sea state cci

System Specification Document (SSD)

version 1.0, 14 November 2019

Contents

Author	Approved	Signature	Date
JF Piolle	Fabrice Ardhuin Ellis Ash		14/11/2019
ESA Acceptance	esa Digitally signed by Craig Donlon Date: 2020.02.19 15:59:58 +01'00'		

Issue	Date	Comments
1.0	14/11/2019	First version submitted for ESA approval

List of Acronyms

CCI	Climate Change Initiative
CMEMS	Copernicus Marine Environment Monitoring Service
CPU	Central Processing Unit
ECV	Essential Climate Variable
ESA	European Space Agency
GDR	Geophysical Data Record
HPC	High Performance Computing
IPF	Input Processor Function?
L1A	Level 1A
L1B	Level 1B
L4	Level 4
LRM	Low Rate Measurement
RA	Radar Altimeters
RR	Round Robin
S3A	Sentinel-3A
S3B	Sentinel-3B
SAR	Synthetic Aperture Radar
SSH	Secure Shell / Sea Surface Height
SWH	Significant Wave Height
WV	Wave (mode for SAR)

1. Introduction

This document presents the System Specification Document (SSD) for **Sea_State_cci**, deliverable 3.2 of the project. This first version is prepared in the first year of the project in order to specify the initial processing system. An updated version will be produced towards the end of each year of the project.

The remainder of this System Specification Document contains sections describing:

- Trade-off analysis
- Engineering methodology
- Cost effectiveness
- Other aspects
- System architecture

2. Trade-off analysis

At the time of system design, little is known on the final choice of IPF (altimeter retracker or SAR Wave Mode processor) that will come out of each round-robin activities. A wide range of different cases is expected, depending on the proposed IPF, and the processing system will therefore be strongly dependant on a trade-off analysis of the selected IPF. The following diagram sketches the decision tree anticipated for this trade-off analysis. We detail in the next section the possible strategies that may be put in place with respect to this analysis for the reprocessing of each mission considered in the CCI Sea State project.

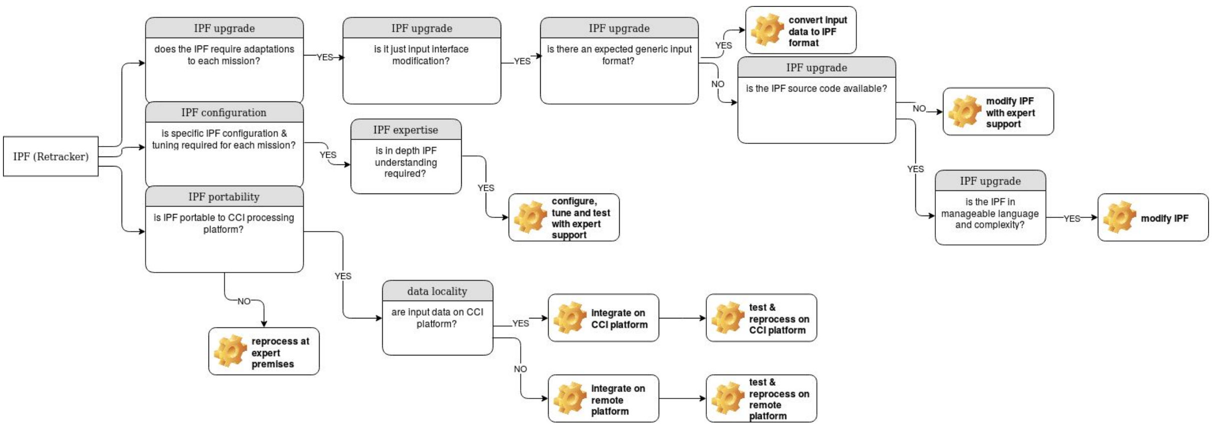

3. Engineering methodology

The CCI Sea State approach on the reprocessing of data climate series is based, in the first phase, on a round-robin exercise that will support the selection of the best algorithm (implemented through a processor or suite of processors) to produce a consistent and stable multi-mission time series of the chosen sea state parameters.

