Cycles for Competitiveness: A View of the Future HPC Landscape

October 6, 2010

Stephen R. Wheat, Ph.D.
Sr. Director, HPC WW Business Operations
Intel, Data Center Group

Legal Disclaimer

- Intel may make changes to specifications and product descriptions at any time, without notice.
- 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
- Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of systems available for purchase.
- Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.
- Intel, processors, chipsets, and desktop boards 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.
- Intel Virtualization Technology requires a computer system with a processor, chipset, BIOS, virtual machine monitor (VMM) and applications enabled for virtualization technology. Functionality, performance or other virtualization technology benefits will vary depending on hardware and software configurations. Virtualization technology-enabled BIOS and VMM applications are currently in development.
- 64-bit computing on Intel architecture requires a computer system with a processor, chipset, BIOS, operating system, device drivers and applications enabled for Intel® 64 architecture. Performance will vary depending on your hardware and software configurations. Consult with your system vendor for more information.
- Lead-free: 45nm product is manufactured on a lead-free process. Lead is below 1000 PPM per EU RoHS directive (2002/95/EC, Annex A). Some EU RoHS exemptions for lead may apply to other components used in the product package.
- Halogen-free: Applies only to halogenated flame retardants and PVC in components. Halogens are below 900 PPM bromine and 900 PPM chlorine.
- Intel, Intel Xeon, Intel Core microarchitecture, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- © 2009 Standard Performance Evaluation Corporation (SPEC) logo is reprinted with permission

High Performance Micro-Architecture for HPC

Moore's Law: Alive and Well at Intel

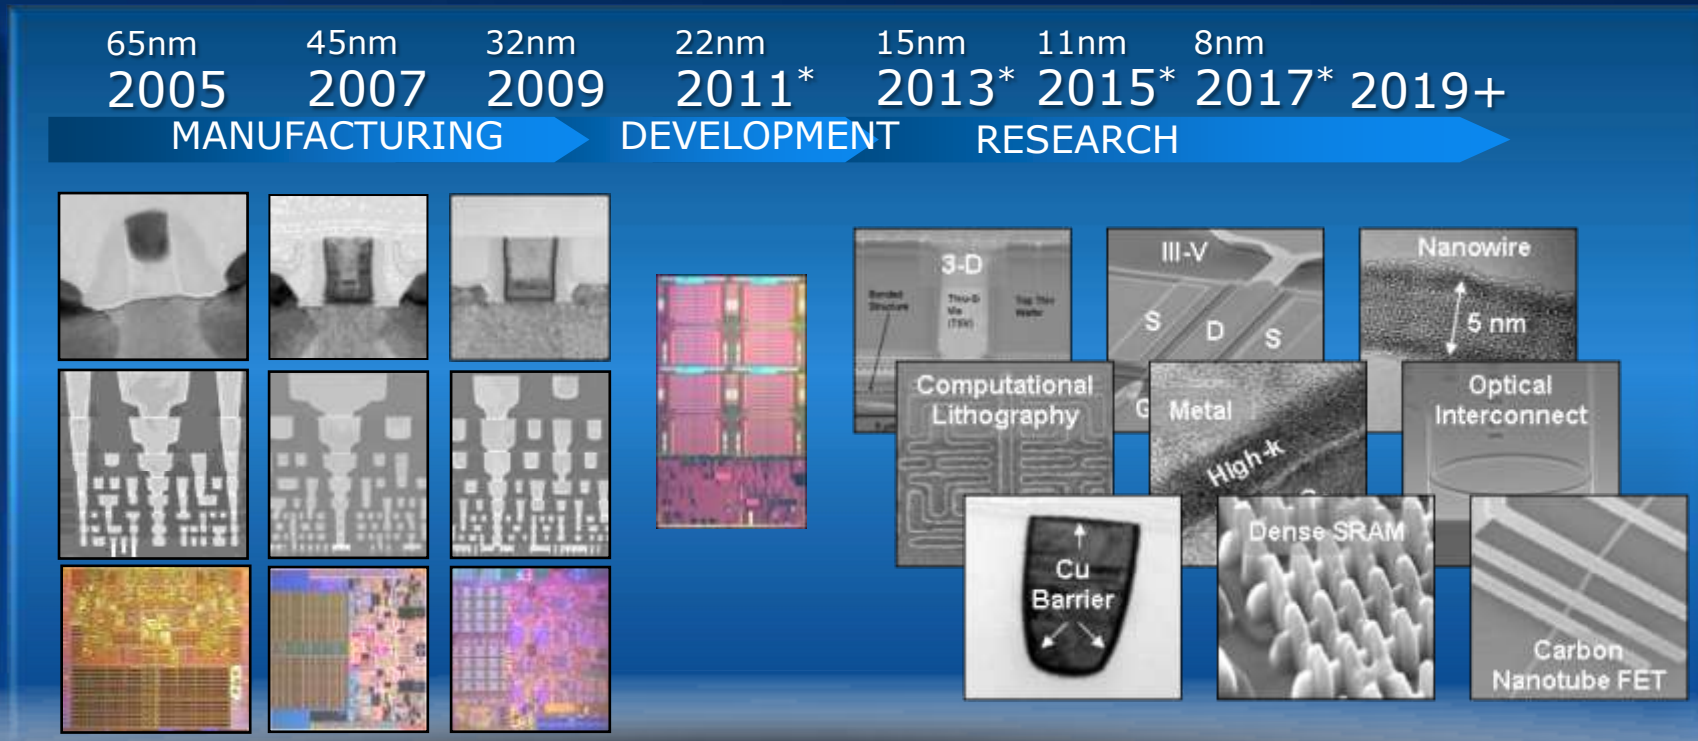

Intel Innovation-Enabled Technology Pipeline is Full

Intel® Xeon® 5600 Energy Efficiency

Building on Xeon® 5500 Leadership Capabilities

Lower Power CPUs

Better performance/Watt
Lower power consumption

- 130W
- 95W
- 80W
- 60W (6C)
- 40W (4C)

Intelligent Power Technology

Integrated Power Gates and Automated Low Power States with **Six Cores**

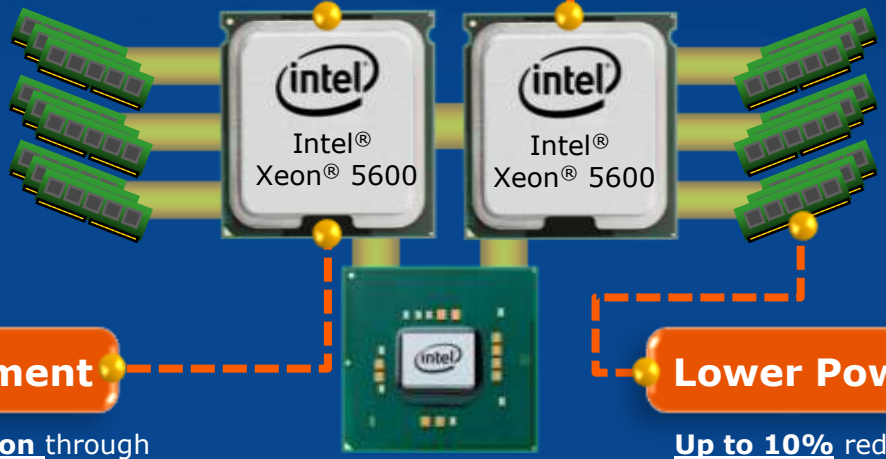

CPU Power Management

Optimized power consumption through more efficient Turbo Boost and memory power management

