PMI RESEARCH & DEVELOPMENT

HPC Enabling R&D at Philip Morris International

Jim Geuther*, **Filipe Bonjour†, Bruce O'Neel†, Didier Bouttefeux†, Sylvain Gubian†, Stephane Cano†, and Brian Suomela†**

* Philip Morris International IT Service Center Sàrl (ITSC)

† Philip Morris Products S.A., Research & Development (R&D)

May 23rd 2013

Presented at the *Forum on HPC in Swiss Industry* (Agenda here)

Why is a tobacco company using HPC?

- R&D for new products that may have the potential to reduce the risks of smoking-related diseases.
- Understand underlying disease mechanisms, biomarker discovery, bioinformatics, text mining, genome sequencing, CFD and molecular dynamics simulations.
- Our HPC infrastructure is hosted in a data center aligned with more traditional enterprise IT requirements (SOX, GLP,…), providing opportunities to leverage existing company standards and best practices.

Service Delivery Model for HPC

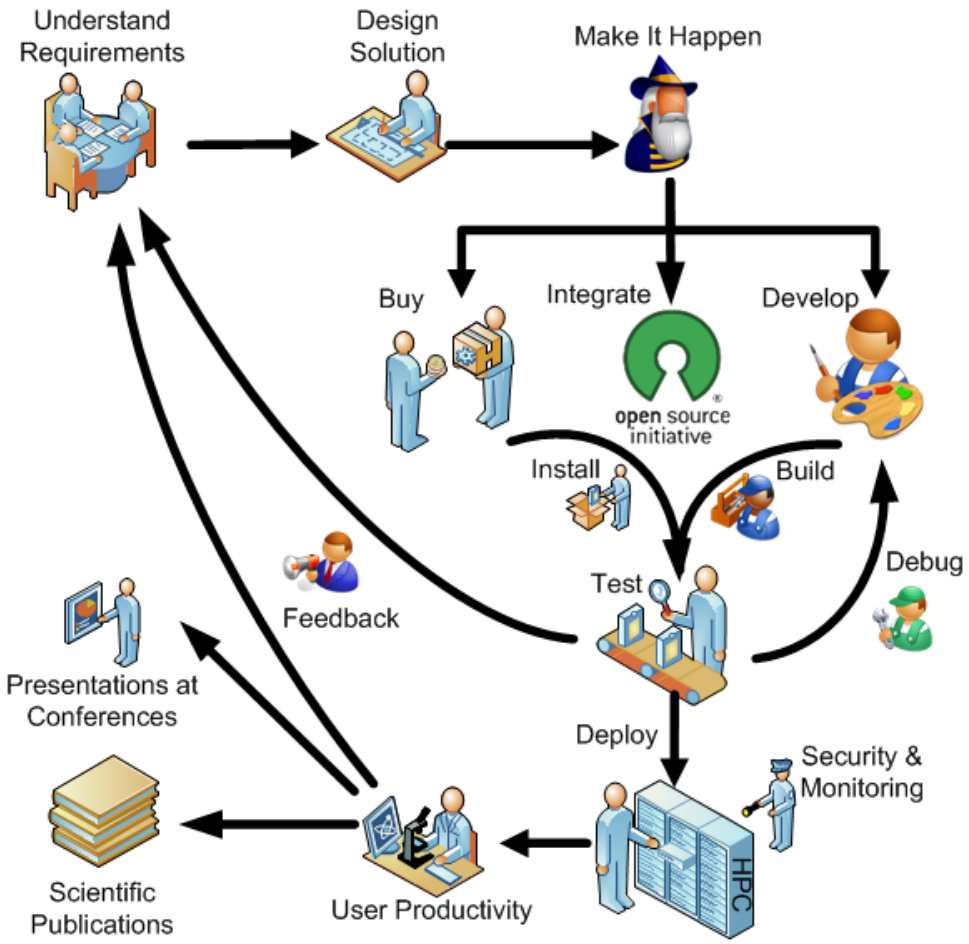

Dedicated HPC support team embedded within R&D:

- Sysadmins
- Software developers

Collaborate with users, anticipate needs, rapidly deliver solutions

2nd level support from shared data center IT support team

HPC Infrastructure Topology

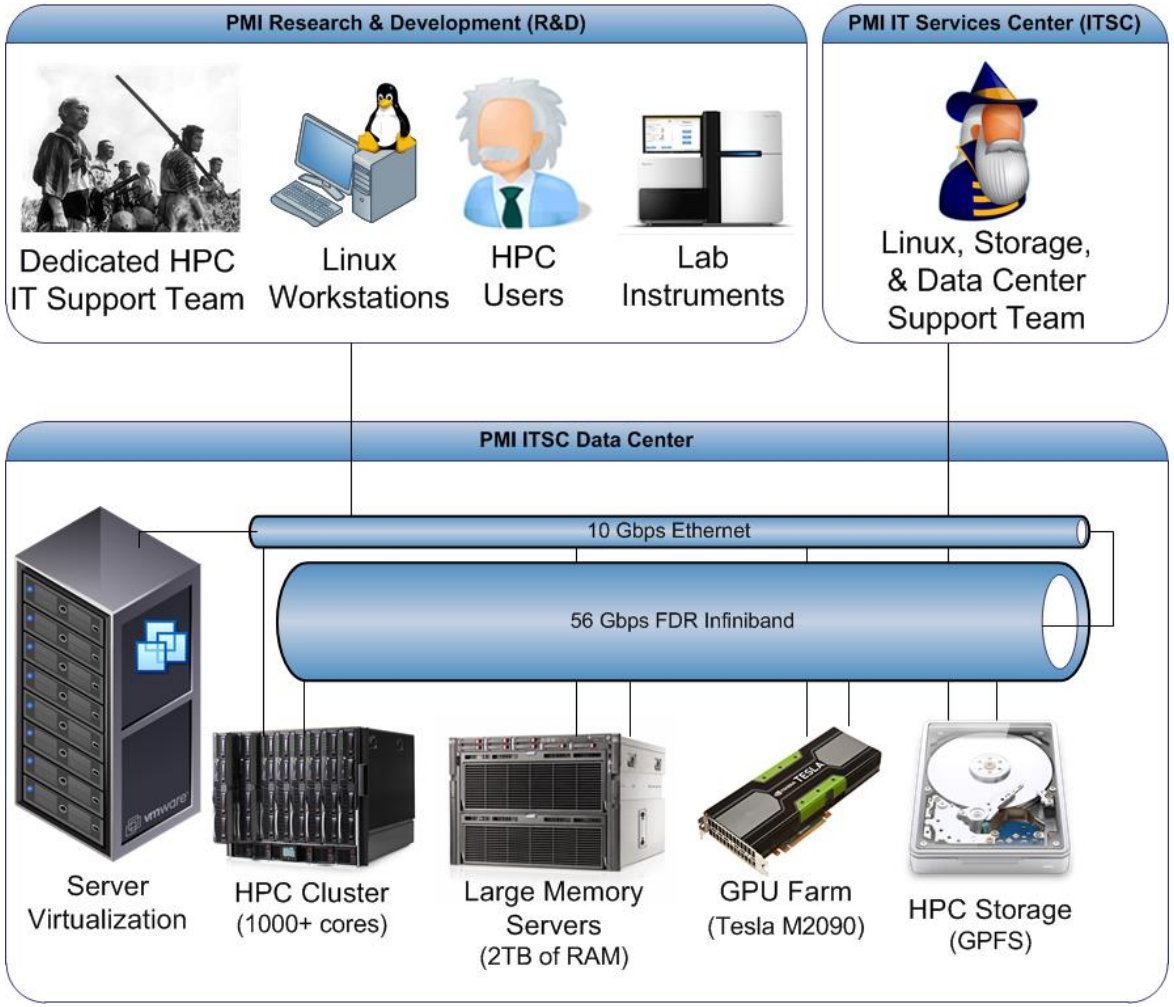

Large memory server requirements

- A growing proportion of our scientific computing workloads have very large memory requirements (>100GB/core).
- Users must wait in the queue for several days for jobs to run on our “Bigfoot” server (2TB RAM, 64 cores).
- We have recently ordered more servers each with 256GB or 768GB RAM.
- A small number of users tend to consume most of these resources.

“Bigfoot” CPU utilization

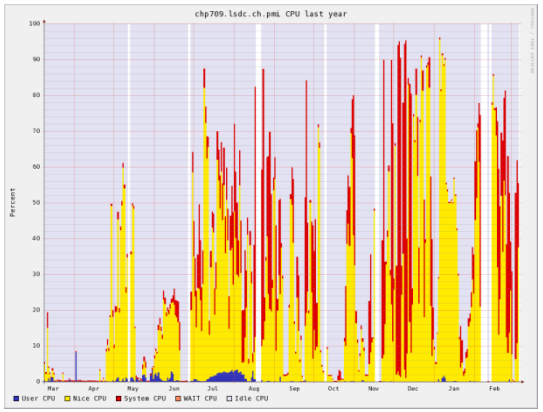

“Bigfoot” Memory utilization

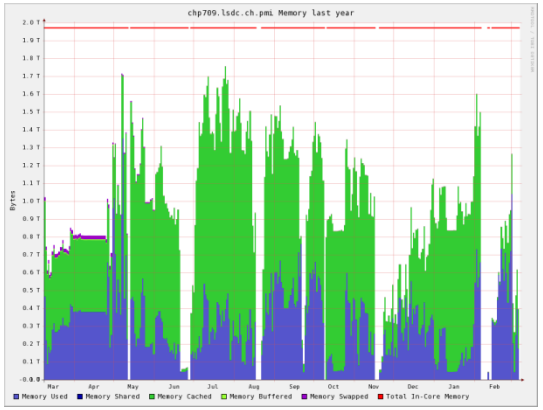

PMI RESEARCH & DEVELOPMENT

HPC cluster node memory requirements

- Our older cluster nodes (2 x E5450) have only 32GB RAM (4GB/core).
- Many jobs need 10GB/core.
- To avoid running out of server memory, some users typically calculate how many extra cores to reserve for single-threaded jobs.
- Cluster utilization appears well below 100% even if the job scheduler queue is full of CPU-intensive workloads.
- Our new cluster nodes (2 x E5-2680) will have 256GB (16GB/core).

Old Cluster CPU utilization

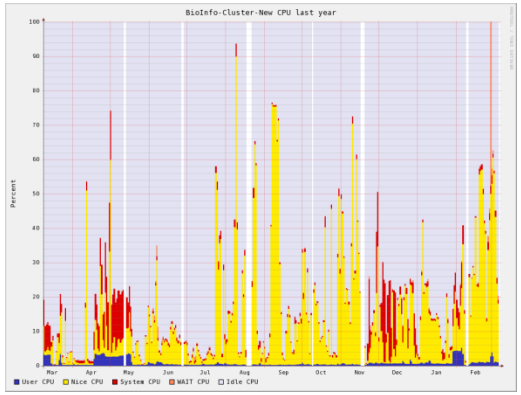

Old Cluster Memory utilization

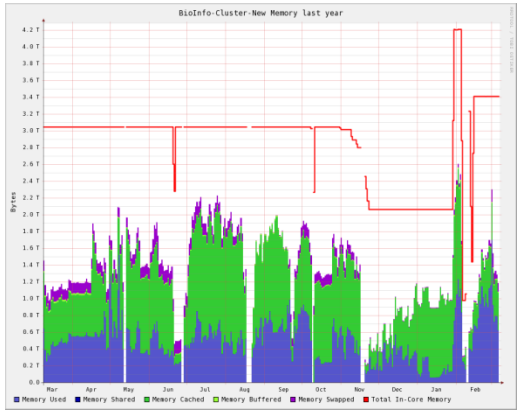

PMI RESEARCH & DEVELOPMENT

Tick, Tock…

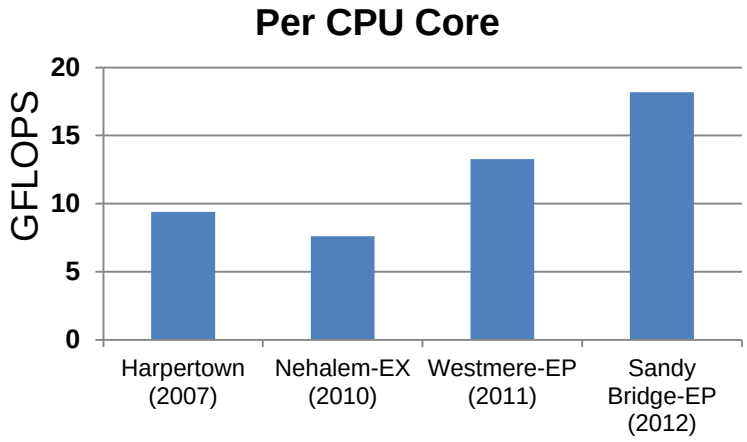

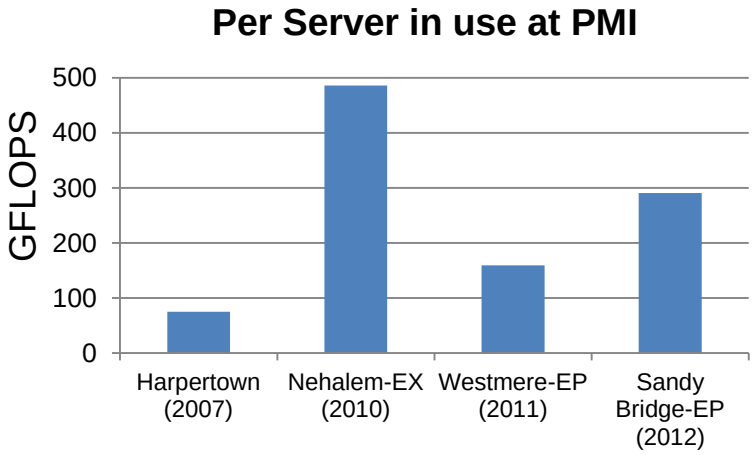

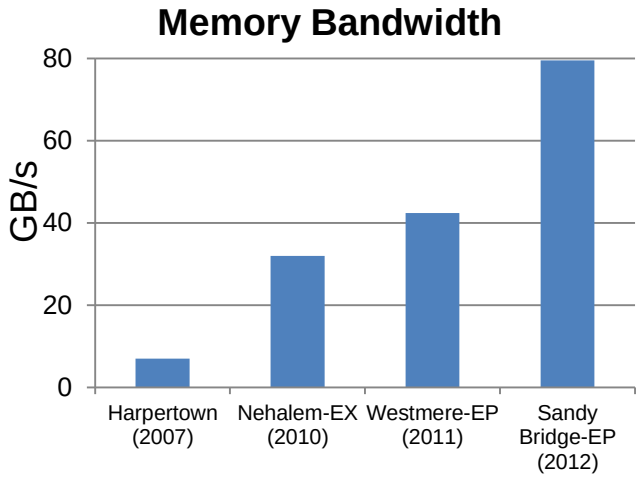

GFLOPS = LINPACK results from Intel and the TOP500 List .
Memory Bandwidth = STREAM benchmarks by Intel and TACC.
Harpertown = 2 x Intel Xeon E5450 (8 cores)
Nehalem-EX = 8 x Intel Xeon X7560 (64 cores)
Westmere-EP = 2 x Intel Xeon X5690 (12 cores)
Sandy Bridge-EP = 2 x Intel Xeon E5-2680 (16 cores)

Upgrades to our HPC cluster

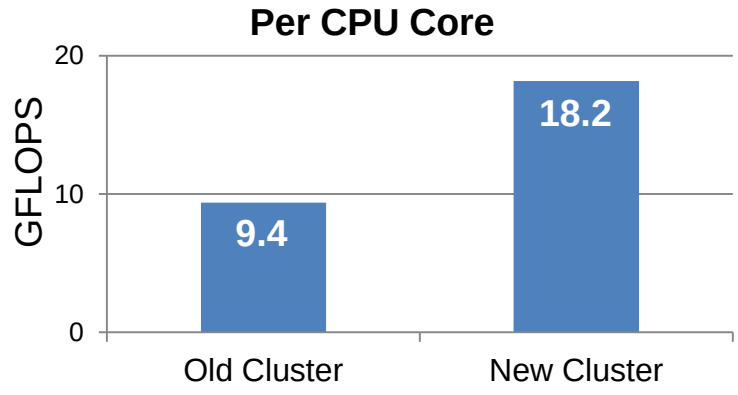

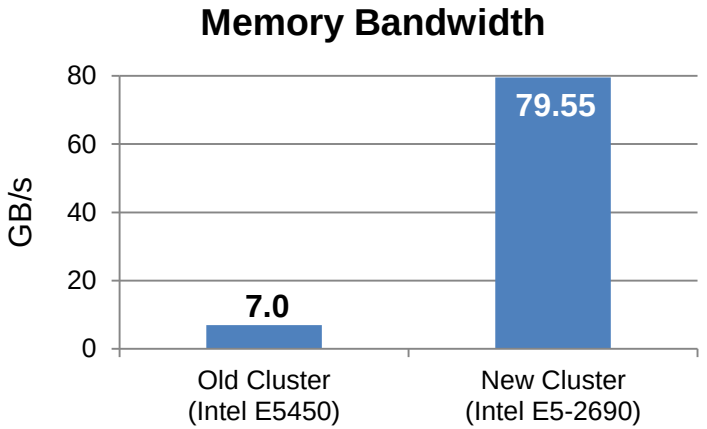

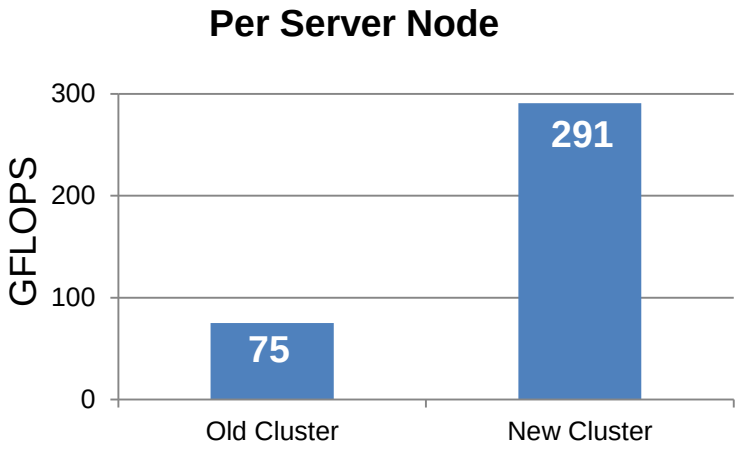

GFLOPS = LINPACK results from Intel and the TOP500 List .
Memory Bandwidth = STREAM benchmarks by Intel and TACC.
Old Cluster = 2 x Intel Xeon E5450 (8 cores)
New Cluster = 2 x Intel Xeon E5-2680 (16 cores)

PMI RESEARCH & DEVELOPMENT

GPU is interesting for certain HPC workloads

- Synthetic benchmarks may not translate as well as expected into real world performance for your application.
- In 2012 we completed a small pilot using:
  - Sixteen NVIDIA Tesla M2090 (Fermi) cards (~0.47 TF / card)
  - One NVIDIA Tesla K20 (Kepler) card (~0.7 TF / card)
- Tested several applications which consume the bulk of our HPC resources (many are developed in-house).
- Up to 2.5x speedup for applications that perform many iterations of large matrix multiplications.
- Code optimization efforts will be required which causes us to delay a decision to scale-up our GPU compute farm.
- Further testing is planned with CFD simulation workloads.

Scientific Publications by PMI R&D (enabled by HPC)

As of April 2013 we have published over 20 papers in peer-reviewed scientific journals and 80 presentations that would not have been possible without HPC:

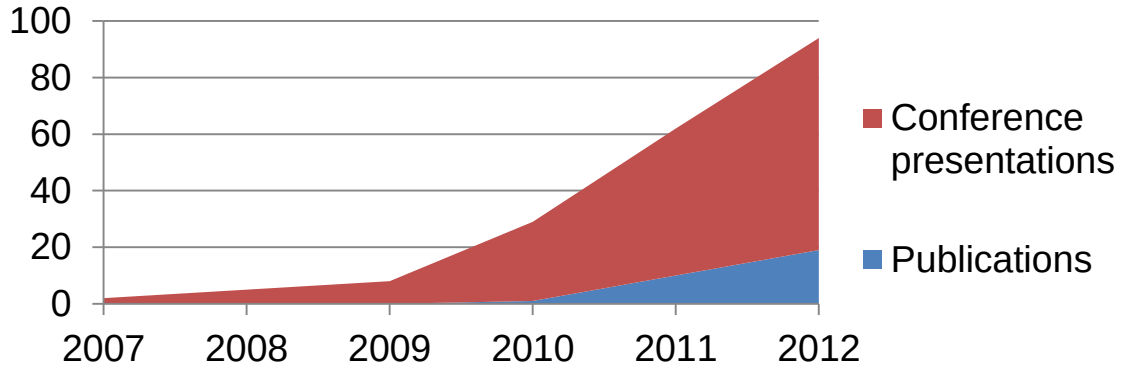

Examples:

- Meyer P, et al. *Verification of systems biology research in the age of collaborative competition*. **Nature Biotechnology** 2011 Sep 8;29(9):811-5.
- Hoeng J, et al. *A network-based approach to quantifying the impact of biologically active substances*. **Drug Discovery Today** 2012 May;17(9-10):413-8.
- Martin F, et al.. *Assessment of network perturbation amplitudes by applying high-throughput data to causal biological networks*. **BMC Systems Biology** 2012 May 31;6:54.
- Laino T, et al. *Mechanisms of Propylene Glycol and Triacetin Pyrolysis* **Journal of Physical Chemistry** 2012, 116 (18): 4602–4609

Thanks for your attention

- Questions?