Choosing 32-bit or 64-bit Linux – a decision making process

Sanhita Sarkar
Oracle Corporation
Sanhita.Sarkar@oracle.com

Topics for Discussion

- What is 64-bit computing?
- The dilemma in IT industry to migrate to the realm of 64-bit computing
- The dependency on various support strategies of the enterprise platform
- Processor architectures: 32-bit Intel Xeon and 64-bit Intel Itanium 2
- Common Design Principles of the 32-bit and 64-bit architectures
- The 64-bit extension technology and advantages
- The AMD64 architecture : Opteron processor and advantages
- The effect of processor architecture on applications
- Making a sound choice: 32-bit or 64-bit
  - By database application type, memory requirements, price/performance goals, user density, flexibility, etc
  - Checklist towards decision making
- Split configuration for multi-tier applications – a solution

What is 64-Bit Computing?

- Number denotes data stream width
  - Registers
  - Internal data paths
- Increased range:
  - Integers and Addresses

    Range goes:

    from 2^32 = 4.3e9

    to 2^64 = 1.8e19

  Factor of 4.3 billion increase
- Addresses
  - 32-bit: 4.3 billion bytes = 4 GB
  - 64-bit: 18 million terabytes
  - Theoretical vs. Practical

Migrating to the realm of 64-bit computing: A dilemma

- Dramatic increase of potential penetration of high-performance 64-bit computing
- Dilemma for IT organizations who have not considered 64-bit computing critical
- Countless infrastructure to migrate
- Unclear performance gains
- Volume of work involved in a comprehensive migration
- Cost and complexity – a cost/benefit analysis required
- So far, infrastructure applications remained on 32-bit and business logic moved to 64-bit platforms
- Complex questions:
  - **Is this enterprise a good candidate to move to 64-bit and is this an appropriate time to migrate?**
  - **Which version of operating system, drivers, database, etc are appropriate for migration?**
  - **Are the 64-bit versions supported/available to meet the migration needs of the enterprise?**
  - **What services are required to assist in a successful transition and how much will it cost?**

The Stack

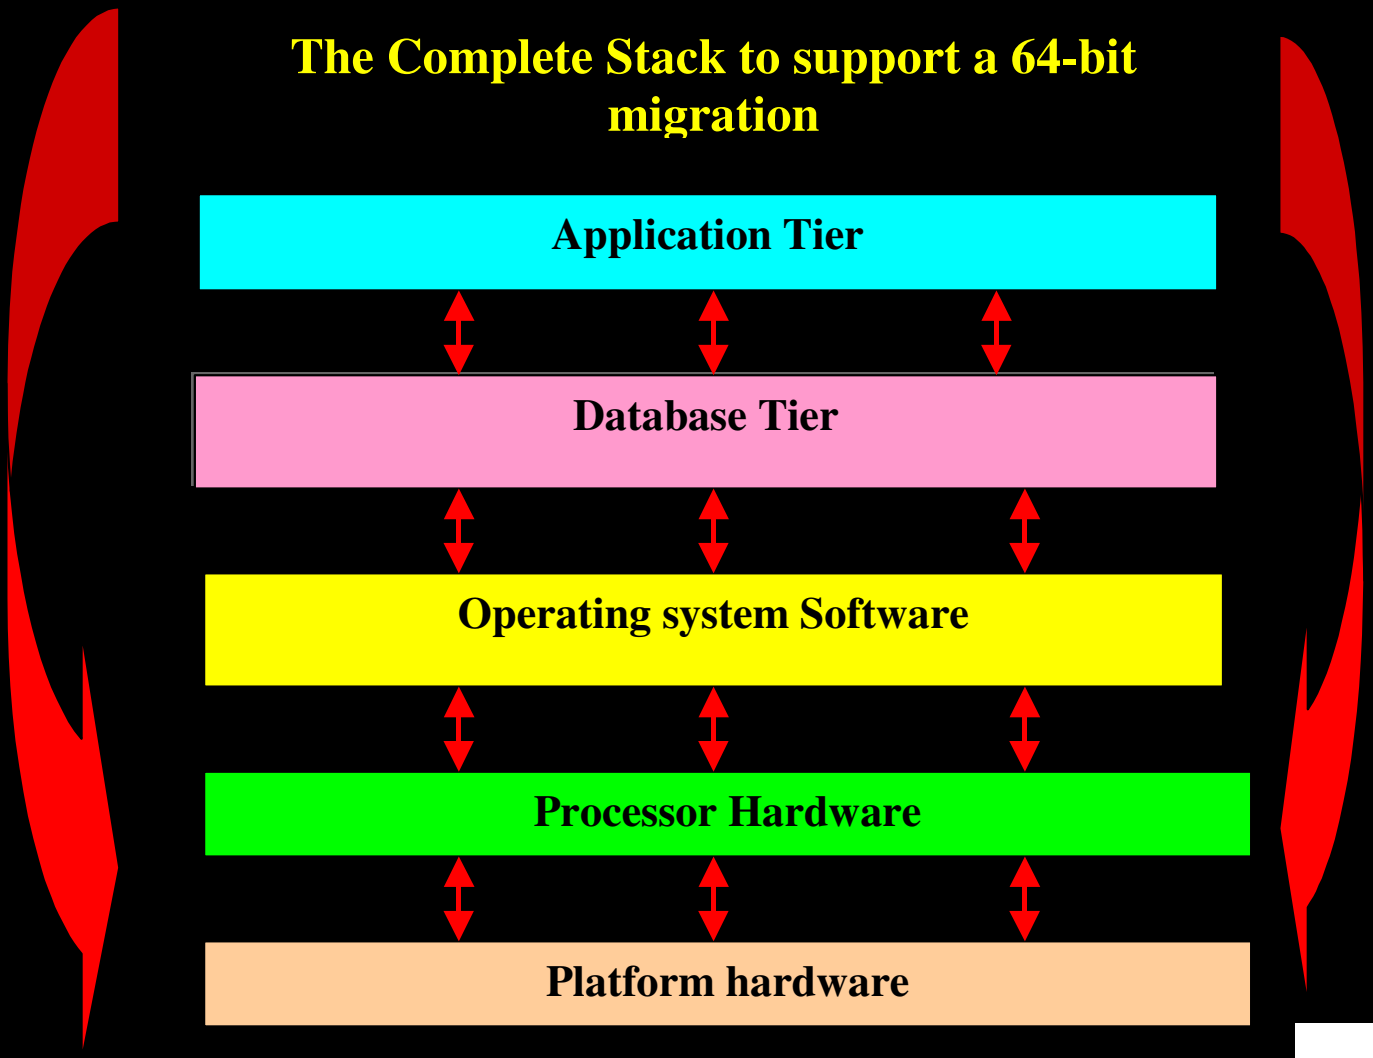

The IA-32 architecture

Features	Benefits
Hyper-Threading technology	Increases the number of processed transactions, user density, faster response times for enterprise applications
Integrated Three-Level Cache Architecture 4MB, 2MB or 1MB iL3 Cache Level 2 Advanced Transfer Cache (512 KB) Level 1 Execution Trace Cache	Provides higher throughput than previous-generation cache for large server workloads. Higher bandwidth path to memory increases throughput for large server workloads
Intel NetBurst Microarchitecture Rapid Execution Engine SIMD Extensions 2 (SSE2)	2x clock speed for integer computations; 144 new instructions improve response times for media servers, secure transactions and next-generation Web services
400 MHz System Bus	Higher throughput when accessing memory and I/O devices for improved server headroom and scalability. Allows for faster data transfer and improved performance
Available at 3 GHz w/4MB iL3 Cache, 2.80 GHz w/2MB cache, 2.20 GHz w/2MB iL3 Cache, 1.90 GHz w/1MB iL3 Cache	One of the fastest MP-capable processor from Intel
PCI-X Available	A faster and higher throughput I/O subsystem maximizes server and I/O performance

The IA-32 architecture: Benefits

- The Intel Xeon processor MP is designed specifically for 4-way, (and above), multi-processor (MP)-based servers
- The Intel Xeon processor MP is now available up to 3 GHz and 4 MB iL3 cache
- Provides a seamless upgrade path for higher performance
- Extensive platform validation with leading OEMs ensure reliability and provide a wide range of scalable solutions that are proven to support business-critical needs of the enterprise
- The Intel Xeon processor MP is ideal for the demanding application tier, plus offers the performance and headroom required for the medium-size database tier

IA-64 : A High Performance Architecture for Enterprise-Class Computing

Evolution of IA-64 architecture

- CISC and RISC architectures – Limitations
- The EPIC concept
- Evolution of IA-64 architecture based on Itanium
- Requirement of efficient compilers running on Itanium
- Predication, Control and Data Speculation, Software Pipelining, Data Prefetching
- A fully predicated architecture providing Fault Tolerance and Reliability

Itanium2 Chip Design

- Three-level Cache Organization
- Memory hierarchy
  - Efficient use of the three cache levels
  - Latencies of L1, L2 and L3 caches
- Large number of registers and larger register width
- Large number of floating point units
- Large Cache, prefetching, elimination of branch predictions

