

# Victor\* agent

**Andrii Tykhonov, Fernando Bareiro Megino, Stephane Jezequel, Simone Campana, Angelos Molfetas, Vincent Garonne, I. Ueda**

\* Cleaning agent to select dataset replicas to be deleted

November 29, 2010  
CERN

# Outline

---

- Introduction
- Functionality of the agent
  - Selection of overfull sites
  - Cleaning algorithms
- Problems
- Summary and outlook

# Introduction

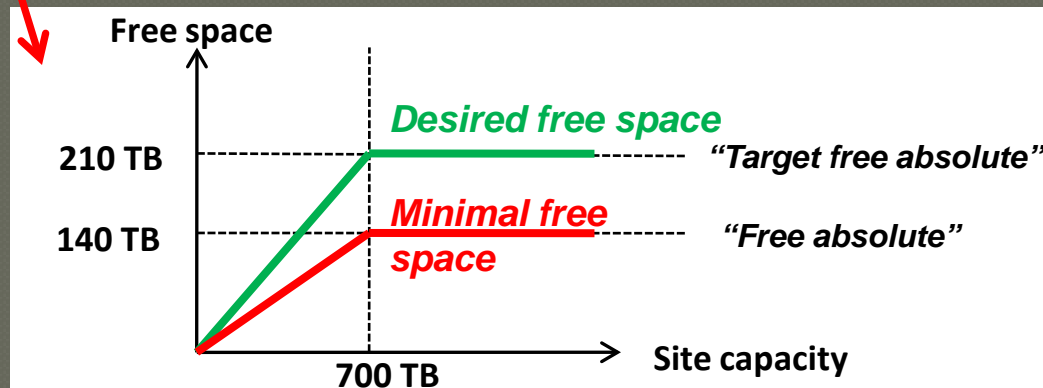
---

- “**Victor**” - dq2 agent, aimed to clean grid storage spaces.
- Victor doesn't delete anything – but it declares dataset replicas to be deleted by central deletion system
- Victor sets the lifetime of the replicas (to 1 day), and replicas type to “archived = ToBeDeleted”
- The central deletion system delete the replicas when the lifetimes are expired
- Currently it runs on DATADISKS, MCDISKS and SCRATCHDISKS

# Functionality: *Select sites, which are overfull*

- Currently victor runs once per day
- It considers site as “overfull” according to a predefined thresholds, which are specified for each spacetoken separately:

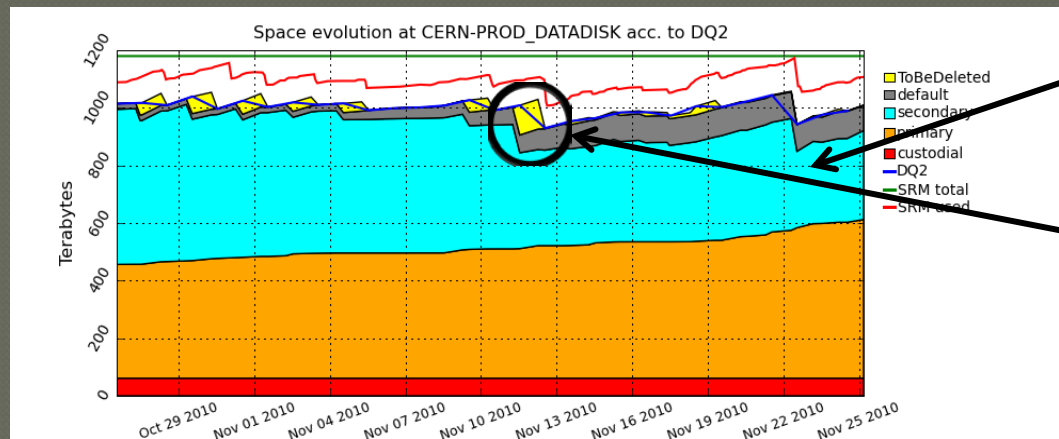
Spacetoken	Free ratio	Free absolute	Target free ratio	Target free absolute	Big site threshold
default	10%	25.0 TB	20%	50.0 TB	250.0 TB
MCDISK	20%	140.0 TB	30%	210.0 TB	700.0 TB
SCRATCHDISK	25%	20.0 TB	30%	24.0 TB	80.0 TB
DATADISK	24%	400.0 TB	30%	500.0 TB	1666.7 TB



