IBM Information ON Demand 2010

Intelligent Analytics and Optimization for Smarter Business

Smart Database: Lower the Cost of Data with Improved Service Level

Daniel Ling
Senior Sales Specialist
Software Group
IBM HK

IBM Information On Demand 2010

Agenda

- Explore various way on failover and scale out load balance for database
- DB2 failover and scale out design
- Introduce DB2 pureScale Active-Active share disk

Critical IT Applications Need Reliability and Scalability

- **Local Databases are Becoming Global**
  - Successful global businesses must deal with exploding data and server needs
  - Competitive IT organizations need to handle rapid change

Customers need a highly scalable, flexible solution for the growth of their information with the ability to easily grow existing applications

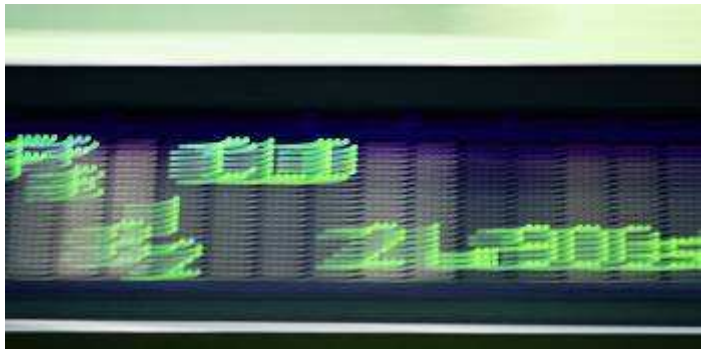

- **Down-time is Not Acceptable**
  - Any outage means lost revenue and permanent customer loss
  - Today's distributed systems need reliability

IT Needs to Adapt in Hours…Not Months

- **Handling Change is a Competitive Advantage**
- **Dynamic Capacity is not the Exception**
  - Over-provisioning to handle critical business spikes is inefficient
  - IT must respond to changing capacity demand immediately, not months later

Businesses need to be able grow their infrastructure without adding risk

- **Application Changes are Expensive**
  - Changes to handle more workload volume can be costly and risky
  - Developers rarely design with scaling in mind
  - Adding capacity should be stress free and not require application changes

1) Traditional Server Based Failover

DB2 automation with built in cluster manager
(Can also use HACMP, MSCS, Sun, Veritas, Heartbeat, etc)

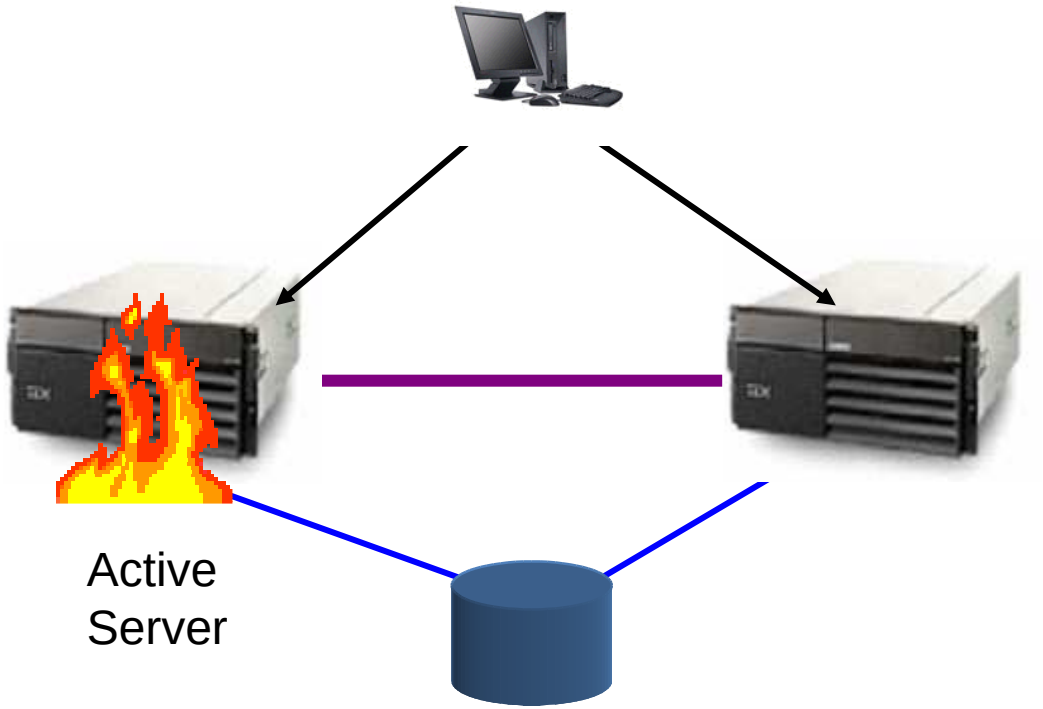

Pros:
- Inexpensive local failover solution
- Protection from software or server failure
- DB2 9.5 integrated with TSA cluster manager

Cons:
- No protection from disk failure
- No protection from site failure
- Failover times vary from 1 to 5+ minutes

2) Database log shipping / mirroring design

- **Redundant copy of the database to protect against site or storage failure**
- **Support for Rolling Upgrades**
- **Failover in under 15 seconds**
  - Real SAP workload with 600 SAP users – database available in 11 sec.
- **100% performance after primary failure**

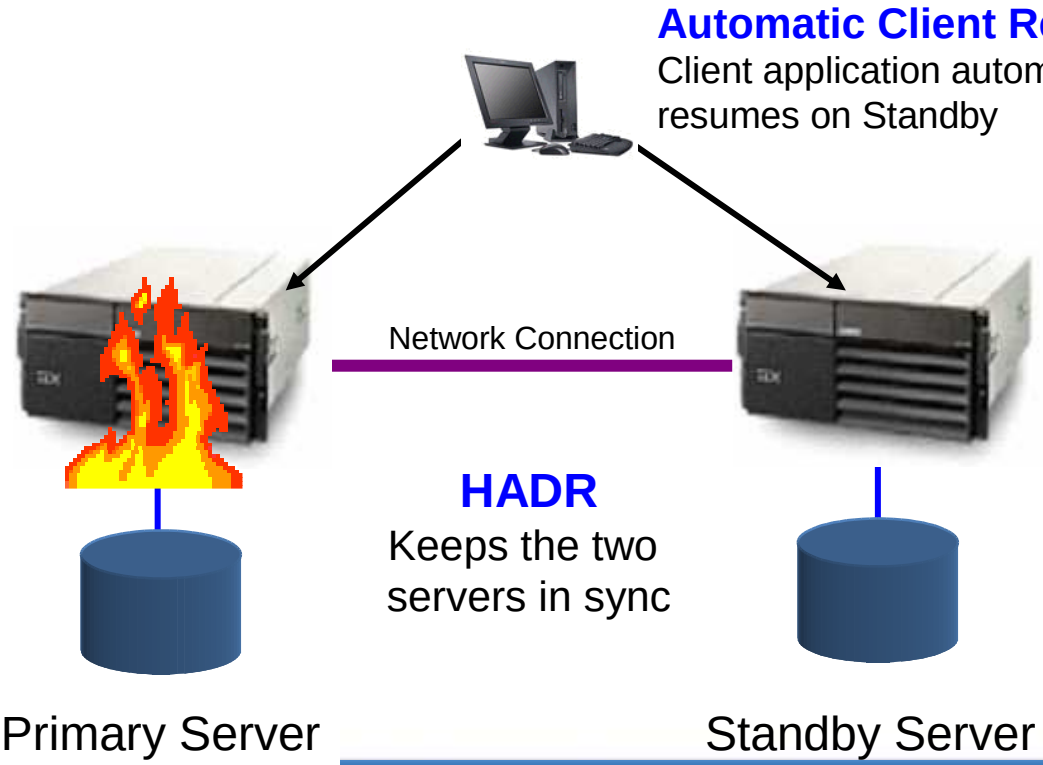

Synchronization Modes

Synchronous, Near Synchronous, Asynchronous

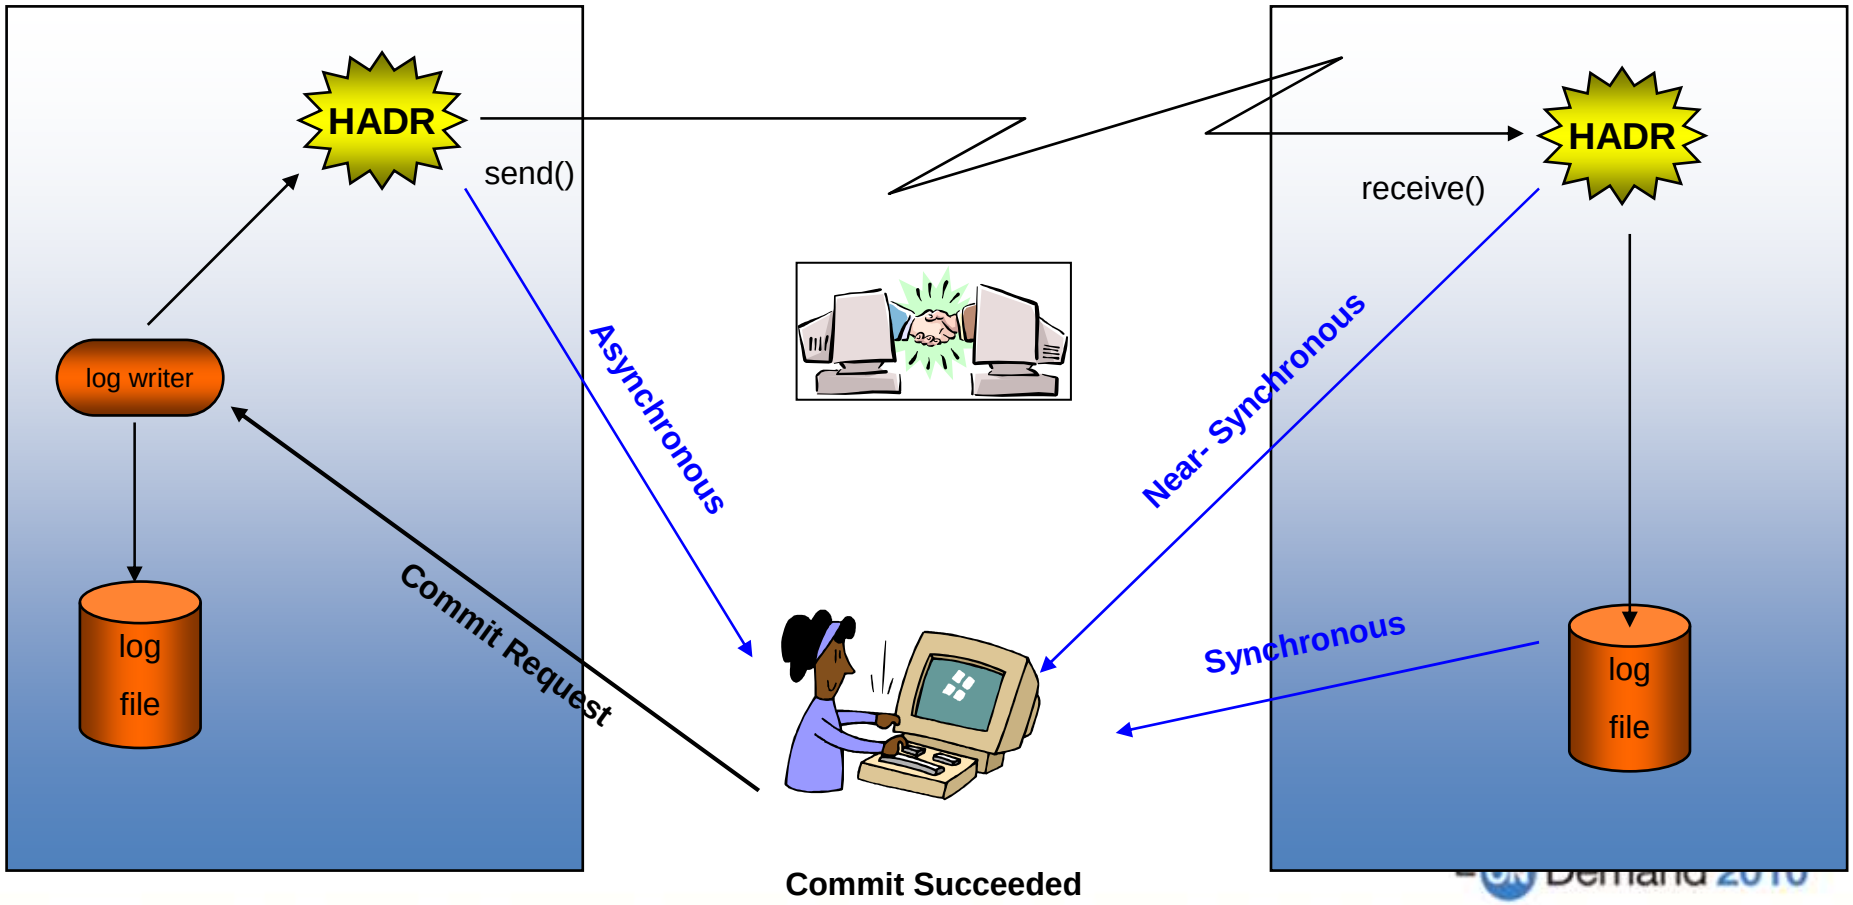

Pros & Cons

Pros:

- DB2 Workgroup and DB2 Enterprise included HADR (FREE of charge)
  - Oracle Data Guard need charge and only available on Oracle EE
- DB2 9.7 allow read on Standby Server make it an active-active design
- Support Rolling Upgrade for uninterrupted upgrade
- Protected from software, server, disk, site failure (no distance limitation)
- Simple to setup, monitor and perform automated failover
- Failover time in the range of 30 sec or less

Cons:

- Two full copies of the database, redundant storage
- Only primary server can perform update

