

The Catalog Search Interface

BY STEVE PRYOR

Storage administration often involves more than simply maintaining proper levels of device utilization. In many installations, the storage administrator wears more than one hat, including those of system programmer and all-around troubleshooter. If a system problem arises, particularly one whose solution requires information about the attributes of datasets or the contents of catalogs, the storage administrator is often called upon to use his knowledge, skills, and tools to attack the problem.

Sometimes obtaining information from ICF catalogs can be problematic. Even something as simple as obtaining a list of dataset names that match a particular naming convention can be difficult, since the IDCAMS LISTCAT function does not allow for much pattern masking. Determining the catalog status of a series of datasets (catalogued, uncatalogued, multi-volume, etc.) or their attributes (alias names, alternate index path names, stripe count, and so on) may require tedious manual examination of many pages of LISTCAT output.

Storage administrators or others who need to obtain dataset names or other information from the catalog have been forced to either read the catalog directly or make do with the limited functionality provided by the CAMLST LOCATE macro. Software vendors have long been able to take advantage of the direct access to catalogs provided by SVC 26 (CTGPL LOCATE), but this interface is complex, difficult to use, and is not documented.

Beginning (officially) with DFSMS version 1.4, the Catalog Search Interface (FMID H00B100) provides users with a way to retrieve catalog information based on criteria such as "wild card" matching for

dataset names. The Catalog Search Interface parameter list is actually a simplified version of the parameter list used by CTGPL LOCATE. The user provides some selection criteria for the type of entries to be located (cluster names, GDG names, alias names, etc.), and optionally, additional information that is to be returned for each entry, such as volume serial number, creation date, high-used RBA, and so on.

The ability of the Catalog Search Interface to use masking in its queries is an important facility. So important, in fact, that the Catalog Search Interface module IGGCSI00 was actually included in DFSMS version 1.2, when APAR OW09581 was introduced to provide support for dataset name masking in the HBACKDS, HMIGRATE, HRECALL, etc., commands in DFSMSHsm. Using CSI, the storage administrator can supply a key of 'PAYROLL.*.MAR1?.*.YR1998', for example, and obtain a list of all datasets that match the specified name. The CSI masking criteria is the same as that used within the DFSMS ACS routines for dataset name masking (i.e., '**' indicates zero or

more qualifiers, '*' indicates one qualifier or a partial qualifier, and % indicates a single character).

The Catalog Search Interface parameter list consists of three words pointed to by register one, that point to areas for the return code, selection criteria, and a work area into which CSI returns the requested information. The parameter list is shown in Figure 1. The selection criteria includes the types of catalog entry desired (such as 'A' for non-VSAM datasets, 'B' for GDGs, 'C' for cluster, and 'D' or 'I' for data or index components) along with the entry name mask. A catalog name may be specified, though the default is to search all the necessary catalogs in standard order.

Following the selection criteria, a list of up to 100 eight-character names of catalog fields for which information is to be returned can be specified. For example, specifying VOLSER, OWNERID, and ATTR1 in the field names would retrieve the volume serial number, owner, and flags for the SPEED, UNIQUE, and REUSE attributes for each dataset selected.

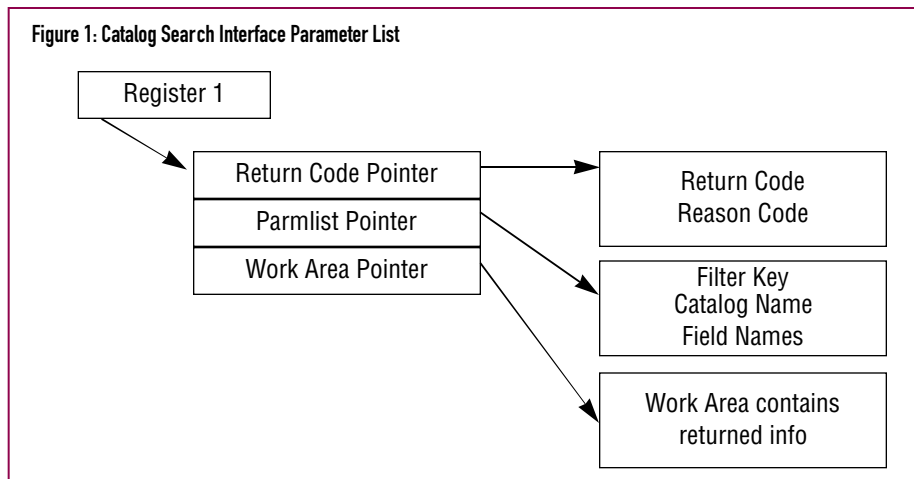
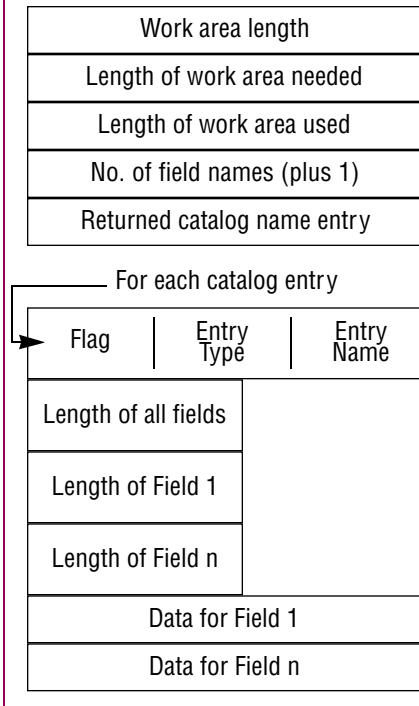


Figure 2: Data Returned From Catalog Search Interface




The returned information is placed into the work area, with the length of each set of returned fields preceding the data for each

entry. An entry for each catalog searched precedes data from the catalog. Since the actual amount of data that will be returned is unpredictable, the work area may not be large enough to hold all of the required information. In that case, a flag can be set in the CSI parameter list that indicates processing is to begin again at the point where it left off when the service is invoked again. The format of the work area is shown in Figure 2.

The Catalog Search Interface is also considerably faster than accessing the catalog via IDCAMS. Where a LISTCAT against a large catalog may take several hours, obtaining similar information via the Catalog Search Interface may take only minutes or seconds.

Selecting entries from ICF catalogs by key can be particularly useful for storage administrators who need to obtain a list of datasets that match some criteria. In the event of a disk drive failure, for example, it might be useful to know what datasets were catalogued to the volume that failed and thus might need to be restored or whose catalog entries might need to be removed. The Catalog Search Interface

includes just such a sample program, IGGCSIVS. This sample program, along with several others, and a REXX EXEC that invokes the Catalog Search Interface, are documented in Appendix D of the *DFSMS/MVS Version 1.4 Managing Catalogs* manual, SC26-4914.

Managing data centers at the end of the millennium means making the best use of all of the available information when planning or problem solving. The Catalog Search Interface provides a welcome new tool that allows storage administrators and others to obtain that information more easily. 

NaSPA member Steve Pryor has more than 15 years of experience in storage management, disaster recovery, software development, and technical support. Steve can be contacted via the Internet at pryor@atlanta.com.

©1998 Technical Enterprises, Inc. For reprints of this document contact sales@naspa.net.