

# **Intel® Itanium® Architecture Roadmap and Technology Update**

**Dr. Gernot Hoyer**  
Technical Marketing EMEA

# Intel® Itanium® Architecture Growth

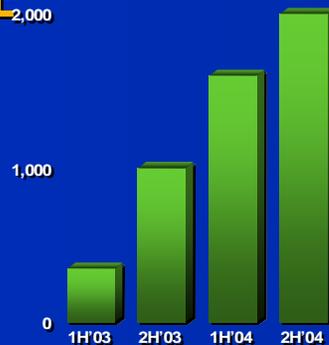
## MARKET

- Over 3x revenue growth Y/Y\*
- More than 10x growth\* in shipments of large SMP systems (64+)

IDC Worldwide Quarterly Server Tracker, August 2004

## SOFTWARE

- Over 100% Y/Y growth
- 2004 forecast of 2000 applications reached TODAY!



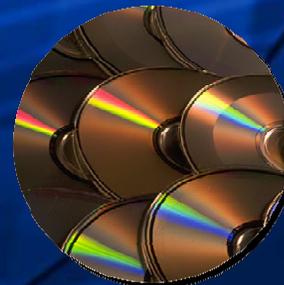
## HARDWARE

- # OEM server models keep growing
- |           | 2002 | 2003 | 2004 |
|-----------|------|------|------|
| 2P, 4P    | 20   | 50   | 70   |
| 8P - 128P | 5    | 15   | 20   |
- 8 of 9 RISC vendors selling Intel Itanium-based servers

## END-CUSTOMERS

- 38 of Global 100 companies using Intel Itanium-based servers today
- High profile wins: General Mills, Pfizer, Thomson Financial, Procter and Gamble, The Weather Channel, First American Title, Motorola

intel



Other names and brands may be claimed as the property of others.

# Broad Ecosystem Support

## Application Choice



- >2000 native applications
- 32-bit application support with IA-32 Execution Layer

## Operating System Choice



- Windows\*, Linux\*, and Unix\* support

## System Vendor & Platform Choice



- >70 vendors selling today
- 20 large SMP systems
- 2-way to 256-way systems
- 8 of 9 RISC vendors support

***Broad industry support  
Choice that you don't get with RISC***

# Itanium® Processor Family: A Strategic Product for Intel

Design teams working on  
more than 6 future  
processors



# Intel® Itanium® Processor Family

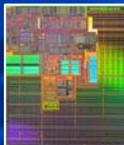
2001



2002



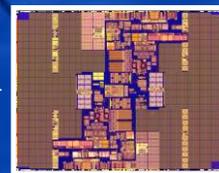
2003



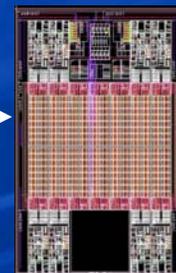
2004



2005



future



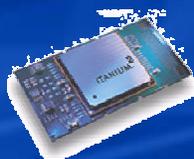
Itanium® 2  
Madison\*\*

Madison9M

Montecito\*\*

Tukwila\*\*

←----- Common Platform -----→



800MHz  
4MB L3-Cache  
460GX Chip-set  
OEM Chip-sets  
180nm

1GHz  
3MB iL3-Cache  
E8870 Chip-set  
OEM Chip-sets  
180nm

1.5GHz  
6MB iL3-Cache  
E8870 Chip-set  
OEM Chip-sets  
130nm

1.6GHz  
9MB iL3-Cache  
E8870 Chip-set  
OEM Chip-sets  
130nm

>>1.6GHz  
24MB iL3-Cache  
Dual-Core  
E8870  
OEM Chip-sets  
90nm

Multi-Core  
Developed by  
former Alpha  
team

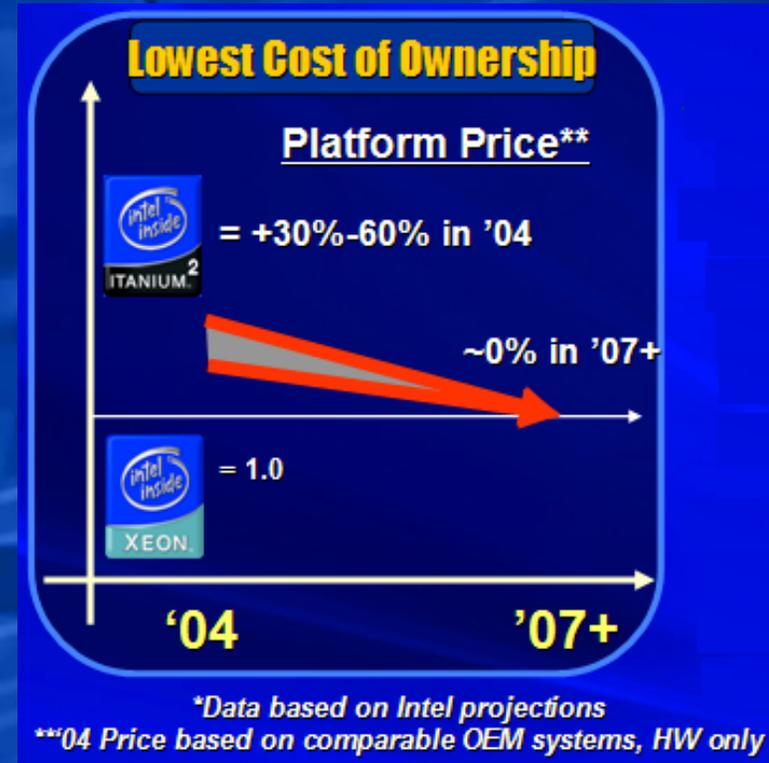
\*\*codename

All features and dates specified are targets provided for planning purposes only and are subject to change.



# Common Platform & Infrastructure

- Today: Itanium® Processor exceeds RISC performance & price / perf
- Today: Itanium® platform delivering superior price / performance vs Intel® Xeon™ Processor on transaction processing
  - 30% more transactions at 10% incremental cost of hardware platform/ OS / database\*\*\*
- '07: Itanium® platform cost reduced to parity with Intel® Xeon™ processor-based platforms
  - Common platform components to lead to common platform infrastructure over time



# Choice and Flexibility for Evolving Enterprise Servers

## Current Architecture or Solutions

RISC Architecture

## Transition Benefits

Exceptional performance – choice of operating systems, software and hardware vendors TODAY

## Architecture of Choice



Mission critical 64-bit architecture – Premier performance, reliability, scalability; cost effective vs. RISC

Target Applications  
Database, ERP, BI, HPC

IA-32 Architecture

64-bit support via Intel® EM64T, great performance for 32-bit applications



Mainstream 64-bit architecture; price – performance – reliability

Target Applications  
MP: SCM, CRM, BI, ERP  
DP: HPC, Application Server, Workgroup E-Commerce, Portals, Firewall/Security, Workstation apps

