Quick Start Guide

Version 3.51
Part number: 0315-OA-QG-E-351
Second edition: March 2015

GRAU DATA

Introduction

About this Quick Start Guide

In this document several terms are used to describe the product but they all refer to the same software. OPENARCHIVE (OA) is the open source version of the Archive Manager (AM) from GRAU DATA.

“OA” is used in this document instead of the former term “IVD” (Infinistore Virtual Disk) from a time when GRAU sold integrated systems including server, disk and tape libraries. “OA” and “IVD” are interchangeable. They are different names for the same product.

Technically “IVD” is still used for object names in the code and for the names of the binaries. The names of the binaries may change in the future, but the term will still occur in the code.

This document was written mainly to describe the handling of the preinstalled Tutorial system inside the VirtualBox-Session, which you can download from www.openarchive.net.

For detailed information about the installation of OA please see the Installation Guide, which is available on www.openarchive.net.

What is OPENARCHIVE?

OPENARCHIVE (OA) is an archiving software which manages millions of files on standard hardware like FC or SCSI disk systems and/or LTO-based tape libraries from any vendor.

OA supports a full-tiered storage concept consisting of performance disks (e. g. SAN), archive disks (e. g. SATA) and tape (see configuration examples under 2.1).

OA has been especially developed to archive data of data-intensive applications like SAP R/3, DMS, e-mail, CAD/CAM, digital imaging, as well as for vertical applications such as audio/video, prepress and medical archiving.

OPENARCHIVE software provides a powerful set of rules, which enables efficient and automated data archiving and retrieving.

OA systems provides logical partitioning, thus allowing each application or user to have a separate logical and physical partition with its own set of rules.

For more detailed information, see www.openarchive.net.

OPENARCHIVE System Requirements

Productive System

For detailed information about the required system configuration to run OA, please see the FAQ on www.openarchive.net.

Tutorial System

The Tutorial in this document references an already installed Debian 6 (Squeeze) System containing the Archive Manager. This system is downloadable as a pair of disk images for the open source virtualizer “VirtualBox” by Oracle.

To run the session, the following is needed:

- System with VirtualBox 4.3 (or higher) installed
- min. 512 MB RAM
- min. 5 GB free disk space
- the disk images: oatutorialda(.vdi) (ca. 1,7 GB, can become up to max. 32 GB)
  oatutorialdb(.vdi) (ca. 40 MB, can become up to max. 64 GB)

For further information about handling disk images in Virtual Box and creating a session, please see the Virtual Box documentation.

OPENARCHIVE Overview

OPENARCHIVE (OA) is a mass-storage-oriented archiving software based on a client-server technology.

It enables transparent and automated data management with fast data access, economical long-term storage and data protection. Data archiving and retrieval can be automated and customized for each end user by implementing logical partitions of the system.

The efficiency of the OA implementation is achieved with a storage capacity that consists of primary and secondary storage space. These two types of storage space embody a successful combination of the speed of disk storage and the capacity and flexibility of tape media.

In an OA implementation secondary storage space stores all data and the primary storage space is used as cache for recently used files. Policy-driven archiving and retrieval processes make it possible for OA administrators to tune the system for best performance based on the type of data and the way applications handle that data.

OA software manages both primary and secondary storage space. It is highly modularized and implemented with a series of intercommunicating components that are responsible for data management. For detailed information about these components, see chapter 1.5.

OPENARCHIVE Configuration 1

Two local copies

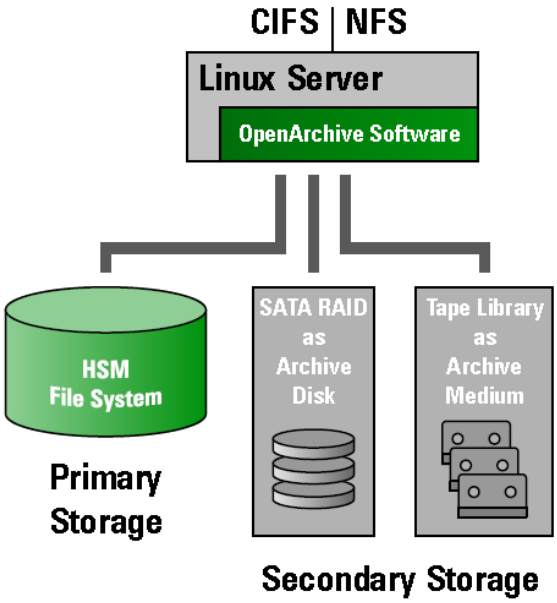

OPENARCHIVE Configuration 2

Two local copies and two remote copies

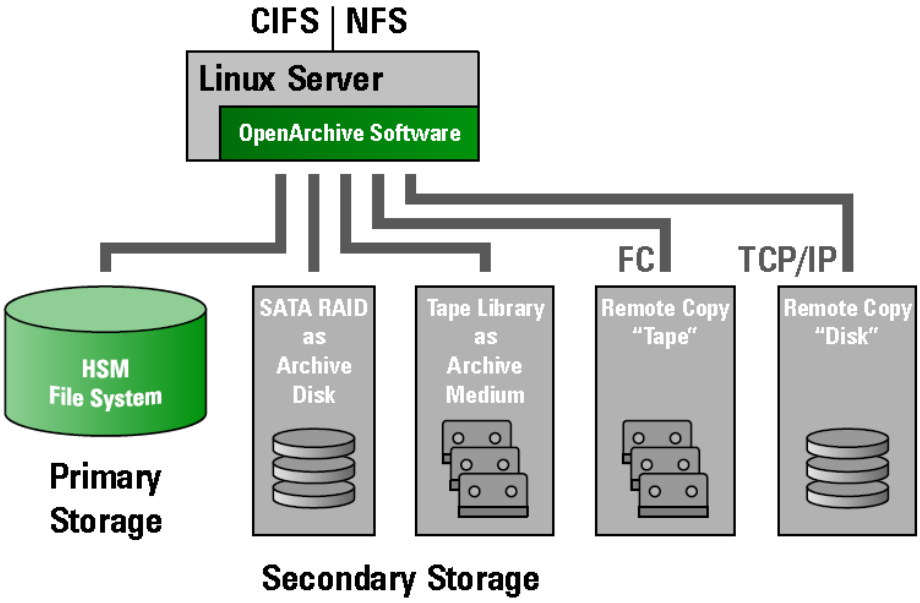

For details on supported and required hardware and software, see the FAQ on www.openarchive.net.

OPENARCHIVE Basics

Infinite file system (HSMFS)

The most significant characteristics of OPENARCHIVE (OA) are a virtually infinite file system and very transparent access to the data through the file system.

OA gives you access to virtually unlimited storage capacity. OA constantly controls the files in use. You define the way OA controls these files with migration, release, and deletion policies. Additionally, OA ensures ultimate data availability and safety with data replication.

It also enables multiple copying of migrated files, thus satisfying the need for additional backup copies and avoiding recall failures.

Library - tape compatibility

OA software supports tape libraries from all major vendors.

OA software supports Sony AIT drive technologies for AIT-2, AIT-3, Sony SAIT drive technology for SAIT-1 media, and LTO drive technologies for LTO Ultrium 1, LTO Ultrium 2, LTO Ultrium 3, LTO Ultrium 4, LTO Ultrium 5 and future LTO technologies.

With AIT and SAIT drives and media, OA software also supports special memory chips, called Memory-In-Cassette(MIC), which are integrated in the AIT and SAIT tape cartridges and enable the corresponding drives to locate files much faster than with conventional tape cartridges. With LTO Ultrium drives and media, similar functionality is provided by Medium Auxiliary Memory (MAM).

For a list of the supported tape libraries, tape drives, and tape media, see the FAQ on www.openarchive.net.

