

Table of Contents

Storage accounting	1
Usage Record used.....	1
Architecture.....	1
Implementation.....	2
Data visualization.....	2
.....	4
.....	4
.....	4

Storage accounting

The following page contains the results of the Storage Accounting activity in progress at INFN-CNAF. This activity started at the beginning of March 2011 and is being carried out by Andrea Cristofori, Enrico Fattibene and Paolo Veronesi.

Usage Record used

The Usage Record (UR) utilized is generated in a way that allows a per file accounting. The fields that are part of the UR used are described on the paragraph "Architecture".

Architecture

The Accounting System utilized to collect, send, store and publish data is DGAS. Data are collected at a site level, sent and stored to the designed HLR for that particular site (site or multi site HLR).

sysDefStorageAccounting is the table that contains the data sent with the UR on the HLR. The first six columns describe the table attributes. The remaining columns contain the following fields: *GlueSchemaField* correspond to the Glue Schema parameter used; *Description* describe the meaning of the field; *AlwaysNULL* tells if the field, in the current implementation, is expected to be always NULL.

Field	Type	Null	Key	Default	Extra	GlueSchemaField	Description	AlwaysNULL
ID	bigint(20)		MUL	NULL	auto_increment			
RecordIdentity	char(64)		PRI				Hash of site, VO, SA, and TimeInstant values	
GlobalFileId	char(64)	YES		NULL			Unique file identifier (LFC file ID)	yes
LocalFileId	char(64)	YES		NULL			Local path of file	yes
GlobalGroup	char(64)	YES		NULL			Grid site name	no
GlobalUsername	char(64)	YES		NULL			Unique identifier of the user	yes
LocalUserId	char(64)	YES		NULL			Local user ID	yes
Charge	int(10)	YES		NULL			The charge associated with the storage utilization	yes
Status	char(64)						Status of a request to the storage system	yes
Host	char(64)					GlueChunkKey : GlueSEUniqueID	Storage Element host	no
SubmitHost	char(64)	YES		NULL			The file source host	yes

ProjectName	char(64)	YES		NULL		GlueSAAccessControlBaseRule	VO name	no
ProjectPartition	char(64)	YES		NULL		GlueSALocalID	Storage Area name	no
StorageType	char(64)	YES		NULL			StoRM, dCache, DPM, etc.	
ProtocolType	char(64)	YES		NULL			The protocol used in the event	yes
Network	int(10)	YES		NULL			The network utilization for the operation	yes
Disk	int(10)	YES		NULL		GlueSAUsedOnlineSize	Used space	no
TimeDuration	int(10)	YES		NULL			Time between 2 measures	no
TimeInstant	int(10)	YES		NULL			Time of the measure	no
ServiceLevel	char(64)	YES		NULL			Persistent, volatile, etc.	

Implementation

URs are, in the current implementation, generated by retrieving accounting data from the Information System. A dedicated script is run once a day from a UI and collects data for all the Italian sites. The URs are then sent to a test HLR.

Data visualization

HLRmon development server (hlrmon-dev.cnaf.infn.it) collects data from the configured HLR (at the moment the HLR used is the dgas-test-vm01.cnaf.infn.it test instance) typically once a day, aggregates them per day and stores the results in its internal DB. The table (*sedata_new*) that contains these data is described in the left half of the following table. The remaining columns correspond to: the corresponding field in the HLR table (if applicable) and the description of the field.

Table *sedata_new* schema

Field	Type	Null	Key	Default	Extra	HLRTableField	Description
ID	bigint(20)	NO	PRI	NULL	auto_increment		
YYMMGG	date	NO		0000-00-00		TimeInstant	Measurement day
Site	varchar(45)	YES		NULL		GlobalGroup	Grid site name
SEName	varchar(45)	YES		NULL		Host	Storage Element
VOName	varchar(45)	YES		NULL		ProjectName	VO
UsedSpace	bigint(20)	YES		NULL		Disk	Used space
Class	varchar(45)	YES		NULL		ProjectPartition	Storage Area name
last_mod_time	timestamp	NO		CURRENT_TIMESTAMP			

HLRmon does not store only information related to URs but it also stores, in another table, data related to the total and free space of each Storage Area (SA), the total space and the used space for each Storage Element

(SE) taken directly from the Information System. This table (*storage_info_system*) is described in the left half of the following table. The remaining columns correspond to: the corresponding field in the Glue Schema (if applicable) and the description of the field.

Table *storage_info_system* schema

Field	Type	Null	Key	Default	Extra	GlueSchemaField	Des
ID	bigint(20)	NO	PRI	NULL	auto_increment		
YYMMGG	date	NO		0000-00-00			Measu day
Site	varchar(45)	YES		NULL			Grid s
SEName	varchar(45)	YES		NULL		GlueSEUniqueID	Storag Elemen
SETotalSpace	varchar(45)	YES		NULL		GlueSETotalOnlineSize	SE tot
SEUsedSpace	bigint(20)	YES		NULL		GlueSEUsedOnlineSize	SE use
SATotalSpace	bigint(20)	YES		NULL		GlueSATotalOnlineSize	SA tot
SAFreeSpace	bigint(20)	YES		NULL		GlueSAFreeOnlineSize	SA fre
Class	varchar(45)	YES		NULL			Storag name
last_mod_time	timestamp	NO		CURRENT_TIMESTAMP			

HLRmon shows data aggregated per site, Storage Element and Storage Area. For each of these aggregation keys, and for the interval of time desired, a set of charts with the temporal trend is produced. Charts are available at:

<https://hlrmon-dev.cnaf.infn.it:8443/hlrmon/report/storage.php>

The same information can also be accessed in a tabular form, with the possibility to save it in xls format:

https://hlrmon-dev.cnaf.infn.it:8443/hlrmon/report/table_storage.php

A per-VO aggregation is not possible at the moment because a Storage Area can be shared among different VOs. In this case the Glue Schema allows the publication of the space used in the Storage Area, but not of the space used by each VO. For this case is not possible to know the space used by each VO.

-- Main.Enrico Fattibene - 2011-07-07

This topic: DGAS > StorageAccounting

Topic revision: r11 - 2011-07-15 - EnricoFattibene



Copyright © 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback