

Minutes of the 3rd SDO Pipeline meeting, Brussels 17th September 2010

JSOC : Rick Bogart

JSOC Status

There is 3 NetDRMS servers at Stanford.

hmidb is for data processing

Slony replicates all records from hmidb to hmidb2 (for data distribution)

hmidb, hmidb2 share the same SUMS, access to same data

SUMS is unstable because it is still in development, that is the reason of the 3rd server

Stanford has 23 partitions of 30TB each, which amounts to about 700TB. It is already almost full.

The test data will not be archived, only raw telemetry data is archived.

Under consideration:

Have a 4th server; NetDRMS client node with large disk farm. But I/O requirement for the system is so high that it is not cheap to add new disk.

But they may set up a remote site with less requirement on I/O (no more tape), and hence more disk space where data can be on line, just as in European data centers.

The tape archive will always be accessed as a last resort.

Published JSOC Data Series

archive flag = 1, meaning archive to tape : e.g. MDI are archived

archive flag= -1: when SUMS data are no more online, they will be destroyed from JSOC. So this is a problem if this has been mirrored by other sites.

Rick showed meta-data series mirroring by sites.

Forward processing, back tracking as getting from test data to final data

Test data will be replaced by definitive data.

New series:

Aia.dark

Aia.flat

Aia.cal

Aia_lev1

Aia.lev1p5

Aia.synoptic

HMI low level: =observable hmi.lev1
Hmi.lc_45c
HMI high level product (what was promised)
Hmi.fsVbinned
Hmi.rdVflows_fd15_frame: results of inversion

Other: gong.mrvzi, gong.site_vzi, mdi data

NetDRMS status:

NetDRMS 2.4 beta released, corresponds to JSOC 5.10
2.4beta has a bug affecting archiving to tape.

It is expected that 2.4 will include JSOC Mirroring Daemon.

The version 3.0 will corresponds to JSOC 6.0 - pipeline production release. The main difference between 2.4 and 3.0 is that the infrastructure will be stabilized.

Planned changes to NetDRMS code 2.4:

Fix slony log parser to handle bad insert statements.
Fix hard coding issue in sum_rm
Fix remaining sum_svc memory affecting NSDO but not JSOC, but still needs to be identified.

Inclusion of DELETE statements in slony log parser

Future release:

Non-slony mechanism, for remote site to provide DRMS data for local ingest

Extension of DRMS data naming scheme to incorporate site name, jsoc:hmi.M_45s : sums

Multiple segments series e.g. AIA combining different wavelengths
Browse segments, JPEG

At this moment: process forward 5 days per days. The test data go away, so data centre must delete the test-data and recreate DRMS, so the recnum will not come by chronological order.

Elie : Problem for people who want to know which data are in the SUMS, which data are they subscribing to.

The tape archive is ok for data that is accessible directly. But once tape archive is full, it needs to go on the shelves, and then forget about it. So data that are often accessed should be on a disk cache.

Coordination with big other archive, have retention based on what the other archive has.

HMI data will be archived

Automated reprocessing : when? No date set up but this has to be within 6 months (SAO has 9 months cache)

When AIA data has expired, needs to be reprocessed from level 0 data whenever required.

Igor JMD

JMD requirements

JMD should be as stand alone as possible, use DMRS/SUMS interface but not DRMS system resources. So if the JMD behaves badly, it does not affect DMRS operations.

Make it as resilient as possible, should be removable, and should be self-contained.

Software used:

Jetty webserver 6.1

Derby embedded db 10.5.3.

Java 1.6

Development within Eclipse IDE --- very handy according to Igor. i.e. when you type a class name, you have a list of all methods available for you.

The wiki is at <http://vso.tuc.noao.edu/VSO/w>

Authoritative node: AN; Peer node: PN; your node: YN

1. Slony log replication from AN to YN

2. Jsoc_fetch query from YN to PN

(peer node = SAO for ROB, ROB for UCLan, ...)

3. Once you know where the data is located, you do data scp-ing to the attributed peer node

Slony replication;

- triggers in replicated table series, write sunum, recnum, series-names to the sunum_queue table

JMD structure

- JMD reads from the sunum_queue table the sunums to process and delete them as soon as they are written to the embedded db.

Current issues

Is "invalid" a final state ?

Distinguishing between a non existing sunum and a time out!

JMD stops writing to log file

Jetty embedded data base issues: size, sometimes network connection drops temporarily

Discussion:

Slony logs, different replication, what happens when one of the nodes fails, and other nodes are taking data from other sites. What happens when the failed node comes back. ?

SDO data centers:

UCLan:

When trying to stop sum_svc, it reports that there is an active users.

To answer share ideas we could set up a forum where people who have problem can post their questions and other their solutions.

What type of architecture is useful for everybody?

Resilience about data centres : scheme of who gets data from whom, even in case of problem.

IAS:

Postgres trigger removed, but then needs to do certain things manually like request download of selected data every hour in the crontab.

How to reprocess backward when new data series will arrive? What strategy to adopt?

Tools developed for monitoring which data were download: see as a function of wavelength which data was downloaded.

IAS: MPS interested in getting AIA data from IAS - still TBC

Discussion on data transfer

See document "SDO data transfer and storage Continued".

Rick:

NetDRMS used to process data coming from the spacecraft, so requirements are very stringent

Retention time = 60 days, build archive then after 60 days

ROB is using sums partitions but it may not be maintained. Sum_svc not really trusty since so many development is still needed

sum_svc is critical for pipeline at JSOC, so Stanford is a bit touchy about it.

