



Alice Computing



Federico Carminati
The Grid of the Americas
February 8-11, 2011




WHY HI COLLISIONS?


- Study QCD at its natural energy scale $T = \Lambda_{\text{QCD}} = 200$ MeV by creating a state of matter at high density and temperature using high energetic heavy ion collisions.




NA49



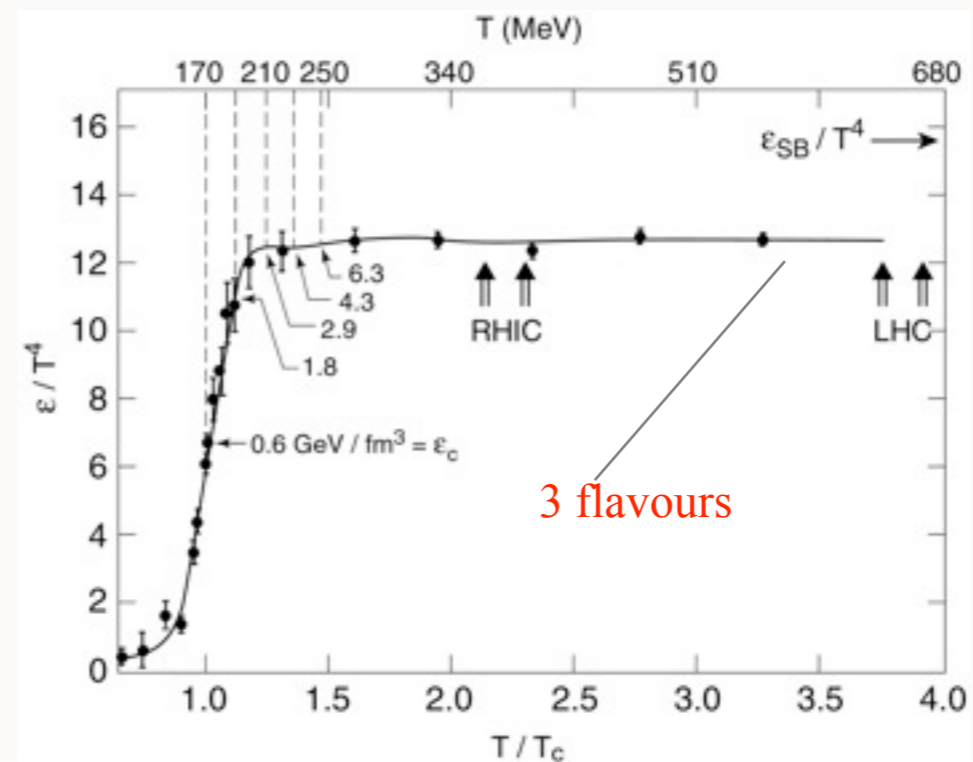
STAR



ALICE



- Indication of trans. HG to QGP at $T_c \cong 170$ MeV $\epsilon_c \cong 1$ GeV / fm³
- Phase trans. or crossover?
- Intermediate phase of strongly interacting QGP?
- Chiral symmetry restoration ?
- Constituent mass \rightarrow current mass



HISTORY OF HIGH-ENERGY A+B BEAMS

- **BNL-AGS: mid 80's, early 90's**

O+A, Si+A 15 AGeV/c $\sqrt{s_{NN}} \sim 6 \text{ GeV}$

Au+Au 11 AGeV/c $\sqrt{s_{NN}} \sim 5 \text{ GeV}$

- **CERN-SPS: mid 80's, 90's**

O+A, S+A 200 AGeV/c $\sqrt{s_{NN}} \sim 20 \text{ GeV}$

Pb+A 160 AGeV/c $\sqrt{s_{NN}} \sim 17 \text{ GeV}$


- **BNL-RHIC: early 00's**

Au+Au $\sqrt{s_{NN}} \sim 130 \text{ GeV}$

p+p, d+Au $\sqrt{s_{NN}} \sim 200 \text{ GeV}$

- **LHC: 2010 (!)**

Pb+Pb $\sqrt{s_{NN}} \sim 5,500 \text{ (2,750 in '10-'12) GeV}$



ALICE Collaboration
 ~ 1/2 ATLAS, CMS, ~ 2x LHCb
 ~1000 people, 30 countries,
 ~ 80 Institutes

Total weight	10,000t
Overall diameter	16.00m
Overall length	25m
Magnetic Field	0.5Tesla

8 kHz (160 GB / sec)

level 0 - special hardware

200 Hz (4 GB / sec)

level 1 - embedded processors

30 Hz (2.5 GB / sec)

level 2 - PCs

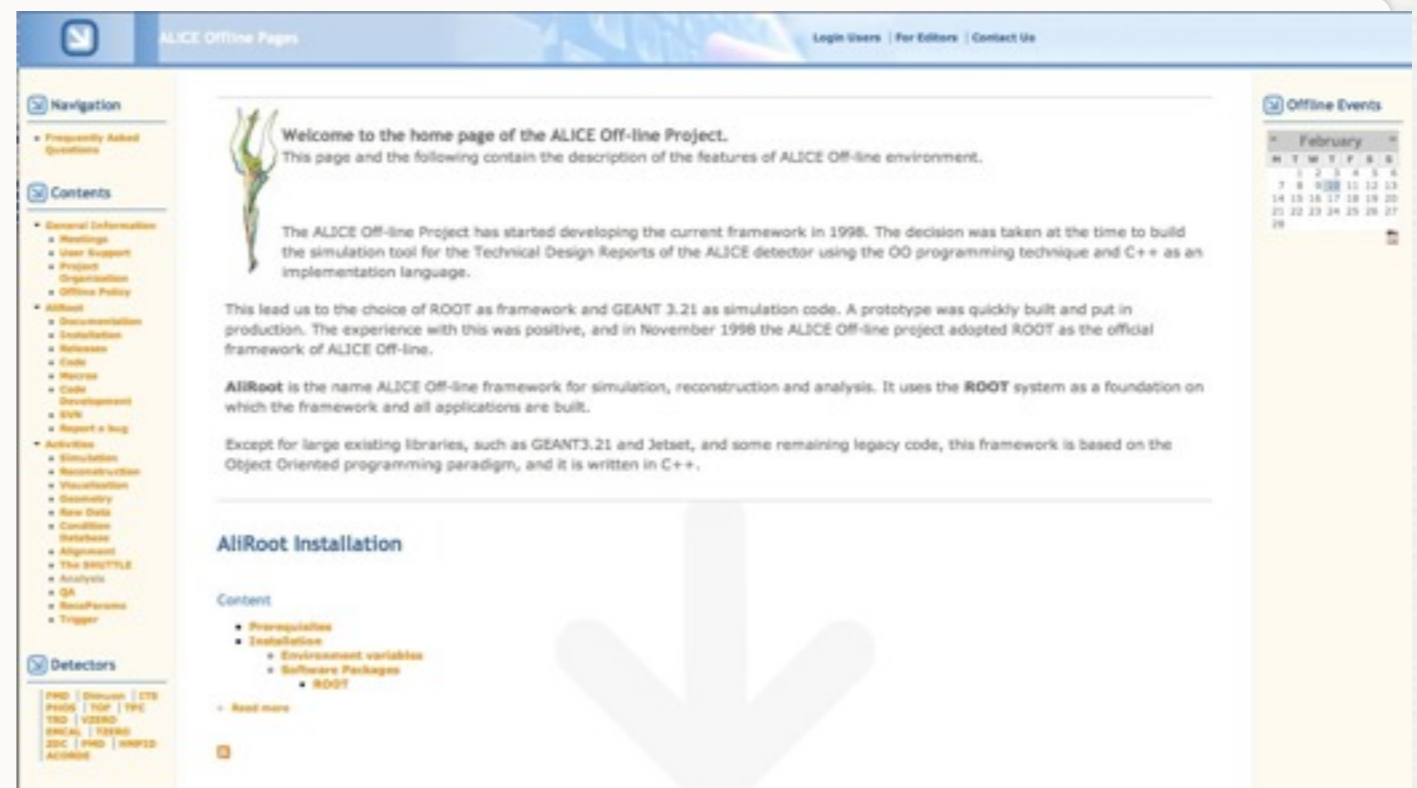
30 Hz
(1.25 GB / sec)

data recording &
offline analysis

A full pp programme
 Data rate for pp is 100Hz@1MB

ORGANISATION

- Core Offline is CERN responsibility
 - Framework development
 - Coordination activities
 - Documentation
 - Integration
 - Testing & release
 - Resource planning
- Each subdetector is responsible for its own offline system
 - It must comply with the general ALICE Computing Policy as defined by the Computing Board
 - It must integrate into the AliRoot framework



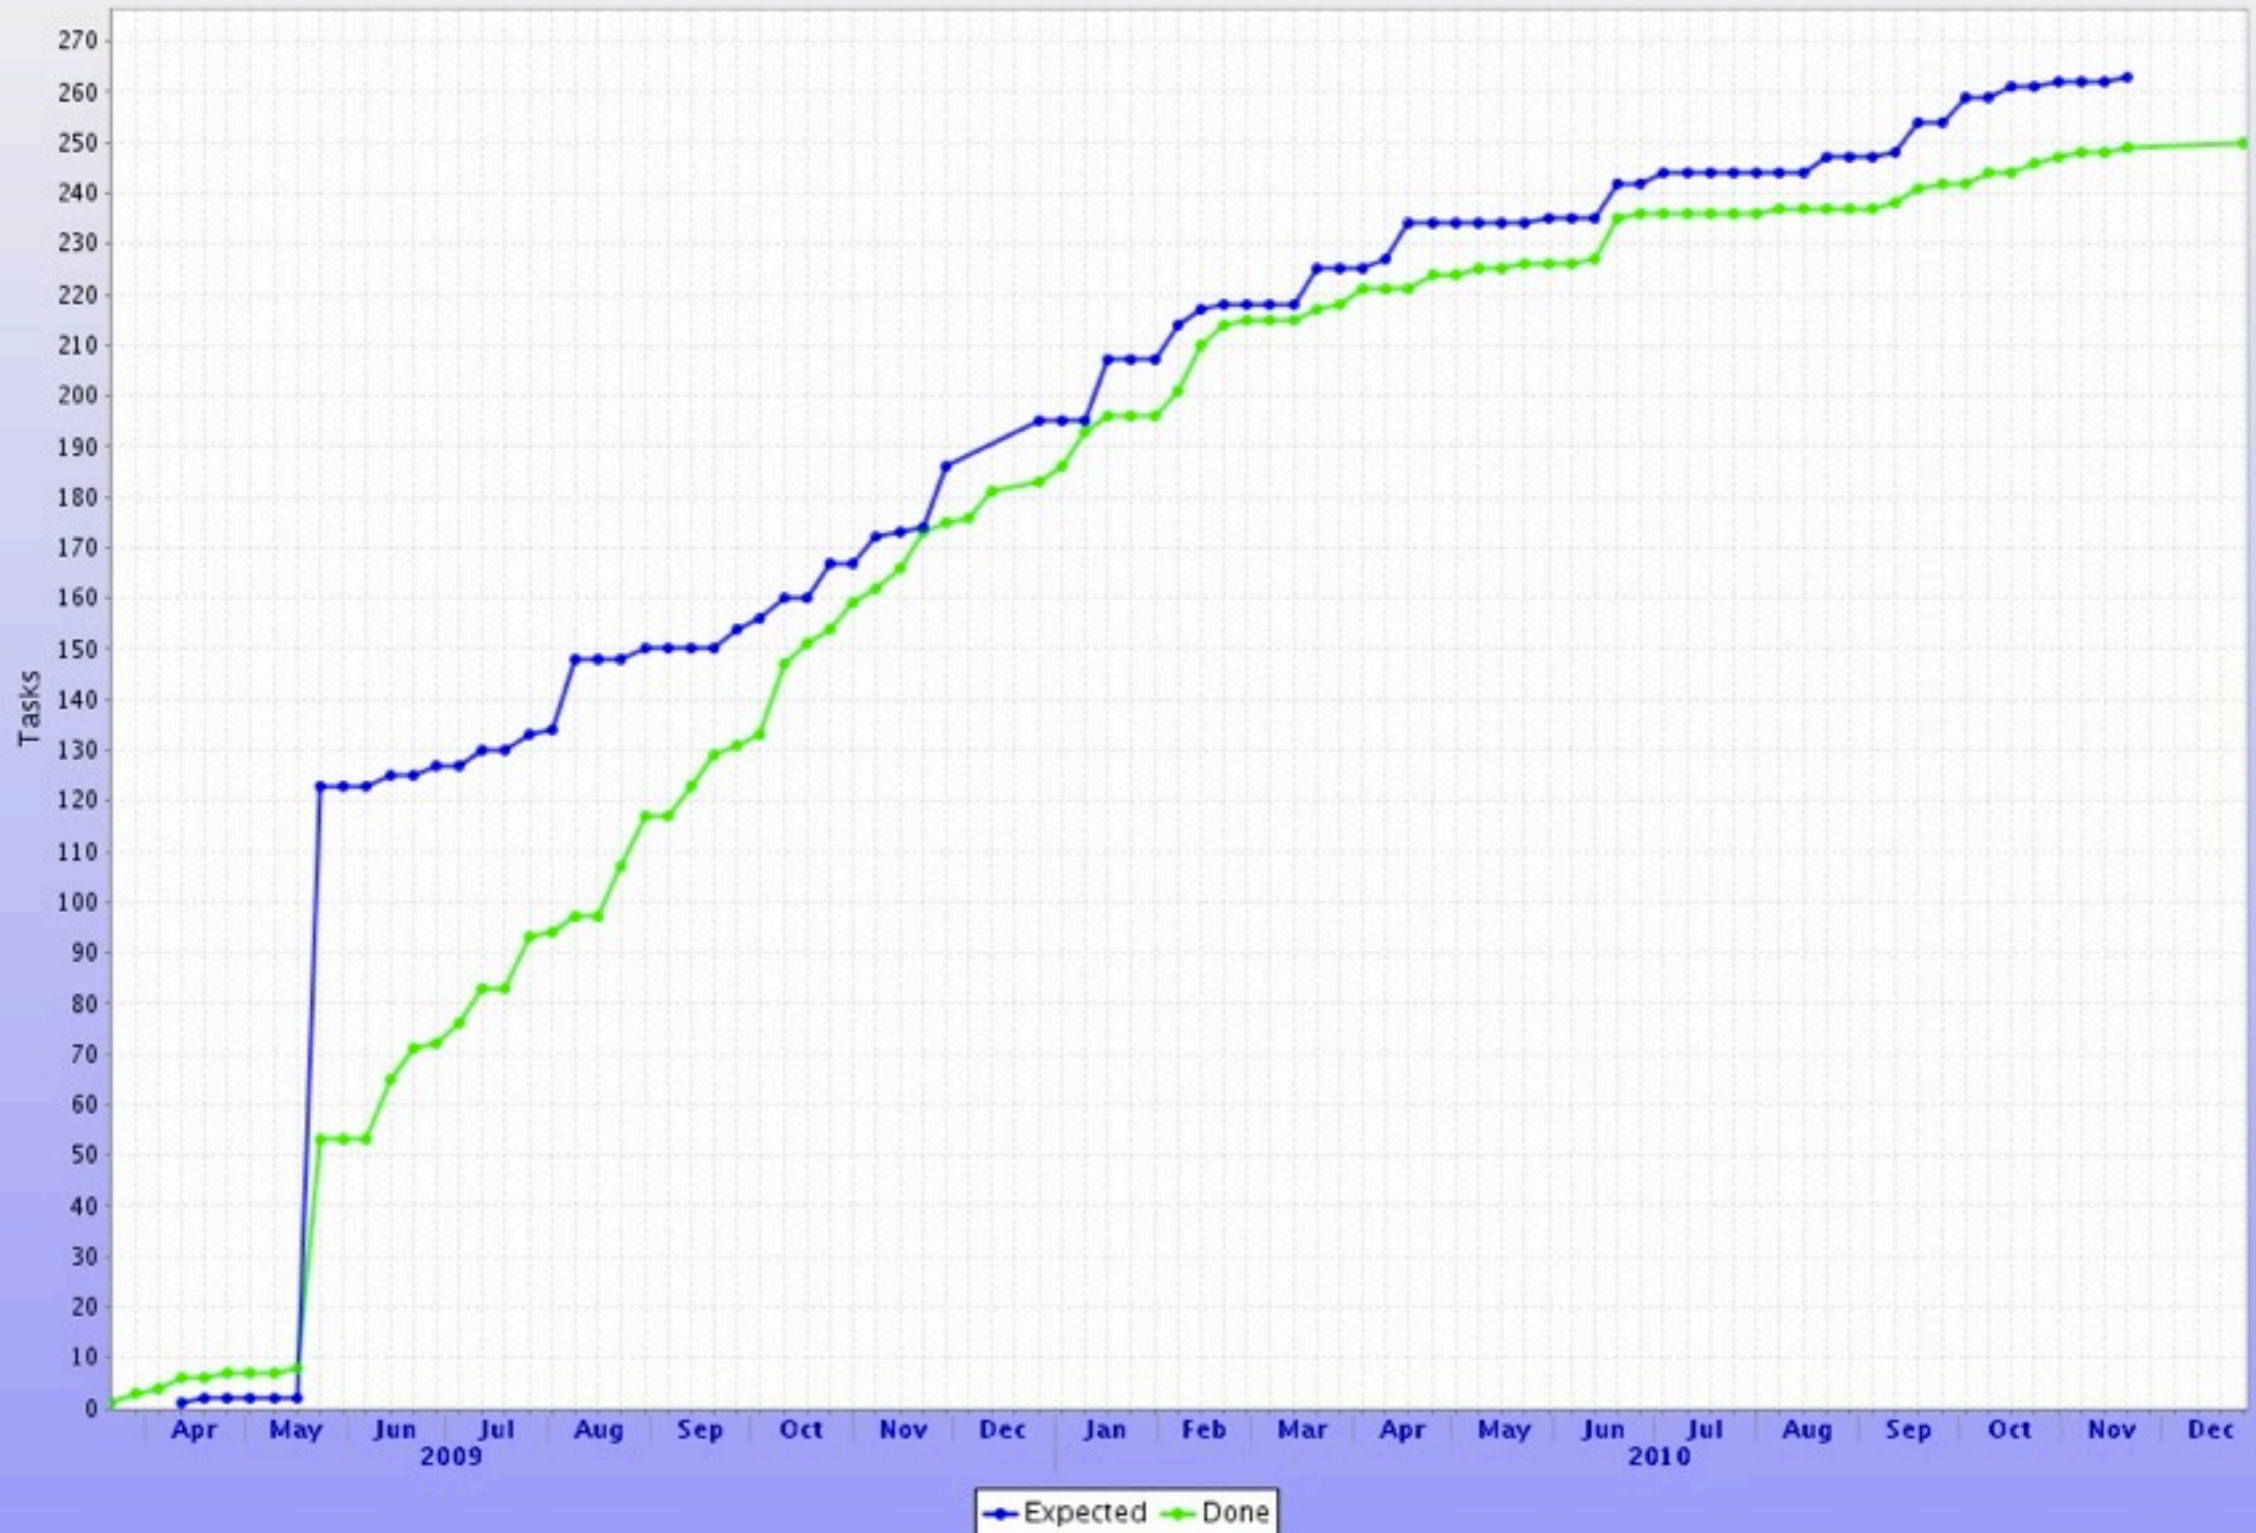
<http://aliweb.cern.ch/Offline/>

PLANNING

IN PREPARING FOR BATTLE I
ALWAYS FOUND PLANS
USELESS BUT PLANNING
ESSENTIAL
GEN D.EISENHAUER

PLANNING

Total tasks – Week View



PLANNING

Total tasks - Week View

270



Previous Results

157 matching items - Items 1 to 50

Next Results



Item ID	Summary	Submitted On	Assigned To	Submitted By
#78018	Change needed in 4.20 release for AliZDCReconstructor.cxx	2011-02-10 11:36	hristov	coppedis
#78008	port in alien OCDB T0 calibration file with fixed bug	2011-02-10 09:37	rgrosso	alla
#77887	CCUP2-B-NOPF-ALLNOTRD does not appear in ESD	2011-02-07 15:54	lietava	kskjerda
#77757	Issues with CMake and detector algorithms	2011-02-03 16:22	hristov	hristov
#77629	PACKLDFLAGS of last package of a module overrides all previous	2011-01-31 21:58	hristov	richterm
#77628	Project internal library dependencies not correctly propagated from ELIBS	2011-01-31 21:44	hristov	richterm
#77607	AliSimulation crash when using AliGenCocktailAfterBurner	2011-01-31 15:59	akisiel	ngutierr
#77591	Need a single factfile for cmake	2011-01-31 11:01	nyastreb	fca
#77590	cmake is linking directly the libraries instead of using -l...	2011-01-31 10:57	hristov	fca
#77589	Setting optimisation with cmake	2011-01-31 10:55	hristov	fca
#77497	Provide svn post commit action running root macros	2011-01-27 21:38	hristov	morsch
#77495	Increase tracking eta acceptance in pass2 of LHC10h	2011-01-27 21:30	miranov	morsch
#77478	VZERO: changes in the reco of Pb-Pb data	2011-01-27 13:34	hristov	cheshkov
#77435	CDB objects affecting geometry	2011-01-26 14:36	kowal2	rgrosso
#77401	AliESDCentrality must be reinstated - older PbPb data unreadable	2011-01-25 23:59	morsch	cholm
#77400	Bad change in AliESDCaloTrigger	2011-01-25 23:41	guemane	cholm
#77350	Incorrect geometry in RAW OCDB	2011-01-25 08:10	masera	lbetev
#77328	GRP objects for the RAW OCDB 2011	2011-01-24 14:34	rgrosso	rgrosso
#77326	different definition of primary particles (gammas from pi0 decays) in pp and PbPb	2011-01-24 13:46	hristov	amarin
#77321	request to store friends for TRD-triggered events	2011-01-24 12:41	rbailhac	misko
#77288	SIGFPE with hijing fastgen	2011-01-21 23:15	morsch	pchrist
#77281	Update of AliESDpid::NumberOfSigmasTOF method	2011-01-21 18:52	belikov	decaro
#77189	Needed calibration types for EMCAL in RAW OCDB 2011	2011-01-19 17:57	dsilverm	rgrosso
#77166	Consider installing scripts, data, etc.	2011-01-19 13:40	hristov	cholm
#77082	What to do with scripts loading other scripts	2011-01-17 14:53	fca	cholm
#77078	Make SVN ignore generated CMakeLists.txt	2011-01-17 14:08	hristov	cholm
#76954	Add lhpdf5.5.1/src/binreloc.c to CSRCS	2011-01-13 09:28	fca	cholm
#76933	Fix access operator	2011-01-12 14:44	matevz	cholm
#76700	Circular dependency STEER and CDB caused by AliMathBase	2011-01-03 11:28	hristov	richterm
#76531	Segmentation fault from VZERO reconstruction	2010-12-17 16:06	cheshkov	ivana
#76462	Addition of the V0 track parameters to AOD	2010-12-15 17:34	amarin	amarin
#76387	TOF calibration in preparation of pass2 of LHC10e-g	2010-12-13 16:11	decaro	schutz

(157 open items, 2703 total)


Expected Done

RESOURCES

- Sore point for ALICE computing

		2011 (RRB year)				2012 (RRB year)				2013 (RRB year)			
		T0	CAF	T1	T2	T0	CAF	T1	T2	T0	CAF	T1	T2
CPU (KHEP06)	Requested	81.8	22.5	144.0	128.8	87.4	13.8	194.4	138.3	78.5	17.3	132.3	140.6
	Pledged	48.3	13.7	71.5	107.1	51.0	18.2	74.9	117.0	-	-	-	-
	Missing	-41%	-39%	-50%	-17%	-42%	32%	-61%	-15%	-	-	-	-
Disk (PB)	Requested	6.8	0.5	8.7	8.6	9.6	1.3	13.3	11.2	11.2	1.6	15.0	12.7
	Pledged	5.5	0.5	5.5	8.0	6.6	0.6	6.5	11.1				
	Missing	-19%	-17%	-36%	-7%	-32%	-55%	-51%	-2%	-	-	-	-
MSS (PB)	Requested	9.7		28.7		14.1		46.6		16.7	-	61.5	
	Pledged	6.8		8.0		7.8		11.1					
	Missing	-30%	-	-72%	-	-44%	-	-76%	-	-	-	-	-

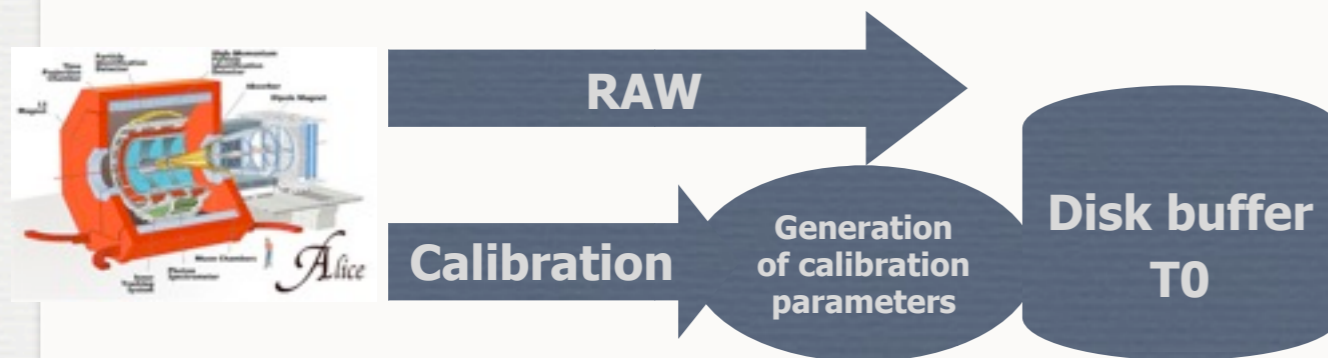
COMPUTING MODEL – PP



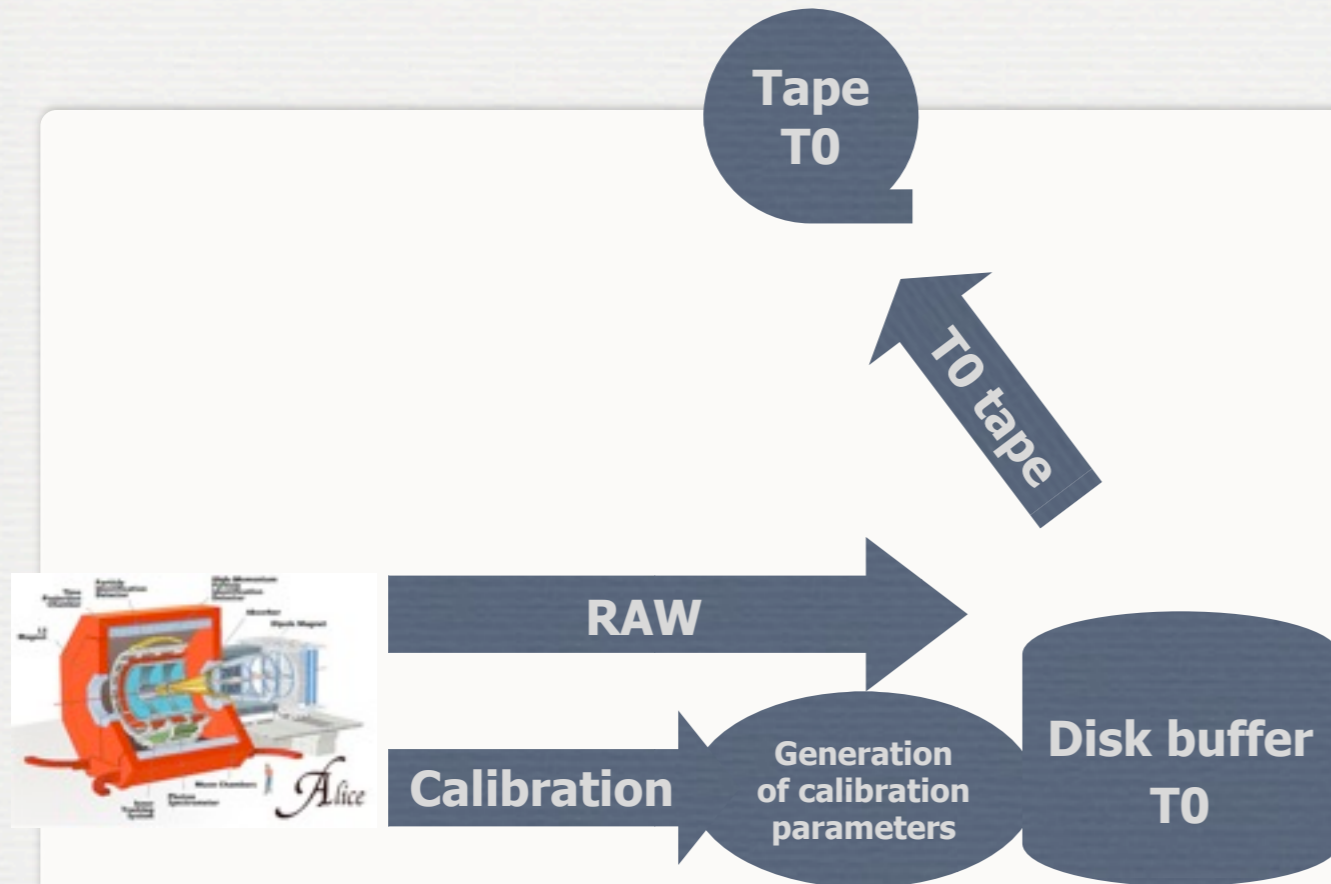
COMPUTING MODEL – PP



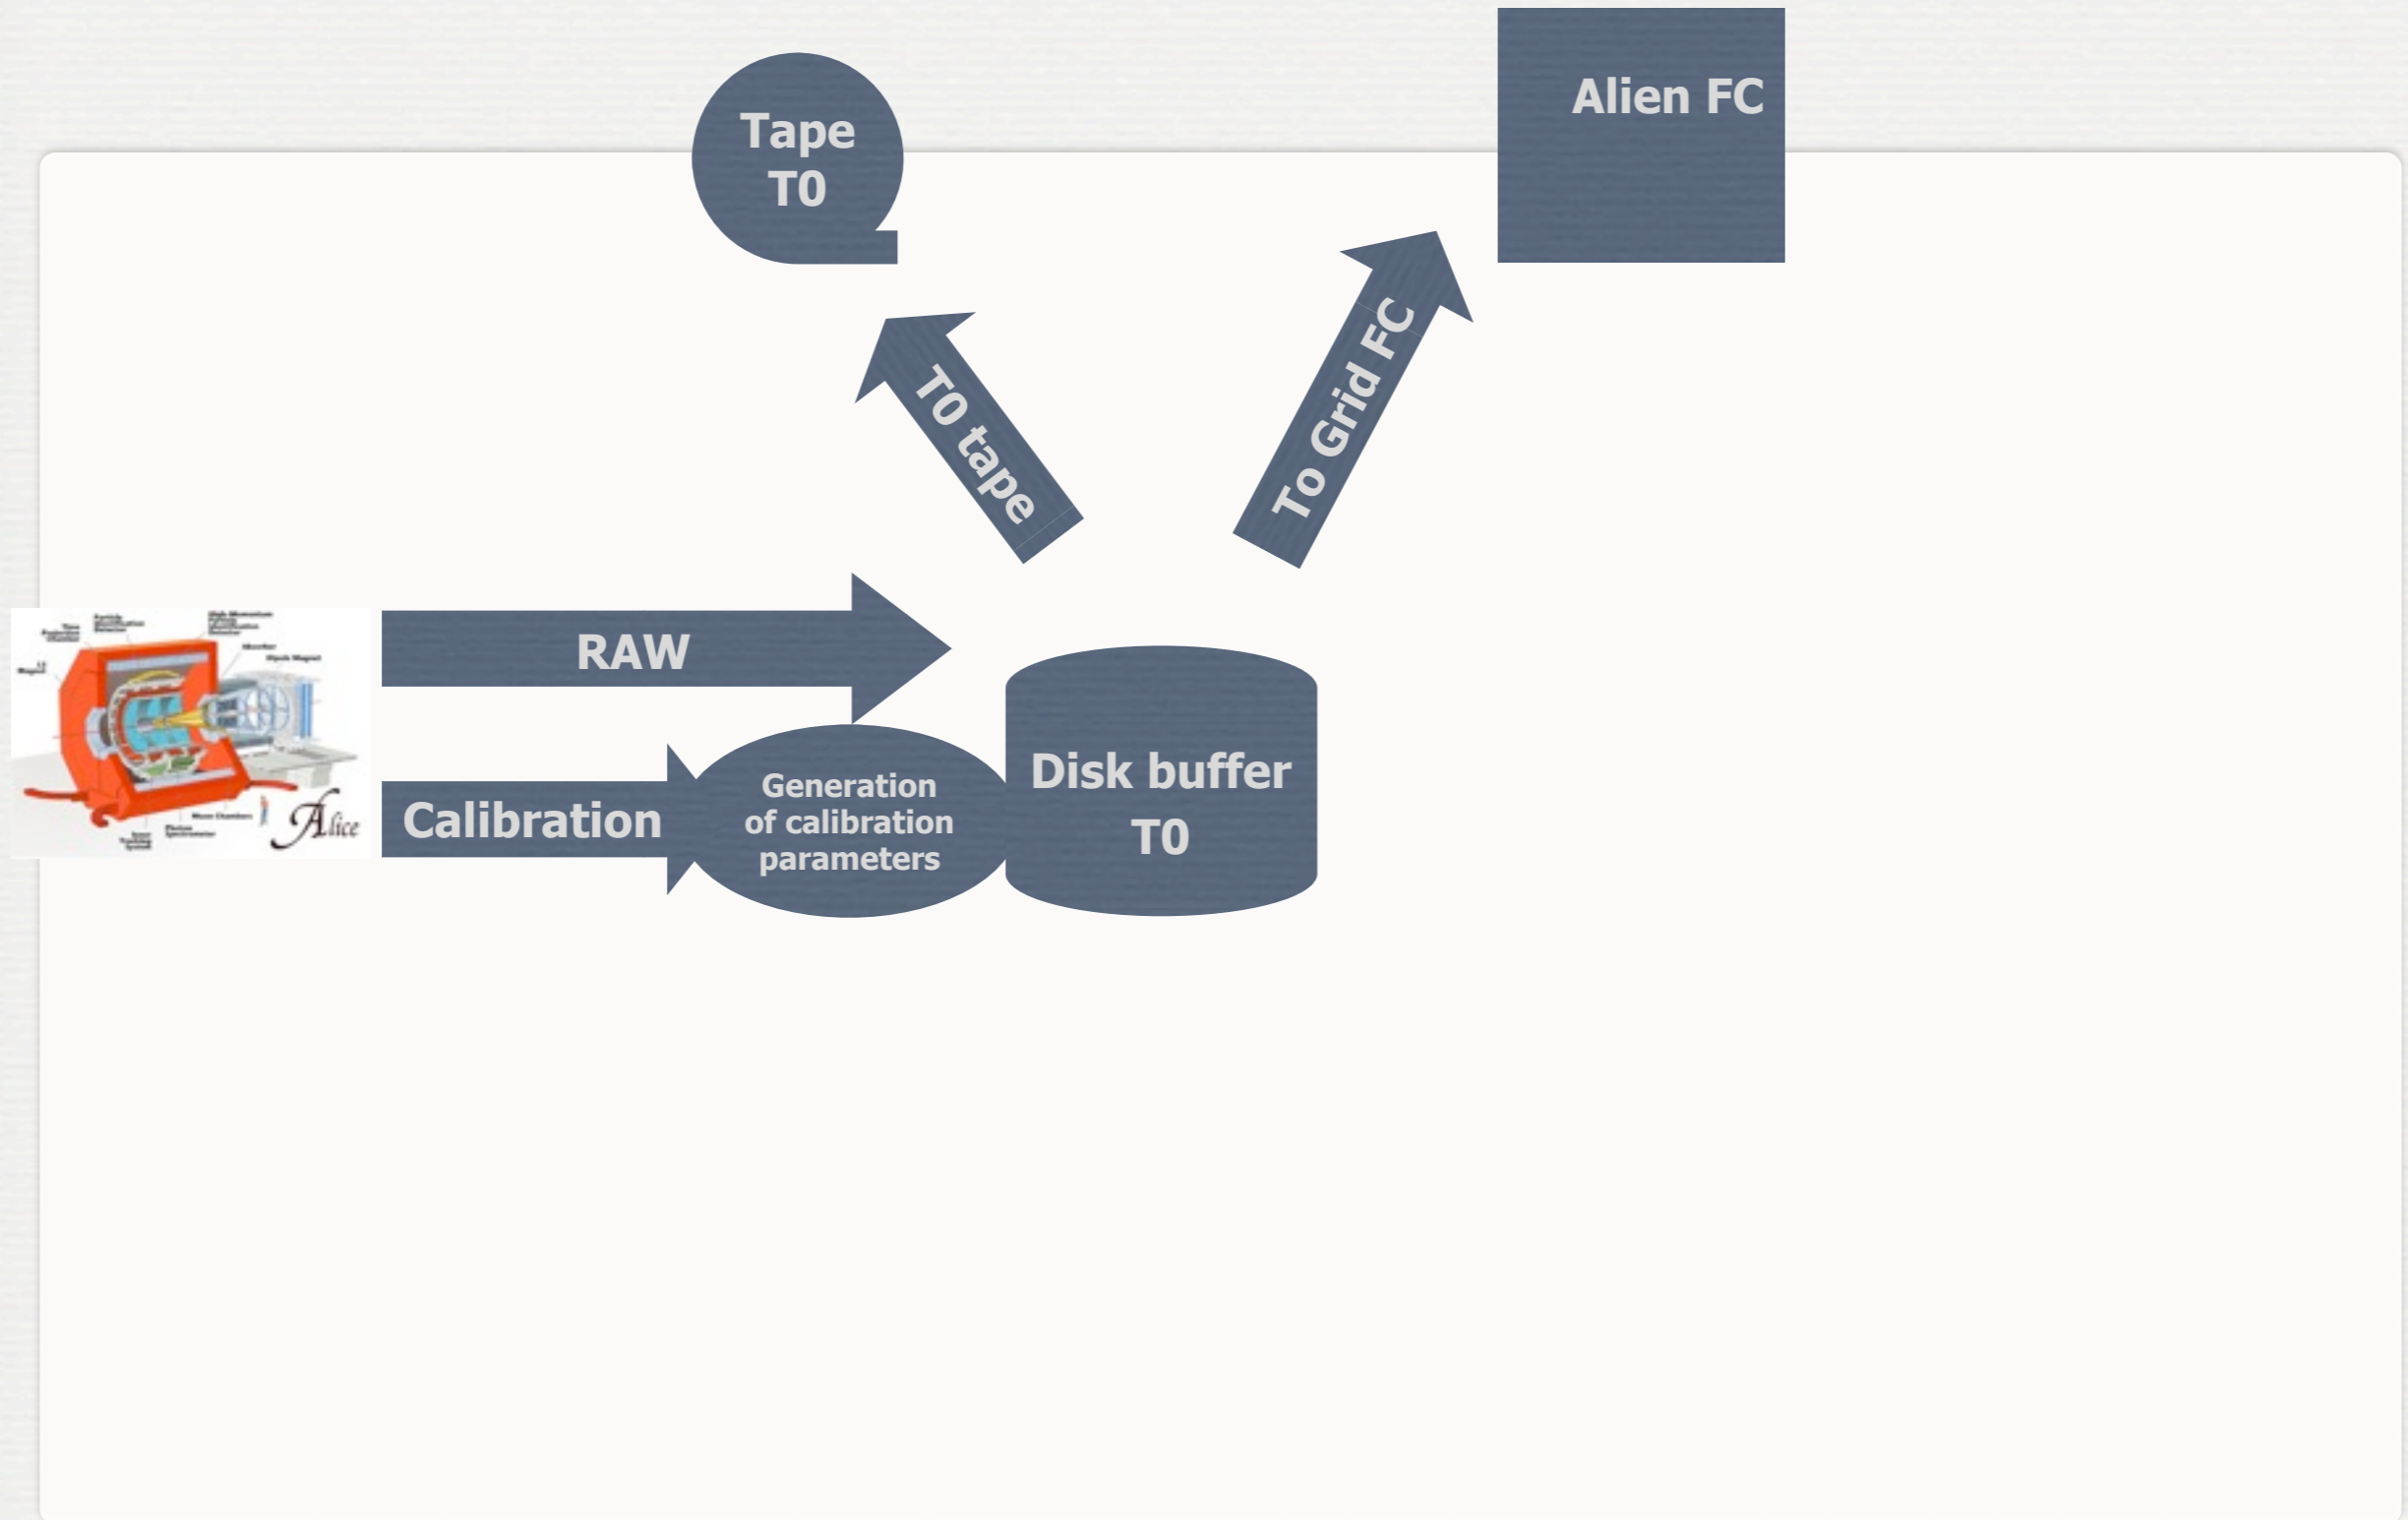
COMPUTING MODEL – PP



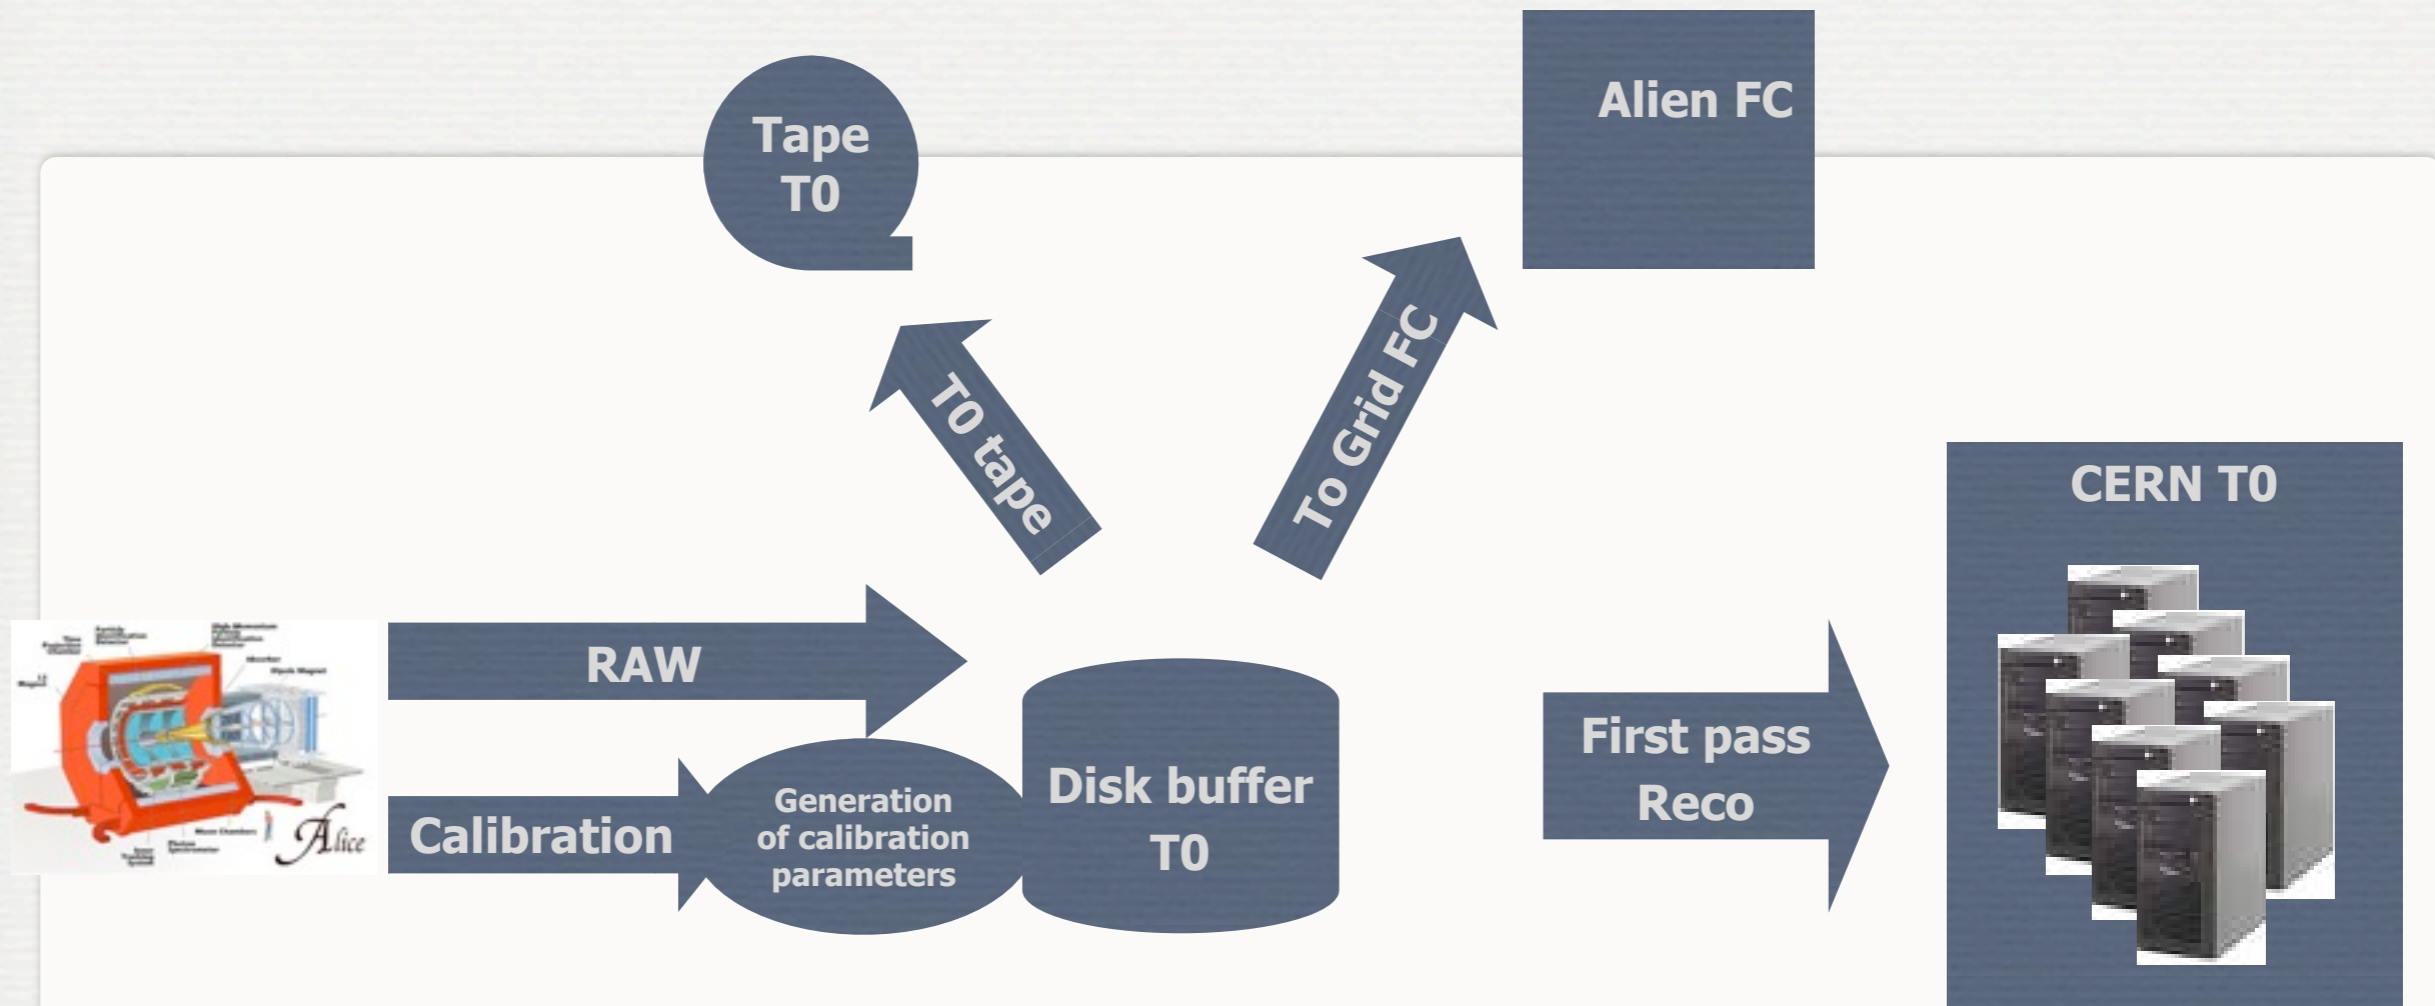
COMPUTING MODEL – PP



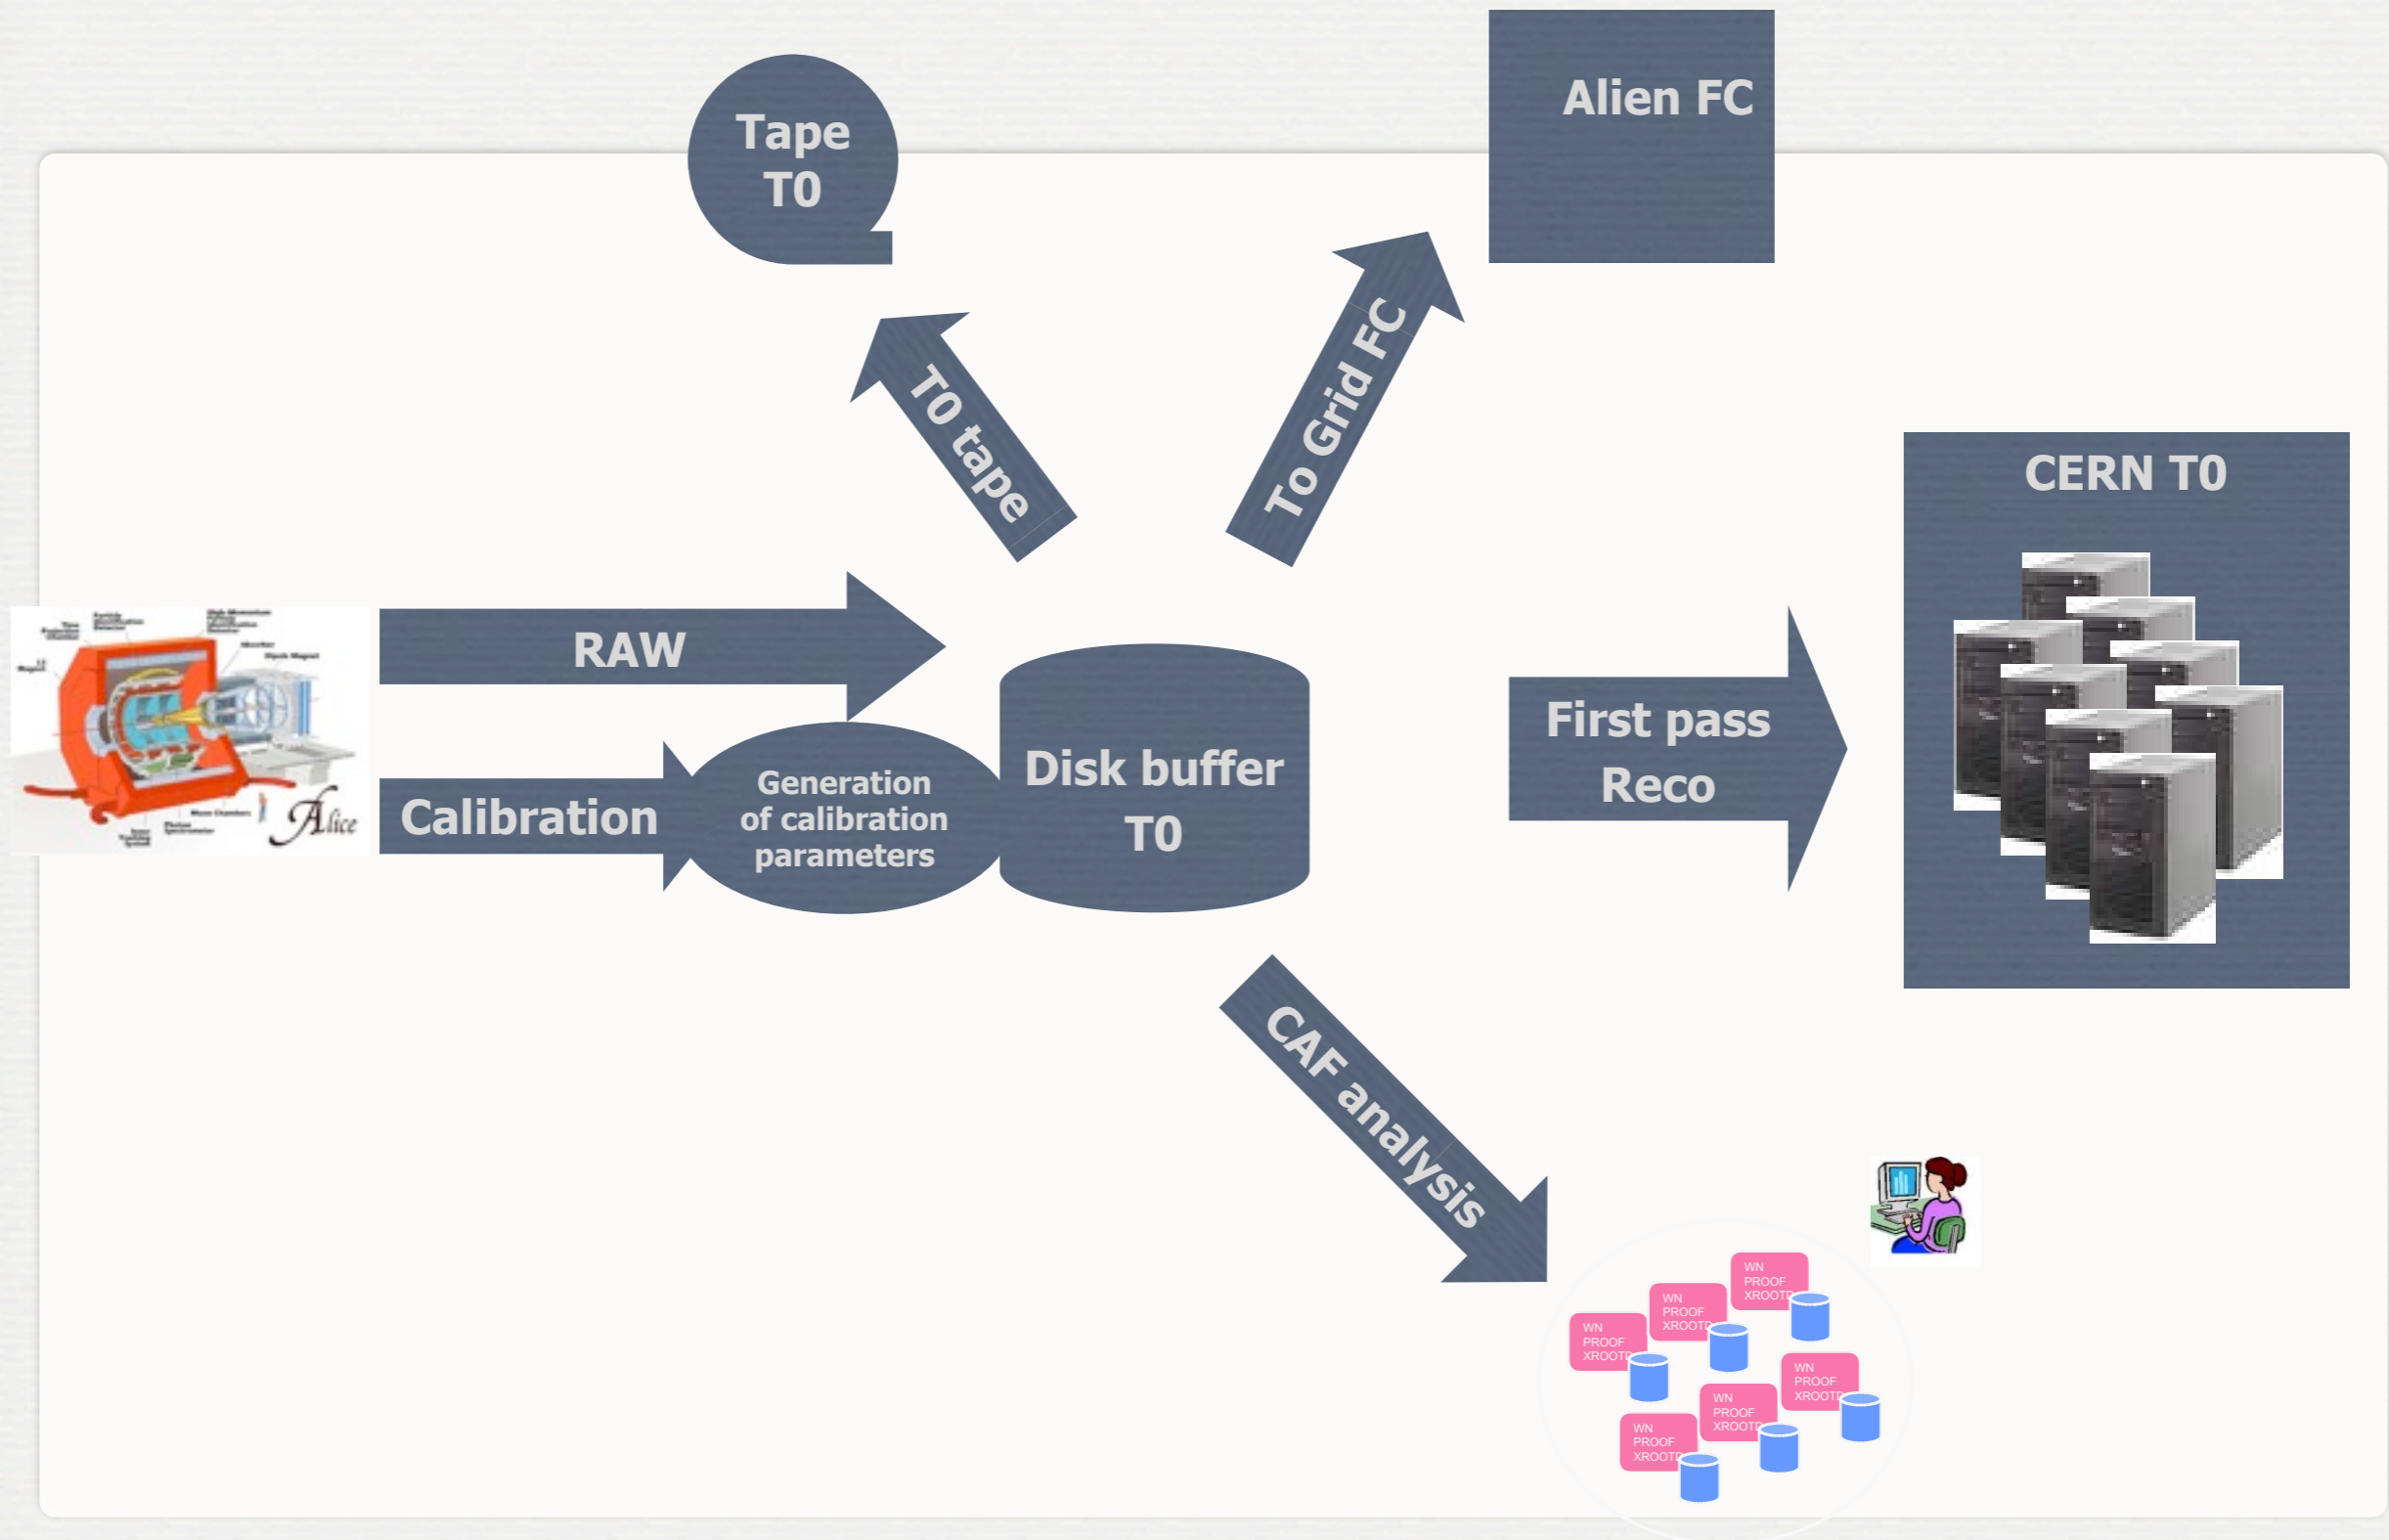
COMPUTING MODEL – PP



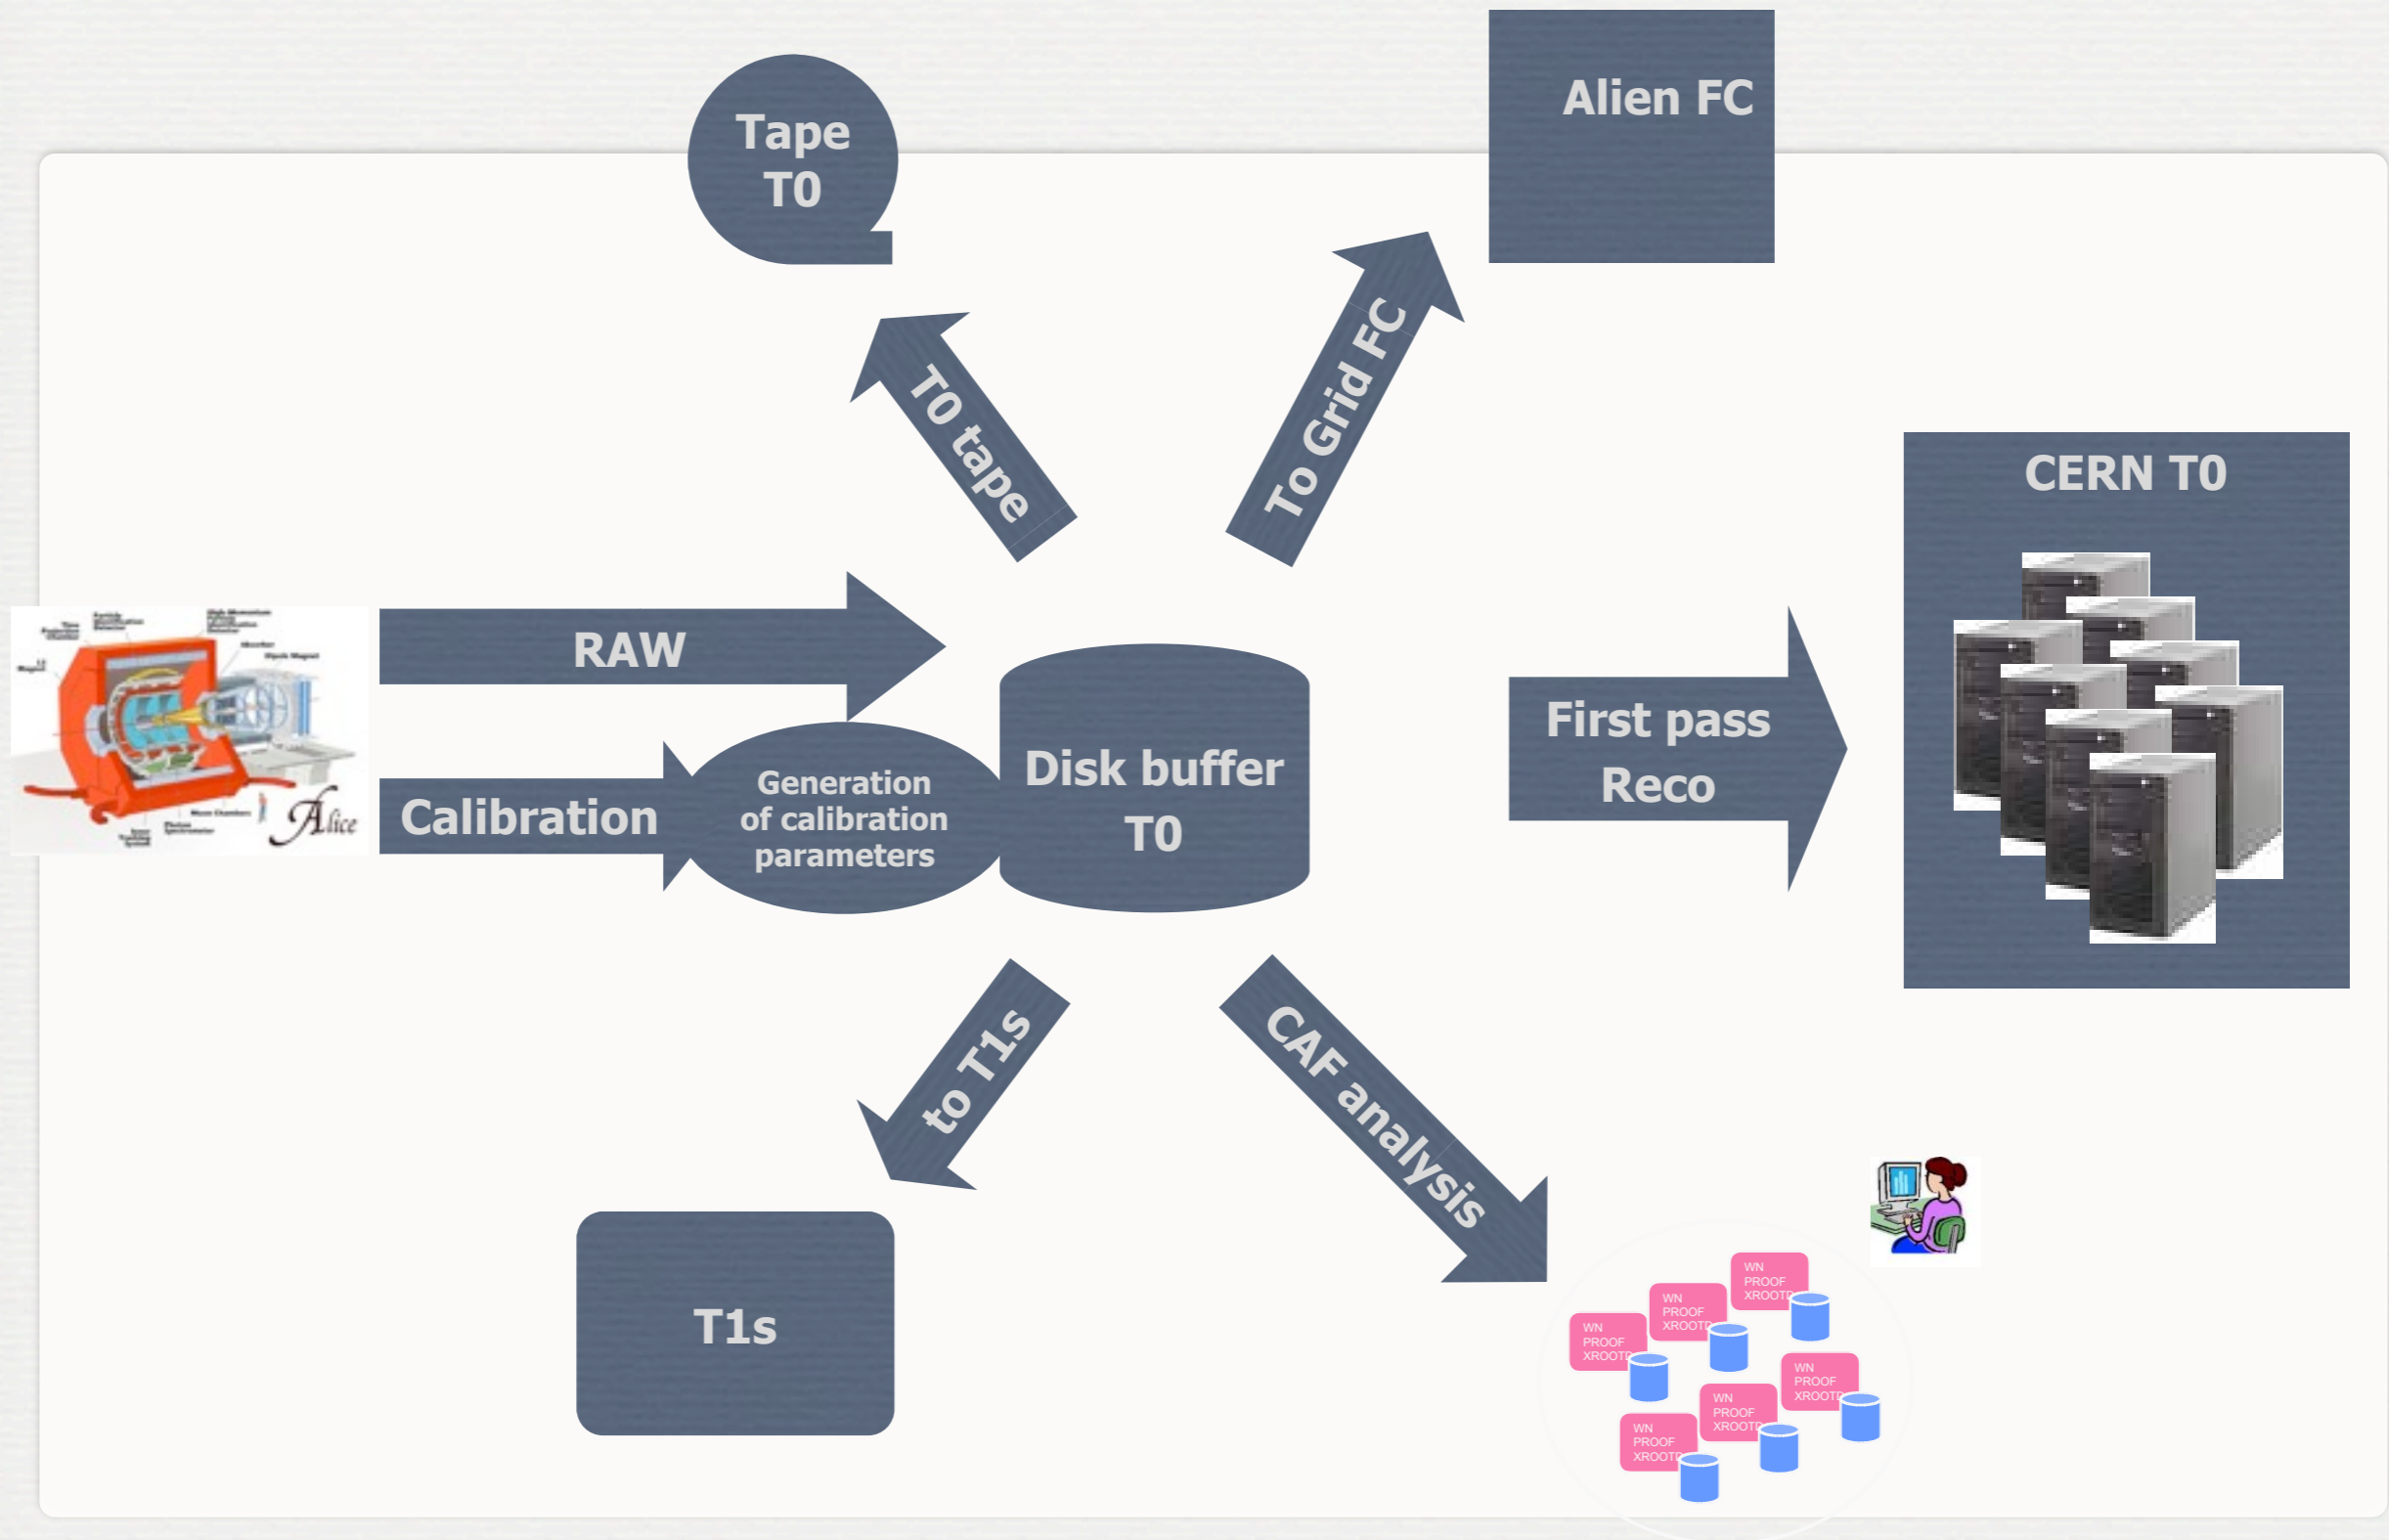
COMPUTING MODEL – PP



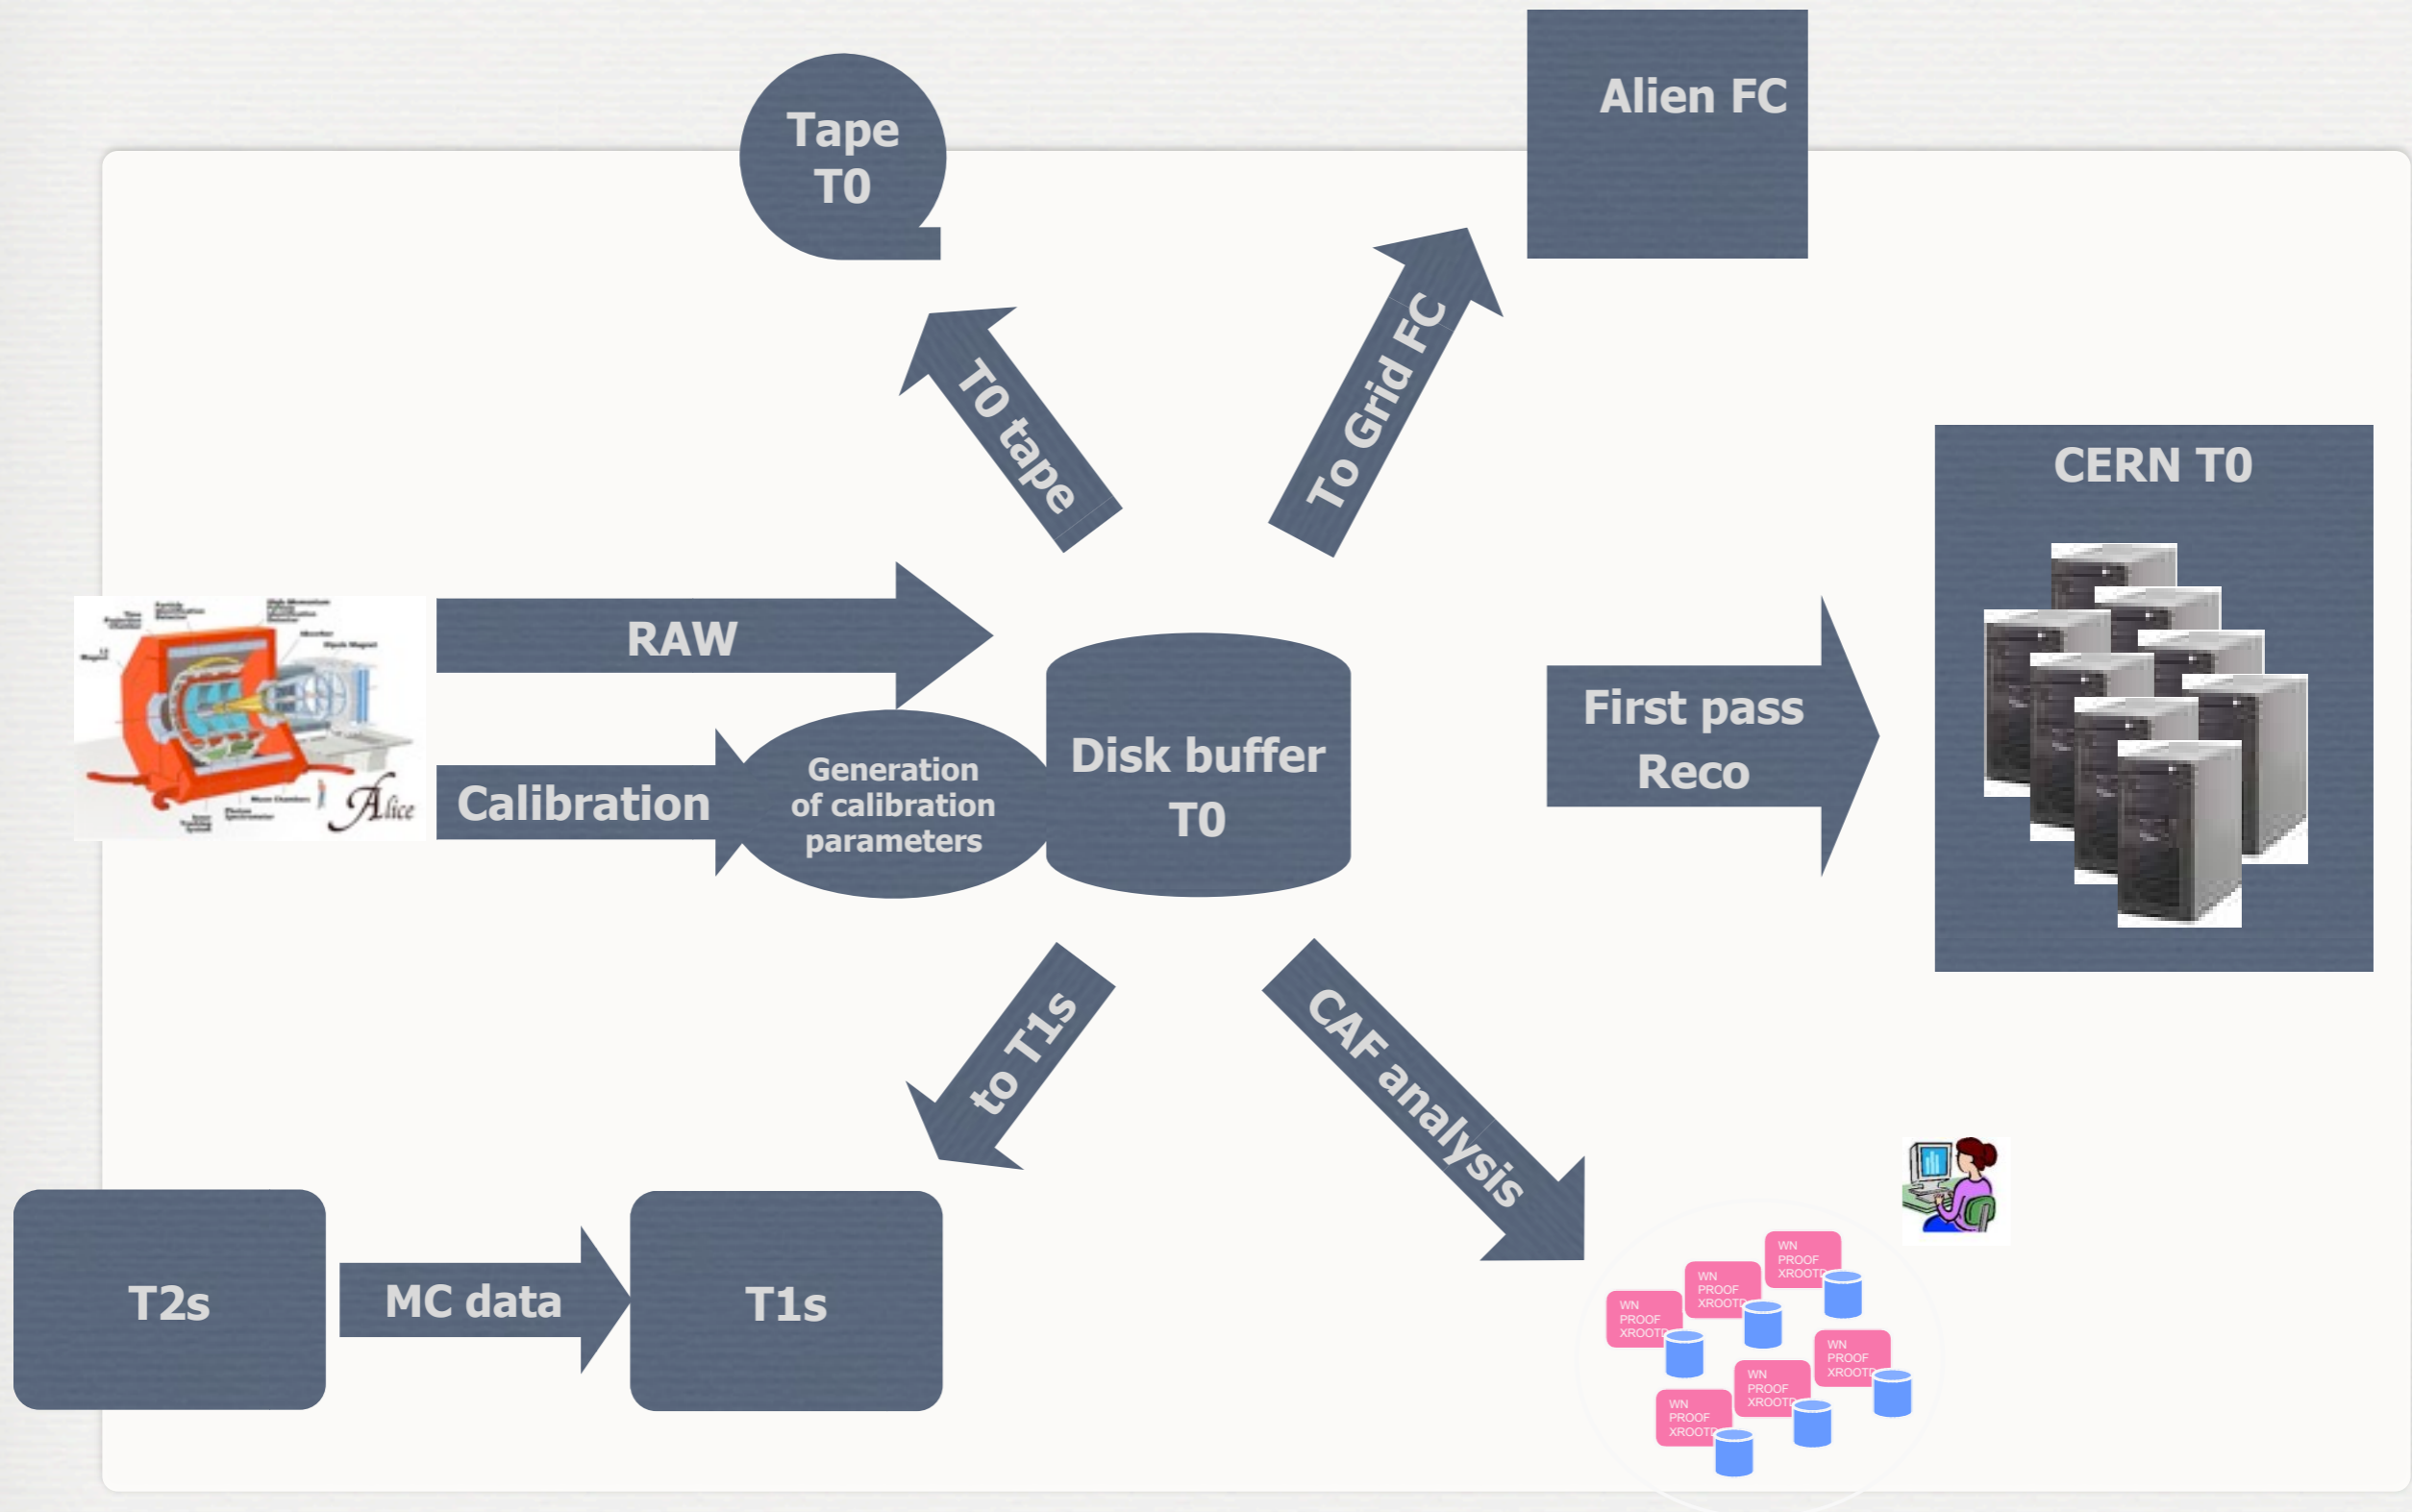
COMPUTING MODEL – PP



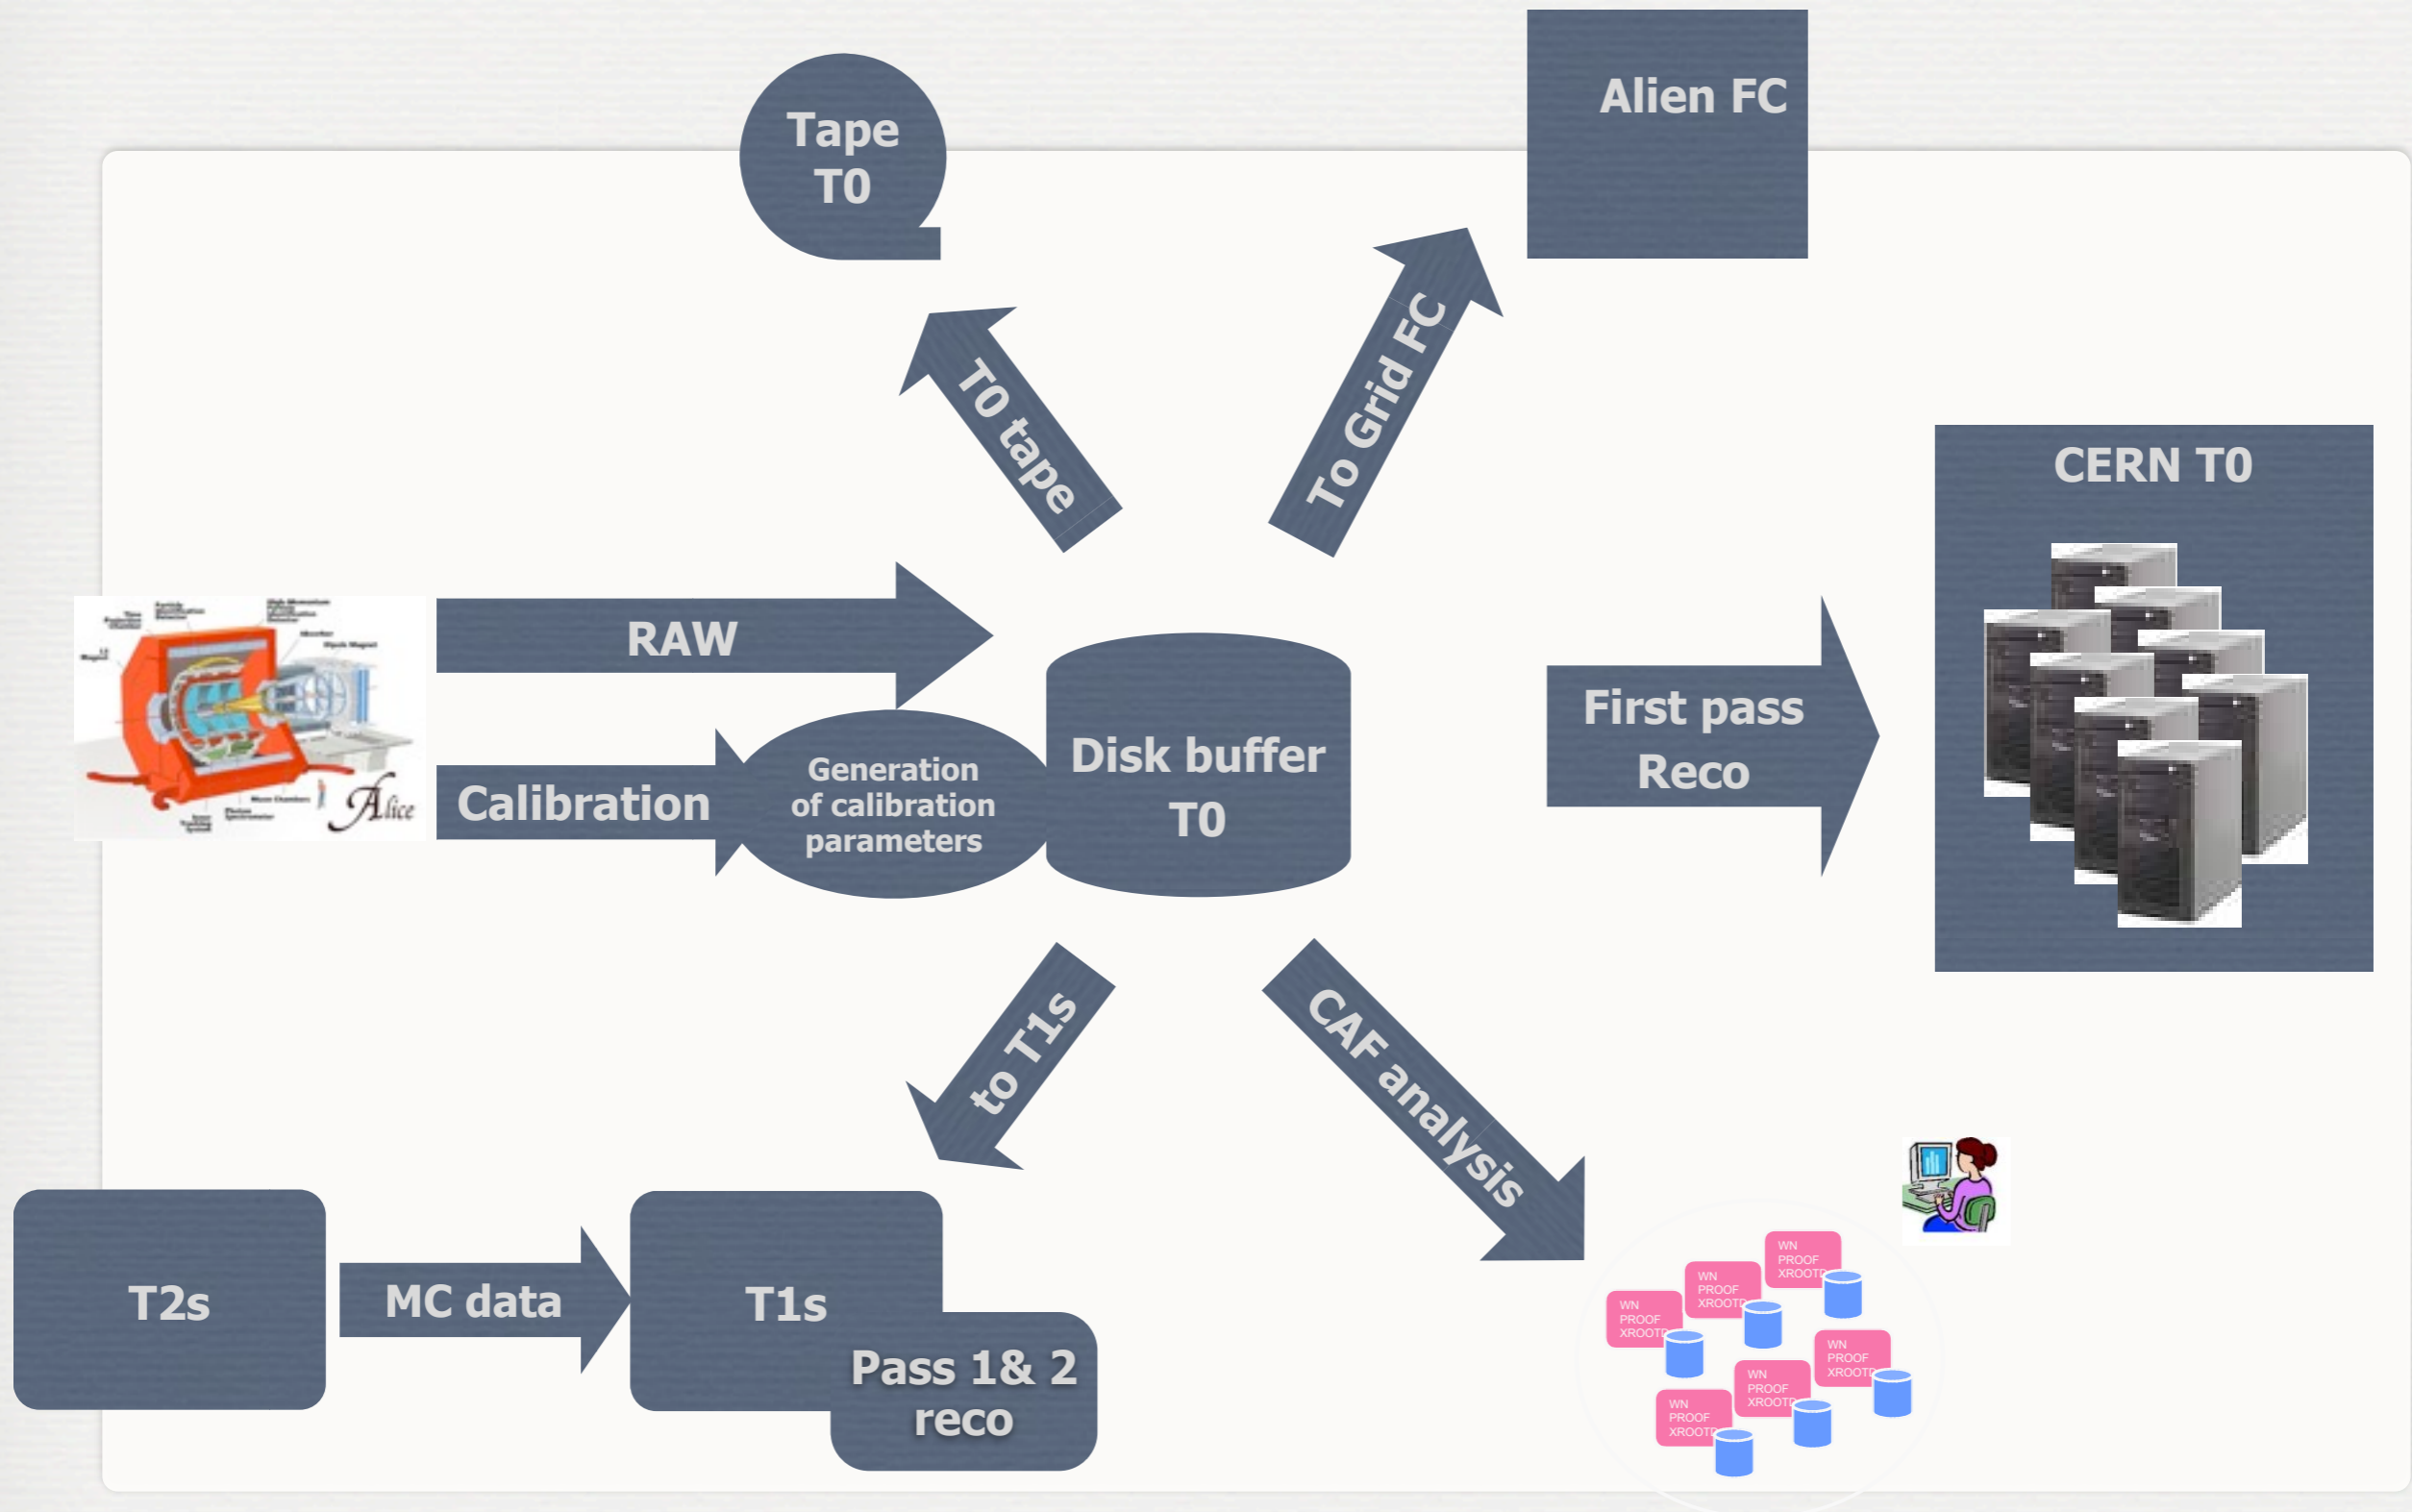
COMPUTING MODEL – PP



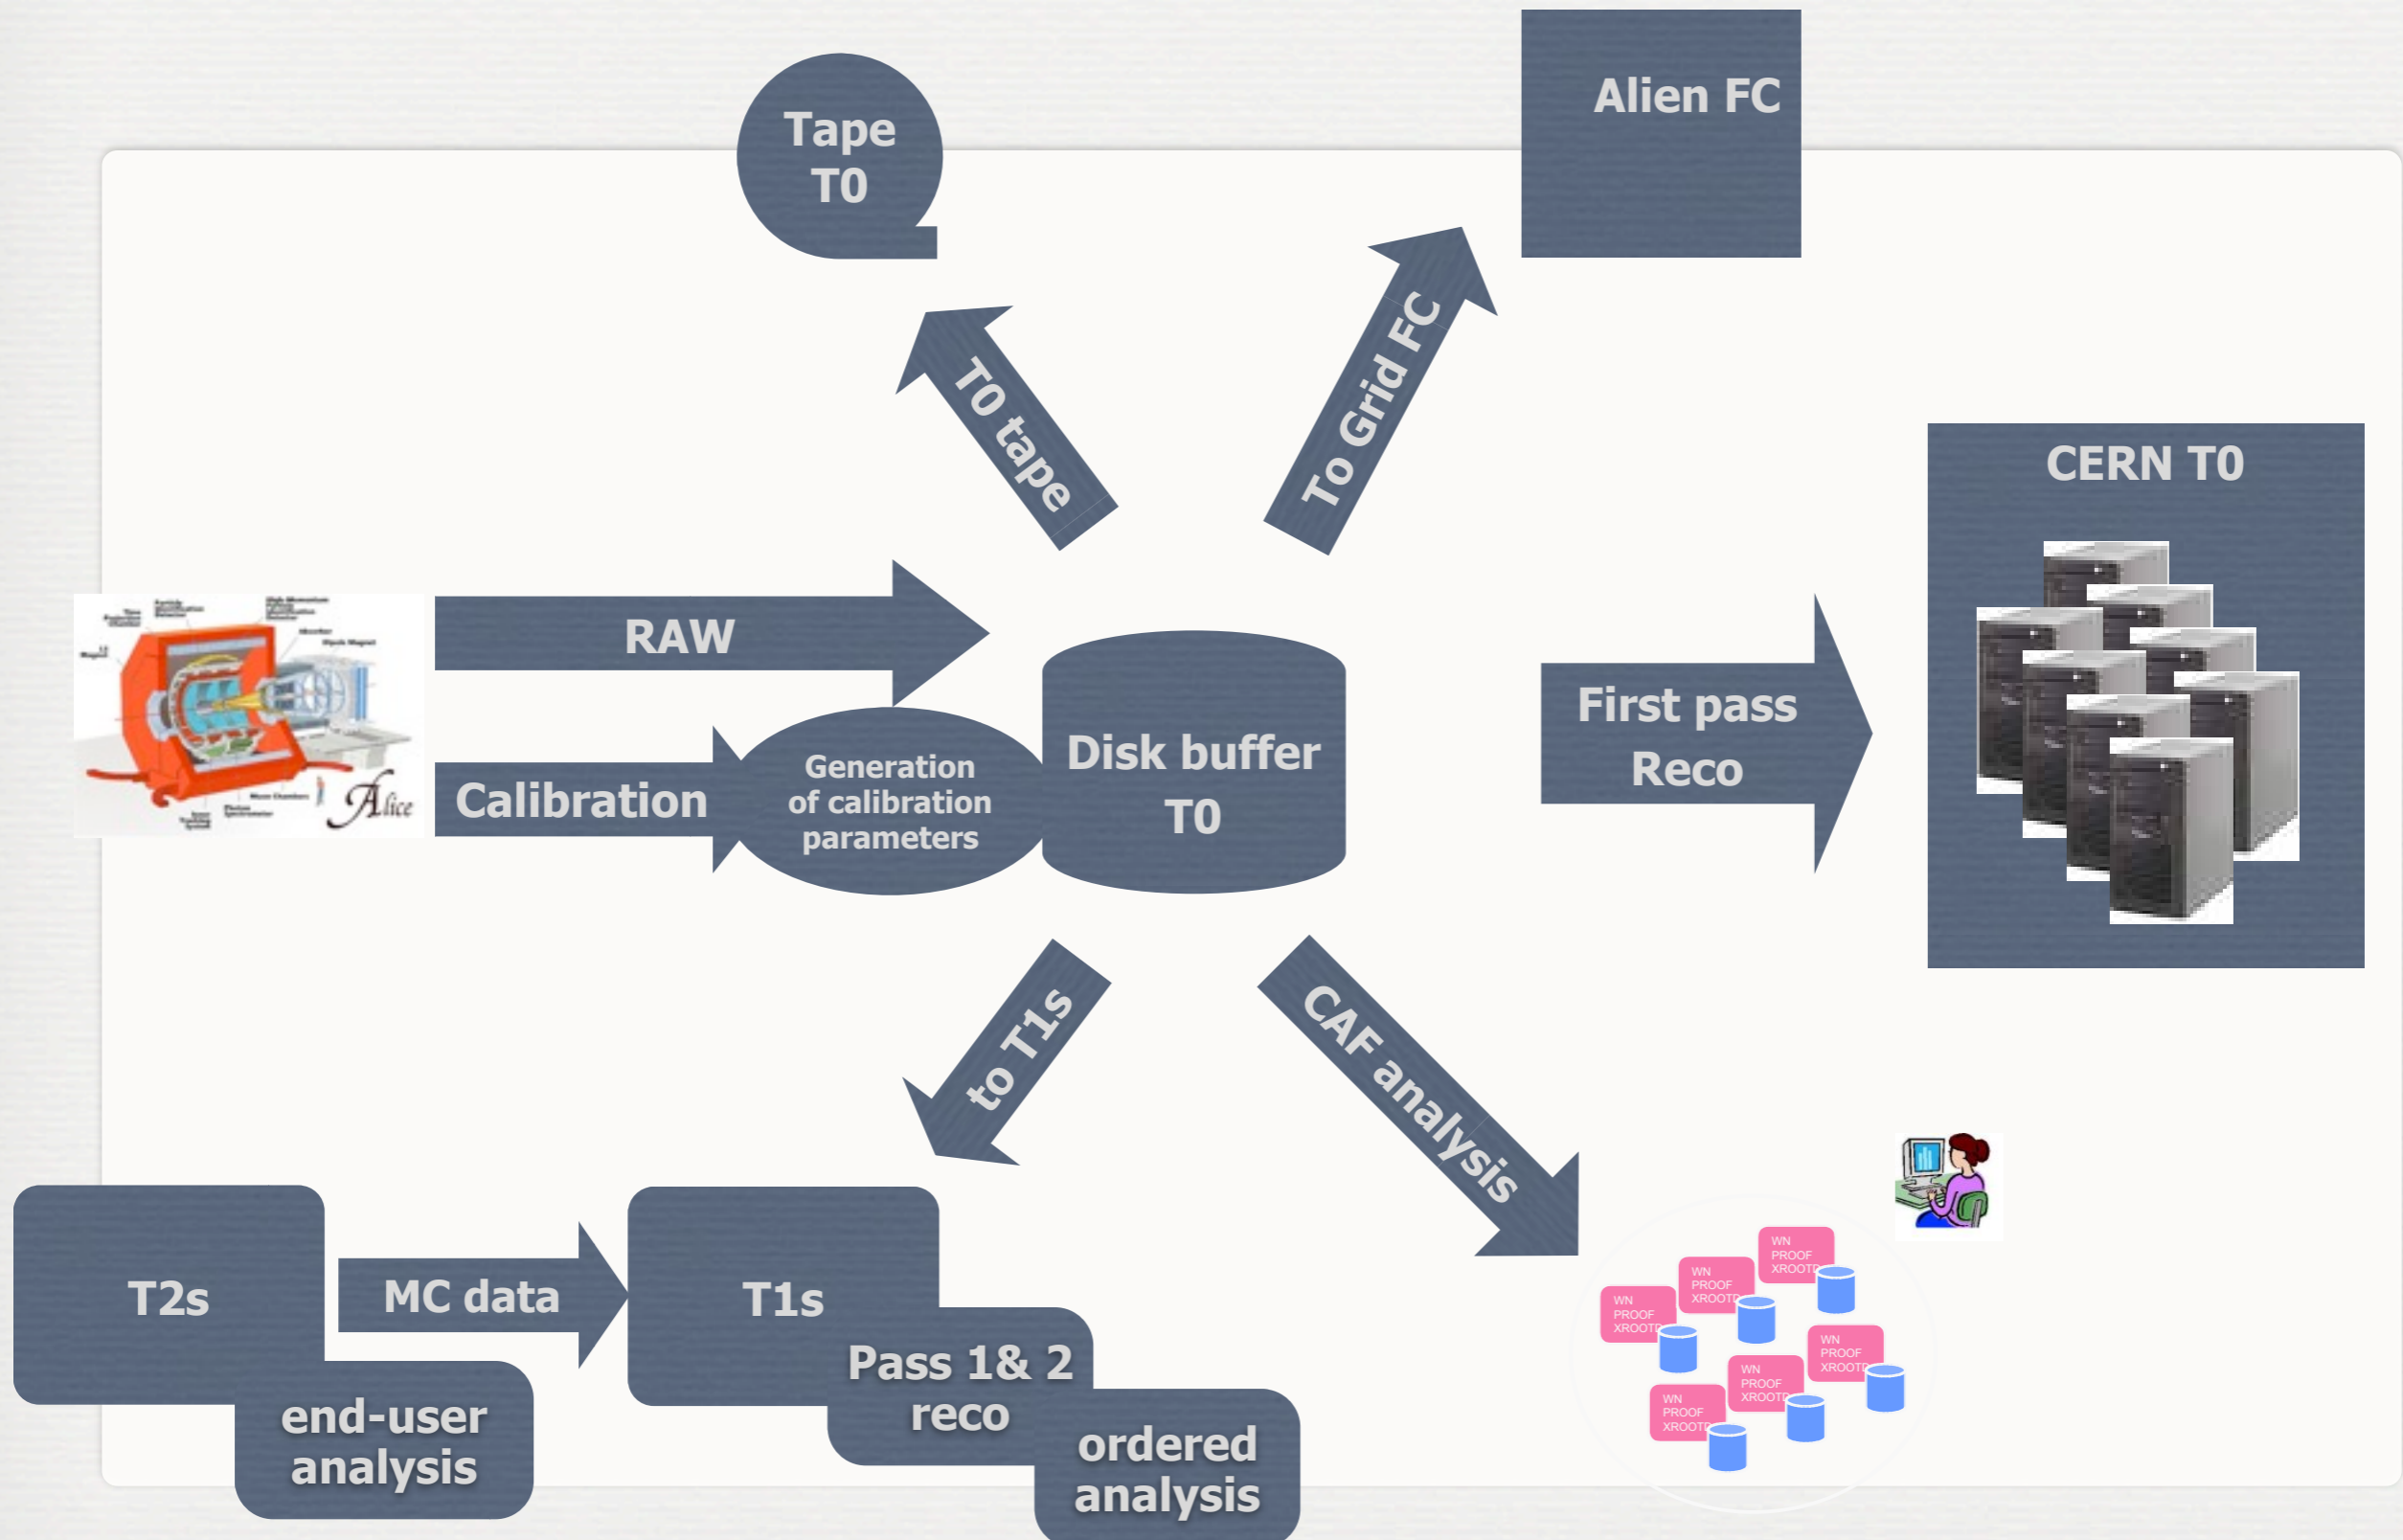
COMPUTING MODEL - PP



COMPUTING MODEL - PP

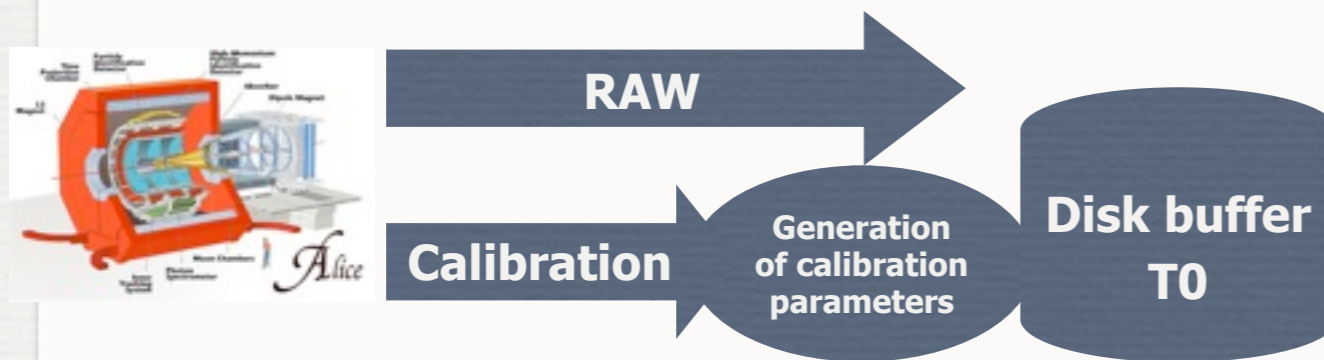


COMPUTING MODEL - PP

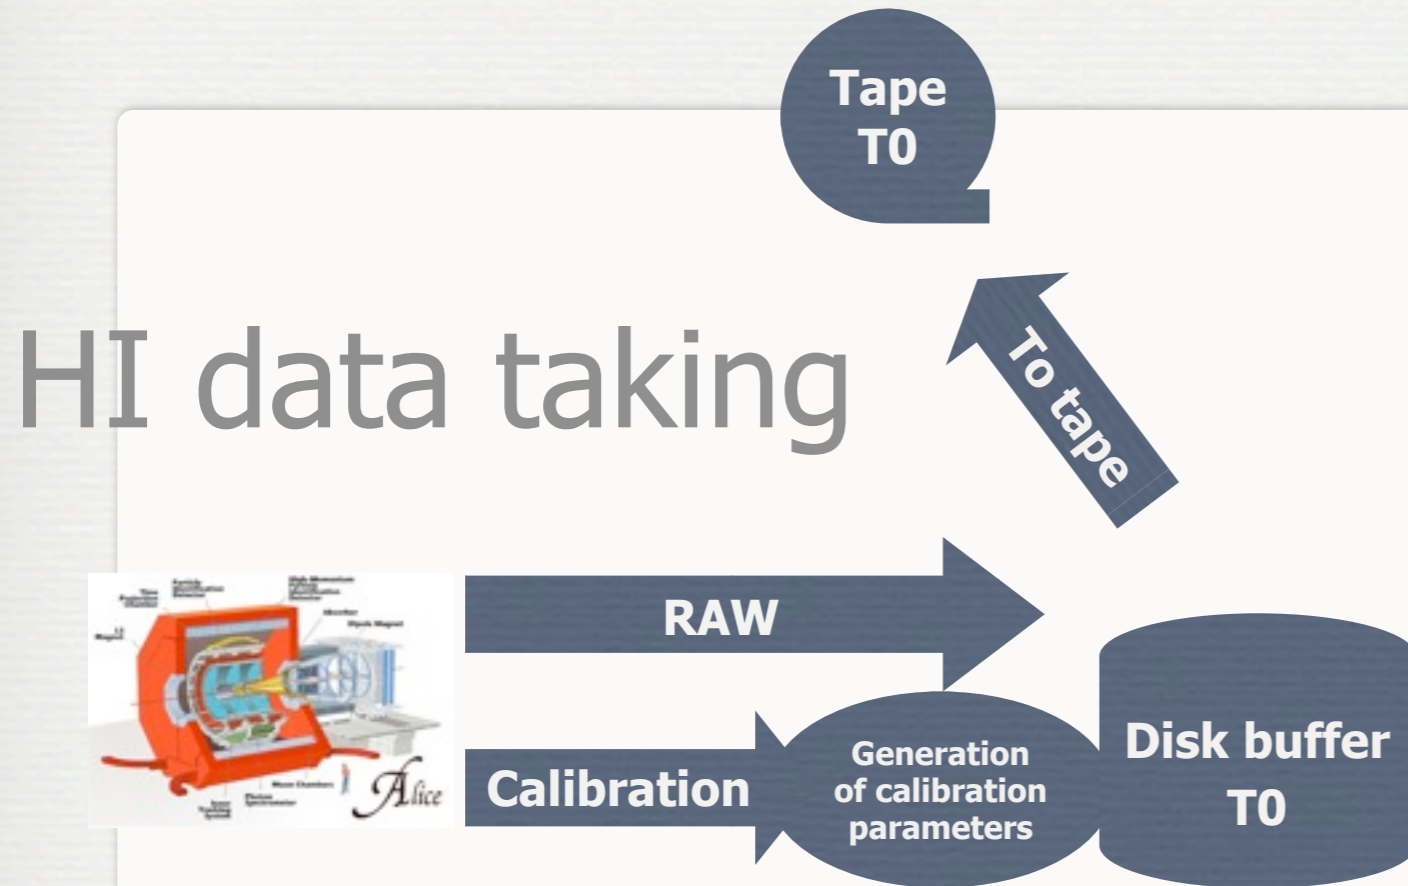


COMPUTING MODEL - AA

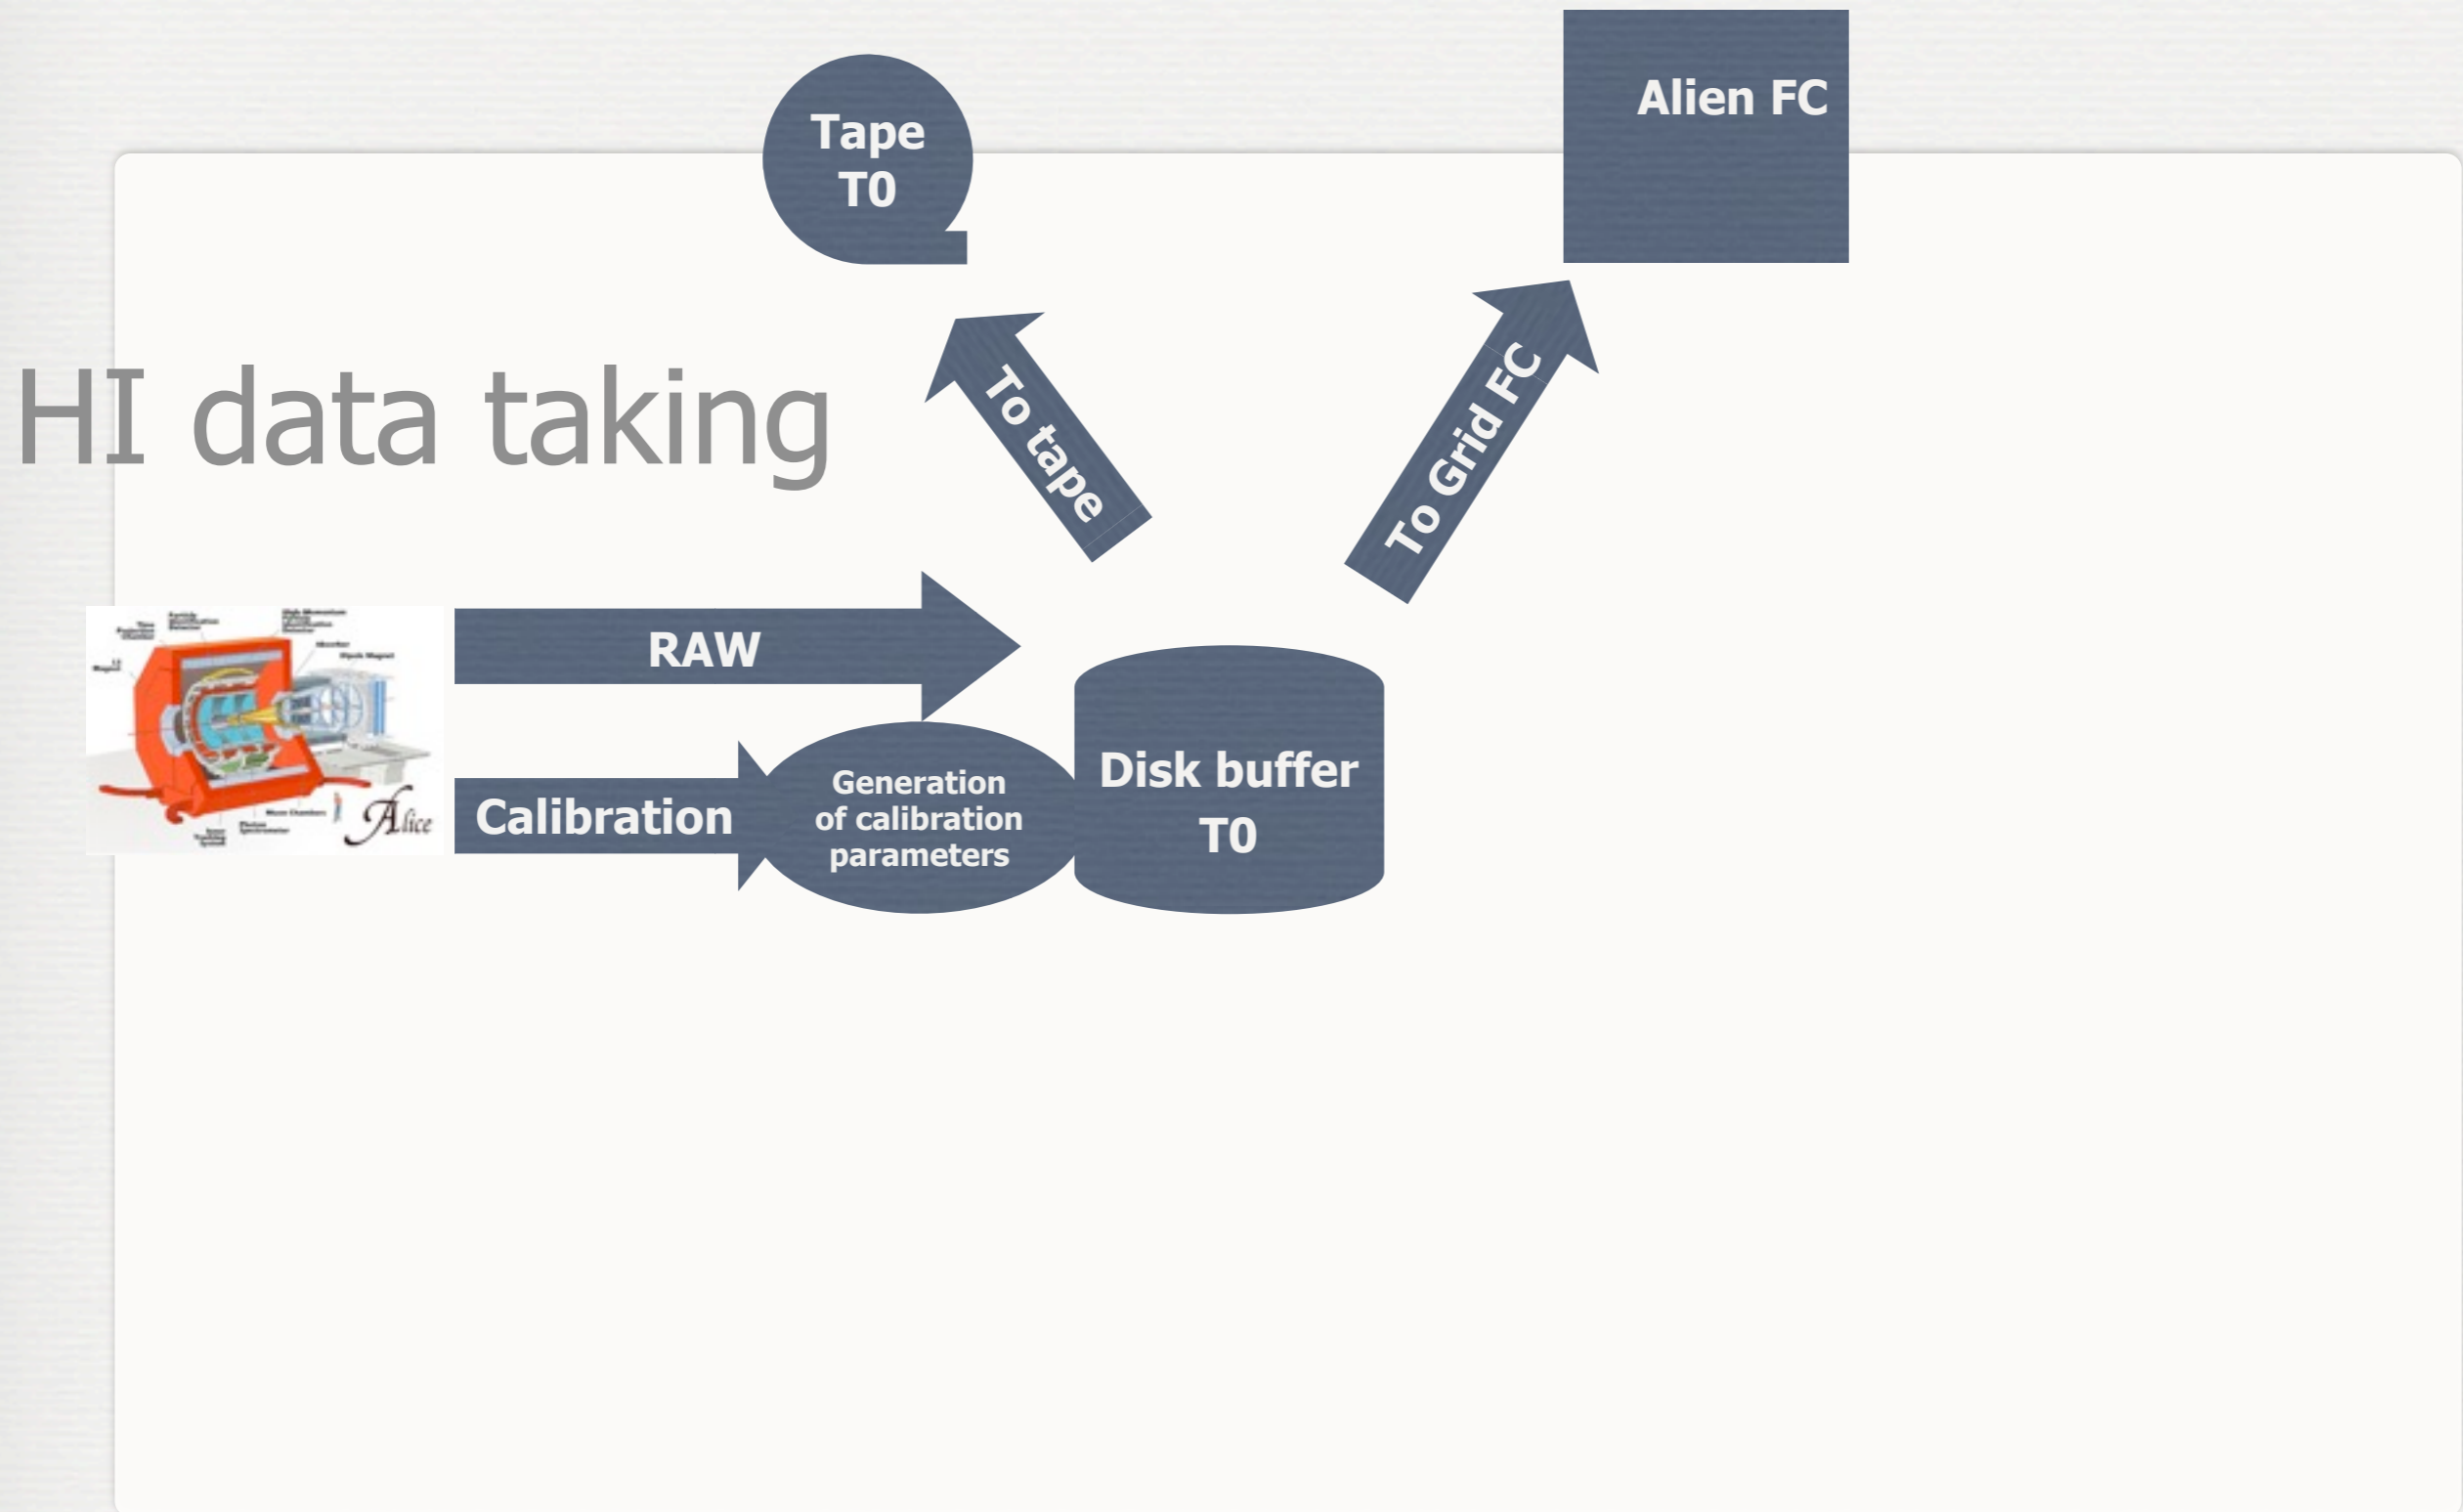
HI data taking



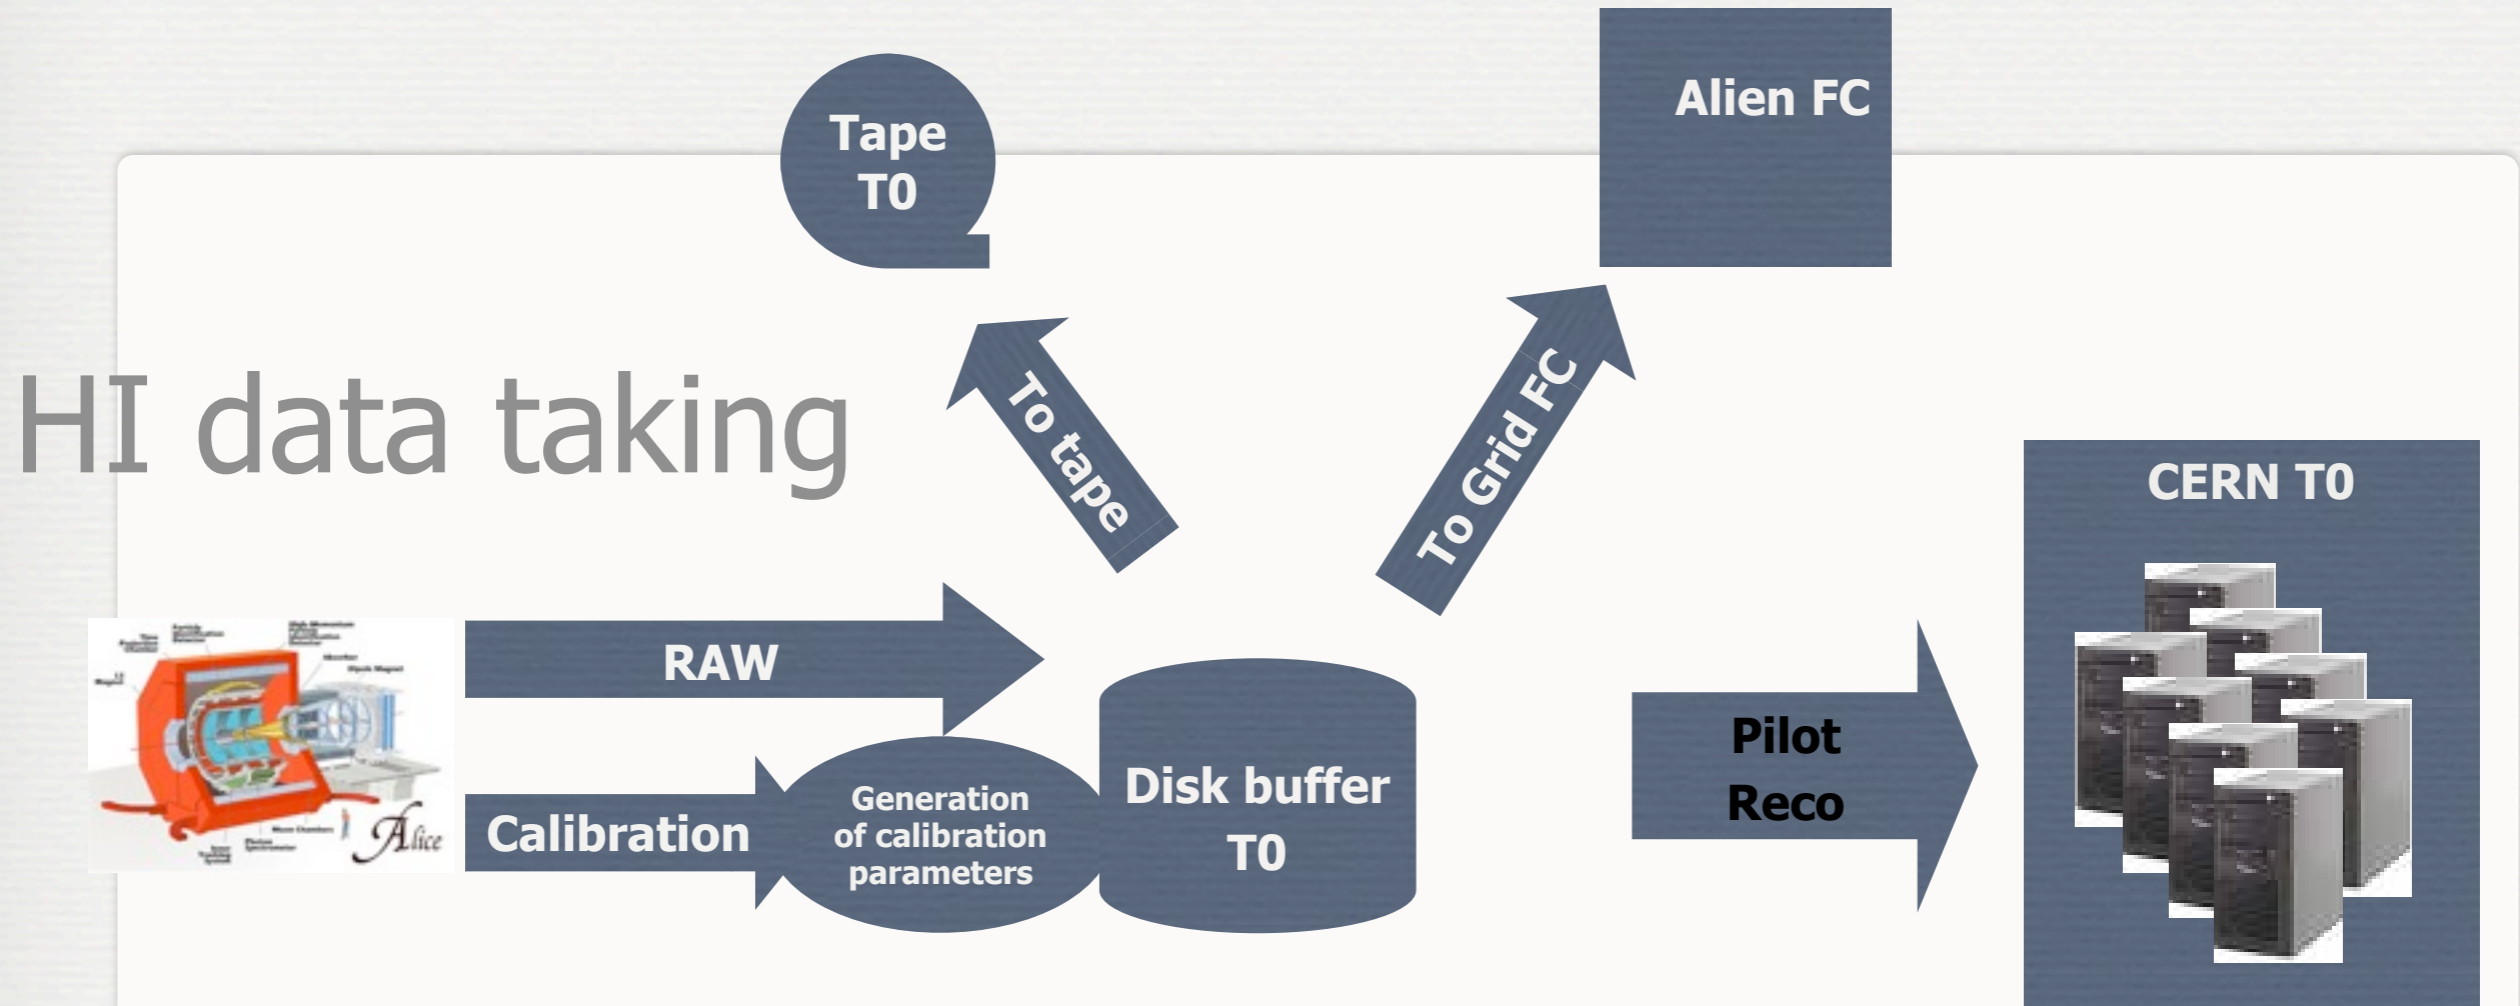
COMPUTING MODEL - AA



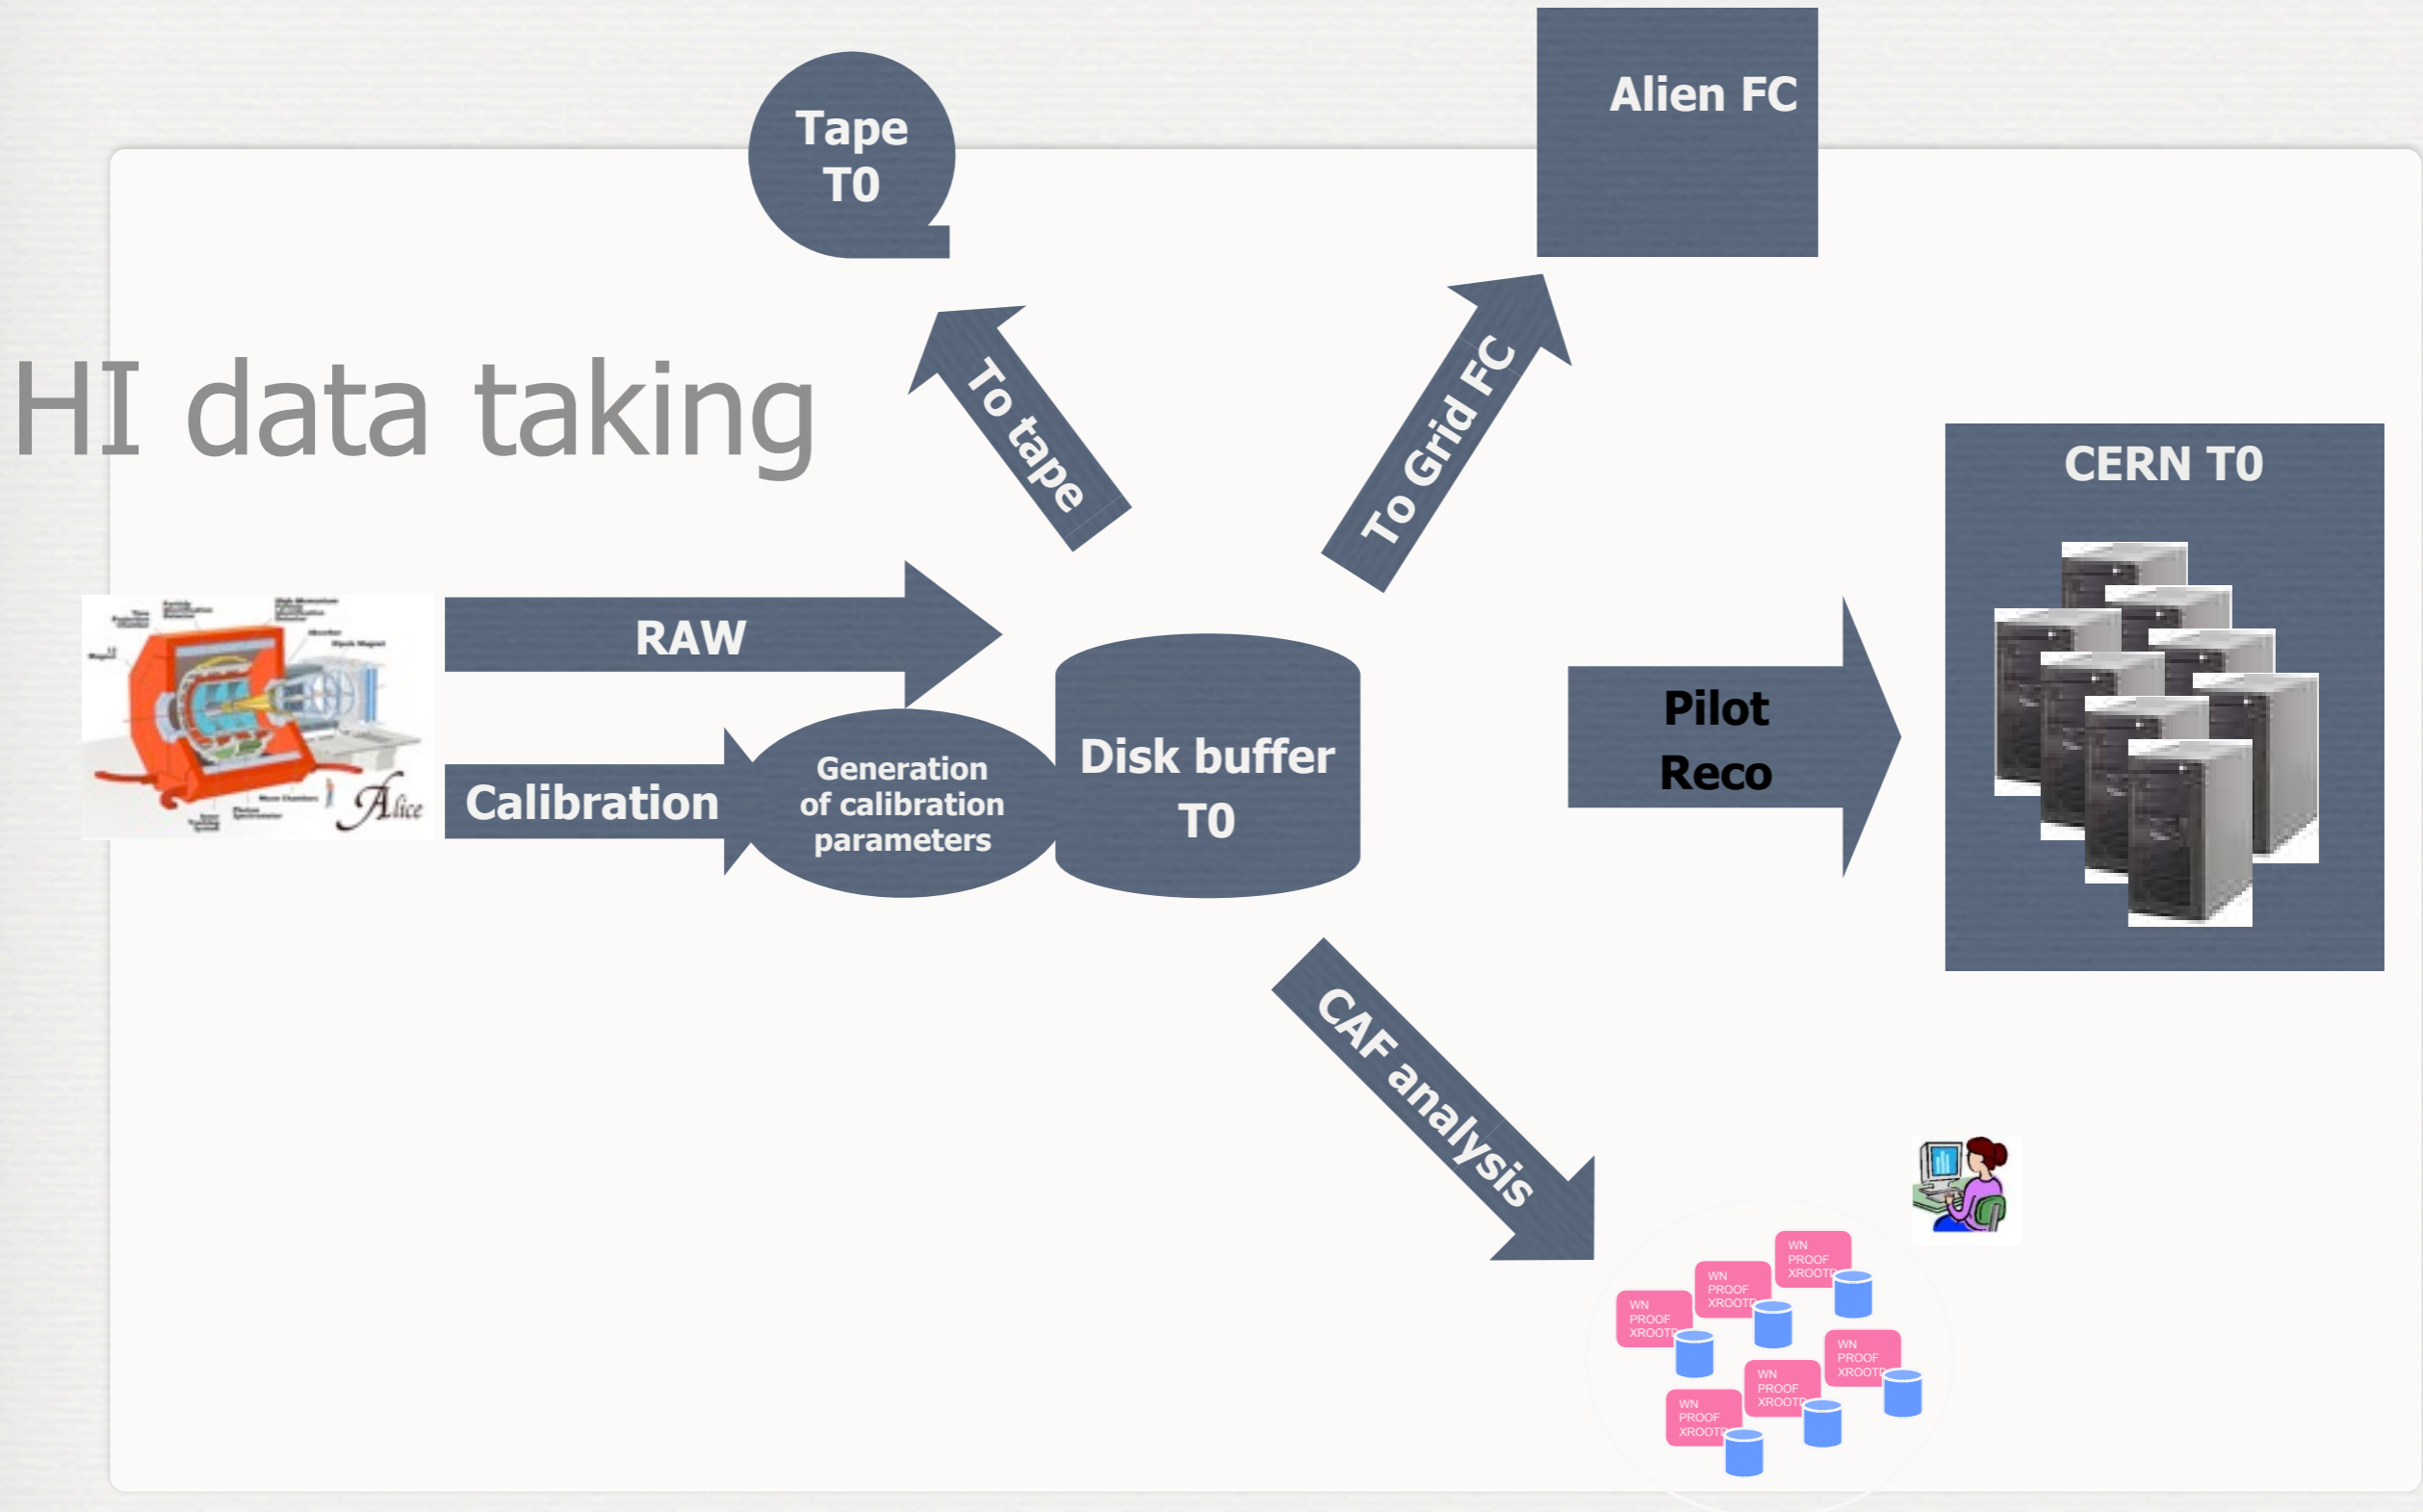
COMPUTING MODEL - AA



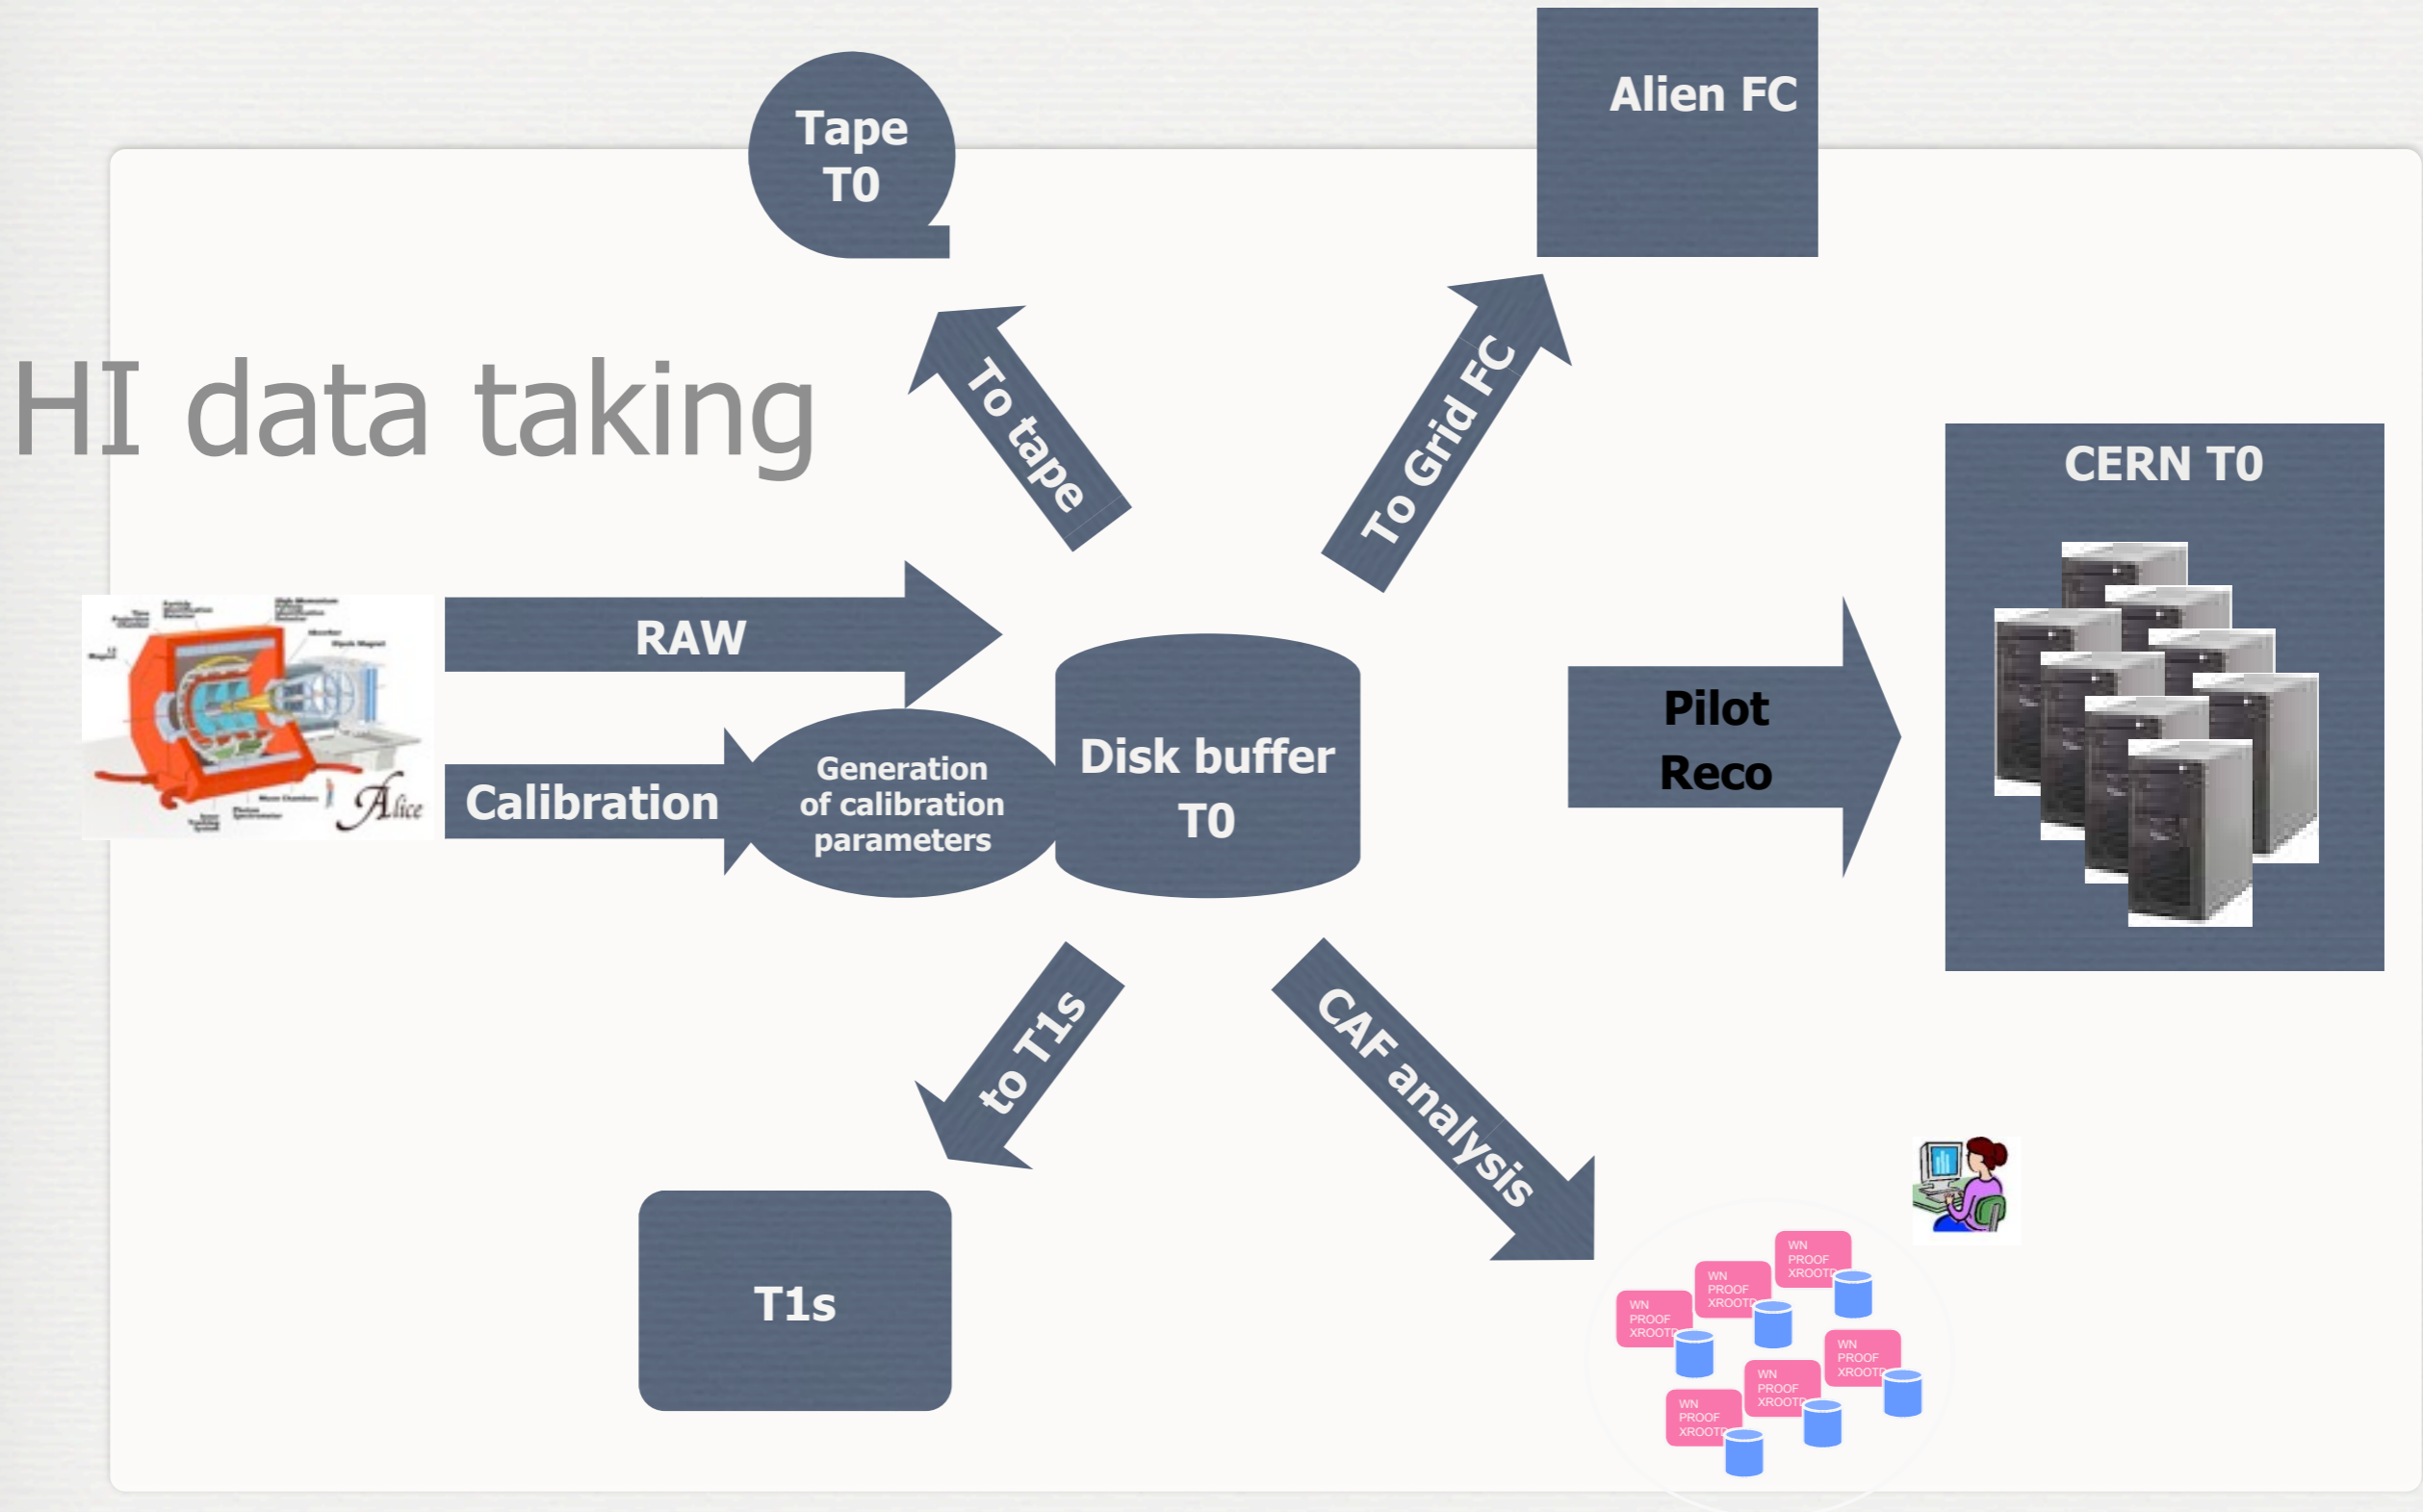
COMPUTING MODEL - AA



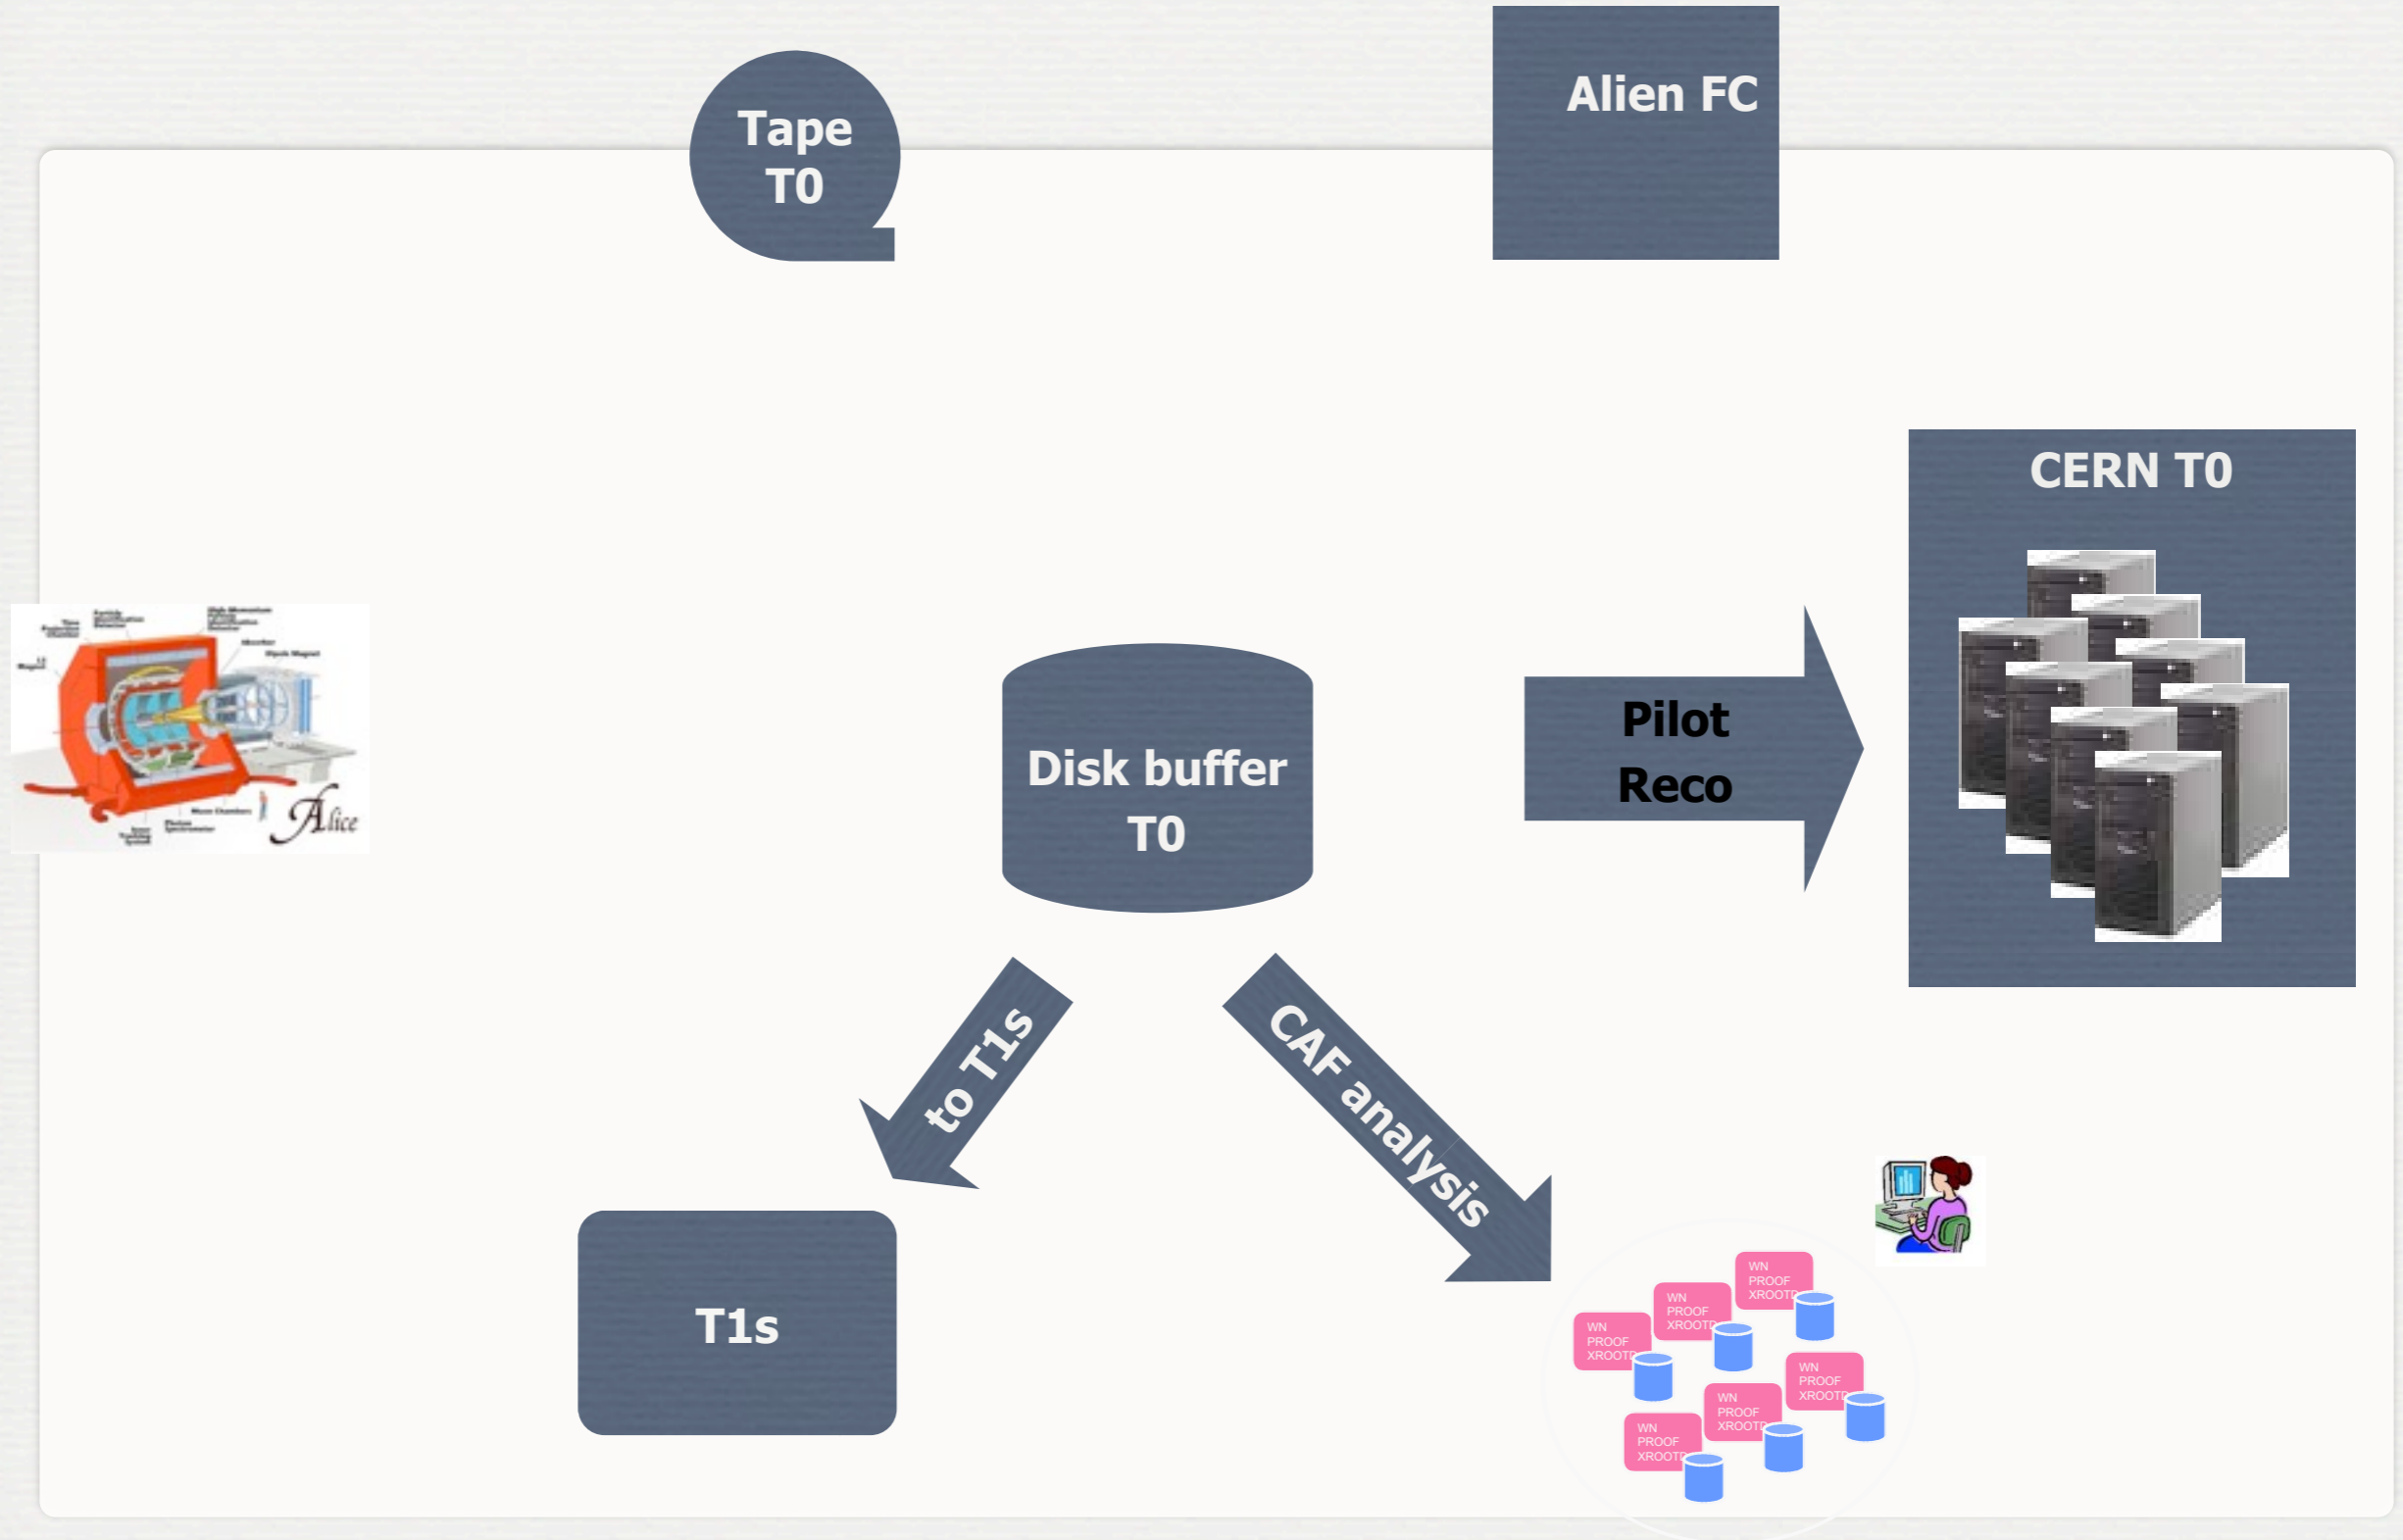
COMPUTING MODEL - AA



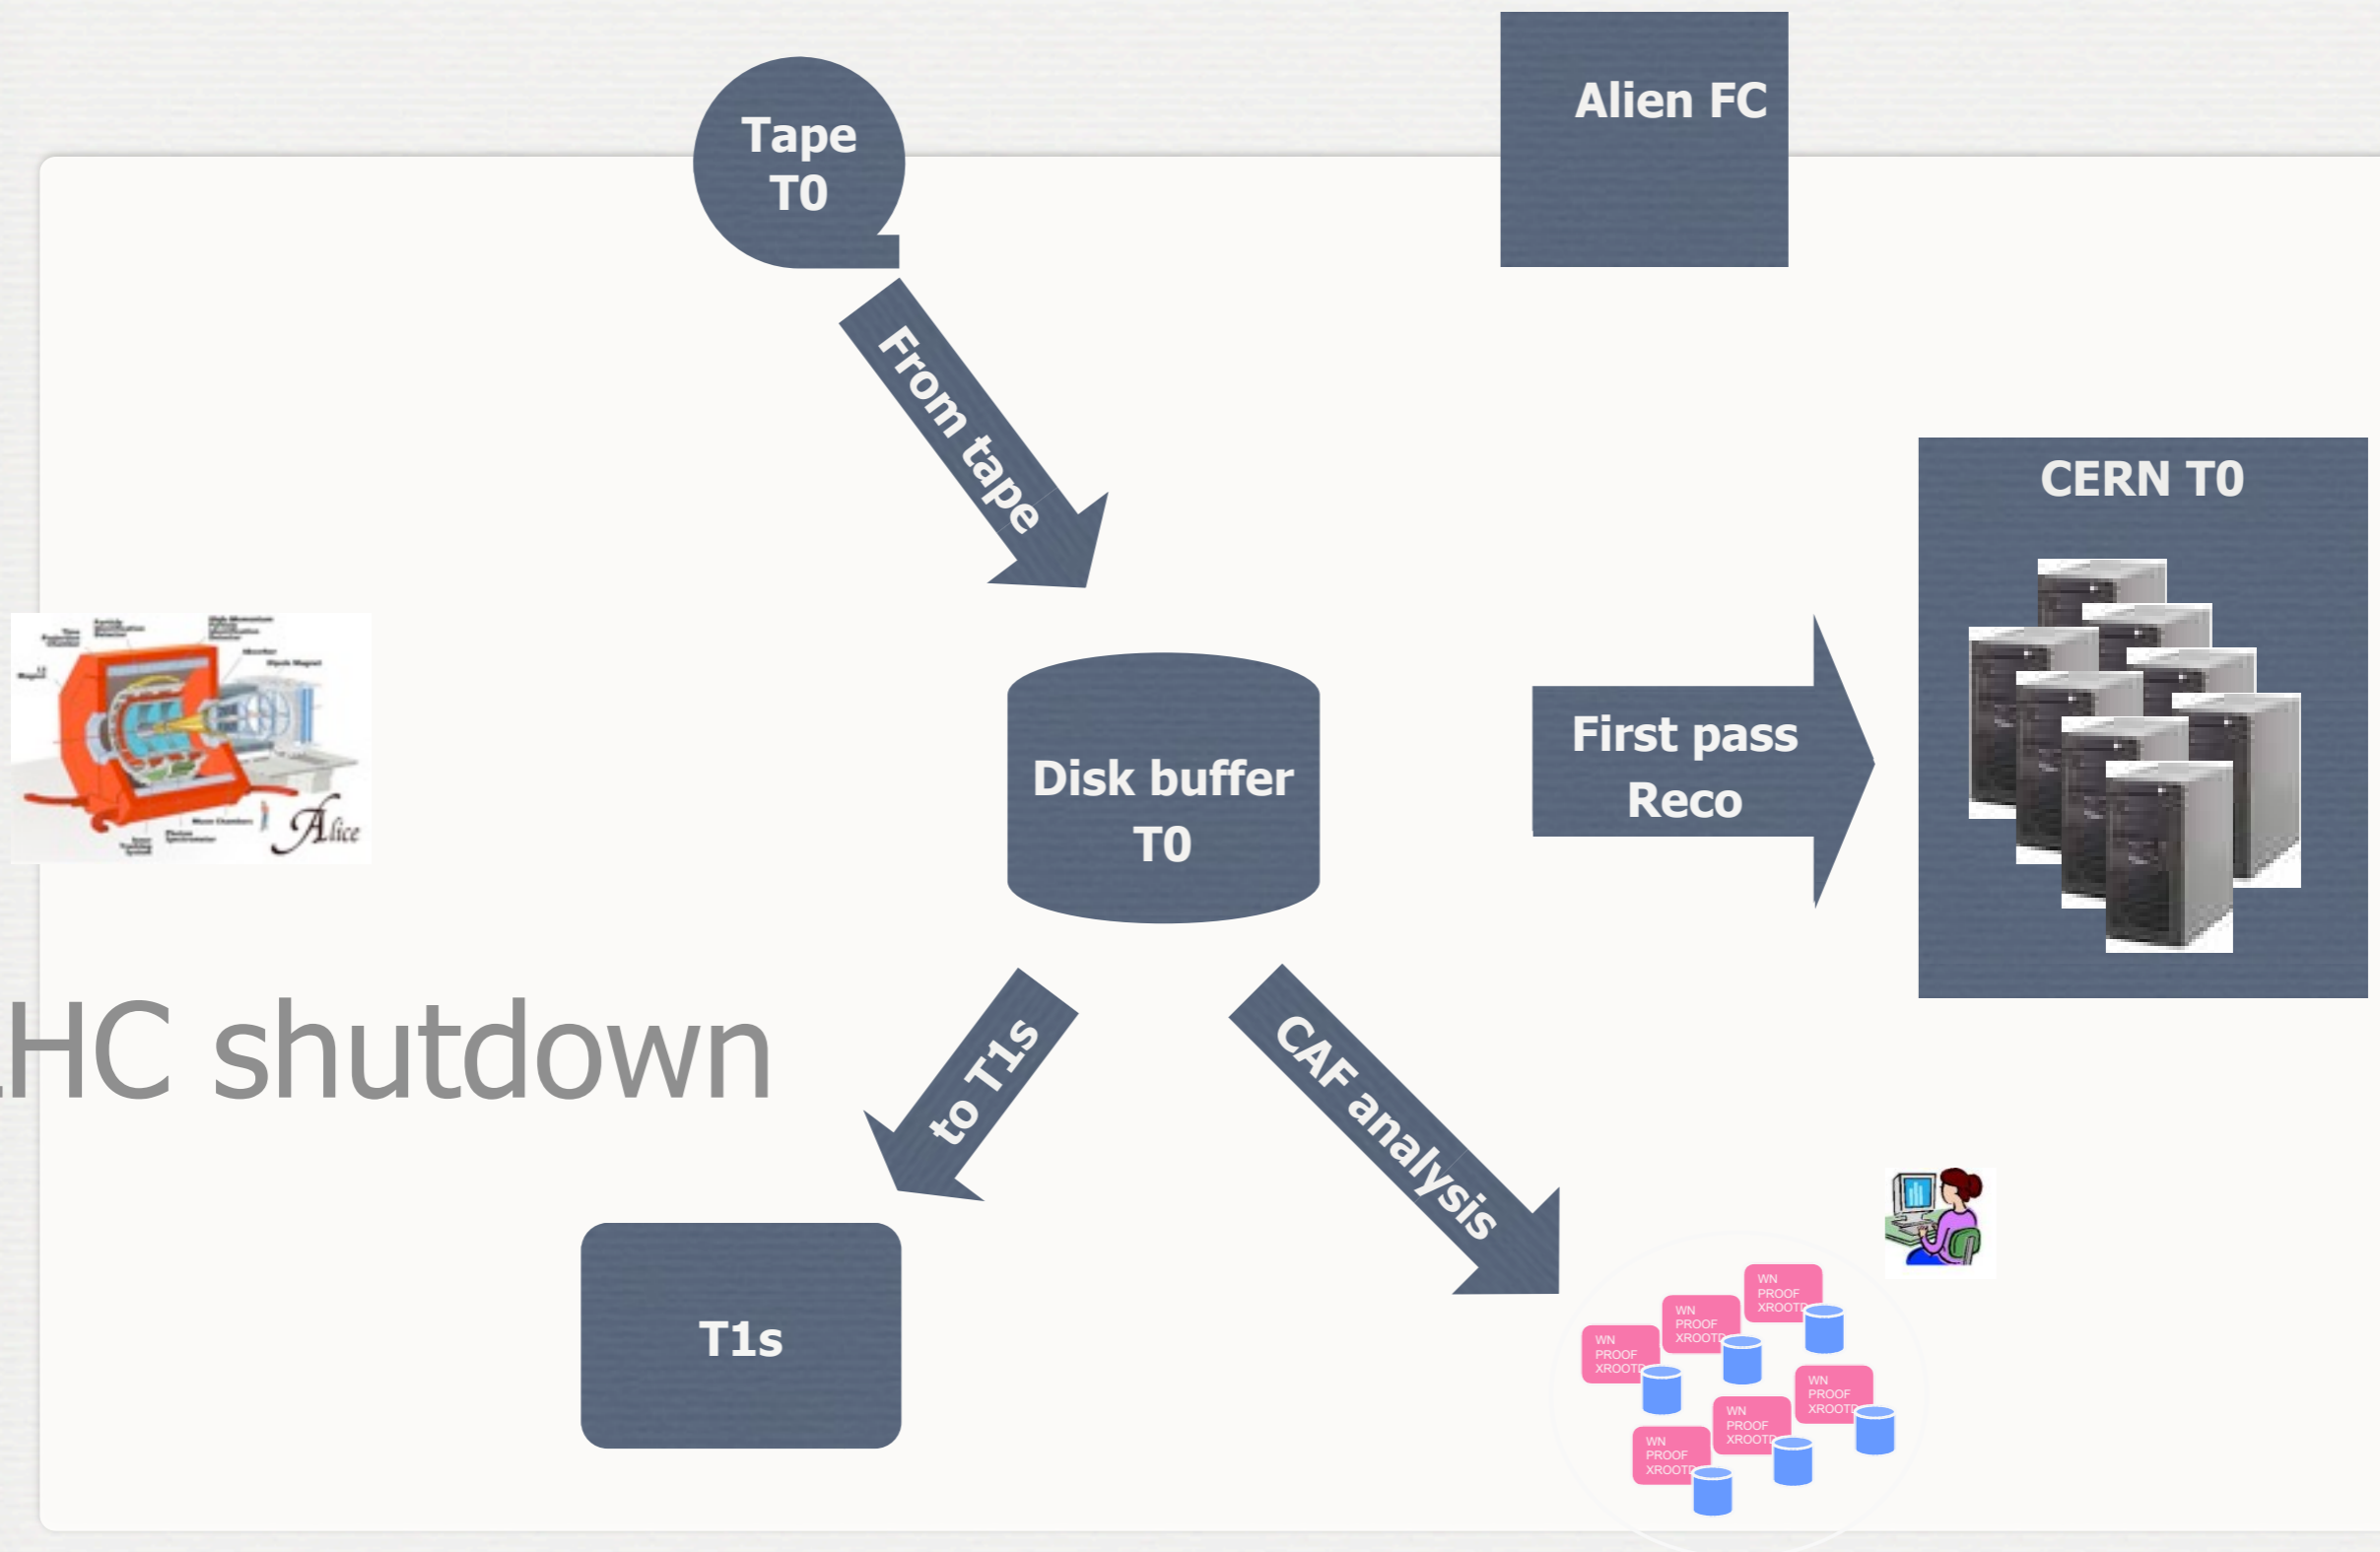
COMPUTING MODEL - AA



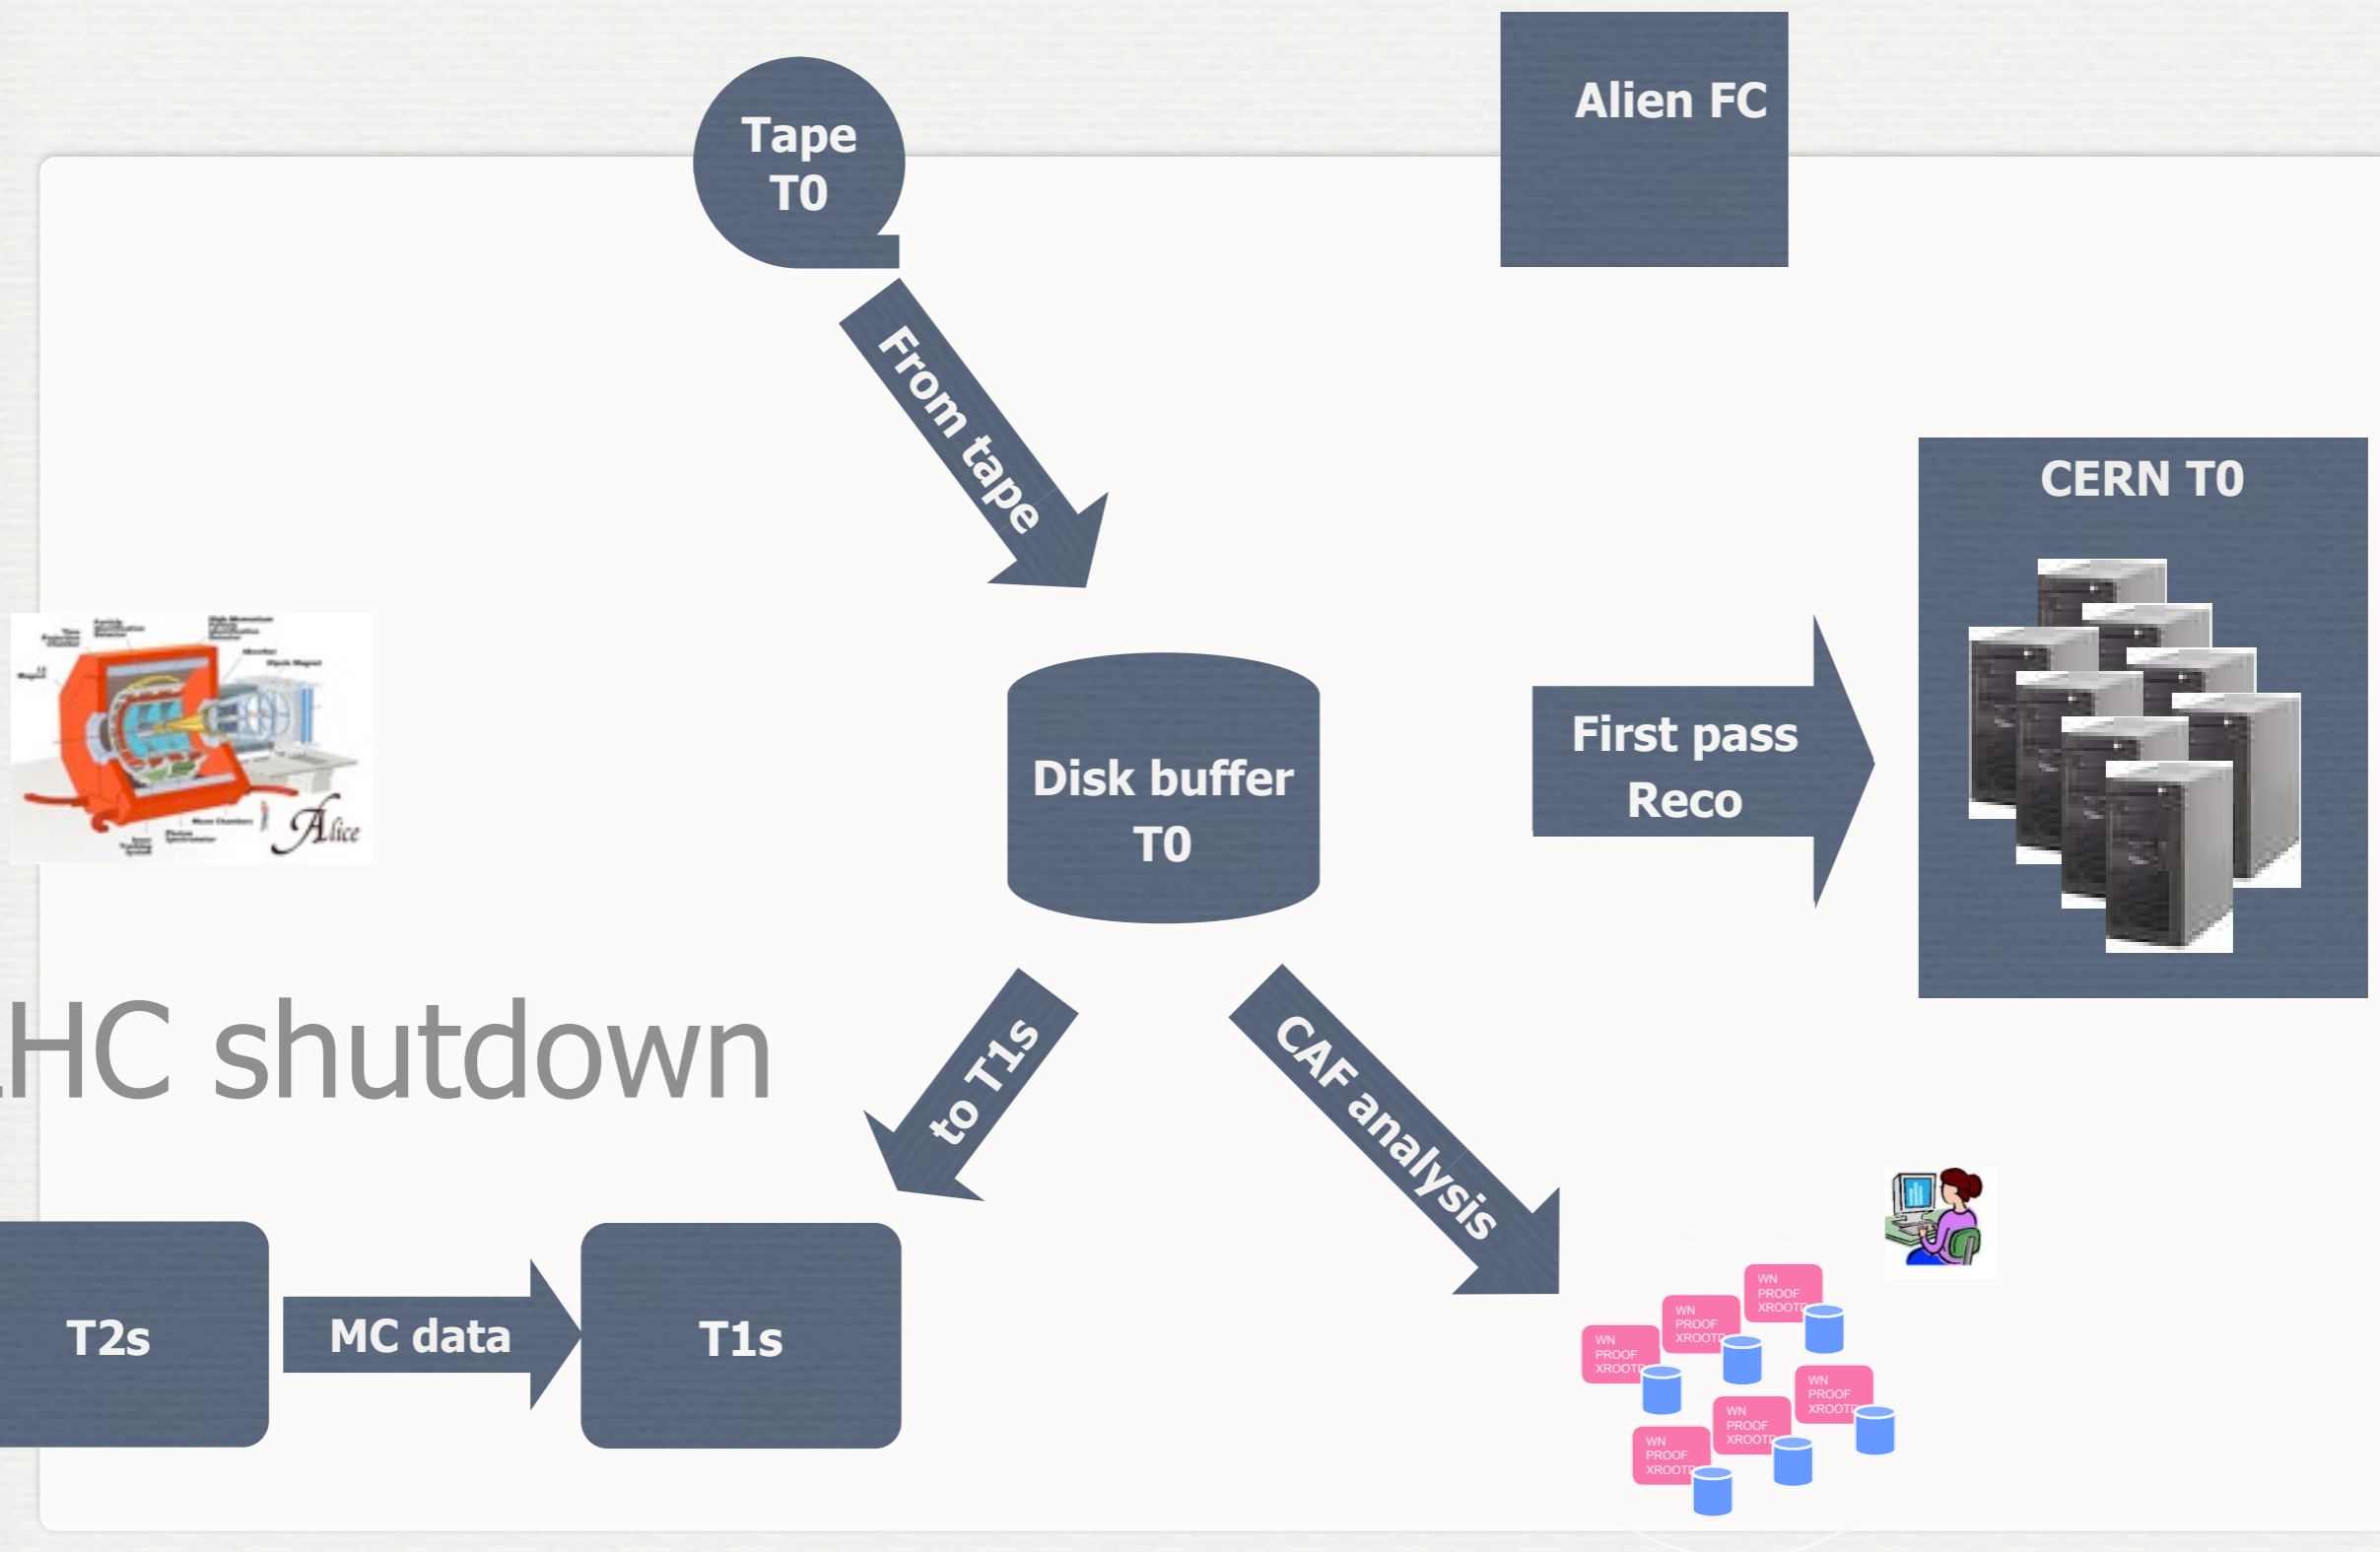
COMPUTING MODEL - AA



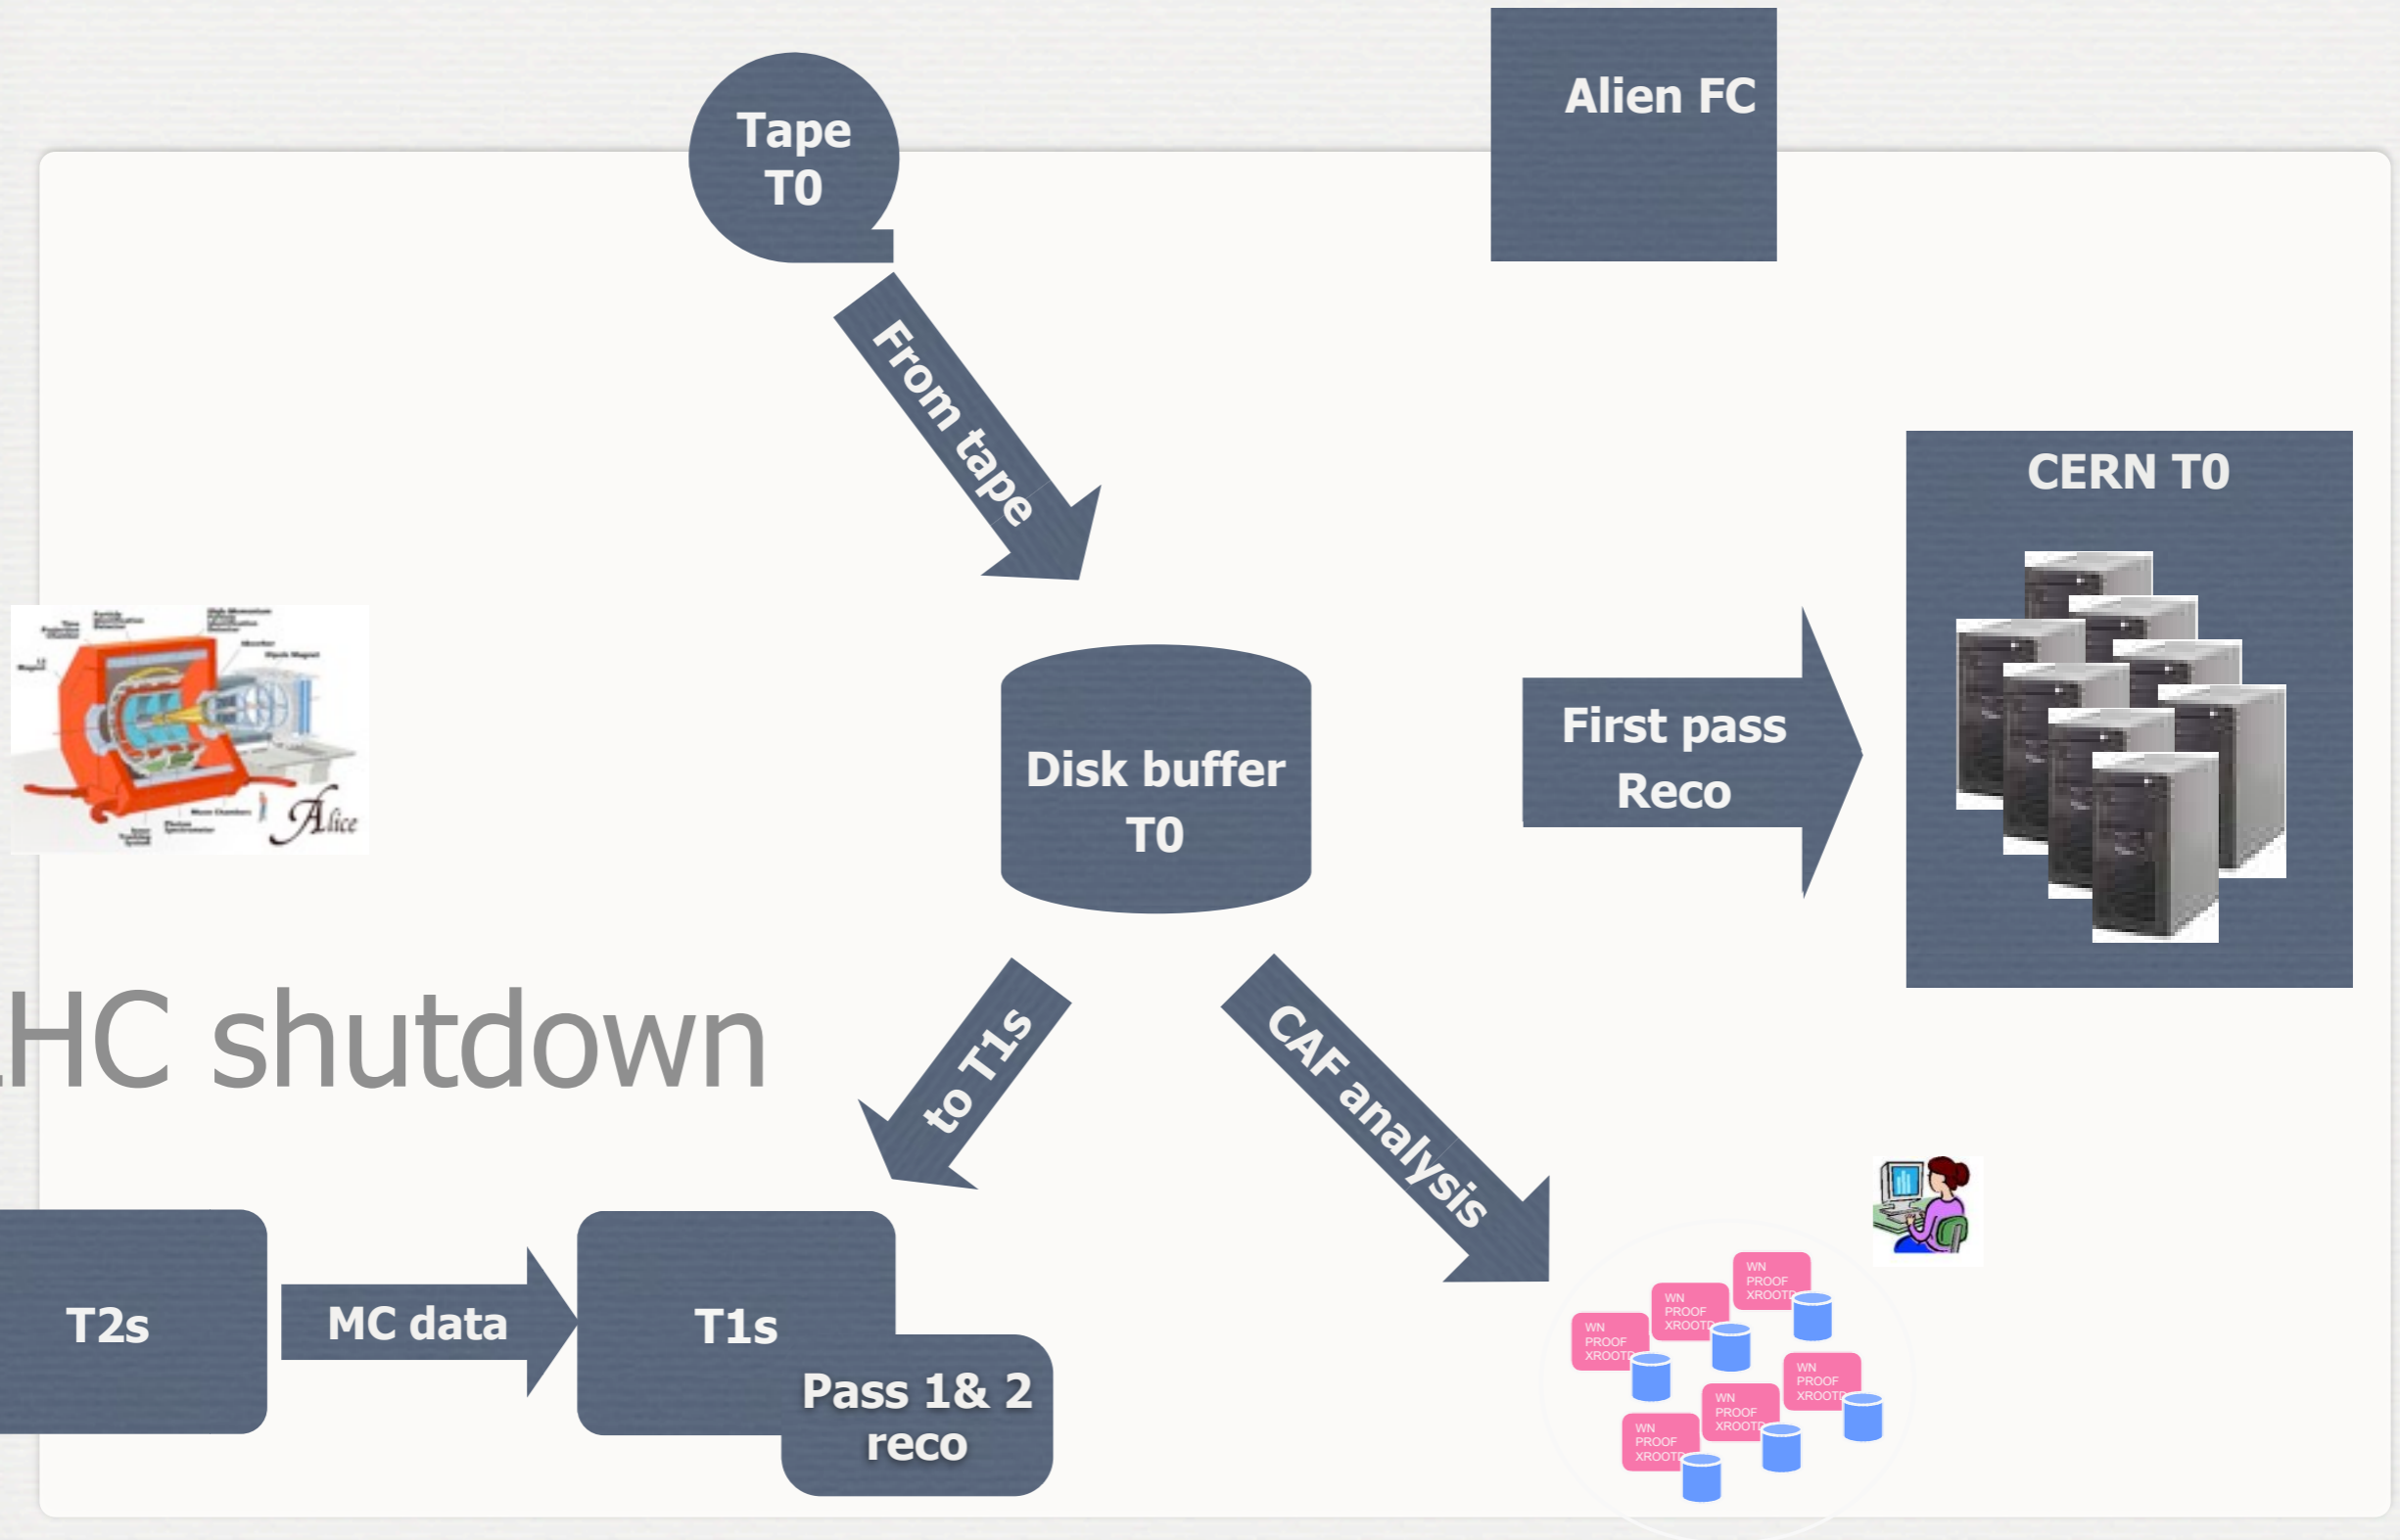
COMPUTING MODEL - AA



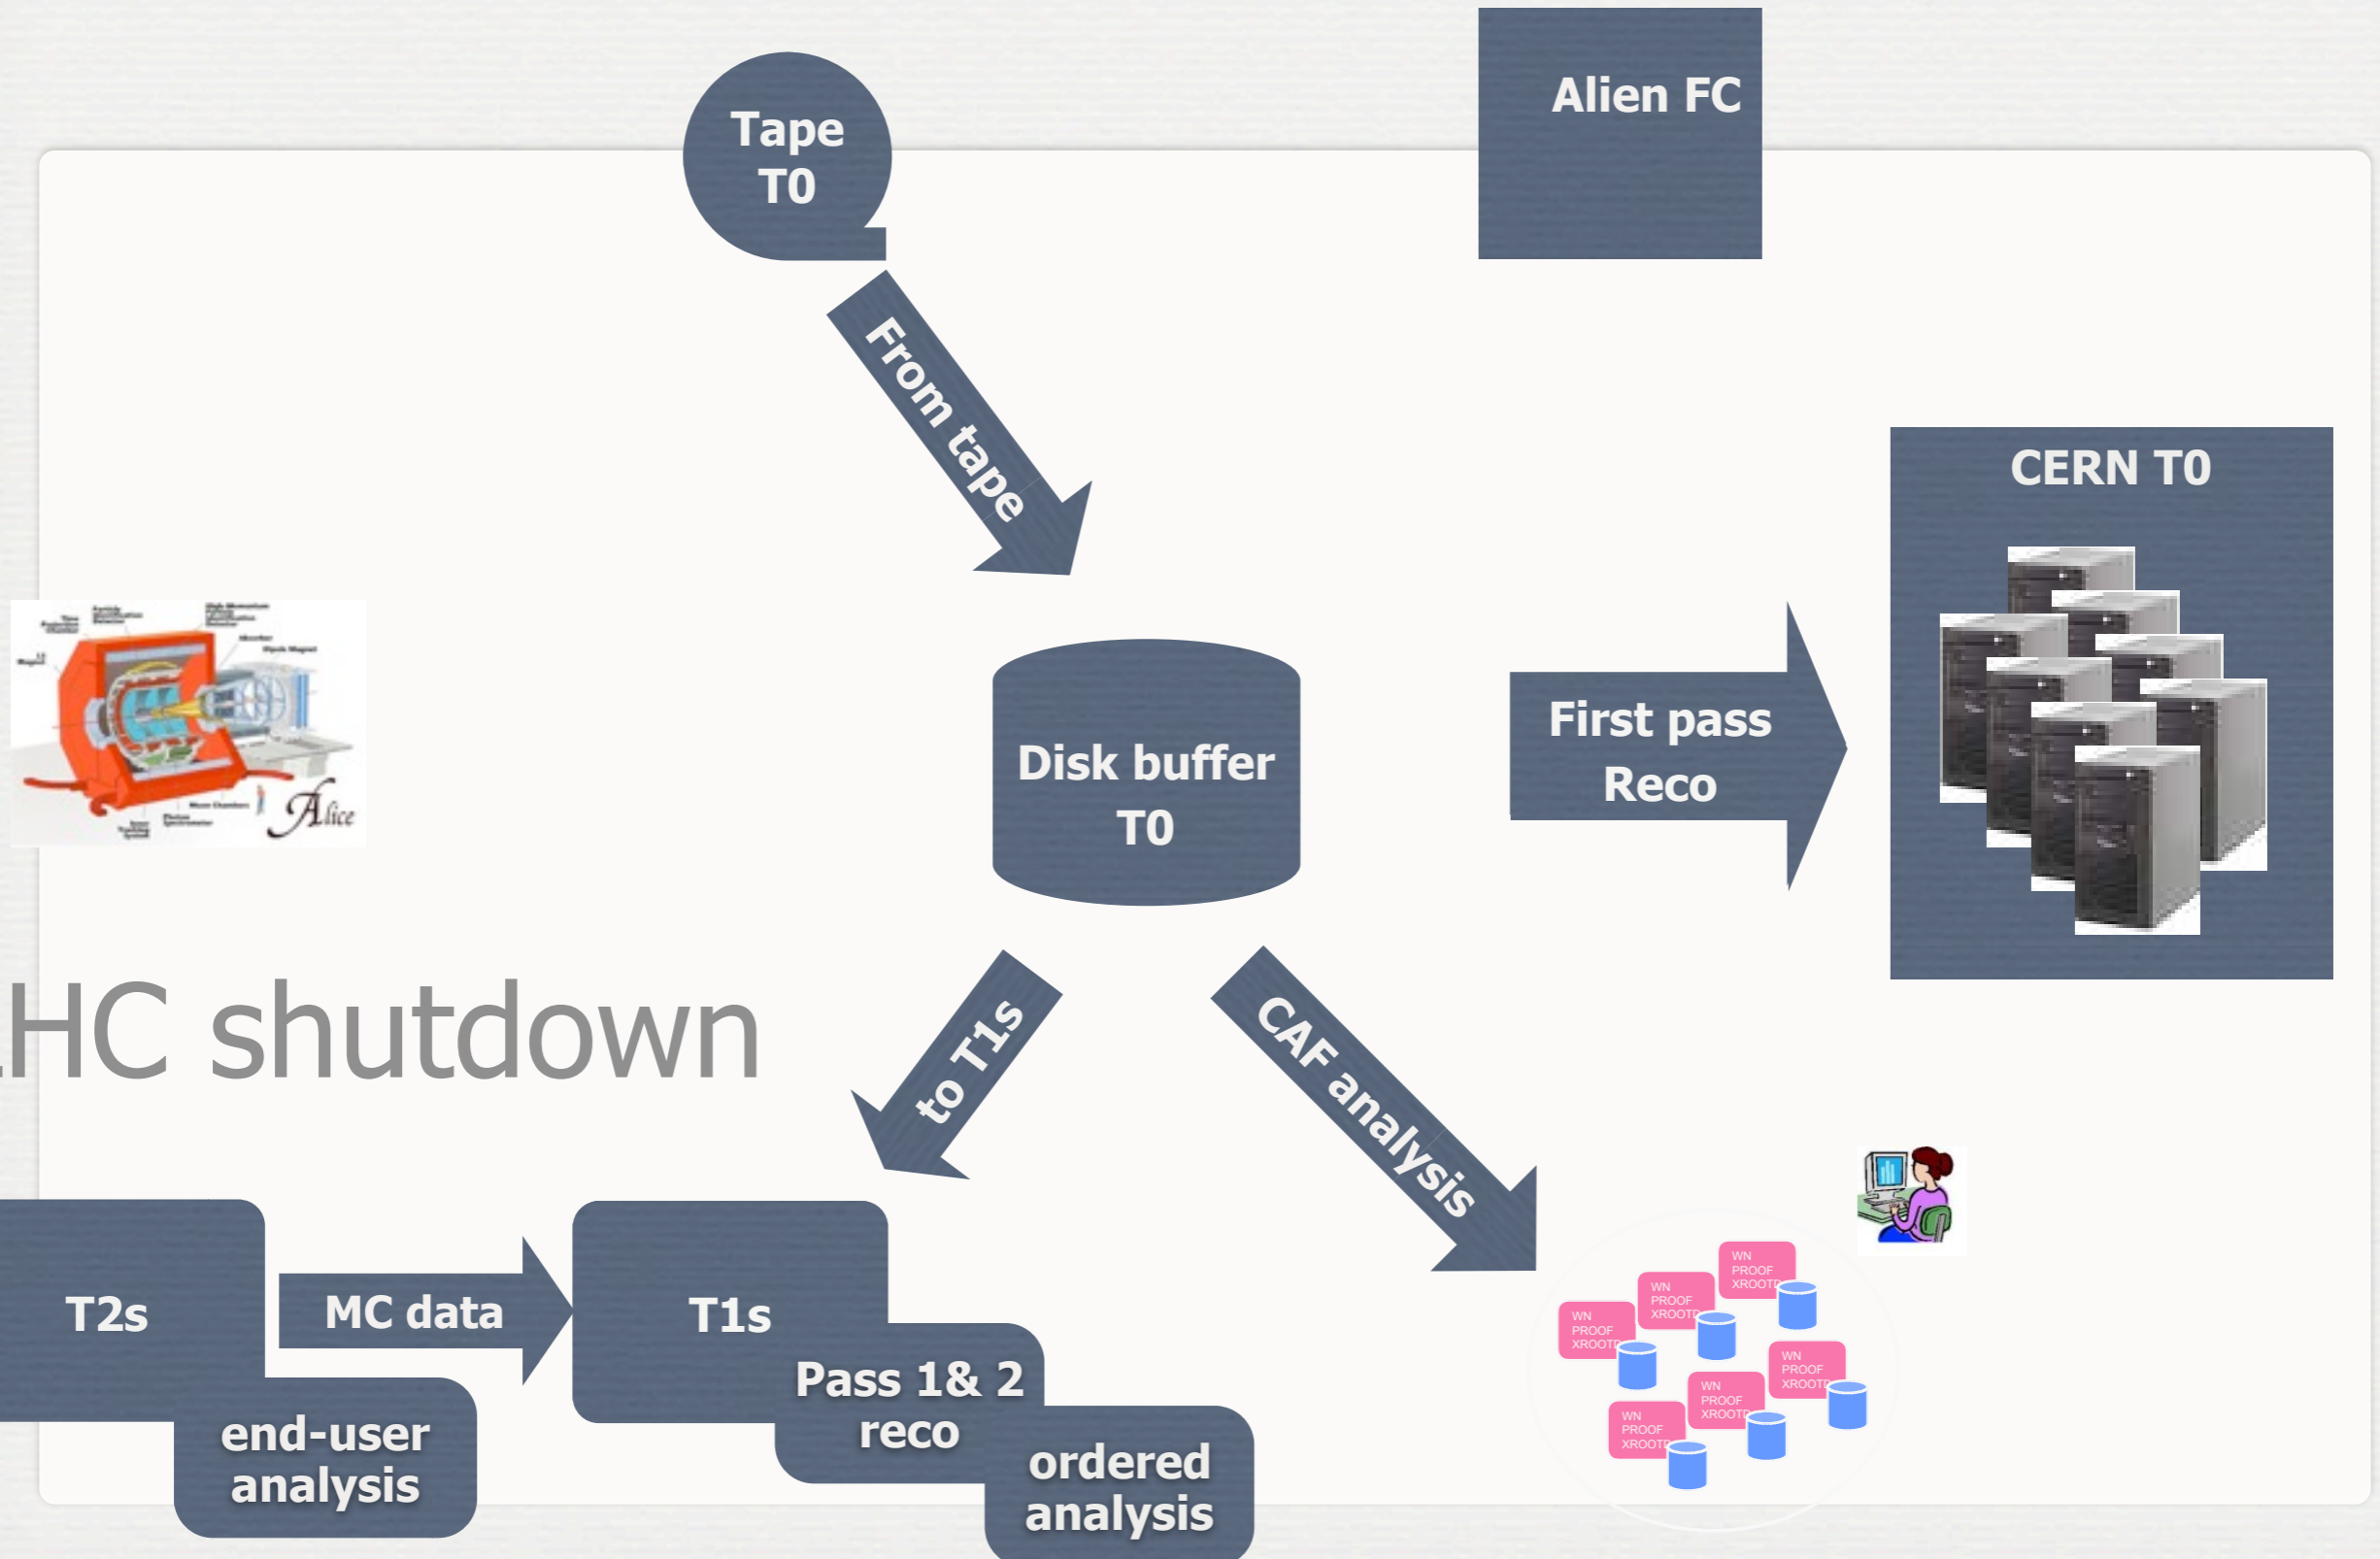
COMPUTING MODEL - AA



COMPUTING MODEL - AA

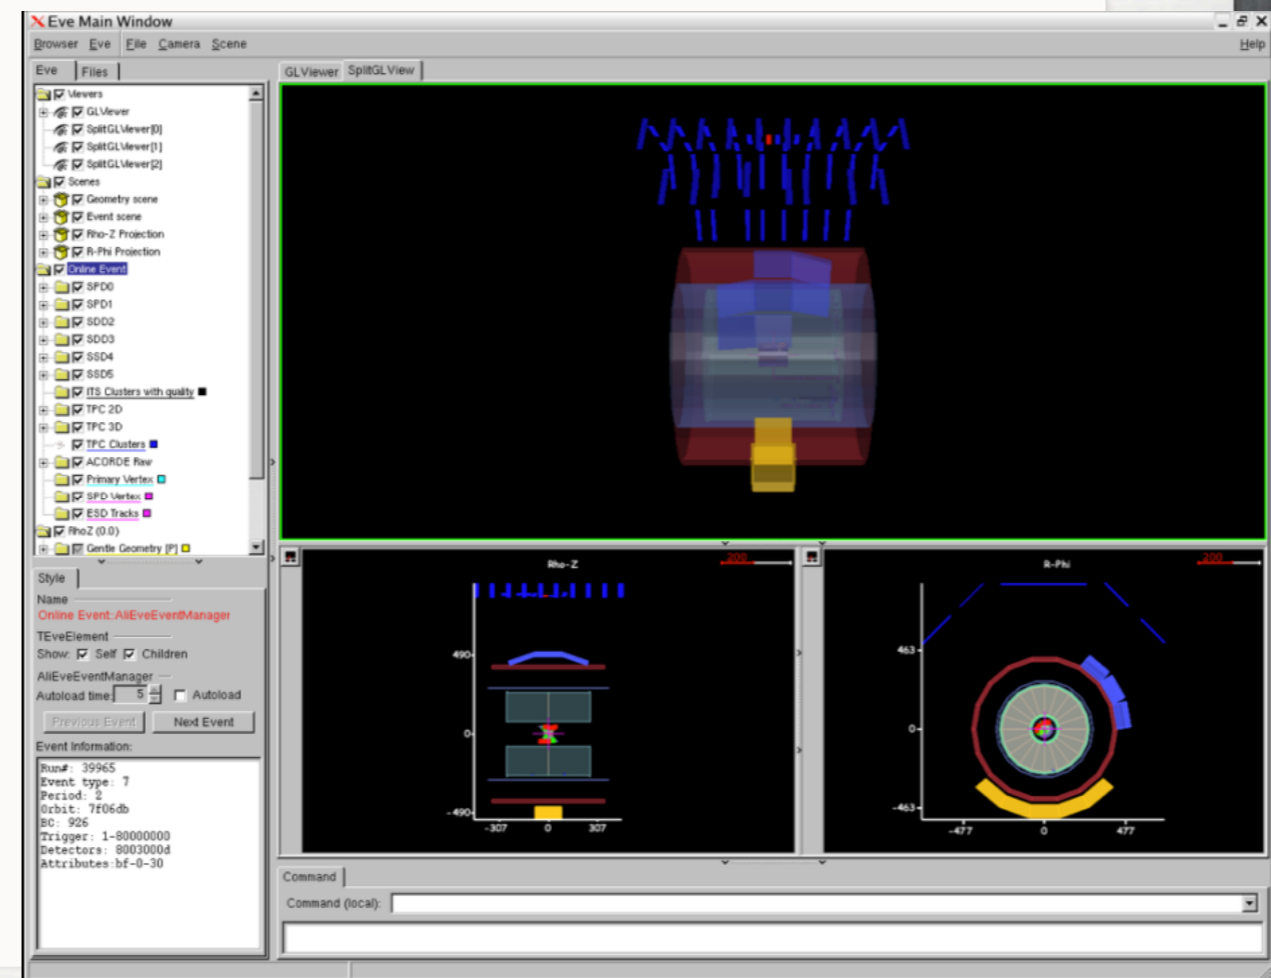
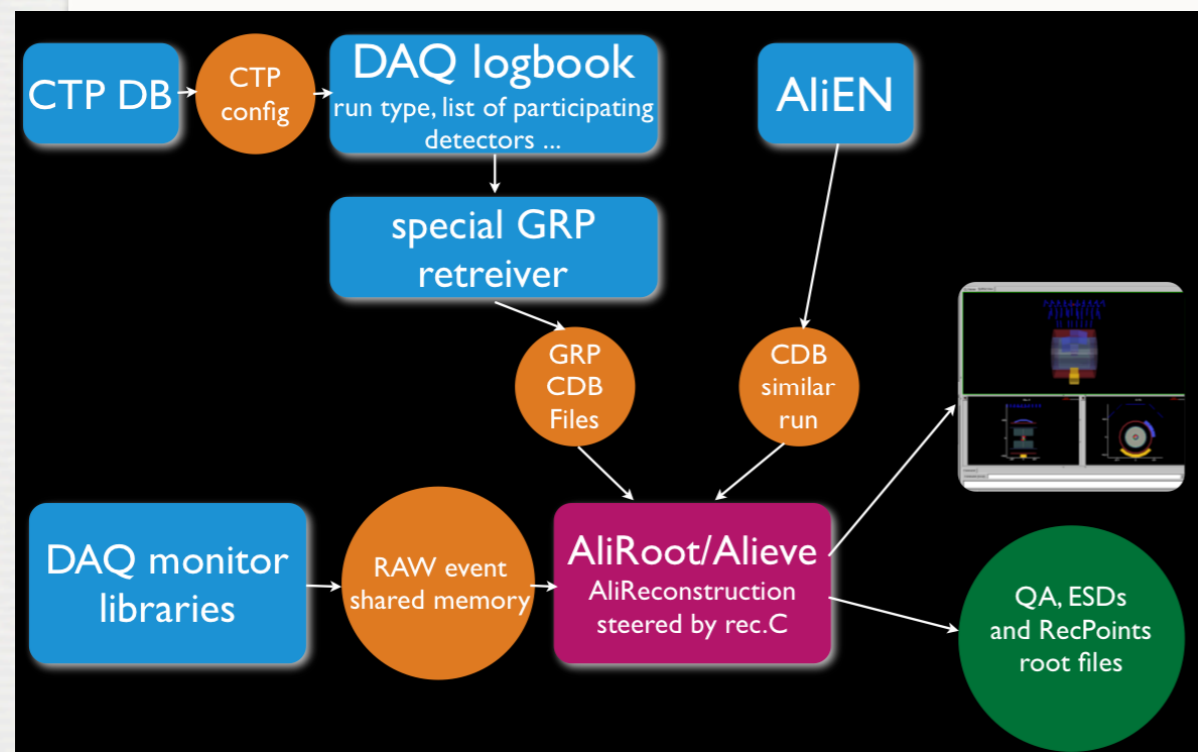
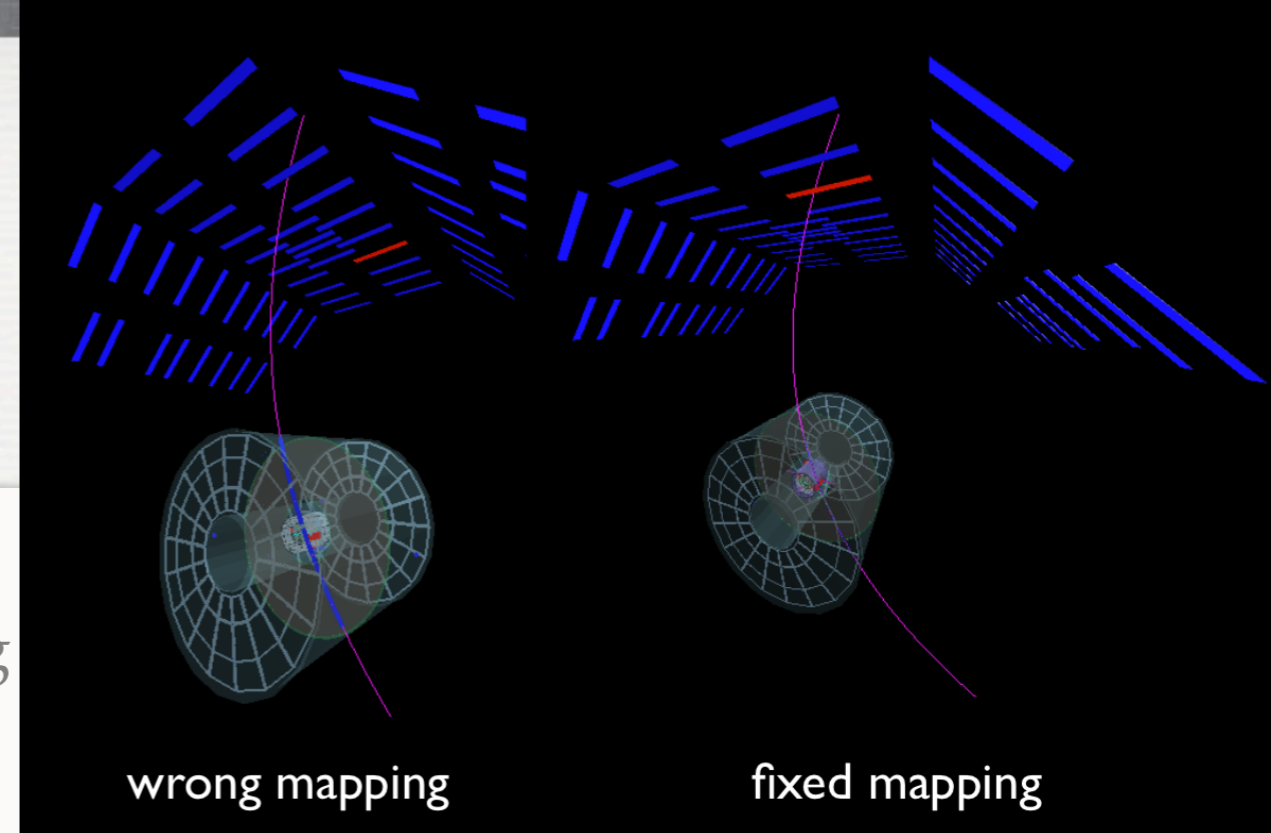


COMPUTING MODEL - AA

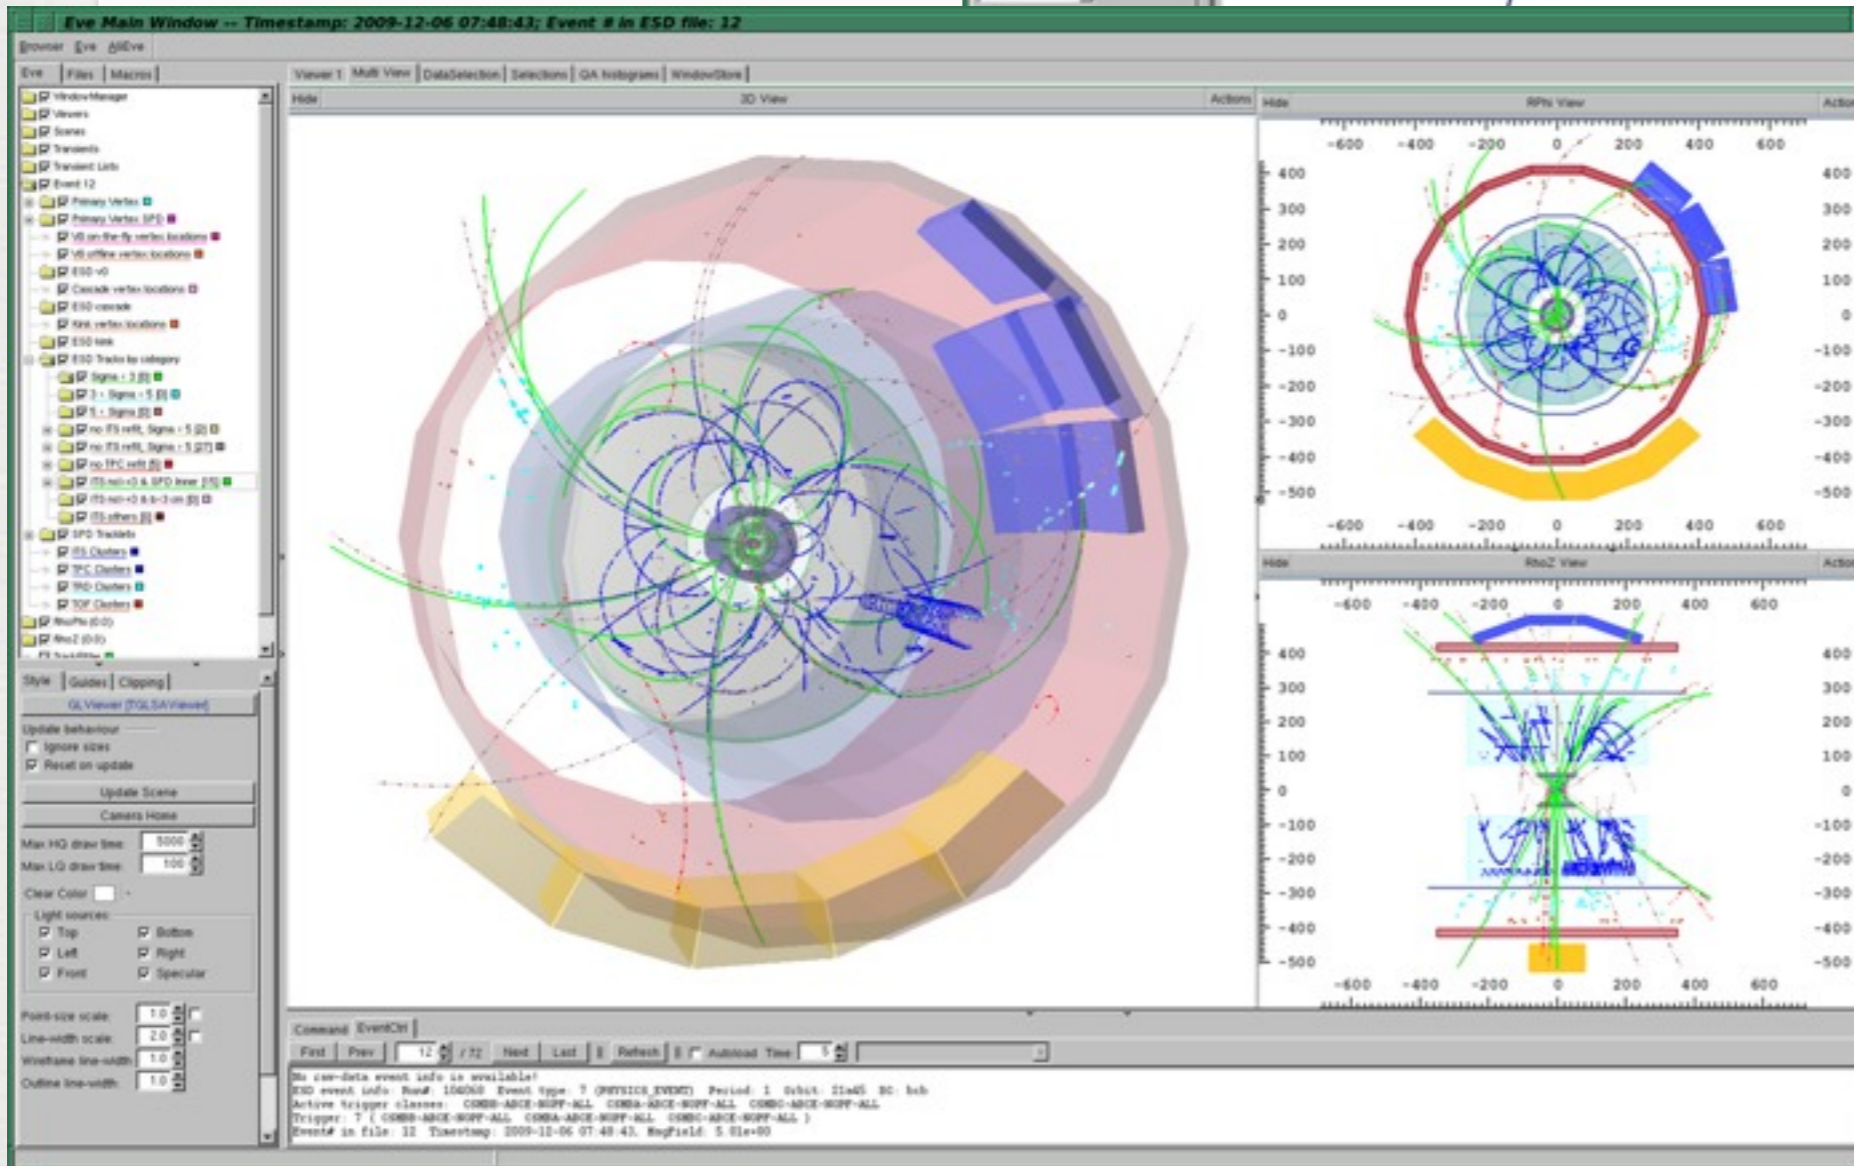
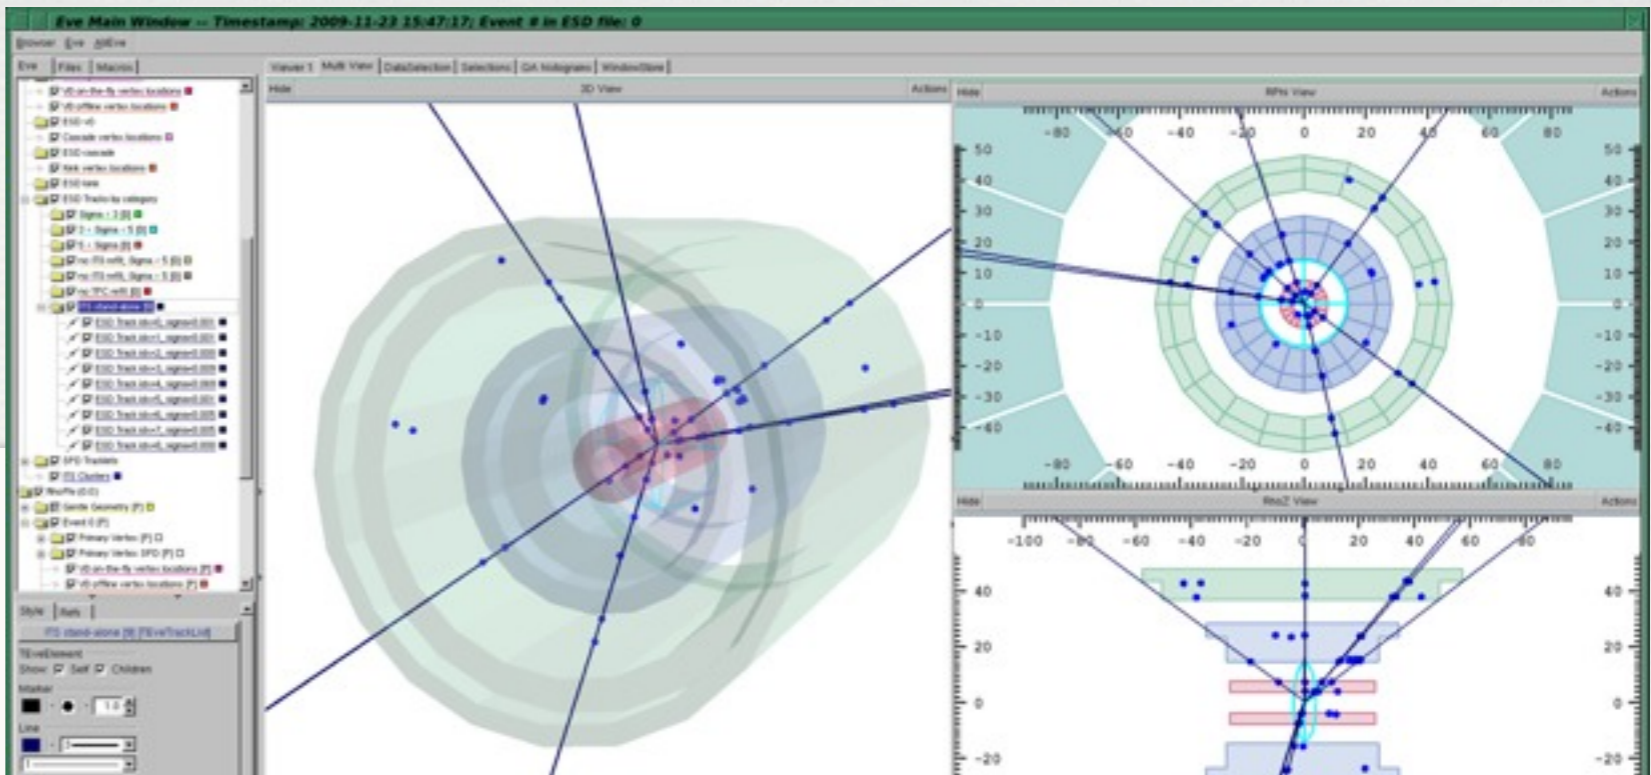


PROMPT RECONSTRUCTION

- Very useful for high-level QA and debugging
- Integrated in the AliEVE event display
- Full Offline code sampling events directly from DAQ memory


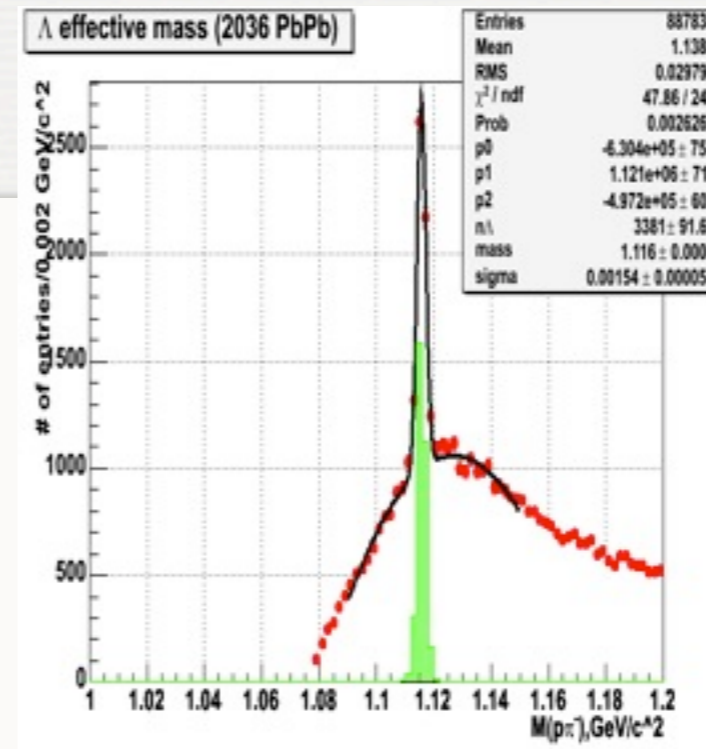


VISUALISATION




ALICE ANALYSIS BASIC CONCEPTS

- Analysis Models
 - Prompt data processing (calib, align, reco, analysis) @CERN with PROOF
 - Batch Analysis using GRID infrastructure
 - Local analysis
 - Interactive analysis PROOF+GRID
- User Interface
 - Access GRID via AliEn or ROOT UIs
- PROOF/ROOT
 - Enabling technology for (C)AF
 - GRID API class TAliEn
- Analysis Object Data contain only data needed for a particular analysis
 - Extensible with Δ -AODs
- Same user code local, on CAF and Grid
- Work on the distributed infrastructure has been done by the ARDA project



ANALYSIS TRAIN



- AOD production will be organized in a 'train' of tasks
 - To maximize efficiency of full dataset processing
