Experience the GRID Today

with Oracle*9i* RAC

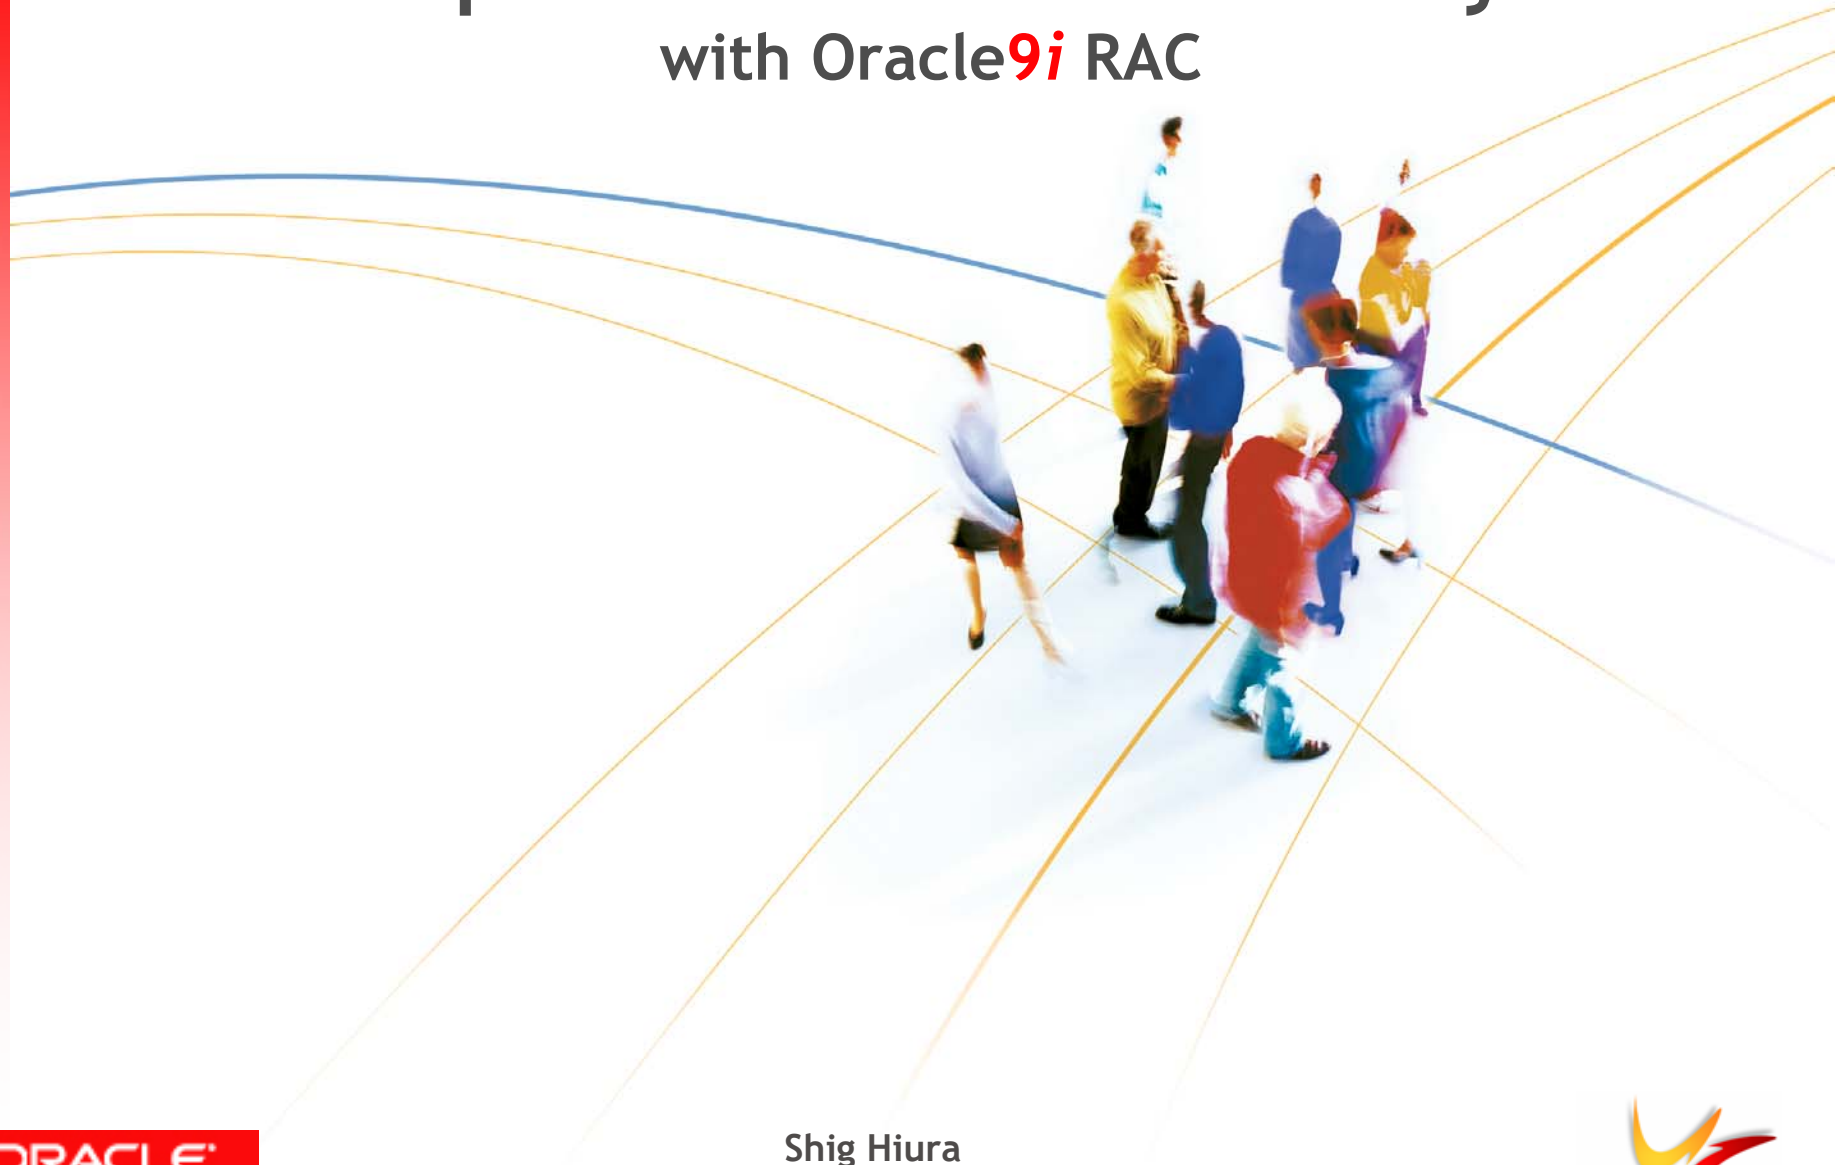

Shig Hiura
Pre-Sales Engineer
Shig_Hiura@etagon.com

ORACLE PartnerNetwork MEMBER PARTNER

ETAGON

Agenda

- Introduction
  - What is the Grid
  - The Database Grid
- Oracle9*i* RAC Technology
- 10g vs. 9*i*R2 Comparison
- Benefits of RAC on Linux
- Summary

ETAGON

The best thing about the Grid is that it is unstoppable.

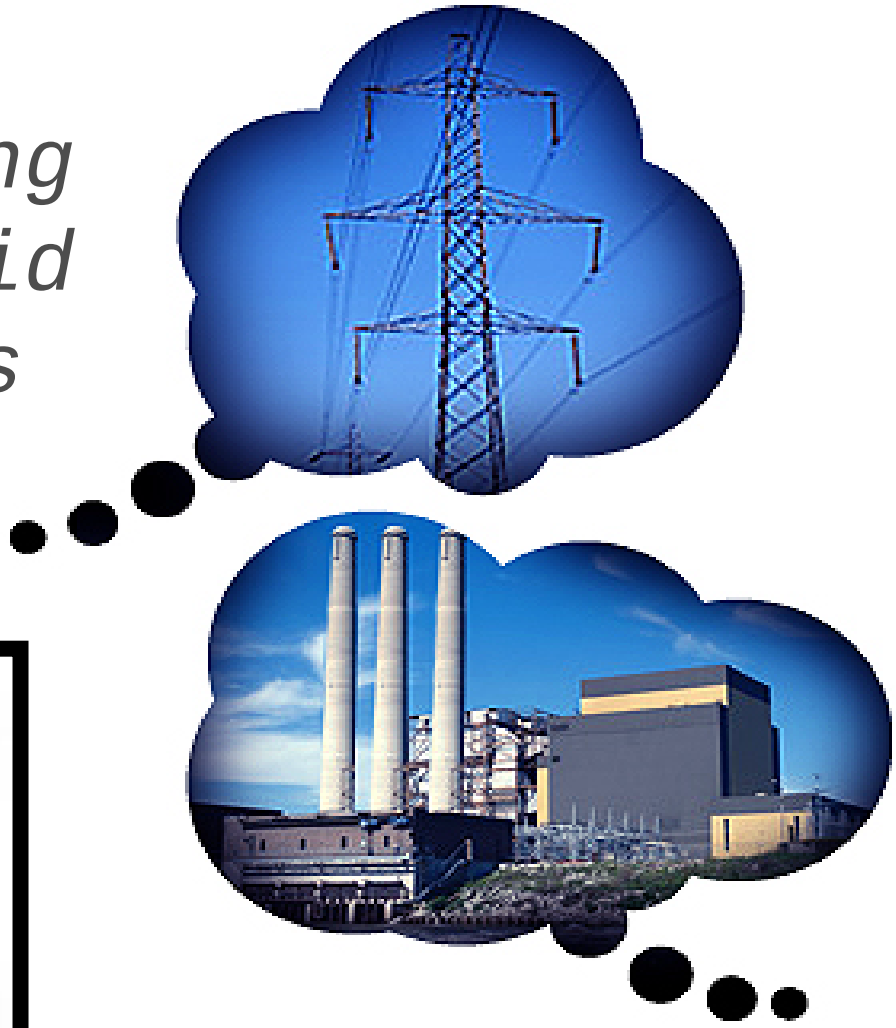

The Economist, June 21, 2001

ETAGON

Key IT Imperatives

- Highest Quality of Service
  - Maximum Availability
  - Maximum Scalability
- Most Efficient Management
  - Do much more …
- Lowest Cost
  - … with much less

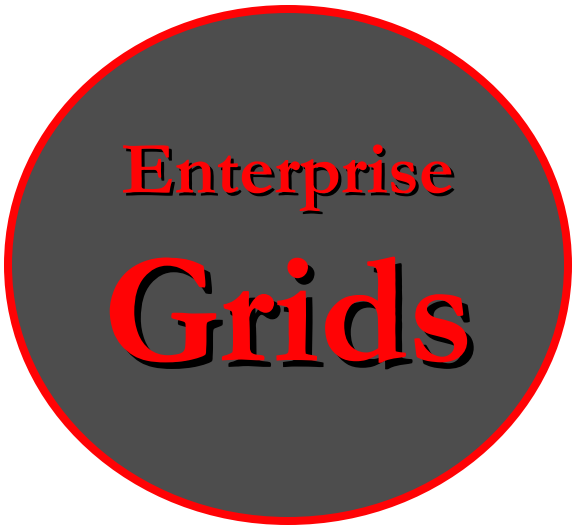

ETAGON

Grid Vision

- Computing as a utility
  - A network of clients and service providers
- Client-side: Simplicity
  - Request computation or information and receive it
- Server-side: Sophistication
  - Availability, reliability, security
  - Capacity on demand, load balancing
- Virtualization
  - Storage
  - Networks
  - Computing Resources

Benefits of the Grid

- Better information faster
  - Perform more work with fewer resources
  - Spread work across resources
  - Access to resources on demand
- Faster response to changing business priorities
  - Instantly and dynamically realign IT resources as business needs change
- Reduced IT costs
  - Improve utilization of existing resources
  - Utilize less expensive commodity platforms

ETAGON

15 December, 2003

Reducing IT Costs

Large Dedicated Server

Oracle Grid Computing

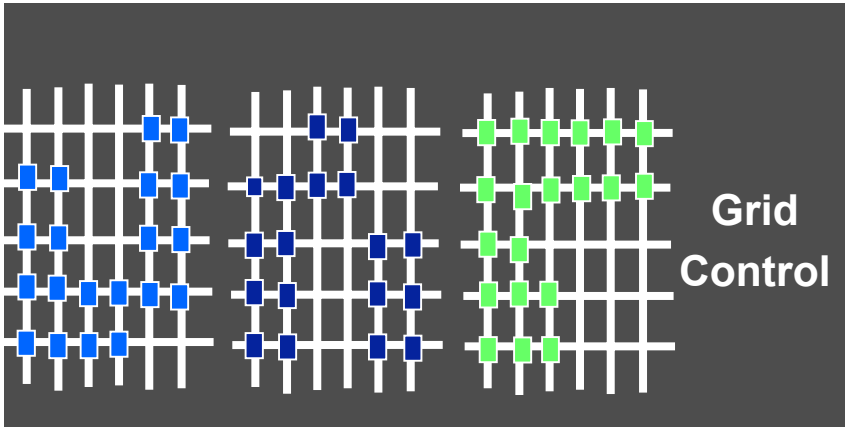

- Expensive components
- High incremental costs
- Configured for peak
- Single point of failure
- Enterprise service at high cost

- Low cost modular components
- Low incremental costs
- Capacity on demand
- Fault tolerant
- Enterprise service at low cost

ETAGON

Technology Trends

- Blades: Every vendor offering them
  - Huge cost advantages
  - Software vendors have to enable usage
  - IBM BladeCenter, HP Proliant BL-Series, etc.
- Linux: Fastest growing OS
  - Commodity OS
  - Ready for blades today
  - Linux and blades naturally complement each other

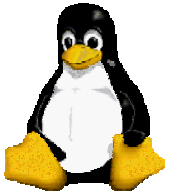

- NAS, SAN, and Infiniband provide storage access from any blade

ETAGON

Roadmap

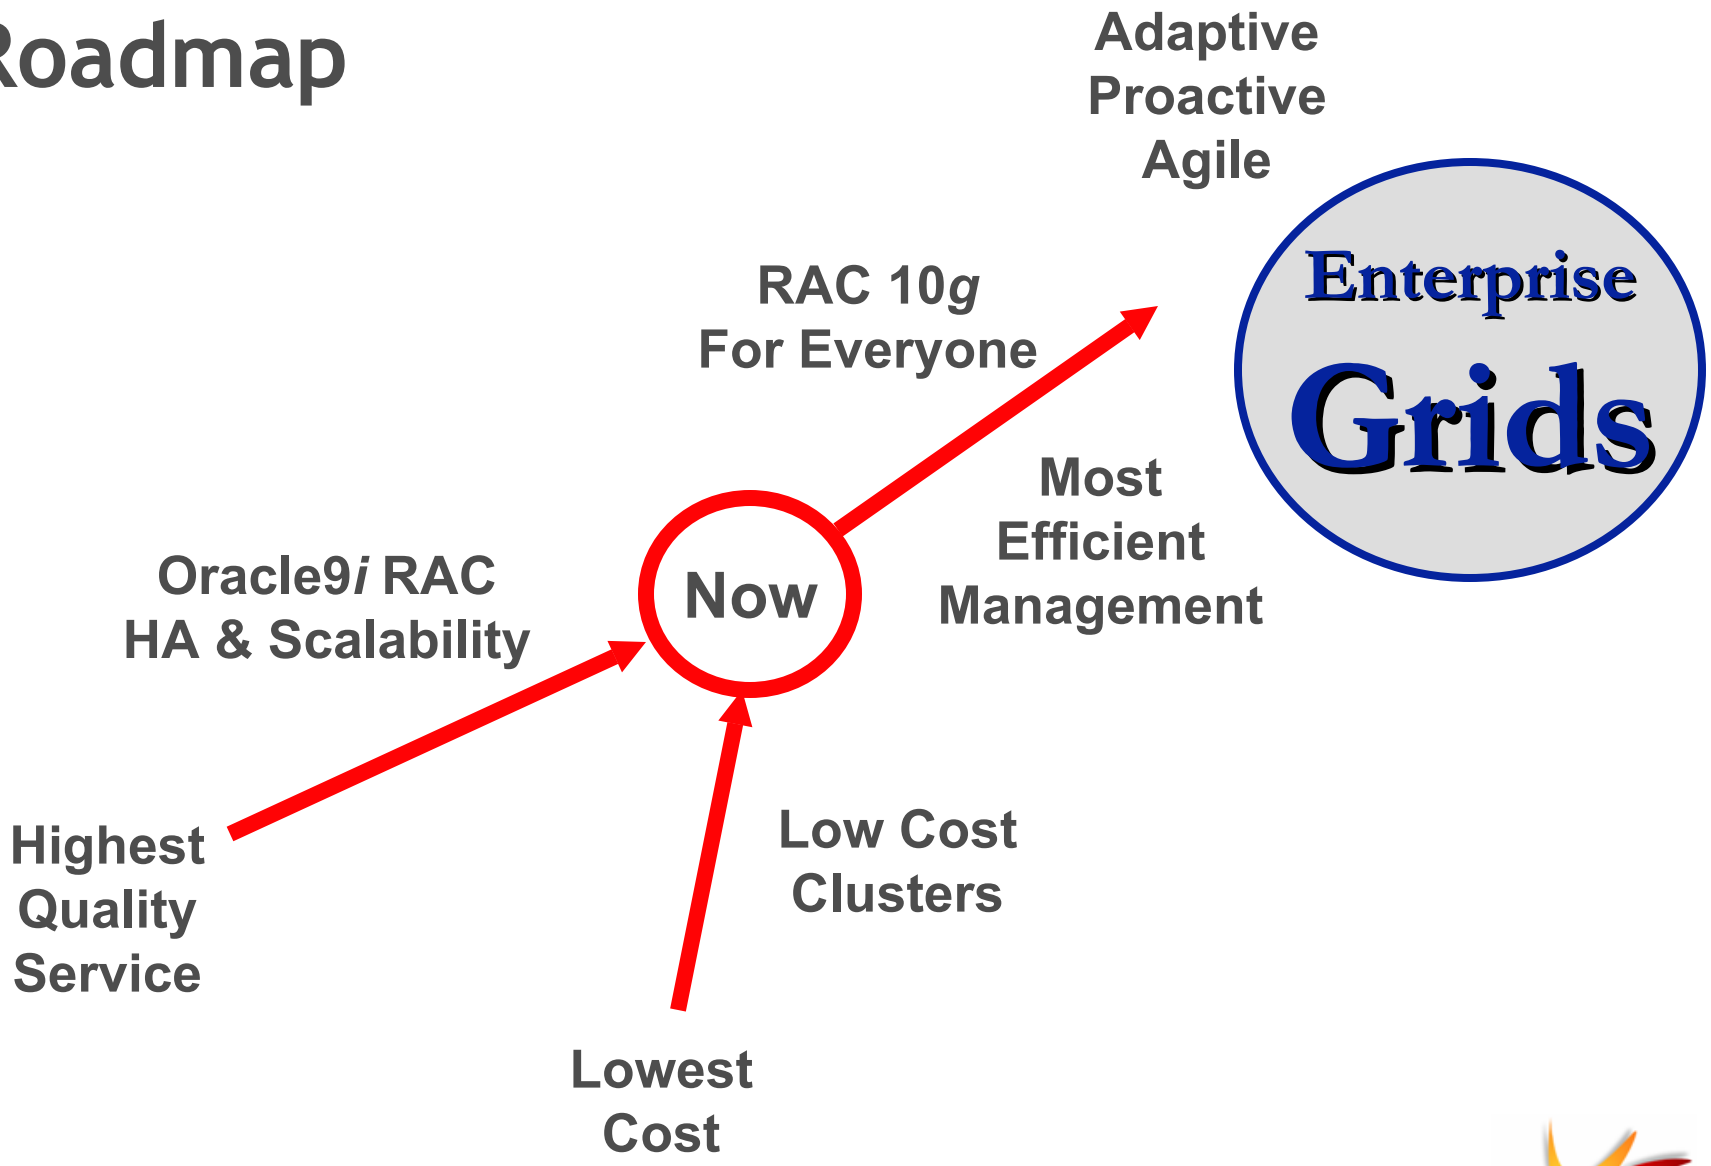

ETAGON

Grid Computing Eliminates Islands of Computation

IT Problem		Grid Solution
Separate servers	→	Shared servers
High h/w & s/w costs	→	Low cost components
Configured for peak	→	Capacity on demand
Single point of failure	→	Fault tolerant
Rigid	→	Flexible
Difficult to manage	→	Unified management

Coordinated use of many small servers acting as one large computer.

