

Jaybird JCA/JDBC Driver



Release Notes v 2.1.1

Table of Contents

[General Notes](#)

[Supported Firebird versions](#)

[Specification support](#)

[What's new in Jaybird 2.1](#)

[Events](#)

[Support for multiple JNI libraries](#)

[Default holdable result sets](#)

[Other features](#)

[What's new in Jaybird 2.0](#)

[Refactorings](#)

[Updatable result sets](#)

[Firebird management interfaces](#)

[Incompatible changes since Jaybird 2.0](#)

[Jaybird JDBC extensions](#)

[JDBC 3.0 compatibility](#)

[Distribution package](#)

[License](#)

[Source Code](#)

[JDBC URL Format](#)

[Pure Java](#)

[Using Firebird client library](#)

[Embedded Server](#)

[JDBC connection properties](#)

[Using Type 2 and Embedded Server driver](#)

[Configuring Type 2 JDBC driver](#)

[Configuring Embedded Server JDBC driver](#)

[JDBC 3.0 Compatibility](#)

[JDBC 3.0 deviations and unimplemented features](#)

[Jaybird Specifics](#)

[Result sets](#)

[Using java.sql.ParameterMetaData with Callable Statements](#)

[Using ResultSet.getCharacterStream with BLOB fields](#)

[Heuristic transaction completion support](#)

[Connection pooling with Jaybird 2.0](#)

[Usage scenario](#)

[Connection Pool Classes](#)

[org.firebirdsql.pool.FBConnectionPoolDataSource](#)

[org.firebirdsql.pool.FBWrappingDataSource](#)

[Runtime object allocation and deallocation hints](#)

[Documentation and Support](#)

[Where to get more information on Jaybird](#)

[Where to get help](#)

[Reporting Bugs](#)

[Corrections/Additions To Release Notes](#)

General Notes

Jaybird is JCA/JDBC driver suite to connect to Firebird database server. Historically Borland opened sources of type 3 JDBC driver called InterClient. However due to some inherent limitations of Firebird client library it was decided that type 3 driver is a dead end, and Firebird team developed pure Java implementation of wire protocol. This implementation became basis for Jaybird, pure Java driver for Firebird relational database.

This driver is based on both the new JCA standard for application server connections to enterprise information systems and the well known JDBC standard. The JCA standard specifies an architecture in which an application server can cooperate with a driver so that the application server manages transactions, security, and resource pooling, and the driver supplies only the connection functionality. While similar to the JDBC 2 `XADataSource` idea, the JCA specification is considerably clearer on the division of responsibility between the application server and driver.

Supported Firebird versions

Jaybird supports Firebird 1.0.x, Firebird 1.5 and Firebird 2.0 SuperServer and Classic regardless of the platform on which server runs when type 4 JDBC driver is used. Type 2 and embedded server JDBC drivers require JNI library, precompiled binaries for Win32 and Linux platforms are shipped in the default installation, other platforms require porting/building JNI library for that platform.

Driver no longer supports InterBase servers due to a Firebird-specific changes in database attachment parameters that are send to the server.

Specification support

Driver supports following specifications:

JDBC 3.0	Driver passed complete JDBC compatibility test suite, though some features are not implemented. It is not officially JDBC compliant, because of high certification costs.
JDBC 2.0 Standard Extensions	Jaybird provides implementation of following interfaces from <code>javax.sql.*</code> package: <ul style="list-style-type: none">• <code>ConnectionPoolDataSource</code> implementation provides connection and prepared statement pooling.• <code>DataSource</code> implementation provides seamless integration with major web and application servers.• <code>XADataSource</code> implementation provides means to use driver in distributed transactions.
JCA 1.0	Jaybird provides implementation of <code>javax.resource.spi.ManagedConnectionFactory</code> and related interfaces. CCI interfaces are not supported.
JTA 1.0.1	Driver provides implementation of <code>javax.transaction.xa.XAResource</code> interface via JCA framework and <code>XADataSource</code> implementation.
JMX 1.2	Jaybird provides MBean to manage Firebird server and installed databases via JMX agent.

What's new in Jaybird 2.1

List of bugs fixed since JayBird 2.1.0

JDBC-92	DatabaseMetaData becomes unusable when the statement that serves the result set is closed.
JDBC-87	Backported feature of upcoming Jaybird 2.2 release that allows specifying the <code>octetsAsBytes</code> connection property and receive <code>byte[]</code> when using the <code>ResultSet.getObject()</code> method
JDBC-85	Fixed returning of the error messages.
-	Bug reported in Firebird-Java group. <code>ClassCastException</code> when obtaining <code>byte[]</code> from BLOB fields in cached result set.

Jaybird 2.1 introduces following new features:

Events

Events is one of the unique features in the Firebird RDBMS and allows asynchronous notification of the applications about named events that happen in the database. The information about how to use this feature can found in free IB 6.0 documentation set as well as in The Firebird Book by Helen Borrie.

The interfaces and classes for the event support can be found in `org.firebirdsql.event` package, which includes:

- `EventManager` interface to register for the synchronous and asynchronous notification about the events in the database;
- `EventListener` interface which has to be implemented by the application that wants to participate in the asynchronous notification;
- `DatabaseEvent` interface which represents the object that will be passed to the `EventListener` notification method;
- Implementation of the above interfaces: `FBEventManager` and `FBDatabaseEvent`.

Please note, that each instance of `FBEventManager` will open a new socket connection to the Firebird server on the port specified by the Firebird.

Similar to other JDBC extensions in Jaybird, the interfaces are released under the modified BSD license, the implementation of the code is released under LGPL license.

Support for multiple JNI libraries

Until this release only one client library could be loaded in the single JVM. That could be either embedded Firebird library (`fbembed.dll/libfbembed.so`), or Firebird client library (`fbclient.dll/libfbclient.so`) or Fyracle client library. This could lead to possible problems, for example, if embedded Firebird was used first, JDBC driver would access database file directly instead of using the local IPC protocol if only path to the database was specified. It was not possible to change this without restarting the JVM.

Now Jaybird is able to correctly load arbitrary number of shared libraries that implement the ISC API and forward the requests correctly depending on the type of the driver being used. Note, this feature requires new JNI library (`jaybird21.dll/libjaybird21.so`)

Note, the JNI library (jaybird21.dll/libjaybird21.so) was changed in a incompatible manner between Jaybird 2.1 beta 1 and Jaybird 2.1 RC1 releases. Please ensure to use the latest version.

Default holdable result sets (closed ResultSet in auto-commit mode)

New connection property that allows to create holdable result sets by default. This is needed as an workaround for the applications that do not follow JDBC specification in regard to the auto-commit mode.

Specifically, such applications open a result set and, while traversing it, execute other statements using the same connection. According to JDBC specification the result set has to be closed if another statement is executed using the same connection in auto-commit mode. Among others the OpenOffice Base users have problems with the restored JDBC 3.0 compatibility in Jaybird 2.0.

The new property is called:

- `defaultResultSetHoldable` as connection property for JDBC URL or for `java.sql.DriverManager` class and no or empty value should be assigned to it; it has an alias `defaultHoldable` to simplify the typing;
- `isc_dpb_result_set_holdable` as a DPB member;
- `FirebirdConnectionProperties` interface was with two methods extended (the `isDefaultResultSetHoldable()` and `setDefaultResultSetHoldable(boolean)`) which is also reflected in the `FBConnectionPoolDataSource` and `FBWrappingDataSource` classes.

Note, the price for using this feature is that each holdable result set will be fully cached in memory. The memory occupied by it will be released when the statement that produced the result set is either closed or re-executed.

Other features

`StatisticsManager` now provides methods to fetch statistics for the specified tables only (like the `-t` switch in `gstat`) and to provide also analysis of average record and version length.

What's new in Jaybird 2.0

Jaybird 2.0 underwent significant internal changes that are not directly visible to the client code, and not necessarily introduced new features. However those changes significantly increase code modularity, stability and maintainability.

Refactorings

- Significant refactoring of the `org.firebirdsql.gds.*` package. Implementations of the interfaces defined in this package are now dynamically loaded from the classpath and provide possibilities to extend driver with custom code without changing the driver itself.
- JCA code was significantly changed, bringing a long awaited stability when critical errors happen into that vital part of code. There is no more so-called “fatal” errors that can cause driver to recycle connections and transactions.
- Refactoring of the transaction and object life-cycle management in the JDBC code. Main reason for this change was the need to remove result set caching in auto-commit mode and make it JDBC specification compliant, but it has also positive effect on the code stability and maintainability.
- Unified connection property handling in JCA, driver manager and pooling code. This solves a longstanding issue with different configuration possibilities in different modes.

- Refactorings to support JDK 5.0.

Updatable result sets

Jaybird provides now support for the updatable result sets. Feature allows Java application to update current record using the `updateXXX` methods of `java.sql.ResultSet` interface. Updates are performed within the current transaction using a best row identifier in `WHERE` clause. This sets the following limitation on the result set “updatability”:

- the `SELECT` references single table;
- all columns not referenced in `SELECT` permit `NULLS` (otherwise `INSERTS` will fail);
- the `SELECT` statement does not contain `DISTINCT` predicate, aggregate functions, joined tables or stored procedures;
- the `SELECT` statement references all columns from the table primary key definition or an `RDB$DB_KEY` column.

Firebird management interfaces

Jaybird provides full support of the Firebird Services API that allows Java applications to perform various server management tasks:

- database backup/restore on remote server; it is possible to performs metadata-only backups, switch the garbage collection during backup off, restore databases with no validity constraints or active indices, etc.
- database maintenance, e.g. database shutdown, sweep, changing the forced writes settings, changing SQL dialect of the database, shadow management, etc.
- retrieving database statistics including header page statistics, system table statistics, data page statistics and index statistics.
- user management, including adding, modifying, and deleting user accounts.

Incompatible changes since Jaybird 1.5

- Signatures of all `org.firebirdsql.gds.*` classes have changed. This was done in preparation of the official Firebird Java API. After the API release all signature changes will be frozen and correctly versioned.
- `FirebirdConnection.getBlob()` method signature changed to return `java.sql.Blob`, previously it returned `org.firebirdsql.jdbc.FirebirdBlob`. Object that is returned can safely be casted to `org.firebirdsql.jdbc.FirebirdBlob` interface. The change is caused by the forthcoming JDBC 4.0 specification to avoid conflict in next versions.
- Jaybird 1.x allowed to execute queries in auto-commit mode, loop through the obtained result set and to execute other statements through the same connection. This is no longer possible in Jaybird without changing the result set holdability. The behavior was changed in order to conform the JDBC 3.0 specification, chapter 10.1 “Transaction Boundaries and Auto-commit”. Applications willing to use this feature have to set the result set holdability to `ResultSet.HOLD_CURSORS_OVER_COMMIT`. Please note, that this option assumes that result set type is `TYPE_SCROLL_INSENSITIVE` and is completely cached on the client side.

Jaybird JDBC extensions

Jaybird provides extensions to some JDBC interfaces. JDBC extension interface classes are released under modified BSD license, on “AS IS” and “do what you want” basis, this should make linking to these classes safe from the legal point of view. All classes belong to `org.firebirdsql.jdbc.*` package. Table below shows all JDBC extensions present in Jaybird with a driver version in which

the extension was introduced.

JDBC extensions			
Interface	Since	Method name	Description
FirebirdDriver	2.0	newConnectionProperties()	Create new instance of FirebirdConnectionProperties interface that can be used to set connection properties programmatically.
		connect(FirebirdConnectionProperties)	Connect to the Firebird database using the specified connection properties.
FirebirdConnectionProperties	2.0	see JDBC connection properties section for more details.	
FirebirdConnection	1.5	createBlob()	Create new BLOB in the database. Later this BLOB can be passed as a parameter into PreparedStatement or CallableStatement.
	1.5	getIsCEncoding()	Get connection character encoding.
	2.0	getTransactionParameters(int isolationLevel)	Get the TPB parameters for the specified transaction isolation level.
	2.0	createTransactionParameterBuffer()	Create an empty transaction parameter buffer.
	2.0	setTransactionParameters(int isolationLevel, TransactionParameterBuffer tpb)	Set TPB parameters for the specified transaction isolation level. The newly specified mapping is valid for the whole connection lifetime.
	2.0	setTransactionParameters(TransactionParameterBuffer tpb)	Set TPB parameters for the specified transaction isolation level. The newly specified parameters are effective until the transaction isolation is changed.
FirebirdDatabaseMetaData		getProcedureSourceCode(String)	Get source code for the specified stored procedure name.
		getTriggerSourceCode(String)	Get source code for the specified trigger name.
		getViewSourceCode(String)	Get source code for the specified view name.
FirebirdStatement	1.5	getInsertedRowCount() getUpdatedRowCount() getDeletedRowCount()	Extension that allows to get more precise information about outcome of some

JDBC extensions			
			statement.
	1.5	<code>hasOpenResultSet()</code>	Check if this statement has open result set. Correctly works only when auto-commit is disabled. Check method documentation for details.
	1.5	<code>getCurrentResultSet()</code>	Get current result set. Behaviour of this method is similar to the behavior of the <code>Statement.getResultSet()</code> , except that this method can be called as much as you like.
	1.5	<code>isValid()</code>	Check if this statement is still valid. Statement might be invalidated when connection is automatically recycled between transactions due to some irrecoverable error.
	2.0	<code>getLastExecutionPlan()</code>	Get execution plan for the last executed statement.
FirebirdPreparedStatement	2.0	<code>getExecutionPlan()</code>	Get the execution plan of this prepared statement.
	2.0	<code>getStatementType()</code>	Get the statement type of this prepared statement.
FirebirdCallableStatement	1.5	<code>setSelectableProcedure(boolean selectable)</code>	Mark this callable statement as a call of the selectable procedure. By default callable statement uses "EXECUTE PROCEDURE" SQL statement to invoke stored procedures that return single row of output parameters or a result set. In former case it retrieves only the first row of the result set.
FirebirdResultSet	2.0	<code>getExecutionPlan()</code>	Get execution plan for this result set.
FirebirdBlob	1.5	<code>detach()</code>	Method “detaches” a BLOB object from the underlying result set. Lifetime of “detached” BLOB is limited by the lifetime of the connection.
	1.5	<code>isSegmented()</code>	Check if this BLOB is segmented. Seek operation is not defined for the segmented

JDBC extensions			
			BLOBs.
	1.5	<code>setBinaryStream(long position)</code>	Opens an output stream at the specified position, allows modifying BLOB content. Due to server limitations only position 0 is supported.
<code>FirebirdBlob.BlobInputStream</code>	1.5	<code>getBlob()</code>	Get corresponding BLOB instance.
	1.5	<code>seek(int position)</code>	Change the position from which BLOB content will be read, works only for stream BLOBs.
<code>FirebirdSavepoint</code>	2.0	interface is equivalent to the <code>java.sql.Savepoint</code> interface introduced in JDBC 3.0 specification, however allows using Firebird savepoints also in JDBC 2.0 (JDK 1.3.x) applications.	

JDBC 3.0 compatibility

Jaybird includes number of fixes that allow it pass JDBC 3.0 compatibility suite. It successfully passes 1216 tests, 60 tests were excluded, because they are either not applicable to Firebird or fail due to some server problems (math rounding issues, limitations of NUMERIC data type, etc.).

Distribution package

Jaybird driver has compile-time and run-time dependencies to JCA 1.0, JTA 1.0.1, JAAS 1.0 and JDBC 2.0 Optional Package. Additionally, if Log4J classes are found in the class path, it is possible to enable extensive logging inside the driver.

Following file groups can be found in distribution package:

- `Jaybird-2.1.0.jar` – archive containing JCA/JDBC driver and JMX management class. It requires JCA 1.0, JTA 1.0.1, and JAAS 1.0.
- `Jaybird-pool-2.1.0.jar` – archive contains implementation of connection pooling and statement pooling interfaces.
- `Jaybird-full-2.1.0.jar` – merge of `firebirdsql.jar` and `mini-j2ee.jar`. This archive can be used for standalone¹ Jaybird deployments.
- `Jaybird-2.1.0.rar` – resource archive ready for deployment in JCA-enabled application servers.
- `lib/jaas.jar` – archive containing JAAS 1.0 classes.
- `lib/log4j-core.jar` – archive containing core Log4J classes that provide a possibility to log into the file.
- `lib/mini-j2ee.jar` – archive containing JCA 1.0, JTA 1.0.1 and JDBC 2.0 Optional Package classes.
- `jaybird21.dll` – Windows version of the JNI library for Type 2 and Embedded Server drivers.

¹ You have to ensure that your class path contains JAAS 1.0 classes when using JDK 1.3.x.

- `libjaybird21.so` – Linux version of the JNI library for Type 2 and Embedded Server drivers.

License

Jaybird JCA/JDBC driver is distributed under the GNU Lesser General Public License (LGPL). Text of the license can be obtained from <http://www.gnu.org/copyleft/lesser.html>. Using Jaybird (by importing Jaybird's public interfaces in your Java code), and extending Jaybird by subclassing or implementation of an extension interface (but not abstract or concrete class) is considered by the authors of Jaybird to be dynamic linking. Hence our interpretation of the LGPL is that the use of the unmodified Jaybird source does not affect the license of your application code.

Even more, all extension interfaces to which application might want to link are released under dual LGPL/modified BSD license. Latter is basically “AS IS” license that allows any kind of use of that source code. Jaybird should be viewed as an implementation of that interfaces and LGPL section for dynamic linking is applicable in this case.

Source Code

Source code can be obtained from the CVS at SourceForge.net. The CVSROOT is `:pserver:anonymous@cvs.sourceforge.net:/cvsroot/firebird`, the module name is `client-java`. Alternatively source code can be viewed online at <http://cvs.sourceforge.net/viewcvs.py/firebird/client-java/>

JDBC URL Format

Driver provides different JDBC URLs for different usage scenarios:

Pure Java

```
jdbc:firebirdsql:host[/port]:/path/to/db.fdb  
jdbc:firebirdsql://host[:port]/path/to/db.fdb
```

Default URL, will connect to the database using type 4 JDBC driver. Best suited for client-server applications with dedicated database server. Port can be omitted (default value is 3050), host name must be present.

First format is considered official, second – compatibility mode for InterClient migration.

Using Firebird client library

```
jdbc:firebirdsql:native:host[/port]:/path/to/db.fdb  
jdbc:firebirdsql:native://host[:port]/path/to/db.fdb
```

Type 2 driver, will connect to the database using client library (either `fbclient.dll` or `gds32.dll` on Windows, and `libfbclient.so` or `libgds.so` on Linux). Requires correct installation of the client library.

```
jdbc:firebirdsql:local:/path/to/db.fdb
```

Type 2 driver in local mode. Uses client library as in previous case, however will not use socket communication, but rather access database directly. Requires correct installation of the client library.

Embedded Server

```
jdbc:firebirdsql:embedded:/path/to/db.fdb
```

Similar to the Firebird client library, however `fbembed.dll` on Windows and `libfbembed.so` on Linux are used. Requires correctly installed and configured Firebird embedded server.

JDBC connection properties

Table below contains properties that specify parameters of the connections that are obtained from this data source. Commonly used parameters have the corresponding getter and setter methods, rest of the Database Parameters Block parameters can be set using `setNonStandardProperty` setter method.

Property	Getter	Setter	Description
database	+	+	Path to the database in the format [host/port:]/path/to/database.fdb
type	+	+	Type of the driver to use. Possible values are: <ul style="list-style-type: none">• PURE_JAVA or TYPE4 for type 4 JDBC driver• NATIVE or TYPE2 for type 2 JDBC driver• EMBEDDED for using embedded version of the Firebird.
blobBufferSize	+	+	Size of the buffer used to transfer BLOB content. Maximum value is 64k-1.
socketBufferSize	+	+	Size of the socket buffer. Needed on some Linux machines to fix performance degradation.
buffersNumber	+	+	Number of cache buffers (in database pages) that will be allocated for the connection. Makes sense for ClassicServer only.
charSet	+	+	Character set for the connection. Similar to <code>encoding</code> property, but accepts Java names instead of Firebird ones.
encoding	+	+	Character encoding for the connection. See Firebird documentation for more information.
useTranslation	+	+	Path to the properties file containing character translation map.
password	+	+	Corresponding password.
roleName	+	+	SQL role to use.
userName	+	+	Name of the user that will be used by default.

Property	Getter	Setter	Description
useStreamBlobs	+	+	Boolean flag tells driver whether stream BLOBs should be created by the driver, by default “false”. Stream BLOBs allow “seek” operation to be called, however due to a bug in gbak utility they are disabled by default.
useStandardUdf	+	+	Boolean flag tells driver to assume that standard UDFs are defined in the database. This extends the set of functions available via escaped function calls. This does not affect non-escaped use of functions.
defaultResultSetHoldable	+	+	Boolean flag tells driver to construct the default result set to be holdable. This prevents it from closing in auto-commit mode if another statement is executed over the same connection.
tpbMapping	+	+	TPB mapping for different transaction isolation modes.
defaultIsolation	+	+	Default transaction isolation level. All newly created connections will have this isolation level. One of: <ul style="list-style-type: none"> TRANSACTION_READ_COMMITTED TRANSACTION_REPEATABLE_READ TRANSACTION_SERIALIZABLE
defaultTransactionIsolation	+	+	Integer value from <code>java.sql.Connection</code> interface corresponding to the transaction isolation level specified in <code>isolation</code> property.
nonStandardProperty	<code>getNonStandardProperty(String)</code>	+ <code>setNonStandardProperty(String)</code> <code>setNonStandardProperty(String, String)</code>	<p>Allows to set any valid connection property that does not have corresponding setter method. Two setters are available:</p> <p><code>setNonStandardProperty(String)</code> method takes only one parameter in form “<code>propertyName[=propertyValue]</code>”, this allows setting non-standard parameters using configuration files.</p> <p><code>setNonStandardProperty(String, String)</code> takes property name as first parameter, and its value as the second</p>

Property	Getter	Setter	Description
			parameter.

Using Type 2 and Embedded Server driver

Jaybird 2.1 provides type 2 JDBC driver that uses native client library to connect to the databases. Additionally Jaybird 2.1 can use embedded version of Firebird relational database allowing to create Java applications that does not require separate server setup.

However type 2 driver has also limitations:

- Due to multi-threading issues in Firebird client library as well as in embedded server version, it is not possible to access them from different threads simultaneously. When using client library only one thread is allowed to access connection at a time, however it is allowed to access different connections from different threads. Client library in local mode and embedded server library on Linux do not allow multithreaded access to the library. Jaybird provides necessary synchronization in Java code, however corresponding mutex is local to the classloader that loaded Jaybird driver. **Care should be taken when deploying applications in web or application servers: put jar files in the main library directory of the web and/or application server, not in the library directory of the web or enterprise application (WEB-INF/lib directory or in the .EAR file). This issue will be fixed in Firebird client library after merge with Vulcan project.**
- Current implementation of the JNI library allows that either client library or embedded server library is loaded into memory. The first library being loaded wins, attempt to load second library will produce an error. Usually this is not a problem, since embedded server library can also act as a client library, so it is enough to load embedded server library first. This limitation will be removed in Jaybird 2.1, implementation already exists, however it was contributed already at RC stage; adding it to the release would delay it for few months.

Configuring Type 2 JDBC driver

Type 2 JDBC driver requires JNI library to be installed and available for Java Virtual Machine. Precompiled binaries for Windows and Linux platforms are distributed with Jaybird:

Please note that Jaybird 2.1 provides update to the JNI libraries to support new features. It is not compatible with JNI library for Jaybird 2.0.

- `jaybird21.dll` is precompiled binary for Windows platform. Successfully tested with Windows 2000 and Windows XP SP1, but there should be no issues also in other Win32 OS. Library should be either copied into the directory specified in `%PATH%` environment variable, or made available to JVM using the `java.library.path` system property.
- `libjaybird21.so` is precompiled binary for Linux platform. It must be available via the `LD_LIBRARY_PATH` environment variable, e.g. copied into `/usr/lib/` directory. Another possibility is to specify path to the directory with Jaybird JNI library in `java.library.path` system property during the JVM startup.
- Other platforms can easily compile the JNI library by checking out the Jaybird sources from the CVS and using `./build.sh compile-native` command in the directory with checked out sources.

After making Jaybird JNI library available to the JVM application has to tell driver to start using this by either specifying `TYPE2` or `LOCAL` type in the connection pool or data source properties or using appropriate JDBC URL when connecting via `java.sql.DriverManager`.

Configuring Embedded Server JDBC driver

Embedded Server JDBC driver uses same JNI library and configuration steps for the type 2 JDBC driver.

There is however one issue related to the algorithm of Firebird Embedded Server installation directory resolution. Firebird server uses pluggable architecture for internationalization. By default server loads `fbintl.dll` or `libfbintl.so` library that contains various character encodings and collation orders. This library is expected to be installed in the `intl/` subdirectory of the server installation. The algorithm of directory resolution is the following:

1. `FIREBIRD` environment variable.
2. `RootDirectory` parameter in the `firebird.conf` file.
3. The directory where server binary is located.

When Embedded Server is used from Java and no `FIREBIRD` environment variable is specified, it tries to find `firebird.conf` in the directory where application binary is located. In our case application binary is JVM and therefore Embedded Server tries to find its configuration file in the `bin/` directory of the JDK or JRE installation. Same happens to the last item of the list. In most cases this is not desired behavior.

