WHITE PAPER

TECHNICAL WHITEPAPER
GRIDSCALE® DATABASE VIRTUALIZATION SOFTWARE

Typical enterprise applications are heavily reliant on the availability of data. Standard architectures of enterprise applications have a single point of failure and a performance bottleneck at the database tier. The database typically consists of a large SMP system that has been architected for peak load and can only be scaled vertically by adding additional resources to the SMP system (e.g. adding additional CPUs or memory). Availability considerations for the database typically consist of an identical – and normally idle – backup database system with data synchronization between the systems.

The GRIDSCALE database virtualization solution from xkoto enables the database tier to be architected in the same manner as Web and application tiers – a load balanced pool of servers that provides unmatched reliability and scalability at a lower cost. This whitepaper provides an overview of current enterprise architectures and describes the benefits of using GRIDSCALE for business-critical databases.

CONTENTS

INTRODUCTION

Server consolidation. Provisioning. Utility computing. Grids. The data center is constantly innovating in order to drive down costs while improving the delivery of business services. “Big Iron” environments have been a mainstay in IT operations, but most businesses now have a standing order to rollout upgraded or new applications using lower cost, clustered commodity systems whenever possible. The challenge is to retain all the performance and reliability of “Big Iron” with these commodity components. The GRIDSCALE® database virtualization solution from xkoto makes this possible.

LEGACY ENTERPRISE APPLICATION ARCHITECTURES

Traditional enterprise application architectures are typically divided into three logical layers: Presentation, Application (or Business Logic) and Data:

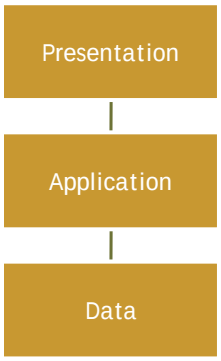

Figure 1 – Typical logical divisions of enterprise applications.

Legacy reference implementations of these applications typically consist of horizontally scaled systems at the Presentation and Application layers, and a single large, vertically scalable, SMP system for the Data layer. To mitigate single point of failure risks associated with the single SMP system at the Data layer, one of several data replication methods can be used to provide a failover system (Figure 2).

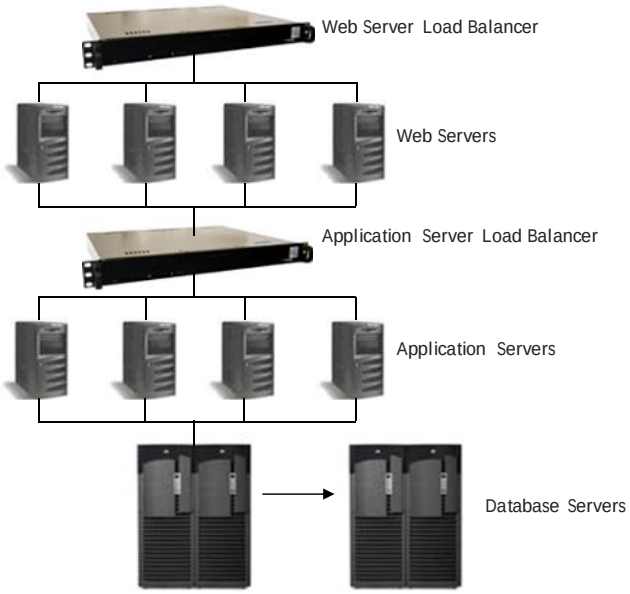

Figure 2 – Legacy reference implementation of enterprise applications

At the Web/Presentation and Application layers, systems are horizontally scaled and load balanced to increase availability and performance. In horizontally scaled architectures, each system has its own independent resources such as processors, memory, and disk. This "shared nothing" architecture provides increased availability. To increase scale and performance of horizontal architectures, new systems are added to the cluster.

By comparison, the vertically scaled systems commonly used at the Data layer have shared processors, memory and disk. To increase application availability, an additional system is used for fail-over purposes. To increase performance, additional resources must be added to the system. If additional resources cannot be added to the system, a larger (and more expensive) system must be purchased.

While it is standard practice to use horizontally scaled architectures for the Presentation and Business Logic layers, vertical SMP systems dominate the Data tier. The inability to maintain data consistency across multiple servers without sacrificing performance has historically limited the viability of using horizontal architectures at the Data tier.

Attempts to solve this issue include:

Two Phase Commit

With Two Phase Commit, the transaction must be successfully committed on all database servers. However, if one of the database servers is off-line or unavailable, the transaction will fail on all the database servers.

Synchronous Transaction Replication

In a Synchronous Transaction Replication model, all transactions are sent to all database servers for simultaneous processing. However, in this case the overall

performance of the cluster is limited by the slowest server because the transaction is not complete until every server in the cluster has completed processing.

Master-Slave Architectures

Master-Slave Architectures have a primary database server with one or more backup copies (known as slaves). The main drawback in this model is that the slave systems only serve as backup systems and do not provide any performance gains. The slave systems also do not provide a real-time consistent view of the data.

Oracle Real Application Clusters (RAC)

While Oracle RAC allows database processing to be moved to commodity clusters, it maintains a single copy of the database which must be implemented in a Storage Area Network (SAN) or multi-ported storage. The data is usually replicated using third-party hardware or software RAID technologies. Oracle RAC is expensive to deploy (licensing fees and SAN/RAID capital/operating costs) and is vulnerable in a disaster scenario since it uses centralized storage.

DATABASE VIRTUALIZATION ARCHITECTURE

To increase availability and performance while driving down the cost of enterprise applications, the ultimate goal is to architect the data layer identically to the other layers – a shared nothing grid of servers that are load balanced for redundancy and scalability (Figure 3).

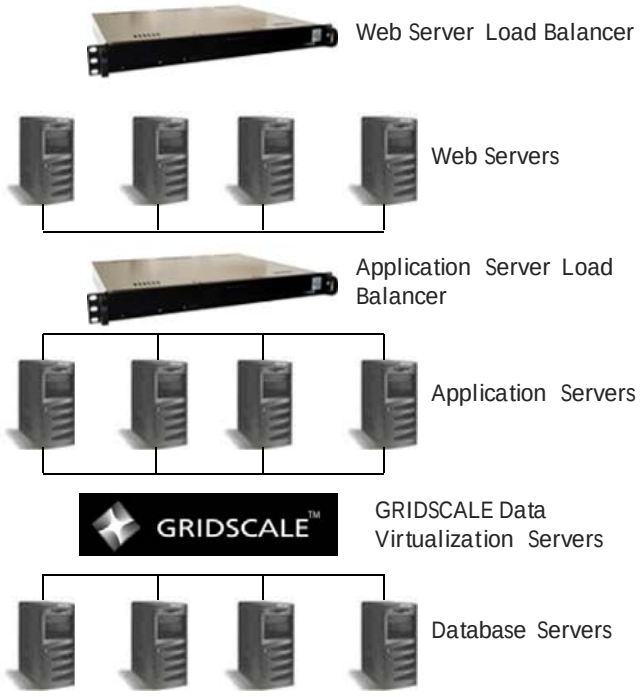

Figure 3 - Next Generation Enterprise Application Architecture

The advantages of a horizontally scaled architecture include:

- Continuous Availability
  In a horizontal grid of servers, one or more systems may fail without affecting the availability of the service. For the Data tier, this enables continuous data access during planned and unplanned outages. Administrators can perform rolling upgrades and other maintenance tasks which previously required planned outages while the application remains available.

- On-Demand Scalability
  Additional servers can be easily added or removed on-demand to ensure sufficient processing power for SLAs without the overhead of having low resource utilization rates.

- Lower Costs
  Lower cost commodity servers are well suited for horizontally scaled architectures. This provides a significantly lower capital and operating cost compared to SMP systems.

- Business Agility
  Centrally managed pool of active/active databases provide enterprises with enhanced ability to change applications more quickly to meet business requirements, to change deployment characteristics and to provision new databases based on application need.

GRIDSCALE allows customers to create a virtualized active/active grid of shared nothing, horizontally scaled database server.

GRIDSCALE DATABASE VIRTUALIZATION

"Storage replication issues go back 30 years. xkoto's innovation was to look at the same problem everyone else was seeing, but rather than look at the output side, ask, 'What if we looked at the input side; what's going into the database?' [This simple idea] is diabolically clever."

- Dan Kusnetzky, Kusnetzky Group

GRIDSCALE enables the clustering of active/active database servers for full redundancy and on-demand scalability. All the database servers in the grid are active, fully utilized and can be dispersed geographically to mitigate disaster scenarios. GRIDSCALE's patent-pending technology maintains complete data consistency across database servers in the cluster and load balances read transactions to increase performance and scalability.

