Integration of iRODS with IBM® Spectrum Archive™

A flexible tiered storage archiving solution

Nils Haustein – IBM European Storage Competence Center, Germany
Mauro Tridici – Euro-Mediterranean Center on Climate Change, Italy

Agenda

- Challenges with HSM managed file systems

  Solution with iRODS and IBM Spectrum Archive Enterprise Edition

© IBM Corporation 2020 and © CMCC Foundation 2020

HSM managed file systems

- File system resides on different storage tiers (Flash, SSD, Disk, Tape, Object Storage)
  - Placement rules define storage tier where files are stored
  - Migration rules define next storage tier and criteria
  - Placement an migration is fully automated by HSM
- User access via standard file system interfaces
  - Such as SMB, NFS, S3 and POSIX
- User can see and access all files transparently
  - Regardless if files are on disk or tape
  - Upon access files located on tape are copied to disk transparently

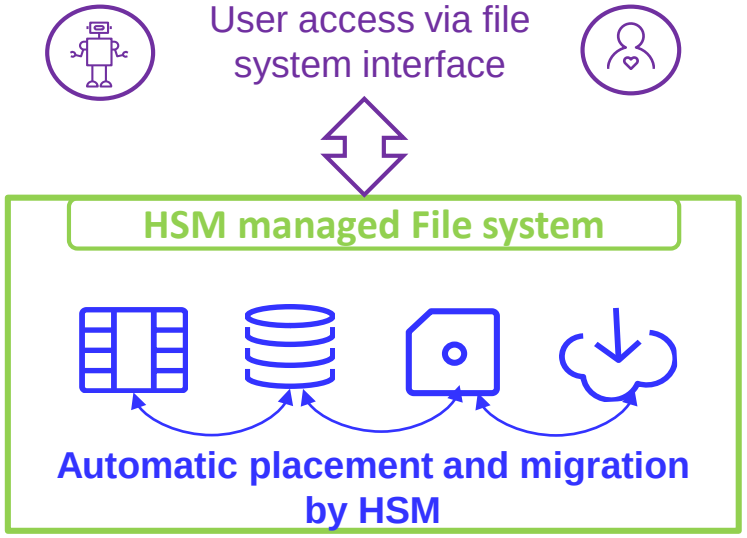

© IBM Corporation 2020 and © CMCC Foundation 2020

Challenges with tapes in HSM managed file systems

- Tapes as storage tier in HSM file systems are blessing and curse
  - Files on tape are visible in file system name space 
  - Access to files on tape cause transparent recall that takes some time 
- It gets worse if many files are accessed and recalled transparently
  - Files will be recalled from tape individually, without sorting
  - Causes many tape mounts and start-stop tape motion 
- Unfortunately the user does not easily see in the file system if the file is on tape
- User or group quota does not normally apply for the tape tier

Links: Blog Article

Smiley Photos by Unknown Author is licensed under CC BY

© IBM Corporation 2020 and © CMCC Foundation 2020

Conclusions

- Not all files are suited to be stored on tapes, for example
  - Files that are frequently accessed – cause many recalls
  - Files that have a short lifecycle – not worth to be migrated to tape
  - Many small files – make tapes performing slow

- Standard filesystems are tape agnostic
  - User has no simple way to see if a file is on disk or tape
  - Files are recalled from tape in the order of access and not in the storage order
  - After accessing files located on tape the user does not see the progress of the recall
  - There is no user and group quota limit for the data stored on tape only

© IBM Corporation 2020 and © CMCC Foundation 2020

Agenda

Challenges with HSM-managed file systems

- Solution with iRODS and IBM Spectrum Archive Enterprise Edition

© IBM Corporation 2020 and © CMCC Foundation 2020

iRODS

- The iRODS software is a data management layer
  - Sits above the storage and below domain-specific applications
  - Virtualizes data stored in heterogenous storage systems in single name space
  - Maintained by iRODS consortium as open-source software

- iRODS enables users to:
  - Access, manage, and share data across any type or number of storage systems through iRODS APIs (iCommands, REST, WebDAV, Python, C++, Java)
  - Automate workflows through powerful rules and microservices
  - Search and find data through descriptive metadata and query tools

Links: iRODS

© IBM Corporation 2020 and © CMCC Foundation 2020

Integration of iRODS with IBM Spectrum Archive

- iRODS can integrate with HSM managed file systems provided by IBM Spectrum Scale™ and IBM Spectrum Archive™
  - HSM managed file system provided by IBM Spectrum Scale is exported via NFS and defined as iRODS storage resource
  - Files stored in Spectrum Scale file systems are migrated to tape by Spectrum Archive based on policies
- iRODS provides powerful rules to intercept data operations and execute workflows
  - These workflows can be used to address the challenges of HSM managed file systems

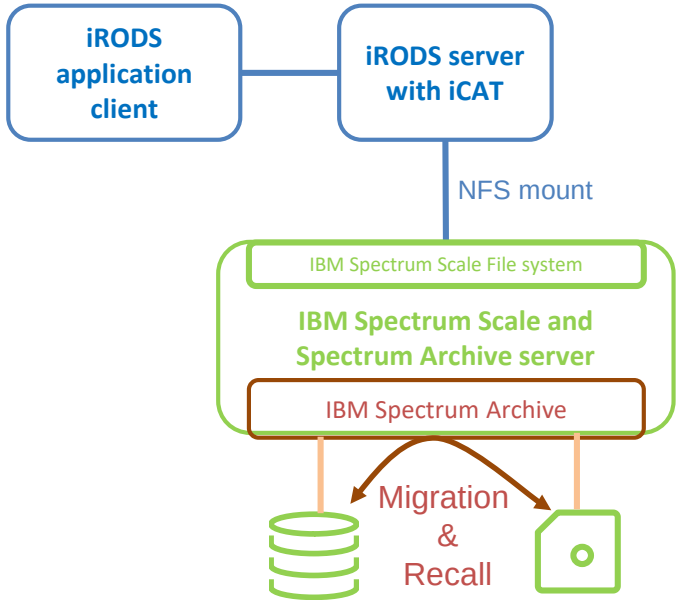

© IBM Corporation 2020 and © CMCC Foundation 2020

iRODS integration in action 1/2

- Recall files accessed through iRODS in a tape optimized manner
  - Use custom iRODS rule to intercept open request, check file state and if file is migrated then add file to queue and present a message to the user
  - Files in queue are recalled in a tape optimized way by a scheduled process on IBM Spectrum Archive

```
$ iget -f file1
file /archive/home/mia/col1/file1 is still on tape, and queued for recall.
```

- Display file migration state
  - New user command (`ifilestate`) that executes a custom iRODS rule to determine file state

```
$ ifilestate file1
Level 0: file /archive/home/mia/col1/file1 is MIGRATED

$ ifilestate -a
Level 0: file /archive/home/mia/col1/file0 is NOT migrated
Level 0: file /archive/home/mia/col1/file1 is MIGRATED
```

Links: Github Project | Blog Article

© IBM Corporation 2020 and © CMCC Foundation 2020

iRODS integration in action 2/2

- Set quota for the entire file system including the tape tier
  - Leverage integrated iRODS rules and automation

```
# Enable quota by editing the file /etc/core.re and adding the following line:
acRescQuotaPolicy {msiSetRescQuotaPolicy("on");

# Create and load delayed rule (/etc/irods/quota.r) that periodically checks quota
irule -F /etc/irods/quota.r -r irods_rule_engine_plugin-irods_rule_language-
instance

# set the quota limits for the user
iadmin suq user1 buffer 2147483648
iadmin suq user1 total 2147483648
```

- Automatically add file header information to the metadata catalog
  - Leverage iRODS rule to get informed after file close, harvest and add metadata using customized tools
  - Example project for NETCDF files: https://github.com/d-w-moore/extract_netcdf_header_msvc

© IBM Corporation 2020 and © CMCC Foundation 2020

Summary

- With iRODS you can address key challenges of HSM managed file systems:
  - Prevent transparent recalls and perform tape optimized recalls instead
  - Allow user to display the file state
  - Apply quota setting for the entire file system, including tape storage
  - Automatically insert metadata based on file type and content

- Furthermore, the capabilities of iRODS allow you to build and manage your archive efficiently
  - Global names space across different storage systems and locations
  - Metadata governance enables users to find archived data
  - Rule engine with flexible policies allow custom integrations and workflows
  - Standard interfaces for data access and administration enable efficient data sharing and collaboration

© IBM Corporation 2020 and © CMCC Foundation 2020

https://www.ibm.com

https://www.cmcc.it/

Disclaimer

This information is for **iRODS user group 2020**, publication beyond this scope is forbidden

This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information is provided for information purposes only as a high level overview of possible future products. PRODUCT SPECIFICATIONS, ANNOUNCE DATES, AND OTHER INOFORMATION CONTAINED HEREIN ARE SUBJECT TO CHANGE AND WITHDRAWAL WITHOUT NOTICE.

USE OF THIS DOCUMENT IS LIMITED TO SELECT IBM PERSONNEL THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.

Important notes:

IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

IBM makes no warranties, express or implied, regarding non-IBM products and services, including but not limited to Year 2000 readiness and any implied warranties of merchantability and fitness for a particular purpose. IBM makes no representations or warranties with respect to non-IBM products. Warranty, service and support for non-IBM products is provided directly to you by the third party, not IBM.

All part numbers referenced in this publication are product part numbers and not service part numbers. Other part numbers in addition to those listed in this document may be required to support a specific device or function.

MHz / GHz only measures microprocessor internal clock speed; many factors may affect application performance. When referring to storage capacity, GB stands for one billion bytes; accessible capacity may be less. Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and the population of all hard disk drive bays with the largest currently supported drives available from IBM.

IBM Information and Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both: the e-business logo, IBM, IBM Spectrum Scale, IBM Spectrum Archive

Other company, product, and service names may be trademarks or service marks of others.

CMCC disclaimer

The Foundation Euro-Mediterranean Centre on Climate Change has its registered office and administration in Lecce and other units in Bologna, Venice, Capua, Sassari, Viterbo and Milan. The CMCC Foundation doesn’t pursue profitable ends and aims to realize and manage the Centre, its promotion, and research coordination and different scientific and applied activities in the field of climate change study.

© IBM Corporation 2020 and © CMCC Foundation 2020