Based on the above analysis and the selected processor(s), different steps and implementation paths are anticipated, with respect to two main aspects:

Processing target

It was originally expected to have all reprocessing activities running on the same CCI platform (at Ifremer) but the preliminary trade-off analysis leads to possibly different scenarii, in particular using other remote processing platforms:

- some input datasets may be too large to be transferred to the CCI processing platform: in particular the L1A/L1B data from Sentinel-3 proved to be very complex, time consuming and anyway too voluminous to restore from their original archive (at ESA our Eumetsat) and transfer electronically to Ifremer. This is a case where other processing platforms are to be sought and used, close to the original data.
- the processing time of a single orbit is highly variable depending on the selected processor for the reprocessing: from the requirement analysis, ratio of 1 to 30 are expected for altimeter LRM retrackers, possibly x50 for SAR mode. Overall CPU time estimation ranges from hundreds of thousands to millions of hours… Completing the

reprocessing of all expected missions within the time frame of the CCI project may therefore require to split the processing on different platforms too
- there is no clear commitment for Round-Robin participants to deliver an executable processor (only the algorithm is to be provided). Re-implementation of a processor would be too time consuming and would require an important validation/intercomparison effort. There may therefore be legal or practical blocking issues to transfer and integrate a processor on CCI (or another) platform: this is another possible case where the processing may have to be moved to a platform where the processor is already available and integrated within an IPF. The reprocessing responsibility may then have to be transferred to another entity/partner within the project.

External support

The outcome of the Round-Robin step will be an algorithm, implemented within an executable processor. At this stage it is not foreseen to re-implement a dedicated processor from an algorithm, as it is estimated as too costly and time consuming. It is therefore expected that in the engineering methodology for the CCI reprocessing activity that support from the processor expert will have to be sought to address possibly different aspects:
- adaptation of processor to other format and content: each satellite mission comes with a different format and content for the instrumental level observation data (L1). This is especially true on the considered CCI time frame that spans over 30 years. It is therefore very likely that either of the following approaches will have to be performed:
  - reformatting a mission's L1 data to the format and content expected by the processor, which can be done without expert support as long as this information is accurately specified and documented
  - adapting the processor to a mission format and content: depending on the language used in the processor implementation, the availability of source code,... this may not be doable by anybody else than the author (person or group) having implemented the processor: this is a case where support from this author will be mandatory.
- configuration, adaptation of specific processing steps and fine tuning of the processing to each mission: all instruments have specifics that will will have to be addressed when applying the selected processor (sensing band, mispointing, sensor drift,...). This will require support from the experts behind the specification and design of the selected processor.

4. Cost effectiveness

As demonstrated in the above analysis, there are at the project start many uncertainties regarding the reprocessing activities, having no prior knowledge of the choice of processor(s) to generate the CCI climate data records.

In terms of processing and infrastructure cost, effectiveness will be achieved through the following drivers:

- relying on existing institutional or institution funded processing resources: internal supercomputers at involved agencies (Ifremer, possibly CNES) or new generation cloud computing platforms funded by Copernicus, ESA, etc… will be sought after in order to support the financial cost without any additional charge for the CCI project
- exploit the locality of the input data: process the data as close as possible to where they are stored whenever possible, instead of doing costly and time consuming data transfer. In particular remote platform such as the ESA Sarvator may be used for the reprocessing of at least some ESA mission (such as Sentinel-3). CNES supercomputer may also be an option for some missions, in addition to the Ifremer platform
- exploit the processing time availability: we demonstrated that the overall CPU time required to process all mission may be overwhelming for a single processing facility, depending on the choice of the processor; as for data locality, we may use different platforms to distribute the processing in order to take advantage of available institutional (free of charge) processing resources and being also more effective in the usage of the existing facilities resources
- leverage on existing parallel initiatives, such as other reprocessing activities with the same or similar retracker performances, to avoid unnecessary waste of processing resources and manpower

5. Other aspects

5.1 Security measures

The CCI Sea State project does not hold any sensitive data and no specific action is undertaken to address such aspect.

Each considered processing platform for the project is however operated in fully secured environments with access restrictions and controls for physical persons, identification and authentication for remote network connection.

5.2 Conformance to EU GDPR

The CCI Sea State project does not hold any personal data. As such, no action is needed to comply to EU GDPR.

6. System Architecture

Keeping long and massive mission archives alive by raising the level of data revisiting through multiple applications, demonstration products or services, or extensive data reprocessing are major requirements for projects such as the CCI Sea State in order to combine numerous sources of data, test new ideas and perform reprocessing over a significant amount of data.

The system architecture for the project will therefore rely on existing infrastructures providing these capacities.

Among the key technical drivers identified to provide an efficient and cost-effective access to historical massive multi-mission archives are :

- fast and online access to massive collections of data
- avoiding data duplication and transfer to users
- minimizing time from algorithm development to processing
- allowing fast and easy to manage large scale reprocessing
- improving data storage and management

There may be several target infrastructures fitting these criteria and the CCI Sea State project may use one or several of them, depending on their respective strengths and limitations for a particular mission retracking. Possible infrastructures to be used include:

- Ifremer “Datarmor” platform, originally intended for most of the processing
- ESA “SARvatore”, which may prove more adapted for ESA missions (thanks to the availability of the input data) such as Sentinel-3A & B, CryoSat-2 or Envisat
- ESA CNES, which hosts all Jason mission data and has already integrated retrackers

Contacts have been engaged with all these possible resource providers and will have to be discussed more deeply once the project has settled the processing scope (missions and retrackers to be used).

Ifremer “Datarmor” platform

Physically, the platform is based on the Datarmor platform operated by Ifremer IT department (refer to the facility section in the management proposal). The current storage capability is about 12 PB (20 PB by 2020) and the available capacity largely exceeds the need for CCI products.

The reprocessing framework for CCI will therefore make use of direct access to the complete CCI input data archive on disk, physically located within the cluster, and distribution of the processing over multiple nodes of the cluster. The task of managing and distributing the processing jobs is alleviated by the use of batch tools implemented by CERSAT (and used for all its reprocessing works) such as Gogolist.

Prun is a tool which aims to ease the execution and monitoring of embarrassingly parallel processings. It is a wrapper which submits jobarrays to batch schedulers (torque/maui, oar, pbspro...) and manages the output logs and progress reports, with an easy monitoring.

Jobarray submission to batch schedulers is generally easy, but monitoring the job progress and accessing to the error logs of a few tasks among thousands is the same as finding a needle in a haystack... This wrapper was created to avoid spending more time in manual “logfile-mining” (grep, tail...) than the processing time itself. Some other needs it addresses:

- listing management avoids too big jobarrays (which crashes some batch schedulers)
- list the tasks in errors to ease their reprocessing
- allowing several execution modes : sequential, streaming (pipe)

- allowing multiple batch scheduler as backend
- having tasks status reporting and real-time monitoring
- having job history, meta-data & log files organization

This tool is easily handled by users and has been used successfully in other projects involving multiple remote partners such as OceanFlux. Online help is available but direct support par CERSAT engineering team is also provided.

The general scenario for the reprocessing of the altimeter or SAR data with a dedicated processor provided by a project partner will be as follows:

Objective:
A partner wants to reprocess or analysis a complete time series of data, which requires large processing resources and accessing a huge volume of data.

Scenario:
The user has a processor or analysis script ready to be used. It may be in python (preferably) or any other language available on the *Datarmor* platform.

If a language or a dependency is not installed on *Datarmor*, it can be added by the facility operation team. If specific system requirements are mandatory (for example, a specific linux distribution and version), a virtual server or a docker image can be installed but we encourage users to rather implement and test their processors on the default server operating system of *Datarmor*.

The user connects to the platform by SSH to a virtual machine on the cluster. It uploads its processor. The Ifremer system team takes it from here and test it on sample data, comparing the result with the partner to ensure they are identical. Integration and testing issues are solved by direct interaction between the Ifremer team and the partner providing the processor.

If the same task (processor) is to be run over a large amount of files, the processing can be easily distributed over different nodes of *Datarmor* using *gogolist*, a user friendly batch manager. The operator just needs to provide the list of files he wants to process, the processor command line, possibly the minimal resource requirements (ex: memory): the batch manager takes in charge the execution of the processor over each file and reports to the user on processing progress, errors. Unprocessed files and logs are clearly reported back to the user at the overall processing completion.

The results can be analysed remotely by the project partners by connecting through SSH to *Datarmor*.

There are two major activities at Ifremer which deal with data services, namely Coriolis/SISMER for in situ and CERSAT for satellite. Their respective operational components are currently being merged. They provide the evidence of Ifremer operational capacity to host the processing part of the CCI Seat State Service. The software for the CCI data processing system will also be based on the existing multi-mission (re)processing framework available from Ifremer for satellite data, the majority of which has been developed within the scope of operational or demonstration services, such as Medspiration, GlobWave, OSI SAF or CMEMS, to implement a real-time operational capacity under strong timeliness and robustness requirements by oceanographic and meteorological agencies. Ifremer is also an expertise and processing center for SMOS mission (CATDS, PI-MEP, SMOS-Storm,) and upcoming French-Chinese wave observation mission CFOSAT.

The operational center is operated by a support technical team (CLS), whereas the Ifremer engineers monitor the data processing, ensuring quality control, interfaces with the different entities involved in these operational activities or with other collaborating data centers, as well as maintenance or improvement of the operating system.

To perform its archiving and processing activities, as well as to support its research team in their activities, CERSAT has implemented a set of resources that will also support the project.

To cope with the reprocessing requirements of the CCI Sea State, the service will be operated on the new Ifremer supercomputer, *Datarmor*. *Datarmor* extends the capabilities of the previous cluster for:

- CPU intensive or memory demanding applications and processing
- massive data storage

It consists of:

- 11088 cores - 426 Tflops (128 GB memory and 28 cores per node) for HPC applications
- 240 cores, 5 TB RAM for non MPI processing
- 10 servers for web services
- large data storage capacity
    - 500 TB Lustre for HPC
    - 1.5 PB Luster storage for reference data
    - 5 PB GPFS storage for project and work data
    - 100 TB for services and web applications
    - 32 TB per home directory

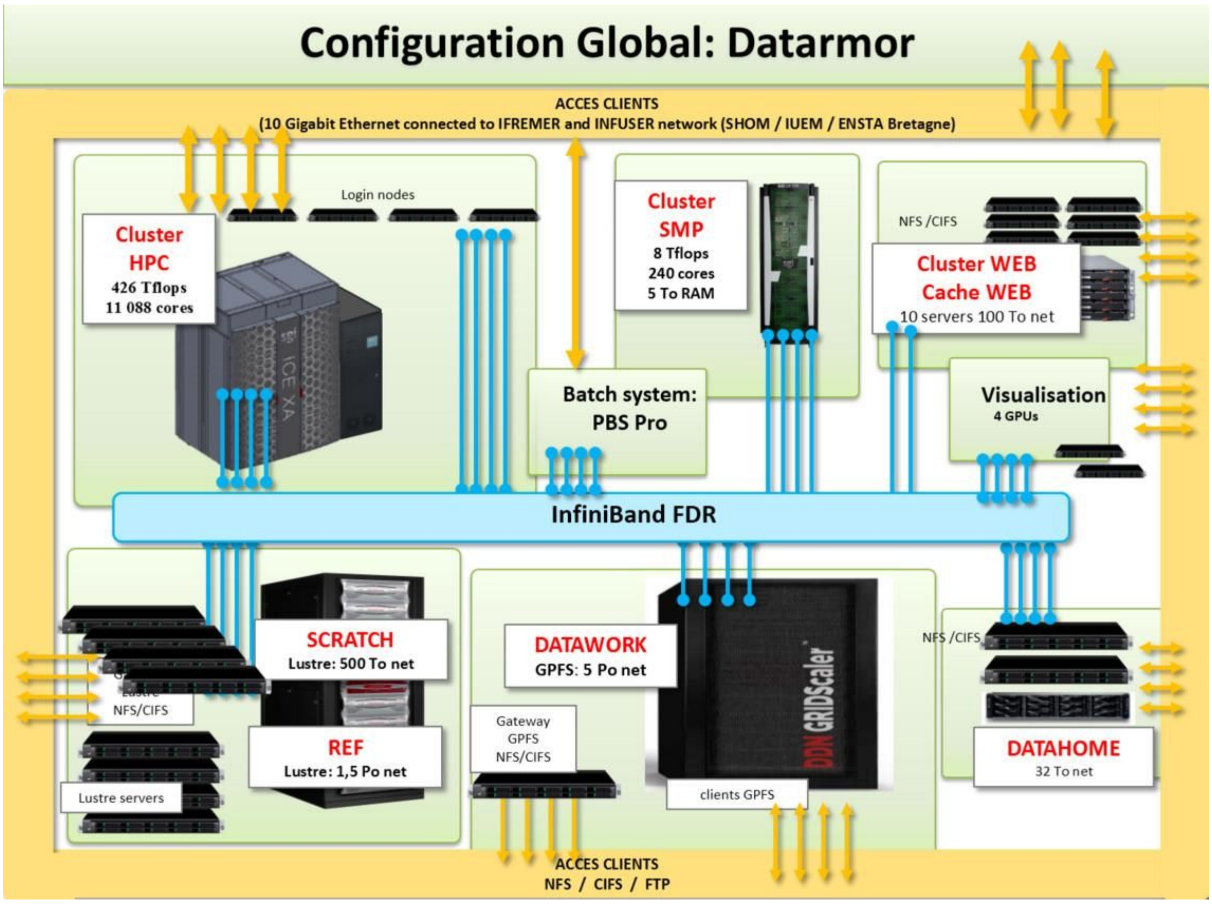

1.5 PB storage is currently reserved for CERSAT processing in this infrastructure which is deemed to be sufficient for the CCI Sea State requirements.

ESA SARvatore

The Grid-Processing-On-Demand Sentinel-3 & CryoSat SAR/SARin Processing service is a web platform that provides the capability to process online and on demand Sentinel-3 SAR and CryoSat SAR and SARin data, from L1a (FBR) data products until Level-2 geophysical data products, with different algorithms then Ground Segment products. The service, coined **SARvatore** (**SAR** **V**ersatile **A**ltimetric **T**oolkit for **O**cean **R**esearch & **E**xploitation), can process data over any surface.

The service is open, free of charge and accessible online from everywhere. To access the service, users need EO-SSO (Earth Observation Single Sign-On) credentials. Registration can be done at https://earth.esa.int/web/guest/general-registration). Afterwards, an e-mail to the G-POD team (eo-gpod@esa.int) shall be sent requesting the activation of the CryoSat-2/Sentinel-3 service for the EO-SSO user account.

In particular it hosts the highly voluminous L1A and L1B data for Sentinel-3 SRAL. In addition, the platform provides engineering support to help the integration and deployment of the processors the CCI Sea State may need to run on the platform.

CNES HAL

The CNES HPC center is already used for some altimetry reprocessing campaigns for AVISO project. It is also a possible target or complement to our project for addressing the Jason mission series and possibly others implemented on this platform (like ERS or Envisat for FCDR project).