MOST
Multimedia Oriented Systems and Technologies s.r.l.

MOSTdoc & Centera Integration: Implementation details

MOST
Multimedia Oriented Systems and Technologies s.r.l.

Contents

MOST
Multimedia Oriented Systems and Technologies s.r.l.

Overview

This paper describes the integration of Centera and MOSTdoc, a data archiving product. The result is a fast, secure and reliable solution that delivers thousands of documents per day in a distributed environment inside one of the biggest italian banks. The main idea behind the integration is the use of the CDRom format (the ISO 9660 standard), as the main content of the C-clip on the Centera. The rationale is that this is a powerful format, and we were able to perform a fast and easy move of the content of existing juke boxes.

The idea can lead also to some interesting proposals, like the use of the digital signature, a standard feature of MOSTdoc, to logically seal the content of a C-clip.

Before Centera: MOSTdoc, a powerful document archiver

MOSTdoc is a well established product for data archiving. Aimed at the Cold and Scanned images needs, it finds its best use in institutions like banks, enterprises, government agencies that must safely archive many documents that can be addressed using some well defined indexes. Following this approach, for search purposes there is no need to access the document itself, a relational DB is enough. A pointer to the document is the only requirement for a later extraction of the document.

One hidden advantage of this idea is that the data used to index the document and the physical location of the document are coupled in a loose way and, moreover, documents can be moved elsewhere without disrupting the application. This concept is similar to HSM, but it is implemented inside the product, using application dependent aging policies.

For many historical reasons the CD Juke Box has been widely used in the past for long term archiving: low media costs, well known and universally adopted file system format (ISO9660), transparent use of the physical media in any PC in the office, even without the Juke Box.

In order to take advantage of the intrinsic usability of the media, we developed a structure, based on JavaScript and HTML, that mimics the behavior of the on line user interface offered by MOSTdoc. So, any CD extracted from the Juke Box can be queried using the usual Web browsers.

Why: Replace the CD Juke Box

Even the apparent advantages of the JB solution were not enough to deal with the drawbacks of the hardware: poor reliability, low transfer speed, less than optimal storage capacity.

So we agreed with our customers that we could integrate the Centera, instead of the Juke Box, within our product, with a few goals in mind:

- leverage on the flexibility of the Centera deployment (ethernet instead of scsi means no single, dedicated server)
- closely match the interface speed, avoiding application bottlenecks
- minimize the internal changes in our product
- avoid any data processing during the conversion of the existing data from the Juke Box to the Centera

The solution, in our opinion both elegant and fast, is to use the ISO9660 images as the allocation unit on the Centera (the *container* in EMC terminology)

The advantages are clear: the Centera can be seen as another kind of juke box, with minimal impact

on the software.

Our software already knows different types of storage backends, completely transparent to the data used to search the documents, so the Centera has become one of the possible storage choiches, among BLOBS, File System, Mainframe Data Sets, Magneto Optic, CDRom or Virtual Juke boxes, without any change to the user look and feel.

The Front End interface is very flexible too: MOSTdoc can behave as a Web based application, an IMAP server for E-mail archives, a Web-Dav Repository, an XML-HTTP application to application server.

How: The details

Application tables hold the pointer to all physical documents by means of indirect pointers. The intermediate table holds the current storage details, and can be changed according to the Information Lifecycle needs. It is quite common for the documents to appear inside a blob and, after a while, to be consolidated in an ISO9660 image. When using removable media devices (f.i. with the CD Juke Box), two more intermediate tables are involved, just to hold the media information (location, label).

CD content is produced by our application, that creates a directory tree that is then translated in a ISO9660 image; Iso images are then handled, via a custom Inter Process Communication, by a Juke Box Daemon, living on the computer to which the JB is connected. The daemon deals with obvious chores like CD mastering, media load and store, mount and unmount of the file system, file extraction.

CD labels are uniquely defined, and these identifiers are held both in the MOSTdoc tables and in the label field of the ISO image.

Centeratool 1.0

So, in order to comply with hardly believable time constraints, we had to deploy a step by step implementation, carefully minimizing the changes to the existing (and working) code of our product.

These steps were:

1. to provide a one to one replacement of a physical juke box, with a dedicated server that handles the Centera, perform a copy of the desired media on the hard disk, mount it, extract a copy of the desired file, and keep a cache of recently accessed iso images
2. to provide simple tools for an ordered transfer of the iso images from the CD Juke to the Centera
3. to deliver an optimized version, that is able to extract directly a file from the iso9660 structure

So, our first effort was to divide the problem in smaller chunks:

- a standalone program that deals with the Centera api (Centeratool)
- some changes in our internal structures (the Centera clipId is larger than our existing CD identifiers), and software. Changes were small, and limited mainly to the juke box daemon, that instead of issuing SCSI commands to the jukebox, had to spawn the Centera tool with the right commands
- tools to safely move iso images from the existing jukebox to the Centera

We delivered on time, and were ready to refine the project, while our customer happily moved the documents inside the Centera, and began to extract pdf and tiff files from the Centera.

Speed was obviously non optimal, though bearable: a typical full iso image of about 700 Mbyte took about 50 to 90 seconds to travel from the Centera to the client hard disk (with a higher than expected CPU load on the API side). These figures are slightly higher compared to the typical operation cycle of a juke box, that take about 30 seconds: unload a media from the drive, park it in the silo, pick the desired media, load it, spin it up.

But we were a clear winner over the worst case scenario of the physical juke box: a shelf based media not only takes minutes to complete, but also involves a manual operation, with high risks of scratching the surface of the media, loading the wrong item, etc.

Centeratool 1.2

With the confidence of 6.000 CDs migrated to the Centera in less than a month, having only to deal with broken medias, we were able to plan the obvious evolution of the project: access speed

There were two possible solutions:

1. use the Centera to hold single files
2. access directly the file inside the iso image, using the Centera api, and by implementing procedures that can deal with the iso9660 file structure

Our customer was in favour of the first solution, that was simpler and looked more compliant with new italian regulations about data archiving, but we opposed a few remarks:

- small files are mirrored on the Centera, so space would be lost
- the existing CD on the Centera had to be reopened, and single files had to be saved one by one on the Centera
- data backup, disaster recovery, and in general, any activity that involves copies of data from one place to another, work better with bigger chunks
- the italian law regulates the so called “substitute archiving”, i.e. a CD Rom digitally signed with a smart card, and time stamped from a regulatory agency. Combining the retention capability of the Centera and the strong authenticity warranties offered by the digital signature is an exciting opportunity to be implemented in the near future

So we added the knowledge of the iso9660 file system (with Joliet extensions) to the Centeratool, and, by navigating the Table of Contents and directory structures of the iso9660, access times to 100KBytes files where reduced to about one second.

Plus: The virtual JB

The reader already knows that we have a strong belief in data consolidation using well known strucures. As a matter of fact we used the same ideas on other storage devices, for what we called the *Virtual Juke box*. Roughly speaking we hold iso images on a convenienent storage (Raid, Nas, San), and use the ability of performing a “loop” mount on the file. The virtual juke box deamon hides completely the fact the Juke Box is just a bunch of hard disks.

Our Centera solution capitalize on this experience, but it's obviously better:

- transparent backup and disaster recovery provided by EMC
- faster access (Nas are somehow slow when performing loop mounts on big files)
- the clip structure could potentially hold a lot of useful information about the CD

MOST
Multimedia Oriented Systems and Technologies s.r.l.

And now: new ideas

We believe that combining Centera with MOSTdoc we can provide clever solutions to our customers' needs, but often it is better to refine the solution with the customer, instead of blindly developing an idea. So we would like to cohoperate with someone in the following areas:

- general purpose JB clone: developing a SCSI device driver that instead of speaking with a real juke box, will issue fake commands to the Centera
- digitally signing of existing Cds: when moving the CD from the existing Juke Box, it would be intriguing to add a digital signature of it, allowing a convenient check of the CD content

Appendix A. Centeratool synopsys

Centeratool(n) Centeratool programming manual Centeratool(n)

NAME

Centeratool - standalone Centera client

SYNOPSIS

Centeratool command [arguments]

DESCRIPTION

Centeratool is a command line program that allows an easy interface to EMC Centera functions

USAGE

info

Cluster information on stdout

capabilities

Restituisce in stdout le capabilities del cluster, relative al profilo corrente dell'aplicazione.

get CLIPID OUTFILE

Writes in *OUTFILE* the first blob contained in the given *CLIPID*

isoget *CLIPID PATH OUTFILE*

Given an iso9660 image in the first blob of *CLIPID*, writes in *OUTFILE* the file *PATH*. *PATH* can be a pathname inside the iso, optionally with joliet syntax. Rockridge extensions are not supported

put *FILE*

Creates a new clip, holding the given *FILE*. the new CLIP-ID appears on sdtout

dput *DIR OUTFILE*

Foreach plain file in *DIR*, creates a new clip. *OUTFILE* will hold the Clip-ID list

delete *CLIPID*

Deletes a C-Clip and its blobs

recover *OUTFILE*

retrieves every clip on the Centera. clip-ids and timestamp are written in *OUTFILE*

WARNING: use with caution, it is very slow

AVAILABILITY

Linux, Solaris

AUTHOR

V.Gionco - MOST s.r.l.

Centera is a trademark of EMC

MOST March 2004 Centeratool(n)