Features of GRIDSCALE include:

CONTINUOUS AVAILABILITY

A key feature of horizontally scaled servers is their resiliency. By sharing the load across multiple servers, the loss of any one server does not have an impact on the availability of the service. With GRIDSCALE, this powerful feature is now available for the Data tier. In addition, since servers can be taken off-line without impact, customers can also perform rolling upgrades on the database servers for true continuous availability.

HORIZONTAL SCALABILITY

Performance tests show that with each new database server added to the cluster, application performance scales in a near-linear fashion, particularly for read-intensive applications. With GRIDSCALE you can scale your application horizontally across low-cost servers and virtual machines. Plus you can provision and de-provision database servers to accommodate the changing needs of your applications.

DATA INTEGRITY

GRIDSCALE is the only database virtualization solution of its kind to support both synchronous and asynchronous data replication between commercial databases in a pool. GRIDSCALE replicates all write statements (inserts, updates, deletes and DDL) in real time to every database server in the grid simultaneously. In its default asynchronous mode, GRIDSCALE does not have to wait for every response before signaling statement completion. The response from the first database server to finish processing the statement is sent back to the application ensuring the fastest response time. The other database servers continue to process the statement normally thereby ensuring that data is replicated on all servers.

GRIDSCALE can also operate in synchronous mode, whereby it waits for responses from all database servers before returning a response to the application. Synchronous replication is especially useful in high performance, write-intensive environments, or where network latency can make it difficult for remote servers to stay up-to-date.

HETEROGENEOUS PLATFORM SUPPORT

GRIDSCALE supports databases running on multiple platforms, including mixed operating systems or mixed databases. For example, with GRIDSCALE, you can run

a pool of DB2 on Linux and SQL Server on Windows. It provides IT architects with unprecedented flexibility to design applications for availability and performance, to migrate platforms or database versions, and to replicate data between disparate systems.

APPLICATION TRANSPARENCY

GRIDSCALE uses standard interfaces to control communication between applications and underlying databases. GRIDSCALE comes with JDBC and ODBC compliant drivers, and also supports .NET to minimize application changes or changes to the underlying database definition. Also, GRIDSCALE offers a non-intrusive, cost-effective Assessment Service to provide compatibility assurance for your applications.

SUPPORT FOR MULTIPLE DATA CENTERS

GRIDSCALE eliminates the need for expensive and intensive data replication strategies between production and disaster recovery (DR) sites. With GRIDSCALE you can solve both HA and DR issues with a single solution. Using a highly efficient communication protocol, GRIDSCALE allows database servers to span multiple, geographically dispersed data centers. As a result, you can load balance across multiple data centers, tap into idle capacity at DR sites and ensure that your DR database is always up-to-date.

EASY-TO-USE, COMPREHENSIVE MANAGEMENT

GRIDSCALE's Web-based console and command-line interface provide straightforward management and administration of GRIDSCALE and database servers. Diagnostic and fault event notification for GRIDSCALE and the underlying managed environment can be forwarded to an enterprise management system for processing.

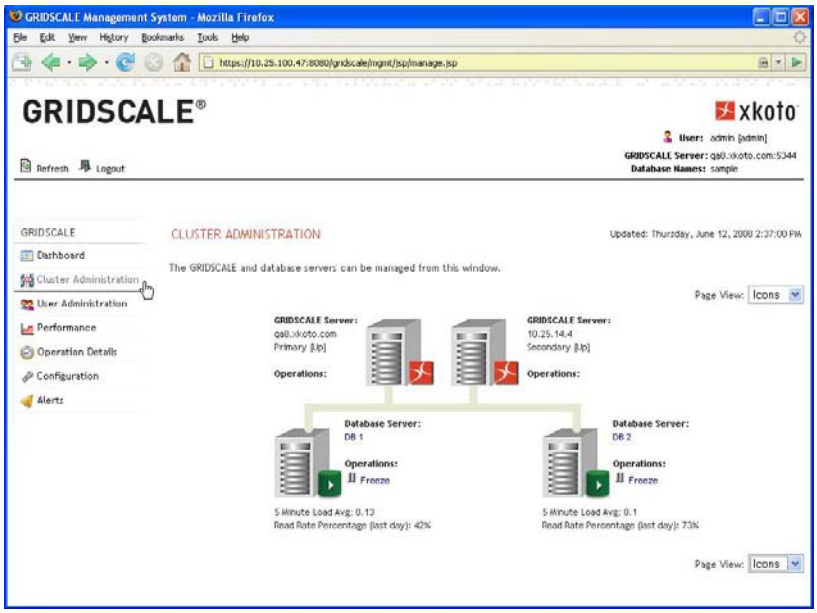

Figure 4 – GRIDSCALE Management Console

GRIDSCALE ARCHITECTURE

"InGrid is a home security company that deploys an innovative distributed digital wireless grid system. With their database receiving inputs from customer systems and serving up real-time status information, reliability is everything. With InGrid's customer base set to grow significantly, xkoto provides the non-stop security they demand."

- Drew DeNardo, InGrid Systems

The GRIDSCALE architecture consists of a set of client drivers, the GRIDSCALE server(s) and a DB Connector on each database server. Typically customers do not need to change their application code or databases in order to use GRIDSCALE.

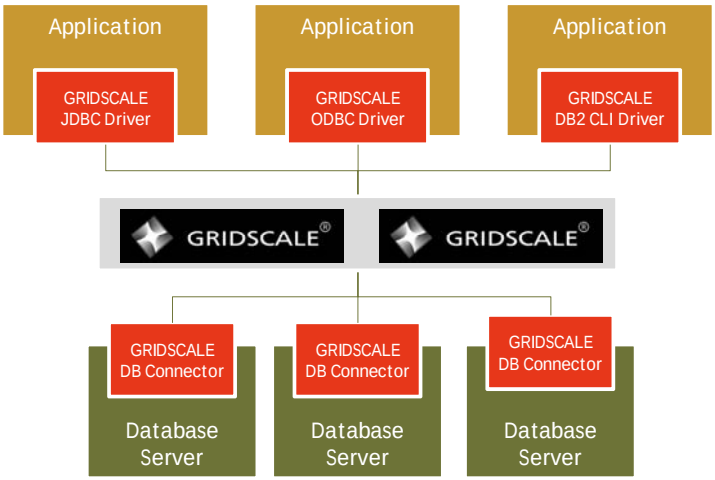

Figure 5 - GRIDSCALE architecture

GRIDSCALE SERVER

The GRIDSCALE server appears as a single fully functional database to all applications. Applications are unaware that any one of several database servers may be used to fulfill their requests. The GRIDSCALE server can also be clustered for highly availability.

GRIDSCALE employs a shared-nothing architecture with a configurable number of database servers for continuous availability. Each database server operates completely independently of the others, maintaining its own complete and consistent copy of the database.

Real -Time Asynchronous Statement Replication

GRIDSCALE supports both synchronous and asynchronous data replication to ensure data consistency.

In asynchronous mode, GRIDSCALE replicates all write statements (inserts, updates, deletes and DDL) in real time to every database server in the grid simultaneously, but does not wait for every response before signaling statement completion. The response from the first database server to finish processing the statement is sent back to the application ensuring the fastest response time. The other database servers continue to process the statement normally thereby ensuring that data is replicated on all servers.

This design provides maximum throughput and availability without the "hold up" behavior of synchronous replication or the lag of Master-Slave solutions.

GRIDSCALE can also operate in synchronous mode, whereby it waits for responses from all database servers before returning a response to the application. Synchronous replication is especially useful in high performance, write heavy environments, or where network latency can make it difficult for remote servers to stay up-to-date.

Intelligent Query Load Balancing

GRIDSCALE intelligently load balances read requests, by selecting the database server that will return up-to-date data the fastest at the time of the request and sending the read request to only that server. In the event of database failure, GRIDSCALE automatically and transparently reroutes outstanding read requests against the failed server to another available server. With Intelligent Query Load Balancing, GRIDSCALE achieves linear scalability when additional nodes are added to the database server grid.

Recovery Log

GRIDSCALE maintains an on-disk queue for SQL transactions that need to be applied to offline databases. This allows GRIDSCALE to provide point in time recovery in the event of a database server failure or planned outages.

Management Tools

A Web and command line interface is provided to allow for easy management and administration of GRIDSCALE and the database servers. Using the management interface, users are able to add/remove, enable/disable and monitor database servers, view alerts, and generate reports on performance.

Management of the database itself can be accomplished using any ODBC, JDBC or DB2 CLI-compliant database management application connected through GRIDSCALE.

xdsql™

xdsql is the command-line SQL interface to GRIDSCALE. xdsql allows customers to easily distribute SQL statements to all the databases in the grid using a command line interface. It provides convenient features such as command recall and in-line command editing.

