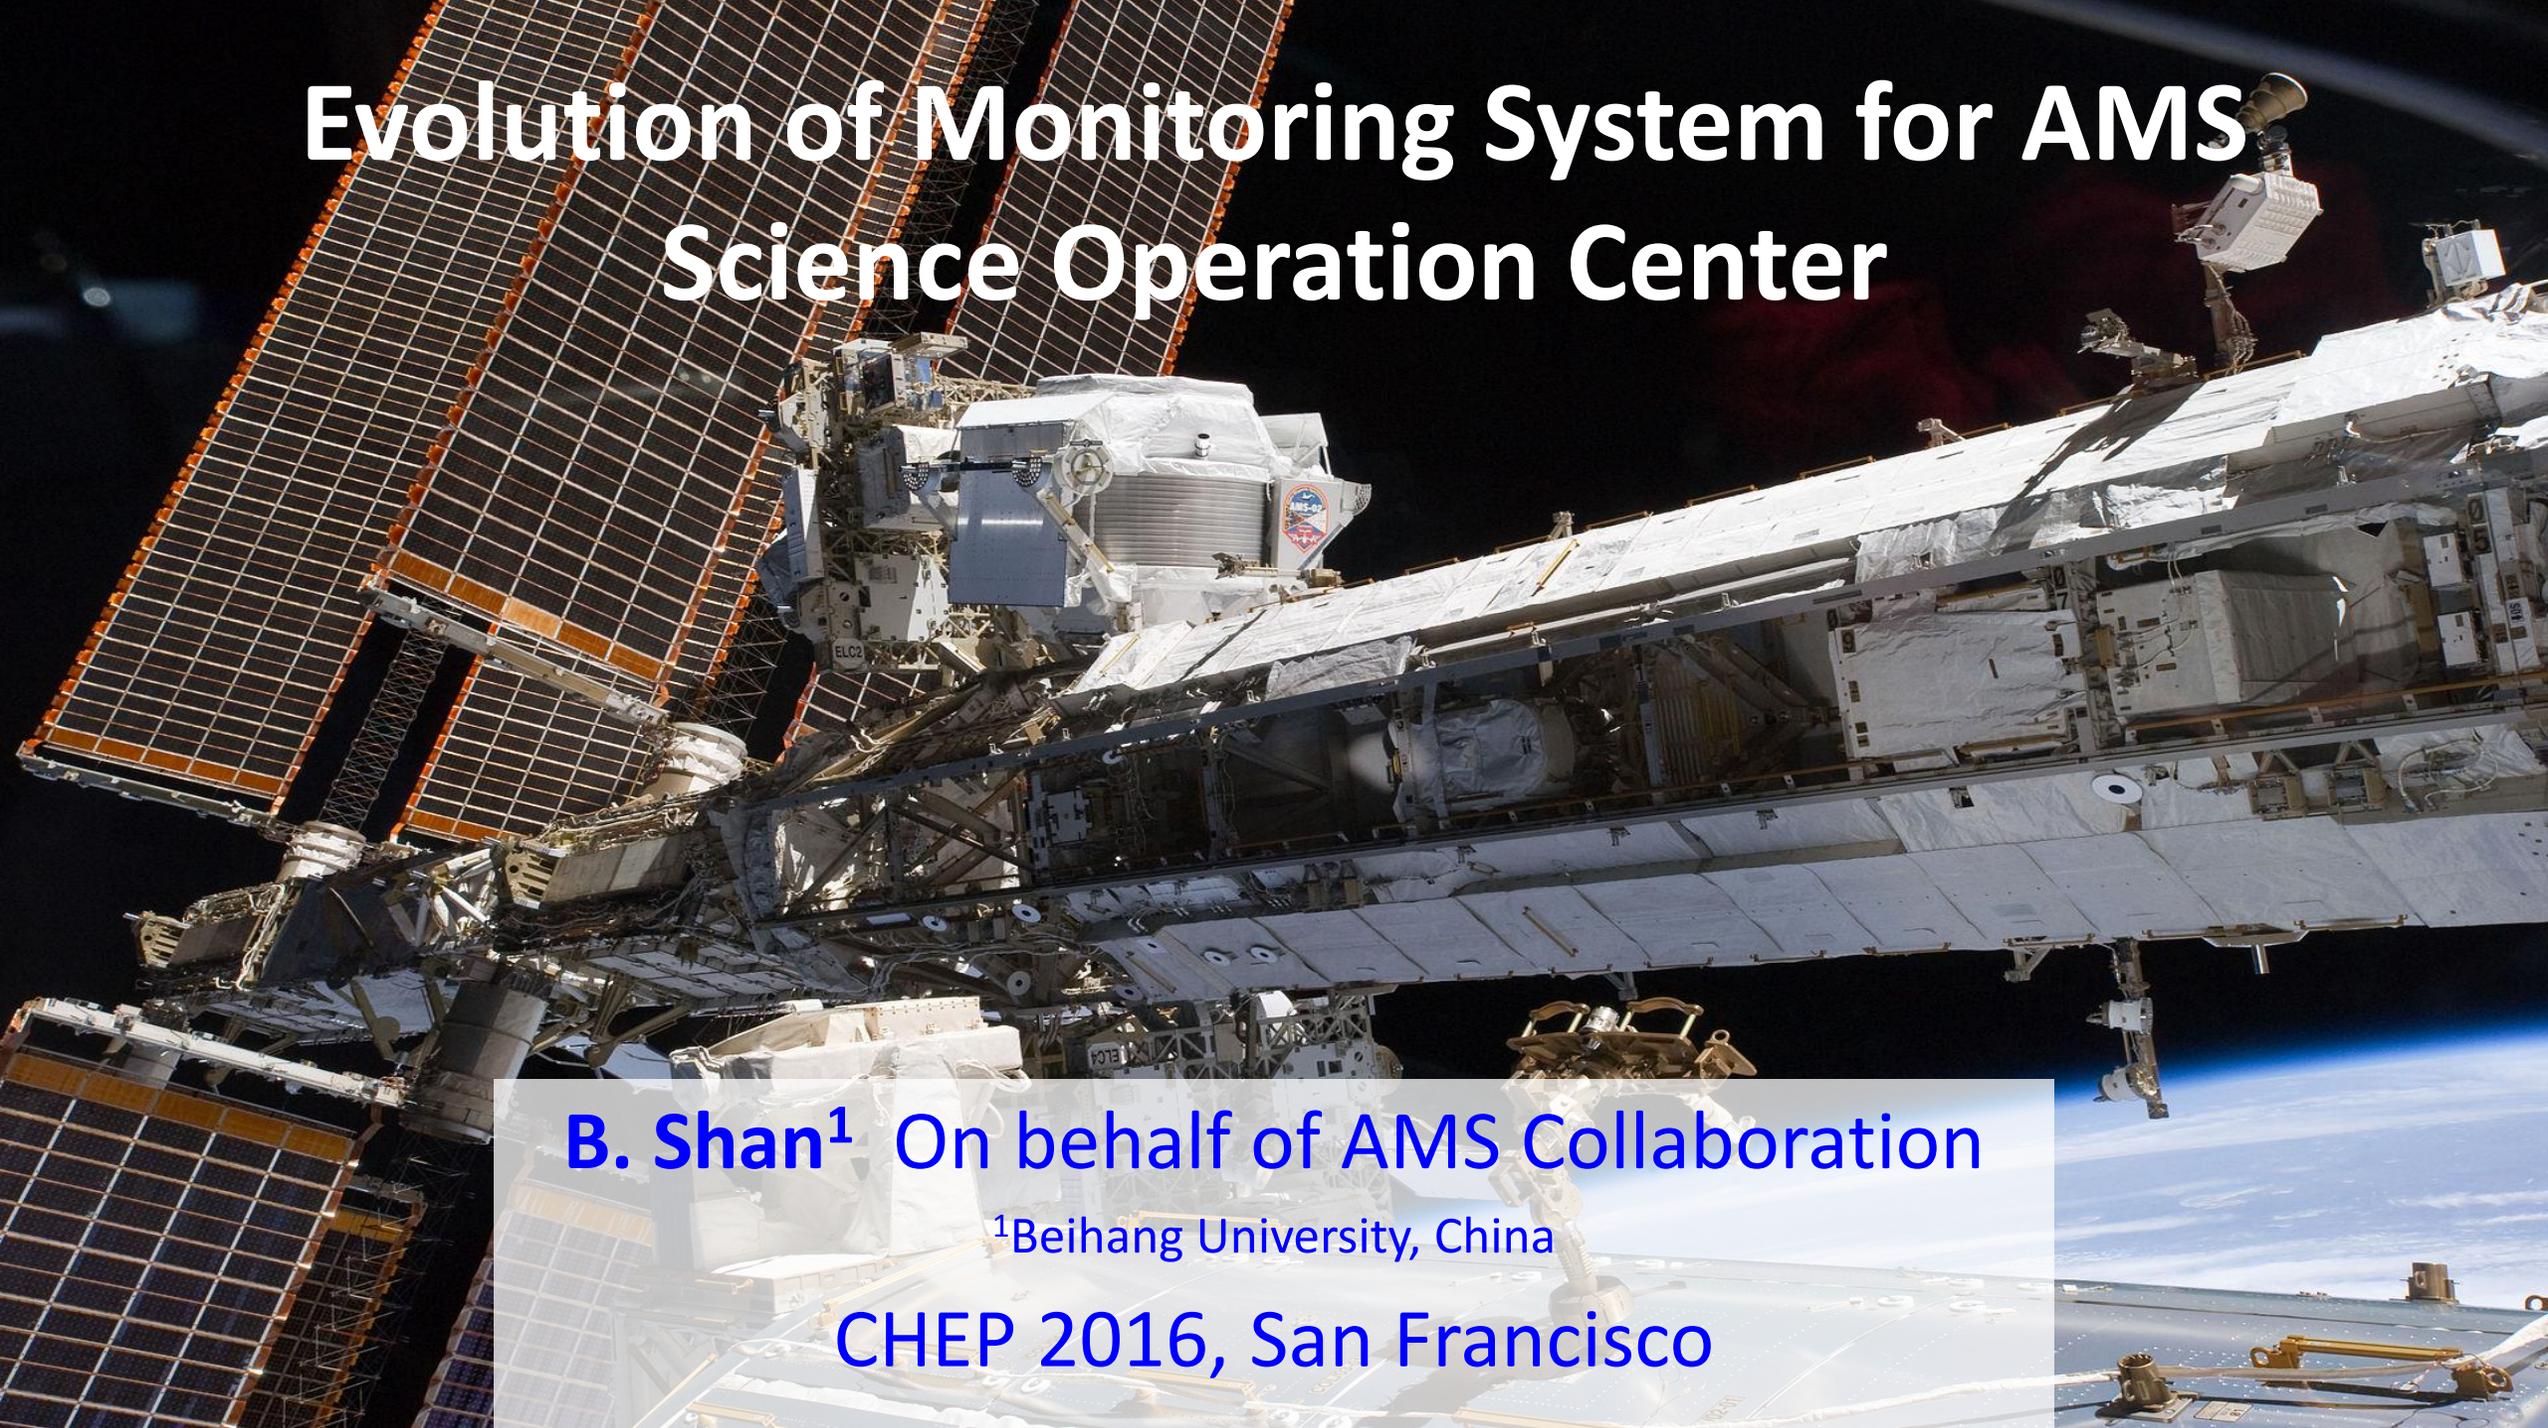


Evolution of Monitoring System for AMS Science Operation Center



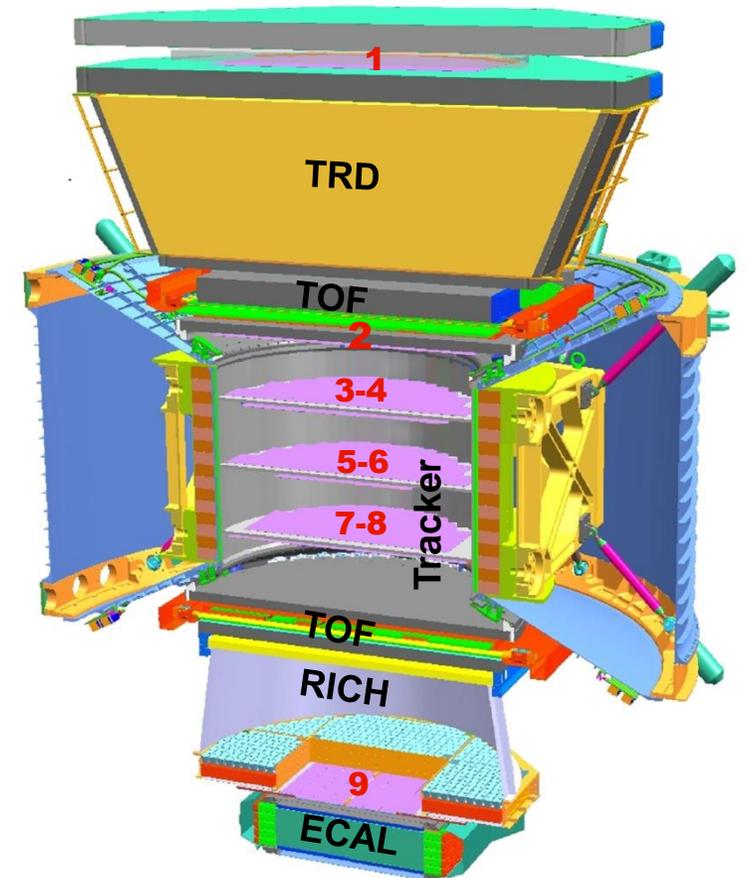
B. Shan¹ On behalf of AMS Collaboration

¹Beihang University, China

CHEP 2016, San Francisco

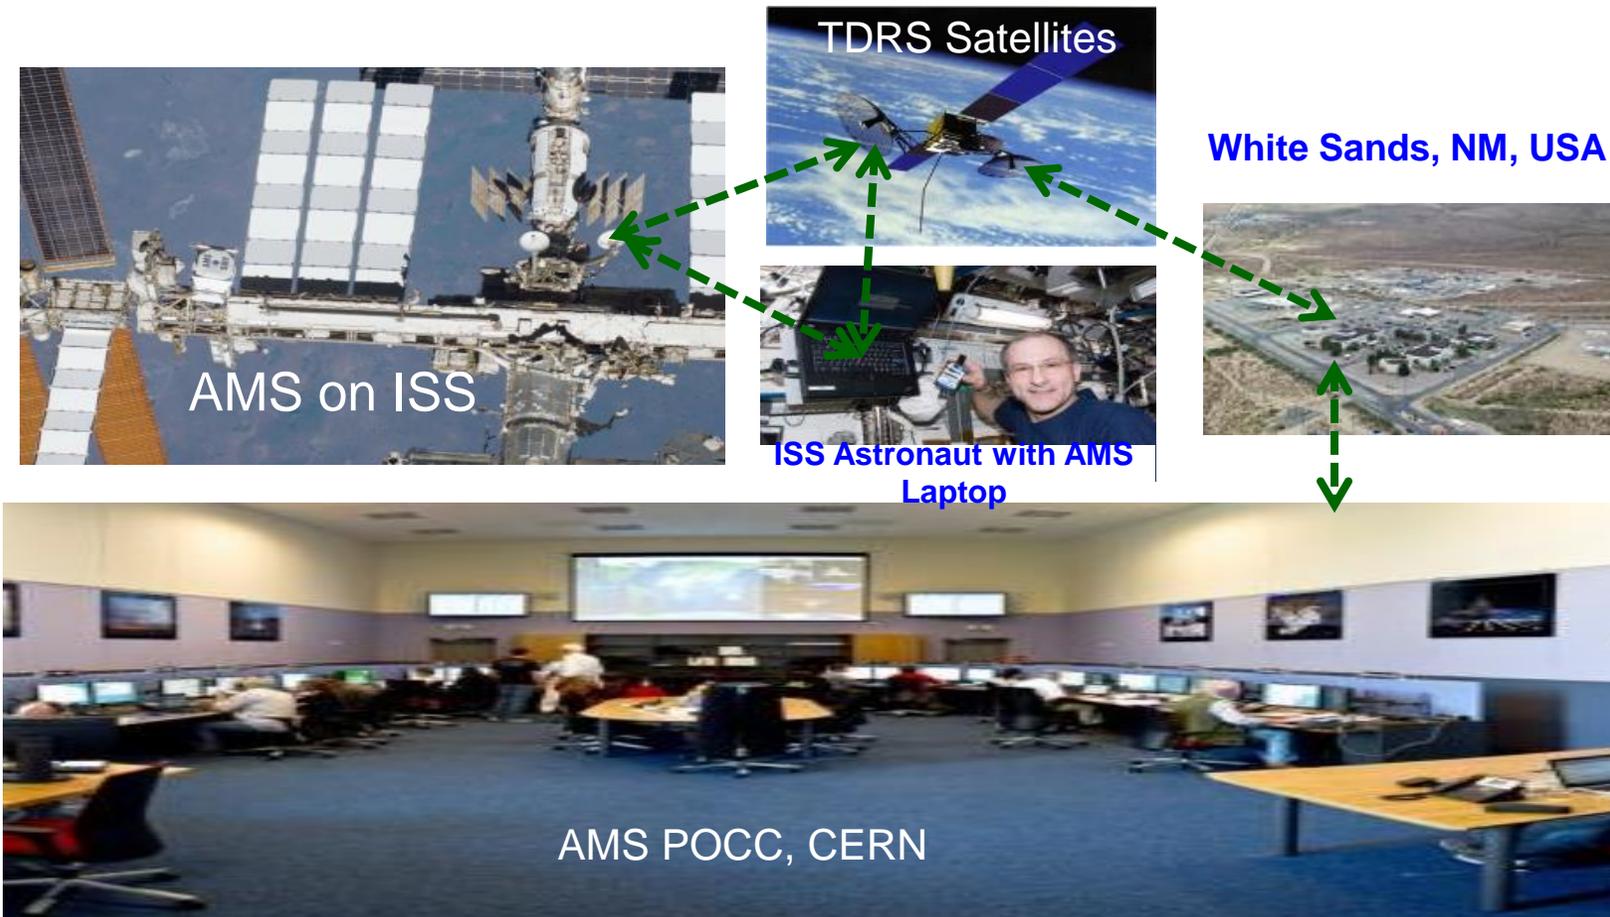
AMS Experiment

- The Alpha Magnetic Spectrometer (AMS) is a high energy physics experiment on board the International Space Station (ISS), featured:
 - Geometrical acceptance: **$0.5 \text{ m}^2 \cdot \text{sr}$**
 - Number of ReadOut Channels: **$\approx 200\text{K}$**
 - Main payload of Space Shuttle **Endeavour's last flight** (May 16, 2011)
 - **Installed on ISS on May 19, 2011**
 - 7x24 running
 - Up to now, over **60 billion events** collected
 - Max. event rate: **2KHz**



AMS Data Flow

- Data transferred via relay satellites to Marshall Space Flight Center, then to CERN, nearly real-time, in form of one-minute-*frames*
- **Preproduction**: Frames \rightarrow runs (**RAW**): 1 run = $\frac{1}{4}$ orbit (~ 23 minutes)
 - Run number is the UNIX time of the starting of the run



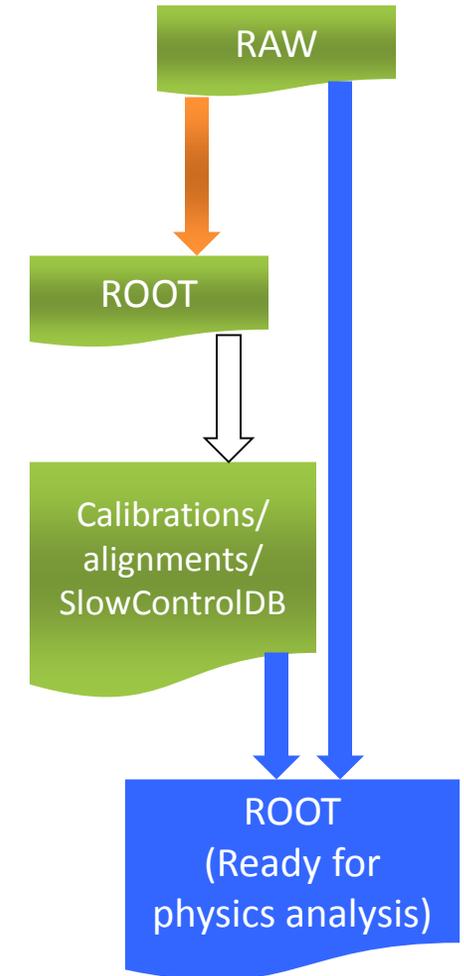
AMS Flight Data Reconstruction

- **Standard Production** ↓

- Runs 7x24 on freshly arrived data
- Initial data validation and indexing
- Produces Data Summary Files and Event Tags (ROOT) for fast events selection
- Usually be available within 2 hours after flight data arriving
- Used to produce various calibrations for the second production as well as quick performance evaluation

- **Pass-N Production** ↓

- Every 3-6 months incremental
- Full reconstruction in case of software major update
- To use all the available calibrations, alignments, ancillary data from ISS, and slow control data to produce physics analysis ready set of reconstructed data



Science Operation Center (SOC) of AMS

Responsibilities

- Processing of the AMS science data:
 - Preproduction
 - Standard production
 - Pass-N production
 - Monte-Carlo simulation
- Management of SOC own production farm
 - 21 hosts, 264 cores
 - 300 TB storage
- Operations on AMS resources at CERN
 - LSF, EOS, CASTOR, AFS, CVMFS, ELOG, PDB-R, etc.

Science Operation Center (SOC) of AMS

Responsibilities

- Processing of the AMS science data:

- **Frame arriving**



Frame Monitor

- Preproduction

- Standard production

- Pass-N production

- Monte-Carlo simulation

- Management of SOC own production farm

- 21 hosts, 264 cores

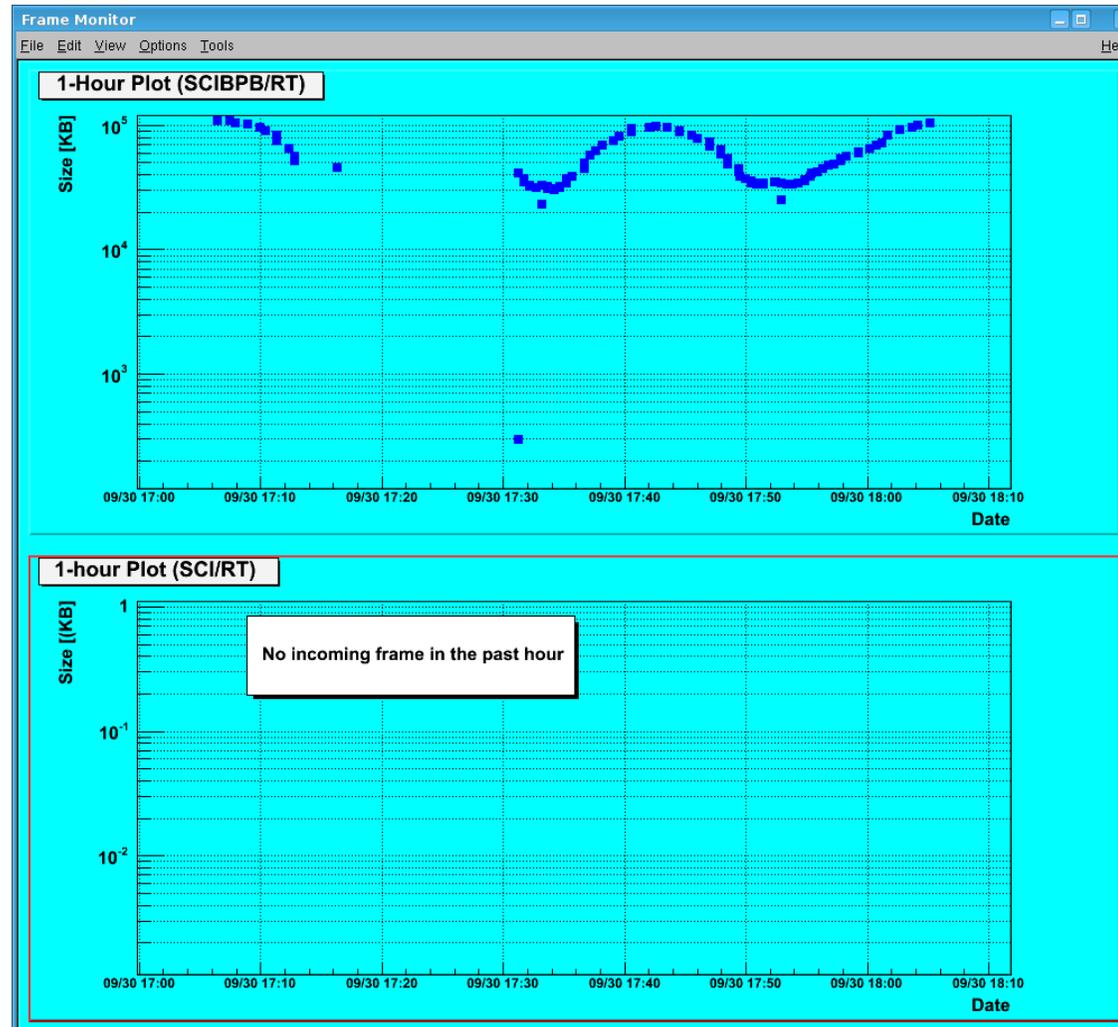
- 300 TB storage

- Operations on AMS resources at CERN

- LSF, EOS, CASTOR, AFS, CVMFS, ELOG, PDB-R, etc.

Frame Monitor

- To monitor the frame data statistics arrived at SOC



Science Operation Center (SOC) of AMS

Responsibilities

- Processing of the AMS science data:
 - Frame arriving
 - **Preproduction**
 - Standard production
 - Pass-N production
 - Monte-Carlo simulation
- Management of SOC own production farm
 - 21 hosts, 264 cores
 - 300 TB storage
- Operations on AMS resources at CERN
 - LSF, EOS, CASTOR, AFS, CVMFS, ELOG, PDB-R, etc.



Data files webpage

Data files web page

- A web page to list the RAW/ROOT files



Run	FilePath	StartTime	Events	Type	SizeMB	OrigFiles
1475410459 / 57f0fa1b / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475410459	02/10/2016 14:14:18	623159	SCI 5 19 4 / OK	1249	SCIBPB/RT/2527/328 SCIBPB/RT/2527/351
/eosams/Data/AMS02/2014/ISS.B1070/std/1475410459.00000001.root						
1475411870 / 57f0ff9e / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475411870	02/10/2016 14:37:49	472328	SCI 5 19 4 / OK	1005	SCIBPB/RT/2527/351 SCIBPB/RT/2527/374
/eosams/Data/AMS02/2014/ISS.B1070/std/1475411870.00000001.root						
1475413238 / 57f104f6 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475413238	02/10/2016 15:00:36	783507	SCI 5 19 4 / OK	1704	SCIBPB/RT/2527/374 SCIBPB/RT/2527/398
/eosams/Data/AMS02/2014/ISS.B1070/std/1475413238.00000001.root						
1475414648 / 57f10a78 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475414648	02/10/2016 15:24:06	602434	SCI 5 19 4 / OK	1207	SCIBPB/RT/2527/399 SCIBPB/RT/2527/421
/eosams/Data/AMS02/2014/ISS.B1070/std/1475414648.00000001.root						
1475416015 / 57f10fcf / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475416015	02/10/2016 15:46:53	591960	SCI 5 19 4 / OK	1197	SCIBPB/RT/2527/421 SCIBPB/RT/2527/445
/eosams/Data/AMS02/2014/ISS.B1070/std/1475416015.00000001.root						
1475417426 / 57f11552 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475417426	02/10/2016 16:10:24	504729	SCI 5 19 4 / OK	1046	SCIBPB/RT/2527/445 SCIBPB/RT/2527/469
/eosams/Data/AMS02/2014/ISS.B1070/std/1475417426.00000001.root						
1475418793 / 57f11aa9 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475418793	02/10/2016 16:33:11	705976	SCI 5 19 4 / OK	2330	SCIBPB/RT/2527/469 SCIBPB/RT/2527/492
/eosams/Data/AMS02/2014/ISS.B1070/std/1475418793.00000001.root						
1475420204 / 57f1202c / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475420204	02/10/2016 16:56:43	10285	SCI 5 19 4 / SHORT	22	SCIBPB/RT/2527/492 SCIBPB/RT/2527/493
1475420371 / 57f120d3 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475420371	02/10/2016 16:59:29	577538	SCI 5 19 4 / OK	1149	SCIBPB/RT/2527/495 SCIBPB/RT/2527/515
/eosams/Data/AMS02/2014/ISS.B1070/std/1475420371.00000001.root						
1475421571 / 57f12583 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475421571	02/10/2016 17:19:30	596795	SCI 5 19 4 / OK	1205	SCIBPB/RT/2527/515 SCIBPB/RT/2527/539
/eosams/Data/AMS02/2014/ISS.B1070/std/1475421571.00000001.root						
1475422982 / 57f12b06 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475422982	02/10/2016 17:43:01	592658	SCI 5 19 4 / OK	1200	SCIBPB/RT/2527/540 SCIBPB/RT/2527/562
/eosams/Data/AMS02/2014/ISS.B1070/std/1475422982.00000001.root						
1475424349 / 57f1305d / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475424349	02/10/2016 18:05:48	619638	SCI 5 19 4 / OK	2916	SCIBPB/RT/2527/562 SCIBPB/RT/2527/586
/eosams/Data/AMS02/2014/ISS.B1070/std/1475424349.00000001.root						
1475425760 / 57f135e0 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475425760	02/10/2016 18:29:19	629312	SCI 5 19 4 / OK	1254	SCIBPB/RT/2527/587 SCIBPB/RT/2527/610
/eosams/Data/AMS02/2014/ISS.B1070/std/1475425760.00000001.root						
1475427127 / 57f13b37 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475427127	02/10/2016 18:52:06	605692	SCI 5 19 4 / OK	1221	SCIBPB/RT/2527/610 SCIBPB/RT/2527/633
/eosams/Data/AMS02/2014/ISS.B1070/std/1475427127.00000001.root						
1475428538 / 57f140ba / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475428538	02/10/2016 19:15:37	731797	SCI 5 19 4 / OK	1444	SCIBPB/RT/2527/633 SCIBPB/RT/2527/656
/eosams/Data/AMS02/2014/ISS.B1070/std/1475428538.00000001.root						
1475429905 / 57f14611 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475429905	02/10/2016 19:38:24	774414	SCI 5 19 4 / OK	2886	SCIBPB/RT/2527/656 SCIBPB/RT/2527/680
/eosams/Data/AMS02/2014/ISS.B1070/std/1475429905.00000001.root						
1475431316 / 57f14b94 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475431316	02/10/2016 20:01:55	632200	SCI 5 19 4 / OK	1259	SCIBPB/RT/2527/680 SCIBPB/RT/2527/703
/eosams/Data/AMS02/2014/ISS.B1070/std/1475431316.00000001.root						
1475432684 / 57f150ec / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475432684	02/10/2016 20:24:42	617910	SCI 5 19 4 / OK	1237	SCIBPB/RT/2527/703 SCIBPB/RT/2527/727
/eosams/Data/AMS02/2014/ISS.B1070/std/1475432684.00000001.root						
1475434094 / 57f1566e / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475434094	02/10/2016 20:48:12	914319	SCI 5 19 4 / OK	1779	SCIBPB/RT/2527/728 SCIBPB/RT/2527/750
/eosams/Data/AMS02/2014/ISS.B1070/std/1475434094.00000001.root						
1475435461 / 57f15bc5 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475435461	02/10/2016 21:10:59	836275	SCI 5 19 4 / OK	1637	SCIBPB/RT/2527/750 SCIBPB/RT/2527/775
/eosams/Data/AMS02/2014/ISS.B1070/std/1475435461.00000001.root						
1475436873 / 57f16149 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475436873	02/10/2016 21:34:32	681502	SCI 5 19 4 / OK	1351	SCIBPB/RT/2527/775 SCIBPB/RT/2527/798
/eosams/Data/AMS02/2014/ISS.B1070/std/1475436873.00000001.root						
1475438240 / 57f166a0 / f114	/afs/cern.ch/ams/Offline/RunsDir/SCI/1475000000/1475438240	02/10/2016 21:57:18	639563	SCI 5 19 4 / OK	1276	SCIBPB/RT/2527/798 SCIBPB/RT/2527/821
/eosams/Data/AMS02/2014/ISS.B1070/std/1475438240.00000001.root						

Science Operation Center (SOC) of AMS

Responsibilities

- Processing of the AMS science data:
 - Frame arriving
 - Preproduction
 - **Standard production** 
 - Pass-N production
 - Monte-Carlo simulation
- Management of SOC own production farm
 - 21 hosts, 264 cores
 - 300 TB storage
- Operations on AMS resources at CERN
 - LSF, EOS, CASTOR, AFS, CVMFS, ELOG, PDB-R, etc.



Production Monitor

- Monitoring/managing standard production

The screenshot displays the 'AMS Production Monitor' application window. The interface includes a menu bar (File, Preferences, Tools), a toolbar (UpdStatus, OnlyActHosts, LockUpd, UseFilter), and a main data area. On the left, a tree view shows the system hierarchy under 'Data', including 'amsprodserver -v5' and various hostnames like 'lxplus6.cern.ch'. The main area is divided into two sections: 'Producer_Runs' and 'Producer_NTuples'.

Producer_Runs Table:

Serv	Run	Uid	Time	FirstEvent	LastEvent	Priority	History	Fails	DataMC	Host	Status
19	5	1475239725	2016.09.30 17:35:47	1	1093	0	ToBeRerun	0	1	lxplus6.cern.ch	Finished
20	5	1475239722	2016.09.30 17:37:28	1	1090	0	ToBeRerun	0	1	pcamsr1.cern.ch	Finished
21	5	1475239706	2016.09.30 17:31:40	1	1085	0	ToBeRerun	0	1	pcamsj2.cern.ch	Finished
22	5	1475242468	2016.09.30 17:58:41	1	494	3	ToBeRerun	0	1	pcamss0.cern.ch	Finished
23	5	1475245246	2016.09.30 17:58:49	1	494	3	ToBeRerun	0	1	pcamss0.cern.ch	Finished
24	5	1475238281	2016.09.30 17:18:53	1	598304	1	ToBeRerun	0	1	scamsnd5.cern.ch	Processing
25	5	1475234134	2016.09.30 16:24:05	1	637795	1	ToBeRerun	0	1	scamsfs0.cern.ch	Processing
26	5	1475227165	2016.09.30 15:11:10	1	927286	1	ToBeRerun	0	1	scamsnd4.cern.ch	Processing
27	5	1475232723	2016.09.30 16:23:59	1	729810	1	ToBeRerun	0	1	scamsfs1.cern.ch	Processing
28	5	1475229944	2016.09.30 16:24:04	1	919063	1	ToBeRerun	0	1	pcamsn0.cern.ch	Processing
29	5	1475235501	2016.09.30 17:19:03	1	854183	1	ToBeRerun	0	1	pcamsn0.cern.ch	Processing
30	5	1475231356	2016.09.30 16:24:23	1	580171	1	ToBeRerun	0	1	pcamsp1.cern.ch	Processing
31	5	1475228576	2016.09.30 15:11:15	1	918753	1	ToBeRerun	0	1	scamsnd1.cern.ch	Processing
32	5	1475236914	2016.09.30 17:18:56	1	591937	1	ToBeRerun	0	1	scamsnd0.cern.ch	Processing
33	5	1475224385	2016.09.30 13:12:30	1	1010848	1	ToBeRerun	0	1	scamsnd0.cern.ch	Processing
34	5	1475239733	2016.09.30 17:18:48	1	500425	1	ToBeRerun	0	1	scamsnd5.cern.ch	Processing
35	5	1475223018	2016.09.30 13:12:24	1	1147670	1	ToBeRerun	0	1	scamsnd4.cern.ch	Finished

Producer_NTuples Table:

Serv	Run	Time	FirstEvent	LastEvent	Name	CRC	Size(MB)	Status	
1	5	1475242468	2016.09.30 17:59:15	1	494	pcamss0.cern.ch:/eosams/Data/AMS02/2014/calib.v5/calib_pm.v5/1475242468/1475242468.0000000...	2095736743	0	Validated
2	5	1475245246	2016.09.30 17:59:16	1	494	pcamss0.cern.ch:/eosams/Data/AMS02/2014/calib.v5/calib_pm.v5/1475245246/1475245246.0000000...	1637200217	0	Validated
3	5	1475223018	2016.09.30 18:00:48	1	1147670	scamsnd4.cern.ch:/eosams/Data/AMS02/2014/ISS.B1070/std/1475223018/1475223018.00000001.r...	2719937555	14031	Validated
4	5	1475224385	2016.09.30 18:08:40	1	0	scamsnd0.cern.ch:/dat0/local/logs/nt/data/SCI/1475224385.00000001.root	0	8613	InProgress
5	5	1475225798	2016.09.30 17:37:07	1	620998	scamsnd1.cern.ch:/eosams/Data/AMS02/2014/ISS.B1070/std/1475225798/1475225798.00000001.r...	3141013866	5656	Validated
6	5	1475227165	2016.09.30 18:08:42	1	0	scamsnd4.cern.ch:/dat0/local/logs/nt/data/SCI/1475227165.00000001.root	0	4697	InProgress
7	5	1475228576	2016.09.30 18:08:49	1	0	scamsnd1.cern.ch:/dat0/local/logs/nt/data/SCI/1475228576.00000001.root	0	10601	InProgress
8	5	1475232723	2016.09.30 18:08:47	1	0	scamsfs1.cern.ch:/dat0/local/logs/nt/data/SCI/1475232723.00000001.root	0	4475	InProgress
9	5	1475234134	2016.09.30 18:08:41	1	0	scamsfs0.cern.ch:/dat0/local/logs/nt/data/SCI/1475234134.00000001.root	0	4362	InProgress
10	5	1475231356	2016.09.30 18:08:51	1	0	pcamsp1.cern.ch:/dat0/local/logs/nt/data/SCI/1475231356.00000001.root	0	2443	InProgress
11	5	1475229944	2016.09.30 18:08:56	1	0	pcamsn0.cern.ch:/dat0/local/logs/nt/data/SCI/1475229944.00000001.root	0	2567	InProgress
12	5	1475239733	2016.09.30 18:08:26	1	0	scamsnd5.cern.ch:/dat0/local/logs/nt/data/SCI/1475239733.00000001.root	0	1053	InProgress
13	5	1475238281	2016.09.30 18:09:03	1	0	scamsnd5.cern.ch:/dat0/local/logs/nt/data/SCI/1475238281.00000001.root	0	1205	InProgress
14	5	1475236914	2016.09.30 18:08:46	1	0	scamsnd0.cern.ch:/dat0/local/logs/nt/data/SCI/1475236914.00000001.root	0	958	InProgress
15	5	1475235501	2016.09.30 18:08:50	1	0	pcamsn0.cern.ch:/dat0/local/logs/nt/data/SCI/1475235501.00000001.root	0	1014	InProgress
16	5	1475239714	2016.09.30 17:35:00	1	1089	pcamsf4.cern.ch:/eosams/Data/AMS02/2014/laser/std/1475239714/1475239714.00000001.root	3204239912	16	Validated
17	5	1475239716	2016.09.30 17:32:35	1	1089	pcamsf4.cern.ch:/eosams/Data/AMS02/2014/laser/std/1475239716/1475239716.00000001.root	284771160	12	Validated

