

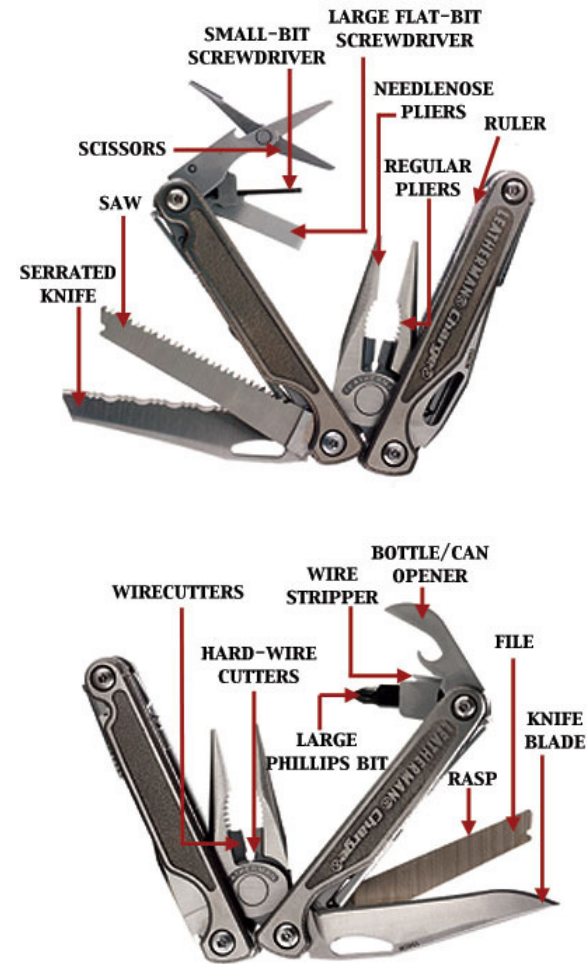
# **OSG Storage Operations Toolkit**

**Production-Driven, Community-Contributed**

**Abhishek Singh Rana**

**OSG Storage Forum, Fermilab, July 1 2009**

# Concept



Leatherman Charge Ti, with 17 tools and lightweight titanium handles, \$100.

Photo Courtesy: Charles Schiller, Copyright: Time Inc., Source: This Old House magazine

Product: <http://www.leatherman.com/multi-tools/full-size-tools/charge-tti.aspx>

# Concept



## OSG Storage Operations Toolkit Contents and Functionality

| Utility Name |                         |   |  | Data Removal Capability | Load Exhaustive Run | Tool Abrasive Level | Automatic Cron |
|--------------|-------------------------|---|--|-------------------------|---------------------|---------------------|----------------|
| RPM 1        | SRM PNFS Reclaimer      | osg_dc_srm_space_reclaimer                        |  | Y                       | Y                   | Y                   | Y              |
| RPM 2        | Core Operations         | osg_dc_analyze_disk_usage_of_all_pools            |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_analyze_disk_usage_of_pnfs_dir             |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_analyze_num_replicas_disk_pnfs_of_pnfs_dir |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_cancel_stuck_restores                      |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_cleanup_broken_transfers                   |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_cleanup_disk_files_not_in_pnfs             |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_compare_checksum_disk_pnfs_for_all_files   |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_create_graph_of_poolmanagerconf            |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_disk_usage_du                              |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_find_files_in_pnfs_not_on_disk             |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_find_files_on_a_pool                       |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_find_num_replicas_for_all_files            |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_find_path_for_pnfsid                       |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_find_pnfsid_for_path                       |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_find_pools_with_file                       |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_get_all_poolgroups_system_summary          |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_get_html_dump_of_fs_usage_of_pnfs_dir      |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_move_files_from_full_pool_to_another_pool  |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_query_all_transfer_rates                   |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_query_srm_transfer_rate                    |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_verify_discrepancy_disk_usage_of_all_pools |  | Y                       | Y                   | Y                   | Y              |
| RPM 3        | Physical File Manager   | osg_dc_PFM  |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_retire_a_pool                              |  | Y                       | Y                   | Y                   | Y              |
| RPM 4        | dCache Chronicle        | osg_dc_chronicle                                  |  | Y                       | Y                   | Y                   | Y              |
| RPM 5        | Development API: Python | Development API for dCache: Python                |  | Y                       | Y                   | Y                   | Y              |
| RPM 6        | Pool Local Operations   | osg_dc_automatic_restart_offline_pool             |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_pool_usage_nagios_plugin                   |  | Y                       | Y                   | Y                   | Y              |
|              |                         | osg_dc_check_certificates_CRLs                    |  | Y                       | Y                   | Y                   | Y              |

Naming convention - osg\_dc: dCache, osg\_bmc: BeStMan, osg\_xrd: xRootD, osg\_hdp: Hadoop

**Protective Measure:** A site can choose to set its **Toolkit Site Cautionary Level** to: **low**, **medium**, or **high**. Abrasive tools may not run until the Site level is set to **low**. Out of the box level is **medium**.