Strong trends to stabilise the system since it is critical for their pipeline.

The main developers are Art (DRMS) and Jim (SUMS)

There is not much documentation, we should write by asking question to Art and Jim.

Pablo Alingery:

Problem encountered: 1 sum_num that is missing, and retention time is set to a particular value.

David Boyes :

There is several level of open source:

RedHat is not open, but there is three levels: (one stable, one that can change)

NetDRMS: is more fluid, more prone to change, but DRMS does not go through main changes

Ideas for future development

Start the design from blank canvas : use a file system: tell where your data is, using file names

Rick:

SUMS: virtual data to hinder between tape and disc storage

DRMS protocol of data and sums, DRMS support different protocols ; there is one generic protocols. It does not know anything to the SUMs, no API provided

Create data series with generic protocols - but then cannot use DRMS as API for the fits file

Complication in the NetDRMS : we have only user access, we do not change anything, read sums, read Postgres, and build a fits file

aia-prep: takes all information from DRMS, and hence can update

Distribution of data to end-users:

Flatfield : level1 is with flatfield correction, but flatfield degrades quite quickly, in 2 years from now, the new data will be photospherically completely wrong --- will need to wait for the updated calibration

VSO :

VSO programmatic interface: query VSO without going to web interface

This way you can;

Embed VSO calls in existing code

Perform complex queries (you can do less complex query on the gui)

- engage in querying providers that offer specific extended interfaces e.g. XRT

Vso_search

Vso_get

IDL > doc_menu, 'vso_search'

Gives you lots of info

Vso_search(data='2010-05-01',provider='jsoc')

Extent= 'FULLDISK' , inst='aia' , phos

Wave= '304 Angstrom

Vso_search('2010/4/30'

Specific query

Interaction : VSO be a service for HEK and Helioviewer

VSO queries HEK for features and events : HEK treats as any of the other providers

HEK can use VSO to get the data

VSO uses HEK as AIA cutout data provider - results that can be provided locally.

What still needs to be done in VSO?

Finish new web interface, Keith : shared with Helioviewer

Handle movies /JPG images
Handle more complex queries

Fully distribute searches / Geo location
VSO installation at other sites
Solve tar on the fly problem : email for later downloads

Implement thumbnails JPG/JPEG2000
Parallel downloads

Helioviewer - Keith Hugitt

Web version: no feature tracking
Jhelioviewer: has a feature tracking algorithm
Back-end development
Modular architecture; common interface, validation, and documentation
Easy to extend
Can choose which modules to support
Current modules: web-client JHelioviewer

Software used: Kakadu, GD for color tables, ImageMagick, FFmpeg - movie generation

Database uses MySQL
Helioviewer interaction (better movie time estimation and status information)

Interaction with SSW cut-out service support, VSO interaction

Server requirements; Ubuntu 10.04 is ok. Idem for Centos. Windows is too difficult - forget about it.

Optimizing server Hardware for Helioviewer
XFS: can support much larger volume than other FS

JPIP server; works fairly well. Look into alternative open source options for generating JPEG2000 movies

Browse Clients
Helioviewer
New functionality:
AIA support
VSO querying
HEK interaction
Normalize movie generation, using H.264, tar.gz

Bugs report, feature request, version control : launchpad

Mailing list ,
launchpad.net/helioviewer
Helioviewer-dev

Jack Ireland's presentation:

Fits keywords, use of metadata that also contains all info about JPEG2000

Want to include all data set, needs some compliance requirements: filename specification, storage directory structure, FITS keywords , XML key-value pair minimal set, information is available online, but will be made much clearer shortly.

JP2Gen: name of the code that converts AIA to JP2K maybe of use for other data set than SDO

AIA:
IDL > HV_AIA_
Distributing JPEG2000 files?
From AIA:
Get JP2 from Stanford or from Lockheed

JP2Gen for data set other than SDO, e.g. SWAP alongside other SDO

www.helioviewer.org/wiki
- contains information on JP2Gen, but is about 6 months out of date

<https://launchpad.net>
What do each data center do wrt Helioviewer? Are we participating in their bug-reporting, etc,..?

How to run locally Helioviewer ? How to have all the functionality of Helioviewer be present at local sites (link with HEK, with VSO)? How can we help them?

Neal Hurburt:

- Level1.5_nrt
<http://jsoc.stanford.edu/data/aia/images> : 2 hours latency

Export request - when data will be older

LMSAL data mining hardware:
300TB user archive
22 nodes SGI XE Cluster for EDS

LMSAL cache
Stream aia.level1p5_nrt
Cut out service moves subset to 'archival' area

Support from JSOC:
read_sdo.pro : point to JSOC, some of these function should point to other sites

SSW AIA support: aia_get_response.pro

Result : aia_index2center.pro

Helio Coverage Registry (HCR) captures also all cutout services, all campaign plan

<http://lmsal.com/sungate> : for details

or

<http://sdowww.lmsal.com> : for results

Cutouts: about 30% are reused

Since May, > 900 cutout requests, > 2TB exported

Hosting cutout service

Requirements; lookdata web service, SSW web service framework, Pseudo FS : SSW is looking for something of that type

STEPS : install ssw_cutout, configure, set observer = site

A bit tricky to set up the web service framework

If IDL is in server mode, it is fairly easy to implement cut-out service