

# Time to Backup: A Firebird/Interbase Backup scheduling service

Michaël Van Canneyt

17th February 2004

## 1 Features

"Time to Backup" is a suite of programs that allow to set up a scheduled backup of Interbase or Firebird databases on multiple hosts. It has the following features:

- On Windows NT (2000 or XP) the scheduler works as a service.
- On Linux, the service works as a simple daemon, and can be run by any user.
- A GUI configuration tool allows to set up various backup sets and to control the scheduler. This tool works on Windows and Linux.
- A separate service or daemon can be installed to allow remote registration: The GUI tool can connect to the remote administration daemon to manage the configuration on a remote host.

The tools come with full sources, written in Object Pascal.

## 2 System requirements

On the windows platform, Windows NT 4, Windows 2000 and Windows XP are supported. Windows 2003 Server is also expected to work. Windows 95, 98 or Millenium edition are *not* supported, because they do not support services. Currently, no plans exist to support these platforms.

On Linux, any recent platform which has at least glibc 2.2, Qt and X installed should work. This means SuSE 7.0 and higher, Redhat 7 and higher, and mandrake 7.2 or higher. If only the daemon is running (no X) only glibc 2.2.3 is needed.

Obviously, on all platforms the Interbase or Firebird client must be installed. Any version of Firebird should work, Interbase 6.0 or higher is required.

To compile the sources, Delphi 5,6 or 7 is required on Windows, and Kylix 2 or 3 on Linux. Furthermore Indy 9 is required if the remote configuration module is needed.

The plugin for IBAdmin needs IBAdmin version 4 at least and needs Delphi 6 or higher.

## 3 Installation on Windows

The setup package required will perform all that is needed to install the backup scheduling service, manager, and the remote configuration service. These packages should be installed

and run as a user which has administrative rights: The user should be able to modify the registry below the HKEY\_LOCAL\_MACHINE root key.

If the backup scheduler wasn't started automatically by the setup procedure, the configuration manager can be used to start the service: The status of the scheduling service is displayed at the bottom of the screen. In the lower right corner, a 'Start' button can be found which will start the service if it is currently not running.

If the remote configuration tool is not started, it can be started with the service manager of windows. This is located in the control panel:

- On Windows NT, the icon 'Services' must be used.
- On Windows 2000 and XP, the 'Services' Icon is located under the 'Administrative tools' icon.

The 'RemoteFibss' service is the remote configuration manager. It should display a status of 'Running'.

## 4 Installation on Linux

The installation on linux is a matter of running the `setup.sh` script which is contained in the installation tarball. It will ask the location where the programs must be installed, and will extract all needed files there. It will then proceed to create a symbolic link in the path.

All directories must be writable by the user installing the package; No checks are made whether this is so, so if some directory cannot be created or is not writeable, some ugly error messages will be displayed.

If the installation is performed as the root user, the script can attempt to install a startup script. The startup script will be installed in the directory

```
/etc/init.d
```

If it fails to locate this directory, the script will not be installed.

If installed successfully, a symbolic link will be installed:

```
\usr/sbin/rcfibss
```

By default, the fibss service is started in runlevels 3 and 5.

Finally, the script will offer to start the service. If not running as root, it will not check whether the service is started successfully. If the runlevel script was installed, that will be used to start the service, in which case a status message should be printed.

## 5 Configuration data on Windows

The configuration data is stored in the Registry on Windows, under the HKEY\_LOCAL\_MACHINE root key, in key

```
\Software\FIBSS\fibss
```

The manager program edits the data under this key. Using the '-R' command-line switch, both the scheduler and the manager tool can be instructed to use HKEY\_CURRENT\_USER as the root key, but this is not recommended.

## 6 Configuration data on Linux

If the scheduler runs as root, it will look for the configuration file (a file in Windows .ini file style) in the /etc directory:

```
/etc/fibss.cfg
```

If the scheduler runs as another user, then the file will be searched in the user's home directory, for example:

```
/home/michael/.fibss.cfg
```

Note the initial dot character.

When running, the scheduler will create a '.pid' file, which contains the Process ID of the scheduler. This is used by the manager program to notify the scheduler of changes in the configuration file. If the scheduler runs as root, then the file is called

```
/var/run/fibss.pid
```

For any other user this file is located in the home directory, for example:

```
/home/michael/.fibss.pid
```

Both the manager and sheduler have command-line options to override the defaults and use specific configuration files:

```
fibsmgr -c /usr/local/fibss/fibss.cfg -p /usr/local/fibss/fibss.pid
```

would use the configuration file /usr/local/fibss/fibss.cfg and PID file /usr/local/fibss/fibss.pid.