Therefore, if application uses character encodings, UDFs or wants to fine-tune server's behavior through the configuration file, the `FIREBIRD` environment variable must be specified and point to the installation directory of the Embedded Server, e.g. current working directory.

JDBC 3.0 Compatibility

As it was mentioned before, Jaybird JCA/JDBC driver passed Sun JDBC CTS 1.3.1 test suite. All tests except those that do not apply to Firebird RDBMS succeeded. However driver is not officially JDBC-compliant because certification procedure is too expensive.

JDBC 3.0 deviations and unimplemented features

The following optional features and the methods for their support are not implemented:

- `java.sql.Array` data type is not yet supported.
- `java.sql.Blob` does not implement following methods:
 - `position(Blob, long)` and `position(byte[], long)`; Firebird does not provide any server-side optimization for these calls, client application must fetch complete BLOB content from the server to do pattern search.
 - `truncate(long)`; Firebird does not provide such functionality on the server side, application must fetch old BLOB from the server and pump old content into a newly created BLOB.
- `java.sql.Clob` data type is not yet supported.
- `java.sql.Connection`
 - `getCatalog()` and `setCatalog(String)` are not supported by Firebird server.
 - `getTypeMap()` and `setTypeMap(Map)` are not supported.
 - `prepareStatement(String, int)`, `prepareStatement(String, int[])` and `prepareStatement(String, String[])` that return auto-generated keys are not implemented, because this functionality is not provided by server.
- `java.sql.PreparedStatement`
 - `setObject(int index, Object object, int type)` Target SQL type is

determined from the class of the passed object and corresponding parameter is ignored.

- `setObject(int index, Object object, int type, int scale)` Same as above, type and scale are ignored.
- `java.sql.Ref` data type is not supported by Firebird server.
- `java.sql.SQLData` data type is not supported by Firebird server.
- `java.sql.SQLInput` is not supported.
- `java.sql.SQLOutput` is not supported.
- `java.sql.Statement`
 - `cancel()` is not supported by Firebird server.
 - `execute(String, int)`, `execute(String, int[])`, `execute(String, String[])`, `executeUpdate(String, int)`, `executeUpdate(String, int[])`, and `executeUpdate(String, String[])` that return auto-generated keys are not implemented, because of server incapability.
 - `getGeneratedKeys()` is not implemented because server does not support this feature yet. As soon as support for this feature is available, driver will be updated.
- `java.sql.Struct` data type is not supported by server.

The following methods are implemented, but deviate from the specification:

- `java.sql.Statement`
 - `get/setMaxFieldSize` does nothing, Firebird server does not support this feature.
 - `get/setQueryTimeout` does nothing, Firebird server does not support this feature.
- `java.sql.PreparedStatement`
 - `setObject(int index, Object object, int type)` Target SQL type is determined from the class of the passed object and corresponding parameter is ignored.
 - `setObject(int index, Object object, int type, int scale)` Same as above, type and scale are ignored.
- `java.sql.ResultSetMetaData`
 - `isReadOnly()` always returns false
 - `isWritable()` always returns true
 - `isDefinitivelyWritable()` always returns true

Jaybird Specifics

Jaybird driver has also some implementation-specific issues that should be considered during development.

Result sets

Jaybird behaves differently not only when different result set types are used but also the behavior depends whether connection is in auto-commit mode or not.

`ResultSet.TYPE_FORWARD_ONLY` result sets when used in auto-commit mode are completely cached on the client before the execution of the query is finished. This leads to the increased time needed to execute statement, however the result set navigation happens almost instantly. When

auto-commit mode is switched off, only part of the result set specified by the fetch size is cached on the client.

`ResultSet.TYPE_SCROLL_INSENSITIVE` result sets are always cached on the client. The reason is quite simple – Firebird API does not provide scrollable cursor support, navigation is possible only in one direction.

`ResultSet.HOLD_CURSORS_OVER_COMMIT` holdability is supported in Jaybird only for result sets of type `ResultSet.TYPE_SCROLL_INSENSITIVE`. For other result set types driver will throw an exception.

Using `java.sql.ParameterMetaData` with Callable Statements

This interface can be used only to obtain information about the IN parameters. Also it is not allowed to call the `PreparedStatement.getParameterMetaData` method before all of the OUT parameters are registered. Otherwise the corresponding method of `CallableStatement` throws an `SQLException`, because the driver tries to prepare the procedure call with incorrect number of parameters.

Using `ResultSet.getCharacterStream` with BLOB fields

Jaybird JDBC driver always uses connection encoding when converting array of bytes into character stream. The BLOB SUB_TYPE 1 fields allow setting the character encoding for the field. However when the contents of the field is sent to the client, it is not converted according to the character set translation rules in Firebird, but is sent “as is”. When such field is accessed from Java application via Jaybird and character set of the connection does not match the character encoding of the field, conversion errors might happen. Therefore it is recommended to convert such fields in the application using the appropriate encoding.

Heuristic transaction completion support

Current JCA implementation does not support `XAResource.forget(Xid)` method. It might be important in cases when distributed transaction that was some time in-limbo was either committed or rolled back by database administrator. Such transactions appear to Jaybird as successfully completed, however XA specification requires resource manager to “remember” such transaction until the `XAResource.forget(Xid)` is called.

