

# CATALOG BASICS SERIES

By Blair Svihra

## Catalog Components Part 1



## Catalog Basics | Part 1 – Catalog Components

Today's modern day catalog structure consists of three components:

- the Basic Catalog Structure (BCS),
- the VSAM Volume Data Set (VVDS), and
- the Volume Table of Contents (VTOC).

It could be argued that the catalog environment could be expanded to include a discussion including any one of the five major Tape Management Catalogs (TMCs) and/or IBM's Hierarchical Storage Manager (HSM) offering, but these topics can be discussed in detail in a later whitepaper. A catalog allows access to a data set without having to know its physical location. Let's look at all three components and how they are related.

The BCS is primarily a data set locator. Vital information needed to open and process a data set (AMDSB information, extent information, data set attributes, and so on) is in the VTOC and VVDS on the volume where physical DASD data sets reside. BCS information includes,

the data set name,

- volume and device type the components reside on,
- ownership information,
- data set security information (if applicable),
- possible VVDS and VTOC record pointers for each component, and
- System Managed Storage (SMS) class information (if the data set is SMS managed).

There will always be four entries in a BCS after it's created:

- the three records that describe the catalog:
  - cluster (self-describing record),
  - data and
  - index components, and
- a VVDS entry for the volume on which the catalog resides.

The BCS is a pure key sequenced data set (KSDS). It can be opened and processed as a normal KSDS outside of catalog management. A typical data set name can be up to 44 characters in length. The key length in all ICF catalogs is 45 bytes. The 45th byte is used for extension record processing within catalogs. Forty-four bytes are used for the data set name in the VTOC and VVDS. The cluster name of every ICF catalog is 45 bytes of binary zeros. This forces a catalog's self-describing record to a fixed location in the BCS – relative byte address (RBA) of 0. This allows system initialization to access a catalog directly by channel program I/O before the catalog address space has been initialized. The name specified in the IDCAMS 'DEFINE' is the name of the catalog's data component, not the cluster name. The index of a BCS is generated automatically and referred to as the CATINDEX. There are a few flavors of the index name that IBM has used since ICF catalogs were first introduced in the early 80's.

The VVDS is a special entry sequenced data set (ESDS). Its name, SYS1.VVDS.Vvolser, is a reserved name. The volser in the VVDS name must match the volser in the unit control block (UCB) for the device on which it resides. The VVDS contains:

- extent information,
- high allocated/high used values,
- complete AMDSB information,
- data set attributes,

## DINO WHITEPAPER: Catalog Components

- SMS class information,
- and so on for all VSAM data sets on a volume.

SMS managed nonVSAM data sets on a volume have entries in the VVDS, but their content mainly consists of SMS class information. SMS constructs reside in both the BCS and VVDS (duplication of information). The reason for this is discussed in a later whitepaper titled "Catalog Diagnostics." VSAM volume records (VVR's) describe VSAM components on a volume. NonVSAM volume records (NVR's) describe SMS managed nonVSAM entries on volume. NVR records will only be present on the first volume of a multi-volume SMS managed nonVSAM data set. NonVSAM nonSMS managed data sets on a volume do not have any entries in the VVDS, only the VTOC. VVCR, VVCM and VVCN are other control records that can be present in the VVDS. The VVCR is always located at RBA 0 and contains a space map detailing the physical layout of the VVDS. It also contains the names of the first 36 catalogs that have or have had data sets on the volume. VVCN records are created to describe additional catalogs that have or have had entries on the volume. Additional VVCM records can exist to map the physical layout of the VVDS if the size of the VVCR space map is not adequate to map the VVDS.

The VTOC has been around much longer than the VVDS. Volumes may or may not have VTOC indexes – SYS1.VTOCIX.volser to further enhance access to VTOC information by DADSM. Allocate one on your volume with IBM's ICKDSF and enable the index. There is an overlap of information that resides in the VTOC and VVDS for some components on a volume. For example, extent information for VSAM components that physically reside on a volume is recorded in both repositories. The VTOC acts as a space repository for the volume while the VVDS extent information is used to map the VSAM object when it's opened. If there is an extent discrepancy between the VVDS and the VTOC, overlays or missing data can occur. Periodic diagnostics on DASD volumes ensure this condition is detected before serious problems occur.