Data requested [sec ago]: 60

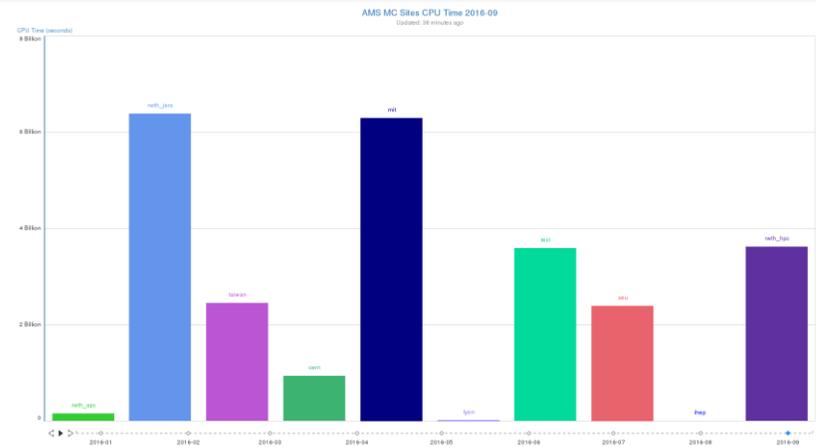
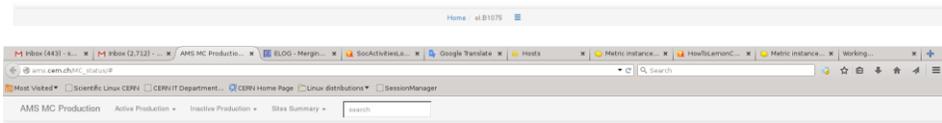
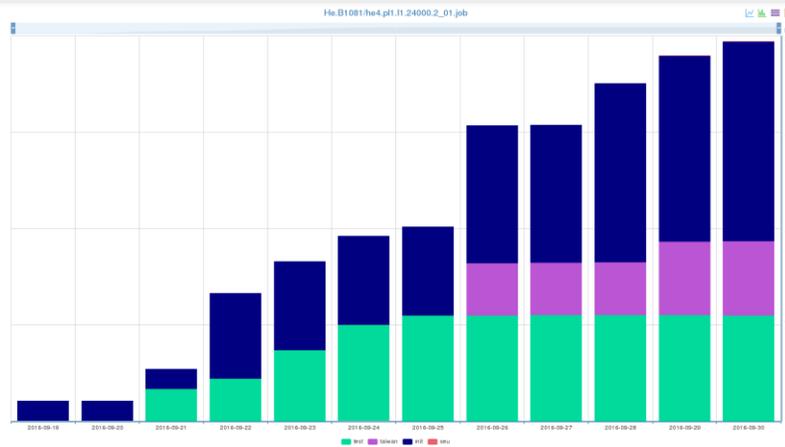
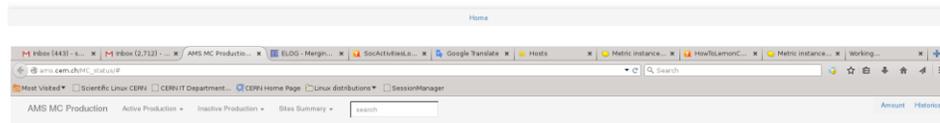
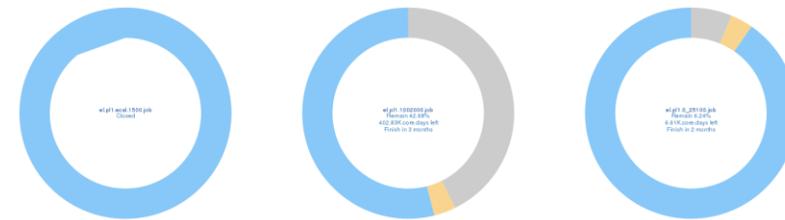
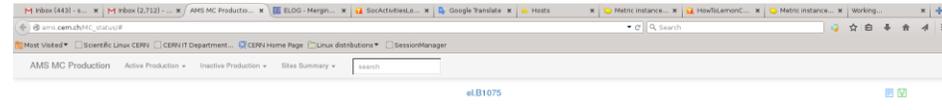
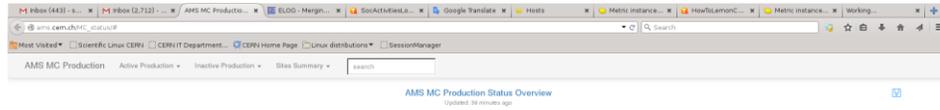
Science Operation Center (SOC) of AMS

Responsibilities

- Processing of the AMS science data:
 - Preproduction
 - Standard production
 - Pass-N production
 - **Monte-Carlo simulation** 
- Management of SOC own production farm
 - 21 hosts, 264 cores
 - 300 TB storage
- Operations on AMS resources at CERN
 - LSF, EOS, CASTOR, AFS, CVMFS, ELOG, PDB-R, etc.

MC Production Monitoring Tool

- Web-based MC progress monitoring (Poster-284)



Science Operation Center (SOC) of AMS

Responsibilities

- Processing of the AMS science data:
 - Frame arriving
 - Preproduction
 - Standard production
 - Pass-N production
 - Monte-Carlo simulation
- Management of SOC own production farm
 - **21 hosts, 264 cores**
 - **300 TB storage**

```
graph LR; A[21 hosts, 264 cores  
300 TB storage] --> B[NetMonitor]
```
- Operations on AMS resources at CERN
 - LSF, EOS, CASTOR, AFS, CVMFS, ELOG, PDB-R, etc.

NetMonitor

- NetMonitor is a monitoring tool without GUI to poll the health status of SOC own computing facilities, including:
 - Production servers/hosts
 - NFS servers
 - NFS mounted volumes on NFS clients
 - Software: production server/Frame decoder/Raw validator
- In case of issues, NetMonitor will send out warning messages to the system administrators

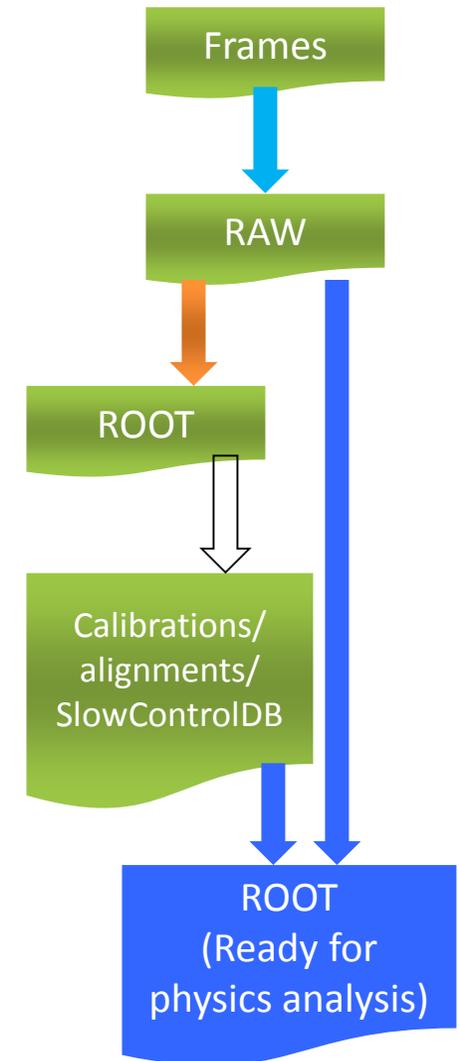
Integrating to the new monitoring infrastructure of CERN

- The above mentioned monitoring tools
 - Powerful
 - “Fragmented”, lack of a “glance view”
 - Most have no direct access from outside of CERN
- Design principles:
 - Unified interface: including all related elements which should be monitored
 - Knowledge base: to define rules for automatic trouble shooting

The key criteria for production health:

Delay

- Frame delay (last frame arrived)
 - If the frames are arriving normally
- RAW delay (latest validated RAW)
 - If frame decoder/raw validator works normally
- ROOT delay (latest validated ROOT)
 - If production server/producers/root validator work normally
- FRESH RUN delay (validated largest RUN number)
 - If job requesting works normally

