Oracle9*i*®
Real Application Clusters
HP Itanium2 Servers

Steve Gonzalez

Senior Solution Architect
Hewlett-Packard

HP WORLD 2003
Solutions and Technology Conference & Expo

Agenda

- Multi-boot Itanium2 Servers
- Oracle9*i* ® RAC Functionality
- Cluster Technologies
- The Demo

O/S Support

Which of the following operating systems are supported by HP on Itanium2 servers?

- HP-UX 11i
- OpenVMS
- CPM
- Windows Server 2003
- RSX-11m
- Red Hat Enterprise Linux AS
- All of the above
- None of the above

O/S Support

Which of the following operating systems are supported by HP on Itanium2 servers?

- [x] HP-UX 11i
- [x] OpenVMS
- [ ] CPM
- [x] Windows Server 2003
- [ ] RSX-11m
- [x] Red Hat Enterprise Linux AS
- [ ] All of the above
- [ ] None of the above

Itanium2 Boot Path

Management Processor
(Console Interface)

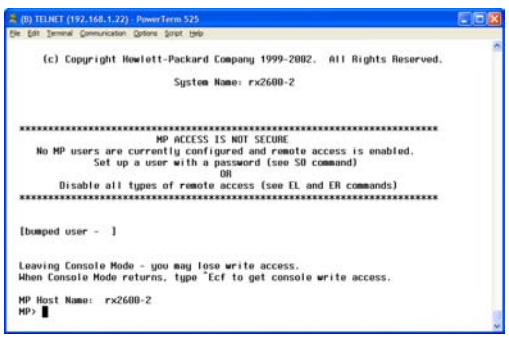

Extensible Firmware
Interface (EFI)

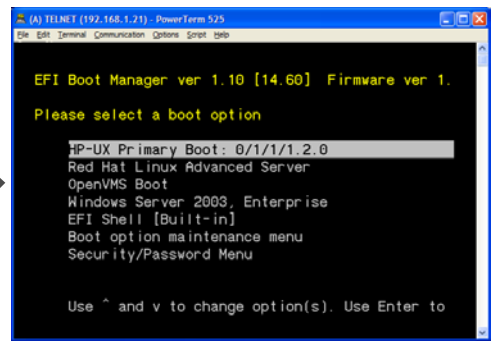

HP-UX 11i

OpenVMS

OpenVMS

redhat

Microsoft Windows Server 2003 Datacenter Edition

Microsoft Windows Server 2003 Enterprise Edition

HP Integrity rx2600

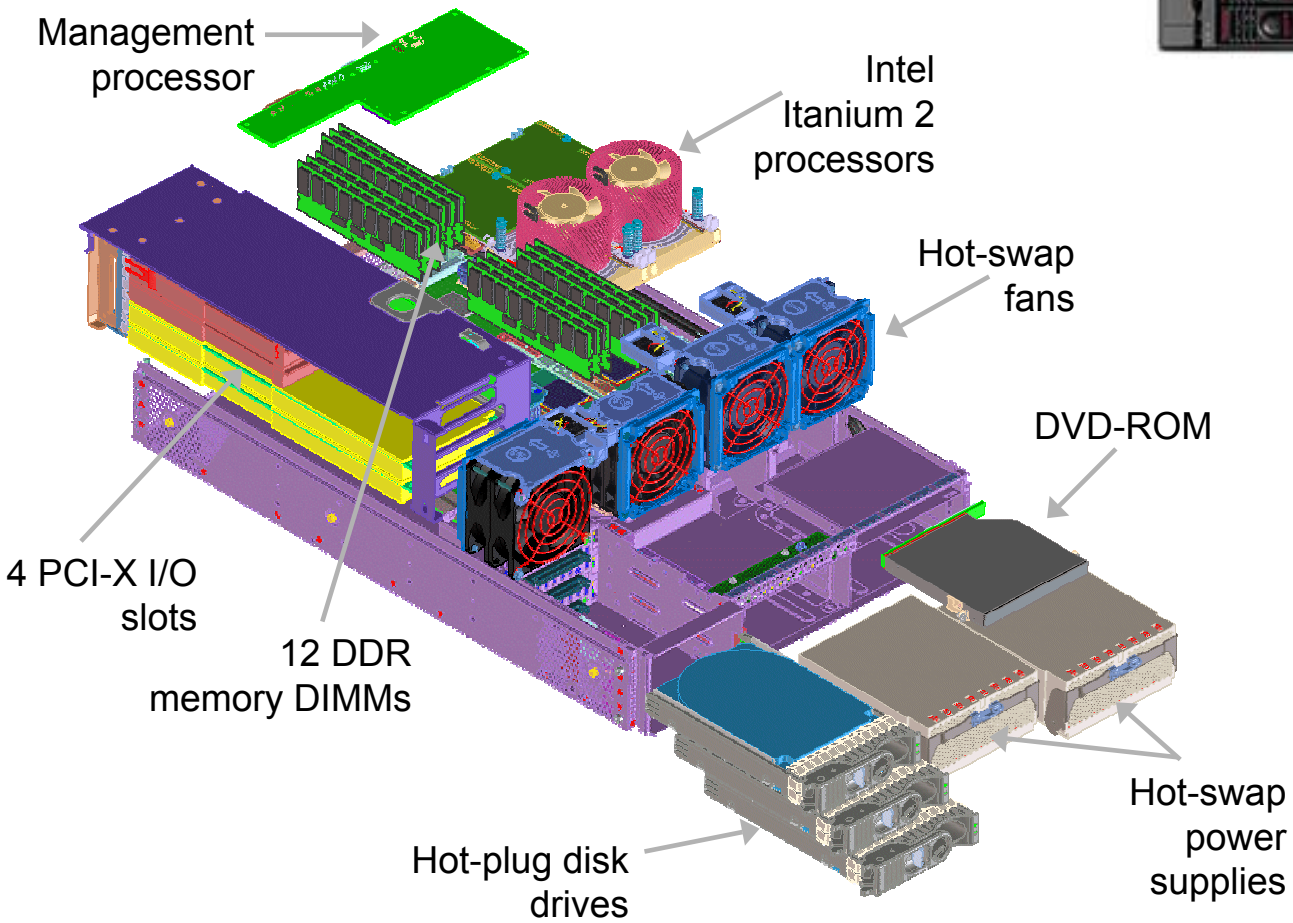

Processors and chipset

- 1 or 2 Intel Itanium 2 processors
- HP zx1 Chipset
- 128-bit, 200 MHz bus
- 6.4 GB/s system bandwidth

Memory

- 1 GB to 24 GB
- PC2100 ECC chip spare DDR
- 128-bit, 266 MHz DDR
- 8.5 GB/s bandwidth

I/O subsystem

- 4 PCI-X full-length slots
  - One 133 MHz (1 GB/s)
  - Three 133 MHz (512 MB/s)
- Ultra320 SCSI, Gigabit, 100Base-T USB, VGA, serial
- 4.0 GB/s bandwidth

HP Integrity rx5670

HP WORLD 2003
Solutions and Technology Conference & Expo

Standard

- 1 to 4 Intel Itanium 2 processors
- HP zx1 Chipset
- 6.4 GB/s system bandwidth
- 12.8 GB/s memory bandwidth
- 10 PCI-X I/O slots
- Independent I/O channels
- 4.0 GB/s I/O bandwidth
- Gigabit LAN and Ultra160 SCSI
- 7U (5 per 2-meter rack)
- Rack-optimized and pedestal
- Management processor
- HP-UX, Windows, and Linux support

Growth and connectivity

- Up to 96 GB error correcting and checking DDR memory
- 4 internal hot-plug disks
- DVD or DAT
- VGA, USB, RAID

High availability

- Hot-plug disks
- Redundant, hot-swap fans and power
- Memory scrubbing and page de-allocation
- CPU failure de-allocation
- Memory chip spare
- HP MC/Serviceguard support

Investment protection

- In-box upgrade from . . .
  - rp5470, rp5430
  - rp5450, rp5400
  - HP 9000 L-Class

Itanium2, 1GHz – 87,741 tpmC @ $5.03
Madison, 1.4GHz – 121,065 tpmC @ $4.97

EFI Boot Selection Menu

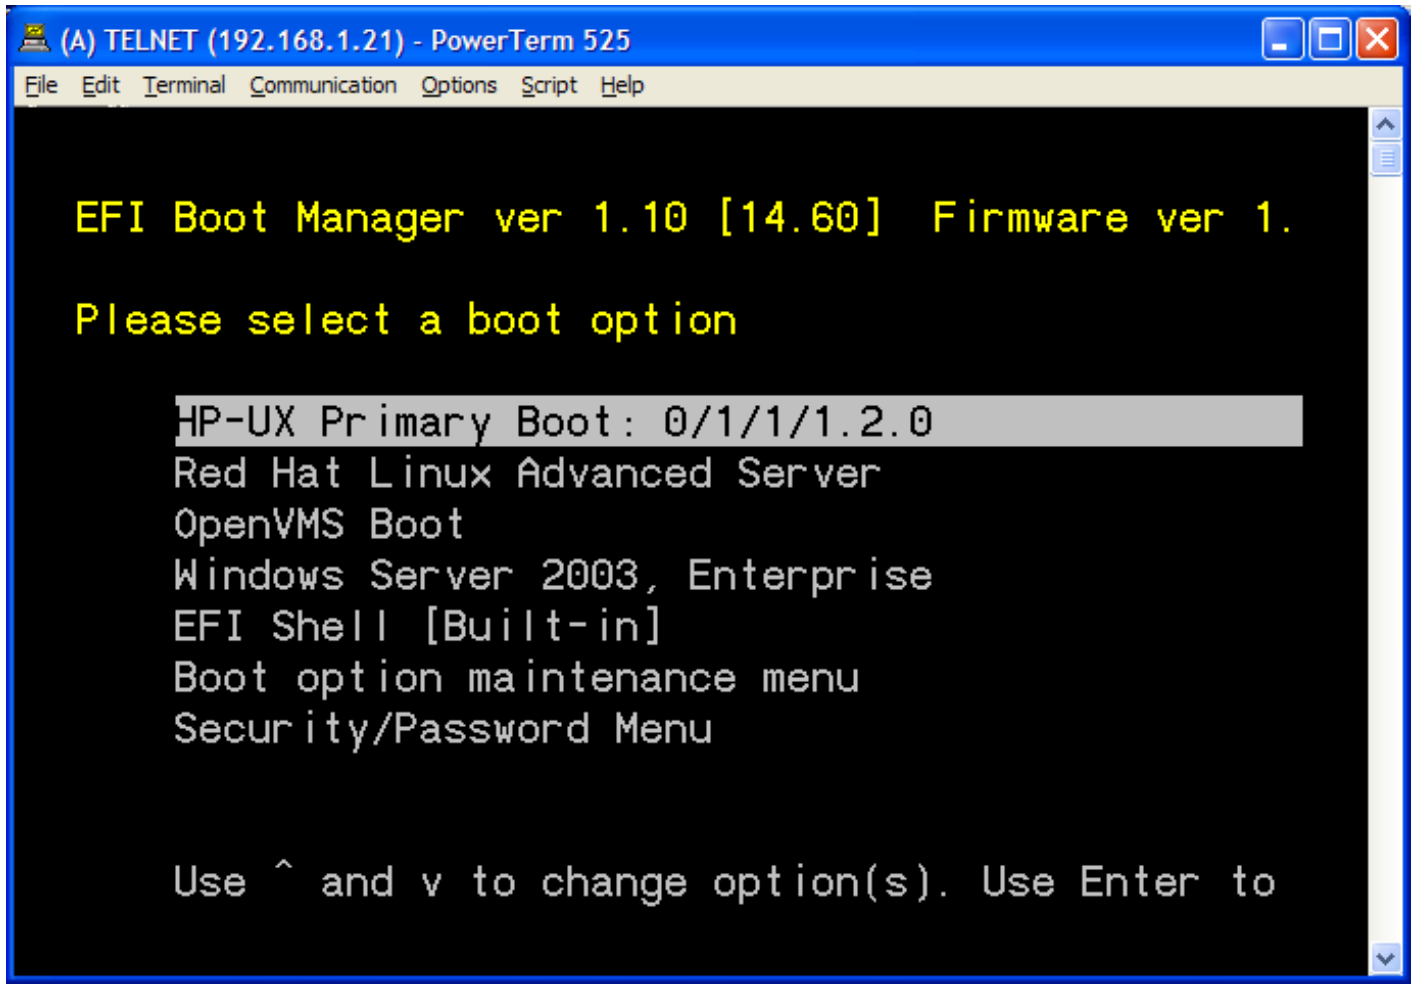

64-way MP on HP Itanium2

Microsoft Windows Server 2003

HP WORLD 2003 Solutions and Technology Conference & Expo

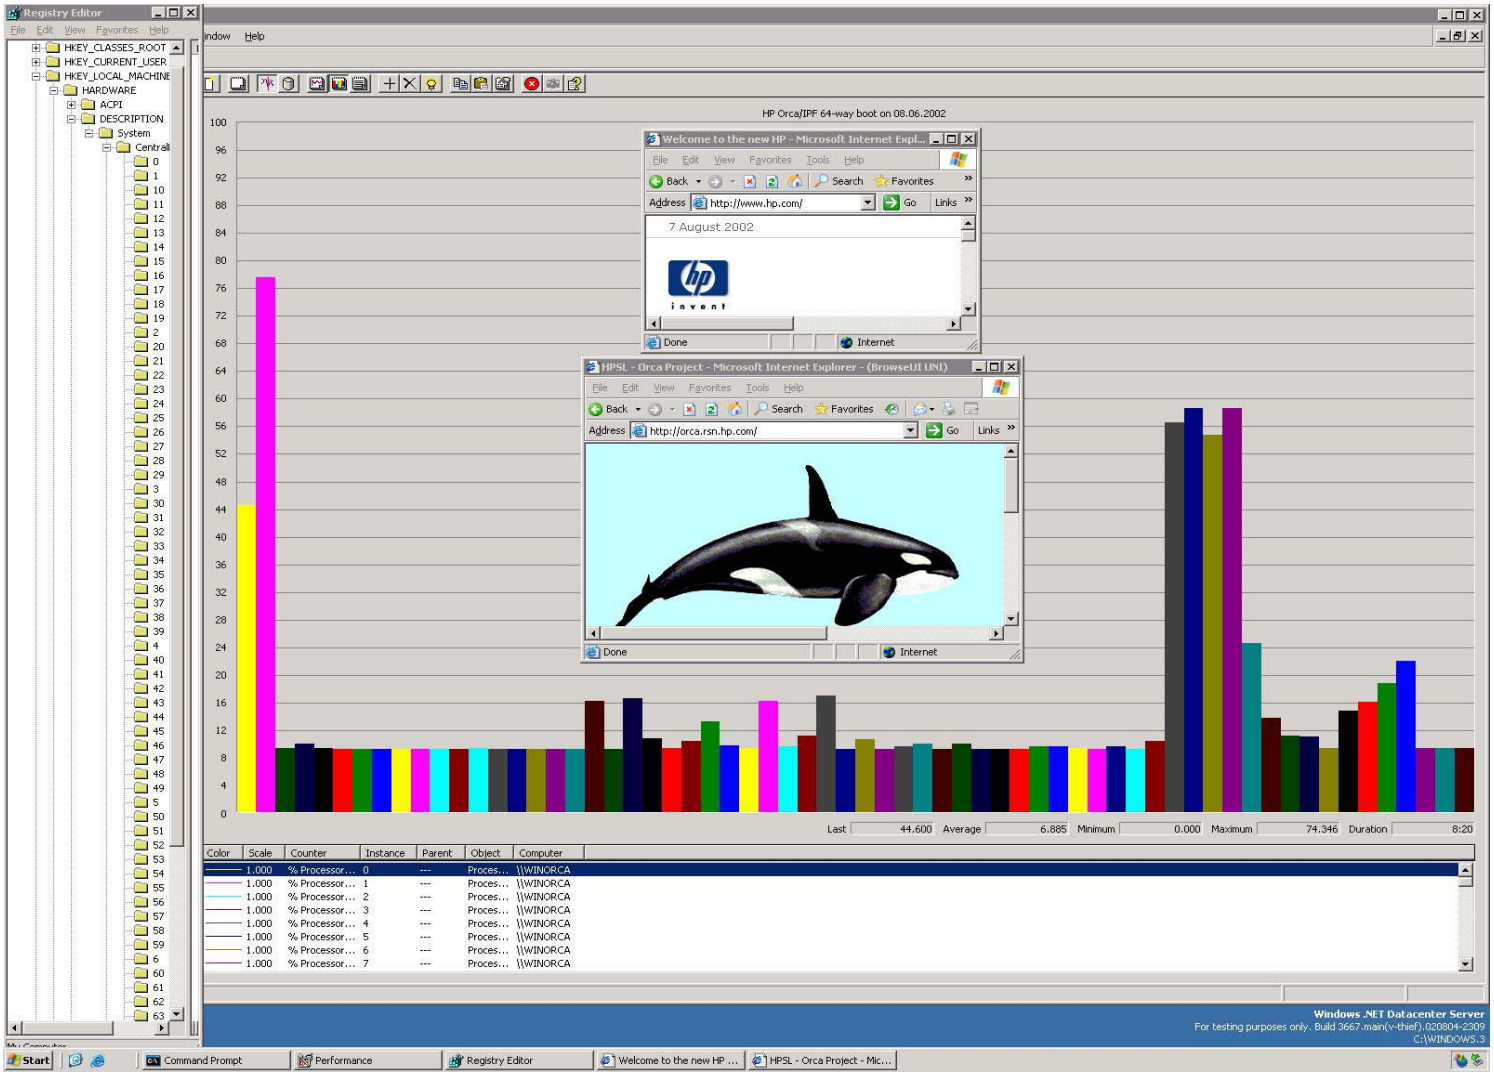

64-way MP on HP Itanium2

Microsoft Windows Server 2003

HP WORLD 2003 Solutions and Technology Conference & Expo

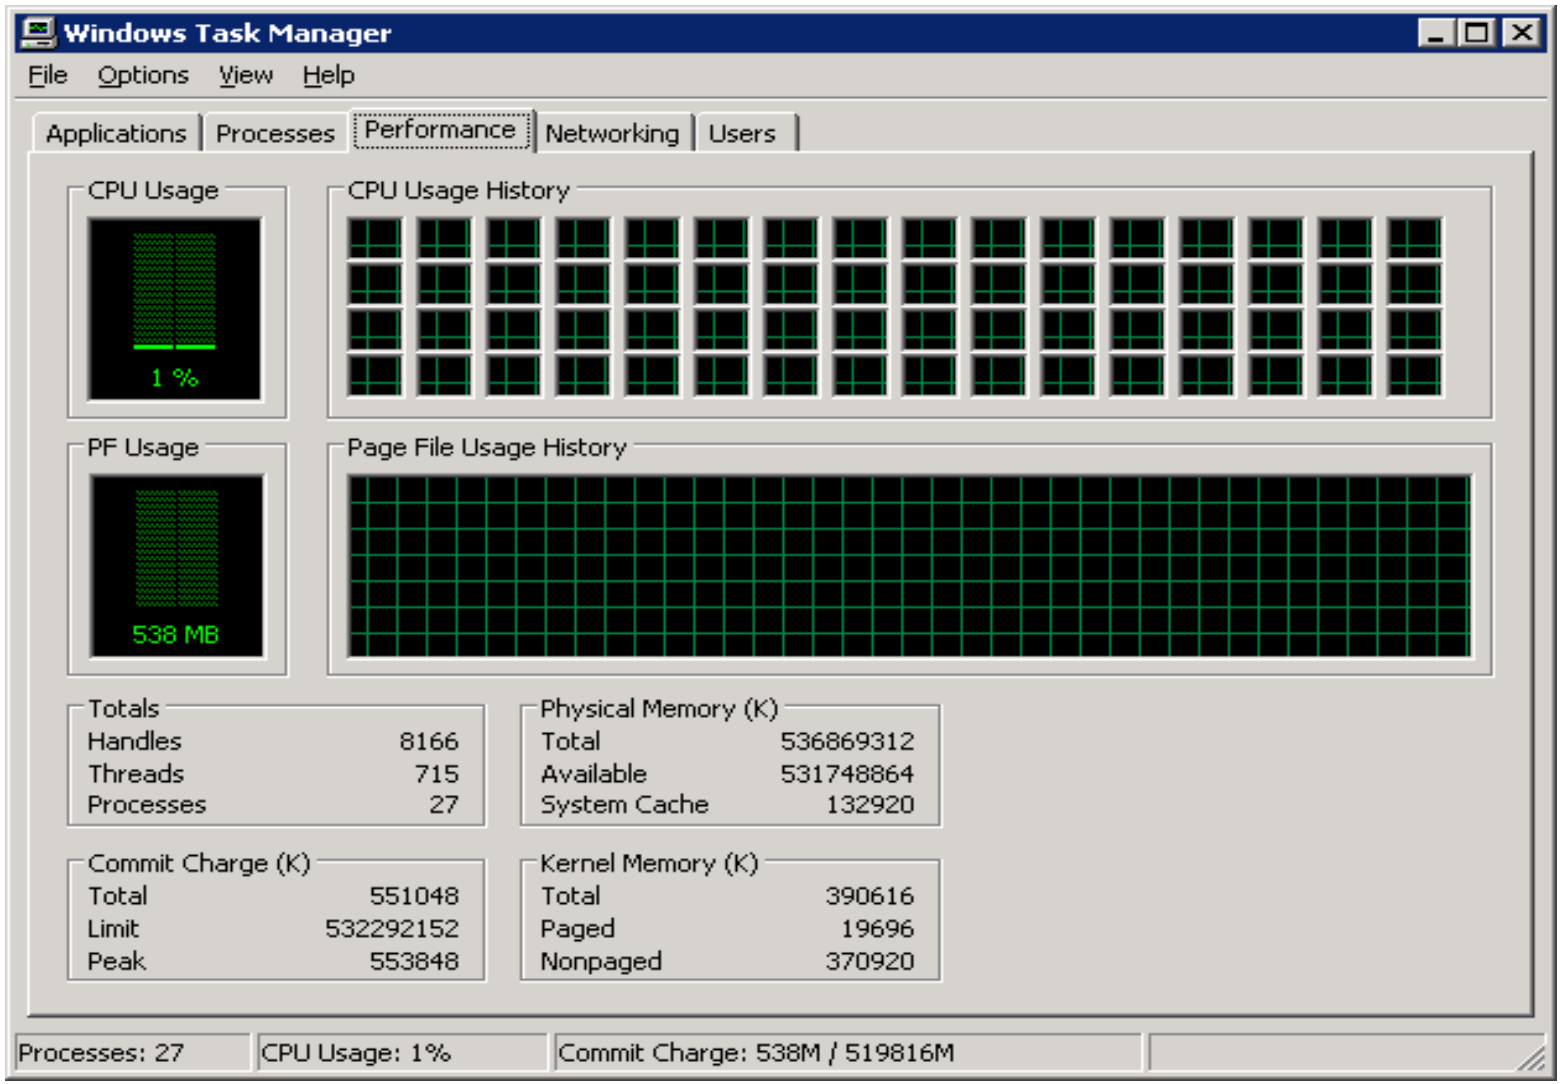

Real Application Clusters

HP WORLD 2003
Solutions and Technology Conference & Expo

grow your DATA

grow your USERS

grow PROCESSING POWER

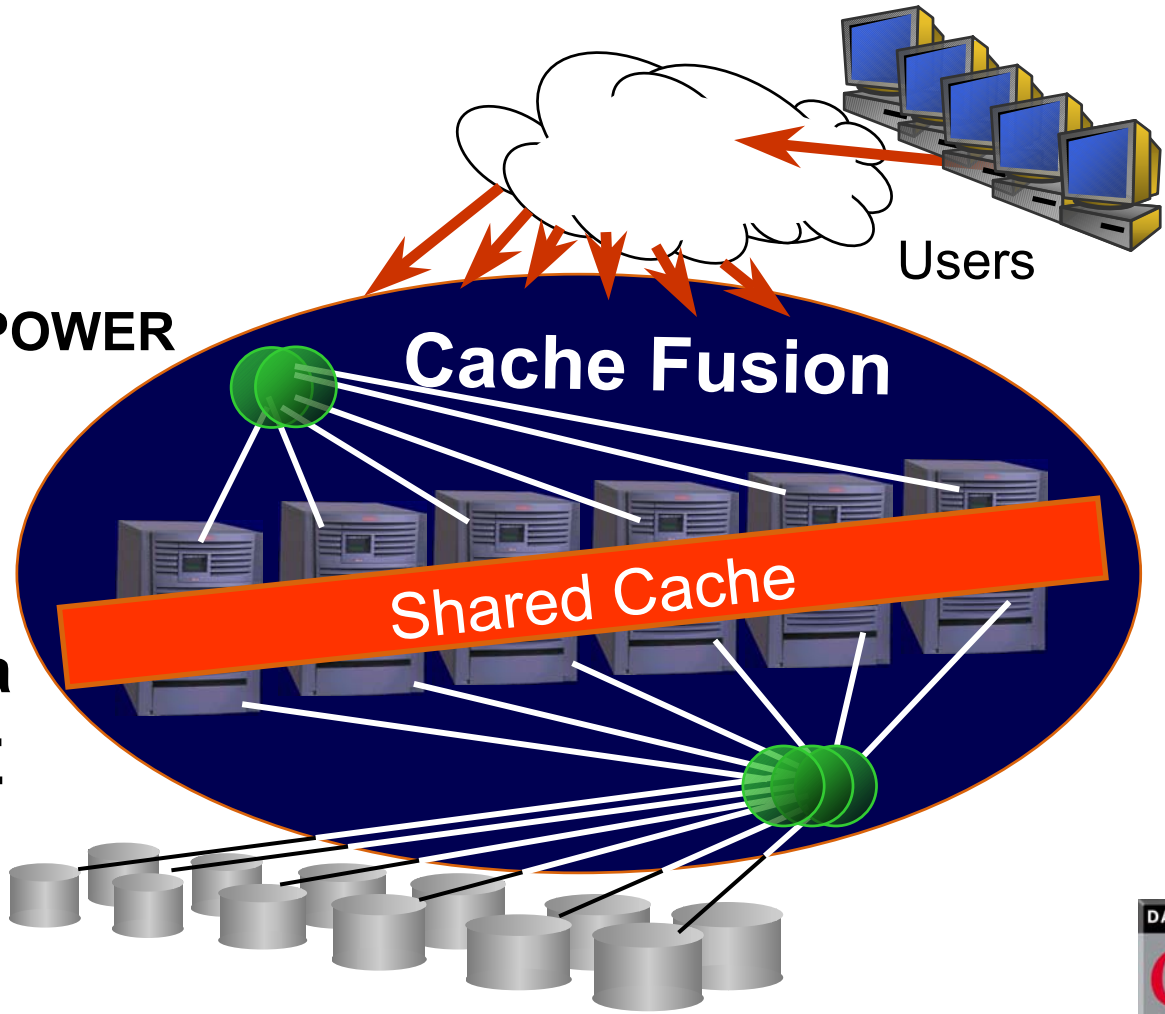

Cache Fusion:

Performance of a SHARED CACHE

ORACLE

DATABASE 9i CLUSTER

Real Application Cluster – what is it?

- Many instances of Oracle running on many nodes