- For “small” sites thresholds are calculated as a fraction of total space, while for “big” sites they are defined as fixed absolute values in TB.
- It is also possible to define a fixed absolute thresholds for particular sites

# Functionality: *Select replicas to delete*

- Selection of unpopular datasets for DATADISK / MCDISK:
  - Select “secondary” datasets which were not accessed during last month
  - If it's not enough to clean site → consider replicas accesses <10 times during last month, etc.
  - The last step is: ..... <1000 times \*



“Secondary”

“To be deleted”

- \* A single job accessing a dataset with 1000 files can be counted as 1000 accesses

# Functionality:

## *Select replicas to delete*

---

- SCRATCHDISK:

- Clean any replica (older than 15 days) until the desired free space will be reached.

**For both algorithms:**

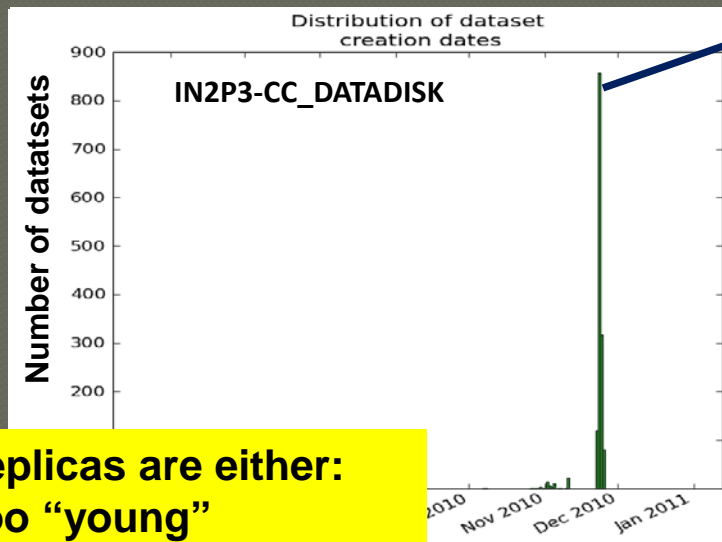
- “ToBeDeleted” replicas are taken into account as a part of free space
- replicas are cleaned in order of their age – **from oldest to newest**
- all considered replicas must be older than 15 days

# Problems: Not cleaned sites

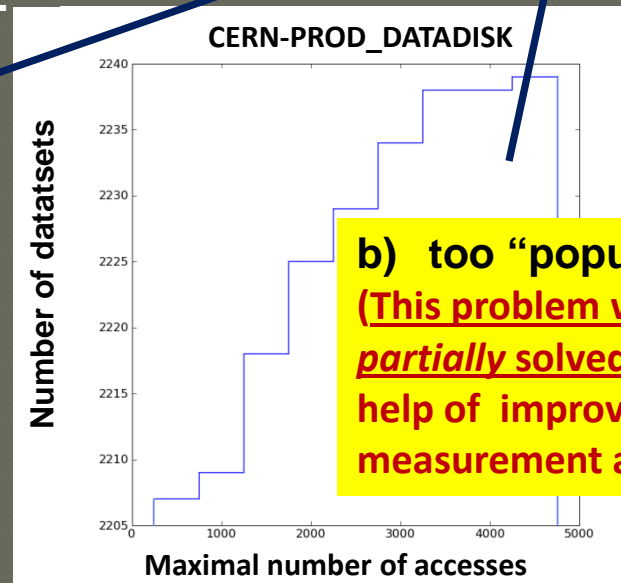
- [http://atlas-ddm-cleaning.cern.ch/victor/diskspacemonitor\\_deletion.html](http://atlas-ddm-cleaning.cern.ch/victor/diskspacemonitor_deletion.html):

