



Firebird White Paper

The Pros and Cons of virtualising SQL and other databases

Part 1

Source: <http://blog.krollontrack.co.uk/concepts-explained/the-pros-and-cons-of-virtualising-sql-and-other-databases-part-1/>

Advantages

The same as with the current virtualisation of desktop applications, servers or storage, the operation of databases on virtual machines also offer some undeniable advantages. In addition to the optimal utilisation of the hardware employed, the resulting savings in electricity, avoidance or limitation of purchases of new equipment and a simplification in the management of the database solution, virtualisation provides still more advantages. Particularly in the case of very resource-hungry database systems, such as data mining for CRM, ERP or BI. Further advantages include:

- Virtual databases can be 'live migrated' from one physical server to another without service interruption in case of hardware issues
- Cost-effective and easy implementation of high-availability solutions
- Flexible, dynamic and automated deployment of new system instances and resources when needed (scalability)
- The possibility of agile database development, such as the use of different virtual machines with different database systems or versions promotes the development or testing of agile software development within the try-and-error principle. Systems with different degrees of updating can be adjusted, changed or deleted, without risking to compromise "functional" databases.
- Improved availability by separating virtual machines from each other, in case of issues with one VM, the rest of the VMs can continue operating without any sacrifices in performance.

Drawbacks

It is no wonder that the virtualisation of databases is increasing. But despite these advantages, it can also lead to problems should their implementation be carried out too fast and without sufficient planning.



Because there are several things to consider when implementing them. Problems can occur especially in case of:

- Databases generally require a lot of resources, whether in a real or a virtualised system. Virtualised database systems based on Microsoft SQL Server and similar, need powerful processors, lots of memory and above all great storage capacity, so that any data can be quickly processed by the system, just like real databases. Without sufficient hardware, the database will most likely suffer a significant degradation in performance.
- In some cases, such as old Oracle databases, previous database licenses cannot be transferred 1:1 to a virtualised system, as the fees are related to the (potential) performance of the system and not to what is actually used. It is therefore important before starting a transition to consider first how many instances and processors are to be actually used, in order to get a comparison between the cost of a physically existing database server and its virtual counterparts.
- Databases are by nature complex, and virtualisation changes nothing in this, which is why it is critical to have skilled and knowledgeable staff to operate them. And with virtualisation, database administrators (DBAs) have to face a new technology level which further adds to the complexity. Thus, if no distinction is made by the company between IT admins who are responsible for virtualisation, and DBAs, then it falls to one type of employee to learn two complex disciplines.
- Lack of exchanges or cooperation between IT Admins and DBAs often causes problems as many database administrators have no real access to the depths of the virtualisation layer, as this is managed by IT administrators. In case of problems with the database caused by an anomaly in the VM or virtual system, there are often long delays in resolving the issue.

This last point is especially something that can be improved, as we have seen in our data recovery experience. As with normal failures of virtual systems, the reasons for the disappearance of virtualised databases are mainly:

- reformatted VMware Data store volumes,
- damaged VMFS data store volumes,
- damaged guest file systems,
- corrupt virtual files (VMDK/VHD) or
- accidentally deleted file systems (VMDK or VHD).

Thus it is not exclusively hardware failures or defects that lead to the failure or disappearance of virtual data or databases, but in many cases this is due to human error. And very often, the responsible DBA appears to be simply overwhelmed when something unexpected happens.

But what should you decide now: virtualise your databases or not? In the second part of this article we provide the answer to this question and also important clues on what to watch out for when considering database virtualisation.



Part 2

Source: <http://blog.krollontrack.co.uk/concepts-explained/the-pros-and-cons-of-virtualisation-sql-or-other-databases-part-2/>

In the first part of our article on database virtualisation, we concerned ourselves with its pros and cons. In this second part, we are going to deal with the question of whether this technology should be implemented and what to watch out for. But first of all: what can actually happen if a virtual database disappears? The following, very real case from our data recovery lab, which clearly demonstrates such a situation:

When something actually goes wrong: a bank without a database!

Database virtualisation, like anything, doesn't always run smoothly and considering its complex nature, there might be difficult challenges in data restoration in the event of a failure. We experienced this when an internationally recognised bank came to us for help.

The bank's full customer and transaction virtual database was lost. The reason? After maintenance was performed, the VMware ESX server with three LUNs (Logical Unit Numbers) refused to work at all and could no longer be booted. The emergency cluster system also failed to work due to the replication link having not been previously disconnected. The database recovery turned out to be much more difficult than anyone had imagined at first. The server's VMFS file system was heavily damaged and had to be reconstructed in several individual steps. Only after this was completed could the affected SQL databases be copied and thus a functioning new customer and transaction database be created.

So, what now? Database virtualisation: yes or no?

It is a fallacy to believe that you can cope with the growing volume of data only by virtualising databases. This is absolutely not true. Therefore, you shouldn't virtualise any databases that are already working at almost full capacity in the real world. Before virtualising, you should first analyse the actual daily load behaviour and based on this, the required hardware resources. Only then can you really make sure that server consolidation doesn't lead to a dramatic drop in performance because you are saving too much in hardware.

In addition, deciding whether virtualisation is actually convenient largely depends on the manner in which databases are implemented and used. It isn't always true that the database server, regardless of whether SQL or Oracle, is really only utilised to a certain extent. The frequently mentioned figure of around 30 percent "wasted" utilisation is just for reference and it also applies to those database servers running only a few instances.



However, when it comes to one or more continuously addressed databases in the fields of business intelligence, data mining, online transactions or ERP/CRM, the whole affair should be looked at from quite a different perspective. In these cases it is quite possible that the database server already is undergoing almost full utilisation of its physical resources, i.e. processor, hard drives, SSDs, etc.

Whoever believes that non-existing resources can be conjured almost from nothingness by virtualising a database is much mistaken. Quite to the contrary: even with virtualisation of a database, only the existing hardware can be utilised more efficiently. And whoever chooses to ignore this, not only risks business-critical databases disappearing into virtual Nirvana in case of failure, but often also puts the entire company at risk, because in the end, availability, scalability and speed of a database system should be ensured at all times due to their critical importance for the business.

If something goes wrong there are very few companies able to recover their virtual servers and virtualised databases. Therefore, it is of the utmost importance to have a detailed emergency plan in stock for such eventualities. However, as many system failures and data loss situations are extremely complex and can hardly be solved by the employees themselves, it is advisable to avail yourself of a competent data recovery service when creating an emergency plan. Employing leading data recovery specialists such as Kroll Ontrack, who have already successfully solved many complex data recovery cases of disappeared virtual databases, is often the safest way to get your business-critical data back up and running quickly.