VC3

A Virtual Cluster Service for Community Computation

Lincoln Bryant, Jeremy Van, Benedikt Riedel, Robert Gardner,
Jose Caballero Bejar, John Hover, Ben Tovar, and Douglas Thain

http://www.virtualclusters.org

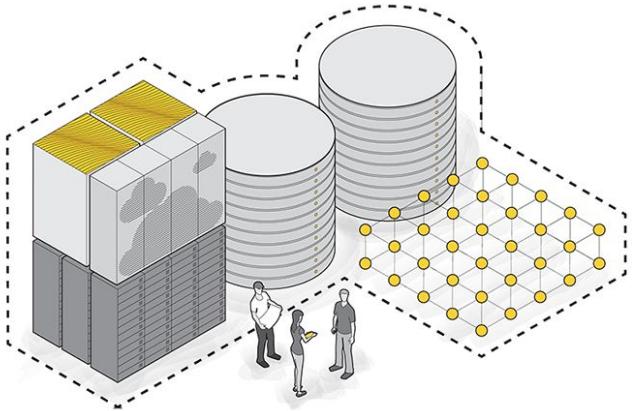

VC3: A platform for provisioning cluster frameworks over heterogenous resources for collaborative science teams

You have developed a complex workload which runs successfully at one site, perhaps your home university.

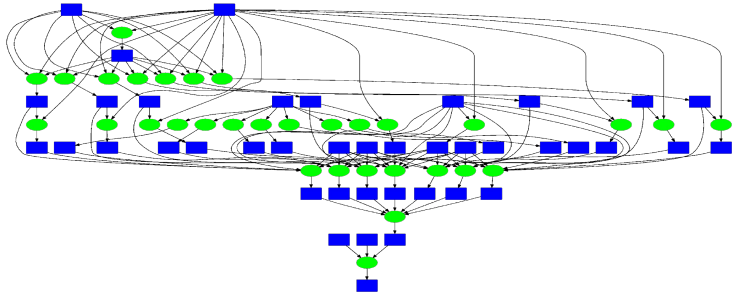

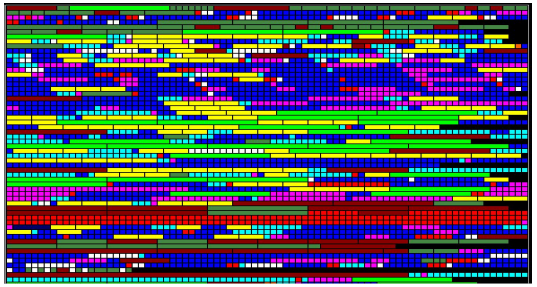

Now, you want to migrate and expand that application to national-scale infrastructure.
And allow others to easily access and run similar workloads.

Leadership HPC Facility

Distributed Computing Facility

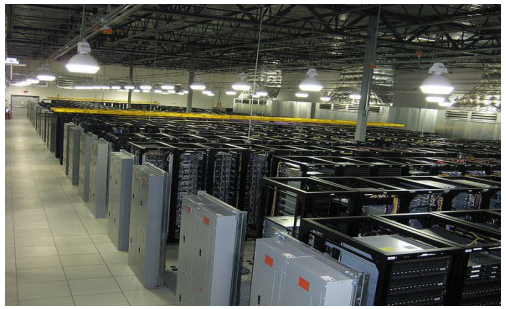

Commercial Cloud

Concept: Virtual Cluster

200 nodes of 24 cores and 64GB RAM
100GB local disk per node
CMS software 8.1.3 and python 2.7
CentOS 6 Operating System
Running the HTCondor batch system.
Shared between Alice, Bill, and Charles.

Virtual Cluster Service

Virtual Cluster Head Node

Virtual Cluster Head Node

Virtual Cluster

Deploy Software

Deploy Software

Deploy Software

Leadership HPC Facility

Distributed Computing Facility

Commercial Cloud

Concept: Virtual Cluster

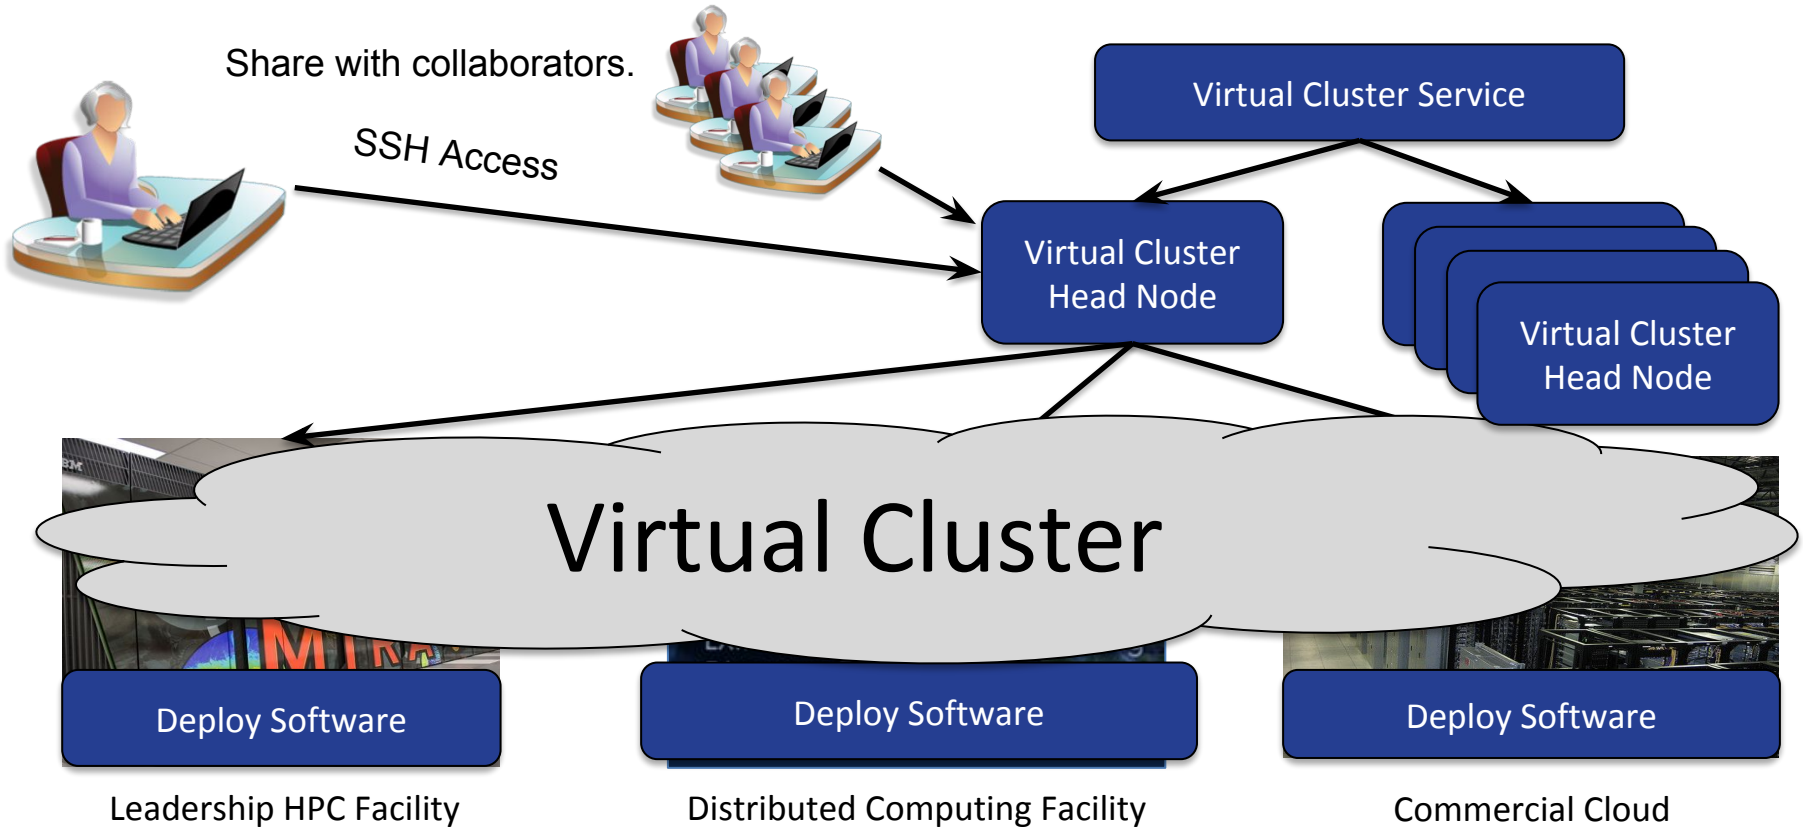

VC3: Virtual Clusters for Community Computation

- **VC3 is an interactive service for creating/sharing/using virtual clusters.**
- A virtual cluster consists of:
  - 1 x head node for interactive access to the cluster.
  - N x worker nodes for executing your workload.
  - Middleware to manage the cluster. (HTCondor, Makeflow, Spark, …)
  - Application software to do real work. (BLAST, Python, etc…)
- A virtual cluster is created using:
  - Your standard accounts/credentials on existing facilities.
  - Plain ssh/qsub access on each facility.
  - Container technology (if available) or user-level software builds (otherwise).
  - (No special privileges or admin access required on the facility.)

Limited Beta Release

Current status: **limited beta release!** We are **looking** for **collaborators** to help us work through bugs and offer feedback.

http://virtualclusters.org

If you have an ambitious goal and feel VC3 may help you get there, please fill out the form below and we'll send an invite:

http://bit.ly/vc3-signup

www.virtualclusters.org

login with a vc3 account

Globus Authentication

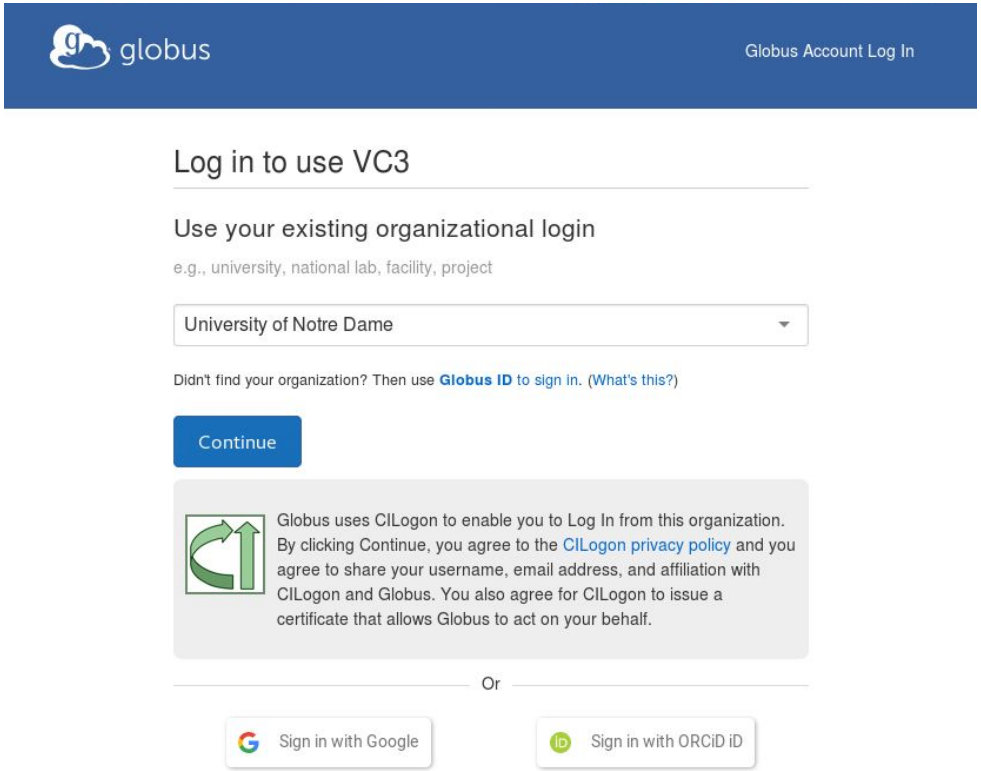

Curated Resources

Resource Profiles

Filter

Name	Organization	Description	Cores	Memory	Storage	Native OS	Features
Cori	National Energy Research Scientific Computing Center (NERSC)	Cori Supercomputer at NERSC	32	4000 MB	10000 MB	suse:v12	Shifter
MWT2	Midwest Tier 2	ATLAS Midwest Tier 2 Center job gateway (UChicago)	4	1000 MB	1000 MB	scientificlinux:v6.9	N/A
Midway	University of Chicago Research Computing Center (RCC)	Midway cluster at the University of Chicago Research Computing Center (RCC)	64	4000 MB	10000 MB	scientificlinux:v6.7	N/A
Stampede 2	Texas Advanced Computing Center (TACC)	Stampede 2 Super Computer	96	2000 MB	10000 MB	centos:v7.4	Singularity
CoreOS	University of Chicago	CoreOS Cluster	4	1000 MB	1000 MB	scientificlinux:v6.9	Singularity
UCT3	University of Chicago	UChicago ATLAS Tier 3	4	1000 MB	1000 MB	scientificlinux:v6.9	N/A
ND CCL	University of Notre Dame Cooperative Computing Lab	ND-CCL login none	4	1000 MB	10000 MB	redhat:v7	Singularity
Bridges	Pittsburgh Supercomputing Center	Bridges Supercomputer at PSC	28	4000 MB	35000 MB	centos:v7.3	Singularity
VC3 Test Pool	VC3	VC3 Test Pool	4	1000 MB	1000 MB	centos:v6.9	N/A
UCLA Hoffman2	University of California, Los Angeles	UCLA Hoffman2	8	1000 MB	10000 MB	centos:v6.9	N/A
OSG Connect	Open Science Grid	Open Science Grid (SL7)	4	1000 MB	1000 MB	Unknown	N/A

Allocations

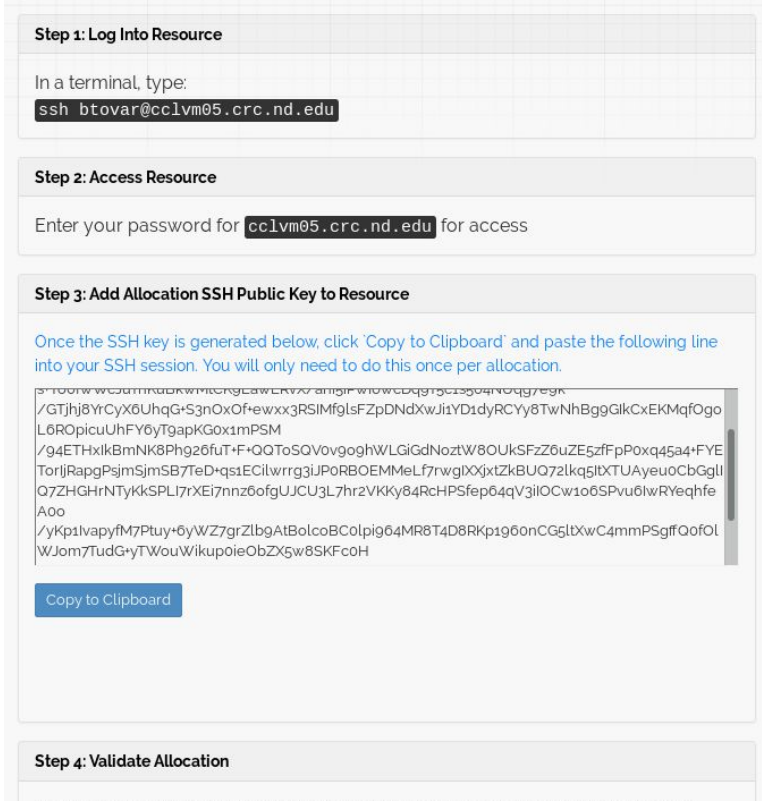

Projects

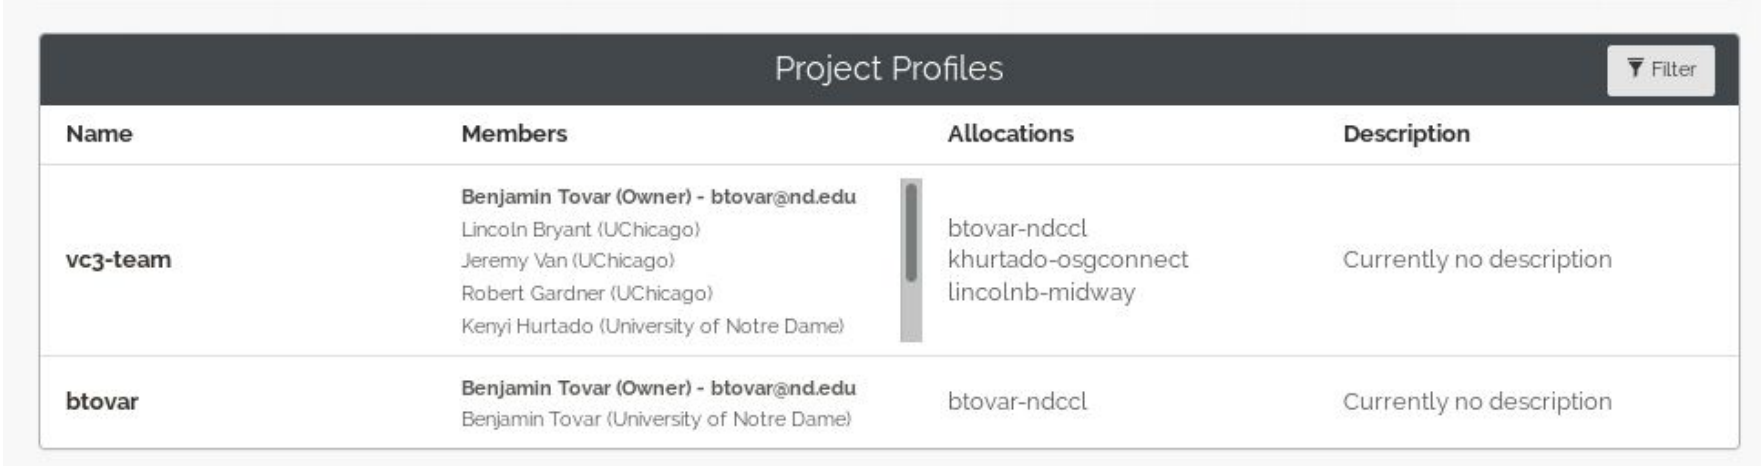
Project Profiles
Filter
Name
Members
Allocations
Description
vc3-team
Benjamin Tovar (Owner) - btovar@nd.edu
Lincoln Bryant (UChicago)
Jeremy Van (UChicago)
Robert Gardner (UChicago)
Kenyi Hurtado (University of Notre Dame)
btovar-ndccl
khurtado-osgconnect
lincolnb-midway
Currently no description
btovar
Benjamin Tovar (Owner) - btovar@nd.edu
Benjamin Tovar (University of Notre Dame)
btovar-ndccl
Currently no description

