2023/06/14 21:12

IBExpert Replication FAQs

This article is also available as a

pdf

What is replication?

Wikipedia: replication (from the Latin replicare meaning "return", "repeat") refers to the storage of data on multiple storage devices, usually to several different locations, and the synchronization of these data sources.

Which Firebird versions can be used for IBExpert replication?

We recommend using the current Firebird 2.5.x or 3.0.x version for synchronous IBExpert replication. Asynchronous IBExpert replication can also be realized with older Firebird versions, but here the data exchange runs via external event-driven scripts, whose runtime behavior may possibly lead to a loss of committed transactions.

We can also create a replication between Firebird 2.5 and 3.0 databases, so you can already work in a multi-master setup with clients who use the Firebird 3.0 advantages like advanced multithreaded support with a second database using the proven Firebird 2.5 version. All operations are working under full transaction control on all replication nodes, so any commit or rollback works globally on all nodes. When using the asynchronous mode, we also implement an unlimited global audit for all records in the database, since all operations are in the transaction protocol with username, timestamp and IP address.

What advantages does IBExpert replication offer compared to a regular backup/restore?

A backup/restore cycle allows you to backup data during runtime, but all transactions that were performed after the backup process was started are lost after the restore. The larger the database, the longer the backup/restore cycle lasts, inevitably leading to a greater potential loss of data. When using IBExpert replication in synchronous mode, Firebird transactions are immediately written to the drive so that in the case of the master suffering a hardware failure, no committed transaction is lost.

What advantages does IBExpert replication offer compared to a shadow?

A Firebird shadow is a constantly updated copy of the database file to a second drive. Depending on the operating system, the shadow can only be located on the same server as the original database file. This means that in the case of server failure, in most cases the shadow is also no longer accessible. Furthermore, a shadow must – even following a brief downtime - be rebuilt from scratch, a process which for large databases can take several minutes.

Does IBExpert replication also work with InterBase®?

Yes, but only asynchronously and with some further restrictions, as we have ceased any further development on this platform, due to the dwindling distribution of InterBase®. We recommend switching to Firebird if demand for high performance and reliability is paramount, and we are more than happy to help you with this.

back to top of page

Does IBExpert replication also work with other database systems?

We already have experience with Oracle, MS SQL and DB2 slave connections. The Firebird database content is transmitted via ODBC into these databases and then synchronized. In general, this replication can be implemented with any ODBC-compliant database to synchronize, for example, the data in a MySQL-based web shop.

The leading system here is always a Firebird database. Writing changes from the other databases can also be imported in batch mode. The external system should have data fields via which the modified records may be transported.

Which operating systems are supported by IBExpert replication?

Generally all operating systems on which Firebird 2.5.x can run are perfectly suitable for IBExpert replication, e.g. Win32, Win64, Mac OSX, Linux etc. IBExpert replication utilizes exclusively Firebird technology for the replication.

Which data can be replicated?

IBExpert replication can replicate all data that can be stored in a Firebird database. We recommend SQL dialect 3, but can, in individual cases, replicate databases based on dialect 1.

Can blobs be replicated?

Yes, even blobs can be replicated fully.

back to top of page

Can metadata be replicated?

Yes, metadata changes can also be replicated, such as CREATE TABLE commands. This capability is particularly important with 24/7 operation and also when using asynchronous replication with laptops.

Which replication methods does IBExpert replication support?

a. Master-slave synchronization: in the case of synchronous master-slave replication, any transaction on the master is immediately transferred to the slave. Write operations can only be performed on the master, read operations on the master and on the slave. If the master fails, the slave can immediately take over the task of the master and also execute write operations. As long as the master can be accessed by the slave, it refuses these write connections. Should the master discover after a failure that the slave has taken over its duties, the former master refuses all connections. Typical application is the scalability for read operations and enhanced reliability.

b. Master multi-slave synchronization: the same procedure as for a., but with multiple read-only slaves. Typical application is the virtually unlimited scalability for read operations.

c. Master-slave asynchronous: with asynchronous master-slave replication, any transaction is initially cached in the master database. Only after completion of the transaction is an attempt made to transfer this in a separate transaction to the slave. If the slave is not available, then the data is initially collected locally in order to then transmit when the connection is resumed. Typical applications are reading chain store applications, where intermittent or slow data lines between master and slave make synchronous transmission impossible.

d. Master multi-slave asynchronous: same procedure as for c, but with multiple read-only slaves. Typical application is the use of laptops that are intermittently connected at unpredictable times. Here the data exchange is initiated by the slave using the pull method.

e. Master-master synchronous: with this method IBExpert replication treats each participating node both as master and as slave. Write operations can be performed without restrictions on each master. To implement this configuration, adjustments may be necessary in the data model because the distributed creation of primary keys and foreign keys must be possible. Typical application is the scalability for read and write operations, and enhanced reliability.

f. Master-master asynchronous: here, the same procedure applies as in e, but each transaction is first cached locally and only then transmitted to the other server. Typical applications are read and write branch applications at two locations, where only intermittent or slow data lines make synchronous transmission impossible.

g. Multi-master synchronous: See e, but with more than 2 masters. Typical application is the almost unlimited scalability for read and write operations.

h. Multi-master asynchronous: See f, but with more than 2 masters. Typical applications are read and write branch applications at more than two sites, where only intermittent or slow data lines make synchronous transmission impossible.

back to top of page

Do the servers have to be constantly connected?

No, not unless this security feature is desired. In synchronous mode, the system can be configured so that if a server fails, it initially remains in the cluster in asynchronous mode and when it becomes available again, is re-synchronized by the systems that are still ongoing, before it can answer client queries again.

How can you recognize the current replication progress?

A query by a stored procedure shows the status of the nodes involved. The result can also be used in case of failure by an external script, to send an email to the administrator, should his intervention be required.

What should be taken into consideration when performing a multi-master replication?

A multi-master replication in synchronous mode is relatively uncritical in use, as any replication conflicts are immediately detected by the transaction security, and competing updates also cause an exception, just as if the operation were being performed in the same database. In asynchronous mode it is necessary to ensure that competing updates are avoided in the data model as far as possible.

Which method is best for laptop replication?

We have positive experiences with several projects using the multi-master asynchronous mode. Following customization of the data model, this mode is used successfully by thousands of systems.

How does replication affect database performance?

Write operations access triggers used by IBExpert replication, which of course affects database performance. With the appropriate hardware, primarily characterized by very fast I/O subsystems (ideally enterprise SSDs), this impact is minimal. We recommend a dedicated gigabit crossover cable network connection, each with their own network cards between the nodes. We advise explicitly against virtualization, as this brings in many factors which play a role that we cannot influence. Virtual systems can rarely achieve the high demands on the required I/O performance. Not without reason do - not just ourselves but - almost all major database vendors of high-performance database systems advise against virtualization.

back to top of page

What is the optimal hardware configuration?

We generally recommend dedicated systems with slimmed-down 64-bit operating systems, in which all the functions that are not used explicitly for the database server are shut down. For Firebird database systems, a NAS (Network Attached Storage) system generally has a negative effect, as Firebird writes primarily small blocks, and for this the latency of the NAS system connection results in considerable performance losses. Generally XEON E3 or E5 CPUs offer very good value for money when used with Firebird. Generally power-saving versions should be avoided, and even too many CPU cores are not always beneficial.

What is the optimal configuration to achieve maximum availability and scalability?

Following the model of Facebook and Google data centers, we recommend preferably several simple replicated systems, rather than replacing these with a high-end server that also usually requires exotic parts.

How secure is replication?

In synchronous mode extremely secure, as all transactions are automatically replicated to all nodes. If one of the nodes raises an exception, then this also applies automatically to the calling client.

How do you prevent multiple trigger-firing?

Replication is transmitted by a specific user. All operations that are performed by this user should be ignored in the trigger source code. This can be automatically solved by simple additions in the first line of any existing triggers.

What happens when replicational conflicts occur?

In synchronous mode this leads, without any restrictions, to the same exception, which occurs even when both competing operations are performed on a single database.

What happens when a replication member is unavailable?

Normally, this partner is then served further in asynchronous mode. If desired, each new transaction can be banned in exceptional cases or until reconnection. But this is rarely required.

back to top of page

Does all data have to be replicated?

No. The IBExpert replication creates replication objects using a stored procedure. This is available in the source code and can exclude any exceptions from the replication.

Can generators be replicated?

No, because generators are not transaction-consistent. But we have several methods to synchronize generators for asynchronous operation, to ensure sufficient security against primary key conflicts. In synchronous mode, if necessary, the generators from a leading master can be given priority.

Can you work in a session with inactive replication?

Use of a special replication user allows commands to also be executed without replication. But this should be used with caution.

Which commands need special treatment?

For bulk operations such as DELETE FROM TABLE, IBExpert replication would normally process each record separately. For such cases, such an instruction can also be transmitted as a command, using a special procedure.

What do I need to do after making alterations to metadata?

If you have followed all the rules (e.g. no table without a primary key), following any metadata changes you only need to call a procedure, which creates all the necessary replication objects. To change the metadata on all systems simultaneously, there are special procedures with which you can run the SQL parallel to all active databases in the cluster. Alternatively, the IBExpert IDE also offers special functions which perform this automatically even in the case of interactive changes.

back to top of page

How can I order IBExpert replication?

You only need a simple master-slave configuration for database use at your business location? Or would you like to offer IBExpert replication to your customers, affiliates or employees? If so, then please contact us for a written quotation for distribution software: sales@ibexpert.biz or +49 4407 3148770.

How does IBExpert replication compare to other Firebird replication solutions?

Our system is based on 15 years of development in enterprise environments, where other solutions could not be used for various technical reasons and conceptual weaknesses. We guarantee low maintenance, stability and high performance if you abide by the rules imparted in our workshop.

This article is also available as a

pdf





IBExpert - http://ibexpert.com/docu/