

KART: A Simple Client/server Interface to Access Different Tape Devices at CINECA Site

Sergio Bernardi
CINECA
s.bernardi@cineca.it

ABSTRACT: KART was developed to provide a controlled and easy access to different tape devices at CINECA site. Based on a client/server architecture, KART provides a set of commands in order to define logical volumes, read from volumes, write to volumes. The set of commands is also accessible via Web page (under development).

A logical volume is made up of physical volumes and provides the user with a simple, transparent way to copy files to and from tapes. In its original implementation the KART package stored the information related to the logical volumes in the home directory of the volume owner. The basic information is:

Introduction.

In 1989 CINECA installed a StorageTek 4400 cartridge silo. The system was originally connected to an IBM/VM machine only and the user community was able to use it through the VMTAPE interface. The station software and the fetch/discard commands were then very popular. In 1991 the StorageTek 4400 was directly attached to the C90. At that point in time there was not any specific software package at CINECA to take advantage of the tape subsystem facility. We decided to start the development of a set of scripts and programs to enhance the accessibility and to take advantage of the direct channel attachment to the Cray system. Unicos was the original development platform for the KART package.

- logical volume name
- list of physical volumes (VSNs)
- list of files stored on volume
- creation date
- logical volume size
- space left
- permissions
- ACL

The KART package was then redesigned to take into account new requirements: the possibility to access the STK 4400 from other systems not directly attached and to access other types of tape devices that might be useful to import and export data.

A server/client design was then chosen to implement the new requirements and to improve the overall portability.

The KART package.

The original design was restricted to a simple basic set of commands to reserve and use two different kinds of tape devices: manual mount 3480/3490 drives and STK 4480 units, especially for import/export operations.

KART allows to organize physical tapes into logical volumes and to keep track of their contents. In addition to that a rather simple Tape Data Base (TDB) was implemented to store and maintain the physical volume ownership information and for accounting purpose.

The TDB implementation is a Unix directory tree structure accessed by a set of maintenance and query commands.

The client/server KART implementation.

The client/server design added new components to the original implementation to reflect new ideas like the following:


- the server may be any system attached to a tape device of some kind
- the client is any system within the LAN
- the client has to be able to talk to many different systems in order to access different tape devices
- the set of data maintained by the KART package should

- be taken out of the user environment
- the TDB should be accessible by many client and kept out of any KART client or server system
- the KART command execution should be carried out by a unique specific user
- the client request should be submitted to the server in a 'secure' fashion

- the TDB (Tape Data Base)
- the KART centralized directory

The basic 'core' KART/TDB implementation was enhanced to take into account the new ideas. The following picture shows the package components and their relationships:

The information kept and managed by the TDB set of commands concerns the physical volume property, media type, and residency. Whenever the user requires the assignment of new physical volumes, the new volumes as well as their characteristics are stored accordingly to the following:



- owner username
- residency (either SILO, CART or VAULT)
- media type (3480, 3490, SD3A, SD3B, SD3C, DAT, EXABYTE, ...)
- compression flag
- uncompressed capacity
- assigned VSN (Volume Serial Number)
- permissions flags

The client/server components talk one another via socket. This is implemented by means of a daemon running on the server side that activates the requested service and the data transfer. That daemon is the **rce** (remote command execution daemon). It allows a user request for a KART service to be executed. The daemon listens on a socket and runs a specific command on behalf of a client. It reads a configuration file that specifies the list of commands that are allowed to be requested and executed.

That information is used by the KART commands to define the data associated to the new logical volume in terms of media type, capacity, residency, block size.

The client side issues a command request that specifies the KART operation to be performed using the **rce** (remote command execution) program. It passes on to the server the requested command plus all the necessary parameters. Once the permission to execute is granted, the **rce** fires the requested action and establishes the necessary data flow connection between the server and client sides.

There is a configuration file used by TDB itself. It contains all the information concerning media types, residency, options, servers, etc. in a 'key-value' pair format:

The KART/TDB internal information.

The package relies on two basic sets of information in order to perform its operations:

```
#
# key fields ..
#

# Default
deftype 3480/200Mb

#
# voltype      tipo          capacita'(Mb)  Compression
#
voltype 3480/200Mb      200           n
voltype 3480/400Mb      400           n
voltype 3490/800Mb      800           y
voltype 3490/1000Mb     1000          y
voltype 9490/800Mb      800           y
voltype 9490/1000Mb     1000          y
voltype SD3A/10000Mb    10000         y
voltype SD3B/25000Mb    25000         y
voltype SD3C/50000Mb    50000         y

#
# voldgn      tipo          res    dgn
voldgn 3480/200Mb SILO   SILO
voldgn 3480/200Mb CART   SILO
voldgn 3480/400Mb SILO   SILO36
```

```

voldgn 3480/200Mb CART SILO36
voldgn 3490/800Mb SILO SILO36
voldgn 3490/800Mb CART SILO36
voldgn 3490/1000Mb SILO SILO36
voldgn 3490/1000Mb CART SILO36
voldgn 9490/800Mb CART TIMB
voldgn 9490/800Mb SILO TIMB
voldgn 9490/1000Mb CART TIMB
voldgn 9490/1000Mb SILO TIMB
voldgn SD3A/10000Mb CART REDW
voldgn SD3A/10000Mb SILO REDW
voldgn SD3B/25000Mb CART REDW
voldgn SD3B/25000Mb SILO REDW
voldgn SD3C/50000Mb CART REDW
voldgn SD3C/50000Mb SILO REDW

```

```

#
dgn      TIMB      origin.cineca.it
dgn      REDW      origin.cineca.it
dgn      SILO      t3e.cineca.it
dgn      SILO36    t3e.cineca.it

#
#

res      CART
res      SILO
res      PARK

#
# alias hosts
#

hostalias sirio.cineca.it  sirio sirio.ypd
hostalias icarus2.cineca.it icarus2 sirio.ypd
hostalias origin.cineca.it  origin

```

The KART centralized information repository is a directory where each logical volume owner has a subdirectory. The latter contains in turn as many directories as the defined logical volumes and each of those has a single file named 'directory' which stores the following data:

- VSNs list
- block size
- list of written files including size, blocks occupied, date and time of copy

The KART information repository still has a limitation due to the present implementation of the logical volume locking mechanism. In order to be managed properly the repository has to be resident on the server attached to the specific device requested. This limitation is going to be removed as soon as a new distributed locking mechanism will be implemented.

The TDB package command set.

The TDB package provides two level of commands in order to manage the data base from the operator side and from the user side.

The user is not allowed to add, delete or modify a physical volume. Those operations are performed by operational scripts reserved to the machine room staff. Entering a new volume into the data base requires to specify basic information such:

- VSN
- media type (3490, SD3, ...)
- ownership (username)

The physical volume residency is decided by the user who can choose among three possible levels:

- SILO (STK library)
- CART (manual mount)
- PARK (vault)

and also request new residency assignement. Whenever the user requires to eliminate one or more physical volumes, the KART volume that contains them is archived and then removed.

The operations the user is allowed to perform on physical volumes are the following:

- physical volumes list
- physical volume ACL list
- physical volume ACL set/reset

Using ACL and permission flags allows the sharing of physical volumes among users.

The KART package command set.

The set of KART commands gives the following functionalities:

- logical volume creation
- logical volume content list
- logical volume removal
- logical volume access control list manipulation
- file write
- file read
- file removal
- optional 'memo' file attachment

To create a new logical volume the user needs to provide only the list of physical volumes' serial numbers and a name:

kart_new my_new_volume XC0010 XC0020 XC0030

The **kart_new** command creates a new directory in the centralized KART repository. The new logical volume is owned by the username that issued the creation request. There is no need for a username to be defined on a KART server.