Automatic error detection

OA automatically detects drive and media-hardware errors. When OA detects such errors, the following happens:

- The drive status is set to error.
- If an error occurs during a migration job, the medium status is set to unreliable, which disables further writing to it.
- If an error occurs during a recall job, the medium status is set to unusable, which disables reading from and writing to it. Neither the drive nor the medium in it is used for the operation until the administrator appropriately changes the drive and medium statuses.

In each of the above cases, an appropriate entry is written to the OA event log. If you find an OA medium marked unreliable or unusable, the OA event log and OA error log need to be searched backwards to identify the root cause of the problem.

Automatic detection of low storage space

The HSM Health Monitor (HHM) utility automatically detects conditions related to low storage space, such as

- low disk space on file systems with OA databases and system files,
- low disk space on HSM file systems,
- low secondary storage space.

If any of the above conditions occurs, an e-mail is sent and a configurable action can be triggered. Likewise, when the condition is not present anymore, an e-mail is also sent and a configurable action can be triggered.

Typical actions that may be triggered involve putting the monitored HSM file systems into Limited Access Mode, notifying the OA users, shutting down the OA system in case of unrecoverable situations, or putting the monitored HSM file system online again in case the condition is no longer present.

NFS and CIFS support

Remote Linux and Windows clients access HSM file systems using standard protocols. Network File System (NFS) is the standard protocol for UNIX and Linux platforms. Common Internet File System (CIFS) is the standard protocol for the Windows platform.

OA supports access to OA partitions from remote locations via NFS and CIFS protocols. HSM file systems that reside on Linux systems can be exported to be accessed from remote clients via the NFS protocol.

Access from remote clients using CIFS protocol is enabled by a specially customized Samba software. Files that are written to or read from an HSM file system through a Samba connection are handled like all other files created or accessed locally on the HSM file system. The customized Samba server notifies remote CIFS clients (especially Windows Explorer on Windows platform) when a file is offline, thus preventing unnecessary recalls of offline files.

For further information about the requirements and limitations of the Samba integration on Linux systems see the FAQ on www.openarchive.net.

OPENARCHIVE Architecture

OA takes complete control of the HSM file systems, thus maximizing their operational efficiency.

Handling of the files residing on the HSM file systems is defined with different policies.

Basically, OA copies new or modified file information from a file on an HSM file system to OA media, releases rarely used files and makes them offline, and recalls offline files when necessary.

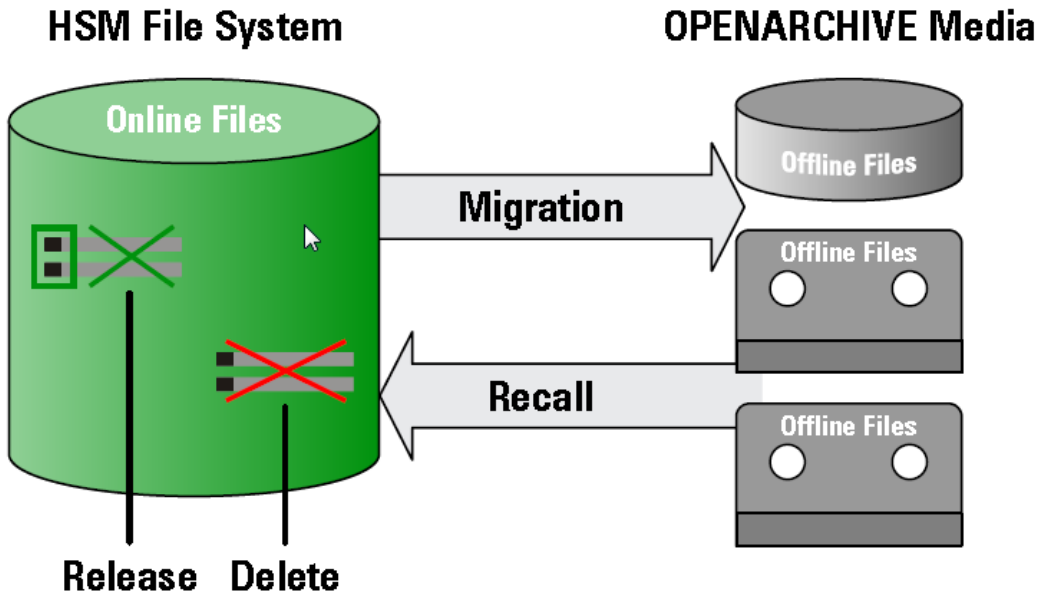

OPENARCHIVE operations

The OA main operations are migration, release, recall and delete. They are configured with user-defined policies. After they are configured, the policies enable automatic, unattended execution of the corresponding operations. However, these operations can also be triggered manually using the **ivdfile** command. For details on manual triggering, see the OA man pages.

Regardless of how they are invoked, the migration and recall operations are represented and controlled in the OA implementation by so-called “OA jobs”. Using the **ivdjob** command, an OA administrator can display information about existing jobs or manipulate them. Examples of situations when such manipulation is needed are the following:

■ The priority of a particular recall job has to be increased, because urgently needed documents have to be recalled from OA media, and there are currently other jobs with higher priorities running in the OA implementation.

■ A library in the OA implementation has run out of OA media and therefore additional media need to be added to it and initialized, in order to get the media available and migration jobs running again.

For more informations about the **ivdjob** command see the OA man pages.

Migration

Migration is a process of copying files from the HSM file system to OA media. If a file is not modified for a configured period of time, it passes the so-called “file age check” and becomes a candidate for migration (migration candidate).

Migration is done in two steps:

1. Files from the migration candidate list are copied from the HSM file system to a special storage space on the OA server, called OA disk buffer. The root directory of the OA disk buffer is

/var/opt/ivd/diskbuf.

2. Data from the OA disk buffer is copied to OA media (for all copies defined for the partition).

It is recommended that the diskbuffer should be at least twice the size of the biggest expected file.

Once the migration job is completed, the temporary location is cleared. The file entry in the Hierarchical Storage Manager (HSM) is moved from the migration candidate list to the release candidate list.

Release

The purpose of the release process is to free space on the HSM file system. During this process, file data is removed from the HSM file system, and only file headers are preserved. Release can only occur after the file has successfully been migrated. Then a release will occur for a migrated or recalled file after the policy is invoked which requires that free space be added to the HSM file system.

Once a file has been migrated to OA media, it becomes a release candidate. OA checks the release candidates for changes. If a file on the release candidate list was not modified, and the space usage on the HSM file system reaches the high watermark, the content (file data) of the file is released from the HSM file system to free disk space. A similar event happens when the space usage on the HSM file system reaches the critical watermark. In this case, however, all migrated files are released, regardless of how much time has passed since their last migration.

Once the release operation is completed, the file is offline and its entry is removed from the release candidate list.

Recall

Recall is a process of copying file contents from OA media back to an HSM file system. It is normally triggered when an offline file is accessed; therefore it is the reverse process to migration.

Recall is done in two steps:

1. File content is copied from the OA media to a special storage space on the OA server, called OA disk buffer.

2. File content from the OA disk buffer is copied back to the HSM file system, thus making the file online again.

Once the recall job is completed, the temporary location is cleared, and the offline file changes its status to online. Its file entry in the Hierarchical Storage Manager is added to the release candidate list. This process is transparent to the user.

A recall is triggered automatically when an offline file is accessed on the HSM file system. However, OA also stores on the OA media older versions of files, referred to as older generations. You can manually recall a particular older generation of a file using the OA command-line interface. Sometimes you may even want to have files recalled before they are actually needed. OA includes functionality to manually recall a single file or a set of files from OA media. You can perform such manual recall using the OA command-line interface.

For details on recalling a single file, see the **ivdfile** entry in the man pages.

