CIF21 DIBBs: Ubiquitous Access to Transient Data and Preliminary Results via the SeedMe Platform

Award Number : 1443083

PI: Amit Chourasia, Co-PI: Michael Norman

San Diego Supercomputer Center, University of California San Diego

Open source ✓

Web based ✓

Cross-platform ✓

What is SeedMe?

SeedMe = Stream Encode Explore and Disseminate My Experiments

SeedMe is a platform that enables easy sharing of transient and preliminary data for a broad research computing community by offering cyberinfrastructure as a service and a modular software stack that may be customized. SeedMe is based on Drupal content management system with a set of building blocks with additional PHP modules and web services clients.

Significant results

✓ Federated authentication	Virtual File Sytem
✓ Small Data Visualization	Small Data API ✓
✓ Pilot project with 580+ registered users 120,000+ content items Integration with four scientific tools	Early demonstration prototype ✓

✓ **Available**

Why build a platform?

Research computing is highly collaborative, distributed and often uses disparate compute resources. Currently available tools do not meet sharing and collaborative needs that must collocate **d**ata, **d**escription and **d**iscussion (3D) and additionally handle transfer, storage and access control. Furthermore, these tools must be cross platform and readily pluggable for automation with existing scientific workflows.

For whom?

Computing researchers
- Collaboration hub
- Personal dashboard

Developers
Integrate with scientific applications

Project repositories

Gateways
Service for data sharing, data publishing, data escrow

CI providers: Offer the platform to your user base

Use how?

As a cloud service
- dibbs.seedme.org
- www.seedme.org

DIY - Run own instance
- On your own hardware
- Condo hardware

Provider run instances
- At your institution
- At national centers
- At public cloud

~~**No lock in**~~

Planned building blocks

Federated authentication + Authorization	Virtual File Sytem	Access Control
Sharing	Search / Index	Microformats
Field Formatters	REST API	Clients (Java, Python) + Command Line
Light Visualization	Rich Text	Discussion

Virtual file system: Sample UI

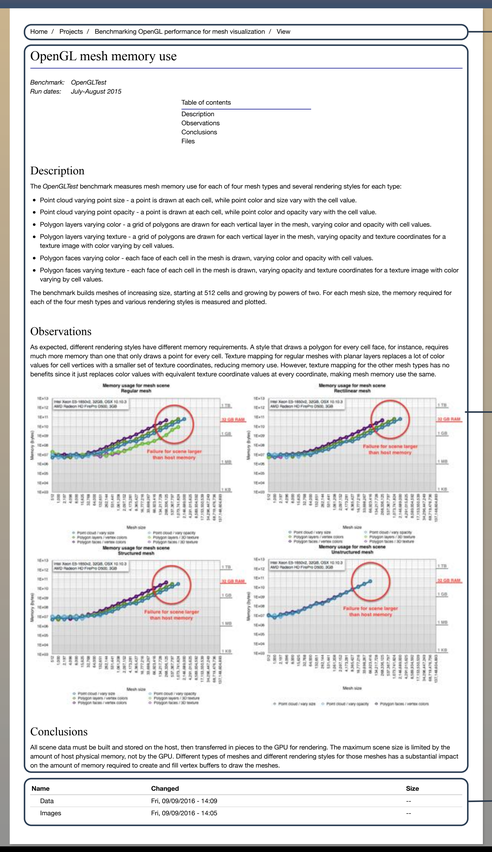

Light visualizations

Example - Image classification work breakdown schedule

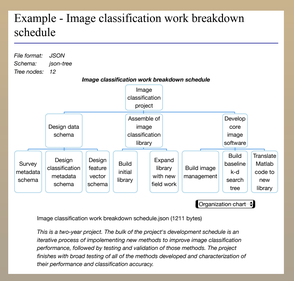

Example - Image preprocessing times for images using multithreaded C++

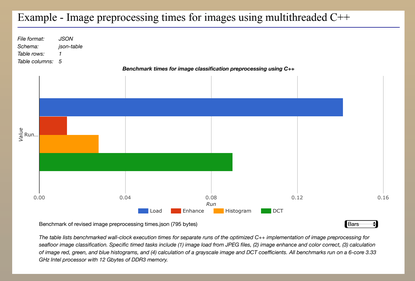

Example - OpenGL mesh memory use as mesh sizes scale up

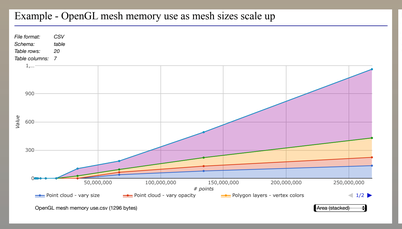

Example - Image preprocessing times for images using Matlab

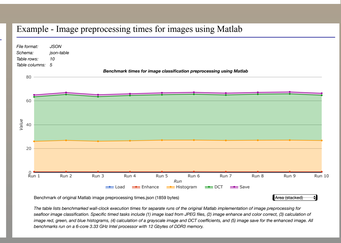

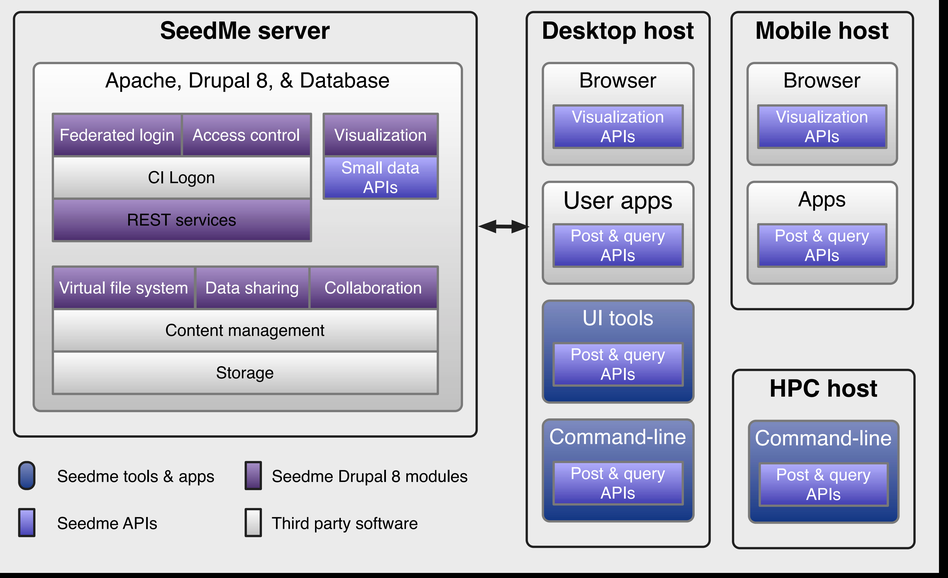

Project information

Website: https://dibbs.seedme.org

- Cross-platform tools, APIs, and Drupal modules
- Post & query data from HPC jobs, workflows, apps, browsers, and command line
- Cloud file system for secure data sharing and collaboration
- Integrated lightweight visualization tools for quick analysis
- Secure access, sharing, and access controls

Acknowledgements: This work is supported by the National Science Foundation under Grant No. ACI-1443083. "Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF."