



SYBASE®

DATABASE VIRTUALIZATION USING CLUSTERS: CHEAPER, EASIER, FASTER, BETTER

PETER THAWLEY
SENIOR DIRECTOR/ARCHITECT
CTO GROUP, SYBASE



Database Virtualization Using Clusters: Cheaper, Easier, Faster, Better

Carl Olofson
Research Vice President
IDC



How We Have Been Managing Databases Up to Now

- From Client/Server to Three-Tier Configurations
- Database Proliferation

Data Center Clutter

- Problems with Dedicated Servers and Storage
- Common Remedies

Virtualizing Databases

- How and Why
- Virtualized Database Clusters and Storage in a Grid Deployment

Recommendations and Conclusions

The Problem

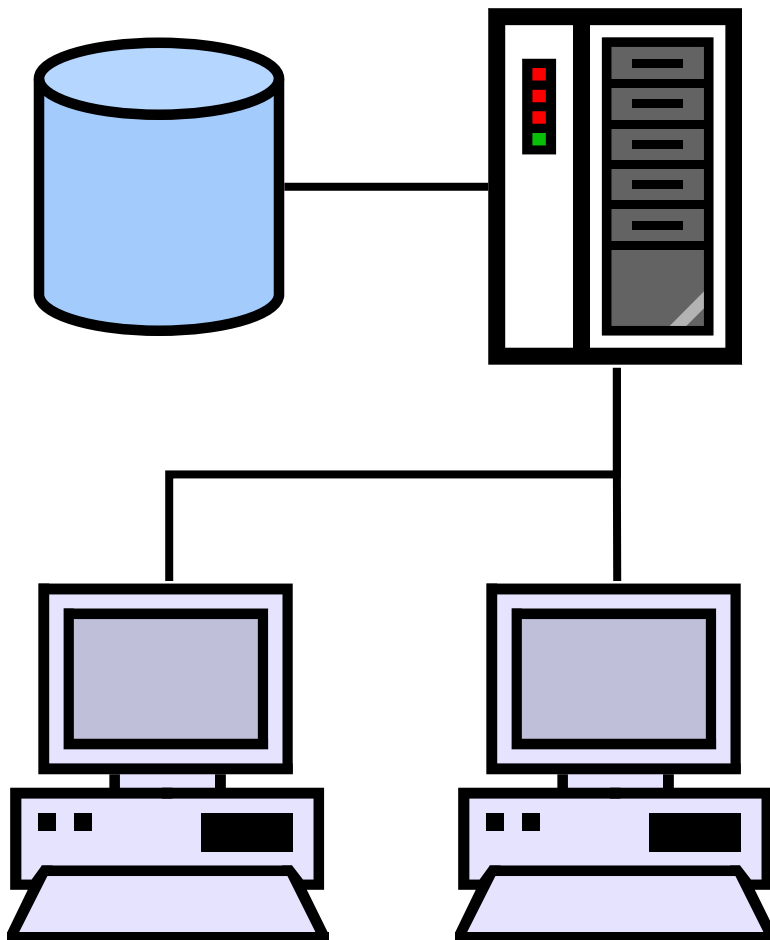
Demand for databases is exploding

- Databases are growing as the business grows
- New databases are needed for business intelligence, business process optimization, master data management, etc.
- Each new database is deployed on its own server or cluster, with its own storage

The result of current practice:

- Each server or cluster is intentionally underutilized, in order to ensure smooth performance during peak times
- Storage is underutilized, to optimize access across volumes, and to allow for growth
- The amount of servers, server clusters, storage volumes, and storage arrays is growing, becoming increasingly expensive to manage

Client/Server Configuration




Simple Configuration

- One for each application
- Application code on the workstation
- DBMS code on the server
- Dedicated Storage

All database applications were set up this way on Unix and PC local area networks, once upon a time