3) Continuous Availability Using Q Replication

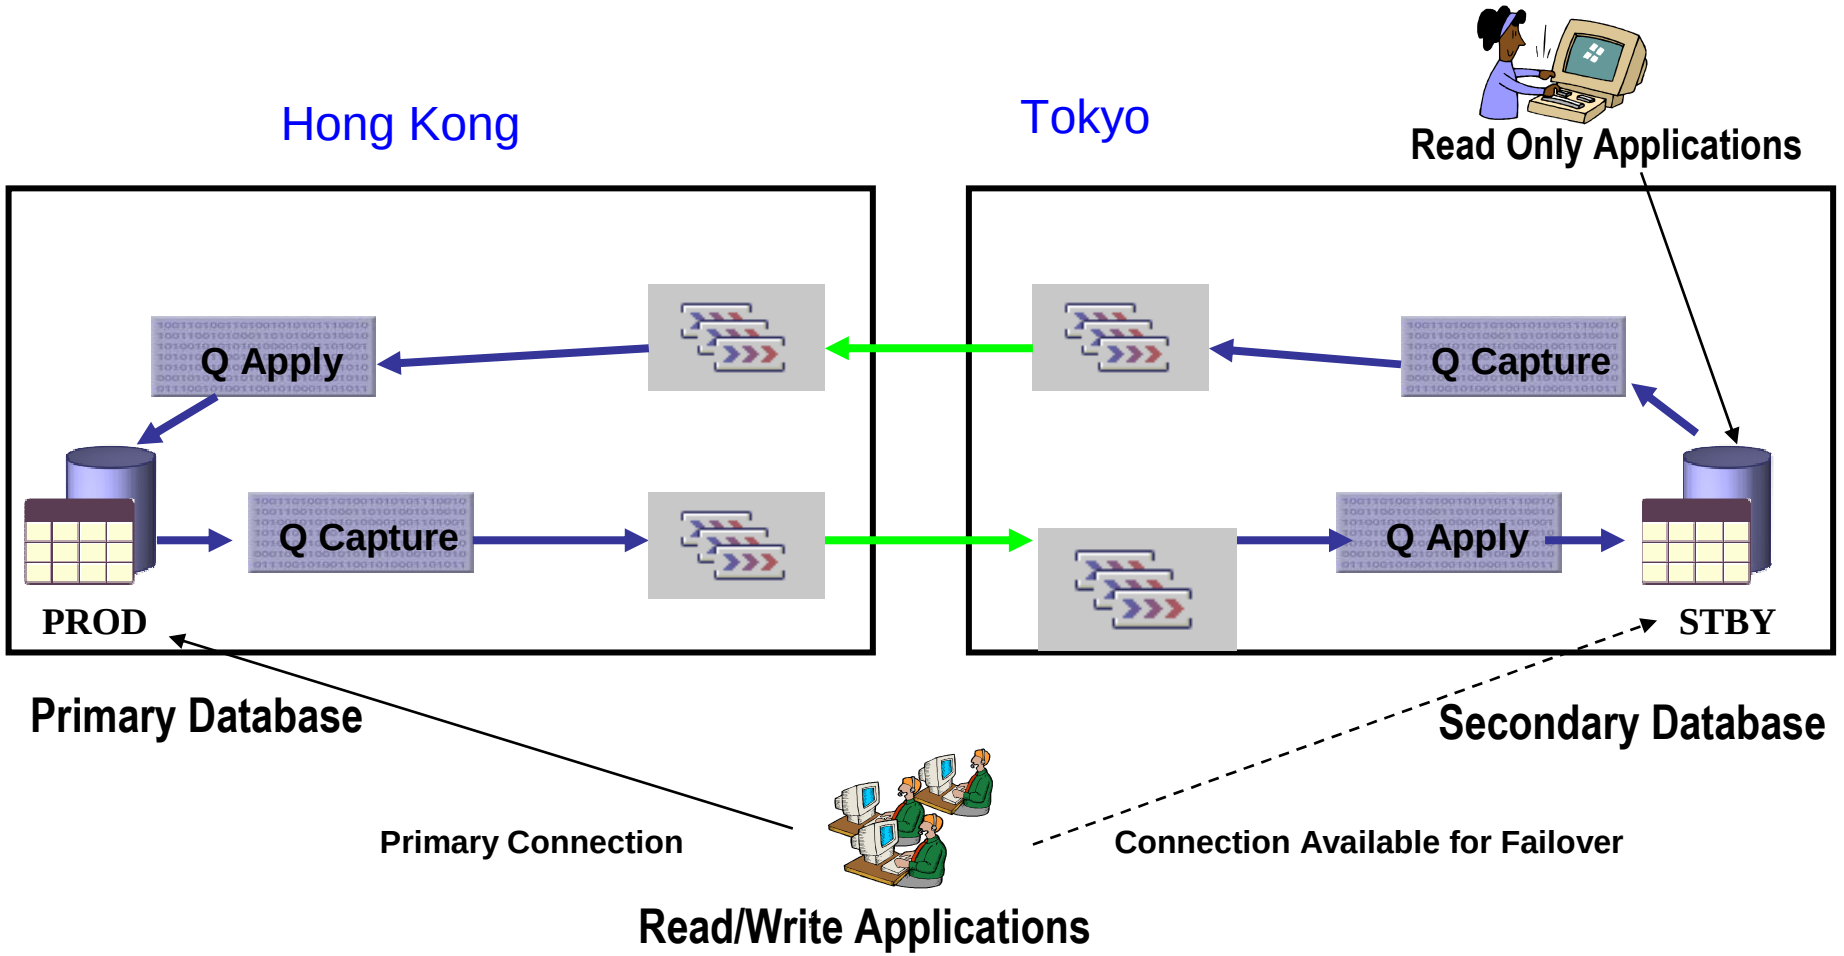

Q Replication provides a solution for continuous availability where the active secondary system is also available for other applications

Pros & Cons

Pros:

- Protected from software, server, storage, site failure
- Failover time is “instant”
- Standby can be full or subset and is fully accessible (read and/or write)
- Multiple standby servers

Cons:

- More complex to setup and monitor (but more flexibility) vs. HADR
- Asynchronous

4) Scale out database by using Database Partitioning

- (not only Range/Table Partitioning)

A single database server having **4T** data

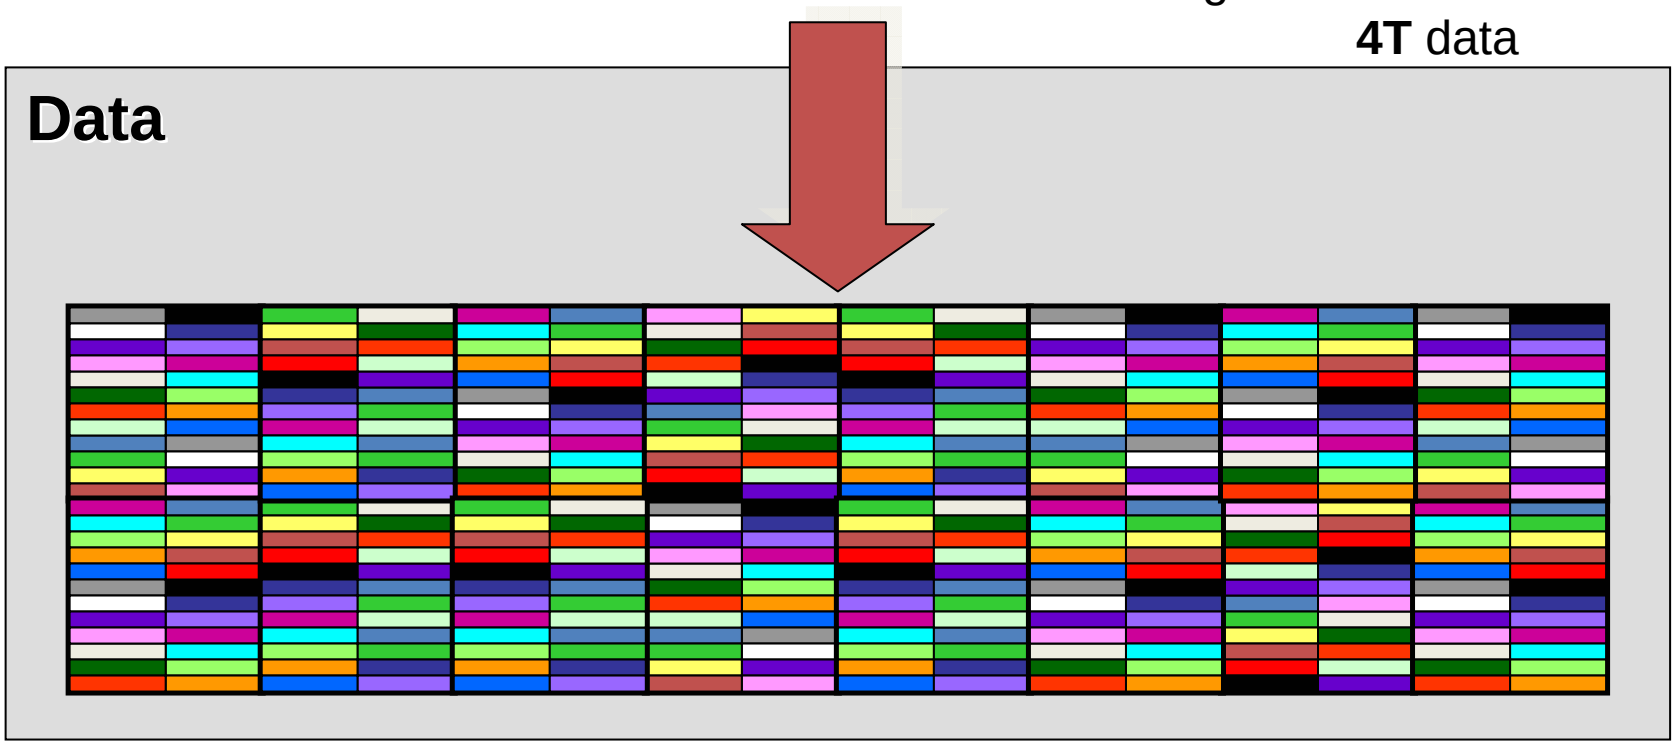

Partition database across servers deliver linear performance scale out

4 database servers having **1T** data each deliver 4 times performance
Application still seeing it as 1 single database

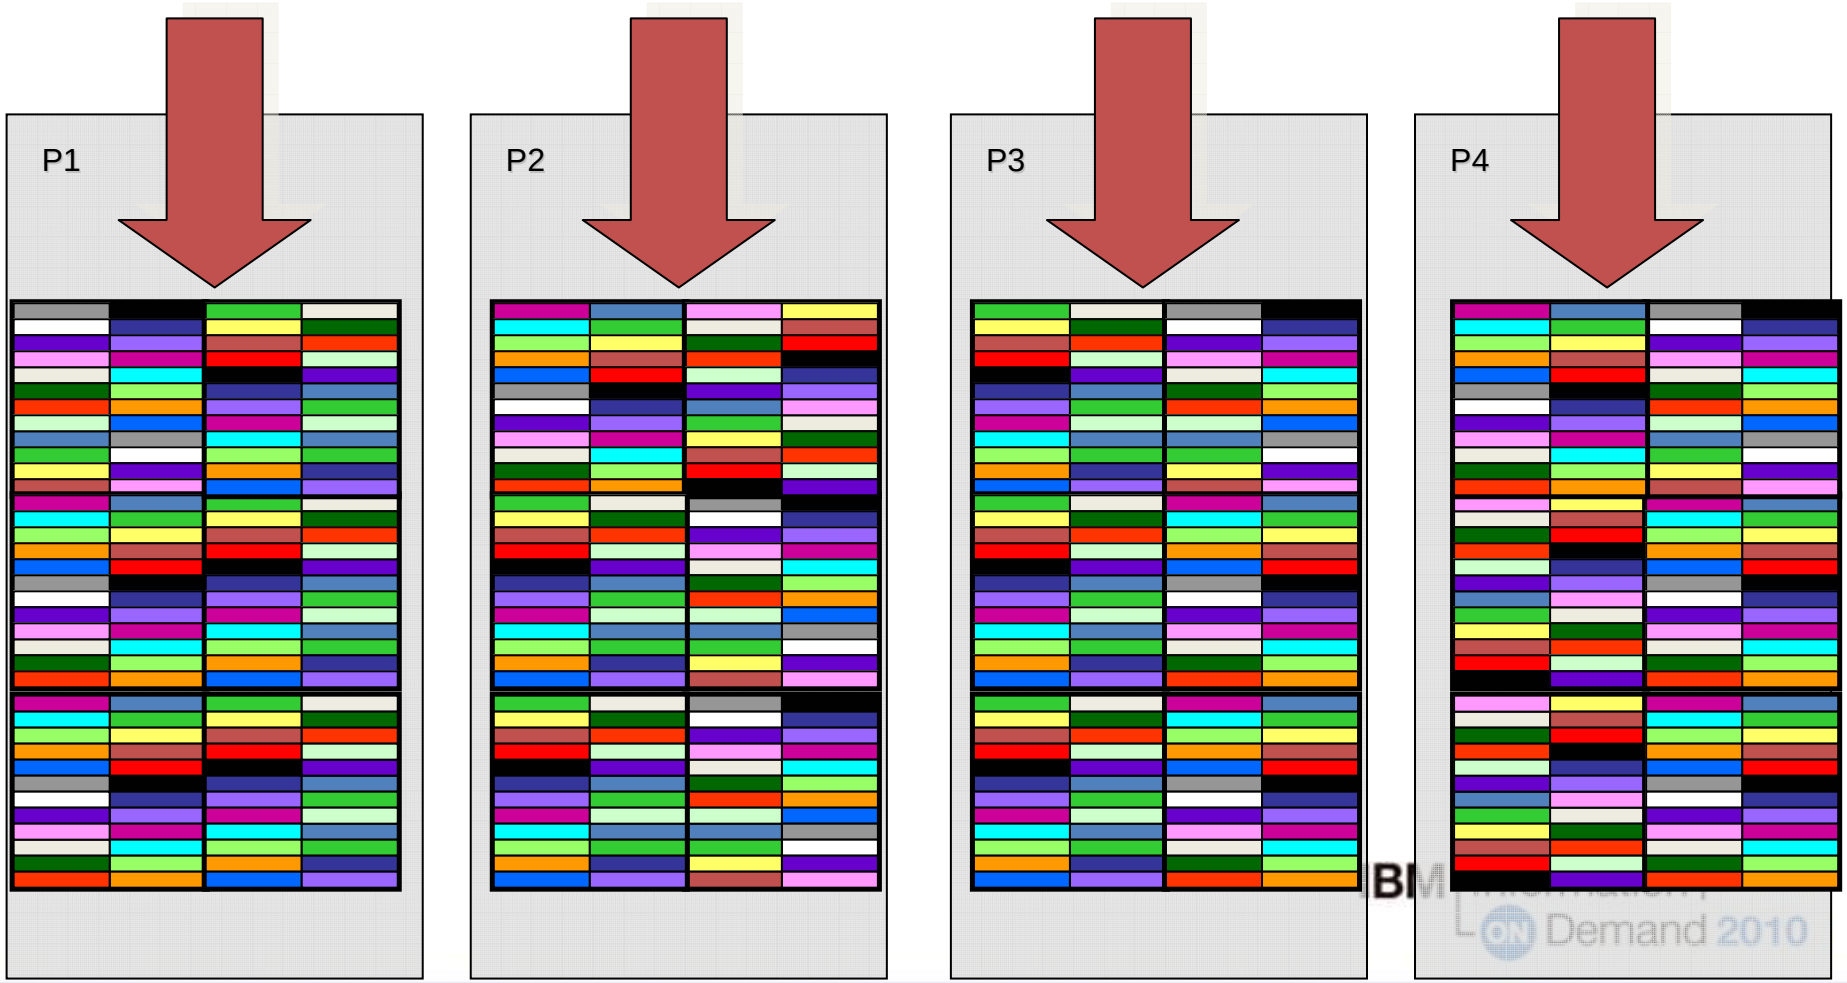

Pros & Cons

Pros:

- A scalable design for Data Warehouse (OLAP) type of database
- Share Nothing – no I/O bottleneck when scale out
- No performance degrade with linear scale out design

Cons:

- Utilize server base failover design only (passive standby)

Introduce DB2 pureScale

- **Unlimited Capacity**
  - Buy only what you need, add capacity as your needs grow
- **Application Transparency**
  - Avoid the risk and cost of application changes
- **Continuous Availability**
  - Deliver uninterrupted access to your data with consistent performance