- Multiple instances share a single physical database
- All instances have common data & controls files
- Each instances has individual log files and rollback segments
- All instances can simultaneously execute transactions against the single database
- Caches are synchronized using Oracle's Global Cache Management technology (Cache Fusion)

Full Cache Fusion

Oracle9i® Cache Fusion increases performance and scalability

- Data is shipped directly over high-speed interconnect
- Minimize disk I/O

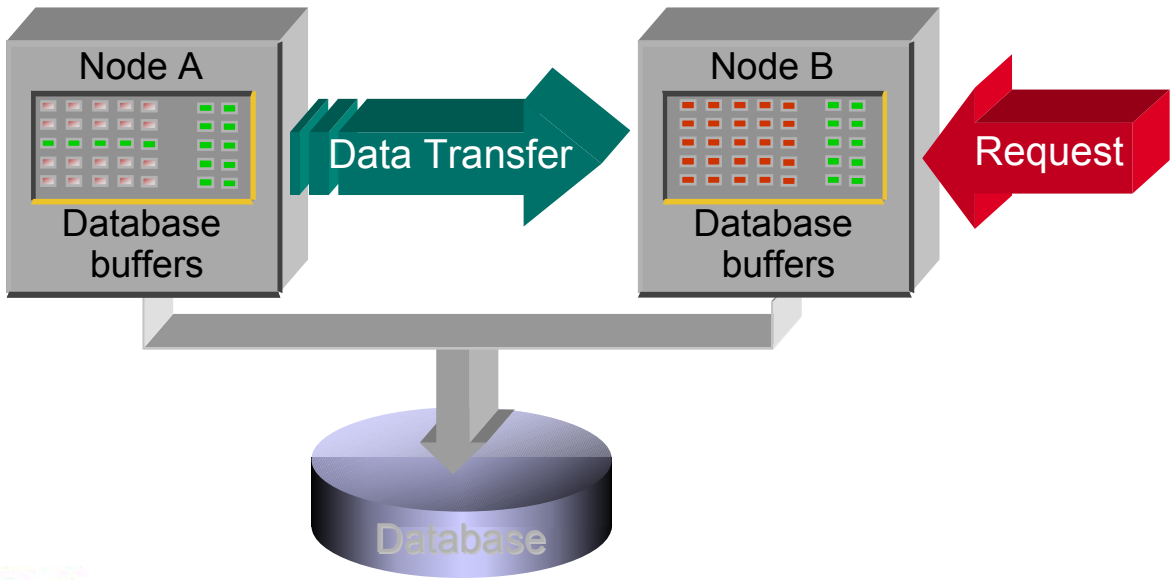

RAC uses a Shared Data Model

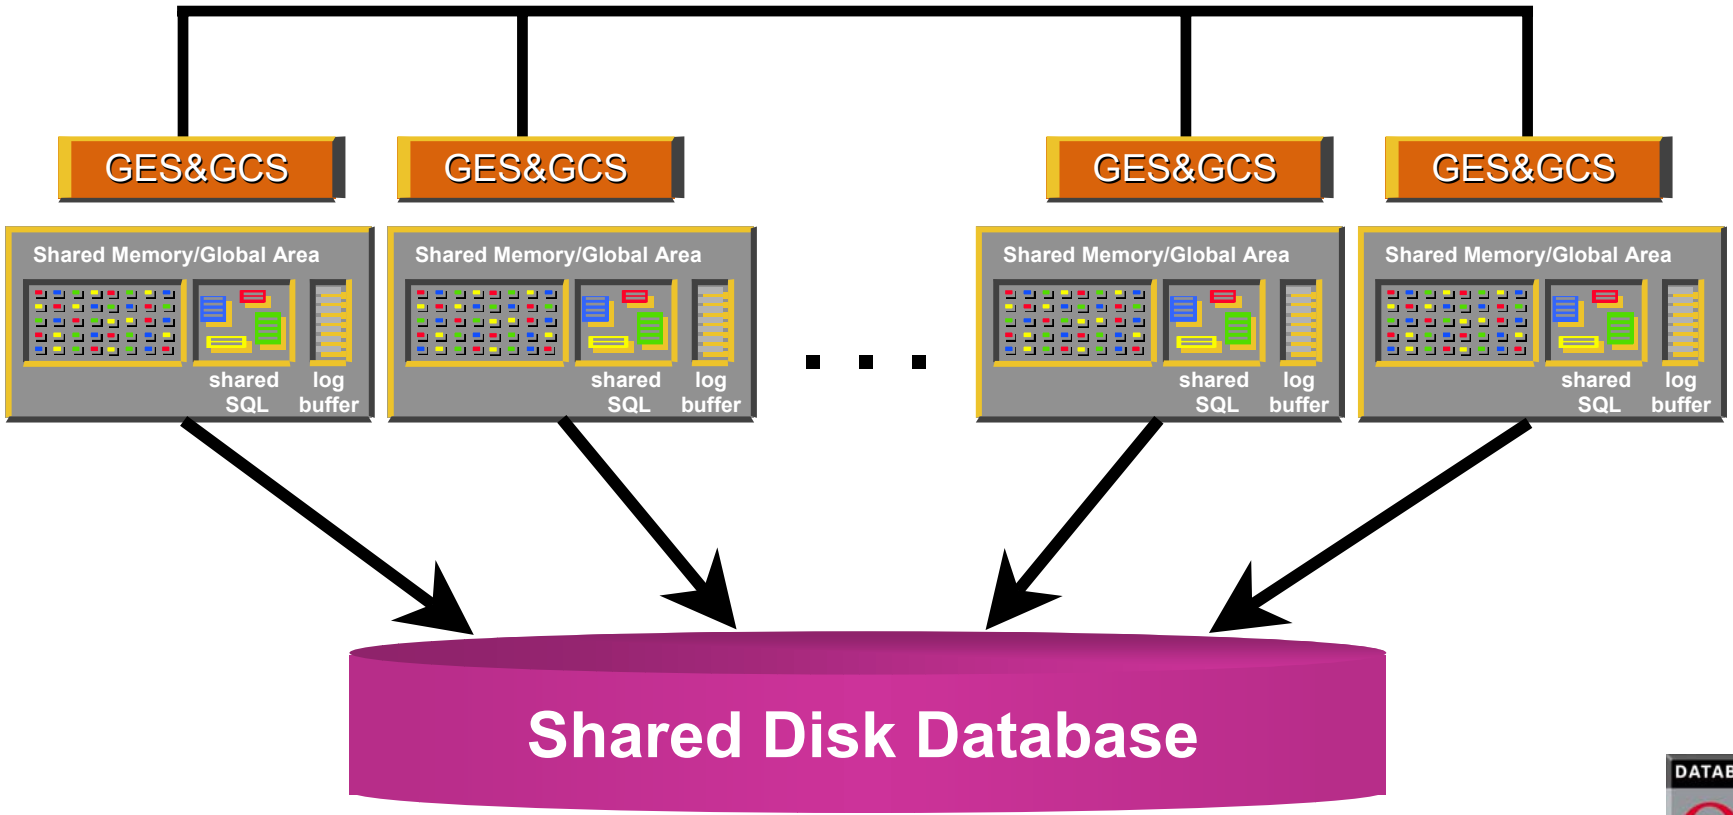

Unlimited Scalability

- Adds incremental I/O bandwidth
- Incremental processing power is immediately available for useful work
- Requires no re-partitioning or re-distribution of data