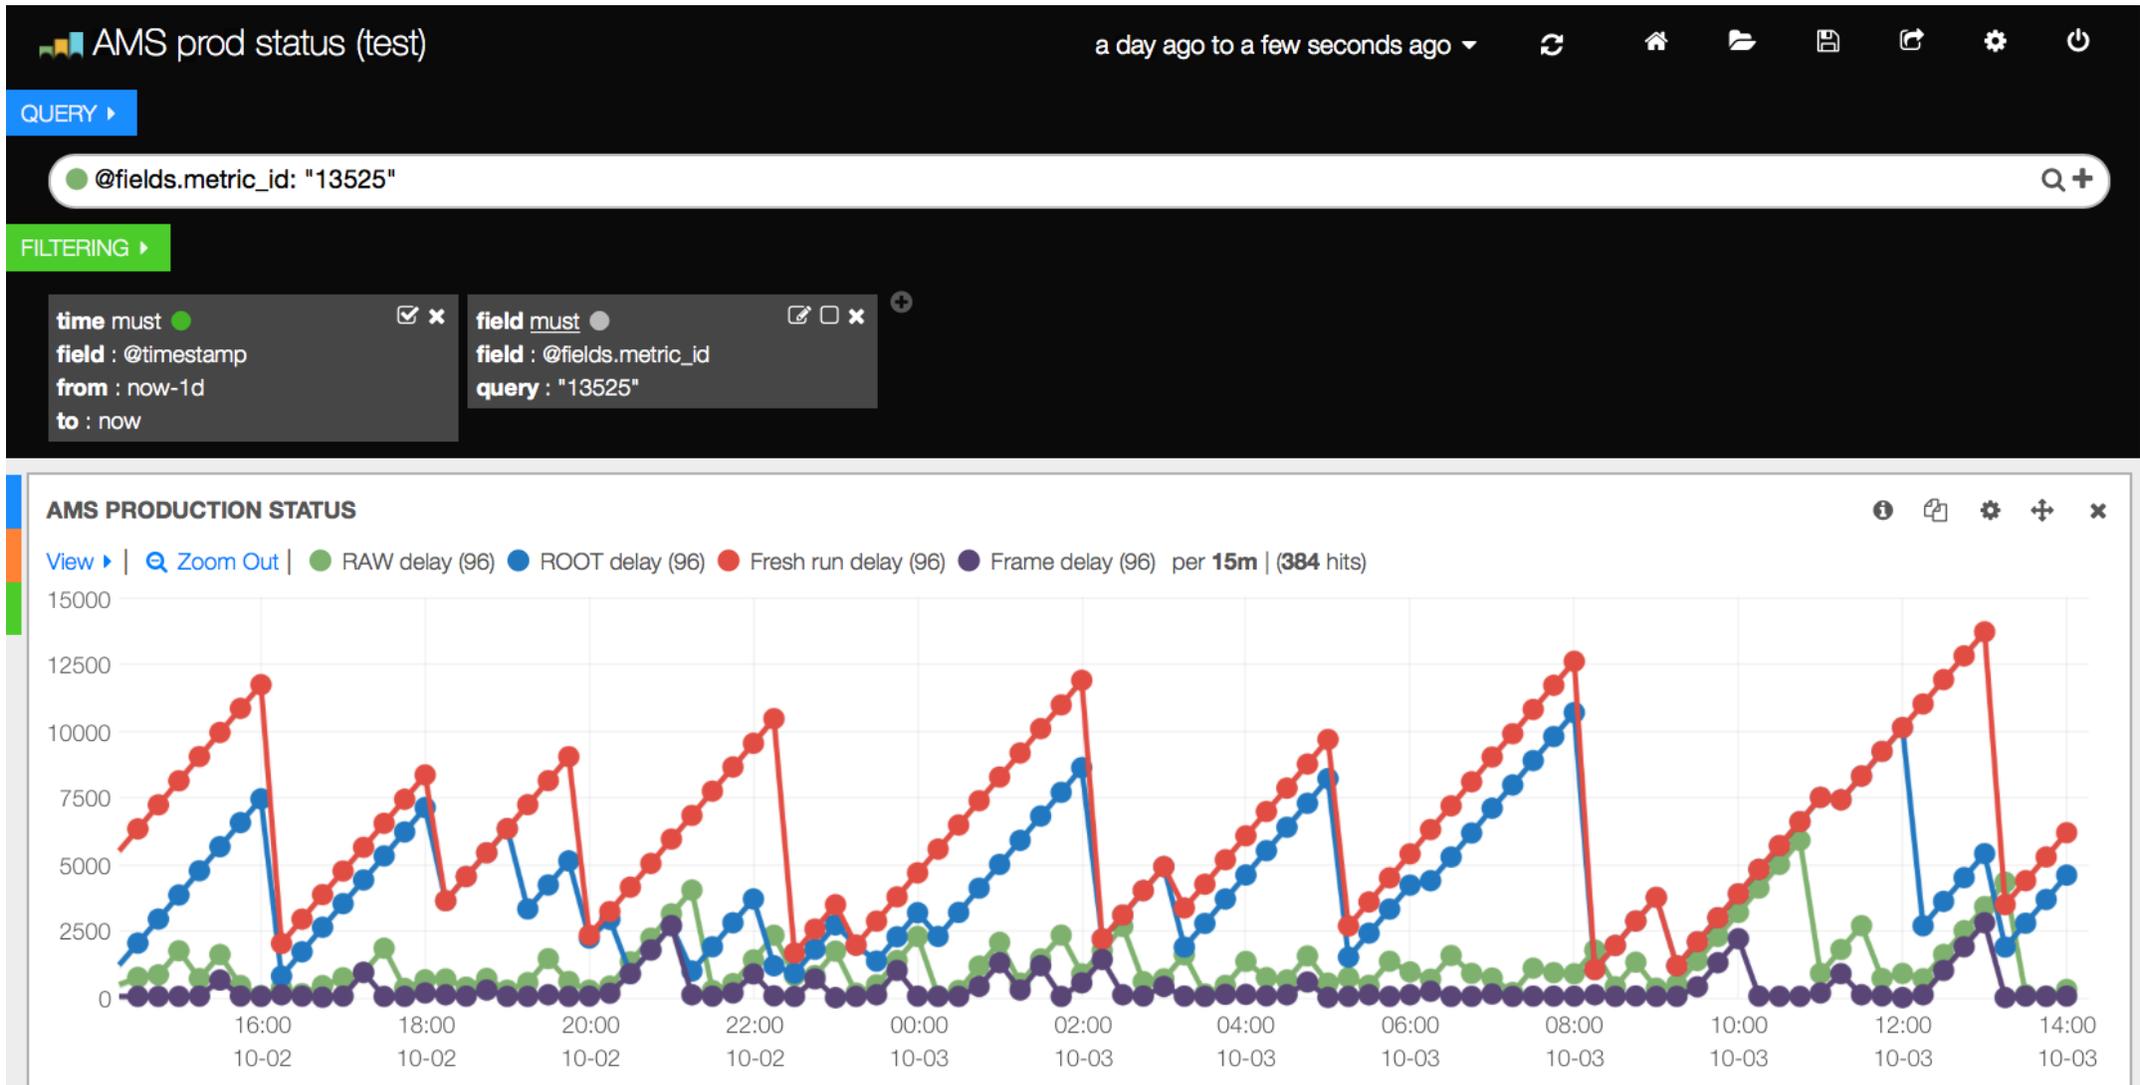


Metric, metric class, sensor

- To describe the delays, a metric class named “amsSOC.proddelay” is defined in “Lemon Metrics Manager”
 - 4 integers to indicate the 4 delay time in seconds
- A metric instance and a sensor are defined based on the metric class
- A puppet based virtual host is created to run this sensor

Web monitoring

- A Kibana page is created to monitor the delay data

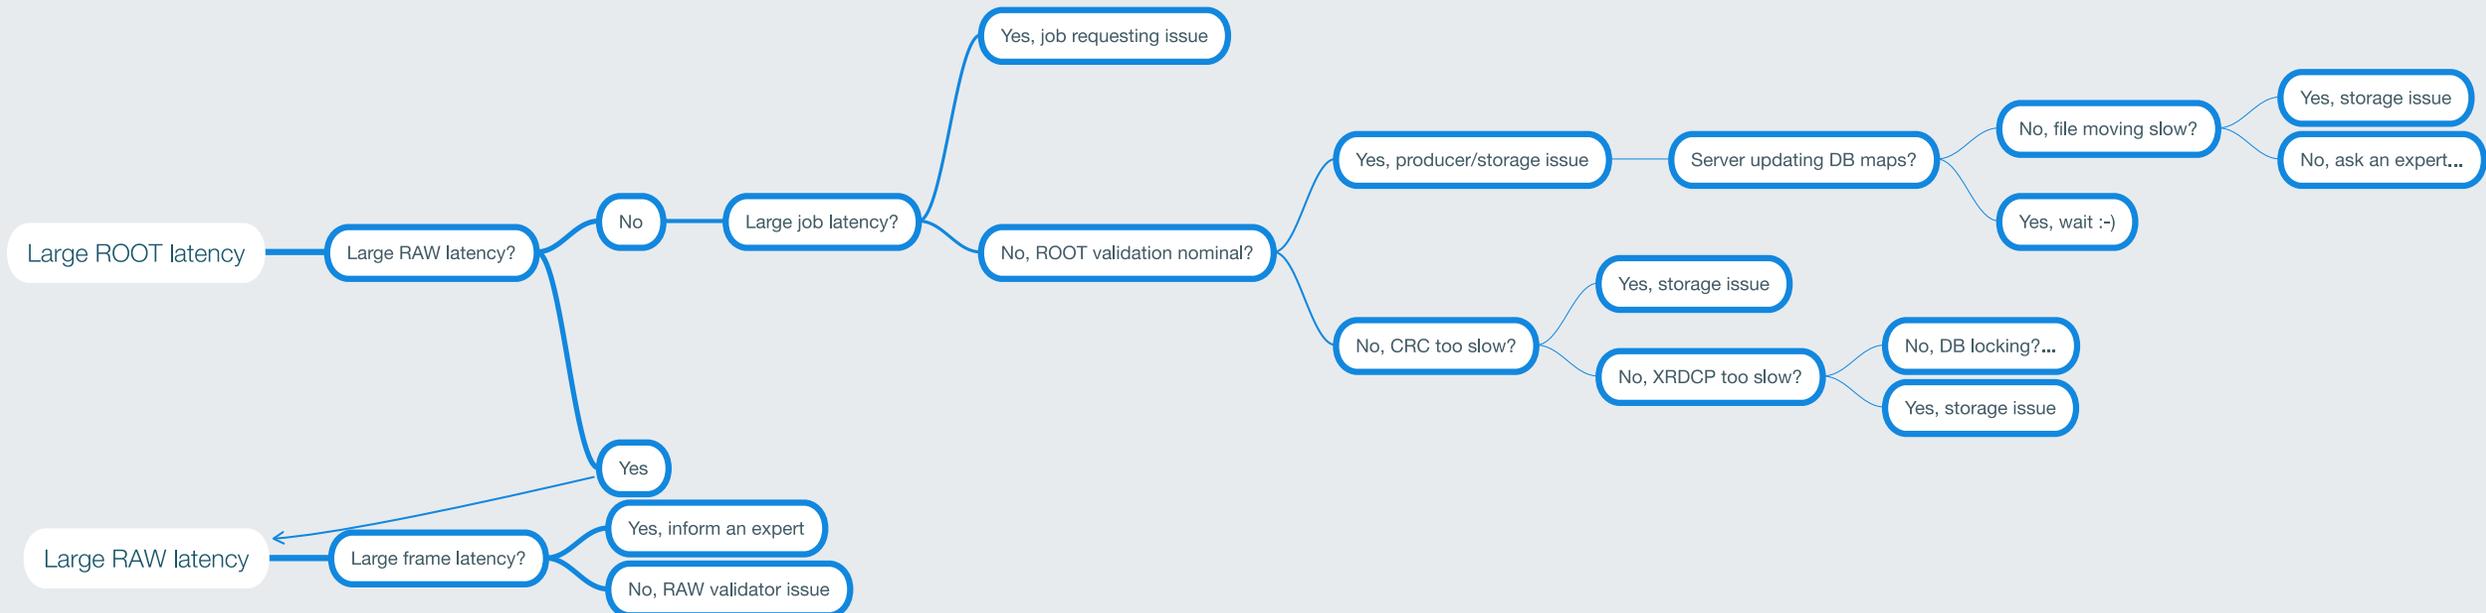


Notification and trouble shooting

- Thresholds are defined for notifications
- Trouble shooting
 - To define rules in the database, which is called rule base
 - Program goes step by step according to the rules defined in the rule base
 - Notifications are sent to the experts about the production status as well as the potential reason(s) given by the diagnostics program
 - After each manual intervention, rule base can be updated by experts to make the diagnostics program smarter than before, hopefully

Notification and trouble shooting

- Example of trouble shooting rule base



Summary

- Legacy SOC monitoring tools provide powerful functions but “fragmented”
- Based on the monitoring infrastructure provided by CERN, a monitoring system for SOC is designed, which provides:
 - Glance view of SOC status
 - Web accessible
 - Automatic notification
- Trouble shooting assistant
 - Experiences from experts are extracted into rules and stored into database, expandable
 - Automatic trouble shooting attempt is carried out in case of issues
 - Manual interventions act as feedbacks to existing rules

Future works

- To manage SOC physical boxes by Puppet, to monitor their health status by sensors, and to integrate the sensor data to the monitoring page
- To add more sensors for SOC programs, like RAW/ROOT validators, job requestor, data movers, etc., and integrate their data to the monitoring page
- To integrate resource availability data of CERN services (EOS/CASTOR/AFS/CVMFS...) into the monitoring page
- To improve the rule base and provide better assistance for trouble shooting