Clouds will win!

FutureGrid Tutorial at PPAM 2011

Torun Poland

September 11 2011

Geoffrey Fox

gcf@indiana.edu

http://www.infomall.org http://www.futuregrid.org

Director, Digital Science Center, Pervasive Technology Institute

Associate Dean for Research and Graduate Studies, School of Informatics and Computing

Indiana University Bloomington

Important Trends

- Data Deluge in all fields of science
- Multicore implies parallel computing important again
  - Performance from extra cores – not extra clock speed
  - GPU enhanced systems can give big power boost
- Clouds – new commercially supported data center model replacing compute grids (and your general purpose computer center)
- Light weight clients: Sensors, Smartphones and tablets accessing and supported by backend services in cloud
- Commercial efforts moving much faster than academia in both innovation and deployment

Big Data in Many Domains

- According to one estimate, we created 150 exabytes (billion gigabytes) of data in 2005. In 2010, we created 1,200 exabytes
- Enterprise Storage sold in 2010 was 15 Exabytes; BUT total storage sold (including flash memory etc.) was 1500 Exabytes
- Size of the web ~ 3 billion web pages: MapReduce at Google was on average processing 20PB per day in January 2008
- During 2009, American drone aircraft flying over Iraq and Afghanistan sent back around 24 years' worth of video footage
  - http://www.economist.com/node/15579717
  - New models being deployed in 2010 will produce ten times as many data streams as their predecessors, and those in 2011 will produce 30 times as many
- ~108 million sequence records in GenBank in 2009, doubling in every 18 months
- ~20 million purchases at Wal-Mart a day
- 90 million Tweets a day
- Astronomy, Particle Physics, Medical Records …
- Most scientific task shows CPU:IO ratio of 10000:1 – Dr. Jim Gray
- *The Fourth Paradigm: Data-Intensive Scientific Discovery*
- *Large Hadron Collider at CERN; 100 Petabytes to find Higgs Boson*

Cost per Megabase of DNA Sequence

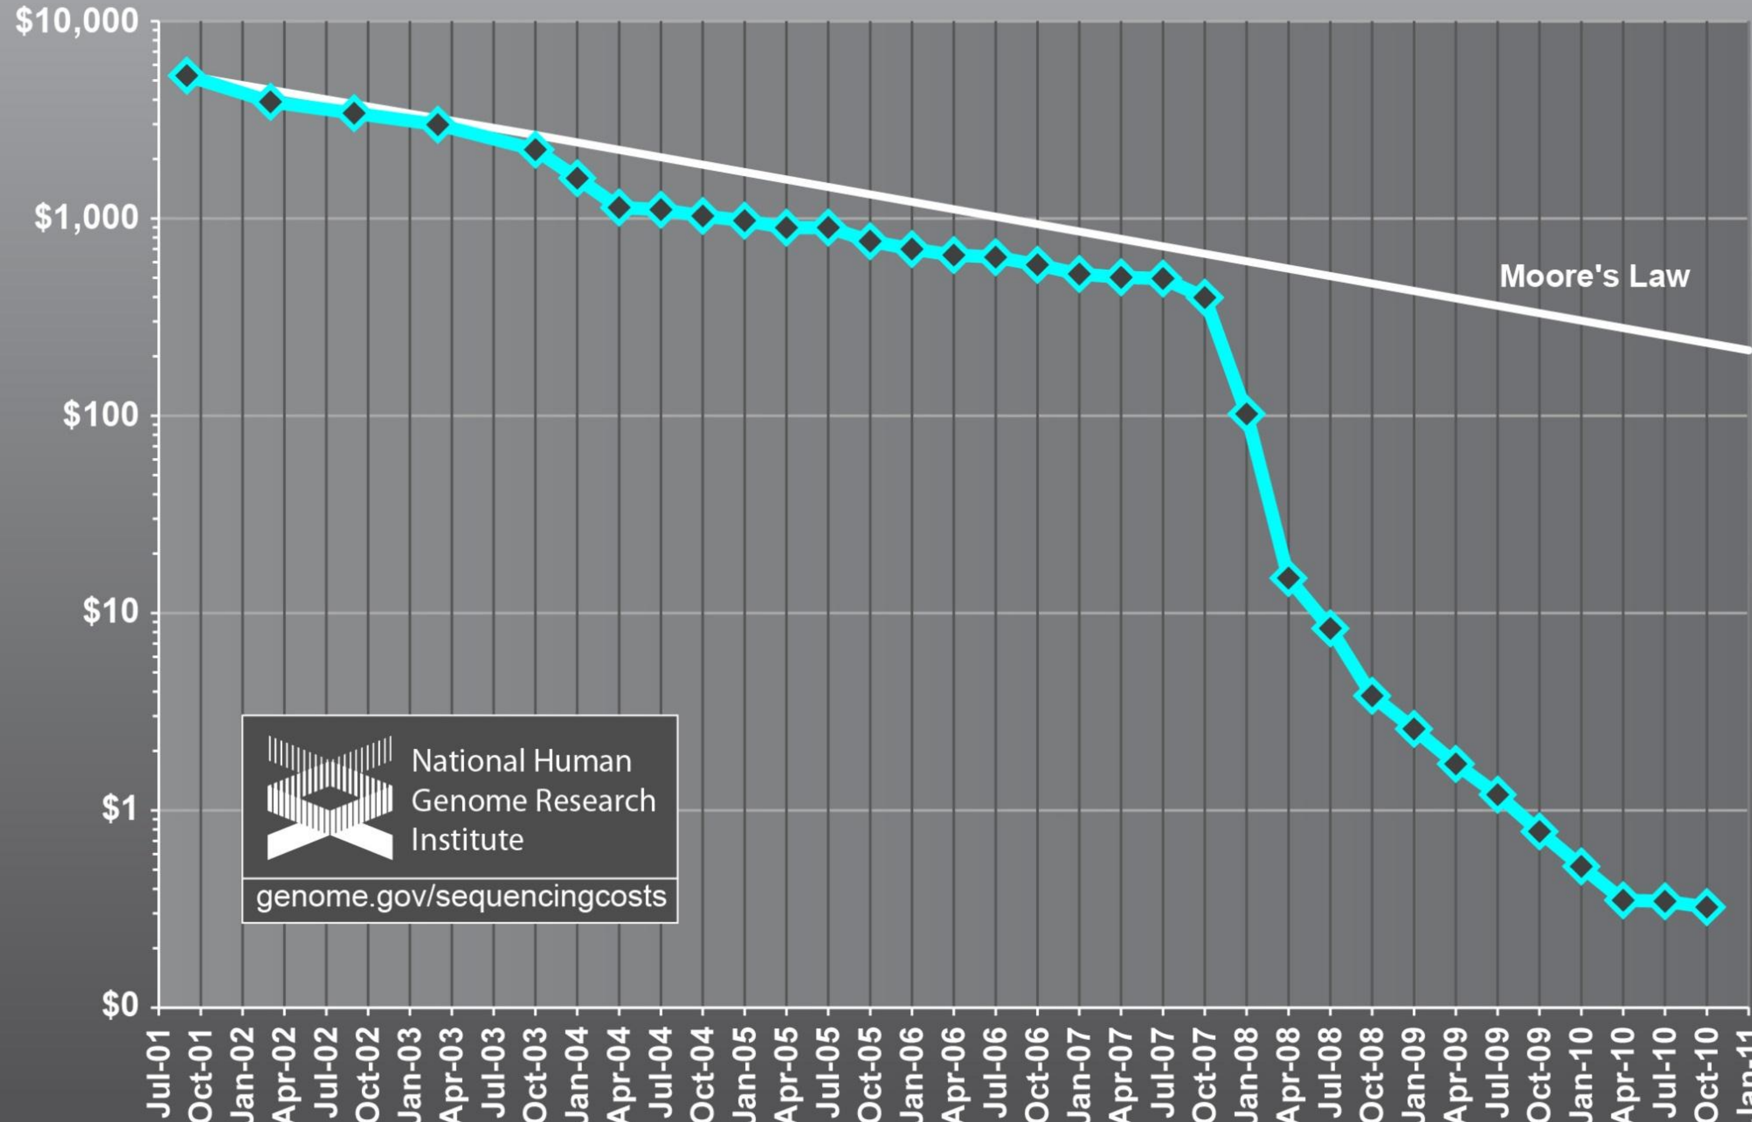

Data Centers Clouds & Economies of Scale I

Range in size from "edge" facilities to megascale.

Economies of scale

Approximate costs for a small size center (1K servers) and a larger, 50K server center.

2 Google warehouses of computers on the banks of the Columbia River, in The Dalles, Oregon

Such centers use 20MW-200MW (Future) each with 150 watts per CPU

Save money from large size, positioning with cheap power and access with Internet

Data Centers, Clouds & Economies of Scale II

- Builds giant data centers with 100,000's of computers; ~ 200-1000 to a shipping container with Internet access
- "Microsoft will cram between 150 and 220 shipping containers filled with data center gear into a new 500,000 square foot Chicago facility. This move marks the most significant, public use of the shipping container systems popularized by the likes of Sun Microsystems and Rackable Systems to date."