Cloud	Tier	Site name	Total space	"ToBeDeleted" space	Free space (before)	Cleaned space	Free space (after)	"Secondary" space: Young/Popular	Site info	Cleaned**		
CANADASITES	Tier2	<a href="#">CA-ALBERTA-WESTGRID-T2 DATADISK</a>	77.00 TB	11.81 TB	16.85 TB	22%	6.86 TB	23.71 TB	31%	<a href="#">html</a> <a href="#">txt</a>	True	
		<a href="#">CA-ALBERTA-WESTGRID-T2 SCRATCHDISK</a>	3.50 TB	0.00 TB	0.83 TB	24%	0.24 TB	1.06 TB	31%	<a href="#">html</a> <a href="#">txt</a>	True	
CERN	Tier2	<a href="#">CERN-PROD DATADISK</a>	1179.63 TB	1.49 TB	73.99 TB	7%	12.51 TB	86.50 TB	8%	184.11 / 116.68 TB	<a href="#">html</a> <a href="#">txt</a>	False
FRANCESITES	Tier1	<a href="#">IN2P3-CC DATADISK</a>	1378.79 TB	11.63 TB	150.83 TB	11%	13.10 TB	163.93 TB	12%	159.67 / 35.26 TB	<a href="#">html</a> <a href="#">txt</a>	False
	Tier2	<a href="#">GRIF-IRFU DATADISK</a>	98.96 TB	4.36 TB	7.37 TB	8%	2.94 TB	10.31 TB	11%	7.57 / 33.23 TB	<a href="#">html</a> <a href="#">txt</a>	False

- Some sites can't be cleaned, because:



Most replicas are either:  
a) too "young"



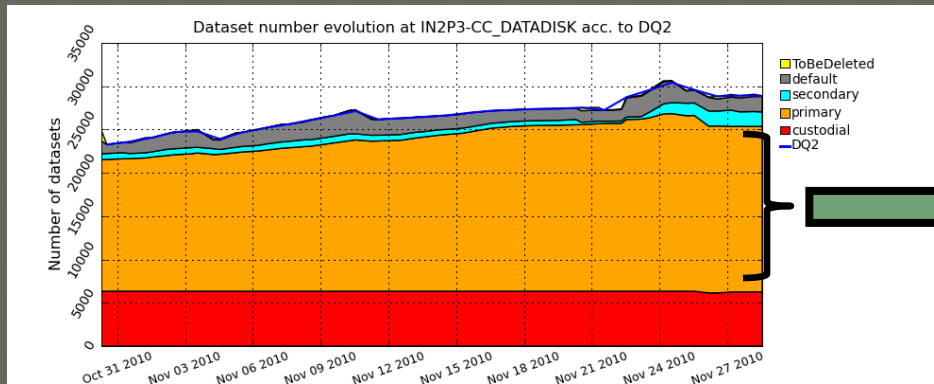
b) too "popular"  
(This problem will be partially solved with the help of improved popularity measurement algorithms)



# Problems:

## What site should do if it remains full?

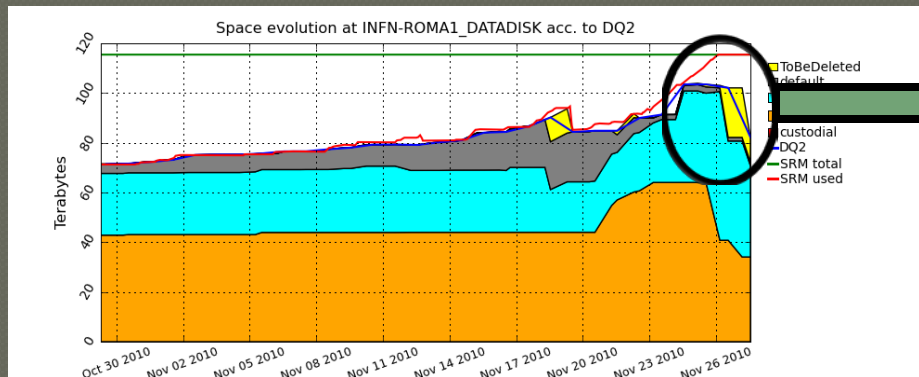
- Check the content of primary replicas:



Most of the replicas are primary → site can't be cleaned

- Check the deletion rate:

[http://bourricot.cern.ch/dq2/deletion/#period=1&menu\\_cloud=ITALYSITES&menu\\_site=INFN-ROMA1](http://bourricot.cern.ch/dq2/deletion/#period=1&menu_cloud=ITALYSITES&menu_site=INFN-ROMA1)



INFN-ROMA1\_DATADISK is not cleaned because deletion rate is not enough

- Check the amount of dark data in your site  
**SCRATCHDISK** when datasets are not frozen)

(there is a problem with

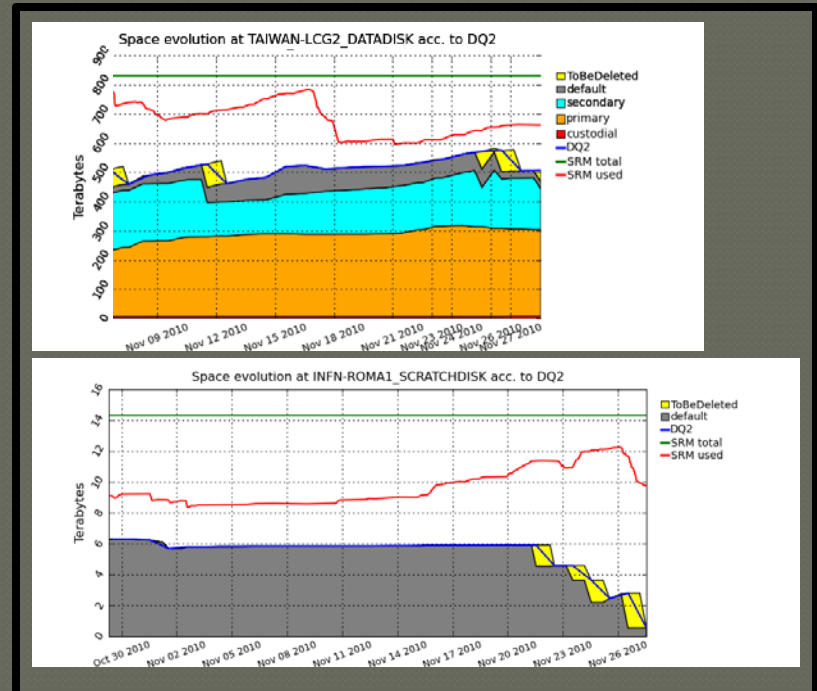
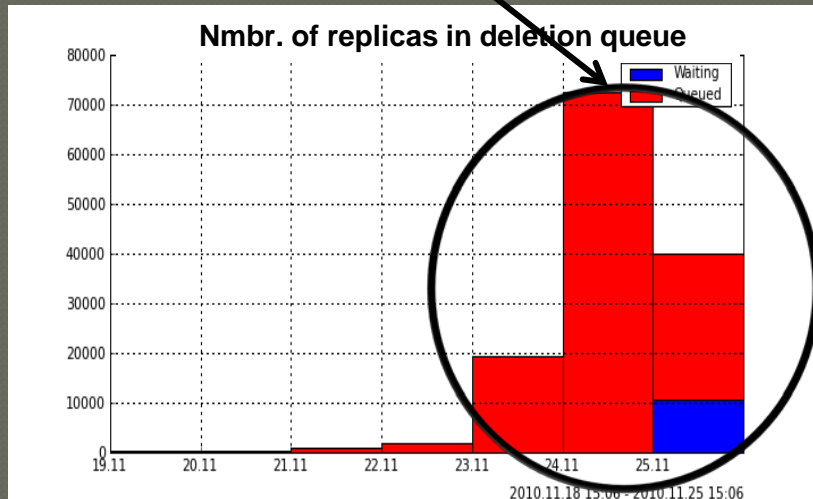


# Problems: Overcleaning of sites

- Overcleaning of sites:

Datasets in the deletion queues are not visible any longer and Victor doesn't know about it

Overcleaning occurs!



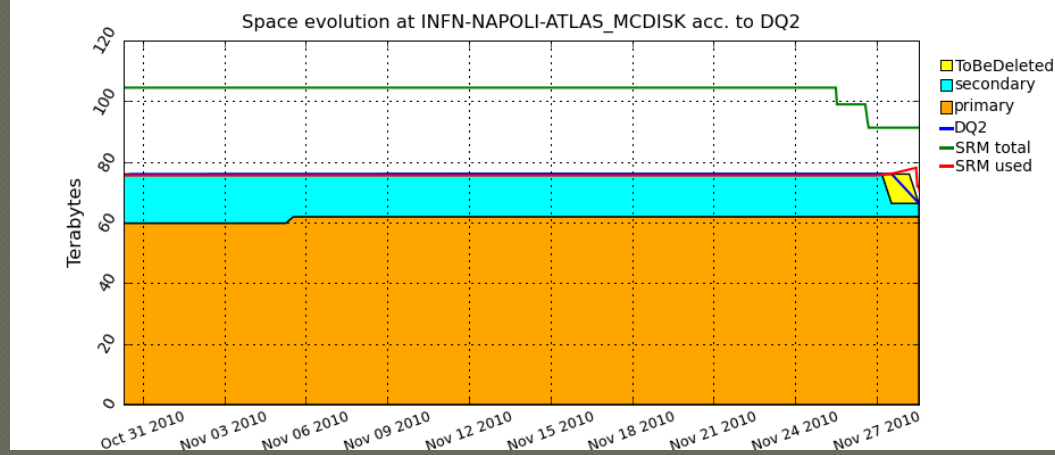
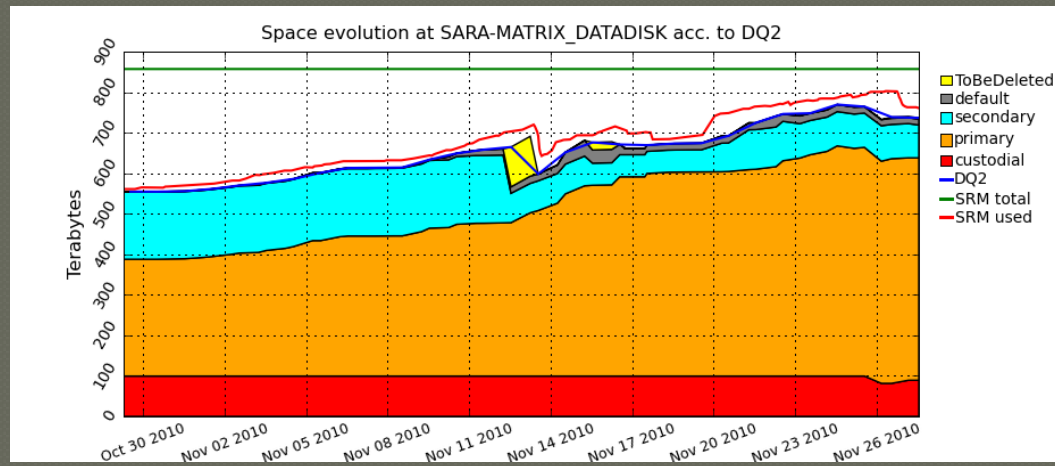
This problem will be solved in a nearest time with the help of new api, which allows victor to account for the data which is in the deletion queue.



# Backup slides

---

# Problems



Sites can't be cleaned because they contain too few of "secondary" replicas.