CA Disk™ Backup and Restore

Overview Guide

r12.5, Second Edition

CA Technologies Product References

This document references the following CA Technologies products:

- CA Allocate™ DASD Space and Placement (CA Allocate)
- CA Datacom®/DB (CA Datacom/DB)
- CA Disk™ Backup and Restore (CA Disk)

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

Contents

Chapter 1: Overview

CA Disk is a flexible, full-featured hierarchal storage management system. This guide introduces you to CA Disk and provides an overview of the basic concepts.

This chapter briefly touches on basic concepts and features of CA Disk. In order for you to understand the basic concepts and features, the chapter first provides a brief explanation of certain system conventions and requirements. It then briefly describes the system functions, and to some degree, function options. For more information about all of these subjects, see the *Installation Guide*, *User Guide*, and *Systems Guide*.

This section contains the following topics:

Basic Conventions and Requirements

Throughout the operation of CA Disk, certain basic conventions and requirements are observed that help make processing consistent. For more information about these conventions and requirements, see the *User Guide*.

Control Statements

Control statements are used to specify processing that CA Disk performs. For more information about how to create and write CA Disk control statements, see the *User Guide*.

Data Set Names and Patterns

The DSN= parameter is probably the most frequently used of all parameters, and the one where the greatest degree of flexibility is usually required. CA Disk allows either full data set names or data set name patterns to be entered in the DSN= parameter or in any CA Disk table that contains a data set name field.

Data Set Allocation

All CA Disk function allocations use SVC 99 for dynamic allocation of non-VSAM data sets. This allows allocation control products such as CA Allocate, or IBM's DFSMS products to redirect allocations.

VSAM definitions are done using SVC 26, and are also candidates for being redirected by a pooling package.

Due to SVC 99 restrictions, CA Disk does not set the DFSMS ACS &ACSENVIR variable. To keep data set allocations consistent, it is not set for VSAM allocations either.

Special Considerations

The following special considerations are briefly mentioned here. For more information, see the *User Guide*.

Types of Data Sets Processed

CA Disk normally processes data by the moving and managing physical data blocks. In some cases, depending on actions requested by the user, CA Disk can process at the logical record level. CA Disk does *full track reads and writes*, processing as much as one cylinder per I/O.

In most CA Disk functions, however, CA Disk is not a *track-image* processing system.

Data is processed at the physical block level to allow CA Disk to manage data without concern for source and target device dependencies, and to allow maintenance functions to be performed during data movement (for example, compress PDS data sets) and validate data set integrity. An exception to this type of processing is the CA Disk volume-level processing for VBACKUP and VRECOVER.

Track Image Processing for Volume-Level Functions

An exception to standard CA Disk processing at the data set level is the volume-level processing for the VBACKUP and VRECOVER functions. These functions back up and restore data at the volume level and use *track image* processing. CA Disk does not examine separate data sets in these functions, and the data is restored exactly as it was backed up, but for an entire volume, not a single data set.

IMS and DB2 Databases

CA Disk can process VSAM data sets created by IMS and DB2 database management software and maintain their integrity.

DBRC Support

CA Disk can also automatically detect that a data set is either part of a database under DBRC control, or that it is an image-copy data set that DBRC knows about. NOTIFY control statements can be produced such that DBRC can be informed of the CA Disk action that has taken place.

IDMS Support

Data sets within IDMS databases are marked as PS (physical sequential) with fixed format records in which the logical record length and block size are equal, and they can be copied or moved like any other sequential data set.

CA OPS/MVS Support

Unicenter CA OPS/MVS delivers the tools to streamline data center operations and consists of a base product and various optional features. The base product, which is a formal z/OS subsystem, runs in a number of z/OS address spaces.

