CA Disk™ Backup and Restore

Installation Guide

r12.5, Second Edition

Third Edition

CA Product References

This document references the following CA Technologies products:

- CA 1® Tape Management (CA 1)
- CA Allocate™ DASD Space and Placement (CA Allocate)
- CA ACF2™ for z/OS (CA ACF2)
- CA Auditor for z/OS
- CA Datacom®/DB (CA Datacom/DB)
- CA Disk™ Backup and Restore (CA Disk)
- CA Graphical Management Interface (CA GMI)
- CA Mainframe Software Manager (CA MSM)
- CA Service Desk
- CA Tape Encryption
- CA TLMS® Tape Management (CA TLMS)
- CA Top Secret® for z/OS (CA Top Secret)
- CA Vantage™ Storage Resource Manager (CA Vantage SRM)

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

This guide describes how to install and implement CA Disk.

This section contains the following topics:

Audience

A systems programmer, who installs CA Disk and does the basic configuration, typically performs the steps in this guide.

The systems programmer should be an experienced mainframe technician with SMP/E installation process experience. Knowledge of your mainframe storage systems, storage re lated software, and security setups are essential for installing and configuring CA Disk.

How the Installation Process Works

CA Technologies has standardized product installations across all mainframe products. Installation uses the following process:

- Acquisition—Transports the software to your z/OS system.
- Installation using SMP/E—Optionally creates a CSI environment and runs the RECEIVE, APPLY, and ACCEPT steps. The software is untailored.
- Deployment—Copies the target libraries to another system or LPAR.
- Configuration—Creates customized load modules, bringing the software to an executable state.

CA MSM provides a web-based interface to make the standardized installation process easier. Using CA MSM, someone with limited knowledge of JCL and SMP/E can install a product.

As a best practice, we recommend that you install mainframe products and maintenance using CA MSM.

Note: If you do not have CA MSM, you can download it from the Download Center at http://ca.com/support. Follow the installation instructions in the CA Chorus Software Manager documentation bookshelf on the CA Chorus Software Manager product page.

You can also complete the standardized installation process manually using pax files that are downloaded from http://ca.com/suppor http://ca.com/suppor, http://www.t or a product DVD.

To install your product, do the following tasks:

1. Prepare for the installation by confirming that your site meets all installation requirements (see page 13).
2. Acquire the product using one of the following methods:
   - CA MSM
   - Pax-Enhanced Electronic Software Delivery (ESD)
   - Order a DVD. To do so, contact your account manager or a CA Technologies Support representative.
3. Install the product based on your acquisition method.

   Note: If a CA Recommended Service (CA RS) package is published for your product, install it before continuing with deployment.
4. Install the CA Common Services using the pax files that contain the CA Common Services you need at your site.

   Install all CA Common Services contained in the Required CA Common Service bundle.

5. Apply maintenance, if applicable.
6. Deploy your target libraries.
7. Configure your product.

Chapter 2: Preparing for Installation

This section describes what you need to know and do before you install the product.

This section contains the following topics:

Hardware Requirements

CA Disk requires IBM (or compatible) processors running under OS/390 operating systems.

Software Requirements

CA Disk requires any IBM supported release of z/OS. Other operating systems are formally supported only if they are fully compatible with MVS.

CA Common Services Requirements

CA Disk requires CA Common Services for z/OS release 2.2 Service Pack 1 or higher. The following CA Common Services are used with CA Disk:

- CAIRIM
- CA LMP

CAISDI

Storage Requirements

Disk space estimates for all libraries are documented in the program directory file on the installation tape. Ensure enough disk space to load and install CA Disk. Also, decide on a naming convention for the CA Disk data sets.

z/OS System Resources Requirements

Ensure that the following z/OS System resources requirements are met before starting the installation:

- APF-Authorized Libraries
- User SVC
- Auto-Restore
- JES Awareness
- Linklist Library
- LPA List Library
- Help Library
- Parameter Library

APF-Authorized Libraries

Establish the following APF-authorized libraries for CA Disk:

- Your CA Disk load library
- Your CA Disk LPA List Library. This library is named in LPALST*xx* or on an LPA statement in PROG*xx,* therefore it does not need to be added to IEAAPF*xx*.

Establishing the APF-authorized libraries in advance permits you to test executions immediately after downloading the distribution tape. Otherwise, you have to wait for an initial program load (IPL) to authorize the libraries, or you can specify the following operator command to authorize the libraries dynamically:

T PR0G=*nn*

User SVC

Make an entry for the CA Disk SVC in the IEASVC*xx* member of your system PARMLIB. It should be a type 3 or type 4 SVC, enabled for interrupts. Designating as either restricted or non-restricted is optional.

Auto-Restore

Ensure that approximately 41 KB of Extended Common Service Area (ECSA) is available in your operating system if you are planning to install the CA Disk auto-restore function. For more information about installing the CA Disk auto-restore catalog management hook, see The Two Auto Methods.

JES Awareness

Auto-restores are automatically bypassed for all jobs starting with JES*nnnnn*. This is done to automatically exclude JES2 and JES3 from invoking an auto-restore. The system hangs if JES2 or JES3 invokes an auto-restore before starting JES.

Linklist Library

Add the CCUWLINK library to the LNKLST*xx* member of your system PARMLIB. If you plan to copy the CCUWLINK library to an existing linklist library, ensure that approximately 45 KB of disk storage are available in that library.

If using CA Disk and CA FAVER for MVS VSAM and CA FAVER EXPORT processing is utilizing the dynamic reload function of the CA Disk interface; ensure that the CA FAVER and CA Disk products are maintained in the LNKLST together or in a concatenated STEPLIB within the CA FAVER executions.

If you do not do this you could recieve the following message:

CSV003I REQUESTED MODULE ADSMI002 NOT FOUND

The above message is accompanied by an S806 ABEND causing the CA FAVER Export to fail.

No changes are required if you are *not* using the dynamic CA FAVER reload function with CA Disk during EXPORT.

LPA List Library

Add the CCUWLPA library to the LPALST*xx* member of your system PARMLIB. If you plan to copy the CCUWLPA library to an existing LPA list library, ensure that approximately 8 KB of disk storage are available in that library. The CCUWLPA library contains the CA Disk SVC and CA-Examine PDM module.

Help Library

If you are installing the TSO interface, either concatenate the CCUWHELP library to the //SYSHELP DD in your TSO logon JCL, or if you plan to copy the CCUWHELP library into your system TSO help library, ensure that approximately 6 KB of disk storage are available for six members that must reside in that library. These members are required for the TSO interface.

Parameter Library

Save a copy of your PARMLIB; so that you can copy your user-defined members after the CA Disk installation process is complete.

Concurrent Releases

You can install this release of CA Disk and continue to use an older release for your production environment. If you plan to continue to run a previous release, consider the following points:

- When installing into an existing SMP/E environment, this installation deletes previous releases. The libraries and DDDEFs from the previous release may be deleted after this release has been accepted in SMP/E.
- If you acquired your product from tape or with Pax-Enhanced ESD, select different target and distribution zones for your new release from where your current release is installed. The new zones use different libraries than your current release.

  Note: CA MSM installs into a new CSI by default.
- Define DDDEF entries in your new zones to point SMP/E to the proper libraries for installation. Ensure that they point to the new release libraries.

Chapter 3: Installing Your Product Using CA MSM

Use the procedures in this section to manage your product using CA MSM. Managing includes acquiring, installing, maintaining, and deploying products, setting system registries, and managing your CSIs. These procedures assume that you have already installed and configured CA MSM.

Note: If you do not have CA MSM, you can download it from the Download Center at the CA Support Online website. Follow the installation instructions in the CA Mainframe Software Manager documentation bookshelf on the CA Mainframe Software Manager product page.

When you have completed the procedures in this section, go to Configuring Your Product.

This section contains the following topics:

Note: The following procedures are for CA MSM r3. If you are using CA MSM r2, see the *CA Mainframe Software Manager r2 Product Guide*.

CA MSM Documentation

This chapter includes the required procedures to install your product using CA MSM. If you want to learn more about the full functionality of CA MSM, see the CA Mainframe Software Manager bookshelf on the CA MSM product page on https://support.ca.com/.

Note: To ensure you have the latest version of these procedures, go to the CA Mainframe Software Manager product page on the CA Support Online website, click the Bookshelves link, and select the bookshelf that corresponds to the version of CA MSM that you are using.

Getting Started Using CA MSM

This section includes information about how to get started using CA MSM.

How to Use CA MSM: Scenarios

In the scenarios that follow, imagine that your organization recently deployed CA MSM to simplify the installation of CA Technologies products and unify their management. You have also licensed a new CA Technologies product. In addition, you have a number of existing CSIs from previously installed products.

- The first scenario shows how you can use CA MSM to acquire the product.
- The second scenario shows how you can use CA MSM to install the product.
- The third scenario shows how you can use CA MSM to maintain products already installed in your environment.
- The fourth scenario shows how you can use CA MSM to deploy the product to your target systems.

How to Acquire a Product

The *Product Acquisition Service (PAS)* facilitates the acquisition of mainframe products and the service for those products, such as program temporary fixes (PTFs). PAS retrieves information about products to which your site is entitled. Then it records these entitlements in a software inventory that is maintained on your driving system.

You can use the PAS component of CA MSM to acquire a CA Technologies product.

Follow these steps:

1. Set up a CA Support Online account.

   To use CA MSM to acquire or download a product, you must have a CA Support Online account. If you do not have an account, you can create one on the CA Support Online website.

2. Determine the CA MSM URL for your site.

   To access CA MSM (see page 27), you require its URL. You can get the URL from your site's CA MSM administrator and log in using your z/OS credentials. When you log in for the first time, you are prompted to create a CA MSM account with your credentials for the CA Support Online website. This account enables you to download product packages.

3. Log in to CA MSM and go to the Software Catalog page to locate the product that you want to manage.

   After you log in to CA MSM, you can see the products to which your organization is entitled on the Software Catalog tab.

   If you cannot find the product you want to acquire, update the catalog (see page 28). CA MSM refreshes the catalog through the CA Support Online website using the site IDs associated with your credentials for the CA Support Online website.

4. Download the product installation packages (see page 29).

   After you find your product in the catalog, you can download the product installation packages (see page 29).

   CA MSM downloads (acquires) the packages (including any maintenance packages) from the CA FTP site.

After the acquisition process completes, the product is ready for you to install or maintain.

How to Deploy a Product

The *Software Deployment Service (SDS)* facilitates the mainframe product deployment from the software inventory of the driving system to the target system. This facilitation includes deploying installed products that are policy-driven with a set of appropriate transport mechanisms across a known topology.

You can use the SDS component of CA MSM to deploy a CA Technologies product that you have already acquired and installed.

Follow these steps:

1. Set up the system registry:
   a. Determine the systems you have at your enterprise.
   b. Set up remote credentials (see page 77) for those systems.
   c. Set up the target systems (Non-Sysplex (see page 55), Sysplex or Monoplex (see page 56), Shared DASD Cluster (see page 57), and Staging (see page 58)), and validate them.
   d. Add FTP (see page 67) information, including data destination information, to each system registry entry.
2. Set up methodologies (see page 105).

3. Create the deployment, which includes completing each step in the New Deployment wizard.

   After creating the deployment, you can save it and change it later by adding and editing systems (see page 122), products (see page 96), custom data sets (see page 98), and methodologies (see page 105), or you can deploy directly from the wizard.

   Note: If you must deploy other products to the previously defined systems using the same methodologies, you must create a separate deployment.

4. Deploy the product, which includes taking a snapshot, transmitting to target, and deploying (unpacking) to your mainframe environment.

After the deployment process completes, the product is ready for you to configure. You may have to perform other steps manually outside of CA MSM before beginning the configuration process.

System Registration

You must add and then validate each system in the enterprise that you are deploying to the CA MSM system registry. You can only send a deployment to a validated system. This process is called registering your system and applies to each system in your enterprise. For example, if you have five systems at your enterprise, you must perform this procedure five times.

Note: After a system is registered, you do not need to register it again, but you can update the data in the different registration fields and re-register your system.

The system registration process contains the following high-level steps:

1. Set up your remote credentials.

   This is where you provide a user ID and password to the remote target system where the deployment will copy the installed software to. Remote credentials are validated during the deployment process. You will need the following information:

   - Remote user ID
   - Remote system name
   - Password
   - Authenticated authorization before creating a remote credential.

   Your system administrator can help you with setting up your remote credentials.

2. Set up your system registry.

   The CA MSM system registry is a CA MSM database, where CA MSM records information about your systems that you want to participate in the deployment process. There is one entry for each system that you register. Each entry consists of three categories of information: general, FTP locations, and data destinations.

   Each system registry entry is one of four different system types. Two reflect real systems, and two are CA MSM-defined constructs used to facilitate the deployment process. The two real system types are Non-Sysplex System and Sysplex Systems. The two CA MSM-defined system types are Shared DASD Clusters and Staging Systems.

   Non-Sysplex Systems

   Specifies a stand-alone z/OS system that is not part of a sysplex system.

   Note: During system validation, if it is found to be part of a sysplex, you will be notified and then given the opportunity to have that system automatically be added to the sysplex that it is a member of. This may cause the creation of a new sysplex system. If you do not select the automatic movement to the proper sysplex, this system will be validated and cannot be deployed.

   Sysplex or Monoplex Systems

   Specifies a *Sysplex* (SYStem comPLEX), which is the IBM mainframe system complex that is a single logic system running on one or more physical systems. Each of the physical systems that make up a Sysplex is often referred to as a *member* system.

   A *Monoplex system* is a sysplex system with only one system assigned.

   Note: Monoplexes are stored in the Sysplex registry tree but with the name of the Monoplex System and not the Monoplex Sysplex name. For example, a system XX16 defined as a Monoplex, with a Sysplex name of LOCAL. It will be depicted in the System Registry as a Sysplex with the name of XX16. This sysplex will contain one system: XX16.

   This system type can help you if you have Monoplexes with the same Sysplex name (for example: LOCAL). Instead of showing multiple LOCAL Sysplex entries that would need to be expanded to select the correct Monoplex system, the CA MSM System Registry shows the actual Monoplex System name at the top-level Sysplex Name.

   Shared DASD Clusters

   Specifies a *Shared DASD Clusters* system, which defines a set of systems that share DASD and it can be composed of Sysplex systems, Non-Sysplex systems, or both. A Staging system cannot be part of a Shared DASD Cluster.

Staging Systems

Specifies a *Staging system*, which is an SDS term that defines a virtual system. A Staging system deploys the deployment to the computer where the CA MSM driving system is located. To use a Staging system, the CA MSM driving system must be registered in the CA MSM System Registry.

Note: A Staging system can be useful in testing your deployments and learning deployment in general. It can also be used if your target systems are outside a firewall. For example, deploy to a Staging system and then manually copy the deployment to tape.

3. Define the FTP location information for every system.

   FTP locations are used to retrieve the results of the deployment on the target system (regardless if the deployment was transmitted through FTP or using Shared DASD). They are also used if you are moving your deployments through FTP.

   To define the FTP location, provide the following:

   URI

   Specifies the host system name.

   Port Number

   Specifies the port number.

   Default: 21.

   Directory Path

   Specifies the landing directory, which is the location that the data is temporarily placed in during a deployment.

4. Define a data destination for every system.

   The data destination is how you tell CA MSM which technique to use to transport the deployment data to the remote system. The following choices are available:

   FTP

   When FTP is selected as the transport mechanism, the deployment data is shipped to the target system through FTP. It is temporarily placed on the target system at the landing directory specified in the FTP Location information section of the System Registry.

   Shared DASD

   When you specify shared DASD, CA MSM uses a virtual transport technique. That is, it does not actually copy the data from one system to the other. Because the two systems share DASD, there is no need to do this. All of the deployment data is kept in USS file systems managed by CA MSM.

   Even though the DASD is shared, the remote system may not be able to find the deployment data in the USS file system. Therefore, CA MSM temporarily unmounts the file system from the CA MSM driving system and mounts it in read-only mode on the remote system.

   For CA MSM to determine where to mount the file system on the remote system, you must specify a mount point location in the data destination. In addition, you can provide allocation information for the creation of the deployment file system, so that when the file system is created on the CA MSM driving system, it will be on the DASD that is shared.

   Data destinations are assigned to Non-Sysplex and Sysplex systems, and Shared DASD Clusters. Data destinations are named objects, and may be assigned to multiple entities in the system registry and have their own independent maintenance dialogs.

   The remote allocation information is used by the deployment process on the remote system, letting you control where the deployed software is placed. By specifying the GIMUNZIP volser, CA MSM adds a volume= parameter to the GIMUNZIP instructions on the remote system. The list of zFS VOLSERs is needed only if both of the following occur:

   - The software you are deploying contains USS parts.
   - You select a container copy option during the deployment process.

   Note: After you have created your systems, you will need to validate them.

5. Register each system by validating that it exists.

   Note: You should validate your Non-Sysplex Systems first, and then your Sysplex or Shared Cluster Systems.

   You start the validation process when you select the Validate button in the Actions drop-down list for a Sysplex System, Non-Sysplex System, and Shared DASD Cluster on that system's System Registry Page. This starts a background process using the CCI validation services to validate this system.

   Note: Staging Systems are not validated. However, you will need to create and validate a system registry entry for the CA MSM driving system if you are going to utilize Staging systems.

   Note: If the validation is in error, review the message log, update your system registry-entered information, and validate again.

   You are now ready to deploy your products.

Deploying Products

After you install software using CA MSM, you still need to deploy it. You can use the deployment wizard to guide you through the deployment process. In the wizard, you can deploy one product at a time. You can also save a deployment at any step in the wizard, and then manually edit and deploy later.

Note: You must have at least one product, one system, and one methodology defined and selected to deploy.

You must complete the following steps in the Deployment wizard before you deploy:

Deployment Name and Description

Enter the deployment name and description using the wizard. The name must be a meaningful deployment name.

Note: Each deployment name must be unique. Deployment names are not case-sensitive. For example DEPL1 and depl1 are the same deployment name.

We recommend that you enter an accurate and brief description of this deployment.

CSI Selection

Select a CSI. A CSI is created for the installed product as part of the installation process.

Product Selection

Displays the products that are installed in the CSI you selected.

Custom Data Set

Custom data sets let you add other data sets along with the deployment. They contain either a z/OS data set or USS paths.

- For a z/OS data set, you need to provide a data set name that is the actual existing z/OS data set and a mask that names the data set on the target system. This mask may be set up using symbolic qualifiers (see page 109) and must be available to CA MSM. During the deployment process, the custom data set is accessed and copied to the target system the same way a target library is accessed and copied.
- For USS paths, you need to provide a local path, a remote path which may be set up using symbolic qualifiers (see page 109) and type of copy. Type of copy can be either a container copy or a file-by-file copy.

You can add a custom data set (see page 99).

Methodology

Methodology is the process by which data sets are named on the target system. A methodology provides the *how* of a deployment, that is, what you want to call your data sets. It is the named objects with a description that are assigned to an individual deployment.

To create a methodology (see page 106), specify the following:

Data set name mask

Lets you choose symbolic variables that get resolved during deployment.

Disposition of the target data sets

If you select Create, ensure that the target data sets do not exist, otherwise, the deployment fails.

If you select *Create or Replace* and the target data sets do not exist, they will be created. If the target data sets exist, *Create or Replace* indicates that data in the existing data set, file, or directory will be replaced, as follows:

Partitioned data set

Create or Replace indicates that existing members in a partitioned data set will be replaced by members with the same name from the source file. Any currently existing member that is not in the source file will remain in the PDS. Any member from the source that does not already exist in the target PDS will be added to the target PDS.

The amount of free space in the PDS should be sufficient to hold the additional content, because no automatic compress is performed.

Directory in a UNIX file system

Create or Replace indicates files in a directory will be replaced by files with same name from the source. Any currently existing directory in a UNIX file system that is not in the source will remain in the UNIX file system.

Sequential data set or a file in the UNIX file system

Create or Replace indicates the existing data set or file and its attributes will be replaced with the data from the source file.

For a VSAM data set (cluster)

Create or Replace indicates that an existing VSAM cluster should be populated with the data from the source file. The existing VSAM cluster must be of the same type as the source cluster (ESDS, KSDS, LDS, or RRDS). In addition, the existing VSAM cluster must have characteristics that are compatible with the source cluster (such as, record size, key size, and key offset). Replace does not verify the compatibility of these characteristics!

Note: You can replace the contents of an existing cluster using the IDCAMS ALTER command to alter the cluster to a reusable state. You must do this before the data from the VSAM source is copied into the cluster using an IDCAMS REPRO command. The REPRO command will use both the REPLACE and REUSE operands, and after you use it, the cluster is altered back to a non-reusable state if that was its state to begin with.

System Selection

Select the system for this deployment.

Preview

Preview identifies the deployment by name and briefly states the products, systems, means of transport, target libraries including source, target and resolution, as well as SMP/E environment and snapshot information. It shows the translated symbolic qualifiers.

Use this option to review your deployment before deploying.

Deploy

Deploy combines the snapshot, transmit, and deploy action into one action. Deploy enables you to copy your CA MSM-installed software onto systems across your enterprise. For example, you can send one or many products to one or many systems. Deploy can send the software by copying it to a shared DASD or through FTP.

Summary

After your products have successfully deployed, you can review your deployment summary and then confirm your deployment. You can also delete a completed deployment.

Confirm

Confirms that the deployment is complete. A deployment is not completed until it is confirmed. After it is confirmed, the deployment moves to the Confirmed deployment list.

How to Maintain Existing Products

If you have existing CSIs, you can bring those CSIs into CA MSM so that you can maintain all your installed products in a unified way from a single web-based interface.

You can use the PAS and SIS to maintain a CA Technologies product.

Follow these steps:

1. Migrate the CSI to CA MSM to maintain an existing CSI in CA MSM.

   During the migration, CA MSM stores information about the CSI in the database.

2. Download the latest maintenance (see page 41) for the installed product releases from the Software Catalog tab.

   If you cannot find a release (for example, because the release is old), you can add the release to the catalog manually and then update the release to download the maintenance (see page 42).

3. Apply the maintenance (see page 45).

Note: You can also install maintenance to a particular CSI from the SMP/E Environments tab.

After the maintenance process completes, the product is ready for you to deploy. You may have to perform other steps manually outside of CA MSM before beginning the deployment process.

Access CA MSM Using the Web-Based Interface

You access CA MSM using the web-based interface. Obtain the URL of CA MSM from the CA MSM administrator.

Follow these steps:

1. Start your web browser, and enter the access URL.

   The login page appears.

   Note: If the Notice and Consent Banner appears, read and confirm the provided information.

2. Enter your z/OS login user name and password, and click the Log in button.

   The initial page appears. If you log in for the first time, you are prompted to define your account on the CA Support Online website.

   Note: For more information about the interface, click the online help link at the top right corner of the page.

3. Click New.

   You are prompted for the credentials to use on the CA Support Online website.

   Important! The account to which the credentials apply *must* have the Product Display Options set to BRANDED PRODUCTS. You can view and update your account preferences by logging in to the CA Support Online website and clicking My Account. You need the correct setting to use CA MSM to download product information and packages.

4. Specify the credentials, click OK, and then click Next.

   You are prompted to review your user settings.

   Note: These settings are available on the User Settings page.

5. Change the settings or keep the defaults, and then click Finish.

   A dialog shows the progress of the configuration task. You can click Show Results to view the details of the actions in a finished task.

Important! If your site uses proxies, review your proxy credentials on the User Settings, Software Acquisition page.

Acquiring Products

This section includes information about how to use CA MSM to acquire products.

Update Software Catalog

Initially, the CA MSM software catalog is empty. To see available products at your site, update the catalog. As new releases become available, update the catalog again to refresh the information. The available products are updated using the site ID associated with your credentials on the CA Support Online website.

If you update the catalog tree and some changes are missing, check your user settings on the CA Support Online website.

Follow these steps:

1. Click the Software Catalog tab.

   Note: The information on the Software Status tab for HIPERs and new maintenance is based on the current information in your software catalog. We recommend that you update the catalog on a daily or weekly basis to keep it current.

2. Click the Update Catalog Tree link in the Actions section at the left.

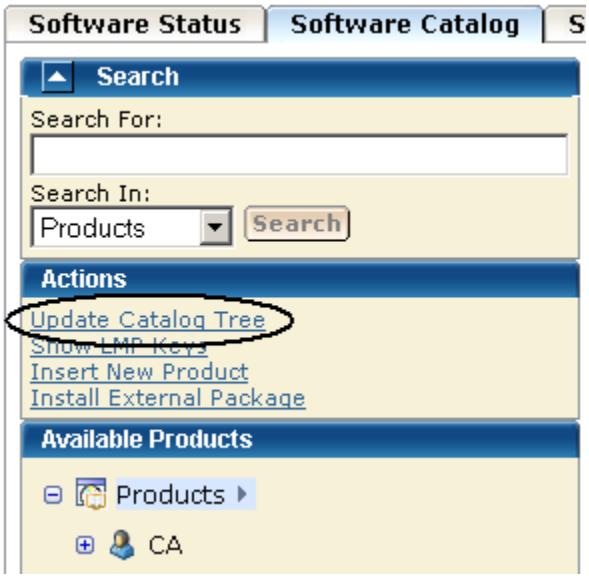

   You are prompted to confirm the update.

3. Click OK.

   A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

   Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

Download Product Installation Package

You can download product packages through the Software Catalog tab. The Update Catalog action retrieves information about the products for your site.

Follow these steps:

1. Verify that your CA MSM login user name is associated with a registered user of the CA Support Online website on the Software Acquisition Settings page.

   CA MSM uses the credentials to access the CA Support Online website.

2. Locate and select the product you want to download by using the Search For field or expanding the Available Products tree at the left.

   The product releases are listed.

   Note: If the product does not appear on the product tree, click the Update Catalog Tree link in the Actions section at the left. The available products are updated using the site ID associated with your credentials for the CA Support Online website. If you update the catalog tree and some changes are missing, check your user settings on the CA Support Online website.

3. Click Update Catalog Release in the Actions column in the right pane for the product release you want to download.

   A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

   Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

   The product packages are downloaded.

Note: You can expand the tree in the right panel by selecting the Products link from the catalog tree. Then, click the vendor link in the right panel. If you select and download multiple products using this method and one of the products cannot be downloaded, the remaining products are not downloaded either. Remove the checks from the products that were processed and repeat the update catalog request.

Migrate Installation Packages Downloaded External to CA MSM

If you have acquired product pax files by means other than through CA MSM, you can add information about these product installation packages to CA MSM from the Software Catalog tab.

Migrating these packages to CA MSM provides a complete view of all your product releases. After a package is migrated, you can use CA MSM to install the product (see page 33).

Follow these steps:

1. Click the Software Catalog tab, and click Insert New Product.

   Note: A product not acquired from the CA Support Online website does not appear in Software Catalog until you perform this step.

   An entry is added for the product.

2. Select the product gen level (for example, SP0 or 0110) for which the package applies.

   The packages for the gen level are listed.

3. Click the Add External Package button.

   You are prompted to enter a path for the package.

4. Specify the USS path to the package you want to migrate, and click OK.

   Information about the package is saved in the CA MSM database.

   Note: To see the added package, refresh the page.

Add a Product

Sometimes, a product is not currently available from the CA Support Online website. For example, if you are testing a beta version of a product, the version is delivered to you by other means. You can add these types of product packages to CA MSM using the Insert New Product action.

Follow these steps:

1. Click the Software Catalog tab, and click the Insert New Product link in the Actions section at the left.

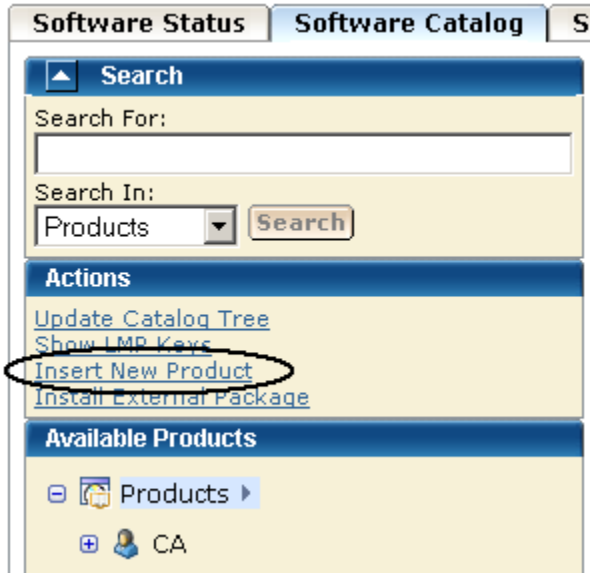

   You are prompted to supply information about the product.

2. Specify the name, release, and gen level of the product, and click OK.

   The product is added to the software catalog.

3. Click the gen level of the product you want to install on the product tree at the left.

   The Base Install Packages section appears at the right.

4. Click the Add External Package button.

   You are prompted to identify the package.

5. Specify the USS path to the package you want to add, and click OK.

   Note: To add several packages from the same location, use masking (see page 32).

   Information about the package is saved in the CA MSM database.

   Note: To see the added package, refresh the page.

Masking for External Packages

Masking lets you add more than one package (see page 31) (or set of maintenance files (see page 43)) from the same location using a pattern (mask). You can use masking for components, maintenance in USS, and maintenance in data sets. You can use masking for files only, not for directories.

Masking: Use the asterisk symbol (*).

- For PDS and PDSE, you can mask members using asterisks.
- For sequential data sets, use the following characters:

  ?

  Match on a single character.

  *****

  Match on any number of characters within a data set name qualifier or any number of characters within a member name or file system name.

  ******

  Match on any number of characters including any number of qualifiers within a data set name.

You can use as many asterisks as you need in one mask. After you enter the mask, a list of files corresponding to the mask pattern appears.

Note: By default, all files in the list are selected. Verify what files you want to add.

Example 1

The following example displays all PDF files that are located in the */a/update/packages* directory:

```
/a/update/packages/*.pdf
```

Example 2

The following example displays all files that are located in the /a/update/packages directory whose names contain *p0*:

```
/a/update/packages/*p0*
```

Example 3

The following example displays all sequential data sets whose name starts with *PUBLIC.DATA.PTFS.*:

```
PUBLIC.DATA.PTFS.**
```

Example 4

The following example displays all members in the PDS/PDSE data set PUBLIC.DATA.PTFLIB whose name starts with *RO*:

```
PUBLIC.DATA.PTFLIB(RO*)
```

Installing Products

This section includes information about how to use CA MSM to install products.

Install a Product

You can install a downloaded product through the Software Catalog, Base Install Packages section. The process starts a wizard that guides you through the installation. At the end of the wizard, a task dynamically invokes the SMP/E and other utilities required to install the product.

Note: If your site uses only one file system (for example, only zFS or only HFS), you can configure CA MSM to use this file system for all installed products regardless of the file system that the product metadata specifies. The settings are available on the System Settings, Software Installation page. The file system type that you specify will override the file system type that the product uses.

Any USS file system created and mounted by CA MSM during a product installation is added in CA MSM as a managed product USS file system. CA MSM lets you enable and configure verification policy that should be applied to these file systems when starting CA MSM. For verification results, review CA MSM output.

These settings are available on the System Settings, Mount Point Management page.

During installation, you select the CSI where the product is to be installed, and specify its zones. You can either specify target and distribution zones to be in the existing CSI data sets, or create new data sets for each zone.

Note: While working with a particular SMP/E environment, the SMP/E environment is locked and other CA MSM users cannot perform any action against it. When the task finishes, logging out from CA MSM, or a CA MSM session is inactive for more than 10 minutes, the lock releases.

Follow these steps:

1. Click the Software Catalog tab, and select the product gen level (for example, SP0 or 0110) you want to install on the product tree at the left.

   Information about the product appears in the Base Install Packages section at the right, for example:

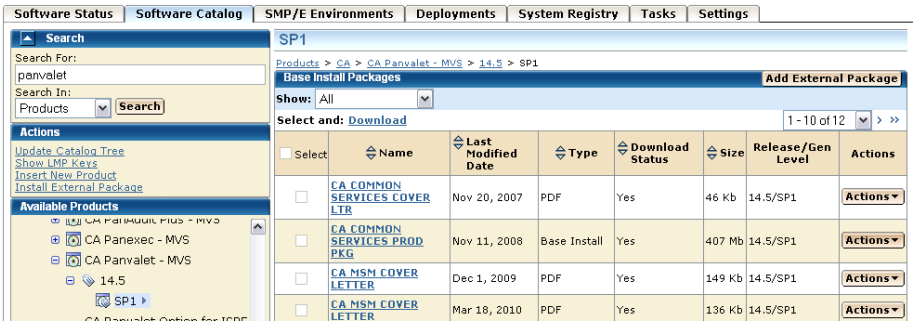

   Note: If a product is acquired external to CA MSM, you can install the product using the Install External Package link. The process starts the wizard.

2. Do one of the following:

   - If the package was acquired using CA MSM, locate the product package that you want to install, click the Actions drop-down list to the right of the package, and select Install.

     or

   - If the package was acquired external to CA MSM, click the Install External Packages link under the Actions section in the left pane, enter the location of the package, and click OK.

   The Introduction tab of the wizard appears.

   Note: An information text area can appear at the bottom of the wizard. The area provides information that helps you progress through the wizard. For example, if a field is highlighted (indicating an error), the information text area identifies the error.

3. Review the information about the installation, and click Next.

   Note: If the license agreement appears for the product that you are installing, scroll down to review it, and accept it.

   You are prompted to select the type of installation.

4. Click the type of installation, and then click Next.

   (Optional) If you select Custom Installation, you are prompted to select the features to install. Select the features, and click Next.

   A summary of the features to install appears, with any prerequisites.

5. Review the summary to check that any prerequisites are satisfied.

   - If no prerequisites exist, click Next.

     You are prompted for the CSI to use for this installation.

   - If prerequisites exist, and they are all satisfied, click Next.

     You are prompted to locate the installed prerequisites. If an installed prerequisite is in more than one CSI or zone, the CSI and Zone drop-down lists let you select the specific instance. After you make the selections, click Next.

     You are prompted for the CSI to use for this installation.

   - If prerequisites are not satisfied, click Cancel to exit the wizard. Install the prerequisites, and then install this product.

     Note: You can click Custom Installation to select only those features that have the required prerequisites. You can click Back to return to previous dialogs.

6. Select an existing CSI, or click the Create a New SMP/E CSI option button, and click Next.

   If you select Create a New SMP/E CSI, you are prompted to specify the CSI parameters (see page 36).

   If you select an existing CSI, the wizard guides you through the same steps. Allocation parameters that you specify for work DDDEFs are applied only to new DDDEFs that might be created during the installation. The existing DDDEFs if any remain intact.

   Note: Only CSIs for the SMP/E environments in your working set are listed. You can configure your working set from the SMP/E Environments tab.

   - If you select a CSI that has incomplete information, the wizard prompts you for extra parameters.
   - If you select an SMP/E environment being used in CA MSM by another user, a notification message appears. You are prevented from performing any actions on the SMP/E environment. You can either wait until the notification message disappears and the SMP/E environment becomes available, or click Cancel to select another SMP/E environment.

   After you select a CSI or specify a new CSI, you are prompted for the target zone to use.

7. Select an existing zone, or click the Create a New SMP/E Target Zone option button. Click Next.

   Note: If you select Create a New SMP/E Target Zone, you perform additional steps similar to the steps for the Create a New SMP/E CSI option. The target zone parameters are pre-populated with the values that are entered for the CSI. You can change them.

If you want the target zone to be created in a new data set, select the Create New CSI Data Set check box and fill in the appropriate fields.

After you select or specify a target zone, you are prompted for the distribution zone to use.

8. Select an existing zone, or click the Create a New SMP/E Distribution Zone option button. Click Next.

   Note: If you selected to use an existing target zone, the related distribution zone is automatically selected, and you cannot select other distribution zone. If you selected to create a new target zone, you create a new distribution zone, and you cannot select existing distribution zone.

   After a distribution zone is selected or specified, a summary of the installation task appears.

   Note: If you select Create a New SMP/E Distribution Zone, you perform additional steps similar to the steps for the Create a New SMP/E CSI option. The distribution zone parameters are prepopulated with the values that are entered for the target zone. You can change them.

   - If you want the distribution zone to be created in a new data set, select the Create New CSI Data Set check box and fill in the appropriate fields.
   - If you want to use the same data set that you have already specified to be created for the target zone, the data set will be allocated using the parameters you have defined when specifying the target zone.

9. Review the summary, and click Install.

   A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

Create a CSI

You can create a CSI while you are installing a product (see page 33). During the process, you are asked to specify the following:

- Data set allocation parameters, which you can then customize for each data set
- Parameters for DDDEF allocation

You can specify data set allocation parameters collectively for all SMP/E data sets, target libraries, and distribution libraries that will be allocated during product installation. You can allocate data sets using one of the following methods:

- Allocate data sets using SMS parameters.
- Allocate cataloged data sets using UNIT and optionally VOLSER.
- Allocate uncataloged data sets using UNIT and VOLSER.

If you allocate uncataloged data sets, you must specify a VOLSER. Based on the value that you enter, CA MSM performs the following validations to help ensure integrity of the installation:

- The value of VOLSER must specify a mounted volume.
- You must have ALTER permissions for the data sets with the entered high-level qualifier (HLQ) on the volume defined by VOLSER.
- To test allocation, CA MSM temporarily allocates one of the uncataloged data sets that should be allocated during the installation.

  1. The data set is allocated with one track for both primary and secondary space.
  2. CA MSM verifies that the data set has been allocated on the specified volume.
  3. The data set is deleted.

  If the data set allocation fails or the data set cannot be found on the specified volume, you cannot proceed with the product installation wizard.

Follow these steps:

1. Click Create a New SMP/E CSI from the product installation wizard.

   You are prompted to define a CSI.

2. Specify the following, and click Next:

   Name

   Defines the name for the environment represented by the CSI.

   Data Set Name Prefix

   Defines the prefix for the name of the CSI VSAM data set.

   Catalog

   Defines the name of the SMP/E CSI catalog.

   Cross-Region

   Identifies the cross-region sharing option for SMP/E data sets.

   Cross-System

   Identifies the cross-system sharing option for SMP/E data sets.

High-Level Qualifier

Specifies the high-level qualifier (HLQ) for all SMP/E data sets that will be allocated during installation. The low-level qualifier (LLQ) is implied by the metadata and cannot be changed.

DSN Type

Specifies the DSN type for allocating SMP/E data sets.

SMS Parameters / Data Set Parameters

Specify if this CSI should use SMS or data set parameters, and complete the applicable fields.

Storage Class (SMS Parameters only)

Defines the SMS storage class for SMP/E data sets.

Management Class (SMS Parameters only)

Defines the management class for SMP/E data sets.

Data Class (SMS Parameters only)

Defines the data class for SMP/E data sets.

VOLSER (Data Set Parameters only)

Defines the volume serial number on which to place data sets.

Note: This field is mandatory if you set Catalog to No.

Unit (Data Set Parameters only)

Defines the type of the DASD on which to place data sets.

Catalog (Data Set Parameters only)

Specifies if you want SMP/E data set to be cataloged.

Note: An information text area can appear at the bottom of the wizard. The area provides information that helps you progress through the wizard. For example, if a field is highlighted (indicating an error), the information text area identifies the error.

Work DDDEF allocation parameters and a list of the data sets to be created for the CSI appear.

3. Specify whether to use SMS or Unit parameters for allocating work DDDEFs for the CSI, and complete the appropriate fields.

   Note: The settings for allocating work DDDEFs are globally defined on the System Settings, Software Installation tab. You must have the appropriate access rights to be able to modify these settings.

4. Review the data set names. Click the Override link to change the high-level qualifier of the data set name and the allocation parameters, and then click Next.

   You are prompted to specify any additional parameters. A new CSI is specified.

Download LMP Keys

When you install a CA Technologies product on z/OS systems, you must license the product on each system that uses the product. You do this by entering CA Common Services for z/OS CA License Management Program (LMP) statements. You can download LMP keys through the Software Catalog tab so that the keys are available for you to enter manually. The Show LMP Keys action retrieves the keys for the products to which your site is entitled.

Follow these steps:

1. Click the Software Catalog tab, and click the Show LMP Keys link in the Actions section at the left.

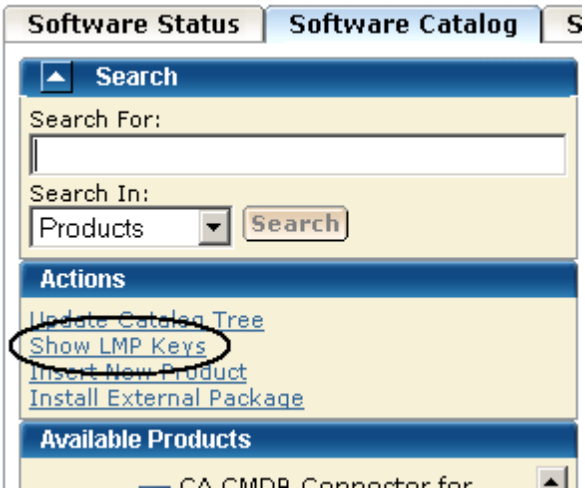

   A list of LMP keys retrieved for the indicated site ID appears.

2. Select the site ID for which you want to list the LMP keys from the Site IDs drop-down list.

   The list is refreshed for the selected site ID.

   If the list is empty or if you want to update the lists, proceed to the next step.

3. Click Update Keys.

   You are prompted to confirm the update.

4. Click OK.

   The LMP keys are retrieved. On completion of the retrieval process, the LMP keys are listed for the selected site.

Note: You can use the Refresh Site IDs button to refresh the information on the page.

Maintaining Products

This section includes information about how to use CA MSM to download and apply product maintenance packages.

How to Apply Maintenance Packages

Use this process to download and apply product maintenance packages.

1. Identify your download method. This section details the steps to use the following download methods:

   - Download Product Maintenance Packages (see page 41)
   - Download Product Maintenance Packages for Old Product Releases and Service Packs (see page 42)
   - Manage Maintenance Downloaded External to CA MSM (see page 43)

   Contact your system administrator, if necessary.

2. Apply the product maintenance package. This section also details the role of USERMODs.

   Note: This section also describes how to back out maintenance that has been applied but not yet accepted.

Download Product Maintenance Packages

You can download maintenance packages for installed products through the Software Catalog tab.

Follow these steps:

1. Verify that your CA MSM login user name is associated with a registered user of the CA Support Online website on the Software Acquisition Settings page.

   CA MSM uses the credentials to access the CA Support Online website.

2. Click the name of the product for which you want to download maintenance on the product tree at the left.

   Maintenance information about the product appears in the Releases section at the right.

3. Click the Update Catalog Release button for the product release for which you want to download maintenance.

   A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

   Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

   The maintenance packages are downloaded.

More information:

Download Maintenance Packages for Old Product Releases and Service Packs (see page 42)

Download Maintenance Packages for Old Product Releases and Service Packs

CA MSM does not retrieve information about old product releases and service packs. If you need maintenance from those releases and service packs, you must add them to the software catalog before you can download the maintenance.

Follow these steps:

1. Click the Software Catalog tab, and click the Insert New Product link in the Actions section at the left.

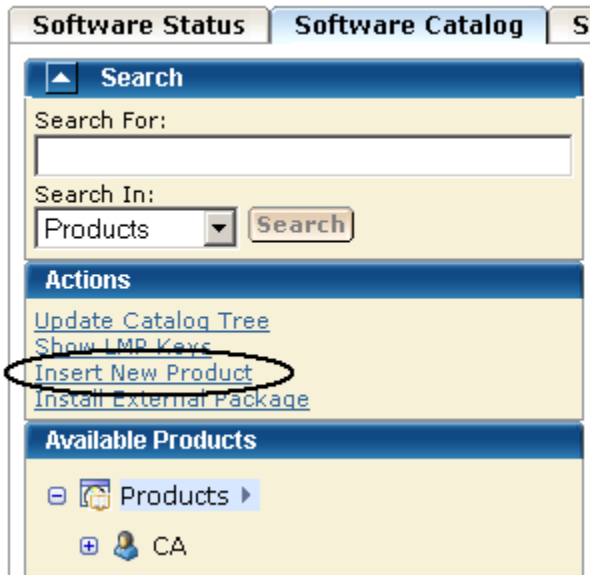

   You are prompted to supply information about the product release.

2. Specify the name, release, and gen level of the product, and click OK.

   Note: Use the same product name that appears on the product tree, and use the release and gen level values as they appear for Published Solutions on the CA Support Online website.

   The product release is added to the software catalog.

3. From the product tree at the left, click the name of the product for which you want to download maintenance.

   Maintenance information about the product appears in the Releases section at the right.

4. Click Update Catalog Release for the added product release.

   Maintenance packages are downloaded. A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

Manage Maintenance Downloaded External to CA MSM

Some maintenance packages, such as unpublished maintenance, APARs, and USERMODs, may be acquired externally to CA MSM. You can add information about these maintenance packages to CA MSM from the Software Catalog tab. The process starts a wizard that guides you through the migration.

Adding these maintenance packages to CA MSM provides you with a complete view of all the maintenance for a product release. After a package is migrated, you can use CA MSM to apply the maintenance (see page 45).

The maintenance package must be located in a z/OS data set or a USS directory. If you use a z/OS data set, it must have an LRECL of 80. If you place the maintenance in a USS directory, copy it in binary mode.

The maintenance is placed as either a single package or an aggregated package that is a single file comprised of multiple maintenance packages. An *aggregated package* is a file that comprises several single maintenance packages (nested packages). When you add an aggregated package, CA MSM inserts all nested packages that the aggregated package includes and the aggregated package itself. In the list of maintenance packages, the aggregated package is identified by the CUMULATIVE type.

When you insert an aggregated package, CA MSM assigns a fix number to it. The fix number is unique and contains eight characters, starting with AM (for Aggregated Maintenance) followed by a unique 6-digit number whose value increases by 1 with each added aggregated package.

Note: If the aggregated maintenance package has the same fix number as one of its nested packages, only the nested packages are added. The aggregated package itself will not be available in the list of maintenance packages.

Follow these steps:

1. Click the Software Catalog tab, and select the product release for which the maintenance applies.

   The maintenance packages for the release are listed.

2. Click the Add External Maintenance button.

   You are prompted to specify the package type and location.

3. Specify the package type and either the data set name or the USS path.

   Note: To add several packages from the same location, use masking (see page 32).

4. Click OK.

   The maintenance package with the related information is saved in the CA MSM database.

   Note: To see the added package, refresh the page.

More information:

Manage Maintenance (see page 45)

View Aggregated Package Details

You can view which nested packages are included in the aggregated package. The information includes the fix number, package type, and package description.

Follow these steps:

1. Click the Software Catalog tab, and select the product release that has the aggregated package whose details you want to view.

   The maintenance packages for the release are listed.

2. Click the Fix # link for the aggregated package.

   The Maintenance Package Details dialog opens.

3. Click the Nested Packages tab.

   A list of nested packages contained in the aggregated package appears.

Manage Maintenance

After maintenance has been downloaded for a product, you can manage the maintenance in an existing SMP/E product installation environment.

Note: While working with a particular SMP/E environment, the SMP/E environment is locked and other CA MSM users cannot perform any action against it. When the task finishes, logging out from CA MSM, or a CA MSM session is inactive for more than 10 minutes, the lock releases.

The following installation modes are available:

Receive and Apply

Receives the maintenance and applies it to the selected SMP/E environment.

Receive and Apply Check

Receives the maintenance and checks if the maintenance can be applied to the selected SMP/E environment.

Receive, Apply Check, and Apply

Receives the maintenance, checks if the maintenance can be applied to the selected SMP/E environment, and applies it if it can be applied.

Receive Only

Receives the maintenance.

The process starts a wizard that guides you through the maintenance steps. At the end of the wizard, a task dynamically invokes the SMP/E and other utilities required to apply the maintenance.

Note: You can also manage maintenance to an SMP/E environment using the SMP/E Environments, Maintenance tab.

Follow these steps:

1. Click the Software Catalog tab, and select the product from the tree at the left.

   Maintenance information appears at the right for the releases you have.

2. Click Update Catalog Release for the release on which you want to apply maintenance.

   The maintenance information is updated.

3. If the information indicates that maintenance is available, click the Release Name link.

   The maintenance packages are listed, for example:

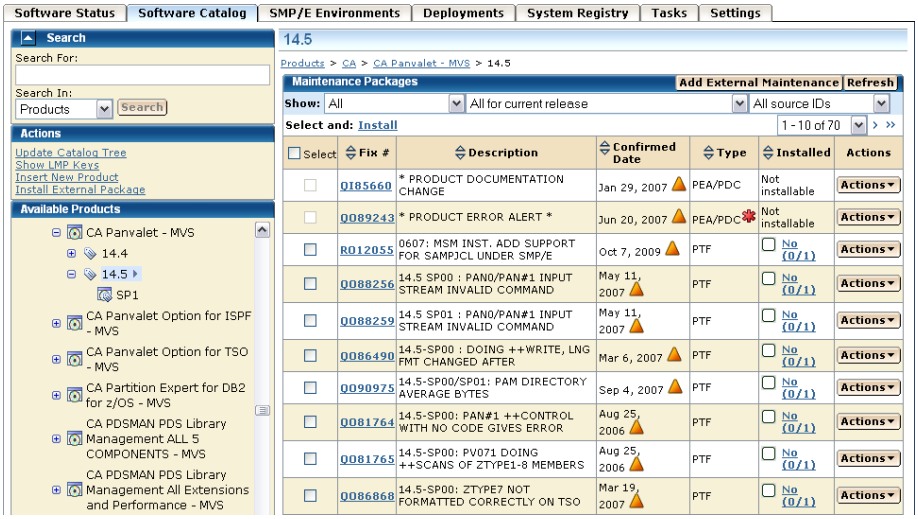

   Red asterisks identify HIPER maintenance packages.

4. Click the Fix # link for each maintenance package you want to install.

   The Maintenance Package Details dialog appears, identifying any prerequisites.

5. Review the information on this dialog, and click Close to return to the Maintenance Packages section.

6. Select the maintenance packages you want to install, and click the Install link.

   Note: The Installed column indicates whether a package is installed.

   The Introduction tab of the wizard appears.

7. Review the information about the maintenance, and click Next.

   The packages to install are listed.

8. Review and adjust the list selections as required, and click Next.

   The SMP/E environments that contain the product to maintain are listed. Only environments in your working set are listed.

9. Select the environments in which you want to install the packages.

10. Click Select Zones to review and adjust the zones where the maintenance will be installed, click OK to confirm the selection and return to the wizard, and click Next.

    Note: If you select an SMP/E environment being used in CA MSM by another user, a notification message appears. You are prevented from performing any actions on the SMP/E environment. You can either wait until the notification message disappears and the SMP/E environment becomes available, or click Cancel to select another SMP/E environment.

11. Select the installation mode for the selected maintenance, and click Next.

    - If prerequisites exist and are available, review them and click Next. CA MSM installs these prerequisites as part of the process. If a prerequisite is *not* available, the wizard cannot continue. You must acquire the prerequisite and restart the process.
    - If HOLDDATA (see page 148) entries exist, review and select them, and click Next.

    SMP/E work DDDEFs of SMPWRK*x* and SYSUT*x*, with their allocation parameters, are listed.

    Note: For more information about SMPWRK*x* and SYSUT*x* data sets, see the *IBM SMP/E for z/OS Reference*.

12. Review the allocation parameters of work DDDEFs, and edit them if necessary to verify, that sufficient space is allocated for them during the maintenance installation:

    Note: Changes in the allocation parameters apply to the current maintenance installation only.

    a. Click Override for a DDDEF to edit its allocation parameters.

       A pop-up window opens.

    b. Make the necessary changes, and click OK to confirm.

       The pop-up window closes, and the DDDEF entry is selected in the list indicating that the allocation parameters have been overridden.

    Note: To update allocation parameters for all DDDEFs automatically, click Retrieve DDDEF. CA MSM provides values for all DDDEFs based on the total size of the selected maintenance packages that you want to install. All DDDEF entries are selected in the list indicating that the allocation parameters have been overridden.

    - If you want to cancel a parameter update for any DDDEF, clear its check box.
    - If you want to edit the allocation parameters for a particular DDDEF after you automatically updated them using the Retrieve DDDEF button, click Override. Make the necessary changes and click OK to confirm, and return to the wizard.

13. (Optional) Review SMP/E work DDDEF and their allocation parameters for the selected SMP/E zones, and click Close to return to the wizard.

    Note: The allocation parameters can differ from the allocation parameters that you obtained using the Retrieve DDDEF button.

14. Click Next.

    A summary of the task appears.

15. Review the summary, and click Install.

    A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

    Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

The task applies the maintenance. You can accept the maintenance (except USERMODs) using the SMP/E Environments, Maintenance tab. As a best practice, CA MSM prevents you from accepting USERMODs.

More information:

Download Product Maintenance Packages (see page 41)
Download Maintenance Packages for Old Product Releases and Service Packs (see page 42)

View Installation Status of Maintenance Package

You can view installation status details of each maintenance package, including a list of SMP/E environments where the package is installed. You can also see the SMP/E environment data sets, and the installation status of the package for each SMP/E environment zone. For example, a maintenance package can be received in the global zone, but applied in a target zone, and accepted in a distribution zone.

Note: The installation status is not available for aggregated maintenance packages, for packages that are uninstallable, and for packages that do not have available SMP/E environments for installation.

Depending on the package status for each zone, you can see available actions for the package. For example, if the package is not received in an SMP/E environment zone, the Install action is available.

Follow these steps:

1. Click the Software Catalog tab, and select the product release that has the maintenance package whose installation status you want to view.

   The maintenance packages for the release are listed.

2. Click the status link in the Installed column for the maintenance package.

   The Maintenance Package Details dialog opens to the Installation Status tab. A list of SMP/E environments with package status per zone appears.

   Note: Click the Actions drop-down list to start the installation wizard for packages that are not yet installed in at least one SMP/E environment zone, or the accept wizard for packages that are not accepted in at least one SMP/E environment zone. Click Install to More Environments to install the maintenance package in one or more SMP/E environments available for the package.

USERMODs

A product USERMOD can be provided as a published maintenance package downloaded during the Update Catalog process. When CA MSM downloads a package including a ++USERMOD statement, it is loaded under the product with a USERMOD type. You can install these packages using CA MSM but cannot accept them because they are not intended to be permanent.

You can create a USERMOD manually, or we can provide an unpublished maintenance package as a USERMOD. In this case, the USERMOD file, which contains the ++USERMOD statement and the body of the USERMOD, must be managed as an externally downloaded package (see page 43).

GROUPEXTEND Mode

CA MSM lets you invoke the SMP/E utility with the GROUPEXTEND option enabled for managing (applying and accepting) maintenance.

Sometimes before you install a maintenance package, you install other maintenance packages first (SYSMODs).

If a SYSMOD - prerequisite for the required maintenance package, has not been applied or cannot be processed, you can install the maintenance package in GROUPEXTEND mode. (For example, the SYSMOD is held for an error, a system, or a user reason ID; it is applied in error; it is not available.) The SMP/E environment where the product is installed automatically includes a superseding SYSMOD.

Note: When applying maintenance in GROUPEXTEND mode, the SMP/E environment *must* receive all SYSMODs that are included in the GROUPEXTEND option.

When you apply maintenance in GROUPEXTEND mode, the following installation modes are available:

Apply Check

Checks if the maintenance can be applied to the selected SMP/E environment in GROUPEXTEND mode.

Apply

Applies the maintenance to the selected SMP/E environment in GROUPEXTEND mode.

Apply Check and Apply

Checks if the maintenance can be applied to the selected SMP/E environment in GROUPEXTEND mode. Then applies it if possible.

For the GROUPEXTEND option, CA MSM does not automatically receive and display maintenance or HOLDDATA prerequisites that must be bypassed when applying the maintenance. Apply check mode lets you check if any prerequisites or HOLDDATA exist and report them in the task output.

You can also use the following similar installation modes to accept maintenance in GROUPEXTEND mode:

- Accept Check
- Accept
- Accept Check and Accept

How Maintenance in GROUPEXTEND Mode Works

We recommend that you apply maintenance in GROUPEXTEND mode in the following sequence:

1. Receive all SYSMODs that you want to include by the GROUPEXTEND option.
2. Run the maintenance in Apply check mode.
   - If the task fails, review SMPOUT in the task output. Review if there are missing (not received) SYSMODs or HOLDDATA that must be resolved or bypassed.
   - If the task succeeds, review SMPRPT in the task output. Review what SYSMODs were found and applied.
3. Run the maintenance in Apply mode, and specify SYSMODs that you want to exclude and HOLDDATA that you want to bypass, if any exist.

   The followings options are available for bypassing HOLDDATA:

   - HOLDSYSTEM
   - HOLDCLASS
   - HOLDERROR
   - HOLDUSER

   Note: For more information about the BYPASS options, see the *IBM SMP/E V3Rx.0 Commands*. *x* is the SMP/E release and corresponds to the SMP/E version that you use.

You can run the maintenance in Apply mode in the same CA MSM session after Apply check mode is completed. The values that you entered for Apply check mode are then prepopulated on the wizard dialogs.

Manage Maintenance in GROUPEXTEND Mode

CA MSM lets you invoke the SMP/E utility with the GROUPEXTEND option enabled for managing (applying and accepting) maintenance.

Note: While working with a particular SMP/E environment, the SMP/E environment is locked and other CA MSM users cannot perform any action against it. When the task finishes, logging out from CA MSM, or a CA MSM session is inactive for more than 10 minutes, the lock releases.

Follow these steps:

1. Click the SMP/E Environments tab, and select the SMP/E environment from the tree on the left side.

   A list of products installed in the SMP/E environment appears.

   Note: If you select an SMP/E environment being used in CA MSM by another user, a notification message appears. You are prevented from performing any actions on the SMP/E environment. You can either wait until the notification message disappears and the SMP/E environment becomes available, or click Cancel to select another SMP/E environment.

2. Click the Maintenance link.

   A list of maintenance packages for the products installed in the SMP/E environment appears.

3. Select the maintenance packages that you want to apply in GROUPEXTEND mode, and click the Apply GROUPEXTEND link.

   The Introduction tab of the wizard appears.

4. Review the information about the maintenance, and click Next.

   The packages that you want to apply are listed.

   Note: Click a link in the Status column for a maintenance package, if available, to review a list of zones. The zones indicate, where the maintenance package is already received, applied, or accepted. Click Close to return to the wizard.

5. Review the packages, and click Next.

   The Prerequisites tab of the wizard appears.

   Important! For the GROUPEXTEND option, CA MSM does not automatically receive and display maintenance or HOLDDATA prerequisites that must be bypassed when applying the maintenance. Apply check mode lets you review if any prerequisites or HOLDDATA exist and report them in the task output. We recommend that you run the maintenance in Apply check mode first.

6. Read the information that is displayed on this tab, and click Next.

   Installation options appear.

7. Specify installation options as follows, and click Next:

   a. Select the installation mode for the selected maintenance.

   b. Review the GROUPEXTEND options and select the ones that you want to apply to the maintenance:

      NOAPARS

      Excludes APARs that resolve error reason ID.

      NOUSERMODS

      Exclude USERMODs that resolve error user ID.

   c. (Optional) Enter SYSMODs that you want to exclude in the Excluded SYSMODs field. You can enter several SYSMODs, separate them by a comma.

   The Bypass HOLDDATA tab of the wizard appears.

8. (Optional) Enter the BYPASS options for the HOLDDATA that you want to bypass during the maintenance installation. You can enter several BYPASS options, separate them by a comma.

9. Click Next.

   A summary of the task appears.

10. Review the summary, and click Apply GROUPEXTEND.

    A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

    Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

    - If you run the maintenance installation in Apply check mode and the task succeeds, review SMPRPT in the task output. Review what SYSMODs were found and applied.
    - If you run the maintenance installation in Apply check mode and the task fails, review SMPOUT in the task output. Review if there are missing (not received) SYSMODs or HOLDDATA that must be resolved or bypassed.

You can accept the maintenance (except USERMODs) in the GROUPEXTEND mode using the SMP/E Environments, Maintenance tab. As a best practice, CA MSM prevents you from accepting USERMODs.

Note: You cannot accept USERMODs in GROUPEXTEND mode. Providing you have not enabled NOUSERMODS option, you can install USERMODs that are prerequisites for the maintenance package being installed.

Back Out Maintenance

You can back out an applied maintenance package (but not an accepted maintenance package) through the SMP/E Environments tab. The process starts a wizard that guides you through the backout.

Note: While working with a particular SMP/E environment, the SMP/E environment is locked and other CA MSM users cannot perform any action against it. When the task finishes, logging out from CA MSM, or a CA MSM session is inactive for more than 10 minutes, the lock releases.

Follow these steps:

1. Click the SMP/E Environments tab, and select the SMP/E environment from which you want to back out maintenance on the tree on the left side.

   Products installed in the environment are listed.

2. Select the product component from which you want to back out maintenance.

   The features in the component are listed.

   Note: You can back out maintenance from all the products in the environment. Click the Maintenance tab to list all the maintenance packages for the environment.

3. Select the function from which you want to back out maintenance.

   The maintenance packages for the feature are listed.

   Note: You can use the Show drop-down list to show only applied packages.

4. Select the packages that you want to back out, and click the Restore link.

   The maintenance wizard opens to the Introduction step.

   Note: If you select an SMP/E environment being used in CA MSM by another user, a notification message appears. You are prevented from performing any actions on the SMP/E environment. You can either wait until the notification message disappears and the SMP/E environment becomes available, or click Cancel to select another SMP/E environment.

5. Review the information about the backout, and click Next.

   The packages to back out are listed.

6. Review and adjust the list selections as required, and click Next.

   Note: To review and adjust a list of zones from where you want to restore the maintenance, click Select Zones. Click OK to confirm the selection and return to the wizard.

   The Prerequisite tab of the wizard appears.

7. Review the prerequisites if they exist, and click Next. CA MSM restores these prerequisites as part of the maintenance backout process.

   A summary of the task appears.

8. Review the summary, and click Restore.

   A dialog that shows the progress of the task opens. When the task completes, you can click Show Results on the Progress tab to close this dialog. The task output browser opens and you can view the action details. Click Close to close the task output browser.

Note: While a task is in progress, you can perform other work. You can click Hide to exit the dialog and view the task status later on the Tasks tab.

Setting System Registry

This section includes information about how to use CA MSM to set the system registry. The *system registry* contains information about the systems that have been defined to CA MSM and can be selected as a target for deployments. You can create Non-Sysplex, Sysplex, Shared DASD Cluster, and Staging systems as well as maintain, validate, view, and delete a registered system, and investigate a failed validation.

View a System Registry

You can view a system registry by using the CA MSM.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, Shared DASD Clusters, or Staging Systems from the tree on the left side.

   Information about the systems that you selected appears on the right side.

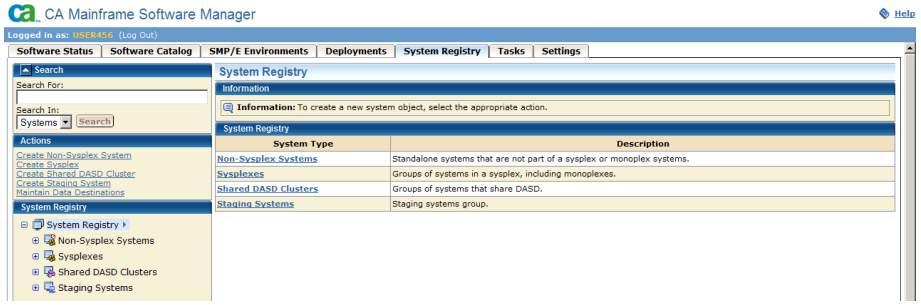

Create a Non-sysplex System

You can create a non-sysplex system registry.

Follow these steps:

1. Click the System Registry tab, and in the Actions section click the Create Non-Sysplex System link.

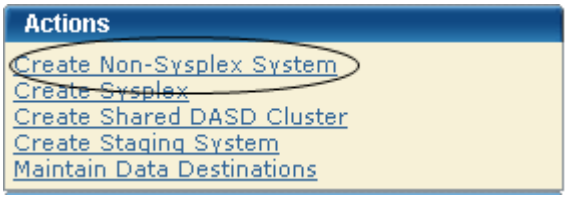

   The New Non-Sysplex System dialog appears.

   Note: The asterisk indicates that the field is mandatory.

2. Enter the following information, and click Save:

   Name

   Enter the non-sysplex system name.

   Limits: Eight characters

   Note: Sysplex and non-sysplex systems can have the same name. Use the Description field to differentiate between these systems.

   Description

   Enter the description.

   Limits: 255 characters

   CCI System ID

   (Optional) Enter the CAICCI system ID.

   Limits: Eight characters

   Note: The *CAICCI system ID* is a unique name for a system that is part of a CAICCI network. If you do not specify one, CA MSM obtains it using a validate action.

   The non-sysplex system is saved, and its name appears in the non-sysplex system list on the left.

   Note: To withdraw this create request, click Cancel.

3. Detail the nonstaging system.

Important! z/OS systems running under VM are treated as being in BASIC mode and not LPAR mode. As a result, the LPAR number is null in the z/OS control block. When the LPAR number is null, the system validation output shows the following message:

```
Property Name: z/OS LPAR Name, Value: ** Not Applicable **.
```

Create a Sysplex or Monoplex

If you have monoplexes with the same sysplex name, you can create a sysplex or monoplex system registry. Monoplexes are stored in the sysplex registry tree but with the name of the sysplex system and not the monoplex sysplex name. For example, you have a system XX16 defined as a monoplex, with a sysplex name of LOCAL. The system registry displays the system as a sysplex, with the name LOCAL. This sysplex contains one system: XX16.

Follow these steps:

1. Click the System Registry tab, and in the Actions section click the Create Sysplex link.

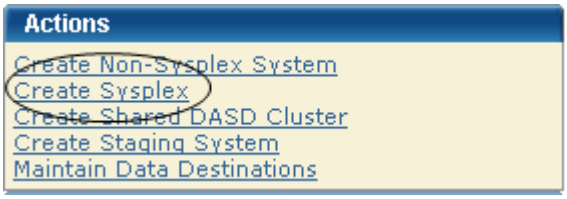

   The New Sysplex dialog appears.

   Note: The asterisk indicates that the field is mandatory.

2. Enter the following and click Save.

   Name

   Enter the sysplex system name.

   Limits: Eight characters

   Description

   Enter the description.

   Limits: 255 characters

   Sysplex and non-sysplex system can have the same name. Use the Description field to differentiate these systems.

   The sysplex system is saved, and its name appears in the sysplex list on the right.

   Note: Click Cancel to withdraw this create request.

   Important! z/OS systems running under VM are treated as being in BASIC mode and not LPAR mode. As a result, the LPAR number is null in the z/OS control block. In this case, the system validation output includes the following message:

   ```
   Property Name: z/OS LPAR Name, Value: ** Not Applicable **.
   ```

3. Right-click the newly added sysplex and select Create Sysplex System to add a system to a sysplex. Repeat this process for each system belonging to this sysplex.

4. Enter the following data items for each system:

   Name

   Enter the sysplex system name.

   Limits: Eight characters

   Note: Sysplex and non-sysplex systems can have the same name. Use the Description field to differentiate between these systems.

   Description

   Enter the description.

   Limits: 255 characters

   CCI System ID

   (Optional) Enter the CAICCI system ID.

   Limits: Eight characters

   Note: The *CAICCI system ID* is a unique name for a system that is part of a CAICCI network. If you do not specify one, CA MSM obtains it using a validate action.

   The non-sysplex system is saved, and its name appears in the non-sysplex system list on the left.

   Note: To withdraw this create request, click Cancel.

5. Detail the nonstaging system.

Create a Shared DASD Cluster

You can create a shared DASD cluster.

Follow these steps:

1. Click the System Registry tab, and in the Actions section click the Shared DASD Cluster link.

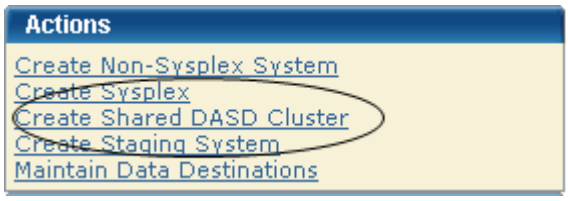

   The New Shared DASD Cluster dialog appears.

   Note: The asterisk indicates that the field is mandatory.

2. Enter the following information, and click Save:

   Name

   Enter the shared DASD cluster name.

   Limits: Eight characters

   Note: Each shared DASD cluster name must be unique and it is not case-sensitive. For example, DASD1 and dasd1 are the same shared DASD cluster name. A shared DASD cluster can have the same name as a non-sysplex, sysplex, or staging system.

   Description

   Enter the description.

   Limits: 255 characters

   The shared DASD cluster is saved, and its name appears in the Shared DASD Clusters section on the right.

   Note: Click Cancel to withdraw this create request.

3. Right-click the newly added DASD cluster name and select Add System or Sysplex to this Shared DASD Cluster. Select the systems or sysplexes that you want to add to the DASD cluster.

Create a Staging System

You can create a staging system.

Follow these steps:

1. Click the System Registry tab, and in the Actions section click the Create Staging System link.

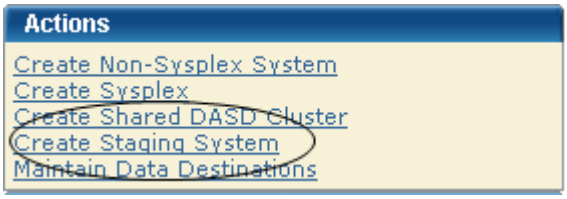

   The New Staging System dialog appears.

   Note: The asterisk indicates that the field is mandatory.

2. Enter the following information, and click Save:

   Name

   Enter the staging system name.

   Limits: Eight characters

   Note: Each staging system name must be unique and is not case-sensitive. For example, STAGE1 and stage1 are the same staging system name. A staging system can have the same name as a non-sysplex, sysplex, or a shared DASD cluster.

   Description

   Enter the description.

   Limits: 255 characters

   The staging system is saved, and it appears in the Staging System Registry on the right.

   Note: Click Cancel to withdraw this create request.

Authorization

CA MSM supports the following authorization modes for the system registry.

Edit Mode

Lets you update and change system registry information.

Note: After the information is changed, you must click Save to save the information or Cancel to cancel the changed information.

View Mode

Lets you view system registry information.

Note: You cannot edit information in this mode.

Change a System Registry

You can change the system registry if you have Monoplexes with the same sysplex name (for example: LOCAL). Instead of showing multiple LOCAL sysplex entries which would need to be expanded to select the correct Monoplex system, the CA MSM System Registry shows the actual Monoplex System name at the top level Sysplex Name.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, Shared DASD Clusters, or Staging Systems from the tree on the left side.

   Information about the systems that you selected appears on the right side.

2. Select the system to change.

   Detailed information about the system appears on the right side.

3. Update the following information as needed. The information that you update is dependent on whether you are changing a Non-Sysplex System (see page 55), Sysplex (see page 56), Shared DASD Cluster (see page 57), or Staging System (see page 58).

4. Depending on the type of system, do one of the following:

   - For Shared DASD or sysplex system only, select the contact system (see page 65), which is the system where the Shared DASD or FTP is located. The FTP location should be set to the contact system URI. The contact system is used for remote credentials.

     For example, if the contact system is set to CO11, FTP location URI is set to XX61 and the remote credentials are set up for CO11, the deployment could fail because your remote credentials might not be the same on both systems (CO11 and XX61) and, because you set the Contact System to CO11 but you are contacting to XX61, a spawn will be started on CO11 but CA MSM will look for the output on XX61 because that is where the FTP location was set.

     Note: Monoplexes are stored in the Sysplex registry tree but with the name of the Monoplex System and not the Monoplex Sysplex name. For example, a system XX16 defined as a Monoplex, with a sysplex name of LOCAL. It will be depicted in the System Registry as a Sysplex with the name of XX16. This sysplex will contain one system: XX16.

     The FTP and DATA Destinations at the system level are not used when the Sysplex is a Monoplex. The only FTP Location and Data Destinations that are referenced are those defined at the Sysplex Level.

   - For Staging systems, enter the GIMUNZIP volume and/or zFS candidate volumes (see page 66).

     The zFS candidate volumes let you specify an optional list of VOLSERs used during the allocation of zFS container data sets for USS parts.

5. Select one of the following actions from the Actions drop-down list in the General bar:

   Cancel

   Cancel this maintenance.

   Save

   Save the changes to this maintenance.

   Validate

   Validate authenticates this entry.

   Note: The validation process is done in steps; each system in this request is validated with the last step summarizing, verifying, and confirming the validation. If the validation fails this step shows how the validation failed. You can investigate the failed validation (see page 63).

   Validation Rules

   - For a Non-Sysplex system, that single system is validated and the last step summarizes, verifies, and confirms the validation.
   - For a Sysplex system, each system within the Sysplex is validated as an individual step and the last step summarizes, verifies, and confirms the validation.
   - For Shared DASD Cluster each Non-Sysplex system is validated, each Sysplex system is validated as described in the Sysplex Rule and the last step summarizes, verifies, and confirms the validation.

   Note: A Staging system is not validated.

   When a system is validated, the status appears in the Status field.

The following are the system validation results:

Validated

Indicates that the system is available, status is updated as valid, and system registry is updated with results from validation.

Validation in Progress

Indicates that the system status is updated to in progress.

Validation Error

Indicates that the system status is updated to error, and you can investigate the failed validation (see page 63).

Not Validated

Indicates that this system has not been validated yet.

Not Accessible

Indicates that the she system has not been validated because it is no longer available or was not found in the CCI Network.

Validation Conflict

Indicates that the system has been contacted but the information entered then different then the information retrieved.

Error Details

When there is a validation conflict, the Error details button appears. Click this button to find the reason for this conflict. You can investigate the failed validation (see page 63).

Note: The error reason resides in local memory. If the message *Please validate the system again* appears, the local memory has been refreshed and the error has been lost. To find the conflict again, validate this system again.

Conflict Details

When a validation is in conflict, the Error details button appears. Click this button to find the reason for this conflict. You can investigate the failed validation (see page 63).

Note: The conflict reason is kept in local memory. If the "Please validate the system again." message appears, the local memory has been refreshed and the conflict has been lost. To find the conflict again, validate this system again.

Failed Validations

Use the following procedures in this section to investigate a failed validation, make corrections, and revalidate:

- Investigate a Failed Validation using the Tasks Page (see page 63)
- Investigate a Failed Validation Immediately After a Validation (see page 64)
- Download a Message Log (see page 64)
- Save a Message Log as a Data Set (see page 65)
- View Complete Message Log (see page 65)

Note: The CA MSM screen samples in these topics use a non-sysplex system as an example. The method also works for a sysplex or a shared DASD cluster.

Investigate a Failed Validation Using Task Output Browser

You can investigate a failed validation, make corrections, and validate it again.

Follow these steps:

1. On the System Registry tab, in the column on the left, find the system with a validation status error and make a note of it.
2. Click the Tasks tab and then click Task History.
3. At the Show bar, select All task, or My task to list the tasks by Owner.

   Note: You can refine the task list by entering USER ID, types, and status.
4. Find the failed validation and click the link in the Name column.

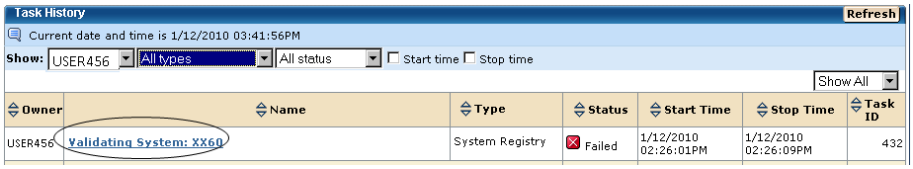

The Validate System Task Output Browser appears.

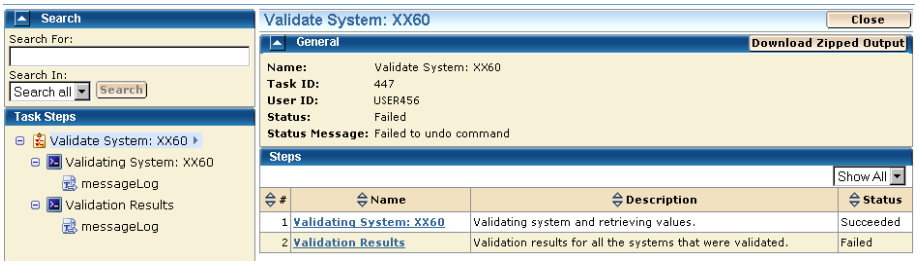

5. Click the Validation Results link to view the results.

6. Click the messages log to review the details for each error.

   Note: You can analyze the error results and can determine the steps that are required to troubleshoot them.

7. Correct the issue and validate again.

Investigate a Failed Validation After Validation

You can investigate a failed validation, make corrections, and validate it again.

Follow these steps:

1. On the System Registry tab, in the column on the left, find the system with a validation status error, and make a note of it.
2. Click Details to see the error details.
3. If the error message prompts you to revalidate the system, click Validate.
4. Click the Progress tab.
5. Click Show Results to view the results.

   The validation results appear.

6. Click the messages logs to review the details for each error.

   Note: You can analyze the error results and can determine the steps that are required to troubleshoot them.

7. Correct the issue and validate again.

Download a Message Log

You can save the message log in the following ways:

- To download a zipped file of all the text messages for this validation, click the Deployment Name on the top left tree. Click the Download Zipped Output button on the General menu bar. Save this file.
- To download as TXT, click the Deployment Name or the Deployment Results on the left tree. Click the Action button on the Message Log bar and click the Download as TXT. Save this file.
- To download as ZIP, click the Deployment Name or the Deployment Results on the left tree. Click the Action button on the Message Log bar and click the Download as ZIP. Save this file.

Save a Message Log as a Data Set

You can save a message log as a data set.

Follow these steps:

1. Click the Deployment Name or the Deployment Results on the left tree. Click the Action button on the Message Log bar, and click the Save as Data Set.

   The Save Output as Data Set dialog appears.

   Note: This information is sent to CA Support to analyze the failed deployment.

   Note: The asterisk indicates that the field is mandatory.

2. Enter the following information and click OK:

   Data Set Name

   Enter a data set name. CA MSM generates a value.

   VOLSER

   For non-SMS data, enter the Volser.

   Example:

   Volser: SYSP01 and SYSP02

   Storage Class

   For SMS Allocation data, enter the Storage Class.

   The message log is saved as a data set.

View Complete Message Log

To view the complete message log for a failed validation, click Show All.

Note: To close the message log, click Close.

Contact System

The *contact system* defines which system the deployment is unpackaged on. That is, which system CAICCI is spawned to run the unpackaging.

When deploying to a shared DASD cluster, sysplex, or both, the deployment is sent to only one system in that configuration, where it is unpackaged. The expectation is that all other systems within that configuration have access to the unpackaged deployment.

For a shared DASD cluster or sysplex, the URI must be the URI of the Contact System. Also, set up Remote Credentials for the contact system, because they are used to retrieve the deployment results.

zFS Candidate Volumes

You can use a *zFS candidate volume* when your environmental setup dictates that zFS container data sets are directed to the specified volume.

When your environmental setup dictates that zFS container data sets are directed to specified zFS candidate volumes, use one or more of the candidate volumes. CA MSM uses the candidate volumes in the IDCAMS statement to create the zFS container VSAM data set.

The zFS candidate volumes are only required if the following statements are true:

- Your deployment has USS parts.
- You are doing a container copy.
- You selected zFS as the container type.
- The remote system requires it.

  Note: Remote system requirement is customer defined.

To allocate and maintain your disk, the following products are recommended:

CA Allocate

CA Allocate is a powerful and flexible allocation management system that lets the Storage Administrator control the allocation of all z/OS data sets.

CA Disk Backup and Restore

CA Disk is a flexible, full-featured hierarchal storage management system.

You can also use the following standard IBM techniques:

- Allocation exits
- ACS routines

If you do not implement any of these options, z/OS needs a candidate list of volumes for placing the zFS archive.

Maintain a System Registry using the List Option

Follow these steps:

1. Click the System Registry tab.

   The System Registry window appears.

2. In the System Registry panel on the right, click the System Type link, and then click the system name.

   The detailed system entry information appears.

Delete a System Registry

Follow these steps:

1. Click the System Registry tab and on the right, in the System Registry panel, select Non-Sysplex Systems, Sysplexes, Shared DASD Clusters, or Staging Systems.

   The system list appears.

2. Select each system registry that you want to delete, click Delete, and then click OK to confirm.

   The system is deleted.

FTP Locations

The FTP (see page 67) Locations lists the current FTP locations for this system. You can add (see page 67), edit (see page 69), set default (see page 70), or remove (see page 70) FTP (see page 67) locations.

An FTP location must be defined for every system. They are used to retrieve the results of the deployment on the target system regardless if the deployment was transmitted through FTP or using Shared DASD. They are also used if you are moving your deployments through FTP. You will need the URI (host system name), port number (default is 21), and the directory path, which is the landing directory. The landing directory is where the data is temporarily placed during a deployment.

Deployment FTP Locations

File Transfer Protocol (FTP) is a protocol for transfer of files from one computer to another over the network.

Define an FTP location for every system if you deploy to specified systems within a sysplex. They are used to retrieve the deployment results on the target system regardless of whether the deployment was transmitted through FTP or using shared DASD. They are also used when you are moving your deployments through FTP. You need the URI (host system name), port number (default is 21), and the directory path, which is the landing directory. The landing directory is where the data is temporarily placed during a deployment.

Add FTP Locations

You can add FTP (see page 67) locations.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, or Shared DASD Clusters from the tree on the left side.

   Information about the systems that you selected appears on the right side.

2. Select the system you want to create FTP locations for.

   Detailed information about the system appears on the right side.

3. Click the FTP Locations tab.

   The FTP Locations window appears.

4. Click Add.

   The New FTP Location dialog appears.

   Note: The asterisk indicates that the field is mandatory.

5. Enter the following information, and click Save:

   URI

   Enter the URI.

   Limits: Maximum length is 255.

   Port

   Enter the Port.

   Limits: Maximum Port number is 65535 and must be numeric.

   Default: 21

   Directory Path

   Enter the Directory Path.

   Limits: Must start with a root directory, that is /.

   The new FTP location appears on the list.

   Note: Click Cancel to withdraw this create request.

More information:

Edit FTP Locations (see page 69)
Delete FTP Locations (see page 70)
Set FTP Location Default (see page 70)

Edit FTP Locations

You can edit FTP (see page 67) locations.

Note: The asterisk indicates that the field is mandatory.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, or Shared DASD Clusters from the tree on the left side.

   Information about the systems that you selected appears on the right side.

2. Select the system you want to change FTP locations for.

   Detailed information about the system appears on the right side.

3. Click the FTP Location tab.

   The FTP Locations window appears.

4. Select the FTP location, click the Actions drop-down list, and select Edit.

   The Edit FTP Location dialog appears.

5. Update the following and click Save:

   URI

   Enter the URI.

   Limits: Maximum length is 255.

   Port

   Enter the Port.

   Limits: Maximum Port number is 65535 and must be numeric.

   Default: 21

   Directory Path

   Enter the Directory Path.

   Limits: Most start with a root directory, that is, /.

   Your changes are saved.

   Note: Click Cancel to close this dialog without saving your changes.

Set FTP Location Default

You can set an FTP (see page 67) location default.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, or Shared DASD Clusters from the tree on the left side.

   Information about the systems that you selected appears on the right side.

2. Select the system you want to set the FTP location default to.

   Detailed information about the system appears on the right side.

3. Click the FTP Locations tab.

   The FTP Locations window appears.

4. Select the FTP location you want to set as the default, and then select Default from the Actions drop-down list.

   Default appears in the Default column, and this location becomes the default FTP location.

   Note: The Default action is not available if only one FTP location is defined.

Delete FTP Locations

You can delete FTP (see page 67) locations.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, or Shared DASD Clusters from the tree on the left side.

   Information about the systems that you selected appears on the right side.

2. Select the system you want to delete FTP locations from.

   Detailed information about the system appears on the right side.

3. Click the FTP Locations tab.

   The FTP Locations window appears.

4. Click the Select box for each FTP location you want to delete, click Remove, and then click OK to confirm.

   The FTP location is deleted from this system.

Data Destinations

The Data Destinations page lists the current data destinations for this system. The following choices are available:

FTP

When FTP is selected as the transport mechanism, the deployment data is shipped to the target system through FTP. The data is temporarily placed on the target system at the landing directory that the FTP Location information section of the system registry specifies.

Shared DASD

When you specify shared DASD, CA MSM uses a virtual transport technique. That is, it does not actually copy the data from one system to the other. Because the two systems share DASD, there is no need to copy the data. All of the deployment data is kept in the USS file systems that CA MSM manages.

Even though the DASD is shared, it is possible that the remote system does not find the deployment data in the USS file system. Therefore, CA MSM temporarily unmounts the file system from the CA MSM driving system and mounts it in read-only mode on the remote system.

For CA MSM to determine where to mount the file system on the remote system, specify a mount point location in the data destination. In addition, you can provide allocation information for the creation of the deployment file system. The file system is created on the shared DASD, on the CA MSM driving system.

Data destinations are assigned to non-sysplex and sysplex systems, and shared DASD clusters. Data destinations are named objects, and can be assigned to multiple entities in the system registry. Data destinations can have their own independent maintenance dialogs.

The deployment process on the remote system uses the remote allocation information and lets you control, where the deployed software is placed. By specifying the GIMUNZIP VOLSER, CA MSM adds a volume= parameter to the GIMUNZIP instructions on the remote system. The list of zFS VOLSERs is needed only if both of the following situations occur:

- The software that you are deploying contains USS parts.
- You select a container copy option during the deployment process.

Note: The FTP and data destinations at the system level are not used when the sysplex is a monoplex. The only FTP locations and data destinations that are referenced are defined at the sysplex level.

Create Data Destinations

You can create data destinations that define the method that CA MSM uses to transfer the deployment data to the target systems.

Follow these steps:

1. Click the System Registry tab, and in the Actions section click the Maintain Data destinations link.

   The Maintains Data Destinations dialog appears.

2. Click Create.

   The New Data Destination dialog appears.

   Note: The asterisk indicates that the field is mandatory.

3. Enter the following information, and click Save:

   Name

   Enter a meaningful name.

   Limits: Maximum 64 characters.

   Note: Each data destination name must be a unique name and it is not case-sensitive. For example DATAD1 and datad1 are the same data destination name.

   Description

   Enter the description.

   Limits: Maximum 255 characters.

   Transmission Method

   Select the transmission method.

   Default: Shared DASD.

   Mount Point

   (Shared DASD only) Enter the mount point directory path, which is a directory path that must exist on the target system. The user that is doing the deployment must have write permission to this directory, and mount authorization on the target system.

   Note: A mount user must have UID(0) or at least have READ access to the SUPERUSER.FILESYS.MOUNT resource found in the UNIXPRIV class.

   Limits: Maximum 120 characters

   Note: SMS is not mutually exclusive with non-SMS. They can both be specified (usually one or the other is specified though). This is where you specify allocation parameters for the deployment on a target system.

Storage Class

(Shared DASD only) Enter the Storage Class.

Limits: Maximum 8 characters

Example: SYSPRG

VOLSER

(Shared DASD only) Enter the Volser.

Limits: Maximum 6 characters

Example: SYSP01 and SYSP02

GIMUNZIP Volume

Enter the GIMUNZIP volume.

Limits: Maximum 6 characters

zFS Candidate Volumes

Enter zFS Candidate volumes (see page 66).

Limits: Maximum 6 characters

The zFS candidate volumes allow the specification of an optional list of VOLSERs used during the allocation of zFS container data sets for USS parts.

The new data destination appears on the Data Destination list.

Note: Click Cancel to withdraw this create request.

Add a Data Destination

You can add current data destinations to an existing system.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, or Shared DASD Clusters from the tree on the left side.

   Information about the systems related to the type you selected appears on the right side.

2. Select the system you want to add data destinations.

   Detailed information about the system appears on the right side.

3. Click the Data Destination tab.

   The Data Destination window appears.

4. Click Add.

   The Pick Data Destination dialog appears.

5. Select the data destinations you want to add and click Select.

   The data destinations are added to the system.

Maintain Data Destinations

You can maintain, delete (see page 76), or create (see page 72) data destinations.

Follow these steps:

1. Click the System Registry tab, and in the Actions section, click the Maintain Data destinations link.

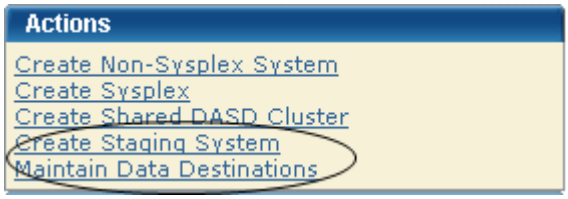

   The Maintains Data Destinations dialog appears.

   Note: A grayed select box indicates that the data destinations is assigned and cannot be removed. It can be edited.

2. Select Edit from the Actions drop-down list for the data destination you want to change.

   The Edit Data Destinations dialog appears.

   Note: The asterisk indicates that the field is mandatory.

3. Update the following and click Save:

   Name

   Enter a meaningful Name.

   Limits: Maximum 64 characters.

   Note: Each data destination name must be a unique name and it is not case-sensitive. For example DATAD1 and datad1 are the same data destination name.

   Description

   Enter the description.

   Limits: Maximum 255 characters.

   Transmission Method

   Select the transmission method.

   Default: Shared DASD.

Mount Point

(Shared DASD only) Enter the mount point directory path, which is a directory path that must exist on the target system. The user that is doing the deployment must have write permission to this directory, as well as mount authorization on the target system.

Note: A mount user must have UID(0) or at least have READ access to the SUPERUSER.FILESYS.MOUNT resource found in the UNIXPRIV class.

Limits: Maximum 120 characters

Note: SMS is not mutually exclusive with non-SMS. They can both be specified (usually one or the other is specified though). This is where you specify allocation parameters for the deployment on a target system.

Storage Class

(Shared DASD only) Enter the Storage Class.

Limits: Maximum 8 characters

Example: SYSPRG

VOLSER

(Shared DASD only) Enter the Volser.

Limits: Maximum 6 characters

Example: SYSP01 and SYSP02

GIMUNZIP Volume

Enter the GIMUNZIP volume.

Limits: Maximum 6 characters

zFS Candidate Volumes

Enter zFS Candidate volumes (see page 66).

Limits: Maximum 6 characters

The zFS candidate volumes let you specify an optional list of VOLSERs used during the allocation of zFS container data sets for USS parts.

The updated data destination appears on the list of data destinations.

Note: Click Cancel to withdraw this change request.

Set a Default Data Destination

You can set a default for a current data destination.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, or Shared DASD Clusters from the tree on the left side.

   Information about the systems you selected appears on the right side.

2. Select the system link to which you want to set the data destination default.

   Detailed information about the system appears on the right side.

3. Click the Data Destination tab.

   The Data Destination window appears.

4. Select the data destination that you want as the default.

5. In the Action field, select Set as Default.

   The word *Default* appears in the Default column.

Delete Data Destinations

You can delete current data destinations that have *not* been assigned.

Important: A grayed selection field indicates that the data destination is assigned and it cannot be deleted. The field can be edited.

Follow these steps:

1. Click the System Registry tab, and select Non-Sysplex Systems, Sysplexes, or Shared DASD Clusters from the tree on the left side.

   Information about the systems that you selected appears on the right side.

2. Select the system where you want to delete a data destination.

   Detailed information about the system appears on the right side.

3. Click the Data Destination tab.

   The Data Destination window appears.

4. Click the Select field for each data destination you want to remove, click Remove, and then click OK to confirm.

   The data destination is deleted from this system.

Remote Credentials

The Remote credentials page sets up remote credentials accounts by owner, remote user ID, and remote system name. Use the Apply button to apply and save your changes.

Important! Remote Credentials are validated during the deployment process when deploying to a nonstaging system. The user is responsible for having the correct owner, remote user ID, remote system name, password, and authenticated authorization before creating a new remote credential.

You can add (see page 77), edit (see page 78), or delete (see page 79) remote credentials.

Add Remote Credentials

Follow these steps:

1. Click the Settings tab, and select Remote Credentials from the tree on the left side.

   Detailed information appears on the right side.

2. In the Remote Credentials Accounts panel, click New.

   The New Remote Credential dialog appears.

3. Enter the following, and click OK:

   Note: The asterisk indicates that the field is mandatory.

   Remote User ID

   Enter a correct remote user ID.

   Limits: 64 characters

   Remote System Name

   Enter a remote system name.

   Limits: Eight characters

   Note: A remote credential default can be set up by creating a remote credential without the system name. This default would be for the user creating these remote credentials only.

   Password

   Enter a correct password.

   Limits: 2 to 63 characters

   Note: The password is case-sensitive. Verify that your password follows the correct case-sensitive rules for your remote system.

Confirm Password

Enter the correct confirm password.

Limits: 2 to 63 characters

Note: The password is case-sensitive. Verify that your password follows the correct case-sensitive rules for your remote system.

The remote credential entry appears on the Remote Credentials list.

4. Click Apply.

   Your changes are applied.

Edit Remote Credentials

You can edit remote credentials.

Important! Remote Credentials are validated during the deployment process when deploying to a nonstaging system. The user is responsible for having the correct owner, remote user ID, remote system name, password, and authenticated authorization before creating a new remote credential.

Follow these steps:

1. Click the Setting tab, and select Remote Credentials from the tree on the left side.

   Detailed information appears on the right side.

2. In the Actions drop-down list, click Edit for the remote credential you want to edit.

   The Edit Remote Credential window appears.

3. Update the following and click OK:

   Note: The asterisk indicates that the field is mandatory.

   Remote User ID

   Enter a correct remote user ID.

   Limits: Maximum 64 characters.

   Remote System Name

   Enter a correct remote system name.

   Limits: Maximum 8 characters.

   Example: RMinPlex

   Note: A remote credential default can be set up by creating a remote credential without the system name. This default would be for the user creating this remote credentials only.

Password

Enter a correct password.

Limits: Minimum 2 characters and Maximum 63 characters.

Note: Password is case sensitive, make sure that your password follows the correct case sensitive rules for your remote system.

Confirm Password

Enter the correct confirm password.

Limits: Minimum 2 characters and Maximum 63 characters.

Note: Password is case sensitive, make sure that your password follows the correct case sensitive rules for your remote system.

The remote credential entry appears on Remote Credentials list.

4. Click Apply

   Your changes are applied.

Delete Remote Credentials

You can delete remote credentials.

Follow these steps:

1. Click the Setting tab, and select Remote Credentials from the tree on the left side.

   Detailed information appears on the right side.

2. In the Actions drop-down list, click Delete for the remote credential you want to delete.

   A Delete Confirmation window appears.

3. Click OK.

   The remote credential is deleted.

Deploying Products

This section includes information about how to use CA MSM to deploy products.

A *deployment* is a CA MSM object that you create to deploy libraries and data sets using a process that copies target libraries defined to SMP/E and user data sets across both shared DASD and networked environments.

Deployment Status

Deployments exist in different statuses. Actions move deployments from one status to another. You can use the following available actions for each of the following deployment statuses.

Under Construction

The user is constructing the deployment.

Available Actions: All but Confirm

Snapshot in Progress

Snapshot is in Progress

Available Actions: Reset Status

Snapshot in Error

Snapshot failed

Available Actions: All but Confirm

Snapshot Completed

Snapshot Succeeded

Available Actions: Delete, Preview, Transmit, Deploy

Note: At this point, no editing, adding, or removing of products or systems is allowed.

Transmitting

The deployment archives are being transmitted using the FTP procedure.

Available Actions: Reset Status

Transmission Error

Transmission Failed

Available Actions: Delete, Preview, Transmit, Deploy

Transmitted

The deployment archives have been transmitted.

Available Actions: Delete, Preview, Deploy

Deploying

The deployment archives are being deployed.

Available Actions: Reset Status

Deploying Error

Deployment failed

Available Actions: Delete, Preview, Deploy

Deployed

The target libraries were deployed.

Available Actions: Delete, Summary, Confirm

Complete

The deployment is complete.

Available Actions: Delete, Summary

Creating Deployments

The deployment creation process consists of the following steps:

1. Initiate deployment creation (see page 82).
2. Define a name and description (see page 82).
3. Select an SMP/E environment (see page 83).
4. Select a product (see page 83).
5. Select a custom data set (see page 84).
6. Select a methodology (see page 84).
7. Select a system (see page 86).
8. Preview and save (see page 86).

Initiate Deployment Creation

You can create a new deployment by using the New Deployment wizard.

To initiate deployment creation, click the Deployments tab, and then in the Actions section, click the Create Deployment link.

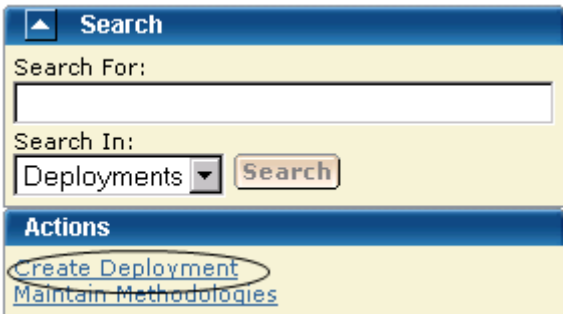

The New Deployment wizard opens to the Introduction step.

Note: When creating a deployment, you can save this deployment at any step in this wizard. This "under construction" deployment is added to the current deployments list. You can maintain this deployment (see page 87) until a successful snapshot has been created.

Define Name and Description

When you create a deployment, you begin by defining the name and description so that it will be known and accessible within CA MSM.

Note: The asterisk indicates that the field is mandatory.

Follow these steps:

1. On the Introduction step, enter a meaningful deployment name.

   Limits: Maximum 64 characters.

   Note: Each deployment name must be unique and it is not case-sensitive. For example, DEPL1 and depl1 are the same deployment name.

2. Enter the description of this deployment.

   Limits: Maximum 255 characters.

3. Click Next.

   The CSI Selection step appears.

Note: When creating a deployment, you can save this deployment at any step in this wizard. This "under construction" deployment is added to the current deployments list. You can maintain this deployment (see page 87) until a successful snapshot has been created.

Select a CSI

After you define the name and description, you select a CSI for the deployment.

Follow these steps:

1. On the CSI Selection step, in CSIs to Deploy, click the CSI you want to select.

   The CSI selections listed are preselected from the SMP/E Environments page.

2. Click Next.

   The Product Selection step appears.

Note: When creating a deployment, you can save this deployment at any step in this wizard. This "under construction" deployment is added to the current deployments list. You can maintain this deployment (see page 87) until a successful snapshot has been created.

Select a Product

After you select a CSI for the deployment, you select a product for the deployment.

Follow these steps:

1. On the Product Selection step, select a product from the list.

   Note: If you cannot select the product or product feature from the list, it is for one of the following reasons:

   - The product or feature is not deployable for the selected CSI.
   - The product feature is part of a product that you must select first.

   If a feature is mandatory for the selected product, the corresponding check box is also selected and disabled, and you cannot deselect the feature from the list.

2. If there is a  text icon in the Text column, click it to read the instructions supplied by CA Support for product, data set, and other necessary information.

3. Click the check box *I have read the associated text*, and click Next. The Next button is disabled until you click the check box.

   Note: If there are no products displayed, the appropriate PTF that enables your products' deployment through metadata has not been installed.

   The Custom Data Sets step appears.

Note: When creating a deployment, you can save this deployment at any step in this wizard. This "under construction" deployment is added to the current deployments list. You can maintain this deployment (see page 87) until a successful snapshot has been created.

Select a Custom Data Set

A *custom data set* is a data set that contains either a z/OS data set or USS path.

Follow these steps:

1. On the Custom Data Sets step, select a custom data set from the list and click Select.

   Note: To add a new custom data set, click Add Data Set and enter the custom data set information (see page 99).

2. Click Next.

   The Methodology Selection step appears.

Note: When creating a deployment, you can save this deployment at any step in this wizard. This "under construction" deployment is added to the current deployments list. You can maintain this deployment (see page 87) until a successful snapshot has been created.

More information:

Add a Custom Data Set (see page 99)

Select a Methodology

After you select a custom data set, you select a methodology, which lets you provide a single data set name mask that is used to control the target library names on the target system.

Follow these steps:

1. On the Methodology Selection step, select a Methodology from the list.

2. (Optional) Click the Create button and enter the new methodology information (see page 106).

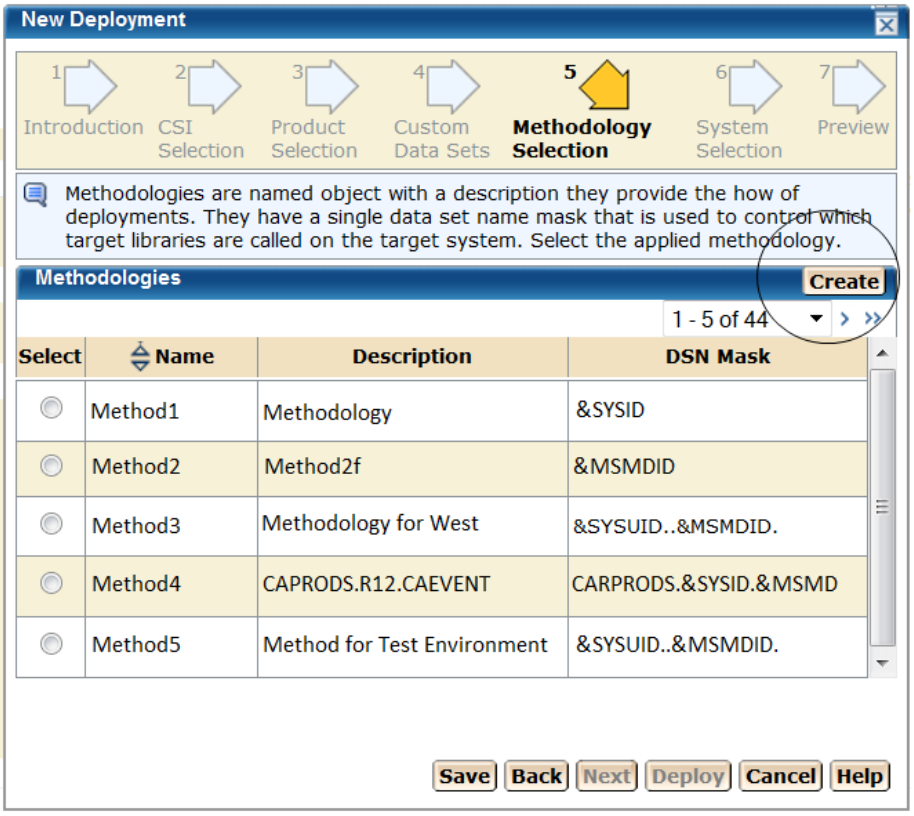

3. Click Next.

   The System Selection step appears.

Note: When creating a deployment, you can save this deployment at any step in this wizard. This "under construction" deployment is added to the current deployments list. You can maintain this deployment (see page 87) until a successful snapshot has been created.

More information:

Create a Methodology (see page 106)

Select a System

After you select a methodology, you select a system.

Follow these steps:

1. On the System Selection step, select the systems to be deployed.

   Note: When two systems have the same name, use the description to differentiate between these systems.

   Sysplex systems are denoted by *sysplex system*:*system name*. For example, PLEX1:CO11, where PLEX1 is the sysplex system, and CO11 is the system name.

2. Click Next.

   The Preview step appears.

Note: When creating a deployment, you can save this deployment at any step in this wizard. This "under construction" deployment is added to the current deployments list. You can maintain this deployment (see page 87) until a successful snapshot has been created.

Preview and Save the Deployment

After you select a system, you are ready to preview the deployment, and then save or deploy it.

- To save the deployment, click Save.
- To set up the deployment, click Deploy.

Note: Click Cancel to exit the wizard without saving.

The Preview identifies the deployment and describes the products, systems, means of transport, and target libraries (including source, target, and resolution), as well as the SMP/E environment and snapshot information.

Important! Data sets may need to be APF-authorized and added to the Link List and Link Pack Area. These data sets are identified in this dialog.

Note: ??? in the Preview indicates that CA MSM has yet to assign this value.

View a Deployment

To view a deployment, click the Deployments tab, and select the current or completed deployment from the tree on the left side. The detailed deployment information appears on the right side.

Change Deployments

You can change deployments any time before you snapshot the deployment.

Important! Each deployment must have at least one product defined, at least one system defined, and a methodology defined.

Follow these steps:

1. Click the Deployments tab. The Deployment window appears.
2. On the right, in the Deployments panel click the current deployment link.

   The detailed deployment information appears.
3. Click the Deployment Name link for the Deployment you want to change.

   This deployment's window appears.

   Change the information on this window as needed. Each deployment name must be unique and it is not case-sensitive. For example DEPL1 and depl1 are the same deployment name.

   Note: The methodology provides the means for deployment. It is used to control the target library names on the target system.

   There are actions that you can perform based on Deployment State (see page 80).
4. To change a methodology, select a methodology from the drop-down list and click Edit.

   The Edit Methodology window (see page 119) appears. The Deployment ID is the value of the MSMID variable.

   Note: You can perform the following actions:

   - You can select (see page 96), add (see page 97), or remove (see page 97) a product.
   - You can select (see page 123), add (see page 123), or remove (see page 124) a system.
   - You can select (see page 98), add (see page 99), or remove (see page 105) a custom data set.
5. Click Save on the Deployment Details window.

6. Click Actions drop-down list to do one of the following:

Preview (Summary)

Note: This action button changes to Summary after a successful deploy.

Generates a list of the following current information:

- Deployment's ID
- Name
- Products
- Systems
- Transport information
- Target libraries including: source, target, and resolved data set names.
- SMP/E environment
- Snapshot path and container

Snapshot

Takes a snapshot of the current deployment.

A *snapshot* of the set of target libraries is taken by CA MSM, by utilizing the IBM supplied utility GIMZIP to create a compressed archive of these libraries, along with a list of applied maintenance. The SMP/E environment is "locked" during this archive creation process to insure the integrity of the archived data.

Transmit

Transmit enables a customer to take their CA MSM installed software and copy it onto systems across the enterprise through FTP, in preparation for a subsequent deployment.

Deploy

Combines the snapshot, transmit, and deploy action into one action.

Confirm (see page 94)

Confirms that the deployment is complete. This is the final action by the user.

Note: A deployment is not completed until it is confirmed. Once it is confirmed the deployment moves to the Confirmed deployment list.

Delete

Deletes deployment and its associated containers, folders, and files. This does not include the deployed target libraries on the end systems. See delete a deployment for a list of deleted files.

Note: A deployment's deletion does not start until it is confirmed.

Reset Status (see page 92)

You can reset a deployment status when the deployment has a status of *snapshot in progress*, *transmitting*, or *deploying*. See reset status (see page 92) for a list of deleted files.

7. Click Save on the Deployment Details window.

   You changes are saved.

More information:

Add a Product (see page 97)
Add a System (see page 123)
Remove a Product (see page 97)
Remove a System (see page 124)
View the Product List (see page 96)
View a System List (see page 123)
Edit a Methodology (see page 119)
Confirm a Deployment (see page 94)

Deployment Maintenance

You can maintain a deployment in the following ways:

- Adding
  - System (see page 123)
  - Product (see page 97)
  - Custom data sets (see page 99)
- Delete
  - Deployment
- Removing
  - System (see page 124)
  - Product (see page 97)
  - Custom data sets (see page 105)

- Editing
  - Maintain deployments (see page 87)
  - Edit a custom data set (see page 102)
  - Edit a methodology (see page 119)
- Viewing
  - System (see page 123)
  - Product (see page 96)
  - Custom data sets (see page 98)

Failed Deployments

When a deployment fails, you investigate, correct, and deploy again. Use the following procedures in this section:

- Investigate a Failed Deployment Using the Tasks Page (see page 90)
- Download a Message Log (see page 64)
- Save a Message Log as a Data Set (see page 65)
- View Complete Message Log (see page 65)

Note: A deployment is processed in steps and in order as listed in the Deployment window. Each step must pass successfully before the next step is started. If a step fails, the deployment fails at that step, and all steps after the failed step are not processed.

More information:

Download a Message Log (see page 64)
Save a Message Log as a Data Set (see page 65)
View Complete Message Log (see page 65)

Investigate a Failed Deployment

When a deployment fails, you investigate, correct, and deploy again.

Follow these steps:

1. On the Deployments Page, in the left hand column, find the deployment with an error and note its name.
2. Click the Tasks tab and then click Task History.

   Note: Click Refresh on the right hand side of the Task History bar to refresh the Task History display.

3. At the Show bar, select All tasks, or select My tasks to only see the tasks assigned to you.

   Note: You can refine the task list further by selecting task and status types from the drop-down lists, and then sort by Task ID.

4. Find the failed deployment step and click the link in the Name column.

   The Task Output Browser appears.

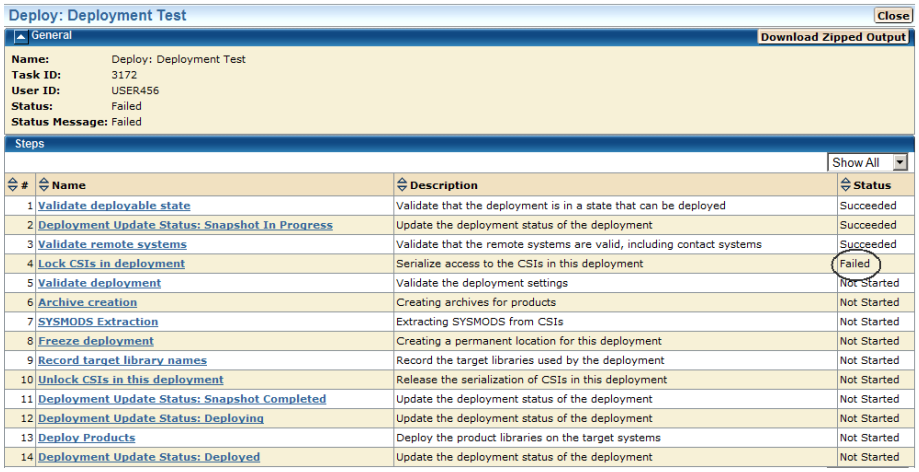

#	Name	Description	Status
1	Validate deployable state	Validate that the deployment is in a state that can be deployed	Succeeded
2	Deployment Update Status: Snapshot In Progress	Update the deployment status of the deployment	Succeeded
3	Validate remote systems	Validate that the remote systems are valid, including contact systems	Succeeded
4	Lock CSIs in deployment	Serialize access to the CSIs in this deployment	Failed
5	Validate deployment	Validate the deployment settings	Not Started
6	Archive creation	Creating archives for products	Not Started
7	SYSMODS Extraction	Extracting SYSMODS from CSIs	Not Started
8	Freeze deployment	Creating a permanent location for this deployment	Not Started
9	Record target library names	Record the target libraries used by the deployment	Not Started
10	Unlock CSIs in this deployment	Release the serialization of CSIs in this deployment	Not Started
11	Deployment Update Status: Snapshot Completed	Update the deployment status of the deployment	Not Started
12	Deployment Update Status: Deploying	Update the deployment status of the deployment	Not Started
13	Deploy Products	Deploy the product libraries on the target systems	Not Started
14	Deployment Update Status: Deployed	Update the deployment status of the deployment	Not Started

5. Click the link in the Name column to view the results, and click on the messages logs to review the details for each error.

   Note: You can analyze the error results and determine the steps required to troubleshoot them.

6. Correct the issue and deploy again.

More information:

Download a Message Log (see page 64)
Save a Message Log as a Data Set (see page 65)
View Complete Message Log (see page 65)

Download a Message Log

You can save the message log in the following ways:

- To download a zipped file of all the text messages for this validation, click the Deployment Name on the top left tree. Click the Download Zipped Output button on the General menu bar. Save this file.
- To download as TXT, click the Deployment Name or the Deployment Results on the left tree. Click the Action button on the Message Log bar and click the Download as TXT. Save this file.

- To download as ZIP, click the Deployment Name or the Deployment Results on the left tree. Click the Action button on the Message Log bar and click the Download as ZIP. Save this file.

Save a Message Log as a Data Set

You can save a message log as a data set.

Follow these steps:

1. Click the Deployment Name or the Deployment Results on the left tree. Click the Action button on the Message Log bar, and click the Save as Data Set.

   The Save Output as Data Set dialog appears.

   Note: This information is sent to CA Support to analyze the failed deployment.

   Note: The asterisk indicates that the field is mandatory.

2. Enter the following information and click OK:

   Data Set Name

   Enter a data set name. CA MSM generates a value.

   VOLSER

   For non-SMS data, enter the Volser.

   Example:

   Volser: SYSP01 and SYSP02

   Storage Class

   For SMS Allocation data, enter the Storage Class.

   The message log is saved as a data set.

View Complete Message Log

To view the complete message log for a failed validation, click Show All.

Note: To close the message log, click Close.

Reset Deployment Status

You can reset a deployment status when the deployment has a status of *snapshot in progress*, *transmitting*, or *deploying*. The message log explains if any containers, folders, and files were deleted during reset.
You can also investigate a failed deployment (see page 63) to see additional details in the message log.
The following statuses may be reset.

Snapshot in progress

Snapshot in progress is reset to *snapshot in error.*

Transmitting

Transmitting is reset to *transmit in error*.

Deploying

Deploying is reset to *deploy in error*.

The following artifacts are reset by status.

Snapshot in Progress

Archive located at Application Root/sdsroot/Dnnnn, where nnnn = Deployment ID automatic number. Application Root is defined in settings under mount point management,

Temp files located at Application Root/sdsroot/Deployment_nnnn, where nnnn = Deployment ID automatic number.

Transmit in Progress

Nothing is reset.

Deploy in Progress

Nothing is reset.

Delete a Deployment

You can delete deployments.

Note: You cannot delete deployments that are currently being deployed.

A deployment deletion must be confirmed before a deletion starts.

Note: If system information was changed, not all files may be deleted. In this case, you many need to delete these files manually. For example, if an FTP transmission was changed to a Shared DASD Cluster or if the remote credentials are incorrect or changed.

The message log explains which containers, folders, and files were deleted during processing and which ones were not deleted. See how to investigate a failed deployment (see page 63) for details on finding the message log.

Note: Target libraries are never deleted.

The following artifacts are deleted by status:

Under Construction

All applicable database records

Snapshot in Error

All applicable database records

Snapshot Completed

Archive located at Application Root/sdsroot/D*nnnn* where *nnnn* = Deployment ID automatic number. Application Root is defined in settings under mount point management.

All applicable database records.

Transmit in Error

Same as Snapshot Completed, plus attempts to delete any transmitted snapshots on target systems.

Transmitted

Same as Transmit in Error.

Deploy in Error

Same as Transmitted.

Deployed

Same as Snapshot Completed.

Complete

Same as Snapshot Completed.

Follow these steps:

1. Click the Deployments tab.

   The Deployment window appears.

2. On the right, in the Deployments panel, click the Current Deployments or Complete Deployments link.

   The detailed deployment information appears.

3. Click the deployment name link, and from the Actions drop-down list, select Delete, and then click OK to confirm.

   The deployment is deleted.

Confirm a Deployment

You can use this procedure to confirm that the deployment is complete.

Note: A deployment is not completed until it is confirmed. After it is confirmed, the deployment moves to the Completed deployment list.

Important! Data sets may need to be APF-authorized and added to the Link List and Link Pack Area. These data sets are identified in this dialog.

Follow these steps:

1. Click the Deployments tab.

   The Deployment page appears.

2. Click Confirm.

   The Confirmation dialog appears.

3. Review the confirmation.

4. Click OK when the deployment is correct.

   Note: Click Cancel to exit this procedure without confirming.

   The Deployment Summary window may contain the following:

   - Deployment's ID
   - Name
   - Products
   - Systems
   - Data Sets actions
   - Transport information
   - Target libraries including: source, target, and resolved data set names.
   - SMP/E environment
   - Snapshot path and container

The following example shows the Data Sets actions, Transport, and Target libraries information.

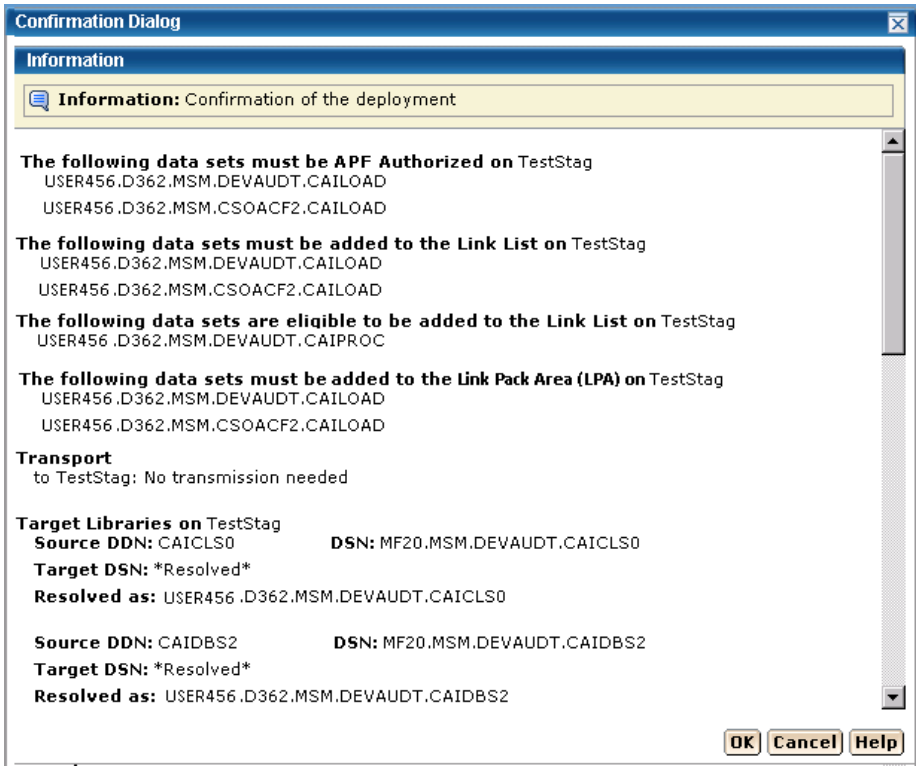

Products

You can view, add, and remove products from a deployment.

View the Product List

You can view a product.

Follow these steps:

1. Click the Deployments tab.
2. Select the current deployment from the tree on the left side.

   The detailed deployment information appears on the right side.

Add a Product

You can add a product to a deployment.

Follow these steps:

1. Click the Deployments tab. The Deployments window appears.
2. On the right, in the Deployments panel click the Current Deployment link.

   A list of current deployments appears.
3. Click the deployment name link.
4. In the Product List panel click Add Products.

   The Add Products wizard appears.
5. Select a CSI and click Next.

   The Product Selection appears.
6. Select a product.
7. If there is a  text icon in Text column, click the text icon to read the instructions supplied by CA Support for product, data sets, and other necessary information.
8. Click the "I have read the associated text by selecting the text icon from the list about" box. This box appears only if there is a text icon.

   Note: You will not be able to click Next until you click this box.
9. Click Next.

   The Custom Data Set Selection appears
10. If needed, select or add a custom data set (see page 99).
11. Click Add Products.

    The Product is added.

Remove a Product

You can remove a product from a deployment.

Note: This product will no longer be associated with the current deployment.

Follow these steps:

1. Click the Deployments tab. The Deployment window appears.
2. On the right, in the Deployments panel click the Current Deployment link.

   A list of current deployments appears.
3. Select the deployment that you want to remove the product from.

4. In the Product List panel, select a product to remove.
5. Click the Remove link.
6. Click OK to the Remove Products confirmation window.

   The product is removed.

Custom Data Sets

You can view, add (see page 99), edit (see page 102), and remove (see page 105) custom data sets from a deployment.

A *custom data set* is a data set that contains either a z/OS data set or USS path.

- For a z/OS data set, you need to provide a data set name that is the actual existing z/OS data set and a mask that names the data set on the target system. This mask may be set up using symbolic qualifiers (see page 109) and must be available to CA MSM. During the deployment process, the custom data set is accessed and copied to the target system the same way a target library is accessed and copied.
- For USS parts, you need to provide a local path, a remote path (which may be set up using symbolic qualifiers (see page 109)), and a type of copy. The type of copy can be either a container copy or a file-by-file copy.

View Custom Data Sets

You can view custom data sets.

Follow these steps:

1. Click the Deployments tab, and select the current deployment from the tree on the left side.

   The detailed deployment information appears on the right side.

   Product Name Sort Arrows

   Click the up arrow to place the product names in alphabetic order or click the down arrow to place them in reverse alphabetic order.

Add a Custom Data Set

You can add custom data sets to a deployment.

Follow these steps:

1. Click the Deployments tab.

   The Deployments window appears.

2. On the right, in the Deployments panel, click the Current Deployment link.

   A list of current deployments appears.

3. Click the deployment name link.

4. In the Custom Data Sets List panel, click Add Data Sets.

   The Add Custom Data Sets dialog appears.

   Note: The asterisk indicates that the field is mandatory.

5. Select a Product from the drop-down list.

   Note: When there are instructions, they are required and supplied by CA Support.

6. Select the Data Set Type, either data set (step 7) or USS (step 10).

   Default: data set

7. For data set, enter the data set name.

   Limits: Maximum 44 characters.

   Note: This is the existing z/OS data set name that you want CA MSM to include in the deployment when it is deployed on the target systems.

8. Enter the data set name mask, click the file icon, and select a symbolic name (see page 109).

   Mask

   This is the mask that will be used to name the data sets that are being deployed. They can contain symbolic qualifiers (see page 109). For example, if you enter CAPRODS.&SYSID, the &SYSID is replaced by its values, and if the SYSID that is being deployed to is XX16, the DSN mask will be CAPRODS.XX16

   Limits: Maximum 64 characters.

   Note: Each deployed target data set is named using the resolved content of the data set name mask followed by the low-level qualifier of the source data set. Appending the low-level qualifier from the source data set helps ensure uniqueness of the final data set name. Verify that the mask that you entered does not exceed 35 characters when it is translated.

The mask consists of one or more qualifiers that are separated by periods. The maximum number of characters is 64, including the periods. While you are entering the mask, CA MSM validates the mask by replacing symbolics with the minimum possible values first, and then with the maximum values. If the validation with the minimum possible values fails, an error message appears at the top of the dialog, and you cannot proceed. If the validation with the maximum values fails, a warning message appears, and you can proceed.

When the mask is translated, it has a maximum length of 44 characters including the periods and the low-level qualifier from the source data set. The low-level qualifier from the source data set has a maximum length of nine characters including a period.

-

Two consecutive periods are required to separate the two masks.

9. Enter the Mask and click OK.
10. For USS data set type, enter the Local Path. The local path is the directory are where files are to be copied from.

    Limit: Maximum 255 characters.

    Note: The asterisk indicates that the field is mandatory.
11. Enter the Remote Path and/or click the file icon and select a symbolic name (see page 109). The remote path is the path where the files are to be copied to.

    Limit: Maximum 255 characters.
12. Select the Type of Copy:

    - If you select Container Copy, proceed to step 14.
    - If you select File-by-file Copy, proceed to step 15, and ensure that the USS path exists on all of the remote systems of this deployment, and that there is sufficient space to hold these target libraries.

    Default: Container Copy
13. Click OK.
14. For Container Copy, enter the container name and/or click the file icon and select a symbolic name (see page 109).

    Limit: Maximum 64 characters.

    Note: It consists of one or more qualifiers separated by periods, and has a maximum input length of 64 characters, including the periods. When it is translated, it has a maximum length of 44 characters, including the periods.

Note: For Container Copy, the following occurs during the deployment process:

a. A file system of the requested type is created.

b. The size of the file system is computed as follows:

- The size off all of the constituent files and directories in the local path are added up as bytes.
- These bytes are converted to tracks and used as the primary allocation value.
- If there is a non-zero percent of free space entered, it is used to calculate the secondary allocation.

c. All of the directories in the mount point are dynamically created.

d. The file system is mounted at the requested mount point.

Note: The mount is not permanent. You will need to update your BPXPARMS to make this mount point permanent.

e. The content from the local path is copied into the newly created and mounted file system.

Note: The asterisk indicates that the field is mandatory.

15. Select the Type of Container from the drop-down list.

16. Enter the Mount Point and/or click the file icon and select a symbolic name (see page 109).

Limit: Maximum 255 characters.

Note: The container is created and it is mounted at a position in the USS file system hierarchy. The place in the hierarchy where it is mounted is known as that containers mount point. Most nodes in the USS file system can be mount points, for any one container.

17. Enter the percentage of Free Space needed.

The percentage of free space is the amount of space to leave in the file system, after the size has been computed. This is done by specifying secondary space on the allocation. For example, the computed space was determined to be 100 tracks. Then 35 would be 35% free space and the space allocations would be in tracks, 100 primary 35 secondary. While 125 would be 125% over and allocation would be in tracks, 100 primary 125 secondary.

Limit: 0 to 1000.

18. Click OK.

The custom data set is added.

Edit a Custom Data Set

You can edit a custom data set.

Follow these steps:

1. Click the Deployments tab.

   The Deployments page appears.

2. On the right, in the Deployments panel, click the Current Deployment link.

   A list of current deployments appears.

3. Click the deployment name link.

4. In the Custom Data Sets List panel, click the Actions drop-down list and click Edit.

   The Edit Custom Data Sets dialog appears.

   Note: The asterisk indicates that the field is mandatory.

5. Select a Product from the drop-down list.

   Note: When there are instructions, they are required and supplied by CA Support.

6. Select the Data Set Type, either data set (step 7) or USS (step 10).

   Default: data set

7. For data set, enter the data set name.

   Limits: Maximum 44 characters.

   Note: This is the existing z/OS data set name that you want CA MSM to include in the deployment when it is deployed on the target systems.

8. Enter the data set name mask, click the file icon, and select a symbolic name (see page 109).

   Mask

   This is the mask that will be used to name the data sets that are being deployed. They can contain symbolic qualifiers (see page 109). For example, if you enter CAPRODS.&SYSID, the &SYSID is replaced by its values, and if the SYSID that is being deployed to is XX16, the dsn mask will be CAPRODS.XX16

   Limits: Maximum 64 characters.

   Note: Each deployed target data set is named using the resolved content of the data set name mask followed by the low-level qualifier of the source data set. Appending the low-level qualifier from the source data set helps ensure uniqueness of the final data set name. Verify that the mask that you entered does not exceed 35 characters when it is translated.

The mask consists of one or more qualifiers that are separated by periods. The maximum number of characters is 64, including the periods. While you are entering the mask, CA MSM validates the mask by replacing symbolics with the minimum possible values first, and then with the maximum values. If the validation with the minimum possible values fails, an error message appears at the top of the dialog, and you cannot proceed. If the validation with the maximum values fails, a warning message appears, and you can proceed.

When the mask is translated, it has a maximum length of 44 characters including the periods and the low-level qualifier from the source data set. The low-level qualifier from the source data set has a maximum length of nine characters including a period.

-

Two consecutive periods are required to separate the two masks.

9. Enter the Mask and click OK.
10. For USS data set type, enter the Local Path. The local path is the directory where files are to be copied from.

    Limit: Maximum 255 characters.

    Note: The asterisk indicates that the field is mandatory.
11. Enter the Remote Path and/or click the file icon and select a symbolic name (see page 109). The remote path is the path were the files are to be copied to.

    Limit: Maximum 255 characters.
12. Select the Type of Copy:
    - If you select Container Copy, proceed to step 14.
    - If you select File-by-file Copy, proceed to step 15, and ensure that the USS path exists on all of the remote systems of this deployment, and that there is sufficient space to hold these target libraries.

    Default: File-by-file Copy
13. Click OK.
14. For Container Copy, enter the container name and/or click the file icon and select a symbolic name (see page 109).

    Limit: Maximum 64 characters.

    It consists of one or more qualifiers separated by periods, and has a maximum input length of 64 characters, including the periods. When it is translated is has a maximum length of 44 characters including the periods.

For container copy the following occurs during the deployment process:

a. A file system of the requested type is created

b. The size of the file system is computed as follows:

- The size off all of the constituent files and directories in the local path are added up as bytes.
- These bytes are converted to tracks and used as the primary allocation value
- If there is a non-zero percent of free space entered, it is used to calculate the secondary allocation.

c. All of the directories in the mount point will by dynamically created.

d. The file system will be mounted at the requested mount point

Note: The mount is not permanent. You will need to update your BPXPARMS to make this mount point permanent.

e. The content from the local path will copied into the newly created and mounted file system.

Note: The asterisk indicates that the field is mandatory.

15. Select the Type of Container from the drop down list.

16. Enter the Mount Point and/or click the file icon and select a symbolic name (see page 109).

Limit: Maximum 255 characters.

Note: The container is created and it is mounted at a position in the USS file system hierarchy. The place in the hierarchy where it is mounted is known as that containers mount point. Most nodes in the USS file system can be mount points, for any one container.

17. Enter the percentage of Free Space needed.

The percentage of free space is the amount of space to leave in the file system, after the size has been computed. This is done by specifying secondary space on the allocation. For example, the computed space was determined to be 100 tracks. Then 35 would be 35% free space and the space allocations would be in tracks, 100 primary 35 secondary. While 125 would be 125% over and allocation would be in tracks, 100 primary 125 secondary.

Limit: 0 to 1000.

18. Click OK.

The custom data set is changed.

Remove a Custom Data Set

You can remove a custom data set from a deployment.

Note: This data set will no longer be associated with the current deployment.

Follow these steps:

1. Click the Deployments tab.

   The Deployment page appears.

2. On the right, in the Deployments panel click the Current Deployment link.

   A list of current deployments appears.

   Product Name Sort Arrows

   Click the up arrow to place the product names in alphabetic order or click the down arrow to place them in reverse alphabetic order.

3. Select the custom data set that you want to remove from this deployment.
4. Click the Remove link.
5. Click OK to the Remove Custom Data Set confirmation window.

   The custom data set is removed.

Methodologies

You can create (see page 106), maintain, edit (see page 119), and delete (see page 121) methodologies from a deployment.

A methodology has the following attributes:

- A single data set name mask that is used to control what target libraries are to be called on the target systems and where these deployment will go.

  z/OS data sets

  z/OS data sets use a data set name mask. The data set name mask is a valid data set name comprised of constants and symbolic qualifiers (see page 109).

  The minimum methodology data consists of a data set mask and a target action. The symbolics in the data set mask are either symbolics defined by CA MSM or z/OS system symbolics.

- Deployment Style information is used to *create only* or *create and replace* a methodology.

  Create Only

  Use C*reate Only* when you are creating a new methodology that does not have any target libraries already associated with a deployment.

  Create or Replace

  Use *Create or Replace* to:

  - Create new data sets and/or files in a UNIX directory.
  - Replace existing sequential data sets or files in a UNIX directory.
  - For partitioned data sets, replace existing members, add new member without deletion of members that are not replaced.

  Note: Using *Create or Replace* would not cause the deployment to fail due to data set name conflicts.

Create a Methodology

You can create a methodology.

Note: The asterisk indicates that the field is mandatory.

Follow these steps:

1. Click the Create button, in the Methodology Selection in the New Deployment wizard.

   The Create a New Methodology dialog appears.

2. Enter the methodology name.

   Limits: Maximum 64 characters.

   Note: Each methodology name must be unique and it is not case-sensitive. For example Meth1 and meth1 are the same methodology name.

3. Enter the description of this methodology.

   Limits: Maximum 255 characters.

4. Enter the data mask name, click the file icon, and select a symbolic name (see page 109).

   Data Set Name Mask

   This is the mask that will be used to name the data sets that are deployed. They can contain symbolic qualifiers (see page 109). For example, assume you enter, CAPRODS.&SYSID. In this case, the &SYSID. will be replaced by its values. If the SYSID that is being deployed to is X16, the DSN mask will be: CAPRODS.X16

   Limits: Maximum 64 characters.

   Note: Each deployed target data set is named using the resolved content of the data set name mask followed by the low-level qualifier of the source data set. Appending the low-level qualifier from the source data set helps ensure uniqueness of the final data set name. Verify that the mask that you entered does not exceed 35 characters when it is translated.

   The mask consists of one or more qualifiers that are separated by periods. The maximum number of characters is 64, including the periods. While you are entering the mask, CA MSM validates the mask by replacing symbolics with the minimum possible values first, and then with the maximum values. If the validation with the minimum possible values fails, an error message appears at the top of the dialog, and you cannot proceed. If the validation with the maximum values fails, a warning message appears, and you can proceed.

   When the mask is translated, it has a maximum length of 44 characters including the periods and the low-level qualifier from the source data set. The low-level qualifier from the source data set has a maximum length of nine characters including a period.

   ■

5. Select a style of Deployment.

Create only

Creates new data sets.

Note: Prior to creating any data sets on the remote system, a check is made, to see if the data sets already exist. The deployment is not allowed to continue if this occurs.

Create or Replace

Creates new data sets if they do not already exist, or replaces existing data sets.

Partitioned data set

Replaces existing members in a partitioned data set with members that have the same name as the source file. Any currently existing member that is not in the source file will remain in the PDS. Any member from the source that does not already exist in the target PDS will be added to the target PDS.

The amount of free space in the PDS will need to be sufficient to hold the additional content, since no automatic compress will be done.

Directory in a UNIX file system

Replaces files in a directory with files with the same name as the source. Any currently existing directory in a UNIX file system that is not in the source will remain in the UNIX file system.

Sequential data set or a file in the UNIX file system

Replaces the existing data set or file and its attributes with the data from the source file.

For a VSAM data set (cluster)

Populates an existing VSAM cluster with the data from the source file.

Note: The existing VSAM cluster must be of the same type as the source cluster (ESDS, KSDS, LDS, or RRDS), and it must have characteristics that are compatible with the source cluster (such as, record size, key size, and key offset). Replace does not verify the compatibility of these characteristics.

To replace the contents of an existing cluster, the cluster is altered to a reusable state by using an IDCAMS ALTER command, if necessary, before the data from the VSAM source is copied into the cluster by using an IDCAMS REPRO command. The REPRO command will use both the REPLACE and REUSE operands. Following the REPRO operation, the cluster is altered back to a non-reusable state if that was its state to begin with.

6. Click Save.

   The methodology is saved.

   Note: Click Cancel to close this dialog without saving.

Symbolic Qualifiers

The data set name mask and the directory path contain the following symbolic qualifiers:

Data Set Name Mask

This is a unique name that identifies each data set. It consists of one or more qualifiers separated by periods, and has a maximum input length of 64 characters, including the periods. When the data set name mask is translated is has a maximum length of 44 characters including the periods.

Directory Path

This is a USS path name, it consists of one or more directory leaves separated by forward slashes, and has a maximum input length of 255 characters including slashes. When the Directory Path is translated it has a maximum length of 255 characters.

Symbolic Substitution

Symbolic substitution, or translation, is a process performed by CA MSM to resolve the mask values specified in the data set name mask and directory path, into real names based upon the contents of the symbolic variables at translation time. A CA MSM symbol is defined in the list of symbols. Each symbol begins with an ampersand (&) and ends with a period (.). For example, the symbol &LYYMMDD. would be completely replaced with its value at translation time, including the ampersand and trailing period. The trailing period is important and is considered part of the symbolic name.

Symbolic Variables

You can use symbolic variables in the construction of a data set name with the value of the symbolic variable to end a data set name segment.

Example: Assume MSMDID is 255.

SYSWORK.D&MSMDID..DATASET

Note: The double periods are necessary because the first period is part of the symbolic name, and therefore does not appear in the translated value.

The final data set name is SYSWORK.D255.DATASET.

Numeric Values

Some CA MSM symbolic names translate to numeric values. In the case where you want to use one of these symbolic variables in your data set name, you may have to precede it with a alpha constant. This is because z/OS data set naming rules do not allow a data set name segment to start with a numeric.

If you wanted to use a date value in your translated data set name, you could use one of the CA MSM defined date symbolic qualifiers such as &LYYMMDD. You must be careful how you construct the data set mask value.

Example: Assume that you want to have a middle level qualifier to have a unique value based upon the date of April 1, 2010.

Mask = SYSWORK.D&LYYMMDD..DATASET, translates to SYSWORK.D100401.DATASET

An incorrect specification of the mask would be:

SYSWORK.&LYYMMDD..DATASET, translates to SYSWORK.100401.DATASET. Because the middle-level qualifier starts with a numeric it is an invalid data set name.

Directory Paths

Symbolic substitution works in the same logical way for directory paths. However, directory paths do not typically have periods in them, so you will typically not see the double dots in directory paths.

Example: Assume the target system is SYSZ.

/u/usr/&MSMSYSNM./deployments translates to /u/usr/SYSZ/deployments.

Preview Example

Note: Before a Product Deployment is deployed, the MSMDID shows as ???. After deployment, the Automatic ID is assigned and this is the MSMDID.

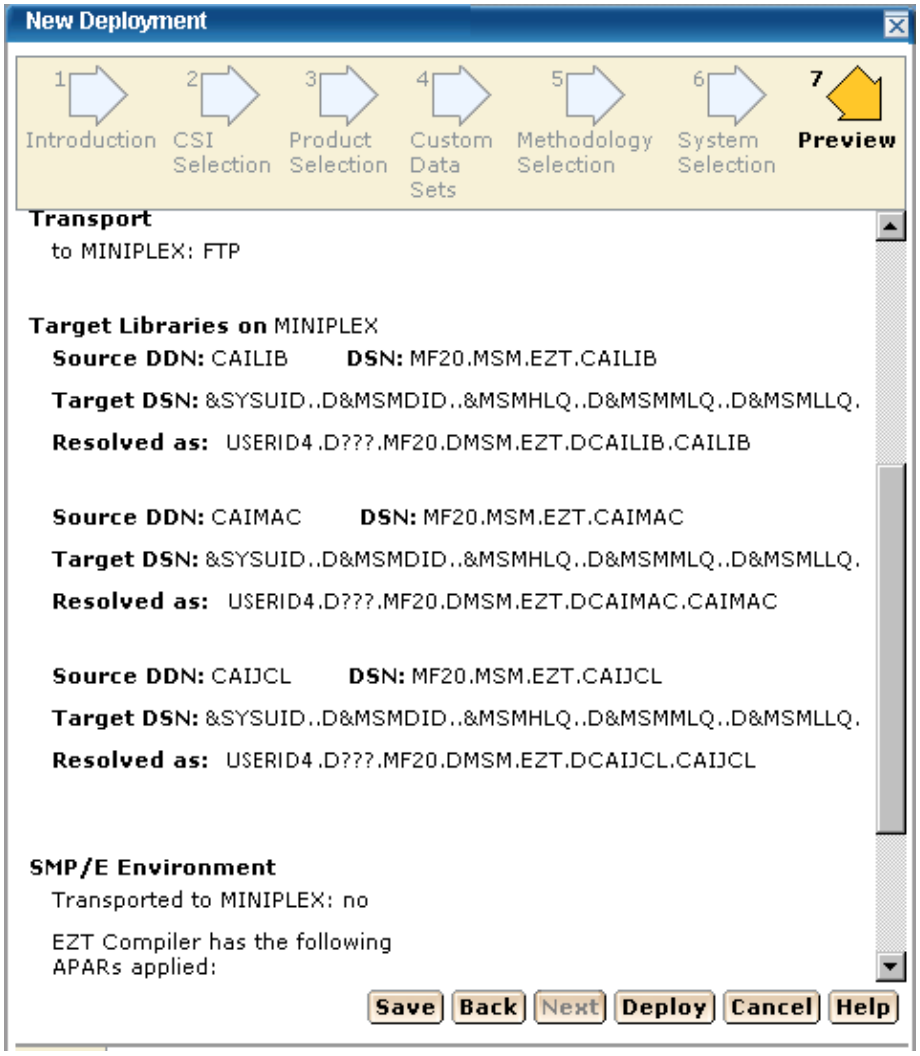

Symbolic Qualifiers

ID and System Information

MSMDID

This is the CA MSM deployment ID.

Limits: This is automatically assigned by CA MSM when the Deploy button is clicked or when a deployment is saved.

MSMMPN

This is the CA MSM Mount Point Name. The value is entered into the mount point name field when adding a custom data set (see page 99) with both the USS radio button and the Container copy radio button set. It is of primary value in remote path.

Note: The Mount Point Name field can contain symbols when it is translated first, the value of the MSMMPN. variable is resolved.

Example: Assume the value of MSMDID is 253 and the user entered the following information.

Mount point name: /u/users/deptest/R&MSMDID./leaf

Remote path: &MSMMPN.

The translated value of &MSMMPN is /u/users/deptest/R253/leaf

MSMSYSNM

This is the CA MSM system object name.

SYSCLONE

This is the shorthand name of the system.

Limits: Maximum 2 characters.

SYSNAME

This is the system name entered when a non-sysplex, sysplex, Shared DASD Cluster, or Staging system is created.

SYSPLEX

This is the system name entered when a sysplex is created.

Note: This symbolic may not be used for a non-sysplex system.

SYSUID

The current user ID.

Target Libraries

MSMHLQ

MSMHLQ is the high-level qualifier for the target library.

Limits: It is the characters before the first period in a fully qualified data set name. The high-level qualifier can be from 1 to 8 characters.

Example: For the data set JOHNSON.FINANCE.DIVISION.SCRIPT, the high-level qualifier is JOHNSON.

MSMMLQ

MSMMLQ is the middle-level qualifier for the target library.

Limits: It is the characters after the first period and before the last period in a fully qualified data set name. The middle-level qualifier size can vary based on the number of qualifiers defined.

Example: For the data set JOHNSON.FINANCE.DIVISION.SCRIPT, the middle-level qualifier is FINANCE.DIVISION.

MSMLLQ

MSMLLQ is the low-level qualifier for the target library.

Limits: It is the characters after the last period in a fully qualified data set name. The low-level qualifier can be from 1 to 8 characters.

Example: For the data set JOHNSON.FINANCE.SCRIPT, the low-level qualifier is SCRIPT.

MSMSLQ

This is the secondary low-level qualifier for the target library and it is the "segment" of the data set name just before the low-level qualifier (MSMLLQ).

Limits: It is the characters after the second to last period and before the last period in a fully qualified data set name. The secondary low-level qualifier can be from 1 to 8 characters.

Example: For the data set JOHNSON.FINANCE.SECOND.SCRIPT, the low-level qualifier is SECOND.

MSMPREF

This is the target library prefix. The target library prefix is the entire data set name to the left of the MSMLLQ.

Example: For the data set JOHNSON.FINANCE.DIVISION.SCRIPT the prefix is JOHNSON.FINANCE.DIVISION.

MSMDLIBN

The deployed library number is a unique number, for each deployed library, within a deployment.

Example: Assume 3 target libraries in a deployment.

```
DSN = USER456.LIBR473.CAIPROC
DSN = USER456.LIBR473.CAILOAD
DSN = USER456.LIBR473.CAIEXEC
```

Assume the methodology specified a mask of:

```
&SYSUID..D&MSMDID..LIB&MSMDLIBN
```

Assume USERID is USER789, and the deployment ID is 877, then the resolved DSNs would be,

```
Deployed library = USER789.D877.LIB1.CAIPROC
Deployed library = USER789.D877.LIB2.CAILOAD
Deployed library = USER789.D877.LIB3.CAIEXEC
```

Local Date and Time

LYYMMDD

This is the local two-digit year.

YY two-digit year

MM two-digit month (01=January)

DD two-digit day of month (01 through 31)

Example: 100311

LYR2

This is the local two-digit year.

LYR2 two-digit year

Example: 10

LYR4

This is the local four-digit year.

LYR4 four-digit year

Example: 2010

LMON

This is the local month.

LMON two-digit month (01=January)

Example: 03

LDAY

This is the local day of the month.

LDAY two-digit day of month (01 through 31)

Example: 11

LJDAY

This is the local Julian day.

LJDAY three-digit day (001 through 366)

Example: The Julian day for January 11th is 011.

LWDAY

This is the local day of the week.

LWDAY is three characters in length. The days are MON, TUE, WED, THR, FRI, SAT, and SUN.

Example: MON

LHHMMSS

This is the local time in hours, minutes, and seconds.

HH two digits of hour (00 through 23) (am/pm NOT allowed)

MM two digits of minute (00 through 59)

SS two digits of second (00 through 59)

Example: 165148

LHR

This is the local time in hours.

LHR two-digits of hour (00 through 23) (am/pm NOT allowed)

Example: 16

LMIN

This is the local time in minutes.

LMIN two-digits of minute (00 through 59)

Example: 51

LSEC

This is the local time in seconds.

LSEC two-digits of second (00 through 59)

Example: 48

UTC Date and Time

Coordinated Universal Time is abbreviated UTC.

YYMMDD

This is the UTC date.

YY two-digit year

MM two-digit month (01=January)

DD two-digit day of month (01 through 31)

Example: 100311

YR2

This is the UTC two digit year.

YR2 two-digit year

Example: 10

YR4

This is the UTC four digit year.

YR4 four-digit year

Example: 2010

MON

This is the UTC month.

MON two-digit month (01=January)

Example: 03

DAY

This is the UTC day of the month.

DAY two-digit day of month (01 through 31)

Example: 11

JDAY

This is the UTC Julian day.

JDAY three-digit day (001 through 366)

Example: The Julian day for January 11th is 011.

WDAY

This is the UTC day of the week.

WDAY is three characters in length. The days are MON, TUE, WED, THR, FRI, SAT, and SUN.

Example: MON

HHMMSS

This is the UTC time in hours, minutes, and seconds.

HH two-digits of hour (00 through 23) (am/pm NOT allowed)

MM two-digits of minute (00 through 59)

SS two-digits of second (00 through 59)

Example: 044811

HR

This is the UTC time in hours.

HR two digits of hour (00 through 23) (am/pm NOT allowed)

Example: 04

MIN

This is the UTC time in minutes.

MIN two-digits of minute (00 through 59)

Example: 48

SEC

This is the UTC time in seconds.

SEC two-digits of second (00 through 59)

Example: 11

Maintain Methodologies

You can edit, replace, or remove (see page 121) methodologies.

Follow these steps:

1. Click the Deployments tab, and in the Actions section click the Maintain Methodologies link. The Maintain Methodologies select window appears.

Note: A grayed select box indicates that the methodology is assigned and cannot be removed. It can be edited.

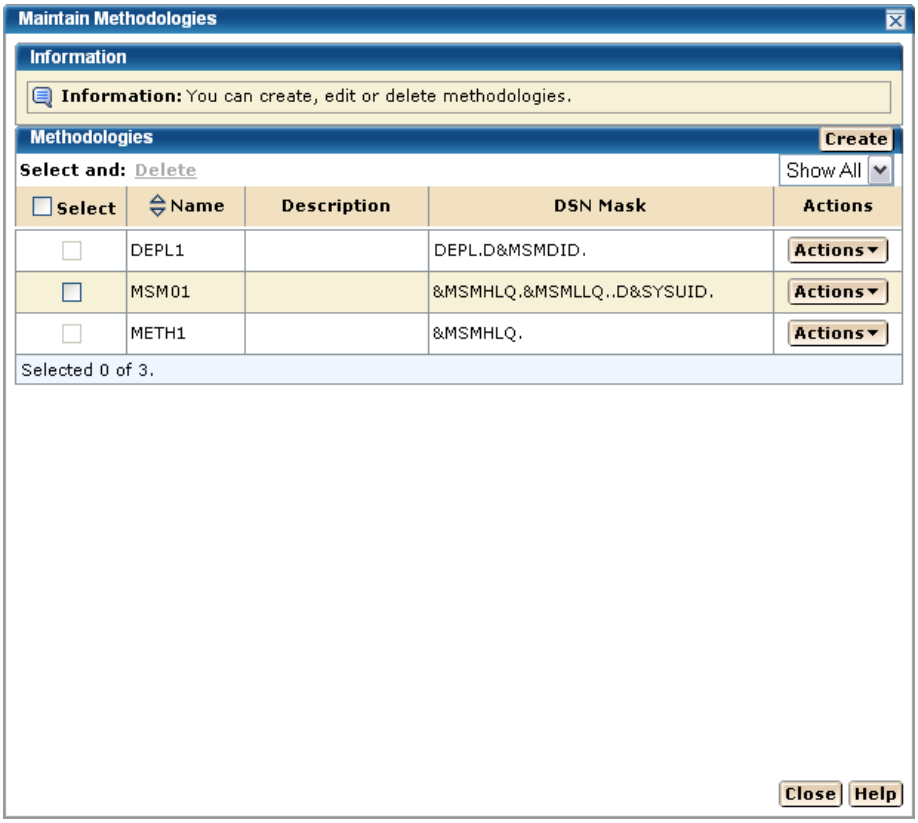

2. Select a methodology. Select Edit from Actions list.

   The Methodology window appears for editing (see page 119).

More information:

Delete Methodologies (see page 121)
Edit a Methodology (see page 119)

Edit a Methodology

You can edit a methodology by updating or modifying any of the fields on the Edit Methodology window.

Follow these steps:

1. Click the Deployments tab, and in the Actions section click the Maintain Methodologies link.
2. Select the methodology that you want to edit, click the Actions drop-down list, and then click Edit.

   The Edit Methodologies dialog appears.

   Note: The asterisk indicates that the field is mandatory.

   As with Add a Methodology, all fields are available to be edited and the details for each field are listed.
3. Enter the Methodology Name.

   Limits: Maximum 64 characters.

   Note: Each methodology name must be unique and it is not case-sensitive. For example, Meth1 and meth1 are the same methodology name.
4. Enter the Description of this Methodology.

   Limits: Maximum 255 characters.
5. Enter the data set name mask, click the file icon, and select a symbolic name (see page 109).

   Data Set Name Mask

   This is the mask that will be used to name the data sets that are deployed. They can contain symbolic qualifiers (see page 109).

   Example: CAPRODS.&SYSID. - in this case the &SYSID. will be replaced by its values. If the SYSID that is being deployed to is XX16 the DSN mask will be: CAPRODS.XX16

   Limits: Maximum 64 characters.

   Note: Each deployed target data set is named using the resolved content of the data set name mask followed by the low-level qualifier of the source data set. Appending the low-level qualifier from the source data set helps ensure uniqueness of the final data set name. Verify that the mask that you entered does not exceed 35 characters when it is translated.

The mask consists of one or more qualifiers that are separated by periods. The maximum number of characters is 64, including the periods. While you are entering the mask, CA MSM validates the mask by replacing symbolics with the minimum possible values first, and then with the maximum values. If the validation with the minimum possible values fails, an error message appears at the top of the dialog, and you cannot proceed. If the validation with the maximum values fails, a warning message appears, and you can proceed.

When the mask is translated, it has a maximum length of 44 characters including the periods and the low-level qualifier from the source data set. The low-level qualifier from the source data set has a maximum length of nine characters including a period.

-

6. Select a Style of Deployment.

Create only

Creates new data sets.

Note: Prior to creating any data sets on the remote system, a check is made, to see if the data sets already exist. The deployment is not allowed to continue if this occurs.

Create or Replace

If you select *Create or Replace* and the target data sets do not exist, they will be created. If the target data sets exist, *Create or Replace* indicates that data in the existing data set, file or directory will be replaced.

Partitioned data set

Create or Replace indicates that existing members in a partitioned data set will be replaced by members with the same name from the source file. Any currently existing member that is not in the source file will remain in the PDS. Any member from the source that does not already exist in the target PDS will be added to the target PDS.

The amount of free space in the PDS will need to be sufficient to hold the additional content, since no automatic compress will be done.

Directory in a UNIX file system

Create or Replace indicates files in a directory will be replace by files with same name from the source. Any currently existing directory in a UNIX file system that is not in the source will remain in the UNIX file system.

Sequential data set or a file in the UNIX file system

Create or Replace indicates the existing data set or file and its attributes will be replaced with the data from the source file.

For a VSAM data set (cluster)

Create or Replace indicates that an existing VSAM cluster should be populated with the data from the source file.

Note: The existing VSAM cluster must be of the same type as the source cluster (ESDS, KSDS, LDS, or RRDS), and it must have characteristics that are compatible with the source cluster (such as, record size, key size, and key offset). Replace does not verify the compatibility of these characteristics!

To replace the contents of an existing cluster, the cluster is altered to a reusable state by using an IDCAMS ALTER command, if necessary, before the data from the VSAM source is copied into the cluster by using an IDCAMS REPRO command. The REPRO command will use both the REPLACE and REUSE operands. Following the REPRO operation, the cluster is altered back to a non-reusable state if that was its state to begin with.

7. Click Save.

   Your changes are saved.

   Note: Click Cancel to close this dialog without saving your changes.

More information:

Symbolic Qualifiers (see page 109)

Delete Methodologies

Follow these steps:

1. Click the Deployments tab, and in the Actions section click the Maintain Methodologies link.

   The Maintain Methodologies select window appears.

2. Select the methodology that you want to delete.

   Note: A grayed select box indicates that the methodology is assigned and cannot be deleted. It can be edited.

3. Click Delete and then OK to the Delete Methodologies confirmation window.

   The methodology is deleted.

Systems

You can view, add, and remove systems from a deployment.

Target System Types

There are two types of *target systems*.

Test Environment

Test Environment target systems isolate untested deployment changes and outright experimentation from the production environment or repository. This environment is used a temporary work area where deployments can be tested, modified, overwritten, or deleted.

Production

Production target systems contain current working product deployments. When activating products in a production target system care must be taken, CA MSM recommends using the following procedure.

1. Copy the product to that target system with the data set names set to private. This allows only those assigned to this area to test these deployed products. The purpose of this first stage is to test or verify that the product is working.
2. Use intermediate test phases for products as they move through various levels of testing. For example you may want to let the application development group as a whole use the product in its test mode prior to moving to production.
3. Move the deployed products to production.

View a System List

You can view a system list.

Follow these steps:

1. Click the Deployments tab, and select the current deployment from the tree on the left side.

   The detailed deployment information appears on the right side.

   System Name Sort Arrows

   Click the up arrow to place the system names in alphabetic order or click the down arrow to place them in reverse alphabetic order.

   Type Sort Arrows

   Click the up arrow to place the types in alphabetic order or click the down arrow to place them in reverse alphabetic order.

   Description Sort Arrows

   Click the up arrow to place the descriptions in alphabetic order or click the down arrow to place them in reverse alphabetic order.

Add a System

You can add a system to a deployment.

Follow these steps:

1. Click the Deployments tab.

   The Deployment page appears.

2. On the right, in the Deployments panel click the Current Deployment link.

   A list of current deployments appears.

3. Click the deployment name link.

4. In the System List panel, click Add Systems.

   The Add Systems window appears.

5. Select a system to add and click OK.

   Note: When two systems have the same name, use the description to differentiate between the systems.

   The Preview window appears, and the system is added.

   Note: Sysplex systems are denoted by Sysplex System:System Name. For example, PLEX1:CO11, where PLEX1 is Sysplex name and CO11 is the system name.

Remove a System

You can remove a system from a deployment.

Follow these steps:

1. Click the Deployments tab.

   The Deployment page appears.

2. On the right, in the Deployments panel, click the Current Deployment link.

   A list of current deployments appears.

3. Select the deployment that you want to remove the system from.

   System Name Sort Arrows

   Click the up arrow to place the system names in alphabetic order or click the down arrow to place them in reverse alphabetic order.

   Type Sort Arrows

   Click the up arrow to place the types in alphabetic order or click the down arrow to place them in reverse alphabetic order.

   Description Sort Arrows

   Click the up arrow to place the descriptions in alphabetic order or click the down arrow to place them in reverse alphabetic order.

4. In the System List panel, select a system you want to remove.
5. Click Remove and then OK to the Remove Products confirmation window.

   The system is removed.

Deployment Summary

The Action button is available after a successful deployment.

Important! Data sets may need to be APF-authorized and added to the Link List and Link Pack Area. These data sets are identified in this dialog.

The Deployment Summary window may contain the following:

- Deployment ID
- Name
- Products
- Systems
- Data Sets actions
- Transport information

- Target libraries including: source, target, and resolved data set names.
- SMP/E environment
- Snapshot path and container

The following example shows the Data Sets actions, Transport, and Target libraries information.

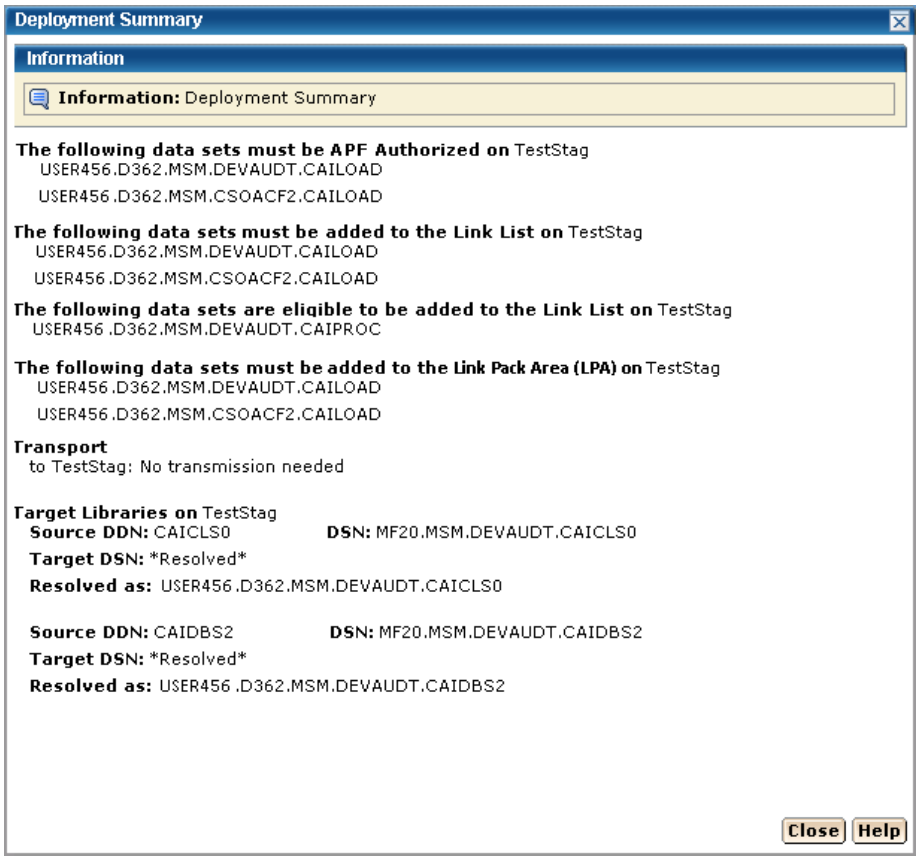

Note: When you have completed the procedures in this section, go to Configuring Your Product.

Chapter 4: Installing Your Product from Pax-Enhanced ESD

This section contains the following topics:

How to Install a Product Using Pax-Enhanced ESD

This section describes the Pax-Enhanced ESD process. We recommend that you read this overview and follow the entire procedure the first time you complete a Pax-Enhanced ESD installation. For experienced UNIX users, the *Pax-Enhanced ESD Quick Reference Guide* has sufficient information for subsequent installations.

Important! Downloading pax files for the SMP/E installation as part of the Pax-Enhanced ESD process requires write authority to the UNIX System Services (USS) directories that are used for the ESD process.

If you prefer not to involve all CA Technologies product installers with z/OS UNIX System Services, assign a group familiar with USS to perform Steps 1 through 4 and provide the list of the unpacked MVS data sets to the product installer. USS is not required for the actual SMP/E RECEIVE of the product or for any of the remaining installation steps.

To install files using Pax-Enhanced ESD, use the following process:

1. Allocate and mount the file system. This process requires a USS directory to receive the pax file and to perform the unpack steps. We recommend that you allocate and mount a file system that is dedicated to Pax-Enhanced ESD and create the directory in this file system. Ensure that all users who will be working with pax files have write authority to the directory.

2. Copy the product pax files into your USS directory. To download files, choose one of the following options:

   - Download a zip file from CA Support Online to your PC, unzip the file, and then upload the product pax files to your USS file system.
   - FTP the pax files from CA Support Online directly to your USS directory.

   Note: Perform Steps 3 through 6 for each pax file that you upload to your USS directory.

3. Create a product directory from the pax file. Set the current working directory to the directory containing the pax file, and create a directory in your USS directory by entering the following command:

   `pax -rvf pax-filename`

4. Use the SMP/E GIMUNZIP utility to create z/OS installation data sets. The file UNZIPJCL in the directory that the pax command created in Step 3 contains a sample JCL to GIMUNZIP the installation package. Edit and submit the UNZIPJCL JCL.

5. Receive the SMP/E package. Use the data sets that GIMUNZIP created in Step 4. Perform a standard SMP/E RECEIVE using the SMPPTFIN and SMPHOLD (if applicable) DASD data sets. Also, specify the high-level qualifier for the RELFILEs on the RFPREFIX parameter of the RECEIVE command.

6. Proceed with product installation. Consult product-specific documentation, including AREADME files and installation notes to complete the product installation.

7. (Optional) Clean up the USS directory. Delete the pax file, the directory that the pax command created, all of the files in it, and the SMP/E RELFILEs, SMPMCS, and HOLDDATA data sets.

More Information:

USS Environment Setup (see page 132)
Allocate and Mount a File System (see page 133)
Copy the Product Pax Files into Your USS Directory (see page 136)
Create a Product Directory from the Pax File (see page 141)
Copy Installation Files to z/OS Data Sets (see page 142)

How the Pax-Enhanced ESD Download Works

Important! To download pax files for the SMP/E installation as part of the Pax-Enhanced ESD process, you must have write authority to the UNIX System Services (USS) directories used for the ESD process and available USS file space before you start the procedures in this guide.

Use the following process to download files using Pax-Enhanced ESD:

1. Log in to https://support.ca.com/, and click Download Center.

   The CA Support Online web page appears.

2. Under Download Center, select Products from the first drop-down list, and specify the product, release, and genlevel (if applicable), and click Go.

   The CA Product Download window appears.

3. Download an entire CA Technologies product software package or individual pax files to your PC or mainframe. If you download a zip file, you must unzip it before continuing.

   For both options, The ESD Product Download Window (see page 129) topic explains how the download interface works.

   Note: For traditional installation downloads, see the *Traditional ESD User Guide*. Go to https://support.ca.com/, log in, and click Download Center. A link to the guide appears under the Download Help heading.

4. Perform the steps to install the product based on the product-specific steps.

   The product is installed on the mainframe.

ESD Product Download Window

You can download CA Technologies product ESD packages multiple ways. Your choices depend on the size of the individual files and the number of files that you want to download. You can download the complete product with all components, or you can select individual pax and documentation files for your product or component.

The following illustration shows sample product files. The illustration lists all components of the product. You can use the Download Cart by selecting one or more components that you need, or selecting the check box for Add All to cart. If you prefer to immediately download a component, click the Download link.

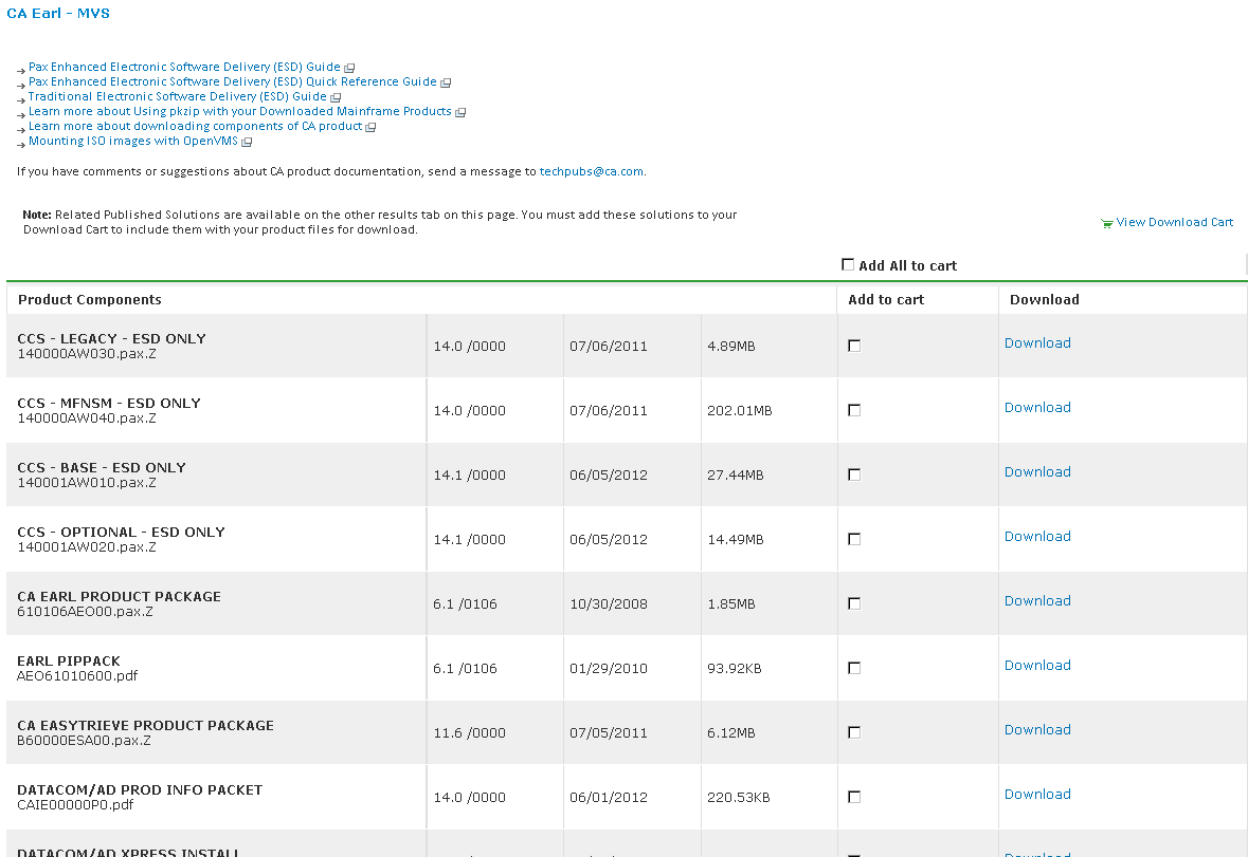

CA Earl - MVS

- Pax Enhanced Electronic Software Delivery (ESD) Guide
- Pax Enhanced Electronic Software Delivery (ESD) Quick Reference Guide
- Traditional Electronic Software Delivery (ESD) Guide
- Learn more about Using pkzip with your Downloaded Mainframe Products
- Learn more about downloading components of CA product
- Mounting ISO images with OpenVMS

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

Note: Related Published Solutions are available on the other results tab on this page. You must add these solutions to your Download Cart to include them with your product files for download.

View Download Cart

Add All to cart

Product Components				Add to cart	Download
CCS - LEGACY - ESD ONLY 140000AW030.pax.Z	14.0 /0000	07/06/2011	4.89MB	☐	Download
CCS - MFNSM - ESD ONLY 140000AW040.pax.Z	14.0 /0000	07/06/2011	202.01MB	☐	Download
CCS - BASE - ESD ONLY 140001AW010.pax.Z	14.1 /0000	06/05/2012	27.44MB	☐	Download
CCS - OPTIONAL - ESD ONLY 140001AW020.pax.Z	14.1 /0000	06/05/2012	14.49MB	☐	Download
CA EARL PRODUCT PACKAGE 610106AEO00.pax.Z	6.1 /0106	10/30/2008	1.85MB	☐	Download
EARL PIPPACK AEO61010600.pdf	6.1 /0106	01/29/2010	93.92KB	☐	Download
CA EASYTRIEVE PRODUCT PACKAGE B60000ESA00.pax.Z	11.6 /0000	07/05/2011	6.12MB	☐	Download
DATACOM/AD PROD INFO PACKET CAIE00000P0.pdf	14.0 /0000	06/01/2012	220.53KB	☐	Download
DATACOM/AD XPRESS INSTALL				☐	

Clicking the link for an individual component takes you to the Download Method page.

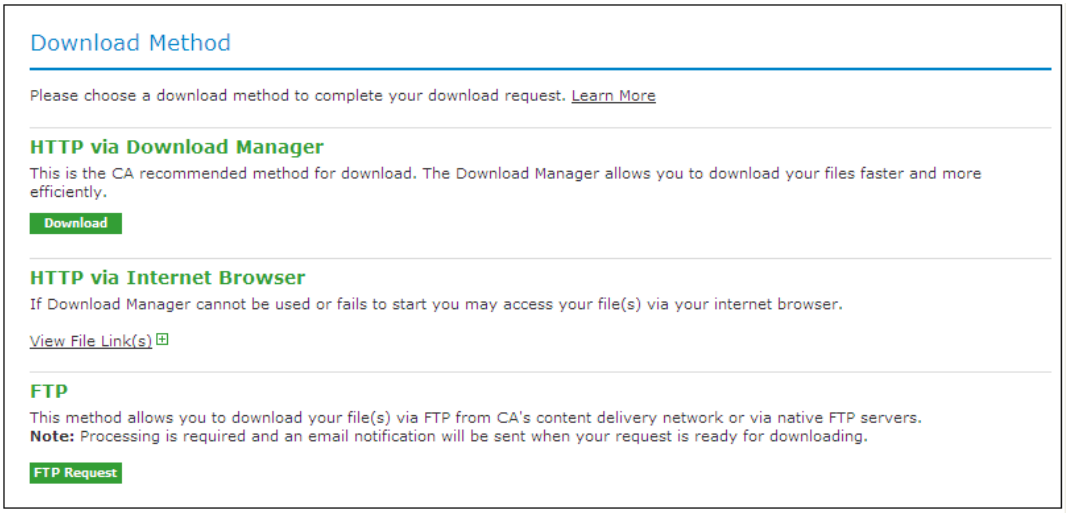

Depending on the size and quantity of ordered product files, the Download Method screen could also have these options:

Note: For mainframe downloads using this HTTP method, click the Learn More link.

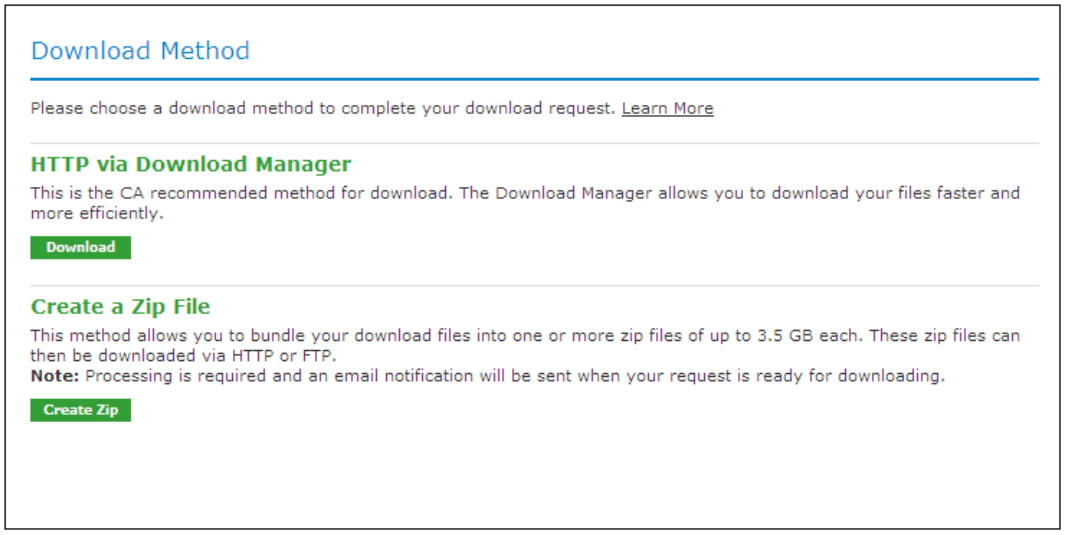

The HTTP method lets you start downloading immediately. The FTP method takes you to the Review Orders page that displays your order, first in a Pending status changing to Ready when your order has been processed.

Preferred FTP uses the new content delivery network (CDN). Alternate FTP uses the CA Technologies New York-based FTP servers.

The Create a Zip File option first creates the zip, and when ready, offers the options that the Zip Download Request examples show in the next illustration.

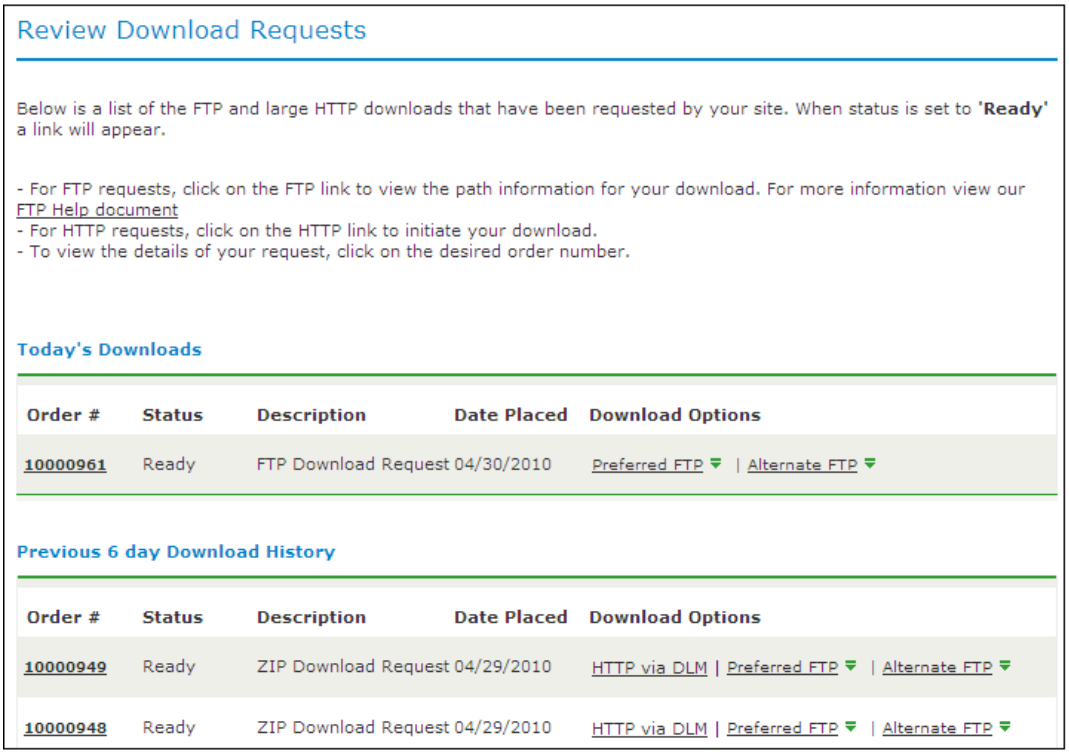

Review Download Requests

Below is a list of the FTP and large HTTP downloads that have been requested by your site. When status is set to **'Ready'** a link will appear.

- For FTP requests, click on the FTP link to view the path information for your download. For more information view our FTP Help document
- For HTTP requests, click on the HTTP link to initiate your download.
- To view the details of your request, click on the desired order number.

Today's Downloads

Order #	Status	Description	Date Placed	Download Options
10000961	Ready	FTP Download Request	04/30/2010	Preferred FTP | Alternate FTP

Previous 6 day Download History

Order #	Status	Description	Date Placed	Download Options
10000949	Ready	ZIP Download Request	04/29/2010	HTTP via DLM | Preferred FTP | Alternate FTP
10000948	Ready	ZIP Download Request	04/29/2010	HTTP via DLM | Preferred FTP | Alternate FTP

USS Environment Setup

You need a UNIX System Services (USS) directory and a file system with adequate space to perform the following tasks:

- Receive product pax files from CA Support Online.
- Perform utility functions to unpack the pax file into MVS data sets that you can use to complete the product installation.

We recommend that you allocate and mount a file system that is dedicated to Pax-Enhanced ESD. The amount of space that you need for the file system depends on the following variables:

- The size of the pax files that you intend to download.
- Whether you plan to keep the pax files after unpacking them. We do not recommend this practice.

We recommend that you use one directory for downloading and unpacking pax files. Reusing the same directory minimizes USS setup. You need to complete the USS setup only one time. You reuse the same directory for subsequent downloads. Alternatively, you can create a directory for each pax download.

Important! Downloading pax files for the SMP/E installation as part of the Pax-Enhanced ESD process requires write authority to the UNIX System Services (USS) directories that are used for the ESD process. The USS file system that is used for Pax-Enhanced ESD must have sufficient free space to hold the directory that the pax command created, and its contents. You need approximately 3.5 times the pax file size in free space to download the pax file and unpack its contents. For example, to download and unpack a 14 MB pax file, you need approximately 49 MB of free space in the file system hosting your ESD directory.

Allocate and Mount a File System

You can use the zSeries File System (zFS) or hierarchical file system (HFS) for ESD downloads.

This procedure describes how to perform the following tasks:

- Allocate a zFS or an HFS.
- Create a mount point in an existing maintenance USS directory of your choice.
- Mount the file system on the newly created mount point.

  Note: You must have either SUPERUSER authority, or the required SAF profile setting to allow you to issue the USS mount command for the file system.
- Optionally, permit write access to anyone in the same group as the person who created the directory.

  Important! USS commands are case-sensitive.

Follow these steps:

1. Allocate the file system by customizing one of the following samples to your site requirements:

   - On a zFS, use the following sample:

```
//DEFINE   EXEC PGM=IDCAMS
//SYSPRINT DD   SYSOUT=*
//SYSUDUMP DD   SYSOUT=*
//AMSDUMP  DD   SYSOUT=*
//SYSIN    DD *
  DEFINE CLUSTER ( +
  NAME(your_zFS_data_set_name) +
  STORAGECLASS(class) +
  LINEAR +
  CYL(primary secondary) +
  SHAREOPTIONS(3,3) +
  )
/*
//FORMAT   EXEC PGM=IOEAGFMT,REGION=0M,
// PARM=('-aggregate your_zFS_data_set_name -compat')
//SYSPRINT DD   SYSOUT=*
//SYSUDUMP DD   SYSOUT=*
//STDOUT   DD   SYSOUT=*
//STDERR   DD   SYSOUT=*
//CEEDUMP  DD   SYSOUT=*
//*
```

   - On an HFS, use the following sample:

```
//ALCHFS EXEC PGM=IEFBR14
//CAESD  DD   DSN=yourHFS_data_set_name,
//            DISP=(NEW,CATLG,DELETE),UNIT=3390,
//            DSNTYPE=HFS,SPACE=(CYL,(primary,secondary,1))
```

   The file system is allocated.

   Note: Ensure that the zFS or HFS data set name that you use conforms to your data set naming conventions for USS file systems. If the allocation of the file system data set fails, it is because of environmental settings not allowing for the allocation. On an HFS, try using the ISPF 3.2 Data Set Utility to allocate your HFS data set.

2. Create a mount point for the file system. This example shows how to create a /CA/CAESD directory in an existing directory, /u/maint. From the TSO OMVS shell, enter the following commands:

   ```
   cd /u/maint/
   mkdir CA
   cd CA
   mkdir CAESD
   ```

   Note: This document refers to this structure as *yourUSSESDdirectory*.

   The mount point is created.

3. Mount the file system by customizing one of the following samples to your site requirements:

   - On a zFS, use the following sample:

     ```
     MOUNT FILESYSTEM('your_zFS_data_set_name')
           MOUNTPOINT('yourUSSESDdirectory')
           TYPE(ZFS)  MODE(RDWR)
           PARM(AGGRGROW)
     ```

   - On an HFS, use the following sample:

     ```
     MOUNT FILESYSTEM('your_HFS_data_set_name')
           MOUNTPOINT('yourUSSESDdirectory')
           TYPE(HFS)  MODE(RDWR)
     ```

   The file system is mounted.

4. (Optional) Set security permissions for the directory. You can use the chmod command to let other users access the ESD directory and its files. For example, to allow write access to the ESD directory for other users in your USS group, from the TSO OMVS shell, enter the following command:

   ```
   chmod -R 775 /yourUSSESDdirectory/
   ```

   Write access is granted.

   Note: For more information about the chmod command, see the IBM *z/OS UNIX System Services User Guide* (SA22-7802).

Copy the Product Pax Files into Your USS Directory

To begin the CA Technologies product installation procedure, copy the product pax file into the USS directory that you set up. Use one of the following methods:

- Download the product pax files directly from the CA Support Online FTP server to your z/OS system.
- Download the product pax file from the CA Support Online FTP server to your computer, and upload it to your z/OS system.
- Download the product file from CA Support Online to your computer. If your download included a zip file, unzip the file, and upload the unzipped pax files to your z/OS system.

This section includes a sample batch job to download a product pax file from the CA Support Online FTP server directly to a USS directory on your z/OS system and sample commands to upload a pax file from your computer to a USS directory on your z/OS system.

Important! The FTP procedures vary due to local firewall and other security settings. Consult your local network administrators to determine the appropriate FTP procedure to use at your site.

Ensure that sufficient free space is available in the USS file system that you are using for Pax-Enhanced ESD to hold the product pax file. If you do not have sufficient free space, error messages similar to the following appear:

```
EZA1490I Error writing to data set
EZA2606W File I/O error 133
```

When the download finishes, the pax file size in your USS directory matches the value in the Size column for the corresponding pax file on the CA Technologies Products Download window.

More Information:

How the Pax-Enhanced ESD Download Works (see page 129)
ESD Product Download Window (see page 129)

Download Using Batch JCL

Use this process to download a pax file from the CA Support Product Downloads window by running batch JCL on the mainframe. Use the sample JCL attached to the PDF file as CAtoMainframe.txt to perform the download.

Important! To simplify the Pax-Enhanced ESD process, the PDF version of this guide includes a sample JCL job that you can copy directly to the mainframe. To access this job, click the paper clip icon at the left of the PDF reader. A window displaying attachments opens. Double-click the file to view the sample JCL.

Note: We recommend that you follow the preferred method as described on CA Support Online. This procedure is our preferred download method; however, we do include the procedure to download to the mainframe through a PC in the next section.

Follow these steps:

1. Supply a valid JOB statement.
2. Replace *yourTCPIP.PROFILE.dataset* with the name of the TCP/IP profile data set for your system. Consult your local network administrators, if necessary.

   The job points to your profile.
3. Replace *YourEmailAddress* with your email address.

   The job points to your email address.
4. Replace *yourUSSESDdirectory* with the name of the USS directory that you use for ESD downloads.

   The job points to your USS directory.
5. Locate the product component to download on the CA Support Product Download window.

   You have identified the product component to download.
6. Click Download for the applicable file.

   Note: For multiple downloads, add files to a cart.

   The Download Method window opens.
7. Click FTP Request.

   The Review Download Requests window displays any files that you have requested to download.

   Note: We send you an email when the file is ready to download or a link appears in this window when the file is available.

8. Select one of the following methods:

   Preferred FTP

   Uses CA Technologies worldwide content delivery network (CDN). If you cannot download using this method, review the security restrictions for servers that company employees can download from that are outside your corporate network.

   Host Name: ftp://ftpdownloads.ca.com

   Alternate FTP

   Uses the original download servers that are based on Long Island, New York.

   Host Name: ftp://scftpd.ca.com for product files and download cart files and ftp://ftp.ca.com for individual solution files.

   Both methods display the host, user name, password, and FTP location, which you then can copy into the sample JCL.

   Note: The following links provide details regarding FTP: the FTP Help document link in the Review Download Requests window and the Learn More link available in the Download Methods window.

9. Submit the job.

   Important! If your FTP commands are incorrect, it is possible for this job to fail and still return a zero condition code. Read the messages in the job DDNAME SYSPRINT to verify the FTP succeeded.

   After you run the JCL job, the pax file resides in the mainframe USS directory that you supplied.

Example: CAtoMainframe.txt, JCL

The following text appears in the attached CAtoMainframe.txt JCL file:

```
//GETPAX   JOB (ACCOUNTNO),'FTP GET ESD PACKAGE',
//          MSGCLASS=X,CLASS=A,NOTIFY=&SYSUID
//*********************************************************************
//* This sample job can be used to download a pax file directly from  *
//* CA Support Online to a USS directory on your z/OS system.         *
//*                                                                   *
//* When editing the JCL ensure that you do not have sequence numbers *
//* turned on.                                                        *
//*                                                                   *
//* This job must be customized as follows:                           *
//* 1. Supply a valid JOB statement.                                  *
//* 2. The SYSTCPD and SYSFTPD JCL DD statements in this JCL may be   *
//*    optional at your site. Remove the statements that are not      *
//*    required. For the required statements, update the data set     *
//*    names with the correct site-specific data set names.           *
//* 3. Replace "Host" based on the type of download method.           *
//* 4. Replace "YourEmailAddress" with your email address.            *
//* 5. Replace "yourUSSESDdirectory" with the name of the USS         *
//*    directory used on your system for ESD downloads.               *
//* 6. Replace "FTP Location" with the complete path                  *
//*    and name of the pax file obtained from the FTP location        *
//*    of the product download page.                                  *
//*********************************************************************
//GETPAX   EXEC PGM=FTP,PARM='(EXIT',REGION=0M
//SYSTCPD  DD   DSN=yourTCPIP.PROFILE.dataset,DISP=SHR
//SYSFTPD  DD   DSN=yourFTP.DATA.dataset,DISP=SHR
//SYSPRINT DD   SYSOUT=*
//OUTPUT   DD   SYSOUT=*
//INPUT    DD   *
Host
anonymous YourEmailAddress
lcd yourUSSESDdirectory
binary
get FTP_location
quit
```

Download Files to Mainframe through a PC

If you download pax or zip files from CA Support Online to your PC, use this procedure to upload the pax file from your PC to your z/OS USS directory.

Follow these steps:

1. Follow the procedures in How the Pax-Enhanced ESD Download Works (see page 10) to download the product pax or zip file to your PC. If you download a zip file, first unzip the file to use the product pax files.

   The pax or zip file resides on your PC.

2. Open a Windows command prompt.

   The command prompt appears.

3. Customize and enter the FTP commands with the following changes:

   a. Replace *mainframe* with the z/OS system IP address or DNS name.

   b. Replace *userid* with your z/OS user ID.

   c. Replace *password* with your z/OS password.

   d. Replace *C:\PC\folder\for\thePAXfile* with the location of the pax file on your PC.

   e. Replace *yourUSSESDdirectory* with the name of the USS directory that you use for ESD downloads.

   f. Replace *paxfile.pax.Z* with the name of the pax file to upload.

   The pax file is transferred to the mainframe.

Example: FTP Commands

This list is a sample of FTP commands to upload the pax file from your PC to your USS Pax-Enhanced ESD directory:

```
ftp mainframe
userid
password
bin
lcd C:\PC\folder\for\thePAXfile
cd /yourUSSESDdirectory/
put paxfile.pax.Z
quit
exit
```

Create a Product Directory from the Pax File

Use the sample job attached to the PDF file as Unpackage.txt to extract the product pax file into a product installation directory.

Important! To simplify the Pax-Enhanced ESD process, the PDF version of this guide includes a sample JCL job that you can copy directly to the mainframe. To access this job, click the paper clip icon at the left of the PDF reader. A window displaying attachments opens. Double-click the file to view the sample JCL.

Follow these steps:

1. Supply a valid JOB statement.
2. Replace *yourUSSESDdirectory* with the name of the USS directory that you use for ESD downloads.

   The job points to your specific directory.
3. Replace *paxfile.pax.Z* with the name of the pax file.

   The job points to your specific pax file.
4. Submit the job.

   The job runs and creates the product directory.

Note: If the PARM= statement exceeds 71 characters, uncomment and use the second form of UNPAXDIR instead. This sample job uses an X in column 72 to continue the PARM= parameters to a second line.

Sample Job to Execute the Pax Command (Unpackage.txt)

The following text appears in the attached Unpackage.txt JCL file:

```
//ESDUNPAX JOB (ACCOUNTNO),'UNPAX ESD PACKAGE ',
// MSGCLASS=X,CLASS=A,NOTIFY=&SYSUID
//*********************************************************************
//* This sample job can be used to invoke the pax command to create   *
//* the product-specific installation directory.                      *
//*                                                                   *
//* This job must be customized as follows:                           *
//* 1. Supply a valid JOB statement.                                  *
//* 2. Replace "yourUSSESDdirectory" with the name of the USS         *
//*    directory used on your system for ESD downloads.               *
//* 3. Replace "paxfile.pax.Z" with the name of the pax file.         *
//* NOTE: If you continue the PARM= statement on a second line, make  *
//*       sure the 'X' continuation character is in column 72.        *
//*********************************************************************
//UNPAXDIR EXEC PGM=BPXBATCH,
// PARM='sh cd /yourUSSESDdirectory/; pax -rvf paxfile.pax.Z'
//*UNPAXDIR EXEC PGM=BPXBATCH,
//* PARM='sh cd /yourUSSESDdirectory/; pax                            X
//*            -rvf paxfile.pax.Z'
//STDOUT DD SYSOUT=*
//STDERR DD SYSOUT=*
```

Copy Installation Files to z/OS Data Sets

Use this procedure to invoke the SMP/E GIMUNZIP utility to create MVS data sets from the files in the product-specific directory.

Follow these steps:

1. Locate and read the product readme file or installation notes, if applicable, which resides in the product-specific directory that the pax command created. This file contains the product-specific details that you require to complete the installation procedure.

   You have identified the product-specific installation details.

2. Use ISPF EDIT or TSO ISHELL to edit the UNZIPJCL sample job. You can perform this step in one of the following ways:

   - Use ISPF EDIT. Specify the full path name of the UNZIPJCL file.
   - Use TSO ISHELL. Navigate to the UNZIPJCL file and use the E line command to edit the file.

   The job is edited.

3. Change the SMPDIR DD PATH to the product-specific directory created by the pax command.

   Your view is of the product-specific directory.

4. If ICSF is not active, perform the following steps:

   a. Change the SMPJHOME DD PATH to your Java runtime directory. This directory varies from system to system.

      Note: The default Java location is the following:

      `/usr/lpp/java/Java_version`

   b. Perform one of the following steps:

      - Change the SMPCPATH DD PATH to your SMP/E Java application classes directory, typically /usr/lpp/smp/classes/.
      - Change HASH=YES to HASH=NO on the GIMUNZIP parameter.

   One of the following occurs: ICSF is active, or you are using Java.

5. Change all occurrences of y*ourHLQ* to the high-level qualifier (HLQ) for z/OS data sets that the installation process uses. We suggest that you use a unique HLQ for each expanded pax file to identify uniquely the package. Do *not* use the same value for y*ourHLQ* as you use for the SMP/E RELFILEs.

   All occurrences of y*ourHLQ* are set to your high-level qualifier for z/OS data sets.

6. Submit the UNZIPJCL job.

   The UNZIPJCL job completes with a zero return code. Messages GIM69158I and GIM48101I in the output and IKJ56228I in the JES log are acceptable.

   GIMUNZIP creates z/OS data sets with the high-level qualifier that you specified in the UNZIPJCL job. You use these data sets to perform the product installation. The pax file and product-specific directory are no longer needed.

   Note: For more information, see the IBM *SMP/E for z/OS Reference (SA22-7772)*.

Receiving the SMP/E Package

If you are installing the package into a new SMP/E environment, use the sample jobs included with the product to set up an SMP/E environment before proceeding.

At this point, complete the SMP/E RECEIVE using files on DASD that the UNZIPJCL job created. Consult the product sample JCL library that contains a sample job customized to receive the product from DASD. Specifically, you must specify the following values:

- DASD data set names for SMPPTFIN and SMPHOLD (if applicable)
- The HLQ that you used in the UNZIPJCL job on the RFPREFIX parameter on the RECEIVE command

How to Install Products Using Native SMP/E JCL

The following steps describe the process to install products using native SMP/E JCL:

1. Allocate product data sets and SMP/E data sets.
2. Create SMP/E CSI.
3. Receive base functions.
4. Apply base functions.
5. Accept base functions.
6. Configure the product according to your site requirements.

Prepare the SMP/E Environment for Pax Installation

The members in this procedure prepare the data sets, initialize the zones, and create the DDDEFs for CA Disk. External DDDEF data sets are required. The default is NULLFILE.

For information about the members, see the comments in the JCL.

To prepare the SMP/E environment for your product

1. Customize the macro SDMSEDIT with your site-specific information and then copy the macro to your syslib location. Replace the rightmost parameters for each ISREDIT CHANGE macro command. Each time you edit an installation member, type SDMSEDIT on the TSO command line, and press Enter to replace the defaults with your specifications.

   The macro is ready to customize your *yourHLQ*.SAMPJCL members.

   Note: Set the DASD HLQ to the same value specified for *yourHLQ* for the unzip to DASD ESD JCL.

   Note: The following steps include instructions to execute the SDMSEDIT macro each time you open a new SAMPJCL member. To edit all SAMPJCL members simultaneously, read and follow the instructions in the SDMEDALL member.

2. Open the SAMPJCL member SDM1ALL in an edit session and execute the SDMSEDIT macro from the command line.

   SDM1ALL is customized.

3. Submit SDM1ALL.

   This job produces the following results:

   - The target and distribution data sets for CA Disk are created.
   - Unique SMPLTS, SMPMTS, SMPSCDS, and SMPSTS data sets for this target zone are created.

4. Open the SAMPJCL member SDM2CSI in an edit session and execute the SDMSEDIT macro from the command line.

   SDM2CSI is customized.

5. Submit SDM2CSI.

   This job produces the following results:

   - The CSI data set is defined.
   - The SMPPTS and SMPLOG data sets are allocated.
   - The global, target, and distribution zones are initialized.
   - The DDDEF entries for your product are created.
   - The DDDEFs for the required SMP/E data sets are created.

Run the Installation Jobs for a Pax Installation

Submit and run these SAMPJCL members in sequence. Do not proceed with any job until the previous job has completed successfully.

To run the installation jobs

1. Submit the *yourhlq*.SAMPJCL member SDM3RECD to receive SMP/E base functions.

   CA Disk is received and now resides in the global zone.

2. Customize and submit the *yourhlq*.SAMPJCL member SDM4APP to APPLY SMP/E base functions.

   Your product is applied and now resides in the target libraries.

3. Customize and submit the *yourhlq*.SAMPJCL member SDM5ACC to ACCEPT SMP/E base functions.

   Your product is accepted and now resides in the distribution libraries.

Clean Up the USS Directory

Important! This procedure is optional. Do not use this procedure until you complete the entire installation process.

To free file system disk space for subsequent downloads after downloading and processing the pax files for your CA Technologies product, we recommend removing the files from your USS directory and deleting unnecessary MVS data sets. You can delete the following items:

- Pax file
- Product-specific directory that the pax command created and all of the files in it
- SMP/E RELFILEs, SMPMCS, and HOLDDATA MVS data sets

  These data sets have the HLQ that you assigned in the UNZIPJCL job.

Note: Retain non-SMP/E installation data sets such as *yourHLQ*.INSTALL.NOTES for future reference.

Follow these steps:

1. Navigate to your Pax-Enhanced ESD USS directory.

   Your view is of the applicable USS directory.

2. Delete the pax file by entering the following command:

   `rm paxfile`

   paxfile

   Specifies the name of the CA Technologies pax file that you downloaded.

   The pax file is deleted.

3. Delete the product-specific directory by entering the following command:

   `rm -r product-specific_directory`

   product-specific_directory

   Specifies the product-specific directory that the pax command created.

   The product-specific directory is deleted.

Note: You can also use TSO ISHELL to navigate to the pax file and product-specific directory, and delete them using the D line command.

Apply Maintenance

CA Support Online may have maintenance and hold data that have been published since the installation data was created.

To apply maintenance

1. Check CA Support Online and download any PTFs and hold data published since this release was created.
2. Transfer the downloaded files to two separate FB 80 sequential data sets. Use one data set to contain the PTFs and the other to contain the hold data.

   The PTFs and hold data become accessible to the *yourhlq*.SAMPJCL maintenance members.

3. Edit and submit the SDMSEDIT macro.

   The *yourhlq*.SAMPJCL members SDM6RECP, SDM7APYP, and SDM8ACCP are customized.

4. Customize the SDM6RECP SMPPTFIN and SMPHOLD DD statements to reference the FB 80 data sets for the PTFs and hold data.

5. Submit SDM6RECP.

   The PTFs and hold data are received.

6. Submit SDM7APYP.

   The PTFs are applied.

7. (Optional) Customize and submit *yourhlq*.SAMPJCL member SDM8ACCP.

   The PTFs are accepted.

   Note: You do not have to submit the job at this time. You can accept the PTFs according to your site's policy.

Note: We recommend that you check for available maintenance; however, you may find that none is available.

HOLDDATA

When you apply maintenance, you typically encounter SMP/E HOLDDATA. We use HOLDDATA to notify your SMP/E system of SYSMODs that have errors or special conditions. We support system and external HOLDDATA.

Note: When you have completed the procedures in this section, go to Configuring Your Product.

Note: When you have completed the procedures in this section, go to Configuring Your Product.

Chapter 5: Configuring Your Product

This chapter contains topics on modifying the CA Disk system libraries. Customization of the CA Disk system can be performed before Deployment. However, modifications made on the remote systems can be overlaid by subsequent deployments.

Additionally, the CA Disk SMP/E CSI will not be deployed to the remote system. Although CA Disk can be modified outside of SMP/E on a remote system, SMP/E usermods are only possible before Deployment. Note, however, that CA Disk SVC Zaps are modifications to the z/OS system, and should be performed at each remote site.

To deploy without risk of overlay, perform all CA Disk customization at the central site prior to creating deployments.

This section contains the following topics:

Activate the CA Disk System

Following is the process for activating the basic CA Disk system. The instructions include:

- CA Disk Parameter Library (PARMLIB)—Identify FILES and PARMLIB data sets to CA Disk functions.
- CA Disk JCL Procedures—Modify symbolic parameters in JCL procedures and customize JCL procedures.
- CA Disk FDS—Create a subfiles definition member in PARMLIB and initialize the FILES data set.

CA Disk Parameter Library (PARMLIB)

The control parameters that CA Disk uses, provide the following information for processing:

- How CA Disk operates
- The format of reports
- Data sets that CA Disk is to skip processing for
- Provide user-dependent information to CA Disk functions

These parameter lists are specified as members of a partitioned data set known as the parameter library (PARMLIB). The PARMLIB implementation allows CA Disk to tailor its operation to your requirements. The contents of a PARMLIB member can be easily changed. After they are changed, the next execution of CA Disk uses them.

The parameter library is ready to use as it was loaded from tape. Do not change the supplied members. The supplied members are parameter lists that are used for internal CA Disk system functions. The members are replaced with each new release. You can place parameter information, for your installation and users, in a new member of the parameter library that you create.

For more information about common CA Disk parameter lists and the purpose of each, see CA Disk Libraries in the *Systems Guide.*

For the remainder of the installation process, only the following two members of PARMLIB are of primary concern:

- SYSPARMS

  SYSPARMS indicate your processing options to CA Disk. See the Sysparms section in the Systems Guide for all available system parameters.

  For example, the default retention period for data sets being archived is 30 days. To change the 30 days to 120 days, create a member that is named SYSPARMS in the PARMLIB data set. Add one line that contains ARCRETPD0120 or one line with RETRETPD0120. Optionally you can add both lines.

- FILEDEFN

  FILEDEFN defines the internal attributes of the CA Disk FILES data set (FDS).

  Because the CA Disk parameter library is a partitioned data set, it can be updated like any other PDS.

Identify FILES and Parmlib Data Sets

Several functions of CA Disk must know the names of the FILES data set and parmlib data set used within the function. The names of these data sets are made available to CA Disk through the SET command.

Additionally, indicators can be set to allow the use of preallocated FILES and parmlib data sets. Alternate FILES and parmlib data sets can by allocating them before executing the desired function.

To provide the correct data set names and permit alternate data sets to be preallocated, use the following utility. A copy of this JCL can be found in the member name DMSSET in the CCUWJCL library.

```
//MODIFY   EXEC PGM=ADSMI002,PARM=ADSDM338
//STEPLIB    DD  DISP=SHR,DSN=CA.DISK.LOADLIB
//SYSLIB     DD  DISP=SHR,DSN=CA.DISK.LOADLIB
//ABNLDUMP   DD  DUMMY
//CMDPRINT   DD  SYSOUT=A
//MSGPRINT   DD  SYSOUT=A
//PARMLIB    DD  DISP=SHR,DSN=CA.DISK.PARMLIB
//SYSPRINT   DD  SYSOUT=A
//SYSUDUMP   DD  SYSOUT=A
//SYSIN      DD  *
SET .....
```

DMSSET updates the load module DMSPARMS. Run the DMSSET job against your CA Disk SMP/E target library. This job prevents problems with the module being overlaid when copying from target to production libraries.

SET Command Syntax

To specify the data set names and options that are used during CA Disk functions, use the SET command. Only one command is accepted per execution. The following command is an example of the format:

```
SET FILESDSN=,PARMSDSN=,ALTFILES,ALTPARMS,PARMLIBD,NOALTFILES,NOPARMS, NOPARMLIBD
```

FILESDSN=

Specifies the name of the CA Disk FILES data set used in the CA Disk environment. FILESDSN= is a required parameter.

PARMSDSN=

Specifies the name of the CA Disk parmlib data set used in the CA Disk environment. PARMSDSN= is a required parameter.

ALTFILES

To allow users to allocate DDNAME=FILES to their ISPF/TSO session before executing CA Disk command processors, specify this simple parameter. ALTFILES allows users to use other than the default FILES data set.

ALTPARMS

To allow users to allocate DDNAME=PARMLIB to their ISPF/TSO session before executing CA Disk command processors, specify this parameter. ALTPARMS allows users to use other than the default parmlib data set.

PARMLIBD

If your online environment already uses a ddname of PARMLIB (for something other than CA Disk), specify this parameter to cause CA Disk to use the alternate ddname of PARMLIBD.

NOALTFILES

To use the default FILES data set for all ISPF/TSO sessions, specify this parameter.

NOPARMS

To use the default parmlib data set for all ISPF/TSO sessions, specify this parameter.

NOPARMLIBD

To cause CA Disk to use the ddname of PARMLIB for the default parmlib data set, specify this parameter.

CA Disk JCL Procedures

CA Disk provides a set of PROCLIB JCL members for all functions. These members have the following effects:

- Save time in the initial JCL preparation
- Reduce execution failures due to improper JCL
- Enhance the installation, testing, and continued use of the CA Disk functions

Modify some of these procedures to increase the size of the temporary and software files, depending on the size of your FILES data set.

The following table is a list of the members in the procedure library and a description of the functions of each member.

PROCLIB JCL Member	Function Description
ARCHIVE	Explicit archive
BILLING	Non-VSAM DASD billing (select, accumulate)
COMPRES	PDS compression
DEFVOLS	List tapes needed for deferred restores
DERASE	Delete deferred archive/restore requests
DIM	Dynamic Installation Manager
DMS	DSCL-invoked functions
DMSAR	Auto-restore
DMSGTF	Diagnostic procedure used for problem resolution
DMSLDSD	RACF profile list utility
DMSLRAC	List alternate RACF profile cross-reference

DMSPOOL	Tapepool update utility
DMSPROF	CA Disk RACF profile maintenance utility
DMSUTIL	Miscellaneous diagnostics
EXTEND	DASD billing (extend totals and clear)
FMS	DSCL Recover
FRECOVER	FDS forward recovery
IXCATLG	Catalog archived DSNs for auto-restore
IXMAINT	Index maintenance utilities
IXUPDATE	Index update utilities
LISTD	List archive index
LISTREQ	List deferred archive/restore requests
MERGE	Merge unexpired archive data sets
MERGCOPY	Copy new archives created by merge
MIGRATE	Sequential migration to tape
PDS2SEQ	Merge PDS members into a sequential file
RACFCHK1	RACF profile synchronization utility
REBUILD	Archive index rebuild
RECOVER	Implicit recovery
RELOAD	FDS recovery
REORG	FDS reorganization
REPORT	Special purpose non-VSAM reports
RESTART	Restart after REORG shutdown
RESTORE	Explicit restore
SMFRPT	Data set report based on SMF data
THRSHMGR	Volume threshold manager
TSTAR	Test a new/updated auto-restore function
UNLOAD	FDS backup
VSAMBILL	VSAM DASD billing (select, accumulate)
XCOPY	Disaster recovery extract utility

All of these procedures make consistent use of the following symbolic parameters and their recommended defaults:

Symbolic Parameters	Recommended Default	Explanation
DCAI	CAI.DISK.ADRRSPNN	Qualifier for CA Datacom/AD CAILIB data sets
DCUS	CAI.DISK.ADRRSPNN	Qualifier for CA Datacom/AD CUSLIB data sets
P =	CA.DISK.CCUWJCL	Library for control statements
Q =	CA.DISK	Qualifier for CA Disk data sets
S =	*	SYSOUT class default
W =	SYSDA	WORKFILE default unit name (must have a storage volume)
U =	Null	TSO userid

If you use the recommended P and Q symbolic parameters, the JCL procedures generate the following data set names:

Data Set Names	Description
&DBPFX..DMS&DBID	CAI.DISK.DMSnnn The data portion of the Datacom database that contains the CA Disk FILES data
&DBPFX..IXX&DBID	CAI.DISK.IXXnnn The index portion of the Datacom database that contains the CA Disk FILES data
&DCAI..CAILIB	CAI.DISK.ADRRSPNN.CAILIB Contains Datacom modules that are not customized by the installation
&DCUS..CUSLIB	CAI.DISK.ADRRSPNN.CUSLIB Contins CA Datacom/AD modules that are customized by the installation
&DCUS..CXX	CAI.DISK.ADRRSPNN.CXX Contains information maintained by the Datacom MUF
&Q..PARMLIB	CA.DISK.PARMLIB Contains the system parameters and installation options
&Q..LOADLIB	CA.DISK.LOADLIB Contains the executable load modules
&Q..FILES	CA.DISK.FILES Contains the index of archived or backed up data sets
&Q..ARCHPRIM	CA.DISK.ARCHPRIM Data set name used for the primary archive/backup tapes

&Q..ARCHCOPY	CA.DISK.ARCHCOPY Data set name used for the duplicate copy archive/backup tapes
&Q..MERGPRIM	CA.DISK.MERGPRIM Data set name used for the primary tapes created by merge
&Q..MERGCOPY	CA.DISK.MERGCOPY Data set names used for the duplicate copy tapes from merge
&Q..MERGTERT	CA.DISK.MERGTERT Data set names used for the tertiary copy tapes from merge
&U.&Q..WORKFILE	CA.DISK.WORKFILE Data set names used for the PDS compression work file
&P	CA.DISK.CCUWJCL Data set names used for the library containing needed control statements

The JCL procedures also use the following control statements (DD Statements). By default, they are retrieved from the CA Disk installation library. You can copy them to any source image library. To name the appropriate library, modify the symbolic parameter P.

Control Statements (DD Statements)	Procedure that uses the statements
DMSORT1	■ ARCHIVE ■ DMS ■ DMSLRAC ■ MERGE ■ MIGRATE ■ RECOVER ■ REPORT ■ RETAIN
DMSORT2	MERGE
DMSORT3	EXTEND (DASD Billing)
DMSORT4	COMPRES
DMSORT5	DMS
DMSORT6	Sort for non-DSNINDEX subfile records
DMSORT7	Sort for DSNINDEX subfile records only
DMSORT8	DMS
DMSPASS1	REPORT
DMUNLOAD	MERGE

DMUNLOD2	RACFCHK1
GTFPOPT	DMSGTF
RELOAD	■ REORG ■ RESTART ■ Reload subfiles
SMFSORT	SMFRPT

Step 2. Modify Symbolic Parameters in JCL Procedures

Modify each JCL procedure manually using ISPF edit or any other editor. Two utilities are provided in the Installation library that lets you edit globally the JCL procedures. The following utilities and macros are available:

PROCUNLD

To unload the procedures into a sequential file in IEBUPDTE format, use PROCUNLD against your procedure library. To modify and change the procedures, use ISPF edit or another editor. For example, to change the first and second-level qualifiers of your CA Disk target libraries, you can use the following ISPF edit command:

```
CHANGE CA.DISK TO SYS2.CA.DISK ALL
```

You can also use the PROCLIB Customization Wizard – the CUSTPROC Edit Macro which helps you customize the unloaded Proclib.

PROCRELD

To reload the tailored procedures in your target procedure library, use PROCRELD against your procedure library.

CUSTPROC

To make a mass change of the sequential file the PROCUNLD utility produces, use CUSTPROC. CUSTPROC, also named the CA Disk PROCLIB Customization Wizard, is in the DMSCLIB library that SMP/E loaded. To use the CLIST as an Edit Macro, include it in a library in your SYSPROC concatenation. If you load the CLIST before installing the DMSCLIB data set, copy CUSTPROC into the //SYSPROC concatenation. The Wizard also requires the use of several panels in the CA Disk Panel Library. You can include the library in your default Panel library concatenation. You can also supply it when the Wizard requests the library name.

To invoke the Wizard, enter CUSTPROC in the command line of an active edit session of the procedures that PROCUNLD unloads. The Wizard displays an ISPF panel. Type values in the fields to replace the current values for the following fields:

- Symbolic parameters Q=, P=, S=, W=
- Default steplibs and Parmlibs
- CA Datacom/AD Library names and the symbolic parameters DCUS=, DCAI=

  The names are used as the prefixes to the Datacom CUSLIB and CAILIB data sets.
- Whether you want the Datacom libraries included as part of the standard STEPLIB concatenation.
- Enter the names of FILES data sets, or databases, that you want to use as your MFILES concatenation.

  The data sets are added automatically to the appropriate procedures. To access fields that are not displayed on the screen, vertically scroll the panel.

Rules for PROCLIB Customization Wizard

Observe the following about the PROCLIB Customization Wizard:

- PROCLIB members that are not part of the standard CA Disk installation will be excluded (from the ISPF EDIT Exclude command) and will not be updated. At the end of the customization they will still be excluded but the Wizard doesn't delete them.
- The ISPF panel that contains the parameters specifications is larger than a full screen and it may be necessary to scroll down to enter all of the necessary values.
- Within the updated procedures where steps execute programs that are neither CA Disk programs nor Datacom programs (such as SORT, for example), neither the STEPLIB nor the PARMLIB data set concatenations will be changed. Thus if SORT requires a STEPLIB the STEPLIB in the SORT step will be unchanged.
- ALL of the supported symbolic parameters in the procedures will be changed to the same value – the value that you enter on the parameter panel – regardless of the value that is currently in the procedures. If you want some procedures to have a different value for W=, for example, it is necessary to perform additional customization on those procedures after you have run the customization Wizard.
- CUSLIB and CAILIB have the default values &DCUS..CUSLIB and &DCAI..CAILIB. It may cause JCL errors in some procedures if you change the value of the names so that they do not include the &DCUS and &DCAI symbolic parameters and requires you to manually modify any procedures that use &DCUS or &DCAI to remove the symbolic parameter.

- The option to include the Datacom libraries should be selected only after CA Datacom/AD has been installed and customized for CA Disk. You will receive JCL errors using the customized procedures if the libraries do not exist at the time the job is run. If you have not yet installed Datacom, or if you plan to NOT use Datacom databases as FILES, the option should be unselected. The Customization Wizard can be rerun in the future against the already customized procedures if your needs change in the future. Any procedures that had additional manual customization to change any of the supported parameters to different values may require that customization to be redone.
- The names entered into the STEPLIB Library Names should only be CA Disk library names. Specifically the Datacom library names should not be entered in here.
- The STEPLIBS for programs that are supported will be updated based on which libraries they need. Most of the Datacom programs require only the Datacom libraries. Some of the CA Disk programs do not need the Datacom libraries because they do not reference the FILES datasets so the STEPLIB will only contain CA Disk libraries.
- The names entered into the PARMLIB Library Names should be the customized library names to be used in the procedures.
- The names in the MFILES Data Set Names should be the names of the FILES data sets that you want to be searched, and in the order that you want them searched, for procedures that support MFILES.
- If the MFILES data set concatenation is in the supported procedures it is replaced with the data sets you have specified, or removed if you have not specified any FILES data set names. If the MFILES data set concatenation is not defined in the supported procedures and the MFILES data set names are specified, the MFILES concatenation is inserted immediately after the last data set in the step.

Type CANCEL in the command field and press the Enter key if you decide that you do not want to proceed with the customization at this time. If you press the END key, you will be prompted to confirm you are leaving the dialog.

Once you have filled in the fields with the values you desire, type GO and press the Enter key. That will start scanning the procedures and make the requested changes. The following is a sample illustration of a filled in panel.

```
SETUP ------------------ PROCLIB Customization Wizard -------------- CA Disk --

Command ===>


Enter GO to start processing or CANCEL to exit without making changes.


More:     +

     PROC Parameters:

       Q= . . . . . CAI.DISK

       P= . . . . . CAI.DISK.CCUWJCL

       DCUS=  . . . CAI.DISK.ADRRSPNN

       DCAI=  . . . CAI.DISK.ADRRSPNN

       S= . . . . . *                W= . . . . . SYSDA


     CA Datacom/AD Library Names:

       CUSLIB . . . &DCUS..CUSLIB

       CAILIB . . . &DCAI..CAILIB


       Enter "/" to select option

        /  Include CA Datacom/AD Libraries in STEPLIB Concatenation


     STEPLIB Library Names:

              . . . &Q..LOADLIB

              . . .

              . . .

              . . .

              . . .
```

```
PARMLIB Library Names:

         . . . &Q..CCTUPARM

         . . .

         . . .

         . . .


MFILES Data Set Names:

         . . . &Q..FILES1

         . . . &Q..FILES2

         . . . &Q..FILES3

         . . .

         . . .

         . . .

         . . .

         . . .

         . . .

         . . .

         . . .

         . . .

         . . .

         . . .
```

Step 3. Customize JCL Procedures

The PDS Compress function uses a work file to perform the compression. It is allocated by the //COMPWORK dd statement in the distributed COMPRES procedure, and must be big enough to hold the largest PDS selected for compression. The supplied primary and secondary space allocation quantities are adequate for processing most PDSs, even when limiting the work file to a single volume. However, to compress a very large PDS, the space available for the work file on one volume cannot be sufficient. If you view this as a potential problem, update the COMPRES proc directly and change the value of the WRKVOLS= parameter, which defaults to 1. A larger value permits the work file to extend to multiple volumes as needed. Set it to the maximum number of volumes you want to make available to the work file. Your installation's configuration limits the value you can assign to this parameter.

The distributed procedures for ARCHIVE, DMS, RECOVER, RESTORE, and RETAIN all include dd statements for tape drives that have been commented out. Although these dd statements can be *uncommented* to cause system allocation routines to allocate the proper device when the job is initiated, CA Disk default processing dynamically allocates them when and as needed. Dynamic allocation is recommended because the actual need for the device depends upon the CA Disk parameters that are specified for the job. Simulated functions probably do not need the drives, but when submitted in live mode, they do. Similarly, the CA Disk ARCHIVE and RESTORE facilities have options to queue the requests rather than executing them immediately. Queuing the requests does not require the drives, but later processing of the queues does.

TAPE is used for the tape unit name and is not included as a symbolic parameter. You can update it to the appropriate value for your installation.

The region size in most of the procedures has been set at 5120 KB. Monitoring the actual memory use during the evaluation period can permit further reductions, based on the specific processing at your installation.

CA Disk FILES

CA Disk uses an assortment of records stored in a receptacle, called the FILES, to record information related to the various functions you are executing. The information stored is divided up into sections called subfiles. The following is a list of each subfile and a description of its function:

DSNINDEX

An index of all data sets in the archives.

ARCHVOLS

An index of tape or disk volumes that contain archived data sets.

DMSPOOLS

Volser names of tapes (assigned to pools) that can be used during archive or backup runs.

ARCHCMDS

A queue for user requests to archive or backup specific data sets.

RESTCMDS

A queue for user requests to restore specific data sets.

RETEXCLD

Re-archive grace periods for restored data sets.

DASDSPCB

DASD space billing records.

MIGRECAT

Re-catalog information for sequential migration backup tapes.

DMSPARMS

TSO dynamic (immediate) restore information.

RACFENCD

RACF profile name cross-reference.

The next two steps are required for all new users of CA Disk. If you already have an FDS from an 8.1 release or above, you can continue to use that FDS; however, you can create and initialize a new FDS for testing purposes before you install this release into your production environment. If you have a FDS from an 8.0D release or below, see Files Data Set Conversion. If you are using Files Database(FDS), the next two steps can be skipped.

Step 4. Create a Subfiles Definition Member in Parmlib

To create the FDS, first define its subfile characteristics to CA Disk. A sample set of subfile definition entries is supplied in the parameter library. The only variable in each entry is the number of records (capacity) you expect to need in each subfile. The minimum is one. For initial use, the supplied capacities should be more than sufficient, and provides the best performance. Copy the sample definitions (member FDSAMPLE in parmlib) into a new member of parmlib named FILEDEFN. This new member is used to initialize your FDS.

Step 5. Initialize the FDS

After your subfile definition member FILEDEFN has been created as outlined above, run the following JCL to allocate and initialize the FDS. Your output consists of a status report based upon your definitions. You can find a sample in the Sample FDS STATUS Report in the Systems Guide.

Special note for RACF users: By default, CA Disk bypasses a data set if the VTOC entry for it indicates it is RACF-indicated. Details on instructing CA Disk to process these data sets are presented in the topic Installing the RACF Security Interface. Therefore, if the FDS becomes RACF-indicated as soon as you allocate it, you can receive message 0725 saying that the data set has been bypassed (and therefore not initialized). Either unprotect the FDS for the duration of the initialization run, or specify sysparm RACFSUPP with a value of Y, which allows the function to complete.

The following JCL can also be found as member FILEINIT in the installation library:

```
   Menu  Utilities  Compilers  Help
 ------------------------------------------------------------------------------
 BROWSE                         CCUWJCL(FILEINIT)     Line 00000000 Col 001 080
 Command ===>                                                Scroll ===> CSR
******************************** Top of Data ********************************
//JOBNAME  JOB (ACCT INFO)
//*****************************************************************
//* SAMPLE JCL TO INITIALIZE (FORMAT) THE CA DISK                 *
//* FILES DATA SET.                                               *
//*                                                               *
//* NOTE: CHECK DSNAMES, ADD VOLSER IF DESIRED                    *
//*****************************************************************
//FILEI   EXEC PGM=ADSMI000,PARM=ADSDM100,REGION=2048K
//STEPLIB   DD DISP=SHR,DSN=CAI.DISK.CCUWLOAD

//ABNLDUMP  DD DUMMY
//CMDPRINT  DD SYSOUT=*
//FILES     DD DISP=(,CATLG,DELETE),
//             DSN=CAI.DISK.FILES,
//             VOL=SER=,                  <=== SUPPLY VOLUME
//             UNIT=SYSALLDA,
//             DCB=(DSORG=DA),
//             SPACE=(CYL,10,,CONTIG)
//MSGPRINT  DD SYSOUT=*
//PARMLIB   DD DISP=SHR,DSN=CAI.DISK.CCUWPARM

//SYSPRINT  DD SYSOUT=*
//SYSUDUMP  DD SYSOUT=*
******************************** Bottom of Data ********************************
```

Step 6. Verify CA Disk Enqueue Usage

CA Disk issues several types of enqueues during processing. Some enqueues are the result of normal system services such as dynamic allocation. Others are issued for internal CA Disk functions. For more information about the enqueues, dequeues, and reserves issued by CA Disk, see the section ENQ/DEQ/Reserve Usage. If you use CA Disk in a multi-system environment, make sure that all enqueues are propagated across all systems. Failing to do so can result in problems later.

Note: When a reserve against a CA Disk FDS is converted to an enqueue, both SCOPE=SYSTEMS and sysparm RSUPPRES set to a Y need to be in place for a GRS environment to function properly.

Activate CA Disk Features

Activating CA Disk Security features and Interfaces consists of:

- Selecting password-indicated data sets
- Activating the CA Disk system security
- Preparing Security Packages to work with CA Disk
- Installing the CA Disk Security Interface.

Each of these features is described in the following sections.

Below are instructions for activating the following CA Disk features:

- System Security Features and Interfaces (see page 164)
- VSAM Support (see page 180)
- ISPF Support (see page 182)
- TSO Support (see page 187)
- DASD Billing (see page 189)
- DBRC Interface (see page 189)

Activate CA Disk Security Features and Interfaces

Activating CA Disk Security features and Interfaces consists of; selecting password-indicated data sets, activating the CA Disk system security, preparing Security Packages to work with CA Disk, and installing the CA Disk Security Interface. Each of these features is described more detailed in the following section.

Step 1. Select Password-Indicated Data Set Support Options

In IBM standard password support, when a password-indicated non-VSAM data set is opened, the operating system prompts the TSO user (for online applications) or the master console operator (for batch jobs) for the password to the data set in question. Authorities to continue is based on the response to that prompt. For non-VSAM data sets, the password indicator bit is the DS1IND10 bit set (bit x'10' at offset 93 x'5D') in the format-1 DSCB.

CA Disk has special code to avoid prompting the operator for the password on OPEN and SCRATCH of non-VSAM user data sets. CA Disk uses this special code only when sysparm PASSWORD is specified with a value of Y, when the PASSWORD parameter is included on the command (if it has an optional PASSWORD parameter), and when CA Disk is running as an APF-authorized task. Most TSO and ISPF sessions run non-APF-authorized, so CA Disk does not exempt them from authority checking.

The security package named SECURE replaces the non-VSAM password-checking feature of the operating system. SECURE does not result in prompting of the operator for open of data sets, so this special CA Disk code is not needed.

There are two sub steps for installing the password-indicated data set support options. The following are the sub steps:

- If your installation does not use the security package SECURE, see the setting of sysparm PASSWORD. If you want CA Disk to process password-indicated data sets, specify the sysparm value as Y.

  If you do not want CA Disk to process password-indicated data sets, it can default to N.

  If your installation uses the security package SECURE, specify sysparm SECUSUPP with a value of Y. With SECURE, the value of the sysparm PASSWORD does not matter.

- See the setting of sysparm PASSNEWN. If you want CA Disk to allow the renaming of password-indicated data sets, specify the sysparm value as Y. If you do not want CA Disk to allow the renaming of password-indicated data sets, let it default to N.

Step 2. Activate CA Disk System Security

Security administrators should review CA Disk sysparms SPFUSRIDn and TSOUSRIDn in the *Systems Guide*, which apply to the ISPF and TSO functions.

See the Security Processing section in the *Systems Guide*. This documentation describes several security features along with the CA Disk security philosophy. Your installation's security administrators should determine the features applicable to your environment. You can activate the features now, or you can add them later.

Step 3. Prepare Security Packages to Work With CA Disk

If you are planning to run CA Disk and a security package on your system without purchasing a CA Disk Selectable Unit for the security package, follow the tasks outlined in the following step. The tasks show the minor changes needed to let CA Disk work properly with your security package.

Step 4. Install CA Disk Security Interface

If you have purchased one of the CA Disk security package Selectable Units, or if you are planning to run CA Disk under one of the security packages CA ACF2, CA Top Secret Security, or IBM's RACF, follow the appropriate parts of this installation step.

Review the setting of sysparm SECURVOL in your installation. Of special importance is its use in the volume-level functions of VBACKUP and VRECOVER. If it is left at its default value of Y, volume-level checking will be performed. This applies to all three CA Disk security interfaces.

Activate Miscellaneous Security Features

Certain CA Disk parmlib members contain information that you do not want your users to override. The SYSPARMS, SPFOPTNS, and DMCOPTNS members are good examples. This section describes how to restrict CA Disk to a list of authorized parmlib data sets. You can activate this protection now or later.

Parmlib Security — PARMAUTH

To restrict your users to a list of authorized CA Disk parmlibs or to activate the Storage Administrator's FACILITY Class Profiles, activate the Security Feature by installing user exit USERMOD5 as follows:

1. Locate the source for the Parmlib Security Feature in member PARMAUTH, located in the library associated with the CCUWSAMP DDDEF. The following illustration is a sample source for PARMAUTH:

```
PARMAUTH TITLE 'CA Disk SYSTEM PARAMETER DATA SET SECURITY'
*********************************************************************
         COMPILE ASEM=RENT,LKED=RENT                                 *
*                                                                    *
* DESCRIPTION:                                                       *
*    THIS IS A SAMPLE USERMOD USED TO TAILOR SYSPARM DATA SET        *
*    SPECIFICATION SECURITY.                  .                      *
*                                                                    *
*    READ THE PARMAUTH-MACRO PROLOG FOR SPECIFICATIONS OF OPTIONS    *
*    YOU MAY OVERRIDE.                                               *
*                                                                    *
PARMAUTH PARMAUTH SECURITY=NO,                                       X
               SECURLIB=SYS1.PARMLIB,                                X
               SECURTBL=ZDMSPARM,                                 @001X
               STGADMIN=NO,                                       @001X
               STGADLIB=SYS1.PARMLIB,                             @001X
               STGADTBL=ZDMSSTGA                                  @001
         END
```

2. To ensure that the changes you make to PARMAUTH are protected during future CA Disk installs or maintenance, copy this member into the source library associated with the //USERASM dd statement in USERMOD5.

3. Customize PARMAUTH as follows:

   - Activate the Parmlib Security Feature by specifying, YES to the SECURITY= parameter. The default is NO, which deactivates the feature.
   - Alter the SECURLIB= parameter as required. The value you specify is the name of the data set that stores the list of authorized parmlibs. The default is SYS1.PARMLIB. However, this can be any highly protected library available to CA Disk, TSO, and ISPF users with an LRECL of 80.

- Alter the SECURTBL= parameter as required. The value you specify is the name of the member that stores the actual list of authorized parmlibs. The default is ZDMSPARMS.
- Activate Storage Administer FACILITY Class Profiles by specifying, YES or LIBRARY to the STGADMIN= parameter. The default is NO, which deactivates the feature.
- Specifying YES causes CA Disk to search for the specified Table in the active PARMLIB with the name specified in the STGADTBL parameter. The member must reside in an authorized PARMLIB and requires that the SECURITY parameter be specified as YES.
- Specifying LIBRARY to the STGADMIN parameter will cause CA Disk to search for the specified Table in the Library specified in STGADLIB. This option does not require PARMLIB Security to be activated.

  Note: If you are currently running with PARMLIB Security deactivated and you want to activate Storage Administration FACILITY Class Profiles; you can use the LIBRARY option. The YES option requires a set of profiles in each Parmlib concatenation. Though this is very flexible, it is harder to administer and can be misused.
- Alter the STGADLIB= parameter, as required. Specify the name of the partitioned data set that stores the list of Functions and the corresponding FACILITY class profile. The default is SYS1.PARMLIB. However, CA Disk, TSO, and ISPF users with an LRECL of 80 can read any highly protected library, which is available.
- Alter the STGADTBL= parameter as required. Specify the name of the member that stores the actual list of Functions and the corresponding FACILITY class profile.
- The default is ZDMSSTGA. The distribution PARMLIB provides a sample member under the name SAMPZADM. The format of the table is documented in the chapter "PARMLIB" in the *Systems Guide*.
- Save your work by issuing SAVE at the TSO command line.

4. Create member to store your authorized parmlibs. The data set must match that specified for SECURLIB=, and the member must match that specified for SECURTBL=.

   If you allowed SECURLIB= to remain at its default value, you must create the member ZDMSPARM in your cataloged SYS1.PARMLIB data set.

   To create a list of authorized parmlibs, use the sample SAMPZDMS in PARMLIB.

To prevent a security exposure, only cataloged CA Disk parmlibs can be authorized. If a user creates an uncataloged parmlib, then tries to use the parmlib using the VOL=SER= parameter, CA Disk notes that the parmlib is uncataloged or incorrectly cataloged, issues a descriptive message and abends.

Note: If you are installing a new release and have created a test parmlib, include this test parmlib in the list. If you want to keep your production users from using your test parmlib, instruct your security package to restrict access.

Install USERMOD5

USERMOD5 must be installed using SMP/E. Sample JCL is provided for you in member USERMOD5 of the install library. Install USERMOD5 as follows:

1. Locate the source for USERMOD5 in the install library. The following illustration is a sample SMP/E JCL for USERMOD5:

```
EDIT                                          CCUWJCL(USERMOD5)   01.00  Columns
00001 00072
Command ===>                                                    Scroll ===> CSR
********************************************************************************
//JOBNAME  JOB (ACCT INFO)
//*
//* ********************************************************************
//* * INSTALL DISK PARMLIB AND SYSPARM                                  *
//* *   SYSPARM SECURITY                                                *
//* *   PARMLIB SECURITY                                                *
//* *                                                                   *
//* ********************************************************************
//*********************************************************************
//*  CHANGE CAISMPE TO YOUR SMPE PREFIX                                 *
//*  CHANGE USERID.DISK.ASM TO YOUR MODIFIED SOURCE LIBRARY             *
//*  CHANGE CAIT0 TO YOUR TARGET ZONE NAME                              *
//*  SEE THE NOTE BELOW REGARDING THE SMP/E MCS STATEMENTS              *
//*********************************************************************
//SMP     EXEC PGM=GIMSMP,REGION=5120K,
//             PARM='CSI=CAISMPE.CSI'
//USERASM   DD DISP=SHR,DSN=USERID.DISK.ASM
//SMPCNTL   DD *
SET BOUNDARY(GLOBAL).
RECEIVE S(SDU1815) .
SET BOUNDARY(CAIT0) .            /* <    CA DISK TARGET */
APPLY   S(SDU1815).
//SMPPTFIN  DD *
++USERMOD(SDU1815).
++VER(Z038) FMID(CCUWC50) .
++SRC(PARMAUTH) DISTLIB(ACUWSAMP) DISTMOD(ACUWMOD0) TXLIB(USERASM).
++SRC(SYSPAUTH) DISTLIB(ACUWSAMP) DISTMOD(ACUWMOD0) TXLIB(USERASM).
/*
```

1. Customize USERMOD5 by supplying the appropriate information for the following:
   - Job card information
   - Global and Target CSI names
   - USERASM data set name

2. Submit USERMOD5.

   After a successful APPLY run, CA Disk attempts to read your SECURLIB= data set. If you allowed SECURLIB= to remain at its default value and your security package protects SYS1.PARMLIB, you must perform one additional step to make the list of authorized parmlibs available to CA Disk. To accomplish this, do the following:

   - From the CA Disk load library, grant READ access to all programs for each CA Disk users.

   Note: If you have users running CA Disk TSO or ISPF functions under a TSO session that is not APF-authorized, these functions must obtain a shared enqueue to read SYS1.PARMLIB. If the shared enqueue can be obtained quickly, this should not be a problem as the read is very brief. However, if a shared enqueue cannot be obtained, the job goes into a wait state until the enqueue is available. CA Disk tasks running APF-authorized do not use an enqueue on SYS1.PARMLIB, and are not affected.

Activate CA Disk Storage Administration FACILITY Class Profiles

Storage Administration FACILITY Class Profiles allows an installation to control access to sensitive data. This feature is optional. The Storage Administration FACILTY Class Profiles feature removes the necessity for storage administrators to have security access to the data sets. Instead, it allows them access to the Storage Management Functions. The Storage Management Functions have unlimited access to the data sets.

To understand this concept, consider the following scenario. User 1 is a storage administrator whose job function is to do the backups, archives, and restores. User 1 is allowed to run the commands ARCHIVE, BACKUP and RESTORE in storage administration mode. Any CA Disk jobs that User 1 submits, that perform those specific commands, will execute in Storage Administration mode and User 1 will be able to process any data set. Any other job that User 1 submits, uses the standard security system profiles to verify access to a referenced data set, preventing User 1 from viewing or modifying sensitive data. In Storage Administration mode, CA Disk also allows the installation to restrict access to sensitive command operands, such as NEWNAME and NEWHLQ.

Note: Storage Administration FACILITY Class Profiles do not restrict access to commands, operands or data sets. The commands and operands can be used regardless of a user's access, or lack thereof, to the FACILITY Class Profiles. The Facility Class Profiles allows users designated as Storage Administrators and access to the FACILITY (command, operand or function) in use to use that FACILITY on data sets regardless of security system access restrictions.

Some operands of commands are individually authorized: such as NEWNAME and NEWHLQ. This is due to the potential to gain access to data that would normally not be accessible but would be so if the data set name were changed. Any command, which uses these operands, goes through two qualifications. First, the command itself must be authorized to the user. If the command is authorized, and one of these protected operands is specified, the user must be authorized. Both must be true for the command to execute in Storage Administration mode.

DMSAR is a special case in regards to the Storage Administration FACILITY Class Profiles. The function of DMSAR is authorized individually rather than the commands that DMSAR executes. DMSAR is a special case where RESTORE commands are executed, but the function being authorized is $AUTORES. In the profile list, $AUTORES would be specified rather than the RESTORE command. This function name will only be used for a procedure that is actually performing an auto-restore. So the userid that the procedure is executing will need to be permitted access to the resource defined in the profile list for $AUTORES.

A series of steps activates the use of Storage Administration FACILITY Class Profiles. The first step can be executed in any order, but the final step should always be the last one executed so that access to Storage Management mode is not accidentally granted to users who should not have it.

To activate the use of Storage Administration FACILITY Class Profiles

1. The Security Administrator has to permit each user who is to have access to Storage Management FACILITY Class Profiles READ access to the resource defined in the Sysparm SMSSTGAD. This designates the user(s) as a Storage Administrator.

   Note: If this resource is undefined, by default the user is a Storage Administrator. If a user is defined as a Storage Administrator, CA Disk will check the access to the Storage Administration FACILITY Class profile for the function.

2. Create a Storage Administration FACILITY Class profile list in a secured source PDS. A sample list with all of the supported functions, commands, and operands is in PARMLIB called SAMPZADM. The default resource name is STGADMIN.DMS.STGADMIN.command [.operand]. A command will not be able to execute in Storage Administration mode, if there is no definition for it in the profile list.

3. The profiles referenced in the profile list created in #2 above now have to be defined to the security system and the storage administrators have to be permitted access to them.

   - Protect all Storage Administration profiles as a blanket.

     For example, if you have CA Top Secret and are using the default profile names, issue the command:

     ```
     TSS ADDTO (deptacid) IBMFAC (STGADMIN.DMS.STGADMIN.**)
     ```

     If you have RACF, issue the command:

     ```
     RDEFINE FACILITY STGADMIN.DMS.STGADMIN.** +
     UACC (NONE) NOTIFY (security administrator)
     ```

     where *security administrator* is the user ID of a person to whom optional violation messages are to be sent.

   - If you want to allow your storage management group of users to be able to execute most or all of the CA Disk functions in Storage Administration mode, you can instruct your security package to allow that.

     If you have CA Top Secret, you can issue the command:

     ```
     TSS PERMIT (acid) IBMFAC (STGADMIN.DMS.STGADMIN.**) ACC (READ)
     ```

     where *acid* is the acid to whom you want to be able to use run-time sysparm overrides.

     If you have RACF, you can issue the command:

     ```
     PERMIT STGADMIN.DMS.STGADMIN.** CLASS (FACILITY) +
     ACCESS (READ) ID (storagemanagementgroup)
     ```

     where *storagemanagementgroup* is the group name (or list of userids) of your storage administrators.

   - If you want to restrict your storage management group from executing certain functions in Storage Administration mode, you can individually protect them.

     For example, if you have CA Top Secret, issue the commands:

     ```
     TSS ADDTO (deptacid) IBMFAC (STGADMIN.DMS.STGADMIN.command)
     TSS PERMIT (acid) BMFAC (STGADMIN.DMS.STGADMIN.command) + ACC (READ)
     ```

     where *acid* is the acid to you want to be able to use run-time sysparm overrides.

     If you have RACF, issue the commands:

     ```
     RDEFINE FACILITY STGADMIN.DMS.STGADMIN.command + UACC (NONE) NOTIFY
     (security administrator)
     PERMIT STGADMIN.DMS.STGADMIN.command CLASS (FACILITY) + ACCESS NONE)
     ID(storagemanagementgroup)
     ```

4. The PARMAUTH member has to be updated to activate the feature. Parameter STGADMIN has to be set to 'YES' or 'LIBRARY' to activate the feature. If STGADMIN is set to 'YES' it will also be necessary to specify the SECURITY parameter on the PARMAUTH macro as 'YES' and supply a list of authorized PARMLIBS. This is a security feature so that a profile list isn't used from an unsecured PARMLIB. The parameter STGADTBL will then need to be updated with the name of the PDS member that contains the profile list. If STGADMIN is set to 'LIBRARY' it will be necessary to supply both of the parameters STGADLIB and STGADTBL. Parameter STGADLIB will have to be specified with the data set name of the PDS and STGADTBL with the member name that contains the profile list. It is easier to administer a single list of profiles, so the use of the LIBRARY parameter is suggested.

Grant or Deny Access to System-Installed Intercept Modules

It may be necessary or desirable to either grant or deny access to modules that are installed as system intercepts. The DMSAR (Auto-Restore function) requires the access to load and activate several modules to install the intercepts that provide the Auto-Restore functions. There are also several tracing modules, which are often used to diagnose CA Disk, you might want to limit access to.

DMSAR requires access to the following modules:

- ADSAR010
- ADSAR026
- ADSHS001
- DIMCH400

There is also an IGGDASU2 exit that may be installed and its name is ADSDASU2. For more information about the purpose of the exit, see Installation of IGGDASU2 User Exit.

Following are the Tracing modules:

- ALLOCTRC
- ATLGTRC
- DYNALLOC
- PROGMTRC
- SCRTHTRC

By defining profiles in your security system and denying or permitting access to those profiles, you can grant or deny access to modules that are installed as system intercepts. The profiles are of the format CSVDYLPA.ADD.*modname* and CSVDYLPA.DELETE.*modname*, where *modname* is either a distinct module name or a module name pattern.

Important! If your installation does not deny access to these resources and you do not need to deny access to the tracing modules, you can skip the rest of this topic.

To grant access to a module name, you need to execute several commands. For DMSAR to work properly, access to several resources must be granted.

Example 1: Deny Access to DIMCH400

These example commands deny access to a module named DIMCH400, based on your security system.

Note: You must have specific authorization to issue these commands.

TOPSECRET Commands:

- TSS ADDTO (deptacid) IBMFAC (CSVDYLPA.ADD.DIMCH400)
- TSS ADDTO (deptacid) IBMFAC (CSVDYLPA.DELETE.DIMCH400)

RACF Commands:

- RDEFINE FACILITY CSVDYLPA.ADD.DIMCH400 +
- UACC (NONE) NOTIFY (security administrator)
- RDEFINE FACILITY CSVDYLPA.DELETE.DIMCH400 +
- UACC (NONE) NOTIFY (security administrator)

Example 2: Grant Access to DIMCH400

These example commands grant access to the DIMCH400 module by userid WXY0005.

TOPSECRET Commands:

- TSS PERMIT (WXY0005) IBMFAC (CSVDYLPA.ADD.DIMCH400) ACC (UPDATE)
- TSS PERMIT (WXY0005) IBMFAC (CSVDYLPA.DELETE.DIMCH400) ACC (UPDATE)

RACF Commands:

- PERMIT CSVDYLPA.ADD.DIMCH400 CLASS (FACILITY) +
- ACCESS (UPDATE) ID (WXY0005)
- PERMIT CSVDYLPA.DELETE.DIMCH400 CLASS (FACILITY) +
- ACCESS (UPDATE) ID (WXY0005)

You can also grant permission to access these FACILTY class resources to a group of people by replacing the userid WXY0005 by a group id or by multiple userids.

Note: It is necessary to perform the previous commands for each module to which you wish to grant or deny access.

Example 3: Define access rules for ACF2:

This example commands define access rules for ACF2:

```
ACF
Set Resource(FAC)
COMPILE *
$KEY(CSVDYLPA) TYPE(FAC)
ADD.DIMCH400 UID(wxy0005) SERVICE(UPDATE) ALLOW
DELETE.DIMCH400 UID(wxy0005) SERVICE(UPDATE) ALLOW
<enter>
STORE
END
```

Running this resource definition denies access to all CSVDYLPA resources except for the user WXY0005. User WXY0005 will be able to access CSVDYLPA.ADD.DIMCH400 and CSVDYLPA.DELETE.DIMCH400.

Note: ACF2 is RULE based and all of the rules for a particular $KEY must be included when that $KEY is recompiled. Any rules that are not included will not be retained. The rules can include wildcard characters in the userid and the resource names, and can use identification criteria other than userid. For more information about defining the access rules, see the CA ACF2 documentation and contact your ACF2 administrator.

Note: If the FACILITY resource class is specified as resident in the GSO INFODIR record, any rule changes or additions can be activated by issuing the following operator command:

F ACF2,REBUILD(FAC)

System Parameter Override Security — SYSPAUTH

CA Disk sysparm overrides are controlled through use of the //SYSPARMS dd statement. Use of this dd statement is described in the Overriding Sysparms Instream section in the Systems Guide. Users can override sysparms only if SYSPARMO is specified with a value of Y in the SYSPARMS member of your CA Disk parmlib.

If you do not have a security package that is compatible with SAF, you cannot limit access to sysparm overrides for a subset of users, or for a subset of sysparms.

If your security package is SAF-compatible, and if you want to restrict your users to a subset of CA Disk sysparms that they can override, or restrict sysparm overrides to a certain group of users, you can activate the System Parameter Override Security Feature by installing user exit USERMOD5 according to the following procedure.

To activate the System Parameter Override Security Feature by installing user exit USERMOD5

1. Locate the source for the System Parameter Override Security Feature in member SYSPAUTH, located in the library associated with the CCUWSAMP DDDEF. The following illustration is a sample source for SYSPAUTH:

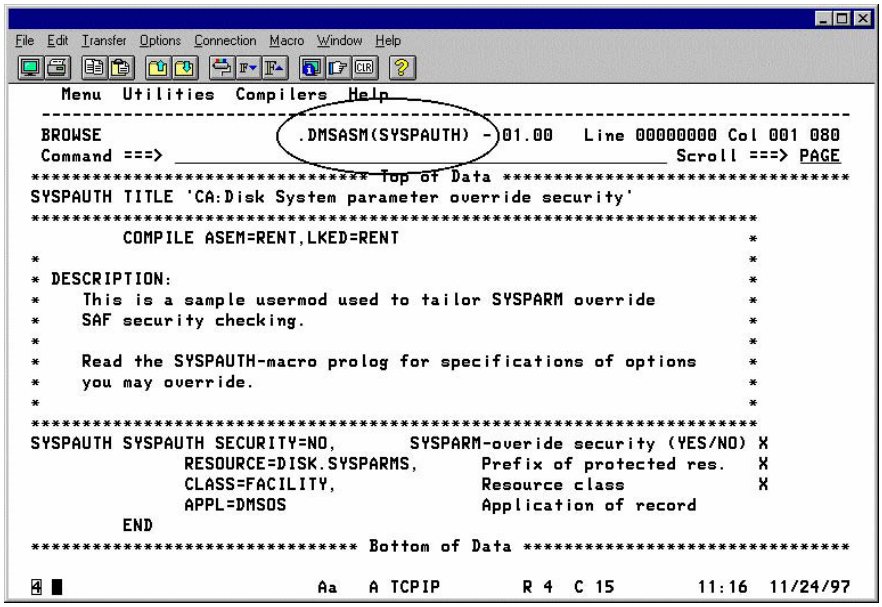

2. To ensure that the changes you make to SYSPAUTH are protected during future CA Disk installs or maintenance, copy this member into the source library associated with the //USERASM dd statement in USERMOD5.

3. Customize SYSPAUTH as follows:

   a. Activate the System Parameter Override Security Feature by specifying YES to the SECURITY= parameter. The default is NO, which deactivates the feature.

   b. Alter the RESOURCE= parameter as required. The value you specify is the prefix of the resource. The default is DISK.SYSPARMS.

   c. Alter the CLASS= parameter as required. We recommend you specify FACILITY for CA ACF2 and RACF security packages, or IBMFAC for CA Top Secret.

   d. The value specified for the APPL= parameter is recorded by SMF. The default value is DMSOS.

4. Save your work by issuing SAVE at the TSO command line.

5. Activate the FACILITY (IBMFAC if you have CA Top Secret) class of your security package. For example, if you have RACF, issue the command:

   ```
   SETROPTS CLASSACT(FACILITY) GENERIC(FACILITY) +
   GENCMD(FACILITY)
   ```

   When SYSPARMO is specified with a Y, CA Disk uses the FACILITY (IBMFAC if you have CA Top Secret) class of your security package to determine if the user has READ access to the resource DISK.SYSPARMS.*sysparmname* before allowing the override.

   Note: The check for READ access is with the FACILITY (IBMFAC if you have CA Top Secret) class, not the DATA SET class. Therefore the name does not refer to the name of an actual data set; there could be other rules on the DATA SET class that refers to data sets with that same name. RACF users must execute this command before creating the following generic profile.

6. Protect all sysparm overrides as a blanket. For example, if you have RACF, issue the command:

   ```
   RDEFINE FACIITY DISK.SYSPARMS.* +
   UACC(NONE) NOTIFY(security administrator)
   ```

   Where *security administrator* is the userid of a person to whom optional violation messages are to be sent.

   If you have CA ACF2, issue the command:

   ```
   TSS ADDTO(deptacid) IBMFAC(DISK.SYSPARMS.*)
   ```

7. To allow your storage management group of users to override most or all of the system parameters, you can instruct your security package to allow that. For example:

   - If you have RACF, issue the command:

     ```
     PERMIT DISK.SYSPARMS.* CLASS(FACILITY) +
     ACCESS(READ) ID(storagemanagementgroup)
     ```

     Where *storagemanagementgroup* is the group name (or list of userids) of your storage administrators.

   - If you have CA Top Secret, you can issue the command:

     ```
     TSS PERMIT(acid) IBMFAC(DISK.SYSPARMS.*) ACC(READ)
     ```

     Where *acid* is the acid to whom you want to be able to use run-time sysparm overrides.

8. To prevent your storage management group from overriding certain system parameters, you can individually protect them. For example:

   - if you have RACF, issue the commands:

     ```
     RDEFINE FACILITY DISK.SYSPARMS.sysparmname +
     UACC(NONE) NOTIFY(securityadministrator)
     PERMIT DISK.SYSPARMS.sysparmname CLASS(FACILITY) +
     ACCESS(NONE) ID(storagemanagementgroup)
     ```

   - If you have CA Top Secret, issue the commands:

     ```
     TSS ADDTO(depacid)IBMFAC(DISK.SYSPARMS.sysparmname)
     TSS PERMIT(acid) BMFAC(DISK.SYSPARMS.sysparmname) +
     ACC(READ)
     ```

     Where *acid* is the acid to whom you want to be able to use run-time sysparm overrides.

9. If you want to allow all users to override certain system parameters, such as ARCONAME and ARCCNAME, you can individually allow that. For example:

   - If you have RACF, issue the commands:

     ```
     RDEFINE FACILITY DISK.SYSPARMS.sysparmname +
     UACC(READ) NOTIFY(securityadministrator)
     ```

   - If you have CA Top Secret, issue the command:

     ```
     TSS ADDTO(deptacid)IBMFAC(DISK.SYSPARMS.sysparmname)
     ```

10. With this security feature activated, CA Disk allows the override of a sysparm only if CA Disk receives a return code of less than 8 from SAF using the macro:

    ```
    RACROUTE REQUEST=AUTH,
    RELATED='DMS/OS AUTH CHECK',
    APPL='DMSOS ',
    ATTR=READ,
    CLASS='FACILITY',
    ENTITY=DISK.SYSPARMS.sysparmname
    ```

    Where *sysparmname* is the 8-character name of the system parameter. CA Disk issues a message and ignore overrides of those sysparms to which the user does not have sufficient authority.

    Note: CA Top Secret translates this FACILITY class check to IBMFAC.

11. Submit USERMOD5 according to the instructions under USERMOD5.

Other Usermods

All other usermods can be installed using the installation instructions for USERMOD5. The following is a list of such usermods:

USERMOD6

Used to tailor MTF function, command, and parameter security. SAF is used for security checking. Read the MTFLOCK-macro prologue for specifications of options.

Source: MTFLOCK

USERMOD6

Used to tailor PET function/command/ parameter security. SAF is used for security checking. Read the PETLOCK-macro prologue for specifications of options.

Source: PETLOCK

USERMOD9

TSO exclusion and exemption table.

Source: TSOEXMPT

USERMOD9

DSN exclusion and exemption table.

Source: DSNEXMPT

USERMOD9

JOB exclusion and exemption table.

Source: JOBEXMPT

USERMOD9

PGM exclusion and exemption table.

Source: PGMEXMPT

USERMODB

When a data set is cataloged to the CA Disk pseudo volume, the catalog management hook (ADSAR010) is designed to return the SVC26 requestor the volser of MIGRAT. If the program name or pattern is listed in this table, the CA Disk pseudo volser will be returned instead. Read the SCRNLIST-macro prologue for specifications of options you can override.

Source: ADSUMODB

USERMODC

Used to specify CA Disk Auto Restore DB2/ASYNC recall delay time. Read the ADSUMODC-macro prologue for specifications of options you can override.

Source: ADSUMODC

USERMODD

Used to modify the CA Disk DDNAME for programs running in the TSO/ISPF environment. Read the TASKLIB macro prologue for specifications of options.

Source: TASKLIB

USERMODE

Used to specify the page line-count, limit for the CA Disk message processor. Read the LINEMAX-macro prologue for specifications of options you can override.

Source: LINEMAX

USERMODE

Used to specify an exit routine for the CA Disk message processor. Read the SYSOUTEX-macro prologue for specifications of options you can override.

Source: SYSOUTEX

USERMODF

Used to activate CA Disk catalog master password security. Read the CATLGPSW-macro prologue for specifications of options you can override.

Source: ADSVS939

USERMODG

Used to tailor MTF message diagnostics for MTF002I. Read the MTF002BY macro prologue for specifications of options you can override. In summary, use MTFMSG=YES to always issue the message; MTFMSG=NO to bypass the message unless MTF DD is active.

Source: ADSUMODG

USERMODK

Used to specify whether CA Disk should wait for RECON data set in case it is being used. Read the ADSUMODK module prologue for definition of values that can be specified.

Source: ADSUMODK

Activating MTFDEBUG DD Statement Security Feature

Using the //MTFDEBUG dd statement is not documented anywhere else in CA Disk documentation, and is intended for use by CA Disk technical staff only. The use of this dd statement can nevertheless be considered a security exposure.

CA Disk allows the use of the //MTFDEBUG dd statement only if it receives a return code of less than 8 from SAF using the macro:

```
RACROUTE REQUEST=AUTH,

RELATED='DMS/OS AUTH CHECK',
APPL='DMSOS ',
ATTR=ALTER,
CLASS='DATA SET',
ENTITY=DISK.USING.MTFDEBUG.DD
```

There is no CA Disk control over this check. If you do nothing, the return code is always less than 8, and use of the //MTFDEBUG dd statement is allowed. If you do not have a security package that is compatible with SAF, you cannot control access to this dd statement.

To activate the CA Disk //MTFDEBUG dd statement security feature

1. Through your security package, grant ALTER access to the data set name *DISK.USING.MTFDEBUG.DD*. A data set of this name does not need to exist; the name need only be significant to your security package.
2. If your SAF-compatible security package has protection-by- default and you want to allow certain users to use //MTFDEBUG dd statements, instruct your security package to grant these users ALTER access to *DISK.USING.MTFDEBUG.DD*.

   If your SAF-compatible security package does not have protection-by-default and you want to allow only certain users to use //MTFDEBUG dd statements, instruct your security package to deny access to *DISK.USING.MTFDEBUG.DD*, (universal access of NONE) and then grant ALTER access to the authorized users.

   CA Disk issues a message and ignore //MTFDEBUG dd statements from users without sufficient authority to that resource.

Activating VSAM Support

Perform the following procedure to activate the VSAM support.

To activate the VSAM support

1. Specify sysparm VSAMSUPP with a value of Y to activate the VSAM support.
2. A CA Disk master password is available to allow processing of any password-protected cluster. For security reasons, it is supplied and documented in a separate enclosure. See the documentation for use in your installation.
3. See the sysparm VSONLINE if your installation uses ICF catalogs and you have VSAM data sets defined on offline DASD volumes. For more information, see the sysparm description for VSDSPACEn in the *Systems Guide*.

4. IDCAMS uses the IDCUT1 and IDCUT2 dd statements in the RESTORE, RECOVER and DMSAR procedures to dynamically allocate workspace when building a VSAM alternate index. Use one of two approaches to control where this space is allocated:

   - Modify the data set names to direct the allocation to a desired VSAM or ICF catalog, and modify the VOL=SER to point to volume(s) where work space can be allocated.
   - Supply the necessary sysparms to cause CA Disk to dynamically allocate the needed dd statements prior to invoking BLDINDEX processing.

   If you select the first method (supplying the JCL statements directly), see the sysparms VSDEFCAT, VSBIXPSW and VSBIXCAT for consideration. The main drawback to this approach is that it forces all restore and recover jobs that use the same data set names for the work files to be single-threaded (that is, each job must wait until the previous restore has completed and released the exclusive enqueues on the work file data set names). This can cause major processing delays and frustration if there are many restore jobs waiting to be serviced. In addition, the enqueues are left outstanding even if no VSAM data sets need to be restored.

   If you select the second method (to leave the IDCUTx dd statements out of the JCL), simultaneous restore jobs can take place (as long as they don't require the same archive tape). When CA Disk restores a VSAM alternate index and determines that BLDINDEX processing needs to be invoked, it checks to see if the IDCUTx dd statements are allocated. If they are not, it generates a unique name for each of the dd statements and allocate them accordingly. Two sysparms govern the allocation of these dd statements.

   VSIDCUTPxxxxxxxx—used to specify the high-level node CA Disk is to use for the work file name. For instance, if VSIDCUTPLABS is specified, the generated work file name begins with LABS. If a value is not specified, the high-level node of the first alternate index being built is used. This sysparm is most useful in installations that have not converted to ICF and therefore needs to have the work data set placed on specific volume(s) owned by the VSAM catalog pointed to by the high- level node of the data set name.

Use SIDCUTVxxxxxx...yyyyyy to specify the volume(s) to be used for the work data set. Up to ten volumes can be specified in the list. To use more than one volume, make sure that each of the volumes is specified as exactly six characters (for example, *VSIDCUTVWRK800WRK802WRK806*). The names that CA Disk generates have the form:

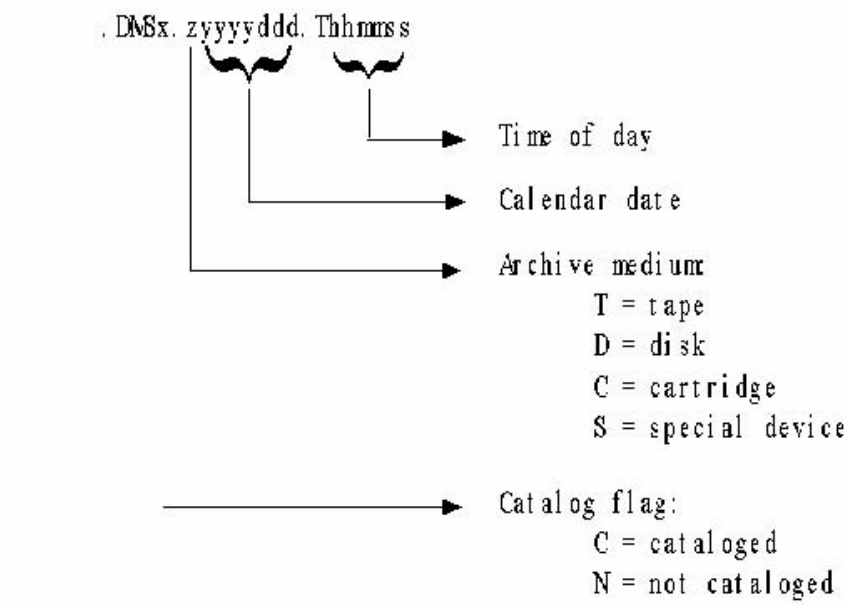

5. If your installation has vendor products--or internally developed application programs--that generate non-standard VSAM entry-sequenced data sets (ESDS) or linear data sets (LDS), see the sysparms VSACCESS and VSARCFMT. Products known to have these non-standard formats are the IBM DB2 database management system and certain products from MSA. These data sets must be archived in a control interval image copy format rather than in logical record format.

6. If your installation is using ICF VSAM support and you want to make use of the last use date support, you must make certain updates to your system to activate the date stamping mechanism (if your installation has not already done so). Date stamping can be requested only for VSAM clusters that are defined in ICF catalogs.

   Note: For more information, see the section CA Disk VSAM Date Stamp (see page 239).

Activating ISPF Support

The CA Disk ISPF support 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 interface to the user in an interactive environment.

After it is implemented, the CA Disk dialog management panels can be used to perform both foreground and background CA Disk functions. Through a simple modification to the ISPF primary option panel, the ISPF user can select CA Disk as a function in the same way other functions are selected. Once the CA Disk selection menu is entered, the user has access to almost all functions of CA Disk.

Authorizing the ISPF Dialog under CA ACF2

It is necessary to install a SAFDEF statement to allow the CA Disk ISPF dialog to continue functioning unauthorized.

Note: For more information, see the section Installing CA Disk Under CA ACF2.

Dynamic Menu-Formatting Feature

To protect against unauthorized use of a function, a dynamic menu-formatting feature is provided. This feature allows the CA Disk installation coordinator to specify what functions of CA Disk each user is authorized to use. When the CA Disk selection menu is entered, users are presented only the menu items for the functions for which they are authorized. Other functions are not accessible. Each user can be given a customized list of authorized functions. Entering control statements in a member of parmlib specifies all options.

Note: For more information, see the section Defining ISPF and DSCL User Options (see page 275).

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

The CA Disk ISPF function FRESTOR interfaces to the TSO RESTORE command processor. Specifying this option, as well as all others to be given to end users, is documented later in this section.

All CA Disk dialog management panels have associated HELP text panels. If at any time the user needs additional information on how to proceed or what to do, the HELP PF key can be pressed. The HELP text can be browsed until the user is ready to continue. The key takes them back to the panel they were originally processing.

Formatting Used for Online Reporting

Several dictionaries in the CA Disk parmlib govern the formatting used for online reporting. One of the options is the capability of including SMS constructs and attributes in the ISPF online reports. Users who want to change these dictionaries should have a working knowledge of the ISPF online reporting feature of CA Disk, be familiar with the ISPF section in the *User's Guid*e, covering the online report facility, and have had hands-on experience with the product. You must also define a report definition library to store your user-defined reports.

Note: For more information, see the section ISPF Online Reports in the *User Guide*.

To run the ISPF reports interactively, we recommend that the TSO region size be at least 3000K. Without the foreground report generation, a region size of 2000K should suffice for the report definition process.

Although the online reporting facility of CA Disk provides a lot of flexibility in defining user reports, it does not allow much tailoring as far as how individual fields are printed (column headings, numeric editing patterns, and so on.). You can tailor these values on an installation-wide basis, however, by modifying the data dictionary in parmlib.

Actually two different dictionaries are used for each report type, a low-level dictionary maps field names to data items in the F1 and F4 DSCBs. These dictionaries are called FMT1FLDS and FMT4FLDS and should not be modified by the user. They only provide a mapping mechanism between the high-level dictionary and the low-level data and have no control over report formatting options. The dictionaries that do control formatting of data fields are called FMT1DICT and FMT4DICT. The FMT1DICT member maps the fields used on the F1-DSCB reports and the FMT4DICT member is used for the volume summary reports (F4 DSCBs).

The ISPF online reports can include SMS constructs and attributes. If you have SMS managed data sets and choose to include these new fields in a report, you must rename the following members within the CA Disk parmlib data set:

- Rename member FMT1FLDS to FMTOFLDS
- Rename member FMT1DICT to FMTODICT
- Rename member FMTSFLDS to FMT1FLDS
- Rename member FMTSDICT to FMT1DICT

Note: If these renames were done while operating under a previous release of CA Disk and there were reports generated from the SMS constructs, you must issue these renames again prior to generating any reports.

Be careful when changing any fields that modify the width required to print a field if report definitions have been generated that use these fields. The column positioning values are calculated at report definition time and a change in the length can cause the output report to be incorrect (overlapped fields, and so on). If this occurs, the problem can be corrected by going through the report modification panels for each report using the affected field(s). No updates need to be made; the field adjustments are automatically made as a by-product of the update process.

Note: For detailed information about each of the fields in the dictionary, see the section ISPF Online Reports in the *User Guide*.

The ISPF support is shipped to the user on the CA Disk distribution tape or the in the ESD download as four separate files. The following is a list of the file contents:

ISPF Dialog Manager Panel

Contains the panel definitions for all CA Disk dialog manager panels and contains all HELP facility panels relating to the function and menu panels.

ISPF Dialog Manager Skeleton JCL

Contains skeleton JCL used for building JCL that is submitted for processing from CA Disk ISPF applications.

ISPF Dialog Manager Message

Contains error messages to be printed by the dialog manager when the CA Disk application detects an abnormal condition.

ISPF Dialog Manager Load Modules

Contains all CA Disk load modules. The dialog manager to obtain all CA Disk load modules.

Connecting the ISPF Libraries to Your System

Allocate the CA Disk dialog manager libraries to your TSO session by performing the following procedure. This must be done before invoking ISPF.

To allocate the CA Disk dialog manager libraries to your TSO session

1. Determine how your installation allocates the data sets referred to by the following ddnames to the TSO session prior to invoking the ISPF facility. They can be allocated by including them in the TSO LOGON JCL procedure or through a CLIST that is executed sometime before ISPF is invoked.

   CCUWPNL0

   DDName Library: ISPPLIB

   CCUWSKL0

   DDName Library: ISPSLIB

CCUWMSG0

DDName Library: ISPMLIB

CCUWLOAD

DDName Library: ISPLLIB

2. Update the allocation of these libraries to concatenate the CA Disk ISPF libraries after them. The names of these libraries correspond to the ddnames to which they should be concatenated.

3. Check the block sizes of the libraries being concatenated. The block size of the first library must be greater than or equal to the second, or errors occur. This can be corrected by either reblocking the libraries or simply putting DCB=BLKSIZE= a larger value on the dd statement for the first data set of the concatenation.

   For example, JCL before libraries are concatenated:

```
//ISPPLIB DD DISP=SHR,DSN=ISP.vrm.ISPPLIB
//ISPSLIB DD DISP=SHR,DSN=ISP.vrm.ISPSLIB
//ISPMLIB DD DISP=SHR,DSN=ISP.vrm.ISPMLIB
//ISPLLIB DD DISP=SHR,DSN=ISP.vrm.ISPLLIB
```

   JCL after CA Disk libraries are concatenated:

```
//ISPPLIB DD DISP=SHR,DSN=ISP.vrm.ISPPLIB
//        DD DISP=SHR,DSN=CA.DISK.CCUWPNL0
//ISPSLIB DD DISP=SHR,DSN=ISP.vrm.ISPSLIB
//        DD DISP=SHR,DSN=CA.DISK.CCUWSKL0
//ISPMLIB DD DISP=SHR,DSN=ISP.vrm.ISPMLIB
//        DD DISP=SHR,DSN=CA.DISK.CCUWMSG0
//ISPLLIB DD DISP=SHR,DSN=ISP.vrm.ISPLLIB
//        DD DISP=SHR,DSN=CA.DISK.LOADLIB
```

Customizing the CA Disk ISPF Support

Use one of the following methods for installing the ISPF interface:

Method #1

This method has the advantage of reducing EXCPs associated with CA Disk ISPF and thereby reducing execution time. It has the disadvantage of increasing the number of concatenated libraries to ISPLLIB and thereby increasing the number of directory blocks searched for non-CA Disk load modules. Placing the CA Disk load library last in the ISPLLIB concatenation can minimize this disadvantage. To use this method, perform the following steps:

1. Alter all TSO/ISPF logon CLISTS by placing the CA Disk load library last in the ISPLLIB concatenation.

2. Add the following line to ISR@PRIM:

```
8,'CMD(ADSSP203) NEWAPPL(ISR)'
```

Method #2

This method has the advantage of decreasing the number of concatenated libraries to ISPLLIB and thereby decreasing the number of directory blocks searched for ISPF functions. It has the disadvantage of increasing the number of EXCPs associated with CA Disk ISPF functions. This increase is incurred as CA Disk modules are loaded. This load process requires a directory search for each individual module, which is not required in Method #1. To use this method, perform the following steps:

1. Alter all TSO/ISPF logon CLISTs by adding an allocation for the CA Disk load library to DMSOSLIB file.
2. Add the following line to ISR@PRIM:

```
8,'PGM(ADSSP202) PARM(ADSSP044) NEWAPPL(ISR)'
```

Activating TSO Support

Follow these instructions to activate the support.

Step 1. Make the Command Processor Available to TSO

Ensure that the dialog manager programs are available to the TSO session. This process has been partially completed, since you have already installed the TSO Command Processors, and control program ADSSP202 into one of the system linklist libraries.

To complete the process, modify your ISPF/TSO logon procedure to allocate a ddname of DMSOSLIB that points to the CA Disk load library; that is:

```
//DMSOSLIB DD DISP=SHR,DSN=CA.DISK.LOADLIB
```

Also, add to your ISPF/TSO logon procedure a ddname of PARMLIB that points to your CA Disk parameter library:

```
//PARMLIB DD DISP=SHR,DSN=CA.DISK.PARMLIB
```

This step is required to give the TSO Command Processors access to other CA Disk programs. The CA Disk control module ADSMI002 always looks for the DMSOSLIB library first, and finds all of the needed CA Disk modules there without searching any other libraries. Your other ISPF functions never waste time searching the CA Disk directory looking for one of their programs.

TSO command processors are available to perform the functions in the following list of modules:

DARCHIVE

Queue a request to archive a data set.

DRESTORE

Queue a request to restore a data set.

DERASE

Erase a queued archive or restore request.

LISTDMS

List the CA Disk index of archived data sets.

LISTREQ

List the status of queued archive/restore requests.

RESTORE

Restore a data set immediately (uses dynamic tape allocation). This command processor must also be designated as privileged.

Step 2. TSO Help Text

The Help text members were placed into the CA Disk TSO Help library during the installation. Ensure that this library is allocated to your TSO session.

Step 3. TSO Dynamic Restore Command Processor

If you want to permit the use of the dynamic RESTORE command, see the instructions that follow to activate this feature. For more information on the dynamic RESTORE command, see RESTORE—Immediate (Dynamic) Restore in the *User Guide*.

The TSO dynamic RESTORE command processor uses a subfile called DMSPARMS, defined within the FILES. This subfile was described in the FDSAMPLE member in parmlib.

The dynamic restore command processor must be designated as privileged or authorized in the MVS environment, due to its use of the ENQ and ALLOCATE SVCs. Update your SYS1.PARMLIB(IKJTSOxx) member to include RESTORE. After the changed member has been activated by an IPL, any TSO user can do dynamic restores. TSO mount authority is required if the data set being restored is archived to tape.

Note: Standard ISPF does not allow authorized TSO commands to be entered from its menu option 6. Unless local ISPF modifications have been made to permit such executions, you must exit ISPF before using this command.

Activating DASD Billing

Follow these instructions to activate the support.

1. A RESTART data set is required for DASD billing. It is used to clear the billing file and in proper recovery after an abend. It is referenced by the //RESTART dd statement in the EXTEND procedure, and defaults to DSN=&Q..RESTART.
2. Create the RESTART data set, and allocate it with the following attributes:

```
//RESTART  DD DCB=(DSORG=PS,RECFM=F,LRECL=96,BLKSIZE=96),
//         SPACE=(CYL,(2,2))
```

Activating the DBRC Interface

Using the DBRC interface is optional. You can skip this step unless you use DBRC (Database Recovery Control) and are running MVS with IMS Release 4.1 or above. Even if you are, this step is still optional.

The performance, ease-of-use, data compression and other attributes of CA Disk provide an alternative to using the IBM image-copy program for backing up IMS databases. The CA Disk backups can be written to disk, tape, or both and managed with the flexibility standard to all CA Disk backups. If you still want to use IBM image-copy. Putting the image copies on disk and then using the CA Disk Sequential Migration to Tape function to free the disk space can obtain benefits. The degree of operator intervention, primarily for tape handling, can be reduced significantly.

The purpose of the DBRC interface in CA Disk is to provide an automated means of notifying DBRC whenever CA Disk takes any action against a DBRC-registered data set. This includes the direct processing of database data sets themselves, such as in backup functions, or in moving IBM image-copy data sets that have been placed on disk.

If you activate the DBRC interface, CA Disk will extract the database names and their associated data sets from the DBRC RECON data set and keep the list in memory. When CA Disk is taking a backup (or archive) copy of one of these data sets, a control statement for the IBM DBRC update utility (program DSPURX00) is created to notify DBRC that a *user image-copy* has been taken. Similarly, the restore (or recover) functions generate control statements to notify DBRC that the recovery has been performed.

If you take IBM image-copy backups and place them on disk instead of tape, the Move/Copy and Sequential Migration to Tape functions can encounter them (intentionally or otherwise) and move them to new locations. Instead of extracting the database data set names as described above, these functions extract the image-copy records from the DBRC RECON data set. When an image-copy data set is moved, a control statement to change the volume list in the DBRC image-copy record is produced.

The DBRC utility program itself can be executed as either the last step of the job that produces the control statements, or later.

Since databases are often managed separately, it is desirable to supply the information as outlined in the next section, but only for the specific jobs designed to need it. Instead of making changes directly to the CA Disk JCL procedures and activating the support for all jobs, this information can be provided only for the specific jobs intended; for example, the database backup and recovery jobs.

The RECON data sets are dynamically allocated based upon the names provided in the IMS resident load library. CA Disk determines the current RECON data set and uses it. Alternatively, you can provide them in the JCL through //RECON1, 2 and 3 dd statements.

The generated control statements and any messages relating to the DBRC processing are printed in a separate output listing for easier review. //DBRCPRT is dynamically allocated for this purpose, but can be supplied in the JCL if desired. If it appears that the control statements are not being generated, supply the //RECONDMP in the JCL to produce a listing of the data sets from the RECON data set. Match the listing against the data sets processed by CA Disk. If the data set appears in both the listing and the CA Disk job, verify member DBRCDSNS of parmlib.

The control statements produced are of the following general format. For fast path databases with replication, messages are issued to indicate that CA Disk has processed an area data set. However, no control statements are produced. If the data set is given a new name when the backup copy is taken, the new name appears in the UDATA field.

For backup:

```
NOTIFY.UIC DBD(DBAHSP00) DDN(DBAHSP01) -
RUNTIME(862991635000) UDATA('*** DMS ***')
```

For recovery:

```
NOTIFY.RECOV DBD(DBAHSP00) DDN(DBAHSP01) -
RCVTIME(862991635000) CURRENT
```

For migration of image-copy data sets:

```
CHANGE.IC DBD(DBAHSP00) DDN(DBAHSI01) -
RECTIME(862671605243) VOLLIST(WRK800)
```

To Activate the DBRC Support

To activate the DBRC support perform the following procedure.

To activate the DBRC support

1. Provide CA Disk with the name of the IMS resident load library in sysparm DBRCRLIB. CA Disk gets the names of the RECON data sets from here. The library is dynamically allocated when it is needed.

   Alternatively, provide a //RESLIB dd statement to identify the library. The ARCHIVE, RETAIN, RESTORE, RECOVER, DMS, CONFIGR and MIGRATE procedures potentially needs the library.

2. Provide CA Disk with the release number of IMS you are using by specifying sysparm DBRCRLSE. The default value is for Release 1.3.0. Enter the value without periods: 130.

   The format of the RECON data sets changed in IMS Release 2.0.0., so for CA Disk to determine the format, the release is needed.

3. Create member DBRCDSNS in the CA Disk parmlib data set.

4. Enter the following in member DBRCDSNS: explicit, pattern names, or both of IMS data sets , image-copy data sets or both for which CA Disk DBRC processing is to be done.

   A sample is provided in member SAMPDBRC of parmlib. If the name of the data set processed does not have a matching entry in this member, no further DBRC processing takes place.

5. Ensure that the naming conventions in member DBRCDSNS provide for easy identification. When they do, this feature provides a more efficient method of determining that DBRC processing is not needed.

   A pattern of / causes every data set to be examined against the list of names extracted from the RECON data set, while an empty member effectively turns off the DBRC support.

   Note: Removing a pattern name or by supplying a pattern that never matches any data set names, also inactivates the DBRC support.

6. Provide a disk data set in which DBRC control statements can be written and identify this data set to CA Disk with a //DBRCCARD dd statement in those functions where it is needed. The data set can be any sequential or partitioned data set that supports 80-byte fixed-length records.

7. Provide JCL for the DBRC update utility with the //SYSIN dd statement pointing to the same data set as in step 6 above.

8. Activate the support by specifying sysparm DBRCSUPP with a value of C.

Activating the CA Disk CA Auditor for z/OS Support

The CAIXXGN1 module is provided for those customers who do also have CA Auditor for z/OS (previously known as CA Examine) and want to include CA Disk to the products audited. If you will be using the CA Disk CAIXXGN1 module, it must reside in the currently active LPA.

Activating the Unicenter® Service Desk Support

Use of the Unicenter Service Desk (USD) interface feature of CA Disk is optional. In order to activate this feature you must have installed the CA Common Services CAISDI components CAISDI/soap and CAISDI/els as described in the *CA Common Services r11 Getting Started Guide*. The event members used by CA Disk must be copied from the Install library to the event library, CAI.CAIEVENT. CAI.CAIEVENT is created as part of the CA Common Services CAISDI/els component installation. A sample job, XGNCOPY, can be found in the Install library. If necessary, edit the member and change the data set names to reflect the names used in your installation and submit the job.

The CAI.CAIEVENT library includes control members used to define CA products to the CAISDI/els interface. Each event member contains text and other control information used in opening USD request tickets by one of the CA products supporting this interface. CA Disk event members begin with GN. The product control member is named XGNCNTL and can be found in the Install library.

For information on activating and customizing the USD interface, see the *CA Common Services r11 Unicenter Service Desk Integration Guide*. For further information on the CA Disk events that can result in USD request tickets being automatically opened, see the section Customizing the Unicenter Service Desk Support.

Note: It is possible to prevent specific CA Disk request tickets from being opened. If any request ticket is not desired, you may simply remove or rename the associated event member in CAI.CAIEVENT.

Activating the CA Tape Encryption CA Disk Interface

You can encrypt and decrypt all archive tapes when the CA Tape Encryption Interface to CA Disk is installed and activated. The BTE Interface with CA Disk is activated through sysparms BTEITAPE (PRIM, COPY or BOTH) and BTEDATCL (DATACLAS).

Note: For more information on these sysparms, see the *Systems Guide*.

CA Disk interfaces to the CA Tape Encryption started task through the DCB-exit that indicates that encryption should be performed. The OPEN intercept will still need to check SMS ACS to determine which encryption method to use.

Tailoring and Other Considerations

Review the following topics in the Tailoring Options section of the *Systems Guide:*

- Archive/Backup Considerations
- Processing PDSs that Contain Anomalies
- Implementing Support for StorageTek Redwood
- Implementing Support for IBM Magstar
- ISPF Custom Reports
- User-Specified Condition Codes

Note: To implement any special processing controls indicated for your installation environment, see the Special Considerations and General Restrictions sections of the *User Guide*.

Archive/Backup Considerations

CA Disk permits numerous tailoring options for archiving and backup functions. Some of the more important options to consider are discussed in this section.

Activating Data Compression

As an optional feature, CA Disk can compress data as it is being written to the archive or backup data sets, whether they are on tape or disk.

If you activate data compression, the default technique is technique number 0. CA Disk-supplied technique number 0 typically reduces the space used by 40 to 50 percent, which yields the same reduction in the amount of tape used. However, if you archive to disk instead of tape, the disk savings achieved can be closer to three or four to one, or even much higher. This is due to the optimized storage techniques in the archives, as well as wasted space in the original data sets.

CA Disk users who are also licensed to use CA's CA Compress product, can specify technique number 1 to use CA Compress as the compression method. CA Disk supplies a formal interface to use this compression method. Instructions for implementation are listed in Using CA Compress within the Archives.

As a direct result of compression, the number of I/Os to the archive media is also reduced by the same percentage, which can provide a significant improvement if I/O contention is a problem in your installation.

To achieve these benefits, however, expect the CPU time to increase. For more information, see the section Data Compression Techniques in the chapter "Algorithms."

The following sysparms are used by data compression. See them for applicability in your installation:

Sysparm	Description
DCDATACP	turns data compression off or on
DCCOMPTC	indicates the compression technique number
DCCMPEXn	exit for compression technique n
DCDCPEXn	exit for decompression technique n
DCDSNDEX	user exit for deciding if data compression is to be done for each data set, and what technique is to be used
DCEXCTBL	exclusion table name
DCINCTBL	inclusion table name (overrides DCDATACP"N")
DCRTSTAT	turns compression statistics off, on, or on with details for each data set compressed

Using CA Compress within the Archives

CA Disk users who are licensed to use CA's CA Compress product can use CA Compress to compress the CA Disk archives. If CA Compress is installed on your system, follow the following steps to implement this compression method.

1. Create a CA Compress File Description Table (FDT) to use when archiving/backing up data. To do this, create a CA Disk backup tape (NOT a disk data set) from a normal CA Disk backup run, but without using data compression (that is, sysparm DCDATACP specified with a value of N). Use this backup tape as input to the CA Compress File Prepass Utility. For specific instructions on creating the FDT, see the *CA Compress for MVS User Guide*.

2. Convert the FDT to load module format using the CA Compress FDTLOADR Utility. Give the FDT an 8-character name ("DMSOSFDT" is recommended) and place it in the CA Disk load library. For specific instructions on using the CA Compress FDTLOADR Utility, review the *CA Compress for MVS User's Guide*.

3. Specify sysparm SHRNKFDT with the value of the FDT name.

   Note: The default value for this sysparm is *DMSOSFDT*. If you have specified this name as the name of the CA Compress FDT in the load library (step 2 previous), it is not necessary to specify this sysparm.

   If you specified a different name, provide your name as the value for this sysparm in member SYSPARMS in the CA Disk PARMLIB.

4. Link the CA Compress interface with CA Disk using the following JCL.

   Note: If you install a new version of CA Compress, you must relink this interface. The JCL and control statements are provided in member RELKSHRK of the installation library.

```
//JOBNAME  JOB (ACCT INFO)
//* ************************************************************
//* * RELINK CA Disk MODULES FOR SHRINK COMPRESSION             *
//* ************************************************************
//S1LKED   EXEC PGM=IEWL,PARM=(LIST,LET),REGION=3072K
//SYSPRINT  DD  SYSOUT=*
//SYSLMOD   DD  DISP=SHR,DSN=CA.DISK.LOADLIB
//SYSLIB    DD  DISP=SHR,DSN=CA.DISK.ACUWMOD0
//SHRINK    DD  DISP=SHR,DSN=CA.DISK.COMPRESS.LOADLIB
//SYSUT1    DD  UNIT=SYSDA,SPACE=(1024,(50,20))
//SYSLIN    DD  *
 INCLUDE SYSLIB(ADSUT410)
 INCLUDE SHRINK(SHRKEXPD)
 ENTRY ADSUT410
 NAME ADSUT410(R)
 INCLUDE SYSLIB(ADSUT411)
 INCLUDE SHRINK(SHRKEXPD)
 ENTRY ADSUT411
 NAME ADSUT411(R)
/*
```

5. Concantenate the CA Compress load library to the CA Disk load library within the following PROCS:

   - DMS
   - DMSAR
   - FMS
   - RECOVER
   - RESTORE

   This ensures the CA Compress utilities are available to CA Disk, if needed.

6. Exempt all CA Compress-controlled data sets from compression during archival or backup processing. This can be accomplished either by placing the data set names or patterns in a data compression exclusion table as a PARMLIB member and specifying sysparm DCEXCTBL with the name of the PARMLIB member, or by using the CA Disk user exit DCDSNDEX.

   The reason for excluding these data sets is that files already compressed are not compressed any further when archived or backed up. Excluding them from compression during archive or backup processing decreases CPU overhead and improves performance.

7. Specify the following required data compression sysparms:

Sysparm	Description
DCDATACP	activate data compression
DCCOMPTC1	use CA Compress as the data compression method for the CA Disk archives

8. Optionally, you can specify the following data compression sysparms:

Sysparm	Description
SHRNKFDT	to specify the FDT name if other than the default of *DMSOSFDT*
DCDSNDEX	8 character user exit for deciding if data compression is to be done for each data set, and what technique to be used
DCEXCTBL	exclusion table name (cannot be used with DCINCTBL)
DCINCTBL	inclusion table name (cannot be used with DCEXCTBL)
DCRTSTAT	turns compression statistics off, on, or on with details for each data set compressed

Note: You cannot change or delete the FDT once you have compressed CA Disk archive files using CA Compress Data.

Specifying the Archive Medium for Archive/Backup

By default, CA Disk dynamically allocates both the primary and duplicate copy to tape. The DD statement used for the primary is //ARCHIVE0 and for the duplicate copy is //ARCHIVEC. Sysparms ARC0TYPE and ARCCTYPE allow you to change the defaults from tape to either disk or 3480 cartridge units. Specify these sysparms with a value of DISK, TAPE or 3480, and place them in the SYSPARMS member of the PARMLIB data set.

There are 2 ways to indicate that no duplicate copy is to be made:

1. Specify *NULL* to sysparm ARCCTYPE as *NULL*
2. In JCL, specify *DUMMY* to the //ARCHIVEC DD statement

For example, to direct the primary copy to disk and the duplicate copy to a 3480 cartridge unit, specify ARC0TYPEDISK and ARCCTYPE3480.

You can also specify the archive medium through JCL. If tape or 3480 devices are specified using JCL, the devices allocated at job initiation are used throughout the run; that is, the device are not deallocated at the end of each tape, but only at the end of the job. The dynamic allocation option can be made to behave in much the same way by specifying sysparm DYNUKEEP. The only exception is when more than five volumes are needed to contain a single archive data set, in which case dynamic allocation is used.

You can also specify archival to disk in JCL by allocating any disk device to the //ARCHIVE0 DD statement. For example:

```
//ARCHIVE0 DD UNIT=SYSDA,VOL=SER=anyvol,DISP=SHR
```

Notice that this only indicates to CA Disk to use disk archive. The volume (or data set) named on the DD statement has no meaning to CA Disk, and dynamic allocation of the ARCHVOL is still done.

Specifying the Archive Medium for Merge

By default, CA Disk dynamically allocates both the primary and duplicate copy to tape, and the tertiary copies are suppressed. The DD statements used for the primaries are //ARCHIVP1-5, the duplicate copies are //ARCHIVC1-5, and the tertiary copies are //ARCHIVC1-5. Sysparms MERPxTYP, MERCnTYP, and MERTnTYP allow you to change the defaults to tape, disk, or 3480 cartridge units as needed. Specify these sysparms with a value of DISK, TAPE, 3480, or DYNn/DYnn, and place them in the SYSPARMS member of the PARMLIB data set.

You can use the following two methods to indicate that no duplicate copy is to be made:

- Specify NULL to sysparm MERCnTYP as NULL.
- In JCL, specify DUMMY to the //ARCHIVC1-5 DD statements.

Because the tertiary copy is suppressed by default, MERTxTYP would need to be specified or the //ARCHIVT1-5 dd statement would have to be included. For example, to direct the primary to disk and the duplicate copy to a 3480 cartridge unit, specify MERP1TYPDISK and MERC1TYP3480.

You can also specify the archive medium through JCL. If tape or 3480 devices are specified using JCL, the devices allocated at job initiation are used throughout the run; that is, the devices are deallocated only at the end of the job, but not at the end of each tape. The dynamic allocation option can be made to behave in much the same way by specifying the sysparm DYNUKEEP. The only exception is when more than five volumes are needed to contain a single archive data set, in which case dynamic allocation is used.

You can also specify archival to disk in JCL by allocating any disk device to the //ARCHIVC1 DD statement. For example:

```
//ARCHIVC1 DD UNIT=SYSDA,VOL=SER=anyvol,DISP=SHR
```

Notice that this only indicates to CA Disk to use disk archive. The volume (or data set) named on the DD statement has no meaning to CA Disk, and dynamic allocation of the ARCHVOL is still done.

Using Multiple Types of Media

This section describes how CA Disk can utilize multiple types of media (that is, 3480s, 3490Es, SILOs, and so on), sysparms to use, IDRC considerations, data compression, and MERGE processing.

Selecting the Output Media

The user selects the media CA Disk writes to by specifying a value to the following sysparms, which are often entered as overrides in CA Disk jobs or kept in the SYSPARMS member in PARMLIB:

Function	Primary Sysparm	Copy Sysparm	Tertiary Sysparm	Notes
ARCHIVE	ARC0TYPE	ARCCTYPE	NA	NA
BACKUP	ARC0TYPE	ARCCTYPE	NA	NA
MERGE	MERPnTYP	MERCnTYP	MERTcTYP	For details, refer to LIMITS= in the User Guide.
XCOPY	MERP1TYP	MERC1TYP		

The following table shows the media options that can be specified with the previous system parameters. The primary and its copy can be directed to different media. These two output media types are synchronized so that if one is shorter than the other and fills up, both are closed and new tapes are mounted and/or new disk archive data sets are allocated.

Other sysparms that are in effect when the respective media option is used are shown in the following table:

Media Option	Type of Device	Unit for Writing or Reading	Unit for Auto-Restore	Tape Length
TAPE	Round tape (default)	DYNTUNITtape	ARESUNITtape	TAPEFEET2300
3480	3480/3490 cartridge	DYNCUNIT3480	ARESUNIC3480	CARTFEET0494
DYN1	Special device type 1	DYN1UNIT	ARESUNI1	TAPEFEET2300
DYN2	Special device type 2	DYN2UNIT	ARESUNI2	TAPEFEET2300
DYN3	Special device type 3	DYN3UNIT	ARESUNI3	TAPEFEET2300
DYN4	Special device type 4	DYN4UNIT	ARESUNI4	TAPEFEET2300
DYN5	Special device type 5	DYN5UNIT	ARESUNI5	TAPEFEET2300
DYN6	Special device type 6	DYN6UNIT	ARESUNI6	TAPEFEET2300
DYN7	Special device type 7	DYN7UNIT	ARESUNI7	TAPEFEET2300
DYN8	Special device type 8	DYN8UNIT	ARESUNI8	TAPEFEET2300
DYN9	Special device type 9	DYN9UNIT	ARESUNI9	TAPEFEET2300
DY10	Special device type 10	DY10UNIT	ARESUN10	TAPEFEET2300

Media Option	Type of Device	Unit for Writing or Reading	Unit for Auto-Restore	Tape Length
DY11	Special device type 11	DY11UNIT	ARESUN11	TAPEFEET2300
DY12	Special device type 12	DY12UNIT	ARESUN12	TAPEFEET2300
DY13	Special device type 13	DY13UNIT	ARESUN13	TAPEFEET2300
DY14	Special device type 14	DY14UNIT	ARESUN14	TAPEFEET2300
DY15	Special device type 15	DY15UNIT	ARESUN15	TAPEFEET2300
DY16	Special device type 16	DY16UNIT	ARESUN16	TAPEFEET2300
DY17	Special device type 17	DY17UNIT	ARESUN17	TAPEFEET2300
DY18	Special device type 18	DY18UNIT	ARESUN18	TAPEFEET2300
DY19	Special device type 19	DY19UNIT	ARESUN19	TAPEFEET2300
DY20	Special device type 20	DY20UNIT	ARESUN20	TAPEFEET2300
DISK	DASD volumes	n/a	n/a	n/a
NULL	Disables Copy sysparms	n/a	n/a	n/a

Specifying the ESOTERIC Unit Name

The media options previously shown relate to system parameters that provide the esoteric unit name CA Disk uses to allocate the proper device for writing, reading and auto restoring. These system parameters are placed in the SYSPARMS member in PARMLIB and should not be changed unless the esoteric name changes, or a media type has been totally removed from CA Disk control. The DYN*n*/DY*nn* sysparm provides the flexibility of defining additional device types to CA Disk that can be installed at various sites.

When a tape is created, CA Disk stores the device type in the ARCHVOLS record along with a flag if the DYN1, through DY20 option was used. With this information, CA Disk retrieves the esoteric unit name from the DY__UNIT or ARESUN__ sysparm so the proper device can be allocated.

Naming the Archives

In theory, the names you assign to archive data sets are of no consequence to CA Disk. CA Disk keeps track of the archive data sets regardless of their names. However, other issues dictate that naming conventions be observed, including the IBM operating system and tape management concerns. CA Disk generates unique archive data set names, providing flexibility in managing the archives and allowing them to be cataloged. Unique data set names also eliminates contention during dynamic allocation that would occur if non-unique data sets were used.

CA Disk sysparm defaults generate unique archive data set names. The base names are taken from either sysparm ARC0NAME or ARCHIVE0 DD statement (for the primary copy) and sysparm ARCCNAME or ARCHIVEC DD statement (for the duplicate copy). If both copies are being produced, the base names should be different—such as DMS.ARCHPRIM and DMS.ARCHCOPY.

The following figure shows Generated 22 Character DSNAME on specifying up to 22 characters for the base name:

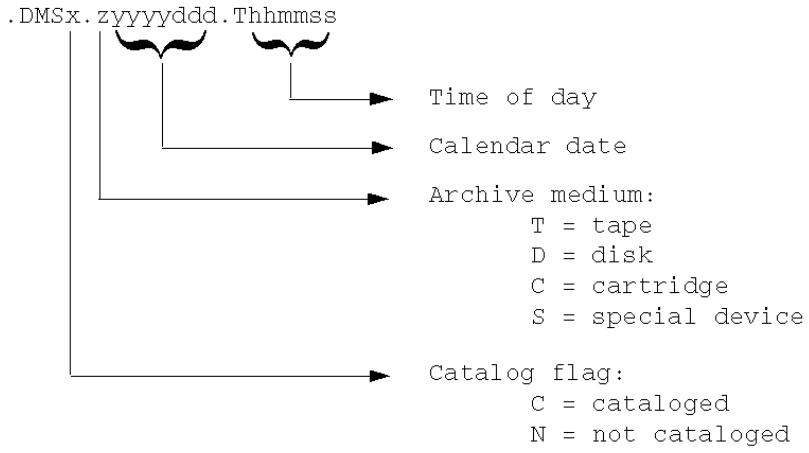

The previous calendar date is the current date in Julian format, and the time of day is in hours/minutes/seconds. Therefore, the archive data set name provides the date and time it was created, and is always unique.

For disk archive data sets, this generated name is always used and is always cataloged. For tape and cartridge archive data sets, the default is to use the generated name but not to catalog it. You can tell CA Disk to catalog the tape or cartridge archive data set name by specifying a value of C for sysparm ARCTNAME.

JCL Overrides and DSNAME Usage

As just explained, CA Disk defaults to using dynamic allocation for the archive data sets. For tape and 3480 devices, however, you can override dynamic allocation by supplying the needed information on the appropriate DD statements in the JCL. CA Disk creates the data set name from the DSNAME you provide in the JCL unless you omitted it completely or specified a temporary name (DSN=&&anyname). In both of these cases, CA Disk uses the appropriate sysparm (ARC0NAME, ARCCNAME, MERPnNAME, MERCnNAME, or MERTnNAM) to create the data set name.

Assigning an Expiration Date to the Archive Data Set

By default, CA Disk assigns the never expire date of 99365 to each of the archive volumes. This guarantees that an archive volume is never expired before all of the data sets that it contains. Therefore, your only concern is the retention period being assigned to each data set being archived. This is recommended if you are using the EDM interface of your Tape Management System.

When CA Disk determines that all data sets on an archive volume have expired, it automatically expires the volume as well. You can, however, specify a different expiration date through dynamic allocation with sysparm DYNEXPDT, or in JCL through the LABEL= parameter. Append either Eyyddd or Rddddd, where E indicates a Julian date follows and R indicates a 5-digit retention period follows.

For more information, see Year 2000 Considerations in the *User Guide* and Customizing the CA Disk Tape Management Support in the *Installation Guide*.

Archiving to Disk: Requirements and Recommendations

Traditionally, backup and archive copies of data sets have been directed to tape. CA Disk also provides the option of writing backup and archive data sets to disk. Using disk devices rather than tape or cartridge units for archive data sets eliminates the need for an operator to mount a tape when the backup or archive task is performed or when restoring data. This is particularly attractive for users implementing the auto-restore capability of CA Disk.

Note: When archiving to disk, you should consider activating software compression to save on DASD space. To activate software compression, see the sysparm description of DCDATASPn in the chapter "Sysparms."

The following summary outlines the steps necessary to implement archival to disk. Where necessary, an expanded explanation is provided for some of the summary items.

1. Specify archival to disk by one of the following methods:
   - If your ARCHIVE and DMS JCL procedures do not contain //ARCHIVE0 or //ARCHIVEC DD statements, specify sysparm ARC0TYPE with a value of DISK for the primary, and/or ARCCTYPE with a value of DISK for the copy, to allocate them dynamically.
   - Supply JCL statements for the previous ddnames that causes allocation to any disk device. CA Disk deallocates and then allocates the device as needed to perform the archive. For example, the following DD statement activates archival to disk for the primary copy:

     ```
     //ARCHIVE0 DD UNIT=3380,VOL=SER=anyvol,DISP=SHR
     ```

2. Specify the base name for the disk archive data sets using the following sysparms. CA Disk appends a unique name to the end of the base name, based on the current date and time.

   ```
   ARC0NAME—for the name of the primary
   ARCCNAME—for the name of the copy
   ```

3. Map the names you gave previously to proper disk pools (members in your PARMLIB data set) from which target volumes are selected. This mapping is specified with entries in member POOLDEFS of the PARMLIB data set. The general form of the entries is:

   ```
   'VOLU3380 DMS/'
   ```

4. Place target volumes in your defined pools by creating the PARMLIB member by the same name as your pool, in this case "VOLU3380", and inserting statements such as:

   ```
   VOL=(vol001,vol002,vol/)
   ```

5. Specify an expiration date using sysparm DYNEXPDT, or by JCL. CA Disk assigns the default expiration date of 99365 if none is specified.

   For more information, see the section Year 2000 Considerations in the *User Guide*.

6. Determine if the default values blocksize is appropriate. If not, override the default with the sysparm ARCDSKBZ.

7. Determine if the default manner of calculating primary and secondary space allocation is appropriate. If not, override the default by specifying sysparms:

   ```
   SPACEPRImmm—megabytes to allocate for primary
   SPACESECmmm—megabytes to allocate for secondary
   ```

8. Determine if default values for index maintenance, merge and rebuild functions are appropriate, and if not, specify sysparm to override the defaults.

9. Exclude Disk ARCHVOLS from CA Allocate EOV processing or any other product designed to add a new DASD volume to an allocation to avoid x37 type abends.

Specifying Archival to Disk

CA Disk is distributed with default sysparms that cause archive and backup copies of data sets to be written to tape, creating both a primary and a duplicate copy tape concurrently. You can tell CA Disk to write archive and backup copies of archive data sets to disk instead. You can write both the primary and duplicate copies to disk, or assign them to different media. You can also *dummy* out the duplicate copy. Indicate your choice of archive media either through JCL or CA Disk sysparms ARC0TYPE and ARCCTYPE.

Step 2—Naming the Disk Archive Data Set

For further information about this step, see Naming the Archives.

Step 3—Map Data Set Names to Diskpools

Candidate volumes for archiving to disk must be provided using diskpools, which reside as members in the PARMLIB data set. The name of the diskpool is simply the member name. CA Disk determines which diskpool (member) to use in the same manner in which tapepool names are determined. That is, the POOLDEFS member of PARMLIB is used to associate (map) a data set name to a poolname. For example:

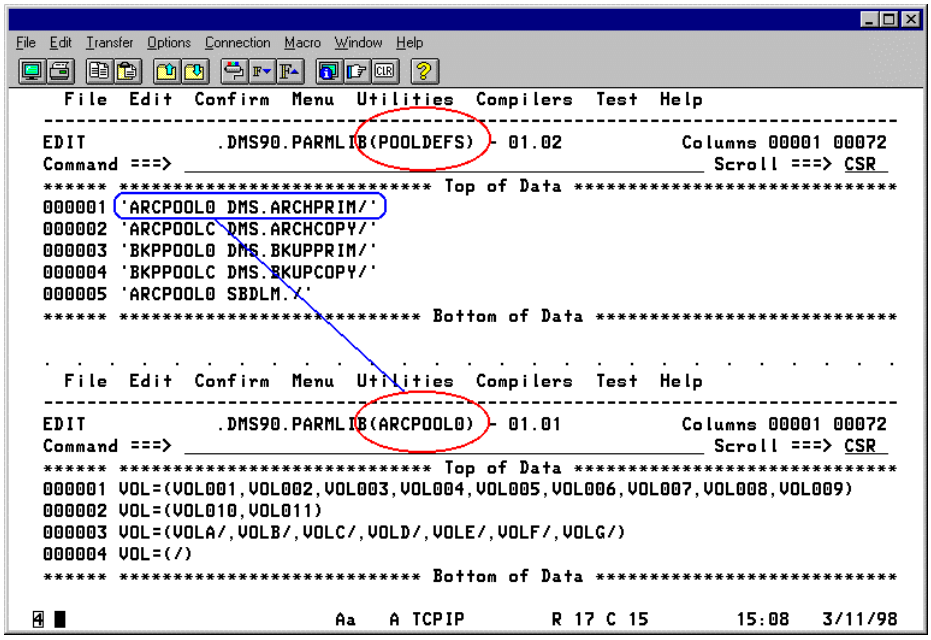

Step 4—Place Target Volumes in Defined Pools

Next, designate your candidate volumes in diskpool=ARCPOOL0 by creating a member in the PARMLIB data set with that name and inserting entries. For an example, examine the lower portion of the previous figure.

Either explicit volumes or volume patterns can be used. However, the total number of volumes in the pool cannot exceed 300. Do not try to continue the volume list on multiple lines (that is, continuation lines are not supported). Instead, enter multiple VOL= statements.

If both the primary and the duplicate copy are being directed to disk, each pool must contain volumes of the same device type(s) as the other pool; that is, if pool 1 has both 3380s and 3350s, pool 2 must also have some 3380s and 3350s. This is required because a volume from pool 1 is selected first, and then pool 2 is searched for a volume of the same device type.

By default, CA Disk automatically separates the copy from the primary; that is, if a volume appears in both pools, CA Disk does not select the same volume for the copy as is used for the primary. Obviously, this is to prevent the loss of both copies if the disk pack is destroyed, whether it is due to oxidation, head crashes, or other reasons. To allow both copies to go to the same volume, specify sysparm ARCSEPC0 with a value of N.

Step 5—Specifying an Expiration Date

For further information about this step, see Assigning an Expiration Date to the Archive.

Step 6—Specifying Blocksize

When archiving to disk, CA Disk automatically adjusts the optimal blocksize value for the device type of the volume; that is, half track blocking on 3375s and 3380s, and full track blocking on other devices. You can override the defaults by specifying sysparm ARCDSKBZ. The blocksize selected for a disk device is also used for any tape or 3480 copy that is being created concurrently.

If you have more than one device type in your archive pool, use sysparm ARCDSKBZ to specify a blocksize appropriate for all devices; that is, one that provides optimal performance and space usage on all the devices.

Step 7—Specifying Space Allocation

CA Disk determines the amount of space to allocate (the amount of space it needs on a diskpool volume) in the following manner. It assumes that quite frequently the primary copy is on disk and the duplicate is on tape or a 3480 cartridge. As stated earlier, the target volume leads CA Disk to the device type, which in turn dictates the blocksize to be used. This, coupled with the density and length of tape being used, allows an easy and very accurate calculation of the tape capacity. (Because the ends of 3480 tapes cannot be clipped off, the most common variable in the calculation has been removed.)

The optimal disk allocation is the exact same capacity as its backup medium, such that when one is filled and both must be closed, nothing is wasted on the other medium. To do this on an exact equivalence, however, would require very large amounts of free space to be available. Therefore, CA Disk selects 1/16th of that value as default primary and secondary space allocations. In many cases, an archive run does not need a full tape, and one or two extents are more than enough space. However, when a large archive run is made and the diskpool volumes contain plenty of free space, CA Disk obtains as many as possible of the 16 extents before closing the archive data set. This maintains high tape usage as well. Any excess (unused) disk space is, of course, released immediately when the data set is closed.

CA Disk also considers the size of the data sets being archived when determining how much space to allocate. If the first data set to be archived is larger than the default primary space, the primary space is reset to the size of the data set. If free space equal to this new primary value is not available but the default value is, CA Disk attempts to archive the input data set to a multivolume output data set. This is done only as a last resort and you should not use a Storage Class that is defined with guaranteed space because unused space is only released in the active volume at CLOSE time (there is a possibility that all preallocated volumes would not be used because the data is written in compressed format).

Once the target archive data set is allocated and the first data set to be archived is copied to it, CA Disk continues to copy additional data sets (the second through nth) as long as there is sufficient space to hold them. To determine whether sufficient space is available, CA Disk examines the unused space in the current extent, and checks to see if additional extents can be obtained. If there is not sufficient space, the current disk archive data set is closed and a new one is allocated (like swapping to a new tape).

If your analysis or practical experience indicates that the default method of calculating space allocations should be changed, specify sysparms SPACEPRI and SPACESEC with a three-digit number representing the number of megabytes of disk space to allocate for each.

When archiving to disk, you normally want the archive data sets to be large, as restoring from a large data set is as fast, or faster, than restoring from several small data sets. CA Disk provides options for balancing the need to conserve storage space with the need to keep overhead to a minimum. If you set the primary space value (sysparm SPACEPRI) to the size you want to use for staging, then set the secondary space value (sysparm SPACESEC) to 000, CA Disk creates smaller archive data sets, but more of them.

Step 8—Exclude Disk ARCHVOLS from CA Allocate EOV Processing

Disk ARCHVOLS should be excluded from CA Allocate EOV processing or any other product designed to add a new dasd volume to an allocation to avoid x37 type abends.

Step 9—IXMAINT, MERGE, and REBUILD

To accommodate archival to disk, the following options are available within the following functions:

IXMAINT

Parameter SCRATCH=YES/NO is available on the DSNDELETE, VOLDELETE, and PURGE commands. DSNDELETE and VOLDELETE default to SCRATCH=YES, which means to scratch and uncatalog the disk archive data set when the archvols record is being deleted. PURGE defaults to SCRATCH=NO, because it is frequently used just prior to running REBUILD. If a disk archive data set is scratched, there is nothing from which to rebuild. If you do not intend to follow PURGE with a REBUILD command, specify SCRATCH=YES. Otherwise, you are creating *disconnected* archive data sets on disk that are never used or deleted by any other means.

MERGE

The following parameters have been added.

TYPE=(DISK,TAPE,3480)—specify any of the three types listed to restrict the merge to archvols of those types.

CREDT=date—only those archvols created on or before the specified date are merged forward.

DAYSOLD=dddd—an "alternative form" of the CREDT parameter that limits the merge to those archvols that were created dddd or more days ago.

REBUILD

Rebuild accepts the archive data set to be read as input from either tape or disk. The supplied JCL procedure requires that you provide the correct //ARCHIVES DD statement pointing at the tape or disk data set to be rebuilt. No input cards are required. If they are provided, only the expiration date parameter is used.

Creating and Using Dual Primary Volumes

Differences of Primary/Copy Volumes vs. Dual Primary Volumes

The typical CA Disk backup/archive process creates duplicate outputs on tape or disk referred to as the primary and copy archive volumes (or, ARCHVOLS).

An optional process that can be utilized during backup/archive is to create Dual Primary Volumes. Using this option, the backup/archive process creates PRIMARY1 and PRIMARY2 (P1 and P2) volumes on tape or disk.

Here are some distinctions between these two processes:

Primary/Copy	Dual Primary Volumes
The primary and copy volumes contain exactly the same data.	The two primary volumes contain exactly the same data.
The primary/copy volume relationship does not change for the life of the volumes.	The P1 and P2 volumes are not connected or related in any way.
For each data set processed, one DSNINDEX record is created.	For each data set processed, a P1 DSNINDEX record and a P2 DSNINDEX record is created.
A primary ARCHVOLS and a copy ARCHVOLS record are created for the output volumes.	A P1 ARCHVOLS record is created for the P1 output volume and a P2 ARCHVOLS record is created for the P2 output volume.

Primary/Copy	Dual Primary Volumes
The DSNINDEX record contains the volser (or key) of the primary ARCHVOLS record.	The P1 DSNINDEX record contains the volser (or key) of the P1 ARCHVOLS record and, the P2 DSNINDEX record contains the volser (or key) of the P2 ARCHVOLS record.
The primary ARCHVOLS record contains the key of the copy ARCHVOLS record.	The P1 ARCHVOLS record has no field pointing to the P2 ARCHVOLS record, nor does the P2 ARCHVOLS record point to the P1 ARCHVOLS record.
The primary and copy archvol cannot be managed separately.	The P1 DSNINDEX and P1 ARCHVOLS records can be managed independently of the corresponding P2 records.
The copy ARCHVOL is kept until the primary ARCHVOL expires, at which time the copy ARCHVOL is also expired.	The P2 DSNINDEX and P2 ARCHVOLS records can expire before, after or at the same time as the P1 records.

In an execution of CA Disk, the option to create Dual Primary Volumes can be utilized or Primary/Copy volumes can be created. These are mutually exclusive processes in the job step being executed. The Copy Utility can be used to create copies of the Dual Primary Volumes.

- All of the P1 and P2 DSNINDEX and ARCHVOLS records are maintained in the FILES, along with the entries for the primary/copy ARCHVOLS and DSNINDEX records.
- The DSNINDEX records and the primary ARCHVOLS records created in the primary/copy process are considered as P1 records (for example, the *P2* indicator flag is *off*).

There is no distinction - and no method to distinguish - between primary DSNINDEX records and ARCHVOLS created in the primary/copy process and the *P1* DSNINDEX records and ARCHVOLS created in the Dual Primary Volume process.

Although the Dual Primary Volume feature provides very flexible options in managing dual output images and volumes, the user must be aware that there can be a significant increase in the number of DSNINDEX records maintained in the FILES data set.

Before using the Dual Primary Volume feature, the following factors must be considered when using an FDS:

- Ensure that the physical size (number of cylinders allocated) of the FDS is adequate to contain the expected increase of DSNINDEX records.
- Ensure that the record (or, indexing) capacity for the DSNINDEX subfile is adequate (specified in member FILEDEFN in the CA Disk PARMLIB).
- The runtime of IXMAINT can be longer because of the additional records created by this feature.

Dual Primary Volume Feature Overview

The following are the functions where the P1 and/or P2 records can be selected or viewed:

RESTORE/RECOVER

Selects P1 and/or P2 DSNINDEX records.

IXMAINT

Selects P1 and/or P2 records for maintenance processing.

MERGE

Selects P1 or P2 volumes.

LISTD

Displays both P1 and P2 DSNINDEX records.

LISTD and LISTV

Shows the setting of the P2 flag field.

REBUILD

Restores the P2 indicator for the DSNINDEX and ARCHVOLS records when the P2 information is available.

Merging of two or more FDS

Includes the P2 indicator for the sort sequence.

Except for the functions previously mentioned , all other CA Disk functions operate without consideration for the P2 record indicator.

Dual Primary Volume Details by Function

Archive/Backup

Dual Primary Volumes are created when specifying the following statement in the DSCL:

```
SET  PRIMARY2
```

When SET PRIMARY2 is specified, a P1 and P2 DSNINDEX record is created for each data set processed, and a P1 and P2 ARCHVOLS record is created for each respective output volume.

The Primary2 Volume data set name and unit information are taken from the ARCCNAME and ARCCTYPE sysparms.

Different expiration dates can be applied to the P1 and P2 DSNINDEX and ARCHVOLS records. The P1 DSNINDEX and ARCHVOLS records can have expiration dates that are greater than, less than or equal to the P2 records expiration dates.

- Parameters RETPD and EXPDT apply to the P1 DSNINDEX records, except as noted in the following.
- Parameters RETPD2 and EXPDT2 apply to the P2 DSNINDEX records as they are created.
- If P2 DSNINDEX records are being created and RETPD2 or EXPDT2 is not specified, the expiration date is taken from the RETPD or EXPDT parm, if specified.
- If P2 DSNINDEX records are being created and neither RETPD/EXPDT or RETPD2/EXPDT2 parameter is specified, the retention period for both the P1 and P2 DSNINDEX records is taken from sysparm RETRETPD.

For SMS data sets, if the Management Class is present and not overridden (for example, sysparm SMSMCBYP is set to N or B), the RETPD/EXPDT and RETPD2/EXPDT2 parameters are ignored and the expiration date of the P1 and P2 DSNINDEX records is set to 99365. Those P1 and P2 DSNINDEX records continue to be managed using DFSMS and cannot be managed separately.

The P1 and P2 ARCHVOLS records can be assigned different expiration dates:

- Sysparm DYNEXPDT applies to the P1 ARCHVOLS records, except as noted in the following.
- Sysparm DYNEXPD2 applies to the P2 ARCHVOLS records as they are created.
- If P2 ARCHVOLS records are being created and DYNEXPD2 is not overridden in the SYSPARMS member or in the job, the expiration date for both the P1 and P2 ARCHVOLS records is taken from sysparm DYNEXPDT.

Here is an example of using the new parameters and the results:

```
SET PRIMARY2
SCAN REALVOLS
BACKUP RETPD=30,RETPD2=15
```

Sysparms:

```
DYNEXPDTE99365
```

The P1 DSNINDEX records expire in 30 days.
The P1 ARCHVOLS record does not expire until all of its DSNINDEX records have been deleted.
The P2 DSNINDEX records expire in 15 days.
The P2 ARCHVOLS record does not expire until all of its DSNINDEX records have been deleted.

Here are more examples of creating Dual Primary Volumes and retention periods:

Command (Today=2000.325, RETRETPD=30)	P1 DSN EXPDT	P2 DSN EXPDT
BACKUP	2000.355	2000.355
BACKUP RETPD=5	2000.330	2000.330
BACKUP RETPD2=5	2000.355	2000.330
BACKUP RETPD=5,RETPD2=10	2000.330	2000.335
BACKUP EXPDT=2000.360	2000.360	2000.360
BACKUP EXPDT2=2000.364	2000.355	2000.364
BACKUP EXPDT=2000.360,EXPDT2=2000.364	2000.360	2000.364

Restore/FMS Recover/Auto Restore

Sysparm RESPRISQ indicates which DSNINDEX records are to be searched, and in which order, for use in the Restore/Recover functions. The values that can be specified are:

- 1—Search only the P1 DSNINDEX records
- 2—Search only the P2 DSNINDEX records.
- 12—Search for P1 DSNINDEX records first but allow use of the P2 record if no P1 is found. (default).
- 21—Search for P2 DSNINDEX records first but allow use of the P1 record if no P2 is found.

These values can affect the version number of the DSNINDEX records selected. See the description of sysparm RESPRISQ in the Systems Guide for more complete information.

IXMAINT

A parameter PRIMARY2 for the DSNDELETE command indicates whether the P2 records (both DSNINDEX and ARCHVOLS) are to be processed. The values that can be specified are:

- YES—The P2 records are processed as individual records (default)
- NO—The P2 records are not processed

ONLY—Only the P2 records are processed

By default, IXMAINT treats P2 DSNINDEX records as individual records (for example, data sets can be expired from P2 volumes without affecting any P1 data sets or volumes) except in COPY processing. In COPY processing, if IXMAINT deletes the P1 record, the P2 record is also deleted. COPY processes are: the COPIES parameter on the DSNDELETE command; SMS backup copies and SMS GDG copies.

Using the following example:

```
DSNDELETE COPIES=5
```

If there are six pairs of P1 and P2 DSNINDEX records, the oldest pair is deleted. If there are only four pairs, none are deleted.

MERGE

A simple parameter **PRIMARY2** for the Merge command indicates that only P2 volumes are selected. The default is that only P1 volumes are selected.

LISTD (in ISPF or Batch)

A simple parameter **both** for the LISTD command allows the display of the most current P1 DSNINDEX record and its corresponding P2 DSNINDEX record, if both are available. The default is that the most current record is displayed, which can be a P1 record or it can be a P2 record if there is no corresponding P1 record. For example, if using the commands:

LISTD DSN=A.B.C./,BOTH

The most current P1 record and corresponding P2 record are displayed.

LISTD DSN=A.B.C./

The most current record is displayed, whether P1 or P2.

The P2 flag field is added to the reports and is an option in the FIELDS parameter on the LISTD command.

LISTV (in ISPF or Batch)

The P2 flag field is added to the reports.

REBUILD

A simple parameter PRIMARY2 indicates that the DSNINDEX and ARCHVOLS records are rebuilt as P2 records. The P2 flag is set and the appropriate expiration dates are used.

VBACKUP

Dual Primary Volumes cannot be created in the VBACKUP function at this time.

XCOPY

XCOPY cannot process PRIMARY2 volumes at this time.

User-Specified Condition Codes

CA Disk assigns a step completion condition code to every step. The code accompanies a message and indicates the circumstances under which the step finished. In most cases, this CA Disk-generated condition code is the best value for your installation. However, you can override these generated condition codes if your installation has special requirements. This override allows you to select the code that accompanies a message upon the completion of a step. These user-specified condition codes are optional.

Note: A similar option is available for those instances where NO CATALOG ENTRIES WERE FOUND exists when DSCL is searching for candidate data sets. For more information, see the sysparm CCDRESETn in the chapter "Sysparms."

To specify user condition codes, you must create a member in PARMLIB called CCSET. This member is not provided with the system, although a member called SAMPCSET is provided as a sample, and can be copied to create your initial CCSET. User condition codes are then specified by an entry in this member.

When using this option, several technical points must be kept in mind with respect to user condition code processing.

1. Once an N (no-override) is encountered during processing, CA Disk determines sets and assigns condition codes for all the message number entries that follow without regard for the overrides specified in member CCSET.
2. Whenever a message is issued during step processing, the highest condition code is set without regard to its origin.
3. Before CA Disk returns control to the operating system, a check of the CA Disk-generated return code is made. If the CA Disk return code is odd and the user condition code option is in effect, then the user override condition code is incremented by one. For example, if the user specified an override condition code of 4 and the CA Disk generated condition code is odd, then the user override condition code of 4 is incremented to 5. This is needed to produce sorted report listings.
4. No duplicate message number entries are allowed in member CCSET. If duplicates are detected at table initialization, the program abends with a user 200.

   Each CCSET entry is composed of a message number, a state flag (O=override or N=no-override), and one of five possible condition codes. The message number is the first four positions followed by a blank. The next two positions consist of the state flag and a condition code respectively.

   To add, change, or delete entries in the CCSET member, use your online editor or other conventional means for updating a PDS member in the following manner:

   1. Create a member named CCSET in the PARMLIB data set by copying and renaming the sample member SAMPCSET.

2. Specify the message display control entry in CCSET. This must be the FIRST entry in CCSET and is required. This entry controls the display of the message, which notifies users that the user condition code option is in effect. It is recommended that you try the default first which displays the message in BATCH and TSO. The format for $FLGXyz is as follows:

Field	Meaning		
$FLG	$FLG		
X	BLANK		
y	BATCH FLAG	B BLANK	display message under BATCH do not display message
z	TSO FLAG	T BLANK	display message under TSO do not display message

In order to suppress the display of the notify message under BATCH or TSO processing, simply blank out the respective flag. If you want to suppress the display of the notify message under BATCH and TSO processing, blank out both flags.

5. Specify the message number, state flag and condition code enclosed by apostrophes in the CCSET member. An example can be:

```
'3179 0E'
```

The format for message number, state code, and condition code is shown in the following table:

Field	Meaning
MMMM	message number
B	blank
s	state flag
c	condition code

The valid entries for MMMMBsc are shown in the following table:

Field	Possible Values
MMMM	any valid four-digit CA Disk message number
B	must be blank

Field	Possible Values
s	**O**—override **N**—no-override
c	**G**—0, good, successful completion I4, informational message **R**—8, resource error **E**—12, error occurred during CA Disk processing **C**—16, command error, DSCL command error detected

Suggested System Parameters

System Parameters for Reports

The Data Set Utilization report can be produced in three different sequences. By default, it is produced in sequence by date-of-last-use. The following DSUTIL report sorting options are available:

Sysparm	Description
DSUTILSQ	In data set name sequence
DSUTILCF	By date-of-last-use within index

Both the explicit and the implicit archiving functions can produce reports in one to three sequences simultaneously. By default, they are produced in both data set name and data set name within volume sequences. See the sysparm ARCHSORTudv in the chapter "Sysparms" to select a different option.

The sequential migration to tape function has this same ability. See the sysparm MIGRSORTndv in the chapter "Sysparms" to adjust its defaults.

The SMF report requires its input from a sequential data set. If your SMF records reside in a VSAM cluster, add an IBM IDCAMS utility step to REPRO the records into a sequential data set to be passed to CA Disk.

System Parameters for Messages

The following sysparms message suppression options affect the number of informational messages that are printed for their respective functions and should be kept in mind when initially testing and setting up CA Disk runs:

Sysparm	Description
DSCLMSGS	DSCL Data Set Selection Messages
RLSEDIAG	Idle Space Release messages
MIGBYPAS	Sequential Migration to Tape messages
VCBYPASS	Move/Copy messages

System Parameter for Lines Per Page

The number of lines per page on CA Disk-generated output can easily be changed from its default of 58. For details, see the sysparm description for RPTLINES in the chapter "Sysparms."

User Exit for All Sysout Print Lines

For more extensive monitoring or modifications of printed output, such as special routing for certain messages see the section SYSOUTEX-SYSOUT Exit in the chapter "User Exits."

System Parameter for FDS Logging

The following three sysparms provide support for the FDS logging capability:

Sysparm	Description
FILOGNAM	For specifying the name of the log data sets and to activate the feature.
FILEUNIT	For specifying the unit type allocated for the log data sets.
FILESPAC	For specifying the number of blocks allocated to the log data sets.

System Parameters for Archive/Backup

The following command ARCHIVE/BACKUP sysparms are commonly specified and should be reviewed for applicability in your installation:

Sysparm	Description
ARCBLKSI	for processing data sets that have BLKSIZE=0
ARCDSORG	for processing data sets with unknown DSORGs
ARCEMPTY	for processing empty data sets
ARCMODEL	for processing model DSCBs
ARCRETPD	change retention period for explicit archives
RETRETPD	change retention period for implicit archives

System Parameters for Restore/Recover

By default, CA Disk issues a catalog locate for each data set being explicitly restored, unless a specific target volume is given. For volume recovery functions, these catalog locates are bypassed. Another locate is done in preparing the restore report. Both restore and recovery tasks normally skips any preallocated version of a data set. Common RESTORE/RECOVER sysparms to review in order to change these defaults are the following:

Sysparm	Description
RESCHCAT	for volume selection for non-VSAM during Restore
RECCHCAT	for volume selection for non-VSAM during Recover
PREALLOC	controls whether or not to overwrite preallocated data sets

Implementing Support for Masstor M860

Masstor Systems Corporation's M860 device is sysgen'd as either a 3420 or a 3480 device. The M860, however, treats the cartridge as either the 3420 or 3480 tape device as it was sysgen'd. Just as native 3420 and 3480 devices require different unit names to prevent requests for conventional tapes to be mounted on a 3480 drive, the M860 also requires this separation.

Providing the unit name of the type of device you want to use is very straightforward when new *tapes* are being created. CA Disk needs to save this distinction, such that the restore, recover, or merge functions can dynamically request the proper type of device. If this is not done, CA Disk issues requests to mount M860 cartridges on either 3420 or 3480 tape drives!

Distinguishing the M860 Device

CA Disk supports the M860 as a distinct device only using dynamic allocation. If archives are written to the M860 by referencing it on a DD statement in your JCL, it appears that it is a 3420 or 3480 and is recorded as such in the ARCHVOLS record (which is subsequently used by restore, recover, and merge processing to allocate what is thought to be the proper type of device).

You inform CA Disk that an M860 device is being used and should be tracked in the ARCHVOLS as requiring special allocation by providing the proper value for sysparms ARC0TYPE, ARCCTYPE, MERPnTYP, MERCnTYP, and MERTnTYP. The standard values for these sysparms are TAPE, 3480, or DISK. For the M860, a value of DYN1 must be specified. This causes the *DYN1 special unit flag bit* to be set in the ARCHVOLS record. Subsequent CA Disk dynamic allocation processing examines this bit and, if on, uses the unit name provided by sysparm ARESUNI1 for auto-restore, and sysparm DYN1UNIT for all other functions. (DYN2 through DY20 options are also provided, but are explained later.)

The following table indicates which sysparms you use to provide the unit name for each device type value of sysparm ARC0TYPE. Sysparms beginning provides unit names for the auto-restore function with *ARES*, while all other functions take their unit names from sysparms beginning with DYN*n*/DY*nn*. The following table shows possible M860 values for ARCOTYPE.

Value	Sysparms for Providing Your Unit Names		
TAPE	DYNTUNITuuuu	ARESUNITuuuu	—for 3420s
3480	DYNCUNITvvvv	ARESUNICvvvv	—for 3480s
DYN1	DYN1UNITxxxx	ARESUNI1xxxx	—for special unit 1

Additional Implementation Needs for the M860

CA Disk tapepools should be created that consist solely of M860 cartridges, and all cartridges within each tapepool must physically be within the same M860 unit.

If you have multiple M860 units and a unique unit name or unit address can be used to reference each of them, you should use DYN1 to map allocations to tapepools defined in one unit, DYN2 for the second, DYN3 for the third, and so on. DYN2 through DY20 works in exactly the same manner as described for DYN1, and together provide the ability to identify and access up to three different groups of devices with special allocation requirements.

- **M860 Unit 1**—Use DYN1 when archiving, and use archive data set names that map to tapepools in this unit. Sysparm DYN1UNITuuuu provides its unit name/address.
- **M860 Unit 2**—Use DYN2 when archiving, and use archive data set names that map to tapepools in this unit. Sysparm DYN2UNITvvvv provides its unit name/address.
- **M860 Unit 3**—Use DYN3 when archiving, and use archive data set names that map to tapepools in this unit. Sysparm DYN3UNITwwww provides its unit name/address.

Implementing Support for StorageTek Redwood

The Redwood device supports cartridge capacities of 10GB, 25GB, or 50GB of uncompressed data. The cartridge used by this device has the same size and same characteristics as existing 3480/3490 cartridges, with two obvious differences:

1. The actual tape exits the cartridge from a different corner
2. To prevent these cartridges from being mounted into a true 3480/3490 device, a notch in the corner of each cartridge is flipped on

CA Disk acknowledges this device in quite the same way it did 3420 and 3480 devices, through an esoteric unitname. Once a unique esoteric is assigned to your Redwood devices, CA Disk begins storing this name in the ARCHVOL records. This allows CA Disk to dynamically request the proper device type during RESTORE, RECOVER, XCOPY, or MERGE processing.

For an overview description on the capacities the Redwood device has, see StorageTek Redwood Device in the *User Guide*.

Distinguishing the Redwood Device

CA Disk supports the Redwood as a distinct device only with dynamic allocation. If archives are written to the Redwood by referencing it on a DD statement in your JCL, it appears that it is a 3420 or 3480 and is recorded as such in the ARCHVOLS record (which is subsequently used by restore, recover, and merge processing to allocate what is thought to be the proper type of device).

You inform CA Disk that a Redwood device is being used and should be tracked in the ARCHVOLS as requiring special allocation by providing the proper value for sysparms ARC0TYPE, ARCCTYPE, MERPnTYP, MERCnTYP, and MERTnTYP. The standard values for these sysparms are TAPE, 3480, or DISK. For Redwood, a value of DYN1 must be specified. This causes the *DYN1 special unit flag bit* to be set in the ARCHVOLS record. Subsequent CA Disk dynamic allocation processing examines this bit and, if on, uses the unit name provided by sysparm ARESUNI1 for auto-restore, and sysparm DYN1UNIT for all other functions. (DYN2 through DY20 options are also provided, but are explained later.)

The following table indicates which sysparms you use to provide the unit name for each device type value of sysparm ARC0TYPE. Sysparms beginning provides unit names for the auto-restore function with *ARES*, while all other functions take their unit names from sysparms beginning with DYN*n*/DY*nn*. The following table shows possible REDWOOD values for ARCOTYPE.

Value	Sysparms for providing your unit names		
TAPE	DYNTUNITuuuu	ARESUNITuuuu	—for 3420s
3480	DYNCUNITvvvv	ARESUNICvvvv	—for 3480s
DYN1	DYN1UNITxxxx	ARESUNI1xxxx	—for special unit 1

Additional Implementation Needs for Redwood

CA Disk tapepools should be created that consist solely of Redwood cartridges, and all cartridges within each tapepool must physically be within the same Redwood unit.

If you have multiple Redwood units and a unique unit name or unit address can be used to reference each of them, you should use DYN1 to map allocations to tapepools defined in one unit, DYN2 for the second, DYN3 for the third, and so on. DYN2 through DY20 work in exactly the same manner as described for DYN1, and together provide the ability to identify and access up to three different groups of devices with special allocation requirements.

- **Redwood Unit 1**—Use DYN1 when archiving, and use archive data set names that map to tapepools in this unit. Sysparm DYN1UNITuuuu provides its unit name/address.
- **Redwood Unit 2**—Use DYN2 when archiving, and use archive data set names that map to tapepools in this unit. Sysparm DYN2UNITvvvv provides its unit name/address.
- **Redwood Unit 3**—Use DYN3 when archiving, and use archive data set names that map to tapepools in this unit. Sysparm DYN3UNITwwww provides its unit name/address.

Implementing Support for IBM's Magstar

Magstar supports cartridge capacities of 10GB and 20GB of uncompressed data. The cartridge used by this device is unique and cannot be used by any other IBM tape device.

CA Disk acknowledges this device using the use of esoteric unit names. Once a unique esoteric is assigned to these devices, CA Disk begins storing this esoteric name in each ARCHVOLS record created on the device. This procedure allows CA Disk to dynamically request the proper input device during XCOPY, RESTORE, RECOVER, or XCOPY processing.

For an overview description on the capacities the Magstar device has, see IBM's Magstar 3590 Tape Device in the *User Guide*.

Distinguishing the Magstar Device

CA Disk supports the Magstar device as a distinct device that can only be accessed through dynamic allocation. If archives are written to a Magstar device by referencing the device on a DD statement in JCL, the Magstar device appears as a 3420 or 3480 and is recorded as such in the ARCHVOLS record (which is subsequently used by MERGE, RESTORE, RECOVER, and XCOPY processing as input to the proper device type).

To inform CA Disk that a Magstar device is being used is accomplished by assigning the proper value to the following system parameters:

- ARC0TYPE
- ARCCTYPE
- MERPnTYP
- MERCnTYP
- MERTnTYP

The standard values for these sysparms are TAPE, 3480, CART, or DISK. To use a Magstar device, a value of DYN1 must be specified. This causes the *DYN1 special unit flag bit* to be set in the ARCHVOLS record. Subsequent CA Disk dynamic allocation processing examines this bit and, if on, uses the unit name provided by one of two system parameters:

- ARESUNI1—for auto-restore
- DYN1UNIT—for all other functions

Additionally, sysparm CARTCALC must be specified with a value of "Y" in order for the full length of the tape to be used. As an example, if you were to GEN into your MVS system the esoteric "3590" to represent Magstar, the following sysparm settings would direct BACKUP/ARCHIVE processing to your Magstar devices:

```
//SYSPARMS DD *              <= POINT TO MAGSTAR
DYN1UNIT3590                 <= USE DYN1
ARC0TYPEDYN1                 <= USE DYN1
ARCCTYPEDYN1                 <= PRIM TAPE NAME
ARC0NAMEyour.primary.name    <= COPY TAPE NAME
ARCCNAMEyour.copy.name       <= USE WHOLE TAPE
CARTCALCY
/*
```

Additional Implementation Needs and Notes for Magstar

Magstar tape cartridges are unique in that they cannot be used in any other IBM or compatible tape system. For this reason, though not required, CA Disk tapepools should be created that consist solely of Magstar cartridges.

Performance Tips

Report Processing

See the sysparms for the reporting function listed earlier that control the use of the LOCATE macro. This macro causes a high level of system overhead. Elapsed time values can be reduced significantly if the locate function is suppressed for runs that process a high number of data sets.

PDS Compress

Sysparm IOTRACKS also affects PDS directory processing. Tests have also shown that for directories in the range of one track or less (which is 46 directory blocks or less on a 3380), setting IOTRACKS to 1 can improve performance. If tests verify this in your environment, supply this value as a permanent sysparm override to the PDS Compress function.

Previous recommendations to exclude SMP/E data sets still being kept as PDSs (rather than VSAM data sets) are no longer applicable. The revisions allow all partitioned data sets to be processed equally well.

FDS

Because the FDS contains the archive indexes, it is an active and critical data set when doing implicit archiving, backup, or any index maintenance function. ENQ and RESERVE macros are also issued against the data set and its containing volume. Improperly selecting the containing volume can result in performance degradation of CA Disk or other jobs trying to access the volume.

For important information, see the sysparm description for RSUPPRESn in the chapter "Sysparms."

I/O Buffering and Memory Requirements

A sysparm that can have a dramatic effect on CA Disk performance and the amount of memory that is required is IOTRACKS. IOTRACKS controls the maximum number of tracks read or written at one time by the CA Disk EXCP access method, which is used for all but exception case handling of non-VSAM data sets. See the sysparm description for IOTRACKSxx in the chapter "Sysparms."

Ideally, optimum performance is achieved when IOTRACKS allows a cylinder to be processed by a single I/O request; for example, 15 tracks for 3380 devices. The default value for IOTRACKS accomplishes this. Test different values of this sysparm to find the best to suit your environment. As with most performance issues, there can be several factors to consider, such as the memory constraints, I/O contention, and page and swap activity.

First consider the size of these buffers themselves. A 15-track buffer on a 3380 is (15)*(47,476) or 712,140 bytes. CA Disk double buffers, which means 1,424,280 bytes (or roughly 1.4 megabytes). For the Move/Copy function, which is reading and writing to disk at the same time, this means 2.8 megabytes merely for basic I/O data buffers. Memory is required for many other items as well, such as space for the loaded programs themselves and buffers for reading VTOCs, PDS directories, and control information from the CA Disk parameter library. Many of these memory tables reside above the 16 MB line. A five- to six-megabyte region is usually sufficient for all CA Disk jobs, including Move/Copy. If real memory is very limited, however, this region requirement can be too high, and must be lowered by reducing the value of sysparm IOTRACKS.

Buffers are acquired based upon the values specified, but a single I/O request never exceeds a cylinder in size or crosses a cylinder boundary. If the number of tracks in a data set (or technically the number of tracks in each extent) is less than a cylinder in size, part of the buffer always remains unused (wasted). It follows that if the average data set size in tracks is less than the value specified for the appropriate sysparm, you can set the sysparm to a lower value to reduce memory constraints without noticeably increasing the number of I/Os (which degrades performance).

If the amount of available real memory is limited, having this sysparm set to a high value can cause jobs to be swapped frequently and/or paging rates to be high, which degrades performance instead of improves it.

If the *to* and *from* volumes are on the same channel when using the Move/Copy function to relocate data sets, the job can perform better by setting IOTRACKS to a much smaller value, perhaps even 1. This is due to the inherent channel contention in this particular case, combined with the paging and swapping problem previously described. Again, testing within your environment demonstrates whether this case warrants special treatment.

Another sysparm to see when considering memory requirements is IOMAXRECnnnnnn. This sysparm sets the maximum allowable record size for data sets to be processed by CA Disk. It is used to acquire buffer space to contain the records during backup and data compression/decompression. The higher this value is set, the greater the memory requirements. Set this sysparm to the lowest value possible that still allows CA Disk to process. The default value is 65,000.

Customization Options

The CA Disk product provides the DASD administrator with a great deal of flexibility. In the pages that follow, several of the more commonly used customization options are discussed.

Whenever a CA Disk customization option directly changes the operating system, it is highly advisable to apply the change as a usermod using SMP/E. At the same time, any change you make to CA Disk itself should be considered as a usermod. You should also make these changes using SMP/E.

Each customization step discussed includes the proper JCL and instructions to install with SMP/E. After you have installed CA Disk with SMP/E, make sure that you have run the ACCEPT step prior to installing any customizing options. This allows for the easy removal of these options, if problems should occur.

Observe the following:

- Several SMP/E format usermods have sample JCL, which should NOT be modified. This sample JCL is always proceeded by a ++JCLIN control card. This sample JCL is used by SMP/E to indirectly determine the libraries to be updated. The sample JCL does not directly identify these libraries. Modification of this sample JCL can cause unpredictable results during installation of the usermod.
- Always verify the target and distribution zones into which you are installing each usermod. Usermods which modify your operating system--for example, the CA Disk SVC, VSAM date stamp and so on--should be installed into the same zones as your operating system. Usermods that modify CA Disk, such as auto restore zaps, and so on, should be installed into the same zones as CA Disk. Improper selection of target and distribution zones during the installation of user modifications can cause unpredictable results during installation.

As part of your planning for installing one or more customizing options, verify that the JCL provided conforms to your installation's JCL procedures. The Assembler JCL provided uses step names of C and L. Your Assembler procedure can use ASM and LKED. SMP/E JCL uses the DDNAME SMPCNTL.

The following customizing options are discussed in this section:

- The CA Disk SVC and applicable zap
- The CA Disk VSAM date stamp module
- The auto-restore function
- The CA Disk tape management facility
- Setting ISPF customization options
- Defining ISPF user options
- Defining TSO customization options

The CA Disk SVC

CA Disk SVC updates a data set VTOC entry, recording the last used date, for example. The SVC has the following advantage over the IBM SU 60 update code: exemption entries are provided such that the product management tasks themselves do not cause data sets to appear used.

Zaps to one or more IBM modules must be applied to call the CA Disk SVC to do the updating.

The CA Disk SVC maintains the SU 60 change bit and last used date fields. SVC also maintains extra CA Disk defined fields and provides the facilities to overcome the SU 60 deficiencies.

The two fields that IBM SU 60 maintains appear to provide the information that you need to help manage disk storage. CA Disk functions can examine the two fields and can execute without the CA Disk SVC being installed. The SVC is required for your implementation for the following reasons:

- The CA Disk storage management tasks open user data sets to perform the management functions, such as, backup, migration, and compression. IBM SU 60 open hook marks these data sets as being used whenever the management functions are run. The IBM SU 60 prevents other functions, such as archive, from finding and processing data sets that are truly inactive. A potentially large number of data sets appear used and are kept on disk because of storage management jobs routinely being run. The CA Disk SVC exempts the management jobs from causing these updates.

- The CA Disk SVC can also record the date whenever the change bit is turned on (mod date). This record is useful in report information. SVC also provides a convenient and low-overhead means of eliminating redundant backup copies of data sets (that full pack dumps or other techniques often create). For more information about recording the date, see "Backup and Archive Consideration" in the *User Guide*.
- The CA Disk data set utilization (DSU) report loses much of its usefulness because it reports on the additional fields that CA Disk SVC maintains.

Note: HFS data sets, DB2, and certain other applications bypass normal open processing and do not maintain a last modification date.

SVC Requirements

If you are going to install CA Disk SVC, provide a user SVC of type 3 or 4. Enable the SVC for interrupts and designate as either restricted or nonrestricted.

Note: For more information, see User SVC (see page 14).

SVC Integrity and Security

Before any fields are modified in the DSCB, the SVC verifies that all INPUT parameters contain an F1-DSCB. SVC also checks to ensure that the caller is in key 0 and in supervisor state.

VTOC Fields Maintained by the CA Disk SVC

If the IBM modules are zapped to execute the CA Disk SVC, the following fields are maintained in the F1-DSCB for each data set:

1

Last use date.

Format: ydd

Bytes: 3

Offset (DEC): 75 SU 60

2

Last modify date.

Format: ydd

Bytes: 3

Offset (DEC): 70

3

Job name or account code.

Format: char

Bytes: 8

Offset (DEC): 62

4

Count of updates to VTOC entry.

Format: bin

Bytes: 2

Offset (DEC): 73

E5

Change bit.

Format: bit

Bytes: x'02'

Offset (DEC): 93 SU 60

6

SVC mode.

Format: char

Bytes: 1

Offset (DEC): 103

Optionally, you can specify Fields 1 and 2 as a 4 byte packed format (ccyydddF). If they are specified, adjust some of the field offsets to prevent overlaying other format 1 DSCB fields. The macro settings within the CA Disk SVC determine Field 3. You can elect to maintain an accounting code or job name. The accounting code measures from the first accounting field on the JOB statement for the job or user accessing the data set. You can also select how to maintain this field:

- The first time the data set is updated.
- Each time the data set is modified.
- Each time the data set is used.

CA Disk SVC Modes

The CA Disk SVC is distributed with a control switch called SVCMODE that affects which fields are maintained and how often.

SVCMODE=4

This mode of operation makes the minimum number of updates to the VTOC entry (F1-DSCB) for a data set, yet maintains all of the fields. The SVC source is distributed in this mode. It updates the F1-DSCB only if:

- The last used date needs to be changed.
- The change bit needs to be set on (change bit turned on when data set opened for output), and MODDT is also updated

Note: (1) causes only one update per day, and (2) causes only one per day unless the change bit is being turned off more often than that (for example, running incremental backup two or more times per day).

Field 4 shows the number of times the F1-DSCB has been updated, not necessarily the number of times a data set has been opened.

SVCMODE=5

This mode of operation causes updates under exactly the same conditions as mode 4, but only the SU 60 fields (1 and 5) and the mode field itself (6) is maintained. This is an SU 60 look-alike mode, with the benefit of the exemption tables.

SVCMODE=6

This mode of operation causes the F1-DSCB to be updated every time the data set is opened. Field 4 shows the number of times the data set has been opened.

This technique causes many more updates of the VTOC entry, and is known to cause occasional problems with partitioned data sets, made evident by duplicate TTR directory entries and loss (overlay) of a member update. This usually occurs only for PDSs that are accessed very heavily and by two or more concurrent users. Shared DASD environments without a *cross-system enqueue package* are especially susceptible.

Tailoring the CA Disk SVC

Source code for the CA Disk SVC is supplied in member ADSMVS60 of the CCUWSAMP library. The four exclusion tables are supplied in members ADSOPPGM, ADSOPDSN, ADSOPJOB, and ADSOPVOL. The tables can be used to exclude from DSCB updating any combination of data sets, DASD volumes, jobnames and programs. The program exclusion capability is recommended for use when running CA Disk jobs. This technique prevents a CA Disk management job from causing a data set to appear to have been used. The table code has each table initialized with either a real or dummy entry to illustrate how entries are defined. The CA Disk SVC is provided with the tables already linked. If they are modified, they must be re-assembled and linked into the CA Disk SVC.

Note: The activating zap is such that the CA Disk SVC *replaces* the MVS SU 60 update code (that is, if the SVC is activated, SU 60 decisions are deactivated).

ADSOPPGM Member

The program exclusion table is initialized to bypass DSCB updating for any program starting with *ADSMI*, as the CA Disk control module does. Patterns are not permitted.

The following illustration is a sample of the ADSOPPGM source:

Note: The LENGTH parameter within the ADSOPPGM source must be left untouched. This is the maximum length for each program name and is set at 8.

```
   Menu  Utilities  Compilers  Help
 ------------------------------------------------------------------------------
BROWSE                         CCUWSAMP(ADSOPPGM)    Line 00000000 Col 001 080
 Command ===>                                                 Scroll ===> CSR
******************************** Top of Data ********************************
ADSOPPGM TITLE 'OPEN SVC EXEMPTION LIST BY PROGRAM NAME'
*********************************************************************
         COMPILE ASEM=RENT,LKED=RENT                                 *
*                                                                    *
* DESCRIPTION:                                                       *
*    PGM EXCLUSION/EXEMPTION LIST                                    *
*                                                                    *
* MAINTENANCE HISTORY                                                *
*--------------------------------------------------------------------*
  MHIST DEFINE
  MHIST 10/16/00,MJB,@000,381172,'EXTERNALIZE EXCLUSION LISTS'
  MHIST 04/10/07,THC,@001,711573,'UPDATE EXCLUDE TABLE'
  MHIST DEFEND
*********************************************************************
*
* SCRNLIST MACRO PARAMETER DESCRIPTIONS:
*
*    SCREEN=     SCREEN PARAMETER LIST. MULTIPLE VALUES CAN BE
*                CODED BETWEEN PARENTHESIS AND SEPARATED BY
*                COMMAS. PARENTHESIS ARE NOT NEEDED WHEN ONLY ONE
*                VALUE IS SPECIFIED.
*                EXAMPLES:
*                SCREEN=(ADSMI,ISR,IDC) IN THIS CASE 3 VALUES
*                ARE PASSED
*                SCREEN=ADS IN THIS CASE 1 VALUE IS PASSED
*
*    LENGTH=     LENGTH OF EACH SCREEN ENTRY.
*                LENGTH=8 IS PROGRAM NAME LENGTH
*
*    PLEASE NOTE THAT WHEN MULTIPLE LINES ARE NEEDED, A NON-BLANK
*    CHARACTER SHOULD BE CODED IN COLUMN 72 AS A CONTINUATION
*    INDICATOR FIELD, AND THE CONTINUATION LINE SHOULD START IN
*    COLUMN 16.
*
*    READ THE SCRNLIST MACRO PROLOGUE FOR A COMPLETE LIST OF
*    AVAILABLE PARAMETERS.
*
*********************************************************************
         SPACE 3
ADSOPPGM SCRNLIST LENGTH=8,                                          X
               SCREEN=ADSMI              EXCLUDE OURSELF
*********************************************************************
* USE THE FOLLOWING TO EXCLUDE DFHSM & DFDSS (MODIFY APPROPRIATELY)
*********************************************************************
*              SCREEN=(ADSMI,              EXCLUDE OURSELF           X
*              ARC,                              DFHSM               X
*              ADR)                              DFDSS            @001
*
         END
******************************** Bottom of Data ********************************
```

ADSOPDSN Member

The data set name exclusion table is initialized to bypass updating the DSCB for temporary data sets and IEHMOVE temporary data sets. Patterns are not permitted but TEMP is used to designate system temporary data sets.

The following illustration is a sample of the ADSOPDSN source:

```
   Menu  Utilities  Compilers  Help
 ------------------------------------------------------------------------------
 BROWSE                       CCUWSAMP(ADSOPDSN)    Line 00000000 Col 001 080
 Command ===>                                          Scroll ===> CSR
******************************** Top of Data ********************************
ADSOPDSN TITLE 'OPEN SVC EXEMPTION LIST BY DATA SET NAME'
*********************************************************************
         COMPILE ASEM=RENT,LKED=RENT                                 *
*                                                                    *
* DESCRIPTION:                                                       *
*    DSN EXCLUSION/EXEMPTION LIST                                    *
*                                                                    *
* MAINTENANCE HISTORY                                                *
*--------------------------------------------------------------------*
  MHIST DEFINE
  MHIST 10/16/00,MJB,@000,381172,'EXTERNALIZE EXCLUSION LISTS'
  MHIST DEFEND
*********************************************************************
*
* SCRNLIST MACRO PARAMETER DESCRIPTIONS:
*
*    SCREEN=     SCREEN PARAMETER LIST. MULTIPLE VALUES CAN BE
*                CODED BETWEEN PARENTHESIS AND SEPARATED BY
*                COMMAS. PARENTHESIS ARE NOT NEEDED WHEN ONLY ONE
*                VALUE IS SPECIFIED.  USE ONLY DATA SET NAMES OR
*                PREFIXES. USE ¬TEMP TO EXCLUDE SYSTEM TEMPORARY
*                DATA SETS (MATCHES DSN=SYS?????.T??????.?????./)
*                NO PATTERNS ARE USED.
*                EXAMPLES:
*                SCREEN=DMSOS.GLH1  IN THIS CASE 1 VALUE IS PASSED
*                SCREEN=(DMSOS,TEMP,SYSZ) IN THIS CASE 3 VALUES
*                ARE PASSED
*
*    LENGTH=     LENGTH OF EACH SCREEN ENTRY.
*                LENGTH=44 IS DATA SET NAME LENGTH
*
*    PLEASE NOTE THAT WHEN MULTIPLE LINES ARE NEEDED, A NON-BLANK
*    CHARACTER SHOULD BE CODED IN COLUMN 72 AS A CONTINUATION
*    INDICATOR FIELD, AND THE CONTINUATION LINE SHOULD START IN
*    COLUMN 16.
*
*    READ THE SCRNLIST MACRO PROLOGUE FOR A COMPLETE LIST OF
*    AVAILABLE PARAMETERS.
*
*********************************************************************
         SPACE 3
ADSOPDSN SCRNLIST LENGTH=44,                                         X
               SCREEN=(¬TEMP,    SYS?????.T??????.?????./            X
               **SYSUT)          IEHMOVE TEMPORARY DATASETS
*
         END
****************************** Bottom of Data ******************************
```

Note: The LENGTH parameter within the *ADSOPDSN* source must be left untouched. This is the maximum length for each data set name, and is set at 44.

ADSOPJOB Member

The jobname exclusion table is initialized to empty. Patterns are not permitted. The following illustration is a sample of the ADSOPJOB source:

```
VIEW                                    - 01.02          Columns 00001 00072
Command ===>                                              Scroll ===> CSR
****** ***************************** Top of Data ******************************
000100 ADSOPJOB TITLE 'OPEN SVC EXEMPTION LIST BY JOB NAME'
000200 ********************************************************************
000300          COMPILE ASEM=RENT,LKED=RENT                                *
000400 *                                                                   *
000500 * DESCRIPTION:                                                      *
000600 *    JOB EXCLUSION/EXEMPTION LIST                                   *
000700 *                                                                   *
000800 * MAINTENANCE HISTORY                                               *
000900 *-------------------------------------------------------------------*
001000   MHIST DEFINE
001100   MHIST 10/16/00,MJB,@000,381172,'EXTERNALIZE EXCLUSION LISTS'
001200   MHIST DEFEND
001300 ********************************************************************
001400 *
001500 * SCRNLIST MACRO PARAMETER DESCRIPTIONS:
001600 *
001700 *    SCREEN=     SCREEN PARAMETER LIST. MULTIPLE VALUES CAN BE
001800 *                CODED BETWEEN PARENTHESIS AND SEPARATED BY
001900 *                COMMAS. PARENTHESIS ARE NOT NEEDED WHEN ONLY ONE
002000 *                VALUE IS SPECIFIED. USE EITHER FULL OR PARTIAL
002100 *                JOB NAMES BUT NO PATTERNS.
002200 *                EXAMPLES:
002300 *                SCREEN=(BACK,REL,ARC) IN THIS CASE 3 VALUES
002400 *                ARE PASSED
002500 *                SCREEN=DMS IN THIS CASE 1 VALUE IS PASSED
002600 *
002700 *    LENGTH=     LENGTH OF EACH SCREEN ENTRY.
002800 *                LENGTH=8 IS JOB NAME LENGTH
002900 *
003000 *    PLEASE NOTE THAT WHEN MULTIPLE LINES ARE NEEDED, A NON-BLANK
003100 *    CHARACTER SHOULD BE CODED IN COLUMN 72 AS A CONTINUATION
003200 *    INDICATOR FIELD, AND THE CONTINUATION LINE SHOULD START IN
003300 *    COLUMN 16.
003400 *
003500 *    READ THE SCRNLIST MACRO PROLOGUE FOR A COMPLETE LIST OF
003600 *    AVAILABLE PARAMETERS.
003700 *
003800 ********************************************************************
003900          SPACE 3
004000 ADSOPJOB SCRNLIST LENGTH=8,                                        X
004100                SCREEN=
004200 *
004300          END
```

Note: The LENGTH parameter within the *ADSOPJOB* source must be left untouched. This is the maximum length for each data set name, and is set at 8.

ADSOPVOL Member

The volume serial exclusion table is initialized to empty. Patterns are not permitted. The following illustration is a sample of the ADSOPVOL source:

```
VIEW                                    - 01.03          Columns 00001 00072
Command ===>                                               Scroll ===> CSR
****** ***************************** Top of Data ******************************
000100 ADSOPVOL TITLE 'OPEN SVC EXEMPTION LIST BY VOLSER NAME'
000200 ********************************************************************
000300          COMPILE ASEM=RENT,LKED=RENT                                *
000400 *                                                                   *
000500 * DESCRIPTION:                                                      *
000600 *    VOL EXCLUSION/EXEMPTION LIST                                   *
000700 *                                                                   *
000800 * MAINTENANCE HISTORY                                               *
000900 *-------------------------------------------------------------------*
001000   MHIST DEFINE
001100   MHIST 10/16/00,MJB,@000,381172,'EXTERNALIZE EXCLUSION LISTS'
001200   MHIST DEFEND
001300 ********************************************************************
001400 *
001500 * SCRNLIST MACRO PARAMETER DESCRIPTIONS:
001600 *
001700 *    SCREEN=     SCREEN PARAMETER LIST. MULTIPLE VALUES CAN BE
001800 *                CODED BETWEEN PARENTHESIS AND SEPARATED BY
001900 *                COMMAS. PARENTHESIS ARE NOT NEEDED WHEN ONLY ONE
002000 *                VALUE IS SPECIFIED.
002100 *                EXAMPLES:
002200 *                SCREEN=(WRK,DS2,IDC) IN THIS CASE 3 VALUES
002300 *                ARE PASSED
002500 *                SCREEN=TSOP IN THIS CASE 1 VALUE IS PASSED
002700 *
002800 *    LENGTH=     LENGTH OF EACH SCREEN ENTRY.
003200 *                LENGTH=6 IS VOLSER LENGTH
003300 *
003400 *    PLEASE NOTE THAT WHEN MULTIPLE LINES ARE NEEDED, A NON-BLANK
003500 *    CHARACTER SHOULD BE CODED IN COLUMN 72 AS A CONTINUATION
003600 *    INDICATOR FIELD, AND THE CONTINUATION LINE SHOULD START IN
003700 *    COLUMN 16.
003800 *
003900 *    READ THE SCRNLIST MACRO PROLOGUE FOR A COMPLETE LIST OF
004000 *    AVAILABLE PARAMETERS.
004100 *
004200 ********************************************************************
004300          SPACE 3
004400 ADSOPVOL SCRNLIST LENGTH=6,                                         X
004500                SCREEN=
004660 *
004670          END
```

Note: The LENGTH parameter within the *ADSOPVOL* source must be left untouched. This is the maximum length for each volume serial, and is set at 6.

After your exclusion lists have been defined, an assembly process can be applied with SMP/E in the form of a usermod. Installation library member USERMODM provides an example of this process.

The following is a sample of USERMODM:

```
VIEW                                        - 01.00              Columns 00001 00080
Command ===>                                                       Scroll ===> CSR
****** ********************************* Top of Data ***********************************
000001 //JOBNAME  JOB (ACCT INFO)
000002 //*
000003 //* ********************************************************************
000004 //* *        INSTALL EXCLUSION TABLES FOR OPEN SVC                      *
000005 //* *                                                                   *
000006 //* ********************************************************************
000007 //SMP     EXEC PGM=GIMSMP,REGION=5120K,
000008 //             PARM='CSI=%%CSI_PREFIX..GLOBAL.CSI'
000009 //SYSPRINT  DD SYSOUT=*
000010 //SMPRPT    DD SYSOUT=*
000011 //SMPOUT    DD SYSOUT=*
000012 //SMPHOLD   DD DUMMY
000013 //USERASM   DD DISP=SHR,DSN=USERID.ASM
000014 //SMPCNTL   DD *
000015 SET BOUNDARY(GLOBAL).
000016 RECEIVE  S(%%USERMOD_PREFIX.1819).
000017 SET BOUNDARY(%%TARGET_ZONE).
000018 APPLY S(%%USERMOD_PREFIX.1819) ASSEM .
000019 //SMPPTFIN  DD *
000020 ++USERMOD(%%USERMOD_PREFIX.1819).
000021 ++VER(Z038) FMID(%%PRODUCT_FMID) .
000022 ++SRC(ADSOPPGM) DISTLIB(ADMSASM) DISTMOD(ADMSLOAD) TXLIB(USERASM).
000023 ++SRC(ADSOPDSN) DISTLIB(ADMSASM) DISTMOD(ADMSLOAD) TXLIB(USERASM).
000024 ++SRC(ADSOPJOB) DISTLIB(ADMSASM) DISTMOD(ADMSLOAD) TXLIB(USERASM).
000025 ++SRC(ADSOPVOL) DISTLIB(ADMSASM) DISTMOD(ADMSLOAD) TXLIB(USERASM).
000026 /*
```

Installing the CA Disk SVC

To install the CA Disk SVC, perform the following procedure:

1. Based on the previous discussion; determine which SVC mode of operation you want to use. If you select mode 4, the default mode, skip the remainder of this substep and proceed to substep 2. If you select mode 5 or mode 6, you must change the supplied source code for the variable SVCMODE from C'4' to C'5' or C'6', corresponding to the mode you selected.

2. Decide whether the CA Disk SVC maintains the job name or the accounting code. Then, update the source code in member ADSOPJBM provided in CCUWSAMP library to reflect your selection. The default is to keep the creating JOBNAME. If this is appropriate, skip to substep 3.

   The macro, $JOBNM, has two parameters: FLD and TYPE. The parameter FLD defines the field to maintain. Valid values are ACCT or JOB. To maintain a job name, specify this parameter with a value of JOB. The parameter TYPE defines when this field is updated based on SVCMODE. Valid values are CREATE, MOD, or USE. To have the field updated each time the data set is accessed, set this parameter to a value of USE and set SVCMODE toThe following is an example of the use of this macro:

   ```
   $JOBNM FLD=ACCT,TYPE=CREATE
   ```

   With the macro coded in this way, the CA Disk SVC maintains the first eight bytes of the creating job's accounting for each data set.

   See Installation library member USERMODN to assemble and link ADSOPJBN into the CA Disk SVC.

3. Determine which SVC number to use for the CA Disk SVC. The CA Disk SVC is normally installed as number 244. 244 is used by some IBM system products as well as other packages. If SVC 244 is already assigned in your system, select a different number for the CA Disk SVC.

   If CA Disk can use SVC 244, proceed to step 4.

   Ensure that the SVC is in the IEASVCxx member of your system parmlib. It should be a type 3 or type 4 SVC, enabled for interrupts. Designating as either restricted or non-restricted is optional.

4. Determine whether any of the fields maintained by the CA Disk SVC will cause a problem at your site. Problems can occur if there are other products that maintain fields in the F1-DSCB. Problems can also occur if packed format is used for Last Modified and Last Use dates, because 4 bytes will be needed to store each one of these dates overlaying other fields in the Format 1 DSCB unless the involved field offsets are adjusted accordingly. If any of the other fields causes a problem, modify the CA Disk SVC to either move the field or not maintain the field.

   To move a field, take the equate label from the following table and modify the SVC source to equate the label to the new value (minus 44) and set the corresponding CA Disk sysparm from the following table to the same value (not minus 44).

   To not maintain a field, take the equate label from the following list and modify the SVC source to comment out any code that uses the equate to examine or change the DSCB, then set the corresponding CA Disk sysparm from the following list to a value of 000 in your parmlib data set:

   SVC Equate Label: DMSUSEDT

   Last opened date (SU 60).

   Sysparm Name: DSCBLUSD

   Default Value: 75

   SVC Equate Label: DMSDSIND

   Change bit x'02' (SU 60).

   Sysparm Name: n/a

   Default Value: 93

   SVC Equate Label: DMSMODDT

   Last modified date.

   Sysparm Name: DSCBLMOD

   Default Value: 70

   SVC Equate Label: DMSSVCMO

   CA Disk SVC mode.

   Sysparm Name: DSCBSVMD

   Default Value: 103

SVC Equate Label: DMSJOBNM

Job name or accounting code.

Sysparm Name: DSCBJBNM

Default Value: 62

SVC Equate Label: DMSOPCNT

Open count (two bytes).

Sysparm Name: DSCBOPCD

Default Value: 73

The value of each SVC equate label (minus 44) must be equal to the value of the corresponding CA Disk sysparm.

5. A diagnostic facility exists within the SVC that can display applicable data used to determine if DSCB updates are to be made to maintain fields within the DSCB. These diagnostic messages are automatically generated if the job name is equal to the job name identified in field DIAGNAME. The default name is DMSTEST. To change the job name used to generate diagnostic messages, update the DIAGNAME field.

   Note: A change to the DIAGNAME (DMSTEST) will require re-assembling the ADSMVS60 module, (USERMOD1) after updating the source provided in the CCUWSAMP library.

   Similar diagnostic facilities are provided for the VSAM date stamp module and auto-restore catalog management hook. If you change the job name for the SVC, consider changing the job names used for the other diagnostic facilities. This is discussed where applicable for the other functions later in this section.

6. If you have installed a version of the CA Disk SVC from a release prior to Release 8.1, you can have compatibility problems when using CA Disk SVC. See the CA Disk SVC section for a discussion of these problems and how to correct them.

7. The CA Disk SVC is now installed into the CA Disk LPA List library. It does not need to be copied to SYS1.LPALIB. You only need to add the CA Disk LPA List library to your Systems LPA list.

   A copy of the CA Disk SVC with default values is placed in the CA Disk LPA List library at product installation time. If the default values are appropriate for your installation, you do not need to reassemble the SVC.

   To change any of the default values, make your changes to the source found in member ADSMVS60 of the CCUWSAMP. Then run job USERMOD1 in the installation library to reassemble and link the SVC into the CA Disk LPA List library.

   If you are running MVS 2.2.0 or higher, you do not have to change the name of the CA Disk SVC to IGC00*nnn* in order to install it. However, install the SVC by adding the following line to the IEASVCXX member of your SYS1.PARMLIB:

   ```
   SVCPARM 244,REPLACE,TYPE(4),APF(NO),EPNAME(ADSMVS60)
   ```

8. To activate the SVC, it must be called as part of the processing for each disk data set. Applying a zap to the proper IBM module for your operating system does this. The following list contains the IBM module names to be zapped depending on IBM product level. Apply 1 or more of the following:

   IBM Module: IFG0194A

   CSECT: IFG0194E

   DFP/DFSMS Level: All levels

   MODEL: ZUZMODEL

   USERMOD: USERMOD2

   IBM Module: IGC0001I

   CSECT: IFG0196W

   DFP/DFSMS Level: All except DFP 2.4 and below

   MODEL: ZOZMODEL

   USERMOD: USERMODH

   IBM Module: IGGDADSM

   CSECT: IGGDAU01

   DFP/DFSMS Levels:

   - DFSMS 1.2 w/UW28103
   - DFSMS 1.3 w/UW28104
   - DFSMS 1.4 and above

   MODEL: ZZZMODEL

   USERMOD: USERMODI

   IBM Module: IGC00020

   CSECT: IFG0201R

   Levels:

   - DFSMS 1.2 w/UW28103
   - DFSMS 1.3 w/UW28104
   - DFSMS 1.4 and above

   MODEL: ZYZMODEL

   USERMOD: USERMODJ

There are many levels of the IBM modules, and each requires a slightly different form of the zap. Zaps to match these maintenance levels are provided in the installation library. Use the one in which all the VERs match your version of the IBM modules. To accomplish this, locate your Operating System level in the preceding table and do the following:

- Get a zap dump of the IBM module
- Use SMP to determine the RMID of this same module

  Note: Try to match the RMID to a corresponding member in the CA Disk installation library named ZUA*nnnnn* or Z1*nnnnn*. If you cannot find a supplied zap that matches your Operating System level, use a model zap and adjust it to fit. In the MODEL column of the preceding Table, corresponding to your system level, is the name of a member in the installation library. If you need assistance in making the adjustments, contact the CA Disk Technical Support Center.

- Make sure that all of the VERs in the zap; match your dump of the module.

9. If the CA Disk SVC number you are using is not the default of 244 (x'F4'), you must now adjust the zaps to be sure they execute the correct SVC; for example, for SVC 245, change 0AF4 to 0AF5 within the body of the zap itself.
10. In the USERMOD column of the preceding table corresponding to your system level is the name of a member in the installation library. You can use this member as sample JCL to SMP/E install the IBM zap. This JCL installs a usermod onto your operating system.
11. IPL the system to load the modified IBM modules and the CA Disk SVC into storage. For MVS, do an MLPA or CLPA to initialize the link pack area.

The CA Disk VSAM Date Stamp

Last use date support for VSAM is maintained by IBM module IDATMSTP in their ICF VSAM support. CA Disk provides a sample date stamp routine for you to use as member IDATMSTP in the installation library. The advantage of using the CA Disk version of IDATMSTP is that it provides exemption entries so that the CA Disk management tasks themselves do not cause VSAM data sets to appear used, similar to the support in the CA Disk SVC for non-VSAM data sets.**here

Note: The VSAM Date Stamp module updates the change bit (DS1IND02 bit in the Format-1 DSCB).

Modifications may be necessary to the source for the CA Disk VSAM data stamp module to initialize exclusion tables and to optionally change the diagnostic job name.

To provide the exemption lists support, CA Disk replaces the operating system's generic IDATMSTP module with its own copy that accommodates customized exemption lists by data set, job name, volume serial, or program. Note that the CA Disk SVC and IDATMSTP exemption lists are separate, and can contain different entries.

The CA Disk VSAM date stamp module contains a diagnostic facility, which displays the data, which CA Disk uses to decide if the last use date will be updated.

These diagnostic messages are generated automatically if the job name matches the job name identified in field DIAGNAME in ADSTMTST CSECT. The default name is DMSTEST. If you cannot use the job name DMSTEST, update the DIAGNAME field in CSECT ADSTMTST. The source code for ADSTMTST is in the CCUWSAMP library. Similar diagnostic facilities are provided for the IBM SVC module and auto-restore catalog management hook. If you change the job name for the VSAM date stamp module, consider changing the job names used for those functions.

CA Disk provides default exclusion lists, which bypass the VSAM date stamp update during CA Disk operations. To make additional entries, find the correct member in the CCUWSAMP library, and make your changes. The following is a list of CCUWSAMP members and the related exclusion list:

ADSTMDSN

Exclusion List: Data Set Name

ADSTMJOB

Exclusion List: Job Name

ADSTMPGM

Exclusion List: Program

ADSTMVOL

Exclusion List: Volume Serial

After you have made your modifications, you are ready to install date stamp support. Even if you have not made any changes, you should install the default date stamp module, to exempt CA Disk jobs from updating the reference date.

Installing the IDATMSTP Module

Use member USERMOD3 in the CA Disk installation library to replace the operating system's copy of IDATMSTP with the default CA Disk version. You must install the default version first, even if you are installing customized exclusion lists as well. SMP/E assembles and properly links the module into the correct library, regardless of your operating system level.

If you modified the exclusion lists or the job name used to generate diagnostics, use member USERMOD4 to install the customized lists. For further changes to the exclusion lists, you can use USERMOD4 to update the lists directly. You will not need to execute USERMOD3 again unless the operating system has made changes to IDATMSTP.

Testing the Installation of the IDATMSTP Module

You should now test the IDATMSTP module to see if it has been installed successfully. Submit a job that will update an ICF VSAM data set, using the jobname that you specified within the ADSTMTST CSECT. Diagnostic messages will be printed to show that IDATMSTP has been entered and what action, if any, will be taken. A return code will also be printed when IDATMSTP is exited. A return code of zero indicates that the last use date will not be updated.

The Auto-Restore Function

One of the more important reasons an installation purchases any data storage management system is to minimize the amount of DASD space allocated to data sets that are rarely or never accessed. The amount of DASD space allocated can be accomplished in CA Disk by both explicit and implicit archival, based on last used date. Typically, the more constrained an installation is for DASD space, the shorter the time allowed for a data set to be unused prior to archival. As the unused time window shrinks, more data sets are archived that do get accessed infrequently. This is especially common for data sets that get referenced only on a monthly or yearly basis.

Although this archival scheme accomplishes one of the DASD manager's primary functions--ensuring enough DASD free space exists for day-to-day operations--it can also create a less than amicable relationship with the data center's users if not implemented properly. For example, if a system is set up to archive all data sets not used within two weeks, those data sets that are only accessed once a month are archived between each job cycle.

The user must either respond to this problem by restoring all required data sets prior to a job run, or artificially referencing the data sets periodically to avoid having them archived. In the first case, an extra workload is created for the user, and in the second case the DASD manager is prevented from doing an effective job. Although the DASD manager could exempt these data sets from being archived by CA Disk, this cannot be a good alternative since the data sets could then remain on DASD long after the system that uses the data sets is removed from production. And if an installation is critically short of free DASD space, it cannot have the luxury of letting seldom- used data sets tie up valuable DASD space between job cycles.

An automatic restore capability helps to solve this potential conflict between the DASD manager and the end user by allowing data sets to be archived by the DASD manager and then restored automatically by CA Disk if they are required by an application later. This solves the DASD manager's problem of keeping free space available and also relieves the user of the burden of restoring required data sets before cyclical job runs.

The Two Auto—Restore Methods

CA Disk has provided two methods to perform auto-restores. One method is the S213 Abend Exit (F1-DSCB Not Found) and the other method is a Catalog Management hook.

New users of CA Disk are encouraged to install the Catalog Management hook and NOT install the S213 Abend Exit. If you have installed the S213 Exit from previous releases, you are encouraged to remove it. See the S213 Abend Exit section for documentation about installing and removing the S213 Abend Exit.

Catalog Management Hook (IGG026DU Module)

IBM provides the IGG026DU module in catalog management to initiate auto-restores. This module is not really an exit, but it hooks into the front end and back end of catalog management processing. Since many catalog management requests do not require auto-restores, this module must decide dynamically which requests to intercept and which to pass on.

DFHSM SVC Hook (IGX00024 Module)

The IGX00024 module is installed in conjunction with the Catalog Management Hook and comprises the second method. It is used primarily to intercept DFHSM auto recall requests that should be processed by CA Disk.

Several program products, DB2 for example, issue catalog locates and look at the returned volser to anticipate the need for a data set to be restored. The program recognizes the special volser MIGRAT as the volser DFHSM uses when it archives a data set. If this information shows that the data set is archived, the program tries to directly invoke DFHSM to restore the data set. Because CA Disk functions differently from DFHSM, the products do not recognize that a data set is archived by CA Disk and can encounter errors later. This hook, in conjunction with the Catalog Management hook, provides a means of overcoming this problem.

This hook is designed to work on your system, whether or not you have DFHSM. It is automatically installed with the Catalog Management Hook.

A catalog locate is usually issued by a routine trying to establish the existence of a data set. If the locate is successful, the data set exists; if it fails, it either does not exist or it is not cataloged. The key for the catalog management hook to go to work is when the locate is issued from a selected function (dynamic allocation, for example) and the return from catalog management is successful, but the data set is cataloged to the CA Disk pseudo-volume. (The pseudo-volume--the CA Disk default name is ARCIVE--is an imaginary volume to which CA Disk optionally re-catalogs a data set when it is archived and scratched, to help identify the data set as being archived.) When these conditions are met, CA Disk automatically restores the data set and then returns to the requestor of the local of the real volume to which it was restored. In this way, the locate requestor is never aware that a restore operation took place. We discuss how to catalog data sets to the pseudo-volume later in these customization instructions.

Programs that check for DFHSM's MIGRAT volser can have problems during the catalog management process of auto restore. For these programs, CA Disk does not actually auto-restore the data set. It passes back the MIGRAT volser contained in the catalog entry, instead of the ARCIVE volser. This action allows the programs to issue a request to DFHSM to restore the data set. Then, CA Disk intercepts the request with the DFHSM SVC Hook and invokes an auto-restore.

CA Disk can recognize that the DFHSM request is for a CA Disk archived data sets by doing its own catalog locate. If the ARCIVE volser is returned, CA Disk was the program that archived the data set.

In the preceding description of the catalog management hook, we said that only specific functions cause auto-restores to take place. For example, it is not good to automatically restore all data sets referenced by an IDCAMS LISTCAT job. Instead, these locates are ignored by the catalog management hook so that the output listing show a volume of ARCIVE for the data set. However, CA Disk intercepts locate requests from dynamic allocation, job initiation, and certain ISPF/TSO requests.

If you use IDCAMS to delete a data set that has been cataloged to the CA Disk pseudo-volume, the data set is uncataloged from the pseudo-volume instead of being auto-restored. The return code for the IDCAMS program is set to indicate the delete was successful. If the data set cannot be uncataloged (possibly for security reasons), CA Disk attempts to auto-restore the data set, so that IDCAMS can attempt the delete.

You see in the installation section that CA Disk does not actually replace IFG0EX0A, IGG026DU and IGX00024 modules because there are other software products that also replace these modules. To allow the user to install both products, we have decided to call our modules by different names. In the installation process, you dynamically or statically relink the appropriate IBM programs and define their entry points as the CA Disk modules. We bring up this topic now because it is important if your installation is also using IBM's program product DFHSM.

Conceptual Diagram of CA Disk Auto-Restore Methods

The CA Disk catalog management hook uses CSECT ADSAR010 as part of module IGG026DU, and ADSHS001 as part of the module IGX00024. The CA Disk F1-DSCB missing exit uses CSECT ADSAR004 as part of module IFG0EX0A. After installing the appropriate modules, you have the following conceptual arrangement. The following illustration is a sample Auto-Restore Method:

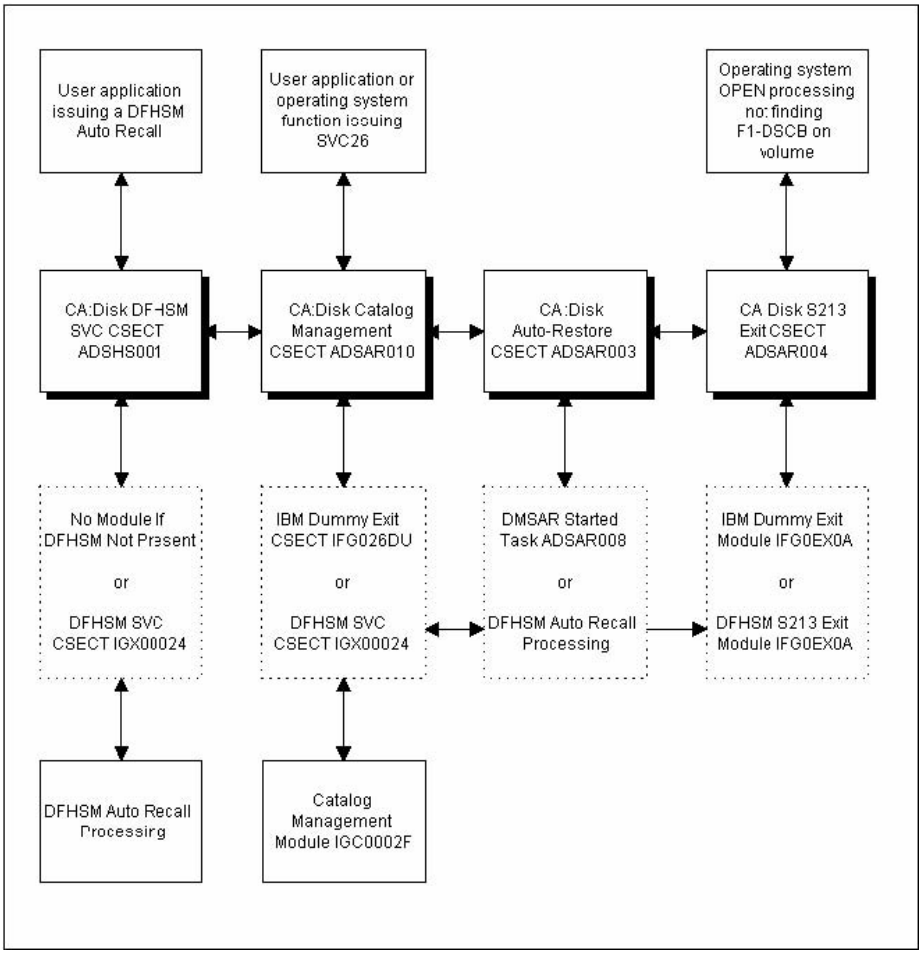

The previous illustration shows that control is passed from the user doing a catalog locate to the catalog management hook and then to either the DFHSM hook or a dummy module. From there, control is passed to catalog management. Catalog management performs its specified function and returns control to the next module up the chain, giving DFHSM a chance to auto recall the data set if appropriate. DFHSM returns control to the next module up the chain. Note that if the dummy module is installed, it just returns control up the chain without any other processing.

CA Disk now gets control. If DFHSM had successfully recalled the data set, it now appears to CA Disk that the data set exists on a real volume. If the data set is cataloged to the CA Disk pseudo volume, CA Disk takes one of two actions. If the program issuing the SVC 26 is recognized as one that needs the special handling provided by the CA Disk DFHSM SVC hook, it returns a volser of MIGRAT. If CA Disk does not recognize the calling program, CA Disk attempts to auto-restore the data set by using the DMSAR started task. If either DFHSM or CA Disk auto-restored the data set, it appears to the user that the data set exists on a real volume.

This design allows auto-restores from either the CA Disk archives or DFHSM migration volumes on the same system without interference. Other software products such as CA Allocate and CA ACF2 also front- and back-end catalog management.

The illustration shows that control is passed from the user invoking a DFHSM Auto Recall to the DFHSM SVC hook. Then, control is passed to the catalog management hook. From this point on, the catalog management hook process is as documented above. When the auto-restore is complete, control is returned to the DFHSM SVC hook. If the data set was not archived by CA Disk, control is passed to the DFHSM SVC for its own Auto Recall processing.

Again, by examining the illustration you can see that if the F1-DSCB missing exit is invoked, CA Disk gets control first in its exit as CSECT ADSAR004 of module IFG0EX0A. If it cannot restore the data set, it calls the CSECT IFG0EX0A to attempt the restore before returning to the user.

Installing the Auto-Restore Function

The auto-restore function applies to MVS installations only, and should be implemented only after a thorough review of the following installation procedure. Be sure to complete each step in the process before continuing to the next. The following discussion applies to all three hooks.

To install the Auto-Restore Function

1. See the procedure named DMSAR in the proclib library shipped with your CA Disk system. In particular, see the following items:

   Note: If you rename the DMSAR catalog procedure, you must also change the value of sysparm ARTASKNM (the default name is DMSAR). Otherwise, the support does not work.

   - Reviews the dispatching priority assigned to the procedure and makes any necessary adjustments based on your installation's requirements.
   - Supply the proper data set name for the STEPLIB.

- If you want dynamic tape allocation, leave the asterisk in column 3 of the ARCHIVES DD statement. If you do not want dynamic allocation to be performed by CA Disk, remove the asterisk. The preferred method is to let CA Disk dynamically allocate the tape device, since this allows requests that do not require tape devices to be processed immediately, even if no tape devices are available. This condition occurs for those data sets that have been archived to disk, have no record in the CA Disk archive index, or are rejected by the user screening exit ARESPREX.
- If you put an ARCHIVES or ARCHIVER DD statement in your JCL, do not specify the DSN= parameter. If you specify a data set name, no concurrent auto-restores can take place. In addition, you cannot be able to run auto-restores at the same time as CA Disk backup, archival, restore or recover jobs.
- If the archives reside on disk rather than tape, the appropriate disk archive data set is dynamically allocated with no operator intervention required.
- See the IDCUT1 and IDCUT2 DD statements. These statements are discussed in the Activating VSAM Support section in the chapter “Installation.”
- If your installation maintains multiple FILES and you want to make them all available to the auto-restore function, use the MFILES DD statement instead of the FILES DD statement. Add each files data set name to the MFILES concatenation, up to 256 data sets are supported. The order of the concatenation determines the order of search for the data set. The first data set that contains an index record for the archived data set is used, even if another files data set contains a more recent version of the data set. Therefore, it is very important that the FILES be specified in the proper order. If you are using all FDBs in the MFILES DD statement, then the databases are logically joined and searched as single entity.

  Note: The FILES DD statement is ignored if there is also an MFILES DD statement in the procedure. Note that concatenated files data sets are not supported with the FILES DD statement.
- Ensure that the DMSAR proc is in a proclib accessible during an MVS START command; for example, SYS1.PROCLIB.
- Users with security packages should note that since the CA Disk auto-restore started task cannot be identified by a JOB statement or a LOGON procedure, most security packages have a Started Procedures Table to associate a started task name with a user ID and possible group name. If your security package protects resources accessed by started tasks and you would like to auto-restore data sets protected by your security package, you must create an entry in your Started Procedures Table to associate the CA Disk auto-restore started task name DMSAR with a user ID and possible group name.
- For RACF, the Started Procedures Table is documented in the IBM manual RACF Installation Reference and in the IBM manual Systems Programming Library: RACF.

- Now that you have the DMSAR started task identified to your security package, anyone can auto-restore user data sets in the CA Disk archives. Although this causes no security exposure (after an auto-restore completes, OPEN, SCRATCH, RENAME or IDCAMS processing prevents user access to data sets to which they are not authorized), two methods have been provided which control this effect.

Front-End Method

This method is always in effect and works well if your security package does not use strict volume rules (for example, a user might have access to a data set on one volume, but does not have access to it on another volume), and you do not make extensive use of discrete RACF profiles.

CA Disk auto-restore processing uses the System Authorization Facility (SAF) to determine if the user has READ access to the data set name in question. Unfortunately, the auto-restore front-end processing does not have actual volume or discrete RACF profile information available. If queried from the catalog management hook, CA Disk passes to SAF the CA Disk pseudo- volume volser along with the data set name. If queried from the S213 exit, CA Disk passes to SAF the volume on which the OPEN was attempted along with the data set name. CA Disk allows the auto-restore to proceed only if it receives a return code of less than 8 from SAF.

If you are an CA ACF2 user, you must activate your SAF interface to let the SAF call issued by CA Disk work properly. Also, you must add the following entry to your SAFPROT GSO list:

```
SAFPROT.DMS CLASSES(DATA SET) CNTLPTS(ADSAR010)
SUBSYS(DMS)
```

Back-End Method

If you use the CA Disk RACF security interface, or if you use the CA Disk/CA Top Secret Security interface and have implemented the facility activated by sysparm ARSECURE, you can use this additional method.

Note: For more information, see the section Installing the CA Top Secret Security Interface in the chapter "Installation".

To install this method, specify sysparm ARSECURE with a value of R or Y. For more information, see the sysparm description for ARSECURE*n* in the *Systems Guide*.

2. See the documentation for sysparms ARESUNIT, ARESUNIC, and ARESUNI*n* (if multiple types of cartridge devices are used). If the defaults are unacceptable, specify their values in your production sysparms.
3. See the documentation for sysparm ARTAPEOK. This sysparm controls the degree to which tape mounts will be accepted--if at all--from auto-restore tasks.
4. See the documentation for user exit ARESPREX. This user screening exit can selectively reject auto-restore requests, and can also perform logging activities.

5. If you are installing the catalog management hook and the DFHSM SVC hook, you can use pool support as well. Pool support allows you to dynamically decide the volume to which to restore a data set, based on a combination of pool definitions and an optional user exit. See the Setting Up DASD Pools for Auto-Restore section. This does not have to done at installation time. You can install the basic auto-restore facilities first and add pool support later.

6. The catalog management hook and DFHSM SVC hook of CA Disk are installed with a started task. A system IPL is not required. You install the hooks by issuing the following command from the operator's console:

   ```
   S DMSAR,DMSAR=INSTALL
   ```

   Output is directed back to the operator's console. When the interface is successfully installed, this started task will end.

   Implementation Notes:

   - DMSAR will attempt to install and identify several modules in CSA and will use the CSVDYLPA function to identify those modules. If your installation has restricted access to the CSVDYLPA.ADD.*modname* and CSVDYLPA.DELETE.*modname* resources, it may be necessary to define multiple FACILITY class resources in your security system and grant access to those resources for DMSAR to execute properly. The resource names are of the format CSVDYLPA.ADD.*modname* and CSVDYLPA.DELETE.*modname*, where *modname* will be:

     - ADSAR010
     - ADSHS001
     - ADSAR026
     - DIMCH400

     If access to these resources is not granted, the DMSAR INSTALL fails with a DIM427I message. For more information, see Granting or Denying Access to System-Installed Intercept Modules.

   - You must issue this command on each CPU for which you want CA Disk's auto-restore function implemented.
   - The CA Disk catalog management hook and DFHSM SVC hook are removed each time you IPL. Therefore, you must use the above procedure to reinstall the hooks when you IPL. Consider installing the command in the automatic start-up procedures.

- If you run CA Allocate, issue the start command for DMSAR before the start command for CA Allocate.

  You can check if the CA Disk catalog management hook and DFHSM SVC hook are installed by issuing the operator command:

  ```
  S DMSAR,DMSAR=STATUS
  ```

  or simply:

  ```
  S DMSAR
  ```

  You can remove the CA Disk catalog management hook and DFHSM SVC hook by issuing the operator command:

  ```
  S DMSAR,DMSAR=REMOVE
  ```

  If you do issue this removal command, there are several processing steps to be aware of:

  After a remove command is accepted, the SVC table entries for SVC 26 and the DFHSM SVC are modified to point to their original addresses. This disables the hooks. The removal task waits eight seconds so that current requests can complete, then DMSAR deletes its hooks from CSA memory. The hooks are completely removed and the started task terminates.

- If you plan to use a CA Disk pseudo-volume name other than ARCIVE, it must be specified through sysparm RECATVOL. Do not select a pseudo-volume name that is a real volume name, or likely to become one. If your pseudo-volume name is a real volume, you cannot reference your data sets on that volume through the catalog. Set the pseudo-volume name as the value for the sysparm RECATVOL.

- For more information, see the sysparm description for PSUDEVTP*n* in the *Systems Guide*.

- An auto-restore diagnostic facility exists within both the Catalog management hook and the DFHSM SVC hook which can display applicable data used to determine if an auto restore is to be performed. These diagnostic messages are available if you have installed the Test Auto-Restore Interface as described next, and if your job name is equal to the test jobname (default is TESTDISK). If you want to change this default jobname, simply zap the jobname you prefer. For more information, see the TSTHOOKS member in the INSTALL library and installation instructions for the Test Auto-Restore Hook.

  Similar diagnostic facilities are provided for the CA Disk SVC module and VSAM date stamp module. The default jobname for the other diagnostic facilities is DMSTEST, so you can make this the same name for all functions. The ability to change the job names for the other functions is discussed where applicable.

Test Auto-Restore Hook Interface

This function lets you install a new or different release of the Auto-Restore operating system hook simultaneously with your production Auto-Restore hook. The Test Auto-Restore Hook only operates on jobs/tasks that have a certain predefined name or DD statement. It also allows you to reproduce problems and receive diagnostics using sysparm ARDIAGNM. Your production Auto-Restore hook continues to operate on all other jobs. This lets you safely test changes or upgrades to your Auto-Restore function while you leave your production CA Disk system running.

How It Works

CA Disk Auto-Restore Hook has program logic at their entry points which determines if each instance of the hook is to run or not. This determination can be based on either a specific job/task name or DD statement. In the case of the Test Auto-Restore Hook, care has been taken to ensure that the production hook runs only when the test hook does not run, and vice versa.

The Test Auto-Restore Hook only operates on jobs/tasks, which have a certain predefined name as their job/task name, or on jobs, which include a certain pre-defined DD statement name. This is configurable using the following system parameters:

ARJOBEXC

The value of this system parameter is the pre-defined job name. It can also be a task name (that is, a userid). For detailed information, see sysparm description for ARJOBEXC _jobname in the *Systems Guide*.

ARDDNEXC

The value of this system parameter is the pre-defined DD statement. For detailed information, see the sysparm description for ARDDNEXC_ddname in the *Systems Guide*.

Components of the Test Auto-Restore Hook Interface

The Test Auto-Restore Hook Interface consists of the following components:

TSTHOOKS member in the INSTALL library (Hook customization job)

This job creates the Test Auto-Restore Hook. This job copies the production hook modules to test names and then ZAPs them so that they use the test job name and test Auto-Restore PROC. Run this job even if you just want to use the default options.

TSTAR member in the PROC library (Test Auto-Restore PROC)

This PROC is the same as the DMSAR PROC, except that Test Auto-Restore has modified it for use. This PROC is started by the Test Auto-Restore hook in order to perform a test Auto-Restore.

TSTEXMPT member in the INSTALL library (Hook exempt job)

This job refreshes the production Auto-Restore Hook to begin excluding the job/task name , the DD statement name or both based on the value specified for sysparms ARDDNEXC and ARJOBEXC.

TSTINSTL member in the INSTALL library (Hook install job)

This job dynamically installs the Test Auto-Restore Hook and begins to accept auto-restores from the job/task name, DD statement name, or both based on the value specified for sysparms ARDDNEXC and ARJOBEXC. You do not need to IPL the system to install this hook.

TSTSTATS member in the INSTALL library (Hook status display job)

This job displays the status of the CA Disk Auto-Restore operating system hook on the system console.

TSTREMOV member in the INSTALL library (Hook remove job)

This job dynamically removes the Test Auto-Restore Hook from your system. You do not need to IPL your system to remove this hook.

SVC 26 Hook Module

Module TSTAR010 (a modified version of module ADSAR010).

SVC 109 Hook Module

Module TSTHS001 (a modified version of module ADSHS001).

Preparing to Use the Test Auto-Restore Hook Interface

There are several things which you must do before you use the Test Auto-Restore Hook feature.

To prepare to use the Test Auto-Restore Hook feature

1. Tailor the following jobs provided in the INSTALL or PROC libraries on your CA Disk distribution tape: TSTHOOKS, TSTINSTL, TSTREMOV, TSTSTATS, TSTEXMPT, and TSTAR. You need to tailor the job accounting, load library name (DD=STEPLIB), and the parameter library name (DD=PARMLIB).
2. Tailor the Q= parameter in the TSTAR PROC. This parameter provides a data set name prefix for the CA Disk libraries in your system.

3. You must select the job name, the DD statement, or both you want to have trigger the Test Auto-Restore feature. This is accomplished through the ARJOBEXC and ARDDNEXC system parameters. After selecting a job name, DD statement or both, supply these sysparms to the TSTEXMPT member in the form on a sysparm override as shown in the following illustration:

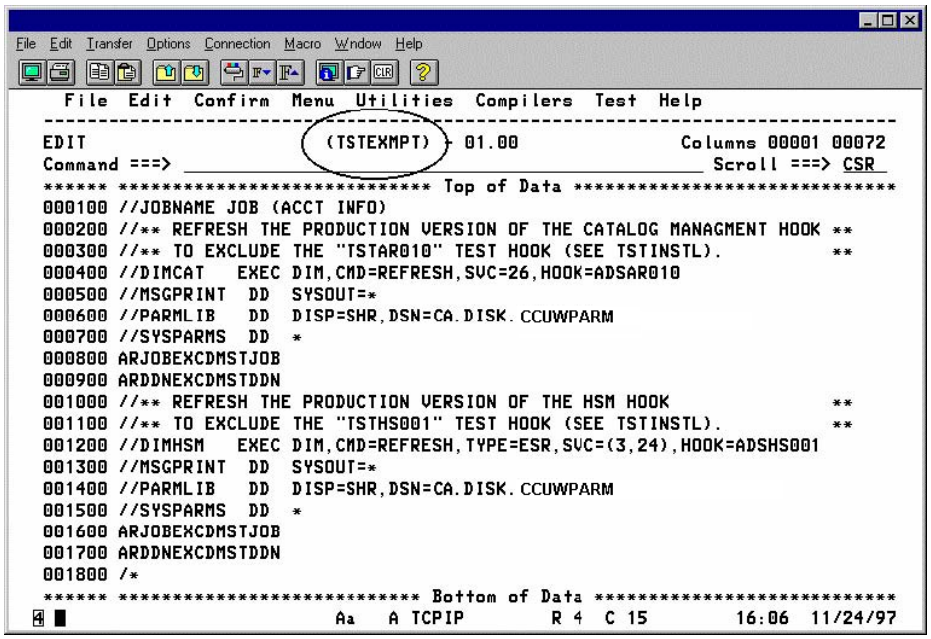

When this job is executed, it refreshes the production Auto-Restore Hook to begin excluding the job/task name , the DD statement name or both based on the value specified for sysparms ARDDNEXC and ARJOBEXC.

Important! The same values you specify for sysparms ARJOBEXC and ARDDNEXC in the TSTEXMPT job, must also be specified in the TSTINSTL job. Otherwise, unpredictable results occur.

Installing the Test Auto-Restore Hook Interface

The Test Auto-Restore Hook Interface can be installed by carefully following the following step by step instructions.

To install the Test Auto-Restore Hook Interface

1. Locate and submit the TSTHOOKS member you customized. This job creates the Test Auto-Restore environment by copying your production Auto-Restore modules to test names.
2. Locate and submit the TSTEXMPT member you customized. This job refreshes the production Auto-Restore Hook to begin excluding the job/task name , the DD statement or both name based on the value specified for sysparms ARDDNEXC and ARJOBEXC.

3. Locate and submit the TSTINSTL member you customized. This job dynamically installs the Test Auto-Restore Hook and begins to accept auto-restores from the job/task name , DD statement name or both based on the value specified for sysparms ARDDNEXC and ARJOBEXC.

Operating the Test Auto-Restore Hook Interface

One of the objectives of this interface is to allow you to activate an Auto-Restore environment at a release level different from that of your production system. The Test Auto-Restore system you just installed only operates on the job/task name , DD statement name or both that match the values you set for ARDJOBEXC and ARDDNEXC sysparms.

For example, with sysparm ARDDNEXC set to a value of TESTAR, the Test Auto-Restore Hook Interface is invoked for all jobs displaying the TESTAR DD statement as shown in the following illustration:

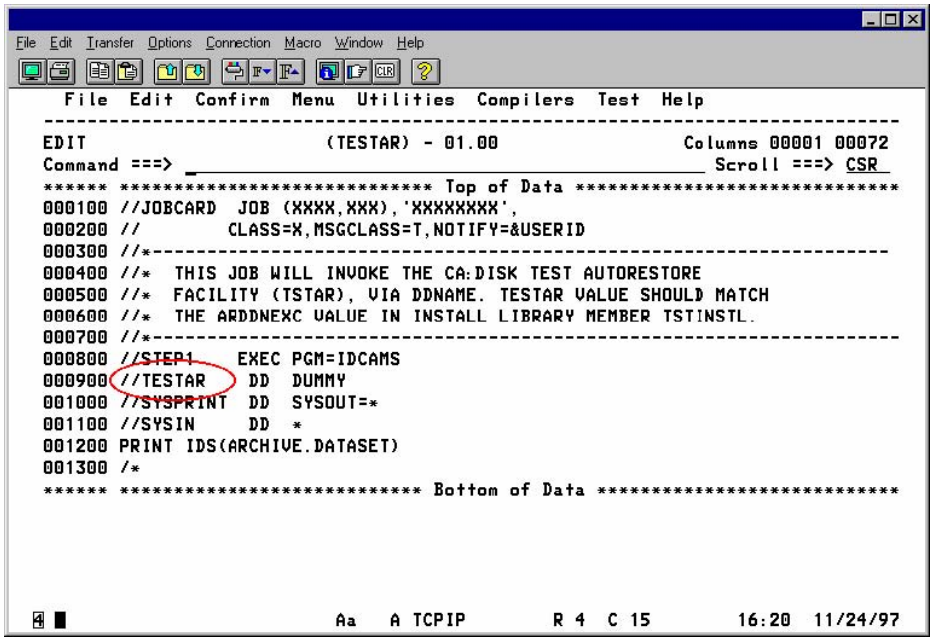

Note: The ARDDNEXC only operates successfully on archived data sets that are referenced dynamically.

The convenient quality of this environment is that you can now thoroughly test this new Auto-Restore system without impacting your production system.

Sysparm ARDIAGNM

There is a diagnostic feature in this interface that can help you solve problems related to Auto-Restore. To activate this feature, simply assign a job name value to sysparm ARDIAGNM. Setting this sysparm instructs the Test Auto-Restore Interface to issue diagnostic messages to each job that enters the system with this job name.

The following illustration shows some of the typical messages routed to SYSLOG when the ARDIAGNM sysparm is activated:

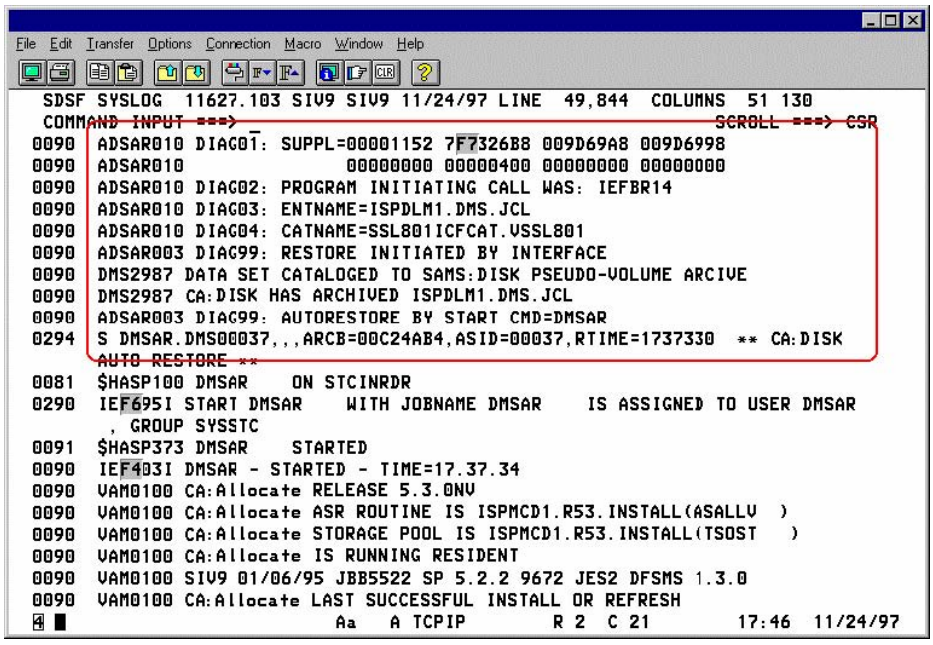

Note: For detailed information, see the sysparm description ARDIAGNM_jobname in the *Systems Guide*.

Member TSTREMOV

After your testing is complete, shutdown the Test Auto-Restore Hook Interface. Doing so reverses what was excluded from your production Auto-Restore while the Test Auto-Restore Hook were active. The Test Auto-Restore Hook Interface can be removed by carefully performing the following step-by-step instructions.

To remove the Test Auto-Restore Hook Interface

1. Locate and submit the TSTREMOV member you customized. This job dynamically removes the Test Auto-Restore Hook.
2. Issue the REFRESH command to your production Auto-Restore system. This command sets sysparms ARJOBEXC and ARDDNEXC back to their default values, which deactivates them. The following is an example of the REFRESH command:

   ```
   S DMSAR,DMSAR=REFRESH
   ```

Setting Up DASD Pools for Auto-Restore

This section applies only if you have installed—or plan to install—the catalog management hook and DFHSM SVC hook. Using the pool support is completely optional.

The purpose of the pool support in the Auto-Restore facility is to allow the maximum flexibility possible in determining the target volume to which the data set is restored. Although this capability is not essential to the basic function of Auto-Restore, it does relieve the DASD administrator of a lot of the manual effort required to manage free space distribution among packs. Perhaps the best way to explain its advantages is to show what management is like without it.

Without pool support, any data set that is Auto-Restored is required to go back to the volume from which it came. Two problems become obvious:

- What happens if the volume no longer exists?
- What if there is not enough space for the data set to fit or the available space is to fragmented?

The answer is that the restore fails and the task that requested the restore fails. This is not too serious if a TSO/ISPF user caused the restore, but it can cause problems if a production job abends because it could not gain access to the data set.

Pool support within the Auto-Restore facility is accomplished separately for VSAM and non-VSAM data sets.

For VSAM Data sets:

For VSAM data sets, pooling is technically complex, and must be accomplished by CA Allocate, CA's Volume Allocation Manager, or some other allocation control product. It is not done within CA Disk itself. If CA Allocate or some other allocation control product is not used, Auto-Restore must place VSAM data sets back on the volumes from which they were archived. If you do not have an allocation product available, consider using RESPRIEX—Screen Restore Requests. See the *Systems Guide*.

For Non-VSAM Data sets:

For non-VSAM data sets, pooling can be accomplished through either or both of two methods. The first method is by using CA Allocate or some other allocation control product. The second method is accomplished within CA Disk Auto-Restore facility, and can be used alone or in conjunction with CA Allocate.

Several advantages become apparent when volume pooling is implemented, such as:

- The data set can be restored to any volume in the pool that has sufficient free space. Each volume in the pool is searched until one with sufficient free space is found, or until it is determined that none of the volumes in the pool has enough free space.
- If a volume (or a string of volumes) becomes unavailable, specify an entry in the volume pool that redirects all restores to that volume to a different pool of volumes.
- You can reserve a set of special volumes used only for Auto-Restores as its own pool. This allows you the capability of managing this pool differently than perhaps your TSO or production packs.
- You can code your own user exit to decide the pool to which to restore, based upon criteria unique to your installation.

Pool support can be thought of as the algorithm used to determine the destination volume for a data set being Auto-Restored. It consists of several possible steps, depending on options you specify. The following is a schematic of the process.

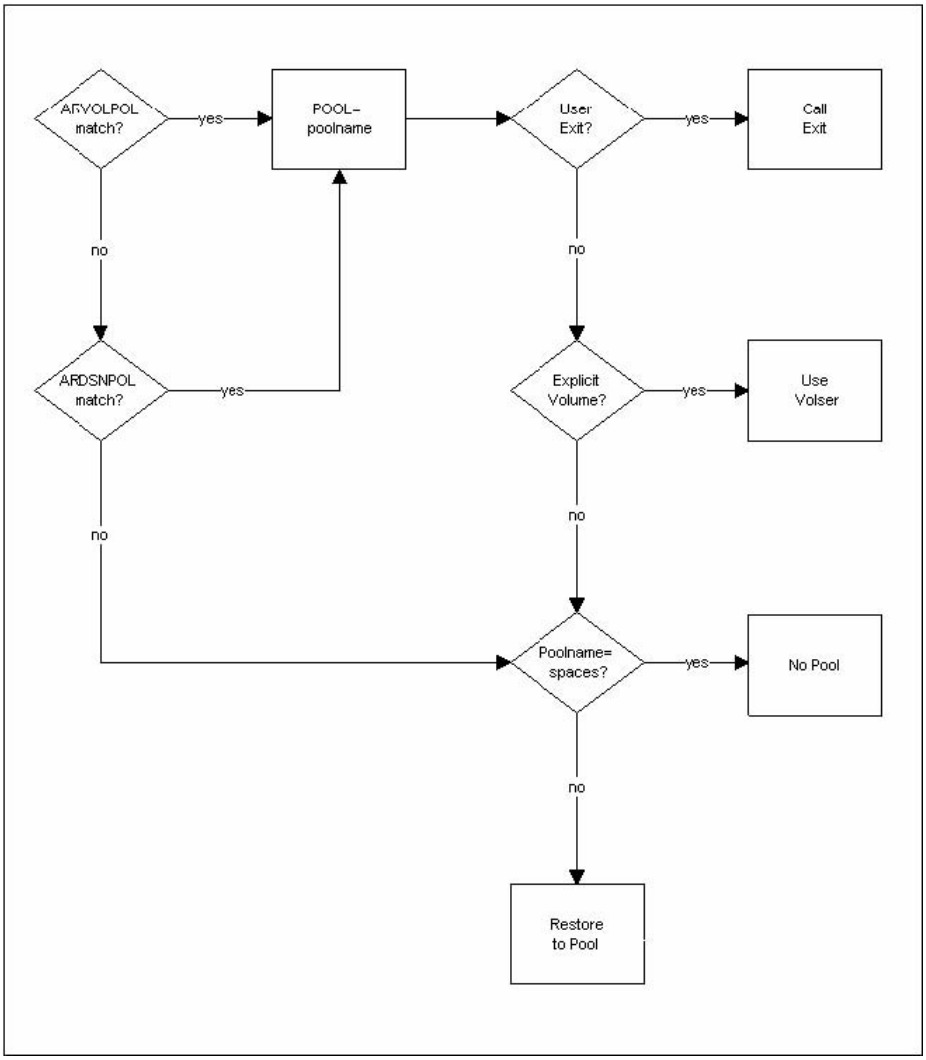

The following illustrations illustrate typical examples of how the syntax works for DASD Pools. These illustrations are then followed by a series of explanatory notes to help describe the illustration.

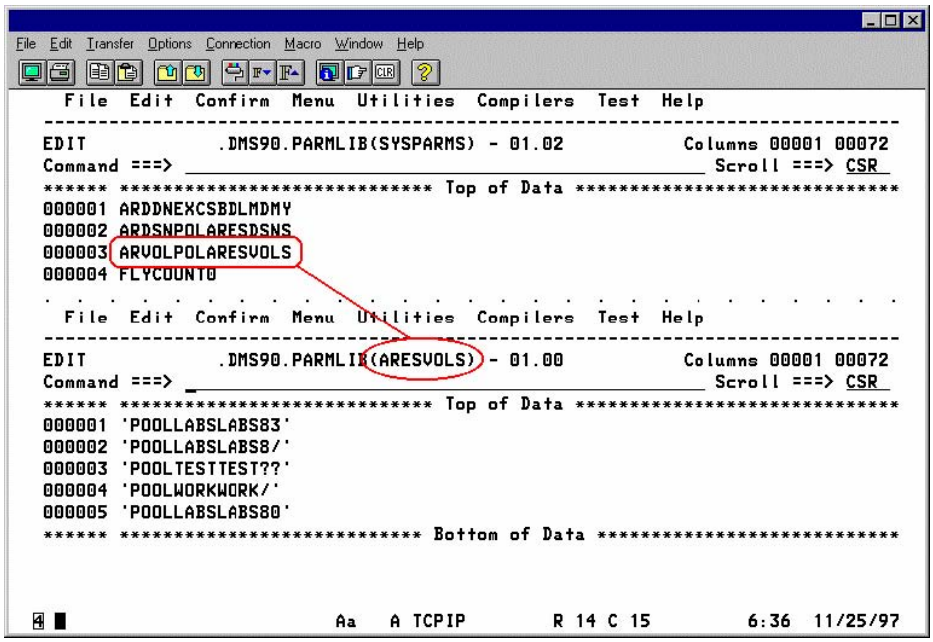

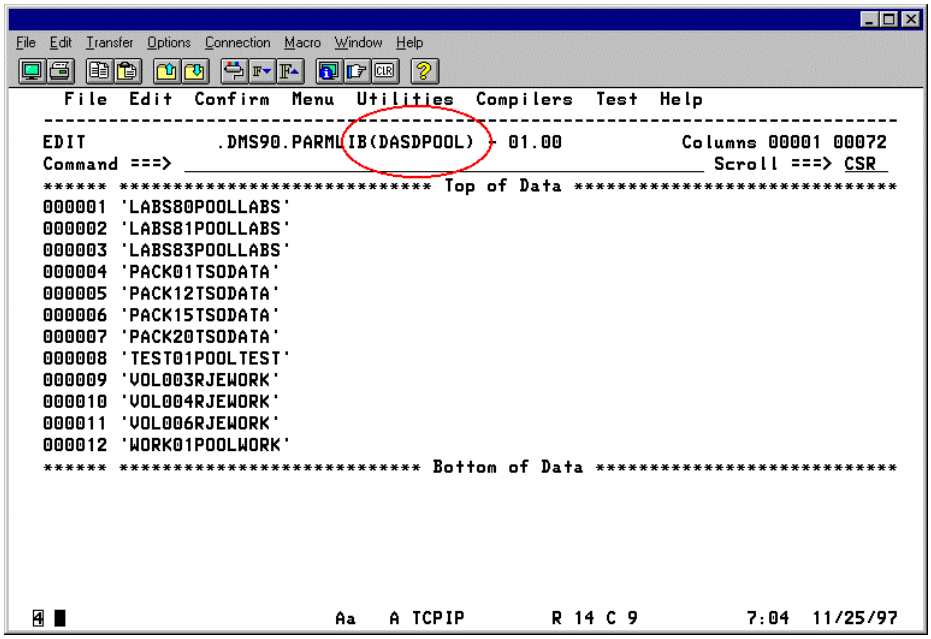

Notes for setting up DASD pool support for DMSAR:

- All pools referred to in ARVOLPOL and ARDSNPOL must be defined in the member DASDPOOL in your parmlib data set. Extensive rules for coding this member are found in the DASD Pool List in the *Systems Guide*.
- All standard CA Disk pattern-matching capabilities are supported in both ARVOLPOL and ARDSNPOL members. The first entry that causes a match causes its corresponding pool name to be used, even if there are other entries in the member that can match. This means start each of these members with the most specific entries first and have the least specific entries entered last.
- CA Allocate or another allocation control product supports Pool capability for VSAM data sets, not by the methods just described.

To use CA Allocate or another allocation control product for non-VSAM data set pooling, create an ARVOLPOL with at least one volume in it (a volume that is normally empty or used for temporary data sets only). This is for two reasons:

- If the original volume is offline or does not have enough space, CA Disk fails the restore without attempting to allocate the data set. No allocation control product gets control. By supplying an ARVOLPOL, you provide CA Disk with an online volume to which it can attempt to make allocations.
- An ARVOLPOL gives you a means of seeing how well your allocation control product rules are doing. By monitoring the volumes defined in this ARVOLPOL, you can determine which data sets are not being controlled by your rules. You can update your rules to prevent similar problems from occurring later.

Going back to the schematic, we start at Step (a) to determine if a source volume pool match is found. There are actually two steps to this process, which also applies to Step (b). First CA Disk retrieves sysparm ARVOLPOL. If you have not specified this sysparm, a no-match condition is assumed. If you do specify the sysparm, the value you supply for it must be the name of the member in your parmlib that contains the pool entries. For instance, if a value of ARVOLPOLVOLPOOL is specified, you must code your entries for this pool in the member VOLPOOL in your parmlib. Assuming this pool is defined, CA Disk compares the data set's original source volume to each explicit or pattern volume in the table. If a match is found, the pool name associated with that entry is used and processing continues to Step (d). If no match is found, or if sysparm ARVOLPOL is not specified, processing goes on to Step (b).

Step (b) follows the same general logic as Step (a). You notify CA Disk of the presence of a data set name pool by specifying sysparm ARDSNPOL with the member name that contains the pool associations. In this pool, CA Disk compares each entry to the data set name being restored. As soon as an explicit or pattern name is matched, that pool name is used and processing continues with Step (d). If no match is found, the pool name is initialized to spaces (Step (c)), indicating that no pool name is used in the restore.

We are now at Step (e), which is an optional user exit that can be coded. If you want to write a special pool selection routine, this is where you add it. CA Disk calls the exit (Step (f)) when sysparm ARPOOLEX is specified with the module name to be called. For more information, see the User Exit description for ARPOOLEX—Auto-Restore DASD Pool Exit in the *Systems Guide*. With this exit, you can see if any pool was selected, based on the entries in the ARDSNPOL and ARVOLPOL members. If no match was found or neither of the members was defined, the pool name is blanks (Step (i)); otherwise, it contains the name of the pool that is used (Step (k)). You can change the pool name to another valid pool name or it can be blanked out to suppress pool support for the restore. You also have the option of passing back an explicit target volume (Step (g)), which overrides any pool specification (Step (h)). The illustration shows that there is no requirement to use the pool support that CA Disk provides—just by coding the user exit you can customize pool support.

If, after the processes above, you exit with the name of a pool, this pool must be defined in the member DASDPOOL in your parmlib. The eligible volumes to which the data set can be restored are all the entries in this member that have the same pool name.

At this point, read the Restore/Recover section in the *User Guide*. A detailed description of the rules CA Disk follows in determining the volume to which a data set is restored.

Customizing the TSO/ISPF Auto-Restore Environment

This section applies only if you have installed the catalog management hook. If you have installed the hook but have not attempted any restores from a TSO/ISPF environment, do so at this time. It is much easier to conceptualize the modifications about to be described if you are familiar with how the support works in its default mode.

In its default mode, the TSO/ISPF user is prompted for answers to the following questions:

- DO YOU WANT TO RESTORE THE DATA SET? (yes/no) If the user replies anything other than *no*, the next question is asked.
- DO YOU WANT AN IMMEDIATE OR A DEFERRED RESTORE? (I/D) If the user responds anything other than *deferred*, the next question is asked.

- DO YOU WANT TO WAIT FOR THE RESTORE TO COMPLETE? (yes/no) If the user answers *no*, the terminal will be unlocked and the user is notified when the restore is complete. If the user responds anything other than *no*, the terminal remains locked until the restore is complete, unless a tape mount is required (and allowed by sysparm ARTAPEOK). If a tape mount is required, the user is prompted one more time with the following question:
  - DATA SET IS ARCHIVED TO TAPE. DO YOU STILL WANT TO WAIT? (no/yes) If the user responds *yes*, the terminal is remain locked until the restore is complete. If the user responds anything other than *yes*, the terminal will unlock and the user will be notified when the restore is complete. The Auto-Restore task begins before the user responds to this prompt, and the response has no effect on the progress of the Auto-Restore job.
  - A WAIT RESPONSE WILL SUSPEND YOUR SESSION. DO YOU WANT TO WAIT (Y?N)? This message allows the TSO user another option in the wait process in instances where another task is restoring the data set.

This series of questions can be very confusing to some users. To suppress any or all of the questions, you can. Installing user exit *USERMOD8* by performing the following procedure does this.

To install user exit USERMOD8

1. Locate the source for ADSUMOD8 in the library associated with the CCUWSAMP DDDEF. The following illustration is a sample source for ADSUMOD8:

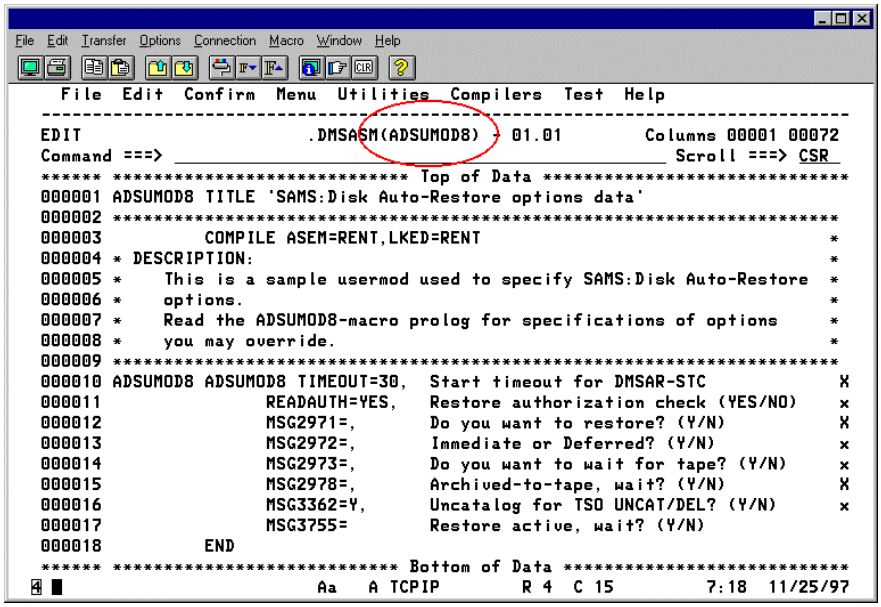

2. In order to ensure that the changes you make to ADSUMOD8 are protected during future CA Disk installs or maintenance, copy this member into the source library associated with the //USERASM dd statement in USERMOD8.

3. Customize ADSUMOD8 as follows:

   a. The TIMEOUT= parameter defaults to 30 seconds before the STC passes control back to the USERID. After this wait period is exceeded, message 2977 is issued and the TSO user (or console operator in the case of a batch job) is asked if the Auto-Restore should be terminated. A *Y* response terminates the Auto-Restore. Any other response extends the wait period an additional 30 seconds. If your installation is prone to resource conflicts that slow task initiation, consider increasing this wait period and eliminate this additional prompt. Increasing the value of this parameter can do this.

   b. The READAUTH= parameter defaults to YES forcing DMSAR to invoke a SAF check for Read authority based on the USERID invoking the Auto-Restore prior to starting the DMSAR STC without regard to the setting of sysparm ARSECURE.

   c. Specifying NO bypasses this security prior to starting the DMSAR STC. In CA ACF2, or depending upon the setting of ARSECURE sysparm in a RACF or CA Top Secret Security environment, the data set could be restored when the USERID requesting the restore had no read authority to the data set. Although this causes no real security exposure, it does waste resources. It is not recommended that customers modify this value in the ADSUMOD8 source unless they encounter a specific security situation where they need this ability.

   d. For more details of security protection during Restore for all security interfaces, see the section Security Processing in the *Systems Guide*.

   e. The MSG*nnnn*= parameters are used to customize the default responses to the TSO prompts. Specify *Y* to pre-answer the TSO response as YES. Specify a *N* to pre-answer the TSO response as NO. Specifying the *null parameter* causes auto-restore to prompt the TSO user for a response. For MSG=2972, specify an *I* to pre-answer an immediate restore, a *D* to pre-answer a deferred restore, or the *null parameter* causes auto-restore to prompt the TSO user for a response.

4. Save your work by issuing SAVE at the TSO command line.

USERMOD8

USERMOD8 must be installed using SMP/E. Sample JCL is provided for you in member USERMOD8 of the install library. Install USERMOD8 in the following manner:

1. Locate the source for USERMOD8 in the library associated with the CCUWSAMP DDDEF. The following illustration is sample SMP/E JCL for USERMOD8:

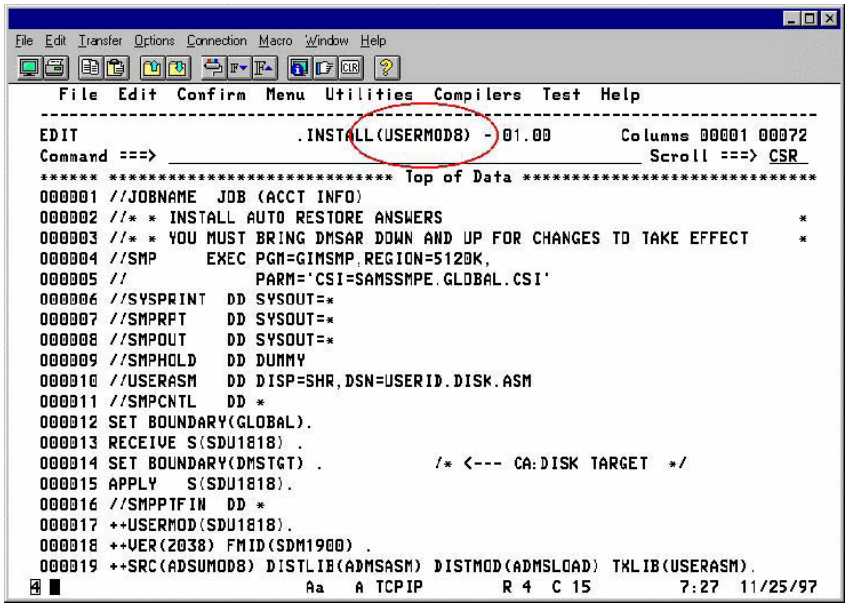

2. Customize USERMOD8 by supplying the appropriate information for the following:
   - jobcard information
   - Global and Target CSI names
   - USERASM data set name
3. Submit USERMOD8
4. Refresh the Auto-Restore hook on all systems for this change to take effect. Issue the following command to accomplish this:

   ```
   S DMSAR,DMSAR=REFRESH
   ```

Auto-Restore Implementation Guidelines

The hardest task in implementing the Auto-Restore function is in determining how much time is allowed to elapse between the time a data set is last referenced and when it becomes a candidate for archival. If this time window is too large, few archives take place and a severe shortage of DASD space probably occurs. On the other hand, if the time window is too small, too many archives take place. If this happens, the cost of archiving and restoring the data sets probably exceed any benefits realized by the temporary DASD space savings.

Obviously there are too many variables to determine an absolute rule in calculating the best time to archive a data set based on last use date. The best method appears to be by trial and error. First, pick a period of time that seems reasonable based on current knowledge of the user environment. Then set up a monitoring period and keep totals on the number of auto- restore requests being processed on a per-shift basis. If this number stays sufficiently low during the trial evaluation, and no shortage of DASD space is evidenced during this period, then you probably have struck upon a good balance between archiving and restoring. Your evaluation period should run through your center's month-end processing cycle to verify that this load does not produce an unreasonable number of restore requests. If it does, your archival scheme probably needs to be re-evaluated.

One issue that must be considered on the archival side is the increased catalog usage due to the need to recatalog data set to the CA Disk pseudo-volume so they can be automatically restored.

One facility that can help in the maintenance of your catalogs is the specification of sysparm UNCATPSU with a value of Y. This causes both MERGE and IXMAINT to issue an uncatalog (if the data set is cataloged and does not have a F1-DSCB on the source volume) when the last DSNINDEX record for the data set is deleted from the archive index. We highly recommended that this sysparm be specified with a Y if you are using the Auto-Restore function.

The established CA Disk customer faces the problem that all data sets that were archived under an old system of *scratch, uncatalog* are not accessible by the Auto-Restore function. To assist these customers, a procedure has been implemented to catalog data sets that are found in the archive index but that are no longer cataloged and do not have a F1-DSCB on the source DASD volume. This procedure, named IXCATLG, is supplied with CA Disk and is documented in the IXCATLG Utility section.

VSAM Alternate Indexes

When a VSAM cluster is archived and recataloged to the CA Disk pseudo-volume, alias entries are also defined for each of the alternate index and path names, as well as for each of the components' names for the cluster. This enables the auto- restore function to be invoked when you reference the base cluster, any of its components, an alternate index, or path name associated with the base cluster.

Restore processing removes the catalog entry to the pseudo- volume and all of the aliases associated with it, and redefines the correct VSAM catalog entries. Should the restore process fail, CA Disk attempts to re-create the initial status by recataloging the cluster to the pseudo-volume and associating the proper alias names to it.

Restrictions

If the catalog management hook and DFHSM SVC hook are installed, a reference to an uncataloged archived data set in a batch job with volume and unit information hard coded in the JCL fails with a JCL error.

Auto-Restore Delayed by Merge

Use SYSPARM MERRELTM*nn*, where DMSAR can be delayed by a Merge job. Specify Merge release time MERRELTM of 01 to 99 minutes (default is 00) to be used in both Merge and DMSAR. During Merge, if MERRELTM is non-zero and there is room on the current volume for the next Archvol, the time interval since the last output volume swap is compared to MERRELTM. If that time interval is greater than or equal to MERRELTM, a test ENQ for DMSWVOLS is issued for that volume. If DMSAR has been waiting with a DMSWVOLS ENQ for that volume, the volume is released by Merge. DMSAR only issues the DMSWVOLS ENQ if MERRELTM is non-zero; the actual value of MERRELTM is picked up in Merge. Tuning can be done within Merge with MERRELTM from 01 to 99 minutes, depending on individual needs.

Propagating Enqueues

In multiple CPU environments with cross-system enqueue packages, you must ensure that all enqueues are propagated across all CPUs. For auto-restores in particular, enqueues for DMSAUT0, DMSAUT1, DMSAUT2, DMSAUT3, and DMSAUT4 must be propagated accurately. These enqueues ensure that data integrity is maintained.

Note: Failure to propagate these enqueues can lead to data loss problems. For example, jobs executing in a multi-system JES complex can corrupt a data set being auto-restored on another system in the same JES complex when the DMSAUTx enqueues are not propagated correctly.

Precautions in Changing Immediate to Deferred Requests

Sysparm ARTAPEOK (with a value of D) and user exit ARESPREX (by setting the return code to H'1') both provide the ability to change an immediate restore request into a deferred request; that is, place the Auto-Restore request in the restore queue rather than process it immediately. The sysparm option automatically defers a request only if it is for a TSO user and a tape mount is required. Batch jobs and restores from disk archives are not affected. The user exit is more general, however, and can defer any of the various request types.

You should consider that there is a slight exposure before using either of them. As you can see from the following description, the exposure is extremely low under most circumstances, especially with the sysparm option, but indiscriminate use of the user exit option can cause a problem.

The concern exists only if two jobs , TSO users or both happen to want the exact same archived data set at the *same* time. CA Disk Auto-Restore processing for the second request detects that another Auto-Restore task (serving the first request) is already processing the same data set. The second job is then placed in a wait state (if it's batch), or informed that the restore is in progress, and to try again later (if it's a TSO user).

If the sysparm option is used and the first request is from TSO and the second one from batch, the batch job will fail because of its assumption that the Auto-Restore for the TSO user was actually restoring the data set, when in fact it just placed it in the queue.

If the user exit is used to force a request to be queued instead of restoring it immediately, the results for the second job of two simultaneous requests can be as just described, without regard to the TSO or batch status of either requester.

To summarize, the exposure exists only if two jobs/users attempt to access the same archived data set at the same time. This probability is very low. If the sysparm option is used, the probability is further reduced by the fact that the first job must be a TSO request, and the second a batch job.

Rejecting Auto-Restore Requests Through Exclusion Tables

The same user exit mentioned above, ARESPREX, can also be used to reject an Auto-Restore request. The exit has access to considerable data regarding the restore itself, and therefore provides a great deal of flexibility in the decisions that can be made, and consequently should be considered first if you have needs in this area.

Another method also exists providing limited flexibility, and requires the assembly and linking of exclusion tables into a CA Disk module. ADSAR003 is responsible for starting the DMSAR Started Task. This module is statically linked into ADSAR010, the catalog management hook. Before ADSAR003 even starts the restore task, it has the ability to screen the request and reject it, based upon either the dsname that is needed, the name of the executing program, the jobname, or the TSO userid. These exclusions can be provided in several CSECTS statically linked into ADSAR010 and available to ADSAR003.

To implement these modifications, first customize the assembly source for exemption table members DSNEXMPT, JOBEXMPT, PGMEXMPT, and TSOEXMPT, which are located in your CCUWSAMP library. The following source shows a sample PGMEXMPT; all other exemption tables follow the same format.

```
PGMEXMPT TITLE 'AUTO-RESTORE EXEMPTION BY PROGRAM NAME TABLE'
*********************************************************************
*         COMPILE ASEM=RENT,LKED=RENT
*
* DESCRIPTION:
*
*    PGM EXCLUSION/EXEMPTION TABLE
*
*********************************************************************
*
* SCRNLIST MACRO PARAMETER DESCRIPTIONS:
*
*    SCREEN=    SCREEN PARAMETER LIST. MULTIPLE VALUES CAN BE
*               CODED BETWEEN PARENTHESIS AND SEPARATED BY
*                       COMMAS. PARENTHESIS ARE NOT NEEDED WHEN ONLY ONE
*                       VALUE IS SPECIFIED.
*                       EXAMPLE:
*      SCREEN=(ADSMI/,ISR/,IDC/) IN THIS CASE 3 VALUES
*                       PASSED WITH WILDCARD CHARACTER / AT THE END OF
*               EACH VALUE.
*
*    LENGTH=    LENGTH OF EACH SCREEN ENTRY. THIS VALUE IS USED
*                        WHENEVER ANYTHING OTHER THAN TYPE=VAR IS
*                        SPECIFIED. THE DEFAULT VALUE IS 1.
*                        EXAMPLE:
*                        LENGTH=8 IN THIS CASE LENGTH IS 8 BYTES.
*
*    PLEASE NOTE THAT WHEN MULTIPLE LINES ARE NEEDED, A NON-
*    BLANK CHARACTER SHOULD BE CODED IN COLUMN 72 AS A
*    CONTINUATION INDICATOR FIELD, AND THE CONTINUATION LINE
*    SHOULD START IN  COLUMN 16.
*
*    READ THE SCRNLIST MACRO PROLOGUE FOR A COMPLETE LIST OF
*    AVAILABLE PARAMETERS.
*
*********************************************************************
         SPACE 3
PGMEXMPT SCRNLIST LENGTH=8,
               SCREEN=
*
         END
*************************** Bottom of Data ***************************
```

Once you have altered the exemption table members to your needs, submit USERMOD9 JCL to complete the customization (sample JCL may be found in the CA Disk INSTALL library). To implement the change, you must re-install the DMSAR hooks (Catalog Management and DFHSM SVC) running on your system by first issuing the REMOVE option, and then issuing the INSTALL option of DMSAR.

The following is a sample of USERMOD9:

```
//JOBNAME  JOB (ACCT INFO)
//*
//* ********************************************************************
//* *        INSTALL EXCLUSION TABLES FOR AUTO-RESTORE                 *
//* *                                                                  *
//* ********************************************************************
//SMP     EXEC PGM=GIMSMP,REGION=5120K,
//             PARM='CSI=CAISMPE.CSI'
//SYSPRINT  DD SYSOUT=*
//SMPRPT    DD SYSOUT=*
//SMPOUT    DD SYSOUT=*
//SMPHOLD   DD DUMMY
//USERASM   DD DISP=SHR,DSN=USERID.ASM
//SMPCNTL   DD *
SET BOUNDARY(GLOBAL).
RECEIVE  S(SDU1819).
SET BOUNDARY(CAIT0).
APPLY S(SDU1819) ASSEM .
//SMPPTFIN  DD *
++USERMOD(SDU1819).
++VER(Z038) FMID(CCUWC50) .
++SRC(PGMEXMPT) DISTLIB(ACUWSAMP) DISTMOD(ACUWMOD0) TXLIB(USERASM).
++SRC(DSNEXMPT) DISTLIB(ACUWSAMP) DISTMOD(ACUWMOD0) TXLIB(USERASM).
++SRC(JOBEXMPT) DISTLIB(ACUWSAMP) DISTMOD(ACUWMOD0) TXLIB(USERASM).
++SRC(TSOEXMPT) DISTLIB(ACUWSAMP) DISTMOD(ACUWMOD0) TXLIB(USERASM).
/*
```

Installation of IGGDASU2 User Exit

This user exit is provided for those who want to make use of SMS GDG's and do not have HSM. A problem can be encountered where LOGREC records are being cut when a SMS GDG rolls out of the sphere and is cataloged to the CA Disk pseudo volser ARCIVE. The problem is with scratch processing for a volume that does not exist. The User exit, IGGDASU2, is provided by IBM to keep logrec records from being cut.

If you plan to use a CA Disk pseudo-volume name other than ARCIVE, it must be specified through sysparm RECATVOL. Do not select a pseudo-volume name that is a real volume name, or likely to become one. If your pseudo-volume name is a real volume, you cannot reference your data sets on that volume through the catalog. Set the pseudo-volume name as the value for the sysparm RECATVOL. For more information, see the sysparm description for RECATVOLvvvvvv in the *Systems Guide*.

CA Disk provides the user exit ADSDASU2. This module gets control during DADSM scratch processing. It looks at the parameter list passed and if the volser is ARCIVE, the request is ignored, indicating to the caller that processing was successfully completed.

If this exit is needed, install it as part of the automatic processing during IPL. ADSDASIN is the installation program for this user exit. Specify PARM=I in the EXEC statement to dynamically install the exit. Use PARM=R to remove the exit for testing purposes or if problems are encountered. Run this exit as its own job in the IPL deck and execute it before DMSAR.

The following is sample JCL to install the user exit. A copy of this JCL is found in the INSTALL library member ADSDASIN.

```
   Menu  Utilities  Compilers  Help
 ------------------------------------------------------------------------------
 BROWSE                        CCUWJCL(ADSDASIN)     Line 00000000 Col 001 080
 Command ===>                                          Scroll ===> CSR
******************************** Top of Data ********************************
//JOBNAME  JOB (ACCT INFO)
//*
//*********************************************************************
//*                                                                   *
//*                  INSTALL ADSDASU2 SCRATCH EXIT                    *
//*                                                                   *
//*********************************************************************
//STEP1   EXEC PGM=ADSDASIN,PARM=I       (USE PARM=R TO REMOVE EXIT)
//STEPLIB   DD DISP=SHR,DSN=CAI.DISK.CCUWLOAD

//PARMLIB   DD DISP=SHR,DSN=CAI.DISK.CCUWPARM

//MSGPRINT  DD SYSOUT=*
//SYSUDUMP  DD SYSOUT=*
******************************** Bottom of Data ********************************
```

Customizing the CA Disk Tape Management Support

Tapes have traditionally been managed by the expiration dates written in the tape labels. Standard IBM support for date- protected data sets requires operator intervention to rewrite a protected tape before its expiration date.

Note: For more information, see Year 2000 Considerations in the *User Guide*.

Tape management systems usually extend this support in several ways. They interpret certain expiration dates as codes to indicate the type of control governing the use of the tape. A master control file, or *tape management catalog*, records this information, and controls access to each tape. If the control file indicates a tape as an available scratch tape, it allows it to be rewritten without operator intervention, regardless of the expiration date in the label.

For example, CA 1 uses an EXPDT=99000 to mean that a tape is eligible as a scratch tape when it becomes uncataloged. As long as the tape is cataloged, it is protected from being used as a scratch tape. This is commonly known as *catalog control*. An EXPDT=99365 is often used as the code to mean *permanently protected* -- the tape cannot be reused as a scratch tape unless the control file is updated to change this status. If you are using LDATE/ddd (to retain ddd days after the tape was last used) be aware that CA 1assigns an expiration date of 98ddd, which might be confusing since it, shows as 1998ddd in some of our reports. Consult the documentation for your tape management system to find the various means of control it provides through expiration date codes.

IBM also recognizes 99365 as a never expire date. This means that for standard applications, operator intervention is required to rewrite such tapes. However, IBM has also provided an exit to allow operator intervention to be avoided. CA Disk optionally makes use of this exit, based upon the value of sysparm TAPEFSCR.

CA Disk uses standard techniques to open and write tape data sets, which causes tape labels and their expiration dates to be created by normal means. As distributed, CA Disk dynamically allocates all tapes that are needed and assigns them an expiration date of Julian 99365 (from the default value of sysparm DYNEXPDT). This value guarantees the integrity of the data by ensuring that a tape never expires before the data that it contains. CA Disk determines when all of the data on each tape has expired, and at that point returns it as an eligible scratch tape.

CA Disk also provides an option to catalog each tape data set that it creates. By specifying expiration dates and catalog options correctly, CA Disk is fully compatible with all major tape management systems. As mentioned above, sysparm DYNEXPDT controls the expiration date for dynamically allocated tapes. If you supply DD statements for the output tapes, the expiration date you supply in the JCL is used.

A description of several implementation options relating to tape management in general is presented in the Tapepool Considerations section in the *Systems Guide*. CA Disk provides three methods for controlling tapes. Select one of these methods and follow the instructions for installing the tape management support.

Method 1 — Controlling Tapes Through the EDM

CA 1 provides an External Data Manager (EDM) interface to allow other system software products to manage their own tapes. CA Disk is one such system. It is through this interface that the greatest amount of control with the least amount of risk is available. See the CA 1 guides for a complete description of the External Data Manager Interface. The advantages of using this method to control CA Disk-created tapes over any other methods are as follows:

- Activation of the EDM interface for CA Disk is greatly simplified.
- Changes within CA 1 or CA Disk do not require the reinstallation of the interface.
- A single point of control is established to determine when a tape should be scratched.
- CA 1 does not attempt to prematurely scratch a CA Disk- owned tape.
- CA 1 does not allow other programs to overwrite a CA Disk- owned tape.
- If CA Disk creates a tape and then later abends, the tape are not scratched, as it normally would without the EDM facility.
- Future changes in CA Disk or CA 1 do not increase the likelihood of tapes being scratched prematurely.

When CA Disk is identified as an EDM to CA 1, any archive/backup tapes created by CA Disk will be managed by CA Disk (that is, CA Disk will inform CA 1 when to scratch the tape). CA 1 exempts these tapes from its normal processing. Any tapes created through sequential migrate--or UNLOAD to tape--processing is not considered externally managed, however, and is scratched by CA 1 based on the tape's expiration date.

A tape can be identified as externally managed by its expiration date (99365) and an indicator flag within the volume's TMC record. The actual expiration date of the tape is kept in the ARCHVOLS record within the FILES. Normally this is 1999365, if the default value of sysparm DYNEXPDT is used. If this is so, the date i scratched when the last data set on the tape expires. When CA Disk determines that all data o the tape has expired, it expires the tape through the EDM interface. Directly changing the expiration date for a tape must be done through the appropriate CA Disk commands, not through CA 1. Only when a tape is expired by CA Disk should the status within the TMC change.

The activation of EDM support is done in two parts: First, sysparm TMSCTLEX must be set to indicate to use the CA Disk EDM program as the tape management interface. Set this sysparm to ADSTH014. Second, CA Disk must be identified as the EDM within CA 1 This is accomplished by modifying the TMOEDMxx member of your CAI.PPOPTION data set. The easiest method of identifying CA Disk as an EDM is by program name. For example:

```
EDM=DISK,PGM=ADSMI002
EDM=DISK,PGM=ADSMI302
```

“DISK” is an example of an EDM name. Any name with 4 characters can be used.

ADSMI002 – The program that creates or expires tapes in all CA Disk processes .

ADSMI302 – For problem determination, Technical Support may ask you to add PET Tracing to your job (Program Event Tracing). So, this entry insures that the output tapes in the abended job will be EDM managed and will not expire prematurely.

Note: For more information on the External Data Manager Interface, see the CA 1 guides.

If you have previously created archive/backup tapes with CA Disk, you also need to change their status within the TMC to indicate that they are also externally managed. This action is required because CA Disk attempts to scratch the tapes through the EDM interface as soon as the sysparm TMSCTLEX is set.

Convert the CA 1 volume records currently in use by CA Disk to EDM management. Run a CA Disk LISTV to get a list of all the volsers currently in use. Then, set up the following utility to convert all of the fields for each volser as shown in this example setup:

```
//STEP1  EXEC PGM=TMSUPDTE
//TMSRPT DD SYSOUT=*
//SYSIN  DD *
VOL 123456,NODSN,NOCHAIN
REP FLAG3=20
REP EDMID=DISK
REP VOLSEQ=1
REP 1STVOL=HEXZEROS
REP NEXTVOL=HEXZEROS
REP PREVVOL=HEXZEROS
REP EXPDT=PERM
```

“EDMID=xxxx” is the value specified for CA Disk in the TMOEDM00 member.

Note: If the conversion of existing tapes is not done properly, CA 1 begins issuing TMSTVEXT-08 messages. Circumvent this problem by executing the TMSUPDTE Utility, changing flag *3* to *20*. For more information about this utility, see your CA 1 guides.

Procedures that are more detailed can be obtained from CA 1 Technical Support for the release of CA 1 you are running.

Other CA Disk sysparms, which control the data set name, expiration date, and catalog action of CA Disk tapes can be set as described in the following section. The expiration date on the internal label of an externally managed tape is not affected by the EDM facility.

Method 2 — Controlling Tapes by Expiration Date

As distributed, CA Disk defaults to assigning expiration date of 99365 to all tapes, but does not catalog them. This technique is intended for those installations that do not have a tape management system. When CA Disk determines that all data on the tape has expired, it is released and made eligible as a scratch tape. This is the recommended method for controlling the CA Disk tapes.

Review if sysparm TAPEFSCR should be specified. CA 1 users should not specify it, but it is recommended for all other users.

To control tapes by a real expiration date, change the expiration date to be assigned using sysparm DYNEXPDT, described in the *Systems Guide*, or your JCL to specify a true expiration date. The value 99000 can be used to place the tapes under catalog control as described in Method 3 (see page 273).

By assigning a real date for expiration of the tape, there is a risk that the tape could be released before the data set records are expired so data could be lost.

You can choose to also catalog the tapes that are actually under expiration date control if it provides you with any additional benefit. Neither CA Disk nor CA 1 makes use of the catalog entry to locate tapes.

To have CA Disk catalog the tapes it creates, specify sysparm ARCTNAME with a value of C, rather than specify DISP=(NEW,CATLG) in the JCL. CA Disk generates a unique name for each tape it creates, and then catalogs it. You should also specify sysparm UNCATARC with a value of Y, to have CA Disk uncatalog each tape when it returns it to the scratch pool.

Note: For other applicable rules, see the section Assigning Tape Expiration Dates in the *Systems Guide*.

Method 3 — Controlling Tapes by Catalog Status

Consider this technique if your tape management system supports an option to designate tapes as in-use as long as they are cataloged, and as available scratch tapes when they are uncataloged. This technique is appropriate for CA TLMS or CA 1, both from CA. When CA Disk determines that all data on a tape has expired, CA Disk releases the tape and makes it eligible to be a scratch tape by uncataloging it.

To implement catalog control, you must use the expiration date that your tape management system defines for that purpose. For tapes dynamically allocated by CA Disk, you must specify this value in sysparm DYNEXPDT. If at some point you decide to override dynamic allocation by providing JCL for the output tapes, you must provide this value in the LABEL=EXPDT=yyddd parameter for the tape DD statements.

To have CA Disk catalog the tapes it creates, do not specify DISP=(NEW,CATLG) in the JCL. Instead, specify sysparm ARCTNAME with an appended value of C. CA Disk generates a unique name for each tape it creates, and then catalog it. You should also specify sysparm UNCATARC with a value of Y, to have CA Disk help keep your system catalog clean by uncataloging each tape when it returns it to the scratch pool.

Important! In the event you lose the catalog that your CA Disk tapes are cataloged in, your tape management system can scratch all CA Disk tapes. Take appropriate steps to prevent this from occurring.

Method 4 — Controlling Tapes Using CA TLMS/EDM

This support provides a module for the TMSCTLEX user exit that interfaces with the EDM developed for CA TLMS. To enable the support, sysparms TMSCTLEX and DYNEXPDT must be set to the values ADSTH017 and E99365 respectively. See also instructions in the CA TLMS documentation supplied by CA.

Note: For detailed information regarding this support, see ADSTH017—Interface for CA TLMS in the *Systems Guide*.

CA Disk Interfacing Considerations

It is a good idea to run the DSNDELETE command of IXMAINT (documented in the DSNDELETE Command and Parameters section of the *User Guide*) nightly before the tape management scratch and clean functions are run. This makes the expired CA Disk tapes available as scratch tapes.

For users who do not use the EDM interface, if a job abends while writing a tape, most tape management systems consider that output tape a scratch as a default. To CA Disk, however, the partial tape is a good tape and must be kept. Take special precautions to prevent your tape management system from marking CA Disk output tapes as scratch tapes after an abend. This can be done manually, or for CA 1, you can make use of one of the exits within the CA 1 system. A CA Disk user has supplied a sample for this exit in member SLI035 of the user mod library. Consult your local CA 1 support staff or CA directly, if necessary, if you have further questions regarding their exit.

Limiting ISPF User Access to Data Sets

Specify sysparm SPFUSRID, described under SPFUSRIDn in the *Systems Guide*, with a value of Y if ISPF functions are to allow users to process only those data sets with names prefixed with their TSO user ID. This feature can be overridden for selected users by specifying their user IDs in the TSOUSERI member of parmlib, described in TSO/ISPF USERID Authorization List in the *Systems Guide*.

ISPF Function-Generated JCL

CA Disk generates JCL to execute various functions in the background environment. All generated JCL executes standard CA Disk JCL procedures that are distributed with the system. It is assumed that these procedures are installed in an accessible procedure library and are compatible with those originally shipped with the system.

Defining ISPF and DSCL User Options

CA Disk dynamically builds its own menus of available ISPF and DSCL functions. This is done to provide a flexible way for CA Disk installation coordinators to control the use of CA Disk in their installation.

When a user enters the CA Disk option on the ISPF primary option menu, CA Disk determines the user ID of the user and goes to an authorization list stored as a member in the CA Disk parmlib data set. This authorization list indicates which users can use CA Disk and what functions of the support they are authorized to use. Based on this information, a menu is presented to the user indicating only those functions that are authorized for their use.

At this point of the ISPF customization, members SPFOPTNS and DMCOPTNS should be created in the parmlib data set by copying in members SAMPSPFO and SAMPDMCO respectively, and then reviewing and modifying as needed.

SPFOPTNS and DMCOPTNS parmlib members contain three fields of information those are coded free form on each list entry. The list contains all the needed entries as needed. The three fields are: TSO USERID, INCLUDE/EXCLUDE Specification and Function Specification.

Field 1 — TSO USERID

This field is a TSO USERID or prefix (partial user ID ending with a slash). This field must start in the first position of the entry after the starting single quote. The slash (/) can be used to indicate groups of users to which the entry applies. A single slash indicates that the entry applies to all users. This field must be followed by one or more spaces.

Field 2 — INCLUDE and EXCLUDE Specification

This field is a one-character field indicating the specification type: *I* for include, *E* for exclude. It must follow the user ID field and be separated before and after with at least one space.

Entries are specified to allow users to process certain options (include) or to disallow users from using options (exclude).

This can be useful by specifying an include statement for a user indicating that all options are allowed. Then by specifying an exclude statement for one option, the user is allowed access to all options except one. This prevents the need to specify every function that a user can use.

Any number of statements can be specified for a user or group of users. One or more include statements must be specified to allow a user access to a function. Any exclude statement for a function overrides an include statement for it; that is, an exclude statement applied to all user IDs will deny access to the specified function regardless of any other include statements.

Field 3 — Function Specification

This field is a list of CA Disk functions that a user or user group is authorized to use. Function codes are seven characters and can be abbreviated to the first three characters. Commas, not by spaces, must separate functions. Up to 150 characters of functions can be specified using the continuation rules of CA Disk parmlib entries. The last function code is followed by a single quote.

The following is a list codes and their valid specifiable functions:

ALL

Indicates that all functions are to be assumed for this specification. In specifying the ALL function, all features are assumed. For any function that has not been licensed, subsequent use of the ISPF option to invoke it can result in abends looking for routines you do not have. If applicable, this value should be specified alone in field 3 of the statement.

ISPF Function: SPECIAL FUNCTION

BARCHIV

Generate JCL for batch job submission.

ISPF Function: ARCHIVE/BACKUP - SUBMIT BATCH JOB TO ARCHIVE DATA SET

BDARCHI

Generate JCL for batch job submission.

ISPF Function: ARCHIVE/BACKUP - QUEUE A REQUEST (VIA BATCH)

FDARCHI

ISPF foreground function: allow user to enter a deferred archive command.

ISPF Function: ARCHIVE/BACKUP - QUEUE A REQUEST ONLINE

BDRESTO

Generate JCL for batch job submission.

ISPF Function: RESTORE - QUEUE A REQUEST (VIA BATCH)

BRESTOR

Generate JCL for batch job submission.

ISPF Function: RESTORE - SUBMIT BATCH JOB TO RESTORE DATA SETS

FDRESTO

ISPF foreground function: allow user to enter a deferred restore command.

ISPF Function: RESTORE - QUEUE A REQUEST ONLINE

FRESTOR

Invoke TSO RESTORE command.

ISPF Function: RESTORE - EXECUTE RESTORE REQUEST (VIA TSO)

LDSINDX

ISPF foreground function: allow user to list entries only.

ISPF Function: LIST - ARCHIVE/BACKUP INDEX ENTRIES

DDSINDX

ISPF foreground function: allow user to list and delete entries for archived data sets.

ISPF Function: LIST - (OR DELETE) ARCHIVE/BACKUP INDEX ENTRIES

CDSINDX

ISPF foreground function: allow user to list and change expiration date values for archived data sets.

ISPF Function: LIST - (OR CHANGE) ARCHIVE/BACKUP INDEX ENTRIES

ADSINDX

ISPF foreground function: allow user to list change expiration dates and delete entries for archived data sets. (If the delete function is requested for an index record associated with a RACF-protected data set for which a discrete profile has been saved, authorization of the ISPF session is required. This can be accomplished if you have a user SVC to obtain authorization, and indicate it to CA Disk by sysparm RADELSVC. If this cannot be provided, give your users the CDSINDX function instead of this one.)

ISPF Function: LIST - (DELETE OR CHANGE) ARCHIVE/BACKUP INDEX ENTRIES

LQUEUED

ISPF foreground function: allow user to list deferred archive and restore requests previously entered.

ISPF Function: LIST - QUEUED ARCHIVE/RESTORE REQUESTS

DQUEUED

ISPF foreground function: allow user to list and delete deferred archive and restore requests previously entered.

ISPF Function: LIST - (OR DELETE) QUEUED ARCHIVE/RESTORE REQUESTS

EDSKDSK

Generate JCL to invoke the Move/Copy function using the COPY command.

ISPF Function: MIGRATE - MOVE A SPECIFIC DATA SET TO A NEW VOLUME

EPDSCOM

Generate JCL to invoke the PDS Compress function for specific data sets.

ISPF Function: COMPRESS - COMPRESSA SPECIFIC PDS

ERELOAD

Generate JCL to invoke the PDS Compress RELOAD function.

ISPF Function: COMPRESS - RESTART COMPRESS OF SPECIFIC PDS

IREPORT

Generate JCL to invoke the REPORT function.

ISPF Function: REPORTS - SCAN VTOCS: GENERATE FIXED FORMAT REPORTS

IARCHIV

Generate JCL to invoke the implicit archive (RETAIN) function.

ISPF Function: ARCHIVE/BACKUP - SCAN VTOCS TO SELECT DATA SETS

FVRPT

Generate JCL to invoke the explicit VSAM (FIND) function to generate VSAM reports only.

ISPF Function: REPORTS - FIND SPECIFIC VSAM CLUSTER OR GROUP

FVSAM

Generate JCL to invoke the explicit VSAM (FIND) function.

ISPF Function: ARCHIVE/BACKUP - FIND SPECIFIC VSAM CLUSTER OR GROUP

IVRPT

Generate JCL to invoke the implicit VSAM function to generate VSAM reports only.

ISPF Function: REPORTS - SCAN CATALOGS TO SELECT VSAM CLUSTERS

IVSAM

Generate JCL to invoke the implicit VSAM function.

ISPF Function: ARCHIVE/BACKUP - SCAN CATALOGS TO SELECT VSAM CLUSTERS

IRCOVER

Generate JCL to invoke the implicit recovery function.

ISPF Function: IMPLICIT RECOVERY - RESTORE DATA SETS FROM A VOLUME

IXMAINT

Generate JCL to invoke the archive index maintenance procedure (IXMAINT).

ISPF Function: MAINTENANCE - DELETE EXPIRED ENTRIES IN ARCHIVE INDEX

VOLDISP

Display information about CA Disk archive volumes in the foreground ISPF environment.

ISPF Function: LIST - ARCHIVE VOLUME INFORMATION

TAPEPOO

Invoke the interactive CA Disk tapepool management functions.

ISPF Function: MAINTENANCE - UPDATE ARCHIVE TAPE POOLS ONLINE

RELEASE

Generate JCL to invoke the idle space release function.

ISPF Function: RELEASE - SCAN VTOCS: RELEASE IDLE SPACE

IDSKDSK

Generate JCL to invoke the implicit Move function.

ISPF Function: MIGRATE - SCAN VTOCS: MOVE DATA SETS TO NEW VOLUMES

IPDSCOM

Generate JCL to invoke the implicit PDS Compress function.

ISPF Function: COMPRESS - SCAN VTOCS: COMPRESS ANY PDS IF ITS NEEDED

IPRELOA

Generate JCL to restart the PDS Compress RELOAD function.

ISPF Function: COMPRESS - SCAN VTOCS: RESTART AFTER AN ERROR

ISEQMIG

Generate JCL to invoke the implicit sequential migration to tape function.

ISPF Function: MIGRATE - SCAN VTOCS: MOVE PS DATA SETS TO TAPE

ALTDISP

When this option is selected it will display a panel that shows the names of the CA Disk data sets that will be used for CA Disk requests.

ISPF Function: MISCELLANEOUS - DISPLAY RELEASE, FILES, AND PARMLIB

SYNCHEK

This utility function causes CA Disk to check the syntax for any member of the CA Disk parmlib data set. This is normally used after a user has modified some member in parmlib.

ISPF Function: MISCELLANEOUS - SYNTAX CHECK A MEMBER OF PARMLIB

ONLIRPT

Generate online reports using the DSCL selection language. After defining a report online it can be produced either online or through a batch job.

ISPF Function: REPORTS - DEFINE YOUR OWN EXECUTE ONLINE (OPTIONAL)

FILDUMP

Generate online hexadecimal dump of subfile records in the FILES.

ISPF Function: MISCELLANEOUS - DISPLAY FILES data set RECORD IN HEX

DSCL

Generate JCL for batch job submission.

ISPF Function: GENERATE DSCL COMMANDS

ONLIRPT

Generate a REPORT command with the SPFRPTS parameter.

ISPF Function: REPORT USING SPFRPTS

IREPORT

Generate a REPORT command.

ISPF Function: REPORT USING DSCL

VBACKUP

Generate a VBACKUP command.

ISPF Function: BACKUP A PHYSICAL VOLUME

Since there are more options available than menu lines on a screen, it is sometimes necessary for CA Disk to implement two levels of menus. This occurs when more than 15 options are given to any user. If a user has fewer than 15 options, all options will be displayed on one selection menu. If more than 15 options are specified, the first menu will display a list of functional categories from which the user must select. Once this has been specified, a second menu will be displayed with the appropriate functions listed.

When specifying user options, the CA Disk installation coordinator considers carefully what options to give each user. Several of the options are similar to each other, and are supplied to give the installation flexibility in distributing functions. Every user does not need every function available. The distributed sample members SAMPSPFO and SAMPDMCO attempts to eliminate this duplication of function, but users have different requirements and need to have different options specified.

In addition, remove any options to which you do not want your users to have access. If your installation has not implemented the CA Disk archive/backup tape pool facility, it should also be removed from the menu display.

The sample options members SAMPSPFO and SAMPDMCO supplied with your system are already set up with entries that can be further tailored for your installation. If you have not already done so, copy them in to create members SPFOPTNS and DMCOPTNS. Two sample entries exist in the member SAMPSPFO. The first entry specifies the ISPF functions that are normally given to all CA Disk users. The second entry specifies the ISPF functions that are to be given to only those persons responsible for DASD management.

The two entries are shown in the following illustration:

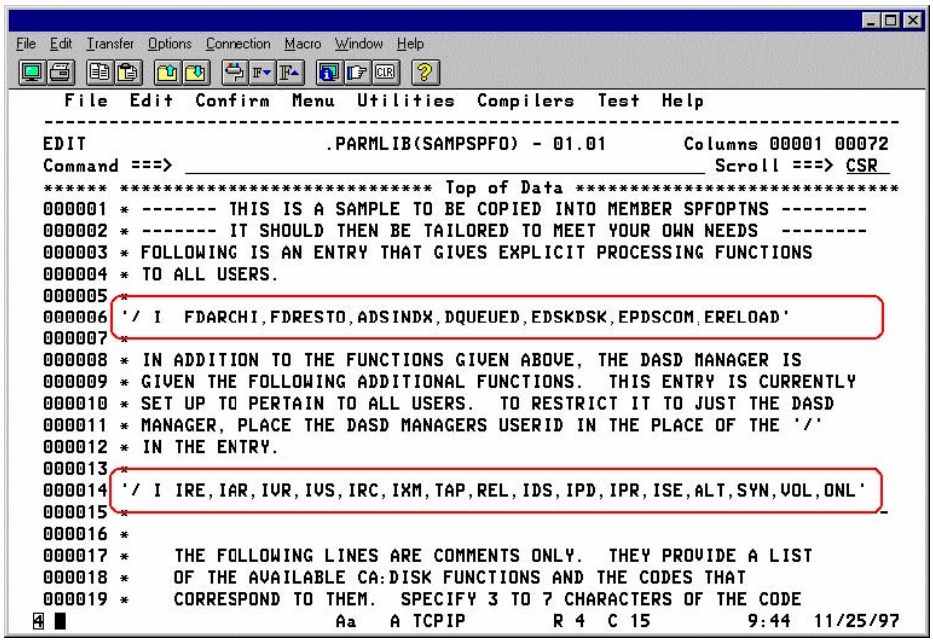

ISPF Return Function Restriction

The CA Disk ISPF function does not fully support the ISPF RETURN or JUMP facility. Users can use the facility to go from CA Disk panels to non-CA Disk panels, but cannot specify multiple menu level numbers when entering CA Disk functions. For example, *=8* and the <RETURN> key can be used to get to the CA Disk ISPF menu, but *=8.4* cannot be used to get to a specific CA Disk ISPF function. However, the ISPF command delimiter character can be used in place of the period. When the command delimiter character is a ";", entering *=8;4;2* selects the fourth option of the condensed menu, then the second option from the selection menu.

Customizing the CA Disk TSO Support

CA Disk can be customized to support TSO. The following Tape Units Allocated Concurrently, Tape Unit Name, Limiting TSO User Access, and TSO Screening Exits for Archive and Restore are some of the custom features that can be initiated.

Tape Units Allocated Concurrently

By default, only one tape unit (one restore) can be allocated by dynamic restore at any given time. Sysparm TSOTULMT can be specified with a numeric value to permit two or more concurrent allocations (restores).

```
TSOTULMT05      <=== tape unit limit set to five
```

Tape Unit Name

The default unit name for allocating a tape drive is TAPE. If this unit name does not cover the devices, use the following sysparm to provide the correct unit name.

```
TSOTUNITnnnnnnnn        <=== where "nnnnnnnn" is the correct name
```

Limiting TSO User Access

The DARCHIVE and DRESTORE commands allow TSO users to archive and restore any data set. To restrict most users to archive and restore only those data sets prefixed with their user IDs, sysparm TSOUSRID must be specified with a value of Y. Access to all data sets can then be given to selected users by placing their user IDs in member TSOUSERI of the parmlib data set.

To allow DARCHIVE requests to have access against offline volumes, sysparm TSOVOLOF must be specified with a value of Y. The default (N) enforces that the data set requested must be found online or the request is rejected. Update parmlib member SYSPARMS with entries TSOUSRIDY and TSOVOLOFY, and create member TSOUSERI with the proper user IDs to accomplish both of these entries.

TSO Screening Exits for Archive and Restore

Exits are also available that allow an installation to screen archive and restore requests. These exits are documented under RESCRNEX*mmmmmmmm* in the *Systems Guide* and can be referenced there if more restrictions that are elaborate are needed.

Both user ID restrictions (sysparm TSOUSRID = Y) and the TSO user screening exits can be used at the same time. The user ID screening check is made first. The user exit module never sees the request if the ID requirement is not met.

Customizing the CA Disk ARCGIVER Support

The ARCGIVER module is provided for those customers who do not already have this module supplied by the vendor of their operating system. If you have the ARCGIVER module that was supplied by your operating system vendor, ignore this version of ARCGIVER, and continue to use the ARCGIVER that is already on your system. If you will be using the CA Disk version of ARCGIVER, you must IPL with CLPA for the change to take effect.

Customizing the Unicenter Service Desk Support

The following sections describe the Service Desk Integration (CAISDI) and how you customize the Unicenter Service Desk Support.

Appendix A: Installing CA Disk Under CA ACF2

To run effectively on a system using CA ACF2, follow these steps. The steps apply even if you do not have the CA Disk Security Interface for CA ACF2 installed.

To process the data sets managed by CA Disk, setup your production CA Disk runs to have sufficient access. The simplest way to accomplish this access is to run these CA Disk tasks with a Logon ID that has the non-CNCL attribute.

Many users with CA ACF2 select instead to set up a Logon ID as a data administrator.

To install CA Disk under CA ACF2:

1. Create a Logon ID that has the MAINT attribute.
2. Create a CA ACF2 GSO options record to allow that Logon ID access using the program ADSMI002 from your CA.DISK.LOADLIB.

   ```
   MAINT    LIBRARY(CA.DISK.LOADLIB)
            LID(logonid)
            PGM(ADSMI002)
   ```

   If you already have a MAINT record, append a qualifier to the record name to generate a unique RECID.

   ```
   MAINT.DMS LIBRARY(CA.DISK.LOADLIB)
             LID(logonid)
             PGM(ADSMI002)
   ```

   Except for the TSO command processors, most of the time CA Disk runs as the program name ADSMI002. The remaining usage is either with the program name of ADSMI000 or ADSMI302.

   The only time to use ADSMI302 is if Customer Support requests a PET trace for debugging a problem. The ADSMI000 program is used where CA Disk does not control the retention of the tape such as Sequential Migrate or FDS Unload. Tapes created with these functions should not be under EDM control.

   Both ADSMI000 and ADSMI302 need a CA ACF2 GSO option.
3. Run these CA Disk production jobs with your Logon ID.

   Other alternatives are possible.

   You can process (archive, restore, migrate, and so on) data sets you are authorized for.

Also, to accomplish its function, CA Disk must update format-1 DSCBs in volume VTOCs. For example, backups must be able to turn off the change bit DS1IND02. Restore and Move/Copy processing must update certain fields that CA Disk maintains. See the CA Disk SVC (see page 226) for a list of those fields. To avoid S913-38 abends as CA Disk tries to do this function, allow CA Disk to open VTOCs for INPUT and OUTPUT.

To run any CA Disk task with a Logon ID that does not have the nonCNCL attribute, add the CA ACF2 access rule. Adding the rule lets CA Disk programs update your VTOCs.

```
$KEY(SYSVTOC)
- UID(*) LIB('CA.DISK.LOADLIB') PGM(ADSMI002)
READ(A) WRITE(A)
```

If you plan to use the CA Disk TSO command RESTORE, add the CA ACF2 access rule:

```
$KEY(SYSVTOC)
- UID(*) LIB('CA.DISK.LOADLIB') PGM(RESTORE)
READ(A) WRITE(A)
```

Note: READ(A) also implies EXEC(A). If the CA Disk load library is in the Linklist, replace LIB('CA.DISK.LOADLIB') with LIB('SYS1.LINKLIB'). WRITE(A) can be replaced with WRITE(L).

Except for the TSO command processors, CA Disk always runs as the program name ADSMI002.

When running all of your CA Disk tasks with Logon IDs that have the non-CNCL attribute, do not worry about CA ACF2 protection for VTOCs.

4. Determine if a user is authorized to perform Storage Administration functions. CA Disk uses the RACROUTE parameter STATUS=ACCESS under ISPF. CA ACF2 does not allow the use of STATUS=ACCESS by an unauthorized ISPF Dialog. Instruct CA ACF2 to allow the access and not abend the request with an S047 abend in module ACF9C000, by installing the following SAFDEF statement:

   ```
   INSERT SAFDEF.adsds1 RB(ADSMI002) PROGRAM(ADSMI002) NOAPFCHK
   RACROUTE(REQUEST=AUTH,CLASS=FACILITY,STATUS=ACCESS,ENTITY=STGADMIN.DMS.-)
   ```

Installing the SAFDEF statement lets the CA Disk ISPF Dialog determine access to the FACILITY class profile without causing a violation or the S047 Abend. Change the operand of the STGADMIN.DMS ENTITY parameter if the default CA Disk FACILITY names controlled by sysparm SMSSTGAD are not used. The default value of SMSSTGAD is STGADMIN.DMS.STGADMIN which is the value ENTITY= in the sample SAFDEF.

This section contains the following topics:

Installing the CA ACF2 Interface (see page 287)

Installing the CA ACF2 Interface

Perform the following procedure to install the CA ACF2 Interface.

To install the CA ACF2 Interface:

1. Activate the CA Disk CA ACF2 Interface by specifying sysparm ACF2SUPP with a value of Y in the SYSPARMS member of the parmlib data set.
2. Review access to data set names VTOC.volser and DMSOS.Vvolser. If you use the *SELECT VTOCS DSCL* statement, CA Disk backs up the VTOC of each volume processed, tracking this information by the esoteric data set name VTOC.*volser*, where *volser* is the volume on which the VTOC resides.

   If you create volume-level backups with the VBACKUP command, CA Disk backs up each volume, tracking this information by the esoteric data set name DMSOS.V*volser*, where *volser* is the volume being backed up.

   BACKUP, VBACKUP, and IXMAINT functions each query any CA Disk Security Interfaces for authority to process these names.

   If you plan to use the SELECT VTOCS DSCL statement, or the VBACKUP command, make sure that your BACKUP, VBACKUP, and IXMAINT functions can each access this fictitious data set name.
3. Examine the following special considerations for implementation of CA ACF2 in multiple-CPU shops.

If the ACF2SUPP sysparm is specified with, a value of Y and CA ACF2 is not active on the CPU on which CA Disk is running, CA Disk issues an appropriate message and abend. This prevents CA Disk from processing data sets that is normally protected by the security system. This situation sometimes occurs in installations that have CA ACF2 installed only on some of the CPUs that CA Disk is running on.

To avoid this problem, specify sysparm ACF2FORC with a value of Y to indicate that CA Disk continues processing even if CA ACF2 is not installed on the system on which it is running. If ACF2FORC is specified as Y and CA ACF2 is not active on the system, no security checking is performed

Appendix B: Installing CA Disk Under CA Top Secret

Regardless of whether or not the CA Disk Security Interface for CA Top Secret is installed, several items must be done for CA Disk to run effectively on a system using CA Top Secret.

1. Set up your production CA Disk runs to have sufficient access to process all data sets managed by CA Disk. The simplest way to accomplish this is to run these CA Disk tasks with a user ACID that has all authority to all data sets.

   Other alternatives are also possible. Individual users can process (archive, restore, and so on.) for only those data sets they are authorized too.

2. Decide if you must set sysparm ARSECURE. If you do not intend to set CA Disk sysparm ARSECURE, proceed to step 4.

   If you intend to set CA Disk sysparm ARSECURE, according to the supplier of CA Top Secret, you must set up CA Disk as an CA Top Secret Facility. Perform the following steps:

   a. Choose an 8-character Facility Name, a unique Facility Code for the TSSUTIL report, and a 7-character user ACID name.