Lower Power DDR3 Memory

Up to 10% reduction in memory power[1]

Intel® Xeon® 5600 delivers greater platform Energy Efficiency

- 1 Based on voltage reduction from 1.50V to 1.35V, using Power (Watts) = Current x Voltage

Lower power CPU SKU options for Xeon® 5600

Xeon® 5500 → Xeon® 5600

Xeon® 5600 (Westmere-EP) SKUs

Xeon® 5500 SKUs

Instruction Set Innovation Continues in Sandy Bridge CPUs

Intel® Advanced Vector Extensions (Intel® AVX)

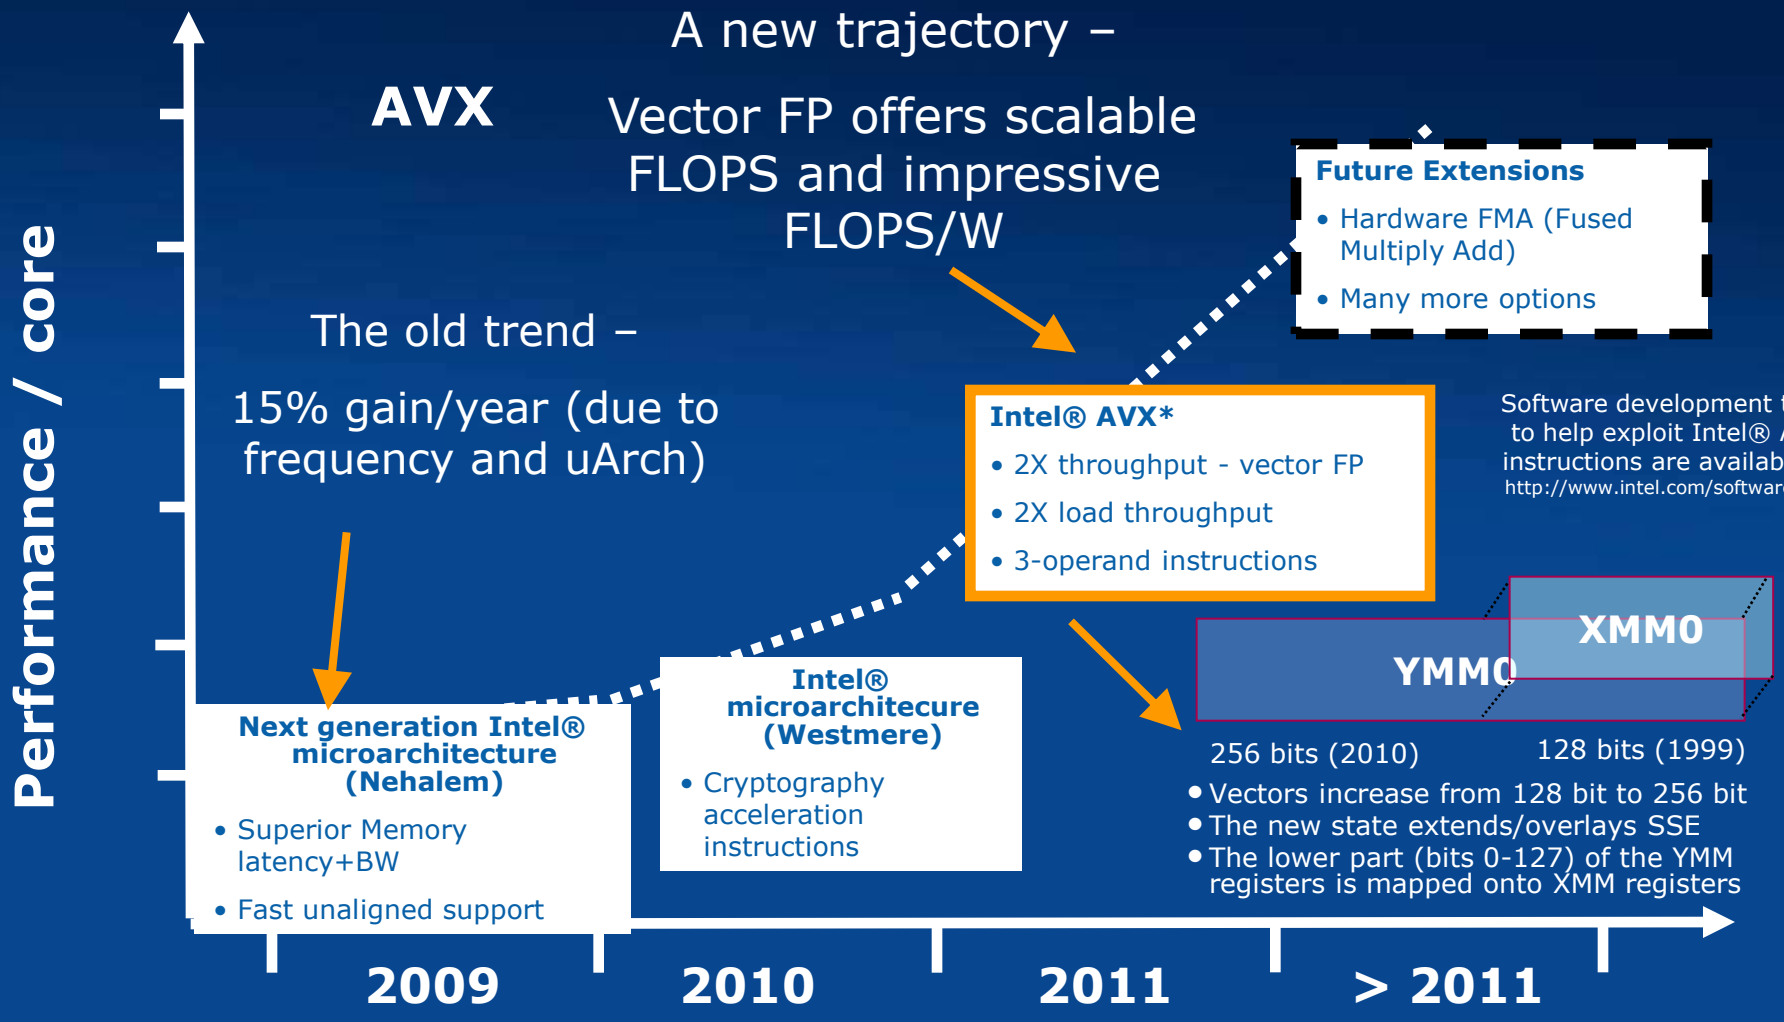

Intel® AVX is a general purpose architecture building upon SSE

Intel® Advanced Vector Extensions Overview

- New instructions to boost FP performance
- Extend existing FP vector instructions to 256-bits
- Full details at: http://www.intel.com/software/avx
- Requires at least a recompile of code

KEY FEATURES	*BENEFITS*
• Wider Vectors – Increased from 128 bit to 256 bit	• Up to 2x peak FLOPs output with good power efficiency
• Enhanced Data Rearrangement – Use the new 256 bit primitives to broadcast, mask loads and permute data	• Organize, access and pull only necessary data more quickly and efficiently
• Three and four Operands, Non Destructive Syntax – Designed for efficiency and future extensibility	• Fewer register copies, better register use for both vector and scalar code
• Flexible unaligned memory access support	• More opportunities to fuse load and compute operations