Deletion

Deletion is a process of obsoleting and removing obsolete files from HSM file systems. You can use this process for various purposes, including Information Lifecycle Management (ILM) and freeing disk space of redundant files.

The criteria that induce file obsolescence on HSM file systems are defined in the deletion policy. Each HSM file system has its own deletion policy.

OPENARCHIVE components

HSM file system

An HSM file system (HSM FS) is a hard disk area reserved for storage of user files and directories (user data), and managed by the OA implementation. The HSM file system is used to store the metadata (names, attributes, and permissions of user data) and user data that is currently online.

The HSM file system is based on the Ext3 file system; the HSM attributes are stored in Ext3 extended attributes.

Hierarchical Storage Manager (HSM) is a component of the HSM file system. As a part of an OA client, it handles HSM file system events and manages the migration and release candidate lists. Each mounted HSM file system has one OA HSM.

In a distributed OA implementation, the HSM file systems, their Hierarchical Storage Managers (HSMs), and the File System Event Manager are located on external OA clients.

OA partition

OA uses OA partitions to split the storage into multiple entities, meaning that one OA implementation can have more than one OA partition. This way, OA provides the possibility for a complex organization of data management.

Each OA partition is related to a single HSM file system and has its own policies and its own File System Catalog (FSC). This enables you to configure individual policies (migration, release, deletion, and resource allocation) for a particular OA partition, according to the characteristics of files and their intended use.

Each configured OA partition is controlled by one Partition Manager (PM), which creates and manages OA jobs for migration, recall, and deletion operations. The Partition Manager is located on the OA server. It is included in the OA server installation package.

File System Catalog

The File System Catalog (FSC) is a database related to a single OA partition. It consists of a Data Location Catalog (DLC) and a Name Space Catalog (NSC).

The Data Location Catalog contains information about the locations of files on the OA media. It stores a full history of file locations. A file entry is made in the Data location catalog after the file has been migrated.

The Name Space Catalog (NSC) contains metadata of files on an HSM file system. It only stores the latest generation metadata, which includes directory structure, standard attributes and additional platform-dependent metadata. This data enables recovery of an HSM file system.

For HSM file systems located on Linux OA clients, NSC also stores ownership information of directories and files.

File System Catalog is located in the **/var/opt/ivd/part/[PartitionName]/fsc** directory.

OA media pool

An OA partition must have at least one media pool assigned to it. An OA media pool is configured to group media of the same type, which have the same characteristics.

Media pools are configured regarding their purpose and correspond to the specifics of OA partition data and OA media types, e. g. disk or tape media pool. Therefore, one media pool cannot be assigned to multiple partitions, thus assuring that files from different OA partitions remain separated. An OA partition can have assigned multiple media pools which do not need to be of the same type. This means that disk and tape media pools can be combined on the same OA partition.

By assigning multiple media pools to one OA partition, multiple copies of OA data are written to the media.

OA medium

OA supports different types of tape media, as well as disk media for data storage. Each OA medium, either a medium cartridge or disk medium, has an entry and properties written in the Resource Management Database (RMDB). Each OA medium is assigned to one OA media pool.

Depending on the particular medium type, an OA medium can be formatted and further divided into partitions that hold medium volumes. The Resource Manager identifies the medium with a barcode.

OA library

An OA library is any library of the supported tape libraries which is attached to the consolidated OA system or OA server in an OA implementation. An OA implementation can have one or more OA libraries. Each OA library is a repository of library slots, where medium cartridges are physically located, with a built-in changer device to move the media around. It has one or more connected drives.

Each configured OA library has one Library Agent running. The Library Agent handles actions in the OA library, such as loading and unloading media, and library inventory rescan.

OA implementations which only use OA disk media do not have OA libraries.

OA drive

An OA drive is a tape drive inside an OA library, configured in the OA implementation. It is used to read to and write from OA media.