 - To optimize CPU/IO
 - Using the analysis framework

USER ANALYSIS ACTIVITIES

- Generally successful
- User jobs priorities are well mastered in AliEn
- Simple priority scheduling seem to work well, will be expanded soon to “pay for what you use” principles
- Storage remains a weak point
 - Only the lack of it – available amount does not allow full (as per computing model) replication of data
- As of today, ~250 registered, ~150 active Grid users
 - Not counting the MC/RAW production and CAF


ANALYSIS ON THE GRID

Production info				Jobs status									
		- Any -											
ID	Tag	Status	Done%	Cfg	Out	Total	Done	Active	Waiting	Runs	Output events	Production description	Comment
782	QA45_LHC10hPbPb	Running	0%			7905	0	149	7600	146 (136833 - 139517)		QA45_LHC10hPbPb: PWG1 QA train	
775	FILTER_p-p_034_LHC10c	Completed	100%			123	123			2 (120822 - 121040)	6,880,957	FILTER_p-p_034_LHC10c: tenders w. V0, TOF corrections -> AODstd(+jets), vertexing, muons	
774	FILTER_p-p_034_LHC10b	Completed	100%			23	23			2 (117116 - 117220)	3,497,142	FILTER_p-p_034_LHC10b: tenders w. V0, TOF corrections -> AODstd(+jets), vertexing, muons	
773	FILTER_PbPb033_LHC10h	Completed	99%			12920	12910			130 (136833 - 139517)	14,607,156	FILTER_PbPb033_LHC10h: tenders w. TOF corrections, centrality, AODstd(+jets), vertexing_highmult	
771	FILTER_PbPb032_LHC11a3	Completed	98%			2116	2078			5 (137161 - 137243)	211,770	FILTER_PbPb032_LHC11a3: centrality, stdAOD(+jets)/vertexing	
769	QA46_LHC10e_p-p_Stage3	Completed	88%			27	24			28 (127819 - 130848)	27,227,839	QA46_LHC10e_p-p_Stage3: PWG1 QA train_FinalMerging	
768	QA46_LHC10e_p-p_Stage2	Completed	94%			295	280			29 (127819 - 130848)		QA46_LHC10e_p-p_Stage2: PWG1 QA train_Merging	
767	QA46_LHC10e_p-p_Stage1	Completed	91%			537	491			30 (127819 - 130848)		QA46_LHC10e_p-p_Stage1: PWG1 QA train_Merging	
766	QA46_LHC10e_p-p	Running	88%			1421	1261	2	1	32 (127817 - 130848)	28,512,577	QA46_LHC10e_p-p: PWG1 QA train	
765	QA44_LHC10epp_Stage3	Completed	100%			40	40			1 (129519 - 129519)	343,872	QA44_LHC10epp_Stage3: PWG1 QA train (no TPC)_FinalMerging	
764	QA44_LHC10epp_Stage2	Completed	85%			7	6			1 (129519 - 129519)		QA44_LHC10epp_Stage2: PWG1 QA train (no TPC)_Merging	
763	FILTER_PbPb031_LHC10h	Running	62%			8945	5615	971	1519	72 (137135 - 139314)	4,437,720	FILTER_PbPb031_LHC10h: tenders w. TOF corrections, centrality, AODstd(+jets), vertexing_highmult	
761	QA44_LHC10epp_Stage1	Completed	100%			13	13			1 (129519 - 129519)		QA44_LHC10epp_Stage1: PWG1 QA train (no TPC)_Merging	
760	QA44_LHC10epp	Completed	95%			40	38			1 (129519 - 129519)	343,872	QA44_LHC10epp: PWG1 QA train (no TPC)	
758	QA43_LHC10h_Stage3	Completed	100%			62	62			62 (137608 - 139517)		QA43_LHC10h_Stage3: PWG1 QA train_FinalMerging	
757	QA43_LHC10h_Stage2	Completed	89%			222	198			67 (137608 - 139517)		QA43_LHC10h_Stage2: PWG1 QA train_Merging	
756	QA43_LHC10h_Stage1	Completed	96%			1583	1533			68 (137161 - 139514)		QA43_LHC10h_Stage1: PWG1 QA train_Merging	
727	QA43_LHC10h	Running	71%			7238	5198	278		69 (137161 - 139517)		QA43_LHC10h: Physics selection, centrality and QAsym	
726	QA37_LHC10h8_spcsim_Merging	Completed	0%			1	0			1 (137161 - 137161)		QA37_LHC10h8_spcsim_Merging: PWG1 QA train_Merging	

ANALYSIS ON THE GRID

Production info				Jobs status									
		- Any -											
ID	Tag	Status	Done%	Cfg	Out	Total	Done	Active	Waiting	Runs	Output events	Production description	Comment
782	QA45_LHC10hPbPb	Running	0%			7905	0	149	7600	146 (136833 - 139517)		QA45_LHC10hPbPb: PWG1 QA train	
775	FILTER_p-p_034_LHC10c	Completed	100%			123	123			2 (120822 - 121040)	6,880,957	FILTER_p-p_034_LHC10c: tenders w. V0, TOF corrections -> AODstd(+jets), vertexing, muons	
774	FILTER_p-p_034_LHC10b	Completed	100%			23	23			2 (117116 - 117220)	3,497,142	FILTER_p-p_034_LHC10b: tenders w. V0, TOF corrections -> AODstd(+jets), vertexing, muons	
773	FILTER_PbPb033_LHC10h	Completed	99%			12920	12910			130 (136833 - 139517)	14,607,156	FILTER_PbPb033_LHC10h: tenders w. TOF corrections, centrality, AODstd(+jets), vertexing_highmult	
771	FILTER_PbPb032_LHC11a3	Completed	98%			2116	2078			5 (137161 - 137243)	211,770	FILTER_PbPb032_LHC11a3: centrality, stdAOD(+jets)/vertexing	
769	QA46_LHC10e_p-p_Stage3	Completed	88%			27	24			28 (127819 - 130848)	27,227,839	QA46_LHC10e_p-p_Stage3: PWG1 QA train	


Running jobs per user




PRODUCTION OF RAW

- Successful despite rapidly changing conditions in the code and detector operation
 - 74 major cycles
 - $7.2 \cdot 10^9$ events (RAW) passed through the reconstruction
 - Processed 3.6PB of data
 - Produced 0.37TB of ESDs and other data

RAW Production Cycles											Processing requests ▶
Production	Description	Status	Run Range	Raw data			Reconstructed			Events	
				Runs	Chunks	Size	Chunks	Size			
LHC10h_ZDC	LHC period LHC10h - ZDC CALIBRATION_EMD	Running	138619 - 139324	2	627	45.96 GB	618	98%	109.7 GB	242%	5,829,923
