



# What's Beyond Superdome and Marvel: HP's Future Server

**Terry C. Shannon**

Consultant and Publisher

Shannon Knows High Performance Computing

# The Fine Print

- **The following text reflects legal requirements**
  - I am not nor ever have been an HP Employee
  - I am not being compensated to deliver a marketing pitch
  - This presentation reflects my opinion, not HP's opinion
  - No NDA material is contained in this presentation
  - I strive for accuracy, but offer no guarantees
  - Trust but verify: always get a second opinion
  - Make no purchasing decisions based on this session
  - Always check with HP before planning your purchases
  - Above all, enjoy the presentation
  - Comments? Email [terry@shannonknowshpc.com](mailto:terry@shannonknowshpc.com)

# Session Agenda

- Where HP is today
  - Current Superdome, PA-RISC and IPF Technology
  - Anticipated Improvements in the next several years
- What HP must do to remain competitive in ~2008
- Gotchas...
  - HP can't reinvent the wheel
  - Critical technology rules of thumb
  - Design criteria
- Getting from 2004 to 2008
- SKHPC's speculation about the HP server of 2008
- And some thoughts on what to do with all that power!

# Before You Even Ask...

- I know what you're wondering
- The answer is ITANIUM
  - I've NEVER been wrong on a session given on 16 August
  - Then again, I've never given a talk on my birthday before
  - But neither Opteron nor Nocona lie beyond Superdome
- Trust me now, verification is 72 hours away
- See my second session (ID 3641) on 19 August
  - **“Why IPF... and Why HP”**
  - 0800 to 0915: Early, but worth attending

# Future Server, Current Progeny

- Today's BCS enterprise systems portend the future
  - Alpha
  - Superdome
- A quick review of Alpha and glueless interconnects
- Superdome in more detail
- Superdome's role in the future server
- And some things you can do with all that performance!

# HP Superdome Family... Today



**16-way**

2 to 16 CPUs  
32 to 128 DIMM slots  
24 to 48 PCI slots  
1 to 4 nPartitions



**32-way**

4 to 32 CPUs  
32 to 256 DIMM slots  
24 to 48 PCI slots  
1 to 4 nPartitions



**64-way**

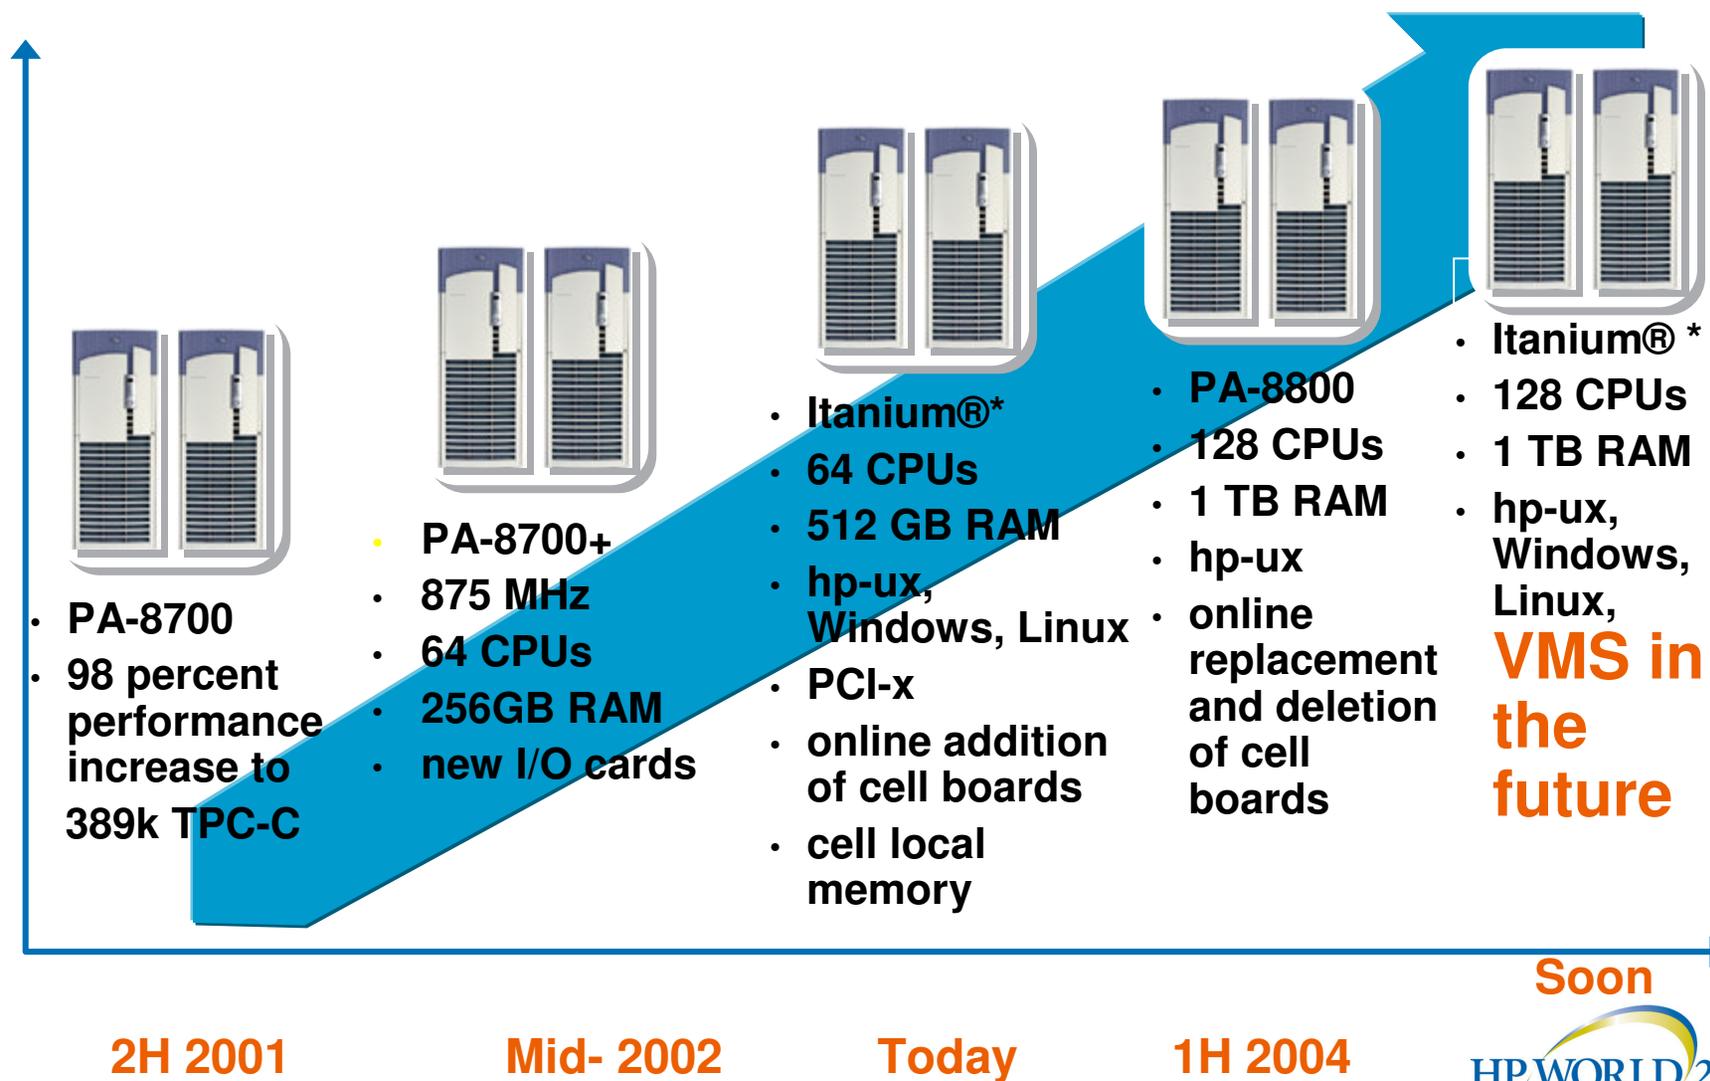
8 to 64 CPUs  
64 to 512 DIMM slots  
48 to 96 PCI slots  
1 to 8 nPartitions



**64-way with  
Expanded I/O**

8 to 64 CPUs  
64 to 512 DIMM slots  
48 to 192 PCI slots  
1 to 16 nPartitions

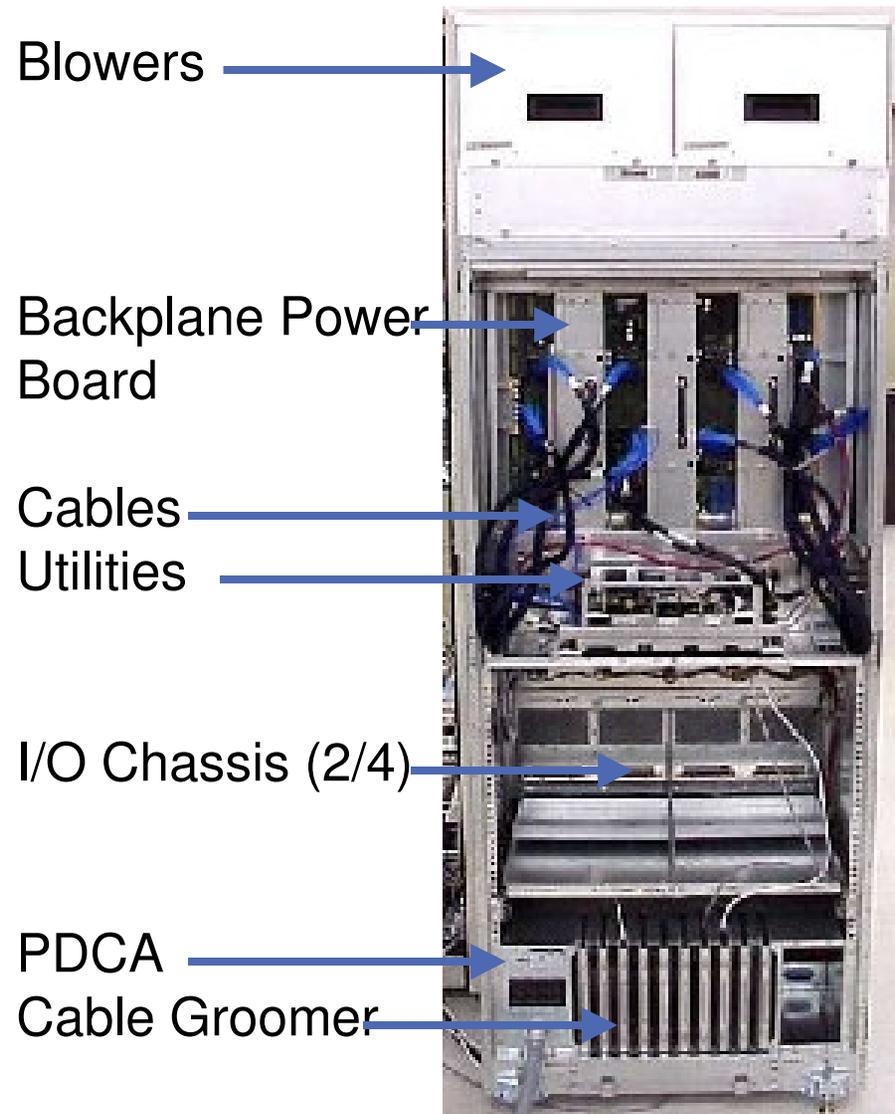
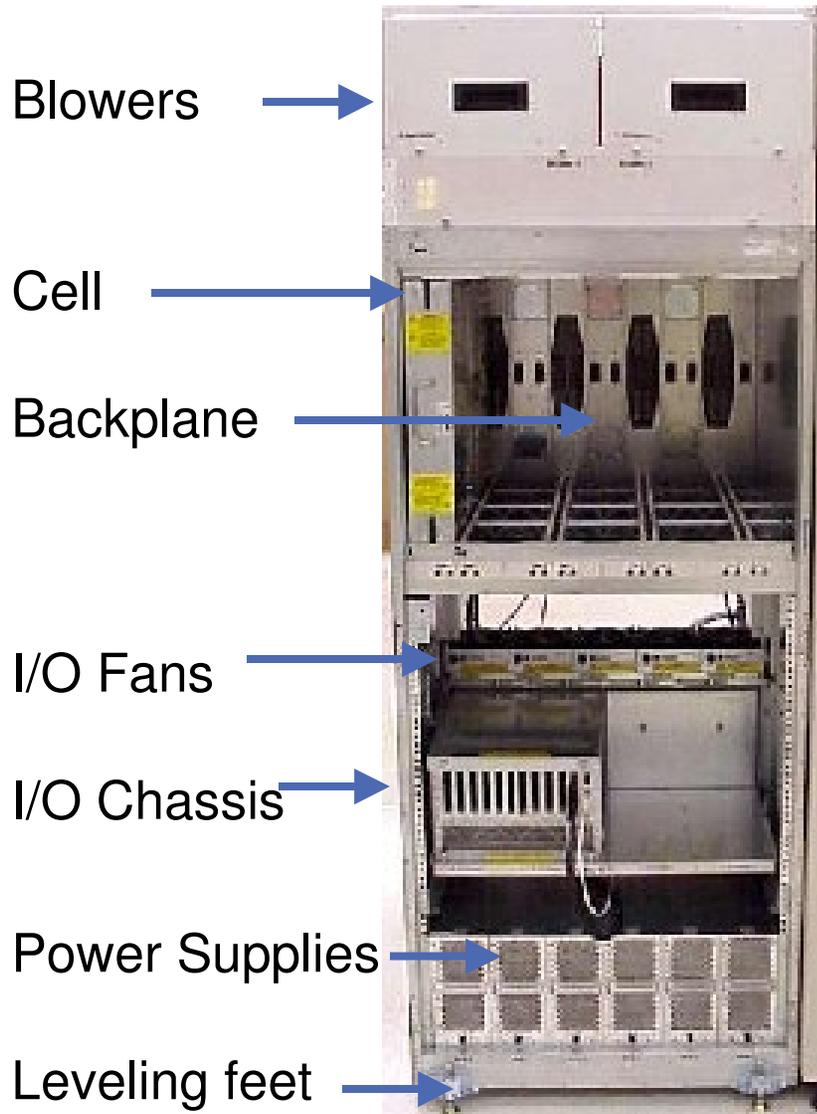
# HP Superdome: Built for the Long Haul



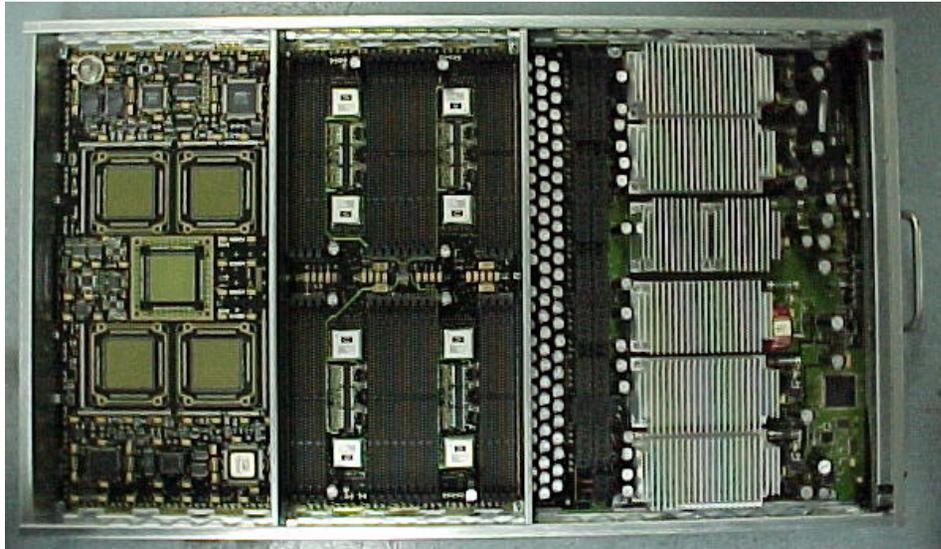
\*based on the Itanium® processor available at that point in time



# Superdome system view

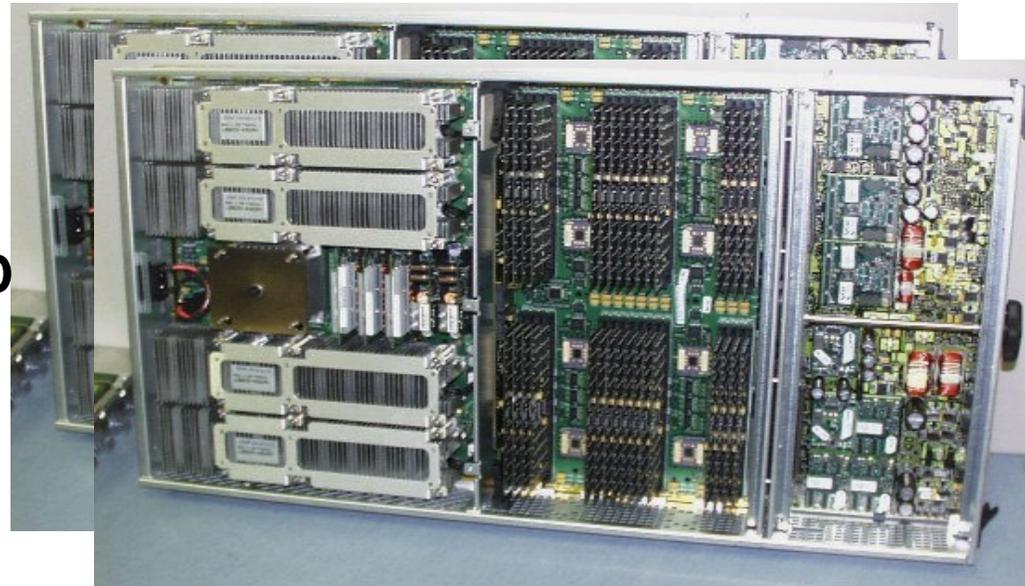


# An Inside Look at Superdome



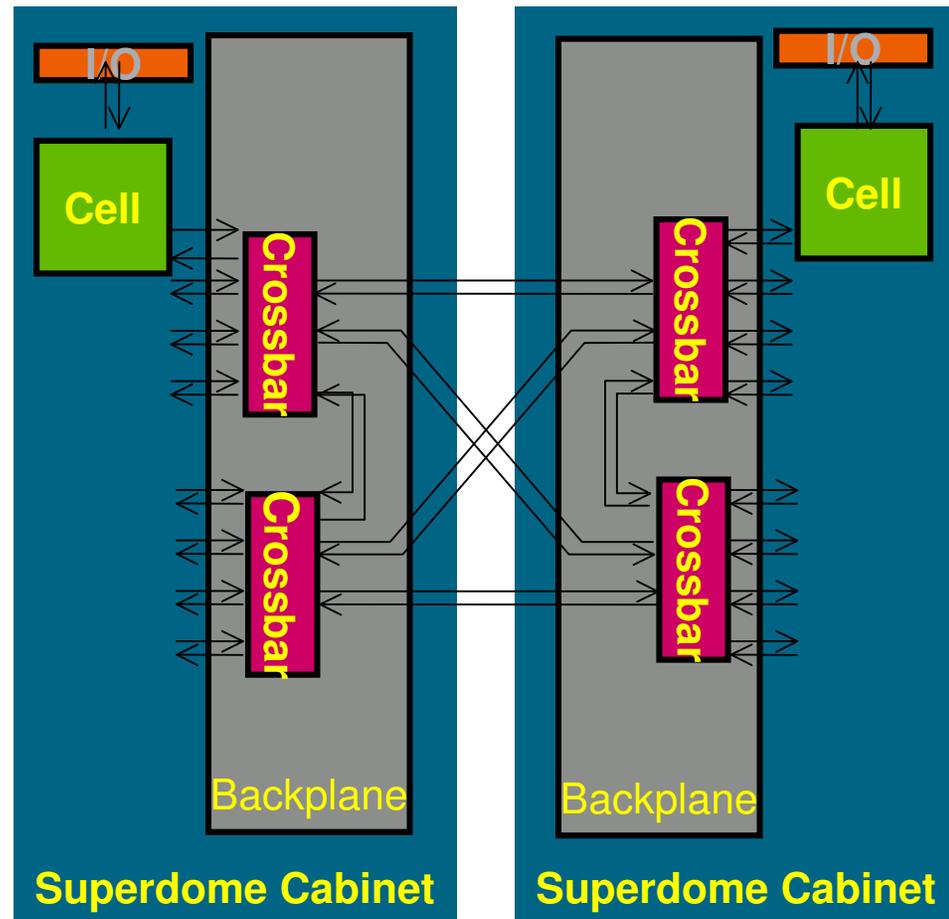
**PA-RISC CELL BOARD**

**IPF CELL BOARD**



# Superdome Interconnect Fabric: Crossbar Mesh

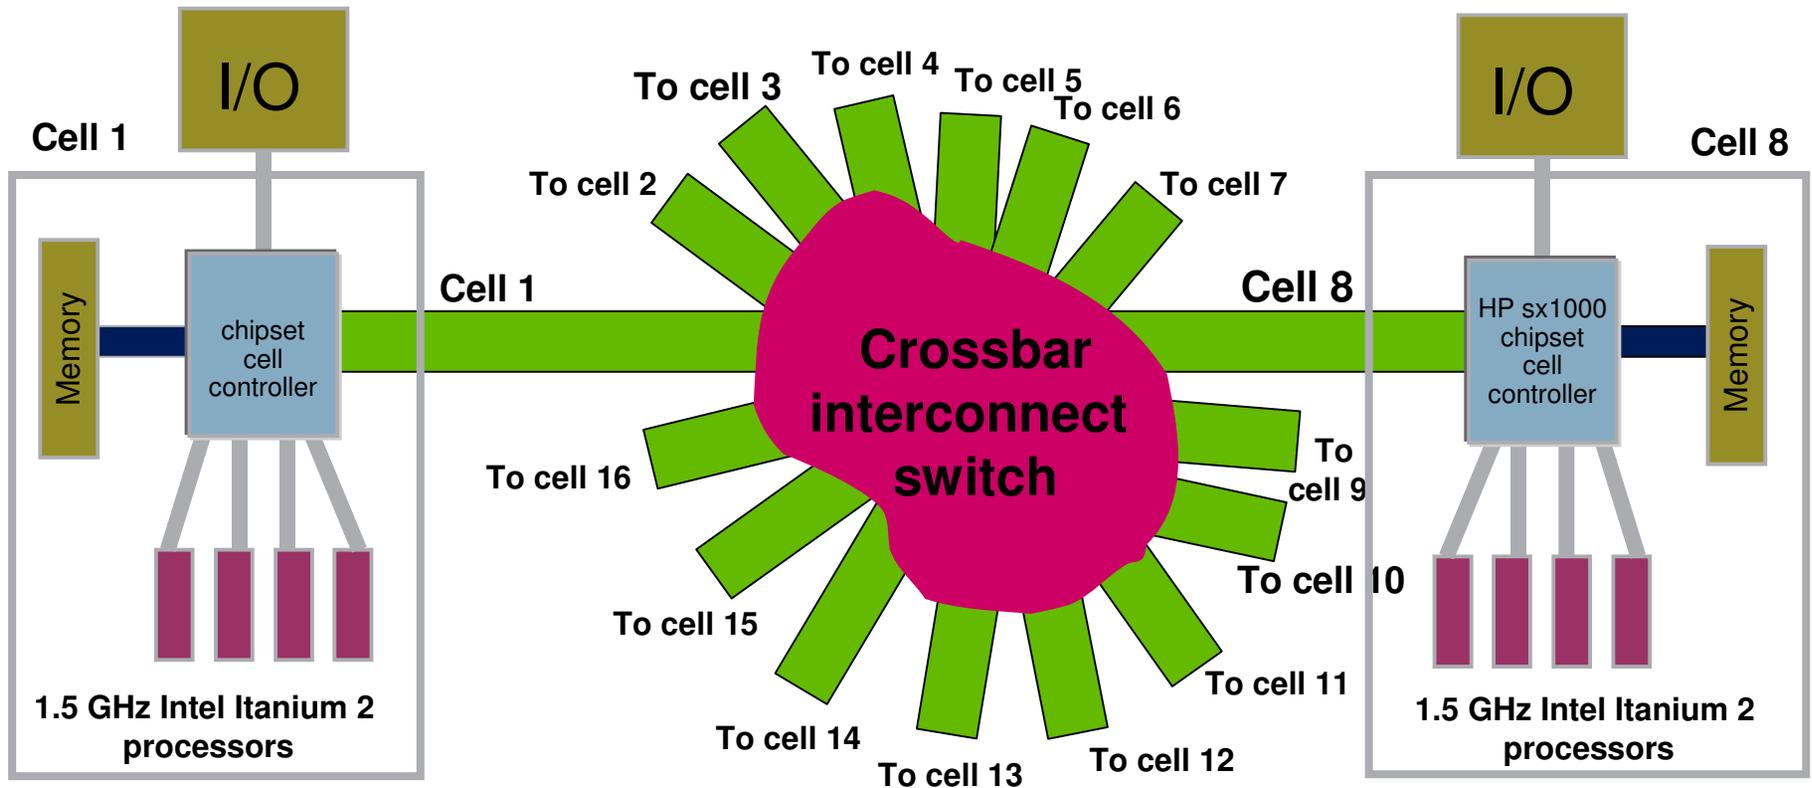
- **Fully-connected crossbar mesh**
  - four crossbars
  - four cells per crossbar
- **All links have equal bandwidth and latency**
  - minimizes latency
  - maximizes usable bandwidth
- **Implements point-to-point packet filtering and routing network**
  - allows hardware isolation of all faults
- **interconnect 16 cells with 3 latency domains**
  - cell local ~200ns
  - crossbar local ~300ns
  - remote crossbar ~350ns



# Chipsets Count in Cellular Architectures

I/O card cage—12 PCI-X slots

I/O card cage—12 PCI-X slots



# Inside Superdome: New sx1000 chipset

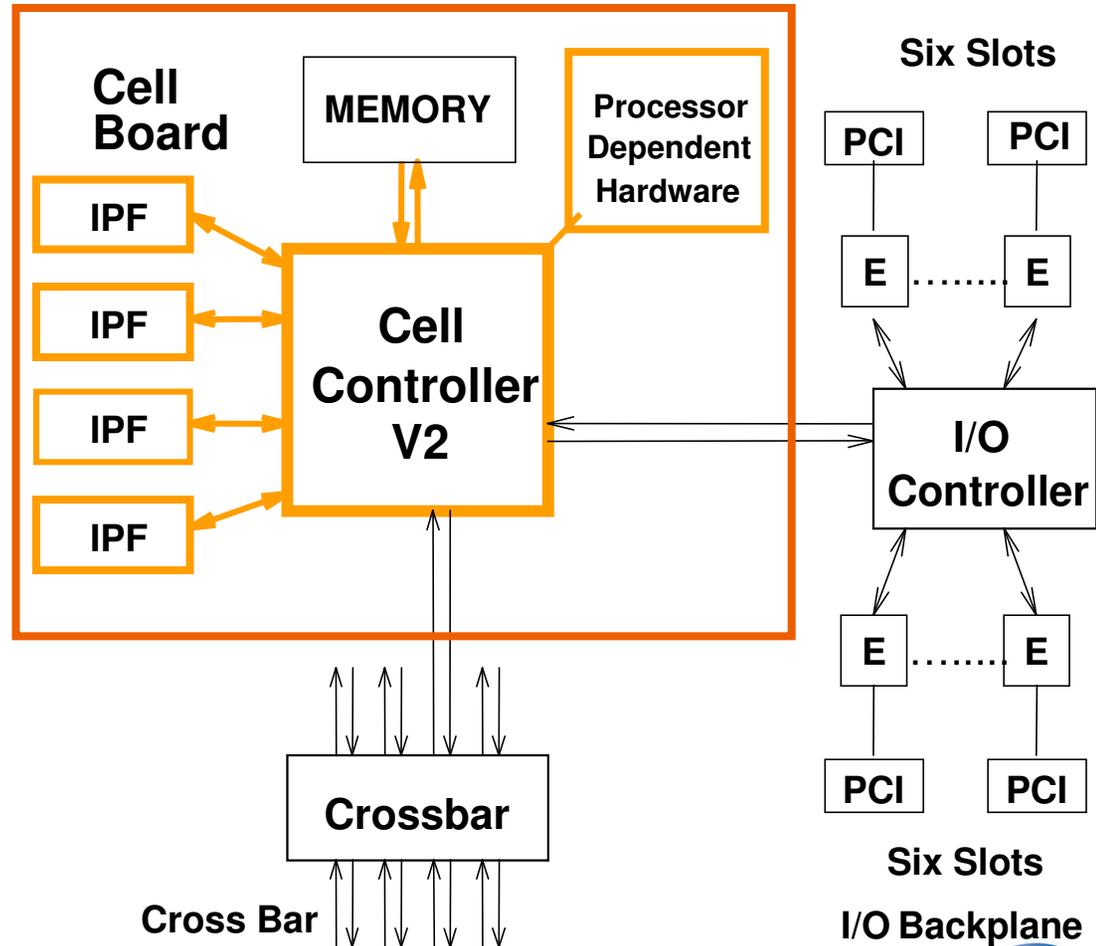
Supports PA 8800, 8900,  
and all Itanium CPUs

Cell controller, cell board,  
and CPUs change  
(orange)

Memory DIMMS, I/O  
connection, and crossbar  
connection remain the  
same

All other system  
infrastructure (frame,  
backplane, I/O chassis,  
etc.) is preserved

Optional PCI-X I/O slots  
can also be used



# The role of the sx1000 in Superdome

- The heart of all HP cellular systems, linking
  - CPUs
  - Memory
  - I/O
  - Crossbar interconnect to other cell boards
- **One chipset per cell board**
  - 4 CPUs per cell board
  - Sx1000 supports 8 CPUs per cell board
  - Enables 128-way Superdome systems

# HP Delivers Dual-Core Before Intel

- PA-RISC PA-8800
  - 2 cores, one chip
  - Each core has own L1 cache
  - 32M shared L2 cache
  - Uses high-bandwidth Itanium system bus
  - Uses same socket and HP chipsets as Itanium 2
- Itanium – HP invents a double-density Madison
  - Hondo project – now known as mx2
  - Two Madison CPUs and HP technology
  - Results – 2 Itanium processors in a single socket
- **Count sockets, not CPUs**

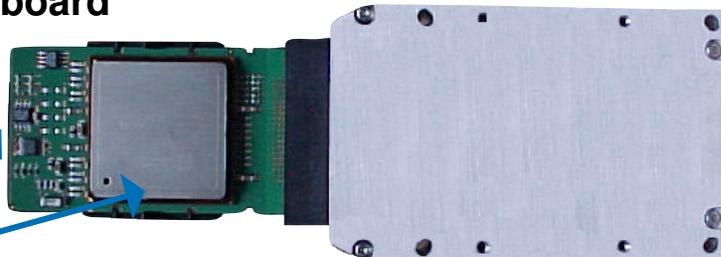
# How Did HP Upstage Intel?

- **Itanium is a joint venture between Intel and HP**
- **HP co-owns certain Itanium intellectual property**
  - HP's High Performance Systems Lab is truly brilliant
  - Richardson, TX: If it's been invented, then make it better
  - **Investigate, innovate, and improve** to achieve leadership
- How HP's 3-I approach rendered "Hondo" the mx2
  - Itanium cartridge packaging was not at maximum density
  - CPU packaged on a carrier board, power at one end
  - The Intel CPU and package could occupy less space
  - HP designed a cache, controller and carrier board
  - HP placed the power supply atop the module and...

# Field-stripping an mx2 module

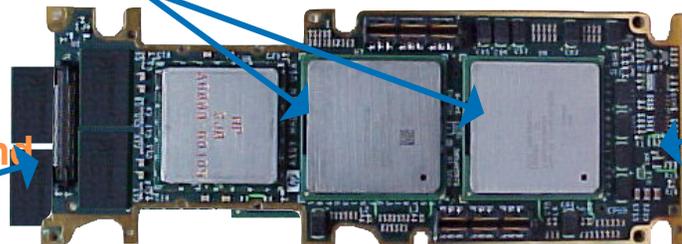
- Mixing is allowed between Madison, HP mx2 module, Madison9M

Intel CPU "carrier" board



Industry-standard Itanium2 power pod (DC to DC power conversion)

Intel CPU chip inside package



HP external cache and controller chip

HP CPU "carrier" board

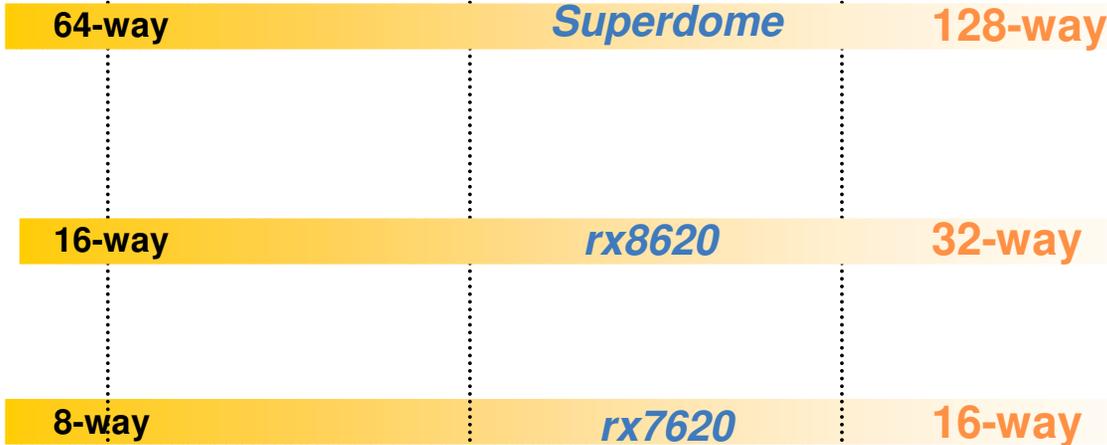


HP power solution  
(goes on top rather than  
on the end)

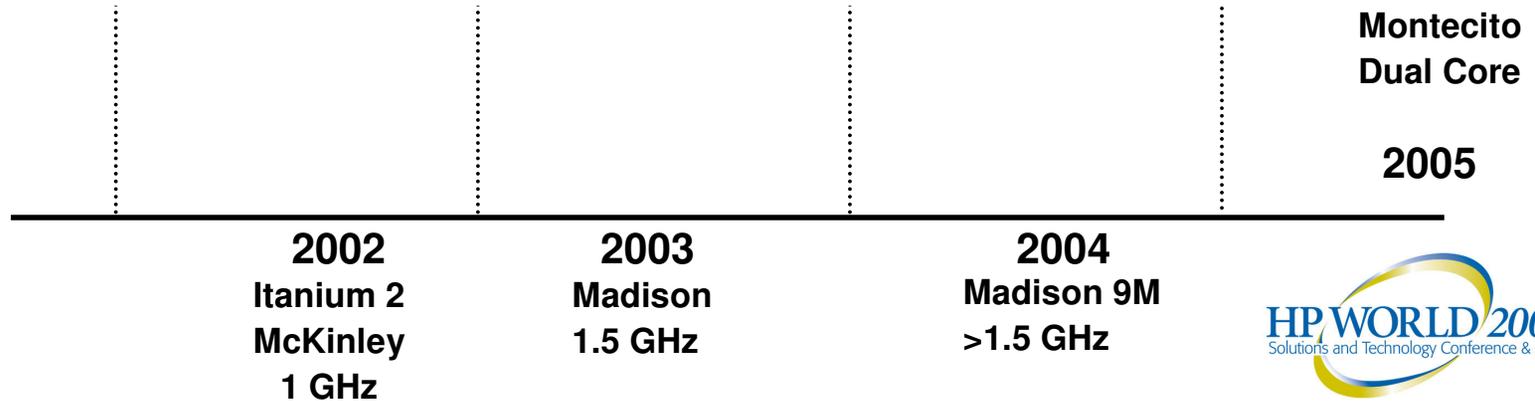
# HP Integrity Cellular System Progress



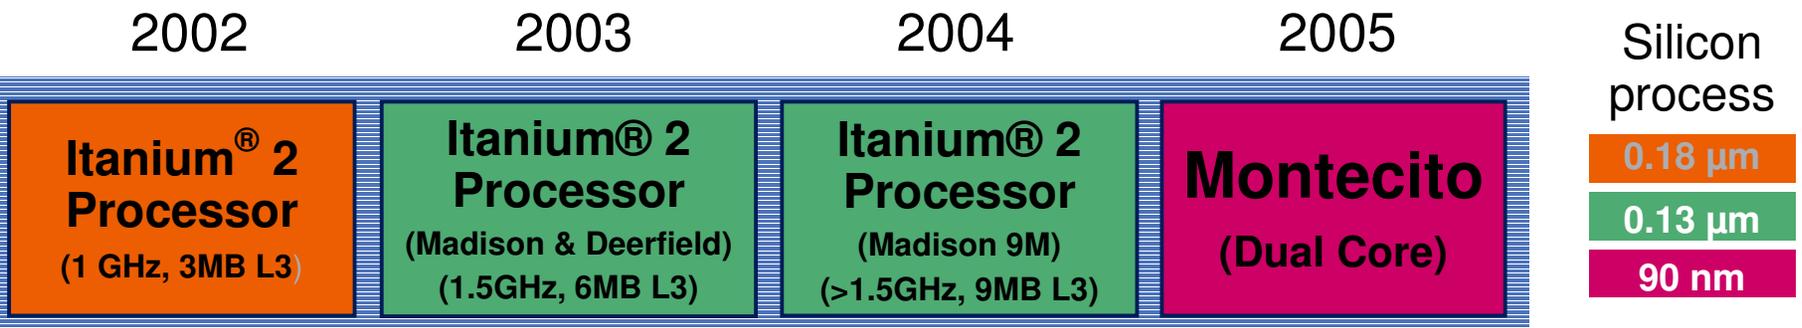
scalable (cell-based) servers



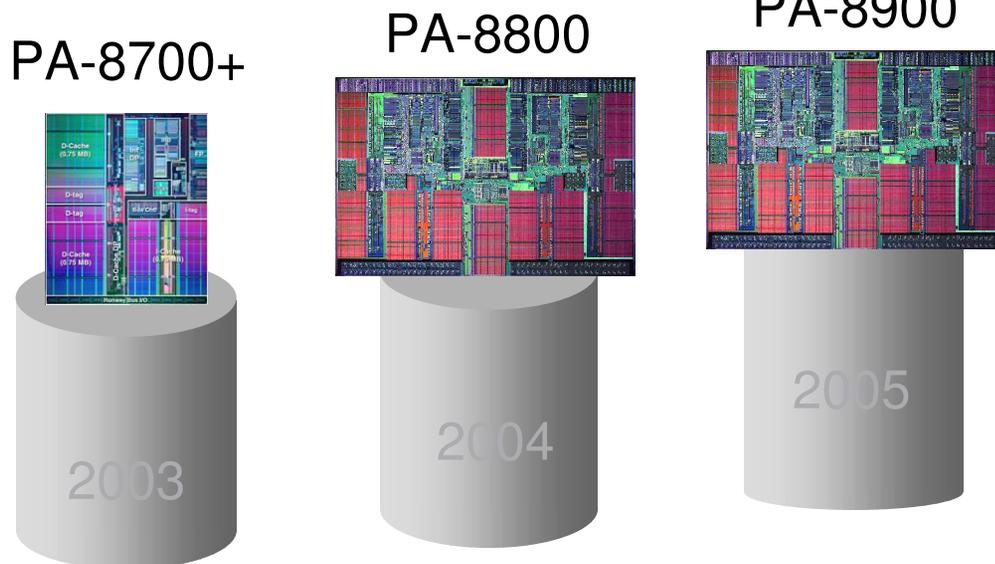
Focus is on Cellular Systems, CPU Count Reflects mx2 CPU



# Power Behind the Platforms... Near Term



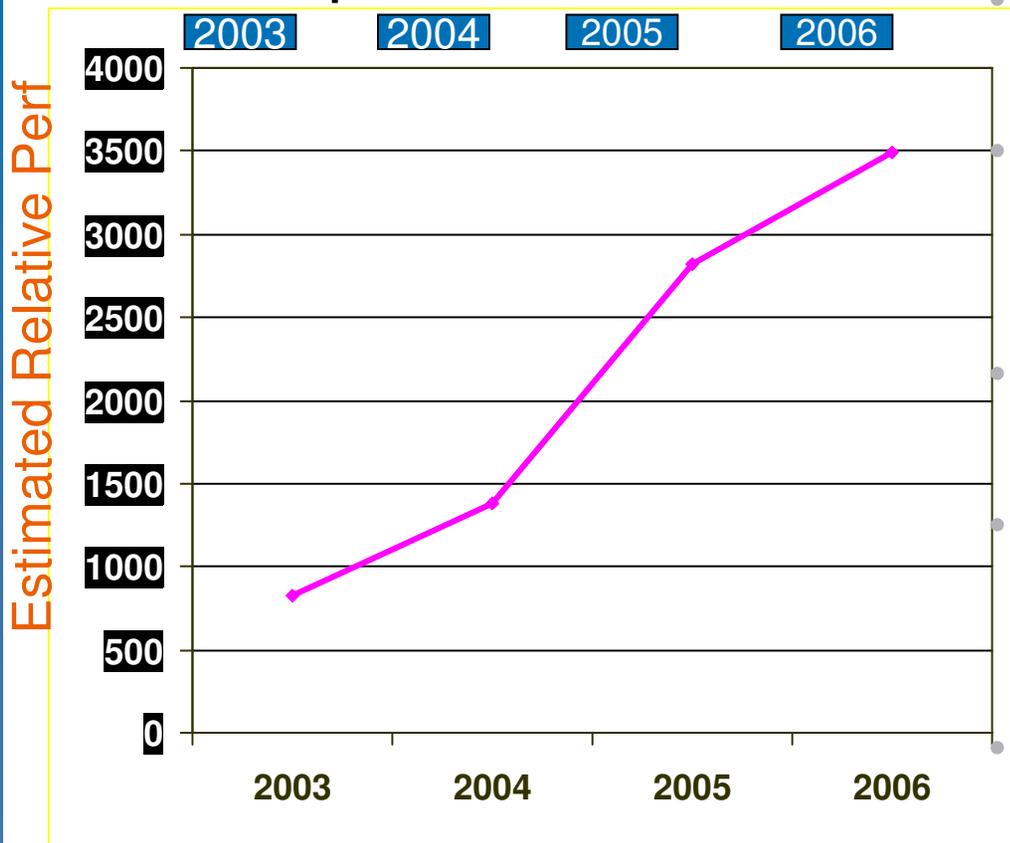
## PA-RISC family



- In 2005, the **PA-8900** will cap the long, successful line of HP PA-RISC CPUs.
- PA-8900 provide several new **enhancements over PA-8800**:
  - **Faster** clock frequency
  - **Larger** L2 cache
  - **Greater** performance

# Superdome Performance Projections

Estimated System OLTP performance



Superdome performance modeled in example.

## Integrity Enhancements:

- Processors follow Intel roadmaps
- Dual CPU motherboard doubles IPF CPU count in 2004
- Mixed Madison CPUs supported
- In-cab interconnect upgrades to sustain scalability
- In-cab upgrades from PA-RISC

## And Now It's Time for Something New...

- **Alpha, PA-RISC, and MIPS will run out of gas** in the 2006-2007 timeframe and become obsolescent also-rans.
- The **NED already has plans** for a future system
- **HP's High Performance Group** must follow suit
- The **Alpha Abdication** killed Compaq's next-generation "Avalanche" and "Snowball" high performance efforts.
- **So the ball is in HP's court** now...
- **Reliance on Itanium technology takes processor performance out of the picture.**
- Hence, HP will differentiate **above the CPU level...**

# Itanium, Integrity, and the Future

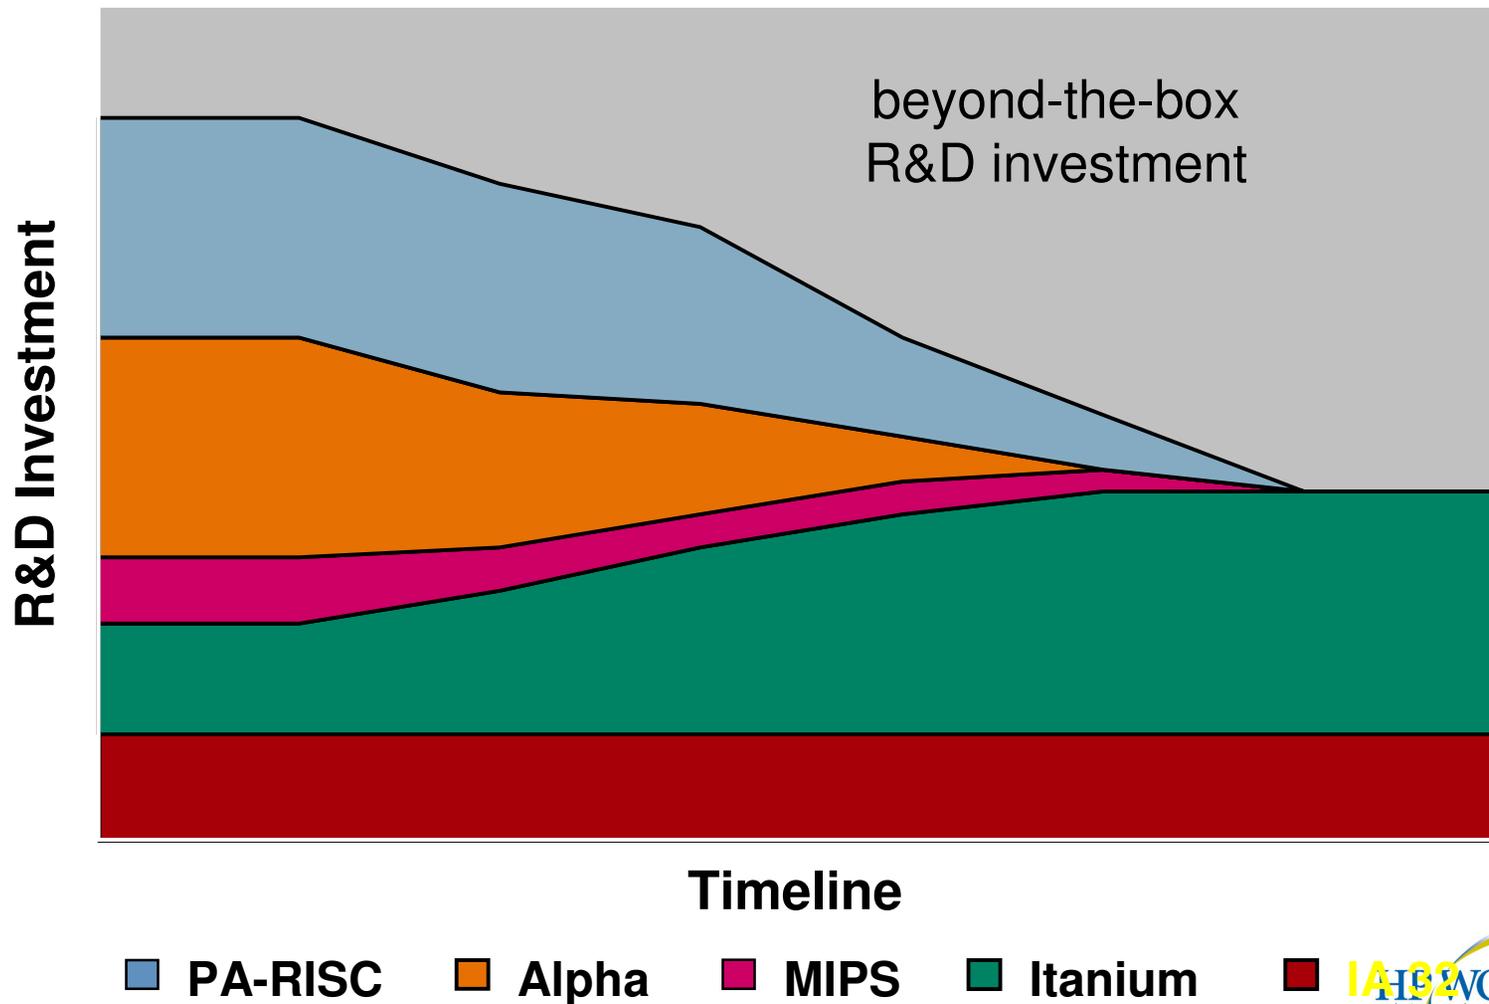
- 64-bit computing imposes new demands
  - Performance
  - Memory
  - Scalability
  - Reliability
- HP Integrity systems meet these demands
  - Today
  - And even more so in 2008

# Integrity Matters... Omit Nothing, Improve Everything!

- Keep the system running 24 x 7
  - Redundancy, N+x components
  - Hot-swappable fans, blowers,
  - Power supplies, backplanes, etc.
- Emphasize error correction
  - CPU cache, chipkill
  - Parity protection for CPUs and I/O
  - ECC for all fabric and I/O paths
  - Redundant memory and A/C power
- If it breaks, fix it fast
  - Diagnostics
  - Fault Isolation



# Focusing Innovation Beyond the CPU

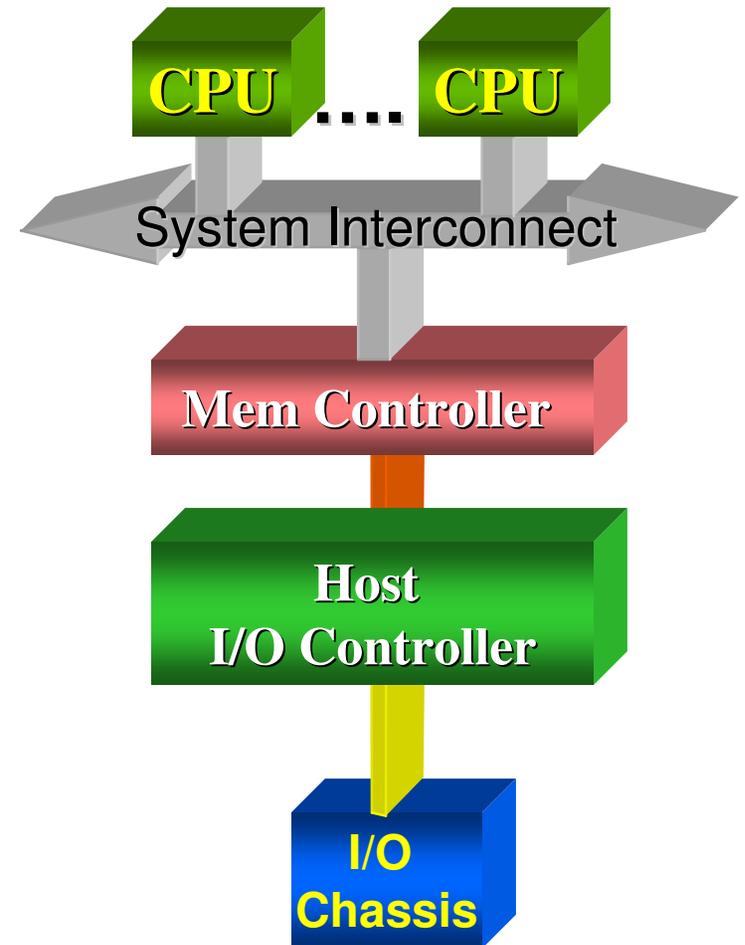


■ PA-RISC    ■ Alpha    ■ MIPS    ■ Itanium



# Future Design Considerations

- Standardize on **Itanium**
- Utilize **standard fabrics and interconnects**
- Leverage component **scalability to yield truly balanced systems**
- Incorporate latest **standard technologies**
- Ensure Customer **investment protection**
- Deliver value added **differentiation**



# Gating Factors in Future System Design

- Key customer concerns
  - Economics
  - Price and price/performance
  - Low TCO
  - Investment protection
  - Interoperability
  - Reliability
  - Scalability to handle any workload
  - Manifestation of AE strategy
  - Grid and UDC enabled

# Gating factors in future server design

- Management issues
- HP must provide an attractive alternative to competitive management products and plans
  - IBM's Eliza and e-Services offerings
  - Sun's nascent N1 strategy
- HP's alternatives include
  - OpenView and SIM
  - Virtualization technology
  - AE and UDC management tools
  - And more to come, from HP Labs and acquisitions

# Gating factors in future server design

- Heterogeneous OS is critical to server consolidation
- Essential to AE and UDC plans
- Most customers run multiple OSes and platforms
- HP must accommodate
  - UNIX
  - Linux
  - OpenVMS
  - NSK
  - Windows

# HP's future server game plan

- Note: the following information is based on SKHPC's analysis of current products, announced roadmaps, and public data. As such, this material is conjectural. For now...
- HP's Goal: Design and Develop **The Mother of all Enterprise Servers**
- *SKHPC* assessment of product development and attributes
  - **Produced by ~50-person IPF development section** as part of the High Performance Systems Lab (HPSL) in Hardware Systems Technology Division (HSTD). Distributed in SHR, MRO, Colorado, Richardson
  - Two teams are addressing specific issues

# HP's future server game plan, ctd..

- The two teams
  - **Advanced development team** focused on roadmap linkage, definition, architecture, topology, performance work, as well as base technology assessment and proofs of concept
  - **An implementation team** to deliver an SPU (system processing unit) consistent with current HSTD development practices including boards, signal integrity, mechanical, power, thermal, utilities, etc. The team will leverage existing HP ASICs, CPUs, diagnostics, I/O options, designs and components wherever possible.

# HP Superdome successor, 2007-2008

- *SKHPC's* vision
  - CPU counts to ~256 processors which allows HP to deliver SC-like systems in a box partitioned all the way down to individual processes using the continuum of partitioning technologies
  - An order of magnitude improvement in availability with proactive diagnosis tools that anticipate possible failures and automatically correct through re-configuration or early replacement of failing components
  - Incredibly large cache and memory sizes allow you to run almost any app in memory at the fastest possible speed
  - Easily clustered, interoperable, and a component of the HP UDC (Built on inexpensive industry standard components and able to run all HP OSes and the full set of industry applications)

# Aims and goals of “Son of Superdome”

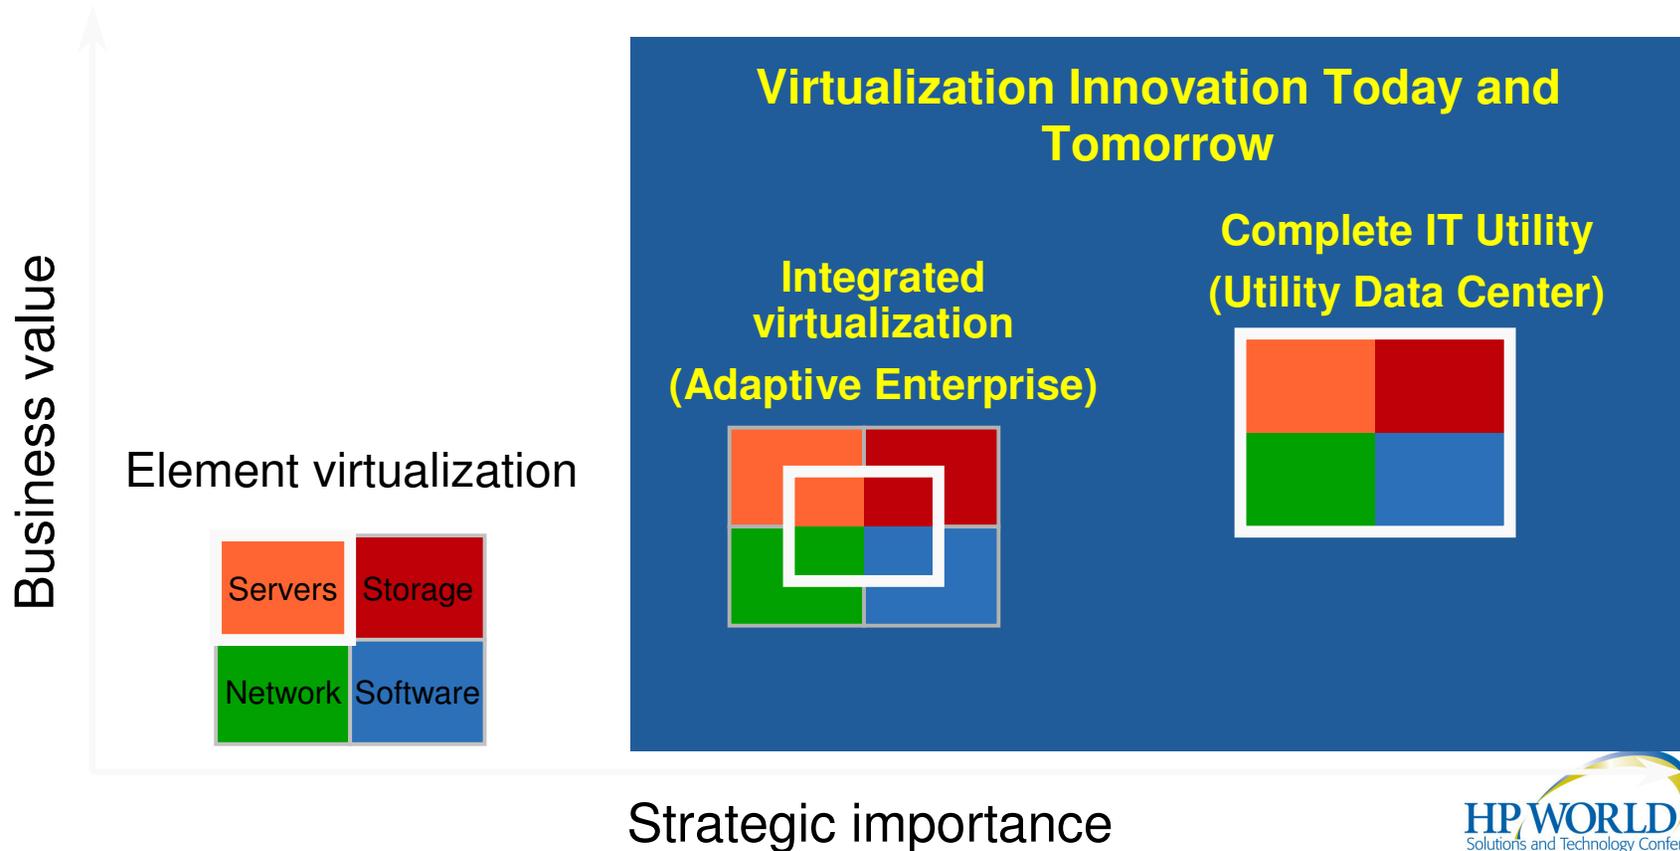
- Strategy: innovate and expand existing technology
- Results: A next-generation system whose design and appearance is **Superdome-centric with Alpha attributes**
- Likely physical characteristics:
  - Appearance: Superdome-like but much larger
  - Design: Superdome-centric with Alpha attributes
- Results and Goals:
  - Superdome successor with substantially higher performance and scalability than the incumbent
  - Deliver the first, true, open-systems alternative to proprietary mainframe and parallel cluster technology. And, a new ability to open a large can of industry-standard whoop-ass on IBM.

# Managing the server of the future

- How do you manage what does not yet exist?
  - Most of the tools and building blocks already exist
  - HP can run a UDC today
  - The Adaptive Enterprise suite is being fleshed out
  - OpenView is being extended
  - HP's partition magicians have a big bag of tricks today
- Probable management environment
  - Take the OpenView foundation available now
  - Add partitioning, COD, and AE tools available in 2007
  - The results: the next-generation system “control panel”
- Thanks, and have fun visualizing the future!

# A Quick Look at Virtualization

- An approach to IT that pools and shares resources so utilization is optimized and supply automatically meets demand

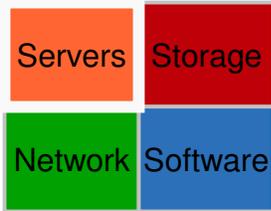


# A Few Words About AE and UDC

- HP's Adaptive Enterprise is a Key Differentiator
  - AE is an adaptable standards-based IT framework that enables users to automatically and dynamically allocate and reallocate their IT infrastructures in response to changing patterns of utilization and changing business, while eliminating the need to “overprovision” by purchasing standby equipment for “spikes”. It's a work in progress.
- HP's Utility Data Center goes beyond AE, and delivers IT resources on demand.
- HP's UDC is reality today (3 locations), and is far more comprehensive than rival “IT as a Utility” efforts from Sun (N1) and IBM's e-Services offerings.
- All offerings rely on virtualization and advanced management tools.

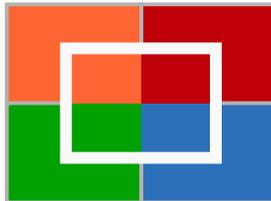
# HP Virtualization Solutions

## Element Virtualization



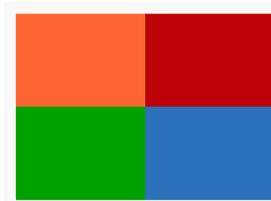
- **New: Pay-per-use (PPU)** now available for imaging and printing and HP Integrity Superdome
- **New: The family of HP Serviceguard high-availability clustering software** adds several new offerings
- **Enhanced HP-UX Workload Manager (WLM)**, the intelligent policy engine for Virtual Server Environment

## Integrated Virtualization



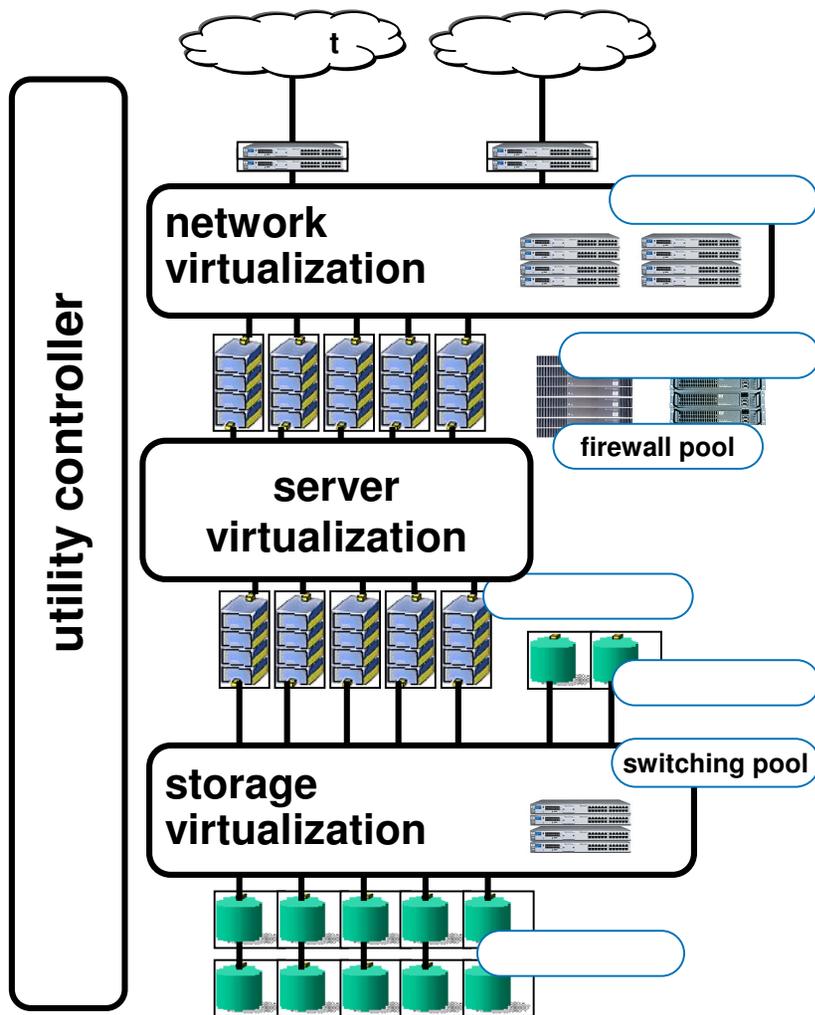
- **New: Blade PC** powers the **HP Consolidated Client Infrastructure**
- **Enhanced BEA and Oracle** integrated with HP Virtual Server Environment
- **New: ServiceGuard extension virtualizes resources across data centers** up to 100kms apart with a single Oracle RAC database
- **New: iCOD for Blades** activate servers as needed by customer

## Complete IT Utility



- **New: Tiered Messaging on Demand**, part of the **HP On Demand Solutions portfolio**
- **New: Automated Service Usage**, part of **HP Integrated Service Management**

# What is HP's Utility Data Center?



## HP UTILITY DATA CENTER

- virtualized pools of resource for instant ignition
- failover protection and data replication to
  - protect servers, storage and network
- wire-once fabric
- utility controller software for service definition and creation

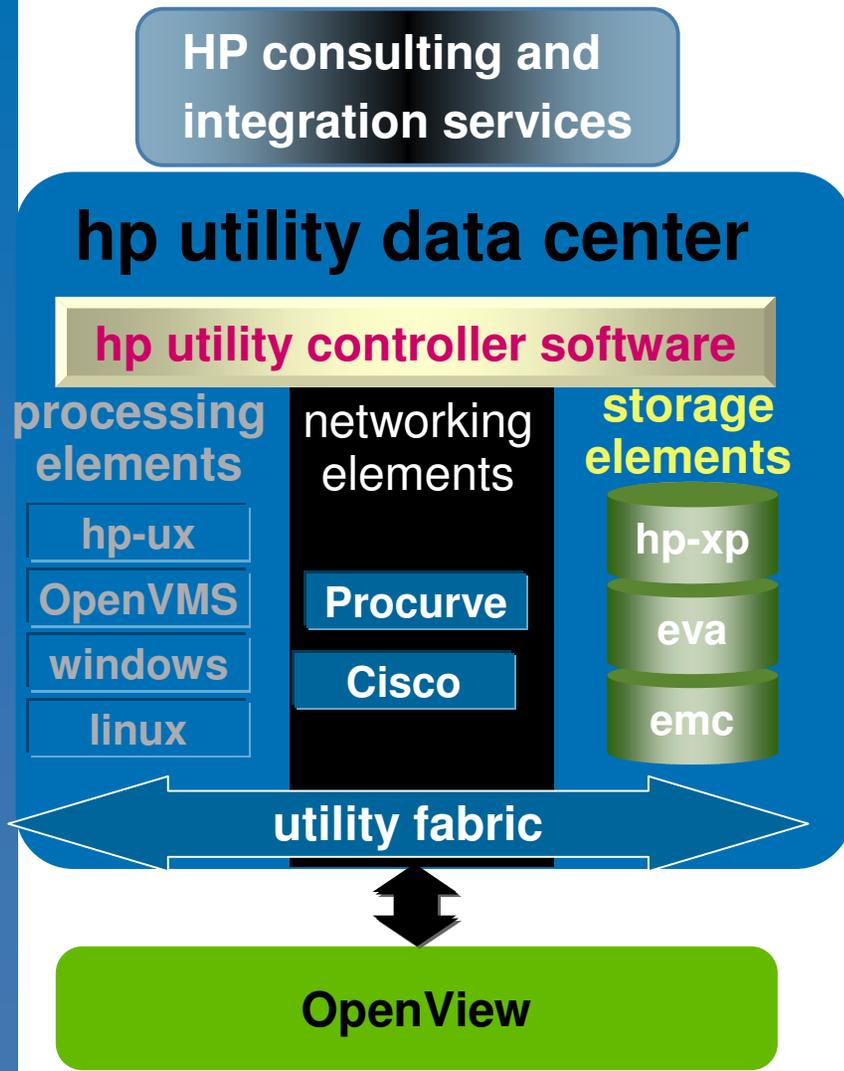
**New applications and systems can be ignited within minutes**

**Server, storage and network utilization approaches 100%**

**Resources are 'virtualized' and optimize themselves to meet your service level objectives**

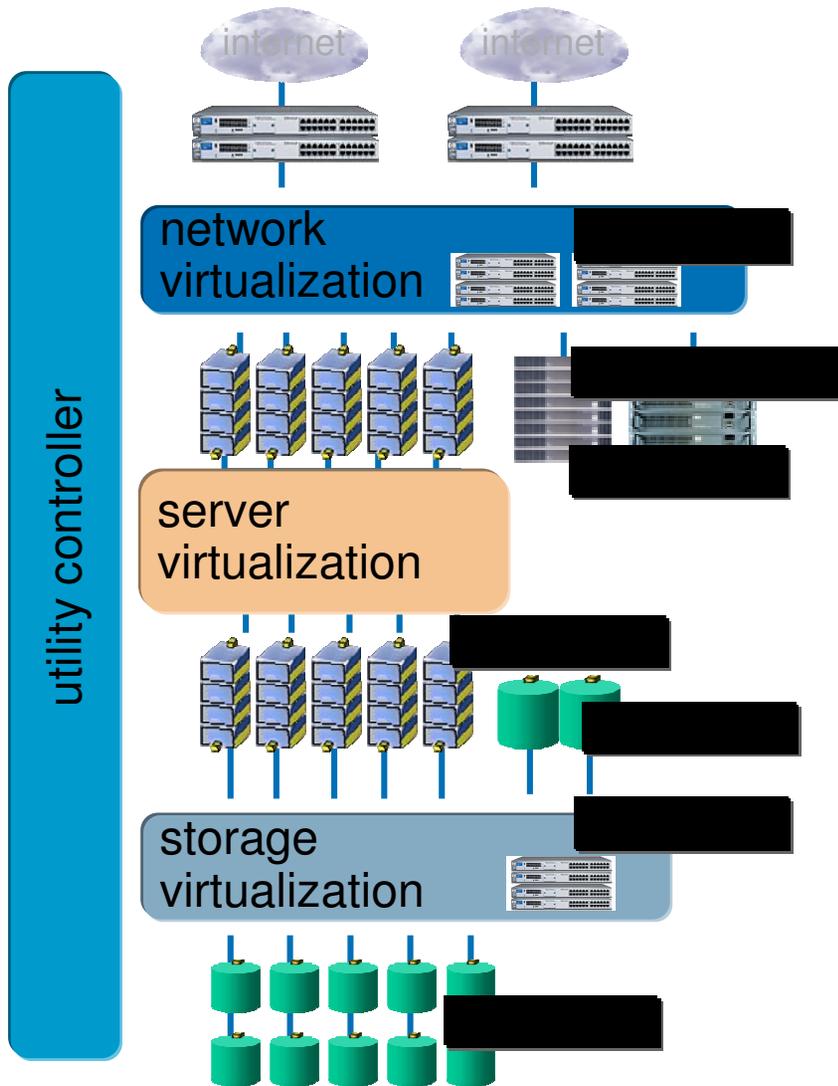
**Administrative and operational overhead is minimized**

# HP Utility Data Center Components



- **Virtual Server Pools**
  - Heterogeneous server environments
  - HP servers optimized for UDC
  - Protect your current investments
- **Virtual Network Pools**
  - Standards-based VLANs
  - Flexible and robust network infrastructure
- **Virtual Storage Pools**
  - HP XP and EVA storage offer flexible 'network-based' virtualization
  - Integration with OpenView for storage management
  - EMC Symmetrix
- **Utility Controller Software**
  - Manages service templates
  - Integrates with HP software: resource, workload and failure mgt.

# HP UDC: Improving Asset Utilization



**Adaptive management solution enabling virtual provisioning of application environments to optimize asset utility**

- **Wire it up just once**
  - network, storage, and server components wired once
- **Virtualize asset pool**
  - All components can be allocated and reallocated
- **Easily reconfigure**
  - simple user interface allows administrators to architect and activate new systems using available resources



# HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

