



# The New Generation of HP Integrity Servers and Oracle perspective

Tomislav Alpeza  
Presales Consultant BCS/SWD





# What's announced..

## Superdome 2



The ultimate mission-critical consolidation platform

## Integrity Server Blades

World's first scale-up blades built on the industry's #1 blade infrastructure



Delivered with Integrity


## BladeSystem Matrix with HP-UX



Ultimate platform for shared services, now enhanced for mission-critical environments

## rx2800 i2 Integrity Server

2-socket, 8-core scalability in 3x less compute density - without sacrificing RAS



## HP-UX 11i v3

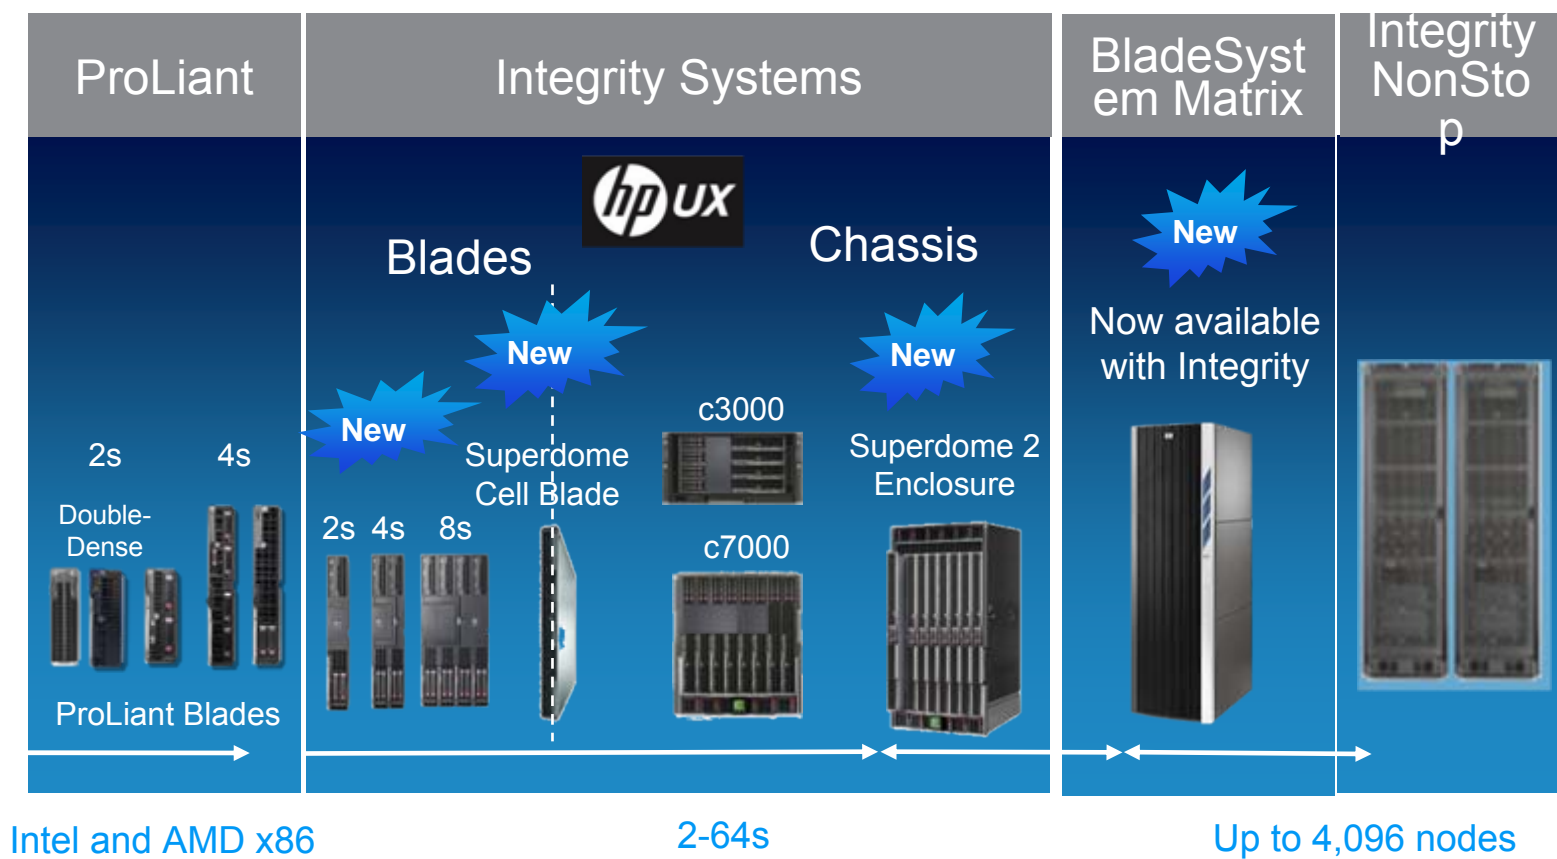
The mission-critical foundation for converging resiliency and optimization





# Unified blade architecture from x86 to Superdome

Simplify by consolidating applications on a common platform










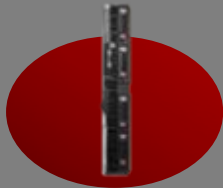
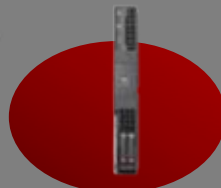

