HPC Technology Trends

High Performance Embedded Computing Conference

September 18, 2007

David S Scott, Ph.D.
Petascale Product Line Architect
Digital Enterprise Group

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).

Rev. 7/19/06

Real World Problems Driving Petascale & Beyond

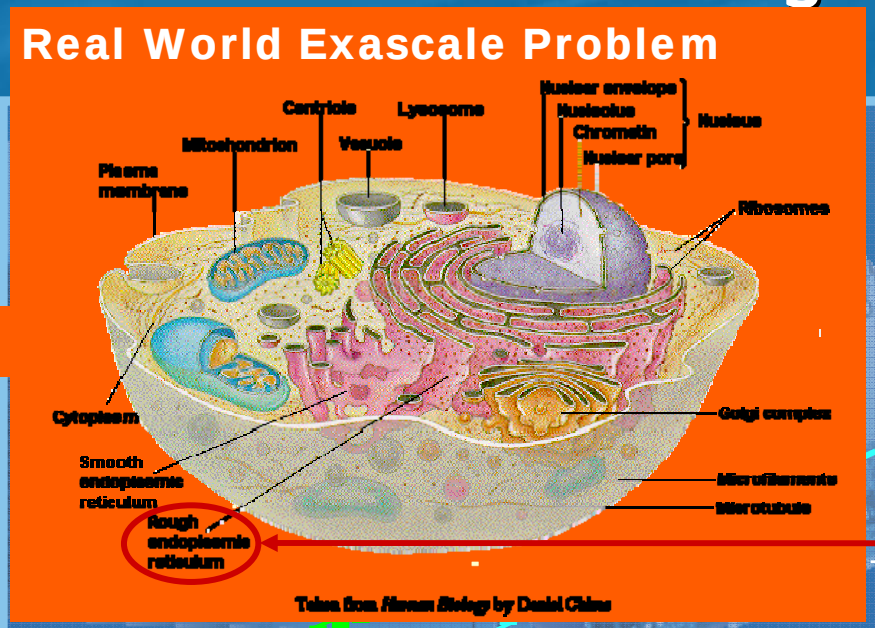

1 ZFlops
100 EFlops
10 EFlops
1 EFlops
100 PFlops
10 PFlops
1 PFlops
100 TFlops
10 TFlops
1 TFlops
100 GFlops
10 GFlops
1 GFlops
100 MFlops

1993 1999 2005 2011 2017 2023 2029

SUM
Of Top500

#1

What we can model today with <100TF

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

intel

Silicon Future

90nm
2003

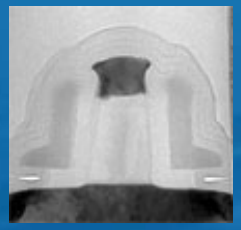

65nm
2005

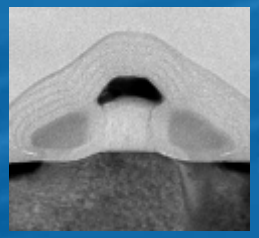

45nm
2007

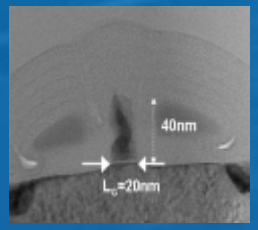

2005 – 2012

32nm
2009

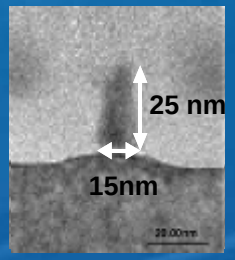

22nm
2011

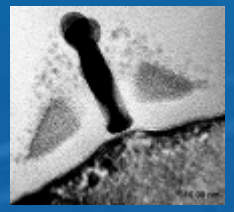

16nm
2013

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

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

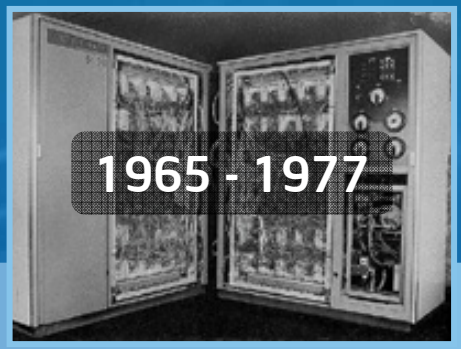

Intel ENDEAVOR
464 Intel® Xeon® Processors 5100 series, 6.85
Teraflop MP Linpack, #68 on top500

~2008 Beyond

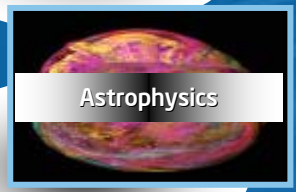

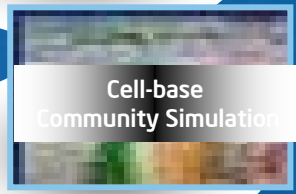

PetaScale Platforms

Yesterday's Supercomputing is Today's Personal Computing

Intel Design & Process Cadence

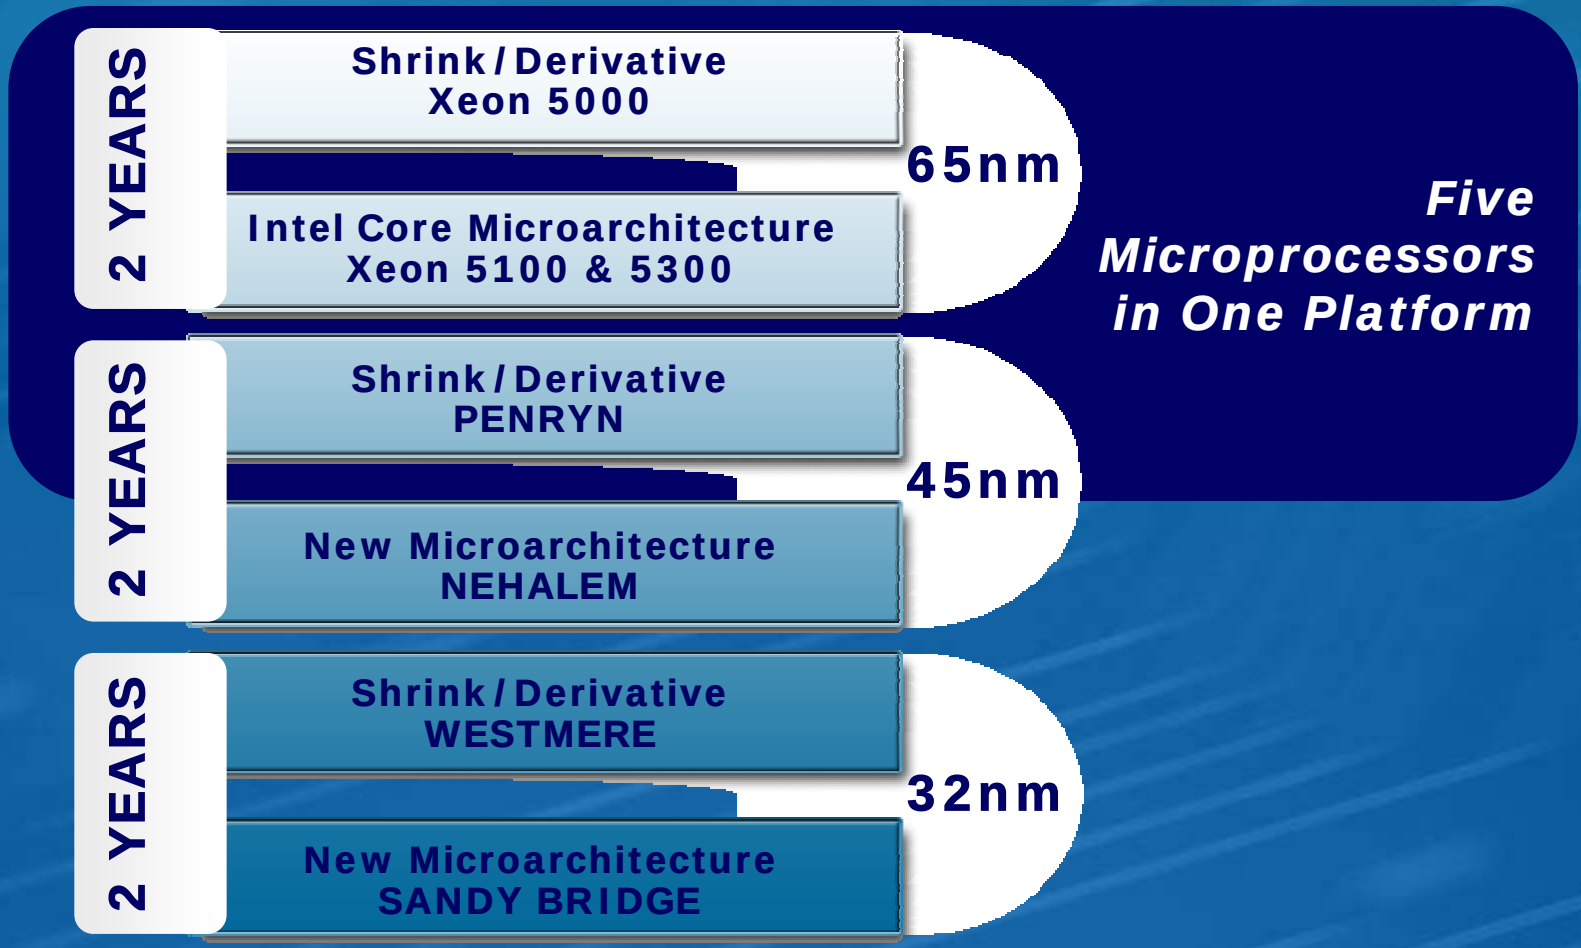

All dates, product descriptions, availability and plans are forecasts and subject to change without notice.

Multi and Many Core

- Multi-core is the current progression of multiple cores on a processor die
- Many core is a discontinuity of putting many simpler cores on a die.
- Jim Held will be talking about Intel's research in this area tomorrow.

Increasing I/O Signaling Rate to Fill the Gap

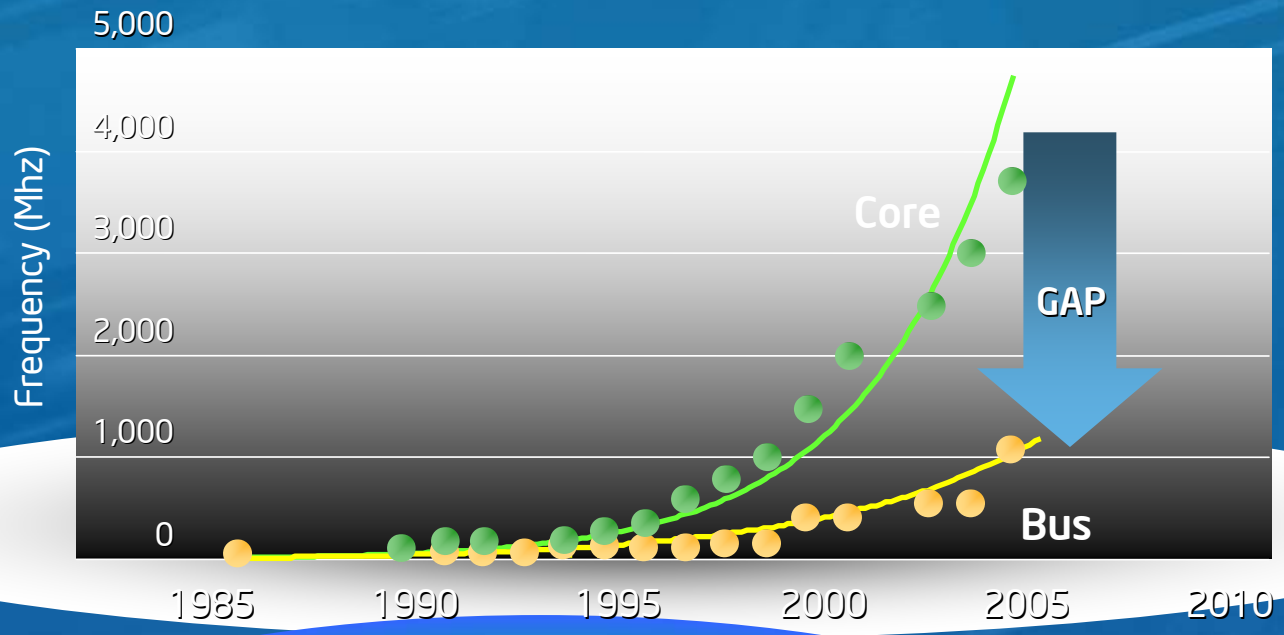

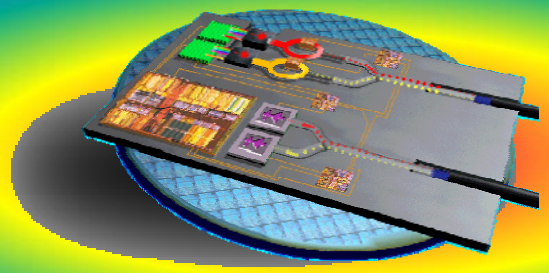

Silicon Photonics

Source: Intel

Increasing Memory Bandwidth
to Keep Pace

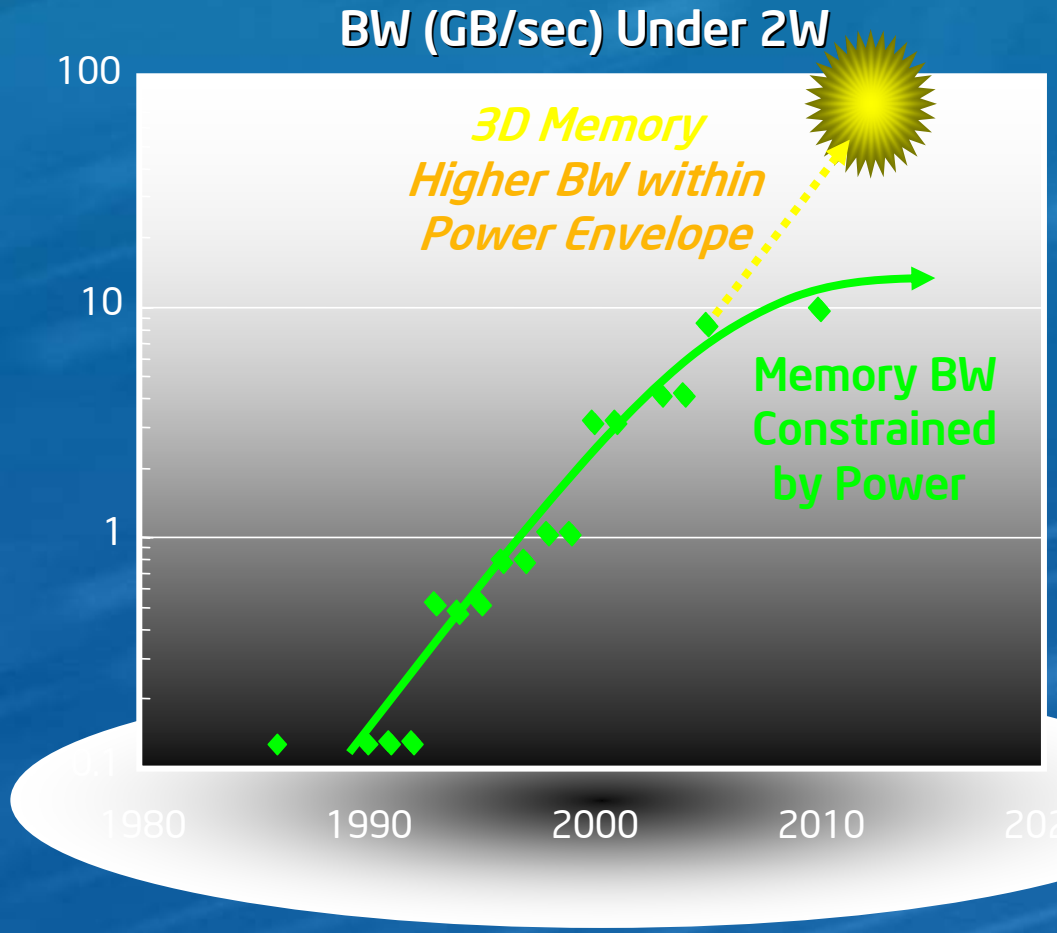

3D Memory Stacking

Power and IO Signals Go Through DRAM to CPU

Thin DRAM Die

Through DRAM Vias

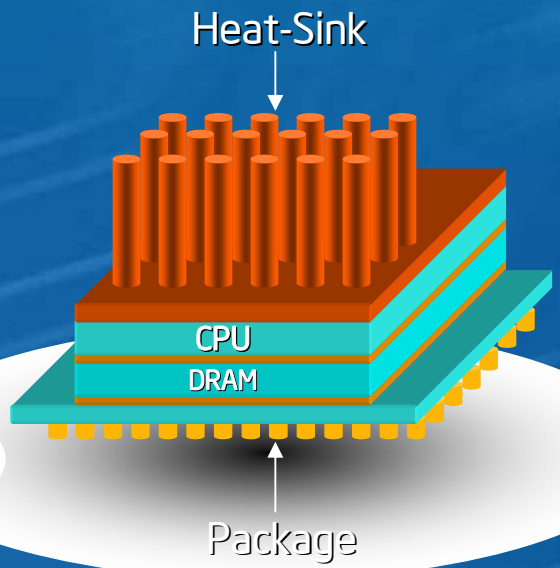

Source: Intel

Power and Cooling Cost Today

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

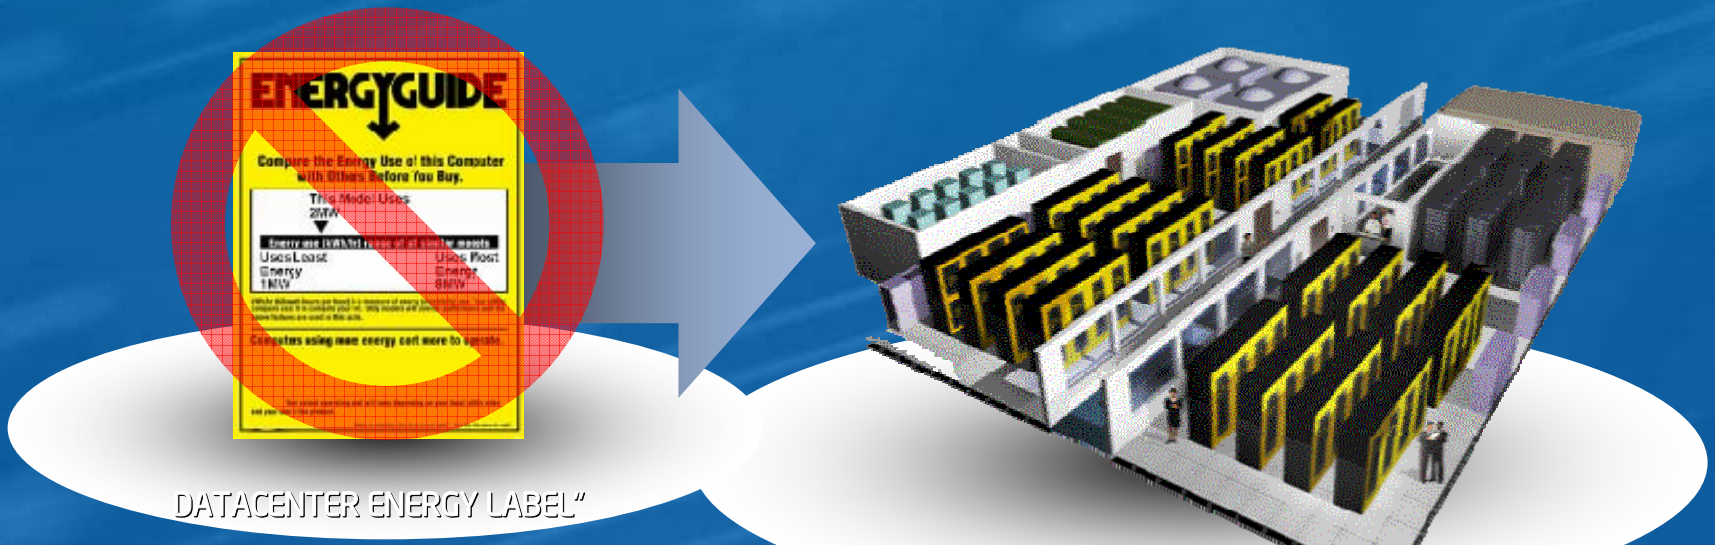

DATACENTER ENERGY LABEL"

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

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

Power Management: From Transistors to Facilities

DELIVERY INEFFICIENCY

CUMULATIVE DATACENTER POWER DELIVERY EFFICIENCY

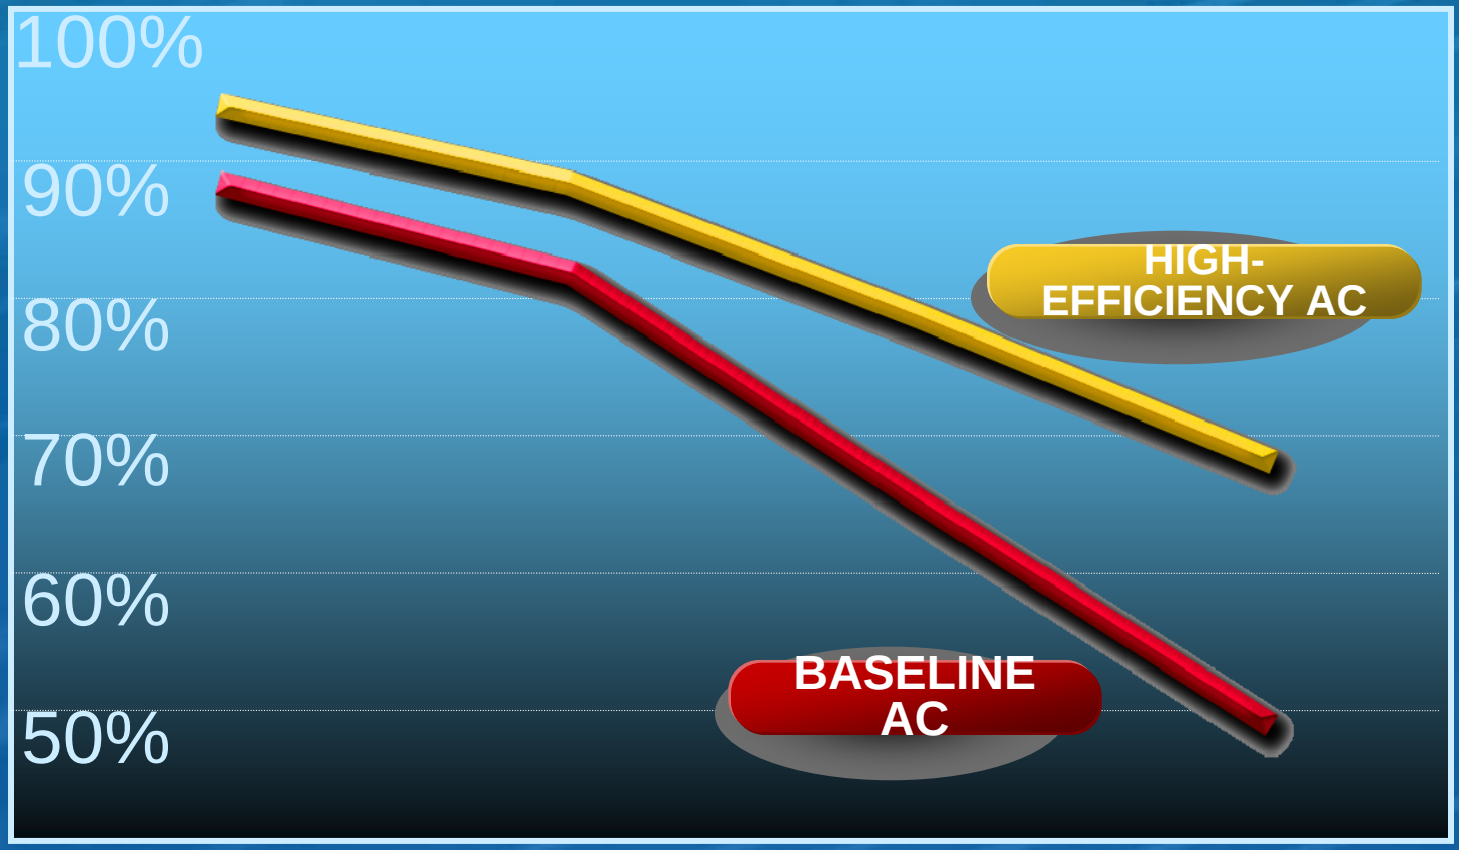

Source: Intel

CONVERSION OVERKILL

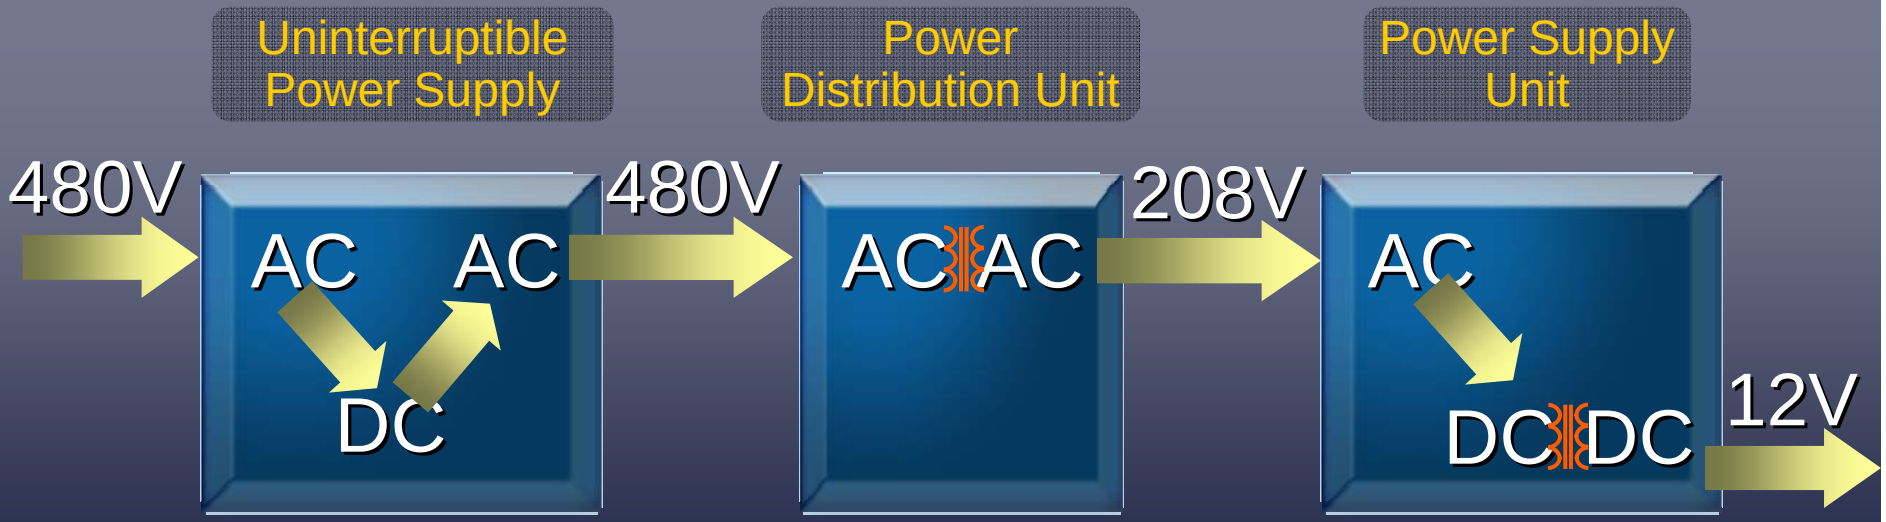

CONVENTIONAL DATACENTER
AC POWER DISTRIBUTION

SIMPLIFIED **DISTRIBUTION**

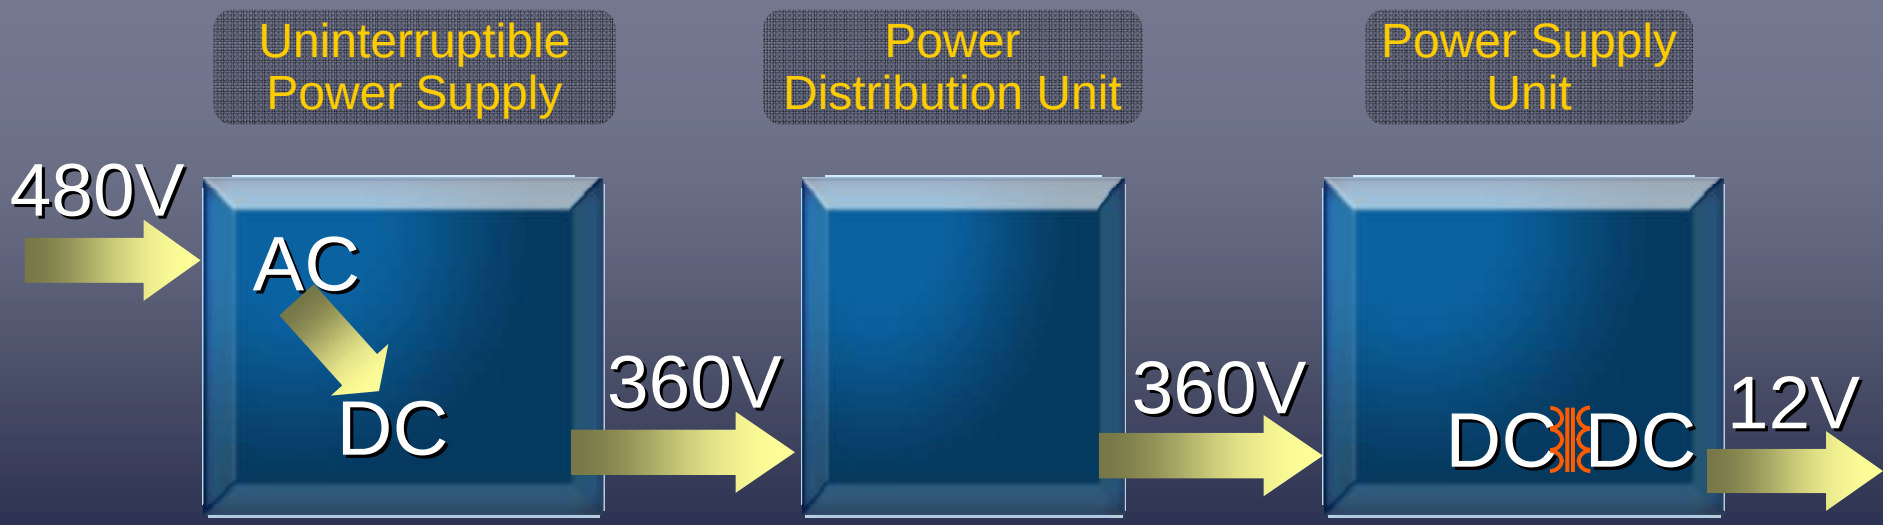

HIGH VOLTAGE DC POWER DISTRIBUTION

HIGH VOLTAGE DC
DATACENTER PROTOTYPE

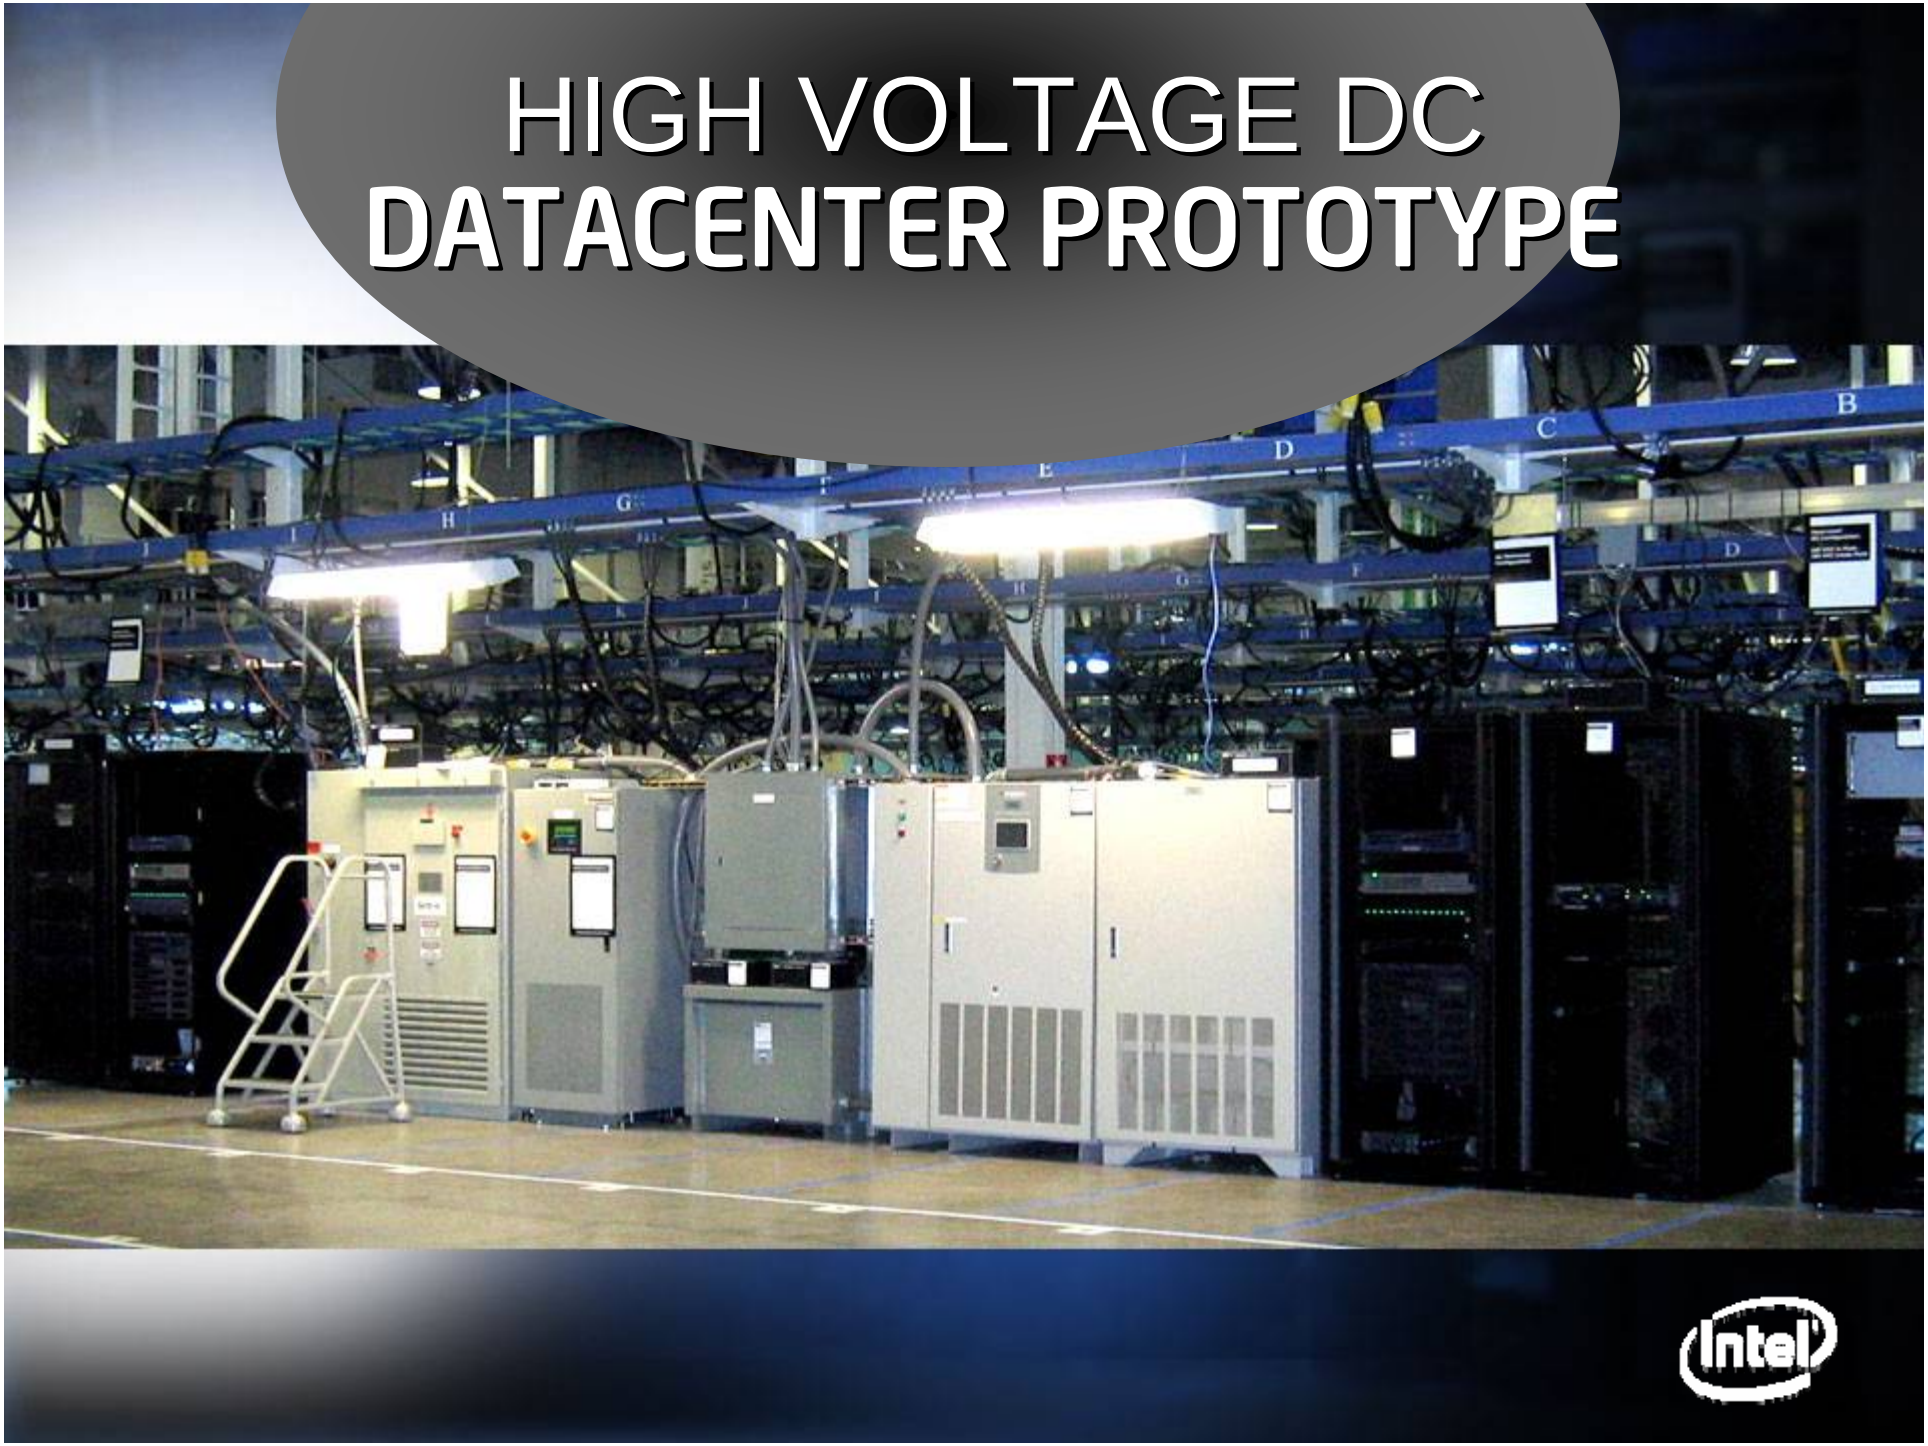

EFFICIENCY **REALIZED**

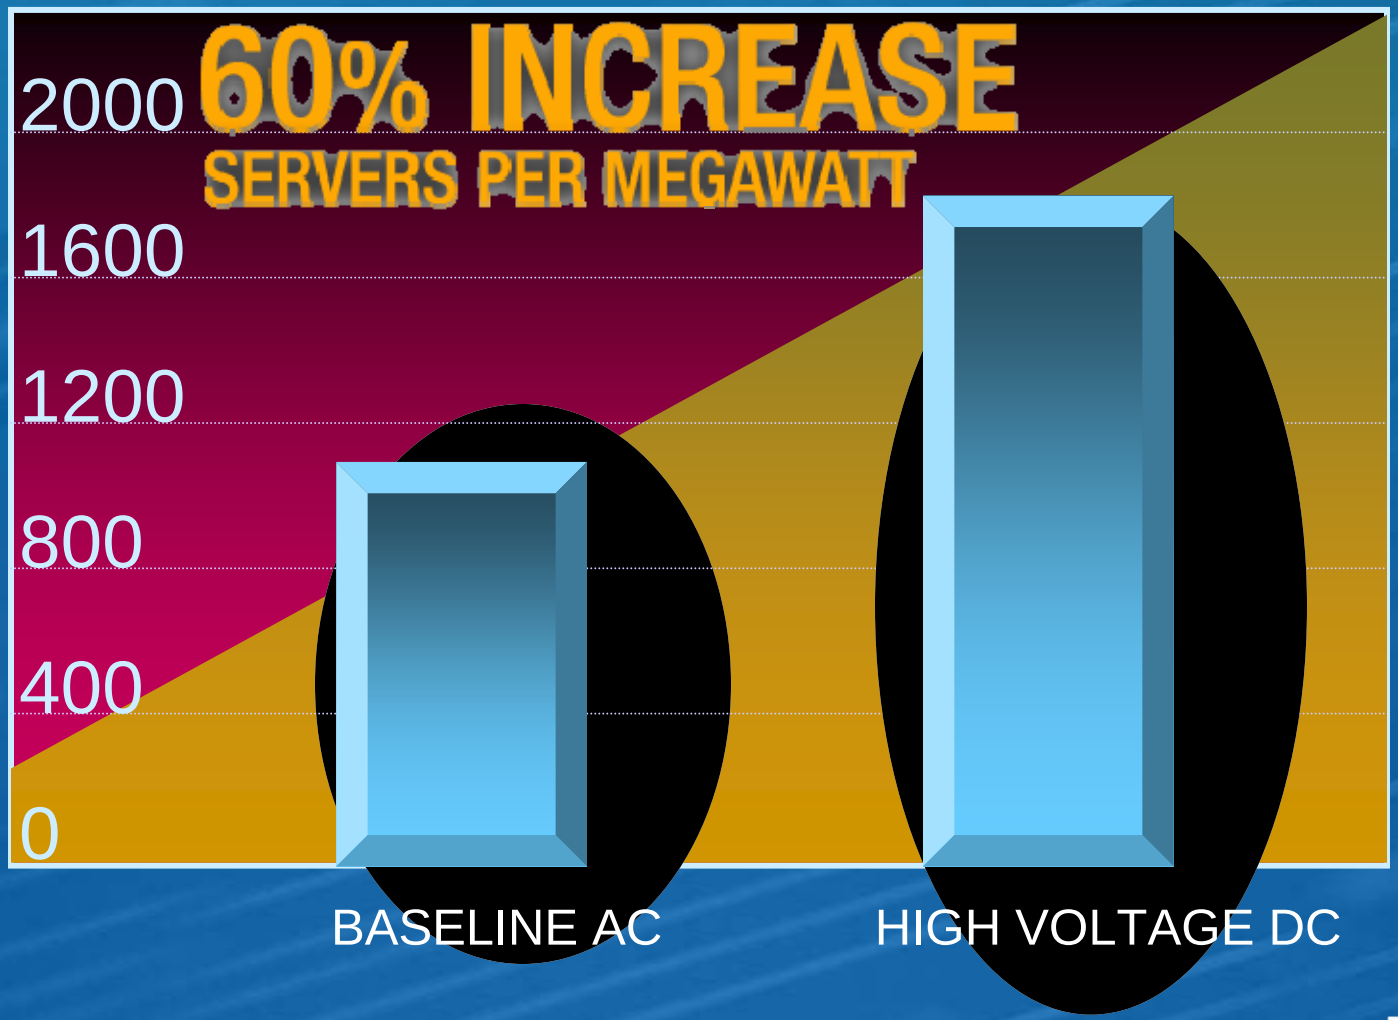

Source: Intel

Reliability Challenge
Billions of Transistors

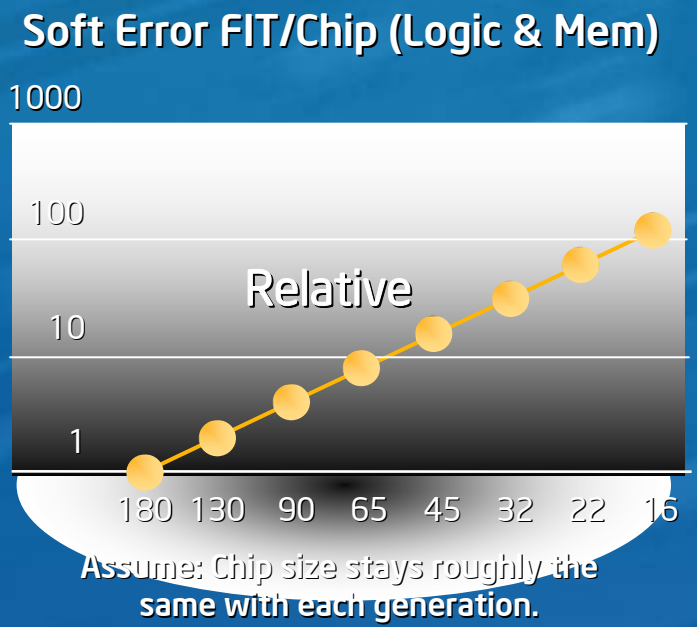

- ***FIT/bit (mem cell)***: expected to be roughly constant
- ***Moore's law***: increasing the bit count exponentially: 2x every 2 years

An exponential growth in FIT/chip

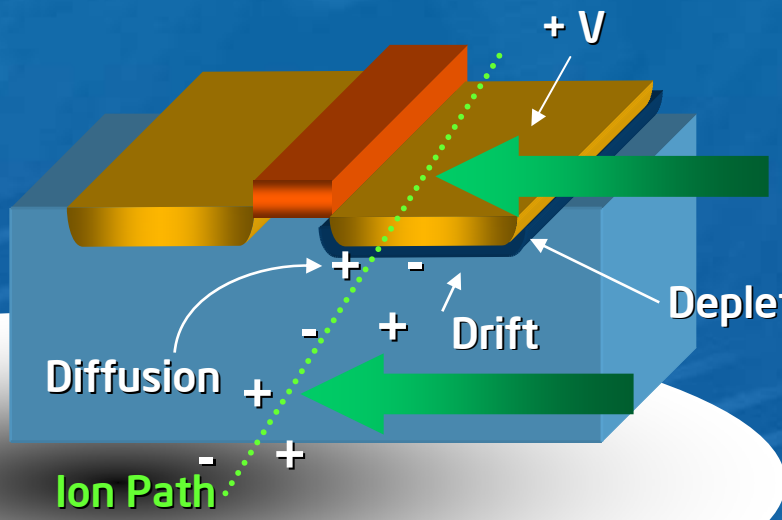

Soft Errors or Single Event Upsets (SEU) are caused by charge collection following an energetic particle strike.

Soft Error: One of Many Challenges

Source: Intel

Innovation with Acceleration

Intel® Architecture with Multi Core

- **General purpose Scalability**
- **Economies of Scale**

Intel® Architecture with Accelerators

- **Special Purpose Performance**
  - Geneseo PCIe extensions
  - QuickAssist Software & Tools

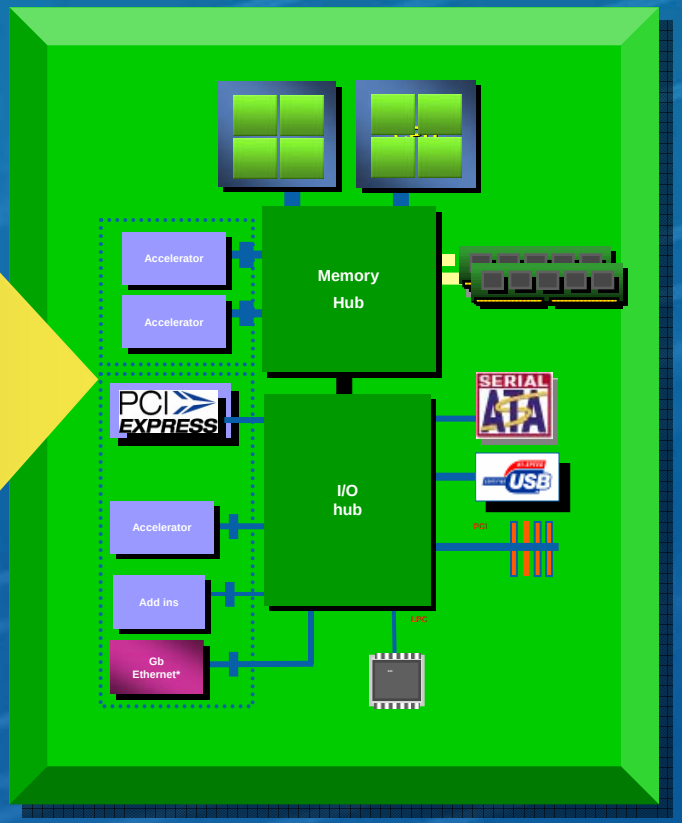

Energy Efficient Performance with Multi-threaded Cores & Accelerators

System bottleneck in supporting Accelerators

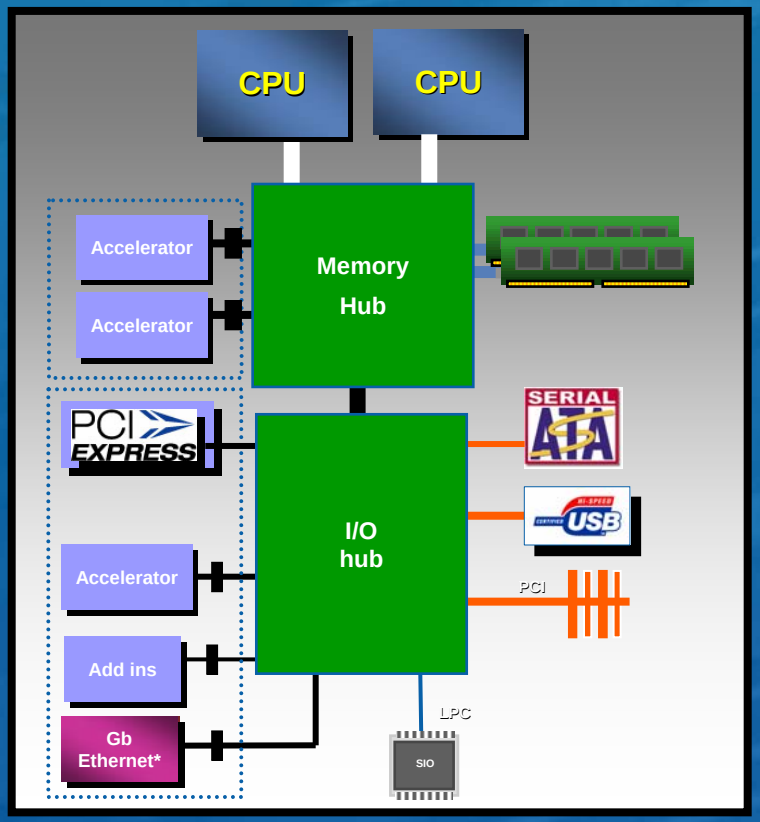

- **Interface Performance**
  - Reduce hardware overhead
  - Lower latency for small packets
  - Higher bandwidth for large packets
- **Programming model and tools**
  - Ease of programming
  - Reduce software overhead; Commands and data movement
  - Status and synchronization
  - Configuration and error handling
  - Virtualization and power management
- **Scheduling & Memory mgmt.**
  - Memory Buffer allocation policy
  - Linear addressing

Software & Platform latencies are 100X the physical IO latency for most accelerators.


intel
Leap ahead™
Questions?