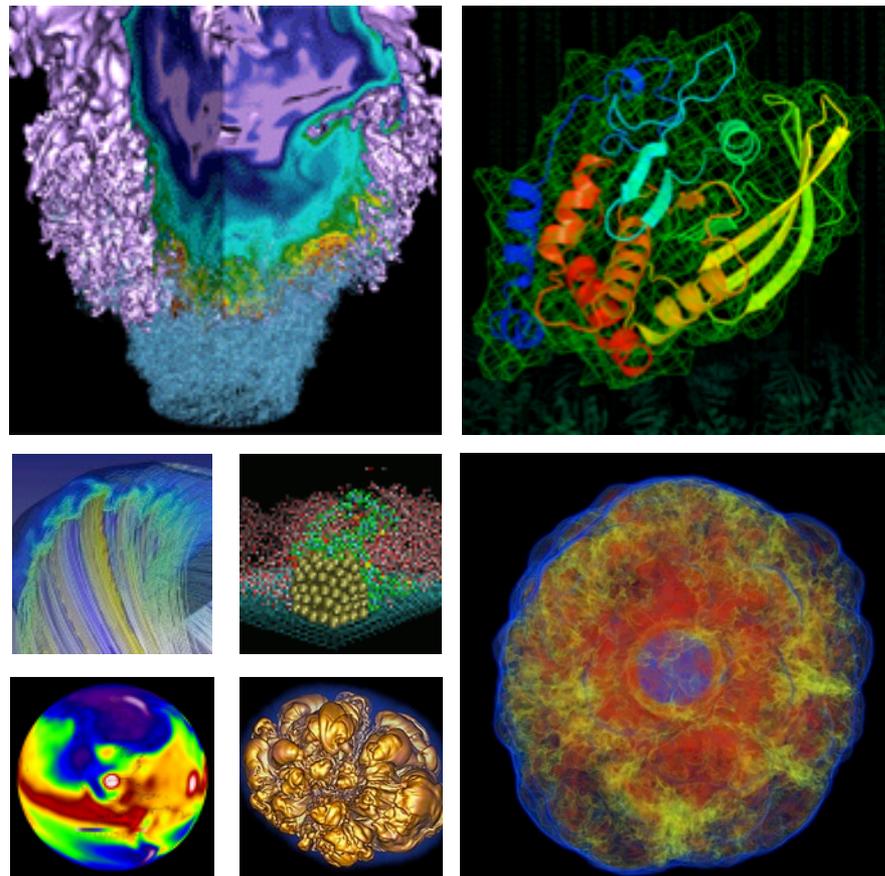


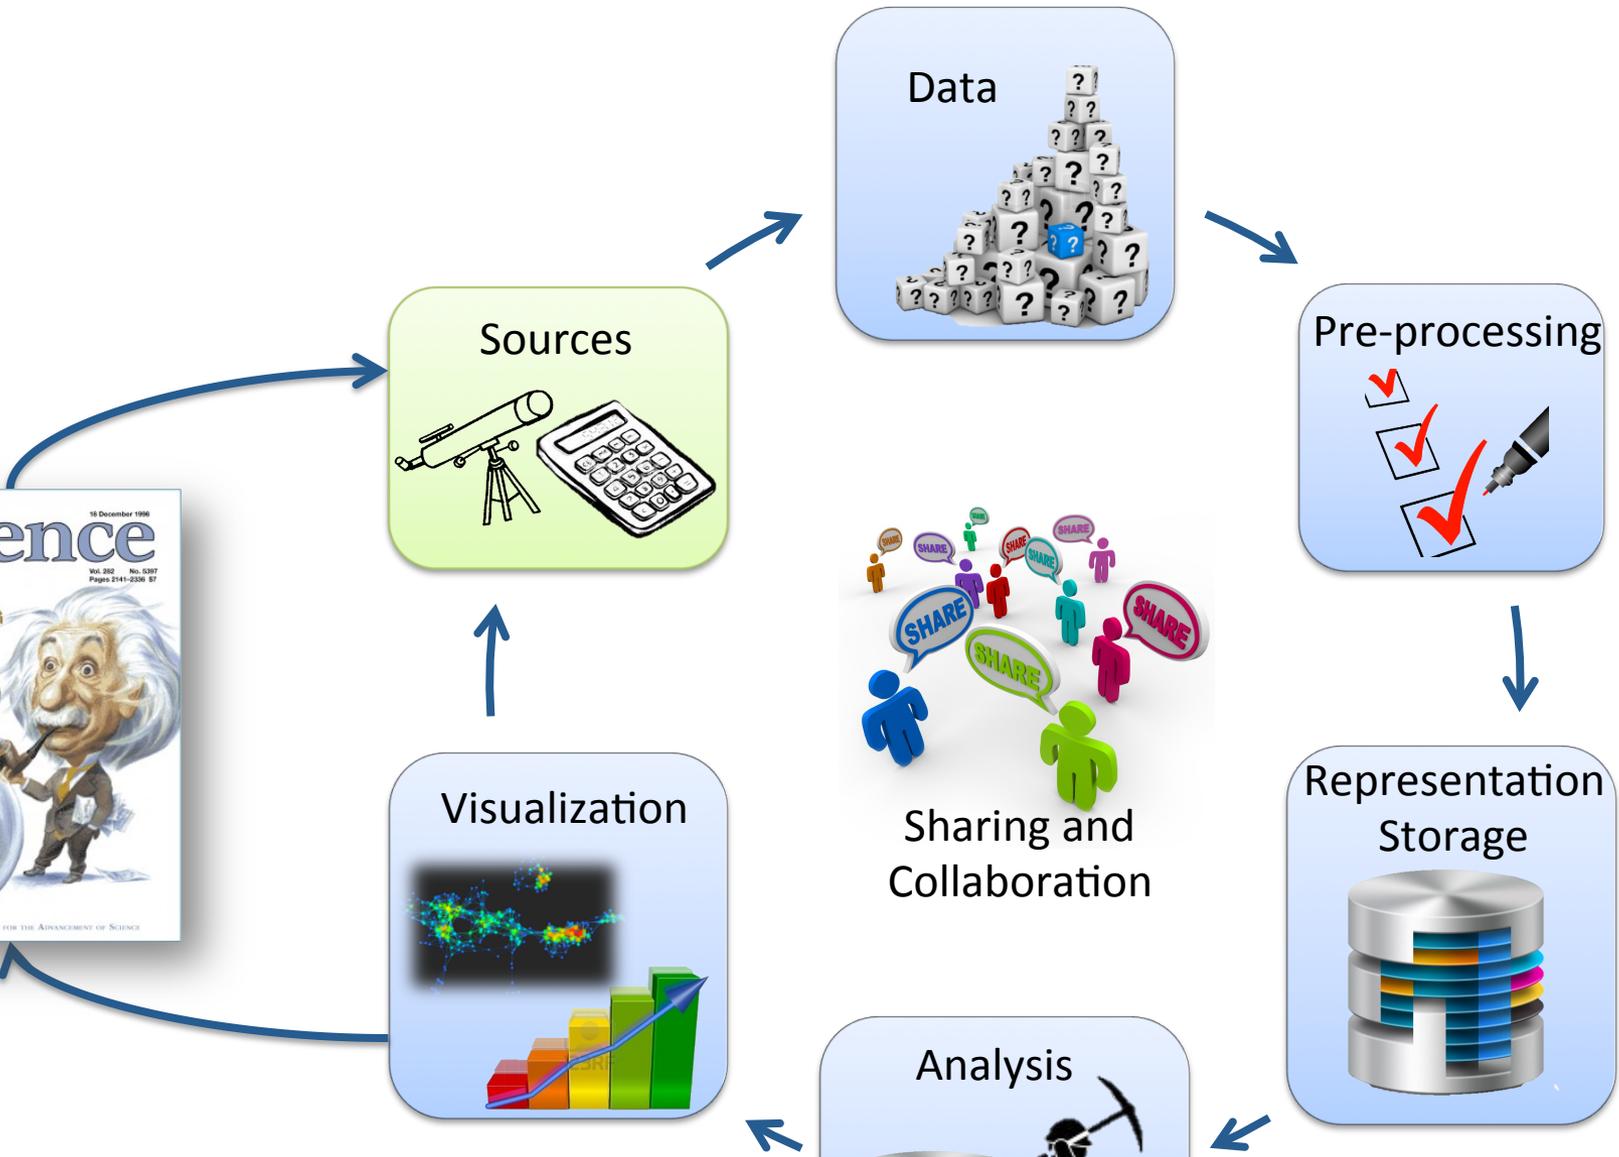
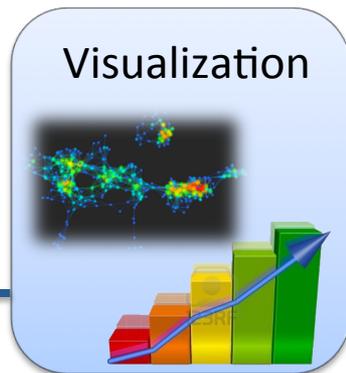
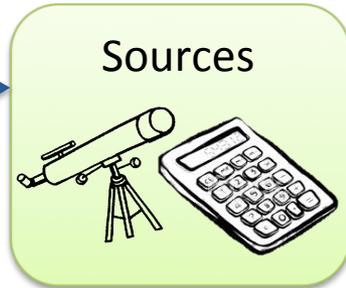
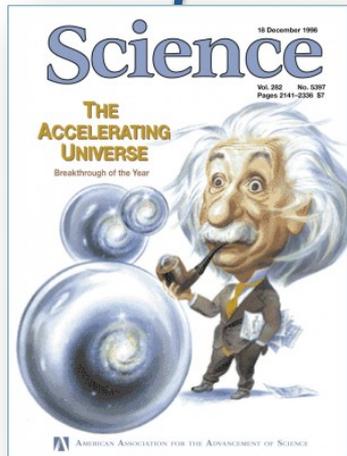
Databases and Data analytic frameworks at NERSC



