

A catalog is a data set that contains information about other data sets. It provides users with the ability to locate a data set by name, without knowing where the data set resides. By cataloging data sets, your users will need to know less about your storage setup.

- Thus, data can be moved from one device to another, without requiring a change in JCL DD statements that refer to an existing data set.

Cataloging data sets also simplifies backup and recovery procedures. Catalogs are the central information point for data sets; all data sets must be cataloged. In addition, all SMS-managed data sets must be cataloged.

DFSMS allows you to use catalogs for any type of data set or object. Many advanced functions require the use of catalogs, for example, the storage management subsystem.

Multiple user catalogs contain information about user data sets, and a single master catalog contains entries for system data sets and user catalogs.

In z/OS, the component that controls catalogs is embedded in DFSMSdfp and is called **Catalog Management**. Catalog Management has one address space for itself named **Catalog Address Space (CAS)**. This address space is used for buffering and to store control blocks, together with code.

The modern catalog structure in z/OS is called the **integrated catalog facility (ICF)**. All data sets managed by the storage management subsystem (SMS) must be cataloged in an ICF catalog.

Most installations depend on the availability of catalog facilities to run production job streams and to support online users. For maximum reliability and efficiency, catalog all permanent data sets and create catalog recovery procedures to guarantee continuous availability in z/OS.

The integrated catalog facility (ICF) structure
An integrated catalog facility (ICF) catalog is a structure that replaced the former MVS CVOL² catalog. As a catalog, it describes data set attributes and indicates the volumes on which a data set is located.

ICF catalogs are allocated by the catalog address space (CAS), a system address space for the DFSMSdfp catalog function.

A catalog consists of two separate kinds of data sets:

- A basic catalog structure (BCS) - the BCS can be considered the catalog.

The basic catalog structure (BCS) is a VSAM key-sequenced data set. It uses the data set name of entries to store and retrieve data set information. For VSAM data sets, the BCS contains volume, security, ownership, and association information.

For non-VSAM data sets, the BCS contains volume, ownership, and association information. When we talk about a catalog, we usually mean the BCS.

A VSAM volume data set (VVDS) can be considered an extension of the volume table of contents (VTOC).

The VVDS can be considered an extension of the volume table of contents (VTOC). The

VVDS is volume-specific, whereas the complexity of the BCS depends on your definitions.

The relationship between the BCS and the VVDS is many-to-many. That is, a BCS can point to multiple VVDSs and a VVDS can point to multiple BCSs.

The VSAM volume data set (VVDS) is a data set that describes the characteristics of VSAM and system-managed data sets residing on a given DASD volume; it is part of a catalog.

The VVDS contains VSAM volume records (VVRs) that hold information about VSAM data sets residing on the volume.

The VVDS also contains non-VSAM volume records (NVRs) for SMS-managed non-VSAM data sets on the volume. If an SMS-managed non-VSAM data set spans volumes, then only the first volume contains an NVR for that data set.

The system automatically defines a VVDS with 10 tracks primary and 10 tracks secondary space, unless you explicitly define it.

Basic catalog structure (BCS)
The basic catalog structure (BCS) is a VSAM key-sequenced data set. It uses the data set name of entries to store and retrieve data set information.

For VSAM data sets, the BCS contains volume, security, ownership, and association information.

For non-VSAM data sets, the BCS contains volume, ownership, and association information.

NOTE: The BCS portion of the ICF catalog contains the static information about the data set, the information that rarely changes.

Every catalog consists of one BCS and one or more VVDSs.

A BCS does not "own" a VVDS; that is, more than one BCS can have entries for a single VVDS.

Every VVDS that is connected to a BCS has an entry in the BCS.

For example, the figure on rftg shows a possible relationship between a BCS and three VVDSs on three disk volumes.

For non-VSAM data sets that are not SMS-managed, all catalog information is contained within the BCS and for other types of data sets, there is other information available in the VVDS.

BCS structure
The BCS contains the information about where a data set resides.

- That can be a DASD volume, tape, or other storage medium.

- Related information in the BCS is grouped into logical, variable-length, spanned records related by key. The BCS uses keys that are the data set names (plus one character for extensions).

One control interval can contain multiple BCS records. To reduce the number of I/Os necessary for catalog processing, logically-related data is consolidated in the BCS.

A catalog can have data sets cataloged on any number of volumes. The BCS can have as many as 123 extents on one volume.

One volume can have multiple catalogs on it. All the necessary control information is recorded in the VVDS residing on that volume.

Master catalog (MCAT)
A configuration of catalogs depends on a master catalog. A master catalog has the same structure as any other catalog. What makes it a master catalog is that all BCSs are cataloged in it, as well as certain data sets called system data sets (for instance, SYS1.LINKLIB and other "SYS1" data sets).

Catalogs offer advantages including improved performance, capability, usability, and maintainability. The catalog information that requires the most frequent updates is physically located in the VVDS on the same volume as the data sets, thereby allowing faster access.

A catalog request is expedited because fewer I/O operations are needed. Related entries, such as a cluster and its alternate indexes, are processed together.

The VSAM volume data set (VVDS)
The VSAM volume data set (VVDS) contains additional catalog information (not contained in the BCS) about the VSAM and SMS-managed non-VSAM data sets residing on the volume where the VVDS is located. Every volume containing any VSAM or any SMS-managed data sets must have a VVDS on it. The VVDS acts as a kind of VTOC extension for "certain types" of data sets.

A VVDS can have data set information about data sets cataloged in distinct BCSs.

1. Data Facility Product (DFSMSdfp) is the heart of the storage management subsystem. It keeps track of all data and programs managed within z/OS.

2. CVOL is the previous generation of MVS Catalogs

3. Volume Table Of Contents, or VTOC, is a data structure, that provides a way of locating the data sets that reside on a particular disk volume.


4. Control Interval is a VSAM grouping of records similar to blocking.

5. Alternate Index is a secondary key structure used to access a base VSAM file.

Catalogs

196
2817

ICF structure 1




ICF structure 2



Basic catalog structure (BCS)

Index	Data component
DSNAME1	DSNAME1 ... VOL001
DSNAME2	DSNAME2 ... VOL002
DSNAME3	DSNAME3 ... VOL001
DSNAME4	DSNAME4 ... VOL003
DSNAME5	DSNAME5 ... VOL003
...	...
DSNAMEi	DSNAMEi ... VOL002

The BCS is a VSAM KSDS



Note: For SMS-managed data sets, JOBCAT and STEPCAT DD statements are not allowed and cause a job failure. Also, they are not suggested even for non-SMS data sets, because they can cause conflicted information. Therefore, do not use them and keep in mind that they have been phased out since z/OS V1R8.

Entry types in the VVDS

- There are three types of entries in a VVDS:
 - VSAM volume control records (VVCR)
 - First logical record in a VVDS
- Contain information for management of DASD space and the names of the BCSs that have data sets on the volume
- VSAM volume records (VVR)
 - Contain information about a VSAM data set residing on the volume
 - Number of VVRs varies according to the type of data set and the options specified for the data set
 - Also included are data set characteristics, SMS data, extent information
 - There is one VVR describing the VVDS itself
- Non-VSAM volume record (NVR)
 - Equivalent to a VVR for SMS-managed non-VSAM data sets
 - Contains SMS-related information

VVDS characteristics

- The VVDS is a VSAM entry-sequenced data set (ESDS) that has a 4 KB control interval size.
- The hexadecimal RBA of a record is used as its key or identifier.
- A VVDS is recognized by the restricted data set name:
 - SYS1.VVDS.vvolser
 - vvolser is the volume serial number of the volume on which the VVDS resides.
- You can explicitly define the VVDS using IDCAMS, or it is implicitly created after you define the first VSAM or SMS-managed data set on the volume.

VVDS SPACE keyword

- Prior to z/OS V1R7, the default space parameter is TRACKS(10,10), which could be too small for sites that use custom 3390 volumes (the ones greater than 3390-9). With z/OS V1R7, there is a new VVDS SPACE keyword of the F CATALOG command, as follows:

F CATALOG, VVDS SPACE(primary,secondary)

- An explicitly defined VVDS is not related to any BCS until a data set or catalog object is defined on the volume. As data sets are allocated on the VVDS volume, each BCS with VSAM data sets or SMS-managed data sets residing on that volume is related to the VVDS.
- VVDS SPACE indicates that the Catalog Address Space are to use the values specified as the primary and secondary allocation amount in tracks for an implicitly defined VVDS. The default value is ten tracks for both the primary and secondary values. The specified values are preserved across a Catalog Address Space restart, but are not preserved across an IPL.

Catalogs by function

- By function, the catalogs (BCSs) can be classified as master catalog and user catalog. A particular case of a user catalog is the volume catalog, which is a user catalog containing only tape library and tape volume entries.
- There is no structural difference between a master catalog and a user catalog. What sets a master catalog apart is how it is used, and what data sets are cataloged in it. For example, the same catalog can be master in one z/OS and user in the other z/OS.

The master catalog

- Each system has one active master catalog. One master catalog can be shared between various MVS images. It does not have to reside on the system residence volume (the one that is IPLed).
- The master catalog for a system must contain entries for all user catalogs and their aliases that the system uses. Also, all SYS1 data sets must be cataloged in the master catalog for proper system initialization.

Identifying the master catalog for IPL

- At IPL, you must indicate the location (volser and data set name) of the master catalog. This information can be specified in one of two places:
 - SYS1.NUCLEUS member SYSCATxx (default is SYSCATLG)
 - SYS1.PARMLIB/SYsn.IPLPARM member LOADxx. This method is recommended.

Determine the master catalog on a running system

- You can use the IDCAMS LISTCAT command for a data set with a high-level qualifier (HLQ) of SYS1 to determine the master catalog on a system. Because all data sets with an HLQ of SYS1 are to be in the master catalog, the catalog shown in the LISTCAT output is the master catalog.SYSRES

User catalogs

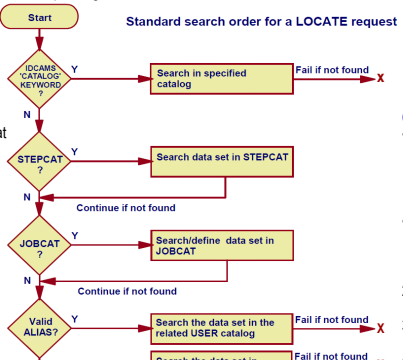
- The difference between the master catalog and the user catalogs is in the function. User catalogs are to be used to contain information about your installation cataloged data sets other than SYS1 data sets. There are no set rules as to how many or what size to have; it depends entirely on your environment.
- Cataloging data sets for two unrelated applications in the same catalog creates a single point of failure for them that otherwise might not exist. Assessing the impact of outage of a given catalog can help to determine if it is too large or can impact too many applications.

Using aliases

- Aliases are used to tell catalog management which user catalog your data set is cataloged in.
- First, you place a pointer to an user catalog in the master catalog through the IDCAMS DEFINE UCAT command. Next, you define an appropriate alias name for a user catalog in the master catalog. Then, match the high-level qualifier (HLQ) of your data set with the alias. This identifies the appropriate user catalog to be used to satisfy the request.

- In the right figure, all data sets with an HLQ of PAY have their information in the user catalog UCAT1 because in the master catalog there is an alias PAY pointing to UCAT1.

- The data sets with an HLQ of DEPT1 and DEPT2, respectively, have their information in the user catalog UCAT2 because in the master catalog there are aliases DEPT1 and DEPT2 pointing to UCAT2.



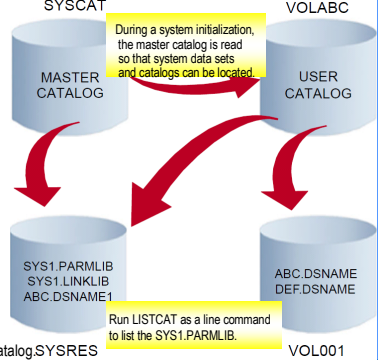
Catalog search order:

- LOCATE** is SVC 26 that calls catalog management asking for a catalog set name search. Most catalog searches are based on catalog aliases. Alternatives to catalog aliases are available for directing a catalog request, specifically the JOBCAT and STEPCAT DD statements; the CATALOG parameter of access method services; and the name of the catalog. JOBCAT and STEPCAT are no longer allowed beginning with z/OS V1R7.
- Search order for catalogs for a data set define request**
For the system to determine where a data set is to be cataloged, the following search order is used to find the catalog:

1. Use the catalog named in the IDCAMS CATALOG parameter, if coded.
2. If the data set is a generation data set, the catalog containing the GDG base definition is used for the new GDS entry.
3. If the high-level qualifier is a catalog alias, use the catalog identified by the alias or the catalog whose name is the same as the high-level qualifier of the data set.
4. If no catalog has been identified yet, use the master catalog.

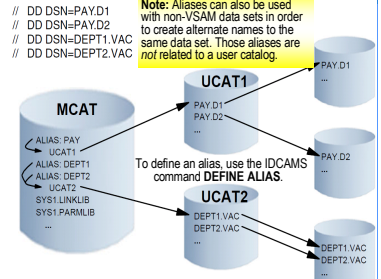
Data Facility Storage Management Subsystem (DFSMS™) is a software suite that automatically manages data from creation to expiration. DFSMS provides allocation control for availability and performance, backup/recovery and disaster recovery services, space management, tape management, and reporting and simulation for performance and configuration tuning.

Important: To minimize update activity to the master catalog, and to reduce the exposure to breakage, only SYS1 data sets, user catalog connector records, and the aliases pointing to those connectors are to be in the master catalog.



Run LISTCAT as a line command to list the SYS1.PARMLIB.

NOTE: Aliases can also be used with non-VSAM data sets in order to create alternate names to the same data set. Those aliases are not related to a user catalog.



To define an alias, use the IDCAMS command DEFINE ALIAS.