ETAGON

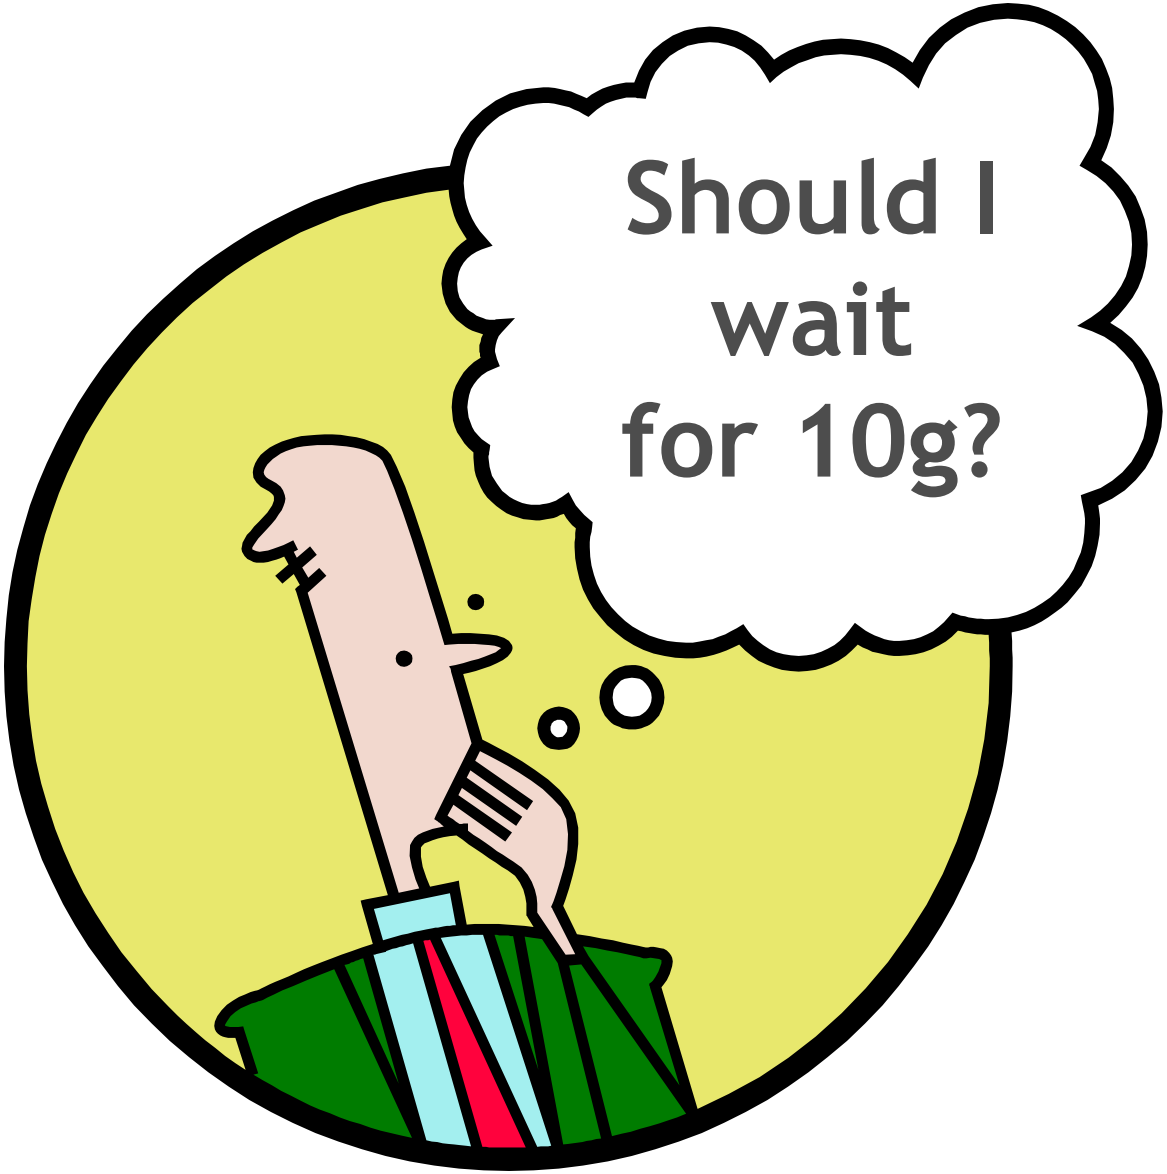

ETAGON

An Oracle9*i* RAC Database

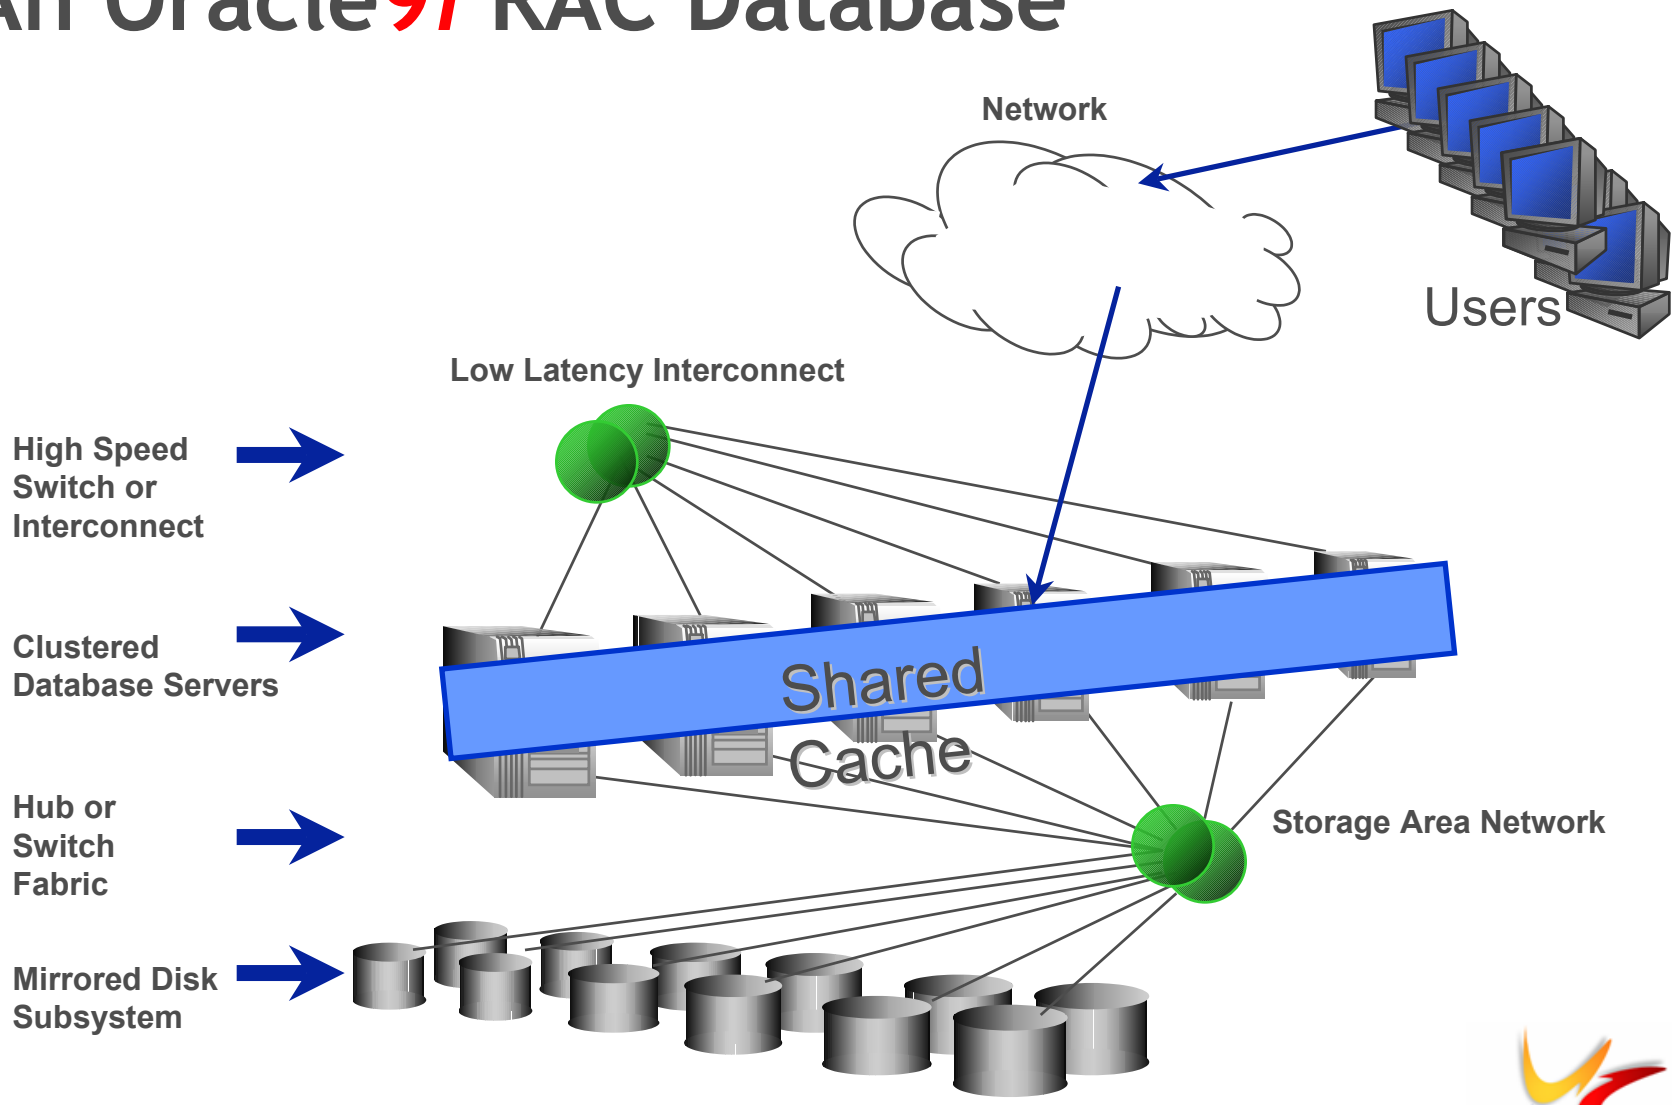

ETAGON

Why Clustering?
Scalability Alternatives

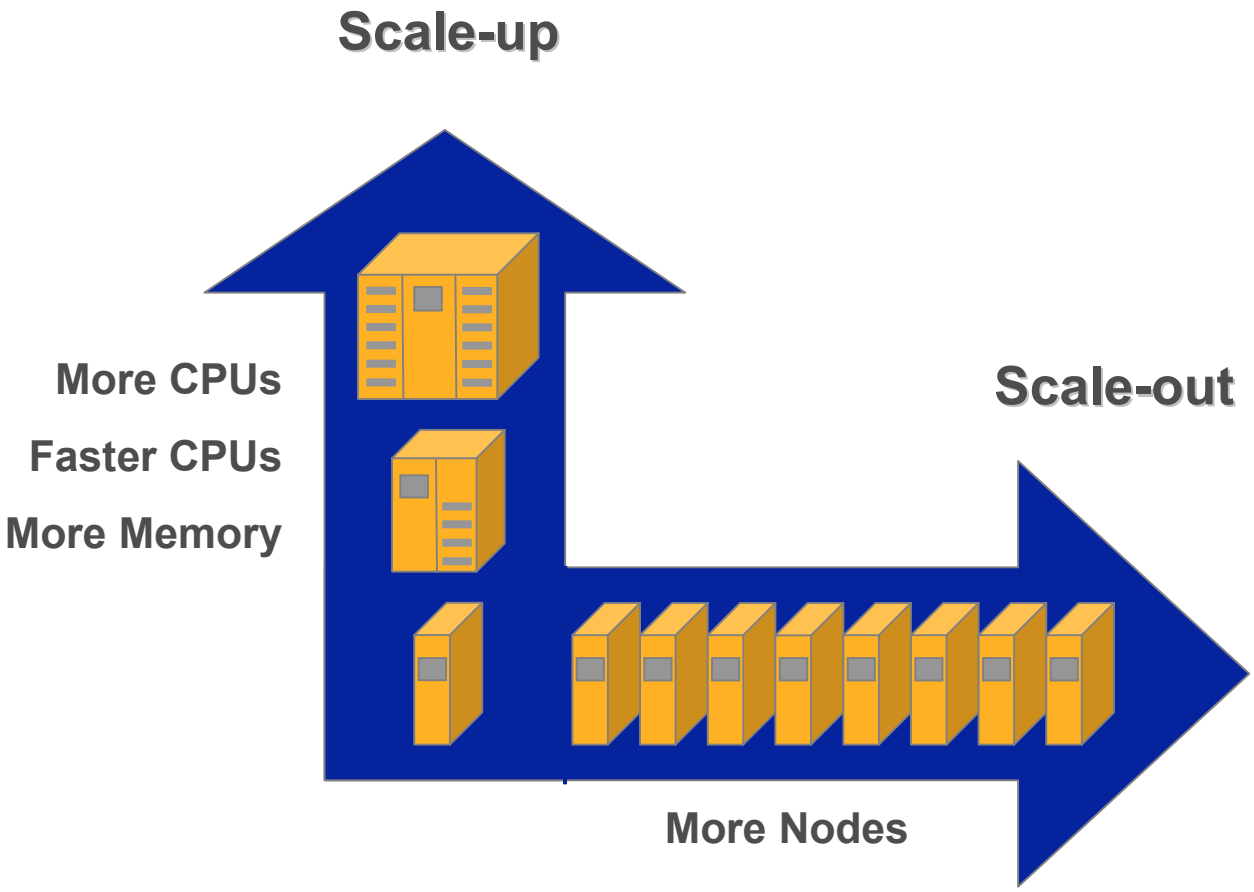

ETAGON

Clusters Reduce Costs

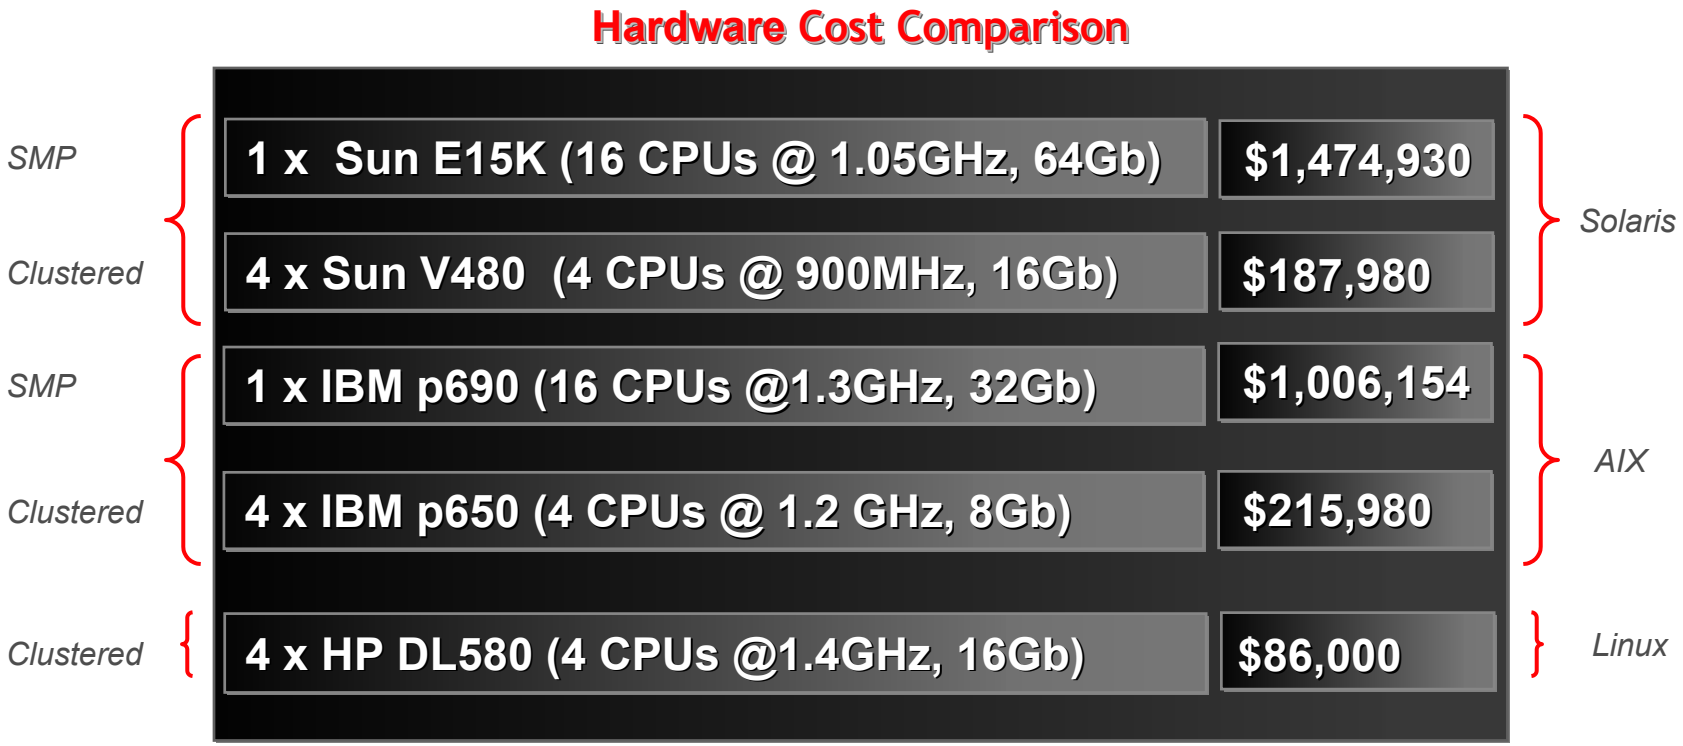

*Source OEM Website - Mid Range SMP vs Clustered hardware

ETAGON

Oracle9*i* RAC Distributions

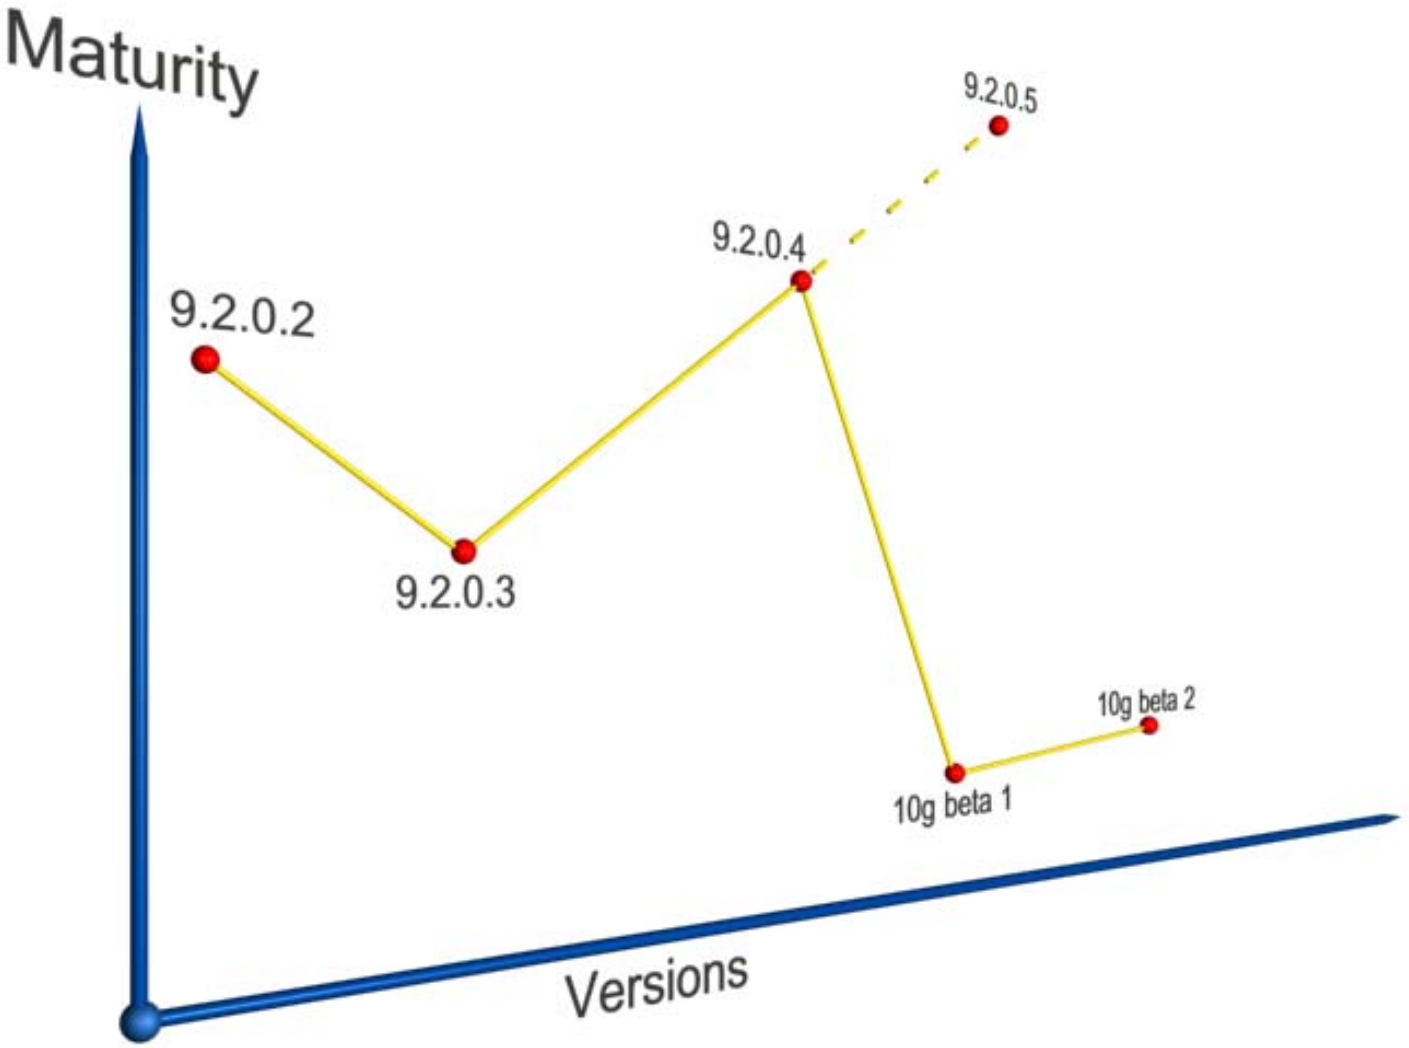

