

Strategic Overview of Docker in Fifteen Minutes

for Middleware And Grid Interagency
Coordination (MAGIC)

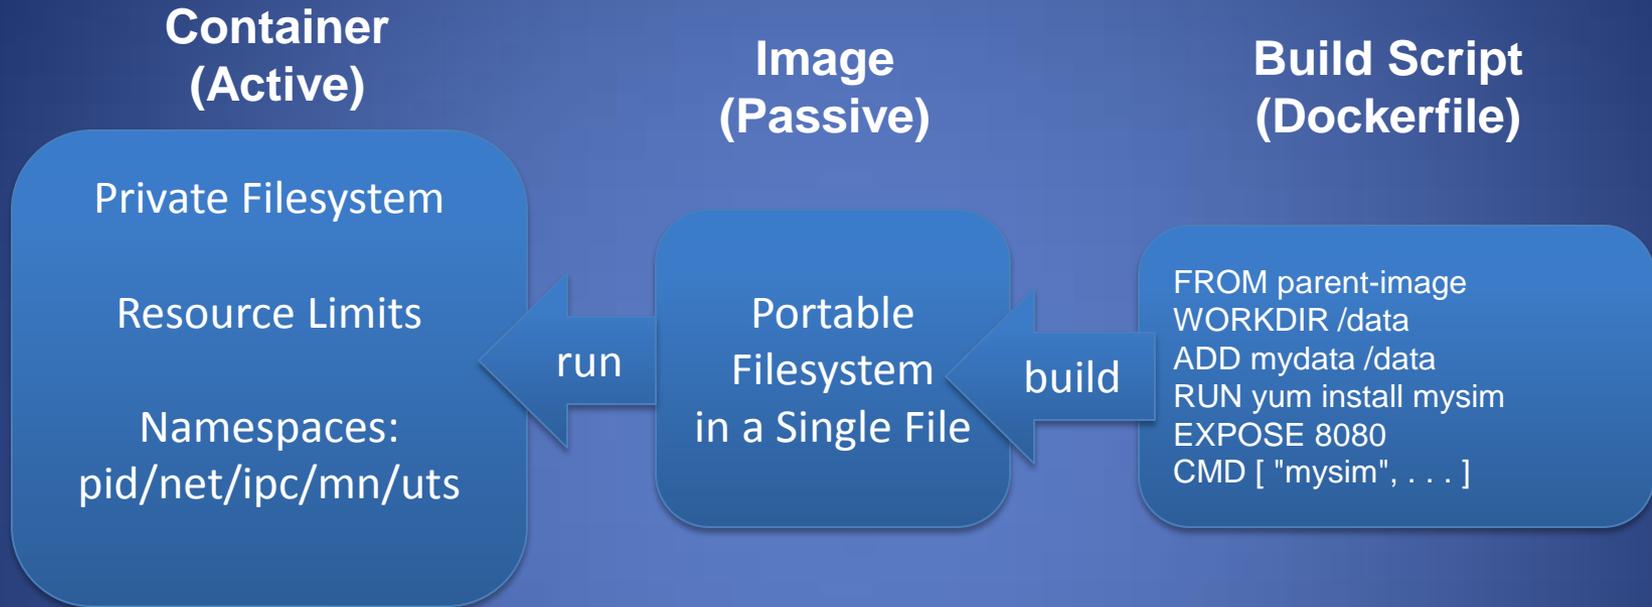
7 February 2018

Prof. Douglas Thain, University of Notre Dame



<http://www.nd.edu/~dthain>
