

Transformative Technologies for Computing

Ian Fisk

August 4, 2013

Previous Transformations

- ✦ The first transformative technology transition for me was the move to x86 and Linux for the field
 - ✦ Suddenly computer costs dropped by a factor of 100 and an operating system we could use to make clusters



- ✦ Now we have copious computing that importantly scales with the amount of money you commit



Game Changes

- ✦ Suddenly computing capacity was copious and scalable
 - ✦ Introduced the graduate student sys-admin



Computing Transformations

- ✦ There are times in computing when suddenly what can be done and how you work changes dramatically
 - ✦ It's not different from turning on a new accelerator or developing a new detector, but it's not always entirely predictable or under our control
- ✦ Sometimes things are only seen as transformative looking back.

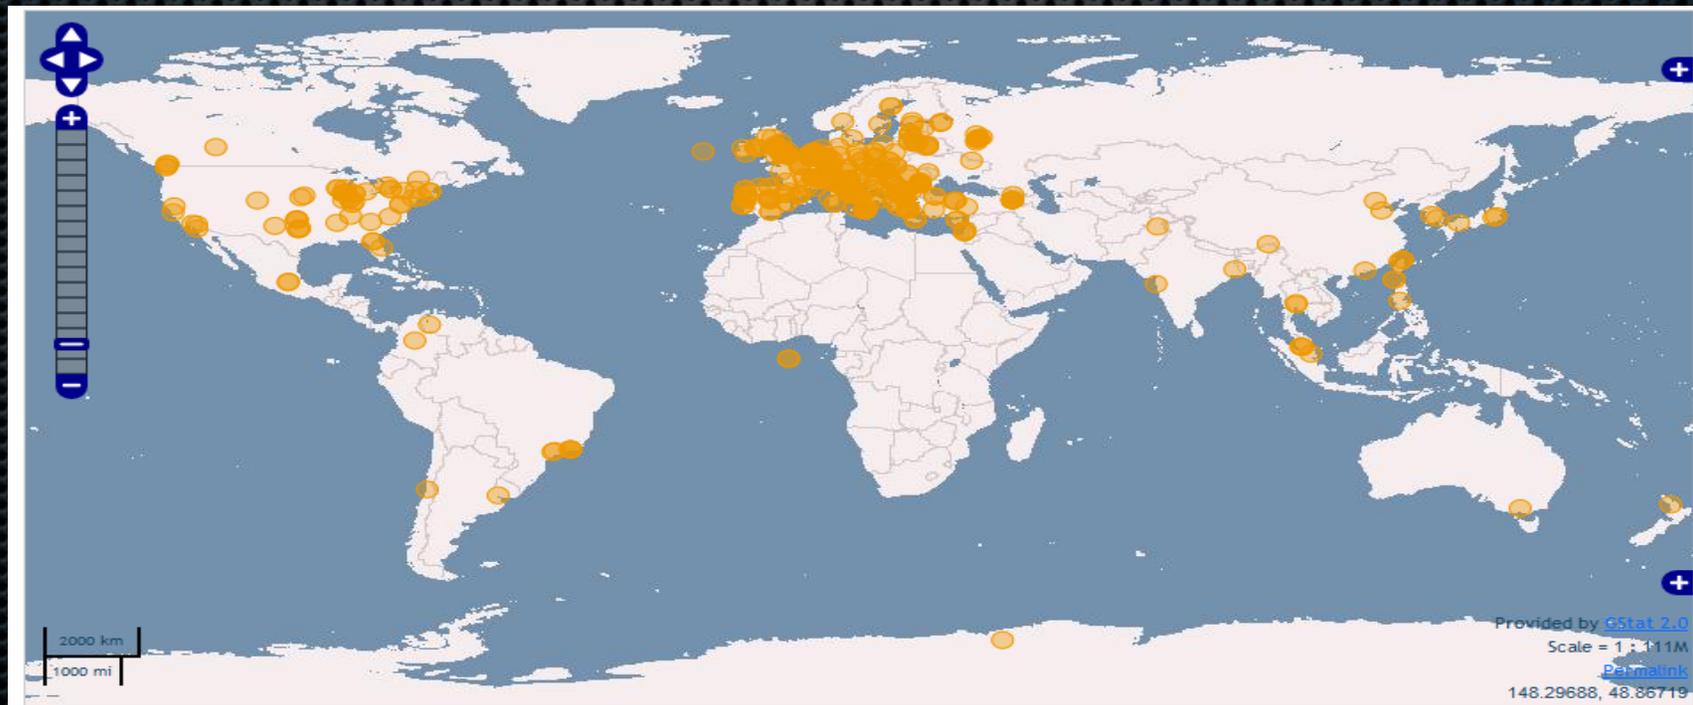
The Grid

- ✦ The grid is a combination of services to allow clusters to be accessed for processing and storage
 - ✦ There was no technical reason all the computing could not be located in one place, but we have more support and more resources this way
- ✦ The grid allows globally distributed systems to function as a coherent resource for solving computing problems
- ✦ Only possible with the development for middleware and services and operations effort



The Network

- ✦ This was made possible by the fact we can move data between sites
- ✦ Data rate out of LHC host lab and between LHC clusters are a factor of 100 higher than the Tevatron at the same time in the life of the experiments



What comes next?

- ✦ We are in the capacity computing regime, and we have a lot of capacity both distributed and locally to communities
- ✦ Solving the distributed data management problem would be a transformative technology change
 - ✦ Eventually we will have to evolve to a Data Intensive Content Delivery Network that delivers data to diverse resources
 - ✦ It will need intelligent caching and transfers replication and clean-up
 - ✦ Users should expect to request samples and have them delivered regardless of location

Flexibility

- ✦ We have ways of packaging the local environment and deploying it even in user space. We can make virtual machines if we have to. We can take very diverse systems and make them look the same to our application
- ✦ Not having to care about data location and controlling the environment is a transformative change
 - ✦ Anything you bring to the table can make a contribution
 - ✦ Local, opportunistic, and temporarily allocated contributions can be a huge increase in capacity or capacity under local control

Data Networks and Packaged Environments

- If everything you have produced is accessible from anywhere you are in a scaleable way and you can bring your environment with you, then all resources are potentially yours
 - There is a transformative technology happening now

AAA++



Data Intensive Content Delivery



Scheduling for Peak

- ✦ About half of what we do is continuous activities where the peak and the average are similar.
 - ✦ Prompt Reconstruction happens shortly after events are collected. Analysis has some fluctuations around conferences, but generally is a sustained load
- ✦ The other half are activities with a goal, but the time is defined by how much computing we have
 - ✦ Reprocessing, simulation, etc.
- ✦ Currently we provision resources for average load

Balancing Resources

Ian Fisk FNAL



- For many parts of the program we do use an average load,
- However there are benefits to growing to peaks that are much larger than the average and then have sustained period of lower than average usage

Peak and Average

- Ideally one makes contributions and budgets for an average utilization, but gets peak delivery
- Depends on working with places that support diverse communities that operate on different timescales. And solving data management
- Requires continued development on resource provisioning interfaces

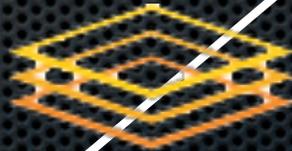
Us



Commercial Providers



Multivo Infrastructures



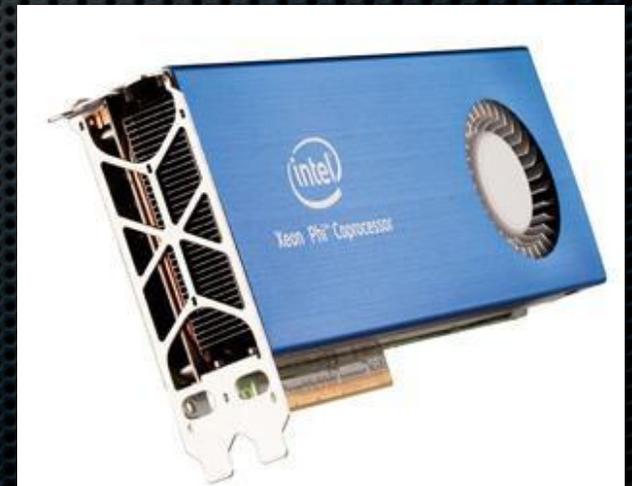
Open Science Grid



Allocations

Specialized Gear

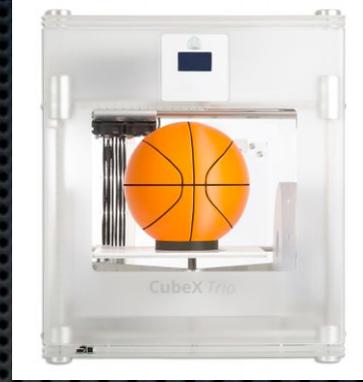
- Specialized systems like GPUs and many core co-processors with enough programming effort have the potential to transform how we compute
 - Matches well to a model where we get an allocation for a defined application



Not Using Resources Costs Too

- ✦ If you assume a processing system has a life of 48 months, a cost of \$2500 to buy, and a lifetime operating cost equal to the hardware cost, then not using it costs \$0.15 an hour
 - ✦ Amazon is more expensive than that, but without additional communities to give unused resources to, preferably with the hope of getting them back, our model has costs too.
- ✦ We would like to use facilities shared with a large and diverse community. Allowing peak delivery to be much larger than average.

Virtual Data



- ✦ If I were a technologist, I would predict within 10 years you will have a 3D printer in your house
 - ✦ Selling you the instructions for making something will be cheaper, more energy efficient and faster than storing things in a warehouse and delivering them when you need them
- ✦ If this is the case, why I am spending millions of dollars storing derived datasets
 - ✦ That could be stored as the instructions needed to make them

Virtual Data Realized

- ✦ CMS wrote about 12PB of raw data during the first run of LHC (including the duplication)
 - ✦ CMS has 60PB of disk and 65PB of tape
 - ✦ Most of what we store is derived data
- ✦ Switching to a model where derived data could be automatically re-derived has advantages
 - ✦ Reduces storage costs and dramatically improves data preservation
- ✦ But it requires among other things the move to peak resource provisioning
- ✦ Achieving it would be transformative

Transformations

- ✦ Transformations don't just happen, they happen because of big investments of effort
 - ✦ Even Linux was a significant code port, a change in mentality, and a cost of operations
 - ✦ Grids were a decade long development and commissioning activity and cost millions
 - ✦ Networking and Connectivity were driven by us
 - ✦ A Data Intensive Content Delivery System and moving to specialized systems achieving virtual data storage will not be easy or cheap but will fundamentally change how we work