



NETWORK CD/DVD SERVERS

WHITE PAPER

AXIS STORPOINT CD E100

Easy Guide to Hard Disk Caching

Created: December 14, 2000
Last updated: December 14, 2000
Rev: 1.1

TABLE OF CONTENTS

<u>INTRODUCTION</u>	<u>3</u>
<u>1 WHAT IS HARD DISK CACHING?</u>	<u>3</u>
<u>2 CACHING APPLICATIONS</u>	<u>4</u>
2.1 CACHING WITH MULTIDISC CD CHANGERS	4
2.2 ACCESS TO MORE CDS THAN DRIVES IN THE TOWER	4
2.3 BOOSTING PERFORMANCE	4
2.4 SECURITY OF CD/DVDS	4
<u>3 CACHING OPTIONS</u>	<u>5</u>
3.1 PERFORMANCE CACHING	5
3.2 ARCHIVE CACHING	5
3.3 MANUAL ARCHIVING	6
3.4 AUTO-ARCHIVING	6
3.5 NO CACHING	6
<u>4 CLEARING THE HARD DISK</u>	<u>6</u>
<u>5 LIMITATIONS</u>	<u>7</u>
<u>6 CONCLUSIONS</u>	<u>7</u>

Introduction

AXIS StorPoint CD E100 is a powerful new CD/DVD server platform that includes the capability for hard disk caching. The following describes the various ways in which the AXIS StorPoint CD E100 handles hard disk caching. Unlike several of its competitors, Axis has a flexible way of handling caching that enhances performance and ease-of-use for the end user and system administrator.

1 What is Hard Disk Caching?

Hard disk caching is the ability for CD-ROM disks to be mirrored (copied) to a hard disk co-located in a CD-ROM tower with CD or DVD drives. The data that has been copied is now accessible directly over the network from the hard disk, rather than from the CD itself.

The primary advantage of this is access speed. When the user clicks on a file icon, the response back to the user is faster than if the file was being pulled from the CD directly. This is very apparent in the case where the CD drive has stopped spinning or "idled down" because it has not been used for a period of time. Data access from a CD requires the drive to spin up and then find the file. This could take a couple of seconds.

In addition, CD disc changers only read from one disc in the changer at a time. If the data a user wants is on a disc that is not in position over the read mechanism, it could take several seconds to position the disc and spin it up to read the file. Hard disk caching avoids this delay by having all the files from all the discs in the changer automatically copied to the hard disk, making them more readily available when needed.

Caching CDs or DVDs to a hard disk aids in the process of disk access, or finding the file a user wants to read. Once the data starts flowing from the AXIS StorPoint CD E100 server onto the network, network traffic and available bandwidth are the primary considerations allowing or preventing the data to arrive at the user's workstation in a timely fashion.

Is it difficult to cache CDs to a hard drive? Not at all. With the AXIS StorPoint CD E100 users can select the various caching options from an intuitive Web interface. After that the caching functions run automatically.

Will the users on the network see any difference? No. Except for the higher performance when reading information from the hard disk, the user will have exactly the same interface to the CD server with or without caching. Where the data is actually coming from is masked from the end user.

2 Caching Applications

Caching CD and DVD data to the local hard disk can yield significant performance improvements in many cases. There are several reasons why a user would want to use hard disk caching:

- Access to data on multi disc CD changers
- Scalability: Access to data from more CDs than the tower has drives
- Needing improved performance accessing multiple files from multiple discs
- Physical security of high value CDs or DVDs

2.1 Caching with Multidisc CD Changers

One of the unique capabilities of the AXIS StorPoint CD E100 server is the ability to browse all discs in a changer without having to cycle through them to read the directories. This is easily done by turning on "Partial" caching through the Web interface. Not until a file is accessed will there be a need for the changer to load the disc into the drive reader. To even further minimize the need of changing discs, the most frequently used discs may be mirrored on to the hard disk in addition to the Partial caching.

2.2 Access to More CDs Than Drives in the Tower

A network-attached CD tower has a limited number of drives installed, typically seven or fourteen. If users have single CD drives, this limits them to seven or fourteen discs. If those users need to have a larger number of discs available on the network and do not want more CD towers, a good option is to archive many CD discs onto the hard disk. The larger the hard disk, the more CDs or DVDs can be copied and simultaneously available to end users. Archive caching is explained in greater detail below.

2.3 Boosting Performance

With a fast CD/DVD drive (e.g. 40x), normally no performance difference will be seen between *reading* from the CD/DVD disc and the hard disk. However when users need to get data from several discs in the tower, the considerably lower seek time of a hard disk will improve performance. A CD-ROM has an access time of 70 to 100 milliseconds whereas a hard disk has an access time generally in the 5 to 10 millisecond range.

A simple test in the Axis lab revealed the following: On a Windows NT 4 workstation with NetBEUI, a 6-MB file was read at 1500 Kbps from both the CD-ROM and the hard disk. In the same environment, 57 small files (1 kB - 166 kB) in the same directory were read at 350 kB/sec from the CD-ROM and 800 kB/sec from the hard disk.

If users have slow CD/DVD drives, the total performance of the CD/DVD solution will be improved by reading data at hard disk speed, even if only one file is accessed, if mirroring is performed.

