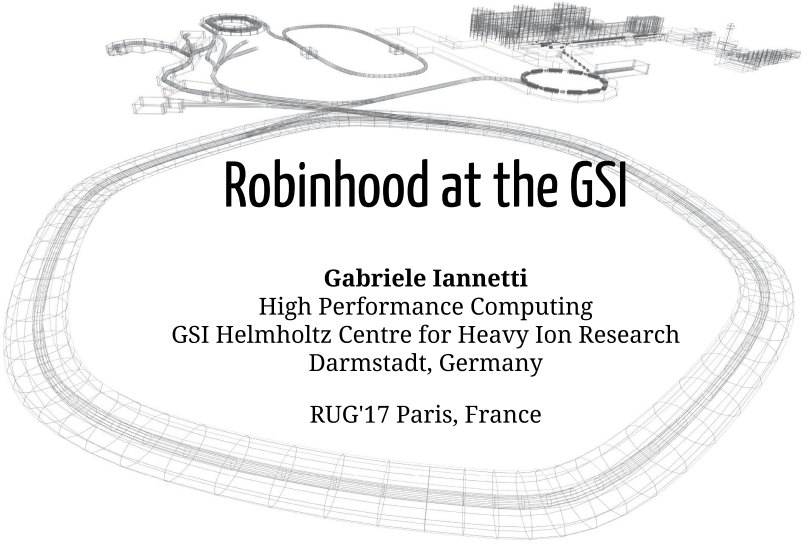

Robinhood at the GSI

Gabriele Iannetti
High Performance Computing
GSI Helmholtz Centre for Heavy Ion Research
Darmstadt, Germany

RUG'17 Paris, France

Agenda

1. Details about our Lustre Installation
2. Versions Used of Robinhood
3. Operation of Robinhood
4. Usage of Robinhood at the GSI
5. Examples for Custom Created Reports
6. Outlook

Details about our Lustre Installation

- Total storage capacity of 14.7 PB
- Pair of active/passive meta data server (2.5.3.90) with manual failover
- 78 file server (2.5.3.90) with ZFS (0.6.3)
- 546 OSTs - 7 OSTs per one OSS
- Running on Debian Wheezy

Versions Used of Robinhood

- Running productive since Robinhood version 3.0 with a Lustre client 2.8 (Lustre client fixed stability issues from minor versions)
- Currently running Robinhood in version 3.1
  (pre-release version already fixed the issue where just the changelogs processed from MDT-0 were saved into the Robinhood database)

Operation of Robinhood

- The Robinhood PE and database are running on two dedicated machines
- Initial full scan of Lustre including accounting information takes 55 hours for almost 465 million entries
- Increased changelog polling frequency to determine more precise file sizes on Lustre after file creation or modification
- Packaging is done for Debian

Usage of Robinhood at the GSI

Overview of Used Storage

- Viewing accounting information per users and groups
- Distribution of different file sizes
- Determination of saved files on OST's

Accounting and Reporting

Realized by custom scripts...

- Daily collection of accounting information on basis of users and groups
- Periodical check for large files >= 1TB
- Creation of weekly storage and large file reports

Examples for Custom Created Reports

- It is available as open source on GitHub at: https://github.com/GSI-HPC/lustre-scripts
- Implementation is done in Python

Example 1: A Large Files Report

- Configured as daily check for large files e.g. 1TB
- Determines the complete file path
- Custom user notifications with a large file overview are sent if enabled
- User information is retrieved from central directory service (LDAP)
- An overview mail is send weekly to a administrators mailing list
- Keeps track of already sent user notifications within a specified interval

Example 2: A Storage Usage Report (I)

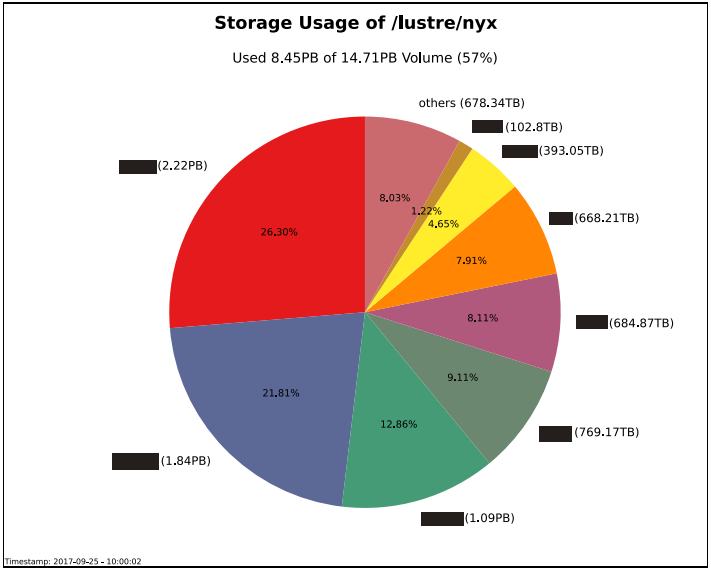

Example 2: A Storage Usage Report (II)

Consists of two program parts:

1. Daily collector that fills up an accounting history table
2. Reporting script that creates a storage usage chart per top x groups

Demonstrates the simplified structure of the ACCT_STAT_HISTORY table:

uid	gid	size	count	date
...	...	...	...	...

Outlook

- Show the trend of storage usage per user and groups based on the accounting information collected
- Archiving data from Lustre to HSM triggered by specified Robinhood policies

Thank you!