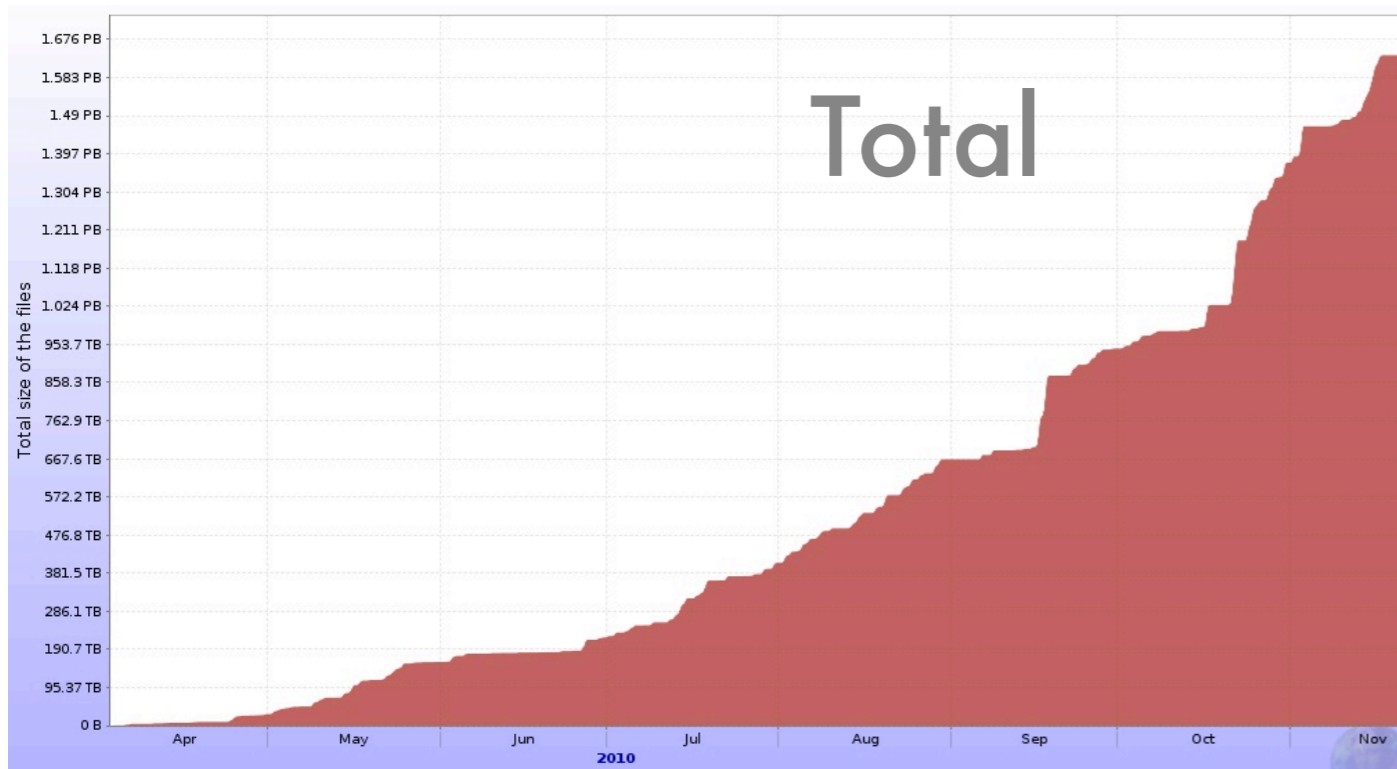


# L'analyse dans Alice


L. Aphecetche

# Prises de données



1,6 PB RAW (pp & PbPb)


220 TB RAW PbPb  
(max. possible cette année  
= 2 PB)



# RAW vers ESDs

- Pour pp reconstruction (pass1 @ T0) quasi-online (i.e. ESDs disponibles quelques heures après la fin d'un run)
- pass2 (@T1) typiquement quelques mois après
- Pour PbPb, un peu plus difficile de garder le rythme, même à la luminosité actuelle... Mais quand même :
  - [arXiv:1011.3916](#) Charged-particle multiplicity density at mid-rapidity in central Pb-Pb collisions at  $\sqrt{s_{NN}} = 2.76$  TeV
  - [arXiv:1011.3914](#) Elliptic flow of charged particles in Pb-Pb collisions at 2.76 TeV

# ESDs vers AODs



- Trains d'analyse
  - Lire les données une fois
  - appliquer autant de tâches d'analyse que possible sur ces données
  - appliqué à l'échelle de l'analyse individuelle ou organisée (seul varie le nombre de wagons)

Production info						Jobs status				
ID	Tag	Status	Done%	Cfg	Out	Total	Done	Active	Waiting	Runs
687	QA35_LHC10h6sim_Merging	Completed	0%			1	0			1 (137161 -
684	FILTERmuon026_LHC10h_Merging	Completed	97%			46	45			41 (-1 -
683	QA35_LHC10h6sim	Completed	100%			402	402			1 (137161 -
682	FILTERmuon026_LHC10h	Running	98%			5901	5811	1	4	42 (137161 -
677	FILTERpass1_025_LHC10h	Technical stop	1%			1146	13			9 (137161 -
674	QA34_LHC10h_Merging	Running	88%			45	40	1		40 (137135 -
673	QA34_LHC10h	Running	95%			5974	5694	13	3	44 (137135 -
669	FILTERpass1_024_LHC10h_Merging	Completed	100%			7	7			7 (137135 -
668	QA33_LHC10h_Merging	Completed	93%			16	15			13 (137135 -
667	FILTERpass1_024_LHC10h	Completed	71%			659	469			7 (137135 -
666	QA33_LHC10h	Completed	99%			2068	2061			13 (137135 -
665	FILTERpass1023_LHC10e_Merging	Completed	100%			162	162			162 (127719 -
661	FILTERpass1_023_LHC10e	Completed	95%			6138	5881			171 (127719 -
660	FILTERpass1_022_LHC10h_Merging	Completed	100%			11	11			10 (136833 -
659	QA32_LHC10h_Merging	Completed	100%			11	11			10 (136833 -
333	QA7_LHC10b	Completed	100%			466	466			
327	QA5_LHC10d_Merging	Completed	85%			7	6			
326	QA5_LHC10d	Completed	98%			535	527			
321	QA5_LHC10c_Merging	Completed	50%			4	2			
320	QA5_LHC10c	Completed	99%			1586	1584			
316	QA4_LHC10c_Merging	Completed	83%			65	54			
315	QA4_LHC10c	Completed	94%			31312	29736			
312	QA3_LHC10c_Merging	Completed	75%			57	43			
311	QA3_LHC10c	Completed	92%			26190	24216			
302	QA2_LHC10c	Completed	25%			1876	479			
300	QA2_LHC10b_Merging	Completed	88%			188	91			
299	QA2_LHC10b	Completed	49%			21554	10667			
295	QA1_LHC10b_Merging	Completed	55%			142	79			
294	QA1_LHC10b	Completed	70%			6702	4752			
291	TR023_LHC10b4	Completed	71%			239	171			
290	TR022_LHC10b3	Completed	80%			422	339			
289	TR021_LHC10b2	Completed	97%			868	843			
288	TR020_LHC10b1	Completed	99%			760	758			
266	QA16_LHC10a6ESD	Completed	98%			342	338			
<b>237 productions</b>			<b>91%</b>			<b>410853</b>	<b>374452</b>	<b>15</b>	<b>42</b>	

