

Louisiana Institute Cyberinfrastructure Development Efforts


Tevfik Kosar

Department of Computer Science & CCT
Louisiana State University

October 31, 2008



CENTER FOR COMPUTATION
& TECHNOLOGY




Statewide Cyberinfrastructure Efforts



- Scheduling, data
- Visualization
- Web Services
- HPC Interfaces & API

- Distributed data management
- Storage, data archival & retrieval


- Sensor Networks
- Cyber Security



**CENTER FOR
SECURE
CYBERSPACE**

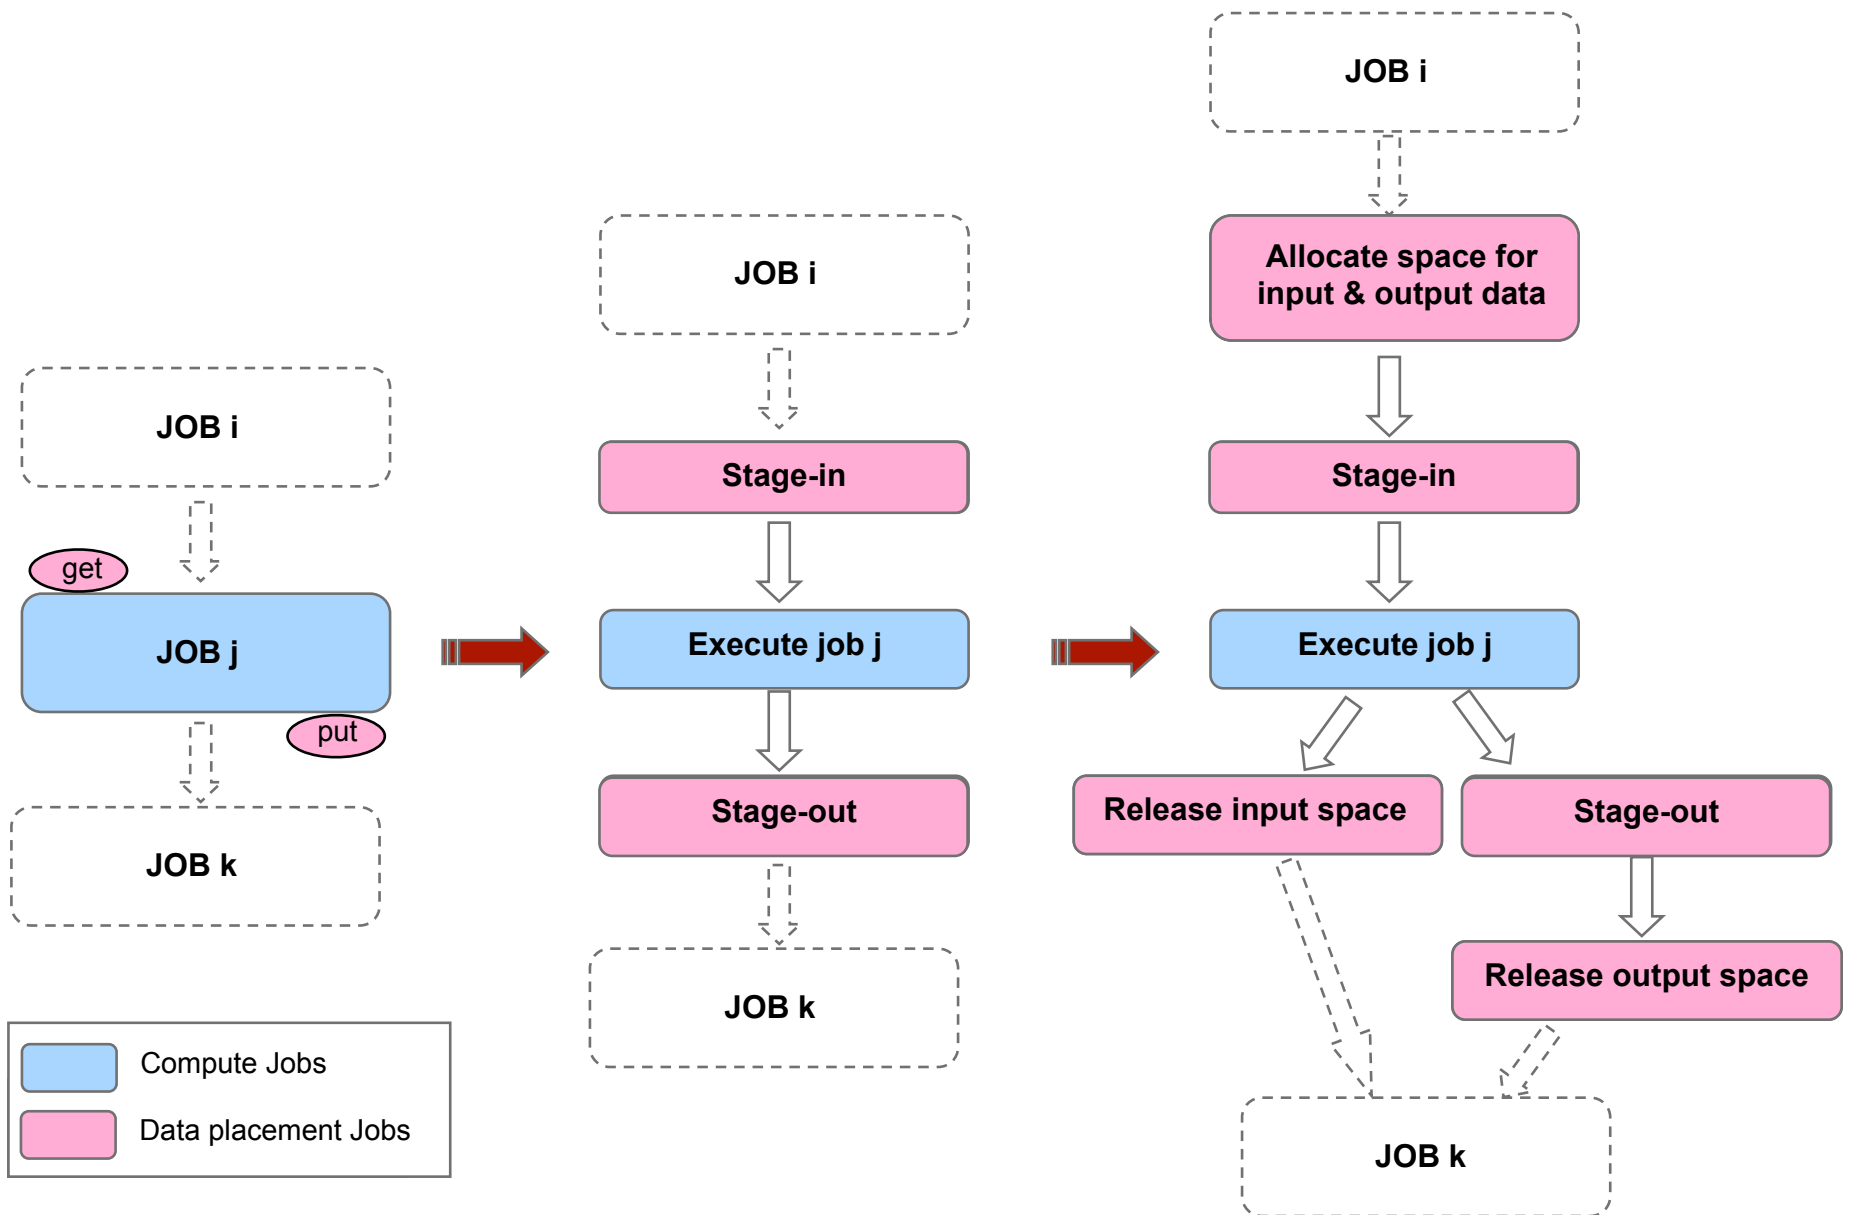
Today's Talk

- ★ End-to-end Workflow Management
- ★ Data Scheduling
- ★ Distributed Data Storage




