



You Probably DO Need RAC

Delivered by:
Matthew Zito, Chief Scientist

156 5th Avenue
Penthouse
New York, NY 10010
P: 646.452.4100
www.gridapp.com

Introduction

- Who I Am
- Basics of Clustering and RAC
- The Value of RAC
- RAC Objections
- Q & A

Setting the Stage

- What is a Cluster
- Clustering Technologies
- Oracle RAC
- What RAC Looks Like

What is a Cluster?

A cluster is a group of servers that act in concert to provide some combination of:

- Improved reliability
- Improved scalability
- Reduced processing costs

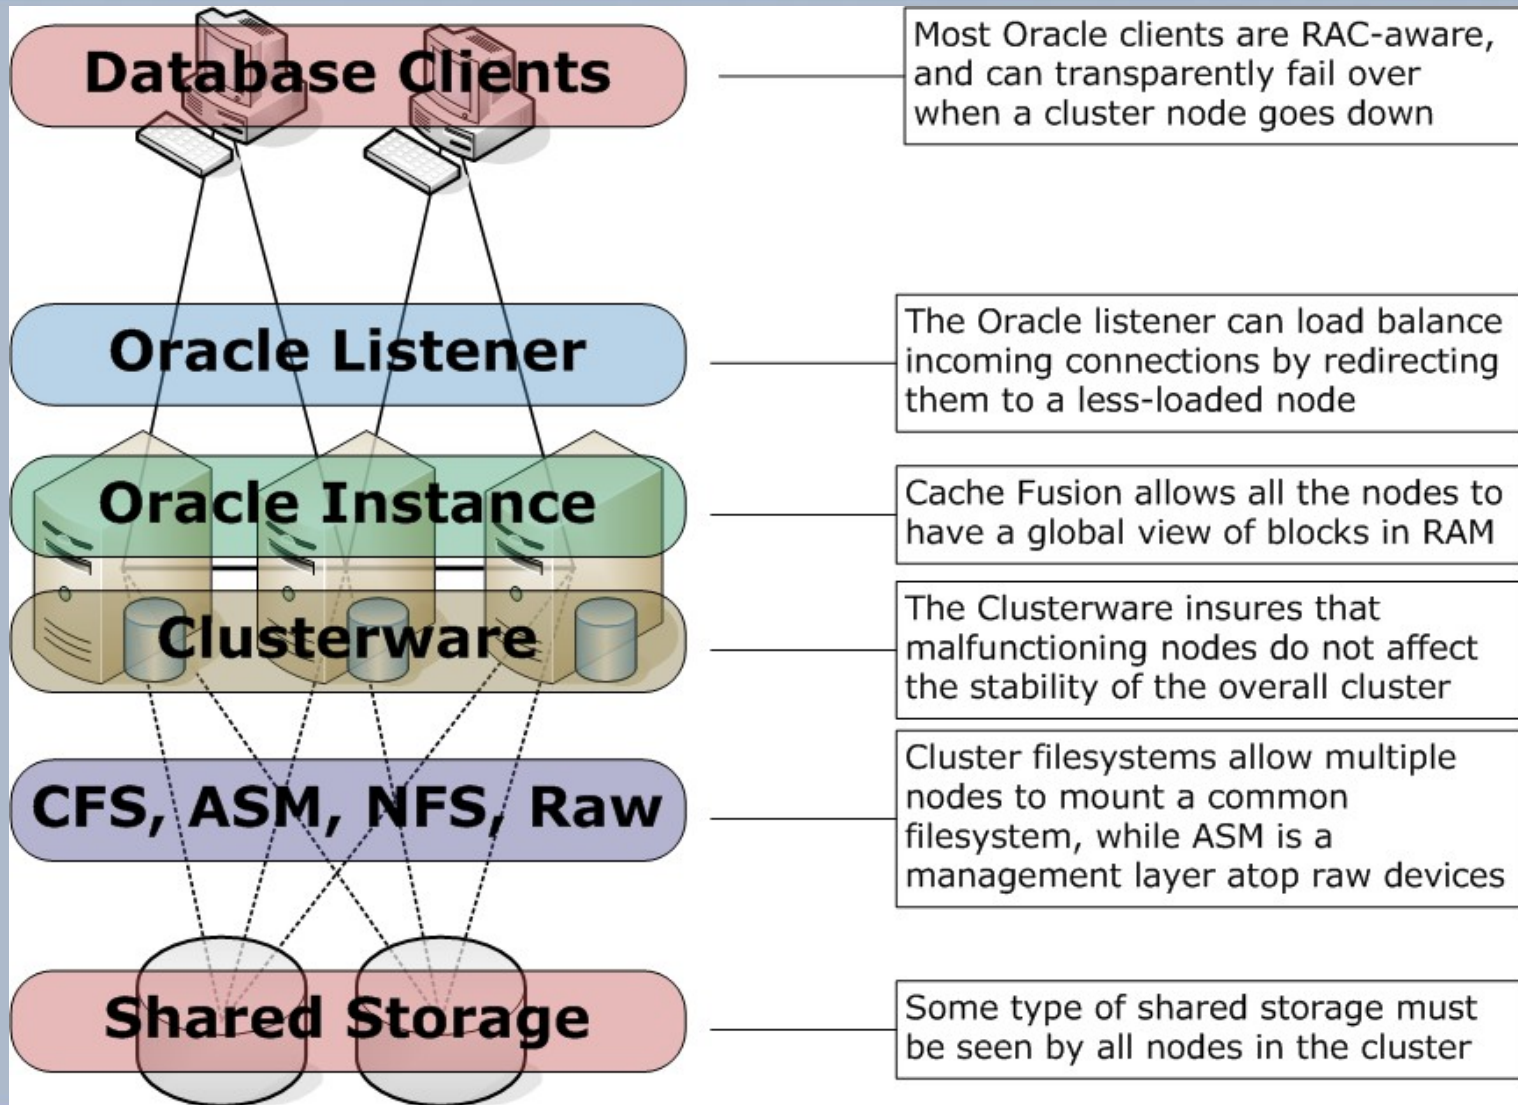
Clustering Technologies

- Shared-nothing – Each cluster node has no common resource with any other node
- Shared-everything – Each cluster node shares a universal pool of common resources
- Shared-something – Each cluster node shares some resources, but not others
- Active-Passive – Only one cluster node at a time is providing any given service

Oracle RAC

- Oracle RAC is a “shared-something” cluster
- RAC nodes share:
 - Application state awareness (Cache Fusion)
 - A single set of on-disk data (ASM, OCFS, etc.)
- RAC provides
 - Enhanced reliability
 - Enhanced scalability
 - Reduced processing cost

What RAC Looks Like



RAC Objections

- RAC is too complicated
- RAC is too expensive
- Commodity hardware isn't reliable enough
- RAC doesn't scale

“RAC is too complicated”

- Older versions of RAC required third-party clusterware except on Linux
 - In 10g, the CRS provides a unified clusterware
 - Users can still choose third-party clusterware based on their particular requirements
- RAC is fairly complicated
 - Adds moving parts
 - Applications can behave differently under RAC
 - Adds skillset requirements
- Automation & standardization can help
 - Make adding and removing nodes seamless
 - Transparently move processing capacity around between clusters
 - Create standardized builds that are easily deployable

“Commodity hardware is not as good as high-end servers”

- RAC mitigates most, if not all, of the downsides of using less reliable hardware
- Small commodity hardware is less likely to suffer from resource contention than a large server
- High-end servers were never that reliable to begin with

“RAC doesn’t scale”

- 10gR2 RAC on Linux has been demonstrated in a lab to scale linearly to more than 10 nodes
- One of GridApp’s customers using 10gR1 RAC has scaled linearly with a production transactional system up to 9 nodes
- Primarily a leftover objection from Oracle 9i RAC

