



Computational Science in Louisiana


Gabrielle Allen (& Dan Katz)
Department of Computer Science
Center for Computation & Technology
Louisiana State University
<http://www.cct.lsu.edu/~gallen>



PITAC Report Summary (June 2005):

“Computational science -- the use of advanced computing capabilities to understand and solve complex problems -- is critical to scientific leadership, economic competitiveness, and national security. It is one of the most important technical fields of the 21st century because it is essential to advances throughout society.”

“Universities must significantly change organizational structures: multidisciplinary & collaborative research are needed [for US] to remain competitive in global science”




Complex problems: Innovations will occur at boundaries



Louisiana Response: Center for Computation & Technology (CCT)


- State commitment for IT in 2002
 - \$25M/year for Vision 20/20, \$9M to LSU
 - University commitment to build new programs
 - Opportunity to build new world class program in interdisciplinary research & education
- Ed Seidel recruited to LSU, created & implemented vision for state-wide collaboration
 - Center for Computation & Technology (2003)
 - Louisiana Optical Network Initiative (2005)
 - CyberTools (2007), LONI Institute (2007)
 - Multidisciplinary Hiring Initiatives (2007)





CCT Plan: Designing a Computational Science Initiative at LSU

- **Cyberinfrastructure:** access to local/national compute, storage, network, visualization resources; support; end-to-end integration.
- **Interdisciplinary research:** university policies on joint appointments, university-wide curricula, appreciation of computational science outcomes (e.g. software, data)
- **Computational research:** faculty who develop or use computation as part of their research; undergraduate and graduate courses to train and educate students





CCT Strategic Plan

- **Mission:** The CCT at LSU is an innovative and interdisciplinary research environment, advancing computational sciences, technologies, and the disciplines they touch. The CCT serves Louisiana through international collaboration, leading progress through revolutionary advancement in academia and industry.
- **Strategic Plan:**
 - Research
 - Education
 - Service
 - Infrastructure
 - Economic Development
- <https://cct.lsu.edu/uploads/CCTStrategicPlan20062010.pdf>

