

Science Gateways @ NERSC

David Skinner

Lawrence Berkeley National Laboratory

Feb 12, 2013

Today at NUG

- What is a science gateway?
- NEWT a web API for HPC
- Examples running at NERSC now
- Futures / Questions

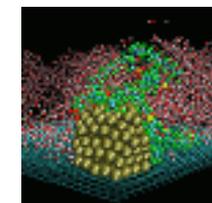
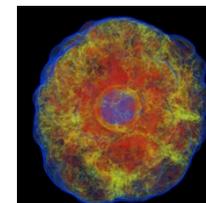
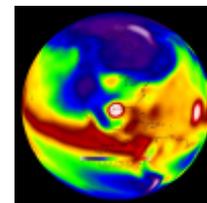
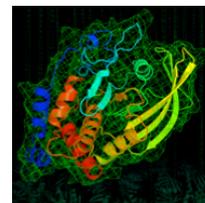
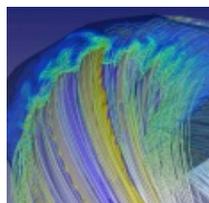
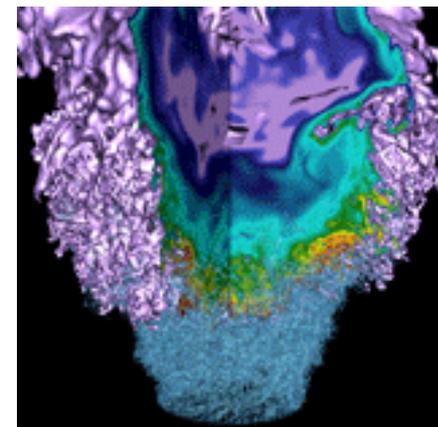
Things people do on NERSC computers and data systems

- Authenticate using NERSC credentials
- Check machine status
- Upload and download files
- Submit a compute job
- Monitor a job
- Get user account information
- Store app data (not scientific data)
- Issue UNIX commands

Q: Do this all in your browser?