OSG Storage Operations Toolkit, with 25+ tools and lightweight titanium handles, \$0.

# Principles

- Driven by production-scale operations.
- Contributed by Community.
- Adopted by larger Community.
- Utility first, design is second.
- Portability of deployment across Sites.
- Not out-of-box, not in the product.

# OSG Storage Operations Toolkit

## Contents and Functionality

### Contents in the Box

| Utility Name |                         |   |   |   |   |
|--------------|-------------------------|---|---|---|---|
| RPM 1        | SRM PNFS Reclaimer      | osg_dc_srm_space_reclaimer                        | Y | Y | Y |
| RPM 2        | Core Operations         | osg_dc_analyze_disk_usage_of_all_pools            | Y |   |   |
|              |                         | osg_dc_analyze_disk_usage_of_pnfs_dir             |   |   |   |
|              |                         | osg_dc_analyze_num_replicas_disk_pnfs_of_pnfs_dir | Y |   |   |
|              |                         | osg_dc_cancel_stuck_restores                      |   |   |   |
|              |                         | osg_dc_cleanup_broken_transfers                   |   |   |   |
|              |                         | osg_dc_cleanup_disk_files_not_in_pnfs             | Y | Y |   |
|              |                         | osg_dc_compare_checksum_disk_pnfs_for_all_files   | Y |   |   |
|              |                         | osg_dc_create_graph_of_poolmanagerconf            |   |   |   |
|              |                         | osg_dc_disk_usage_du                              |   |   |   |
|              |                         | osg_dc_find_files_in_pnfs_not_on_disk             | Y |   |   |
|              |                         | osg_dc_find_files_on_a_pool                       |   |   |   |
|              |                         | osg_dc_find_num_replicas_for_all_files            | Y |   |   |
|              |                         | osg_dc_find_path_for_pnfsid                       |   |   |   |
|              |                         | osg_dc_find_pnfsid_for_path                       |   |   |   |
|              |                         | osg_dc_find_pools_with_file                       |   |   |   |
|              |                         | osg_dc_get_all_poolgroups_system_summary          |   |   |   |
|              |                         | osg_dc_get_html_dump_of_fs_usage_of_pnfs_dir      | Y |   |   |
|              |                         | osg_dc_move_files_from_full_pool_to_another_pool  | Y | Y |   |
|              |                         | osg_dc_query_all_transfer_rates                   |   |   |   |
|              |                         | osg_dc_query_srm_transfer_rate                    |   |   |   |
|              |                         | osg_dc_verify_discrepancy_disk_usage_of_all_pools |   |   |   |
| RPM 3        | Physical File Manager   | osg_dc_PFM  | Y | Y |   |
|              |                         | osg_dc_retire_a_pool                              | Y | Y |   |
| RPM 4        | dCache Chronicle        | osg_dc_chronicle                                  | Y |   | Y |
| RPM 5        | Development API: Python | Development API for dCache: Python                |   |   |   |
| RPM 6        | Pool Local Operations   | osg_dc_automatic_restart_offline_pool             |   |   | Y |
|              |                         | osg_dc_pool_usage_nagios_plugin                   |   |   | Y |
|              |                         | osg_dc_check_certificates_CRLs                    |   |   | Y |

Data Removal Capability  
Load Exhaustive Run  
Tool Abrasive Level  
Automatic Cron

