

Table of Contents

The priorityuser role IS NOW DEPRECATED, OBSOLETE, NOT FUNCTIONING.....	1
Policy and Purpose.....	1
Conveners' Responsibility.....	1
Priorityusers' Responsibility.....	3
Tier-2 Sites' Responsibility.....	4
The StoreResults Service.....	5
Purpose of a migration from /store/group/ to /store/results/.....	5
Creating a StoreResults Request.....	5
Old-Savannah conveners.....	5
Monitoring of the progress of the StoreResults tasks.....	6
Ensure the data is located at global DBS and CMS.PhEDEX DB.....	6
Instructions for groups to use CMS.PhEDEX replication requests from Tier-2 to Tier-2.....	7
Review status.....	7

The priorityuser role IS NOW DEPRECATED, OBSOLETE, NOT FUNCTIONING

Policy and Purpose

PAGs, POGs, and DPGs are allowed to specify a limited number of people who will have special privileges at the Tier-2s, both for CPU as well as write access to storage. The special privileges entail:

- priority batch CPU access to half of the 50% of the pledged resources for analysis at all Tier-2s.
- write access into /store/group/ at all Tier-2s associated to that group.

We assume that on average a lot less than 25% of the CPU resources at the Tier-2s will be used that way. The policy is thus set up for a burst but not average use. We will monitor usage.

The space used in /store/group/ will have to be accounted by the site's system administrators towards the quote the group has at that site.

Conveners' Responsibility

The conveners fill out the following table to indicate their priority users, and then send an email to the AnalysisOps hypernews for information, also mention who has been removed from the list. They do this every time they want to change something. However, we would like the groups to not change the names more often than once every few months. We will try to keep the actual privileges available as implemented in the CMS VOMS in synch with this Twiki table. It will be the responsibility of the people listed there to subscribe to this role using the CMS VOMRS interface. It should take not more than a few working days to get an approval email. The dynamic list of people authorized for the priorityuser role is here

- user's DN is the output of the command: `grid-proxy-info -identity` (this will be used to access the StoreResults status web information, see below)

Phedex group name	DN	given name	email address
AnalysisOps	/DC=org/DC=doegrids/OU=People/CN=Frank Wuerthwein 699373	Frank Wuerthwein	fkw at ucsd dot edu
AnalysisOps	/O=GermanGrid/OU=RWTH/CN=Thomas Kress	Thomas Kress	thomas dot kress at cern dot ch
AnalysisOps	/C=IT/O=INFN/OU=Personal Certificate/L=Trieste/CN=Stefano Belforte	Stefano Belforte	stefano dot belforte at ts dot infn dot it
AnalysisOps	/DC=org/DC=doegrids/OU=People/CN=Sanjay Padhi 496075	Sanjay Padhi	sanjay dot padhi at cern dot ch
AnalysisOps	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=dlevans/CN=638933	David Evans	d dot evans at cern dot ch
AnalysisOps	/O=GermanGrid/OU=RWTH/CN=Manuel Giffels	Manuel Giffels	manuel dot giffels at cern dot ch
tau	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=slehti/CN=408791/CN=Sami Lehti	Sami Lehti	sami dot lehti at cern dot ch
tau		Roger Wolf	

WorkbookGroupActivities < CMSPublic < TWiki

	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=rwolf/CN=667992/CN=Roger Wolf		roger dot wolf at cern dot ch
b-physics	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=griley/CN=731209/CN=Grant Riley	Grant Riley	gvriley at gmail dot com
b-physics	/C=CN/O=HEP/O=PKU/OU=PHYS/CN=Linlin Zhang	Linlin Zhang	Linlin dot Zhang at cern dot ch
b-physics	/DC=org/DC=doegrids/OU=People/CN=Elif Asli Albayrak 560229	Elif Asli Albayrak	eea at cern dot ch
exotica	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=srimanob/CN=628708/CN=Phat Srimanobhas	Phat Srimanobhas	srimanob at cern dot ch
ewk	/C=IT/O=INFN/OU=Personal Certificate/L=Napoli/CN=francesco fabozzi	Francesco Fabozzi	francesco dot fabozzi at cern dot ch
ewk	/C=GR/O=HellasGrid/OU=demokritos.gr/CN=Stilianos Kesisoglou	Stelios Kesisoglou	stilianos dot kesisoglou at cern dot ch
forward	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=mpetek/CN=676136/CN=Marko Petek	marko petek	marko dot petek at cern dot ch
top	/C=DE/O=GermanGrid/OU=RWTH/CN=Felix Hoehle	Felix Höhle	felix dot hoehle at cern dot ch
top	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gbenelli/CN=485942/CN=Gabriele Benelli	Gabriele Benelli	gabriele dot benelli at cern dot ch
tracker-pog	/DC=org/DC=doegrids/OU=People/CN=Boris Mangano	Boris Mangano	boris dot mangano at cern dot ch
tracker-pog	/DC=org/DC=doegrids/OU=People/CN=Brian R Drell 388439	Brian Drell	brian dot drell at cern dot ch
muon	/C=IT/O=INFN/OU=Personal Certificate/L=Bologna/CN=Alessandra Fanfani	Alessandra Fanfani	alessandra dot fanfani at cern dot ch
moun	/C=IT/O=INFN/OU=Personal Certificate/L=Sns/CN=Giovanni Petrucciani	Giovanni Petrucciani	giovanni dot petrucciani at cern dot ch
muon	/DC=org/DC=doegrids/OU=People/CN=Andrew Kubik 945511	Andrew Kubik	andrew dot michael dot kubik at cern dot ch
susy	/C=DE/O=GermanGrid/OU=uni-hamburg/CN=Lukas Vanelderren	Lukas Vanelderren	lukasvanelderren at gmail dot com
susy	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=alkaloge/CN=654885/CN=Alexis Kalogeropoulos	Alexis Kalogeropoulos	Alexis dot Kalogeropoulos at cern dot ch
higgs	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=jfernan/CN=529280/CN=Javier Fernandez Menendez	Javier Fernandez Menendez	javier dot fernandez at cern dot ch
higgs	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=klute/CN=529652/CN=Markus Klute	Markus Klute	markus dot klute at cern

			dot ch
qcd	/DC=org/DC=doegrids/OU=People/CN=Vanessa Gaultney 303904	Vanessa Gaultney	vanessa dot gaultney at cern dot ch
qcd	/DC=org/DC=doegrids/OU=People/CN=Suvadeep Bose 616377	Suvadeep Bose	suvadeep dot bose at cern dot ch
qcd	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=yuanchao/CN=596728/CN=Yuan Chao	Yuan CHAO	yuan dot chao at cern dot ch
jets-met_hcal	/DC=org/DC=doegrids/OU=People/CN=Konstantinos Kousouris 2271	Konstantinos Kousouris	konstantinos dot kousouris at cern dot ch
jets-met_hcal	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=delre/CN=651284/CN=Daniele Del Re	Daniele del Re	daniele dot del dot re at cern dot ch
e-gamma_ecal	/C=UK/O=eScience/OU=Imperial/L=Physics/CN=nikolaos rompotis	Nikolaos Rompotis	nikolaos dot rompotis at cern dot ch
e-gamma_ecal	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=meridian/CN=555986/CN=Paolo Meridiani	Paolo Meridiani	paolo dot meridiani at cern dot ch
b-tagging	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=adamw/CN=369199/CN=Wolfgang Adam	Wolfgang Adam	wolfgang dot adam at cern dot ch
b-tagging	/DC=org/DC=doegrids/OU=People/CN=Pratima Jindal 522031	Pratima Jindal	pratima dot jindal at cern dot ch
generators	/C=IT/O=INFN/OU=Personal Certificate/L=Trieste/CN=Fabio Cossutti	Fabio Cossutti	fabio dot cossutti at cern dot ch
heavy-ions	/DC=org/DC=doegrids/OU=People/CN=Wei Li 344273	Wei Li	wei dot li at cern dot ch
B2G	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=bazterra/CN=667990/CN=Victor Eduardo Bazterra	Victor Bazterra	baites at fnal dot gov
B2G	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=ajkumar/CN=718734/CN=Ajay Kumar	Ajay Kumar	a dot kumar at cern dot ch
B2G	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=jdangov/CN=537562/CN=Jordan Dangov	Jordan Dangov	jdangov at fnal dot gov
B2G	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=jdolen/CN=674229/CN=James William Dolen	James Dolen	james dot dolen at gmail dot com

Priorityusers' Responsibility

- Request permission to use the role
 - ◆ following these instructions
- Create and use a proxy with the priorityuser Role when doing high priority work
 - ◆ following these instructions
 - ◆ that proxy is needed for every crab command, whenever you logout or proxy expires, do it again

- ◆ make sure you add to `crab.cfg`, in the `GRID` section the line
 - ◇ `role = priorityuser`
- Select the appropriate stage-out directory `/store/group/"groupname"` at one of your group's `"T2_XX_YY"`
 - ◆ for the correct spelling of the "groupname" follow the CMS.PhEDEx group convention [?](#)
 - ◆ download `GetStorageCfgFromTFC.py` [?](#) into a CMSSW release area
 - ◆ execute `python GetStorageCfgFromTFC.py --site="T2_XX_YY" --lfn=/store/group/"groupname"` (replace `"T2_XX_YY"` and `"groupname"`)
 - ◆ include the script output in the `USER` section of your `crab.cfg`
- Beware: by default all your shell sessions on the same host use the same proxy, do not try to operate at same time on normal task and priorityuser tasks (which is possible, but it is an expert operation and if you have no idea how to do it properly, it is better you do not do it). Using different roles on different machines (even simply two lxplus nodes) is OK.

Tier-2 Sites' Responsibility

In order to facilitate data transfers of physics results between Tier-2 sites, DDT has organized the commissioning of data transfer links in PhEDEx between Tier-2 sites associated with the same analysis group. Details can be found [here](#). Priorities for link commissioning are suggested by the groups themselves. In practice, if groups have one or two Tier-2 sites with good connectivity in both directions to all of the others in a group, then data can be transferred via these "hub" Tier-2 sites between any of the other Tier-2's with at most 1 "hop" via the hub. Tier-2s should make sure that the `-ignore` and `-accept` options in their `FileExport` and `FileDownload` agents of PhEDEx are set up properly to enable transfers from other Tier-2s in their groups.

Please grant write access for VOMS CMS `role=production` to `/store/results/` and for `role=priorityuser` to `/store/group/`. Please also check your TFC for these new LFN conventions. In addition, for data replications, PhEDEx needs to have read and write access on `/store/results/`.

Please configure your batch system that the share from the CMS pledges is 50% for `role=production`, the other half to be distributed 25% for `role=priorityuser` and 25% for "normal" CMS users. The appropriate way for the setup is to use a mapping to Pool Accounts in all three cases.

The StoreResults Service

Purpose of a migration from /store/group/ to /store/results/

Although group skims or e.g. special group formats like PATs produced by a priorityuser, stored in /store/group// and registered in an analysis local-scope DBS instance can be accessed by any CMS Grid user through Crab, these samples are not official CMS data and are neither transferable to other sites using CMS.PhEDEx nor included in the groups' disk quota accounting. The StoreResults service elevated such a data set if approved by one of the performance or physics groups.

Important:

Only datasets published with > CRAB_2_6_0 EDM can be migrated using the Store Results service. The files may not contain group specific data formats which are not in official CMSSW releases.

Important: At the moment growing datasets are not supported. You can not make another request for elevation with the same dataset name, if you produce more similar data you need to do something like adding _v1, _v2, _v3 etc. to the name.

Creating a StoreResults Request.

We were previously using the Savannah Interface to request and approve T2-T2 dataset migrations, including injection into GlobalDBS and CMS.PhEDEx using the StoreResults service. However Savannah is now **retired**. Therefore you should open a CMS GGUS (https://ggus.eu/?mode=ticket_cms) ticket instead with the following information:

- **Subject:** Something like "Dataset elevation" should do.
- Write in the body of the ticket the following information needed:
 - ◆ **Input dataset:** it has to be located at the **same Tier-2** which is used to finally hold the group data in /store/results/
 - ◇ Please use the dataset name published in your local-scope DBS
 - ◇ Please create **one request** for **each transfer**
 - ◆ **DBS URL:** For example: "phys01" for <https://cmsweb.cern.ch/dbs/prod/phys01/DBSReader>. Do **not** insert the DBS url of the writer instance for example <https://cmsweb.cern.ch/dbs/prod/phys01/DBSWriter>.
 - ◆ **CMSSW release:** which should be used for merging step. In general, use always the version used for the dataset production, if the version is outdated please provide one that is able to read your data.
 - ◆ **Physics Group:** requesting the migration, this will determine the subdirectory below /store/results/ and sets the appropriate Phedex accounting group tag.
- **Type of Problem:** Select "CMS_Central Workflows"
- **CMS Support Unit:** Select "CMS Workflows"

Here is an awesome example of a ticket: [110773](#)

- The name of the dataset will be changed by the StoreResults service. Please see the section below for a description how to find your dataset once the migration is done

Old-Savannah conveners

Physics group	Allowed to approve	Savannah account	Allowed ...	Savannah ...	Allowed ...	Savannah ...	Allowed to ...	Savannah ...
---------------	--------------------	------------------	-------------	--------------	-------------	--------------	----------------	--------------

b-physics	Hermine Woehri	yes	Alessandra Fanfani	yes				
b-tagging	Wolfgang Adam	yes	Andrea Rizzi	yes				
beyond-two-generations	Freya Blekman	no?	Salvatore Rappoccio	yes?	Victor Bazterra	yes?		
e-gamma_ecal	Chris Seez	yes	Paolo Meridiani	yes	Giovanni Franzoni	yes	Tommaso Tabarelli de Fatisi	no
ewk	Markus Klute	yes	Michael Schmitt	yes	Hamel de Montch. Gautier	yes		
exotica	Steve Worm	yes	John Paul Chou	yes				
forward	Krzysztof Piotrkowski	yes	Grzegorz Brona	yes				
heavy-ions	Olga Kodolova	yes	Gunther Roland	no				
higgs	Markus Klute	yes	Javier Fernandez	yes				
jets-met_hcal	Victor Daniel Elvira	no	Philipp Schieferdecker	yes	Frank Chlebana	yes	Vladimir Gavrillov	no
muon	Slava Valuev	yes	Riccardo Bellan	yes				
qcd	Vivian ODell	yes	Nikos Varelas	no				
susy	Chris Silkworth	yes	Alexis Kalogeropoulos	yes	Jim Smith	yes	Dominick Olivito	yes
tau-pflow	Colin Bernet	yes	Simone Gennai	yes				
top	Roberto Tenchini	yes	Roberto Chierici	yes	Stijn Blyweert	yes	Michael Maes	yes
tracker-dpg	Christophe Delaere	yes	Thomas Speer	yes				
tracker-pog	Kevin Burkett	yes	Boris Mangano	yes				
trigger	Sridhara Dasu	yes	Christos Leonidopoulos	yes				

Monitoring of the progress of the StoreResults tasks

At the moment we use two StoreResults instances at RWTH Aachen and Imperial College London, task IDs which can be divided by 3 are handled by IC, the other 2/3 by Aachen. For both instances the status of the tasks can be monitored [↗](#) by the priorityusers. We have imposed an access restriction based on Grid DNs given above to restrict the database queries, if you don't have access please send an Email with your DN to Manuel Giffels and Thomas Kress.

Ensure the data is located at global DBS and CMS.PHEDEX DB

First the StoreResults tool merges the /store/group/ files and stores the output, using file sizes appropriate for a performant data transfer, at the /store/results/ directory and publishes the data set in global DBS. After this step the dataset is visible by a global DBS query [↗](#) and can already be used for user analyses. Usually the original dataset in /store/group/ should then be removed with the help of the local Tier-2 data managers or by using Grid file deletion commands to clean the buffer space. Unfortunately the local-scope DBS dataset

registration is not cleaned automatically in parallel and at the moment we can not suggest an easy way the user can clean-up the local-scope DBS entries. (NB: DBS is deprecated, in favor of DAS. An example DAS query is shown below.)

- Important:

Do not omit the **StoreResults** part for a Global DBS query!

- Assuming your dataset is called

```
/TTbar_1jet/top-ProcessedDSName/USER
```

- DAS query example:

```
dataset=/TTbar_1jet/StoreResults-ProcessedDSName/USER
```

Store Results also injects the dataset into the CMS.PhEEx database, which you can query. Once this is finished after another approval step which is done centrally during the October Exercise, the dataset is available for transfers and included in the quota group accounting.

- CMS.PhEEx query example (Perl syntax): first "select all" or specific Tier-2s in "Nodes Shown" from "Show options" menu, then enter

```
/TTbar_1jet.*/StoreResults.*
```

The groups should inform their users once new datasets in /store/results/ have become accessible for analysis and transfers.

Instructions for groups to use CMS.PhEEx replication requests from Tier-2 to Tier-2

After StoreResults the dataset in /store/results/ has become an official CMS dataset. It can be transferred to other CMS Grid sites like any other dataset located in /store/data/ or /store/mc/. Groups or users can request a replication to other Tiers, which have to be approved by the corresponding Tier data managers. If you want to move the data from one Tier-2 to another, you have to first make a CMS.PhEEx replication request followed by a deletion request for the initial site, Phedex move requests are presently not enabled among T-2s. Please note that if you replicate a /store/results/ data set from another group Y to one of your group X Tier-2 site, the LFN will be /store/results// since CMS.PhEEx preserves the structure. This is however transparent to the user and for the CMS.PhEEx group accounting.

To request a migration search for the dataset in Global DBS and follow the CMS.PhEEx link below the dataset name.

Review status

Reviewer/Editor and Date (copy from screen)	Comments
JohnStupak - 9-September-2013	Review with minor changes
Luis Contreras - 4-March-2014	Update for DBS transition

-- ThomasKress - 06-Aug-2010

This topic: CMSPublic > WorkBookGroupActivities

Topic revision: r60 - 2019-06-18 - StefanoBelforte



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
Ideas, requests, problems regarding TWiki? Send feedback