      Create the facility by using the following CA Top Secret Control Options:

      ```
      FAC(USERx=NAME=DMSOS)
      FAC(DMSOS=PGM=ADS)
      FAC(DMSOS=ACTIVE,NOABEND,NOASUBM,NOAUDIT,AUTHINIT)
      FAC(DMSOS=ID=facilitycode)
      FAC(DMSOS=NOINSTDATA,KEY=8,LCFCMD,LOCKTIME=0)
      FAC(DMSOS=NOLUMSG,LOG(MSG))
      FAC(DMSOS=SUAS,NORNDPW,RES,SIGN(M))
      FAC(DMSOS=NOSHRPRF,NOSTMSG,NOTSOC,WARNPW)
      FAC(DMSOS=NOXDEF)
      ```

   b. Build an ACID for the CA Disk auto-restore started task. If you call the ACID DMSAR, you can use the following TSO command:

      ```
      TSS CREATE(DMSAR) NAME('DMS AUTO-RESTORE')-
      FAC(STC) TYPE(USER) PASS(NOPW) -
      DEPT(deptname) MASTFAC(DMSOS)
      ```

c. CA Disk uses the RACROUTE parameter STATUS=ACCESS under ISPF to determine if the user is authorized to perform Storage Administration functions. Due to CA ACF2 not allowing the use of STATUS=ACCESS by an unauthorized ISPF Dialog it is necessary to instruct CA ACF2 to allow the access and not abend the request with a S047 abend in module ACF9C000, by installing the following SAFDEF statement:

```
INSERT SAFDEF.adsds1 RB(ADSMI002) NOAPFCHK
RACROUTE(REQUEST=AUTH,CLASS=FACILITY,STATUS=ACCESS,ENTITY=STGADMIN.DMS.-)
```

This will allow the CA Disk ISPF Dialog the ability to determine access to the FACILITY class profile without causing a violation or the S047 Abend. The operand of the ENTITY parameter (STGADMIN.DMS.-) may need to be changed if the default CA Disk Facility names are not used. The default value of SMSSTGAD is STGADMIN.DMS.STGADMIN which is the value ENTITY= in the sample SAFDEF shown previously.

d. Connect the user ACID to the CA Disk Facility in the CA Top Secret STC record. If you call the ACID DMSAR, you can use the following TSO command:

```
TSS ADDTO(STC) PROC(DMSAR) ACID(DMSAR)
```

e. Connect each of your end-user ACIDs that might use auto-restore to the CA Disk Facility in the CA Top Secret STC record, or if you select to call the facility DMSOS, you can use the following TSO command:

```
TSS ADD(DMSAR) FAC(DMSOS)
```

This section contains the following topics:

Installing the CA Top Secret Security Interface (see page 290)

Installing the CA Top Secret Security Interface

To install the CA Top Secret Interface, perform the following procedure.

To install the CA Top Secret Interface

1. Activate the CA Disk CA Top Secret Security Interface by specifying sysparm TOPSSUPP with a value of Y in the SYSPARMS member of the parmlib data set.
2. Review access to data set names VTOC.volser and DMSOS.Vvolser. If you use the SELECT VTOCS DSCL statement, CA Disk will back up the VTOC of each volume processed, tracking this information by the esoteric data set name *VTOC.volser*, where *volser* is the volume on which the VTOC resides.

   If you create volume-level backups with the VBACKUP command, CA Disk backs up each volume, tracking this information by the esoteric data set name *DMSOS.Vvolser*, where *volser* is the volume being backed up.

   BACKUP, VBACKUP, and IXMAINT functions each query any CA Disk Security Interfaces for authority to process these names.

If you plan to use the SELECT VTOCS DSCL statement, or the VBACKUP command, make sure that BACKUP, VBACKUP, and IXMAINT functions can each access this fictitious data set name.

3. Examine the following special consideration for implementation of CA Top Secret.

   Under most versions of IBM operating systems, OPEN, SCRATCH, and RENAME processing will query CA Top Secret for authorization, regardless of the setting of the RACF-indicator bit. This feature is called *always call*. Data sets cataloged in ICF catalogs also cause a query of CA Top Secret for authorization, regardless of the setting of the RACF-indicator bit.

   Under some operating systems, data sets not cataloged in ICF catalogs queries CA Top Secret only if the RACF-indicator bit is on. For non-VSAM data sets, the RACF-indicator bit is the DS1IND40 bit (bit x'40' at offset 93 x'5D') located in the data set's format-1 DSCB.

   CA Disk security processing normally queries CA Top Secret for authorization, regardless of the setting of the RACF-indicator bit. If you do not have the *always call* feature of the operating system and you do not use ICF catalogs, specify sysparm TOPSALWZ with a value of N in the SYSPARMS member of the parmlib data set.

Appendix C: Installing CA Disk Under IBM RACF

Regardless of whether you intend to install the CA Disk Security Interface for RACF, there are several items that you must consider for CA Disk to run effectively on a system using IBM's RACF.

- Set up your production CA Disk runs with sufficient access to process all data sets managed by CA Disk.

  The simplest way to accomplish this is to run these CA Disk tasks with a user ID that has the OPERATIONS attribute. Other alternatives are also possible.

- Your individual users can process (archive, restore, migrate, and so on) those data sets for which they are authorized.

- RACF does not give different DASDVOL authorities based on the program name being run. With the exception of the TSO command processors, CA Disk always runs as the program name ADSMI002.

- If you plan to run the DSCB Update function with a user ID that has the OPERATIONS attribute, or do not plan to run the utility, you do not need to worry about RACF DASDVOL rules.

This section contains the following topics:

Installing the RACF Security Interface

Thee following is the procedure for installing the CA Disk RACF Security Interface. Steps 1 through 4 are required only for users who have discrete RACF profiles at their shop. Due to RACF's continuing support of discrete profiles, we recommend that all users follow each step.

To install the CA Disk RACF Security Interface

1. Store CA Disk-Saved Discrete Profiles.

   Prepare for CA Disk-saved discrete profiles to be kept in IBM RACF data set. This step is optional but review it for possible applicability.

CA Disk-saved discrete profiles are created and maintained through standard RACF macros (see the section Security Processing in the Systems Guide). RACF places the new profiles in the RACF data set selected by the user. While most users elect to keep all data set profiles in a single RACF data set, it can be useful (for example, to improve RACF performance by reducing contention between CA Disk RACF and other RACF requests) to separate CA Disk profiles from standard RACF data set profiles. This can be done by using the CA Disk saved discrete profile data set name prefix as an identifier to CA Disk profiles (see the next step). Then specify the CA Disk prefix in the RACF RANGE TABLE (ICHRRNG) to indicate to RACF where to place CA Disk profiles.

2. Specify the name for CA Disk-Saved Discrete Profiles.

   Provide a prefix (first qualifier) for the data set name of CA Disk saved discrete profiles by specifying sysparm RACFUSID with a 1- to 8-byte name. This prefix identifies the CA Disk profiles and allows them to be placed on a RACF data set apart from the standard RACF data set. It is a RACF restriction that this RACFUSID value represents a user ID or group ID. We recommend using a user ID, not a group ID, for the value for this sysparm. It is meaningful to you in identifying CA Disk profiles. We recommend using DMSOS.

   To avoid having RACF RACDEF processing update the PERMIT list, ensure that the user ID does not have the GRPACC attribute.

   This RACF user ID is able to restore any data set for which CA Disk has saved a discrete profile. To prevent this exposure of unauthorized use of this RACF user ID, it can be revoked using the TSO command:

   ```
   ALTUSER racfusid REVOKE
   ```

   Revoking the RACF user ID in this manner does not prevent its use for CA Disk discrete profile support.

3. Specify the volume for CA Disk-Saved Discrete Profiles.

   Provide a volume name to be associated with CA Disk-saved discrete profiles by specifying sysparm RACFDVOL with a character volume serial number. Due to RACF restrictions, this volume must be a real DASD volume. Once specified, it must not change. Therefore, select a volume that is always online.

4. Review the RACF Utility function.

   Review the description of the utility documented in the section Management of CA Disk-Saved Profiles in the Systems Guide. You do not need to use the utility at CA Disk installation time, but be aware of the utility's existence. After the CA Disk RACF Security Interface has been implemented, the utility can be run periodically, prior to running the CA Disk IXMAINT function. See the IXMAINT Utility section in the *User Guide*.

5. Set sysparm RACFSUPP and RACFPROC.

   Activate the CA Disk RACF Security Interface by specifying sysparm RACFSUPP and RACFPROC with a value of Y in the SYSPARMS member of the parmlib data set.

6. Review access to special CA Disk data set names.

   Review access to data set names VTOC.volser and DMSOS.Vvolser. If you use the *SELECT VTOCS DSCL* statement, CA Disk backs up the VTOC of each volume processed, tracking this information by the esoteric data set name VTOC.*volser*, where *volser* is the volume on which the VTOC resides.

   If you create volume-level backups with the VBACKUP command, CA Disk backs up each volume, tracking this information by the esoteric data set name DMSOS.V*volser*, where *volser* is the volume being backed up.

   BACKUP, VBACKUP, and IXMAINT functions each query any CA Disk Security Interfaces for authority to process these names.

   If you plan to use the SELECT VTOCS DSCL statement, or the VBACKUP command, make sure that BACKUP, VBACKUP, and IXMAINT functions can each access this fictitious data set name.

7. Special RACF considerations.

   Examine the following special consideration for implementation.

   Under most versions of IBM's operating systems, OPEN, SCRATCH and RENAME processing queries RACF for authorization regardless of the setting of the RACF-indicator bit. This feature is called always call. Data sets cataloged in ICF catalogs also cause a query of RACF for authorization, regardless of the setting of the RACF-indicator bit.

   Under some operating systems, data sets not cataloged in ICF catalogs will query RACF only if the RACF-indicator bit is on. For non-VSAM data sets, the RACF-indicator bit is the DS1IND40 bit (bit x'40' at offset 93 x'5D') located in the data set's format-1 DSCB.

   CA Disk security processing normally queries RACF for authorization, regardless of the setting of the RACF- indicator bit. If you do not have the *always call* feature of the operating system and you do not use ICF catalogs, specify sysparm RACFALWZ with a value of N in the SYSPARMS member of the parmlib data set.

8. Reviewing applicable sysparms.

   See the following sysparm descriptions for possible use in your installation. For more information, see the *Systems Guide*.

   - RACFALLO
   - RACFBKUP
   - RACFDVL2
   - RACFMDSN
   - RACFMODL
   - RACFMVOL
   - RACFNEWN
   - RACFPDSW
   - RACFPRED
   - RACFSEQM
   - RACFVCAV

Appendix D: S213 Abend Exit

This exit is not recommended because some installations, particularly those using SMS, are less likely to have hard coded volsers in JCL. In addition, using this exit can result in restore errors because the time at which it is called does not allow data set allocations to change volumes. These changes are more likely to occur in an SMS environment.

F1-DSCB Not Found S213 Abend Exit (IFG0EX0A Exit)

The recommendation to not use this exit is:

- The methods that cause the exit to be invoked are not generally practical and they are not compatible within an SMS environment. These methods are:
  - Archiving data sets and not changing the volume indicator in the catalog to reflect that they are archived.
  - Coding the VOLSER in JCL to allocate to data sets that are already in existence.
- The S213 exit has the following significant limitations:
  - Only non-VSAM data sets can be auto-restored.
  - No pooling capabilities.
  - Restores cannot be invoked by ISPF.
  - Installation requires an IPL.
  - False invocation of this exit has been known to occur.

The Catalog Management hook (see the Two Auto Methods section) has the following advantages over this exit:

- VSAM and non-VSAM data sets can be auto-restored.
- Restored data sets are pooled by CA Disk, CA Allocate or DFSMS.
- Restores can be invoked by ISPF.
- Installation, removal, or both of the Catalog Management hook is accomplished dynamically with a started task.

The first exit that CA Disk supported was the *F1-DSCB-not-found* user exit in OPEN and EOV processing. The exit is also called the *S213 abend* exit. It is invoked during OPEN processing (non-VSAM) when the data set being opened is cataloged to a volume, but no F1-DSCB exists on that volume for that data set. It is also invoked by batch jobs when the volume of a data set is hard-coded in the JCL, but does not exist on the volume.

At this point, the exit is invoked as a last-ditch effort to save the job from a system 213 abend. If the data set is restored to the volume by the exit, it can signal OPEN to retry the DSCB search again; since it just restored the data set, there should be one out there! If it couldn't restore the data set, the job gets the S213 abend. The dummy exit that IBM supplies with the base MVS system merely sets the return code to tell OPEN to go ahead and abend. The exit name that gets called is IFG0EX0A.

Why two different methods? Because neither of them alone can initiate auto-restores at all times when they are needed. The IFG0EX0A exit does not get called for VSAM data sets, nor does it get invoked for TSO/ISPF functions-- the catalog management hook does, however. So why not implement just the catalog management support? If you only did that, you cannot automatically restore data sets in batch jobs that had the volume hard-coded in the JCL, or when OPEN is processing a data set that is cataloged to a real volume on which the data set doesn't reside (F1-DSCB missing!). Conceptually, think of the two methods as complementary functions. Whenever processing for a data set is going to occur, two primary questions must be answered:

- What volume does the data set reside on?
- What are its attributes on the volume?

To answer the first question, catalog management is invoked with a locate option, which asks where the data set is currently located. It searches its catalogs until it finds the entry for the data set. With this information, OPEN can allocate the volume and then perform an obtain for the F1-DSCB to get the data set attributes and extent information. By understanding the order of information retrieval, you understand how these methods work, and also why the catalog management hook allows more flexibility.

Assume in the above example that the data set is cataloged to a real volume instead of the CA Disk pseudo-volume. If the data set was being opened for processing, a catalog locate is issued to determine the location of the data set. Since the data set is cataloged to a real volume, the catalog management hook lets the locate pass through without any modifications. OPEN processing allocates the volume(s) for the data set. It issues an obtain for the F1-DSCB that describes that data set. If this obtain fails, the IFG0EX0A exit gets control to determine if it can restore the data set. If it can restore the data set at this point, however, it must go back to the volume to which it is cataloged. OPEN has already allocated the volume and is expecting the data set on that volume only. This is why the catalog management hook allows more flexibility--it intercepts the locate before OPEN allocates the volume(s). It can therefore put the data set back to any volume it wants, as long as it passes back the correct volume list to the requester of the locate.

Installing the S213 Abend Exit

Installing the S213 exit requires an IPL. Perform the following procedure to install the S213 abend exit.

To install the S213 abend exit

1. If you installed module IFG0EX0A into your LPALIB from a release of CA Disk prior to Release 7.1, remove it as follows.
   a. Assemble and link member IFG0EX0A, contained in the CA Disk installation library, into your SYS1.LPALIB.
   b. If you also have DFHSM and want CA Disk and DFHSM to auto-restore data sets, use the IFG0EX0A supplied by DFHSM.
2. If you did not install module IFG0EX0A into your LPALIB from a release of CA Disk prior to Release 7.1, continue here.
3. Sample JCL is provided in the installation library member USERMOD7.
4. After completing the relink, run the following JCL to verify your changes.

   ```
   //LIST     EXEC PGM=AMBLIST
   //SYSPRINT  DD  SYSOUT=*
   //SYSLIB    DD  DISP=SHR,DSN=SYS1.LPALIB
   //SYSIN     DD  *
       LISTLOAD MEMBER=(IFG0EX0A)
   ```

   Make sure that the correct version of IFG0EX0A has been link-edited. If DFHSM is not installed, the IFG0EX0A CSECT in module IFG0EX0A should consist of only two instructions:

   ```
   SR 15,15
   BR 14.
   ```

5. IPL your system with either the CLPA or MLPA option. The auto-restore function is now active. If MLPA is used, be sure to add the proper entry to SYS1.PARMLIB member IEALPAxx and specify it during the IPL.

Note: When changing modules that reside in LPA, it can be common practice to link the new version into a linklist library other than SYS1.LPALIB, and then use the MLPA ability to test the new module. If all testing goes well, you might use IEBCOPY to copy and replace the old version in SYS1.LPALIB, and then schedule an IPL with a CLPA. This all works, but only if you specify the entire alias names as well as the true member name in the SELECT MEMBER=(...list...) statement for IEBCOPY. Omitting the alias names can cause the IPL with CLPA to fail with unpredictable symptoms.

If archived data sets are cataloged to the CA Disk pseudo-volume, they can be restored by the catalog management hook to any one of a pool of volumes. If the data set is recataloged to a real volume, only the S213 exit of OPEN will be invoked and the restored data set has to go back to that specific volume. The real volume to which archived data sets are recataloged can be thought of as a common *staging* volume for auto-restores. The volume can be monitored for activity and easily managed to ensure enough space to handle the restores.

If the S213 exit is installed, the following restrictions apply:

- The IFG0EX0A exit does not get invoked for VSAM data sets, and therefore they must be restored manually.
- The data set must be restored to the same volume to which the 213 was invoked. If sufficient free space does not exist on the volume, the restore fails and the originating task terminates with a system 213 abend.
- The IFG0EX0A exit is not entered from TSO and ISPF functions.

Appendix E: Message Check Utility

CA Disk Message Check is a utility centralizing messages from different jobs for viewing and decision-making from one location. The application uses SYSOUT EXIT to filter messages and write them to data sets (message data sets). As a system administrator you can then run batch jobs to consolidate all message data sets into a single work data set and report on it.

Comparing JOBNAME does message filtering and CA Disk messages with two user-defined tables with JOBNAME/prefixes and CA Disk message number/word/strings. The two tables are PDS members and can be updated at any time through ISPF panels. The selected messages are written to dynamically allocated message data sets. DSN of a message data set is determined by the user specified prefix (through the application ISPF panel) and JOBNAME. When consolidating message data sets, messages are merged and sorted into generation data sets and all message data sets are then deleted. Message check does not alter any CA Disk job output.

Message check supports the following DD:

- MSGPRINT
- SYSPRINT
- CMDPRINT
- RELPRINT
- MOVPRINT

Installing CA Disk Message Check

To install the CA Disk message check perform the following procedure.

To install the CA Disk message check

1. Create a PDS that is used to hold user defined tables (for JOBNAME's, message number/word/string's). The PDS can be RECFM=FB and LRECL=80 with 1 directory block. You need update access to maintain the tables. See Step 7.
2. Define up to 4 GDG base entries for a consolidation run. Use LIMIT (255) for each of them. For example:

```
DEFINE GENERATIONDATAGROUP -
NAME(CA DISK.MSGCHECK.CONSOLID.GDG.ONE) -
LIMIT(255))
```

3. Insert the call (shown in the following) to the application into an ISPF selection panel. The panel can be your primary ISPF panel.

   ```
   11,'CMD(%SMSPC) NEWAPPL(SMSP)'
   ```

   The corresponding text may look like option 11 in the following list of ISPF Options:

ISPF Option	Primary	Option Menu
0	Settings	Terminal and user parameters
1	View	Display source data or listings
2	Edit	Create or change source data
3	Utilities	Perform utility functions
4	Foreground	Interactive language processing
5	Batch	Submit job for language processing
6	Command	Enter TSO or Workstation commands
7	Dialog Test	Perform dialog testing
8	LM Facility	Library administrator functions
9	IBM Products	IBM program development products
10	SCLM	SW Configuration Library Manager
11	SMSP	CA Disk Message Check
IP	IPCS	IPCS Problem Analysis Services
S	SDSF	System display and Search Facility
R	BookMgr/Read	IBM BookManager/Read MVS
TP	Third Party	Third Party Products
M	MVS	MVS Support Menu
N	Network	Network Product & Support Menu
D	Database	Database Product & Support Menu

   Enter X to Terminate using log/list defaults.

4. Initialize the global environment of the application by entering the following command through the dialog:

   ```
   Service. (ISPF 7.6)
   SELECT CMD(%SMSPCOK) NEWAPPL(SMSP)
   ```

For example:

```
------- CA Disk Message Check Application Enter required field -------

Command ===>


CA Disk Message Log Output (DSN Prefix will be appended with Jobname)

data set Prefix ===>

DASD Unit ===> can be sysallda


CA Disk Load Library - specify only the default here

data set Name ===>


Check Jobs/Steps with DD-Stmt. SYS$EXIT only? ===> (Y/N Yes/No=all Jobs)

Print Summary Lines from CA-Disk? ===> (Y/N Yes/No)

Propagate CA-Disk Messages to CA-Vantage? ===> (0, 1 or 2)

0 = No, 1 = only recognized Error Messages, 2 = all checked Messages


Library Name ===>

for the Tables for this application - Tables are generated automatically.

(if you want to change this DSN please press HELP for instructions)


Press END KEY to save, Enter CANCEL command to cancel
```

Data set Prefix

The DSN prefix for all message data sets.

Check Jobs/Steps with DD-Stmt.

Indicating Yes performs the message check only if SYS$EXIT DD SYSOUT=* is in JCL.

Note: You still need to define JOBNAME table (see Step 9) to make this option become effective.

Print Summary Lines from CA-Disk

Indicating Yes includes report summary lines (for example: total data sets archived …).

Library Name

The PDS name created in Step 1. Tables (members) with initial values are created.

These fields can be modified through the panel defined in Step 3.

5. Assemble/link the module SYSEXIT$.

   Submit the member USERMODL in the INSTALLIB to activate the message check exit.

6. Ensure that the exit program VANSDM00 resides in the CCTULINK library if you want to propagate filtered messages to CA Vantage. Repeat this step if any CA Vantage PTF changes VANSDM00.

7. **Note:** For additional information, see the section Usermods – Exits to Other Products in the *CA Vantage Reference Guide*.

8. Selecting the option you defined in Step 3 gives you the following panel:

   Note: Storage Management Experts should use this application.

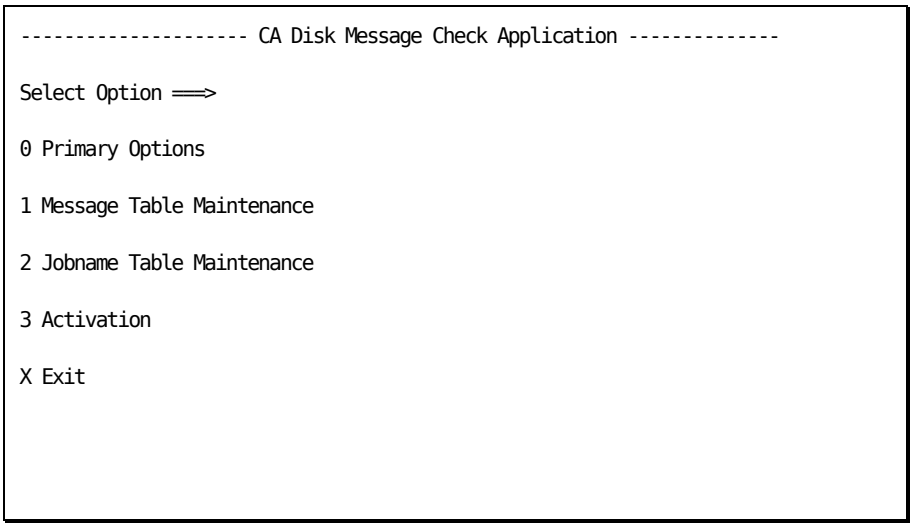

Option 0

Update fields by selecting this option. To change the PDS created in Step 1, you need to:

- Allocate a new PDS with the same attributes
- Copy members (tables with initial values created in Step 4) from the original PDS to the newly allocated PDS
- Update the library name (the last fields) in this panel

Option 1

To specify message number/word/strings for filtering process. See Step 8 for more details.

Option 2

To specify JOBNAME's for filtering process. See Step 9 for more details.

Option 3

To activate any update made in the previous 3 options.

9. To Insert or Update a CA Disk Message, select option 1 from the panel in the CA Disk Message Check Application screen to specify a message number, a specific word in message text, or a string in message text to identify messages you don't want to keep. You can add, delete, or change any entry in this table but make sure there is at least one entry in the table. The following is an example of inserting a new entry:

```
------------------ Insert or Update CA Disk Message ---------------------

Command ===>

Message Number      ===>    9999    4 Byte CA Disk Message Number

Check   ===>        W       S=String, W=Word, N=Normal (check only Number)

Offset or Counter   ===>    6       Position for "String" or nth "Word" in Message

String or Word      ===>    PDSE

Comment  ===>        Idle space release bypassing PDSE


Enter a CA Disk Message Number and additional check options above:

N: Check Message Number only (normal check)

Other checks:

S:  Check for a string starting on column n of the Message Text

    Allowed Range: 1 to 110 (Character 1 is the first character of the Text)

    The Message Number and leading blanks are not counted here.

W:  Check for a Word (counts all words in the Message Text and checks the

    The word with the given number)

    Allowed Range: 2 to 64 (word one has to be searched as a string)

        Press END KEY to terminate
```

To interpret the resulting message text:

```
ADSDM341 9999 DATA SET = DDD.SSS.NNN BYPASSED, PDSE DATA SET NOT COMPRESSED
```

9999 indicates the informational message issued by the PDS compress process. Specify W to check for a specific word, the word is the 6th word in message text (PDSE) Don’t count the module name - ADSDM341 and message number - 9999; any character(s) between two blanks is counted as a word), and the specific word to search is PDSE.

Brief comments are allowed to identify this particular message.

Note: We use 9999 as an example because CA Disk issues various 9999 messages and some of them are not informational but critical errors. For this same message, you may change the entry to check for a string such as 'DATA SET =' while setting the offset field to 1.

Words (a word string in a message text) or message numbers can be entered with N to identify a specific message. A 4-digit message number is required for each entry.

10. To Insert or Update a CA Disk Jobname or Jobname Prefix - Select option 2 from the panel in the CA Disk Message Check Application screen to specify generic or fully qualified JOBNAME's in this table. These identify which jobs are eligible for the message check process. You can add, delete, or change any entry in this table. At least one entry in the table is required. The following is an example of inserting a new entry:

```
------------Insert or Update Disk Jobname or Jobname Prefix -----------------

Command ===>

    Jobname or Prefix   ===> dsktest    (see below)

    Type                ===> p          (N=full qualified Jobname, P=Jobname

                                         Prefix)

    Comment             ===>            CA Disk test jobs

    Include/Exclude     ===>            I (I=Include, E=Exclude this Job)

    Enter a Jobname      full qualified, Type is N )

                         Or a

    Generic Jobname      (which means you specify a prefix instead a full

                         qualified Jobname and Type must be P)

All Jobs which are excluded by these table entries are bypassed for all

Message checks ¦

                        Press END KEY to terminate
```

To validate a message check against all CA Disk jobs with JOBNAME DSKTEST* complete the following steps:

a. Enter the JOBNAME prefix in the Jobname or Prefix field (dsktest).

b. Enter P in the type field (dsktest is the prefix, not full job name).

c. Enter I to include these job(s) in the message check process.

If you have specified Y for Check Jobs/Steps with DD-Stmt. SYS$EXIT only? ===> (Y/N Yes/No=all Jobs) in Step 2, the DD SYS$EXIT must be in every JCL of those jobs included for message checking.

11. To Update or Change the tables, select option 3 from the panel in the CA-Disk Message Check Application screen to activate the two tables you have updated in Step 8 (Insert or Update CA Disk Message)and Step 9 (Insert or Update CA Disk Jobname or Jobname Prefix) Clearing the first Job Card in this panel sets the entire JCL to defaults.

```
---------- JCL to activate the CA Disk Message and Jobname Tables ------

COMMAND ===>


JOB Statement Information:         (verify before proceeding)


    ===> //       ISPAAA1A JOB (ACCOUNT),'NAME'

    ===> //       NOTIFY=ISPSXC1,CLASS=A,

    ===> //       MSGLEVEL=(1,1),MSGCLASS=T,

    ===> //       COND=(0,LT)

    ===> //*

    ===> //*

    ===> //*

Press ENTER to proceed

                        Use END Command to exit
```

Change the JOB card and review the entire job stream. There are three steps in this job: assemble, link, and copy.

```
--------------- JCL to activate the CA Disk Message and Jobname Tables ----------

COMMAND ===>

EXEC/DD Statement Information (1 of 3) (verify before proceeding)

    ===> //ASSEMBLY EXEC PGM=ASMA90,PARM='NODECK,OBJECT,LIST'

    ===> //*

    ===> //SYSPRINT DD SYSOUT=*

    ===> //SYSLIB       DD DISP=SHR,DSN=SYS1.MACLIB

    ===> //SYSUT1       DD UNIT=SYSDA,SPACE=(CYL,(1,1))

    ===> //SYSLIN       DD UNIT=SYSDA,SPACE=(CYL,(1,1,1)),

    ===> //     DSN=&TEMP(SMSP$TAB),DISP=(,PASS),

    ===> //     DSORG=PO,BLKSIZE=3200

    ===> //*

            //SYSIN     DD *     (will be inserted here)

                Press ENTER to proceed

                Use END Command to exit
```

```
-------------- JCL to activate the CA-Disk Message and Jobname Tables -----------

COMMAND ===>

EXEC/DD Statement Information (2 of 3)   (verify before proceeding)


    ===> //BIND  EXEC PGM=IEWL,PARM='AC=1,LIST,LET,RENT'

    ===> //SYSPRINT       DD SYSOUT=*

    ===> //SYSUT1       DD UNIT=SYSDA,SPACE=(CYL,(1,1))

    ===> //SYSLMOD       DD DISP=(,PASS),DSN=&LOAD,

    ===> // SPACE=(CYL,(1,1,1)),DSORG=PS,RECFM=U,

    ===> // BLKSIZE=15476,UNIT=SYSDA

    ===> //SYSLIB       DD DISP=(OLD,DELETE),DSN=&TEMP(SMSP$TAB)

    ===> //SYSLIN       DD DDNAME=SYSIN

    ===> //*

    //SYSIN DD *        will be inserted here

        INCLUDE SYSLIB(HYPSMTAB)         will be inserted here

        NAME SYSEXTAB(R)        will be inserted here

    /*          will be inserted here

            Press ENTER to proceed

                        Use END Command to exit
```

If you do not see the following copy step, clear the first JCL statement, press enter, and update the job card again.

```
//COPY EXEC PGM=IEBCOPY,PARM=REPLACE
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DISP=(OLD,DELETE),DSN=&LOAD
//SYSUT2 DD DISP=OLD,DSN= The.PDS.Created.In.Step1
//SYSUT3 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT4 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
```

Once you have all three steps in place, enter the SUB command and check zero return code of this job. The tables initialized in Step 1 will be replaced.

Implementation

Once you have the application installed, and the message and jobname table updated, you can start running your CA Disk jobs with message check in effect. If you specified Y to SYS$EXIT in the primary options panel, you need to add //SYS$EXIT DD SYSOUT=* to the JCL.

For each CA Disk job qualified for the message check process, a unique message data set will be allocated (and cataloged) or extended with the data set name constructed from the data set prefix specified in the primary options panel (in Step 4 or Step 7 of Installation) and JOBNAME of the job.

For example, if you specified CA-DISK.MSGCHK as the data set prefix and job DSKTEST1 is qualified for message check process, the new sequential data set CA Disk.MSGCHK.DSKTEST1 will be allocated and possibly some messages will be written into it. For subsequent jobs with the same jobname DSKTEST1, more messages will be added into the data set.

Eventually you will have many message data sets and should do the consolidate run. The following procedure is an example of the consolidate run.

To do the consolidate run

1. Modify and submit the member MSGCHKJ1.

```
//DOIT  EXEC PGM=IKJEFT01,DYNAMNBR=500,
// PARM='%SMSPC005 #DSN.PREFIX.SPECIFIED.IN.ISPF
//SYSTSPRT  DD SYSOUT=*
//SYSUT2    DD DSN=
    #INSERT.YOUR.GDGNAME.ONE(+1),
//*  Ø
//
   DISP=(,CATLG),SPACE=(189,(5,5),RLSE),AVGREC=K,
//  LRECL=189,RECFM=FB,
//  UNIT=SYSALLDA
//  DCB=(A.VALID.DCB)
//*  Ø
//SYSPROC   DD DSN= #INSERT.SYSPROC.DSN,DISP=SHR
//SYSTSIN   DD DUMMY Ø
```

Where;

#DSN.PREFIX.SPECIFIED.IN.ISPF

Is the Data set prefix specified in the primary options panel.

#INSERT.YOUR.GDGNAME.ONE

Is the GDG base defined in Installation Step 2.

A.VALID.DCB

Creates a model DSCB (PS,FB,LRECL=189).

#INSERT.SYSPROC.DSN,DISP=SHR

Is the Data set where the application's REXX reside.

This job will merge all message data sets into one sequential data set for the next step. All message data sets will be deleted after this point.

The REXX EXEC allocates (concatenates) all message data sets before merging them into a single data set, you may need to run the consolidate jobs more often if the total number of message data sets is more than 255.

2. Modify and submit the member MSGCHKJ2.

```
//SORT2    EXEC PGM=ICEMAN
//SYSPRINT  DD SYSOUT=*
//SYSOUT    DD SYSOUT=*
//SYSUDUMP  DD SYSOUT=*
//SORTIN   DD DSN= #INSERT.YOUR.GDGNAME.ONE,
//         DISP=OLD
//SORTOUT  DD DSN= #INSERT.YOUR.GDGNAME.TWO(+1),
//         DISP=(,CATLG,DELETE),
//         RECFM=FB,LRECL=189,
//         SPACE=(10,(2,1),RLSE),AVGREC=K,
//         UNIT=SYSALLDA
//         DCB=(A.VALID.DCB)
//*
//*-------------------------------------------------------------------
//SYSIN    DD *
ALTSEQ CODE=(F402,F501)
SORT FIELDS=(51,1,BI,A,39,4,CH,A,36,2,CH,A,33,2,CH,A,43,8,CH,A,
              1,8,CH,A,51,1,CH,A,53,1,AQ,D,54,2,BI,A)
//*
```

```
//SORT3    EXEC PGM=ICEMAN
//SYSPRINT  DD SYSOUT=*
//SYSOUT   DD SYSOUT=*
//SYSUDUMP  DD SYSOUT=*
//SORTIN   DD DSN= #INSERT.YOUR.GDGNAME.TWO(+1),
//         DISP=SHR
//SORTOUT  DD DSN= #INSERT.YOUR.GDGNAME.THREE(+1),
//         DISP=(,CATLG,DELETE), Ø GDG base defined in Installation Step 2
//         RECFM=FB,LRECL=189,
//         SPACE=(100,(2,1),RLSE),AVGREC=K,
//         UNIT=SYSALLDA
//         DCB=(A.VALID.DCB)
//*
//*-------------------------------------------------------------------
//* SELECT ERROR MESSAGES
//*-------------------------------------------------------------------
//SYSIN    DD *
OPTION COPY
INCLUDE COND=(51,1,CH,EQ,C' ')
//*
//PRINT4   EXEC PGM=IEBPTPCH
//SYSPRINT  DD SYSOUT=*
//SYSUT1   DD DSN= #INSERT.YOUR.GDGNAME.THREE(+1),
//         DISP=SHR
//SYSUT2   DD SYSOUT=*,CHARS=X15N,
//         LRECL=145,BLKSIZE=145,RECFM=FA
//SYSIN    DD *
PRINT MAXFLDS=99,MAXLINE=68
TITLE      ITEM=('ERROR MESSAGES FROM CA-DISK',11),
               ITEM=('- ALL JOBS -',42),
               ITEM=('LISTING FOR STORAGE MANAGEM. TEAM',75),
               ITEM=('(ROOM:....)',109)
TITLE      ITEM=('JOBNAME',1),
               ITEM=('MESSAGE IDENT',11),
               ITEM=('E R R O R M E S S A G E ',28)
RECORD     FIELD=(8,1,,1),
               FIELD=(132,58,,11)
//*
```

```
//PRINT5   EXEC PGM=IEBPTPCH
//SYSPRINT       DD SYSOUT=*
//SYSUT1T  DD DSN= #INSERT.YOUR.GDGNAME.TWO(+1),
//         DISP=SHR
//SYSUT2   DD SYSOUT=*,CHARS=X15N,
//         LRECL=145,BLKSIZE=145,RECFM=FA
//SYSIN    DD *
PRINT MAXFLDS=99,MAXLINE=68
TITLE      ITEM=('ERROR MSG. AND SUMMARY LINES ',10),
           ITEM=('- ALL JOBS, COMPLETE LIST ',42),
           ITEM=('PLEASE DISTRIBUTE TO STOR. MGMT. ',75),
           ITEM=('(ROOM ....)',109)
TITLE      ITEM=('JOBNAME',1),
           ITEM=('STEPNAME',10),
           ITEM=('PROCSTEP',19),
           ITEM=('DATE + START TIME ',28),
           ITEM=('MESSAGE - OR -',49),
           ITEM=(' SUM. LINE ',66)
RECORD     FIELD=(8,1,,1),
           FIELD=(8,9,,10),
           FIELD=(8,17,,19),
           FIELD=(10,33,,28),
           FIELD=(8,43,,39),
           FIELD=(102,58,,49)
//*
//COPY5    EXEC PGM=ICEGENER
//SYSUT1   DD DSN= #INSERT.YOUR.GDGNAME.ONE,
//         DISP=(OLD,DELETE,KEEP)
//SYSUT2   DD DSN= #INSERT.YOUR.GDGNAME.FOUR(+1),
//         DISP=(,CATLG,DELETE),
//         SPACE=(500,(2,1),RLSE),AVGREC=K,
//         UNIT=SYSALLDA
//         DCB=(A.VALID.DCB)
//*
//*
//*
//*SCRATCH2       DD DSN= #INSERT.YOUR.GDGNAME.TWO,
//*        DISP=(OLD,DELETE,KEEP)
//*SCRATCH3       DD DSN= #INSERT.YOUR.GDGNAME.THREE,
//*        DISP=(OLD,DELETE,KEEP)
//SYSIN    DD *
//SYSPRINT       DD SYSOUT=*
```

Where under //SORT2;

`#INSERT.YOUR.GDGNAME.ONE`

Is the GDG base defined in Installation Step 2.

`#INSERT.YOUR.GDGNAME.TWO`

Is the GDG base defined in Installation Step 2.

A.VALID.DCB

Creates a model DSCB (PS,FB,LRECL=189).

Where under //SORT3;

#INSERT.YOUR.GDGNAME.TWO

Is the GDG base defined in Installation Step 2.

#INSERT.YOUR.GDGNAME.THREE

Is the GDG base defined in Installation Step 2.

A.VALID.DCB

Creates a model DSCB (PS,FB,LRECL=189).

Where under //PRINT4;

#INSERT.YOUR.GDGNAME.THREE

Is the GDG base defined in Installation Step 2.

Where under //PRINT5;

#INSERT.YOUR.GDGNAME.TWO

Is the GDG base defined in Installation Step 2.

Where under //COPY5;

#INSERT.YOUR.GDGNAME.ONE

Is the GDG base defined in Installation Step 2.

#INSERT.YOUR.GDGNAME.FOUR

Is the GDG base defined in Installation Step 2.

Appendix F: The IBM ARCGIVER Interface

The ARCGIVER Interface lets you talk to DFSMShsm. You can ask DFHSM to Back up, Migrate, Recall, or Recover your data sets, and execute DFHSM commands. This interface lets programs divorce data set allocation or catalog management from the recall of an archived or migrated data set. It also lets the program request asynchronous processing so that it can do other work while the request is being executed.

The use and coding of the ARCH macros are described in the IBM manual DFSMShsm Managing Your Own Data. The expansion of these macros includes a LINK to the ARCGIVER module. ARCGIVER builds an MWE (HSMese for Management work element) from the parameters passed to it and in turn passes that MWE to CA Disk. ARCGIVER issues an SVC 109 with the address of the MWE as a parameter. The storage management product then schedules the process requested. A limitation of this interface is that it does not provide support for data set name patterns. Only explicitly named data sets can be processed.

DB2 also uses ARCGIVER, asynchronously restoring data spaces that have been migrated while continuing to process SQL requests from other sources.

Overview

Appendix G: DFSMShsm to CA Disk Conversion

This section contains the following topics:

Conversion Considerations

Read this Appendix carefully, before attempting DFSMShsm to CA Disk conversion. Several sections on strategy and diagrams on process flow and terminology differences can help in your conversion. We have included a brief section on CA Vantage and its possible use to aid you in your conversion.

Conversion Strategy

You do not need to convert all your DFSMShsm data immediately to start using CA Disk. However, any CA Disk system that shares the Files Data Set or Files Data Base must be at r12 QO85050 or higher before starting. You can run all of the conversion process in simulation for testing and validation. Once you have installed and validated CA Disk, begin the conversion from DFSMShsm to CA Disk.

How DFSMShsm to CA Disk Conversion Process Works

Understanding the DFSMShsm to CA Disk Conversion process can help you customize the process for your organization or troubleshoot issues.

To convert DFSMShsm to CA Disk:

1. Prepare all CA Disk jobs for the daily/weekly/monthly process.
2. Start production with CA Disk.
3. Review Conversion JCL Members.
4. Start the Conversion.

5. Verify the Conversion.

Convert DFSMShsm to CA Disk

To convert DFSMShsm to CA Disk:

1. Prepare all CA Disk jobs for the daily/weekly/monthly process. (All functions with CA Disk software instead DFSMShsm.)
2. Stop all DFSMShsm-tasks except AUTOMATIC RECALL.
3. Review and submit the CREATEFD member.

   This job creates records from the BCDS and MCDS Control Data Sets and loads them into your CA Disk Files Data Set. The job creates your base Files Data Set used in production. CA Disk can do record maintenance over both DFSMShsm and CA Disk data.

   Command Syntax

   ```
   CREATEFD SGIFOFFL=
   ```

   SGIFOFFL=

   (Optional). When the SMS data set source volume is offline and you cannot determine a Storage Group, use this command to assign a new storage group.

   After CREATEFD has executed, you can submit RPTJCL00, RPTJCL01, RPTJCL02, RPTJCL05, and RPTJCL06 to verify that the data sets are ready for conversion.
4. Update the CA Disk Sysparms member for the DSNDLPEX user exit for a value of ADSHC010. The system issues a DELETE to DFSMShsm when a DFSMShsm pseudo DSNINDEX record is deleted. This exit cleans up and frees the resources that are used in the BCDS/MCDS before being converted.

   Update the CA Disk Sysparms member for the RESCRNEX user exit for a value of ADSHC011. When the system finds a DFSMShsm pseudo DSNINDEX record during restore, this exit issues a RECOVER/RECALL to DFSMShsm to restore the data set.
5. Update the CA Disk Sysparms member for the RESCRNEX user exit for a value of ADSHC011. When a DFSMShsm pseudo DSNINDEX record is found during restore, this exit issues a RECALL to DFSMShsm to restore the data set.
6. Update the CA Disk Sysparms member for the DINXUFEX user exit for a value of ADSHC012. After conversion, this exit issues a delete to DFSMShsm to clean up and free the resources that are used in the BCDS/MCDS.
7. Start the daily work with CA Disk using the Files Data Set created in Step 3.

8. Review the HSMGEN2 member.

   This job sorts the BCDS records created in Step 3 into VOLSER and file sequence number. The job creates the DFSMShsm and CA Disk commands in VOLSER sequence. The commands are bracketed with comments indicating DFSMShsm VOLSER and space to recover the data sets on the volume. The JCL4ML1 DD statement points to the model JCL used for data sets stored on ML1 and the JCL4ML2 to the model for ML2. Update these models to suit the requirements of your installation (members HSMCONB1 and HSMCONB2 of the INSTALL library). A SYSIN command controls the number of DFSMShsm volumes that are selected for processing.

   Command Syntax

   ```
   CONVERT     VOLCOUNT=nn,SIM,ML=,NEWHLQ=,STORGRP=,SMSVOL=,
               SMSUNIT=,NSMSVOL=,NSMSUNIT=,BYPERRORS=
   ```

   VOLCOUNT=*nn*

   Specify a one- to two-digit number indicating the number of DFSMShsm volumes to be converted. 99 indicates you want to convert all of the volumes.

   SIM

   Specifies a simulation run. In simulate mode, CA Disk produces messages and reports as if processing had taken place in LIVE mode, but no data sets are altered.

   To execute the JCL in LIVE mode after verifying that the results from the SIMULATE run are as expected, remove this parameter.

   ML=

   Specify either 1 or 2 to select either ML1 or ML2 for conversion. Default is 2.

   NEWHLQ=

   A high-level qualifier to be assigned to each data set selected. Default is H2DCON.

   STORGRP=

   Enter the Storage Group name that contains the volumes where the data sets to select reside. If the backup source volume no longer exists, add it to the proper Storage Group definition for this parameter to be applicable. Optionally, you can specify the SGIFOFFL= parameter in the CREATEFD JCL to assign a new Storage Group to the backed-up data set.

   SMSVOL=

   Specify the target VOLSER that the conversion process uses to retry the RECOVER of SMS-managed data sets in case of failure to RECOVER to the original VOLSER.

   SMSVOL= can only be specified together with SMSUNIT=.

SMSUNIT=

Enter the unit type to be used with the SMSVOL= parameter. The supported values are 3380, 3390, and 9345.

NSMSVOL=

Specify the target VOLSER that the conversion process uses to retry the RECOVER of non-SMS managed data sets in case of failure to RECOVER to the original VOLSER.

NSMSVOL= can only be specified together with NSMSUNIT=.

NSMSUNIT=

Enter the unit type to be used with the NSMSVOL= parameter. The supported values are 3380, 3390, and 9345.

BYPERRORS=

Specify one of VSAM, NONVSAM, ALL or NONE to bypass the error and continue processing if errors are found for any of them to stop processing after the first error. This parameter is optional. Default is NONE.

9. Review the HSMGEN1 member.

   This job sorts the MCDS records, created in Step 3, into VOLSER and file sequence number. This job creates the DFSMShsm and CA Disk commands in VOLSER sequence. Messages are printed indicating the DFSMShsm VOLSER and the space to recover the data sets on the volume. The DD statement JCL4ML1 points to the model JCL used for data sets stored on ML1 and JCL4ML2 to the model for ML2. Update these models to suit the requirements of your installation (members HSMCONM1 and HSMCONM2 of the INSTALL library). A SYSIN command controls the number of DFSMShsm volumes that are selected for processing.

   Decrease processing time by splitting the MCDS records into separate data sets and running multiple HSMGEN1 jobs concurrently. Multiple jobs mean more space that is used and more ML1 and ML2 units that are used during processing.

   Two methods exist to run multiple jobs:

   a. Create separate jobs that are based on the full VOLSER list and then cut it into pieces.

      We executed two jobs in parallel. Each job creates a number of records (15000 in the sample) that are treated and submitted with the next job. The samples are named HSMMIGnn. You can find them in our CCUWJCL library.

      Create parallel JCLs which manage the VOLSER one by one. Each new submit removes the first VOLSER on the list. Each JCL submitted once automatically starts the other. Only HSMCOV01 creates the complete VOLSER list. Then a loop starts with HSMCOV02 and HSMCOV03. The loop continues until there are no more tapes to convert.

Command Syntax

```
CONVERT VOLCOUNT=nn,SIM,ML=,STORGRP=,SMSVOL=,
        SMSUNIT=,NSMSVOL=,NSMSUNIT=,BYPERRORS=
```

VOLCOUNT=*nn*

Specify a one- to a two-digit number indicating the number of DFSMShsm volumes to be converted. Specifying 99 converts all volumes.

SIM

Specifies a simulation run. In simulate mode, CA Disk produces the normal messages and reports as if processing had taken place in LIVE mode. CA Disk does not alter any data sets.

After verifying that the results from the SIMULATE run are as expected, remove this parameter to execute the JCL in LIVE mode.

ML=

Specify either 1 or 2 to select either ML1 or ML2 for conversion. Default is 2.

STORGRP=

Enter the Storage Group name that contained the volumes where the data sets to select resided.

SMSVOL=

If RECALL to the original VOLSER fails, specify the target VOLSER that the conversion process uses to retry the RECALL of SMS managed data sets.

SMSVOL= can only be specified together with SMSUNIT=.

SMSUNIT=

Enter the unit type to be used with the SMSVOL= parameter. The supported values are 3380, 3390, and 9345.

NSMSVOL=

If RECALL to the original VOLSER fails, specify the target VOLSER that the conversion process uses to retry the RECALL of SMS managed data sets.

NSMSVOL= can only be specified together with NSMSUNIT=.

NSMSUNIT=

Enter the unit type to be used with the NSMSVOL= parameter. The supported values are 3380, 3390, and 9345.

BYPERRORS=

Specify one of VSAM, NONVSAM, ALL, or NONE to bypass the error and continue processing if errors are found for any of them to stop processing after the first error. This parameter is optional. The default is NONE.

10. Evaluate the results in SIM mode from the HSMGEN1 and HSMGEN2. Using the SIM mode, you can see how much space you require to recover the data sets on the DFSMShsm volumes.

    - The allocated tracks from each data set are added and the total is printed at the end of each processed volume.

      Specify one of VSAM, NONVSAM, ALL or NONE to bypass an error and continue processing if errors are found for any of them to stop processing after the first error. This parameter is optional and NONE is the default. To process the backup versions of a data set, a pseudo HLQ is needed because the data set can exist. Although a new HLQ is used to allocate the data set, all of the information CA Disk requires from the BSDS is maintained in the DSNINDEX record. Only the ARCHVOL data is updated. Since a new HLQ is used, your allocation product can use this HLQ to direct the DFSMShsm recovery to a stand-alone volume group. Doing this relieves any space constraints in the original groups.

    - Submit both jobs in LIVE mode as many times as required to convert the backups and archives from DFSMShsm. Once the commands for a volume have been generated, it is marked as converted so you can rerun the jobs as many times as needed.

11. Verify the conversion by executing RPTJCLA3, RPTJCL03, and RPTJCL04 report jobs. They are based on BCDS reporting compared with Files data set entries.

12. You can shut down HSMGEN1 or HSMGEN2 conversion jobs early and gracefully by using HSMREPLY. A graceful shutdown allows the jobs to complete all current work in progress. The jobs then produce the normal reports of error conditions and the work is successfully completed before early job termination.

    With HSMREPLY active, message “4570 REPLY 'END' TO TERMINATE HSMGEN# PROCESSING” is issued to the operator at the beginning of the job. This message remains on the console until the job finishes. Replying END stops the job early.

MIGRATION Program Details

The following are the migration programs:

ADSHC002

Generate BACKUP-DSNINDEX and ARCHVOLS based on BCDS.

ADSHC003

Generate ARCHIVE-DSNINDEX and ARCHVOLS based on MCDS.

ADSHC004

Generate RECOVER and BACKUP JCL.

ADSHC005

Generate RECOVER and ARCHIVE JCL.

ADSHC010

Generate HSM-DELETE commands. (DSNDPLEX for IXMAINT.)

ADSHC011

Generate HSM-RECALL commands. (RESCRNEX for Restores.)

ADSHC012

Generate HDEL commands after conversion to CA Disk. (DINXUFEX for ARCHIVE.)

ADSHC041

Issue DFSMShsm ARCHRCOV macro and appropriate messages for ADSHC004.

ADSHC142

Update DSNDINEX records for ARCTMDTE.

ADSHC457

Update REPORTS DD headers.

RPTCA001

Generate list from BCDS report.

RPTCA002

Match list from FDS/FDB with BCDS report.

RPTCA003

List dsns that failed to HRECOVER (ARCHRCOV.)

RPTCA004

List dsns that failed to HRECOVER (ARCHRCOV) and have missing DSNINDEX records.

JCL Details

The following are the JCL details:

CREATEFD

Sample JCL

HSMCONB1

Skeleton JCL

HSMCONB2

Skeleton JCL

HSMCONM1

Skeleton JCL

HSMCONM2

Skeleton JCL

HSMGEN1

Sample JCL - ADSHC005

HSMGEN2

Sample JCL - ADSHC004

RPTJCL00

Lists FDS/FDB (dsn, H2DBKP/H2DMIG/volser, arctime, arcdate and flag1) and creates BCDS and MCDS reports to be used by RPTJCL01 and RPTJCL05.

RPTJCL01

Lists the input BCDS copy (dsn, archvol, source volume and backup date).

RPTJCL02

Match RPTJCL00/01 lists. Generate complete input BCDS copy status NOT OK (not ready for conversion) lists.

RPTJCLA3

Creates BCDS report to be used by RPTJCL03.

RPTJCL03

List of dsns that failed to HRECOVER (ARCHRCOV). Dsn, RC=, version and volser.

RPTJCL04

List of data sets that failed to HRECOVER (ARCHRCOV) that are NOT OK not ready to transfer) due to missing DSNINDEX records.

RPTJCL05

Lists the input MCDS copy (dsn, archvol, source volume and backup date).

RPTJCL06

Match RPTJCL00/05 lists. Generate complete input MCDS copy status / NOT OK (not ready for conversion) lists.

Illustration—BCDS Entries

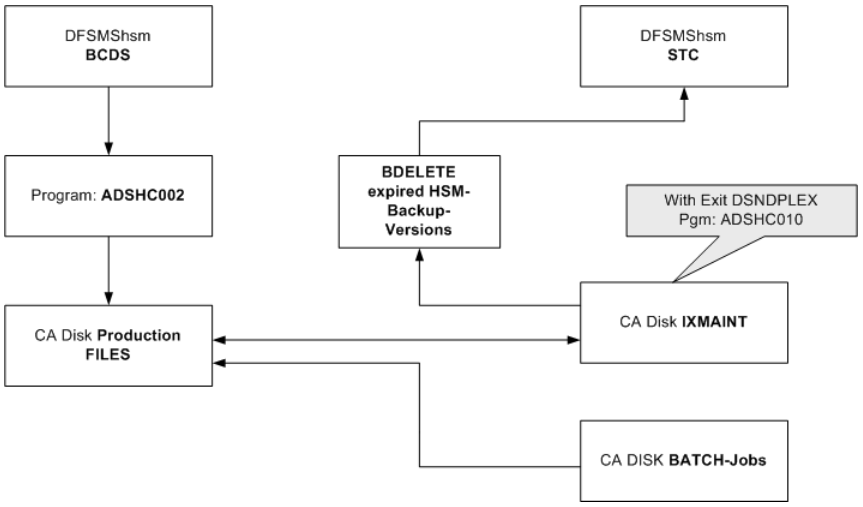

Illustration—MCDS Entries

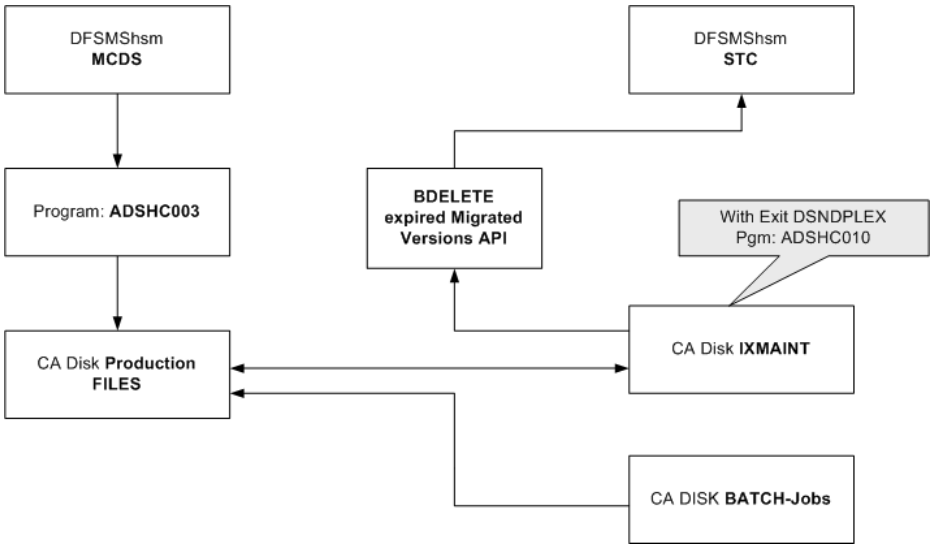

Illustration—Recover and Archive

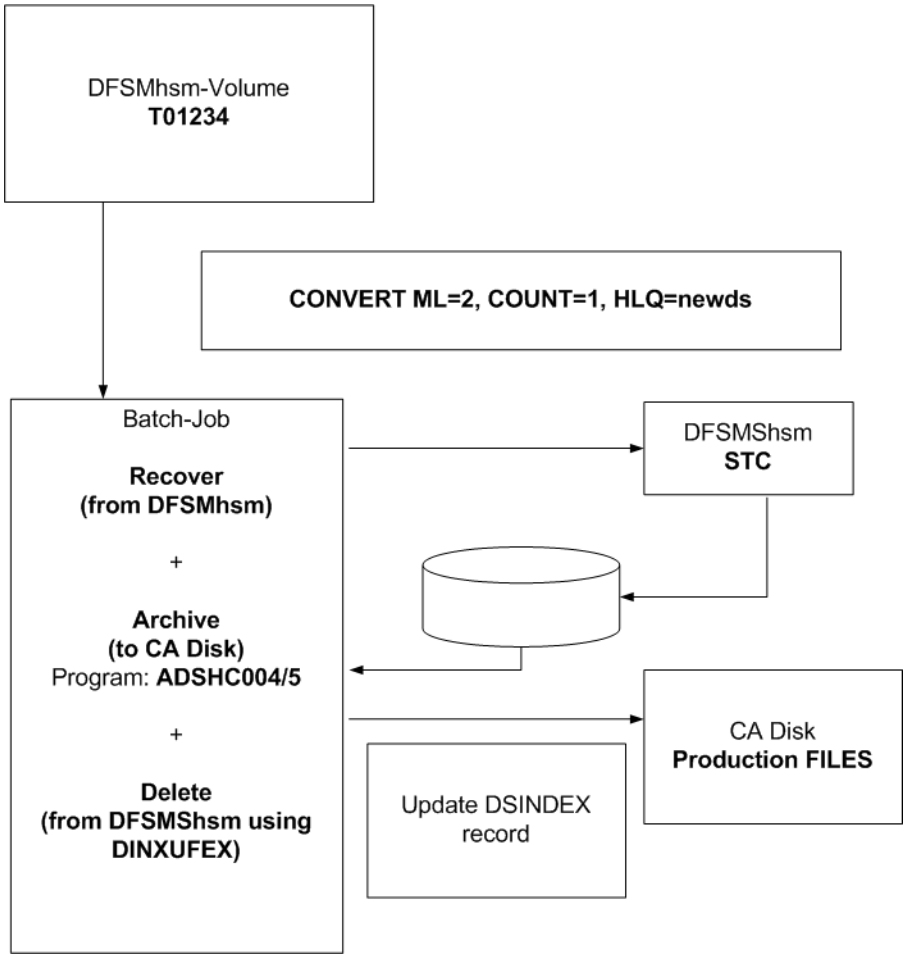

Sample Sysouts

The following screen shows a sample SYSOUT from HSMGEN2 MSGPRINT DD:

```
ADSHC004 4503 PROCESSING STARTED FOR VOLUME ---PIT001 DEVICE TYPE 3030200F
ADSHC041 4523 DATA SET USER.IDEKE01.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.XE49.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.IDEKE01.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.XE49.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.IDEKE01.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.XE49.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.IDEKE01.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.XE49.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC004 4504 PROCESSING COMPLETE FOR VOLUME --PIT001 - TRACKS REQUIRED ---- 423
ADSHC004 4503 PROCESSING STARTED FOR VOLUME ---PIT002 DEVICE TYPE 3030200F
ADSHC041 4523 DATA SET USER.IDEKE01.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC041 4523 DATA SET USER.XE49.TESTCAT IS AN INTEGRATED CATALOG FACILITY CATALOG SO IT WILL BE BYPASSED.
ADSHC004 4504 PROCESSING COMPLETE FOR VOLUME --PIT002 - TRACKS REQUIRED ---- 92
ADSHC004 4513 NUMBER OF VOLUMES PROCESSED ---------- 2
ADSHC004 4505 TOTAL TRACKS REQUIRED TO PROCESS THESE VOLUMES ---------- 515
```

The following screen shows a sample SYSOUT from HSMGEN2 REPORTS DD:

```
2009.236  AUG 24, 2009                LIST OF DATA SETS SUCCESSFULLY RECOVERED BY DFSMSHSM         PAGE    1
MONDAY     2.42 PM                                                                              CA DISK r12

ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSMIG                          VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN0.NONSMS              VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN5.NONSMS              VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSMIG                          VERS=001 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSNC                           VERS=001 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSNG1.SMS                VERS=003 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN0.VERS                VERS=011 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSMIG                          VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN0.NONSMS              VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDS3                            VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN2.SMS                 VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN8.NONSMS.OFFLINE      VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSNG2.SMS                VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSMIG                          VERS=003 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSNC                           VERS=003 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDS3                            VERS=003 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSNG1.SMS                VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSNC                           VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN1.NONSMS.OFFLINE      VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN6.SMS                 VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSNG2.SMS                VERS=003 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN0.VERS                VERS=012 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDSNC                           VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN1.NONSMS.OFFLINE      VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN3.NONSMS              VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN9.NONSMS              VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDS2                            VERS=003 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN2.SMS                 VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN0.VERS                VERS=010 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN0.VERS                VERS=013 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.TEST.KSDS2                            VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN4.SMS.OFFLINE         VERS=001 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN3.NONSMS              VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN8.NONSMS.OFFLINE      VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN0.VERS                VERS=014 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN5.NONSMS              VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN9.NONSMS              VERS=002 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN6.SMS                 VERS=004 TEST=YES
ADSHC041 4500 ARCHRCOV DSN=GARJO24.P810913.TEST.DSN7.SMS.OFFLINE         VERS=001 TEST=YES
ADSHC004
ADSHC004
ADSHC004 TOTAL NUMBER OF DATA SETS RECOVERED           39
```

The following screen shows a sample SYSOUT from HSMGEN1 MSGPRINT DD:

```
ADSHC005 4503 PROCESSING STARTED FOR VOLUME ---PIT001 DEVICE TYPE 3030200F
ADSHC005 4504 PROCESSING COMPLETE FOR VOLUME --PIT001 - TRACKS REQUIRED ---- 2
ADSHC005 4513 NUMBER OF VOLUMES PROCESSED ---------------------- 1
ADSHC005 4505 TOTAL TRACKS REQUIRED TO PROCESS THESE VOLUMES ---------------- 2
```

The following screen shows a sample SYSOUT from HSMGEN1 REPORTS DD:

```
2009.236  AUG 24, 2009                  LIST OF DATA SETS SUCCESSFULLY RECALLED BY DFSMSHSM         PAGE    1
MONDAY     4.14 PM                                                                                CA DISK r12


ADSHC005 4501 ARCHRCAL DSN=GARJO24.P810913.TEST.DSNG1.SMS               TEST=YES
ADSHC005 4501 ARCHRCAL DSN=GARJO24.P810913.TEST.DSNG2.SMS               TEST=YES
ADSHC005
ADSHC005
ADSHC005 TOTAL NUMBER OF DATA SETS RECALLED             2
```

Converting DFSMShsm Tasks to CA Disk Jobs

Using the CA Vantage Graphical User Interface (GMI) can reduce the time needed to analyze the DFSMShsm tasks. There are a number of objects that can be presented for review. The most useful of these would be the Active DFSMShsm Tasks and Patches.

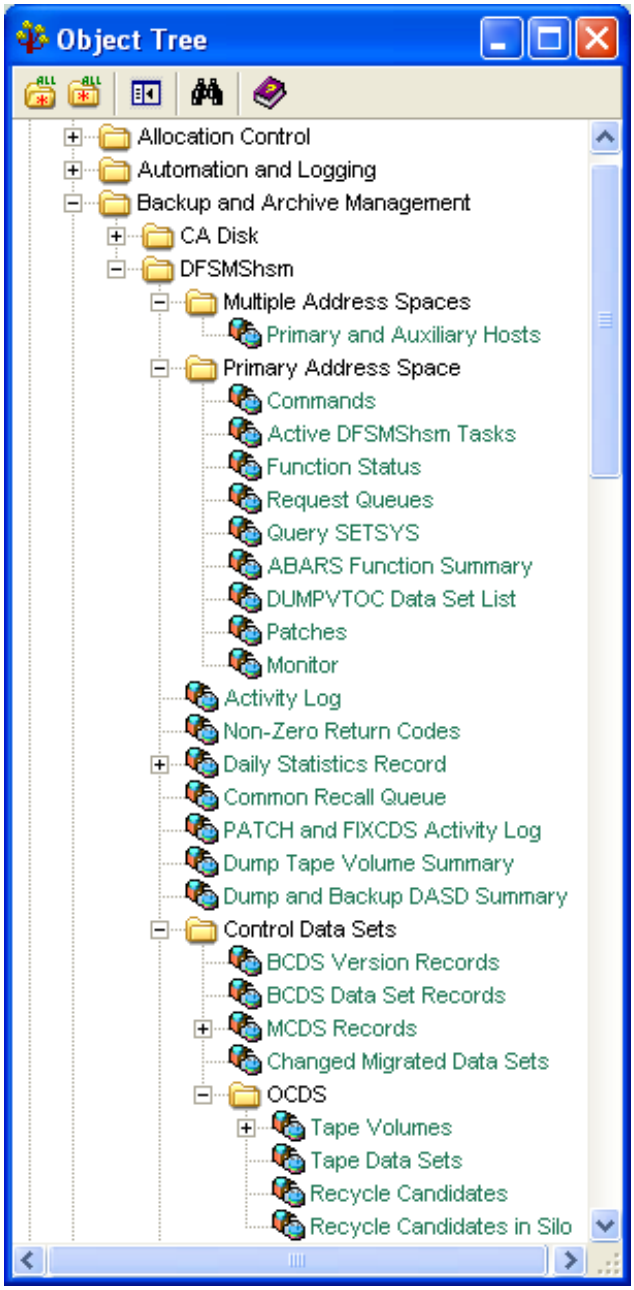

Comparing CA Disk Terminology to DFSMShsm

The following is a list of DFSMShsm terms, the DFSMShsm definition, and the equivalent CA Disk term:

DFSMShs Term	CA Disk Term	Description
Aggregate backup	Backup, including XCOPY	Copy data and control information from a user-defined group of data sets.
Alternate tape volume	Duplicate copy tape	Copy of original tape volume, created simultaneously or subsequent copy process.
Auto-recall	Auto-Restore (DMSAR)	Recall initiated automatically, transparent to the requesting task.
Backup	Backup	Copy a data set residing on a primary disk to a backup volume.
Backup control data set (BCDS)	Files data set (DSNINDEX)	Control file containing information about data set backup versions.
Expiration	IXMAINT	The removal of a user data set from a primary volume or the deletion of a migrated data set version.
Full Volume Dump	VBACKUP	Process using DFSMSdss that backs up the allocated space on a disk volume.
Level 0	Primary disk (DASD)	Online disk volume containing data directly accessible by the user.
Level 1	Disk archive	Disk volume under DFSMShsm control containing data sets migrated from Level 0.
Level 2	Tape	Volume (usually tape) under DFSMShsm control containing data sets migrated from Level 0 or level 1.

Migrate	Archive	Move a cataloged data set from primary disk to ML1 or ML2.
Migration control data set (MCDS)	Files data set (DSNINDEX)	Control file containing information about data sets that scratched after backup.
Offline control data set (OCDS)	Files data set (ARCHVOL)	Control file containing information about migration and backup tapes and about each data set on these tapes.
Recall	Restore	Moving a migrated data set back to Level 0 from Level 1 or 2.
Recover	Restore/Recover	Recovering a single or multiple data sets to level 0 from backup volumes.
	Recycle	Copies all valid data on a tape to a tape spill backup or ML2 volume.