LHC11a_TRD	TRD Krypton calibration data LHC11a - Pass1	Running	140441 - 140602	18	27,425	13.14 TB	26,817	97%	6.302 TB	49%	267,593,892
LHC10h(PH)	LHC period LHC10h (Pb+Pb) - Run 137161, for PH	Running	-	1							
LHC10e(Pass0-reco)	LHC period LHC10e - Pass0 (reconstruction)	Running	127719 - 130850	142	99,656	260.9 TB	92,114	92%	111.2 GB	0%	0
LHC10e(2)	LHC period LHC10e - Pass2 (with Pass0)	Running	127817 - 130848	33	28,788	75.68 TB	25,151	87%	7.173 TB	10%	68,654,658
LHC10h(PMD)	LHC period LHC10h (Pb+Pb) - Run 137161, for PMD	Completed	137161 - 137161	1	3,355	1.547 TB	2,890	86%	242.6 GB	17%	461,866
LHC10h(TOF)	LHC period LHC10h (Pb+Pb) - Run 137161, for TOF	Completed	137161 - 137161	1	3,355	1.547 TB	3,059	91%	314.3 GB	21%	592,742
LHC10h(vis)	LHC period LHC10h (Pb+Pb) - Golden run, visualisation	Completed	137161 - 137161	1	3,355	1.547 TB	339	10%	119 GB	74%	67,986
LHC10h(5+)	LHC period LHC10h (Pb+Pb) - ITS standalone (5+)	Completed	137161 - 137161	1	3,355	1.547 TB	3,246	96%	663 GB	43%	628,205
LHC10h(4+)	LHC period LHC10h (Pb+Pb) - Pass0+Pass1(4+)	Completed	137161 - 137693	8	21,889	26.64 TB	21,145	96%	7.322 TB	28%	5,631,022
LHC10h(Pass0-reco)	LHC period LHC10h - Pass0 (reconstruction)	Completed	137161 - 139517	152	397,303	764.4 TB	369,320	92%	396 GB	0%	0
LHC10h+++	LHC period LHC10h (Pb+Pb) - Pass1+++	Completed	137161 - 137162	2	6,165	2.835 TB	6,133	99%	1.093 TB	38%	1,315,555
LHC10h++	LHC period LHC10h (Pb+Pb) - Pass1++	Completed	137161 - 137161	1	3,355	1.547 TB	3,344	99%	588.7 GB	37%	646,174
LHC10h+	LHC period LHC10h (Pb+Pb) - Pass1+	Completed	137045 - 137161	2	3,415	1.567 TB	3,397	99%	589.5 GB	36%	1,149,065
LHC10h(TPC)	TPC LASER data LHC10h_TPC - Pass1	Completed	136917 - 136917	1	6	2.425 GB	6	100%	2.153 GB	89%	1,215
LHC10h	LHC period LHC10h (Pb+Pb) - Pass1	Running	136833 - 139517	149	255,658	492.5 TB	205,188	80%	82.11 TB	20%	46,007,678
LHC10g(Pass0-reco)	LHC period LHC10g - Pass0 (reconstruction)	Completed	135780 - 136377	40	49,034	128 TB	48,399	98%	27.49 GB	0%	0
LHC10d(2)	LHC period LHC10d - Pass2 (with Pass0)	Completed	122374 - 126437	93	63,525	166.4 TB	61,761	97%	19.67 TB	12%	231,881,404
LHC10g	LHC period LHC10g - Pass1	Completed	135654 - 136377	108	76,367	178.1 TB	73,678	96%	13.93 TB	8%	193,032,901
LHC10f(Pass 1+)	LHC period LHC10f - Pass 1+	Completed	134297 - 134301	3	1,513	3.965 TB	1,495	98%	203.1 GB	5%	6,452,133
LHC10d(Pass0-reco)	LHC period LHC10d - Pass0 (reconstruction)	Completed	122374 - 126437	61	48,522	127.6 TB	47,868	98%	60.72 GB	0%	0



