BRUDEN · OSSG
On-Shore Systems Group

Oracle 10g on OpenVMS

Guy Peleg
Senior Member of the Technical Staff
Director of EMEA Operations
guy.peleg@bruden.com

Agenda

- **Oracle 10g RAC Overview**
- Oracle Considerations
- Success Stories
- OpenVMS Installation tips

Short History Lesson

- The concept was first introduced in Oracle 6
- Prior to 9i used to be called Oracle Parallel Server (OPS)
- In 2004 Oracle acquired Compaq's cluster management architecture
- RAC Originally shipped in Oracle 9i
- Impact of the purchase was seen only in 10g

Real Application Cluster

- Multiple Servers (nodes) act as a single (clustered) server.
- Multiple database instances (one on each server) accessing the same database over shared storage
- Clients connect to the various instances in the cluster access and share the data between instances via the cluster interconnect.

- Sounds Familiar ???? ;-)

.

Typical RAC configuration

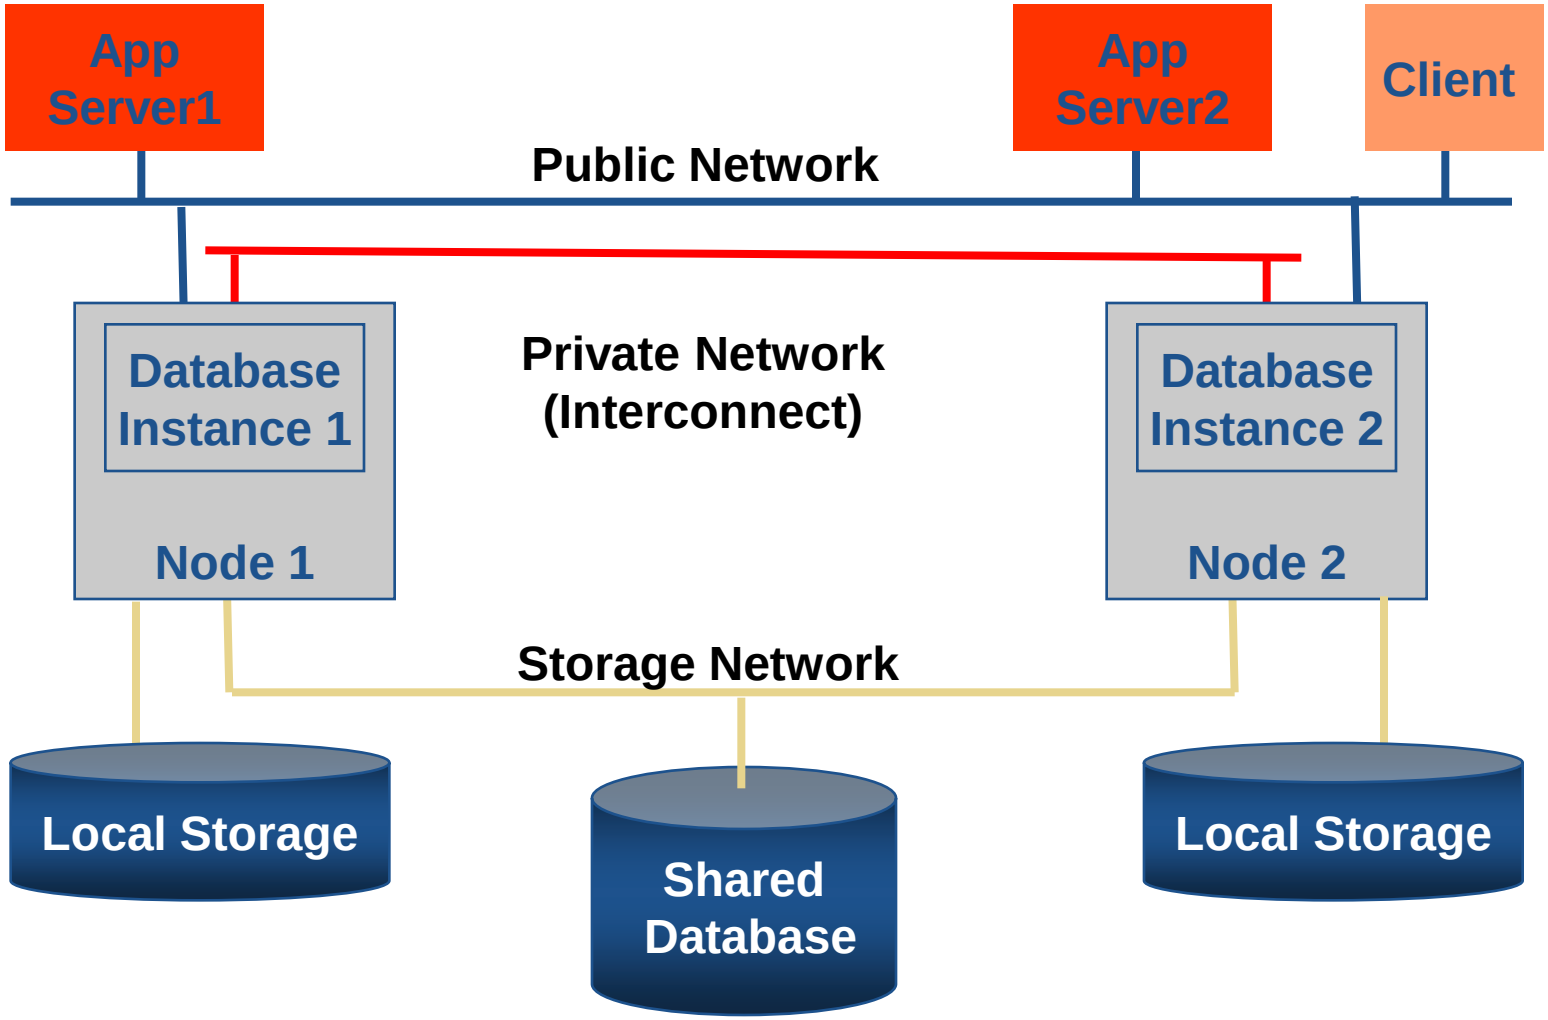

RAC

- Clients typically connect to a virtual TCP/IP address (Virtual database service)
  - Address and service migrate to another node in case of a failover
  - Similar to cluster alias
- Oracle requires (at least one) dedicated cluster interconnect
  - Preferably 1 Gigabit Ethernet using Jumbo frames
  - Interconnect used for
    - Cluster management
    - Locks
    - Cache Fusion

RAC - CRS

- Cluster Ready Services (CRS) is a new feature for 10g RAC
- Provides a standard cluster interface on all platforms
- Should be installed prior to Oracle DB
- Creates 3 processes
  - CRSD
    - Engine for HA related operations
  - OCSSD
    - Cluster membership management
  - EVMD
    - Generate events when “things” happen

RAC & OpenVMS

- Clustering is not a LP on OpenVMS
  - OpenVMS utilizes database specific functions provided by CRS
  - Cluster management is still done by VMS
- BRUDEN performed several RAC installations over the past year, The installation process is not trivial
  - Once installed & configured it just runs….
- TAF is really transparent

Screenshot #1 – CRS Setup

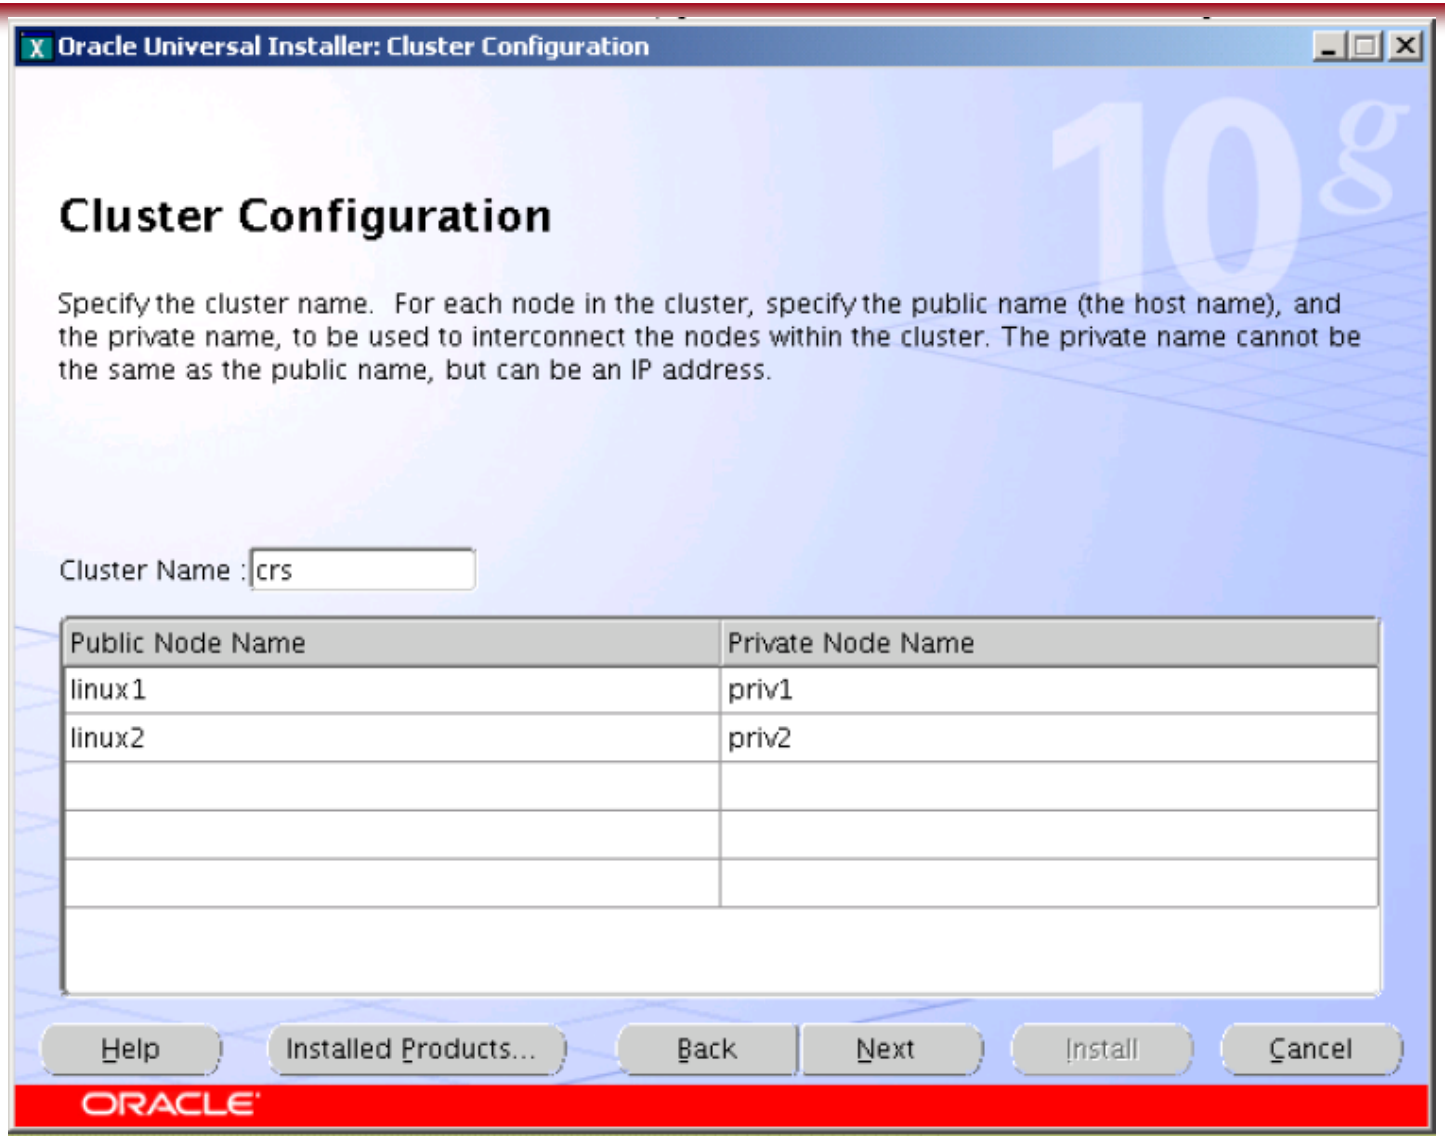

Screenshot #2 – Database installation with RAC

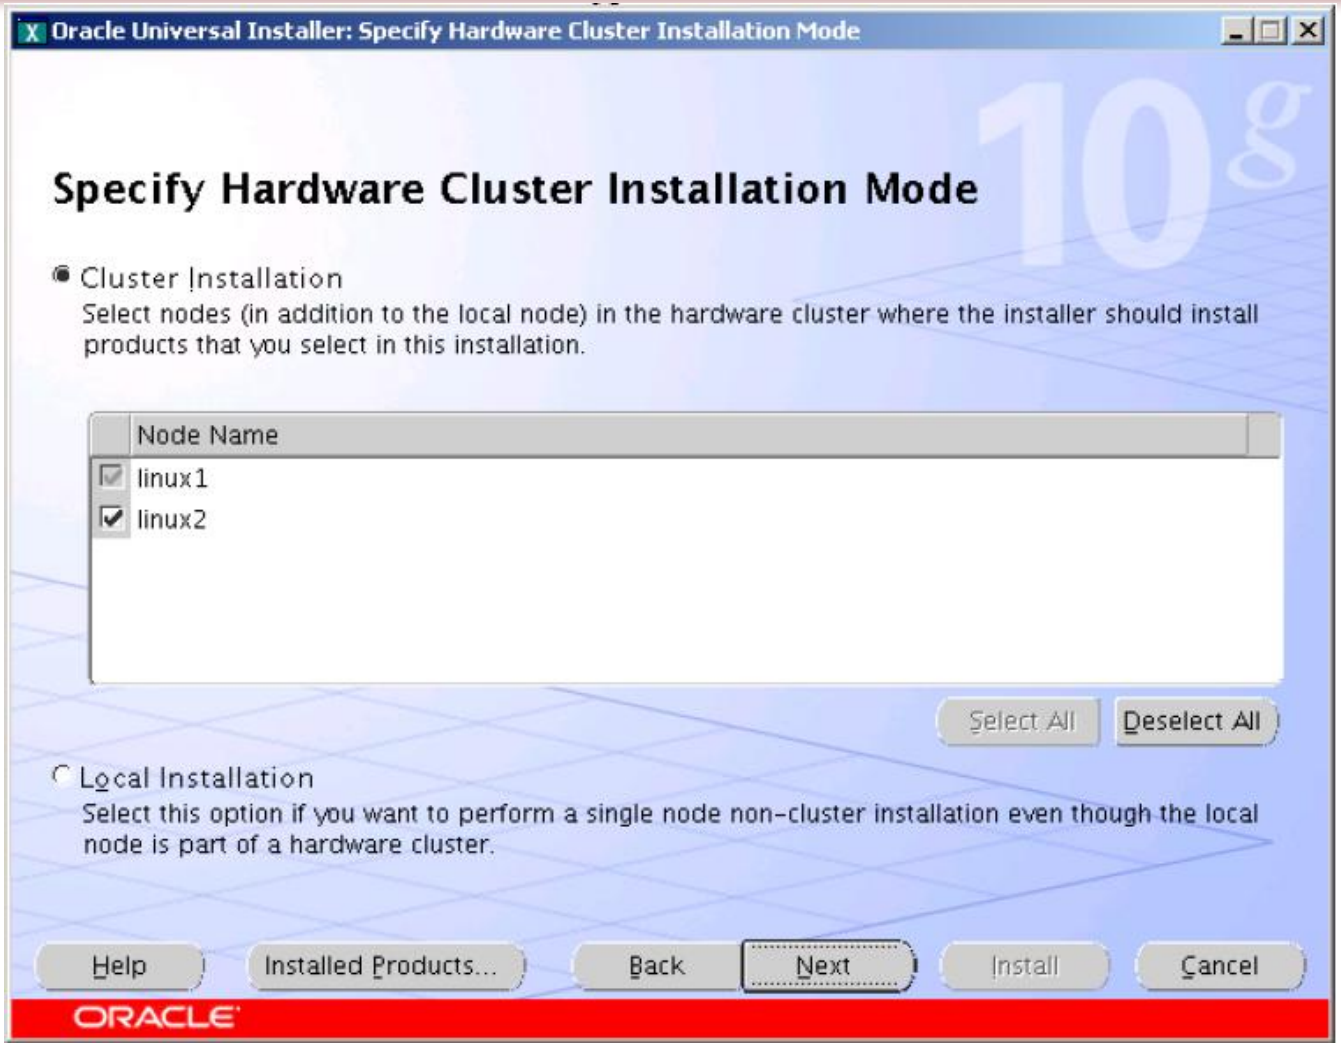

Screenshot #3 – DBCA with RAC enabled

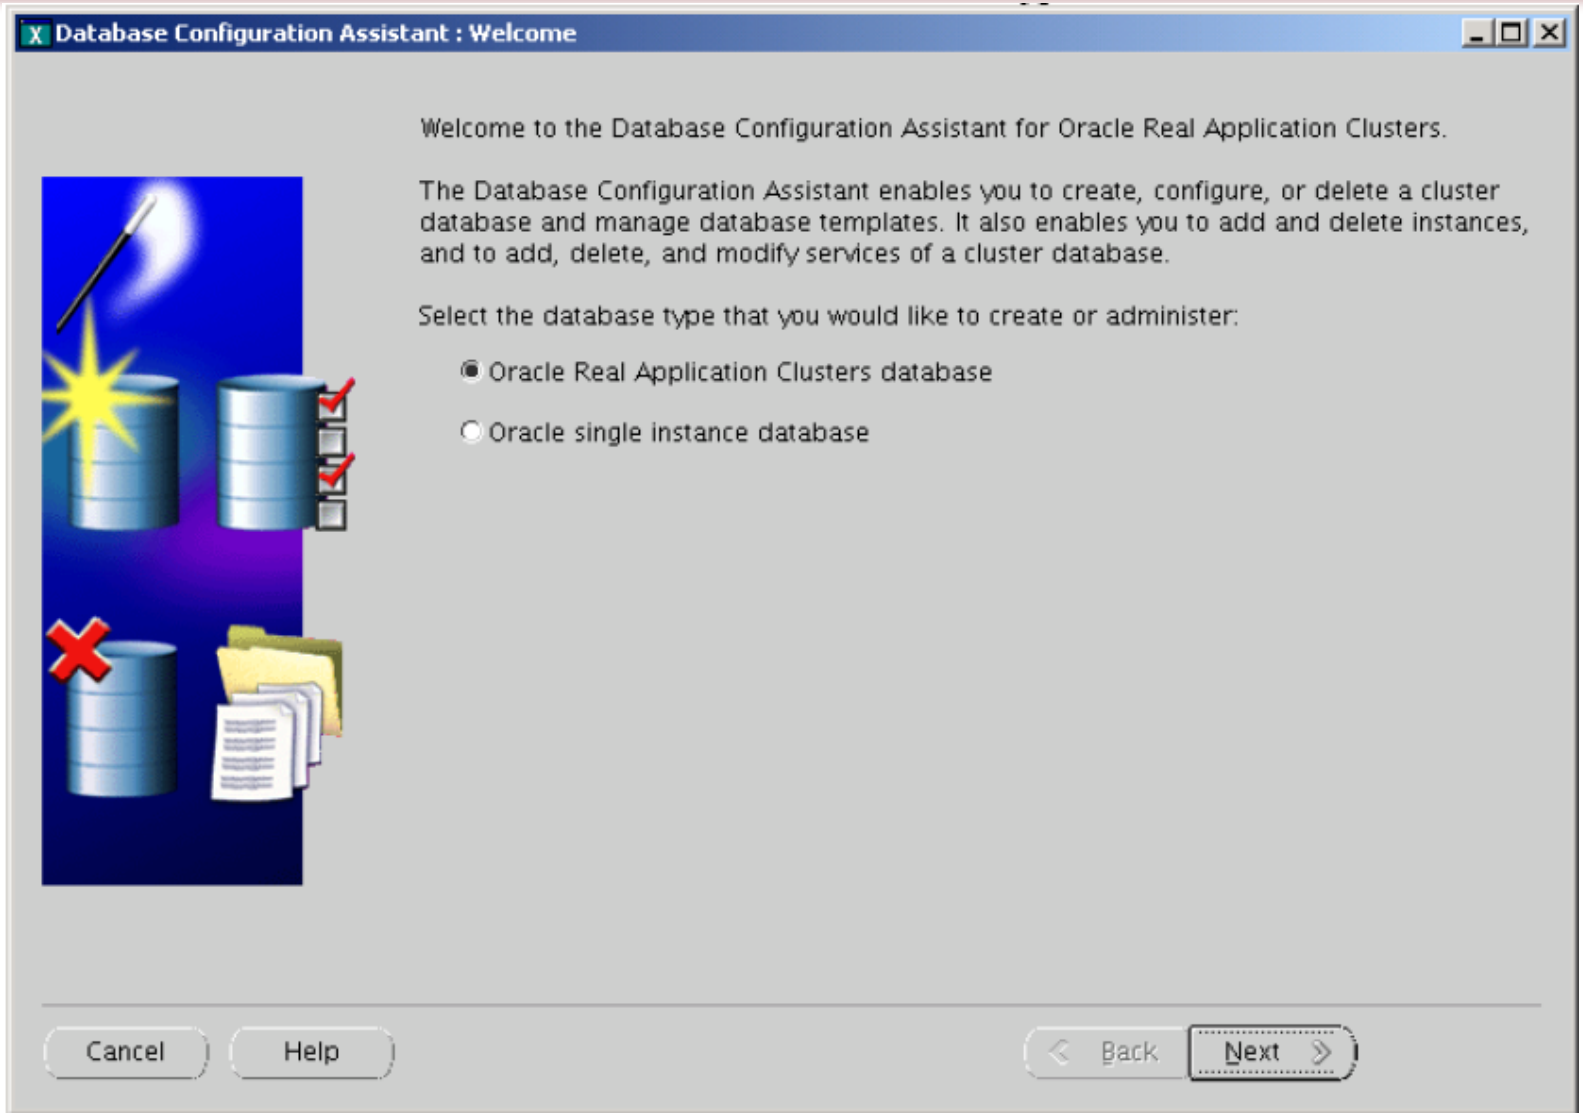

RAC Pros/Cons

- Pros
  - High Availability
  - Scalability
  - Load Balancing
- Cons
  - Increased complexity
  - Increased Cost

Agenda

- Oracle 10g RAC Overview
- **Oracle Considerations**
- Success Stories
- OpenVMS Installation tips

Oracle Considerations

- Oracle 10g is the only version ported to Itanium
- Rule-Based Optimizer not supported with Oracle 10g
  - Cost-Based Optimizer (CBO) is the only supported optimizer
  - May require changes to SQL statements
  - Check initialization parameter OPTIMIZER_MODE
    - Should be set to ALL_ROWS or FIRST_ROWS
- Various performance issues with Oracle ODBC driver due to RBO “leftovers”
  - Most problems fixed (new ODBC driver)
  - Metalink articles available (373129.1)
  - Consider using MS ODBC driver

Performance expectations

- Watch out for long connect time to the database
  - 0.5 seconds instead of 0.05
- According to HP - Performance on Itanium is good
  - HP performed several tests using “out of the box” configurations
  - Itanium outperformed Alpha
    - 20% - 50% faster
    - Lab conditions ;-)
- Oracle builds the product using common sources
  - Performance differences are the impact of the hardware

Oracle comparisons (1000 Iterations)