dthain@nd.edu
@ProfThain

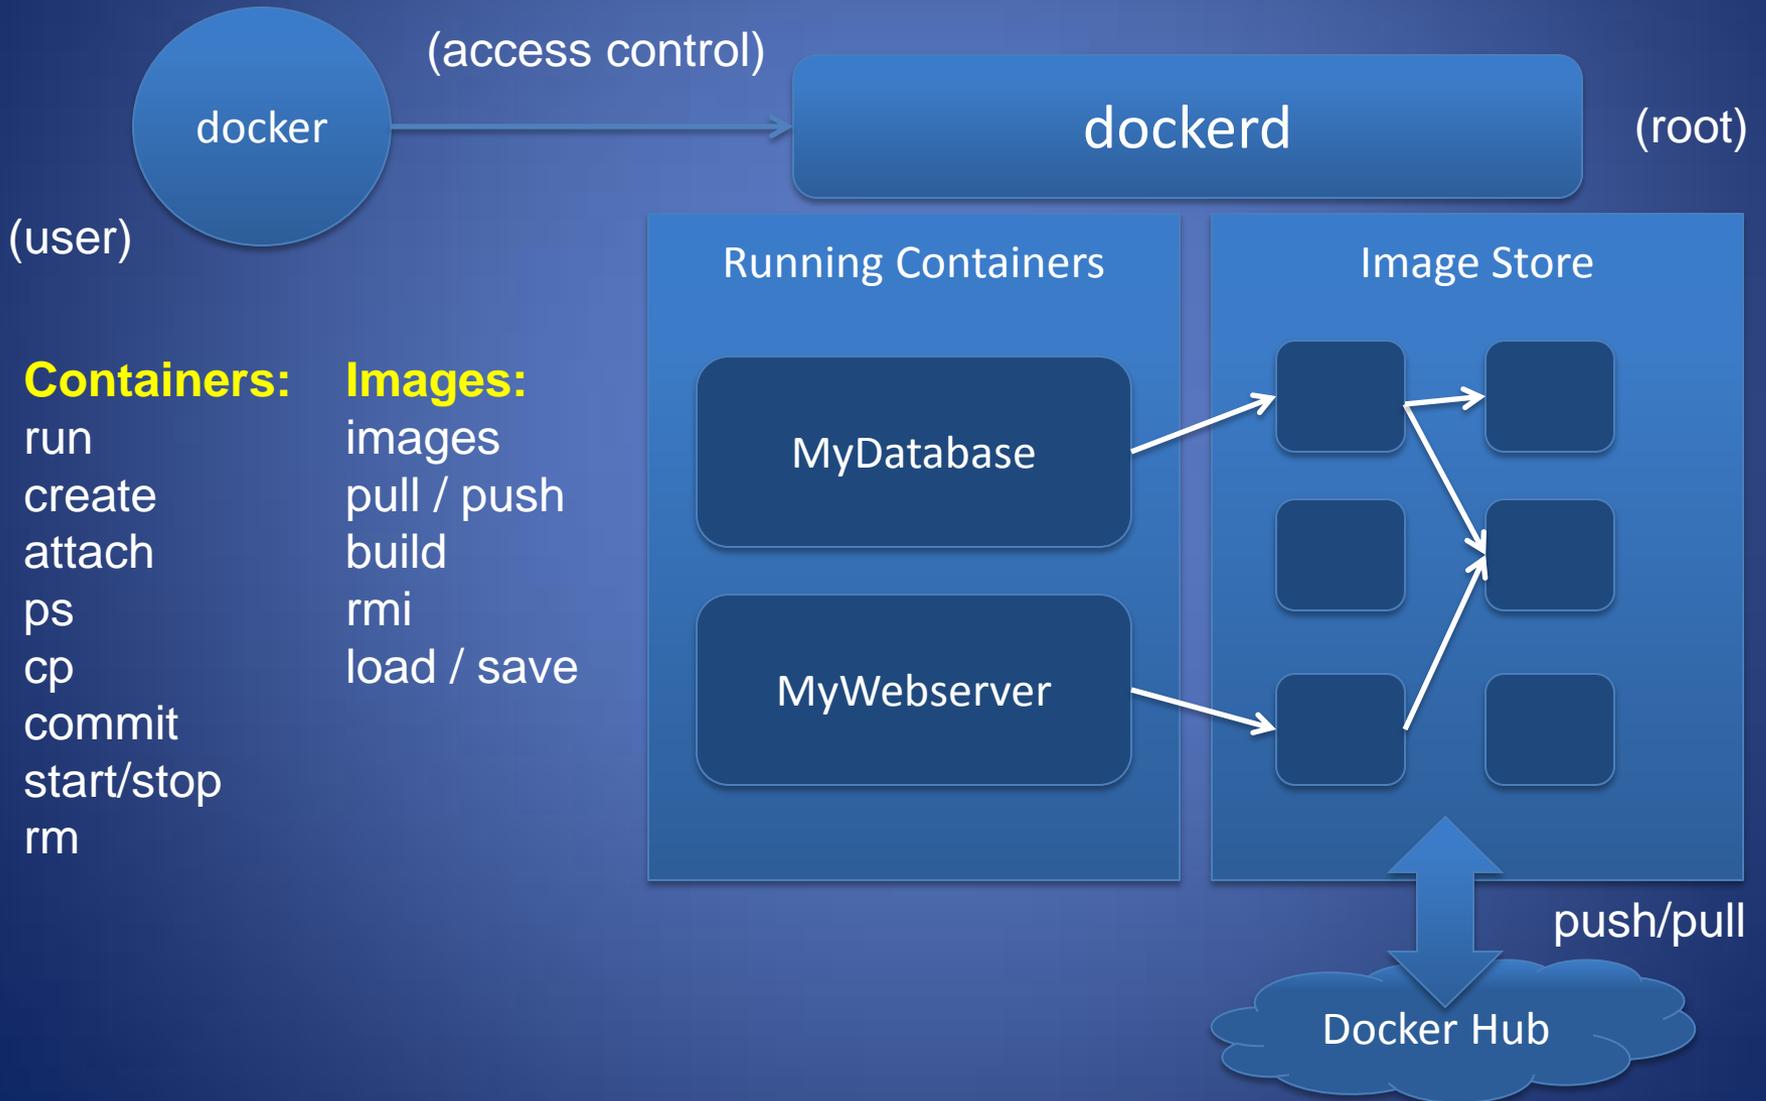
Container Overview



Purpose:

To provide independent applications with the ability to define private, customized environments that are easily distributable.

Docker Structure



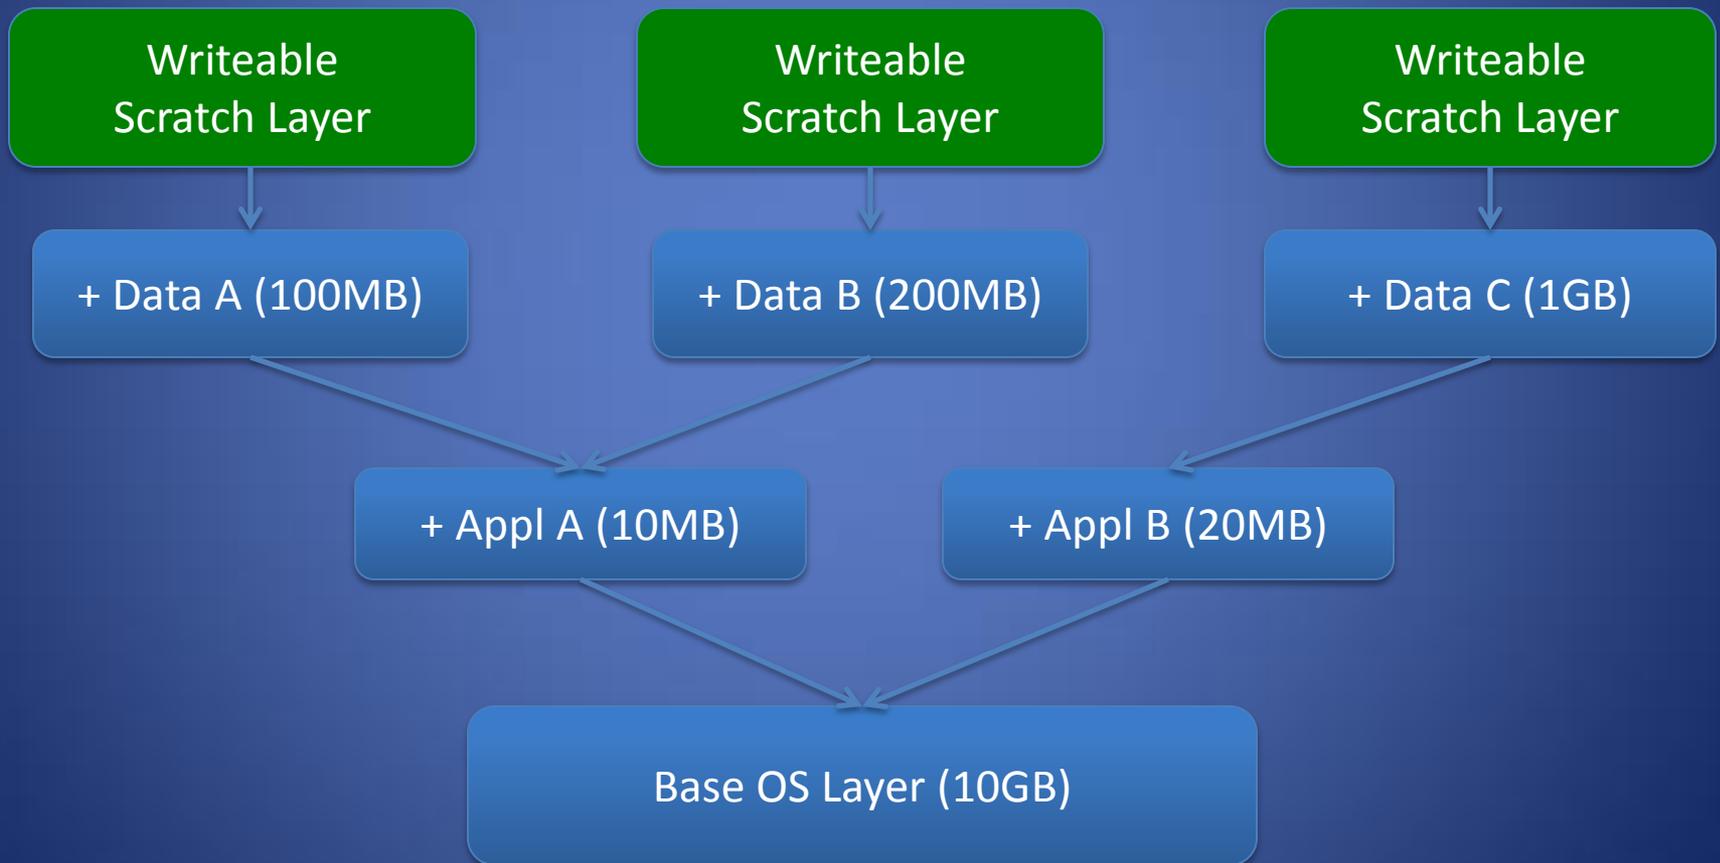
Container Sub-Structure



Dockerfile:

```
FROM rhel7.2
RUN yum install myexe
COPY mydata /data
```

Encourage De-Duplication



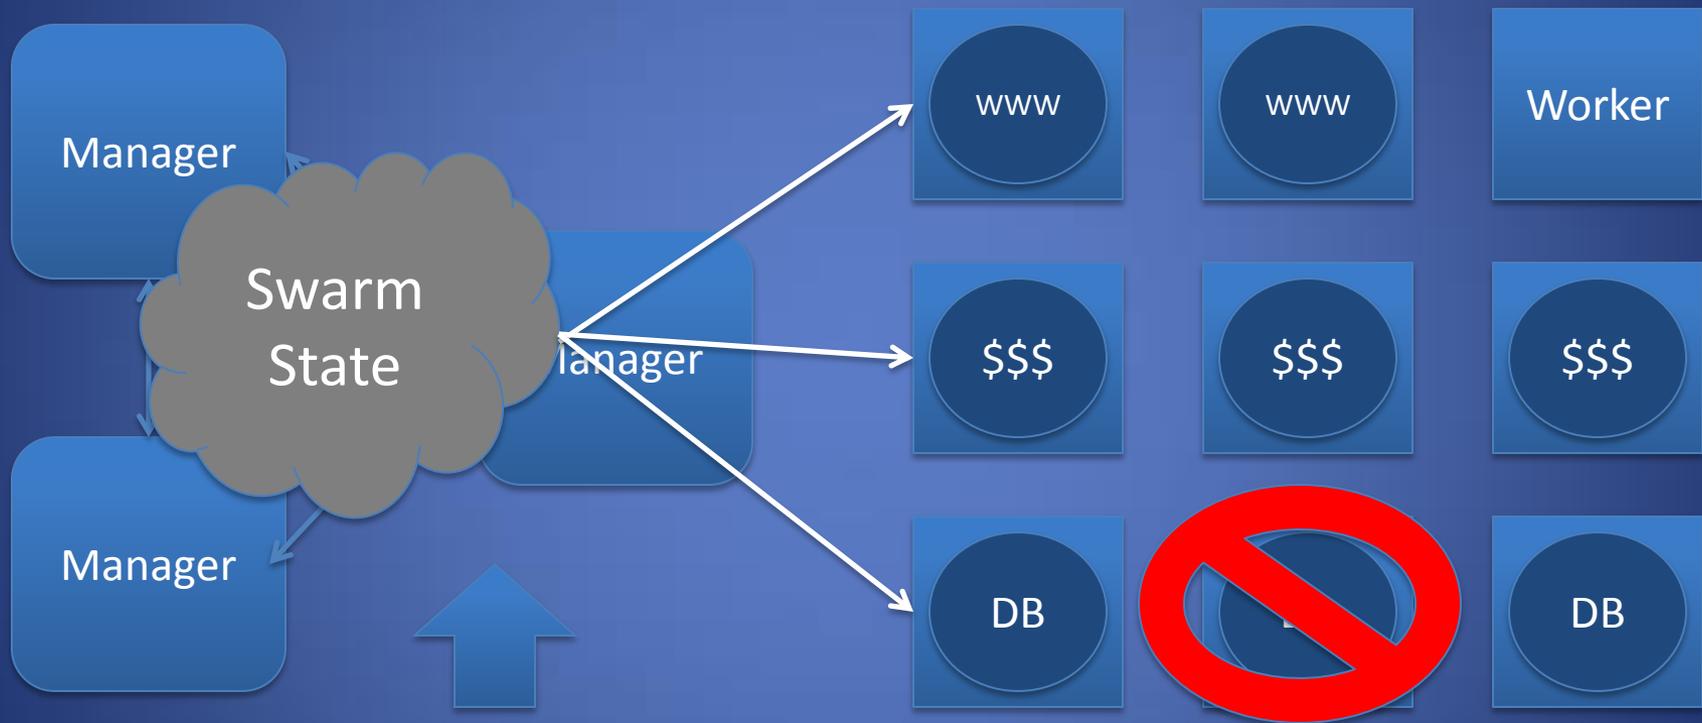
Docker Compose (Static)

- Common case: a web-service that consists of multiple concurrent services.



Persistent state is **outside** the cluster.
Requires external supervision to maintain.

Docker Swarm (Dynamic)



create service www (2)
create service memcached (3)
create service database (2)

Persistent state is **inside** the cluster.
Once created, self-maintaining.

Docker Ecosystem

- Docker containers are used widely in industry and have become an interchange format.
- Google GCE, Amazon ECS, Mesos, Kubernetes all support Docker container execution. (But not necessarily using Docker per se.)
- Running your own container cluster is rather challenging! One opinion:
<https://thehftguy.com/2016/11/01/docker-in-production-an-history-of-failure/>

Problem: Garbage Collection

- Problem: Any authenticated user can consume storage by create/pulling new images, and nothing removes them automatically. Any automated activity will eventually consume all storage and wedge docker.
- A cronjob seems to be the community solution:

```
docker images -q -a | xargs -no-run-if-empty docker rmi
```

- However, this doesn't take into account containers that are "needed" even if not currently run.
- Effective use requires one of several ideas:
 - Use a small number of containers widely.
 - Trust users to delete unneeded containers.
 - Make use cases robust to periodic removal.