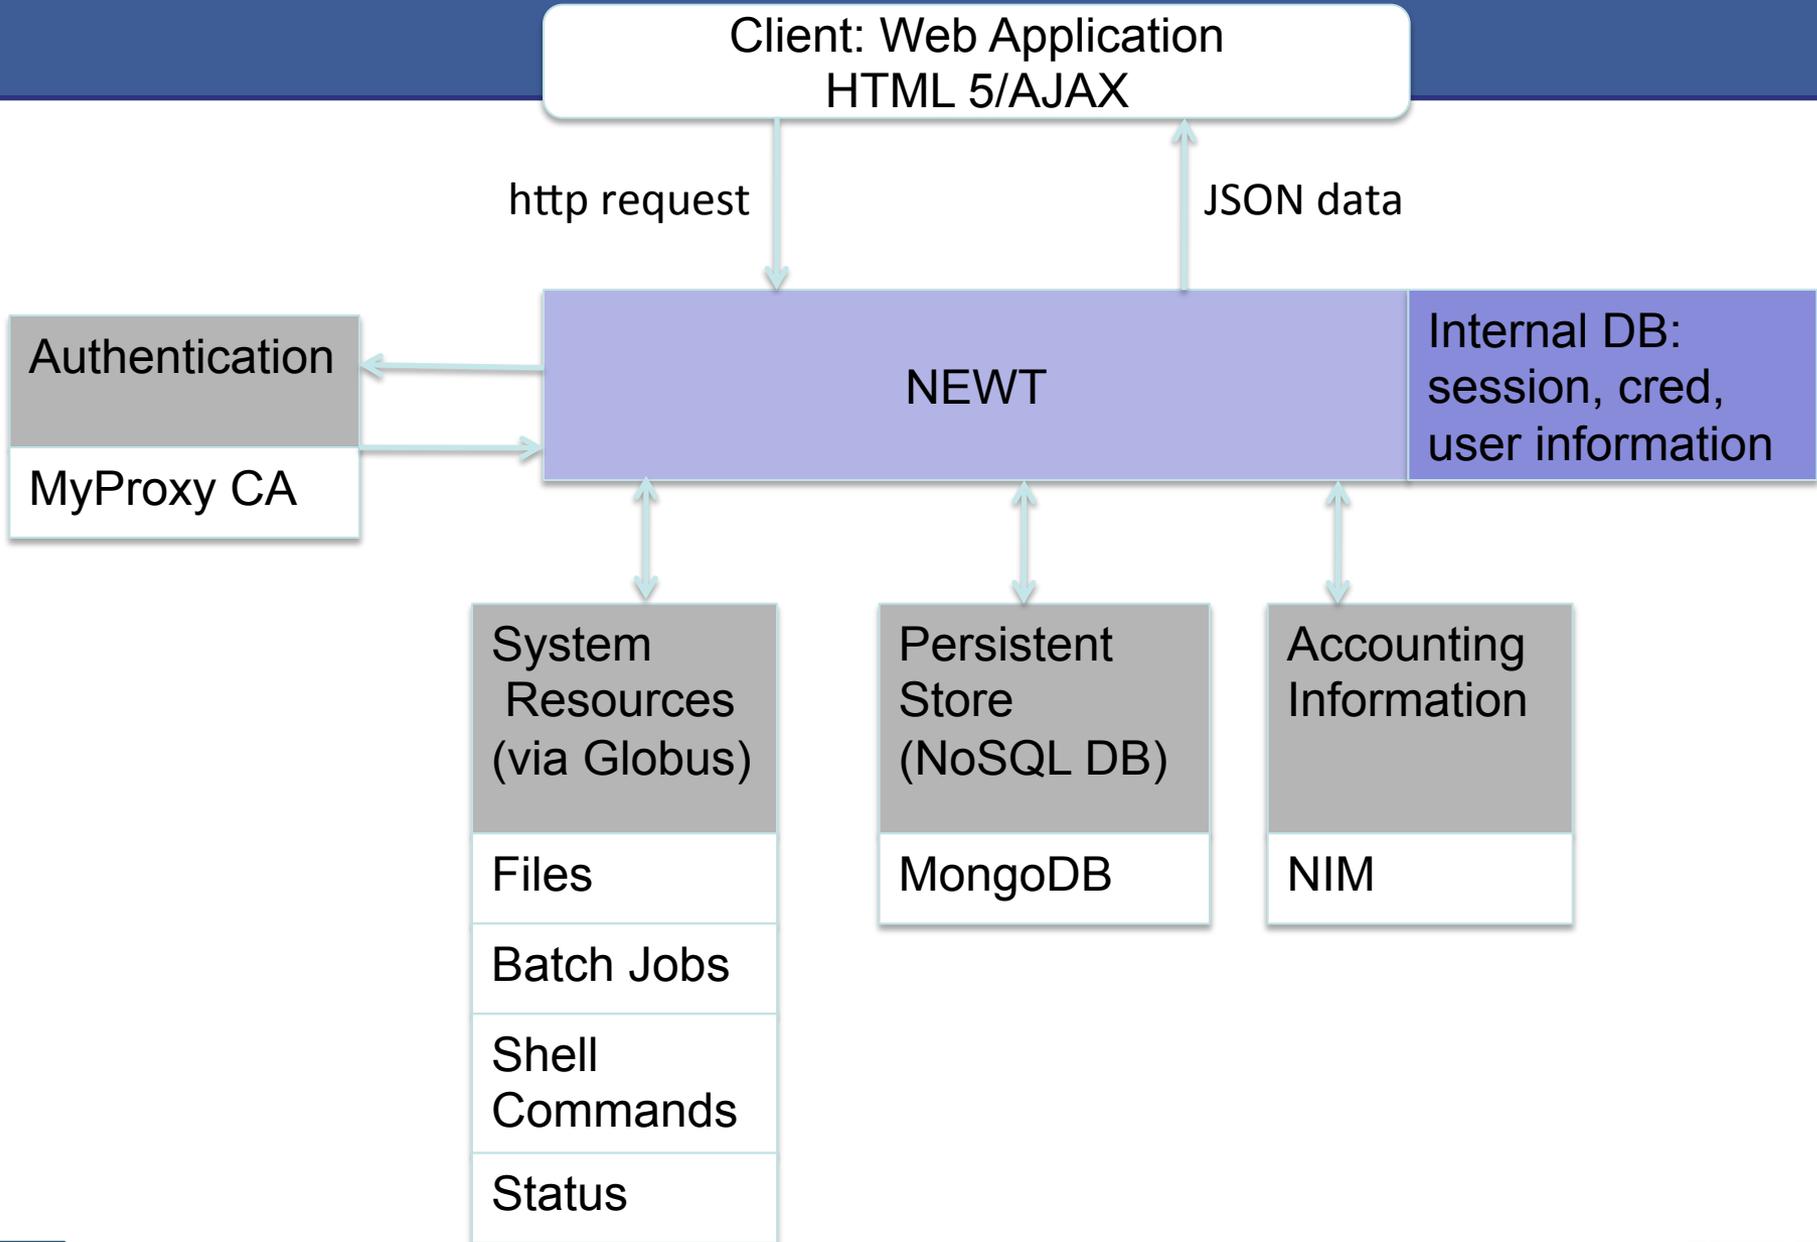
A: Yes with a science gateway!