ETAGON

Oracle9*i* RAC Distributions

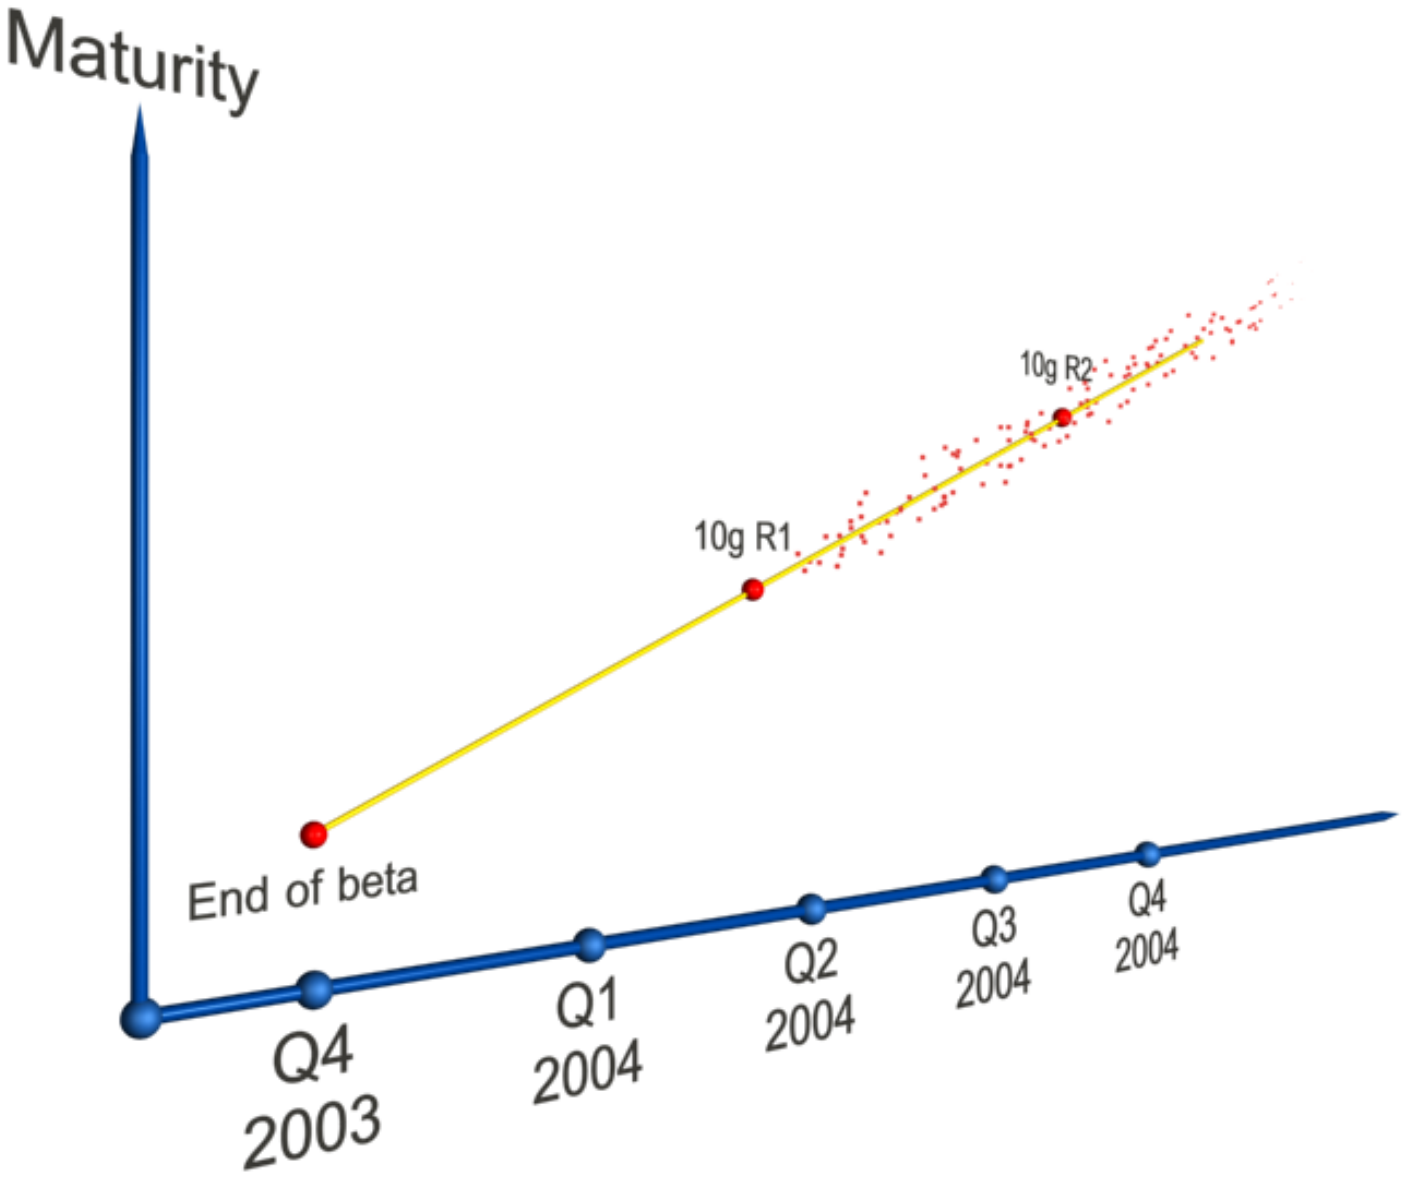

ETAGON

Oracle9*i* RAC HA

Failover Operation	RAC	'Cold'
Reconfigure Group Membership	15 sec	0 sec
Reconfigure Distributed Locks	5 sec	0 sec
Failover Disk Volumes	0 sec	Up to 20 min
Restart Oracle	0 sec	Up to 5 min
Recover Oracle	20 sec	20 sec
Warm Buffer Cache	0 sec	5 + min
Total Failover Time	< 60 sec*	> 30 min