Connection pooling with Jaybird

Connection pooling provides effective way to handle physical database connections. It is believed that establishing new connection to the database takes some noticeable amount of time and in order to speed things up one has to reuse connections as much as possible. While this is true for some software and for old versions of Firebird database engine, establishing connection is hardly noticeable with Firebird 1.0.3 and Firebird 1.5. So why is connection pooling needed?

There are few reasons for this. Each good connection pool provides a possibility to limit number of physical connections established with the database server. This is an effective measure to localize connection leaks. Any application cannot open more physical connections to the database than allowed by connection pool. Good pools also provide some hints where connection leak occurred. Another big advantage of connection pool is that it becomes a central place where connections are obtained, thus simplifying system configuration. However, main advantage of good connection pool comes from the fact that in addition to connection pooling, it can pool also prepared statement. Tests executed using AS3AP benchmark suite show that prepared statement pooling might increase speed of the application by 100% keeping source code clean and understandable.

Usage scenario

When some statement is used more than one time, it makes sense to use prepared statement. It will be compiled by the server only once, but reused many times. It provides significant speedup when some statement is executed in a loop. But what if some prepared statement will be used during lifetime of some object? Should we prepare it in object's constructor and link object lifetime to JDBC connection lifetime or should we prepare statement each time it is needed? All such cases make handling of the prepared statements hard, they pollute application's code with irrelevant details.

Connection and statement pooling remove such details from application's code. How would the code in this case look like? Here's the example

Example 1. Typical JDBC code with statement pooling

```
001  ...
002  Connection connection = dataSource.getConnection();
003  try {
004      PreparedStatement ps = connection.prepareStatement(
005          "SELECT * FROM test_table WHERE id = ?");
006      try {
007          ps.setInt(1, id);
008          ResultSet rs = ps.executeQuery();
009          while (rs.next()) {
010              // do something here
011          }
012      } finally {
013          ps.close();
014      }
015  } finally {
016      connection.close();
017  }
018  ...
```

Lines 001-018 show typical code when prepared statement pooling is used. Application obtains JDBC connection from the data source (instance of `javax.sql.DataSource` interface), prepares some SQL statement as if it is used for the first time, sets parameters, and executes the query. Lines 012 and 015 ensure that statement and connection will be released under any circumstances. Where do we benefit from the statement pooling? Call to prepare a statement in lines 004-005 is intercepted by the pool, which checks if there's a free prepared statement for the specified SQL query. If no such statement is found it prepares a new one. In line 013 prepared statement is not closed, but returned to the pool, where it waits for the next call. Same happens to the connection object that is returned to the pool in line 016.

Connection Pool Classes

Jaybird connection pooling classes belong to `org.firebirdsql.pool.*` package.

Description of some connection pool classes.

<code>AbstractConnectionPool</code>	Base class for all connection pools. Can be used for implementing custom pools, not necessarily for JDBC connections.
<code>BasicAbstractConnectionPool</code>	Subclass of <code>AbstractConnectionPool</code> , implements <code>javax.sql.ConnectionPoolDataSource</code> interface. Also provides some basic properties (minimum and maximum number of connections, blocking and idle timeout, etc) and code to handle JNDI-related issues.
<code>DriverConnectionPoolDataSource</code>	Implementation of <code>javax.sql.ConnectionPoolDataSource</code> for arbitrary

Description of some connection pool classes.	
	JDBC drivers, uses <code>java.sql.DriverManager</code> to obtain connections, can be used as JNDI object factory.
<code>FBConnectionPoolDataSource</code>	Jaybird specific implementation of <code>javax.sql.ConnectionPoolDataSource</code> and <code>javax.sql.XADataSource</code> interfaces, can be used as JNDI object factory.
<code>FBSimpleDataSource</code>	Implementation of <code>javax.sql.DataSource</code> interface, no connection and statement pooling is available, connections are physically opened in <code>getConnection()</code> method and physically closed in their <code>close()</code> method.
<code>FBWrappingDataSource</code>	Implementation of <code>javax.sql.DataSource</code> interface that uses <code>FBConnectionPoolDataSource</code> to allocate connections. This class defines some additional properties that affect allocated connections. Can be used as JNDI object factory.
<code>SimpleDataSource</code>	Implementation of <code>javax.sql.DataSource</code> interface that uses <code>javax.sql.ConnectionPoolDataSource</code> to allocate physical connections.

org.firebirdsql.pool.FBConnectionPoolDataSource

This class is a corner stone of connection and statement pooling in Jaybird. It can be instantiated within the application as well as it can be made accessible to other applications via JNDI. Class implements both `java.io.Serializable` and `javax.naming.Referenceable` interfaces, which allows using it in a wide range of web and application servers.

Class implements both `javax.sql.ConnectionPoolDataSource` and `javax.sql.XADataSource` interfaces. Pooled connections returned by this class implement `javax.sql.PooledConnection` and `javax.sql.XAConnection` interfaces and can participate in distributed JTA transactions.

Class provides following configuration properties:

Standard JDBC Properties

This group contains properties defined in the JDBC specification and should be standard to all connection pools.

Property	Getter	Setter	Description
<code>maxIdleTime</code>	+	+	Maximum time in milliseconds after which idle connection in the pool is closed.
<code>maxPoolSize</code>	+	+	Maximum number of open physical connections.
<code>minPoolSize</code>	+	+	Minimum number of open physical connections. If value is greater than 0, corresponding number of connections will be opened when first connection is obtained.
<code>maxStatements</code>	+	+	Maximum size of prepared statement pool. If 0, statement pooling is

Property	Getter	Setter	Description
			switched off. When application requests more statements than can be kept in pool, Jaybird will allow creating that statements, however closing them would not return them back to the pool, but rather immediately release the resources.

Pool Properties

This group of properties are specific to the Jaybird implementation of the connection pooling classes.

Property	Getter	Setter	Description
<code>blockingTimeout</code>	+	+	Maximum time in milliseconds during which application can be blocked waiting for a connection from the pool. If no free connection can be obtained, exception is thrown.
<code>retryInterval</code>	+	+	Period in which pool will try to obtain new connection while blocking the application.
<code>pooling</code>	+	+	Allows to switch connection pooling off.
<code>statementPooling</code>	+	+	Allows to switch statement pooling off.
<code>pingStatement</code>	+	+	Statement that will be used to “ping” JDBC connection, in other words, to check if it is still alive. This statement must always succeed.
<code>pingInterval</code>	+	+	Time during which connection is believed to be valid in any case. Pool “pings” connection before giving it to the application only if more than specified amount of time passed since last “ping”.

Runtime Pool Properties

This group contains read-only properties that provide information about the state of the pool.

Property	Getter	Setter	Description
<code>freeSize</code>	+	-	Tells how many free connections are in the pool. Value is between 0 and <code>totalSize</code> .
<code>workingSize</code>	+	-	Tells how many connections were taken from the pool and are currently used in the application.

Property	Getter	Setter	Description
<code>totalSize</code>	+	-	Total size of open connection. At the pool creation – 0, after obtaining first connection – between <code>minPoolSize</code> and <code>maxPoolSize</code> .
<code>connectionCount</code>	+	-	<i>Deprecated.</i> Same as <code>freeSize</code> .

org.firebirdsql.pool.FBWrappingDataSource

This class is a wrapper for `FBConnectionPoolDataSource` converting interface from `javax.sql.ConnectionPoolDataSource` to `javax.sql.DataSource`. It defines same properties as `FBConnectionPoolDataSource` does.

Runtime object allocation and deallocation hints

Pool implementation shipped with Jaybird can provide hints for the application where the connection was obtained from the pool, when it was released back to the pool, when the statement was prepared. Such information is written into the log when appropriate system properties are set to `true`. Additionally, when connection or prepared statement is closed twice, driver will throw an SQL exception with an attached stack trace of previous call to `close()` method.

List of properties

Property name	Description
<code>FBLog4j</code>	<p>Enables logging inside driver. This is the essential property, if it is not present or set to <code>false</code>, no debug information is available.</p> <p>When it is set to <code>true</code>, pool automatically prints the following information:</p> <ul style="list-style-type: none"> • When physical connection is added to the pool – <code>DEBUG</code> • When a maximum pool capacity is reached – <code>DEBUG</code> • When connection is obtained from pool – <code>DEBUG</code> • When connection is released back to pool – <code>DEBUG</code> • Whether pool supports open statements across transaction boundaries – <code>INFO</code>
<code>FBPoolShowTrace</code>	<p>Enables logging of the thread stack trace when debugging is enabled and:</p> <ul style="list-style-type: none"> • Connection is allocated from the pool – <code>DEBUG</code> • Thread is blocked while waiting for a free connection – <code>WARN</code>
<code>FBPoolDebugStmtCache</code>	<p>When statement caching is used and debugging is enabled, following information is logged:</p> <ul style="list-style-type: none"> • When a statement is prepared – <code>INFO</code> • When statement cache is cleaned – <code>INFO</code> • When statement is obtained from or returned back to pool – <code>INFO</code>

Documentation and Support

Where to get more information on Jaybird

The most detailed information can be found in the Jaybird Frequently Asked Questions (FAQ). The FAQ is included in the distribution, and is available on-line in several places.

Also a new resource, JaybirdWiki is available at <http://Jaybirdwiki.firebirdsql.org>.

Where to get help

The best place to start is the FAQ. Many details for using Jaybird with various programs are located there. Below are some links to useful web sites.

- The <http://groups.yahoo.com/group/Firebird-Java> and corresponding mailing list Firebird-Java@yahoogroups.com.
- The code for Firebird and this driver are on <http://sourceforge.net/projects/firebird>.
- The Firebird project home page <http://www.firebirdsql.com>.

Reporting Bugs

The developers follow the Firebird-Java@yahoogroups.com list. Join the list and post information about suspected bugs. This is a good idea because what is often thought to be a bug turns out to be something else. List members may be able to help out and get you going again, whereas bug fixes might take awhile.

If you are sure that this is a bug you may report it in the Firebird bug tracker, “Java Client (Jaybird)” at SourceForge.net project area (<http://sourceforge.net/projects/firebird>).

Corrections/Additions To Release Notes

Please send corrections, suggestions, or additions to these Release Notes to the mailing list at Firebird-Java@yahoogroups.com.