Three-Tier Configuration

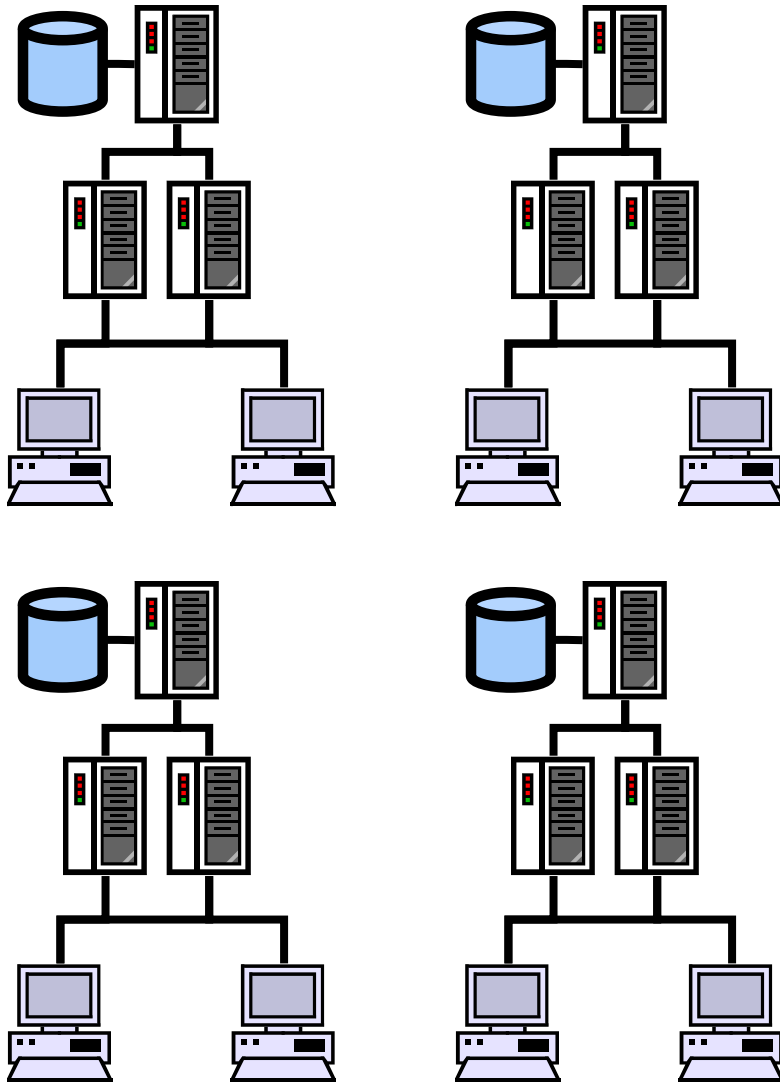


More Complex Configuration

- One DB server for each application
- Business logic on one or more application servers
- User interface code on each desktop, or in Web servers
- DBMS code on the DB server
- Dedicated Storage

Business logic deployment can scale by adding application servers, but DB logic cannot, so DB servers would grow into large SMP boxes.

Adding Many Applications and DBs




Servers and Storage Pile Up

- Each server requires an administrator
- Each database requires a DBA team
- Networking is more complicated
- Floor space, power consumption and air conditioning become problematic

Computer rooms become cluttered with servers and storage, and operations schedules become nightmarishly complex

Adding Many Applications and DBs

Computer room clutter creates...



- Increased costs
 - Hardware
 - Power and cooling
 - Staff time
- Increased risk
 - More complexity means more error
 - Brittle operations means critical jobs may not get done
 - More downtime results

Also, this situation inhibits the deployment of new databases that may be critical to your business!