Itanium vs Itanium 2 Chip Architecture

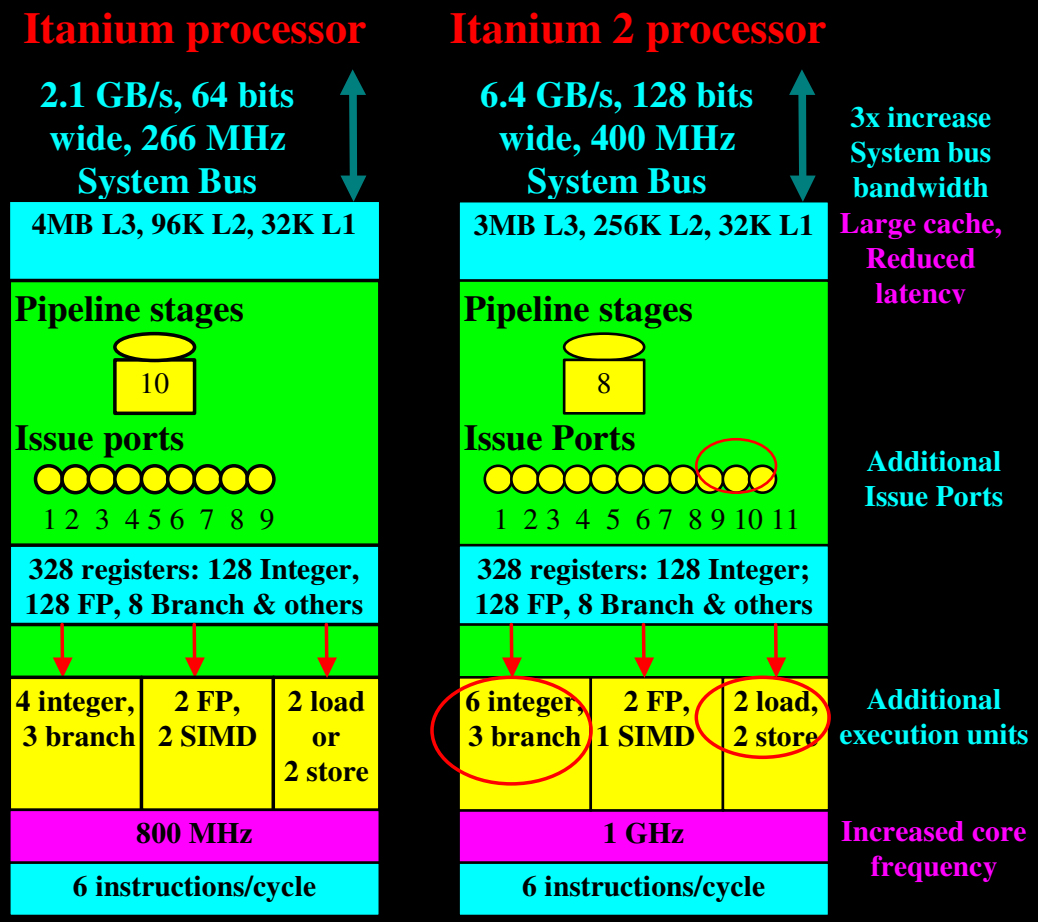

Memory Hierarchy : 3-level Cache Organization

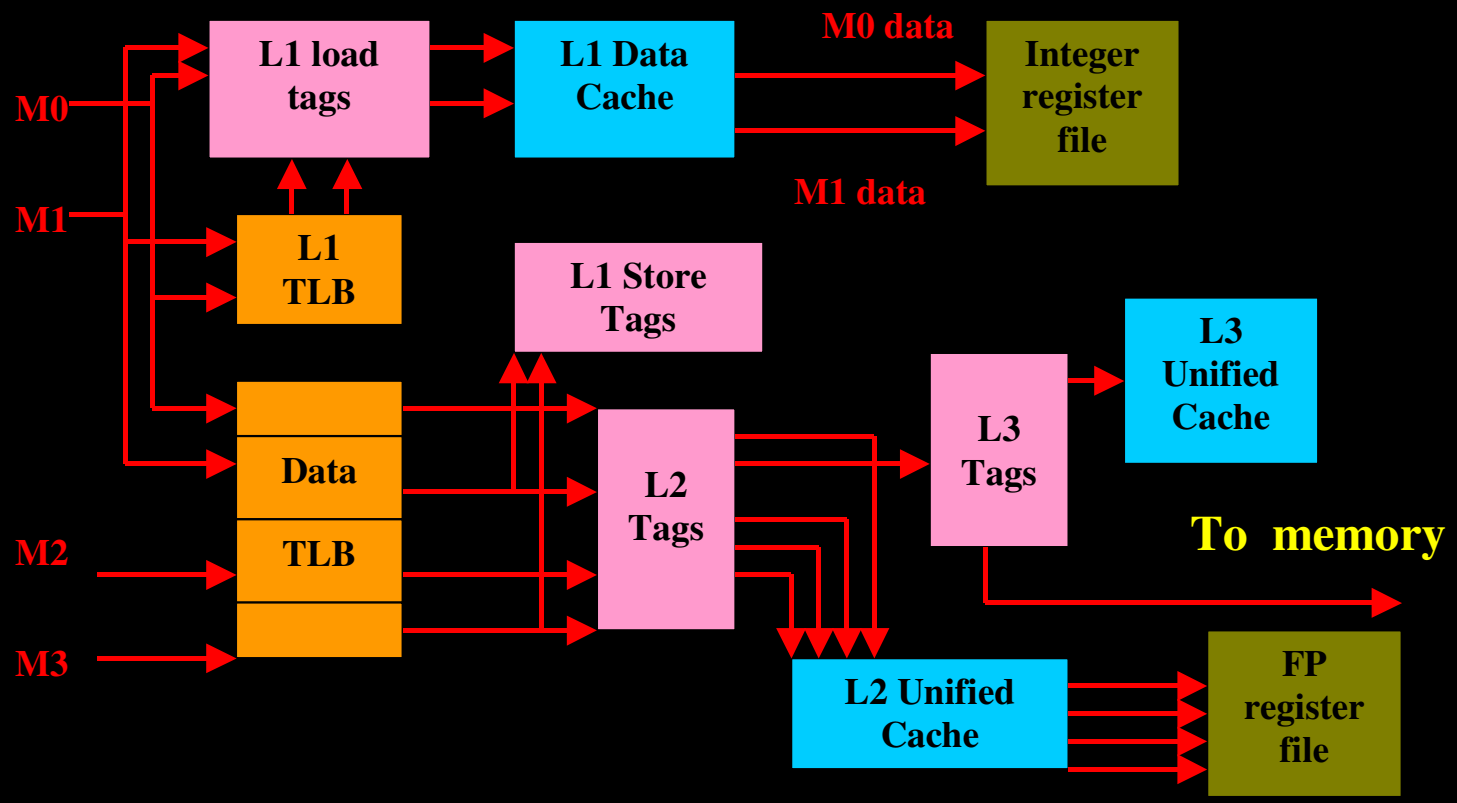

Error Handling by PAL and SAL on Itanium 2

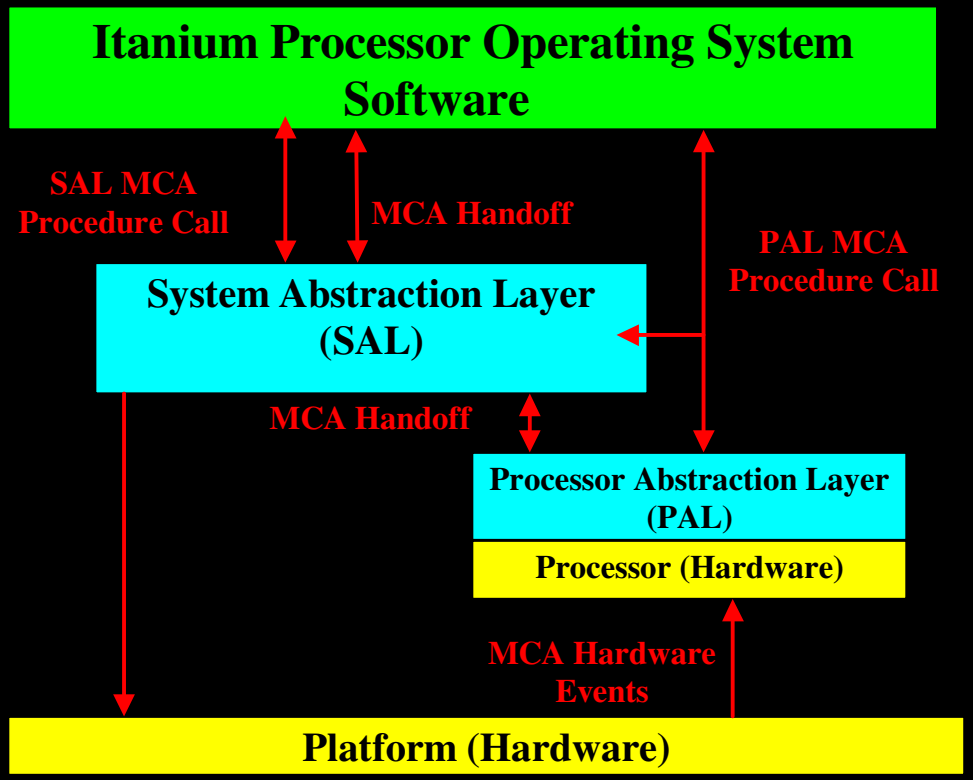

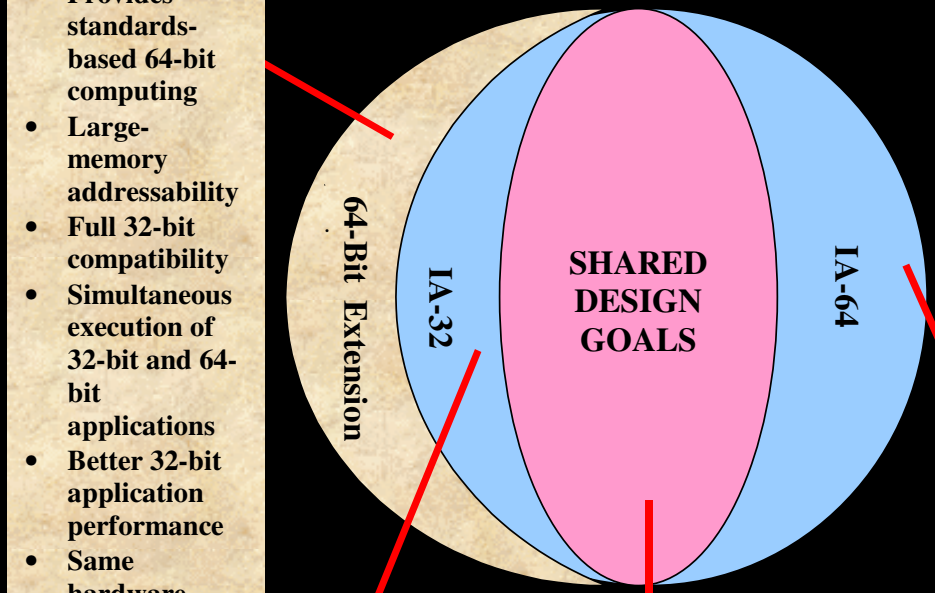

- Provides standards-based 64-bit computing
- Large-memory addressability
- Full 32-bit compatibility
- Simultaneous execution of 32-bit and 64-bit applications
- Better 32-bit application performance
- Same hardware, easy migration