# Examples (osg\_dc = dCache)

|  |  |
|--|--|
| <code>osg_dc_analyze_disk_usage_of_all_pools</code>            | <code>osg_dc_find_num_replicas_for_all_files</code>            |
| <code>osg_dc_analyze_disk_usage_of_pnfs_dir</code>             | <code>osg_dc_find_path_for_pnfsid</code>                       |
| <code>osg_dc_analyze_num_replicas_disk_pnfs_of_pnfs_dir</code> | <code>osg_dc_find_pnfsid_for_path</code>                       |
| <code>osg_dc_cancel_stuck_restores</code>                      | <code>osg_dc_find_pools_with_file</code>                       |
| <code>osg_dc_cleanup_broken_transfers</code>                   | <code>osg_dc_get_html_dump_of_fs_usage_of_pnfs_dir</code>      |
| <code>osg_dc_cleanup_disk_files_not_in_pnfs</code>             | <code>osg_dc_move_files_from_full_pool_to_another_pool</code>  |
| <code>osg_dc_compare_checksum_disk_pnfs_for_all_files</code>   | <code>osg_dc_query_all_transfer_rates</code>                   |
| <code>osg_dc_du</code>   | <code>osg_dc_query_srm_transfer_rate</code>                    |
| <code>osg_dc_find_files_in_pnfs_not_on_disk</code>             | <code>osg_dc_verify_discrepancy_disk_usage_of_all_pools</code> |
| <code>osg_dc_PFM</code>  | <code>osg_dc_automatic_restart_offline_pool</code>             |
| <code>osg_dc_retire_a_pool</code>                              | <code>osg_dc_check_certificates_CRLs</code>                    |
| <code>osg_dc_pool_usage_nagios_plugin</code>                   | Development API for dCache: Python                             |
| <code>osg_dc_chronicle</code>                                  |  |

## The dCache Chronicle<sup>©</sup>

THE DCACHE CHRONICLE

| Thursday, May 14, 2009

| Evening Ed.

Executive Summary: Total disk capacity is 272.0 TB, of which 208.0 TB is in precious usage, 7.0 TB is cached. Disk utilization today is ~79%. Almost 55.0 TB is free. Storage system has 2 PoolGroups. Total pools found now are 90. Analysis of user-storage directory shows 46 users. There are 7 users consuming over a TB. 20 users have no consumption. Filecount is 107821 combined, and worth 23.0 TB in the filesystem. Overall, the users have 353 lost files. Full listing is in next mail.

# Adoption

- Install/configuration is easy and quick.
- In use by many ATLAS and CMS dCache Sites in the US (and Brazil).
- Surprisingly popular among EU Sites. In use across Sites in Germany, UK, Italy, among others.
- Plan to start joint effort with DESY dCache team to facilitate standardization.

# Current Availability - Downloads

## Operations**Toolkit**

Open Science Grid **Storage**

Download



Read



Contribute  
your Tools



### OSG Storage Operations Toolkit

### All RPMs

v2.4\_p2



|                |                              |                         |      |
|----------------|------------------------------|-------------------------|------|
| RPM 1          | SRM PNFS Space Reclaimer     | dCache PNFS Server node | v1.1 |
| RPM 2          | Core dCache Operations       | Any dCache Server node  | v2.5 |
| RPM 3          | Physical File Manager / PFM  | Any dCache Server node  | v2.1 |
| RPM 4          | dCache Chronicle             | Any dCache Server node  | v1.4 |
| RPM 5          | Development API: Python      | Any dCache node         | v2.1 |
| RPM 6          | dCache Pool Local Operations | dCache Pool nodes       | v1.1 |
| <b>Future:</b> |                              |                         |      |
| RPM 7          | Hadoop Chronicle             | Any Hadoop node         | v1.0 |
| RPM 8          | Development API: Perl        | Any dCache node         | v1.0 |

# Current Availability - Documents

## OperationsToolkit

Open Science Grid Storage



### General: Auto-notification of updates, ToolControl.conf, ExaByte Grids

Toolkit now comes with a minor monthly auto-notification feature. It compares the locally installed version number with the currently available more recent version, if any, and notifies the site to do timely upgrades. Note that it is similar to the standard feature in the market, but an automatic download or install is not performed for obvious reasons.

Primary configuration resides in ToolControl.conf. A sample is [here](#).

Aim of this *software* is to simplify or automate or autoengineer operations, so that *humans* can spend more time building ... Exabyte Grids and Clouds.

### RPM 1: SRM PNFS Space Reclaimer

**Requirements:** Storage system - dCache. User - root. PNFS server node with DB access from SRM node.

This utility has its own policy file, and requires careful initial configuration. To enable operations of this utility, see detailed document at [OSG TWiki](#).

### RPM 2: Core dCache Operations

**Requirements:** Storage system - dCache. User - root. Any server node, with PNFS mounted. Preferably, with [Jython](#) installed. Jython is only required by 4 utilities. All other utilities will perform as normal, even without it.

Example, see this [note](#). To install Jython:

```
$ {jdk_dir}/bin/java -jar jython_installer-2.2.1.jar --console
```

Standard install of Jython does not provide all packages. Choose 'all/everything' option.

Deployment of Jython based minor subset, on a node without a local dCache install, requires a local copy of [cells.jar](#) and [dcache.jar](#) in `/opt/d-cache/classes`.

