Adager's Consistency Check

F. Alfredo Rego
Adager Corporation
Sun Valley, Idaho 83353-3000 • *USA*

www.adager.com

Let's discuss an automatic function of Adager that does a lot of work under the covers. Every time Adager opens an IMAGE database it does a very thorough analysis that looks for several potential internal database inconsistencies. Some of these inconsistencies are so subtle that IMAGE itself may not catch them until it is too late and your applications may, inconveniently, abort.

You may ask, "Why is Adager so paranoid?" Because IMAGE databases consist of several files and tables that absolutely, positively must be consistent among themselves. There are various ways (malicious or innocent) to disrupt this tight coupling. For instance, we have seen some cases where careless operators—whose operating-system capabilities exceeded their knowledge of delicate and privileged things—restored only certain datasets from a backup, thereby introducing catastrophic inconsistencies within an IMAGE database. We have also encountered RootFiles that have become corrupted because of a hardware problem or a system failure.

An IMAGE database always consists of a RootFile and one or more datasets.

• The RootFile is a binary representation of the database's schema (a schema is the pseudo-human description of datasets and their field layouts — the equivalents of relational tables and their columns — as well as their performance-enhancement criteria). The RootFile is a compact digital repository of the information that IMAGE needs (IMAGE's metadata) in order to keep track of your information (your data).

• Datasets are binary representations of your information (your data). In addition, datasets contain structural information (metadata) that enables IMAGE to keep track of (and to quickly find) any tiny speck of your information (your data) within a vast

space of billions of *other* pieces of information. Think of looking for (and finding) the proverbial needle in a haystack. Some datasets may, in turn, define additional files (in MPE/iX file space and/or in Posix file space) to improve their performance in terms of time (B-Trees, for instance) and/or space (Jumbo datasets, for instance).

There is a lot of data (and metadata) floating around. And it has to be all perfect. Adager is very strict about following the rules.

RootFile   First Adager checks the RootFile to make sure that it is self-consistent. Upon finding inconsistencies, Adager will report certain problems as errors and other problems as warnings. A corrupted path table would be an error while an automatic master with no paths would be a warning. Adager displays a wealth of options to help you correct these internal discrepancies, before the computer system propagates the damage. This is equivalent to checking a swimming pool to make sure it has water *before* you do a triple-loop dive into it. If Adager finds out that it is "empty" you may choose to "fill it up!"

Dataset files   After making sure that the RootFile appears reasonable, Adager checks all the files corresponding to all of the database's datasets. The file structures as well as the data within the files must be correct, at least globally speaking (Adager will not complain if the price of a widget is $20 instead of $45, due to a data-entry error). The kinds of things that Adager reviews are: file record size, blocking factor, block size, end-of-file, file-limit, file-creator, user labels, global dataset counts, highwater mark, "presence" bitmaps, dataset names, data item names, fields, paths, and so on.

Not necessarily, because Adager's global consistency check — which takes just a few seconds — limits itself to the big picture. For instance, it does not detect broken chains within a dataset. For this purpose, you would use Adager's "Examine Path" function.

In general, if your database passes Adager's global consistency check without a hiccup, you have a quick, warm, fuzzy feeling that things are *not* obviously wrong with your database. However, if Adager finds a global inconsistency, then you probably also have other inconsistencies — such as broken chains — and we recommend doing a full database diagnostic as soon as possible. Such a detailed Adager diagnostic, by its very nature and functionality, will take more time; but, by the same token, its recommendations will be more precise, down to the data entry level: chain pointers, search field values, and so on.

If my database passes Adager's global consistency check with flying colors, is my database OK on all counts?

You should immediately use Adager's Therapy functions.

What if my database does NOT pass Adager's global consistency check?