Intel® Advanced Vector Extensions is the start of the compute density increase

Intel® Many Integrated Core (Intel® MIC) Architecture

- Up to 32 Intel coherent Intel processor cores on 1 silicon die
- Implements all four salient architectural features of Intel® CPUs
  - x86 Cores, Coherent caches, SIMD, SMT threads
- Enables developers to *scale* applications forward to future Intel® MIC products

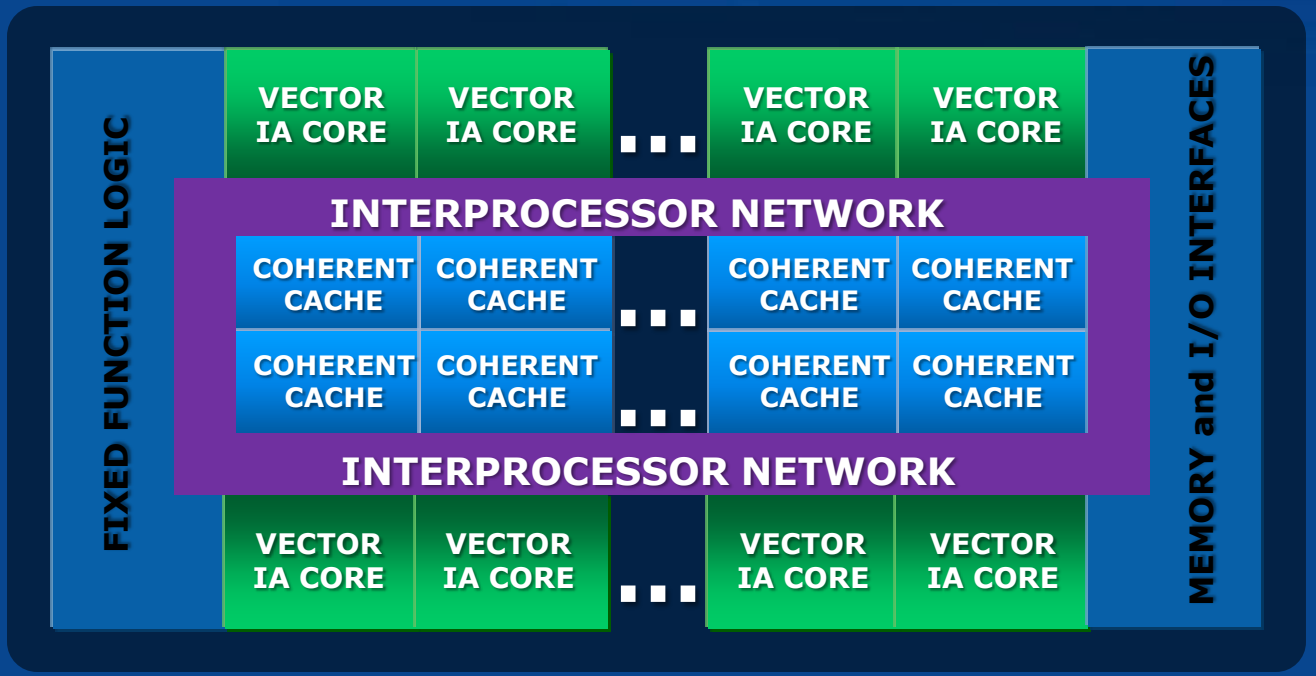

Intel® MIC Architecture Co-Processor

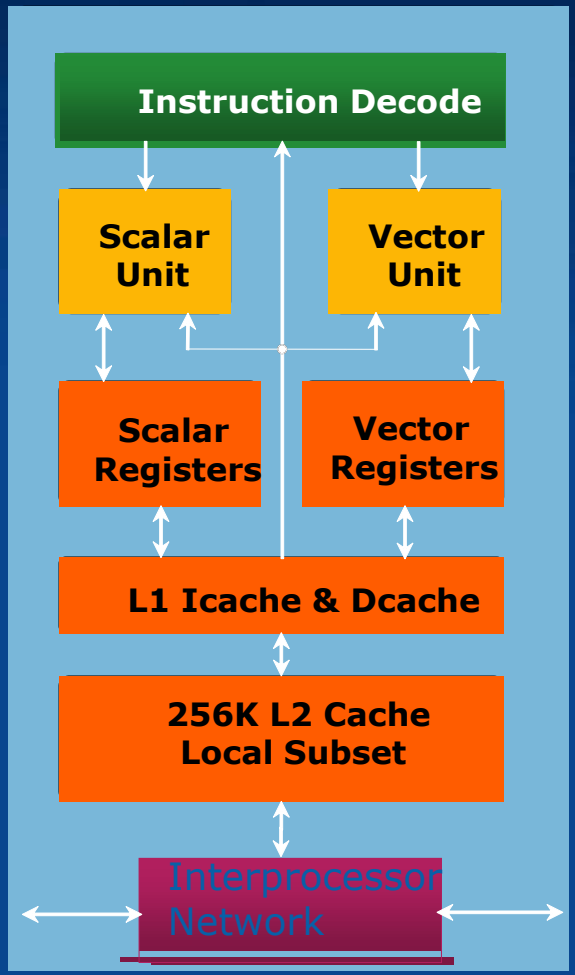

The co-processor core features include:

- Scalar pipeline derived from the dual-issue Intel® Pentium® processor
- Short execution pipeline
- Fully coherent cache structure
- Significant modern enhancements such as multi-threading, 64-bit extensions, and sophisticated pre-fetching
- 4 execution threads per core
- Separate register sets per thread
- Supports IEEE standards for floating point arithmetic
- Fast access to its 256KB local subset of a coherent L2 cache
- 32KB instruction cache per core
- 32KB data cache for each core
- Enhanced x86 instructions set with:
  - Over 100 new instructions,
  - Wide vector processing operations
  - Some specialized scalar instructions
  - 3-operand, 16-wide vector processing unit (VPU)
  - VPU executes integer, single-precision float, and double precision float instructions
- Inter-processor Network with:
  - 1024 bits wide, bi-directional (512 bits in each direction)

10

Program for Intel® Architecture Today and be Intel® Many Integrated Core Architecture Ready