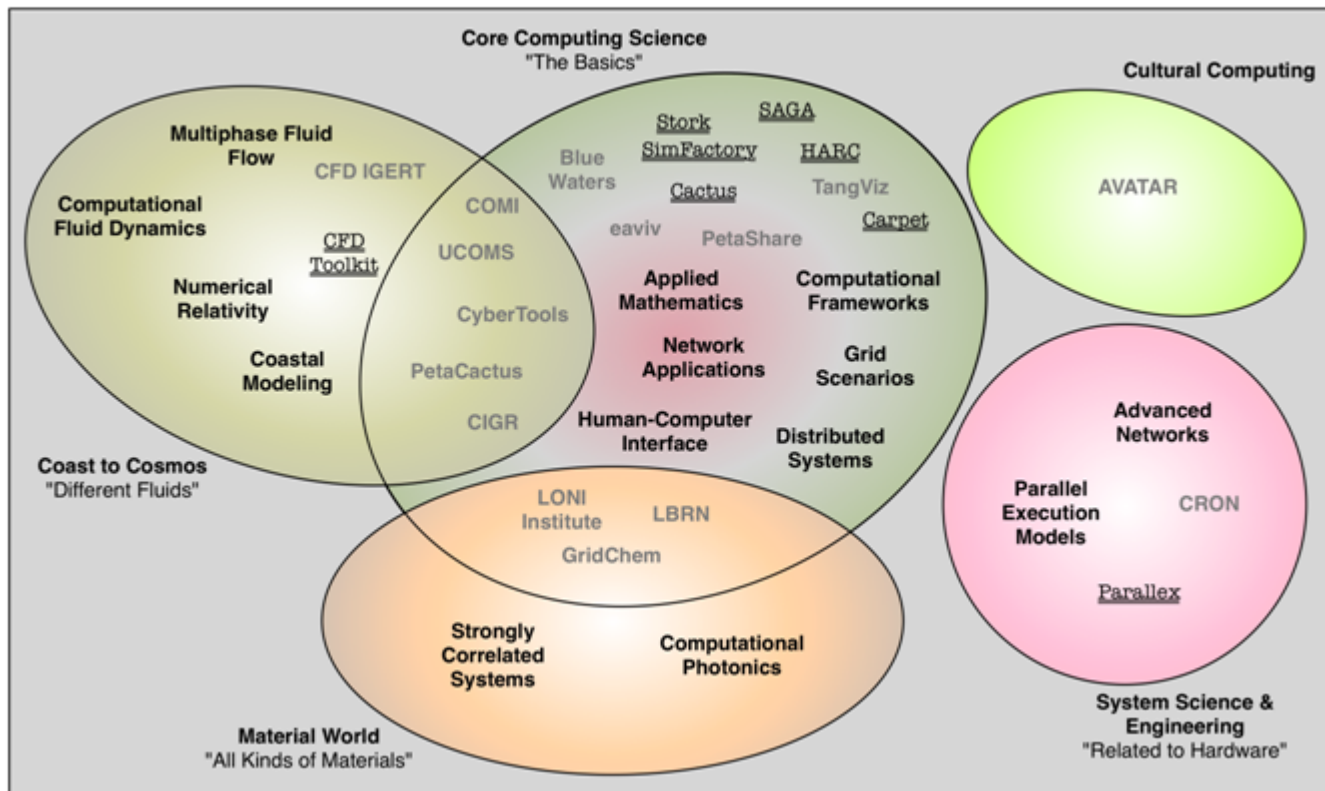


CCT Focus Areas: Faculty Led

Areas chosen to bring together interdisciplinary researchers towards common theme:

Faculty, **Research Staff**, Postdocs, Students


System Science & Engineering / Core Computing Science / Coast to Cosmos / Material World / Cultural Computing





Recruiting Faculty

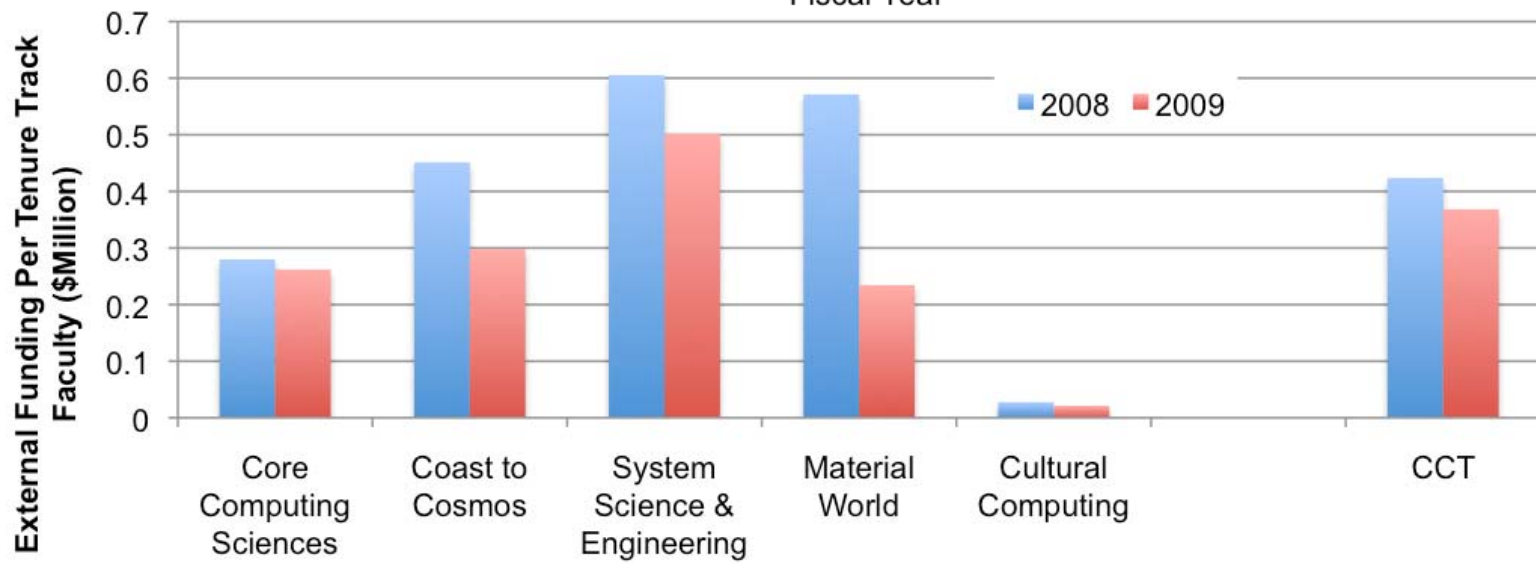

Joint appointments: 50% CCT, 50% home department
Tenure resides in department, but CCT has input in P&T
Reduced teaching load to accommodate higher level service
Also associate faculty from campus





External Funding:

Increasingly able to respond to large solicitations



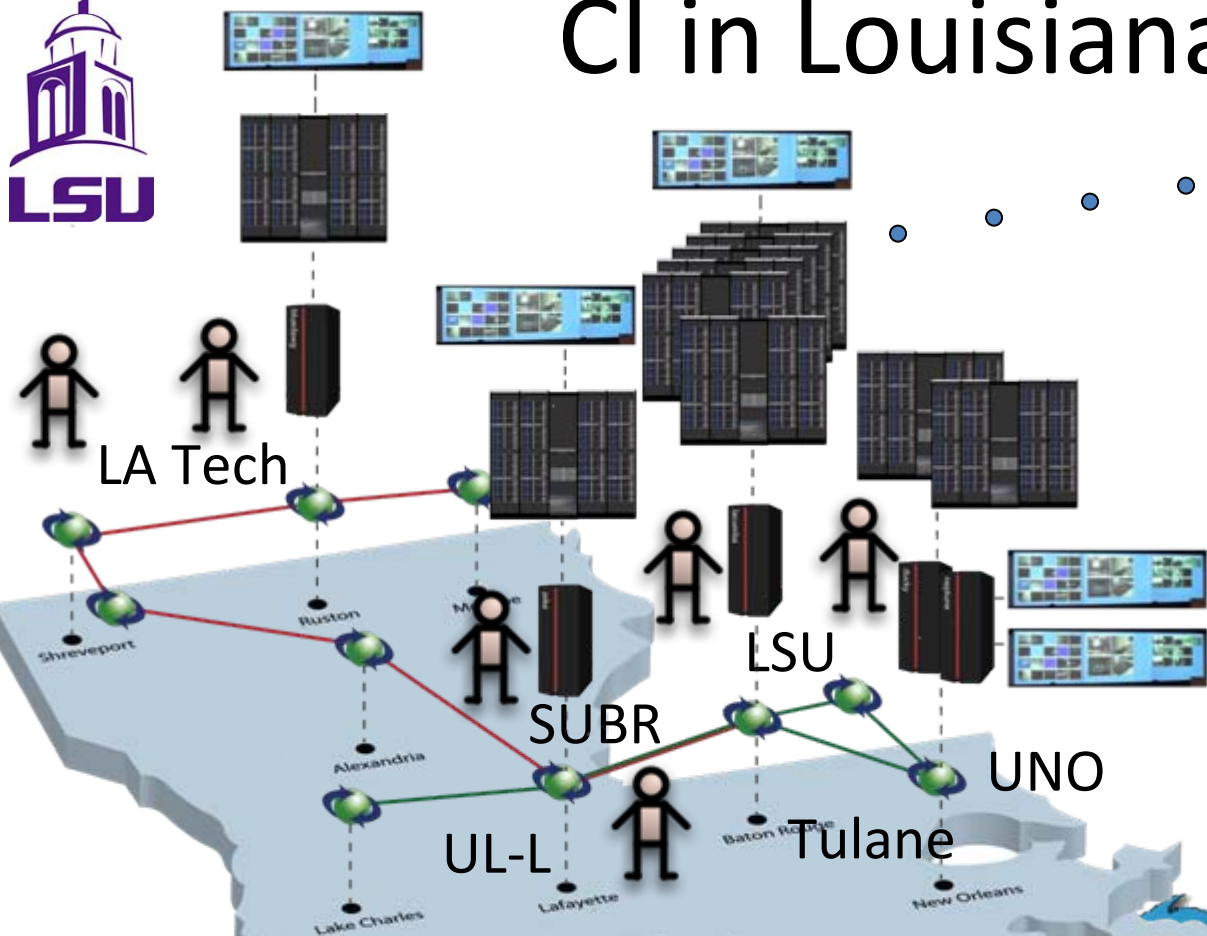


Cyberinfrastructure Development

- Vision: combine research and infrastructure
 - Research
 - Computer science
 - Applications
 - Tools
 - Infrastructure
 - Hardware
 - Operations
 - Policies
- CyD Division at CCT to coordinate this
- CyD staff: PhDs in CS and applications who understand the whole picture and want to grow the ecosystem



CI in Louisiana



TeraGrid

LONI Institute: **People and Collaborations**

Cybertools: **Tools and Services**

LONI: ~100TF IBM, Dell **Supercomputers**

LONI: 40 Gbps **network**

National Lambda Rail





Louisiana Optical Network (LONI)



