

Billing and Pool Queue Plots

Albert L. Rossi

Fermi National Accelerator
Laboratory

Billing /Accounting

dCache Book, Chapter 15.

- dCache has built-in monitoring capabilities which provide an overview of the activity and performance of the installation's doors and pools. There are two options for how this data can be represented and stored:
 - a set of log files written to a known location
 - a database (the billing database).
- These options can be enabled simultaneously. If the database option is selected, the data in those tables will also be displayed as a set of histogram plots on the installation's web page.

Billing Logs

Plain text log files, written to known location (*/var/lib/dcache/billing/yyyy/mm*), named:

- *billing-<yyyy.MM.dd>.log*
- *billing-error-<yyyy.MM.dd>.log*

Four kinds of messages stored:

- mover message (*MoverInfoMessage*)
- remove message (*RemoveFileInfoMessage*)
- door message (*DoorRequestInfoMessage*)
- storage message (*StorageInfoMessage*)

A very detailed description of the formatting of the entries (and how to customize them) can be found in:

/usr/share/dcache/defaults/billing.properties

Note that with 2.6, there is a new boolean attribute on the mover message, 'p2p', denoting a pool-to-pool transfer. This is available for use (and present in the database) but has not been added to the default format for backwards compatibility.

Billing Logs

The summary table based on these logs is still available as a web page at

<http://<httpdhost>:2288/billing>

but has not been incorporated into the new webadmin pages.

dCache Billing

Total Request Overview

Action	Total Request Count	Request Failed
request:door	28	4
store:pool	13645	13645
transfer:pool	30	0
remove:pool	84	0

Pool Statistics

Pool	Mover Transfers	Restores from HSM	Stores to HSM	Total Errors
w-dmsdca02-1	0	0	0	9548
Total	0	0	0	9548

Billing Database

Has been around since 1.9.12.

The schema has been extended, indices and triggers added, and one table (*costinfo*) eliminated.

It now consists of four fine-grained tables (*billinginfo*, *storageinfo*, *doorinfo* and *hitinfo*), along with a corresponding set of aggregate tables (hourly to daily), and a set of views over the fine-grained data (for the sake of guaranteeing reads with results always of a fixed size).

For configuration and customization, see chapter 15 of the *dCache Book*.

A note concerning the hitinfo tables: these record disk cache hits and misses; however, they will remain unpopulated unless

`poolmanager.cache-hit-messages.enabled=true`

Since there is additional cost to sending these messages, it is not enabled by default.

Billing Plots

A fixed set of plots which display

- (Giga)bytes read and written for disk and backend
- Pool-to-pool (giga)bytes written
- Number of read and write transfers for disk and backend
- Number of pool-to-pool transfers
- Connection time (max, min, average)
- Cache hits and misses

The data is grouped into four columns showing 24-hour (by hour), 7-day, 30-day and 365-day (by day).

To generate plots, you need to set

`billingToDb=yes`

`generatePlots=true`

The plots are refreshed at an interval of every five minutes; this can be changed via

`billing.plot.refresh-threshold` (in minutes)

There is a limited set of options for the type of data representation:

`billing.plot.default-style` (**CONNECTED**, **OUTLINE**, **FILLED**)

`billing.plot.default-scale` (**lin**, **log**)

NOTE: The billing infrastructure has been rewritten since 2.2 to enable running the billing cell/service on a different host from the httpd service.



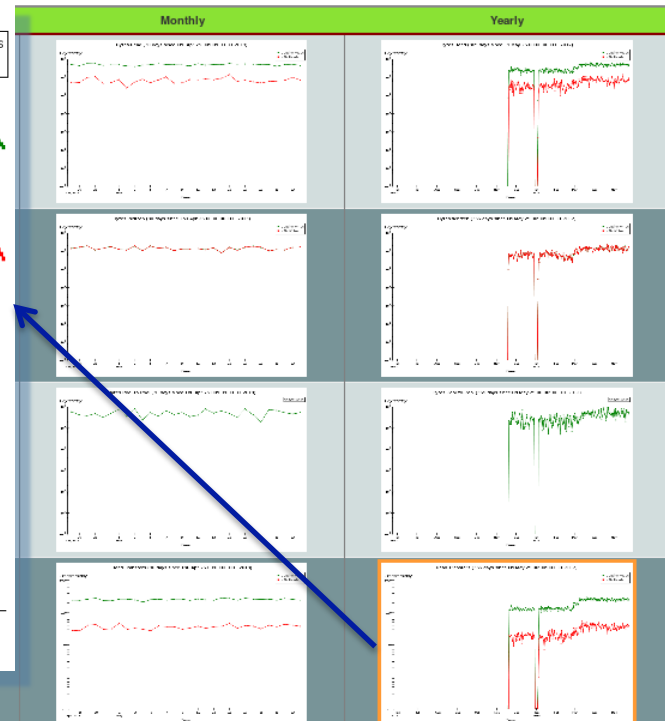
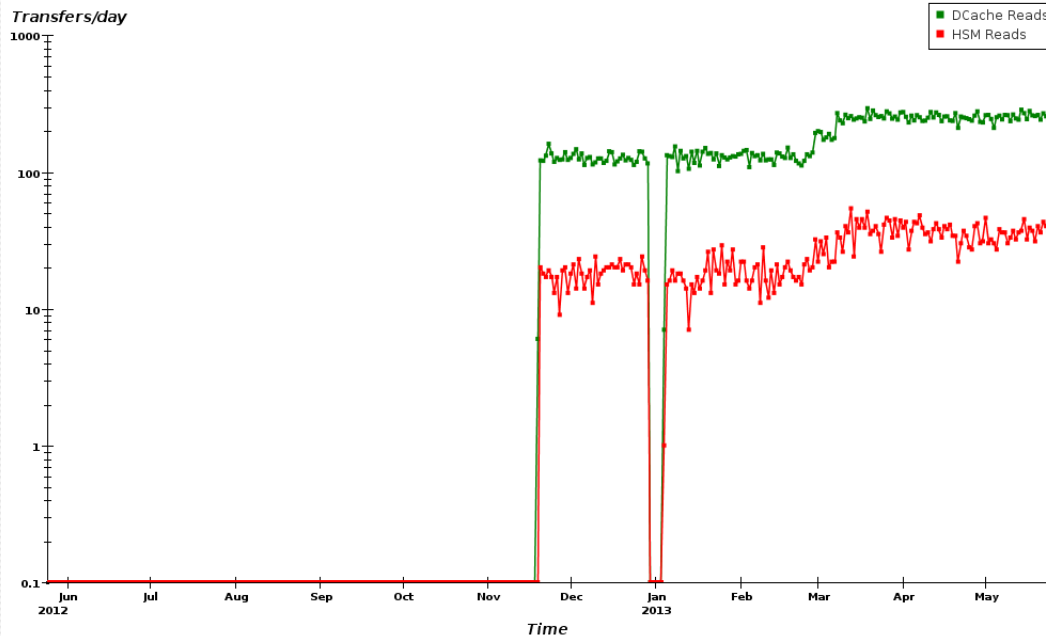
- HOME
- CELL SERVICES
- POOL USAGE
- POOL QUEUES
- POOL QUEUE PLOTS
- POOLGROUPS
- TAPE TRANSFER QUEUE
- ACTIVE TRANSFERS
- BILLING PLOTS**
- POOL SELECTION SETUP
- POOL ADMIN
- CELL ADMIN

Last updated Fri May 24 07:25:49 CDT 2013

CONNECTED, log

BILLING HISTORY PLOTS

Read Transfers (365 days since Fri May 25 00:00:00 CDT 2012)



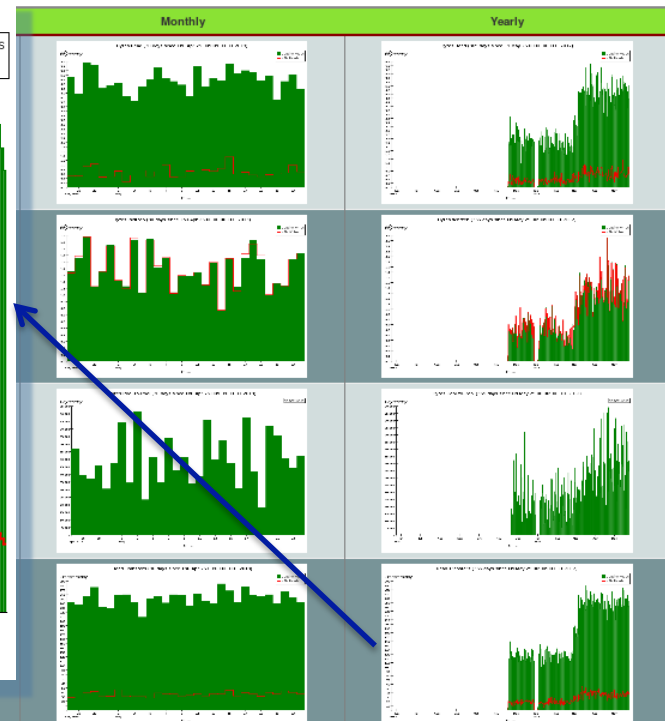



Last updated Fri May 24 07:18:48 CDT 2013

OUTLINE, lin

BILLING HISTORY PLOTS

Read Transfers (365 days since Fri May 25 00:00:00 CDT 2012)



Pool Queue Plots

New feature with 2.6

Based on the pool queue statistics table:

login

HOME CELL SERVICES POOL USAGE **POOL QUEUES** POOL QUEUE PLOTS POOLGROUPS TAPE TRANSFER QUEUE ACTIVE TRANSFERS BILLING PLOTS POOL SELECTION SETUP POOL ADMIN CELL ADMIN
SPACE TOKENS INFO XML ALARMS

POOL REQUEST QUEUES

Quick Find:

Pool	Domain	Movers			regular			P2P-Server			Stores			P2P-Client			Restores		
		Act	Max	Que	Act	Max	Que	Act	Max	Que	Act	Max	Que	Act	Max	Que	Act	Max	Que
pool3		<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
pool4		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
pool1	otfridDomain	0	100	0	0	100	0	0	10	0	0	0	62	0	10	0	0	0	0
pool2	otfridDomain	0	100	0	0	100	0	0	10	0	0	0	52	0	10	0	0	0	0
Total	Total	0	200	0	0	200	0	0	20	0	0	0	114	0	20	0	0	0	0

Data is stored in round-robin database such as used for network data. To generate plots, set:

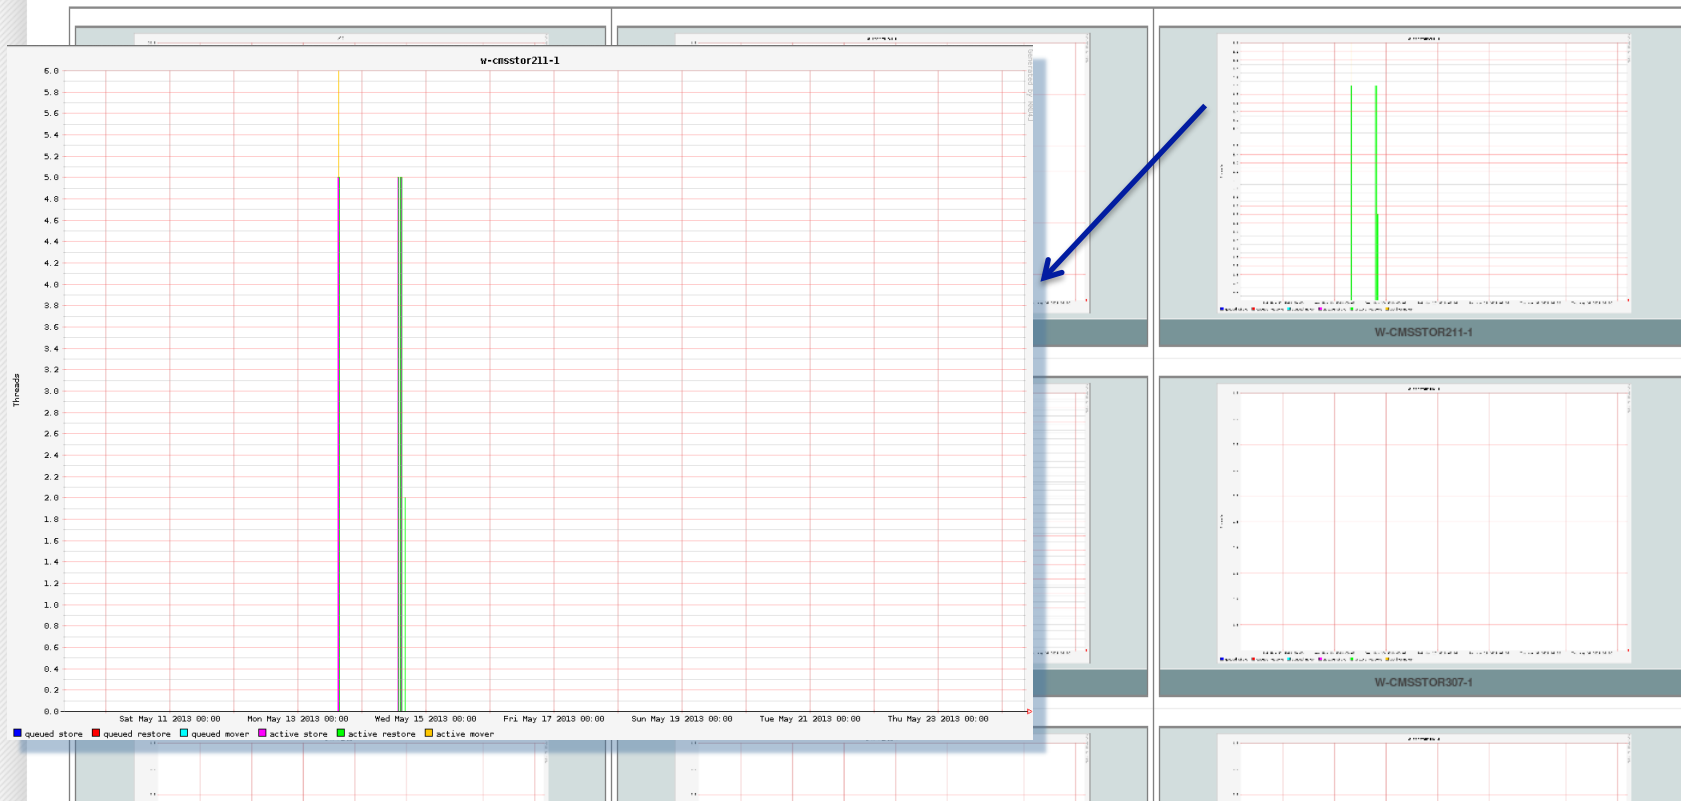
poolplots.enabled=true



FILTER

Sort Order up down Refresh
Match Expression Regular Expression

POOL QUEUE ACTIVITY PLOTS



Pool Queue Plots

Names are ordered
alphabetically, ascending or
descending.

Filter provided for convenient selection of subset
of pools (regular expression – here “CMSSTOR3”)

POOL QUEUE ACTIVITY PLOTS

Grid expands &
contracts dynamically
(max 4 to a row)

The screenshot displays the dCache web interface for Pool Queue Activity Plots. At the top, there is a navigation menu with options like 'SPACE TOKENS', 'INFO XML', and 'ALARMS'. Below this is a 'FILTER' section containing a 'Sort Order' dropdown, radio buttons for 'up' and 'down', a 'Refresh' button, and a 'Match Expression' field with the value 'CMSSTOR3'. A blue arrow points to the 'Match Expression' field, and another blue arrow points to the 'Filter provided for convenient selection of subset of pools (regular expression – here “CMSSTOR3”)'. Below the filter is the main content area titled 'POOL QUEUE ACTIVITY PLOTS', which contains a grid of four empty plots. Each plot has a title like 'W-CMSSTOR307-4'. A blue arrow points to the grid, and a text box says 'Grid expands & contracts dynamically (max 4 to a row)'. At the bottom of the page, there is a search bar with 'Find: SMTP' and navigation buttons like '< Previous', '> Next', 'Highlight all', and 'Match case'.

Questions or Comments Welcome