SciServer

Collaborative data-driven science

Collaborative Analysis of Petabyte-Scale Scientific Data

Jordan Raddick

Johns Hopkins University

Johns Hopkins University

NSF

idies

Motivation

- Much of the scientific process is about data
  - collect, clean, archive, organize, analyze, publish, mirror, distribute, curate…
- Jim Gray: “Often it turns out to be more efficient to move the questions than the data.”

Motivation

- Helen Shen: “Brown's students write explanatory text and intersperse it with raw code and the charts and Sitting in the airplane… Brown can interact with the work..”

***Nature**, 5 November 2014*

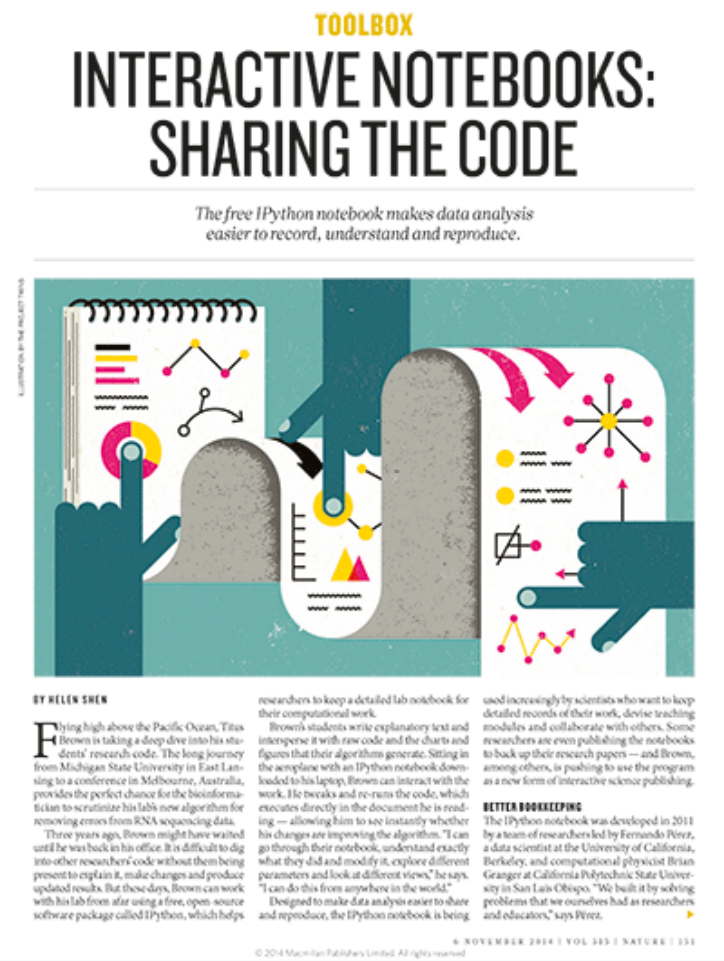

TOOLBOX

INTERACTIVE NOTEBOOKS: SHARING THE CODE

The free IPython notebook makes data analysis easier to record, understand and reproduce.

BY HELEN SHEN

Flying high above the Pacific Ocean, Titus Brown is taking a deep dive into his students' research code. The long journey from Michigan State University in East Lansing to a conference in Melbourne, Australia, provides the perfect chance for the bioinformatician to scrutinize his lab's new algorithm for removing errors from RNA sequencing data.

Three years ago, Brown might have waited until he was back in his office. It is difficult to dig into other researchers' code without them being present to explain it, make changes and produce updated results. But these days, Brown can work with his lab from afar using a free, open-source software package called IPython, which helps researchers to keep a detailed lab notebook for their computational work.

Brown's students write explanatory text and intersperse it with raw code and the charts and figures that their algorithms generate. Sitting in the aeroplane with an IPython notebook downloaded to his laptop, Brown can interact with the work. He tweaks and re-runs the code, which executes directly in the document he is reading — allowing him to see instantly whether his changes are improving the algorithm. "I can go through their notebook, understand exactly what they did and modify it, explore different parameters and look at different views," he says. "I can do this from anywhere in the world."

Designed to make data analysis easier to share and reproduce, the IPython notebook is being used increasingly by scientists who want to keep detailed records of their work, devise teaching modules and collaborate with others. Some researchers are even publishing the notebooks to back up their research papers — and Brown, among others, is pushing to use the program as a new form of interactive science publishing.

BETTER BOOKKEEPING

The IPython notebook was developed in 2011 by a team of researchers led by Fernando Pérez, a data scientist at the University of California, Berkeley, and computational physicist Brian Granger at California Polytechnic State University in San Luis Obispo. "We built it by solving problems that we ourselves had as researchers and educators," says Pérez.

History

- Started with the SDSS SkyServer (2001)
- Goal: instant access to rich content
- 2012: NSF DIBBS to extend/reengineer SkyServer

SkyServer in 2001

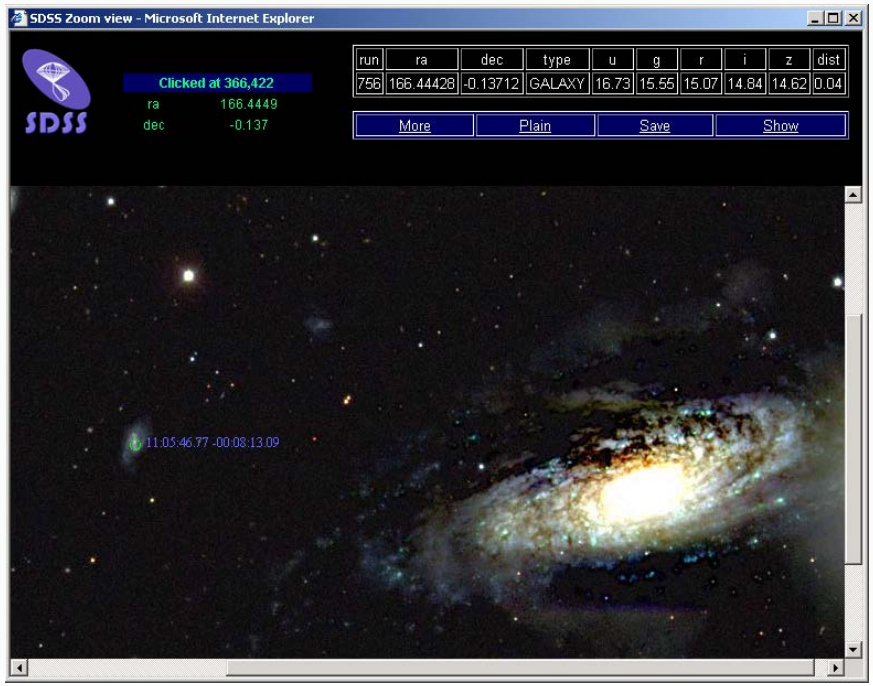

SkyServer today

Data from HPC Simulations

- HPC is an instrument in its own right
- Need public access to the best and latest
  - Largest simulations approach/exceed petabytes
  - Also need ensembles of simulations
- Creates new challenges
  - Access?
  - Data lifecycle?
  - Analysis patterns?
  - Architectural support?
- Our total science data: ~2.5PB
  - Everything is a Big Data problem!