**Train once, certify once, deploy once**






# COMMON MODULAR BUILDING BLOCKS

– Simplicity through standardization

|                            |                                                                                                     |                                                                                                        |                                                                                                                 |                                                                                                   |
|----------------------------|-----------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| Common management          | Matrix operating environment delivered by Insight Dynamics                                          |                     | Onboard Administrator, iLO3  |                                                                                                   |
| Common networking          | Virtual Connect  |                                                                                                        |                                                                                                                 |                                                                                                   |
| Common enclosure           | c3000              | c7000                | Superdome 2 enclosure        | Common spares  |
| Common server architecture | ProLiant blades  | Integrity blades  | Cell blades                |                                                                                                   |

# HP INTEGRITY BL8X0C I2 SERVER BLADES

World's first scale-up blades built on the industry's #1 blade infrastructure



- Up to 8 socket/32 cores Intel® Itanium® processor 9300 series
- Up to 384GB DIMMs
- Up to 16 x 10 GbE (Flex-10) NICs

## Common Architecture from x86 to Superdome

- Mix and match new and existing Integrity, ProLiant and StorageWorks storage blades within the same enclosure
- 2.5x compute density compared to traditional rack mount servers

## Blade Link

- Scale up, out and within, scale more and scale linear
- Combine multiple blades to create 2, 4 and 8 socket systems

## HP Virtual Connect Flex-

- Network scalability and configuration flexibility
- Up to 20x increase in networking bandwidth
- Virtually connect LAN, SAN, facilities, etc.

**Flexible mission-critical server blades combined with the efficiency of HP BladeSystem to accelerate IT effectiveness**

# NEW INTEGRITY SERVER BLADES OVERVIEW

## Management

- Integrity iLO 3
- Integrity iLO 3 Advanced Pack
- Integrated VGA console
- c-Class Onboard Administrator (firmware v2.0)

## I/O subsystem per blade

- 5X I/O BW increase over previous generation
- Integrated p410i RAID controller
- 2 dual-port 10GbE Flex-10 NICs
- 3 PCIe G2 mezzanine slots

## Operating system support

- HP-UX 11i v3
- OpenVMS 8.4 (CQ3 2010)
- Windows 2008 R2 (Q3/Q4 2010)

3-year,  
next-day, on-site  
base warranty

## Processors and chipset

- Intel Itanium 9300-series processors (Tukwila family)
- Intel E7500 Scalable Memory Buffer (Mill Brook)
- Intel E7500 IOH (Boxboro)
- Intel ICH10 south bridge

## Memory

- ~6X memory BW increase over previous generation
- 24 PC3-8500 DIMM sockets
- 192 GB capacity per blade with 8GB DIMMs



## Form factor

- Full-height c-Class form factor
  - Single wide BL860c i2 — 2S
  - Double wide BL870c i2 — 4S
  - Quad wide BL890c i2 — 8S
- Supported in c3000 and c7000 enclosures

## High availability

- Memory double chip spare
- Redundant hot-plug power
- Redundant hot-swappable fans
- Internal SAS RAID
- Processor deallocation on failure

## Additional I/O options

- Two hot-plug SFF SAS HDDs per blade
- Partner blade support (BL860/870) – disk, tape

# Blade Link

Linear scalability with industry's first 2-4-8 socket UNIX server blades

Blade Link combines multiple blades into a single, scalable system



**Scale**

Up, Out and Within

**Scale**

More

Only 8-Socket UNIX blade in industry standard blade enclosure

**Scale**

Linear

System resources grow evenly across CPU, memory, I/O, and etc

| CPU    | 2s/8c     | 4s/16c             | 8s/32c             |
|--------|-----------|--------------------|--------------------|
| Memory | 96GB      | <b>X 2 =</b> 192GB | <b>X 2 =</b> 384GB |
| LAN    | 4 x 10GbE | 8 x 10GbE          | 16 x 10GbE         |
| HDDs   | 2 Slots   | 4 Slots            | 8 Slots            |

**8 socket system at 2x the performance in half the footprint**



# BUSINESS BENEFITS OF BLADE SCALE ARCHITECTURE

THE TRUE VALUE OF A MISSION-CRITICAL INFRASTRUCTURE  
Business outcomes

Up to 100% application uptime

>2x improvement in resource utilization

Reduce deployment times by up to 50%

30% less power per core

Half the data center footprint

Integrity server blades based on Blade Scale Architecture

Per-socket performance increase

- Up to 2.7x SPECint\_rate\_base2000
- 3.8x SPECfp\_rate\_base2006
- 2.5x SPECsfs2008\_nfs.v3
- 9.1x STREAM Triad

## System Performance



- 2 – 9X on Benchmarks
- Up to 40x for our historic customers



# Introducing Superdome 2

The ultimate mission-critical consolidation platform



- Up to eight 2s cell blades (16s/enclosure)
- 32 DIMM sockets (2TB with 8GB DIMMS)
- Up to 24 mezz. & 96 stand up I/O slots
- 18U in standard HP Rack
- 64s capable with 4 base enclosures
- Programmable active door display

## Common Architecture from x86 to

- Common spares: power supplies, fans & I/O
- Modular, front-back serviceable racks
- Common management for entire infrastructure
- Zero-to-managed in minutes

## Superdome 2 Crossbar Fabric for Extreme Scalability and Reliability

- Independent I/O scaling to meet any workload
- Only Unix system with end-to-end multi-pathing
- Boosts infrastructure reliability by 450%

## Power-on-Once Technology

- **Superdome 2 Analysis Engine:** proactive error detection and prescriptive recommendations
- **End-to-End Transaction Retry:** transactions tracked, retried and rerouted to completion
- **Online Optimization and Repair:** tool-free serviceability + single-click firmware upgrades

**Extreme scalability, on-demand modularity and unquestioned reliability**



# SUPERDOME 2

– The ultimate mission-critical consolidation platform



### Blades/ Processors

- Up to eight 2S Superdome blades per enclosure

### Memory

- 2TB memory capacity (per enclosure, 8GB DIMMS)

### Form Factor

- 18U in standard HP rack

### Networking

- 32 integrated 10Gb Ethernet NICs
- 8 switch bays
- 1GE Manageability LAN

### I/O Slots

- 24 mezzanine slots
- Up to 96 stand-up I/O slots with IO expansion

### Management


Onboard Administrator: HP Insight display

### Partitioning

- 8 nPars (up to 32 with 4 enclosures)
- vPars (16 per nPar), HPvm

### Bandwidth

- 1.2TB/s memory BW, .8TB/s I/O BW (64 socket)



**SD64**  
32 Cell blades  
(36U, dual 19" racks)

### Tukwila

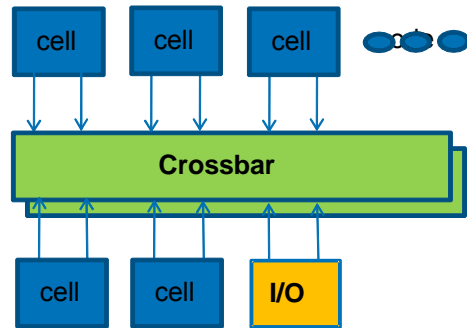
|                     |                |
|---------------------|----------------|
| <b>64</b>           | <b>Sockets</b> |
| <b>256</b>          | <b>Cores</b>   |
| <b>512</b>          | <b>Threads</b> |
| <b>8* TB</b>        | <b>Memory</b>  |
| <b>128 Internal</b> | <b>10 GbE</b>  |
| <b>96 Internal</b>  | <b>PCI-E</b>   |
| <b>96 IOX</b>       | <b>PCI-E</b>   |



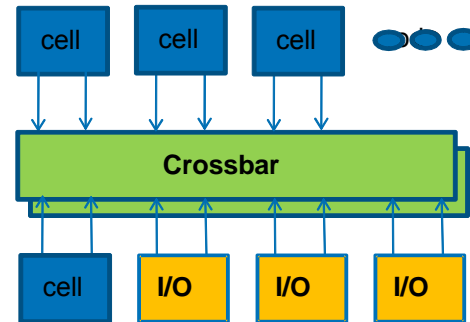


# SUPERDOME 2 CROSSBAR FABRIC

Independent I/O scaling with fault tolerant crossbar



Compute Intensive



I/O Intensive

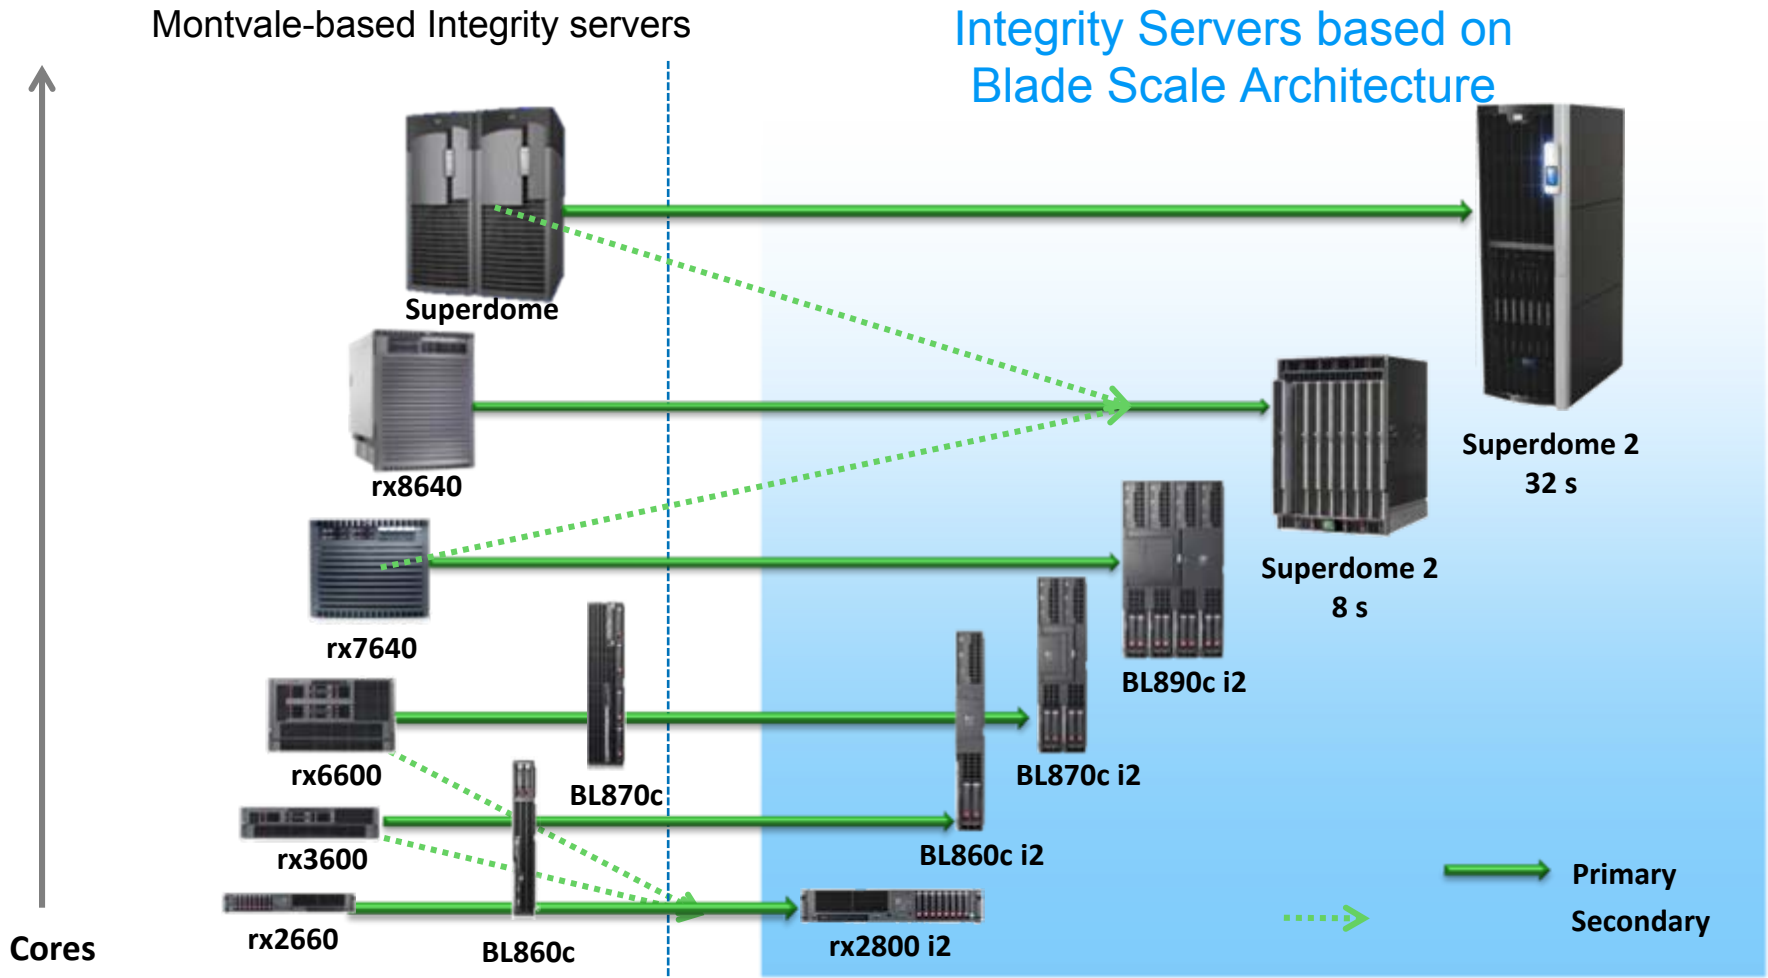
## Independent Scaling

- Flexibly adjust compute and I/O mix
- Only Unix server to scale I/O independently from CPU
- Enables optimal configuration for every workload

## Fault Tolerant Crossbar

- Multiple crossbars enable fully redundant data path
- Crossbars blade modules provide simple hot-plug serviceability

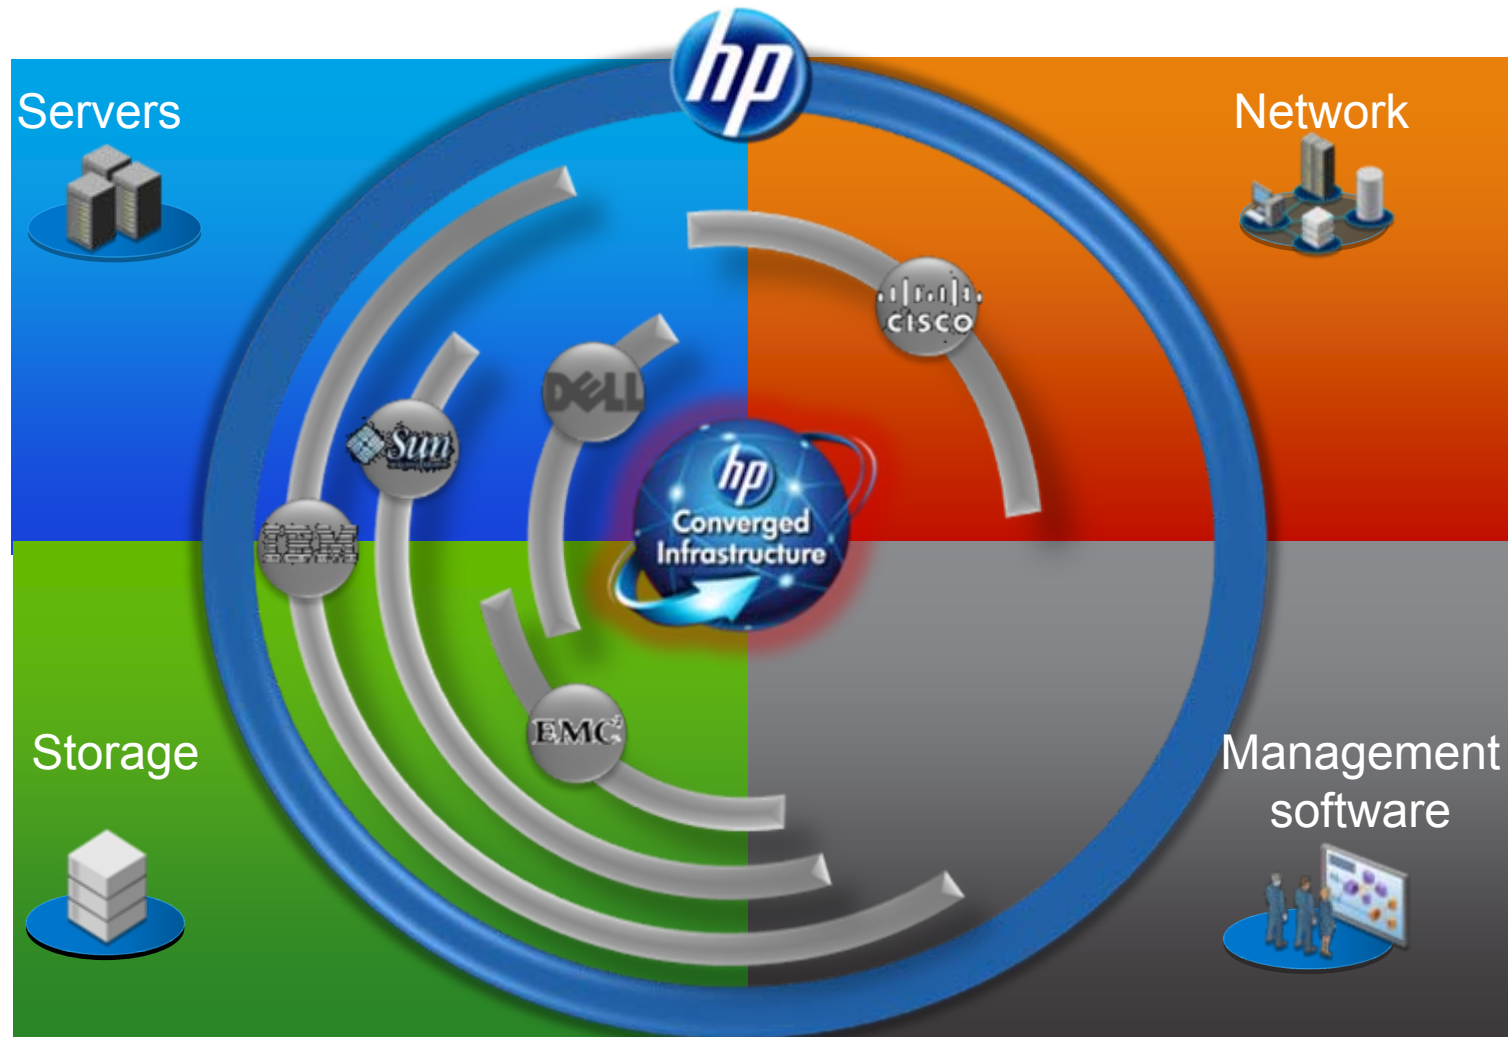
# Positioning the new vs. current servers



>2x performance per socket and much better price/performance

# The HP Advantage

The **Most Complete** IT Portfolio



# HP-UX 11i SOFTWARE POLICIES

– Simple, predictable, sustainable investments

## Licensing



Now it's per socket

- Easy to understand
- Fewer licenses to manage


## Pricing



Not based on core counts

- Prices will be the same, or slightly higher, than today's Montvale socket price (which means per core price is 50% less)

## Investment protection




Maintains the value

- Customers can “trade-in” existing software licenses for new licenses when upgrading servers\*
- “A socket for a socket” – 2 per core licenses = 1 per socket licenses
- Retains the value of prior investments

\* Software support contract required

# IBM & ORACLE SOFTWARE LICENSING REQUIREMENTS FOR HP INTEGRITY SERVERS:


- As of February 16th, 2010 Oracle is charging IBM Power7 customers one full license per processor core, consistent with their POWER6 license charge. All Intel processors, including Itanium, still have a core factor of .5.<sup>1</sup>
- IBM charges a 20% higher license fee for its own software (including DB2) running on POWER7 (POWER 770/780) systems vs. Power5 or Itanium-based Servers.<sup>2</sup>



**Oracle:**  
 $0.5 \text{ Core Factor} * 4 \text{ Cores} = \mathbf{2.0 \text{ Licenses}}$

**IBM Software:**  
 $100 \text{ PVU's} * 4 \text{ Cores} = \mathbf{400 \text{ PVU's Required}}$

|                    | IBM SW        | Oracle |
|--------------------|---------------|--------|
| IBM Power7         | Up to 120 PVU | 1.0    |
| RISC (Dual Core)   | 100 PVU       | .75    |
| Intel Itanium 93XX | 100 PVU       | .5     |



**Oracle:**  
 $1.0 \text{ Core Factor} * 8 \text{ Cores} = \mathbf{8.0 \text{ Licenses}}$

**IBM Software:**  
 $120 \text{ PVU's} * 8 \text{ Cores} = \mathbf{960 \text{ PVU's Required}}$

Licenses for IBM P6 cost:

- 200% more per core than Itanium running Oracle
- 20% more per core than Itanium running IBM SW

<sup>1</sup> <http://www.oracle.com/corporate/contracts/library/processor-core-factor-table.pdf> (02/16/2010)  
<sup>2</sup> [http://www-01.ibm.com/software/lotus/passportadvantage/pvu\\_licensing\\_for\\_customers.html](http://www-01.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html)



# ORACLE DATABASE AND MIDDLEWARE FOR HP-UX



## NGIS Blade Scale Architecture

- Oracle Database and OFM versions
  - If supported on HP-UX 11iV3 → supported for NGIS
  - No Oracle certifications required
- HP APS lab has been using NGIS servers for several months
  - testing Oracle Database (single-instance and RAC)
  - testing OFM (on physical servers and HP VMs)
- “NGIS support for Oracle RDBMS” best practices white paper in preparation
- Oracle per-core database pricing (0.5) the same as current Integrity platforms





## NGIS Blade Scale Architecture

### – For major “managed” applications

- Focused on E-business suite, PeopleSoft, Siebel, J.D. Edwards, Retail
- No certifications needed
- Applications supported on 11iV3 will run on NGIS at launch
- New upcoming releases of these applications will be fully supported on 11iV3

### – Results from testing

- EBS seeing about 10% performance increase per core
- Scaling testing shows more than 2X performance per socket

### – Other applications

- Check availability matrix



# ORACLE ON HP-UX 11.31



– BL8x0-i2 servers are just another “HP-UX 11.31 Itanium” platform for Oracle

- Database 10gR2 requires “-ignoreSysPreReqs” option to Universal Installer
- HP-UX 11.31.1003 contains all the HP-UX patches required for currently shipping Oracle products
  - Exception: PHKL\_40381 for Database 11.2.0.1
- Kernel parameter recommendations remain unchanged

– Java Support

- Base HP-UX install includes 1.4.2, 5.0, and 6.0
- JVM 1.4.2 included with Oracle 10gR2 is outdated; runInstaller may require “-jreLoc /opt/java1.4” option to function properly.

# ORACLE ON HP-UX 11.23

– BL8x0-i2 servers require HP-UX 11.31

- HP-UX 11.23 support available as HPVM Guest O/S
- Patches required on guest: PHCO\_40685 and PHKL\_40684

– Oracle Database 11.2.0 is NOT supported on HP-UX 11.23






# PREVIEW: TWO SIMPLE RULES REMEMBER THESE IF NOTHING ELSE!

1. Be consistent in configuring Oracle and HP-UX
  - Ensure sufficient CLM or disable ccNUMA optimizations
  - Especially with older Oracle/HP-UX versions!
2. We highly recommend that you disable Oracle's ccNUMA optimizations when dynamic resource allocation will be used; otherwise:
  - Need to ensure survival of all "original localities"
  - Dynamically adding resources tends to de-optimize Oracle optimizations



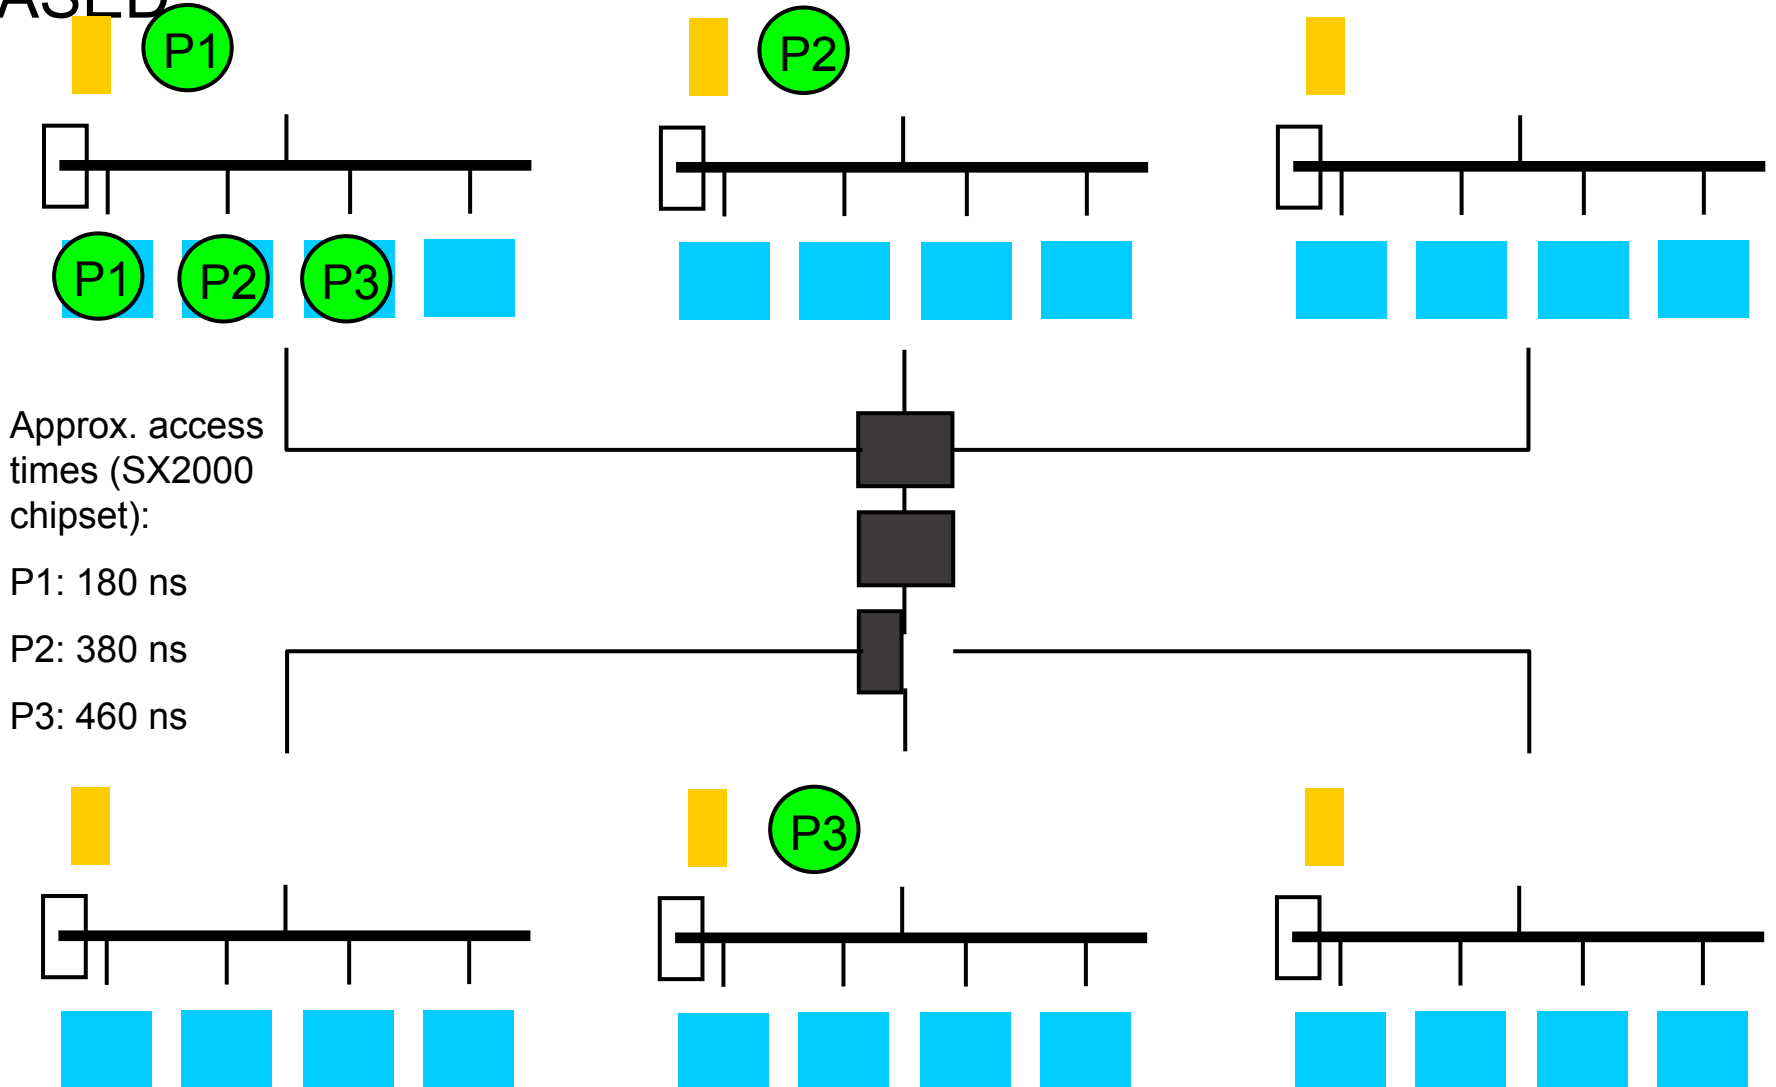
# TRADITIONAL SERVER DESIGN

## UNIFORM MEMORY ACCESS



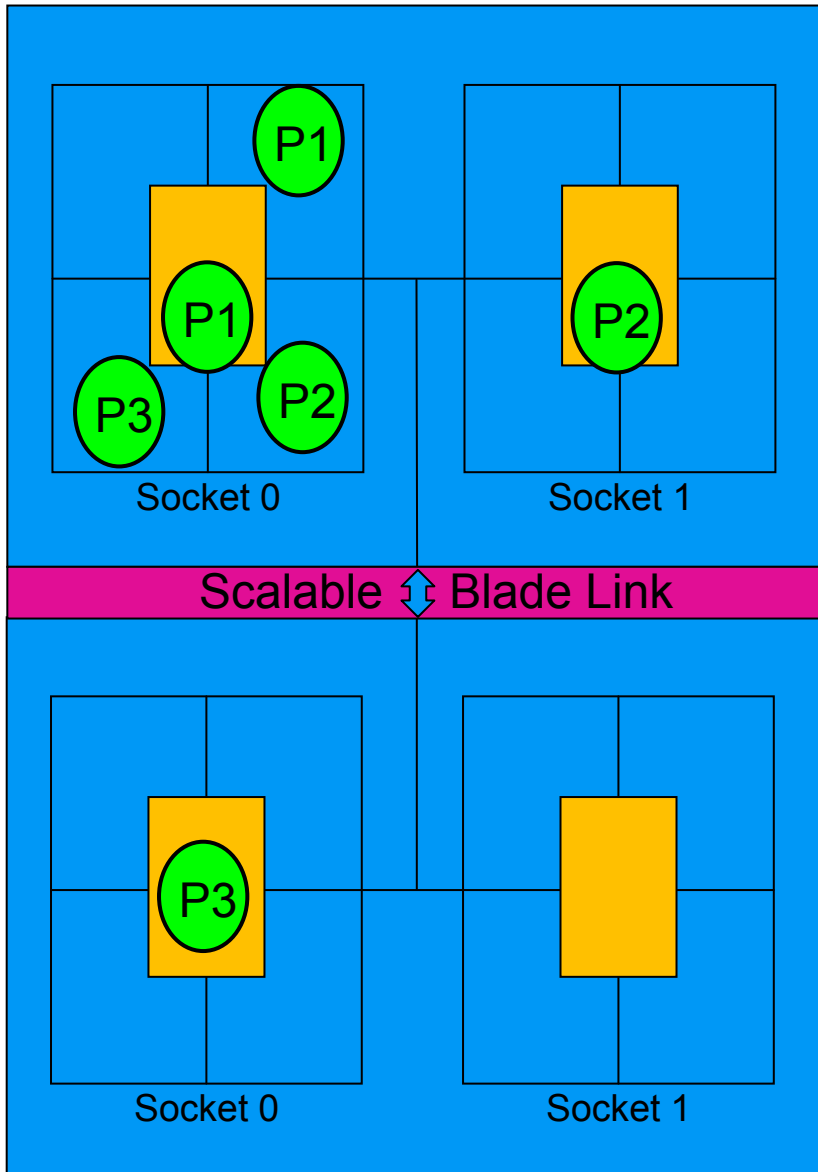
# SCALING UP WITHOUT THE BUS BOTTLENECK

## NON-UNIFORM MEMORY ACCESS (NUMA) – CELL-BASED



# CCNUMA IN A BLADE ARCHITECTURE

## 2<sup>ND</sup>-GENERATION INTEGRITY BLADES



- Each socket is a locality domain (LDOM)
- In one-blade or two-blade server, every LDOM at most one “hop”
  - P1: ~175ns
  - P2: ~245ns
  - P3: ~263ns



# ORACLE CCNUMA OPTIMIZATIONS

## ORACLE DATABASE – SUPPORTED VERSIONS

- 10gR2 was the first to support NUMA optimizations on HP-UX
  - 10.2.0.3 : optimizations enabled by default
  - 10.2.0.4 : optimizations enabled by default BUT DO NOT WORK (bug 9668940). Optimizations should be explicitly disabled.
  - 10.2.0.5 (future) : optimizations off by default; 9668940 will be fixed.
- 11gR1 (11.1.0.6 and 11.1.0.7) : optimizations on by default
- 10gR2, 11gR1: no supported way to disable optimizations
  - Oracle recommends patch 8199533 to switch optimizations off by default (10.2.0.5: not necessary)
- 11gR2: *NEW* supported init.ora parm to control optimizations
  - 11.2.0.1 : optimizations disabled by default
  - 11gR2 generates messages in alert log
- *Best Practice: explicitly set optimizations state in init.ora!*

# ALWAYS MATCH HP-UX & ORACLE SETTINGS!

## Optimize or don't optimize

| hp-ux<br>Oracle | 11i v2                                                                                                                          | 11i v3                                                                                                                                                                          |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 10gR2/11gR1     | Ensure adequate CLM & <code>_enable_NUMA_optimization=true</code> (or <code>set to false and config most memory as ILM</code> ) | Ensure adequate CLM & <code>_enable_NUMA_optimization=true</code> (or <code>set to false and config system in non-LORA mode – most memory ILM</code> )                          |
| 11gR2           | N/A<br>(Oracle 11gR2 does not support HP-UX 11i v2)                                                                             | <code>_enable_NUMA_support=true</code> ; Ensure that <code>LORA_MODE=1</code> (config enough CLM) (or <code>config system in non-LORA mode – config most memory as ILM</code> ) |



# REFERENCE MATERIALS

- <http://www.hp.com/go/integrity>
  - Select “HP Integrity server blades” link at left
- <http://support.oracle.com>
  - “Platform Notes” section for HP-UX Itanium certification page contains specific language about Tukwila support
- <http://www.hporacle.com>
  - Success stories about Oracle on Integrity server blades
  - Reference architectures for server blade deployments

Outcomes that matter.



# Backup Slides

# ON JANUARY 27<sup>TH</sup> ORACLE TOOK CONTROL OF SUN AND DELIVERED THIS MESSAGE

## Oracle + Sun: Transforming The Industry

- Complete, engineered, and integrated systems
- Innovation across the stack
- Open standards
- Industry solutions
- Highest level of customer service

**Improving the way you buy,  
run and manage business systems.**

# WITH SUN ELLISON WANTS TO BUILD ANOTHER “IBM OF THE 60’S” VISION FOR THE NEW ORACLE



## Oracle + Sun Complete, Open, Integrated Systems



- Engineered to work together
- Tested together
- Certified together
- Packaged together
- Deployed together
- Upgraded together
- Managed together
- Supported together

# ORACLE'S "GO FORWARD" STRATEGY – DO IT ALL

But When and for Whom?

**Only Oracle Delivers Complete Systems**

|                  | Oracle | IBM | Microsoft | HP | SAP |
|------------------|--------|-----|-----------|----|-----|
| Vertical Apps    | ●      |     |           |    |     |
| Horizontal Apps  | ●      |     | ●         |    | ●   |
| Middleware       | ●      | ●   | ●         |    | ◐   |
| Database         | ●      | ●   | ●         |    | ◐   |
| Operating System | ●      | ●   | ●         | ●  |     |
| Virtualization   | ●      | ◐   | ●         | ◐  |     |
| Servers          | ●      | ●   |           | ●  |     |
| Storage          | ●      | ●   |           | ●  |     |
| Management       | ●      | ●   |           | ●  |     |

# EXADATA V2 – 1<sup>ST</sup> ORACLE PRODUCT REFLECTING ORACLE/SUN HIGH VALUE APPLIANCE STRATEGY

- Based on low cost x86, servers - not SPARC
- Uses Oracle Enterprise Linux (OEL) , not Solaris
- Software costs are astronomically high – \$2,624,000 per rack
- 22 servers per rack to manage (14 storage, 8 database)