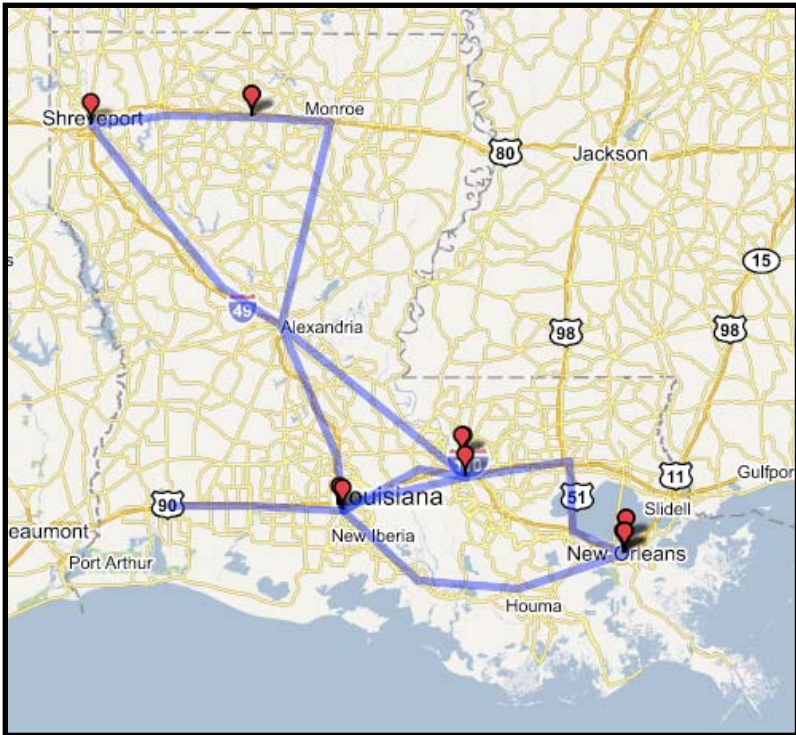
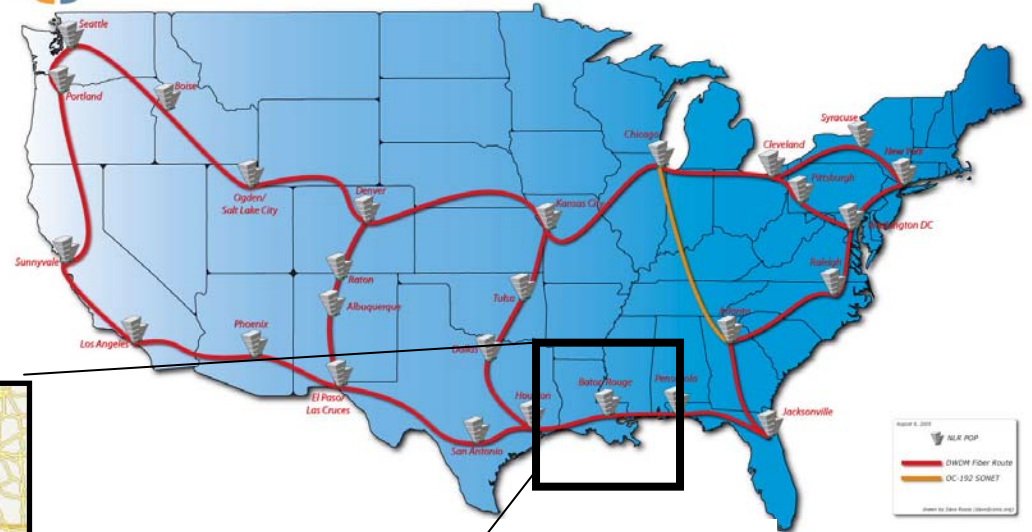
State initiative (\$50M) to support research (2004):

40 Gbps optical network + NLR

Connects state universities, health science centers



National LambdaRail™ Infrastructure



Compute resources: ~100 Tflops across state

Data resources ~500TB with NSF PetaShare

LONI customers: MS universities, K12, hospitals, LPB



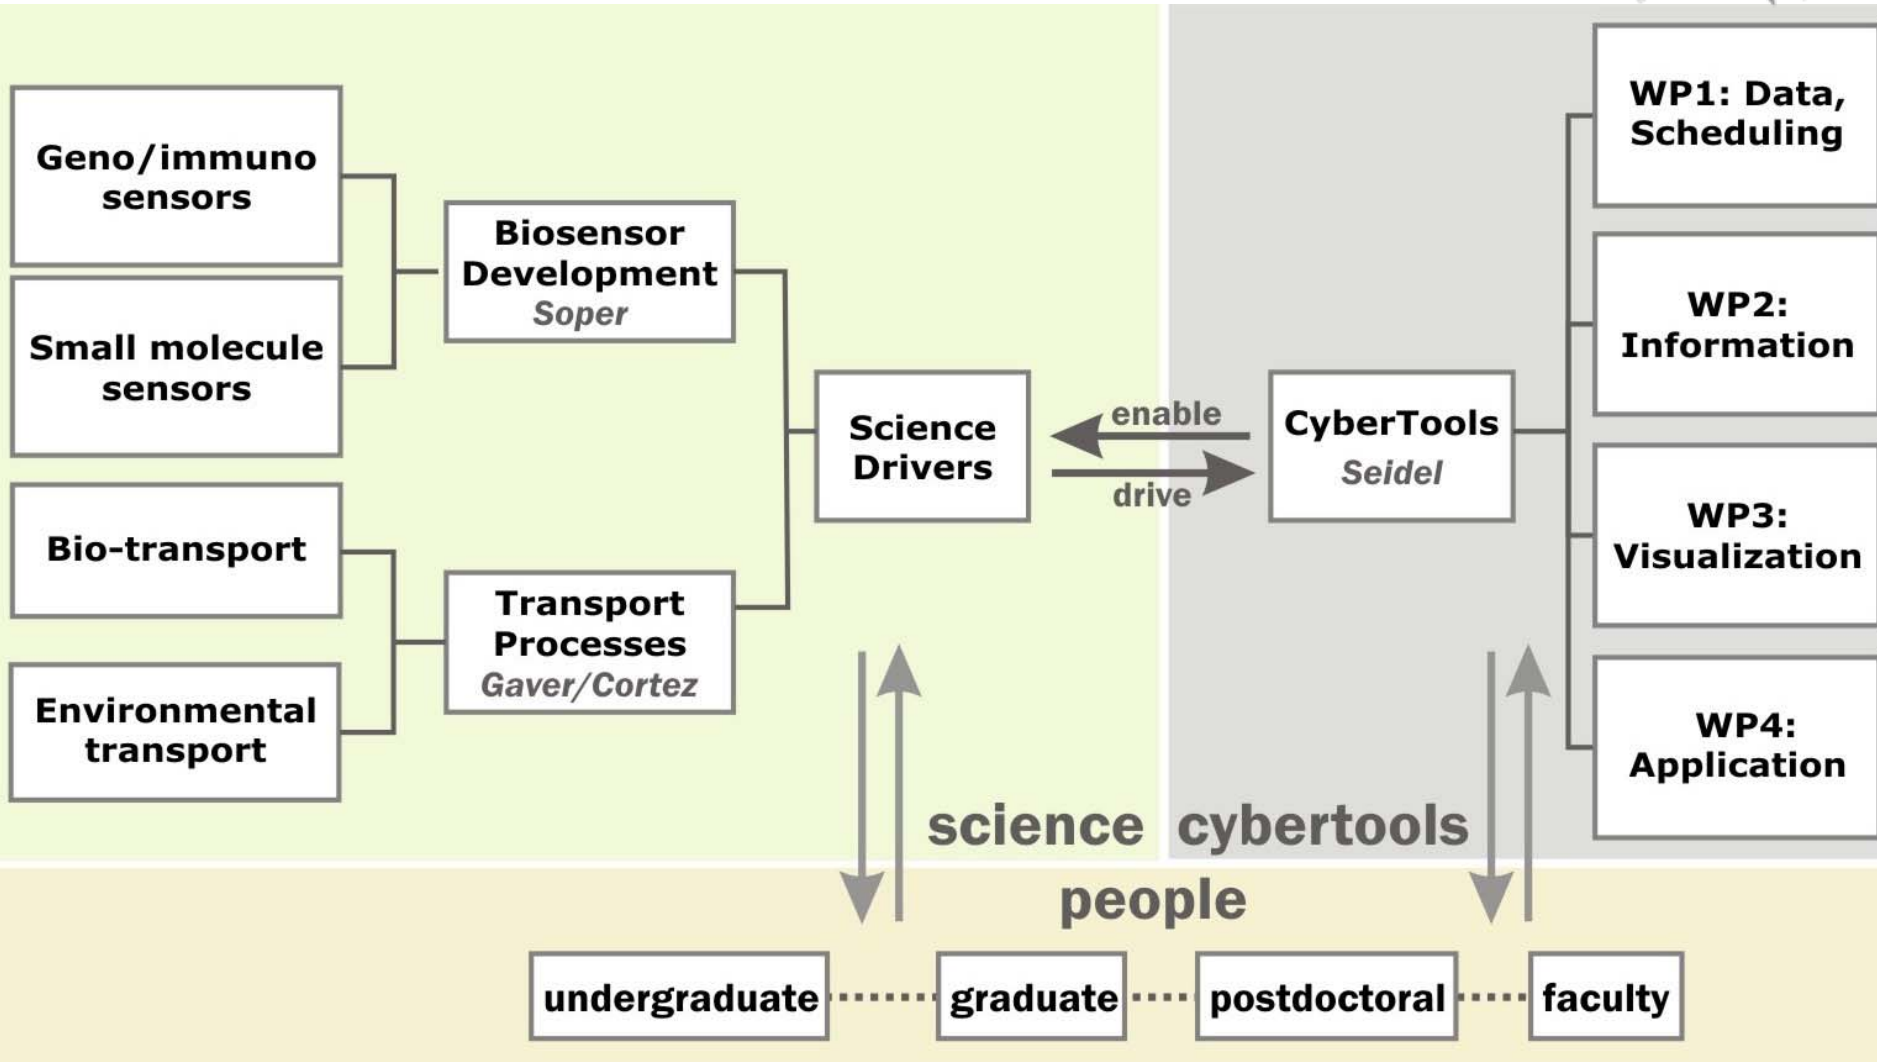
Support: HPC@LSU

- Partnership between LSU CCT and LSU ITS to support and maintain LSU HPC resources and provide user training
 - HPC@LSU also supports/maintains LONI, through support from Louisiana BoR
 - HPC@LSU also supports TeraGrid, through support from NSF



NSF EPSCOR RII

CyberTools Project: Enabler and Driver





Cybertools Infrastructure



- WP1: Scheduling and Data Services
 - Infrastructure deployment, high availability, scheduling, data archiving and retrieval, metadata
- WP2: Information Services and Portals
 - Information services (infrastructure, apps, experiments), application interfaces for scientists, portals for information gathering
- WP3: Visualization services
 - Data/viz integration, HD streaming viz, advanced viz facilities, integration with application toolkits
- WP4: Application toolkits
 - Toolkits to support simulation codes (CFD, MD, other), application managers. SAGA interfaces



Overview

People

Science Drivers

CS Research

Publications

Documents

Downloads

Meeting Notes

Mailing Lists

Announcements

Media Coverage

PetaShare FAQ

DIDC 2010 Workshop



A Distributed Data Archival, Analysis and Visualization
Cyberinfrastructure for Data-intensive Collaborative
Research.

Announcements:

PetaShare Storage is Online
April 21, 2008

PetaShare storage is now online
accepting allocation proposals.

[\[Read More\]](#)

Overview

PetaShare is an NSF sponsored project which responds to the urgent need of scientists who work with large-scale data generation, sharing and collaboration requirements. PetaShare aims to enable domain scientists to focus on their primary research problem, assured that the underlying infrastructure will manage the low-level data handling issues.

PetaShare employs a very novel approach to solve the distributed data sharing and management problem. Unlike existing approaches, PetaShare treats data storage resources and the tasks related to data access as first class entities just like computational resources and compute tasks, and not simply the side effect of computation. The key technologies that are being developed in this project include data-aware storage systems, data-aware schedulers, and cross-domain metadata scheme which take the responsibility of managing data resources and scheduling data tasks from the user and perform these tasks transparently.

An initial prototype of PetaShare is deployed at seven Louisiana campuses: Louisiana State

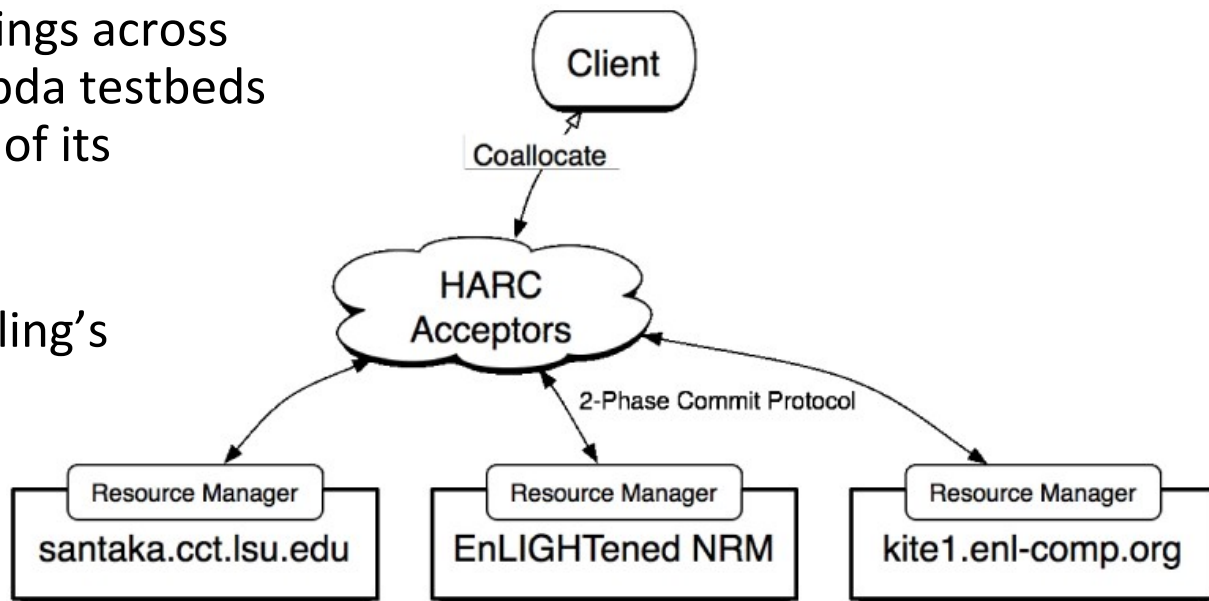
Sponsored by:





HARC: Highly Available Resource Co-scheduler

- Extensible, open-sourced co-allocation system
- Can already reserve:
 - Time on supercomputers (advance reservation), and
 - Dedicated paths on GMPLS-based networks with simple topologies
- Uses Paxos Commit to atomically reserve multiple resources, while providing a highly-available service
- Used to coordinate bookings across EnLIGHTened and G-lambda testbeds in largest demonstration of its kind to date
- Used for setting up the network for Thomas Sterling's HPC Class which goes out live in HD





Vish

The Vish Visualization Environment

Register | Sign In

Home | Screenshots | Forum | Issues | Statistics | Development




The Vish Project

Project Information

VISH is an abstraction interface to build visualization modules independent from a certain application. There is a reference implementation that allows to build a standalone application.

Category:	visualization
Operating System:	Windows, Linux
License:	Light & License for free academic and personal use
Programming Language:	C++
Project Description:	CGAP

Welcome to the VISH Project!



Viz Tangibles



Above left: co-located and distributed users collaboratively manipulate a 3D visualization in an AccessGrid meeting using viz tangibles interaction devices

Right: an “interaction tray” is used together with RFID-tagged “tangible menus” to access and manipulate scientific visualizations





Simple API for Grid Applications

SAGA

A Simple API for Grid Applications

[Site Map](#) [Accessibility](#) [Contact](#)

only in current section

[Home](#)

[Software](#)

[Documentation](#)

[Publications](#)

[Community](#)

[News](#)

[Projects](#)

[Log in](#)

You are here: [Home](#)

Software

- [C++ Libraries](#)
- [JAVA Libraries](#)
- [Python Modules](#)

About

- [The Team](#)
- [Funding & History](#)
- [OGF SAGA WG](#)

Publications

- [Papers](#)
- [Presentation Slides](#)
- [OGF Documents](#)

Welcome


SAGA is an API that provides the basic functionality required to build distributed applications, tools and frameworks so as to be independent of the details of the underlying infrastructure. SAGA can be used to provide simple access layers for distributed systems and abstractions for applications and thereby address the fundamental application design objectives of Interoperability across different infrastructure, Distributed Scale-Out, Extensibility, Adaptivity whilst preserving simplicity.


If you want to learn more about SAGA and how it's being used by scientific applications across numerous Grid infrastructures, watch this short introductory [video clip](#) produced for the Supercomputing Conference 2009.

Implementations

Currently, two native open source implementations and several language bindings for the SAGA API standard ([GWD-R.90](#)) are available:


News

 [SAGA C++ 1.4.1 Released](#)
Apr 09, 2010

 [SAGA C++ 1.4 Released](#)
Dec 01, 2009

[More news...](#)

Upcoming Events

 [SAGA Tutorial](#)
Web streaming, access grid, and Johnston 218 (CCT LSU),
Apr 28, 2010

[Previous events...](#)

[Upcoming events...](#)



Cactus Framework

The screenshot shows the homepage of the Cactus Framework website. The header is blue with the 'cactus code' logo in white and green. Navigation links include 'Home | Contact', 'About', 'Media', 'Demo', 'Download', 'Documentation', 'Community', and 'Internal'. The main content area is divided into two columns. The left column contains a 'Welcome' section with three paragraphs of text. The right column features a 'Google Custom Search' box, a 'Recent News' section with four news items, and a search button.

Home | Contact

cactus code

About Media Demo Download Documentation Community Internal

Welcome

Cactus is an open source problem solving environment designed for scientists and engineers. Its modular structure easily enables parallel computation across different architectures and collaborative code development between different groups. Cactus originated in the academic research community, where it was developed and used over many years by a large international collaboration of physicists and computational scientists.

The name Cactus comes from the design of a central core ("flesh") which connects to application modules ("thorns") through an extensible interface. Thorns can implement custom developed scientific or engineering applications, such as computational fluid dynamics. Other thorns from a standard computational toolkit provide a range of computational capabilities, such as parallel I/O, data distribution, or checkpointing.

Cactus runs on many architectures. Applications, developed on standard workstations or laptops, can be seamlessly run on clusters or supercomputers. Cactus provides easy access to many cutting edge software technologies being developed in the academic research community, including the [Globus](#) Metacomputing Toolkit, [HDF5](#) parallel file I/O, the [PETSc](#) scientific library, [adaptive mesh refinement](#), web interfaces, and advanced visualization tools.

Google Custom Search

Recent News

7 November 2009
[Whisky Retreat V](#)

17 March 2009
[Webcast: *From Black Holes to Gamma-Ray Bursts*, hosted by SiCortex](#)

3 February 2009
[Cactus 4.0 beta 16 released](#)

2 December 2008
[Cactus team demonstrates Alpaca tools at SuperComputing 2008](#)

16 August 2008



LONI Institute

- \$15M 5-year project (2007-2012)
 - \$7M BoR, \$8M from universities
- Create bold new inter-university superstructure
 - New faculty (12), staff (6), students (36); train others. Focus on CS, Bio, Materials, but all disciplines impacted
 - Promote research at interfaces for innovation
- Draw on, enhance strengths of all universities
 - Solve complex problems through collaboration & computation
 - Much stronger recruiting opportunities for all institutions
 - Statewide interdisciplinary education & research program
- Create University-Industry Research Centers (UIRC)s
- Transform our state
 - Such committed cooperation between sites extraordinary



Conclusions

- Cyberinfrastructure has played a crucial role in Louisiana in building an environment able to support and nurture high quality research
- Advances were needed in many areas:
 - Networks, Computing, Tools & services, Applications, People
- Real results:
 - REU program in computational science, 16 undergrads
 - 2 of 16 Blue Waters undergraduate internships to LA
 - PetaApps, PRAC, CAREER awards
 - Louisiana is developing cyberinfrastructure for its own researchers, and for the nation (Cactus, PetaShare, SAGA, HARC, Vish,)



References

- Louisiana: A Model for Advancing Regional e-Research through Cyberinfrastructure,
 - D. S. Katz, G. Allen, R. Cortez, C. Cruz-Neira, R. Gottumukkala, Z. D. Greenwood, L. Guice, S. Jha, R. Kolluru, T. Kosar, L. Leger, H. Liu, C. McMahon, J. Nabrzyski, B. Rodriguez-Milla, E. Seidel, G. Speyrer, M. Stubblefield, B. Voss, S. Whittenburg, *Phil. Trans. R. Soc. A* 28, June 2009, vol. 367, no. 1897, 2459-2469, (2009).
- Computational Science, Infrastructure and Interdisciplinary Research on University Campuses: Experiences and Lessons from the Center for Computation & Technology,
 - G. Allen, D. S. Katz,, *CCT Technical Report Series, CCT-TR-2010-1*, 2010 (Full version soon)