HP Integrity Server

DECUS 2004
20.-22. April 2004

Jörg Demmler
Technologie Consultant
joerg.demmler@hp.com

Agenda

- Integrity Server
  - Strategische Serverplattform für das Enterprise Computing
- Prozessortechnologien
  - Die Itanium Architektur
- Chipsets und Prozessormodule
  - HP zx1 Chipset & sx1000 Chipset
  - HP mx2 Dual Prozessormodul
- Low- & Mid-range Server rx1600, rx2600, rx4640
- Mid-range Server rx7620, rx8620
- Superdome

HP Processor Roadmap

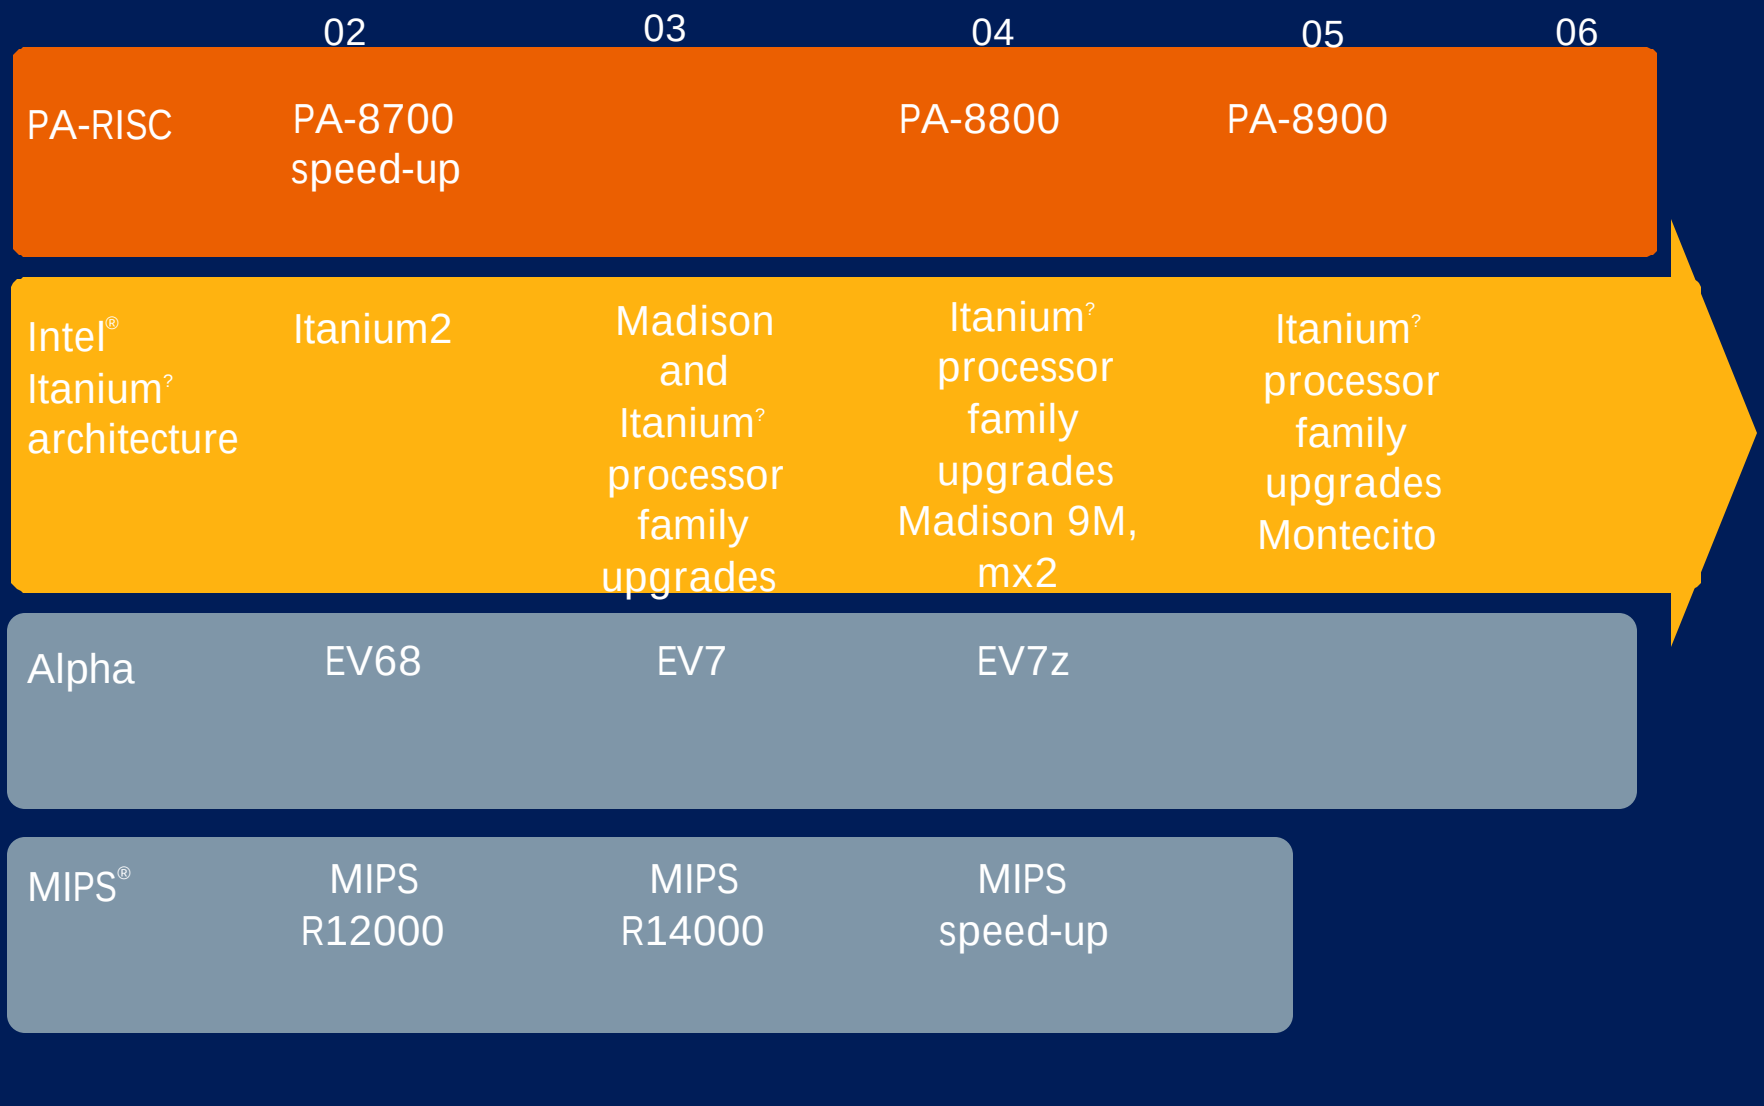

HP's Industry Standards-Based Server Strategy

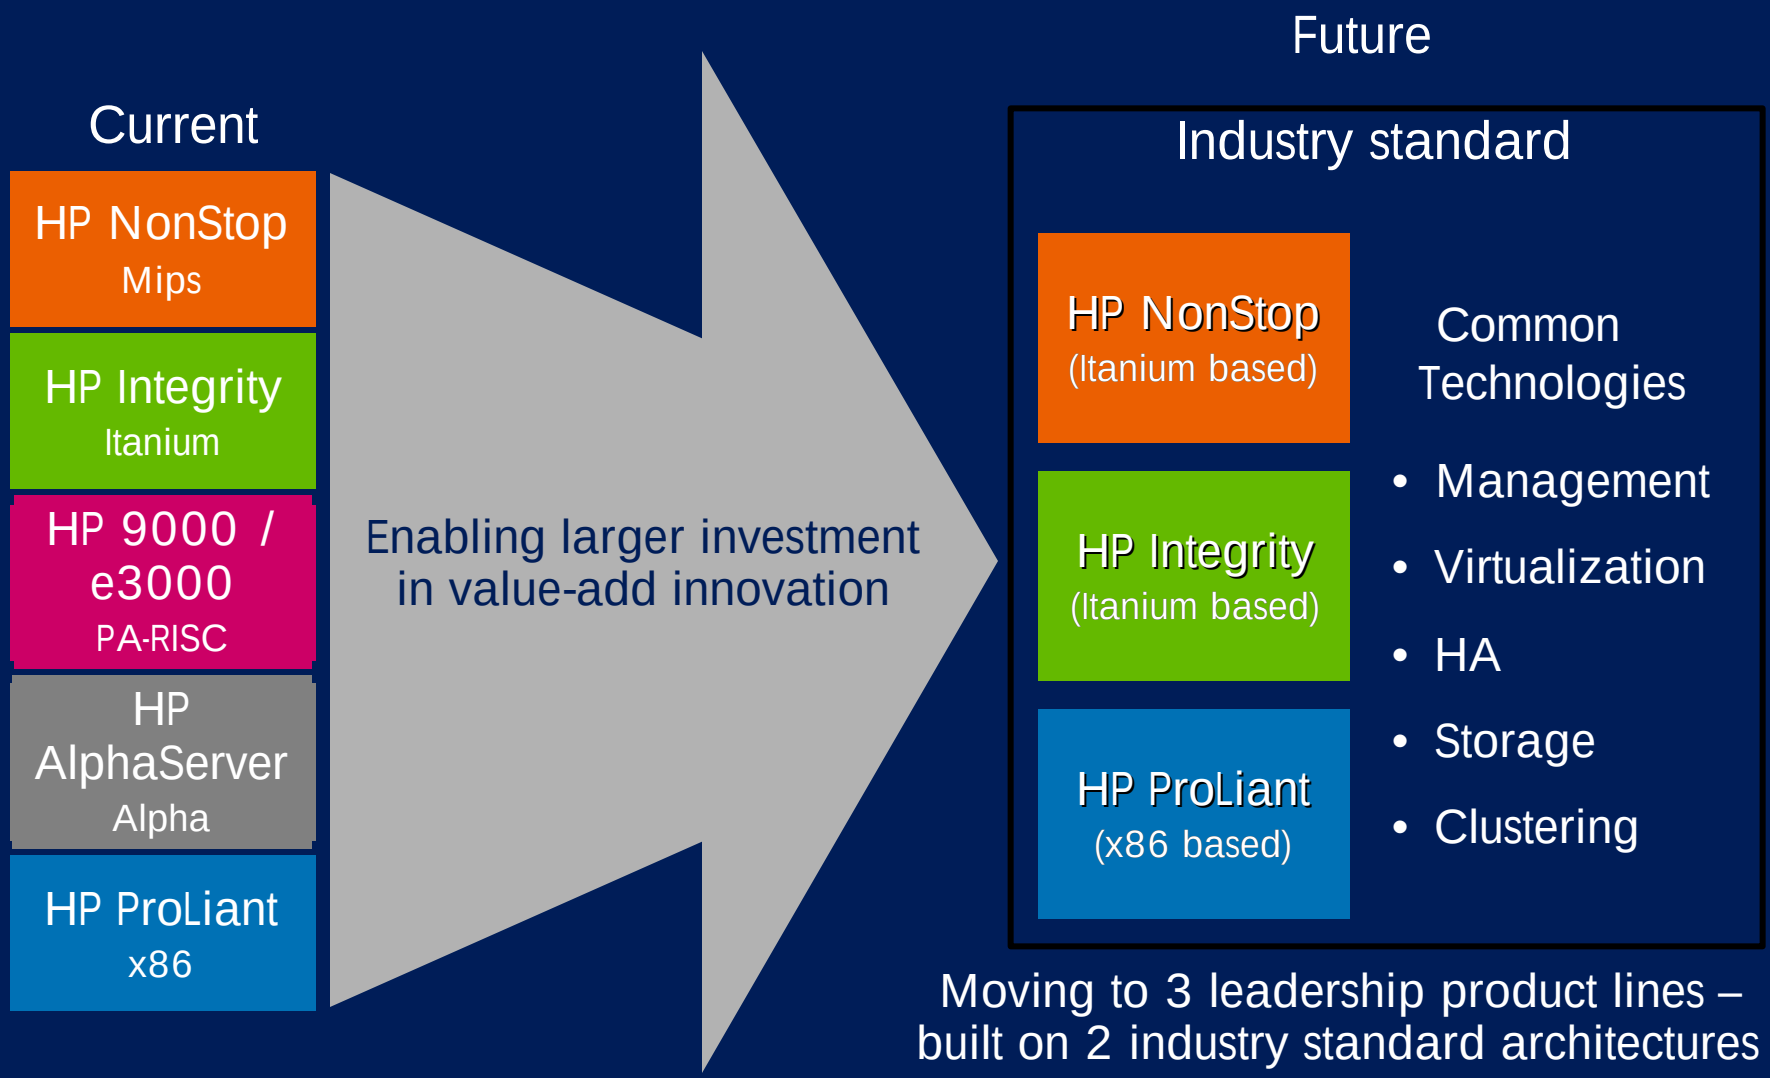

Integrity Server
Die Systemplattform für zukünftige Anforderungen

- Can I stack multiple applications under one operating system?
  Yes
- Can I run a multi operating system environment in one consolidated system?
  Yes
- Can I change operating systems easily?
  Yes
- My main application will no longer be developed under UNIX. New versions are on Linux. Am I in trouble?
  No

Strategische Planung: Bis 2005 zählt Konsolidierung zu den dringlichsten Aufgaben des IS-Managements im Kampf gegen Kosten und Serverwildwuchs (Wahrscheinlichkeit: 0,7).

Warum?