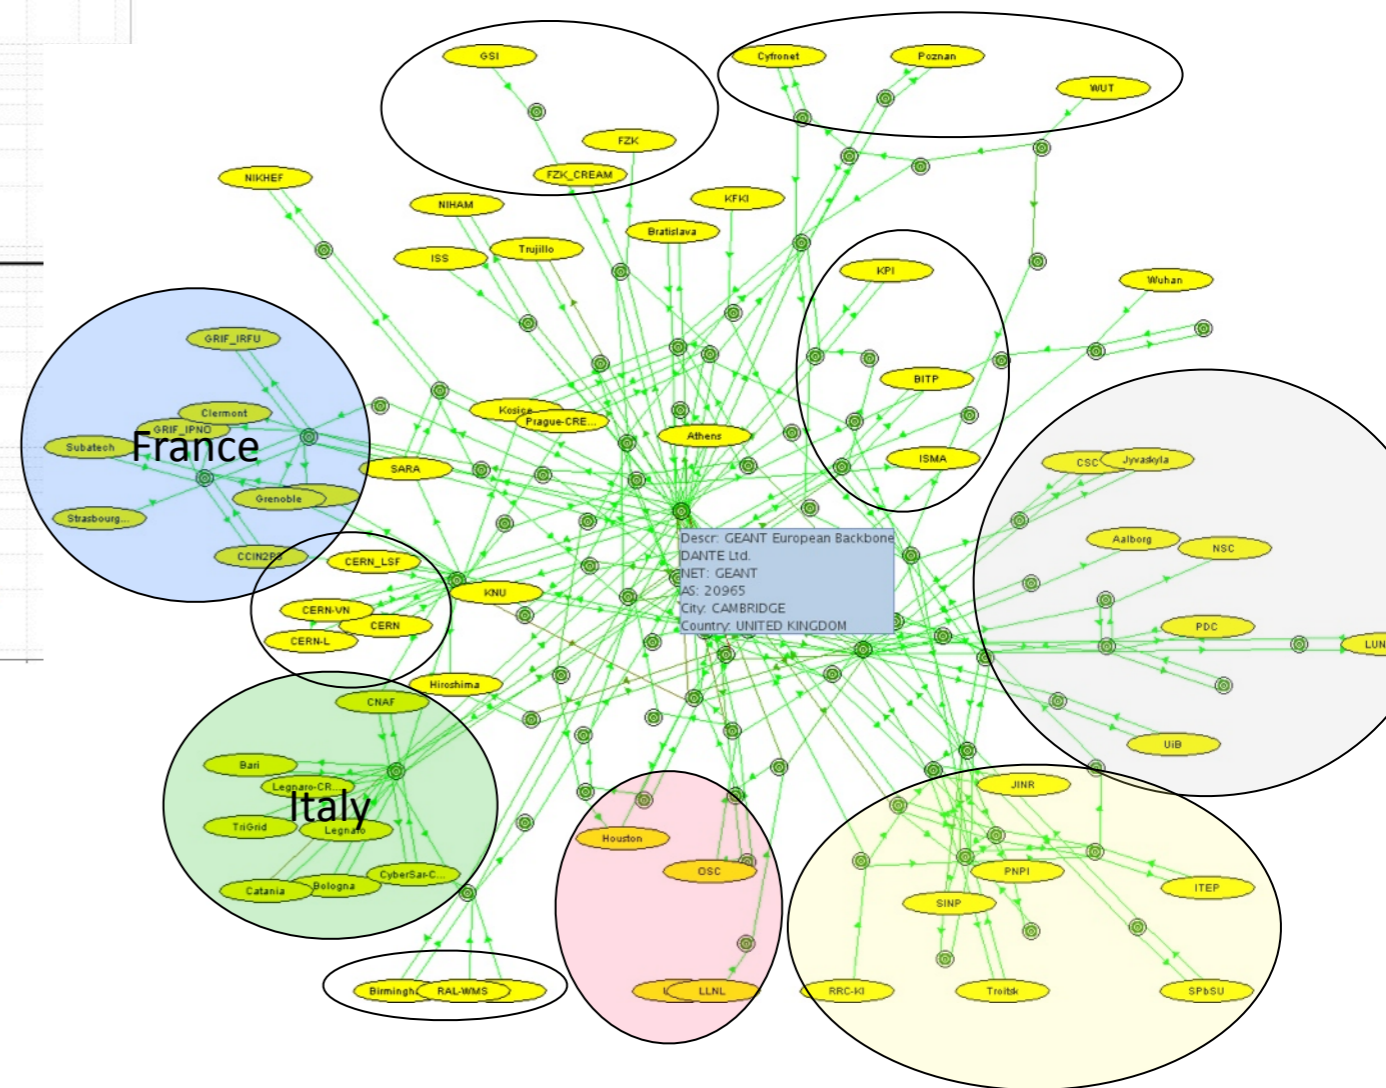
