

Hiding Virtual Computing and Supercomputing inside a Notebook: GISandbox Science Gateway & Other User Experiences

Eric Shook

Domain Champion for GIS, XSEDE

Department of Geography, Environment and Society

University of Minnesota

eshook@umn.edu

My Background

- Started as a Computer Science graduate student working with Open Science Grid and the LHC (eventually as a grid and systems administrator)
- Switched to Geography and now my research is cyberinfrastructure-based geographic information science and systems (CyberGIS)
- Helped develop two XSEDE Science Gateways: *GISolve* and the *CyberGIS Gateway*
- Now developing my own science gateway: *GISandbox*

My Perspective

- My background as a grid and cluster administrator and science gateway developer has influenced my perspective as a user
- I can appreciate the difficulties in developing usable systems (and their documentation!)
- I can *really* appreciate bugs, early prototypes, and associated difficulties with shared systems
- But as a user, I appreciate an easy-to-use system. :)

My Quest

It should be as easy as
or easier than using my
desktop computer*

* Whether it be supercomputers, virtual computers, containers, clouds, ...



My Quest

Use Cases

GISandbox

Considerations



Singularity Containers

Comparing Spatial Computing Systems

- A small interdisciplinary team aiming to compare the performance of several established and cutting-edge systems for processing big spatial data
 - Spatial databases
 - Spark-based systems
 - Parallel programming languages for spatial data
- Perfect use case for containerization
 - Setup and configure the container and software first
 - Run each system on XSEDE, profiling the execution
 - Tweak system configurations to improve performance
 - Wash, rinse, and repeat



Early User Experience on Bridges

- Bridges supports Singularity (common across XSEDE)
- Documentation for Singularity is minimal
 - A quick search on other XSEDE systems yielded similar results
- Bridges did not have a base image (Shout out to Roberto and PSC admins for creating one for me)
- Even with a base image **we hit a showstopping issue: permissions**.
 - We could not tweak the systems software or rebuild the image
 - Required root-level permission, which Singularity does not allow

IMPORTANT: This is ***not*** a negative reflection on Bridges, because I believe I would have had this same experience on any XSEDE machine. I just happen to use Bridges heavily so it was the best testbed for me.

Early User Experience on Bridges

That said ...

XSEDE and PSC support for helping me explore Singularity was phenomenal

PSC and Bridges folks responded to my request immediately and helped me consider my options

It's just that some users may not know to ask...

Build and Configure Containers On-Site

For containers to be successful on XSEDE

I believe users *must* be able to build and configure them directly on the supercomputer

- Tedious to copy, tweak, rebuild, and copy back
- Optimizations are difficult if hardware is different


Containers currently work for stable software (read traditional HPC community), but are unworkable for rapidly evolving software (read many non-traditional HPC communities). I argue containers should best serve these new communities – it should make their lives easier.

Paleoscape Model

Paleoscape Model and Human Origins

Simulate Climate and
Vegetation during the
Last Glacial Maximum
(~140,000 years ago)

Altered coastline,
climate, flora,
and fauna



Republic of
South Africa

Simulate Humans
Using Agent-based
Models




Paleoscape

Coupled model approach

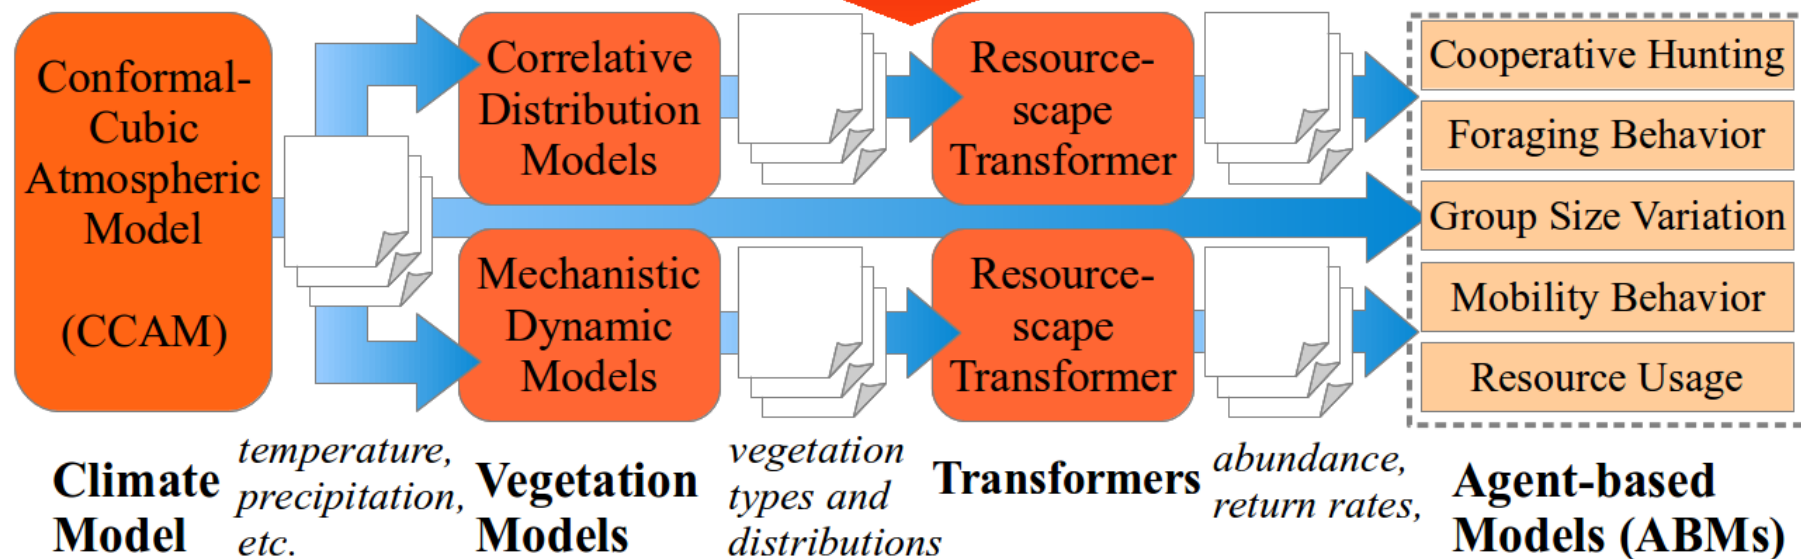
XSEDE
Supercomputers



Climate States



Computational Workflow



(Shook, et al. 2015) won “Best Accelerating Discovery Paper” at XSEDE’15

The Challenge for Many Users: Terminals

```
Using keyboard-interactive authentication.
XSEDE Authentication
password:
Last login: Mon Jun 26 14:08:51 2017 from 196.34.85.98
***** W A R N I N G *****
You have connected to br005.pvt.bridges.psc.edu

This computing resource is the property of the Pittsburgh Supercomputing Center.

It is for authorized use only. By using this system, all users acknowledge
notice of, and agree to comply with, PSC policies including the Resource Use
Policy, available at http://www.psc.edu/index.php/policies. Unauthorized or
improper use of this system may result in administrative disciplinary action,
civil charges/criminal penalties, and/or other sanctions as set forth in PSC
policies. By continuing to use this system you indicate your awareness of and
consent to these terms and conditions of use.

LOG OFF IMMEDIATELY if you do not agree to the conditions stated in this warning

Please contact remarks@psc.edu with any comments/concerns.

***** W A R N I N G *****
[shook@br005 ~]$ ls
files      paleoscape-model  slurmscript.sh      tocopy.zip  workflow
netlogo    pylon-paleoscape  test-submit-scripts wdir_FGOALS
[shook@br005 ~]$ module list
Currently Loaded Modulefiles:
  1) psc_path/1.1      2) slurm/17.02.5    3) intel/17.4
[shook@br005 ~]$
```

Command-line
Interface

Batch Queuing
System

Split architecture:
Head Nodes and
Compute Nodes

...

