



Dynamic Applications on Grids

Computing ---> Data, Metadata, Networks

Ed Seidel

Center for Computation &
Technology

Louisiana State University



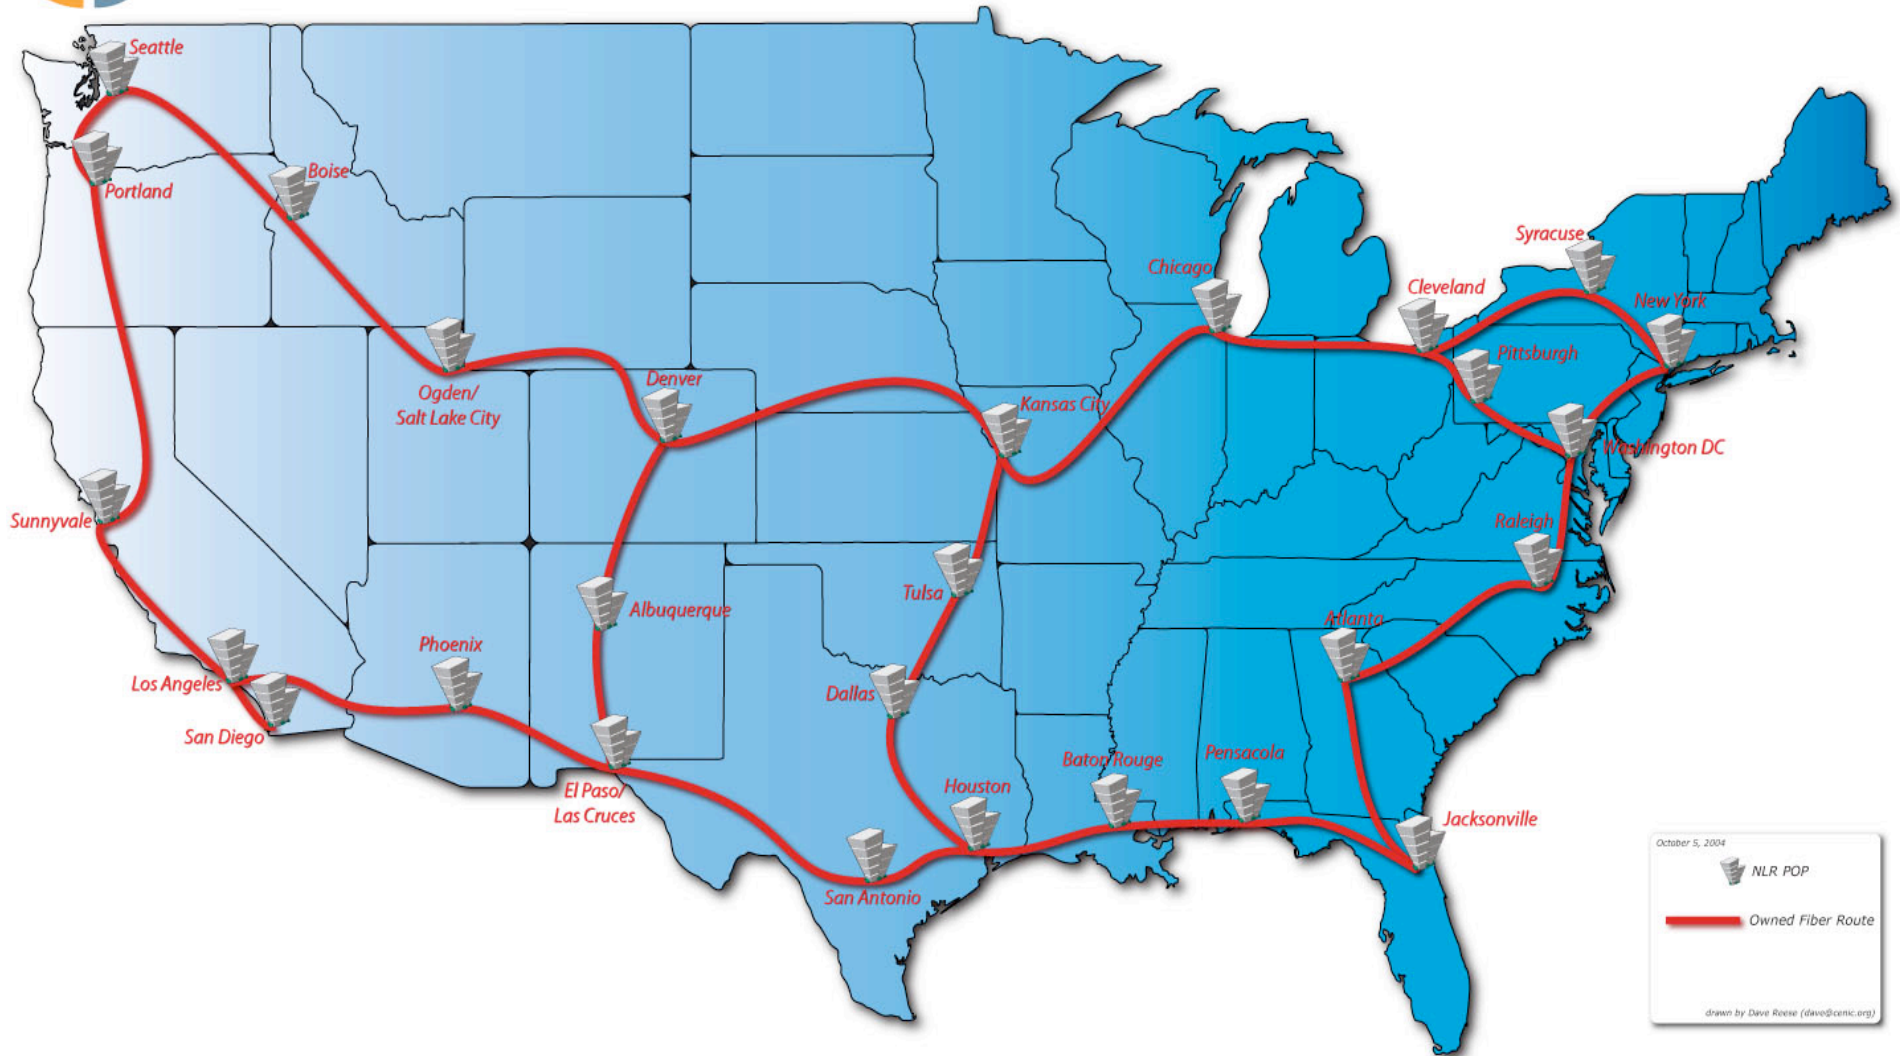


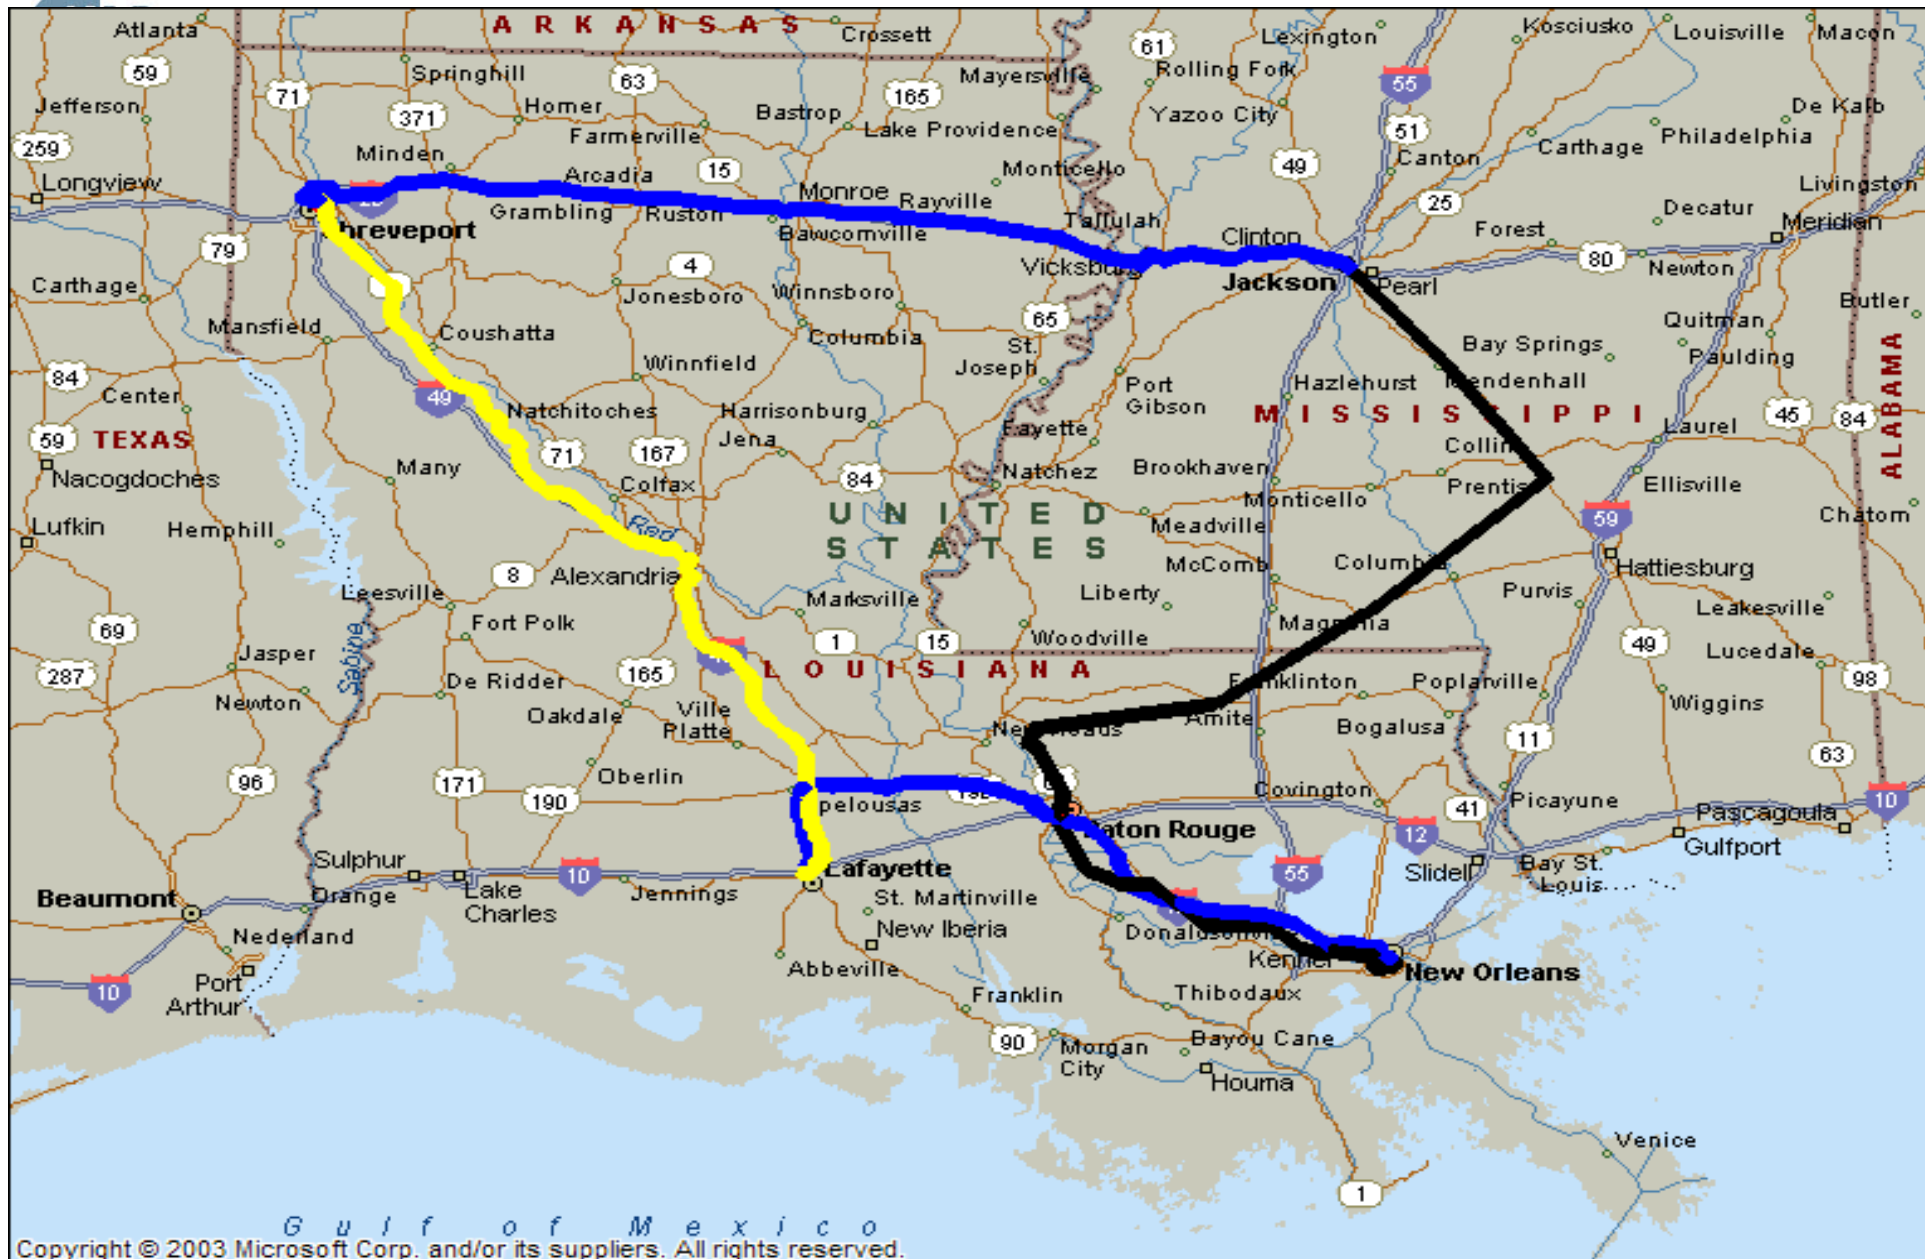
Optical Networks

- New technology for decade: 1000x faster than typical regional networks: Louisiana statewide network (LONI)
- National Lambda Rail
 - \$100M USA Optical Network
 - Backbone for next gen. research
 - Locally funded!! No Federal \$\$
- Other countries: \$100'sM
 - Canada, Poland, Holland, Czech, etc; linked!
- 2 dozen+ states committed
- Southern US states investing to be competitive for federal funding, industry



