Part D

Enterprise Manager

D.1 Introduction

The Enterprise Manager (EM) is the programmer's main tool for table and database maintenance operations. The EM allows a developer to

- Perform routine maintenance of NexusDB table structure,
- View and maintain data in tables,
- Maintain and execute SQL scripts,
- Perform live backup and restore operations, and,
- Export and import data in CSV format.

The EM is an MDI application allowing multiple table, SQL and restructure views to be opened at the same time. The use of cut and paste operations between all views is fully supported.

This tool is intended for development use only. Due to its ability to allow direct access to the data in NexusDB tables, it is not recommended for distribution to your clients. As the full source for the EM is provided for registered users, it is straightforward for developers to implement their own application specific table utilities.

D.1.1 Booting up as single user

When the EM is booted up and there are no NexusDB Servers broadcasting (or registered) the screen shown in Figure 1 is displayed. The server displayed is the embedded server that is compiled into the EM. The red cross indicates that the server is currently inactive. To activate the server, just double-click on the red cross or server name (in this case, Internal Server) label itself, or press Enter when the server is the active selection.

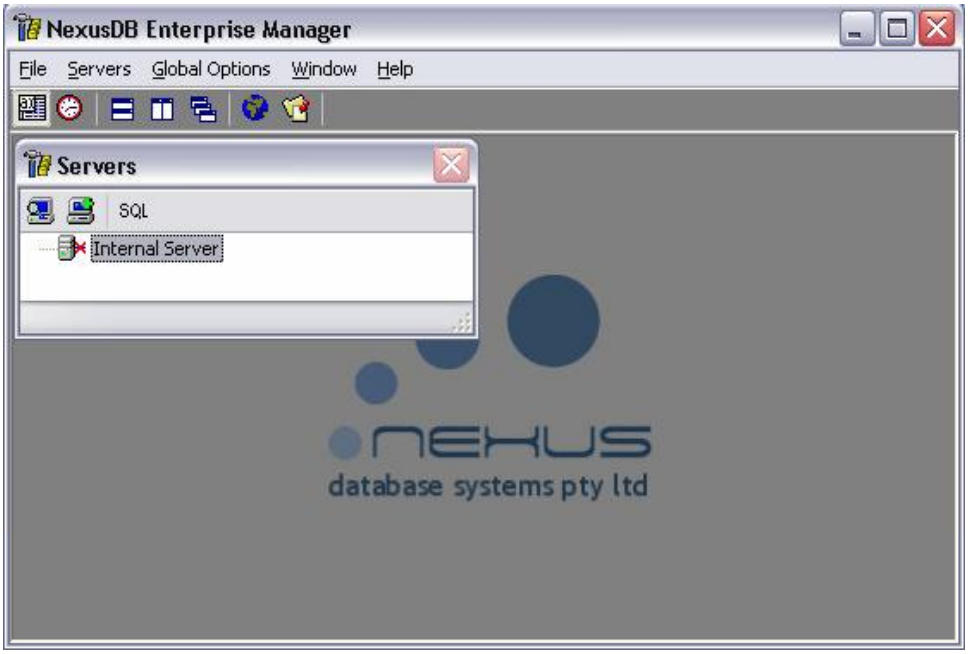

Figure 15: Enterprise Manager–Initial display with no external servers detected

Before we investigate connecting to servers, and other actions, we shall first discuss the Global Options settings.

D.1.2 Global Options Menu

There are some global options that can be set in the EM. These options are available via "Global Options" from the main menu when you first boot-up as shown in Figure 2 below.

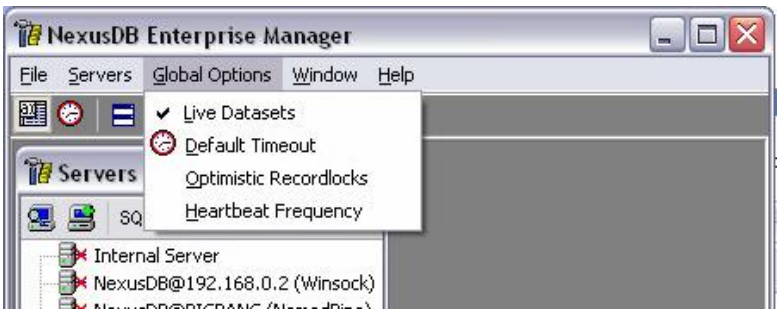

Figure 16: Enterprise Manager–Global Options

The options available are:

1. Live Datasets – determines if data can be modified or not
2. Default Timeout – sets timeout for dataset operations
3. Optimistic Recordlocks – implements optimistic record locking with external servers
4. Heartbeat Frequency – how often EM tells the server(s) its still alive

Live Datasets toggles Read-only mode for all SQL and table browser windows. NOTE: To protect against inadvertently modifying data, the default (when the EM is first used) is Read-only = True. If you want to modify data whilst in the EM, then you will need to change this setting to Read-only = False.

The current state of this option is shown in the first button of the toolbar. The default position for this button is shown in the Figure 3.

Figure 17: Toolbar icon with Live Datasets disabled

If you enable Live Datasets (allows data to be modified whilst in the EM) then the state of this button is as shown in Figure 4. This is the usual state selected by developers for this option.

Figure 18: Toolbar icon with Live Datasets enabled

Default Timeout – The default timeout value is set to 10 seconds as shown in Figure 5. This value is used as the default when table browsers are opened and for SQL queries.

Figure 19: Dialog for checking and modifying the timeout setting

Note that the value is presented in milliseconds (there are 1,000 milliseconds per second). If you set this value to 0 milliseconds, then the timeout is ignored (has, to all intent purposes, an infinite value). This can be handy when you are testing a particularly complex query which may take many minutes to evaluate however, generally speaking, you would want the timeout to be a finite value to guard against deadlock situations.

If you set this value to –1, then the value used will be the value of the next highest level component in the connection chain to the currently selected server. For the internal server, that is also 10 seconds. For external servers the value may be different.

Optimistic Locks – Determines whether table browser windows use optimistic or pessimistic locking when performing edits.

Heartbeat Frequency (ms) – How often EM sends heartbeat messages to the connected server(s) to inform them that its still alive. If the connection is lost for the duration of 2 heartbeats, the server closes the connection and frees up server-side resources.

D.2 The Servers window

When started, the EM will come up with the Servers window already open as shown in Figure 6 below. The Server window is special in that it always remains visible on screen at all times. It cannot be closed, minimized, or maximized, but you can position and resize it any way you like. This window is where most tasks can be initiated or performed.

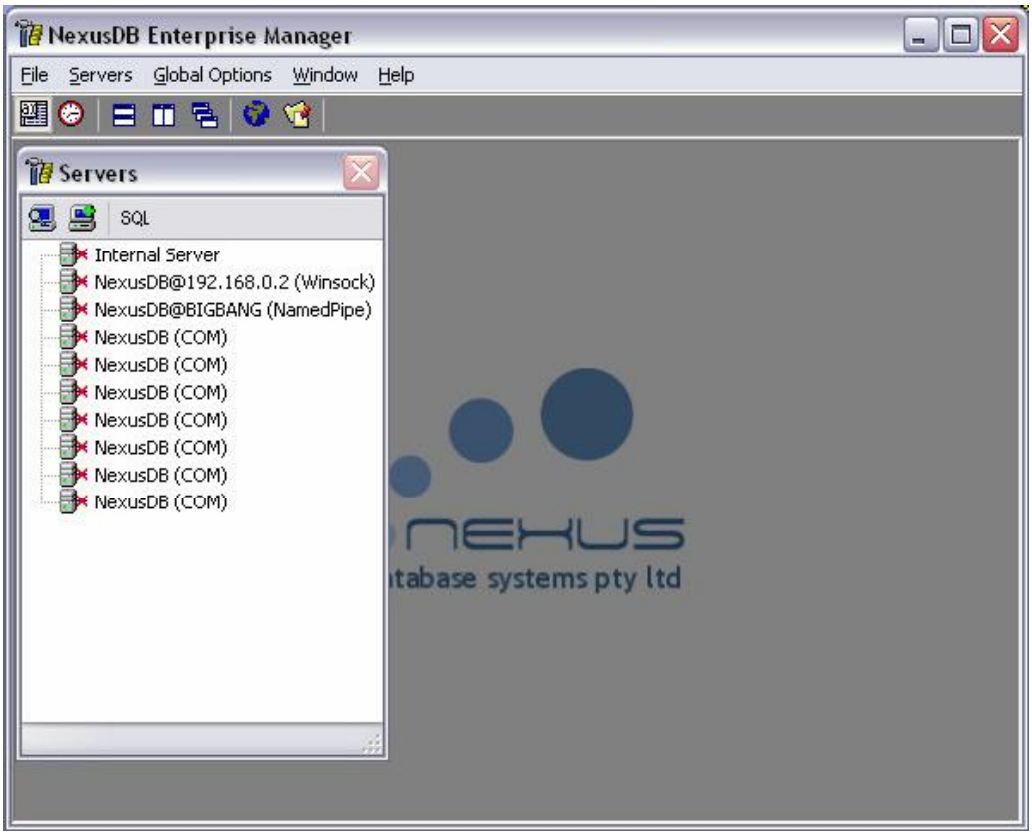

Figure 20: Enterprise Manager – Servers Window on startup

The EM scans for servers using broadcast and lists all servers found, on all available protocols. With the exception of the Internal Server, which uses an internal server engine, servers are named in this format: "servername@address (protocol)". Figure 1 above shows a typical screen on startup of the EM in which a server has been found at 192.168.120.8 using Winsock transport. The small red **x** next to each server located indicates that a connection to that server has not yet been established.

In Figure 7 below, two external server connections are listed (one using Winsock transport, one using NamedPipe). The internal server connection has been started (to start, simply double click on either the server icon or name) and displays several aliases. One of those aliases ("d") has been expanded to show all the tables in that alias.

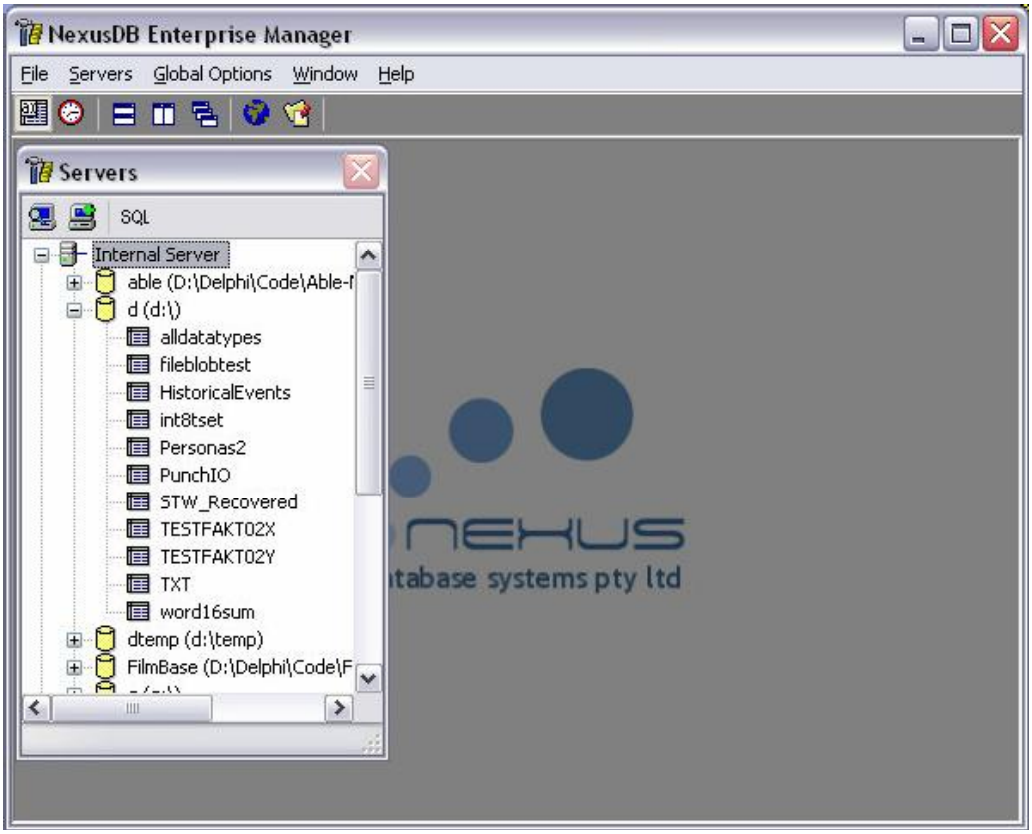

Figure 21: Internal Server has been started and aliases displayed

D.2.1 Servers Menu

The Servers window toolbar and the "Servers" item on the main menu have the items shown in Figure 8. This menu set allows you to maintain connections at the database server level.

Figure 22: Server options available

Refresh Server List (F5) – If you start up an external NexusDB Server, this will refresh the list to make it appear. Similarly with closing down servers and refreshing aliases in external servers.

Register Server – Opens a dialog where you can add servers manually. This must be done if the server can't be reached by the normal broadcast method (for instance, if the server is on a different subnet, or only reachable via the Internet). See section D.6 in this chapter for details.

Run SQL Queries (F4) – To open a SQL window on the active database or table.

Let's now run through the pop-up (right-click) menus in the server tree, which is where most of the important commands in the EM are found. Right-click a node in the server tree. Depending on the type of node (server, database or table), the menu displayed has these items (note: not all items have associated icons):

D.2.2 Server Popup Menu

Once the focus is set on a particular server in the main window area, a popup menu to manipulate that particular server is activated. The items available in this popup menu are shown in Figure 9. This menu set allows you to manipulate database aliases and work on the currently active server.

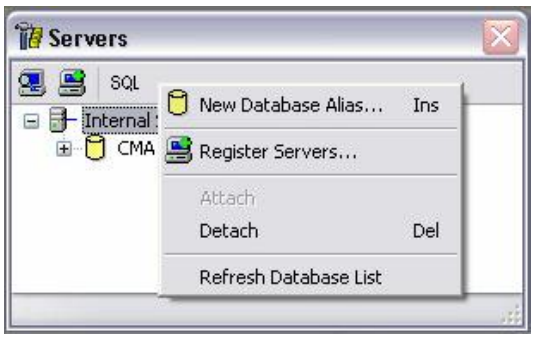

Figure 23: Popup Menu options available for each displayed server

New Database Alias – To create a new alias on the active server in the list, enter an alias name and a directory path in the dialog as shown in Figure 10. Keep in mind that the directory path is always relative to the machine the NexusDB Server is running on (which could be another machine). If the path is not on the local machine, then the EM can't create it for you.

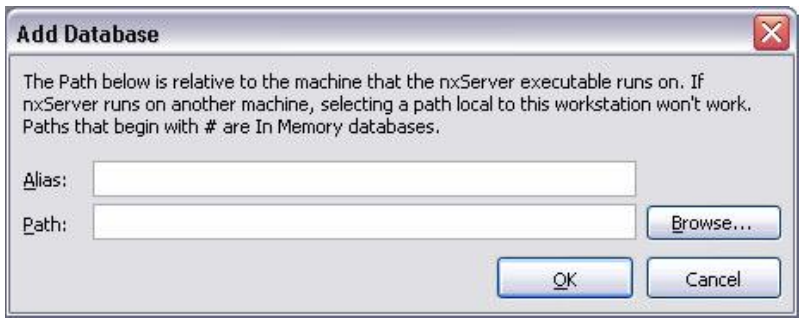

Figure 24: Dialog for creating a new alias associated with the active server

Register Server – this option brings up the dialog for registering servers with the EM. This is discussed fully in section A.6 The Registered Servers Dialog.

Attach – Use this command to open the connection to a server (if you right-click to open the pop-up menu, the EM will automatically try to establish a connection). If the connection is made active, the server icon in the tree will display a green "link active" line on its right-hand side. If no connection exists, the server icon shows a red "x" at the right-hand side.

Detach – Use this command to close a connection to a specific server.

Refresh Database List – If the alias list is changed outside of the EM (that is, directly on a NexusDB Server), this option will refresh the list here.

D.2.3 Database Popup Menu

On setting the focus to a particular alias in the currently selected Server, a further popup menu to manipulate that alias and tables therein is available for use. The items in this popup menu are shown in Figure 11. This menu set allows you to manipulate the selected database alias and to work on the tables within that alias.

Figure 25: Popup Menu options available for each displayed alias

New Table (Ins) – Open the Create table dialog. See section D.5 for details about creating a new table.

 **New Database Alias** – Create a new alias on the server this database belongs to. This is the same dialog linked to the similar menu option for the popup menu associated with servers.

 **Delete Database Alias** – Remove an alias. The command only removes the alias from the server's list of aliases, no tables or directories are physically deleted.

Rename Database Alias – Give the alias another name.

Start Transaction – Starts a transaction on the database. The database node in the tree will get a circled "T" marker to indicate that a transaction is in progress:  All normal table- and SQL operations will now operate in the context of this transaction. **Caution**: keep in mind that using transactions on a database where other people are working, will lock others out from doing their work. Use with extreme caution in production environments!

Commit Transaction – Commits all edits done since the Start transaction command.

Rollback Transaction – Rolls back all edits done since the Start transaction command

Refresh Table List – Update the list of tables owned by the database, in case something external to EM changes the list.

 **Pack All Tables** – Compact and reindex all tables in the current database.

Set Password (All Tables) – Allows you to give all tables in a database the same password. Note carefully: this password feature **only** works within the EM. It is only intended to stop end-users who have the EM executable from using it to view the content of your tables. Keep in mind that any person with a NexusDB license can simply disable the password checking in the EM source, and view your tables.

Live Backup – Run a backup from one server/database to another, even while the database is in use. See the TnxBacupController (page !!!) for more info on how the backup operates.

Restore – Copy back tables that was previously backed up. This command also works if you just want to copy tables from one server/database to another.

CSV Import – Import CSV data into a new NexusDB table. Supports most variants of CSV and date formats. See section D.7 for detailed info on the import wizard.

SQL **Run SQL Queries (F4)** – Opens a SQL window for querying on the active database. See more info on the query window in section D.4.

D.2.4 Table Popup Menu

Drilling down even further in our focus, we can now concentrate on individual tables within an alias on a particular Server. Setting the focus to a particular table in the currently active alias (in the selected Server), takes us to a further popup menu to manipulate that table. The items in this popup menu are shown in Figure 12. This menu set allows you to manipulate the selected table in the database alias and to invoke SQL command scripts within that alias that default to that table.

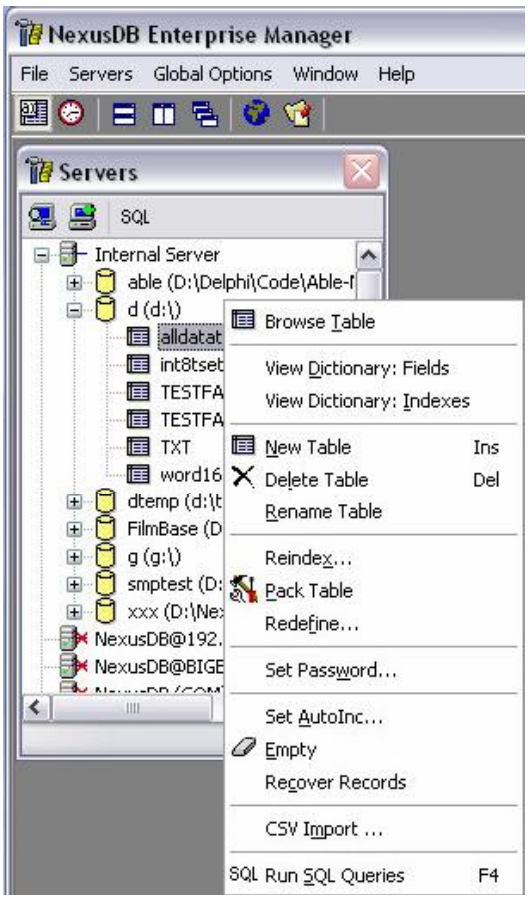

Figure 26: Popup Menu options available for each table displayed in an alias

View Table – Open the Table browser window. See section D.3 for details about this window.

View Dictionary: Fields – Open the Create/Restructure window in read-only mode, with the Fields tab active. See section D.5 for details about this window.

View Dictionary: Indices – Open the Create/Restructure window in read-only mode, with the Indices tab active. See section D.5 for details about this window.

 **New Table (Ins)** – Open the Create table window. See section D.5 for details about this window.

 **Delete Table (Del)** – Deletes table permanently from disk.

Reindex – Choose an index in the dialog that you want to reindex, and press OK. If you want to reindex all Indices, use the Pack command.

 **Pack Table** – Compact and reindex all Indices.

Redefine – Opens the Create/Restructure window with the current contents of the table. See section D.5 for details about this window.

Set Password – Allows you to give the table a password. Note carefully: this password feature **only** works within the EM. It is only intended to stop end-users who have the EM executable from using it to view the content of your tables. Keep in mind that any person with a NexusDB license can simply disable the password checking in the EM source, and view your tables.

Set AutoInc – Allows you to change the last used autoinc value for a table. The next insert operation will generate the value (newvalue+1) in the autoinc field.

 **Empty** – Delete all the records in the table. This command is not reversible.

Recover Records – If you have a table whose internal structure has been damaged, this function will attempt to extract all recoverable data into a new "recovered" table. Records that were not fully recoverable are copied into a "failed" table.

CSV Import – Import CSV data into a new NexusDB table. Supports most variants of CSV and date formats. See section D.7 for detailed info on the import wizard.

 **Run SQL Queries (F4)** – Opens a SQL window for querying on the active database. See more info on the query window in section D.4.

D.3 The Table Browser window

When you want to view the plain data in your tables, or check results of filters, ranges etc, the Table Browser window will do this and more.

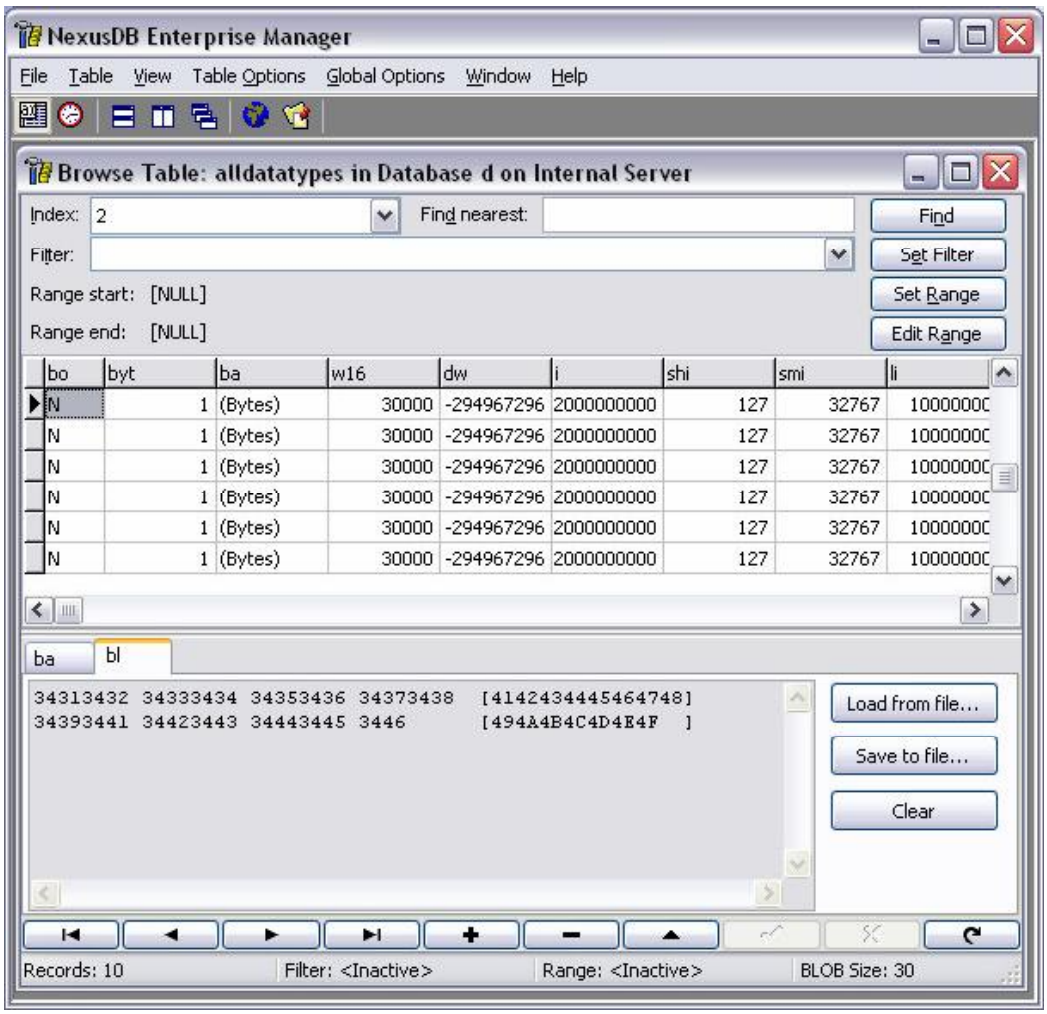

Figure 27: Viewing the data contained in a table

Select the index you want to sort by in the Index ComboBox at the top. To use the Find and Set Range functions, which both require a searchable index., select an index other than the Sequential Access Index.

The Set Range dialog is dynamically configured to have one line of settings for each field in an index. Set start and end values for each field to be ranged, and press OK. The fields and options correspond to the SetRange() and KeyExcusive parameters that you would use in normal table-handling code.

The grid displays all fields in the table, with default VCL formatting. The data is directly editable, unless the main menu Live Datasets option is off (it is off by default).

If the table contains Blob fields, then the bottom area of the window will display the blobs in the active record (see picture above, the Blob field contains an image). This can be turned off in the View menu for browsing speed.

When made the active window, the Table Browser window adds a few menus to the main menu. These are discussed in the next section.

D.3.1 The Table menu

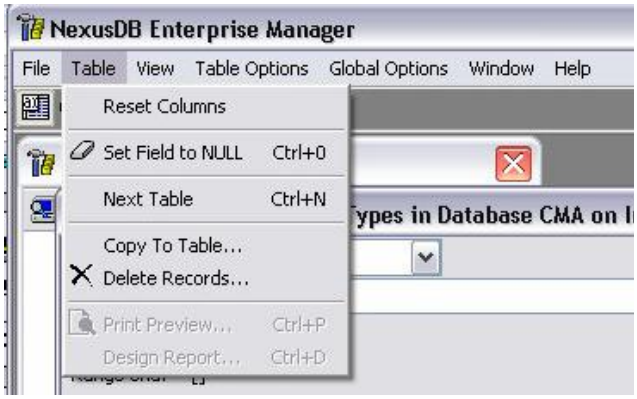

Figure 28: Table menu options

Reset Columns – Ensures that columns have reasonable min/max widths.

Set Field to Null (Ctrl+0) – Clears the value of the active field in the grid.

Next Table (Ctrl+N) – Closes the active table, and opens the next table in the database (in alphabetical order) in the same window.

Copy To Table – Copies the contents of the table to another table, respecting any active filters and ranges.

Delete Records – Deletes the records in the table, respecting any active filters and ranges. Records not selected by filters or ranges remain.

Print Preview (Ctrl-P) – If the Report Engine DLL is installed, opens the table in print preview. See section D.8 for more on the Report Engine DLL. In Figure 14 above this, and the next, option are disabled as the Report Engine DLL could not be found for the current installation.

Design Report (Ctrl-D) – If the Report Engine DLL is installed, opens the report designer. See section D.8 for more on the Report Engine DLL.

D.3.2 The View menu

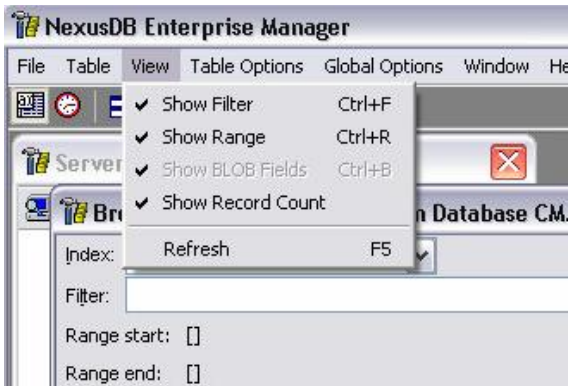

Figure 29: The view menu options

Show Filter (Ctrl-F) – Shows or hides the Filter area of the Table Browser.

Show Range (Ctrl-R) – Shows or hides the Range area of the Table Browser.

Show Blob Fields (Ctrl-B) – Shows or hides the Blob display area of the Table Browser.

Show Record Count – Shows or hides the display of RecordCount in the statusbar.

Refresh (F5) – Performs a table.Refresh.

D.3.3 The Table Options menu

Figure 30: The table options menu options

Debug Log – Turn on logging of errors to the debug log.

 **Set Timeout** – Default timeout is 10 seconds. For operations on large table, it might be necessary to increase this setting.

D.4 The SQL window

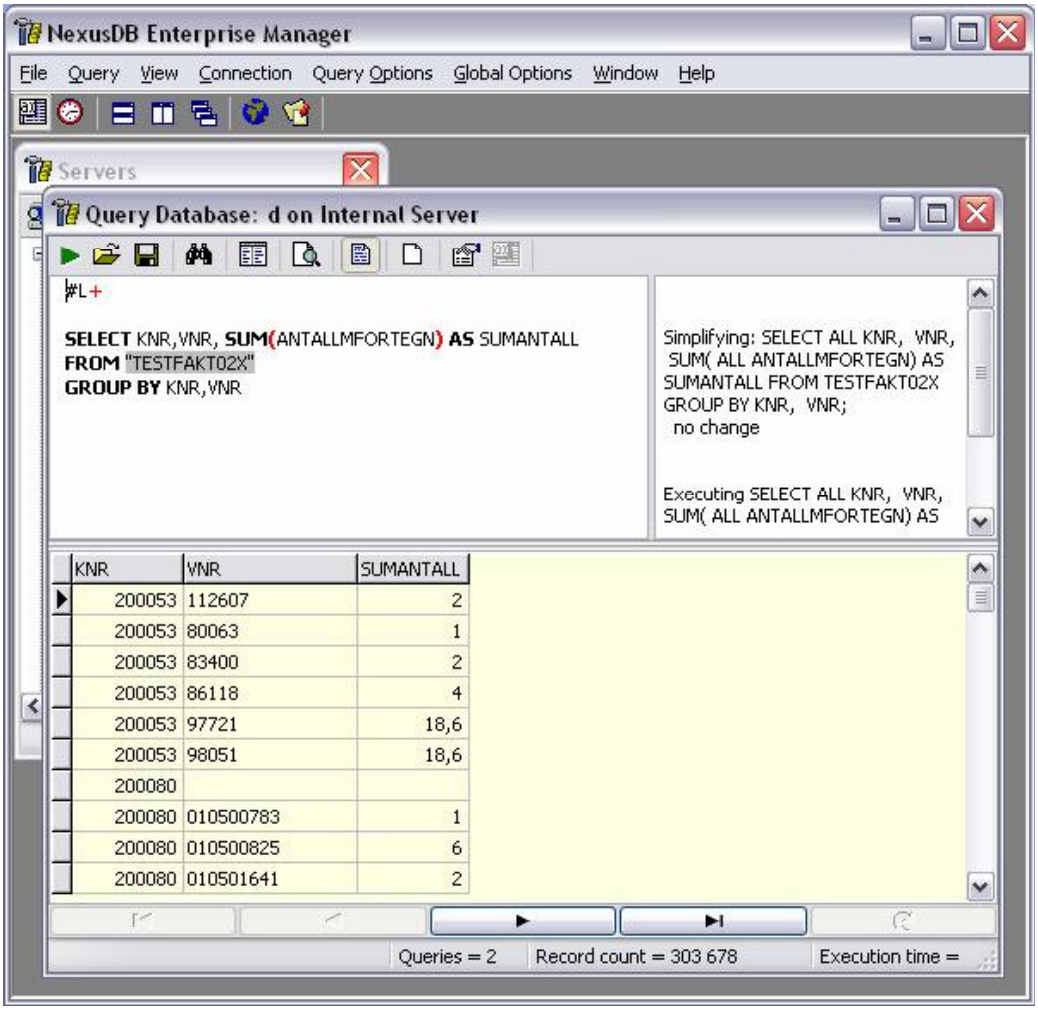

Figure 31: Enterprise Manager – SQL Window

This is where SQL queries can be tested and executed. The SQL script input area accepts scripts of any length. Multi-statement scripts must have each single statement terminated with ';'. If the statement is a single SELECT statement, or the last statement in a multi-statement query is a SELECT, then the query result set displays in the grid after execution. If the last statement is not a SELECT, the grid will be empty and instead the statusbar will display how many records were affected by the last non-SELECT statement. If live (editable) queries is turned on in the local and global options, and the query syntax conforms to what is supported for live queries, the grid will be editable, otherwise it will be read-only. If you want to change values in a non-live query, either use UPDATE statements, or open the table in a table browser window (see section D.3).

The query browser window allows you to have several SELECT statements open on the same database at once. This is a handy way to compare different variants of a query to see which records

they return. Use the Create New Connection function to add statements. The list of statements is displayed in the combobox in the toolbar.

To view NexusDB's SQL execution plan for a query, a log window can be activated. Note that in order to return the necessary log, the query must include the #L+ option. It is autoinserted at the top of the query text when you turn on the log via the menu.

If the query resultset contains Blob fields, then the bottom area of the window will display the blobs in the active record. This can be turned off in the View menu for browsing speed.

The following menus are added to the main menu when an SQL Browser window is active:

D.4.1 Query Menu

Figure 32: The Query menu options

Execute (Ctrl+E or F9) – Opens or runs an SQL statement. If the statement contains parameters, a dialog will pop up to get parameter values before execution starts.

Open (Ctrl+O) – Read in SQL scripts stored on disk.

Save (Ctrl+S) – Save SQL scripts to disk.

Find (Ctrl+F) – Search for a particular word in the current SQL script.

Parameter values (Ctrl+M) – Open the parameter values dialog manually.

Copy To Table – Copy the active result set to a new table.

Print Preview (Ctrl+R) – Opens the current active result set in the Report Engine print preview. See section D.8 for more info about the Report Engine DLL.

Design Report (Ctrl+T) – Opens the Report Engine design view. See section D.8 for more info about the Report Engine DLL.

D.4.2 The View menu

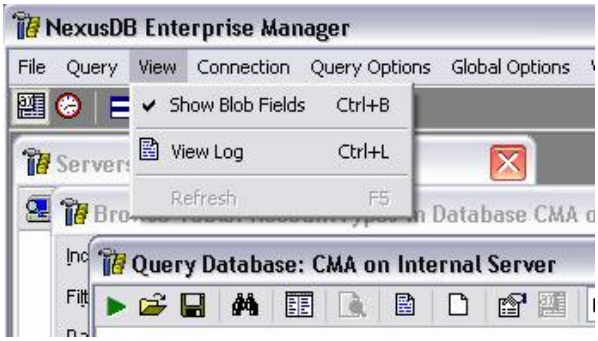

Figure 33: The SQL View menu options

Show Blob Fields (Ctrl+B) – Shows or hides the Blob display area of the SQL Browser.

View Log (Ctrl+L) – Shows or hides the display of the log (also called execution plan).

Refresh (F5) – Reruns the active query.

D.4.3 The Connections menu

Figure 34: The SQL Connection menu options

New (Ctrl+N) – Opens a new SQL query, leaving the existing one in the list of open queries.

Remove (Ctrl+K) – Closes the currently active SQL query.

D.4.4 The Query Options menu

Figure 35: The SQL Connection menu options

Live Dataset (Ctrl+D) – Determines whether the next SELECT statement attempts to return a live (editable) dataset. Turning it on does not guarantee that the result set is editable.

Properties (Ctrl+P) – Opens the query property dialog. In this dialog you can set global defaults for font and timeout (used for all subsequently opened queries), query name (used for the list of active connections), and timeout for the active query.

Debug Log – Turn on logging of errors to the debug log.

D.5 Create/Restructure/View Table structure window

Use the Table Restructure window to view and edit the table dictionary. If you select the View Fields or View Indices options, this screen is in read-only mode. The table can only be restructured by the Redefine menu option. When you select any of these options, the screen shown in Figure 22 below is displayed.

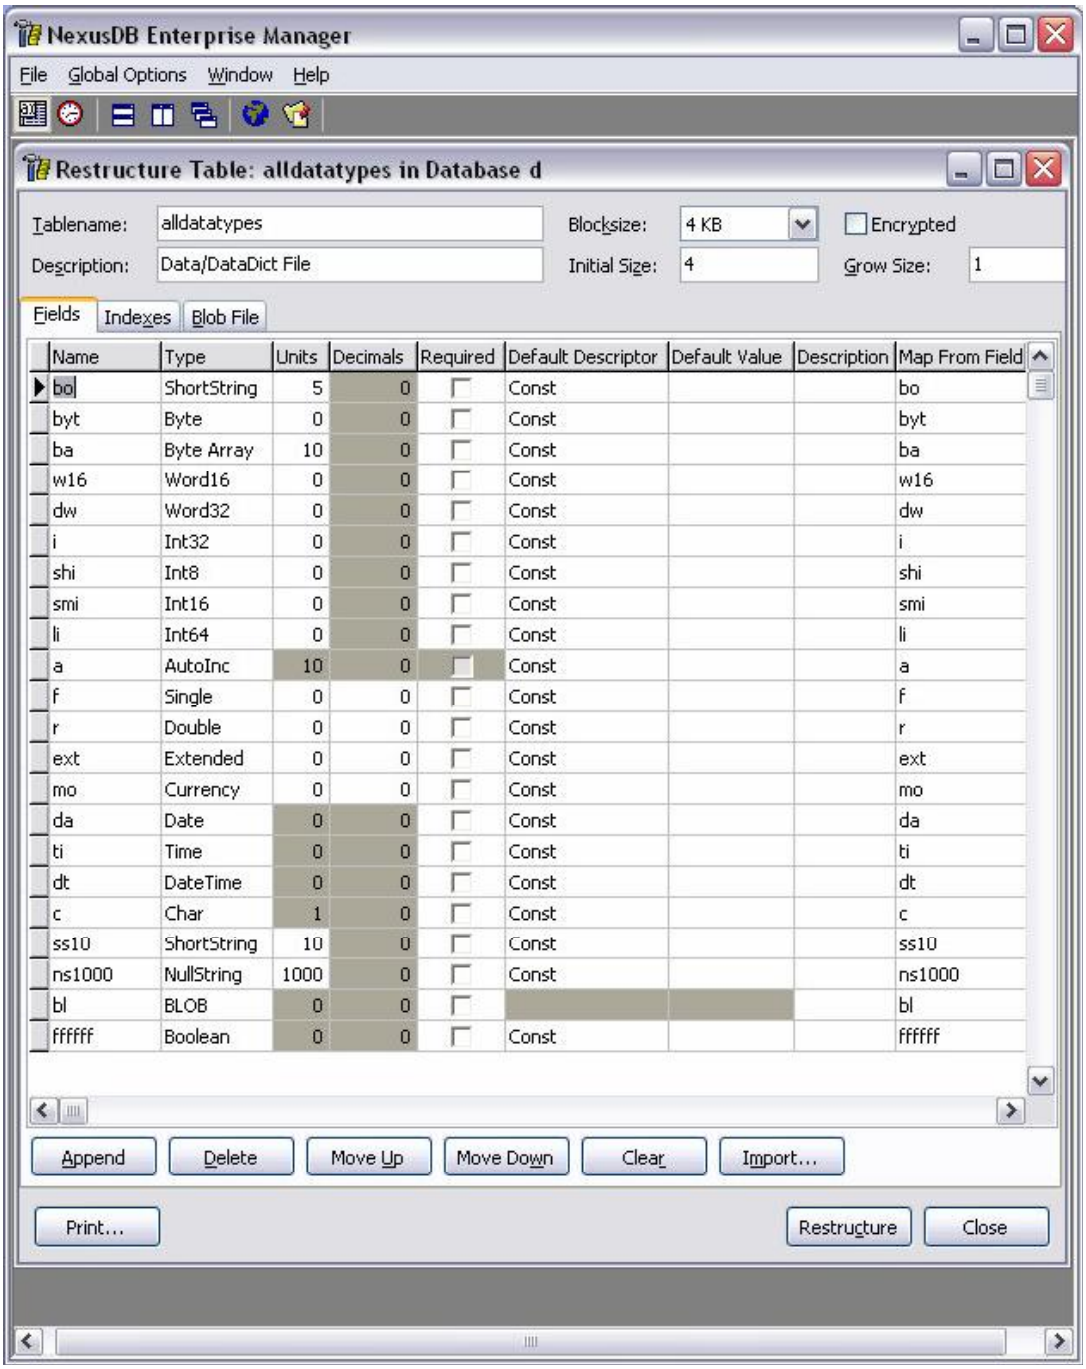

Figure 36: Enterprise Manager – Table Structure Window

The top panel of the window contains a number of header items :

- Tablename can be any valid windows filename. It is limited only by the OS in size and characters (typically up to 255 "standard" filenaming alphanumeric characters).

- Description is stored in the table header and is essentially unlimited. NexusDB ignores any text in this field.

- Blocksize is restricted to values of 4, 8, 16, 32 and 64 kilobytes. The server engine will ensure that each block can hold at least one record and conforms to index block page requirements as well. As a heuristic you should bear in mind that larger block sizes take longer to read/write to disk.

- Initial Size is the number of blocks the table is initially created with. The default for Initial Size is 4, which is the minimum number of blocks NexusDB needs to totally define a table.

- Grow Size is the number of blocks it grows by on disk when more space is needed in the file. To reduce disk fragmentation you should use a large value for Grow Size, especially if you have a table with a large number of indices and/or is being continuously appended to.

- Turning on the Encrypted checkbox instructs the server engine to encrypt the table on disk (it will still, however, be readable to anybody who has a copy of the EM). Tables with Encrypted checked instructs the server engine to encrypt each block as it is written to disk. Blocks read will also be decrypted. The encryption is controlled server side. You do not have to provide any keys. This merely stops others from reading the table by any means other than through a server engine. This is intended to prevent end-users from viewing a table outside of either the EM or your applications. It is not intended to replace the very secure encryption techniques supplied by other vendors.

The larger part of the window is used by three tabs: Fields, Indices, Blob File. The Fields tab lists the fields that make up the table, along with their properties:

- A field Name and a field Type are required. Some types (strings, bytearray) also require a value for size which is entered in the Units column. Field names can be of any conceivable length with no restriction on characters.

- The Decimals column can be used for floating point fields. It is not, however, mandatory for these fields.

- Check the Required tickbox if the field must have a non-null value.

- When inserting records, fields can have non-null default values assigned. For numeric and string fields, set the Default Descriptor to "Const", and enter the required initial value in the Default Value column.

- Description is reserved for your own use and is virtually unlimited in size. The engine simply stores to the table dictionary whatever you enter here.

- Map from Field is only visible when you restructure tables. It displays the field map (on a field by field basis as shown) used when the table is restructured.

The indices tab is where maintenance of the indices is performed. On clicking this tab the window shown in Figure 23 below is displayed.

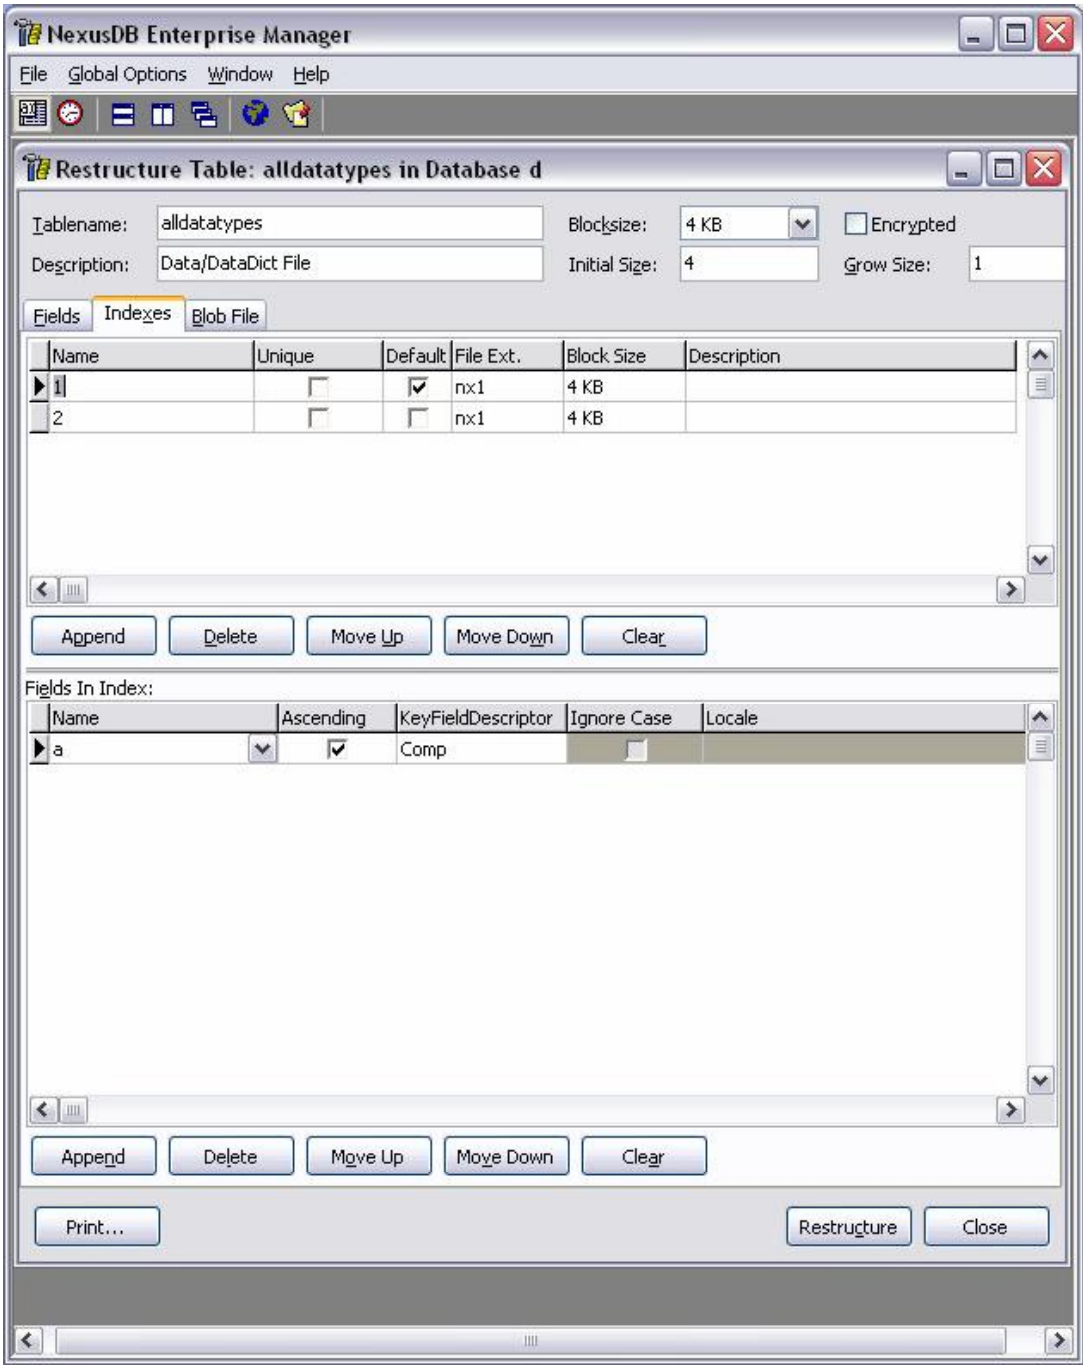

Figure 37: Enterprise Manager– Index creation and maintenance

The Indices tab allows you to set and define indices for your table from the fields entered. This consists of 2 parts for each index. The index header information is entered into the top grid on this tab. The fields for each index are entered in the lower grid. Note that you must enter at least one

field for each index as you proceed. (That is, you can not simply enter all the indices and then come back and add the fields for each index as a 'second pass'.)

Each index must be given a name (again, you choice for length) as a minimum entry in the top grid. If all keys in this index are unique, then check the Unique column. Only one index can be checked as the Default index. (This index is selected by the server engine if you do not set an IndexName, or IndexFieldNames, when opening a table in your applications. If no index is flagged as the Default, then the Sequential Access Index becomes the Default index.).

To save the index in an external file, provide a valid file extension name in the File Ext column. (Normal filename extension conventions apply.) If you direct the index to be stored externally, then you may set a block size for this external file in the Block Size column.

The Description column is a optional place for you to enter whatever text you would normally like to put comments related to each index. Again, this is a virtually size-unlimited field.

The fields in each index contain the following values:

A valid (existing) field name must be entered. This is selectable from a dropdown list.

NexusDB allows you to set a field based sort order. The default is Ascending, but you are not restricted to having all fields in an index as either ascending or descending.

The value entered into the KeyFieldDescriptor describes which sorting engine you will be using for this field. The default is Comp (which is for non-text fields), Text is the non-ANSI comparisons (which is basically just a byte by byte comparison) and ExtText is the ANSI comparison.

Check the Ignore Case field if you don't want case sensitivity in your sorting.

The other values that can be set for each field of an index are only used for the ANSI (ExtText) comparison. This is discussed elsewhere in the manual, see also the Win32API help file in Delphi – topic : Compare String.

Blobs may optionally be stored in an external file. To do so you merely have to enable the use of an external file and then set the file extension, block size, initial size and grow size. These have the same meanings as for the main file and for external index files.

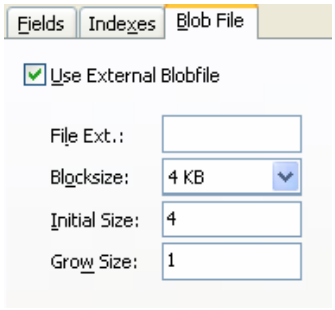

Figure 38: Enterprise Manager – Blob File parameters

D.6 The Registered Servers dialog

Use this option to register a server with the EM that, for instance, you can not reach by broadcast. Enter the details required and click the Add button.

To remove a registered server you need only select it and then click the Remove button.

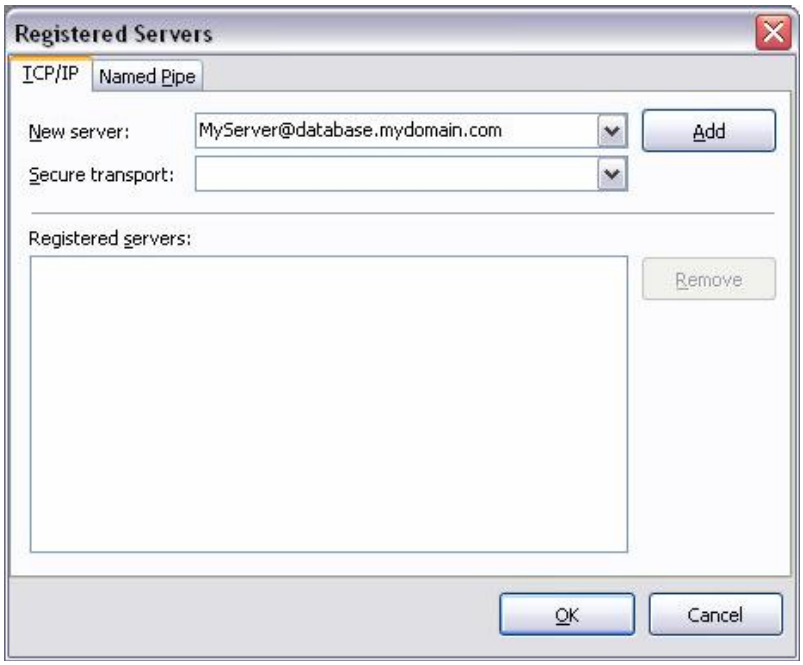

Figure 7: Enterprise Manager – Registered Servers dialog

To use a BlowFish secured server, enter the servername into the New Server box, and select the secure transport in the combobox below it. Press Add, and you have created a securely wrapped transport setup.

When you exit the dialog by pressing OK, any existing connections are closed, and the tree of servers in the Servers window is regenerated.

D.7 The CSV Import wizard

This wizard allow you to import a wide range of textfiles, both delimited and fixed width formats.

To ensure proper handling of date/time formats, make sure to set the date/time options in the Advanced settings dialog, reachable on the last page of the import wizard, before pressing the "New Table" button on that wizard page.

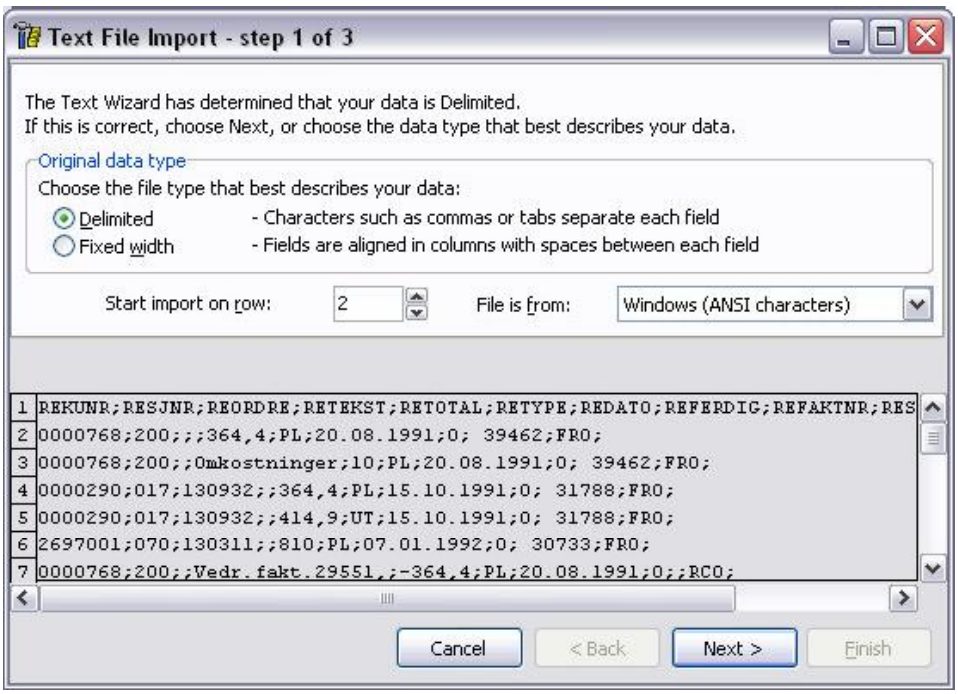

Figure 8: Enterprise Manager – CSV Import Wizard page 1

The first page of the wizard displays the unformatted contents of the start of the file. The wizard also tries to analyze the start of the file to determine what type of file it is. Most of the time it will guess correctly, but you should check that the Delimited or Fixed width selection matches the file contents.

Also, make sure to change the "Start import on row" setting if there are blank lines or descriptive column names at the top of the file, to avoid importing blank or erroneous data.

Note that if "Start import on row" is larger than 1, then the column descriptions (if any) will be used to create fieldnames in the "New table" function on page 3.

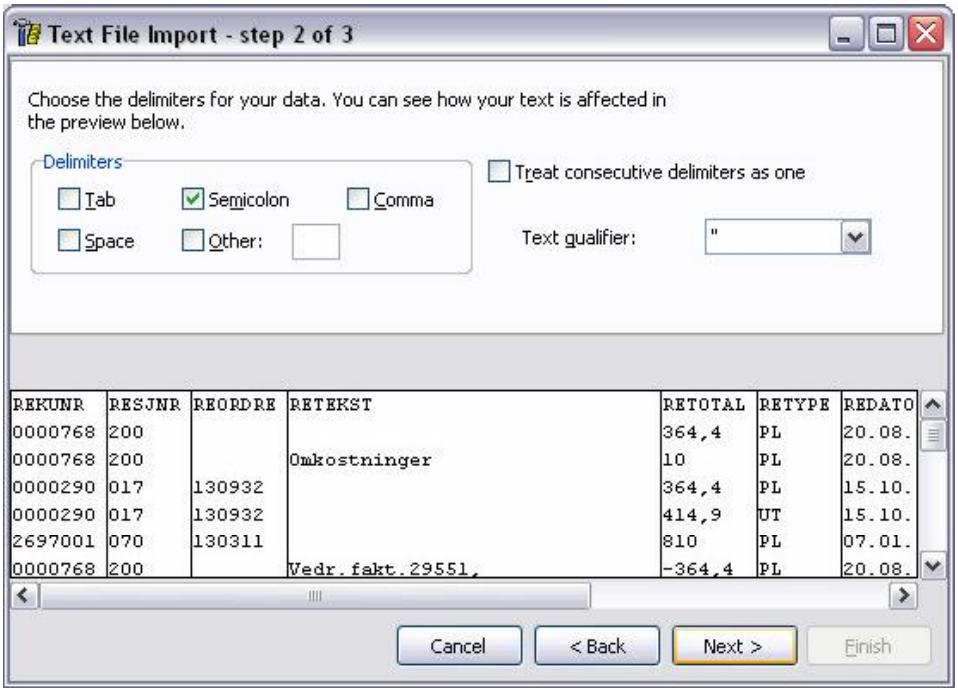

Figure 9: CSV Import Wizard page 2

In the second wizard page, select all delimiters used in the file you are importing from. The Other setting can be used if an uncommon delimiter has been used.

The preview of the columns changes to indicate how the options affect the import.

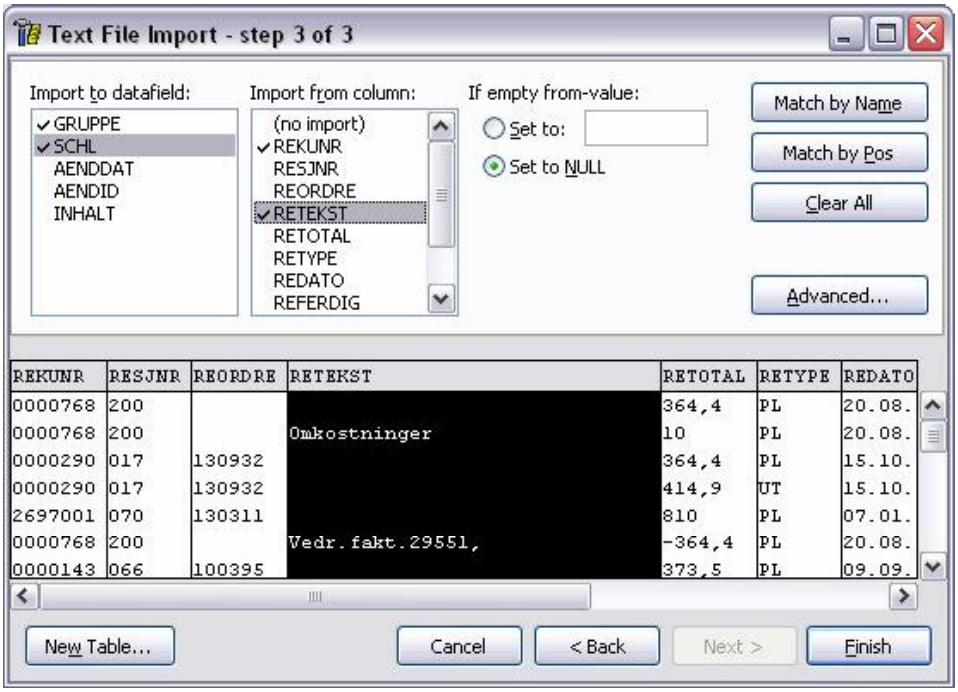

Figure 10: CSV Import Wizard page 3

On the third and last wizard page, set formatting options, and link table fields with import columns.

The Import to datafield list contains the list of fields in the table you are importing to, or a generic numbered name (field 1 etc) if you haven't created a table yet.

The Import from column lists the columns in the text file. If there are column descriptions in the first line of the file, and you chose "Start import on row" to be larger than 1 on the first wizard page, then the descriptions are used, otherwise generic numbered column names are displayed.

To link a field in the table with a column, click the field in the fieldlist first, and then click the column you want to import into the field. If you make a wrong selection and want to remove the column, click the (no import) item at the top of the column list. Checkmarks appear next to fields and columns that have been selected.

The column preview will display the chosen column and invert its colour. You may select columns by clicking either the Import from column list, or directly in the preview.

Before linking individual fields and creating a new table, you should open the Advanced settings dialog and set the options to match what is contained in the CSV file. Also, to speed up importing, it is a good idea to use a higher number for Records per transaction (recommend 1.000-10.000).

The New Table button will bring up the Create Table dialog, with fields, fieldtypes and units preset from the values detected in the CSV file. Ensure that text fields are large enough to hold the longest text you have in a column etc. The detection process only looks at the start of a file, so there might be longer text strings in a column than what is detected from the first part of the file. If you successfully Create the table, the new table is now the one being imported into.

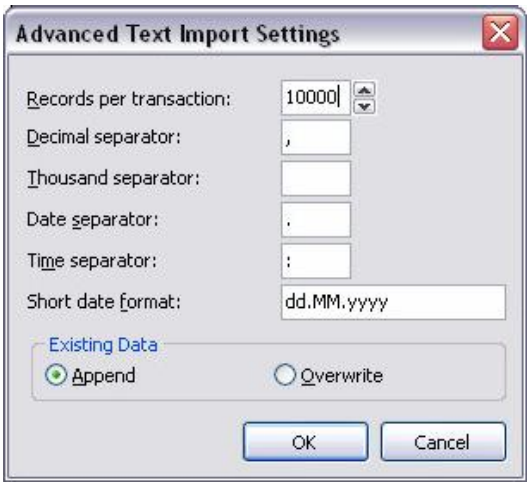

Figure 11: CSV Import Wizard – Advanced settings dialog

D.8 The Report Engine DLL

The Report Engine DLL is used for all printouts. The EM will recognize if a Report Engine DLL is present in the same directory as the EM exe. If nxEMReportEngine.DLL is present, then the following print-related items will be enabled:

1. The Table Browser – Print Preview and Design Report menu options will be enabled
2. SQL Query Browser – same options as above are available under Query
3. Dictionary Browser – the Print button on the lower left-hand side of the window is enabled

The DLL downloadable from the Nexus Website is based on FastReports. Full source for the DLL is included if you want to add support for another report engine.