



HADOOP – Usage in Medspiration

F Paul, JF Piollé - IFREMER

Nephelae big data platform

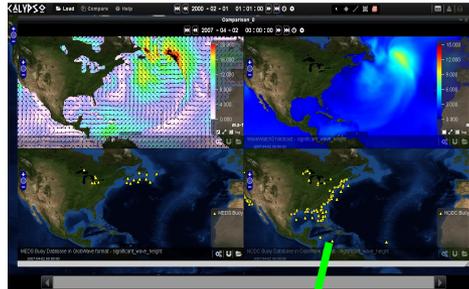


ERS
ENVISAT
LandSat
MSG
GOES series

series NOAA
series METOP

AQUA
SMOS
ADEOS
JASON
TOPEX
CryoSat
GFO
AltiKa

TRMM
QuikSCAT
OceanSat
HY2



Analysis, comparison and synergy tools



Stereo video camera

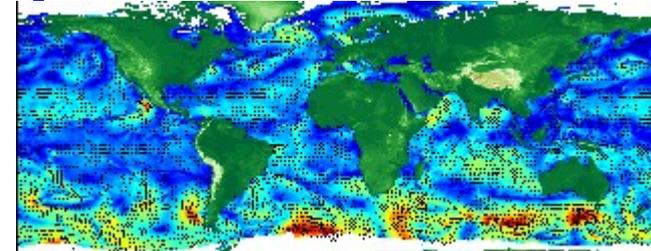


Buoys, floats



Nephelae

Weather, ocean, wave models



Liste des Jobs Gogoliat:

ID	Workspace	Jobname	User	Start	Stop	Duration	State	Tasks	Infos
31	workspace_20120208_000001	Test	Test	2012-02-08 12:32:04	2012-02-08 13:38:17	00:06:13	OK	1 (Completed: 1)	OK (137) (Completed: 0)
32	workspace_20120208_000002	Test	Test	2012-02-08 13:32:10			OK	1 (Completed: 1)	OK (137) (Completed: 0)
33	workspace_20120208_000003	Test	Test	2012-02-08 13:32:16			OK	1 (Completed: 1)	OK (137) (Completed: 0)
34	workspace_20120208_000004	Test	Test	2012-02-08 13:38:16			OK	1 (Completed: 1)	OK (137) (Completed: 0)
35	workspace_20120208_000005	Test	Test	2012-02-08 13:38:40	2012-02-08 13:39:40	00:01:00	OK	1 (Completed: 1)	OK (137) (Completed: 0)
36	workspace_20120208_000006	Test	Test	2012-02-08 13:39:45	2012-02-08 13:39:46	00:00:01	OK	1 (Completed: 1)	OK (137) (Completed: 0)
37	workspace_20120208_000007	Test	Test	2012-02-08 13:37:38	2012-02-08 13:37:38	00:00:00	OK	1 (Completed: 1)	OK (137) (Completed: 0)
38	workspace_20120208_000008	Test	Test	2012-02-08 13:47:43	2012-02-08 14:11:24	00:23:41	OK	1 (Completed: 1)	OK (137) (Completed: 0)
39	workspace_20120208_000009	Test	Test	2012-02-08 13:53:59	2012-02-08 14:13:00	00:19:01	OK	1 (Completed: 1)	OK (137) (Completed: 0)
40	workspace_20120208_000010	Test	Test	2012-02-08 13:53:59	2012-02-08 14:03:07	00:09:08	OK	1 (Completed: 1)	OK (137) (Completed: 0)
41	workspace_20120208_000011	Test	Test	2012-02-08 13:53:59	2012-02-08 14:03:07	00:09:08	OK	1 (Completed: 1)	OK (137) (Completed: 0)

Processing tools

- 1.5 PB
- 600 processing cores
- 2.5 TB memory

Nephelae in support to GHRSSST



Need at Ifremer/CERSAT for a (at least partial) copy of GHRSSST archives :

- ✓ Global (re)analysis
- ✓ HR-DDS, MDB & MMDB generation
- ✓ User application requiring intensive data access and processing (OceanFlux project, uses 3 archives of SST)

GHRSSST Archive :

- ✓ complete GHRSSST L2P datasets, converted to NetCDF4
- ✓ European L3 and L4 products
- ✓ historical OSI SAF data (Goes, Seviri) converted to GDS
- ✓ new OSI SAF datasets (incl. L2P METOP, L3P NAR NOAA-19 & METOP, L3P GLOB METOP in GDS v2)
- ✓ To be extended for the purpose of felyx demonstration
- ✓ **Total : 22 TB** => makes transfer to user for large scale applications impossible
- ✓ Mid-term archive for MyOcean (latest months of compressed NetCDF-3 L2P) remains accessible too.