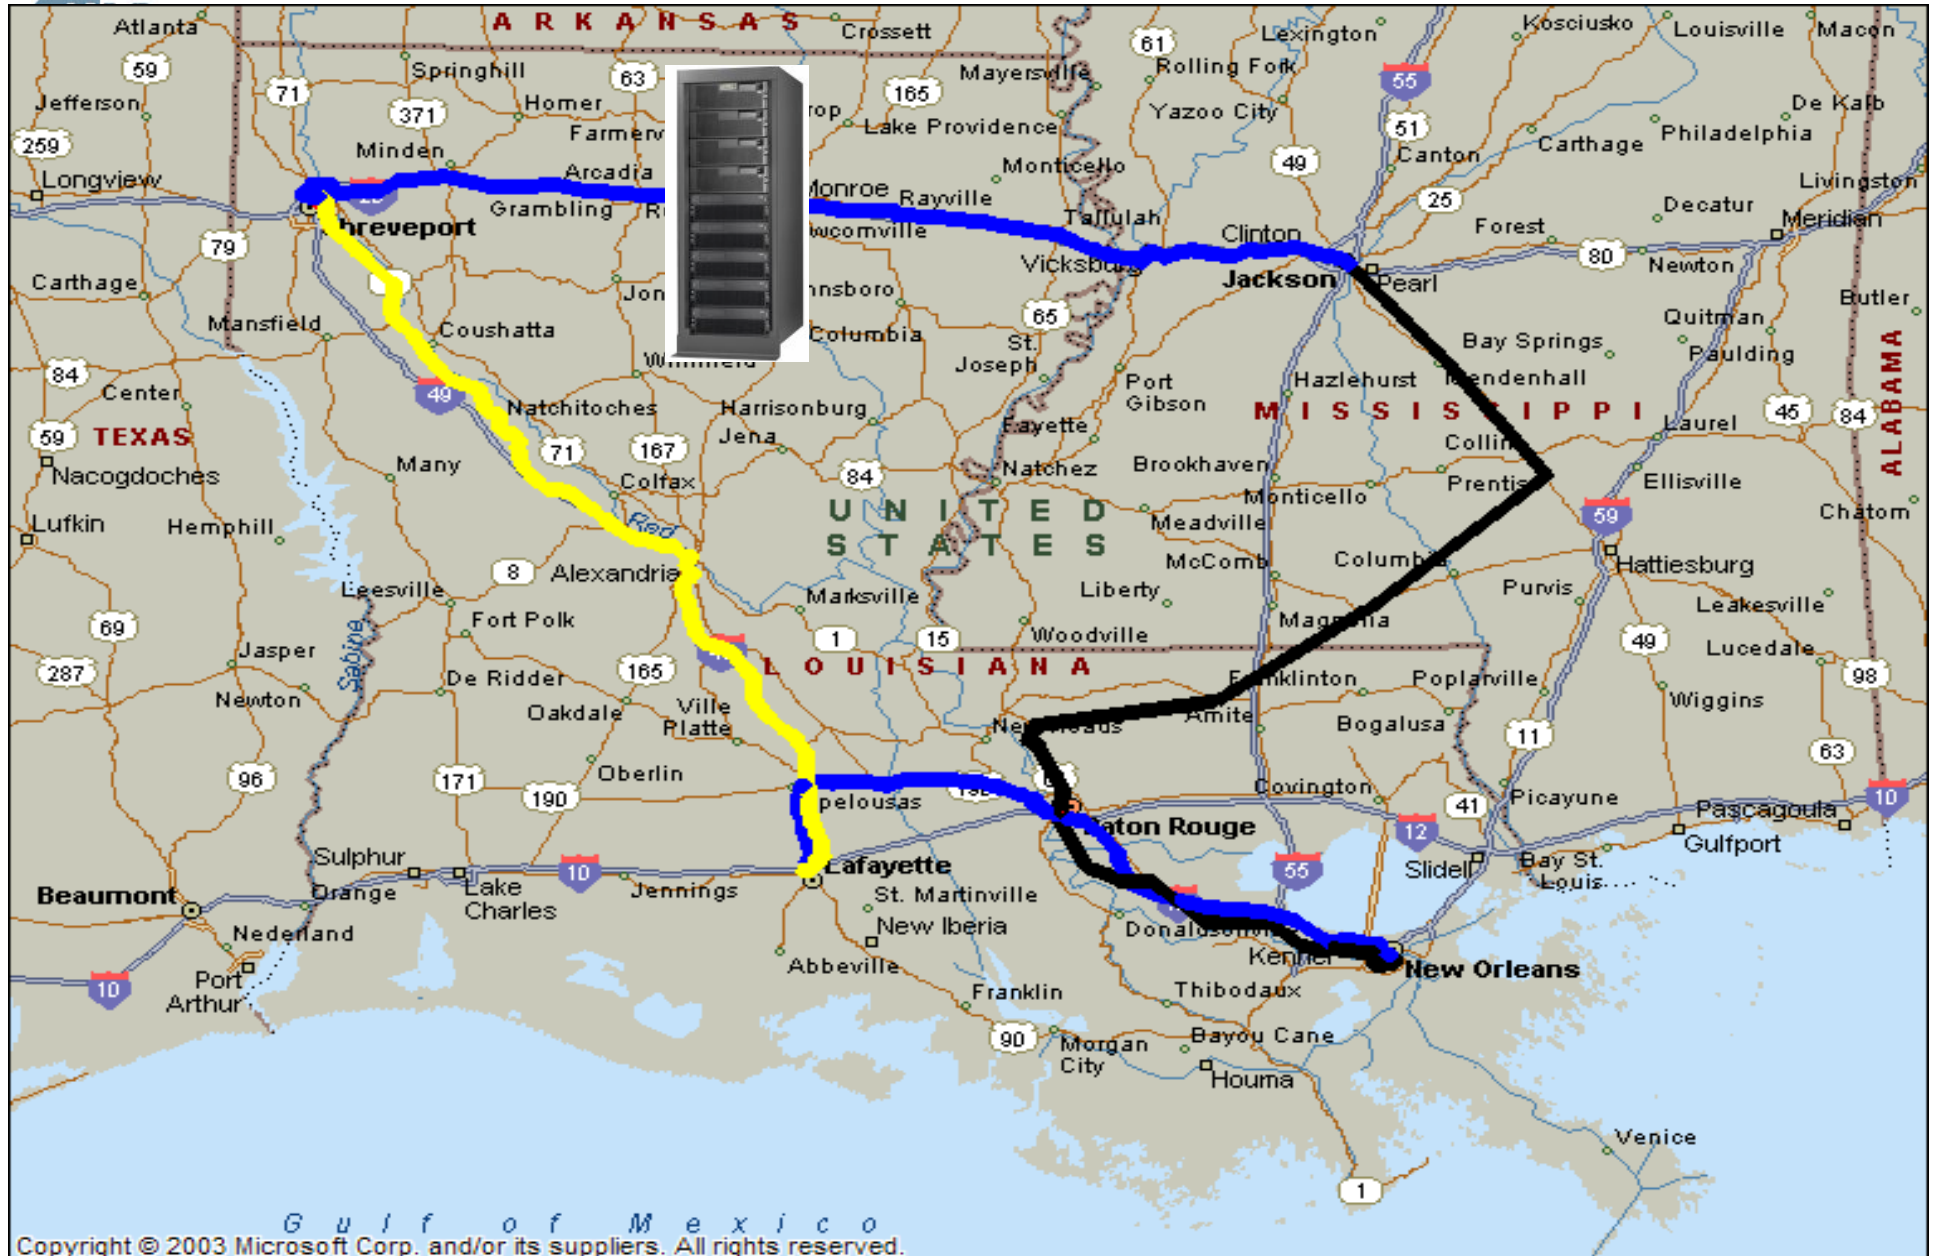
National LambdaRail Architecture

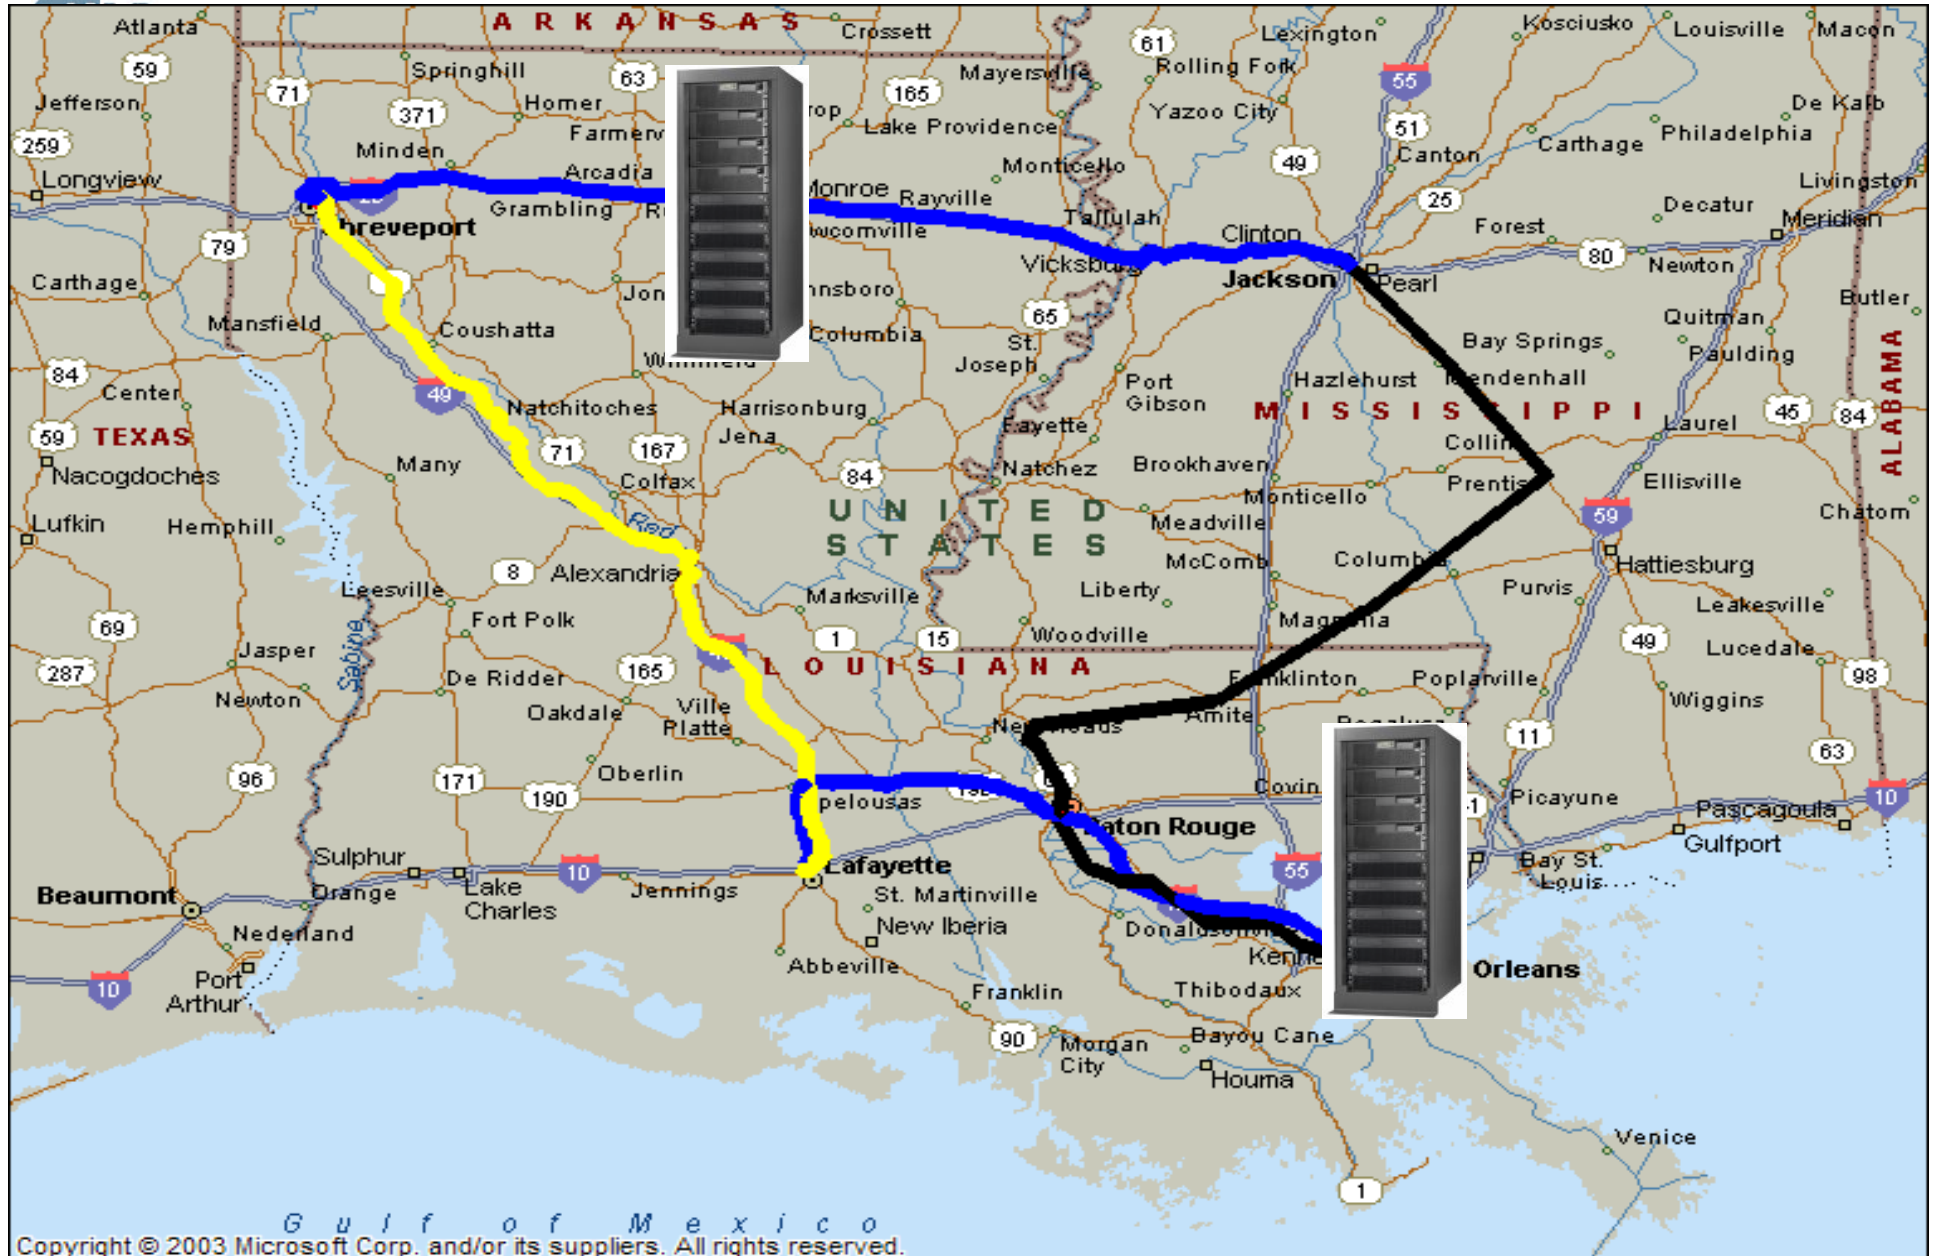


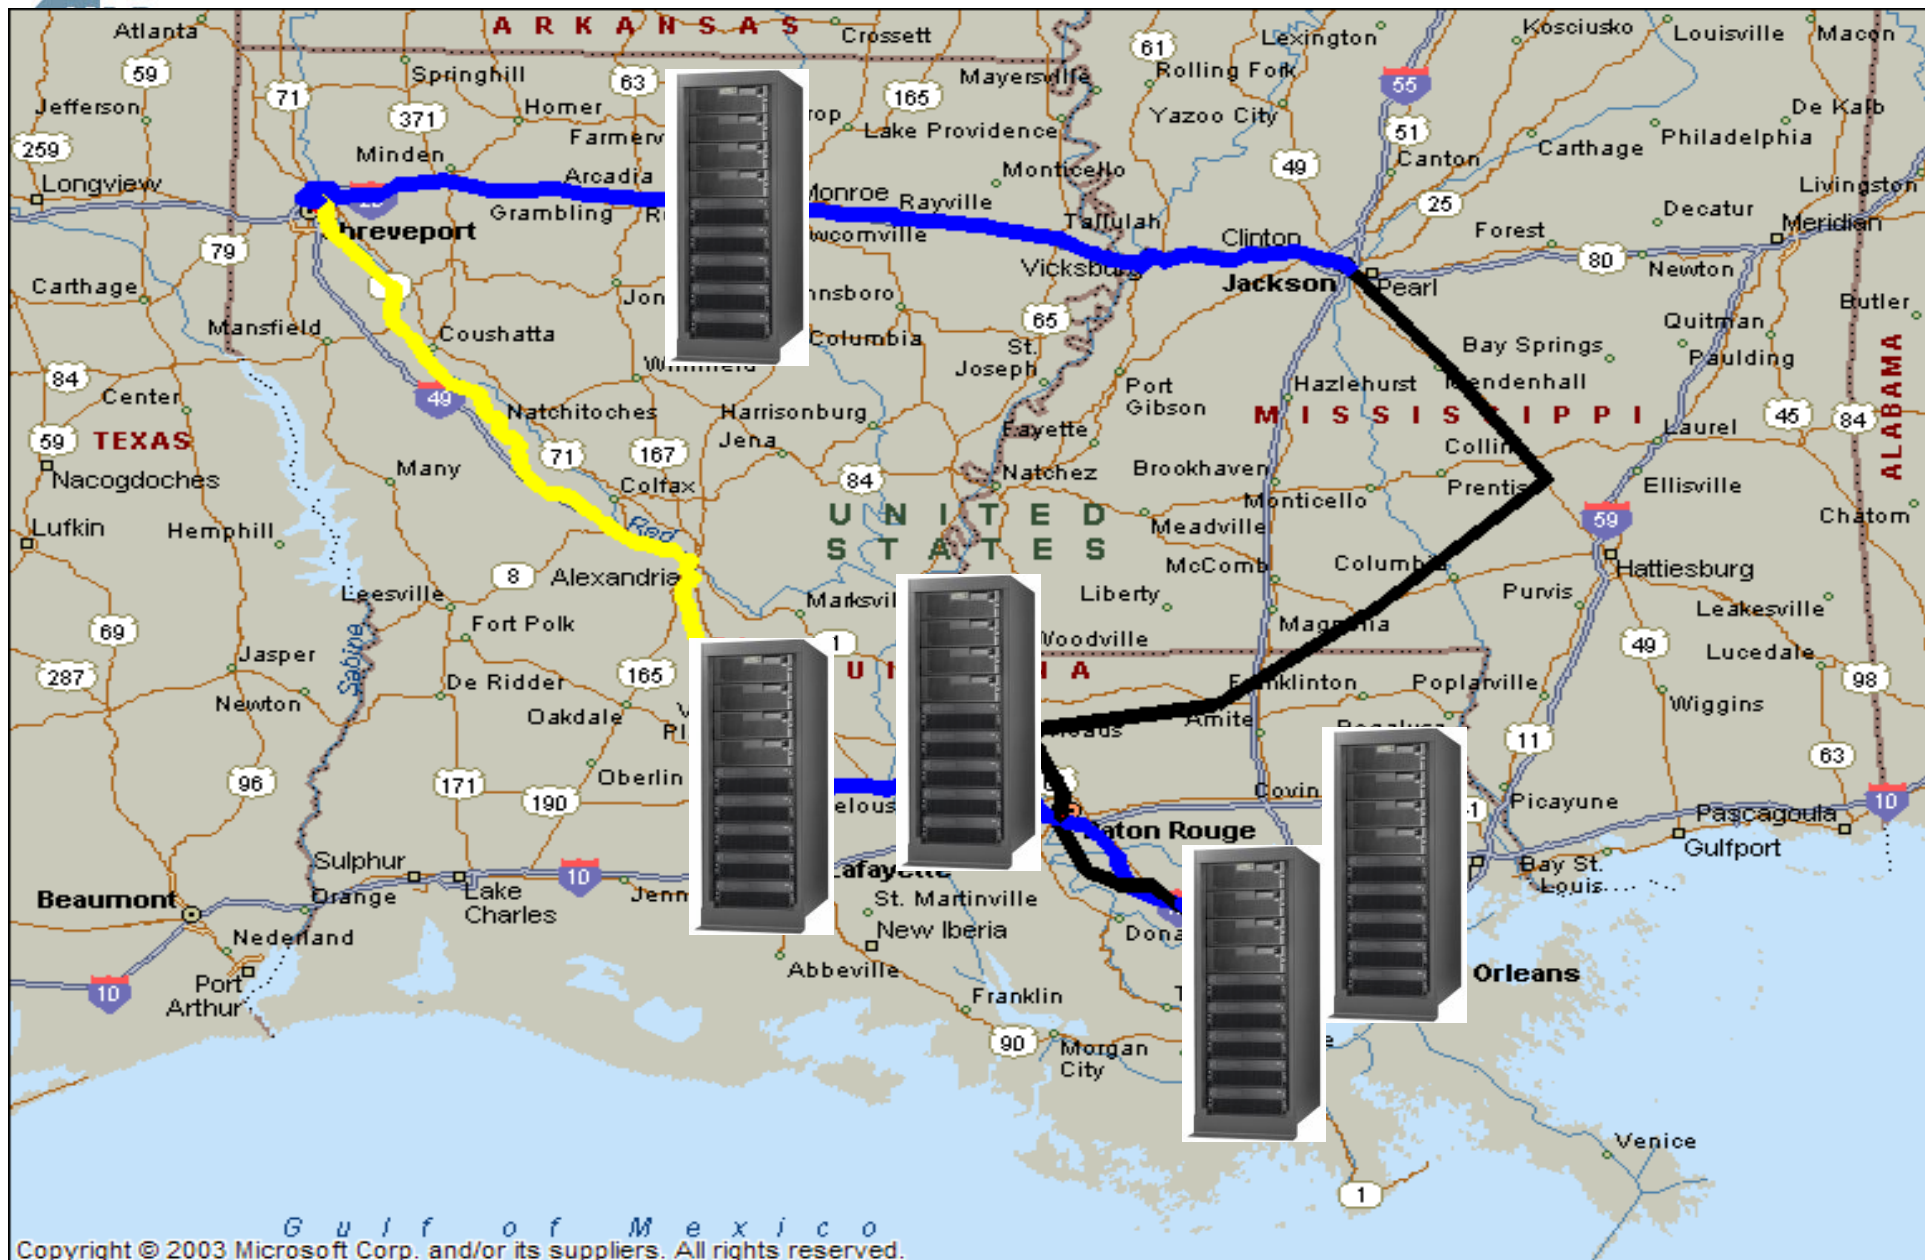


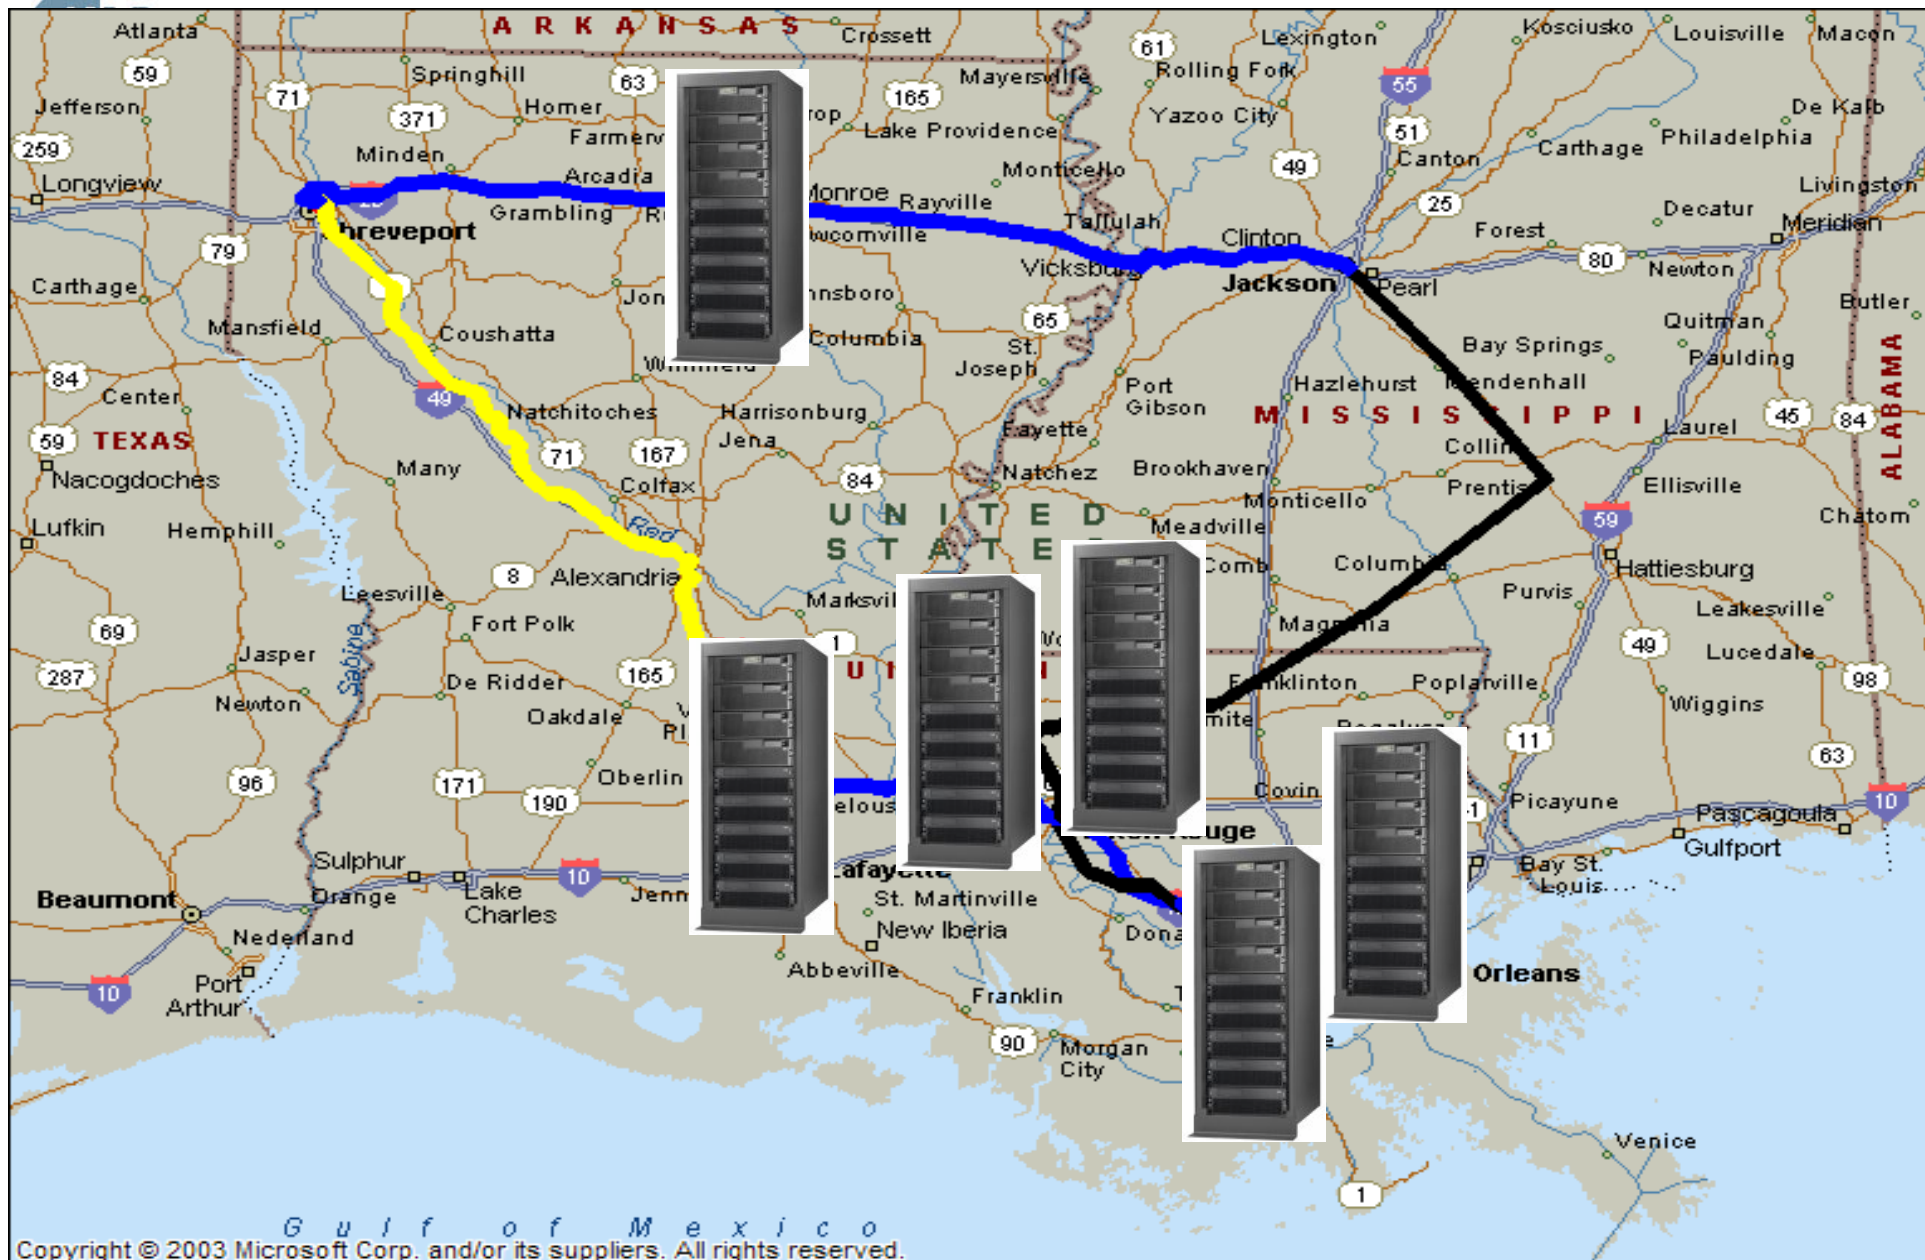
Gulf of Mexico

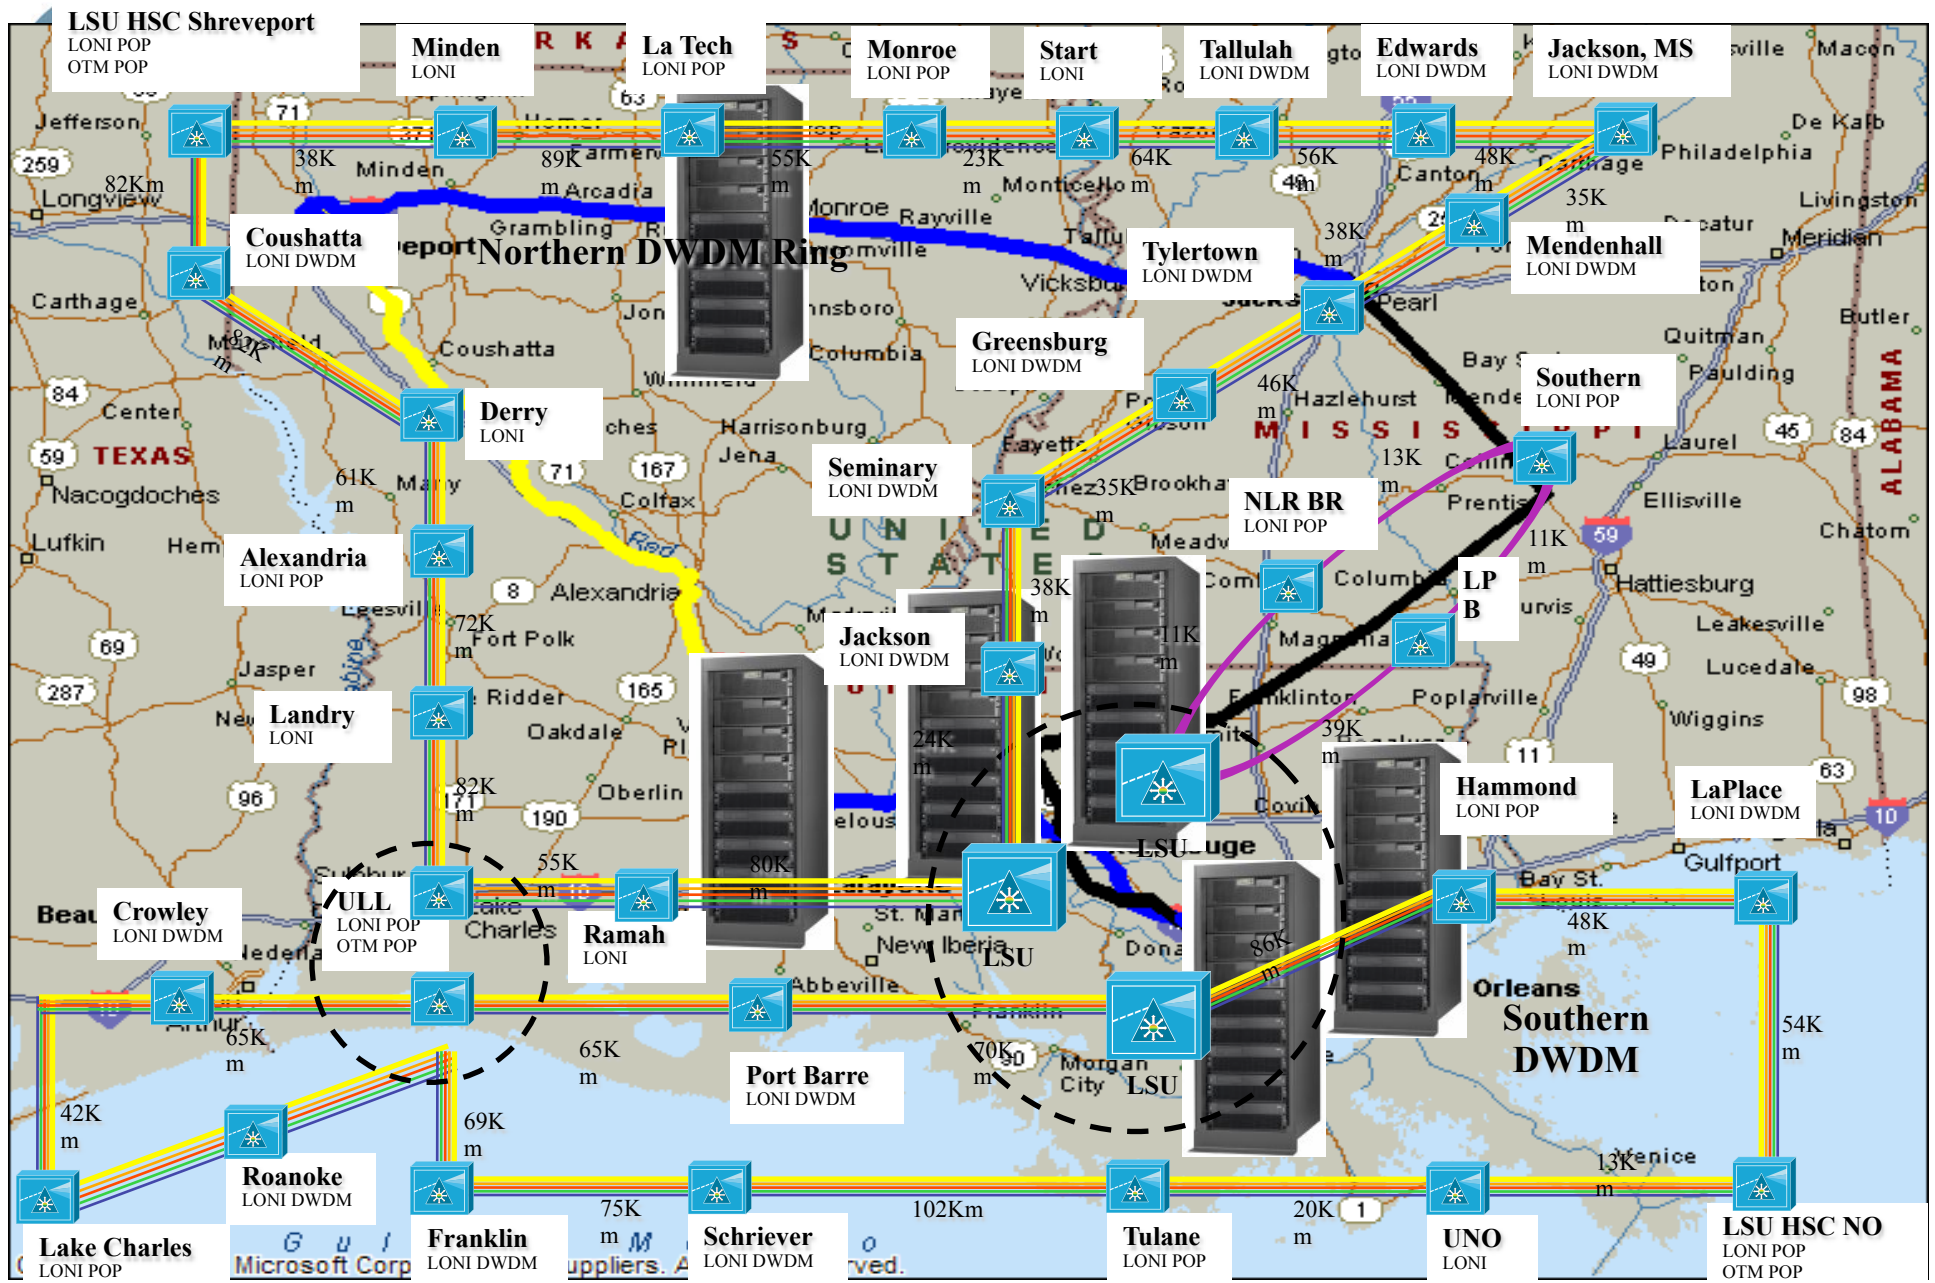
Copyright © 2003 Microsoft Corp. and/or its suppliers. All rights reserved.












2004

International Lambdas




SunLight



UKLight

NorthernLight

CERN



STARLIGHTSM

European lambdas to US
-10Gb Amsterdam—Chicago
-10Gb London—Chicago
-3Gb CERN — Chicago

Canadian lambdas to US
-10Gb Chicago-Canada-NYC
-10Gb Chicago-Canada-Seattle

US lambda to Europe
-7Gb Chicago—Amsterdam


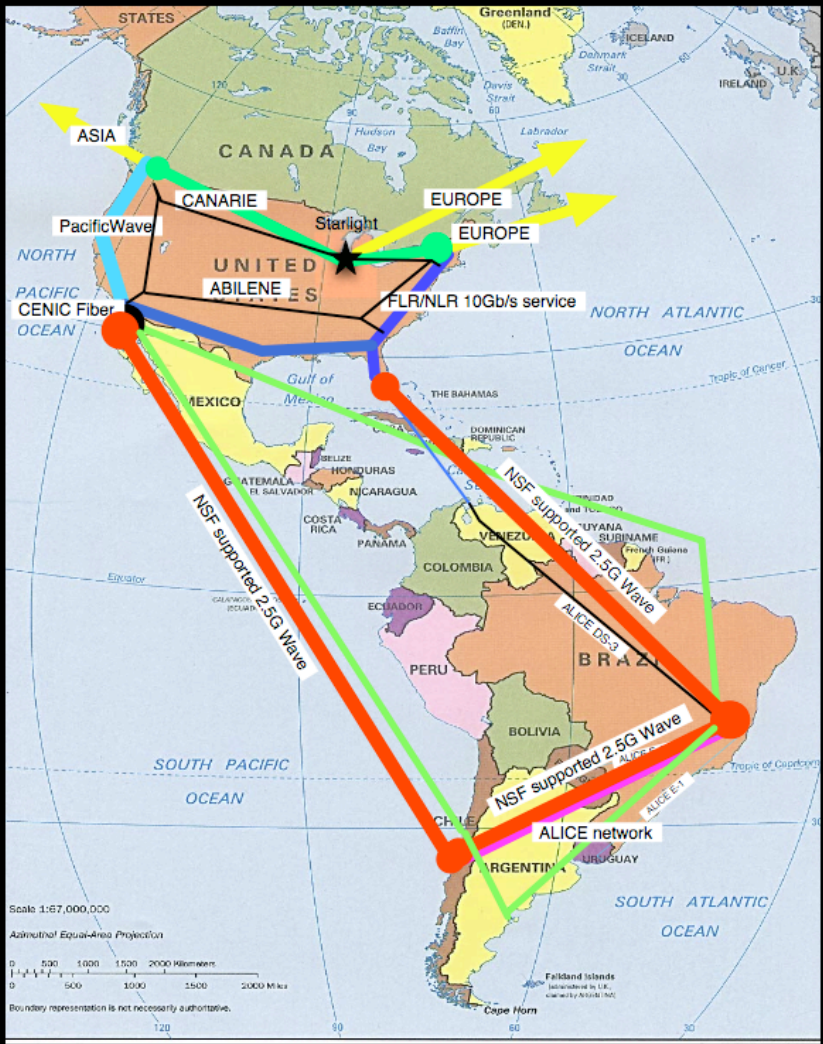
US/Japan lambda
-10Gb Chicago—Tokyo

European lambdas
-10Gb Amsterdam—CERN
-2.5Gb Prague—Amsterdam
-2.5Gb Stockholm—Amsterdam
-10Gb London—Amsterdam

IEEAF lambdas (blue)
-10Gb NYC—Amsterdam
-10Gb Seattle—Tokyo

Source: DeFanti

2004 International Lambdas



European lambdas to US
 -10Gb Amsterdam—Chicago
 -10Gb London—Chicago
 -3Gb CERN — Chicago

Canadian lambdas to US
 -10Gb Chicago-Canada-NYC
 -10Gb Chicago-Canada-Seattle

US lambda to Europe
 -7Gb Chicago—Amsterdam

US/Japan lambda
 -10Gb Chicago—Tokyo

European lambdas
 -10Gb Amsterdam—CERN
 -2.5Gb Prague—Amsterdam
 -2.5Gb Stockholm—Amsterdam
 -10Gb London—Amsterdam

IEEAF lambdas (blue)
 -10Gb NYC—Amsterdam
 -10Gb Seattle—Tokyo

Source: DeFanti




cdt What do we want to do with this?

- Collaboration
 - Distributed communities share resources: GWEN, GEON, NEES, etc.
 - “Shared Cyberinfrastructure”: data, code, tools, resources, simulations...
- “Standard Things”
 - Task Farming, Resource Brokering, Remote Steering, Managing Data (!)
- “New Scenarios”
 - Apps abstracted, become services
 - Dynamic apps decide their future, find their services (data, resources, applications, people): DDDAS
 - Apps distributed, spawned, task farmed, controlled, monitored with other apps/people



cct

Collaborations for Complex Problems



NASA Neutron Star Grand Challenge

- 5 US Institutions
- Attack colliding neutron star problem



NSF Black Hole Grand Challenge

- 8 US Institutions
- 5 years
- Attack colliding black hole problem



EU Astrophysics Network

- 10 EU Institutions
- 3 years
- Continue these problems

Examples of Future of Science & Engineering

- Require large scale data, simulations, beyond reach of any machine
- Require large geo-distributed cross-disciplinary collaborations
- Require Grid technologies



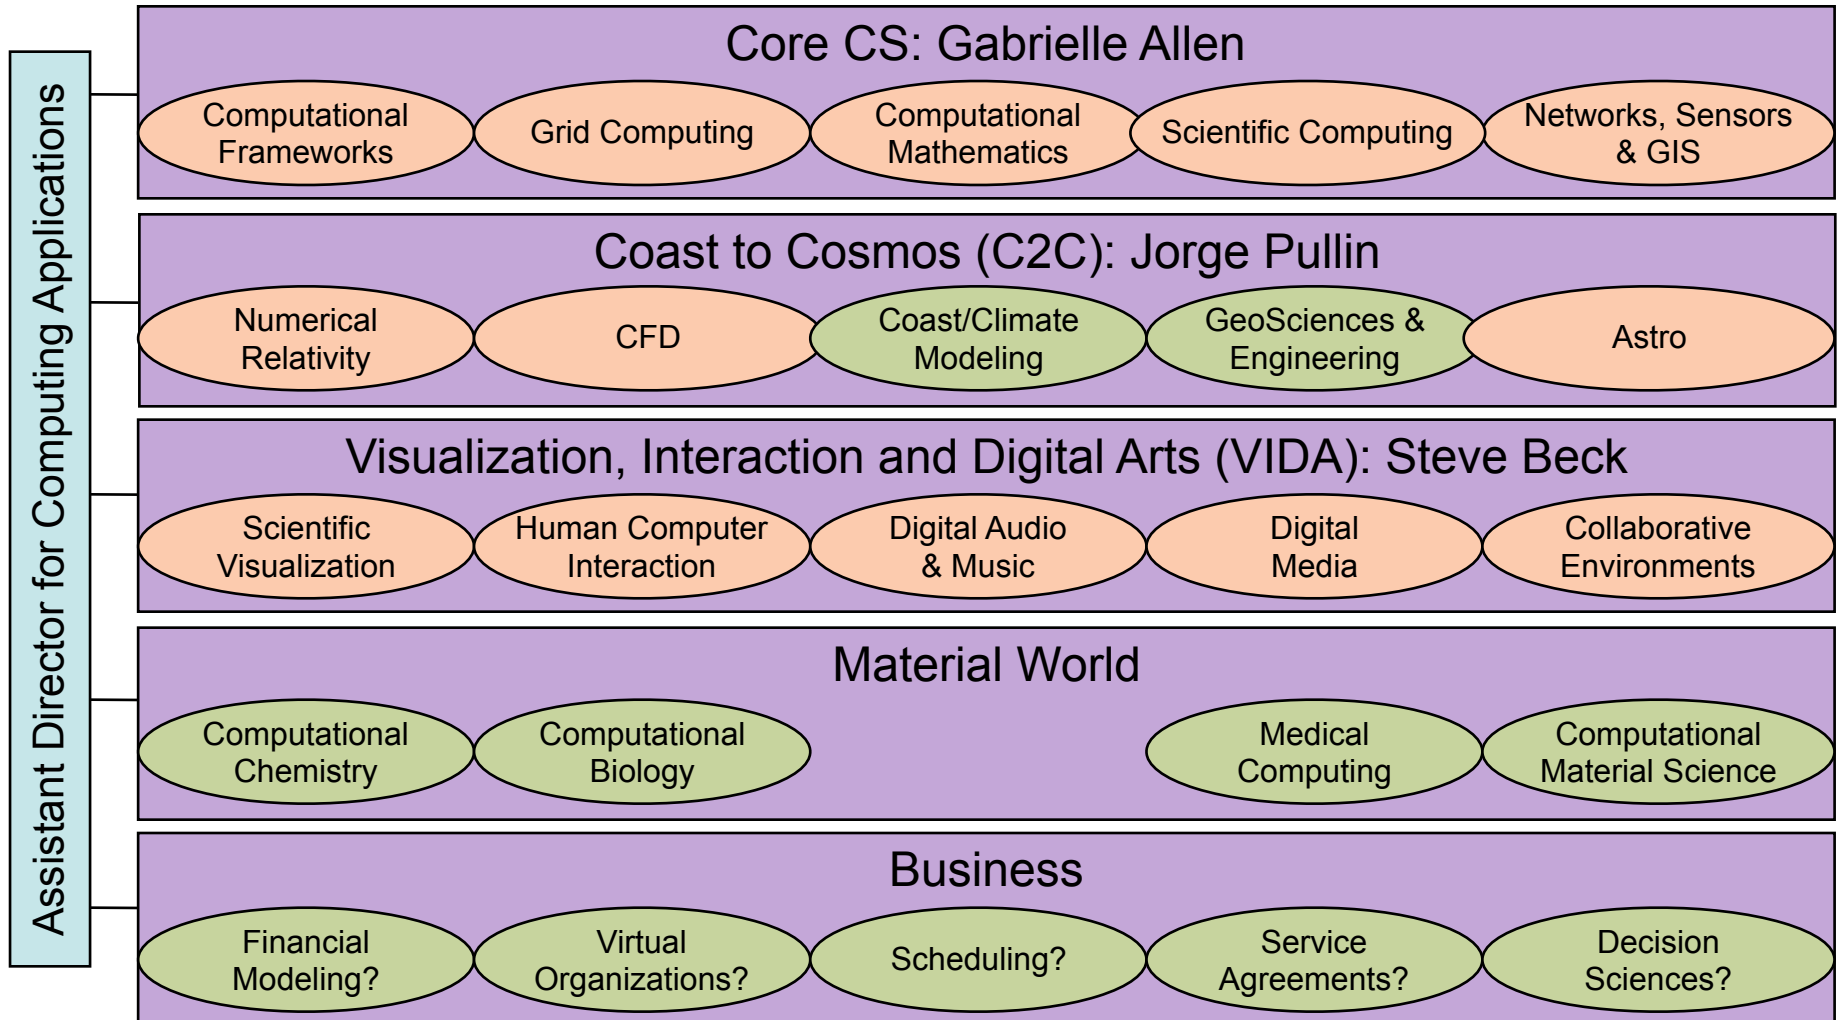
Recommendations from PITAC



- Universities must significantly change organizational structures: multidisciplinary & collaborative research to remain competitive in global science.
- Federal investments must rebalance to:
 - *Software*: create reliable, easy to use, scalable software that will enable scientists to focus on discovery: *Software Crisis*
 - *Hardware*: develop, prototype, evaluate new hardware architectures to deliver larger peak and sustained performance at the petaflop level for scientific apps
 - *Data*: focus on data-intensive solutions to address the coming data explosion with advances in sensors and sensor networks

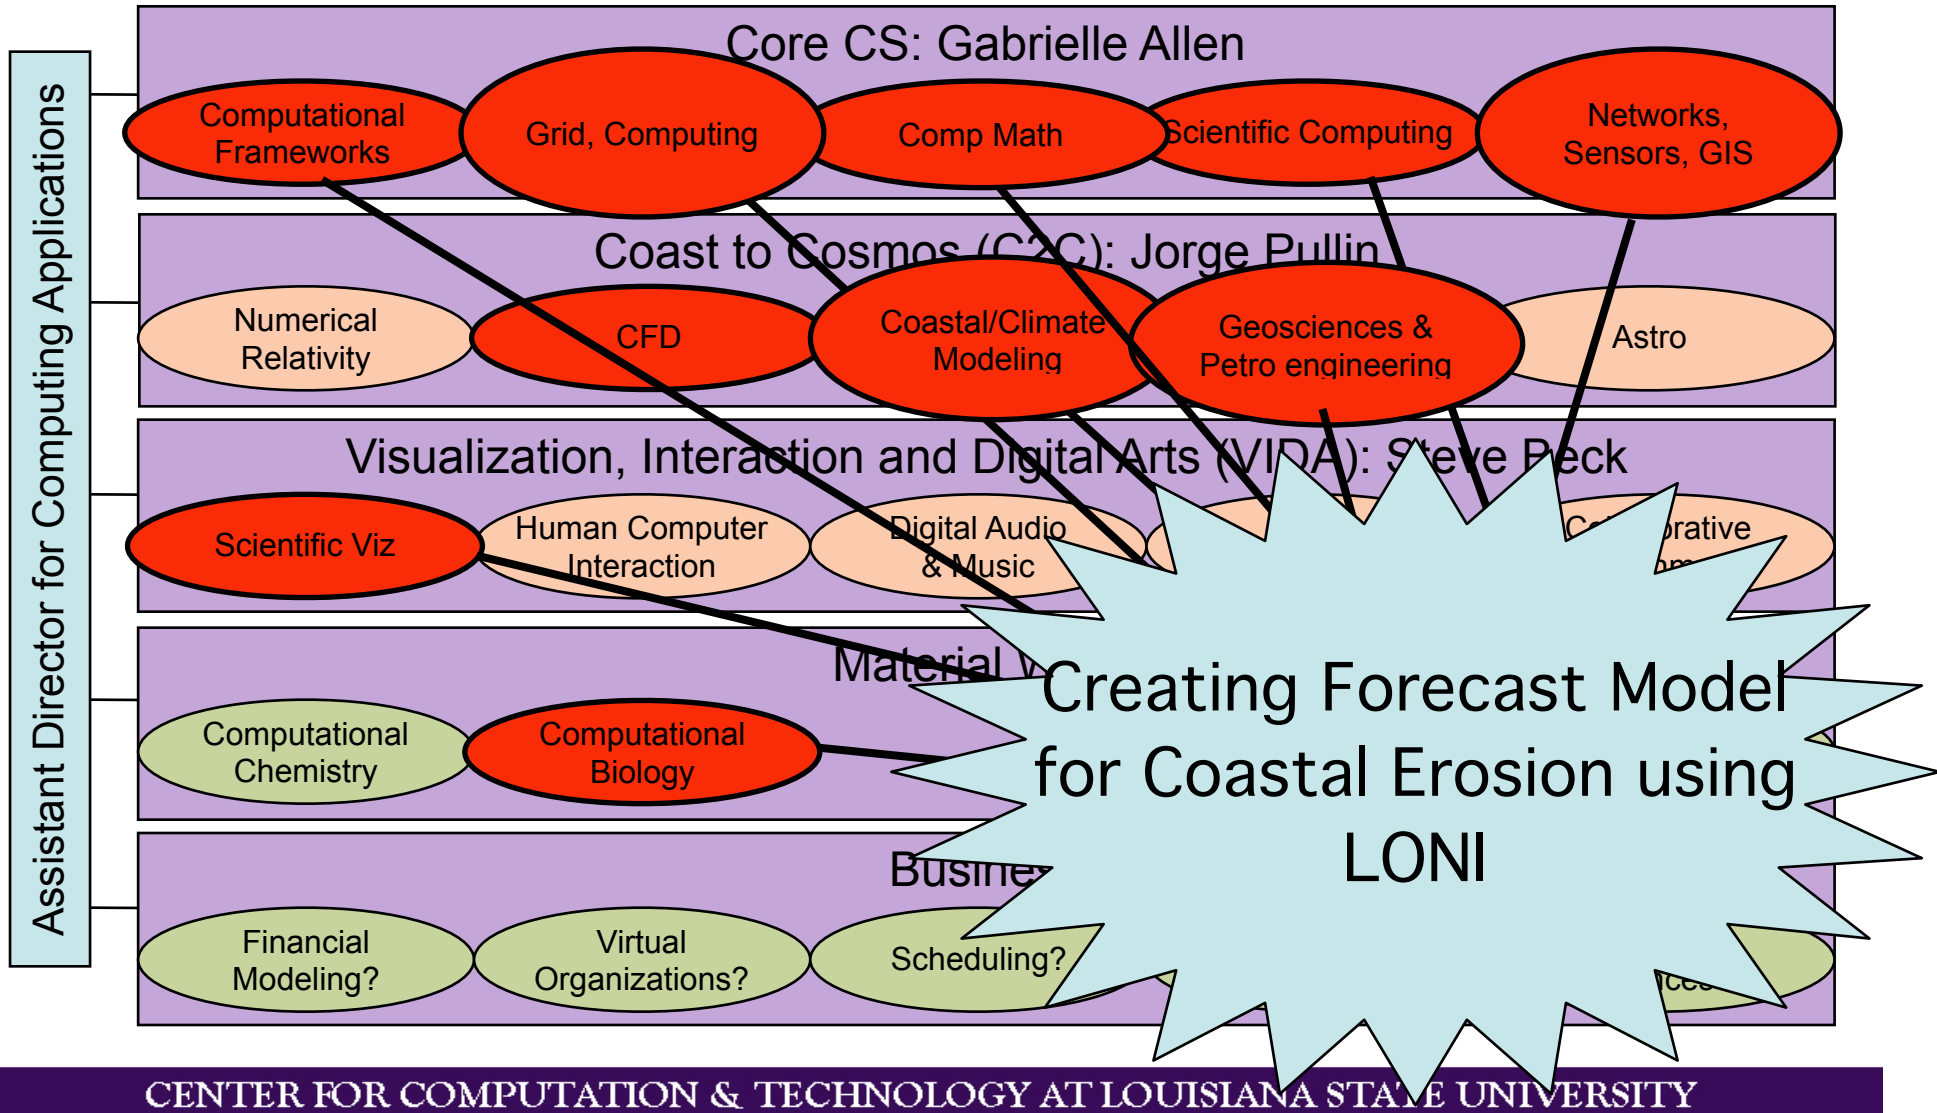


The NCSA-CCT Interdisciplinary Model






The NCSA-CCT Interdisciplinary Model





cct

Huge Black Hole Collision Simulation





cct

Huge Black Hole Collision Simulation





cct Issues for Complex Simulations

- Huge amounts of data needed/generated across different sites
 - Smarr: supercomputers are merely petabyte generators
 - How to retrieve, track, manage data across Grid? Data Archives! Metadata!
 - In this case, had to fly Berlin to NCSA, bring data back on disks!
- Many components developed by distributed collaborations
 - How to bring communities together?
 - How to find/load/execute different components?
- Many computational resources available
 - How to find best ones to start?
 - How to distribute work effectively?
- Needs of computations *change* with time!
 - How to adapt to changes?
 - How to monitor system?
- How to interact with experiments? Coming! Grids evolve: DDDAS
- Complex infrastructure accessed through common toolkits
 - Everyone should not reinvent the wheel (especially physicists)



Cactus Framework: Enabling User Communities in Advanced HPC

www.cactuscode.org

- Toolkit for HPC, Grid, Collaborative Applications
- Abstract interfaces for everything
 - Parallelism, I/O, AMR, elliptic solvers, etc
 - Other packages, toolkits (Grace, Petsc, Samrai, HDF, Carpet, etc)
- Advanced capabilities
 - Streaming data, communication, Viz
- Grid capabilities
 - Globus, GAT, Portals, etc
- Toolkits: Einstein Toolkit, CFD Toolkit, Bioinformatics Toolkit
- International developer/user base
 - We have openings for developers!!



Cactus Framework: Enabling User Communities in Advanced HPC

www.cactuscode.org

Numerical Relativity/Astrophysics

Other Applications

AEI numrel

WashU

UNAM

Soton

RIKEN

Chemical Engineering
(U.Kansas)

Goddard

Monash

Thessaloniki

Pitt

PSU

Climate Modeling
(NASA, Utrecht)

Tübingen

TAC

SISSA

Portsmouth

UT

EU Astrophysics
Network

NASA

Arizona

Caltech

Bio-Informatics
(Chicago)

Brownsville

S. Africa

Early Universe
(LBL)

LSU

UNAM

Benchmarking

CFD
(KISTI, LSU)

Garching

CS groups
worldwide



Plasma
(Princeton)

Astrophysics
(Zeus/Innsbruck)



Dynamic Adaptive Distributed Computation

(T.Dramlitsch, with Argonne/U.Chicago)



SDSC IBM SP
1024 procs
 $5 \times 12 \times 17 = 1020$

NCSA Origin Array
 $256 + 128 + 128$
 $5 \times 12 \times (4 + 2 + 2) = 480$



These experiments:

- Einstein Equations (but could be any Cactus application)

Achieved:

- First runs: 15% scaling
- With new techniques: 70-85% scaling, ~ 250GF

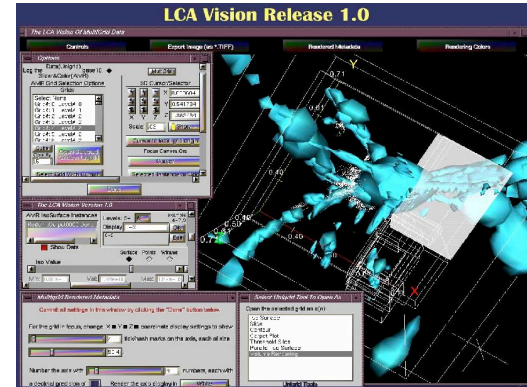
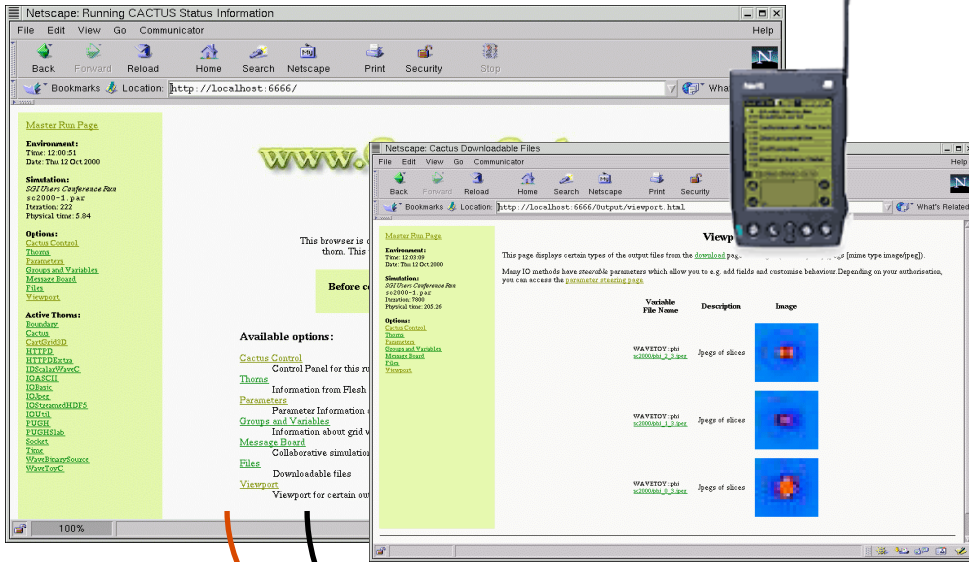
Dynamic Adaptation: Number of ghostzones, compression, ...

Won

"Gordon Bell Prize" (Supercomputing 2001, Denver)



Interacting with Jobs using Cactus




Any Viz Client:
LCA Vision, OpenDX
Changing any steerable parameter

- Parameters
- Physics, algorithms
- Performance

Remote Viz data
Streaming HDF5
Autodownload

Remote Viz data
Amira





Notification and Information



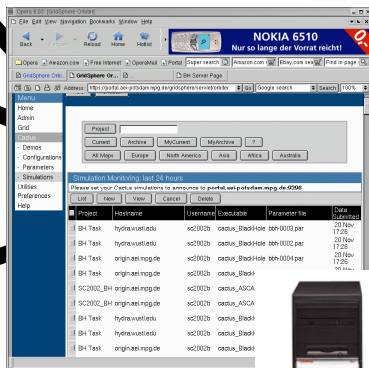
Replica Catalog



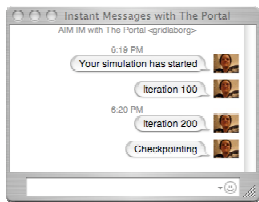





SMS Server




GridSphere Portal



"The Grid"



Mail Server





Computers as Petabyte Generators

(Black holes and more...)

- Crude vacuum BH collisions today
 - 10^{15} Flops, Tbytes output, vastly downsampled!
- Desired BH, NS scenarios in few years
 - 10^{20} Flops, multiple orbits, adaptive meshes, hydro, 1 day on 50TF machine, 25TB+ output per run
 - May be physically distributed due to grid activities
- Need metadata to describe simulations
- Real time scheduling across multiple resources for distributed computing
 - Lambda provisioning on demand: spawning (analysis, steering), migration, interactive viz from distributed collaborations
- Parameter Space! 10^3 - 10^6 simulations!



New Grid Applications: be creative

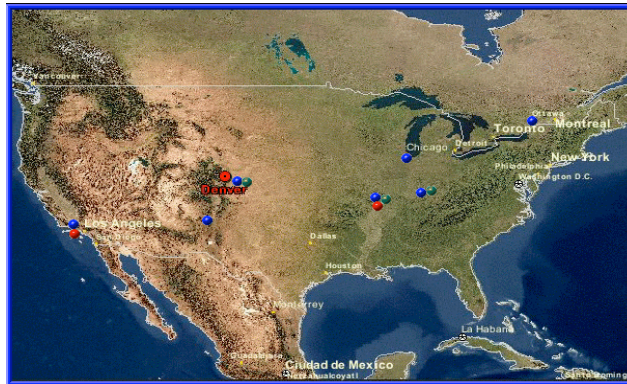
- Intelligent Parameter Surveys, Monte Carlos
 - May control other simulations!
- Dynamic Staging: move to faster/cheaper/bigger machine ("Grid Worm")
 - Need more memory? Need less?
- Multiple Universe: clone to investigate steered parameter ("Grid Virus")
- Automatic Component Loading
 - Needs of process change, discover/load/execute new component somewhere
- Automatic "Look Ahead", convergence testing
 - spawn off and run coarser resolution to predict future, study convergence
- Spawn Independent/Asynchronous Tasks
 - send to cheaper machine, main simulation carries on
- Routine Profiling
 - best machine/queue, choose resolution parameters based on queue
- Dynamic Load Balancing: inhomogeneous loads, multiple grids
- DDDAS: injecting data into the above, feed back to experiment
- Model-Model coupling: very important in DDDAS



Already did a lot of this!

Spawning on ARG Testbed

Main Cactus BH Simulation starts here



User only has to invoke “Spawner” thorn...

All analysis tasks spawned automatically to free resources worldwide

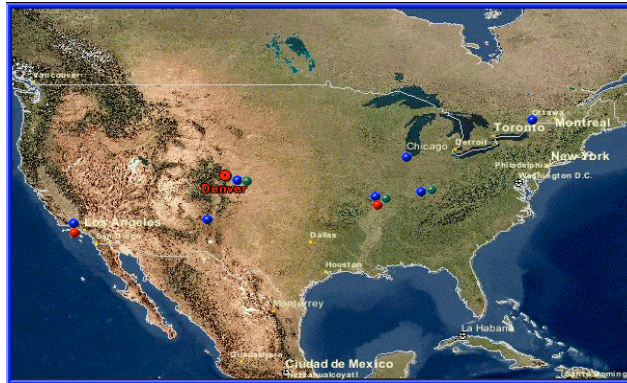





Already did a lot of this!

Spawning on ARG Testbed

Main Cactus BH Simulation starts here



User only has to invoke
“Spawner” thorn...

All analysis tasks spawned
automatically to free resources
worldwide

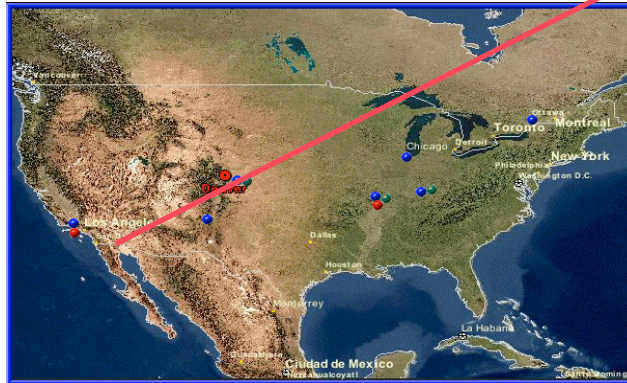
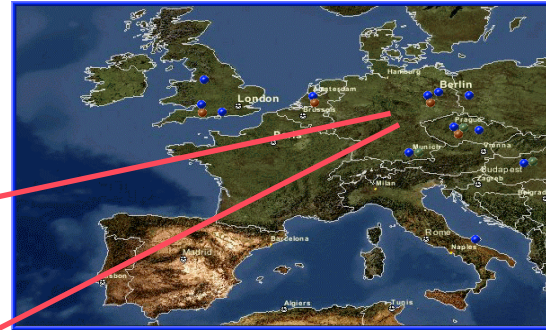




Already did a lot of this!

Spawning on ARG Testbed

Main Cactus BH Simulation starts here



User only has to invoke
“Spawner” thorn...

All analysis tasks spawned
automatically to free resources
worldwide

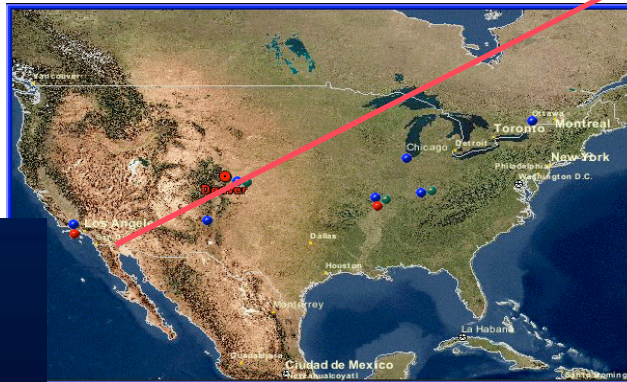
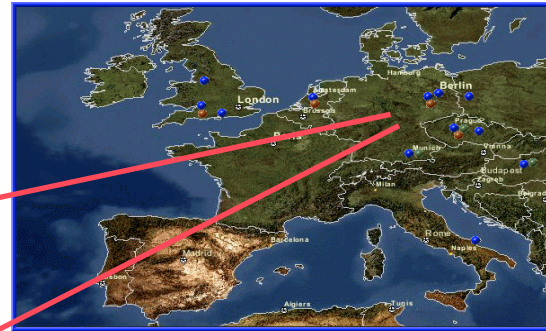





Already did a lot of this!

Spawning on ARG Testbed

Main Cactus BH Simulation starts here



All analysis tasks spawned automatically to free resources worldwide



User only has to invoke “Spawner” thorn...



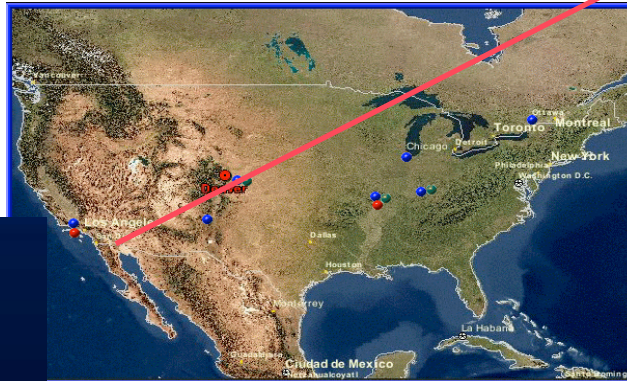
Already did a lot of this!

Spawning on ARG Testbed


Main Cactus BH Simulation starts here



All analysis tasks spawned automatically to free resources worldwide



User only has to invoke "Spawner" thorn...





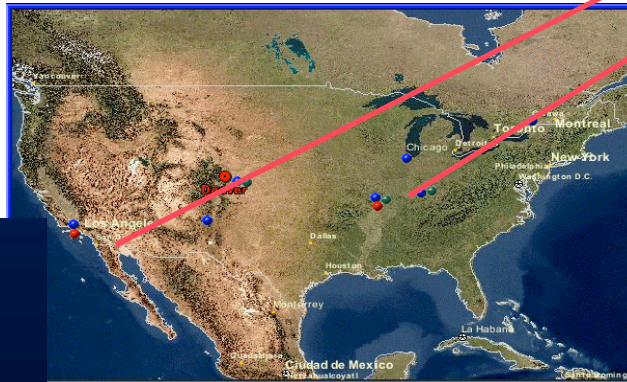
Already did a lot of this!

Spawning on ARG Testbed

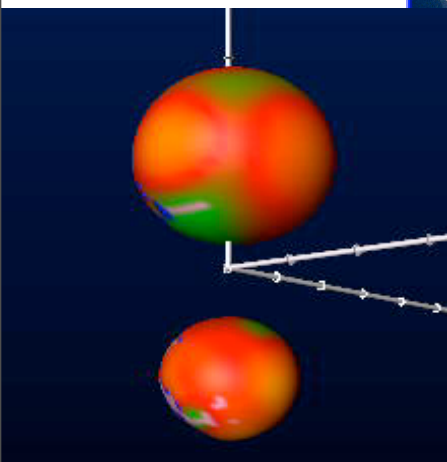

Main Cactus BH Simulation starts here



All analysis tasks spawned automatically to free resources worldwide

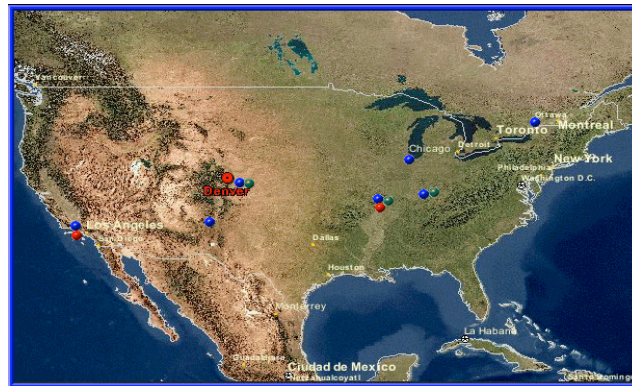


User only has to invoke "Spawner" thorn...





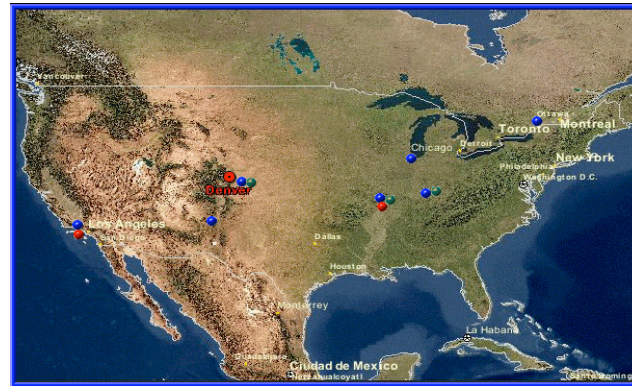
Task Farming, Spawning & Migration





Task Farming, Spawning & Migration

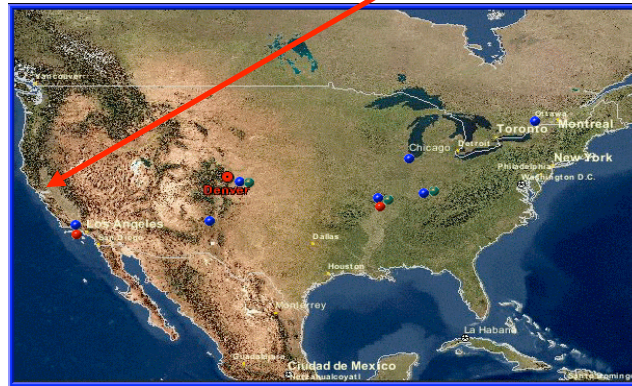
Main Cactus BH Simulation starts in Berkeley





Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



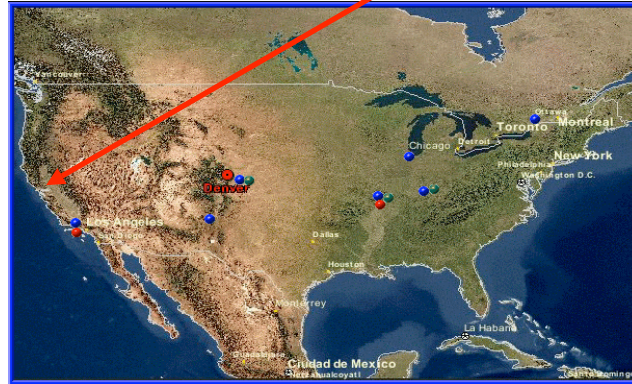


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters



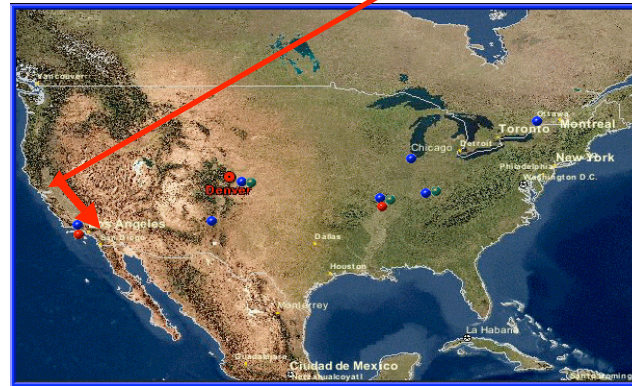


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley




Dozens of small jobs sent out to test parameters



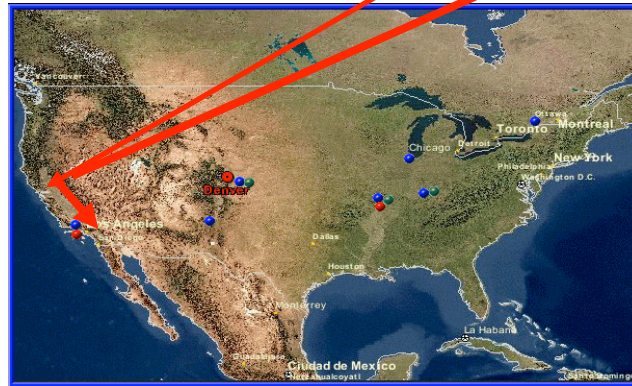


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters



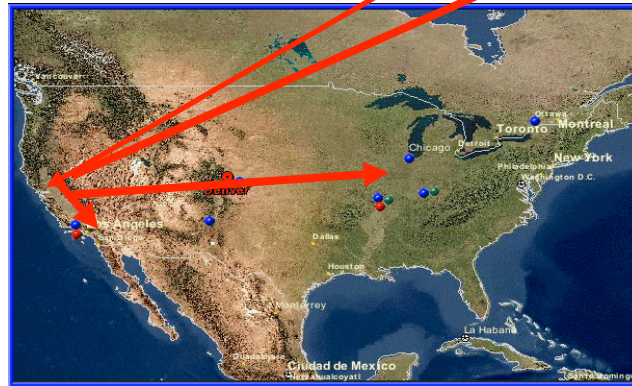


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters



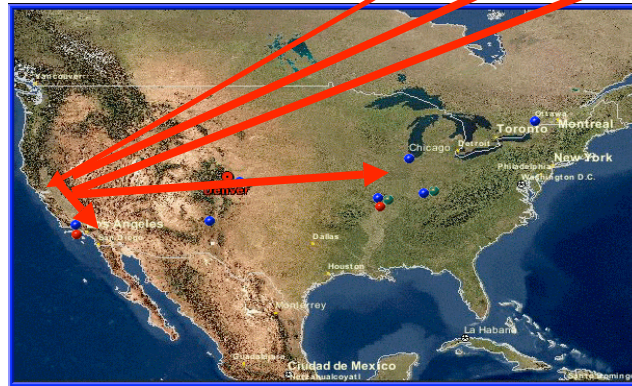


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters



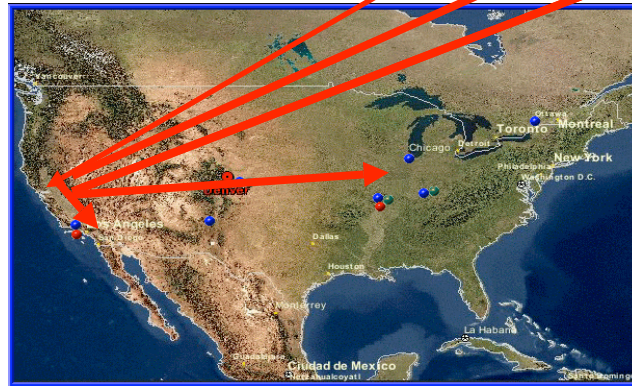


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



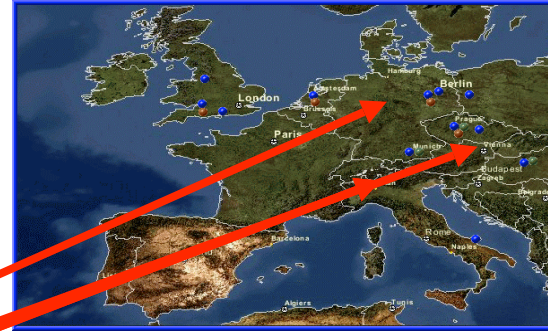
Dozens of small jobs sent out to test parameters



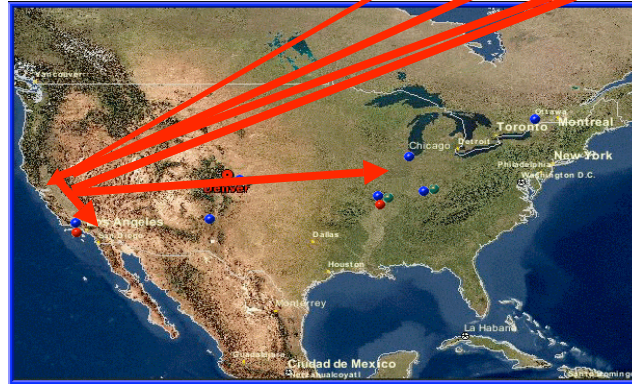


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley




Dozens of small jobs sent out to test parameters



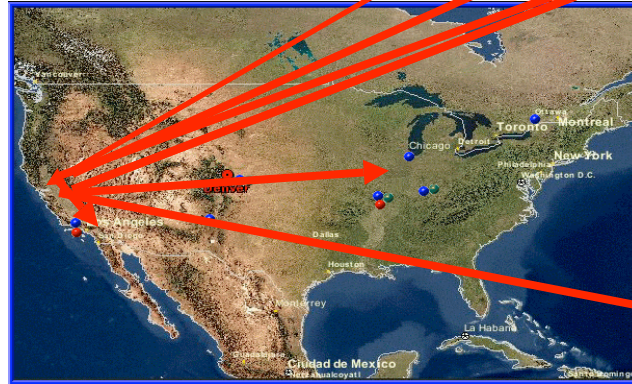


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters



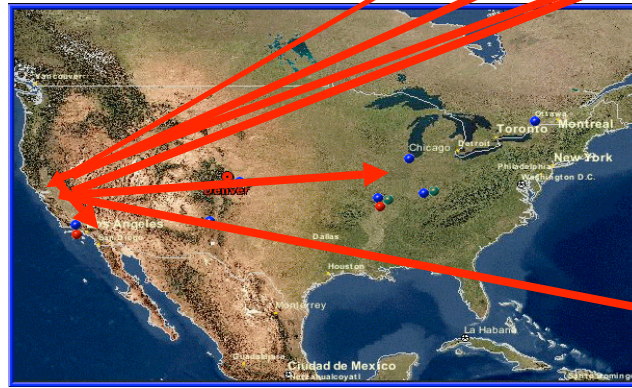


Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters



Data returned for main job



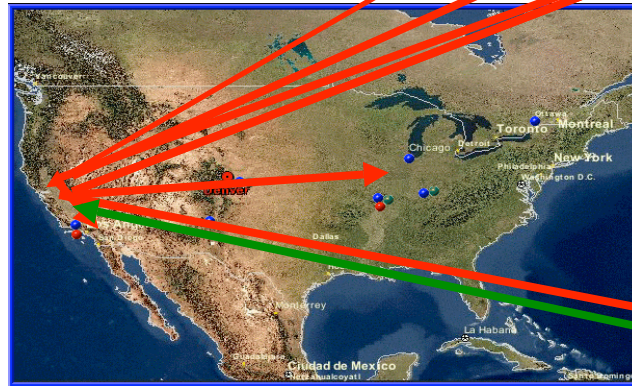


Task Farming, Spawning & Migration


Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters



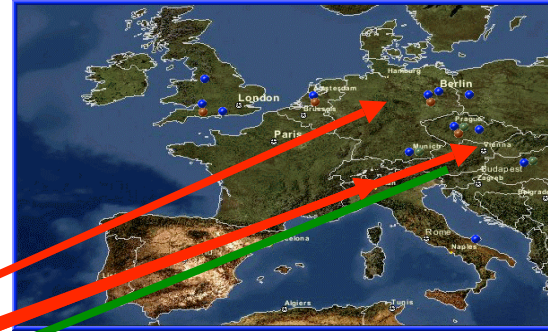
Data returned for main job



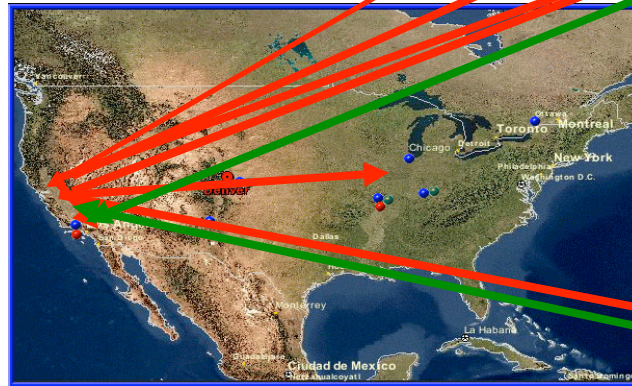


Task Farming, Spawning & Migration


Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters




Data returned for main job



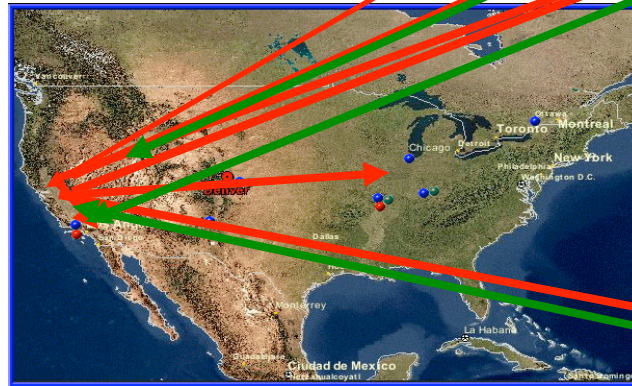


Task Farming, Spawning & Migration


Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters




Data returned for main job



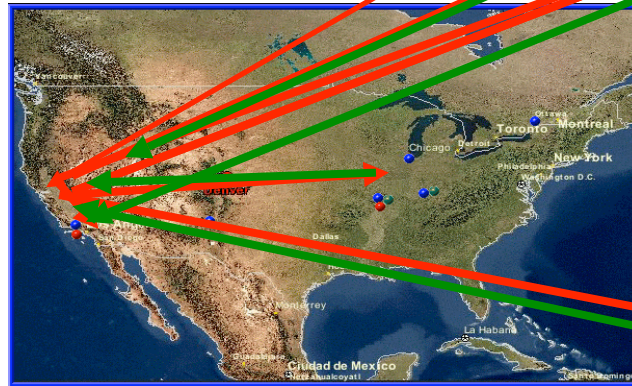


Task Farming, Spawning & Migration


Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters




Data returned for main job






Task Farming, Spawning & Migration


Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters




Data returned for main job
Huge job generates remote data
to be visualized in Baltimore



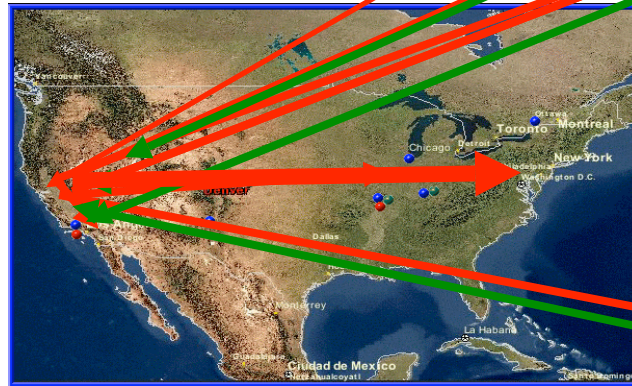


Task Farming, Spawning & Migration


Main Cactus BH Simulation starts in Berkeley



Dozens of small jobs sent out to test parameters




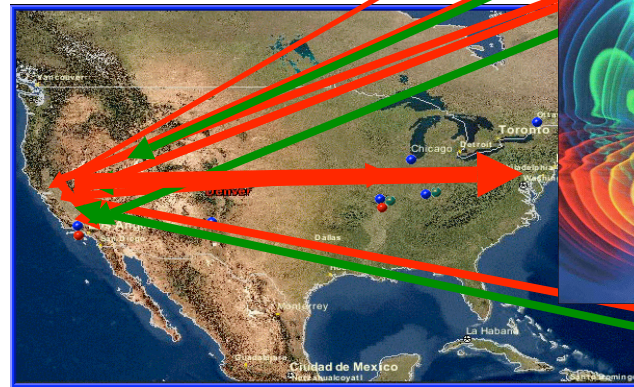
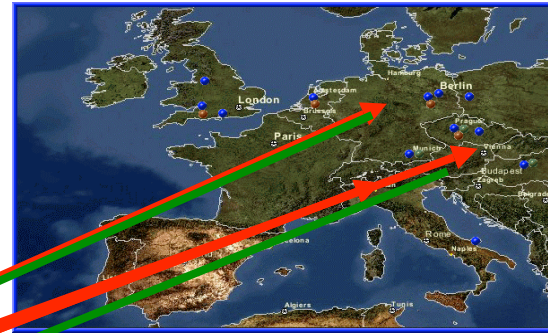
Data returned for main job
Huge job generates remote data to be visualized in Baltimore





Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley



tens of small jobs sent to test parameters

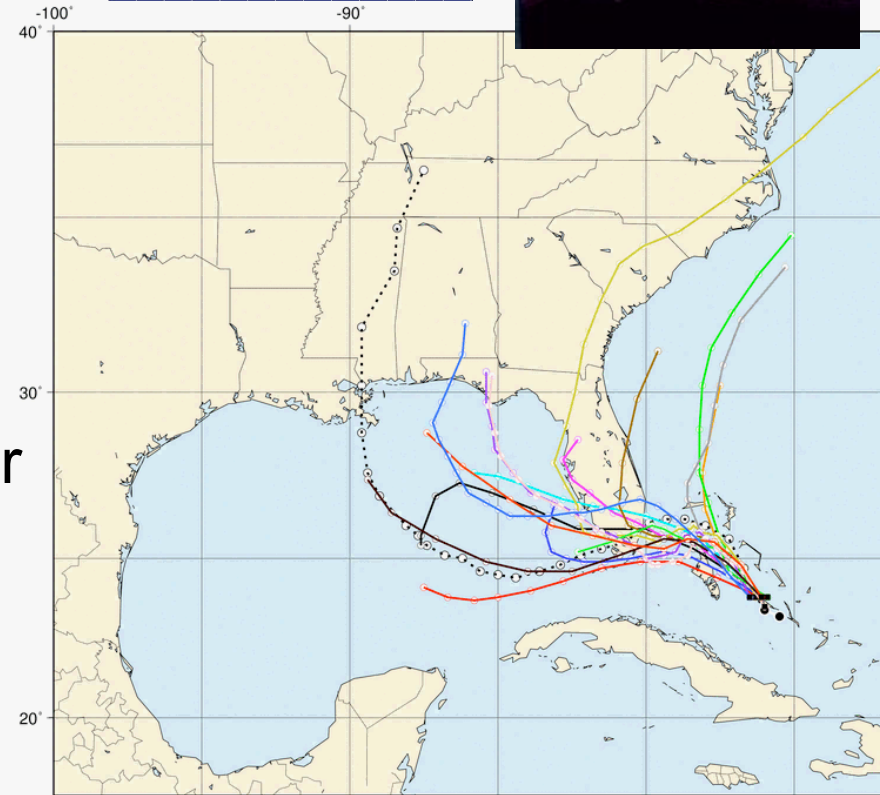
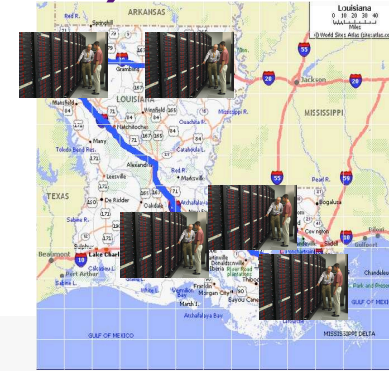


Data returned for main job
Huge job generates remote data to be visualized in Baltimore




DDDAS Scenarios (Gabrielle Allen)

- UCoMS (Petroleum Engineering)
 - Deploy sensor networks across Gulf
 - Data collected to provide input to simulations, tasks farmed out
 - Results collected (transmitted back)
 - <http://www.ucoms.org>
- SCOOP, DynaCode (Coastal Modeling)
 - Data coming in from sensors all over Gulf Realtime Operational Grid
 - Feeds in to models on Grid sites
 - Algorithms invoked dynamically

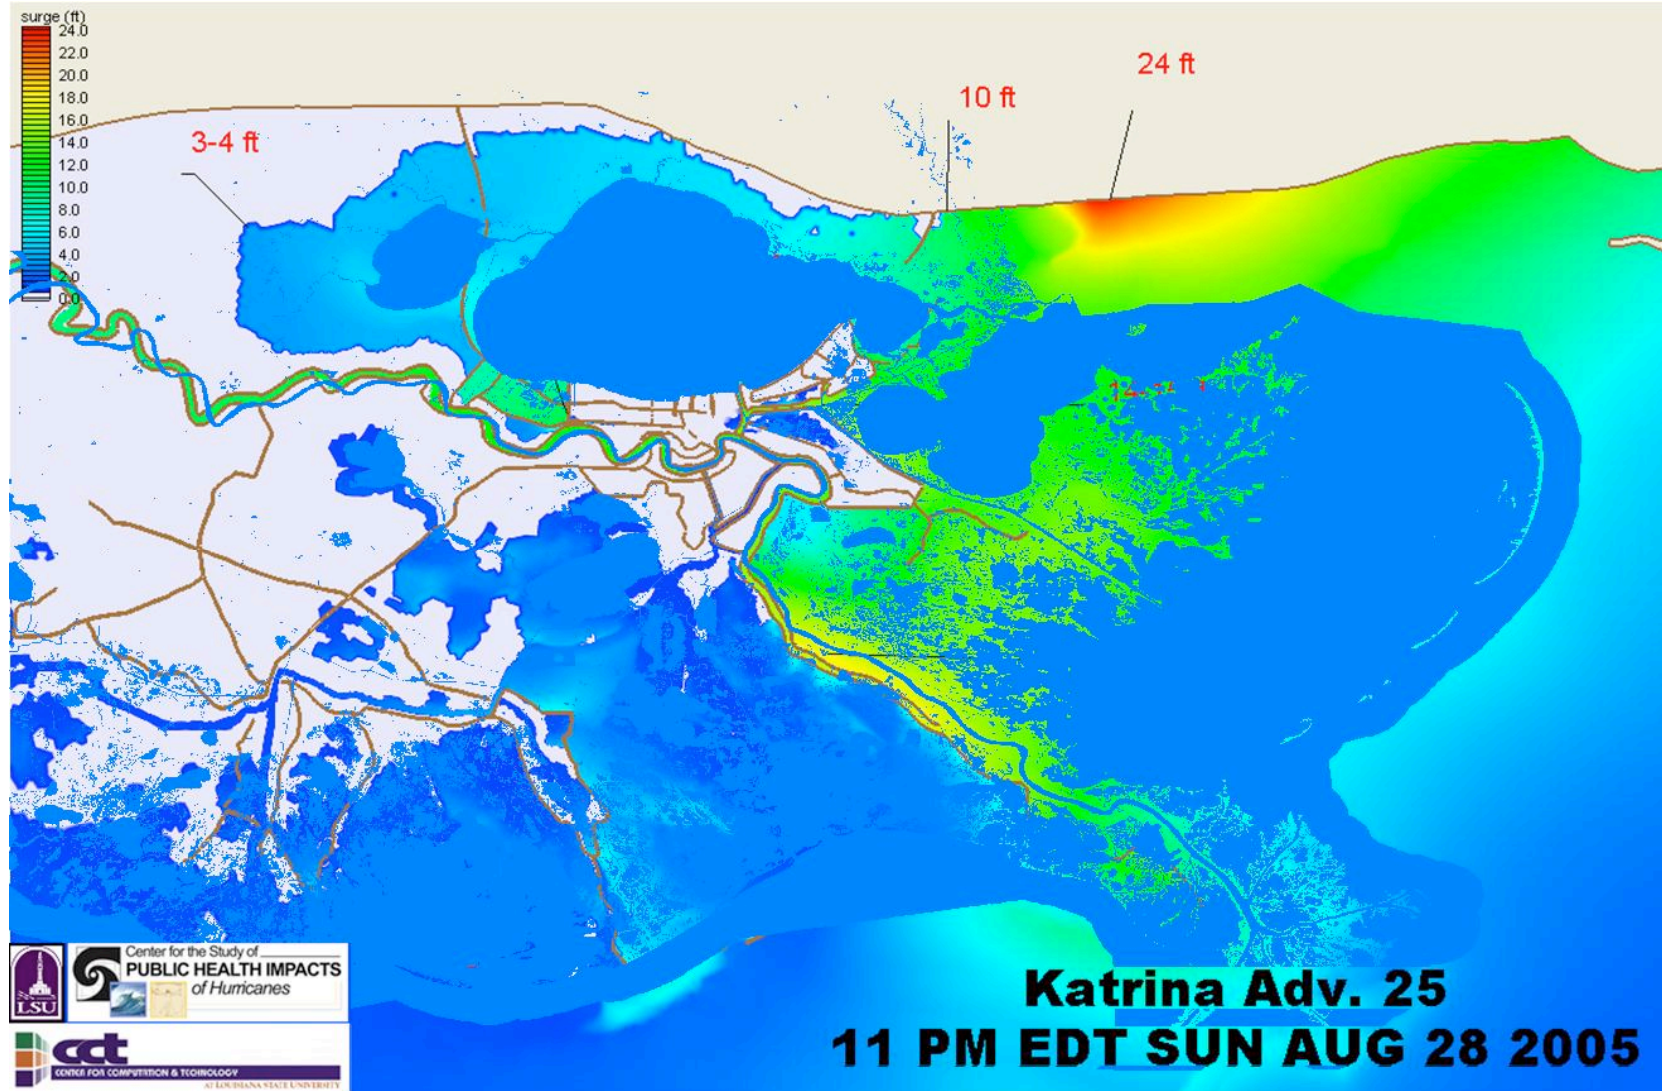






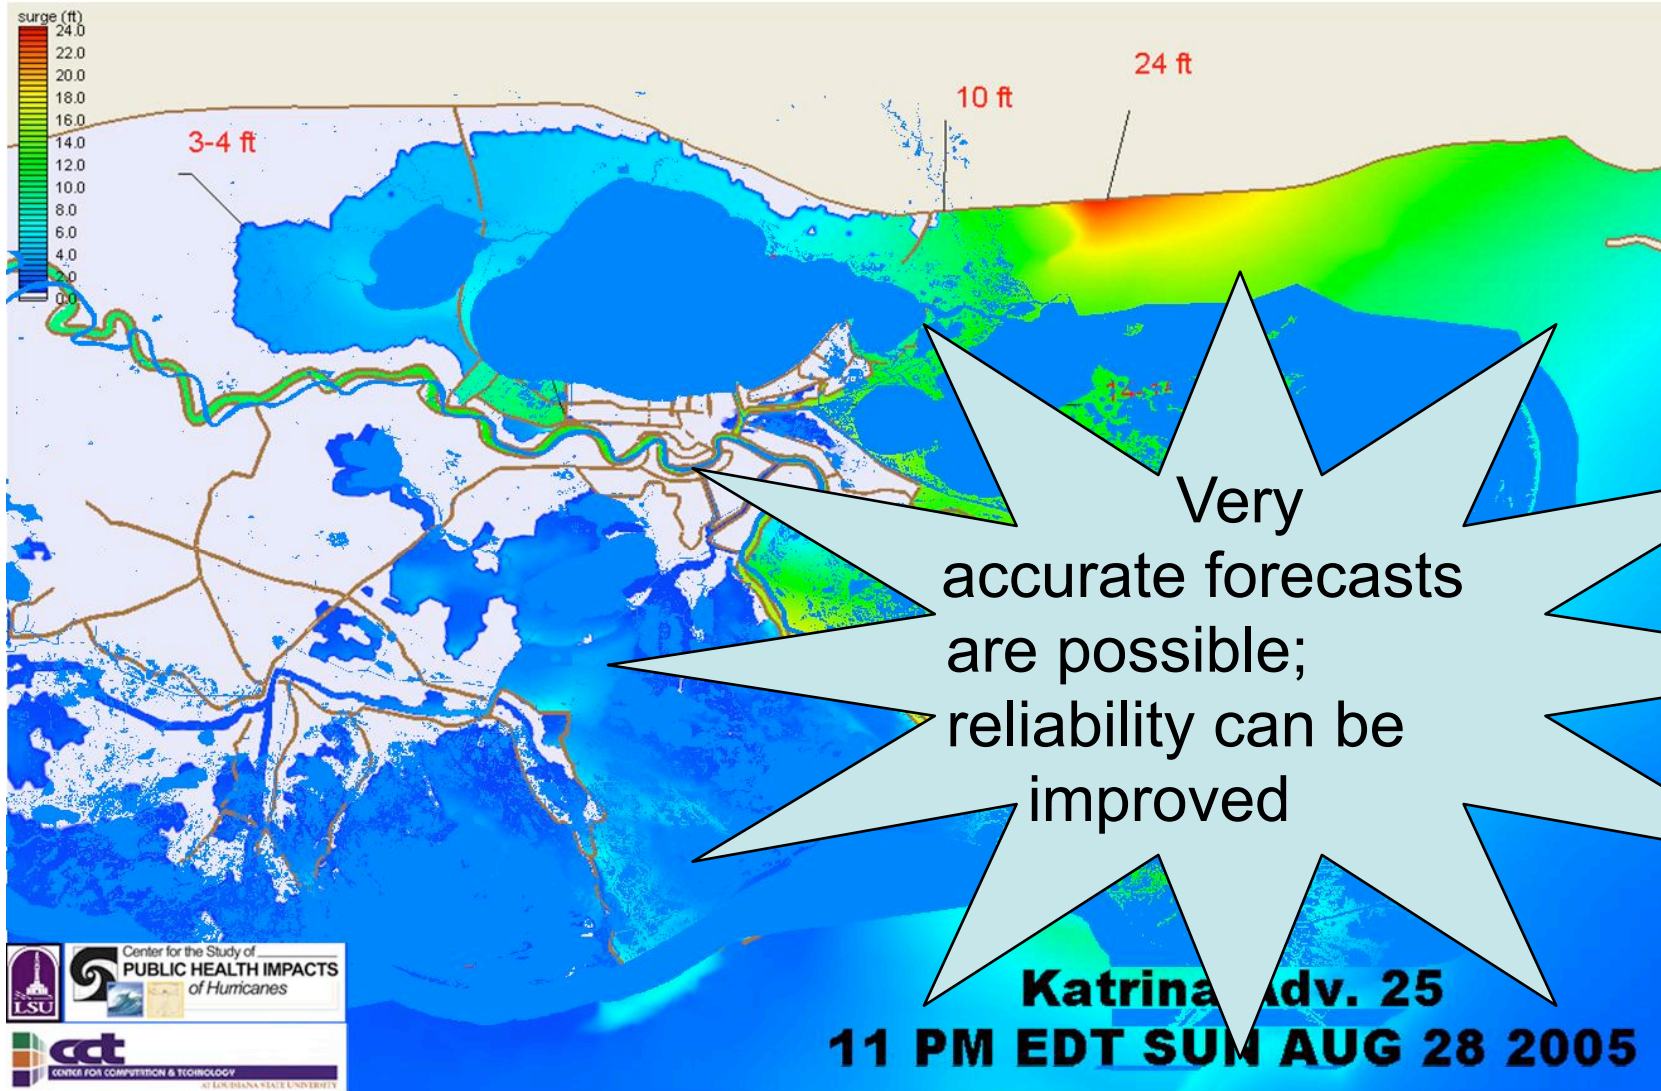


Emergency Forecast







Emergency Forecast





cct

Emergency Forecast



Used to notify officials, FEMA, State, Governor. Email, web, phone: can do better!

Very accurate forecasts are possible; reliability can be improved

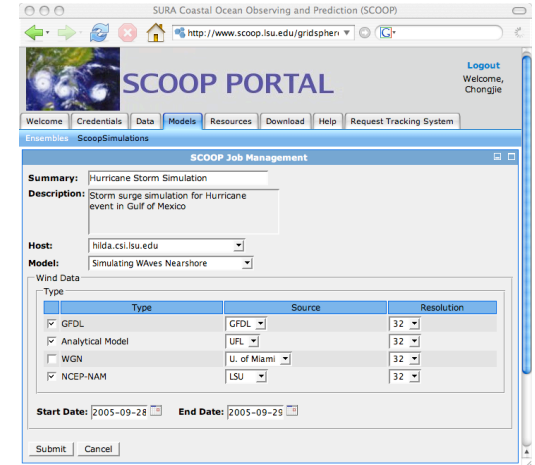


Katrina adv. 25
11 PM EDT SUN AUG 28 2005





SURA Coastal Ocean Observing Program (SCOOP)

- Integrating data from regional observing systems for realtime coastal forecasts in SE
- Coastal modelers, computer scientists
 - couple models
 - provide data solutions
 - deploy ensembles of models on the Grid
 - assemble realtime results with GIS



University of Alabama at Huntsville, University of Florida, GoMOOS, Louisiana State University, University of Miami, University of Maryland, University of North Carolina, Texas A&M, Virginia Inst of Marine Sciences


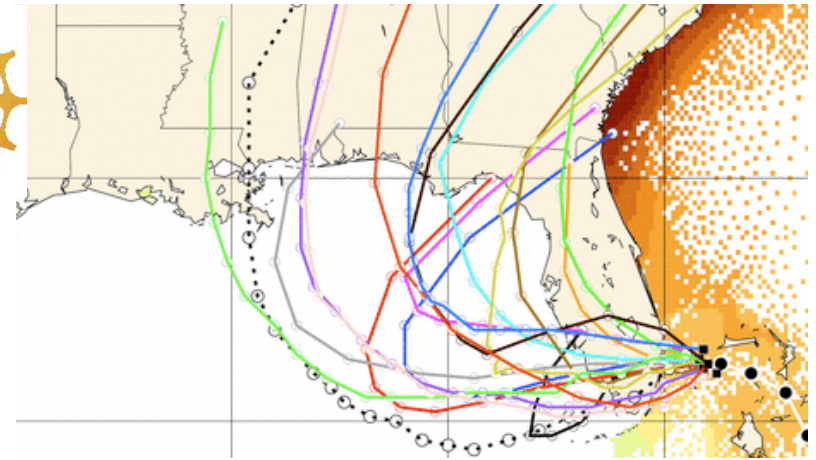






DynaCode



- Focus on scenarios:
 - Hurricane ensemble modeling
 - Coupling ocean circulation, storm surge, wave generation models for the Gulf
 - Notifications from NHC trigger customized ensemble hurricane models (surge/wind/wave), sensors verify, guide dynamic ensembles
 - Event driven, dynamic component framework with algorithm selection, optimization tools, workflow, data assimilation, result validation with sensor/satellite.
 - Ecological restoration and control
 - Breton Sound diversion, control structure to allow Mississippi to flow into wetlands
 - Coupled models (hydrodynamic, salinity, geomorphic, sediment) control diversion, sensors/wind fields inject real time data.




Dynamic, Event Driven Computing

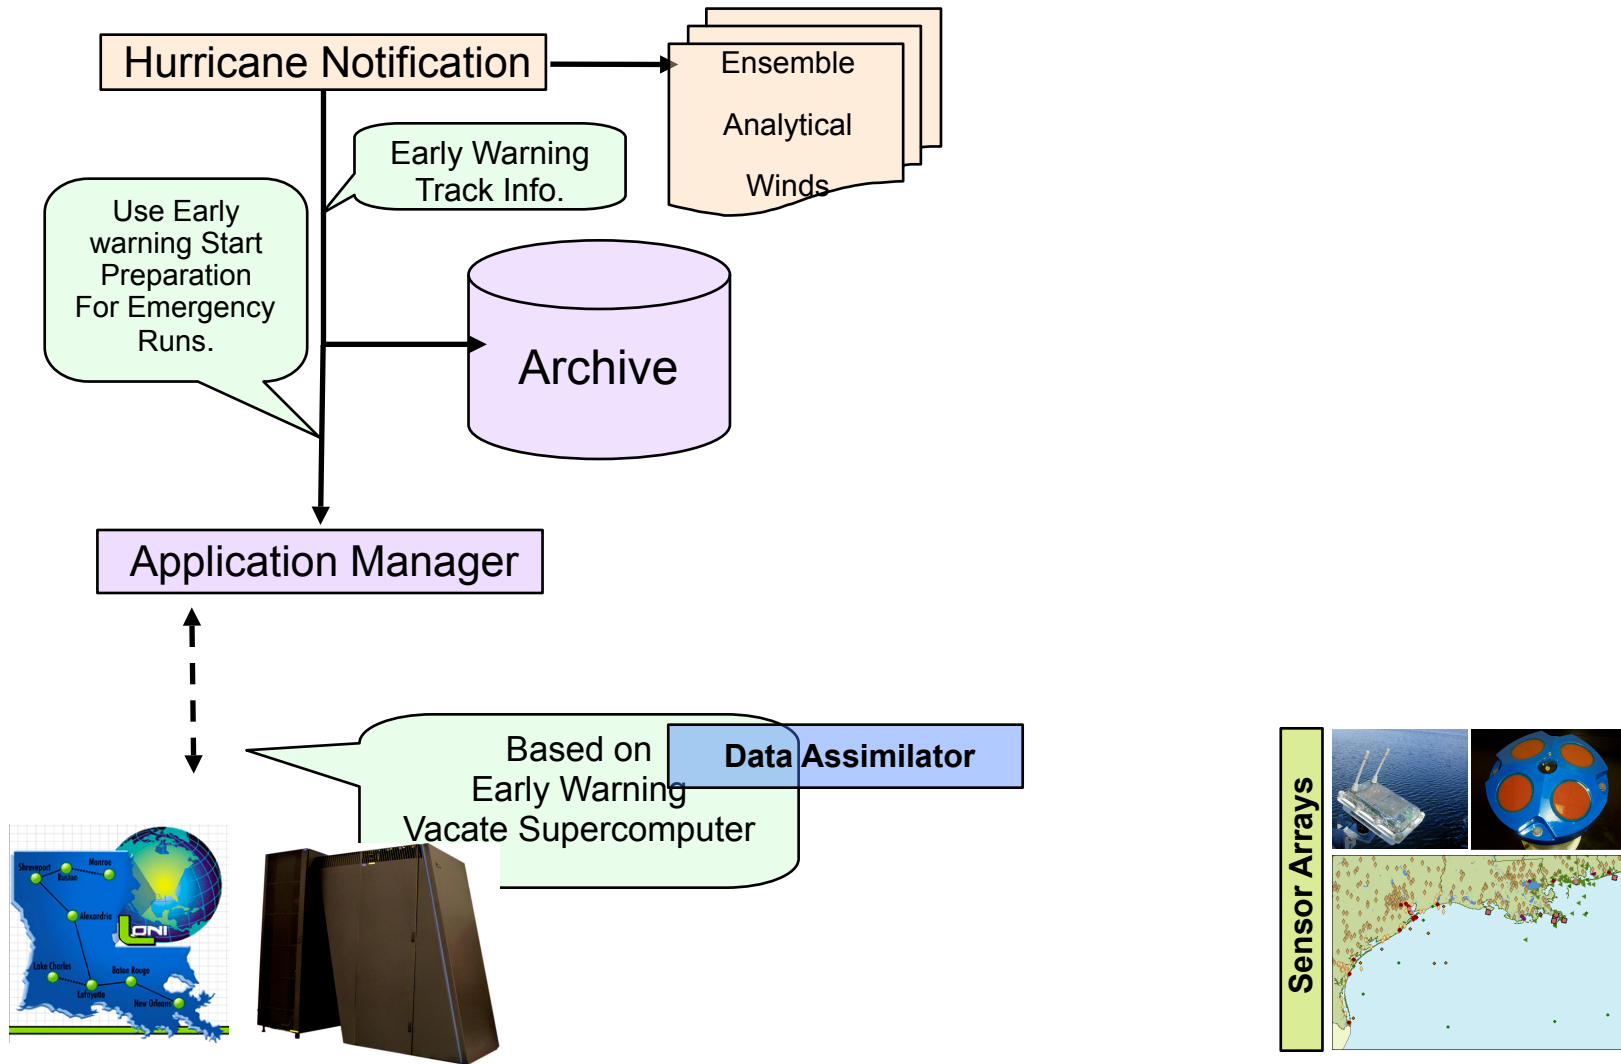


Application Manager

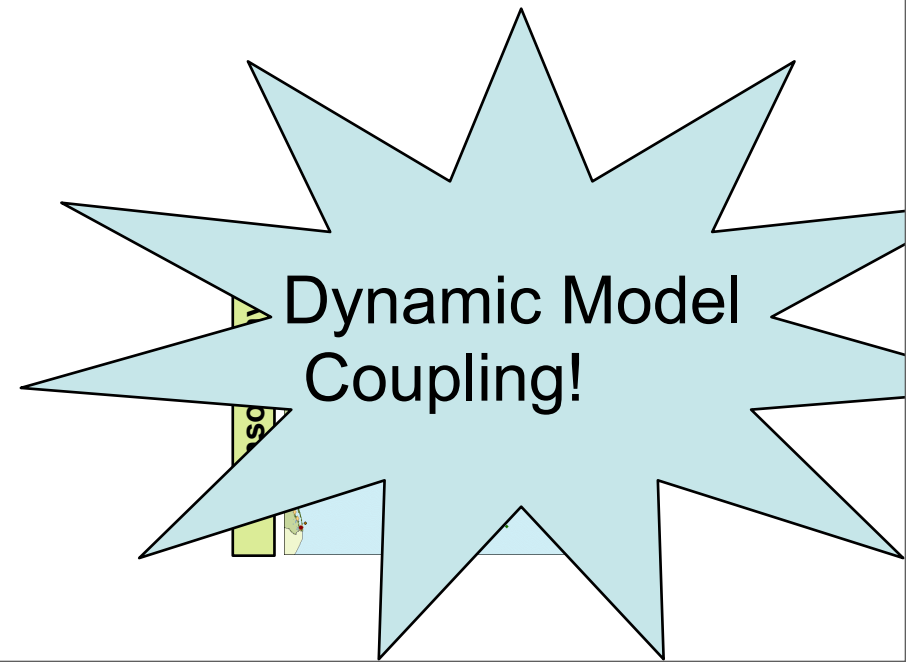
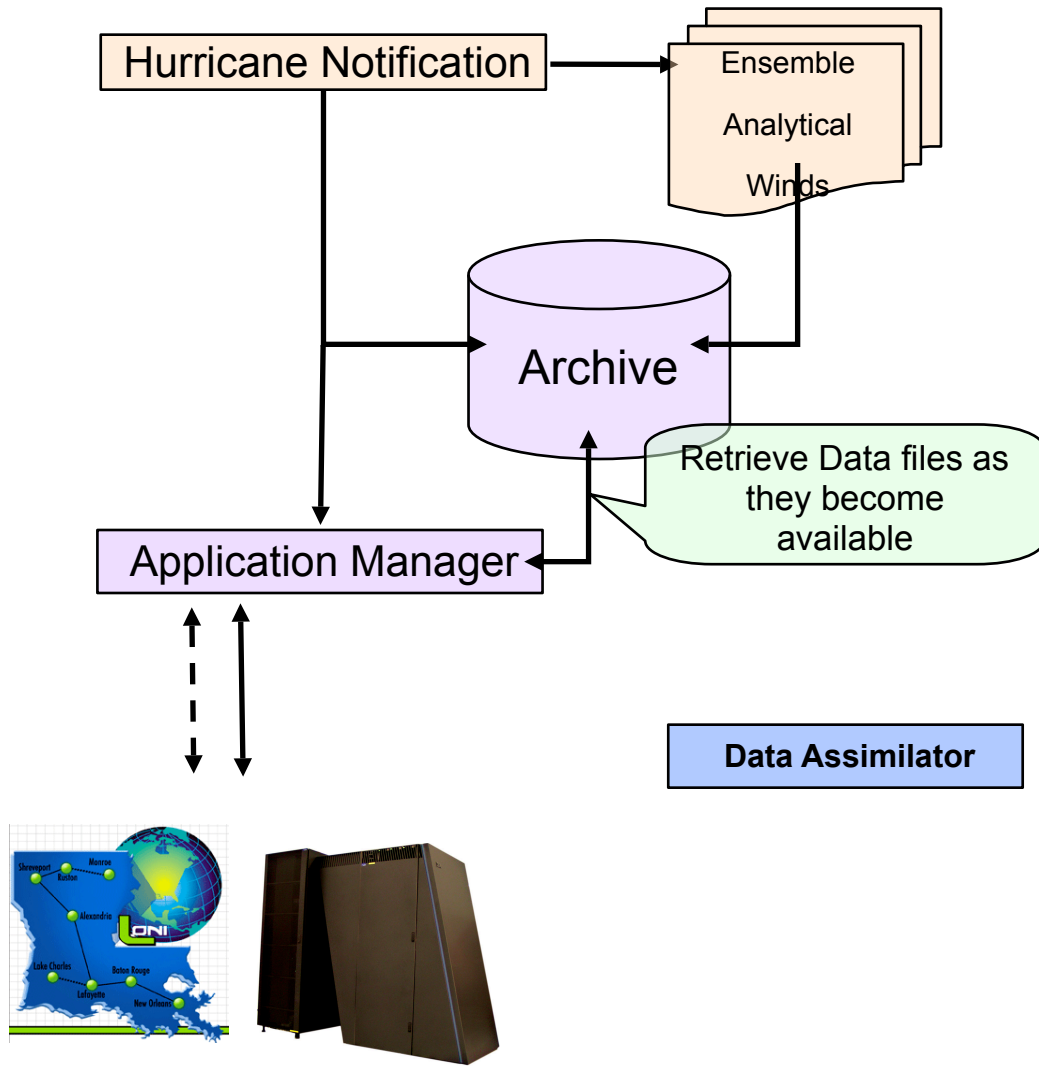
Data Assimilator



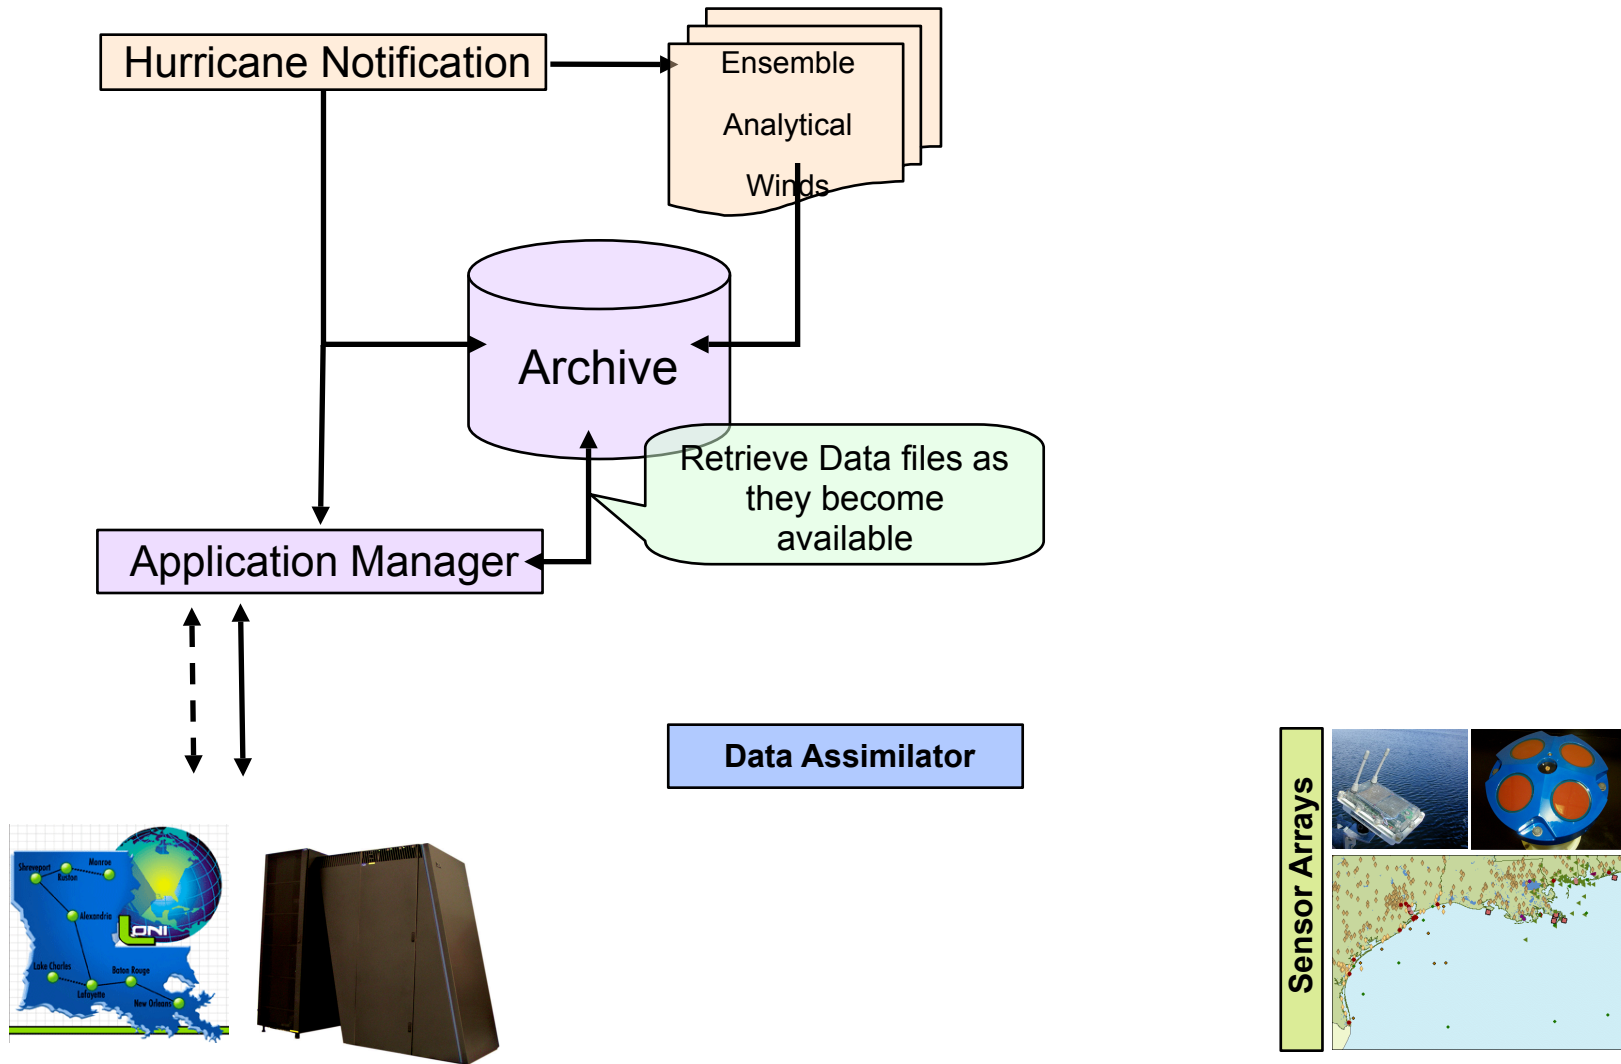
Dynamic, Event Driven Computing



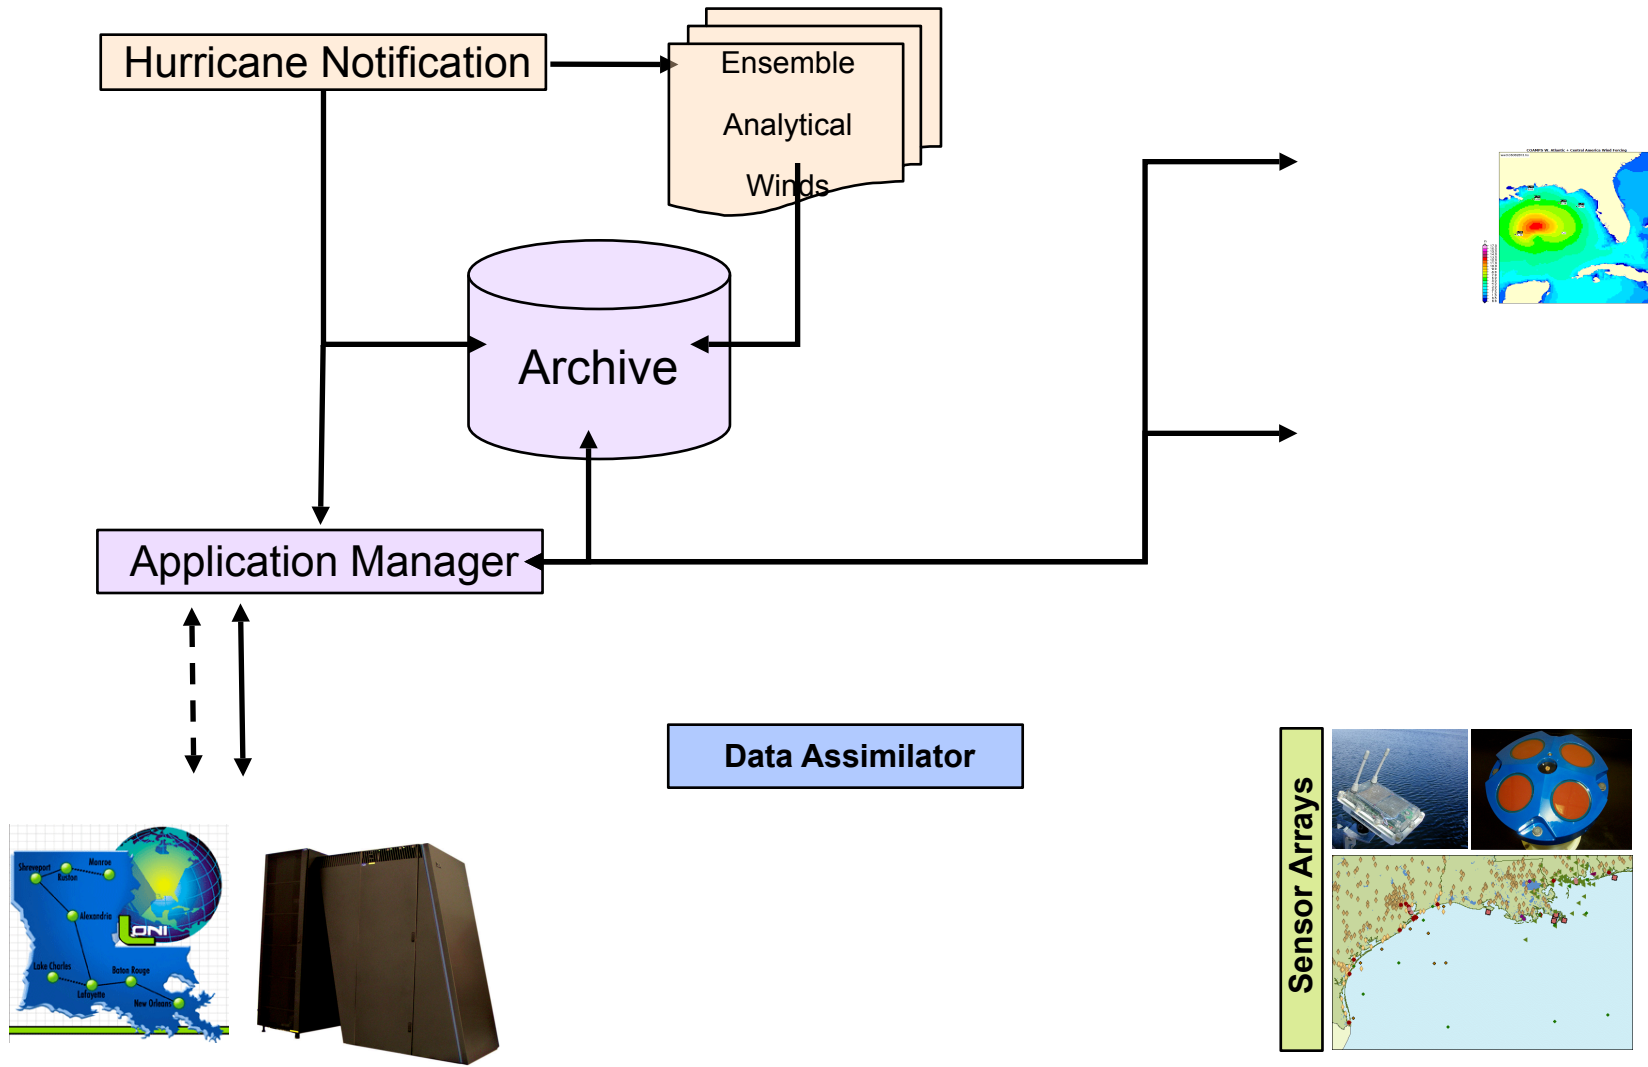
Dynamic, Event Driven Computing



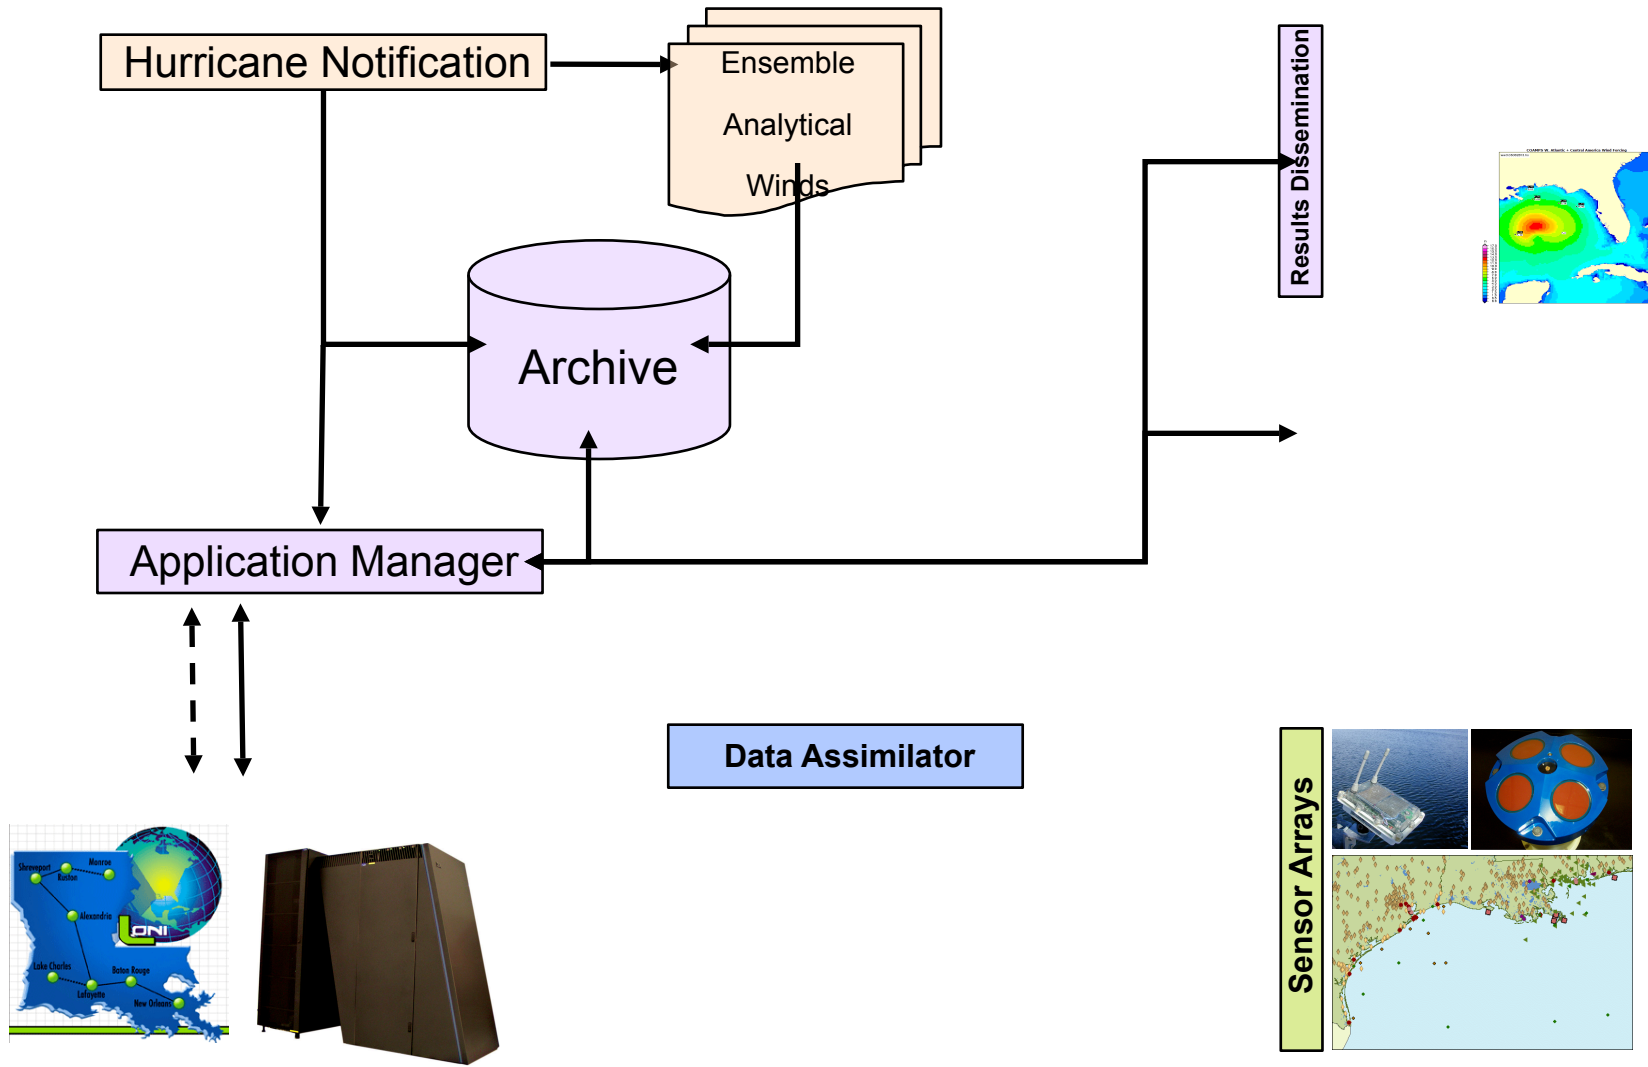
Dynamic, Event Driven Computing



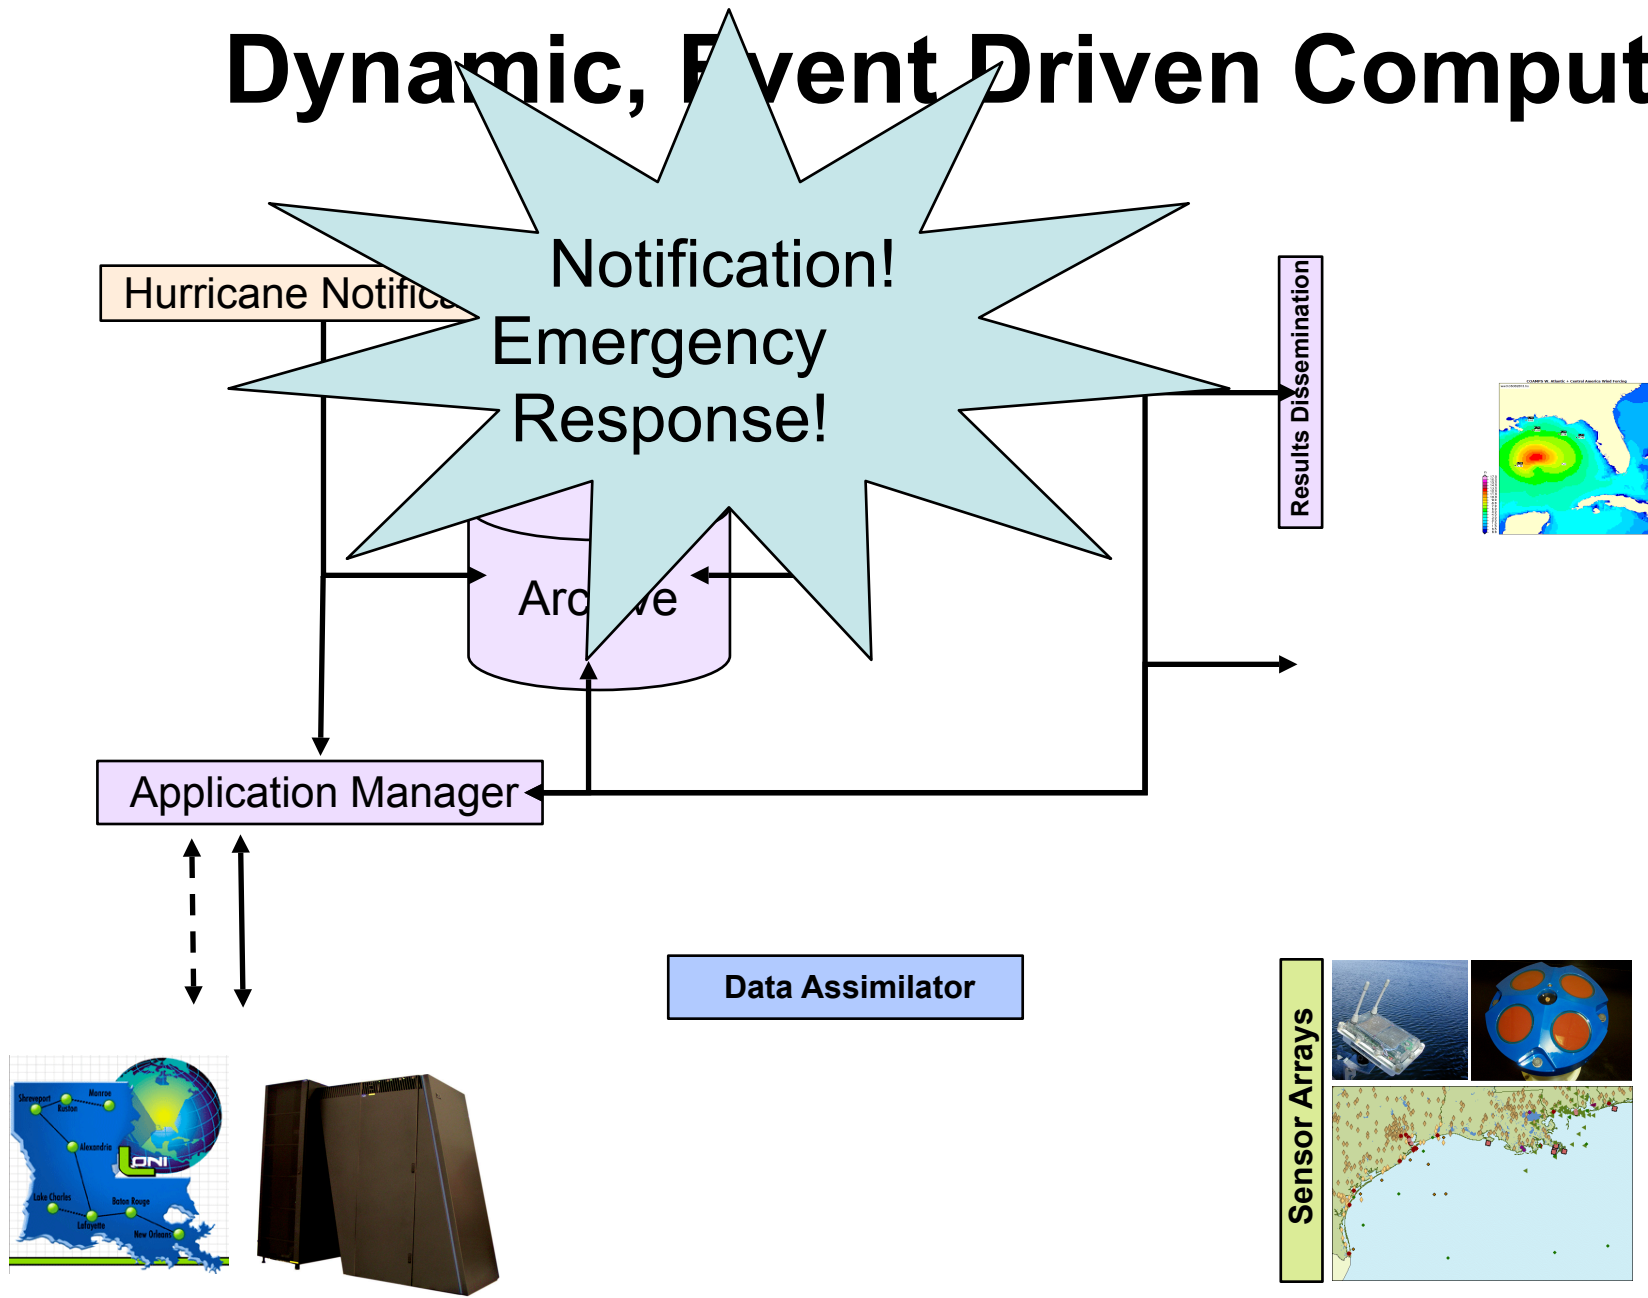
Dynamic, Event Driven Computing



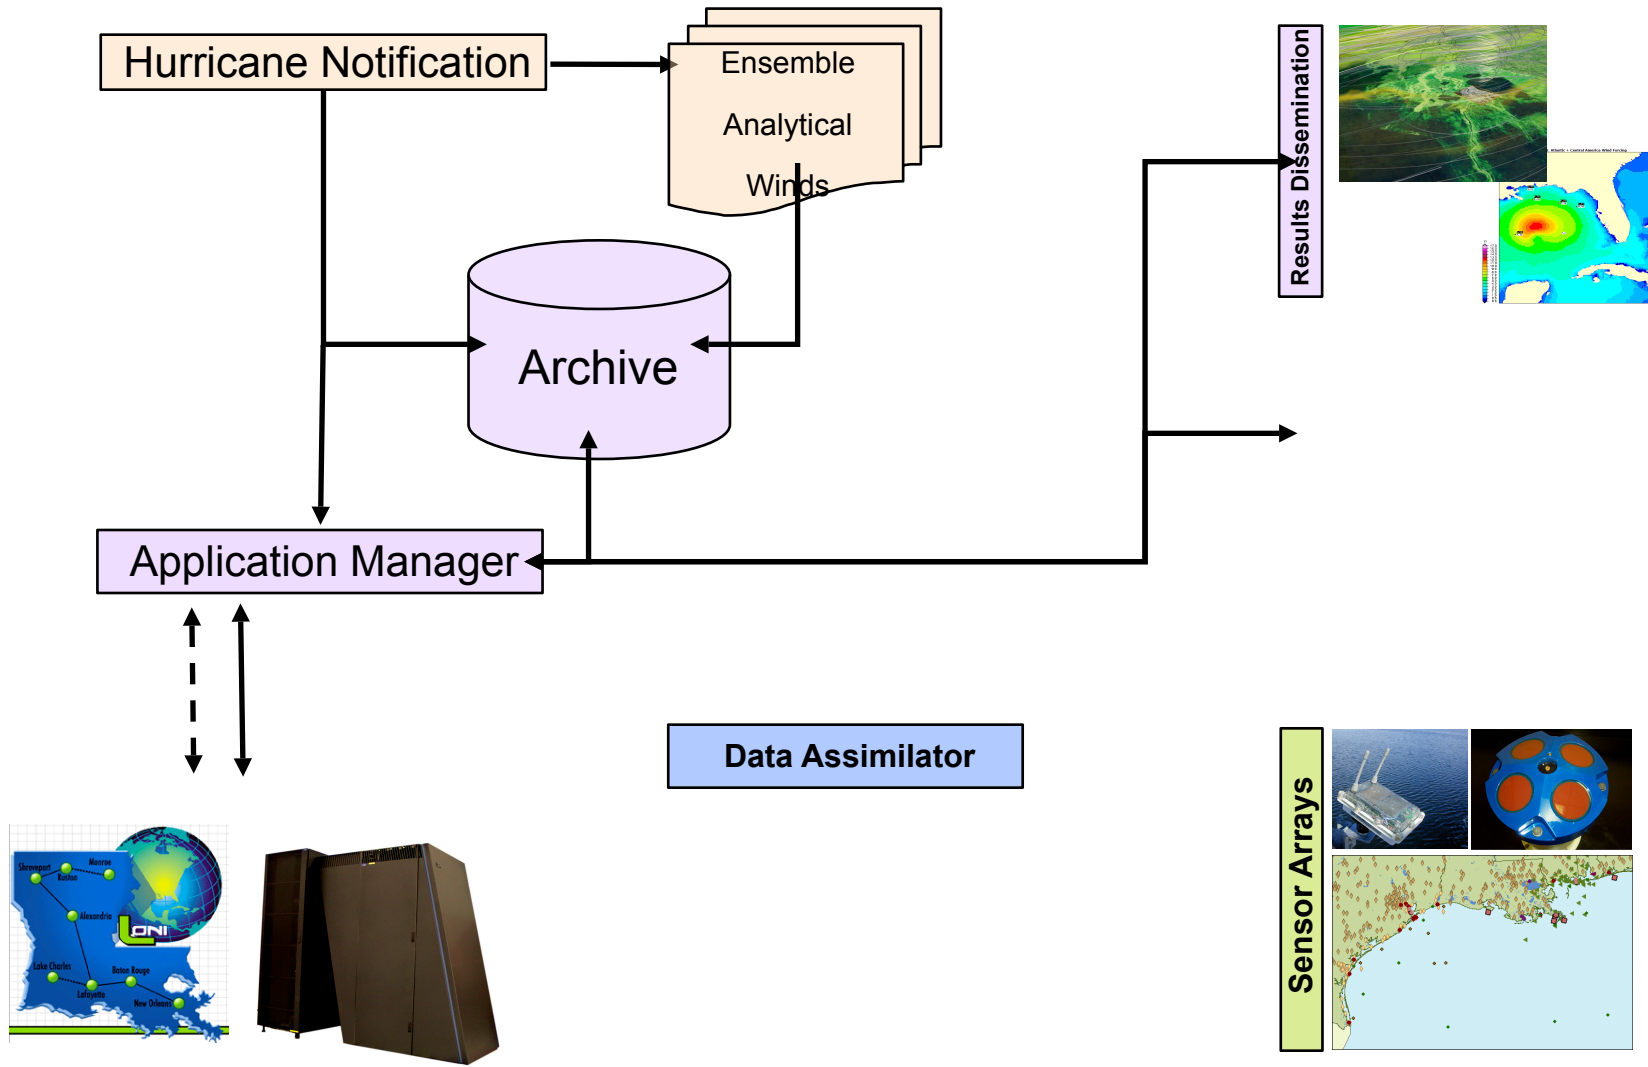
Dynamic, Event Driven Computing



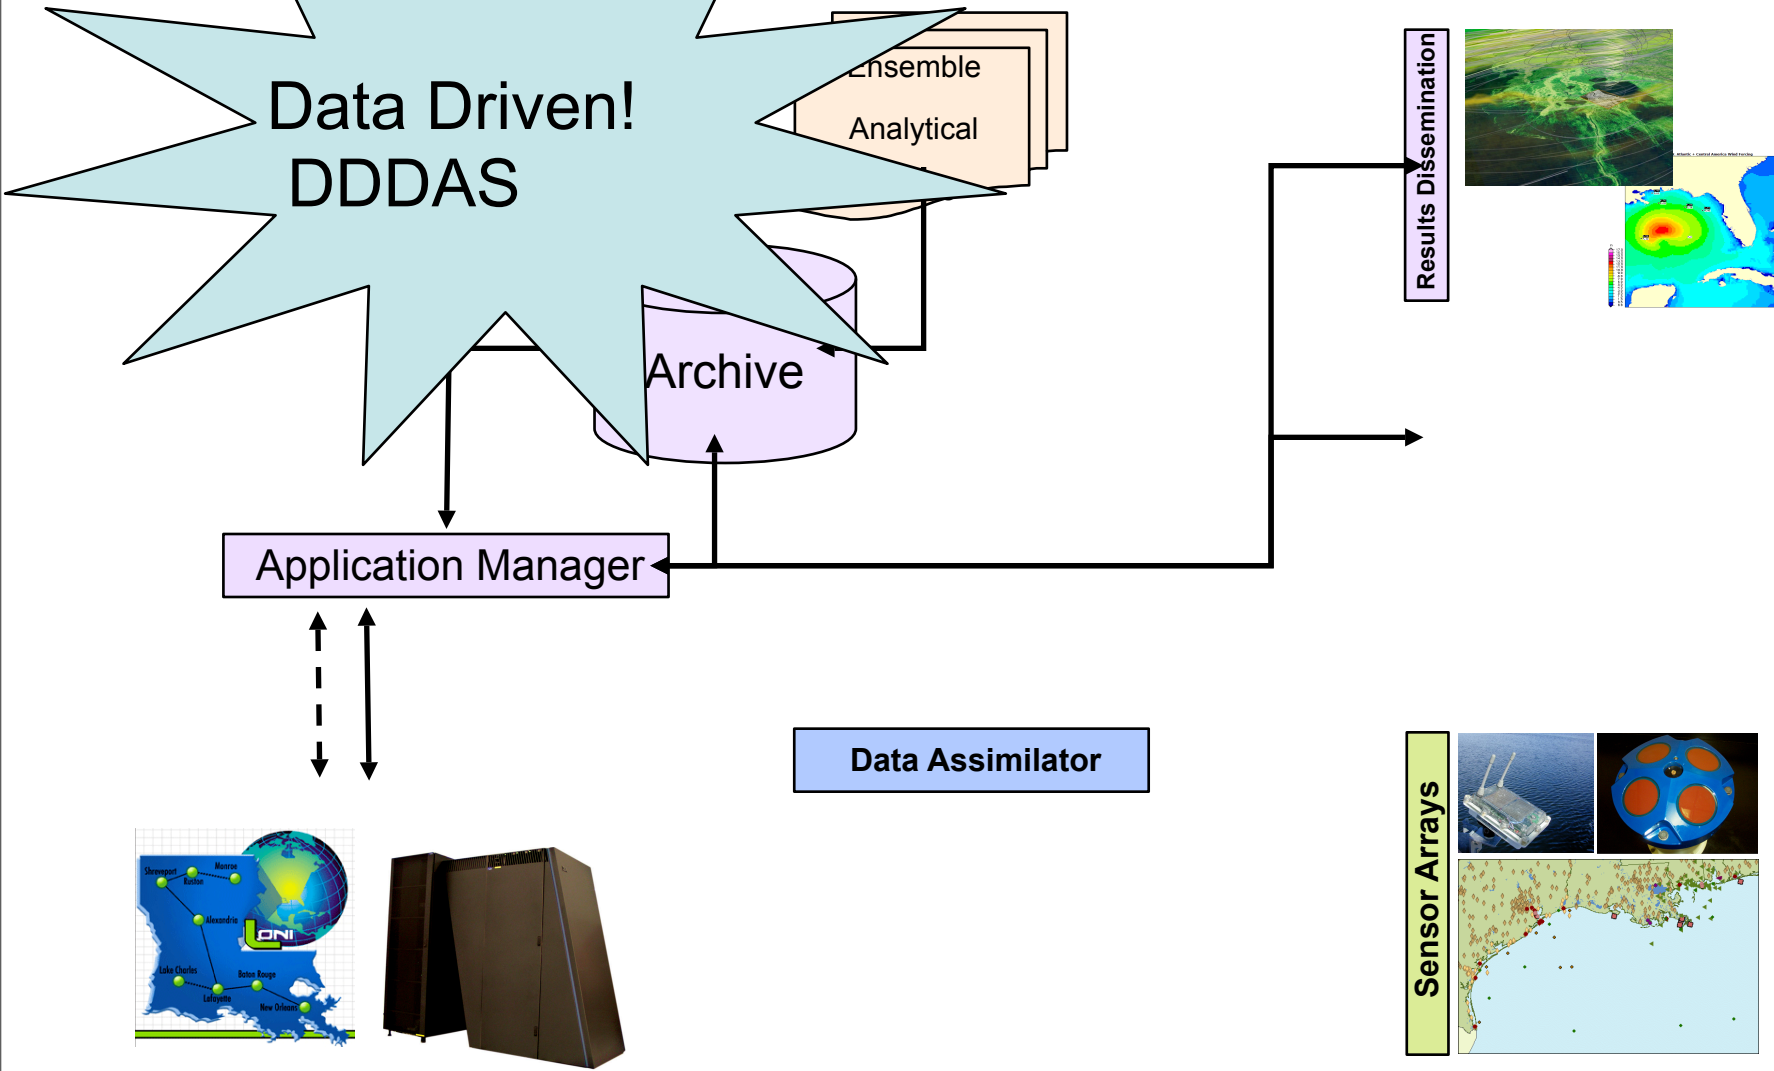
Dynamic, Event Driven Computing




Dynamic, Event Driven Computing



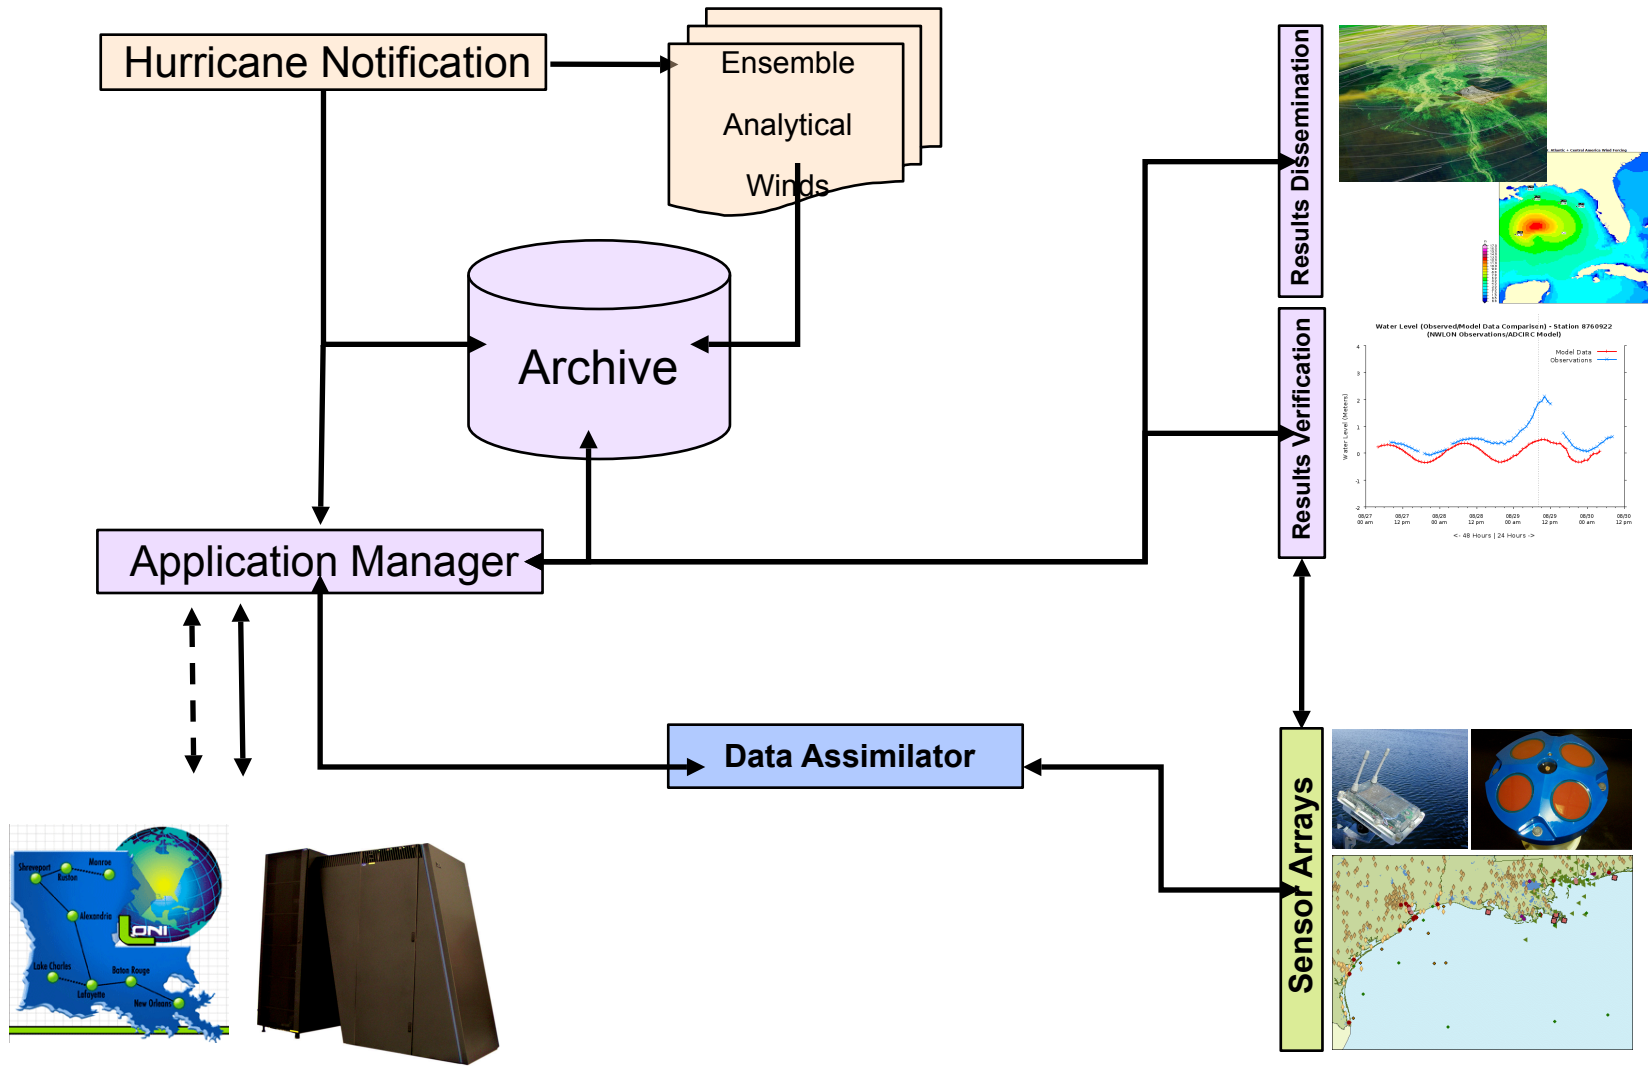
Dynamic, Event Driven Computing



Dynamic, Event Driven Computing

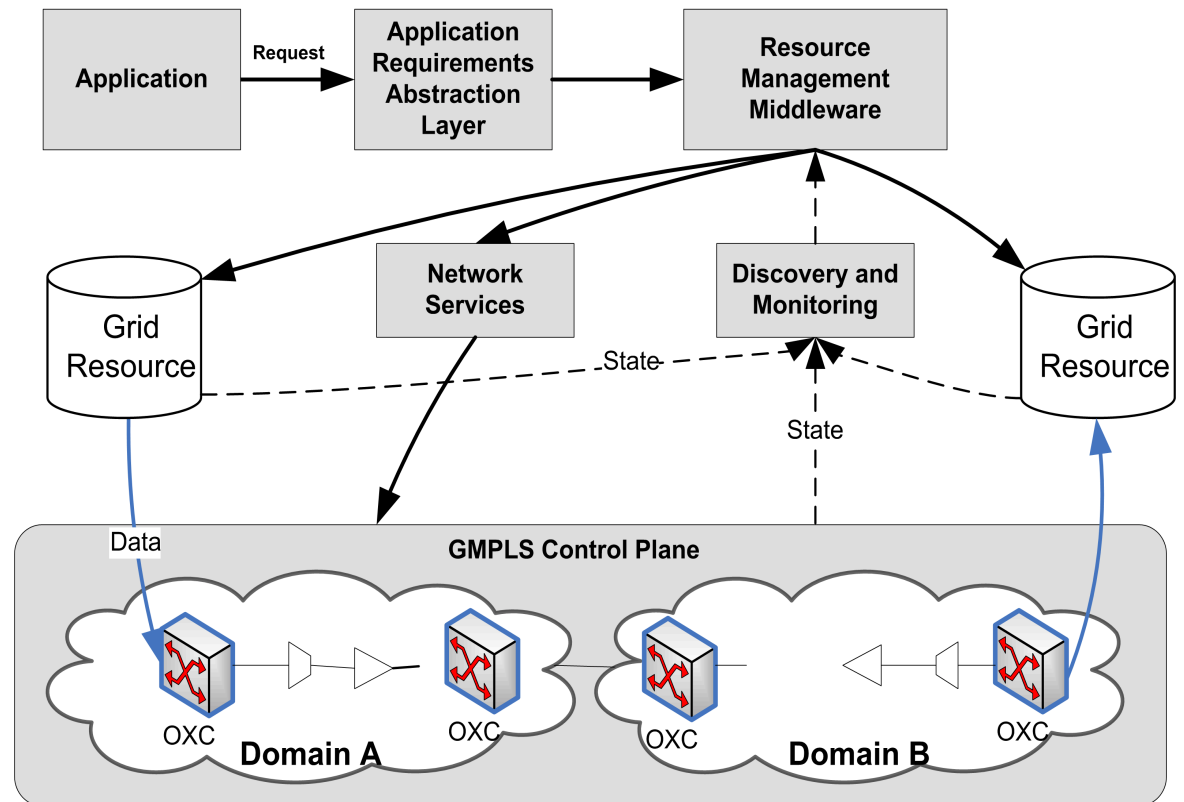


Dynamic, Event Driven Computing



enLIGHTened computing

- Coordinated use of compute and network resources
- Define the network through the set of services that it can offer
- MCNC (lead), LSU, RENC1, NCSU, Cisco, AT&T, Calient Networks
- 10Gig capacity testbed running over LONI and NLR

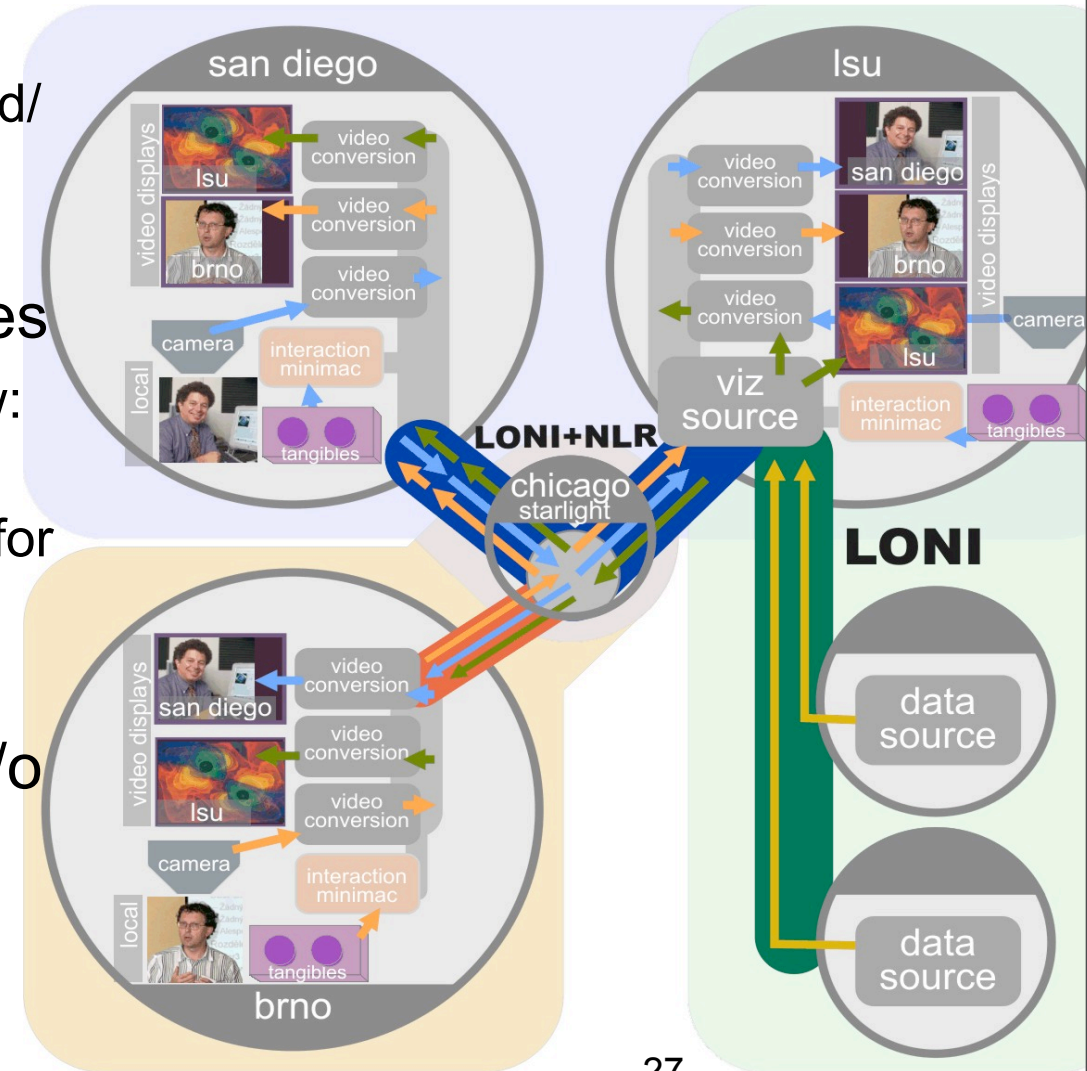




iGrid 2005

Distributed, Collaborative, Viz

- Scenario
 - TBs data, must be interactively viz'd/analyzed by international collaboration
- Distributed Data and Viz Services
 - Data larger than memory, disk slow: streaming data servers
 - Co-scheduling machines, network for collaborative event
 - Lambda Provisioning
- High Resolution: not science w/o details, analysis
 - Uncompressed HD stream: near latency-free remote viz at 1.5Gbit
 - Video for colleagues







All Require a Common Infrastructure

- Common Needs Driven by the Science/Engineering
 - Large Number of Sensors / Instruments
 - Daily Generation of Large Data Sets
 - Data is on Multiple Length and Time Scales
 - Automatic Archiving in Distributed Federated Repositories
 - Large Community of End Users
 - Multi-Megapixel and Immersive Visualization
 - Collaborative Analysis From Multiple Sites
 - *Complex Simulations Needed to Interpret Data*
- Will need Optical Networks
 - Communications → Dedicated Lambdas
 - Data → Large Peer-to-Peer Lambda Attached Storage
 - FAQ: Will usage policies allow scientists to use them?

Source: Smarr

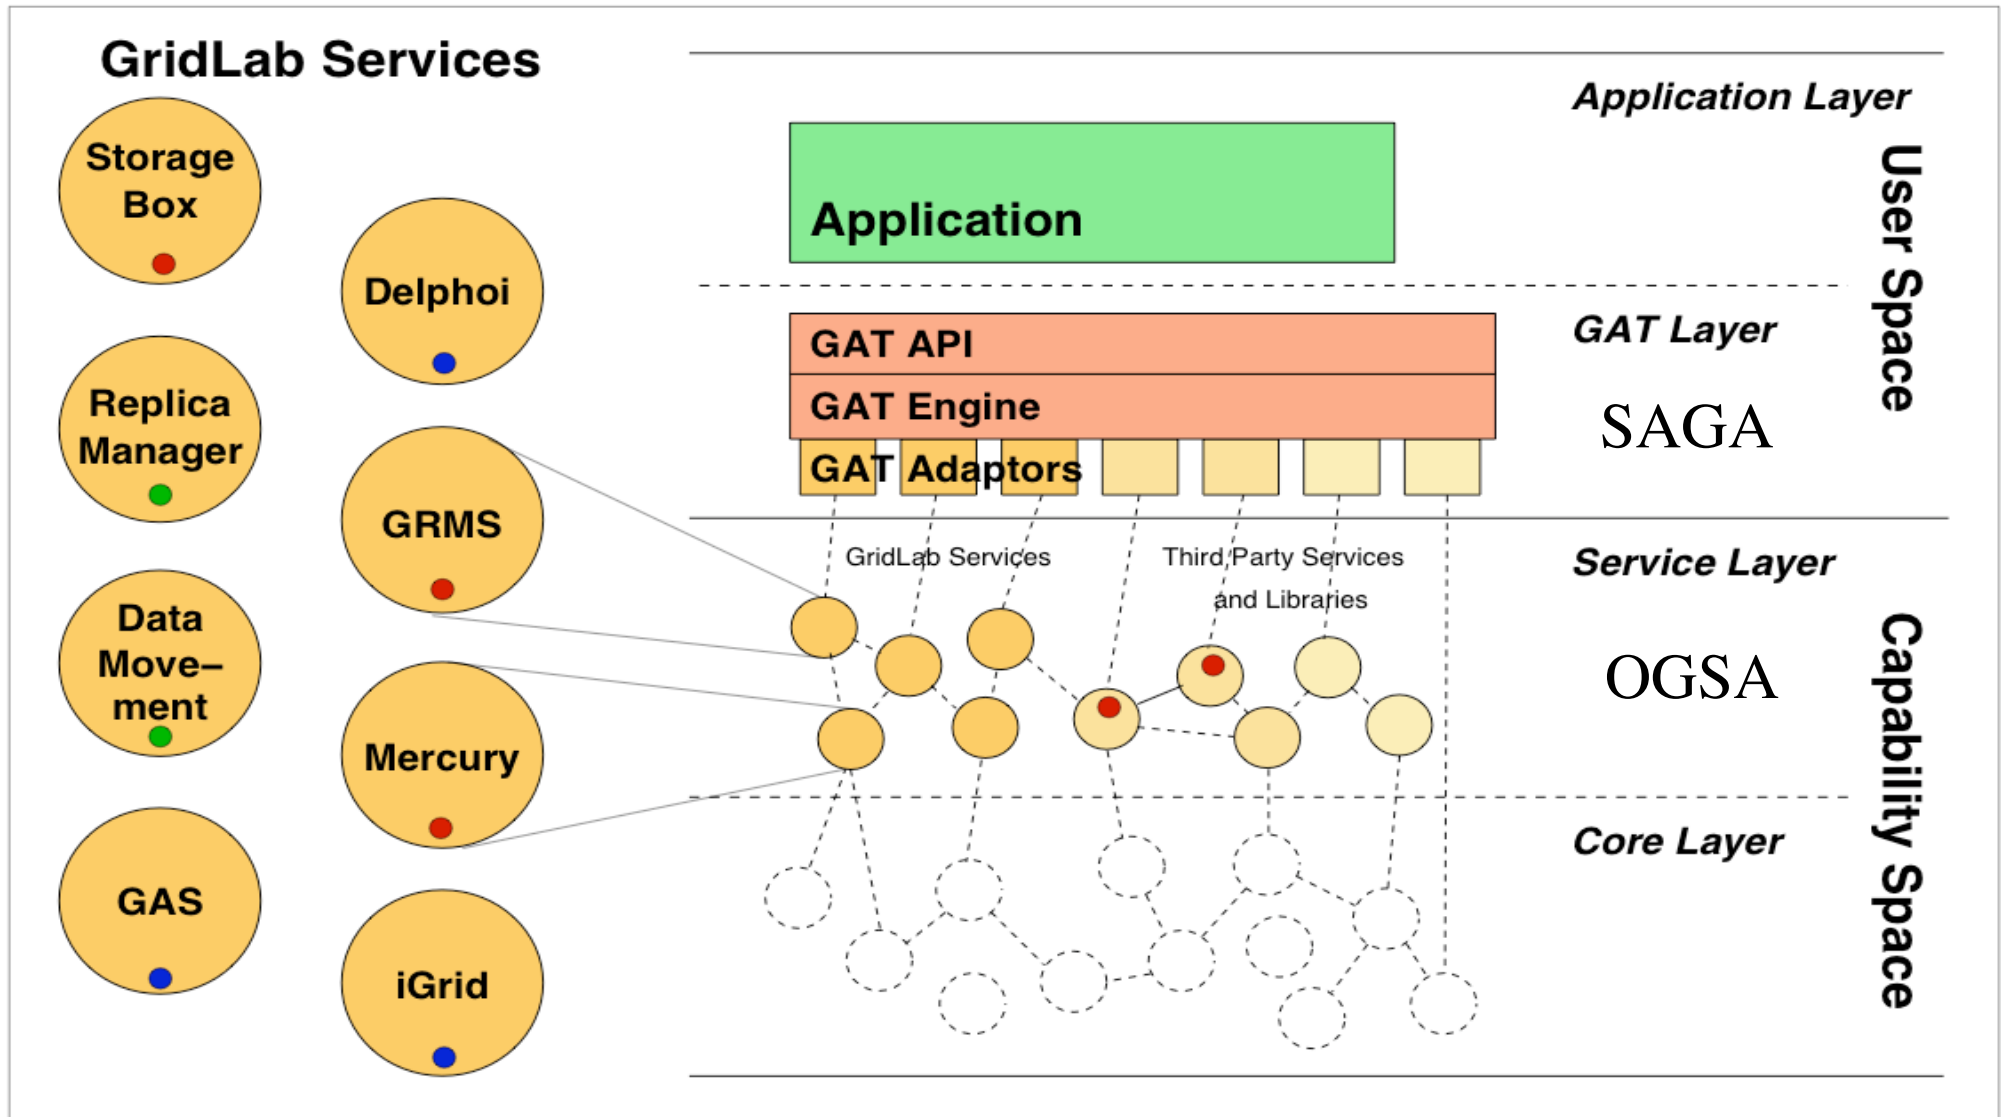


Developing Advanced Applications on Grids



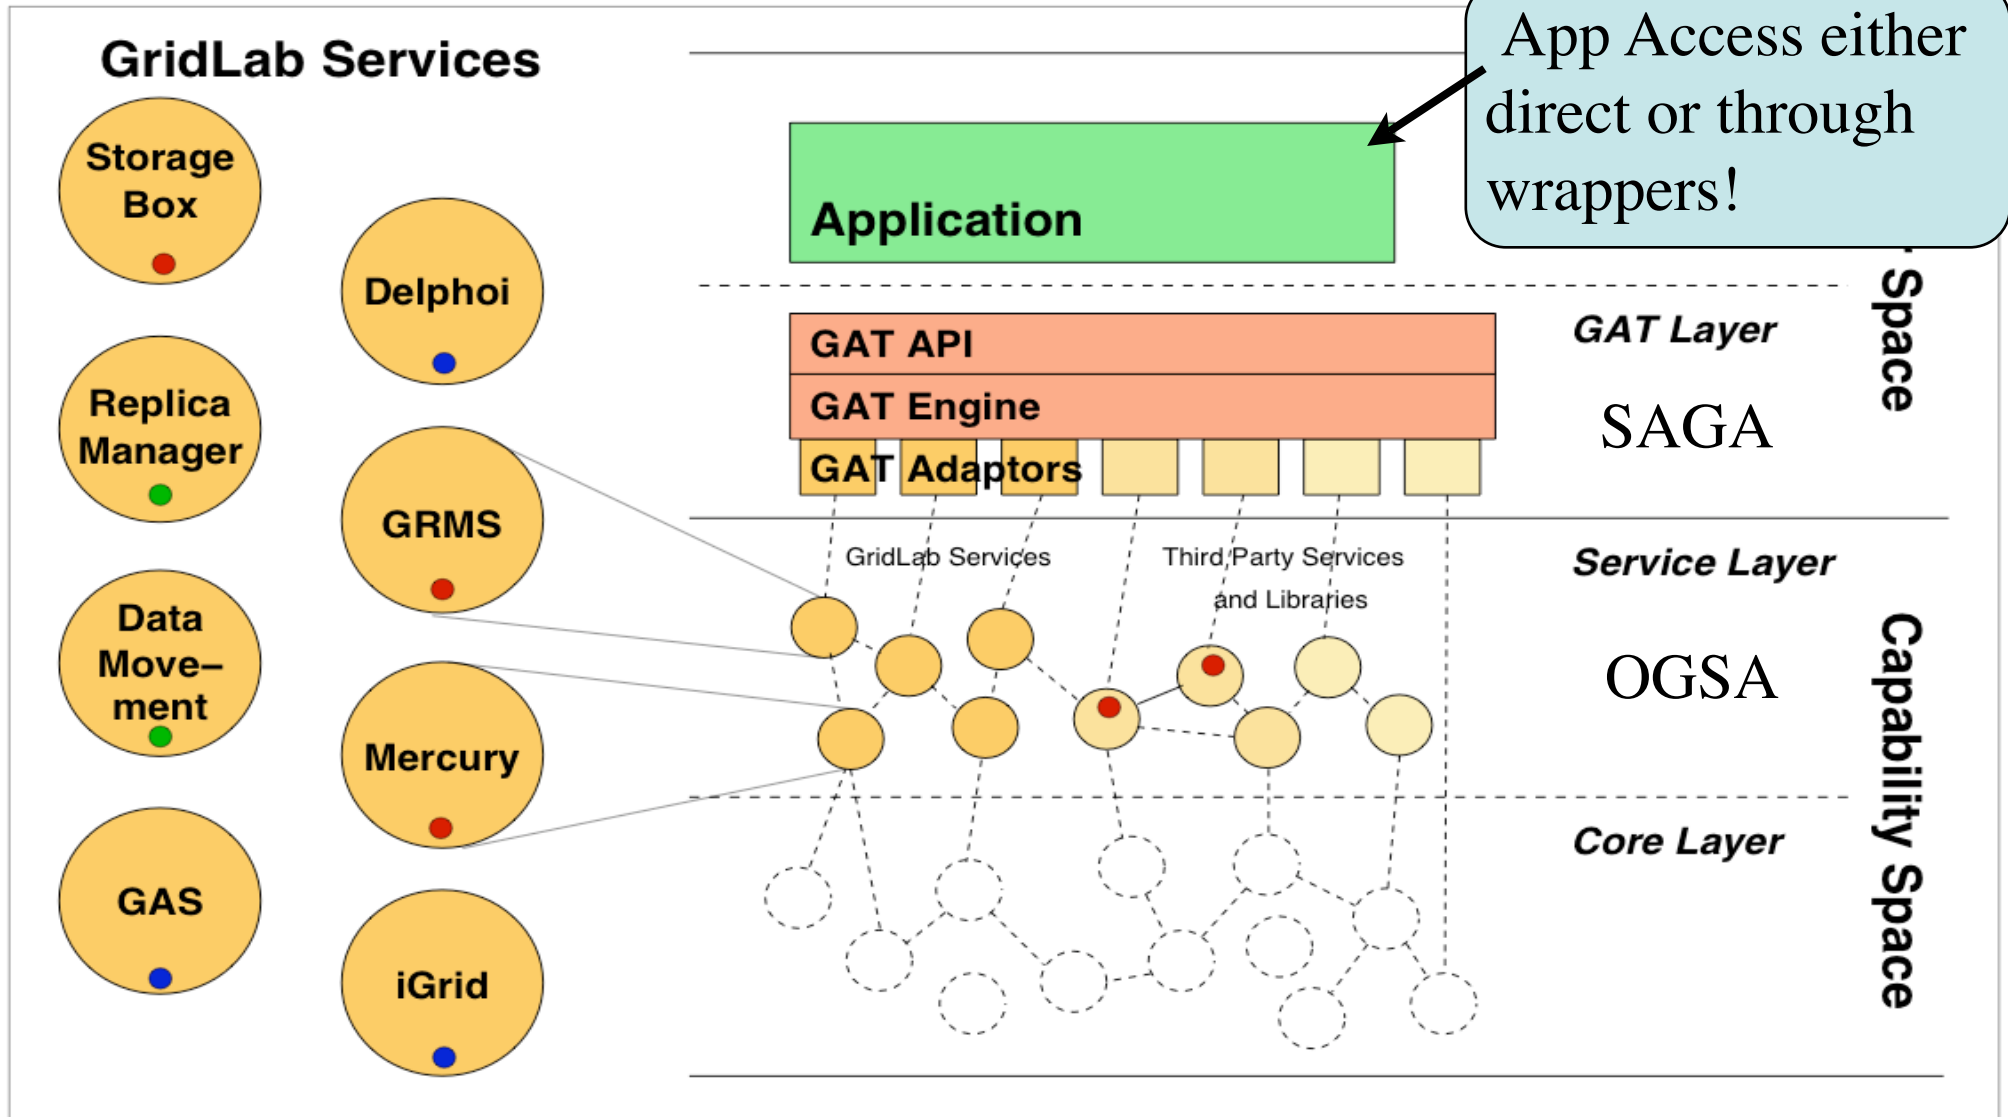


GridLab Architecture





GridLab Architecture





How to build Grid Apps?

Application

“Is there a better resource I could be using?”

“Can I get a lambda for my data?”

The
Grid



How to build Grid Apps?

Application

“Is there a better resource I could be using?” “Can I get a lambda for my data?”

SOAP

WSDL

CORBA

OGSA

Other

Monitoring

Security

Profiling

Information

Logging

Notification

Data Management

Resource Management

Application Manag

Migration

GLOBUS (v1, 2, 3, 4, ...)

UNICORE

Other Grid Infrastructure?



Hiding Complexity with Toolkits

Application

“Is there a better resource I could be using?”
GAT_FindResource()

GAT

The Grid



Example: Remote File Copying

Application

"Copy my file from there to there .."

SOAP

WSDL

Corba

OGSA

Other

Monitoring

Security

Profiling

Information

Logging

Notification

Data
Management

Resource
Management

Application
Manager

Migration

GLOBUS

Other Grid
Infrastructure?



Copy a File: GASS

```
int RemoteFile::GetFile (char const* source,
                        char const* target) {
    globus_url_t          source_url;
    globus_io_handle_t   dest_i_o_handle;
    globus_ftp_client_operationattr_t source_ftp_attr;
    globus_result_t      result;
    globus_gass_transfer_requestattr_t source_gass_attr;
    globus_gass_copy_attr_t source_gass_copy_attr;
    globus_gass_copy_handle_t gass_copy_handle;
    globus_gass_copy_handleattr_t gass_copy_handleattr;
    globus_ftp_client_handleattr_t ftp_handleattr;
    globus_io_attr_t      io_attr;
    int                   output_file = -1;

    if ( globus_url_parse (source_URL, &source_url) != GLOBUS_SUCCESS ) {
        printf ("can not parse source_URL \"%s\"\n", source_URL);
        return (-1);
    }

    if ( source_url.scheme_type != GLOBUS_URL_SCHEME_GSIFTP &&
        source_url.scheme_type != GLOBUS_URL_SCHEME_FTP &&
        source_url.scheme_type != GLOBUS_URL_SCHEME_HTTP &&
        source_url.scheme_type != GLOBUS_URL_SCHEME_HTTPS ) {
        printf ("can not copy from %s - wrong prot\n", source_URL);
        return (-1);
    }
    globus_gass_copy_handleattr_init (&gass_copy_handleattr);
    globus_gass_copy_attr_init (&source_gass_copy_attr);

    globus_ftp_client_handleattr_init (&ftp_handleattr);
    globus_io_fileattr_init (&io_attr);

    globus_gass_copy_attr_set_io (&source_gass_copy_attr, &io_attr);
    globus_gass_copy_handleattr_set_ftp_attr
        (&gass_copy_handleattr,
         &ftp_handleattr);
    globus_gass_copy_handle_init (&gass_copy_handle,
                                   &gass_copy_handleattr);

    if (source_url.scheme_type == GLOBUS_URL_SCHEME_GSIFTP ||
        source_url.scheme_type == GLOBUS_URL_SCHEME_FTP ) {
        globus_ftp_client_operationattr_init (&source_ftp_attr);
        globus_gass_copy_attr_set_ftp (&source_gass_copy_attr,
                                        &source_ftp_attr);
    }
    else {
        globus_gass_transfer_requestattr_init (&source_gass_attr,
                                                source_url.scheme);
        globus_gass_copy_attr_set_gass(&source_gass_copy_attr,
                                        &source_gass_attr);
    }

    output_file = globus_libc_open ((char*) target,
                                    O_WRONLY | O_TRUNC | O_CREAT,
                                    S_IRUSR | S_IWUSR | S_IRGRP |
                                    S_IWGRP);
    if ( output_file == -1 ) {
        printf ("could not open the file \"%s\"\n", target);
        return (-1);
    }
    /* convert stdout to be a globus_io handle */
    if ( globus_io_file_posix_convert (output_file, 0,
                                       &dest_i_o_handle)
        != GLOBUS_SUCCESS) {
        printf ("Error converting the file handle\n");
        return (-1);
    }

    result = globus_gass_copy_register_url_to_handle (
        &gass_copy_handle, (char*)source_URL,
        &source_gass_copy_attr, &dest_i_o_handle,
        my_callback, NULL);
    if ( result != GLOBUS_SUCCESS ) {
        printf ("error: %s\n", globus_object_printable_to_string
            (globus_error_get (result)));
        return (-1);
    }
    globus_url_destroy (&source_url);
    return (0);
}
```



Copy a File: CoG/RFT

```
package org.globus.ogsa.gui;

import java.io.BufferedReader;
import java.io.File;
import java.io.FileReader;
import java.net.URL;
import java.util.Date;
import java.util.Vector;
import javax.xml.rpc.Stub;
import org.apache.axis.message.MessageElement;
import org.apache.axis.utils.XMLUtils;
import org.globus.*
import org.gridforum.ogsi.*
import org.gridforum.ogsi.holders.TerminationTimeTypeHolder;
import org.w3c.dom.Document;
import org.w3c.dom.Element;

public class RFTClient {
    public static void copy (String source_url, String target_url) {
        try {
            File requestFile = new File (source_url);
            BufferedReader reader = null;
            try {
                reader = new BufferedReader (new FileReader (requestFile));
            } catch (java.io.FileNotFoundException fnfe) { }
            Vector requestData = new Vector ();
            requestData.add (target_url);
            TransferType[] transfers1 = new TransferType[transferCount];
            RFTOptionsType multirftOptions = new RFTOptionsType ();

            multirftOptions.setBinary (Boolean.valueOf (
                (String)requestData.elementAt (0)).booleanValue ());
            multirftOptions.setBlockSize (Integer.valueOf (
                (String)requestData.elementAt (1)).intValue ());
            multirftOptions.setTcpBufferSize (Integer.valueOf (
                (String)requestData.elementAt (2)).intValue ());
            multirftOptions.setNotpt (Boolean.valueOf (
                (String)requestData.elementAt (3)).booleanValue ());
            multirftOptions.setParallelStreams (Integer.valueOf (
                (String)requestData.elementAt (4)).intValue ());
            multirftOptions.setDcau (Boolean.valueOf (
                (String)requestData.elementAt (5)).booleanValue ());

            int i = 7;
            for (int j = 0; j < transfers1.length; j++)
            {
                transfers1[j] = new TransferType ();

                transfers1[j].setTransferId (j);
                transfers1[j].setSourceUrl ((String)requestData.elementAt (i++));
                transfers1[j].setDestinationUrl ((String)requestData.elementAt (i++));
                transfers1[j].setRftOptions (multirftOptions);
            }
        }
    }
}
```

```
TransferRequestType transferRequest = new TransferRequestType ();
transferRequest.setTransferArray (transfers1);

int concurrency = Integer.valueOf
    ((String)requestData.elementAt(6)).intValue();

if (concurrency > transfers1.length)
{
    System.out.println ("Concurrency should be less than the number"
        "of transfers in the request");
    System.exit (0);
}
transferRequest.setConcurrency (concurrency);

TransferRequestElement requestElement = new TransferRequestElement ();
requestElement.setTransferRequest (transferRequest);

ExtensibilityType extension = new ExtensibilityType ();
extension = AnyHelper.getExtensibility (requestElement);

OGSIServiceGridLocator factoryService = new OGSIServiceGridLocator ();
Factory factory = factoryService.getFactoryPort (new URL (source_url));
GridServiceFactory gridFactory = new GridServiceFactory (factory);

LocatorType locator = gridFactory.createService (extension);
System.out.println ("Created an instance of Multi-RFT");

MultiFileRFTDefinitionServiceGridLocator loc
    = new MultiFileRFTDefinitionServiceGridLocator();
RFTPortType rftPort = loc.getMultiFileRFTDefinitionPort (locator);
((Stub)rftPort)._setProperty (Constants.AUTHORIZATION,
    NoAuthorization.getInstance());
((Stub)rftPort)._setProperty (GSIConstants.GSI_MODE,
    GSIConstants.GSI_MODE_FULL_DELEG);
((Stub)rftPort)._setProperty (Constants.GSI_SEC_CONV,
    Constants.SIGNATURE);
((Stub)rftPort)._setProperty (Constants.GRIM_POLICY_HANDLER,
    new IgnoreProxyPolicyHandler ());

int requestid = rftPort.start ();
System.out.println ("Request id: " + requestid);

}
catch (Exception e)
{
    System.err.println (MessageUtils.toString (e));
}
}
```



Copy a File: GAT/C++

```
#include <GAT++.hpp>
```

```
GAT::Result RemoteFile::GetFile (GAT::Context context,  
                                std::string source_url,  
                                std::string target_url)  
{  
    try  
    {  
        GAT::File file (context, source_url);  
        file.Copy      (target_url);  
    }  
    catch (GAT::Exception const &e)  
    {  
        std::cerr << "Some error: " << e.what() << std::endl;  
        return e.Result();  
    }  
    return GAT_SUCCESS;  
}
```



GGF SAGA-WG

A. Merzky, T. Goodale, S. Newhouse, et al

- GAT evolves into GGF standard
 - Numerous attempts to address: GAT most ambitious, but also CoG, DRMAA, GridRPC, GridCPR, many others
- SAGA: Simple API for Grid Applications
 - Bringing all these efforts together through single API spec
 - Chicago, Berlin, Brussels, LSU, Berkeley, Seoul, Chicago, Boston
- GGF focussing now on standardization
 - SAGA API spec done
 - Much momentum in SAGA now



Finally

- Optical Networks, Grids promise new ways of computing
 - More than computing: Network, Data services critical
- Complex applications require comprehensive services
 - Event Driven: many examples. Hurricanes, earthquakes, tornados, explosions, ...
 - DDDAS: http://www.nsf.gov/cise/cns/dddas/2006_Workshop/exec_summ.pdf
 - Much of this can work now as prototypes
- Standards developing
 - 15 years ago: parallel computing drove interconnects, HPF, MPI
 - Now: 2 levels...OGSA grid services, SAGA for apps
- Cactus Computational Toolkit: www.cactuscode.org
- GGF “Simple API for Grid Applications” (SAGA)