* Oracle RAC 10g failover time < 8 sec

ETAGON

Oracle9*i* Cluster Stack vs. Oracle 10*g*

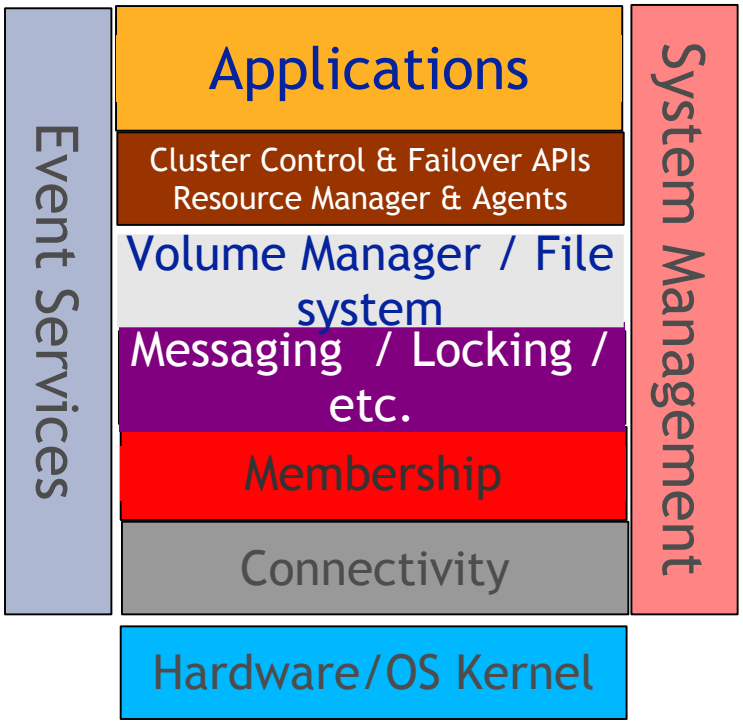

ETAGON

10g moving to an Integrated clusterware

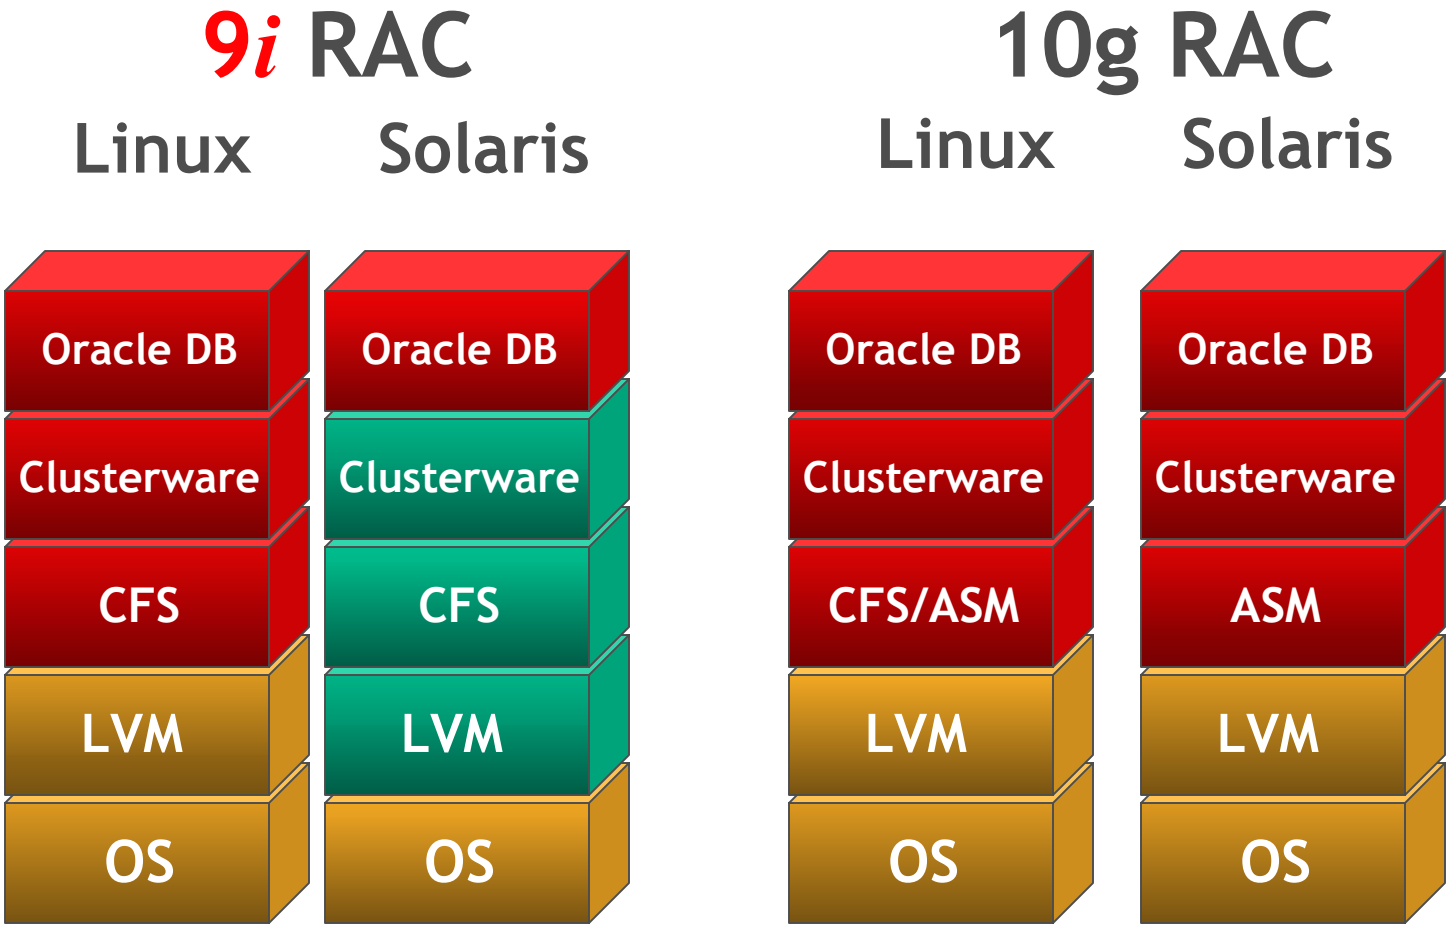

ETAGON

Automatic Storage Management (ASM)

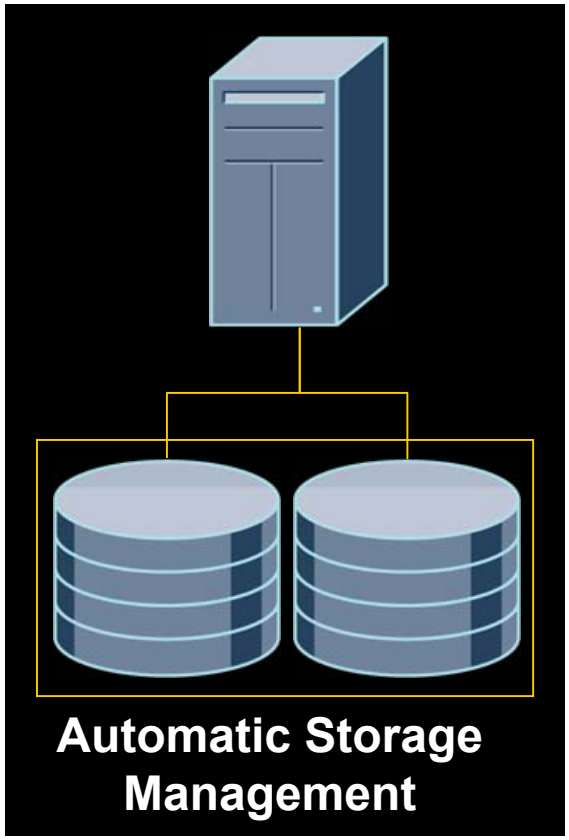

- Eliminates need for conventional file system and volume manager
- Capacity on demand
  - Add/drop disks online
- Automatic I/O load balancing
  - Stripes data across disks to balance load
  - Best I/O throughput
- Automatic mirroring
- Virtualizes the storage resources

ETAGON

ASM – How it Works

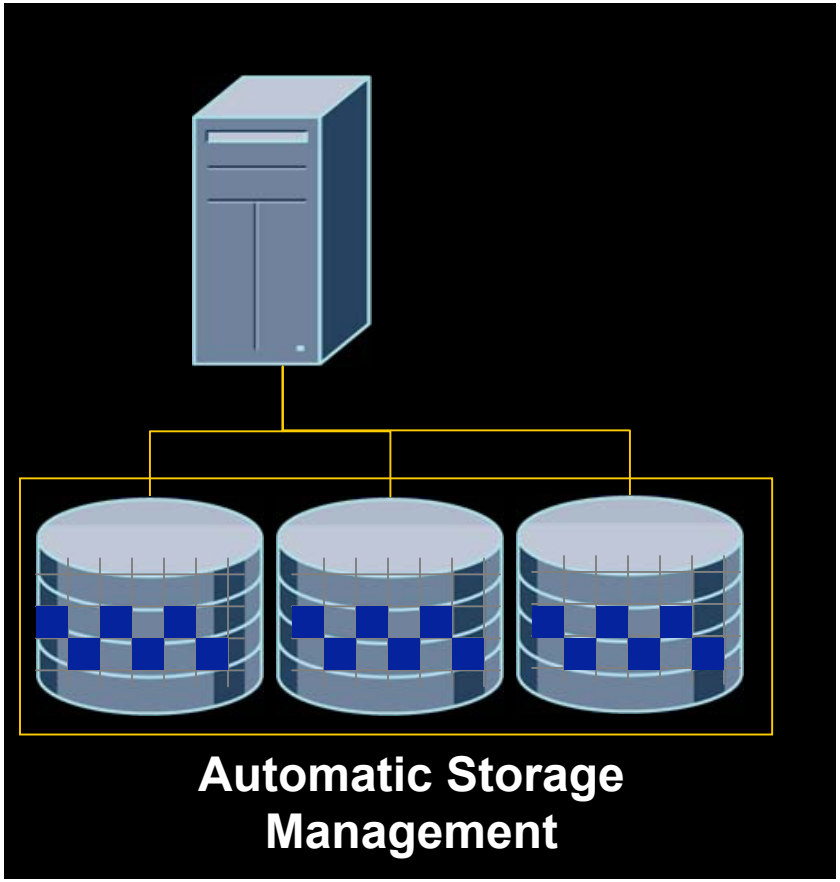

- No volumes: just a pool of storage.
- Partitions total disk space into uniform 1 megabyte extents.
- Online add/remove of disk with automatic load balancing.

ASM Architecture

RAC or Non–RAC Databases

Clustered Servers

Clustered Pool of Storage

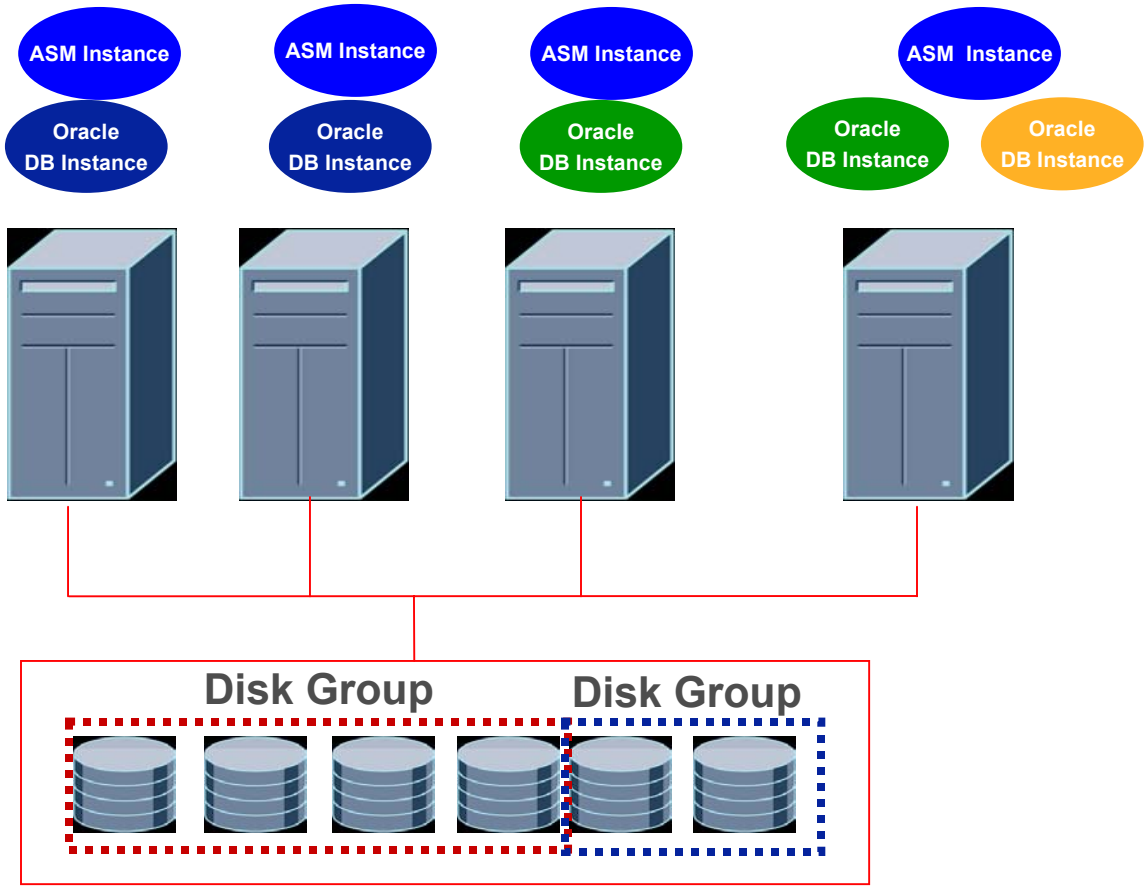

ETAGON

The Operational Stack

TODAY

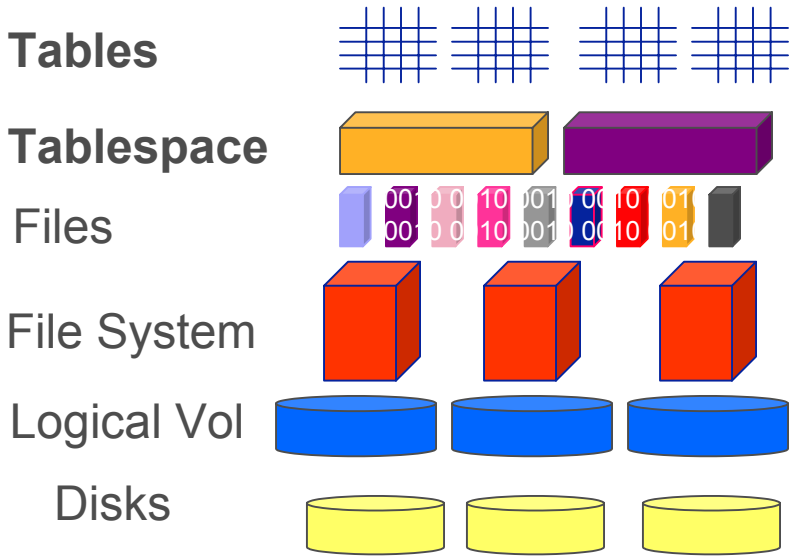

ASM

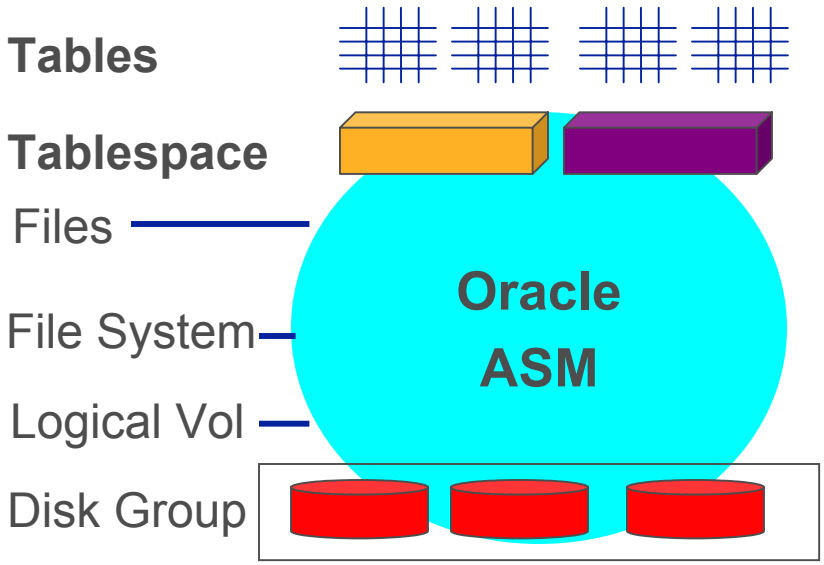

ETAGON

RAC 10g and ASM Synergy

- Simplified Cluster Environment
- Easier Management
- Higher Availability
- Smoother Growth
- Adaptable Resources

