Microsoft

EXTREME COMPUTING GROUP

Defining the future.

Cloud Computing for Research

Roger Barga

Cloud Computing Futures, Microsoft Research

Trends: Data on an Exponential Scale

- Scientific data doubles every year
  - *Combination of inexpensive sensors + exponentially faster computing*
  - *Cost of acquiring data has dropped close to zero*
- Changes the nature of research computing
- Cuts across disciplines (eScience)
- Becoming increasingly harder to extract knowledge
- ***Got data, now what?***
- ***And it is really is about data***, not the FLOPS (going faster)…
  - Very extended distribution: data sets on all scales!
  - When data collection does grow large, not able to analyze.
  - Tools are limited, must dedicate resources to build analysis tools (and this doesn't help complete the actual research).

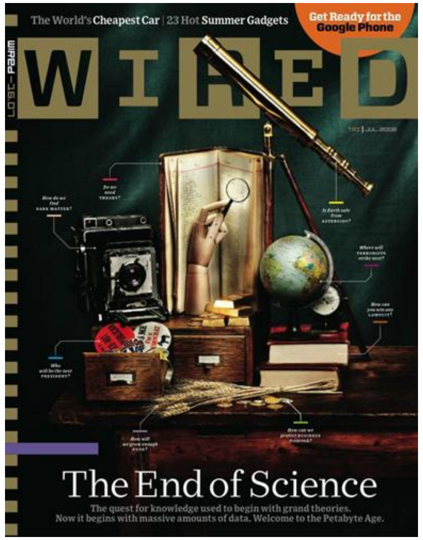

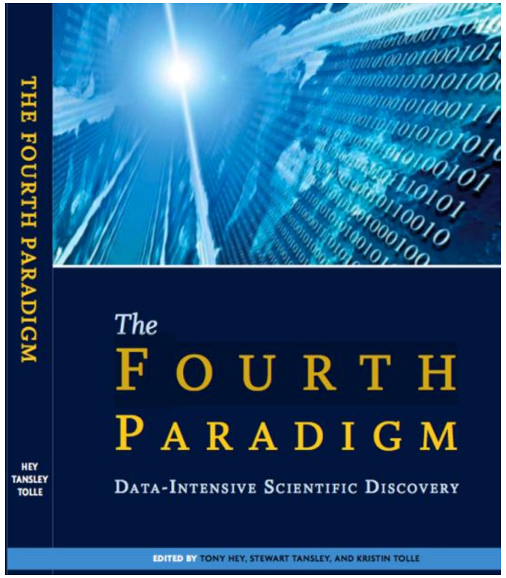

Trend: Lack of Broad Access

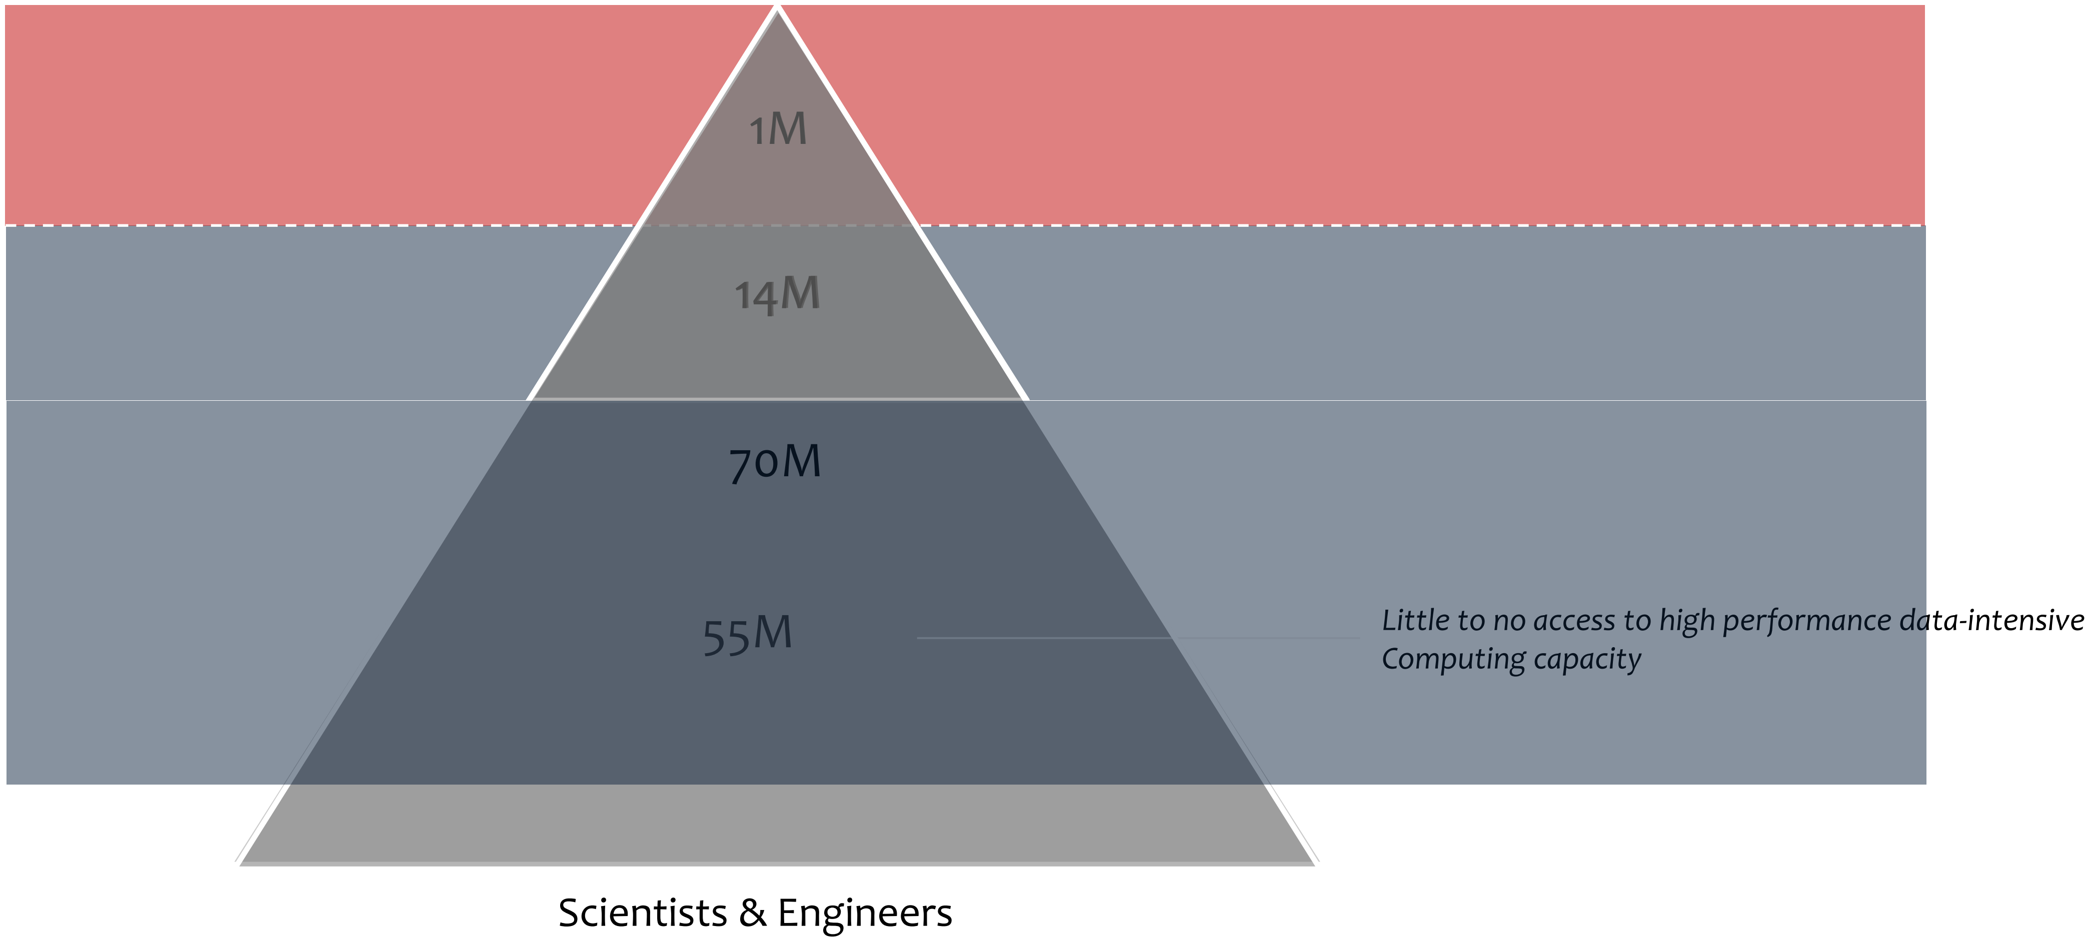

53% of technical organizations are forced to scale down their advanced problems 'to fit' within their technology limitations, and 57% of companies have problems they can't solve with existing computers. US Council on Competitiveness, 2010.