## 7 Scheduling backups

When it starts, the backup scheduler reads the backup time from its configuration data. It then goes to sleep till the specified time arrives. If it cannot find a backup time, it will give an error message and exit. It is therefore important that a default configuration is set up before starting the scheduler.

When the manager changes the backup time, the scheduler is notified of this. If the configuration data is modified manually, this notification must be done manually too (by stopping and restarting the scheduler, for instance).

On Windows the scheduler can be notified by sending it a custom service control code 129. On Linux, sending the service a HUP signal will cause it to re-read its configuration data.

The unit ctrlBackup in the sources contains the necessary routines for this.

## 8 General configuration

The scheduler has some general configuration settings:

**Backup time** the time of the day when backups are made. All backups are started, one after the other, at this time.

**Backup directory** if no directory is specified for a backup file, this directory is used.

**Log directory** if no directory is specified for a log file, this directory is used.

**backup extension** if no backup extension is given for the backup file, then this extension is used.

## 9 Backup configurations

After this, a series of backups can be defined: Each backup is identified by its name. A sensible name should be used, do not include special characters in this name. A special name is '(default)' which is used to specify default settings for the backup configurations.

For each backup, the following can be specified:

**Name** Name of the backup configuration (required).

**Server** Name of the Firebird/Interbase server to connect to. Leaving this empty assumes that Interbase is running on the local machine.

**Database path** the database path for the database to be backed-up (required). This must be a path on the server.

**Backup filename** this is the name of the backup file (required). (multiple backup files are not yet supported in the manager)

**Username** Username to use when connecting to the Firebird/Interbase server (required).

**Password** Password to use when connecting to the Firebird/Interbase server (optional).

**Log file** Name of a log file to which a log of the backup operation will be written (optional).

**Options** various options that can be set for an Firebird/Interbase backup. Refer to the documentation of these products for the meaning of these options.

As many backup configurations as required can be created. They will be handled one by one.

Should the backups take more than 24 hours, then the scheduler will *not* wait for a previous backup to finish, but will simply wait another 24 hours.

## 10 Using defaults and Templates

If a lot of similar backups must be made, it makes sense to use defaults: For all values except the database path, a default value can be specified using the 'Defaults' button.

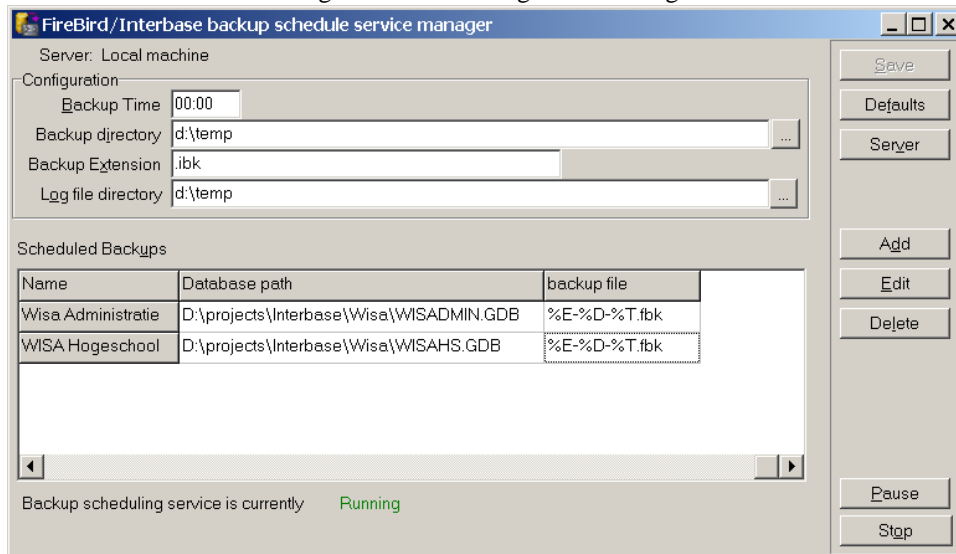
To be able to distinguish backup files from various backup operations, the backup filename and log filename can contain placeholders. A placeholder exists of a special token (a % character followed by a letter) which will be replaced with a specific value:

**%E** will be replaced by the backup entry name.

**%D** will be replaced by the date of the backup.

**%T** will be replaced by the time the backup is started.

Figure 1: Main configuration dialog



%N will be replaced by a unique number which is incremented for each subsequent backup file with the same name. This will only work if the backup file is on the same system as the backup scheduler, as the scheduler will check whether the filename exists.

for example, given the template %E-%N. ibk and an entry name of 'employees', the scheduler will check whether the files

```
employees-1. ibk
employees-2. ibk
...
employees-N. ibk
```

The first number for which no backup file exists will be used.

