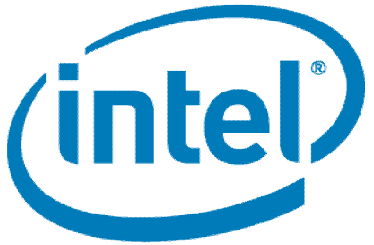

Intel Processor Roadmap:

The Road To PetaScale Computing

HPC @ Intel

Stephen R. Wheat, Ph.D.
PetaScale Strategy Director,
Server Platforms Group
Intel Corp.

March 13, 2007

Risk Factors

Today's presentations contain forward-looking statements. All statements made that are not historical facts are subject to a number of risks and uncertainties, and actual results may differ materially. Please refer to our most recent Earnings Release and our most recent Form 10-Q or 10-K filing available on our website for more information on the risk factors that could cause actual results to differ.

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. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations (http://www.intel.com/performance/resources/limits.htm).

Legal Disclaimer

Notice: This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information. Contact your local Intel sales office or your distributor to obtain the latest specification before placing your product order.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

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

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Itanium® and Intel® Xeon® processors may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

The code names Woodcrest & Clovertown & Tulsa presented in this document are only for use by Intel to identify products, technologies, and services in development, that have not been made commercially available to the public, i.e., announced, launched or shipped. They are not "commercial" names for products or services and are not intended to function as trademarks.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at http://www.intel.com.

3

Market Dynamics*

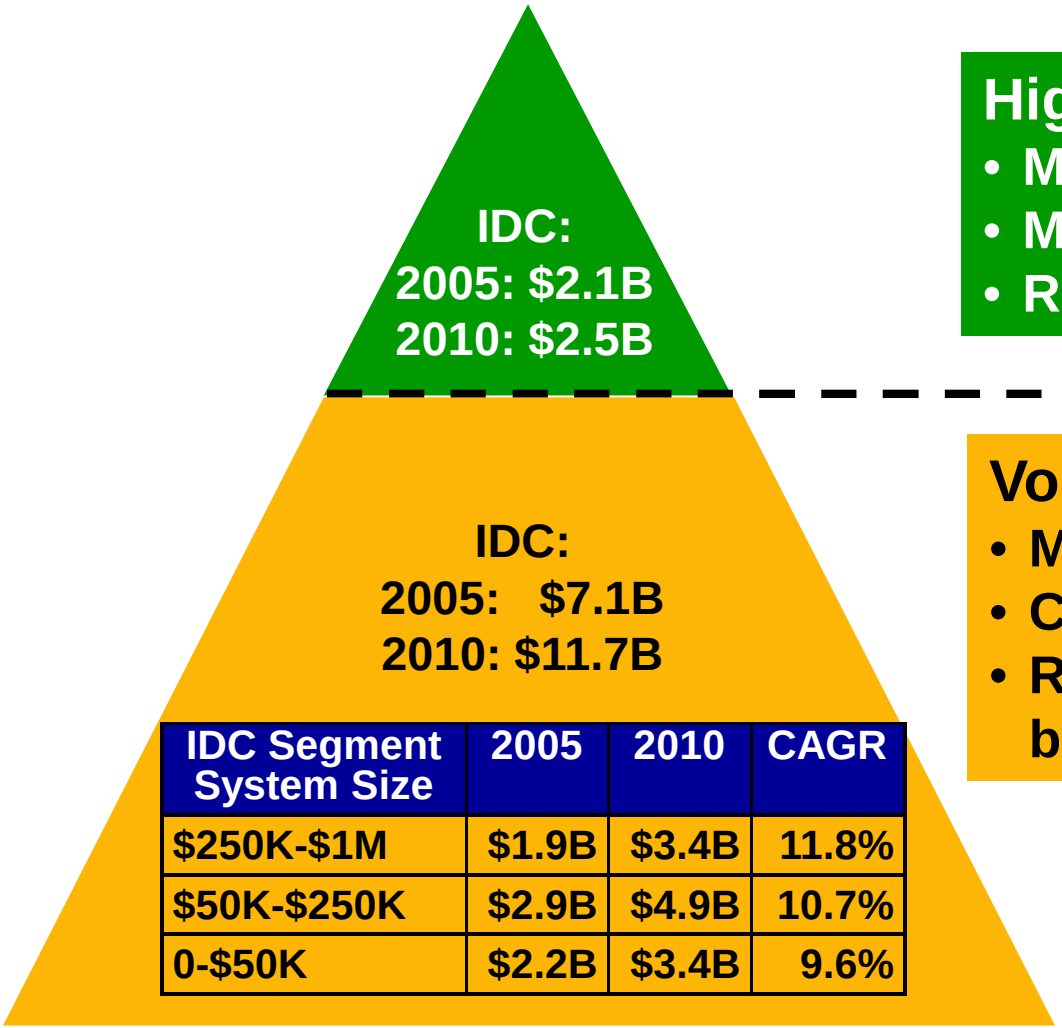

IDC Segment System Size	2005	2010	CAGR
$250K-$1M	$1.9B	$3.4B	11.8%
$50K-$250K	$2.9B	$4.9B	10.7%
0-$50K	$2.2B	$3.4B	9.6%

High End (>$1M system)
- Maps to "Top 50"
- Mainly capability & constellations
- Risk takers, early adopters

Volume
- Mainly capacity
- Clusters >50% of revenue
- Rapid growth since 2002 - driven by broader commercial HPC usage

Total segment > $10.0B in 2006
Forecast >$14.2B in 2010

Approaching 20% of overall server segment, and growing faster

* All data from IDC Worldwide Technical Computing Systems Revenue 2006-2010 #201733 May 2006.

Yesterday, Today, and Tomorrow in HPC

ENIAC
20 Numbers in Main Memory

ASCI Red
(word fastest on top500 till 2000)
First Teraflop Computer,
9298 Intel Pentium® II Xeon Processors

CDC 6600 – First successful Supercomputer 9MFlops

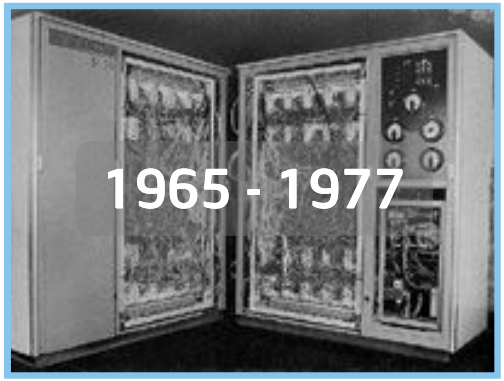

Intel ENDEAVOR
464 Intel® Xeon® Processors 5100 series, 6.85 Teraflop MP Linpack, #68 on top500

~2008 Beyond

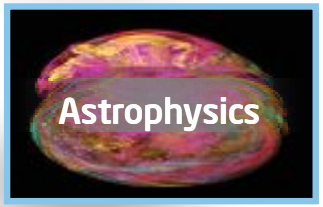

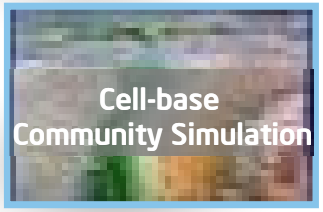

PetaScale Platforms

Yesterday's Supercomputing is Today's Personal Computing

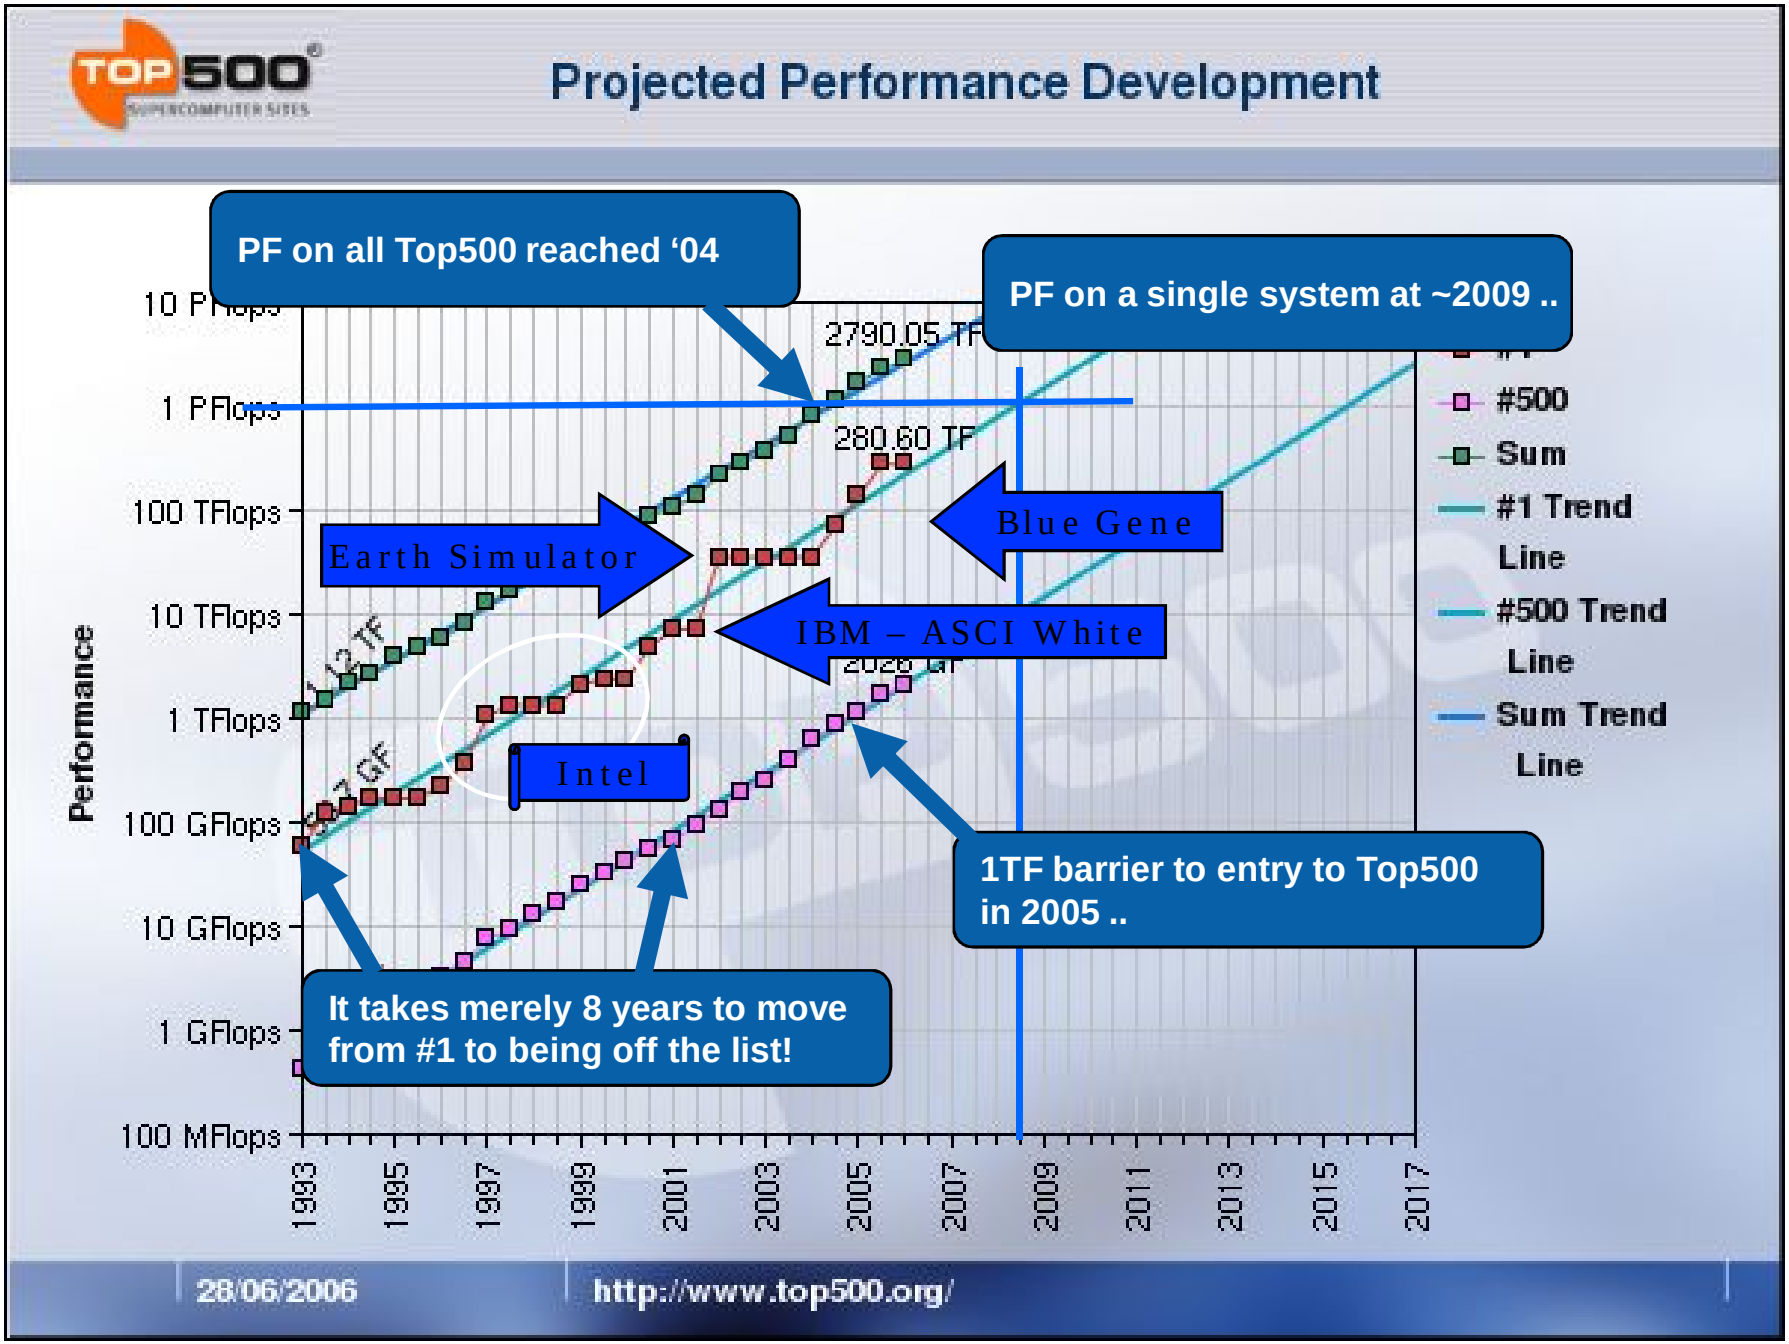

Real World Problems Driving PetaScale & Beyond

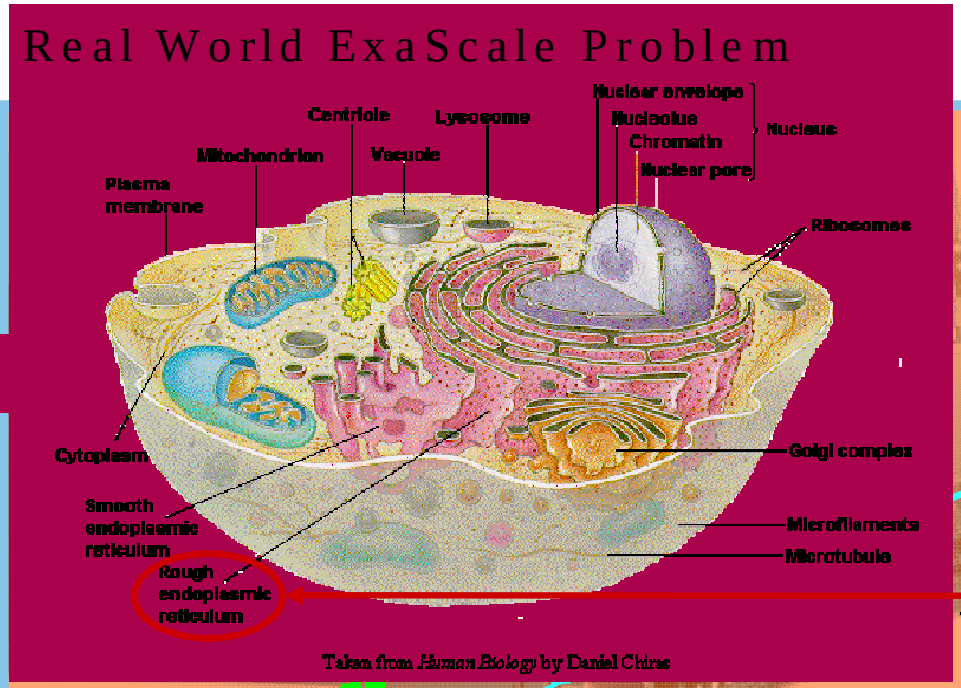

Example real world challenges:

- Full modeling of an aircraft in all conditions
- Green airplanes
- Genetically tailored medicine
- Understand the origin of the universe
- Synthetic fuels everywhere
- Accurate extreme weather prediction

Silicon Future

90nm
2003

65nm
2005

45nm
2007

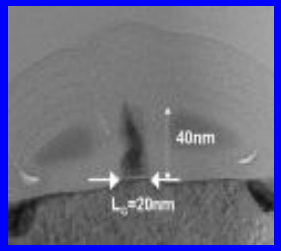

2005 – 2012

32nm
2009

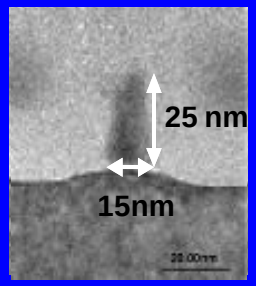

22nm
2011

16nm
2013

11nm
2015

8nm
2017

2013 – 2017

New Intel technology generation every 2 years
Intel R&D technologies drive this pace well into the next decade

Roadmap

Research

Source: Intel

All dates, product descriptions, availability and plans are forecasts and subject to change without notice.