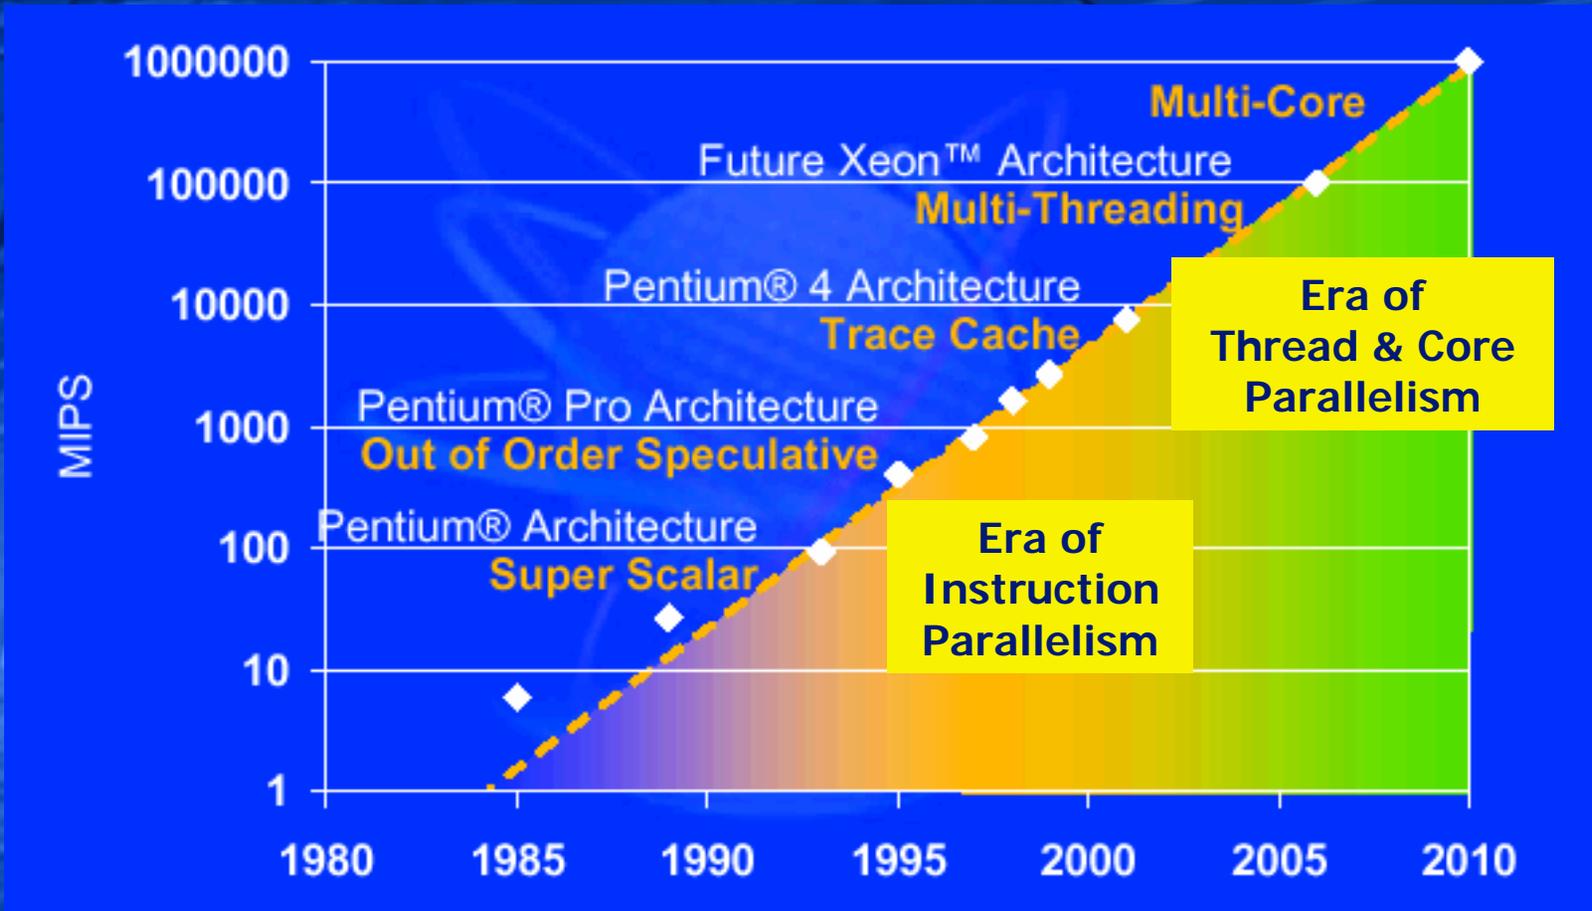


\* Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

# **Future technologies: Processor and Platform Enhancements**

- **Multiple Cores and Thread**
- **Virtualization**
- **Power Management**

# Intel Processor Technology Trends

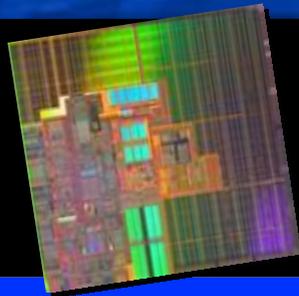


All features and dates specified are targets provided for planning purposes only and are subject to change without notice

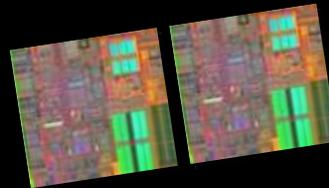


# Enterprise Multi-Core Transition

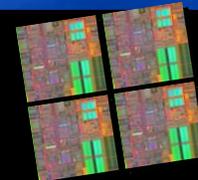
## ...dual core a natural evolution



**TODAY**  
Single Core

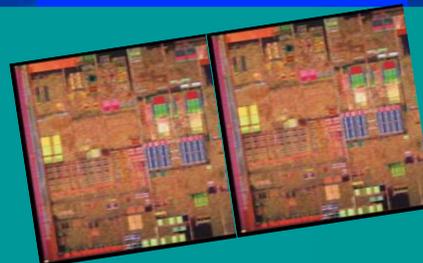


**2005-2006**  
Dual Core



**4 or more cores**  
+ cache

**Future**  
Multi-Core



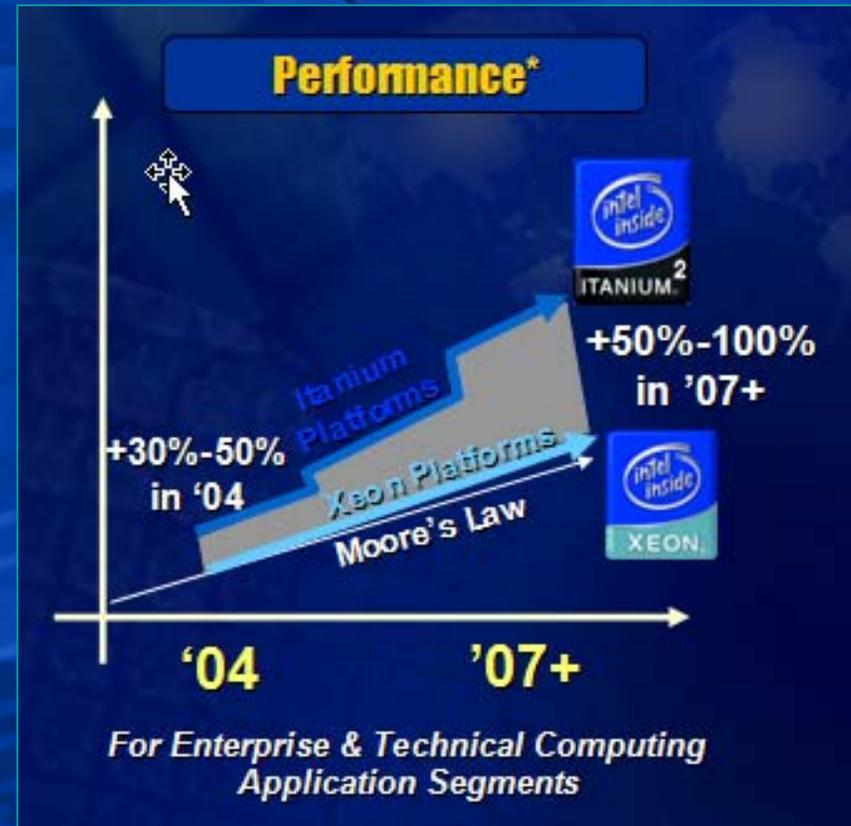
**2 or more cores**  
+ cache



# Itanium<sup>®</sup> Architecture

## Optimized for Multi-Core

- Architecture: Parallelism and many registers to keep data on-chip
- Core size: Smaller than IA-32, up to 2x more cores per die on than on IA-32



All products, dates, and figures are preliminary and are subject to change without notice.

\* EPIC is Itanium's architecture, "Explicitly Parallel Instruction Set Computing"

All features and dates specified are targets provided for planning purposes only and are subject to change without notice.

# Multi-Threading Approaches

Sharing resources

High utilization  
High complexity  
Medium latency hiding

Single thread/cycle

Multiple thread/cycle

Medium utilization  
Low complexity  
Medium latency hiding

Temporal

Event

A	A	A	B	B	
A	A	B	B	B	
A	A	B	B	B	
A	A	A	B	B	B
A	A	B	B	B	B

A	A	A			
B	B	B	B	B	
A	A	A	A		
B	B	B			
A	A	A			
B	B	B	B	B	

A	A	A	A		
A	A	A	A	A	
B	B	B			
B	B	B	B		
B	B	B	B	B	

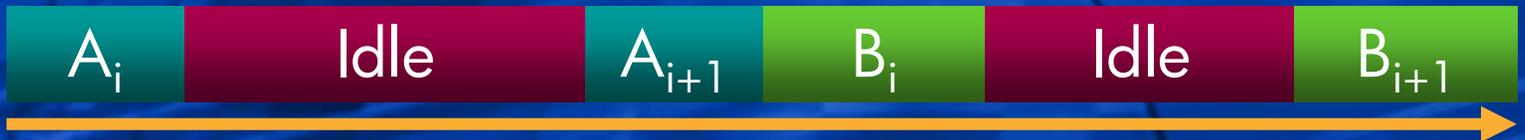
Medium utilization  
Low complexity  
High latency hiding

“Event” for the core and “Multiple” for the caches

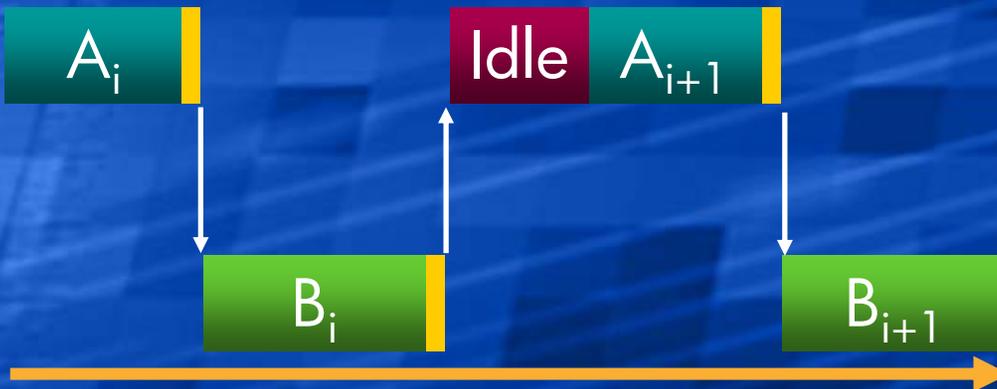


# Montecito Multi-Threading

## Serial Execution



## Montecito Multi-threaded Execution



Multi-threading decreases stalls and increases performance

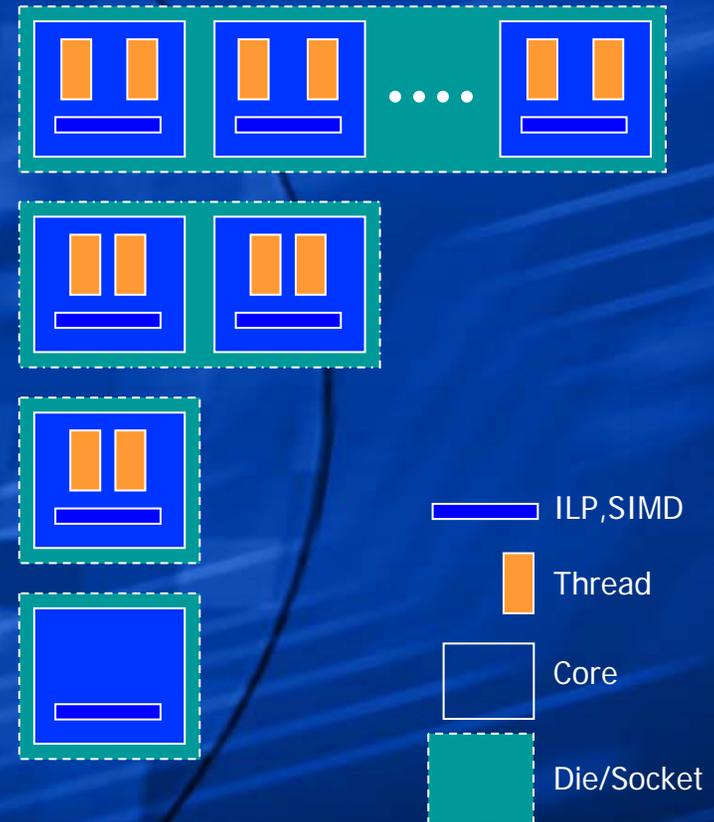
# Dynamic Thread Switching

- **Optimal**
  - Determine when execution is stalled for long latency operations
- **Practical**
  - Predict that a long latency event will stall execution
  - Hysteresis to avoid needless switches
- **`hint@pause` gives software control**

Effective solution allowing streaming and access clumping

# Multi-Level Parallelism

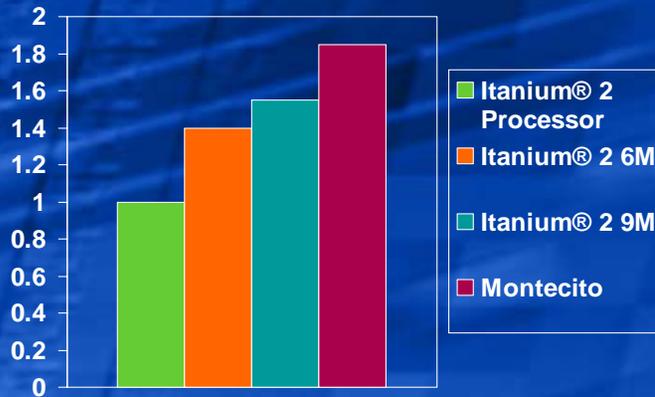
- **MULTI-CORE**
- **DUAL-CORE**
- **MULTI-THREADING**
- **SINGLE CORE (ILP, SIMD)**



# Performance Innovations

- Itanium® Processor Performance Strategy: Increased performance/thread, then increased number of threads

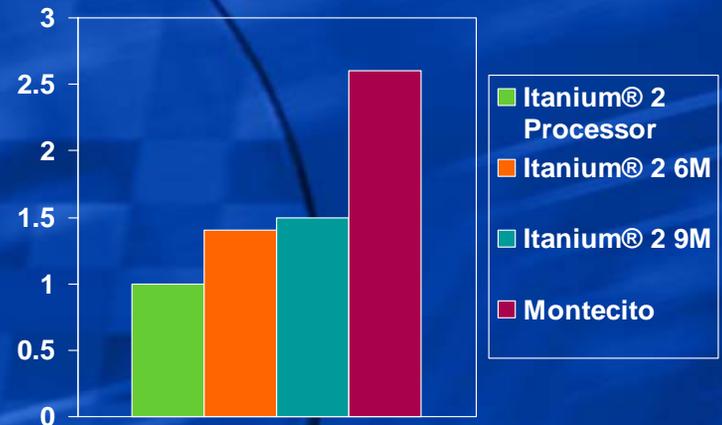
## Increased Performance/Thread



Floating Point Performance (Single thread) - Relative Performance

- Driven by:
  - Increased frequency
  - Increased L3 cache
  - Increased bus speed

## Multi-threaded Performance



4S Transaction Processing Relative Performance

- Driven by:
  - Dual core Montecito
  - Multi-threading support in Montecito

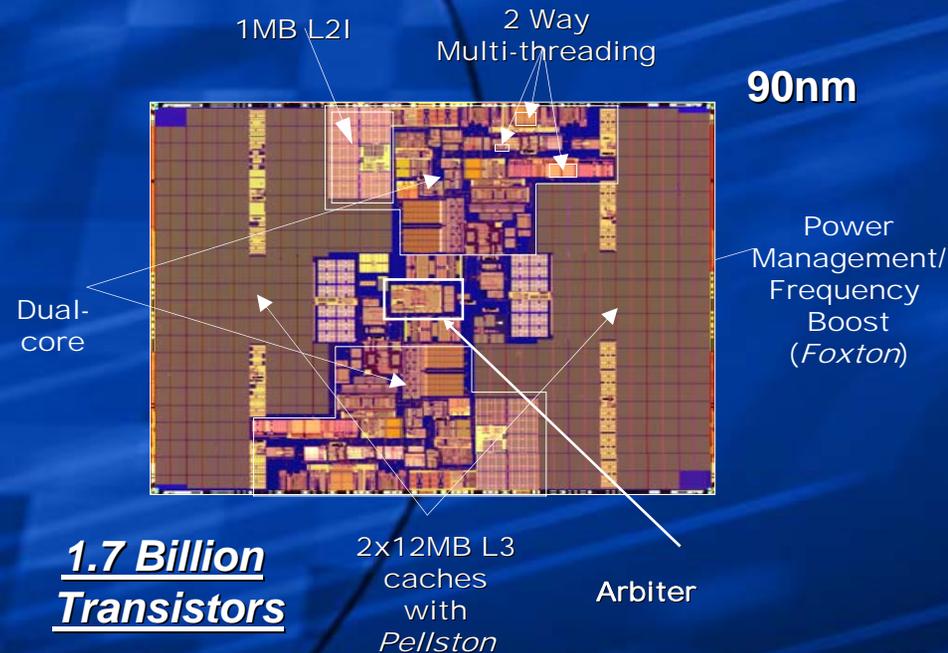
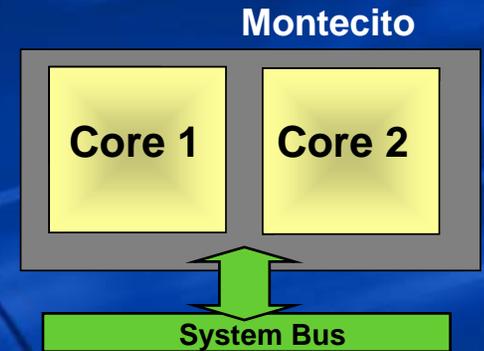
**Montecito: 4 virtual processors**

\*Third-party marks and brands are the property of their respective owners; Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. All products, dates, and figures are preliminary and are subject to change without any notice. Copyright © Intel Corporation 2004.



# First Implementation of Key Features: Montecito

- Key Processor Features
  - Intel's first dual-core processor
  - Intel's first processor with >1 billion transistors
  - 24 MB L3 cache
  - Multi-threading
  - Compatible with existing Itanium 2-based systems
  - Pellston, Foxtan and Silverdale Technology
- Targeting H2'2005



**Multiple cores, Multiple threads  
and L3 Cache on ONE die**

# Montecito Technologies

- **Pellston Technology**

- Automatically disables cache lines in the event of a hard cache memory error
- Allows processor and system to continue normal operation

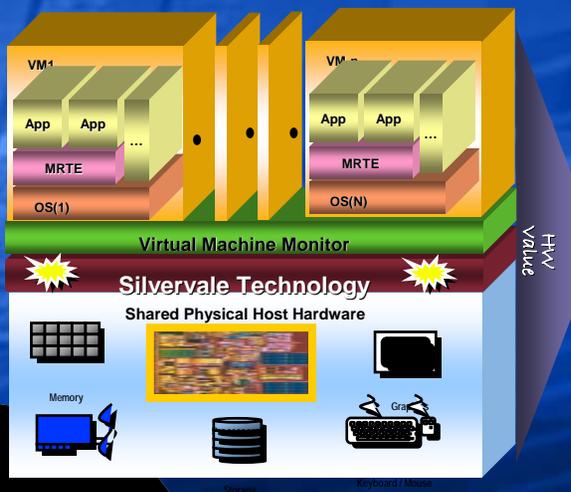
**Improve reliability / uptime**

- **Foxtan Technology**

- Processor boosts performance dynamically based on application power consumption (up to 10% freq.)
- Largest performance impact expected on transaction processing

**More performance /  
no platform modifications**

- **Silverbale Technology**



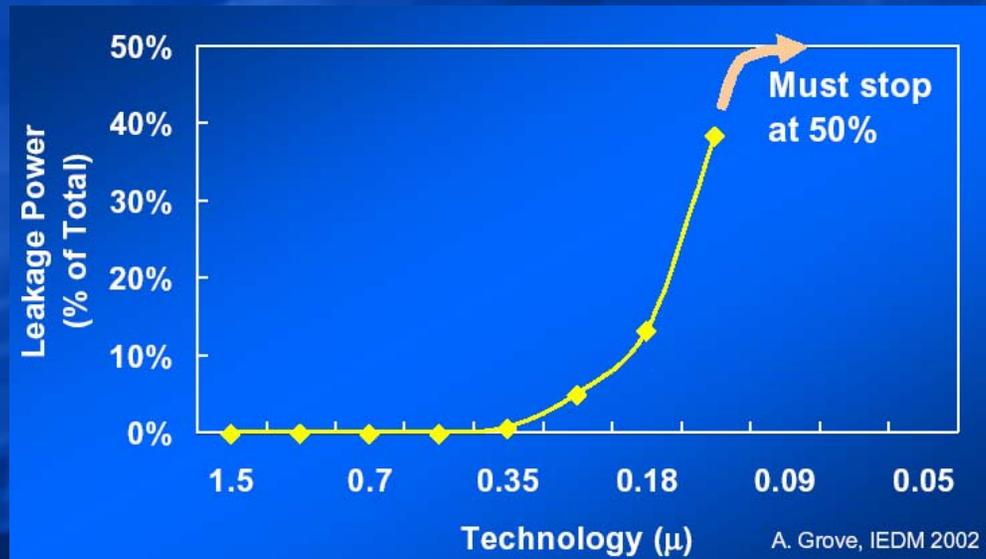
- Provides the hooks for a hardware supported virtualisation
- Enables the Virtual Machine Monitor to be more robust and shows a higher performance

**Next level of Server Consolidation**



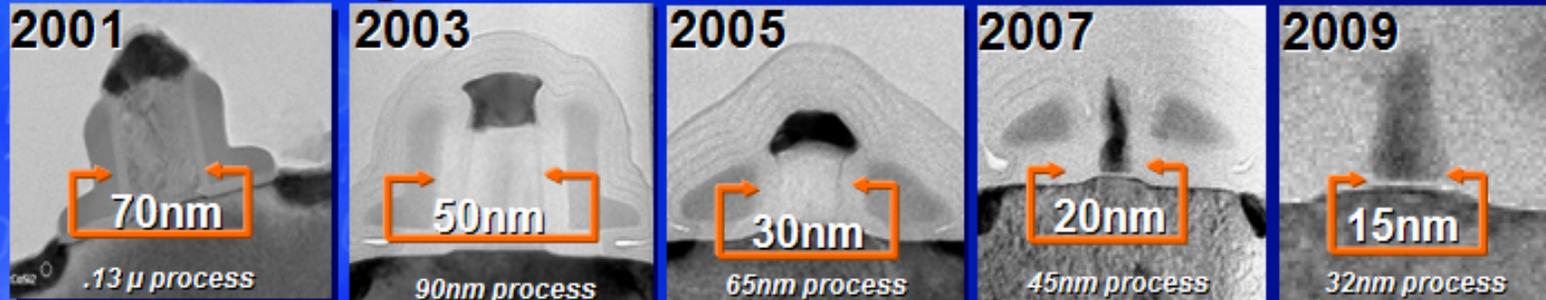
# Designing for Power

- Leakage current increases exponentially as process size decreases linearly
- Business as usual is not an option – 45nm CPU might need 1kW

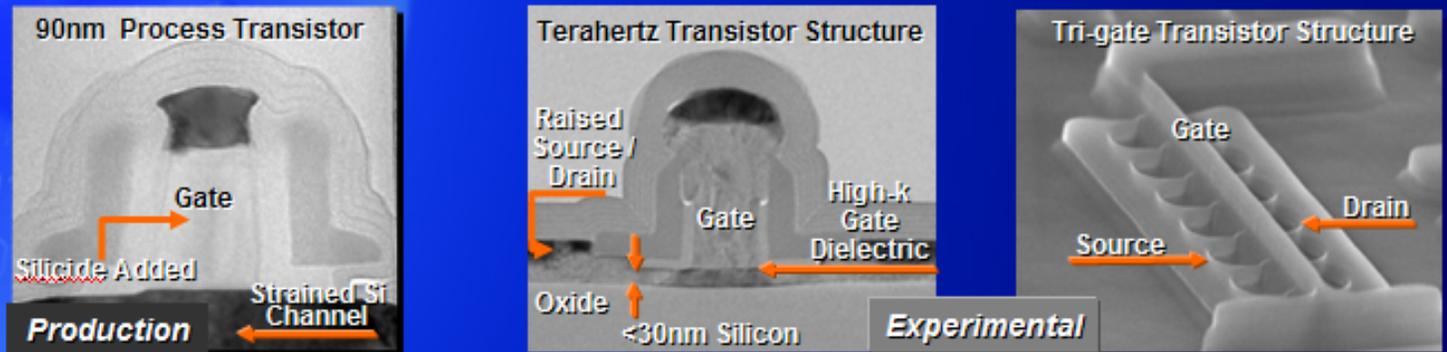


# Transistor Design Roadmap

## Intel Gate Lengths



## Intel Transistor Structures



Extending Moore's Law:  
Continued World Leading Transistor Scaling and  
Novel Structures for Low Power / High Performance

# Power Reduction Features

- In addition to on-going improvements such as voltage scaling, new power reduction features are planned for each process generation

	90nm 2003	65nm 2005	45nm 2007	32nm 2009
Strained silicon	✓	✓	✓	✓
Sleep transistors		✓	✓	✓
High-k/ metal gate			✓	✓
Tri-gate transistors				✓

# Power Management Innovation

Growing Capabilities over TIME

## Reduced Power via Management

- Demand Based Switching (**DBS**)
- Aggressive use of **C1E** state

## Power Prediction/Monitoring

- Report Configuration Power (**PConfig**)
- Monitor Power (**PSMI**)

SILICON TRANSISTORS PACKAGES

HEAT SINKS

SYSTEMS

FACILITIES

## Lower power cores

- Process (**65nm**)

Proliferate Benefits

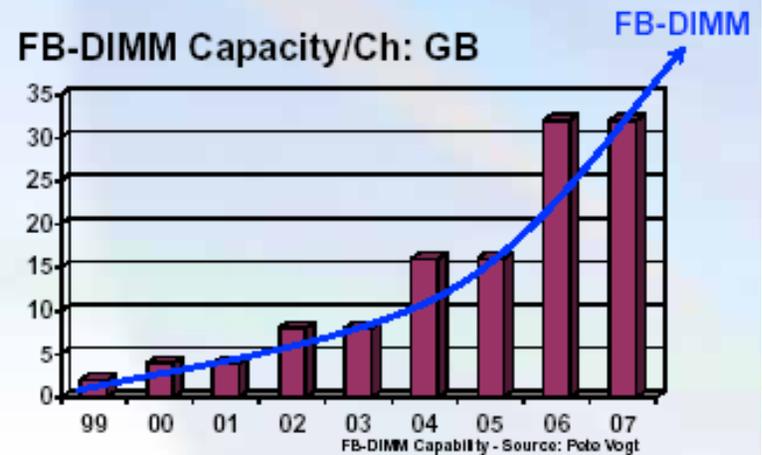
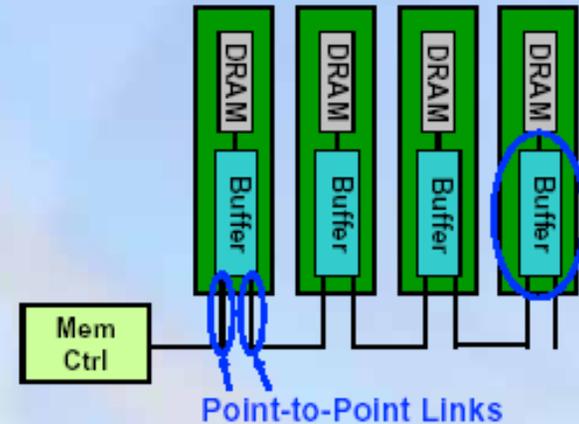
# Advances in Memory Technology

<b>DDR2 Features</b>	<b>IT Benefits</b>
<b>DRAM Technology for Multiple Generations</b>	<b>Increased Platform Longevity</b>
<b>Higher Bandwidth</b>	<b>Higher System Performance</b>
<b>Lower Power</b>	<b>Increased Server Density</b>
	<b>Increased Memory Density</b>
<b>On Die Termination</b>	
<b>Four DIMMs per Channel</b>	<b>Lower cost memory configurations</b>

Note: Comparison relative to DDR333

# Fully-Buffered DIMM Memory

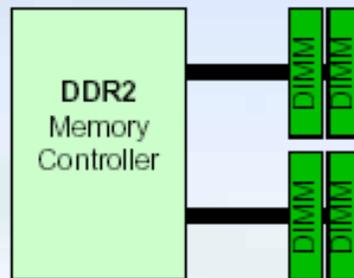
- FB-DIMM buffers the DRAM data pins from the channel and uses point-to-point links to eliminate the stub bus
- FB-DIMM capacity scales throughout DDR2 & DDR3 generations



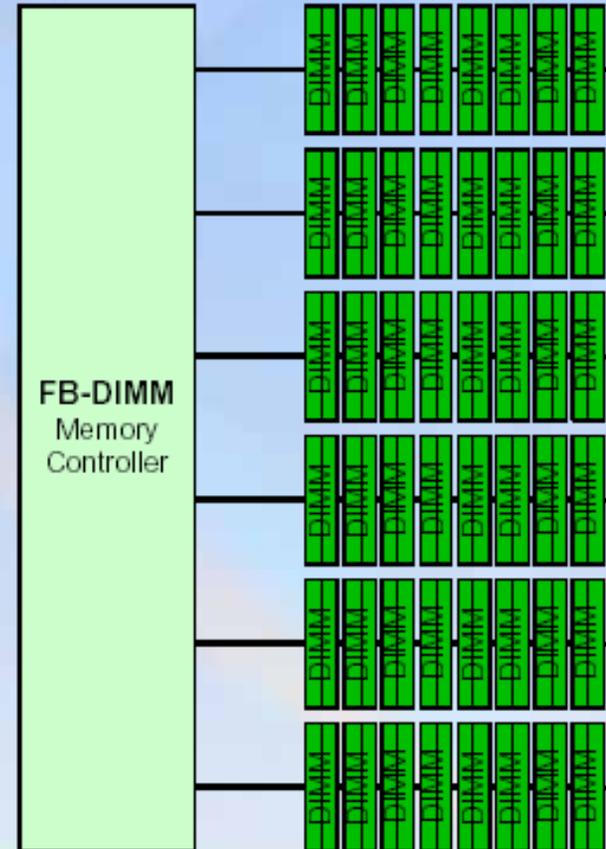
# DDR2 vs. FB-DIMM

## Capacity Comparison

- 24x capacity
  - 8GB vs. 192GB
- ~4x bandwidth
  - ~10GB/s vs. ~40GB/s
- ~Lower pin count
  - ~480 vs. ~420



8GB with 1Gb x4 DRAMs  
~10GB/s of BW w/DDR2-800  
(only 2 ranks per channel)

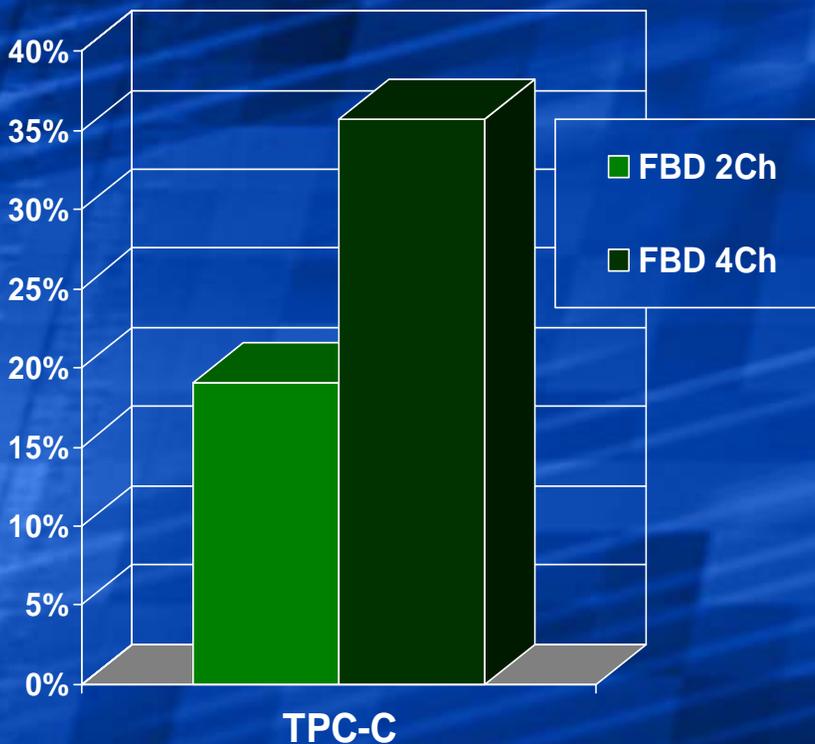


192GB with 1Gb x4 DRAMs  
~40 GB/s of BW w/DDR2-800  
(2 ranks per DIMM)

# FBD & Dual Core Analysis

Simulation -- Theoretical Analysis

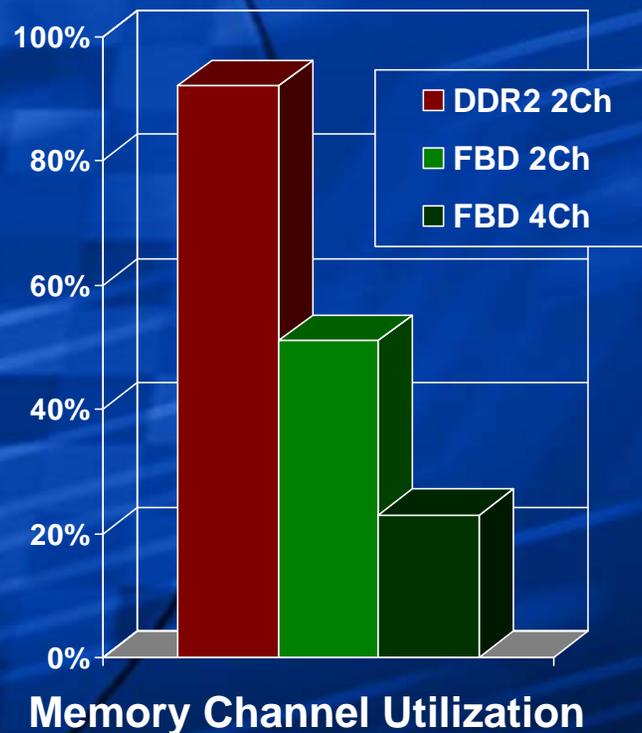
## FBD Performance Advantages Significant with Dual Core CPUs



Gains vs  
DDR2 2Ch

## FBD Addresses Mem Channel Demands of Dual Core CPUs

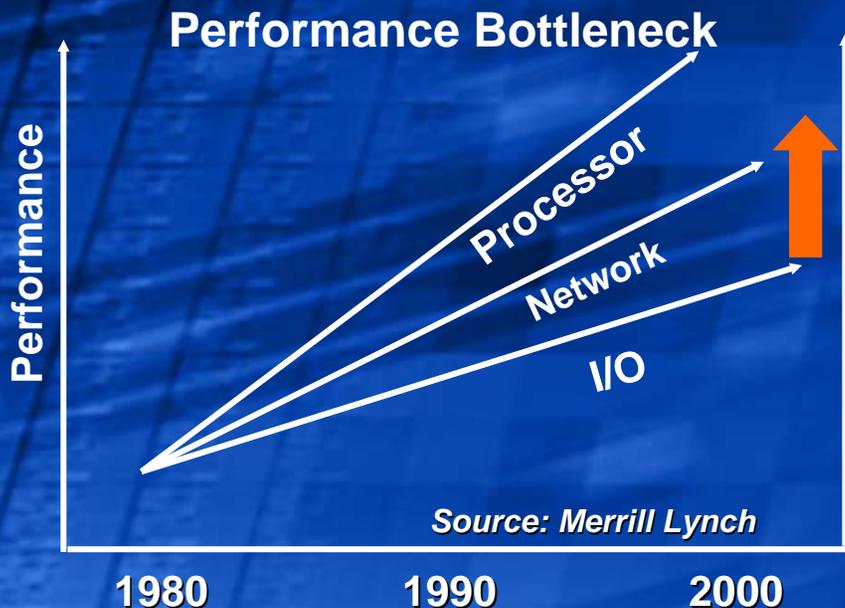
(lower is better)



Unleashing Dual Core Performance Requires FBD



# PCI Express\* for I/O Advantages



## I/O Breakthrough

Keep pace with rest of platform

## Lower Cost

Standards-based & Serial I/O drive inherently lower product development & manufacturing, and TCO

## Investment Protection

“Future proofing” through 10 year roadmap. Increased RAS.

## Rapid Enterprise Adoption

Majority of OEMs/ODMs with PCIe slots, compelling '04 adapter availability, and strong IHV development plans

*“IDC expects PCI Express to be a leading contender in keeping the future IT infrastructure fed with lots of I/O delivered quickly and managed securely.”*

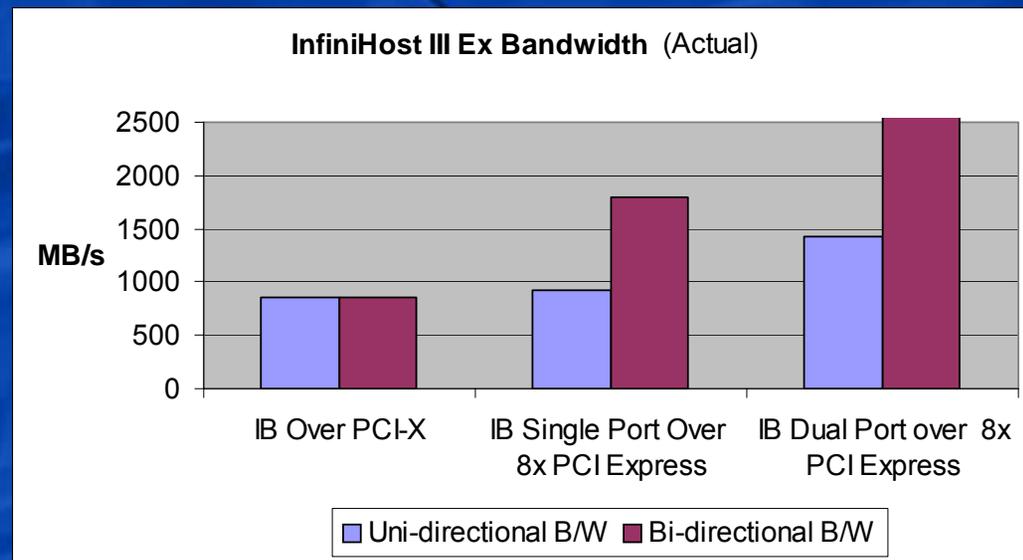
- IDC, Vernon Turner, June 2003

# PCI Express\* Performance

## Mellanox Technologies Data



- Mellanox Measurements:
  - Realizing over 2.9x the bandwidth of PCI-X 133
  - 20% reduction in latency
  - reduced CPU overhead
- Additional performance possible with tuning



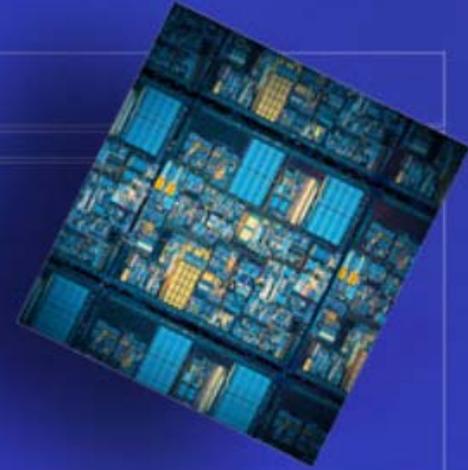
\*\*Source: Mellanox Technologies(March 04)

Full report available at :

[http://www.mellanox.com/products/shared/InfiniHost\\_III\\_EX\\_Launch.pdf](http://www.mellanox.com/products/shared/InfiniHost_III_EX_Launch.pdf)

PCI-Express delivers significant performance improvement over PCI-X

Vielen Dank !