2.4 Security of CD/DVDs

Many CDs and DVDs represent high value content, sometimes worth thousands of dollars. By Archive caching images of valuable CD/DVD discs, those discs may be ejected and put in a physically secure environment e.g. in a fire-proof safe.

3 Caching Options

There are two main ways of caching in the AXIS StorPoint CD E100 server:

- Performance
- Archiving

3.1 Performance Caching

For Performance caching, two options exist:

- **Partial** - Information about directory structure and files, but no actual data, is being cached. This will speed up directory browsing and finding files, especially with the use of disc changers.
- **Mirror** - A complete image of a CD/DVD disc is being created on the hard disk, and thus increasing performance, especially with multiple users or frequent access of multiple files.

When a new CD/DVD disc is inserted into a drive, the process of storing directory and file information (*Partial*) or a complete image of the disc (*Mirror*) is immediately started.

If a new disc is replacing an existing disc, the first disk is "removed" from the list of available discs and cannot be accessed by the end user. However, as long as there is space available, the cached data will remain on the hard disk, and if the disc is inserted again the cached data becomes immediately available to the end user.

3.2 Archive Caching

Archive caching is the same as the *Mirror* option for Performance caching, but with one important distinction. Archive caching keeps the mirrored data on the hard disk available to the user, even if the CD/DVD disc is ejected. The archiving can be done manually or automatically.

The number of discs that can be archived varies with the size of the discs and the size of the hard disk. For example, an 30 GB hard disk used as a caching drive can store about 45 full length CD-ROMs, either from swapping out discs using single CD drives or from discs loaded into multi disc changers.

Since the AXIS StorPoint only uses the amount of space required for the particular files, and not the arbitrary 650 MB of a typical full-length CD, it uses the hard disk space very efficiently. The number of CDs cached is dependent on how much data is being copied, not on a fixed number of "full capacity" CDs.

Note that automatic performance caching (Partial and Mirror) can be enabled for all discs and drives in addition to setting archiving for individual ones.

3.3 Manual Archiving

Each disc is individually selected via the **ARCHIVE** checkbox in the Web interface. By clicking this, the disc will be copied to the hard disk and be available even if the original disc is removed from the drive.

Archiving a 650-MB CD takes around six minutes. If the CD has already been mirrored, the archiving procedure is done in seconds.

3.4 Auto-archiving

With automatic caching, *Auto-archiving*, the Administrator can identify the CD or DVD drive (or drives) in the tower that will automatically archive the discs that are inserted. Essentially the drive is setup to cache any disc that is placed in this drive. Instead of selecting the individual discs, the Administrator selects a particular drive that will apply caching to all discs in that drive.

3.5 No Caching

By setting the global **CACHE MODE** parameter to **NONE**, automatic caching is disabled by default. CDs or DVDs in the tower will function as if no hard disk was available, with data coming off the CDs and DVDs directly.

4 Clearing the Hard Disk

The hard disk can become full if Archive caching is used, since this is the only caching mode that will allow data to be persistent on the disk after the original disc has been removed. For the other caching options the oldest cached data will be automatically removed if hard disk space is needed for a more recently inserted disc.

The Administrator can be notified if the hard disk is full through:

- **E-mail notification** - By specifying an e-mail address and e-mail server, the AXIS StorPoint CD will automatically send an e-mail to that address to notify of system events such as a full hard disk.
- **Event log** - By looking in the **EVENT LOG**, the Administrator can find information that a disk has become full.
- **Statistics** - By viewing **STATISTICS**, the Administrator can see how much hard disk space is available for caching.

If the hard disk becomes full of archived discs, one or more of these will have to be removed through the Web administration interface before any new data can be cached.

If space is not made available for caching, the CD/DVD contents will be available directly from the CD/DVD disc. Furthermore, it is also possible to clear the entire hard disk cache, whereby all cached data will be deleted.

5 Limitations

While hard disk caching can improve performance in many instances, it is not a panacea for all types of data access from network-attached CDs or DVDs. It depends on the application, number of users and the network structure. Administrators need to look at where they may be creating network "data flow" bottlenecks.

Generally speaking, large numbers of users trying to get data from the same hard disk simultaneously may be slower than having data distributed among several CD-ROM drives. The rule of thumb is that more than 10 users accessing data *simultaneously* makes caching slower. One hundred users could access data on a daily basis and still not meet this criterion. The key word here is "simultaneous" access, which might be rare unless there are a very large numbers of users.

6 Conclusions

For the applications where caching is desired, AXIS StorPoint CD offers ease of use and at the same time maximum flexibility. All that is required to get going with hard disk caching is a hard disk attached to the AXIS StorPoint CD. Performance caching (*Partial*) will then be done automatically for all discs used with the AXIS StorPoint CD solution.

If a disc is removed when Mirroring caching is used, the space on the hard disk is automatically made available for the next disc to be inserted. And if there is no more space to mirror a new disc, it will only be partially cached if possible, otherwise not cached at all. It will, of course, be available to the users on the network, and the Administrator does not have to worry about the server – he or she will get a notifying e-mail. The server is still fully functional at all times, and all discs can be shared.

This should be compared to other available caching solutions where all caching is done manually, both if it's a matter of mirroring a disc or removing a cached disc from the hard disk.

Other solutions only allow for automatic caching and the Administrator has no flexibility in deciding what kind of caching should be done and for what discs and drives. The AXIS StorPoint CD E100 server allows Administrators to fully decide what level of caching to use and where.