GRIDSCALE CLIENT DRIVERS

GRIDSCALE includes ODBC drivers, fully compliant Type 4 and 2 version 3.0 JDBC drivers, and a DB2 CLI driver for native IBM DB2 and Microsoft SQL Server databases. These drivers provide the necessary communication between the application and GRIDSCALE. They also make it possible to use GRIDSCALE without modification to application code or the underlying databases. The GRIDSCALE database drivers are enhanced to support communication with a cluster of database servers while presenting a consistent view of the database to applications.

DB CONNECTOR

The xkoto DB Connector is installed on each database server. Currently GRIDSCALE supports both IBM DB2 and Microsoft SQL Server. Additional databases will be supported over time. Instead of providing a generic interface for all database vendors,

the DB Connector allows GRIDSCALE to support applications that utilize features unique to a specific database vendor.

DATABASE ROUTING COMMUNICATION PROTOCOL

All components of the GRIDSCALE solution communicate using xkoto's lightweight and efficient Database Routing Communication Protocol (DRCP) over standard Ethernet connections. This low-bandwidth protocol allows GRIDSCALE to communicate with database servers regardless of location, and keep them in sync and up to date. The DRCP protocol can also be encrypted to protect the data from the application to the database server.

GRIDSCALE BENCHMARK

GRIDSCALE benchmarking conducted at IBM's Linux Integration Center Lab in Austin, Texas has confirmed the results below.

E-COMMERCE BENCHMARK

The e-Commerce Benchmark is similar to TPC-W, a transactional system that defines a series of workflows suitable for evaluating e-commerce systems. These benchmarks utilize "typical" transactions that are modeled after the operations of an online retail store.

The Shopping Mode of the Benchmark depicts a user selecting items from an online store and placing them in their "shopping cart". The benchmark utilizes 80% read-only queries and 20% write transactions.

THE RESULTS

This test shows that GRIDSCALE provides over 85% scalability with near-linear performance for every database server that is added (1-7 nodes shown) to the cluster. Applications with similar read/write characteristics should demonstrate similar scalability benefits, and more read-intensive applications can scale even faster. Based on these and other performance tests, GRIDSCALE is projected to deliver near-linear performance for up to 30 database nodes.

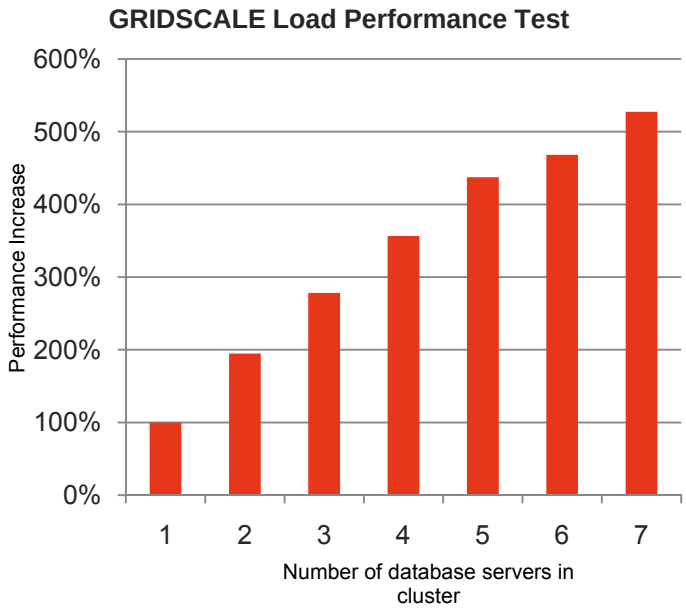

Figure 6 – Scalability benchmark results using GRIDSCALE with 5 database servers.

GRIDSCALE FOR IBM DB2

GRIDSCALE for IBM DB2 has been specifically built to integrate with DB2 APIs and interfaces. GRIDSCALE for IBM DB2 has also been tested by IBM at its lab in Markham, Canada for compatibility. Specific features for DB2 include the following:

DB2 CLI CLIENTS

GRIDSCALE for IBM DB2 provides a Call Level Interface (CLI) for DB2. This allows native DB2 applications to connect directly to GRIDSCALE transparently.

SQL PARSER FOR DB2

GRIDSCALE for IBM DB2 is 100% compatible with the SQL dialect used by DB2. This allows existing DB2 applications to use GRIDCALE for DB2 without any code changes.

XDB2

xdb2 is a cluster-friendly equivalent command line processor (CLP) to IBM's DB2 CLP. xdb2 allows customers to issue SQL and DB2 maintenance commands (e.g. load, runstats, etc.) through GRIDSCALE so that the command is applied to all the database servers. This allows the GRIDSCALE cluster to be managed like a single instance of DB2 reducing management costs.

DB2 DATA WAREHOUSES

GRIDSCALE for IBM DB2 provides availability for business-critical data warehouses for round-the-clock operation. You can use GRIDSCALE for IBM DB2 to provide disaster protection for warehouses in different geographic data centers.

In addition, GRIDSCALE enables you to perform maintenance on warehouses without outages. For example, if you plan on repartitioning a database across multiple database servers, GRIDSCALE for IBM DB2 allows you to accomplish this without an outage.

TECHNICAL SUPPORT BY IBM

GRIDSCALE for IBM DB2 is officially supported by IBM. As a business running GRIDSCALE with DB2, you'll have one-stop, 24x7 support from IBM and the convenience and confidence of having a single-point of contact for support issues.

Technical support for other database platforms is available from xkoto and its authorized partners.

SOLUTION COMPARISON

This table shows how GRIDSCALE compares to other clustering and high availability (HA) solutions for DB2:

Solution	GRIDSCALE® Database Virtualization	WebSphere Replication Server (Q-Rep)	HACMP	HADR
Hardware Utilization	All Active	All Active – however, requires manual load balancing	Active-Passive	Active-Passive
Horizontal Scalability	Yes	Yes, however it requires manual load balancing	No	No
Automatic Failover	Yes	No	Yes	No - Requires other cluster solution such as Tivoli Systems Automation (TSA)
Protection Level	Database	Table	Server	Database
Points of Failure	None	None	Disk	None
Consistent Data View	Yes	No	Yes	Yes
Easy to Use/Maintain	Yes	No	No	Yes

GRIDSCALE FOR MICROSOFT SQL SERVER

GRIDSCALE Database Virtualization software is also availble for Microsoft SQL Server. The GRIDSCALE Server runs on Windows Server 2003 or higher and supports SQL Server 2005/2008. With GRIDSCALE for SQL Server, you can escape the scalability limitations of SQL Server clusters, patch and maintain database servers without outages, and eliminate the cost and complexity associated with traditional clusteirng solutions. Plus, since all database instances are active-active, you eliminate the waste assocated with to passive standby and witness server hardware.

GRIDSCALE for SQL Server currently supports JDBC, ODBC connectivity. xkoto will be adding .NET and ADO connectivity shortly.

SOLUTION COMPARISON

This table shows how GRIDSCALE compares to other clustering and high availability solutions for SQL Server:

	Database Virtualization with GRIDSCALE®	Database Clustering (e.g, MSCS)	Database Mirroring
Single point of failure	No	Yes	No
Automatic failover	Yes	Yes	Yes, with Witness server.
Disaster Recovery (DR) capabilities	Yes	No	Yes
Allows for rolling upgrades to hardware and Operating System	Yes	Limited – requires single point of failure during maintenance	Limited – requires single point of failure during maintenance
Rolling upgrades to database software	Yes	No	No
Rolling upgrades to data	Yes	No	Limited – requires single point of failure during maintenance
Horizontal scalability	Yes, for read heavy applications.	No	No
Performance overhead	Additional network hop introduces slight latency.	None	Synchronous mode introduces network latency.

CONCLUSION

This whitepaper provided a brief overview of current 3-tier architectures and the differences between horizontally and vertically scaled systems. While it is standard practice to use horizontally scaled architectures for the Presentation and Application tiers, the Data tier is almost always implemented using a large vertically scaled SMP system.

Horizontal architectures powered by database virtualization provide significant advantages in resiliency, on-demand scalability and most of all, capital and operating cost. The driving force behind the creation of GRIDSCALE is the desire to take advantage of horizontal architectures at the data tier.

The GRIDSCALE database virtualization solution uses groundbreaking technology to offer full advantage from commodity systems at the data tier.

FIND OUT MORE FROM XKOTO.COM

Technical Specifications – http://www.xkoto.com/products/technology.php

Customer Successes – http://www.xkoto.com/customers

xkoto Database Virtualization Weblog – http://www.xkoto.com/blog

Frequently Asked Questions - http://www.xkoto.com/products/faq.php

White Papers, Case Studies and Downloads - http://www.xkoto.com/resources/

Evaluation Copy of GRIDSCALE - http://www.xkoto.com/products/demo.php

ABOUT XKOTO

Based in Waltham, MA with research and development offices in Toronto, Canada, xkoto, Inc. is a fast growing enterprise software company that provides database virtualization solutions for global businesses.

For more information, visit http://www.xkoto.com.