SENDING JOBS TO DATA



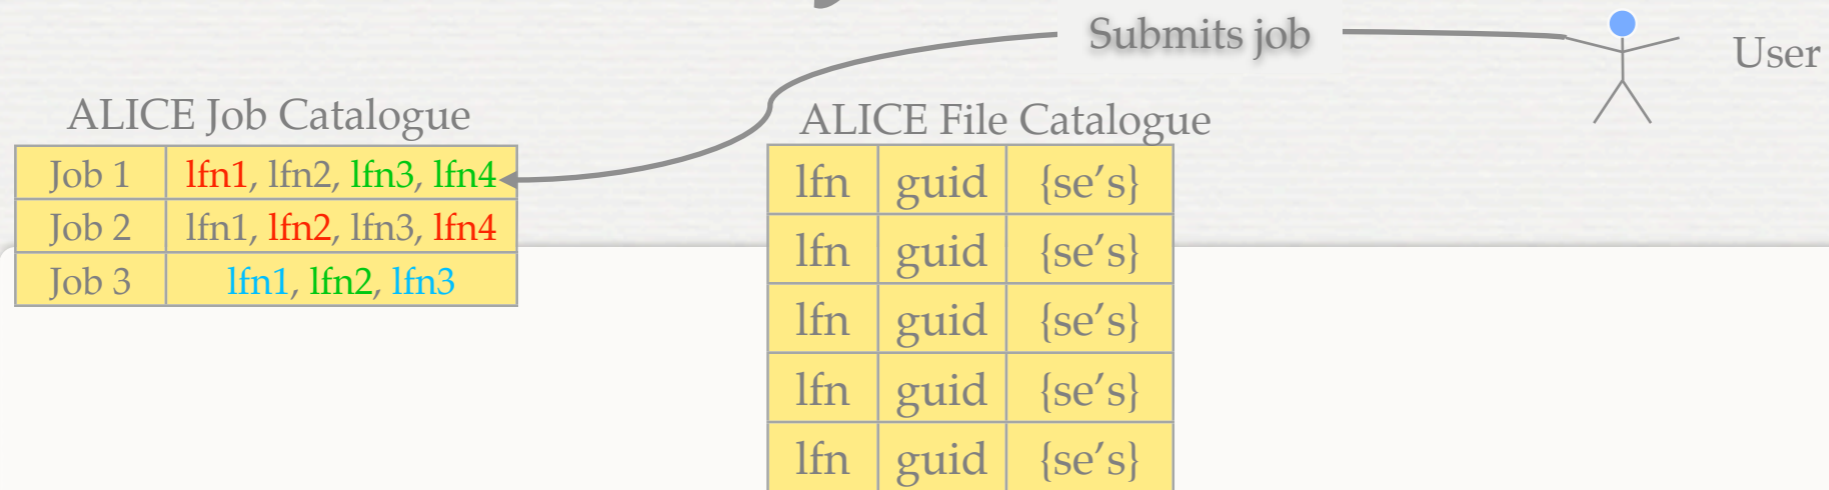
SENDING JOBS TO DATA



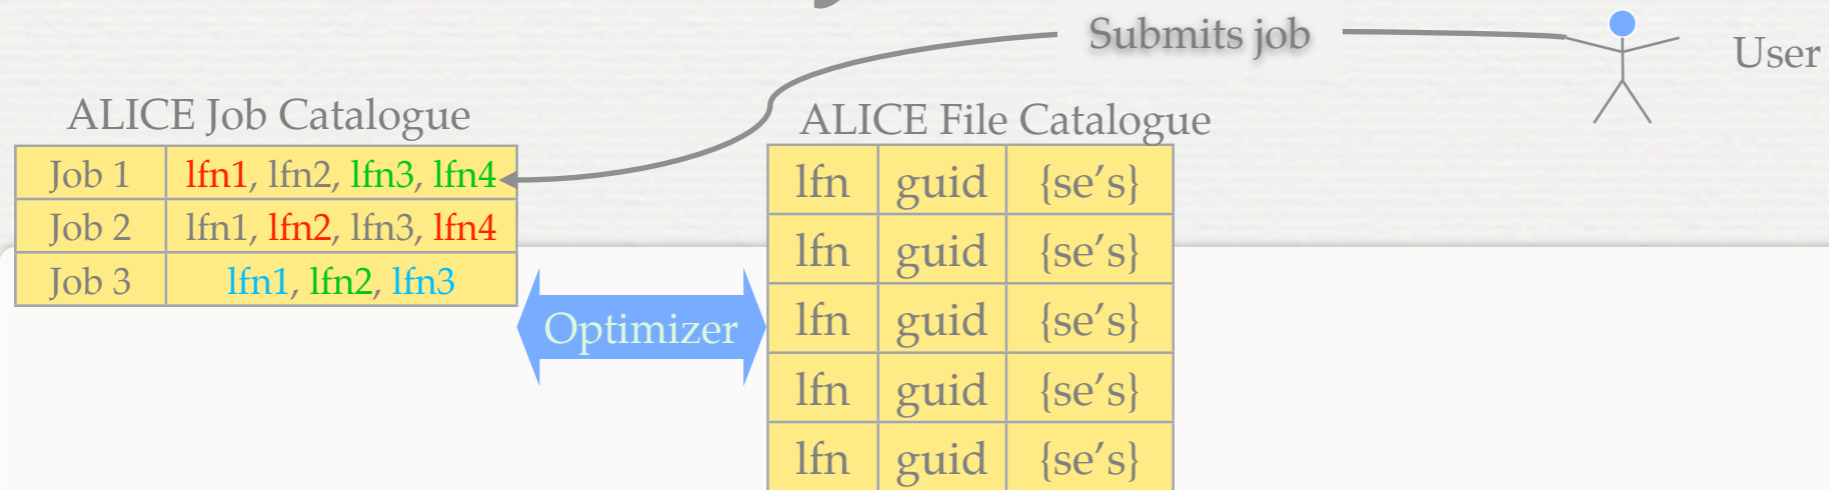
ALICE Job Catalogue

Job 1	lfn1, lfn2, lfn3, lfn4
Job 2	lfn1, lfn2, lfn3, lfn4
Job 3	lfn1, lfn2, lfn3

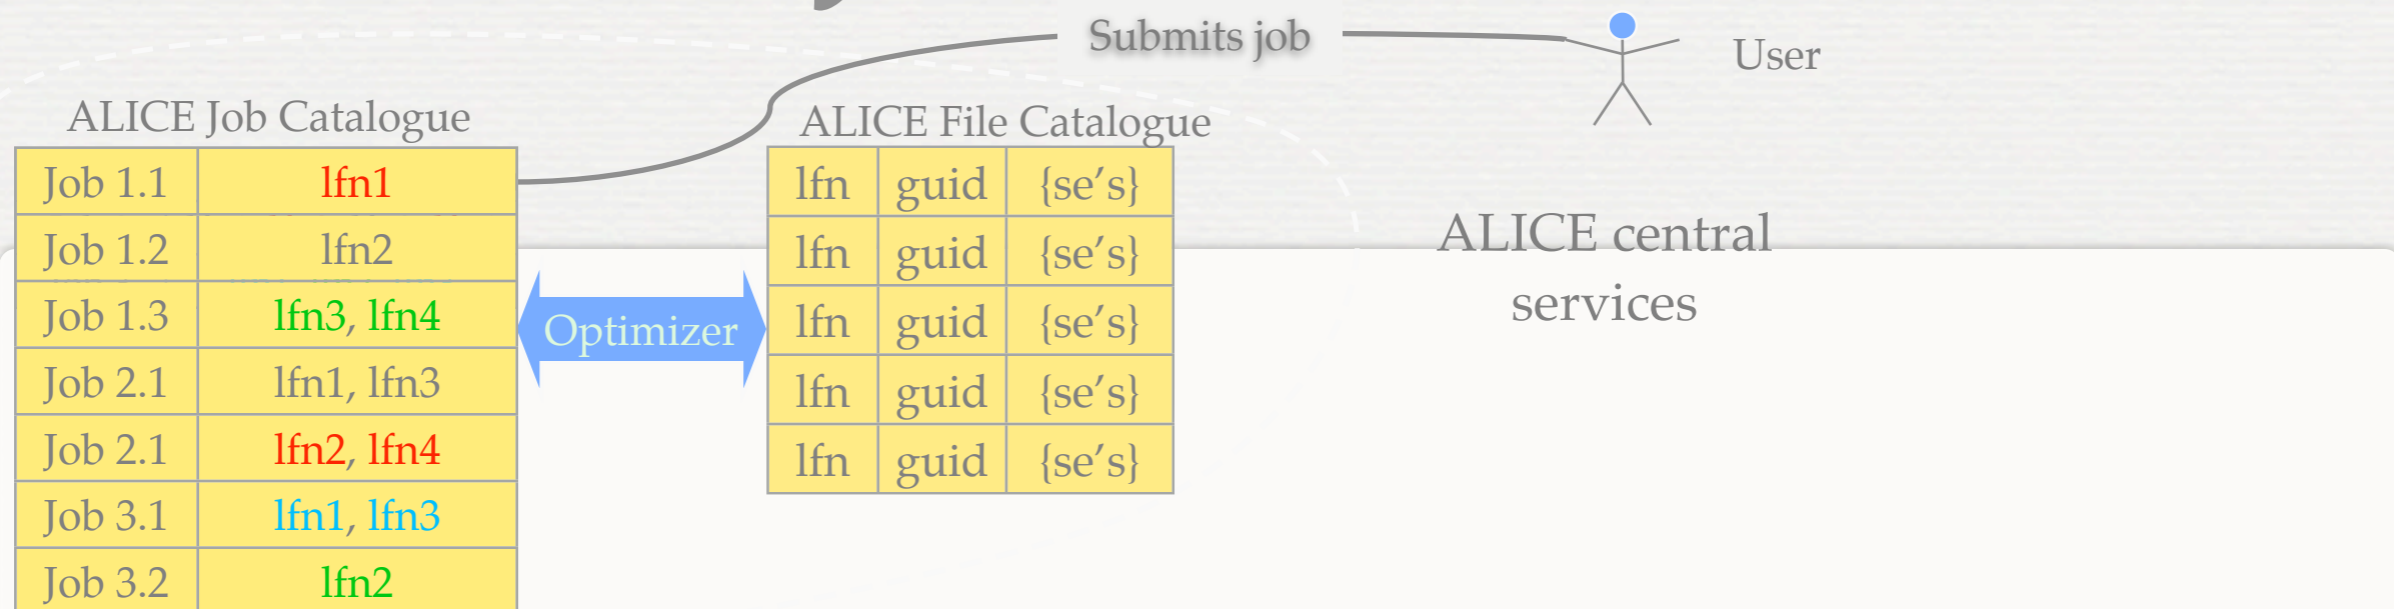
SENDING JOBS TO DATA



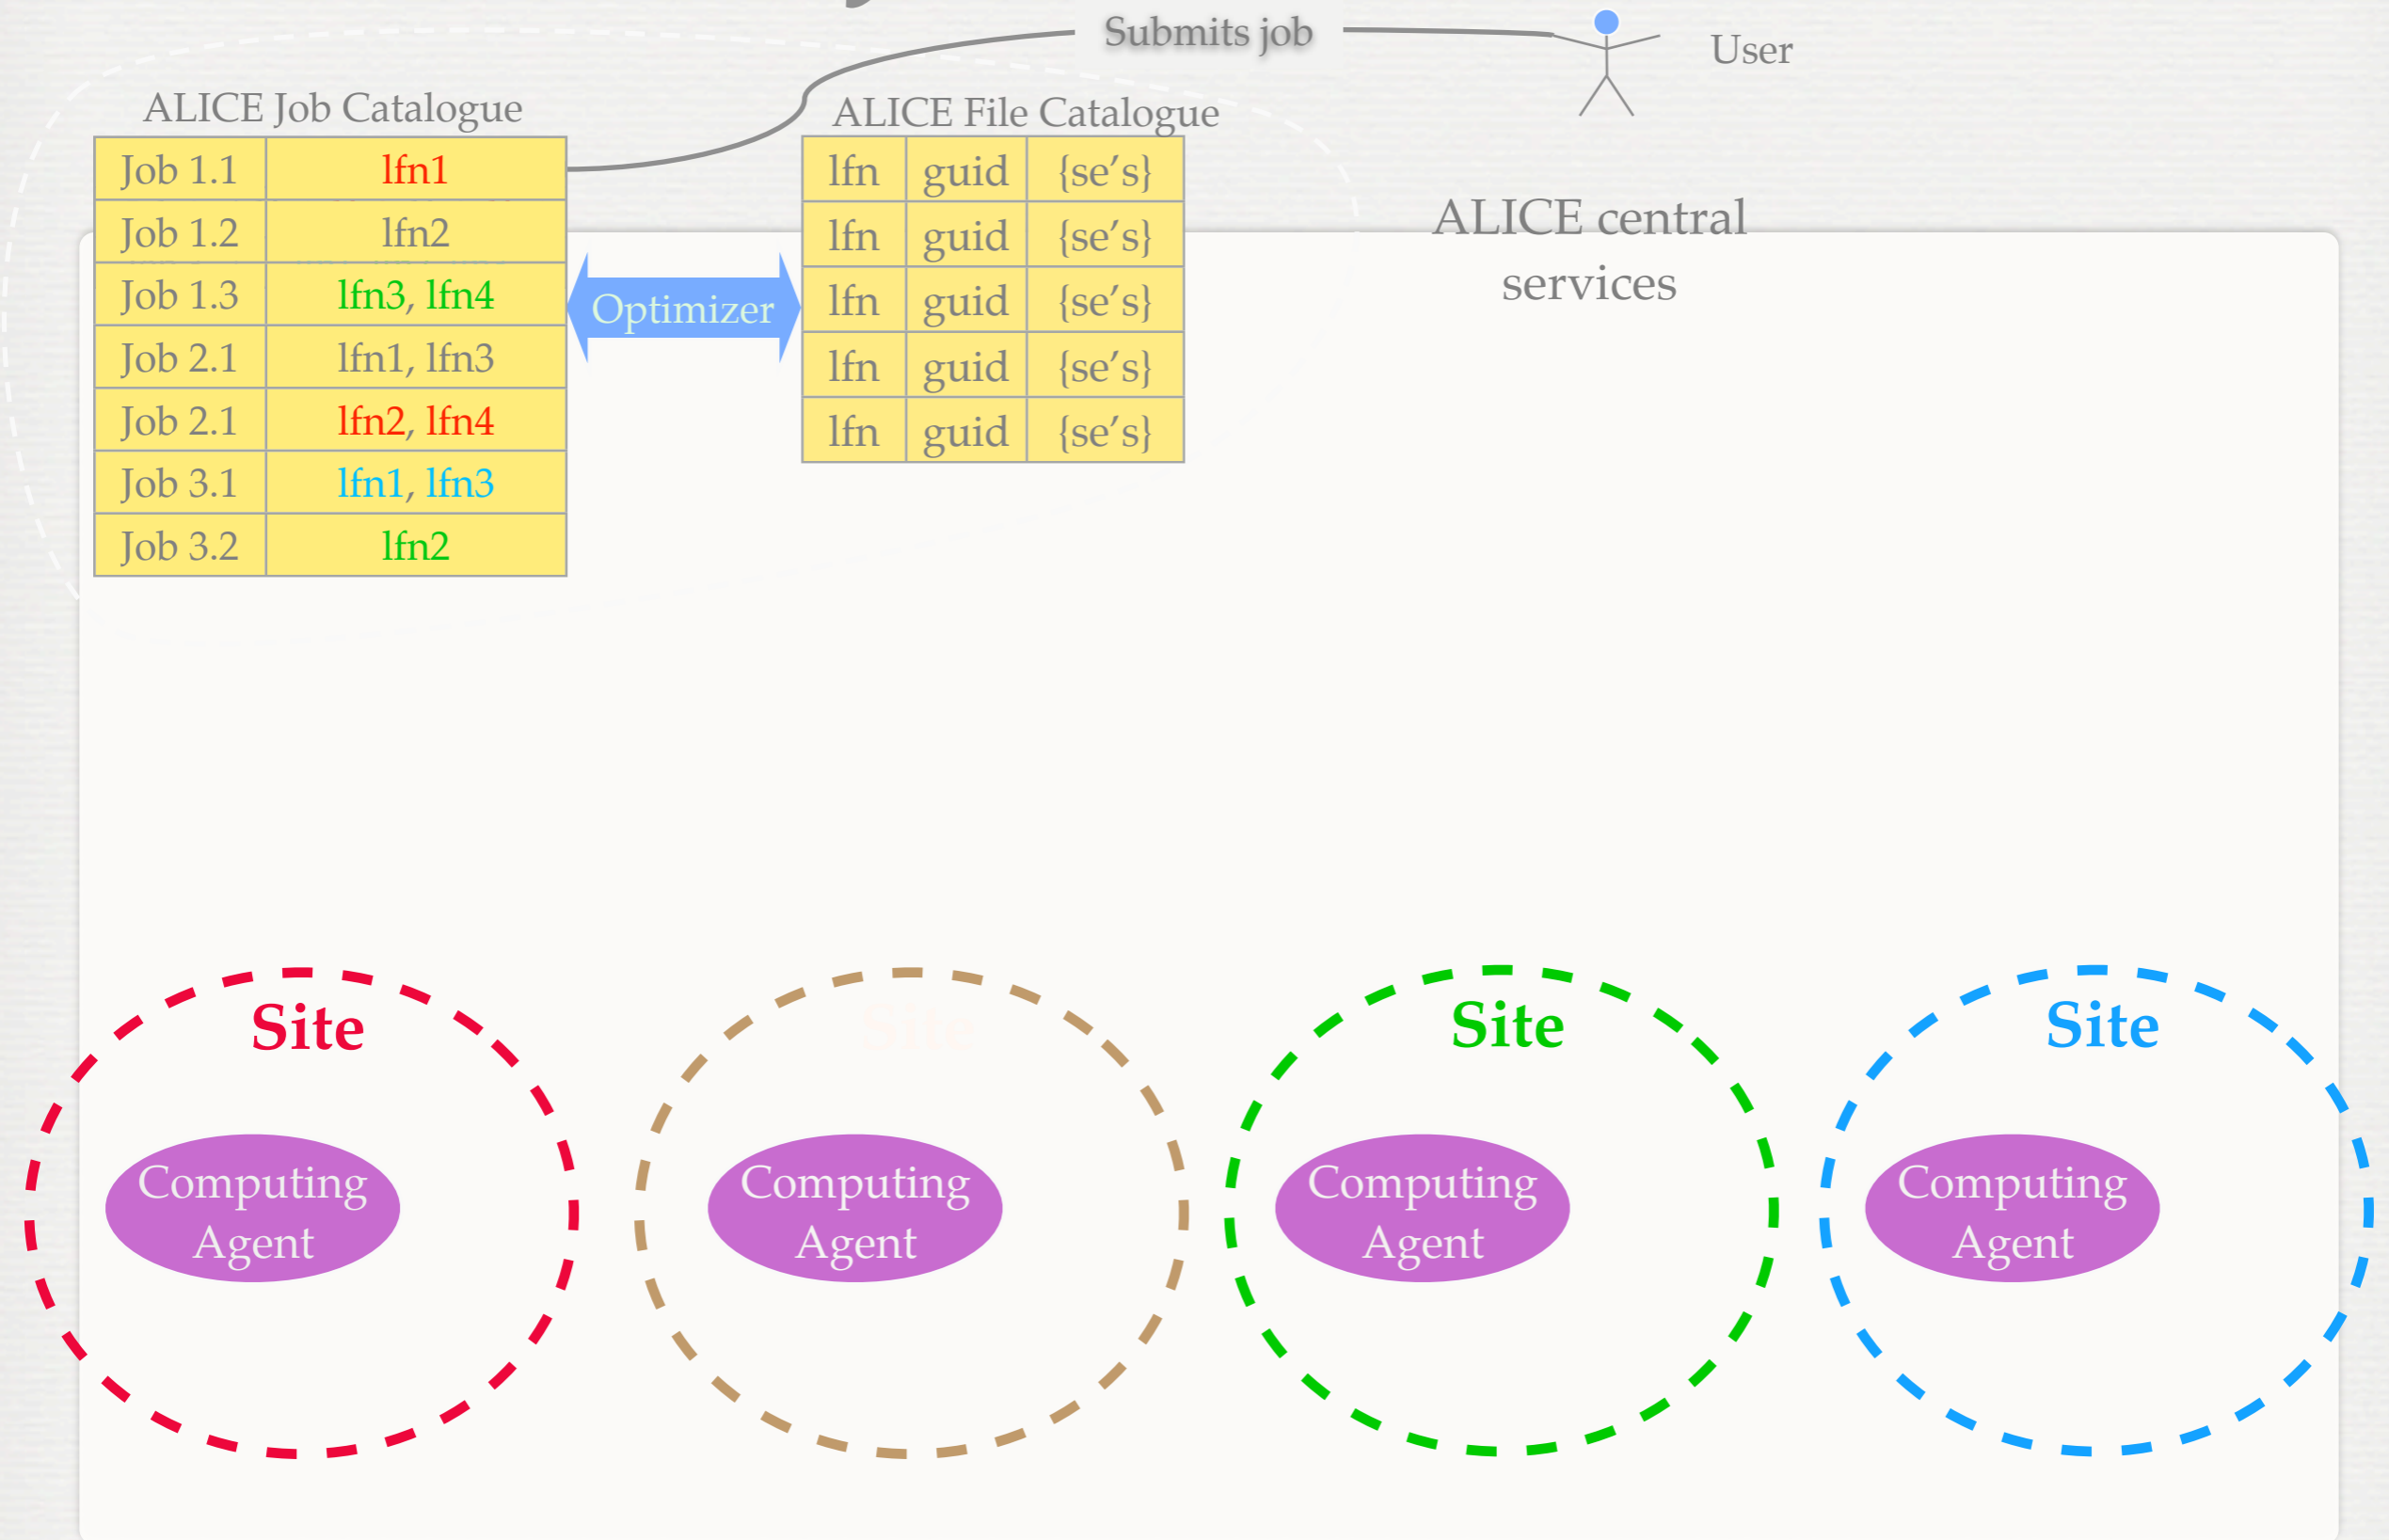
SENDING JOBS TO DATA



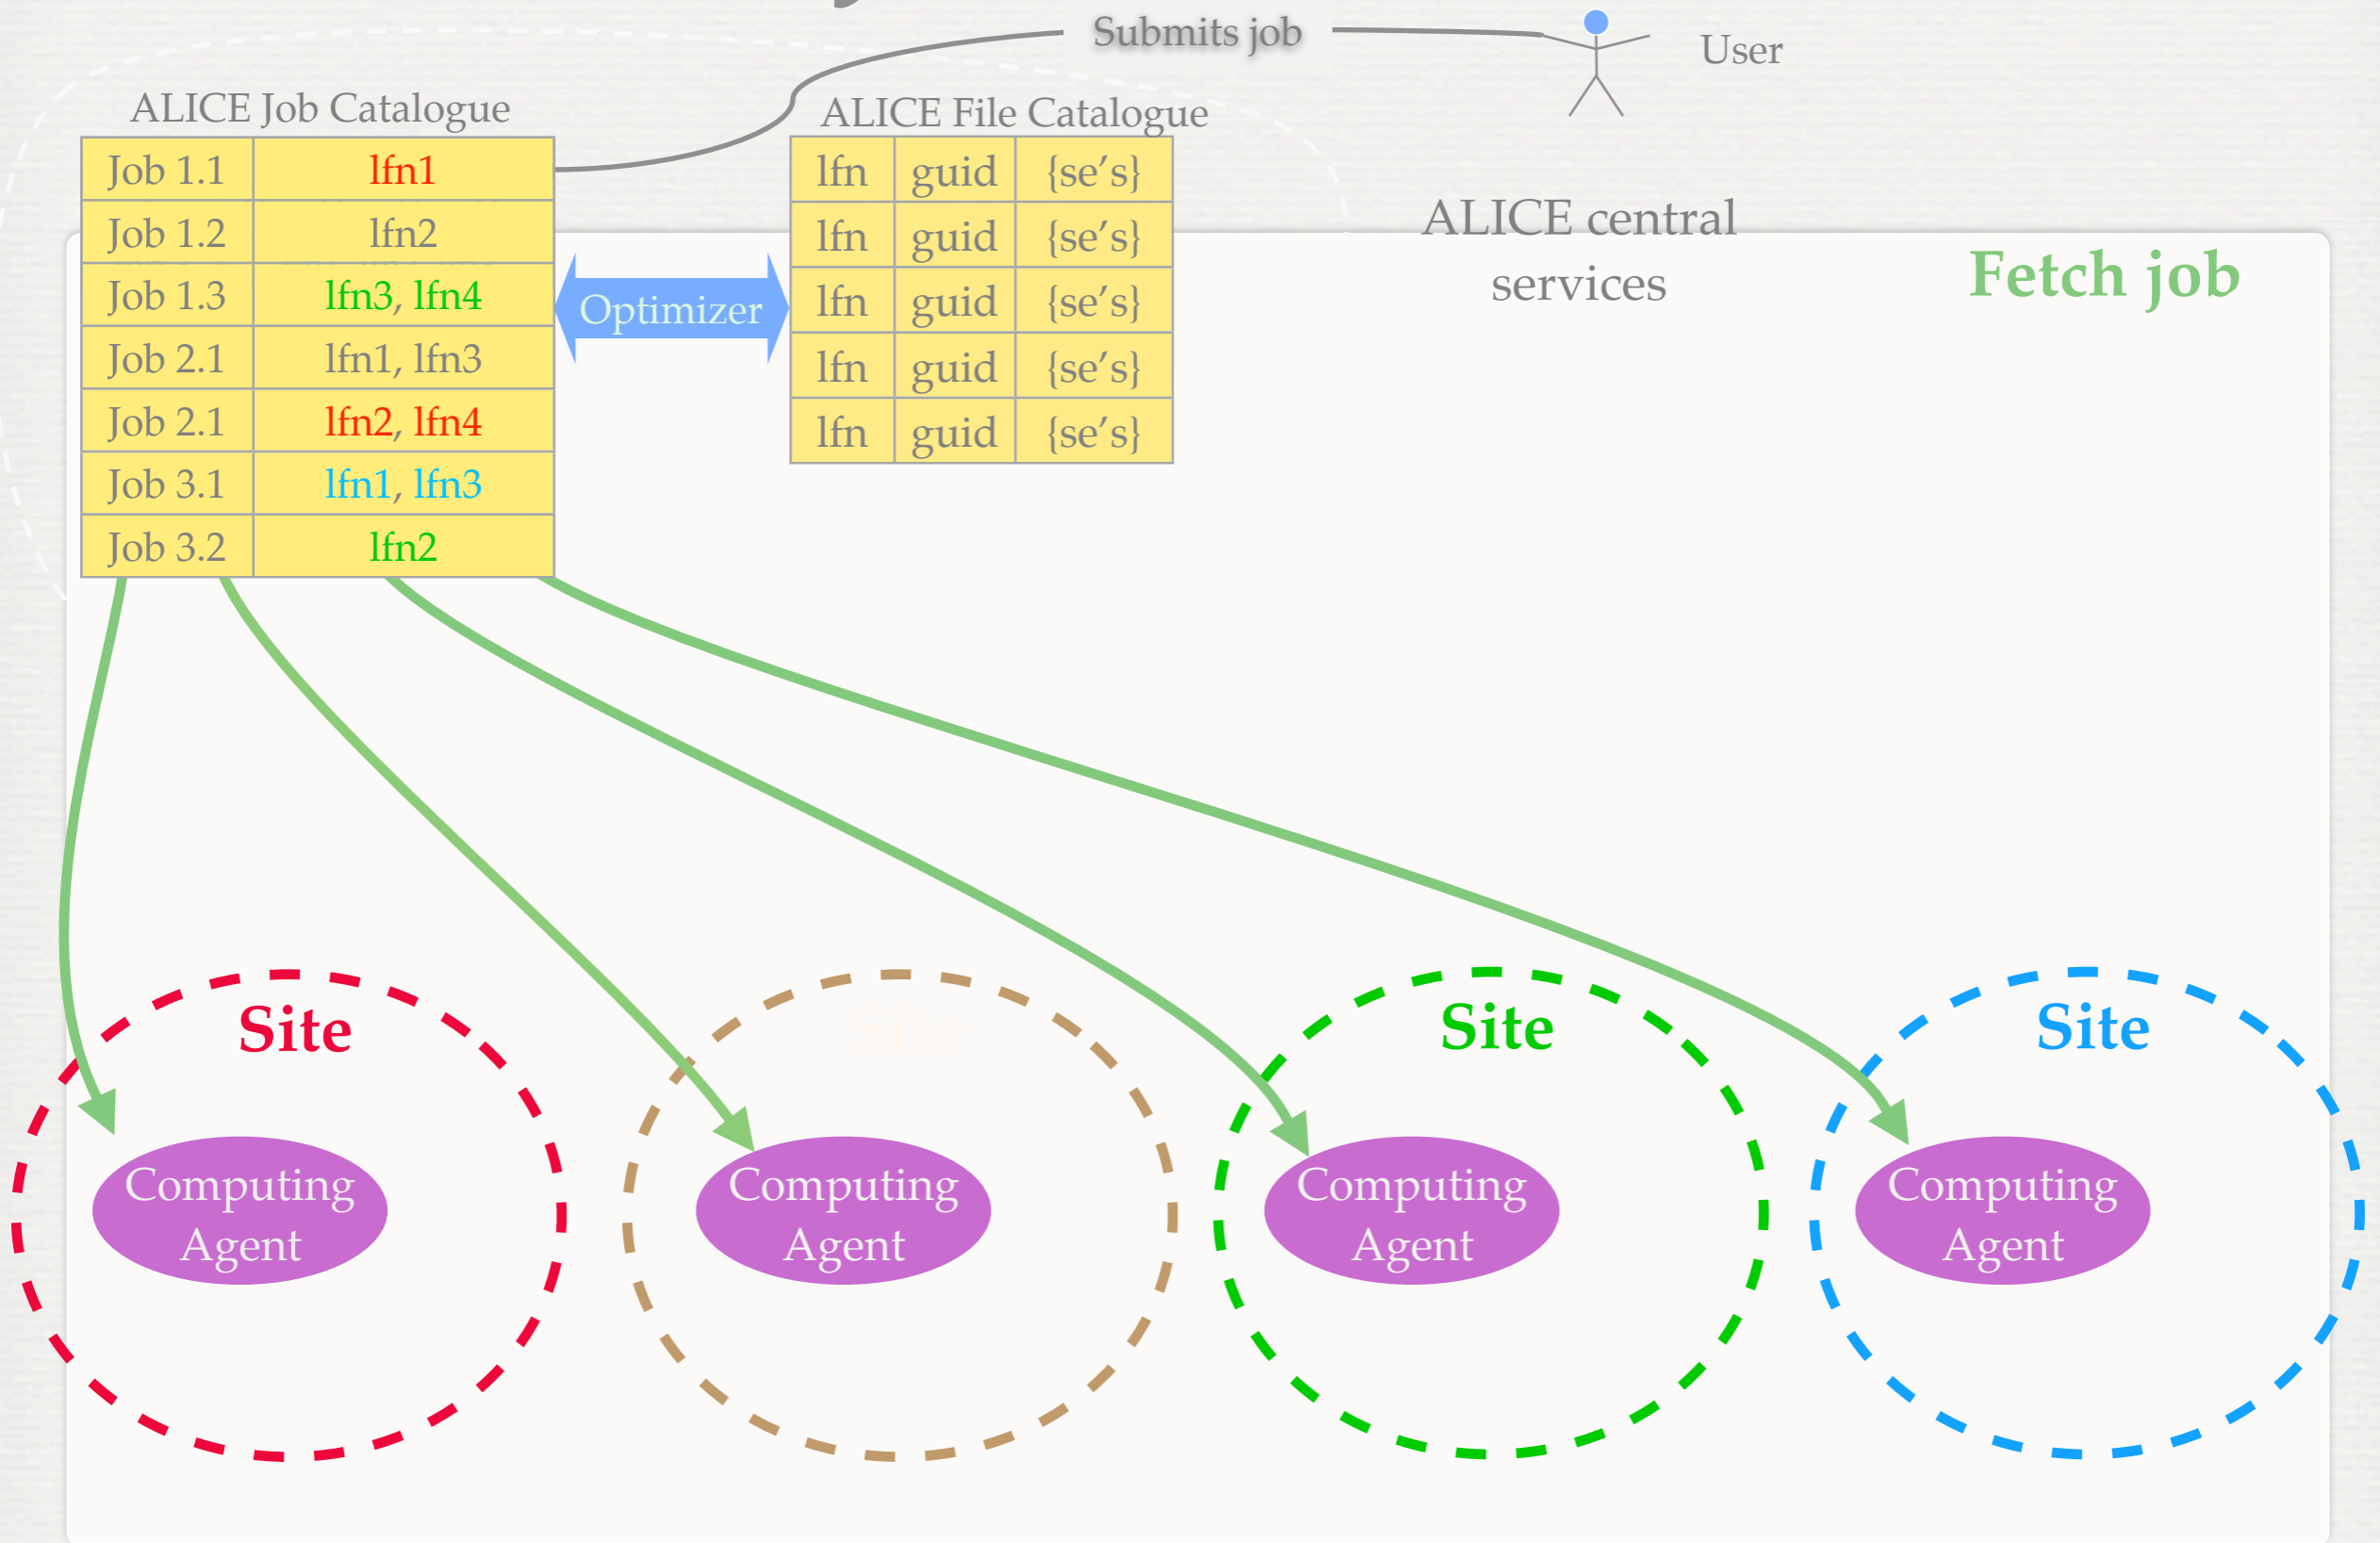
SENDING JOBS TO DATA



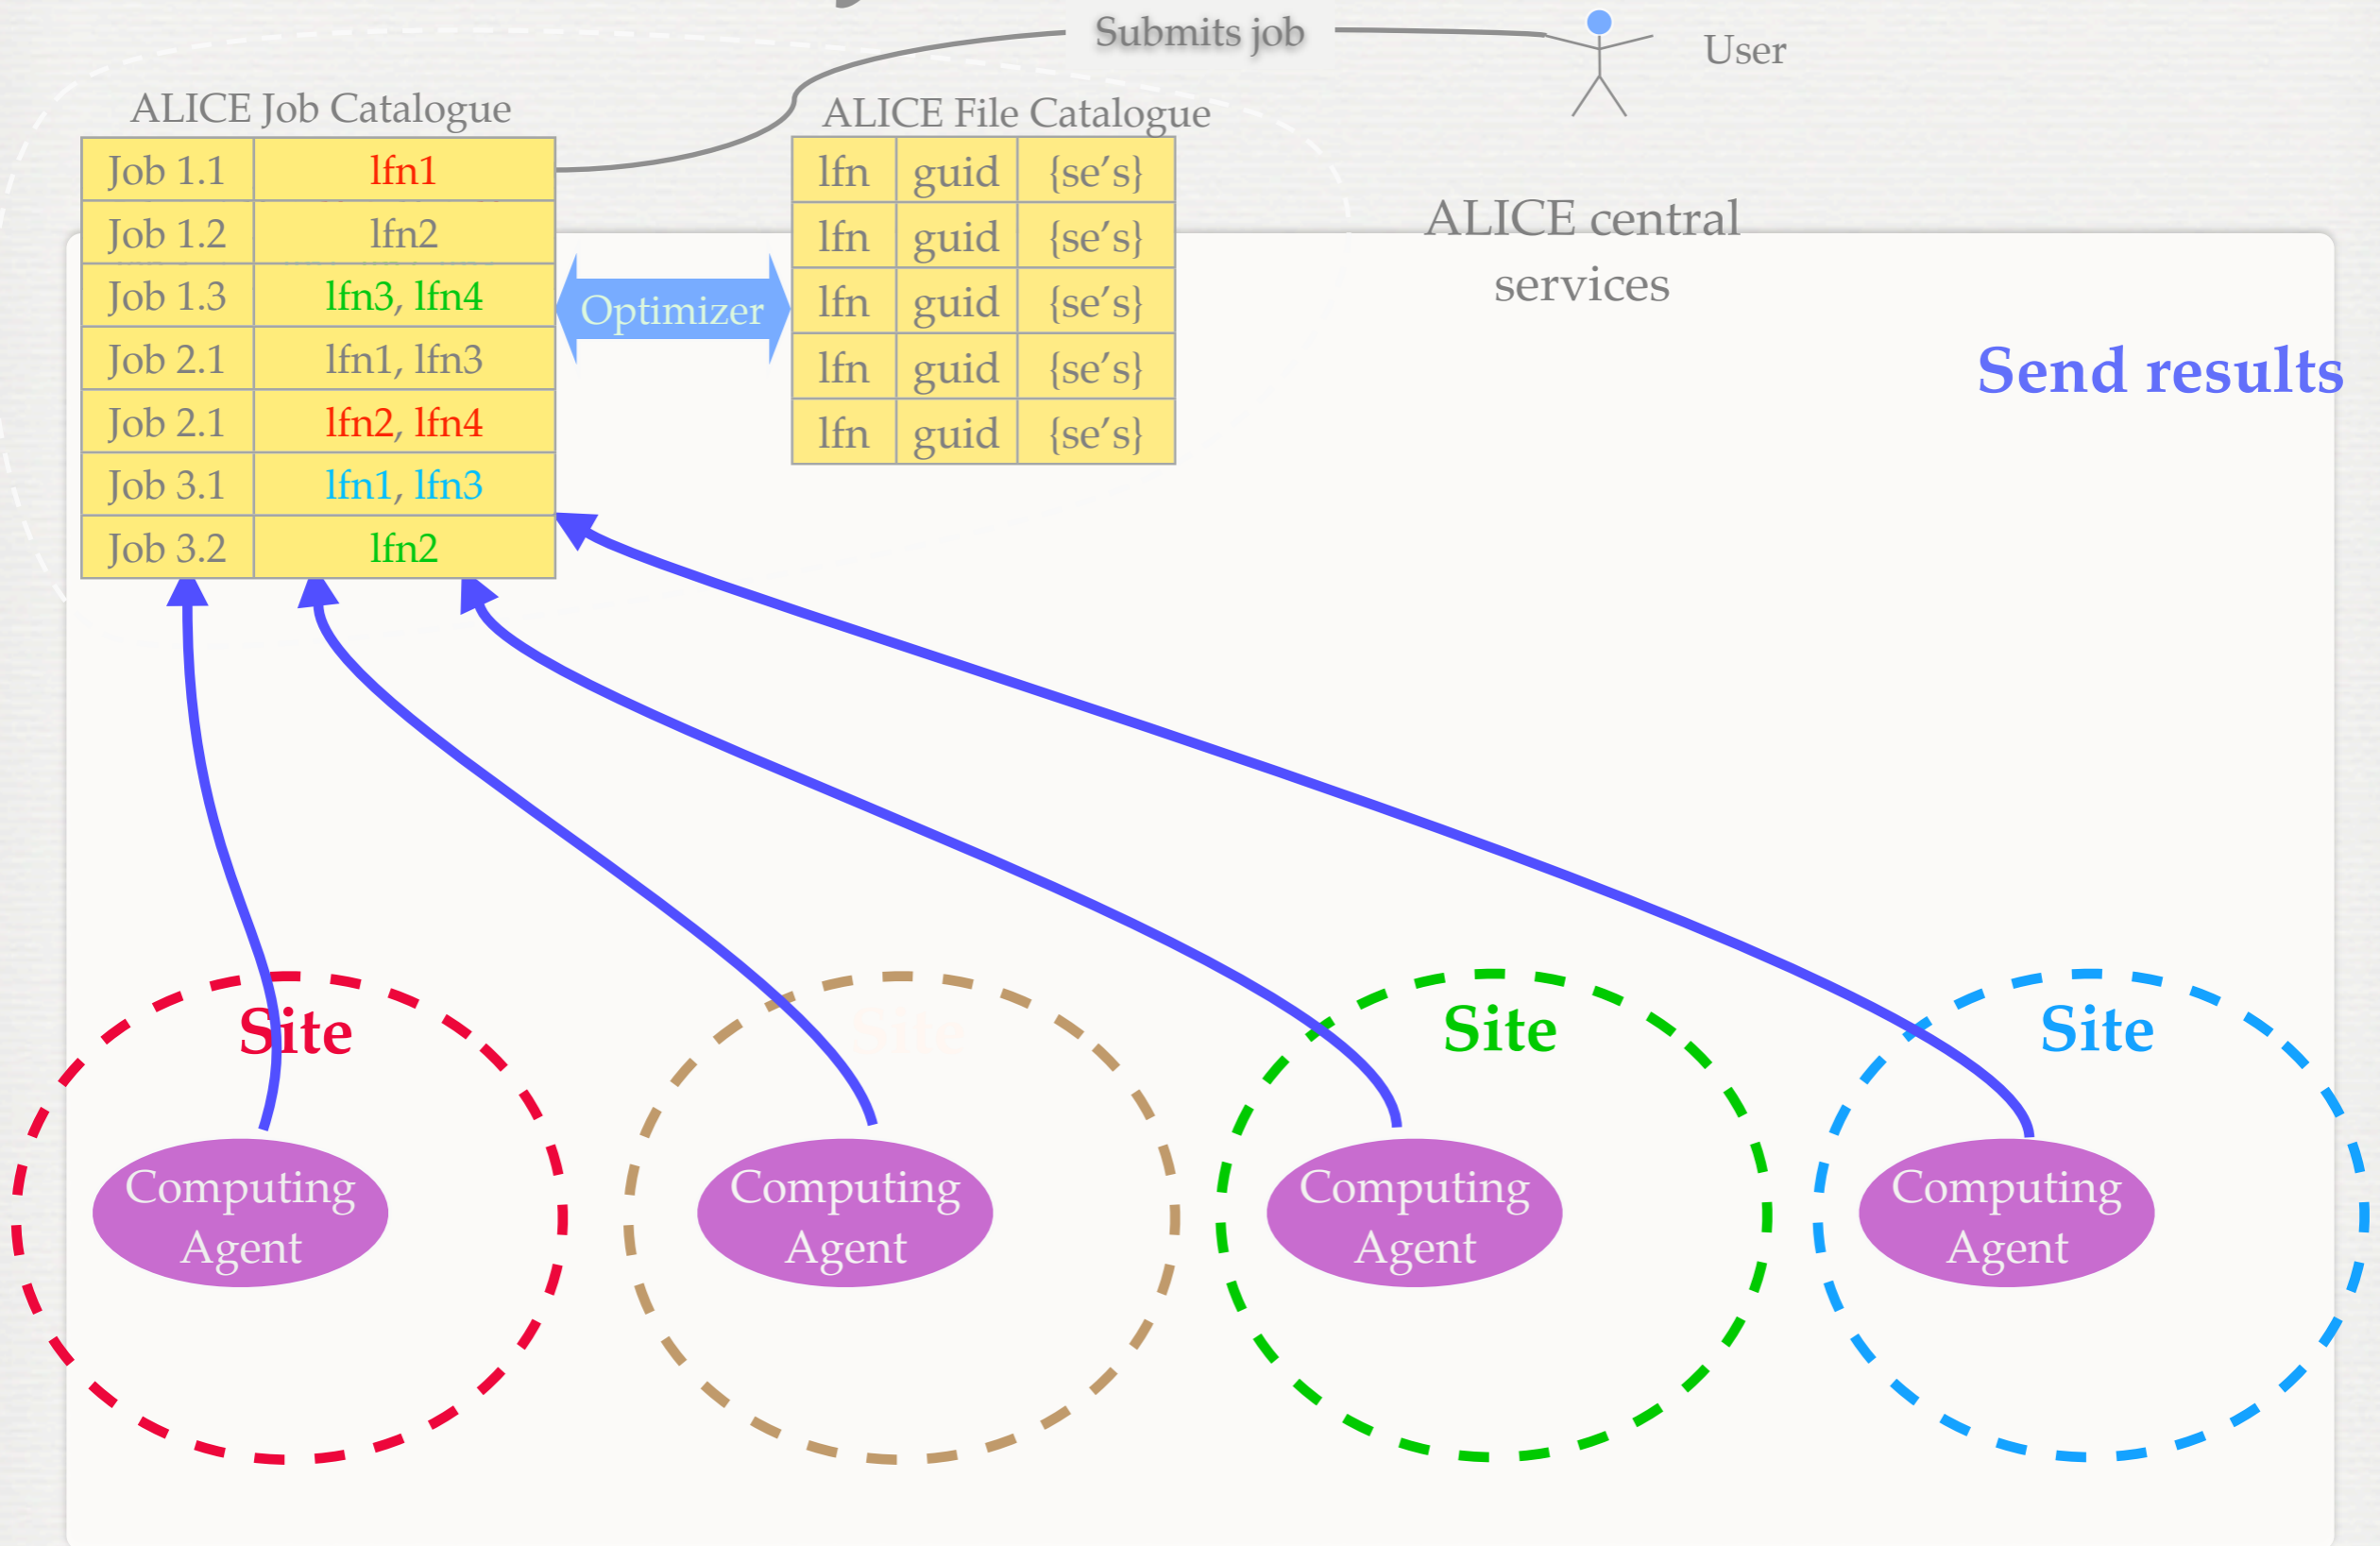
SENDING JOBS TO DATA



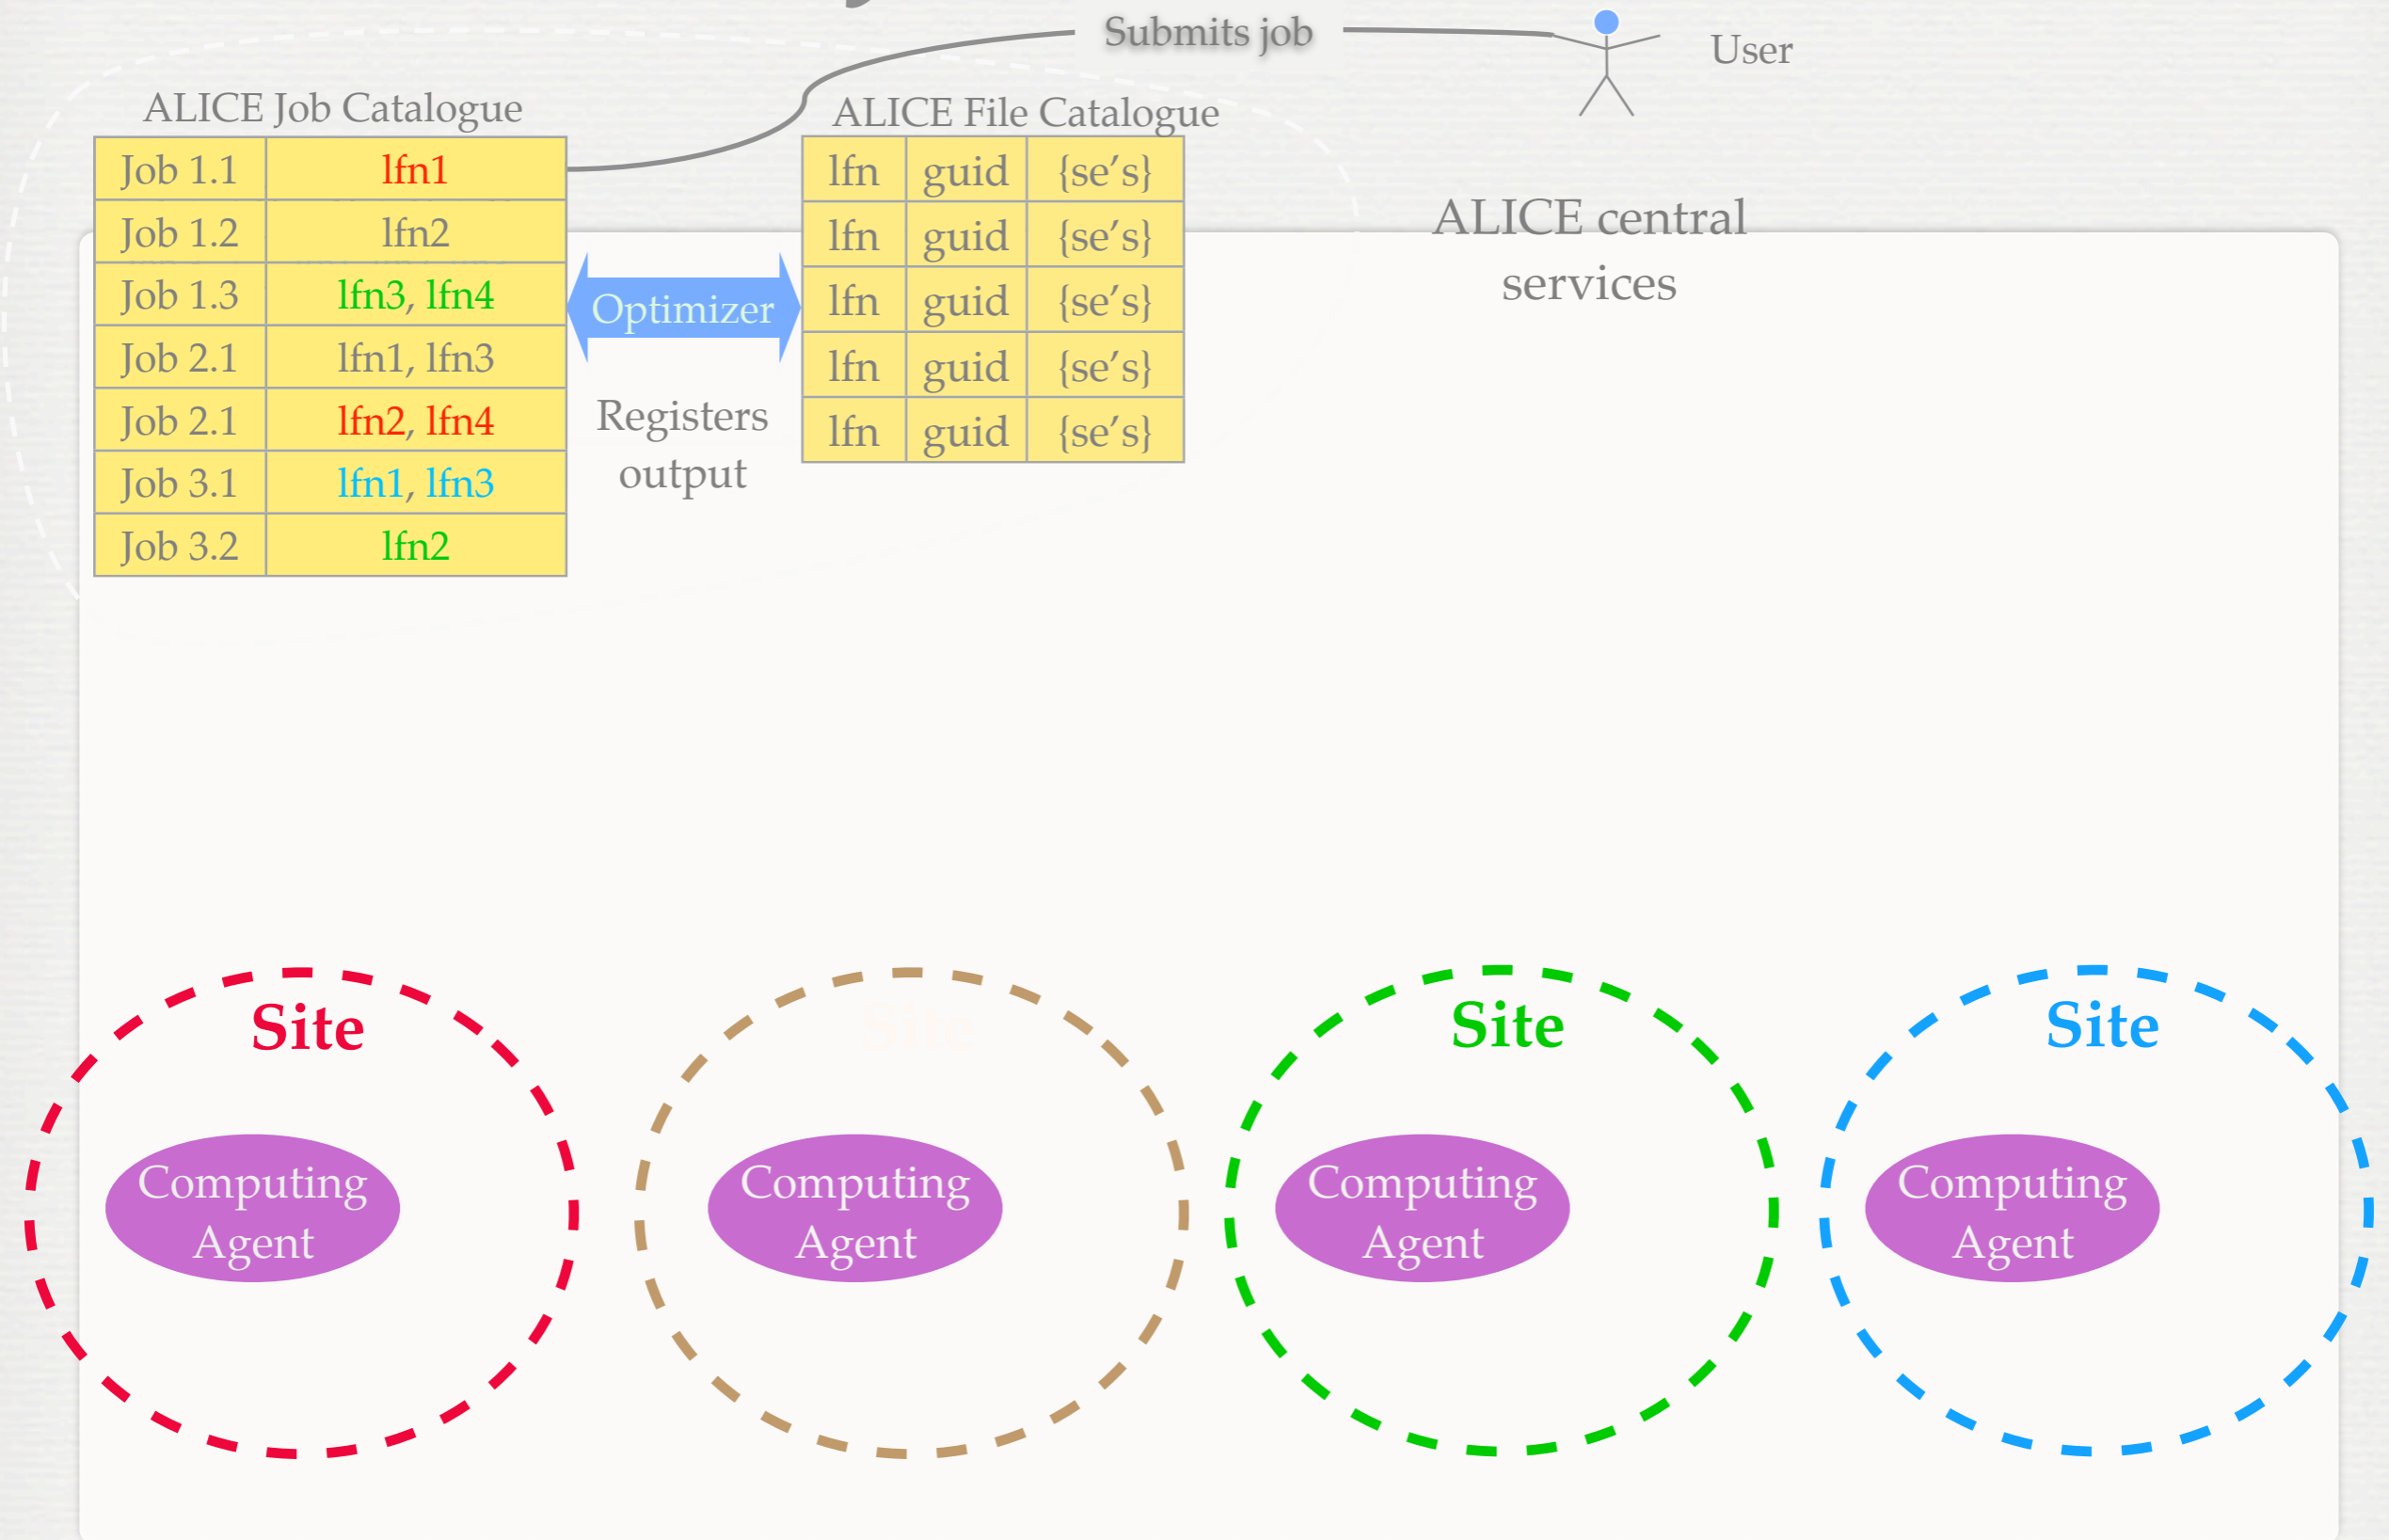
SENDING JOBS TO DATA



SENDING JOBS TO DATA



SENDING JOBS TO DATA



Disk

VOBOX::SA

xrootd
(manager)

xrootd
(server)

xrootd
(server)