CA Disk will notify CA OPS/MVS when the Catalog SVC intercepts are installed, removed or refreshed so that any installation dependant event processing can be performed. The events for which notification is performed are STARTING, UP, STOPPING and DOWN. The STARTING and STOPPING events are sent when the installation or removal process is started, and the UP and DOWN events are sent when the process has completed successfully. If the installation process starts and then an error that prevents the successful completion is encountered, a STOPPING and a DOWN event are sent rather than an UP event.

CA Disk requires that the CA OPS/MVS parameter APIACTIVE be set to ON for the API event to be processed.

CA Disk and DFSMS

CA Disk has been developed to provide system managed storage for all data within an enterprise, whether under the control of SMS or not. Using a combination of the traditional rules for managing data, along with the SMS constructs where appropriate, storage administrators can use CA Disk to design sophisticated storage management schemes that will meet the needs of multiple users, expedite routine application production processing, and slow the ever-increasing demand for additional costly hardware storage devices.

The following flowchart reveals the logic flow used when CA Disk processes an SMS managed data set. It will assist you in understanding how the CA Disk SMS sysparms process an SMS managed data set.

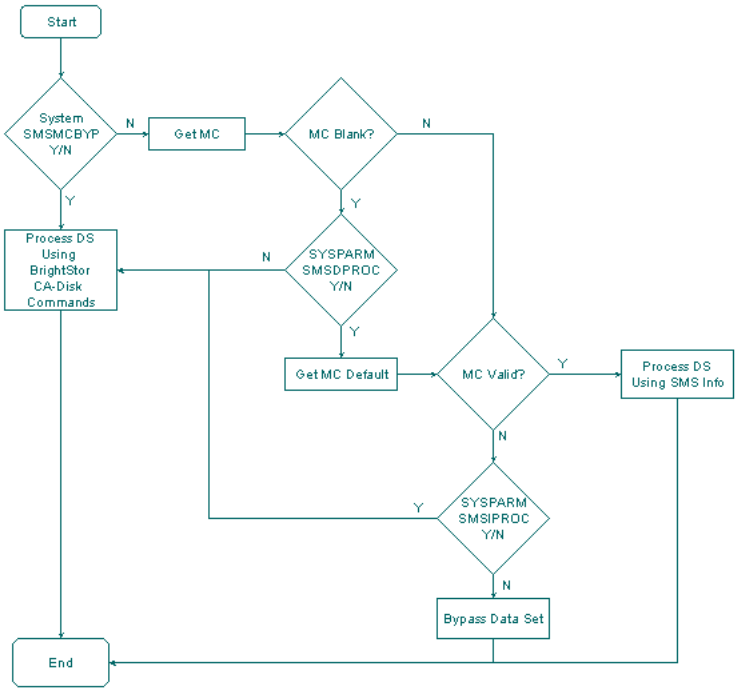

Diagnostics

To aid in the determination of why CA Disk acted in a particular way with SMS-controlled data sets, decision making diagnostic messages are written to a DD statement called SMSPRINT.

These diagnostic messages are also available while in simulate mode, and when included with the DATE= parameter, examples of future processing [if conditions remain unchanged] can be provided.

CA Disk Features

The following sections describe the main features of CA Disk . For more information, see the *User Guide*.

Data Storage Command Language (DSCL)

CA Disk Data Storage Command Language (DSCL) provides a comprehensive palette of commands for efficient, common processing of mixtures of non-VSAM data sets and VSAM clusters.

Within a single execution you can report on, back up, archive, delete, expire, move, copy, and release idle space in non-VSAM data sets and VSAM clusters. You can do this based on scans of VTOCs of disk volumes, scans of ICF, VSAM or explicitly named OS CVOL catalogs, or the results of a catalog superlocate.

You also have a multiple set of criteria to select data sets for, and exclude data sets from, processing as desired. A multiple set of action parameters lets you control exactly what is done to these data sets.

CA Disk preprocesses your command stream to find any *overlapping requests*, that is, those that request processing of the same volumes or catalogs. Such requests are grouped together to allow all needed actions to be performed during a single pass of volume VTOCs, catalogs, or catalog superlocates.