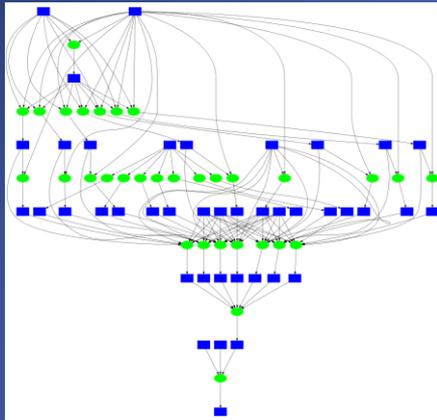
Problem: Adding Docker to HPC

- Idea sounds simple: You already have a batch system, just put Docker on every node.
- **Very** challenging to get this right:
 - Docker has very specific kernel requirements, consumes large amounts of local space, resources conflict with each other, independent scheduling.
- Underlying challenge: Docker really wants to be your core resource manager. It's fighting with the ones you already have: batch system, parallel file system..
- (Shifter, Singularity, CharlieCloud, and others are working to address this space directly.)

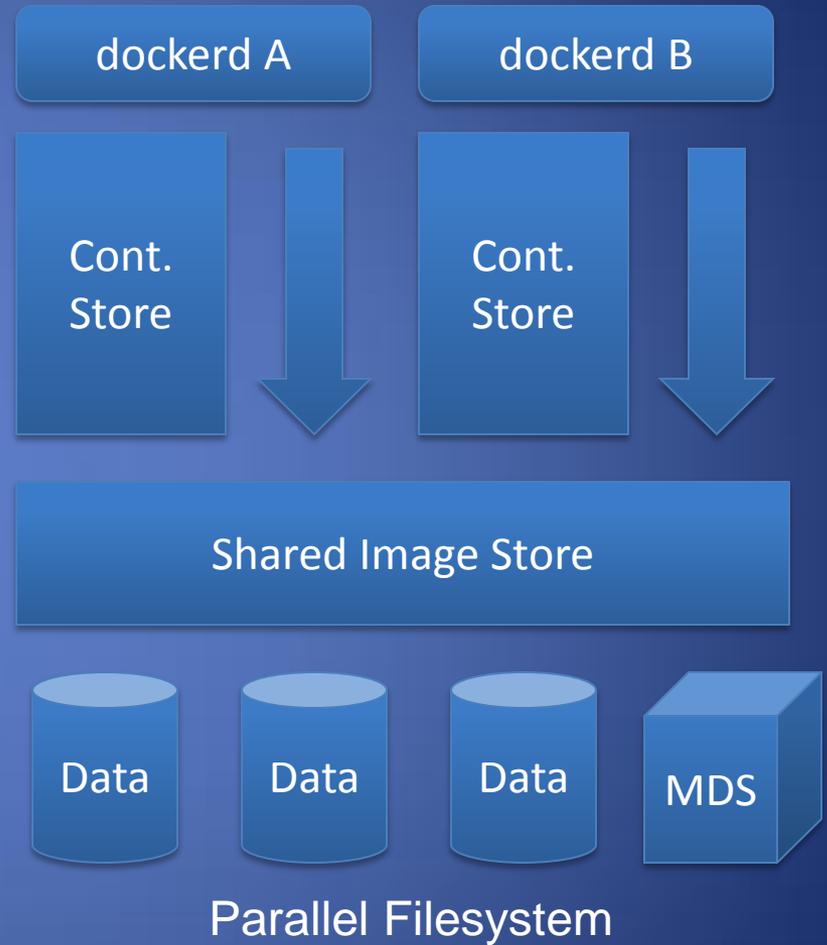
Takeaways

- Designed for deploying micro services, which is different than executing HTC/HPC jobs.
- Docker container format is now standard.
- Docker works very well within its own ecosystem, but is quite difficult to integrate into existing facilities.
- Garbage collection remains unsolved prob.
- Very sensitive to underlying kernel tech.

