRA	C5		200002240000	201312312355
MOD03	MODIS_TERRA	C5		200002240000	201312312355
MOD04_L2	MODIS_TERRA	C5		200002240010	201312312355
MYD02SS1	MODIS_AQUA	C5		200207032200	201312312355
MYD03	MODIS_AQUA	C5		200207030000	201312312355
MYD04_L2	MODIS_AQUA	C5		200207040045	201312312355
SCCD/SCCN	Aqua-FM3	Edition3	300301	200206	201304
SCCD/SCCN	Aqua-FM3	Edition3	300300	200207	201112
SCCD/SCCN	Aqua-FM3	Edition3	300303	201305	201306
SCCD/SCCN	Aqua-FM4	Ed3-NoSW	300300	200504	201112
SCCD/SCCN	Aqua-FM4	Ed3-NoSW	300301	201201	201304
SCCD/SCCN	Aqua-FM4	Ed3-NoSW	300303	201305	201306
SCCD/SCCN	Aqua-FM4	Edition3	300301	200206	200206
SCCD/SCCN	Aqua-FM4	Edition3	300300	200207	200606
SCCD/SCCN	Terra-FM1	Edition3	300301	200002	201304
SCCD/SCCN	Terra-FM1	Edition3	300300	200003	201112
SCCD/SCCN	Terra-FM1	Edition3	300303	201305	201306
SCCD/SCCN	Terra-FM1	Edition4	400403	200002	200912
SCCD/SCCN	Terra-FM2	Edition3	300301	200002	201304
SCCD/SCCN	Terra-FM2	Edition3	300300	200003	201112
SCCD/SCCN	Terra-FM2	Edition3	300303	201305	201306
SCCD/SCCN	Terra-FM2	Edition4	400403	200002	200912