Automatic Service Provisioning

- Hands-free allocation and re-allocation of servers to workloads (services) based on rules
  - Automated routing of service requests to appropriate server with lowest load
  - On server failure, automated re-allocation of surviving servers to services

ETAGON

Automatic Service Provisioning

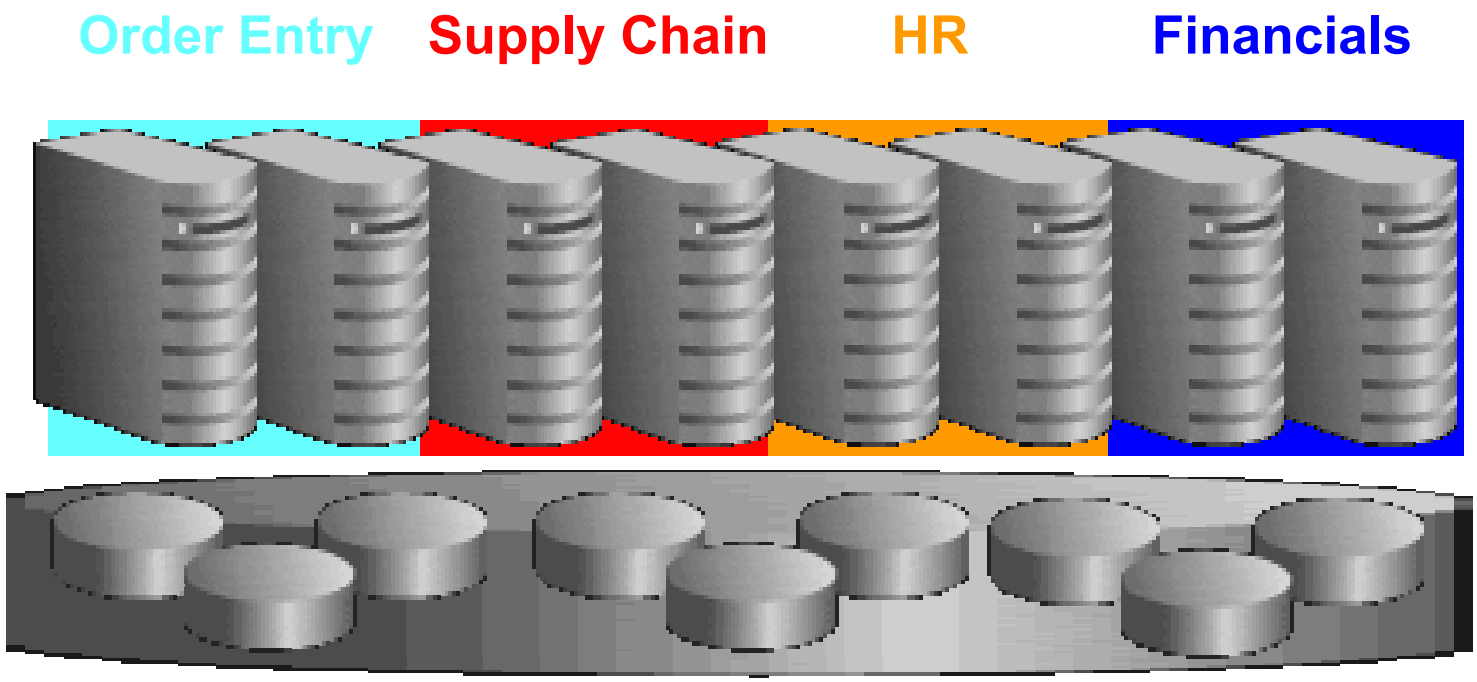

Normal Server Allocation

ETAGON

Automatic Service Provisioning

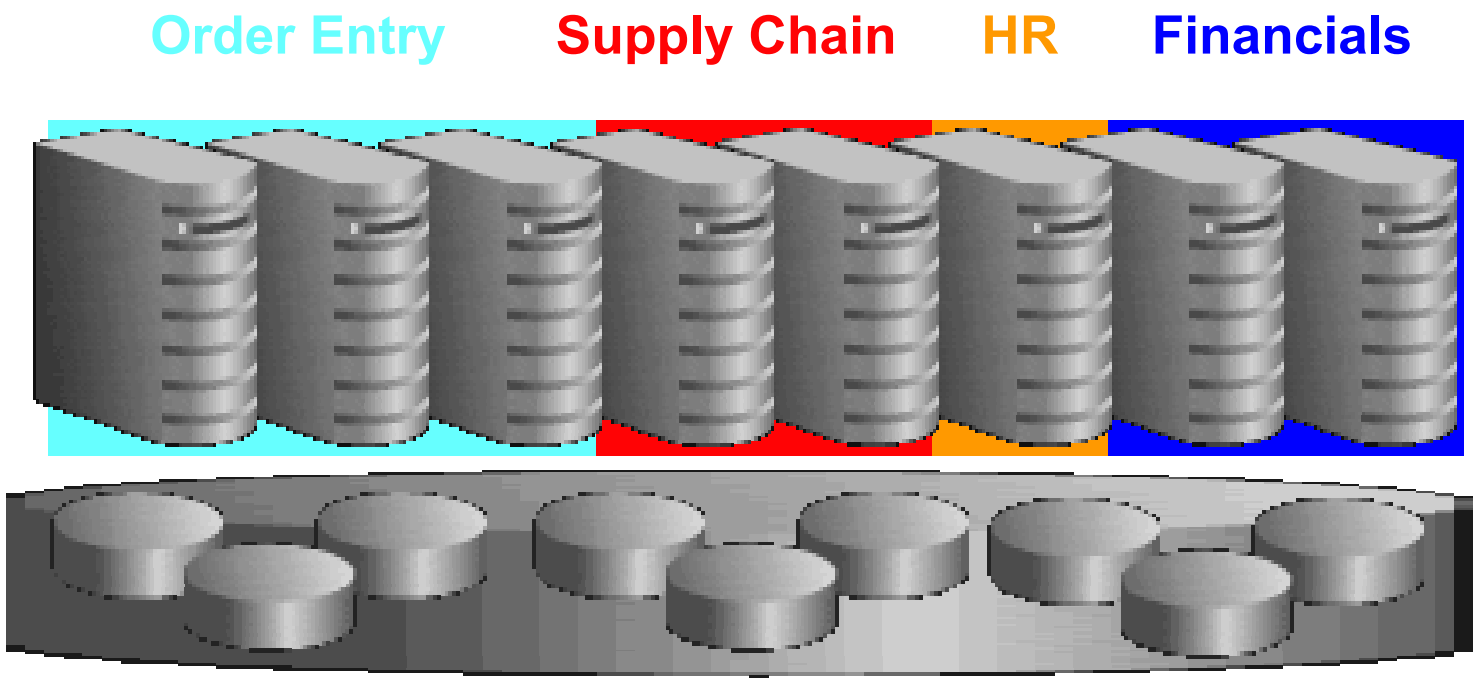

End of Quarter

ETAGON

Automatic Service Provisioning

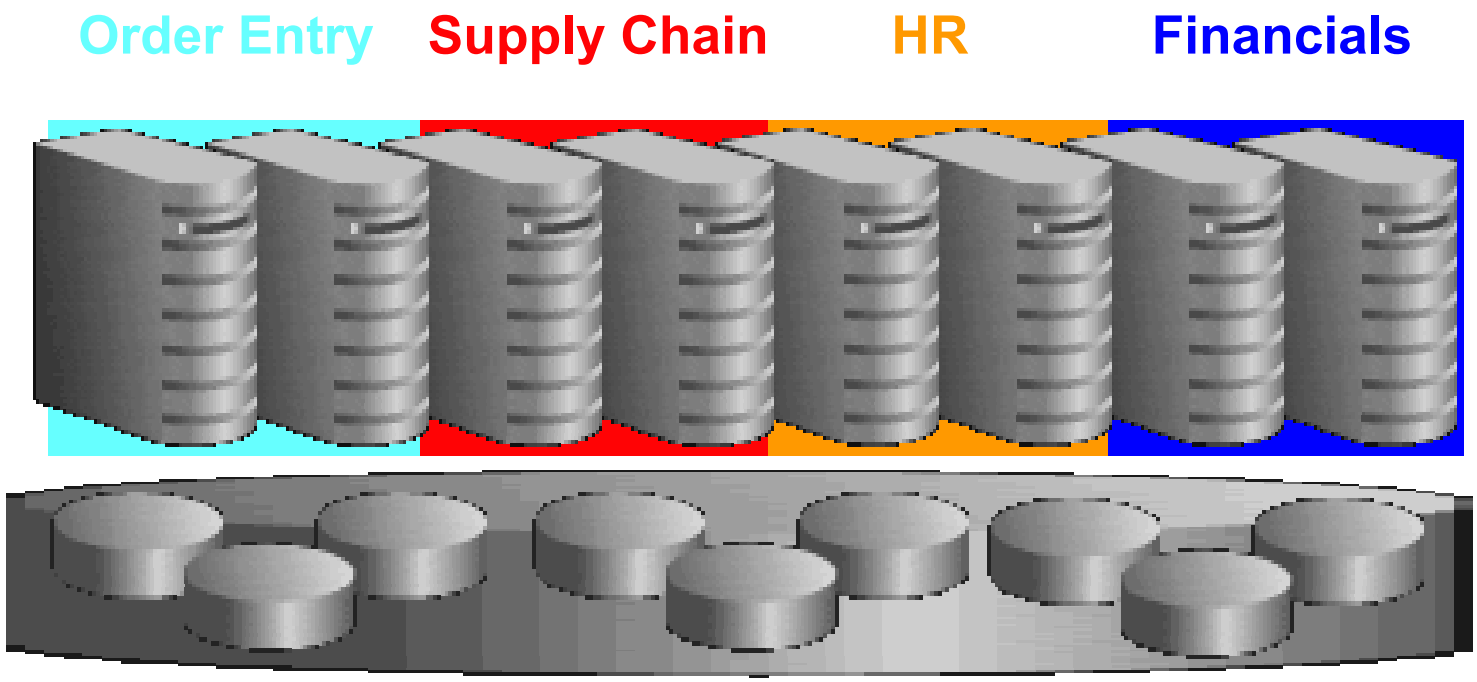

Normal Server Allocation

ETAGON

15 December, 2003

Automatic Service Provisioning

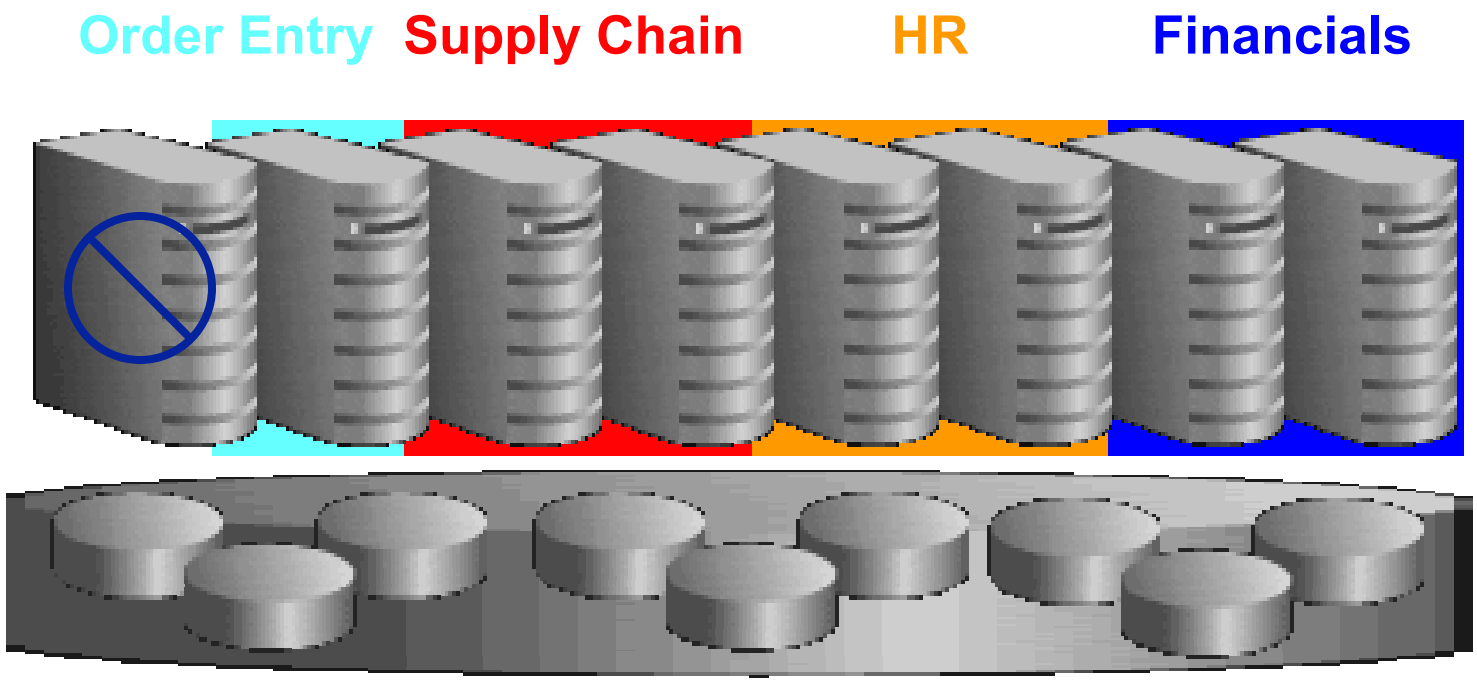

Server Fails

ETAGON

Automatic Service Provisioning

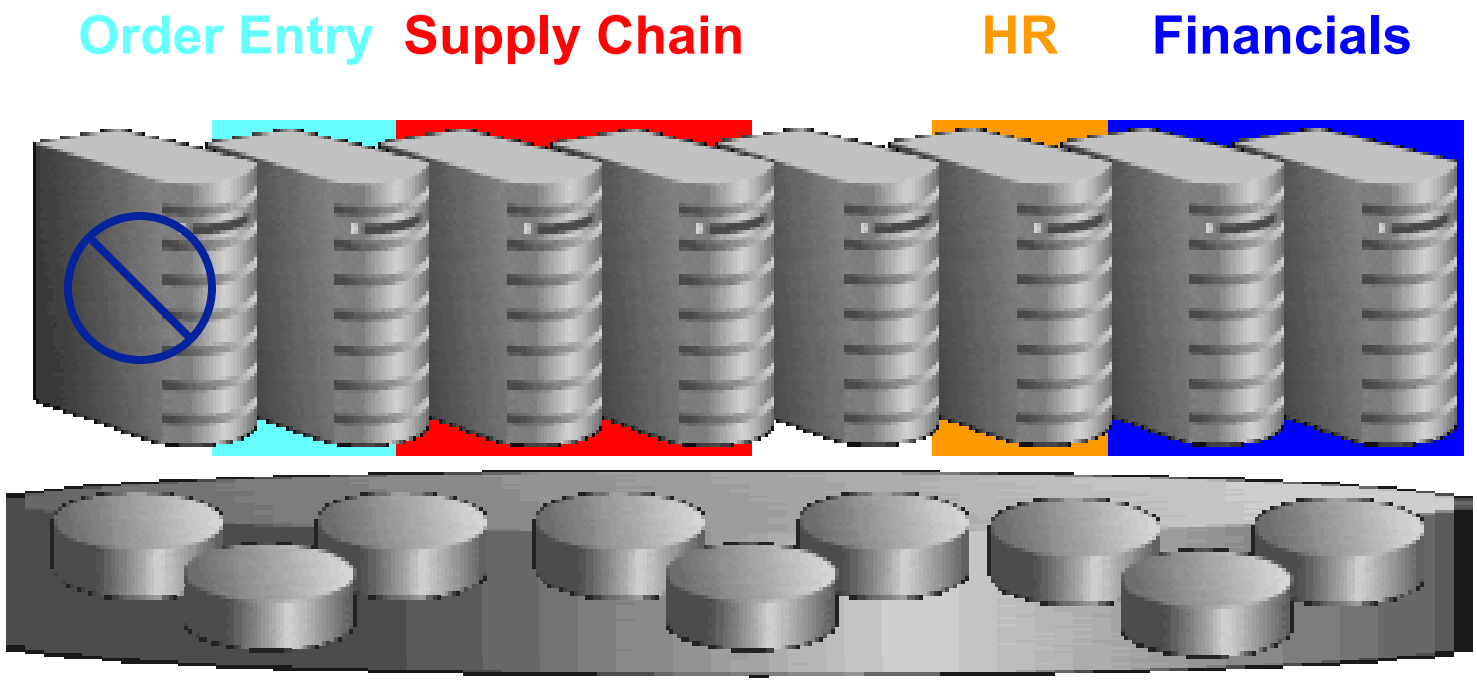