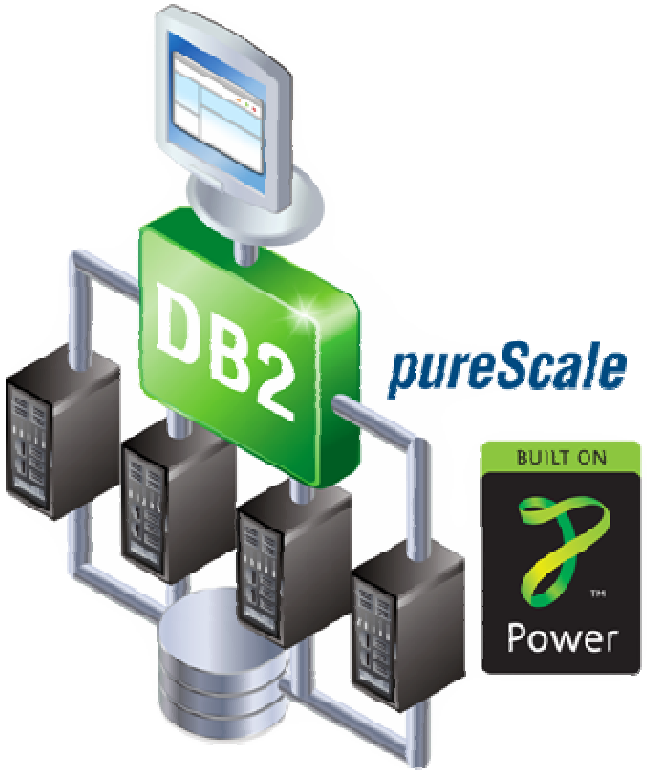

DB2 pureScale Architecture

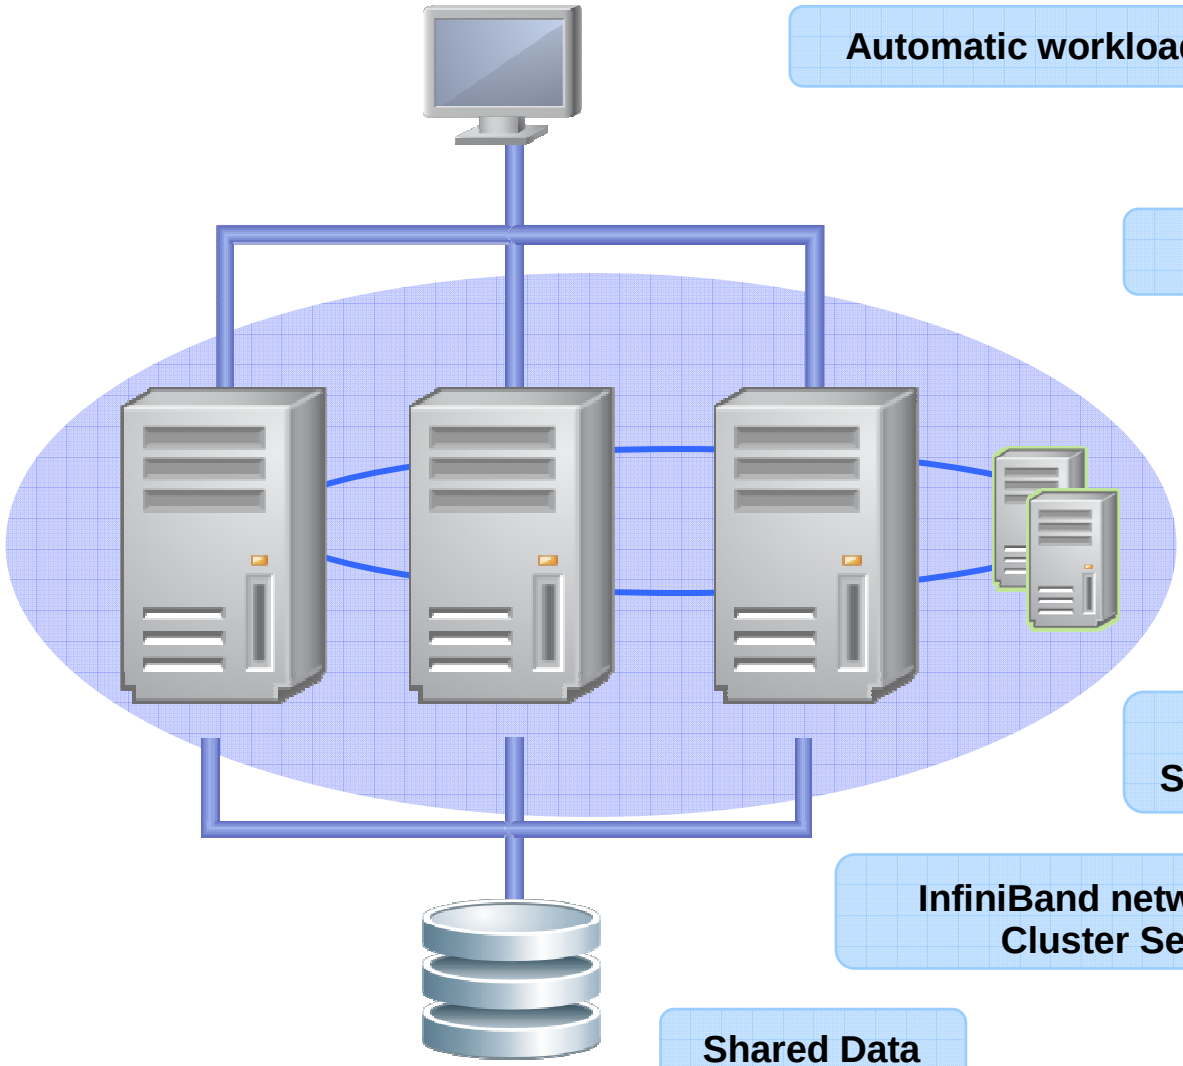

Automatic workload balancing

Cluster of DB2 members

Leverages the global lock and memory manager technology from System z

Integrated Tivoli System Automation

InfiniBand network & DB2 Cluster Services

Shared Data

Unlimited Capacity

- DB2 pureScale has been designed to grow to whatever capacity your business requires

- Flexible licensing designed for minimizing costs of peak times

Issue:
All year, except for two days, the system requires 3 servers of capacity. But for those 2 days it needs 4 servers

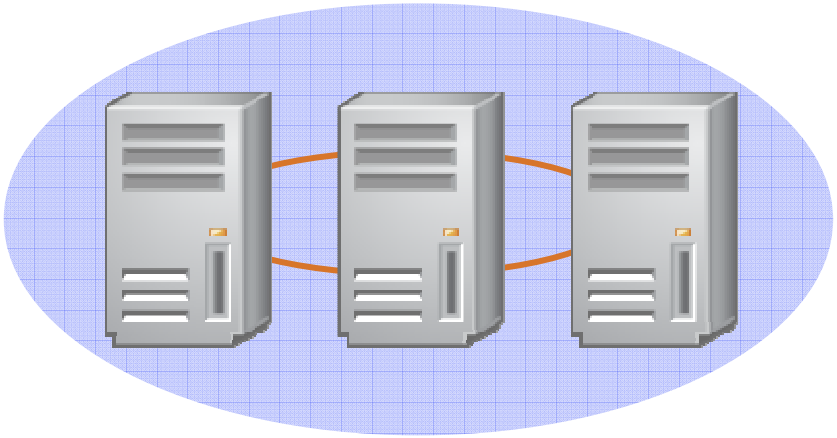

Unlimited Capacity

- DB2 pureScale has been designed to grow to whatever capacity your business requires

- Flexible licensing designed for minimizing costs of peak times

Solution:
Use DB2 pureScale and add another server for those two days, and only pay sw license fees for the days you use it.

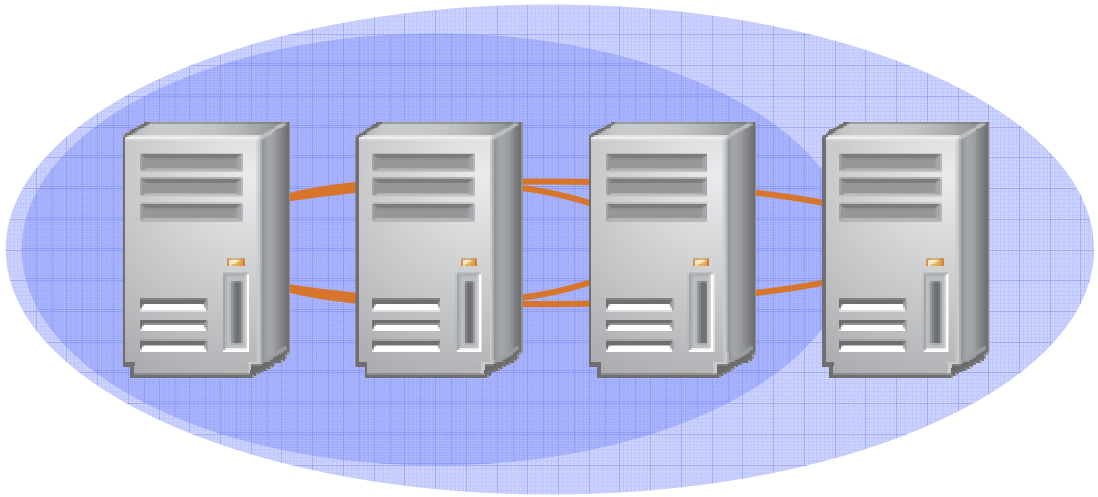

Unlimited Capacity

- DB2 pureScale has been designed to grow to whatever capacity your business requires
- Flexible licensing designed for minimizing costs of peak times
- Only pay for additional capacity when you use it even if for only a single day

Need more… just deploy another server and then turn off DB2 when you're done.

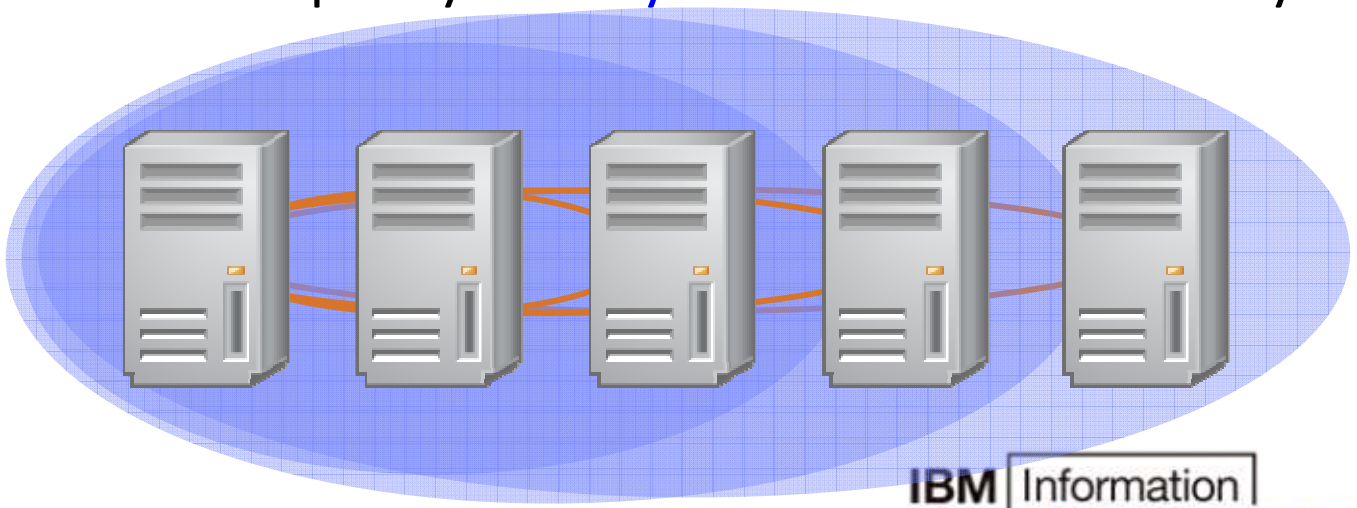

Scalability Results

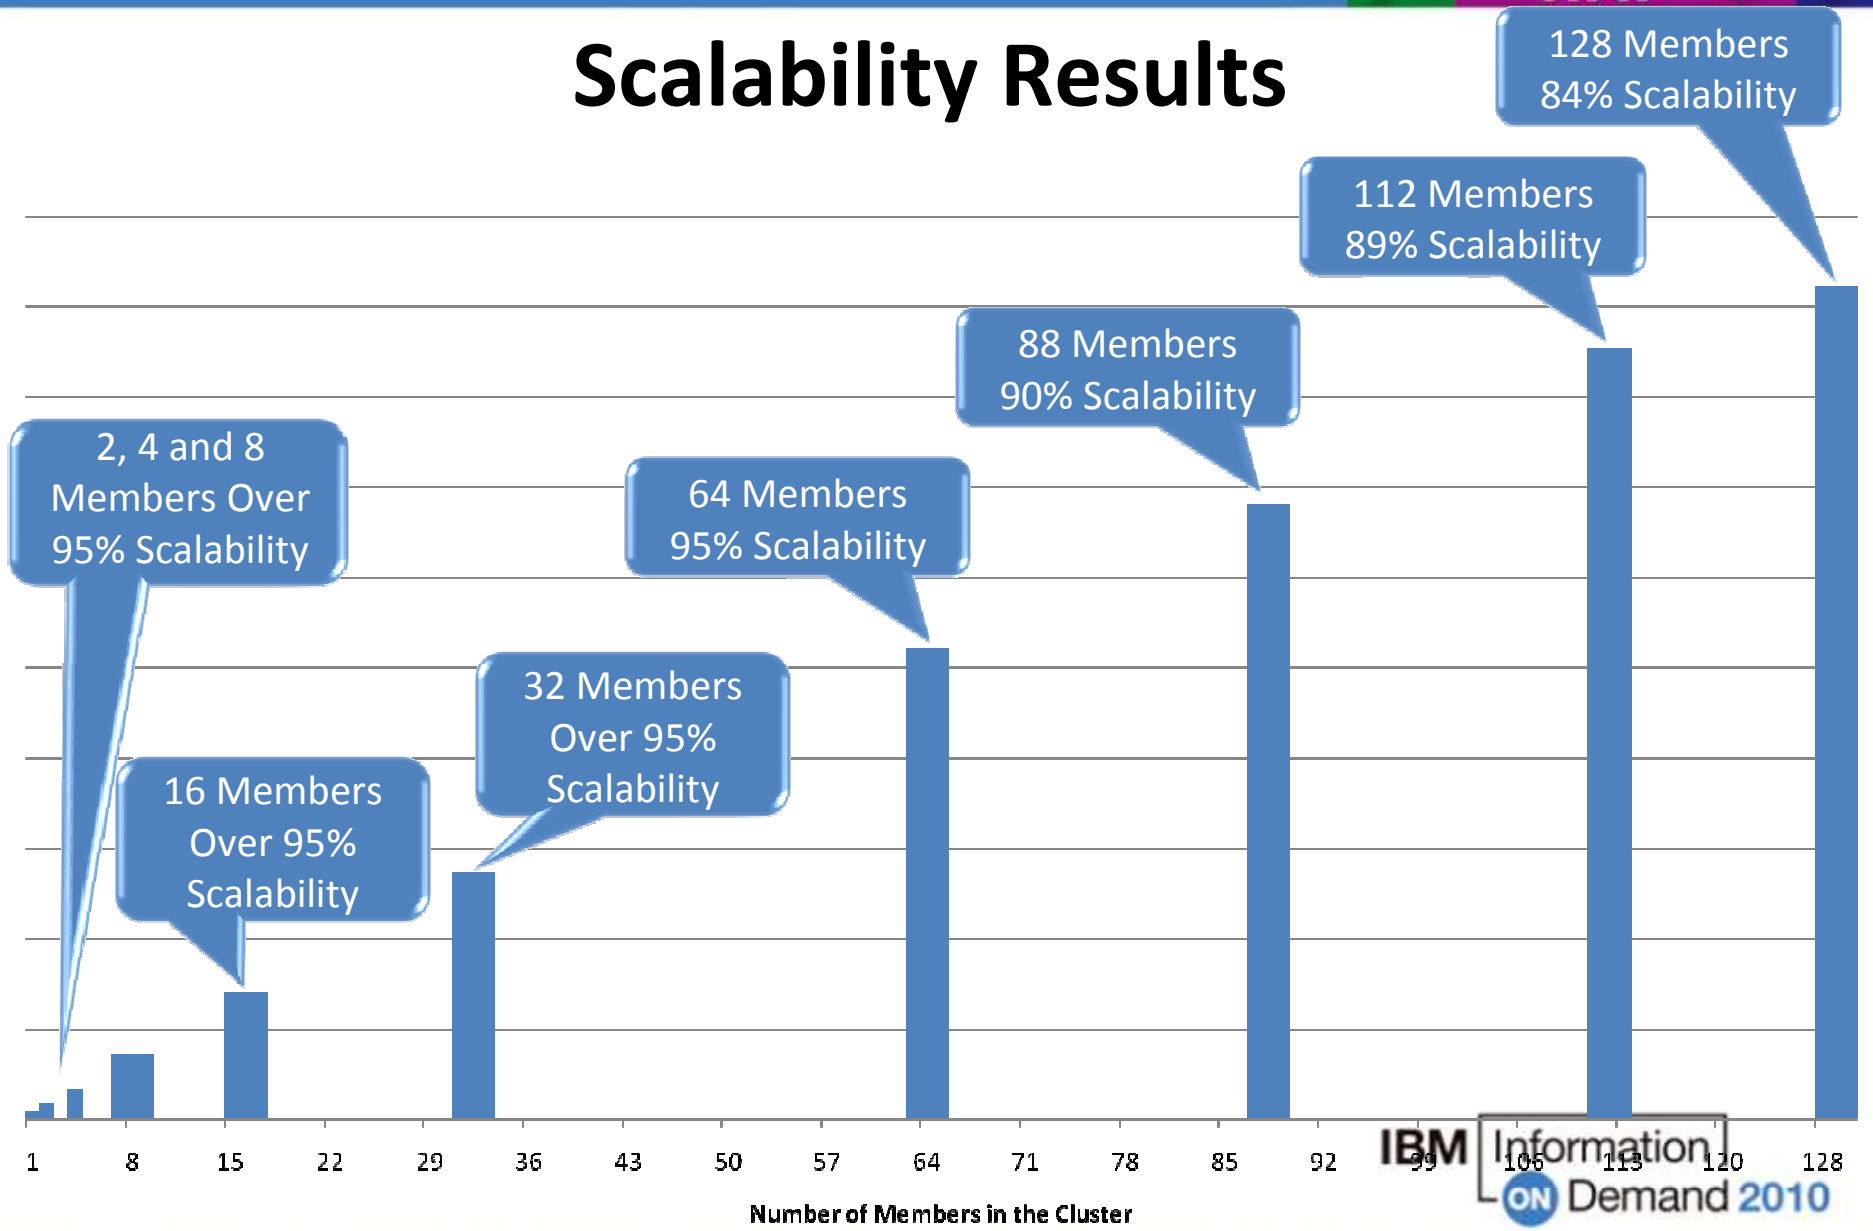

Application Transparency

Take advantage of extra capacity instantly

- No need to modify your application code
- No need to tune your database infrastructure

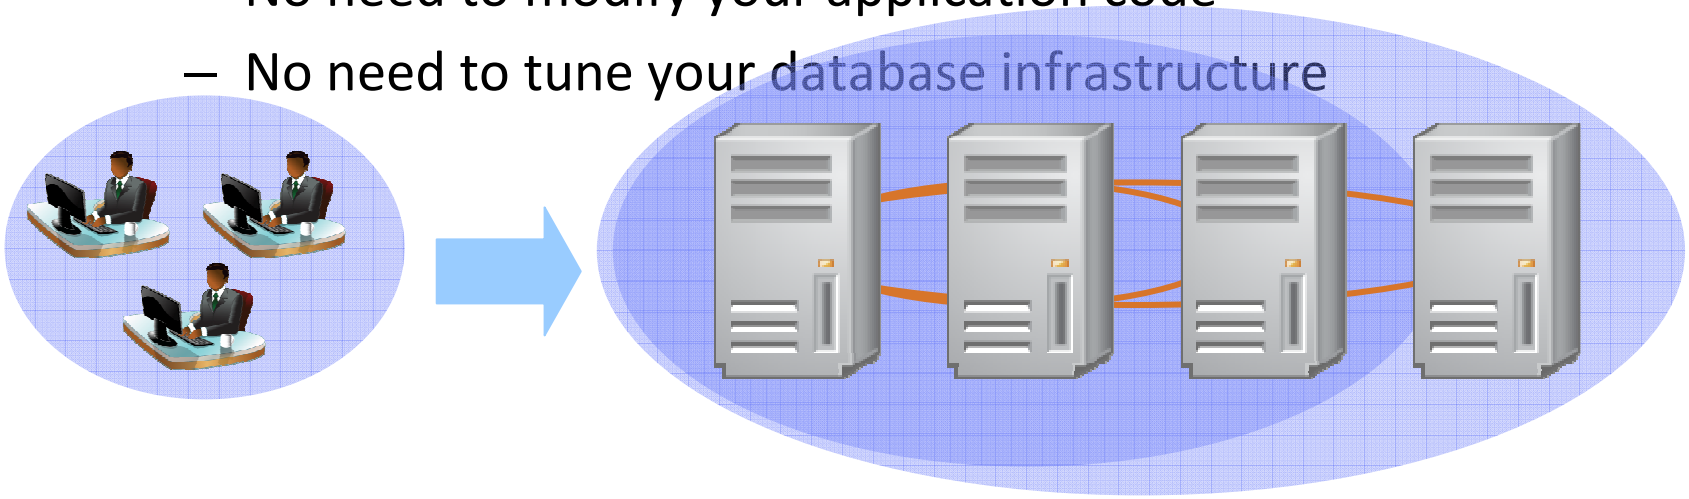

Your DBAs can add capacity without re-tuning or re-testing

Your developers don't even need to know more nodes are being added

DB2 pureScale is Easy to Deploy

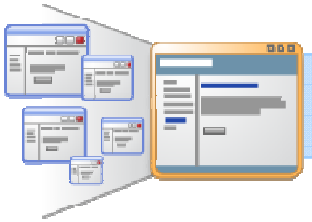

Single installation for all components

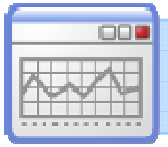

Monitoring integrated into OPTIM tool set

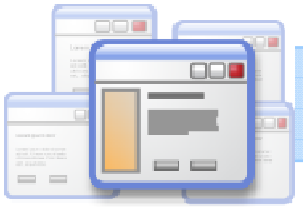

Single installation for fixpaks and updates

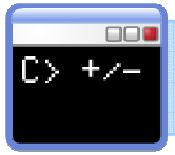

