



# Itanium-based Servers

**Jerry Huck**  
**HP Fellow**

© 2004 Hewlett-Packard Development Company, L.P.  
The information contained herein is subject to change without notice



# Agenda

- Integrity Server Technology
- Linux MLIB capability
- System management

# Intel® Itanium® Processor Family Roadmap

## Multi-Processor (MP) Capable

<b>Itanium® 2 Processor (Madison)</b> 1.5GHz, 6M 1.4GHz, 4M 1.3GHz, 3M	<b>Itanium® 2 Processor (Madison 9M)</b> 1.7GHz, 9M	<b>Montecito</b> <i>Dual Core, 24M, Multi-threading 90nm Technology</i>	<b>Tukwila</b> <i>Multi Core, Developed with ex-Alpha team</i>
400MHz bus	400MHz, 667MHz Q1 '05	400MHz, 667MHz bus	

## Dual Processor (DP) Capable

<b>Itanium® 2 Processor</b> 1.6GHz, 3M, DP (May) 1.4GHz, 3M, DP (Apr)	<b>Itanium® 2 Processor (Fanwood)</b> 1.6GHz, 3M, DP	<b>Millington</b> <i>DP, Montecito-based</i>	<b>Dimona</b> <i>DP, Tukwila-based</i>
400MHz bus	400MHz, 533MHz bus	400MHz, 667MHz bus	

## DP Capable, Lower Power

<b>LV Itanium® 2† Processor (Deerfield)</b> 1.0GHz, 1.5M, DP	<b>LV Itanium® 2 Processor (LV Fanwood)</b> 1.3GHz, 3M, DP	<b>LV Millington</b> <i>DP, Low Voltage, Montecito-based</i>	<b>LV Dimona</b> <i>DP, Low Voltage, Tukwila-based</i>
400MHz bus	400MHz bus	400MHz bus	

1H'04

2H'04

2005

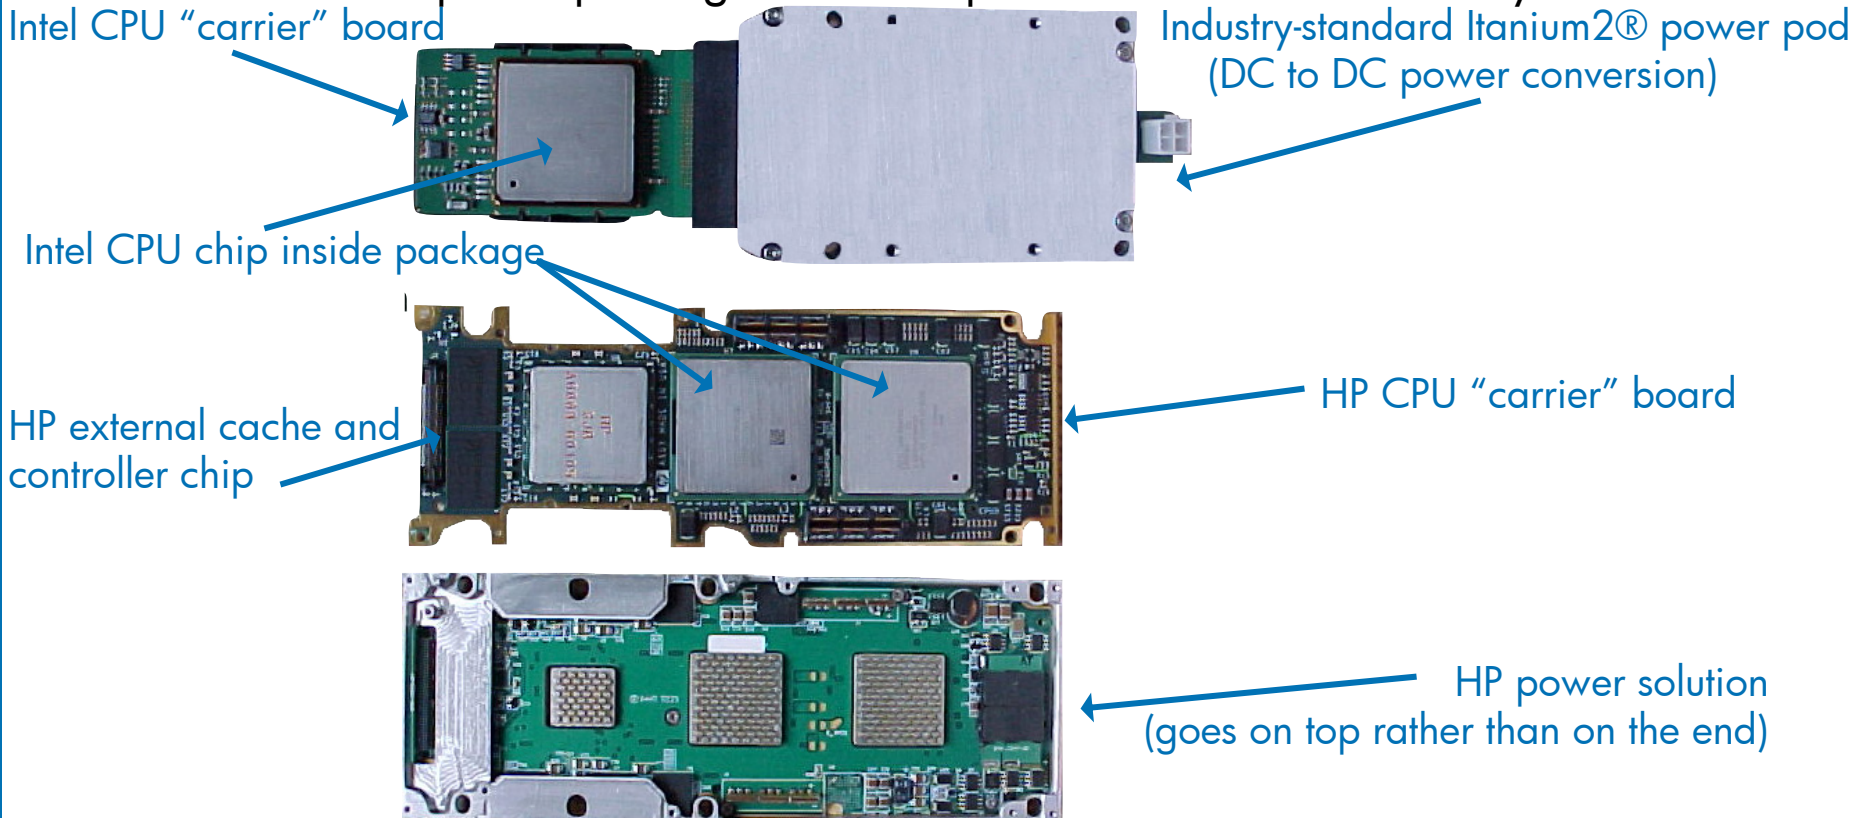
Next Generation



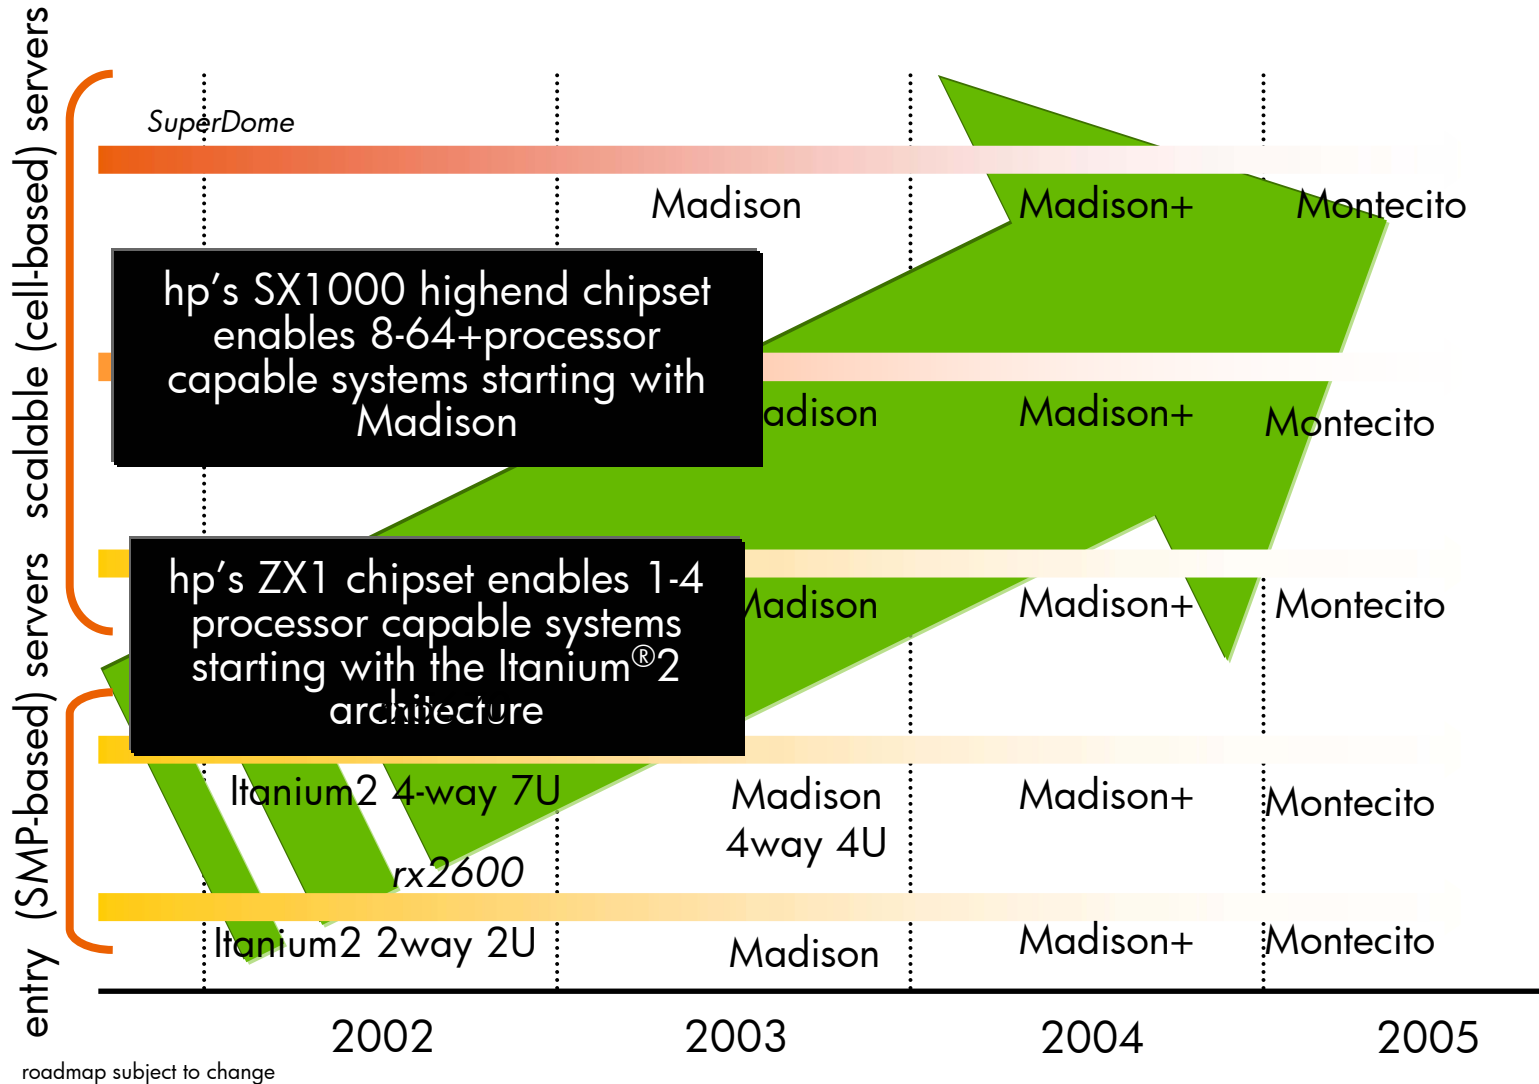
† Low Voltage Intel® Itanium® 2 processor