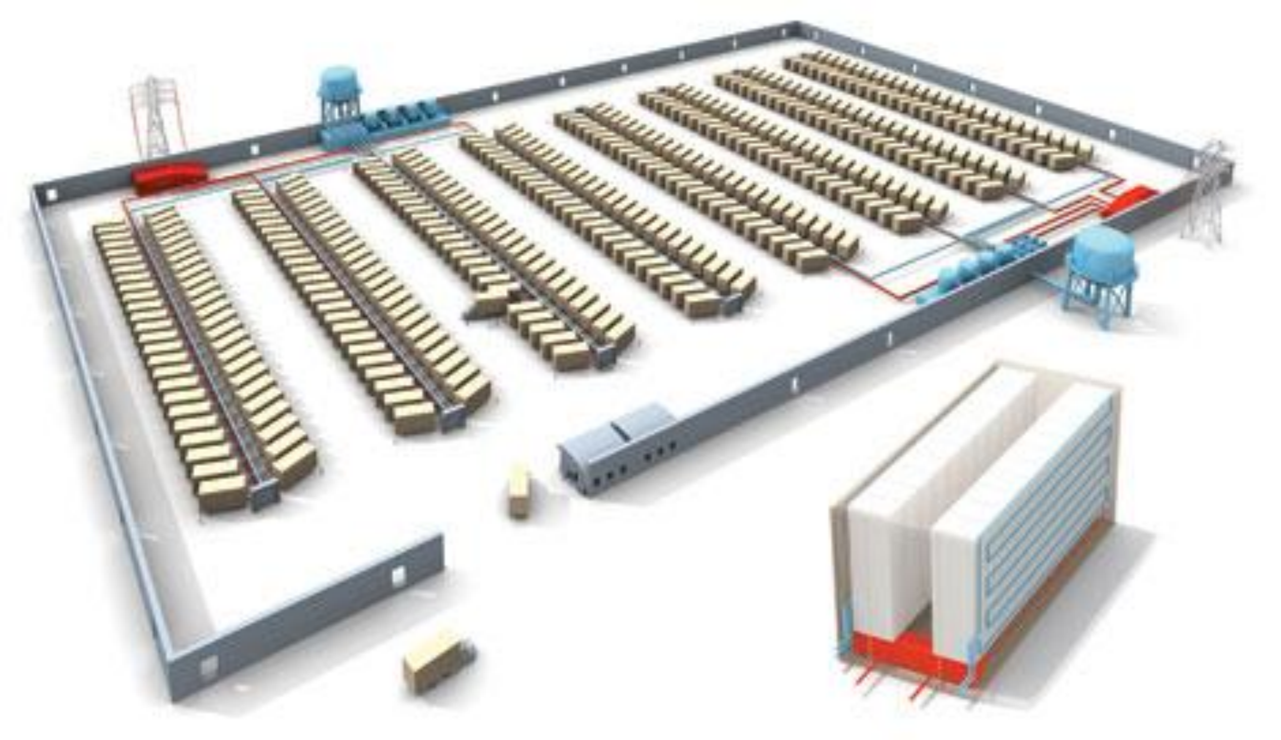

X as a Service

- SaaS: Software as a Service imply software capabilities (programs) have a service (messaging) interface
  - Applying systematically reduces system complexity to being linear in number of components
  - Access via messaging rather than by installing in /usr/bin
- IaaS: Infrastructure as a Service or HaaS: Hardware as a Service – get your computer time with a credit card and with a Web interface
- PaaS: Platform as a Service is IaaS plus core software capabilities on which you build SaaS
- Cyberinfrastructure is "Research as a Service"

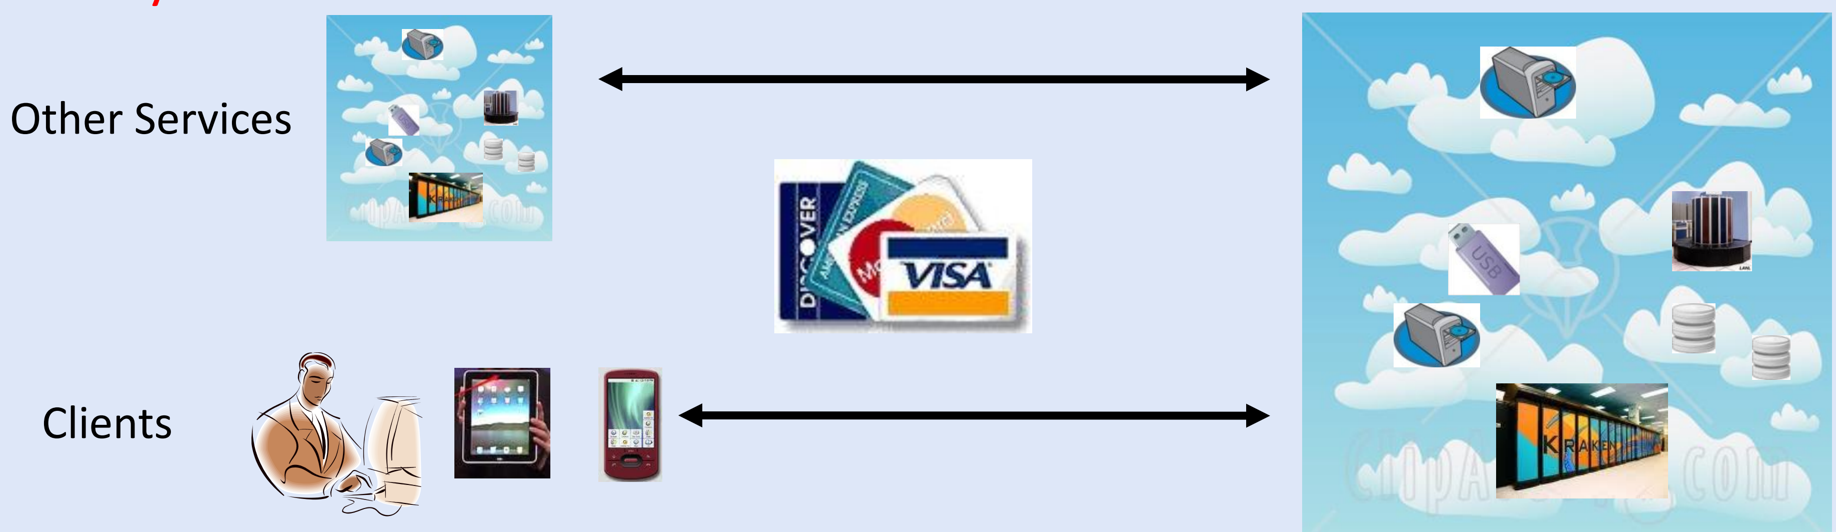

Sensors as a Service

Cell phones are important sensor

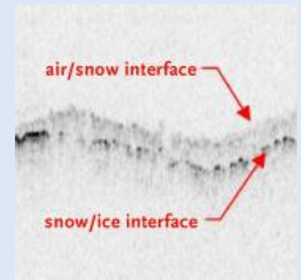

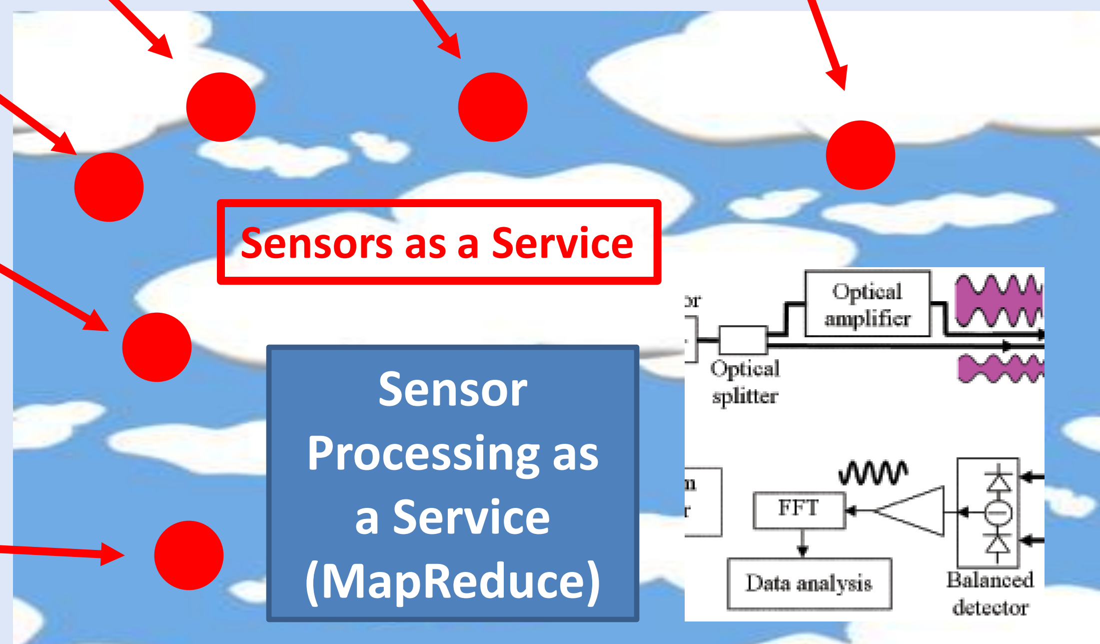

Clouds and Jobs

- Clouds are a major industry thrust with a growing fraction of IT expenditure that IDC estimates will grow to $44.2 billion direct investment in 2013 while 15% of IT investment in 2011 will be related to cloud systems with a 30% growth in public sector.
- Gartner also rates cloud computing high on list of critical emerging technologies with for example “Cloud Computing” and “Cloud Web Platforms” rated as transformational (their highest rating for impact) in the next 2-5 years.
- Correspondingly there is and will continue to be major opportunities for new jobs in cloud computing with a recent European study estimating there will be 2.4 million new cloud computing jobs in Europe alone by 2015.
- Cloud computing is an attractive for projects focusing on workforce development. Note that the recently signed “America Competes Act” calls out the importance of economic development in broader impact of NSF projects

benefit	years to mainstream adoption			
	less than 2 years	2 to 5 years	5 to 10 years	more than 10 years
Transformational		**Cloud Computing** **Cloud Web Platforms** **Media Tablet**	3D Printing Context Delivery Architecture Extreme Transaction Processing	Autonomous Vehicles Human Augmentation Mobile Robots Terahertz Waves
High	Mobile Application Stores Predictive Analytics	Activity Streams E-Book Readers Electronic Paper Interactive TV Internet Micropayment Systems Location-Aware Applications Private Cloud Computing Social Analytics	Augmented Reality Internet TV Virtual Assistants Wireless Power	Mesh Networks: Sensor
Moderate	Consumer-Generated Media Pen-Centric Tablet PCs	3D Flat-Panel TVs and Displays Biometric Authentication Methods Gesture Recognition Idea Management Microblogging Speech Recognition Video Telepresence	4G Standard Public Virtual Worlds Speech-to-Speech Translation Video Search	Computer-Brain Interface
Low				Tangible User Interfaces

As of August 2010

Philosophy of Clouds and Grids

- Clouds are (by definition) commercially supported approach to large scale computing
  - So we should expect Clouds to replace Compute Grids
  - Current Grid technology involves “non-commercial” software solutions which are hard to evolve/sustain
- Public Clouds are broadly accessible resources like Amazon and Microsoft Azure – powerful but not easy to optimize and perhaps data trust/privacy issues
- Private Clouds run similar software and mechanisms but on “your own computers”
- Services still are correct architecture with either REST (Web 2.0) or Web Services
- Clusters still critical concept

Grids MPI and Clouds

- **Grids** are useful for **managing distributed systems**
  - Pioneered service model for Science
  - Developed importance of Workflow
  - Performance issues – communication latency – intrinsic to distributed systems
  - Can never run large differential equation based simulations or datamining
- **Clouds can execute any job class that was good for Grids** plus
  - More attractive due to platform plus **elastic** on-demand model
  - **MapReduce easier to use than MPI for appropriate parallel jobs**
  - Currently have performance limitations due to poor affinity (locality) for compute-compute (MPI) and Compute-data
  - These limitations are not "inevitable" and should gradually improve as in July 13 2010 Amazon Cluster announcement
  - Will probably never be best for most sophisticated parallel differential equation based simulations
- **Classic Supercomputers** (MPI Engines) run **communication demanding differential equation based simulations**
  - **MapReduce and Clouds replaces MPI** for other problems
  - Much more data processed today by MapReduce than MPI (Industry Informational Retrieval ~50 Petabytes per day)

Fault Tolerance and MapReduce

- MPI does “maps” followed by “communication” including “reduce” but does this iteratively
- There must (for most communication patterns of interest) be a strict synchronization at end of each communication phase
  - Thus if a process fails then everything grinds to a halt
- In MapReduce, all Map processes and all reduce processes are independent and stateless and read and write to disks
  - As 1 or 2 (reduce+map) iterations, no difficult synchronization issues
- Thus failures can easily be recovered by rerunning process without other jobs hanging around waiting
- Re-examine MPI fault tolerance in light of MapReduce
  - Twister will interpolate between MPI and MapReduce

Authentication and Authorization: Provide single sign in to both FutureGrid and Commercial Clouds linked by workflow
Workflow: Support workflows that link job components between FutureGrid and Commercial Clouds. Trident from Microsoft Research is initial candidate
Data Transport: Transport data between job components on FutureGrid and Commercial Clouds respecting custom storage patterns
Program Library: Store Images and other Program material (basic FutureGrid facility)
Blob: Basic storage concept similar to Azure Blob or Amazon S3
DPFS Data Parallel File System: Support of file systems like Google (MapReduce), HDFS (Hadoop) or Cosmos (dryad) with compute-data affinity optimized for data processing
Table: Support of Table Data structures modeled on Apache Hbase/CouchDB or Amazon SimpleDB/Azure Table. There is “Big” and “Little” tables – generally NOSQL
SQL: Relational Database
Queues: Publish Subscribe based queuing system
Worker Role: This concept is implicitly used in both Amazon and TeraGrid but was first introduced as a high level construct by Azure
MapReduce: Support MapReduce Programming model including Hadoop on Linux, Dryad on Windows HPCS and Twister on Windows and Linux
Software as a Service: This concept is shared between Clouds and Grids and can be supported without special attention
Web Role: This is used in Azure to describe important link to user and can be supported in FutureGrid with a Portal framework

Amazon offers a lot!

AWS | Products | Developers | Community | Support | Account

Compute
- Amazon Elastic Compute Cloud (EC2)
- Amazon Elastic MapReduce
- Auto Scaling

Content Delivery
- Amazon CloudFront

Database
- Amazon SimpleDB
- Amazon Relational Database Service (RDS)

E-Commerce
- Amazon Fulfillment Web Service (FWS)

Messaging
- Amazon Simple Queue Service (SQS)
- Amazon Simple Notification Service (SNS)

Monitoring
- Amazon CloudWatch

Networking
- Amazon Virtual Private Cloud (VPC)
- Elastic Load Balancing

Payments & Billing
- Amazon Flexible Payments Service (FPS)
- Amazon DevPay

Storage
- Amazon Simple Storage Service (S3)
- Amazon Elastic Block Storage (EBS)
- AWS Import/Export

Support
- AWS Premium Support

Web Traffic
- Alexa Web Information Service
- Alexa Top Sites

Workforce
- Amazon Mechanical Turk

close

Products & Services

Amazon Elastic Compute Cloud (Amazon EC2)

Amazon EC2 Details
- EC2 Overview

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.

Sign Up For Amazon EC2

Traditional File System?

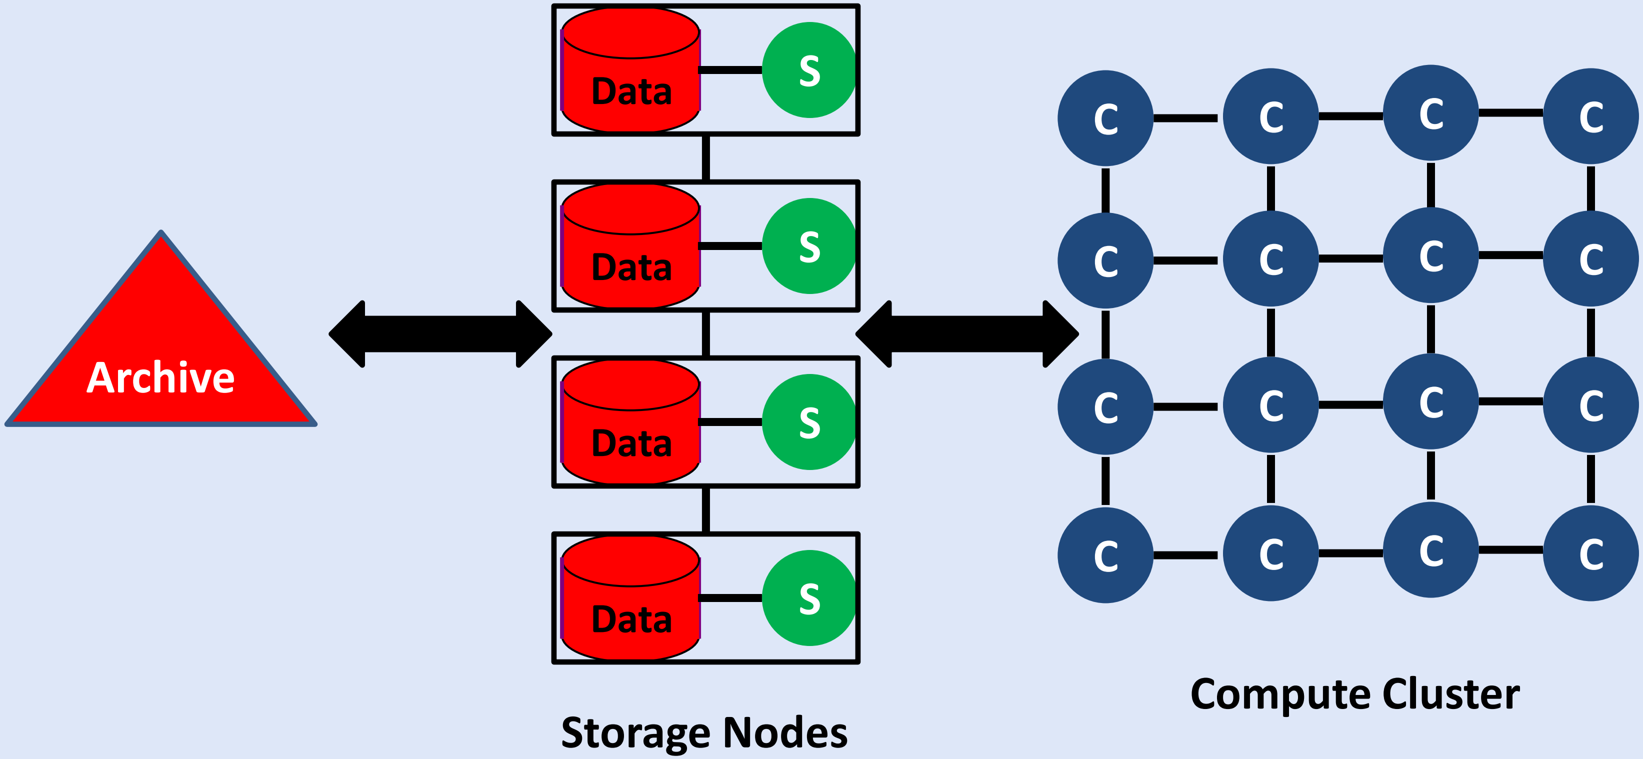

- Typically a shared file system (Lustre, NFS …) used to support high performance computing
- Big advantages in flexible computing on shared data but doesn't "**bring computing to data**"

Data Parallel File System?

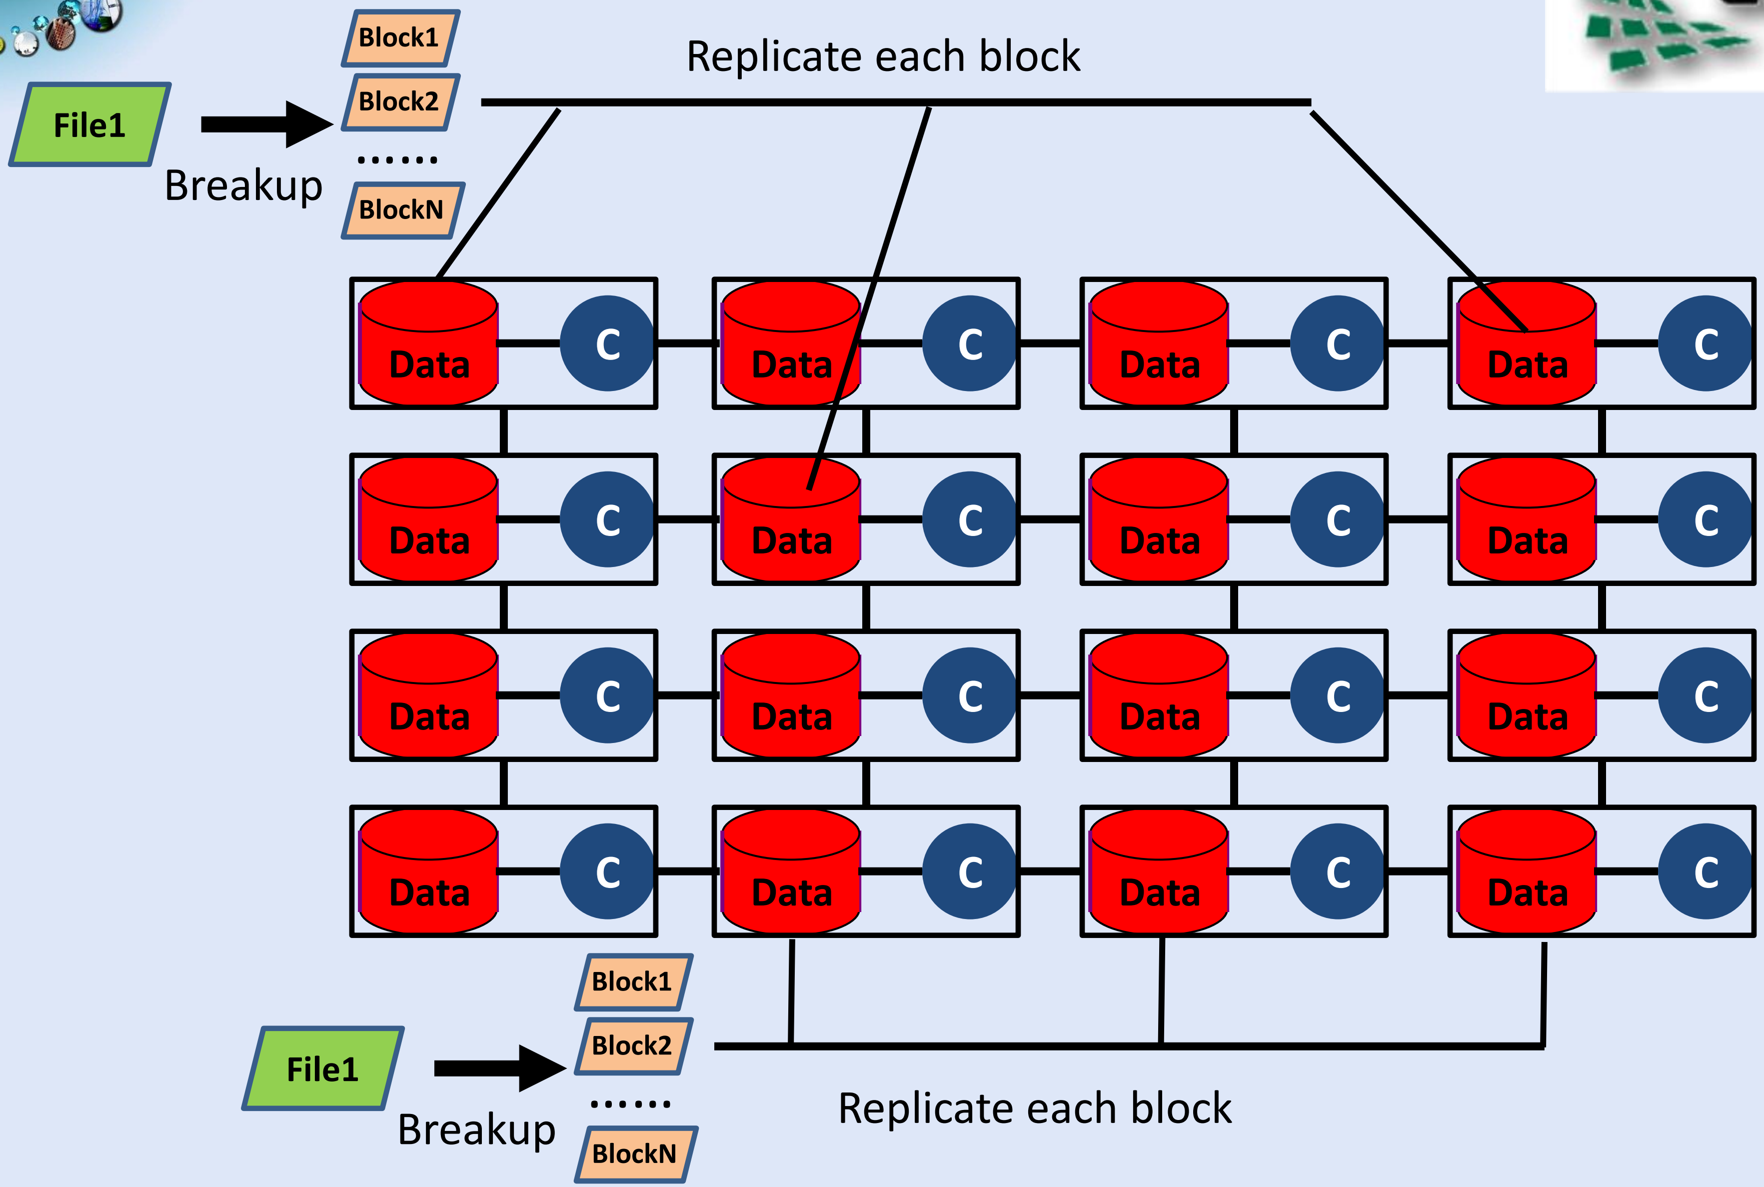

- No archival storage and computing brought to data

Important Platform Capability MapReduce

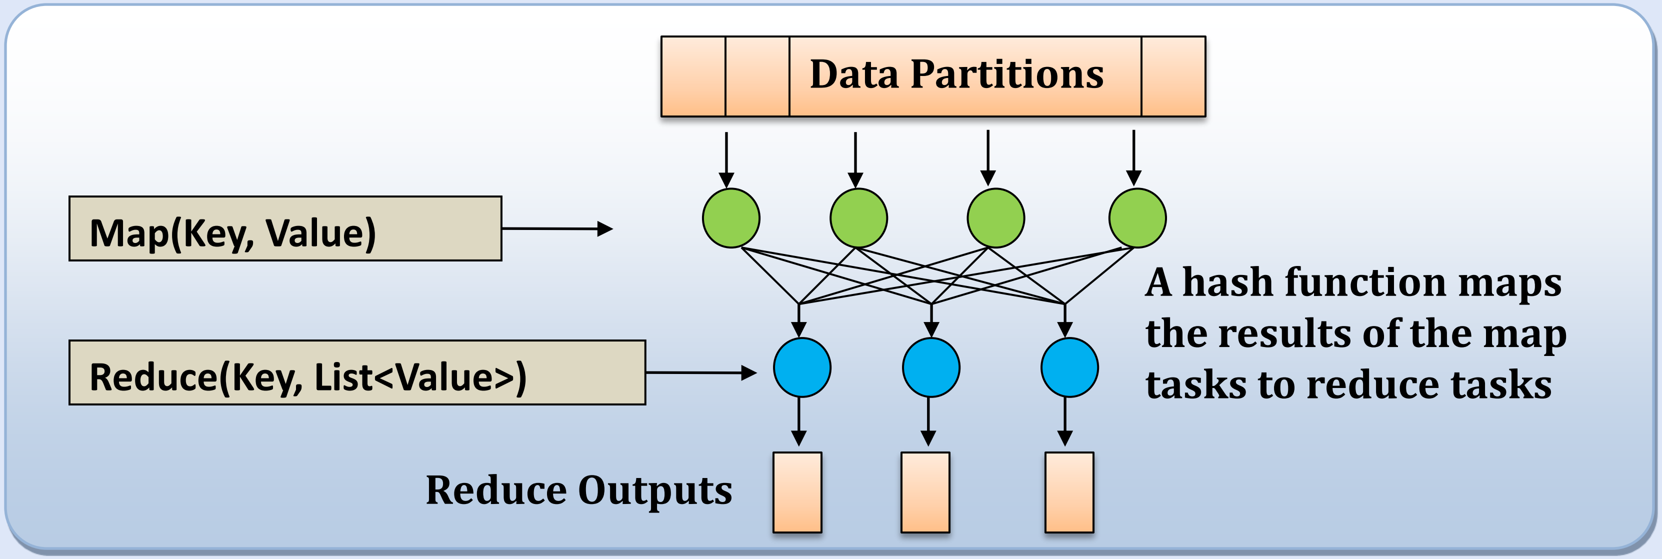

- Implementations (Hadoop – Java; Dryad – Windows) support:
  - Splitting of data
  - Passing the output of map functions to reduce functions
  - Sorting the inputs to the reduce function based on the intermediate keys
  - Quality of service