End-to-end Workflows

Data-aware Workflows

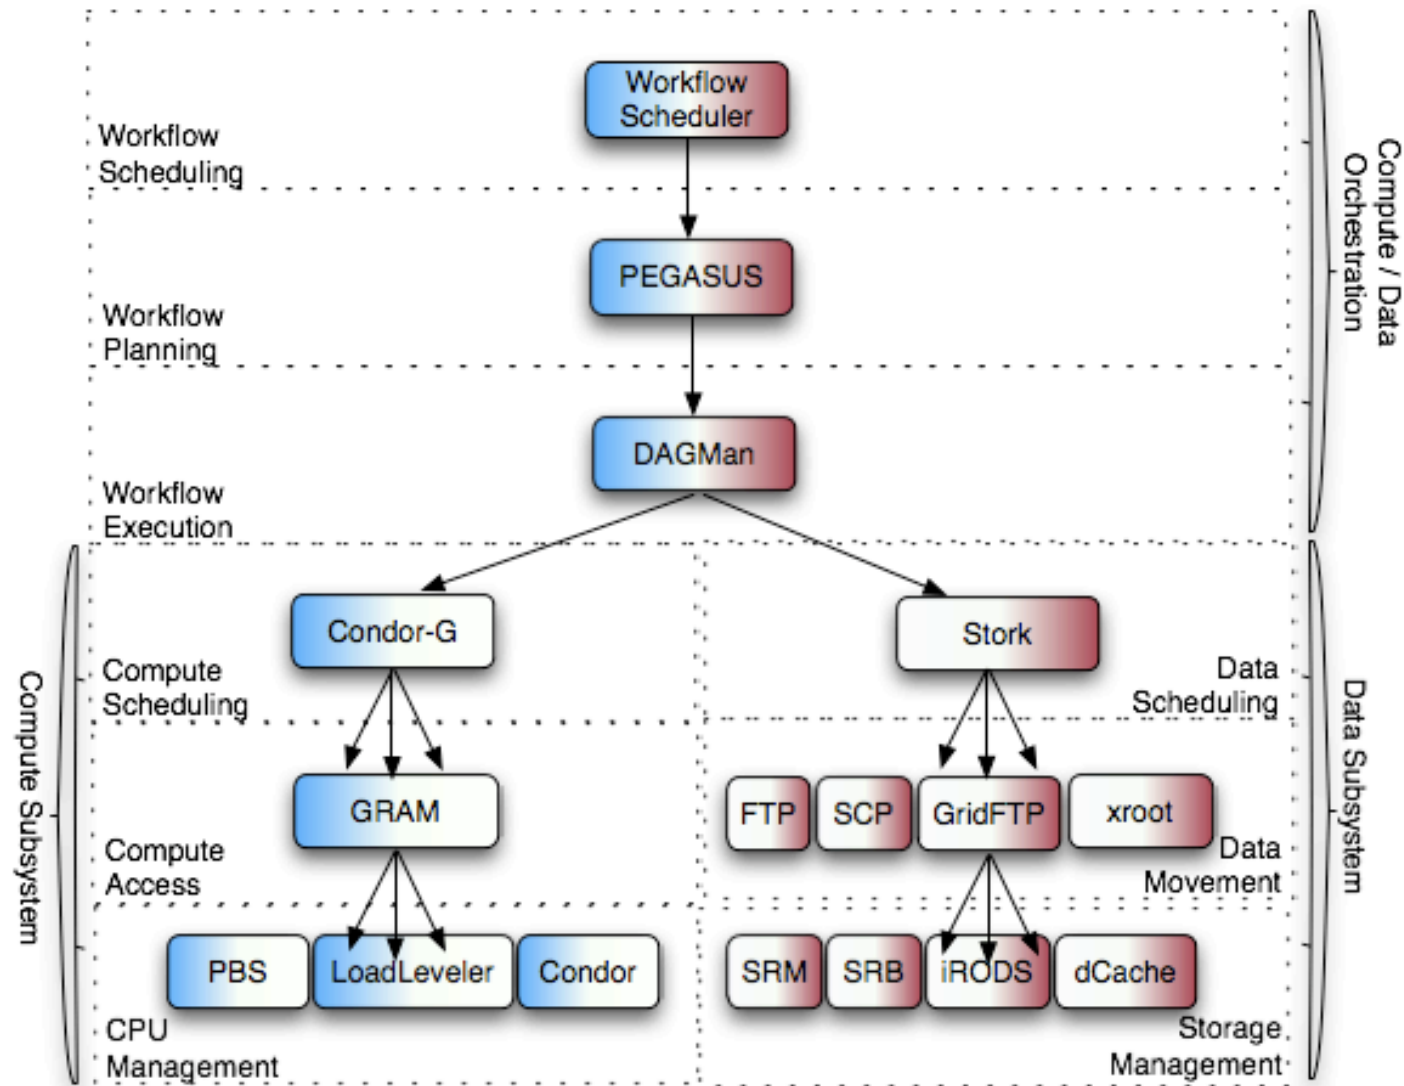


A Less Simpler Workflow

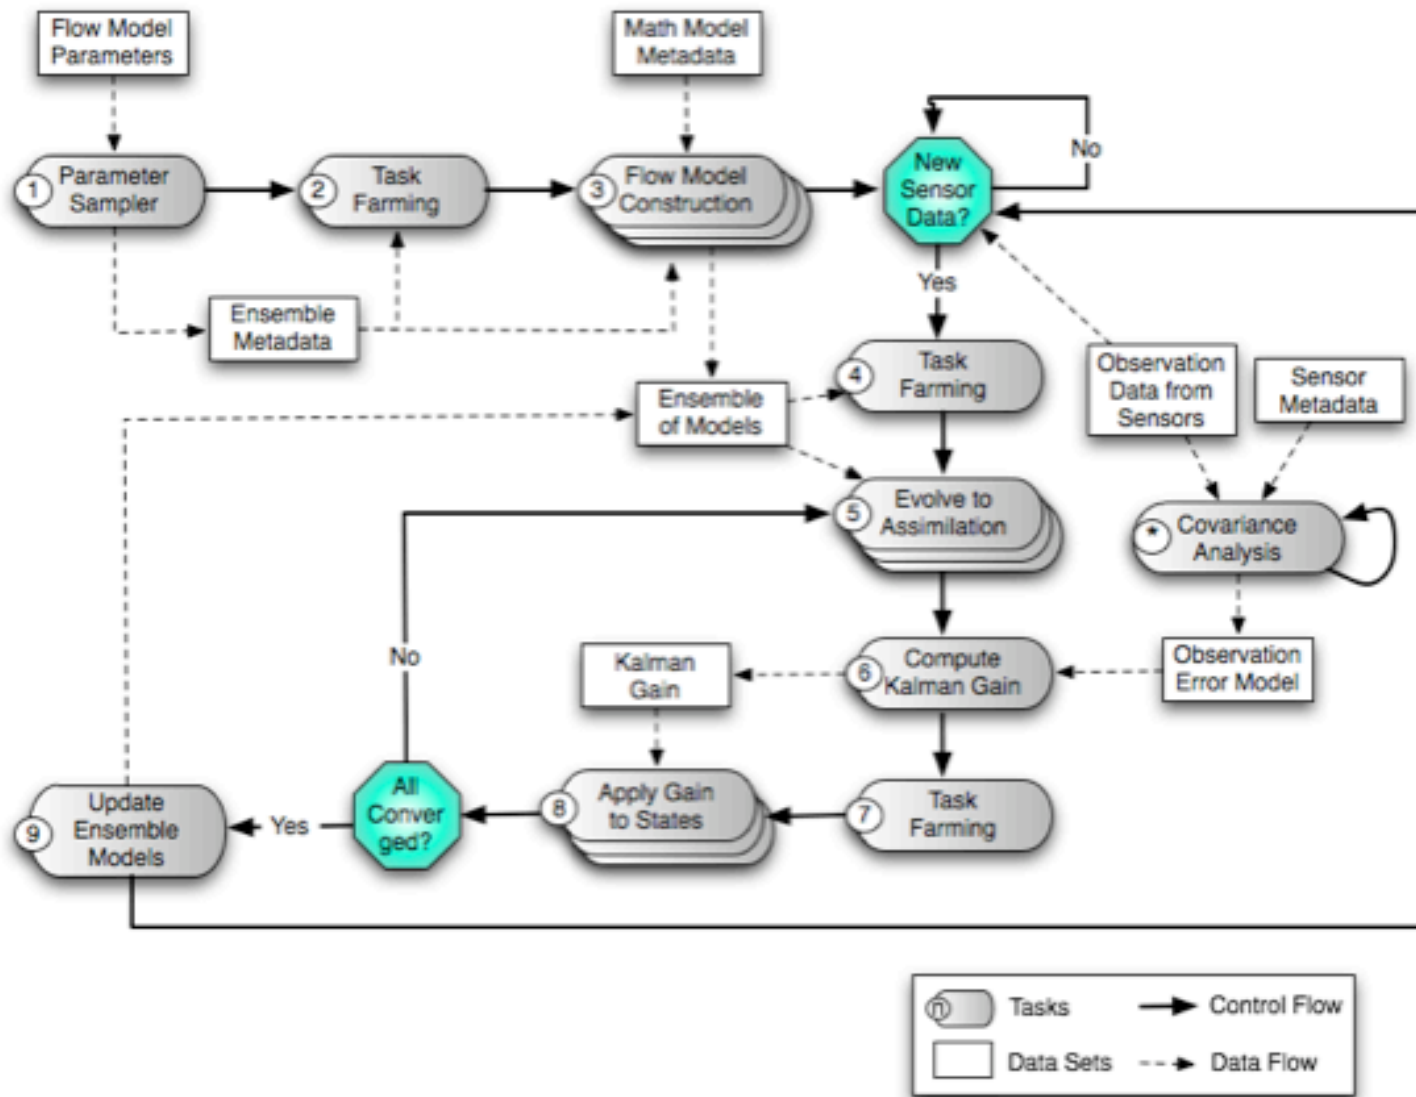


Credits: Montage Project

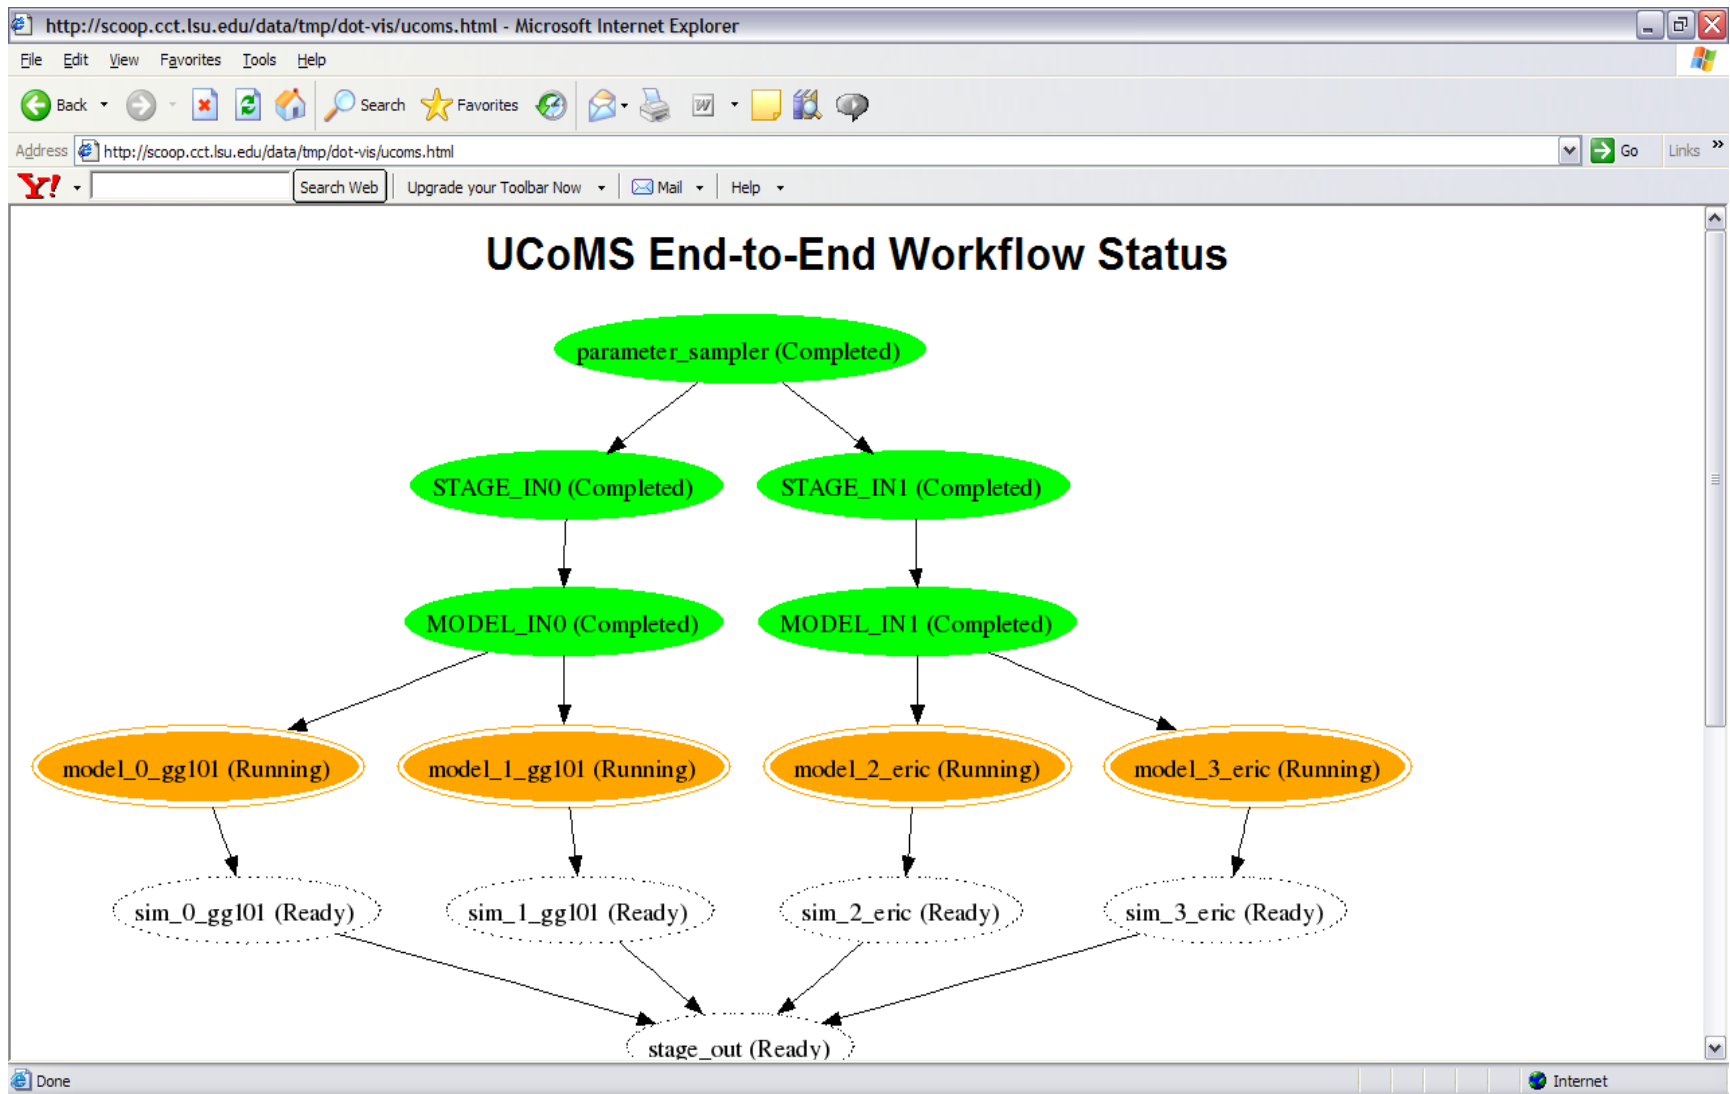
Separation of CPU & IO



UCoMS Workflow




Monitoring Workflows via WEB



Data Scheduling

Separation of CPU & IO



Traditional Schedulers

- ★ Traditional schedulers not aware of characteristics and semantics of data placement jobs

```
Executable = genome.exe
```

```
Arguments  = a b c d
```

```
Executable = globus-url-copy
```


```
Arguments  = gsiftp://host1/f1
```

```
            gsiftp://host2/f2
```

Any difference?