NEWT API



HTTP: the language of HPC?

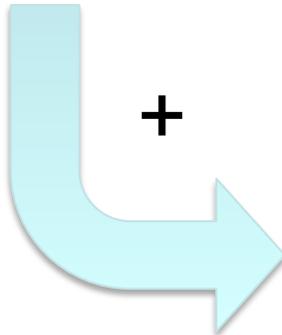
VERB	RESOURCE	DESCRIPTION
POST	/queue/R	Submits POST data to queue on R; returns job id
GET	/file/R/path/	Returns directory listing for /path/ on R
GET	/account/user/U	Returns user account info for U
DEL	/store/DB/DOC	Deletes object DOC in DB



The web augments the command line

Job id	Name	User	Time Use S Queue
997298.sdb	zigzagBwH+600	huangj3	0 H reg_small
997312.sdb	zigzagBwH+100	huangj3	0 H reg_small
997314.sdb	zigzagNwH+600	huangj3	0 H reg_small
997319.sdb	zigzagNwH+100	huangj3	0 H reg_small
1274346.sdb	QD09a	zhaohui	0 H reg_small
1300198.sdb	QD65	zhaohui	0 H reg_small
1340353.sdb	...neKL12h.R2.1	lanxin	0 H reg_small
1399578.sdb	T833N010ppm	helasrag	0 H reg_small
1399579.sdb	T833N010ppm	helasrag	0 H reg_small
1399582.sdb	T833N010ppm	helasrag	0 H reg_small
1399583.sdb	T833N010ppm	helasrag	0 H reg_small

dskinner@hopper02:~>



Command line when you need it, web when you don't

Use HTML+APIs to make science gateways (web apps)

NERSC Username NERSC Password

QSUB Demo

Choose one of the following:

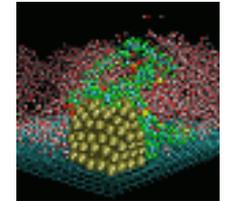
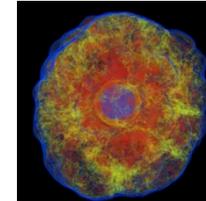
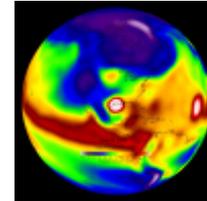
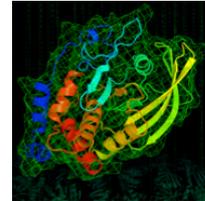
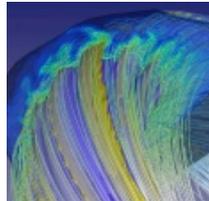
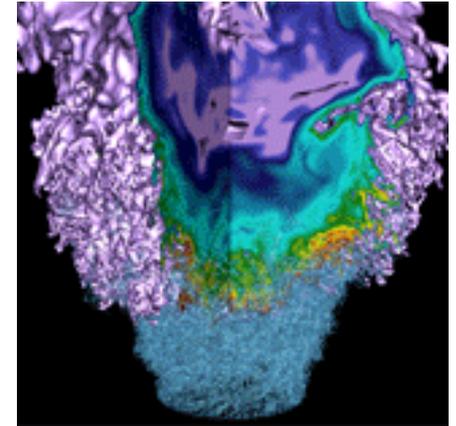
▼ Enter a job file path on franklin:

▶ Submit the batch script below to Franklin:

ID:

Status:

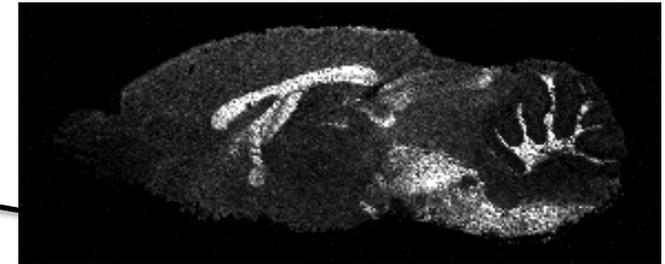
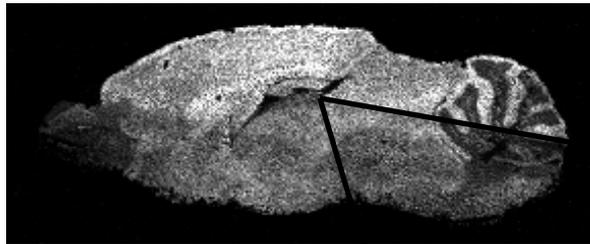
Science Gateways Tour



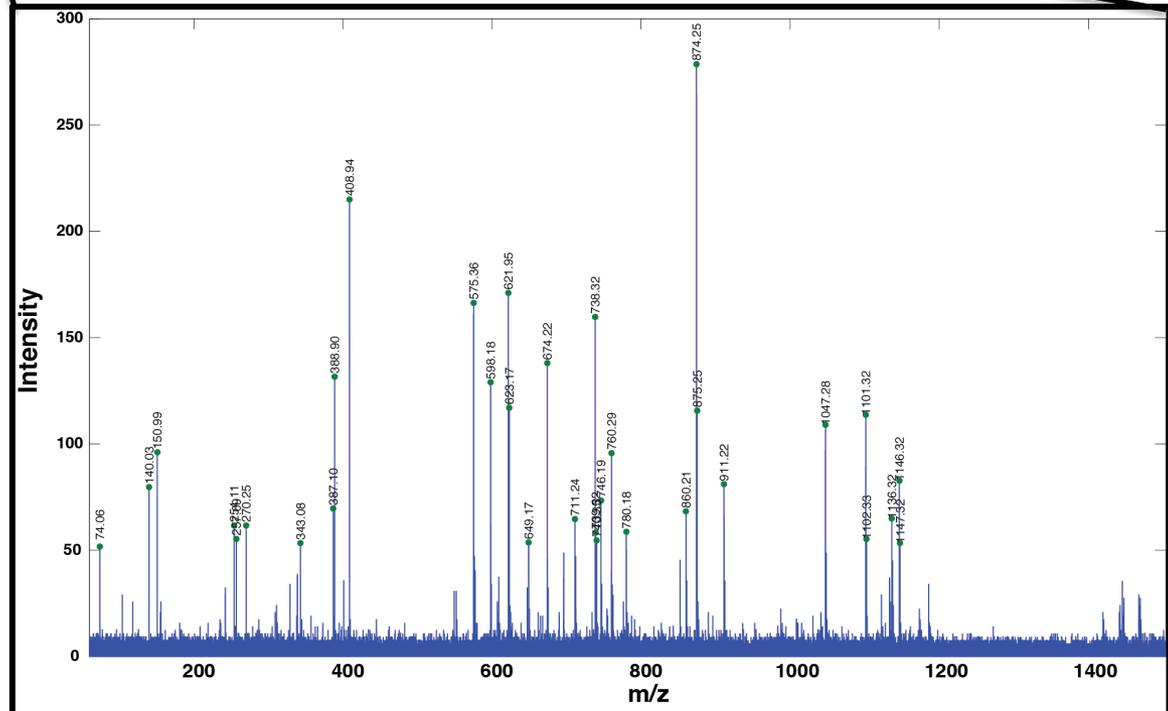
Mass Spectrometry Imaging (MSI)

Each ion is a dimension in the image

There are 1000s of dimensions in each image

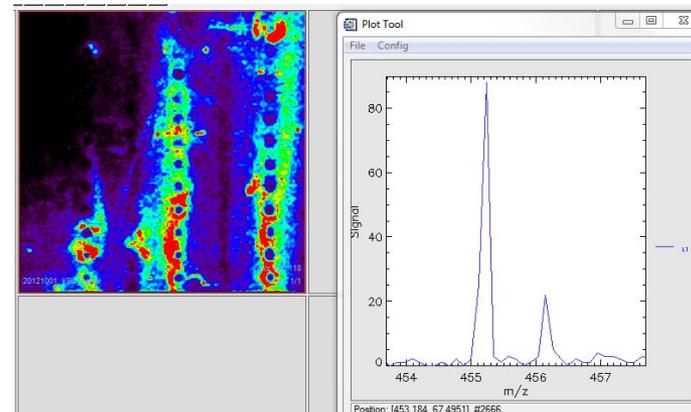


Each pixel
has a
complex
spectrum



Scale and complexity in MSI

- Low-performance, expensive software is a major impediment to driving science using MSI
 - Tissueview (\$10k), minimal features, obscures raw data
 - Fleximaging + ClinProtTools (\$20K), not cross-platform, doesn't overcome file-size limitations
- Typical files are ~10-30 Gbyte
 - If a user successfully opens a file, there is a large amount of information to sort
- People demand rapid access to processed “results” they can feel secure in reporting
 - At this point it requires a specialist to analyze MSI data
 - “Wants” include management of data, complex visualizations, and domain specific analyses



