White Paper Jukebox

(english version)

ASM GmbH & Co. KG
Vosskamp 1 • D-26655 Westerstede • Germany
• Tel ++49-(0)-4488-8496-40
• Fax ++49-(0)-4488-8496-11
• Internet: http://www.asm-jukebox.de
• E-mail: marketing@asm-jukebox.de

With this document we would like to explain some parameters who should be checked before deciding to run a particular jukebox solution.

We have focused on the most important parameters, but there should be still further ones to be considered carefully.

A jukebox can be run with three different types of applications:

- Corporate Critical Applications
- Department Critical Applications
- Person Critical Applications.

Of course for corporate critical applications it is of most importance that the jukebox meets all requirements, especially because system availability (e.g. influenced by fault tolerance and speed of recovery) is not only important for the operation in general. The system performance has basic influence on both, productivity of each employee and the workflow between employees and departments.

If the jukebox and the way the system can be operated in daily administration can not be adapted in its configuration to the individual requirements of a corporation, the jukebox becomes a bottleneck.
Although the jukebox is not defect it costs the company every day real money because productivity and workflow do not reach the level they could reach -> The jukebox as bottleneck blocks the application(s).

We would not like to introduce all technical details of a jukebox to you, but only the most important parameters. The following points should be a checklist for you to distinguish between suitable and non-suitable jukebox system - always based on your individual requirements.

These parameters are as follows:

Jukebox and Server - a Team
Jukebox and Server have to do permanent job sharing. On the one hand the jukebox does nothing without getting a start command (Job) from the server, but on the other hand the jukebox (based on its level of „intelligence“) handles this job then on its own.

This does not mean that there is no permanent communication between jukebox and server. There is a permanent communication of SCSI commands (status reports, confirmations, data flow...) on the bus, but this is not the same for every jukebox and server software.

If you select a jukebox the system should have a 100% conformity to the SCSI command standard. Every special command and routine requires adaptations of the server software - which is all the time a source of faults.

The working basis of a jukebox - the Job
A job is an inquiery of an application, e.g. an archiving-/DMS software, for a particular disk on which the required data is stored. The job can be complete positive or negative. Either by inserting the media into a drive and transmitting the data or with an error message. An error mesage can for example occur if a media slot is empty although it should carry a media. In this case such a system inconsistancy, caused by a database corruption, must be fixed by a reinitialisation of the server.

Important:
No jukebox does two jobs at the same time, but always job after job.

The performance of a jukebox is therefore mainly influenced by how fast it can work off jobs and recover/initialize after a system inconsistancy.

Recovery procedure - how a barcode reader works
In case of a system disarrangement, e.g. power failure or database corruption, the content of the jukebox must be checked. In this process all disk are identified and addressed to a particular slot number.

The process can take, e.g. for 300-500 disk (depending on hardware and software) more than one hour, because every media has to be put into a drive and read to be identified.

If you have a barcode reader in the jukebox this time period can be reduced drastically to only some minutes. Only the barcode on either CD-Caddy or WORM cartridge has to be read to properly identifying the disk and to address it to the slot number.

Barcode readers to contribute essentially to increase system availability, based on functions of the jukebox. They do speed up the recovery of the server and are therefore of real value.

For critical and fault sensitive applications a barcode reader should be standard equipment in a each jukebox system.

As side effect of the barcode label on CD-Caddy or WORM cartridge it is also much easier to identify a disk offline.

Burden Sharing
The aim of burden sharing is to utilize all jukeboxes connected to a server with the same intensity. If the total capacity is split into two or more physically separated units, burden sharing becomes a real administrative problem. The negative effect is that one or more jukeboxes have only few jobs to do and one or more jukeboxes cannot work off all jobs send from the server -> a bottleneck arises!

This situation exists always when the total number of installed drives cannot be

ccessed by all disks. A subset of disks has only access to a subset of drives.

Those negative effects can be limited by checking user statistics of each disk manually. This should be done frequently because such statistics are often influenced by variable factors. Temporarely massive accesses to a particular jukebox can of course not be avoided.

Another solution for this problem is to install a larger jukebox, because in such systems the robotic can access all drives and all disks. Burden sharing is done therefore automatically.

Better two Jukeboxes than one large system?

A basic consideration while planning a jukebox based solution is to split the total number of disks and drives into two or more units. This can make sense but is mainly influenced by the later operating mode in daily operation and the aspect of burden sharing. Also this strategy does not help in case of breakdown of the jukebox server or trouble caused by the management software running on the server.

The operation mode descripes the way of data handling in case of partial system inavailability, e.g. system breakdown or when a jukebox is switched off for service purposes.
Especially this data handling in case of a system breakdown outweights often the assumed advantages of a (jukebox) capacity split. The three major barriers are:

1.
When a disk should be transfered from the switch off system to another one, you have to have free slots or magazines there. This means it must be possible for the system administrator to distinguish between more or less important data on complete single disk.

This is usually a complicated procedure, because the data, e.g. of a tax file, is split onto several disks according to the time structure. As many files are never complete, e.g. health care files, you can never write a file onto a single disk.

But if the administrator can not select data in a jukebox for keep them temporarely offline, he can also not free an slots/magazines for disks transfered from the defect system.

In this case the only way is to keep slots/ magazines in all jukeboxes permanently free for such purposes, or the have an additional jukebox as backup system for any case of hardware problems with a jukebox.

2.
If you could manage to jump over barrier (1), you have to have the ability to export a disk with the jukeboxmanagement software although the disk has been „exported“ manually. This is necessary to be able to import this media properly in either on of the other jukeboxes or the backup system. By handling of magazines you have to keep their limitations in mind.

This export/import in the management software has often to done because the software can usually manage only one copy of a disk and not also further ones. But if you cannot export a disk from a defect jukebox you cannot import it to another. The software than assumes the disk as already online and might conclude that a system inconsistency did occur. This leads than to a reinitialization.

Especially for disks which are not yet fully written, e.g. to be able to add data to files on this disk lateron, this „single“ copy handling is necessary. If it is not done this way you can not avoid data inconsistency, corruption and loss.

3.
If you have also jumped over barrier (2) you have to check the average time you need to transfer disks from one jukebox unit to

another one (and later again the transfer back). This has to be compared with the time you usually need to fix the problem with the hardware.

A permanent advantage of splitting the total volume of disks to two or more jukeboxes is that in case of a hardware defect with one system the others are (usually) still available for the applications/users.

But you have to check whether this advantage (and the likelyhood of its occurance) does not being outweighted by other permanent factors in daily operation, e.g. burden sharing or basic problems with the jukeboxserver configuration.

Internal Scaling

If you start a project you have to estimate the storage volume you need today and how it increases over the time.
A basic question is whether to purchase this volume already with the first installation of the jukebox and to fill it up over the time. Alternatively you can start with one smaller jukebox and add additional ones to the server over the time to cover the increased volume.

If you have decided for the first alternative you can - depending on the manufacturer - scale the system internally. The means you get a complete jukebox, but first of all limited to a particular number of slots. Further slots can than be purchased over the time. This enables you to spread the purchase over more than one budget.

Ratio of drives to number of storage disks

The longer the time period is you store data in a jukebox the more likely it is that working off a job affects more than one disk. Therefore you should have enough reading devices which avoids e.g. that a single job spread over 3 disks has not to be served by only two drives. As well it slows down your system if a job occupies (nearly) all reading devices in the system.

If you do not see a demand for a higher number of reading devices the jukebox should offer an easy upgrade to more devices.

Drive Upgrades and Data Condensation

Several suppliers do offer upgrades for their jukeboxes when new drive or technology generations are available.
Often this new generation offers higher capacities than a former one, e.g. from 5.2 to 9.1 GB WORM. If such an upgrade is not being offered the jukebox can not follow further technological developments.

To make sure that you can enjoy all advantages of such future upgrades already today some details have to be considered. The most important aspect is the data condensation.

This effect is easy to understand because an upgrade to higher density drives is nothing else than pooling the data from two to one system.

Example:
Instead to expand a jukebox with 100 platters and 4 drive of 5.2 GB each by and additional and identical system, the capacity is achieved by upgrading the existing system to 9.1 GB drives.

This can cause a critical bottleneck. On the one hand the data access has been spread in the „old“ system over 2 jukeboxes, 200 disks and 8 drives. On the other hand it is statistically unlike that the data on the new, larger disk can be provided in a way that two former accesses to two different disks are now unified to one.

This means that each drive upgrade - which goes along with an increased capacity - must also caused an increased number of drives in the jukebox (for the example given to 8).

Otherwise the access and response times will become less good with the upgraded jukebox over the time.

Caching and Systemperformance
Caching of data is an essential instrument to increase the performance of a jukebox system. But the effect of caching is limited whenever the size of the cache is (softwarewise) limited or the application and the way of use does not allow any kind of senseful caching algorithm.
A typical case for the later example are e.g. insurance- and tax files. After such a file has been created you cannot predict when you will need it the next time. It can be that the customer is calling already tomorrow again to get some information or to make changes. But it also very likely that you need the data the next time many years from now, or, in case of tax files, sometime in the course of the year.

To make the jukebox not a limiting factor in such applications with „chaotic“ access for the user, the jukebox must have enough access ressources - means drives - available. The lower the ratio of reading devices to media in the jukebox is, e.g. 1:25, the faster jobs can be worked off and data can be supplied. But you should not equip the jukebox with more drives than the robotic can manage in your individual way of operation.

Drive specification and system performance
Beside caching of data the specification of the integrated drives plays an important role and is a decisive factor for the system performance - but „faster“ does not always mean „better“.

If mostly a high number of small files are read from the disks in the jukebox, the specification must be in accordance to this operation mode. Especially for CD-ROM and CD-R this rule is valid. Slower drives, e.g. 8x or 12x are faster in spin-up and spin-down than 24x or 32x drives and also the positioning of the head is faster. Therefore these drives are quicker ready to read and transfer data. Only few drives in the market can already read data during spin-up.
As a 24x or 32x drive is in standard applications slower than a 12x speed drive, the use of such drive makes sense when you have to read normally large files (8 MByte and above).

For reading from CD-Rs an additional factor has to be considered. As the bit structure of information is usually not so well defined on CD-Rs than industrial produced CD-ROMs a high rotation speed of the drive makes it more difficult to read the data properly. Vibration, finger prints and dirt adds to the basic problems. Therefore often a fast drive, e.g. 32x has to spin down to slower speed until the reading can start (in most worst and rare cases to single speed). This process costs precious time and reduces the number of jobs the jukebox can work off per hour and therefore the performance of the overall system.

Jukebox with Media Deliverymechanism - how long takes a job?
A jukebox with media deliverymechanism can basically - assuming you have the right management software - reduce the problems of burden sharing if you run physically separated jukebox systems.

The only thing you have to keep in mind is that the delivery of a disk, e.g. from jukebox A to jukebox C, is one job for all involved jukeboxes. Therefore jukebox A delivers no data for any further inquiery until the disk required from jukebox C has been arrived and inserted in a free drive.

Jukebox C and B are also not able to do any other job regarding their other disks and drives until the disk has been arrived in

jukebox A, because they are involved in this job.

In addition the delivery mechanism can only work if in every jukebox always free slots are available.

Otherwiese you have export a disk from a jukebox to another one (or keep them temporarely offline) to be able to import another one from one of the other jukeboxes. This takes additional time and expands the time period to work off the job.

Magazines or not?
A jukebox can be organised in single slots or magazines. If the system uses single slots it always handles single units of disks. Therefore magazines are often assumed to be easier in handling, especially for offline handling. Unfortunately it exist only few applications where such an organisation is useful.

The normal case of daily system administration is that - in a partially or completely equipped jukebox - the application asks for a disk which is currently not in the jukebox but offline.

In this case the system administrator takes the magazine with the required disk and puts it into the jukebox. This can lead to two situations.

a)
In case that the jukebox is not fully equipped with disks one or more magazines have been left free. After having inserted the new magazine and removed the old one both, jukebox and serversoftware have to check the new status of the complete jukebox or at least the new magazine. This means, e.g. for a 50 disk magazine, that all 50 slots have to be checked and in case a disk is found, this has to be inserted into a free drive to be identified. The disk is than also addressed to a particular slot number.

All this has to be done to work off one single job and to keep the system consistent. If the job has than been finalized the jukebox is ready to serve the next request/job.
This way of administration leads often to a non-availability of the system to all users only because there was a request for offline data. As data of offline disks is usually not in the cache the access to the physical disk cannot be avoided.

b)
In case of a fully equipped jukebox basically the same happens as mentioned under a). But in this case a complete magazine with all disks has to be removed from the system just to satisfy one offline disk request.

To properly remove such a magazine you need to make sure that none of the disks of this magazine is currently inserted in one of the drives. If this is the case the disk might not be able to be moved back to the slot because in the exchange magazine this slots carries already a disk. This leads again to a system inconsistency and a reinitialization of the serversoftware.

In a most negative case one media of each magazine is currently used in one of the drives. This makes the exchange of a magazine impossible and therefore also to work off the request for an offline disk.

In addition you have to keep in mind that the request for a single, rarely used disk, causes the non-availability of many other, more frequently used disks.

The advantage of jukeboxes which are organized in slots is that you have to keep only few slots free to be able to insert offline disks from time to time. In case that all slots are already carrying a disk you need only to export a single media via the mailbox to insert one of the offline disks.

A basic disadvantage of magazines is also that, depending on the architecture of the jukebox, you have to open the jukebox to exchange a magazine. Based on the security features of the jukebox against manual handling of disks (door- or magazine-sensor) this leads to a complete or partial reinitialisation of the jukebox. This process can take up to one hour.
Concering the offline handling magazines cannot be opened manually (especially to avoid disks to fall out of the magazine).
If it is necessary to take a single disk out of the magazine this can than be done either carefully by hand or by inserting the magazine into the jukebox and to export the disk than via the mailbox slot.

Service and Maintainance
Regular maintainance for a jukebox, as device with mechanical and electronical components, should go without saying. Especially for corporate critical applications this is a decisive point.

The service of a jukebox - as for any other equipment - should be organized in a short service chain. Every unnecessary part in this chain costs time, money and leads to misunderstandings.
The direct contact to the service- and maintainance department of the manufacturer should be possible by phone, fax or e-mail.

Spare Part Supply
Every jukebox consists of some components who face wear over the time, e.g. drives. At some of these components this wear can be seen or measured, at others not, e.g. electronic elements.

Therefore it is important to have direct access to the manufacturer in any emergency case. If any spare part is not available at your service partner the manufacturer can send it directly to you.

Stable Product Range
A stable product range over years of a manufacturer is an advantage for several reasons.

1.
If you purchase more than one jukebox over the time all those systems are identical. They are therefore better and easier integratable into your EDP environment. In addition it is easier to avoid and solve any kind of problems with server and jukebox.

2.
A stable product range also usually guarantees an upgrade of „old“ jukeboxes with drives of the latest generation, because also new delivered jukeboxes have those drives installed.

Version 1.4; 31.12.2001