# “Inventing” a dual core Itanium<sup>®</sup> CPU

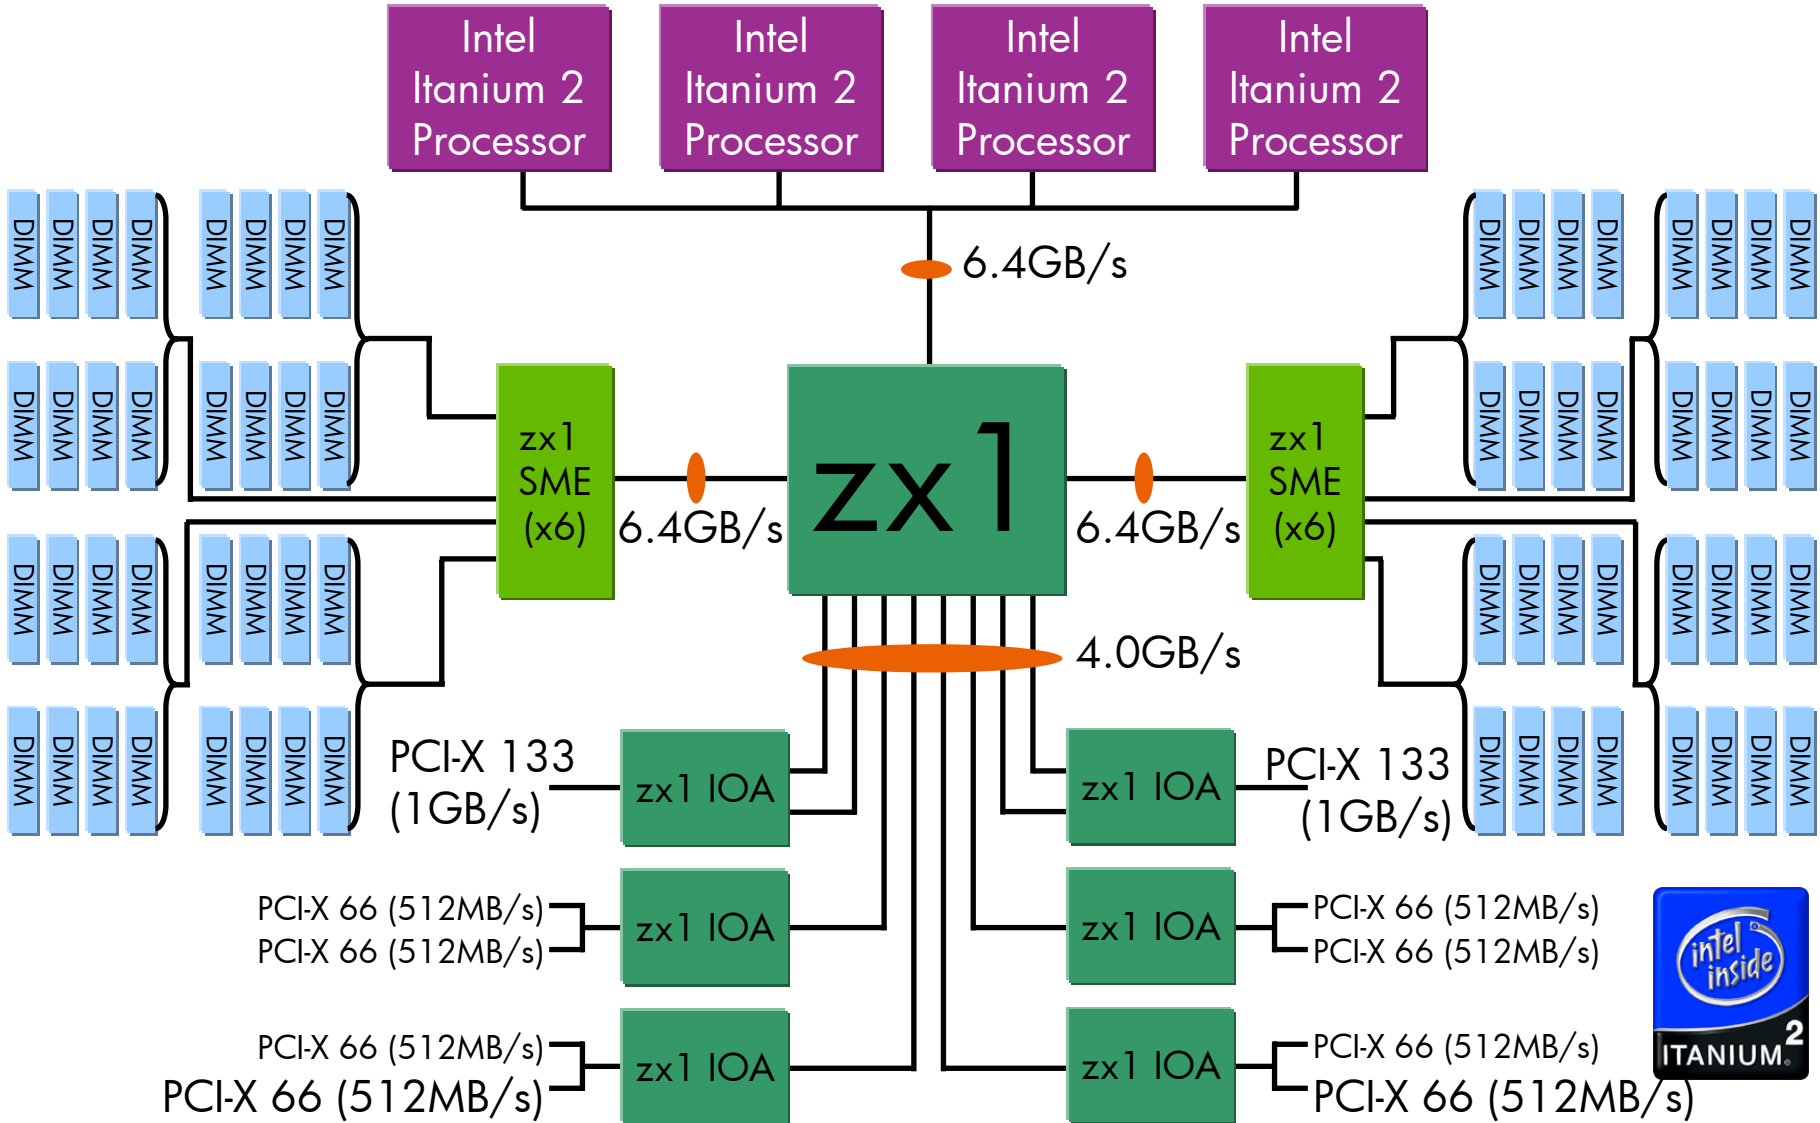
- The standard Intel cartridge packaging is not at maximum density
  - CPU silicon chip is in a package, on a carrier board, with power on the end
  - The basic chip and package could be packed much more densely



# Itanium<sup>®</sup>2-based hp servers: the most scalable roadmap from any vendor



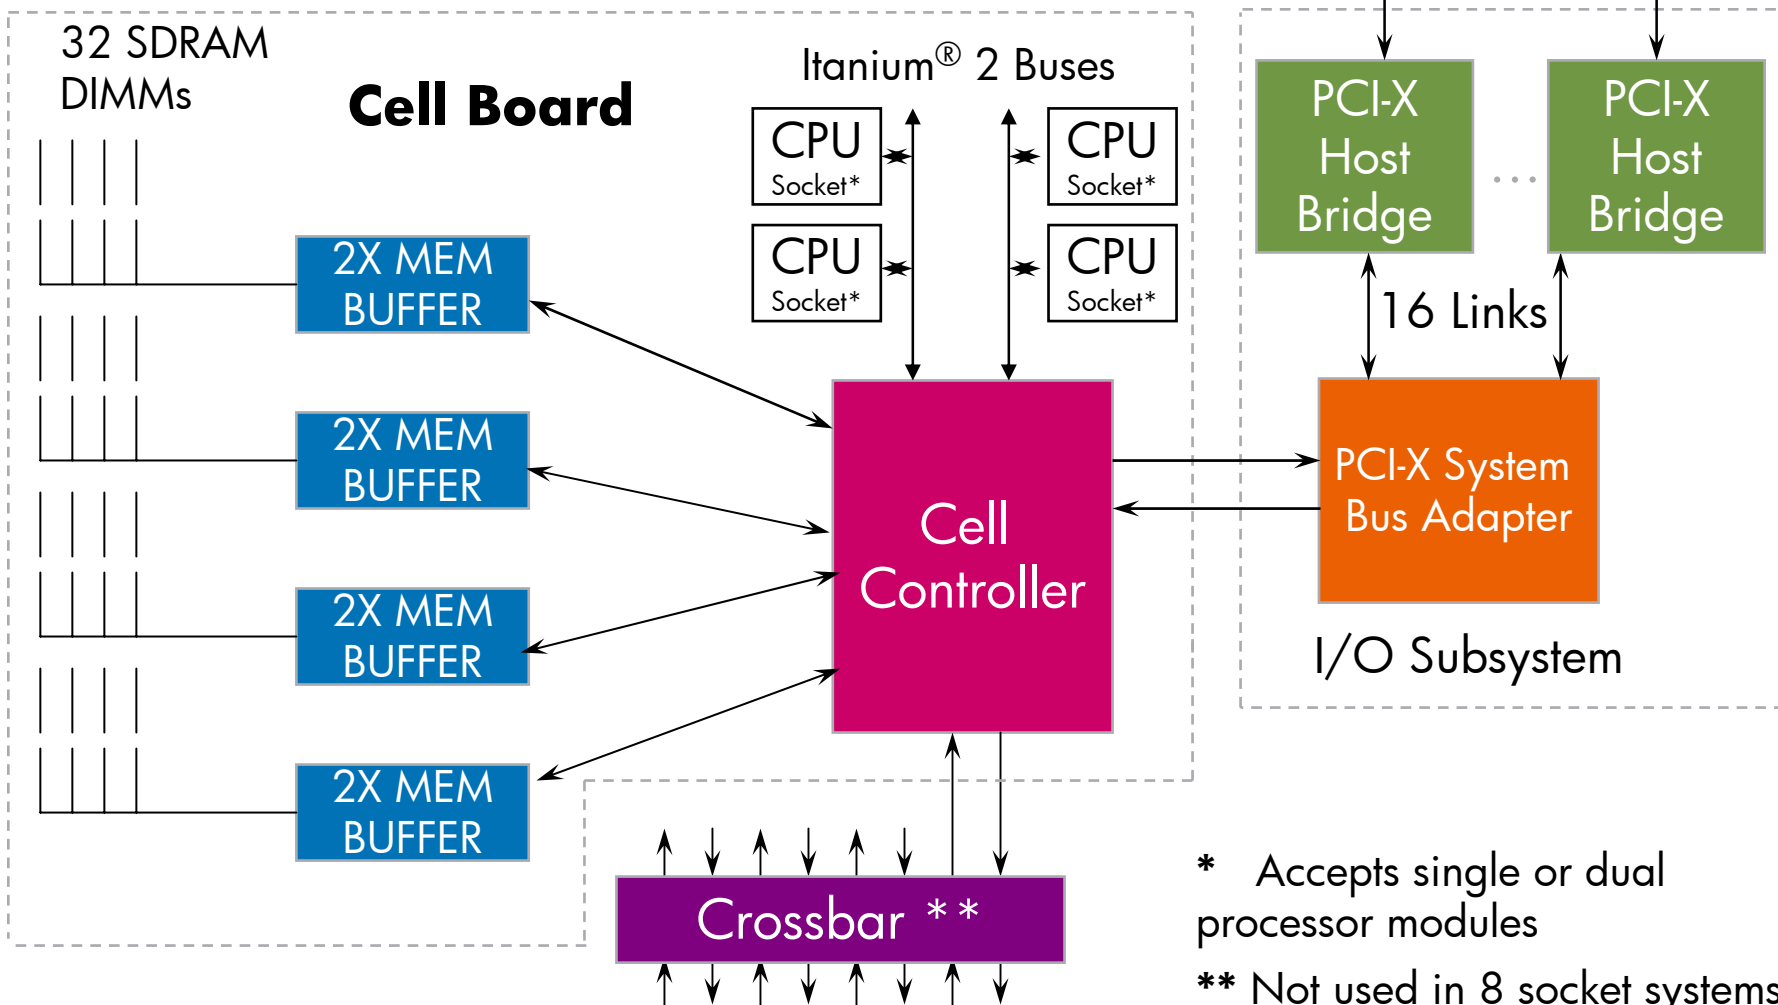
# 4-CPU Itanium 2-based platform block diagram - zx1 chipset



# sx1000 chipset includes five distinct chips



Maximizes application performance from 4 to 128 processors



\* Accepts single or dual processor modules

\*\* Not used in 8 socket systems

# HP Integrity and HP 9000 Server roadmap



Revision 3.1 Mar-04

Current offering



HP Integrity Superdome

4-128P



HP Integrity rx7620-8, rx8620-16

2-32P



HP Integrity rx4640-4

1-8P



HP Integrity rx2600-2

1-4P



HP Integrity rx1600-2

1-2P

	2004	2005	2006
HP Integrity Superdome	CPU: mx2 dual processor module OS: HP-UX11iv2, Windows	CPU: Itanium2 9M OS: HP-UX11iv2, Windows, Linux*	CPU: Itanium2 "Montecito" New Chipset OS: HP-UX 11iv2, v3* Windows, Linux, OVMS
HP Integrity rx7620-8, rx8620-16	CPU: mx2 dual processor module OS: HP-UX11iv2, Windows	CPU: Itanium2 9M OS: HP-UX11iv2, Windows, Linux*,	CPU: Itanium2 "Montecito" New Chipset OS: HP-UX 11iv2, v3* Windows, Linux, OVMS
HP Integrity rx4640-4	CPU: mx2 dual processor module OS: HP-UX11iv2, Windows, OVMS*	CPU: Itanium2 9M OS: HP-UX11iv2, Windows, Linux*, OVMS	CPU: Itanium2 "Montecito" New Chipset OS: HP-UX 11iv2, v3* Windows, Linux, OVMS
HP Integrity rx2600-2		CPU: Itanium2 9M OS: HP-UX11iv2, Windows, Linux*, OVMS	CPU: Itanium2 "Montecito" New Chipset OS: HP-UX 11iv2, v3* Windows, Linux, OVMS
HP Integrity rx1600-2			CPU: Low voltage Itanium2 (+) OS: HP-UX 11iv2, v3*, Windows, Linux, OVMS
			CPU: Next Gen Low voltage Itanium2 OS: HP-UX 11i, Windows, Linux, OpenVMS

Timeframes not to scale

Plans subject to change

All upgrades "in-box" except as noted

\*Not available at initial processor release

25-Sept-2004

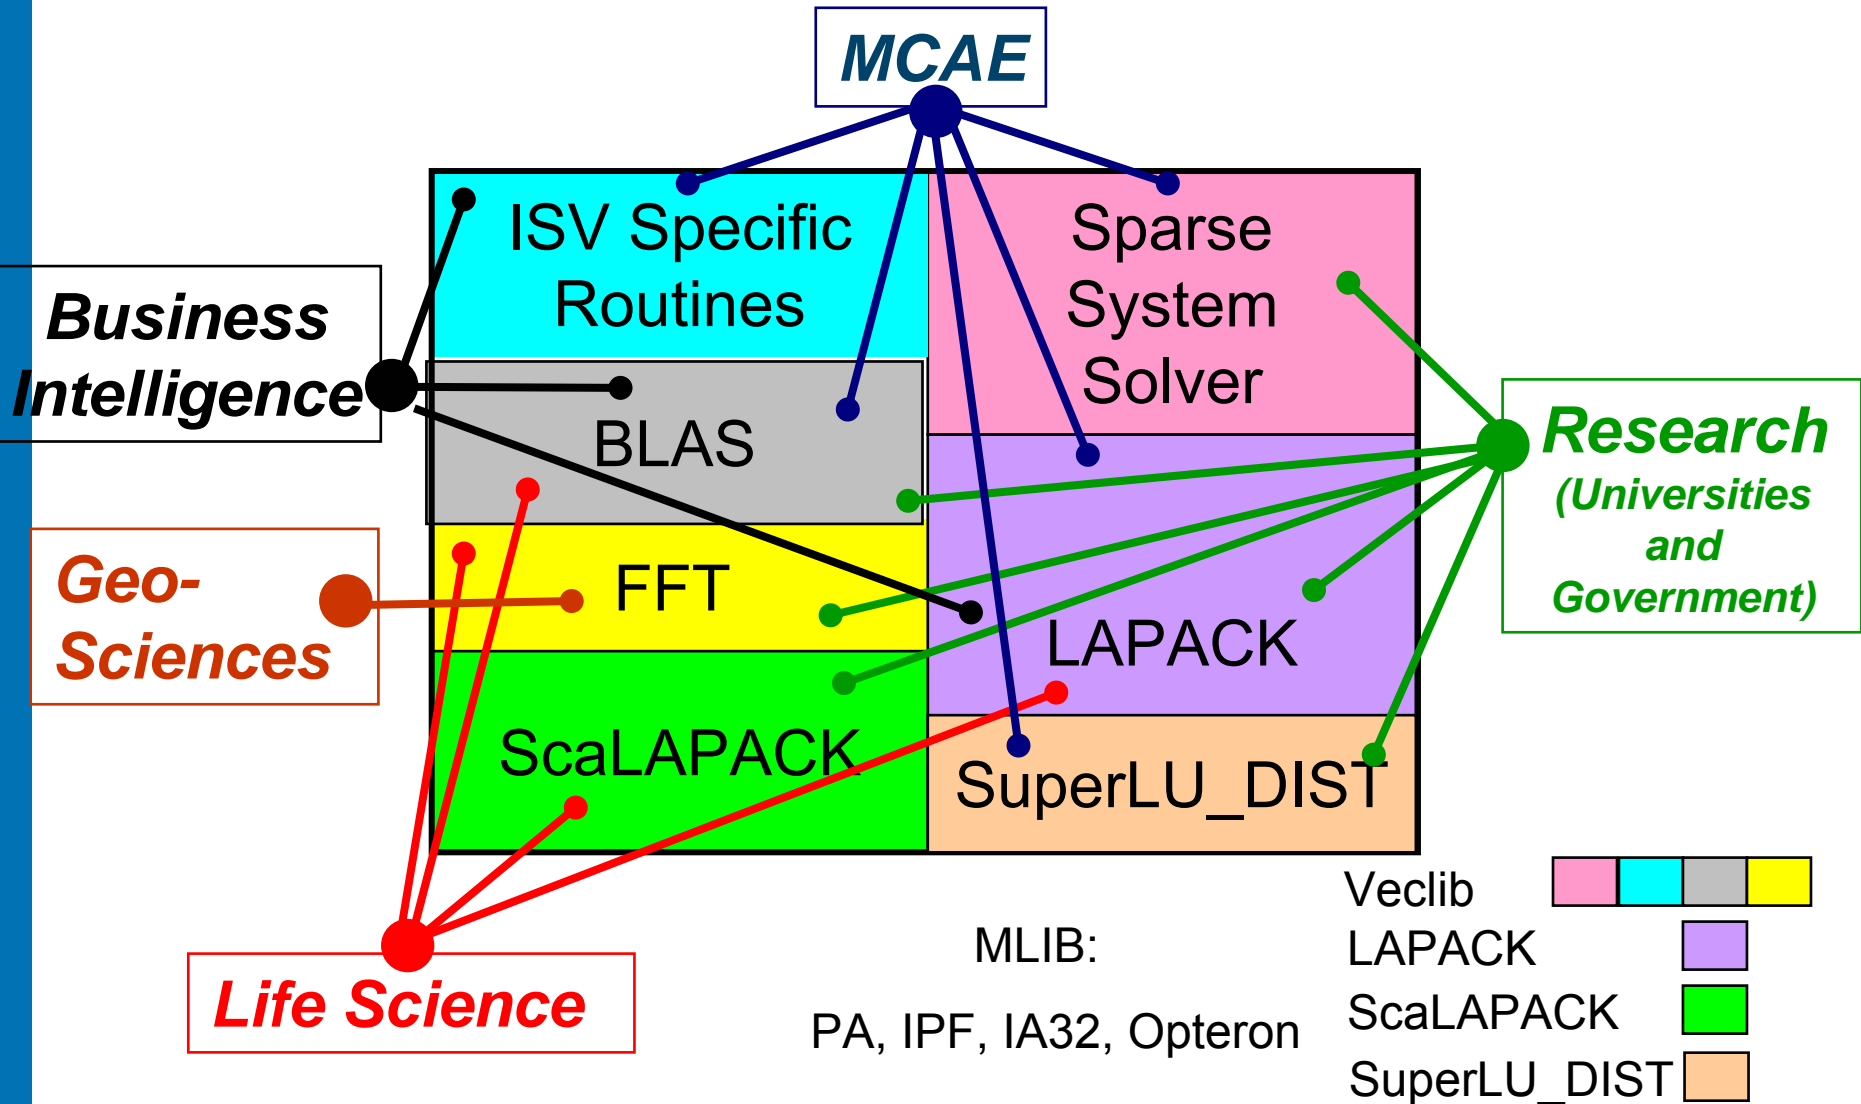
Copyright © 2004 HP corporate presentation. All rights reserved.

# HP MLIB (Math LIBrary)



- HP's Mathematical Software Product – MLIB
- Collection of subroutines providing frequently used mathematical software for scientific and technical computing
- Optimized for high performance on PA-RISC 2.0, Itanium and IA-32 processors
- HP-UX includes 32-bit addressing, 64-bit addressing, 64-bit addressing with 64-bit integers
- Itanium Linux includes 64-bit addressing, 64-bit addressing with 64-bit integers
- Includes shared and archive libraries
- Evaluation copies of MLIB and documentation available at <http://www.hp.com/go/MLIB>
- Contact: [mplibcore@rsn.hp.com](mailto:mplibcore@rsn.hp.com)

# How MLIB Products Are Used

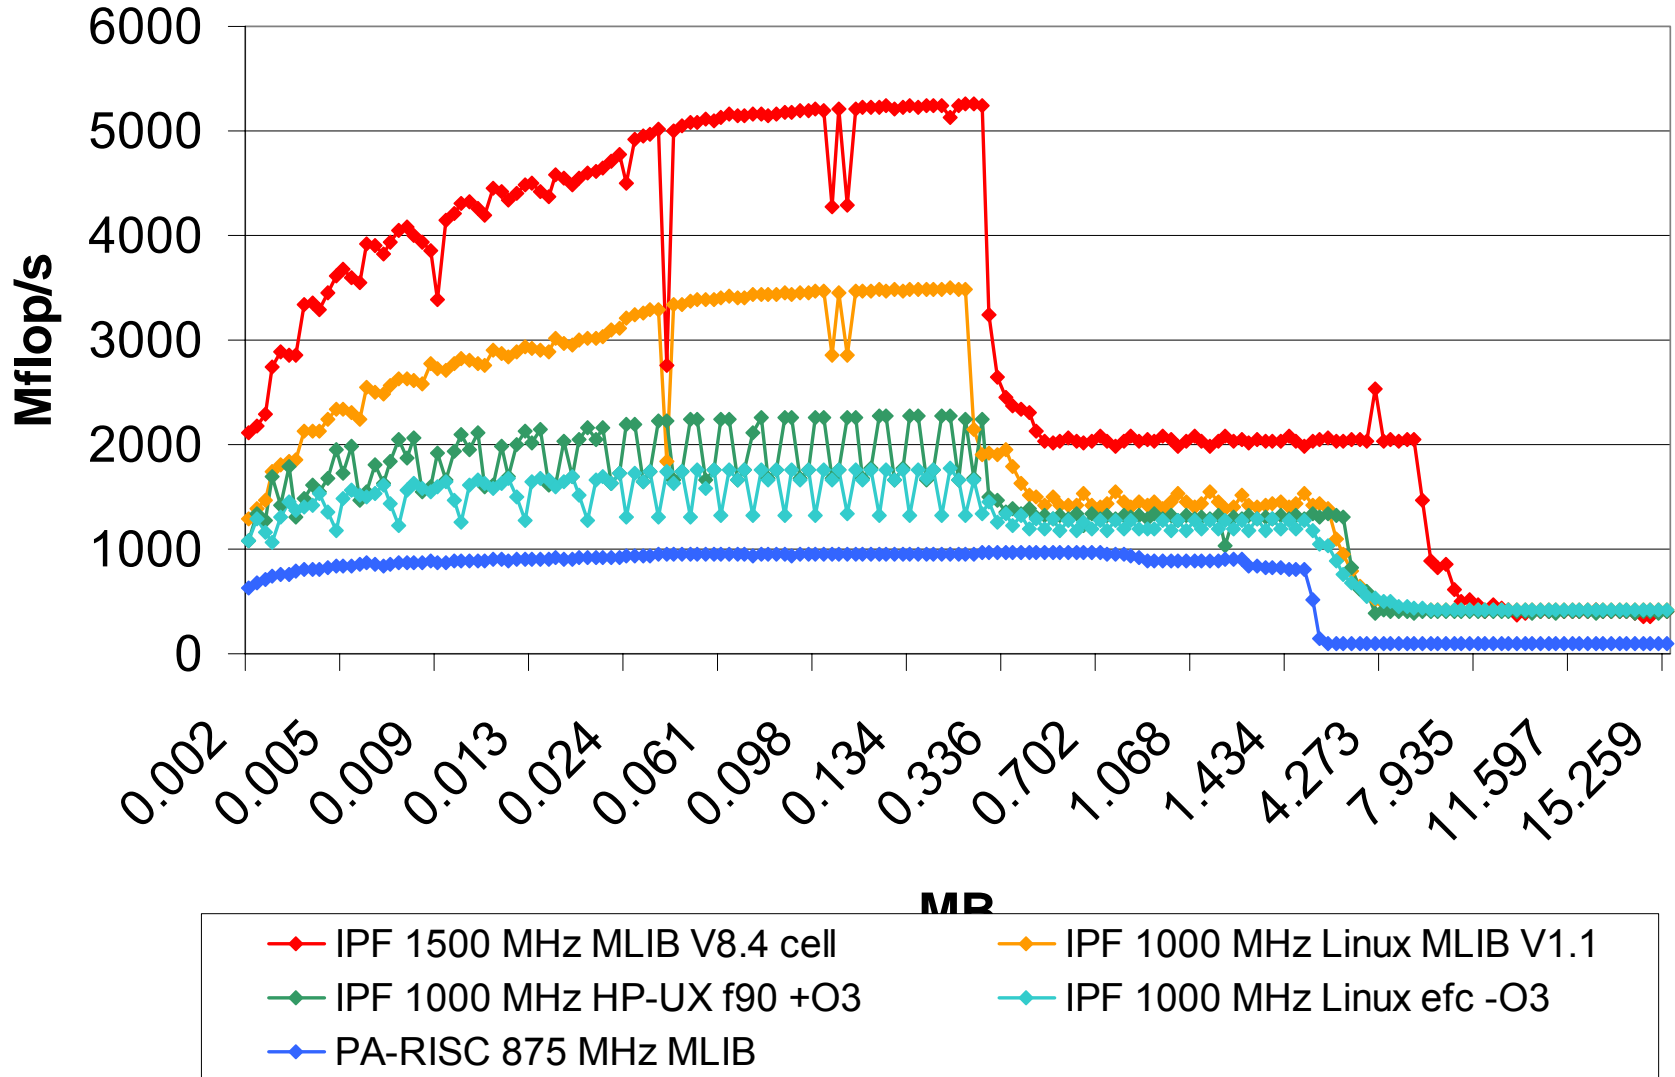


# Performance Comparison

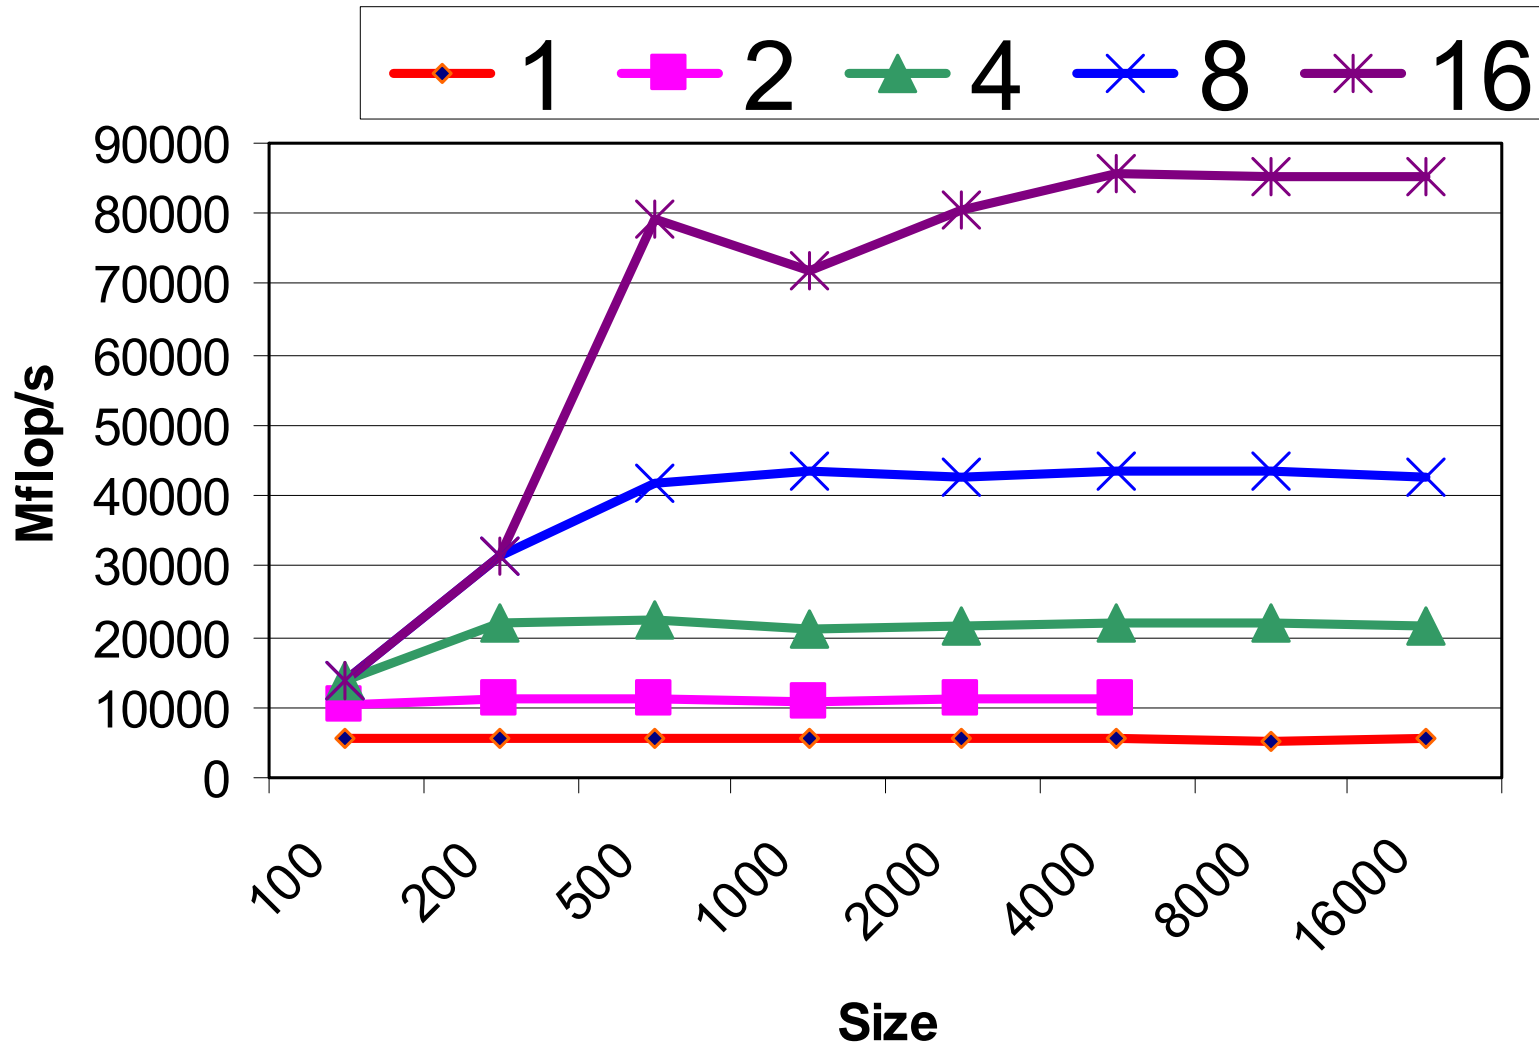
## Processors Featured

<b>Processor</b>	<b>Available</b>	<b>Clock (MHz)</b>	<b>64-bit theo. flt-pt peak (Gflop/s)</b>	<b>cache size</b>
IPF Itanium 2	now	1500	6.0	L2/L3: 256 KB/ 6 MB
IPF Itanium 2	now	1000	4.0	L2/L3: 256 KB/ 3 MB
PA-8700+	now	875	3.5	1.5 MB

# DAXPY performance



# MLIB DGEMM Performance on Integrity SD 1.5GHz - Scale Up



# Linpack NxN results on 1.5 GHz Itanium2 based platforms




- 64 nodes HP rx2600 (128 cpus) / Elan 3  
–637 Gflop/s
- 64 nodes HP rx2600 (128 cpus) / Elan 4  
–673.6 Gflop/s
- 64 nodes HP rx2600 (128 cpus) / Infiniband  
–683 Gflop/s

# MSC.NASTRAN Public Benchmark



## MSC.NASTRAN v2004 Elapsed Serial Time XXAFST (Solution 101; 2,490,516 DOF)



### Conclusions:

- Itanium2 (1.5GHz) HP-UX outperforms IBM 1.7GHz and Altix servers on this large benchmark and others
- Xeon is performance laggard
- IBM Opteron data posted but 2,705 sec is too slow to plot (?)

Reference: v2004 data posted at [www.mscsoftware.com/support/prod\\_support/nastran/performance/v04\\_sngl.cfm](http://www.mscsoftware.com/support/prod_support/nastran/performance/v04_sngl.cfm) as of 20-Feb 2004. XXAFST is Solution 101 run on Propeller Housing model. This benchmark is the largest benchmark (2.5MDOF) offered on the MSC web site.

# Architecture features used for high performance codes written using Assembly



- **Instruction level parallelism**: instruction bundling, templates and dispersal
  - This enables MLIB to have high ILP.
- **Large register files**: 128 general and 128 floating point registers
  - This allows MLIB to unroll loops a lot more
- **Rotating registers**: Acts like short vectors: with each iteration of a loop, data in rotation registers moves to the next register in the set
  - This is not your father's Vector Processor
- **Quadwords load**: read two double words using one instruction
  - This allows program to access twice amount of data per instruction. This is very useful when high cache bandwidth is needed.



# Linux in the Enterprise

# Adaptive management for HP Integrity servers with resource life-cycle management



Enterprise mgmt  
operations/  
service/business

Data center  
virtualization

Central  
Management

Performance mgmt

Workload mgmt

Deployment

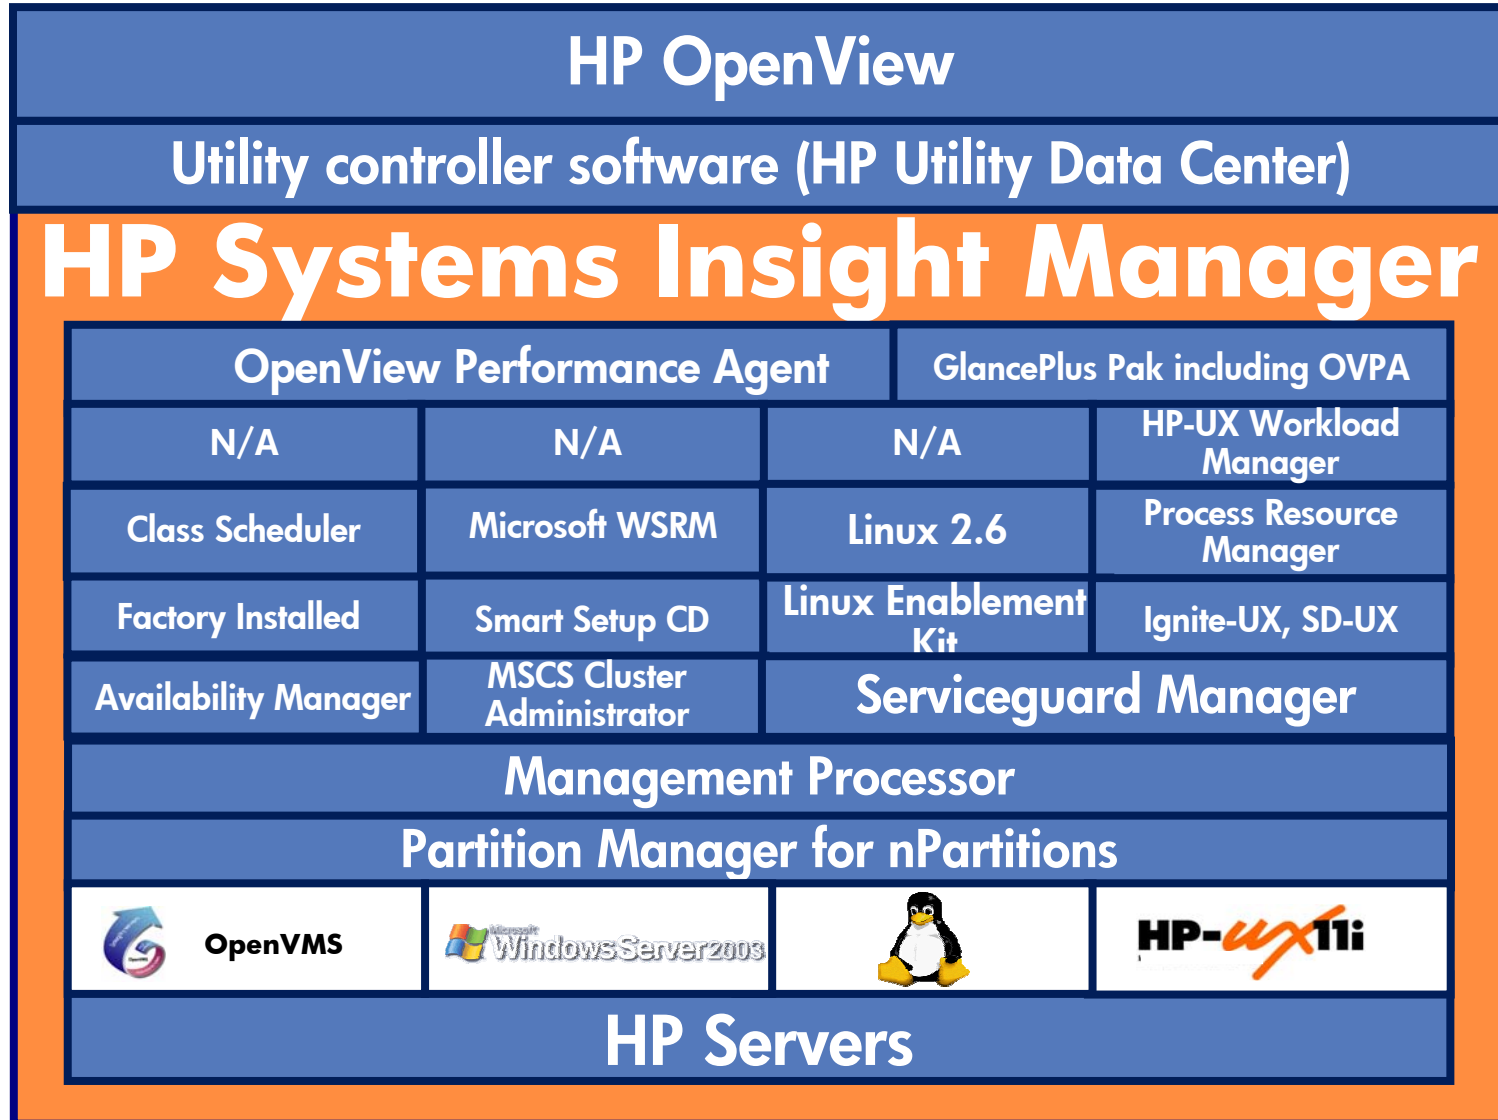
Cluster mgmt

Remote mgmt

Partition mgmt

Operating System

Resources



# Management processor embedded remote management



- Why management processor?
  - Improve productivity
  - Manage securely from anywhere
  - Embedded in the server
- What does it do?
  - Web-based or telnet access anytime, anywhere regardless of server state
  - Power cycle server
  - View event and status logs
  - View text based console
  - SSL encryption



# HP Serviceguard for Linux: Fibre-channel based configurations



## Server support

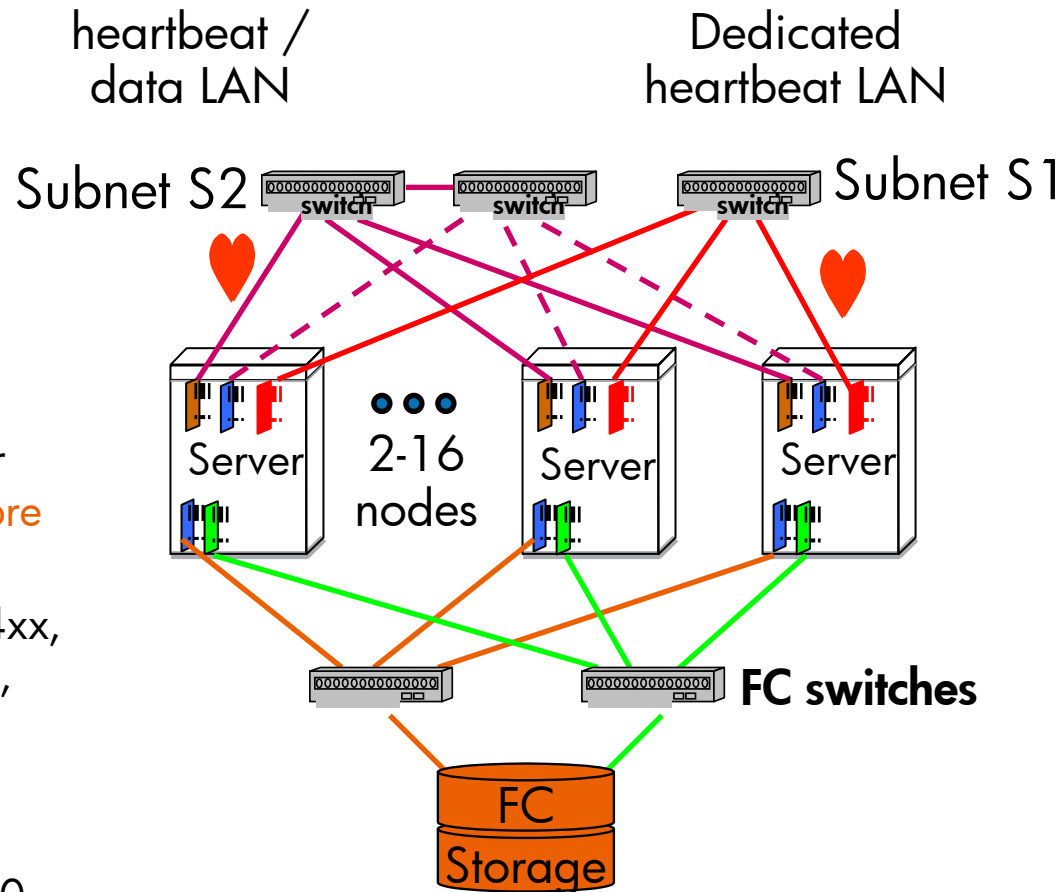
rx1600*	DL360	DL380
rx2600*	DL560	DL580
rx4640*	DL740	DL760
rx5670*	ML350	ML370
	BL20P	BL40P

## Mass storage

- HBA – HP StorageWorks fca2214 for ProLiant, **PCI-X Dual Channel 2Gb Fibre Channel HBA for Integrity A6826A**
- Storage – MSA1000, VA71xx, VA74xx, EVA 3000, EVA 5000, XP48, XP148, XP512, XP1024

## NIC

- NC3134 integrated dual port 10/100
- NC6136 1Gbit daughter Bd
- Any certified server's integrated NIC



\* Serviceguard for Linux support with:  
A.11.15.01 with SLES 8 SP3  
A.11.15.02 with RHEL AS 3 u1

# HP Integrity servers and Linux

## Delivering customer value

- Driving value-added capabilities for the enterprise
  - HP Services
  - Multi-OS, multi-platform systems management
  - Serviceguard for Linux
  - Indemnifying customers against SCO
- Itanium Linux kernel expertise
- Record-setting performance
  - SPECfp2000, TPC-C

## Growing the ecosystem

- Working with leading Linux distributors
  - Red Hat and SUSE LINUX
- Working with leading commercial ISVs
  - Oracle, Sybase, SAP, BEA + hundreds of others



- Entry-level – rx1600, rx2600, rx5670, rx4640
  - Red Hat Enterprise Linux AS 2.1 and 3
  - UnitedLinux 1.0
    - SuSE Linux Enterprise Server 8
- Mid-range servers – rx7620, rx8620
  - Red Hat Enterprise Linux AS 3
  - SUSE LINUX Enterprise Server 9 (2H04)
  - 8 way scalability
- High-end servers – Superdome
  - Red Hat Enterprise Linux AS 3
  - SUSE LINUX Enterprise Server 9 (2H04)
  - 8 way partitions



i n v e n t