# Introduction

This hands-on tutorial has two parts. In the first part, different cost scenarios will be tested by copying files to dCache. In the second part, (hot spot) replication will be tested by copying files out of dCache.

There are five pools of  $\sim 9.87 \text{ GB}^1$  each for this hands-on session, with all but one containing files:

- sfswn050\_small\_1: 17 files<sup>2</sup> of 0.5 GB each,
- sfswn050\_small\_2: 8 files of 1 GB each,
- sfswn050\_small\_3: 24 files of 0.25 GB each,
- sfswn050\_small\_4: 12 files of 0.25 GB each,
- sfswn050\_small\_5: empty.

These pools are mounted in the directories /smallpool\_1 to /smallpool\_5. They all have the same queues and transfer modes enabled.

All pools for this hands-on tutorial are in the pool group store-pools, which is in the link group store-link. This link is assigned to the section CourseSection.

The big pools  $sfswn050_big$  and  $sfswn054_big$  are not needed for this tutorial. So please log on to the admin interface<sup>3</sup> and disable both.

# 1 Costs

## 1.1 Preliminaries

Create a 50 MB file, e.g. via dd if=/dev/urandom of=/tmp/50MB bs=50MB count=1. In addition, create a 500 MB file.

## 1.2 Random distribution

Go to the PoolManager cell in the admin interface and change both cost factors for the CourseSection to zero<sup>4</sup>. Now copy your 50 MB file via dccp into dCache for  $\sim 20$  times (please choose file names that you can remember). Now check if these files have been randomly distributed (they are the newest files in the /smallpool\_X/data directories).

## 1.3 Distribution based on free space

Delete your files from the previous exercise. Then set the space cost factor to 1.0, while leaving the performance cost factor at 0.0. Copy our 50 MB file via dccp into dCache for  $\sim$ 15 times and check their distribution afterwards. One expects all files to end up in sfswn050\_small\_5.

 $<sup>^{1}</sup>$ In this document, 500 MB equals  $10^{9}$  bytes.

<sup>&</sup>lt;sup>2</sup>They are called /pnfs/physik.uni-wuppertal.de/workshop/datasets/smallpools/pool\_1\_500MB\_1 to /pnfs/physik.uni-wuppertal.de/workshop/datasets/smallpools/pool\_1\_500MB\_17.

<sup>&</sup>lt;sup>3</sup>ssh -c blowfish -p 22223 admin@localhost -i /root/.ssh/dcache\_key for password free access <sup>4</sup>Use pm set in PoolManager cell.

### 1.4 Distribution based on transfer load

Please delete your files from the previous exercise. Then set the performance cost factor to 1.0 and space cost factor at 0.0. Once more copy your 50 MB file via dccp into dCache for  $\sim 20$  times, but have these transfers run in parallel (e.g. use ampersand on the shell). Then check their distribution afterwards. One expects files to be more or less uniformly distributed on the 5 small pools.

Afterwards, delete the files from this exercise.

#### 1.5 Using the GUI

Check which transfer modes are allowed and which client queues can be used by the client (i.e.  $maxAllowed \neq 0$ ), e.g. by using the pcells GUI (bee below) or by using cm ls in the PoolManager cell. Which performance cost do you expect for exactly one client transfer via dcap?

Please install the pcells GUI on your notebook if it not has been set up already. The GUI can be downloaded at http://www.dcache.org/downloads/gui/. Please connect to your machine as admin on port 22223 (protocol: ssh1). Now choose the PoolCosts tab, update the tree and switch on automatic updates of the costs every other second. Change the performance cost to 1.0 and transfer your 500 MB to dCache. Which cost do you see in the GUI? Transfer the file once more to dCache and check which cost is returned by cm ls -d in the PoolManager cell.

Please delete your files again.

# 2 Hot-spot replication

### 2.1 p2p parameters in PoolManager

Please ensure that both p2p-allowed and p2p-oncost are set to off for the CourseSection. Then disable dcap access to pool sfswn050\_small\_3 and try to transfer a file from this pool out of dCache.

Now enable p2p-allowed and transfer one file from sfswn050\_small\_3.

#### 2.2 On cost replication using fixed threshold scheme

Please enable p2p-oncost in CourseSection and look up the p2p value in the CourseSection.

Copy all files from pool sfswn050\_small\_1 out of dCache in parallel. Use the GUI to monitor costs and where to files are being replicated.

### 2.3 On cost replication using percentile scheme

Set p2p-oncost in CourseSection to 80% and sweep replicas of files from all pools.

Now start several file transfers of files from pool 2; check if dCache replicates any files. Then start several file transfers of files from pool 1; check again for replication activities.