
Adager's Database Therapy

Ken Paul

Adager Corporation

Sun Valley, Idaho 83353-3000 • USA

<http://www.adager.com>

The checking of the IMAGE root file and its relationship to the other datasets is done proactively by Adager every time it opens the database. In this article I will talk about the steps that you, the database user/administrator, should be taking to avoid structural problems within your IMAGE database and about the Adager functions that you can use to fix database problems.

As most IMAGE users are aware, IMAGE is a very efficient and reliable database management system (DBMS). IMAGE offers additional reliability and integrity through options such as Intrinsic Level Recovery (ILR) under MPE/V, dynamic rollback under MPE/iX and transaction logging under both platforms. Under MPE/iX, IMAGE is also integrated with the Transaction Manager (XM) to insure database integrity.

As a result, most of you have never encountered a broken chain or a damaged IMAGE database. This does not necessarily mean that your databases are 100% structurally sound. Running a complete diagnostic on the database helps in experiencing a warm, fuzzy feeling that things work ok.

Disc drive problems, system failures and operator errors are ways in which IMAGE databases can become damaged. The most common cause for database corruption, however, can usually be traced to human error. Accidental or deliberate partial restores of a database are a main cause of structural problems.

IMAGE reliability

A database models the dynamic behavior of entities and their relationships by means of *entries*. An entry consists of a key

Database Entries

(which uniquely defines the entity or relationship) and a collection of attributes (which give quality and color to the entity or relationship).

A **dataset** is a homogenous collection of entries. There are different kinds of datasets, each optimized for a specific access technique.

To make an IMAGE database functional, we access its entries in a variety of ways, ranging from **serial scans** of entire datasets to **hashing and chaining**. Hashing and chaining are techniques based on direct access to specific addresses so that we may jump directly into the entry or entries which interest us without having to wade through millions of irrelevant entries. An IMAGE database entry has two parts: the data (your information) and the metadata (the information necessary to do hashing and chaining). IMAGE requires privileged mode (enforced by MPE) for access to the metadata.

Diagnostic Tools

Typically, diagnostics will identify faulty entries by entry numbers and/or by the values of the offending search fields. In either case, we must eventually define some search field value that will guide us in trying to find the master dataset's ChainHead. The fact that we find (or fail to find) the ChainHead does not really tell us much if we have not **exhaustively** tested its master dataset beforehand. For instance, even when we find the ChainHead, the involved master dataset synonym chain may have been broken in the past and there might be other "floating" duplicate master ChainHeads with the same search field value. Not finding the ChainHead is an even surer symptom of a faulty master dataset. The conclusion is that a complete certification of the involved master datasets is a pre-requisite for doing any kind of chain therapy.

ALL Adager users (i.e. Adager Model 1 and Adager Model 2) have the EXAMINE PATH and EXAMINE CHAIN functions available within the ADAGER program to diagnose specific paths or chains for problems. For master datasets, Adager's EXAMINE PATH function performs a check of all of the synonym chains. For detail datasets it performs a check of all of the chains in the path between a given master and the detail. If any problem is found EXAMINE PATH will display a description of the problem, including the entry numbers involved, and will also create permanent log files with this information. EXAMINE PATH is available in both compatibility and native modes.

Adager's EXAMINE CHAIN allows you to interactively diagnose specific chains when you know either the search field value or the entry number of one of the members of a chain. EXAM-

INE CHAIN allows you to traverse any chain to review privileged metadata parameters such as pointer values, chain counts and dataset user labels and bitmaps.

Adager users also receive several additional tools to use in checking the integrity of their IMAGE databases. The LIBRARY group of the REGO account contains these programs: DBCHECK, TDBCHECK and NMCHECK.

DBCHECK is a re-written version of HP's DBCHECK program. It has been enhanced with several options to handle the latest versions of IMAGE and can be used on either IMAGE/3000 or TurboIMAGE databases. TDBCHECK is a copy of DBCHECK and is provided for backward compatibility of existing jobstreams. NMCHECK is a Native Mode version of the DBCHECK program for those users on Spectrum machines. These programs allow a single dataset, a range of datasets (e.g. 1/20), several datasets (e.g. 1,5/10,12), or all datasets to be diagnosed. The default is that all datasets are to be diagnosed.

Most HP3000 systems also come with DIOGENES, a DBCHECK-like database diagnostic program supplied by Hewlett-Packard in the Telesup account and intended to be used by its support personnel.

When users call with possible broken chains I ask them the following question:

“Do you know in which set and possibly which path the problem exists?”

If the answer to this question is “YES”, I tell them to use Adager's EXAMINE PATH function on the suspect path(s). If the user answers “NO” to this question I usually have them run Adager's DBCHECK or NMCHECK depending on the type of machine the database resides on.

This is a very common question. Most diagnostic tools do require exclusive access to the database in order to perform the diagnostic. An exception is Adager's EXAMINE PATH function, because it allows other readers (i.e. DBOPEN mode 5 or 6). Only by performing the diagnostic on a database that is not being updated can you be assured that reported errors are not false alarms. Another notable exception is Adager's EXAMINE CHAIN function, which allows both readers and writers to access the database while you perform the diagnostic.

*Which Tool To Use
When?*

*Do I Need To Have
Everyone Log Off The
Database?*

To avoid having to log everyone off the database, I suggest that you restore a copy of the database into a different group or account (or another machine) and run the diagnostic against the copy. If you do not have enough disc space for this copy (or do not have another machine), the next option is to run the diagnostic at night or on the weekend. All Adager diagnostic tools access the database in a read-only mode; consequently you may abort the diagnostic at any time without harming the database if the diagnostic is taking too long.

I cannot stress enough the importance of establishing a schedule of regular database diagnostics. It is much better for you to find a database problem first, before it finds you. And you know that, if it finds you, Murphy's Law states that it will be during month-end or year-end processing.

The whole subject of database therapy is extremely delicate. Fortunately, should you encounter a problem with your IMAGE databases, Adager can help you in three important (and tightly integrated) areas: Diagnosis, Therapy, and Technical Support.

Adager Database Therapy Functions

Adager offers a wealth of database therapy functions, each one geared to the task of correcting a specific database situation. All of the Adager therapy functions are available to both Model-1 and Model-2 Adager users.

Before using any of the therapy functions it is always advisable to plan the course of action by studying the results from the diagnostics. This will always simplify the steps to take, and avoid backtracking. The first thing to determine, of course, is whether there is indeed a database corruption. Sometimes you may be misled into thinking that your database has a broken chain, when the cause for the message may be an improper locking strategy in a read-only application program.

In any therapy process you have to consider the trade-off between the time the process takes and the complexity of the actions, the degree of your hands-on involvement, and the database availability for other users. In most cases Adager offers you more than one therapy method for the same problem.

Fix Dataset

Adager's "Fix Dataset" function deals with individual dataset parameters. Because it does not deal with the path relationships between datasets, this function is very fast. You should use this function when the problem affecting a database is contained within a dataset or to correct a dataset before proceed-

ing to correct its paths to other datasets. For example, use “Fix Dataset” if you get the “lost free space” message in Query.

For master datasets, “Fix Dataset” re-hashes all the entries and rebuilds all the synonym chains in the process. This function deletes duplicate entries. For automatic masters it also removes entries whose detail ChainHeads are all zeroes. At the end of this process Adager updates the dataset’s entry count.

For detail datasets, “Fix Dataset” rebuilds the delete chain, if one exists, and also updates the dataset entry count.

The main thing to remember about the “Fix Dataset” function is that it is not a way to fix broken chains between a master and a detail dataset or within a detail dataset. To repair broken detail chains you should use “Fix Path” or “Fix Chain.”

Adager’s “Fix Path” function rebuilds all of the chains of a given detail path according to the values of the detail’s search field (and according to the values of the detail’s sort field, if the path is sorted). You should use this function when there is a considerable number of broken chains within a given path. Before using this function, however, you should make sure that the associated master dataset does not have any broken synonym chains by using Adager’s “Examine Path” function, since a broken chain symptom may be caused by either bad pointers in the detail dataset, or by bad pointers in the master dataset Chain-Head. Please remember that you should use “Fix Path” to fix detail paths, and “Fix Dataset” to fix broken master synonym chains.

Fix Path

“Fix Path” rebuilds a path by linking the detail entries in the order it finds them during a serial read. This may result in a different chronology of the chains if the path is not sorted. However, if you have a field in the detail dataset which can be used to establish the desired chronology (for instance, order-date), you may use Adager’s “Add Sort” in combination with “Fix Path” to preserve the original chronology.

“Fix Path” automatically adds any missing ChainHeads to the corresponding master dataset.

“Fix Chain” allows you to deal with individual chains, either by search-field value or by entry number.

Fix Chain

On larger datasets, or in cases where the broken chain is very complex, Adager’s “Fix Chain” provides a way to fix a specific broken chain (rather than rebuilding all chains with “Fix Path”).

If you request a master (manual or automatic) dataset, “Fix Chain” will deal with specific synonym chains. If you request a detail dataset, “Fix Chain” will deal with specific detail-path chains, which include the detail ChainLinks as well as the related master ChainHeads.

You can access specific chains by providing the search-field value or the entry number. To obtain them you may want to first run a diagnostic (such as Adager’s “Examine Path”) to obtain a record number or the key value of the broken chain in question.

Because of the unforeseen complexities that “Fix Chain” can encounter when it tries to repair a problem, you can only run it interactively with exclusive access to the database. Certain problems (such as ChainHead count mismatches) can be fixed automatically by Adager but more complex problems (such as inter-leaved chains) must be handled on an individual basis and manually. The “Fix Chain” function is very powerful and allows you to walk up and down chains and to modify in-use bits, pointers, counts and key values as well as the dataset’s UserLabel.

Adager’s “Fix Chain” function has a very clear interface. Nevertheless, because IMAGE database problems occur very seldom, most likely you won’t remember all the internal bits-and-bytes structural information when a problem occurs. I recommend that you try out this function on a test database and if you are not sure how to proceed, give our technical support team a call.

Fixing one broken chain with the “Fix Chain” function may take a few minutes where rebuilding all of the chains of a 3.8 million entry detail with “Fix Path” may take hours. Whichever function you decide to use, it is important to do a full diagnostic first to determine the scope of the problem.

Replace Dataset

If you have lost the file for a dataset (due to a disc crash, for example), Adager’s “Replace Dataset” function creates a brand-new empty file with the appropriate structure.

If the dataset is a master with paths, you should subsequently do a “Fix Path” on each detail path. This adds the master entries (with the appropriate search field values) required to “parent” the “orphan” detail chains, if any. If the replaced dataset is a detail with paths, you should subsequently do a “Fix Path” on each detail path to initialize the master ChainHeads.

This function is best used with automatic master datasets. If an automatic master dataset becomes lost or is severely dam-

aged, the “Replace Dataset” function can easily be used to build an empty automatic master followed by a “Fix Path” to re-populate the entries into the automatic master without any loss of data.

Creative Adager users have found ways of taking advantage of the by-products of some Adager functions to repair their IMAGE databases. On one such instance, Adager users who experienced the IMAGE prefetch problem under MPE/iX used Adager to quickly rebuild the root-file by adding a dummy item and thus obtaining a new, clean root-file that would allow them to continue with their daily operations within a few minutes after experiencing the problem. Others have taken advantage of Adager’s DetPack function to “rebuild” the free-entry list (since DetPack eliminates the delete chain). Since neither “Add Item” nor “Repack Dataset” are intended for database repair work, we suggest that before using any of these functions to repair an IMAGE database, you call the Adager technical support team to discuss your specific case. On the other hand, if you have some examples on how to use the by-products of some other Adager function to do database therapy, we would like to hear your experiences and suggestions.

Creative Adager users

Fixing database problems can be a nerve-racking and high-pressure situation for you. Fortunately, you are not likely to experience this very often, thanks to the reliability and robustness of the HP3000 computer, the MPE operating system, and the IMAGE database management system. Nonetheless, Adager is always ready to help should the need arise. Feel free to call our technical support staff at anytime to get a “second opinion” or to get our help in solving your particular problem. Our technical support team has seen many of the problems that can exist with damaged IMAGE databases and we can give you several options to correct your specific problem.

*We are here to help
you*