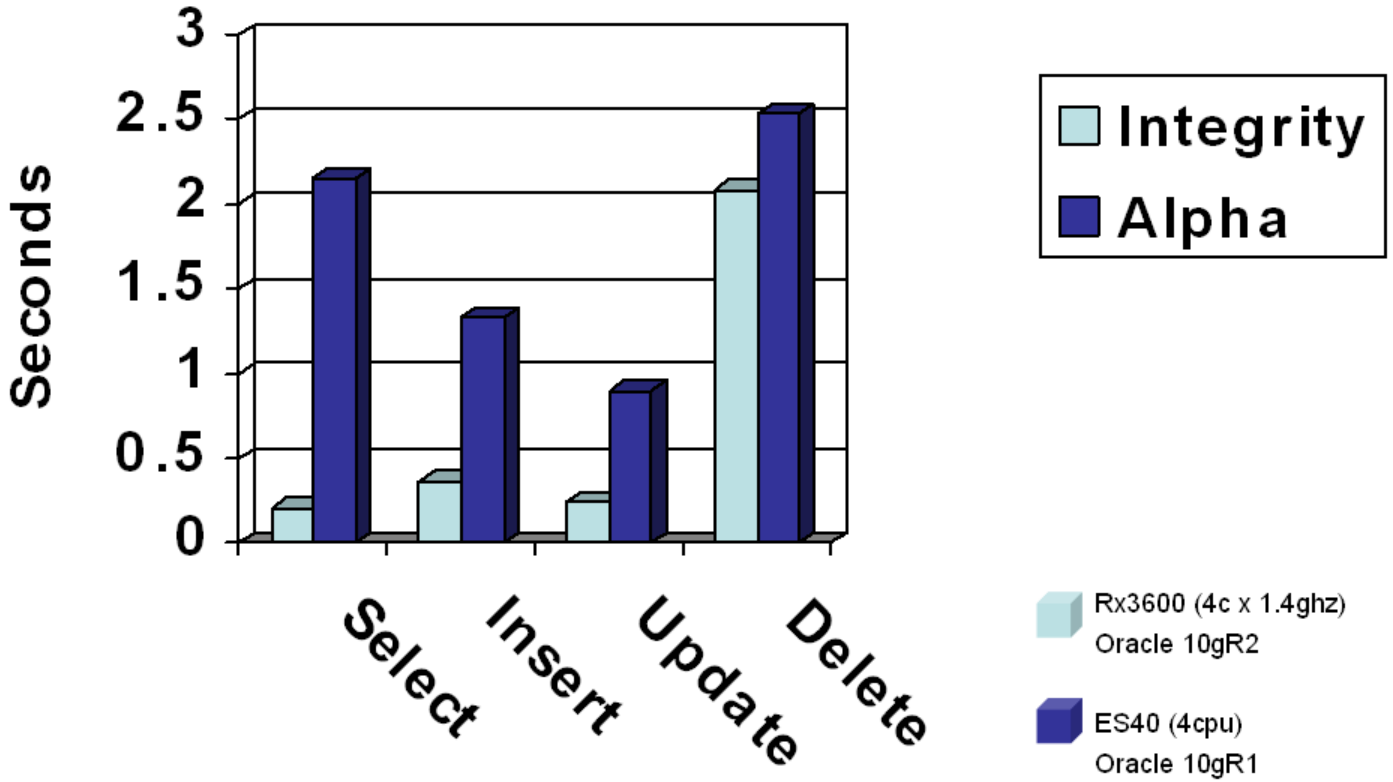

Source: Hewlett Packard

Less is better

Oracle Comparison – 100,000 local iterations

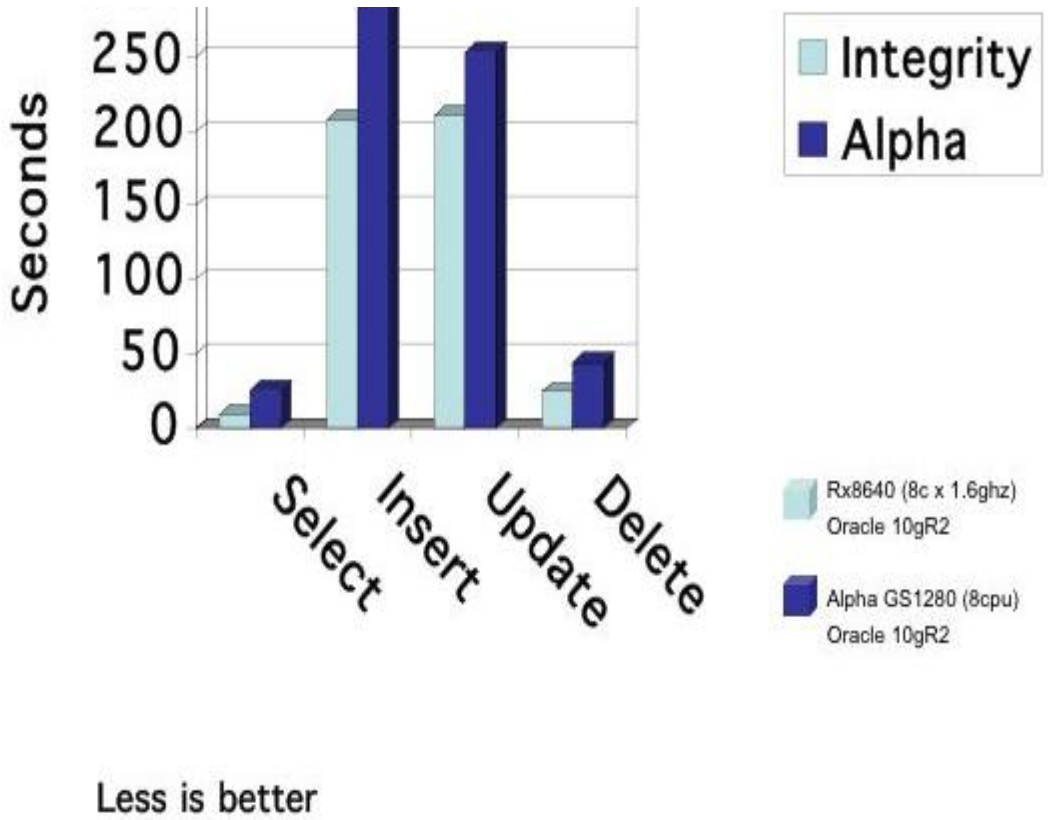

Source: Hewlett Packard

Agenda

- Oracle 10g RAC Overview
- Oracle Considerations
- **Success Stories**
- OpenVMS Installation tips

Oracle on OpenVMS Success Story #1

- Financial Institute in EMEA
- OpenVMS serves as Database Server
- Client application is running on Windows
  - Data is fetched from the database
  - Processed on the PC
  - Stored back in the database
- Hundreds of interactive users during the day
- Critical batch jobs during night
- *Customer needed more performance and throughput*

Hardware Configuration

- Production Environment
  - 8P AlphaServer GS1280 7/1150
  - 1.75MB L2 cache
  - OpenVMS V7.3-2
  - EVA Controller
  - Oracle 8i

- Benchmark Environment
  - 8P/16C RX8640
  - 9MB L3 Cache
  - OpenVMS V8.3
  - EVA Controller
  - Oracle 10gR2

Alpha Vs. Itanium – Initial results

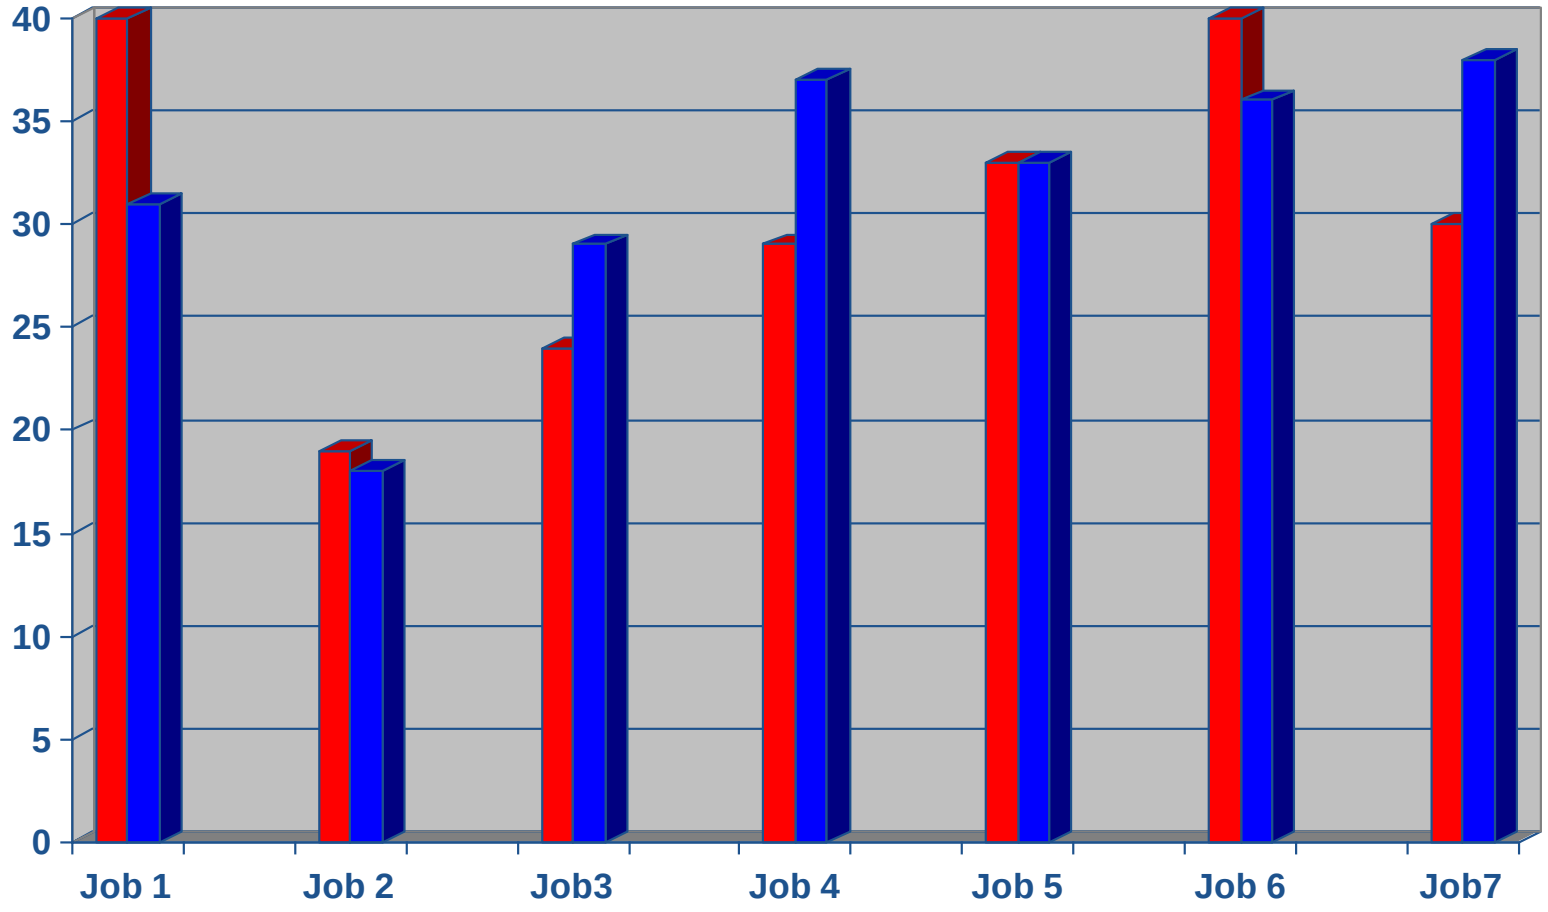

Elapsed time (minutes) to execute 7 jobs
Less is better

Disappointment….

