

Release size and deletion

David Rousseau (LAL-Orsay)



David Rousseau, release size
and deletion, SIT, 2th December
2010

Rel	Production	kJob	Decision	Feedback
16.3.0	Simu/Muon val		Keep	
16.2.0	HI Tier0+ana		Keep	
16.0.2	Fall repro		Keep	
16.0.1		3416	Del 15 th Nov	
16.0.0		207	Del 15 th Nov	
15.9.0		1383	Del 30 th Nov	Used by exotics=>deletion postponed to 30 th nov
15.8.0		1147	Del 30 th Nov	Used by ID alignment=>deletion postponed to 30 th nov
15.6.13	pp Tier0 +ana	270	Keep	
15.6.12	MC10 Geant4	3510	Keep	
15.6.11	evgen	1217	Keep	
15.6.10	MC09 evgen	2837	Keep	
15.6.9	(may repro)	4271	Keep	
15.6.8		122	Del 30 th Nov	
15.6.7	(april repro)	278	Del 30 th Nov	
15.6.6		313	Del 30 th Nov	
15.6.5		638	Del 30 th Nov	
15.6.3	MC09 Geant4	460	Keep	
15.6.1		21	Del 15 th Nov	
15.6.0		0	Del 15 th Nov	
15.5.6		0	Del 15 th Nov	
15.5.4		11	Del 15 th Nov	
15.3.1		23	Del 15 th Nov	Higgs thesis=>keep at Freiburg, GridKA and DESY
14.2.25		39	Del 15 th Nov	Higgs thesis=>keep at Freiburg, GridKA and DESY