Cosmological Simulations

- Mirror of Millennium DB (universe sim)

Raw data:
Particles

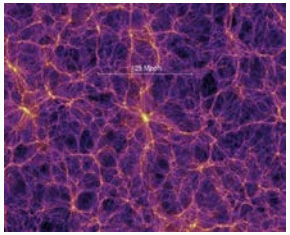

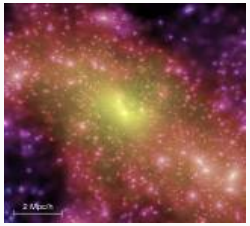

FOF groups
Subhalos

Halo merger trees

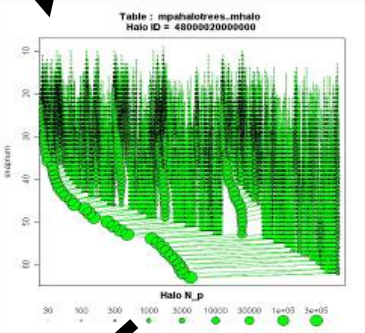

Density fields

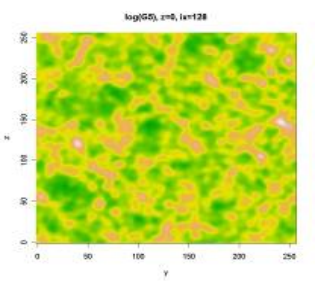

Synthetic galaxies

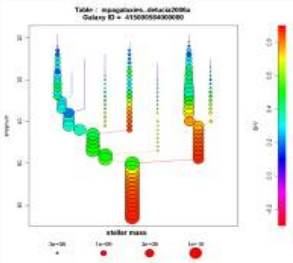

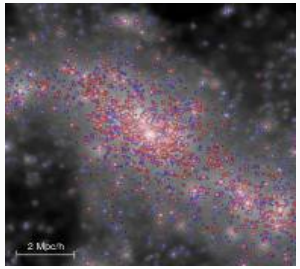

Mock images

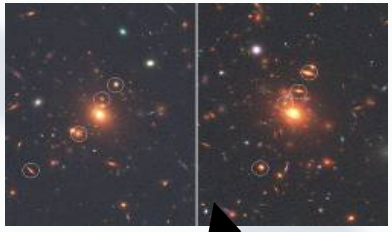

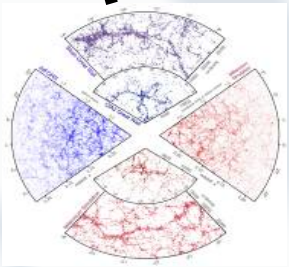

Mock catalogs

Turbulence databases (JHUTB)

- http://turbulence.pha.jhu.edu/

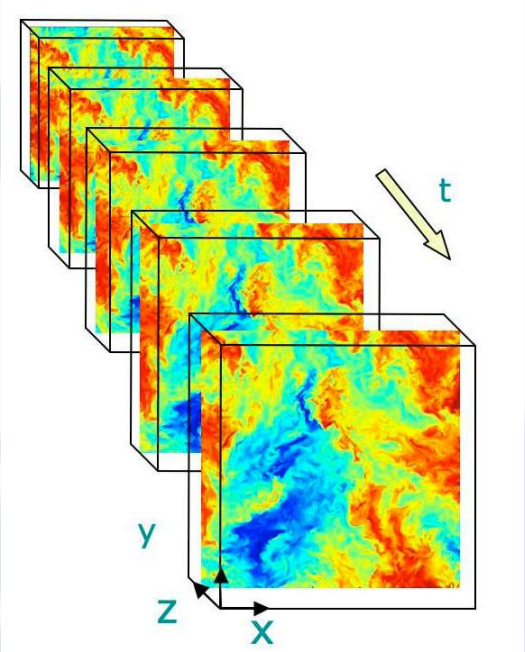

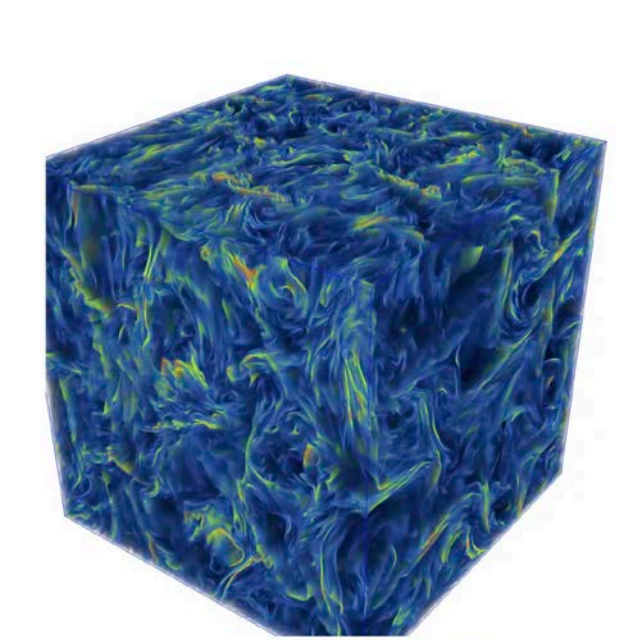

Current field $\nabla \times \mathbf{B}$

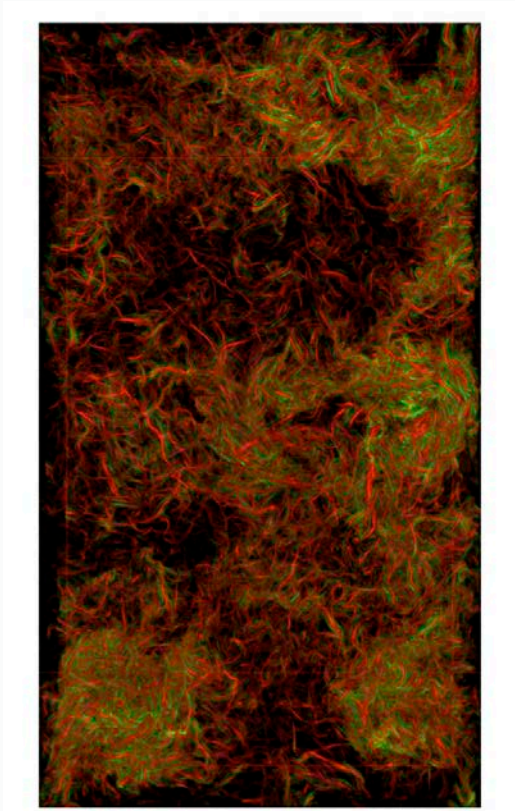

Vorticity $\nabla \times \mathbf{u}$

Next step: Numerical Laboratories

- From data retrieval to data **analysis**
  - No need to download entire datasets
  - Analysis is server-side through web services
- Use virtual sensor metaphor (“the cow in the tornado”)
  - Many access patterns are local in space and/or time
- turbulence.pha.jhu.edu:
  19 trillion points delivered!

Oceanography

Hydrostatic and non-hydrostatic simulations of dense waters cascading off a shelf: The East Greenland case Marcello G. Magaldi, Thomas W.N. Haine