- Initial results did not favor Itanium
  - Itanium ran some jobs slower than Alpha
  - Alpha was running full production load while running the benchmark, Itanium was idle (no production load)
  - Not enough improvement to justify a new purchase
- Customer considered upgrading their existing Alpha
- As Itanium should have outperformed Alpha….the customer still had some hope….

Performance Results Or No Expense

- BRUDEN was called to the rescue

- We started by running a set of general benchmarks to prove the Itanium hardware is capable of outperforming Alpha.

Memory Benchmarks

Memory latency (ns)
less is better

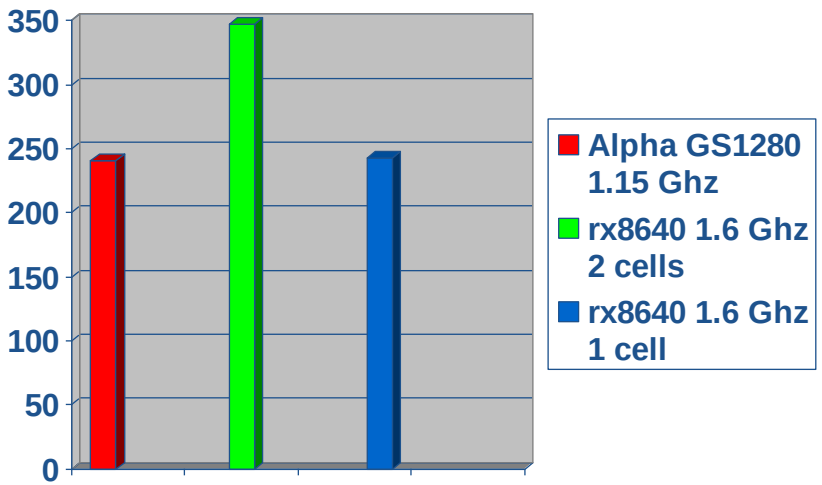

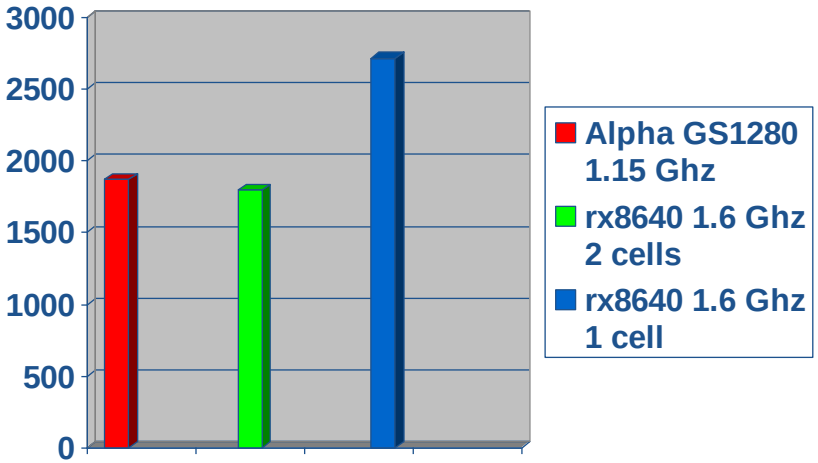

Memory Bandwidth (MB)
More is better

CPU & I/O Benchmarks

Elapsed time for CPU bound application – less is better

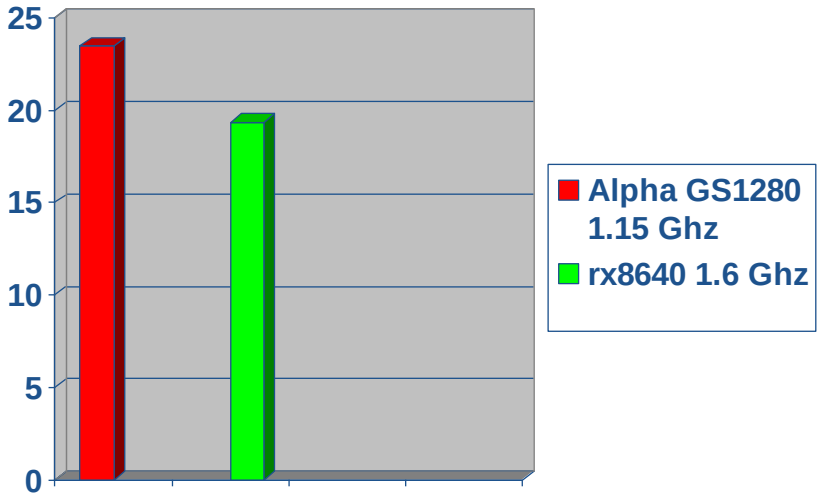

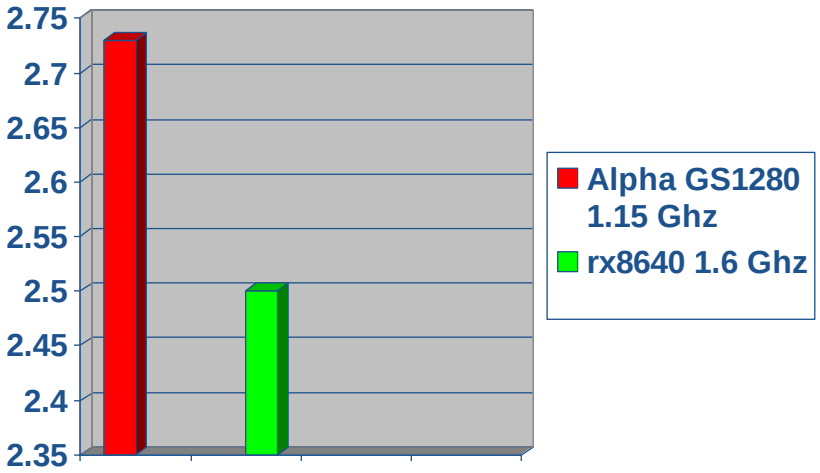

Elapsed time for I/O bound benchmark – less is better

BRUDEN's Approach

- General system tuning
  - Poorly tuned system will impact Oracle's performance
  - Balance I/O paths
  - System parameters
  - Fastpath ports
  - …more
- Mount Data disks /NOCACHE
  - Use SGA instead of XFC
- HyperThreads
- Database parameters
- No changes to SQL statements

Alpha Vs. Itanium – Final results

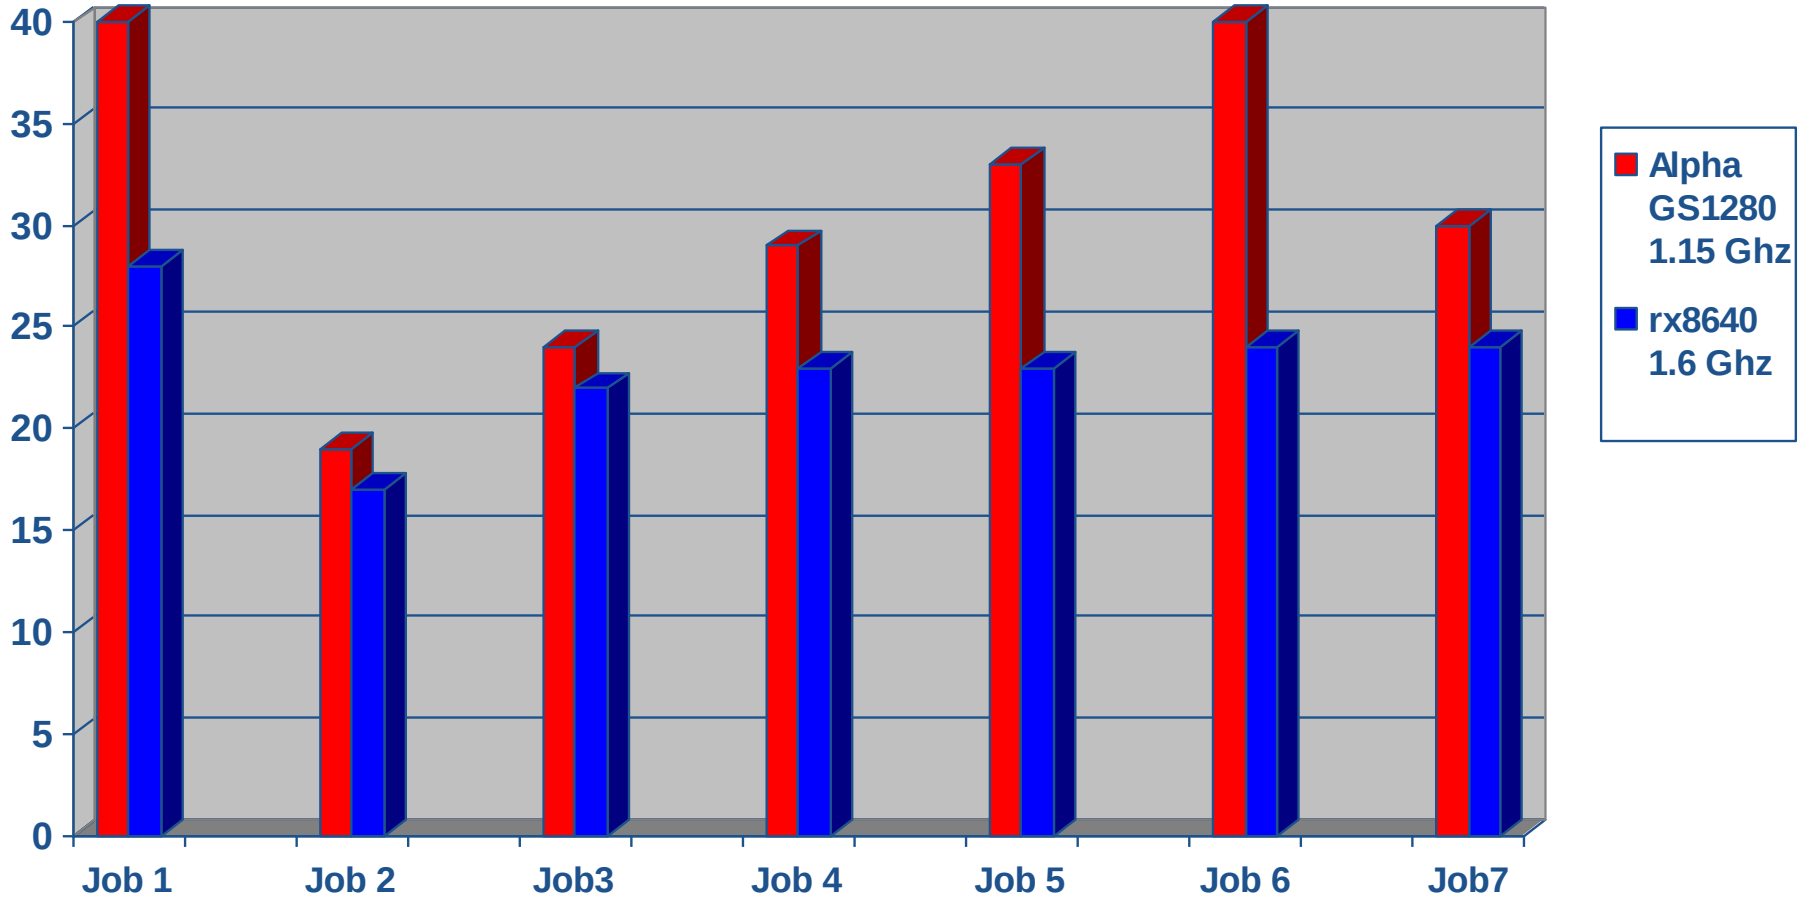

Elapsed time (minutes) to execute 7 jobs – With full production load
Less is better

Alpha Vs. Itanium – Final Results

% of CPU Utilization
less is better

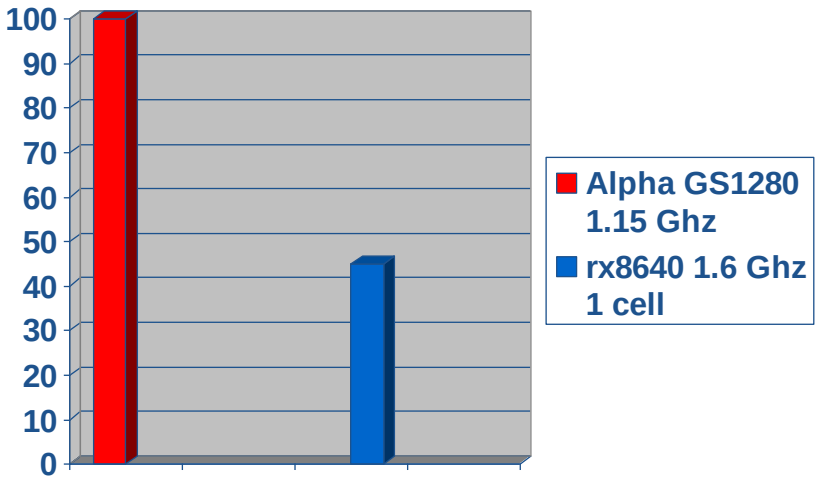

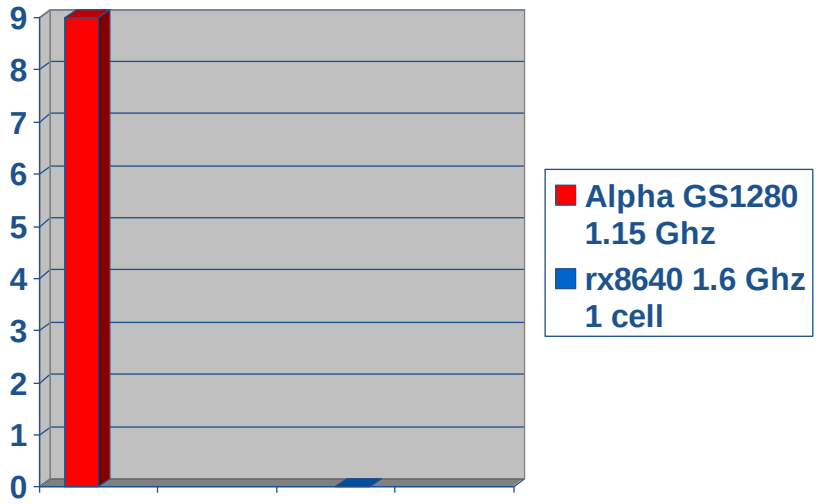

Length of Compute Queue

Less is better

Alpha Vs. Itanium - Summary

- Itanium is running every single job faster
  - 10% - 40% performance increase per job
- Throughput has been **DOUBLED !!**
  - Batch Cycle reduced significantly
  - Improved response time for interactive users

Hyperthreading with Stalls vs Hyperthreading with No Stalls

Serial Execution with Stalls (no Hyperthreading)

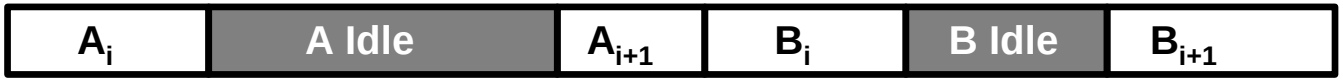

Hyperthreading with Stalls

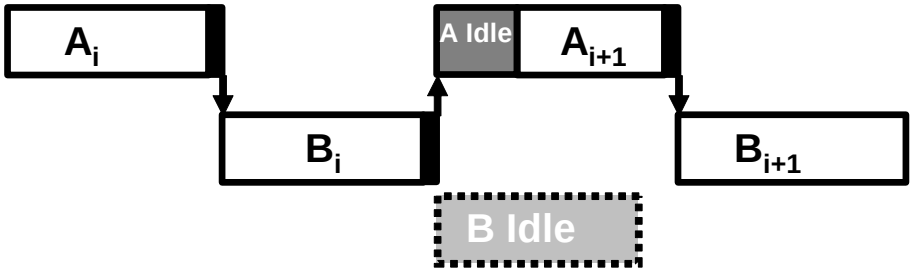

Serial Execution with No Stalls (no Hyperthreading)

Hyperthreading with No Stalls

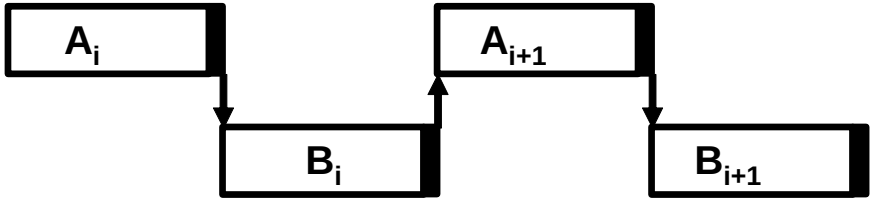

Two Cores vs Hyperthreading (NoStalls)

Serial Execution with No Stalls on Two Cores

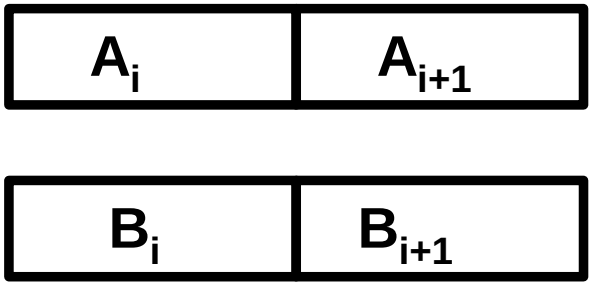

Hyperthreading with No Stalls

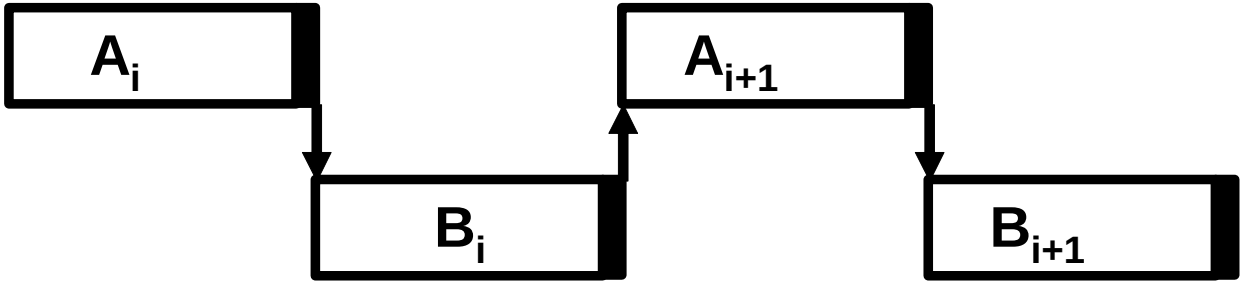

Sample Run on Single Core

Session 1

```
$ set process/aff/perm/set=31
$ r tc_cf
cycle interval = 0.00000000250000 seconds
cycles 10282355676
time to execute == 25.70588919000000 seconds
Wall clock time: 49143ms
System CPU time: 25620ms
```

Session 2

```
$ set process/aff/perm/set=31
$ r tc_cf
cycle interval = 0.00000000250000 seconds
cycles 10353364312
time to execute == 25.88341078000000 seconds
Wall clock time: 49225ms
System CPU time: 25900ms
$
```

Sample Run on Single Core with CoThreads

Session 1

```
$ set process/aff/perm/set=31
$ r tc_cf
cycle interval = 0.00000000250000 seconds
cycles 14092713701
time to execute == 35.23178425250000 seconds
Wall clock time: 35322ms
System CPU time: 18640ms
$
```

Session 2

```
$ set process/aff/perm/set=63
$ run tc_cf
cycle interval = 0.00000000250000 seconds
cycles 14346088453
time to execute == 35.86522113250000 seconds
Wall clock time: 35950ms
System CPU time: 19260ms
$
```

Sample Run on Two Cores

Session 1 (Still affinitized to CPU 31)

```
$
$ r tc_cf
cycle interval = 0.00000000250000 seconds
cycles 9572697013
time to execute == 23.93174253250000 seconds
Wall clock time: 23997ms
System CPU time: 23920ms
$
```

Session 2 (Non-Co-Thread CPU)

```
$ set process/aff/perm/set=32/clear=63
$ stop/cpu 63
%SMP-I-CPUTRN, CPU #63 was removed from the active set.
$
$ r tc_cf
cycle interval = 0.00000000250000 seconds
cycles 9726287048
time to execute == 24.31571762000000 seconds
Wall clock time: 24398ms
System CPU time: 24290ms
$
```

The HyperThreads Advantage

- According to Intel's marketing numbers, HTs provide up to 25% performance increase
- Applications with poor locality are good candidates to benefit from HyperThreads
  - An Oracle application potentially fits the profile
- Not suitable for all applications
  - Biggest degradation in performance we've seen was 14%
- **YMMV !!!!**
  - **Tuning is critical !**

HyperThreads – Impact on Oracle Jobs

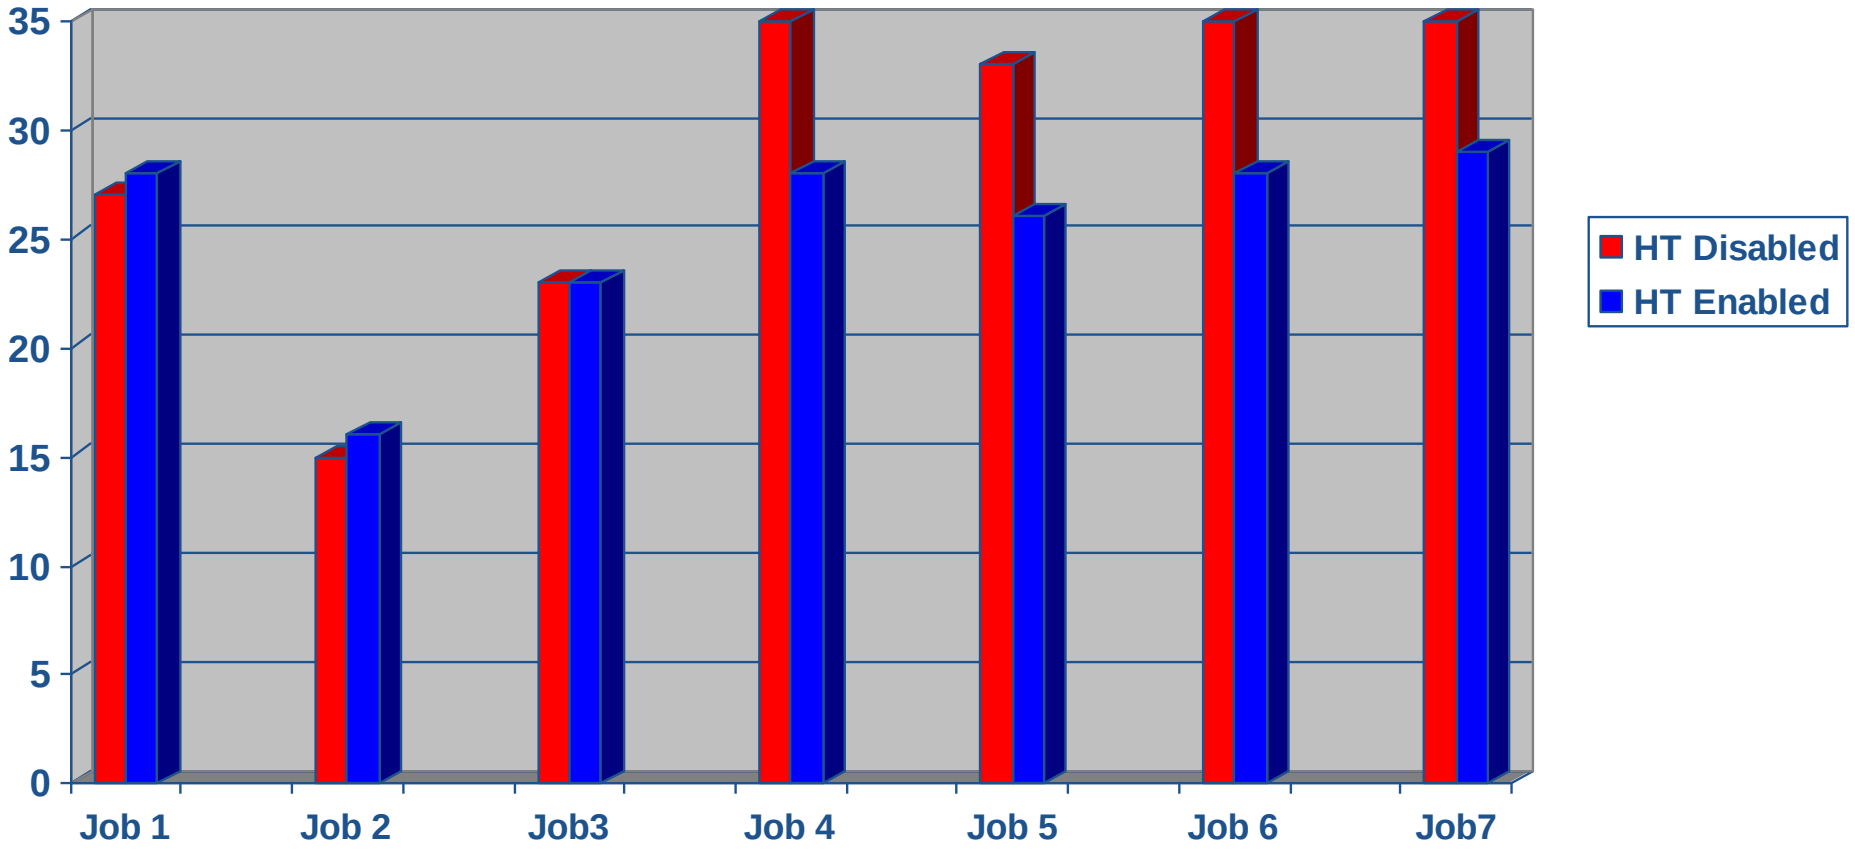

Elapsed time (minutes) to execute 7 jobs
Less is better

Oracle on OpenVMS Success Story #2

- Retail application – EMEA customer
- Several Alpha Servers GS1280 7/1150 & 7/1300
- Oracle 10gR2
- Application is 100% SQL statements