Simple command to add and remove members

Continuous Availability

- Protect from infrastructure outages
  - Handles multiple concurrent failures and keep on going, and going, and going…

Automatic workload balancing

Duplexed secondary global lock and memory manager

Tivoli System Automation automatically handles all component failures

DB2 pureScale stays up even with multiple node failures

Shared disk failure handled using disk replication technology

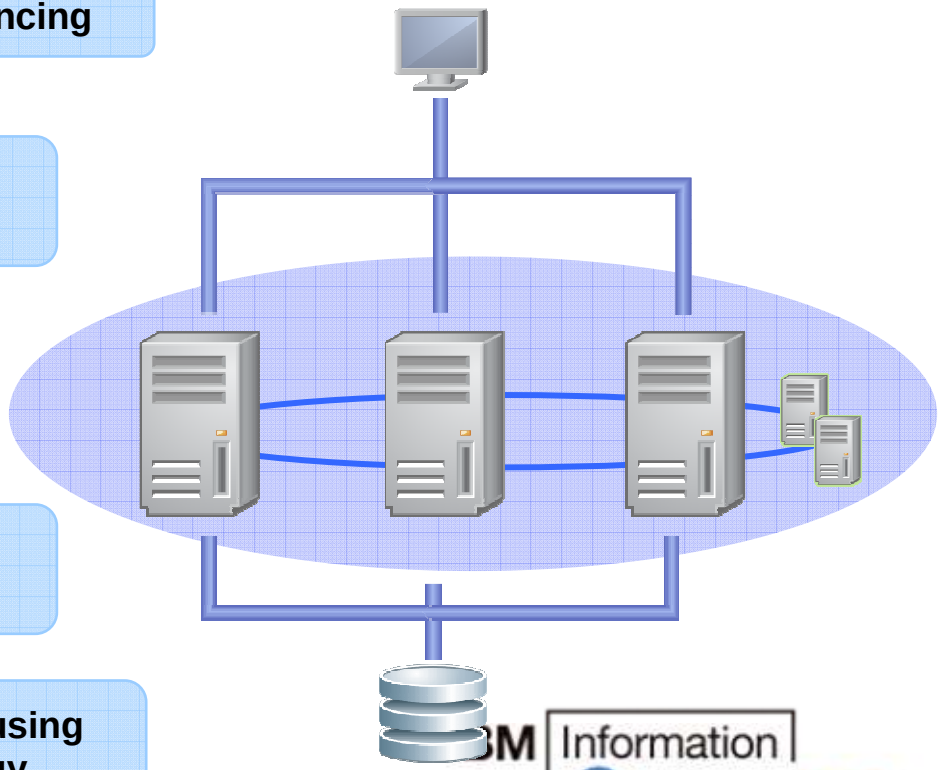

Recover Instantaneously From Node Failure

- Protect from infrastructure related outages
  - Redistribute workload to surviving nodes immediately
  - Completely redundant architecture
  - Recover in-flight transactions on failing node immediately, and automatically