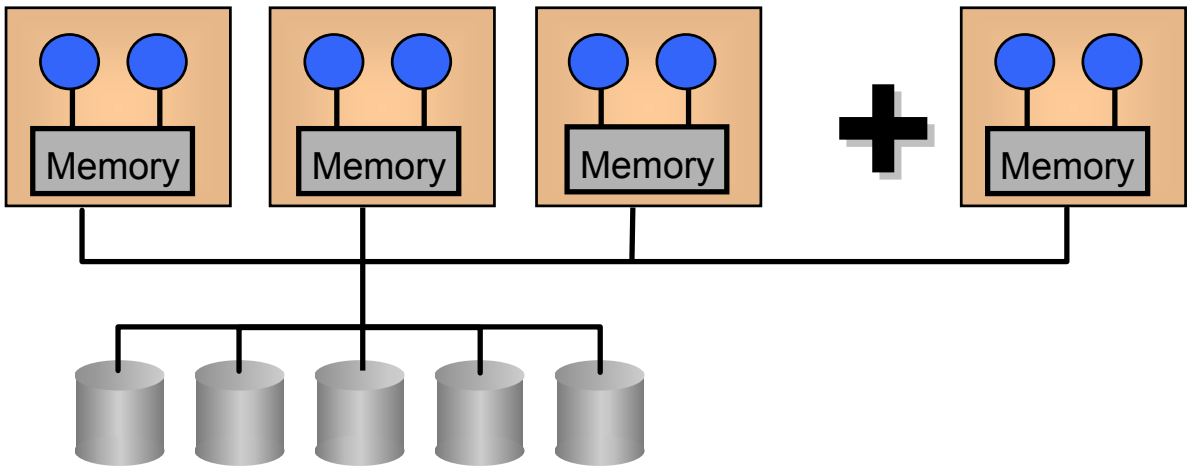

Windows Server 2003 and Oracle9*i*

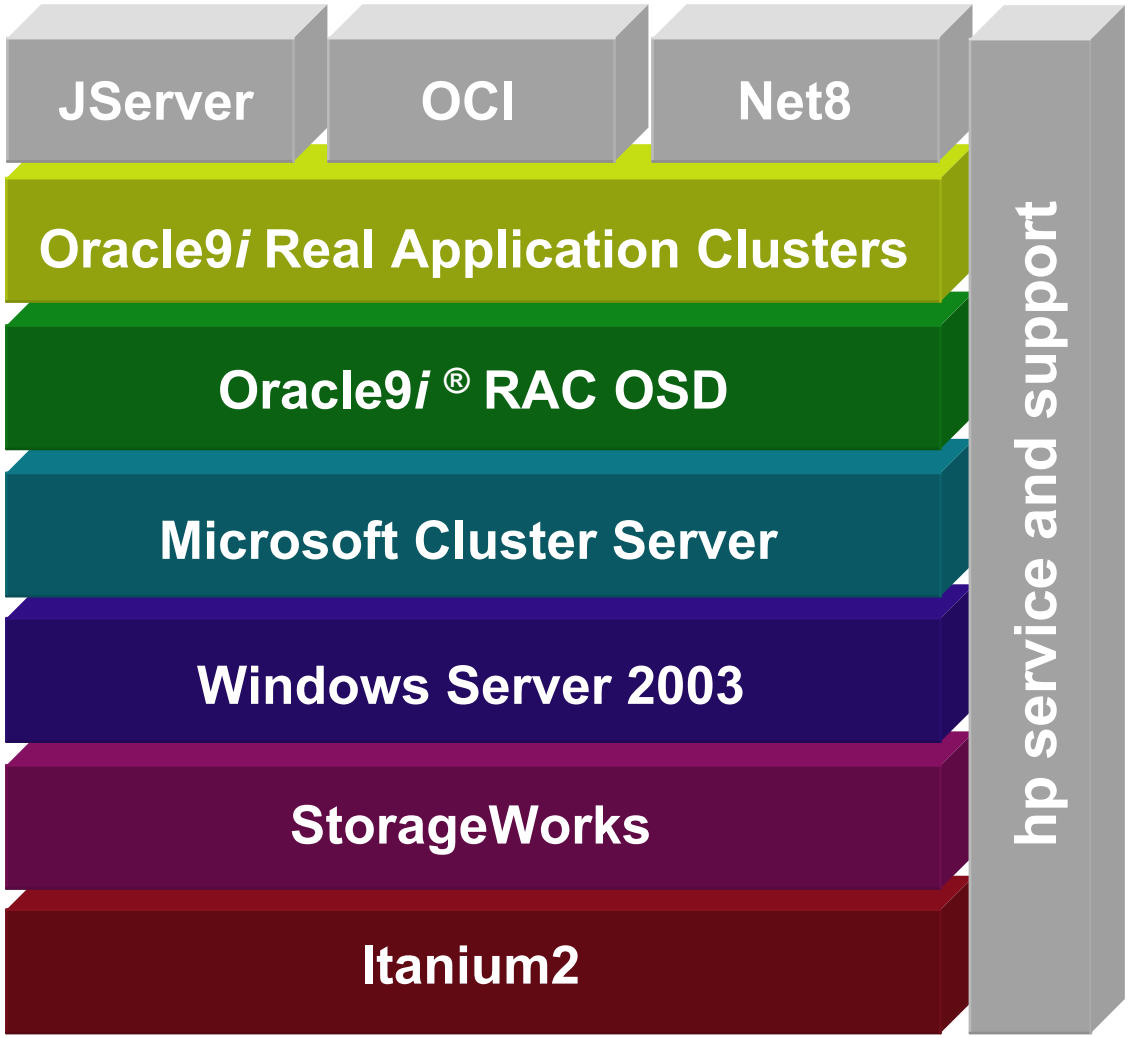

Traditional Clusters

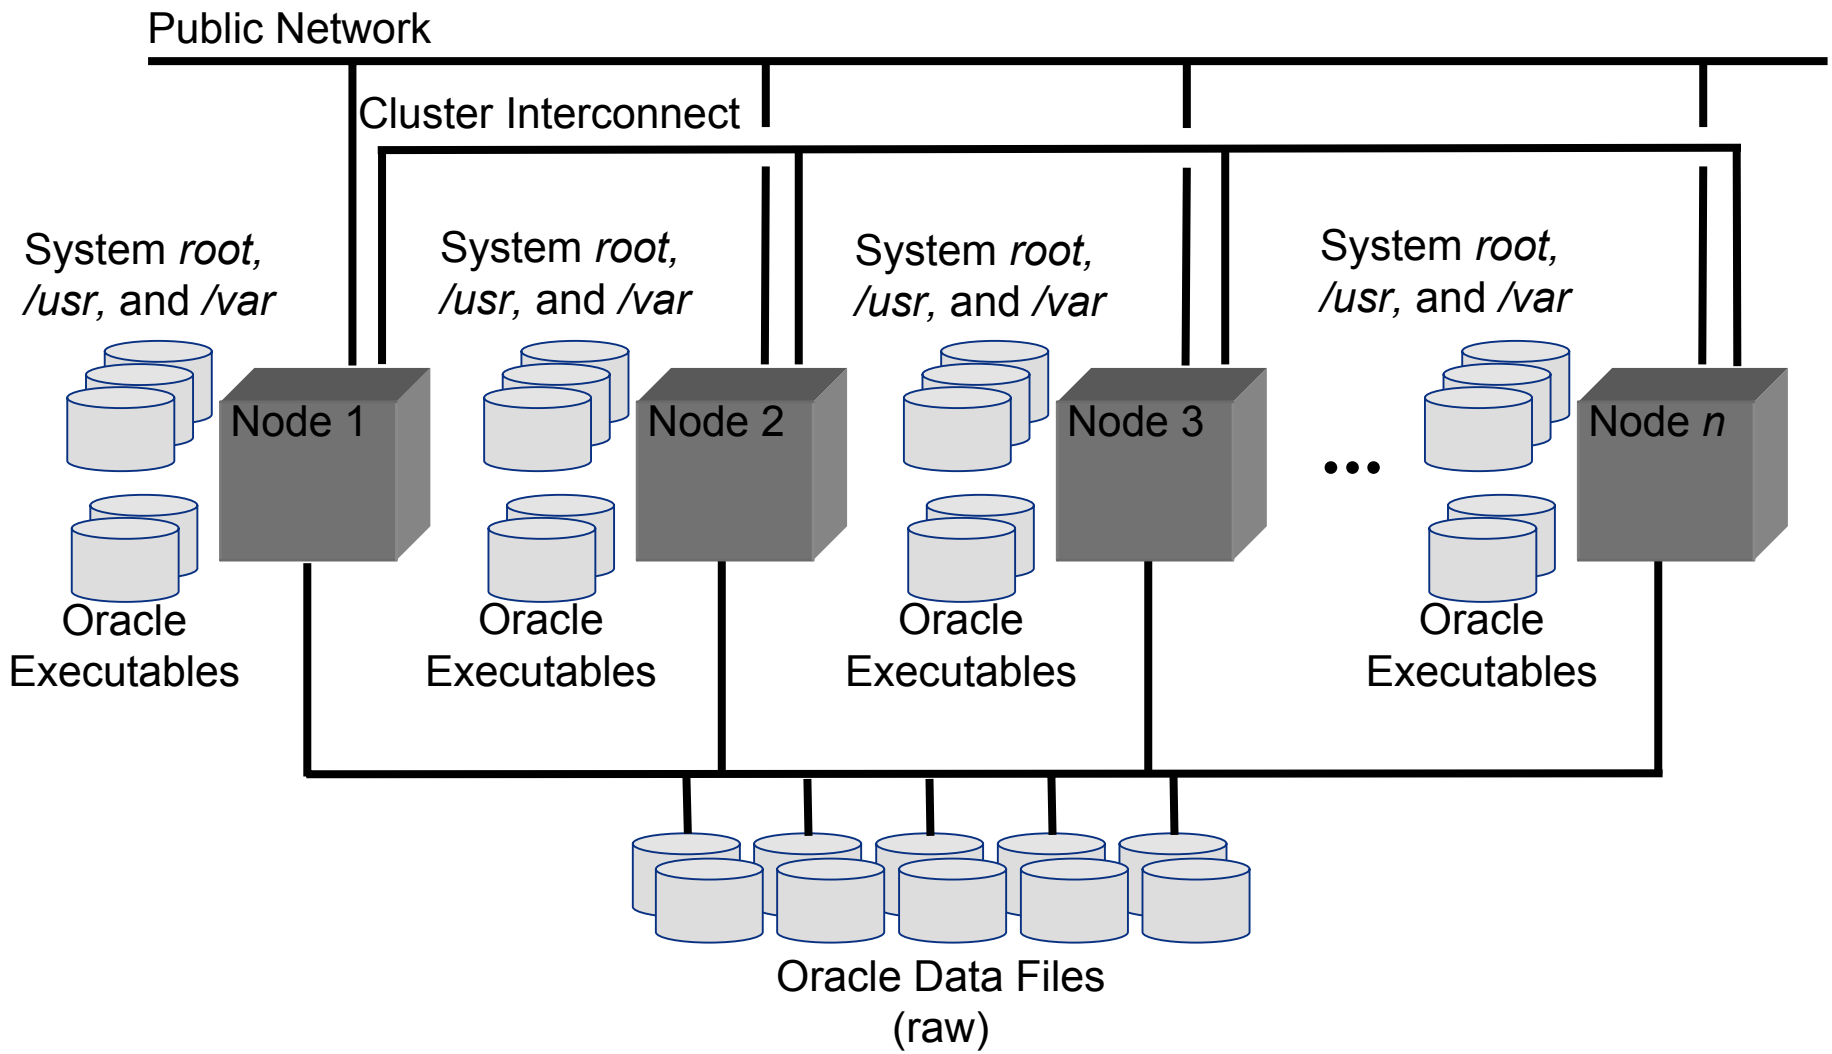

Single System Image (SSI) Clusters

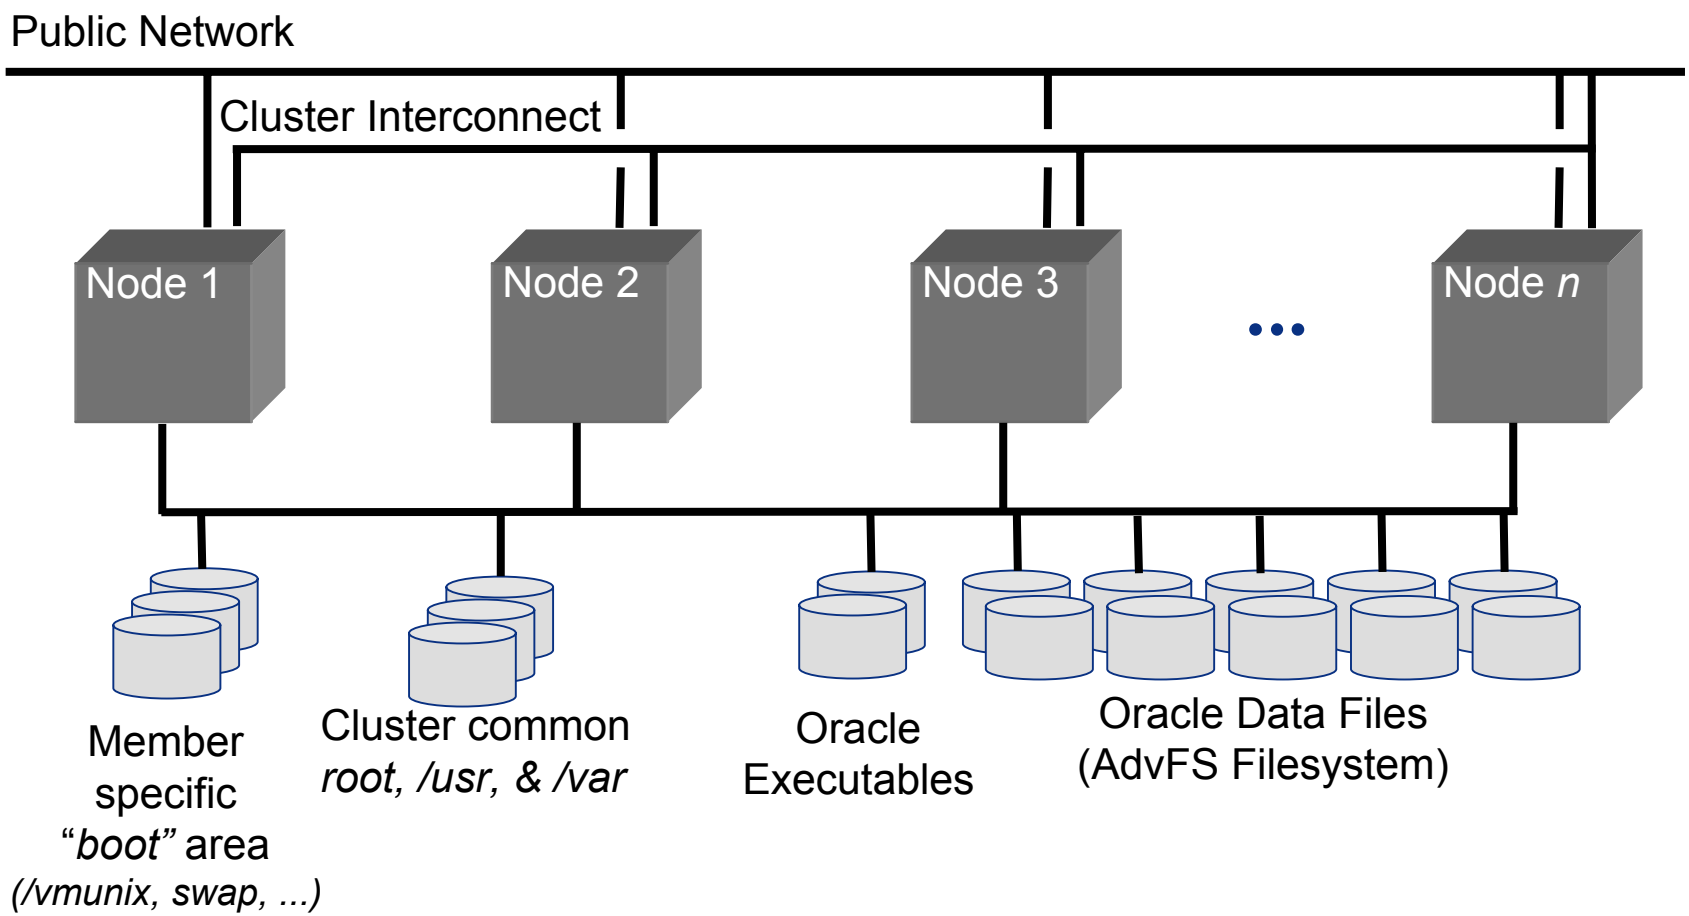

The Demo

- Start the base operating system on both servers
- Launch the cluster manager
- Start Oracle Databases and Instances
- Oracle9*i* Database
  - Parallel Query – single application scalability
  - Calling Circle multi-node, multi-user scalability
  - Transparent Application Facility (TAF) Failover – single instance
  - Calling Circle load-balancing and TAF
  - Oracle Enterprise Manager