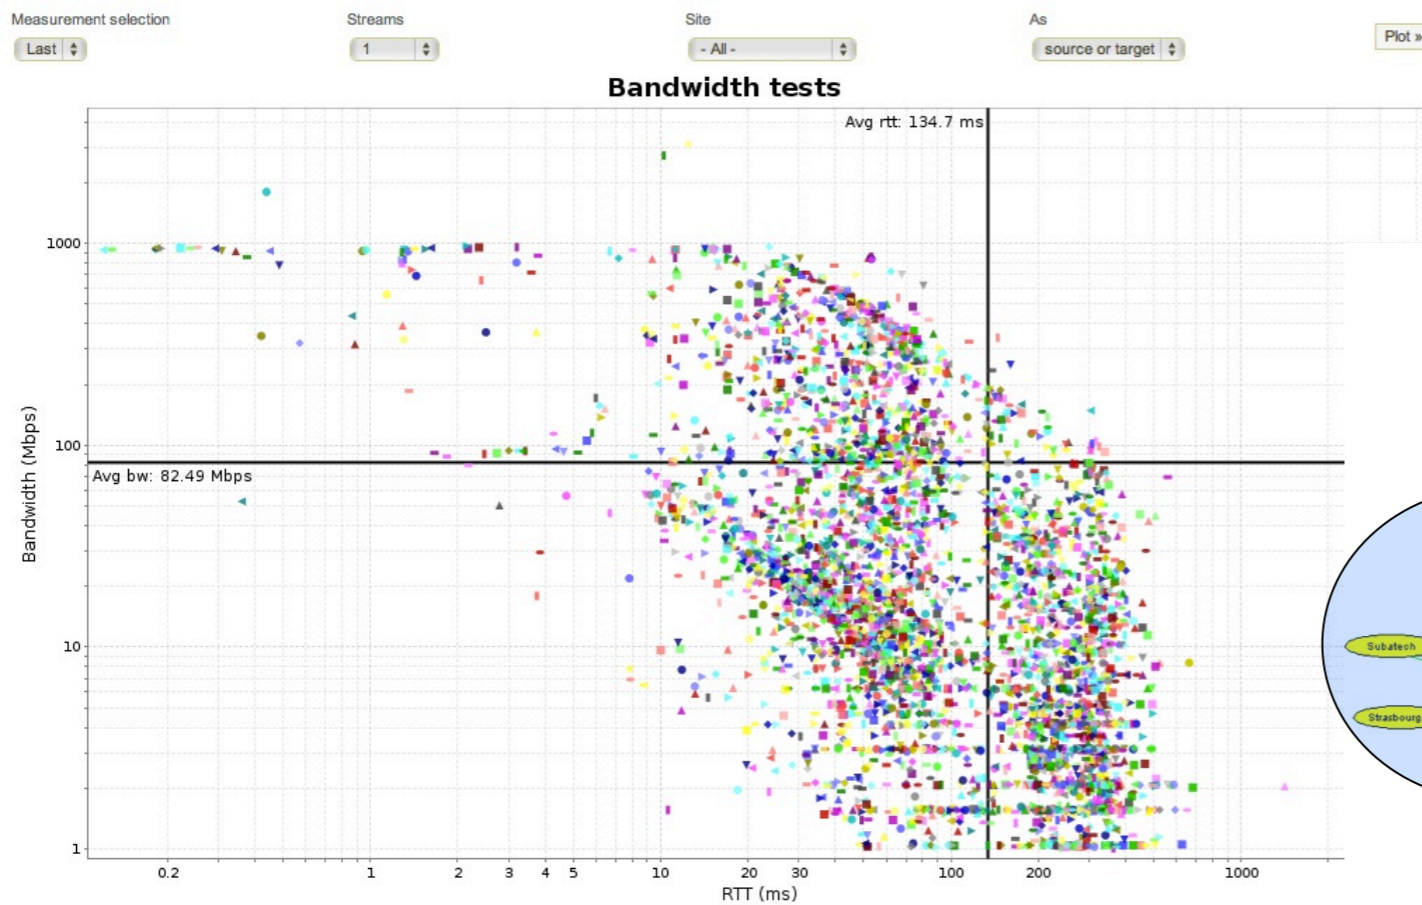
# Analyse facilitée

- Analyser des données devrait être facile pour un physicien
- Une série de petites (ou grosses) aides a été mise en place en ce sens

# Stocker les résultats d'analyse


- L'utilisateur ne choisit pas le(s) SE pour ses sorties, mais uniquement le nombre de répliques souhaitées
- Alien choisit seul le SE le «plus proche» pour chaque job, en fonction de :
  - topologie réseau collectée par MonALISA
  - tests fonctionnels des SE effectués en continu
  - taux d'occupation des SEs

# Topologie réseau



# Modes d'analyse

- Locale
- Proof
- Grille



Un seul code + une macro (steering) pour passer de l'un à l'autre  
= Analysis Framework's Alien plugin (fait partie de AliRoot)



# Alien plugin

```
AliAnalysisAlien plugin;
```

```
plugin.SetAliROOTVersion("v4-21-05-AN");
```

```
plugin.SetAnalysisSource("AliHistogramCollection.cxx AliAnalysisMuMu.cxx");
```

```
plugin.SetAdditionalLibs("libPWG3base.so AliHistogramCollection.h  
«AliHistogramCollection.cxx AliAnalysisMuMu.h AliAnalysisMuMu.cxx»");
```

```
plugin.SetMergeViaJDL(kTRUE);
```

```
plugin.SetGridWorkingDir(gridworkingdir);  
plugin.SetGridOutputDir("output");
```

```
plugin.SetProofCluster("nansafmaster.in2p3.fr");
```

```
plugin.SetProofDataSet("/alice/data/LHC10e_000128175_p1#esdTree");
```

version d'aliroot

Code utilisateur

Grille

Proof

# Suivi des jobs

Currently running jobs - ALICE Grid Monitoring with MonALISA

http://alimonitor.cern.ch/job\_stats.jsp

DDAI - Hébergement Read Later Apple Perso Subatech Alice ADSL Suba/EMN GRID MonAlisa CNRS -Porta...ébergement Livres/CD/DVD Air France Comm

OP Vistars RAW data production requests - ... Currently running jobs - ALICE G...

**ALICE** MonALISA Repository for ALICE

My jobs My home dir Catalogue browser Repository Home Administration Section ALICE Reports Events XML Feed Firefox Too

ALICE Repository

- ALICE Repository
  - Google Map
  - Shifter's dashboard
  - Run Condition Table
  - Production Info
  - Job Information
    - Site views
      - Summary plots
      - Job states
      - Jobs per site
      - Jobs per site table
      - JA LCGStatus
      - Resource usage
    - User views
      - Summary plots
      - Jobs status
      - Grid packages
      - Quotas
    - Task queue
      - Task queue summary
      - Jobs in TQ table
    - Job timings
    - Memory profiles
      - By site
      - Per user
      - Current jobs
    - SE Information
    - Services
    - Network Traffic
    - FTD Transfers
    - CAF Monitoring
    - SHUTTLE

**Jobs in TaskQueue**

pid	owner	first seen	last seen	subjobs	Job states										
					SPLIT	WAITING	STARTED	RUNNING	SAVING	DONE	ERRORS	ERROR_V	ERROR		
65893345	laphecet	16.11.2010 16:08	16.11.2010 23:04	40						100%	40				
64900884										100%	14				
64768331										100%	13				
64746676										100%	1				
64729979										100%	1				
64729930										100%	1				
64703226										100%	12				
63775735										100%	1				
63775707										100%	1				
63712803										90%	10	9%	1		
63775709										100%	1				
63775708										100%	1				
63712805										100%	3				
63712802										100%	12				
63712804										100%	4				
63126303										100%	8				
61942421										100%	1				
61942671										100%	1				
61942670										100%	1				
61942419										100%	1				
61942418										100%	1				
61861641										81%	13	18%	3		
61861642										92%	13	7%	1		
61861622										96%	30	3%	1		
61861621	laphecet	21.10.2010 22:46	22.10.2010 03:30	9						100%	9				
61861620	laphecet	21.10.2010 22:46	22.10.2010 02:07	20						100%	20				

Status of masterjob 65893345

Masterjob 65893345 of laphecet, status : DONE (refresh | JDL | less details | update status)  
 Automatically generated analysis JDL, //alice/cern.ch/user/l/laphecet//Analysis/JpsiPbPb/LHC10h//output/000137161,  
 ROOT v5-27-06b, AllROOT v4-21-05-AN  
 Subjobs: 40

DONE (40)

```

65893346 : trace | trace all | log files (16m 2s running, 1m 4s saving @ ALICE::GRIF_IPNO::LCG)
65893347 : trace | trace all | log files (13m 15s running, 1m 22s saving @ ALICE::GRIF_IPNO::LCG)
65893348 : trace | trace all | log files (16m 7s running, 1m 59s saving @ ALICE::GRIF_IPNO::LCG)
65893349 : trace | trace all | log files (14m 10s running, 0m 35s saving @ ALICE::GRIF_IPNO::LCG)
65893350 : trace | trace all | log files (17m 55s running, 1m 40s saving @ ALICE::GRIF_IPNO::LCG)
65893351 : trace | trace all | log files (10m 58s running, 0m 50s saving @ ALICE::GRIF_IPNO::LCG)
65893352 : trace | trace all | log files (12m 36s running, 1m saving @ ALICE::GRIF_IPNO::LCG)
65893353 : trace | trace all | log files (11m 46s running, 1m 16s saving @ ALICE::GRIF_IPNO::LCG)
65893354 : trace | trace all | log files (0m 35s running, 0m 48s saving @ ALICE::GRIF_IPNO::LCG)
65893355 : trace | trace all | log files (0m 26s running, 0m 32s saving @ ALICE::GSI::LSF2)
65893356 : trace | trace all | log files (8m 8s running, 1m 27s saving @ ALICE::GSI::LSF2)
65893357 : trace | trace all | log files (8m 2s running, 0m 43s saving @ ALICE::GSI::LSF2)
65893358 : trace | trace all | log files (13m 56s running, 2m 4s saving @ ALICE::GRIF_IPNO::LCG)
65893359 : trace | trace all | log files (15m 55s running, 2m 5s saving @ ALICE::GRIF_IPNO::LCG)
65893360 : trace | trace all | log files (7m 52s running, 0m 43s saving @ ALICE::GSI::LSF2)
65893361 : trace | trace all | log files (7m 54s running, 0m 25s saving @ ALICE::GSI::LSF2)
65893362 : trace | trace all | log files (8m 14s running, 0m 42s saving @ ALICE::GSI::LSF2)
65893363 : trace | trace all | log files (8m 36s running, 0m 43s saving @ ALICE::GSI::LSF2)
65893364 : trace | trace all | log files (8m 2s running, 1m 23s saving @ ALICE::GSI::LSF2)
65893365 : trace | trace all | log files (1m running, 1m 14s saving @ ALICE::GSI::LSF2)
65893366 : trace | trace all | log files (8m 8s running, 0m 38s saving @ ALICE::GSI::LSF2)
65893367 : trace | trace all | log files (11m 18s running, 0m 52s saving @ ALICE::GRIF_IPNO::LCG)
65893368 : trace | trace all | log files (8m 3s running, 1m 21s saving @ ALICE::GSI::LSF2)
65893369 : trace | trace all | log files (8m 7s running, 0m 45s saving @ ALICE::GSI::LSF2)
65893370 : trace | trace all | log files (7m 52s running, 1m 14s saving @ ALICE::GSI::LSF2)
65893371 : trace | trace all | log files (8m 19s running, 1m 9s saving @ ALICE::GSI::LSF2)
    
```

# aaf.cern.ch

- Alice Analysis Facility
- Un «packaging» de ferme Proof prenant en compte :
  - gestion des datasets (e.g. staging des données)
  - gestion des packages (les mêmes que sur la AliEn)
  - monitoring

# AAF dans le monde

## ALICE PROOF Clusters

What is this abo

### Cluster list

Name	Online	Status	Cluster			ROOT	Aggregated disk space			AF xrootd		xrootd
			Proof master	Workers	Users	Version	Total	Free	Used	Running	Latest	Version
1. CAF	Green	Stable	alice-caf.cern.ch	52	2	v5-27-06b	80.86 TB	19.3 TB	61.57 TB	1.0.35	1.0.35	20100510-1509_d
2. HAF	Green	Maintenance sin...	proofmaster.internal	64	0	v5-26-00b-6	91.05 GB	11 GB	80.06 GB	1.0.17	1.0.17	20100510-1509_d
3. JRAF	Red			-	-		-	-	-			
4. KAF	Green	Stable	afmaster01.sdfarm.kr	132	0	v5-26-00-proof	1.057 TB	80.1 GB	1002 GB	1.0.22	1.0.22	20100510-1509_d
5. LAF	Red			-	-		-	-	-			
6. SAF	Green	Maintenance sin...	nansafmaster.in2p3.fr	48	0	v5-27-06b	12.07 TB	7.68 TB	4.392 TB	1.0.35	1.0.35	20100510-1509_d
7. SKAF	Green	Stable	skaf.saske.sk	30	0	v5-27-06b	53.72 TB	38.2 TB	15.52 TB	1.0.35	1.0.35	20100510-1509_d
8. SKAF_TEST	Green			-	-		815.9 GB	677 GB	138.9 GB			20100510-1509_d
<b>Total</b>				<b>326</b>	<b>2</b>		<b>148.6 TB</b>	<b>65.93 TB</b>	<b>82.67 TB</b>			

En production : CAF et SKAF  
Bientôt : KAF (Corée), SAF & LAF (?)

