





# Scale-up with virtualised Integrity servers

Ken Surplice

Business Critical Systems  
HP Europe, Middle East, Africa



Want to combine reduced costs, continuous availability and rapid installation of new Oracle instances?



# How can we help each other?

You



Me



# A bit about me



## EMEA Enterprise Technology Centre

The Enterprise Technology Centre is a Oracle Europe, Middle East and Africa (EMEA) facility located in Oracle UK Headquarters at Thames Valley Park, Reading, UK. We have been fully operational since April 1999, and have hosted over 400 customer visits.

Our key mission is to ensure organisations are successful in choosing and deploying the most appropriate technology solutions to meet their business challenges. We achieve this by showcasing Oracle and Partner products and joint technologies and advising on the most appropriate complimentary solutions that meet each organisations unique requirements.

Technical / business presentations and discussions are significantly enhanced by a suite of compelling demonstrations which validate these products and joint technologies.

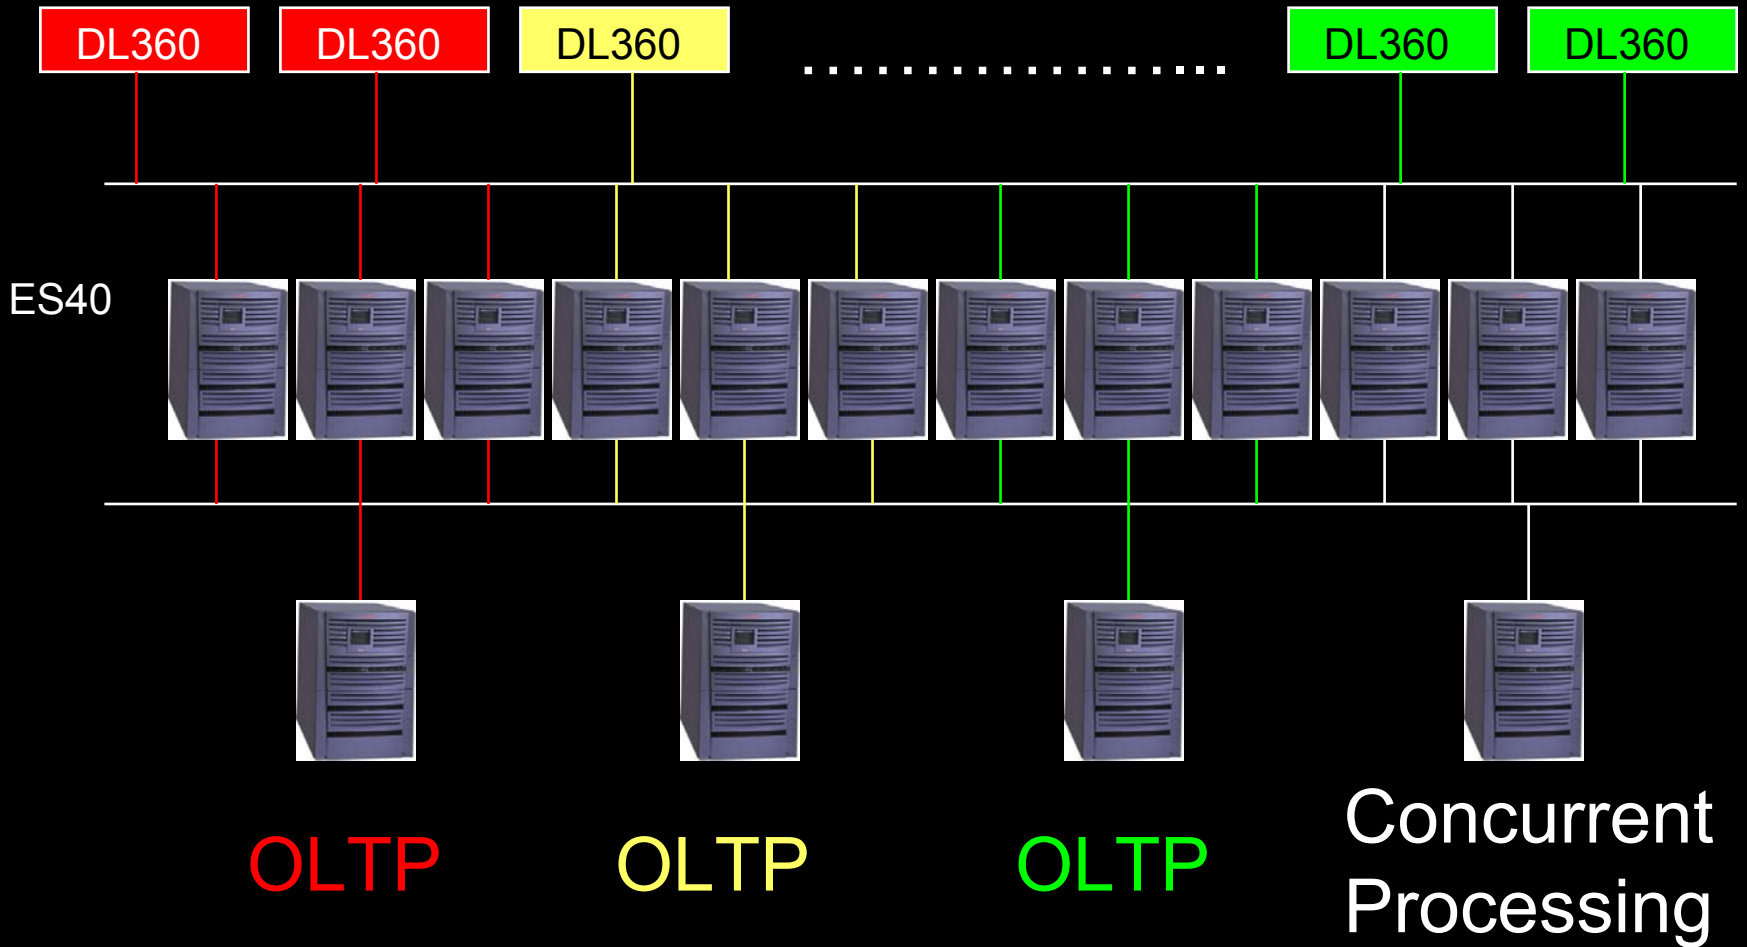
If you are considering High Availability, Disaster Recovery and Consolidation in your Data Centre, we are able to focus on specific areas of Oracle's Maximum Availability Architecture (MAA) and provide proof-points to reinforce the message.

For more information on how the ETC can help you or your customers, [e-mail the ETC](mailto:etc_emea@oracle.com) at [etc\\_emea@oracle.com](mailto:etc_emea@oracle.com)

# etc



# Overall Picture



# 9i Real Application Clusters

- Scalability
- Availability
- Reduce Total Cost of Ownership
  - Hardware Procurement Costs
  - Database Server Consolidation



# You?

Use HP

Use  
Integrity  
Itanium

Use HP  
9000

Use  
Alpha

Use  
Blades

Use  
ProLiant  
x86

Use  
UNIX

Use  
Windows

Use  
Linux



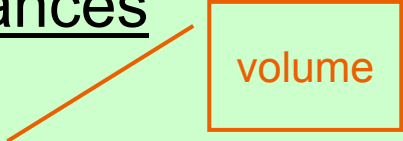
# Objectives for today

## What

- Save you money
- Deliver more to your business
- Simplify
- Standardise
- Get your weekends back


Benefit from Integrity  
(Itanium) servers

## How

- Scale-up, stack-up, to reduce instances
- TCO study 
- Standardise – “goodbye RISC”
- Apply Serviceguard Storage management suite
  - High availability
  - File system simplicity with near raw performance
- Apply common datacentre management

# Common enterprise building blocks

The world's broadest, most robust enterprise offering



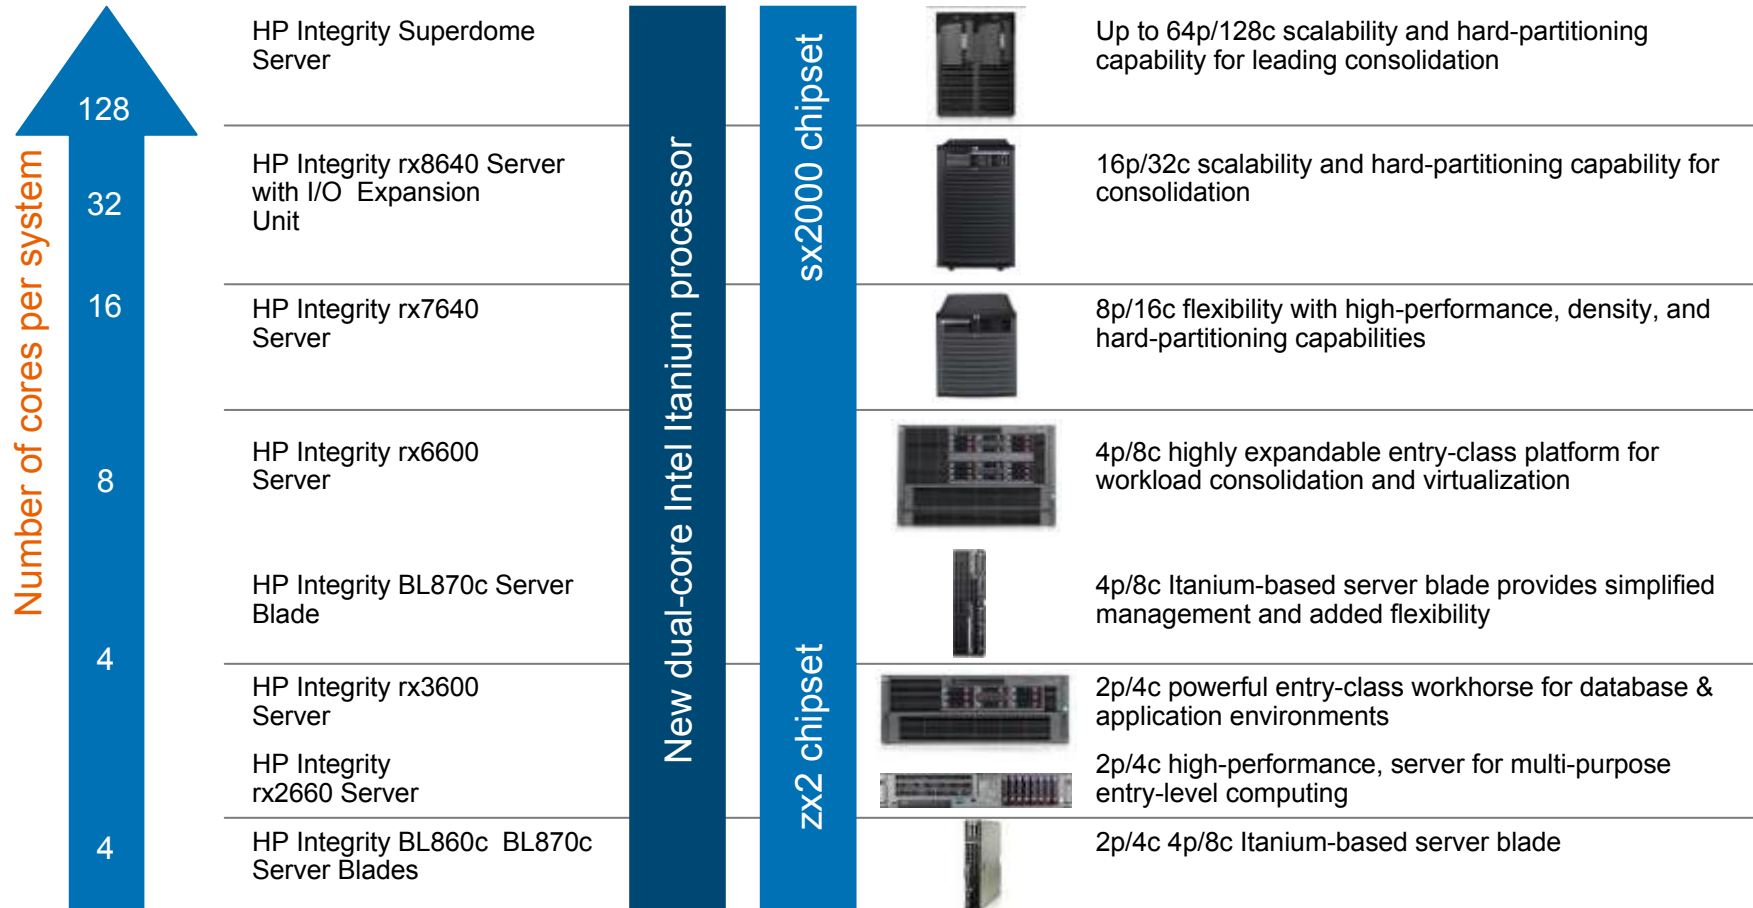


“There is no more important partner for Oracle, and there is no more important platform for Oracle than HP Integrity and Itanium”

---

Larry Ellison  
CEO of Oracle Corporation  
March 2006


# HP Integrity servers: The broadest line of Intel® Itanium® based systems




# How do you choose?



# Your choice: scale-out or scale-up?



- ✓ Cheap
- ✓ Trendy
- ✓ More and more powerful



- ✓ Reliable
- ✓ Predictable
- ✓ Secure
- ✓ Scalable

# Where does your money go?

## Direct one-time Costs

Hardware – 7%  
Software – 9%  
Support – 9%

## + On-going Costs

Communications – 3%  
IT Operations & Administration – 25%  
Facilities – 2%  
Changes – 4%  
Availability – 21%  
Security – 12%  
Application Development – 8%



Lower costs through:

Virtualization; optimized asset utilization, availability & control lowers cost & risk of unplanned downtime

Security; prevent, detect, and contain


Availability; clustering, disaster tolerance & recovery, dynamic tunables

Management; seamless single asset view & control



# Economic Impacts on IT: Management and Power & Cooling Costs Ramp Dramatically

## Worldwide IT Spending on Servers, Power and Cooling, and Management/Administration



IDC - 2007


Many Servers, Much Capacity, Low Utilization = \$140B unutilized server assets



# Data center challenges

- Rapidly respond to business and customer needs
  - Struggle to meet service level agreements for critical workloads
  - Can't implement new projects fast enough
  - Need agility to respond to changing business needs
- Reduce costs of operations and maintenance
  - Too many applications, too much customization
  - Underutilized servers
  - Escalating power and cooling costs
- Enhance quality of service to enable business success
  - Need to ensure business continuity
  - Protect critical resource and data assets
  - Meet Compliance requirements

## IT budget allocation

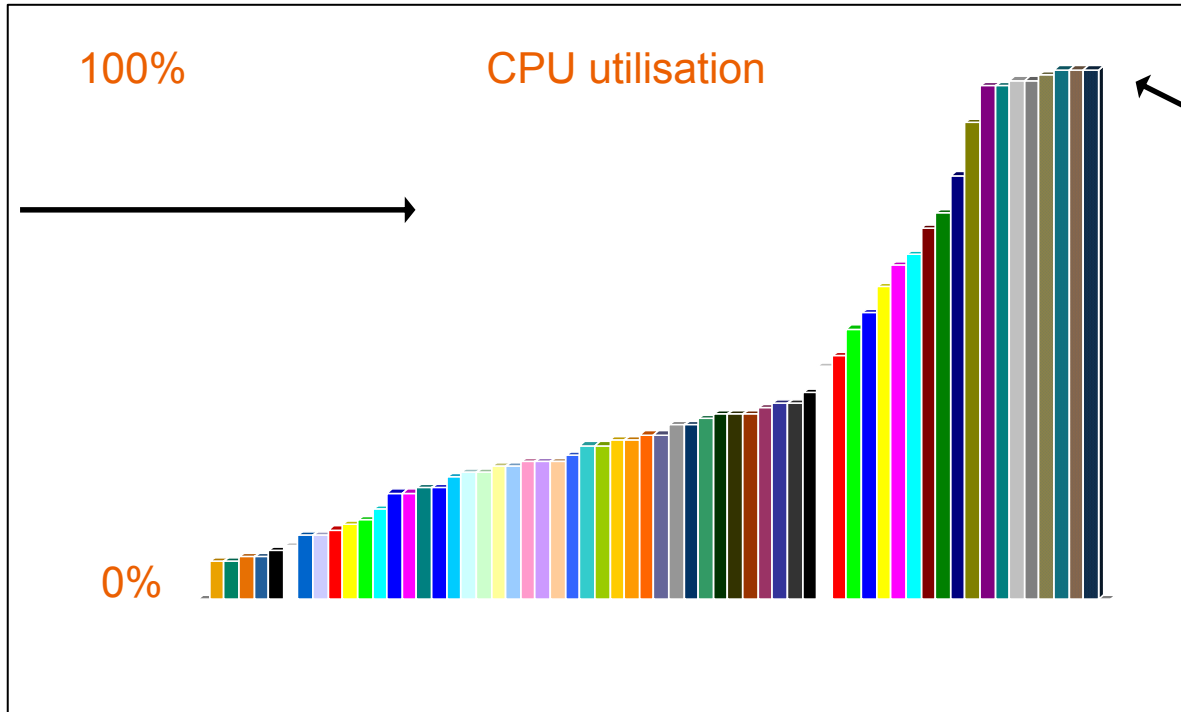


Not enough investment in innovation; too much in maintaining legacy infrastructure

Shift resources from maintenance to innovation

# Plenty of power but in the wrong place

unused  
capacity



Yet these  
systems are  
unable to  
handle the  
load

- Most reports put average utilization at approximately 30%
- Some applications still not able to meet performance requirements

Making the right  
decision

Hints and tips for  
you



# Customer themes today

“Choosing is loosing”

Do you have the same problems as last year?

Does the real problem go unsolved?

If you do not transform, you get more of the same

Automation – too hard vs. off the shelf

Challenge: IT forecast of demand

Service oriented infrastructure

Lights out

Zero touch

From “my needs, my solution” to “my needs; do it for me”

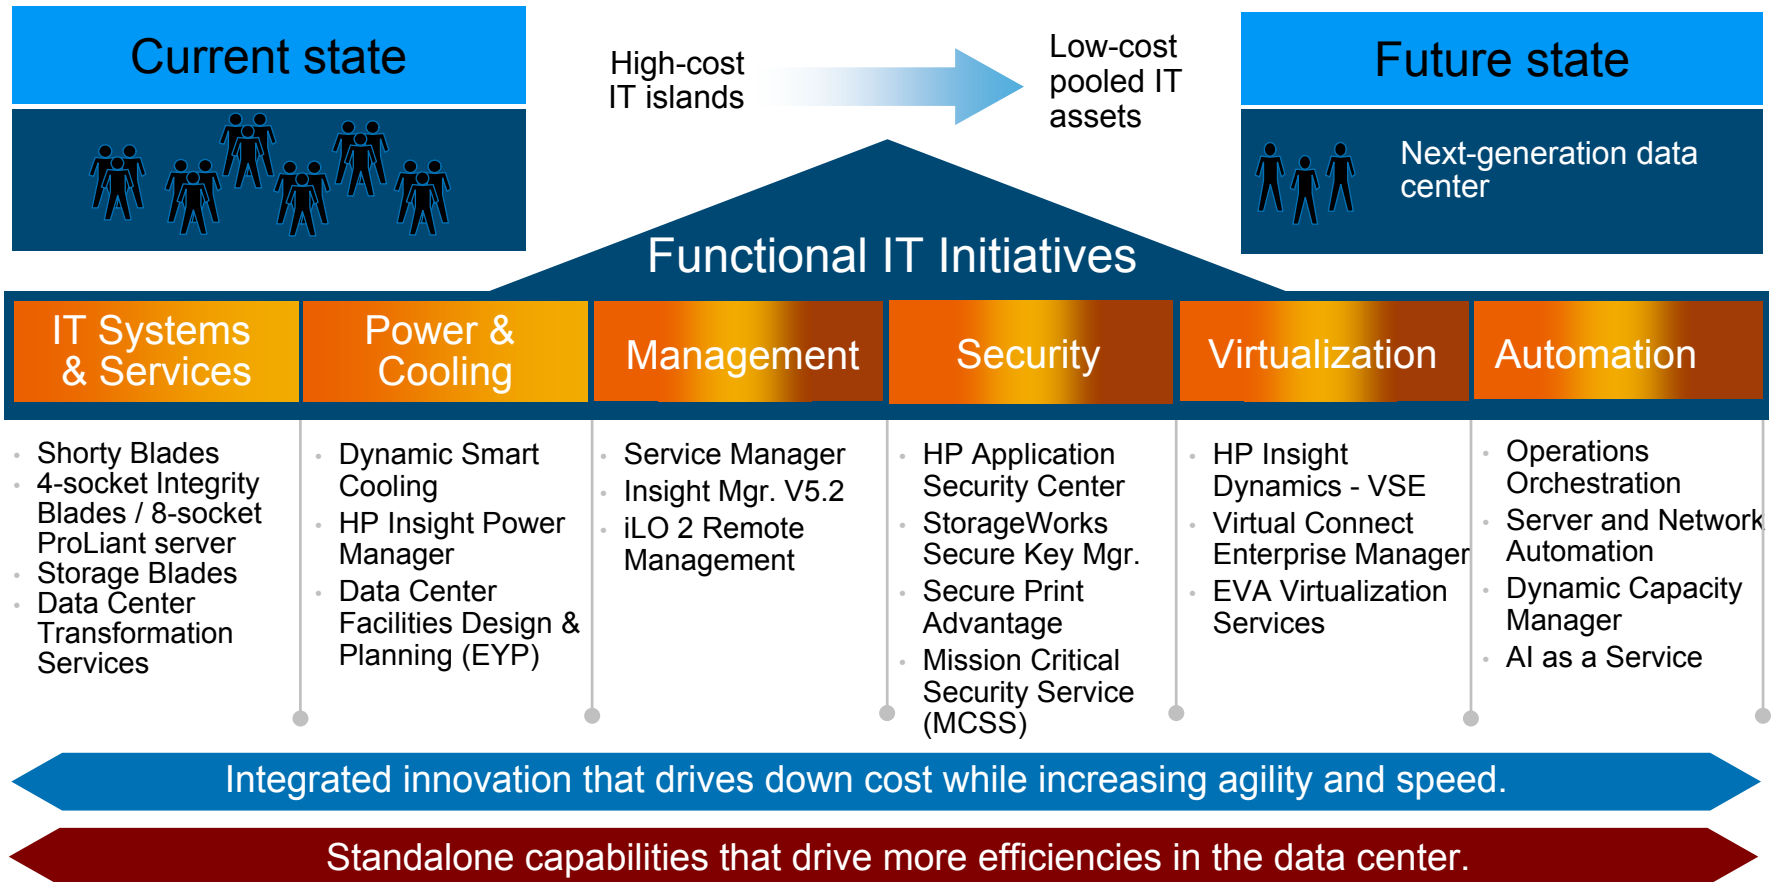
Move to low cost assets

Modular, repeatable environments

High availability and disaster tolerance built-in

# Adaptive Infrastructure same for all servers/operating systems

Avoid “choosing is loosing”



Servers & Storage
  Services
  Software

# Enterprise servers for HP-UX, Windows, Linux, OpenVMS

- Fewer hardware instances
- Fewer software instances

HP BladeSystem

Virtualisation

Polyserve

x86

Scale-up Integrity

Bigger, fewer,  
virtualised instances

Reduced management  
Dynamic workload control

Integrity

Reduce TCO

Further reduced TCO



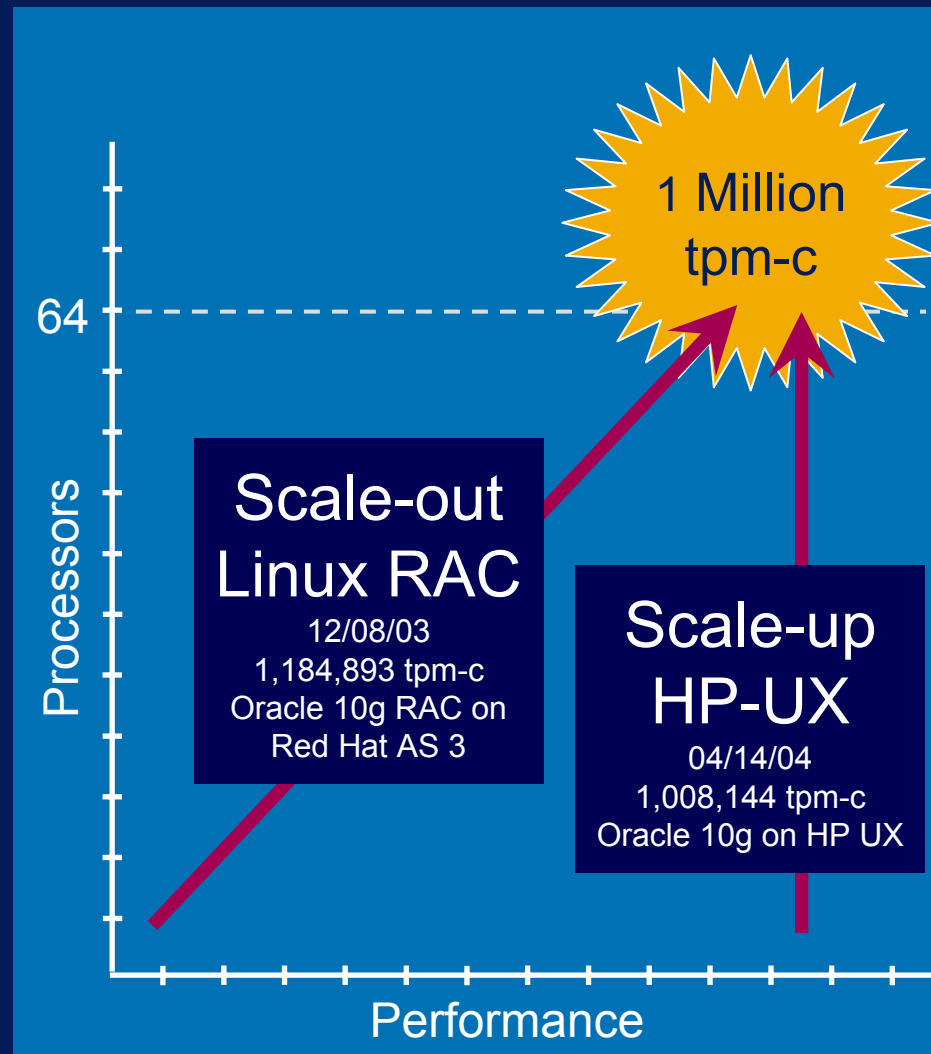
# Case study: tpc-c benchmark configurations