Problems with Dedicated Servers and Storage

Storage and server underutilization

- Servers are typically used at below 40% capacity
- Storage is almost never at more than 60%
- This means throwing money away on unused facilities

The staff is trapped in a cycle of low-value activities

- More database servers means more complicated operations schedules and procedures
- DBAs must tune server and storage usage one by one
- High-value activities, such as database structure optimization and application consultation, are left undone

For Applications

- Clustering and grid-based deployments increase scalability and deployment flexibility, reduce over provisioning and downtime risk
- Virtualization allows multiple applications to operate within a single server, and to be moved easily from server to server
- Application consolidation is sometimes employed to reduce the number of instances
- Can be helpful for smaller reporting, test, and development databases


These Don't Work For Large or Performance-Sensitive Databases Using Generic Tools

- Clustering and virtualization require internal support from the DBMS, or the result can be poor performance and difficulties in tuning
- Database consolidation requires a major and risky project

Creating Pools of Resources

- If storage and servers can be dynamically assigned
 - Servicing the servers becomes separate from managing the database
 - Servers can be clustered with shared storage for automatic failover
 - Mirroring ensures 24x7 operation
- If database servers and clusters can be virtualized
 - Multiple databases can share servers and storage
 - Databases can be dynamically redeployed without changing their internal operating parameters
- Benefits
 - Servers and storage can run at nearly 100% utilization
 - Expanding capacity means adding servers or storage
 - Lower-cost servers and storage can be used
 - Common management tasks can be applied across a range of databases, saving staff time and simplifying operations

Virtualized Database Clusters and Storage in a Grid Deployment



Everything is virtualized

- Databases are kept on virtual volumes
- Database servers are virtualized, and organized into virtual clusters
- Database storage and server resources can be reassigned at will

Benefits

- Networks are simpler, requiring less care and attention
- Everything serves as a standby resource for everything else, optimizing availability
- The resource pool approach fits neatly into any cloud-based database deployment scenario

Conclusions and Recommendations

Conclusions


- Hybrid shared disk clustering, combined with database virtualization can only be done within the DBMS
- The result is that many databases can be concentrated on less storage and fewer servers
- Large databases can be redeployed from expensive SMP boxes to clusters of smaller commodity servers
- Costs are reduced due to lower equipment, power, and air conditioning
- Costs are further reduced because simpler configurations mean less pressure on administrative staffs, and risk of error is also reduced

Conclusions and Recommendations

Recommendations

- Databases should be deployed using the right technology for each job
- Shared disk and shared nothing each has its place in certain workloads
- For database application proliferation, and for reducing the cost of large database applications, a hybrid shared disk approach combined with database virtualization may be the right way to go
- Database administrators should demand of their technology suppliers scalability, flexibility, availability, simplicity, and features that help reduce overall cost and risk

Contact Information



Carl Olofson

Research Vice President

colofson@idc.com



SYBASE®

VIRTUALIZATION RESOURCES IN PRODUCTION DATABASES

IT TAKES MORE THAN HARDWARE AND VIRTUAL MACHINES
TO EXCEED SLAs!

