



Firebird White Paper

Big Data, Replication or "Where are the limits of Firebird?"

Holger Klemt, March 2016

Over the last two years we have implemented a client project, which continues to grow steadily. This project illustrates the amount of data that is possible with a Firebird database on a suitably-configured system environment.

EPOS systems for gastronomy: Multi-master replication

A large gastronomic enterprise and their partner responsible for POS systems approached us in 2014 at the initiative of the software vendor. Their administration software, written in Delphi, was deployed as a self-contained solution at each site. Data changes always occurred only at the local sites, and there were interfaces for some data, but the whole system was not really satisfactory, especially taking into consideration the rate of growth of the corporate group.

The software producer was already deploying Firebird in the decentralized solution. Various techniques and guidelines were discussed and determined in a series of meetings. It was decided to place the new project on a multi-master replication, so that transparent replicated data could be used at all sites in online/offline operation.

With over 100 sites and not always totally reliable Internet connections, a pure cloud solution was not an option. If a manager cannot see the staff rota, because the Internet is not accessible or the central cloud solution for all locations fails for whatever reason, this has irrevocable negative effects on the quality of service in-house.

Due to the growth it is all the more important to be able to schedule employees from neighbouring sites as a temporary help, and be in a position to correctly record their working hours.

We proposed a decentralized and failsafe online/offline cloud. Each site receives a so-called black box as Firebird server in its own local network, which is then used as a database server by the Windows/Delphi program on the workstation which was already in use.

When an Internet connection is available, the box, which runs an Ubuntu Server Linux operating system, establishes a tunnel via SSH to the communication server in the data center. All data material from all sites is collected in two Firebird database servers in the data center, and distributed to all relevant sites using our replication technology implemented according to specific rules.



All sales-relevant transaction data is also transferred from the POS system to these servers, after these have been transferred to the black box by the respective POS system software at the site using the Firebird ODBC driver.

Time-critical replication

So that the corporate management can view all the relevant data as promptly as possible, even though the transaction data is not available in its entirety until early in the morning following the so-called Z-report, there is very little time to transfer all data to the data center. Currently around 1.2 million records from various locations are transmitted daily to the data center in this way. The time window available is 4:30 to 8:00 am, which is only 3.5 hours, or the equivalent of 12,600 seconds.

Therefore a constant rate of around 100 records a second has to be achieved by the upload during this period. In reality however, the uploads are completed in less than an hour. The ODBC upload to the black box performed by the relatively slow POS software is probably currently responsible for the lower upload rate. In simulations our environment manages more than 1000 records per second.

For a number of reasons the transaction data is then distributed in the data center to 3 other servers, and also to 2 servers running at respective administrative sites on their own servers. This means about 6 million records need to be distributed to other databases every morning.

In addition to the transaction data, other information such as working hours, employee information, etc. are stored in the database. This data is replicated almost at real time across all sites around the clock in order to create records, for example, for new staff at the headquarters, who can then record their working hours at their local site. Unlike the transaction data, the master and dynamic data is automatically replicated to all other locations. Even changes of larger data volumes at the end of each month are thus distributed in a matter of seconds to all sites.

Of course, absolute consistency of all database transactions is crucial. When data is sent from one location, its packets must be packaged transaction-safely. In the case of an eventual connection interruption it must be ensured that only the missing information be retransmitted, otherwise the upload capacity of slow DSL lines can quickly be overloaded.

Remote maintenance

Our black boxes are fully controllable via remote maintenance, and the operation and connection on site are both straightforward. The hardware supplied by ourselves has only a power switch and connections for the network cable and the power supply. Network port clearances on the router and similar requirements were not requested and are not necessary due to the SSH technology used. As soon as the black box is connected to a DHCP-enabled Internet router, the SSH tunnel is active, complete with encryption and compression, and the box can be reached by the data center.

Should the black box become offline and not automatically return online as is usual when the Internet connection is reactivated, then the only intervention required by site staff is to disconnect the power cable and reconnect it. How long which black box was or is offline, can be viewed by the customer on a central website and kept track of via email.



Black boxes, which cannot be started up again by simply rebooting, are reported by the customer, sent to us and exchanged with replacement boxes. In accordance with the customer specifications, we ensure that we have a sufficient number of replacement devices, stored by ourselves or at their own locations. So a replacement box is generally available on the next working day following notification.

Metadata updates

The whole system is not a static solution, as the software manufacturer is constantly adapting his Delphi program to meet the customer's requirements. The database structure is therefore highly dynamic and such alterations also have to be transferred to all servers, even when these, for example, are offline for a number of days following an internet failure. In collaboration with the software manufacturer, we use various IBExpert technologies to perform the database synchronization. Other subprojects based on PHP or .NET are also connected directly to the Firebird database.

Logs provide clarity

The key point for all parties involved in the project is the high degree of transparency of our solution. Strict compliance to all rules that we have jointly determined (as already mentioned), is essential for the software development. No further specifications are necessary.

All data manipulation has so far indefinitely been stored in the log database, so that all data alterations can be traced at all times.

Preconfigured hardware

Each site receives one black box completely preconfigured by ourselves, as a local Firebird server for a budget of less than 1,000 €. (Plus additional servers in the data center depending on customer specifications.) For 50 sites a budget of about 60,000 € should be planned.

Our black boxes can also be supplied as a mobile version complete with 12/24 volt connection and UMTS/LTE connection. This allows large quantities of data, such as image and video recordings of IP cameras recorded mobile for the long term and replicated as a thumbnail with the head office to minimize the requisite traffic. If necessary, the recorded data can be retrieved in full quality without any limitations. We can also supply suitable IP cameras.

The complete operation is carried out by ourselves for an annual maintenance fee of 15% of the total project cost, so that neither the end user nor the software manufacturer needs to worry about the technical details.

The hardware supplied by ourselves is available for European customers, who may then put the black boxes into operation outside Europe at their own risk if they wish.

For customers outside the EU, we can supply a disk image for download, made available for our predefined hardware, which can be put into operation on site.

The communication servers are located in Germany at our data center partners and are operated under our supervision.



IBEXPERT KG WHITE PAPER



Key figures concerning data volumes and hardware

- Daily Firebird transactions on the central communication server: 3 million
- Size of the database on the communication server: 45 GB
- Size of the log databases on the communication servers: 250 GB
- Size of the databases on the black boxes: 10-15 GB
- Largest table: 137 million records
- IBExpert Firebird Benchmark communications servers: Drive index 230% CPU index 150%
- Hardware used for the communication servers: IFS Server 6.0 with options
- IBExpert Firebird Benchmark black box: Drive index 50%, CPU index 55 %
- System stability in the data center: the central database is replicated live on 3 systems

Are you interested in this or a similar solution?

Please contact sales@ibexpert.com.