



On-Demand Research Computing

- Infrastructure as a Service –
- Software as a Service –
- Cloud Storage Solutions –

David A. Lifka

Cornell Center for Advanced Computing

lifka@cac.cornell.edu



Cornell CAC: Focused on Service

Consulting

- Assist with new faculty start-up packages
- Benchmarking & performance analysis
- Proposal development & participation
- Custom programming, debugging, parallelization & optimization
- Development and support for scientific workflows
- Custom training > live & web-based
- Strategic partnerships > vendors, national CI, researchers

Computing



- High-performance computing
- Cloud computing

Data storage and management

- Scalable disk storage & archival storage
- Database resources

National cyberinfrastructure partnerships

- NSF XSEDE – *Extreme Science and Engineering Discovery Environment*
- NSF TACC Stampede – *Enabling, Enhancing, and Extending Petascale Computing*
- NSF STCI University of Chicago – *Managing Cloud Usage Allocation and Accounting*





STCI - Experimental TeraGrid Resource (#0844032)

Software as a Service (SaaS)

red cloud with MATLAB provides:




- Seamless access to MATLAB Distributed Computing Server (MDCS) workers from your personal workstation
- Access to NVIDIA GPUs for added performance*
- “Quick Queue” for quick turnaround and scientific gateways
- Received HPC Innovation Excellence Award in 2011 from the International Data Corporation



Requires MATLAB Client with Parallel Computing Toolbox

*Use of GPUs can decrease your code run time providing optimal use of your subscription



512-core simulation of networks of coordinated amino acid variation in Hepatitis C virus, a major cause of liver disease worldwide, enabled by the MATLAB on the TeraGrid experimental resource at CAC





nanoHUB Tools



NanoMOS Simulation

Transport and Bias | Device Description | Simulation Options | Simulate

Load example: A Well-Tempered Double-Gate MOSFET

Transport | Bias

Transport Model: classical ballistic transport

Low Field Mobility (cm/s): 300

Caughey-Thomas Parameter: 2

Electron Saturation Velocity (cm/s): $1e+07$

Device Description >

Transport and Bias | Device Description | Simulation Options | Simulate

Device | Material

Ambient Temperature: 300K

Source/Drain Doping Concentration (/cm³): $2e+20$

Body Doping Concentration (/cm³): 0

Top Gate Length (nm): 9

Bottom Gate Length (nm): 9

Source/Drain Length (nm): 10

Source Extension Length (nm): -4

Drain Extension Length (nm): -4

Source Gaussian Doping Profile Slope (dec/nm): 1

< Transport and Bias | Simulation Options >



NanoNet Simulation

Options: Effect of alignment on network transistor

Tube Parameters | Device Geometry

Tube Density [D] (1/um²): 12

Tube Length [L_S]: 1um

Tube Diameter [d]: 1nm

Coupling Parameter [C_B]: 1

Alignment Parameter: 10, 60, 50

Random seed: 10

Simulate >




Case Study Examples

Analysis of MRI Brain Scans

Ashish Raj and Miloš Ivković

Weill-Cornell Medical College

Given two different regions of the human brain, how interconnected are they?



Indirect Selection in Mutualistic Networks

Andria Salas and Dr. Stuart Borrett

University of North Carolina Wilmington, Center for Marine Science

Hypothesis: the nested structure of mutualistic networks provides ecological systems more resilience to evolutionary change compared to random networks.





CUSat Recontact Analysis

Prof. Mason Peck, Dr. Rajesh Bhaskaran,

Dr. Steven Lantz, Undergrads

Cornell University

Create and deploy an autonomous in-orbit inspection satellite system.





red cloud Motivation

- **Research computing means many different things...**
 - Scientific workflows have different requirements at each step
 - Cloud is only part of the solution
 - Connecting to and from other CI resources is important
- **Nobody likes a bad surprise**
 - Transparency, no hidden costs
 - Need a way to bound financial risk
 - Geographically distributed resources for disaster recovery
 - Bursting for peak periods vs. over investment in dedicate infrastructure
- **Economies of scale**
 - Sharing hardware and software where it makes sense
 - Pay for what you need, when you need it
- **Customized environments for various disciplines**
 - Collaboration tools
 - Data storage & analysis tools
 - Flexibility to support different computing models (e.g. Hadoop)



red cloud Provides

Predictable, Reproducible, Reliable Performance

We publish hardware specifications (CPU, RAM, network) and do not oversubscribe.

