
intel
Leap ahead™
Intel Developer
FORUM

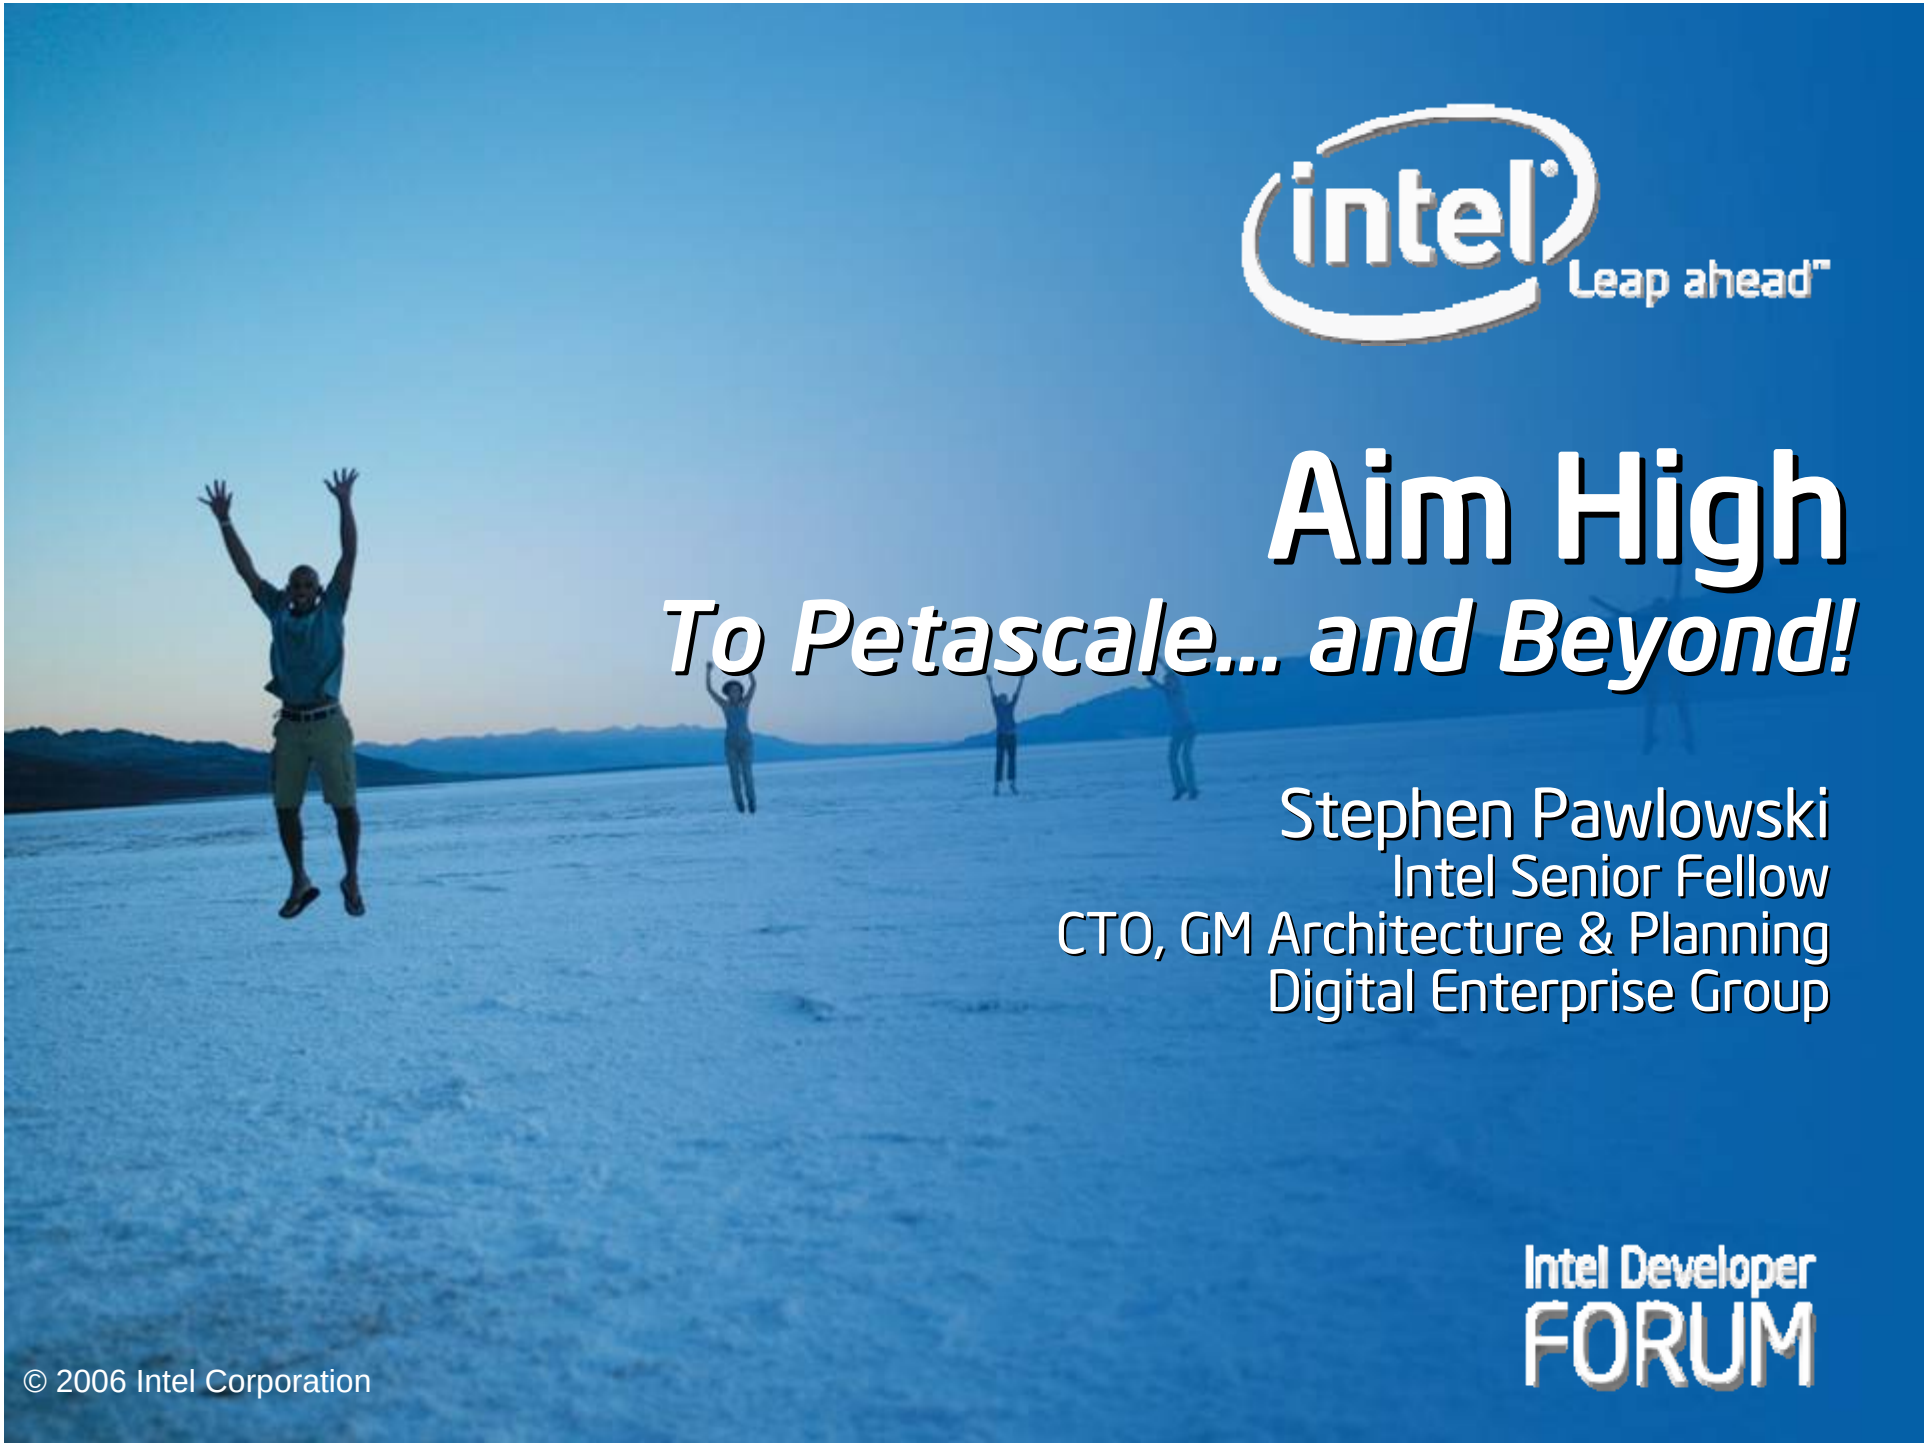

Aim High

To Petascale… and Beyond!

Stephen Pawlowski
Intel Senior Fellow
CTO, GM Architecture & Planning
Digital Enterprise Group

Intel Developer FORUM

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.

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*

Supercomputing Datacenters Today

Supercomputers Solve Complex Problems

- 60% of US profits are in financial services
- As many as ten top tier firms with IT Budgets > $2B/year
- Performance Needs:
  - 60,000 Transactions Per Second
  - 30ms average latency
  - 32GB real time database
- Good Scale out Application
  - Application is amenable to parallelization both data & thread
  - Requires Double Precision FP, high memory BW, …

Sample Application: Collateralized Debt Obligation (CDO) Monte Carlo Simulation

400,000 underlying assets (400 branches, ~1,000 loans each)

Monte Carlo Simulation considers external factors and their interrelationships

400B scenarios (simplified).

- For each asset there are 1M paths. 1,000 prepay * 1,000 loss
- Future goal to compare CDO to another CDO deal (CDO^2) 160T scenarios (400B^2). Considering CDO^3

Wall Street:
Building 2-4 Data Centers Per Company by 2010

Source: Intel

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

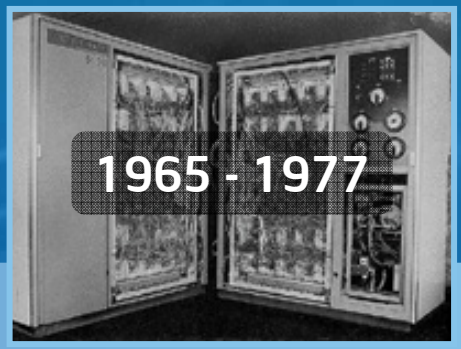

Intel ENDEAVOR
464 Intel® Xeon® Processors 5100 series, 6.85
Teraflop MP Linpack, #68 on top500

~2008 Beyond

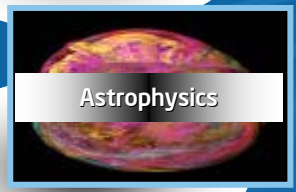

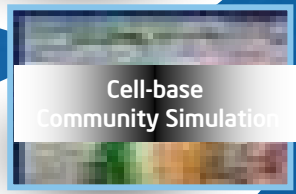

PetaScale Platforms

Yesterday's Supercomputing is Today's Personal Computing

Intel Developer FORUM

intel

... together provide a powerful environment to rapidly develop solutions for computationally challenging problems

- Tyan* Personal Supercomputer
- Intel Quad-Core Xeon Processor
- Mellanox* High Performance InfiniBand Interconnect
- Microsoft* Windows Compute Cluster Server
- Wolfram* grid*Mathematica* Supercomputing Environment

Intel Developer FORUM

*Other brands and names are the property of their respective owners.