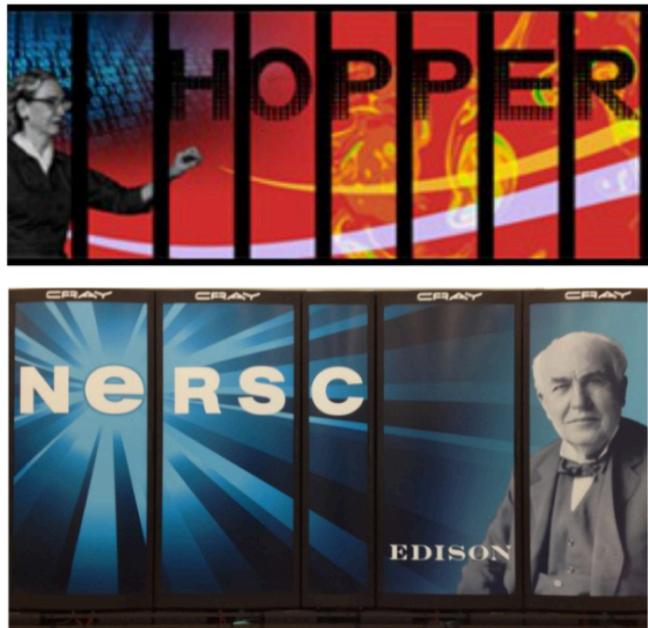
NERSC **40** YEARS
at the
FOREFRONT
1974-2014

Yushu Yao (yyao@lbl.gov)

June 16, 2014



NERSC Data Analytic Services



Big and Diverse Computing Facility
6000+ Users, 700+ Projects
3+ PetaFlops (20+pf more coming)
50+ PB Storage