Agenda

- Integrity Server
  - Strategische Serverplattform für das Enterprise Computing
- Prozessortechnologien
  - Die Itanium Architektur
- Chipsets und Prozessormodule
  - HP zx1 Chipset & sx1000 Chipset
  - HP mx2 Dual Prozessornodul
- Low- & Mid-range Server rx1600, rx2600, rx4640
- Mid-range Server rx7620, rx8620
- Superdome

Trends bei Prozessor-Technologien

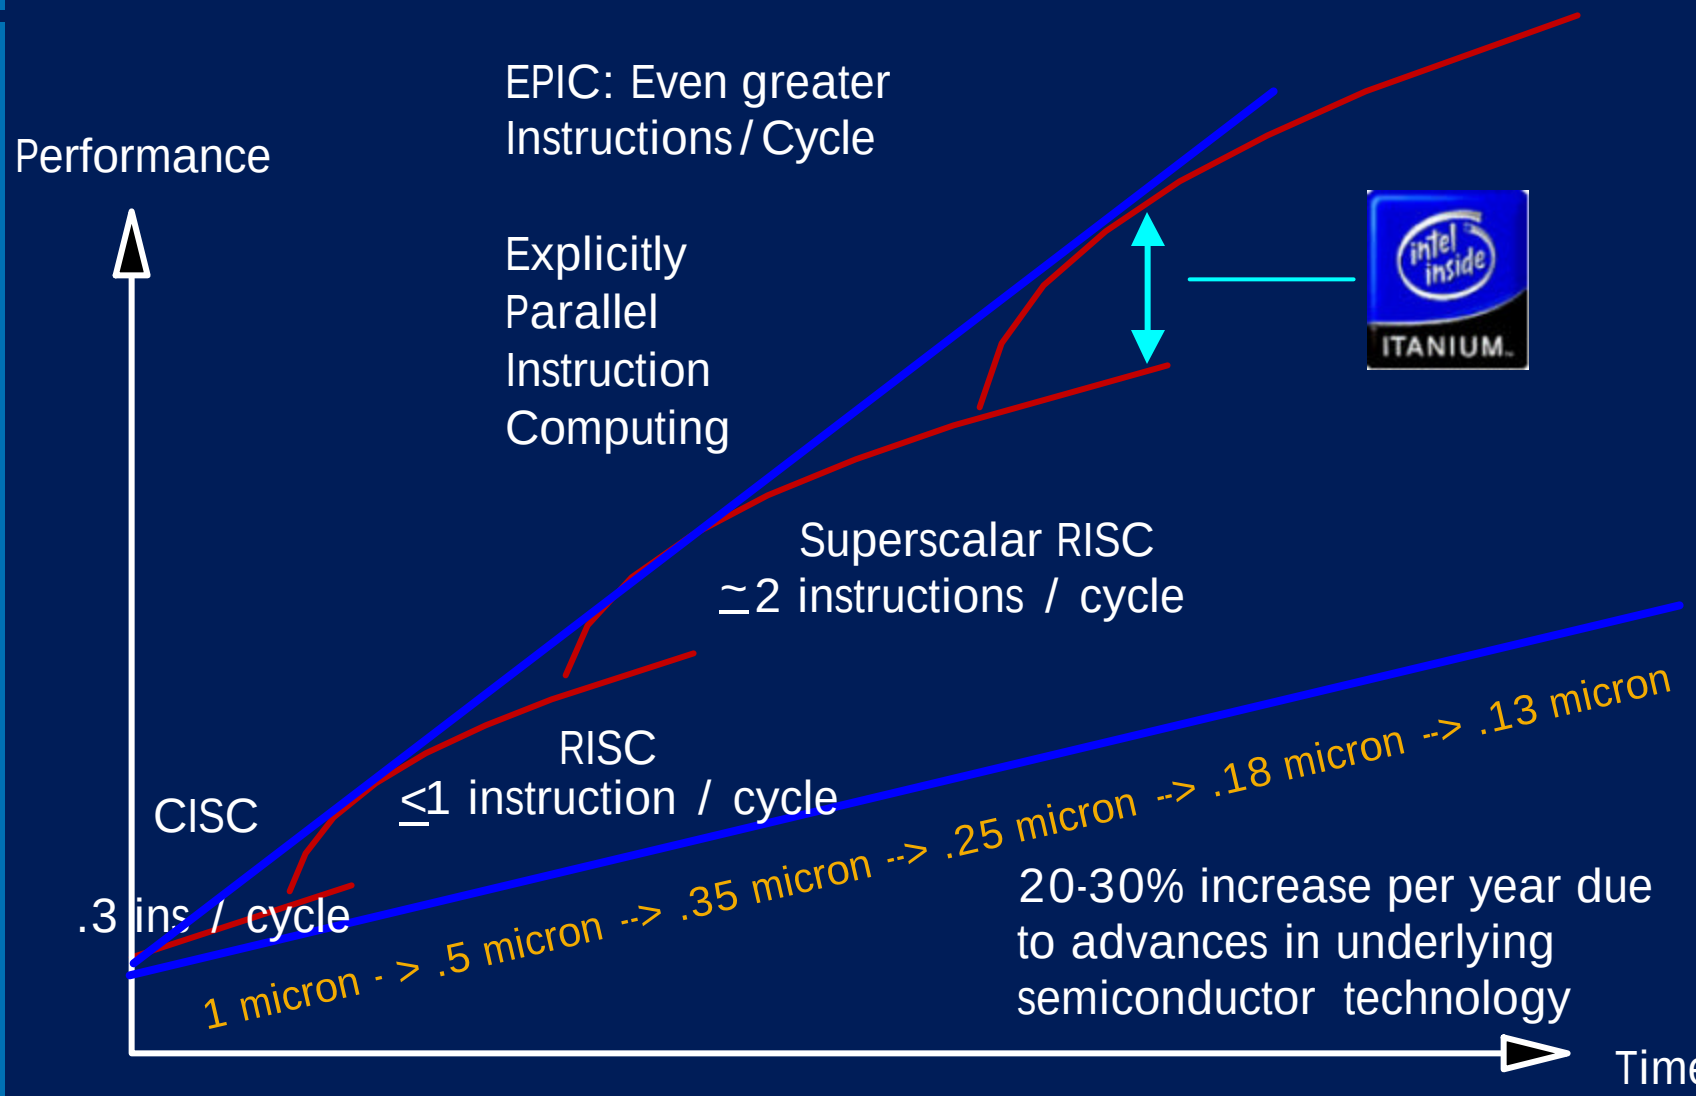

Itanium® Architektur: 4 Key Features

- **Massive resources:** 2* 128 64-bit+ registers. n* Integer Units, m*Floating-Point Units, lots of special registers for branches, predication, loop unrolling etc.
- **Explicit Parallelization:** The compiler 'tells' the processor what can be executed in parallel (and what requires sequential processing)
- **Speculation:** The processor can 'pre-load' data into the caches even if the access is potentially illegal - so it speculates that the data may be valid. Correctness can later be checked in 1 cycle !
- **Predication:** The compiler tells the processor to run (for example) both parts of an 'IF condition' in parallel and then discard the 'wrong part'.

Itanium® Architektur: Explicit Parallelism

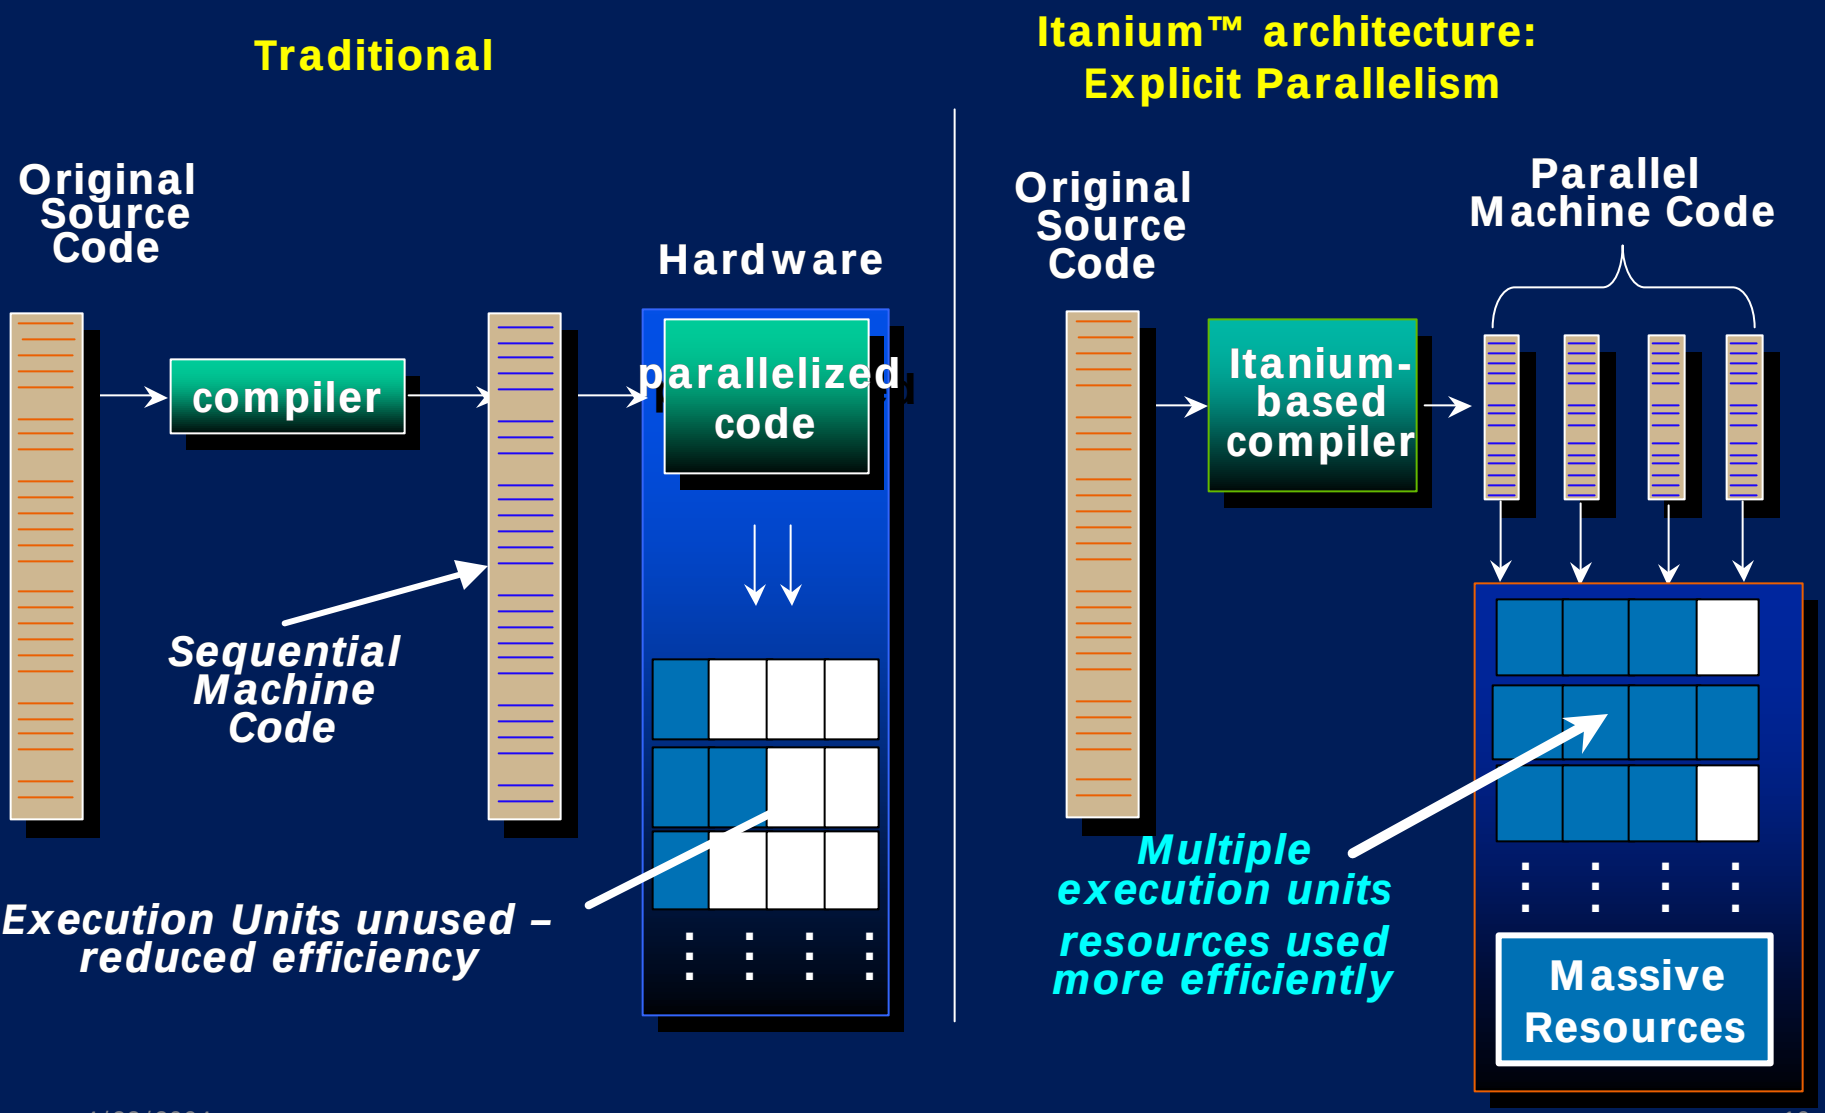

Very Large Instruction Word (VLIW)

- Bundle
  - Set of three instructions (41 bits each)
- Template
  - Identifies types of instructions in bundle and delineates independent operations (through "stops") (5 bits)

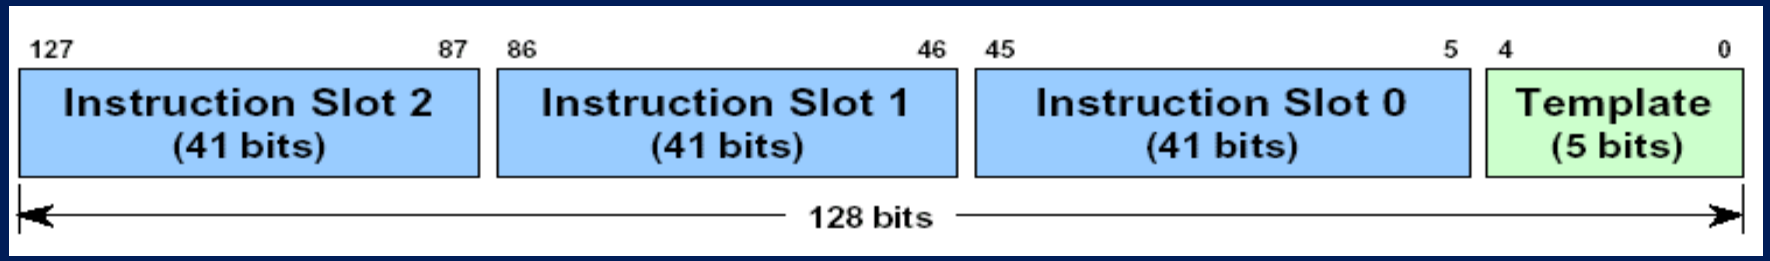

Instruction Handling

Program:

Instruction Groups:
- Explicit group stops
- No RAW or WAW dependencies

Instruction Bundles:
- 3 Instructions and template
- Stops at the end or within

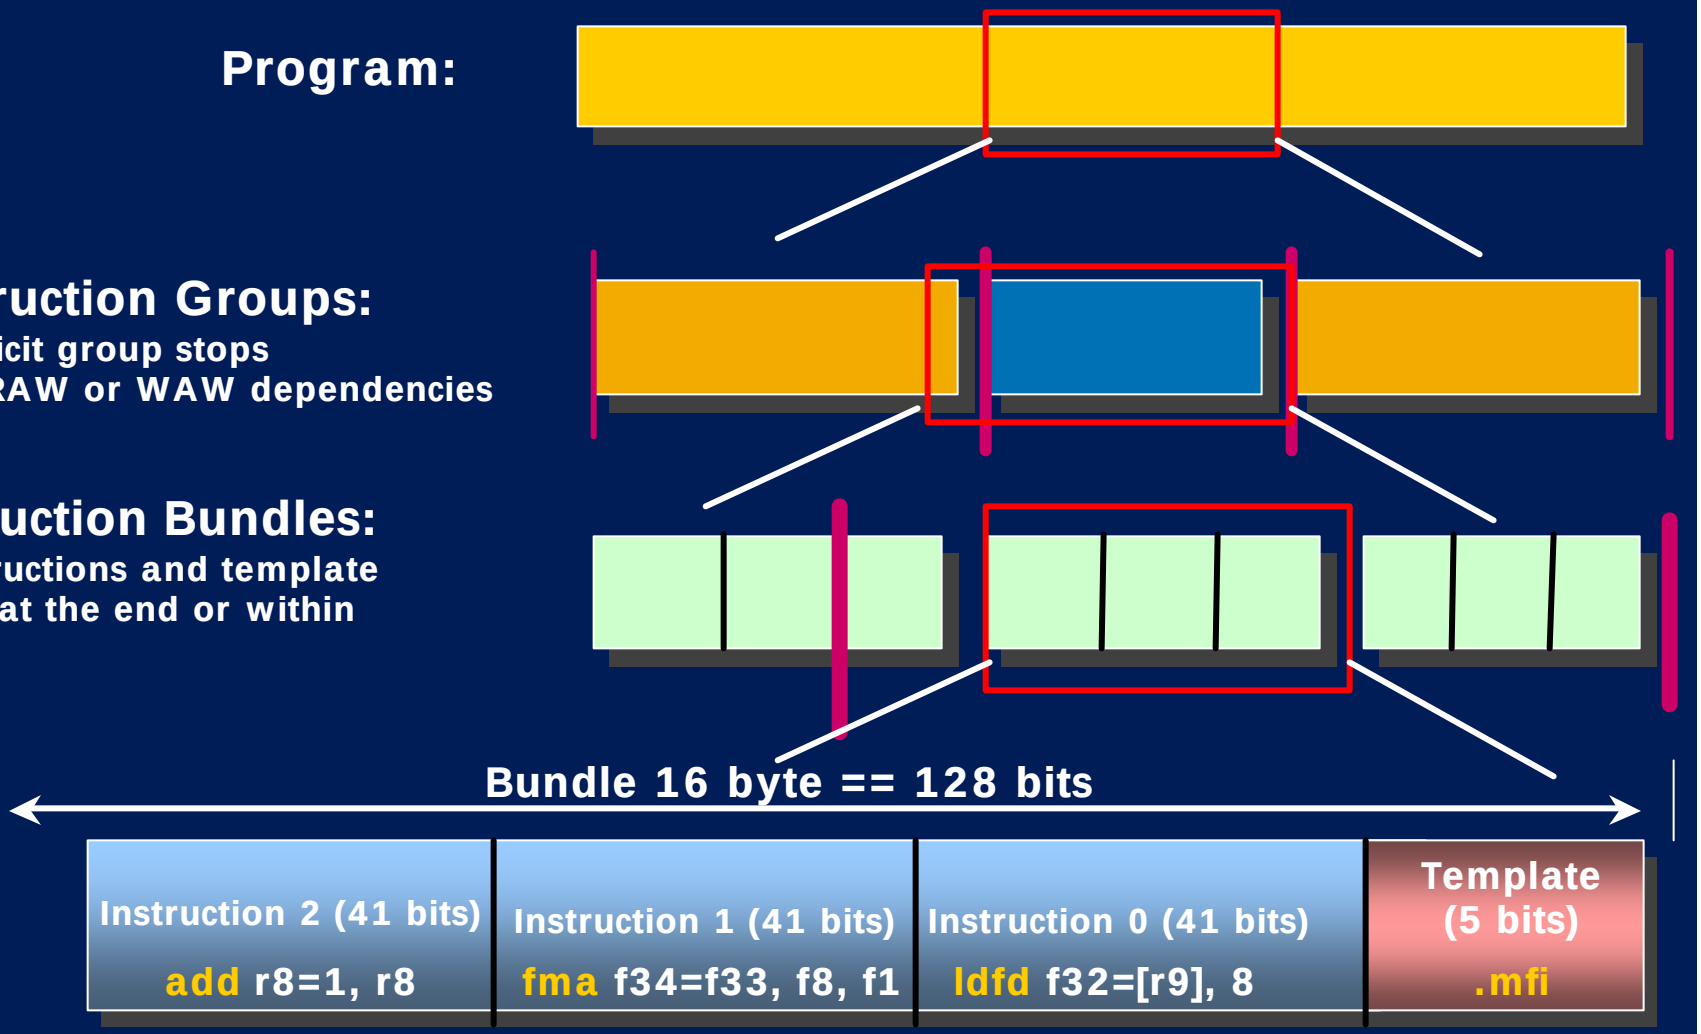

Instruction Handling

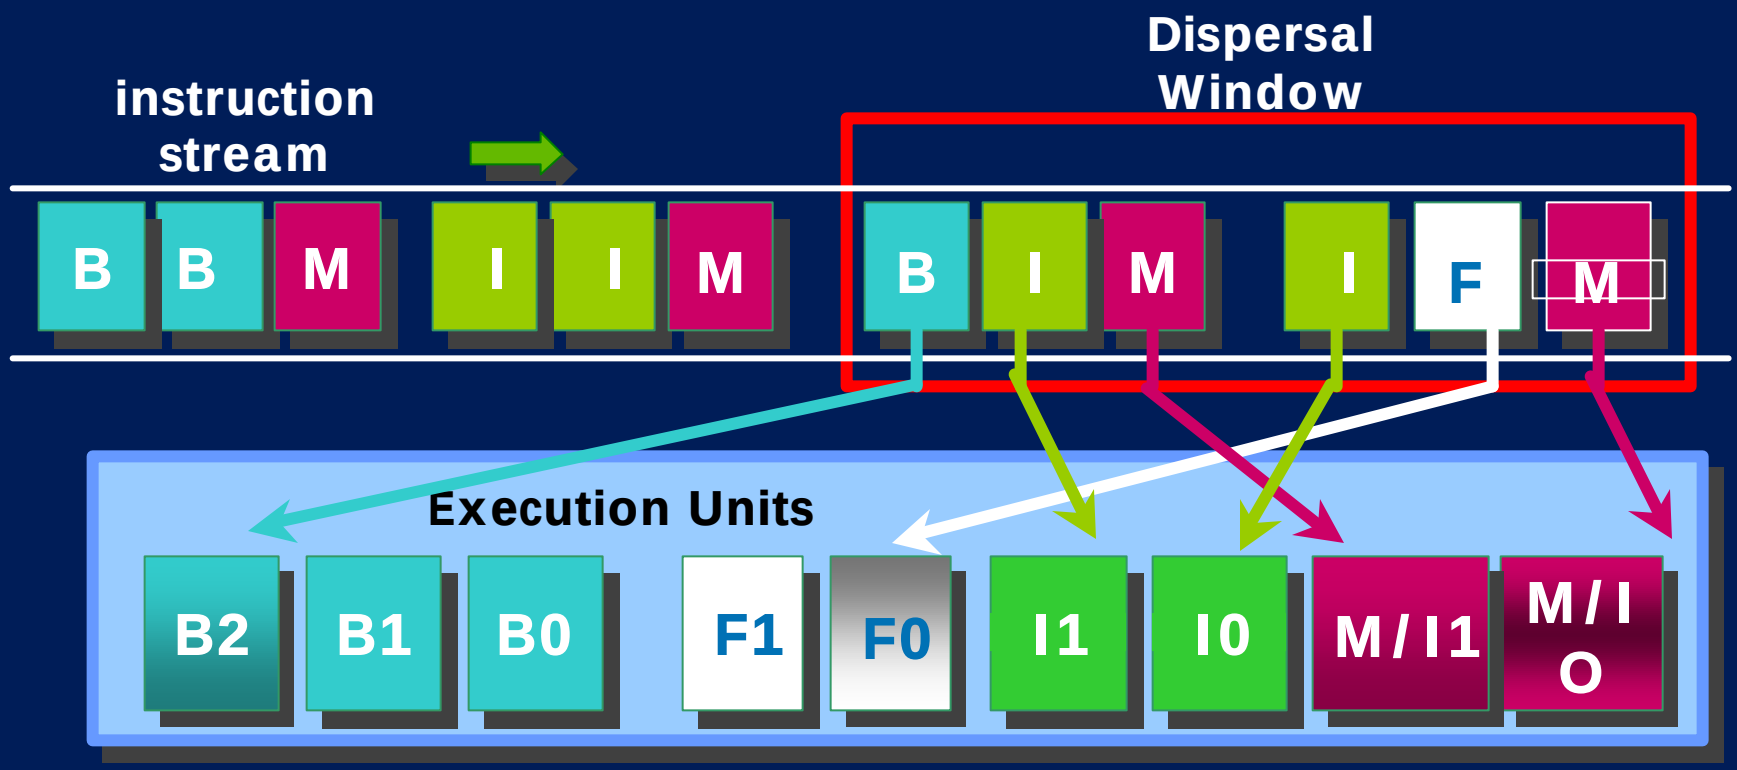

Flexible Issue Capability

Up to 6 instructions executed per clock

Predication

- predication provides the ability to conditionally execute instructions based on computed true / false conditions
- avoids branches
  - predicated instruction either completes or is dismissed
  - (no ops)predicate registers are set by compare / test instructions

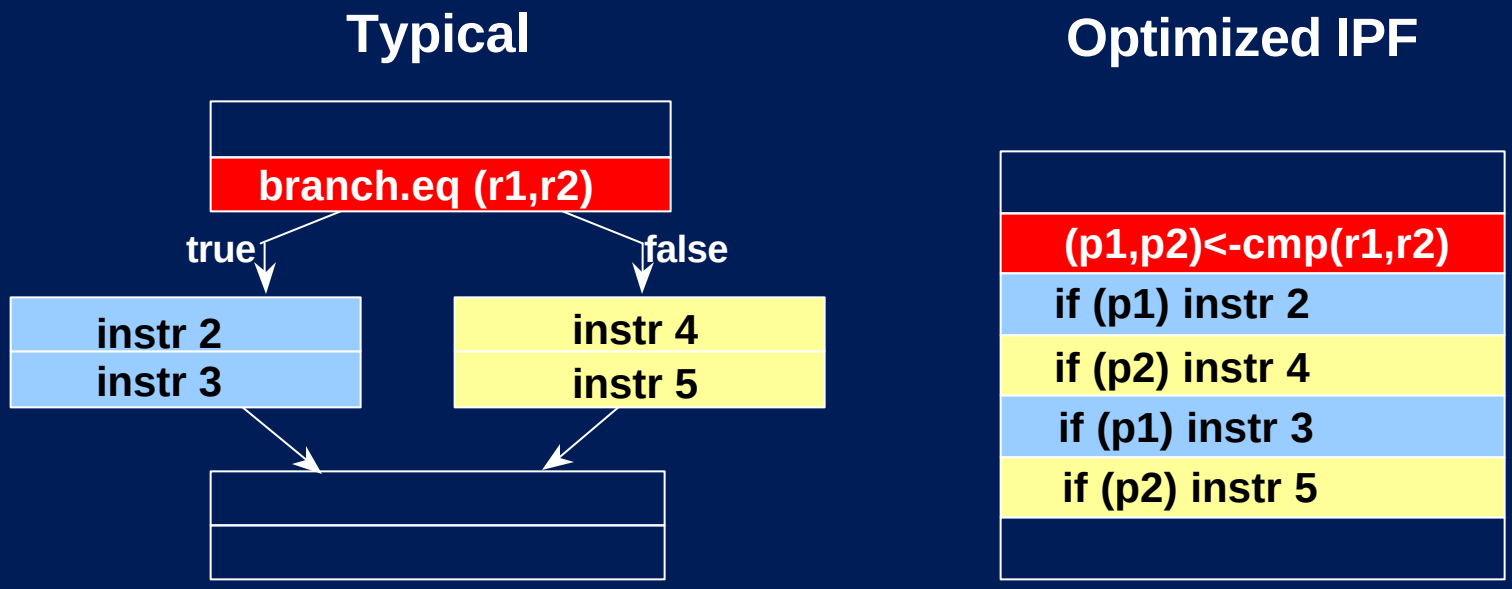

Data Speculation

- allows early execution of loads to hide latency
- advance load before a possible data dependency (load before store)
- speculative load before a branch that guards it

Memory latency can be responsible for 60% or more of processor stalls

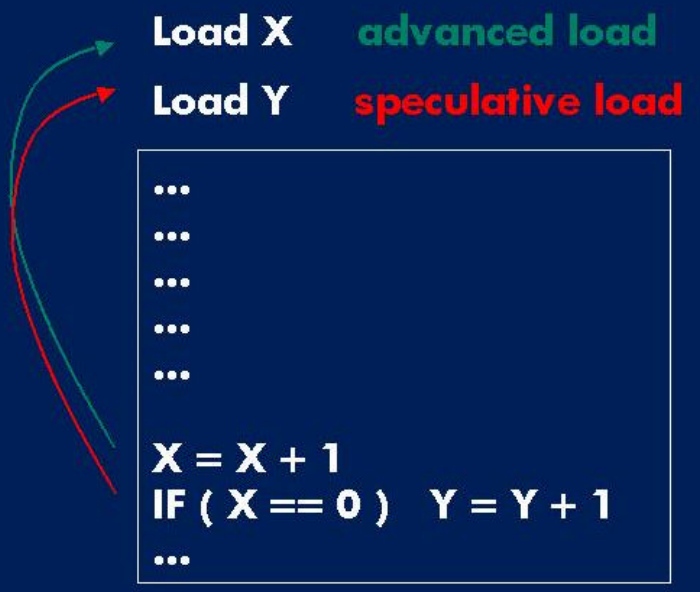

Cache Enhancemnets

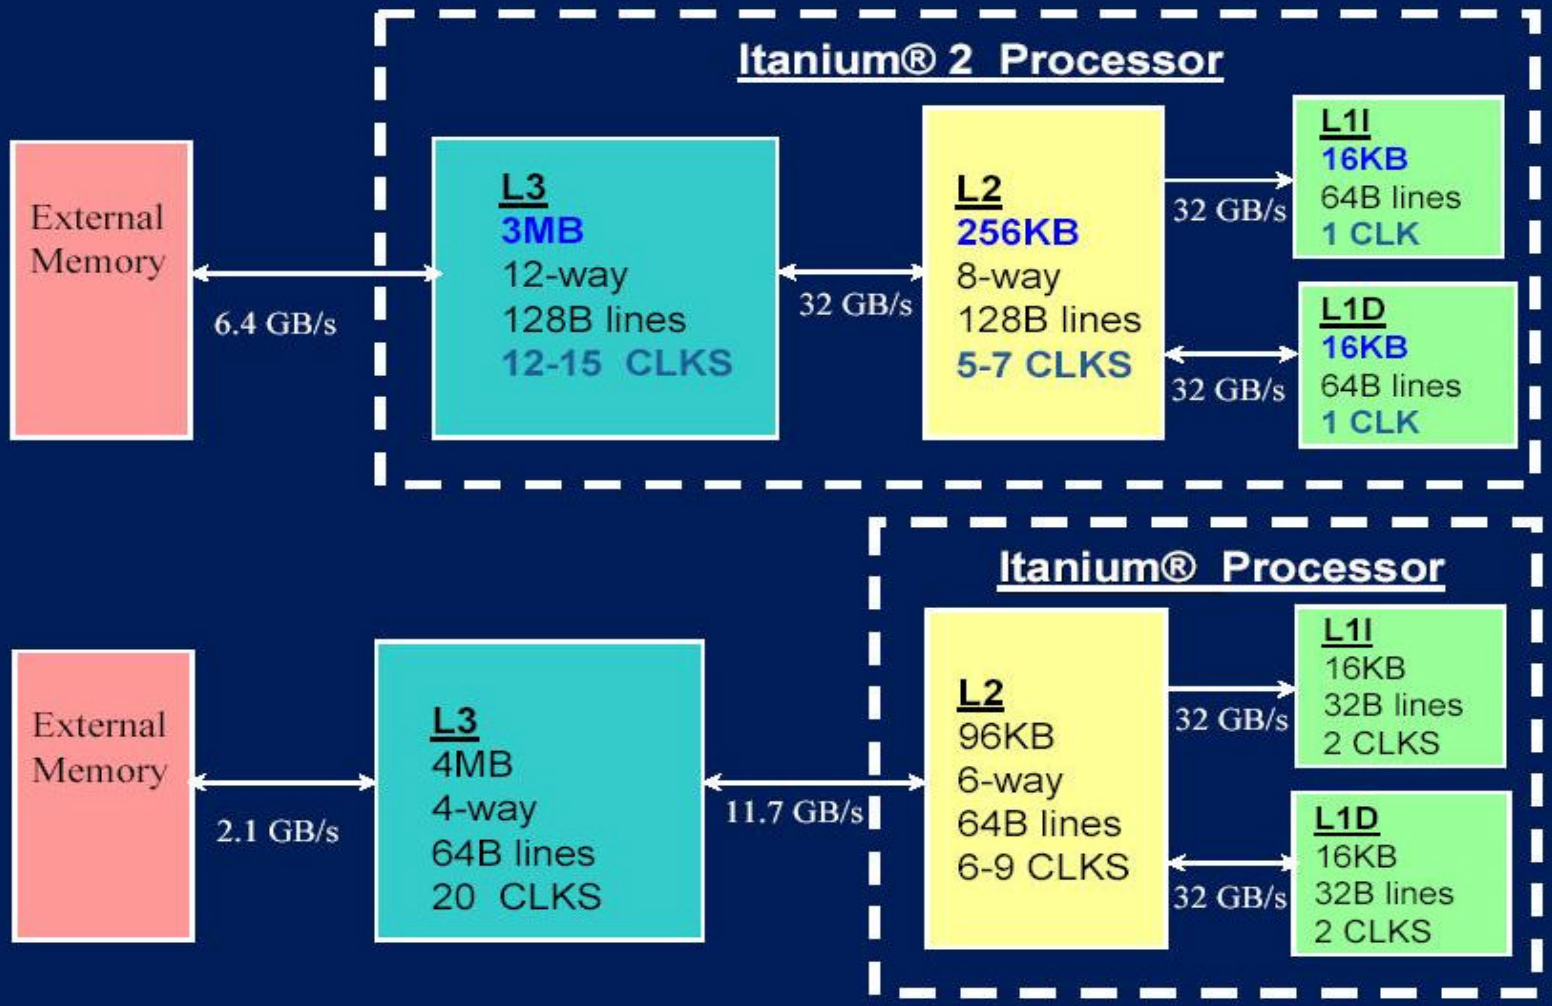

Itanium®-Architektur

Itanium™ Processor

System Bus
64 bits wide
133MHz
2.1 GB/s

Width
2 bundles per clock
4 integer units
2 load or stores per clock

Caches
L1 – 2X16KB - 2 clock latency
L2 – 96K – 12 clock latency
L3 - 4MB external –20 clk
11.7 GB/s bandwidth

Addressing
44 bit physical addressing
50 bit virtual addressing
Maximum page size of 256MB

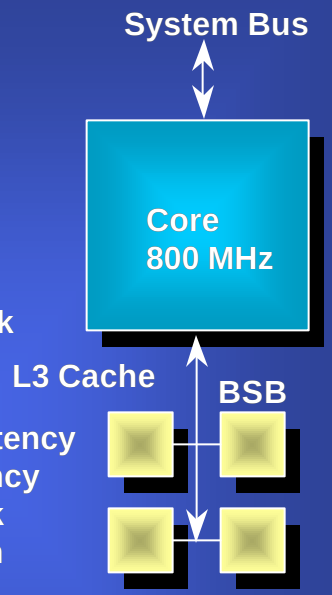

Itanium2 / McKinley / Madison

System Bus
128 bits wide
200MHz
6.4 GB/s

50% Increase in Clock Rate

Width
2 bundles per clock
6 integer units
2 loads and 2 stores per clock

Caches
L1 – 2X16KB - 1 clock latency
L2 – 256K – 5 clock latency
L3 - 3MB / 6MB – 12 clk
32 GB/s bandwidth

Addressing
50 bit physical addressing
64 bit virtual addressing
Maximum page size of 4GB

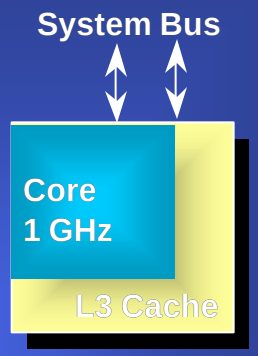

Itanium – Performanceleader in SPECcpu2000

SPECcpu2000 - RISC/EPIC Server Processors

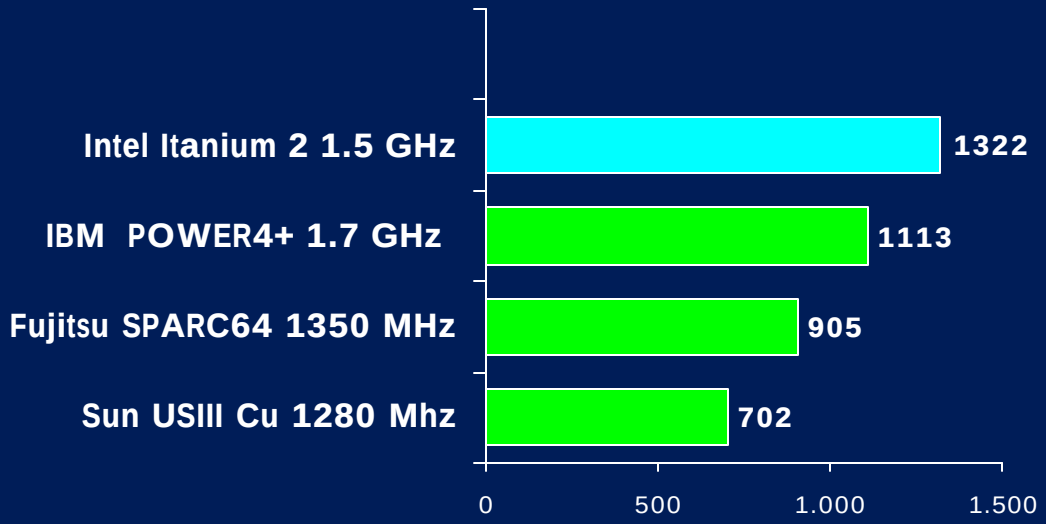

SPECint_base2000

Best SPECint_base2000 for each processor

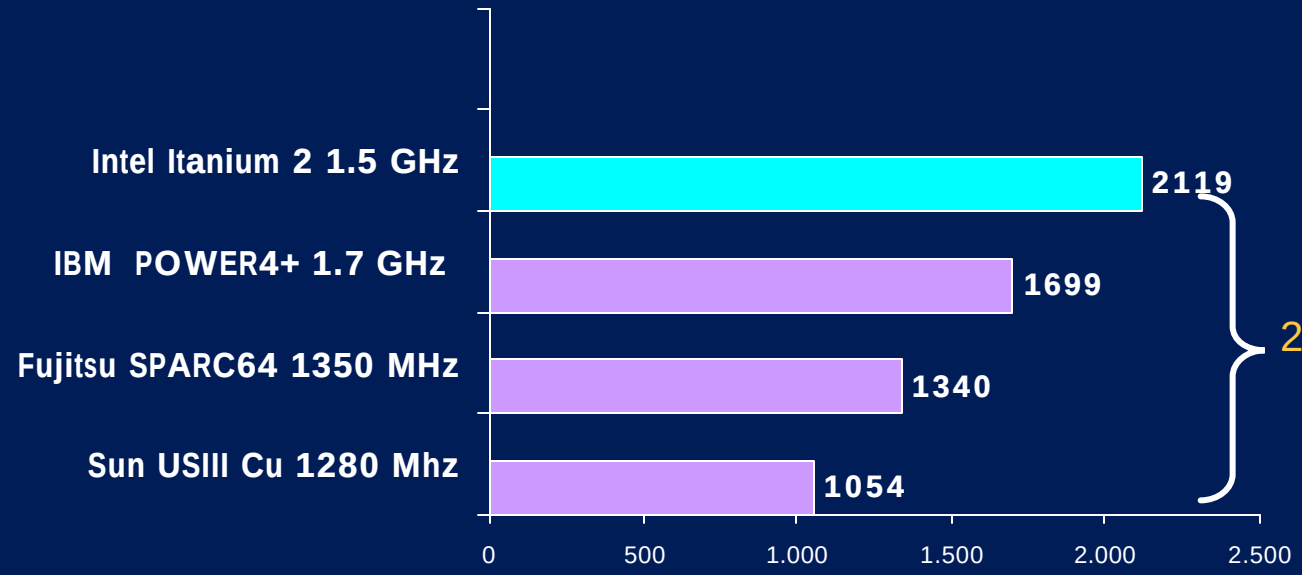

SPECfp_base2000

Best SPECfp_base2000 for each processor

Results as of February 24, 2004
For more information on SPEC, see www.spec.org

Intel® Itanium® Processor Roadmap

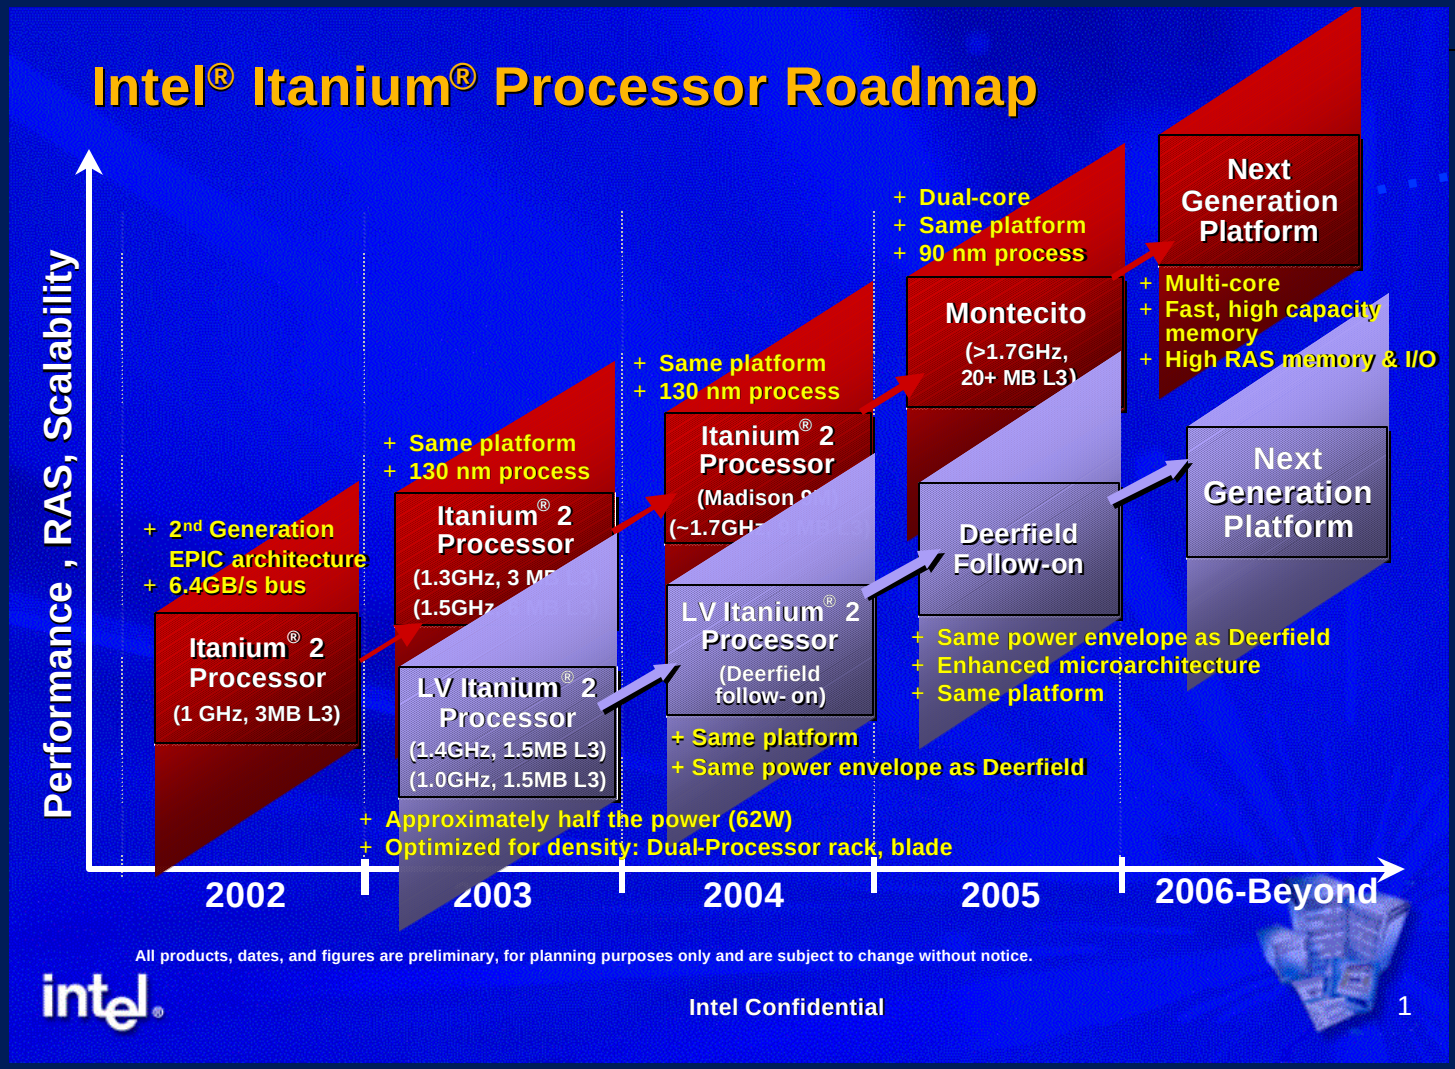

- **HP Integrity rx1600 uses Intel's Low Voltage Itanium 2 Processor**
- **HP Integrity rx2600 uses Intel's Low Voltage Itanium 2 Processor AND Itanium 2 Processor (Madison)**

Agenda

- Integrity Server
  - Strategische Serverplattform für das Enterprise Computing
- Prozessortechnolgien
  - Die Itanium Architektur
- Chipsets und Prozessormodule
  - HP zx1 Chipset & sx1000 Chipset
  - HP mx2 Dual Prozessornodul
- Low- & Mid-range Server rx1600, rx2600, rx4640
- Mid-range Server rx7620, rx8620
- Superdome

What chipsets did HP develop for Itanium® 2 ?

- For workstations and small servers (1-4 CPUs) :

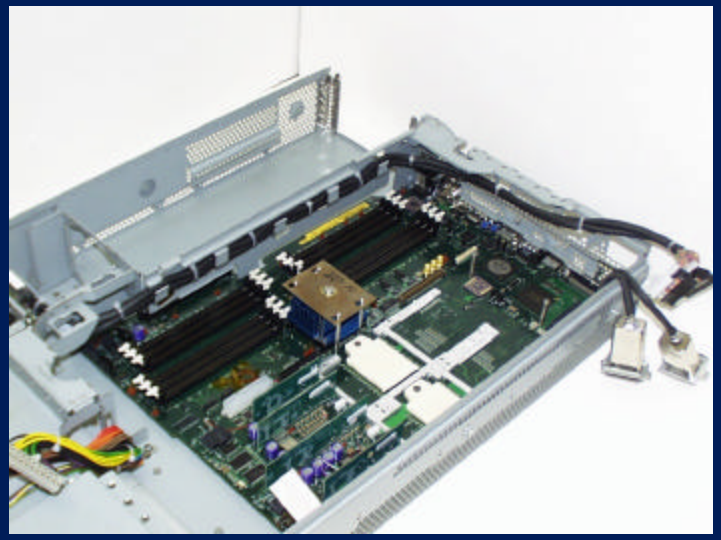

The zx1 chipset

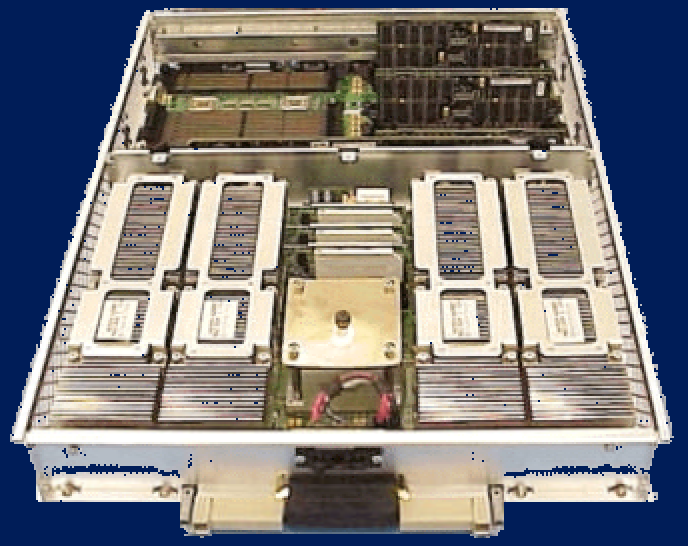

- For medium-sized and big servers

The sx1000 chipset

HP zx1 chipset components

The HP zx1 chipset contains three components :

zx1 mio

- zx1 memory & I/O controller
  - connects to processor bus
  - contains memory controller
  - contains I/O cache controller

zx1 ioa

- zx1 I/O adapter
  - PCI and PCI-X
  - AGP

zx1 sme

- zx1 scalable memory expander
  - optional component used to :
    - increases memory capacity
    - increases memory bandwidth

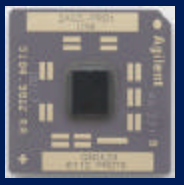

HP zx1 chipset (4-CPU configuration)

System Bus
128 bits wide
400 MT/s
6.4 GB/s

Memory
266 MHz DDR
(Double Data Rate)
High capacity :
48 DIMMs (96GB)

I/O
4 GB/s
PCI-X support

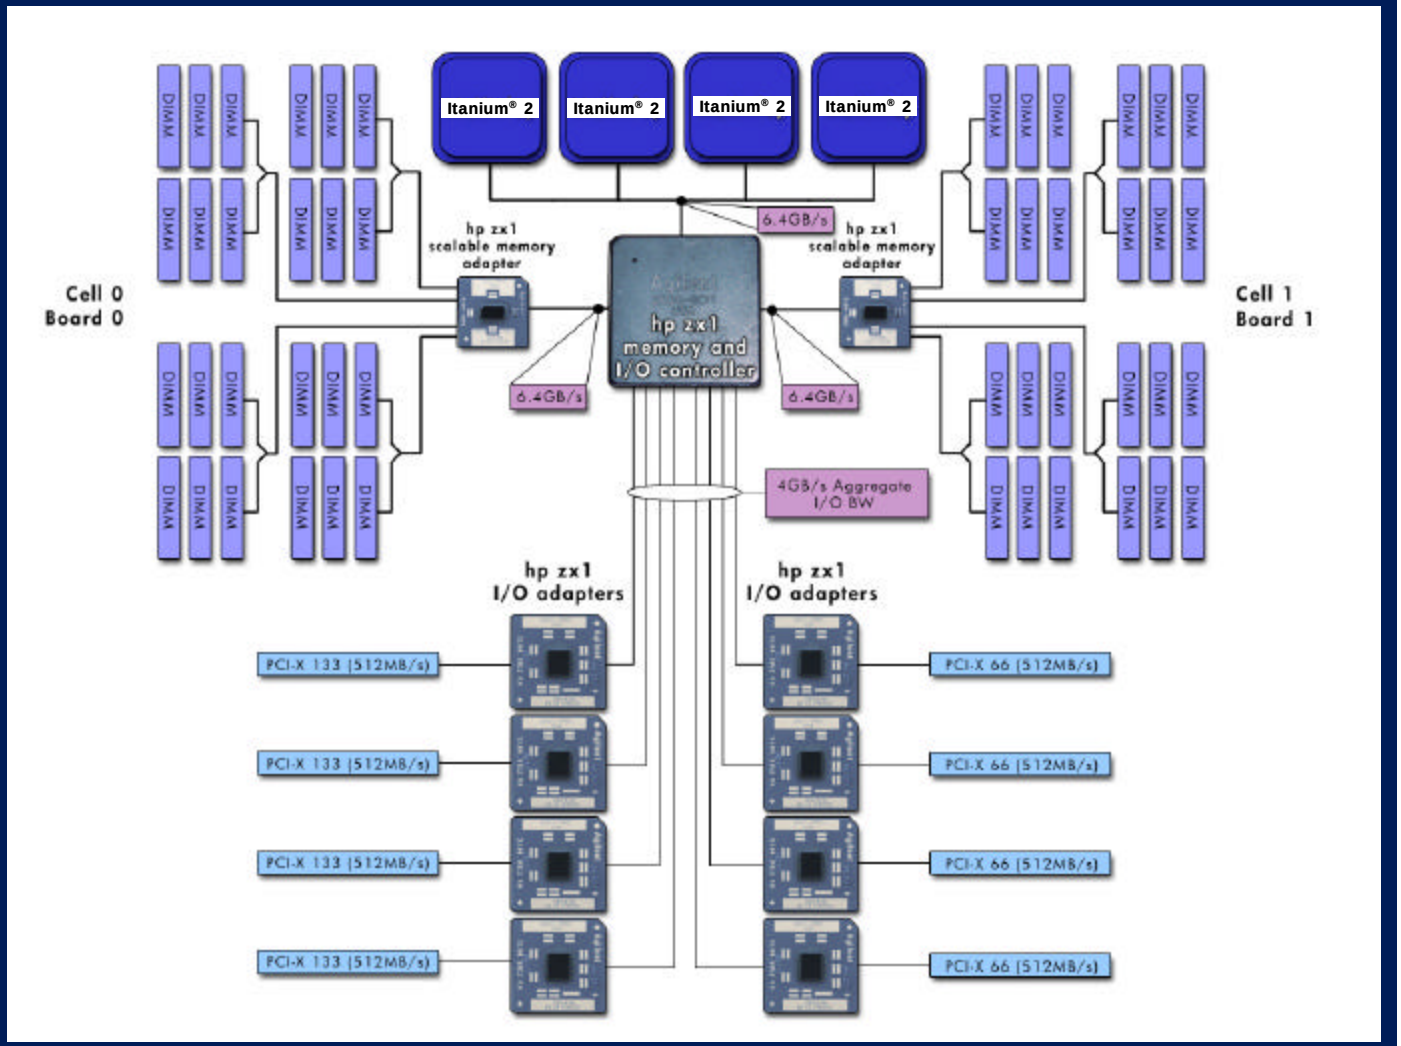

The HP sx1000 chipset

What is the HP sx1000 ?

A new chipset that supports the next generation Itanium® 2 and PA-RISC processors: Madison **and** PA-8800

Where will the sx1000 be used ?

In HP's new high-end and mid-range servers (cell-based)

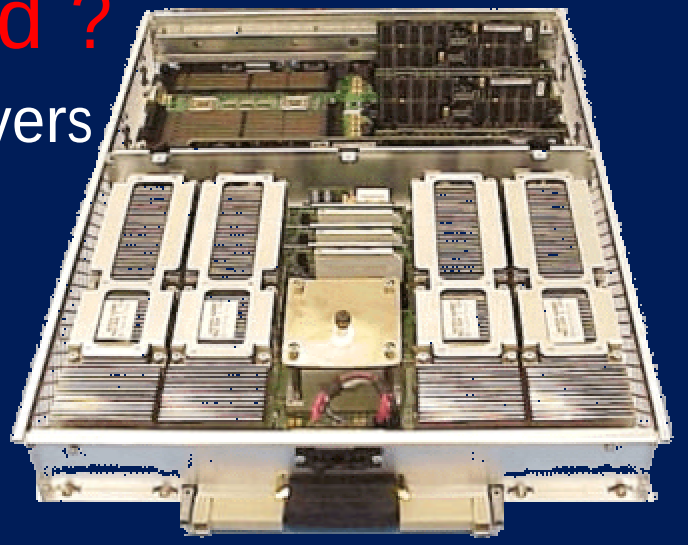

sx1000 chipset components

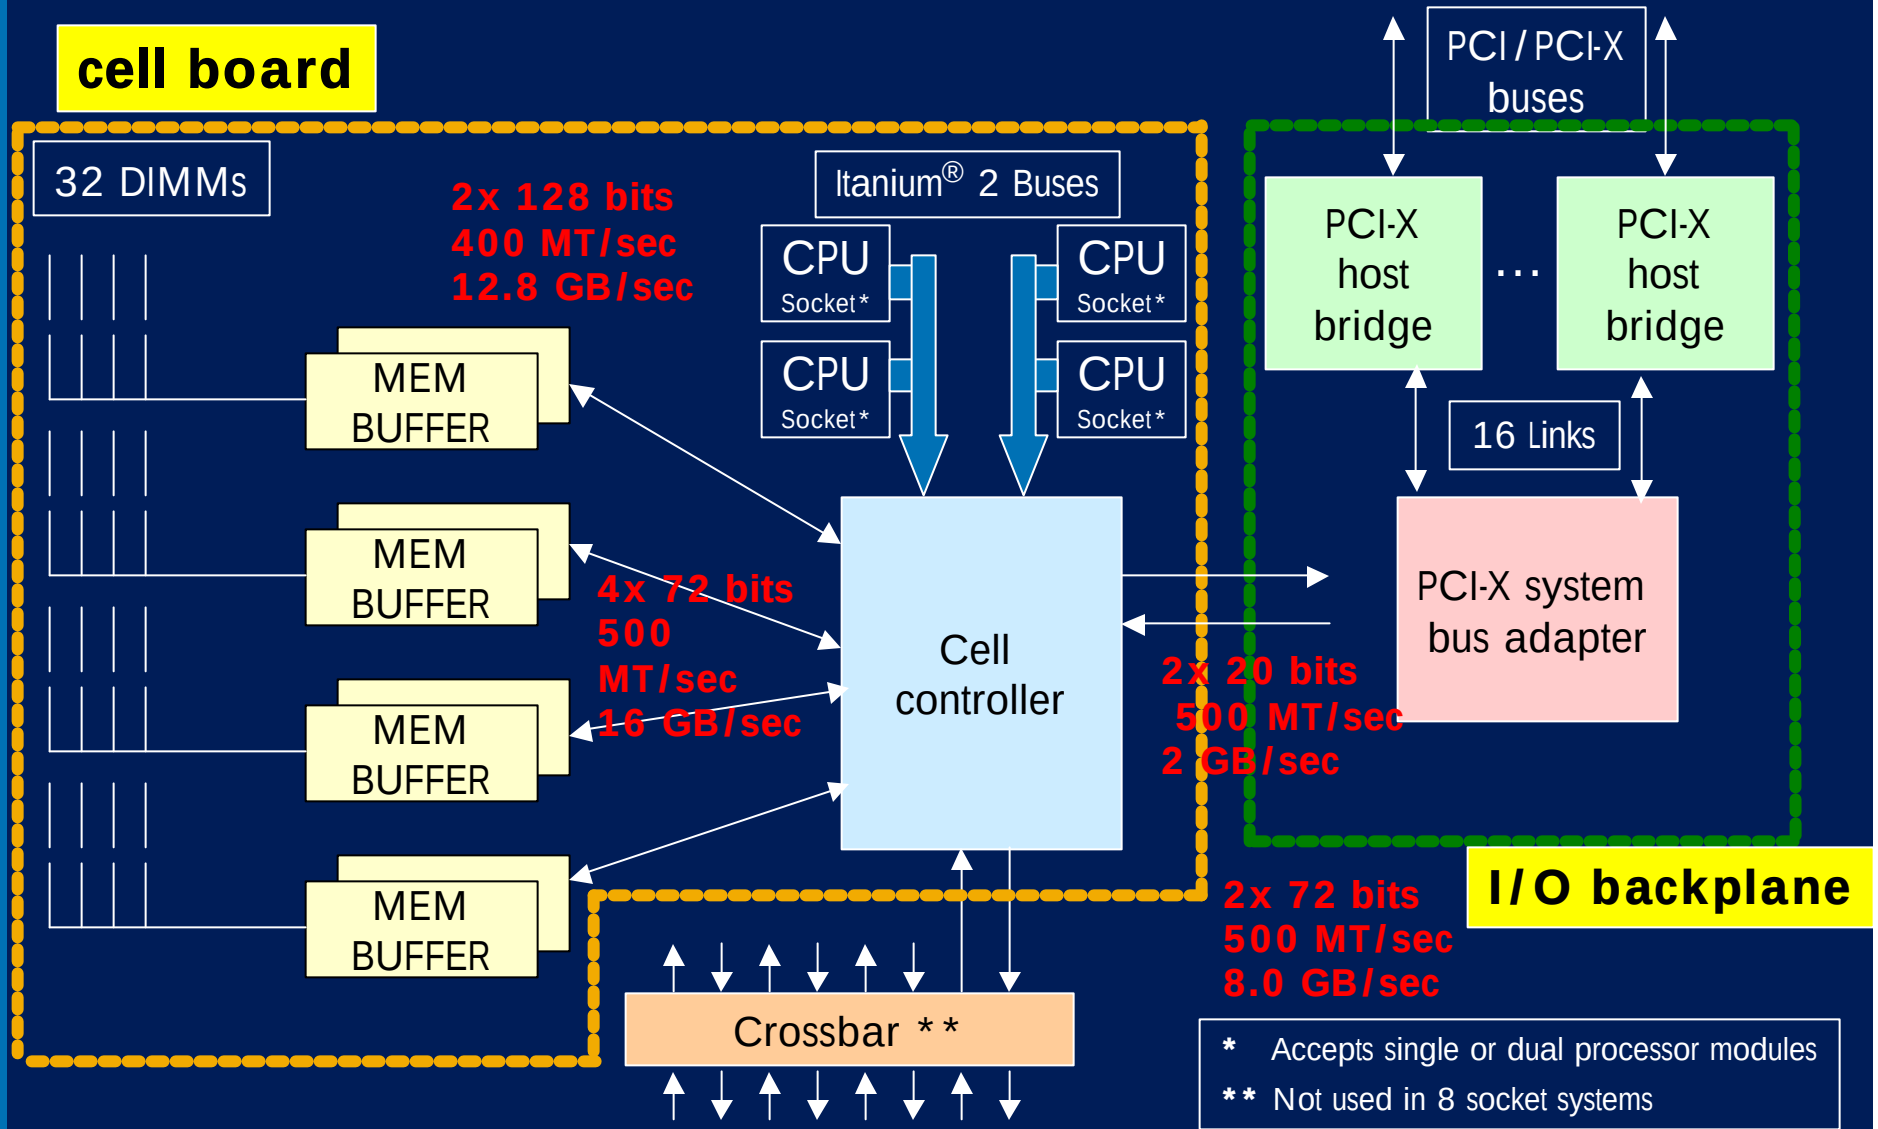

HP mx2 Dual Processor Module

What is the HP mx2?

- An invention by HP which doubles the number of Itanium 2 processors in a server
- A daughter card that combines two future Itanium® processors (Madison) and a 32 MB L4 cache into a single module that is pin-compatible with standard Itanium 2 processor (Madison) sockets

Where will the mx2 be used?

- Across a wide range of Itanium®-based HP systems – from entry-level to high-end servers

When?

- HP will begin shipping mx2 based servers in H1 2004

HP mx2 Dual Processor Module

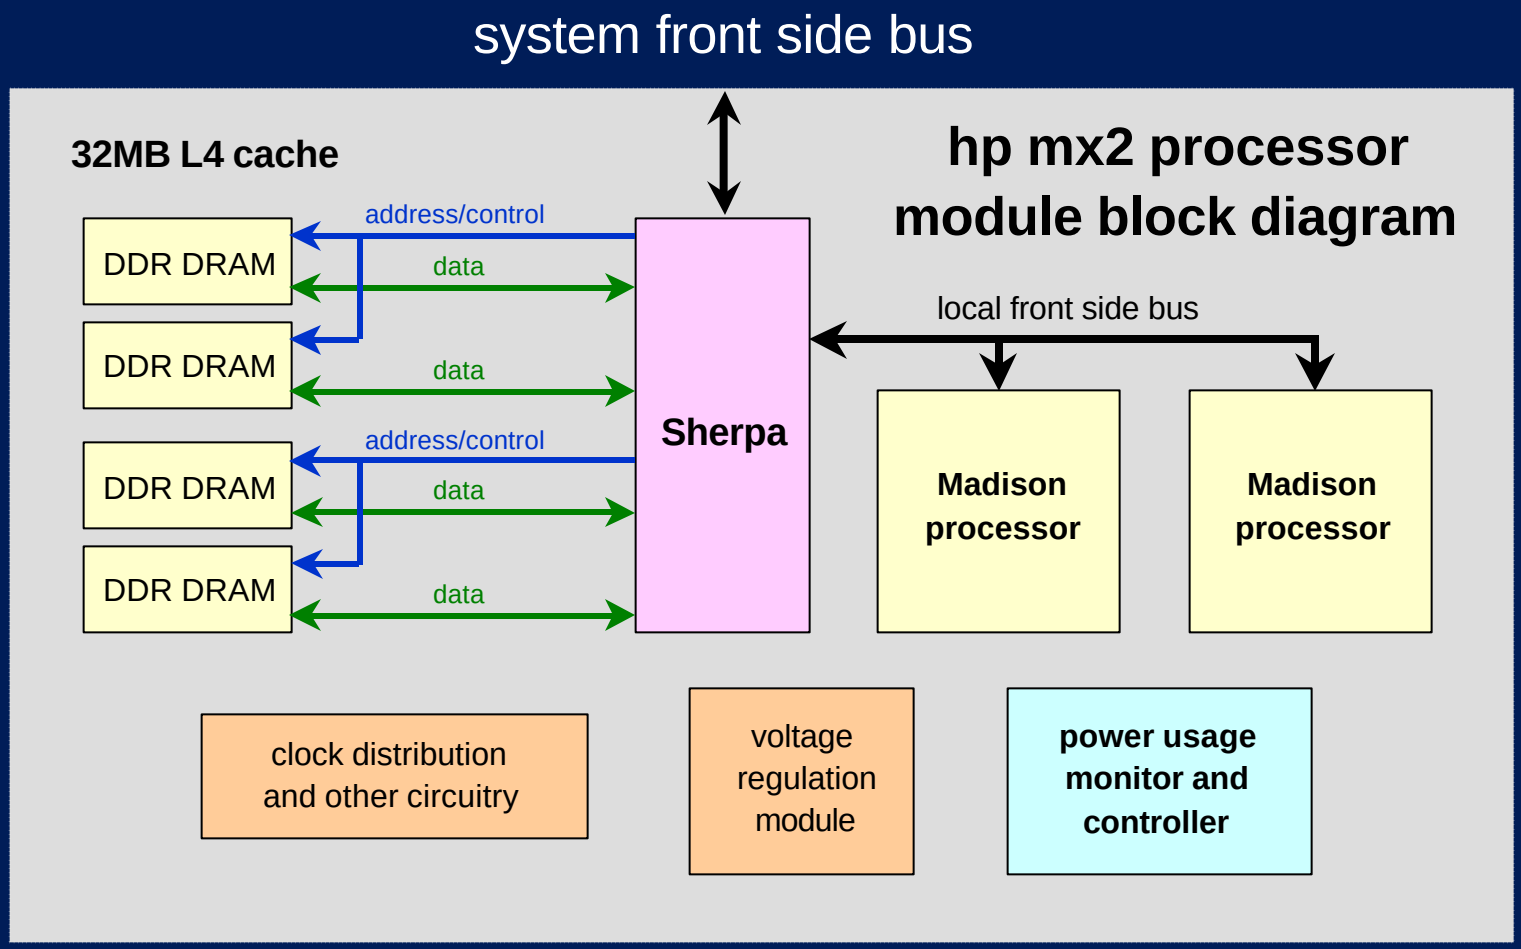

Agenda

- Integrity Server
  - Strategische Serverplattform für das Enterprise Computing
- Prozessortechnologien
  - Die Itanium Architektur
- Chipsets und Prozessormodule
  - HP zx1 Chipset & sx1000 Chipset
  - HP mx2 Dual Prozessornodul
- Low- & Mid-range Server rx1600, rx2600, rx4640
- Mid-range Server rx7620, rx8620
- Superdome

HP Integrity rx1600 - Overview

Processors and chipset

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

Management

- Baseboard management controller with basic control and fault management built in
- Optional extended manageability adds Web/LAN interface with enhanced management capabilities
- Optional 10/100Base-T management LAN

Form factor

- 1 EIA units (U) or 1.75″ height
- 40 servers per 2m rack

Designed for data center and utility closet operation (5–35°C)

Memory

- 512MB to 16 GB (8 slots)
- PC2100 DDR DIMMs w/ ECC
- 128-bit, 266 MHz mem interface (8.5 GB/s peak)
- Chip sparing support

I/O subsystem

- 2 PCI-X 64 bit 133 MHz slots
  - One full length, one half length
- Ultra320 SCSI, Gigabit, 100Base-T, USB, serial built in
- Optional extended built-in I/O: VGA, management LAN, additional serial ports
- 3 GB/s peak I/O bandwidth

High availability

- CPU de-allocation on failure
- Dynamic processor resilience
- Memory chip sparing

Internal peripherals

- 2 hot-plug SCSI HDDs
  - 36 GB (15K rpm)
  - 73 GB (15K rpm)
  - 144GB (10K rpm)
- DVD or DVD/CD-RW

HP Integrity rx2600 - Overview

- One or two **Itanium2™** processors 1.3 or 1.5 GHz or one or two **LV Itanium® 2** processors 1.0 or 1,4 Ghz
- 4 independent PCI-X slots
- 200 MHz processor bus
  - 6.4 GB/s of bandwidth
- 12 memory DIMM slots
  - DDR SDRAM
  - 24 GB of memory
  - 8.5 GB/s of memory bandwidth (peak)
  - multi-bit memory error correction
- System Package
  - 2U, 26 inches deep
  - 3 hot-swap bays for 1” hard disks, 1 built-in DVD ROM
  - Optional redundant power supplies

rx1600 versus rx2600

	HP Integrity rx1600 server 2-way, 1U Itanium server solution	HP Integrity rx2600 server 2-way, 2U Itanium server solution
positioning		
processors	LV Itanium® 2 at 1GHz/1.5M L3 cache	Itanium® 2 1.5Ghz/ 6M L3 cache Itanium® 2 1.3GHz / 3M L3 cache LV Itanium® 2 1.4Ghz/1.5M L3 cache LV Itanium® 2 1Ghz/1.5M L3 cache
memory	up to 16GB DDR SDRAM	up to 24GB DDR SDRAM
bandwidth	6.4 GB/s system; 8.5 GB/s memory; 3.5 GB/s I/O	6.4 GB/s system; 8.5 GB/s memory; 4.0 GB/s I/O
pci-x slots	two 64-bit PCI-X @ 133MHz	four 64-bit PCI-X @ 133MHz
internal storage	up to 292GB	up to 438GB
operating system	HP-UX and Linux at first release Windows and OpenVMS to follow	HP-UX and Linux at first release Windows and OpenVMS to follow
high availability	2 hot-plug disks chip spare technology memory scrubbing and page de-allocation CPU failure de-allocation	redundant, hot-swap fans and power hot-plug disks chip spare technology memory scrubbing and page de-allocation CPU failure de-allocation
application segment	System management, security, software engineering, Internet infrastructure	MP: collaboration, app server ERP, SCM, Biz Intelligence, CRM, workgroup DP: HPTC, system mgmt, security, Internet infrastructure

HP Integrity rx4640

- 4U height, 19” width, ~26” depth
- Up to 4 Itanium2 CPUs at 1.3 GHz w/3MB cache or 1.5 GHz w/6MB cache
- HP *zx1* chipset
- up to 64GB of industry standard DDR memory
- Six 64-bit, PCI-X slots
  (2@133MHz, 4@66MHz)
- Integrated Gbit LAN, Ultra160 SCSI, DVD and LS240 drives
- 2 low profile, hot plug SCSI HDD’s
- Redundant hot-plug PSU’s & fans
- Integrated remote manageability

Announced in Nov 2003

Agenda

- Integrity Server
  - Strategische Serverplattform für das Enterprise Computing
- Prozessortechnologien
  - Die Itanium Architektur
- Chipsets und Prozessormodule
  - HP zx1 Chipset & sx1000 Chipset
  - HP mx2 Dual Prozessornodul
- Low- & Mid-range Server rx1600, rx2600, rx4640
- Mid-range Server rx7620, rx8620
- Superdome

Itanium-based rx7620 & rx8620

1. Significant performance advantage over RISC-based systems enabled by powerful 64-bit Itanium processors and the sx1000 chipset
2. Support for 4 different operating systems – HP-UX11i, Linux, Windows Server, and OpenVMS ('04/'05) enables tremendous application availability
3. Ability to run multiple operating systems in the same server at the same time gives customers unparalleled flexibility and opportunity for consolidation

"For commercial markets that require faster secure transactions and business processing and HPTC markets that require raw performance, the hp servers rx7620 and rx8620 provide the industry's leading application availability, performance, and flexibility to best address today's demanding technical and business workloads."

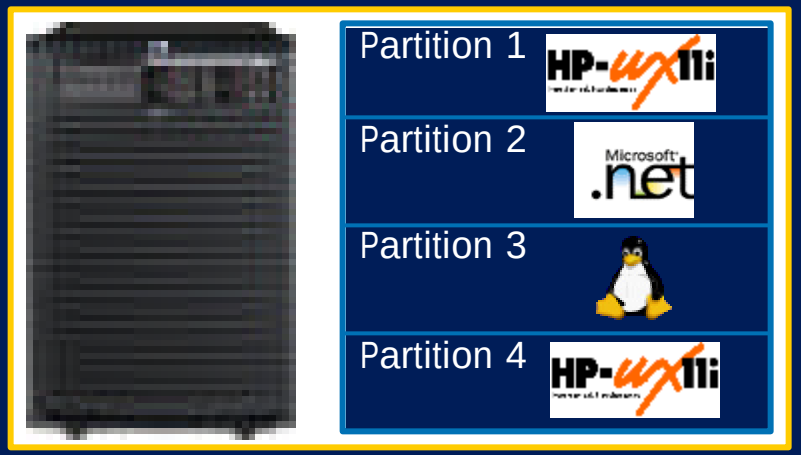

Note: For illustrative purposes only. All operating systems may not be available immediately upon product announcement.

HP Integrity rx7620

- 2- to 8-way Itanium® 2 processors at 1.3 GHz w/3MB cache or 1.5 GHz w/6MB cache
- HP-UX, Windows, VMS and Linux
- 64GB memory
- 15 PCI-X slots
- 1 removable media (DAT/DVD)
- 4 internal HDDs
- dual core I/O
- up to 2 hard partitions
- up to 8 virtual partitions (not at 1st release)

HP Integrity rx8620

What is it?

- A 16-socket server related to HP9000 rp8420
- In-box upgradeable from HP9000 rp8400/8420

Key Product features

1 to 16 CPUs (1.3 and 1.5 GHz Intel Itanium2)

HP sx1000 chipset

Up to 16 PCI-X slots (32 with optional SEU)

Up to 2 servers per 2m rack

Up to 2 hard partitions (4 with optional SEU)

In-box upgradeable to future Itanium processors

Multi-OS support:

- ✓ HP-UX11iv2
- ✓ Linux (1H 2004)
- ✓ Windows Server 2003 Enterprise Edition (1H 2004)
- ✓ Windows Server 2003 Datacenter Edition (1H 2004)
- ✓ OpenVMS (2005)

HP Integrity rx8620 – Front View

Agenda

- Integrity Server
  - Strategische Serverplattform für das Enterprise Computing
- Prozessortechnologien
  - Die Itanium Architektur
- Chipsets und Prozessormodule
  - HP zx1 Chipset & sx1000 Chipset
  - HP mx2 Dual Prozessornodul
- Low- & Mid-range Server rx1600, rx2600, rx4640
- Mid-range Server rx7620, rx8620
- Superdome

Superdome with Itanium ® 2 Hardware Configurations

- 2-way and 4-way cell boards can be mixed within the system.
- 512 MB and 1 GB DIMMs can be mixed within the system.
- Operating systems:
  - HP-UX 11i version 2: July 28, 2003
  - Windows Server 2003: October 2003
  - Linux RedHat Advanced Server 3.0: December 2003
  - OpenVMS V8.x: 2005

	Superdome 16-way	Superdome 32-way	Superdome 64-way
# of 2-way or 4-way cell boards	1–4	1–8	3–16
Itanium 2 1.5GHz CPUs (with 2-way or 4-way cell boards)	2–16	2–32	6–64
Memory (with 512 MB or 1 GB DIMMs)	2–128 GB	2–256 GB	6–512 GB
12-slot I/O card cages	1-4	1-4 1-8 with IOX	1-8 2-16 with IOX
Hot swap I/O slots (OS specific)	12-48	12-48 12-96 with IOX	24-96 24-192 with IOX
Hot swap redundant power supplies	4	6	12
Hot swap redundant fans	4	4	8
I/O fans	5	5	10
nPartitions (npars)	4	8	16

Superdome with Intel® Itanium® 2 architecture

- Performance
  - Crossbar bandwidth: 64 GB/s
  - I/O bandwidth: 32 GB/s (2.0 GB/s per cell)
  - Memory bandwidth: 256 GB/s (16 GB/s per cell)
  - Max latency: 440 ns
- Hardware reliability
  - ECC on all CPU and memory paths
  - Parity-protected I/O data paths
  - Full SDRAM correction (chipkill)
  - Online power and fan replacement
  - N+1 power and fan
  - Dual power sources

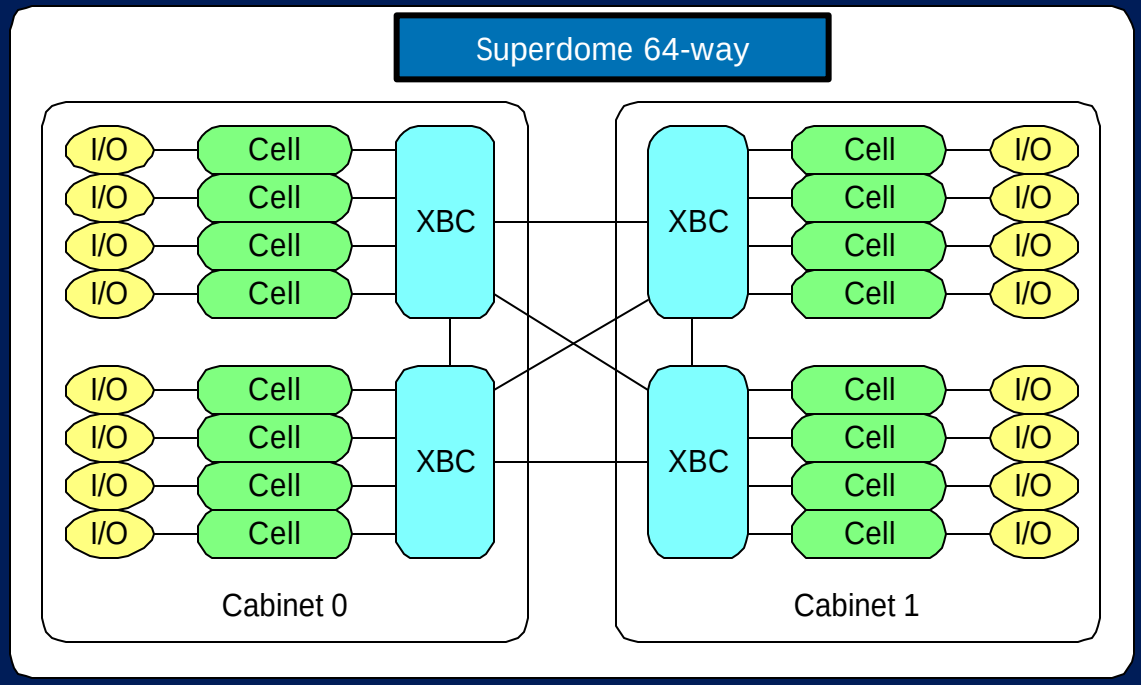

HP Integrity Superdome: World Record in Enterprise UNIX/Oracle, Windows/SQL

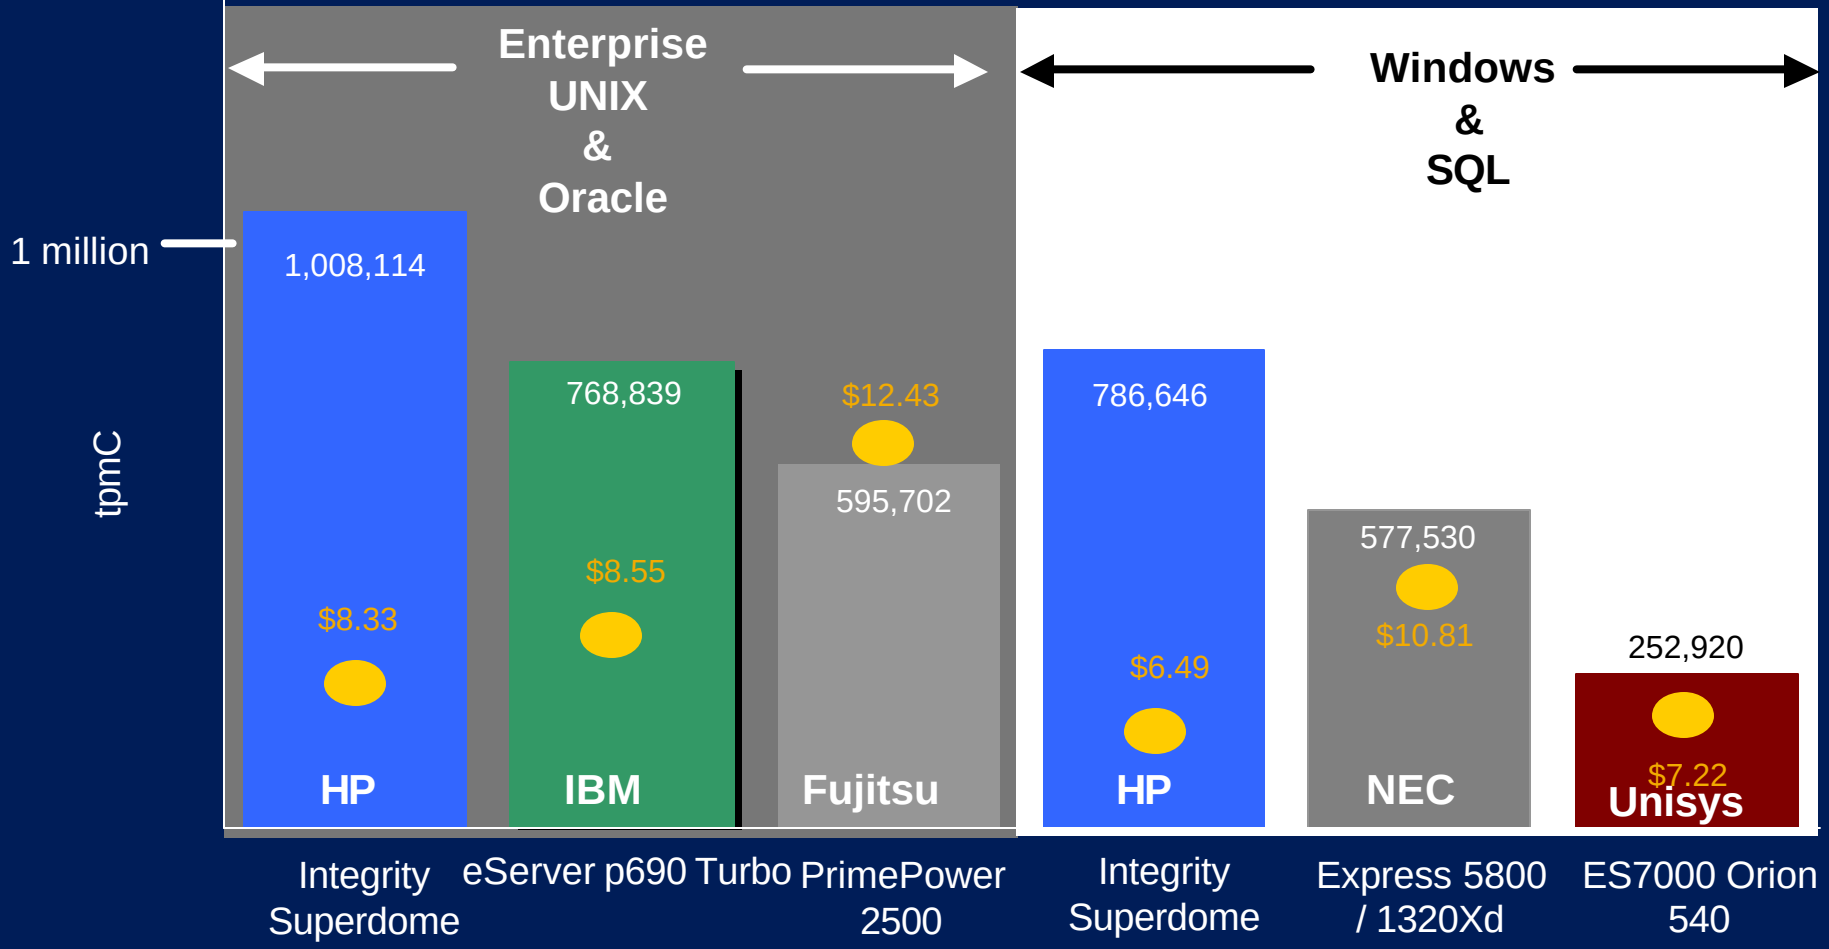

System Model	Operating System/Database	tpmC	$/tpmC	Available	System Model	Operating System/Database	tpmC	$/tpmC	Available
hp Integrity Superdome	HP-UX 11i v2 Oracle 10g	1,008,114	$8.3	3 04/14/04	Fujitsu PrimePOWER2500	Sun Solaris 8 Oracle 10g Enterprise Edition	595,702	$12.43	4/30/04
hp Integrity Superdome	Microsoft Windows Server 2003 Datacenter Edition Microsoft SQL Server 2000 Enterprise Ed. 64bit	786,646	$6.49	10/23/03	NEC Express5800 /1320Xd	Microsoft Windows Server 2003 Datacenter Edition Microsoft SQL server 2000 Enterprise Edition	577,530	$10.81	12/01/03
IBM eServer p690	IBM AIX 5L V5.2, IBM DB2 UDB 8.1 IBM AIX 5L V5.2, Oracle 10g	763,898 768,839	$8.25 $8.55	11/08/03 2/29/04	Unisys ES7000 Orion 540	Microsoft Windows Server 2003 Data Center Edition Microsoft SQL Server 2000 Enterprise Edition	252,920	$7.22	7/22/03

Results as of January 15, 2004
For more information, see www.tpc.org

HP Partitioning

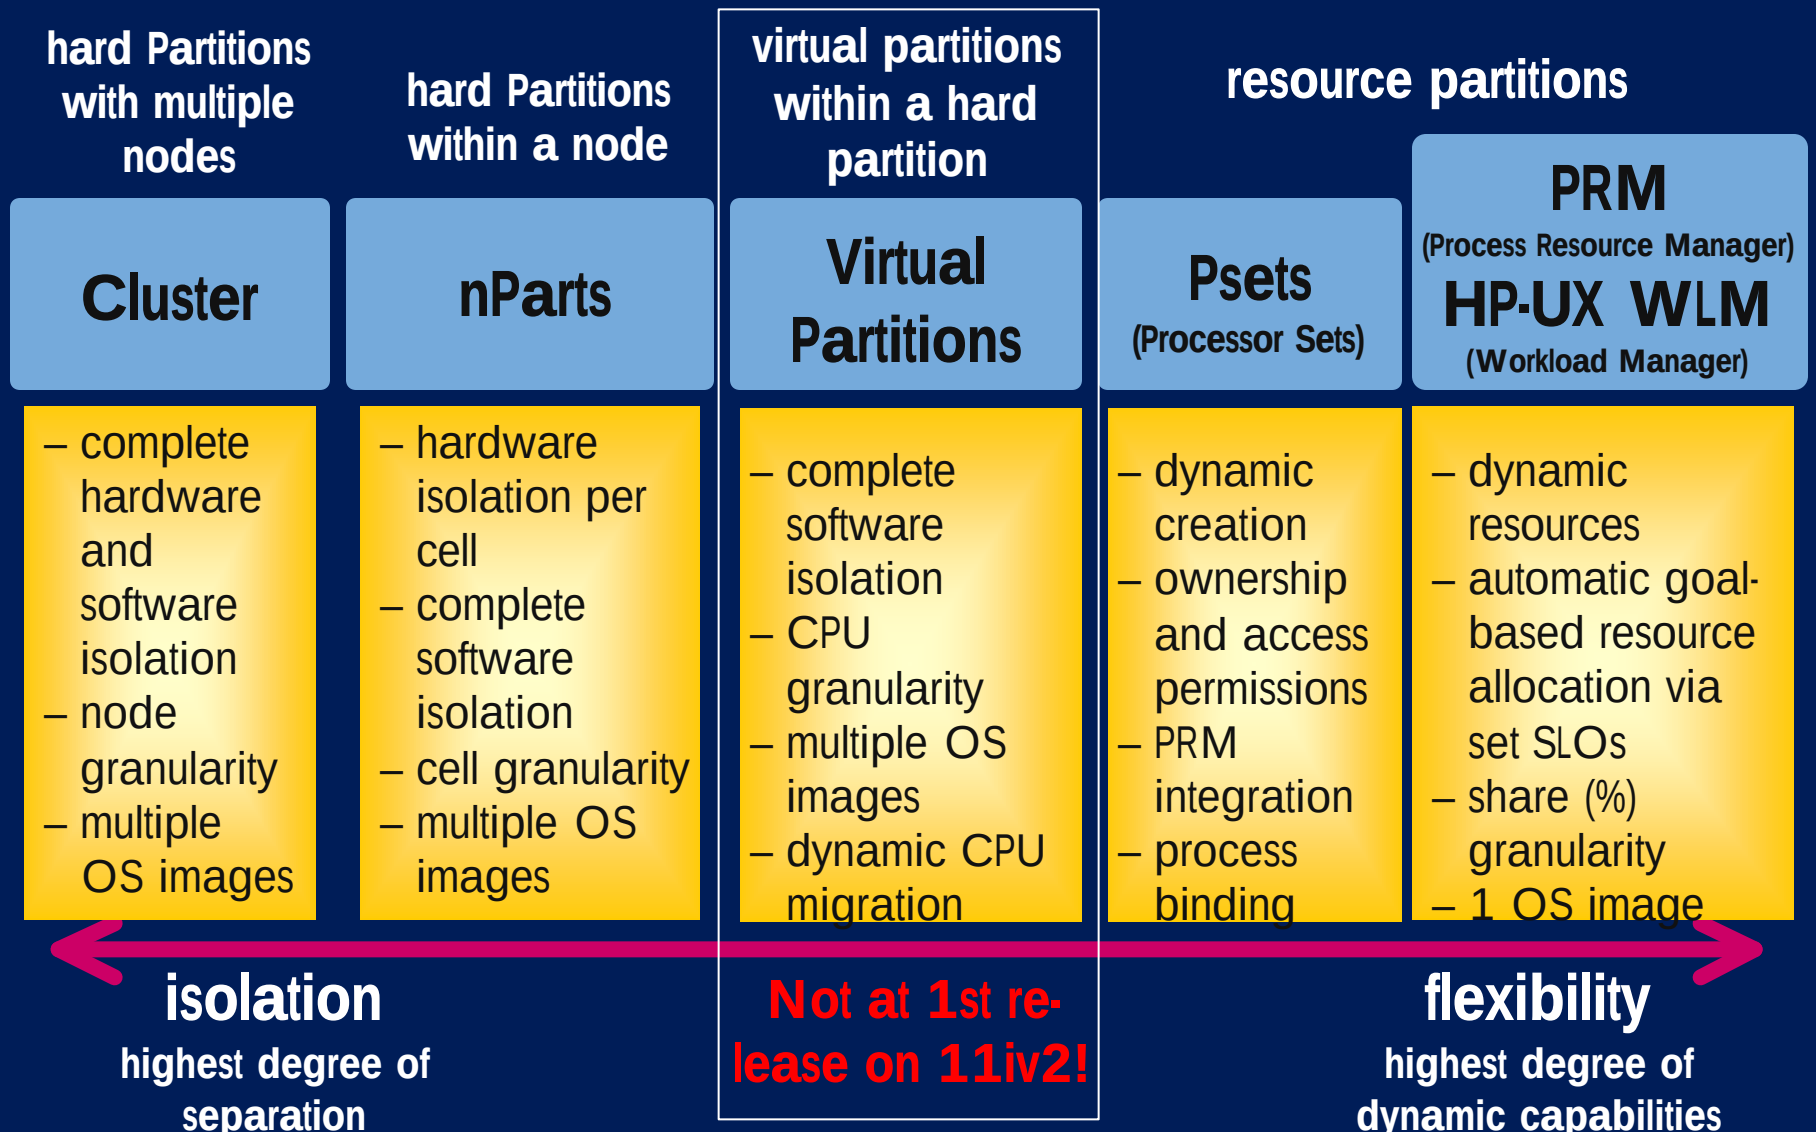

Question and Answers


hp
invent