The above mentioned information is displayed by the command **kart_dir**:

kart_dir my_new_volume

```
-----  
CARTUCCIA:      my_new_volume  
Volumi fisici:      1  
Tipo volumi   : 3480/200Mb  
Compressione : n  
Label[s]:      a00487  
Totale files:  14  
-----
```

| FILE | BYTES | DATA | ORA | COMMENTO |
|--------------|----------|----------|-------|----------|
| sergio_dir | 31109120 | 92/01/30 | 11:17 | |
| lisp_sun | 10818560 | 92/03/18 | 14:52 | |
| movietar | 19112448 | 92/06/01 | 14:42 | |
| cpiosardir07 | 22019584 | 92/08/11 | 09:21 | |
| cpiosar0892 | 21755392 | 92/09/02 | 12:14 | |
| save_pcnet | 51702784 | 92/10/07 | 14:44 | |
| rigel_sergio | 9539072 | 92/11/27 | 12:23 | |
| kcl_tarZ | 1233301 | 93/07/28 | 14:51 | |
| sbrolog_tarZ | 2519177 | 93/07/28 | 14:58 | |
| raycray | 369527 | 93/08/24 | 09:43 | |
| mbox | 729621 | 93/09/14 | 15:09 | |
| dbtree | 720896 | 93/11/21 | 11:36 | |

```
-----  
Mbytes totali:      200  
Mbytes occupati:   191  
Mbytes liberi:      9  
Proprietari   :      abnti01  
Flags         :      NRW  
Residenza     :      SILO  
ACL:  
abnti01              rw  
root                 rw  
-----
```

Dettaglio permessi volumi:

| Proprietari | pr. | rp | wp | Res. | Tipo | Volumi |
|-------------|-----|----|----|------|------------|--------|
| abnti01 | N | R | W | SILO | 3480/200Mb | A00487 |

The user can list all the logical volumes she/he owns by simply issuing:

kart_dir

| Cart | Vol. | Tipo | Files | Mb. | Occ. | Lib | Res. |
|--------|------|--------------|-------|--------|------|-------|------|
| da_t3e | 1 | SD3A/10000Mb | 9 | >10000 | 6833 | >3167 | SILO |

To copy a file to a logical volume the user needs the following command:

kart_put my_new_volume file_name < file_to_copy

or to move an entire directory structure:

tar cf - ./dir_to_copy | kart_put my_new_volume dir_name

To get a file stored in **my_new_volume**:

kart_get my_new_volume file_name > file_on_disk

Sharing KART logical volumes.

A useful feature of the KART package is the possibility to share a logical volume among several users. The package provides a specific command to attach an ACL (Access Control List) and define access permissions for specific users.

The logical volume owner assigns read and/or write permission to another user by means of:

kart_acl -a -u user_1 -p rw my_volume

The logical volume **my_volume** is now read/write accessible by **user_1**.

The ACL and permission flags go down to the physical volume level in the TDB, where access flags and ACL for a single volume may be previously assigned by the command:

```
tdbacl -a -u user_1 -p rw VSN_list
```

The current KART/TDB environment.

Today the KART and TDB packages are installed and operating on the following platforms:

- Unicos/mk
- IRIX 6.x
- Solaris 2.x
- AIX 4.x
- Linux

The servers are located on the systems attached to the newly installed STK 9310 that supports 3480/3490, Redwood and Timberline units via SCSI and BMX connections.

The server systems are:

- T3E (KART)
- Origin 2000 (KART)
- UltraSparc 5 (TDB)

The KART server port from the Unicos environment to IRIX was made in May-June 1998 using a beta version of the TMF subsystem. The programs and scripts port was rather straightforward. The major difference between the two operating system that required some minor change was the fact that IRIX is not 'block size' aware with respect of tape devices as seen by TMF.

Conclusion.

There is still some development activity concerning the KART/TDB. It is mainly related to provide a Web interface to it to take into account those users that access our site from a non Unix platform (MSWindows, NT, ...).