

# Zeus User Meeting

June 6<sup>th</sup>, 2012

# Agenda

- Upcoming downtimes
- Known Issues
- Managing Data on Zeus
- Q&A

# Downtime Schedule

- No schedule downtimes in the near future
- There will be power maintenance on June 20<sup>th</sup>. We intend to use generator power during this time and provide continued access to the system.
  - If this changes, we will let you know.

# NCEP directory links

- The links in the root filesystem (/com, /ptmp, etc) are temporary and will be removed soon.
- We will be providing links for you through a special ncep module. This will be installed in the next week.

Variable Name	Path
NCEPCOM	/scratch2/portfolios/NCEPDEV/rstprod/com
NCEPNCEP1	/scratch1/portfolios/NCEPDEV
NCEPNCEP2	/scratch2/portfolios/NCEPDEV
NCEPPROD	/scratch2/portfolios/NCEPDEV/rstprod
NCEPNWPROD	/scratch2/portfolios/NCEPDEV/rstprod/nwprod
NCEPPTMP	/scratch2/portfolios/NCEPDEV/ptmp
NCEPSTMP	/scratch2/portfolios/NCEPDEV/stmp

Define these paths as variables in your script and reference the paths. DO NOT HARDCODE paths across your scripts. This is not portable.

[https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Using\\_environment\\_variables\\_for\\_paths](https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Using_environment_variables_for_paths)

# Outstanding Issues

- Details of issues and current status can be found at:
  - [https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Known\\_Issues](https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Known_Issues)
  - We'll take a look at this site...

# Managing Data on Zeus

- We have implemented several tools to help users manage data on Zeus
  - Rstprod utilities
  - Account\_params
  - Squotas
  - Sdu

# RSTPROD Utilities

- It is not sufficient to just set the group on restricted files to rstprod to protect them.
  - A lower level interface must be used, Access Control Lists
- To simplify the use of this interface, we have created tools for the users to manage their data:

Tool Name	Description
tag_rstprod	Tag a file as restricted
untag_rstprod	Untag a file as restricted
ls_rstprod	List restricted status of files

[https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Protecting\\_Restricted\\_Data](https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Protecting_Restricted_Data)

# Utility: tag\_rstprod

```
[Joe.Smith@tfel rsttest]$ ls -al
total 44
drwxr-xr-x 2 Joe.Smith global 4096 Mar 27 00:15 dir1
drwxr-xr-x 2 Joe.Smith global 4096 Mar 27 00:15 dir2
-rw-r--r-- 1 Joe.Smith global    6 Mar 27 00:15 some-restricted-file1
-rw-r--r-- 1 Joe.Smith global    6 Mar 27 00:15 some-restricted-file2
[Joe.Smith@tfel rsttest]$ tag_rstprod some-restricted-file1
[Joe.Smith@tfel rsttest]$ ls -al
total 44
drwxr-xr-x  2 Joe.Smith global 4096 Mar 27 00:15 dir1
drwxr-xr-x  2 Joe.Smith global 4096 Mar 27 00:15 dir2
-rw-r-----+ 1 Joe.Smith global    6 Mar 27 00:15 some-restricted-file1
-rw-r--r--  1 Joe.Smith global    6 Mar 27 00:15 some-restricted-file2
```

While the `ls_rstprod` utility will show which files are tagged as restricted, you can also see this with the '+' symbol in the directory attributes. The '+' indicates there are additional ACLs applied to this file.

# Utility: ls\_rstprod

```
[Joe.Smith@tfe1 rsttest]$ ls_rstprod  
dir1 is not restricted  
some-restricted-file1 is restricted  
some-restricted-file2 is not restricted
```

# Rstprod Utilities Notes

- Wildcards are permitted to apply the operation to a set of files.
- Supported Options:
  - -h – Print Help
  - -R – Apply command recursively
  - -v – Verbose
  - --verbose – Display Version
- For tag\_rstprod
  - -g – Set group. This resets the group of the file to the group specified

# Rstprod Utility Implementation Schedule

- 6/5 – Tools are installed and ready for use
  - Start using them!
- 6/11 – One pass will be made of entire system to convert all files with group ownership rstprod.
  - Groups will be reset to the project group of the directory which they exist
  - Files will be tagged via the rstprod utilities to be restricted

End of June when we are confident that the system is working and all users have been notified of the use of these files, we will turn on group scrubbing (reset groups to top level project) and the quotas will properly reflect data contained within the project directory.

# Obtaining disk usage information

- We have 3 programs to help determine how disk is being used by each user/project.
  - Account\_params – Provide high level overview of account information including disk usage and core allocations
  - Squotas – Smart Quotas – Provide quota usage per project and estimates of first level directory usage
  - Sdu – Smart du – Provide Unix du like information from filesystem scans.

# Utility: account\_params

- This tool has been available since users first logged in to Zeus
- Utility provides two pieces of project information
  - Core hours available
  - Disk usage and quota
- See:
  - [https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Getting\\_Information\\_About\\_Your\\_Account](https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Getting_Information_About_Your_Account)

# Utility: quotas

- The quotas tools provide more details as to the use of your assigned disk space on /home and the scratch filesystems.
- The tool also provides some insight as to your projects and which users are using the disk space.
- For project managers, it is a good tool to identify which users are using the most disk space within a project.

# squotas Report

```
[root@fel quotas]# squotas
Super quota report for user Craig.Smith, Wed Jun 6 03:43:33 2012 - usage and quota listed in GB
- RealUse is the current group disk usage as reported by the system
- EstUse is the amount measured by du, may not be accurate if group ownership is incorrect
```

Dir	Quota	RealUse	EstUse	PCT	Last Estimated
/home/Craig.Smith	10	3		30.9%	Wed Jun 6 03:43:33 2012
/scratch1/portfolios/BMC/nescmngmt	<b>225000</b>	<b>86208</b>	<b>86199</b>	38.3%	<b>Mon Jun 4 09:46:27 2012</b>
Christopher			1731	0.8%	Mon Jun 4 09:46:27 2012
djand			82615	36.7%	Mon Jun 4 09:46:27 2012
dnelson			279	0.1%	Mon Jun 4 09:46:27 2012
George			713	0.3%	Mon Jun 4 09:46:27 2012
gwv			47	0.0%	Mon Jun 4 09:46:27 2012
James			206	0.1%	Mon Jun 4 09:46:27 2012
letscher			3	0.0%	Mon Jun 4 09:46:27 2012
Steve			600	0.3%	Mon Jun 4 09:46:27 2012
Todd			1	0.0%	Mon Jun 4 09:46:27 2012
/scratch2/portfolios/BMC/sepp	5000	1987	1987	39.7%	Thu May 31 17:32:20 2012
fim			765	15.3%	Thu May 31 17:32:20 2012
fim_test			1185	23.7%	Thu May 31 17:32:20 2012
GSI			34	0.7%	Thu May 31 17:32:20 2012

Last Estimated: The date when the directory was last scanned.

# Utility: sdu

- sdu == Smart du
  - Running du is slow on large parallel filesystems.
  - A background daemon walks the scratch filesystems and reports disk usage.
  - Sdu provides similar information to the Unix du command, except it uses the disk usage reports generated above.
  - This tool provides more information than squotas, because you can look at any directory on the filesystem.
  - It takes about 3 days to walk the entire file system. We are working to improve this time without impacting the whole system.

Use sdu instead of du when starting to investigate disk usage for yourself and/or your project

# Options to sdu

```
# sdu -h
```

```
sdu - System du - Reports du like information for  
project filesystems based on system data
```

Usage:

```
sdu (options) directory1 directory2 .....
```

Options:

```
-s : Summary Report  
-h : Report numbers in human readable form  
-b : Report numbers in bytes  
-k : Report numbers in kilobytes  
-m : Report numbers in megabytes  
-g : Report numbers in gigabytes  
-t : Report numbers in terabytes  
-q : Quiet, don't report errors
```

# sdu report

```
[root@fel quotas]# sdu -s /scratch1/portfolios/BMC/nescmgt/djand
88707407140166 /scratch1/.../djand          Last Update:Mon Jun  4 09:46:27 2012
[root@fel quotas]# sdu -sh /scratch1/portfolios/BMC/nescmgt/djand
80T          /scratch1/.../djand          Last Update:Mon Jun  4 09:46:27 2012
[root@fel quotas]# sdu -sh /scratch1/portfolios/BMC/nescmgt/djand/*
Skipping file, /scratch1/.../djand/foo.  Only sizes of directories are supported.
184G        /scratch1/.../djand/home      Last Update:Mon Jun  4 09:46:27 2012
80T         /scratch1/.../NOAA_CSC_ESA          Last Update:Mon Jun  4 09:46:27 2012
[root@fel quotas]# sdu -sh /scratch1/portfolios/BMC/nescmgt/djand/NOAA_CSC_ESA
80T         /scratch1/.../djand/NOAA_CSC_ESA    Last Update:Mon Jun  4 09:46:27 2012
[root@fel quotas]# sdu -sh /scratch1/portfolios/BMC/nescmgt/djand/NOAA_CSC_ESA/*
8T          /scratch1/.../NOAA_CSC_ESA/CM-HR   Last Update:Mon Jun  4 09:46:27 2012
7T          /scratch1/.../NOAA_CSC_ESA/FIM     Last Update:Mon Jun  4 09:46:27 2012
1T          /scratch1/.../NOAA_CSC_ESA/GFS     Last Update:Mon Jun  4 09:46:27 2012
76G        /scratch1/.../NOAA_CSC_ESA/gfsenkf  Last Update:Mon Jun  4 09:46:27 2012
[root@fel quotas]# sdu -sh /scratch1/portfolios/BMC/nescmgt/djand/NOAA_CSC_ESA/FIM/*
726M       /scratch1/.../djand/NOAA_CSC_ESA/FIM/fimdata.c4  Last Update:Mon Jun  4 09:46:27 2012
726M       /scratch1/.../djand/NOAA_CSC_ESA/FIM/fimdata.c4  Last Update:Mon Jun  4 09:46:27 2012
3T         /scratch1/.../djand/NOAA_CSC_ESA/FIM/FIMrun  Last Update:Mon Jun  4 09:46:27 2012
25G        /scratch1/.../djand/NOAA_CSC_ESA/FIM/prep   Last Update:Mon Jun  4 09:46:27 2012
```

You don't have to use sdu for all du like reports. When you go down low enough in your directories, it is ok to run du. We are just trying to level out metadata load.



Any Questions?