The way to the Cloud and Firebird

A White Paper by Holger Klemt

During a conversation with a software vendor, whose largest customer complained about the performance of their Firebird database, especially with larger databases, we conducted an IBExpert benchmark on the customer's hardware as part of our IBExpert and Firebird hotline service.

The server used is a high-end HP model, with a number of built-in SSDs, the price range was estimated by the customer at about 35,000 Euros 3-4 years ago.

The customer's motivation to commission us with this analysis, was to make the best possible use of the current hardware, until the planned conversion to cloud-based servers had been completed.

The reference server is our IFS hardware, where the benchmark is completed in about 20-30 seconds.

As we had some waiting time here, there was time for explanations of the physical basics.

Benchmark status after 5 minutes:

The benchmark has not finished yet, but the computer is still working ...

So let's use the time to identify the hardware we are using.

The CPU could be identified as an Intel Xeon E5 CPU, which although it has more than 10 cores, unfortunately only operates them with 2.0 GHz. Modern Xeon E3s operate at significantly more than 4 GHz per core.

Let's take a look at the simple rule of three: multiply the number of cores/threads by the number of GHz and the theoretical total performance of the CPU. If, however, there is only one thread that gets the maximum performance, this figure won't help you.

So we multiply 1 as the number of cores by the maximum GHz figure and arrive at the result that a 4 GHz machine can perform most operations within half the time of a 2 GHz machine.

The customer is using Firebird 2.5 Superserver version, as the software manufacturer has not yet upgraded to Firebird 3.0. The Superserver version is well suited to relieve the I/O subsystem, i.e. to achieve a high speed with a slow HDD/SSD.

The Classic or Superclassic server uses multiple CPU cores only when multiple clients are requesting services concurrently. If only one client requests services, the Classic architecture has neither advantages nor disadvantages. If you compare the Super and Classic versions with the I/O load generated by the IBExpert benchmark, you will see that the benchmark generates 13 GB of read and write operations on the Superserver, and 43 GB of read and write operations on the Classic and Superclassic versions. So if a theoretical read/write performance of the HDD allows 100 MB per second, we get a minimum of 130 seconds for the Superserver and 430 seconds for the Classic versions.

The number of CPU cores is therefore rarely helpful when combined with rotating hard disks. If your software does not generate any CPU load with the Firebird process, it is because this process has to wait for the I/O subsystem. Or in other words, the HDD/SSD combination used is too slow and therefore unsuitable for use by a Firebird server.

Benchmark status after 10 minutes:

The benchmark still doesn't display any partial results, but there are still read/write operations visible in the Task Manager when the columns are explicitly displayed, so we'll keep waiting ...

The avid administrator would probably now like to mention that the HDDs or SSDs were configured as RAID, as that would offer the optimal combination of security and speed.

When using suitable hardware, there is nothing wrong with a RAID. But unfortunately the effects of many RAID controllers are not clear to some administrators. While m.2 SSDs with 4 PCI channels, connected on the mainboard or via PCI, can write data back and forth directly between memory and SSD memory with very simple commands from the chipset without generating a large CPU load, the probability is very high that the RAID controller has its own CPU and this often has the performance class of a Celeron or Atom processor.

If the clients process data on 12 CPU cores with 2 GHz each and would also like to read and write data concurrently, they must query the CPU running at a maximum of approx. 1 GHz, which then attempts to process the queries with 1-2 cores. When it comes to 1 GB-sized parts of the database file, the RAID CPU can hardly do anything; however with a lot of small read and write operations the situation is quite different.

Benchmark status after 13 minutes:

The first interim result for the Drive Index has now been entered in the IBExpert benchmark. It took the server about 740 seconds to write this value.

As we can now estimate that it will take about 15 minutes to get the overall result, we will use the time to explain what the benchmark in the Drive Index represents.

The benchmark database is generated from the script in the IBExpert Benchmark subdirectory. After the database is finished, the cache buffers with the drive index are set to the smallest possible value of 50 in order to use the disk as intensively as possible. Whenever Firebird no longer has any free cache buffers, the data must be swapped from memory back to disk and reloaded as necessary.

For the drive index, this process is induced approx. 700,000 times. With rotating hard disks these values used to be technically estimated based on the average access time, because a disk with 5ms creates approx. 200 I/O operations per second. Modern m.2 SSDs, on the other hand, manage more than 250,000 I/O operations per second, abbreviated as IOPS. So if you're using SSDs with significantly lower IOPS limits, you should reconsider your hardware selection.

The database for the benchmark is fully recreated for the CPU test, however is now set with 5,000 cache buffers, so that instead of 700,000 read-write operations, only about 1,400 read-write operations are required. In this test, the interaction between CPU, motherboard and RAM is much more important.

