



Taking the New Road

John Masters
Marketing Director, EMEA

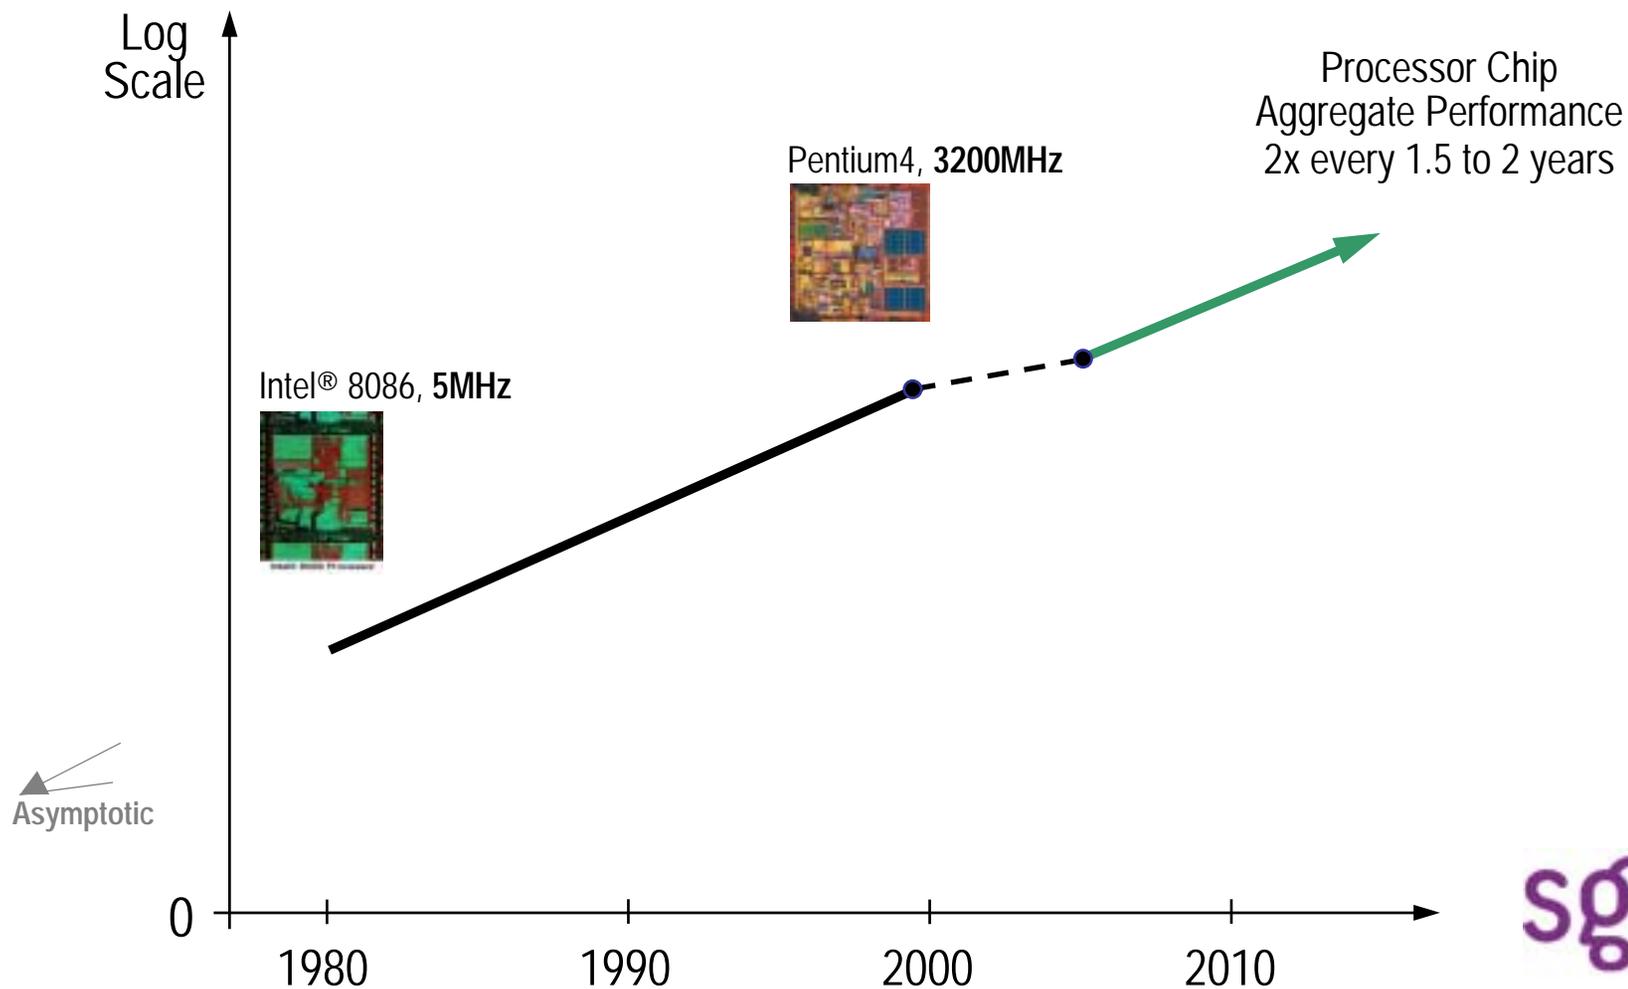
HPC Issues : 2006 feedback

- From Reliance on Clock-Speed to Multicore
- Large-Cluster Issues : Power, Reliability, Weight
- Global shared memory but at Lower Cost
- Vendor viability

HPC Issues : 2006 feedback

- From Reliance on Clock-Speed to Multicore
- Large-Cluster Issues : Power, Reliability, Weight
- Global shared memory but at Lower Cost
- Vendor viability

Change in Processor Progress : Clock-Speed to Multicore



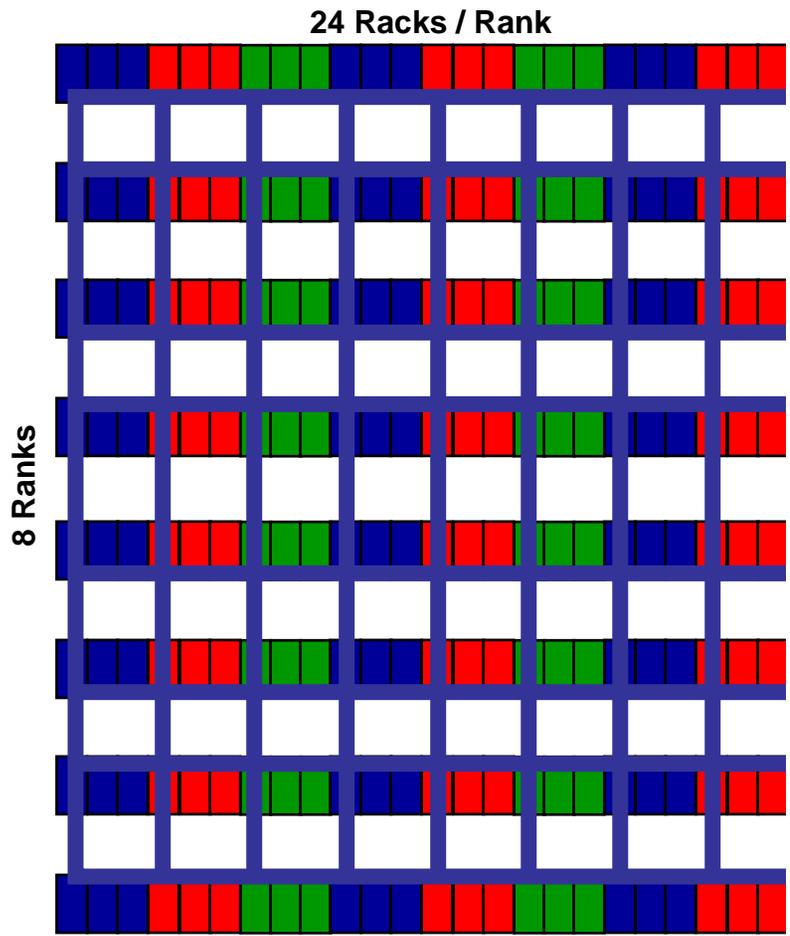
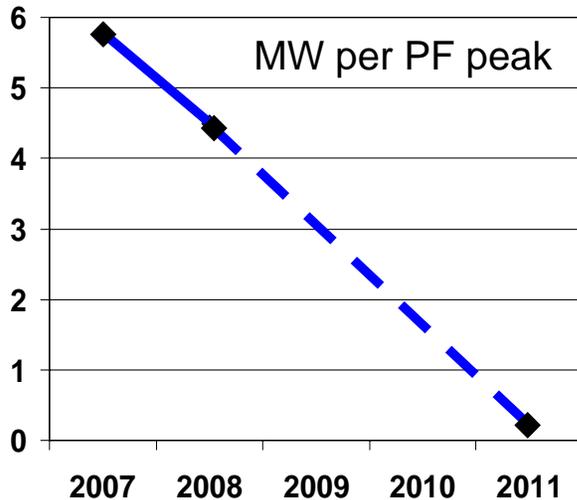
HPC Issues : 2006 feedback

- From Reliance on Clock-Speed to Multicore
- Large-Cluster Issues : Power, Reliability, Weight
- Global shared memory but at Lower Cost
- Vendor viability

Packaging Strategy : Energy Efficiency

Carlsbad '07 : 1.05 PF Peak in 192 Racks

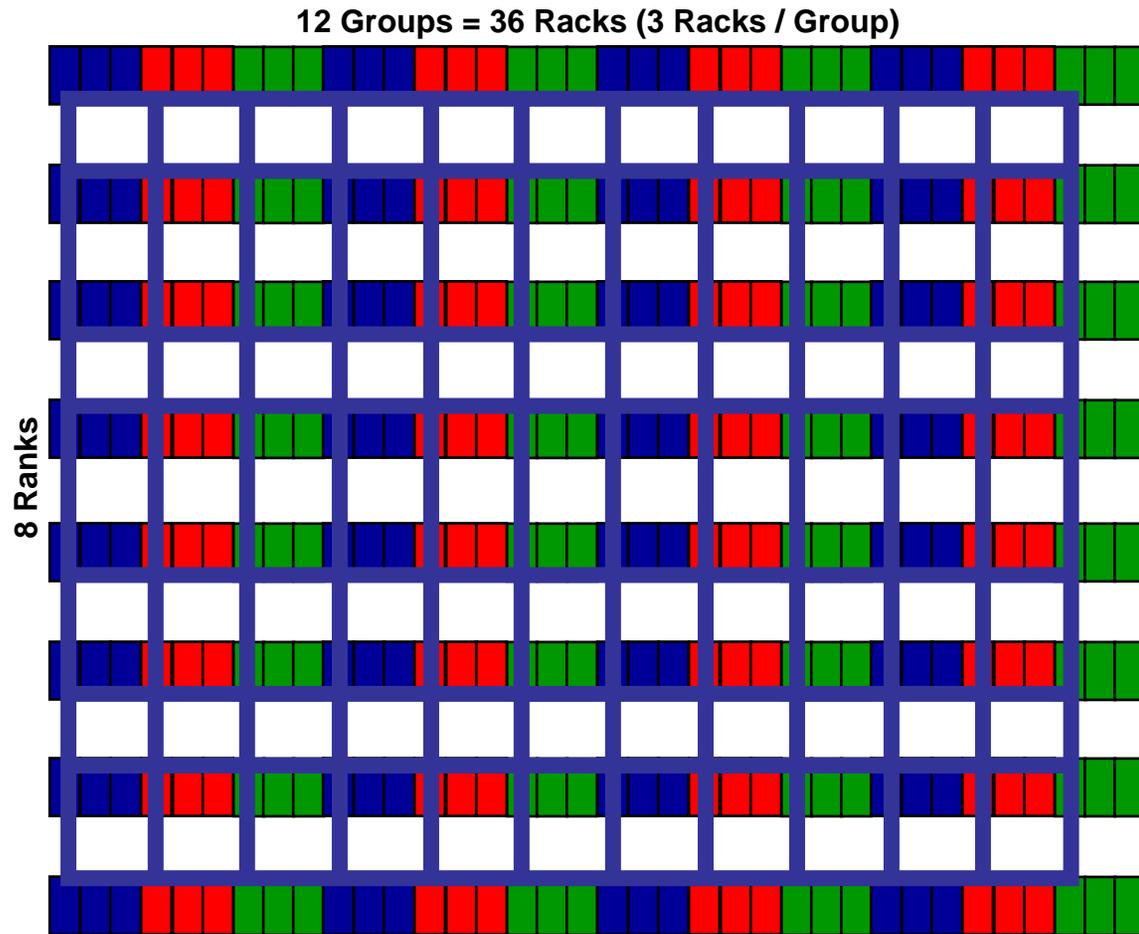
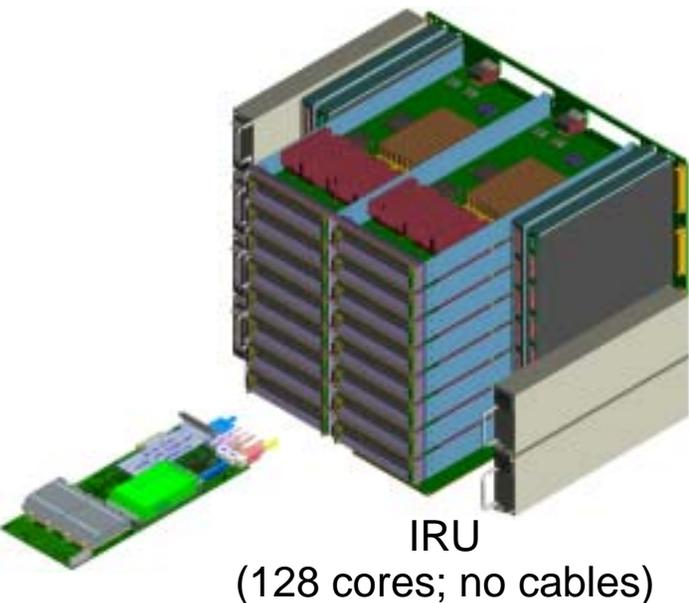
- ~100K cores (2.66 GHz)
- ~6K 24-port IFB_{DDR} switches
- ~55' x 60' footprint
- 6 MW (80% derate)
- **5.8 MW, 1PF peak, '07**
- **4.5 MW, 1PF peak, '08**



Packaging Strategy : Density & Reliability

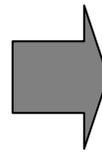
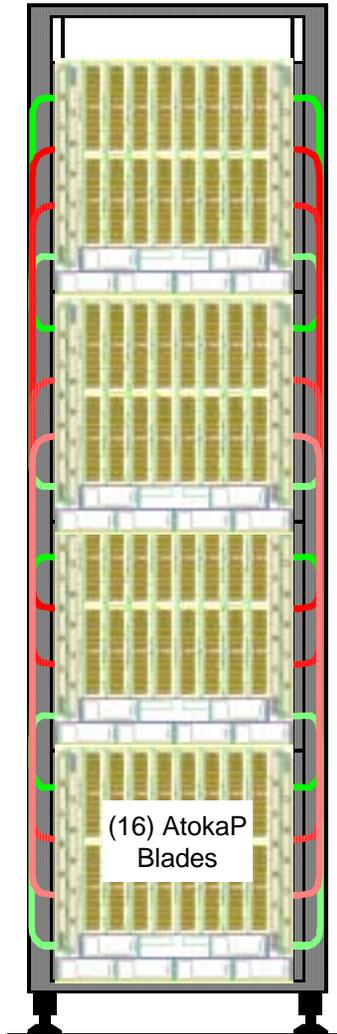
Carlsbad '07 : Cable-Free IRU (much reduced cabling)

- 1.5 PF peak
- Normally = **55,296 cables**
- Via cable-free IRU
= **18,432 cables**



Carlsbad '07 : Density => Weight

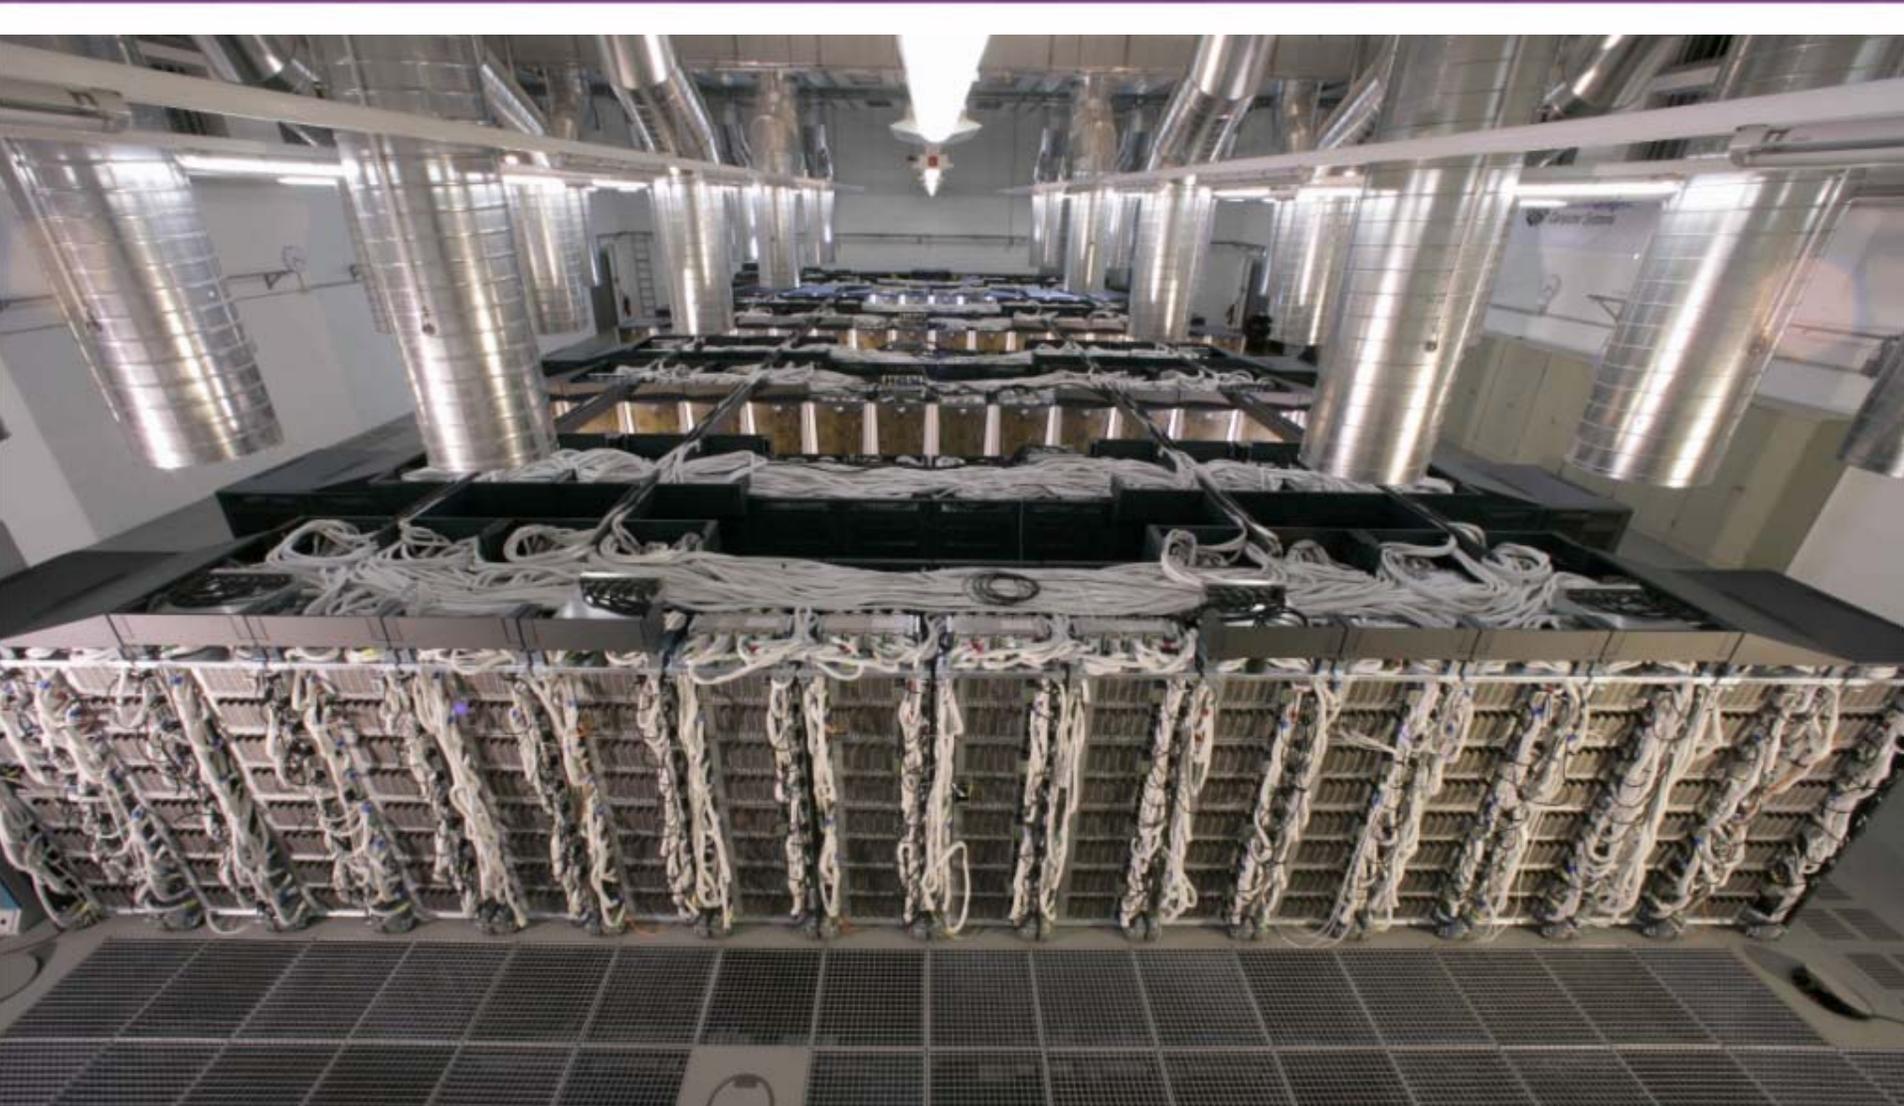
- 2,050 lb
- 19" Rack
- **275 lb / ft²**
(250 lb / ft² limit)



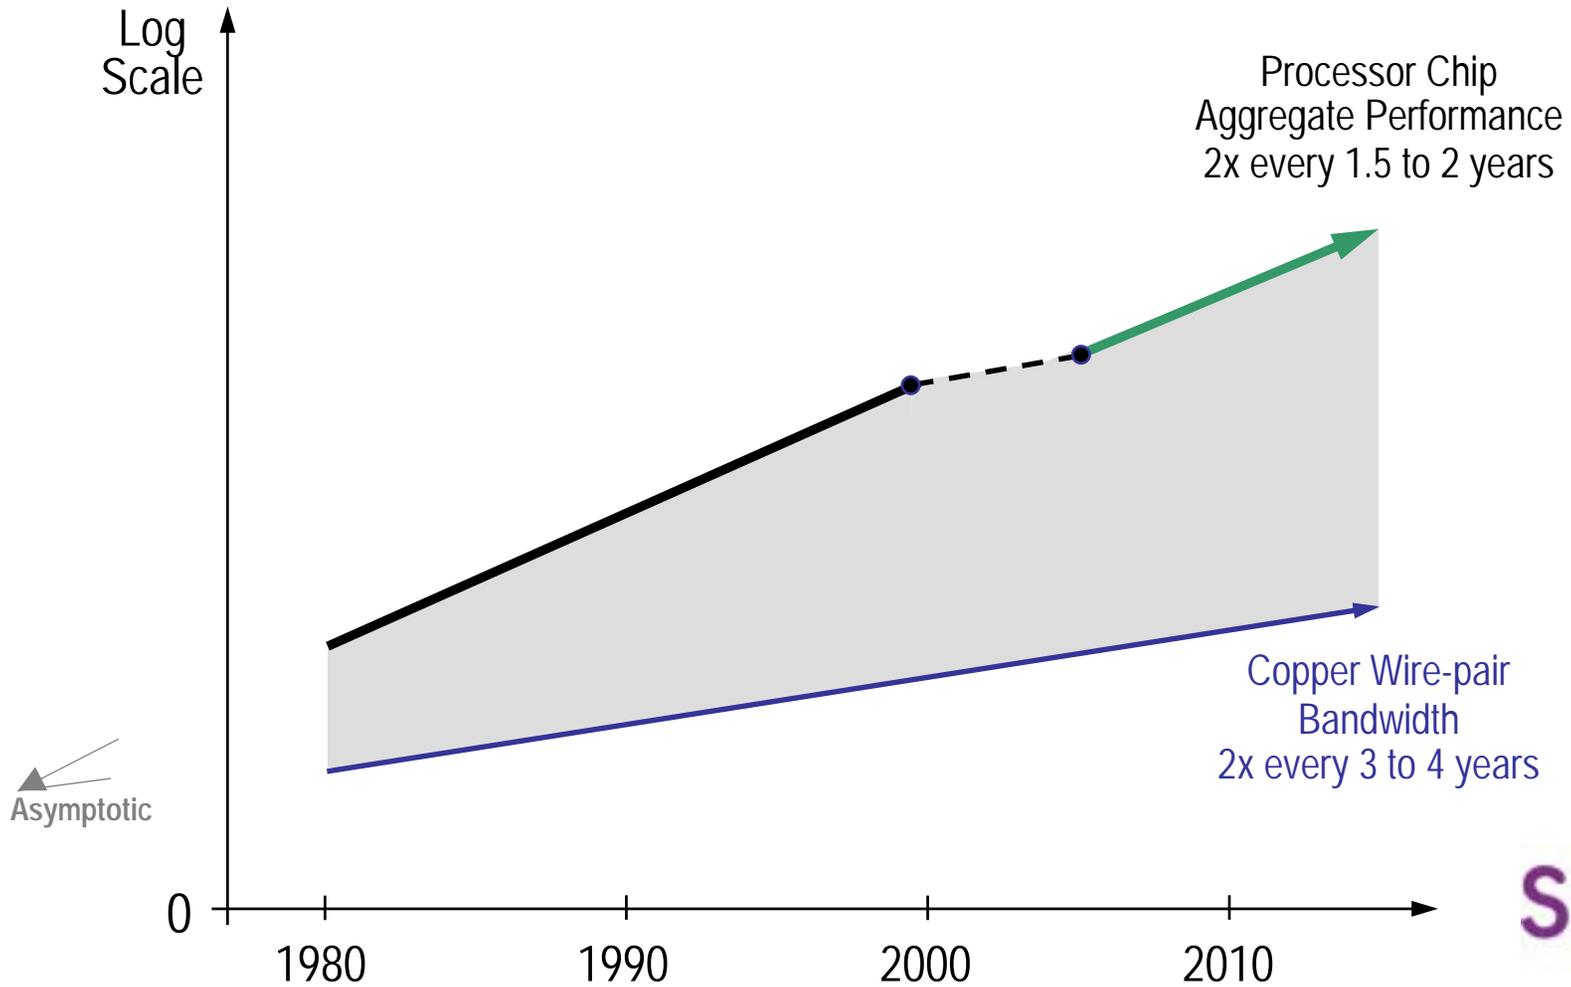
- 2,050 lb
- 24" Rack
- **246 lb / ft²**



Copper : Weight becoming significant



Copper : Weight becoming significant. Why?



HPC Issues : 2006 feedback

- From Reliance on Clock-Speed to Multicore
- Large-Cluster Issues : Power, Reliability, Weight
- Global shared memory but at Lower Cost
- Vendor viability

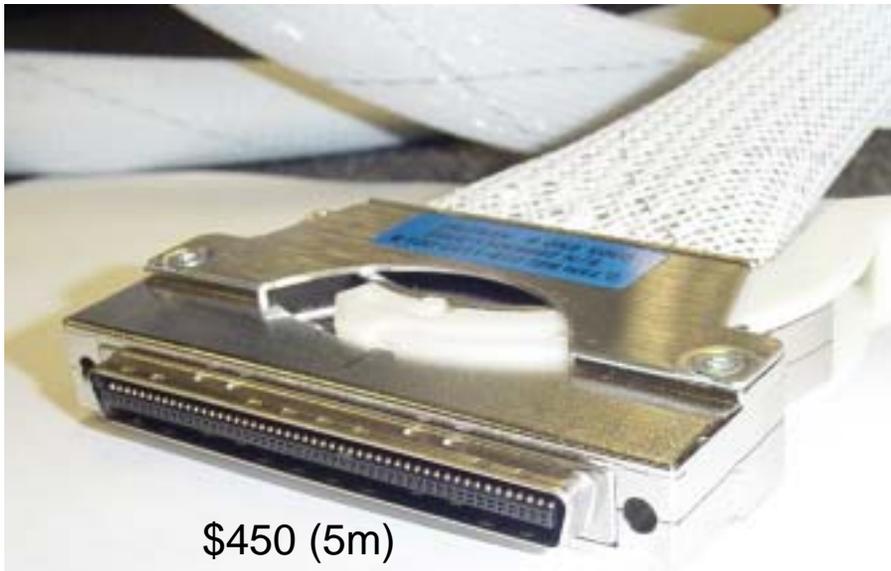
Interconnect Strategy : Reduce Cost & Increase Capability

- **NUMAlink4 (Today)**

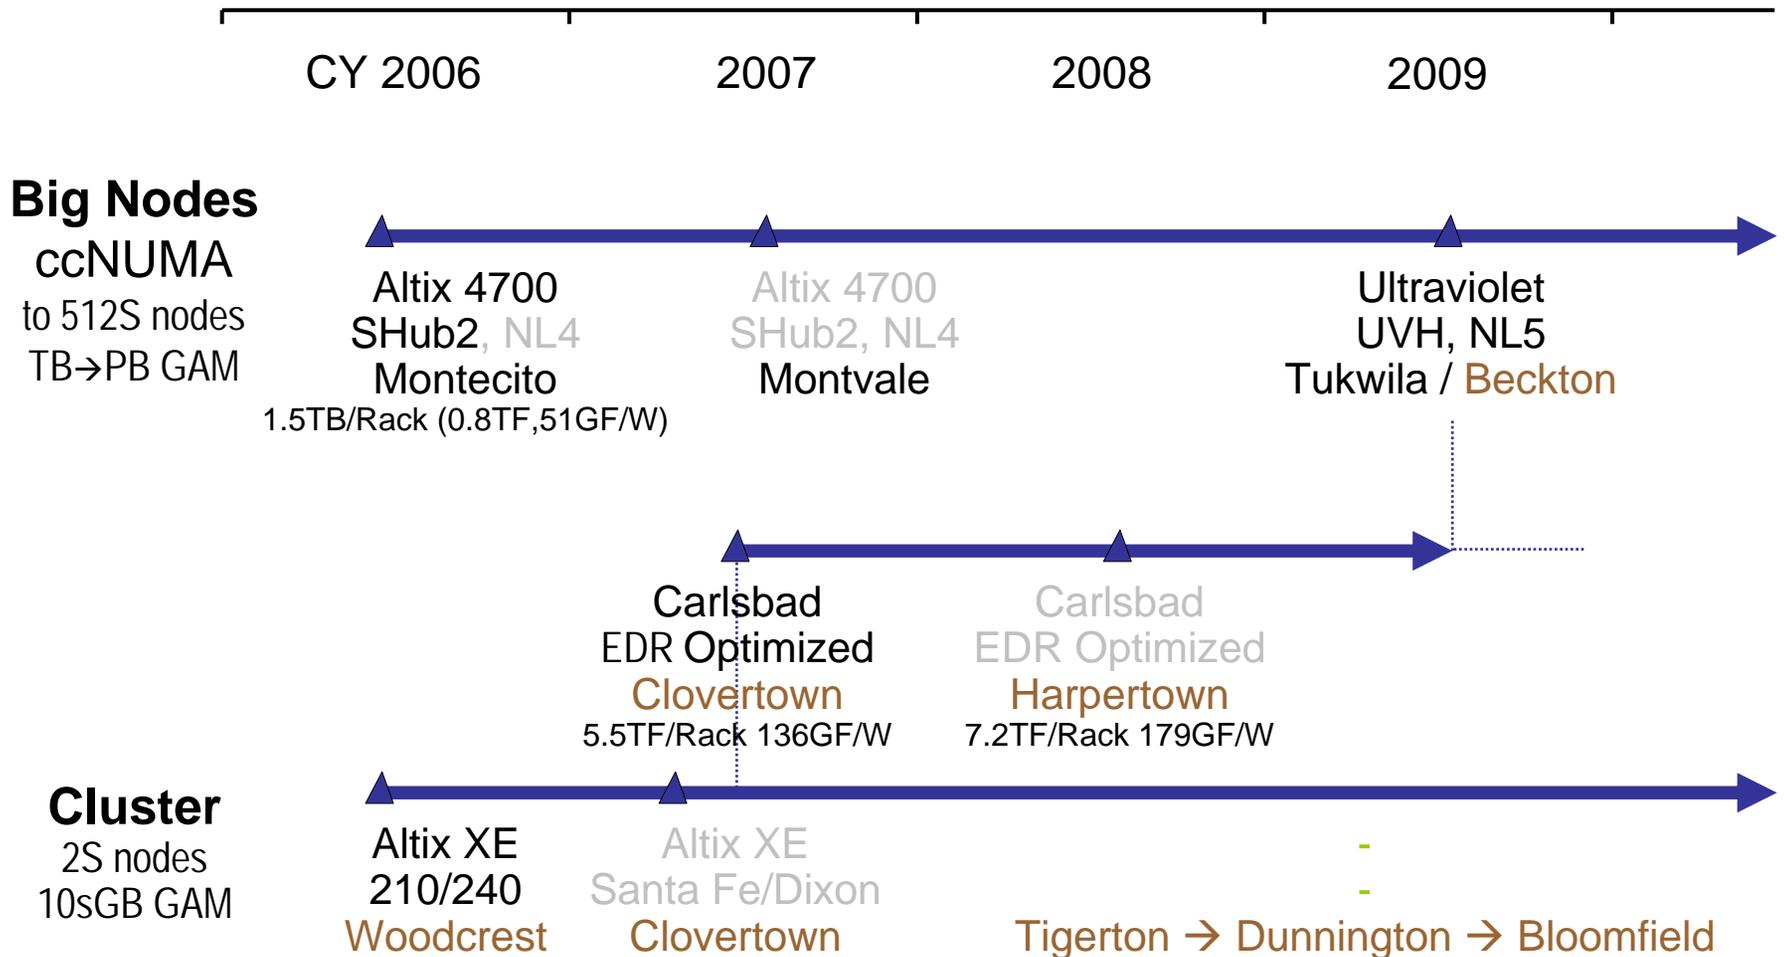
- Custom copper cable
- Custom signalling
- Custom protocol

- **NUMAlink5: hw extension of IFB ('09)**

- COTS Infiniband_{12x} copper cable
- COTS serdes
- Custom protocol (higher capability)

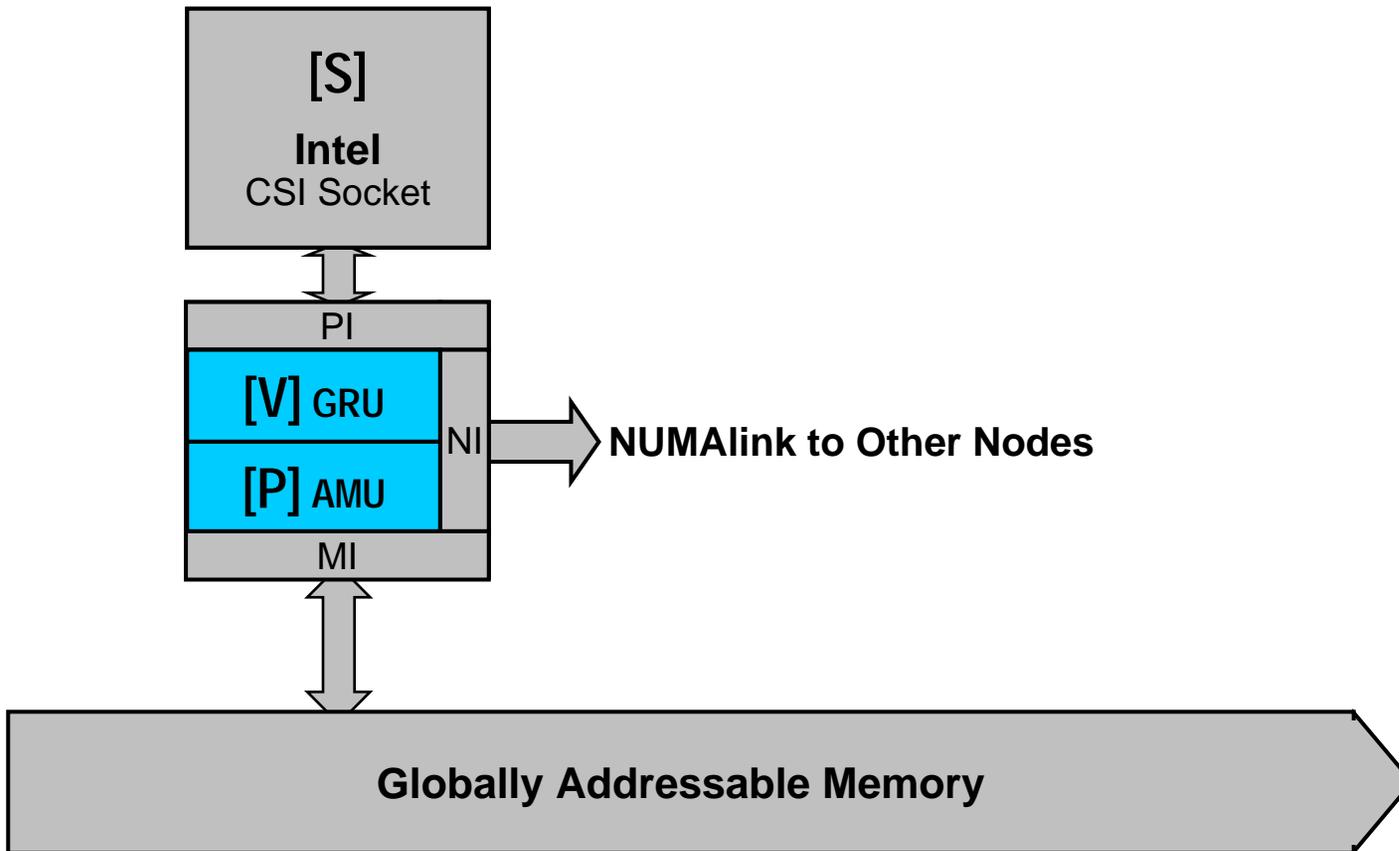


SGI Server Roadmap



Interconnect Strategy : Reduce Cost & Increase Capability

Shub2 (Today) to UV-Hub (2009)



HPC Issues : 2006 feedback

- From Reliance on Clock-Speed to Multicore
- Large-Cluster Issues : Power, Reliability, Weight, IO-Data
- Global shared memory but at Lower Cost
- Vendor viability

The “New SGI” Financially Healthy

- **Financial / Re-Engineering Complete**
 - Emerged from Chapter 11 one month earlier than original plan
 - Recapitalised the Company
 - Eliminated \$250M legacy debt
 - Bolstered our cash position
- **New Leadership Team, New Board of Directors**
- **Costs Reduced \$150M Annually**
- **Stabilised the Business**
- **Most Complete Product Portfolio in Company’s History**
- **Focus on Growth Initiatives**

SGI® Altix® Dual-Core Servers

Altix 450 and Altix 4700

Revolutionary Architecture Leads Next-Wave of Blade-based Platforms

Innovative Blade-to-NUMALink4 Design: Provides unprecedented versatility, density, power efficiency, performance

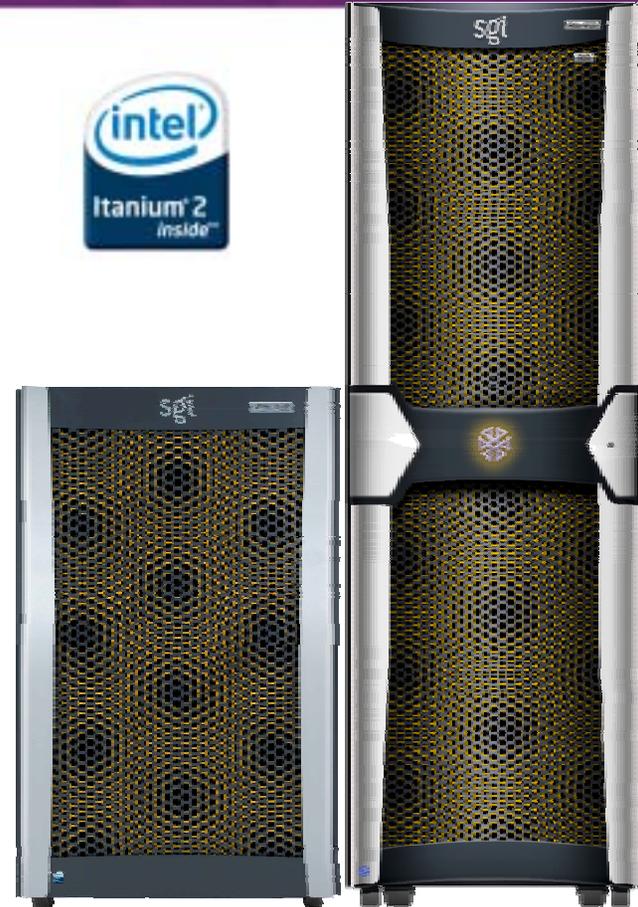
Investment Protection: Processor-only upgrade to future Intel® Itanium® 2 processors or expansion via new blades

Enables Multi-Paradigm Computing: Enhanced integrated computational resources that can be seamlessly accessed as application requires

With the Scalable, Reliable Performance on Linux that Altix is Known for

“With the SGI Altix we’re significantly increasing our capacity for large-scale computations requiring large memory, problems that are even beyond the scope of our existing system, which is hardly small. We’ve run some benchmarks, and jobs that were taking about 50 to 60 hours on the HP supercomputing system actually took about 2 to 6 hours on a similarly configured SGI Altix. I consider the Altix to be a keystone in our infrastructure.”

- Dr. Gordon Springer, Scientific Director, Bioinformatics Consortium, University of Missouri



SGI® Altix® 450

SGI® Altix® 4700

SGI® Altix® XE1200 Clusters

- Competitive price/performance
- Dual Intel® Xeon® Processors
- Linux® OS-based
- Full factory integration for plug 'n' play implementation in target markets
- Technologically advanced architecture (energy efficiency, enhanced software environment)
- Solid roadmap for future differentiation based on extended performance and scalability enhancements

“IDC research studies show that end users frequently find large clusters very complex to manage and they require significant effort to setup and optimize. SGI's approach is to integrate application-optimized clusters in the factory and to support these systems with the same team that has serviced its supercomputers and storage products.”

- Earl Joseph, vice president for IDC's Technical Computing Systems practice



SGI® Altix® XE1300 Clusters

- Dual socket with Dual/Quad core Xeon based processors
- 2 servers in 1U (‘server’ = mobo, disk, 2 x gigE, IB ...)
 - Power per server reduced
 - Half the number of Power Supply Units
- Infiniband integrated onto motherboards
 - Dramatically reduced cost for high performance interconnect***
- 512 cores + 23kW in 1 rack

ROADMAP : XE1300 + Tighter integration = ‘Carlsbad’



HPC Issues : 2006 feedback

- From Reliance on Clock-Speed to Multicore
- Large-Cluster Issues : Power, Reliability, Weight
- Global shared memory but at Lower Cost
- Vendor viability



Questions?