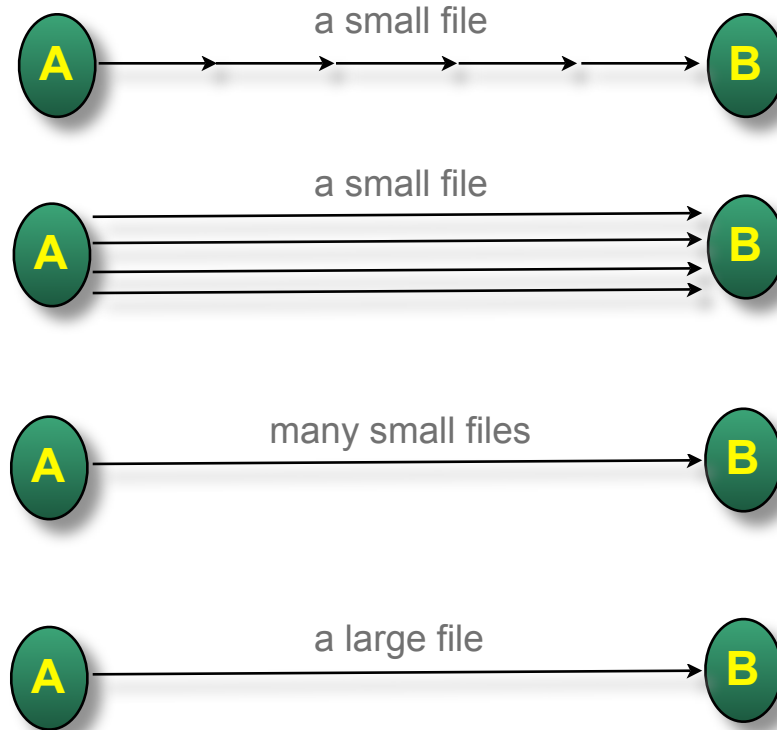
Data-Aware Schedulers

- ★ Type of a job?
 - 🕒 transfer, allocate, release, locate..
- ★ Priority, order?
- ★ Protocol to use?
- ★ Second vs Third party?
- ★ Available storage space?
- ★ Best concurrency level?
- ★ Reasons for failure?
- ★ Best network parameters?
 - 🕒 tcp buffer size
 - 🕒 I/O block size
 - 🕒 # of parallel streams




Stork Transfer Methods

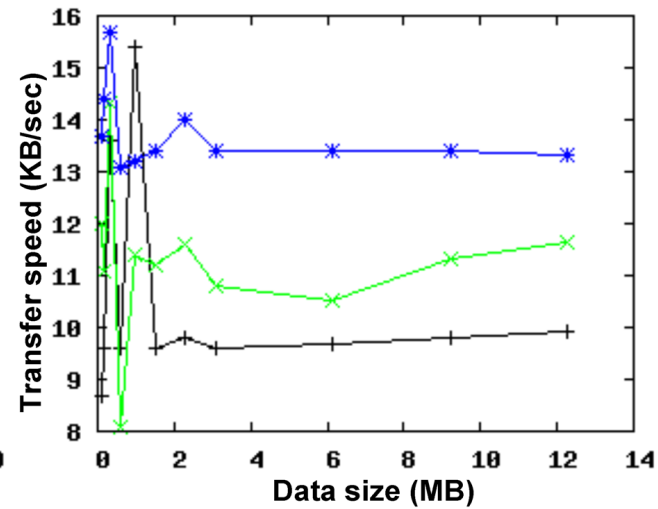
- **regular:**
 - one connection per file, serial transfer
- **multi-connection:**
 - one connection per file, concurrent transfer
- **single-connection:**
 - one connection for all transfers
- **data-fusion:**
 - merge small files into larger chunks



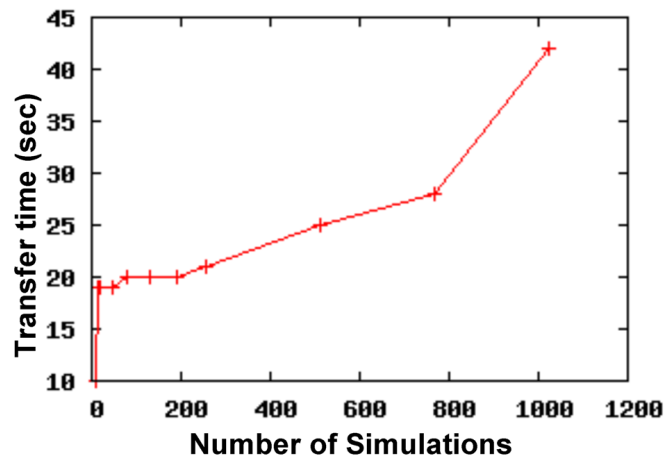
Stork Transfer Results




(a)



(b)




(c)




(d)

regular —+— Stork single-connection —*—
 Stork multi-connection —x— Stork data-fusion —+—

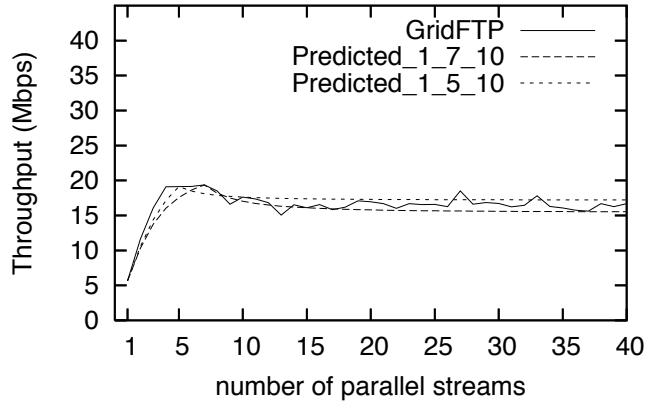
a) Dinda et al Model Prediction Results




b) Averaging Prediction Results




c) Break Function Prediction Results



d) Logarithmic Prediction Results




e) Newton's Method Prediction Results




Dynamic
Tuning

Based on
light-weight
prediction




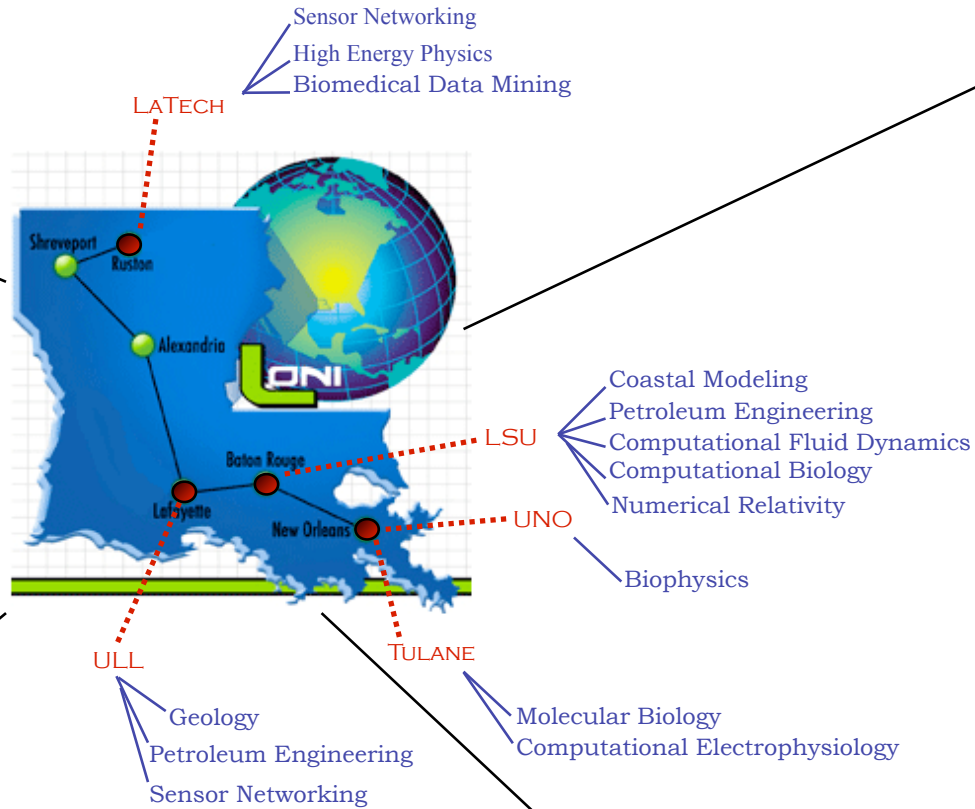

Distributed Storage



- **Goal:** enable domain scientists to focus on their primary research problem, assured that the underlying infrastructure will manage the low-level data handling issues.
- **Novel approach:** treat data storage resources and the tasks related to data access as first class entities just like computational resources and compute tasks.
- **Key technologies** being developed: data-aware storage systems, data-aware schedulers (i.e. Stork), and cross-domain meta-data scheme.
- **Provides** and additional 250TB disk, and 400TB tape storage



Participating institutions in the PetaShare project, connected through LONI. Sample research of the participating researchers pictured (i.e. biomechanics by Kodiyalam & Wischusen, tangible interaction by Ullmer, coastal studies by Walker, and molecular biology by Bishop).



Infrastructure Overview


LaTech

ULL

LSU

UNO

Tulane




~ 100 TFLOPS

~8 TB RAM


250 TB Disk

400 TB Tape



SDSC/NCSA
100s TB

Replica & Meta Data Management



PetaShare Core

POSIX interface:

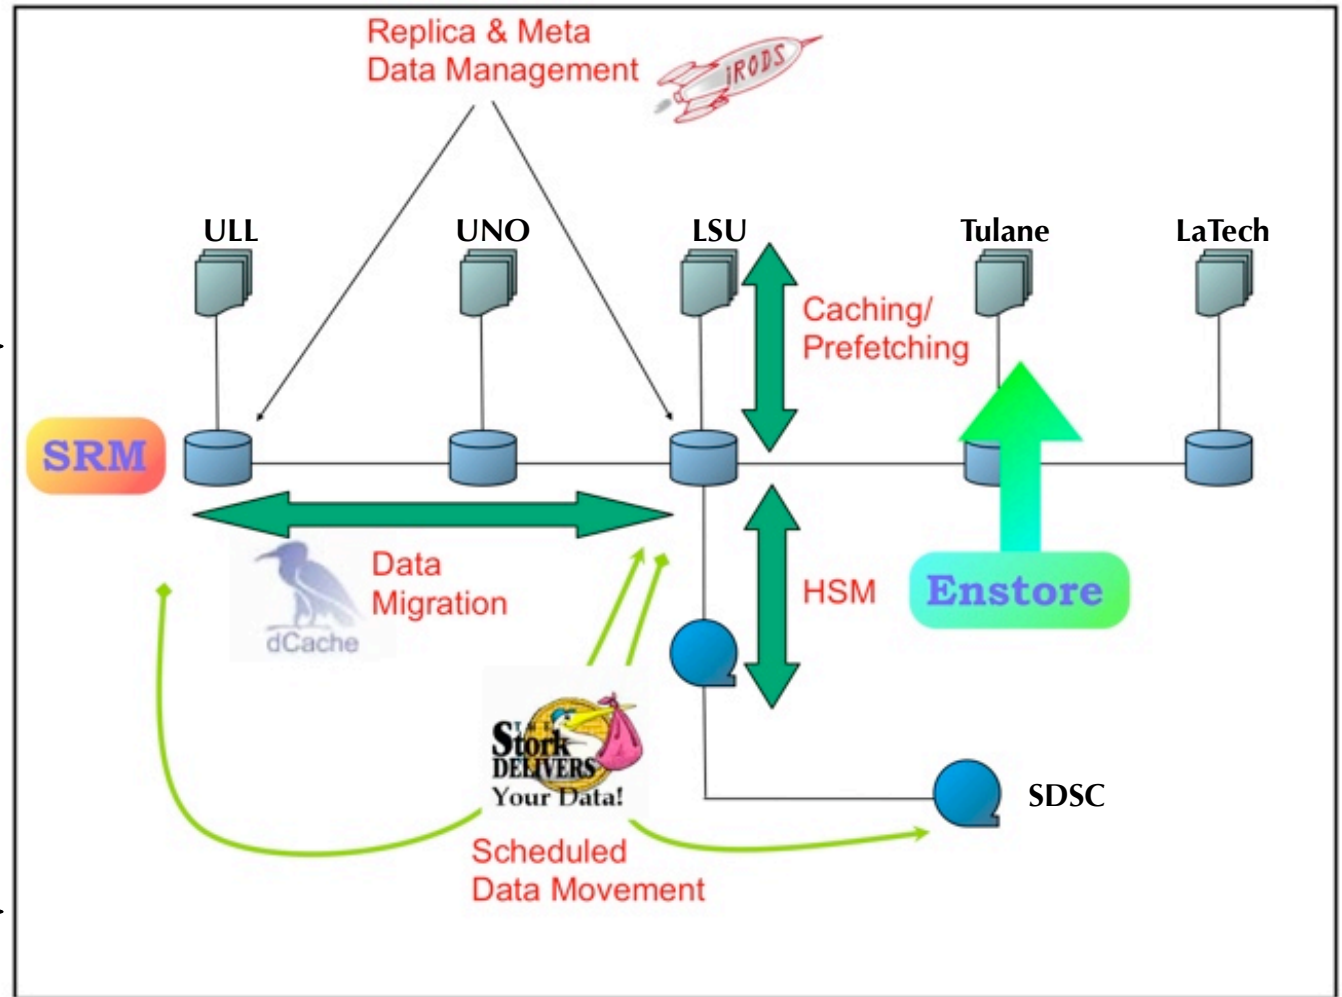
- NO need to change code
- NO relinking
- NO recompiling
- NO privileged access

petashell

user-level virtual FS:

- NO need to change code
- NO relinking
- NO recompiling

petafs



Web interface:

PetaSearch

petashell

- a POSIX compatible shell interface to PetaShare

```
$ petashell
```

```
psh% cp /tmp/foo.txt /petashare/tulane/tmp/foo.txt
```

```
psh% vi /petashare/tulane/tmp/foo.txt
```

```
psh% cp /tmp/foo2.dat /petashare/anysite/tmp/foo2.dat
```

```
psh% genome_analysis genome_data -->
```

```
psh% genome_analysis /petashare/uno/genome_data
```

```
psh% exit
```

```
$
```

PetaSearch

Search Keywords:

- In Archives:
- SCOOP
 - UCoMS
 - NumRel
 - Digital Media
 -

 - All Archives

Accepting Allocation Proposals

1) Title of the Project:

2) Short Description of the Project (Approximately 250 words):

3) Project Web Page (if any):

4) Project PI and Affiliation:

5) Other Senior Personnel:

6) If none of the project participants are associated with PetaShare, please specify a PetaShare contact person:

7) Amount of Storage Allocation Asked:

- In Short Term (first 6 months):

- In Long Term (after 6 months):

8) Preferred Storage Site(s):

LSU (available)

UNO (available)

ULL (soon)

Tulane (soon)

LSUHSC (soon)

LaTech (after May)

Other _____

9) Check if your application includes any of the following:

MPI jobs

Batch jobs

Real-time Visualization

Data Streaming

10) Please specify from which platforms you will be accessing this storage:



Linux

AIX


Windows

Other _____

Send an email to: kosar@cct.lsu.edu




A system driven by the local needs (in LA), but has potential to be a generic solution for the broader community!



For more information on PetaShare:
<http://www.petashare.org>

The Large Hadron Collider (LHC)

A screenshot of a classroom discussion interface. It features a video player at the top showing a person in a white shirt. Below the video is a chat window with a transcript of a discussion. The transcript includes the following text:

Carmen: This is the core they started with, and they want it to look like that. I'm going to first flip it up. Which what happens. Okay. Now I'm going to go back to where we started, and now I'm going to... flip it down. Hmm. What happens each time?
Student: It's the same.
Student: It's... it goes to the same thing.
Carmen: Flipping it up, or flipping. Now, do you think we have to test it on their core square?
Class: Yeah.
Carmen: Alright. Okay, now, are they the same?
Class: Yes. No. Yes.
Carmen: I'm going to flip this one up, maybe, there, I'm gonna flip that one up and I'm going to flip this one down.
Class: Same! Same!

A snippet from a news article titled "Astronomers Detect New Category of Elusive 'Brown Dwarfs'". The text includes "2MASSW J1217-03" and "infrared view". A diagonal banner across the image reads "NATIONAL SCIENCE FOUNDATION".

A collection of astronomical images. On the left is a 3D model of a star cluster with red, green, blue, and yellow stars. On the right is a network diagram with nodes and connecting lines. A red circle highlights a specific feature in a grayscale image.