[www.intel.com](http://www.intel.com)

# Backup

# Long Term Goal: 1M Transactions per Minute

Today



In 2007



With planned performance improvements, a 4-way Itanium®-based server in '07 could deliver equivalent OLTP of a current 64-way system, delivering dramatically

- Lower TCO
- Lower power consumption
- Higher density



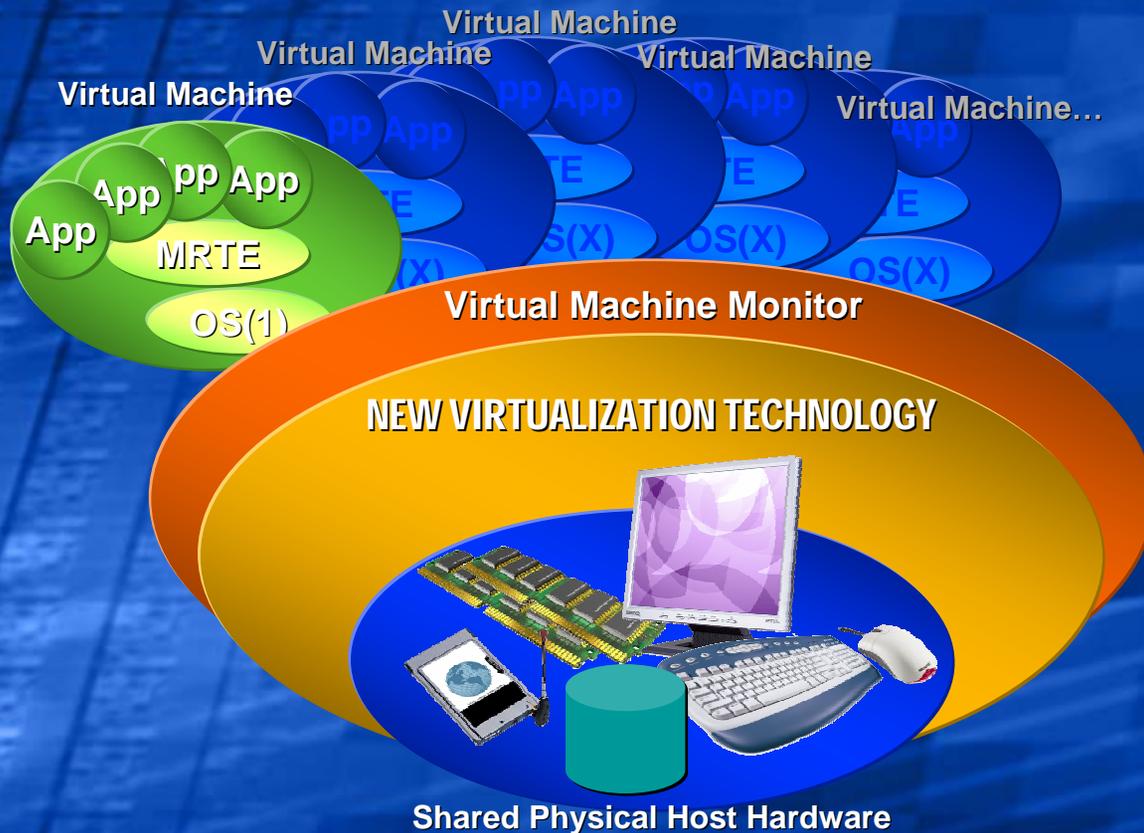
Shown are representations of 64-way system (today) and 4-way system (2007). Not to scale.

All products, dates, comparisons, and information are preliminary and subject to change without notice.



# SILVERVALE TECHNOLOGY

Better Virtualization through OSVs and ISVs



## Virtualization End User Benefits

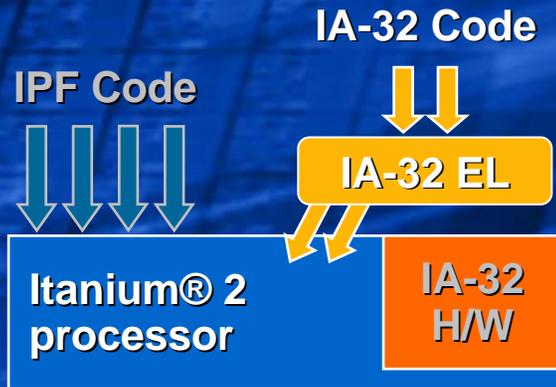
- Reliability
- Efficiency & flexibility
- Security

## Silvervale Technology Benefits

- Choice
- Robustness
- Performance

# IA-32 Execution Layer

## Availability



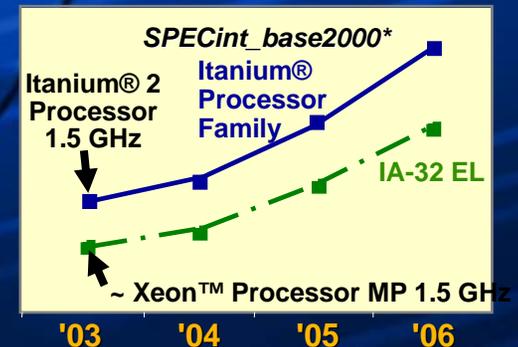
- Historically, support of IA-32 applications has been carried out by on-die hardware
- When using OS with IA-32 EL, support for IA-32 applications is provided by IA-32 EL

Operating System	Available
Microsoft* Windows* Server 2003 Enterprise Edition	✓
Microsoft Windows Server 2003 Datacenter Edition	✓
Microsoft Windows XP Professional 64-Bit Edition	✓
Red Hat* Enterprise Linux 3	✓
SGI* Advanced Linux Environment with ProPack* 3.0	✓
SUSE* Linux Enterprise Server 9	✓
Asianux* 1.0 <sup>1</sup>	2H'04
Red Flag* Advanced Server 4.1 <sup>1</sup>	2H'04
Red Flag DC Server 4.1 <sup>1</sup>	2H'04
Miracle Linux* 3.0 <sup>1</sup>	2H'04

- For Microsoft Windows, IA-32 EL is currently available at [Microsoft Download Center](#), and will ship with Windows Server 2003 SP1 RTM
  - For Linux, IA-32 EL ships or will ship with the OS, as indicated by the availability date
- <sup>1</sup> Asianux 1.0 includes Red Flag Advanced Server 4.1 and Miracle Linux 3.0.

**Gives access to the iA-32 EcoSystem**

Performance Scaling with Future Processors<sup>1</sup>



All products, dates, comparisons and information are preliminary and subject to change without notice.

# SW Tools, Cross Platform Support

Intel Software Development Products		 Windows* Linux*		   Windows* Linux*		 WinCE* Linux*	
		<i>Compilers</i>	C++	●	●	●	●
	Fortran	●	●	●	●	NA	NA
<i>Performance Analyzers</i>	VTune™ Performance Analyzer	●	●	●	●	●	●
<i>Libraries</i>	Integrated Performance Primitives	●	●	●	●	●	●
	Math Kernel Library	●	●	●	●	NA	NA
<i>Threading Tools</i>	Thread Checker	○	○	●	○	NA	NA
<i>Cluster Tools</i>	Trace Analyzer / Collector	NA	●	NA	●	NA	NA

● Currently available

○ Planned



**Single Source Code → Multiple Platforms**  
**Itanium® 2, Xeon™ (EM64T/32Bit) and XScale™ Processor**



# Foxton Technology

- Dynamically adjust Voltage (V) & frequency (f)
  - Exploit full power envelope for all applications
  - Demand Based Switching and flexible power settings
- Large power change → small frequency change ( $P=fCV^2$ )
  - 3% power change with only 1% frequency change
- Monitor/calculate power and temperature
  - Set Voltage to the minimum value needed to support highest frequency
  - Over power and/or temperature results in voltage change
  - Frequency responds to global and local voltage