



## Utopian vision for databases

Published: 14th May 2007

Over the course of the last few months I have written a couple of times speculating on the development of appliances that might be more generally deployed than the very specific products we have seen to date. In particular, in my last article on this subject, I discussed the scope for improving the performance of merchant databases by replacing their horizontal approach to data management with a more vertical MPP (massively parallel processing) based solution that could be packaged as an appliance. Just such an offering has now been released by Dataupia.



As far as the technology is concerned, this is based on Dual Opteron based Blades, each of which has 2Tb capacity on 8 attached RAID drives. Additional units, of the same capacity, can be added on demand and as required. On the software side, Dataupia runs on Linux (which you don't touch) and plugs into the top layer of the Oracle (or DB2 or SQL Server) execution engine and then Dataupia takes over execution of database primitives (parse, joins and so forth) and storage management (which you will no longer need to administer) from the database. There is a utility provided to extract or replicate data from existing storage media and there is a fast indexing facility that runs in the background

even while you are loading or working on the data. As one might expect from an appliance, high availability is built-in, together with back-up and restore capabilities.

Perhaps the most important additional point to note is that the implementation of Dataupia has no impact (except to make them run much faster) on your existing applications. That is, your applications do not need to be changed in any way, they simply run against the Oracle, SQL Server or DB2 database as they always have.

In theory, Dataupia could be used to improve the performance of any type of database installation using one of the merchant databases mentioned. However, in practice, most OLTP systems are not large enough to warrant the uses of such a solution—here, clustering or grid technology is probably more appropriate to speed up the environment. At the other extreme, Dataupia is not aimed at the very high performance of the data warehouse appliance vendors such as Netezza and DATAlegro. In practice, therefore, Dataupia is targeting the enterprise data warehouse (EDW) market, especially where you need or would like to keep large amounts of historic data on line so that you can do things such as better account management, and more accurate forecasting. In practice, the sheer volume of data required to have a longer-term perspective has meant that it has been cost prohibitive for many organisations, but Dataupia makes this much more affordable, delivering multi-terabyte solutions without any application porting costs, and at a fraction of the cost of traditional approaches.

The advent of Dataupia is likely to make as much of a disturbance within the market as Netezza did when it first appeared. In particular, it will strengthen the positions of the merchant databases against any rivals that do not integrate with Dataupia such as Teradata and Sybase, as well as newcomers like H-P. It will also inevitably impact on the success of vendors as they seek to add EDW capabilities to their existing offerings, as well as to new entrants into the market such as Vertica.

Of course, Utopia literally means “no place”: if Dataupia is the success that it promises to be then any vendor that does not either partner with it or develop comparable technology of its own, will be likely to find itself with no place to hide.

Philip Howard  
*Director of Research, Bloor Research*



© Bloor Research 2007