- Balance between cost and performance
- Broadest range of uses (desktop to SMP)
- Environmentally sensitive (heat, form factor, etc)
- Optimized for 8 processors and below

- Improving price/performance
- Broad application support
- Increased memory, bus and I/O bandwidth
- Balanced floating-point and integer performance
- Optimized for SMP

- Highest performance and throughput
- Requires 64-bit application
- Optimized above 4-processor SMP
- Memory resident data sets
- Parallelism in integer and floating-point performance

Design Goals for IA-32, 64-Bit Extensions and IA-64

The 64-bit extension technology: Advantages

- Provides standards-based 64-bit computing and large-memory addressability
- Complete 32-bit compatibility
- Enables to run 64-bit and 32-bit applications simultaneously while improving 32-bit application performance
- Ease of migration to 64-bit applications on the same hardware platform when software support becomes available
- Builds on the existing 32-bit architecture and provides a flexible lower cost platform than Itanium
- Protection of capital investments with a new level of flexibility and price/performance
- Advantageous for mixed-purpose or infrastructure servers running a variety of applications on a single platform
- Applications like directory services, DNS, database front ends, messaging and groupware can benefit from 64-bit extension technology due to better performance and price/performance than Itanium for large number of servers required to deliver such services

Intel processors : 32-bit, w/64-bit extension and 64-bit

	Intel Xeon (32-bit)	Intel Xeon with 64-bit extension	Intel Itanium 2 (64-bit)
32-bit mode	Native	Native	Through emulation layer
64-bit mode	No	Extension technology	Native
System bus	533 MHz, 64 bits wide:up to 4.3 GB/s bandwidth	800 MHz, 64 bits wide: up to 4.3 GB/s bandwidth	400 MHz, 128 bits wide: up to 6.4 GB/s bandwidth
Cache (levels 2/3)	512 KB/up to 2MB	512 KB/up to 2M	512 KB/up to 6MB
Memory addressing	32 bits (4GB)	36-40 bits (1TB)	60 bits (1024 TB or 1 PB)
Error correcting code	No	No	Yes
Lockstep support	No	No	Yes
Corrupted data containment	No	No	Yes
IBM Chipkill memory feature	Yes	Yes	Yes
Memory spares	Yes	Yes	Yes

AMD64 architecture: The *Opteron* processor

- provides compatibility, performance, investment protection and reduced total cost of ownership

The AMD Opteron

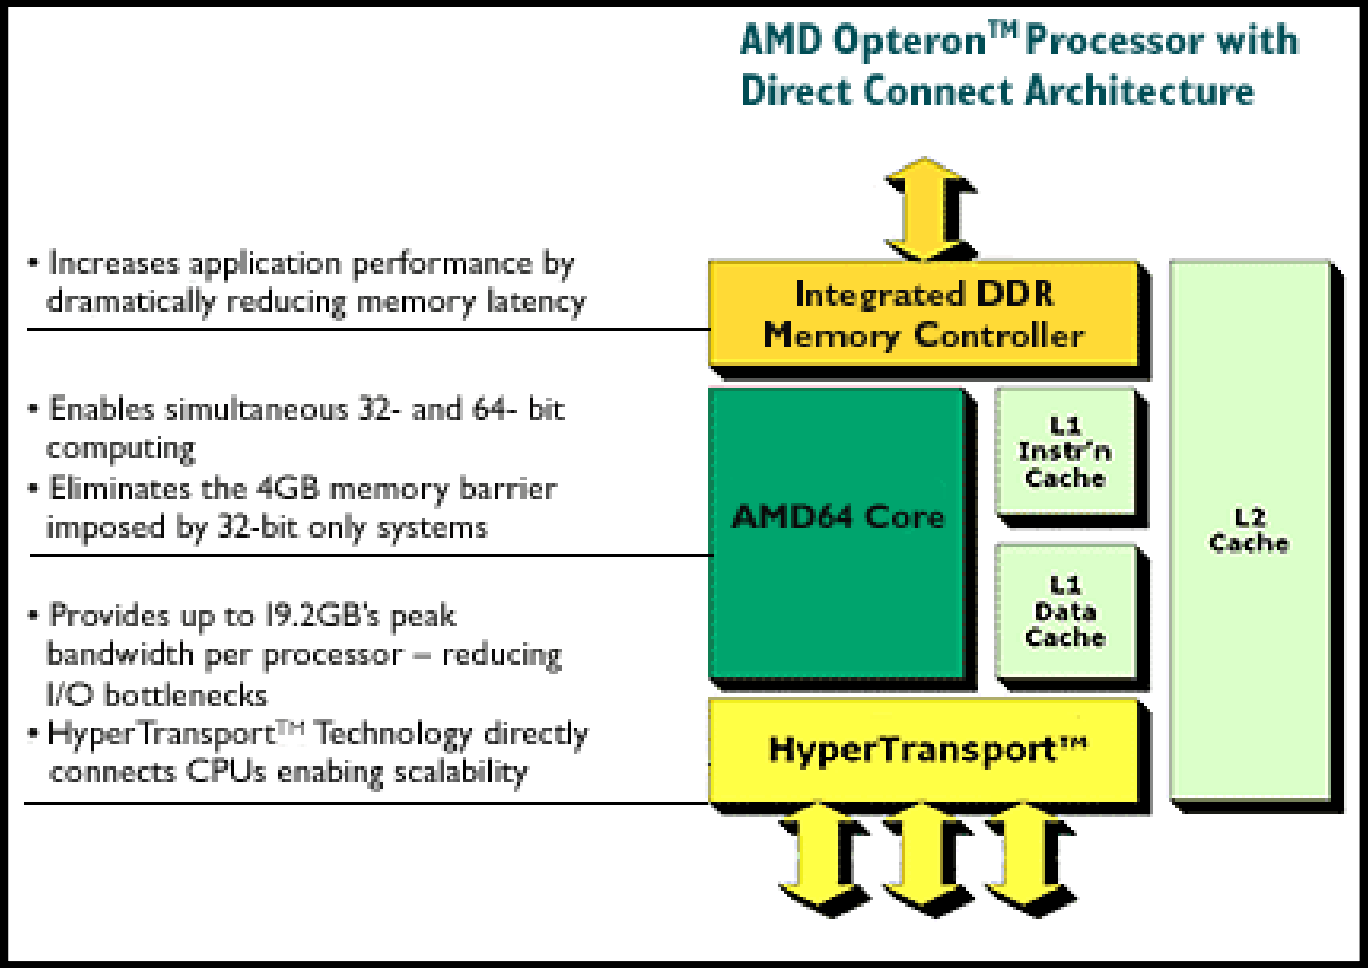

The AMD Opteron contd…..

- Simultaneous 32-bit and 64-bit computing capabilities without sacrificing performance
- Direct Connect Architecture addresses and helps reduce the real challenges and bottlenecks of system architectures
  - Increases memory performance, provides more balanced I/O throughput
  - Allows for more linear symmetrical multiprocessing
- Support of up to three (3) coherent HyperTransport links, providing up to 19.2GB/s peak bandwidth per processor
- 256 Terabytes of memory address space
- Scales from 1-way to 8-way across entire data or compute centers utilizing the same hardware and software infrastructure
- Integrated memory controller reduces latencies during memory access in a SMP server system
  - Memory is directly connected to the CPU providing optimized memory performance
  - Provides low-latency memory access and bandwidth that scales as processors are added
- Low-power processors provides uncompromised performance

The Opteron, Xeon and Itanium processors

	AMD Opteron	Intel Xeon (32-bit)	Intel Itanium 2 (64-bit)
32-bit mode	Native	Native	Through emulation layer
64-bit mode	Native	No	Native
System bus	1.4 – 2.4 GHz, 11.2 – 19.2 GB/s bandwidth	533 MHz, 64 bits wide:up to 4.3 GB/s bandwidth	400 MHz, 128 bits wide: up to 6.4 GB/s bandwidth
Cache (levels 2/3)	1MB/not applicable	512 KB/up to 2MB	256 KB/up to 6MB
Memory bandwidth 2P (4P) system	12.8 GB/s (25.6 GB/s)	4.3 GB/s (N/A)	6.4 GB/s (6.4 GB/s)
SMP capabilities	Up to 8-way	Up to 2-way	Up to 4-way
Direct Connect Architecture	Yes	No	No
HyperTransport Technology	Yes	No	No
Integrated DDR Memory Controller	Yes	No	No

The effect of processor architecture on applications

- Dependent on performance parameters of each application and how the application handles data
- Memory size and data width of the processor execution path are key performance factors for the following application categories:
  - **Sequential versus random requests:** Video coding and streaming versus file-and-print sharing
  - **Logic-based versus load-based requests:** Applications with large, complex algorithms and floating point calculations versus those using simpler ones
  - **Memory set:** 64-bit platforms help improve performance by providing a large addressable memory space for memory-intensive applications

What are 64-bits good for?

- More per-process memory
  - Useful for applications that can run faster by utilizing more memory than the 32-bit user space will allow
- More system-wide memory
  - Useful for systems where total memory usage by all applications exceeds 32-bit limitations
- 64-bit arithmetic

  (and other 64-bit instructions)

But the catch with 64-bit is…

- More memory to account for
  - Additional page table levels
- Larger code size
  - Executable size increase
  - Increased I-cache pressure due to (sometimes) larger instructions/operands
- Larger data size
  - Increased D-cache pressure due to bigger data types, particularly pointers

Drilling down on memory accounting

- What is a TLB?
  - Translation Look-Aside Buffer is a hardware on-chip structure for virtual to physical address translations
  - Supports high performance paged virtual memory system
  - Scarce resource on a processor; limited number of TLB entries
- What is a TLB miss?
  - Miss happens when the virtual page requested isn't covered by the current list of TLB entries
  - TLB miss handling requires reloading of a TLB entry (memory reference) for the virtual page access to continue

Drilling down contd….

- Why is TLB so important w.r.t. 64-bit?
  - Increase in physical memory exacerbates pressure on TLB
  - With 16K page (Linux IA64)
    - 48 GB of process memory requires 3M TLB translations
    - Itanium 2's TLB covers 2 MB of virtual address space
- Huge TLB Support seeks to…
  - Increase TLB coverage and reduce TLB miss rate
  - Reduce process' page table memory requirement
  - Pin data pages in physical memory
- How?
  - Avoid spending execution time in low level TLB (Translation Look-Aside Buffer) miss handling
  - Reduce TLB pressure by making each TLB entry map as large a contiguous memory region as possible
  - Smaller number of large pages rather than large number of small pages

Drilling down contd….

- Linux VM subsystem support of multiple page sizes
  - IA-32 - 4K page size
  - Itanium - user page size configurable at kernel build time: 4K, 8K, 16K or 64K
  - Itanium 2 - concurrent multiple page sizes: 4K, 8K, 16K, 64K, 256K, 1M, 4M, 16M, 256M, 1G and 4G

Payoffs and Results

- For 16 KB to 256 MB page size
  - DTLB miss rate reduced by 65%
  - Data references between misses increases by 280%

- Reduced page table memory
  - For example:
    - 45 GB of shared cache, 8 bytes/PTE
      ~ 3M 16KB regular TLB pages = 22.5 MB
      = 180 256 MB huge TLB pages = 1440 bytes

- Memory pointed to by huge TLB pages is not considered for swapping

- Faster virtual memory access time
    - 16 KB to 256 MB page size
    - 11% improvement for OLTP workload

Additional Optimizations

- Async IO
  - Reduces large number of process context switches due to blocking synchronous I/O
  - Move from I/O helper threads that issues synchronous I/O to submit and completion queues
- Raw vary IO
  - Handle I/O's > 4 KB in single physical I/O request
- Global io_request_lock_reduction
  - Reduce time spent spinning on highly contended lock
  - From 1 global lock to 1 per device
- I/O and Processor improvements
  - number of useable processors supporting up to 16 logical CPUs (or 8 hyperthreaded CPU pairs)
  - Support for a timed semaphore call (semtimedop) reduces time spent spinning on highly contended lock

Always good with 64-bit…

- More registers
- Bigger pages
- Profile-based optimization
- Chip-aware compilers/linkers/tuning tools
- On-chip event counters

Application categories

- Categorize applications to determine which architecture will suit best:
  - 32-bit, 64-bit extension technology, or 64-bit
- Compute-intensive: Vertical and business-critical applications – HPC, database apps, CRM and ERP, e-business applications, etc benefit from 64-bit processors
- Compute/load-balanced: Infrastructure-based apps e.g., internet caching, security, DNS, DHCP, SSL and database front ends. Evaluation for 64-bit benefit should be done
- Standard Infrastructure: Simple file and print sharing, resource sharing and less critical single-use/low volume business apps – No 64-bit benefit

Making the right choice

- Depends upon the characteristics of the processors
  - Xeon (32-bit) family:
    - **Volume optimized**
    - **Ease of use**
    - **Multipurpose**
    - **Widest application acceptance**
    - **Multithread processing (Hyperthreading technology) for higher levels of transaction throughput**
  - Itanium (64-bit) family:
    - **Performance and scalability optimized**
    - **Application tuned**
    - **Task optimized (with increased task-specific features introduced over time)**
    - **Application support dependent**
    - **Extended on-chip memory**
    - **Parallel processing**
    - **Floating point performance**

Making the right choice contd…..

- What the processors are good for…..
  - Xeon (32-bit) family:
    - **Application and Web servers**
    - **32-bit applications**
    - **Small and medium databases (dependent on volume, users, complexity)**
    - **N-way load-balanced servers**
    - **High-density servers (bricks/blades)**
    - **Simple and moderately complex data sets**
    - **Up to 8-processor SMP (with some niche systems beyond 8 way)**
    - **Multithread processing**
    - **Limited and well-defined scale**
  - Itanium (64-bit) family:
    - **Technical and high performance servers and workstations**
    - **64-bit applications focus (32-bit compatible)**
    - **Large databases, especially those constrained by memory residence**
    - **N-way clustered servers**
    - **Up to 32-processor SMP**
    - **Complex data sets**
    - **Multitask processing**
    - **High or unpredictable scalability**
    - **Memory-intensive applications**

Making the right choice contd…..

- **What the processors are good for …..**
  - **AMD Opteron:**
    - For larger memory requirements above 4GB
    - Must manage a large number of concurrent users or application threads
    - Requires real-time encryption and decryption for enhanced security
    - Requires mathematical precision and floating-point performance
    - Requires large, high-power database performance
    - Requires the x86 compatibility
    - Provides digital content creation capabilities such as CAD, CAM, CAE, etc
    - Requires maximum performance for realistic and cinematic consumer experiences including computer games, digital video, and real-time collaboration; and
    - Migrates capabilities previously available on 64-bit workstations to the business, consumer, and hobbyist desktop, including 3D
    - Good for modeling, rendering, animation, simulation, and software development.

Making the right choice contd….

- The enterprise goals
  - Ease of use and flexibility:
    - **Xeon family is optimized to meet broadest range of end-user requirements**
    - **Flexibility to be used across many system types**
    - **Xeon family is optimized for various system configurations – workstations, mid-tier apps or database servers, volume departmental server, ultra-dense Web server.**
    - **Web or application servers, application layers of n-tier architectures in mid-sized deployments (SAP, Siebel, PeopleSoft, firewalls, departmental or business unit databases and general purpose systems not requiring 64-bit high performance)**
  - Enterprise-Level Performance:
    - **Itanium (64-bit) family for high-end technical computing**
    - **Large data sets, where performance would be enhanced by enabling large volumes to reside in processor cache and main memory**
    - **Multithreaded applications where performance is enhanced by true parallelism on the processor**
    - **Data mining against large volumes requiring complex computations**
    - **Workloads that require between 4 and 32 processors in an SMP configuration**

The ultimate decision process

- If there is no strong alignment between the workload type and 64-bit processor design benefits, users should stick to 32-bit platform
  - **Applications on a 32-bit platform within the scalability and performance levels do not need to migrate to 64-bit**
- 64-bit platform will likely to add value if :
  - **The applications have been compiled and optimized in accordance to a 64-bit architecture**
  - **Applications proven to be scalable and reliable on the 64-bit platform**
  - **If scalability is required beyond 8 processors in a single SMP system**
  - **If the application or database uses large memory sets held in processor memory**
  - **If applications use compute-intensive floating point calculations**
  - **If applications are of commercial transaction processing type and if multiple threads in parallel will improve performance**
  - **If server is used as a security gateway where encryption algorithms are processed in real time**
  - **If application is under performing, reaching a wall of scalability and 64-bit optimized version is available, then a migration is beneficial**

Split Configurations of Application and Database Tiers

- provides better manageability and performance and investment protection

Typical Single-Tier Application environment

- Everything on one system
- Easiest and cheapest to implement
- May not provide best performance, flexibility, manageability