Reallocate HR server to Order Entry

ETAGON

Automatic Service Provisioning

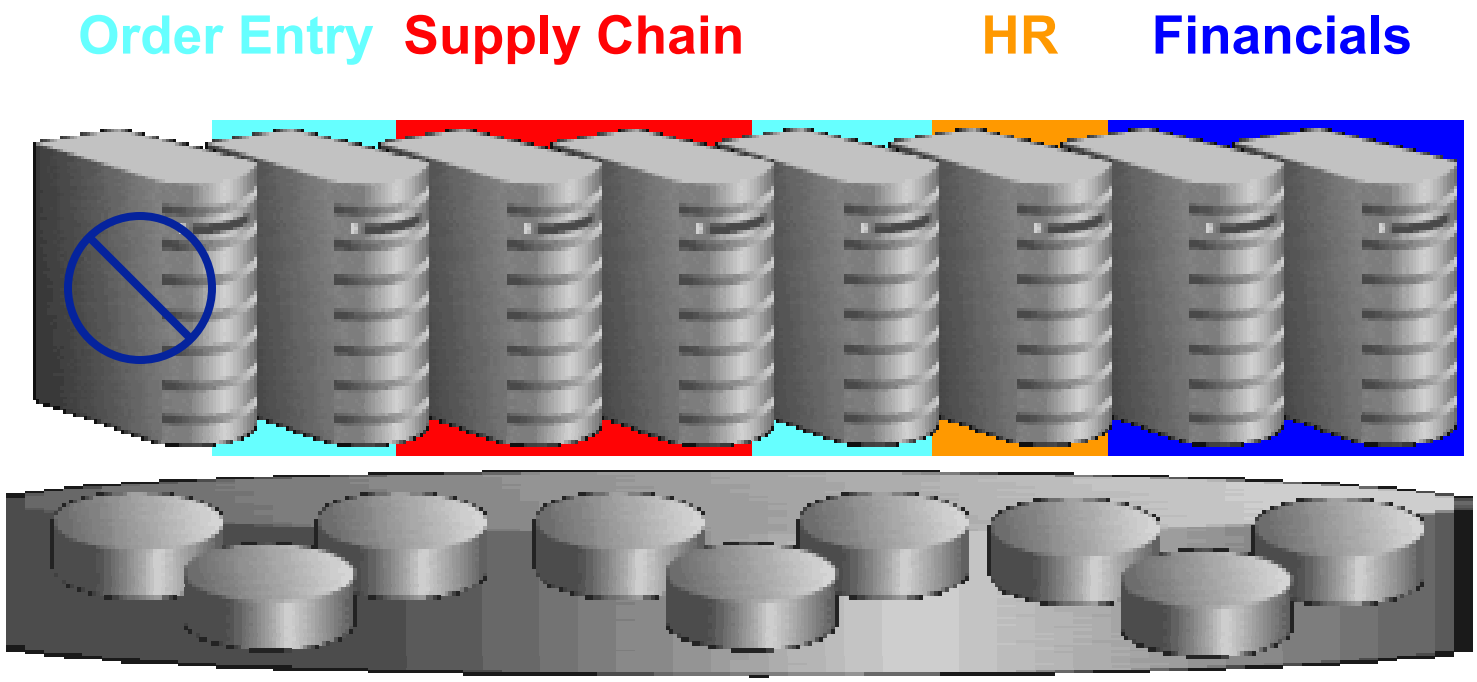

Order Entry on Two Servers

ETAGON

RAC 10g – Services

- Policy based behavior
  - Performance
  - Availablility
- Uses Automatic Workload Repository Statistics
- Created and managed with existing tools
  - OEM
  - DBCA
  - SRVCTL

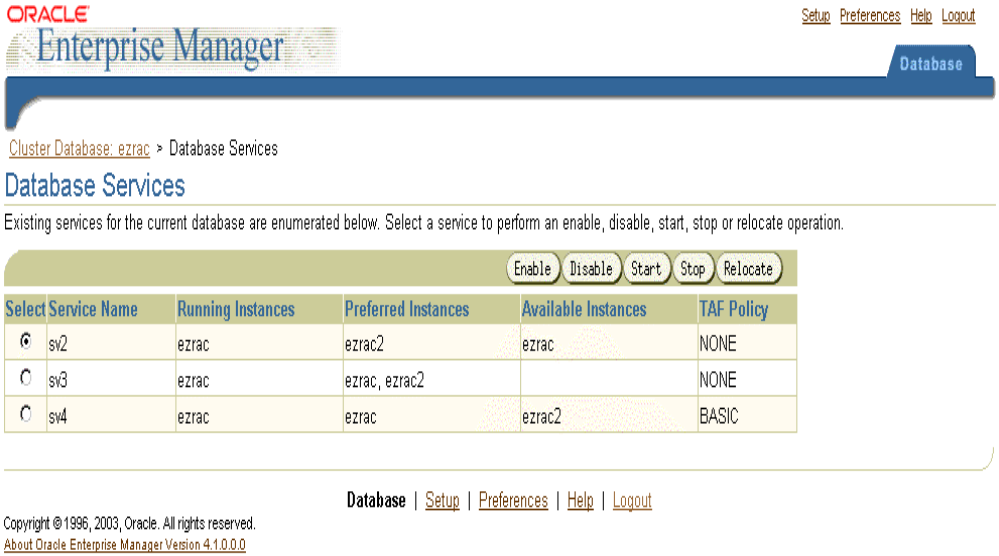

ETAGON

Database Server Consolidation with Oracle*9i* RAC

Mixed Apps and Database Environment

Mixed Apps with common Database but separate data models

Mixed Apps with common Database with 'single source of truth'

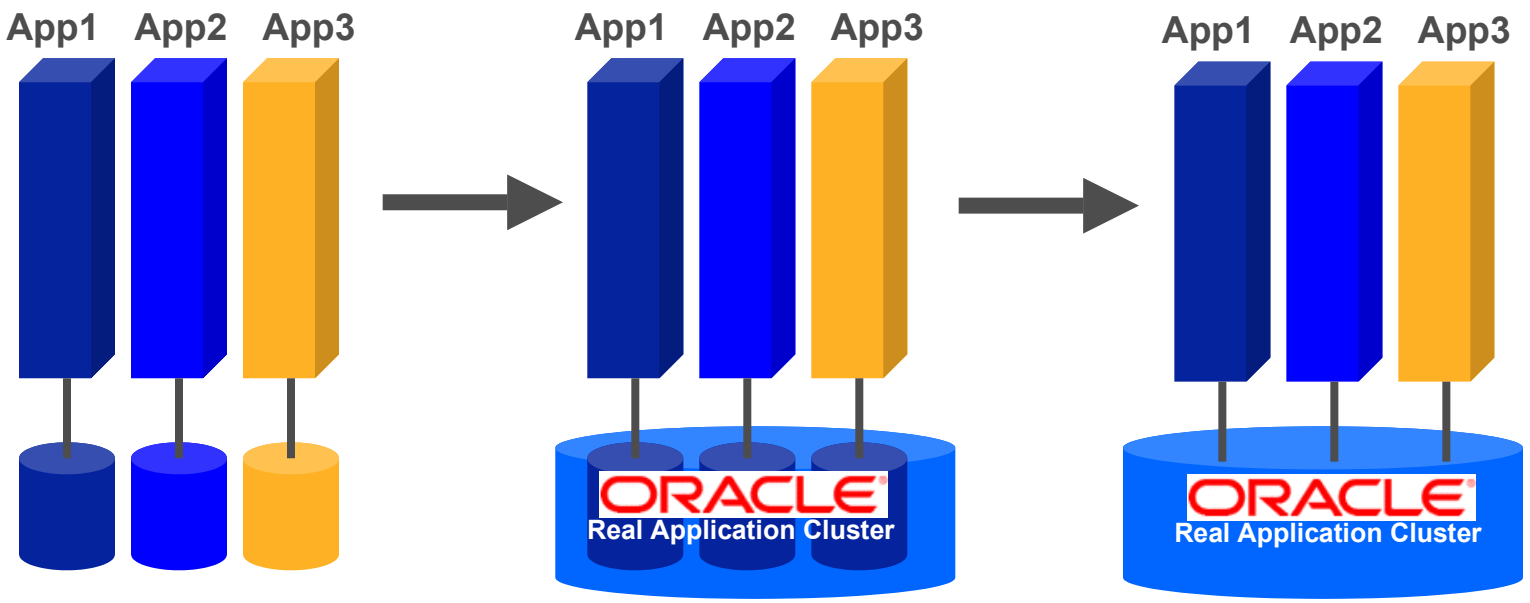

Phase 1 **Phase 2** **Phase 3**

***9i* RAC** will support mixed workloads (OLTP/DSS) within common DB

ETAGON

How many nodes ?