OpenDAP : <http://www.ifremer.fr/opendap/cerdap1/ghrsst/>
FTP : <ftp://eftp.ifremer.fr/ghrsst>

(login : g1e9b2 , password : K0mpkKVW or request to fpaf@ifremer.fr)

Product catalogue at : <http://www.ifremer.fr/vcerdmz1/joomla/data/collections/ghrsst>

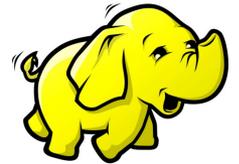
Access to virtual machines and distributed processing capabilities is freely possible on agreement with CERSAT (contact jfpiolle@ifremer.fr). Approach has been tested successfully in multi partner projects over the last two years. Big data and cloud computing are key element for the future sustainability and usage of multi-mission archives.

Hadoop ?



WHAT

open-source software framework
designed to supports data-intensive distributed applications
using large clusters of commodity hardware



MAIN FEATURES

MapReduce Engine
Distributed File System
Smart use of hardware and architecture characteristics for I/O : minimize
network use, disk streaming vs seeking

<http://hadoop.apache.org>

Hadoop & GHRSSST



- Use of Satellite Data Archives (>50TB L2)
- Years of data for a dozen of satellite sensors, several data formats
- Analysis of long time series, intercomparison, multi-variate analysis

Where does HADOOP help ?

- MapReduce : easy to merge datasets and extract metrics (or files)
- Scalable BigData distributed processings : high performances
- Fully Fault-tolerant (hardware, software) : easy to manage
- Native distributed storage ? No... not really suitable for NetCDF

Medspiration - Data metrics



Use case : compute SST and SST Anomaly metrics from high resolution SST datasets over several years and SST climatology, multivariate analysis with other time series

Input Datasets

- SST global maps (4 years, files per day)
- SST Climatology

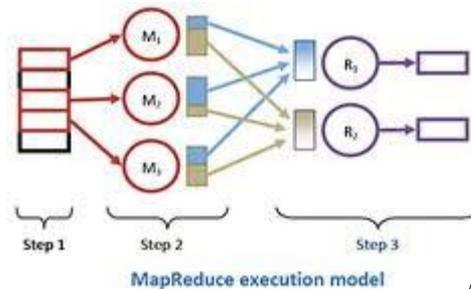
Input Parameters

- time period (years ?)
- geographic zone (all?)
- data threshold

Output

For SST&SST Anomaly,
Per day/month/year :

- timeseries
- histogram
- maps

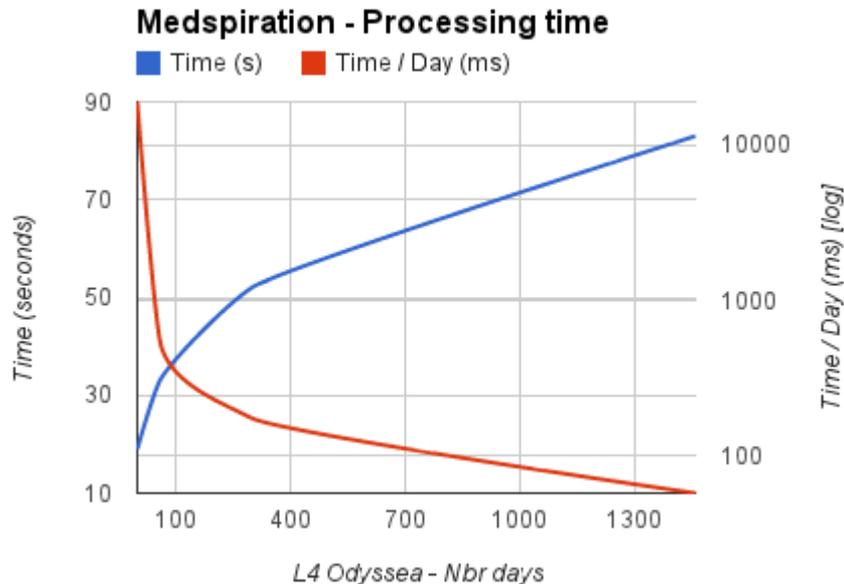


~ 1 minute

Performances



Bench platform : 15 nodes (12 cores @2.4GHz, Dell R510)



What does this means ?

- **4 years (1460 files, 50GB archive) of L4 Odyssea SST data** processed in less than **90 seconds** : 1 single run to compute several statistics timeseries and histograms
- each file processing time \ll 1s
- but ... 20s minimum for whole processing (even limited to 1 file), due to scheduler overhead. Interactivity issue ?
- Medspiration case.... relevant only for long series ?

Other relevant use case examples :

- Compute map for **WaveWatch3 HS** percentiles over **15 years** (1.7TB, 1 point every 3h @0.5° see on the right) : **10 minutes**
- **10 years** (400GB, >50k files) of **Quikscat L2B** to retrieve daily global wind_speed min/max/mean : **2 minutes**
- **6 years** (80GB, > 28k files) of **Ascat L2** to retrieve monthly global wind_speed mean/percentiles : **5 minutes**