Itanium2 Portable 9i RAC Demo Cluster

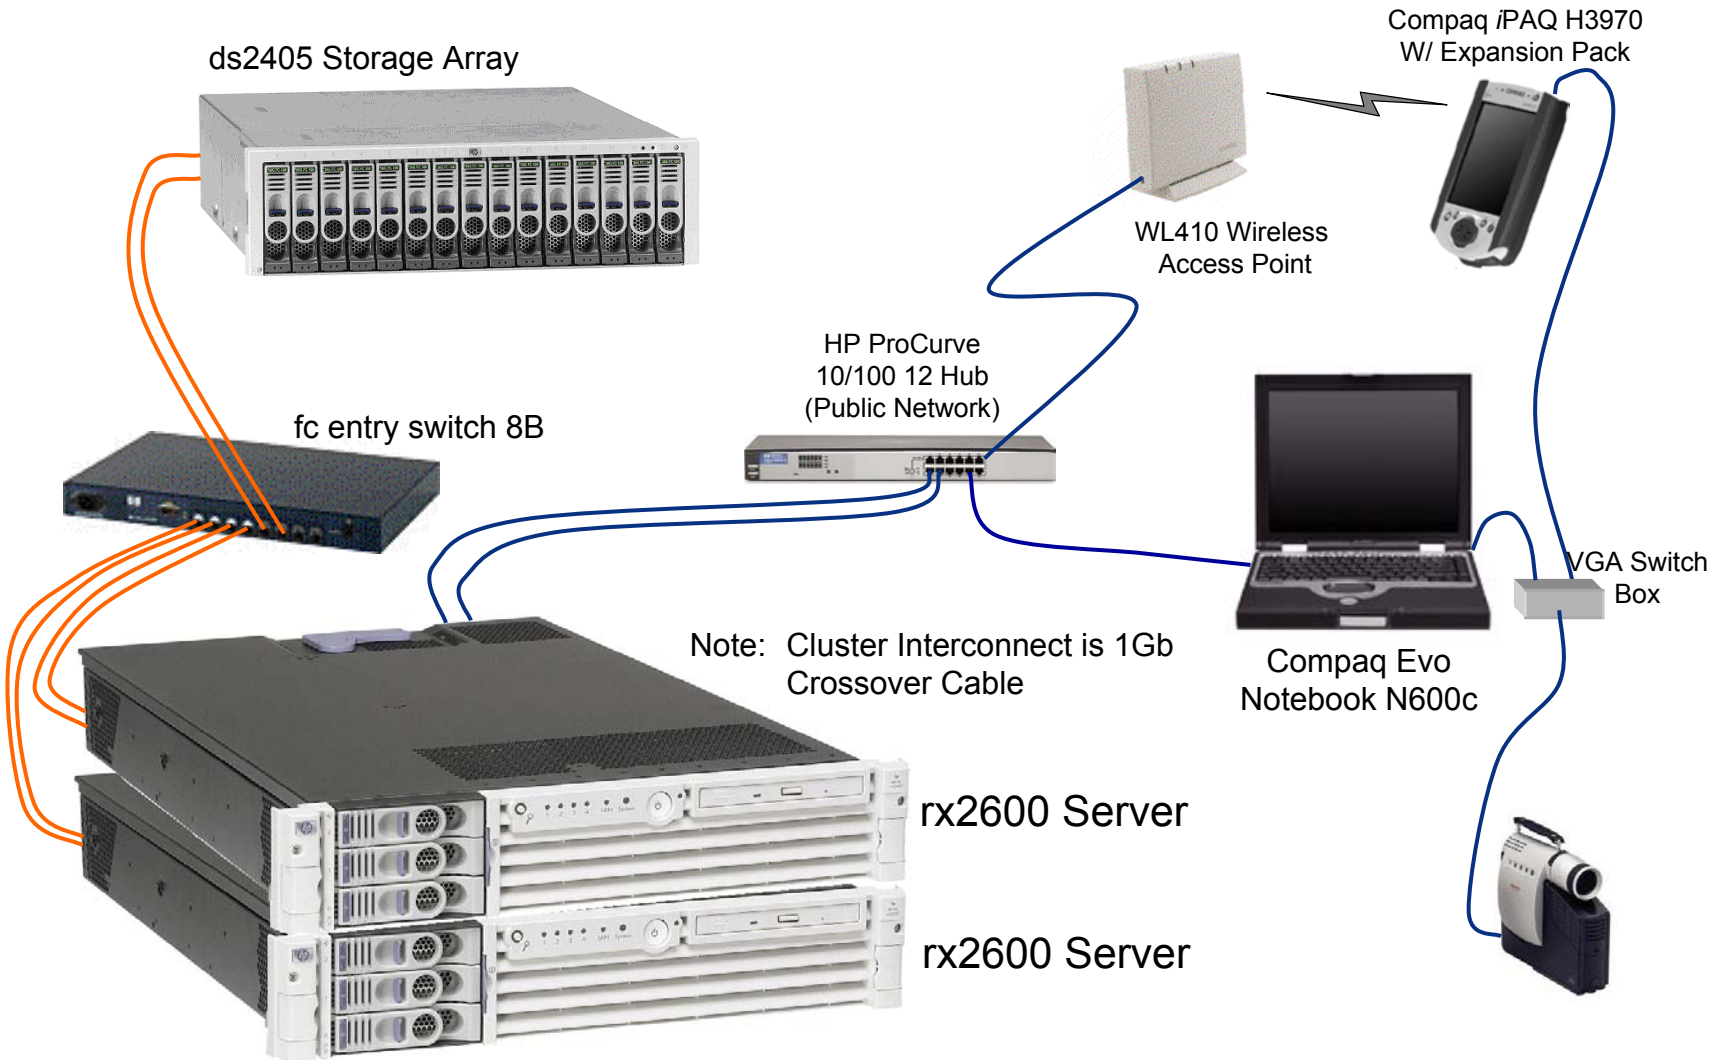

The "Parallel Query" Demo – Single System

```
select
min(total), max(total), avg(total), sum(total), count(*)
from so;
```

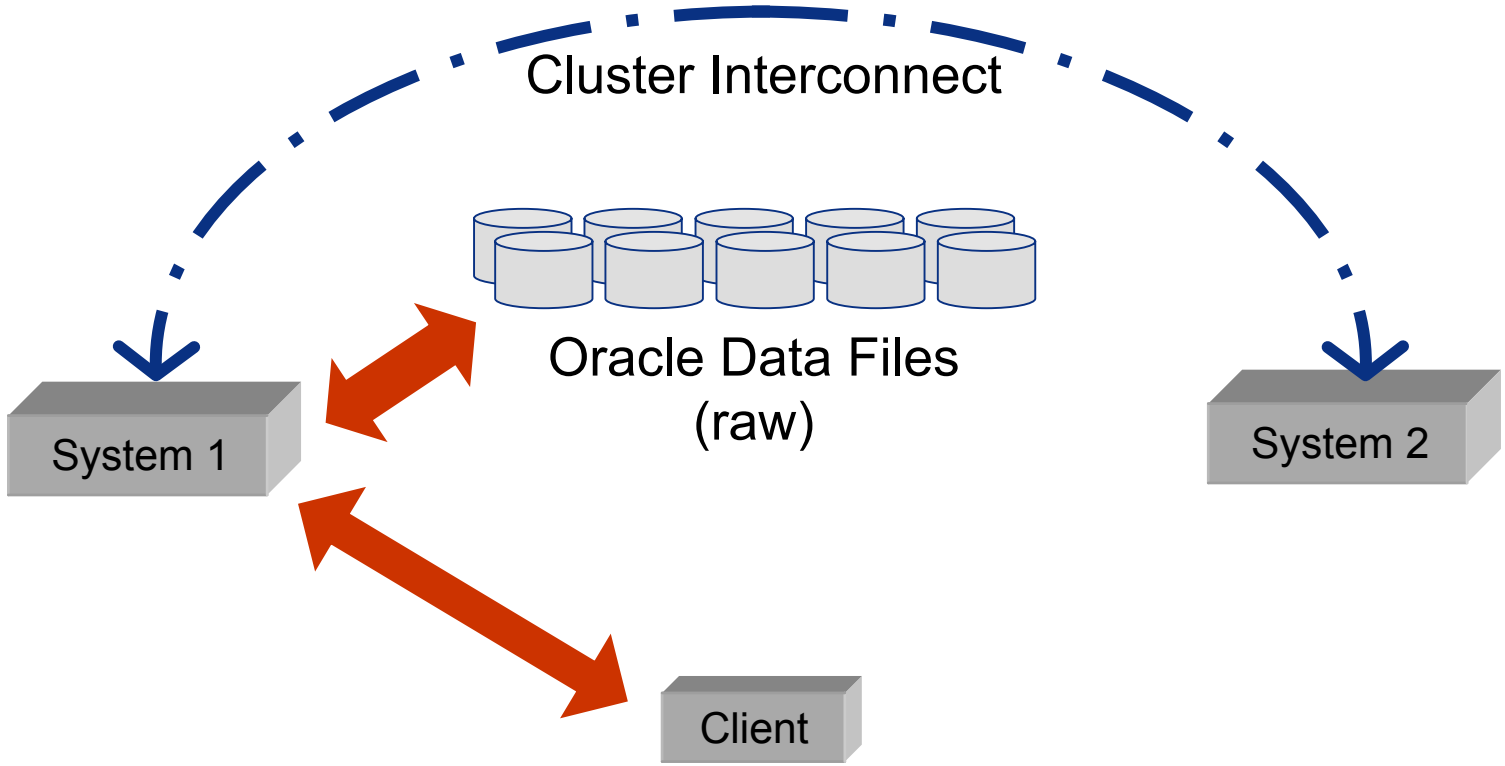

The "Parallel Query" Demo – Multiple Systems


```
select  /*+parallel(so,4,2)*/
min(total), max(total), avg(total), sum(total), count(*)
from so;
```