PETER THAWLEY
SENIOR DIRECTOR/ARCHITECT
CTO GROUP, SYBASE





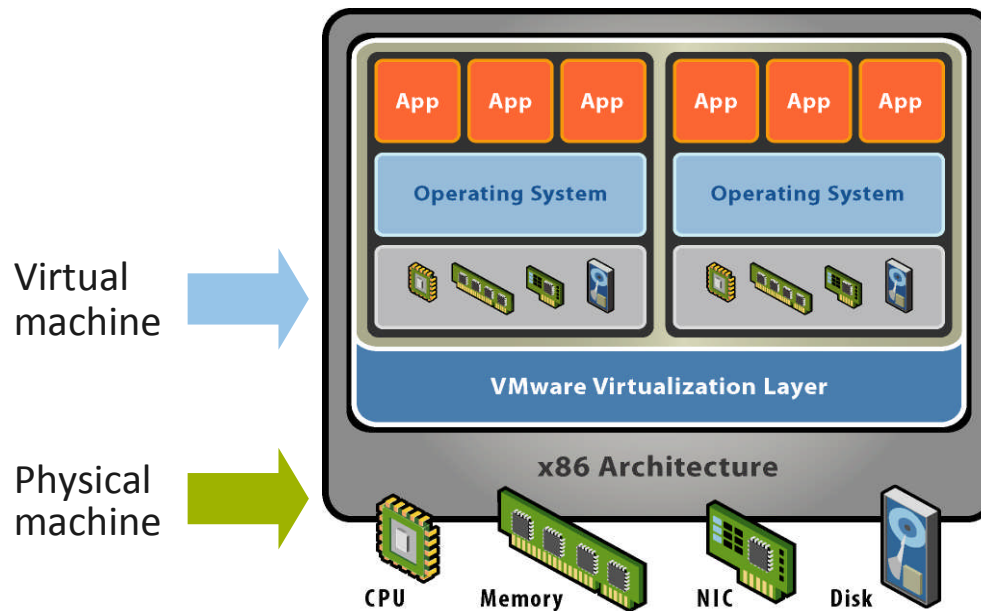
VIRTUALIZATION COMES IN MANY FORMS

Mature concept in many different areas of computing

- **Virtual memory** allows the OS to “fool” the application into thinking it has more memory than it really does
- **Platform virtualization** abstracts an operating system from the underlying HW platform resources
- **Resource virtualization** abstracts specific system resources
 - **Storage virtualization** abstracts physical location to allow flexible movement of files for load balancing, tiering, migration, etc.
 - **Network virtualization** abstracts network addresses and domains to allow applications
- **Data virtualization** abstracts differences in data stores (RDBMS, files, Web services, etc.) and location to provide applications with consistent syntax and semantics
- **Application virtualization** allows the hosting of individual applications on alien hardware/software

VIRTUALIZATION – THE DEFACTO USAGE*

- A level of indirection between hardware and software
- Virtual machines: A full abstraction of a physical machine
 - Safely and efficiently multiplex virtual hardware on physical hardware



- Benefits: Resource sharing, portability, mobility, compatibility

* Slide content courtesy of VMware®, an EMC² company



VIRTUALIZATION MEETS PRODUCTION

General assumptions about which applications/systems are candidates to virtualize

- Consume “small” % of system capacity
 - Adequate capacity of system to “timeshare” simultaneously
- Execute for “short” durations
 - Grid-style application partitioning — may consume entire system capacity, but only for limited and mutually-exclusive amounts of time
- Well-behaved in a virtual machine
 - Shares resources “well with others” — no “idle spinning” or “wasting” of resources others could use
- Not sensitive to [variable] latency
 - Performance overhead diminishing, but still present
 - Virtualization at odds with deterministic performance



PRODUCTION DATABASE INSTANCES

Are production databases candidates for virtualization?

- As always, it depends
 - Smaller, no-business-critical production database systems are being successfully virtualized today — Sybase running 100s of production database instances today
- But production databases for business-critical, tier 0/1 applications generally don't fit previous assumptions
 - Utilization often too high (as a % of capacity)
 - Multi-user and performance requirements often too significant
 - Predictability often valued more than cost of resource sharing
 - Risk of “yet another moving part”



PRODUCTION DATABASE INSTANCES

Are production databases candidates for virtualization?

- As always, it depends
 - Smaller, no-business-critical production database systems are being successfully virtualized today — Sybase running 100s of production database instances today
- But production databases for business-critical, tier 0/1 applications generally don't fit previous assumptions
 - Utilization often too high (as a % of capacity)
 - Multi-user and performance requirements often too significant
 - Predictability often valued more than cost of resource sharing
 - Risk of “yet another moving part”
- QUESTION: Aren't these business-critical, tier 0/1 applications exactly the ones that need the benefits of virtualization most?
 - Flexible resource management that allows resources to be moved around the system in response to capacity needs



VIRTUALIZING RESOURCES WITHIN THE DATABASE

Is this really useful?

- These databases often serve multiple applications or workload profiles
 - Often simultaneously
 - Often with different SLAs
- Consider a trading system with concurrent OLTP and DSS
 - Risk of starving OLTP application high
 - OLTP short duration execution profile before blocking
 - DSS, especially in large memory configurations, don't block often
- How can you “carve up” resources and dynamically bind resources to applications?




RESOURCE MANAGEMENT IN THE 21st CENTURY

- Multiple applications per server is normal
 - Need to ensure “service levels” to end users
 - Resource requirements vary by application, module, time of day, and geography
- Batch and maintenance windows shrinking
- Software parallelism is the brute-force approach to good response time
 - Easy to saturate large hardware with operational decision-support
- Need to provide dynamic flexibility in customizing resource allocation to meet diverse and changing needs
 - “Transaction processing” versus “batch processing”

VIRTUALIZING DATABASE CACHE MEMORY


Treat memory as an optimizable resource



- Database, table, and index assignments to specific caches
- Run-time object to cache assignments and memory movement
- Configurable reduction in contention on buffer synchronization
- Complete control over resource consumption by object
 - Only rational way to support large memory
 - Ensure SLAs for key transactions with guaranteed resource reservation

VIRTUALIZING CPUs, CORES, AND HW STRANDS

Virtual process management



- Customize resources tailored to application environment
 - Affinity threads, procedures, or application to engine(s)
 - Task priorities per thread, application, or stored procedure
- Predictable performance for mixed workloads
 - Guaranteed resource reservation to ensure capacity for SLAs
 - Low-priority work doesn't impact high-priority work
- Dynamic “application to resource” mapping
 - Move resources between applications as needed

VIRTUALIZING NODES IN A CLUSTER


Logical clusters

Clients

Application-specific logical clusters

Physical instances


Databases



SYBASE ASE CLUSTER EDITION

Ensure continuous availability with shared-disk clusters

Handle multiple failures gracefully with faster failover and failback



Balance system load automatically and transparently

Leverage all nodes in the cluster by routing requests automatically, based on business rules



Minimize performance impact of maintenance activities

Use standby nodes to take others offline for routine maintenance



Simplify implementation with virtualized resources

Deploy and configure workloads easily for maximum performance and reliability





ASE CLUSTER EDITION

Get the benefit of a mature “private cloud” infrastructure without the hype

- **Elasticity** – dynamically grow and shrink resources through DBA tools or programmatic interfaces using policy-driven load balancing with no application code required
- **Low TCO** – minimize CAPEX expenditures from higher utilization and virtual resource management
- **Multi-tenancy** – simplify operations while ensuring SLAs at a granular level through elasticity and isolation
- **Reliability** – continuous availability with policy-driven fail-over and fail-back with no application code required
- **Security** – secure infrastructure using traditional enterprise database techniques and granular (column-level) encryption with no application code required



SUMMARY

- Production databases require more than HW and VMs to gain value from virtualization
- Requires RDBMS internalize virtualize both SMP and cluster hardware resources and dynamic resource assignment for applications
 - Instances: Nodes of a shared disk cluster
 - Engines: CPUs, cores, and HW strands (threads)
 - Memory: Named caches and buffer pools
 - Network: Listeners
 - Disks: Devices
- Why?
 - Service-level “predictability”!
 - Flexibility to move resources as dictated by demand or operational needs



Resources to learn more

- For more information, visit <http://www.sybase.com/clusters>
- To inquire about a cluster edition briefing, email clusterbriefing@sybase.com
- For information on the Development Assistance Program, email amit.satoor@sybase.com
- Send us an email with your thoughts and questions

Peter Thawley Senior Director and Architect, Data Management
peter.thawley@sybase.com

Amit Satoor ASE Cluster Edition Product Manager
amit.satoor@sybase.com

Thank you
for
attending.

For more information
www.sybase.com/clusters



SYBASE®