The functions Sequential Migrate to Tape, Merge, Xcopy, Recover, Restore, and PDS Compression continue to use their own unique command sets.

To avoid unexpected execution of partially bad command strings, the DMS procedure takes no action until all commands have been successfully understood, and all sysparm overrides have been validated.

To allow for sorted reports, the DMS proc has separate process, sort and print steps.

Reports

CA Disk reporting facilities provide a comprehensive palette of commands for efficient, common reporting on mixtures of non-VSAM data sets and VSAM clusters.

Within a single execution you can report on non-VSAM data sets, individual VSAM components, ICF catalogs, and VSAM clusters.

You can also define and dynamically change your own reports in a matter of seconds using the ISPF panels. You can then execute them online and review the output, or submit them as batch requests. For more information about online reports, see the *User Guide* and the ISPF help text.

You can do this based on scans of VTOCs disk volumes, scans of ICF or VSAM (but not OS CVOL) catalogs, or the results of a catalog superlocate.

You also have a rich set of criteria to select data sets for and exclude data sets from reporting as desired, and a rich set of reports to let you control exactly what you want to see.

Standard Reports

Standard reports are shipped with CA Disk. They display data fields proven through the years to be highly relevant. These are *fixed format reports* in which the contents and format of the output have been predefined.

Standard reports display data obtained from volume VTOC entries and catalog entries, and are commonly produced from batch execution. When combined with the select and exclude capabilities of the command language (DSCL), you have a high degree of flexibility in obtaining just the information you need.

Standard Reports include:

- **MVDICT, SVDICT, and MAPPER reports—**These are produced for non-VSAM data sets from VTOC information are also produced for VSAM clusters from information in the catalog.
- **Catalog-based reports—**These are formatted identically to the VTOC-based reports; however, the results can differ for multivolume data sets.
- **VTOC-based reports—**These will show only the volumes specified in the selection criteria, whereas the catalog-based reports will show all volumes the data set resides on even if the primary volume is the only volume selected.
- **CLDATA and CIARPT reports—**These will analyze and display information, which is unique to VSAM.
- **ALLOCS, ATTRBS, DISTRS, FREESP, and LAYOUT reports—**These are considered volume-level reports. These reports do not perform data set level processing. By eliminating this overhead, performance for these reports is greatly improved.
- **Online reporting**—This facility gives you the ability to define your own reports. All of the possible reporting fields are displayed for you on the ISPF panels, from which you choose the ones that are needed. You also get to specify the *print position* for each field, sub and grand totals desired, and multiple sorting fields in ascending or descending sequence. When you finish defining the report, you can execute it interactively (online), or submit a batch job to produce it. For more information about these reports, see the *User Guide.* For more information about the function, see the ISPF help text.
- **Multiple reports—**These can be generated during a single scan of a set of volumes. Records for each type of desired report are generated during the scan phase, are sorted into the appropriate sequence, and then printed by the report writer component of CA Disk.

Backup and Archive

CA Disk provides comprehensive, efficient, common processing of mixtures of non-VSAM data sets and VSAM clusters. Within a single execution you can back up, archive, delete, and expire partitioned (PO), partitioned extended (POE), physical sequential (PS), physical sequential extended (PSE), direct access (DA), indexed sequential (ISAM), and OS CVOL catalog non-VSAM data sets. You can also back up, archive, delete and expire VSAM clusters, and back up ICF catalogs.

You can do the backups and archives in either an immediate mode, or in a deferred mode in which the back up or archive is done at a later time.

You can also direct the output to different media types. For instance, by default backups and archives are written to tape. However, by overriding a couple of sysparms, they can easily be directed to disk.

You can do this based on: scans of VTOCs of disk volumes, scans of ICF or VSAM (but not OS CVOL) catalogs, or the results of a catalog superlocate.

You also have a rich set of criteria to select data sets for and exclude data sets from processing as desired, and a rich set of action command parameters to let you control exactly what you want to do.

Incremental Backups

You can perform incremental backups with CA Disk. CA Disk provides incremental backup and recovery at data set, volume, business application and entire complex levels. Incremental backups eliminate the need to perform full-volume backups, without any loss of data or system integrity. But CA Disk goes further. When archived data is needed it is automatically restored.

Basic System-FMS PROC

The objective of adding the FMS PROC to CA Disk was to provide a DSCL-like command language for data set recovery. This PROC provides equivalent functionality available in the original RECOVER PROC, and is equipped with its own palette of commands and parameters.

You also have a rich set of criteria to select data sets for and exclude data sets from processing as desired. CA Disk preprocesses your command stream to find any overlapping requests, that is, those that request processing of the same volumes or catalogs.

To avoid unexpected execution of incomplete or partially bad command strings, the FMS PROC takes no action until all commands have been successfully understood.

To allow for sorted reports, the FMS PROC has separate process, sort, and print steps. Sorted reports are also available in simulate mode.

FMSGEN PROC

This PROC provides the ability of generating multiple restore jobs from a single FMS DSCL request. However, because this PROC uses the combination of a REXX program and a user-defined template, it can produce jobs for a wide variety of purposes, not just restore.

REXX Program—ADSRX010

FMSGEN uses a REXX program, ADSRX010, which combines DSNINDEX records and a user-defined template to generate reports, output files, and JCL statements. Using this technique, you can design a template for virtually any type of output, including the ability to generate multiple restore jobs.

Sample Templates

There are several sample templates located in the INSTALL library that is shipped with CA Disk. These templates utilize variables defined in the REXX program that can be used as is, or customized to suit special needs.

RESTORE and RECOVER

The CA Disk restore function allows you to restore DA, PO, POE, PS, PSE, and ISAM data sets and OS catalogs without the need for preallocation. You can also restore VSAM clusters and ICF catalogs. In most cases, you can even restore to a device type other than that from which the data set was backed up or archived.

- Physical Sequential (PS)
- Direct Access (DA)
- Indexed Sequential (IS)
- Partitioned Organization (PO)
- Access Method—VSAM (AM)

An individual data set or cluster is restored from the RESTORE command. A group of data sets and clusters, an entire volume, or a group of volumes, however, is restored from the RECOVER command.

VBACKUP and VRECOVER

Typical CA Disk processing operates at the data set level; that is, each data set is processed as an individual logical unit, with the data being read and written in logical order. However, in order to provide the capability to perform Stand-Alone Restore, CA Disk has developed the ability to process at the volume level for backup and recover processing; that is, a DASD volume is read and written starting at cylinder 0 track 0, and continuing through the end of the used cylinders.

VBACKUP Function is the CA Disk volume-level backup function invoked through DSCL. The volume-level backup function will create one backup tape (or a set of backup tapes) for each volume. Each tape or set of tapes will contain all of the data, starting at cylinder 0, track 0, and continuing through all used cylinders. CA Disk supports software-based hot backup and an additional hot backup technique provides capabilities of through integration with IBM's Snapshot and Flashcopy features. Hot backup allows file backups to occur while the applications using these files are still running. This additional capability facilitates application-level backups with virtually zero downtime. This functionality enables CA Disk to select and process full backups while the input volume is in OFFLINE status. This feature supports image volumes generated by Hot Backup technologies like EMC Timefinder, IBM Flashcopy, HDS Shadowcopy, and others.

VRECOVER Function is the CA Disk volume-level recovery function. It will overlay the entire volume in extent track order, or a specified range of absolute tracks on the volume. If the data on the volume-level backup is compressed, it will be decompressed during VRECOVER processing using the CA Disk-supplied techniques.