- Depends on the underlying platform
- 2-8 nodes is realistic in 9i
- 10g supports up to 64 nodes
- Scalability – test your own application TODAY

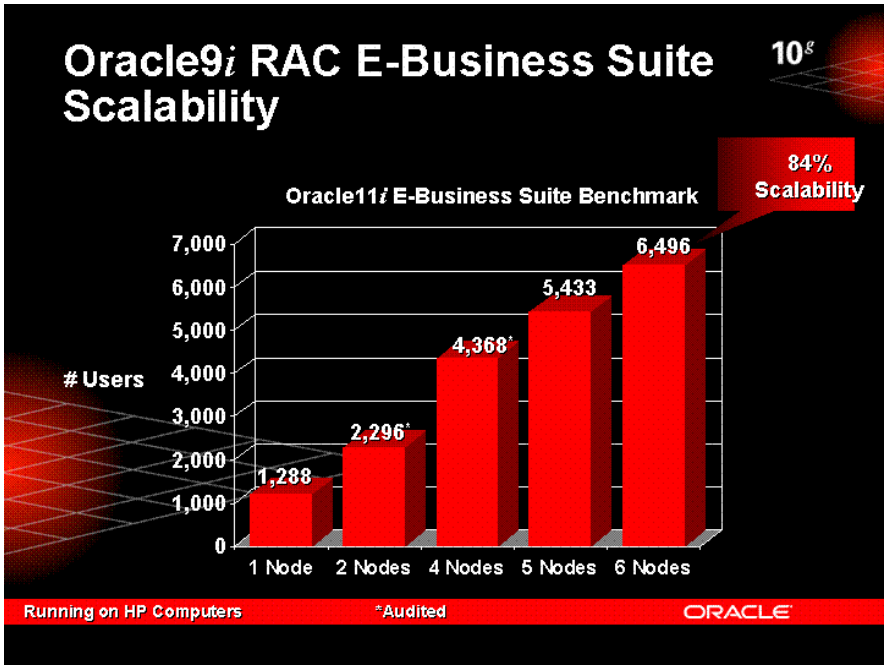

ETAGON

TPC-C 10*g* vs. 9*i*R2 on Superdome

On the same HP PA-RISC Superdome,

10*g* was 28% faster than 9*i*R2

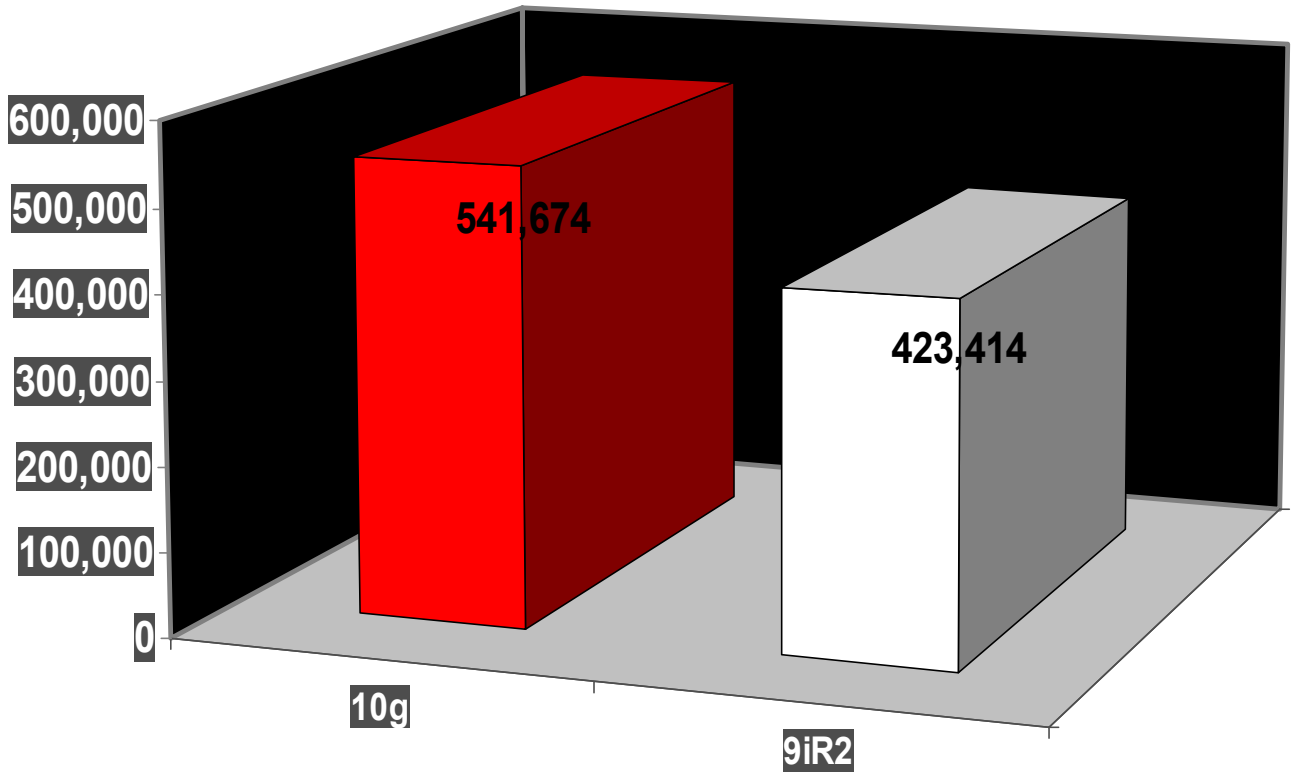

As of September 5, 2003: : Oracle Database 10*g* Enterprise Edition, HP 9000 Superdome, 541,673.76 tpmC, $10.69/tpmC, available 12/31/03. Oracle9i Database Release 2 Enterprise Edition, HP 9000 Superdome, 423,414.41 tpmC, $15.64/tpmC, available 8/26/02. Source: Transaction Processing Council (TPC), www.tpc.org

ETAGON

Rolling Patch Upgrade using RAC

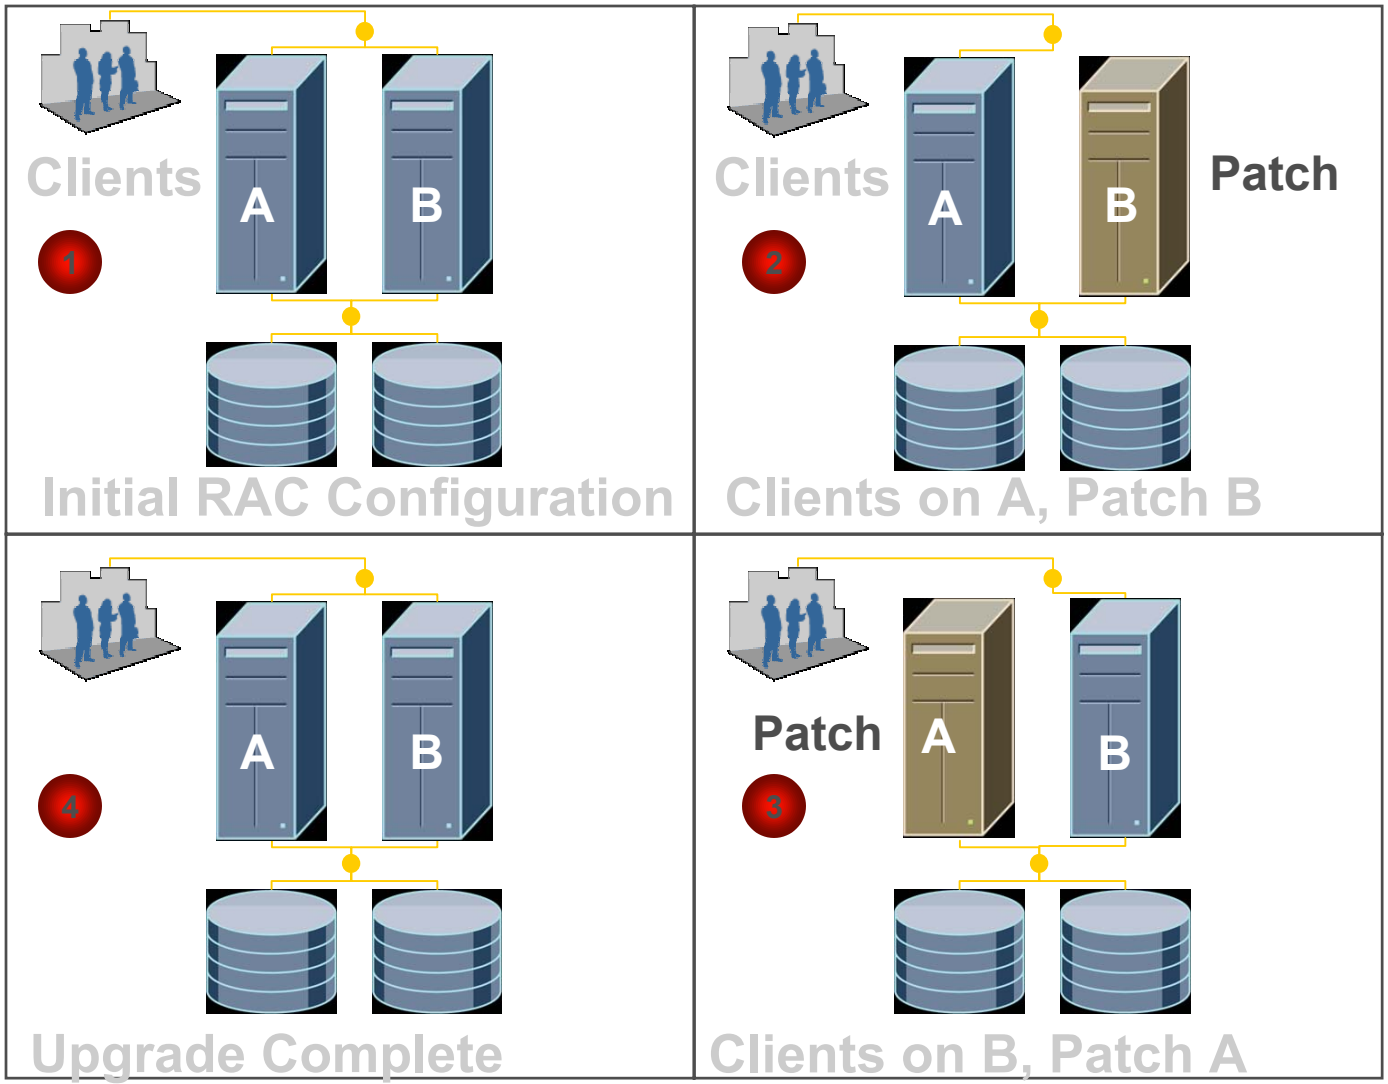

ETAGON

Will I need my old Oracle management tools when I migrate to Oracle 10g?

ETAGON

Oracle9*i* RAC on Linux

- **Clustering consists of 2 Oracle-supplied components**
  - **Clusterware**
    - **Cluster Manager (oracm)**
      - responsible for process level cluster status
    - **hangcheck-timer**
      - monitors the Linux kernel for system hangs
      - resets node from within kernel if abnormal hangs occur
  - **Cluster Files System (OCFS)**
    - overcomes limits and management issues of using RAW
    - comparable performance to RAW IO

Unbreakable Linux - New Support Model

- Oracle takes first call; provides direct support for the operating system
- Oracle has ability to provide relief patches for OS
- Integrated development teams allow adoption of patches in future releases

ETAGON

Summary: Why (and When) RAC?

- **Protection from local site system failures**
- **Why not pure HW cluster?**
  - **Fast Failover**
    - Faster than cold cluster failover solution
- **Scalability**
  - Add and remove nodes transparently – scale-out
- **Cost**

Back to Basics

- What problems are you trying to solve?
- What are the alternatives from a technical/technology perspective
  - Ask the tough questions and test
  - Feature comparison
  - Pay attention to the "fine print" … (unsupported datatype, …)
  - Understand the limitations of the solution and of the implementation
- What is the cost of the solution?
  - Product licensing
  - Additional components needed
  - Don't forget associated cost: application migration, staff training, …

Industry leaders chose Oracle*9i* RAC

9i RAC capabilities for addressing IT challenges

- ✓ High availability
- ✓ High scalability
- ✓ Pay as you grow

9iRAC IS PROVEN

ETAGON

Questions and Answers

?

ETAGON