Petascale to Exascale

Extending Intel's HPC Commitment

Kirk Skaugen

Vice President, Intel Corporation
General Manager, Data Center Group

intel

Congratulations Prof. Dr. Meuer on the 25th Anniversary of ISC

25 Years Also = Intel Beginnings in HPC: "The Cosmic Cube"

- Scalable from 32 to 128 nodes
  - Intel 80286 microprocessor
  - Intel 80287 math coprocessor
  - 512K RAM local memory
  - Ethernet-connected hypercube
- Peak performance of 3.2 MFLOPS

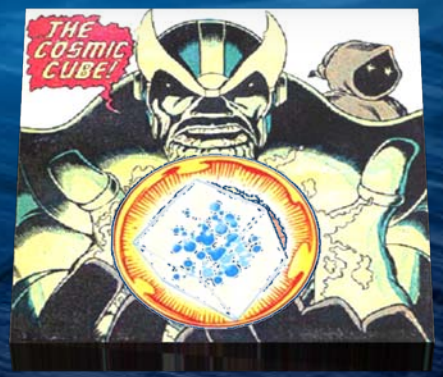

A Rich History of Silicon and Software Innovation for HPC

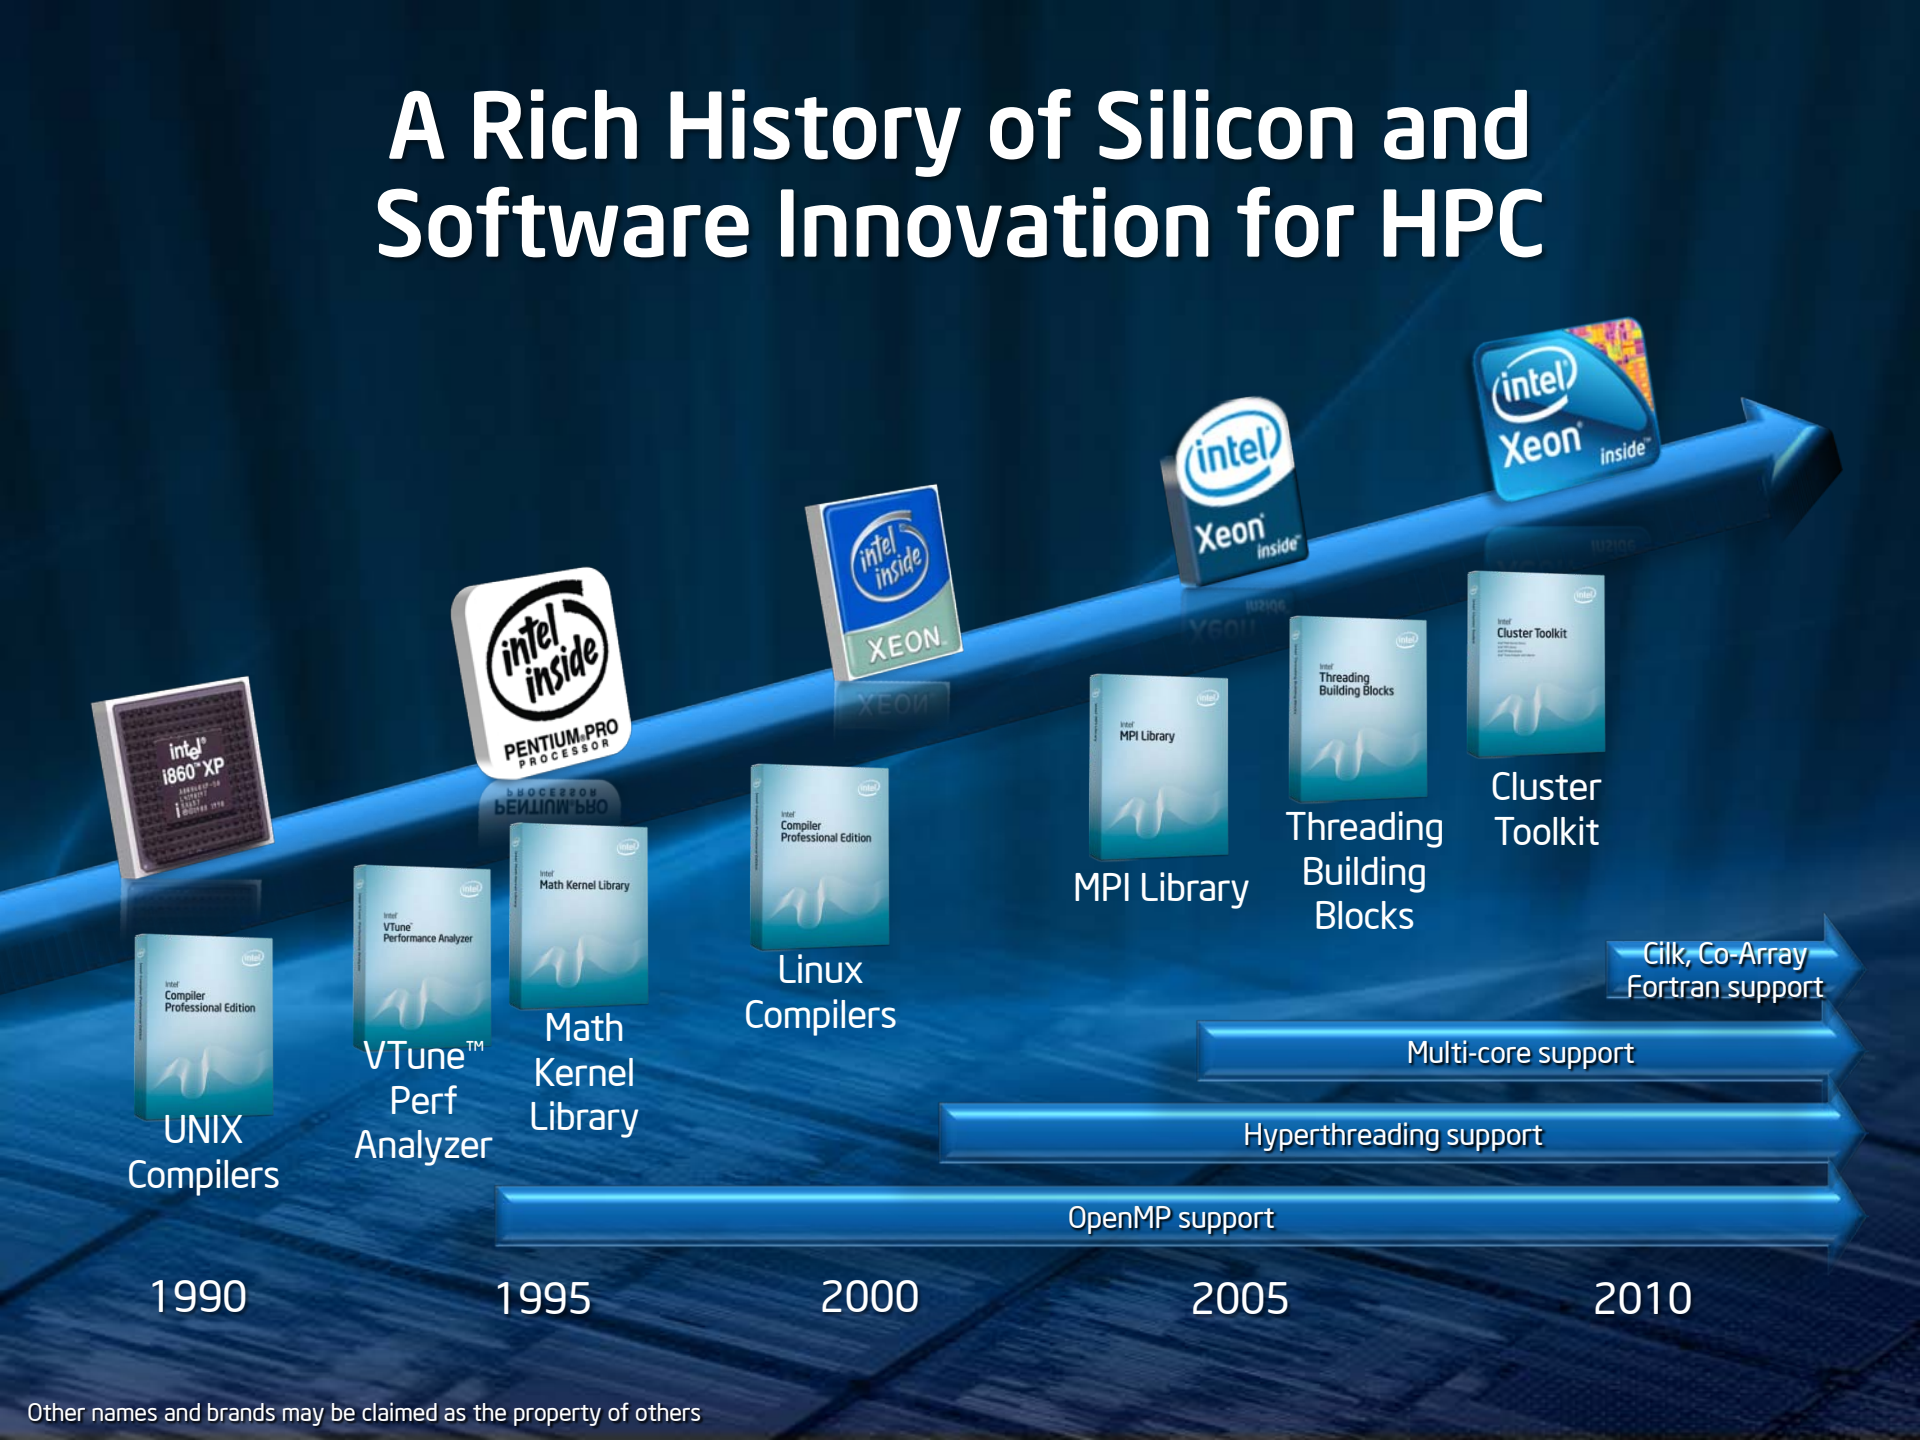

Intel Top 500 Market Adoption

Value Proposition:

Volume economics

IA programming model

Robust ecosystem

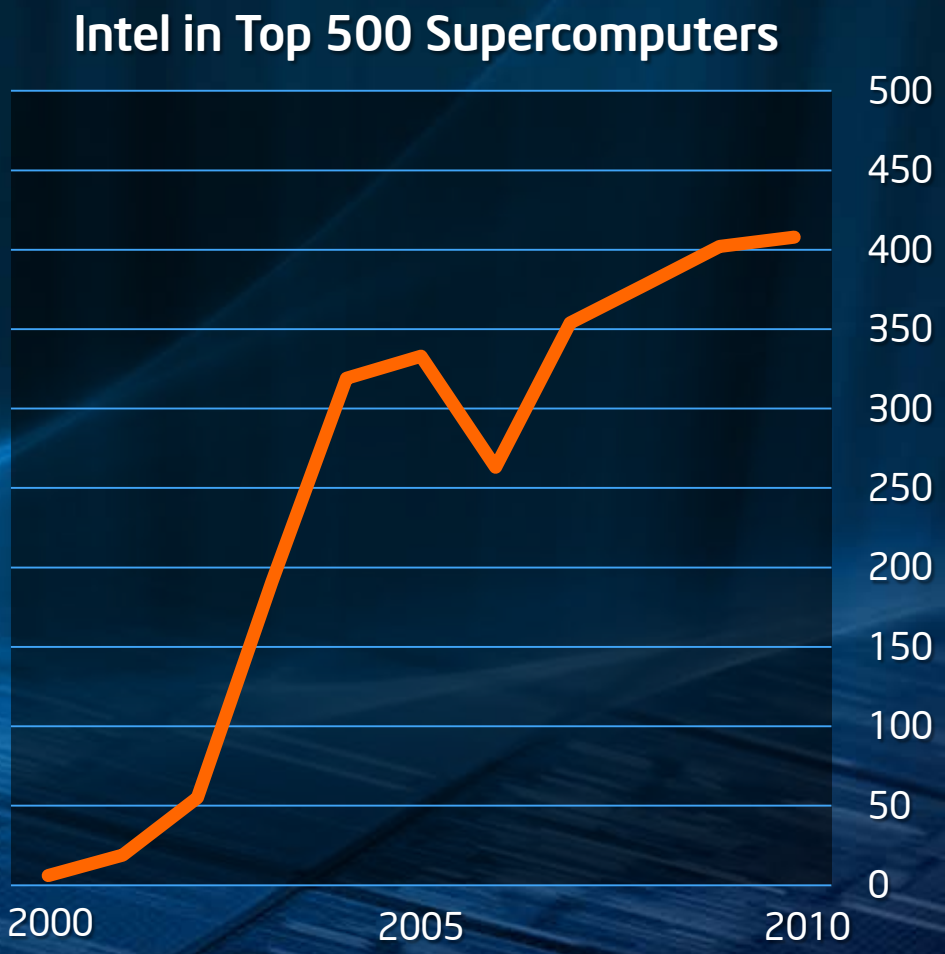

Source: www.top500.org

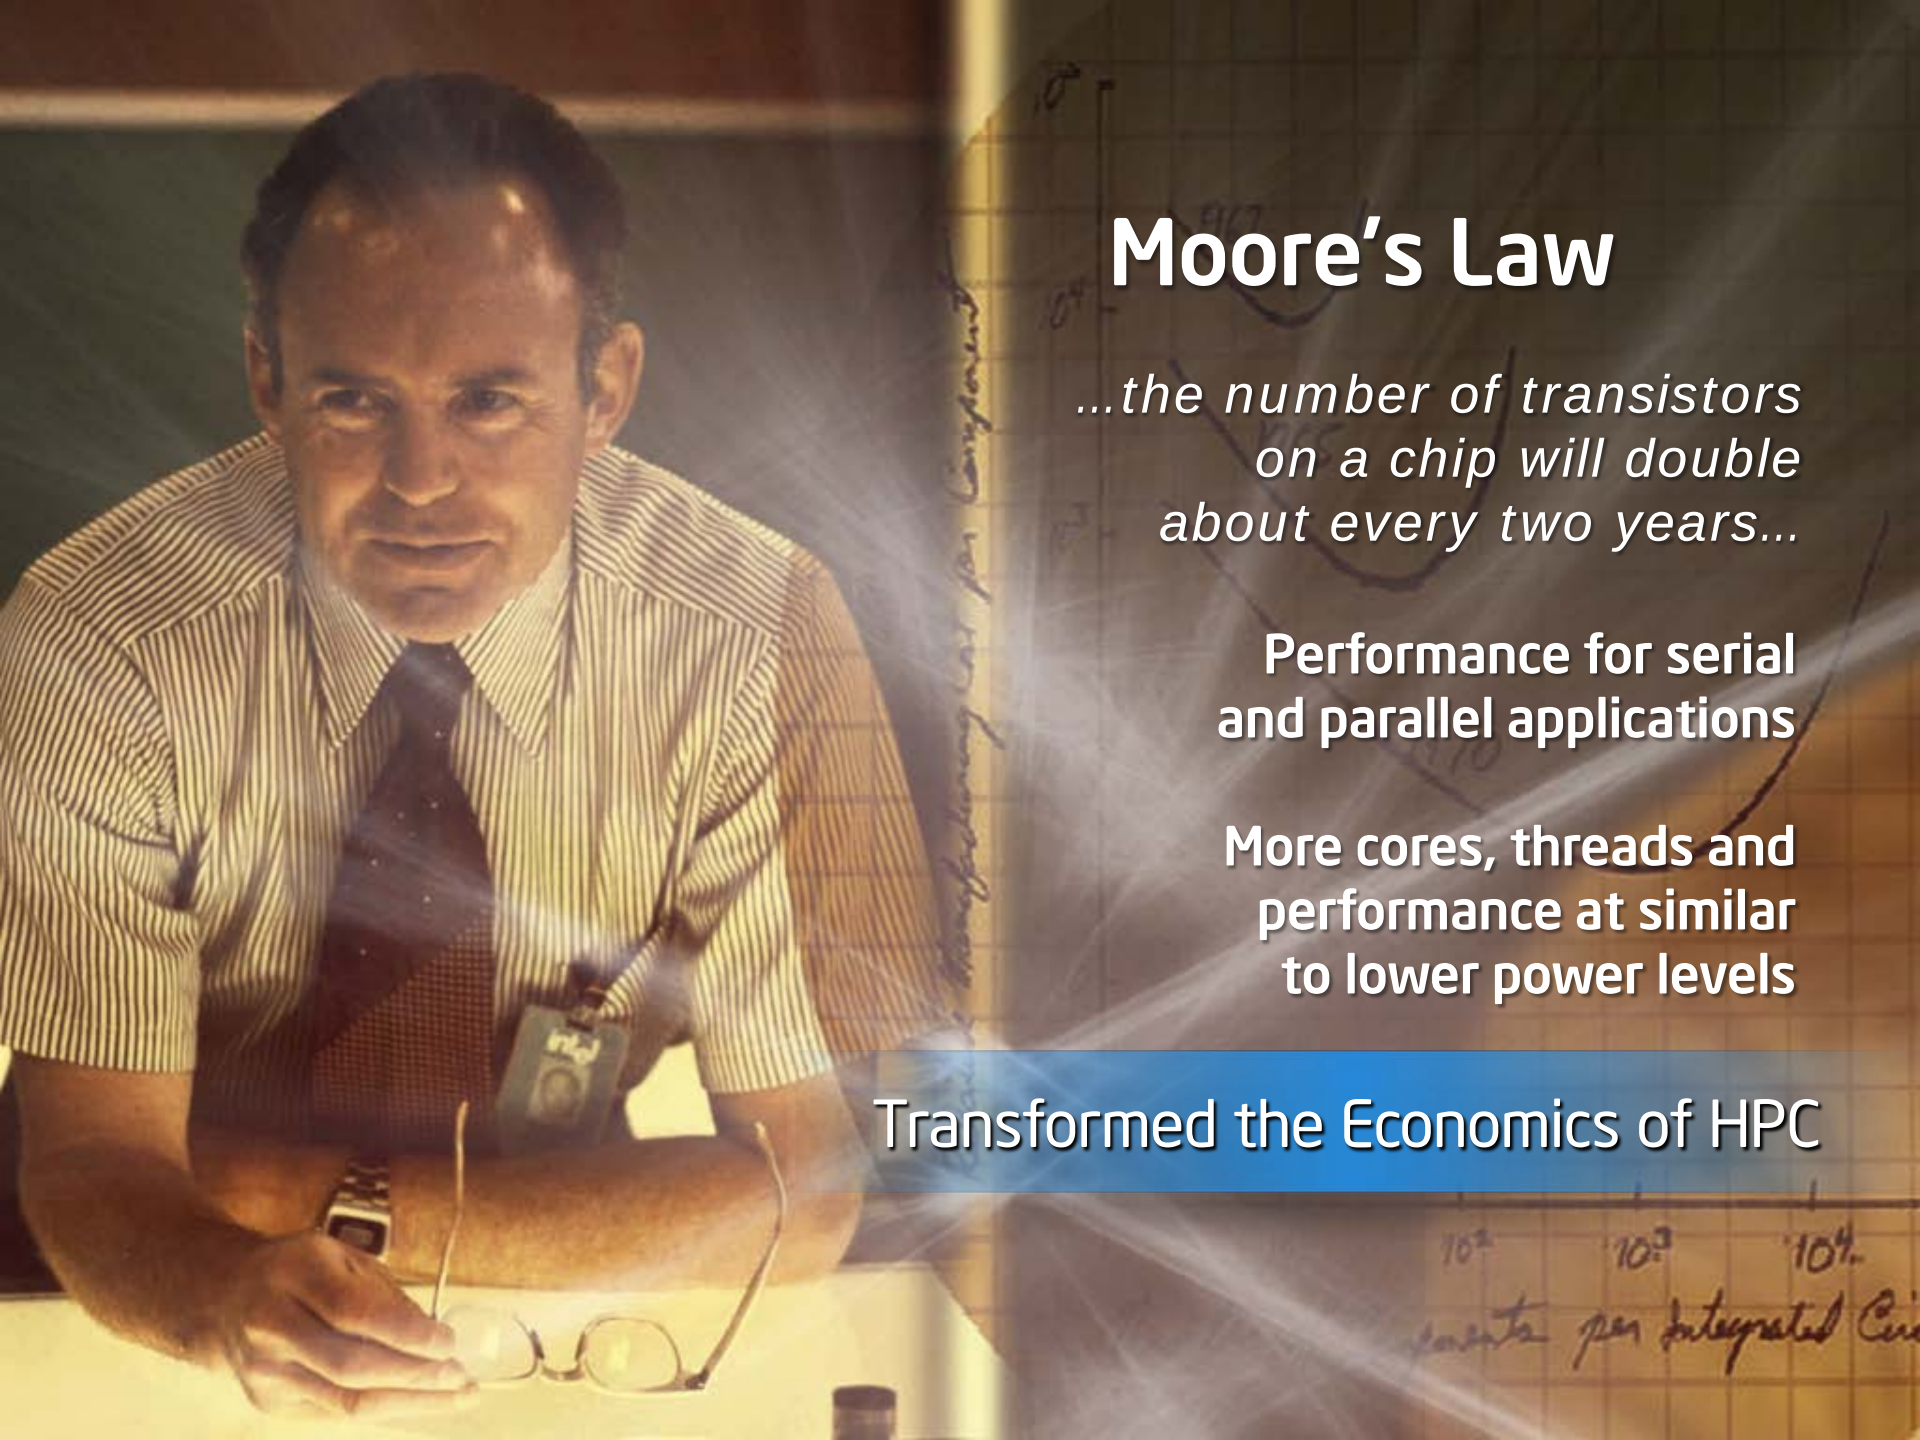

Moore's Law

...the number of transistors on a chip will double about every two years...

Performance for serial and parallel applications

More cores, threads and performance at similar to lower power levels

Transformed the Economics of HPC

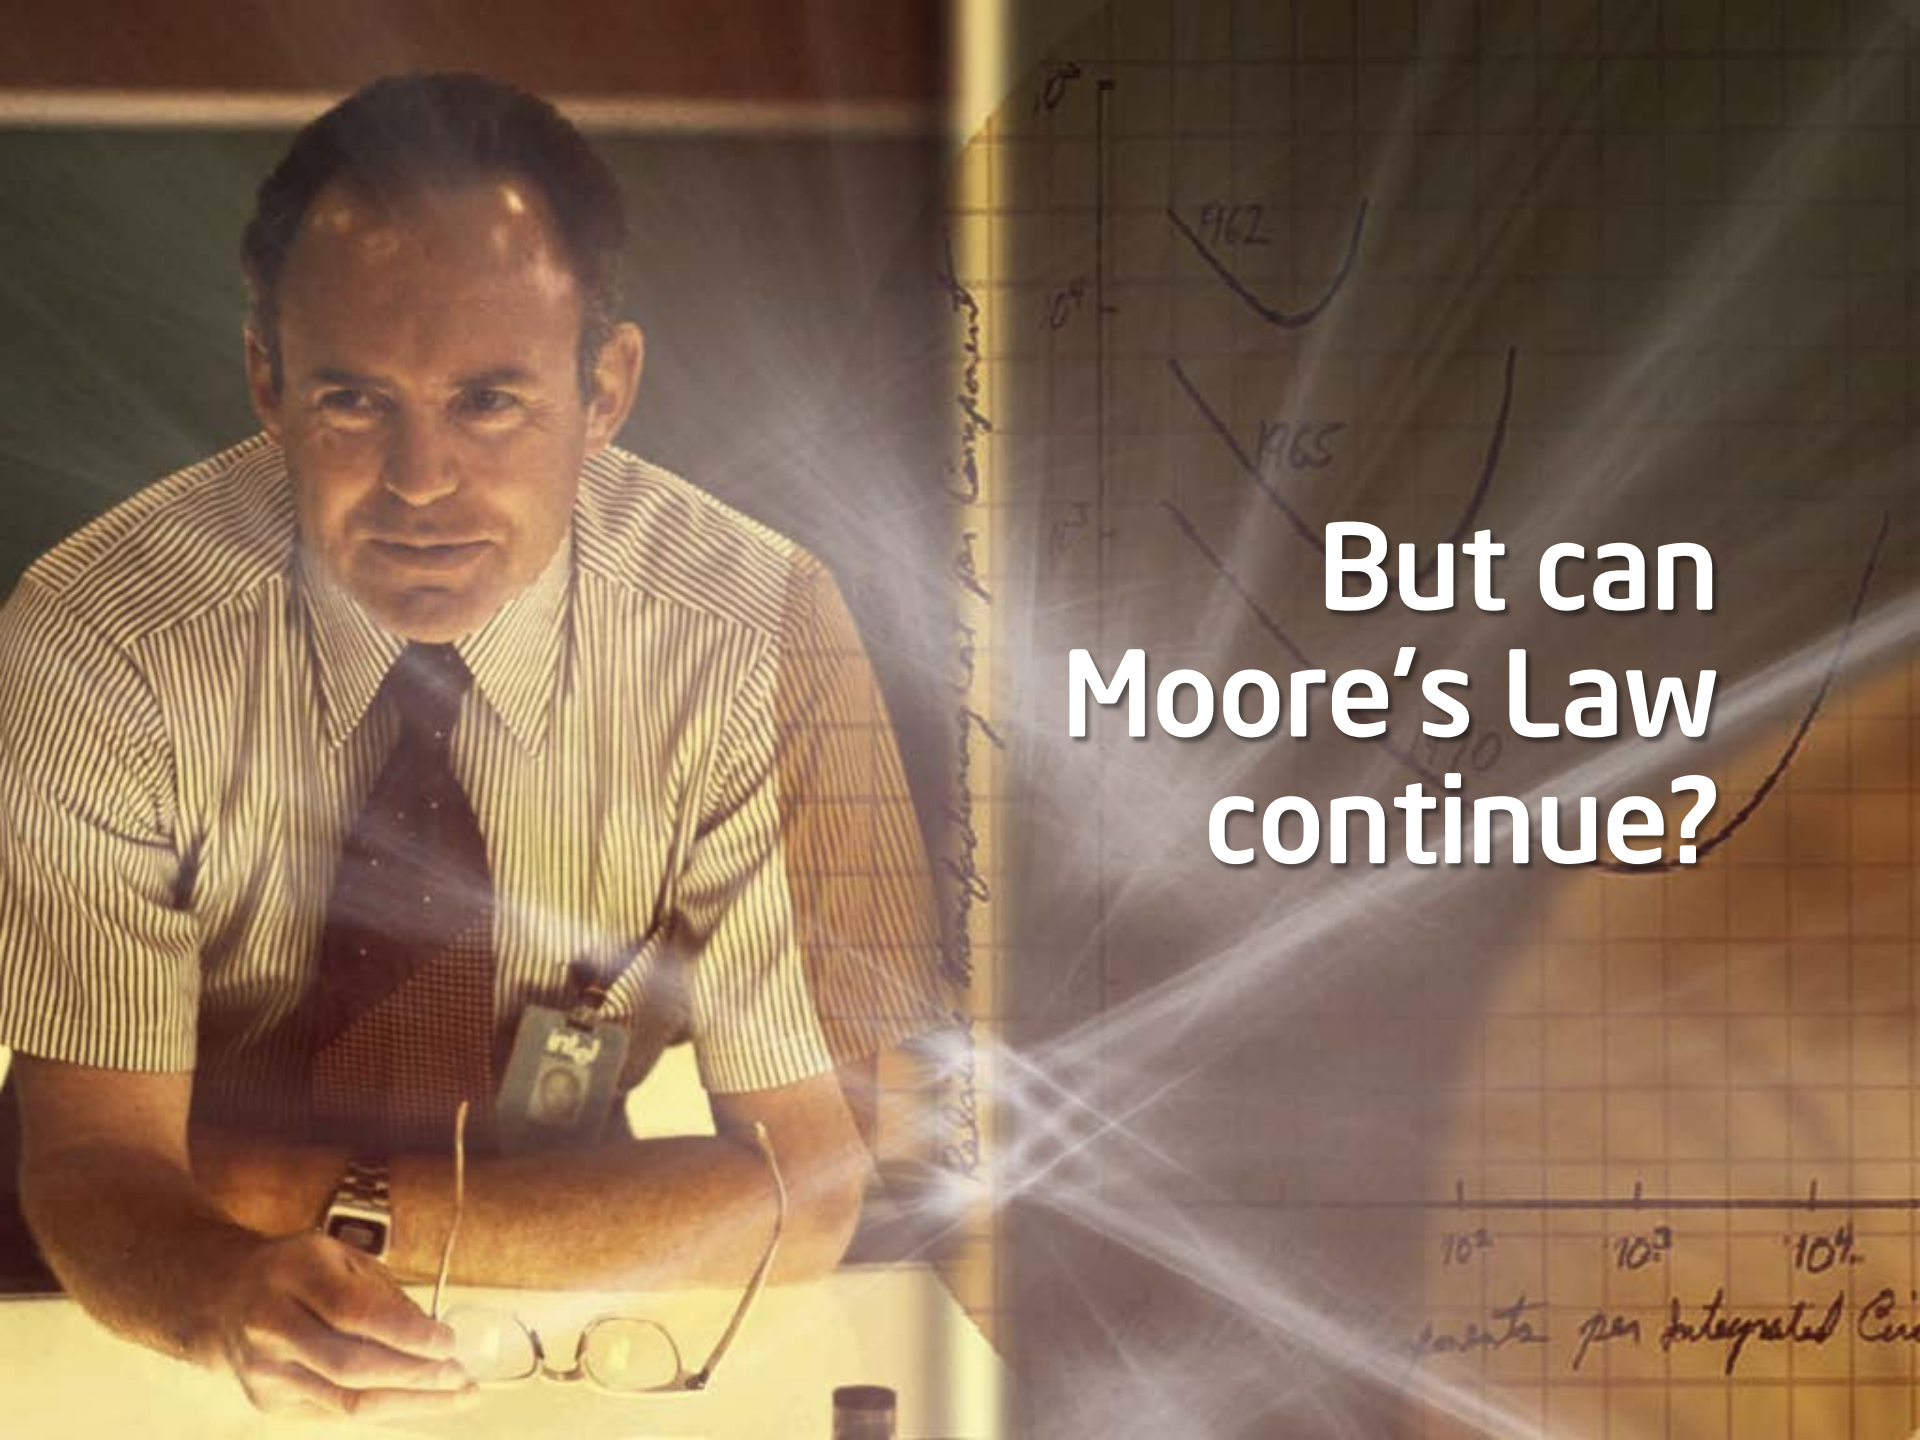
But can
Moore's Law
continue?

Moore's Law: Alive and Well at Intel

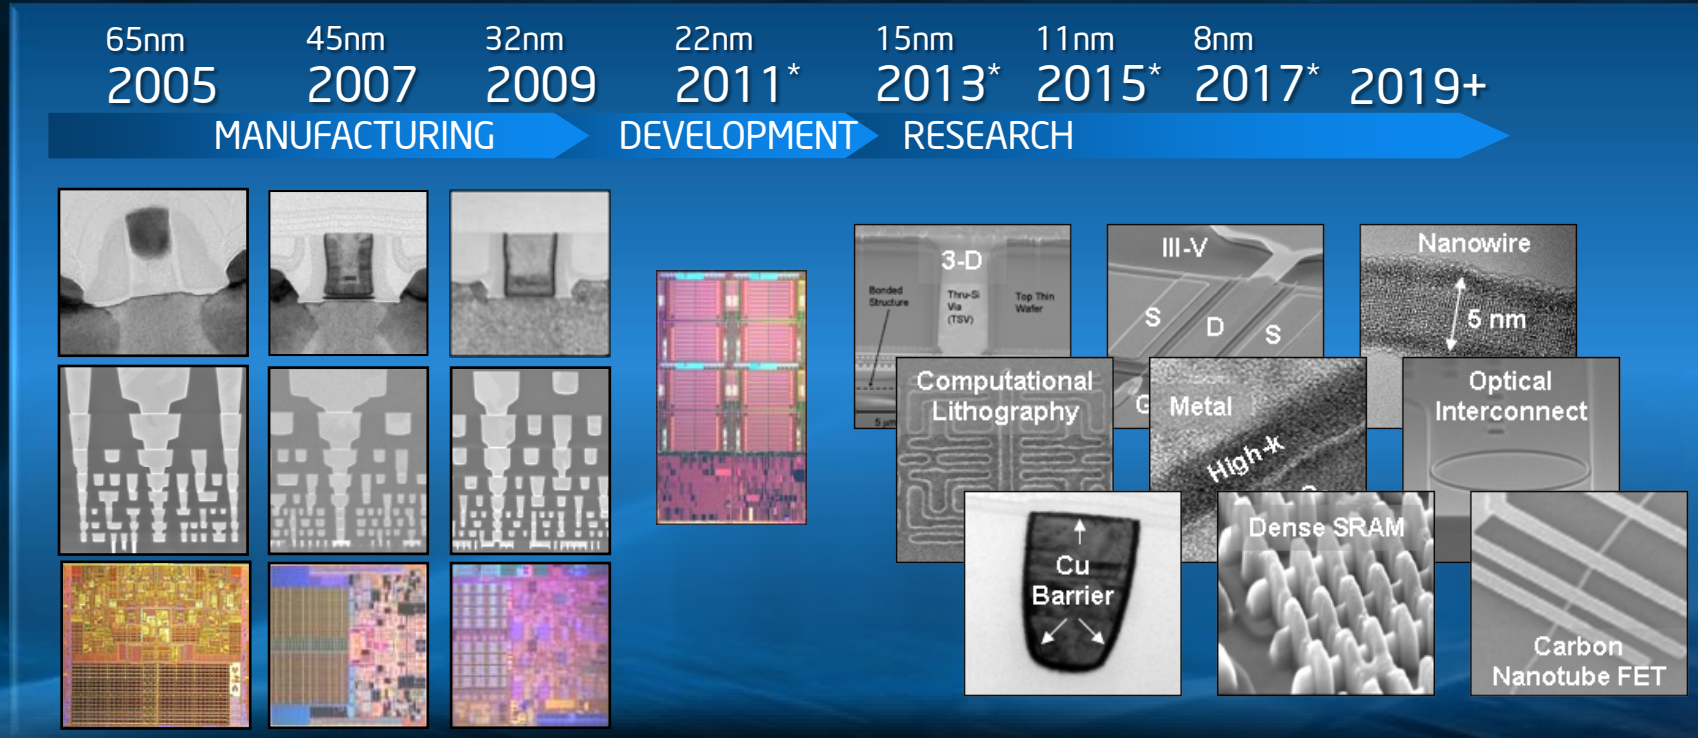

Intel Innovation-Enabled Technology Pipeline is Full

intel

Still an Insatiable Need for Computing

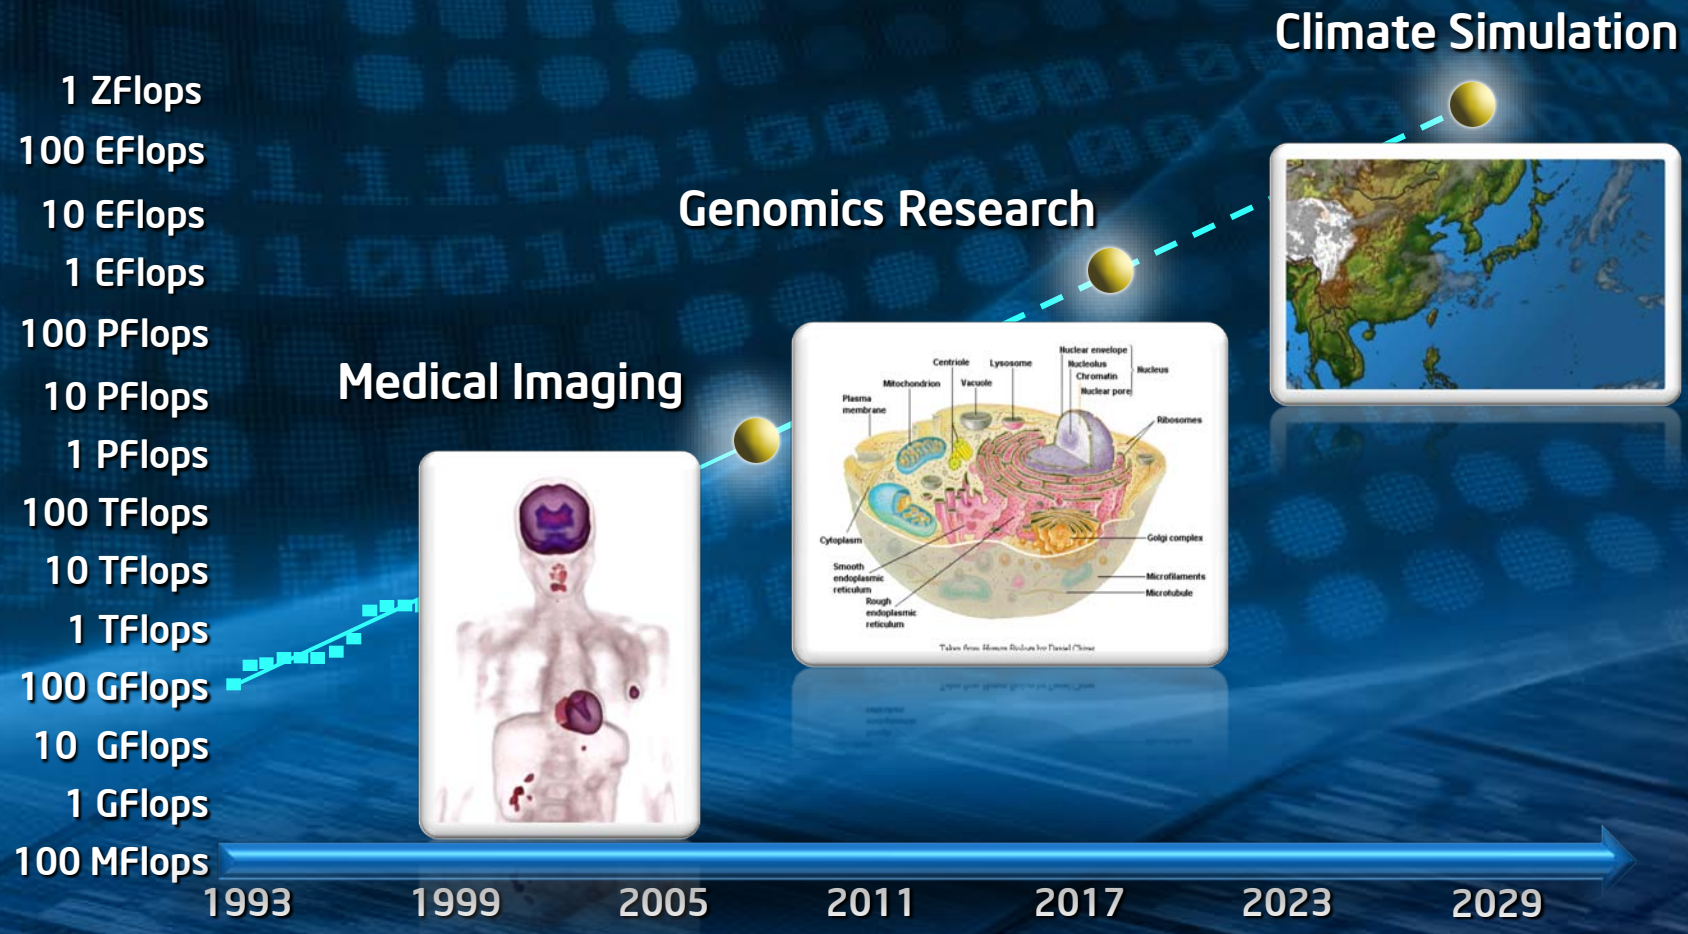

Source: www.top500.org

High Performance Micro-Architecture for PetaScale Deployments

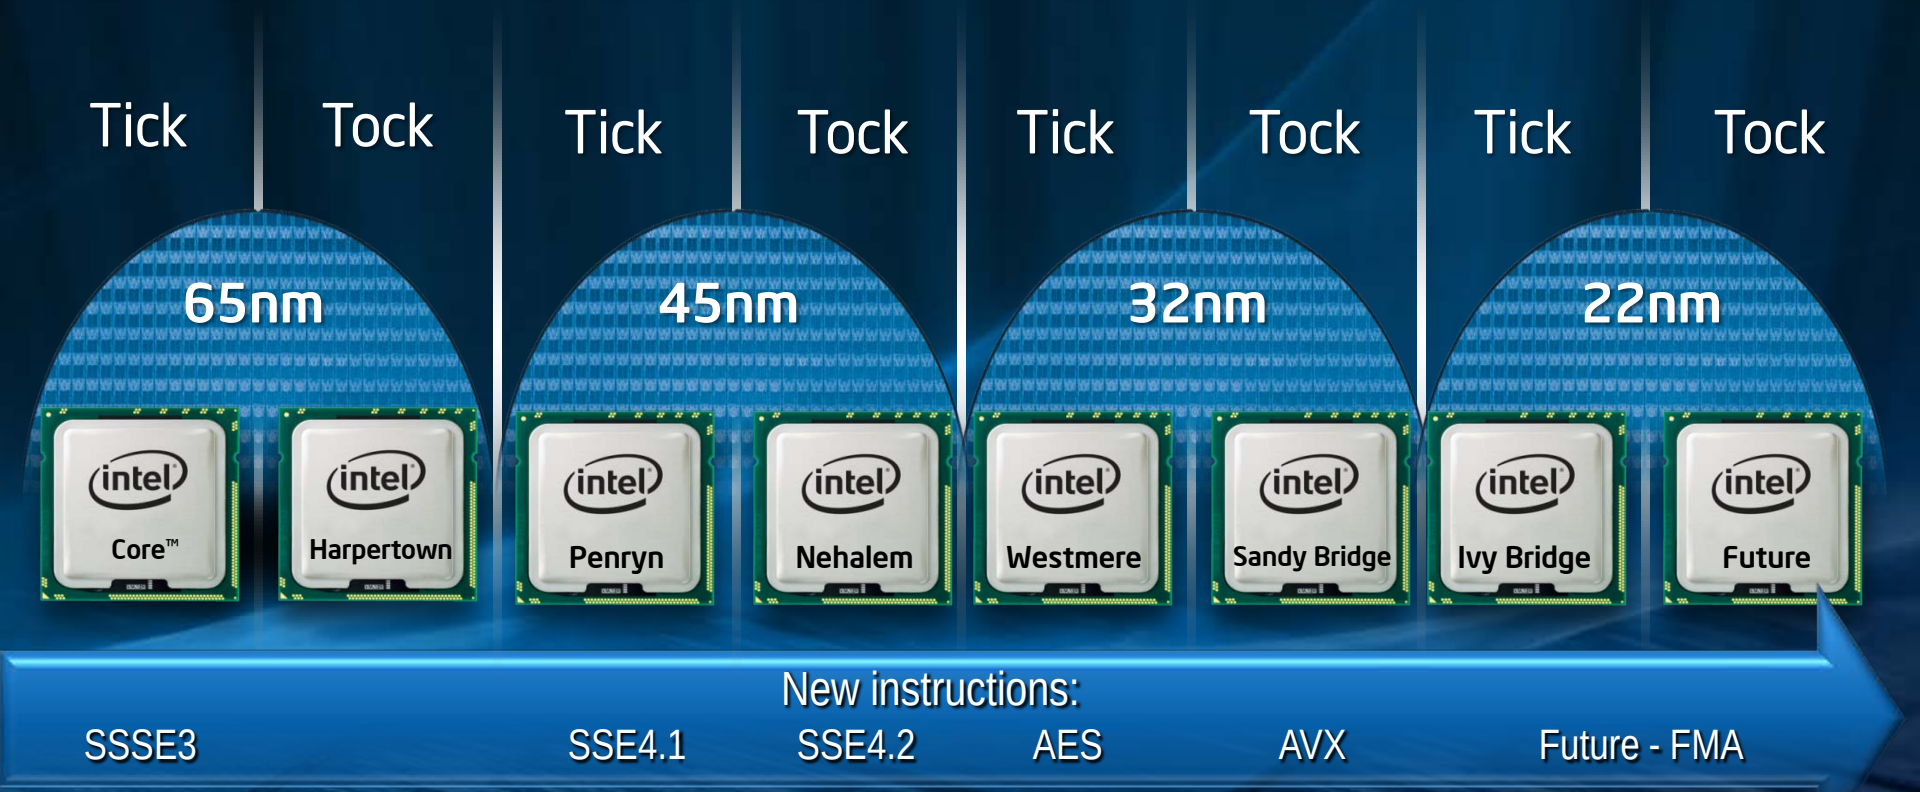

Intel® Xeon® 7500

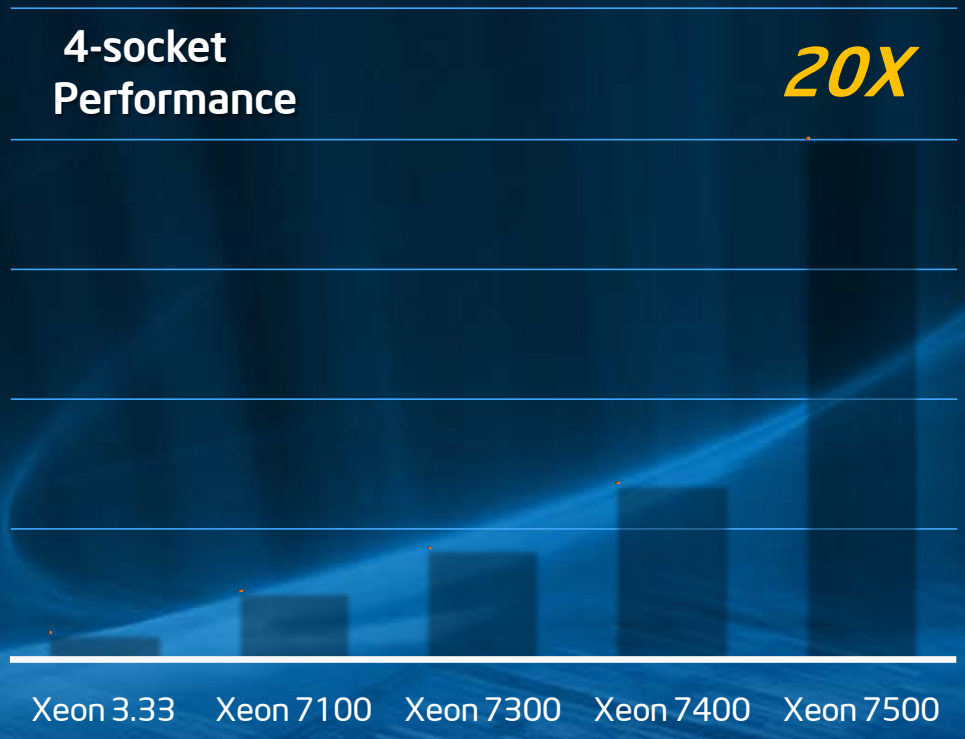

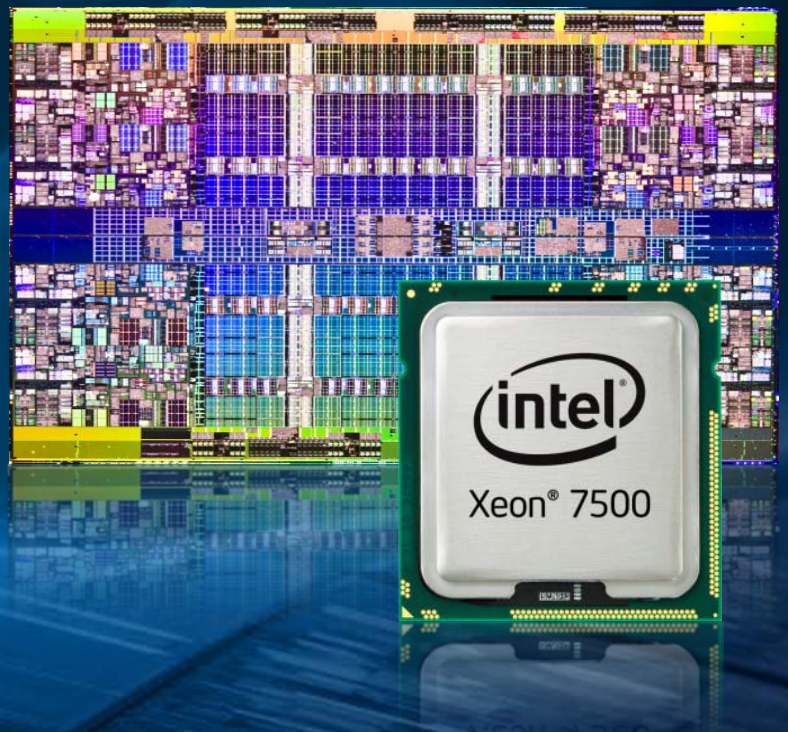

Source: Intel internally measured results 15 January 2010. Each bar represents the score or estimated score of best measured/estimated results on the geometric mean of internal benchmarks (server-side Java*, integer throughput, floating-point throughput, ERP, and OLTP). Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. 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, Go to: http://www.intel.com/performance/resources/benchmark_limitations.htm.
Relative performance is calculated by assigning a baseline value of 1.0 to one benchmark result, and then dividing the actual benchmark result for the baseline platform into each of the specific benchmark results of each of the other platforms, and assigning them a relative performance number that correlates with the performance improvements reported.

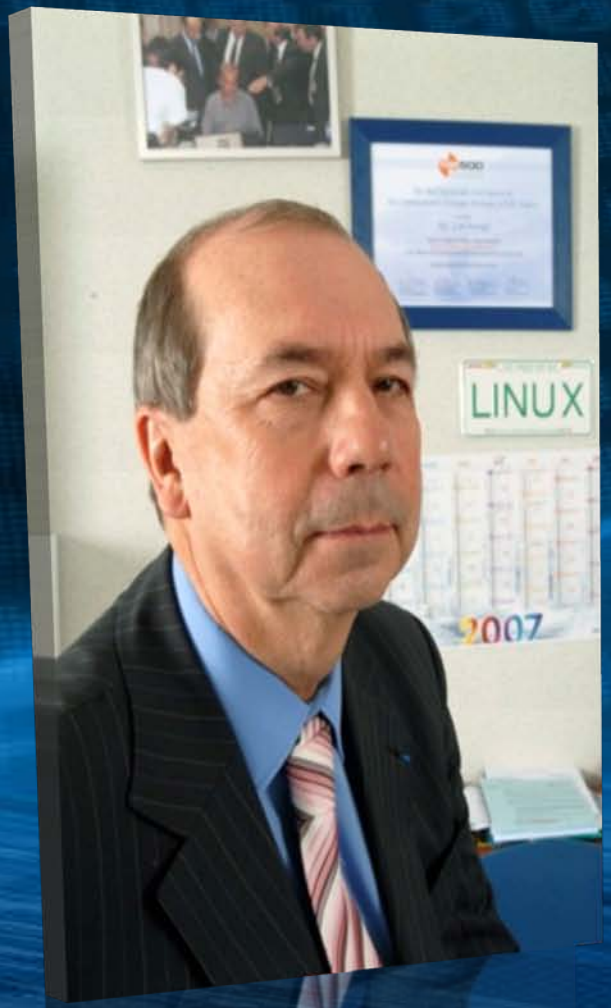

Jean Gonnord

Program Director for Numerical Simulation & Computer Sciences

CEA DAM

energie atomique • energies alternatives

TERA 100

First petaflop/s computer ever designed and built in Europe

Jean Gonnord

Chef de Projet
Simulation numérique

CEA/DAM

Jean Philippe Nominé

Chargé d'affaire HPC
Member of PRACE Technical Board

CEA/DAM

May 27th 2010

Four weeks ahead of initial planning

A significant industrial success

TERA 100 a machine of world records

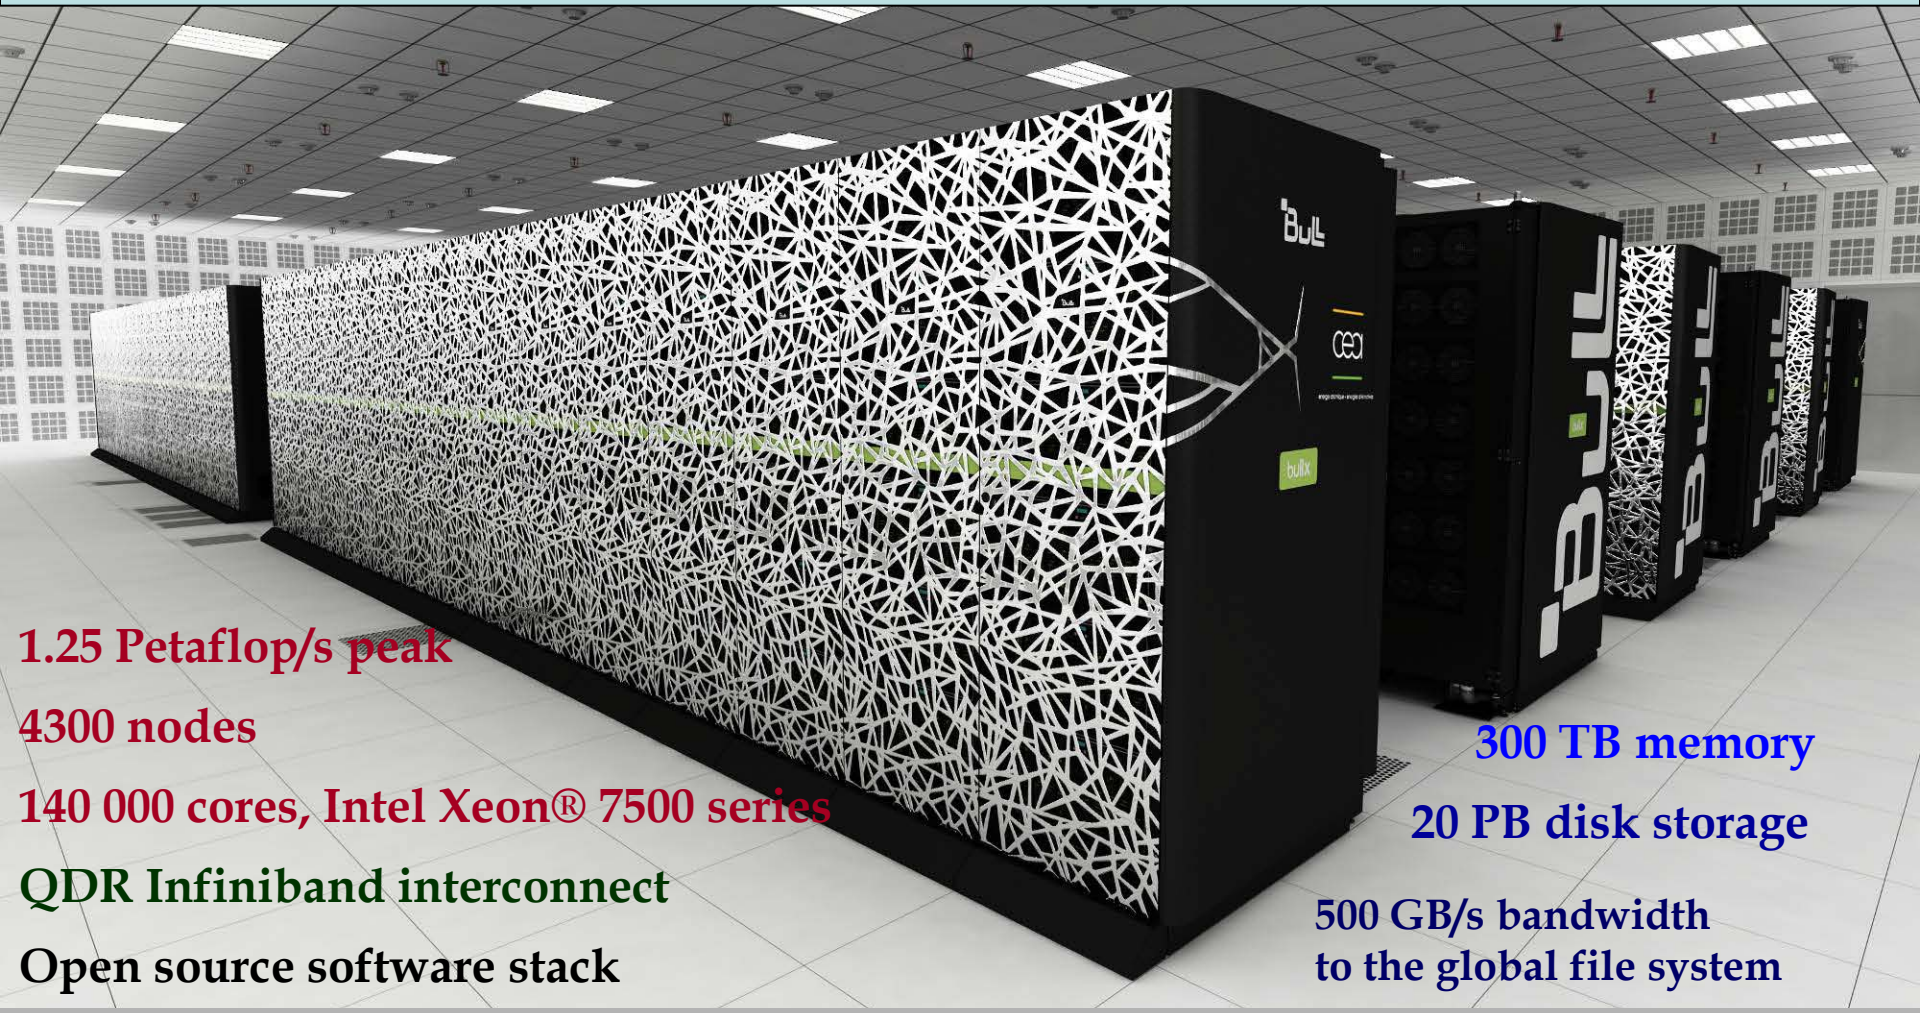

1.25 Petaflop/s peak

4300 nodes

140 000 cores, Intel Xeon® 7500 series

QDR Infiniband interconnect

Open source software stack

300 TB memory

20 PB disk storage

500 GB/s bandwidth
to the global file system

Beside the records,
a production machine, with high level of reliabilty, for CEA strategic needs

TERA 100 a step on the CEA roadmap

CEA : a major actor in the HPC field

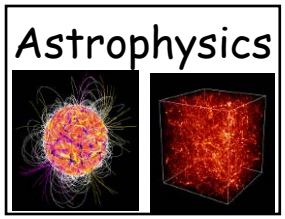

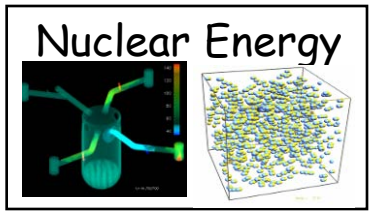

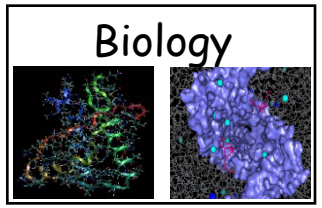

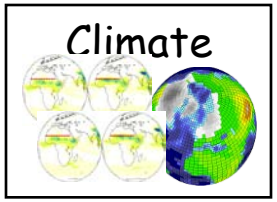

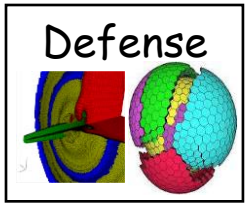

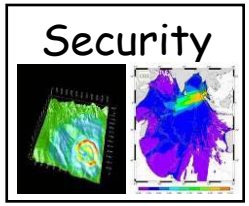

Numerical simulation is an essential tool

TERA 100 a step on the CEA roadmap

CEA : a major actor in the HPC field

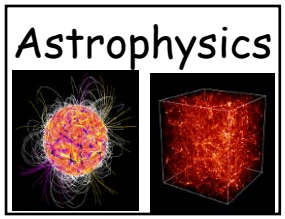

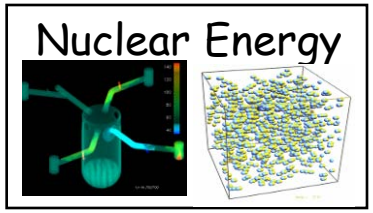

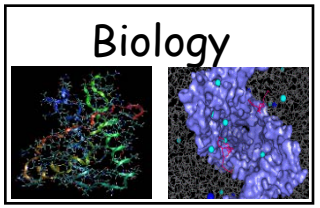

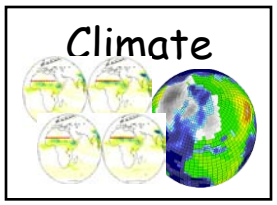

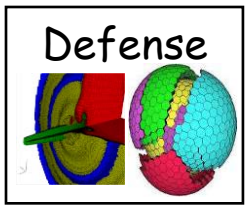

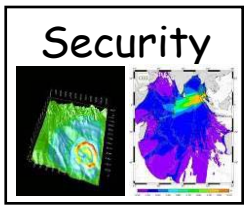

Numerical simulation is an essential tool

2001

N°5 Monde
N°1 EU
4,8 TF
TERA 1

x10

N°5 Monde
N°1 EU
63 TF
TERA 10

x20

N° ? Monde
N°? EU
➢1Pf
➢TERA 100

x30

N° ? Monde
N°? EU
➢30Pf
➢TERA 1000

x30

N° ? Monde
N°? EU
➢ 1 Ef
➢EXA 1

2019

CEA open computing center

CEA classified computing center

TERA 100 a step on the CEA roadmap

CEA : a major actor in the HPC field

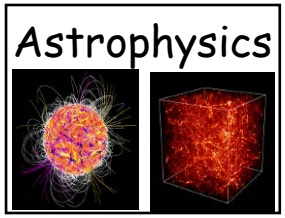

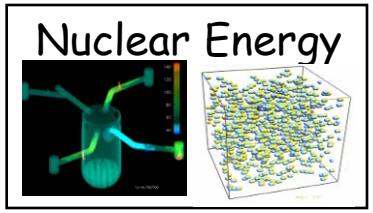

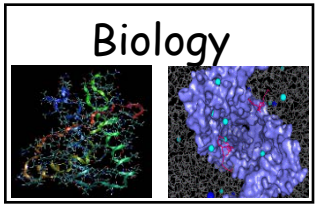

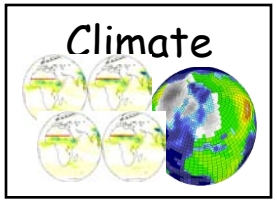

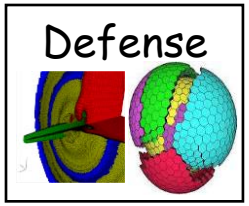

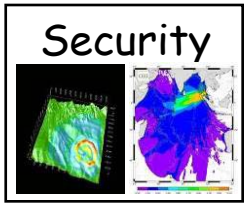

Numerical simulation is an essential tool

2001

N°5 Monde
N°1 EU
4,8 TF
TERA 1

x10

N°5 Monde
N°1 EU
63 TF
TERA 10

x20

N° ? Monde
N°? EU
➢1Pf
➢TERA 100

x30

N° ? Monde
N°? EU
➢30Pf
➢TERA 1000

x30

N° ? Monde
N°? EU
➢ 1 Ef
➢EXA 1

2019

An ambitious roadmap for a strategic goal:

Maintaining the capacity of designing and building very large computing systems in Europe

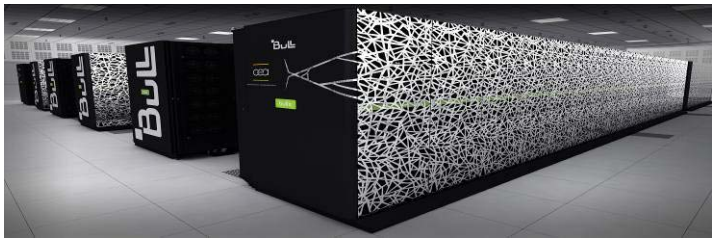

CEA open computing center

CEA classified computing center

CEA/DAM has the operational responsibility of implementing this roadmap

TERA 100 a great thank to INTEL

With a special mention to Richard Dracott

For delivering us on schedule the 18000 NEHALEM-EX chips

For giving us the opportunity of this presentation

Join us tomorrow at 1pm at BULL booth 320 for a drink

We will begin to prepare the future of European HPC with you

energie atomique • energies alternatives

The Next Generation Xeon Processor

Sandy Bridge "Tock"

- Significantly greater performance with higher core-count & Intel® Hyper-threading Technology
- 2x Flops / clock peak using new AVX instructions

Making Petascale Widely Available for Leading Science

intel

Petascale Programming Challenges

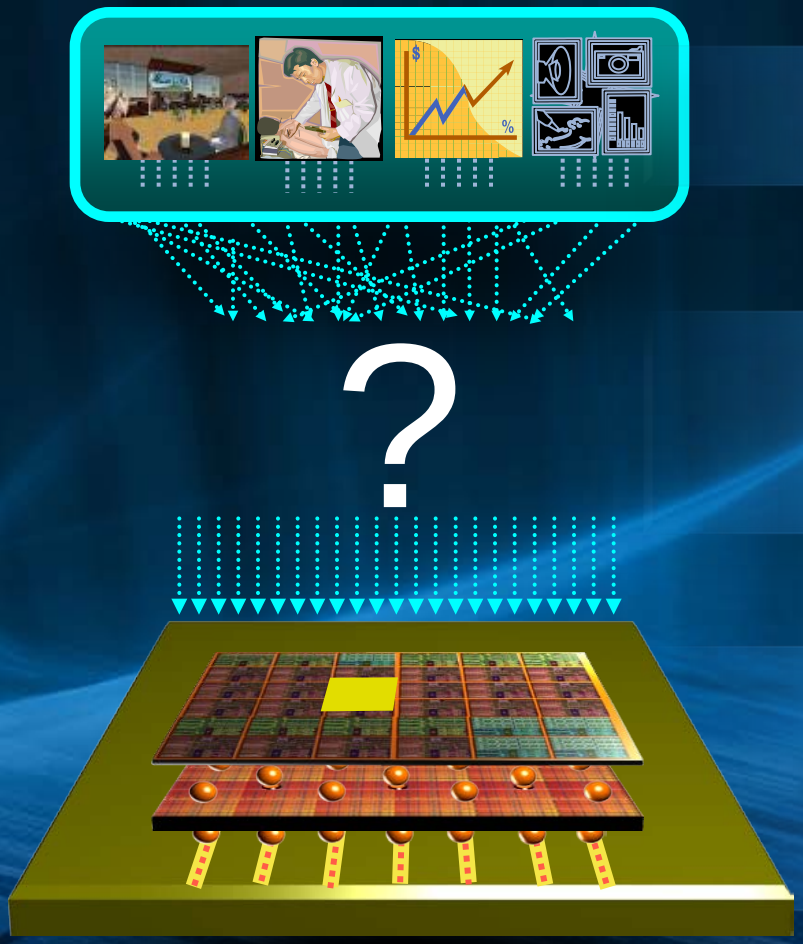

Irregular Patterns and Data Structures

Scale to Multi-Core → Hard
Scale to Many-Core→ Harder

Increasing number of cores & threads
Vector instructions

IA Programming Flexibility

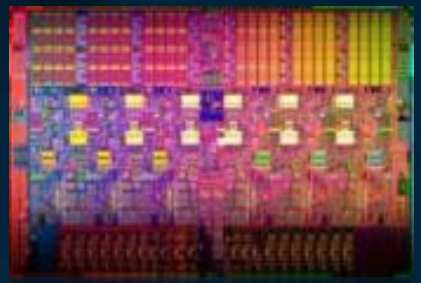

intel® Software

Instruction Parallelism	Data Parallelism	Thread Parallelism	Cluster / Process Parallelism

Serial Code Node Level
Fast Scalar performance, Optimized C/C++,FORTRAN, Threading and Performance Libraries, Debug / Analysis Tools

Parallel Node Level
Multi-core, Multi-Socket, SSE and AVX instructions, OpenMP, Threading Building Blocks, Performance Libraries, Thread Checker , Ct , Cilk

Multi-Node / Cluster Level
Cluster Tools, MPI Checker

Programming choices and standards for range of parallel efficiency

intel

Simplifying Software Development:
Intel® Software Development Tools

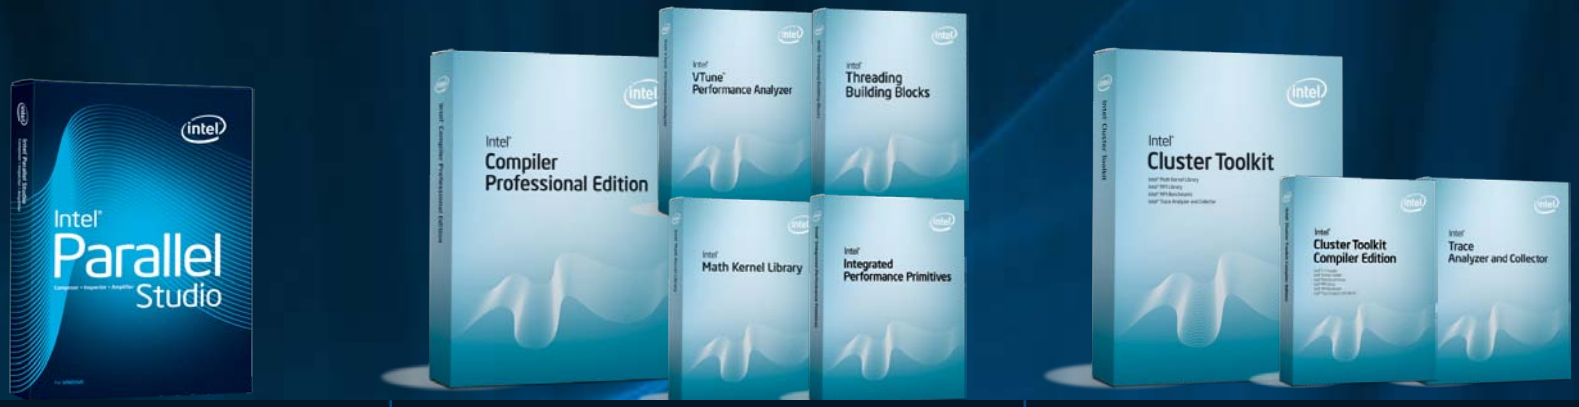

Parallel Studio	HPC Tools	Cluster Tools
Essential Parallelism	Advanced Parallelism	Distributed Parallelism

Tools to preserve your source code investments

Parallel Programming Education

2K universities in 88 countries

4K faculty trained

320K students trained

intel.com/thinkparallel

Source: Evans Data Corporation, Intel

Exascale: *The Next Frontier*

Challenges

- Power – energy / operation of computation, data transport, memory
- Threading software to millions/billions of threads
- Memory/Storage capacity and bandwidth
- Managing high-node count systems in the existence of failures (MTBF)
- Affordability

Energy

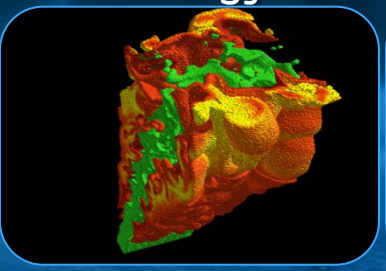

Physics

Biology

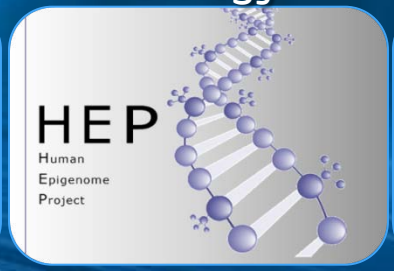

Climate

Astrophysics

Intel committed to solving the challenges of Exascale

intel

Intel Co-Sponsored HPC Labs in Europe

Introducing Today

ExaTec Lab, Paris

Performance and scalability of Exascale applications

ExaCluster Lab, Jülich

Exascale cluster scalability and reliability

Advancing Exascale Computing on Intel Architecture

Dr. Pradeep Dubey

Senior Principal Engineer
IEEE Fellow

Director of the Throughput
Computing Lab, Intel Labs

intel

Intel's Many-Core Research Program

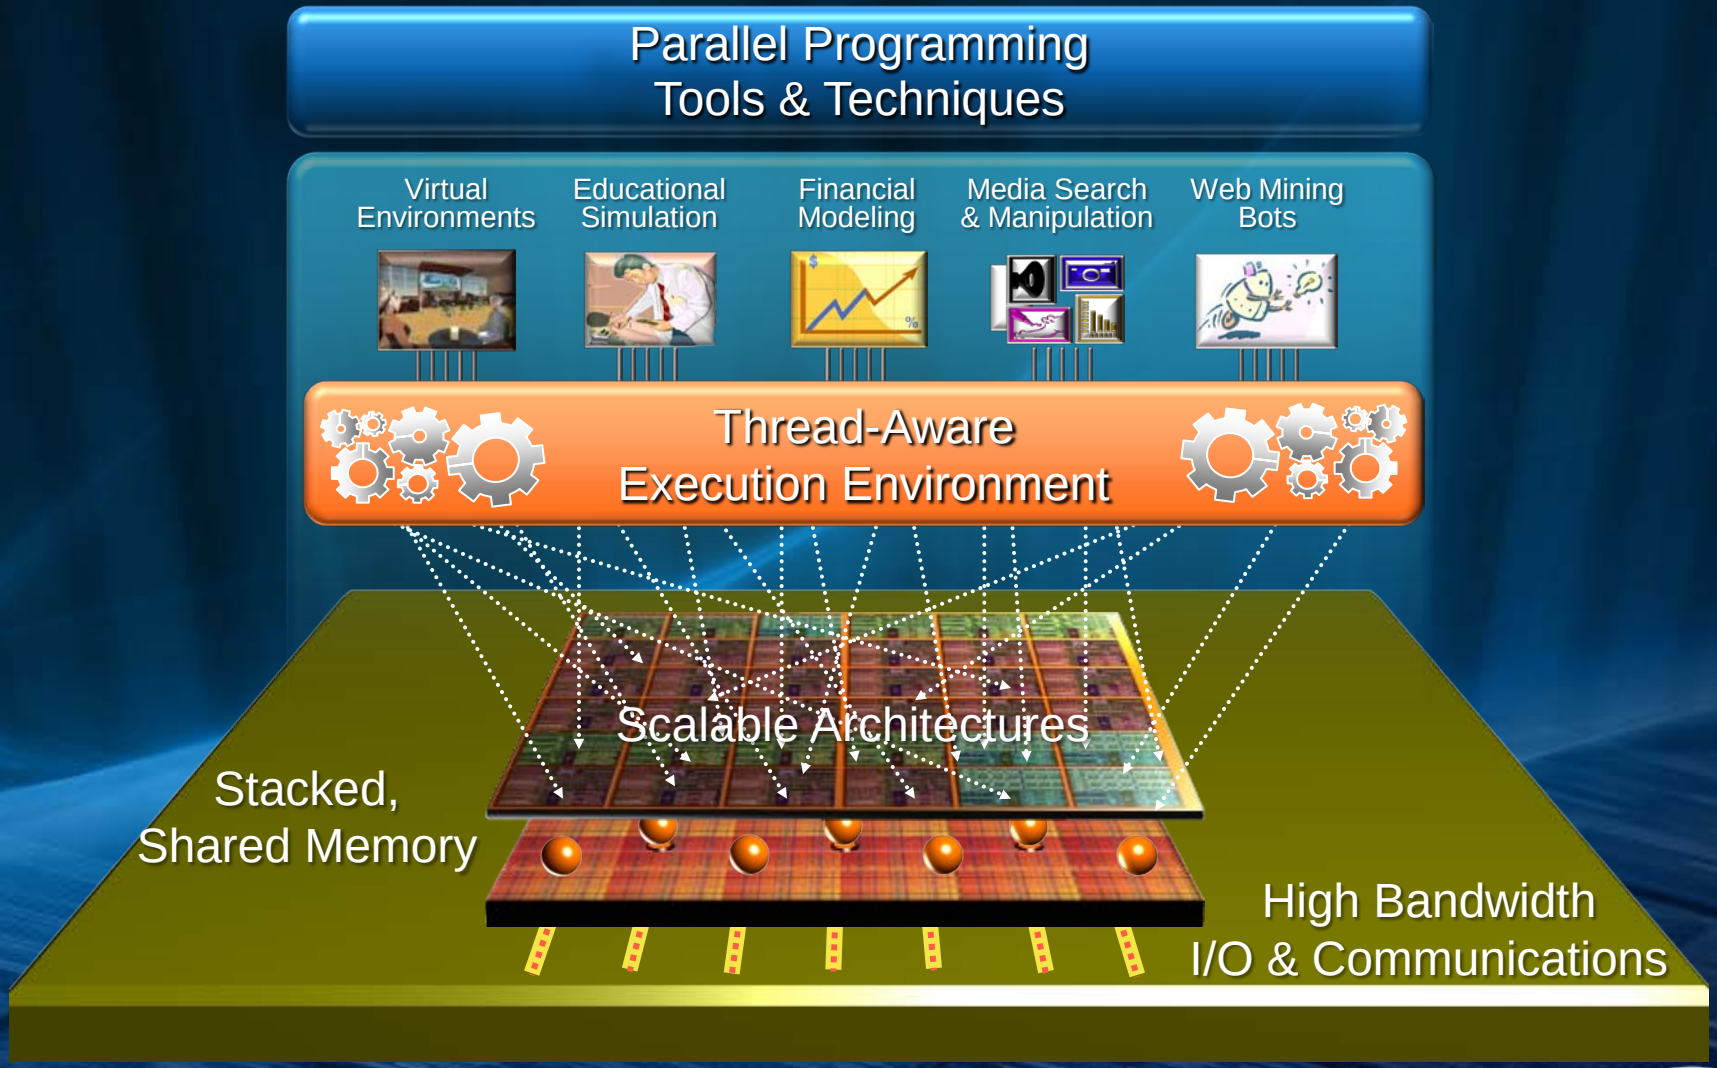

Application-Driven Architecture Research

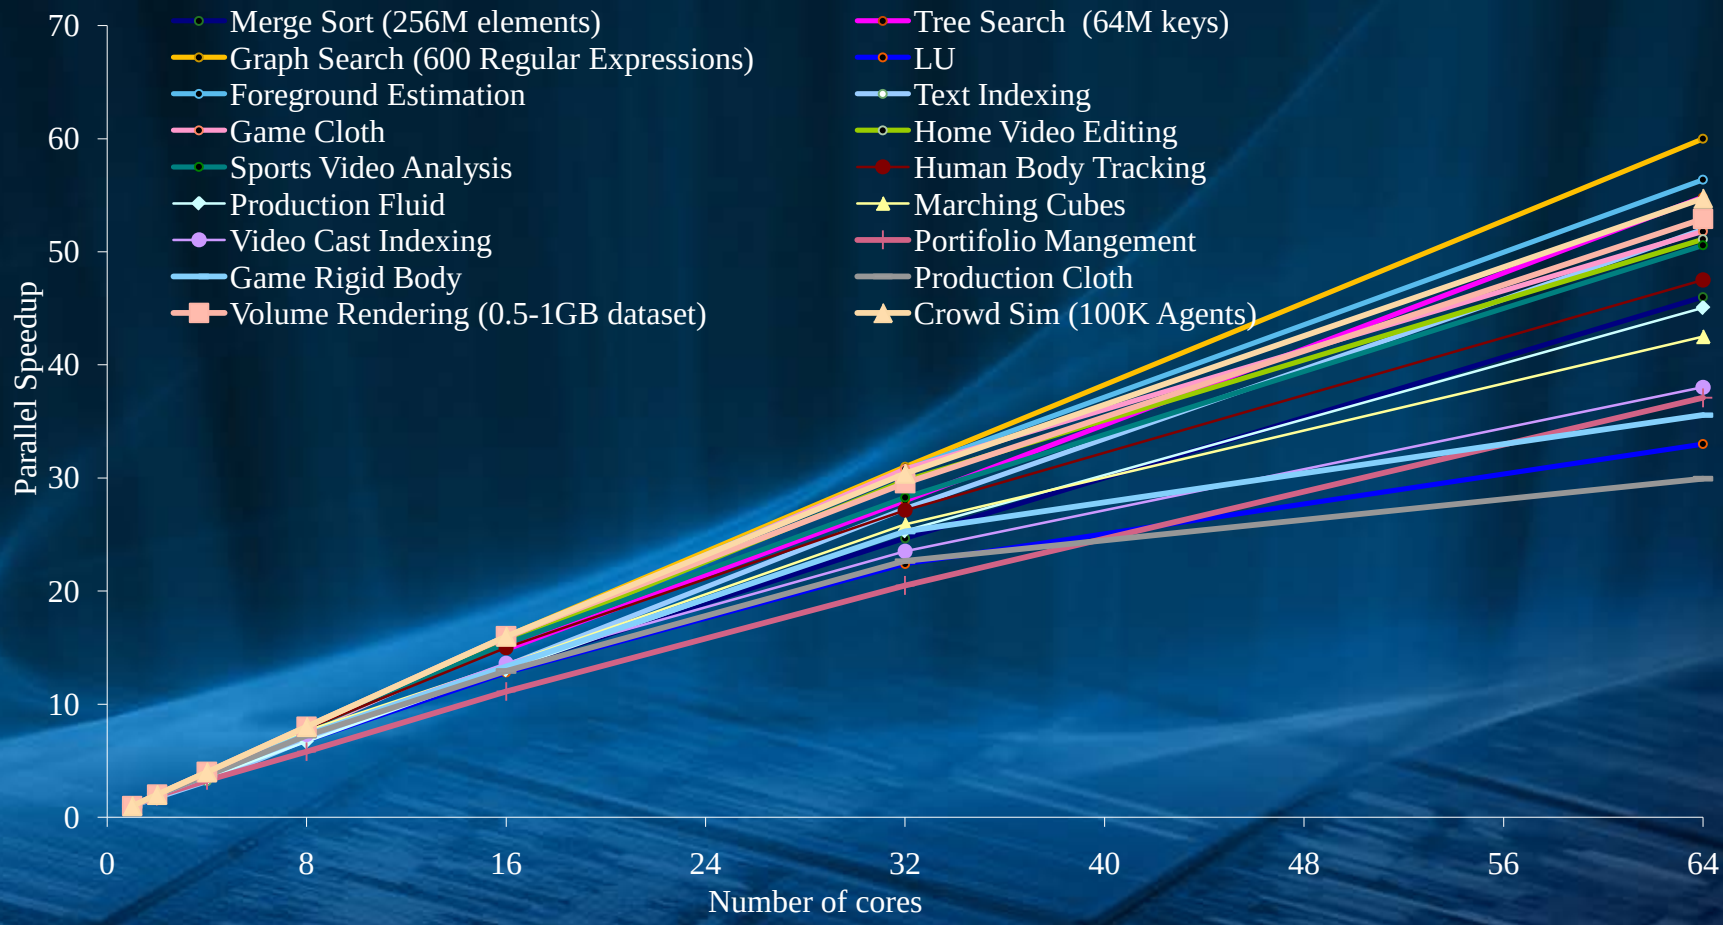

Constantly Evaluating Options for All Workloads

intel

Intel Labs
Parallel Computing Research

Research Processors from Intel Labs

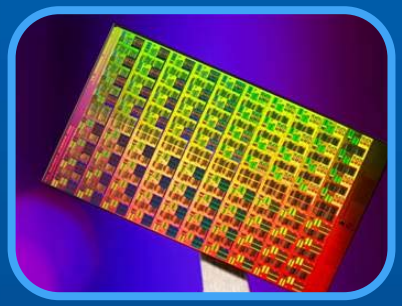

Tera-scale Research Processor
Mar 2007

Single Chip Cloud Computer
Dec 2009

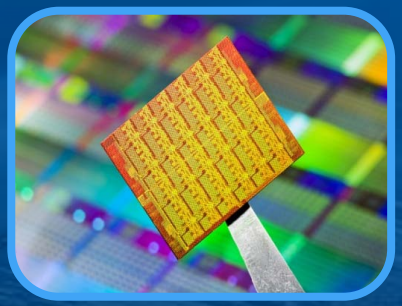

Intel Labs
Parallel Computing Research

Research Processors from Intel Labs

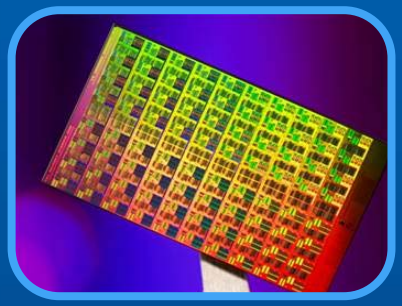

Tera-scale Research Processor
Mar 2007

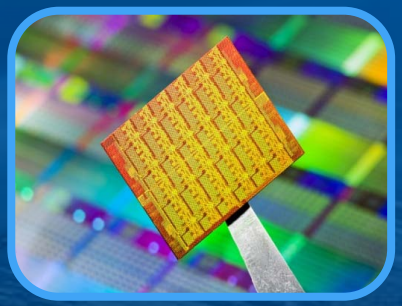

Single Chip Cloud Computer
Dec 2009

1999 - 2006

Origination of Intel's multi-core explorations

Coming Challenges in Microarchitecture and Architecture

"Era of Tera" Keynote at Intel Developer Forum

Recognition, Mining , Synthesis Moves Computers to the Era of Tera

Hundreds of Cores: Scaling to Tera-scale Architecture

Few Cores to Many: A Tera-scale Computing Research Overview

2007

Demonstration of Intel experimental 80-core processor

"Ct" language proposal for Tera-scale Architectures

Intel® C++ STM Compiler Prototype release

Integration Challenges and Tradeoffs for Tera-scale Architectures

Package Technology to Address the Memory Bandwidth Challenge for Tera-scale Computing

Runtime Environment for Tera-scale Platforms

Architectural Support for Fine-Grained Parallelism on Multi-core

Datacenter-on-Chip Architectures: Tera-scale Opportunities and Challenges in Intel's Manufacturing Environment

Media Mining—Emerging Tera-scale Computing Applications

High-Performance Physical Simulations on Next-Generation

Intel Labs
Parallel Computing Research

Research Processors from Intel Labs

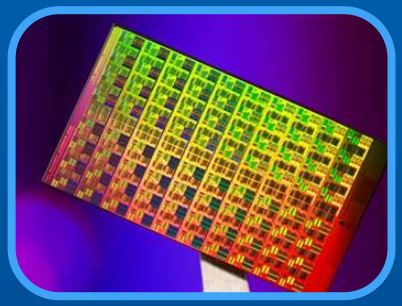

Tera-scale Research Processor
Mar 2007

Single Chip Cloud Computer
Dec 2009

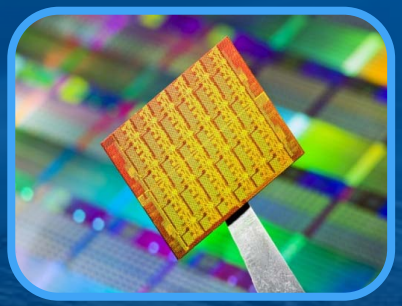

Architecture with Many Cores

Demonstrated Intel McRT ("Manycore Runtime") 11 Issue 03

Carbon: Architectural Support for Fine-Grained Parallelism on Chip Multiprocessors

Physical Simulation for Animation and Visual Effects: Parallelization and Characterization for Chip Multiprocessors

Scaling performance of interior-point method on large-scale chip multiprocessor system

2008

Second Life and the New Generation of Virtual Worlds

Larrabee: A Many-Core x86 Architecture for Visual Computing

Atomic Vector Operations on Chip Multiprocessors

Efficient Implementation of Sorting on Multi-Core SIMD CPU Arch

Convergence of Recognition, Mining, and Synthesis Workloads and Its Implications

Accelerating Video-Mining Applications Using Many Small, General-Purpose Cores

2009

Mapping High-Fidelity Volume Rendering for Medical Imaging to CPU, GPU and Many-Core Architectures

From Research to Realization. Announcing...

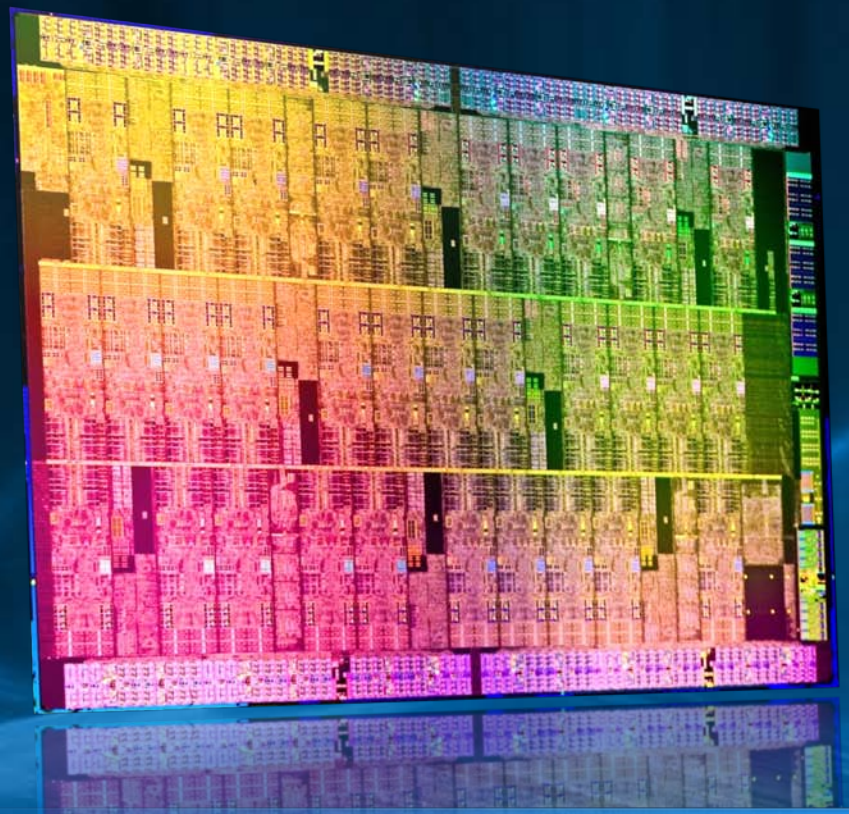

Intel® Many Integrated Core Architecture

The Newest Addition to the Intel Server Family.
Industry's First General Purpose Many Core Architecture

intel

Intel® MIC Architecture:
An Intel Co-Processor Architecture

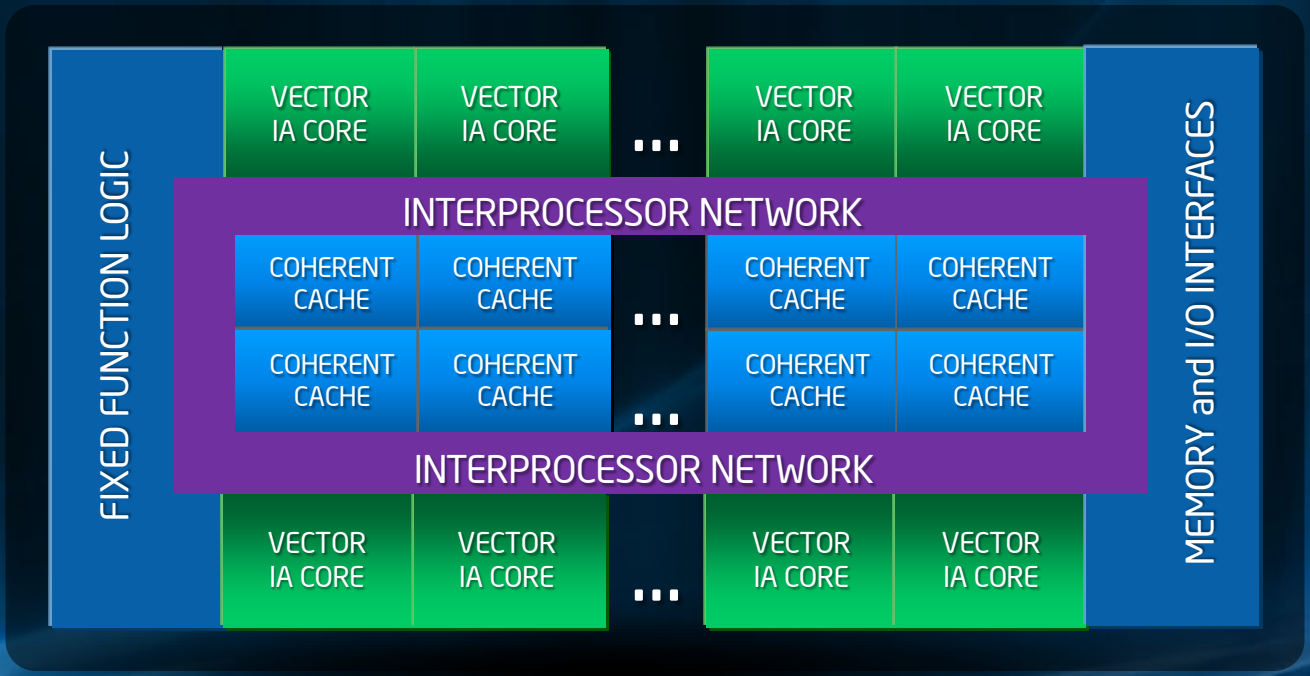

Many cores and many, many more threads

Standard IA programming and memory model

Knights Ferry

- Software development platform
- Growing availability through 2010
- 32 cores, 1.2 GHz
- 128 threads at 4 threads / core
- 8MB shared coherent cache
- 1-2GB GDDR5
- Bundled with Intel HPC tools

Software development platform for Intel® MIC architecture

intel

Intel® MIC Architecture Programming

Single Source

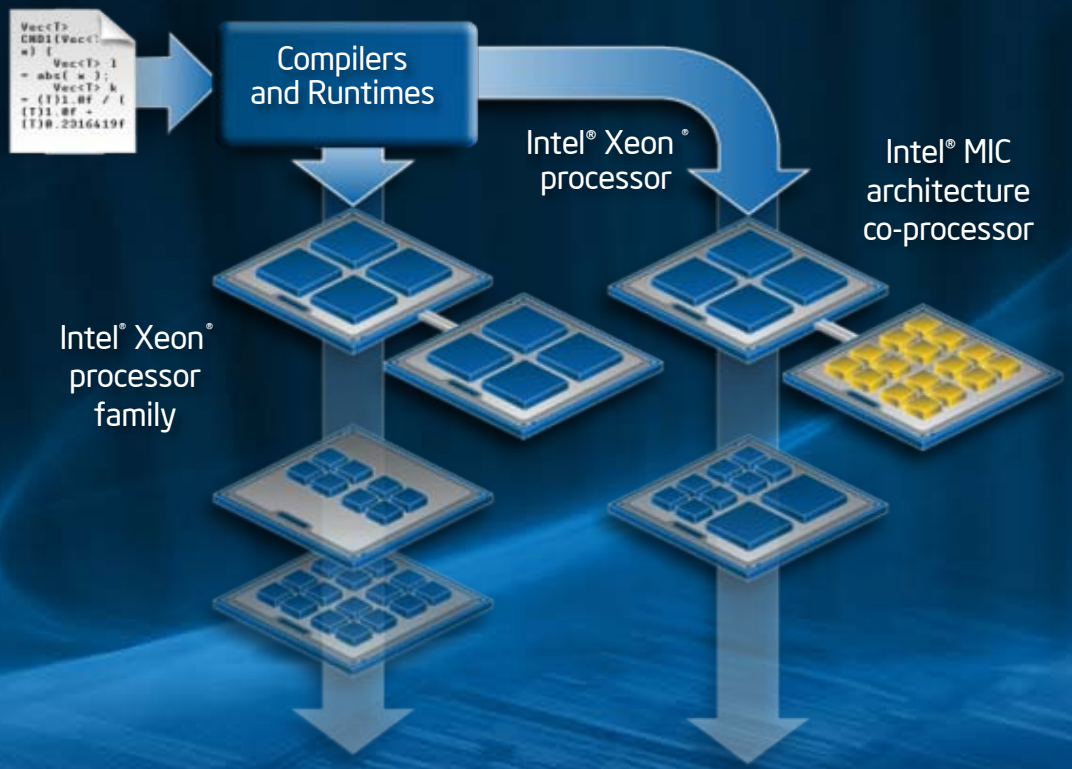

Common with Intel® Xeon®

- Languages
- C, C++, Fortran compilers
- Intel developer tools and libraries
- Coding and optimization techniques
- Ecosystem support

Eliminates Need for Dual Programming Architecture

intel


Knights Ferry Demo
intel

Summary

11/09: Leading performance SGEMM (>1 Teraflop)

11/09: Leading performance SpMVM

Today: Leading performance LU (>½ Teraflop)

The Knights Family

Future Knights Products

Knights Corner

1st Intel® MIC product

22nm process

>50 Intel Architecture cores

Knights Ferry

intel

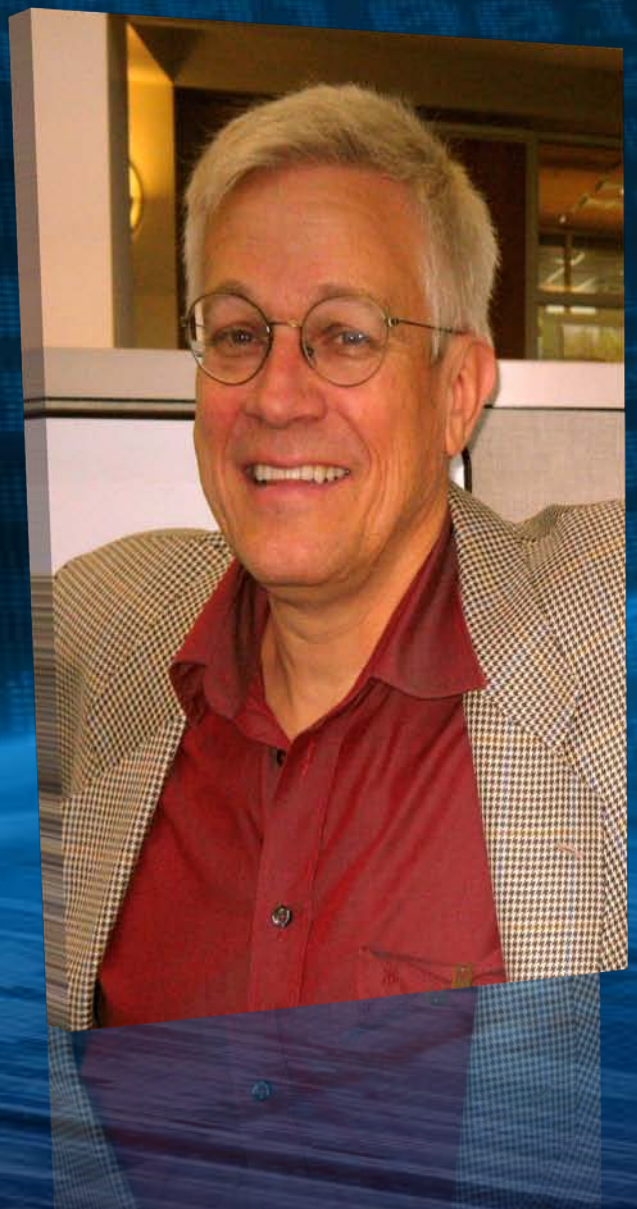

Sverre Jarp

Chief Technical Officer

CERN Openlab

CERN's Large Hadron Collider

- LHC is 27 km in circumference, 100 m underground, and operates at 1.9º Kelvin

- It has now been up and running since November 2009

- World record in beam energy
  - 3.5 T achieved as of 30 March
  - By now (May 2010): Over 1'000'000'000 events recorded

Four experiments, with detectors as 'big as cathedrals':
ALICE
ATLAS
CMS
LHCb

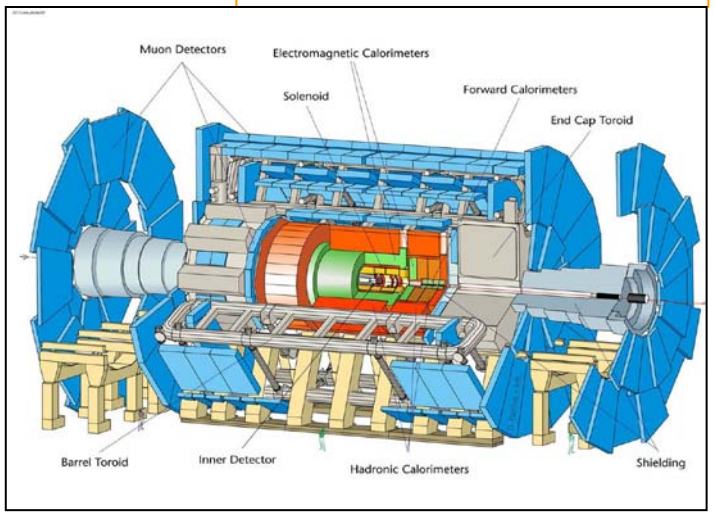

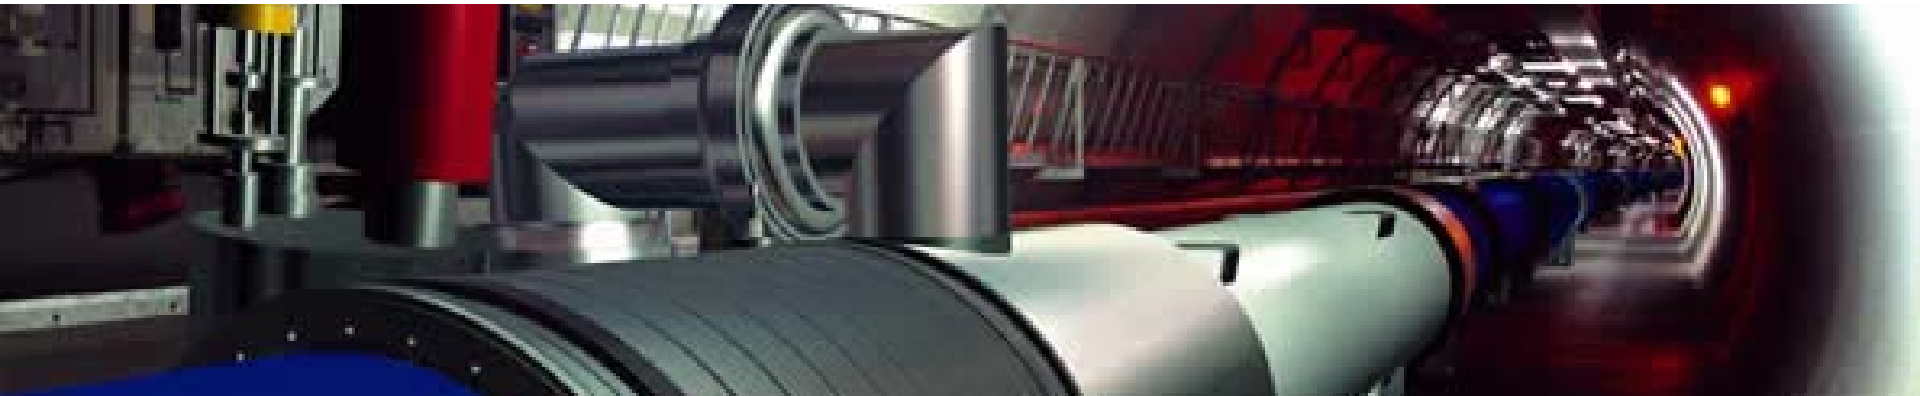

World-wide LHC Computing Grid

- Largest Grid service in the world !

- **Almost 160 sites in 34 countries**
- **More than 200’000 IA processor cores (w/Linux)**
- **20% at CERN**

Data Handling and Computation for Physics Analysis

Online trigger and filtering

High Level Trigger

Selection & reconstruction

Reconstruction

Event Summary Data

10%

processed data

100%

RAW data

Event reprocessing

Batch physics analysis

Analysis (ROOT)

1%

analysis objects
(extracted by physics topic)

Event simulation

Simulation (GEANT4)

Interactive physics analysis

Summary

- At Intel, Moore's Law is alive and well
- Sandy Bridge & AVX drives Xeon family on a new FP trajectory

Broad new Intel Supercomputing investments:

- New: Exascale lab with FZ Jülich, Partec, Intel
- New: Intel® Many Integrated Core (MIC) architecture
- New: Heterogeneous IA HPC tools to simplify road to Exascale
- New: Knights family of co-processors
  - Knights Ferry software development platform
  - Knights Corner product targeting 22nm and >50 Intel Architecture cores

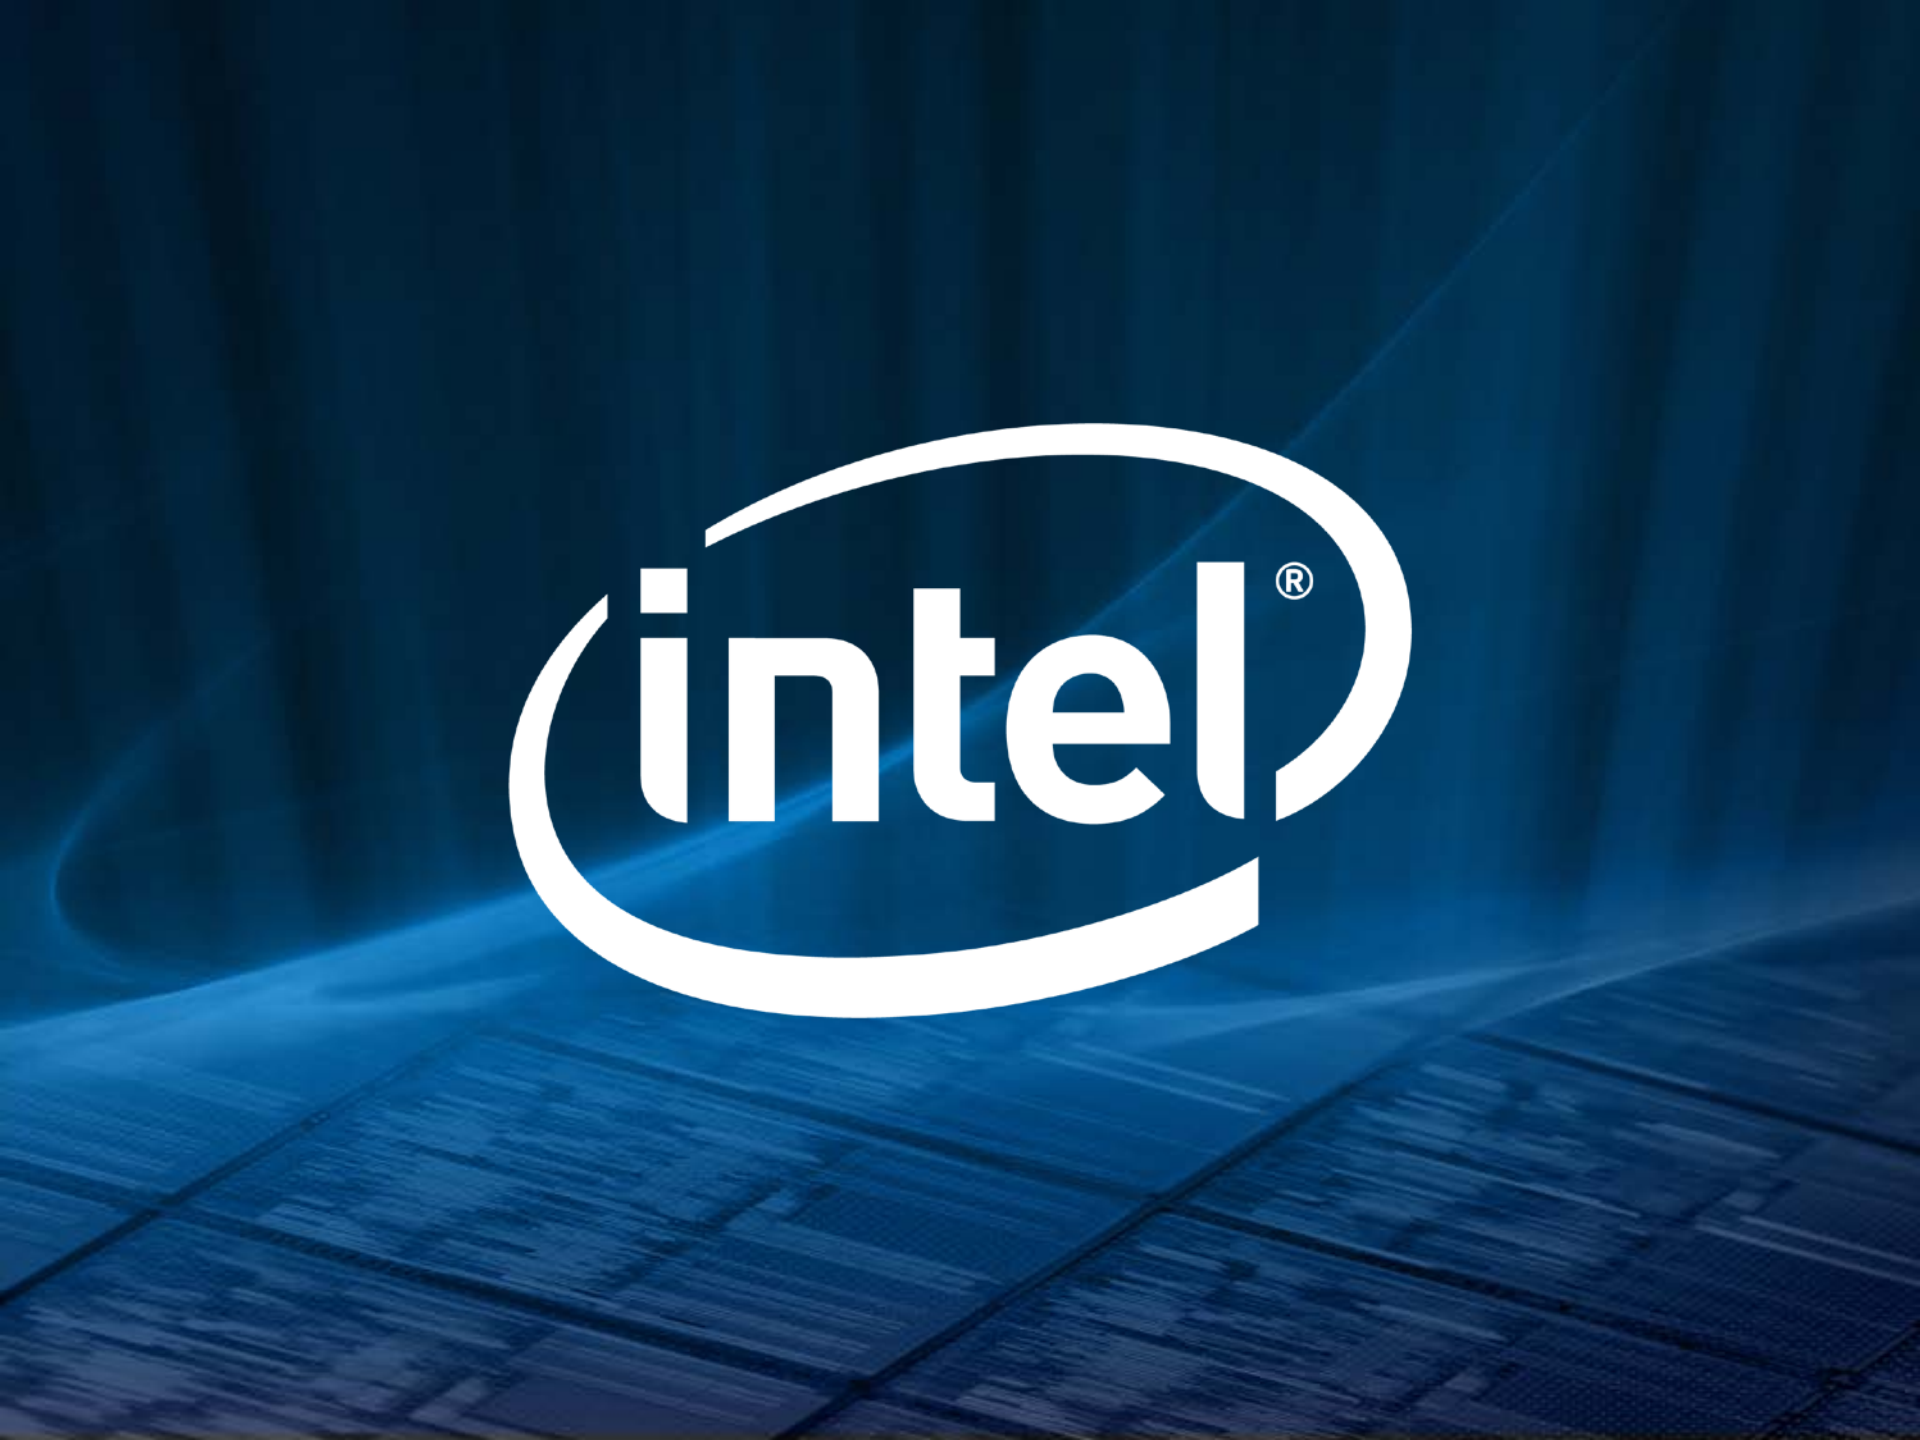
intel®