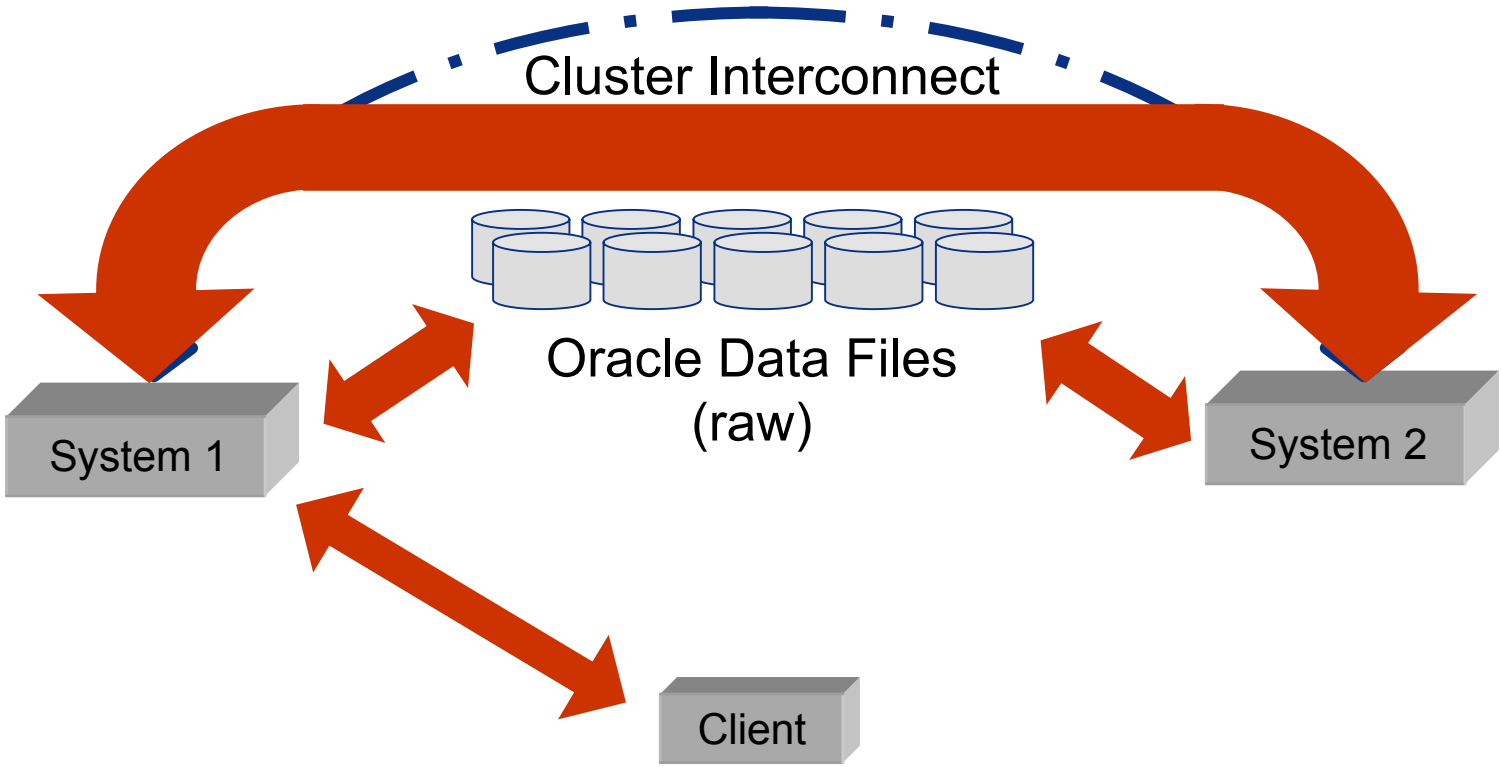

Multi-node Scalability

- Custom Java Based Development
- Simulates a “Calling Circle” application

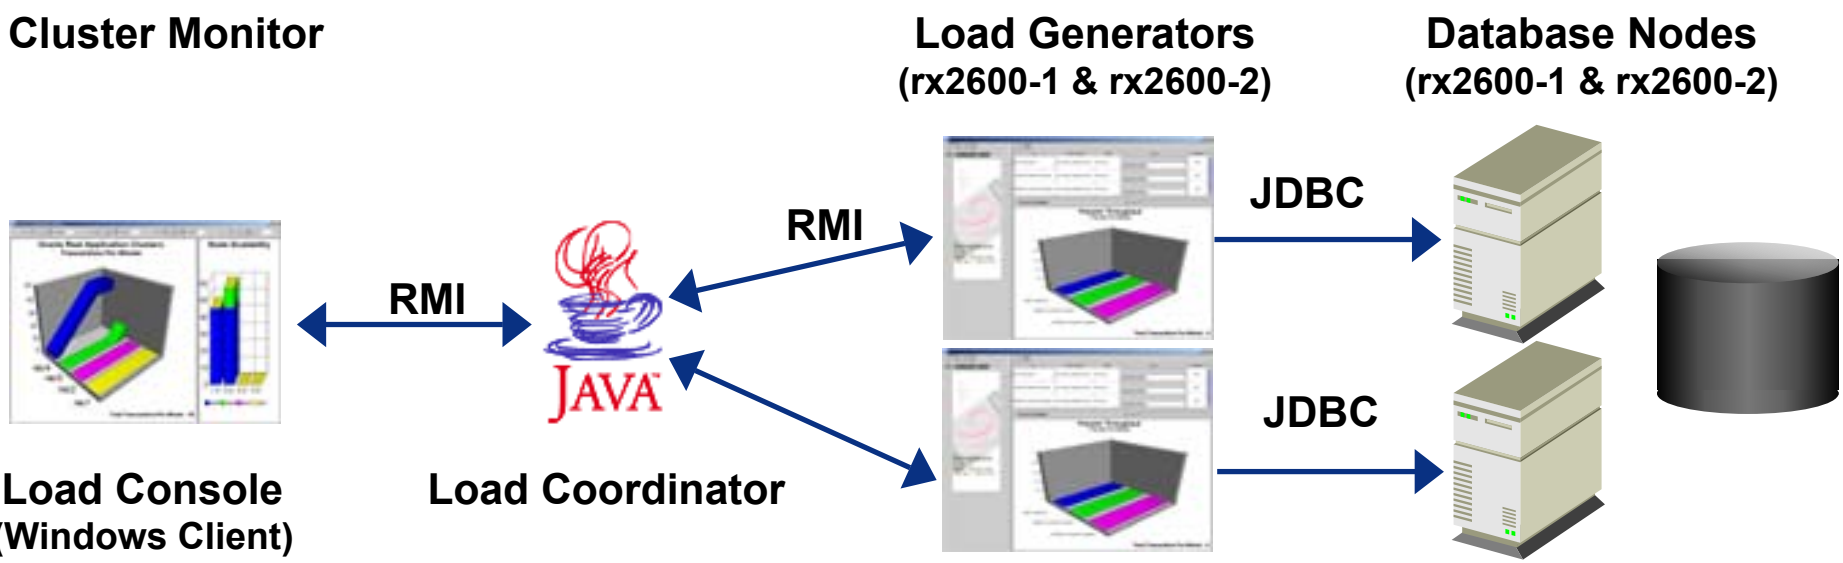

Calling Circle Benchmark

- Modelled from production system
- Simulates SQL generated by Web Application
- Users can create, update or query a calling circle
- Update intensive
- No optimization for RAC in code
- Developed to include aspects known to cause scalability challenges with Oracle Parallel Server
- Features several keys generated from sequence numbers, resulting in contention for right growing indexes

8-node Scalability

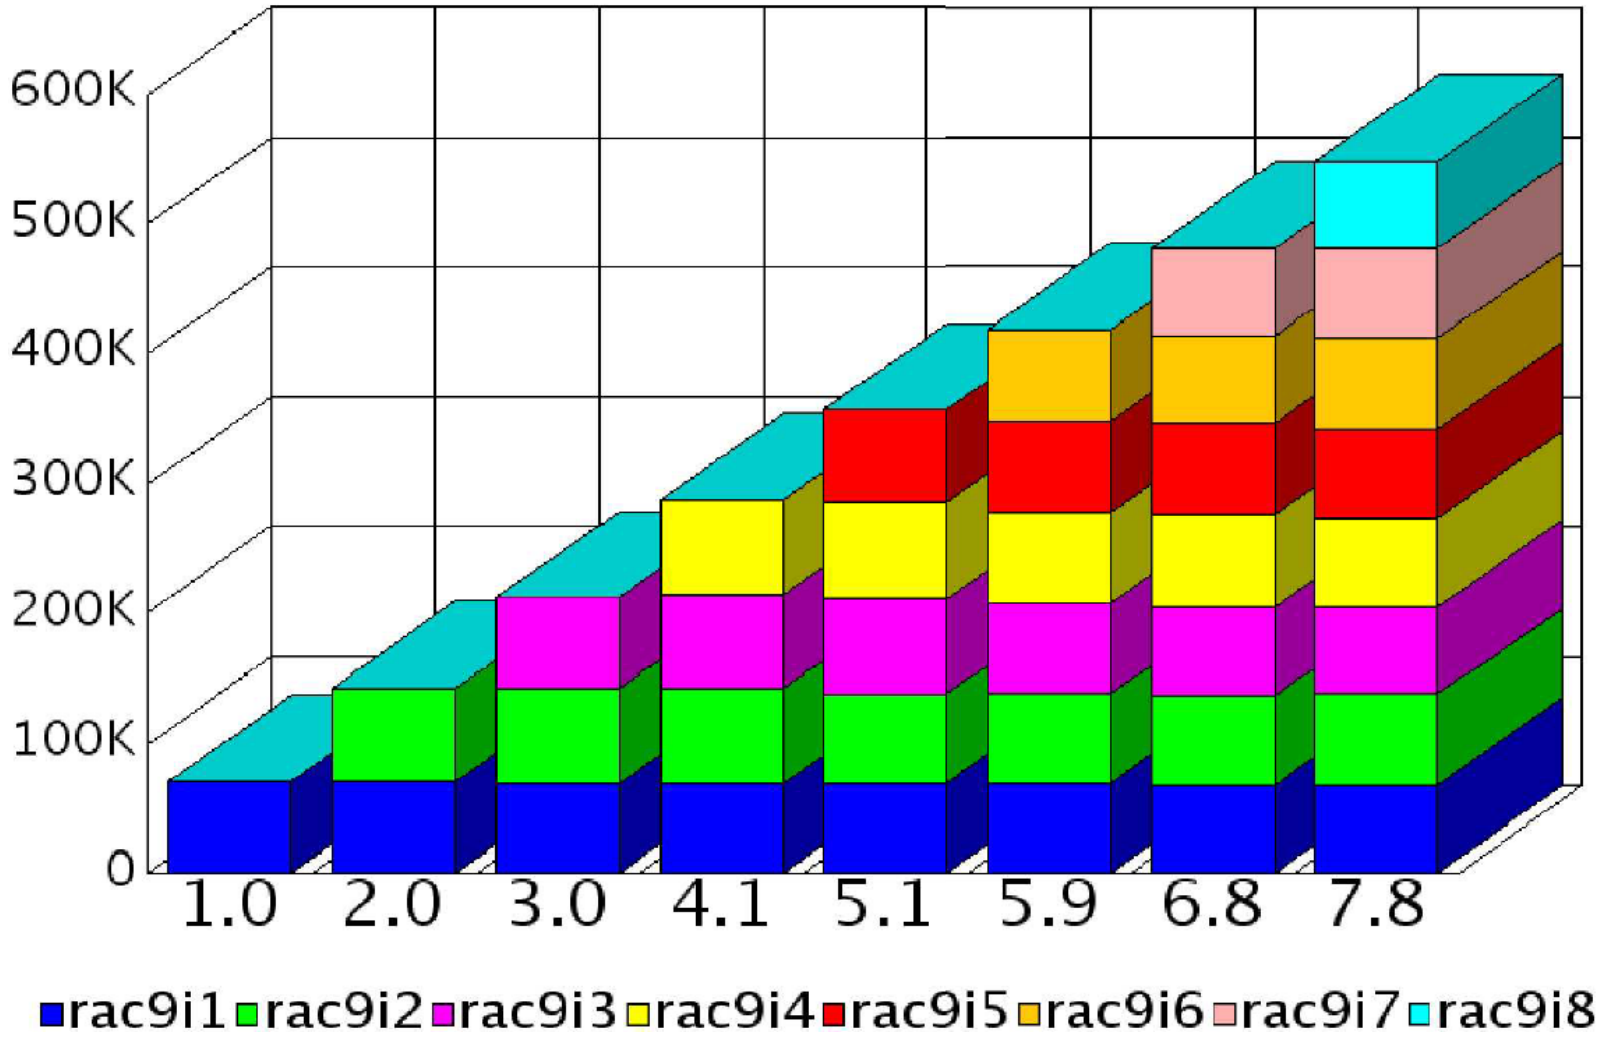

The "fail-over" Demo

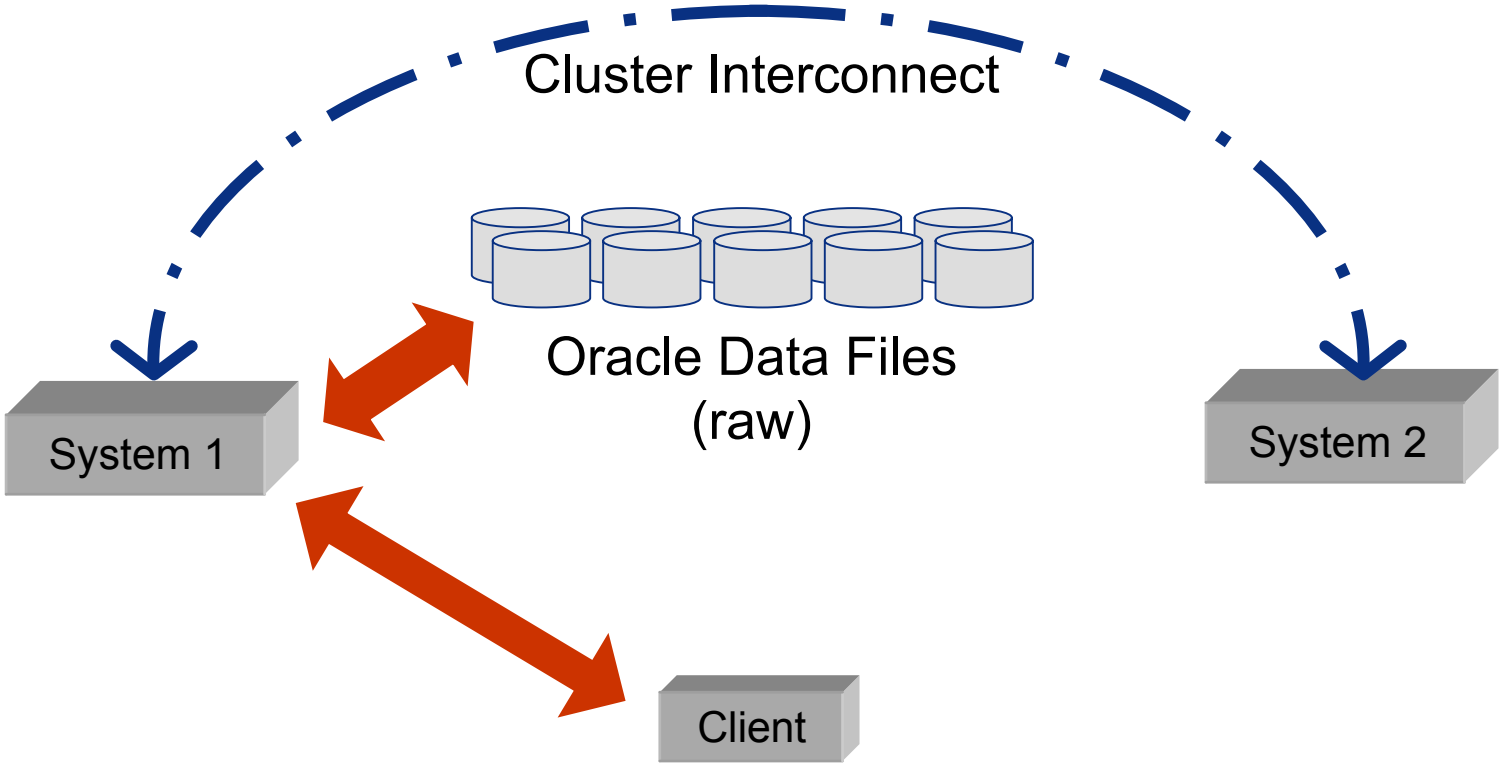

The "fail-over" Demo

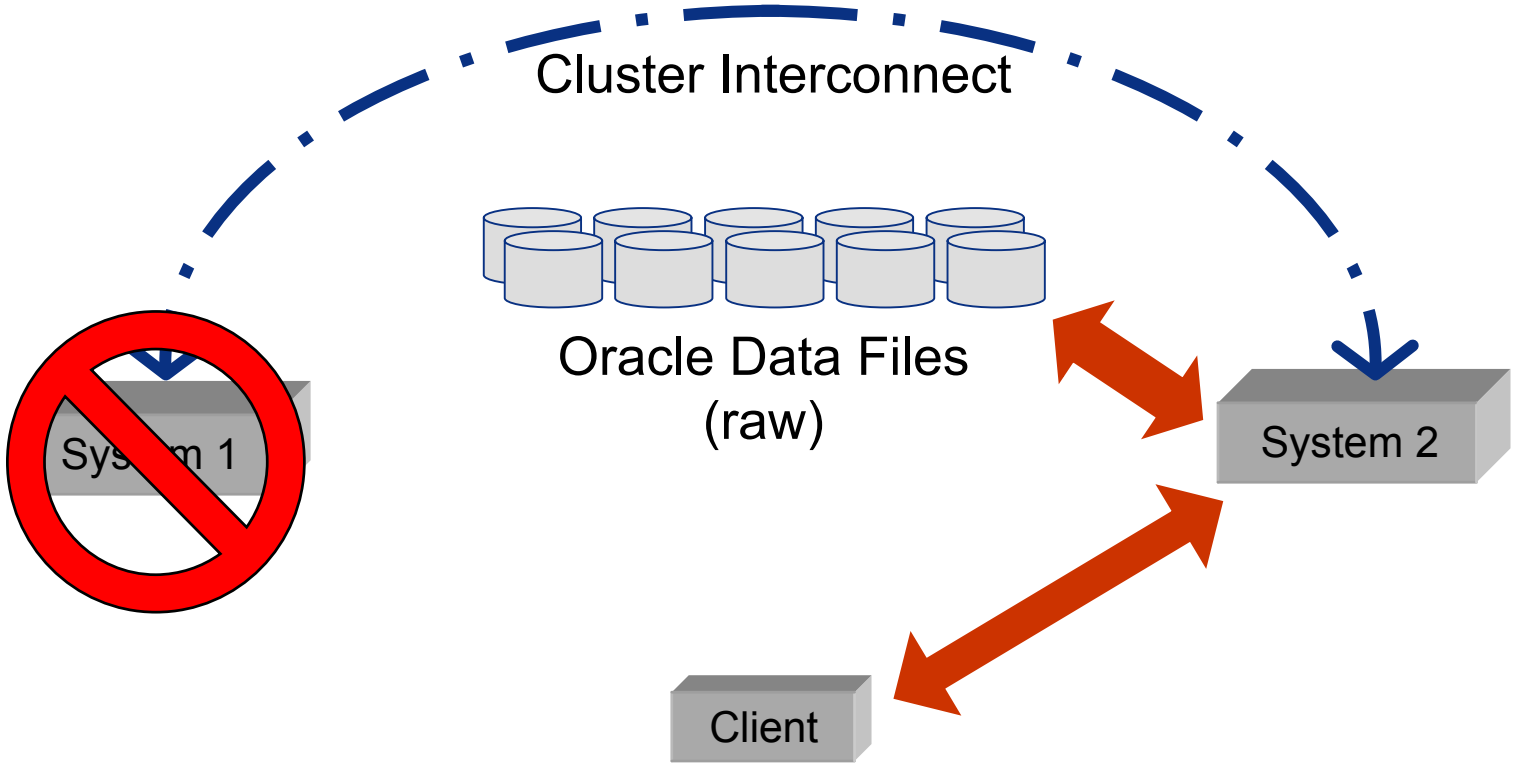

Load-balancing and TAF

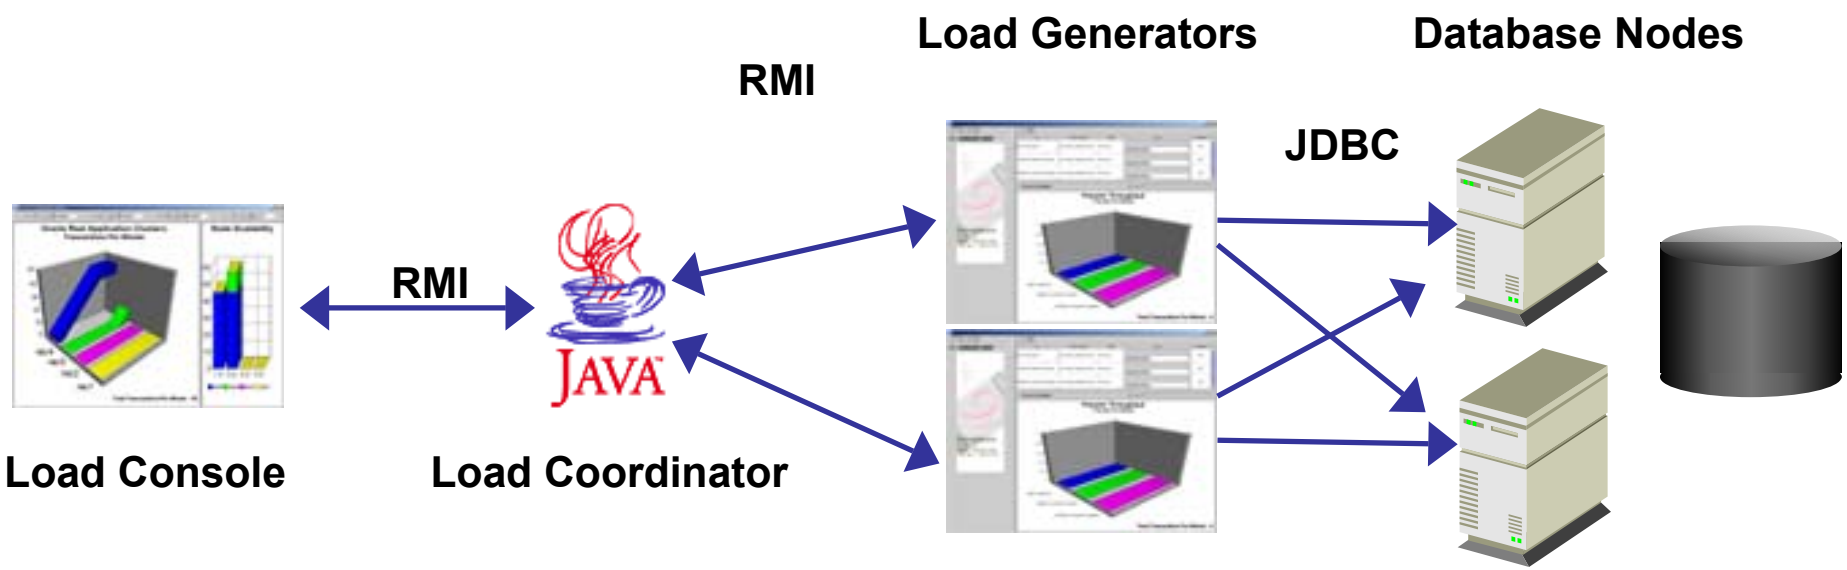

Client-side tnsnames.ora

```
# TNSNAMES.ORA Network Configuration File:
C:\oracle\ora92\network\admin\tnsnames.ora
# Generated by Oracle configuration tools.

MGRTLB =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (LOAD_BALANCE = ON)