Other RAC Objections & Downsides

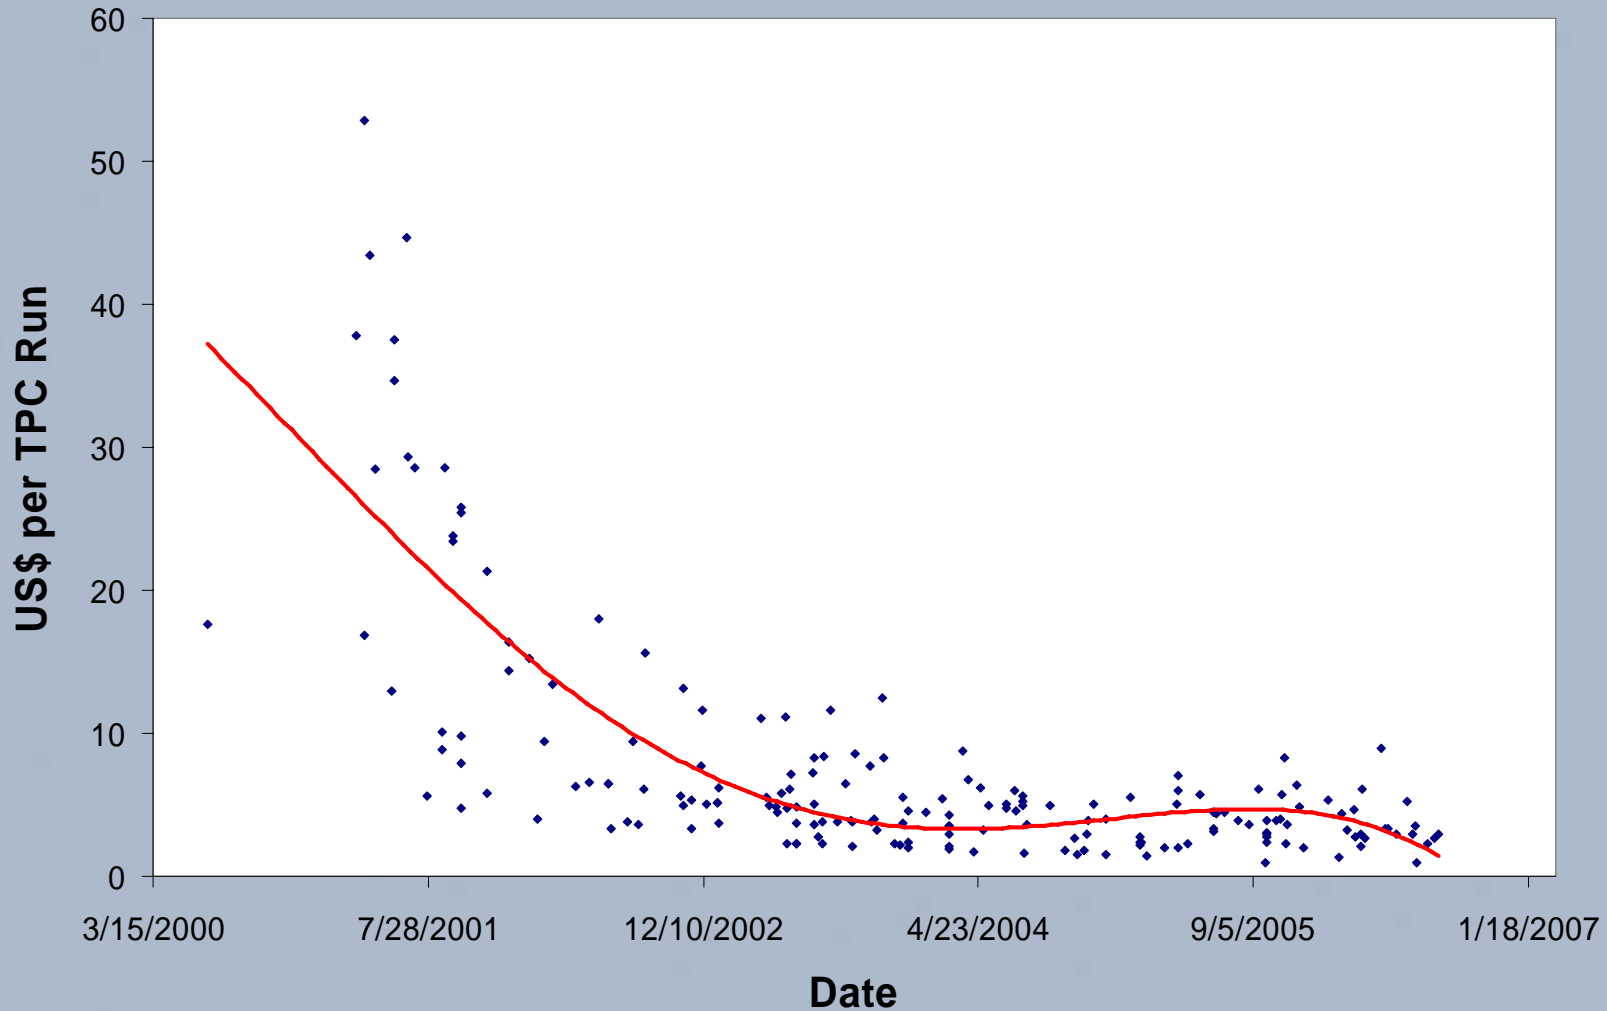
- Oracle sometimes claims that RAC is a seamless migration – it is not
 - Sometimes RAC requires application tuning, or even rewrites of certain functions to optimally take advantage of its features
 - Some applications simply don't run well with RAC
- RAC requires additional training and education
 - This is comparable to that required for Veritas and other active/passive clustering solutions

So Why Do I Probably Need RAC?

You probably need RAC because:

- Commodity processing is cheap
- Downtime is expensive
- Incremental scalability is a huge win
- This is where the world is heading

Processing is Getting Cheaper...



...and Servers are Getting Smaller

- 8/23/06 – IDC Server Shipments
 - Enterprise Servers (>\$500k) – **down 6.9%**
 - Midrange Servers (\$25k-\$499k) – **down 3.5%**
 - Volume Servers (<\$25k) – **up 6.2%**
- Organizations are dramatically downsizing their average deployed server, while dramatically increasing the total number of servers

Server Comparison



1 Sun 6900

32 cores, 128GB RAM

\$1,135,000



5 Dell 6850

40 cores, 160GB RAM

\$139,875

Oracle Licensing Comparison



32 cores UltraSparc IV
Enterprise Edition
\$1,280,000



40 cores Xeon
EE + RAC
\$1,200,000

Total Cost



Hardware: \$1,135,000

Oracle: \$1,280,000

Total: \$2,415,000



Hardware: \$139,875

Oracle: \$1,200,000

Total: \$1,339,875

Is This a Fair Comparison?

- Aren't Xeon cores slower than Sun cores?
- What about a backup server for the non-redundant Sun
- What about power, cooling, rack space?
- What about forklift upgrades?

Reduction of Downtime

- When a node fails in an Oracle RAC cluster, the other nodes continue processing
- Oracle clients can be configured to seamlessly fail over to surviving nodes
- Active/Passive clusters add additional cost to the cluster, require 2x hardware, increase complexity, and create downtimes on failures
- Forrester- 42% of companies found that one hour of downtime costs them ~\$100k

Incremental Scalability

- Buy only what you need today
 - No need to purchase 2x headroom
 - If original configuration isn't fast enough, add more nodes
- Take advantage of the latest technology
 - Moore's law works for you, since you can buy faster servers as released and add them to the cluster
 - Incrementally swap out older, slower hardware for faster hardware

Where the world is headed

- Commodity processing is the way of the future
- Grid Computing (the real kind), is helping organizations create seamless pools of processing power out of servers
- Increasing complexity in the datacenter is driving companies to release products geared to help organizations standardize, simplify, and automate the management of these complex systems

Some Tips for Deploying RAC

- Define a few standard architectures as “certified” RAC builds internally (i.e., 2-node cluster with ASM, 4-node with NFS, etc.)
- Expect to need to invest effort in application tuning to take advantage of RAC
- Get the systems and storage teams involved and educated on RAC
- Test, test, test

GridApp Systems

- Leader in Oracle RAC management and database automation
- Offers the world's first Oracle RAC appliance – the D2500
 - Proven reference architecture
 - Automation to simplify RAC deployment
 - One-click scalability of RAC clusters
 - Free RAC licenses
 - Over 100 deployments
- Headquartered right here in NYC

Summary

- RAC offers huge advantages over traditional single-instance environments
- However, RAC is not the seamless transition Oracle sometimes claims
- Plan, document, test, and reap the rewards



Q&A