Convenient

Need system up and running yesterday.

Need a big fast machine for only a few months, weeks or days.

Need a small server to run continuously.

No Hidden Costs

No cost for network traffic in or out of the cloud.

Fast Access to Your Data

Fast data transfers via 10Gb Ethernet in or out of the cloud at no additional charge.

Globus Online access

Economies of scale

IaaS: Infrastructure

SaaS: MATLAB

Storage: EBS, Globus based Archival Storage

Expert Help

System, application, and programming consulting are available.

Easy Budgeting with Subscriptions

No billing surprises!

Eucalyptus is Amazon API Compatible

Migrate when your requirements outgrow Red Cloud.



2012 – Hurricane Sandy

Disaster brings new opportunity

A Flooded Mess That Was a Medical Gem

By ANEMONA HARTOCOLLIS

Published: November 9, 2012



The federal government's emergency management chief trudged through darkened subterranean hallways covered with silt and muddy water Friday, as he toured one of New York City's top academic medical centers in the aftermath of Hurricane Sandy. The basement of the complex, NYU Langone Medical Center in Manhattan, smelled like the hold of a ship — a mixture of diesel oil and water.





Weill Cornell Medical College Secure Network Extension

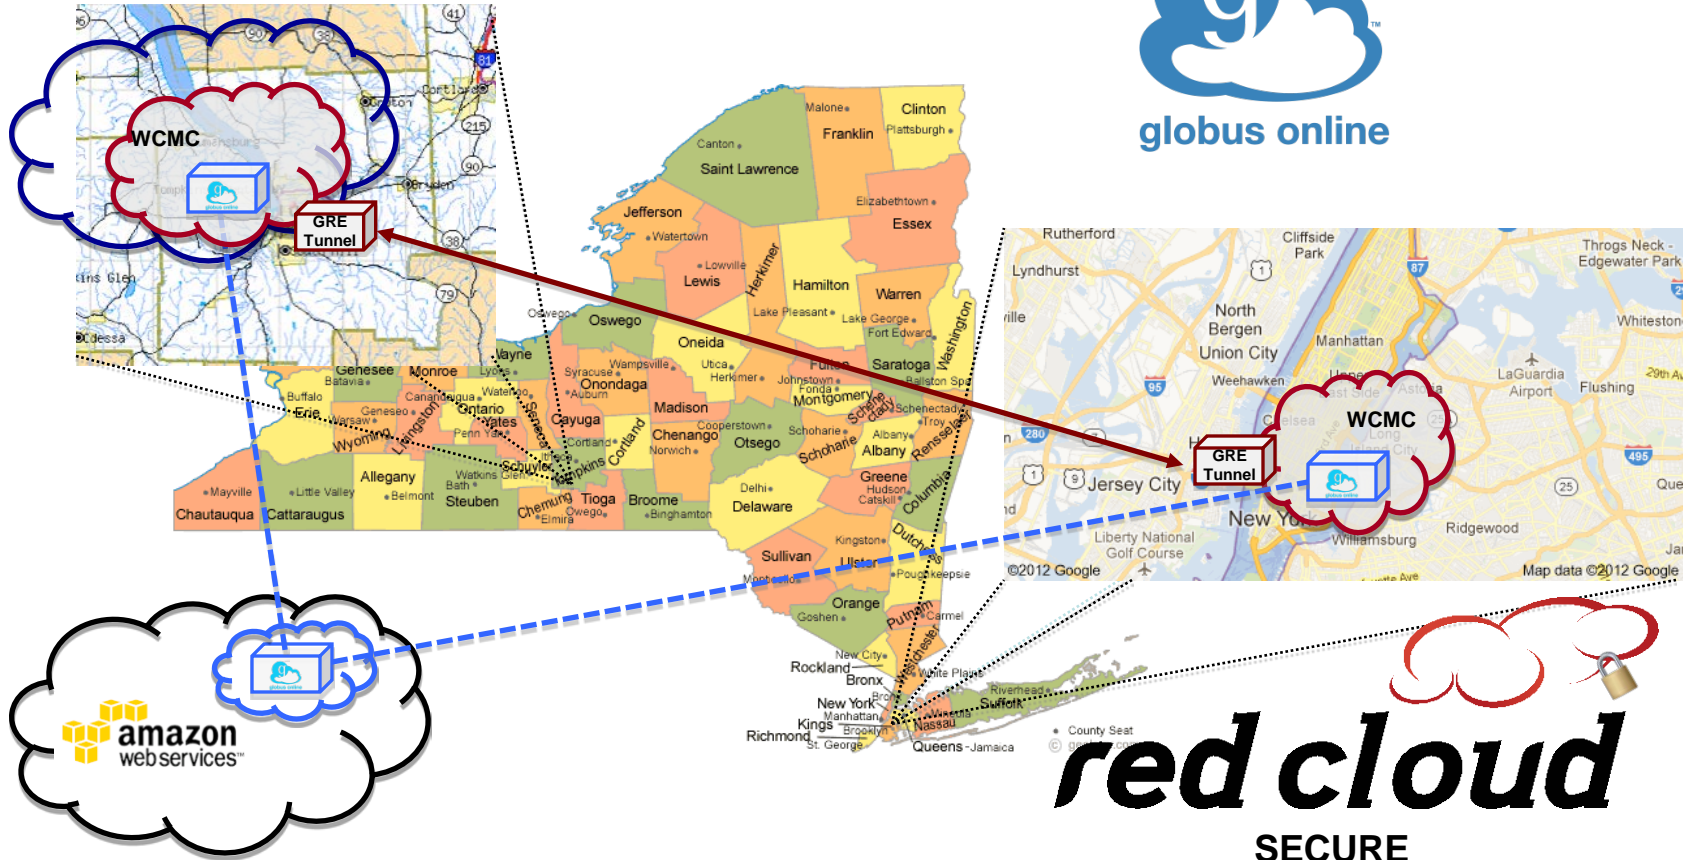
HIPAA Compliant
Co-location Facility
Cornell Ithaca, NY Campus





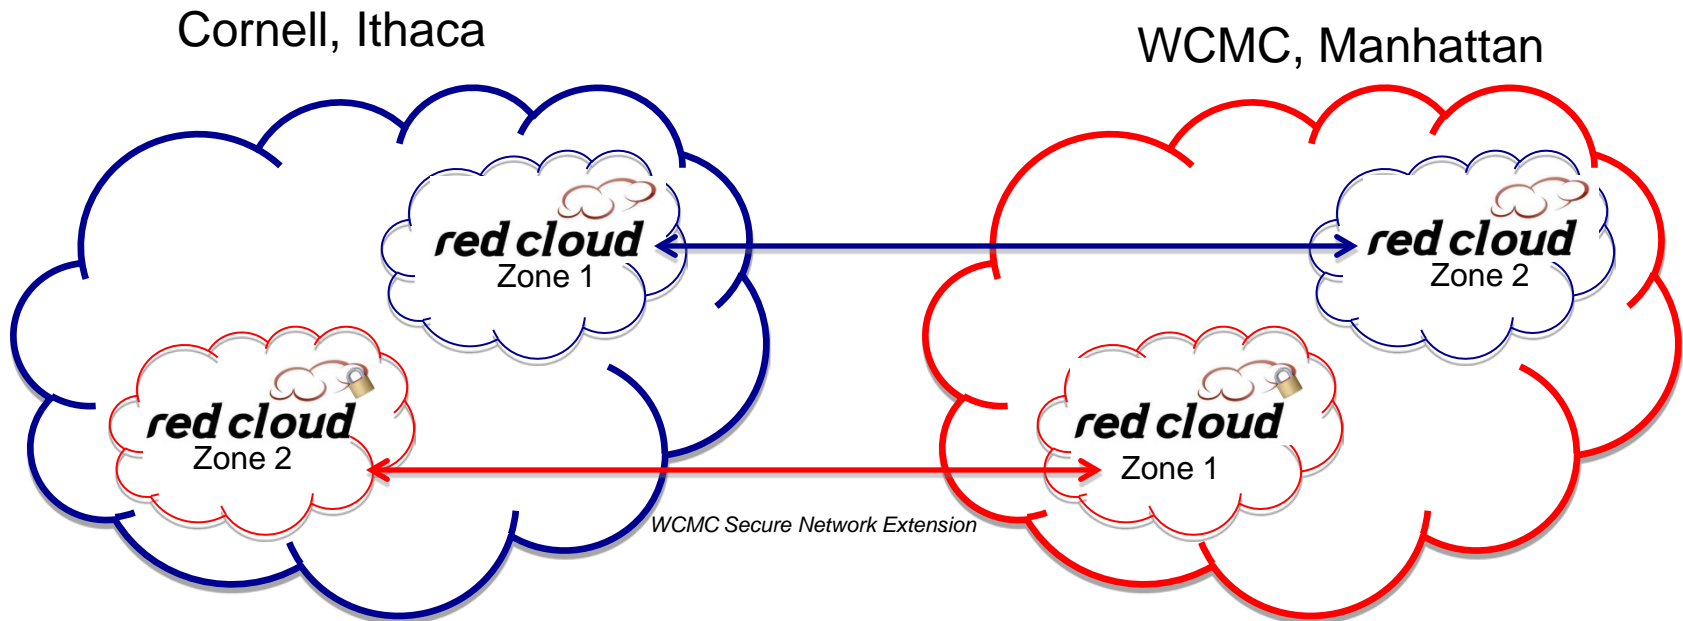
WCMC Firewall Extension to Ithaca

Co-location for Critical Servers





Two Availability Zones for Redundancy



- Independent Management and Administration
 - Zone 1 is the administrative zone for each cloud, providing authentication & authorization
 - Zone 2 is a second geographically distributed set of resources
 - Support for disaster recovery
 - Support for additional users under heavy load or data locality



Red Cloud Virtual Compute Cluster


- Allows any Red Cloud subscriber to submit traditional HPC batch scripts from the CAC Linux login node using the SLURM scheduler.
 - <https://computing.llnl.gov/linux/slurm/>
- Users don't have to know anything about Cloud
 - If you know how to submit a batch script to an HPC cluster and you are ready to go
- Different SLURM job queues will offer different HPC cluster node configurations. Some examples include (but not limited to):

<u>Cores</u>	<u>RAM (GB)</u>	<u>Local Disk (GB)</u>	<u>Ethernet (Gb)</u>
1	4	100	10 shared
2	8	200	10 shared
4	16	400	10 shared
8	32	800	10 shared
12	48	1000	10 shared



Deploy a Hadoop Cluster on Demand





Current Users

- Columbia University
- Cornell University
- George Mason University
- Northwestern University
- Norwegian University of Life Sciences
- Purdue University
- Seoul National University of Science and Technology
- Skolkovo Institute of Science and Technology
- University of Washington
- University of Wisconsin-Milwaukee





Red Cloud Subscriptions

www.cac.cornell.edu/redcloud

Subscriptions

	<u>Cornell University</u>	<u>Other Academic Institutions</u>
1 "Red Cloud"	\$400.00	\$640.00
1 "Red Cloud with MATLAB"	\$600.00	\$960.00
Storage (Current)	\$0.50/GB/year	\$0.80/GB/year

1 subscription = 8,585 core hours.

50GB storage is included with your first subscription.



Thank You!

Questions?