# Can RAC scale? Yes!

## How did we do it?


- HP and Oracle first to break 1 million tpm-c
  - UNIX and Linux
  - Both with same # of Intel Itanium processors
- HP sets the bar for Oracle
- Linux RAC TPC-C
  - 16 nodes x 4 Itanium processors
  - Red Hat Linux AS 3
  - HP StorageWorks MSA100
  - Gigabit Ethernet interconnect



# Oracle single instance database

Thousands of data accesses from memory can occur while one block is fetched from storage

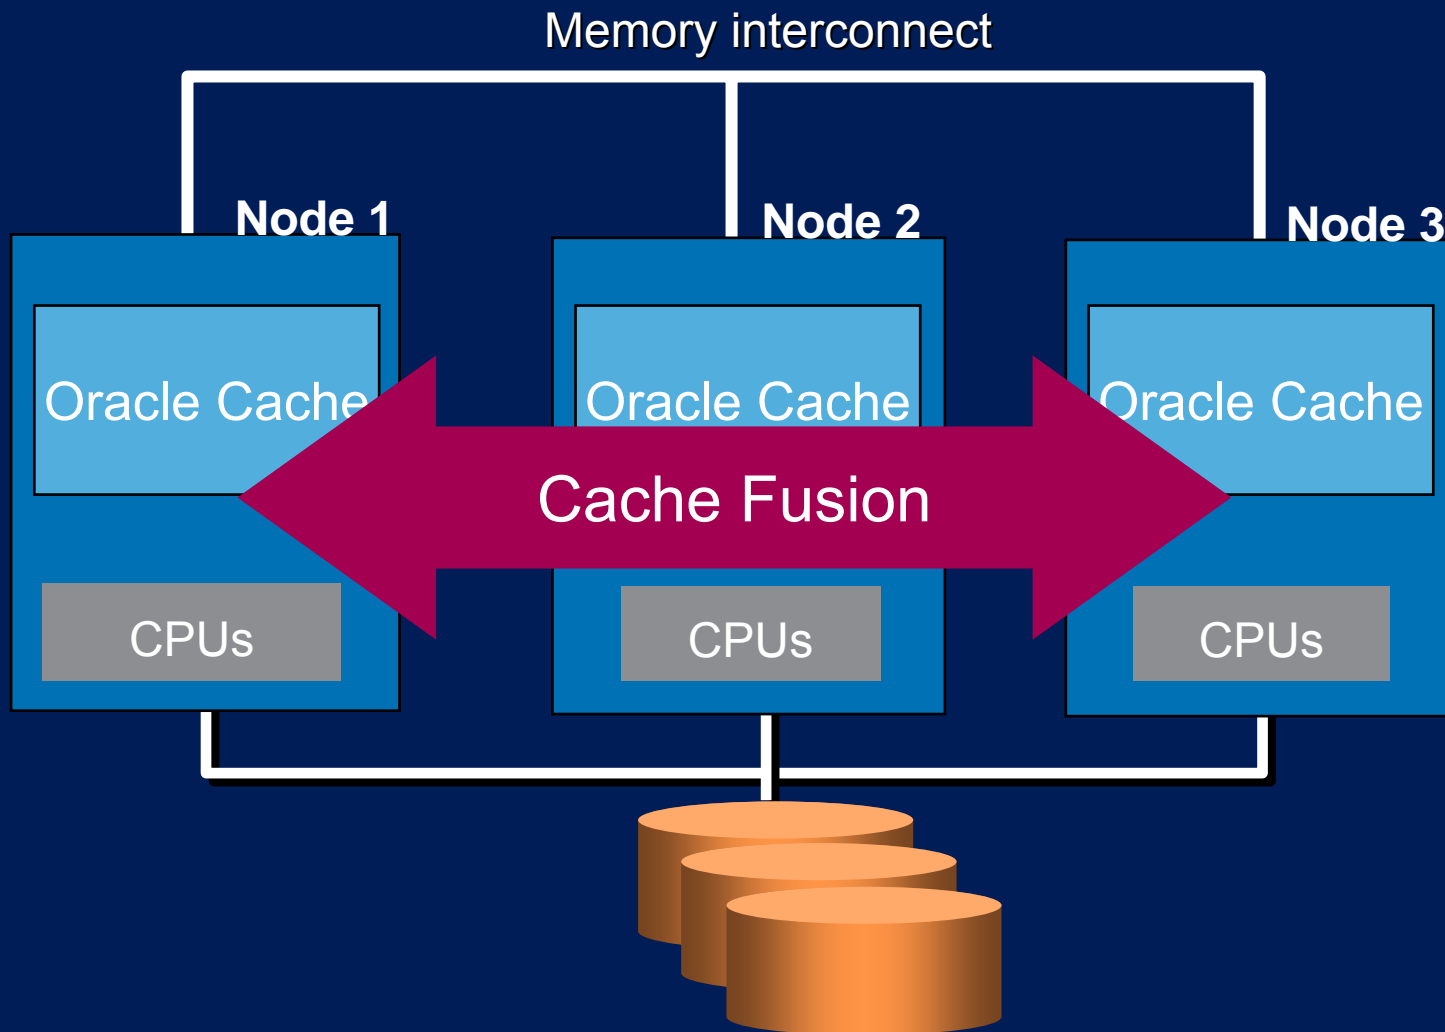
1. First job runs
2. Second job runs
- ...
- n. Next job runs



Data in memory is accessed in nanoseconds  
 $1/1,000,000,000$  second

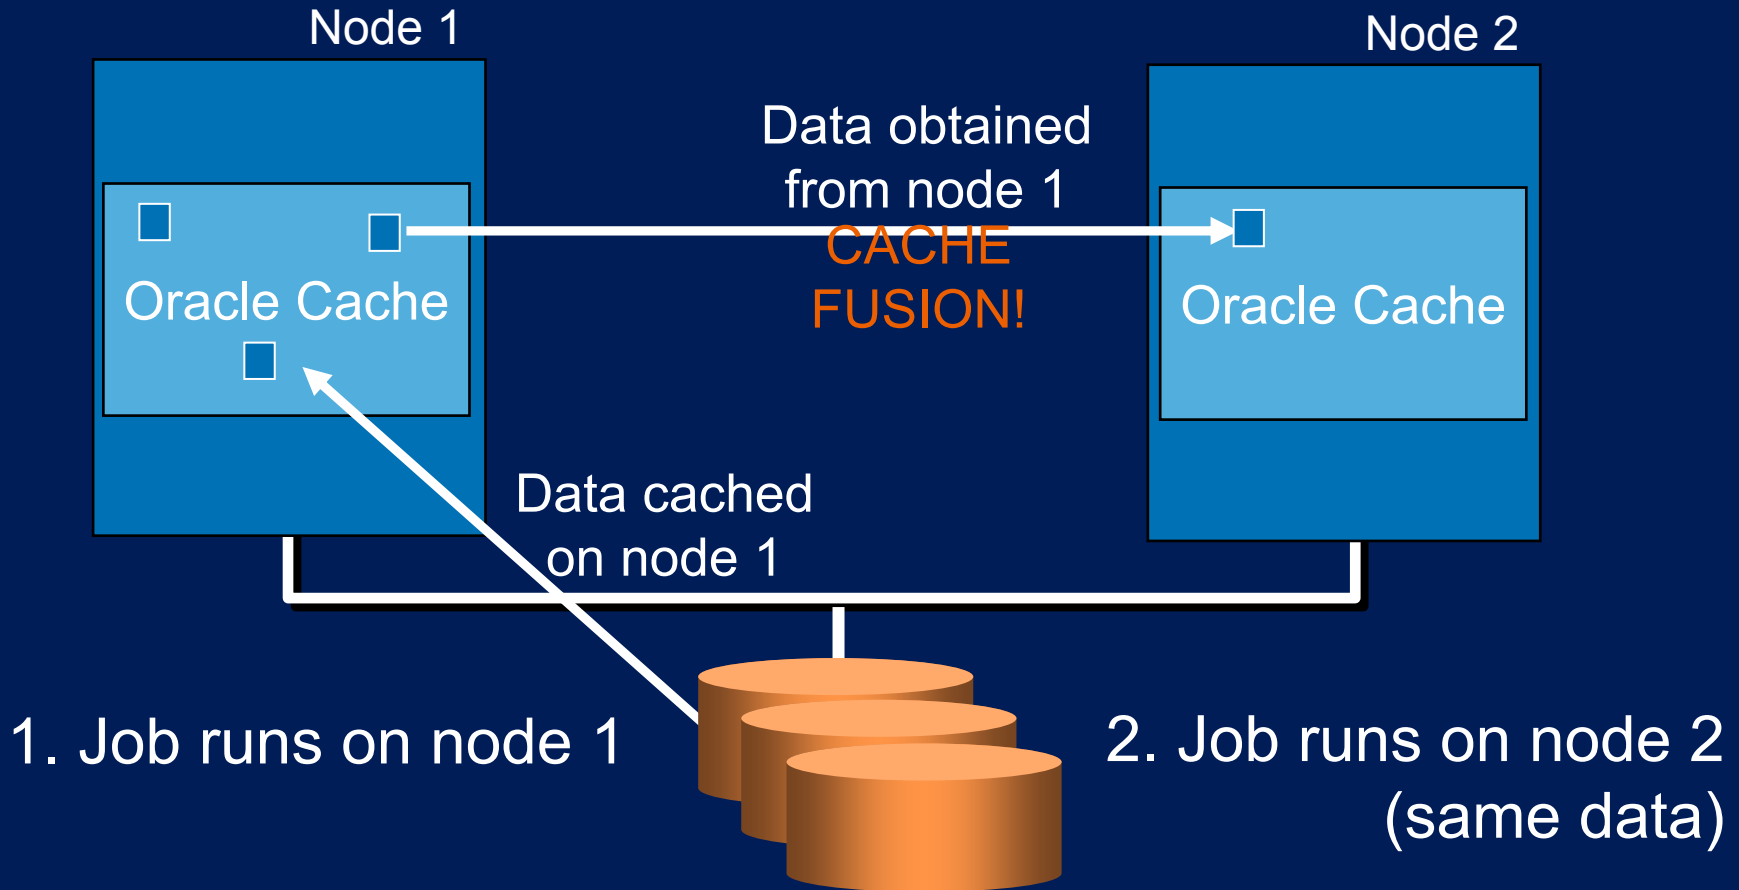
Data on disk is accessed in milliseconds  
 $1/1,000$  second

# Oracle clustered database technology Real Application Clusters



# An example of cache fusion

observe: latency across nodes



# Performance characteristics

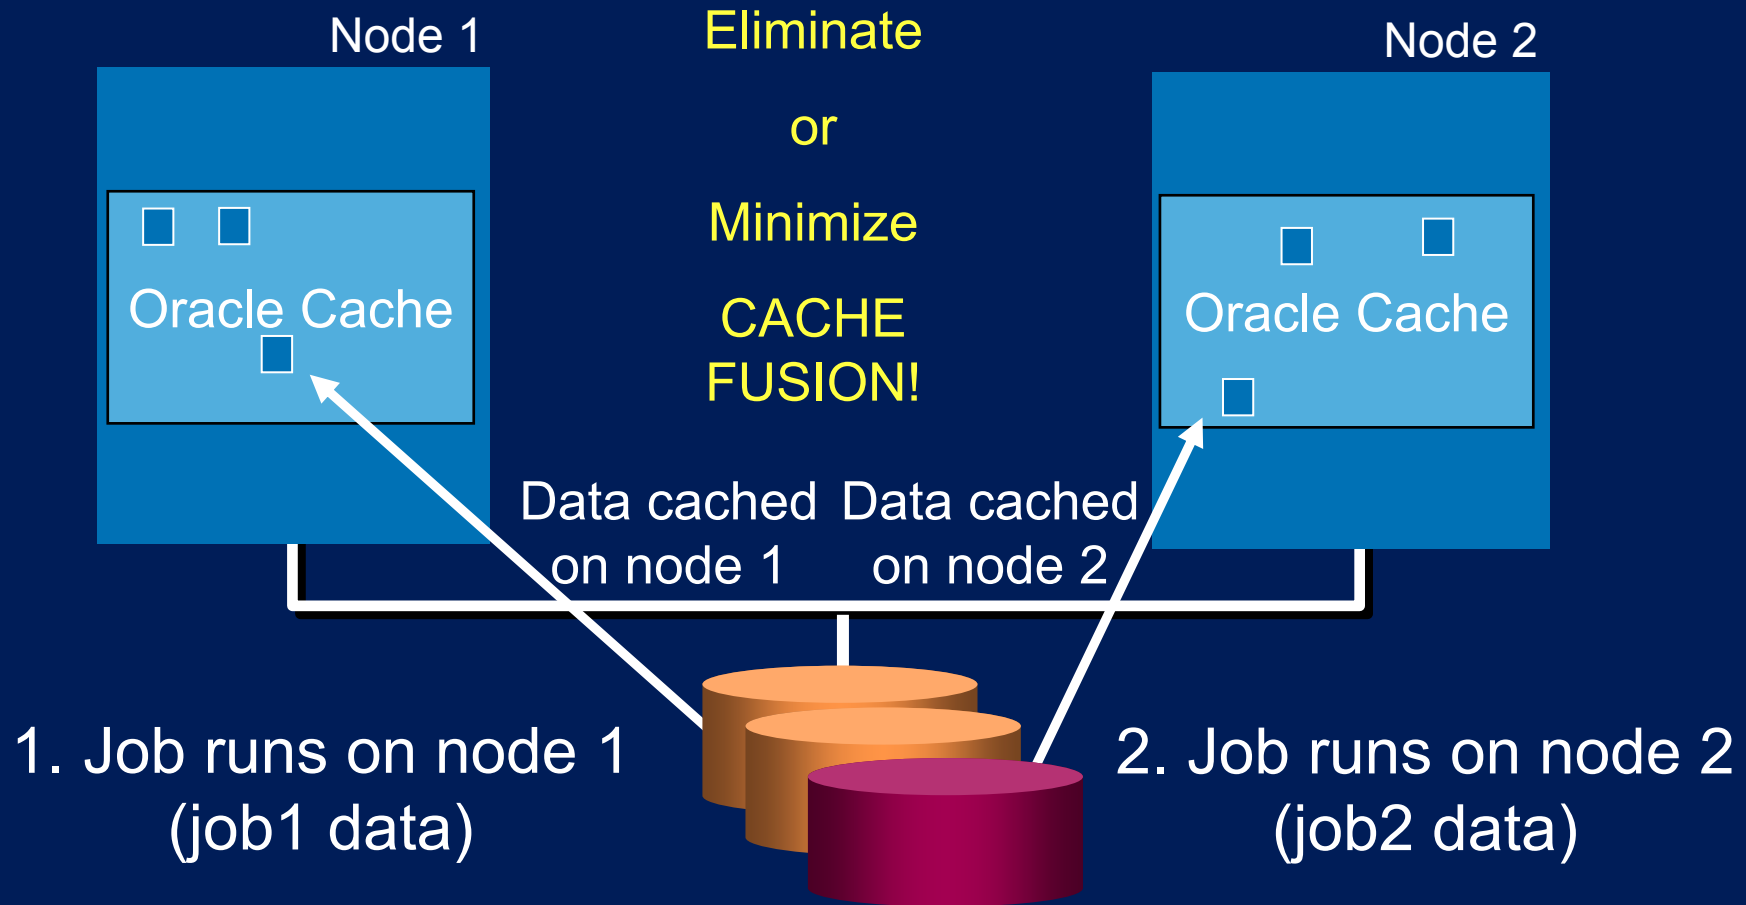
## Cache fusion vs. SMP backplane




Performance characteristic	Cache Fusion		HP-UX internal memory backplane
	1 Gb Ethernet	10 Gb Ethernet	
Data rate	1 Gbps	10 Gbps	320 – 1024 Gbps
Latency	Approximately 60,000 ns <b>136–324x Faster!</b>	Approximately 20,000 ns <b>51–108x Faster!</b>	240 – 440 ns <b>185 – 395 ns</b>
Estimated average CPU utilization	50%	50%	Less than 1%

# Ways to make RAC scale

Make 2 discrete sets of work and data



# tpmC User Response times



# Time and effort to deploy solution

Install & configure Oracle on a scale-up server

## Steps to deploy

- Set up server
- Install NICs for client network
- Connect to network switch
- Connect to SAN switch
- Install OS

## Install Oracle

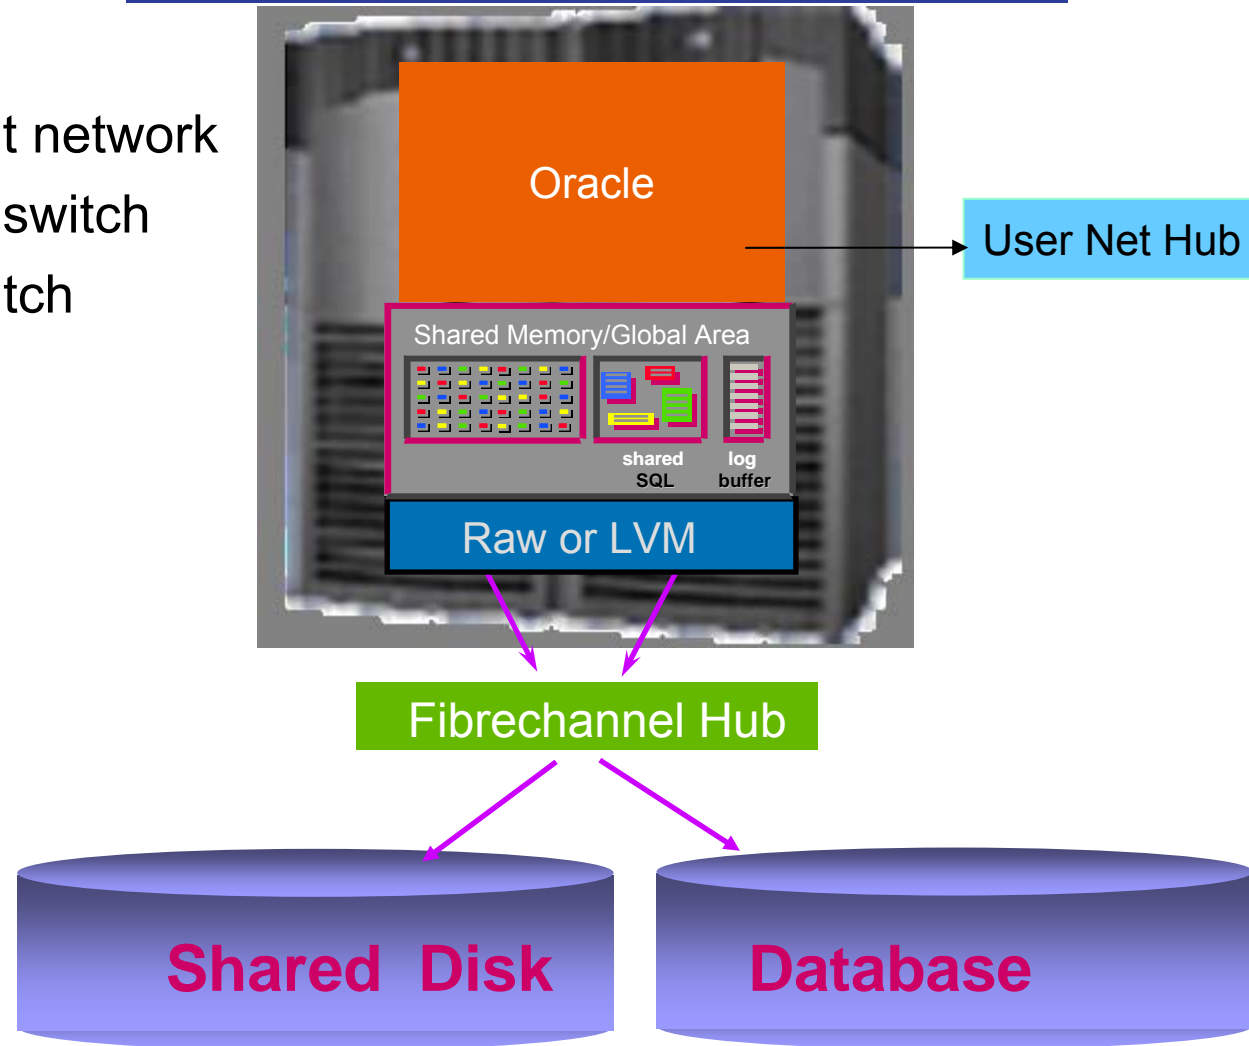
Define Oracle data

...

MAINTAIN

TROUBLESHOOT

Scale-up Integrity 64 processor – TPC-C





# Time and effort to deploy solution

## Install & configure Oracle on a RAC cluster



### Steps to deploy

#### Set up server 1

Install NICs

- two for RAC interconnect
- two for client network

Connect to network switches

Connect to SAN switches

Cross connect SAN switches

Install OS

Define dual pathing to SAN

Define Virtual IPs

...

#### REPEAT FOR SERVERS 2-16

Initialize

- cluster vote disk
- cluster repository

#### Install on server 1

Cluster Software

Advanced Storage Mngt

#### REPEAT FOR SERVERS 2-16

Install Oracle RAC

Define Oracle data


...

#### MAINTAIN

#### TROUBLESHOOT

23 October 2008

### Scale-out RAC 16 server 64 processor – TPC-C



# RAC Application Usage Summary

## Applications that will work

- Partitionable by function e.g. Order Entry and General Ledger
- Partitionable by data e.g. date range or alpha numeric
- Partitionable by users
- Read intensive applications
- ISV applications that have been certified with RAC e.g. Oracle ebusiness suite

## Applications less well suited

- Monolithic application
- OLTP write intensive applications e.g. > 30% writes
- Batch intensive applications that may only run on one node
- ISV applications which have not been certified with RAC i.e. Siebel
- OLAP
- Single threaded application - fastest CPU (e.g. Oracle Supply Chain Planning)

# Application design for scale up / out



Application	Scale Up	Scale Out
SAP R/3 app server	+	+++
SAP BW/APO & Oracle ASCP	+++	
Siebel app server	+	+++
Oracle app server	+	+++
MS Exchange	+	+++
Web Infrastructure (web serving, firewall, proxy, DNS)		+++
Database OLTP	+++	+ (RAC)
Business Intelligence (OLAP)	+++	+ (very few queries suited to scale out)
Java BEA WLS	+	+++

# Typical characteristics of application scalability

## Oracle Real Application Clusters

Applications not suited for RAC Best for Scale Up	Applications that can scale on RAC
<u>Monolithic applications</u> (majority of legacy applications)	Applications can partition by function, data, and users
<u>CPU intensive workloads</u> (such as demand planning)	Discrete, smaller workloads
<u>Write intensive workloads</u> (such as customer billing)	Read intensive workloads (such as data reporting)
Large <u>OLAP/data warehousing workloads</u>	Smaller, partitionable data warehouse workloads
Applications not certified by vendor for RAC	Applications certified by vendor for RAC
Applications that are <u>dynamic</u> —that have “burst” workloads	Applications that are static—that have predictable workloads



# Other considerations



## High Availability

RAC makes the database highly available

BUT

Availability depends on MORE than just the database instance

- RAS of hardware and operating system
- Complexity of the system
- Ability to patch or upgrade the database

## Another option:

- Service Guard for failover of a single instance database
- Reduces planned downtime with rolling upgrades

# Other considerations

## Cost



Cost of hardware and software needs to be reviewed

- Initial PRICE of servers is cheaper

BUT

- Cost of additional network cards and switches
- Higher Oracle software costs



# Oracle licensing & porting strategy

- Licensing:
  - Generally 1 license per core with few exceptions
  - Itanium dual-core CPUs: 0.5 license per core
  - RISC dual-core CPUs: 0.75 license per core

# Virtualisation – where to start?

We don't have the time to do it

Enables very dense computing, but can you cool it?

To move to low cost assets, need an adaptive infrastructure

How do you take savings from infrastructure and move them to applications development?

Everything: IP networks, storage, servers



# Getting real return on your investment

Does your total cost of ownership include:

- ✓ Purchase price
- ✓ Support
- ✓ Staff
- ✓ Facilities and overhead
- ✓ Change costs



How to achieve 20-30% TCO savings:

- ✓ HP Integrity servers
- ✓ Virtual Server Environment
- ✓ HP Blades
- ✓ HP Storageworks, HP Storage Essentials



**Individualized TCO analysis for real world ROI**



# iCAP

- Purchase for average utilization and provision for peak with iCAP cpus
- Create failover nodes that expand to support the failover with iCAP
- What is iCAP
  - Cpus installed in a server but which are not purchased
  - A license to activate cpus for purchase at a later date
  - A conversion to purchase at the price when purchased
  - Refundable iCAP licensing fee at time of purchase
  - **iCAP has no operating environment software costs**
  - **iCAP has no support costs**
  - **iCAP cpus have 5 days of temporary activation with each iCAP cpu**