Performance Results

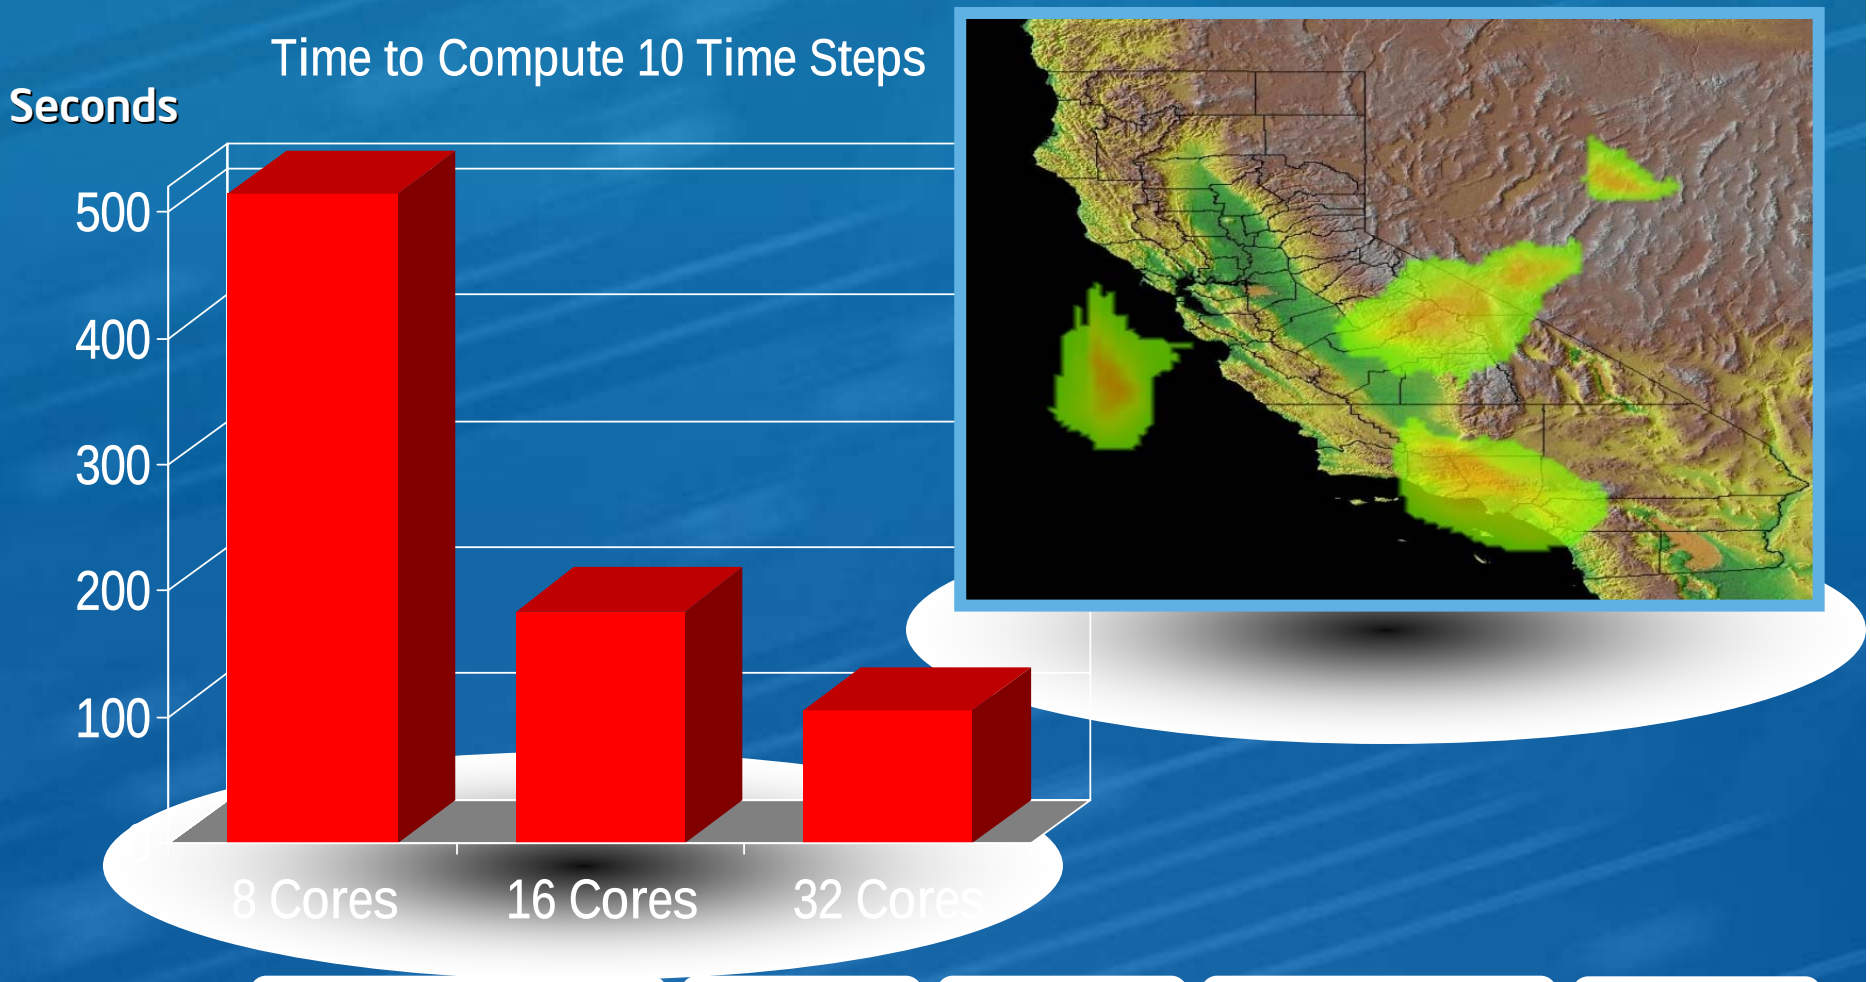

Intel Developer FORUM

Microsoft Windows Compute Cluster Server 2003

TYAN

Mellanox TECHNOLOGIES

WOLFRAM RESEARCH MAKERS OF MATHEMATICA

intel

Performance

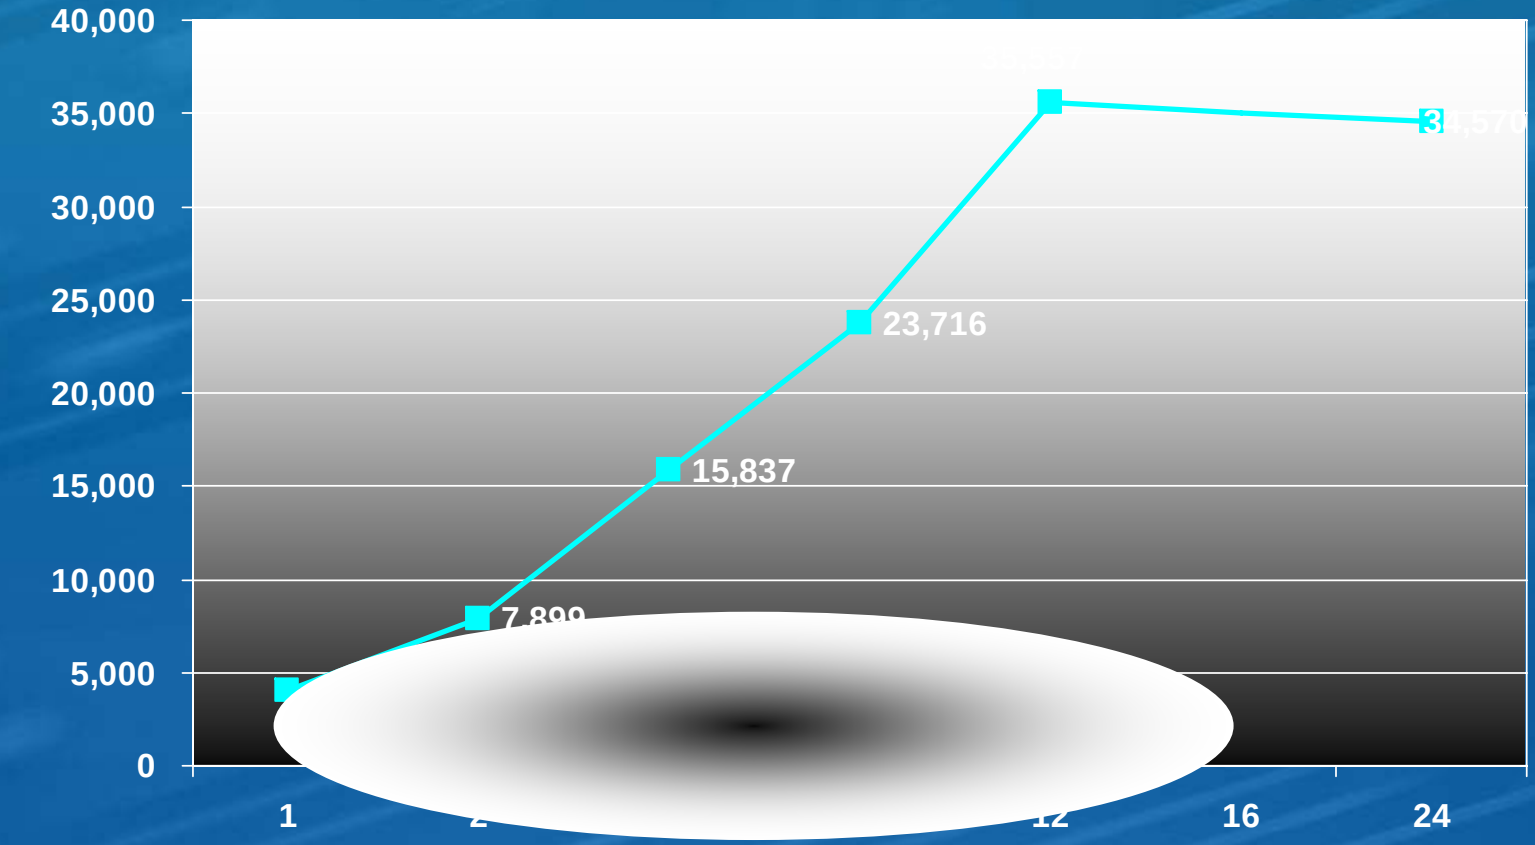

12 x Intel XEON 5160 DC (Woodcrest), 3GHz, 4MB L2; 36/48GB RAM

(Source: ScaleMP™)

Intel Developer FORUM

The Top500: Reaching Petascale

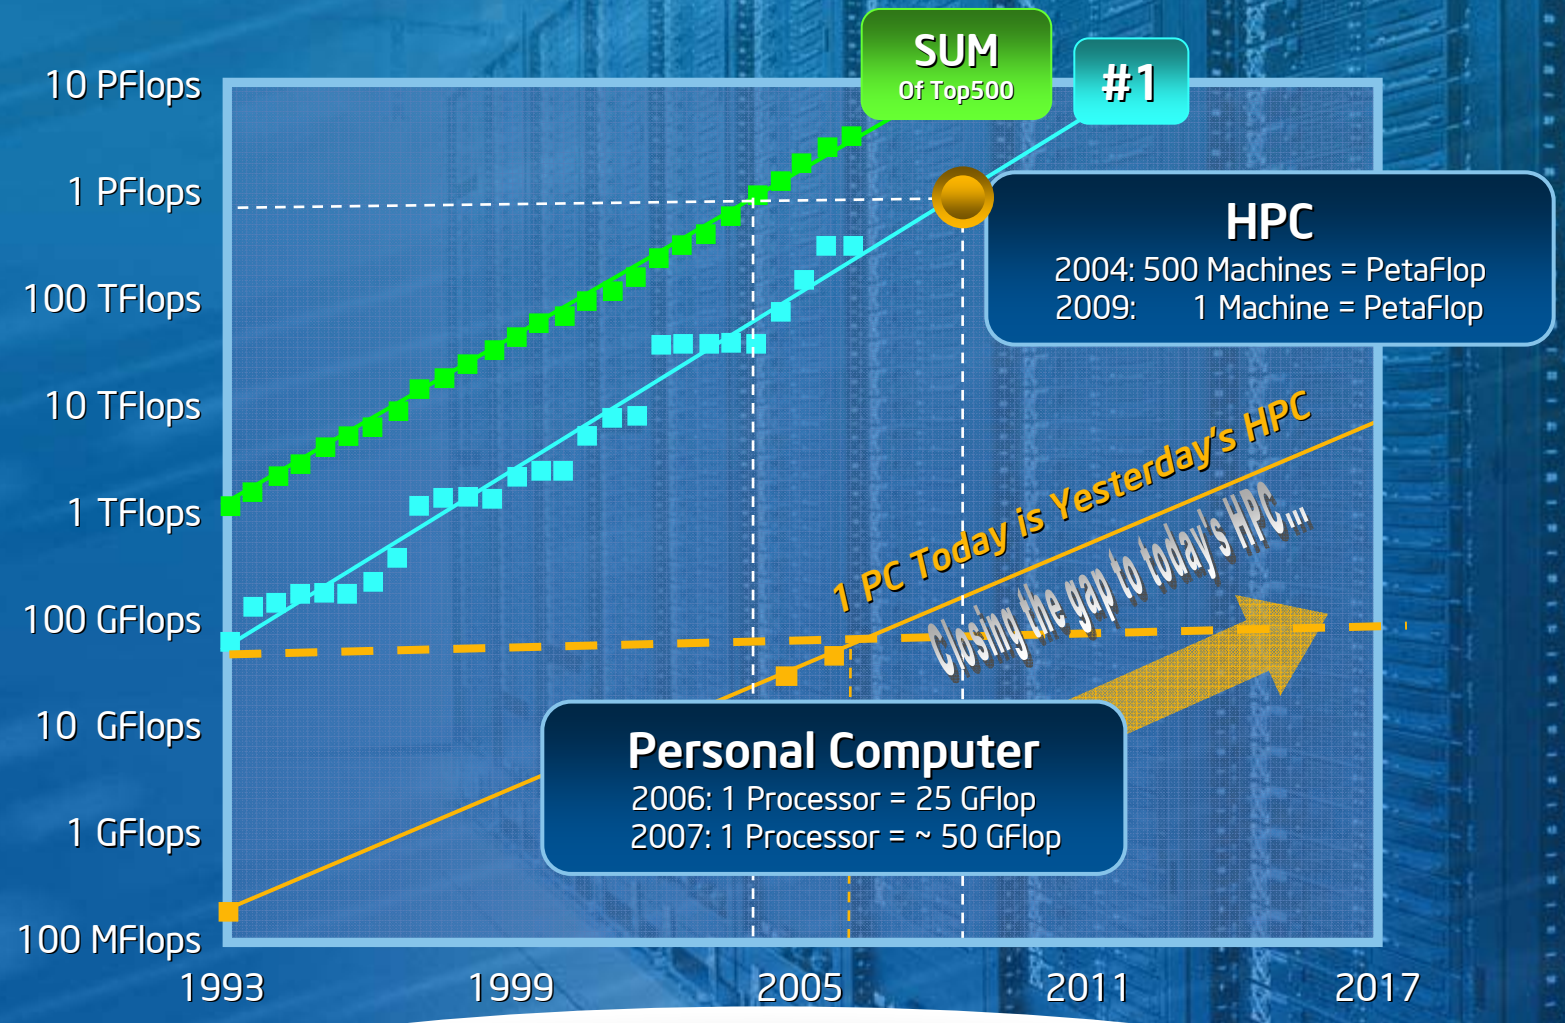

2008: Peak and Linpack PetaFlop – 2011: Sustained

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

Let's Talk About ...

Challenges and Outlook in Building a Petascale Machine

Intel Developer FORUM

Intel

Processor Performance

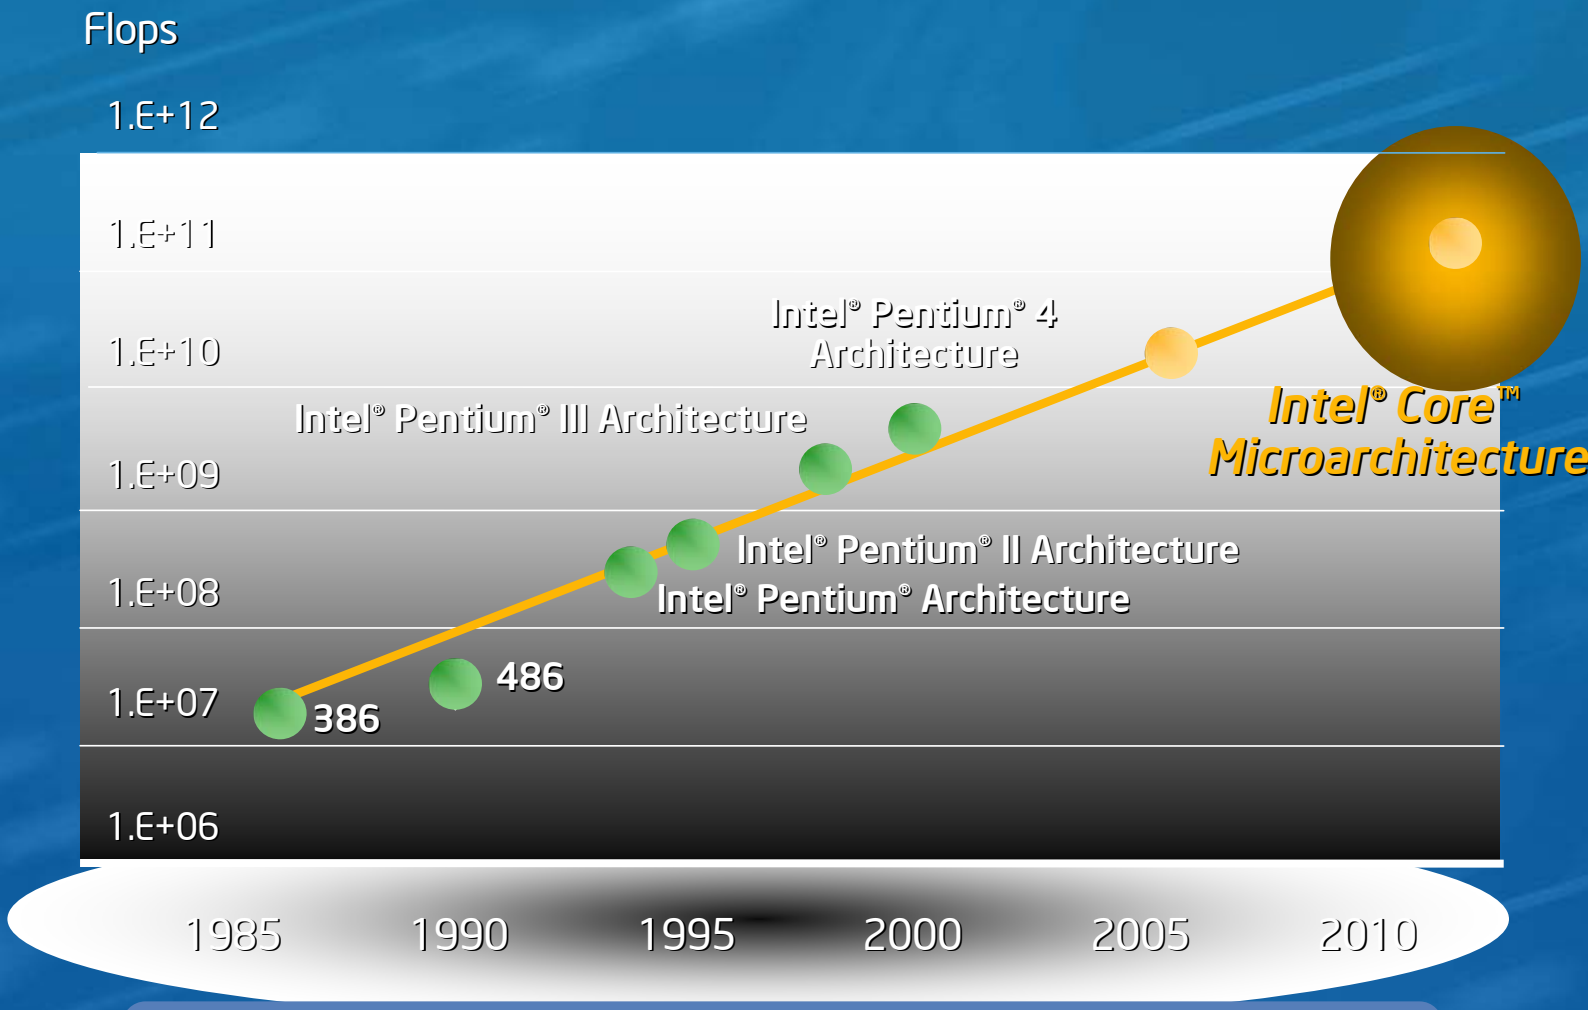

Reaching Petascale with ~100,000 Processors in 2010*

Assuming approx. 100Gflops processors

* Petascale assumes 10's of PF Peak Performance and 1PF Sustained Performance on HPC Applications.

Multi-threaded Cores

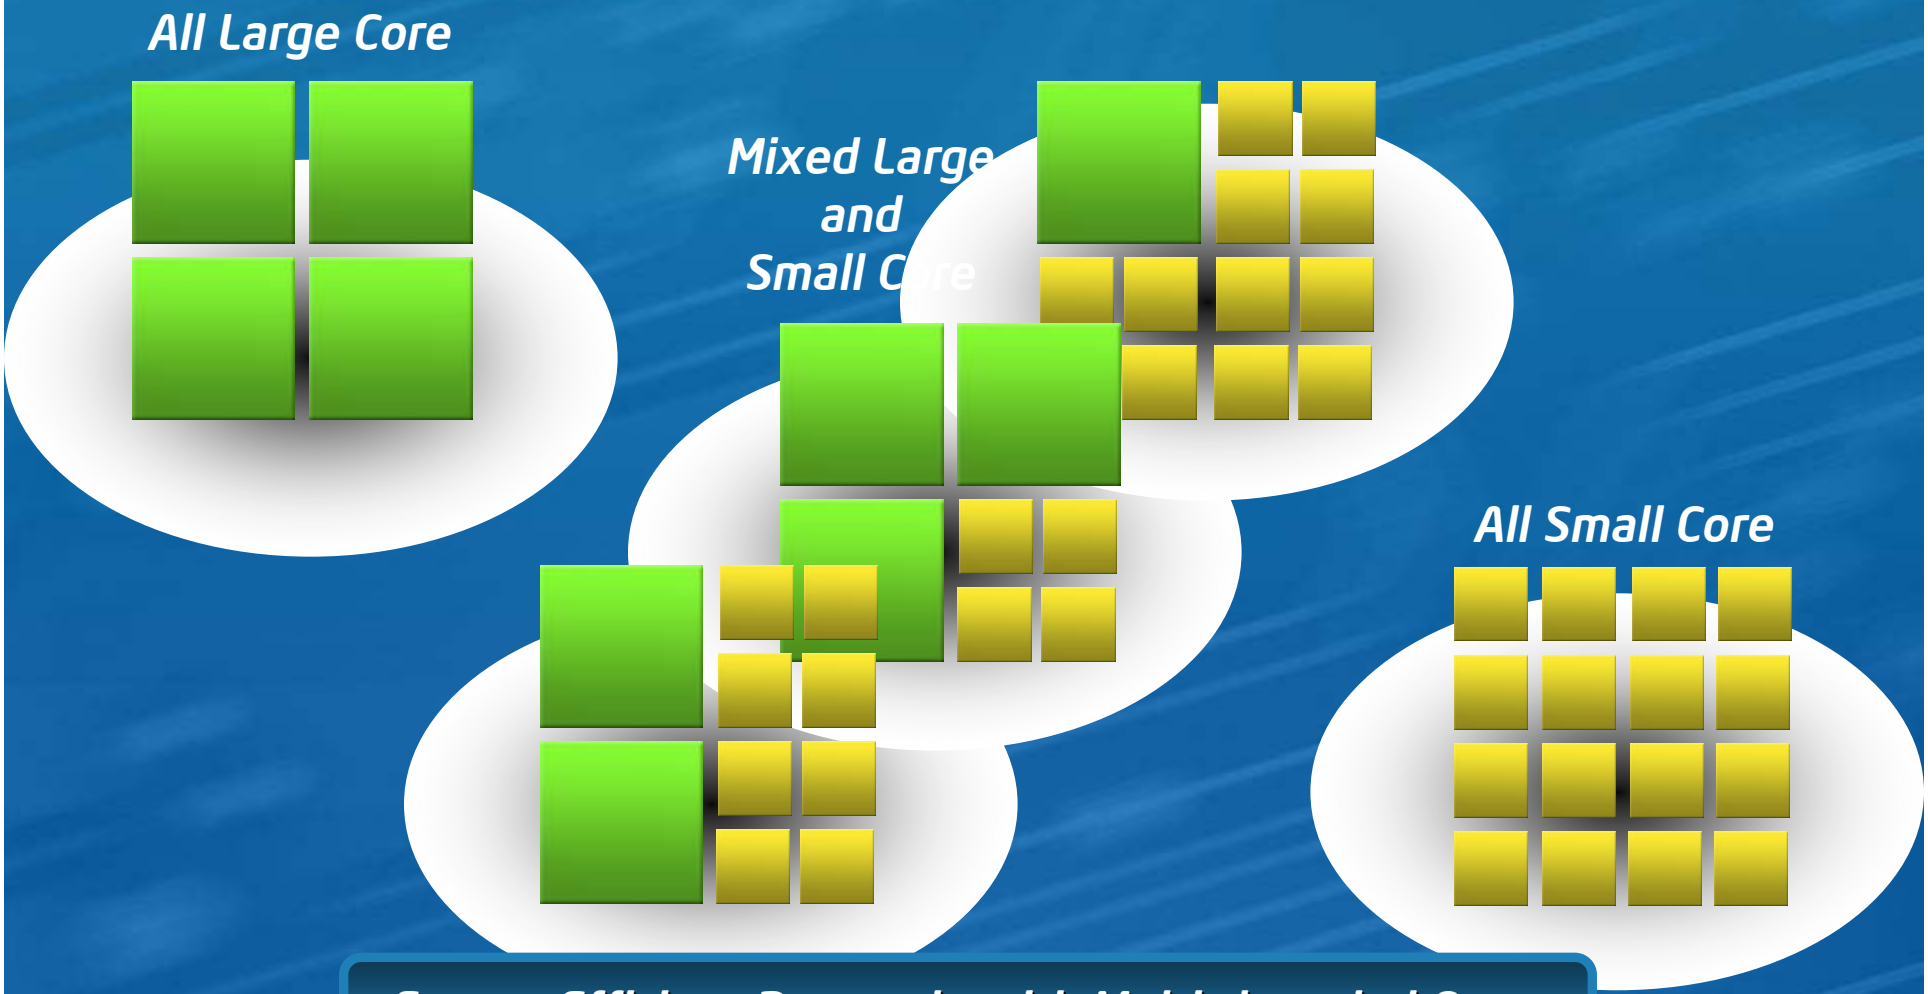

Energy Efficient Petascale with Multi-threaded Cores

Intel Developer FORUM

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

Intel

Increasing Throughput through Parallelism

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

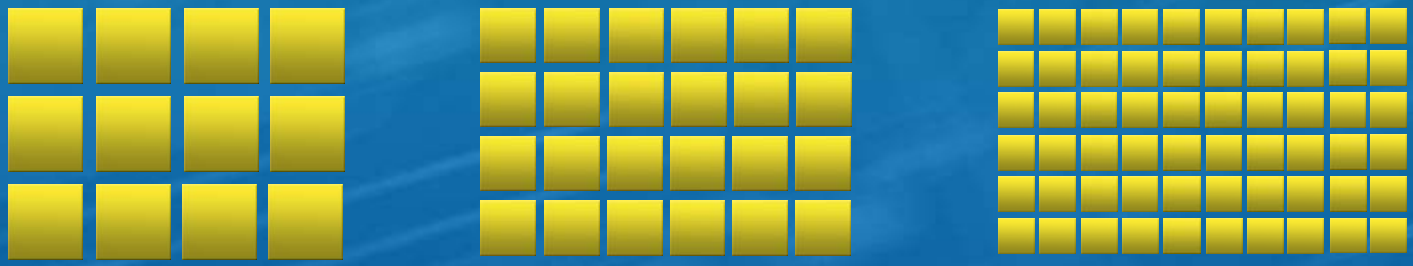

Single Core Performance

Relative Performance

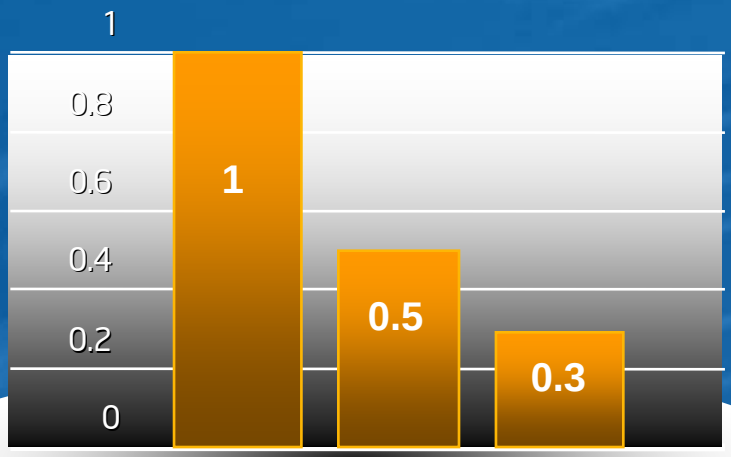

System Performance

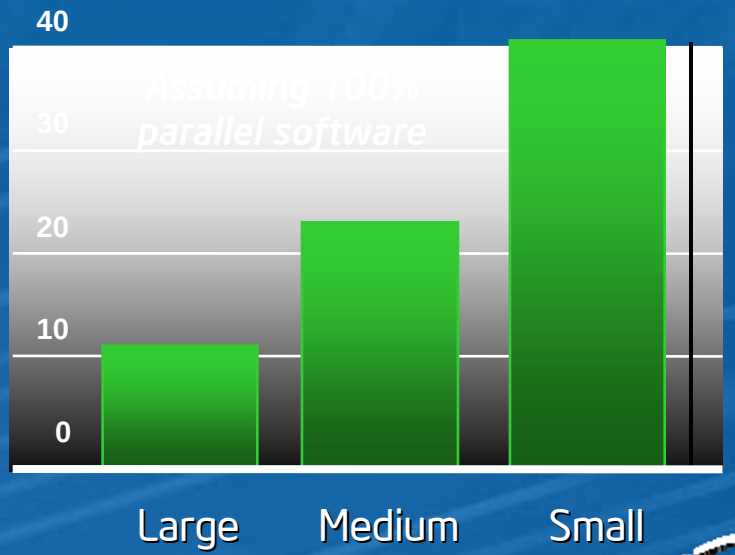

* N = number of cores

Inter-chip Interconnect Challenges
Bandwidth, Link Bandwidth and Power

Topology Effect on Bandwidth

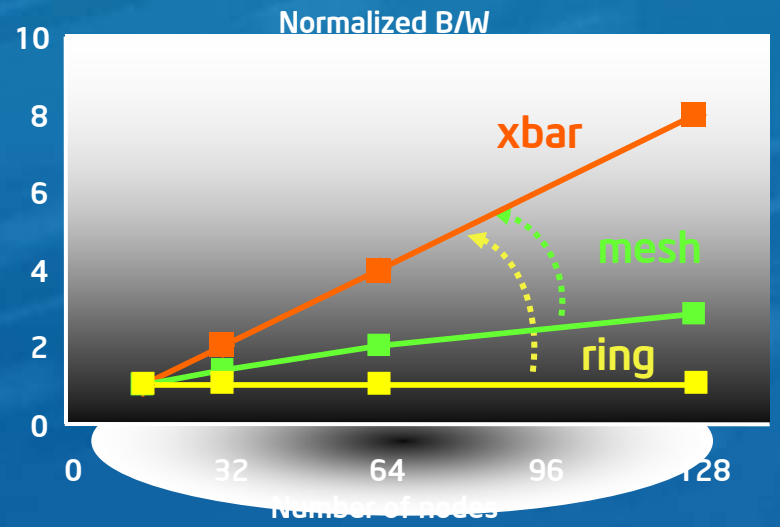

Energy Iso-Bandwidth

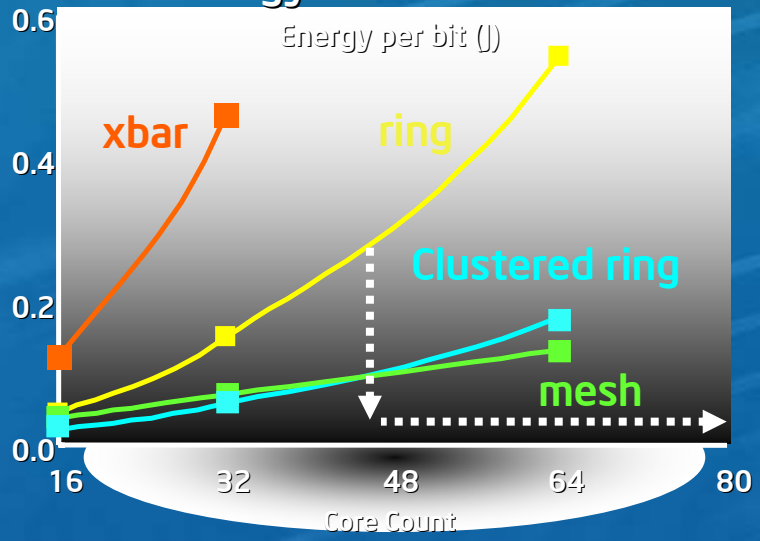

Interconnect Area Iso-Bandwidth

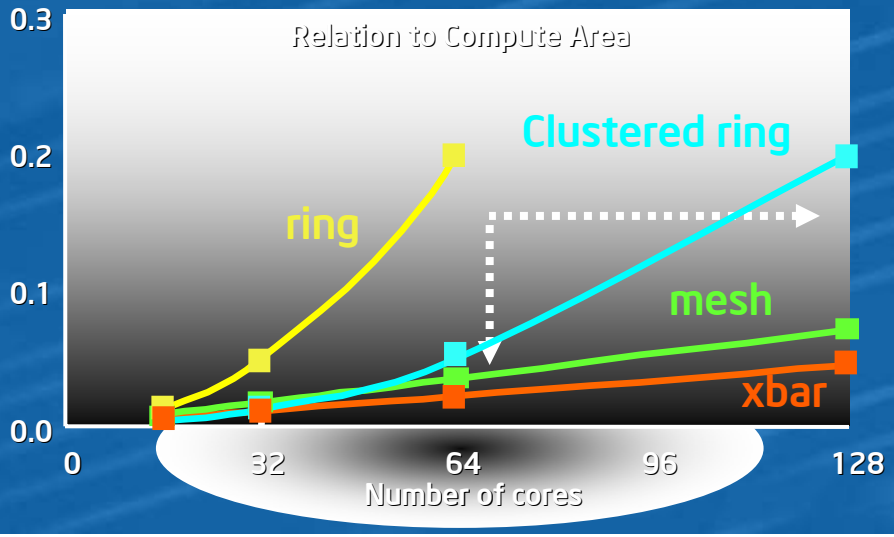

Intel Developer FORUM

intel

Increasing Processor Performance

Through Multi-threaded Cores

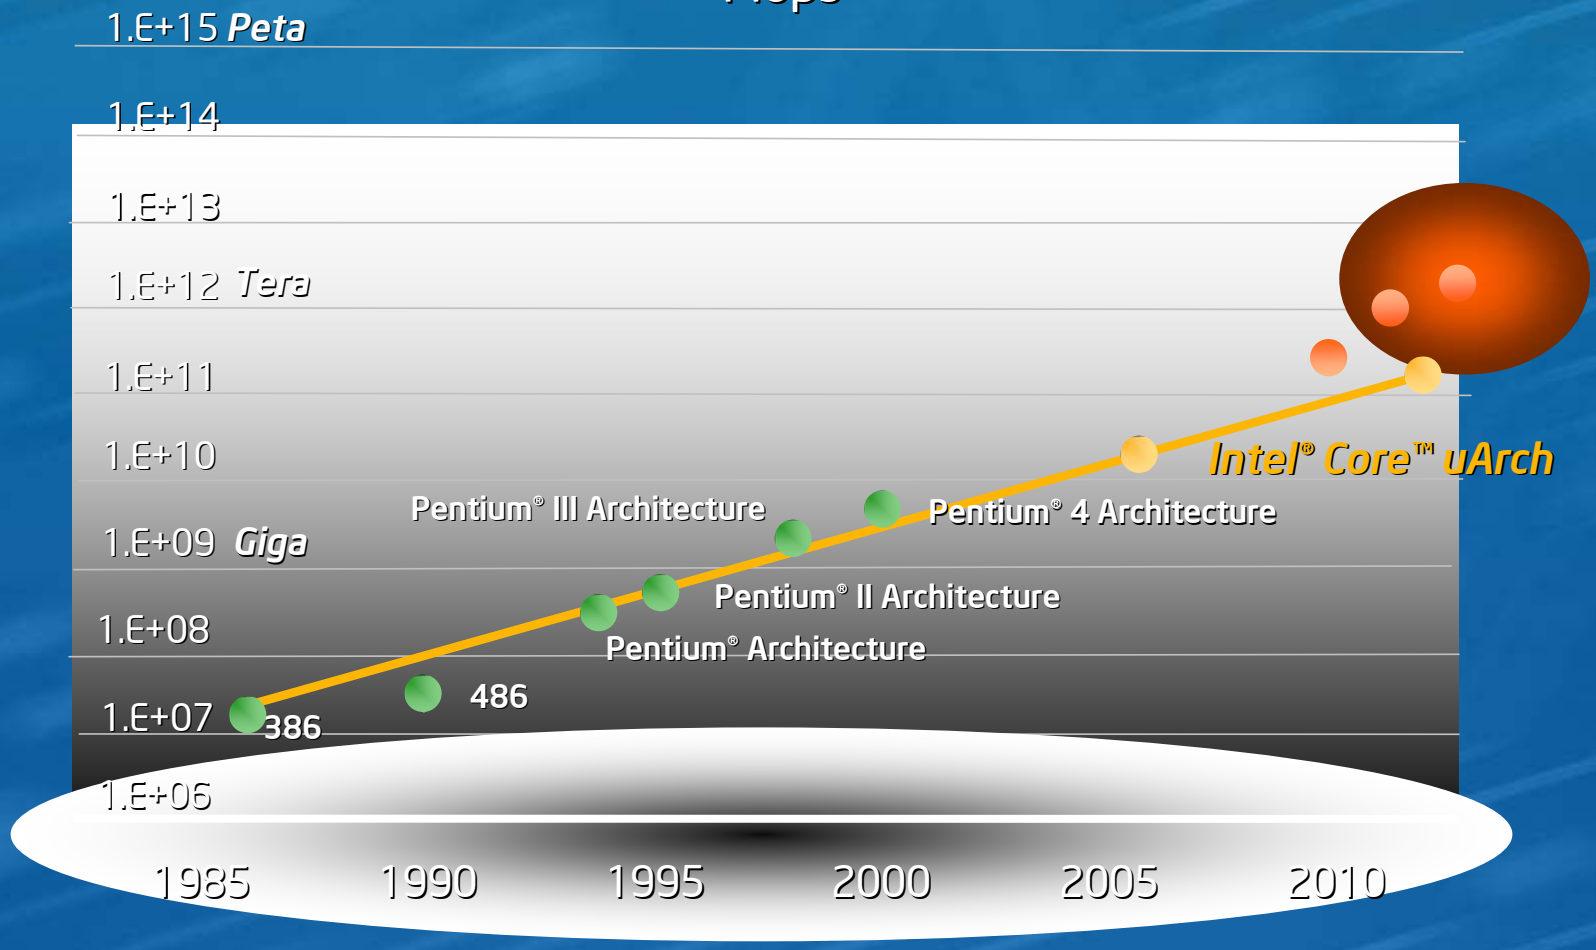

Reaching Petascale with ~5,000 Processors

Inter-Chip Interconnect Performance Today

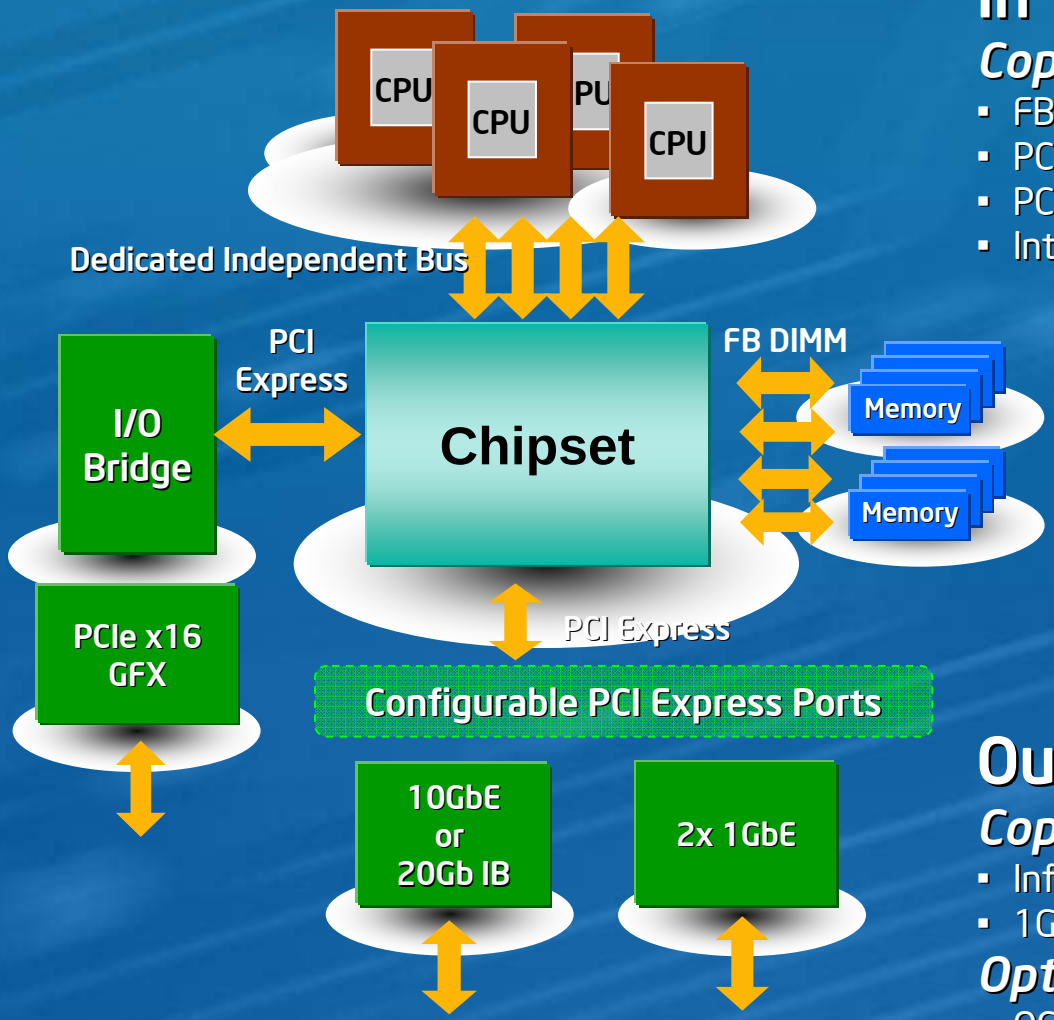

In the Box

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

Memory Performance for Balanced Computing

Byte : Flop Ratio Has Been Consistent and Steady

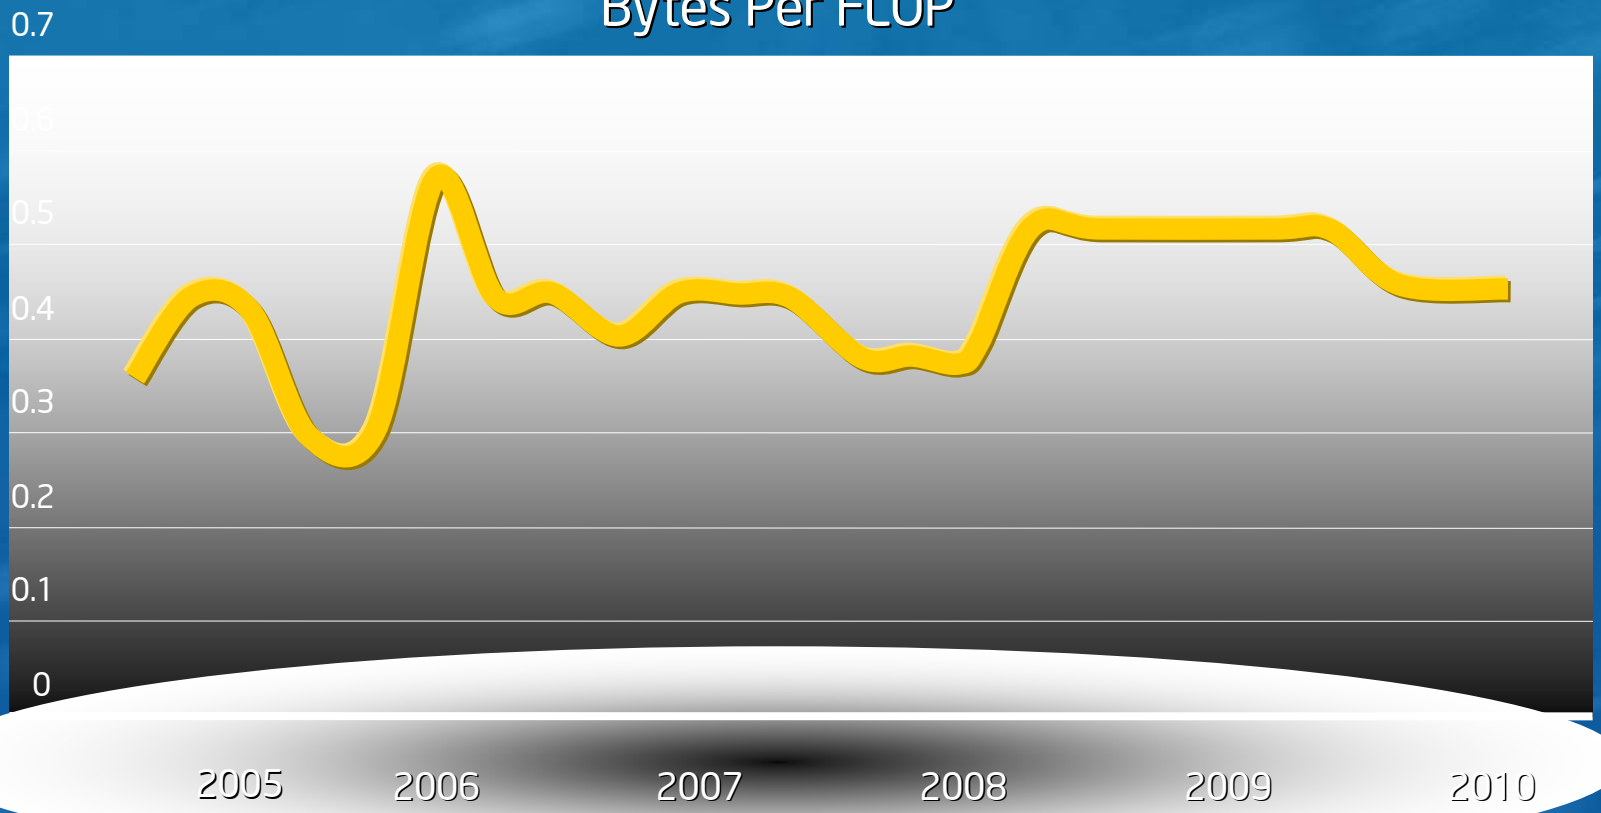

Continuing the Trend for Petascale Performance

Source: Intel

Increasing Memory Bandwidth
to Keep Pace

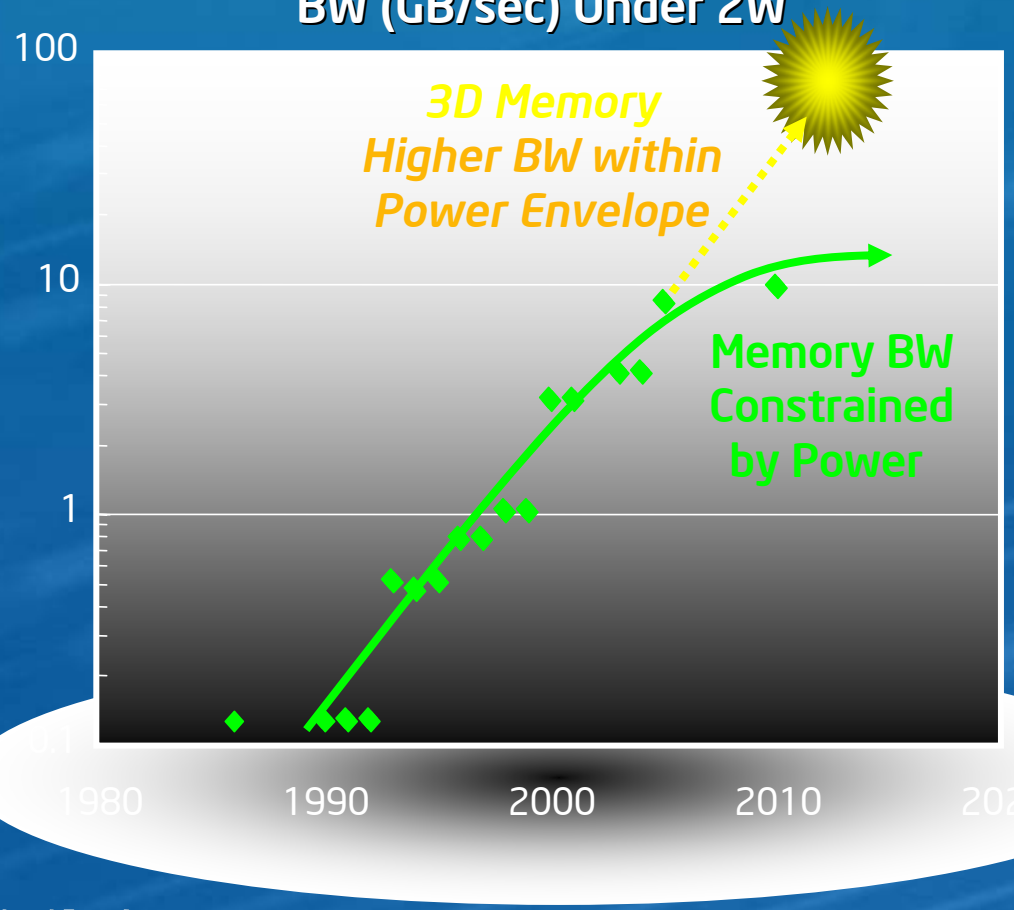

3D Memory Stacking

Power and IO Signals Go Through DRAM to CPU

Thin DRAM Die

Through DRAM Vias

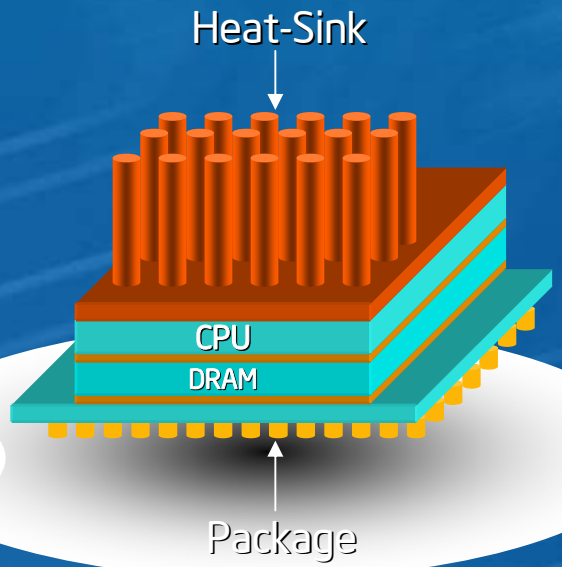

Intel Developer FORUM

Source: Intel

Intel

PCI Express to Meet I/O Demand

Performance, Bandwidth and Functionality

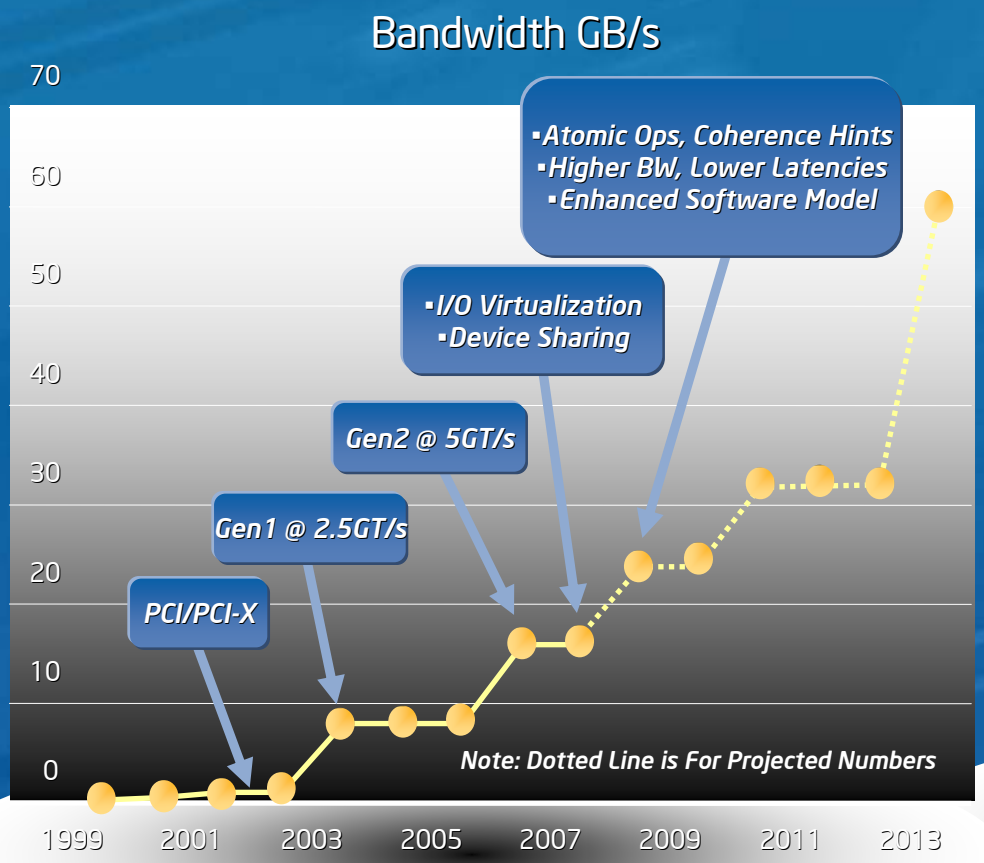

Based on x16 PCIe channel

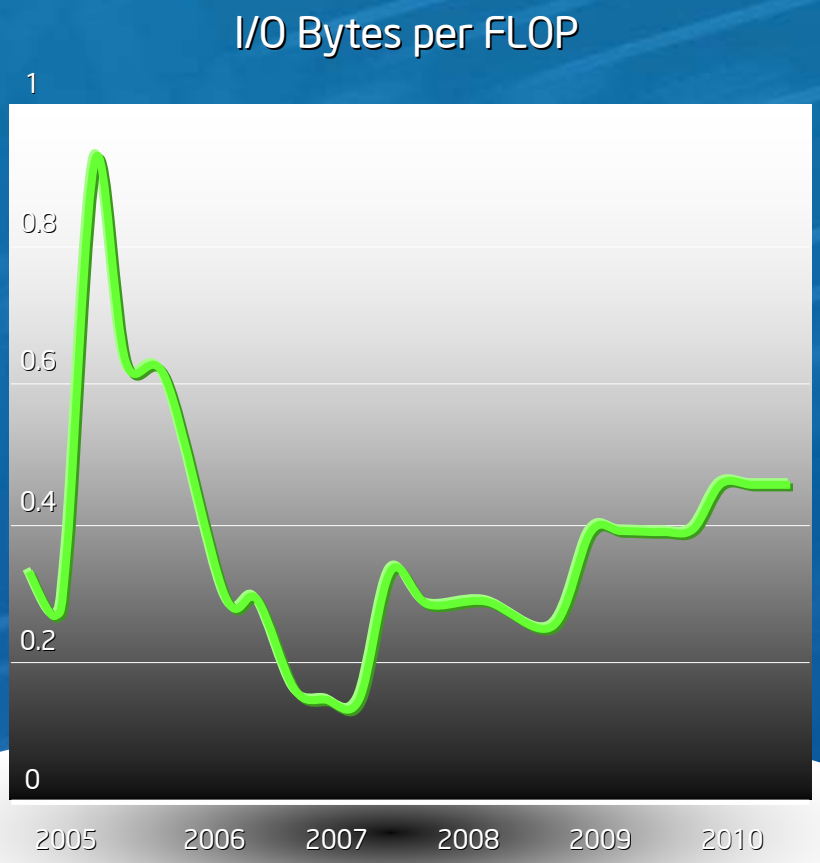

Tracking Moore's Law

Source: Intel

But ... How about

Power and Reliability

Power and Cooling Cost Today

10¢ *Kilowatt Hour* + **9** *Megawatt Datacenter* + *Power Delivery + Cooling* = **$14.6M** *Electricity Costs/Year*

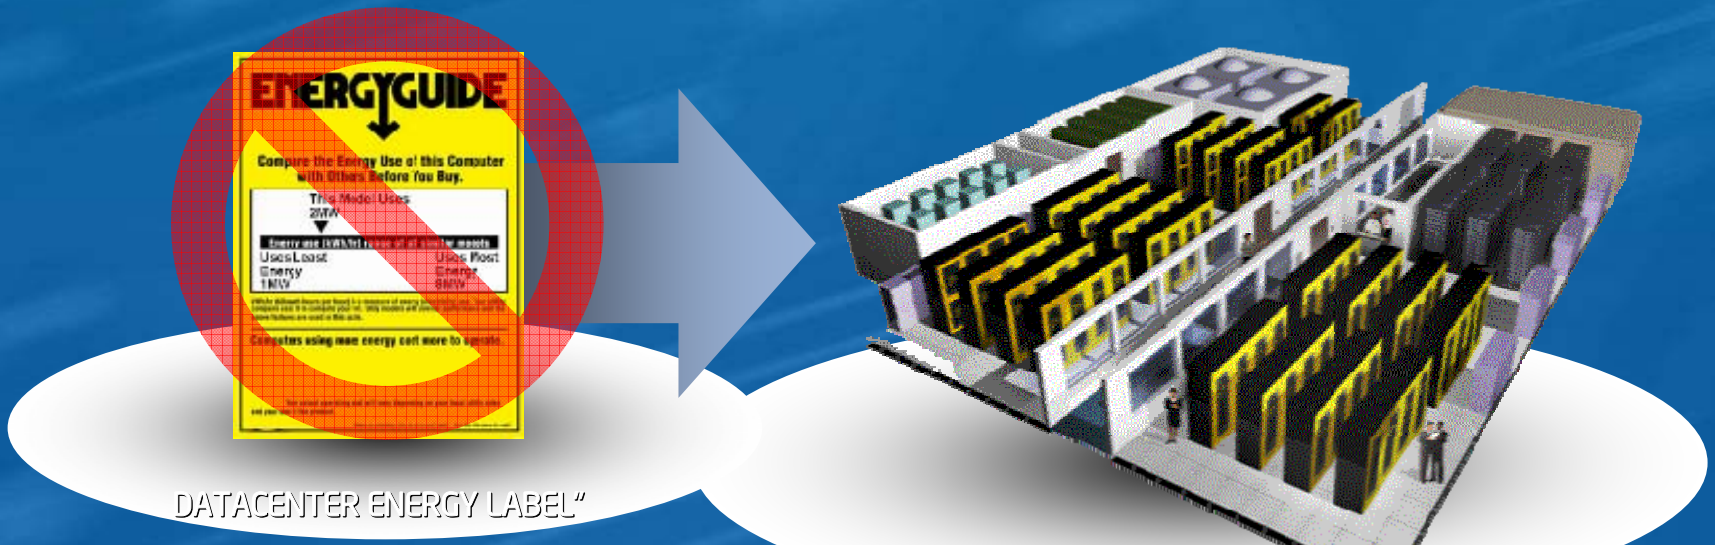

DATACENTER ENERGY LABEL"

Assume: 9MW system power*, 90% power delivery efficiency, cooling Co-efficiency of Performance (COP)=1.5

*Source: HPC Wire "A Petaflop Before its Time," June 28, 2006

Intel Developer FORUM

intel

Managing Power and Cooling Efficiency

Silicon:
Moore's law, Strained silicon, Transistor leakage control techniques, Clock gating

Processor:
Policy-based power allocation
Multi-threaded cores

System Power Delivery:
Fine grain power management,
Ultra fine grain power management

Facilities:
Air cooling and liquid cooling options
Vertical integration of cooling solutions

Reliability Challenge
Billions of Transistors

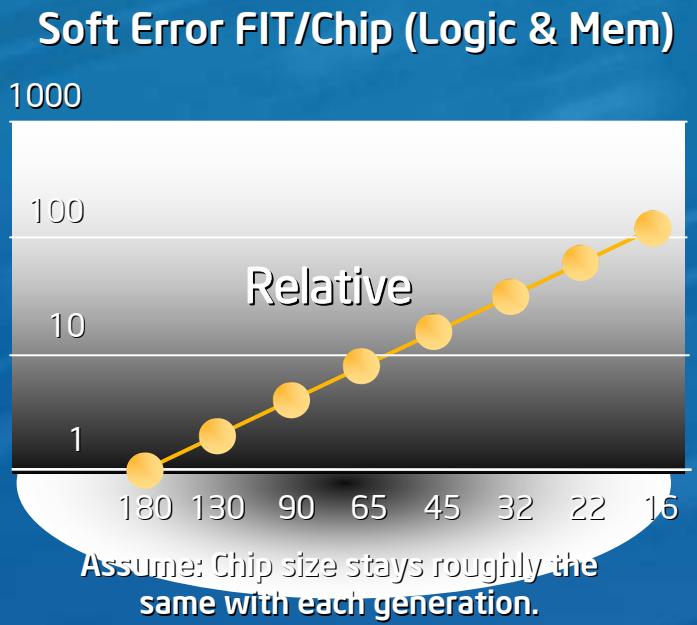

- ***FIT/bit (mem cell)***: expected to be roughly constant
- ***Moore's law***: increasing the bit count exponentially: 2x every 2 years

An exponential growth in FIT/chip

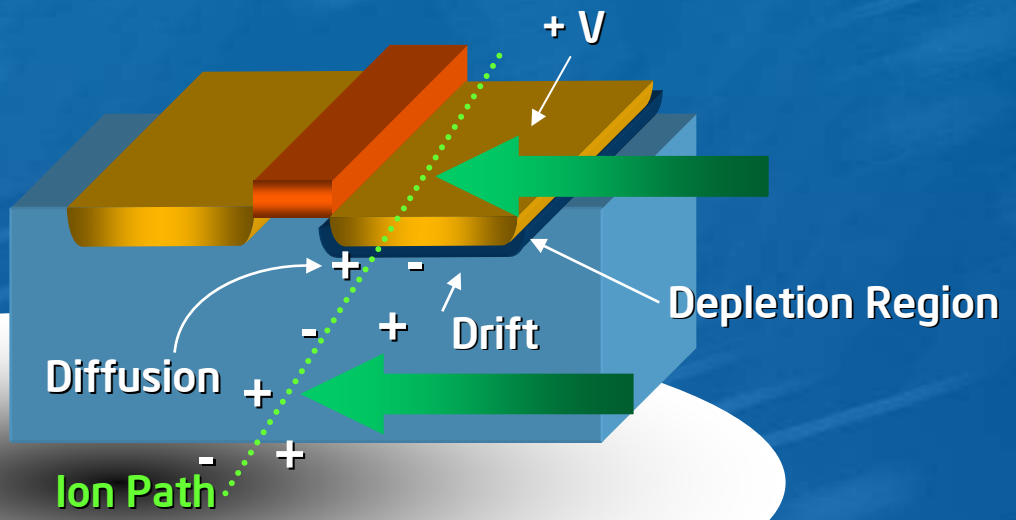

Soft Errors or Single Event Upsets (SEU) are caused by charge collection following an energetic particle strike.

Soft Error: One of Many Challenges

Source: Intel

Intel Developer FORUM

intel

100s & 1000s of Processors

An Example of Datacenter Growth …

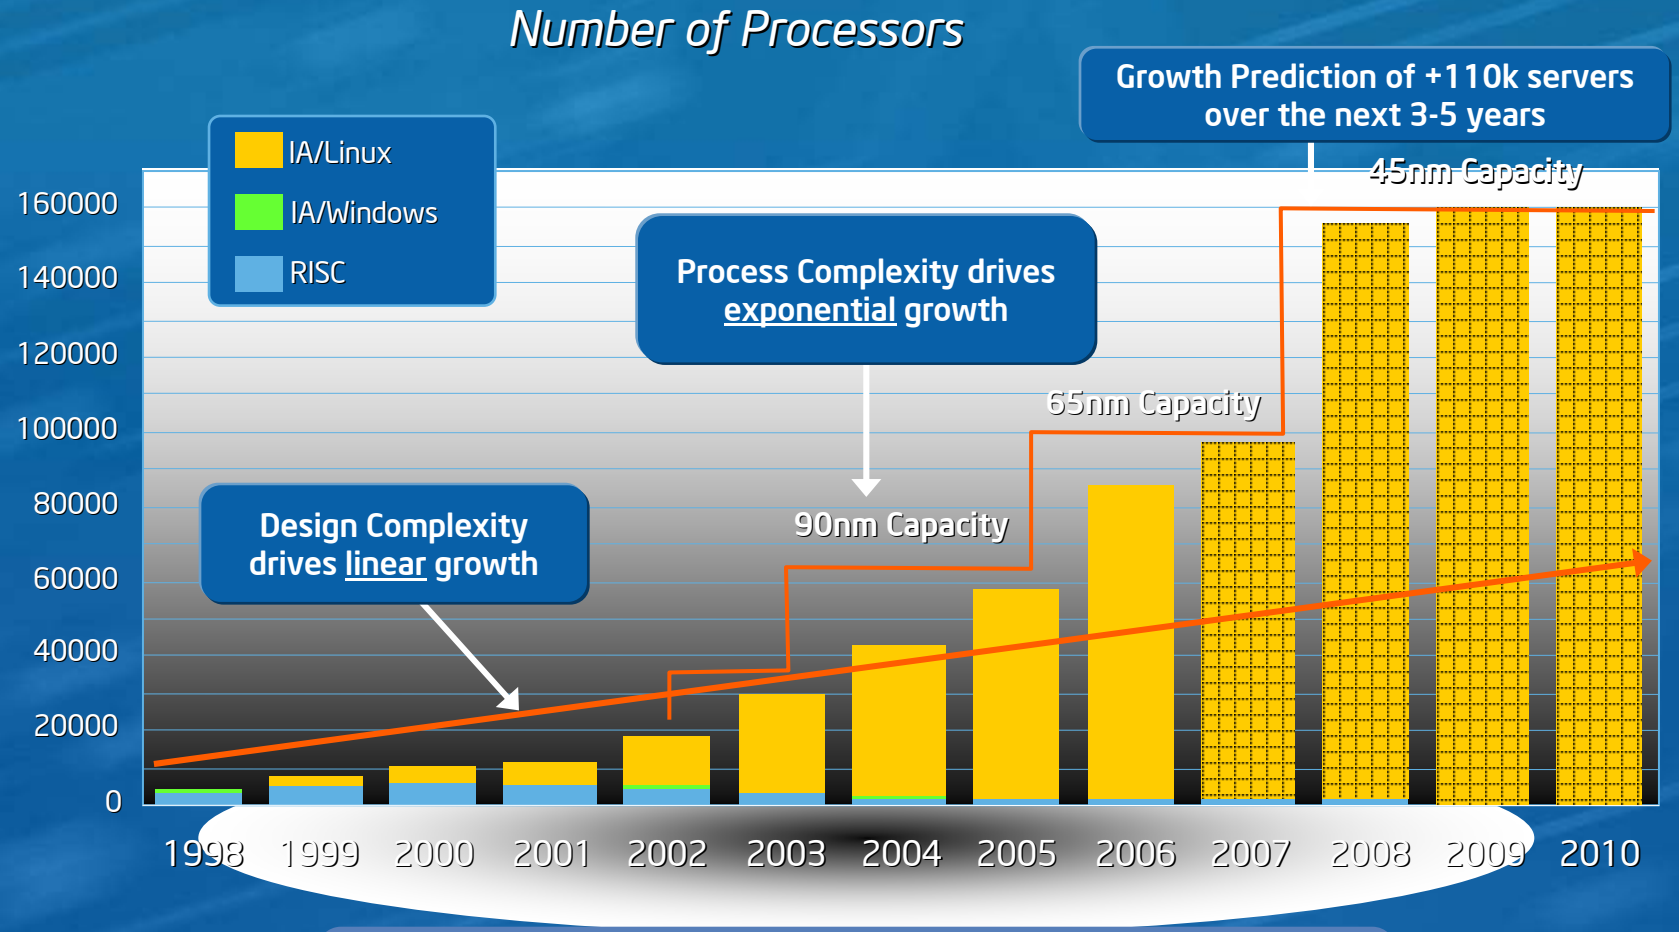

Compute Capacity Growth is Accelerating

Intel Developer FORUM

Source: Intel

Intel

Reliable Systems With Unreliable Components

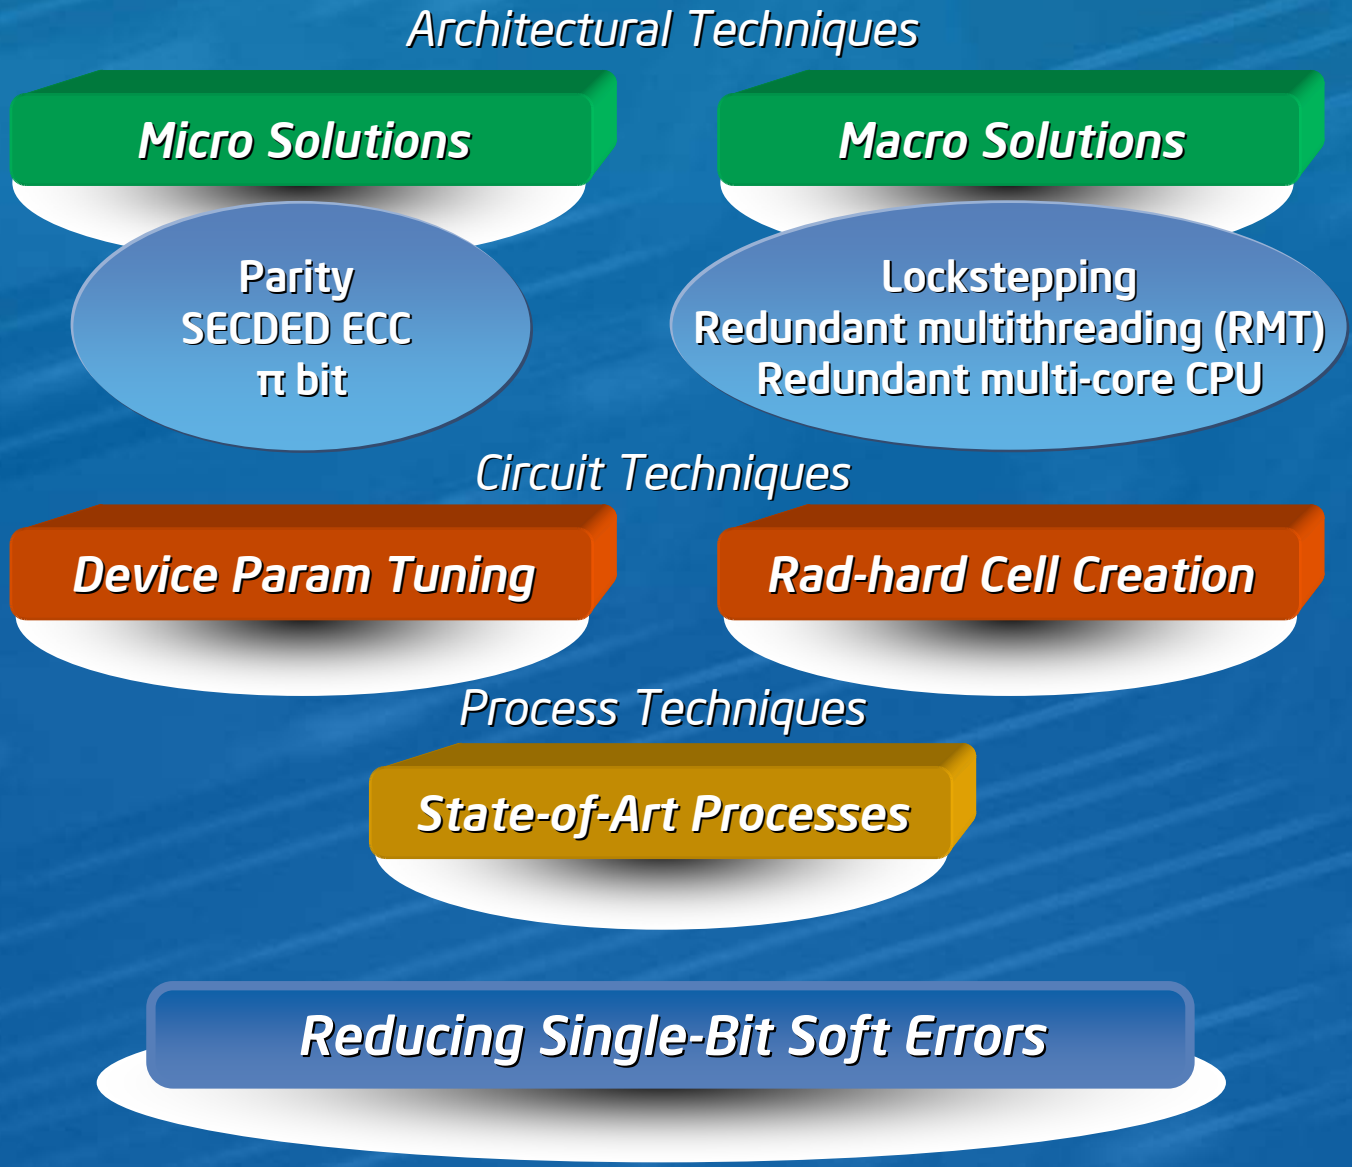

Do we need more than Petascale?

Parallel and programmable
10s to 100s cores and threads
100s and 1000s of processors
100s Gbps chip-to-chip signaling
10s to 100s billions transistors
Dynamic self-test, detect, reconfigure, & adapt

ynamic Analysis:	1 Petaflops
Laser Optics:	10 Petaflops
Molecular Dynamics in Biology:	20 Petaflops
Aerodynamic Design:	1 Exaflops
Computational Cosmology:	10 Exaflops
Turbulence in Physics:	100 Exaflops
Computational Chemistry:	1 Zettaflops

Source: Dr. Steve Chen, "The Growing HPC Momentum in China", June 30th, 2006, Dresden, Germany

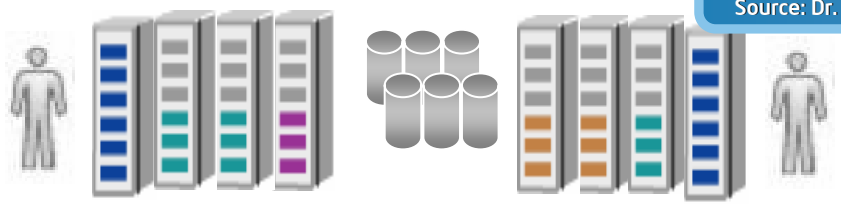

What can we expect!

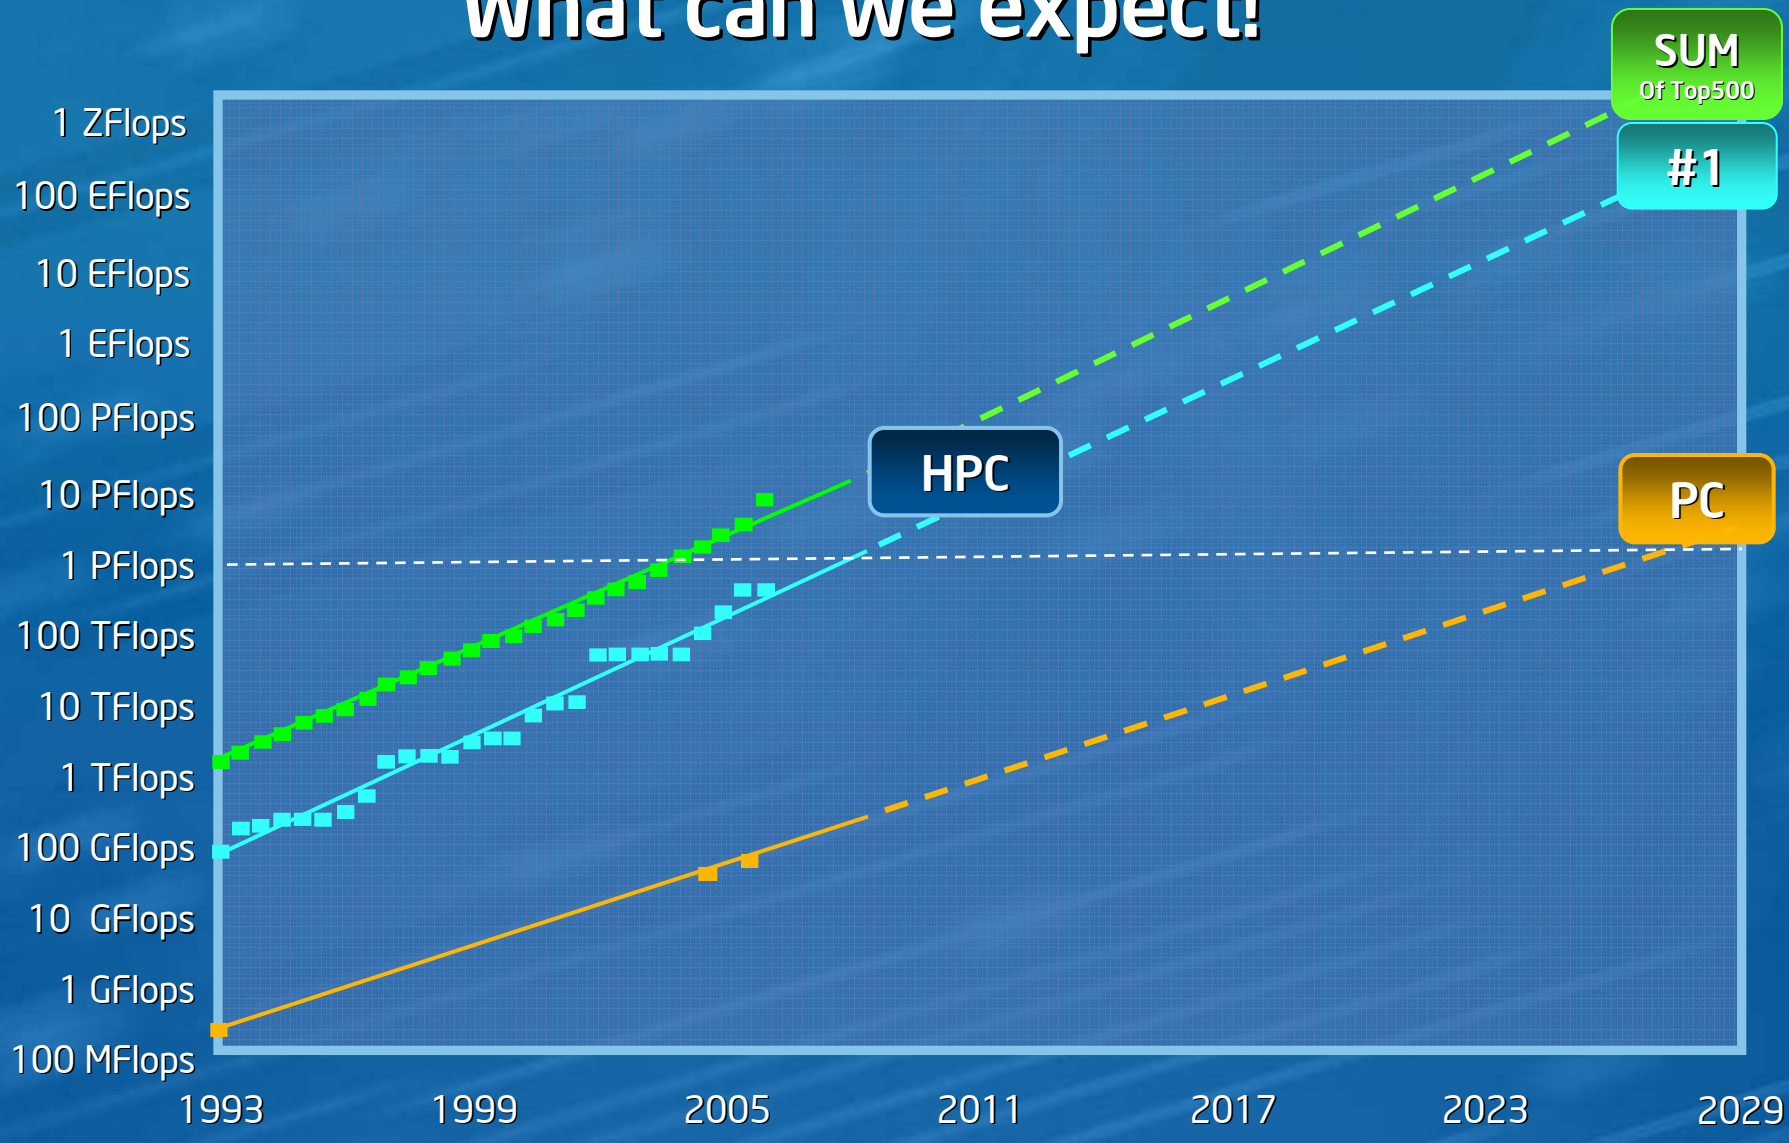

Intel Developer FORUM

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

intel


intel
Leap ahead™

Q&A


intel
Leap ahead™
Intel Developer
FORUM