DPM scalability test plan

Authors: Lana Abadie, Andreas Peters

v 1.2 Feb 14, 2008

Objectives

Evaluate DPM scalability and performance in terms of number of files and clients and file sizes.

Verify the space management functionality (quota).

Evaluate DPM as a Tier-2 storage system suitable for analysis running experiment analysis code.

DPM Configuration

HW Setup – cluster fileserver/dpm-test

- 1 DPM head node : 16 GB of memory + 8 cores
- 27 disk servers (8 cores / 8 GB) – total capaciy 100 TB :
    - 8 x 600 GB XFS
    - 19 x 5 TB XFS
    - Linux version 2.6.9-55.0.12.EL.cernsmp

All machines are 64bits Intel Xeon (2 Ghz core) and single Intel Pro 1000 NIC.

Configuration done via Yaim

Use Yaim for DPM installation and configuration using the latest version from internal repository version 1.6.10. Start with a default setup for DPM.

HEAD-Node Configuration

Depending on the first results of the tests default thread parameters on the head-node will have to be tuned.

The default number of threads for the DPNS is 20: we may need to increase it. The maximum number of threads for the DPM is 100. By default there are 20 slow and 20 fast threads.

We will apply the default parameters for rfio : but we may change some values via rfio_setopt (RFIO IOBUFSIZE, RFIO_READOPT, version of rfio v2 or v3).

Advices from Jean-Philippe Baud may be needed to tune the DPM.

We don’t plan to change the filesystem options.

Type of tests

Test 1 : Performance evaluation DPM

We will start with 640 parallel clients and try to fill the full storage capacity of 100 TB. All tests will be done using rfio protocol.

Storing files

Each client will store concurrently a 1MB file via rfcp. File names will be generated with `uuidgen` and stored in a directory structure with max. 1000 files per directory in the DPM namespace.

Reading files

Each client will read concurrently with rfcp random selected files.

Mixture of requests

Half of the clients will store files, the other half will read files (later additionally add clients deleting files randomly)

Test 2: Space management design

We will test the space token implementation.

We will use different user identities for verification. Each identity will reserve private space. Each identity will write, read and delete files from the DPM. The certification team in GD has been kind to let us use their test user certificates so we can have up to 200 different identities.

Monitoring

Time measurement

We will measure time distributions (measured on clients) needed to store/read/delete files. We will measure typical node parameters on the head node and do profiling with oprofile of the DPM/DPNS & mysql processes.

Test scripts will insert the timestamp, hostname, execution time, error code, error messages of each request in a separate MySQL DB.

Lemon

Lemon will be used to measure the network and cpu load, memory usages

Results

Graphics will be produced

Results will be reported here: https://twiki.cern.ch/twiki/bin/view/LCG/Scalability_DPM