Container Portability In Makeflow



Wharf: Containers on Shared Filesystems



o ccl.cse.nd.edu

The Cooperative Computing Lab

Software | Download | Manuals | Papers

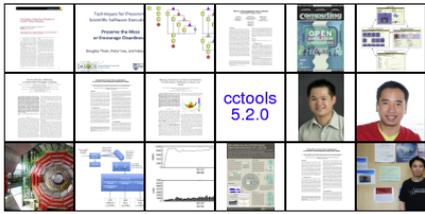
Take the [ACIC 2015 Tutorial](#) on Makeflow and Work Queue

About the CCL

We design [software](#) that enables our [collaborators](#) to easily harness [large scale distributed systems](#) such as clusters, clouds, and grids. We perform fundamental [computer science research](#) in that enables new discoveries through computing in fields such as physics, chemistry, bioinformatics, biometrics, and data mining.

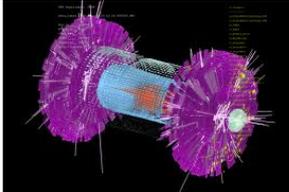
CCL News and Blog

- [Global Filesystems Paper in IEEE CISE \(09 Nov 2015\)](#)
- [Preservation Talk at iPres 2015 \(03 Nov 2015\)](#)
- [CMS Case Study Paper at CHEP \(20 Oct 2015\)](#)
- [OpenMalaria Preservation with Umbrella \(19 Oct 2015\)](#)
- [DAGVz Paper at Visual Performance Analysis Workshop \(13 Oct 2015\)](#)
- [Virtual Wind Tunnel in IEEE CISE \(09 Sep 2015\)](#)
- [Three Papers at IEEE Cluster in Chicago \(07 Sep 2015\)](#)
- [CCTools 5.2.0 released. \(19 Aug 2015\)](#)
- [Recent CCL Grads Take Faculty Positions \(18 Aug 2015\)](#)
- [\(more news\)](#)



Community Highlight

Scientists searching for the Higgs boson have profited from Parrot's new support for the [CernVM Filesystem \(CVMFS\)](#), a network filesystem tailored to providing world-wide access to software installations. By using [Parrot](#), CVMFS, and additional components integrated by the [Any Data. Anytime. Anywhere](#) project, physicists working in the [Compact Muon Solenoid](#) experiment have been able to create a uniform computing environment across participating institutions. Instead of maintaining large software repositories, Parrot is used to download files as needed and efficiently. A pilot project at the University of Wisconsin demonstrated the feasibility of this approach by harnessing 370,000 CPU-hours across access to 400 gigabytes of software in a repository.



- Dan Bradley, University of Wisconsin

http://ccl.cse.nd.edu

@ProfThain

Twitter, Inc.

Search Twitter Have an account? Log in



Douglas Thain @ProfThain

Distributed computing for big data problems in science and engineering.

Notre Dame, IN
nd.edu/~dthain

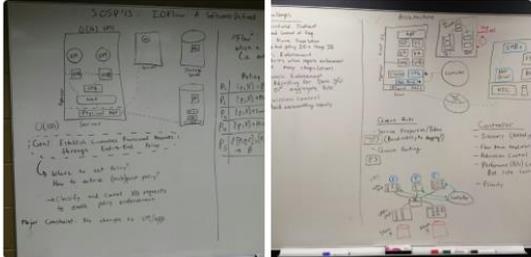
13 Photos and videos



Tweets Tweets & replies Photos & videos

Douglas Thain @ProfThain · Nov 10

My grad students now summarize research papers by preparing a whiteboard in advance. Much better than a slide deck!



New to Twitter?

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development
(NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674,
Fax: 202-459-9673, Email: nco@nitrd.gov, Website: <https://www.nitrd.gov>