With both methods, the time measurements are not performed at the start of the script, but only when the respective database is ready with the corresponding cache buffer settings. As very few commands via SQL over the network execute only short stored procedure calls, the result is fairly identical whether running locally on the server or from another machine in the network. It also depends on the speed of the server running Firebird.

Benchmark finishes after 28 minutes:

The benchmark is ready!

Now the modest final result can be discussed in more detail. The benchmark has calculated 2% in the drive index, 14% in the CPU index and 6% in the thread test.

A test of the same hardware with a RAM disk produces comparatively modest values, but who would seriously want to run his database on a RAM disk?

In comparison: Our IFS servers in the 4,500 Euro price range deliver values of approx. 500%/250%/ 300%. All times measured on our IFS servers are up to 200 times faster!

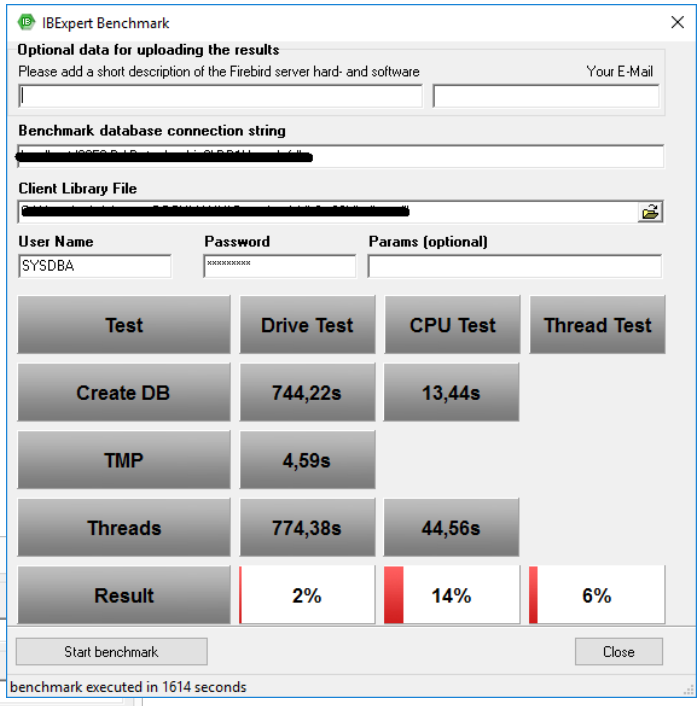

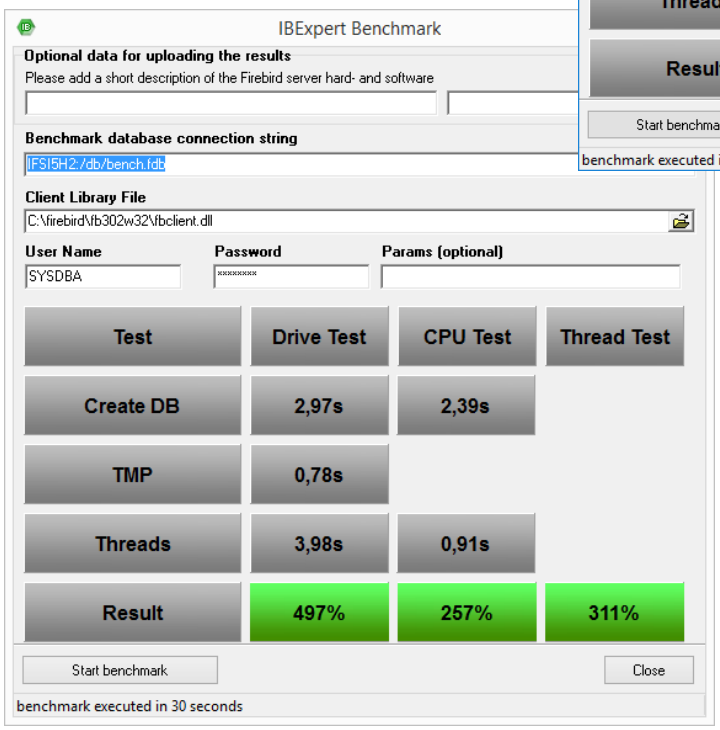

Even with optimal utilization of the customer hardware using just one thread and with sufficient storage space, our IFS server requires only 2.4 instead of 13.5 seconds for the drive test.

Simple processes, such as the nightly backup, take 6 hours on the customer server, whereas on our IFS server it takes just one hour. This advantage is of course even more noticeable with the much more complex database restore. But as long as the backup is finished overnight, many admins are more or less satisfied and should an error occur everything is doubly protected via virtualization. But do you really want to rely on that?