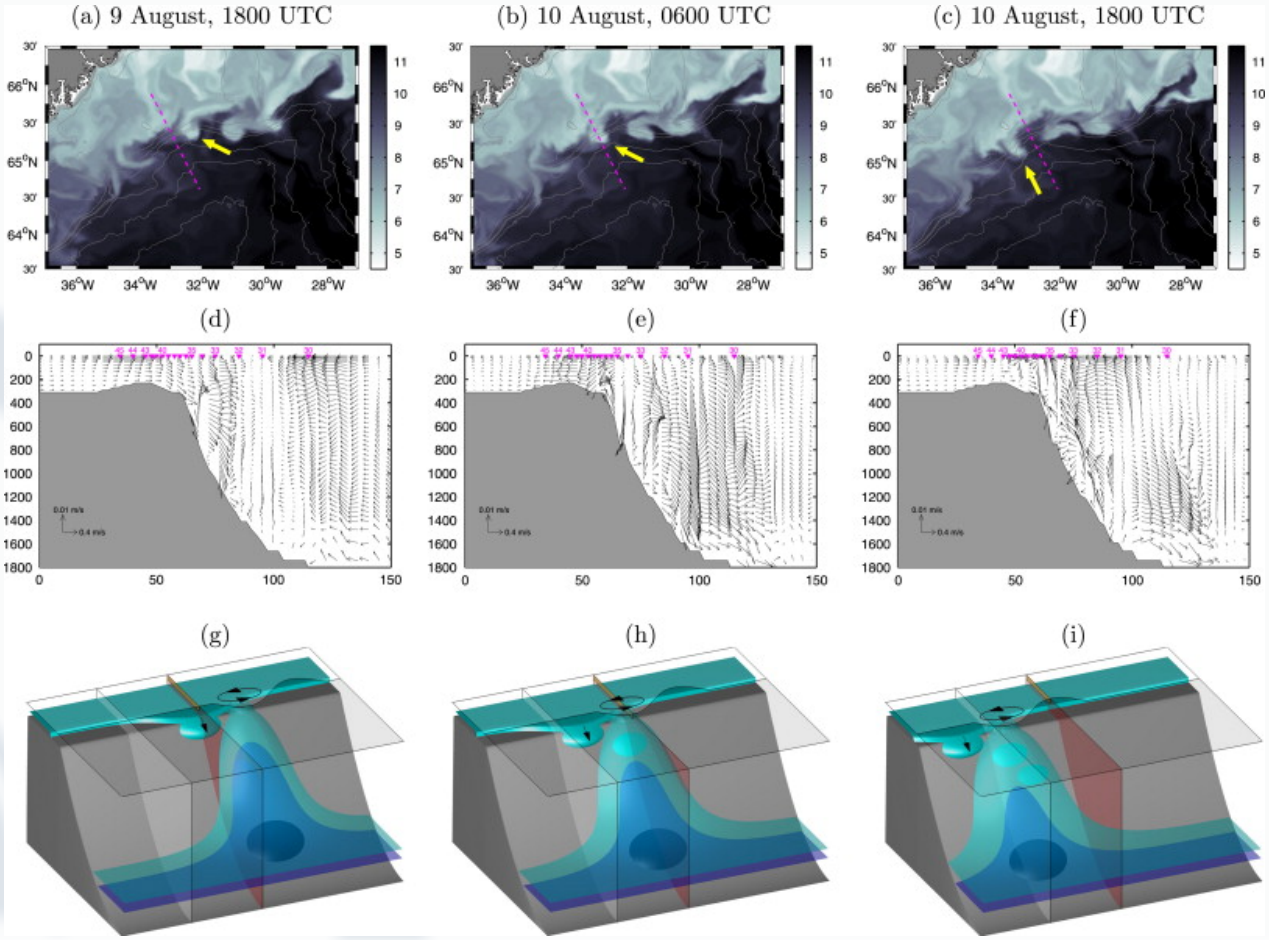

Genomics

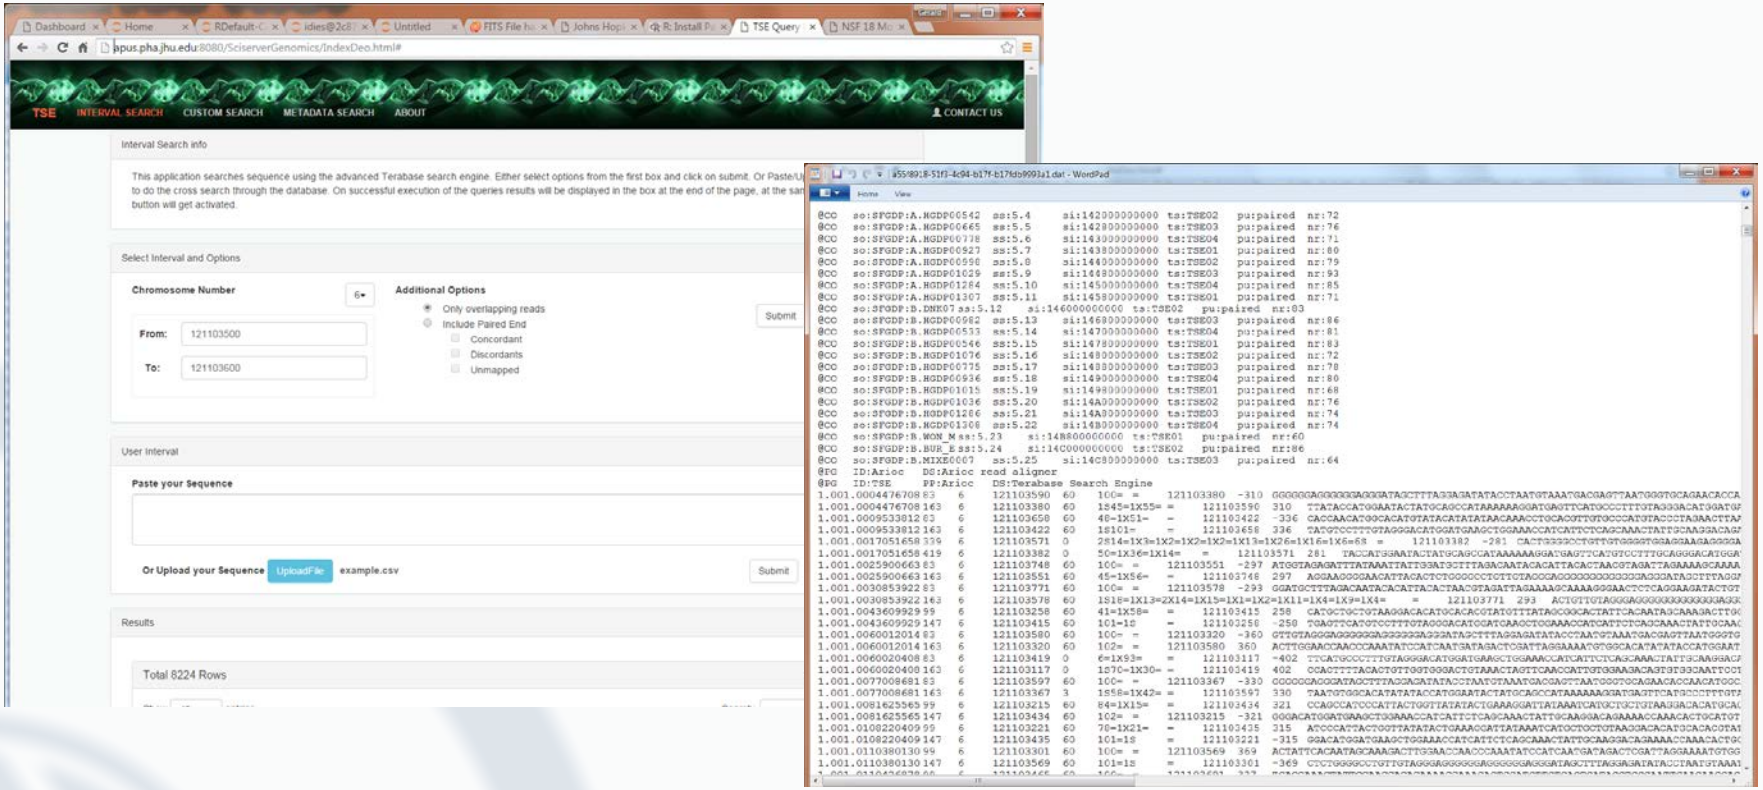

Building a DB of a trillion short reads from Next Gen Sequencing

Materials Science

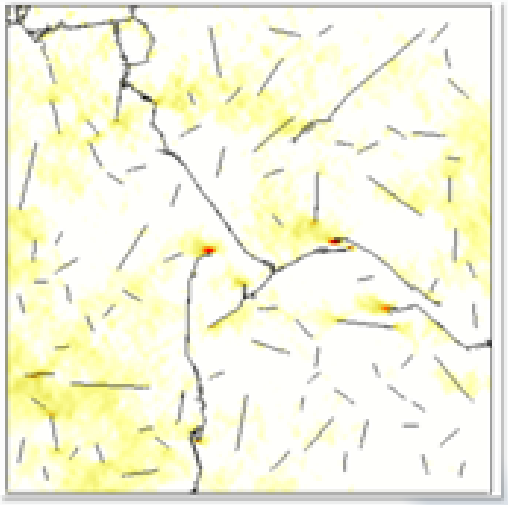

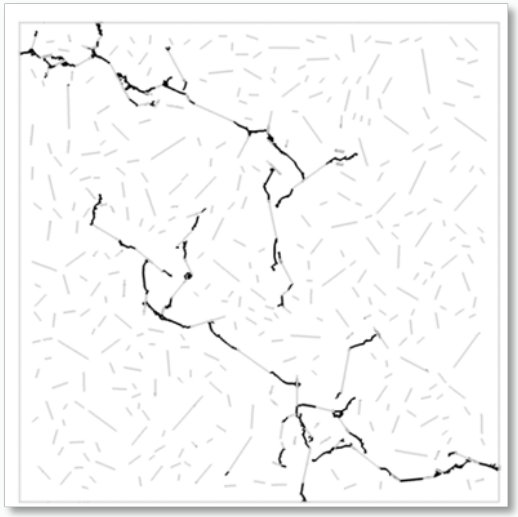

Daphalapurkar, Brady, Ramesh, Molinari. JMPS (2011)

SciServer Project Objectives

- Extend infrastructure to support additional science domains
- Host and serve Petabyte datasets
- Support custom user datasets
- Provide access and query services
- Provide scalable compute services
- Support analyses and data sets too large to handle locally
- Provide collaborative tools for shared analysis

Computations stay CLOSE to the DATA…

SciServer Project Components

Major Components	Supporting Technologies
Core	Microsoft SQL Server
• Login Portal	Open Stack
• CASJobs	Docker
• SciServer Compute	Jupyter
• SciDrive	
Applications	
• SkyQuery	
• SkyServer	
• GLUSEEN	
• Turbulence	

SciServer Project Timeline

Timelines

Year 1 (2013-2014)	Project Setup, Scoping, Planning, Begin Refactoring, SDSS Unification
Year 2 (2013-2014)	Architectural Refactoring – API, Single Sign-on, prototype Compute
Year 3 (2013-2014)	**SciServer System Release, Interactive Compute, Scalable Job Management, Basic Dashboard, Initial Collaborative capabilities**
Year 4 (2013-2014)	Implementation in Science Domains, Educational workbooks
Year 5 (2013-2014)	System Scale out, Data Analytics, Advanced Deployment Scenarios

NOW

SciServer Project Current Plans

Timelines – Year 3

Apr 2016	• SciServer System Release
May 2016	• Interactive Compute • SkyQuery • Gluseen
August 2016	• Prototype Scalable Job Management • Basic Dashboard • Initial Collaborative capabilities
October 2016	• Scalable Job Management • Turbulence • Cosmology
November 2016	• Project 3 year Review

Login Portal

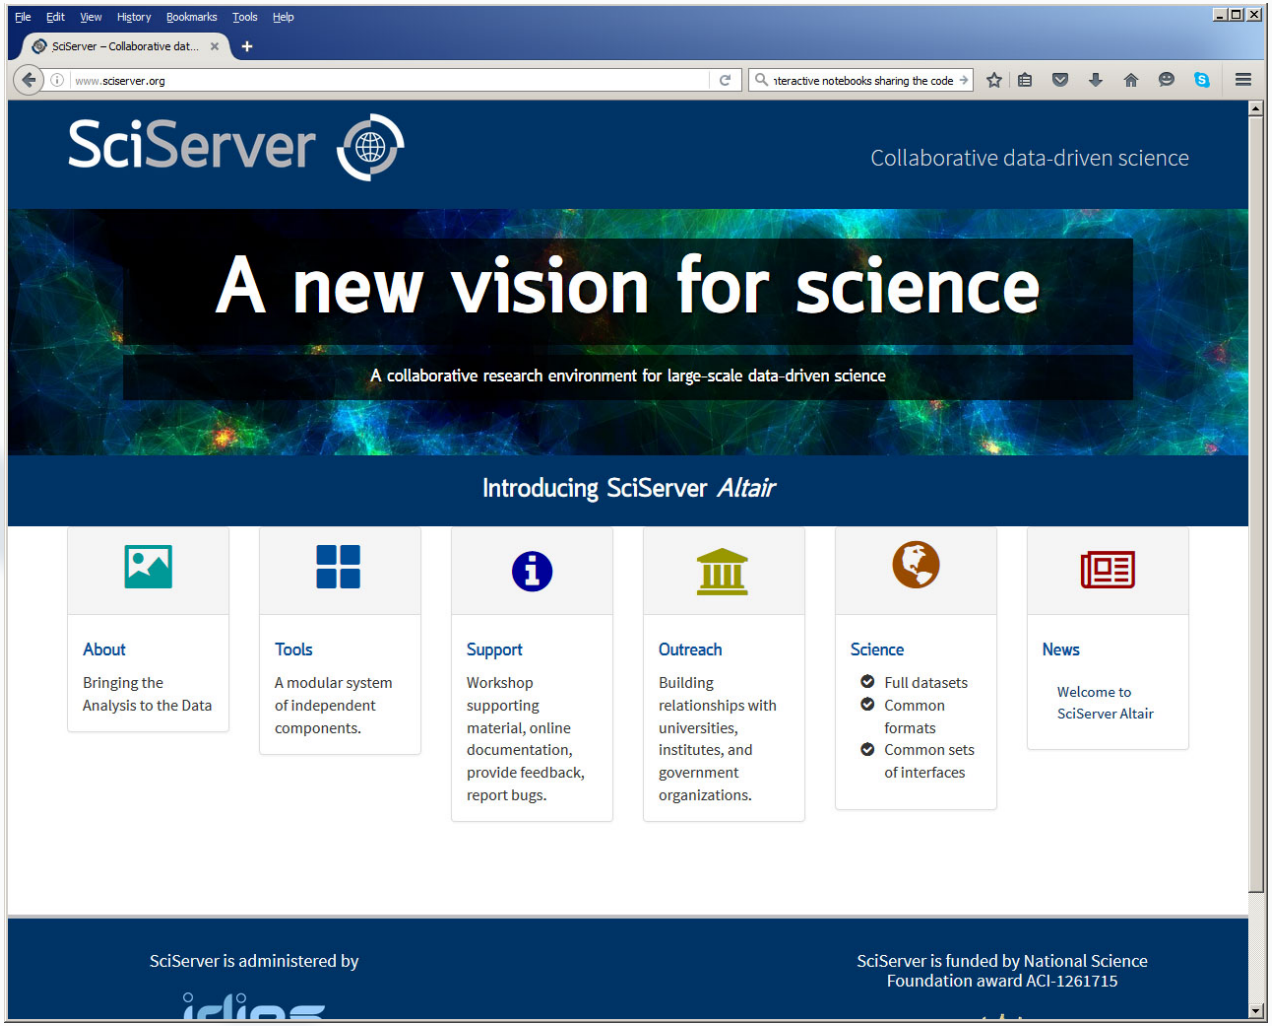

SkyServer

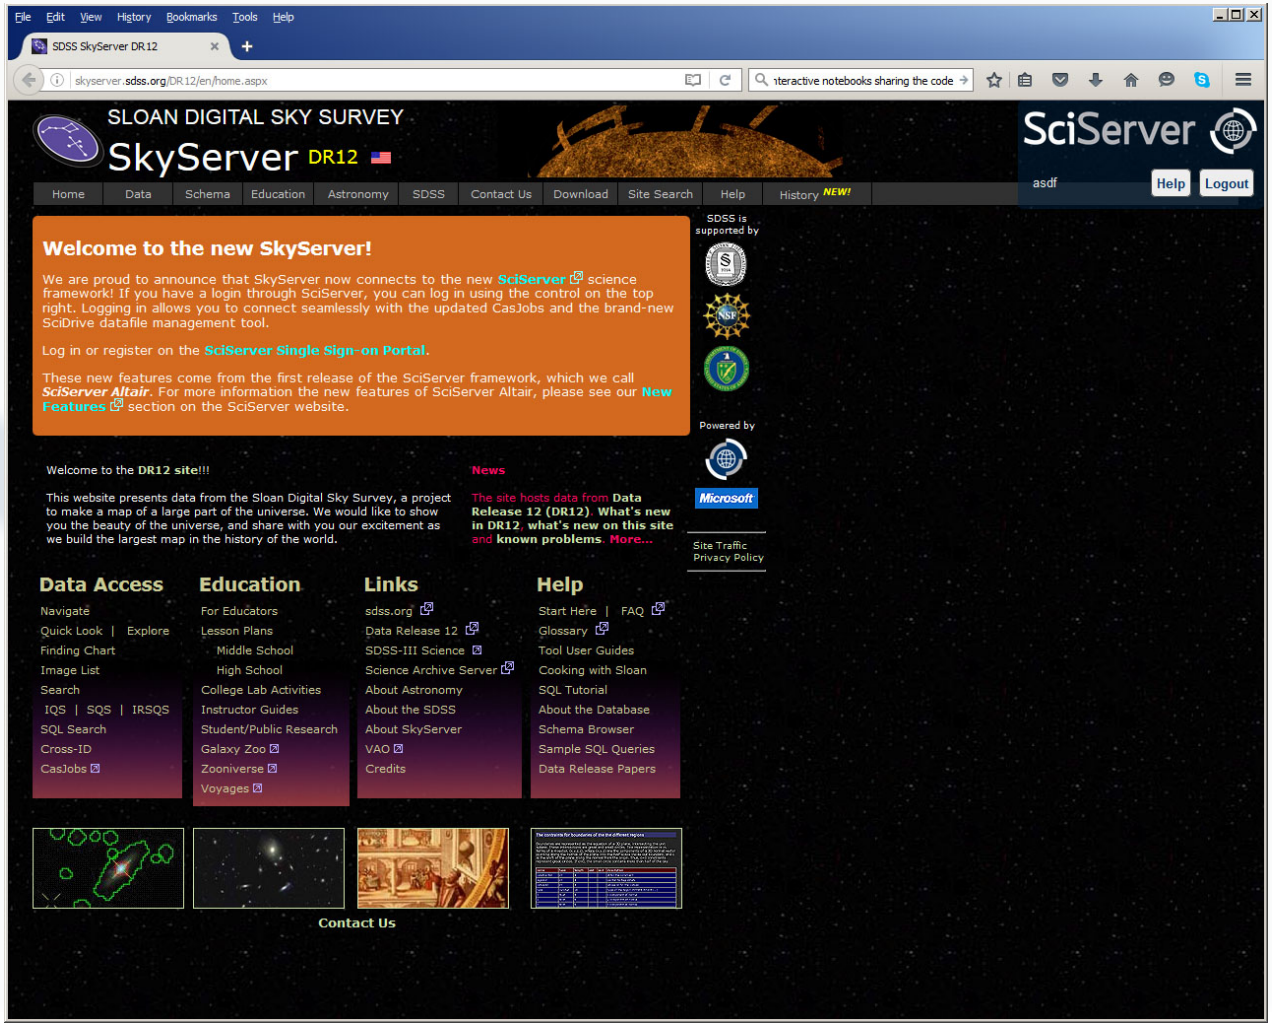
File Edit View History Bookmarks Tools Help
SDSS SkyServer DR12
skyserver.sdss.org/DR12/en/home.aspx
SLOAN DIGITAL SKY SURVEY
SkyServer DR12
SciServer
asdf
Help
Logout
Home
Data
Schema
Education
Astronomy
SDSS
Contact Us
Download
Site Search
Help
History NEW!
Welcome to the new SkyServer!
We are proud to announce that SkyServer now connects to the new SciServer science framework! If you have a login through SciServer, you can log in using the control on the top right. Logging in allows you to connect seamlessly with the updated CasJobs and the brand-new SciDrive datafile management tool.
Log in or register on the SciServer Single Sign-on Portal.
These new features come from the first release of the SciServer framework, which we call SciServer Altair. For more information the new features of SciServer Altair, please see our New Features section on the SciServer website.
SDSS is supported by
Powered by
Microsoft
Site Traffic
Privacy Policy
Welcome to the DR12 site!!!
This website presents data from the Sloan Digital Sky Survey, a project to make a map of a large part of the universe. We would like to show you the beauty of the universe, and share with you our excitement as we build the largest map in the history of the world.
News
The site hosts data from Data Release 12 (DR12). What's new in DR12, what's new on this site and known problems. More...
Data Access
Navigate
Quick Look | Explore
Finding Chart
Image List
Search
IQS | SQS | IRSQS
SQL Search
Cross-ID
CasJobs
Education
For Educators
Lesson Plans
Middle School
High School
College Lab Activities
Instructor Guides
Student/Public Research
Galaxy Zoo
Zooniverse
Voyages
Links
sdss.org
Data Release 12
SDSS-III Science
Science Archive Server
About Astronomy
About the SDSS
About SkyServer
VAO
Credits
Help
Start Here | FAQ
Glossary
Tool User Guides
Cooking with Sloan
SQL Tutorial
About the Database
Schema Browser
Sample SQL Queries
Data Release Papers
Contact Us

CasJobs: Query

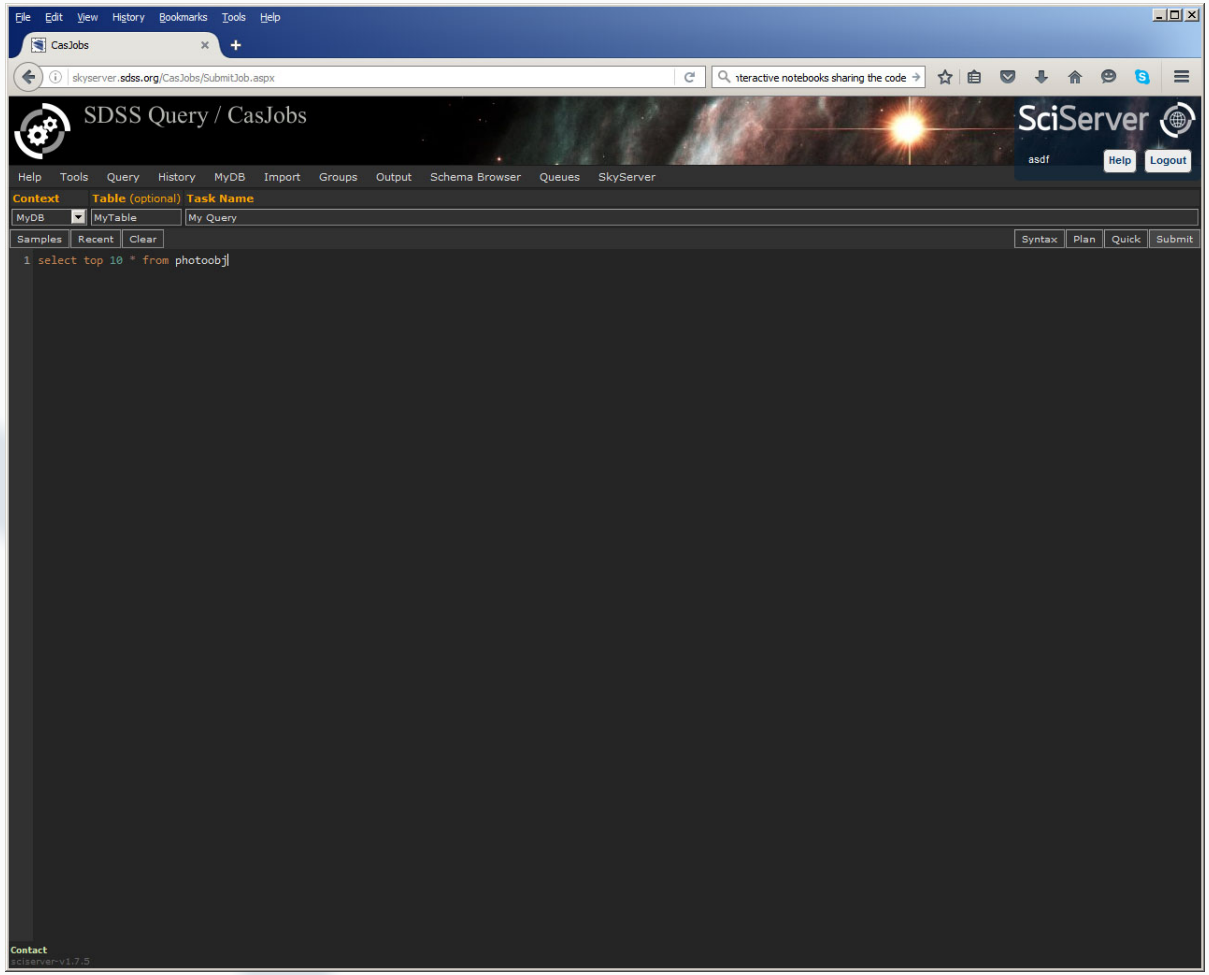

CasJobs: MyDB

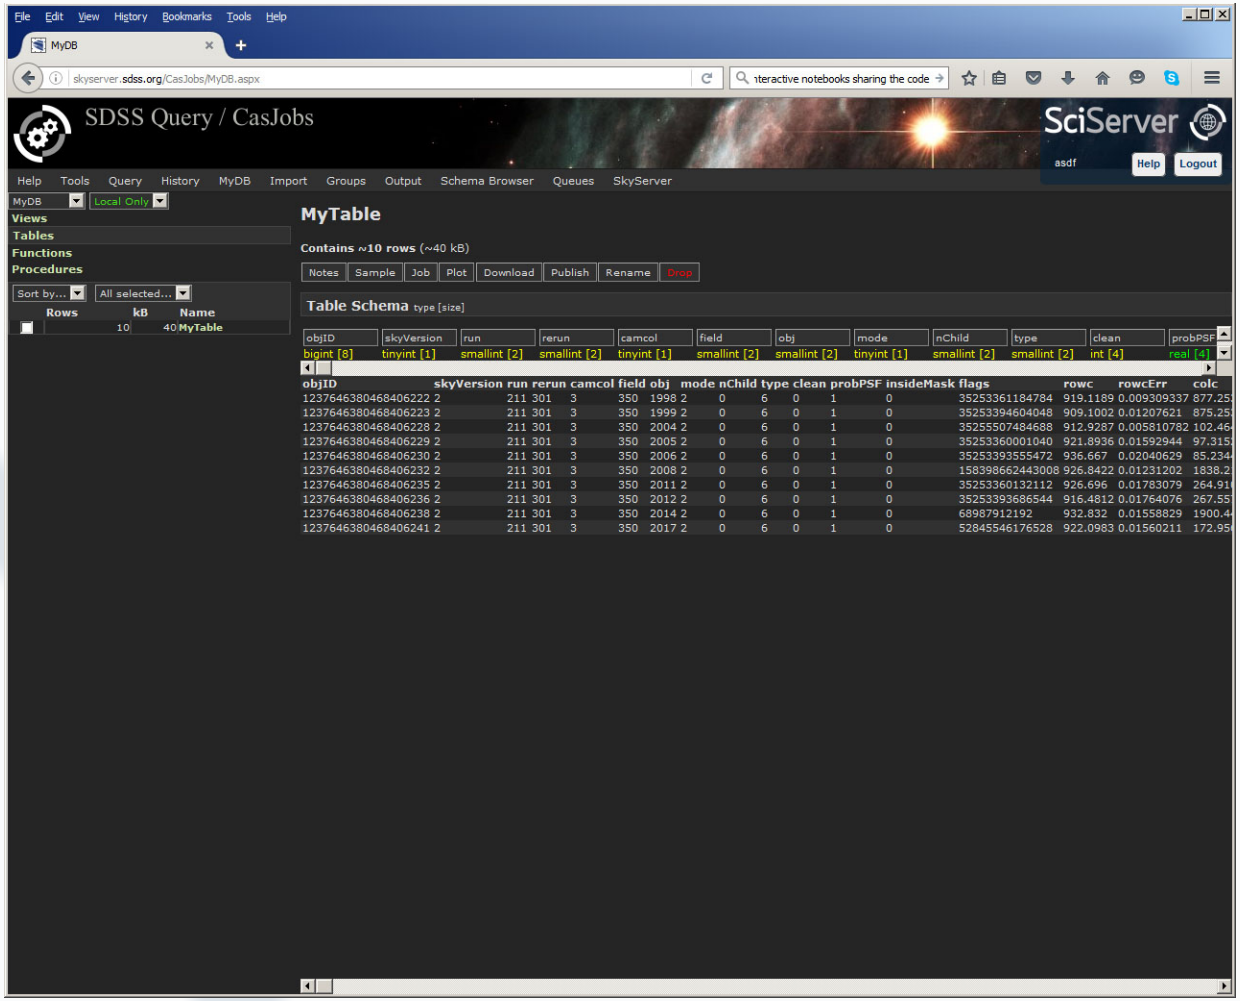

SciDrive

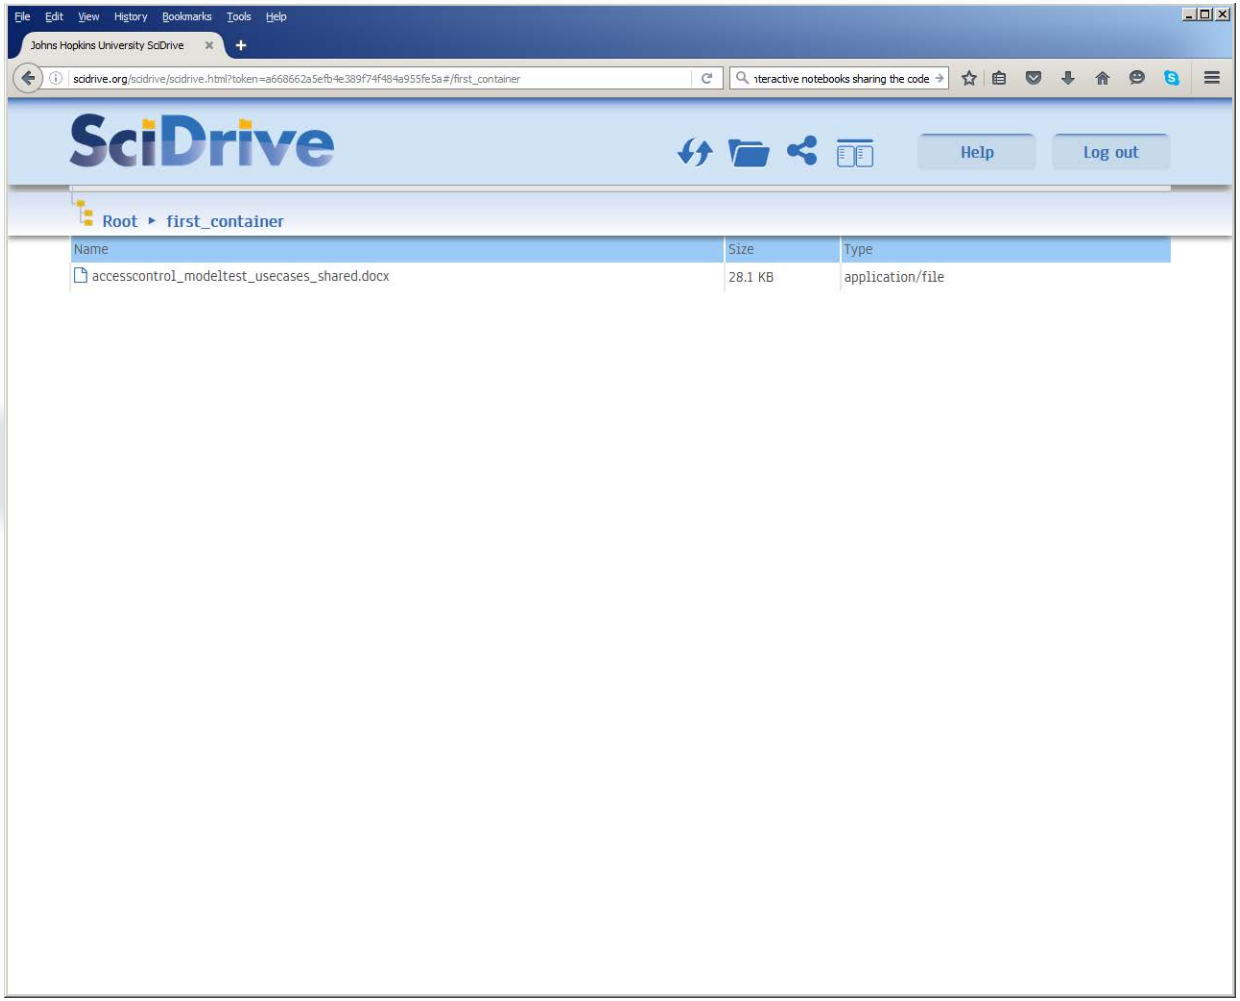
File Edit View History Bookmarks Tools Help
Johns Hopkins University SciDrive
scidrive.org/scidrive/scidrive.html?token=a668662a5efb4e389f74f484a955fe5a#/first_container
SciDrive
Help
Log out
Root ▸ first_container
Name
Size
Type
accesscontrol_modeltest_usecases_shared.docx
28.1 KB
application/file

What is SciServer Compute?

- Interactive Jupyter notebooks hosted inside Docker containers.
- Pre-configured images to create new containers from (R, Python, MATLAB, …).
- High-bandwidth, low-latency access to other SciServer services and data sources through the notebooks.
- Users manage their own containers.

What are Docker Containers?

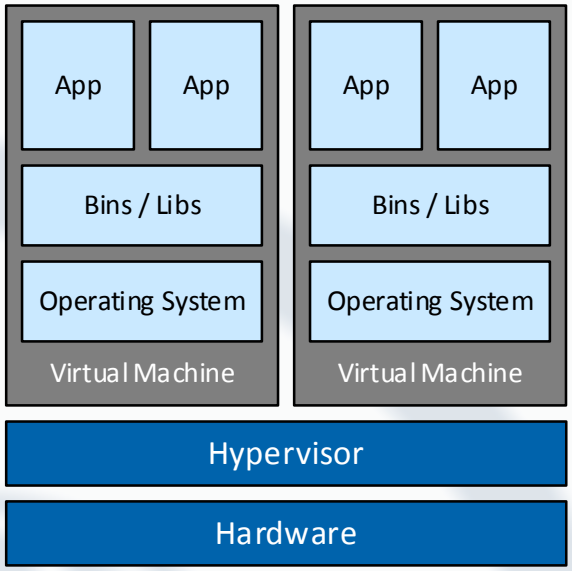

Type 1 Hypervisor

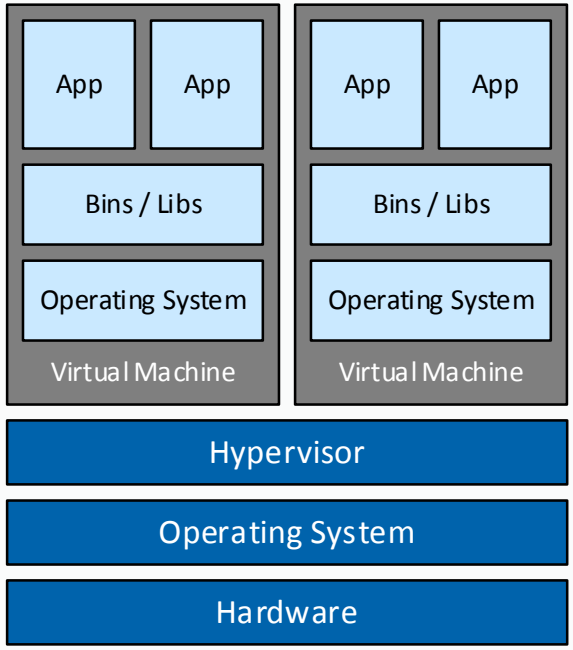

Type 2 Hypervisor

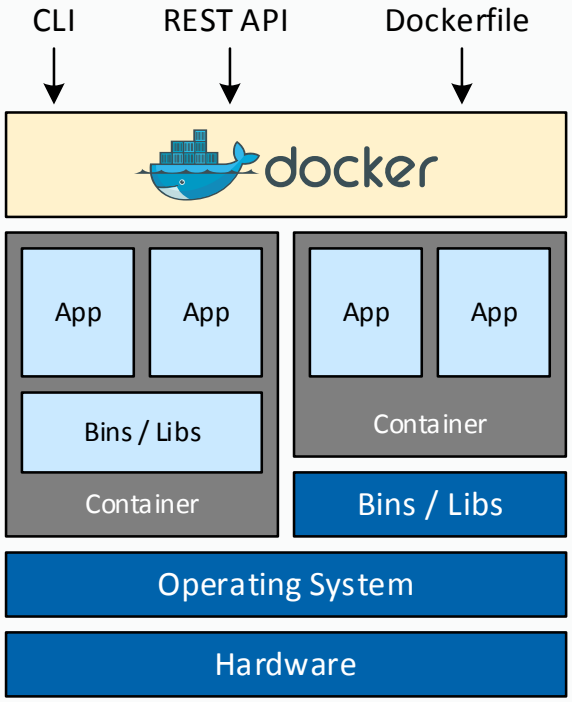

Linux Containers

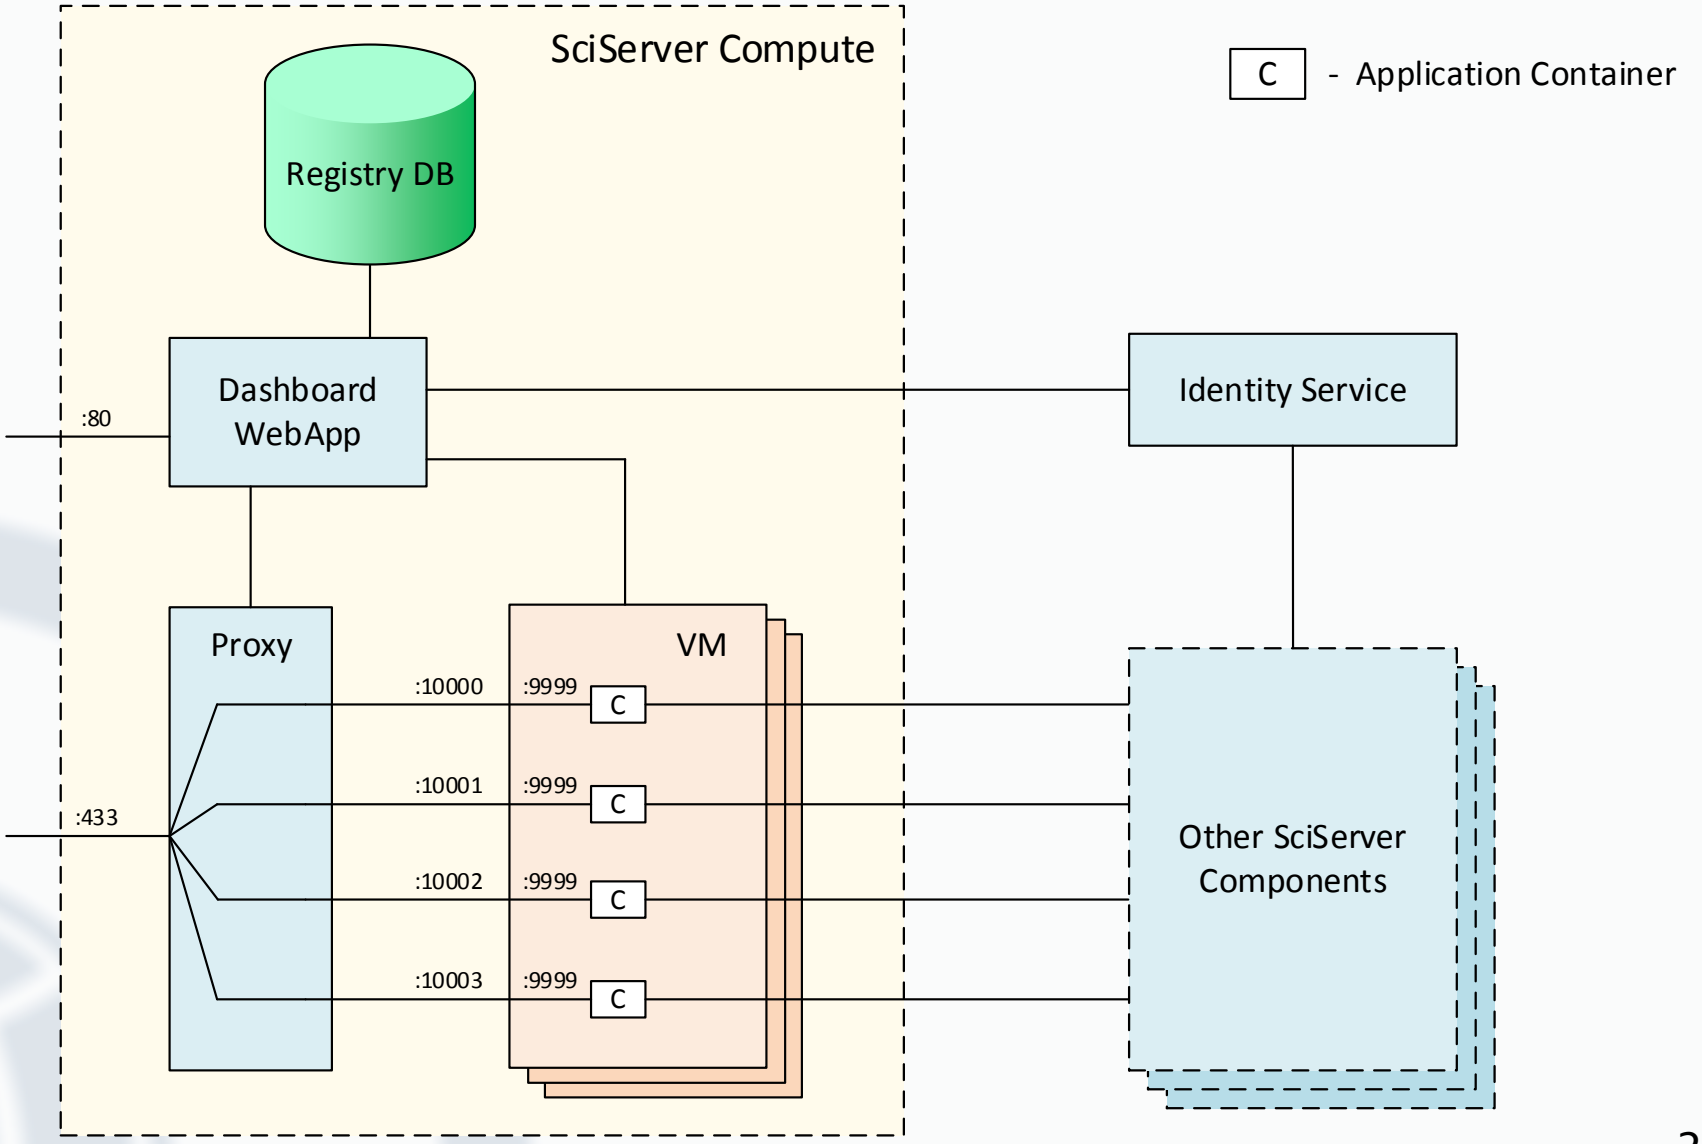
SciServer Compute
C - Application Container
Registry DB
Dashboard WebApp
:80
Identity Service
Proxy
VM
:433
:10000
:9999
C
:10001
:9999
C
:10002
:9999
C
:10003
:9999
C
Other SciServer Components

Data Storage Configuration

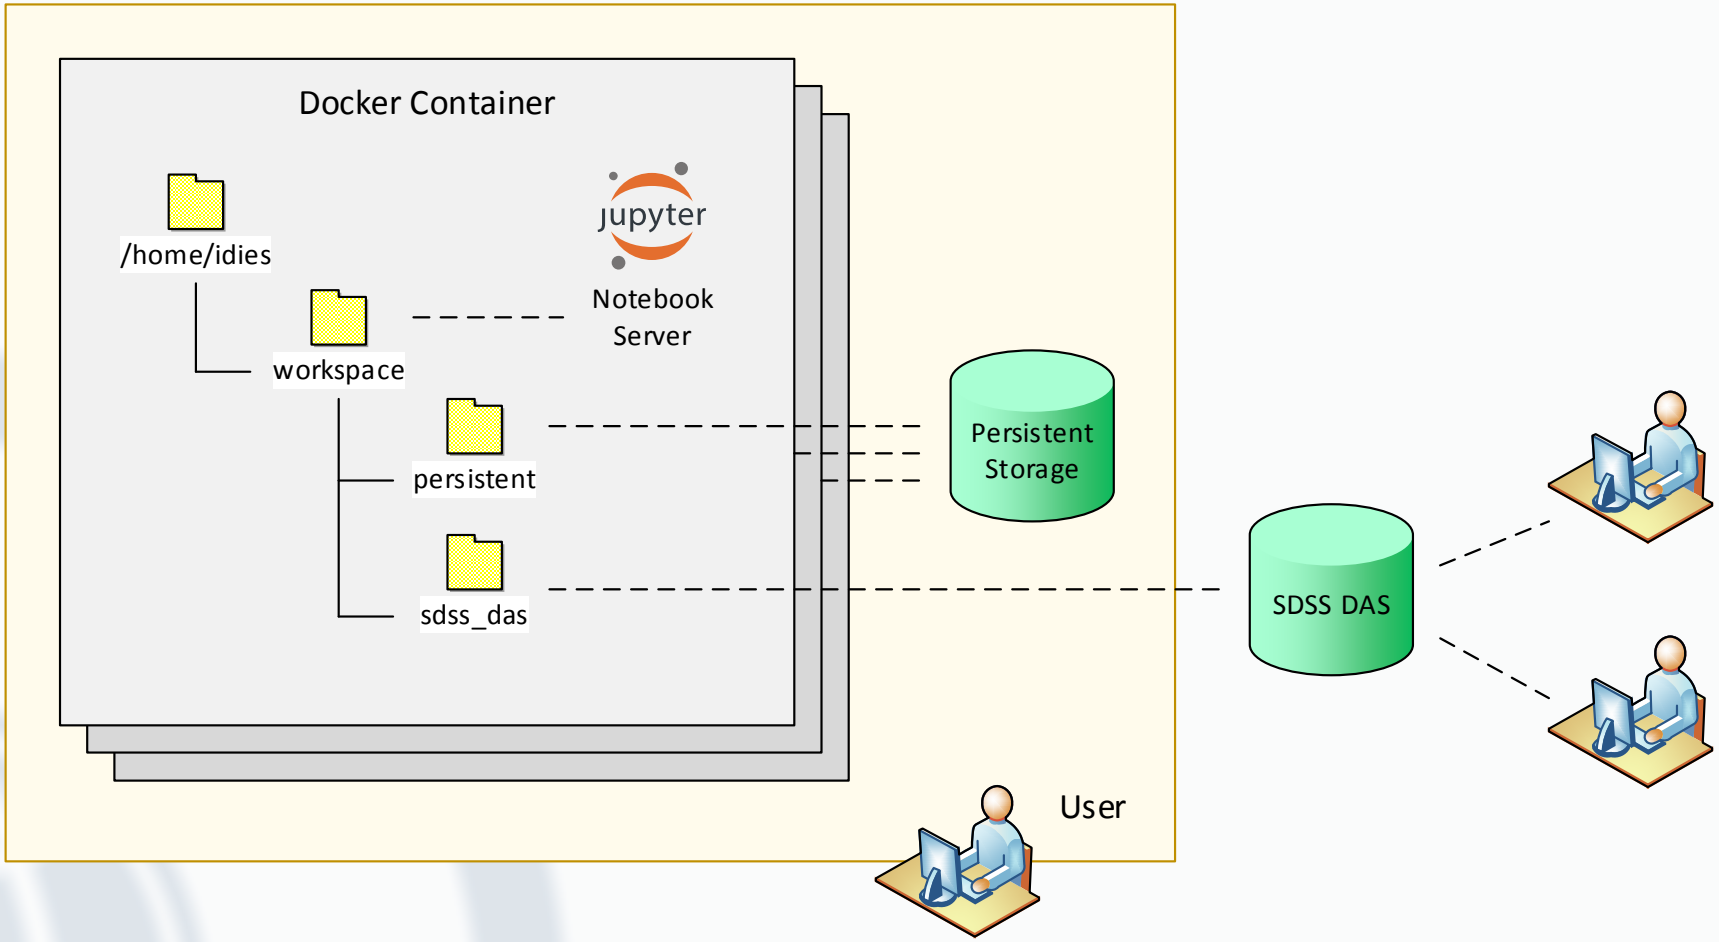

SkyServer Compute Status

- Run asynchronous non-interactive jobs in separate Docker containers. It's meant to be more than just Jupyter notebooks!
- Create new VM nodes on-demand to accommodate growing number of users.
- Provide scratch (temporary) storage space for working with large amounts of data.
- Improve resource management.

Questions?

Architectural Challenges

- Define tradeoffs
  - Data Analytics system is different from supercomputer
  - What is the right balance between I/O and compute?
- Need high bandwidth to big data
  - Computations/visualizations must be on top of the data
  - Must support at least few 100TB per server
  - Petascale: 3 copies for production (or erasure code?)
  - Wide area data movement/backbone is hard
- Lessons from the database world:
  - It is nontrivial to schedule complex I/O patterns
  - For subsets we must use indexing, cache resilient storage