Usability

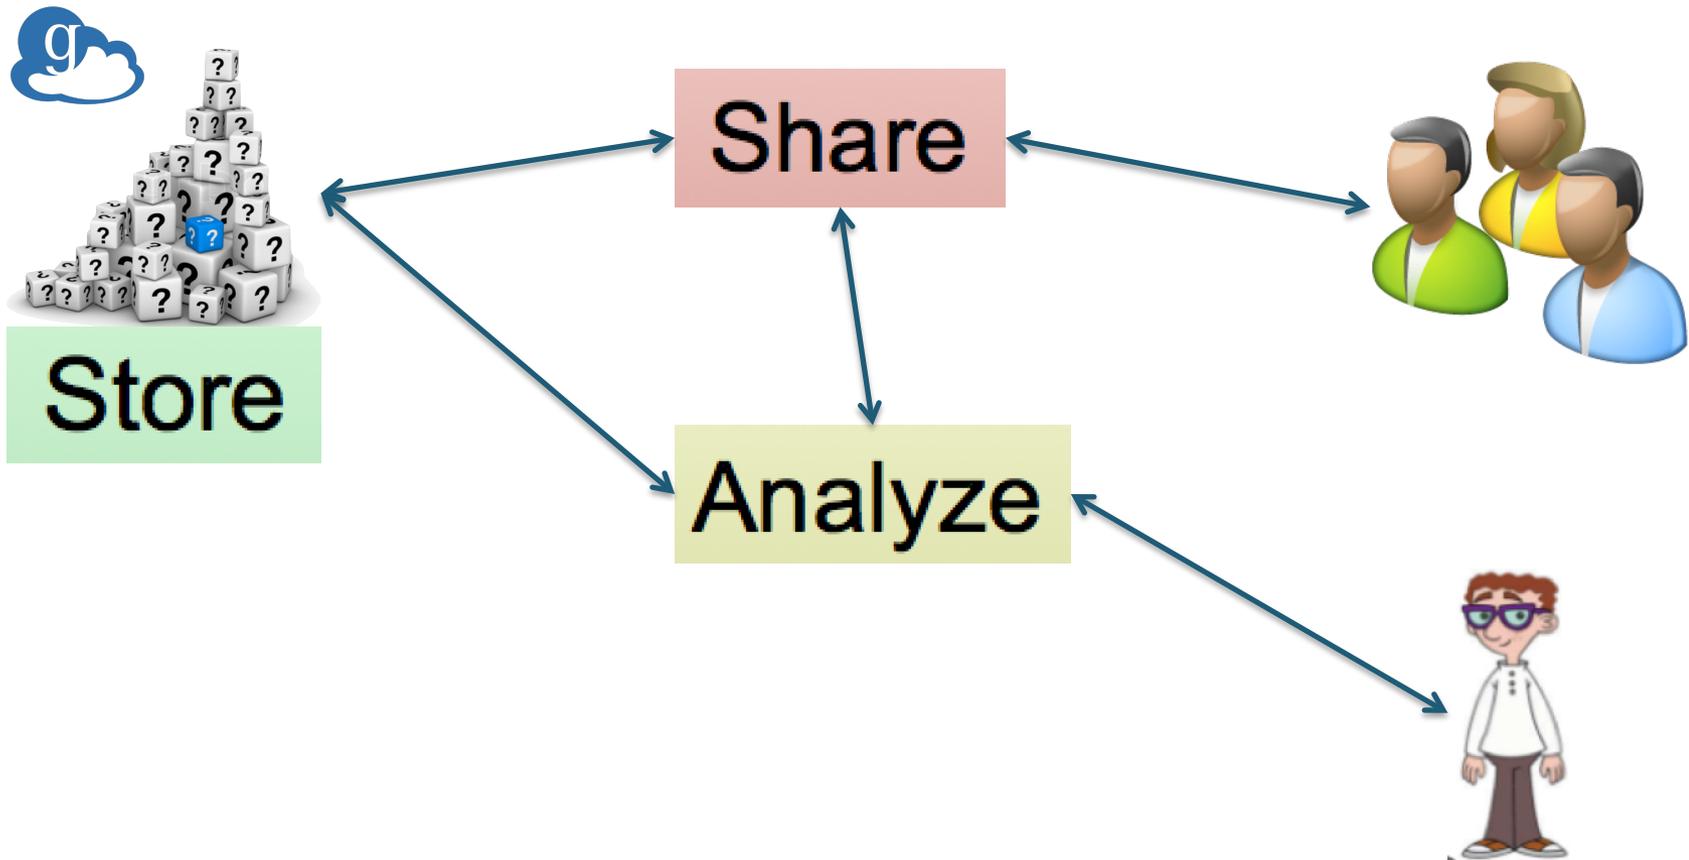
Scale

Management

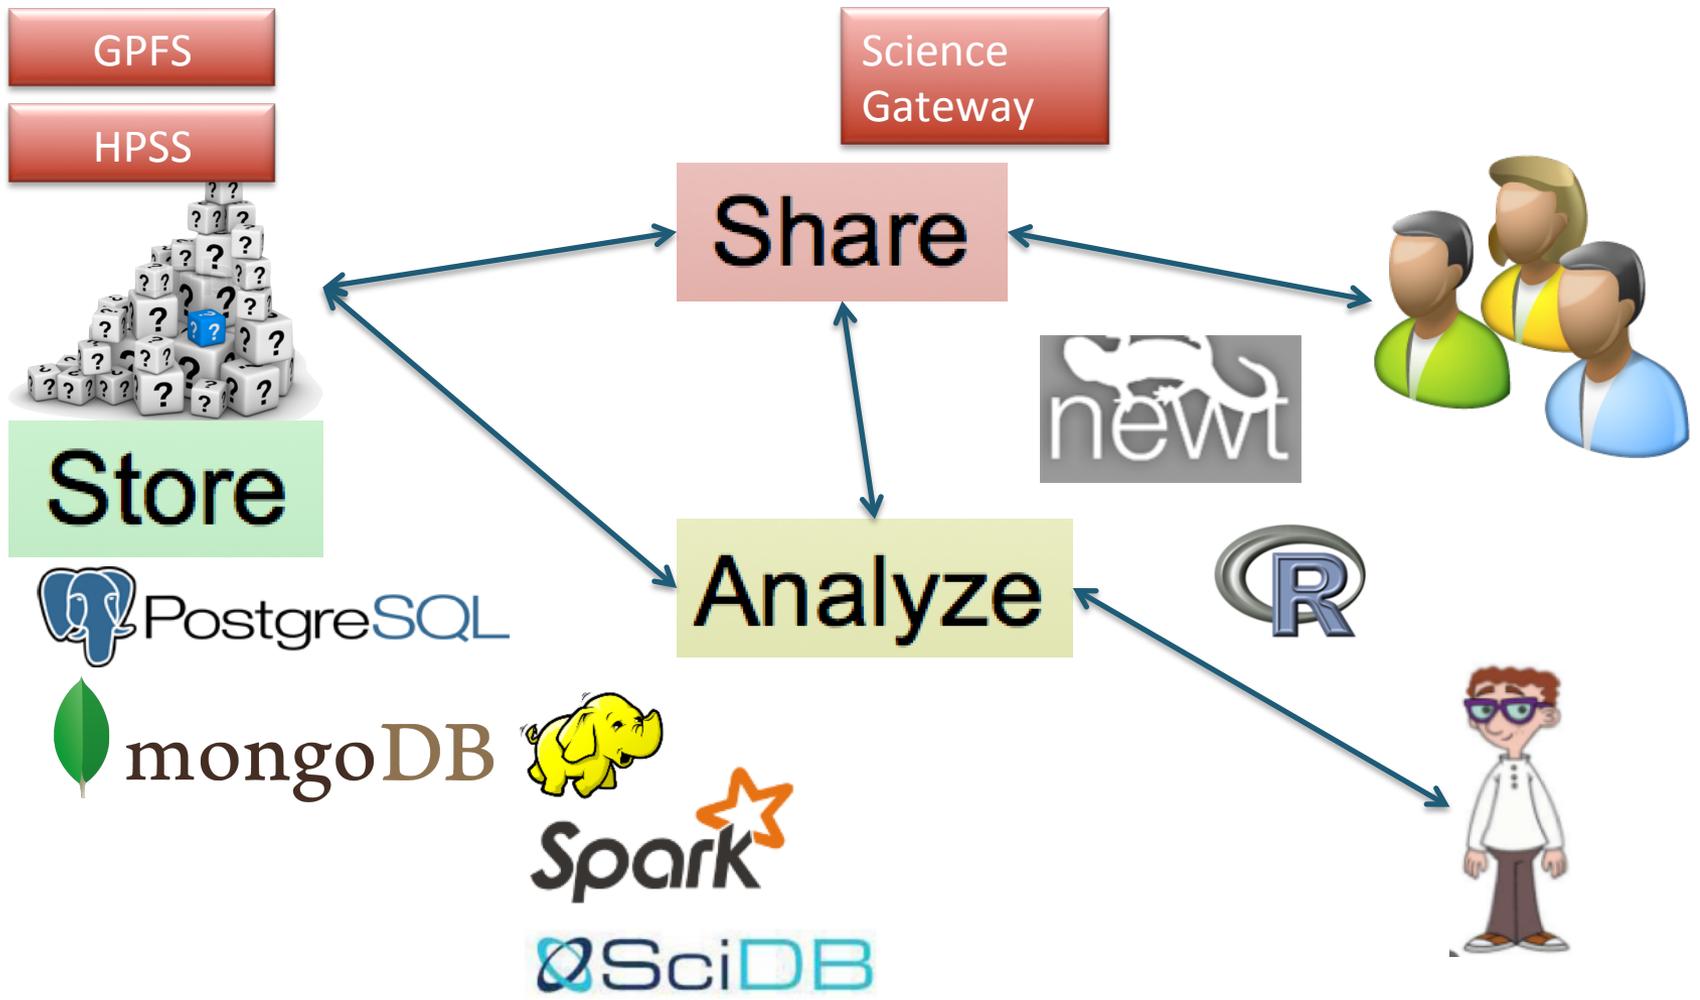
Analysis

Share

Store/Share/Analyze Data At NERSC



Store/Share/Analyze Data At NERSC



Science Gateway Services



- **Publish data on the web**
 - Create a www directory in your project space and put your data on the web
- **Build sophisticated web portals**
 - Build full stack web applications for your science at NERSC using Python/Django, PHP, Ruby on Rails
- **NEWT – the NERSC REST API**
 - Use the NEWT REST HTTP API to access NERSC HPC resources directly from your web apps.
 - Support for Authentication, Jobs, Commands, Files, Persistent Store, NIM, System Information at NERSC

Gateway examples - <http://portal.nersc.gov>

NEWT – <https://newt.nersc.gov>

Science Gateway information <http://www.nersc.gov/users/science-gateways/>

Traditional SQL Database Services



- **PostgreSQL and MySQL**
- **Good For:**
 - Structured, Relational Data
 - Mid-Size, \leq several GB in total
 - Transactional Operations

- **Key-value pair / Text database**
- **Good For:**
 - Un-Structured, Text Data
 - Mid-Size to Large, e.g. 10 GB of Text
 - E.g: for metadata that has ever changing schema

To request a MongoDB/PostgreSql/MySQL database:

<https://www.nersc.gov/users/science-gateways/science-gateway-databases/>

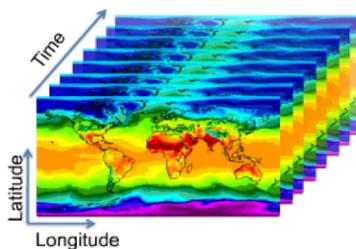
**How about
Large Structured Data?**

Array Like Science Data

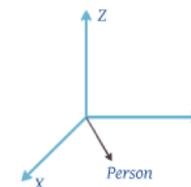
– More common than you think

Climate Simulation Output

-Terabytes of Output Per Run



Brain MRI Image

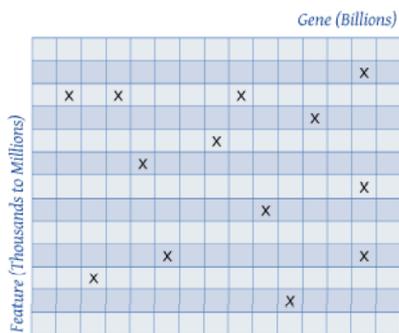


Many of people -> 4th Dimension

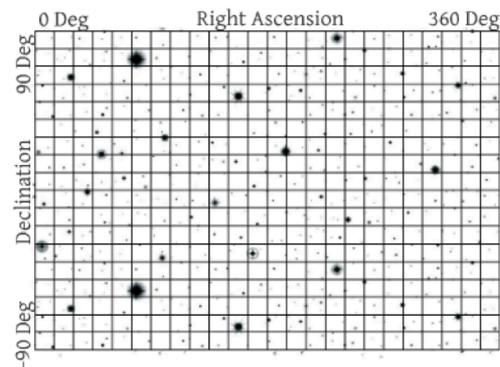


Gene Labeling

- Large Sparse Array



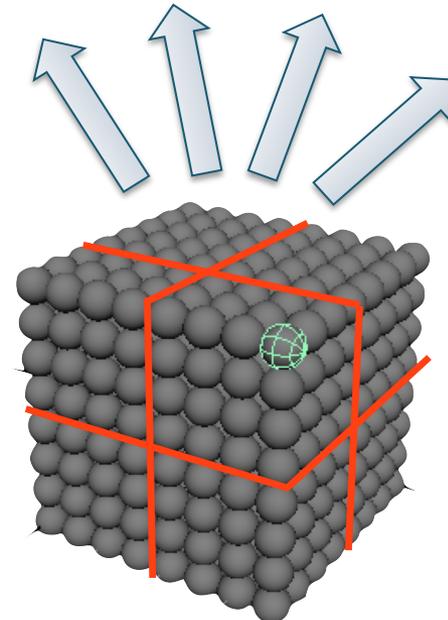
Catalog of Billions of Stars



SciDB, parallel processing without parallel programming

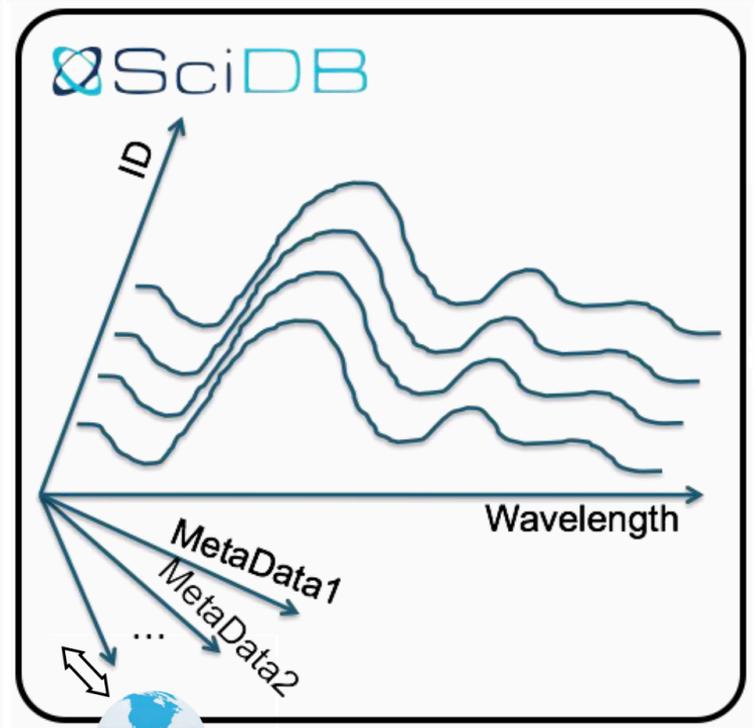
Everything in Arrays

- Locate an element at $O(\text{constant})$
- **Can be very sparse**
- Best for machine/simulation generated structure data
- Good for metadata too
- **Query-like language, auto-parallelization**
- **Do Calculations inside the DB**
 - **Data Analytic Framework**



- **Partner up with Science Teams**
 - Hold their hands to load the 1st batch of data and implement the 1st major analysis operation
- **15+ Science Projects**
- **Complicated Workflows and Algorithms**
- **Multiple Science Domains:**
 - Astronomy, Climate, Bio-imaging, Genomic
- **Multiple Types of Data**
 - Spectrums, Images, Time Series
- **Large amount of data**
 - Normally 100GB-1TB, some has 5+TB

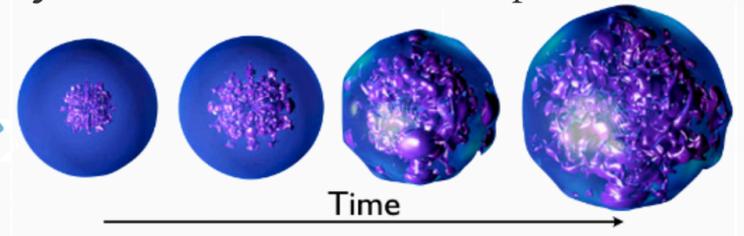
Match Supernova Observation to Simulation



Which spectrum looks like this?

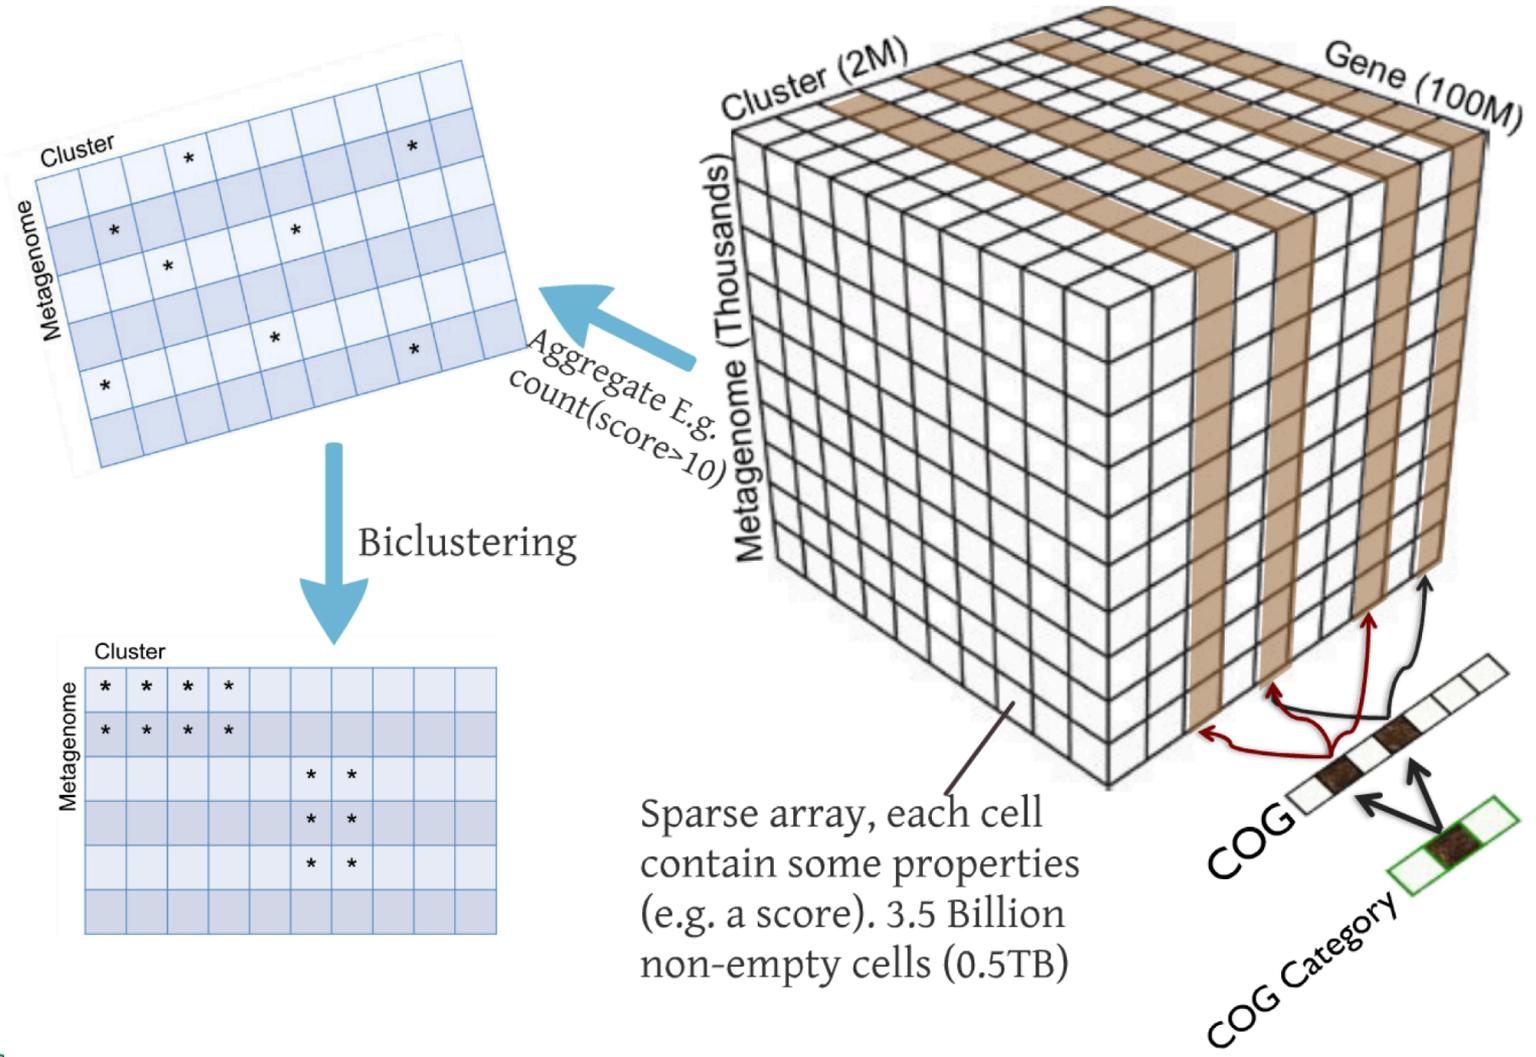


- Simulation of Supernova Explosion produces (many thousands to millions) spectra



- Web user can search/plot like any other gateway
- NEW: Web user can upload a spectrum, and Odetta will find the most "similar" spectrum in the database:
 - Very Compute/IO intensive
- For 55K spectra, SciDB returns result in 10sec. (comparing to 20min in PosgreSQL +Python)
- For 1Million Spectra, SciDB return result in 2min

Metagenome Workflow

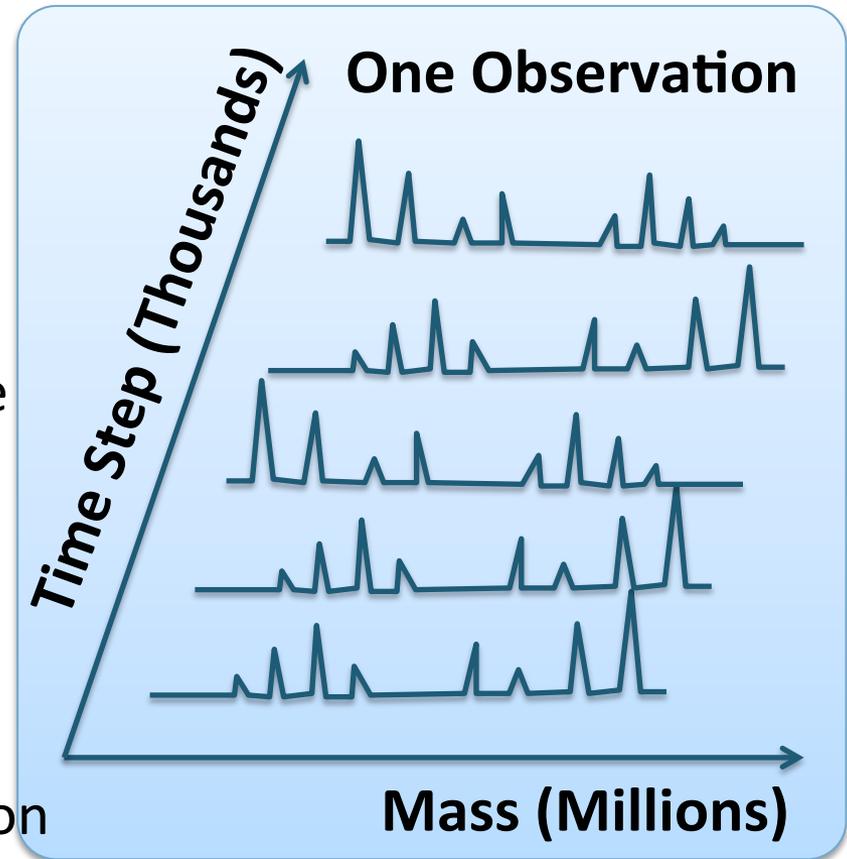


Cluster			*				*
Metagenome	*			*			
		*					
			*				
	*						*

Cluster	*	*	*	*					
Metagenome	*	*	*	*					
					*	*			
					*	*			
					*	*			

LCMS Workflow

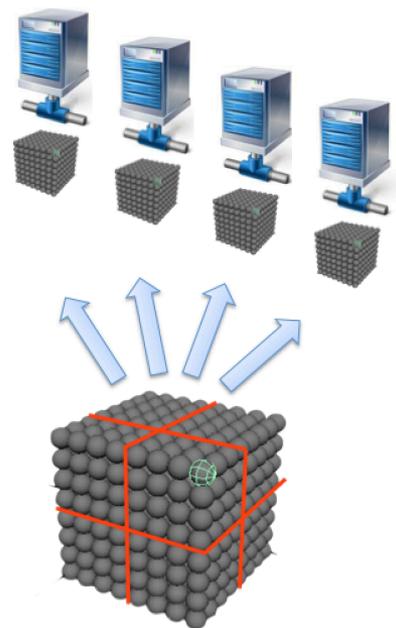
- **Liquid chromatography–mass spectrometry**
- **Each Observation**
 - Thousands of time steps
 - One mass spectrum at each time step
- **100K+ observations**
- **Operations:**
 - Subset/Aggregate observations
 - Aggregate inside each observation
- **100X-1000X faster than MySQL for some operations**



Each run is 50-200MB
100K runs and growing

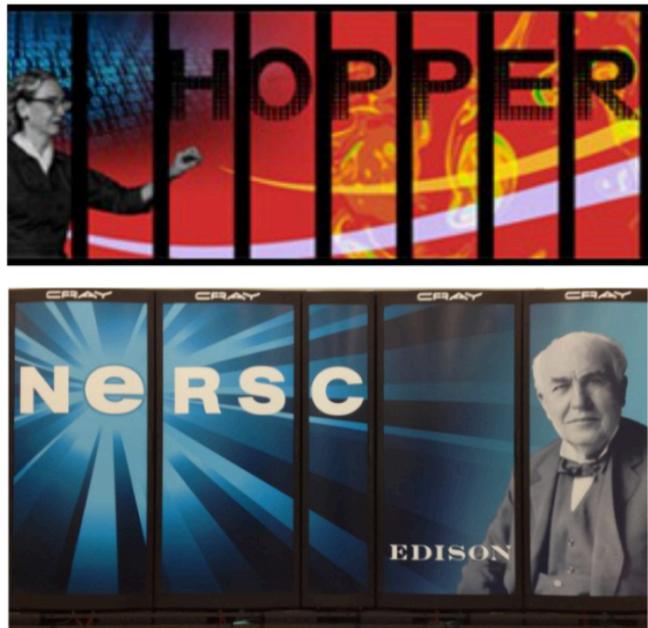
SciDB For High Usability Big Data Analytic

- **Why?** It's painful to manage and analyze terabytes of data. Need a unified solution that's easy to use.
- **What?** SciDB is a parallel database for array-structured data, great for **Terabytes** of:
 - Time series, spectrums, imaging, etc
- The greatest benefit of SciDB is:
 - **Usability:** Use HPC hardware without learning parallel programming and parallel I/O.



SciDB
Distribute a big array on many nodes

NERSC Data Analytic Services



Big and Diverse Computing Facility
 6000+ Users, 700+ Projects
 3+ PetaFlops (20+pf more coming)
 50+ PB Storage



Thank You

email yyao@lbl.gov