Intel Design & Process Cadence

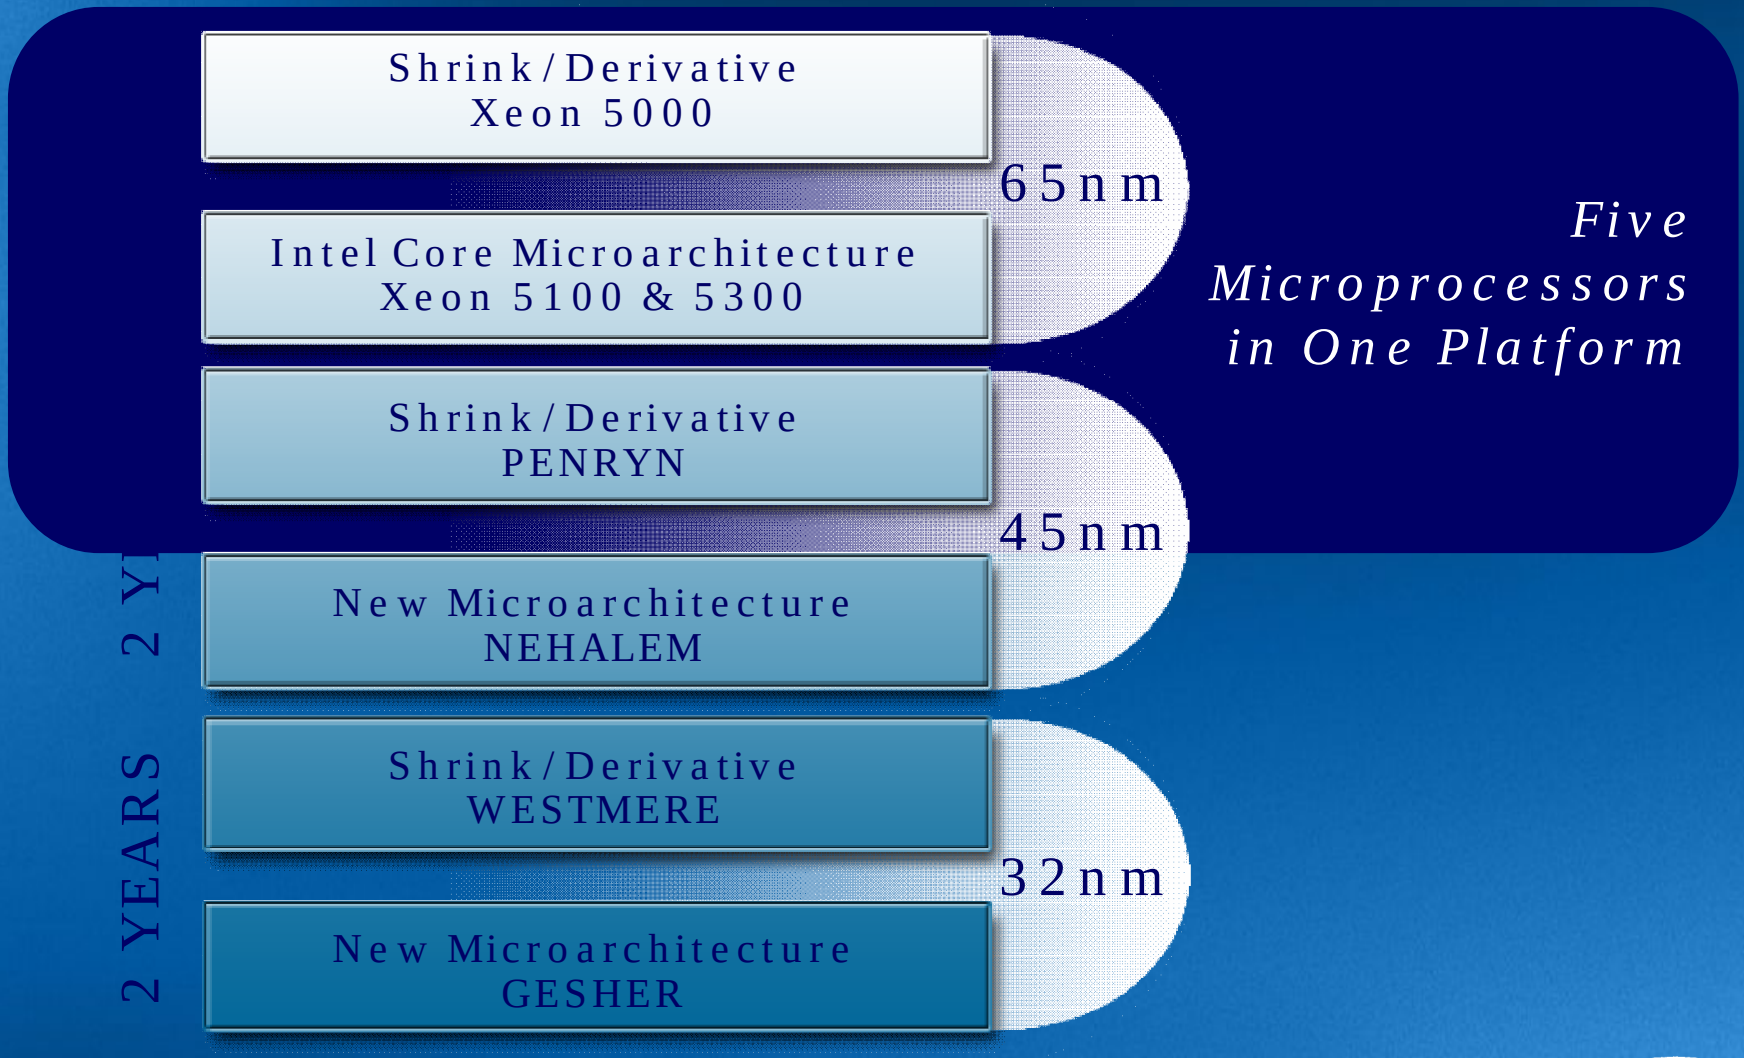

All dates, product descriptions, availability and plans are forecasts and subject to change without notice.

サーバー向けインテル® デュアル/クアッドコア・プロセッサーのロードマップ

	2006	2007	将来
インテル® Itanium® 2 9000番台	インテル® Itanium® 2 プラットフォーム インテル® Itanium® 2 プロセッサー9M デュアルコア Itanium® 2 プロセッサー 9000番台 Montvale** インテル® E8870およびサードパーティー製チップセット		Richford*プラットフォーム Tukwila** Poulson** 将来チップセット
インテl® Xeon® MP 7000 番台	インテル® Xeon® MP プラットフォーム *デュアルコア インテル® Xeon® プロセッサー 7000番台 *デュアルコア インテル® Xeon® プロセッサー 7100番台 E8500/8501 およびサードパーティー製チップセット	Caneland*プラットフォーム Tigerton** Dunnington** Clarksboro* / サードパーティー製チップセット	
インテル® Xeon® DP 5000 番台	*インテル Xeon プロセッサー E7520	インテル® Xeon® DP プラットフォーム クアッドコア インテル Xeon プロセッサー 5300 番台 デュアルコア インテル Xeon プロセッサー 5100番台 *デュアルコア インテル Xeon プロセッサー5000 番台 インテル 5000 P/V チップセット	将来プロセッサー 将来プロセッサー
インテル® UP サーバー	*インテル® Pentium® D プロセッサー E7230	インテル® Xeon® UP プラットフォーム クアッドコア インテル Xeon プロセッサー 3200番台 デュアルコア インテル Xeon プロセッサー 3000番台 インテル® 3000/3010チップセット	将来のUPプラットフォーム 将来プロセッサー 将来プロセッサー 将来チップセット

＊= Intel NetBurst® Microarchitecture
＊＊= 開発コード名

4+ Cores

Where We've Been

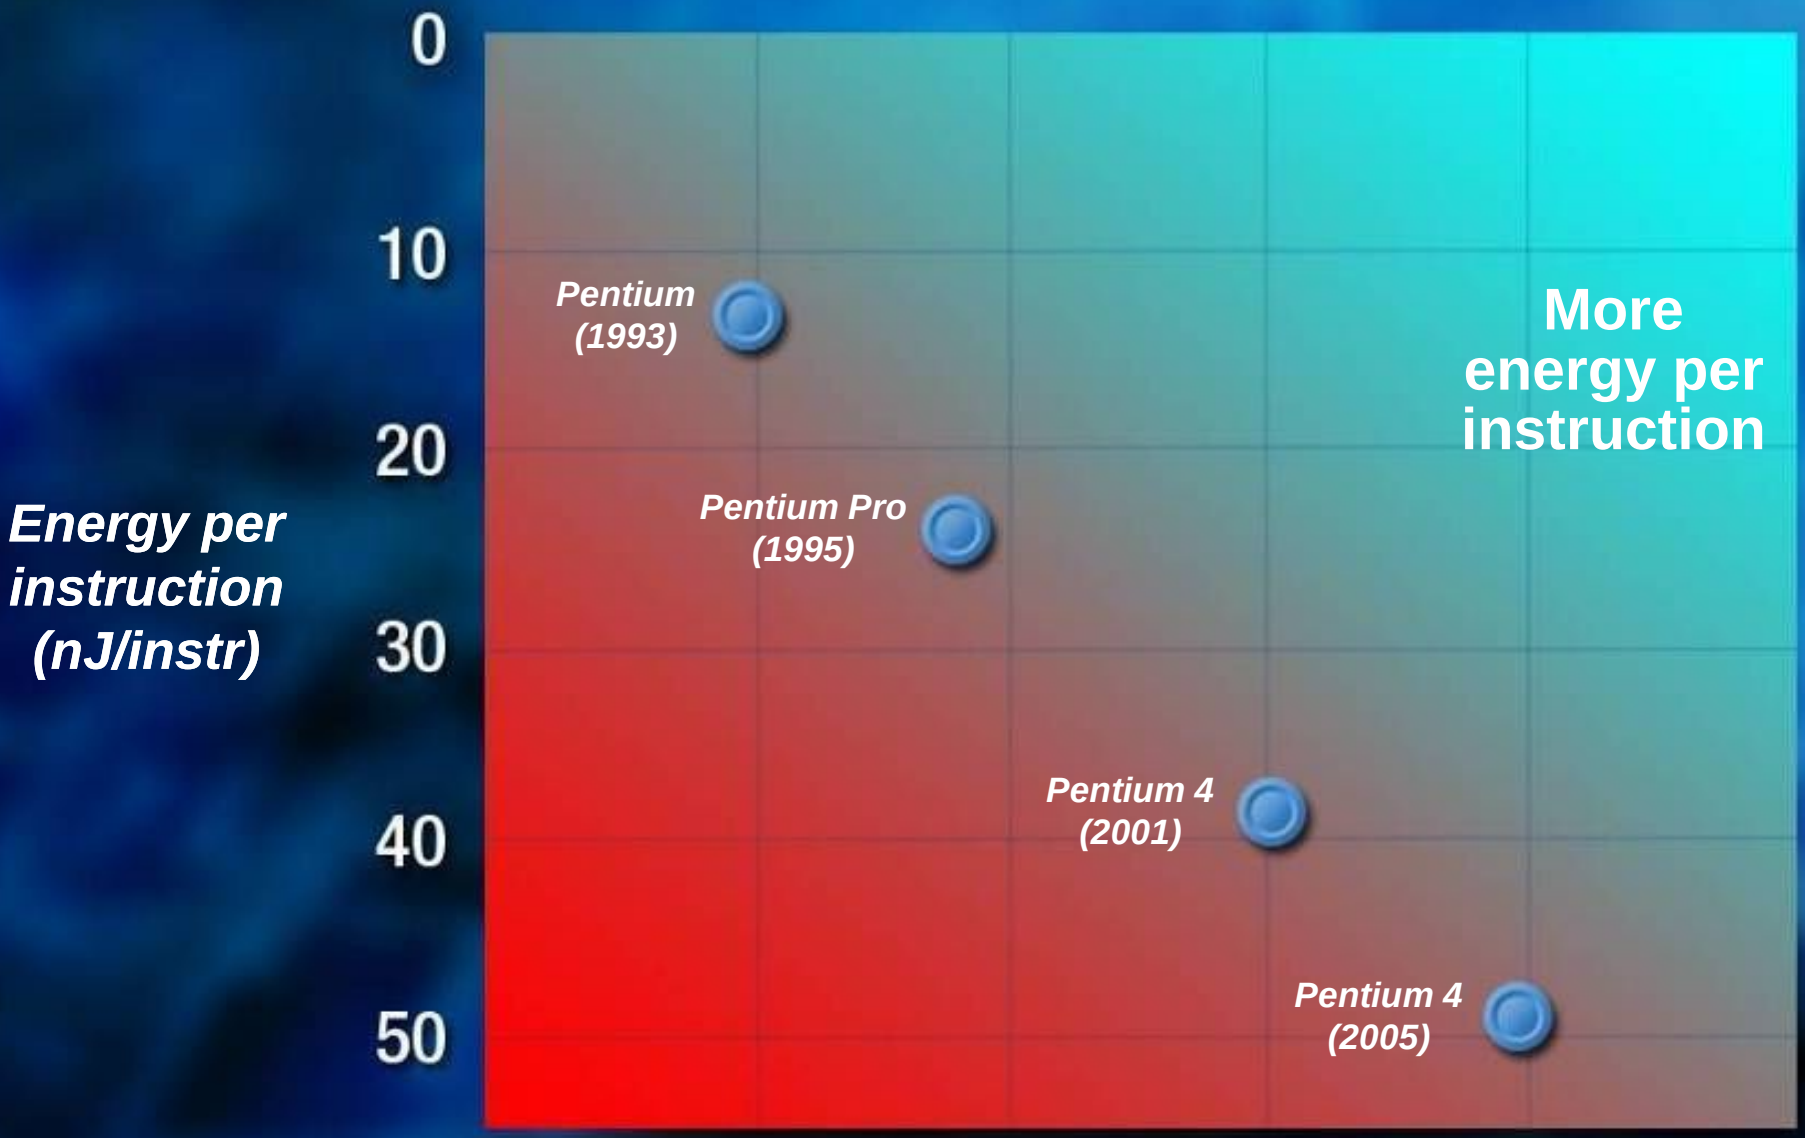

Source: Intel

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. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations (http://www.intel.com/performance/resources/limits.htm).

Where We're Going

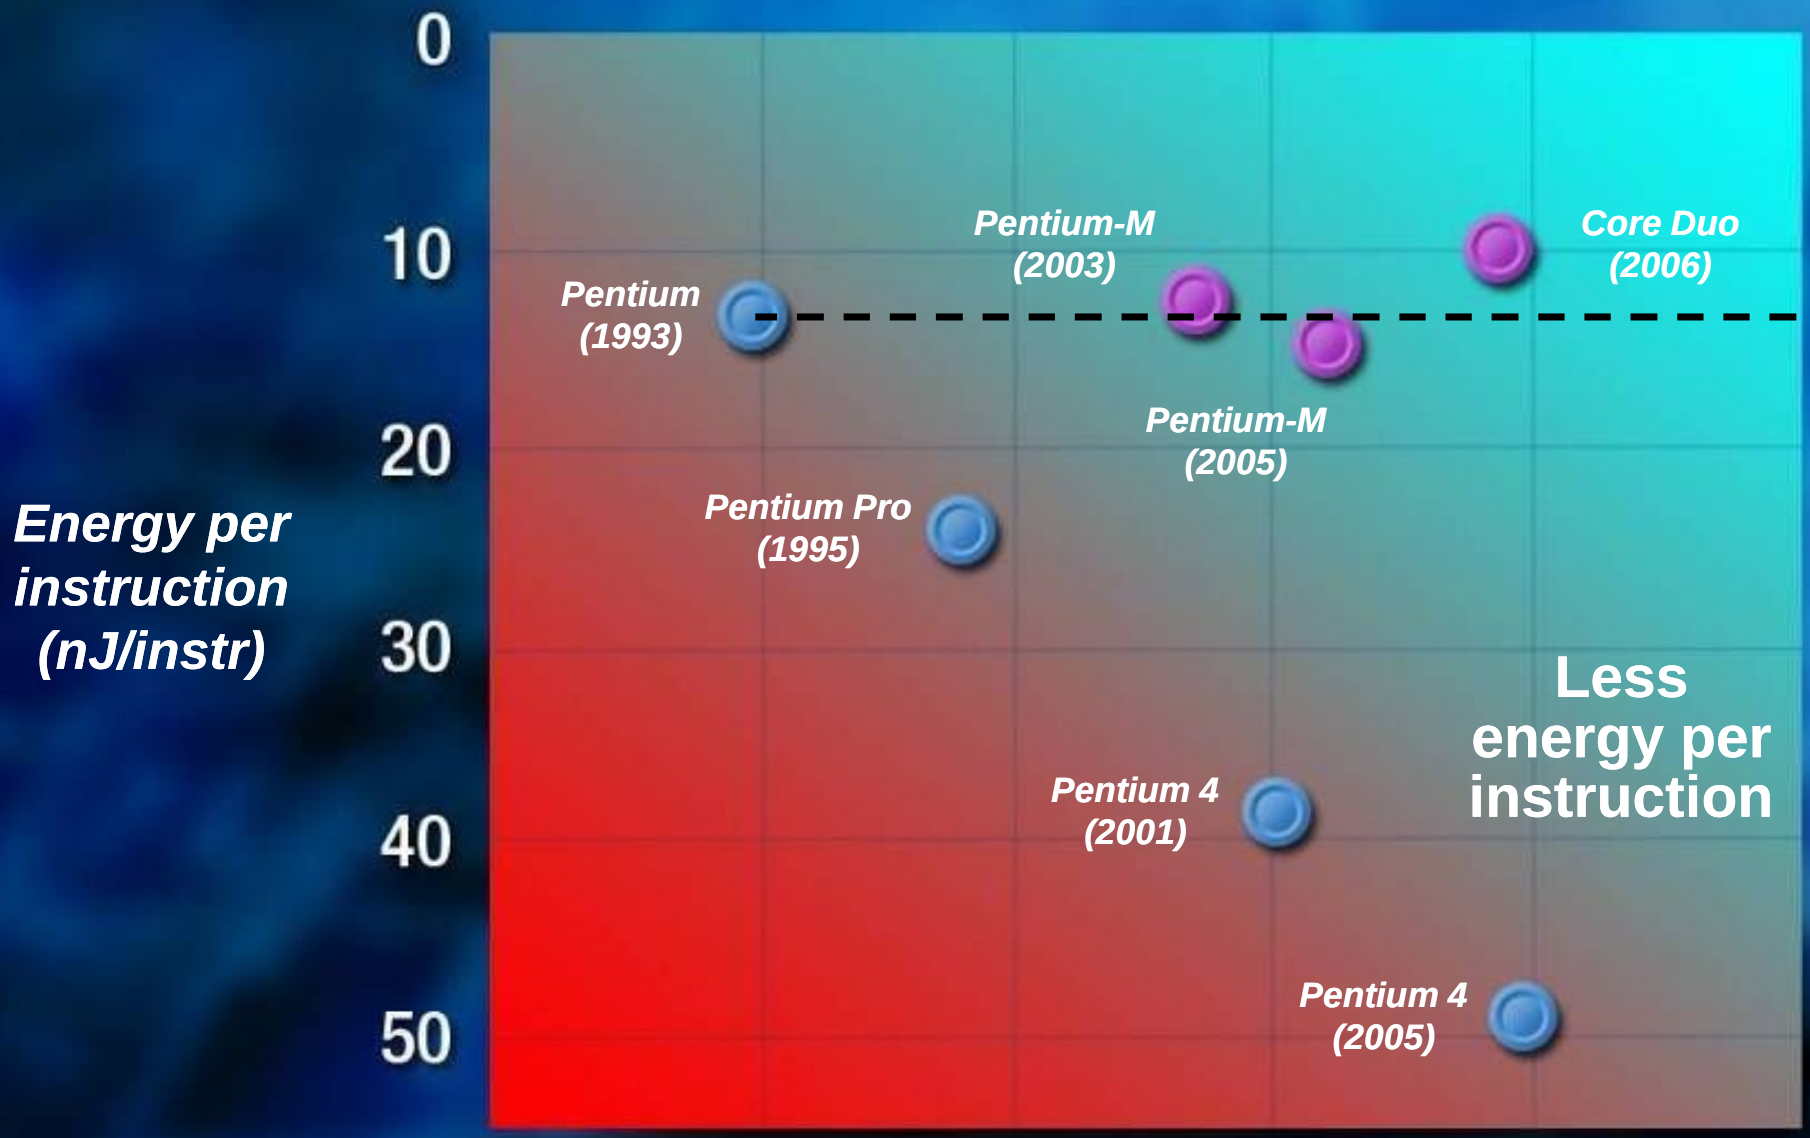

Source: Intel

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. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations (http://www.intel.com/performance/resources/limits.htm).

