How an OpenInsight Database is Loaded into Memory

Loosely speaking, a database is a collection of data available to an application. Some databases are persistent, meaning they are stored on disk (etc.). Some databases are stored in memory, and are cleverly named "in-memory databases". Some databases (their "files") are manipulated directly, like the original dBase and Revelation Software's Linear Hash database (in "local" mode), and some databases are manipulated by a server process, such as Oracle, Microsoft SQL Server, and Revelation Software's LH/NLM and LH/NT Service network products.

OpenInsight's database supports all of these by allowing the databases and storage implementations to plug in as filing systems. Most often (99% of the time) OpenInsight databases use Linear Hash as the persistent storage implementation and an in-memory filing system implementation for the database configuration itself.

The database configuration is composed of three tables:

- SYSVOLUMES The set of volumes that contain the available tables in the database, stored using an in-memory filing system implementation (R TP50);
- SYSTABLES The set of tables that are available in the database, stored using an in-memory filing system implementation (RTP50);
- SYSCOLUMNS The set of columns making up the tables in the database, obtained dynamically (i.e. assembled) from the dictionaries of the tables available in the database.

To understand SYSTABLES and SYSVOLUMES, it is necessary to define what volumes and tables are:

- A volume is a named collection of tables. A volume (via the filing system open-media call) provides access to a volume table, called a media
 map, which defines the entire set of tables accessible in that volume. (REVMEDIA is the name of this table in the Linear Hash filing system.) At
 runtime, a volume is a collection of those tables that are part of the current database.
- A table is a named collection of records. A table is part of a volume. A table is only accessible (i.e. can be opened) if it is attached.

Internally, the OpenInsight database maintains the table and volume information in two dimensioned arrays, **@TABLES**, and **@VOLUMES**. They are also can be opened the tables **SYSTABLES** and **SYSVOLUMES** in the database, just like any native OpenInsight table.

When you login to OpenInsight, the persistent information about the tables and volumes to be attached (stored in the application's **.DBT** file) is read into memory and stored in **@TABLES** and **@VOLUMES**. The application references these memory variables - *not* the .DBT file - when the application is running. If additional tables are attached or detached, or indexes are built, only the memory variables are affected. This means that the database definition stored in the .DBT file can become unsynchronized with the in-memory variables. An unsynchronized database could produce maintenance problems.