Database Structure

When you create a new database, be it single or multiple file, a number of things happen:

- The database file(s) are created;
- The header page is formatted and written;
- The various system tables - RDB$ and MON$ - and associated indices are created - and appropriate pages formatted and written to disc;
- Every page in the database is formatted with a defined page type.

The various page types are described elsewhere in this document.

A database is created and the DBA can specify the page size, or leave it to default. This action creates a database file, or files, with enough space allocated to create all the system tables and indices. New pages will be added to the end of the database file(s) as and when the user creates new tables and/or indices. For example, a brand new database, created on a 32-bit Linux system, with a 4Kb page size allocates a total of 0xa1 pages (161 pages) for the system tables, indices and the various database overhead pages.

Single file databases

A single file Firebird database consists of a number of pages, each the same size, and all held within one file on the underlying file system, be it NTFS, FAT32, EXT3 etc.

The first page in the database is always a header page (page type 0x01 - see Database Header Page) which holds details about the database itself, the page size and so on.

The second page in the database is a Page Inventory Page or PIP (page type 0x02) which details which pages in the database are in use and which are free.

Up until Firebird 3.0, the next page is a Write Ahead Log page (page type 0x0a) but this page is wasted space, if present, and will most likely be dropped from Firebird 3.0 onwards.

The remaining pages consist of Index Root Pages (page type 0x06), Transaction Inventory Pages or TIP (page type 0x03), Pointer Pages (page type 0x04), Index B-tree Pages (page type 0x07), Data Pages (page type 0x05) and so on. There is a discussion of each page type in this documentation.

Multiple file databases

A multiple file Firebird database is almost identical to the single file database except it has more than one file on the underlying file system. Each file has the same page size as the initial file, and each file has a header page (page type 0x01) at the start of the file.

back to top of page

Shadow files

Shadow files are additional files that can be used by single and multiple file databases to assist in recovery after a failure of some kind. They are not helpful in the case of a DROP DATABASE as the shadow file(s), being part of the database, are also dropped!

Shadow files are updated as the database main file(s) are updated and in this manner, the shadows are an identical copy of the database. In the event of a problem, the SYSDBA can manually activate a shadow, or have the Firebird engine activate one automatically.

Unfortunately, if a database write stores corrupt data in the database, the shadow file(s) will be identically corrupted.

Because shadow files are effectively identical copies of the database files, they will not be discussed further.

From:
https://ibexpert.com/docu/ - **IBExpert**

Permanent link:
https://ibexpert.com/docu/doku.php?id=01-documentation:01-08-firebird-documentation:firebird-internals:database-structure

Last update: **2023/07/11 09:20**