Why Multi-Core?

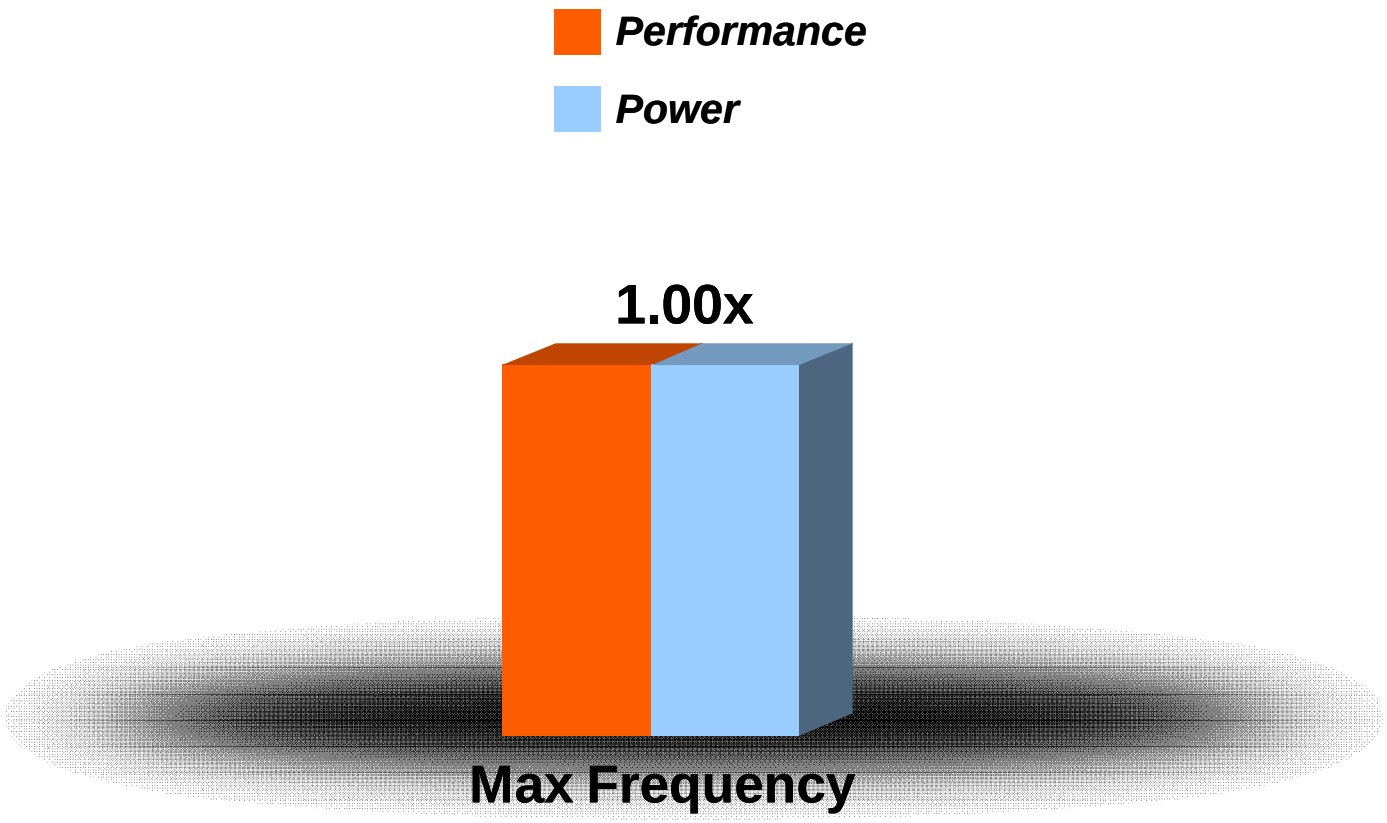

Relative single-core frequency and Vcc

Source: Intel

Over-clocking

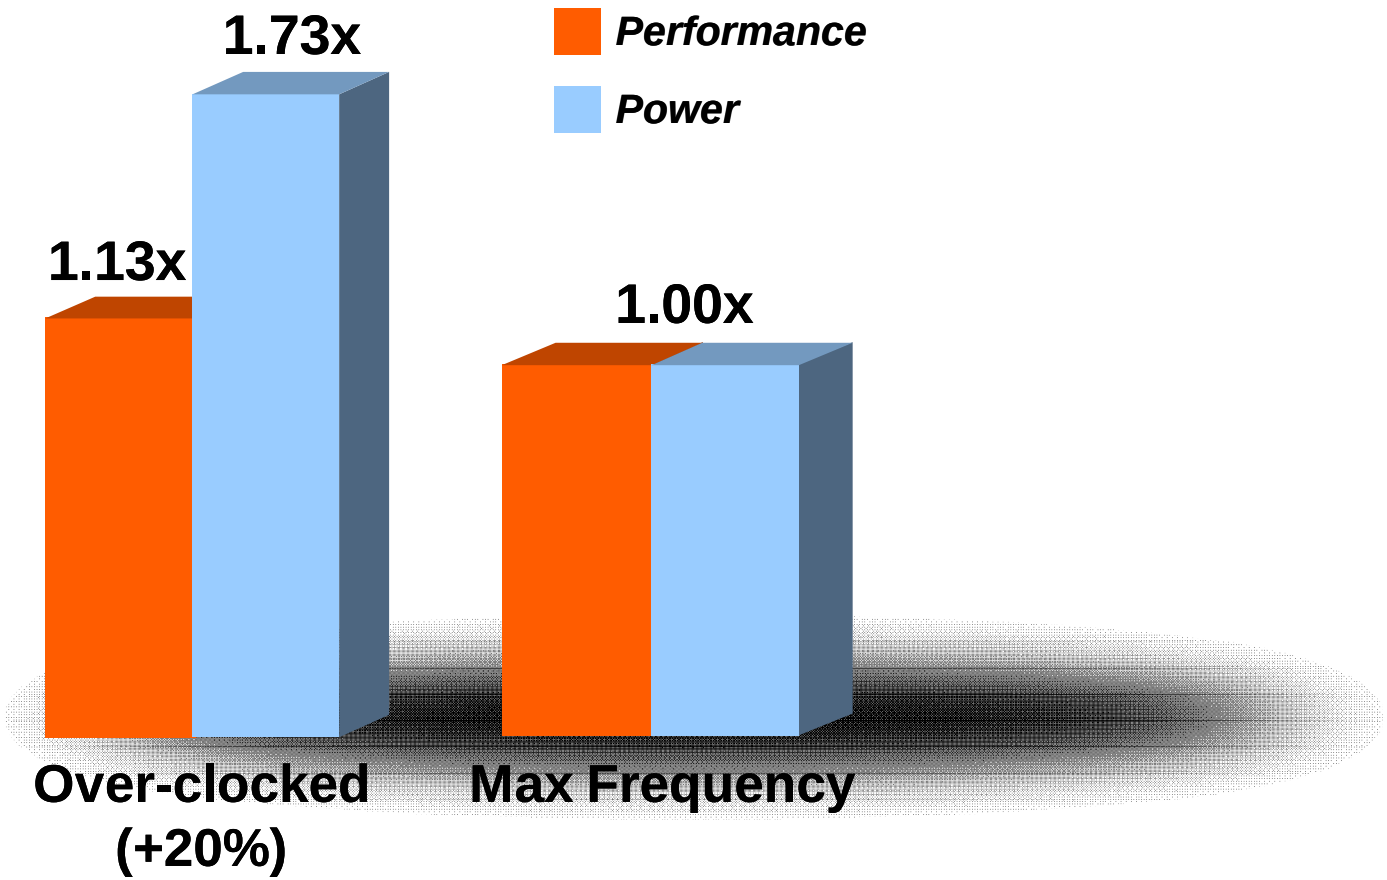

Relative single-core frequency and Vcc

Source: Intel

Under-clocking

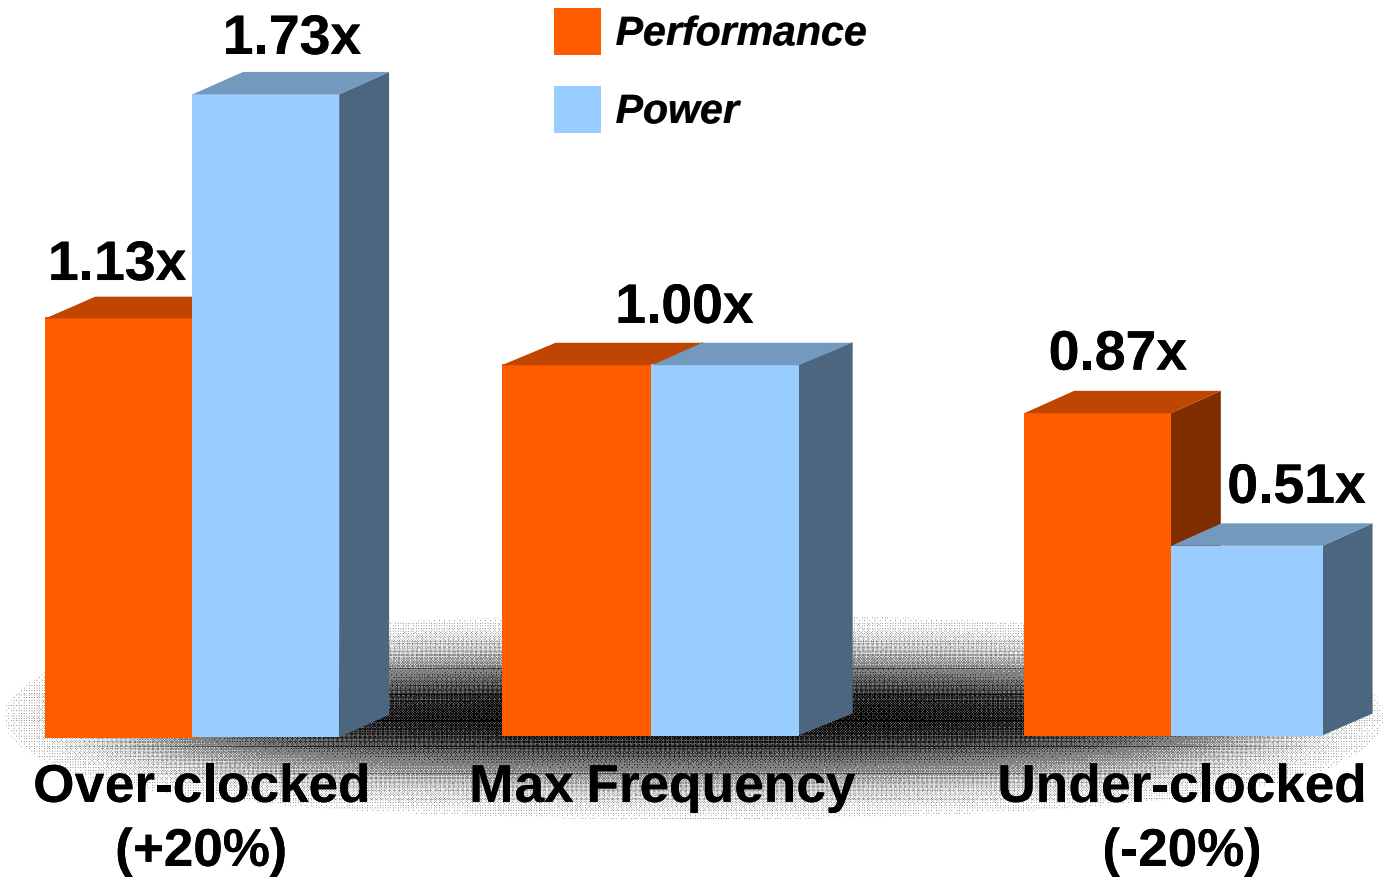

Relative single-core frequency and Vcc

Source: Intel

Multi-Core
Energy-Efficient Performance

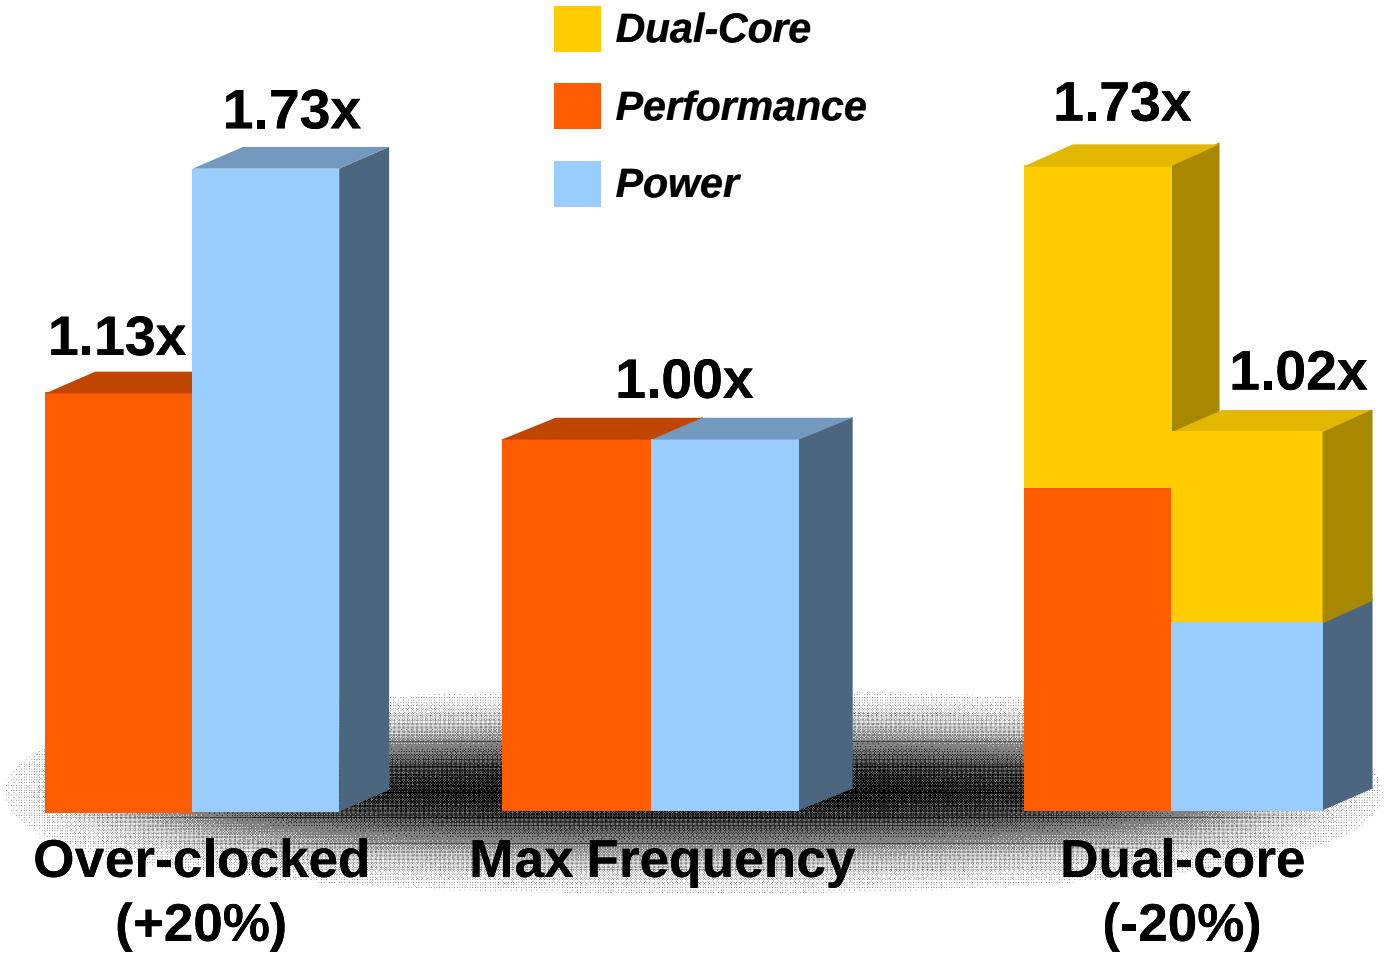

Relative single-core frequency and Vcc

Source: Intel

Intel Server Strategy: Power & Thermals

Silicon Advances

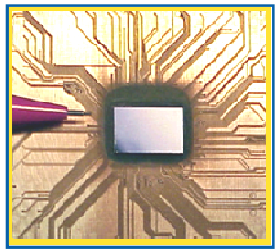

- Process Technologies
- Microarchitecture
- Packaging

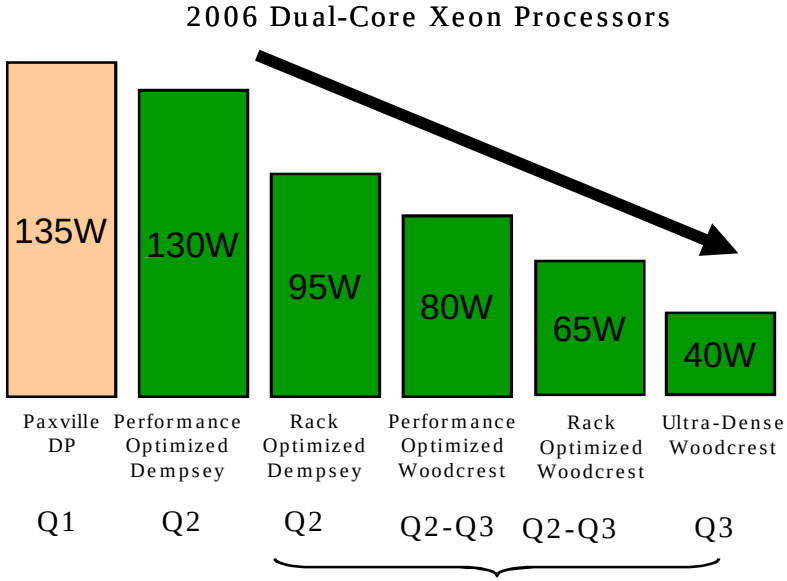

Intel Power Tools

- Up to 24% Energy Savings with Demand Based Switching
- Up to 57% Improved Density with Power Toolkit
- Datacenter Framework

Platform & Architectural Advances

- Multi-Core Processors
- Low Power, High Speed Memory
- Platform / Architectural Flexibility
- Enhanced Utilization (Virtualization)
- Software Optimization

It's MUCH More Than Just the Processor

The World's First IA Quad Core Processor

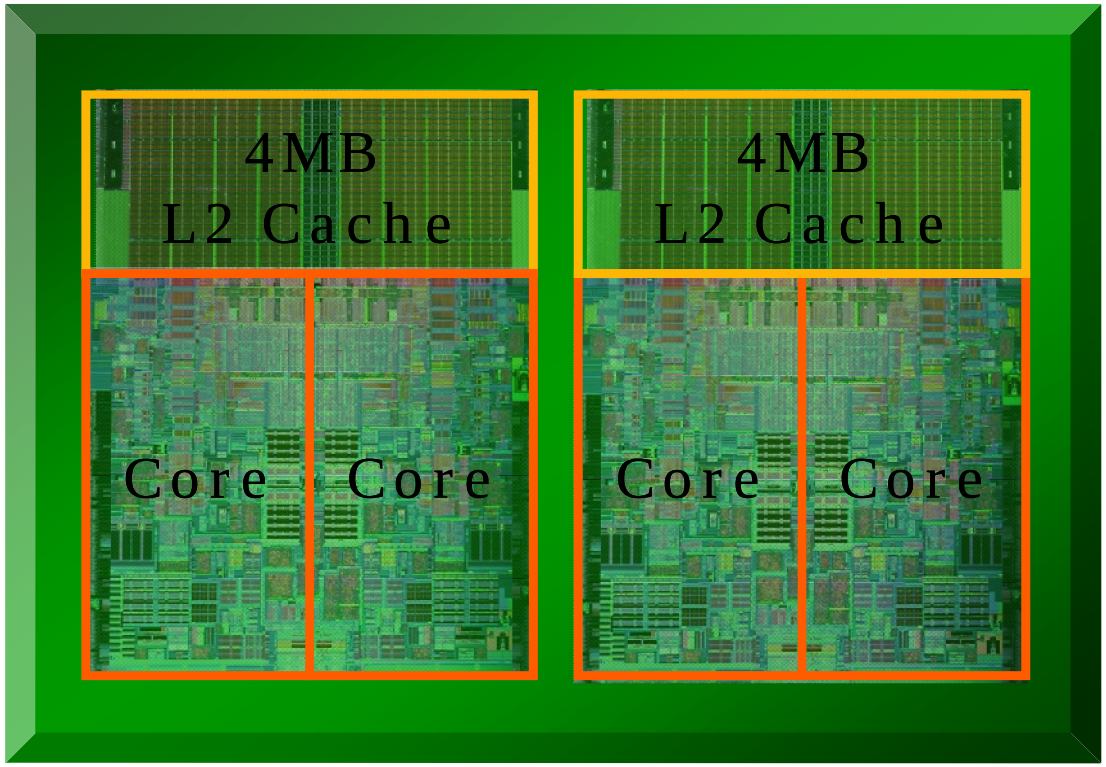

Cache

Die Selection

Compatibility

Cost

Capacity

Customers

1066/1333/1600 MHz

Exceptional Performance And Energy Efficiency

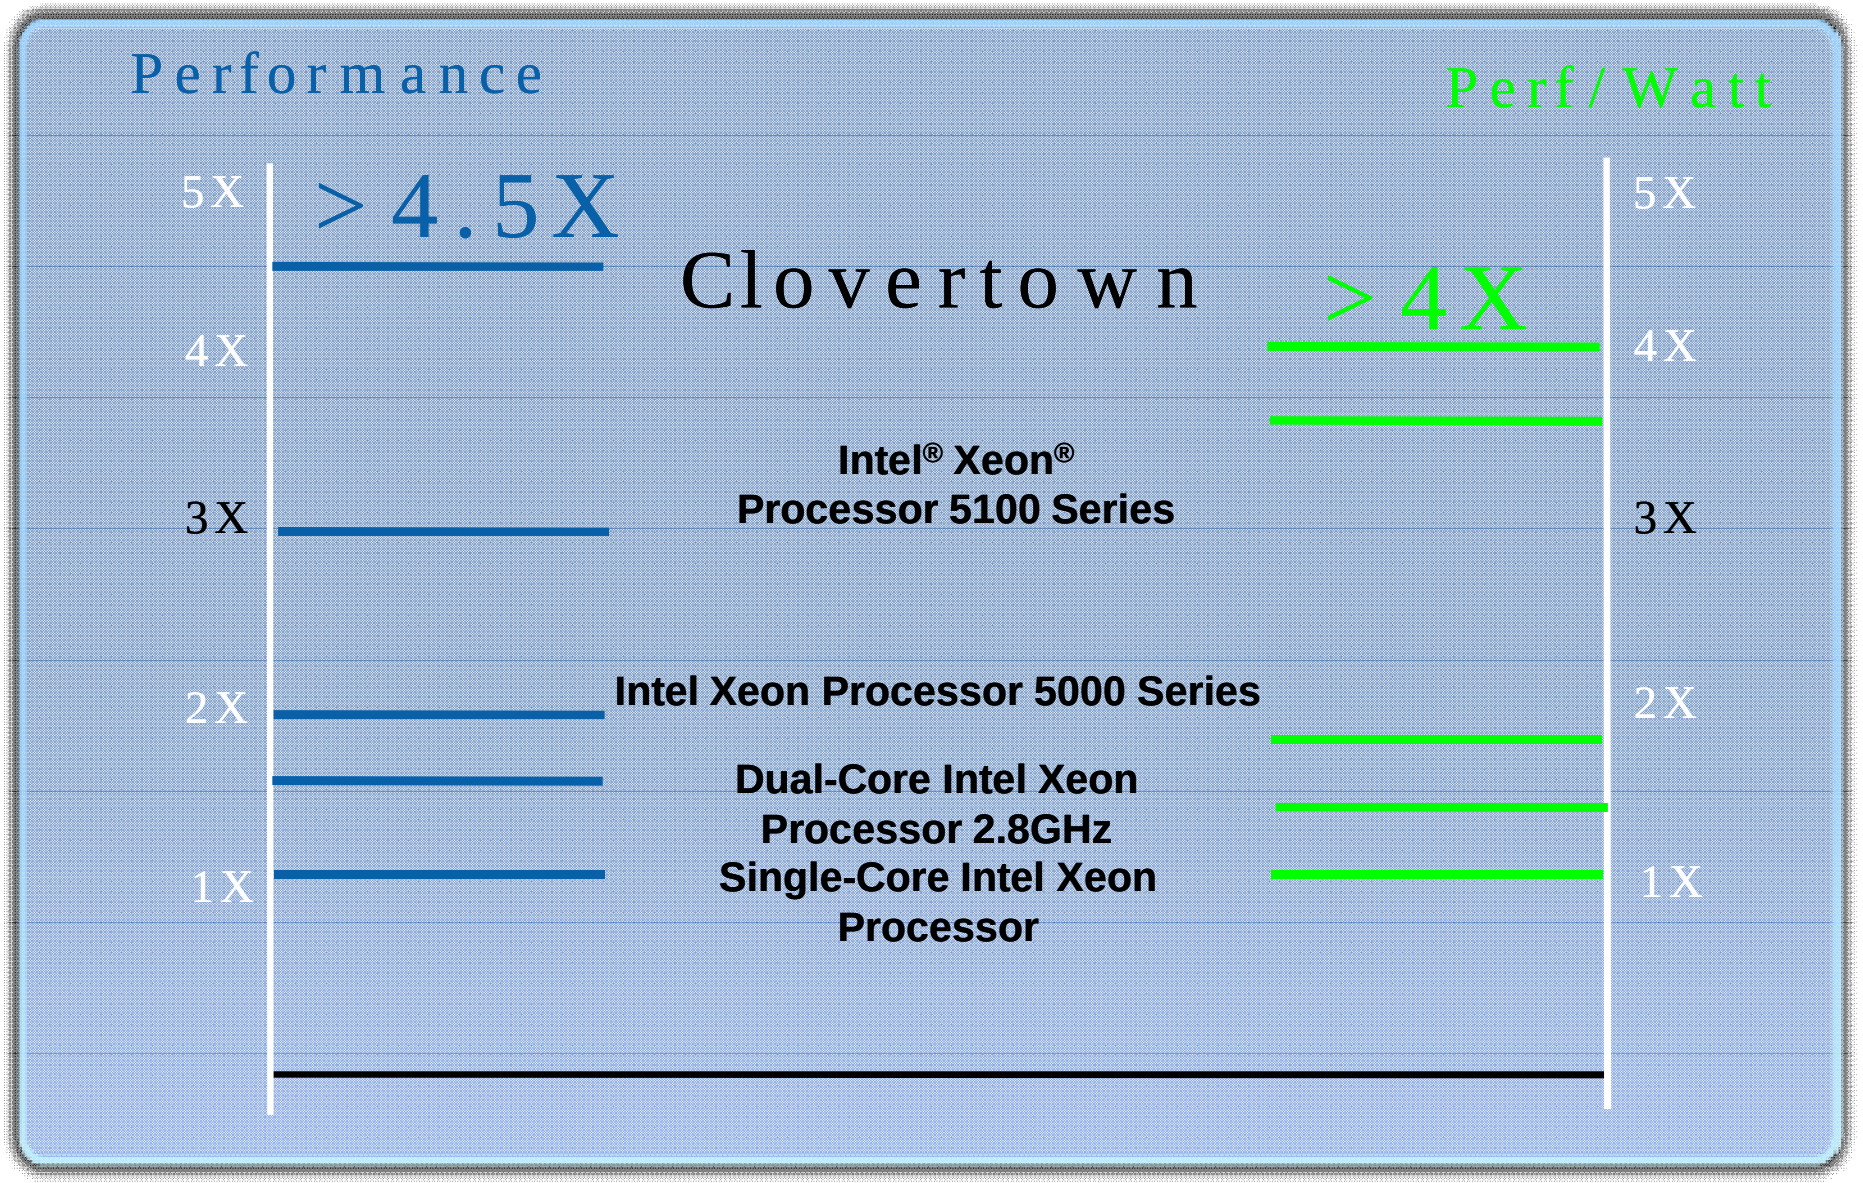

Based on published/measured SPECint_rate_base2000. Intel Xeon Processor 3.60 GHz 4GB; Dual-Core Intel Xeon Processor 5160 3.00GHz 8GB; Clovertown 2.67 GHz 8GB. Perf/Watt based on SPECint_rate_base2000 benchmark w/system power measurements. Intel internal data. Published AMD Opteron 2220 SE (2.8GHz, socket F).spec.org. Other names and brands may be claimed as the property of others

Announcing the Next Generation

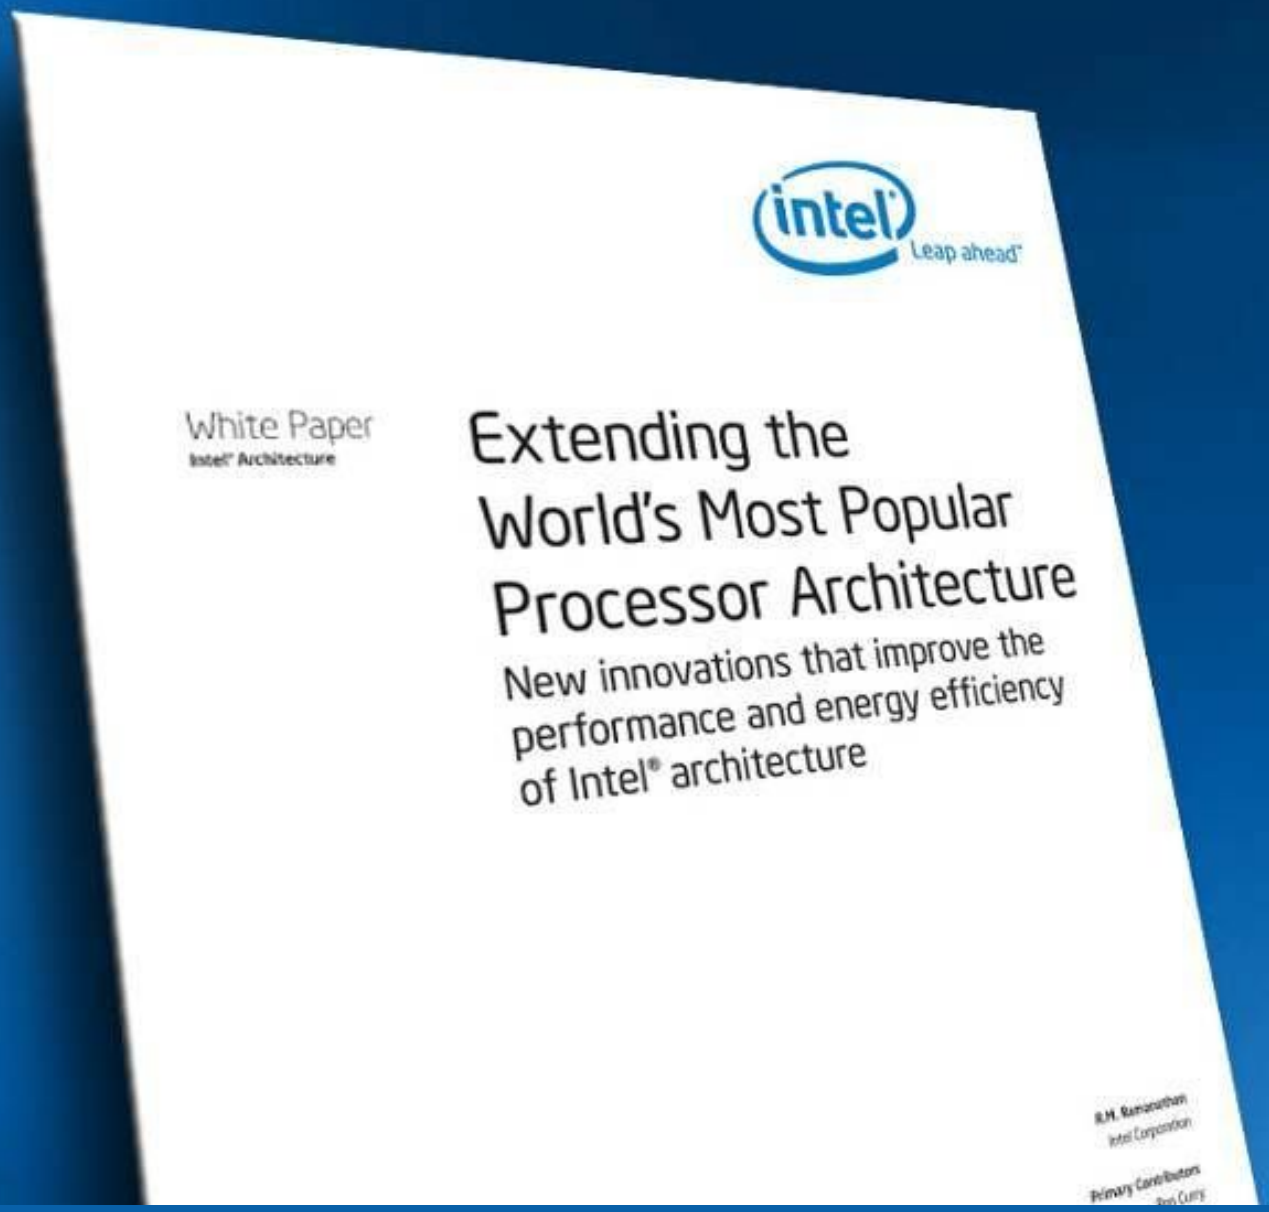

New Instructions

Vectorizing Compiler

Media

String and Text Processing

Application Targeted Accelerators

テラ・フロップのパフォーマンス

80 コア
1 TFLOP@**98** W
256 GB/s bisection

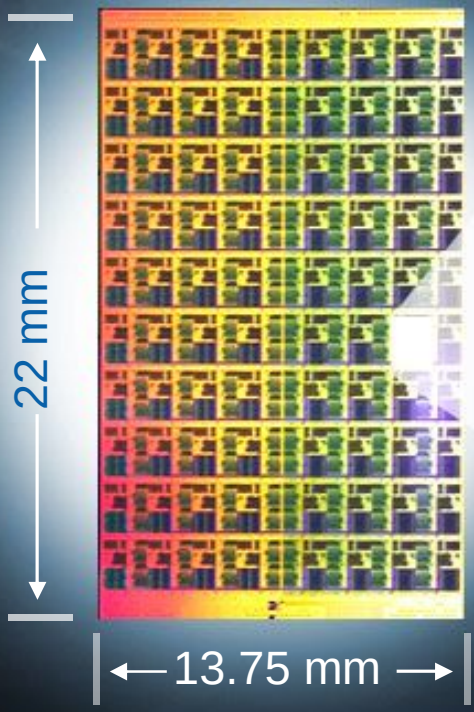

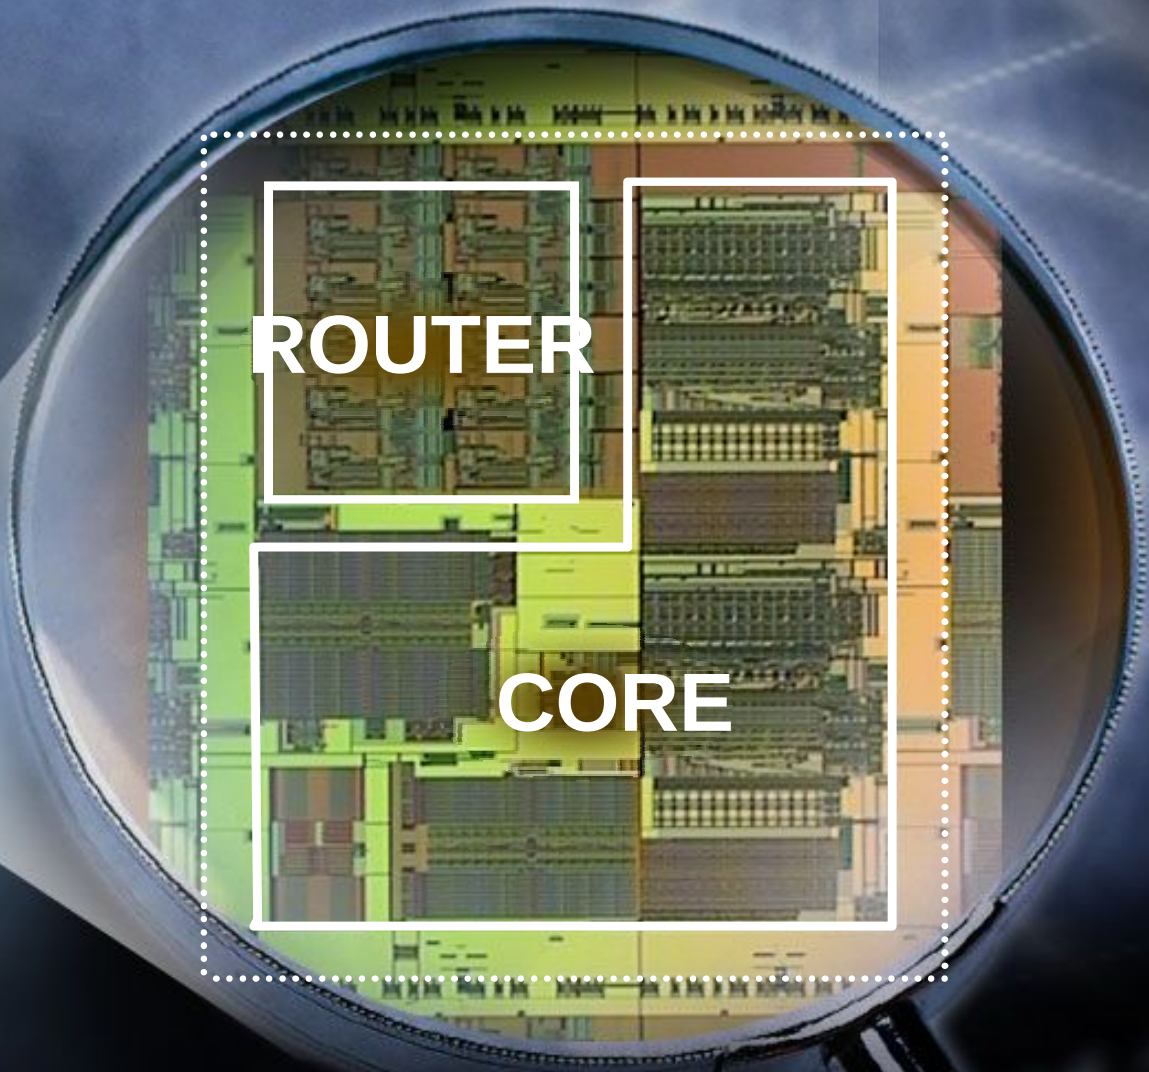

Source: Intel

Processor Performance

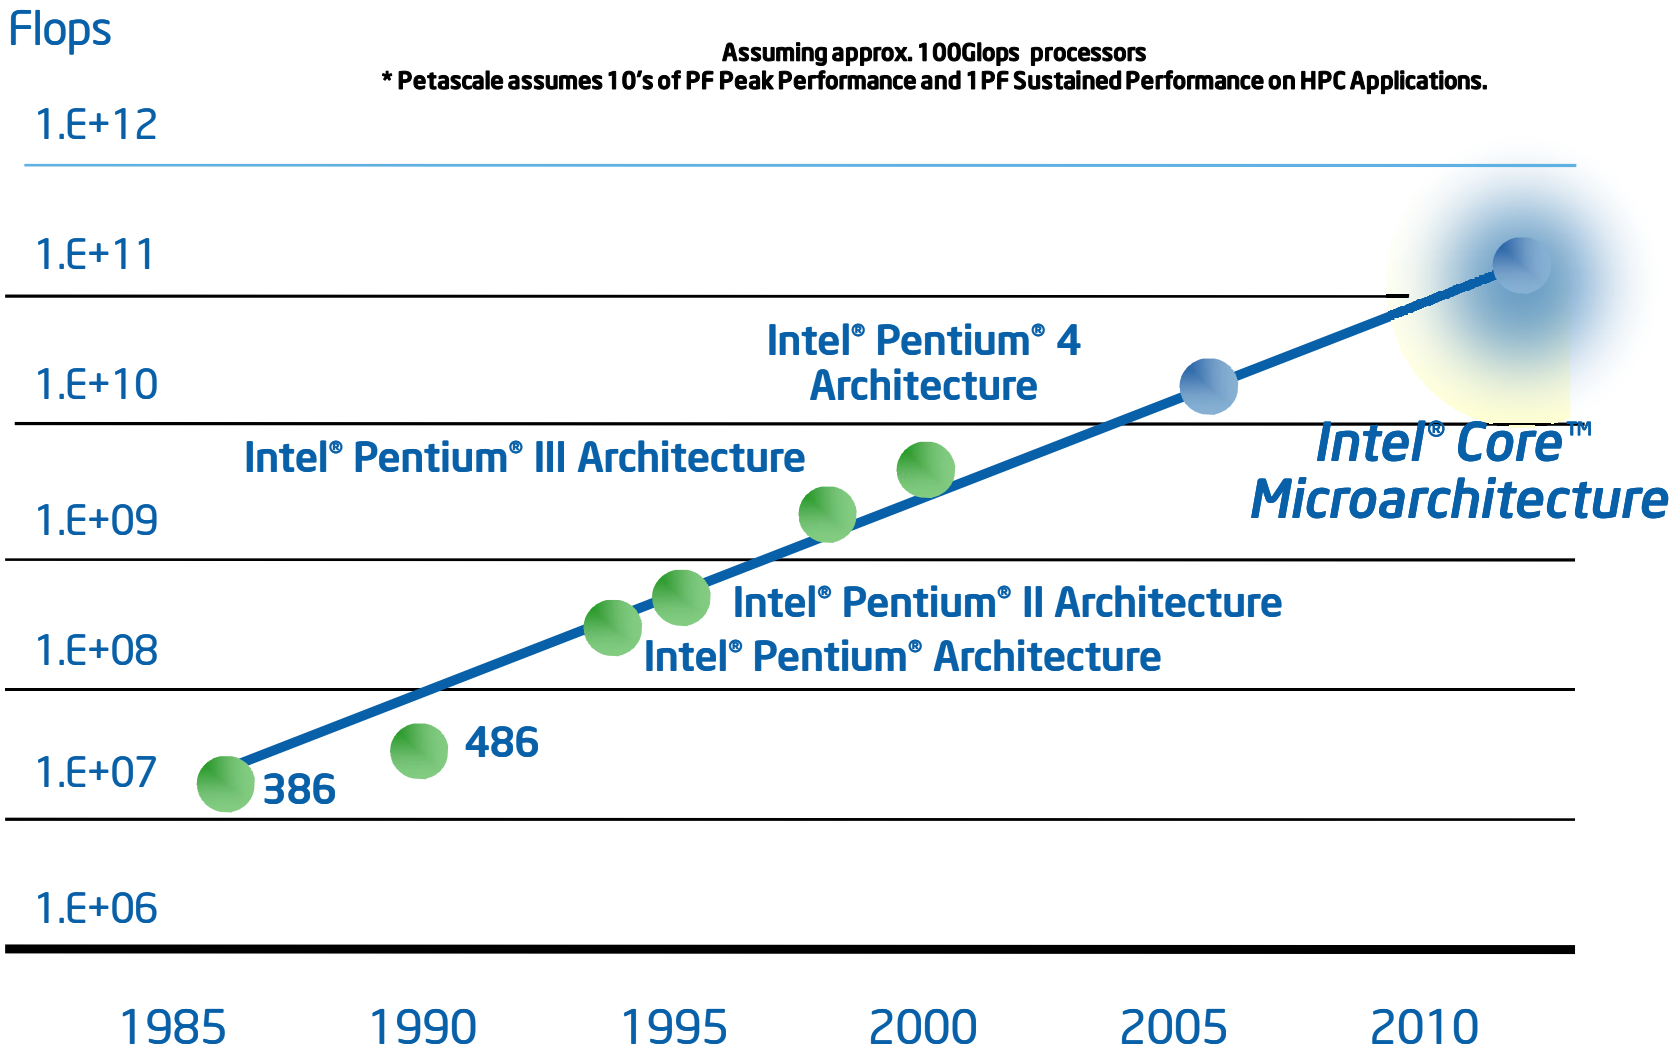

Reaching Petascale with ~100,000 Processors in 2010*

Source: Intel

Multi-threaded Cores

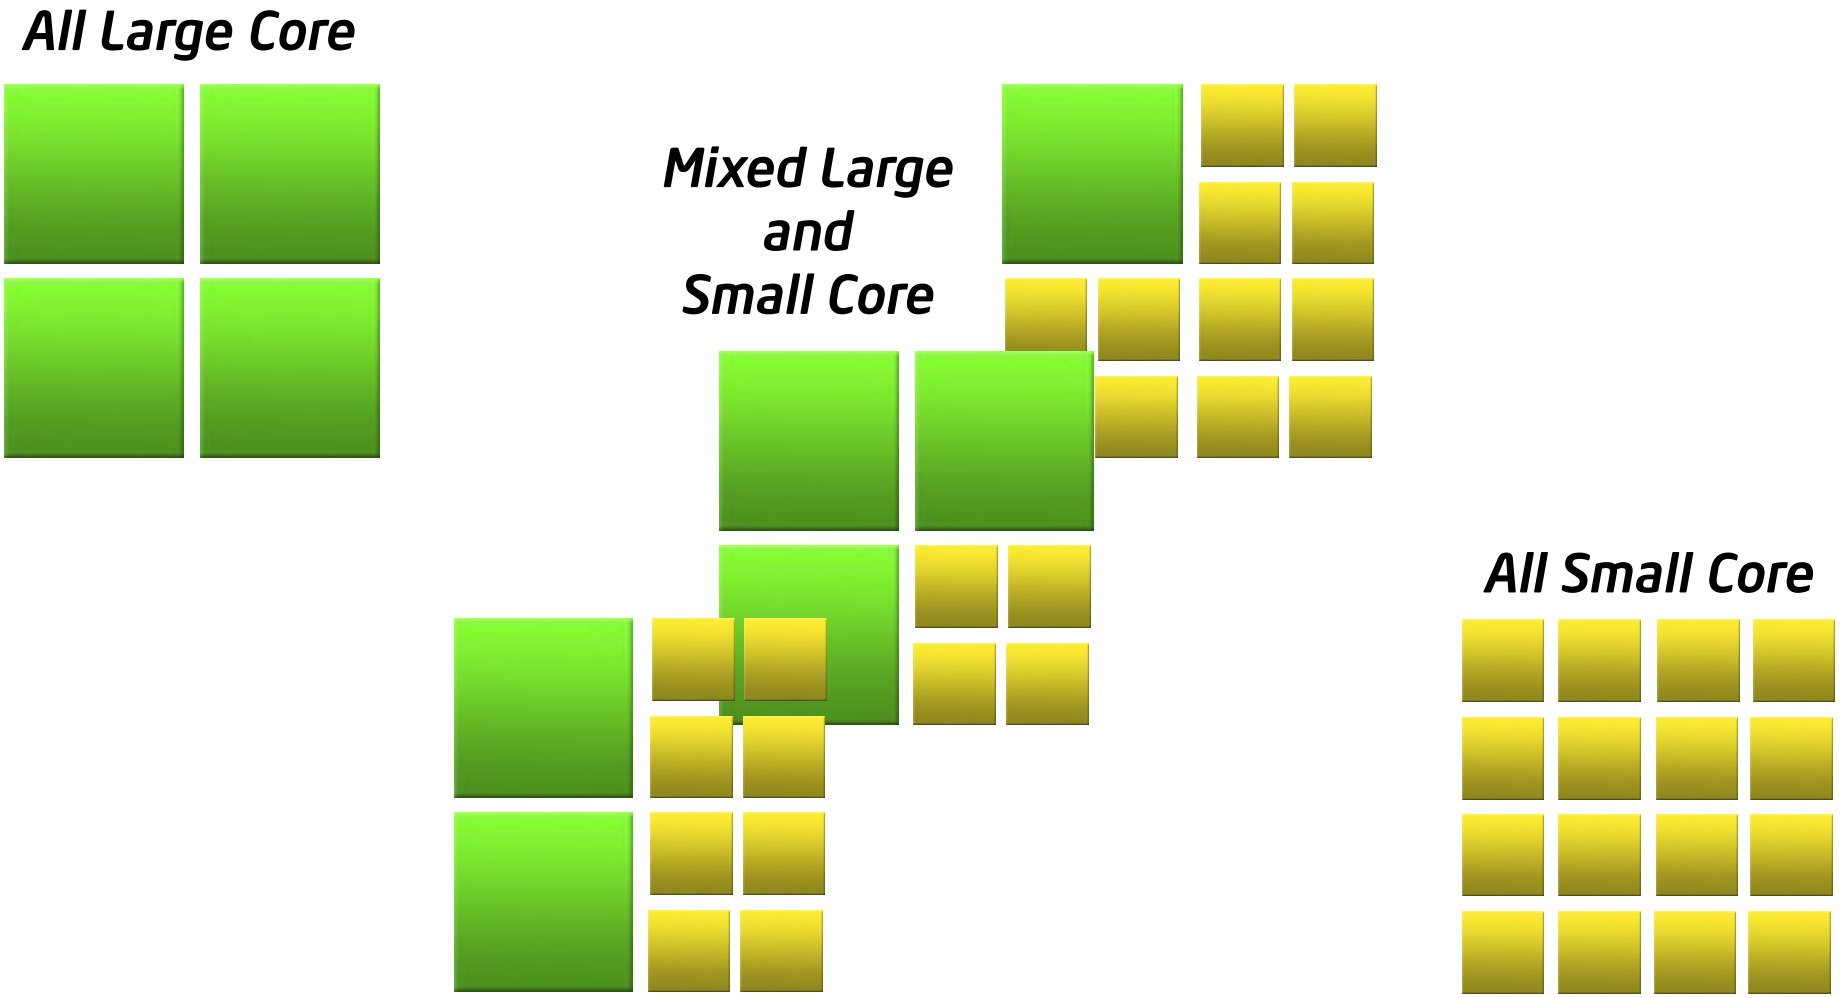

Energy Efficient Petascale with Multi-threaded Cores

Note: the above pictures don't represent any current or future Intel products

Inter-chip Interconnect Challenges Bandwidth, Link Bandwidth and Power

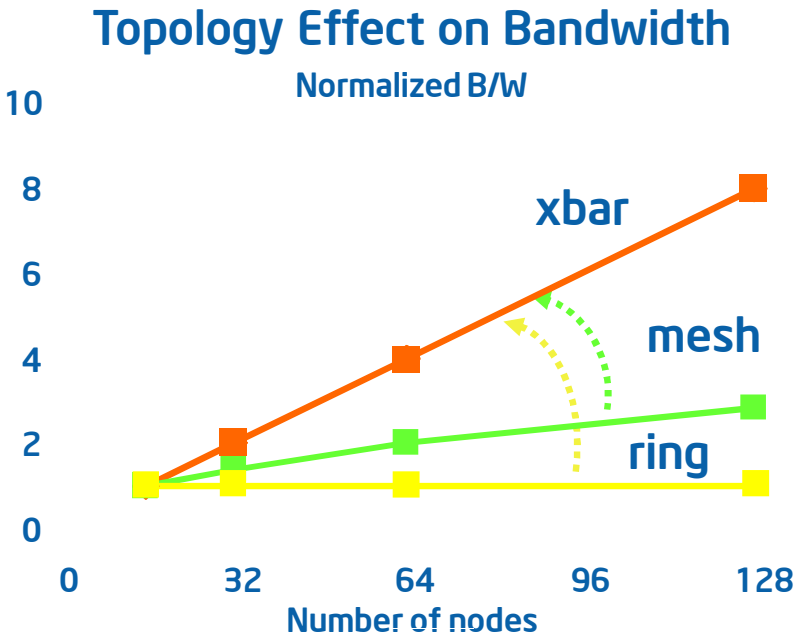

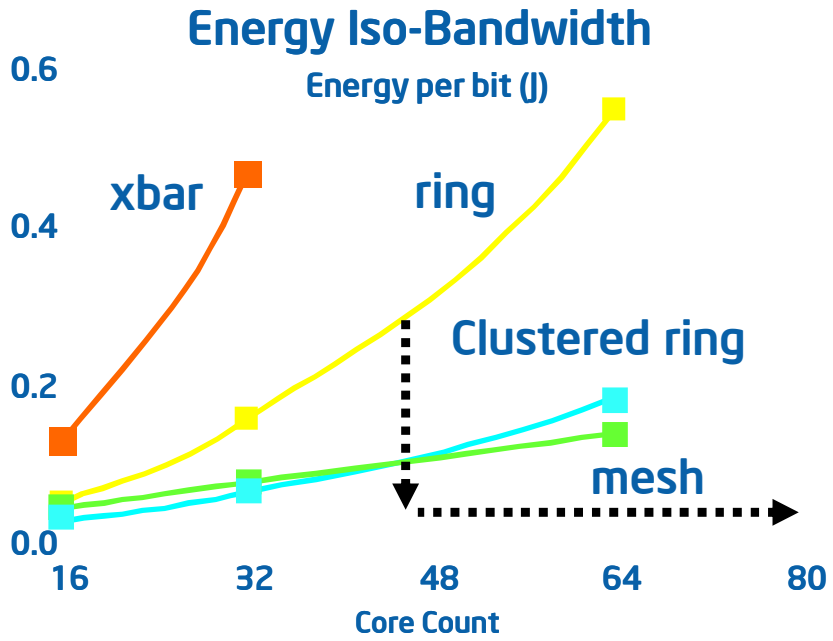

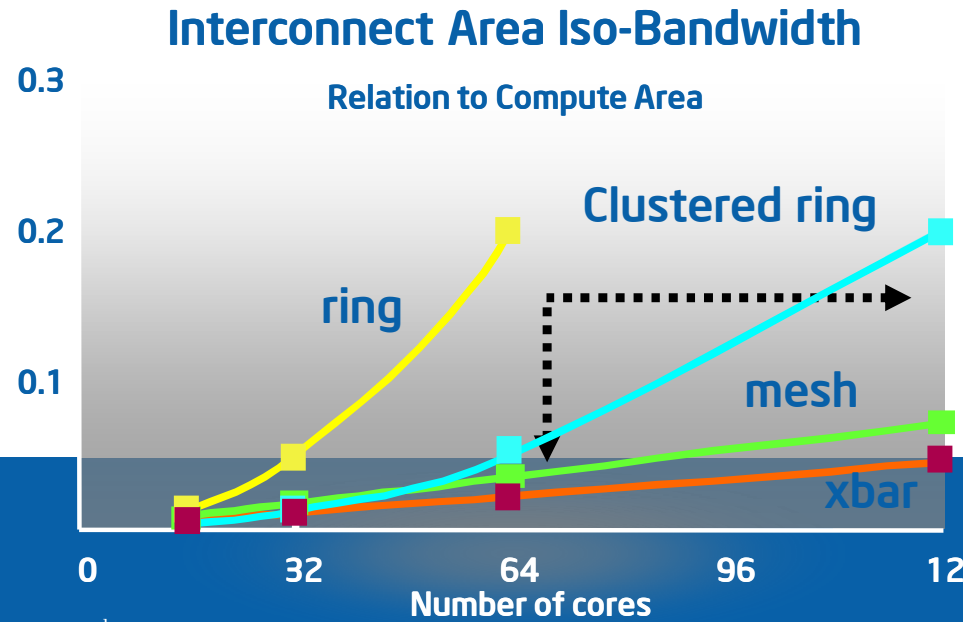

Source: Intel

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. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations (http://www.intel.com/performance/resources/limits.htm).

Increasing Processor Performance

Through Multi-threaded Cores

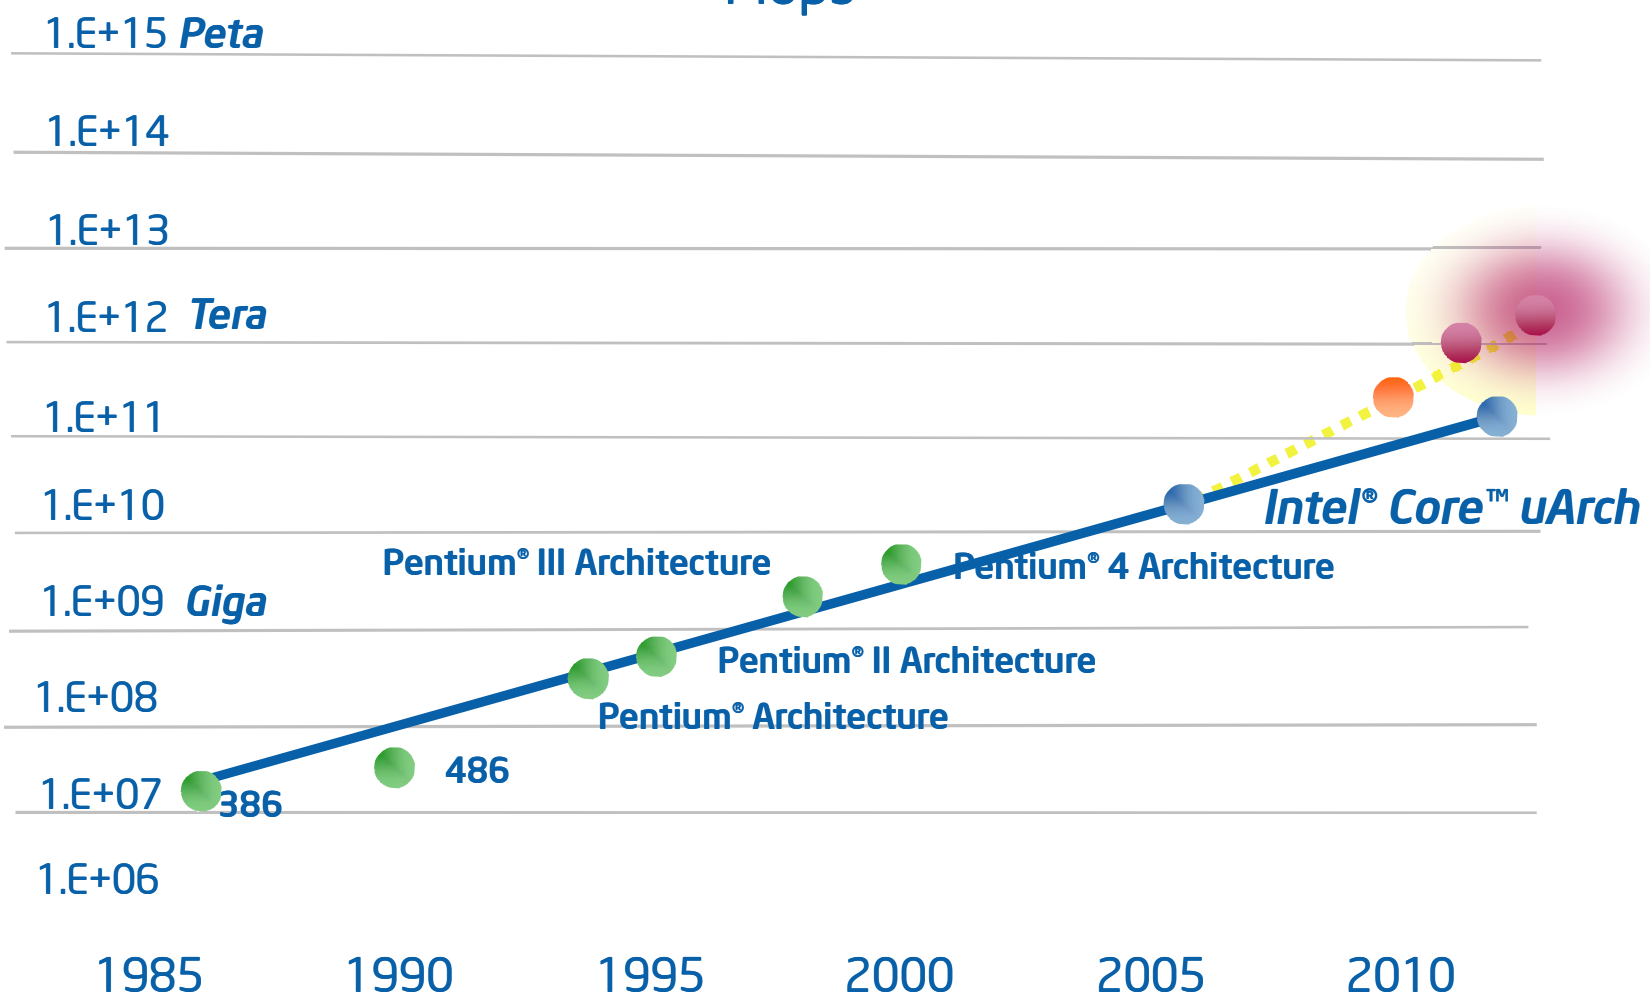

Source: Intel

Reaching Petascale with ~5,000 Processors

Inter-Chip Interconnect Performance Today

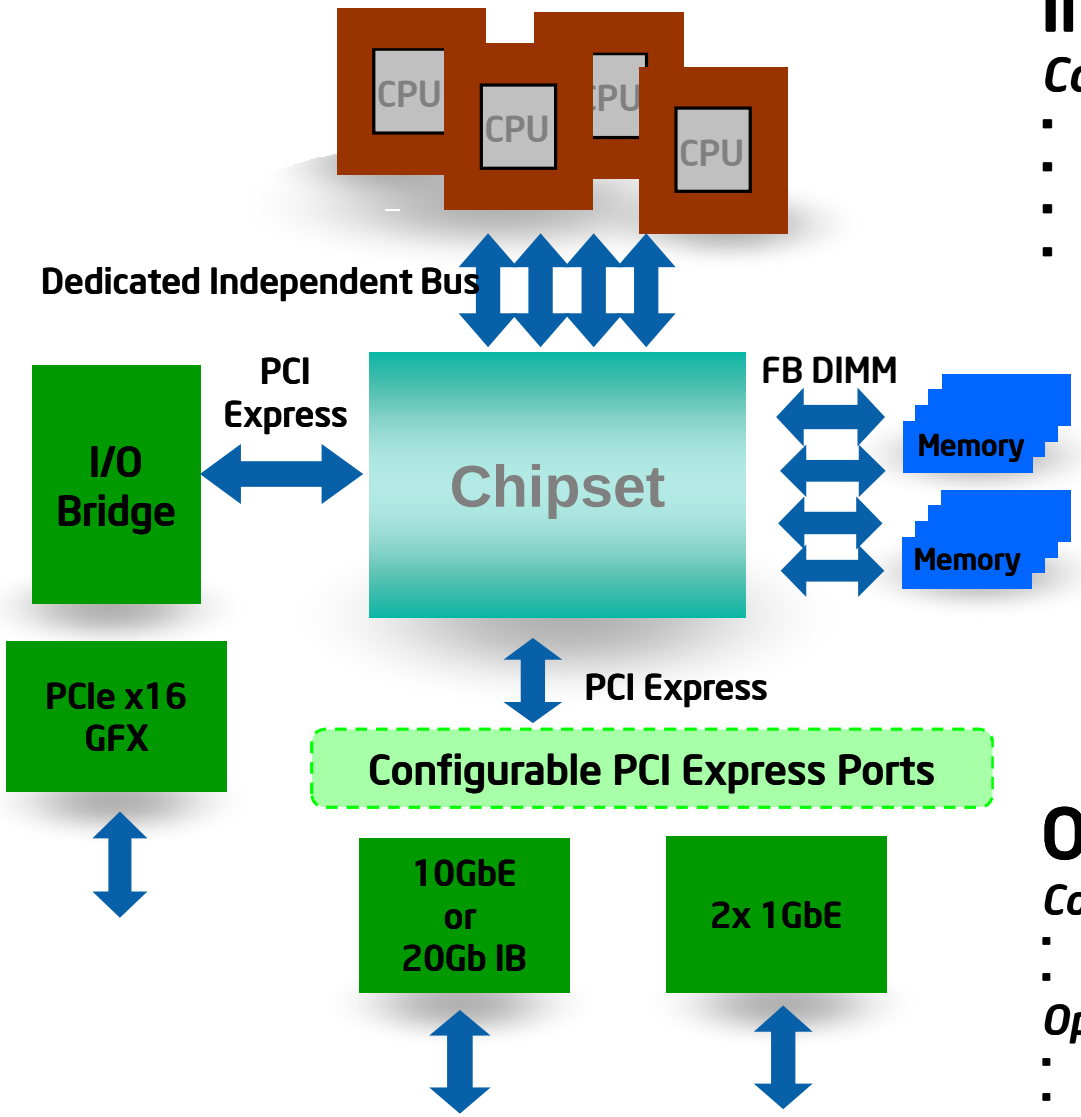

In the Box

Copper-based links:

- FB DIMM – 4GB/s per DIMM channel
- PCI express, gen. 1 – 2.5Gb/s
- PCI express, gen. 2 – 5Gb/s
- Intel Front-side Bus – 17GB/s

Out of the Box

Copper-based links:

- Infiniband DDR x4 - 20 Gb/s
- 1G Ethernet – 1Gb/s

Optical-based links:

- OC-192 (long haul optical) – 10Gb/s
- 10G Ethernet – 10Gb/s

Source: Intel

Silicon Photonics

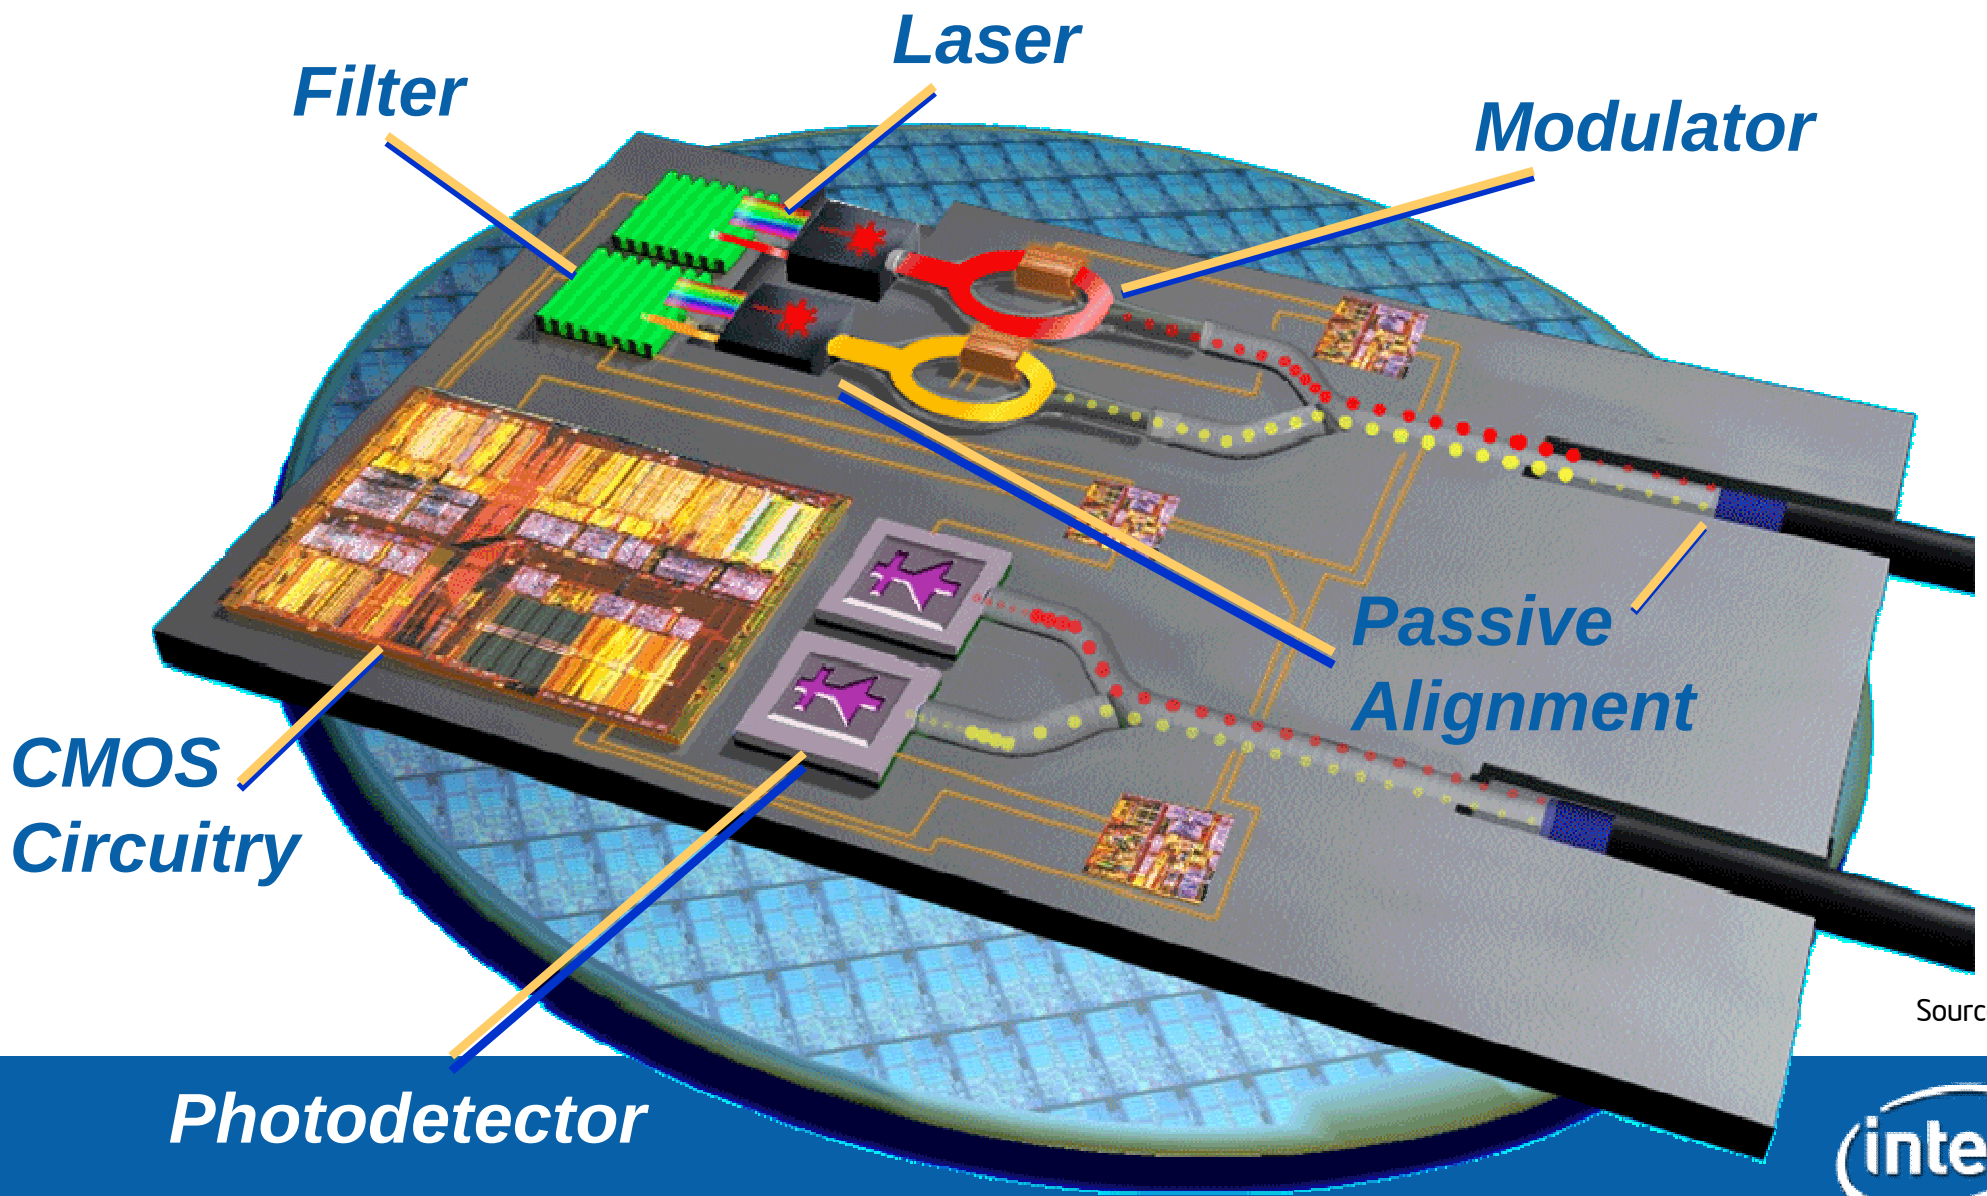

Source: Intel

Memory Performance for Balanced Computing

Byte : Flop Ratio Has Been Consistent and Steady

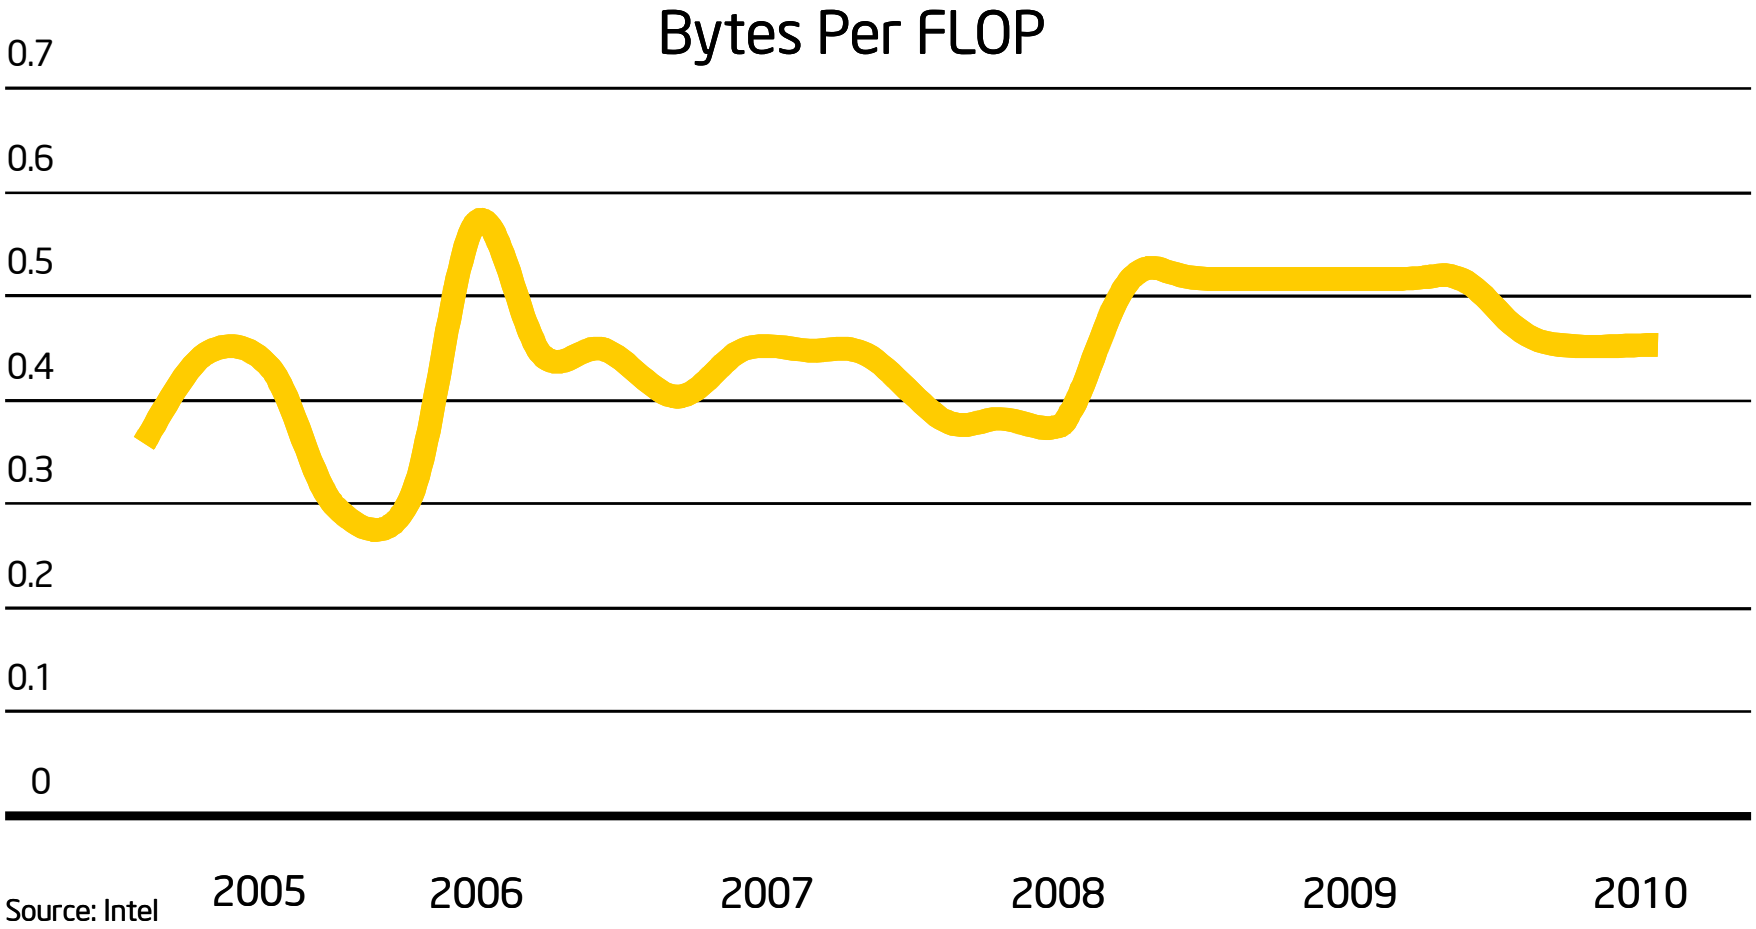

Source: Intel

Source: Intel

Continuing the Trend for Petascale Performance

Increasing Memory Bandwidth *to Keep Pace*

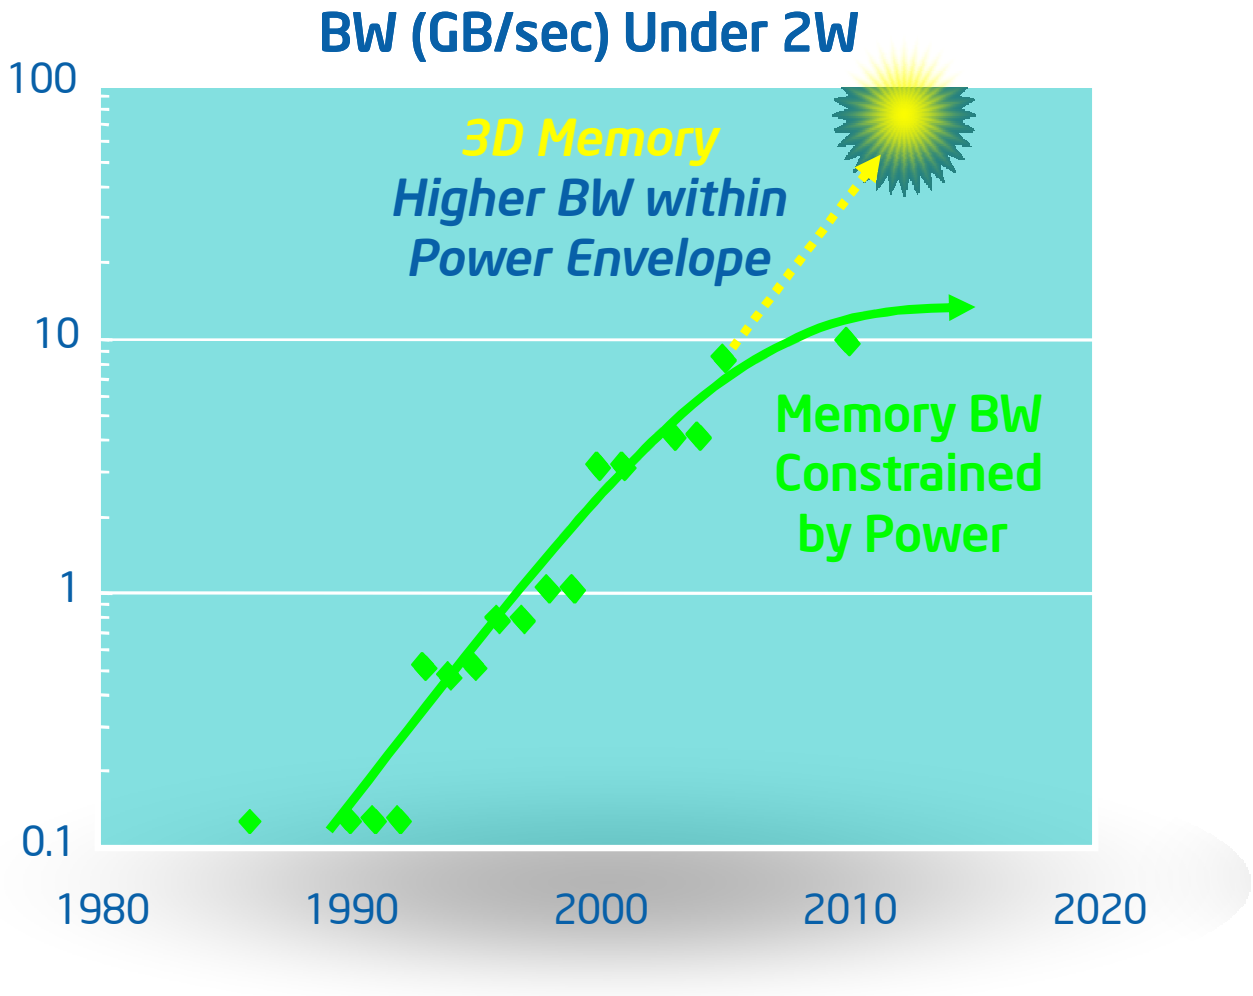

3D Memory Stacking

Power and IO Signals Go Through DRAM to CPU

Thin DRAM Die

Through DRAM Vias

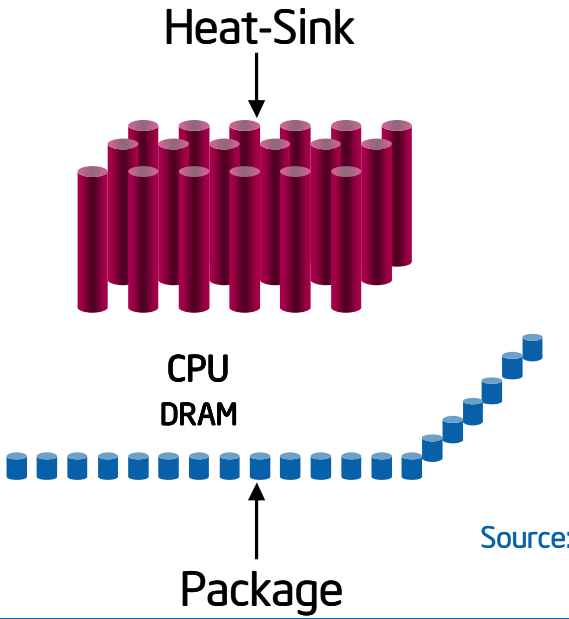

Source: Intel

PCI Express to Meet I/O Demand

Performance, Bandwidth and Functionality

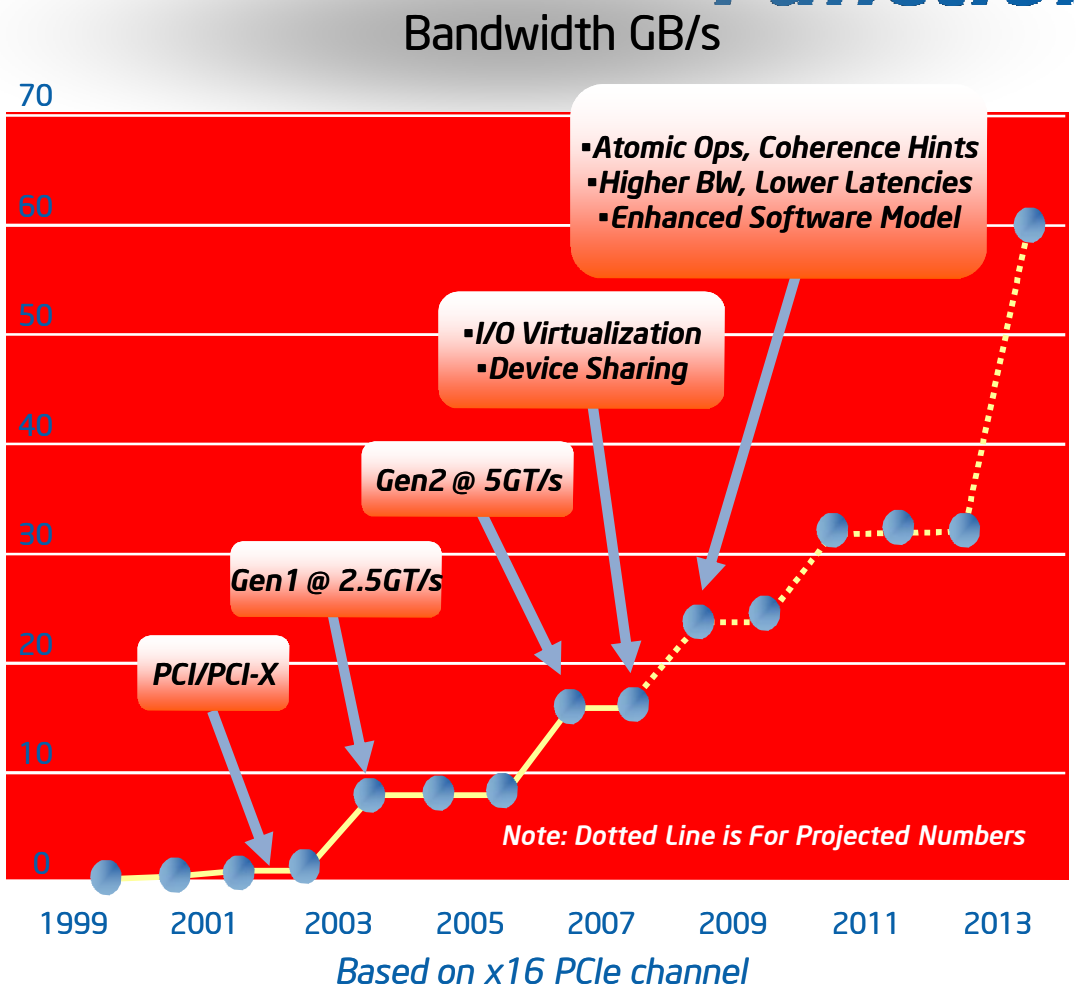

Based on x16 PCIe channel

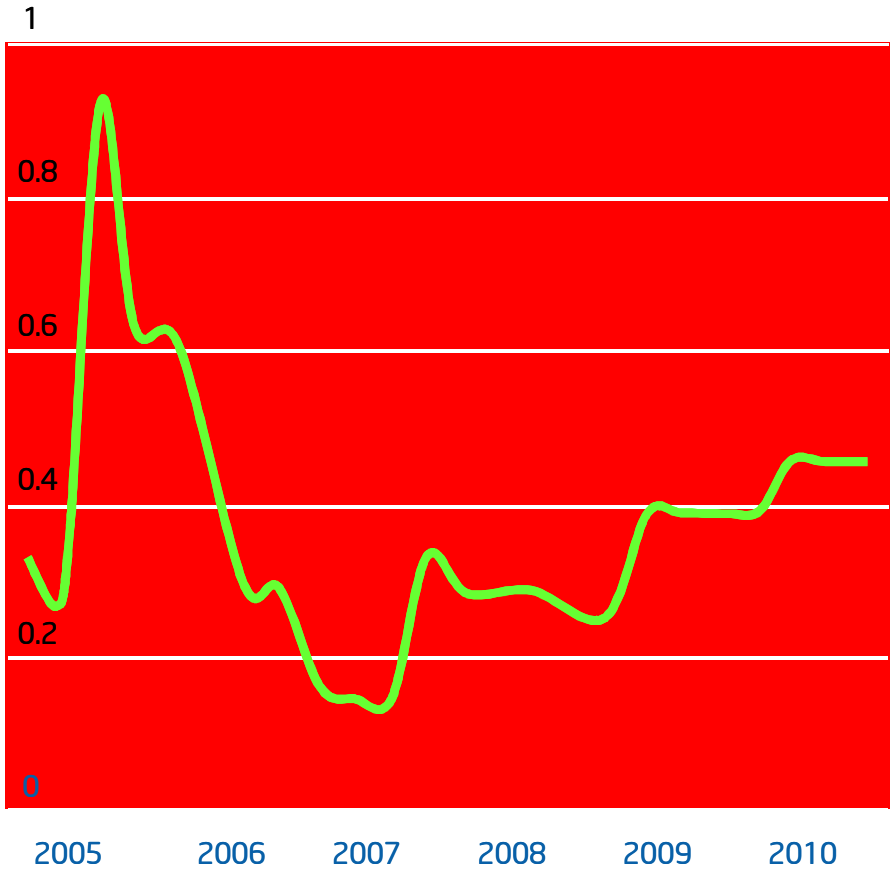

Source: Intel

Tracking Moore's Law

What can we expect!

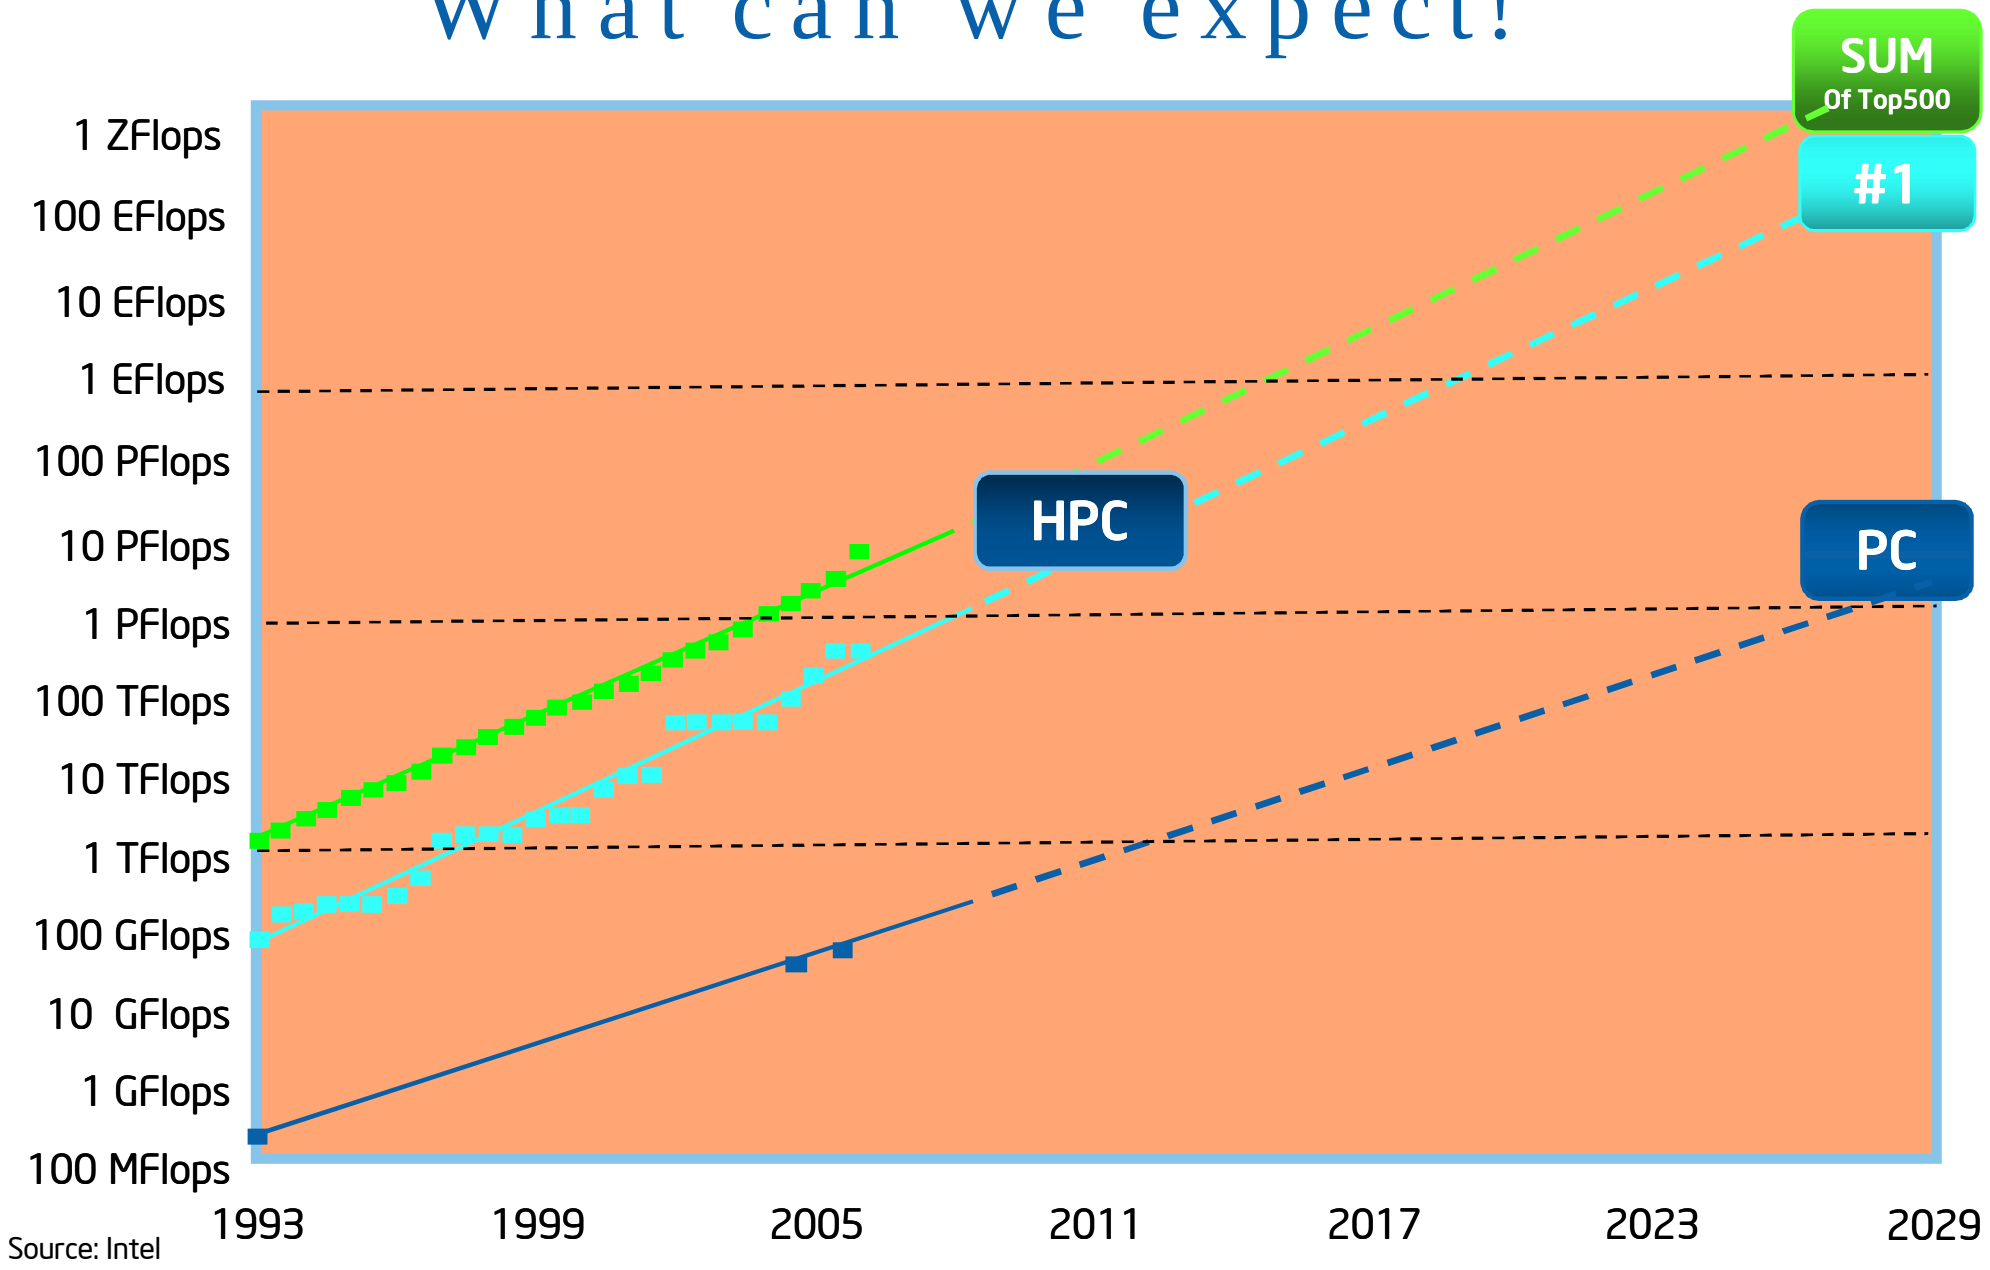

Source: Intel

Source: HPC - www.top500.org, June 2006, Intel


intel®
Leap ahead™

Increasing Throughput through Parallelism

Amdahl's Law: Parallel Speedup = 1/(Serial% + (1-Serial%)/N)*

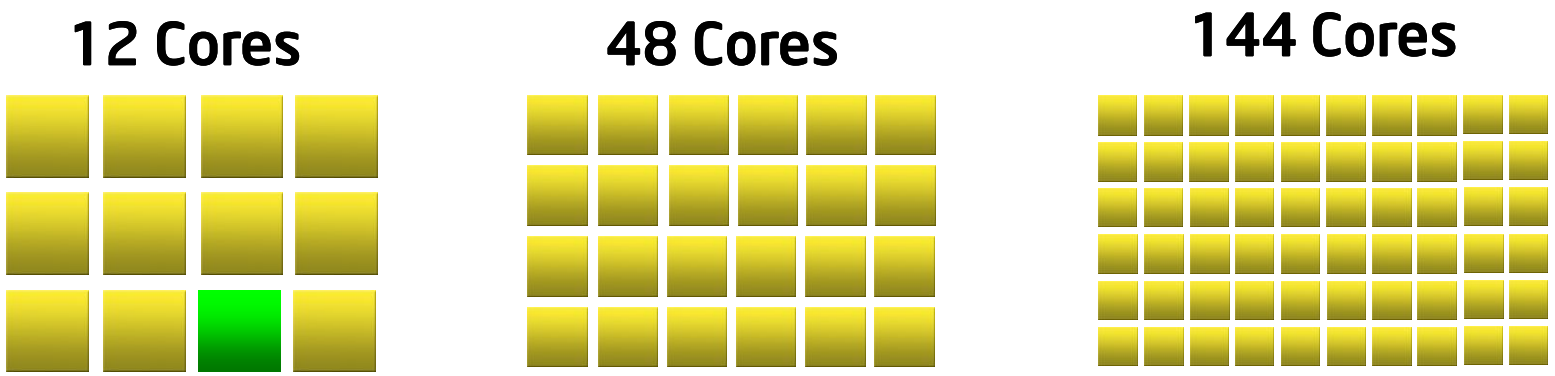

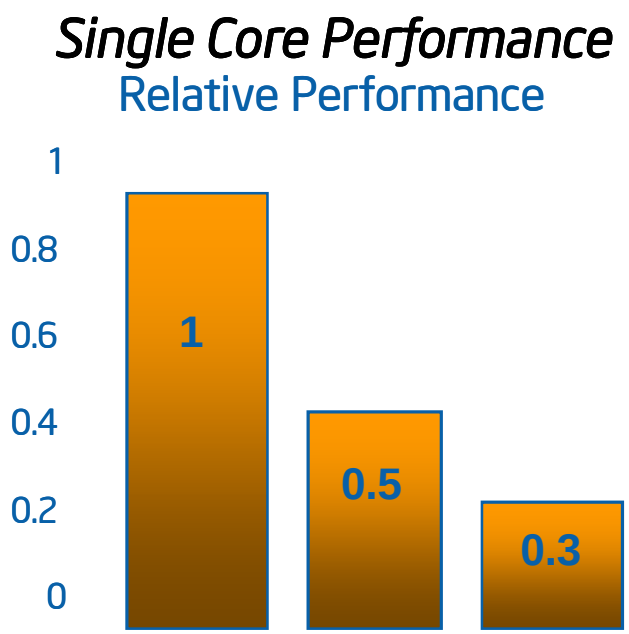

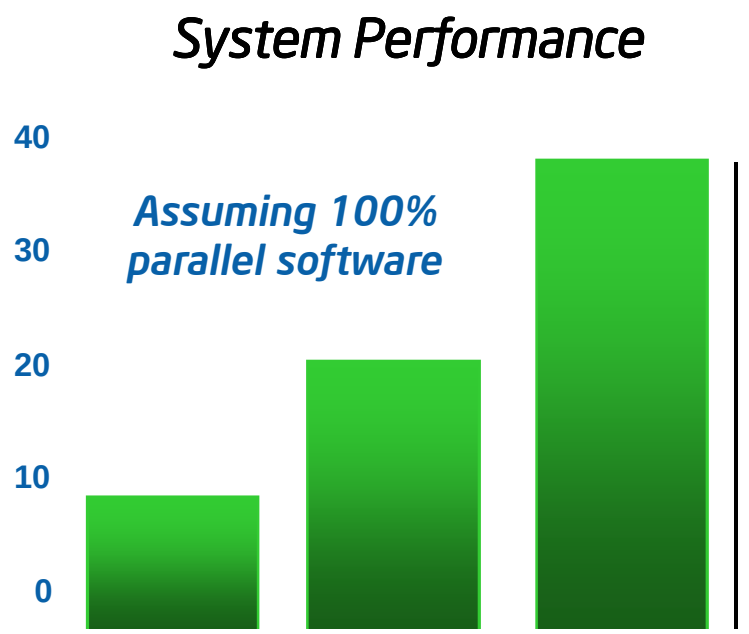

Source: Intel

* N = number of cores

Large Medium Small

Large Medium Small

Definitions

High Performance Computing (HPC) –

A collection of hardware systems, software tools, languages and generic programming approaches which make previously unfeasible applications possible and which is available at an appropriate price.

Peta Scale Computing –

" (the) wide spread use of systems that deliver sustained applications performance a level above a PFlop/s." - Horst D Simon, LBNL 8/24/2006

Impact of Petascale Computing on Medicine

(the intersection of bio/nano/intelligent computing)

Personalized, Community Medicine

- Rapid & accurate diagnosis
- Effective treatments without waste
- Put Medicine on Moore's Law

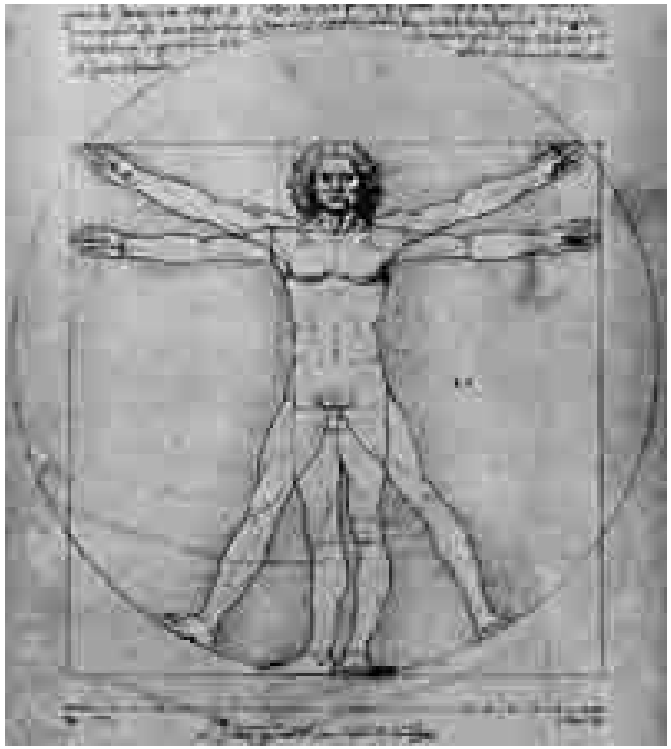

Nano + Bio + Computing

- Revolutionize medical knowledge
- Improve and create treatments
- Keep NM healthy & productive
- Control HC costs at all levels
- Create and provide high paying jobs

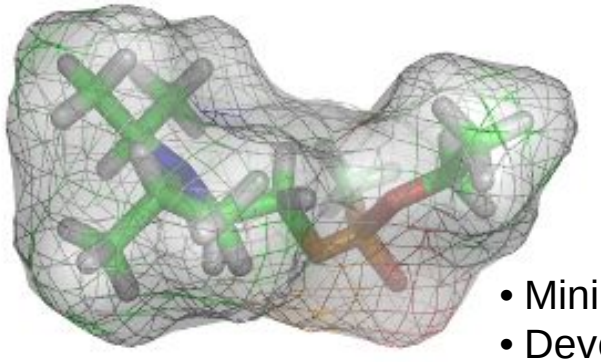

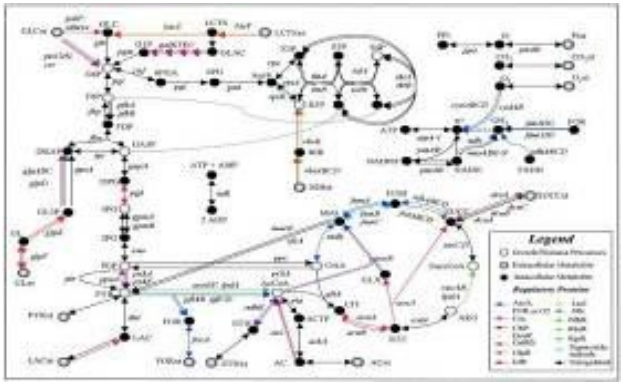

In silico drug design & testing

- Minimize dangerous human or animal testing
- Develop orphan & blockbuster drugs equally effectively
- Improve community medicine, and control costs