High learning curve

Bridges supercomputer at the Pittsburgh Supercomputing Center



My Quest

Use Cases

GISandbox

Considerations

Science Gateways

Science gateways allow science & engineering communities to access shared data, software, computing services, instruments, educational materials, and other resources specific to their disciplines. (sciencegateways.org)

- Lower the barrier to entry for science disciplines
- Common platform for collaborative science
- Scientists use the allocation for the science gateway so there is no need to write an allocation proposal

GISandbox

Play place for researchers and educators to learn about, experiment with, and advance geographic information systems and science (gisandbox.org)

The screenshot displays the Jupyter Notebook interface within the GISandbox environment. At the top, the header reads 'jupyter Notebook Demo' with a 'Last Checkpoint: 4 minutes ago (unsaved changes)' status. The top navigation bar includes 'File', 'Edit', 'View', 'Insert', 'Cell', 'Kernel', 'Widgets', and 'Help'. A toolbar below the navigation bar contains icons for saving, running, and other notebook functions. The main content area shows a 'Text cell' with the text: 'Welcome to Jupyter Notebooks on the GISandbox' and 'This example notebook was written by Eric Shook for the GISandbox. It loads the Cesium visualization package and displays it for interactive visualization.' Below the text cell is a 'Code cell' containing the following Python code:

```
In [6]: import dojo
import cesiumpy

v = cesiumpy.CesiumWidget()
v
```

 The output of the code cell is an '(Interactive) Output' showing a 3D globe of Earth, centered on North America, rendered using the Cesium library. The bottom of the interface shows an empty input field for the next code cell, labeled 'In []:'.

GISandbox User Interface (Jupyter Notebooks)


GISandbox Architecture (10,000 foot view)



While the rest of the code might be the same:

```
In [22]: fig = plt.figure(figsize=(15,7))
#m.fillcontinents(color='gray',lake_color='gray')
m.drawcoastlines()
m.drawparallels(np.arange(-80.,81.,20.))
m.drawmeridians(np.arange(-180.,181.,20.))
m.drawmapboundary(fill_color='white')
cs = m.contourf(x,y,air_c[0,:],20)
plt.title('Monthly mean SAT')
```


Out[22]: <matplotlib.text.Text at 0x1116bd210>



Jetstream
Cloud
Computing
Resource



Credits: https://github.com/koldunovn/python_for_geosciences



Comet supercomputer at the
San Diego Supercomputing Center

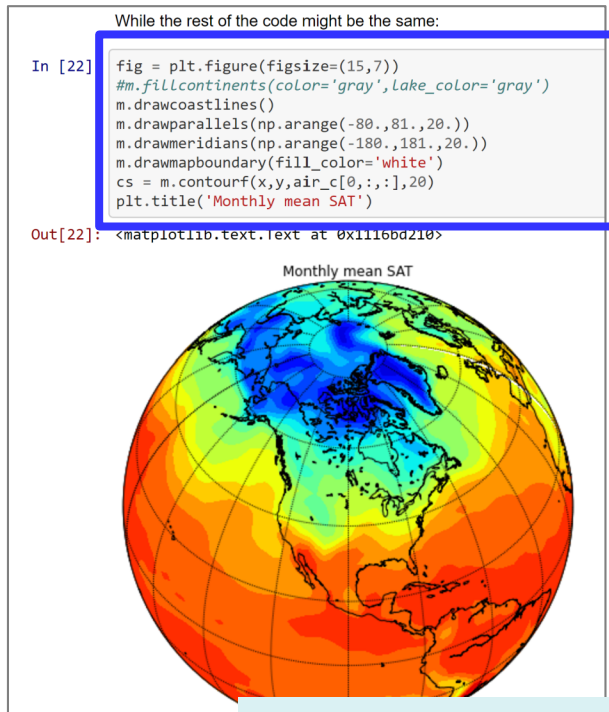
Source: <http://ucsdnews.ucsd.edu>



Bridges supercomputer at the
Pittsburgh Supercomputing Center

Source: <http://insidehpc.com>

GISandbox Architecture (10,000 foot view)



Jetstream
Cloud
Resource



Jupyter
“magic”
command to
run code cells
on Comet or
Bridges
supercomputers



**Comet supercomputer at the
San Diego Supercomputing Center**

Source: <http://ucsdnews.ucsd.edu>



**Bridges supercomputer at the
Pittsburgh Supercomputing Center**

Source: <http://insidehpc.com>

GISandbox Architecture (10,000 foot view)



While the rest of the code might be the same:

```
In [22]: fig = plt.figure(figsize=(15,7))
#m.fillcontinents(color='gr',lake_color='gray')
m.drawcoastlines()
m.drawparallels(np.arange(-80.,81.,20.))
m.drawmeridians(np.arange(-180.,181.,20.))
m.drawmapboundary(fill_color='white')
cs = m.contourf(x,y,air_c[0,:,:],20)
plt.title('Monthly mean SAT')
```

Dedicated Virtual Machine

24 cores
60 GB RAM



Jetstream
Cloud
Resource



1,944 Nodes

24 cores per node
Standard memory: 128 GB
Large memory: 1.5 TB
GPU + Haswell processors

Comet supercomputer at the
San Diego Supercomputing Center
Source: <http://ucsdnews.ucsd.edu>

Hybrid System

Regular memory: 128 GB
Large memory: 12 TB or 3 TB
Supports DBs, Spark, and more

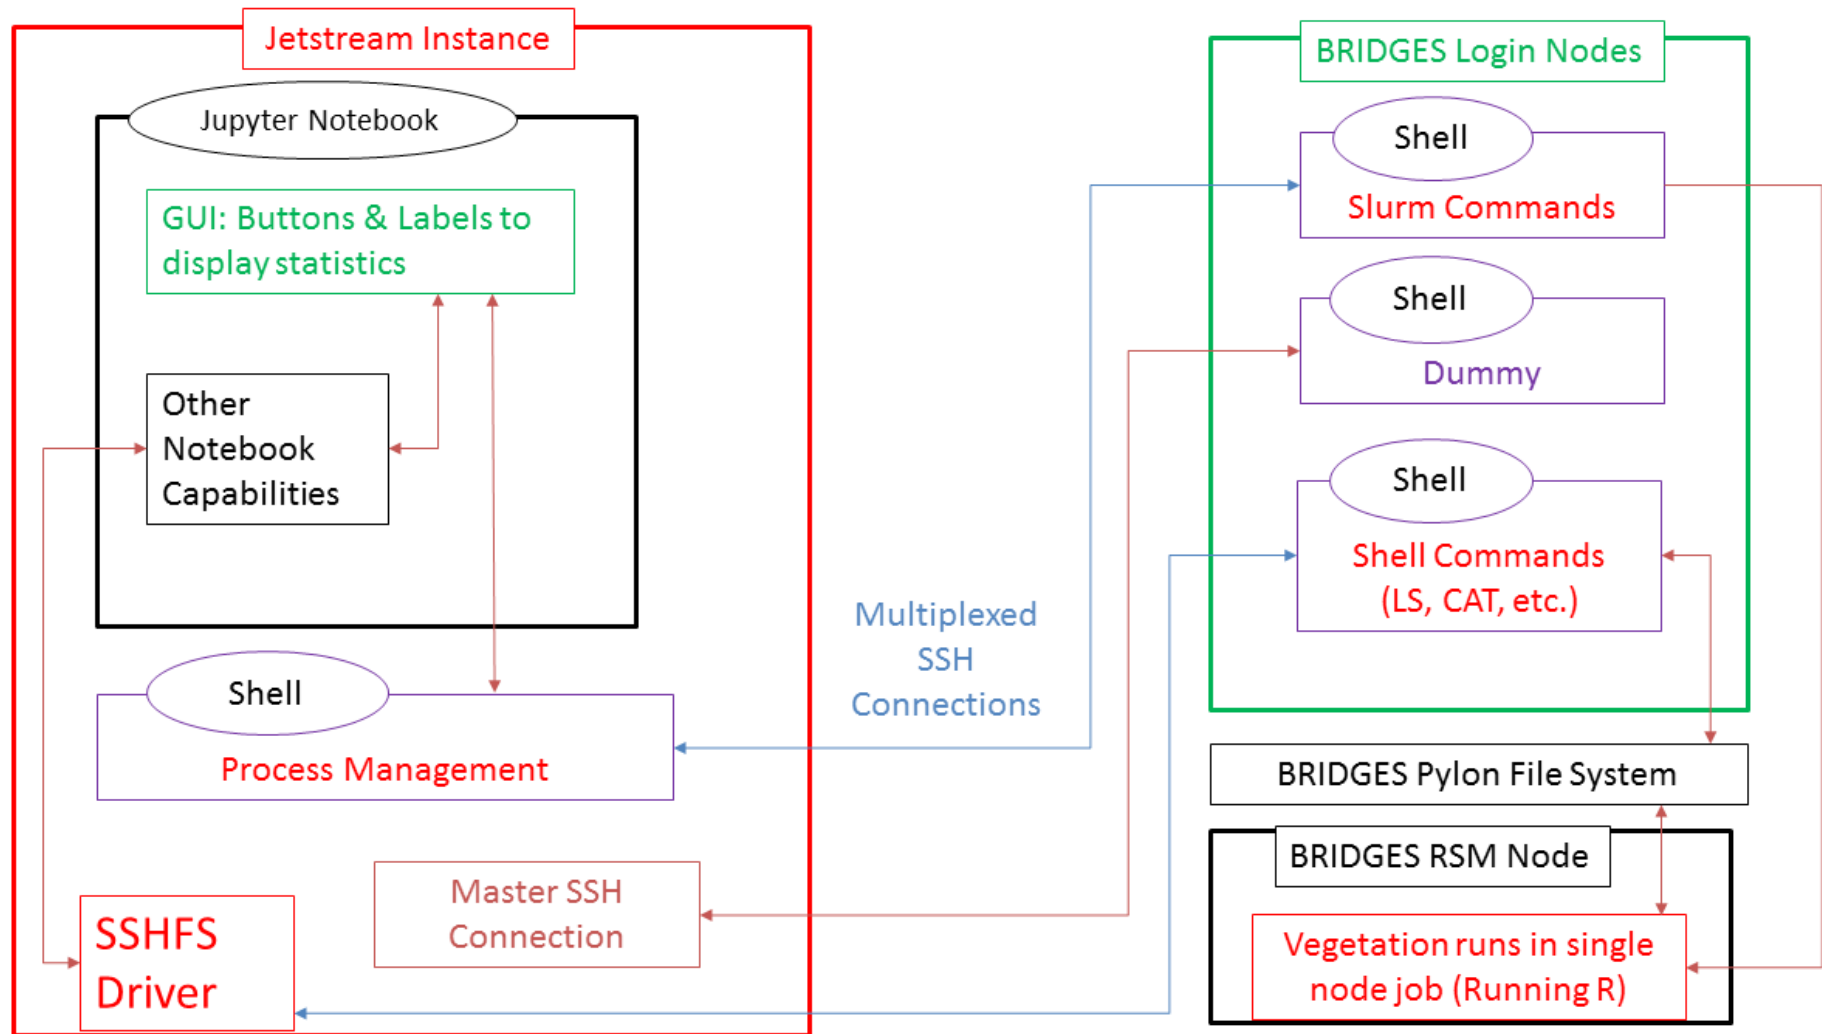
Bridges supercomputer at the
Pittsburgh Supercomputing Center
Source: <http://insidehpc.com>

GISandbox Behind-The-Scenes*

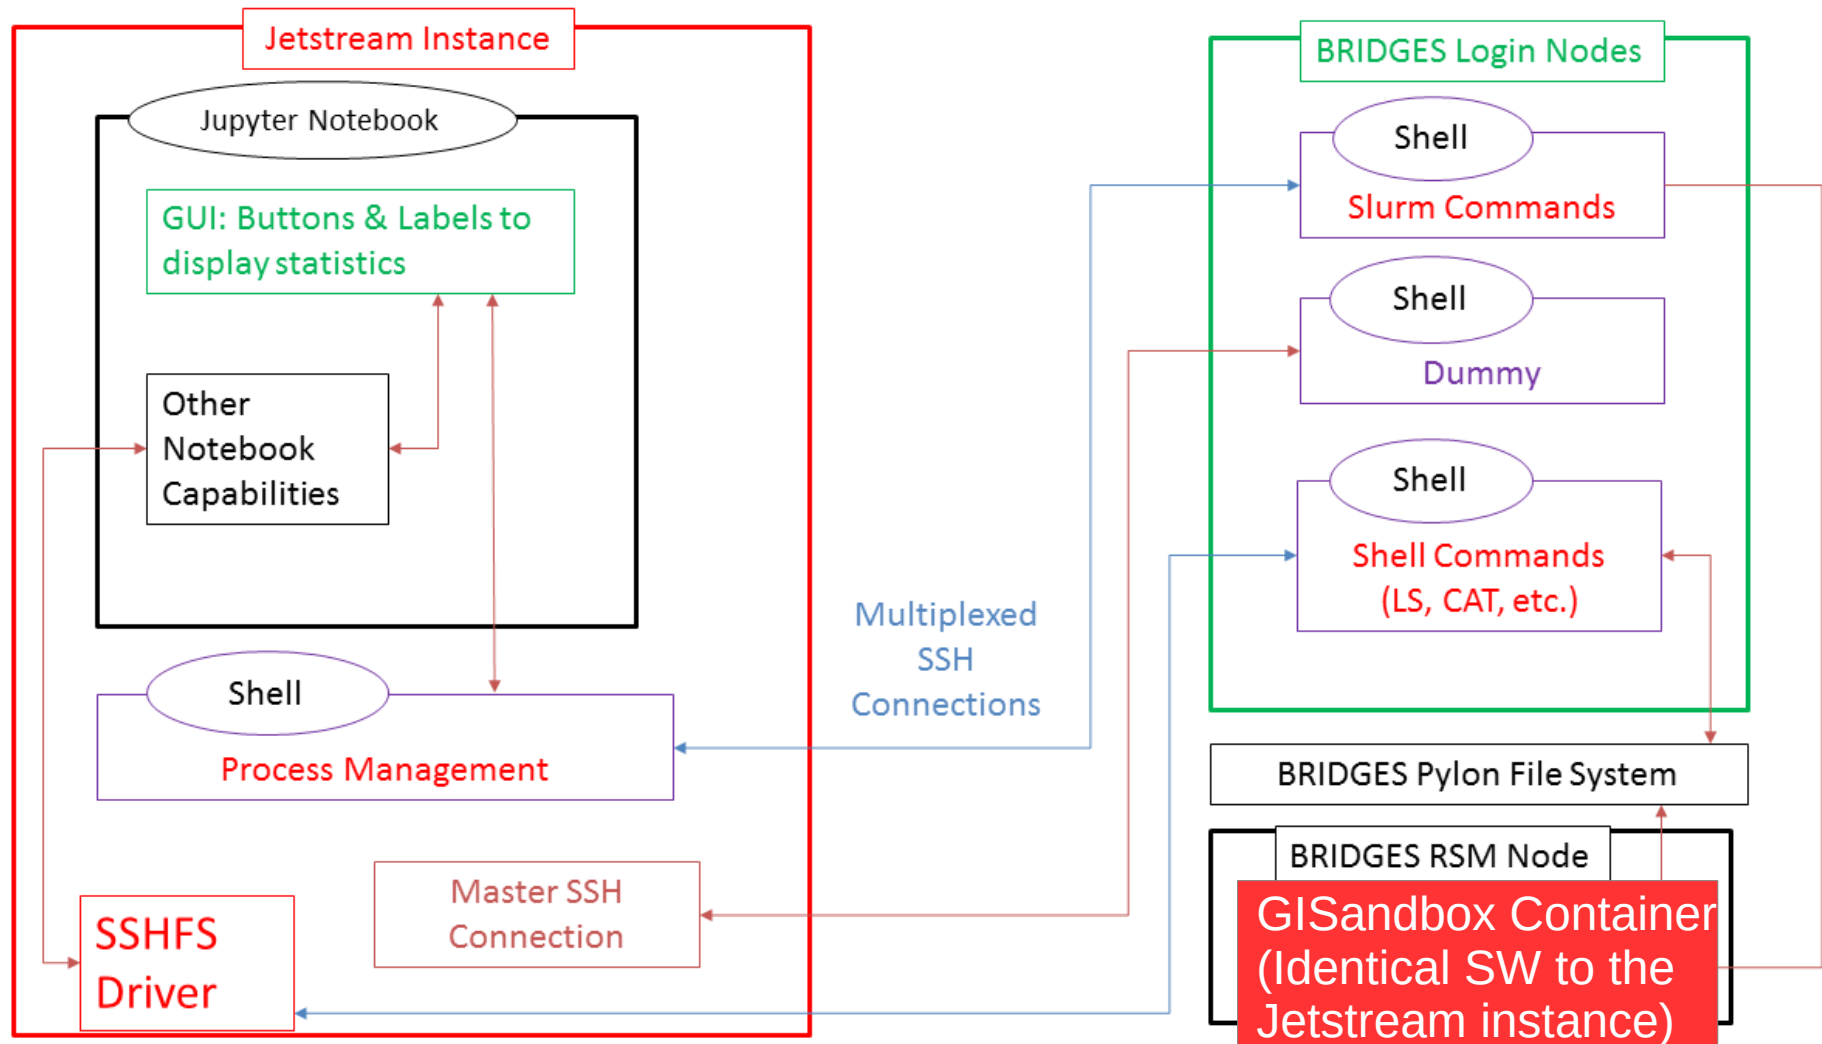
- Jetstream Virtual Machine
 - Ansible scripts to build the system and software
 - Ubuntu
 - JupyterHub + CILogon (to use XSEDE credentials for login)
 - Python 2, Python 3, R
 - GIS and other specific software
 - Jupyter Improvements
 - Code cell → Supercomputing job
- Bridges and Comet
 - CPU & Storage Allocations
 - GIS and other specific software will match the VM

* ECSS Davide Del Vento, Jun Wang, and now Andrea Zonca to the rescue!!

GISandbox Architecture (Details)



GISandbox Architecture (A Dream?)



Right now that is too much work, but it shouldn't be

Slide courtesy: Sergiu Sanielevici and Davide Del Vento



My Quest

Use Cases

GISandbox

Considerations

Considerations From a User Perspective

- Storage
 - Many XSEDE SPs have limited storage allocations
 - Containers may put pressure on XSEDE to increase storage for users and projects
- Reinventing the wheel
 - As I work through these difficulties I am encountering others who are doing the same thing.
 - A mechanism for users and SPs to share experiences (and code) would be helpful.

Considerations From a User Perspective

- Supercomputer DMZ
 - We need a demilitarized zone on each supercomputer to build and configure images.
 - Something like a supercomputer running a VM running a container for isolation.
 - (Systems hat) Security is next to impossible, but (Users hat) for containers to be truly useful we need to work toward figuring this out.
- Collaboration
 - Need to figure out how to share containers and data between containers/users

Thank You

Eric Shook
eshook@umn.edu

I want to thank Davide Del Vento, Jun Wang, Eroma Abeysinghe, Andrea Zonca, and many from PSC for all their help!