XCOPY

The CA Disk Disaster Recovery Extract Utility (XCOPY) creates ARCHVOLS from CA Disk archive and backup tapes. With its unique design, XCOPY creates ARCHVOLS in a way that will not impact your normal production backup window. These newly created ARCHVOLS are identical in file format to their archive and backup counterparts; they just reside in separate FILES.

XCOPY provides its own set of flexible commands and parameters that allow Storage Administrators to copy only a subset of archive and backup data onto the new ARCHVOLS. Data sets can be selected by fully qualifying their data set names, or by selecting a special character and specifying a pattern of a data set name.

CA Disk data compression, activated via the DCDATACP sysparm, has no effect on XCOPY. Data is processed identically as the input volume, whether it was compressed or not.

Note: Output ARCHVOLS can by compressed using IDRC even if the input ARCHVOL was not.

Finally, XCOPY has an internal mechanism that gives the Storage Administrator the ability to maintain different DSNINDEX and ARCHVOL expiration dates from those on their primary archive and backup copies.

XCOPY is an ideal complement to a Data Center's disaster recovery procedure. While running XCOPY, you will benefit in lower CPU usage compared to that of your normal archive and backup production jobs. XCOPY is designed to copy data from an archive and backup ARCHVOL onto a Disaster Recovery ARCHVOL, eliminating all data set and volume selection overhead.

Basic Operation

XCOPY uses archive and backup ARCHVOLS previously created by CA Disk and the FILES they reside in as input to create subsets of the data. It was developed for Data Centers that were required to create subsets of their backup data to take offsite for disaster recovery. By creating a disaster recovery ARCHVOL from a primary archive and backup ARCHVOL, copying the data can be run outside the backup time window without having to quiesce systems relying on the data.

FILES Maintenance

CA Disk provides a convenient means to manage the index entries for archived data sets and the volumes that contain them.

The archival and backup functions assist you in managing your online storage, but the copies in the archives (and the index entries for them) must also be managed.

CA Disk combines the data sets being archived (or backed up) into a single output data set. This output data set is in sequential format, and therefore can be placed on a standard tape, a tape cartridge, or any disk device. Since archive and backup functions have traditionally directed their output to tape, the output data set has also corresponded one for one to an output volume. This has led us to think of them as archive volumes, with CA Disk keeping track of them in an archvols index. But keep in mind that each is really just a sequential data set that can be on either disk or tape.

Whenever a data set needs to be restored, rapid access into the ARCHVOLS data set is desired, regardless of whether it is on disk or tape. To provide this ability, CA Disk also maintains an index to each of the archived data sets. The index entry contains the exact location of each archived data set within the single ARCHVOLS data set.

Listing Records

Through the use of commands CA Disk can provide you with lists of archived data sets:

LISTV command—ARCHVOL records created for each archive data set can be listed out with the LISTV command. By default, this command will list out the entire ARCHVOLS subfile. However, using any combination of 3 subcommands will allow you to list out a selective number of entries.

LISTD command—DSNINDEX records created for each archive data set can be listed out with the LISTD command. By default, this command will list out the first 90 bytes of the record. However, there are 3 ways in which you can list out a selective portion of the 256-byte records. Index records can also be listed out interactively; this method is particularly useful for Flag Field information.

LISTF command—Index records are created for each archive data set. The first 90 bytes of the 256- byte record can be listed with the LISTD command. Information gathered from the format 1 of the data set at archive time, and other additional fields can be listed with the LISTF command.

LISTS command—the LISTS command is used to list SMS data set information that is stored in the index record if a data set is SMS managed.

IXMAINT Utility General Information

This utility provides the capability to process the entire CA Disk archive data set name index and delete specified records. It can delete both expired and unexpired index entries. Through command parameters, you can restrict it to processing the index entries for only certain specified data sets.

Used as a general maintenance tool, IXMAINT can be run on a daily basis with preset commands to perform periodic cleanup of the index files, or it can be used to accomplish specific tasks on an as-requested basis. Its main objective is to delete index entries for expired data sets.

SMS and IXMAINT

IXMAINT uses SMS Management Class information to determine eligibility of data sets for Index Maintenance.

Archive Index Rebuild Utility

A utility is available that can be used to recreate index records from the archive tapes (or archive data sets on disk) that contain the archived data sets. The support makes use of the control record and format-1 DSCB that are written with the data at the time of archival.

Use of this facility would be appropriate in any of the following situations:

- Restoring data sets at a site other than that at which the data sets were archived.
- FDS logging was not in place, or the FILES could not be fully recovered to a particular point in time.
- The index entries for a given volume have either expired or have been PURGED, but the volume has not been written over or scratched, and one or more data sets must be restored.

IXUPDATE Utility

This utility allows you to use commands that provide a facility to update various fields in the CA Disk archive data set name index (DSNINDEX), and archive volume index (ARCHVOLS).

Making Additional Copies of Archive Tapes

Using the copy utility you can create a copy of an archived tape using CA Disk. CA Disk must keep track of multiple copies of archive tapes, so when it makes a copy of an existing archive tape, the volume serial of the new copy is recorded in the index of the volume record for which the copy is made.

MERGE

The MERGE function provides a convenient means to manage the index entries for archived data sets and the volumes that contain them. It helps you to manage the CA Disk archive data sets that have been put on various disk, and tape cartridge.

CA Disk data compression, activated from the DCDATACP sysparm, has no effect on MERGE. Data is handled whether it's compressed or not.

Merge processing has been designed to provide maximum flexibility in managing the archive volumes on disk, tape and other storage media, and at the same time make the best use of the available resources.

Each merge run consolidates the unexpired data sets from one or more tape volumes (or archive data sets on disk) onto new archive volumes, thus freeing up tape volumes, disk space or both. This reduces the space required to store the CA Disk archives, as well as eliminating disk archive data sets with a large percentage of expired data sets.

After the unexpired data sets have been successfully merged forward, the input data sets on disk are deleted, and input tapes are returned to *scratch status* — disk space is reclaimed and tapes are made available for reuse. All expired CA Disk index entries are dropped from the archives and unexpired entries are updated to reflect their locations on the new tape or disk archives.

Keeping some archives on disk can be desirable, because it maintains easy access to the data through auto-restore processing, without operator intervention.

Disk-to-disk merging provides an efficient means to manage these archives, also without operator intervention.

MOVE and COPY

The Move and Copy function provides the capability to migrate data sets directly from one DASD volume to another by using a single, common command format. This frees the user from unnecessary involvement with different data set organizations, attributes, and device types. Data sets are selected through the use of DSCL selection criteria. DSCL provides a powerful and convenient means to configure all or portions of many volumes. Migration capabilities include CA Disk archive to disk data sets.

Move and Copy can process both VSAM and non-VSAM data sets, whether they are SMS-managed or not. It can be used to move or copy non-SMS data sets to SMS volumes and make them SMS managed. ACS rules are used to determine the data set placement for SMS data sets. In some cases, SMS data sets can be redirected back to their original Storage Class or Storage Group. To ensure this will not happen, review your ACS rules carefully.

Idle Space Release

This function of CA Disk releases unused space from partitioned, physical sequential, and VSAM data sets. Data sets are selected for release through the use of DSCL selection criteria. The DSCL action command RELEASE initiates the release process. The amount of unused space released can be controlled so that a portion of the unused space is retained.

PDS Compression

PDS Compression is provided to:

- Recover the dead space within a partitioned data set.
- REDUCE or INCREASE the size of the data set.
- REDUCE or INCREASE the size of the directory of the data set.
- Change the allocation type of a partitioned data set.

The reduction or increase amounts are based on the contents of the data set. PDS compression can be accomplished implicitly by the designation of volumes to be scanned for eligible datasets, or explicitly by the specification of a data set name.

The PDS Compression facility is provided not only to recover the dead space within a PDS, but also to enlarge or shrink the size of both the data set and its directory, based upon their content. The compression is accomplished by copying each data set to a work file, calculating the reallocation requirements, then copying it back. A message is issued each time the copy-back phase of a compress is begun.

Since a single work file is used for compressing all of the datasets, the size of the work file must be large enough to accommodate the compressed form of the largest data set to be encountered.

The work file can be on any DASD device type (it can span multiple volumes) and given any block size, since CA Disk restructures the data to satisfy the target work file.

PDS compression can be accomplished implicitly by the designation of volumes to be scanned for eligible datasets, or explicitly by the specification of a data set name.

The report produced by PDS Compress contains two lines for each data set processed. The first line presents the status of the data set prior to compression. The second line shows its status after compression.

Sequential Migration to Tape

This CA Disk facility migrates sequential data sets from disk devices to tape. The data sets are written and recataloged as standard labeled tape files, multiple files per tape. Any job that previously accessed its disk data set from the catalog now calls for a tape mount and continues running as before.

Sequential data sets can be moved from online storage to tape and recataloged by either explicit or implicit reference. Each DASD data set migrated is written as a standard labeled tape data set and cataloged, enabling any job that formerly referenced the data set through the catalog to continue to do so without change.

DASD Billing

This CA Disk facility provides billing capabilities for the space occupied by non-VSAM data sets and VSAM clusters on DASD. Periodic VTOC scans are made to determine the space occupied by each individual data set. Assigned billing rates are then applied and accumulated. The VTOC scans can be run as frequently as desired. Each run's billing amount is based on current space occupancy assumed to have existed over the period since the previous scan.

The CA Disk DASD Space Billing function accumulates a billing charge against each data set based on two criteria:

- The SPACE OCCUPIED by (allocated to) each data set.
- The NUMBER OF DAYS it occupies that space.

The product of these two values provides a billing quantity in either TRACKDAYS or KILOBYTE-DAYS. The user must supply the billing rate for either one track-day or one kilobyte-day of occupancy. A simple multiplication of the billing quantity times the billing rate yields the billing charge for each data set.

For example, a two-track data set residing on DASD for 50 days provides a 100 track-day billing quantity. If the billing rate is one cent per track per day, the charge for this period is one dollar.

To find the space occupied by each data set (item 1 above), periodic VTOC scans are made and space allocations calculated.

There are two methods available to determine the days of occupancy for data sets. They are referred to as MODE=ELAPSED and MODE=IMMEDIATE.

ISPF

The ISPF support of CA Disk provides an interface from the IBM System Productivity Facility, Dialog Management Services, to the CA Disk product. CA Disk uses ISPF dialog management service panels to give you an easy-to-use interface in an interactive environment.

The CA Disk dialog management functions either guide the user through the process of generating JCL to execute batch CA Disk facilities, or interface directly to the CA Disk foreground applications. No JCL is involved for foreground processing functions, as all requests are performed immediately in the TSO user's region.

All CA Disk dialog manager panels have associated HELP text. If at any time you need information on how to proceed or what to do, the <HELP> PF key can be pressed. The HELP information can be browsed until you are ready to continue. The <END> key will return you to the panel originally being processed.

TSO Command Processor

Specialized support has been provided for the TSO environment in the form of command processors. To minimize the amount of familiarization necessary for TSO users to master the commands, the parse service has been used in the command processors for command and parameter formatting and validation.

The commands for these TSO command processors are described in detail in the *User Guide*. TSO HELP text is available for all command processors.

The following TSO commands are available:

- **DARCHIVE**—Deferred Archive

  Use of this command results in the placement of a record containing all of the specified parameters into the ARCHCMDS file. All pending command records are processed and appropriately flagged at the next execution of the program that disposes of the deferred archive requests.

- **DRESTORE**—Deferred Restore

  This command causes a record containing the specified parameters to be placed in the RESTCMDS file. These records are processed during the next batch execution of the program that disposes of the deferred restore requests.

- **RESTORE**—Immediate (Dynamic) Restore

  Data sets can be dynamically restored in the TSO user region by way of the RESTORE command. The command processor, if needed, dynamically allocates a tape drive, for the period of time necessary to restore the requested data set.

- **LISTDMS**—Listing the Archive Index

  The LISTDMS command selectively lists the contents of the archive index.

- **LISTREQ**—List Status of Deferred Requests

  A TSO command processor is available for displaying the status of requests for deferred archiving and restoring of data sets. The display will include the requests submitted in both TSO and batch environments.

- **DERASE**—Erase a Deferred Request

  The DERASE command deletes deferred archive and deferred restore requests. The DERASE command can also be used in batch runs.

Examples

To help you understand and utilize CA Disk we have provided you with a large number of examples:

- REPORT
- VREPORT
- ARCHIVE
- RESTORE
- Volume-Level BACKUP
- Volume-Level RECOVER
- RECOVER - FMS PROC
- BACKUP
- Incremental BACKUP
- LISTD/LISTV
- IXMAINT
- MERGE
- REBUILD
- IXUPDATE
- Idle Space Release
- Sequential Migrate
- PDS Compression
- Move/Copy
- DASD BILLING
- XCOPY

See the *User Guide* for more details.

CA Datacom

CA Disk is now one of several CA products that are tightly integrated with, and make use of the database capabilities in CA Datacom.

CA Datacom provides a state-of-the-art, high performance database for CA Disk. That is, as a licensed customer of CA Disk, you can use CA Datacom to manage your CA Disk FILES data set.

For CA Disk users, this includes the following abilities:

- By now utilizing the CA Datacom Database, CA Disk will no longer use EXCP I/O to access the FILES dataset, it will use the high performance multi-tasking and multi-threading engine provided by the new database.
- Conversion from one or more FILES dataset to CA Datacom Databases is done without downtime to CA Disk processes. Several FILES datasets can be consolidated into a single database or kept separated. In addition, a mixed implementation with data on FILES dataset and CA Datacom Databases is supported.
- Switching to the new CA Datacom Database will save the once tedious time for reorganizations. The new process is *self-reorganizing* thus many times easier to use.
- Historically CA Disk required the operating system to use the Data Set and Device enqueing, thus single threading the work load. CA Disk will now process using the record level locking provided by the new CA Datacom Database.

CA Vantage GMI

CA Disk is now one of several CA products that make use of the same Graphical Management Interface (GMI) known as CA Vantage GMI. That is, as a licensed customer of CA Disk, you can use CA Vantage GMI and its Windows, Java, and 3270 clients to view and analyze your storage-related data. For CA Disk users, this includes the following abilities:

- View and analyze your FILES online
- View and analyze the Status Report and Subfiles online
- Enhance DMSAR auto-restores with the Auto Restore Manager
- It provides CA Disk activity analysis
- You can Filter/Sort any field
- You can ZOOM from DSNINDEX to ARCHVOLS
- You can ZOOM from DSNINDEX to SMS Class
- It provides Historical Restore analysis
- You can perform queries from data set lists
- You can perform Auto-Restore consolidation
- It provides Thrashing analysis
- It can tell you what tapes contain data sets backed up last night
- It can advise you what data sets will expire in tonight's IXMAINT
- You can find out what data sets have been restored nn times
- You can find out when the data sets were restored
- You can find out how much space is needed to restore a user's archived data
- You can find out if any data sets been archived but not re-cataloged
- You can filter on user data in the DSNINDEX
- You can find out what's unprocessed in the deferred queue

This interface will increase storage management productivity by allowing personnel to gather information and identify problems fast because all the required information is just a click away.

For more information on CA Vantage GMI, see the CA Vantage documentation.