Widespread On Demand Access

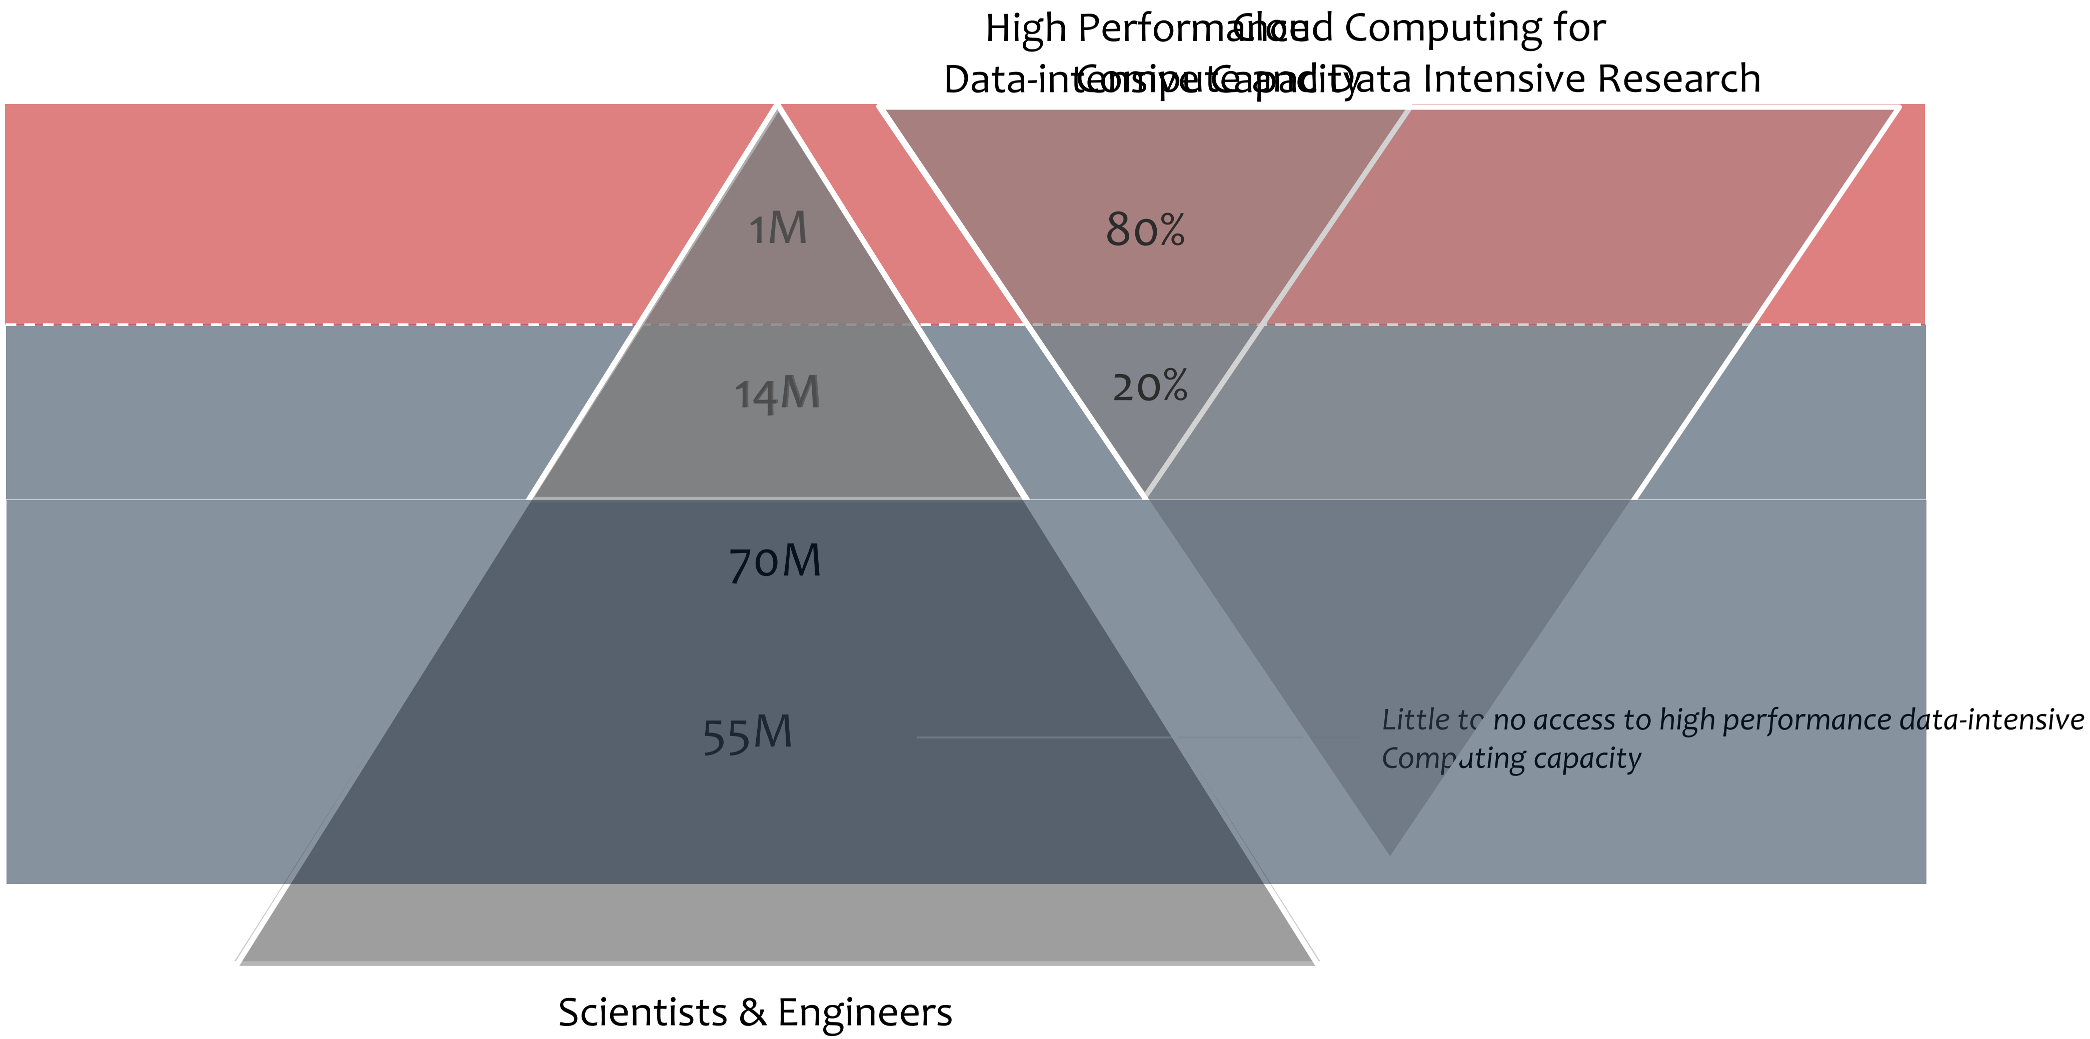

Bridging the Gap with the Cloud

With cloud computing, researchers can:

- Marshal needed storage and compute resources on demand without knowing or caring how it happens
- Access curated reference datasets, *in the cloud*.
- Run key algorithms as *Software as a Service* without having to know coding details or how to install software
- Collaborate and share data and algorithms.
- Use familiar client tools, leverage the cloud as a resource and intellectual amplifier.
- Offer new (higher value) data services for research
  - *Scalable analytics to explore data sets*
  - *Interactive data visualizations*
  - *Result provenance*

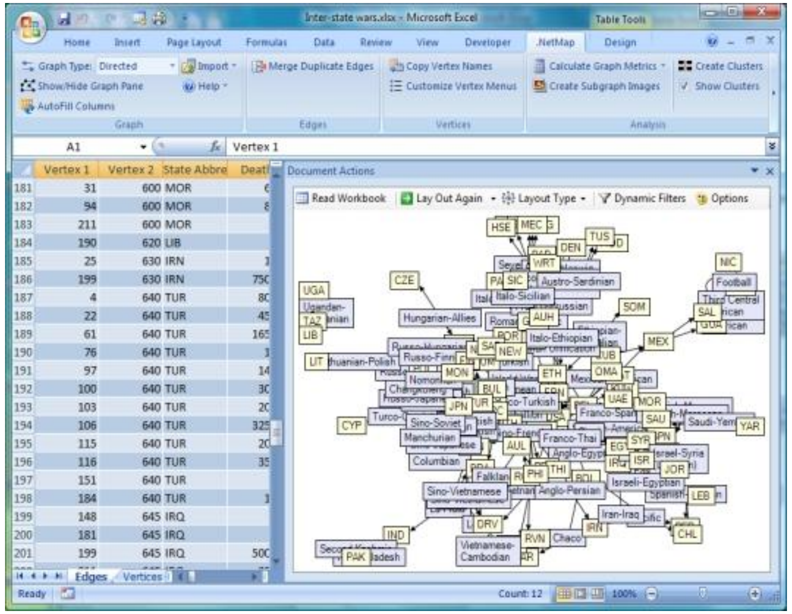

Suggested Application Model

Using queues for reliable messaging

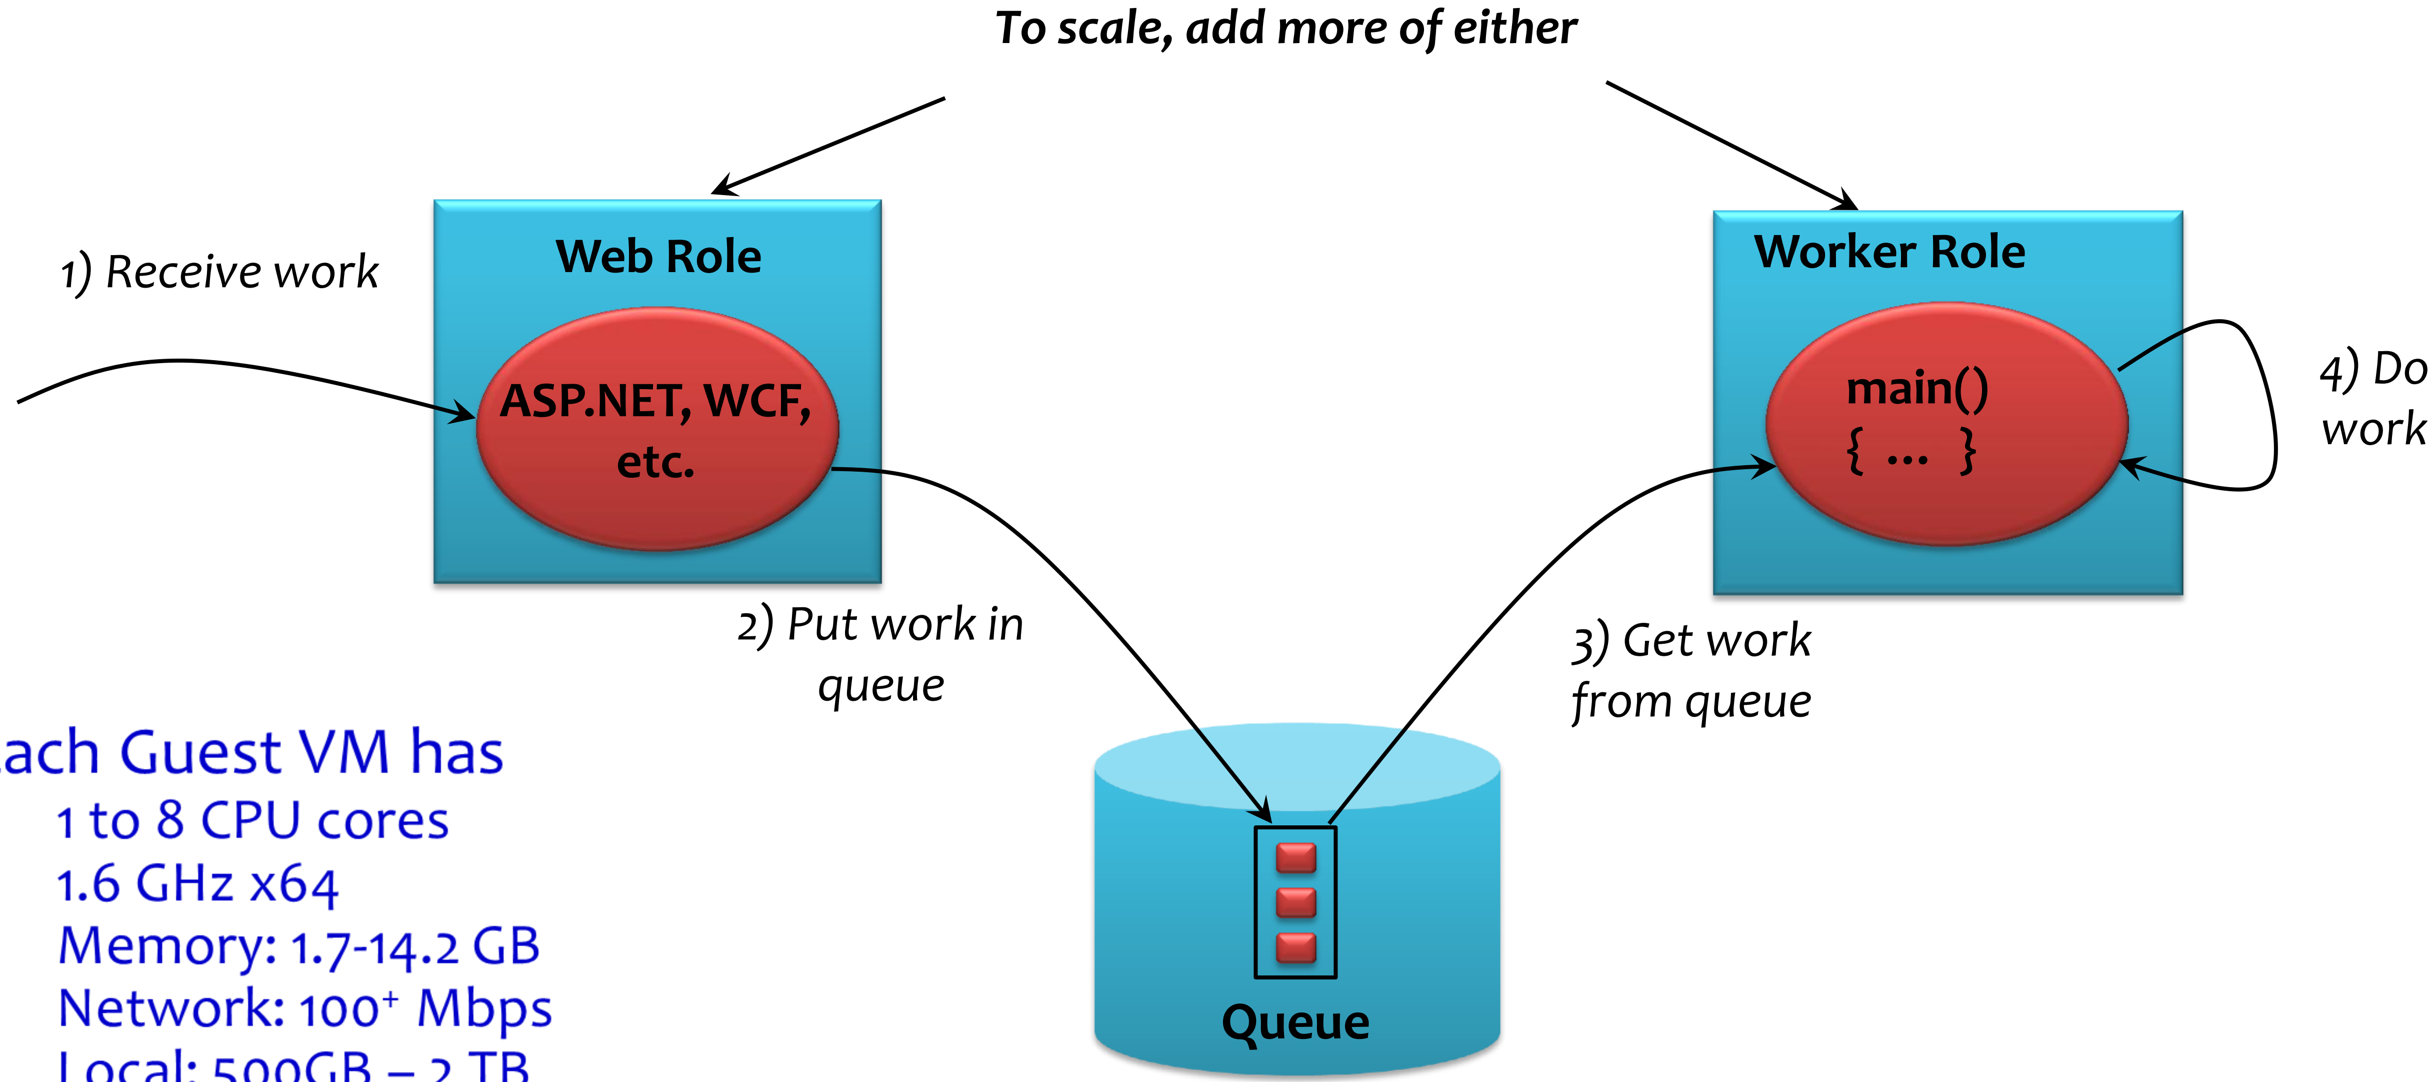

Each Guest VM has

- 1 to 8 CPU cores
- 1.6 GHz x64
- Memory: 1.7-14.2 GB
- Network: 100+ Mbps
- Local: 500GB – 2 TB

Storage Service in Windows Azure

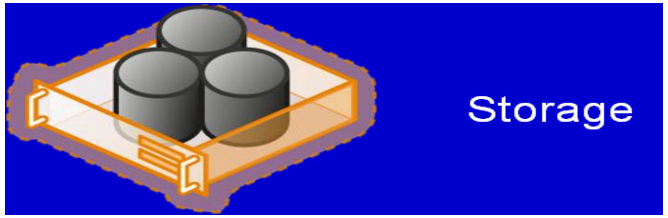

Azure applications can use **Azure storage services** or **SQL Azure**

SQL Azure

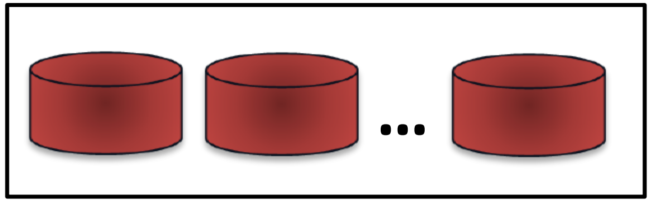

Blobs: large, unstructured data (audio, video, etc)

Tables: simply structured data, accessed using ADO.NET Data Services

Queues: serially accessed messages or requests, allowing web-roles and worker-roles to interact

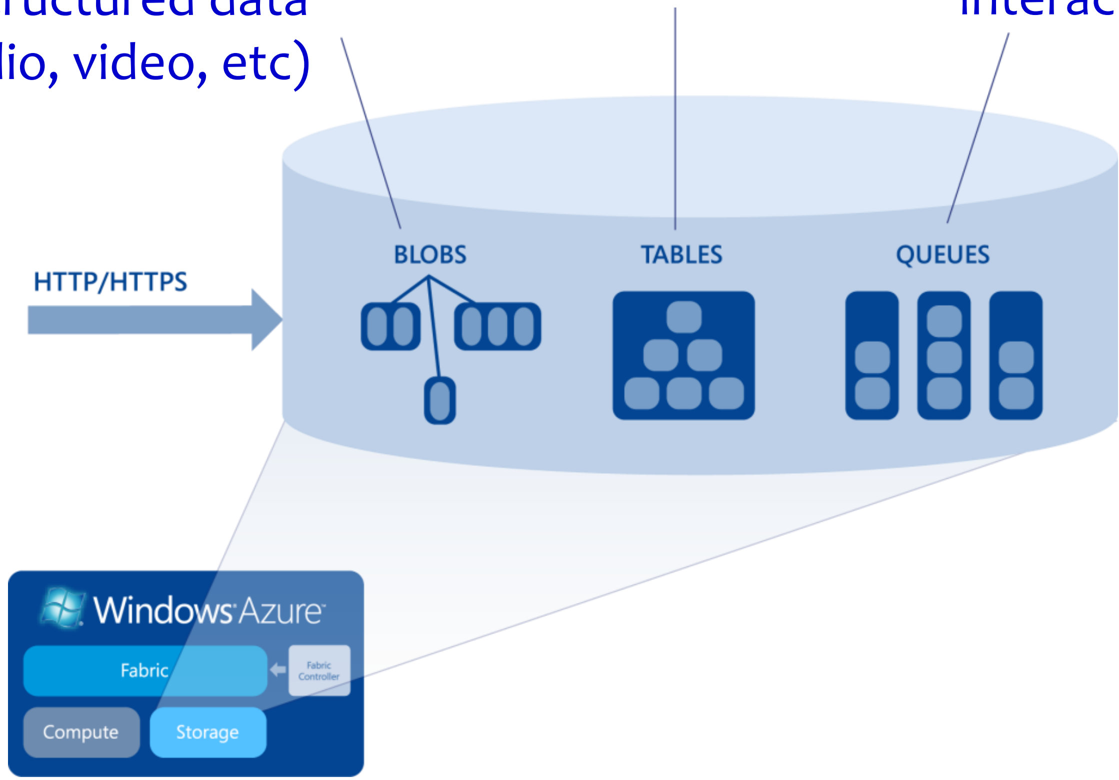

Modern Data Center Network

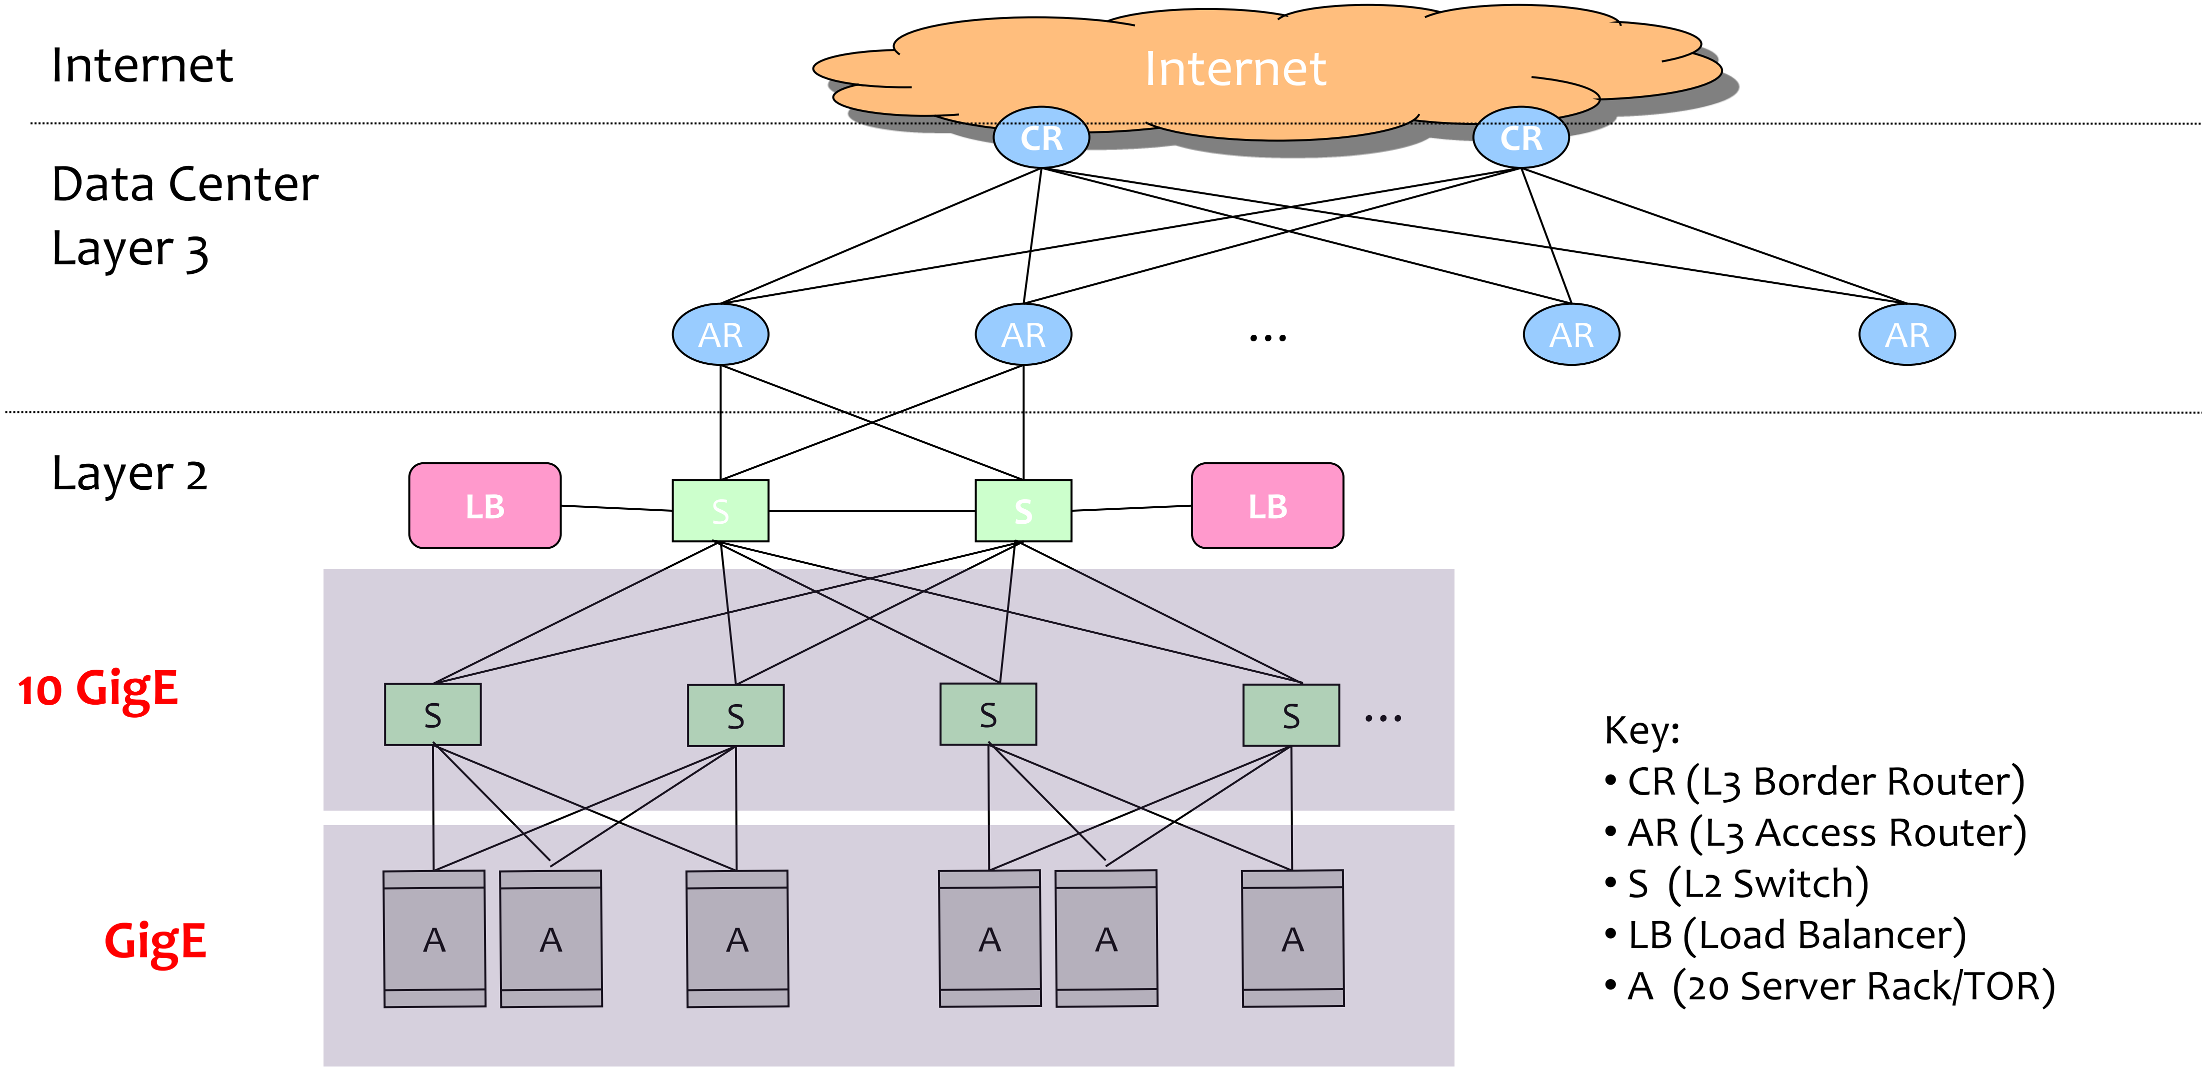

Key:
- CR (L3 Border Router)
- AR (L3 Access Router)
- S (L2 Switch)
- LB (Load Balancer)
- A (20 Server Rack/TOR)

Source: Albert Greenberg and Cisco

Worker Role TCP-Endpoints

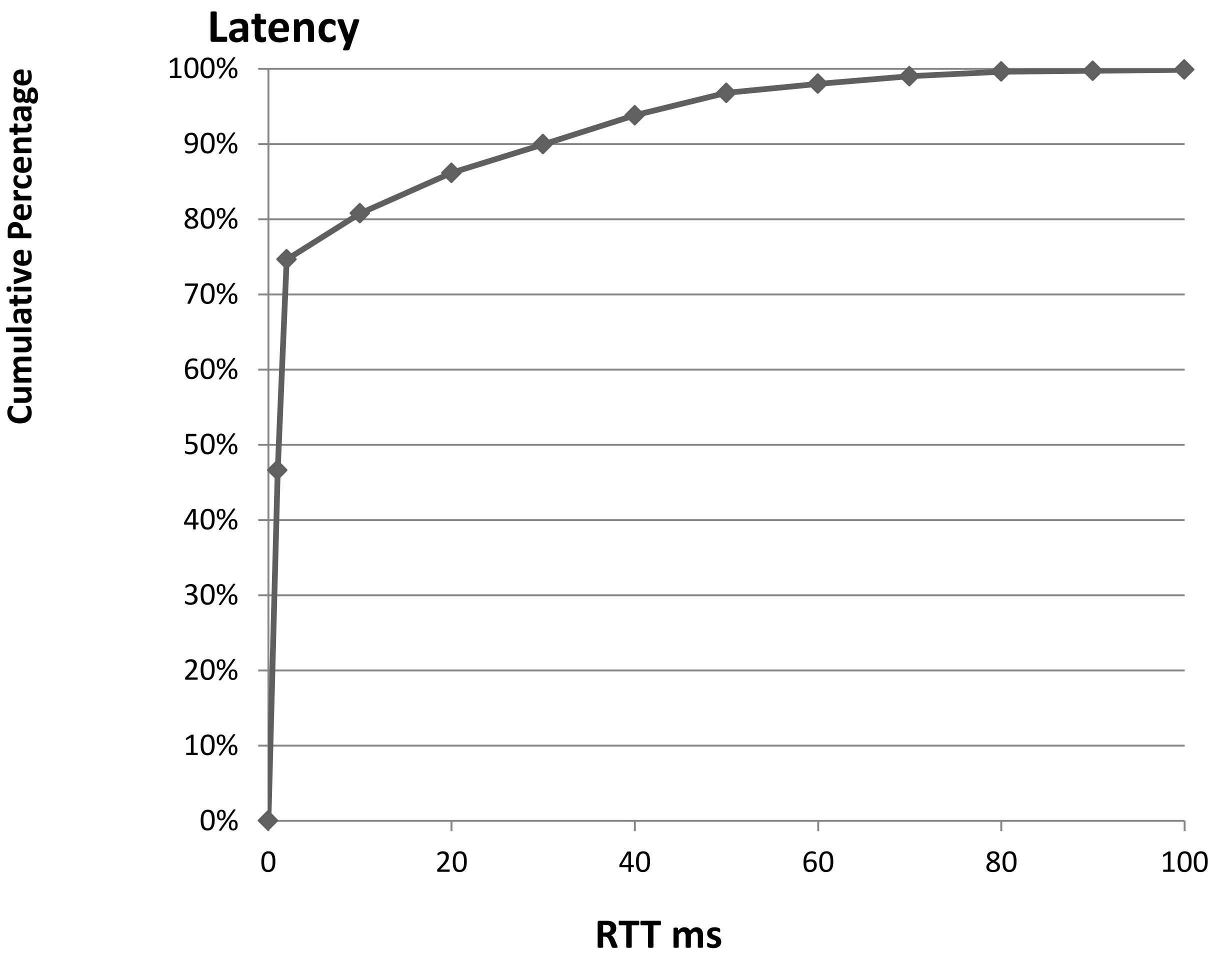

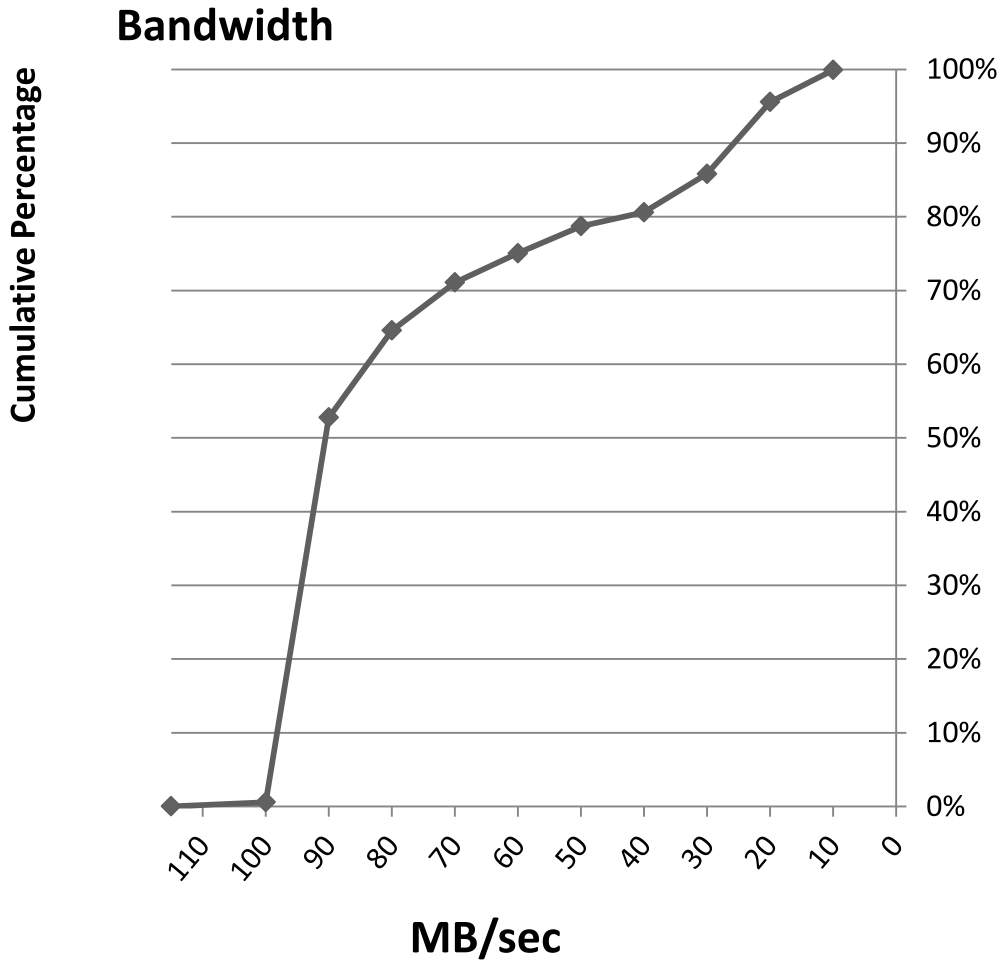

Blank Page - Windows Internet Explorer

http://azuroscope.cloudapp.net

EXTREME COMPUTING GROUP

Defining the future.

Microsoft

http://azurescope.cloudapp.net/

> Azure service, Silverlight frontend, data archived in SQL Azure
> Microbenchmarks, algorithm benchmarks, best practices
> Includes utility to measure local to cloud throughput (**download**)
> Run weekly on datacenters across the world
> New version(s) in development pipeline

Your input welcome…

Demo

Done | Internet | Protected Mode: Off | 100%

Demo

AzureBLAST

Bioinformatics in the Cloud

- Most important app in bioinformatics, exponential data growth
  - Research groups are unable to secure sufficient compute resources
- External research group, blasted ~5K proteins (700K sequences)
  - Normal BLAST job, between 700 – ~1000 CPU hours, couldn't run on NCBI
  - 3 day run was reduced to 30 min, a 30 minutes reduced to 30 sec
  - Total cost $100, publishable result after one hour of cloud computing

  "Reduced a weeks work of research down to one afternoon of computing"
- All-against-all non-redundant protein database (10m sequences)
  - Theoretically 100 billion sequence comparisons, 6 years on 8 core node
  - One of the largest BLAST jobs completed to date;
  - 1.85 billion results

The $24,000 Slide Deck

4000 cores (500 XL VMs), running for over 7 days

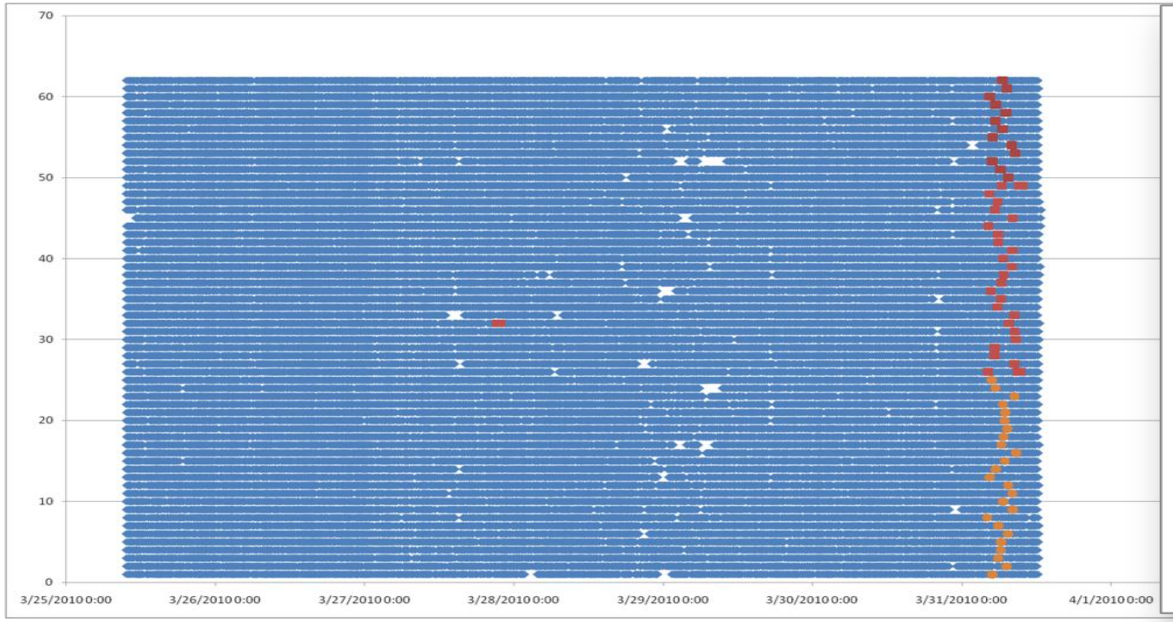

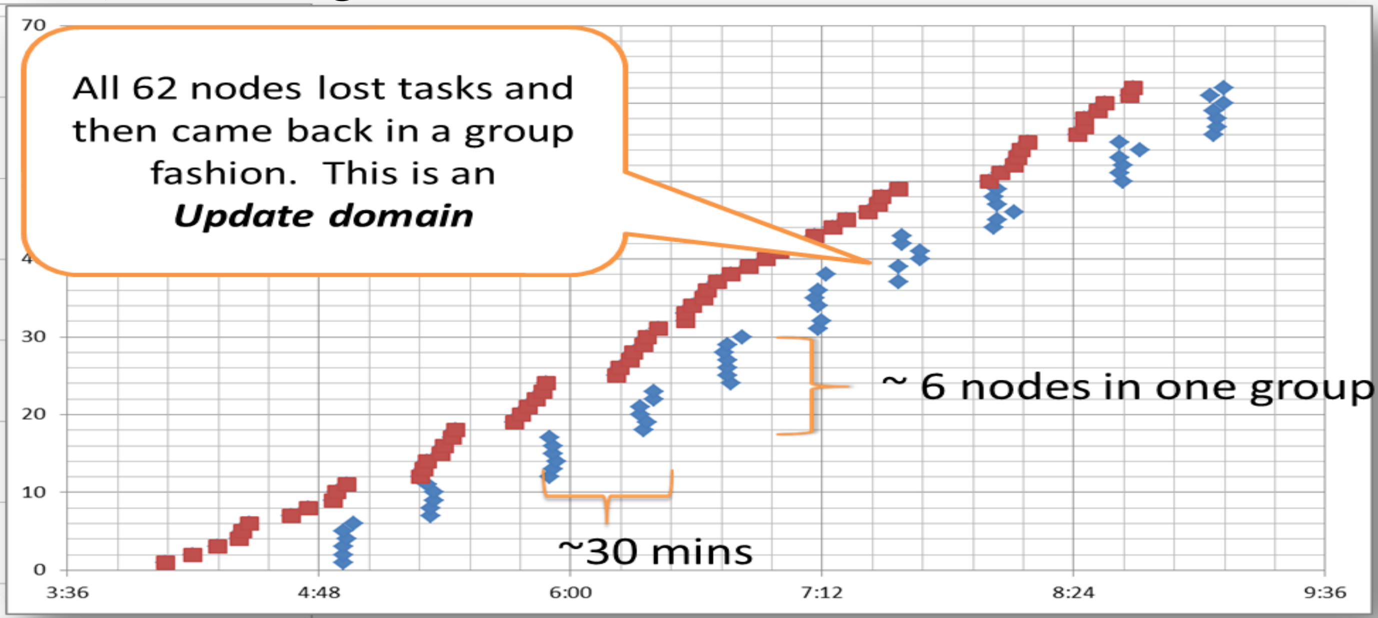

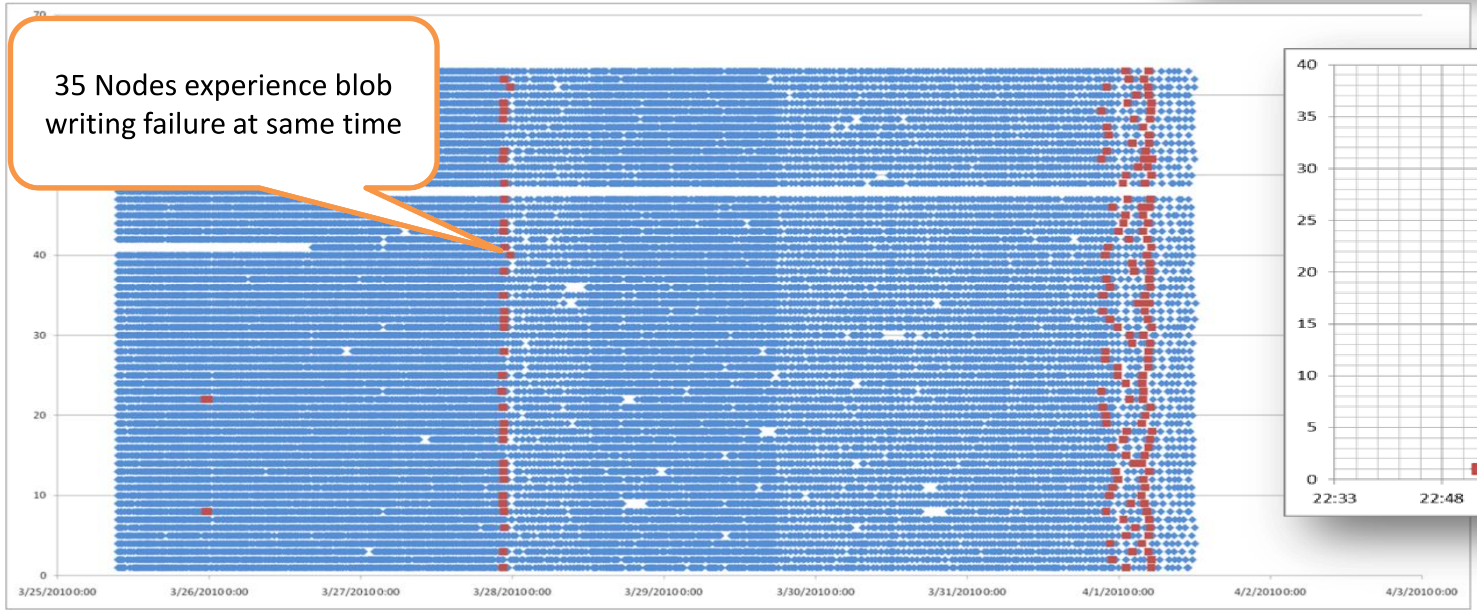

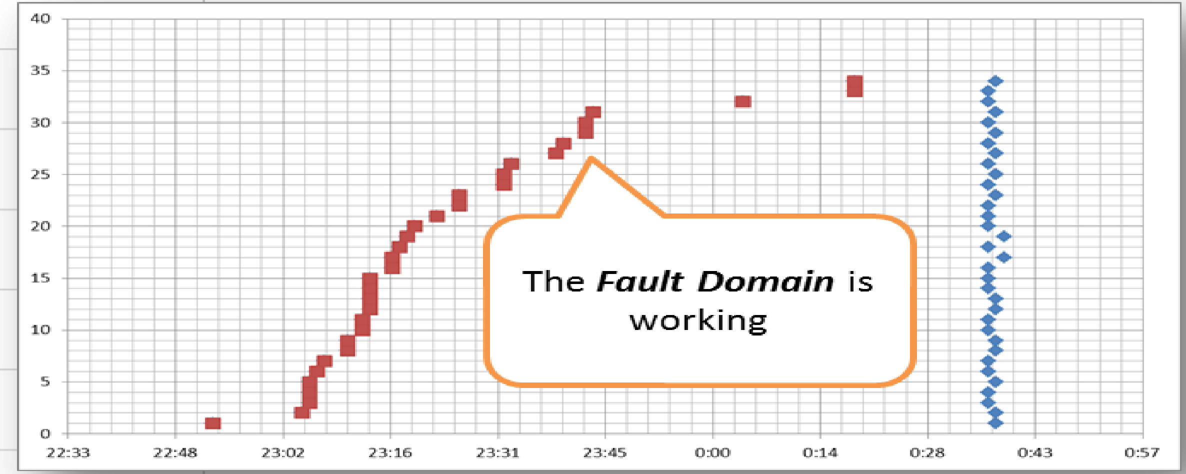

Application Benchmarks Inform Design

Task size vs. Performance

- Benefit of the warm cache effect
- 100 sequences per partition is best choice

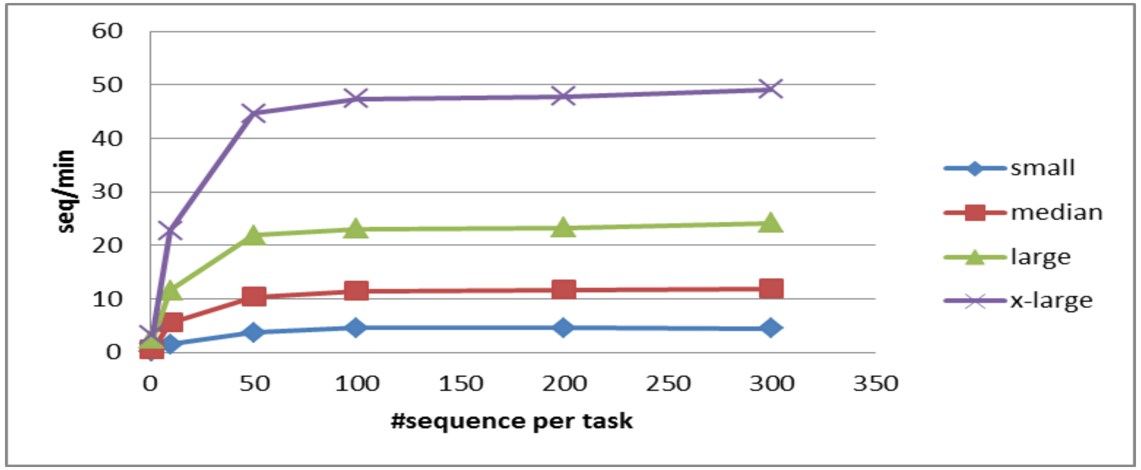

Instance size vs. Performance

- ***Super-linear*** speedup with larger size worker instances
- Primarily due to the memory capability.

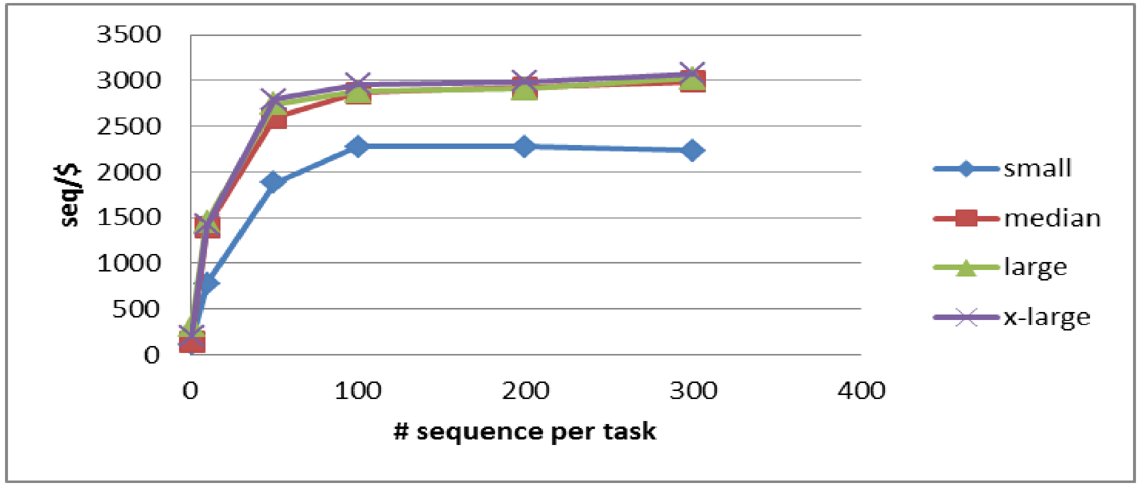

Task Size/Instance Size vs. Cost

- Extra-large instance generated the best and the most economical throughput
- Fully utilize the resource

AzureMODIS – *Computing Evapotranspiration (ET) in the Cloud*

Project with UC Berkeley, Univ Virginia, Indiana University, and LBNL

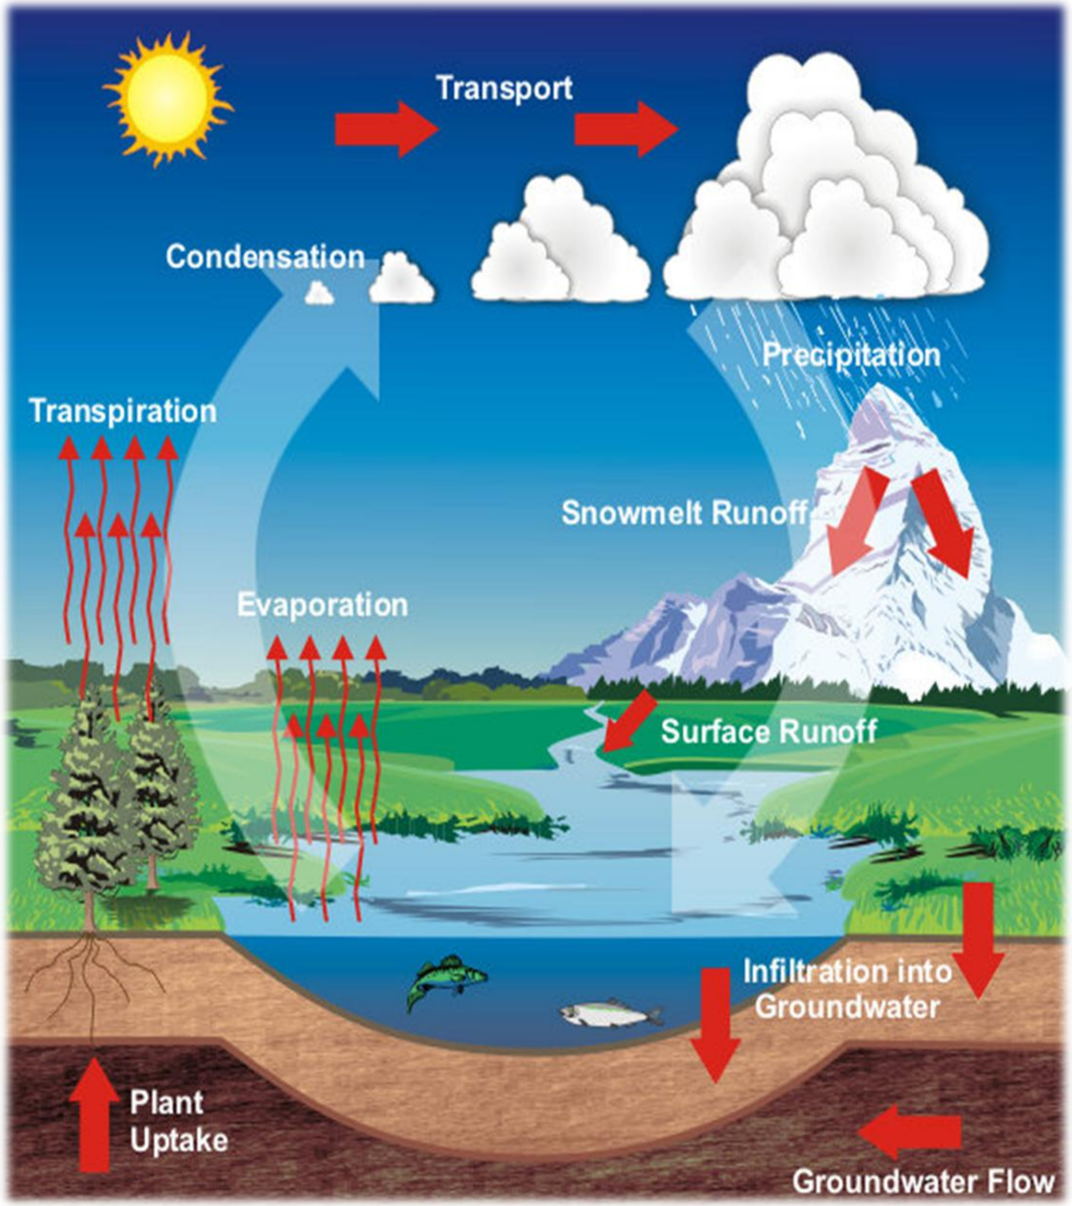

Two MODIS satellites

- Terra, launched 12/1999
- Aqua, launched 05/2002
- Near polar orbits
- Global coverage two days
- Sensitive in 36 spectral bands

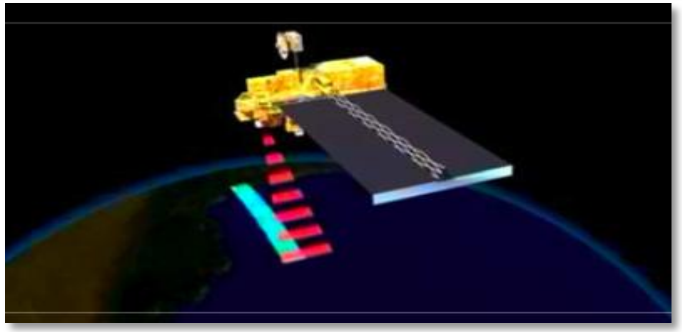

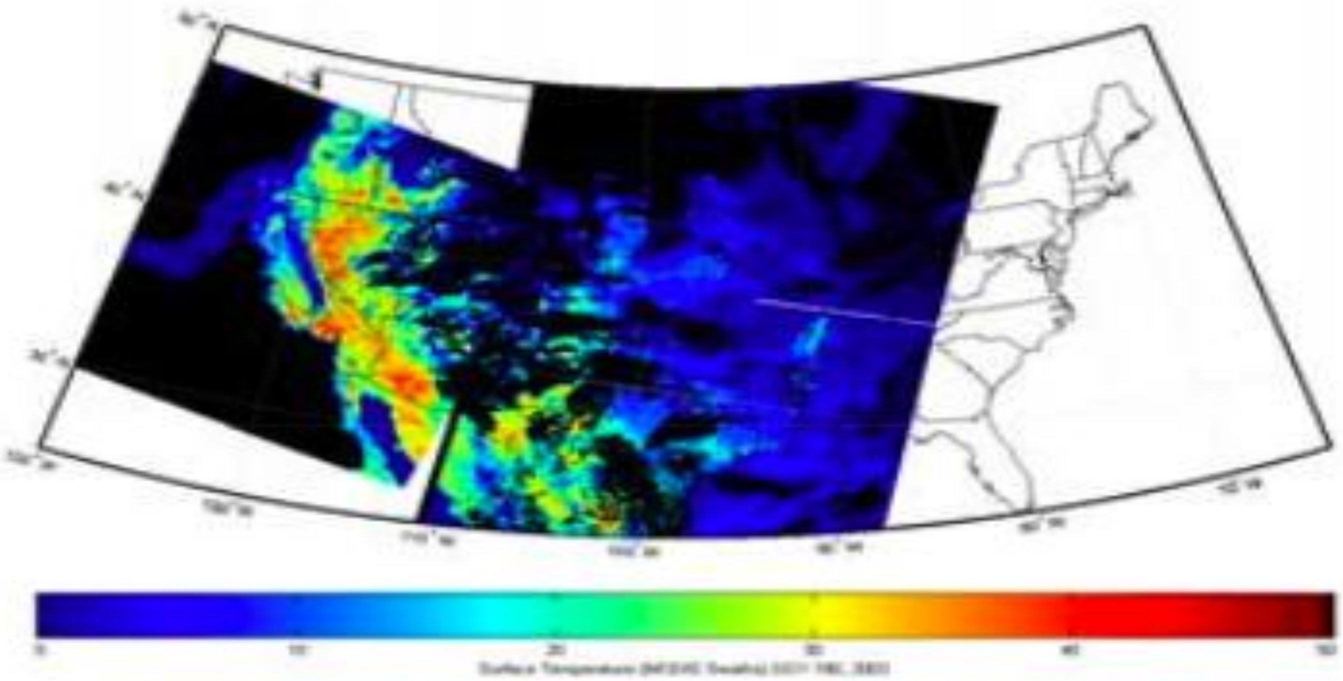

AzureMODIS – *Four stage image processing pipeline*

http://research.microsoft.com/en-us/projects/azure/azuremodis.aspx

Pipeline for download, processing, and reduction of satellite imagery.

- ~35 different data products;
- Atmospheric and land products are in different projections;
  - Re-projection, spatial temporal resolve.
  - Integrate data from different swaths, days,…
- 5 TB data processed (600,000 files)
- 35,000 hours reprojection stage
- 12,000 hours derivation reduction stage
- 3,000 hours analysis state
- Total cost for computing **one US year ET** computation: ~ **$2000**.

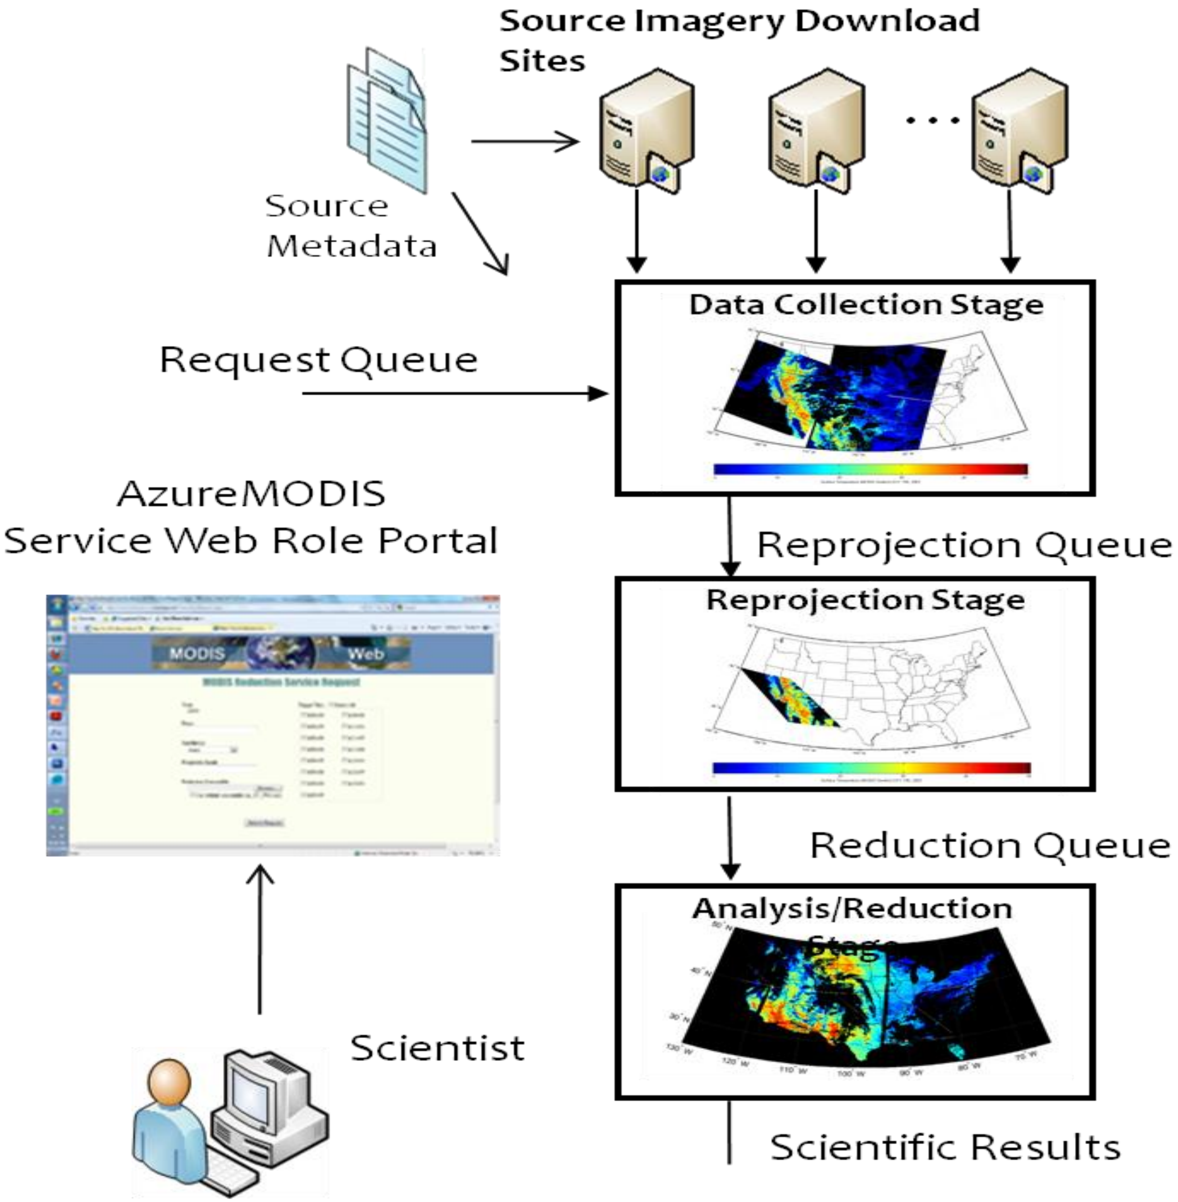

Supporting Smart Sensors and Data Fusion

in collaboration with Ed Lazowska, Bill Howe, Keith Grochow & Mark Stoermer, U of Washington

NSF Ocean Observing Initiative

- Hundreds of cabled sensors and robots exploring the sea floor
- Data to be collected, curated, mined

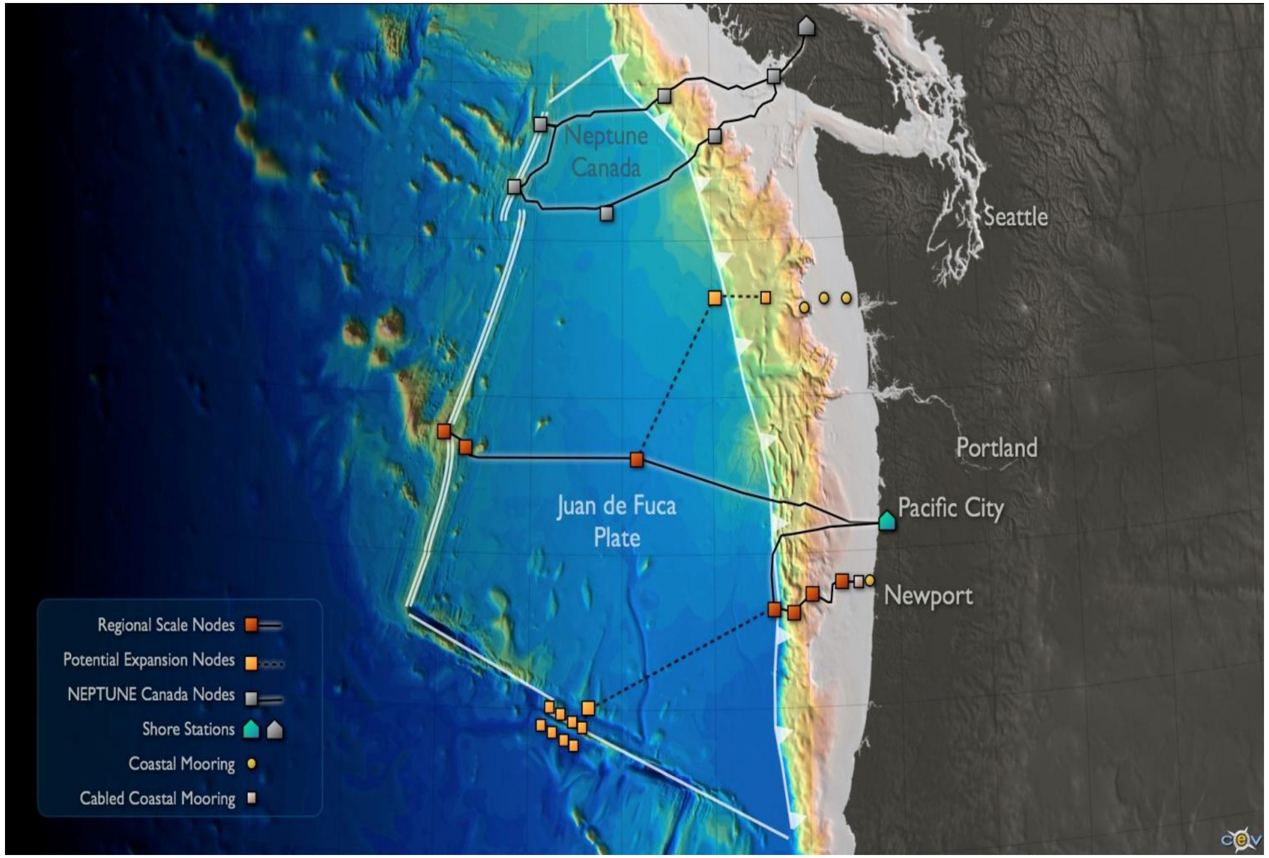

Demo

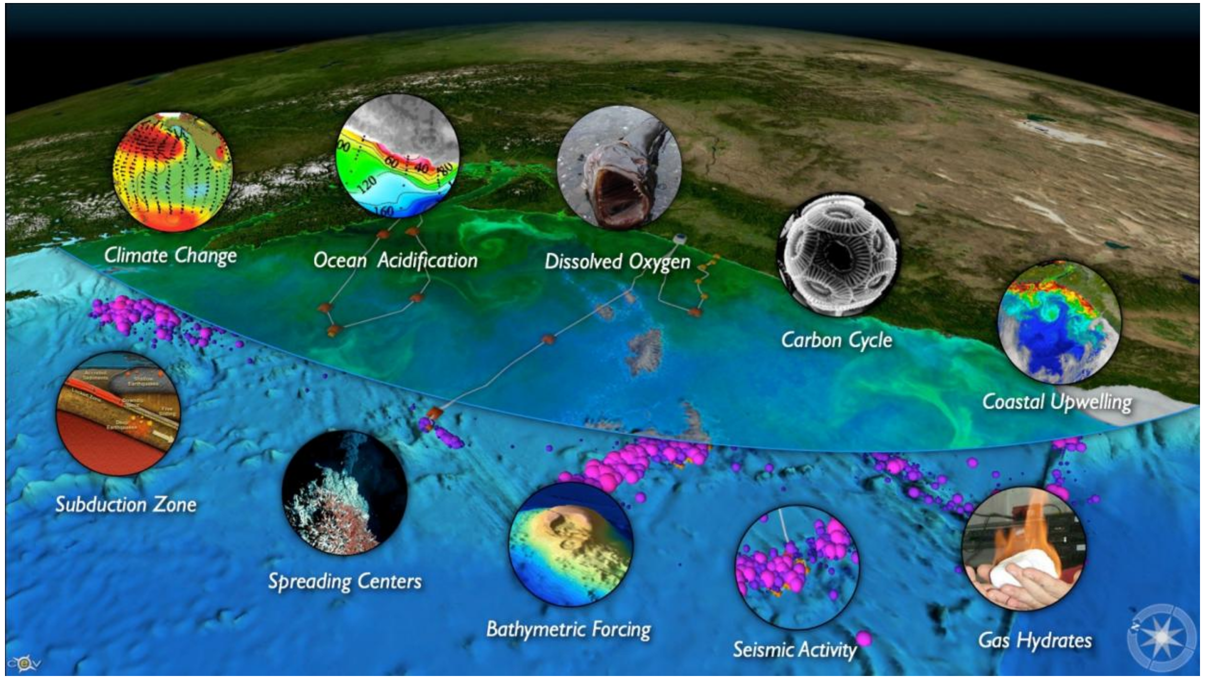

Conclusions

- Clouds are the largest scale computer centers ever constructed and have the potential to be important to both large and small scale research.
- Suitable for "loosely coupled" data parallel applications, but tightly coupled low-latency applications perform poorly on clouds today.
  - *Landscape fast changing, research computing market is significant*
  - *Provide valuable fault tolerance and scalability abstractions*
- Use patterns enormously benefit from databases and data services
  - *Rapidly extract small subsets of large data sets*
  - *Compute aggregates, perform compression, etc.*
  - *Fast sequential read performance is criticial.*
- Software is becoming a new kind instrument
  - *Value added federated data sets*
  - *Simulations*
  - *Hierarchical data replication*

Questions?

barga@microsoft.com
http://research.microsoft.com/en-us/people/barga/

Microsoft Research eXtreme Computing: research.microsoft.com/en-us/labs/xcg/default.aspx

Microsoft Research Cloud Research Team: www.research.microsoft.com/cloud

The Fourth Paradigm: www.research.microsoft.com/fourthparadigm

Microsoft Cloud Services: www.microsoft.com/cloud