Community needs to move beyond MATLAB to scalable software and analysis methods

Project Lead



Project Team



- Fernando Fuentes
- Christina DeBianchi

Project Affiliates

- E. Wes Bethel
- David Skinner
- Peter Nugent
- Trent Northen
- Gary Karpen

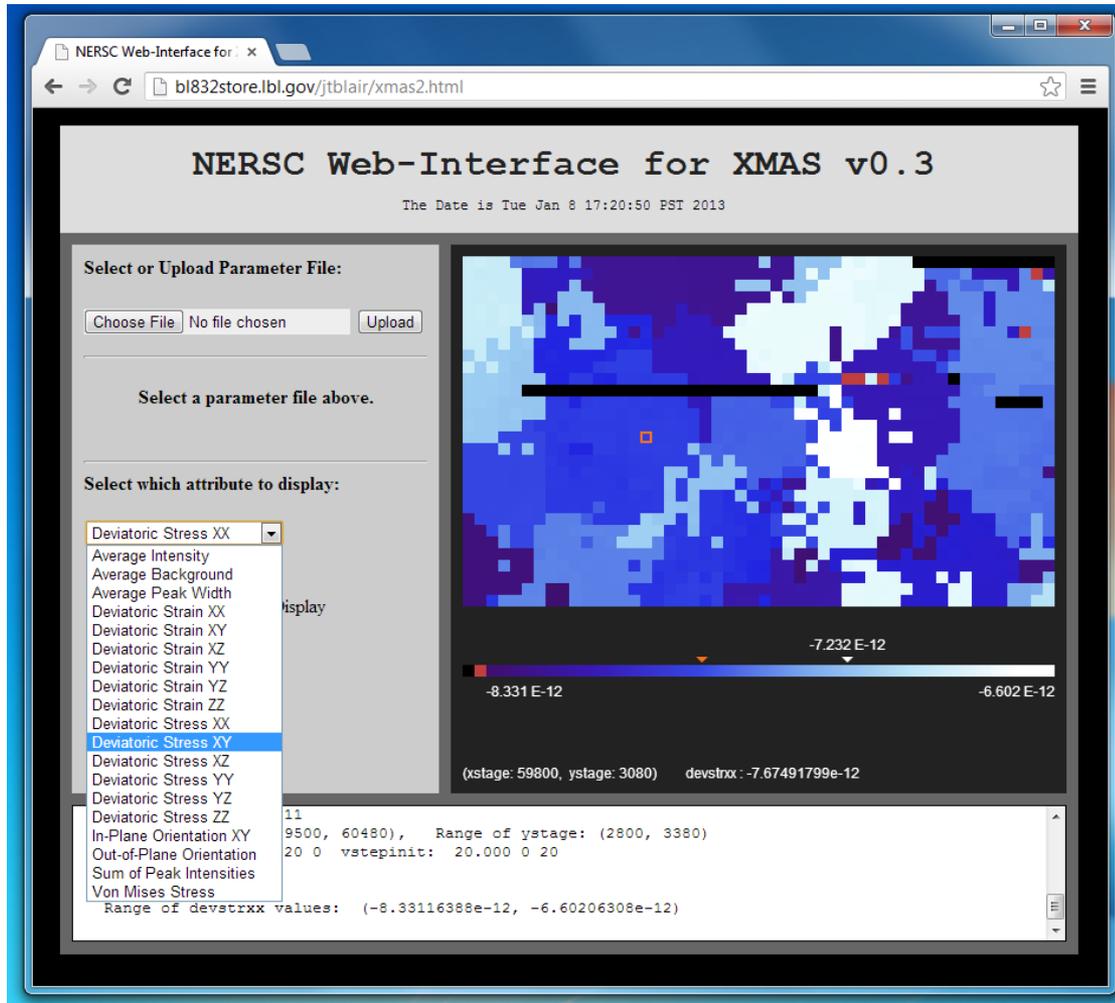
LBLN Imaging
Workshop
February 2.2012

NERSC ERCAP
Allocation Approved
April 4.2012

Introductory
Meeting
March 5.2012

First Complete
System Prototype
October 2012

Science Gateways: the future of X-RAY Data analysis



- Establish HPC conduit to NERSC capabilities for ALS users to analyze data in real time
- Develop new visualization tools for displaying multi dimensional datasets
- Introduce reverse modeling tools for understanding plastic deformation at the mesoscale and guiding experiments

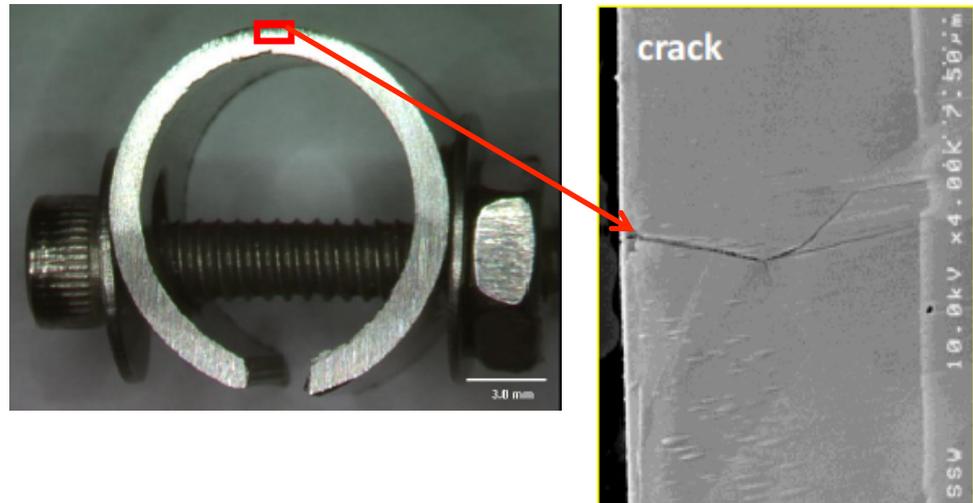
Study of Stress Corrosion Cracking Ni Alloys

Ni-16Cr-9Fe (Alloy 600) boiler tubing is widely used in PWR nuclear power stations.

Water causticity, local mechanical stresses, high temperature gradients → accelerated corrosion fatigue of boiler tubing

“C-ring” stressed to 2% plastic strain by bolt tightening

Stress corrosion cracking (SCC) is the single most important reason that nuclear steam generators are periodically replaced.



ALS Beamline 12.3.2 studies microdi of these and other samples. Each sca generate ~200GB of data which need analyzed.

Faster analysis == Faster understand

Current system deployed on VM

- Virtual Machine with 2 CPUs and 4GB of RAM
 - Easily scalable (within limits)

	Prototype	Alpha Release	Multiple Users	Full scale
#Users	1	10-20 at LBNL	30-60 at 3 Labs	~1000 at 100 Labs
#Data Contributors	1	3	10	~300
Number of Files	6	100	300-500	10,000

	Prototype	Alpha Release	Multiple Users	Full scale
Storage	<1 TB	10 TB	50 TB	1 PB
Virtual Machine	2 CPUs, 4GB	8 CPUs, 64GB, Parallel file system (+flash)
Analysis	1 node 1-3 GPUs	100's nodes 10's GPUs	...	1000's nodes 100's GPUs

OpenMSI Audience

Computing Requirement

- Public facing pages on OpenMSI would potentially have 1000s of users per day

Example

- A data repository for X-ray science, point of reference for data used in publications
- A watering hole for discussions on data standards
- Server side data analysis tools for images of unwieldy size
- Spinning disk and tape storage available
- <http://cxidb.org/>

CXIDB ID 10

Citation Details	
Title:	Femtosecond free-electron laser x-ray diffraction data sets for algorithm development
Authors:	S. Kassemeyer et al.
Journal:	Optics Express
Year:	2012
DOI:	doi:10.1364/OE.20.004149
Experimental Conditions	
Method:	Single Particle X-ray Diffraction Imaging
Sample:	Nanorice
Wavelength:	1.0 nm
Lightsource:	LCLS
Beamline:	AMO
Deposition Summary	
Depositor:	Ilme Schlichting
Contact:	ilme...@mpimf-heidelberg.mpg.de
Deposition date:	2011-08-21
Last modified:	2011-08-21
Data Files	
Diffraction Patterns:	cxidb-10-amo15010-r0067.tar (1.4 GB)

4 more Examples

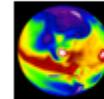
<p>Joint Genome Institute</p> <p>Mostly public web gateways for retrieval and analysis of (meta)genomics data. Many distinct databases organized by domain (fungal, plant, prokaryote, etc.). Oracle.</p>	<p>DeepSkyProject.org</p> <p>Public image web gateway for transient detection, deep field, and coaddition. 13 M images. Private DB connections available to power users. Postgres.</p>
<p>Earth Systems Grid (at NERSC)</p> <p>IPCC AR5 community data sets with fine grained access controls. Target for all datasets is public over the long term. Large dataset subselection. Custom.</p>	<p>MaterialsProject.org</p> <p>Search for materials given a set of properties. We do DFT so you don't have to. Public but throttled access. Many private sector users. Data ingress is by arrangement. MongoDB.</p>

NERSC Science Gateways

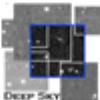
Current List of Web Gateways to Science at NERSC:



The Materials Project



20th Century Reanalysis



DeepSky



Dayabay



QCD



Earth System Grid



CXIDB



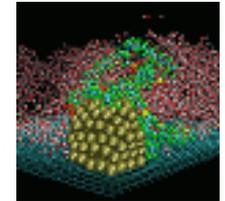
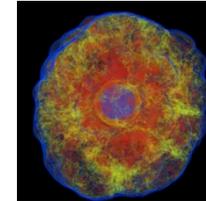
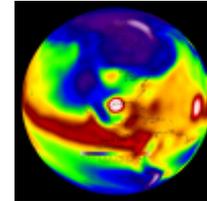
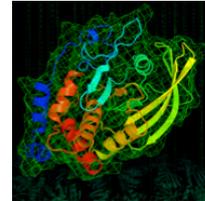
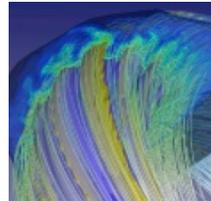
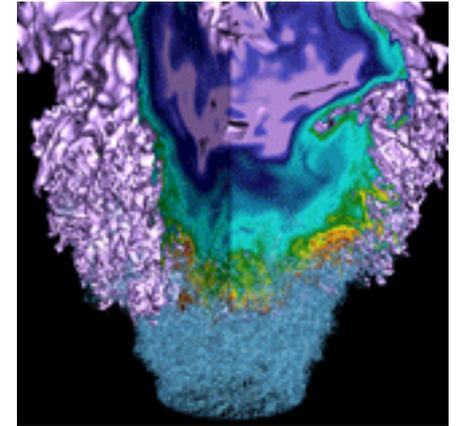
NEWT



NOVA

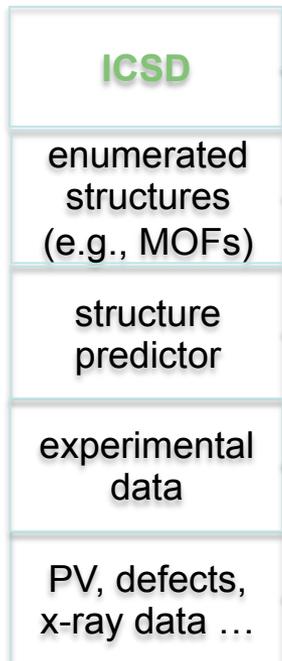
Your Gateway? Contact deskinner@lbl.gov

NEWT Futures



Gateways Broaden the Impact of HPC

computational survey drivers



simulation engines



Hands-off Expertise Automation (A)
 HighThroughput Computing (A)
 Big Data Materials Ontologies (B)
 Probablistic Data Management (B)
 Anomally Detection (C)
 Consistency / V&V (C)
 Query Language / Web Gateway (D)
 Experimental validation is key(!)

(A)

(B)

operational
In 2012

(D)

functional
electronic
materials

phosphors
&
scintillator

carbon
storage

critical
materials...

materials genome

Community
Web
Gateway

(!)

(C)

NEWT 2.0 Fall 2013

NEWT 1.0 :

- NERSC Web Toolkit REST API that covers things people do at HPC centers
- Batch and interactive oriented execution
- HTTP based file/data movement

NEWT 2.0:

- Refactoring around flatter URLs
- Batch, interactive, and pipeline oriented execution
- High Throughput Computing (HTC)
- New approaches for wildcarding
- Data analysis provenance through tagging
- Extensibility: requests to extend/modify API through POSTS
- Take this to the W3C

NEWT 2.0 : Channels for HTC

- Command (runs in seconds @ login node)
- Job (runs hours @ batch)
- Task (runs minutes to hours @ channel)
- Channel {



id: name: odetta.pipelineA_swarp

version: 1.3.2

state : {active, stopped, disabled}

repo: (who to bill), max_run:, max_wall:, max_cores:, max_mem:

machines: (**resources**(s) where this task can run)

users: (**user(s)** who can access this **channel**)

script: {} (code that accepts the parameters and executes)

routing_tags: {} (add these KVP strings to the system routing tag)

}

HTML5 enables C2C apps

- WebSockets (`ws://newt.nersc.gov/my_app`)
 - Now widely accepted web standard
 - Broaden's web browsers beyond sipping data from http and https.
 - Think gridftp or VisIt in a browser
- WebRTC
 - It's here. In the next browser you will download
- How will scientists use this?
 - Faster, richer applications
 - Need to influence W3C standards for web science