AMS Data Production Facilities at Science Operation Center at CERN

V. Choutko[1], A. Egorov[1], A. Eline[1], B. Shan[2]

[1]Massachusetts Institute of Technology, [2]Beihang University

Abstract

The Alpha Magnetic Spectrometer (AMS) is a high energy physics experiment on the board of the International Space Station (ISS). This paper presents the hardware and software facilities of Science Operation Center (SOC) at CERN. Standard Data Production is build around AMS production server - a scalable distributed service which links together a set of different programming modules for science data de-framing and reconstruction. The server has the capacity to manage around 1000 multithreaded jobs, i.e. up to 32K CPUs. Monitoring and management tools with enhanced GUI are also described.

Introduction

The AMS Science Operation Center at CERN receives and stores all AMS science and housekeeping data, ensures full science data reconstruction, calibration and alignment, keeps data available for further analysis at CERN and AMS Regional Centers. Standard Data Production undertaken by SOC computing infrastructure, is split into three stages:

- pre-production - provides de-framing of ~1 minute science data frames into ~23 minutes raw files
- first production - provides initial data validation and indexing, produces data summary files and event tags for second production
- second production - produces analysis ready datasets

SOC computing infrastructure

Includes SOC own computing resources as well as CERN dedicated computing resources, rented by AMS experiment.

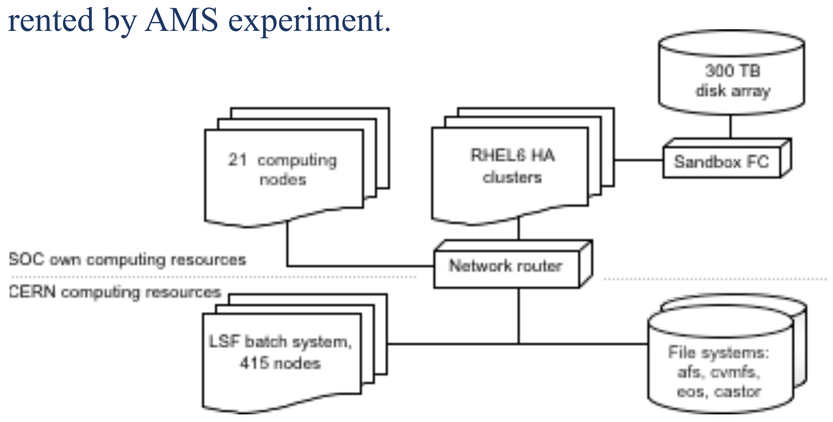

Figure 1. SOC computing infrastructure

Its own computing resources consist of 21 nodes powered by 264 Intel Xeon processors, and run under SLC5/6, CC7 operating systems. It provides computing power for pre-production and first production running 24/7 on freshly arrived data.

CERN dedicated computing resources consist of 415 compute nodes powered by ~6200 Xeon processors, and run SLC6 operating system. They are used primary for physics analysis as well as for second production and MC-simulation. Second production runs every 3-6 months incrementally.

The main amount of disk space is provided by CERN. The EOS file system has been chosen as a main storage to keep raw data and resulting CERN ROOT files. To provide data safety all files are archived to CERN CASTOR tape system.

File system	Capacity	Destination	Access
EOS	4 PB	Main storage to keep raw data and resulting CERN ROOT files	Distributed: CERN, Remote Computing Centres
CASTOR	-	Tape archiver for all raw data and resulting CERN ROOT files	Distributed: CERN, Remote Computing Centres
AFS	25 TB	Conditional database with Time Dependent Values (TDV), logs of execution of Standard Data Production modules, AMS Offline Software repositories and development framework	Distributed: CERN, Remote Computing Centres
CVMFS	10 TB	Fast read-only mirror of AFS	Distributed: CERN, Remote Computing Centres
SOC disk array	300 TB	Redundant disk space	Only inside SOC computing subnetwork

Table 1. SOC disk storage resources

AMS Offline Software has the ability of operator assistant or automatic redirection of IO streams in case of losing disk access, its performance degradation or by other technical reason.

Software facilities of AMS Standard Data Production

The software used on pre-production stage does not require computing power and may run on SOC own front-ends. Unlike pre-production, first and second production use the computing power of ~80 and few thousand processors correspondingly. The software used on first and second production implements clients-server approach. Data processing is done by remote job producers and is managed by scalable service called AMS production server. Other production modules are involved as clients :

- WEB server – prepares job scripts, delivers job execution requests to the server by schedule or by operator request
- job validator – validates finished jobs as well as produced ROOT data files, provides file transfer to disk storage. The list of finished jobs and ready data files is taken from the server
- Oracle Berkeley DB engine – is a run-time mirror of the internal tables of the server
- production GUI - provides monitoring of first and second production

The server links all production modules together by ORBIT2 CORBA interface for inter-process communications. Used ORBIT2 API is extended with orbitcpp bindings and is fully compatible with client software modules written with C++, Python and Perl programming languages.

AMS production server

The node to start the producer is chosen by the server using the criteria:

$$\max((p_{node}-b_{node})*clk_{node})$$

where p_{node} is predefined amount of processors, b_{node} is number of busy processors, clk_{node} is predefined CPU ratio. The number of processors to start multithreaded producer is also chosen basing on node configuration.

On start the producer requests from the server calibration constants and slow control corrections. The state of job processing is periodically sent to the server (once per ~20 sec). On finish the producer uploads the server with event summary files. The server also provides:

- monitoring of job execution status: waiting for submission, pended, running, finished, failed, killed, canceled
- job restarting on job failure, the number of retries is predefined
- job killing on timeout
- logging of the server-producer conversation

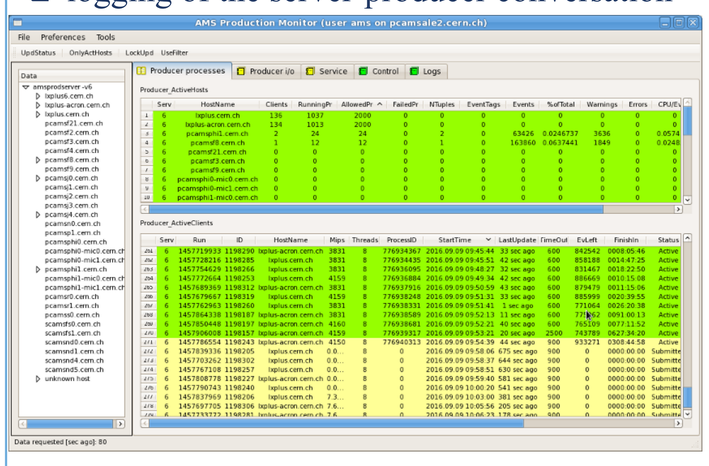

Figure 2. Enhanced production GUI view

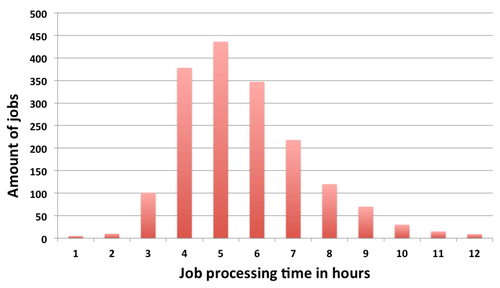

Figure 3. First production in 01-30.Sep.2016

The server has an enhanced production GUI providing the features:

- server configuration tools
- manual job management such as job canceling, restarting or removing from production
- job status monitoring and logging information
- server health monitoring
- pre-production health monitoring
- post-production health monitoring

Conclusions

The current status of SOC computing infrastructure is shown. The operation mode of Standard Data Production is highly automated and normally does not require human intervention. The enhanced GUI simplifies monitoring of Offline Data Processing as well as diagnostics of running jobs. From the beginning of operations of AMS on the board of ISS, more than 120000 run files (~88 millions of physical events) have been processed by Standard Data Production. Figure 3 shows data availability after first production.