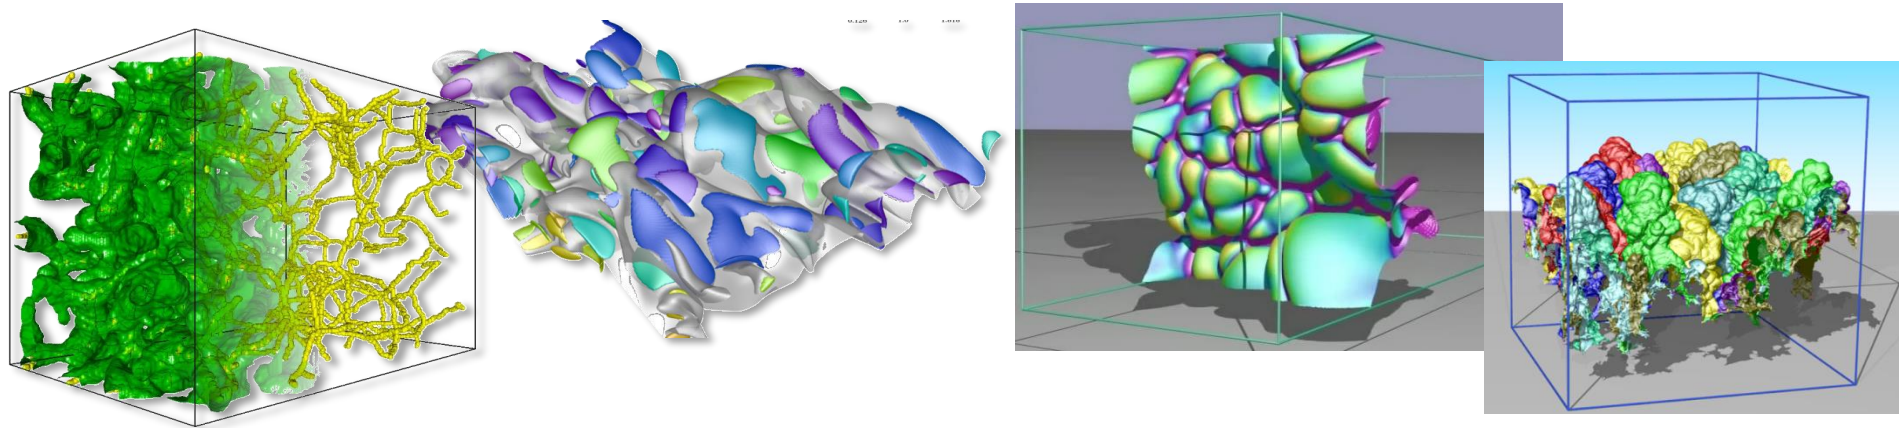
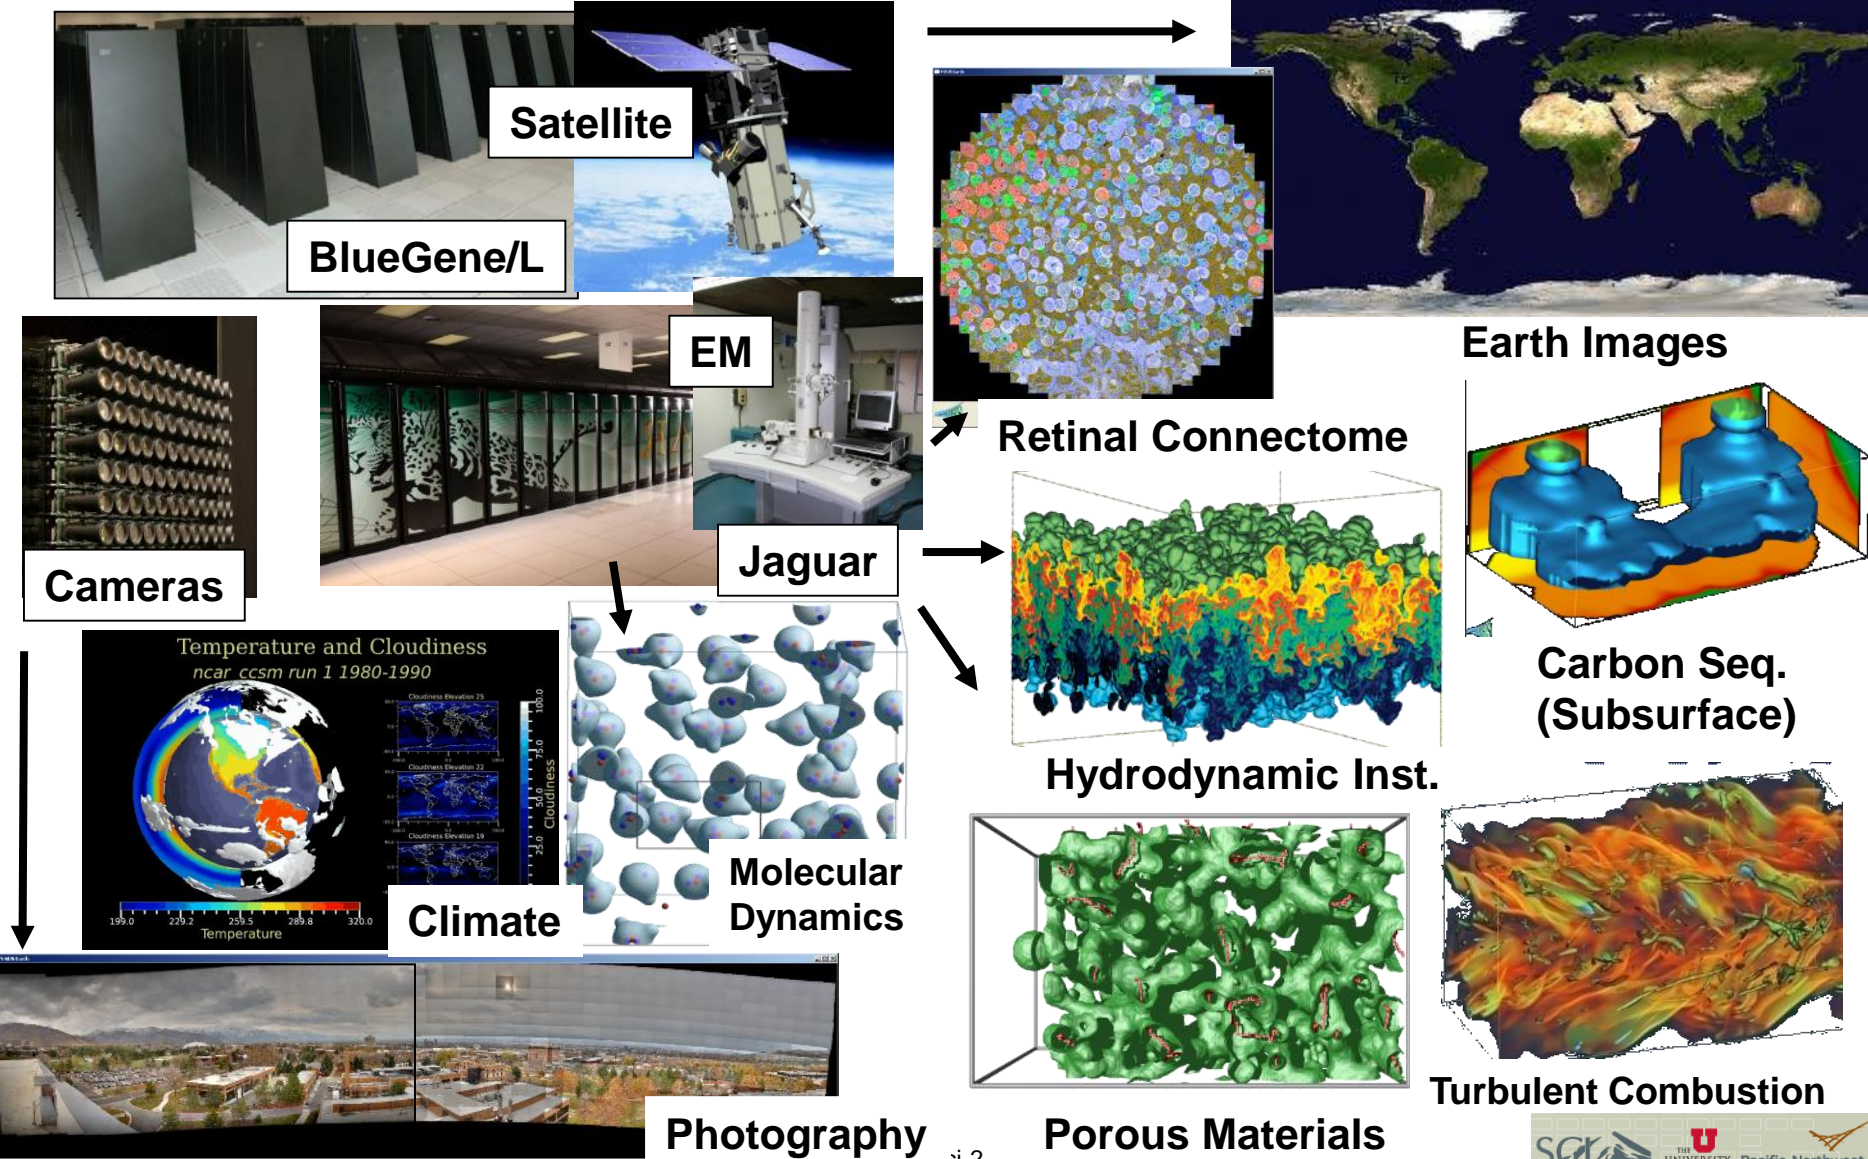


# The OpenVisus Framework for Extreme Data Management, Analysis and Visualization



Steve Petruzza, Aniketh Venkat, Nate Morrival, Giorgio Scorzelli, Valerio Pascucci, and Many Many Many Others

# Massive Simulations and High Resolution Sensing Generate Big Data Challenges and Opportunities



Satellite

BlueGene/L

EM

Jaguar

Cameras

Earth Images

Retinal Connectome

Carbon Seq. (Subsurface)

Hydrodynamic Inst.

Turbulent Combustion

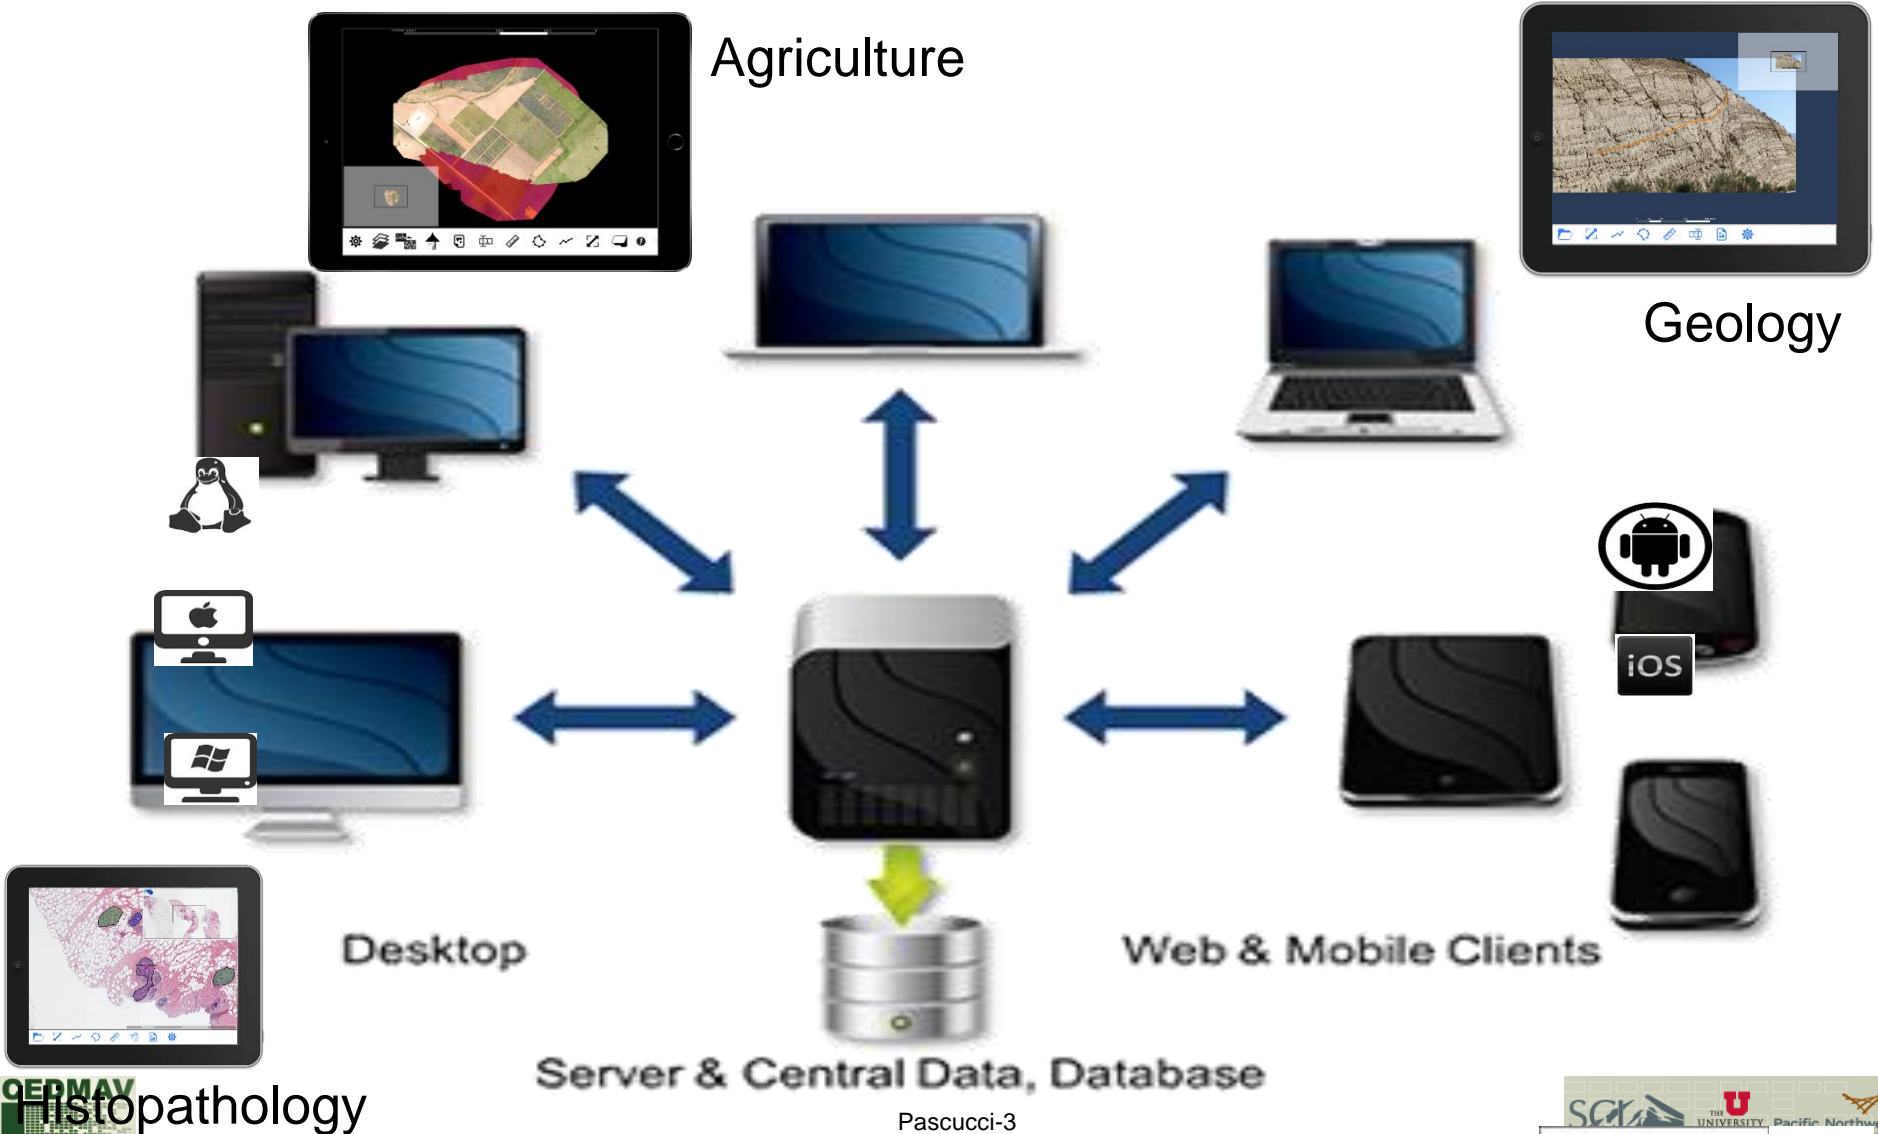
Climate

Molecular Dynamics

Photography

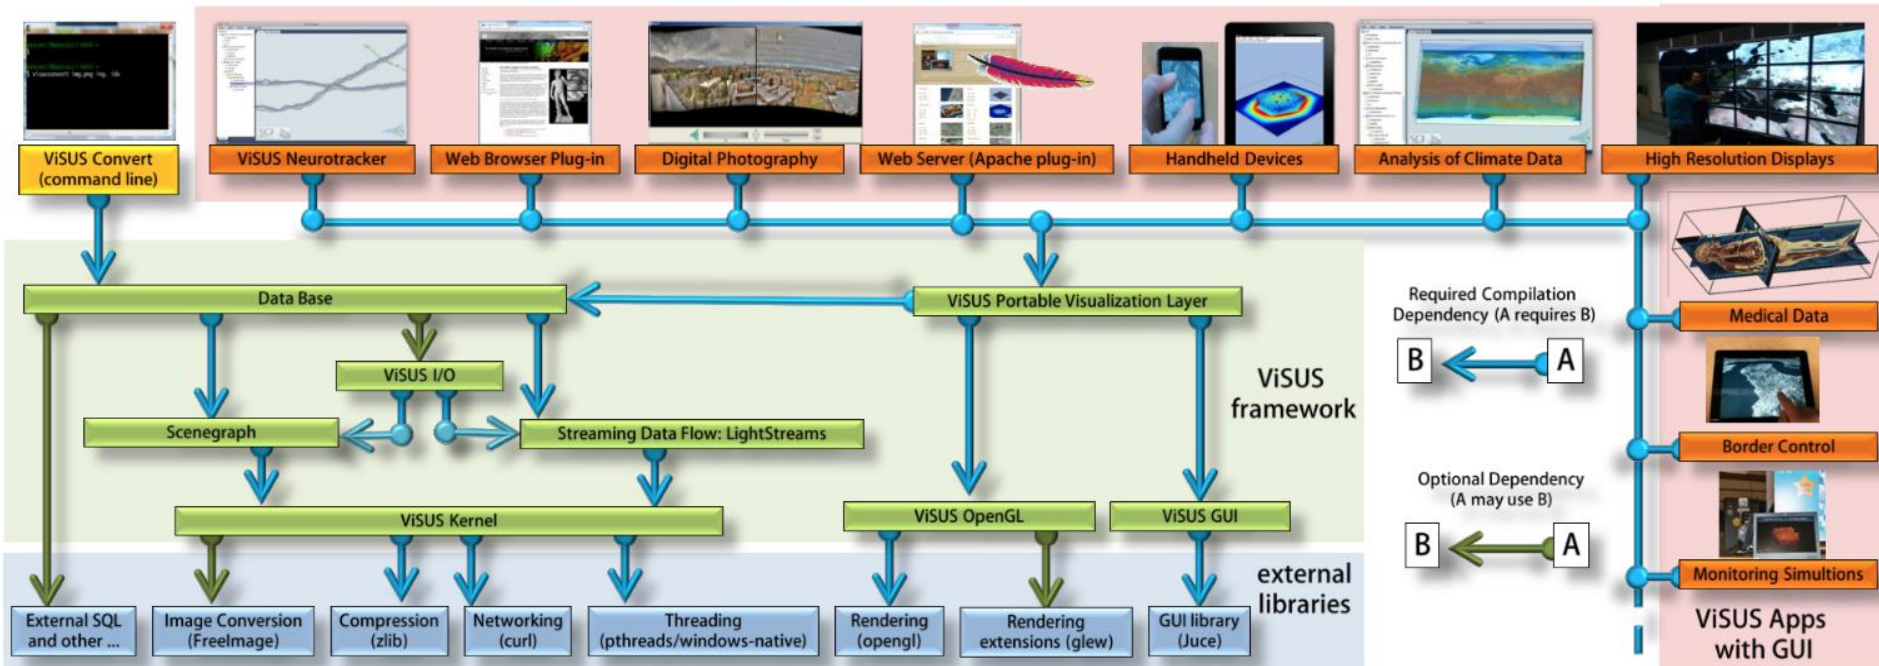
Porous Materials

# OpenViSUS: A Data Intensive Collaborative Analytics and Visualization Platform



# OpenViSUS Software Platform

A set of interoperable components easily customized and deployed for different solutions as needed




# PIDX demo

---

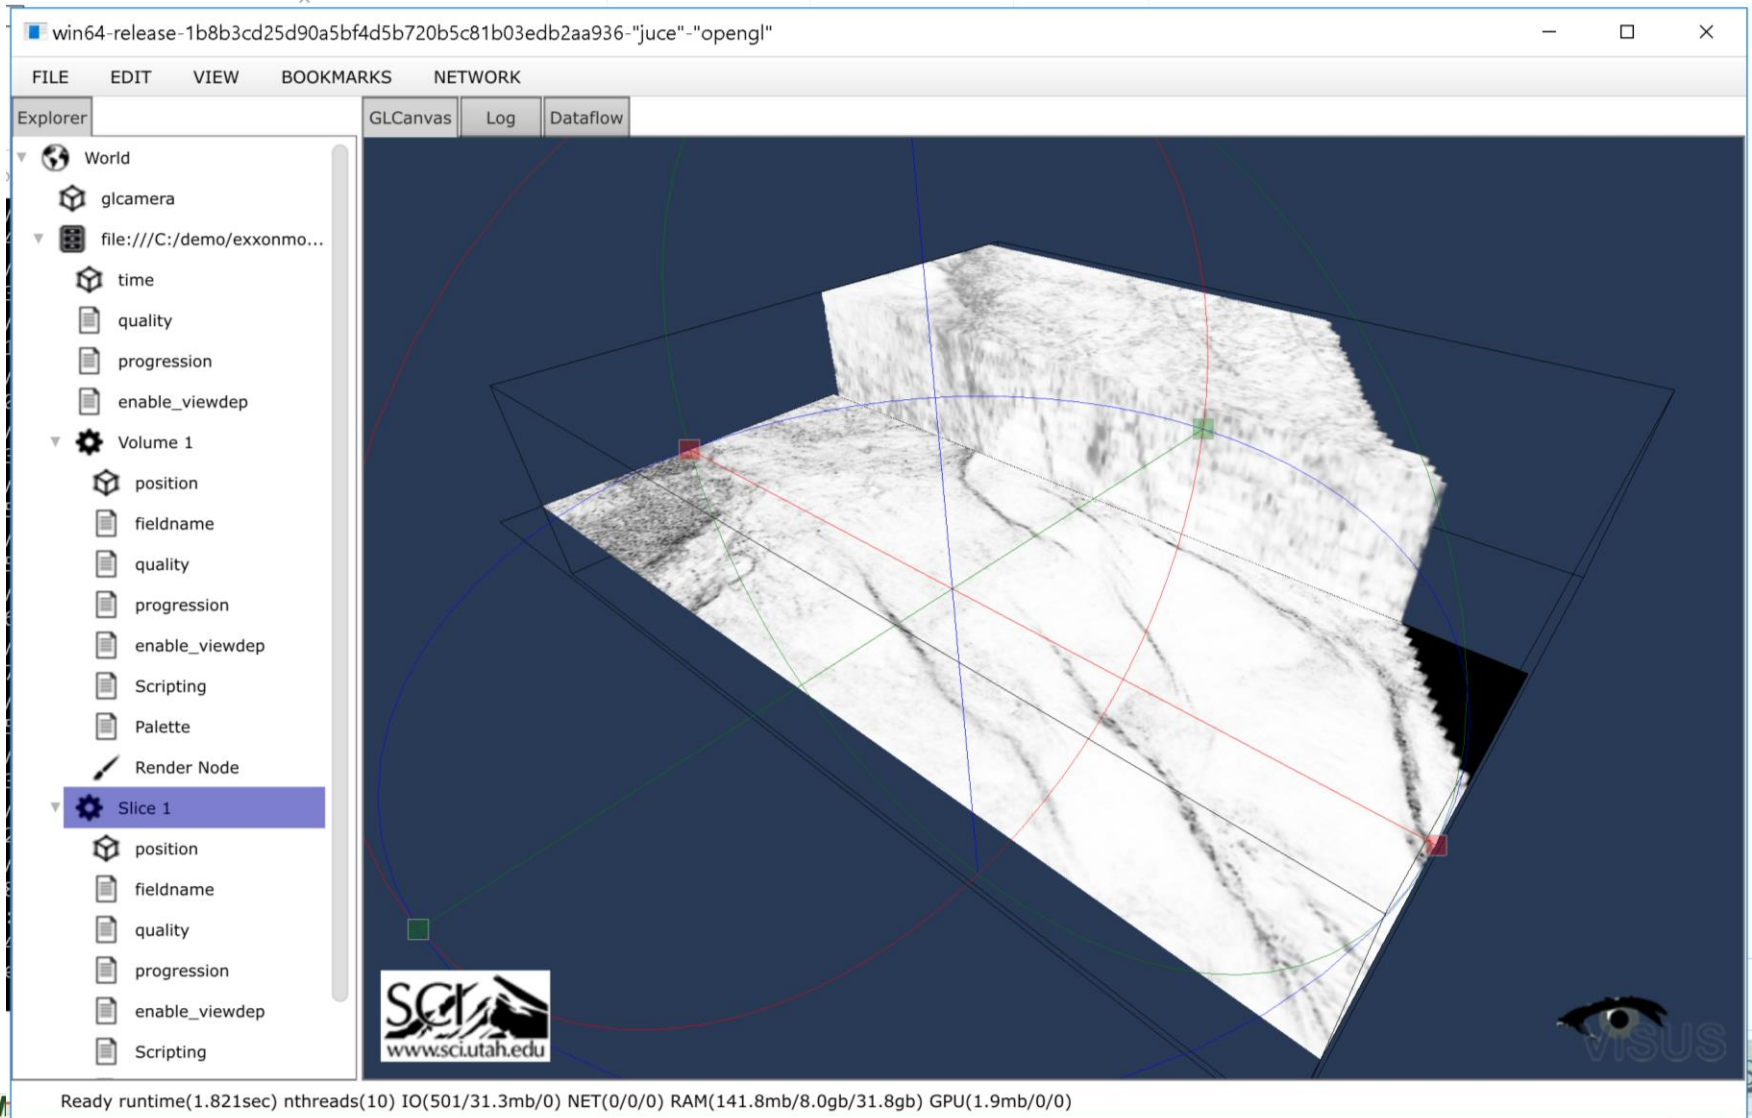
- Reference: <https://github.com/sci-visus/PIDX/>
- Optional: go through the code of the PIDX examples
- Run checkpoint-restart example to produce data in IDX format using PIDX
- Add dataset to the data portal
- Show data from the web viewer of the data portal

# Demo: Interactive Remote Analysis and Visualization of 6TB Imaging Data

- EM datasets of resolution: 130Kx130Kx340




# High Resolution Seismic Models



# Demo: large Scale Geology Data

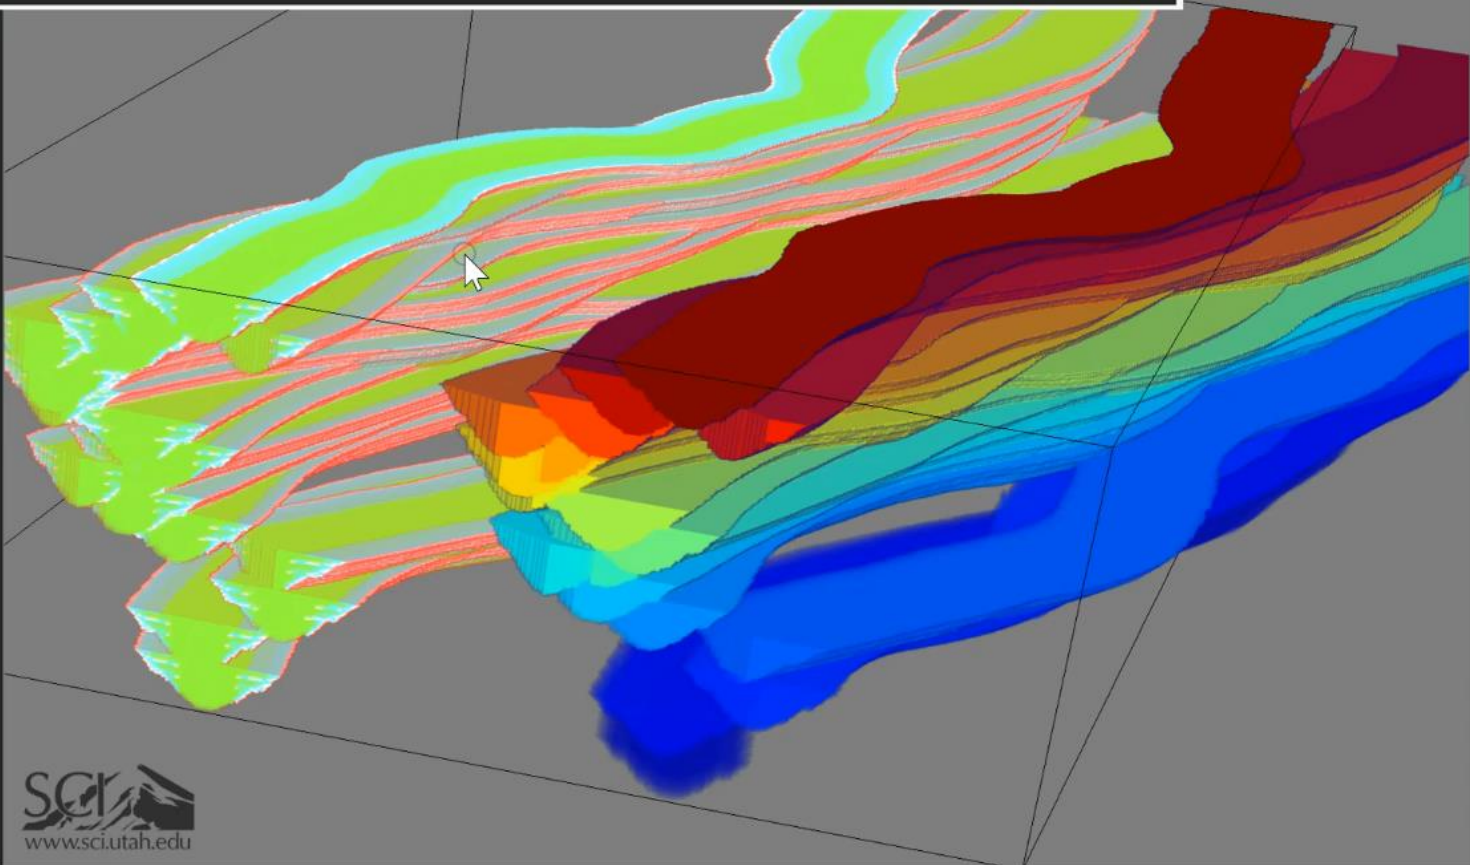
File Edit View Bookmarks

GLCanvas Log Dataflow



Visus is ready RUNTIME(0.398sec) NTHRD(5) IO(0/0/0) NET(5/0/0) RAM(79.8mb/6.9gb/7.9gb) GPU(2.4mb/0/1.5gb)

- fieldname
- render\_bbox
- time
- progression
- quality
- channel
- rock
  - position
  - fieldname
  - enable\_viewdep
  - Transfer function
  - modelview
  - Array render



SCI  
www.sci.utah.edu

Visus is ready RUNTIME(38.948sec) NTHRD(11) IO(2.9kb/183.1mb/0) NET(0/0/0) RAM(1.6gb/7.8gb/7.9gb) GPU(617.5mb/0/1.5gb)


GED



# A Science Cyberinfrastructure Requires Efficient Big Data Management and Processing


- **Advanced data storage techniques:**

- Data re-organization.
- Compression.




- **Advanced algorithmic techniques:**

- Streaming.
- Progressive multi-resolution.
- Out of core computations.




- **Scalability across a wide range of running conditions:**

- From laptop, to office desktop, to cluster of PC, to BG/L.
- Memory, to disk, to remote data access.




# A Scalable Solution for Acquisition and Processing High Resolution Data

Applications: Microscopy for Neuroscience



Visual Cortex (green) of Whole Mouse 318 TB  
Penny-Sized Macaque Brain 31 TB  
Volume 30 TB

(1) Tissue Clearing (31 days)




Large Scale Data Management and High Performance Computing Fully Integrated in a Brain Connectomics Workflow.


(2) Data Source




(3) Preliminary Interactive Analytics



(5) Interactive, Exploratory Assessment and Feedback



(4) Optimal Asynchronous Parallel Processing



# Demo Data Portal: convert and publish

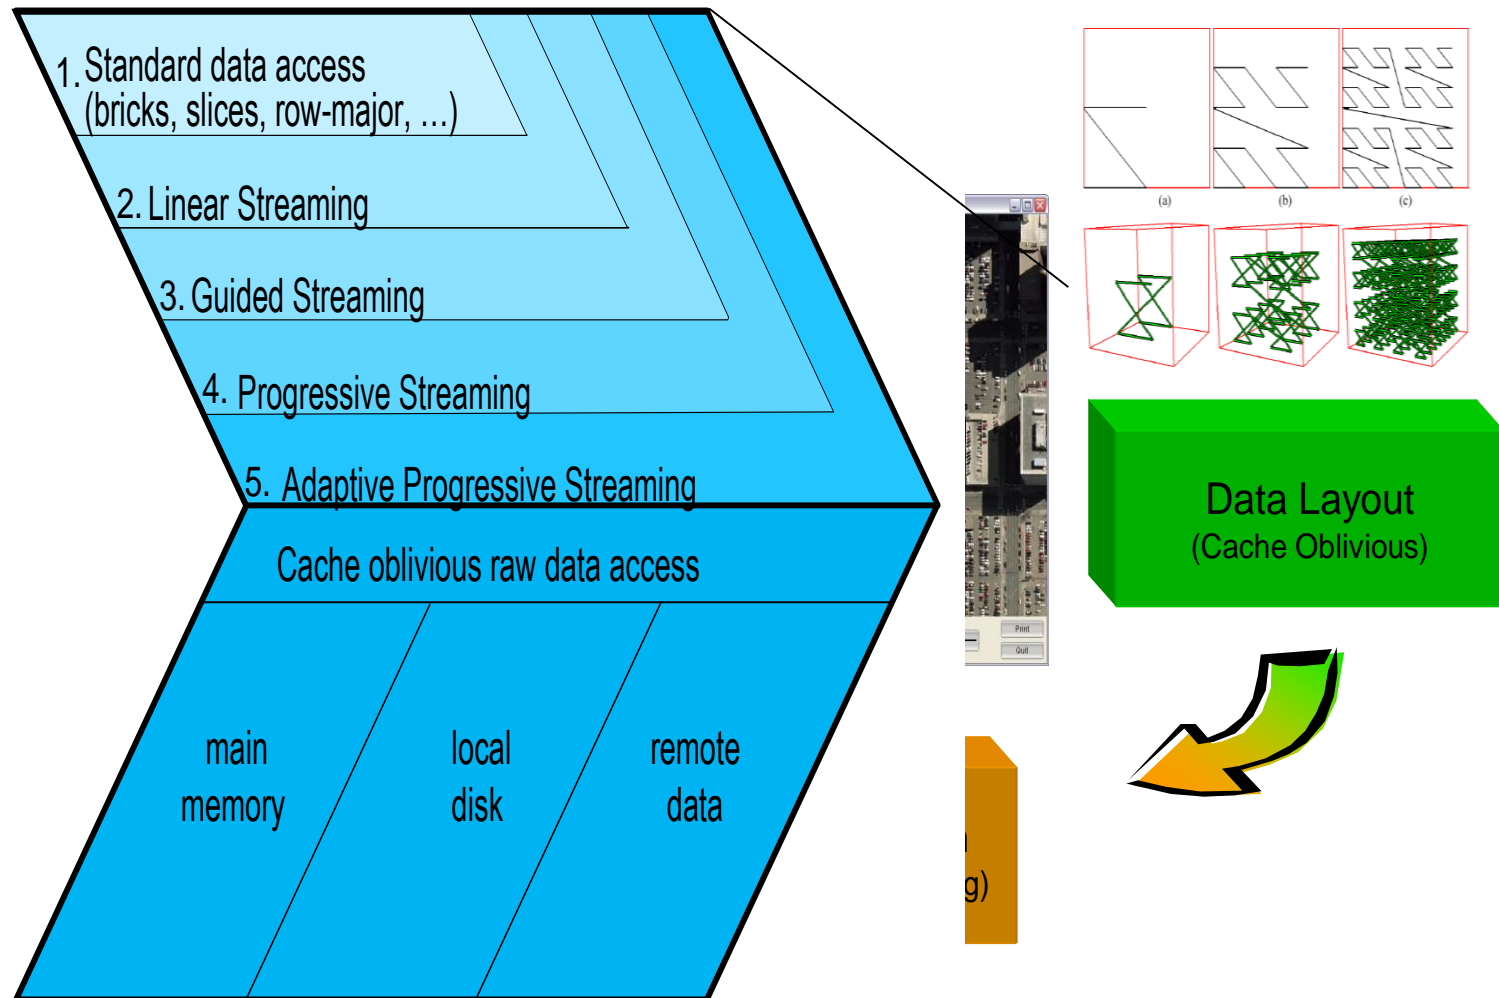
---

- Reference: [https://wiki.visus.org/index.php/ViSUS\\_Data\\_Portal](https://wiki.visus.org/index.php/ViSUS_Data_Portal)
- Sample data (to unzip):
  - Asteroid (500x500x500 float32):  
<https://drive.google.com/open?id=1wHdWynj9jnlf5NrRnzSZyOHbOMwKmltQ>
  - Heart (stack of images):  
[https://drive.google.com/open?id=1NLWe2FVhUUoEaaA8JZ2w61H\\_dT77z1La](https://drive.google.com/open?id=1NLWe2FVhUUoEaaA8JZ2w61H_dT77z1La)

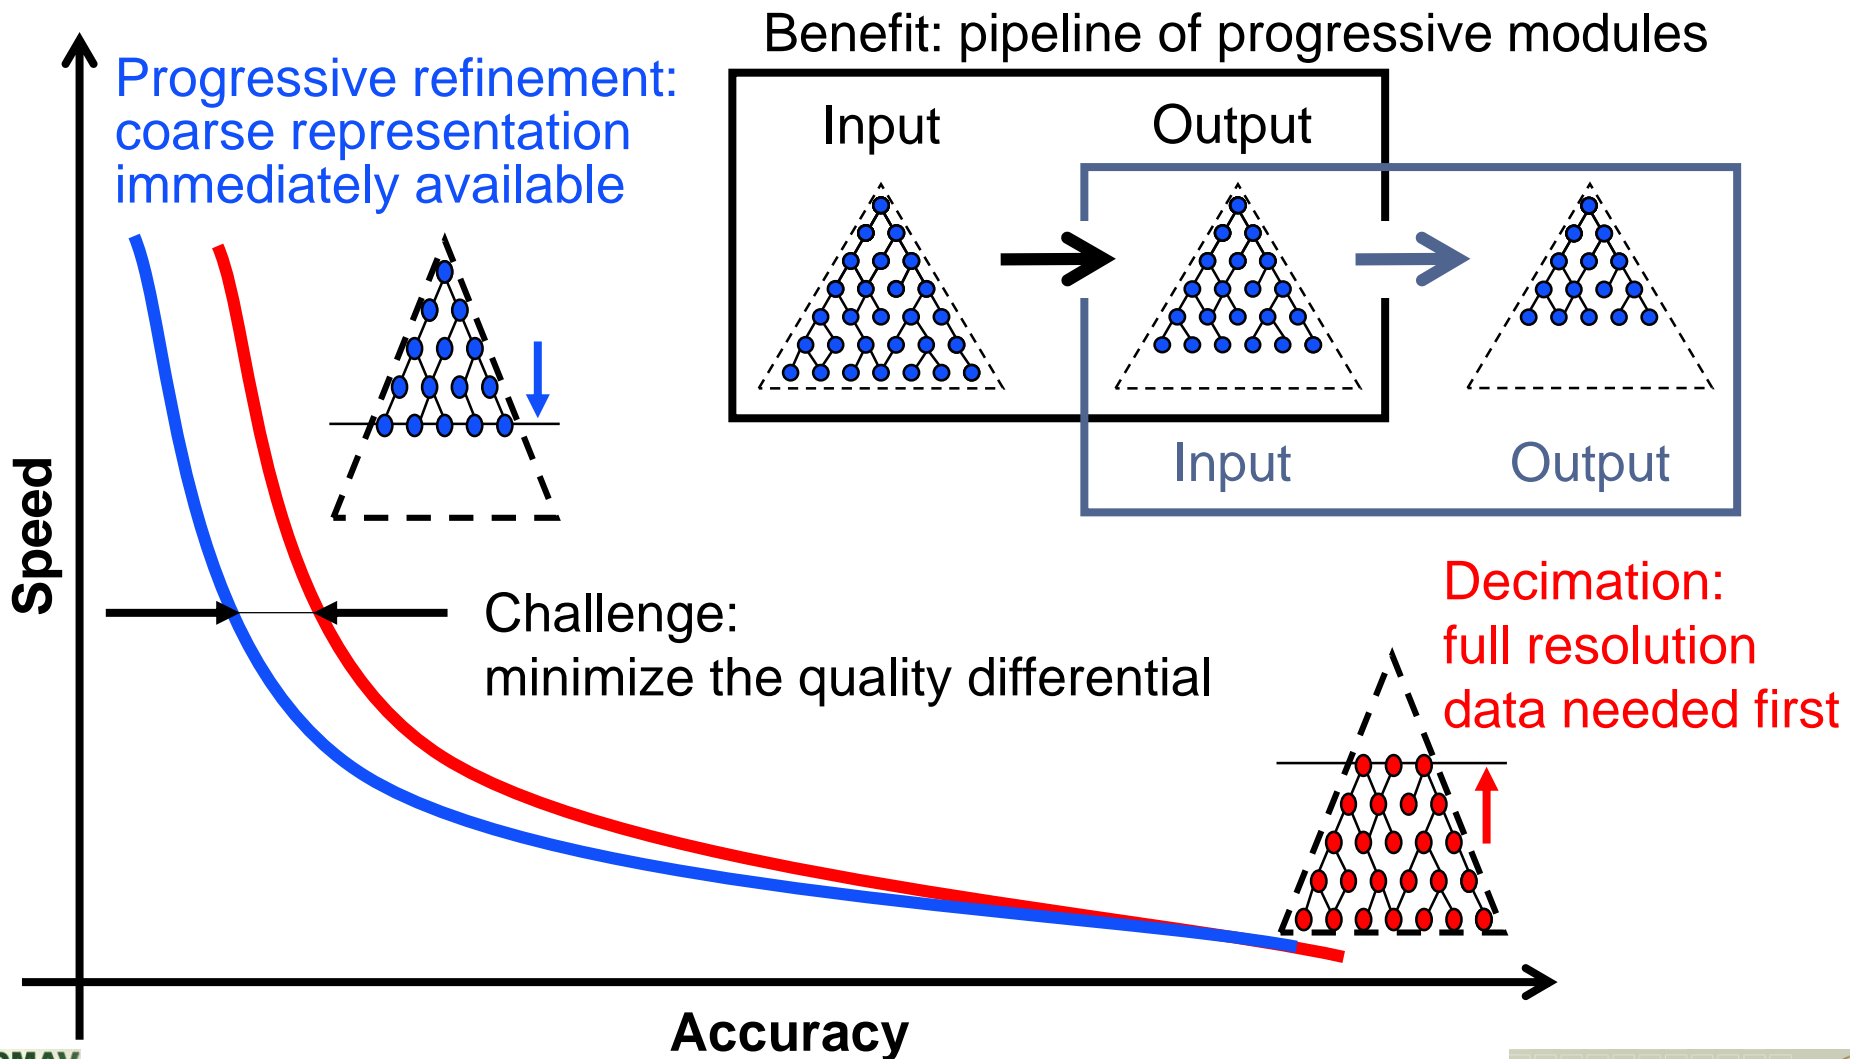
## Demo

- Start/login your Docker installation
- `docker pull visus/dataportal`
- `docker run -p 8080:80 visus/dataportal`
  
- Portal login:
  - User: admin
  - Password: password
- Manage data
  - Asteroid: convert single file (browse, upload, select, insert information, convert)
  - Heath: convert stack of images (browse, upload folder, select folder, convert)
- Add to server
- After 1-2 min data should be available for streaming on the viewer (Explore data)

# We Characterize Algorithmic Classes Based on Effect in a Processing Network

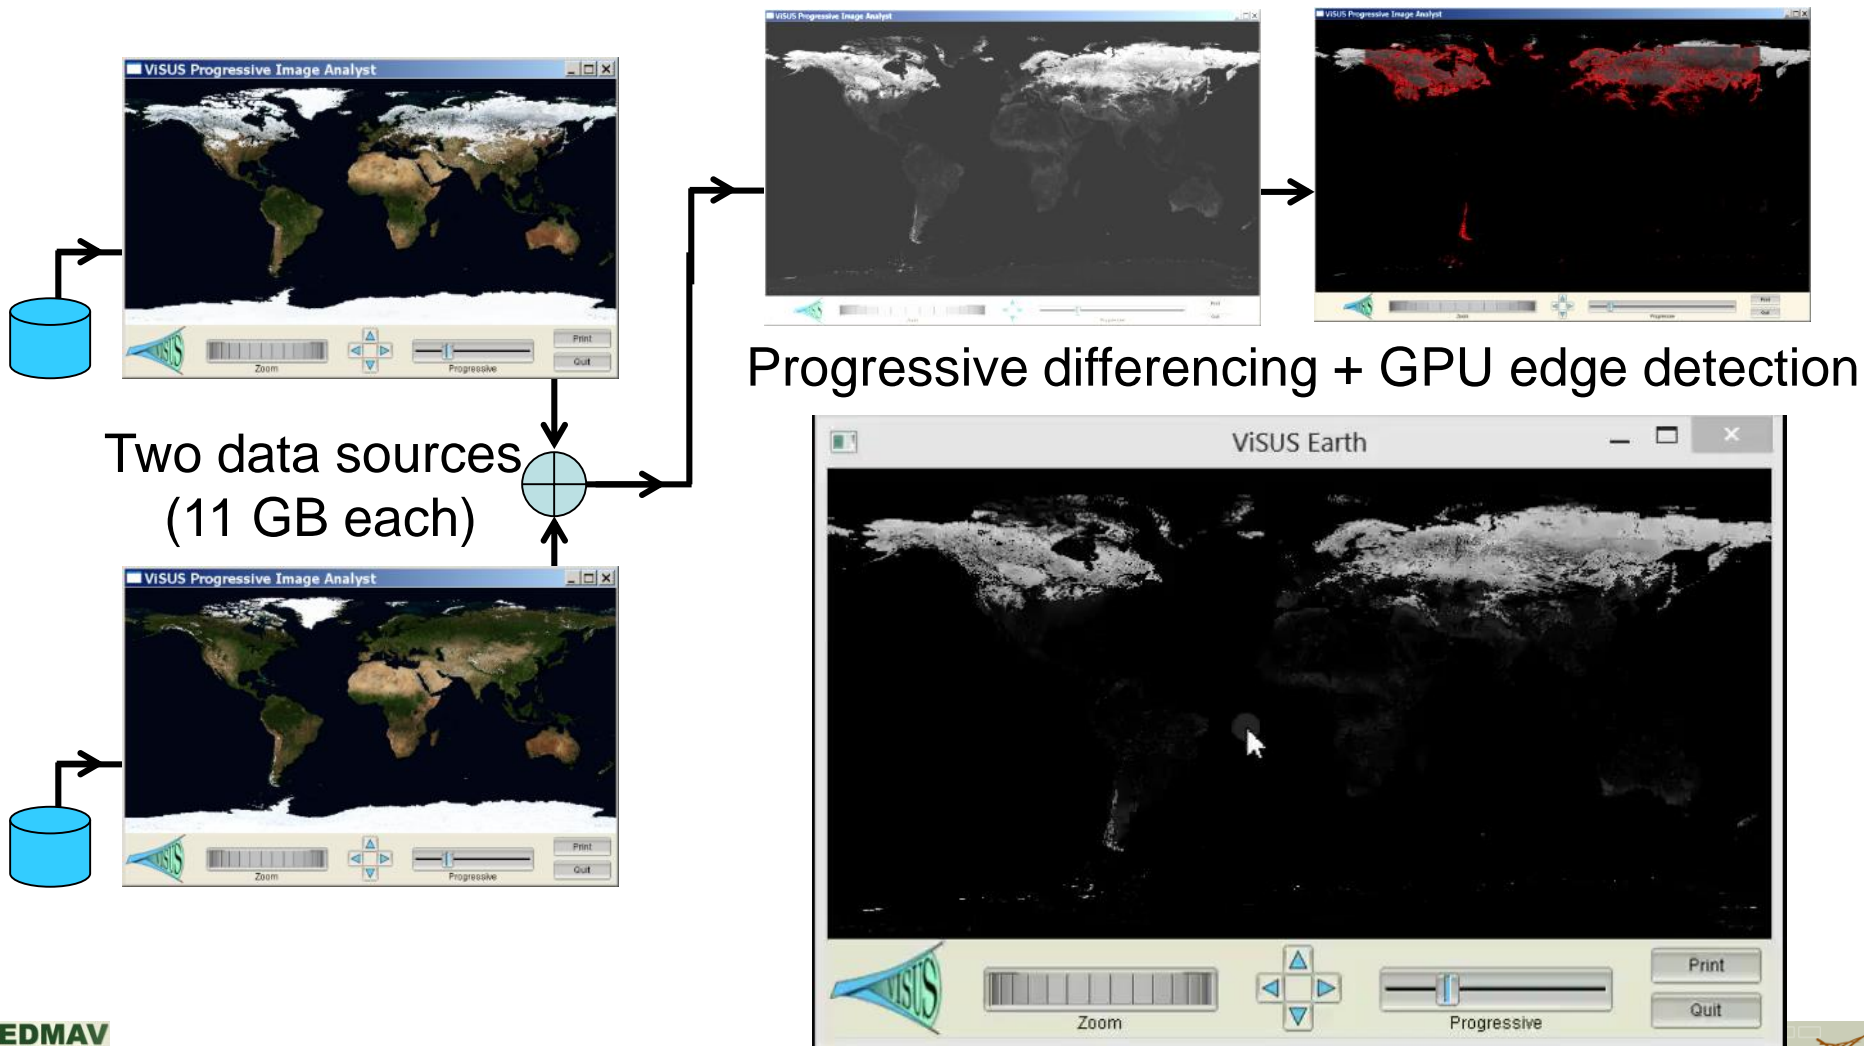


# The use of top-down and bottom-up processes have a strong impact on the data stream




# We Allow Distributed Computations at Different Stages of the Data Stream

- Progressive Image Differencing + Editable GPU filter.

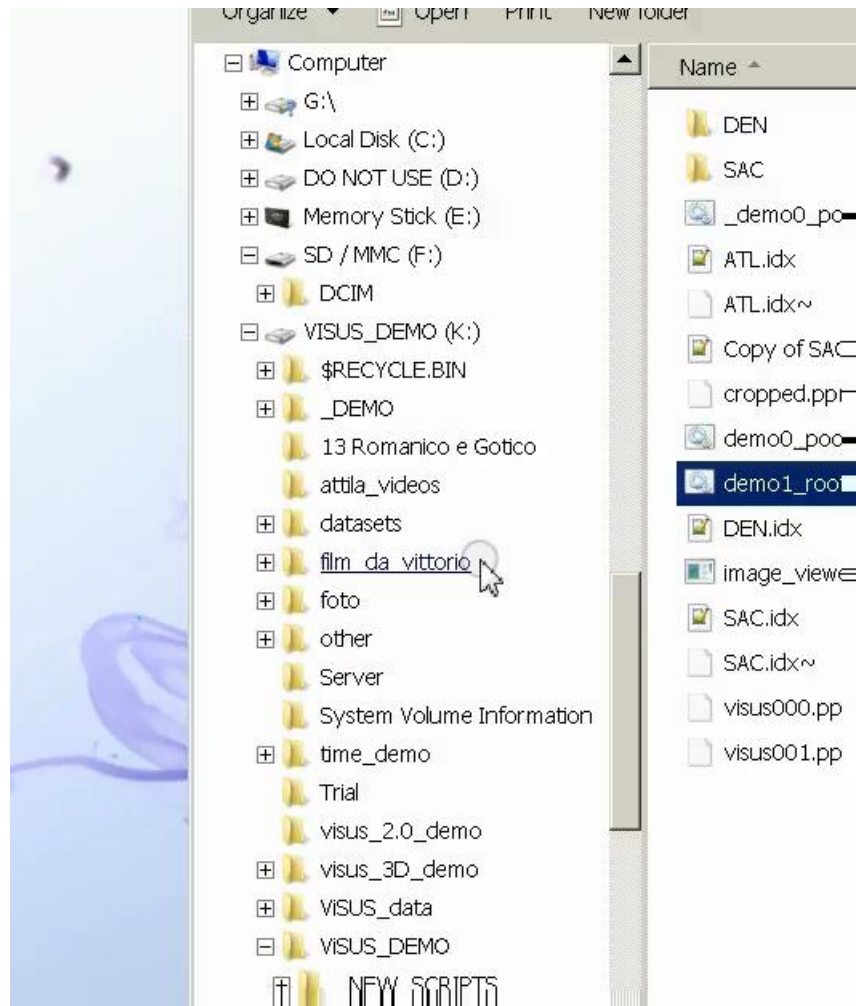
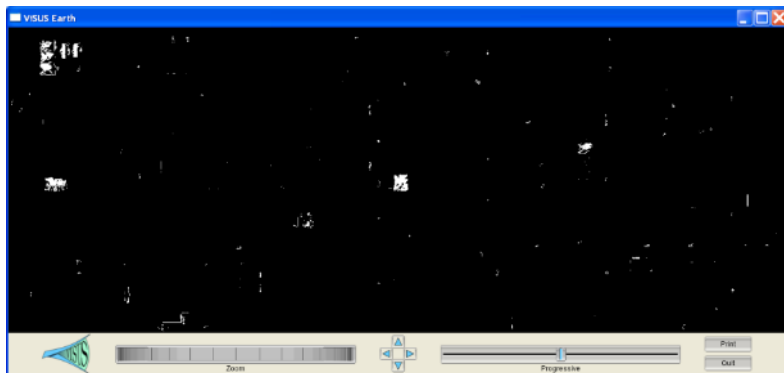


# We are Developing Progressive Scheme for Content Based Image Processing

- **Sample:**



- **Progressive Analysis:**



# Poisson Solver for Image Cloning in Massive Image Collections

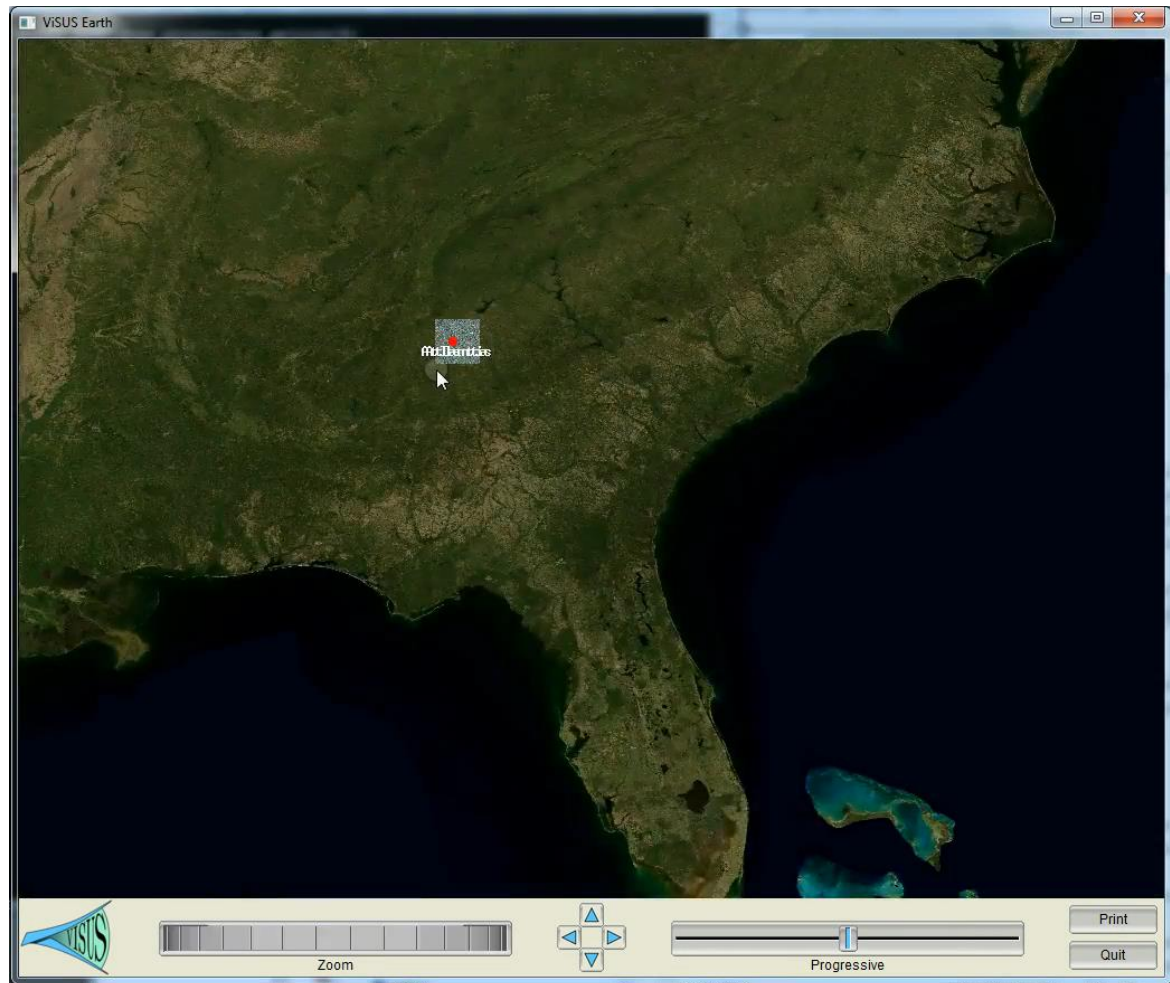

- Color correction of 600+ images in real time





# Poisson solver for composition of massive images

- Pasting a 300GB satellite image of a city in background world map merged in real time



# Server can be wrapped in Apache plug-in Client can be run in a web browser



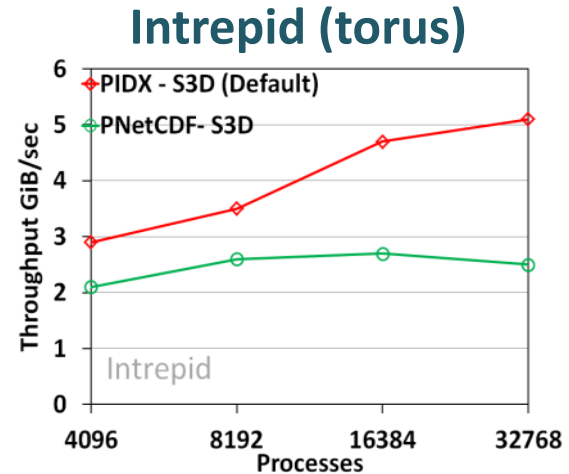
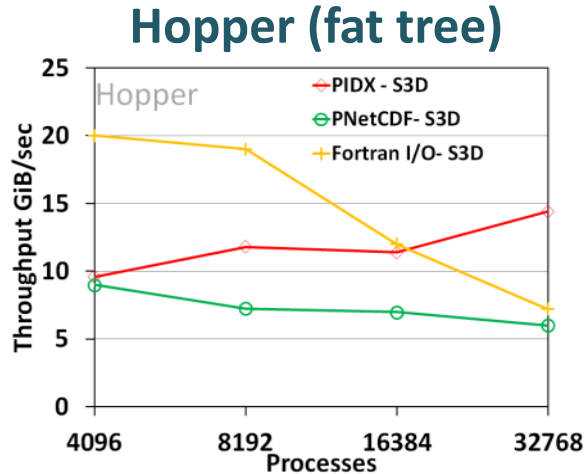
**One billion polygons  
to billions of pixels**

Welcome to the first  
gigapixel, multi-view  
rendering of  
The Digital Michelangelo  
Project's David

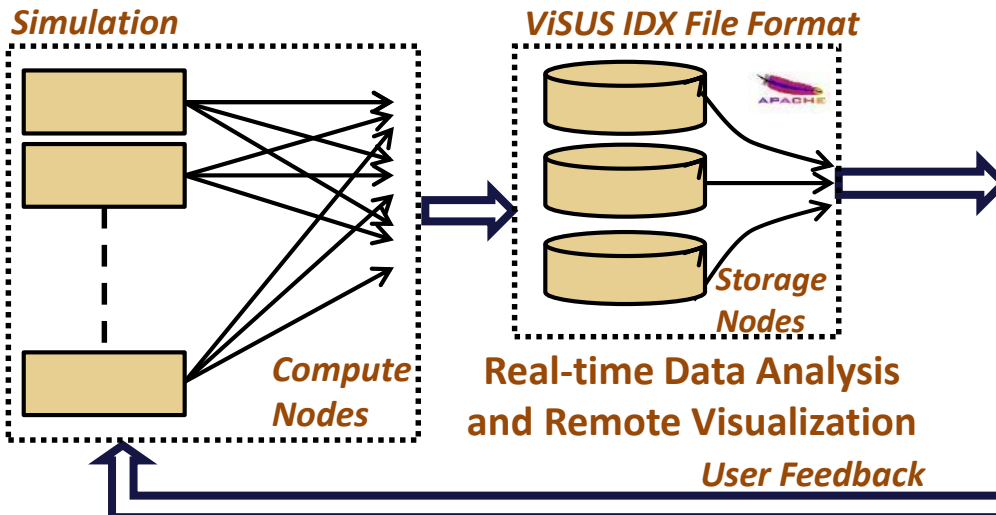
**GIGAPIXELDAVID**

# High Performance Data Movements for Real-Time Monitoring of Large Scale Simulations




Scale simulation dumps to 130K cores with better performance than state of the art libraries while enabling real-time, remote visualization

*End User*



# PIDX High Write Performance on Mira



187GB/s

# Server can be wrapped in Apache plug-in Client can be run in a web browser

One billion polygons to billions of pixels

www.sci.utah.edu/news/60/431-visus.html

THE UNIVERSITY OF UTAH

SCI Home About SCI Research Publications Media Software Collaborators People Contact

The Scientific Computing and Imaging Institute at the University of Utah

SCI Home » News » Projects » One billion polygons to billions of pixels

SCI Home  
About SCI  
Research  
Publications  
Media  
Software  
Collaborators  
People  
Contact  
News  
Seminars & Events  
Archived News and Feature Stories

### One billion polygons to billions of pixels

Welcome to the first gigapixel, multi-view rendering of The Digital Michelangelo Project's David.

The David model consists of 833 million triangles from a laser-scan of the original statue created by Professor Marc Levoy and members of The Digital Michelangelo Project at Stanford university. The model was aligned by Benedict Brown and Szymon Rusinkiewicz using the non-rigid alignment method described in their 2007 SIGGRAPH paper.

Each of the 4 2-gigapixel sized frames (29280 x 70416 pixels) was rendered using the Manta Interactive Ray Tracer. Manta is a highly portable interactive ray tracing environment designed at the SCI Institute to be used on both workstations and super computers. For these renderings, Manta leveraged a recursive 4-level grid to accelerate the rendering. In all, each frame took 30 hours to render using 64 cores each (256 total) of the SCI Institute's 204 core SGI UV 1000 with 2.8TB of RAM and 2.87GHz Intel Xeon X7542 cores. More information on Manta can be found at: [http://manta.wsl.sci.utah.edu/manta/index.php/Main\\_Page](http://manta.wsl.sci.utah.edu/manta/index.php/Main_Page)

The final rendering was stored in the hierarchical, space-filling curve format of the VISUS technology. VISUS intelligently reorganizes the raw data enabling efficient, streaming pipelines that process the information while in movement. The results are then visualized in a progressive environment, allowing for meaningful explorations with minimal required resources. This technology enables real-time management of large datasets on a variety of systems ranging from desktops and laptop computers to portable devices such as iPhones/iPads. VISUS has been deployed in a variety of large data applications such as the monitoring of large scientific simulations and the editing of massive images and panoramas.

The VISUS David viewer is currently available as a Windows web browser plugin (Firefox and Chrome) or as a standalone application for Windows, Mac OS X, or OpenSUSE. Please follow the links below to access the gigapixel David.

- Download and install VISUS application and plugin for Windows
- Download the VISUS application for OS X
- Download the VISUS application for OpenSUSE
- View David via web plugin
- View David via web plugin (for slower connections)

Scientific Computing and Imaging Institute

VISUS

atlantis.sci.utah.edu

## Visus

Streams for Ultimate Scalability

Official Visus site [www.pascucci.org](http://www.pascucci.org) [sci.utah](http://sci.utah)

Privacy Policy | Contact Us

### Project description

A Research Project in Scientific Visualization centered on the development of scalable distributed approaches for the management, streaming and rendering of large surface and volume models. Principal Investigator: Warren Pascucci Other team members: Tilo Lauer, Daniel Lauer, Peter Lindstrom, Main Collaborators: Mark Schmalzer, Donald A. Frame, George Sorokin, David Bremer, Guenter Janke, James Barnes (Ph.D.), Bill Sargent (Ph.D.), Paul Tassi, Steven (Ph.D.), Jose Cabe-Abadillo (SD), Ajay Mathuram (Ph.D.), Vijay Natarajan (Ph.D.), Steven Piro (Ph.D.), Jonathan Truitt (MS)

In the VISUS project (see the featured article in the SIGGRAPH report) we devised data streaming techniques for progressive processing and visualization of large scientific datasets. Our strategy is to split the loading between hierarchical algorithms and progressive multi-resolution data structures to reduce an end-to-end optimized flow of data from the original source, such as remote storage or large scientific simulation, to the rendering hardware. The implementation of this approach, called *streaming for major visualization modalities*, (1) interactive visualization on high resolution devices, (2) interactive visualization on desktop workstations, (3) interactive monitoring of remote productions from a desktop workstation. These innovative, targeted multiple phases in the process of generating and exploring very large simulation datasets where real-time user interaction can increase the productivity of scientists.

### Download Windows Visus Installer

<b>San diego</b>		<b>zkbitt</b>	
Type: RGB	Type: Gray		
Size: 204 Gb	Size: 8 Gb		
Width: 200,000	Width: 2048		
Height: 185,000	Height: 2048		
	Depth: 2048		

<b>Visible male</b>		<b>Nuclear</b>	
Type: RGB	Type: Flow32		
Size: 12 Gb	Size: 1 Gb		
Width: 2048	Width: 800		
Height: 1216	Height: 800		
Depth: 1878	Depth: 100		
	Num fields: 7		


<b>Atlanta</b>		<b>Phoenix</b>	
Type: RGB	Type: RGB		
Size: 295 Gb	Size: 1038 Gb		
Width: 520,001	Width: 720,000		
Height: 527,001	Height: 540,000		

<b>Hamilton</b>		<b>San Francisco</b>	
Type: RGB	Type: RGB		
Size: 169 Gb	Size: 207 Gb		
Width: 346,000	Width: 225,000		
Height: 232,000	Height: 330,000		

<b>Chattanooga</b>		<b>Microscopy</b>	
Type: RGB	Type: RGB		


# A Scalable Solution for Acquisition and Processing High Resolution Data

Applications: Microscopy for Neuroscience



Visual Cortex (green) of Whole Mouse Brain 318 TB  
Penny-Sized Macaque Brain 31 TB  
Volume 30 TB


(1) Tissue Clearing (31 days)




(2) Data Source



(3) Preliminary Interactive Analytics




(5) Interactive, Exploratory Assessment and Feedback



(4) Optimal Asynchronous Parallel Processing



Large Scale Data Management and High Performance Computing Fully Integrated in a Brain Connectomics Workflow.

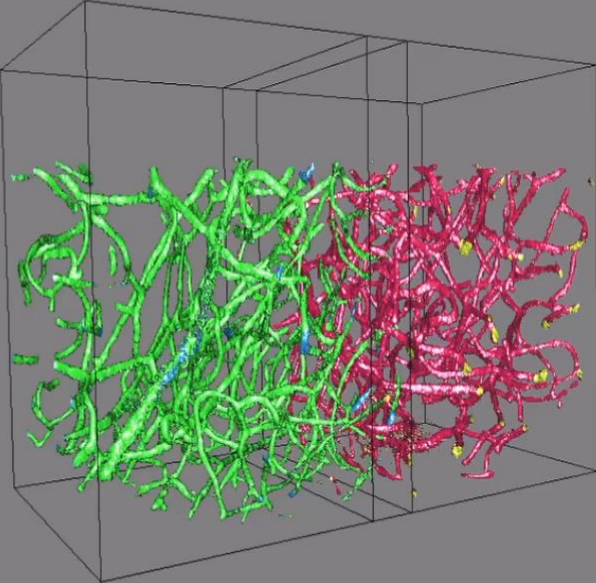


# Remote Monitoring of Data Quality During Acquisition

EDIT VIEW BOOKMARKS NETWORK


GLCanvas Log Dataflow

ord  
enable\_viewdep  
progression  
quality  
glicamera  
file:///D:/research/idx/MM360\_March\_4\_2\_/MM360\_March\_4\_2\_0...  
time  
Volume 1  
position  
fieldname  
quality  
progression  
enable\_viewdep  
Scripting  
Marching cube  
Palette  
Mesh Render  
Palette  
file:///D:/research/idx/MM360\_March\_4\_2\_/MM360\_March\_4\_2\_0...  
time  
Volume 1  
position  
fieldname  
quality  
progression  
enable\_viewdep  
Scripting  
Marching cube  
Palette  
Mesh Render  
Palette



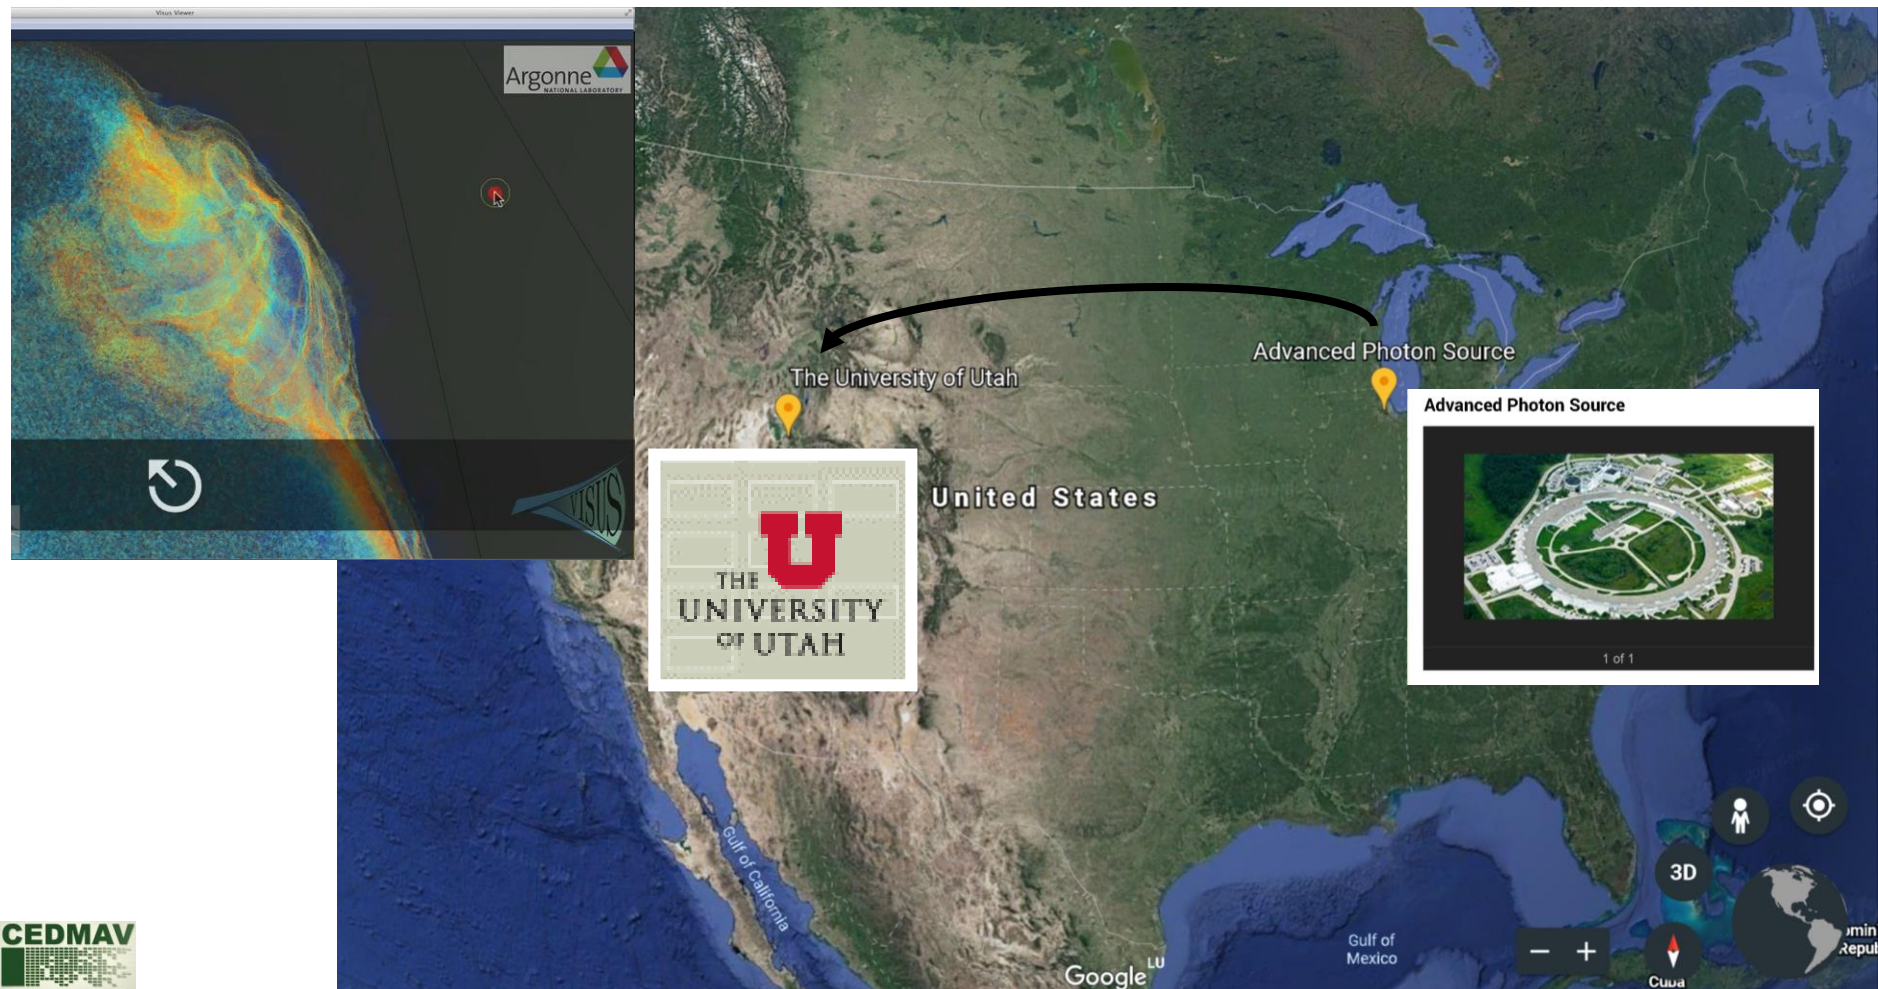
ing. TJOB(4) NJOB(0) nthreads(12) IO(0/0/0) NET(0/0/0) RAM(1.3gb/7.2gb/63.9gb) GPU(7.9mb/0/4.0gb)

SCR www.scri.utah.edu



# High Performance Data Movements for Real-Time Access to Large Scale Experimental Data

- Experiment run at Advance Photon Source at ANL
- Materials Scientists at University of Utah






# Scalable Deployment: Real Time Exploration of 3.5 Petabytes of Weather/Climate Data

## Workflow


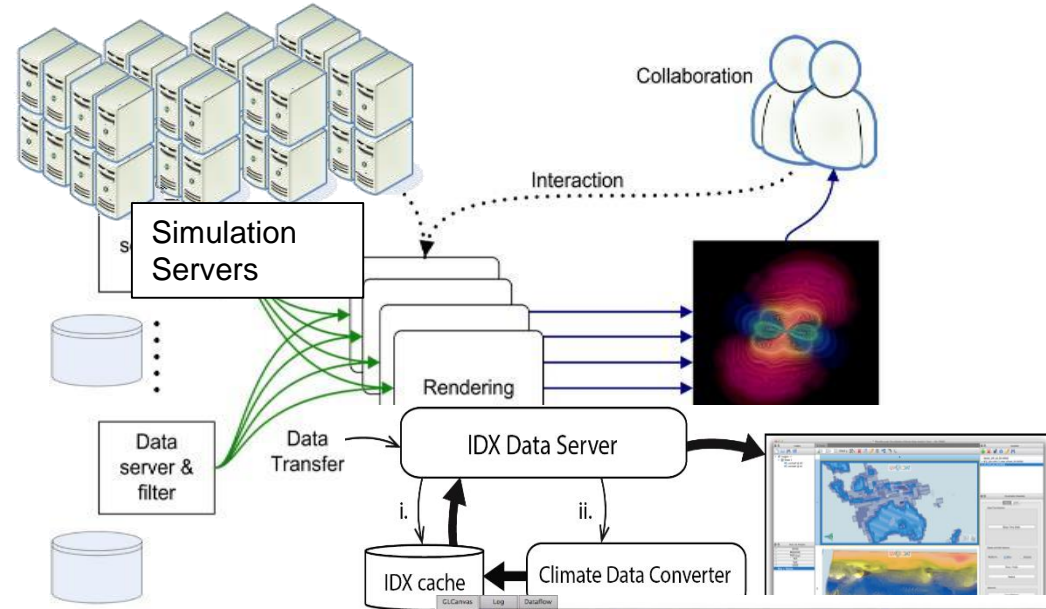
- **Data creation**
  - **Data Management**
- Processing
  - Analysis
  - Visualization



- 7km GEOS-5 “Nature Run” -> 1 dataset, 3.5 PB
- theoretically: openly accessible -> practically: precomputed pics

## Distributed Resources


- 3.5 PB of data store in NASA
- Primary ViSUS server in LLNL
- Secondary ViSUS server in Utah
- Clients connect remotely
- Work without additional HPC resources





[http://atlantis.sci.utah.edu/visus/webviewer/nature\\_2007\\_aer1\\_hourly](http://atlantis.sci.utah.edu/visus/webviewer/nature_2007_aer1_hourly)

# Streaming Analytics and Visualization

Live demonstration from ANL to SLC



Infrastructure that scales gracefully with available hardware resources



# Demo Interactive Data Analysis with Python

---

- Reference: <https://github.com/sci-visus/OpenVisus>
- Jupyter examples: <https://github.com/sci-visus/OpenVisus/tree/master/Samples/jupyter>
- Installation with pip:
  - `python -m pip install --user numpy OpenVisus==1.3.3`

## Demo

- Use standalone viewer and python scripting for interactive analysis
- Use Jupyter notebook to fetch data from a server visualize, analyze and share with other tools (e.g., with the standalone viewer)

# High Resolution Display Platforms for High Resolution Outcrop and Seismic Data

---


---



# ViSUS PowerWall: Installed and Fully Operational in a Few Hours at KAUST

---

---




# ViSUS PowerWall: RIKEN



# We Allow Distributed Computations at Different Stages of the Data Stream

- Progressive Image Differencing + Editable GPU filter.



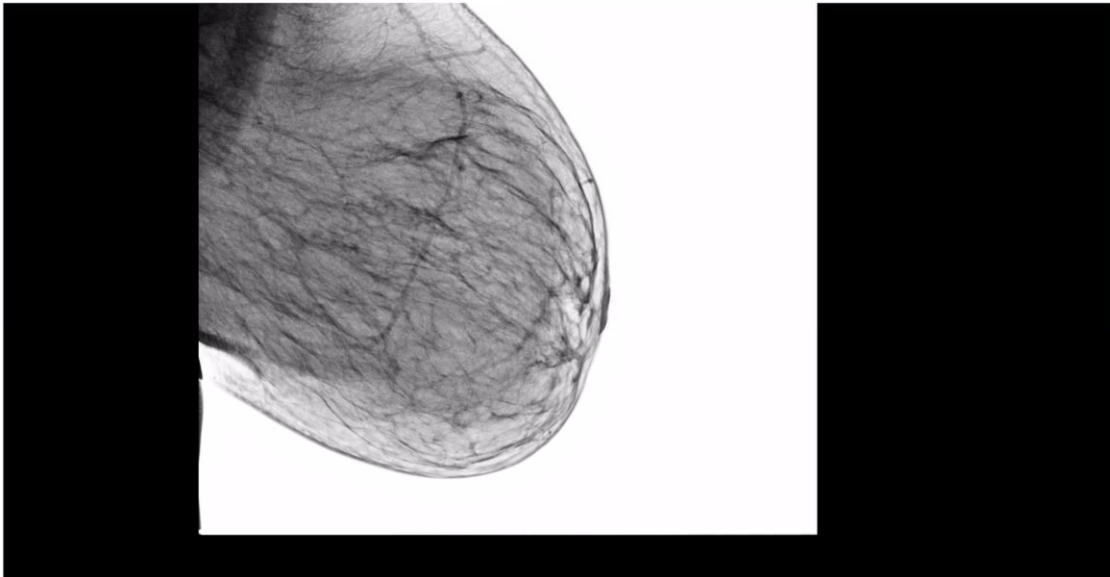
# <https://visoar.org/mamografia/>

Mamografía


https://visoar.org/mamografia/

## Mamografía del hospital Castro Rendon

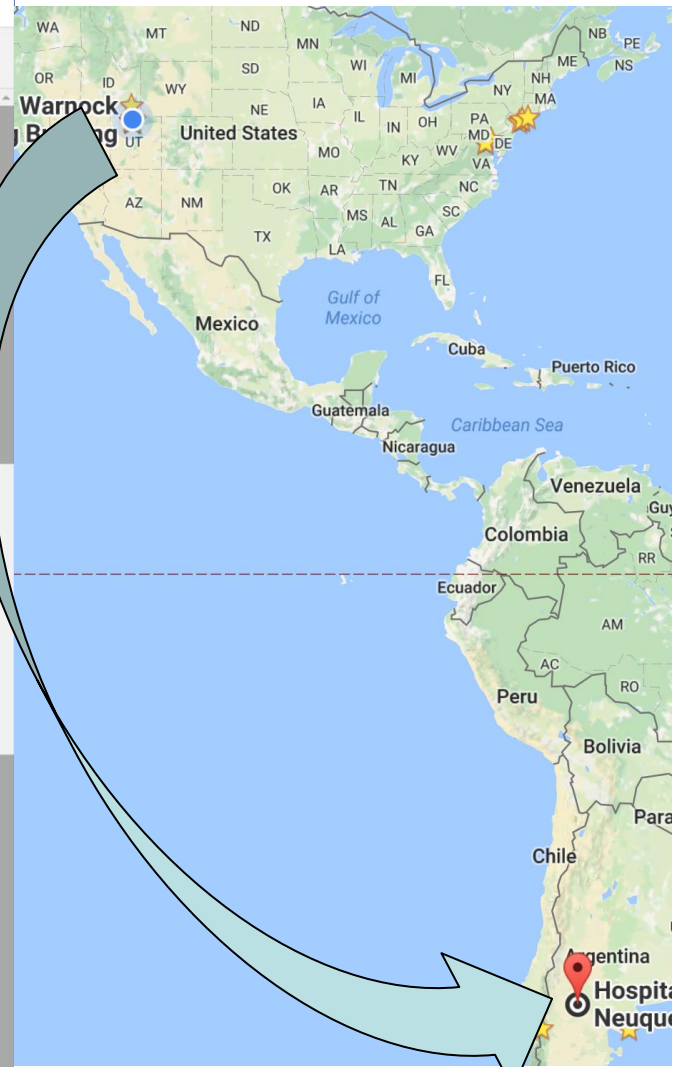
### Visualization of Mamografia



Hospital Provincial Neuquén  
Dr. Eduardo Castro Rendon



© 2016 POWERED BY SCI, U of Utah





# ViSOAR: a Unified Solution for Distribution of Imaging Data in Medical and Geology

## Example of Visualization and Annotation of outcrop and medical data



**Healthcare:** remote access and diagnostics for Doctor and Patients on commodity devices



Mohamed E. Salama, MD, Chief of Hematopathology, Professor of Pathology



# ViSOAR Medical

---

---



**Dr. Rick Ash**  
Neurobiology and Anatomy

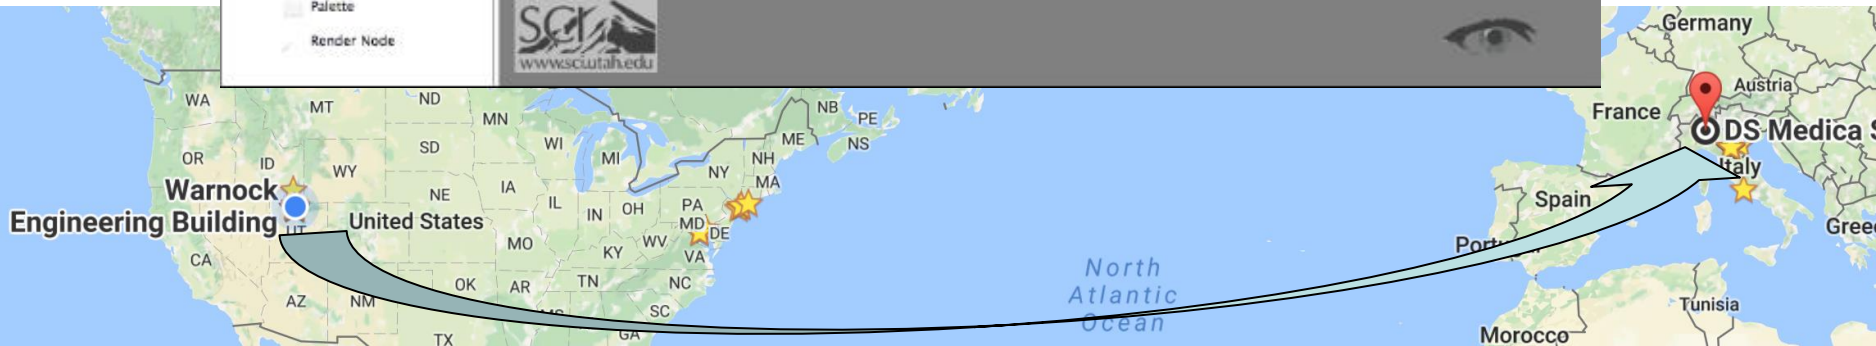

# Teaching Histopathology with a Great Group of Students!!!

---

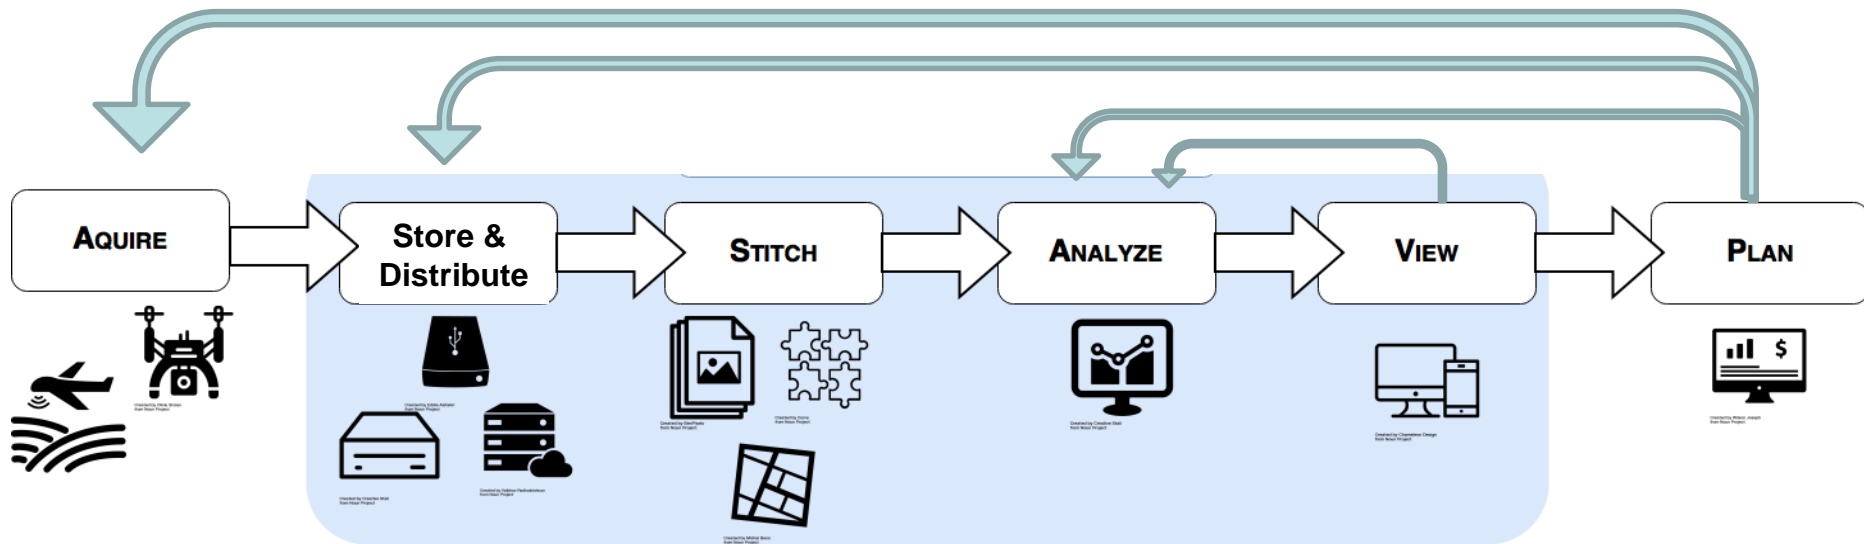


For more information see:  
[www.visus.org](http://www.visus.org)

# Computed Tomography Exploration: Volume Rendering





# An OpenViSUS Pipeline for Dynamic Data-Intensive Agricultural Applications



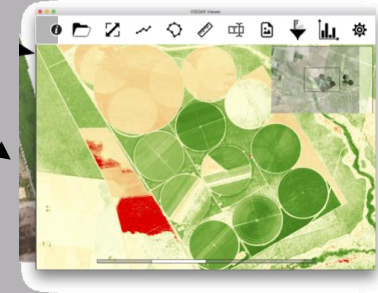
# ViSUS Gigapixel Progressive Ditching for Aerial Imagery




Registration  
& Stitching



Analytics



# ViSUS Gigapixel stitching for aerial imagery



# Farmers want actionable information...



**To achieve this goal and enable wide-spread adoption of aerial imagery use in crop management there is an urgent need for real-time image stitching on high-latency low-bandwidth networks in rural areas**



# OpenVisus references

---

- Main website: <https://www.visus.org/>
- Documentation website: <https://wiki.visus.org/>
- Gitter chat: <https://gitter.im/sci-visus/OpenViSUS>
  
- PIDX code: <https://github.com/sci-visus/PIDX/>
- OpenViSUS code: <https://github.com/sci-visus/OpenVisus>
- Jupyter examples: <https://github.com/sci-visus/OpenVisus/tree/master/Samples/jupyter>