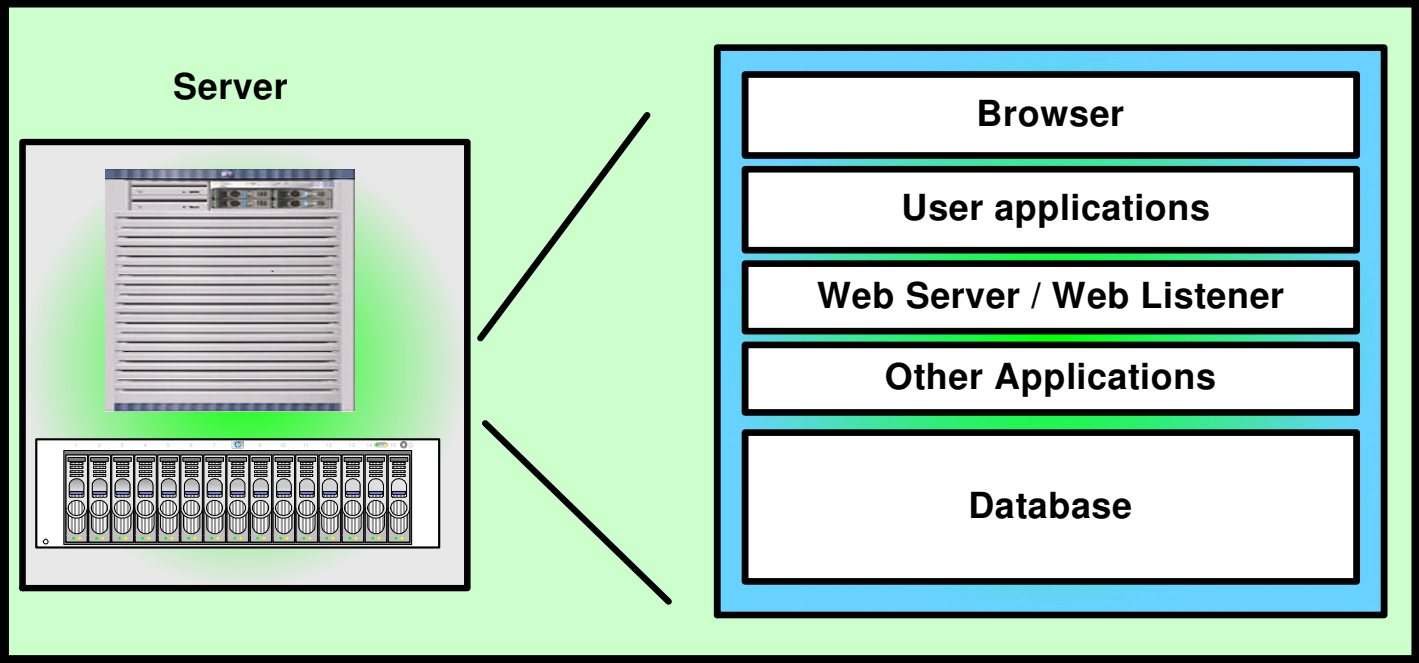

Typical Multi-tier Application environment

- Applications running on different tiers
- Complex to implement
- Better performance flexibility, manageability
- Good for high-end complex applications

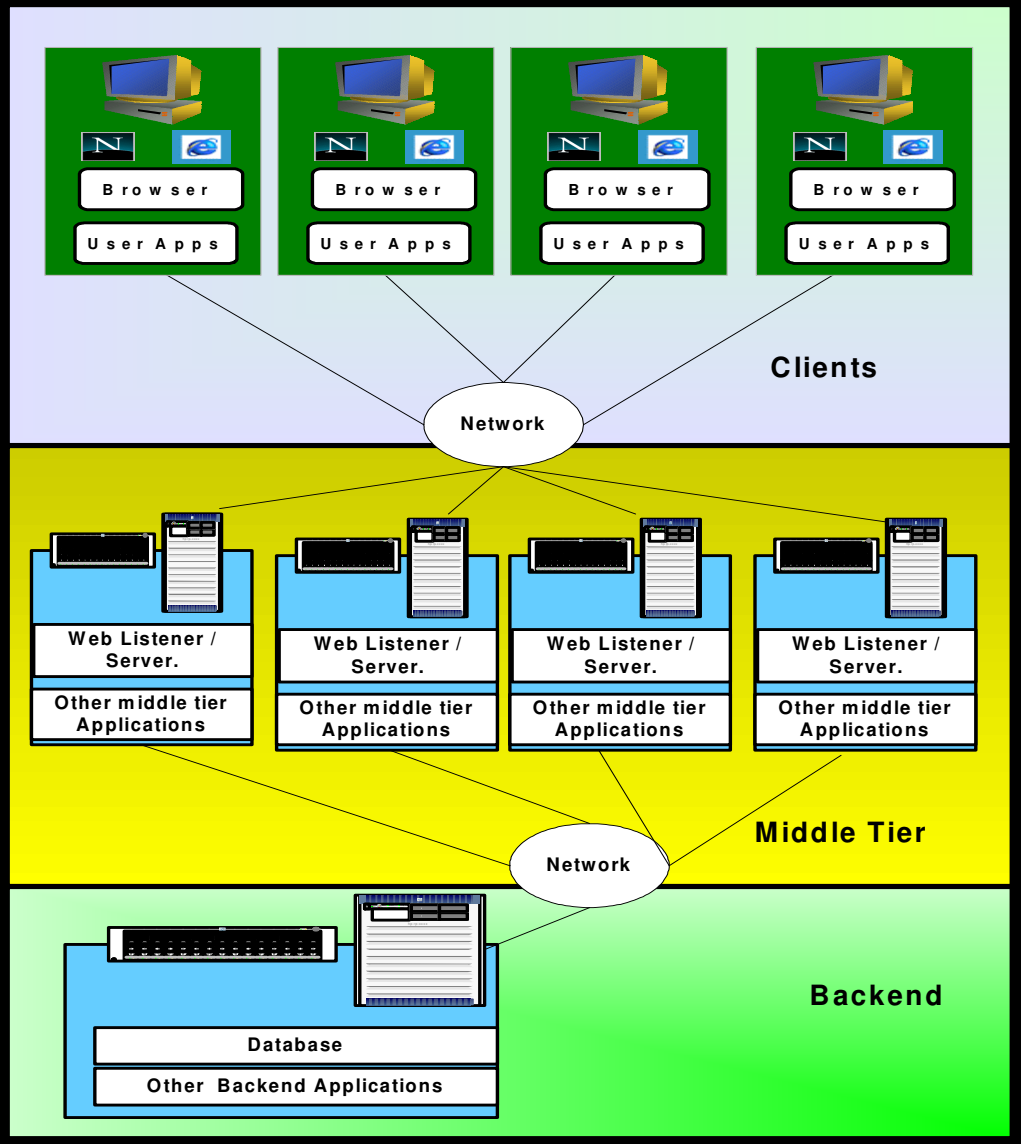

Typical split configuration

- Applications running on a different tier/platform
- Can use most appropriate platform(s) for each tier
- Complex to implement
- Enhanced performance, flexibility, manageability
- Good for high-end applications

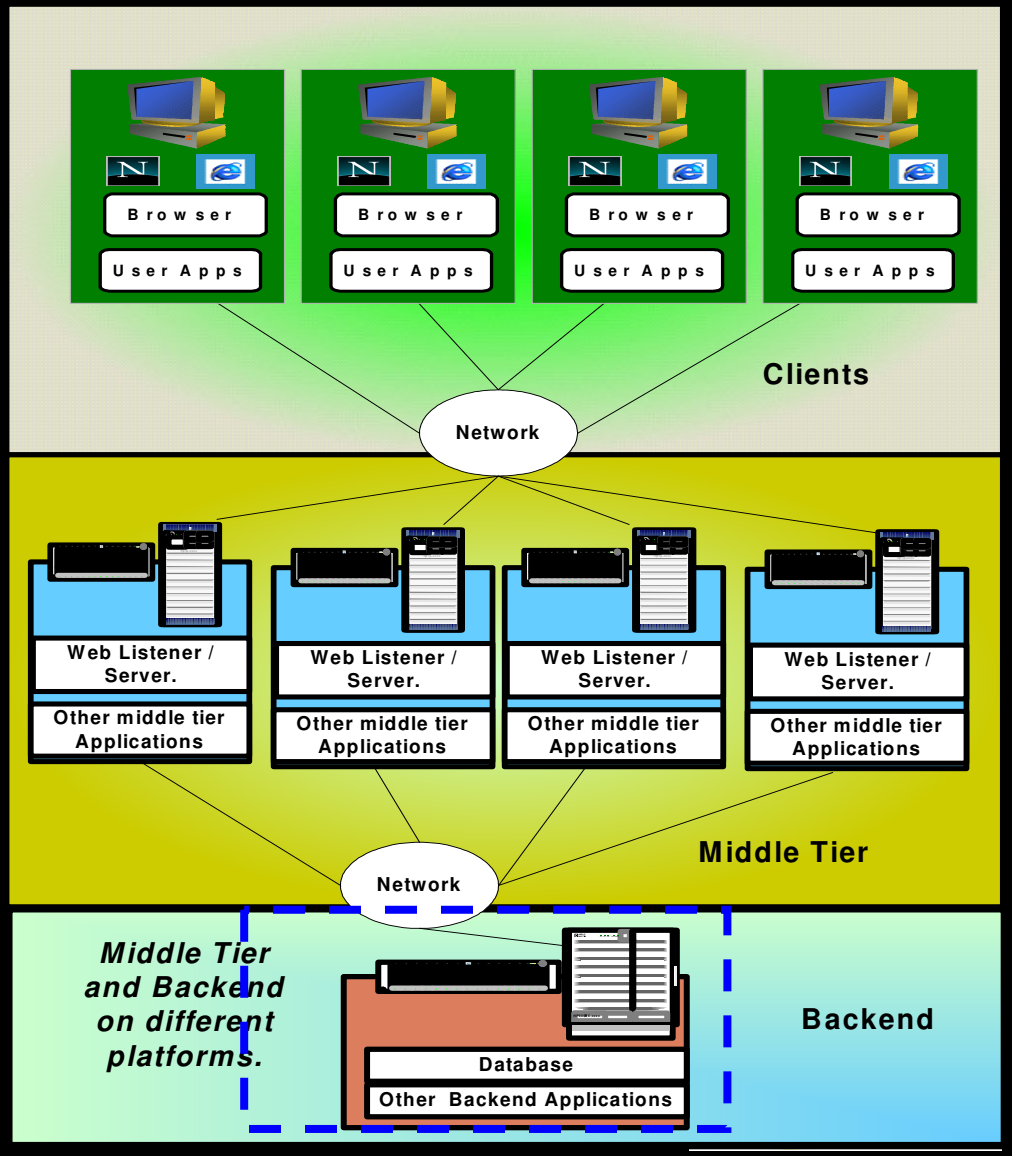

Oracle E-Business Suite (EBS) Overview

- Set of business applications
  - Enables to manage CRM, HRM, ERP, …
- Usually various components interact with single database
- Typically implemented in multi-tier environment

E-Business Suite Components

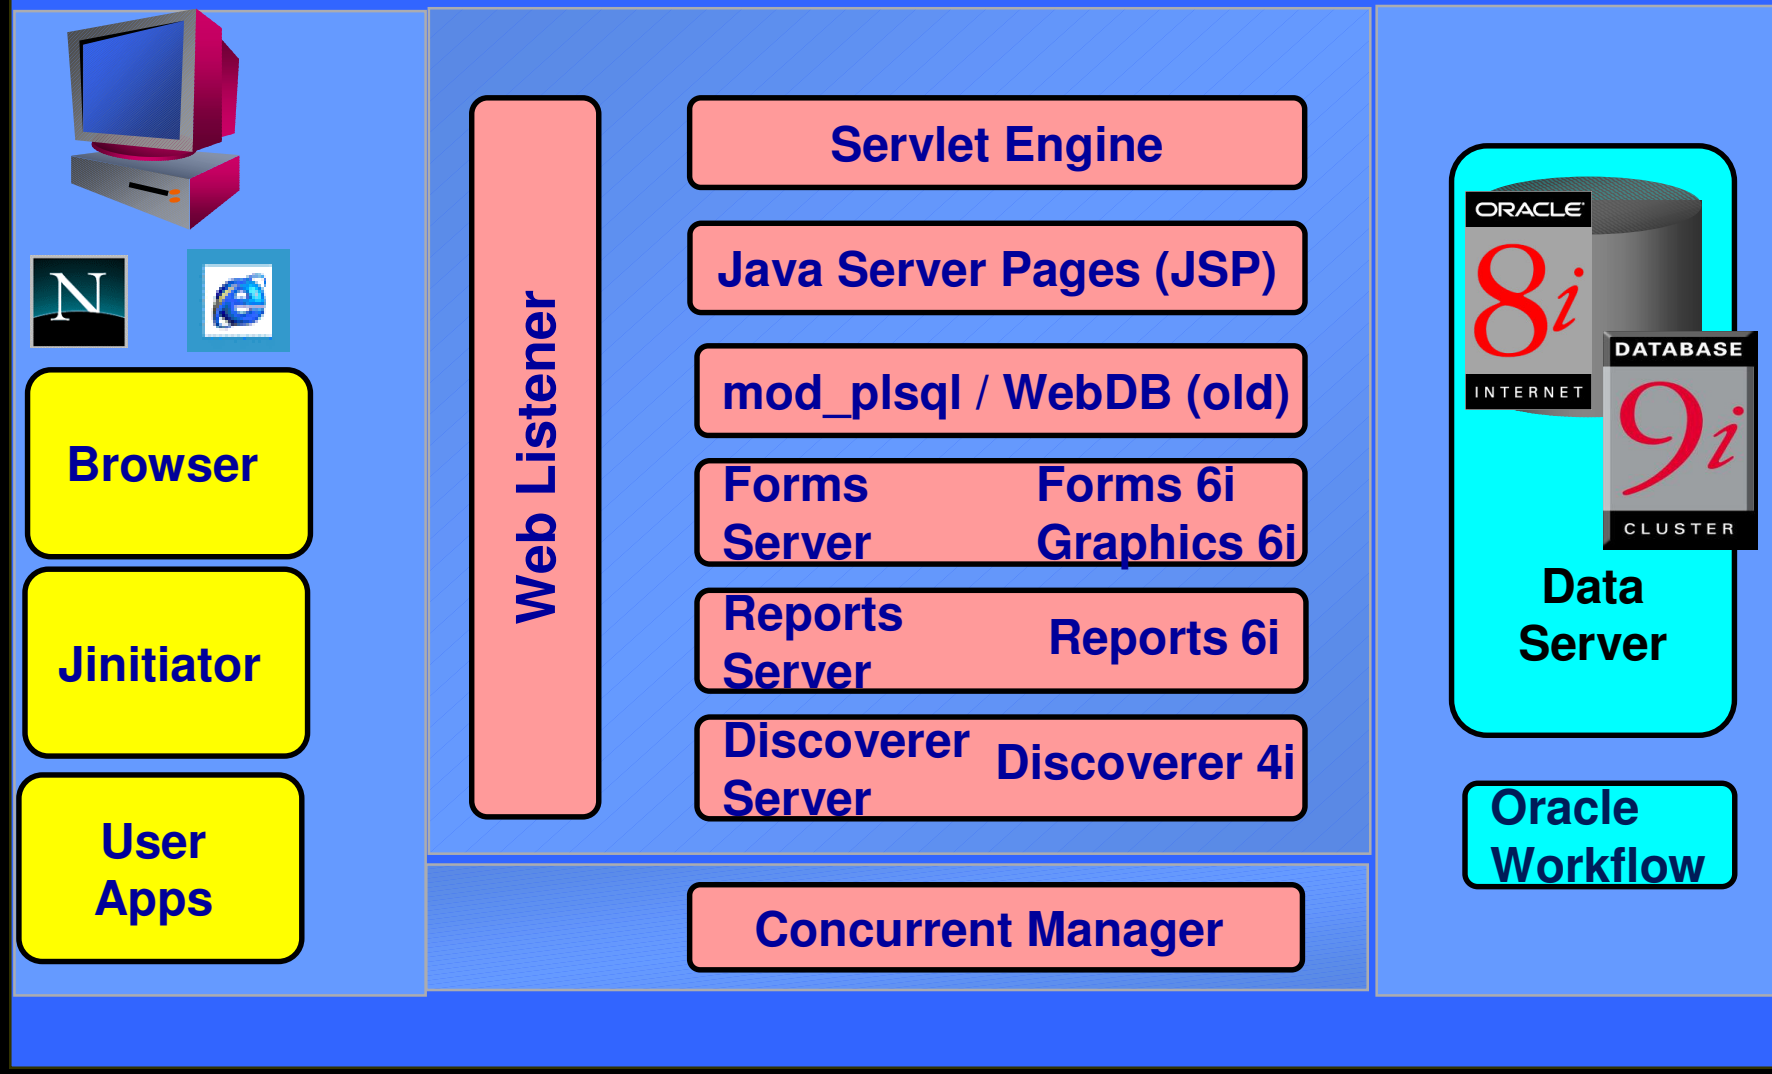

E-Business Suite Implementation

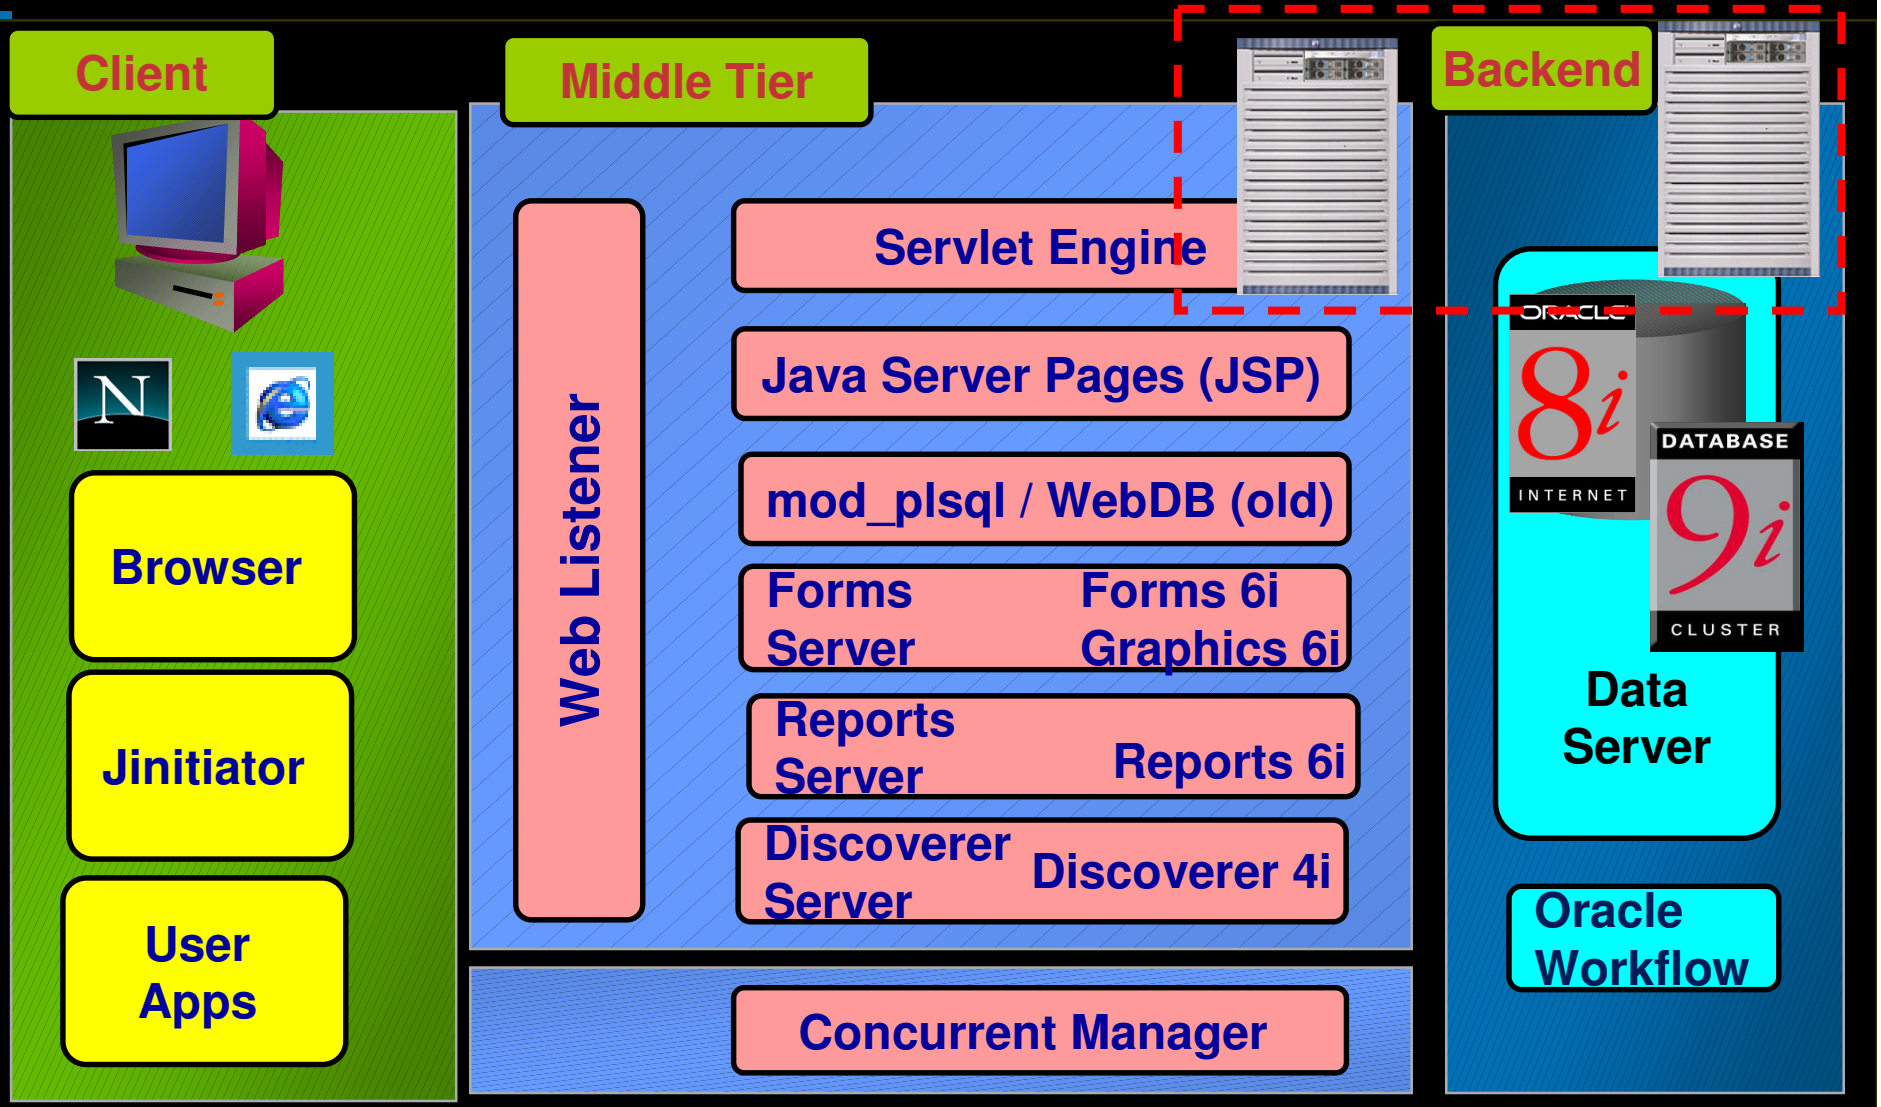

E-Business Suite in Split Config env. using IPF (Single Instance DB)

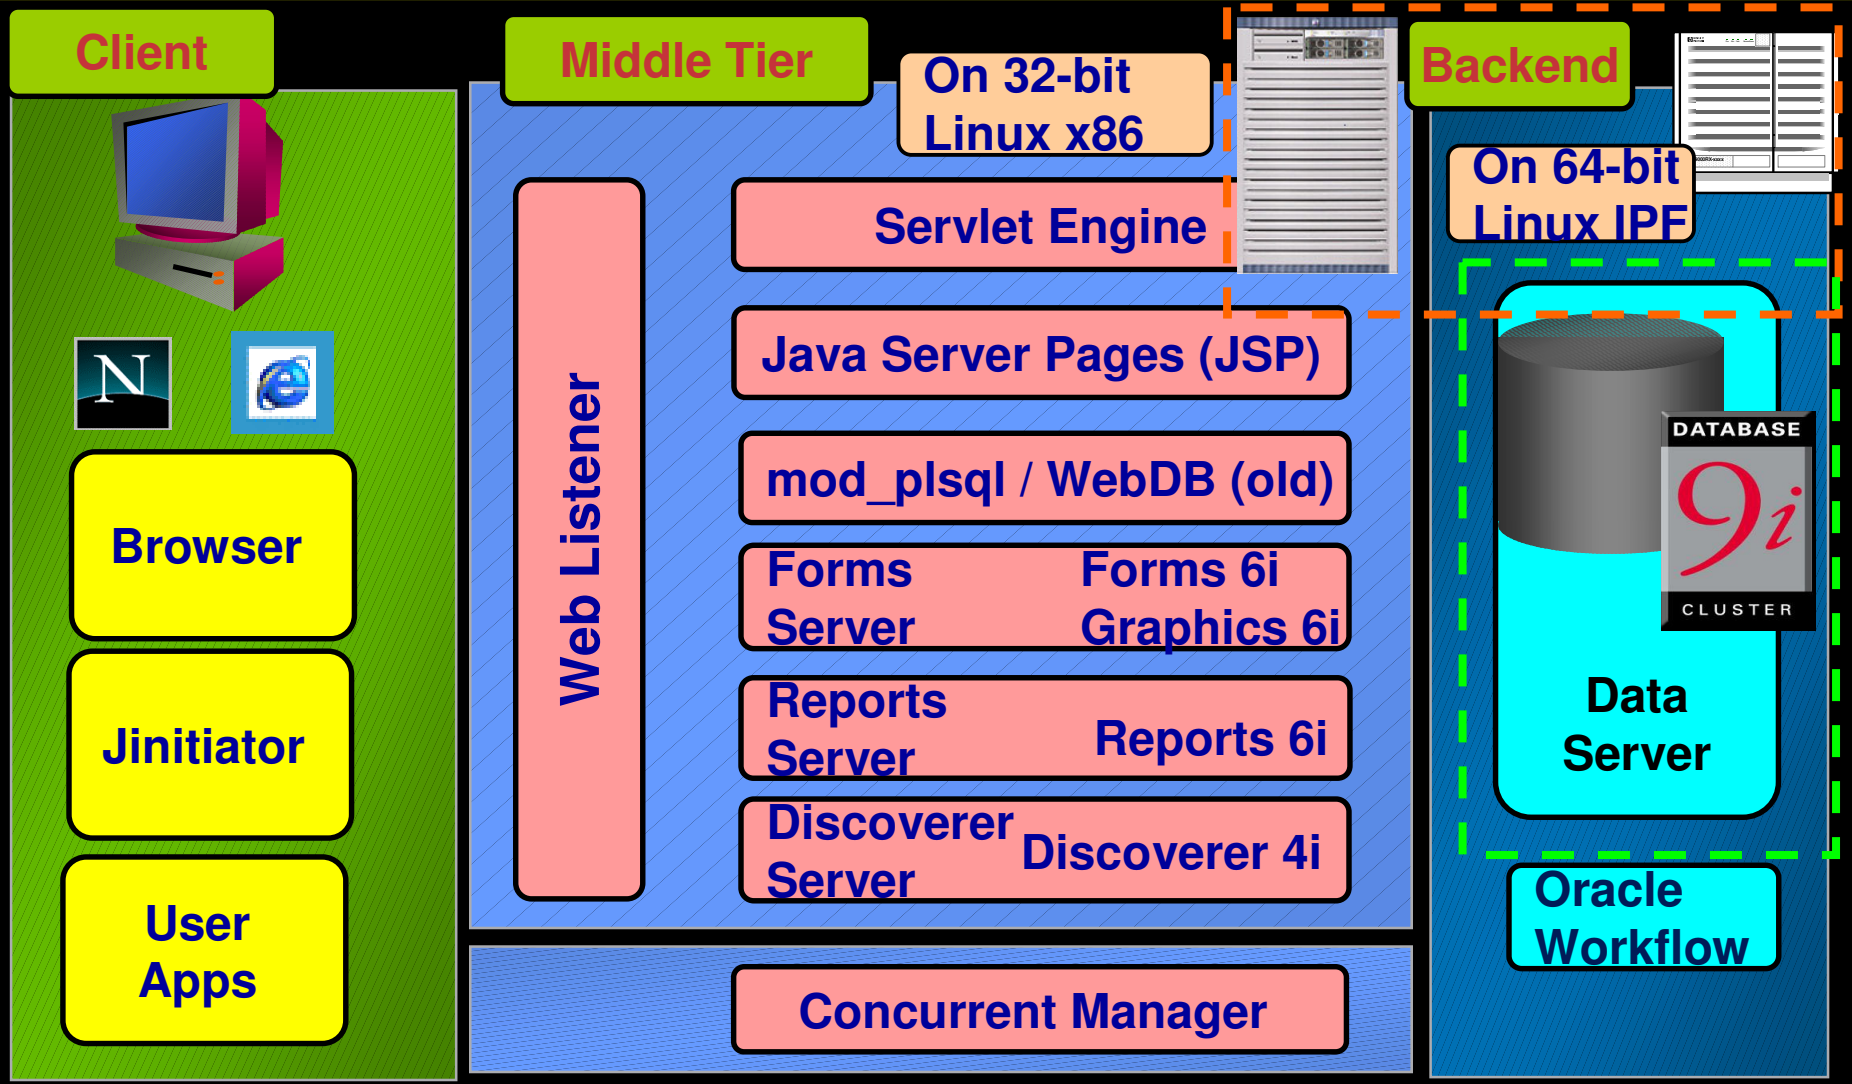

E-Business Suite in Split Config env. using IPF (RAC DB)

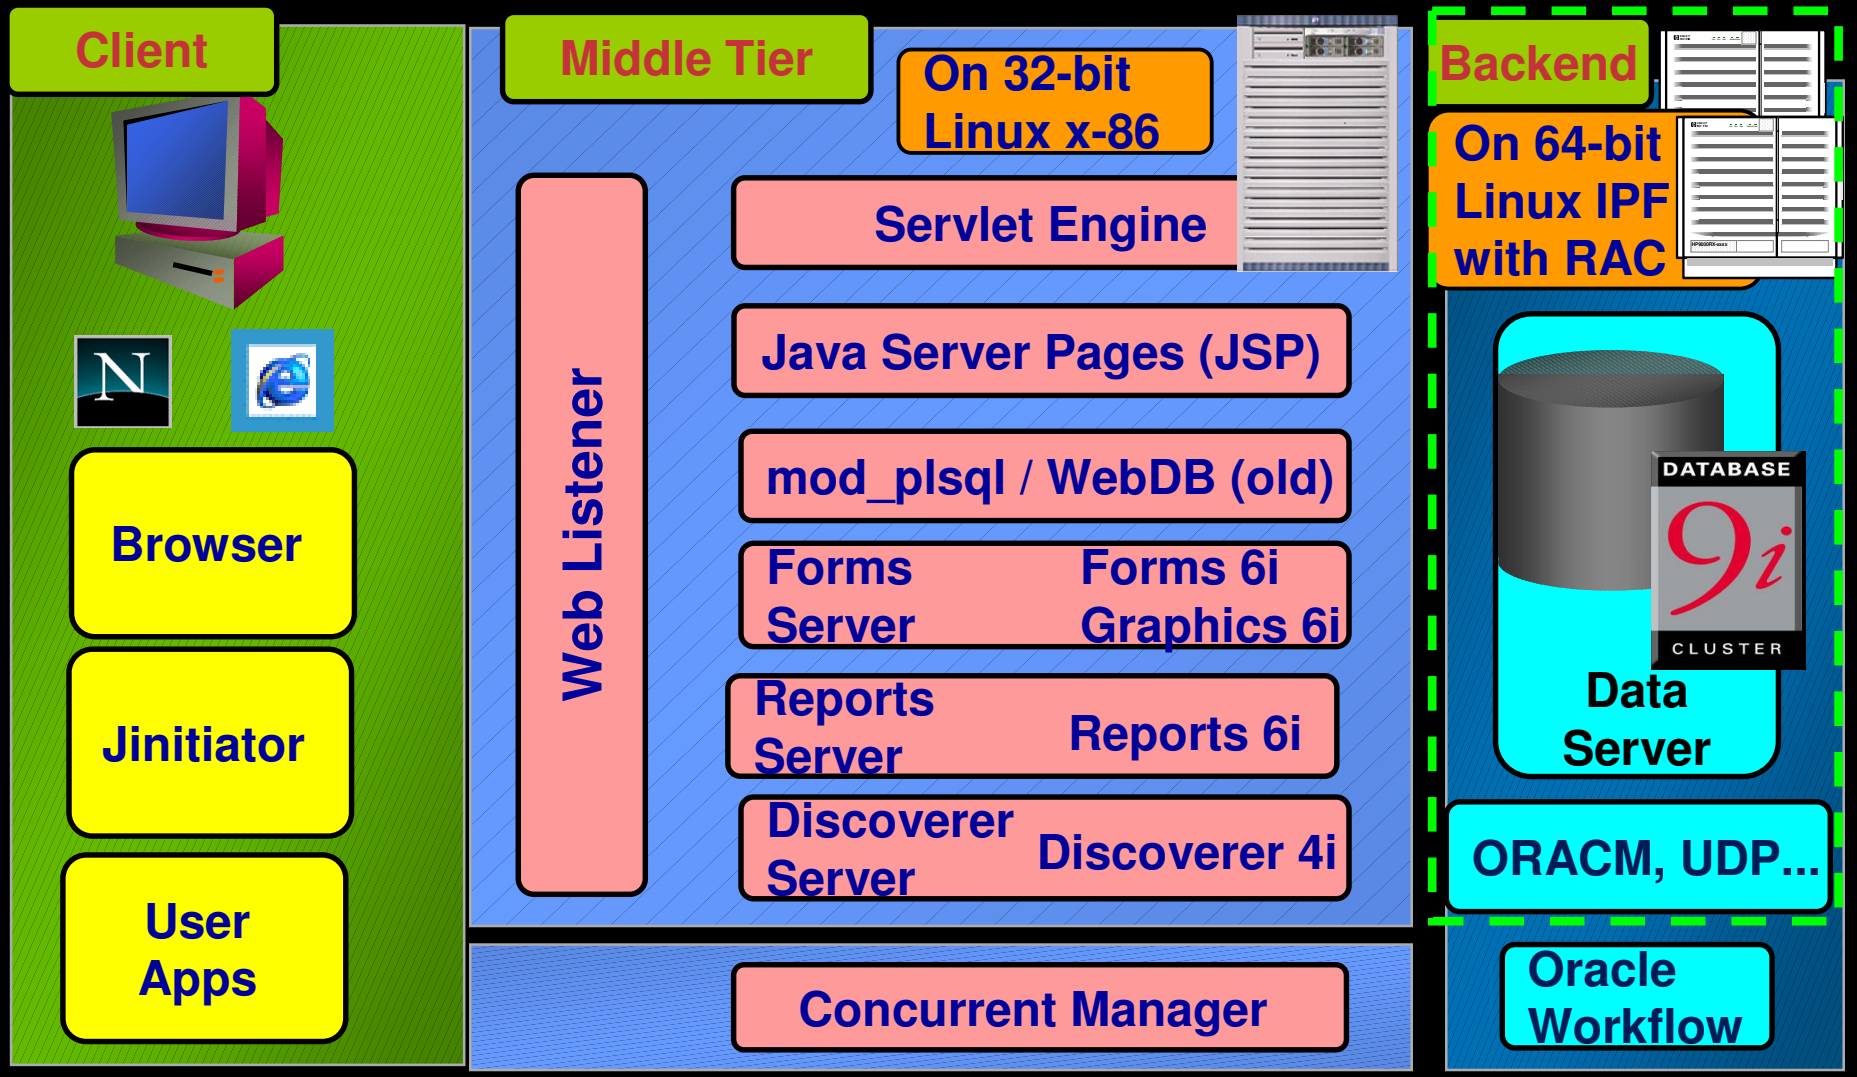

Database tier

- **Setup**
  - H/Ws (Linux IPF Server(s)) for database
  - Network
  - Storage
  - Database Server(s) for IPF
    - Install Linux for IPF and required S/Ws
    - Configure Kernel
    - Install and configure RAC components (if needed)
- **Install database S/W (example: Version 9i-R2)**
- **Configure database**
  - Configuration files (.ora, listener etc)
  - Data files, log files, table space etc
- **Create database**

Middle Tier/Client

- **Setup**
  - H/Ws (Linux 32-bit Server(s)) for apps
  - Network
  - Storage
  - Middle-tier Server(s) – 32-bit Linux x86
    - Install Linux x86 and required S/Ws
    - Configure Kernel
- **EBS**
  - Install
  - Configure EBS configuration files including TNS, etc
- **Client system**
  - Install needed software
  - Configure appropriately
- **Ready to fly**

Why split configuration?

- IPF provides better database performance
- No need to upgrade everything at the same time
  - Can expand gradually
  - Measure, assess and grow
- One can utilize the existing H/W
  - Existing Linux x86 database servers can be utilized in the middle tier
  - 32-bit Applications run with modest performance on x86 server
- Any tier can have mix of x86 and IPF in future

Example split configurations

- Deploying E Business Suite Release 11*i* in a Split Configuration Database Tier model with a combination of Applications supported platforms is supported with no additional certification.
- Currently the Application Supported Platforms are:
  - **Linux x86**
  - **HP Tru64**
  - **HP-UX PA-RISC**
  - **Windows 2000/Windows Server 2003**

  And others……..
- When E-Business Suite (Forms, Reports, concurrent programs, etc) are installed on a supported Application platform, the following additional platforms may be used for the Database Tier:
  - **Red Hat Enterprise AS/ES 2.1 (Itanium)**
  - **Red Hat Enterprise AS/ES 3 (Itanium) Projected**
  - **Red Hat Enterprise AS/ES 3 Linux x86-64 Projected**
  - **SuSE SLES8 (Itanium) Projected**
  - **SuSE SLES8 Linux x86-64 Projected**
  - **HP-UX Itanium**

  And others……
- For detailed information, please refer to *"metalink.oracle.com"*

Work in Progress…..

- Investigate emerging 32-bit and 64-bit Linux platform architectures and solutions with Oracle 10g
- Re-analyze choices for 32-bit, 64-bit extensions and 64-bit technology for Oracle applications and keeping up with the IT trend
- Extending support for split configurations combining various middle-tier and database-tier platforms
- Analyze migration paths to various supported split configurations for better application performance and ease of manageability

Acknowledgments

- Hewlett Packard: Development Alliance Lab
- Hewlett Packard: Performance Benchmarking Group
- Oracle: Product Line Engineering
- Oracle: Release documentation Group
- Oracle: Product Management and Alliance
- Oracle: Server Technologies Division
- Oracle: Platform Technologies Division

"I want to know God's thoughts – the rest are details"

- Albert Einstein

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP