HP Integrity Server

DECUS Herbsttreffen 2004
15.11.2004

Jörg Demmler
Technologie Consultant
joerg.demmler@hp.com

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

Agenda

- Integrity Server – Why?
  - IT Requirements & Trends
  - HP Adaptive Enterprise
  - HP Industry Based Server Strategy
- Itanium® Processor Technologies and Roadmap
- Itanium® Processor chipsets & Module
- HP Integrity Server
- Server consolidation

Integrity Server – Why?
IT-Requirements and Trends

Reduce the Cost and increase the value of IT

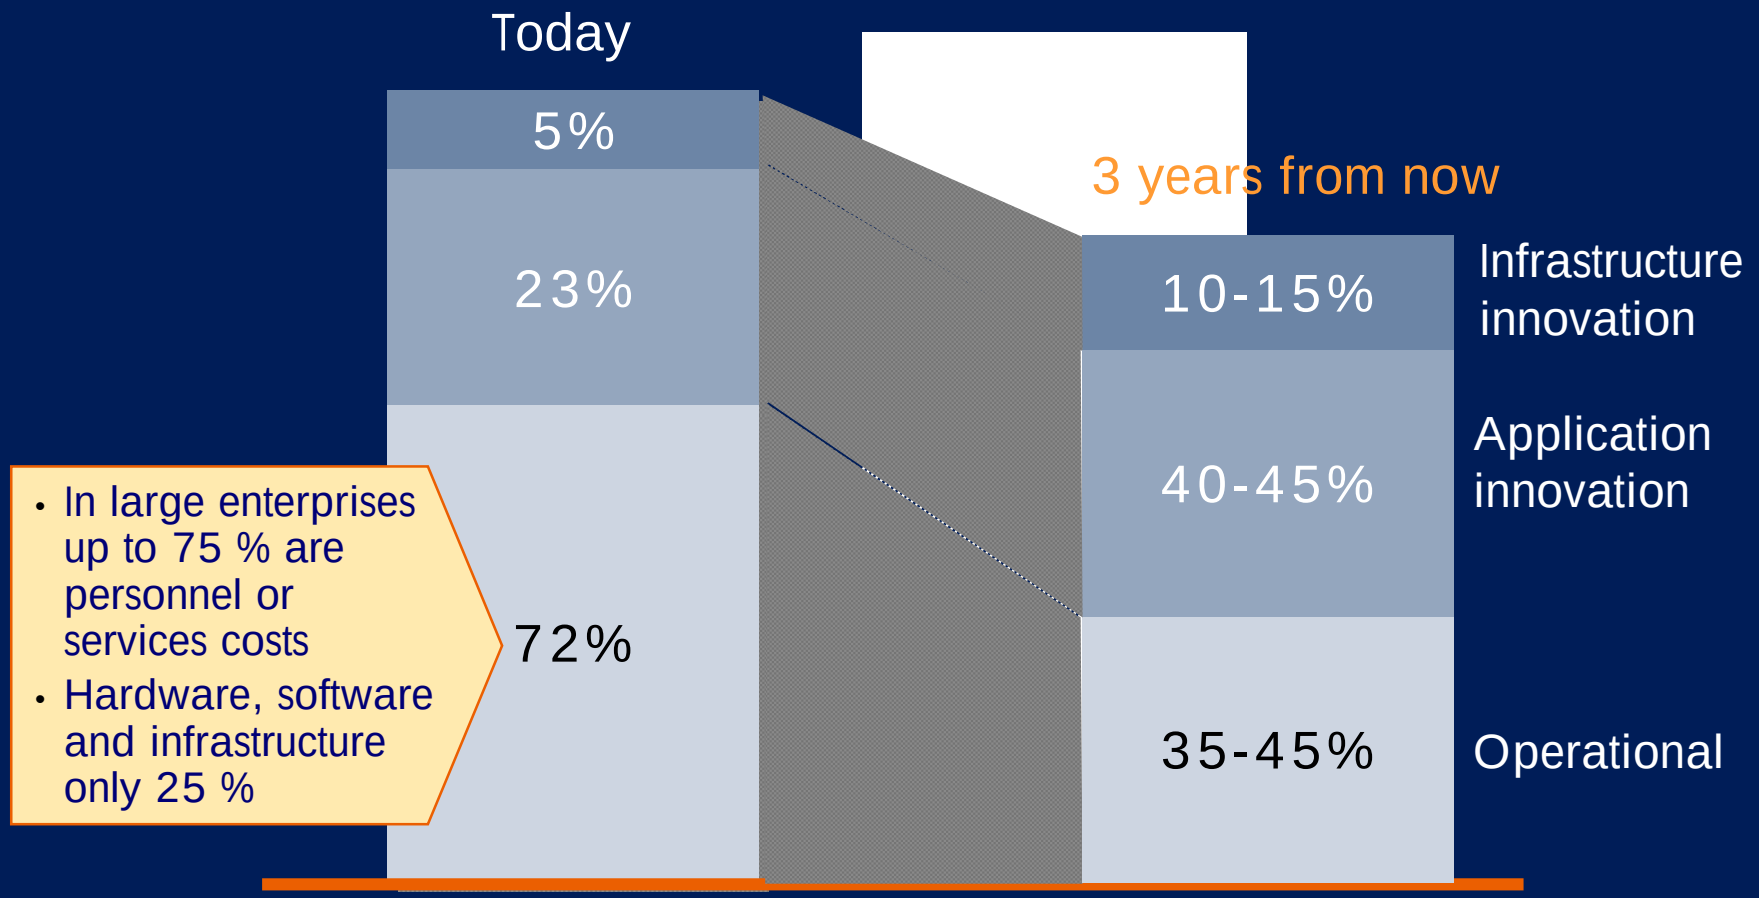

How can we move to a sustainable cost structure?

Integrity Server – Why?
IT-Requirements and Trends

Integrity Server – Why?
Adaptive Enterprise - New Model of Computing

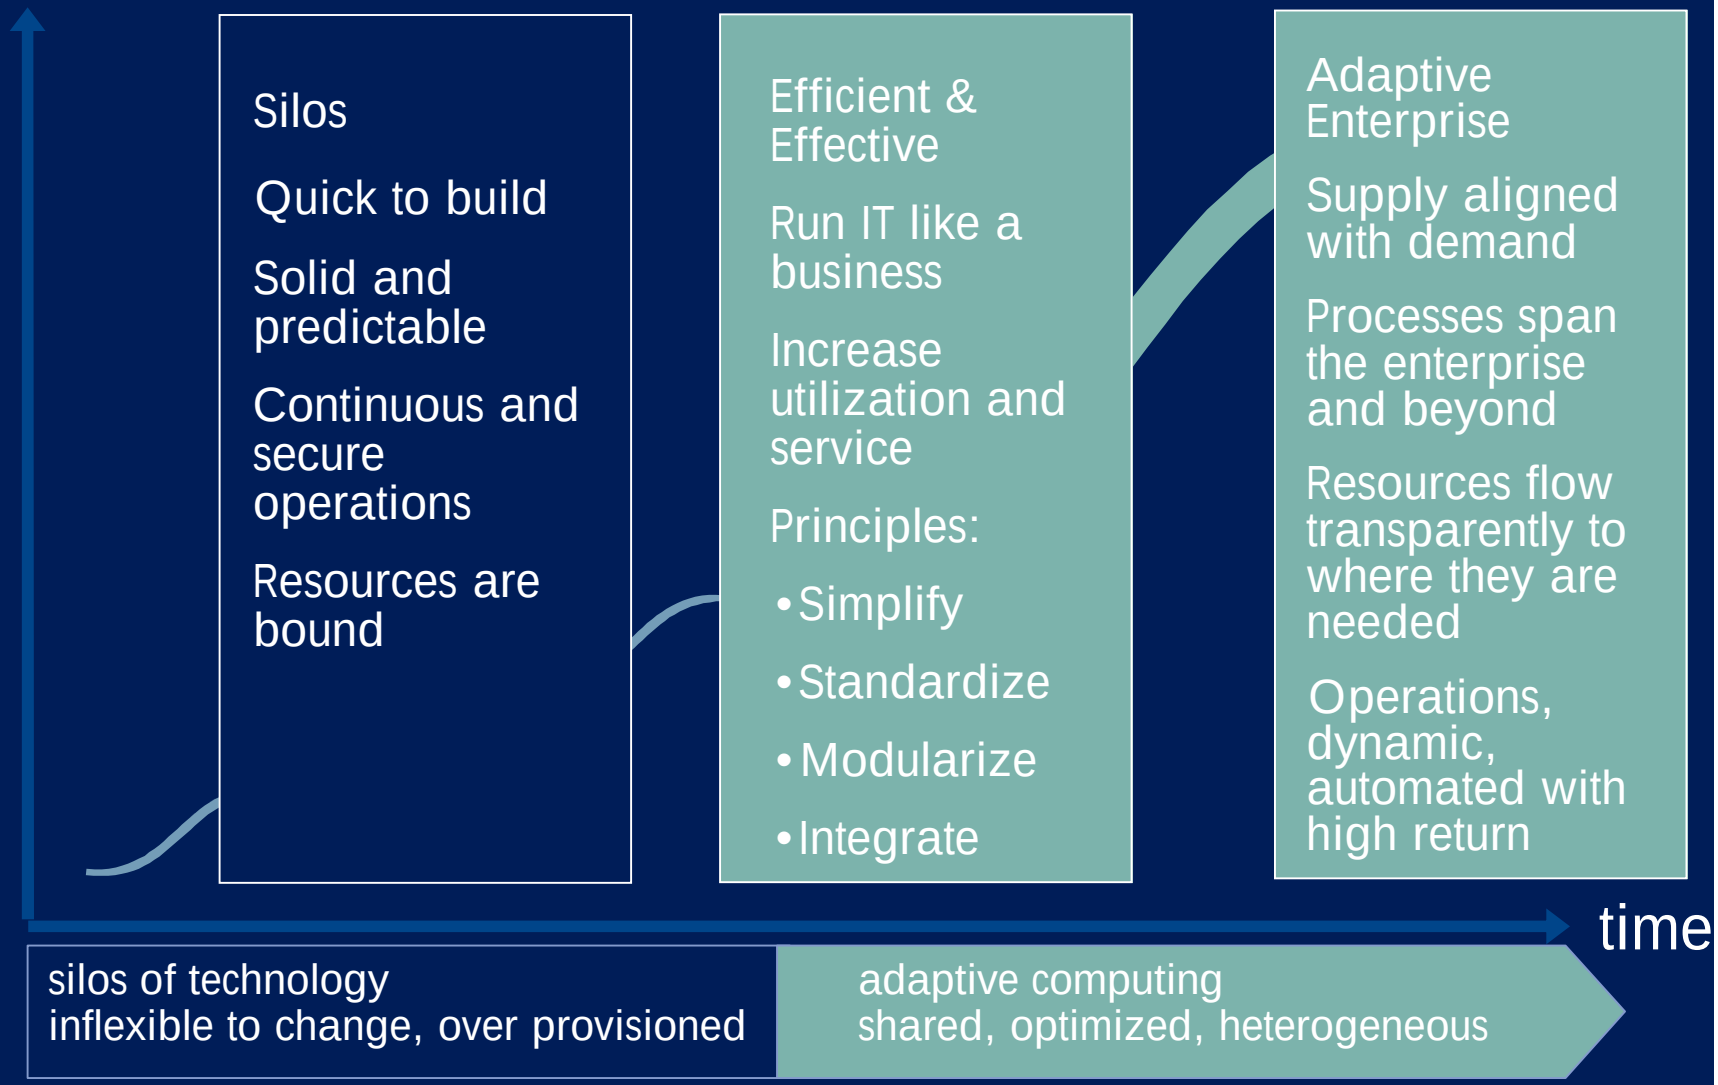

Integrity Server – Why?
HP Industry based Server Strategy

Moving to 3 leadership product lines – built on 2 industry standard architectures

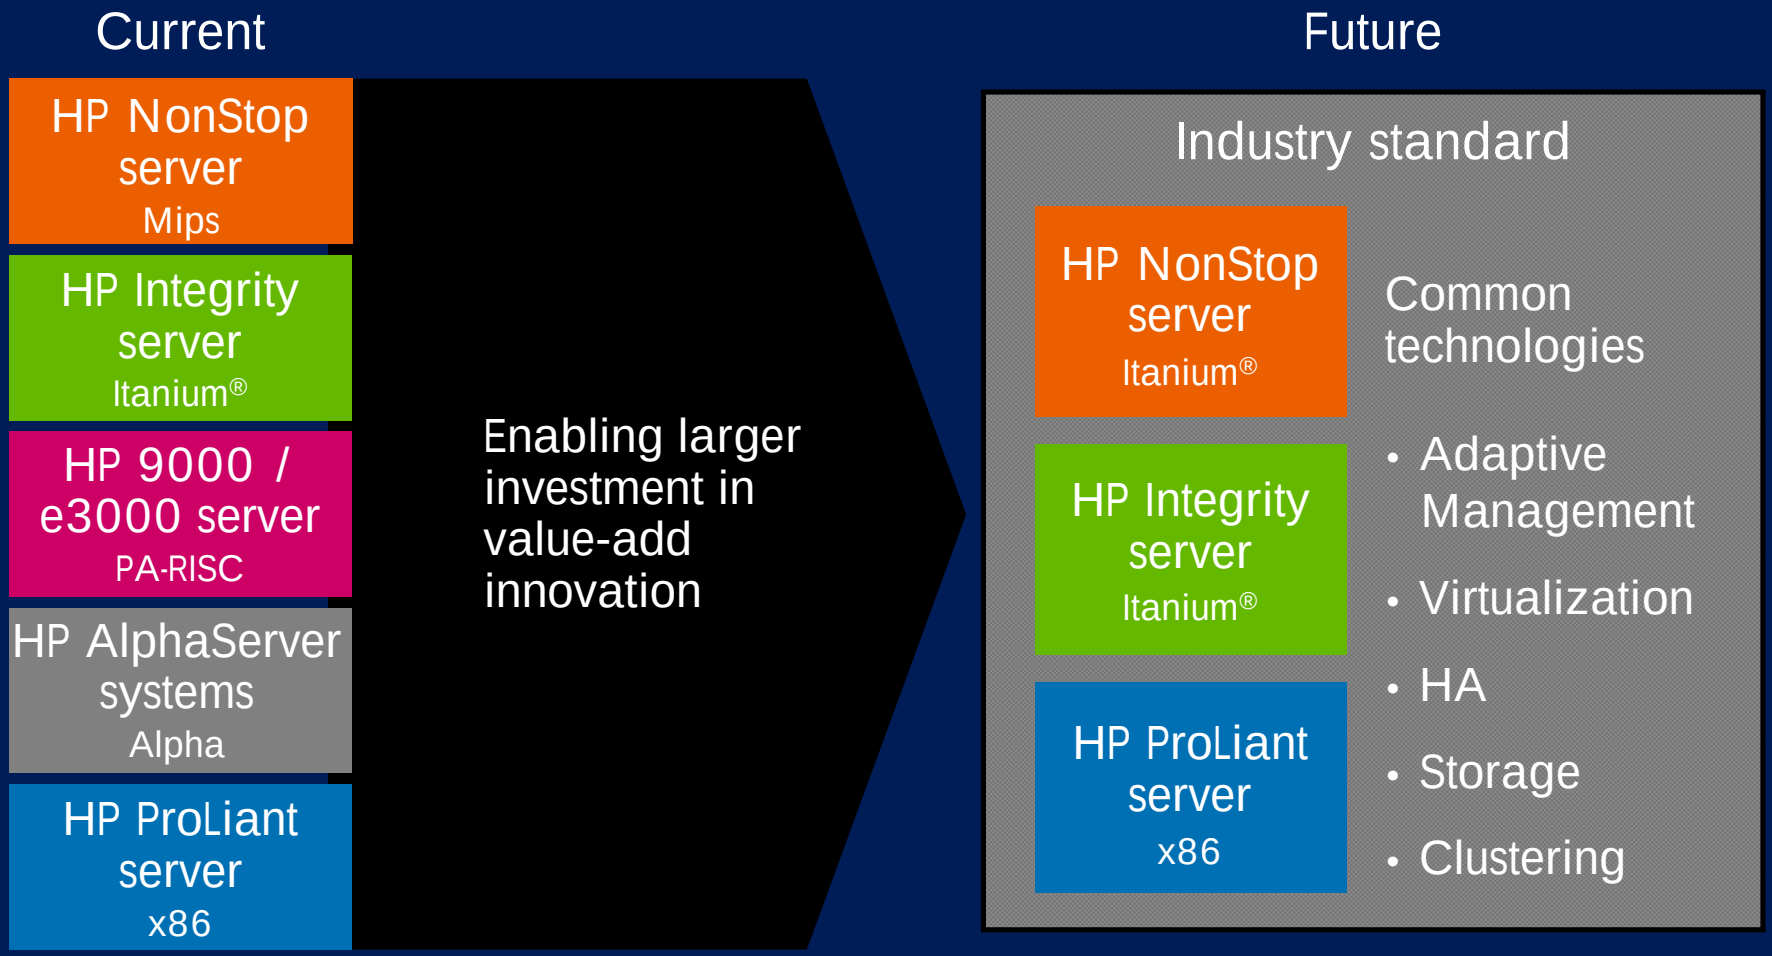

Integrity Server – Why?
HP Industry Based Server Strategy - Positioning

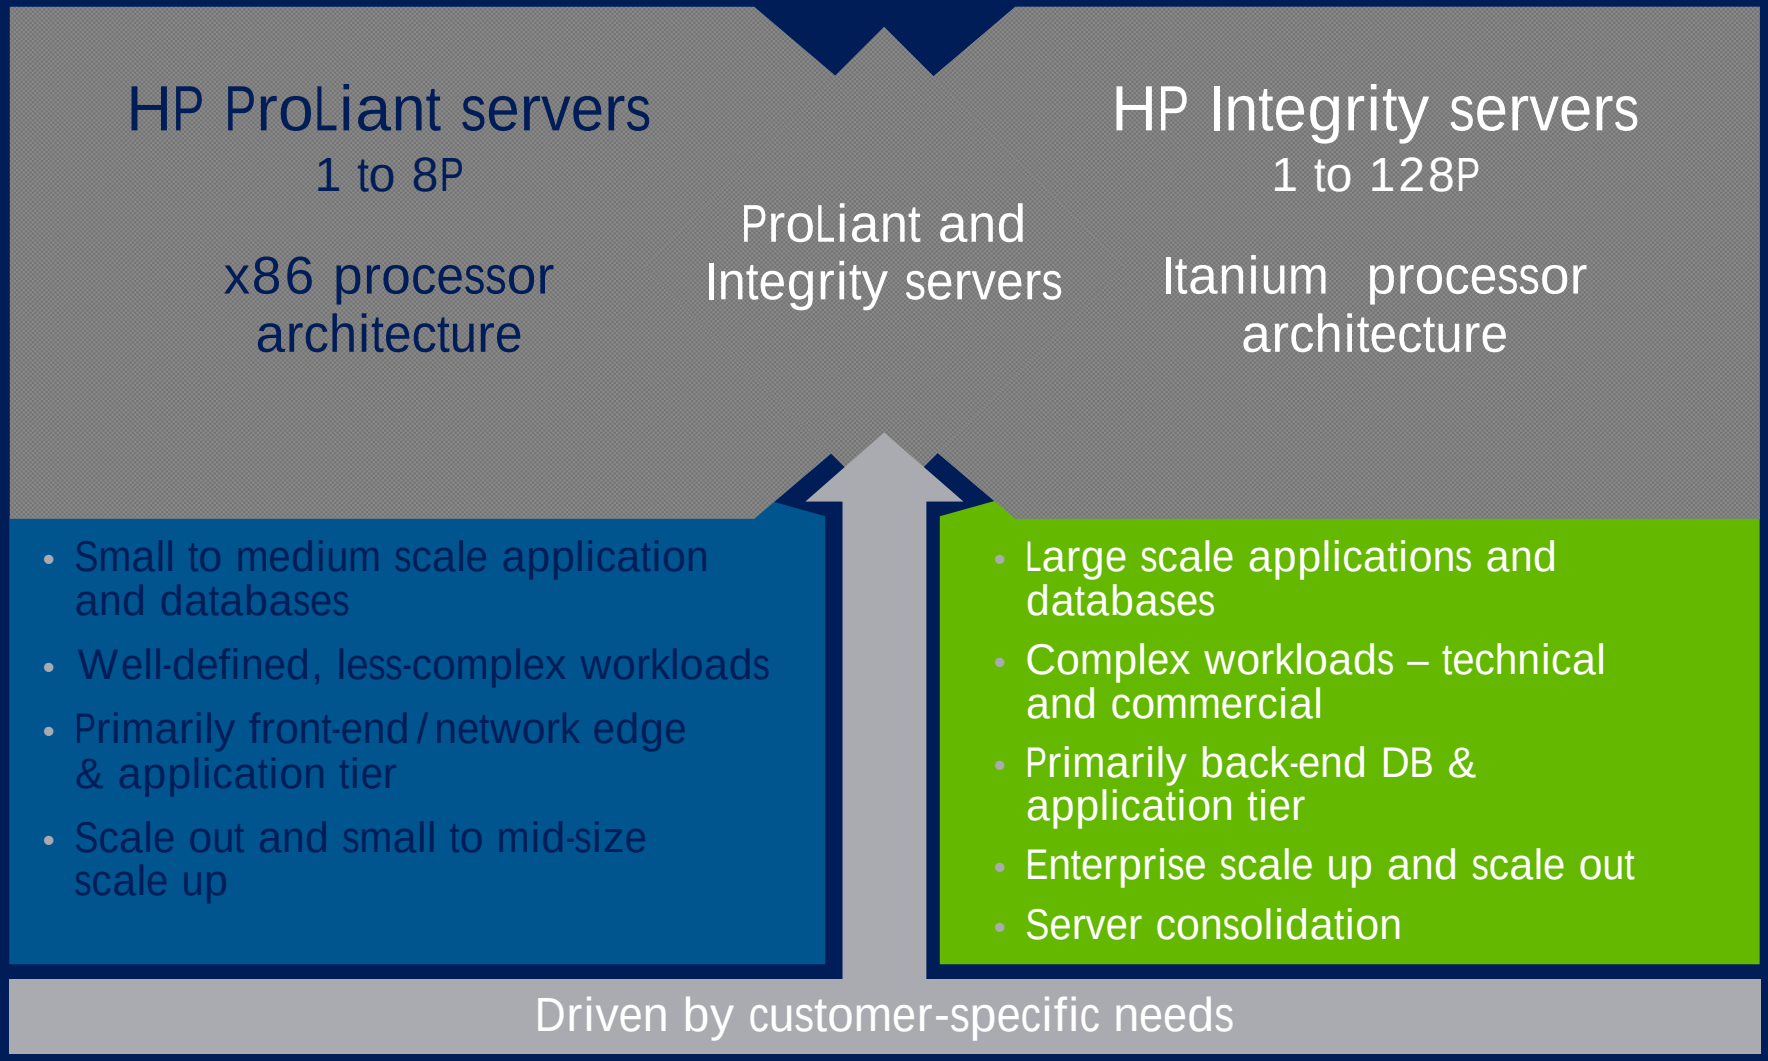

Agenda

- Integrity Server – Why?
  - IT Requirements & Trends
  - HP Adaptive Enterprise
  - HP Industry Based Server Strategy
- Itanium® Processor Technologies and Roadmap
- Itanium® Processor chipsets & Module
- HP Integrity Server
- Server consolidation

Trends bei Prozessor-Technologien

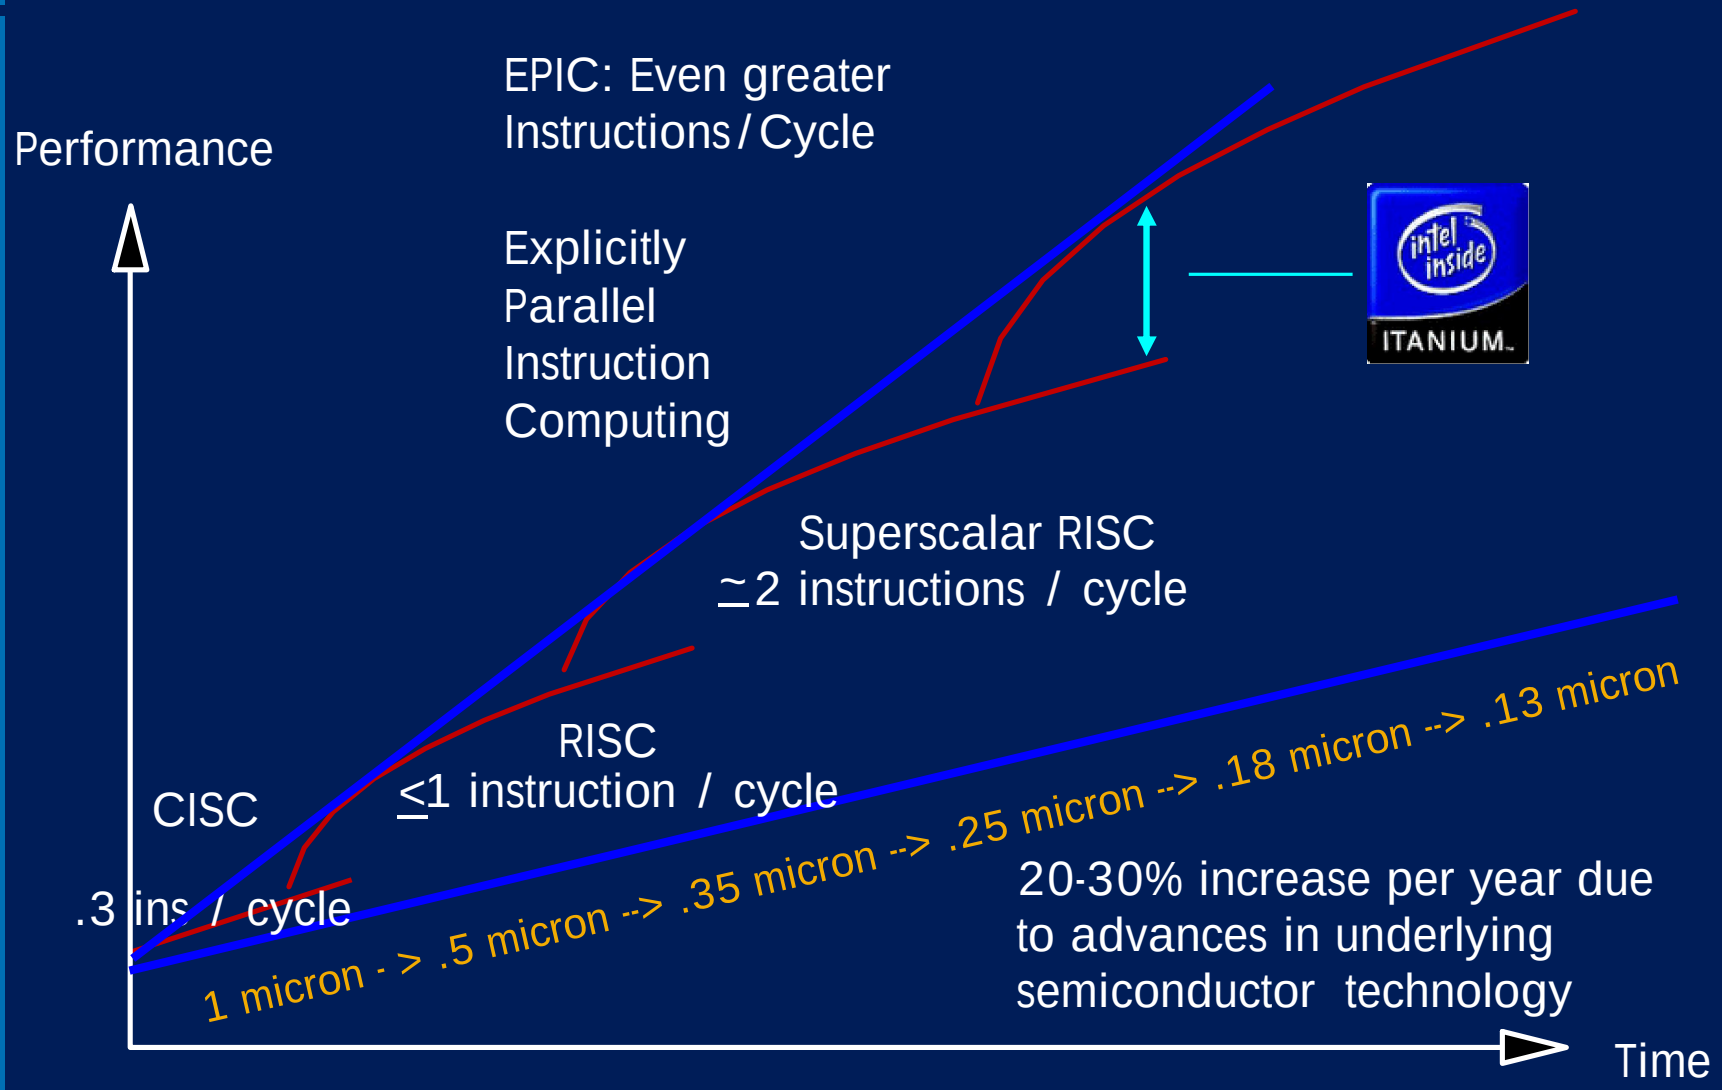

Itanium® Architektur: 4 Key Features

- **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

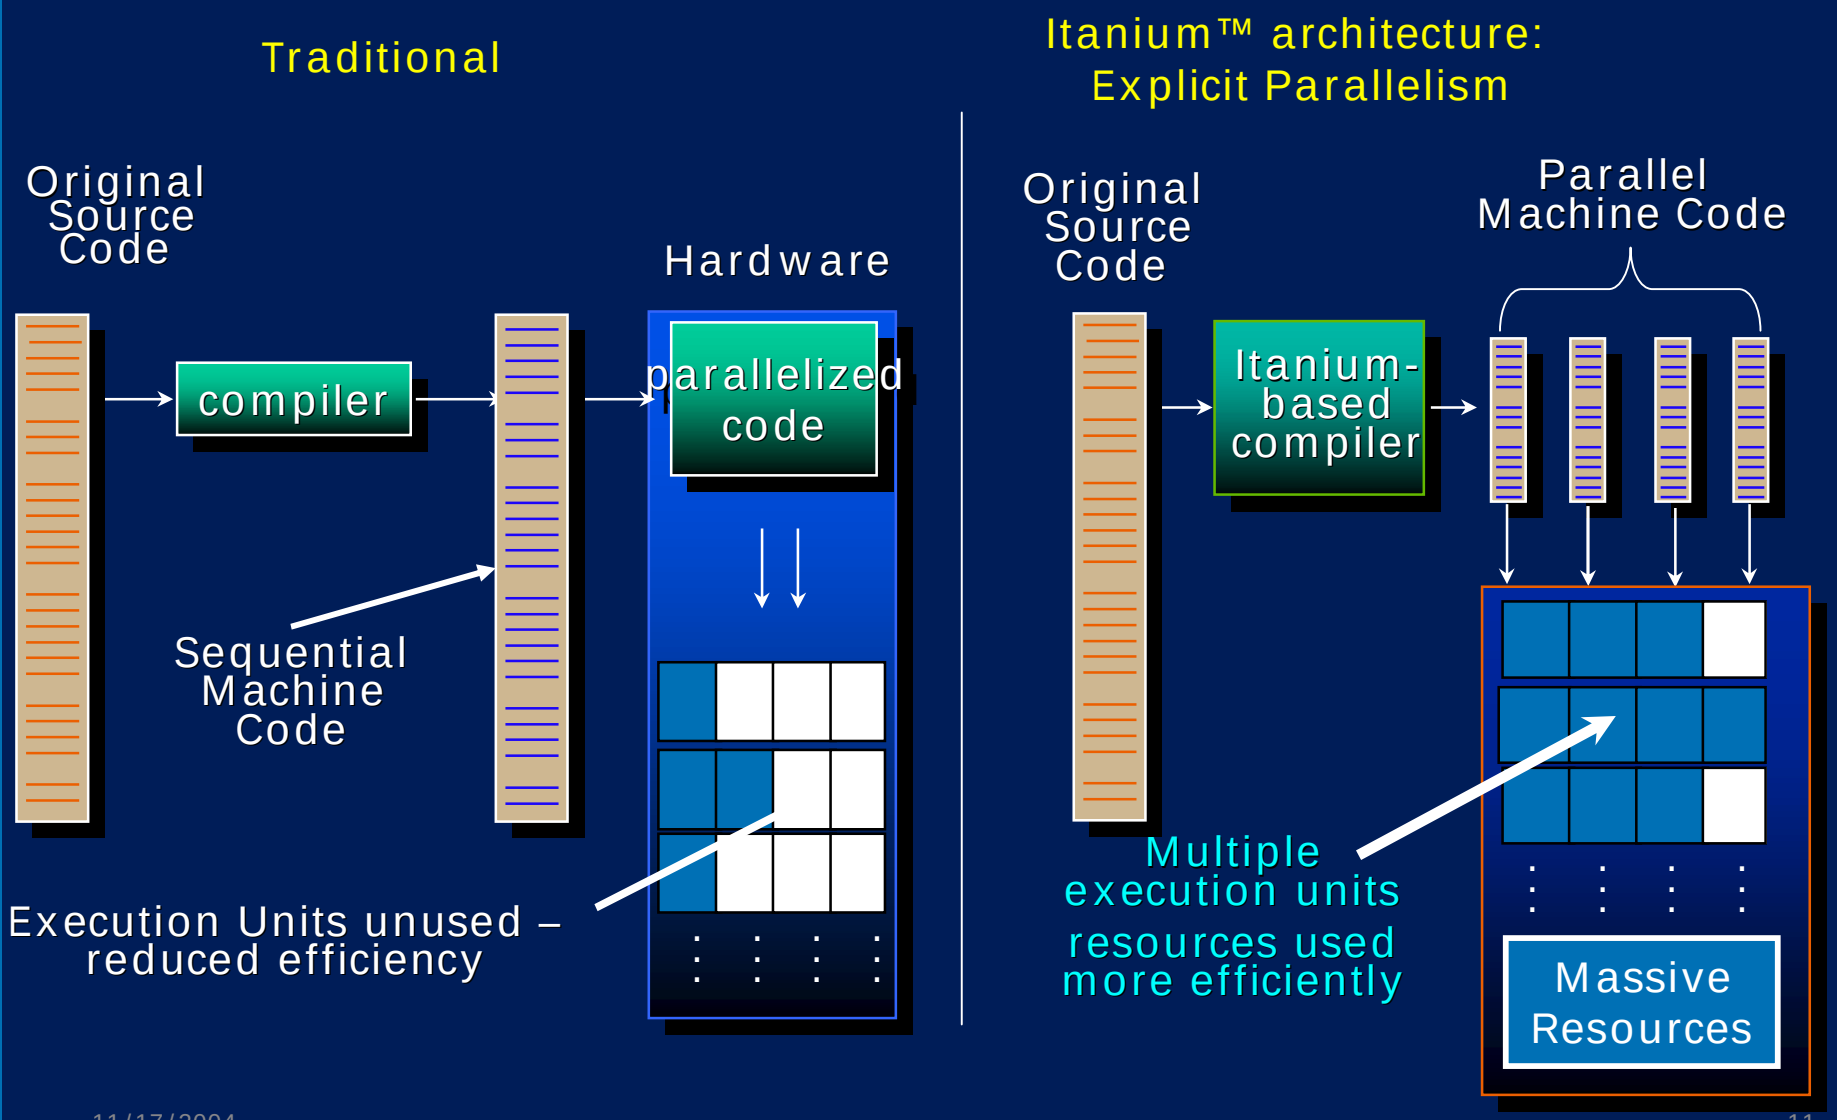

Very Large Instruction Word (VLIW)

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

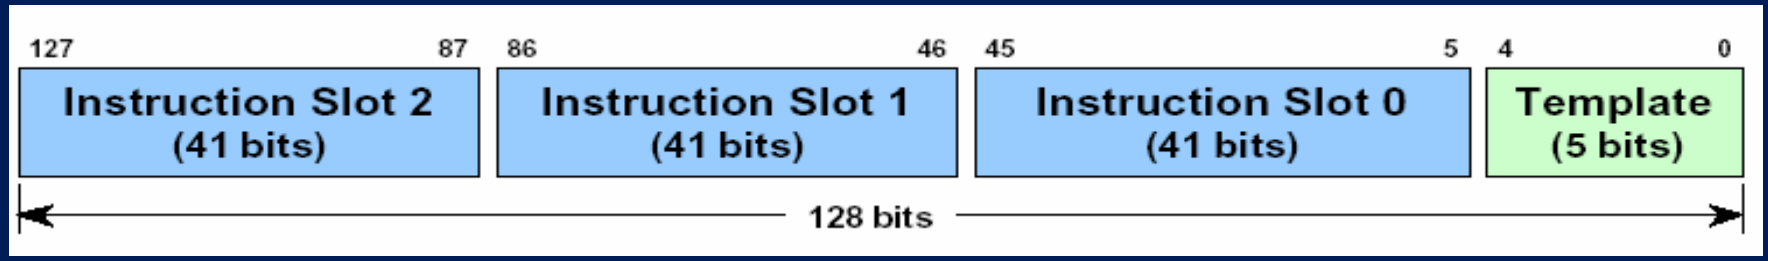

Instruction Handling

Program:

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

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

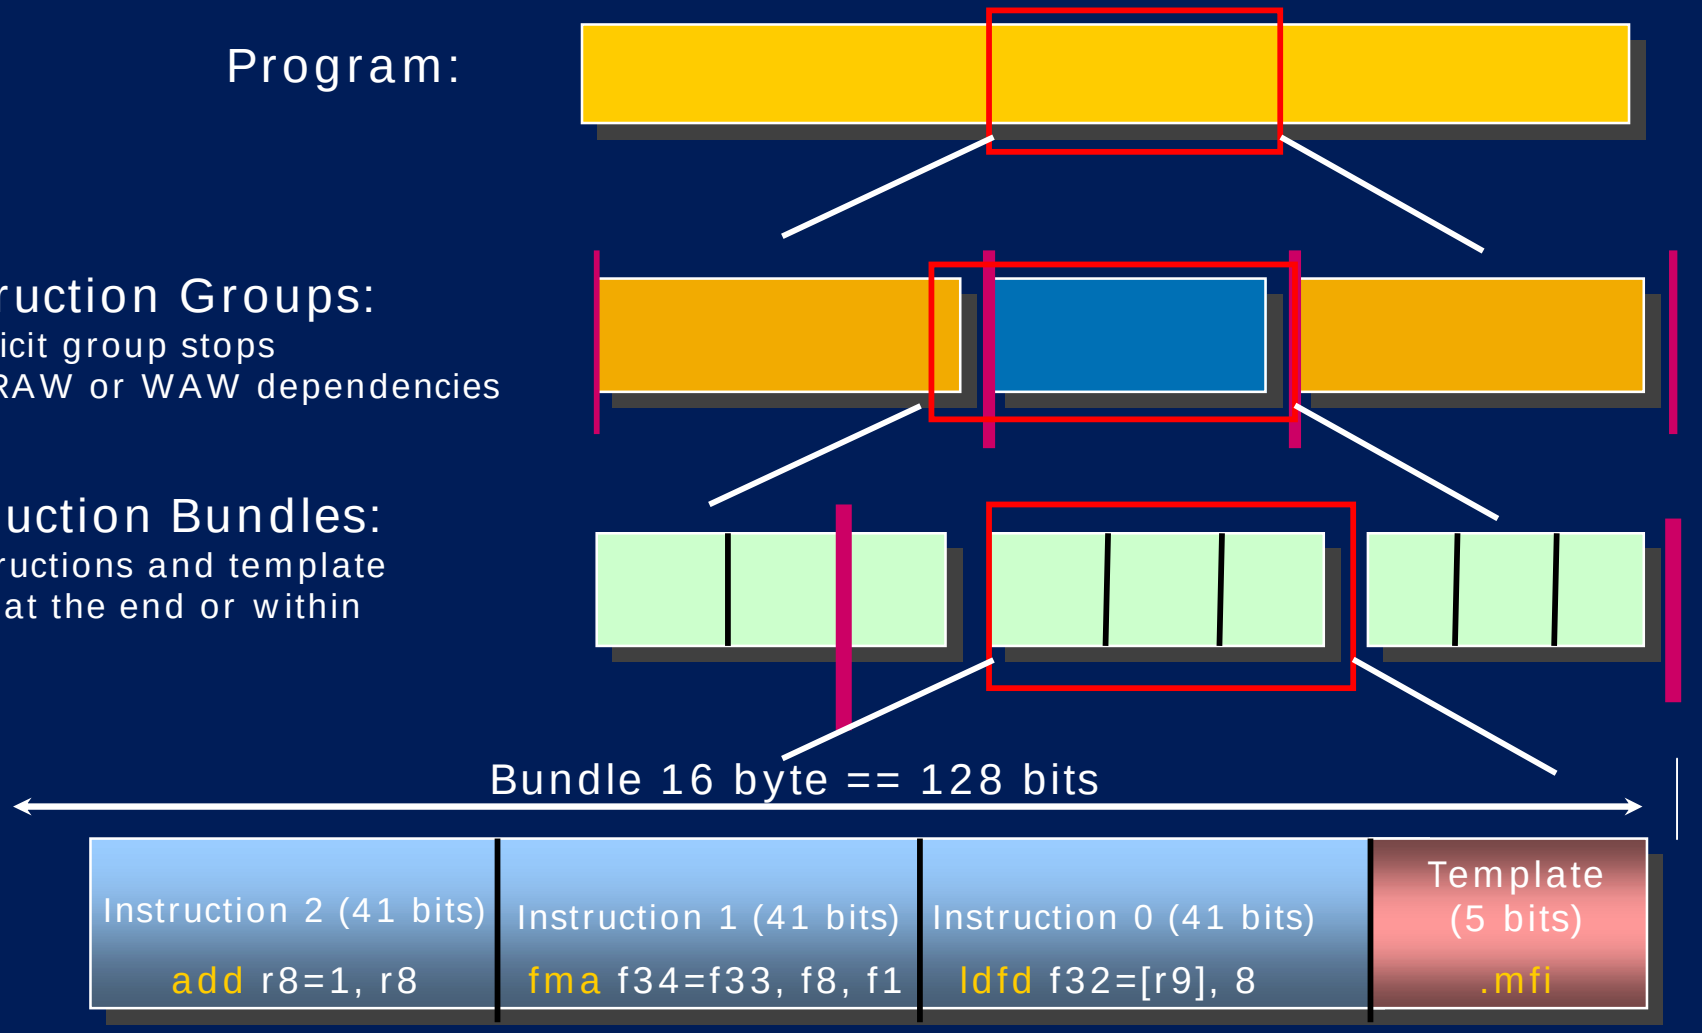

Instruction Handling

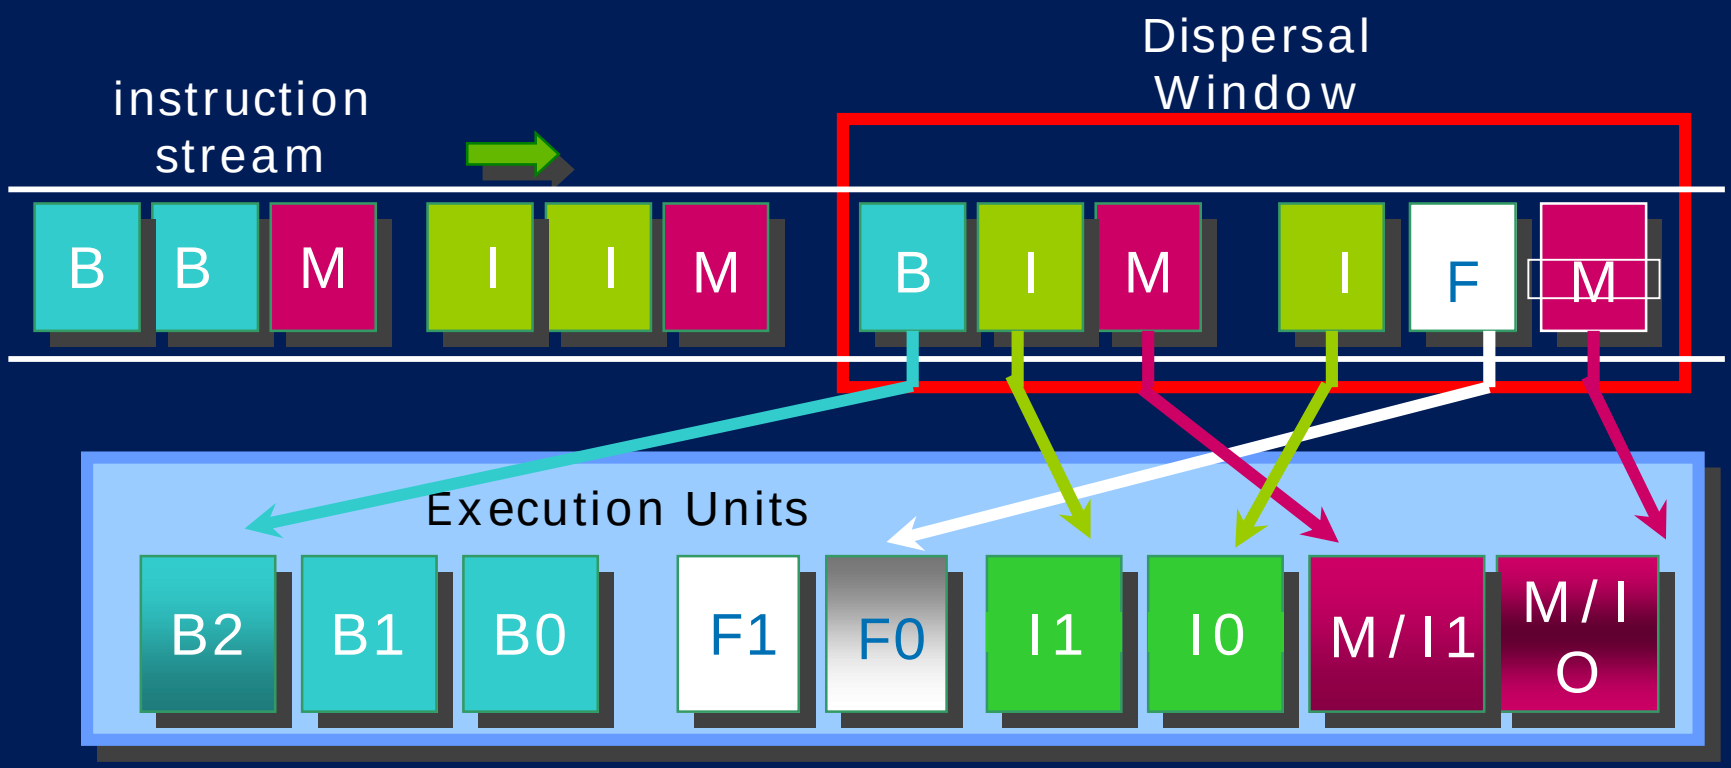

Flexible Issue Capability

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

Typical

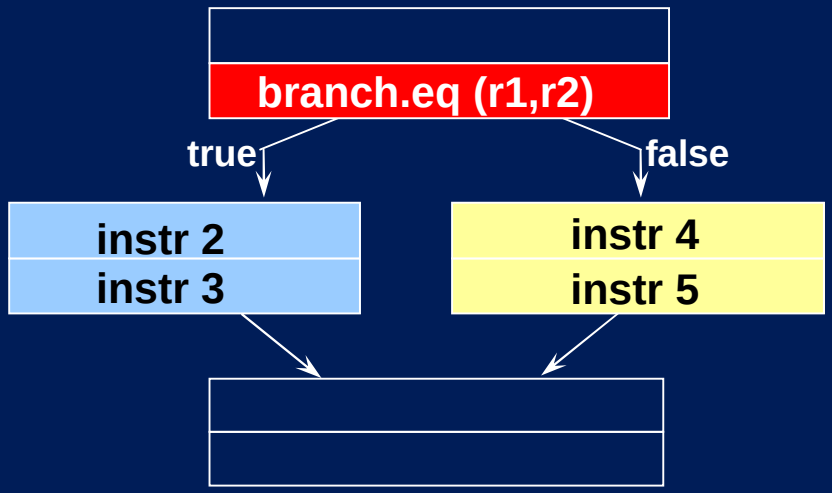

Optimized IPF

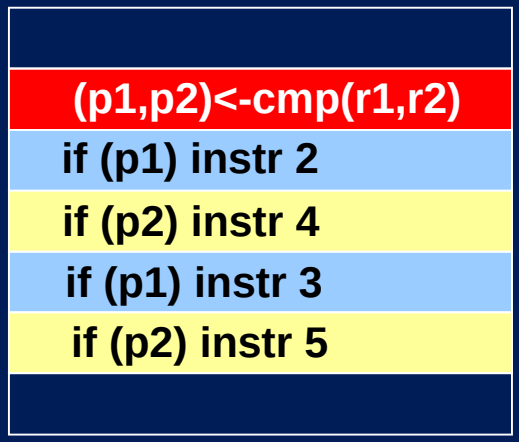

Data Speculation

- 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

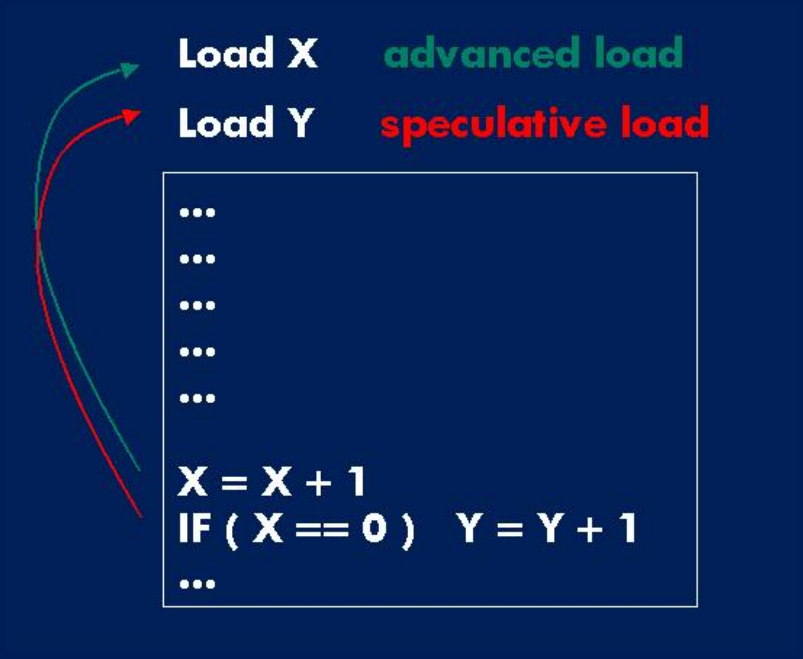

Itanium®2 9M processor
Performance leader

news 08.11.2004 20:02

<< Vorige | Nächste >>

Intels neuer Itanium-2-Prozessor um über 25 Prozent schneller

Nicht so sehr der geringfügig höhere Takt von 1,6 GHz -- gegenüber 1,5 GHz zuvor --, sondern vor allem der um 50 Prozent auf nunmehr maximal 9 MByte vergrößerte L3-Cache beschleunigt den Itanium-2-Prozessor (Madison) in den meisten Applikationen um 25 Prozent und mehr. Mit einem SPECfp-base2000-Wert von 2712 setzt er sich im Number-Crunching-Wettbewerb klar vor dem derzeit schnellsten IBM Power5 (2576) an die Spitze. In der Integer-Performance (SPECint_base2000) muss er sich mit 1590 zwar weiterhin den hauseigenen Workstation/Desktop- und Spiele-Prozessoren und auch der Athlon-FX/Opteron-Konkurrenz unterordnen, aber bei den größeren multiprozessortauglichen Systemen liegt er damit ebenfalls in Front. Hauptkonkurrent Power5 erzielt hier 1398.

Agenda

- IT Requirements and Trends
  - HP Adaptive Enterprise
  - HP Industry Based Server Strategy
- Itanium® Processor Technologies and Roadmap
- Itanium® Processor chipsets & Module
- HP Integrity Server
- Server consolidation

Itanium® Processor chipsets & Module
zx1 & sx1000 chipset

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

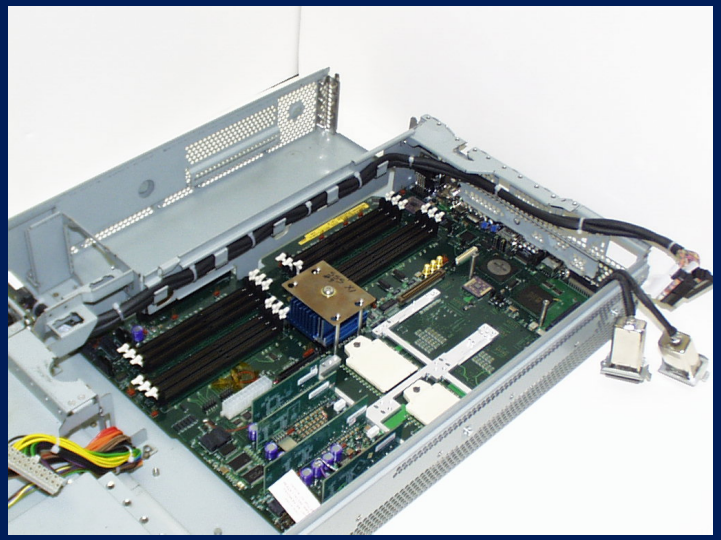

The zx1 chipset

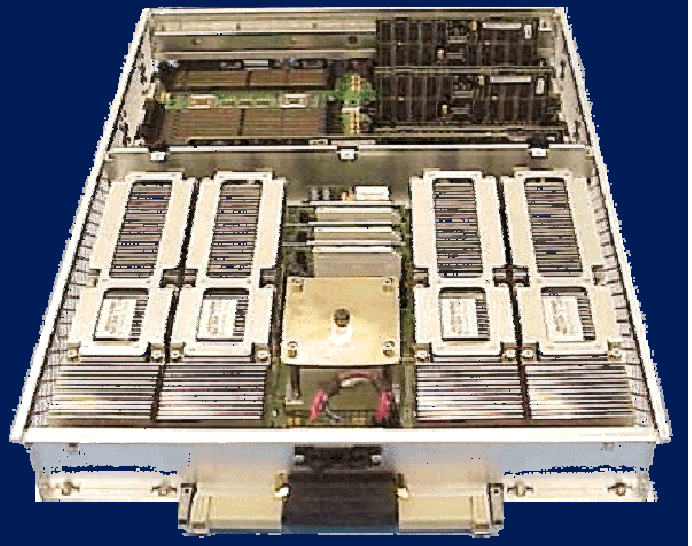

- For medium-sized and big servers

The sx1000 chipset

Itanium® Processor chipsets & Module
HP zx1 chipset

The HP zx1 chipset contains three components :

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

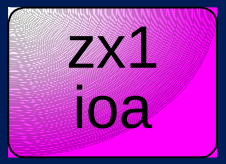

- zx1 I/O adapter
  - PCI and PCI-X
  - AGP

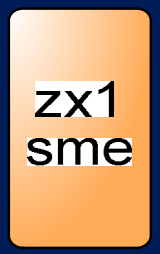

- zx1 scalable memory expander
  - optional component used to :
    - increases memory capacity
    - increases memory bandwidth

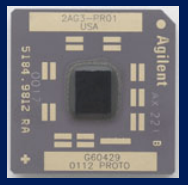

Itanium® Processor chipsets & Module
HP zx1 chipset (4-CPU configuration)

<u>System Bus</u>
128 bits wide
6.4 GB/s

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

<u>I/O</u>
4 GB/s
PCI-X support

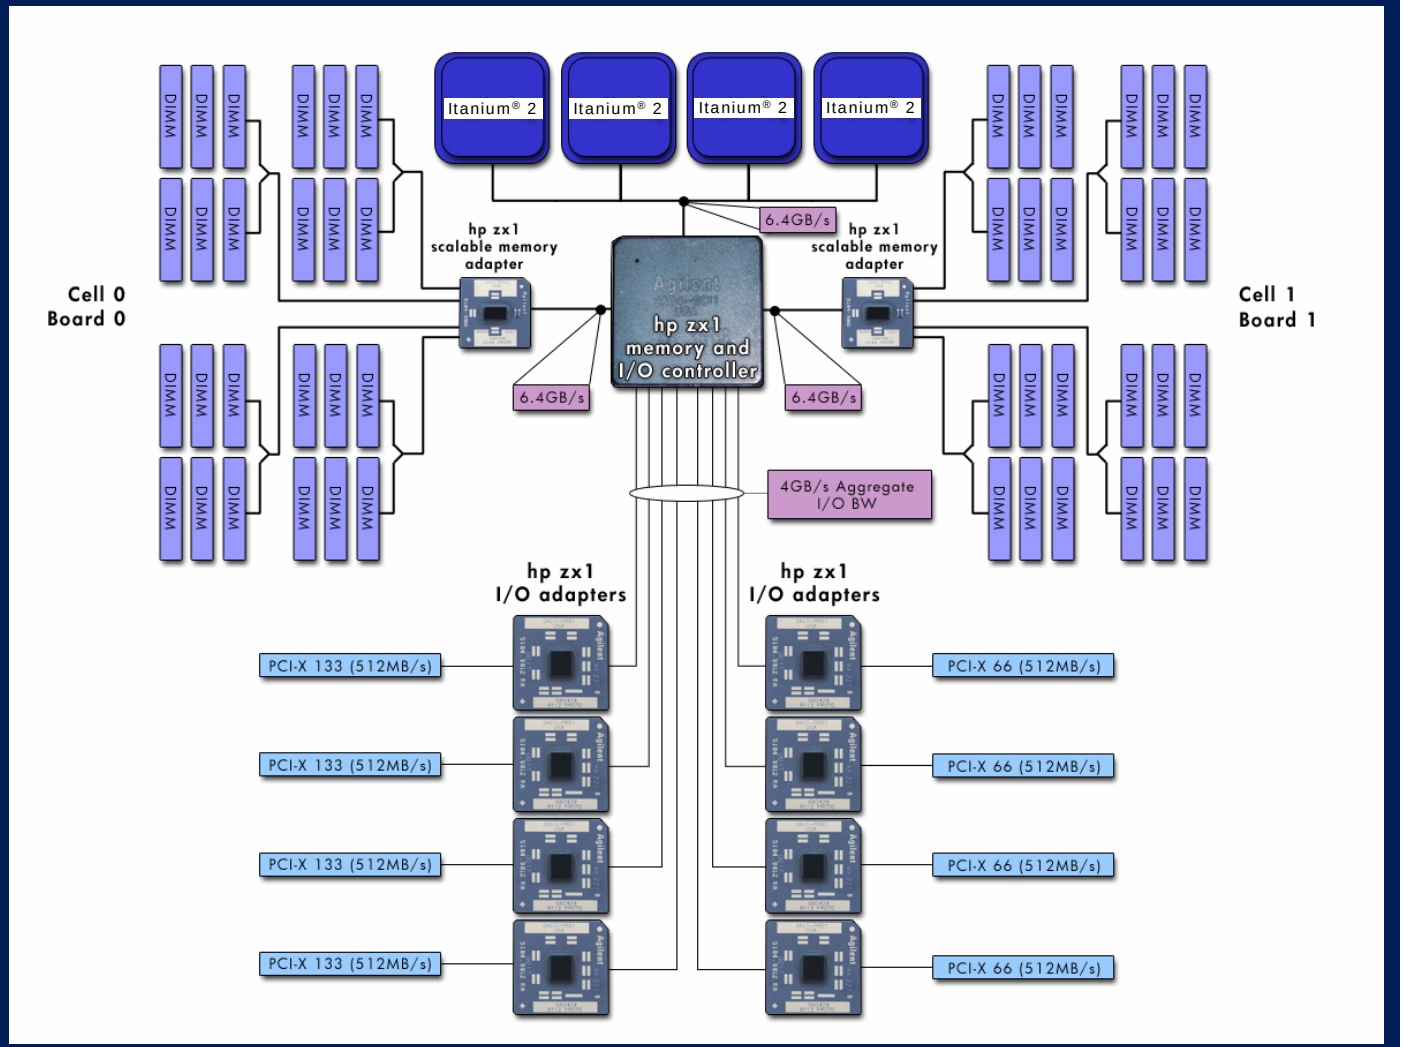

Itanium® Processor chipsets & Module
HP sx1000 chipset

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)

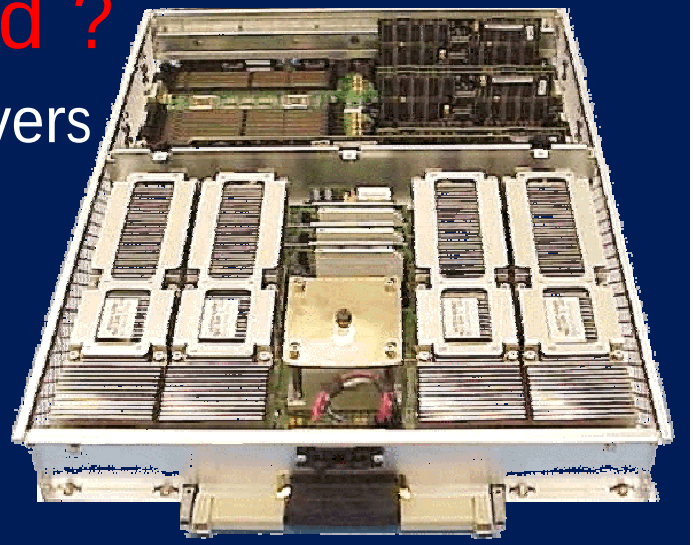

Itanium® Processor chipsets & Module sx1000 chipset

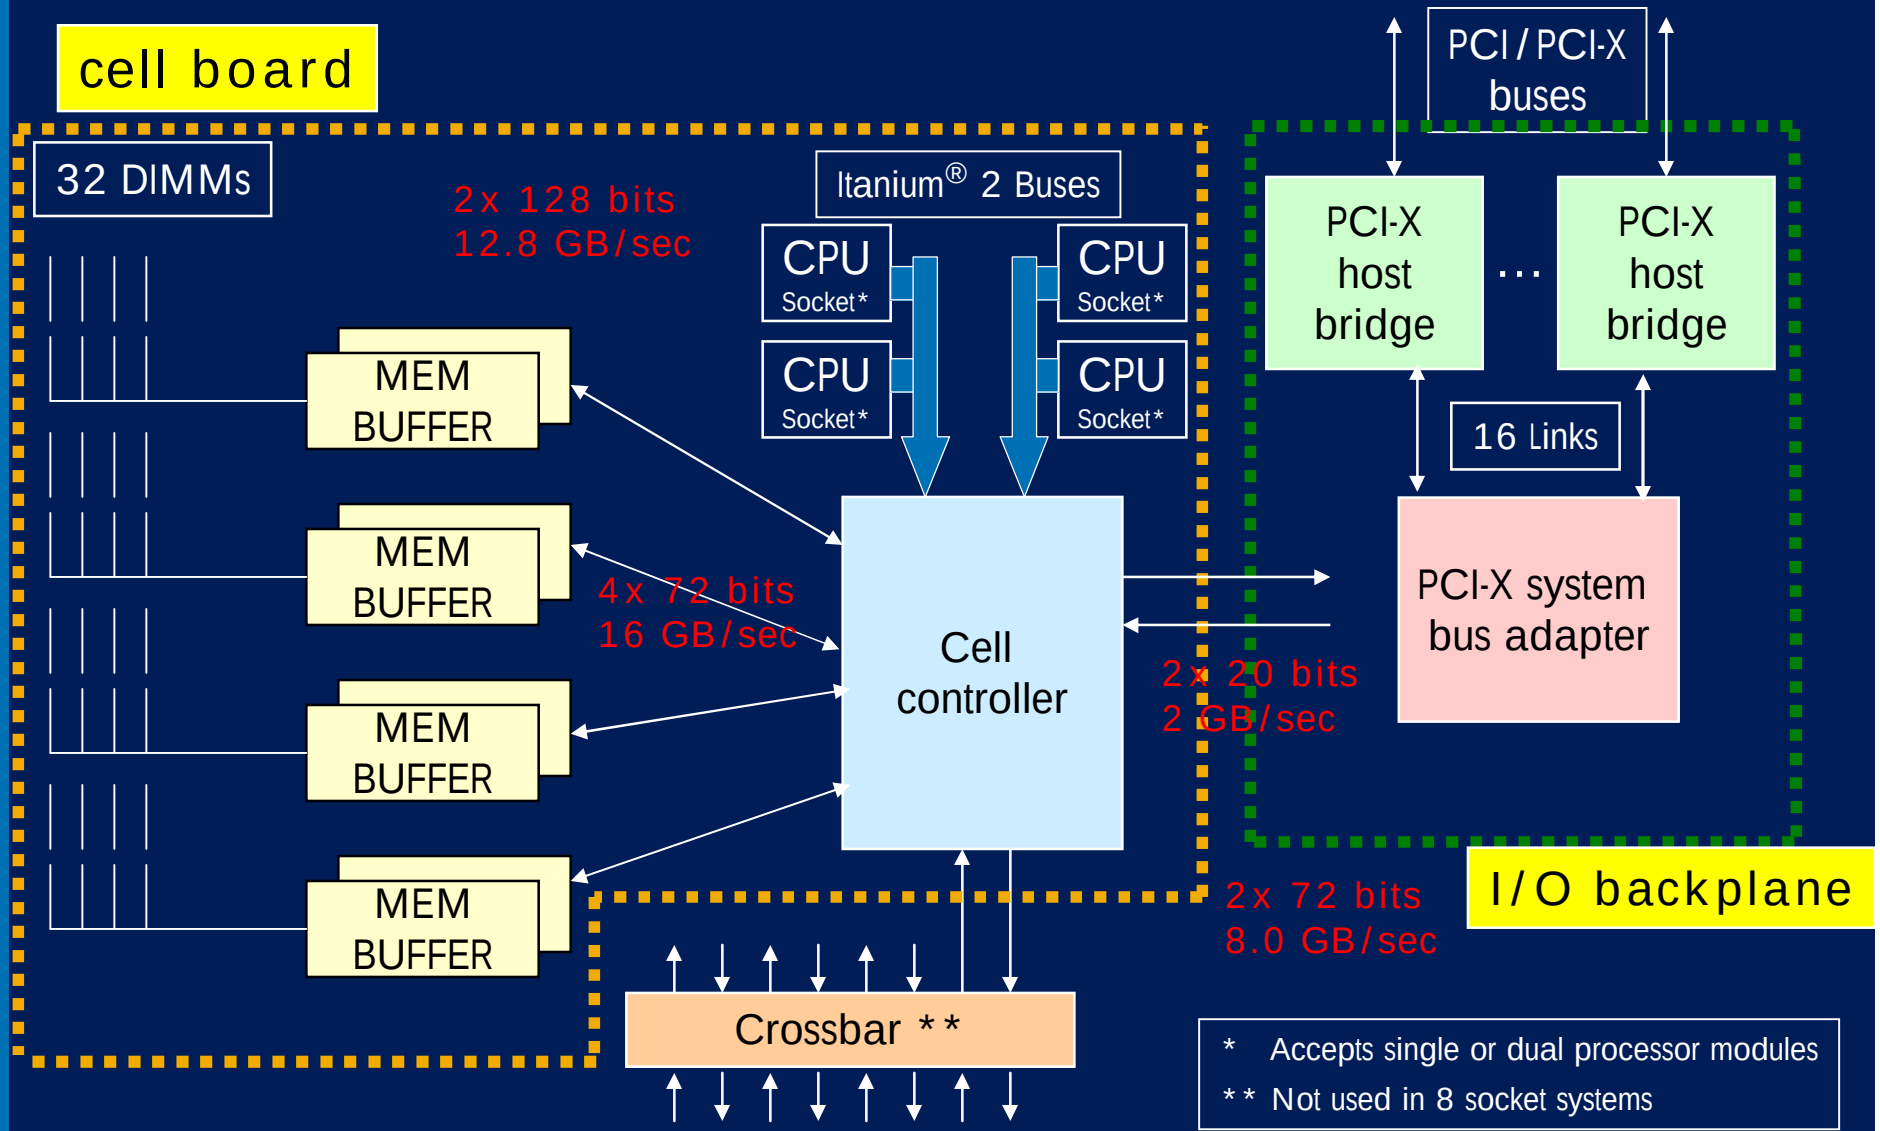

Itanium® Processor chipsets & Module
HP mx2 Dual Processor Module

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

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Itanium® Processor chipsets & Module
HP mx2 Dual Processor Module

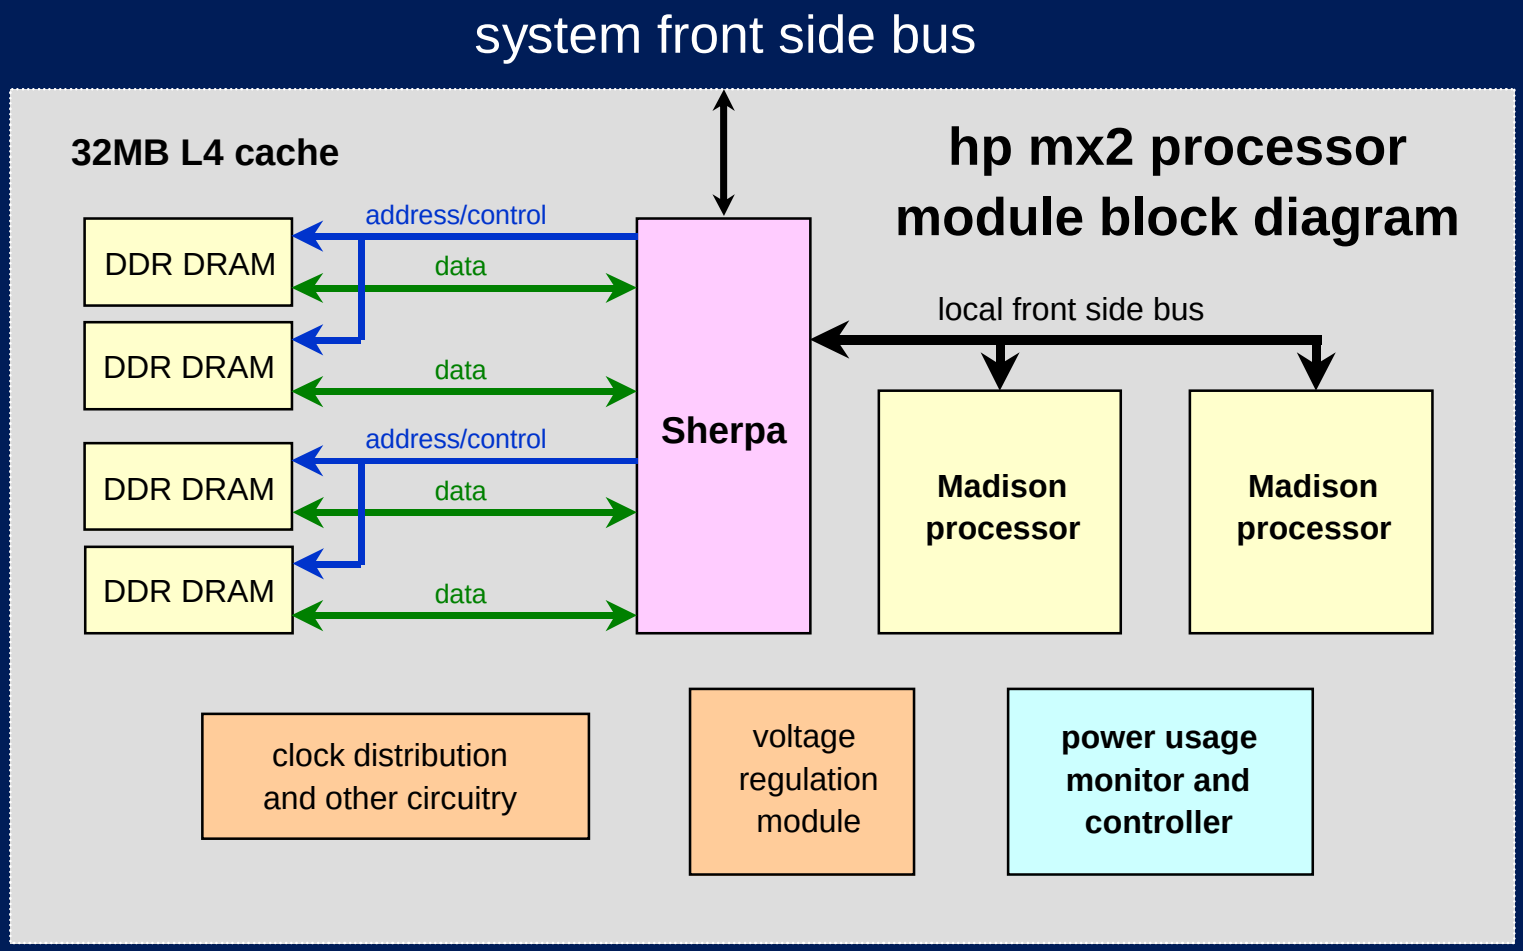

Itanium® Processor chipsets & Module
HP mx2 Gives Itanium 2 a Performance Kick

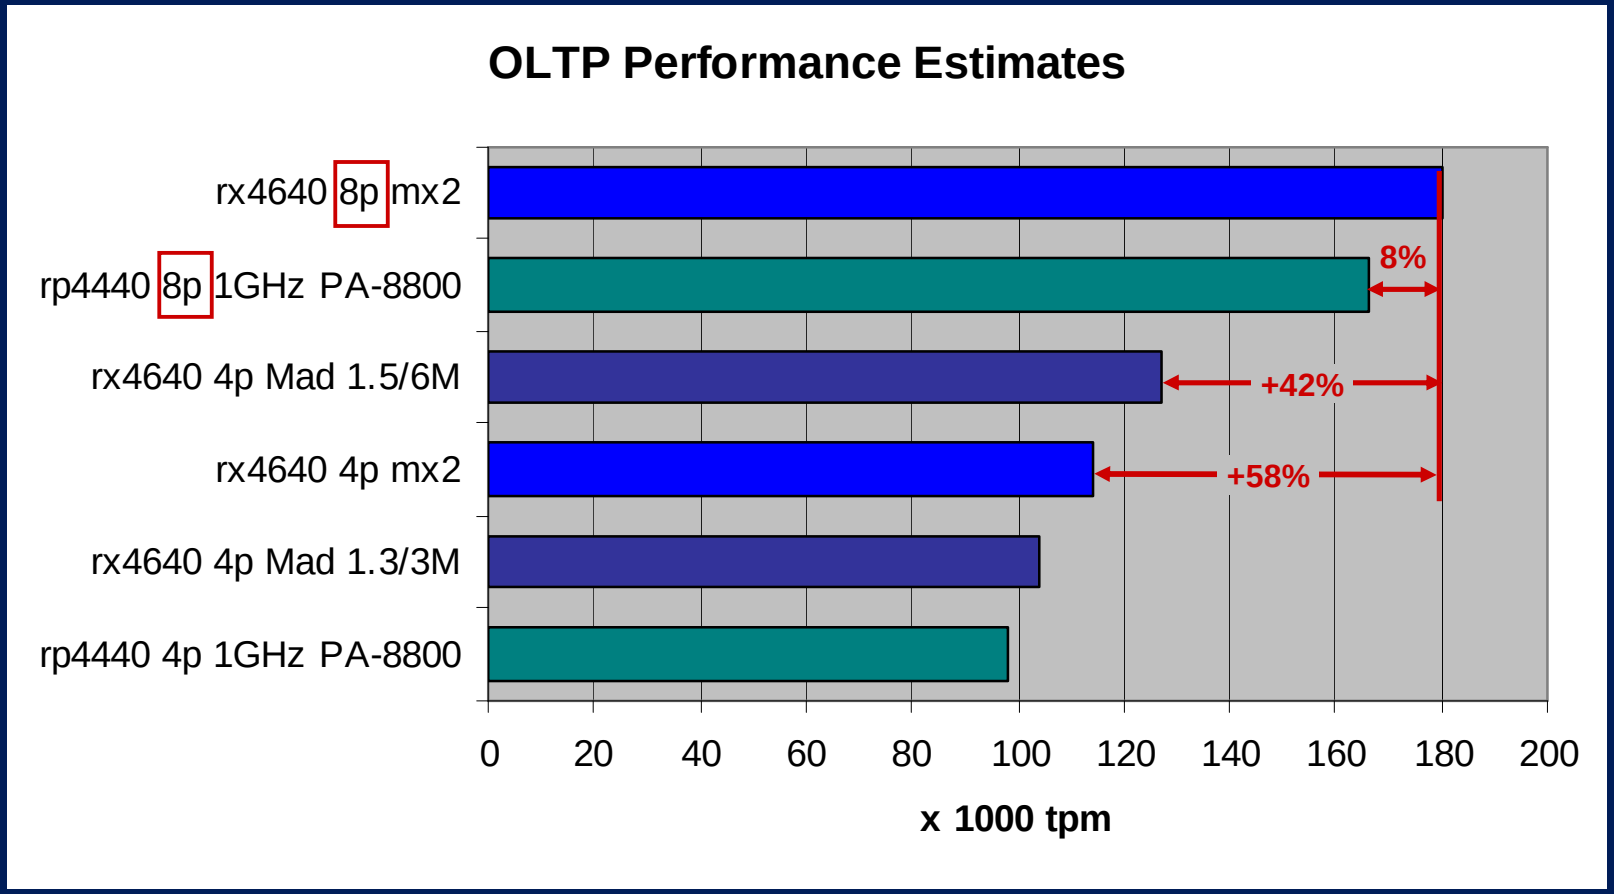

All figures are estimates and assume use of Oracle 10g database. rp4440 estimates assume the use of HP-UX 11i v1; rx4640 estimates assume the use of HP-UX 11i v2.

Agenda

- IT Requirements and Trends
  - HP Adaptive Enterprise
  - HP Industry Based Server Strategy
- Itanium® Processor Technologies and Roadmap
- Itanium® Processor chipsets & Module
- HP Integrity Server
- Server consolidation

HP Integrity Server

HP Integrity

High-end

Itanium 2 6M Processor
mx2 Processor Module

HP Integrity Superdome
up to 128 processors

XC6000
HPTC cluster

Midrange

Itanium 2 6M Processor
mx2 Processor Module

HP Integrity
rx7620-16
(16-way)

Itanium 2 6M Processor
mx2 Processor Module

HP Integrity
rx8620-32
(32-way)

Entry Level

Itanium 2 9M Processor

HP Integrity
rx1620-2*
(2-way)

HP Integrity
rx2620-2*
(2-way)

Itanium 2 9M Processor
mx2 Processor Module

HP Integrity
rx4640-8 (8-way)
dense form factor

HP NonStop

2004

Intel®
Itanium®
architecture

NonStop
Advanced
Architecture

HP ProLiant

XC3000
HPTC cluster

HP ProLiant
ML series

HP ProLiant
DL series

HP ProLiant
BL series

* New core I/O

HP Integrity rx1620-2 Server

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
- New Advanced Remote Server Management (iLO)

High availability

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

Processors and chipset

- New Mad9M Intel® Itanium2®-based "Fanwood" processors:
- 1.6GHz/3MB 267 FSB (533 MT/s)
- 1.3GHz/3MB 200 FSB (400 MT/s)
- HP zx1 Chipset

Memory

- 512MB to 16 GB (8 slots)
- PC2100 DDR DIMMs w/ ECC
- Chip sparing support

Internal peripherals

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

Form factor

- 1 EIA units (U) or 1.75″ height
- 41 servers per 2m rack
- Designed for data center and utility closet operation (5–35°C)

I/O subsystem

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

Operating Systems

- HP-UX 11i v2
- Windows 2003 Ent. Edition
- Linux Red Hat AS3 & SUSE SLES 9
- OVMS v8.2 (Q1 2005)

HP Integrity rx2620-2 Server

Management

- Local, Web, and remote console capability with remote power
- Built-in fault management system
- 10/100Base-T management LAN
- Baseboard management controller
- New Advanced Remote Server Management (iLO)

I/O subsystem

- 4 PCI-X full-length slots
  - One 133 MHz (1 GB/s)
  - Three 133 MHz (512 MB/s)
- 2 Ultra320 SCSI Channels
- 2 Gigabit ports
- 100Base-T USB, VGA, serial

Internal peripherals

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

Processors and chipset

- New 1 or 2 Intel ® Itanium ® 2 Processors 9M:
- 1.6GHz/3MB 200 FSB (400 MT/s)
- 1.3GHz/3MB 200 FSB (400 MT/s)
- HP zx1 Chipset

Memory

- 1 GB to 24 GB
- PC2100 ECC chip spare DDR
- 128-bit, 266 MHz DDR

Form factor

- 2 EIA units (U) or 3.5″ height
- 20 servers per 2m rack
- Designed for data center and utility closet operation (5–35°C)
- Standalone, pedestal option

High availability

- Redundant hot-plug power
- Dual SCSI channels
- CPU de-allocation on failure
- Dynamic processor resilience

Operating systems

- HP-UX 11i v2
- Linux Red Hat AS3 and SuSe SLES 9
- Windows Server 2003 Ent. Edition
- OVMS v8.2 (Q1 2005)

HP Integrity rx4640-8 server

- Processors
  - 1- to 4-way Itanium® 2–based system
    - 1.6 GHz (9 MB cache) – only 4 P-Bundle
    - 1.6 GHz (6 MB cache)
      1.5 GHz (4 MB cache)
  - 1 to 4 HP dual processor mx2 modules (4 to 8 processors)
    - 1.1 GHz (32 MB of L4 cache)
- Up to 64 GB memory
- Six available PCI-X slots
- 10 servers per 2m rack

Unsurpassed OS support

- Runs HP-UX 11i v2, Linux, Windows Server 2003, and OpenVMS (2H 2004)
- Broad choice of applications

Solid high availability

- Redundant, hot-plug power supplies, fans, and hard disk drives
- HP Serviceguard for HA clustering

*not available with mx2 dual processor modules

Unmatched performance density

- Rack-dense 4U chassis
- Blazing fast HP zx1 Chipset

Leading rapid deployment

- Available site prep, installation, and custom configuration

Excellent investment protection

- In-box upgradable to future Itanium processors

Easy to install and service

- Fast, “no-tools” rack installation
- Quick Find LED diagnostic panel accelerates problem isolation
- Designed for “no tools” customer repair with swap-out components

Automated intelligent management

- Consistent multi-OS management & support
- Comprehensive fault, inventory and configuration management
- Dedicated management processor

HP Integrity rx7620-16 & rx8620-32 server

rx7620

(codename: Eiger)

- 2- to 16-way Itanium2 CPUs, 64 GB memory
- 15 PCI-X slots
- 1 removable media device (DAT/DVD)
- 4 internal disk drives
- dual core I/O
- up to 2 hard partitions

rx8620

(codename: Olympia)

- 2- to 32-way Itanium2 CPUs, 128 GB memory
- 16 PCI-X slots
- 2 removable media device (DAT/DVD)
- 4 internal disk drives
- dual core I/O
- up to 2 hard partitions (up to 4 w/ optional SEU)
- optional server expansion unit (SEU)

OS choice

- Runs HP-UX11i, Windows Server, and Linux
- Supports multi-OS in same server
- Unmatched application choice

high availability

- chip-spare technology
- fault-tolerant power compliance
- MC/Serviceguard for in-box and cluster failover

investment protection

- In-box upgradeable to future generations of Itanium processors (including dual-core CPUs)
- in-box upgradeable to future high-performance chipsets

physical specifications

- leading performance density
- "no tools" accessibility and serviceability

utility pricing

- pay-per-use*
- pay-per-forecast*
- instant capacity on demand*
- temporary capacity*
- cell iCOD*

* Initially supported on HP-UX partitions only

rapid deployment

- includes site prep and installation
- Includes custom configuration

HP Integrity Superdome

- Processors
  - 2- to 64-way Itanium® 2–based system 1.5 GHz (6 MB cache)
  - 2 to 64 HP dual processor mx2 modules (4-128 processors)
- 1 TB memory
- 192 PCI-X slots
- Up to 16 hard partitions
- Concurrent Multi-Operating System Support
- Ideal for Multi-OS Consolidation
- Multi-OS: HP-UX 11i v2, Windows Server 2003, Linux and OpenVMS (2005)

HP Integrity Superdome

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
  - Online power and fan replacement
  - N+1 power and fan
  - Dual power sources

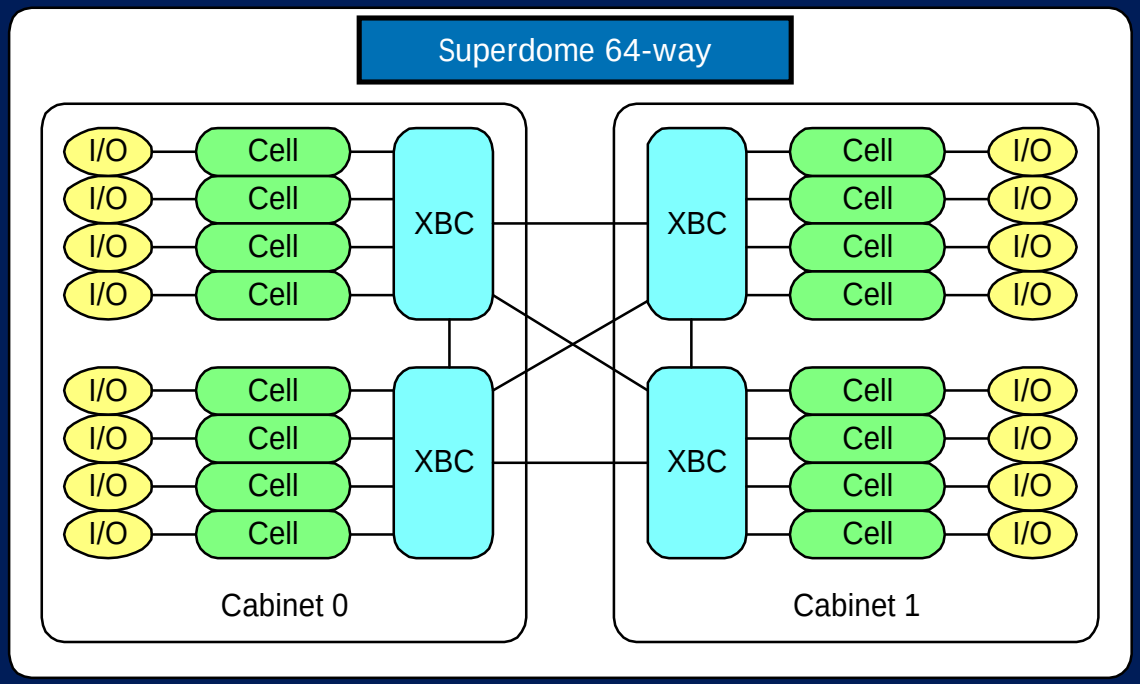

Agenda

- IT Requirements and Trends
  - HP Adaptive Enterprise
  - HP Industry Based Server Strategy
- Itanium® Processor Technologies and Roadmap
- Itanium® Processor chipsets & Module
- HP Integrity Server
- Server consolidation

Server consolidation Element Virtualization

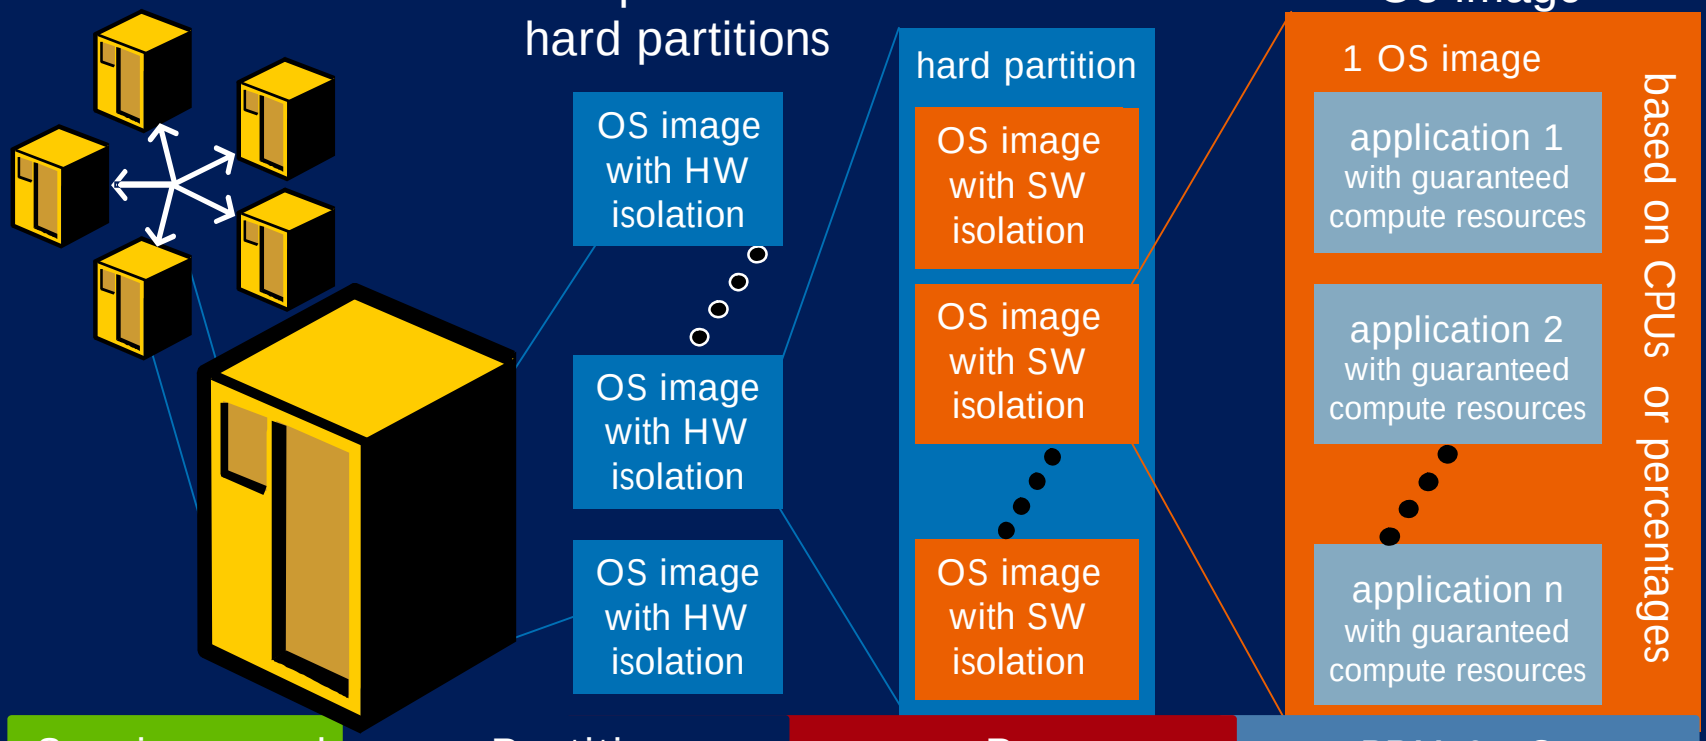

Isolation
Highest degree of separation

Flexibility
Highest degree of dynamic capabilities

Server consolidation
HP Integrity Virtual Machines

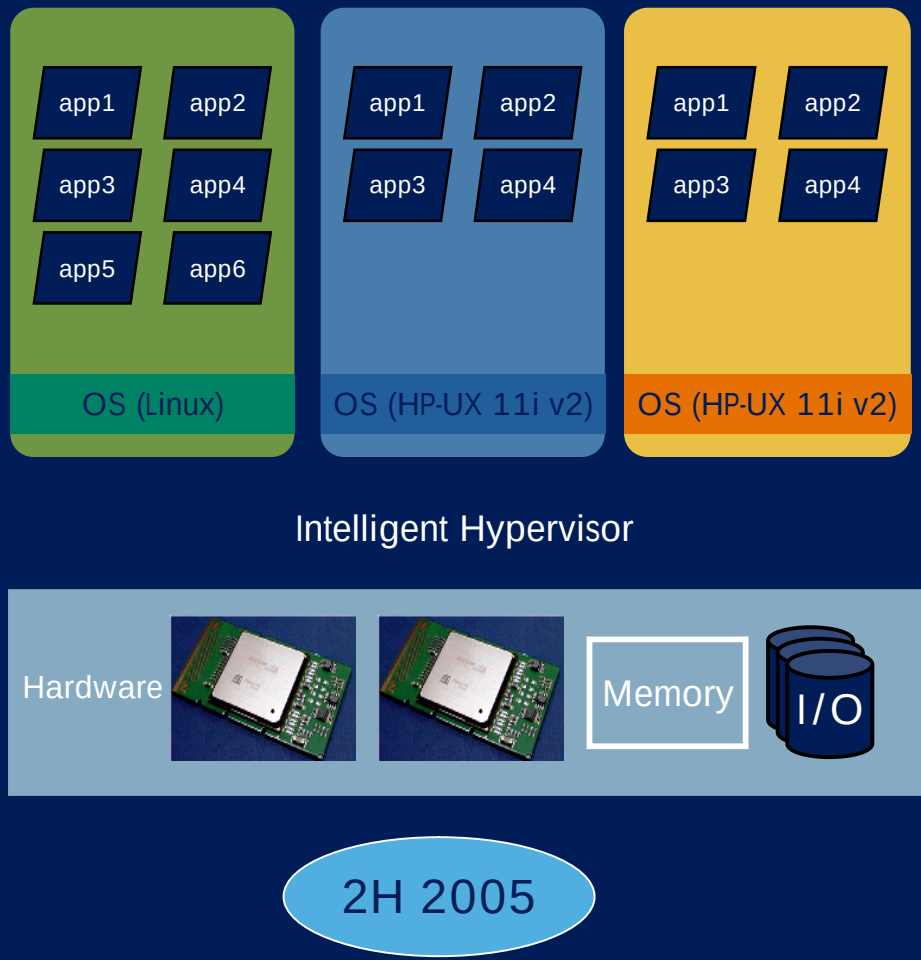

- Sub CPU virtual machines with shared I/O
- Runs on a server or within an nPar
- Dynamic resource allocation built-in
- Resource guarantees as low as 5% CPU granularity
- OS fault and security isolation
- Supports all (current and future) HP Integrity servers
- Designed for multi OS
- VSE integration for high availability and utility pricing

Summary

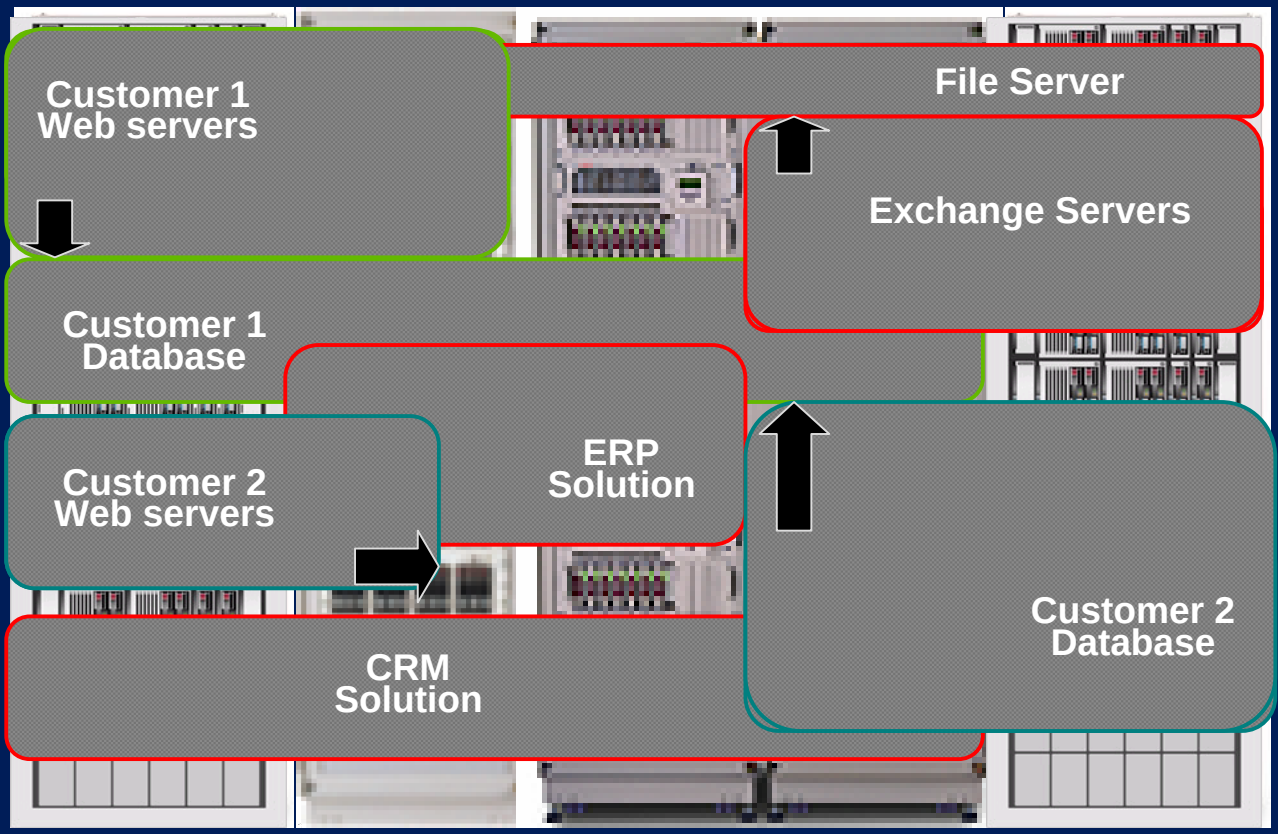

- HP Integrity Server the ideal platform for Adaptive Enterprise

Question and Answers


hp
invent