- Customer benchmarked rx7640
  - Main goal was to reduce the length of the nightly batch cycle

- BRUDEN tuned the rx7640 for optimal performance

Batch job X

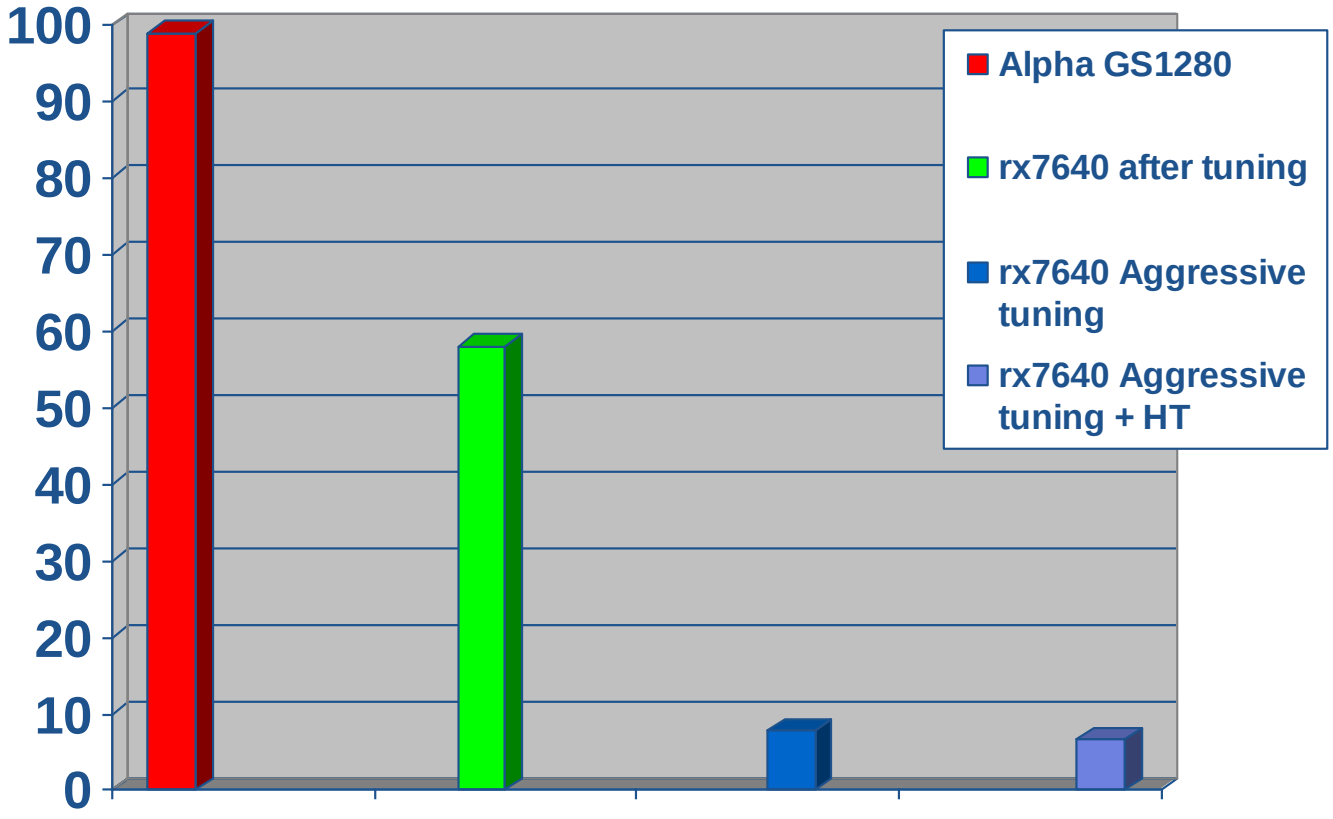

HyperThreads improve Performance by 15.7 %

Minutes to complete a batch job (100% SQL)

Less is better

Simulated Nightly Cycle

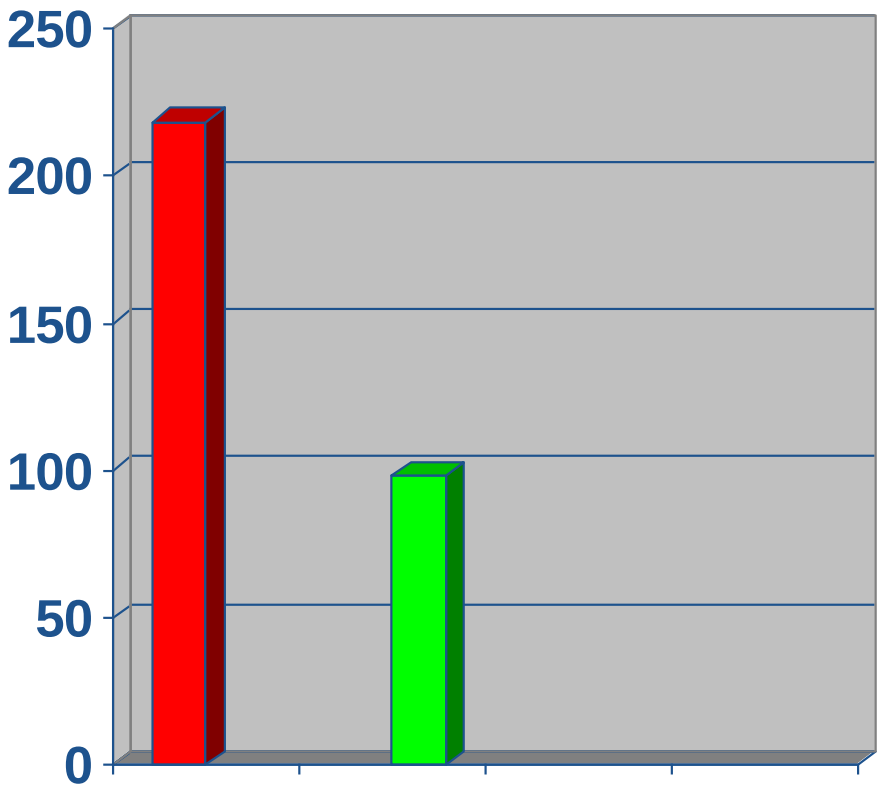

Minutes to complete simulated nightly cycle

Less is better

Alpha is 2.22 times slower than IA64

Agenda

- Oracle 10g RAC Overview
- Oracle Considerations
- Success Stories
- **OpenVMS Installation tips**

RAC Installation tips

- **Read the installation manual**
- Read it again
- SCSNODE must be equal to TCPIP$HOST_NAME
- Identical network interfaces on all RAC nodes
- C++ has to be installed
- All disks should have allocation class assigned
- The “ORACLE” disk and scratch disk must be ODS-5
- All file specifications must use logical volume names
  - avoid using $ sign in the device name
- The free space on the scratch disk may be calculated wrong
  - avoid using “too large” scratch devices
- Watch out for initialization parameters not supported on OpenVMS

Summary

- Oracle 10g RAC increases the availability of your database.
  - Allows scaling outside of the box
- Oracle on OpenVMS Itanium outperforms Oracle on Alpha
  - Memory bandwidth
  - Larger caches on the CPU
  - HyperThreads
- Tuning is critical for achieving optimal performance
  - Don’t run “out of the box”
- Running Oracle on Alpha? Switching to Itanium should be your growth path.

Questions?

BRUDEN-OSSG thanks you for attending this session.

See us at www.BRUDENOSSG.com for:

- *Performance analysis*
  - *(Performance Results Or No Expense)*
- *Oracle Tuning*
- *Porting assistance*
- *Special OPS (OpenVMS Programming Services)*