(other templates may be added later on)

## 11 Configuration Manager

When the configuration manager is started, the main configuration dialog is shown (as in figure 1 on page 5). The following items can be specified:

**Backup time** time of the day when the backup will be run.

**Backup directory** directory where backups will be stored by default.

**Backup extension** extension to use for backup files when no extension is specified. The initial dot (.) character must be specified.

**Log file directory** directory where logs will be written by default.

The 'Save' button will save the configuration. If the service is running, it will be notified of the changed configuration, so that it can re-read the configuration.

The 'Server' button can be used to specify a remote machine on which the scheduling service is running. More about this feature is described in section 12 on page 7.

Figure 2: Backup entry dialog

Hitting the 'Add' or 'Edit' button will open the backup entry dialog. This dialog can be used to create or change backup entries. It looks as in figure 2 on page 6. The following settings can be managed in this dialog:

**Name** Name of the backup entry. This should only be specified when adding a new entry.

**Server name** Name of the Interbase or Firebird server which should make the backup.

**Database path** Filename (on the Interbase/Firebird server) of the database which must be backed up.

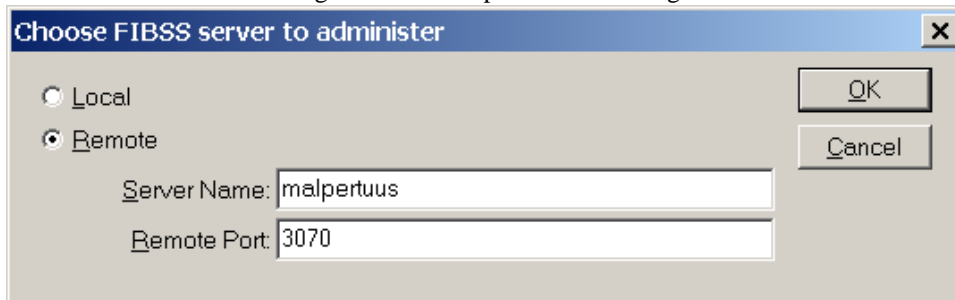
**Backup filename** filename to use when backing up. Templates may be used for the filename. If no path is given, the standard backup path is used.

**Username** username to use when connecting to the database.

**Password** password to use when connecting to the database. The password is not shown when you type it.

**Log file** if set, a log of the backup will be kept in the filename as specified here. If no path is given, the standard backup path is used. Here also, templates can be used in the filename.

Figure 3: Server specification dialog



**Backup options** the various options to use when creating a backup. See the Interbase/Firebird documentation for more information about the various options.

**Schedule options** If the 'Skip entry temporarily' option is checked, this entry will be disregarded when creating backups. If the 'Only backup on these days' option is set, a backup will only be created on the given days.

If any of these entries is not supplied, The corresponding entry from the default settings will be used. The defaults can be set via the 'Defaults' button in the main configuration screen.

Double clicking the grid with backup entries in the main configuration screen will also open the 'Backup Entry' dialog.

The 'Pause/Continue' and 'Stop/start' buttons in the main configuration dialog can be used to pause, respectively stop the backup scheduling service.

## 12 Remote configuration

Remote configuration works through a TCP/IP connection between the manager and a special configuration service. This means that, on the server where the scheduling service is located, a special service can be installed which will manage the configuration data of the backup scheduler on the same machine. It simply listens to requests from a remote client and applies requested configuration changes or provides configuration data to the client.

The manager program can issue calls to the remote configuration manager running on a remote machine. To do this, an IP address and port number must be entered for the remote configuration manager. The address can be entered using the 'Server' button. A dialog as shown in figure 3 on page 7 opens where the IP address and port can be entered.

By default, a connection to port 3070 is used, but this can be changed on the server.