Single Source

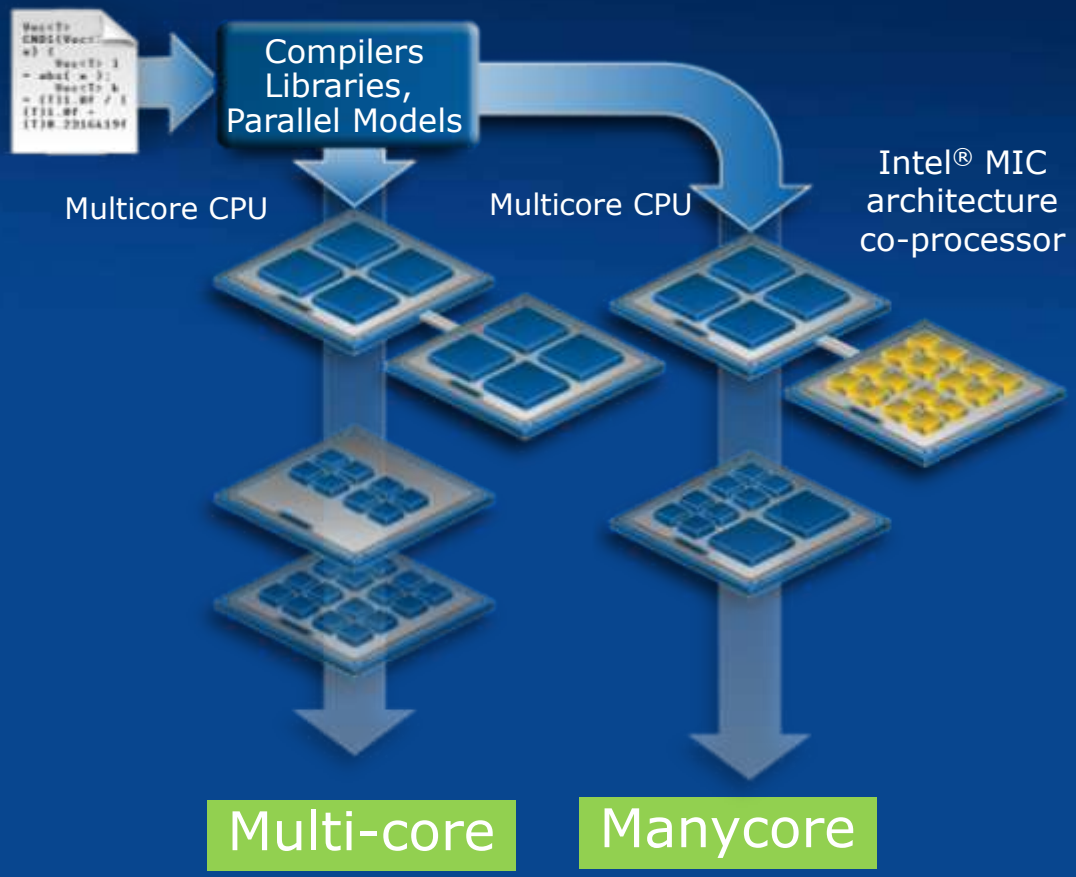

- Full Intel® C / C++ / Fortran compilers and Intel® Math Kernel Library & Intel® Integrated Performance Primitives libraries
- Flexibility of an Intel architecture design allows tools, choice of programming models, and familiar languages
- Programming models that span multi-core Intel architecture and Intel MIC Architecture processors
- Performance acceleration
- Intel architecture ecosystem support