Supported drive types are:

- LTO Ultrium 1, LTO Ultrium 2, LTO Ultrium 3, LTO Ultrium 4 and LTO Ultrium 5 and future LTO technologies
- AIT-2, AIT-3
- SAIT-1

For further information about the supported drive models, see the FAQ on www.openarchive.net.

OA RMDB

The Resource Manager DataBase (RMDB) keeps track of all components. These include:

- Drives
- Media
- Media Locations
- Media Volumes
- Tape Libraries

It also contains information about the relationship to the OA partition. The RMDB is changed whenever the files in the configuration database are changed by using component-modification options such as **--add**, **--modify** and **--remove**.

In case of a disaster recovery, it is crucial to have a recent version of this database. The **ivdbackup** command automatically collects this information. Please refer to the man pages regarding the **ivdbackup** command.

Tutorial: HSM - Using and Operation

Login to the Tutorial System

After starting the virtual machine, the linux login screen appears.

Login is **root**
Password is **oaisgood**

Configuration of the Tutorial System

The tutorial system consists of one managed file system: This means one OAroot partition (two copies for each file, meaning two media pools) and one disk medium in each media pool. Some disk media have already been low-level prepared, but not yet assigned. The partition parameters have been adjusted to fit for a small test system.

The current status of the system can be viewed on the command line.

Command overview

By using the following commands, you can review the components of the OA system.

In order to use OA commands, you have to be logged on to the system as root.

The command options are mostly of the same layout. The option **-h | --help** *gives you an overview of the available options and their usage for almost every OA command.*

List existing Libraries

The command for handling OA libraries is **ivdlibrary**. The option **-l | --list** generates a list of existing libraries and their details.

The Tutorial system does not contain any libraries.

List existing Drives

The command for handling OA drives is **ivddrive**. The option **-l | --list** generates a list of existing drives and their details.

The Tutorial system does not contain any drives.

List existing Media Pools

The command for handling OA media pools is **ivdpool**. The option **-l | --list** generates a list of the existing pools and their details:

```
[root@localhost ~]# ivdpool –l
Pool  Family Type     Partition  BlkSize[KB]     #Vol VSize[MB]
      SysVol      SysVol[MB]
dmpool1   disk  regular     fs01  64          1       100         none n/a
dmpool2   disk  regular     fs01  64          1       100         none n/a
```

List existing Partitions

The command for handling OA partitions is **ivdpartition**.

The option **-l | --list** generates a list of the existing partitions and their details:

```
[root@localhost ~]# ivdpartition –l
Partition   Status          Host              Mountpoint
fs01        mounted   hsmserver.hsmnet    /hsm/fs01
```

The option **-t <partition name>** generates a statistical overview about the corresponding partition.

List existing Media

The command for handling OA media is **ivdmedium**. The option **-l | --list** generates a list of the

existing media and their details:

```
[root@localhost ~]# ivdmedium –l
Barcode    Type Pool        Status          Location   #Vol SysVolNo
dm01            disk  dmpool1   open        n/a        3    n/a
dm02            disk  dmpool2   open        n/a        3    n/a
```

List running Jobs

The command for handling OA jobs is **ivdjob**. The option **-l | --list** generates a list of the running jobs and their details.

View Log File Entries

The logfile **ivd.log** is located in **/var/opt/ivd/log** and can be viewed with any standard text

viewer (e. g. **vi**). To view only the last entries, use **tail /var/opt/ivd/log/ivd.log**.

View status of OPENARCHIVE Services

The command **ivd** with the option **status** generates an overview of the OA services and their statuses.

```
[root@localhost ~]# ivd status
ivd-svc ( pid 2180 ) is running...
```

ivd-rm (pid 2190) is running...
ivd-mif (pid 2202) is running...
ivd-fsevtmgr (pid 2495) is running...

The following services are listed:

■ **ivd-svc**
Launches other OA daemons (services) and OA agents.

■ **ivd-rm**
Manages the Resource Management Database and allocates OA resources (OA drives, OA media, and so on).

■ **ivd-mif**
Handles user requests issued with OA commands and manages OA administrative jobs.