On the one hand do not believe that a virtualization is fail-safe in all circumstances, and that all data can be recovered should an error occur, and, on the other hand and even more important, an error in a file is not always immediately detectable.

A practical example shows that two stored procedures with the same name existed in a database, even though this is technically not normally possible. However the nightly backup didn't show any errors and so it was assumed they were safe.

That the nightly backups were worthless without an error-free restore was however only discovered by our customer when he needed the restore, and realized that all backups of the last 9 months were defective. We were able to solve the problem with direct manipulation of the backup files using a hex editor, however, a complete working day of a company with over 200 employees was practically lost.

This worst-case scenario was a very clear indication that the security concept needed to be redefined.

... Back to the benchmark and a potential future in the cloud ...

Proposals for possible optimizations of the existing system and some immediate measures, such as the adjustment of the cache buffers, affinity, priority, and our recommendation to switch to Firebird 3 as quickly as possible, were immediately adopted and our conversation then returned to the planned conversion to cloud-based servers in order to be able to replace the poor performance by purchasing better and arbitrarily scalable hardware etc.

Since we already have similar, very poor benchmark results for Amazon EC2 cloud systems, we clearly pointed out that the Amazon machines may be useful for many purposes, but before purchasing them, the small print needs to be studied very carefully! Look closely at what Amazon has to offer in the surcharge list for example, so that you don't end up having to spend your entire budget for a new server on Amazon on a quarterly basis to merely attain a modest performance improvement.

The following PDF file (approx. 1,041 pages) is suitable reading; the drive types are described starting on page 803:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-ug.pdf#EBSVolumeTypes

So preferably no cloud? Not at all!

We support dozens of customer fast Firebird servers in data centers. We work together with Hetzner and now have several racks including customer hardware for Firebird. We handle the administration of the Firebird system. Our customers can rent terminal servers for other tasks in the same data center or use their own hardware.

Furthermore, we can also recommend virtual server packages for customers' own projects, which provide good performance for Firebird or Firebird-based applications.

Did you know that a virtual server is available for just 10 Euro per month, including Windows 2012R2 license and 100 GB SSD storage?

With the IBExpert benchmark this server achieves 40-50% performance in all 3 values!

We have been implementing hybrid cloud solutions, in which data is replicated between the cloud and local database servers, for many years. Even millions of new transactions and billions of data records are part of our daily business.

Conclusion

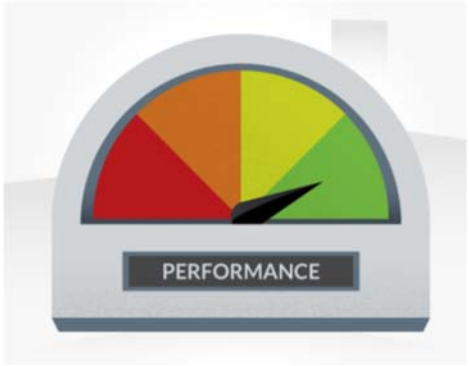

A Firebird server is very demanding as a database server in the I/O performance of the server system used. Just because a certain server XY can do other things quite well does not necessarily make it suitable for Firebird databases. The hardware I/O performance best suited to Firebird is much cheaper than you think. And if configured accordingly, it can be used for much longer than 5 years.

Anyone who knows me in practice, from various lectures or training sessions, or who reads my white papers, knows that I always tend to have a "real life" example at hand:

If a baker should use a 40 ton articulated lorry to try to supply each household in a residential area with a bag of fresh rolls each morning, he would certainly question his choice of vehicle, regardless of whether the lorry was also purchased for other purposes or whether it spends most of the day parked in the yard.

To hire a freight company to supply your customers using such vehicles is obviously not a realistic solution. The baker can also save himself an inquiry to a shipping company that would use inland waterway vessels for this purpose. If the baker were to meet an honest shipowner or freight forwarder, they would reject his enquiry and forego the deal.

However, experience has shown that cloud providers give their potential customers the impression that they can do everything better, faster and more securely at a fraction of the cost. In comparison, the cloud providers can offer you everything from a bicycle messenger to a lunar module, however without the necessary performance and certainly not the necessary experience. Switching from one provider to another is extremely rare and these dependencies are exploited by cloud providers once you get hooked.

Or as someone at the CeBIT trade fair once said so aptly: "If lies really had short legs *[a German adage, similar to “lies will soon catch up with you”]*, then most people would walk around on their buns".

The cloud does not mean loss of control, but together with an independent, competent partner as consultant can be a good addition to the existing infrastructure, or even a replacement. **Stay independent!**

Holger Klemt, January 2019