Eliminate Need for Dual Programming Architecture

11

Projected Performance Development

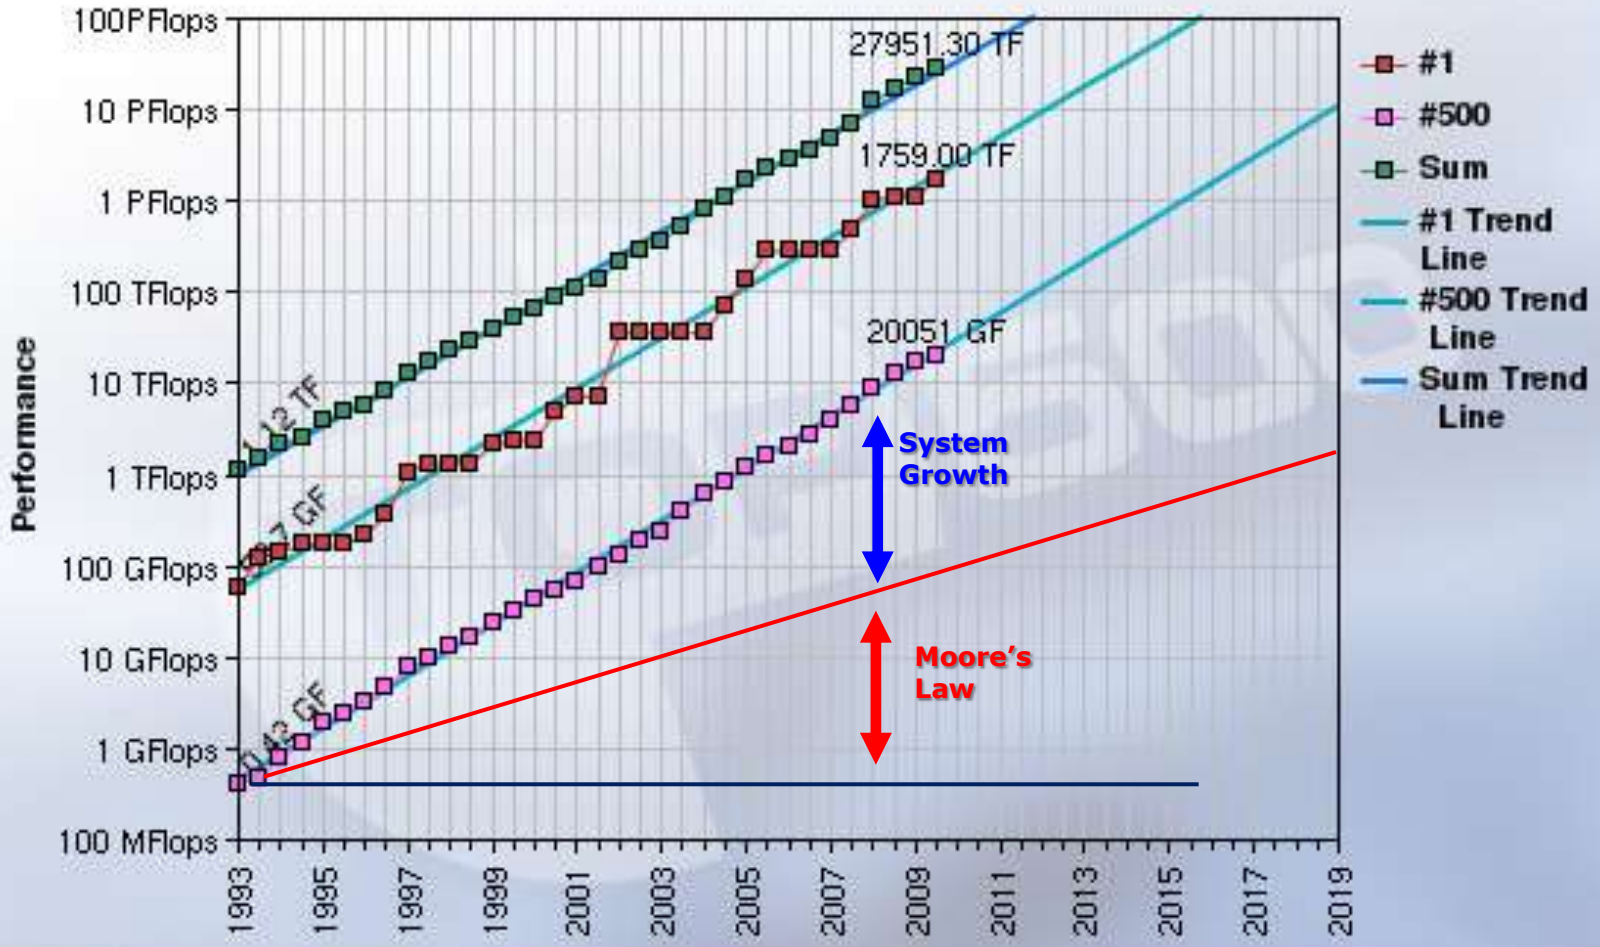

Looking for the Missing Middle

HPC Arenas and Differentiators

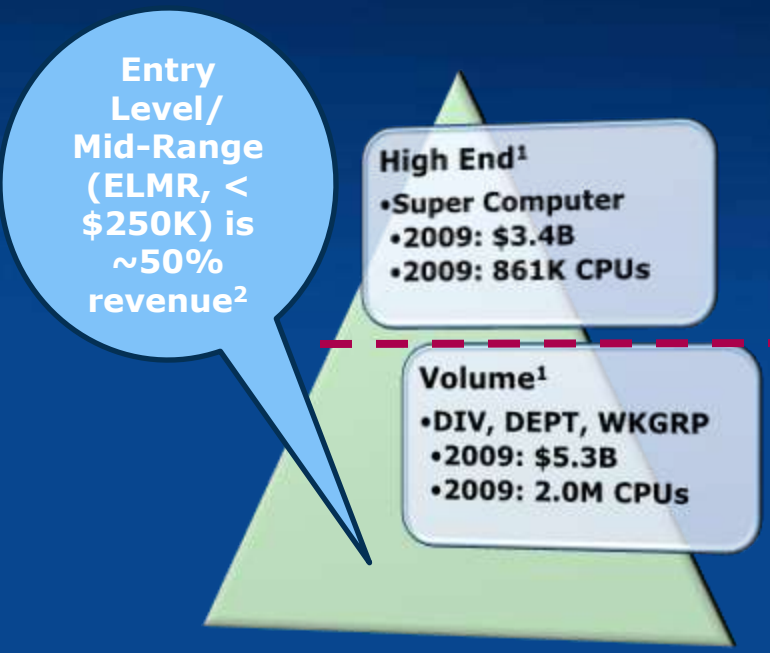

Key Customer Values:
- maximize performance
- early adopters, innovators
- minimize power & cost
- manageability @ scale
- open solutions

Key Customer Values:
- time to productivity
- maximize productivity
- application availability
- robustness of total solution
- true cost of ownership
- ease of acquisition, deployment, mgmt.

HPC is bifurcated into two areas w/several sub-segments
High End HPC and Volume HPC have different requirements

1 – Source: IDC HPC Qview – Q409
2 – Source: InterSect360 Research, Traditional HPC Total Market Model and Forecast, 2009

Implied Perspectives

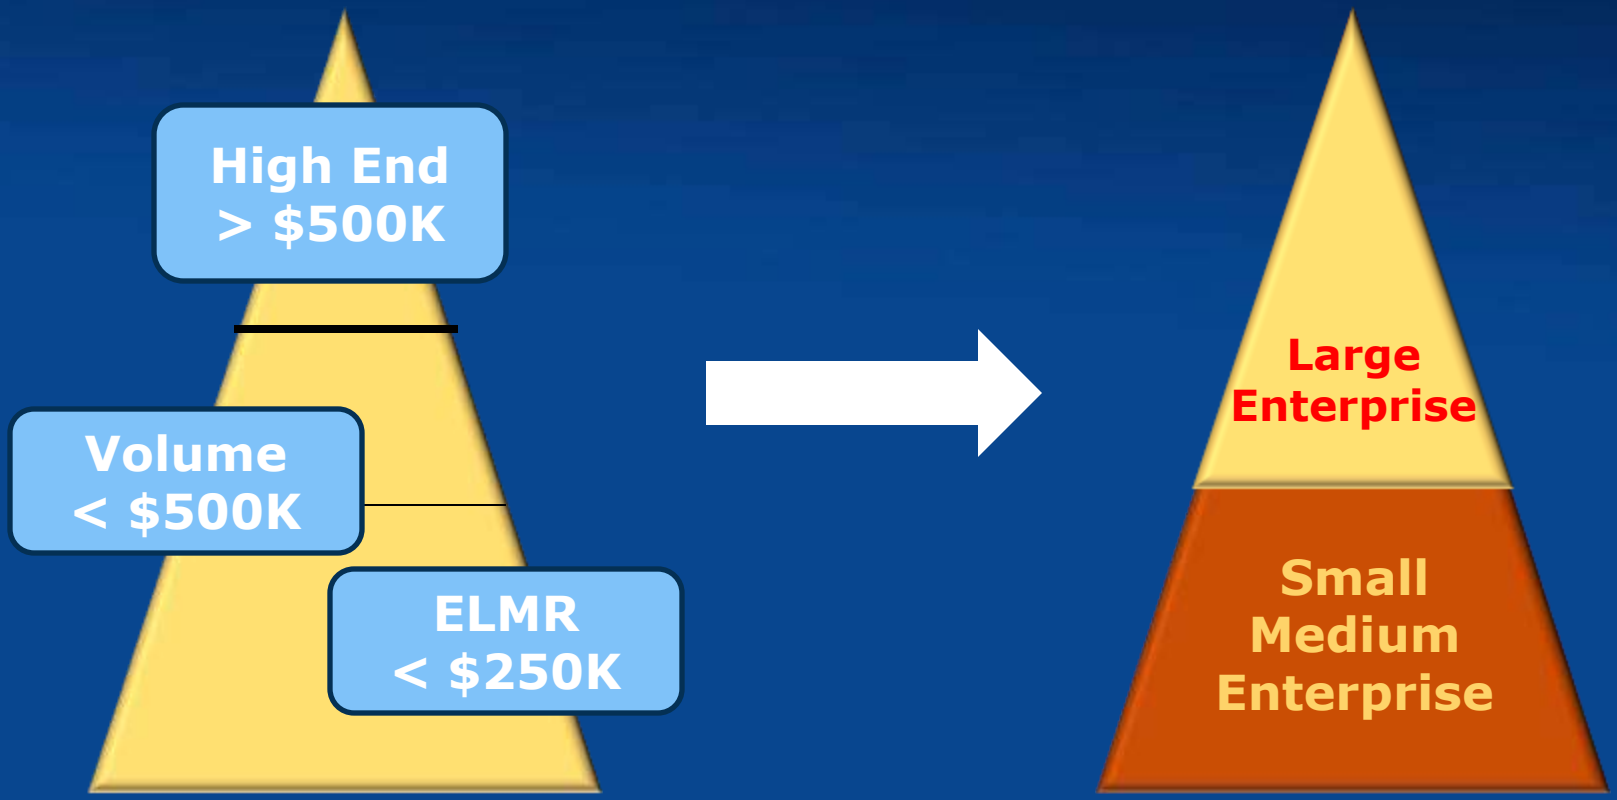

Reality?

- About two-thirds of ELMR-sized (<$250K) systems are upgrades or add-ons to larger systems[1]
- InterSect360 measures that:
  - Of true ELMR systems, 20-25% go to users who also have larger (high-end) systems.
  - so, only 10-15% of said systems go to ELMR users[2]
- IDC sees something similar, with 70%[3] of the <$500K going to the Workgroup, Department, Divisional segments.
  - Needs further visibility/corroboration

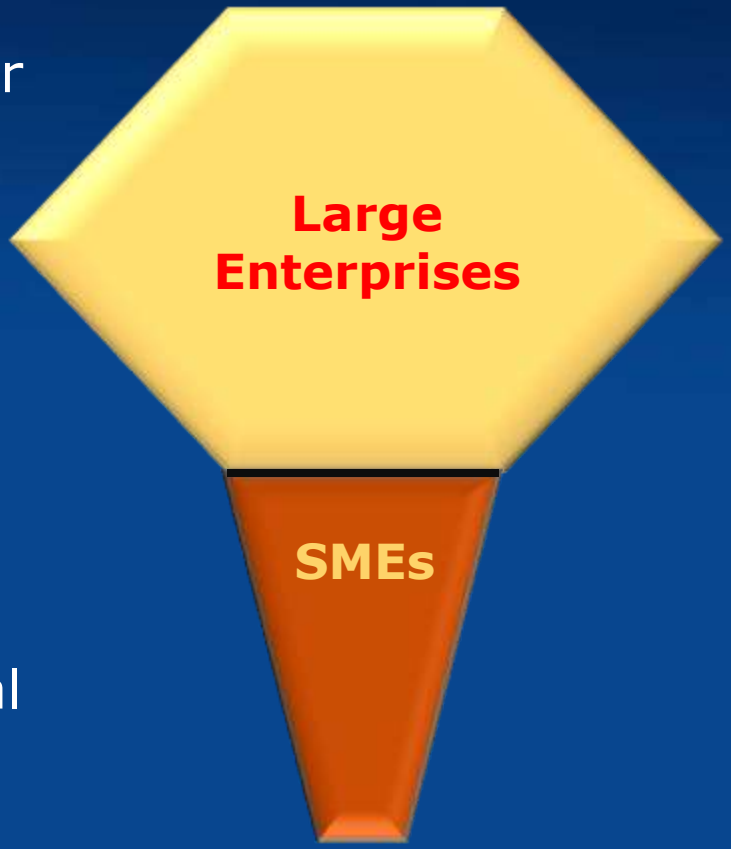

1 – Source: InterSect360 Research, HPC User Site Census: Lifecycles, 2009.
2 – Source: InterSect360 Research, custom user study, 2009.
3 – Source: IDC, personal comms, 2010

A Notional Perspective

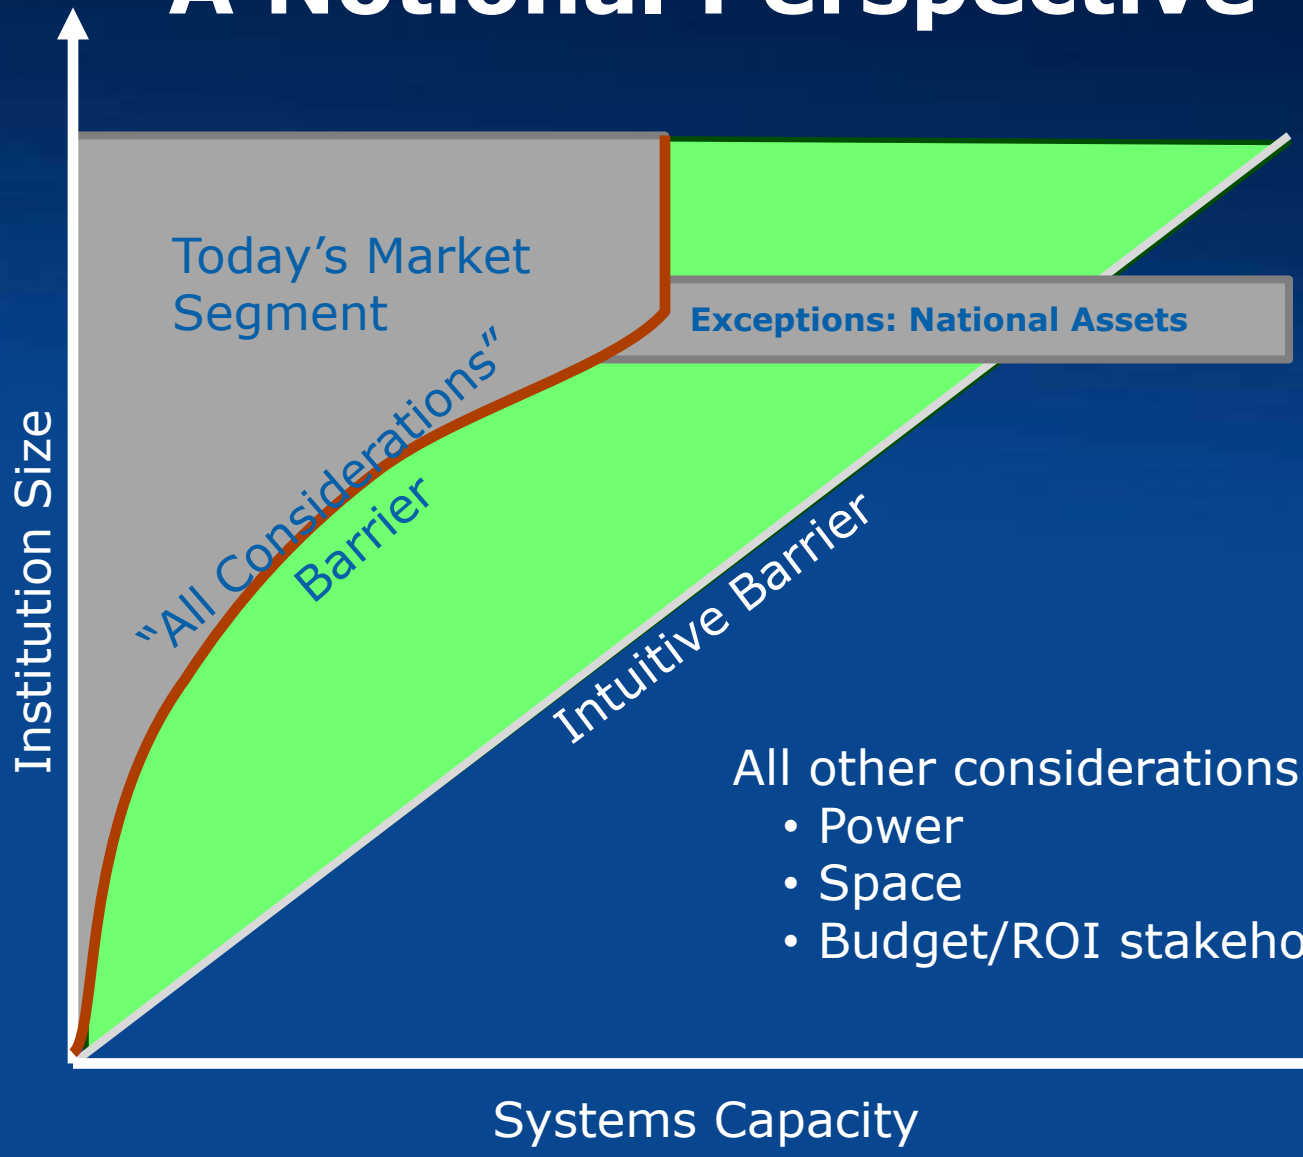

All other considerations:

- Power
- Space
- Budget/ROI stakeholder buy-in

Defining the Missing Middle

From ncms.org

Identifying the Missing Middle

- In a few words, the missing middle is comprised of those institutions that do not use HPC and yet HPC would result in a net-positive ROI to them.
  - Also includes those that use some HPC but not as much as they could

Key Barriers

- The COC/IDC Reveal[1] report concluded that there are three major system barriers stalling HPC adoption:
  - Lack of Application Software
  - Lack of Sufficient Talent
  - Cost constraints
- They noted that these were the same constraints identified four years prior[1,2]
- InterSect360[3] had a similar perspective; that cost is not the top barrier.
  - "You could give companies free HW and SW, and it wouldn't solve these problems:
    - Political will to change a workflow and to build faith in simulation to supplement physical testing.
    - Expertise and knowledge for using scalable systems, and
    - Creation of digital models."

1 – Source: CoC/IDC Reveal report, 2008.
2 – Source: CoC Study of US Industrial HPC Users, July 2004
3 – Source: Addison Snell, InterSect360

Barriers Summarized

This is the Right Time

- Several key events are transpiring, thus making this the first time in history to successfully tackle the missing middle problem
  - Parallel everywhere, and becoming more so each year
    - Recent launch of the Intel® Xeon® Processor 5600 Series
    - And the launch of Nehalem-EX
    - Tick-Tock
  - Everyone in DC focused on jobs
    - A 20th Century work force needs 21st Century job skills
  - International competitiveness increasingly defined by innovation as enabled by modeling and simulation
  - Large-scale industry and USG are motivated to solve the problem
- The environment is aligned for action; but what action?

Action on the Right Issues

- Processors and programming paradigms are not the right issues.
  - We have more than enough performance now and on the roadmap for the foreseeable future
  - Basic software tools are sufficient

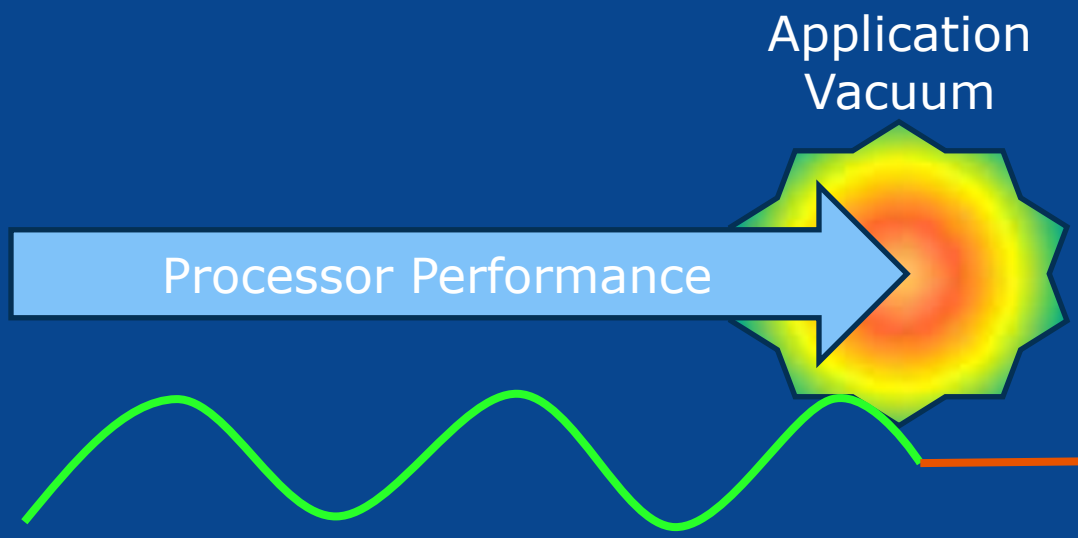

Catching Their Attention

- Must make computational modeling and simulation:
  - Easy to use
    - Application Frameworks
    - End-user specific infrastructure
  - Deliver computational continuity
    - Scaled use
    - Seamless compatibility
    - Affordable access models
  - Easy to see ROI

Putting the Pieces Together

- Need a modeling and simulation supply chain
- The missing middle is comprised of a few communities
  - The actual business community – needing that competitive edge
  - The liaison community – facilitating the tech transfer
  - The application framework community – developing relevant software environments
  - The speculative innovation community – aka academic research teams focused on local industry mod/sim needs

National Digital Manufacturing Strategy

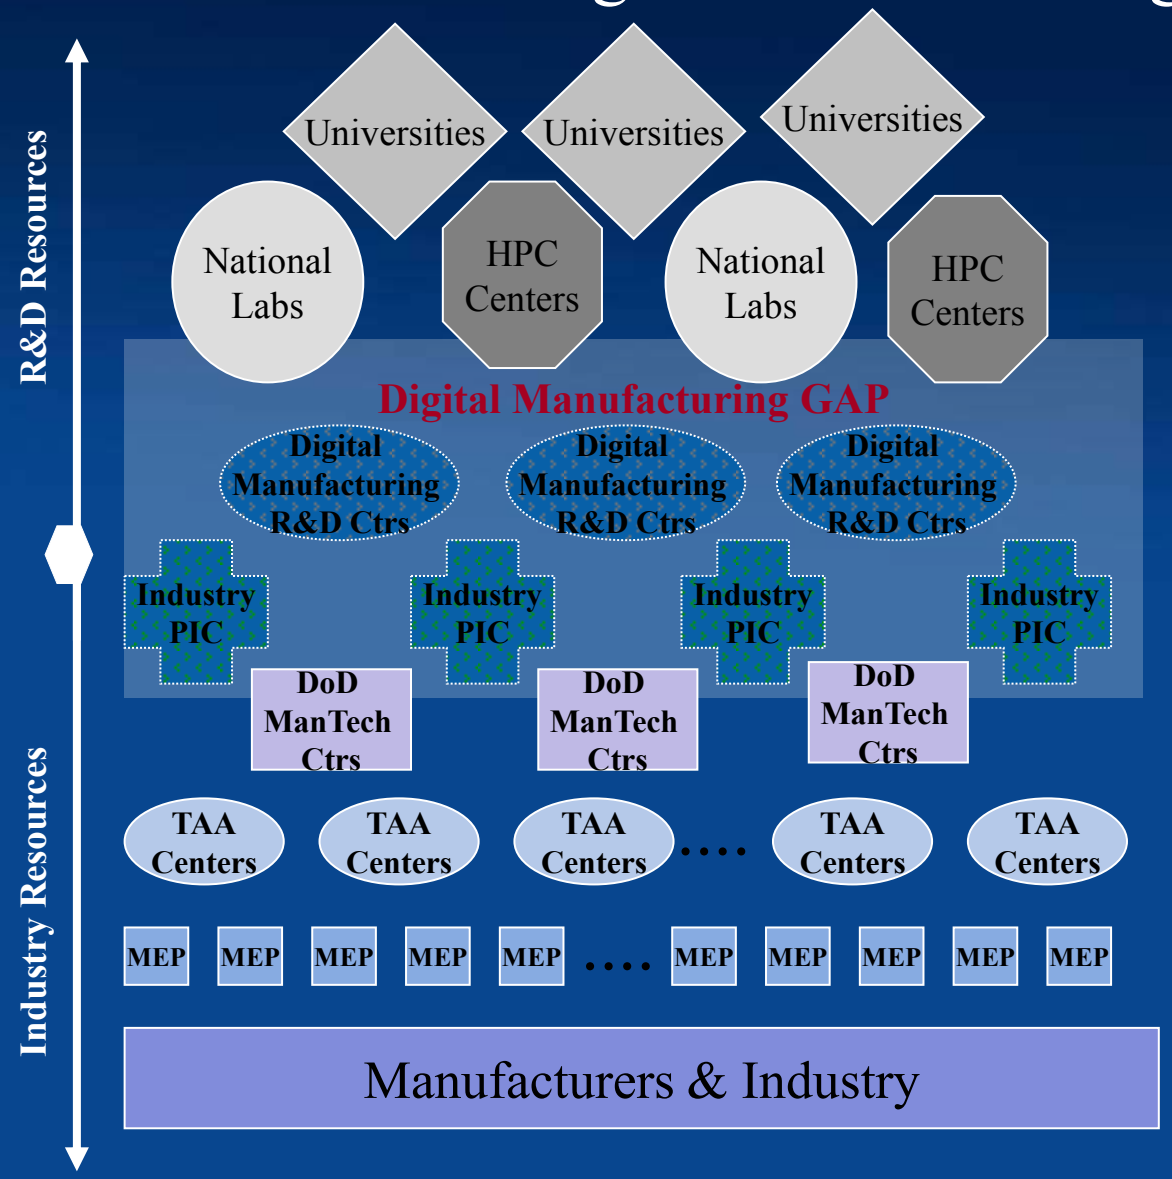

Existing R&D Expertise

- Universities
- National Labs
- DoE Labs
- HPC Centers (i.e. OSC, NCSA, etc.)

Proposed National Manufacturing Innovation Network

- Digital Manufacturing R&D Centers (academic focus)
- Industry Predictive Innovation Collaboration Centers (non-profit e.g. NCMS)

Trade Adjustment Assistance Centers (TAAC)

- Approx. 14 National Centers
- Expand mission beyond trade impacted companies

MEP's (NIST)

- 60+ National Centers
- New focus on Digital Manufacturing

Focused Digital Manufacturing Training

- Community colleges, NAM, Manufacturing web portals

Requires an Industry Effort

- This isn't something Intel can or should do alone
- It will require a concerted and determined effort on everyone's part
- It may require a USG managed agenda
- It won't happen over night

- History is in the making!

- Established to pursue solutions to the barriers facing the Missing Middle in US Manufacturing
  - "Transforming American Manufacturing for Economic Growth"
- Comprised of more than 35 entities, from:
  - Computer OEMs
  - ISVs
  - Academia
  - Manufacturing
  - National Labs
- Recent results:
  - America COMPETES Renewal language for IAWG
  - Further analysis: results released via NCMS on 9/30/2010
  - Industry Recognition Initiative launched at IDC HPC User Forum on 9/14/2010

Is there more?

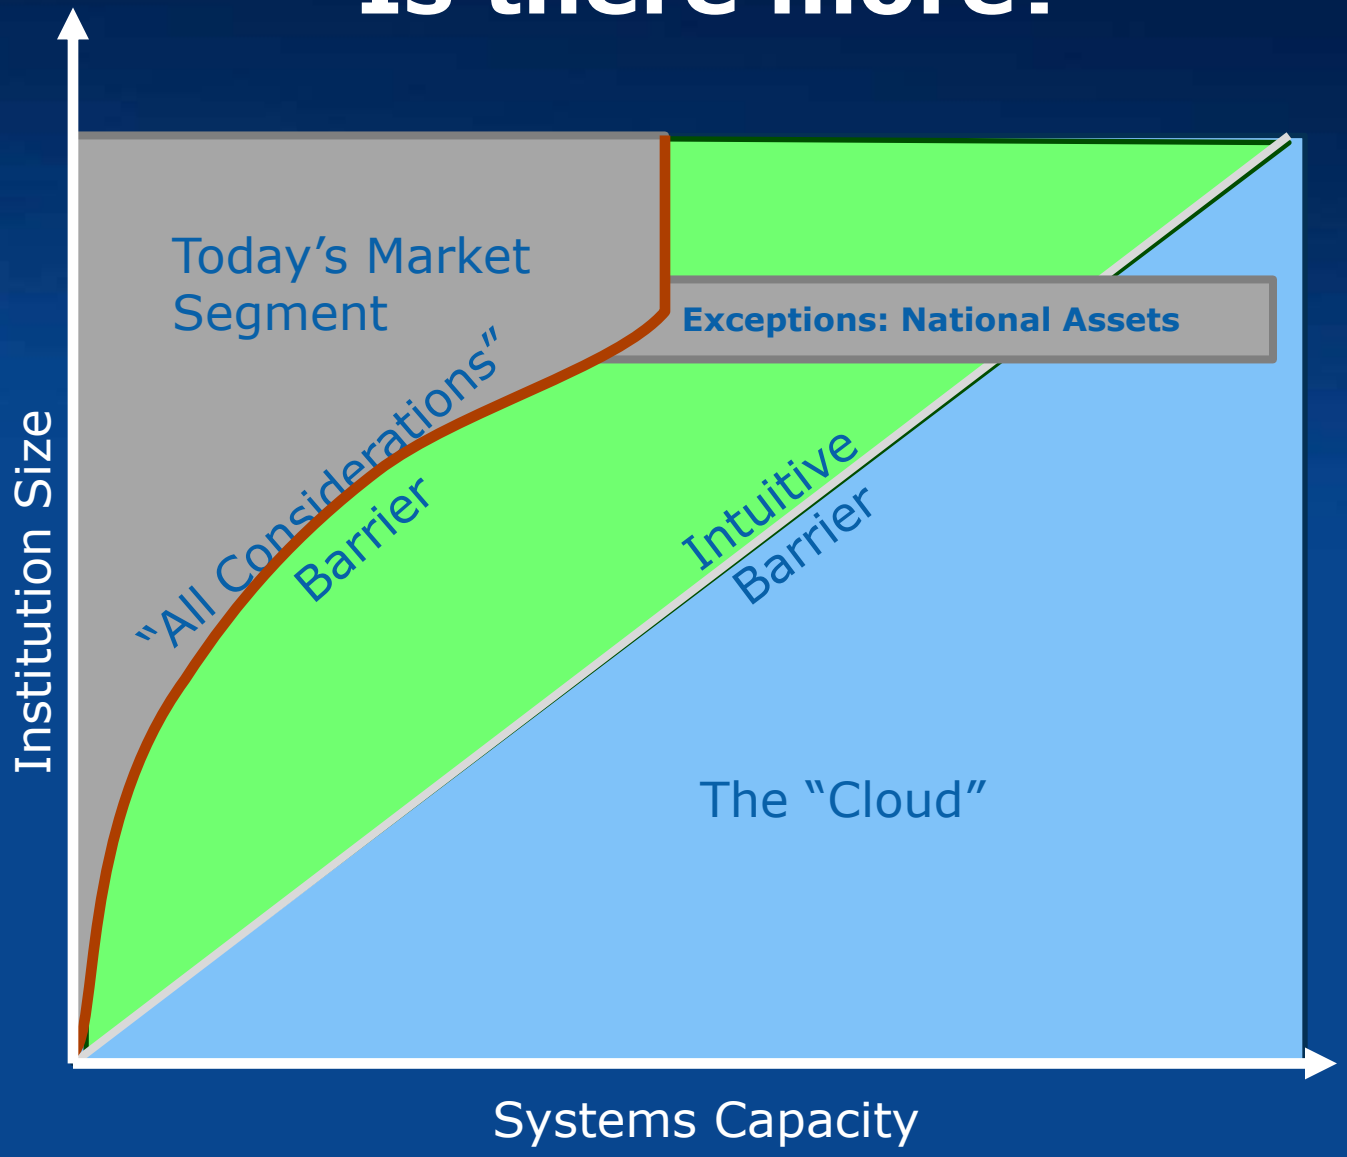

Definition of Success

- When the middle isn't "missing"

Thank you!

HPC @ Intel