



Review of xkoto GridScale

By Dan Kusnetzky, Principal Analyst

INTRODUCTION

After a recent briefing, the Kusnetzky Group will review xkoto's GridScale, place it into the Kusnetzky Group model of virtualization, determine which user segment would be best suited for this technology and present a quick analysis of the company's strengths and weaknesses. The review will also include a quick survey of the opportunities and competitive threats the company faces.

QUICK REVIEW OF XKOTO GRIDSCALE

The GridScale Database Load Balancer uses patent pending technology to enable commercial off-the-shelf databases to run on a cluster of industry standard systems with the same or better reliability and performance as much more expensive, special purpose proprietary systems.

GridScale offers customers a number of benefits including the following:

- The ability to make their database continuously available availability twenty four hours a day every day.
- Increase the scalability of database-based applications by using a grid computing-like environment and spreading the application workload over many machines
- Reduce or eliminate unscheduled downtime
- Offer a more convenient window of time for maintenance and updates
- Synchronization of data across multiple, distributed datacenters
- Increased utilization of computing resources without requiring application changes.

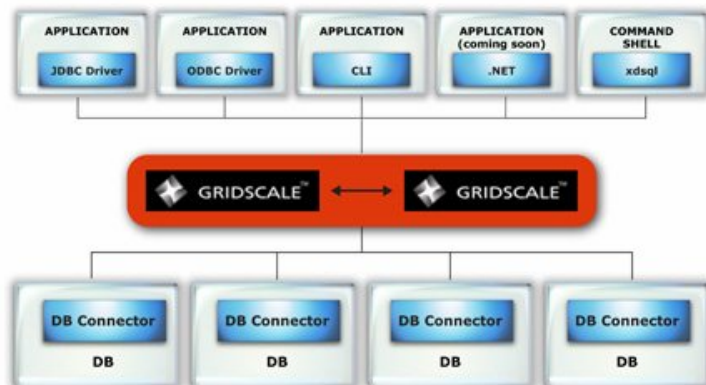
WHAT XKOTO GRIDSERVER ACTUALLY DOES

xkoto GridScale captures the flow of SQL statements before databases have an opportunity to see them.

The statements are replicated using a sophisticated store-and-forward queuing system to other machines so that changes made to the data can be reproduced on all of the systems as appropriate.

The systems in this "application level cluster" don't have to be identical, their storage subsystems don't have to be shared and they only have to see SQL statements that effect data that they actually need to access.

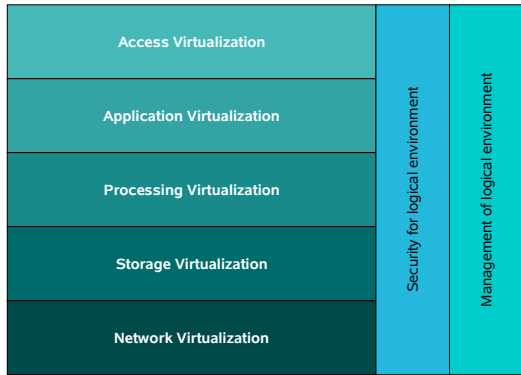
GridScale offers a number of reporting options and administration tools making the software both powerful and easy to use.





WHERE DOES GRIDSCALE FIT IN KG'S REFERENCE MODEL?

Although xkoto refers to GridScale as a "data virtualization" product, it fits best in the application virtualization category. This is due to the fact that it is very specific to creating a virtual environment for database-based applications rather than for every application a system might be supporting. Other applications would not gain the benefits that GridScale provides.



Kusnetzky Group © 2007

Application virtualization is an important component of an effective IT infrastructure. It allows organizations to create an environment that is ideal to fulfill a specific application's requirements even though the physical environment might differ dramatically. This allows the organization to deploy different systems in different datacenters based upon the actual consumption of what that application does. Some datacenters need the most powerful server that is currently available. Others would be best served by a much more modest system.

WHAT SEGMENT OF THE USER POPULATION WOULD FIND THIS MOST VALUABLE?

Many organizations live in an "Internet World" where all of their systems, not just business critical systems, must be available always, must scale enough to handle whatever workload the network imposes and, yet, must be cost effective. Organizations can no longer afford the luxury of over provisioning a datacenter to gain these benefits.

KG expects that organizations in financial services, health care, government and even retail/distribution would find this technology appealing.

SNAPSHOT ANALYSIS

Providing high performance, reliable, scalable access to database data for virtual systems is clearly a challenge for many organizations.

- Keeping database data inside of the virtual server can make the configuration simple but, that approach really isn't scalable or reliable.
- Keeping the data on a storage server can provide higher levels of access and reliability but, it also means deploying parallel database software (so that multiple virtual servers have access to the same data), either storage or database replication software (so that the data can be replicated to several different storage servers for reliability/availability) and, of course sophisticated management software allowing administrators to oversee the operations of all of these different tools.

STRENGTHS

xkoto has looked at the problem of data access, performance, scalability and availability in a new way. Rather than trying to monitor the buffers in the database, monitor output buffers in the operating system, or monitoring the storage subsystem in the hopes of catching changes as they happen and replicate them elsewhere, they've done something new. They're monitoring the input to the database and replicating the SQL commands. The change in thinking appears simple, but the impact is quite profound.

SQL statements are much smaller than blocks or tracks of data. Replication would most certainly be faster and simpler. Tricky coding allowing monitoring and control of the rapidly changing internal database buffers or output buffers for the operating system isn't necessary. It also isn't necessary to scan storage subsystems for changes (a little bit like trying to boil the ocean don't you think?).

WEAKNESSES

xkoto's approach is new and many decision-makers will want to see GridScale in operation before adopting xkoto's product. xkoto has addressed this issue by offering quite an impressive list of reference sites.

The company is also relatively small and would be easy to overlook. Xkoto must be creative to develop a sales and marketing strategy that will get the attention of decision-makers, convince them that GridScale must be considered and provide the necessary materials to convince these decision-makers that GridScale is the approach they simply must have.

OPPORTUNITIES

The Kusnetzky Group is impressed by both the simplicity of xkoto's approach and all of the wonderful things that could be done once the SQL stream can be monitored. We can envision quite a number of future abilities that xkoto could choose to develop based upon this new way to look at the problem.

- It would be possible optimize database operations based upon real application usage without requiring developers to change a single application. The product could "fix up" the SQL statements on the fly before the database engine(s) had a chance to see them.
- It would be possible to improve database security by knowing which systems had access to what data and preventing the injection of malicious SQL statements from systems outside of the cluster.
- With a bit of discipline in the development of database schemas, it would be possible to "cluster" database servers based upon totally different operating systems and database engines.

We could go on and on about the opportunities for new features and services this approach offers. I suspect we're going to hear about partnerships and alliances xkoto will develop with most of the major hardware suppliers, database suppliers, suppliers of management software for virtualized resources.

THREATS

At this point, the competitive treats come from other approaches to the same problem including database replication products, distributed caching products, parallel database products and storage replication products. None of these approaches offer the simplicity and effectiveness of the approach xkoto has developed.

S U M M A R Y

If the organization has faced an intractable problem of how to scale up database-based applications without purchasing expensive parallel database software, configuring complex replication environments or purchasing bigger machines that really needed, it would be wise for their IT decision makers to closely examine GridScale.