# Packages

- Synchronisés avec les packages utilisés par AliEn
- si le code marche sur la grille, il marche sur une AAF (\*)

# Datasets

- Une seule «syntaxe» = alien path, e.g. alien:///alice/data/2010/LHC10h/.../AliESDs.root
- Trois types de datasets
  - static
    - données sont copiées sur le cluster Proof avant tout
  - dynamic
    - données sont copiées à la première analyse
  - alien
    - données accédées directement depuis la grille

# Performances ?

The screenshot displays a Mac OS X desktop environment. At the top, the menu bar shows 'X11 Applications Édition Fenêtre Aide' and system status icons including network, Wi-Fi, and battery. The main workspace contains several windows:

- Terminal Window (SAF):** Shows a series of log messages: 'Retrieving logs: 34 ok, 0 not ok (0 % processed)' through 'Retrieving logs: 49 ok, 0 not ok (100 % processed)'. Below this, it displays a warning about parsing lines, a message 'Mst-0: Number of mergers set dynamically to 7 (for 48 workers)', and a 'grand total: sent 12 objects, size: 43916675 bytes'. It also shows 'I-AliAnalysisMuMuFromESD::Terminate: size after prune histograms = 41.7 MB' and 'AliHistogramCollection : 109 keys and 65 histos'. The path '/ANY/CMBAC-B-NOPF-ALL/137544/N0/ 289605' and 'root [1]' are also visible.
- PROOF Query Progress Dialog:** A modal window titled 'PROOF Query Progress: laphecet@localhost'. It reports: 'Executing on PROOF cluster "localhost" with 48 parallel workers: Selector: AliAnalysisSelector 2616 files, number of events 761514, starting event 0'. A green progress bar is shown. A purple circle highlights the 'Processing time: 12 min 24 sec' and 'Processed: 761514 events (309.54 GB)'. A speedometer on the right shows 'Ev/s' with a needle at approximately 100. Other metrics include 'Initialization time: 4.7 secs' and 'Processing rate: 1022.3 evts/sec (425.5 MB/sec)'. Buttons for 'Show Logs', 'Performance plot', 'Memory Plot', 'Run in background', 'Stop', 'Cancel', and 'Close' are present.
- Code Editor:** Shows C++ code for mask processing:

```
DecodePairCutMask(str->GetUniqueID(), singleTrackMask, pairMask);  
  
Bool_t testi = ( ( maski & singleTrackMask ) == singleTrackMask );  
Bool_t testj = ( ( maskj & singleTrackMask ) == singleTrackMask );  
Bool_t testij(kTRUE);  
  
if ( pairMask > 0 )  
{  
    testij = ( (maski & pairMask) == pairMask ) &&  
            ( (maskj & pairMask) == pairMask );  
}
```

The dock at the bottom contains various application icons, and the desktop background features a grid of icons including 'partition alice', 'usage t2 français', and 'QM MUON fine week'.

# Quelques problèmes cependant

- Sur la grille comme sur Proof
  - Fuites de mémoires
  - Taille (et/ou type) de sortie de l'analyse (histos vs tree par exemple) peut poser problème lors de la phase de merging



# Solutions

- Fuites de mémoire
  - «éduquer» les codeurs à debugger leurs analyses avant de lancer sur proof ou grille
    - outils à disposition : valgrind, Coverity, etc...
  - protéger les machines en tuant les processus déviants...
- Merging
  - bypass complet
    - e.g. sur proof, enregistrer chaque morceau comme dataset qui sera l'input de l'analyse suivante, jusqu'à ce que l'output soit «mergeable»
  - merging plus intelligent ? (pas possible dans tous les cas de figure)

# Organisation nationale ?

- A l'heure actuelle, pas vraiment...
- En théorie, deux AAF en France :
  - SAF
    - «vraie» AAF, commence à produire
  - LAF
    - marche depuis un certain temps, mais n'est pas une AAF...
- Tutorial pour LAF organisé par R. Vernet début décembre au CC.