

Linux in the Enterprise


James Morle, Scale Abilities Ltd





`who am i`

- Founder of Scale Abilities, Ltd
 - Aka Scalabilities, Scalebilities, Scala Abilities, and so forth
- Author of Scaling Oracle8i
 - Now available for free!
- Co-author of Oracle Insights
- Co-founder of BAARF Party
- Co-founder Oaktable Network



Five.



“Not the...” Agenda

- I am not talking about
 - Linux being different to UNIX
 - Windows
 - SQL Optimisation
 - RAID-F of any sort, ever

Linux Positioning

- Capacity
 - Where Linux platforms reside in the ‘power curve’
- Reliability
 - Do they stay up?
- Support
 - What’s the support like?
- Cost
 - Cheap as chips, but how cheap is that?



Capacity – Benchmarks Overview

- Traditional Perception:
 - Linux is for clustering
 - Linux is relatively low powered

Benchmark Comparisons

- Are these servers actually ‘low powered’?
 - TPC-C soundbites:
 - 1999 - 32-cpu Sequent NUMA-Q scored 48,793 tpmC
 - 2005 – 4-cpu (dual core) HP DL585 scored 187,296 tpmC (albeit running Windows/SQL Server)
 - Almost 4x throughput with 1/8th CPUs

Benchmark Comparisons (2)

- Compare to 'still in use' legacy servers:

Year	Vendor	System	tpmC
2002	Compaq	4-cpu ES45/1250	56,375
2001	Sun	14-cpu UE4500	67,103
2001	Sun	64-cpu UE10000	115,395

Capacity – Important Trends

- Database Performance *needs* (hardware only)
 1. FSB speed!
 2. Large Memory footprints/64-bit addressability
 3. Multiple balanced I/O channels
- Both Intel and AMD continue to increase FSB speeds in addition to CPU clock speeds

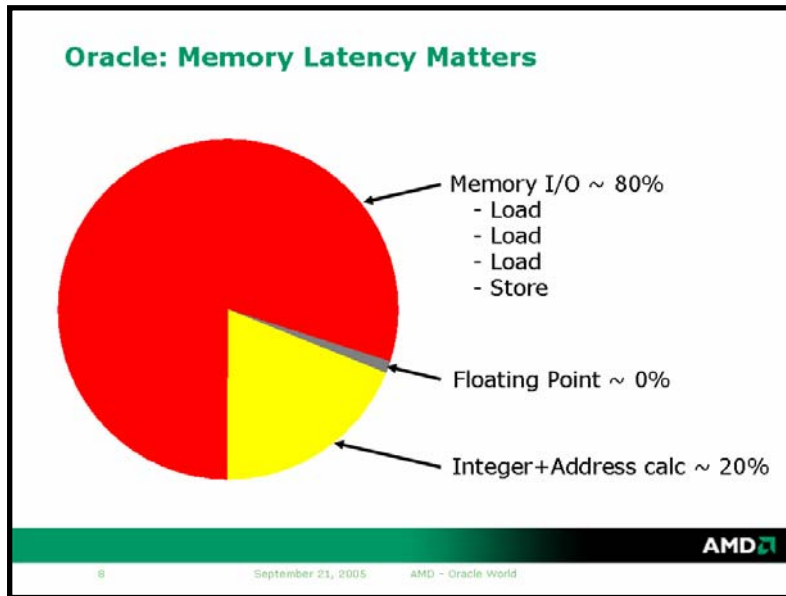


Intel and AMD

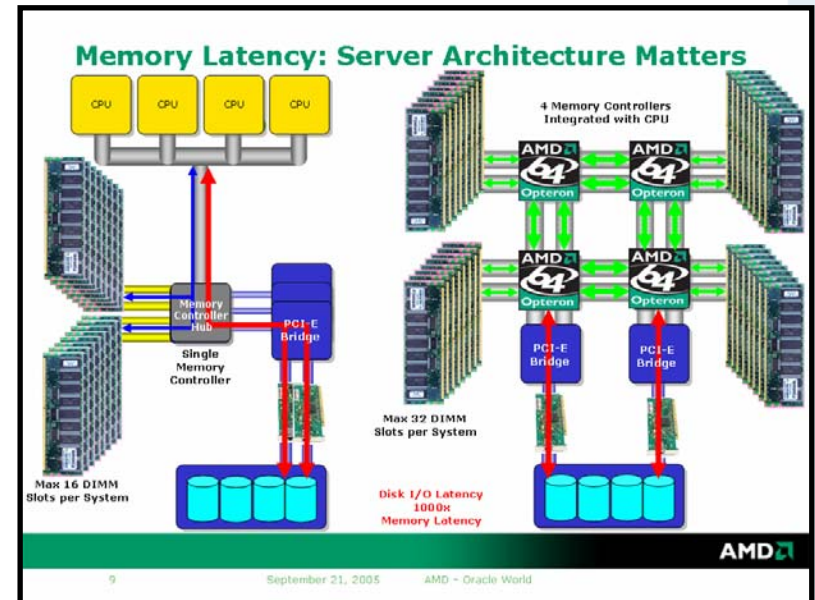
- Xeon
 - IA32/x86
 - Recently extended to support AMD64 model – EM64T
 - Northbridge Memory Controller
- Itanium2
 - IA64
 - EPIC (Explicitly Parallel Instruction Computing)
 - NOT compatible with x86 (IA32)
 - Northbridge Memory Controller
- Opteron
 - AMD64
 - Hypertransport@1GHz

Extract

Presentation given at OOW2005 by Jeff Needham, Scale Abilities, on AMD's booth



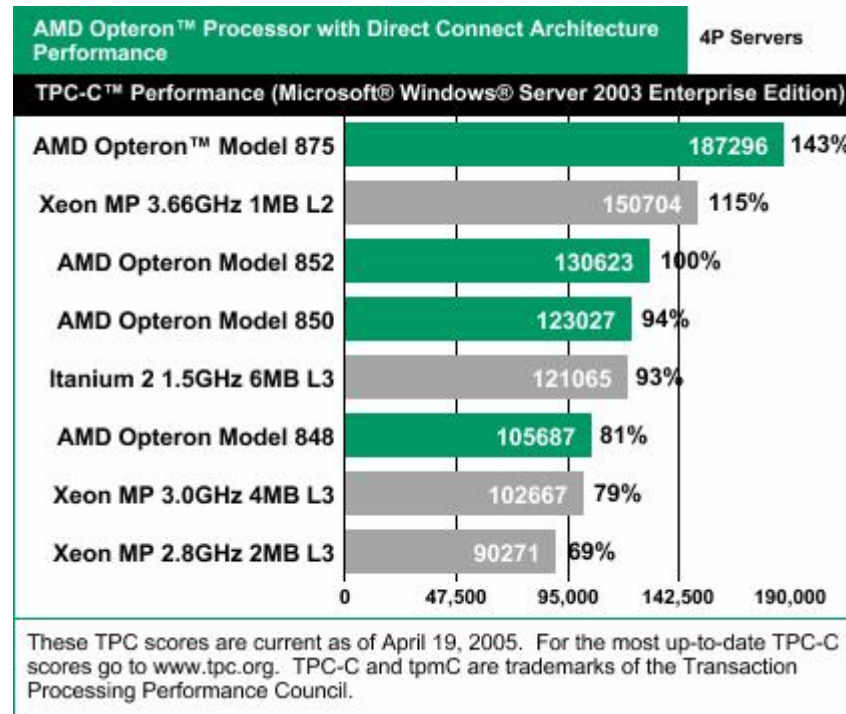
October 2005



10

© Scale Abilities Ltd

AMD/Intel Database Performance



Source: www.amd.com



Capacity – Is Linux only for Clustering?

- Do I need performance greater than four dual-core CPUs?
- Or, do I need availability?
- Or both?

- Don't just leap into clusters, assuming you have a capacity problem!

Reliability: Hardware

- Less parts means fewer failures
- Commodity hardware is very reliable by implication
- The same PCI cards are used across (nearly) all server platforms anyway
 - QLogic
 - Emulex
 - Intel
 - Broadcom
- But hey – where did my console go?

Reliability: Software

- Base Operating System
 - Mostly rock-solid in enterprise kernels
 - Can break ‘around the edges’ at the high end of operating limits
 - Large LUN count
 - Very high/concurrent process counts



Reliability Software: HiAv Products

- NIC/Switch failover
 - ‘bonding.c’ or proprietary
- Storage redundancy
 - Software-based redundancy or hardware-based
- Database Clusters
 - Polyserve Matrix Server with Oracle CRS
 - Red Hat GFS with Oracle CRS
 - Oracle OCFS/ASM with Oracle CRS

HiAv – What is Polyserve?

- Multiple facilities working together
 - Clusterware
 - Node Membership
 - Distributed Lock Manager
 - Virtual Services
 - Clustered Volume Manager
 - Clustered Filesystem
 - Optional: File Serving Option

Support

- Seamless Support
 - Tertiary bilateral support agreements
 - Abstraction layers
- The reality
 - Support is essentially disparate, or at least separate
- *Seamless support can also be a lock-in*
- Certification Matrices
 - Server/OS/HBA/SAN
 - ‘Unbreakable’ support
- Not many experts on the support desk anymore

Support contd.

- Disposable Computing Architecture
 - The box will fail!
 - Design the system to cater for failure
 - ‘Throw away’ hardware disconnects support latency

Support contd.

- Is your vendor really ‘enterprise’ grade?
 - Product lifespan ?
 - Field serviceable?
 - Replacement policy - onsite swap, or attempted fix?
- Memory – who can you buy from?

Pricing – Big Box

- Server Pricing
 - HP DL585 w/4xdualcore CPUs, 4GB memory
 - ~\$25k list
- Database licensing
 - Oracle 10g EE (8x 0.75 CPU licenses for dual core)
 - \$240k list
 - Note: Only EE is permitted with >4 cores
- Total cost: ~\$265k



Pricing – Small Box

- Server Pricing
 - HP DL385 w/2xdualcore CPUs, 4GB memory
 - ~\$7k list
- Database licensing
 - Oracle 10g SE (4x 0.75 CPU licenses for dual core)
 - \$45k list
- Total cost: ~\$52k



Pricing – 2-node ‘Big Box equivalent’ Cluster

- Server Pricing
 - 2 x HP DL385 w/2xdualcore CPUs, 4GB memory
 - ~\$7k list (= \$14k)
- Database licensing
 - Oracle 10g EE (8x 0.75 CPU licenses for dual core)
 - \$240k list
 - Oracle 10g RAC (8x 0.75 CPU licenses for dual core)
 - \$120k list
- Total cost: ~\$374k



Pricing – 2-node ‘Small Box equivalent’ Cluster

– Server Pricing

- 2 x HP DL385 w/1xdualcore CPUs, 4GB memory
 - ~\$6k list (= \$12k)

– Database licensing

- Oracle 10g SE (4x 0.75 CPU licenses for dual core)
 - \$45k list
- Bundled/Free Oracle 10g RAC