Release deletion status

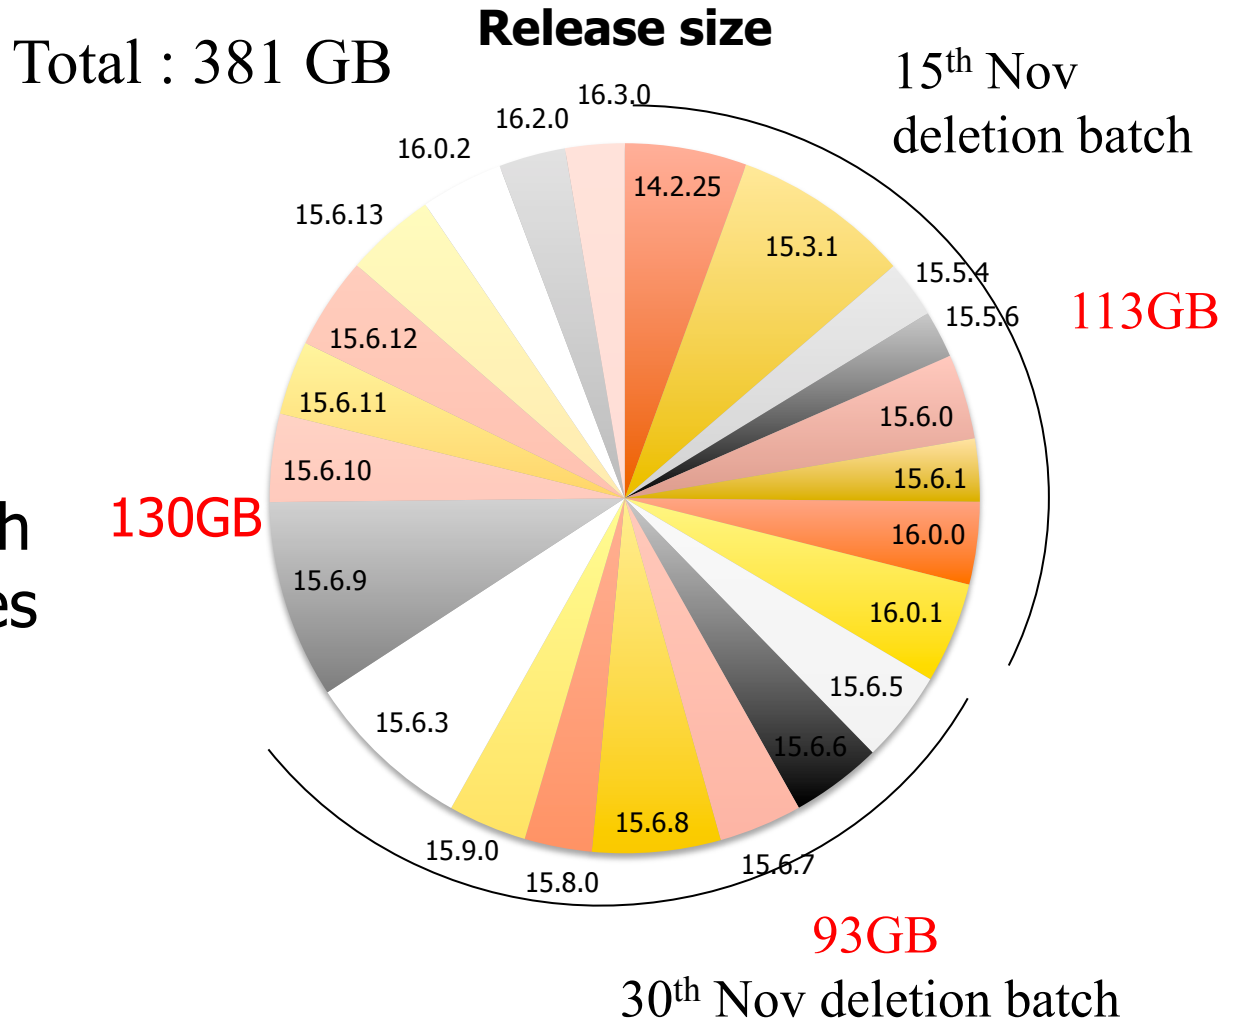


- ❑ Release deletion feedback went rather well (so far!)
 - One group asked for two old releases to be kept at three sites
 - One physics group moved away from 15.9.0
 - One detector group moved away from 15.8.0
- ❑ 15th November batch deleted
- ❑ 30th November on-going
- ❑ CERN Releases area to be cleaned-up (so that it matches the list of grid supported releases) in:<http://atlas-computing.web.cern.ch/atlas-computing/projects/releases/status/>

Release deletion



- ❑ Analysis of CERN releases area beg November
- ❑ "Releases" area at CERN should match release area in sites (?)
- ❑ 15th Nov deletion on going

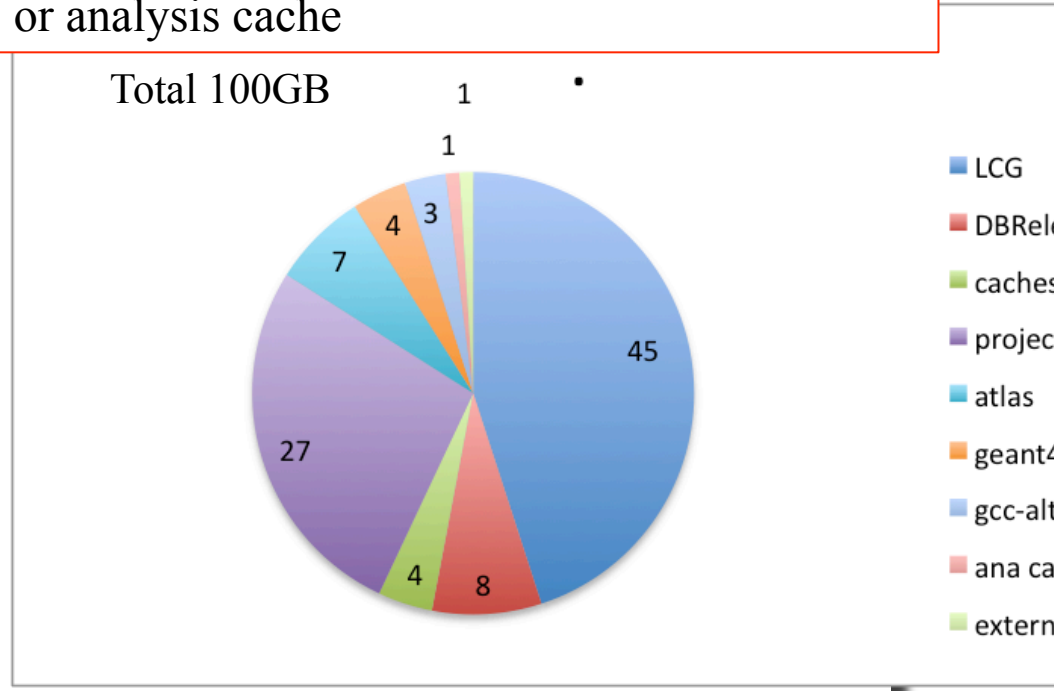
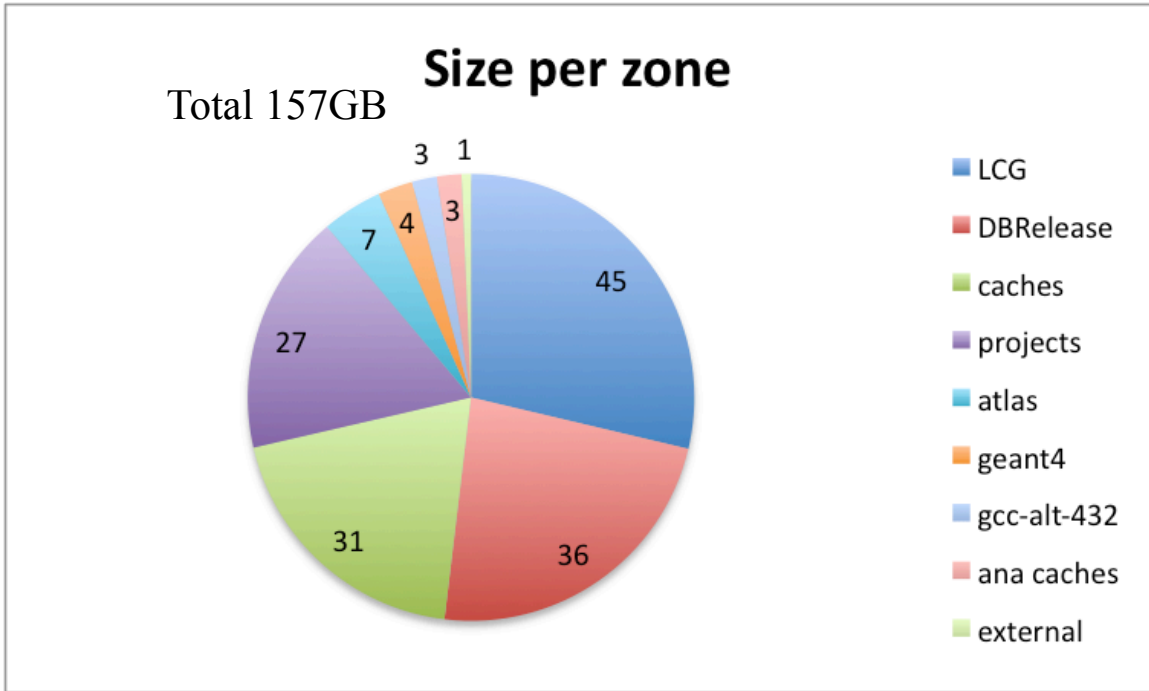


Release size analysis



- ❑ Examined in details release sizes at CERN, for releases that are not scheduled to be deleted (9 releases)
- ❑ Again caveat that release area at CERN might not match release area at sites

Now suppose in each release we only keep the latest DBRelease, AtlasProduction cache, or analysis cache



Conclusion



- ❑ After November clean-up size should be < 200GB (on a quota of 300Gb) for 9 releases
- ❑ But many releases to come
- ❑ Cern vmfs (release caching) to be enabled but will benefit from smaller release as well
- ❑ By default, releases are self contained (no sharing between releases)
- ❑ DBReleases : 36Gb to 8GB just keeping the last one for each release (zero loss of functionality since DBRelease n+1 include DBRelease n). Could even imagine removing the last one for sites where DBRelease is put on a hot disk. Could also imagine having only 1 DBRelease in total (<1GB).
- ❑ AtlasProd caches : 31 to 4 GB keeping only the last one for each release. Doable but might be a management hassle
- ❑ Analysis caches : 3 to 1 GB keeping only the last one for each cache type/release. Not worth the hassle.
- ❑ "atlas" area : 7GB, 0.8Gb per release a few large files : small test RDO file, magnetic field map and especially frozen shower library, maybe some of these files could be put in a separate area, shared between releases. Or should be handled like conditions file
- ❑ LCG area : 45GB, 4.5GB per release !!! Largest : MCGenerator 2GB !, Java 0.9Gb! , LCG-APP 0.5Gb, qt 0.3GB, then many <1Gb