Supercomputing & Big Data
A Convergence?

Randal E. Bryant
Office of Science and Technology Policy

Two Classes of Large-Scale Systems

- **Modern supercomputer**

  - Run programs in hours or days that would require decades or centuries on normal machine
  - Designed for numerically-intensive applications

- **Internet Data Center**

  - Support millions of customers
    - Mostly small transactions
    - + large-scale analytics
  - Designed for data collection, storage, and analysis

Computing Trends

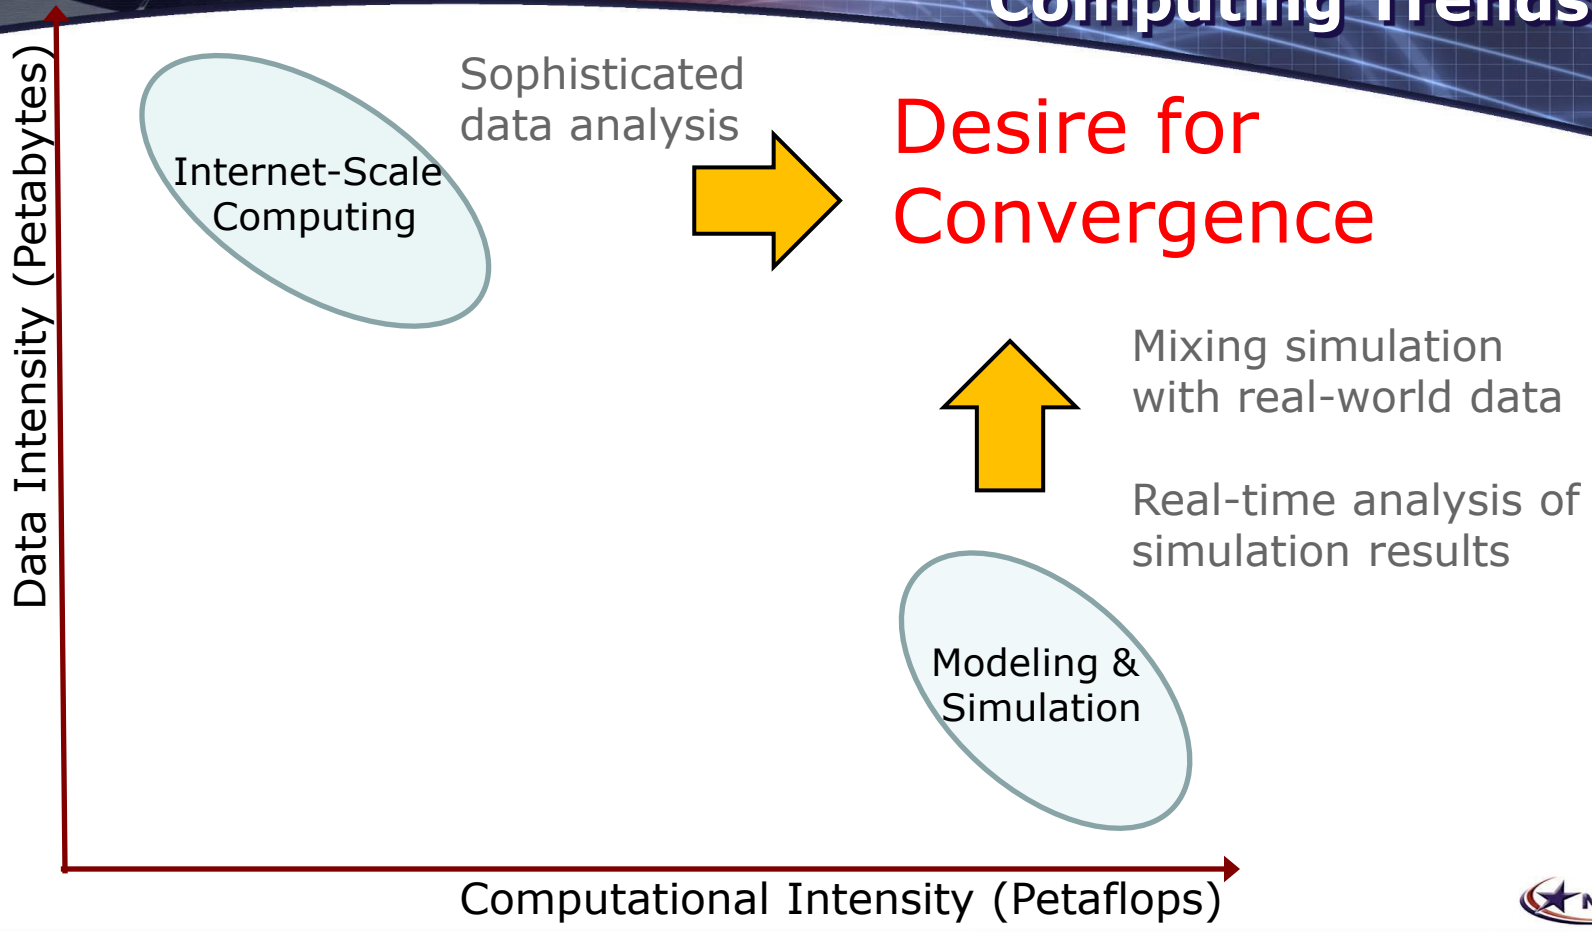

Titan Hardware

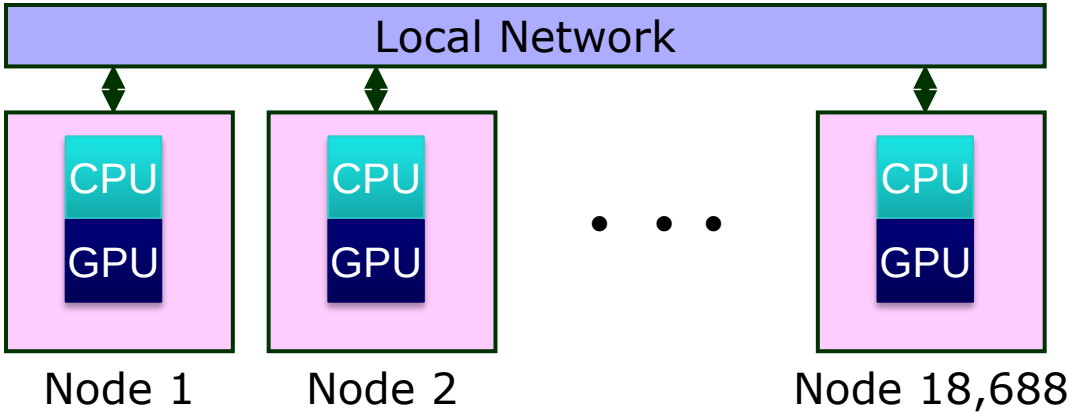

- **Each Node**
  - AMD 16-core processor
  - nVidia Graphics Processing Unit
  - 38 GB DRAM
  - *No disk drive*

Network Cluster

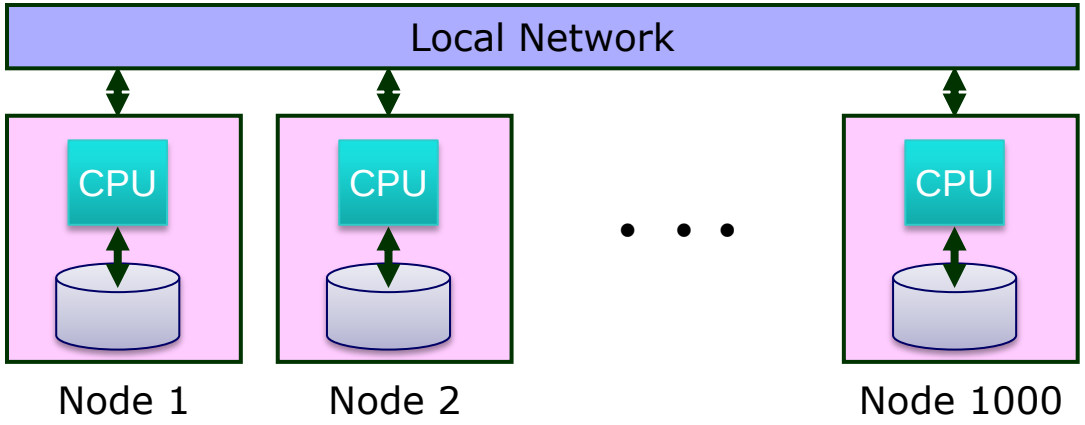

- **Typical Node**
  - 2 multicore CPUs
  - 2 disk drives
  - DRAM
- **Enhancements**
  - Flash memory
  - GPUs
  - Fast / high-bandwidth networking

Challenges for Convergence

Supercomputers	Network Clusters
Hardware	
▪ Customized	▪ Consumer grade
▪ Optimized for reliability	▪ Optimized for low cost
Run-Time System	
▪ Source of "noise"	▪ Provides reliability
▪ Static scheduling	▪ Dynamic allocation
▪ Strive for homogeneity	▪ Tolerate variability
Application Programming	
▪ Low-level, processor-centric model (e.g., MPI)	▪ High level, data-centric model (e.g., Hadoop, Spark)

Issues to Be Discussed

- What can the supercomputing and big-data communities learn from each other?
- Can and should the technologies for big data and high-fidelity HPC simulation really merge?
- What new classes of applications would arise through a convergence?
- What research is needed to enable a convergence?