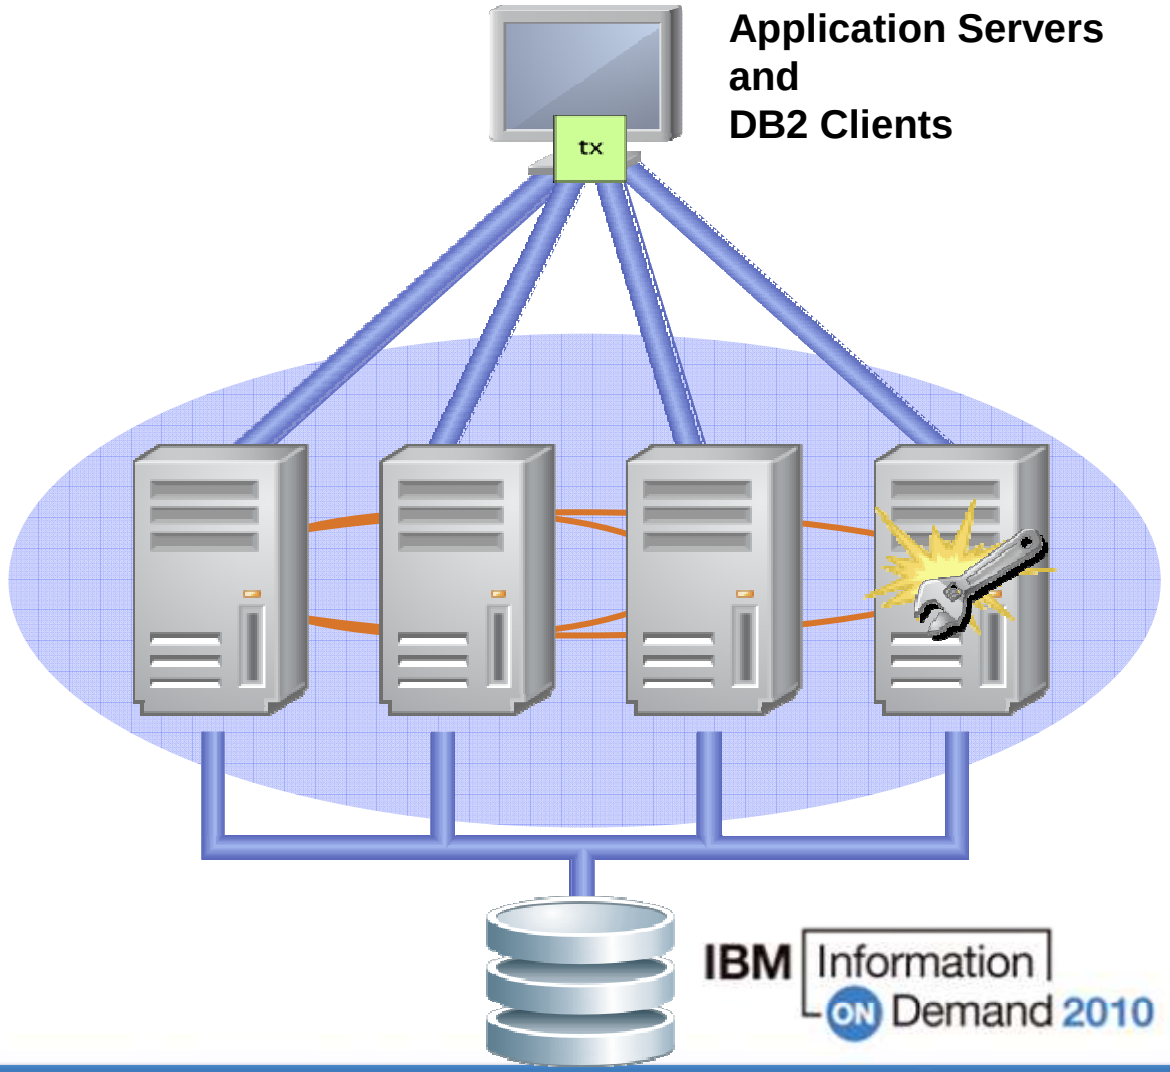

Minimize the Impact of Planned Outages

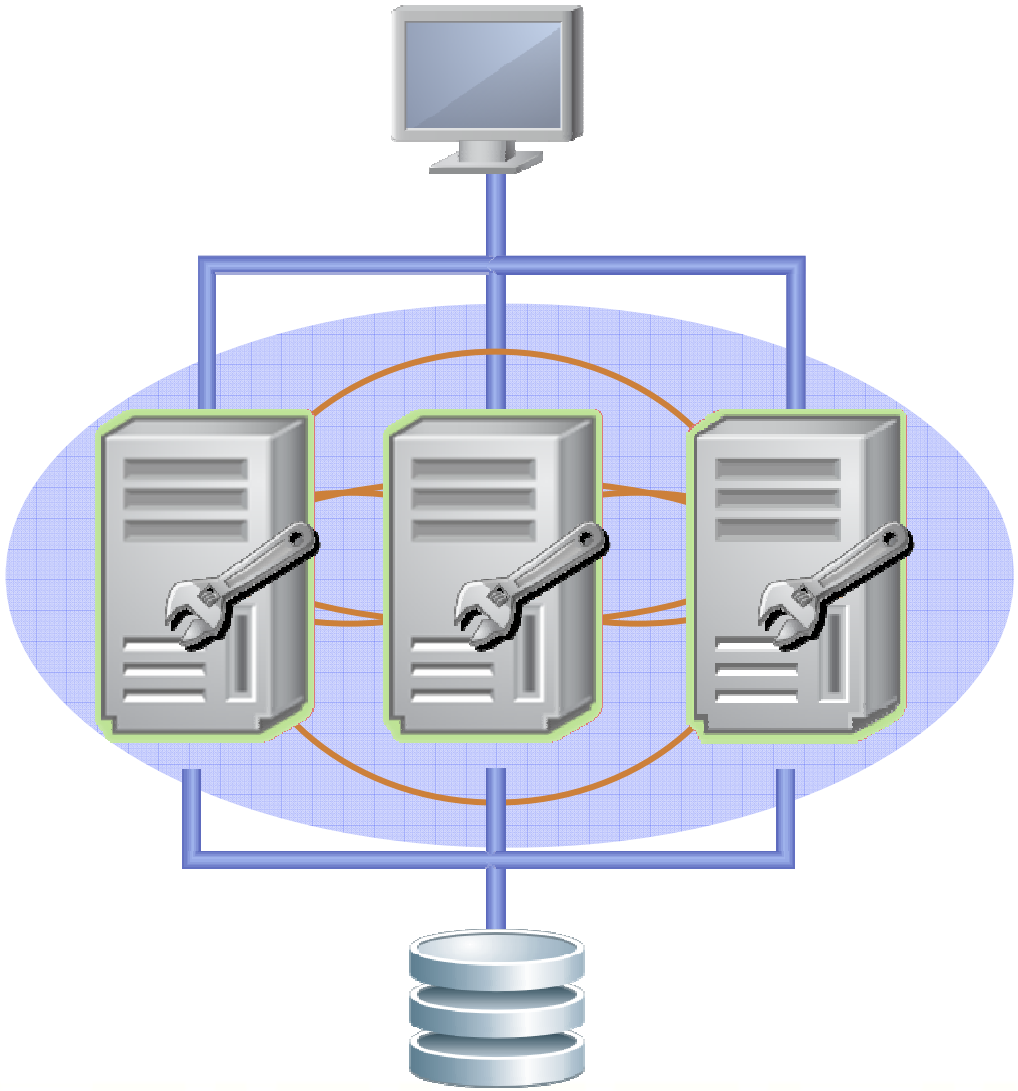

- Keep your system up
  - During OS fixes
  - HW updates
  - Administration

Summary

- There are many different options for database failover and scale out design

- No one option is perfect for all scenarios
  - Now we can said that because we have all the options available and we have no prejudice on any option

Thank You!