      (FAILOVER = ON)
      (ADDRESS = (PROTOCOL = TCP)(HOST = rx2600-1)(PORT = 1561))
      (ADDRESS = (PROTOCOL = TCP)(HOST = rx2600-2)(PORT = 1561))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = MGRT)
      (FAILOVER_MODE =
        (TYPE = SELECT)
        (METHOD = BASIC)
        (RETRIES = 64)
        (DELAY = 4)
      ) ) )
```

Load-balancing and TAF

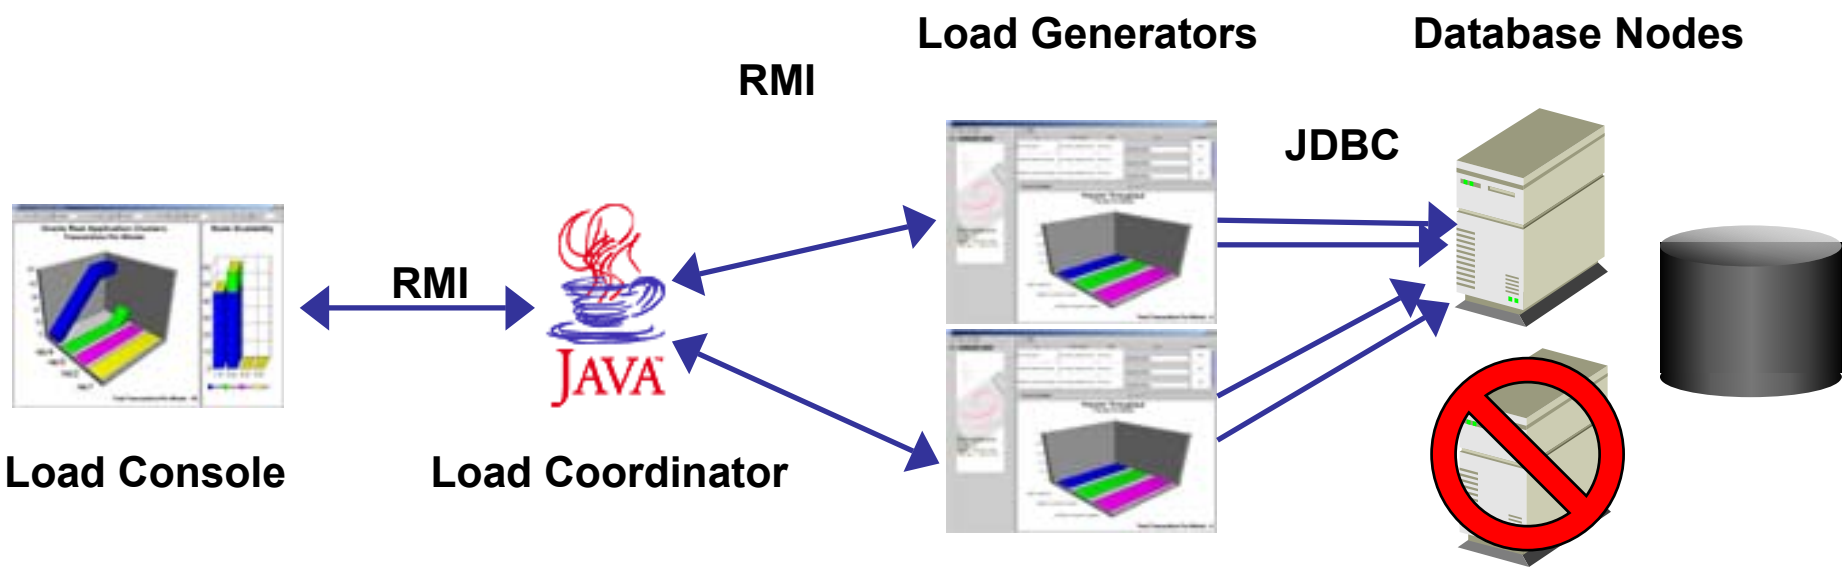

Interex, Encompass and HP bring you a powerful new HP World.