

A Short Introduction to Grid Computing

High Performance Computing Day

Texas A&M University

May 2, 2007

Steve Johnson

Department of Mathematics
and

Institute for Scientific Computation
Texas A&M University

What is a Computational Grid?

- Coordinates resources that are not subject to central control
- Uses open standards
- Non-trivial quality of service is achieved

Ian Foster, "What is the Grid? A Three Point Checklist"

Why use a grid?

- Cycle scavenging
- Event driven computing
- Cumulative cluster computing
- General purpose HPC
- Distributed Data Grids

Cycle Scavenging

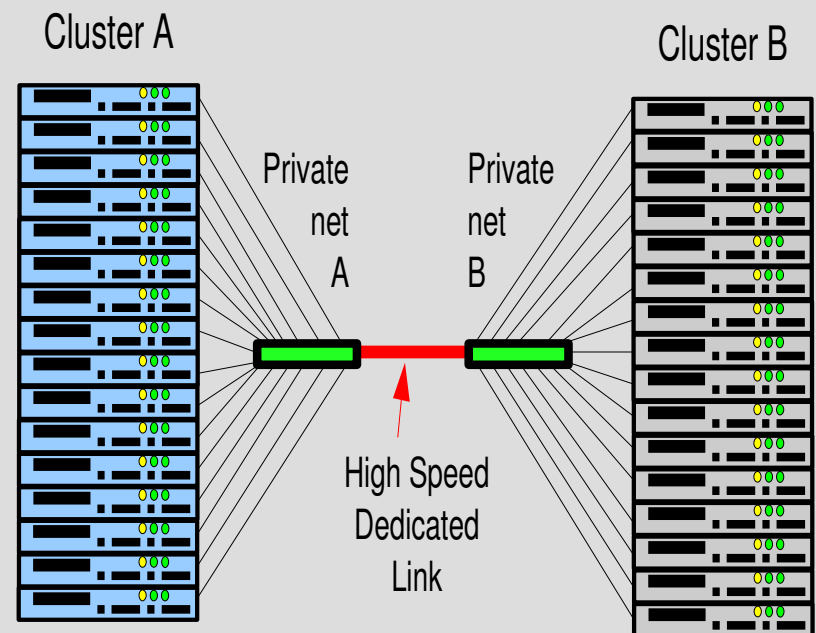
- Utilize “excess” cycles on an otherwise idle resource.
- Job output is not time critical
- [SETI@home](#), [Folding@home](#). Initiated by the owner of the resource.
- Condor (free), United Devices (not free)
 - Schedules jobs on idle workstations
 - Supports checkpoint/restart
 - Usually for single-node non-MPI jobs

Event Driven Computing

- Time critical computations
- Triggered by availability of input data
- Usually runs on HPC resources, not desktops
- A scheduler supporting preemption is desirable for busy resources
- Examples:
 - Weather models
 - Storm surge models
 - High energy physics data analysis

Cumulative Cluster Computing

- Run a single application using multiple, distinct resources
- MPI across two or more clusters
- Not frequently used
 - Difficult to schedule
 - Bandwidth issues
 - Connectivity issues: compute nodes are typically on a private network.

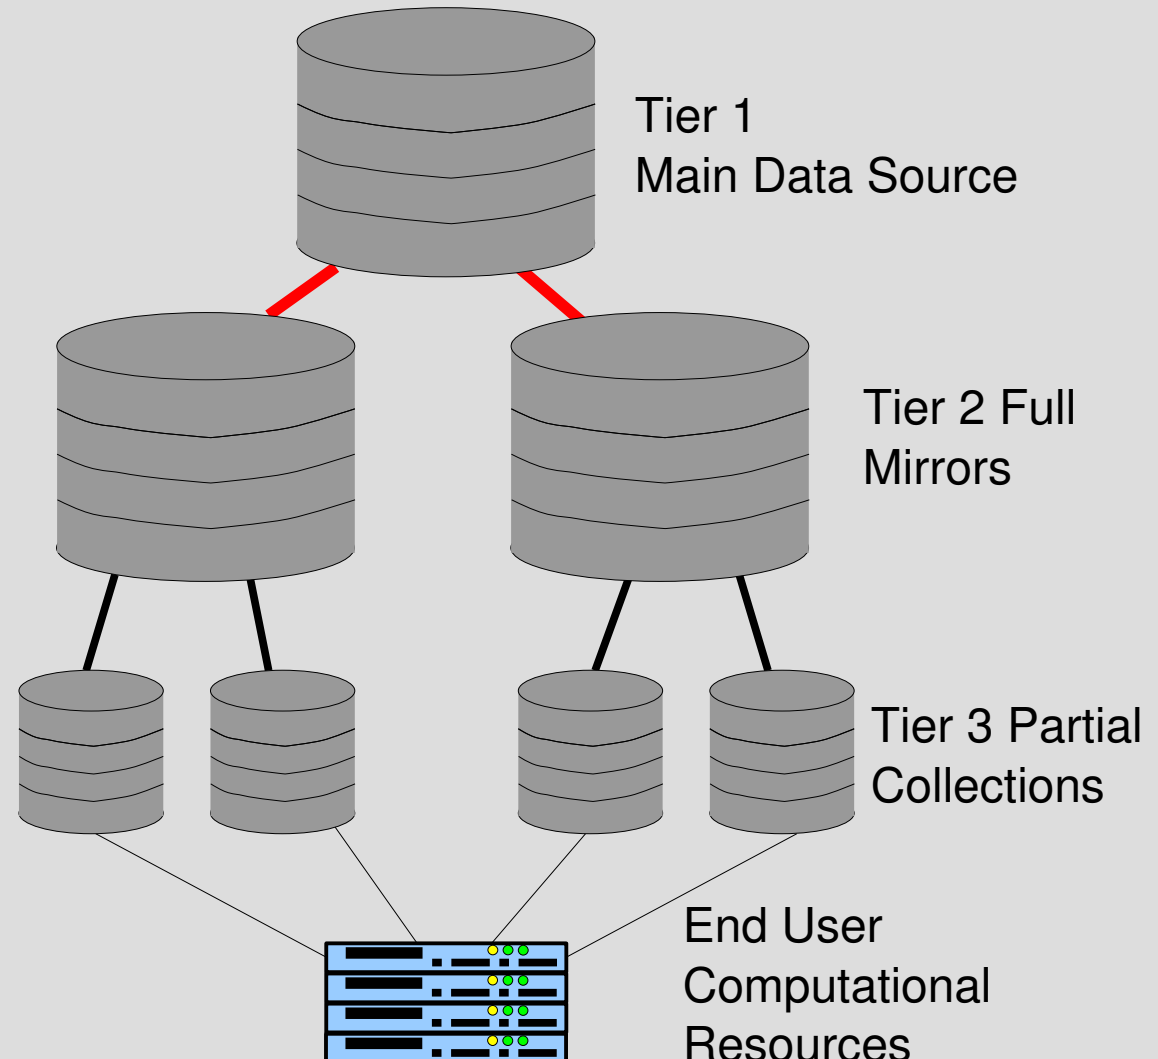


General Purpose HPC

- Use a grid metascheduler to match resources to the requirements of the job.
 - Maintains list of resources, their attributes, and their availability.
- Can also write your own metascheduler if you have knowledge of the resources.

Distributed Data Grids

- Not necessarily a grid solely for computation
- Provides access to data collections and repositories
- End users access lower tier data based on their project affiliations, proximity



Authentication Authorization


- Each user and each resource is issued a certificate
 - Public and private keys
 - User unlocks his/her private key w/ a password
- Trust fabric between grid participants
 - User can trust that the host is who it says it is and vice versa.
- Resource owners have ultimate control over who can access their systems

SURAGrid

- Southeastern Universities Research Association - <http://www.sura.org/suragrid>
- Started as NSF NMI testbed in 2003
- Over 25 participating institutions
- Over 10 Tflops capacity and growing

SURAgriid Participants

SURAgriid Participants
(As of October 2006)



SURAggrid Resources at A&M

- Calclab
 - 191 Dell Optiplex desktops available during off hours. 281 cores, 2TFlops, 1GB mem per core, 3.5TB disk
- Tensor
 - 128 node Opteron cluster. 256 cores, 735GFlops, 1GB-2GB mem per core, 3TB disk, GigE and IB (35 nodes)
- Immersive Visualization Center
 - 25' x 8' x 120 degrees rear-projected curved screen
 - Dell Precision 690, Dual Xeon, 32GB RAM, 2xNVIDIA Quadro 4500

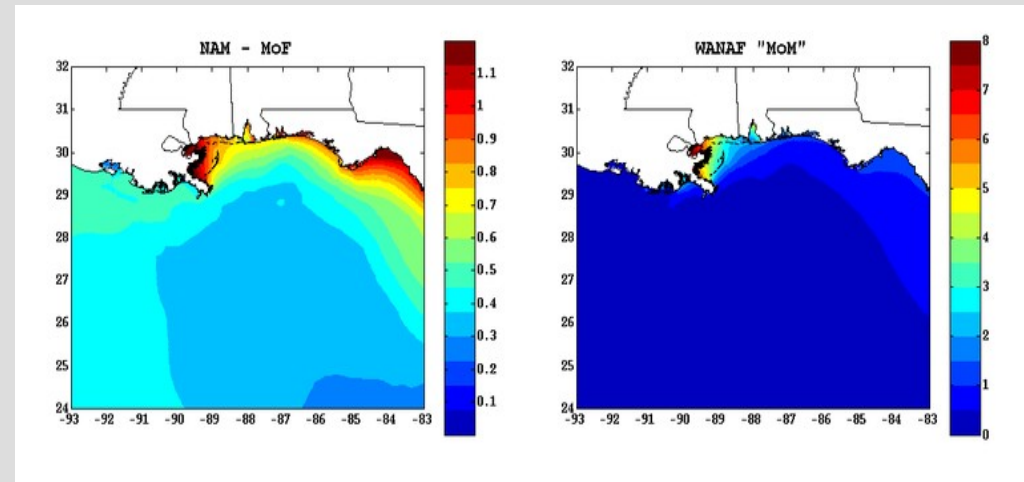
SURAggrid Resources at A&M

(coming soon)

- Hydra
 - IBM p5/575+. 40 nodes, 32GB/node. 21TB GPFS disk, HPS, 640 total cores, 5TFlops.
 - 20% of resources for SURAggrid per IBM/SURA partnership
- TexAQS & Windfarm
 - AMD and Intel clusters, 84 cores, 2GB/core, GigE, 410GFlops
 - support SURA Coastal Ocean Observing and Prediction (SCOOP) program

SURAggrid Applications

- ADCIRC – Coastal Circulation and Storm Surge Model
 - Part of SURA Coastal Ocean Observing and Prediction (SCOOP) program
 - Managed by RENCI
 - Data driven on-demand computing
 - Uses its own internal scheduler, portal



Left: ADCIRC max water level for 72 hr forecast starting 29 Aug 2005, driven by the "usual, always-available" ETA winds.

Right: ADCIRC max water level over ALL of UFL ensemble wind fields for 72 hr forecast starting 29 Aug 2005, driven by "UFL always-available" ETA winds.

Images credit: Brian O. Blanton, SAIC

ADCIRC info courtesy of Howard Lander, RENCI

SURAggrid Applications

- Wavewatch 3 – event driven coastal modeling
 - Part of SCOOP project
 - LSU
- UcoMs – Ubiquitous Computing and Monitoring System
 - Petroleum simulation
 - LSU
- Optimization for Threat Management in Urban Water Systems
 - User submission and simulated event triggered
 - NCSU



SURA Coastal Ocean Observing and Prediction (SCOOP) Program

Simple SCOOP Status Portlet

Today's Run Status

Model Name	Wind Field	Cycle	Grid Field	Start Time	CPU Count	Resource	Status
WW3	WANAFe08-UFL	00	LLMF	Thu 2006.11.30 02:14:59 PM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe10-UFL	00	LLMF	Thu 2006.11.30 02:14:44 PM CST	32	mike4.cct.lsu.edu	
ADCIRC	WANAFe02-UFL	00	ec95	Thu 2006.11.30 02:14:36 PM CST	6	canbc01.louisiana.edu/jobmanager-pbs	
ADCIRC	WANAFe07-UFL	00	ec95	Thu 2006.11.30 02:14:34 PM CST	6	canbc01.louisiana.edu/jobmanager-pbs	
ADCIRC	WANAFe05-UFL	00	ec95	Thu 2006.11.30 02:14:30 PM CST	6	canbc01.louisiana.edu/jobmanager-pbs	
ADCIRC	WANAFe04-UFL	00	ec95	Thu 2006.11.30 02:14:24 PM CST	8	dante0.renci.org/jobmanager-pbs	
ADCIRC	WANAFe03-UFL	00	ec95	Thu 2006.11.30 02:14:24 PM CST	8	ci-team.acis.ufl.edu/jobmanager-pbs	
ADCIRC	WANAFe06-UFL	00	ec95	Thu 2006.11.30 02:14:10 PM CST	0	Unknown	
WW3	WANAFe09-UFL	00	LLMF	Thu 2006.11.30 02:14:10 PM CST	32	mike4.cct.lsu.edu	
ADCIRC	WANAFe01-UFL	00	ec95	Thu 2006.11.30 02:14:02 PM CST	8	gauss.math.tamu.edu/jobmanager-pbs	
ADCIRC	WANAFe11-UFL	00	ec95	Thu 2006.11.30 02:13:54 PM CST	8	gauss.math.tamu.edu/jobmanager-pbs	
ADCIRC	WANAFe10-UFL	00	ec95	Thu 2006.11.30 02:13:01 PM CST	8	ci-team.acis.ufl.edu/jobmanager-pbs	
ADCIRC	WANAFe08-UFL	00	ec95	Thu 2006.11.30 02:13:01 PM CST	0	Unknown	
ADCIRC	WANAFe09-UFL	00	ec95	Thu 2006.11.30 02:13:01 PM CST	8	ci-team.acis.ufl.edu/jobmanager-pbs	
WW3	WANAFe06-UFL	00	LLMF	Thu 2006.11.30 10:38:33 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe03-UFL	00	LLMF	Thu 2006.11.30 10:38:10 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe05-UFL	00	LLMF	Thu 2006.11.30 10:37:44 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe02-UFL	00	LLMF	Thu 2006.11.30 10:37:17 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe01-UFL	00	LLMF	Thu 2006.11.30 10:36:51 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe07-UFL	00	LLMF	Thu 2006.11.30 10:36:13 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe08-UFL	00	LLMF	Thu 2006.11.30 10:33:54 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe11-UFL	00	LLMF	Thu 2006.11.30 10:33:29 AM CST	32	mike4.cct.lsu.edu	
WW3	WANAFe04-UFL	00	LLMF	Thu 2006.11.30 10:33:01 AM CST	32	mike4.cct.lsu.edu	
ADCIRC	WANAFe03-UFL	00	ec95	Thu 2006.11.30 10:32:38 AM CST	8	dante0.renci.org/jobmanager-pbs	

SURAggrid Applications

- Genome Alignment
 - GSU
- Bio-Sim: Bio-electric Simulator for Whole Body Tissues
 - ODU
- Dynamic BLAST
 - Use meta scheduler for distributing BLAST computations
 - UAB

SURAggrid Applications

- Grid enabling of WRF and MM5 to support SCOOP project
 - Weather models
 - Texas A&M
- Grid enabling of QXPPC2D
 - Quasicrystal phason-phonon coupling 2D model
 - wide parameter space
 - Texas A&M
- Distributed Computing Course at Old Dominion
 - CS 775/875: Distributed Systems
 - Spring 2007



SURAggrid at A&M

The Next Steps

- SURAggrid is still under development, but is capable of running some time critical production jobs. E.g., ADCIRC.
- Looking for suitable grid computing apps from TAMU.

Application Development

Early Considerations

- Match resources to your application
 - MPI, SMP, or single-threaded?
 - x86, x86_64, p5, sparc?
 - time requirements?
- Data movement
 - assume no long term storage on remote resource
 - bandwidth requirements?
 - total storage requirements?
 - temporary storage on compute nodes?

Contacts

- SURAGrid web sites
 - <http://www.sura.org/suragrid>
 - <https://gridportal.sura.org>
- Texas A&M contact
 - Steve Johnson
 - steve //AT// isc.tamu.edu
 - 979-845-4267

Questions?