As user root, install the RPM.

```
$ rpm -Uvh osg-dcache-tools-core-operations-VER.i386.rpm
$ Edit /opt/d-cache/toolkit/ToolControl.conf
```

Note the [Cautionary\\_Level](#) tuning. Based on a site's own discretion and level of comfort, this can be set to [high](#), [medium](#), or [low](#). Toolkit comes with a default setting

### RPM 2: Core dCache Operations

**Requirements:** Storage system - dCache. User - root. Any server node, with PNFS mounted. Preferably, with [Jython](#) installed. Jython is only required by 4 utilities. All other utilities will perform as normal, even without it.

Example, see this [note](#). To install Jython:

```
$ {jdk_dir}/bin/java -jar jython_installer-2.2.1.jar --console
```

Standard install of Jython does not provide all packages. Choose 'all/everything' option.

Deployment of Jython based minor subset, on a node without a local dCache install, requires a local copy of [cells.jar](#) and [dcache.jar](#) in `/opt/d-cache/classes`.

As user root, install the RPM.

```
$ rpm -Uvh osg-dcache-tools-core-operations-VER.i386.rpm
$ Edit /opt/d-cache/toolkit/ToolControl.conf
```

Note the [Cautionary\\_Level](#) tuning. Based on a site's own discretion and level of comfort, this can be set to [high](#), [medium](#), or [low](#). Toolkit comes with a default setting of [medium](#). Only when a site decides to choose a setting of [low](#) cautionary level, will the tools with data removal capability become enabled.

```
$ source /opt/d-cache/toolkit/setup.sh
```

This will bring all `osg_dc*` in \$PATH. Type `osg_dc`, then try if TAB toggles and highlights all utilities.

Most site administrators prefer command-line quick reference. Documentation is available with each utility.

For each utility, please use `'-h'` or `'--help'`.

```
osg_dc analyze disk usage of all pools
osg_dc analyze disk usage of pnfs dir
osg_dc analyze num replicas disk pnfs of pnfs dir
osg_dc cancel stuck restores
osg_dc cleanup broken transfers
osg_dc cleanup disk files not in pnfs
osg_dc compare checksum disk pnfs for all files
osg_dc create graph of poolmanagerconf
osg_dc disk usage du
osg_dc find files in pnfs not on disk
osg_dc find files on a pool
osg_dc find num replicas for all files
osg_dc find path for pnfsid
osg_dc find pnfsid for path
osg_dc find pools with file
osg_dc get all poolgroups system summary
osg_dc get html dump of fs usage of pnfs dir
osg_dc move files from full pool to another pool
```

# Current Availability - Contribute

## Operations**Toolkit**

Open Science Grid **Storage**

Download



Read



Contribute  
your Tools



OSG Storage welcomes all creative solutions developed by members of the Site community. Please use this page to directly upload your contributions. Thank you for sharing your expertise.

|                            |   |
|----------------------------|---|
| Your full name (required): | <input type="text"/>  |
| File:                      | <input type="text"/> <input type="button" value="Browse..."/> |
|                            | <input type="button" value="Upload"/>                         |

Valid file formats are .tar, .gz, .zip, .tgz

# Forward Plan

- Aggregation and inclusion of key **Hadoop (osg\_hdp)** utilities. E.g. Hadoop Chronicle.
- Periodic addition of new features, new utilities, bug fixes, new APIs.
- **BestMan** and **xRootD** utilities are welcome from the OSG Community.
  - Current need is not known, or not existent, possibly due to efficient and quick direct interfaces between LBNL team and Sites.

# Contributing Institutions, June 2009

- Brookhaven National Laboratory
- California Institute of Technology
- Cornell University
- Fermi National Accelerator Laboratory
- Istituto Nazionale di Fisica Nucleare, Italy
- Massachusetts Institute of Technology
- Purdue University
- University of California-San Diego
- University of Edinburgh, Scotland
- University of Nebraska-Lincoln
- University of Wisconsin-Madison

# More Information

- Web:

<http://datagrid.ucsd.edu/toolkit/>

- Email:

[osg-storage@opensciencegrid.org](mailto:osg-storage@opensciencegrid.org)

- Contact:

- Brian Bockelman: [bbockelm@math.unl.edu](mailto:bbockelm@math.unl.edu)
- Frank Würthwein: [fkw@ucsd.edu](mailto:fkw@ucsd.edu)
- Abhishek Singh Rana: [rana@fnal.gov](mailto:rana@fnal.gov)