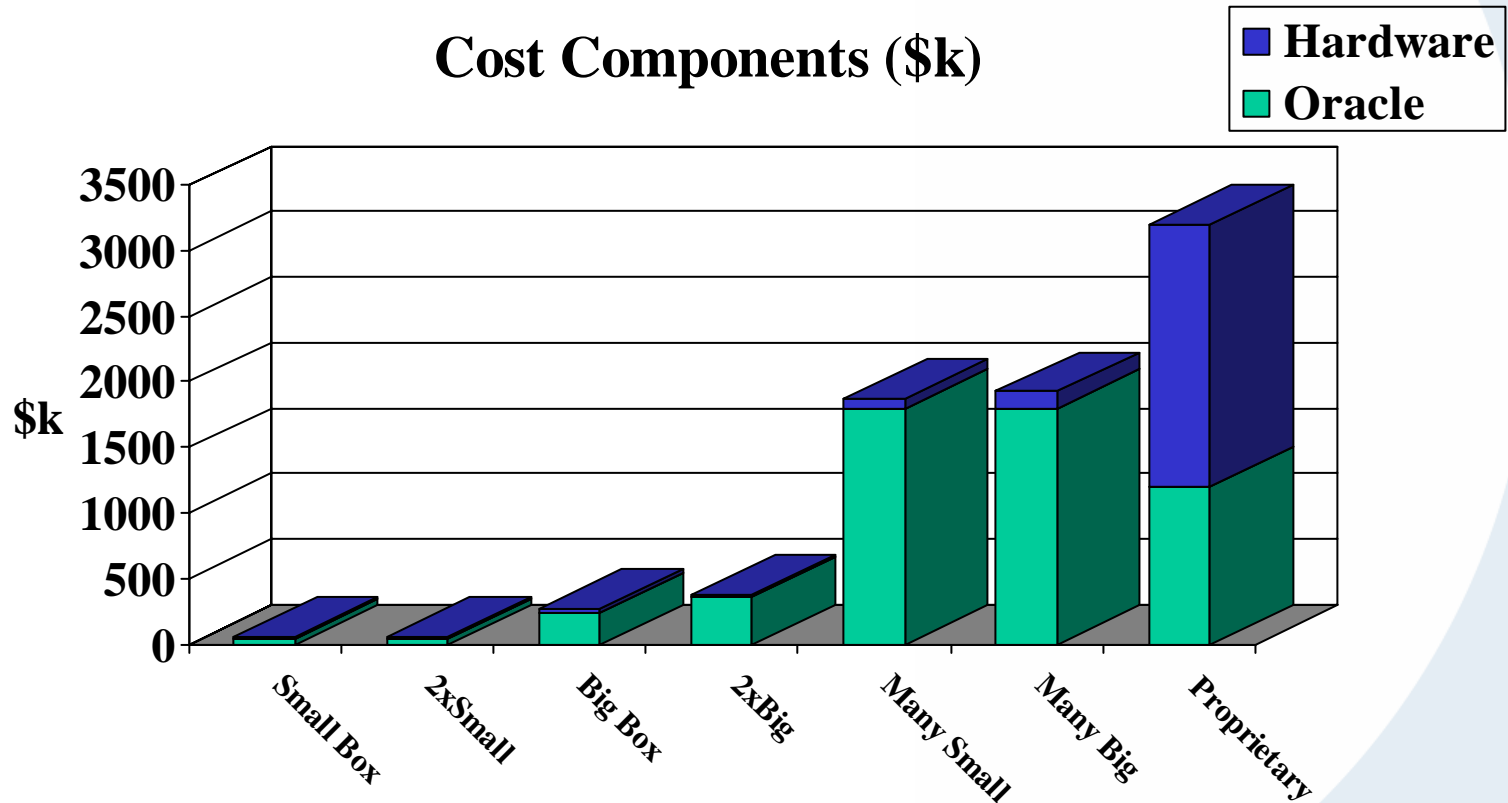
– Total cost: ~\$57k



Pricing – ‘Many Small’ vs ‘Fewer Big’ Cluster

- ‘Many Small’
 - Server Pricing
 - 10 x HP DL385 w/2xdualcore CPUs, 4GB memory
 - ~\$7k list (= \$70k)
 - Database licensing
 - Oracle 10g EE (40x 0.75 CPU licenses for dual core)
 - \$1.2m list
 - Oracle 10g RAC (40x 0.75 CPU licenses for dual core)
 - \$600k list
 - Total cost: ~\$1.87m
- ‘Fewer Big’
 - Server Pricing
 - 5 x HP DL585 w/4xdualcore CPUs, 4GB memory
 - ~\$25k list (= \$125k)
 - Database licensing
 - Oracle 10g EE (40x 0.75 CPU licenses for dual core)
 - \$1.2m list
 - Oracle 10g RAC (40x 0.75 CPU licenses for dual core)
 - \$600k list
 - Total cost: ~\$1.925m

Pricing Summary



Cool Boxes

- My personal ‘hotspots’:
 - Sun x4100
 - 2 Socket Opteron, 1U
 - HP DL385
 - 2 Socket Opteron, 2U
 - HP DL585
 - 4 Socket Opteron, 4U



Enterprise Linux Distros

- Red Hat Enterprise 4
 - Fedora Core
- SLES 9
 - OpenSuSE.org ->10
- White Box Linux
 - Branch of RHEL3
- CentOS-4
 - Built from `ftp://ftp.redhat.com/pub/redhat/linux/enterprise/4/en/os/i386/SRPMS`

Oracle on Linux

- Development Platform as of 10gR2
 - 32-bit currently, 64-bit port
- Oracle contribute to Linux development:
 - OCFS, v1 and v2
 - Firewire storage drivers
- Oracle make good money from Linux



Summary

- Pros
 - Base development platform for Oracle
 - Low hardware costs
 - Commodity support costs
 - Scalable O/S
- Cons
 - Support
 - Oracle is still Expensive
 - ‘Box Bloat’



Contacts

Email: James.Morle@ScaleAbilities.co.uk

Office: +44 1285 644533

Mobile: +44 7867 525555

Skype: morlej

<http://www.scaleabilities.co.uk>



Websites

BAARF Party: *www.baarf.com*

OakTable Network: *www.oaktable.net*

Scale Abilities: *www.scaleabilities.co.uk*

Polyserve: *www.polyserve.com*

SuSE: *www.suse.com*

Red Hat: *www.redhat.com*

White Box Linux: *www.whiteboxlinux.net*