DPM SRM

Castor SRM

xrootd
(server)

MSS

WN

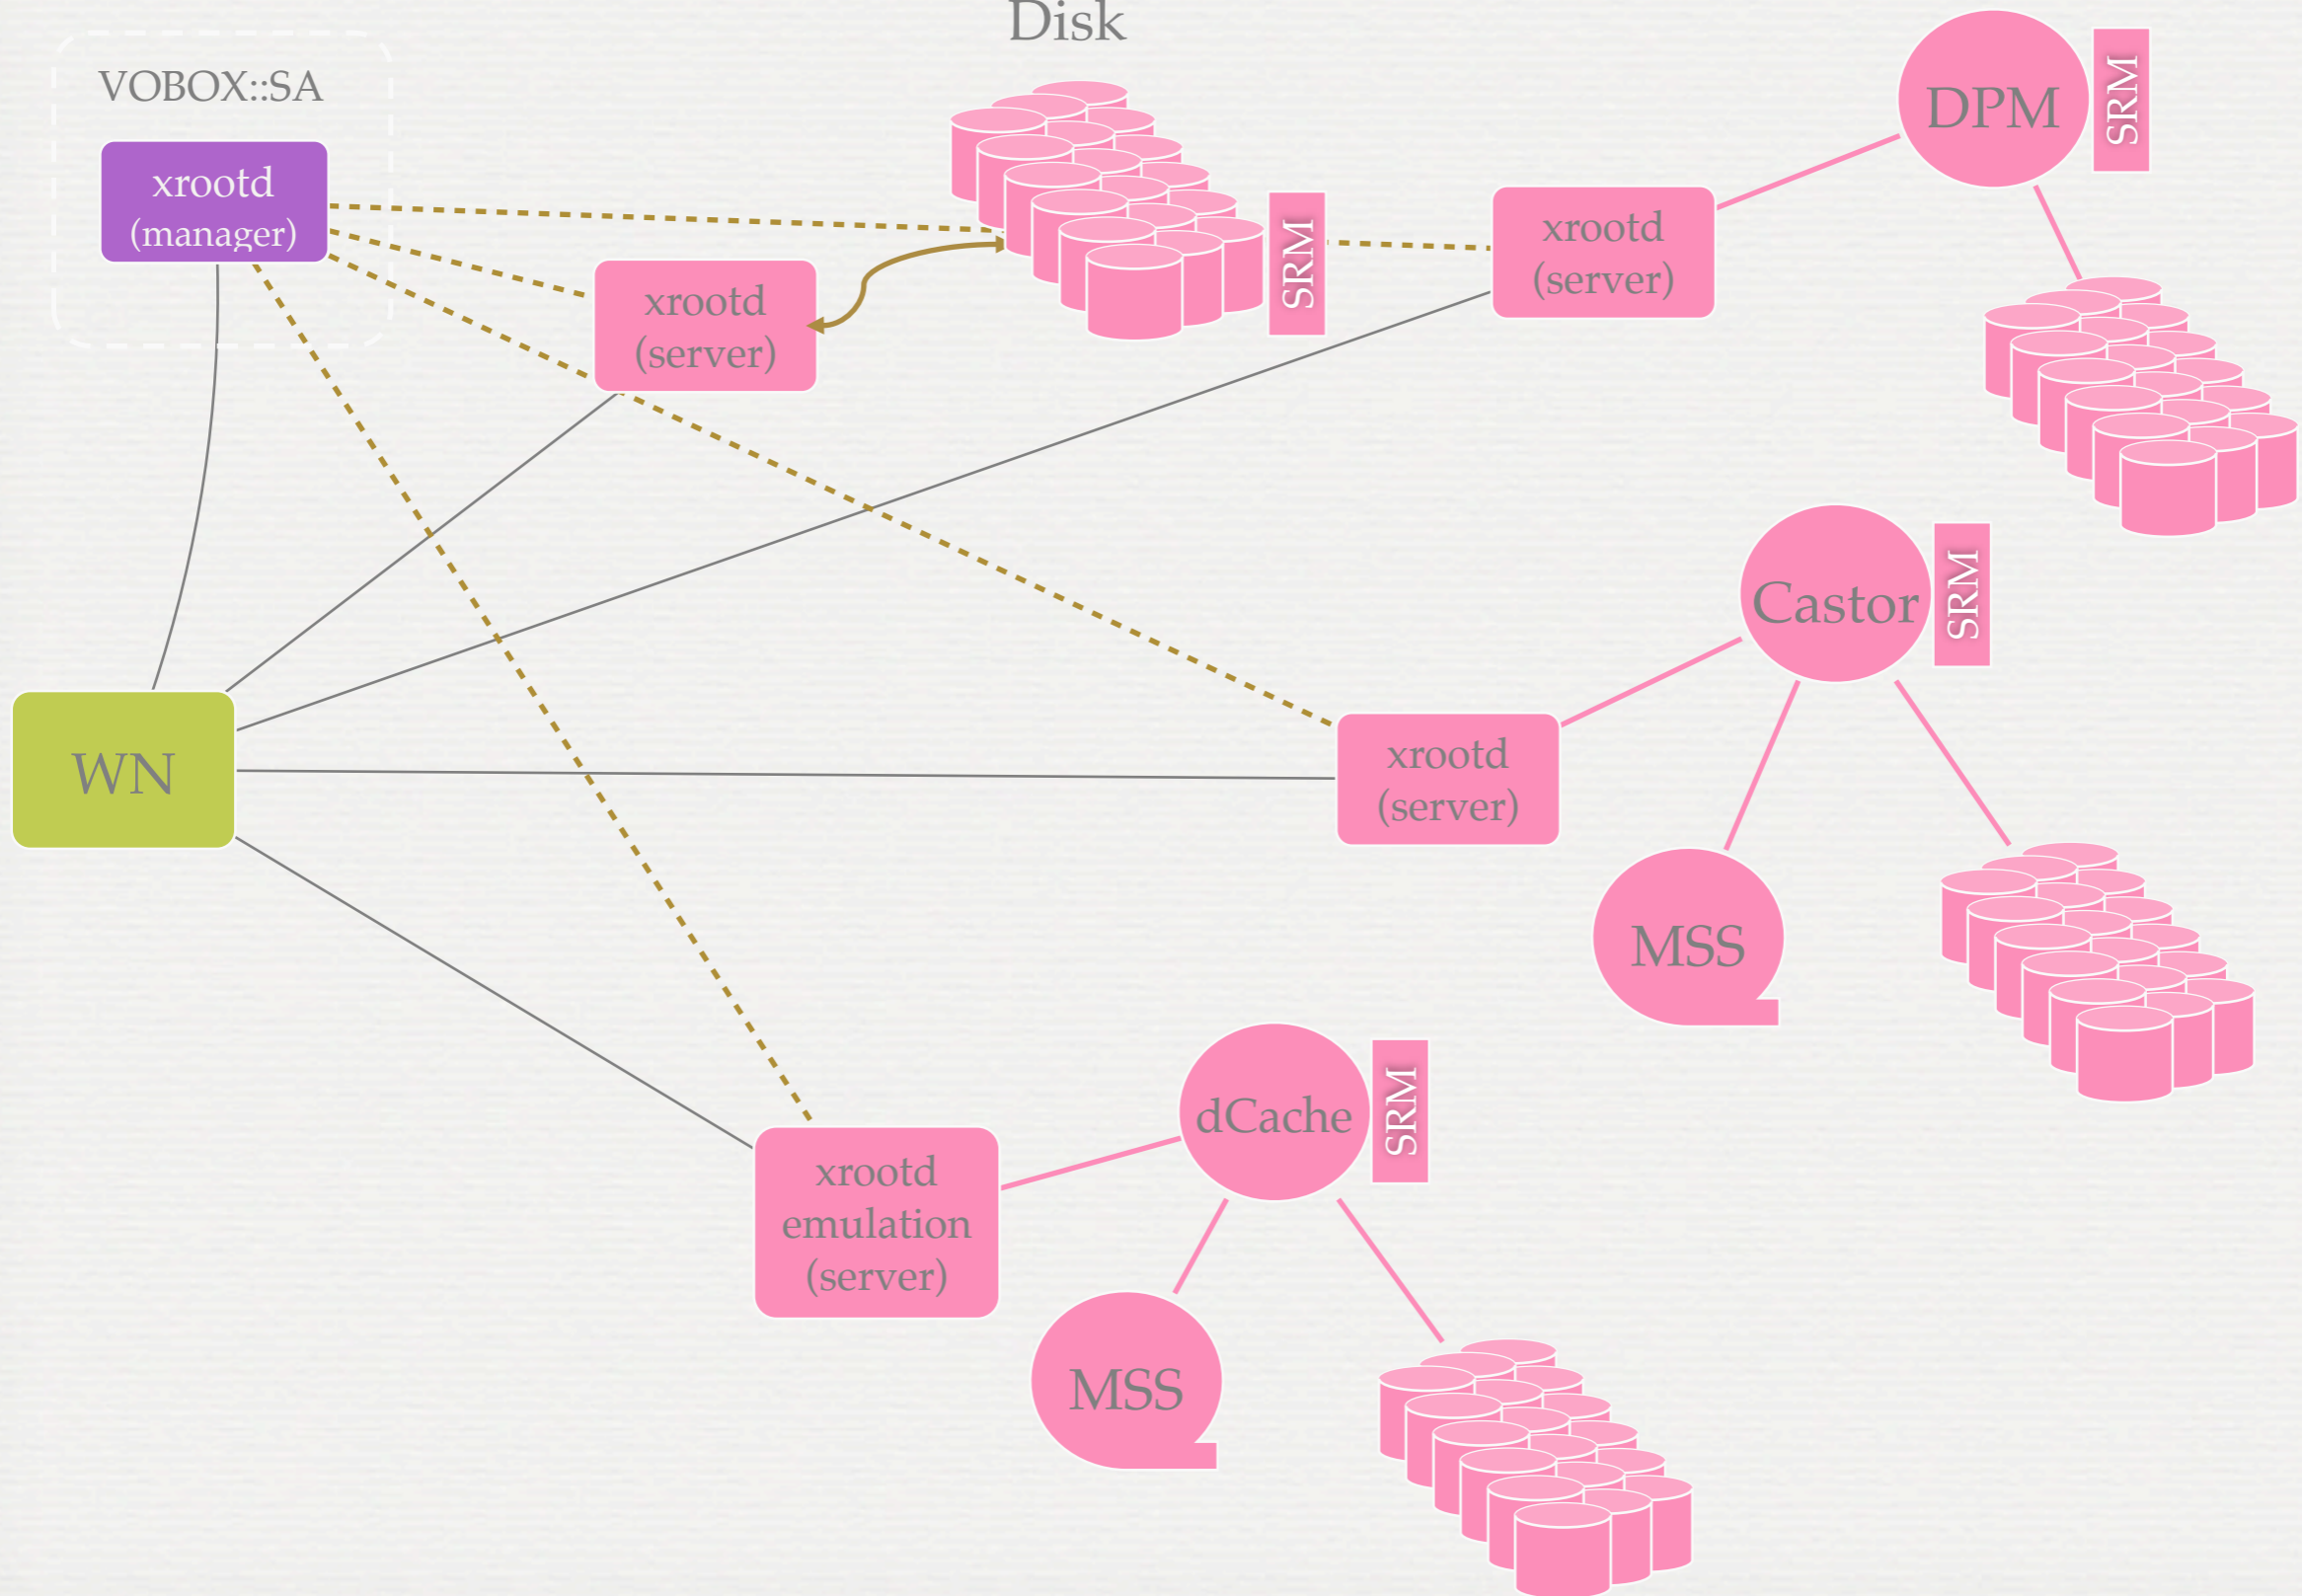
xrootd
emulation
(server)

dCache SRM

MSS

STORAGE STRATEGY

Disk

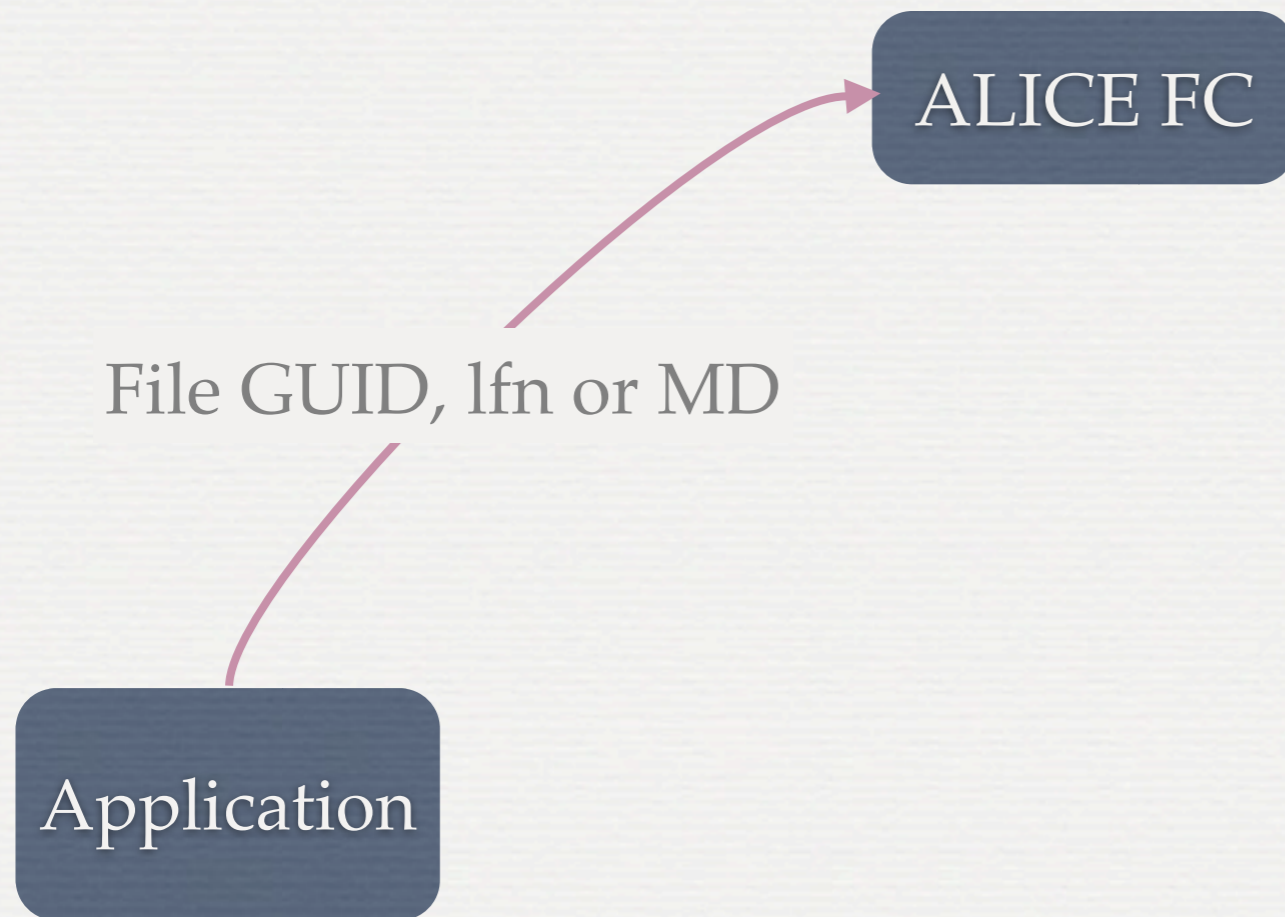


THE ACCESS TO THE DATA

Application

Direct access to data
via TAliEn/TGrid interface

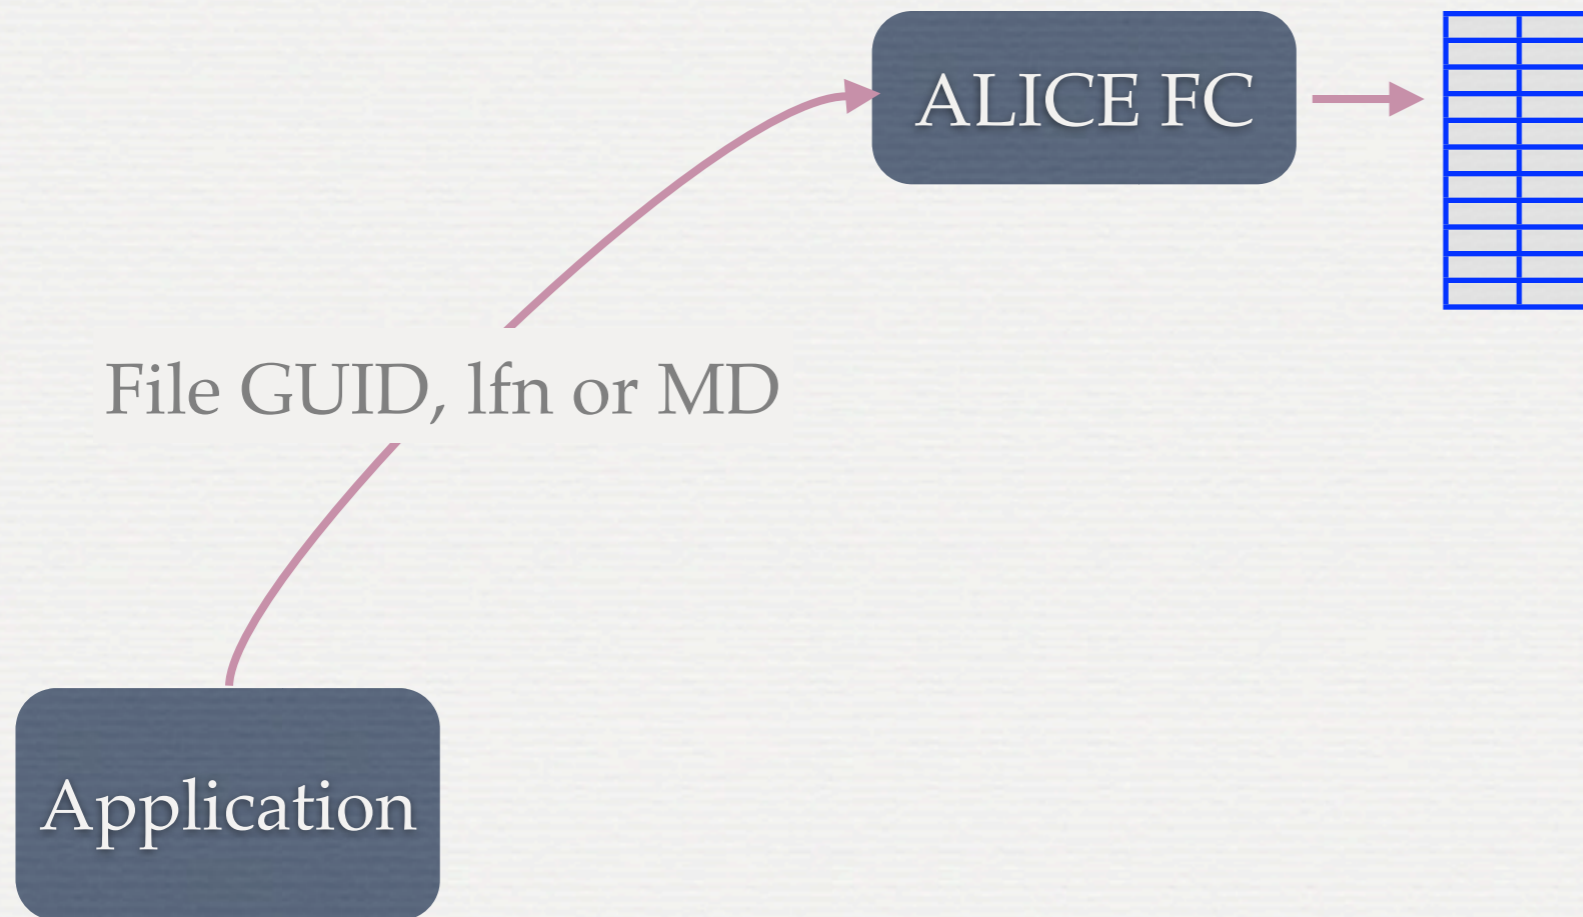
THE ACCESS TO THE DATA



Direct access to data
via TAliEn/TGrid interface

THE ACCESS TO THE DATA

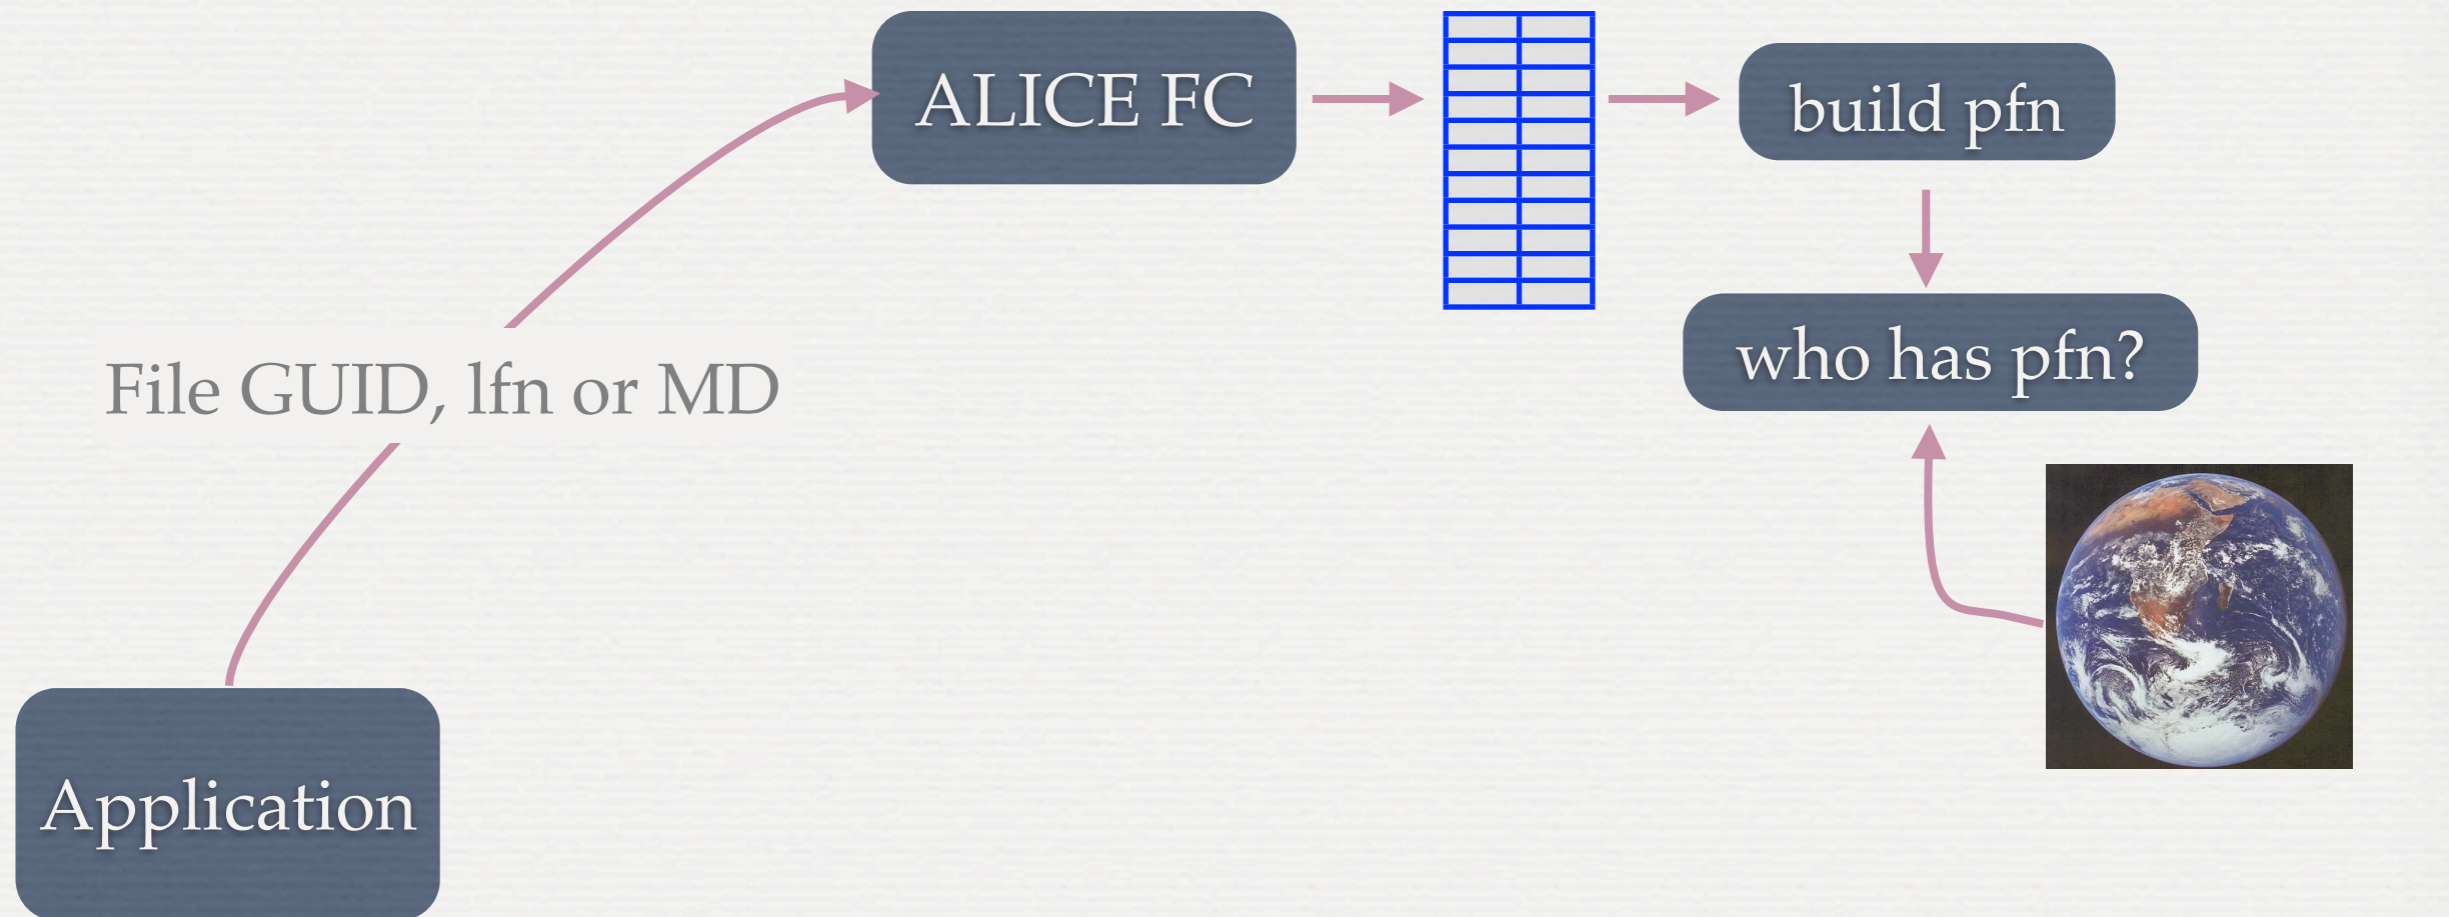
lfn → guid → (acl, size, md5)



Direct access to data
via TAliEn/TGrid interface

THE ACCESS TO THE DATA

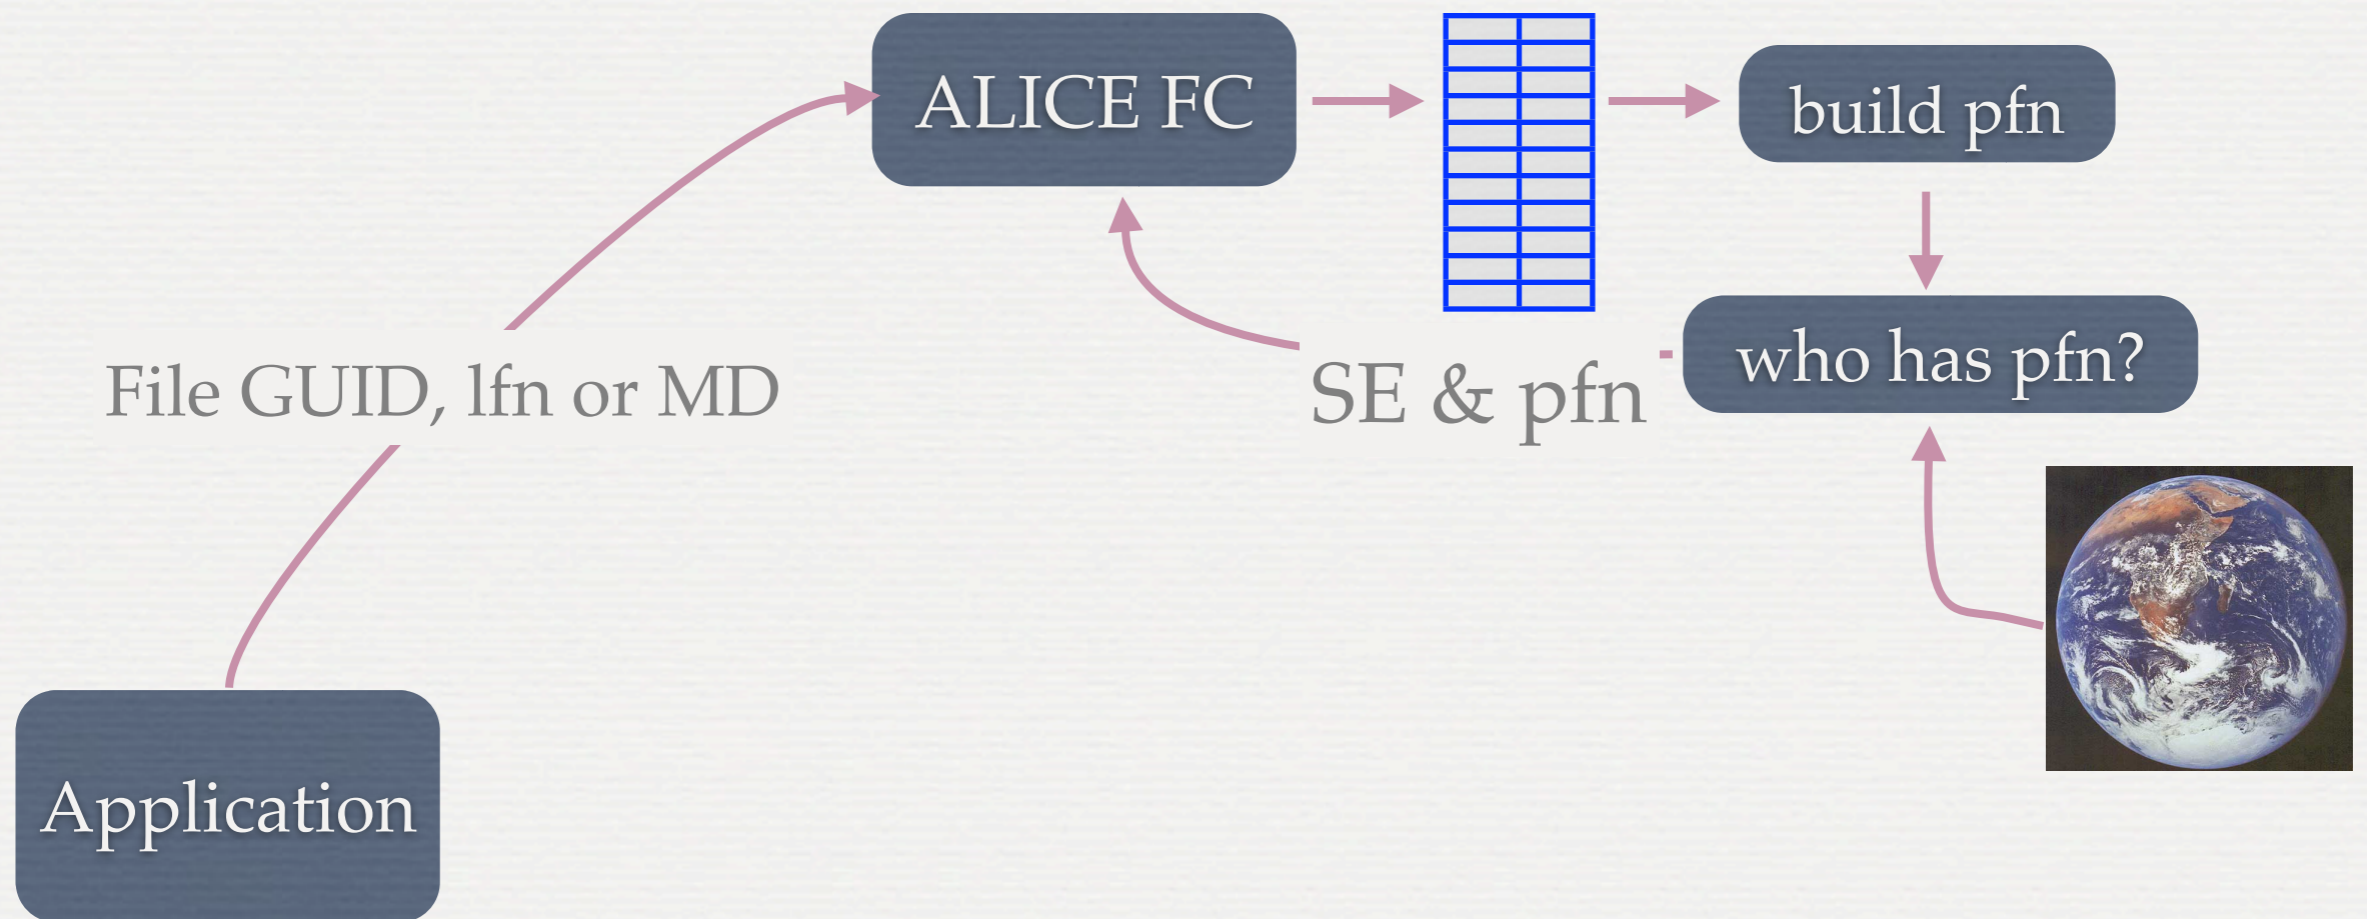
lfn \rightarrow guid \rightarrow (acl, size, md5)



Direct access to data
via TAliEn/TGrid interface

THE ACCESS TO THE DATA

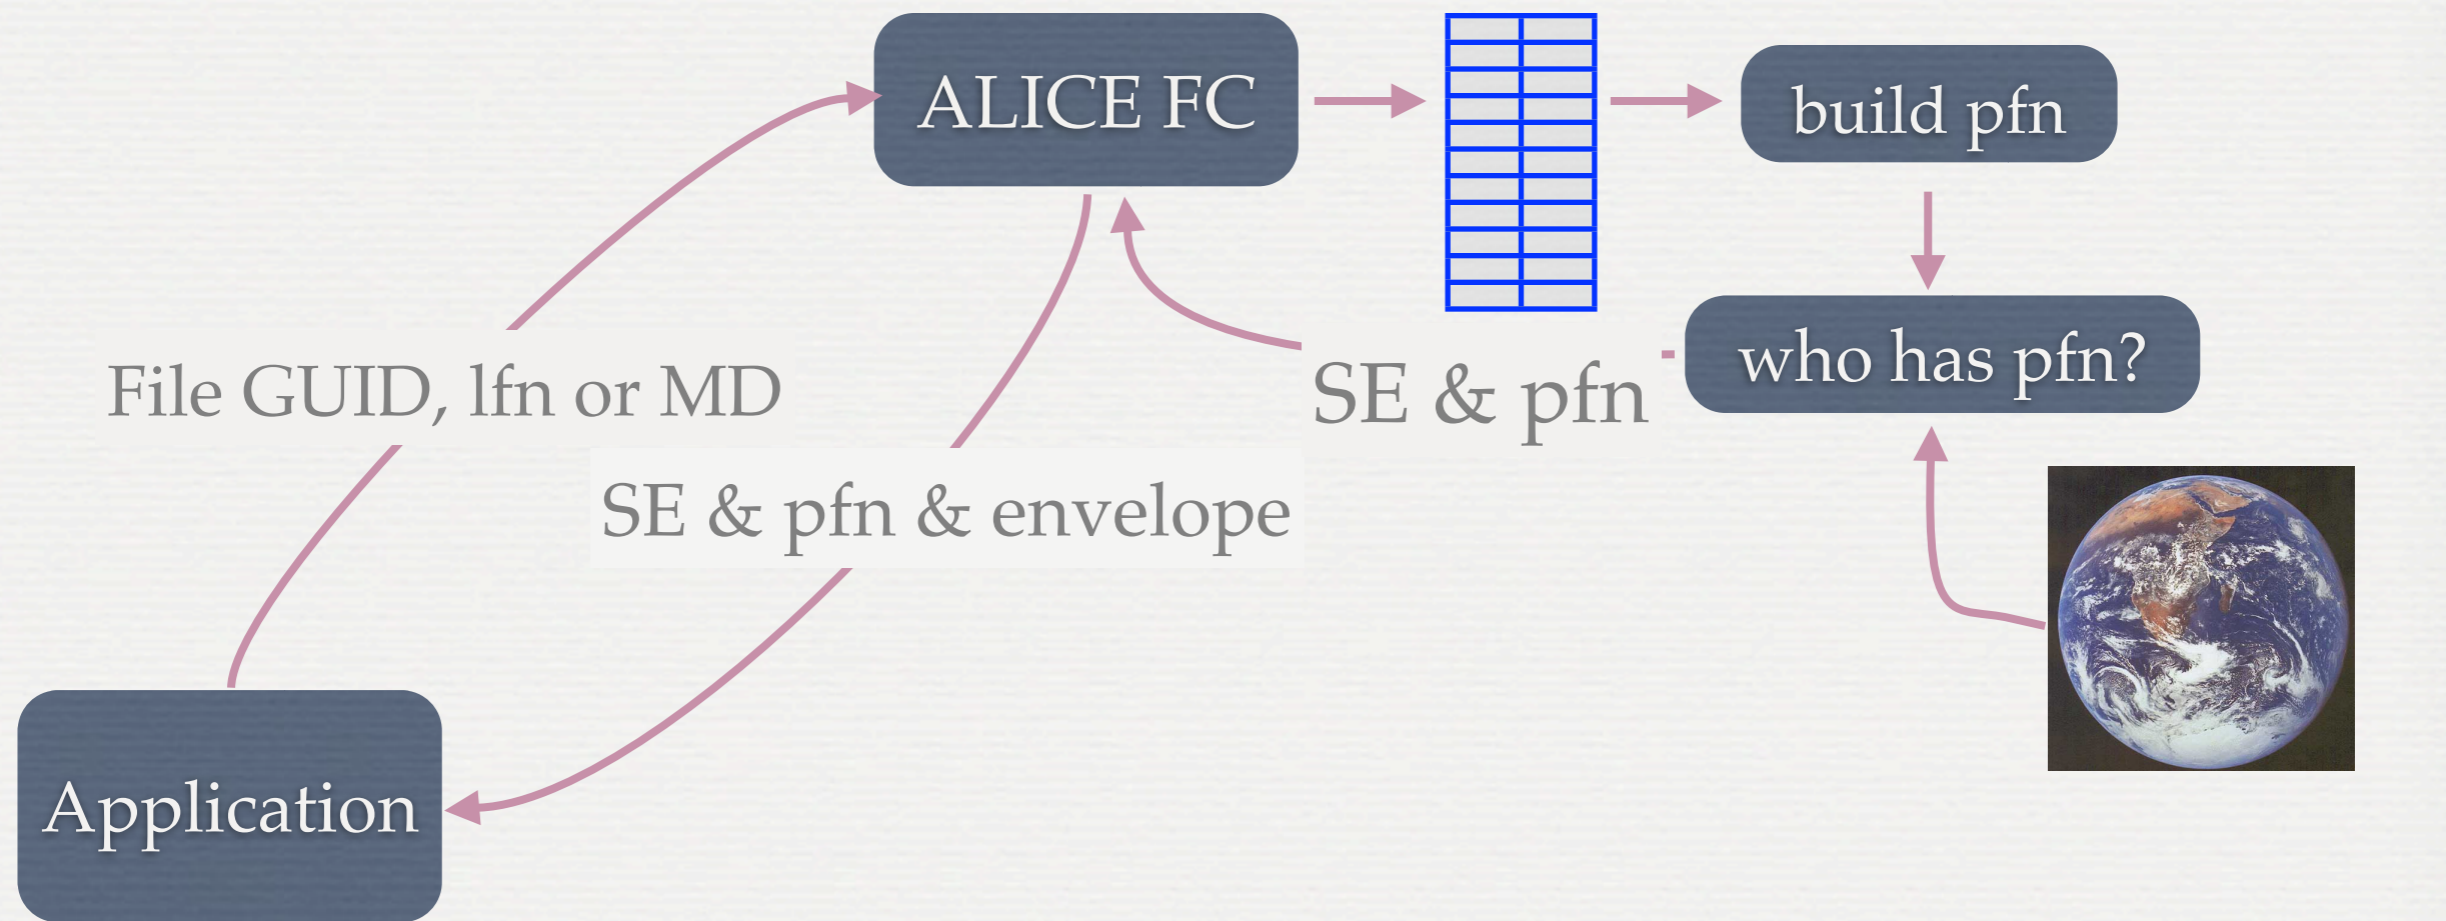
lfn \rightarrow guid \rightarrow (acl, size, md5)



Direct access to data
via TAliEn/TGrid interface

THE ACCESS TO THE DATA

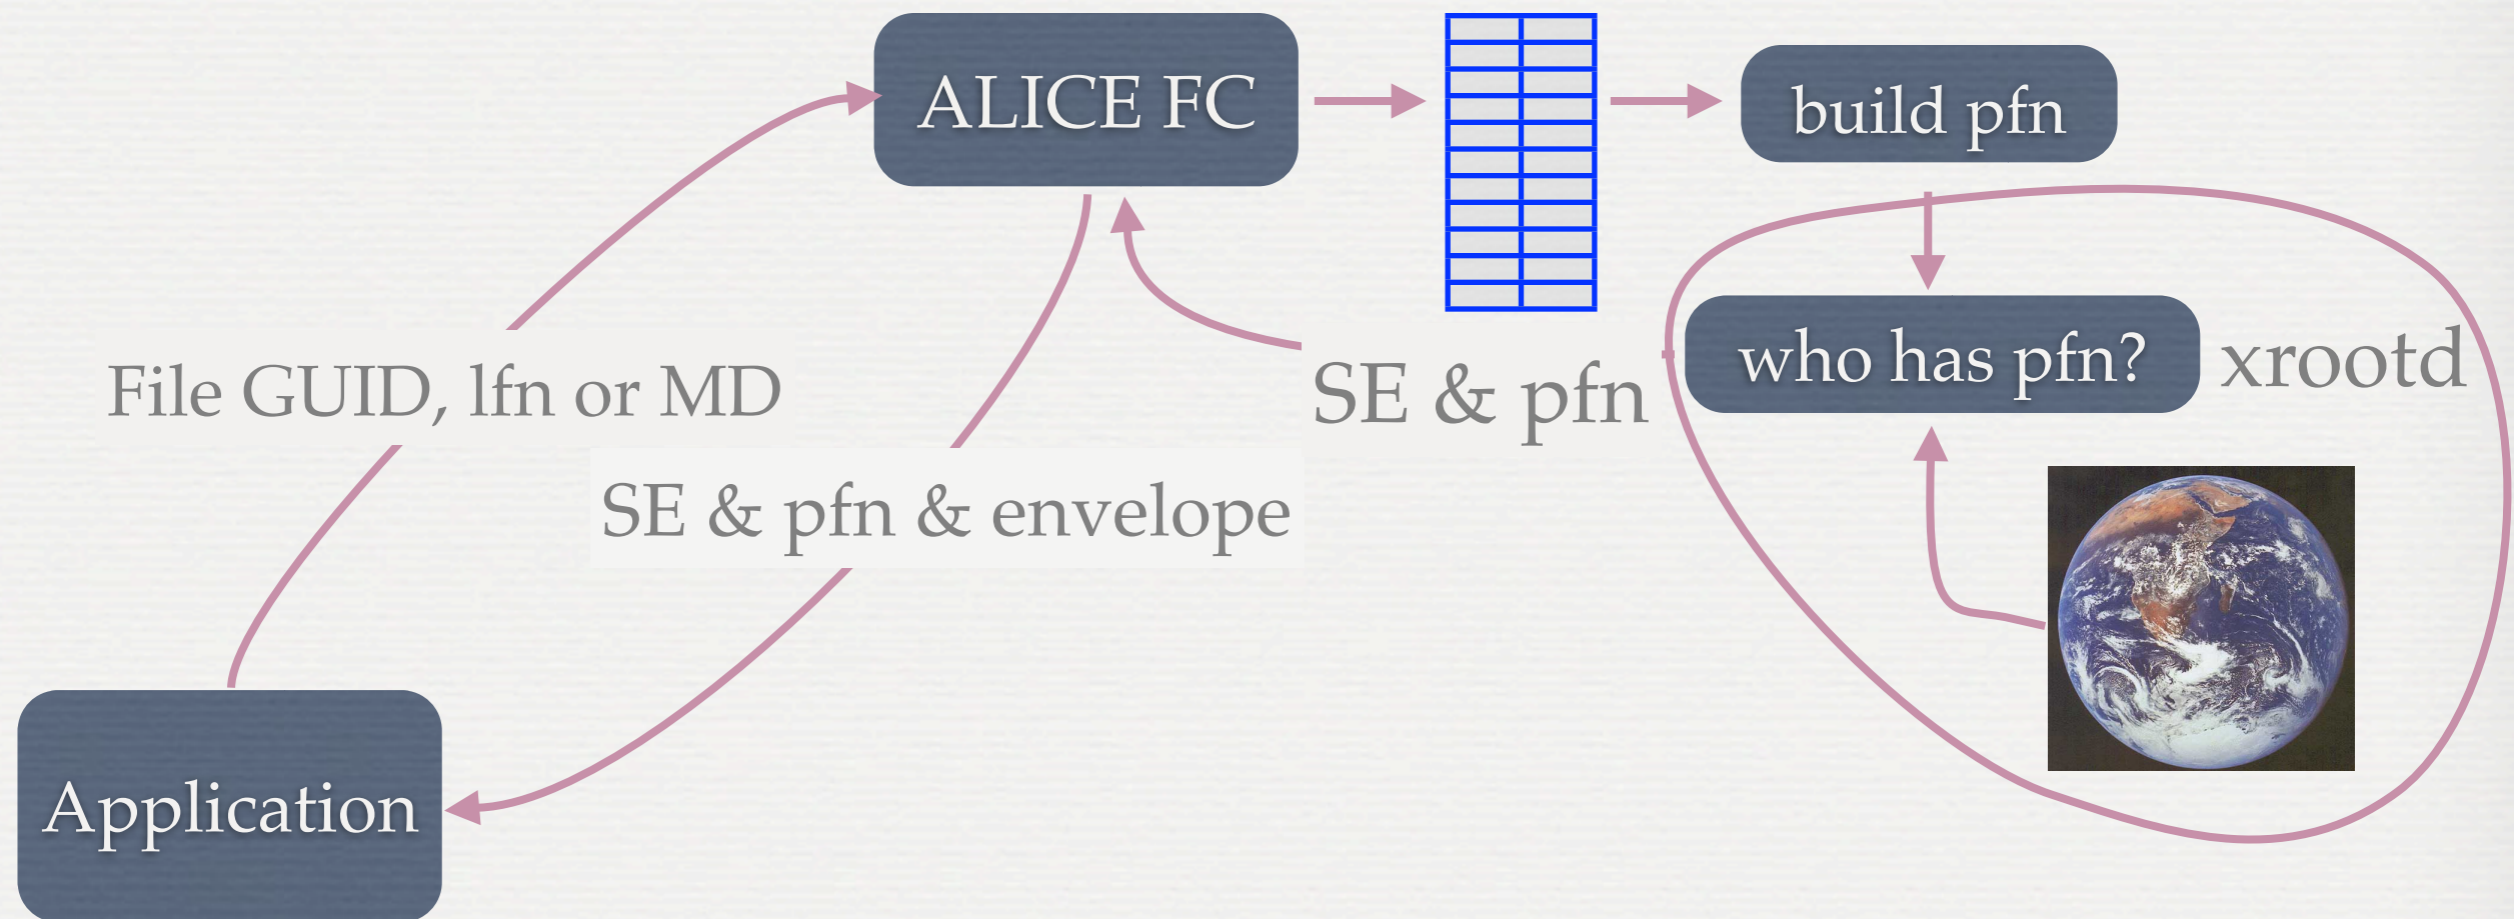
lfn → guid → (acl, size, md5)



Direct access to data
via TAliEn/TGrid interface

THE ACCESS TO THE DATA

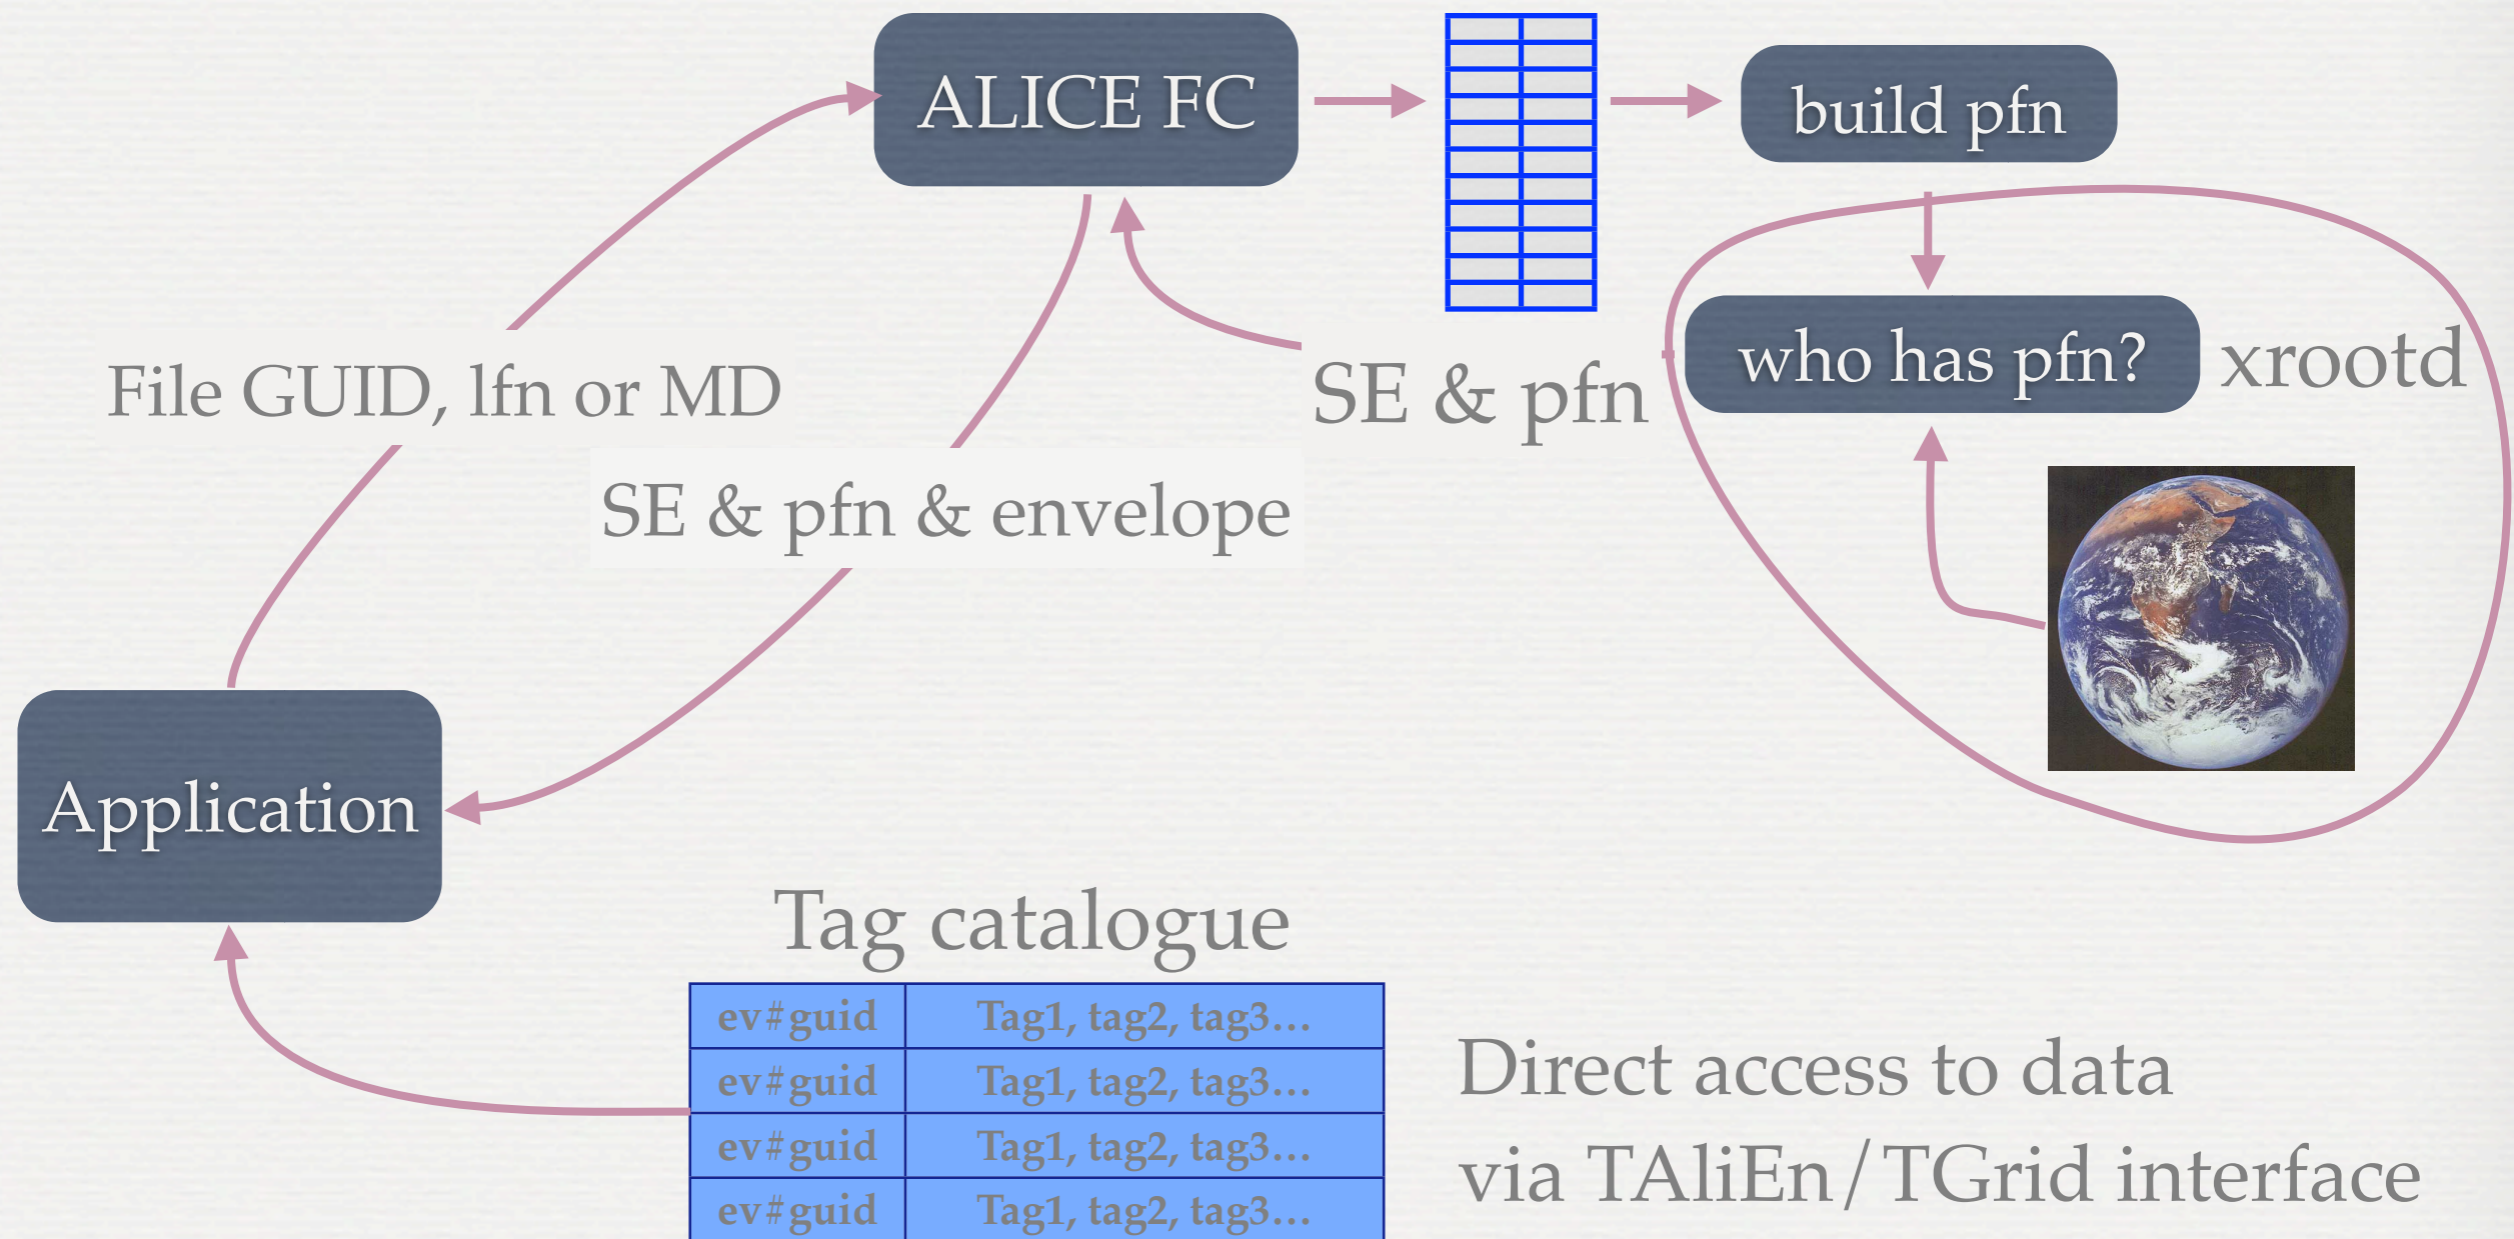
lfn \rightarrow guid \rightarrow (acl, size, md5)



Direct access to data
via TAliEn/TGrid interface

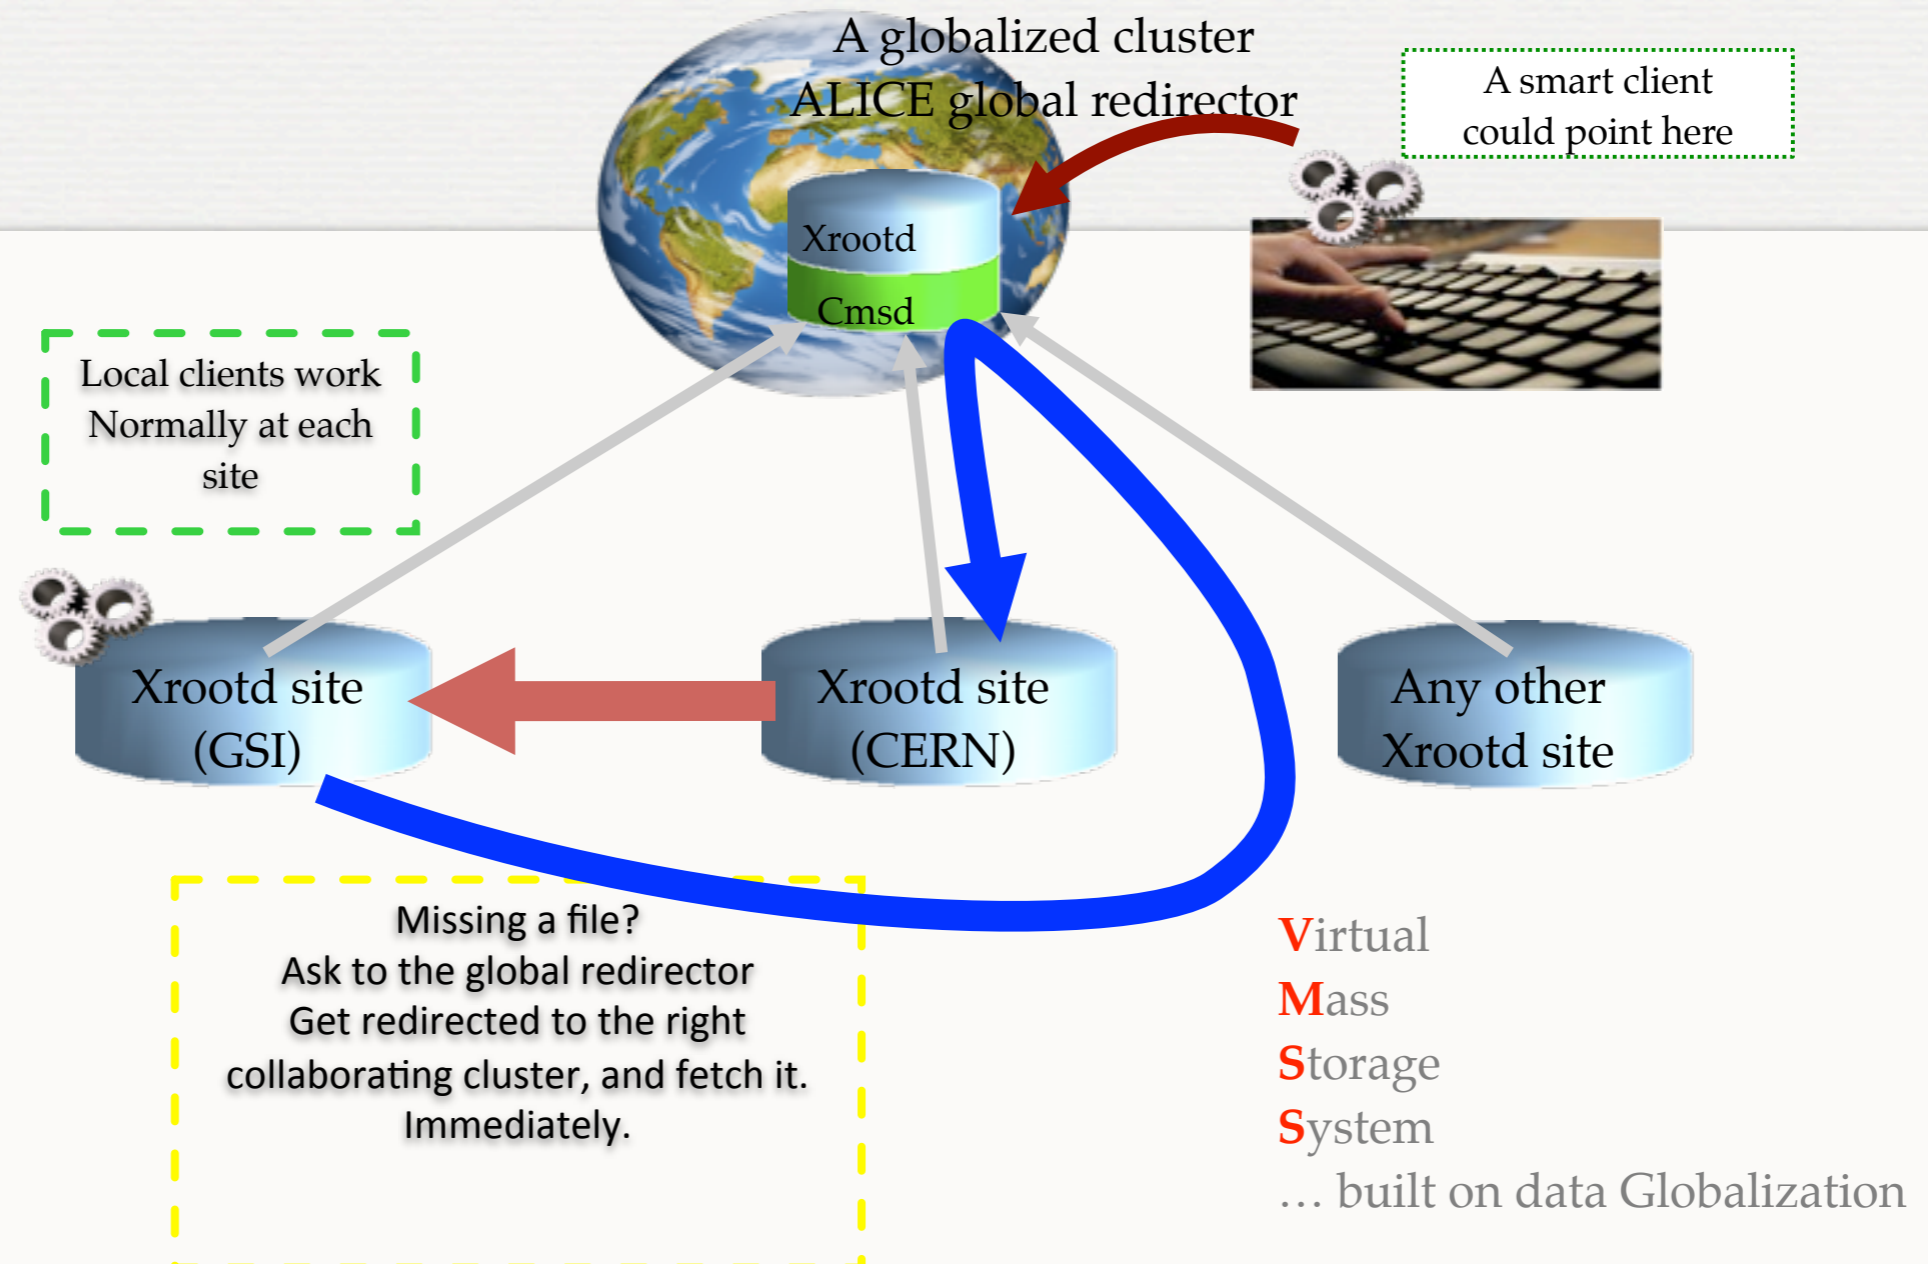
THE ACCESS TO THE DATA

lfn → guid → (acl, size, md5)



Direct access to data
via TAliEn / TGrid interface

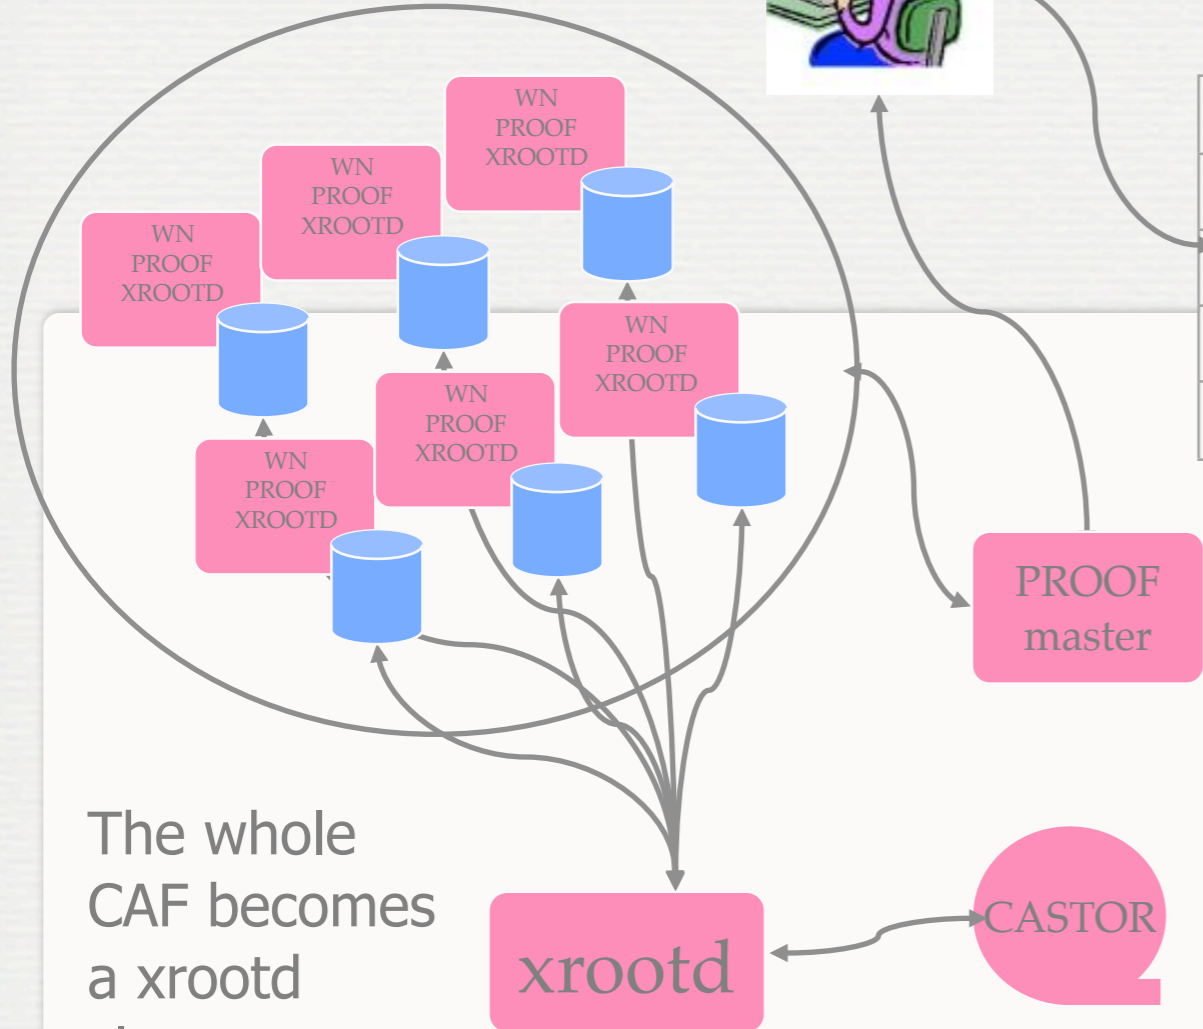
THE ALICE WAY WITH XROOTD



- Pure Xrootd + ALICE strong authz plugin. No difference among T1/T2 (only size and QOS)
- WAN-wide globalized deployment, very efficient direct data access
- Tier-0: CASTOR+Xrd serving data normally.
- Tier-0: Pure Xrootd cluster serving conditions to ALL the GRID jobs via WAN

More details and complete info in ["Scalla/Xrootd WAN globalization tools: where we are." @ CHEP09](#)

CAF



The whole CAF becomes a xrootd cluster

- Powerful and fast machinery – very popular with users
- Allows for any use pattern, however quite often leading to contention for resources

lfn	guid	{se's}
lfn	guid	{se's}
lfn	guid	{se's}
lfn	guid	{se's}
lfn	guid	{se's}


MIT PHOBOS

PROOF in PHOBOS

Maarten Ballintijn / MIT
maartenb@mit.edu


May 24, 2006 – Application Area Meeting

Rare high multiplicity event set




- Burak Altunokur
- Dataset: 11k files, 4.5 TB
- 150 slaves, ~1 hour

Two particle correlation function of minias dAu 200GeV



- Wei Li, Constantin Loizides
- Dataset: 4.5k files, 1.5 TB
- 100 slaves, 75 min


Query Short in different environments



Observed speedup

Expected speedup

History of cluster usage



THE ALICE GRID



search...

About | Development | MonALISA ALICE Grid Monitor | ALICE Collaboration | Contact & Team

- AliEn working prototype in 2002

AliEn

- Single interface to distributed computing for all ALICE physicists

» Home
» User distribution
» Download
» Documentation
» Virtual Organizations
» Events

- File catalogue, job submission and control, software management, user analysis

~80 participating sites now

- 1 T0 (CERN/Switzerland)

- 6 T1s (France, Germany, Italy, The Netherlands, Nordic DataGrid Facility, UK)

- KISTI and UNAM coming (!)

- ~73 T2s spread over 4 continents

~30,000 (out of ~150,000 WLCG) cores and 8.5 PB of disk

Username

Password

Remember Me

Login

- Resources are “pooled” together

• Forgot your password?
• Forgot your username?

- No localisation of roles / functions

- National resources must integrate seamlessly into the global grid to be accounted for

- FAs contribute proportionally to the number of PhDs (M&O-A share)

- T3s have the same role than T2s, even if they do not sign the MoU

<http://alien.cern.ch>

ALL IS IN MONALISA



MonALISA Repository for ALICE



[My jobs](#) [My home dir](#) [Catalogue browser](#) [Repository Home](#) [Administration Section](#) [ALICE Reports](#) [Events XML Feed](#) [Firefox Toolbar](#) [MonaLisa GUI](#)

- ALICE Repository
 - ALICE Repository
 - Google Map
 - Shifter's dashboard
 - Run Condition Table
 - Production Info
 - Run view
 - RAW production cycles
 - RAW activities
 - Analysis train
 - MC production cycles
 - MC production requests
 - Job Information
 - Site views
 - Summary plots
 - Job states
 - Jobs per site
 - Jobs per site table
 - Resource usage
 - User views
 - Summary plots
 - Jobs status
 - Grid packages
 - Quotas
 - Task queue
 - Task queue summary
 - Jobs in TQ table
 - Job timings
 - By site
 - Per user
 - Memory profiles
 - By site
 - Per user
 - Current jobs
 - SE Information
 - Status
 - Traffic
 - Files
 - xrootd
 - CERN Castor2x
 - AF's
 - Services
 - Network Traffic
 - FTD Transfers
 - CAF Monitoring
 - SHUTTLE
 - Build system
 - HepSpec
 - Dynamic charts



● Running jobs ● Running jobs but no ML info ● Site service problem(s) prevents job execution ● No jobs match the site resources ● ML service down & no running jobs


Map options Show xrootd transfers

Jump to: [Europe](#) [North America](#) [South America](#) [Asia](#) [World](#) [Save position and options](#)

Imagery ©2011 TerraMetrics, NASA - Terms of Use

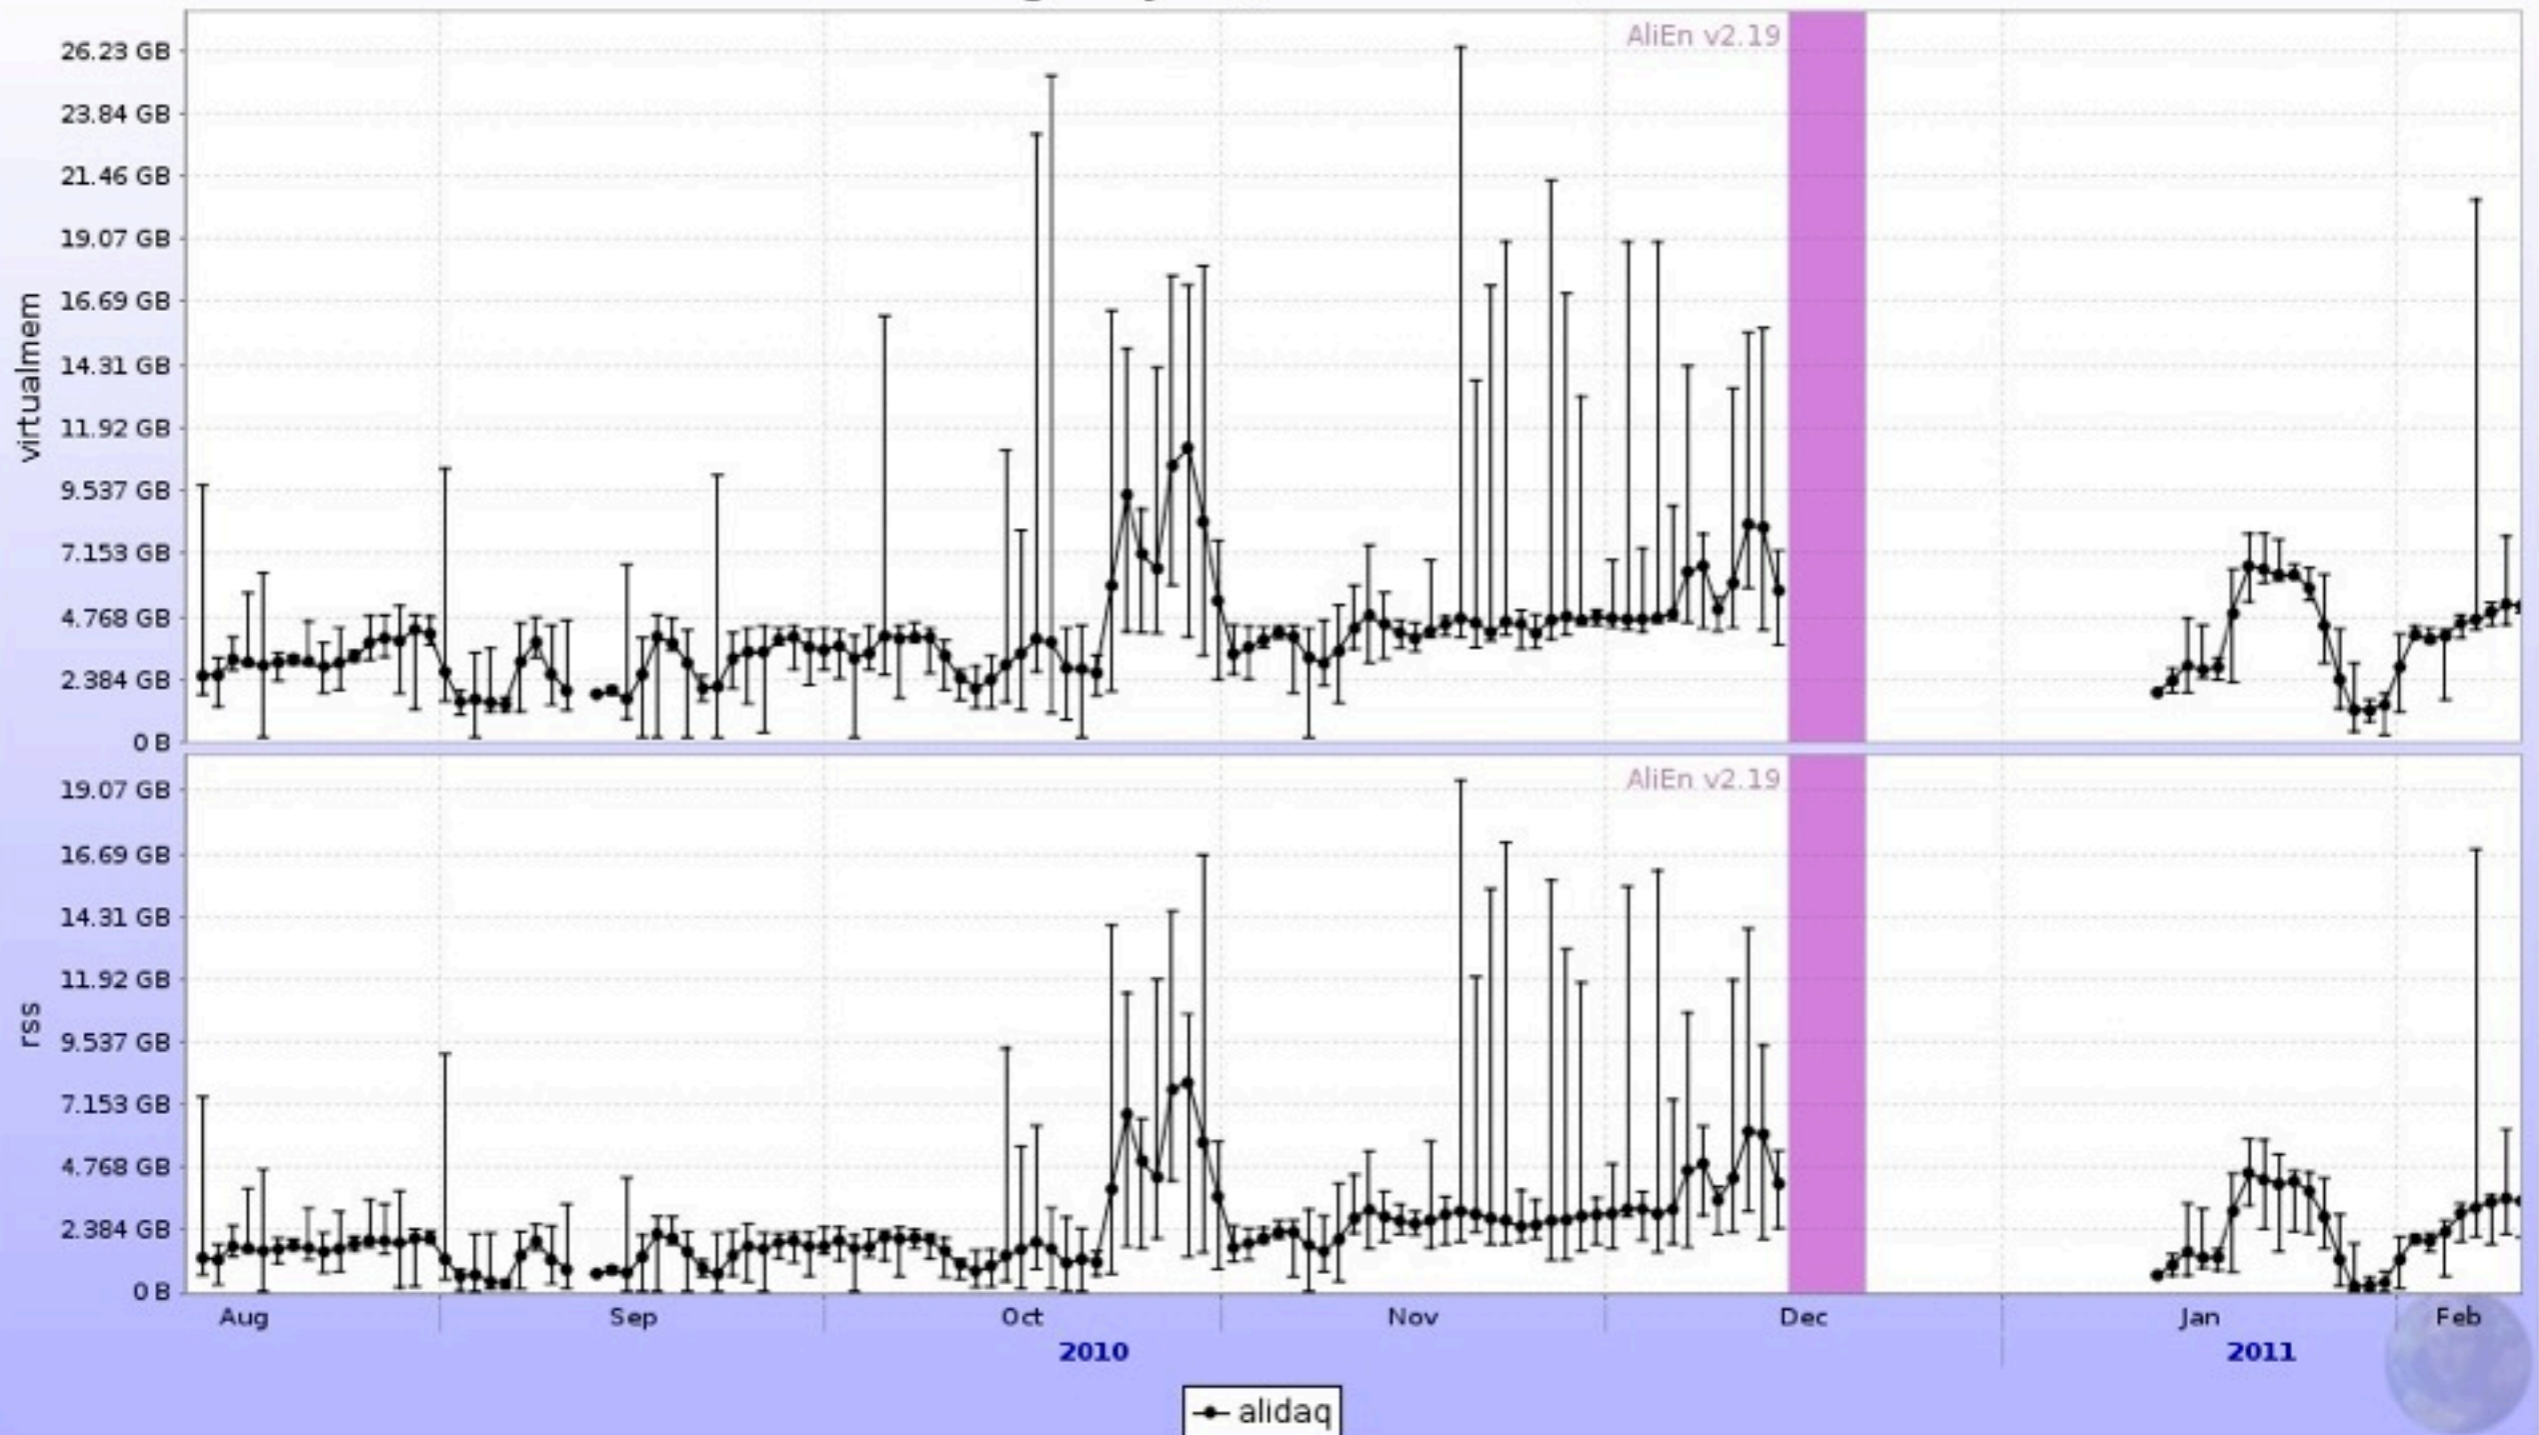
[Find your location](#)

ALL IS IN MONALISA




ALL IS IN MONALISA

Largest job (site=CERN-L)



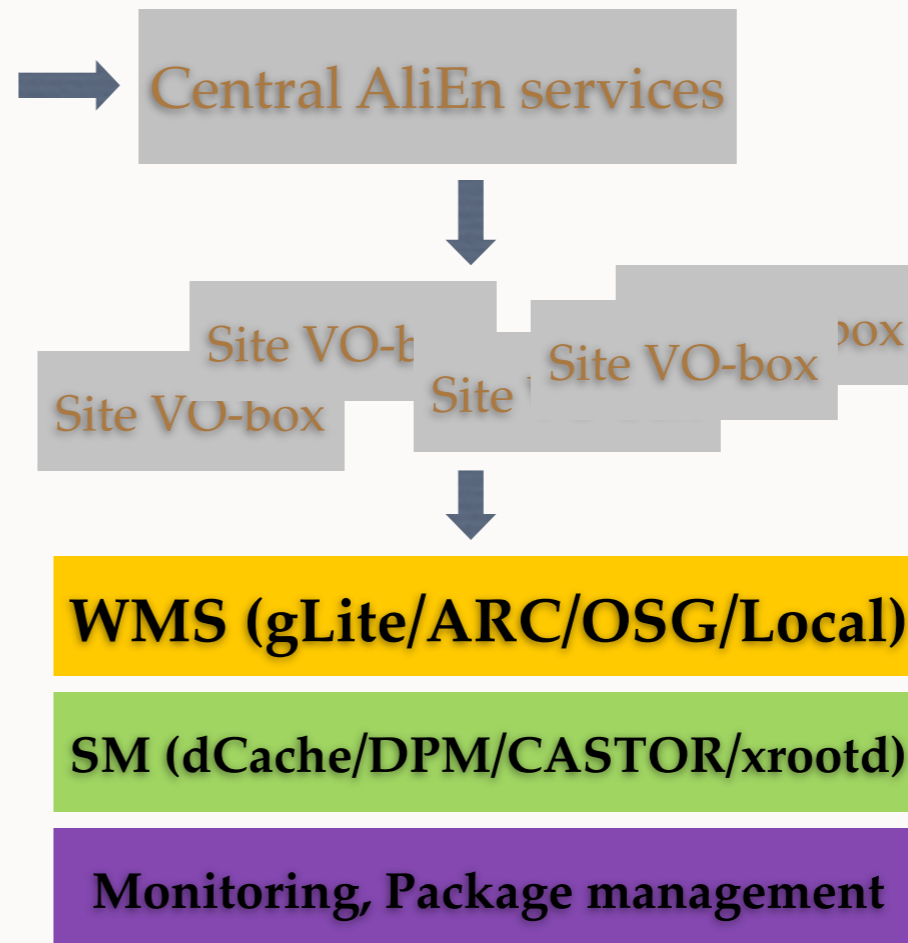
ALL IS IN MONALISA

Running Jobs



- SUM ● Athens ● Bari ● Birmingham ● Bologna ● Bratislava ● Cagliari ● Catania ● CCIN2P3 ● CCIN2P3-CREAM ● CERN ● CERN-CREAM ● CERN-L
- Clermont ● CNAF-CREAM ● COMSATS ● CSC ● CyberSar ● CyberSar-CREAM ● Cyfronet ● DCSC_KU ● Dortmund ● FZK_CREAM ● FZK_glexec
- Grenoble ● GRIF_IPNO ● GRIF_IRFU ● GSI-CREAM ● GSI-SCLAB ● HHLR_GU ● Hiroshima ● IHEP ● IPNL ● ISS ● ISS_LCG ● ITEP ● ITEP-CREAM ● JINR
- KFKI ● KISTI-CREAM ● KISTI_GSDC ● KNU ● Kolkata-CREAM ● Kosice ● KPI ● LBL ● Legnaro ● Legnaro-CREAM ● LLNL ● Loewe_CSC ● LUNARC
- Madrid ● MEPHI ● NIHAM ● NIKHEF ● NSC ● OSC ● PAKGRID ● PDC ● PNPI ● Poznan ● Prague-CREAM ● RAL ● RRC-KI ● SaoPaulo ● SARA
- SPbSU-CREAM ● Strasbourg_IRES ● Subatech ● SUT ● Torino ● Torino-CREAM ● Trieste ● TriGrid ● Troitsk ● Trujillo ● UiB ● UNAM ● Wuhan
- Yerevan

GRID OPERATION PRINCIPLE



- The VO-box system (very controversial in the beginning)
 - Has been extensively tested
 - Allows for site services scaling
 - Is a simple isolation layer for the VO in case of troubles


OPERATION – CENTRAL/ SITE SUPPORT

- Central services support (2 FTEs equivalent)
 - There are no experts which do exclusively support – there are 6 highly-qualified experts doing development/support
- Site services support - handled by ‘regional experts’ (one per country) in collaboration with local cluster administrators
 - Extremely important part of the system
 - In normal operation ~0.2FTEs/site
- Regular weekly discussions and active all-activities mailing lists

If you want to know **EVERYTHING!**
All the questions
you did not dare asking




Software Development in HEP (F.Carminati, CERN)





Virtualisation in GRID computing (P.Buncic, CERN)

Evolution of Parallel Computing in HEP (F.Rademakers, CERN)



Aspects of Internet Law for HEP Software Developers (L.Pinsky, University of Houston)

Databases in High Energy Physics (J.Shiers, CERN)



Towards a globalised data access (F.Furano, CERN)

The planetary brain (G.Galli Carminati, HUG, Geneva University)