■ **ivd-fsevtmgr**
Intercepts mount events for the HSM file system and starts the Hierarchical Storage Manager (HSM).

Assigning Media to a Media Pool

Two media are already applied in the tutorial system, to facilitate an immediate start. For information about applying new media, see chapter 3.

To assign a medium (dm01 & dm02) to a media pool (dmpool1 & dmpool2), use the following command:

ivdmedium -a dm01 -O dmpool1
ivdmedium -a dm02 -O dmpool2

To check if the medium is assigned to the media pool, use the following command:

ivdmedium -l

The media shown by **ivdmedium -l** *are still uninitialized, so the next step is required to make the media available.*

Format and initialize Media

To format and initialize a medium (dm01 & dm02), use the following command:

ivdmedium -f dm01 –i
ivdmedium -f dm02 -i

If the job cannot be finished and **Requesting rescources** *is displayed, see Troubleshooting / Migration Problems on www.openarchive.net for help.*

To check if the medium is formated and initialized, use the following command:

ivdmedium -l

The medium status should be **free** now.

Store Files to the File System

To see the working progress of OA, there have to be some files on the file system that can be handled.

Create a directory (test1) on the HSM file system
mkdir /hsm/fs01/test1

Create a 4 MB file (file1) in the directory (test1)
dd if=/dev/zero of=/hsm/fs01/test1/file1 bs=1024 count=4096

Create two more 64 MB files (file2, file3) in the directory (test1)
dd if=/dev/zero of=/hsm/fs01/test1/file2 bs=1024 count=65353
dd if=/dev/zero of=/hsm/fs01/test1/file3 bs=1024 count=65353

Watching the Migration Process

Check the HSM files statuses:
ivdfile -l /hsm/fs01/test1/*

The files are now shown as “dirty” files. The term “dirty” is rootet in the cache terminology because the files need to be written to the secondary storage.

After the file has been migrated, it is shown as “online”.

Check the fs01 assigned media statuses:
ivdpartition -t fs01

Depended on the OA setting, it takes a while until the migration process handles the files.

Release a file (file1):
ivdfile --release /hsm/fs01/test1/file1

Check current directory size of /hsm/fs01/test1
ls /hsm/fs01/test1 -l

There should be three files visible in the directory (file1 4 MB, file2 64 MB, file3 64 MB), with a combined size of 132 MB. Nevertheless, because file1 is already released, the directory size is only 128 MB.

total 130860
-rw-r--r-- 1 root root 4194304 Feb 14 13:15 file1
-rw-r--r-- 1 root root 66921472 Feb 14 13:18 file2
-rw-r--r-- 1 root root 66921472 Feb 14 13:18 file3

Recall a file (file1):

The released files are automatically recalled when they are touched by the file system:
file /hsm/fs01/test1/file1

If the directory size of /hsm/fs01/test1 is checked again now, it should be 132 MB.

ls /hsm/fs01/test1 –l

Tutorial: Creating Disk Media from scratch

For further information about linux commands such as **pvscan**, **lvcreate**, **lvscan**, *... see the appropriate man pages.*

Locate free disk space
pvscan

Create two new 1 GB logical volumes (dm03, dm04) in the volume group hsm_hdb
lvcreate -L 1G -n dm03 hsm_hdb
lvcreate -L 1G -n dm04 hsm_hdb

Check logical volumes
lvscan

Create an ext3 Filesystem on the new logical volumes (dm03, dm04)
mke2fs -j /dev/mapper/hsm_hdb-dm03
mke2fs -j /dev/mapper/hsm_hdb-dm04

Create an fstab entry

Open **/etc/fstab** in a text editor (e. g. **vi**).

Enter the following text below the lines of the existing logical volumes dm01 and dm02:

/dev/mapper/hsm_hdb-dm03 /var/opt/ivd/dm/dm03 ext3 defaults 1 2
/dev/mapper/hsm_hdb-dm04 /var/opt/ivd/dm/dm04 ext3 defaults 1 2

Create Mountpoints for disk media
Create one folder for each new disk medium:
mkdir /var/opt/ivd/dm/dm03
mkdir /var/opt/ivd/dm/dm04

Mount disk media
mount /var/opt/ivd/dm/dm03
mount /var/opt/ivd/dm/dm04

Check mounted volumes
df -h

Add disk media to OA as unassigned media
ivdmedium -a dm03
ivdmedium -a dm04

View disk media statuses
ivdmedium –l

Tutorial: Creating an OPENARCHIVE partition

Create new media pools (dmpool3, dmpool4)

Create a copy of the disk media template for dmpool3 & dmpool4
cp /opt/ivd/newconfig/ivd-pool-DISK.cfg /tmp/dmpool3.cfg
cp /opt/ivd/newconfig/ivd-pool-DISK.cfg /tmp/dmpool4.cfg

Make the template copy writeable
chmod a+w /tmp/dmpool3.cfg
chmod a+w /tmp/dmpool4.cfg

Edit the template copy (e. g. with vi)

Modify the parameters as follows:
Name = "dmpool3" **(dmpool4 in file dmpool4.cfg)**
MediaFamily = DISK
PoolType = Regular
BlockSize = 64KB
VolumeSize = 100MB

Add a new media pool to OA by using the template
ivdpool -a /tmp/dmpool3.cfg
ivdpool -a /tmp/dmpool4.cfg

Check the existing media pools:
ivdpool -l

Remove the temporary config file
rm /tmp/dmpool3.cfg
rm /tmp/dmpool4.cfg

Creating a partition (fs02)

The first steps for creating a partition must be performed on the command line.

Create the mount point for the partition (fs02)
mkdir /hsm/fs02

Create the HSM file system
Create a new 1 GB logical volume (hsmfs02) in the volume group hsm_hda:
lvcreate -L 1GB -n fs02 hsm_hda

Create the file system on the new logical volume (hsmfs02):
mke2fs -j /dev/mapper/hsm_hda-fs02

Create the fstab entry
Open **/etc/fstab** in a text editor (e. g. **vi**).

Enter the mapper device with the following parameters below the line of the other hsmpartition hsmfs01,

e. g.:
/dev/mapper/hsm_hda-fs02 hsmfs noauto 0 0

Create a copy of the partition template for fs02
cp /opt/ivd/newconfig/ivd-partition.cfg /tmp/fs02.cfg

Make the template copy writeable
chmod a+w /tmp/fs02.cfg

Edit the template copy (e. g. with vi)

Modify the parameters as follows:

Name = "fs02"
Client = "hsmserver.hsmnet"
FileSystemID = "/dev/mapper/hsm_hda-fs02"
Type = "Regular"
Pools = ("dmpool3", "dmpool4")

Add a new partition to OA by using the template
ivdpartition -a /tmp/fs02.cfg

Check if the partition is added:
ivdpartition -l

Mount the partition:
mount /hsm/fs02

View status of the file system:
df -h

Remove the temporary config file
rm /tmp/fs02.cfg

Templates

For each resource type, a default configuration template is available so that you do not need to create configuration files from scratch. The templates already contain most of the required parameters, along with remarks about the appropriate use of configuration variables and details of maximum and minimum values.

When configuring a new OA resource, make a copy of the appropriate template, remove the read-only flag in the copied file’s properties, modify the copy according to the characteristics of the resource, and then create the resource using the copy.

While modifying templates for configuration of new resources, preserve the double quote characters (") that surround the values of some variables. Removing these characters may lead to problems with adding new resources in certain cases.

The templates are located in the following directory on the consolidated OA system or OA server:

/opt/ivd/newconfig

To configure your resources and your policies, use OA commands or the OA Management Console.

Troubleshooting

In case you have any problems with OPENARCHIVE, please see Troubleshooting on www.openarchive.net for information and help.