Launching a Virtual Cluster

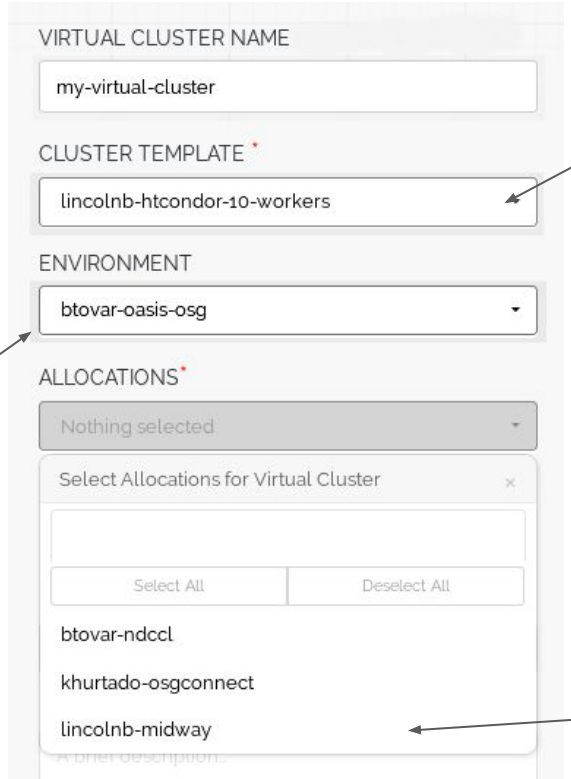

shared cluster definition

workers will have this environment installed

allocations available in this project

Cluster Status

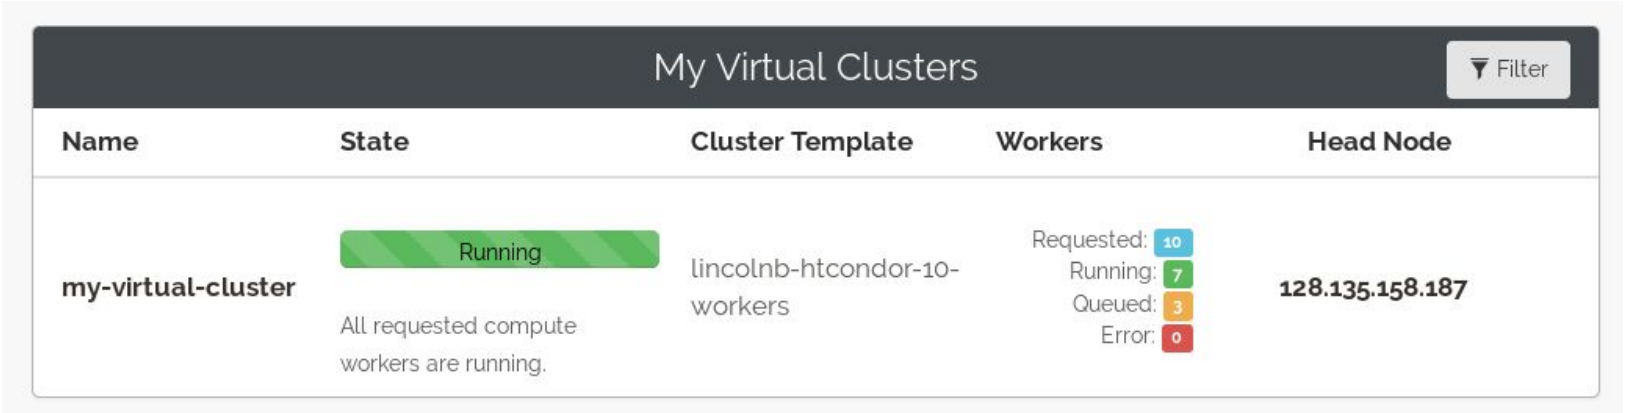

Workers from many sites

```
[btovar@btovar-my-virtual-cluster ~]$ ip addr | grep 128.135.158.187
    inet 128.135.158.187/25 brd 128.135.158.255 scope global dynamic eth0
[btovar@btovar-my-virtual-cluster ~]$ condor_status
Name                                     OpSys      Arch   State     Activity LoadAv Mem   ActvtyTime

slot1@glidein_21791@camd01.crc.nd.edu    LINUX      X86_64 Unclaimed Idle      5.120 4013  0+00:19:37
slot1@glidein_29106@camd01.crc.nd.edu    LINUX      X86_64 Unclaimed Idle      5.120 4013  0+00:19:37
slot1@glidein_91802@camd05.crc.nd.edu    LINUX      X86_64 Unclaimed Idle      5.260 4013  0+00:19:37
slot1@glidein_39133@iut2-c257.iu.edu     LINUX      X86_64 Unclaimed Idle     34.620 3223  0+00:19:48
slot1@glidein_61297@lnxfarm275.colorado.edu LINUX   X86_64 Unclaimed Idle      6.990 3002  0+00:14:36
slot1@glidein_28373@midway091.rcc.local  LINUX      X86_64 Unclaimed Idle      8.170 2013  0+00:19:36
slot1@glidein_71179@midway098.rcc.local  LINUX      X86_64 Unclaimed Idle      7.480 2013  0+00:19:36
slot1@glidein_46364@midway260.rcc.local  LINUX      X86_64 Unclaimed Idle      8.040 2013  0+00:19:35
slot1@glidein_39282@midway324.rcc.local  LINUX      X86_64 Unclaimed Idle      8.750 2013  0+00:19:36
slot1@glidein_39133@uct2-c373.mwt2.org   LINUX      X86_64 Unclaimed Idle     34.080 2415  0+00:19:33

                     Machines Owner Claimed Unclaimed Matched Preempting  Drain

            X86_64/LINUX       10     0       0        10       0          0      0

                   Total       10     0       0        10       0          0      0
[btovar@btovar-my-virtual-cluster ~]$
```

ND

OSG

UChicago

VC3 Monitoring

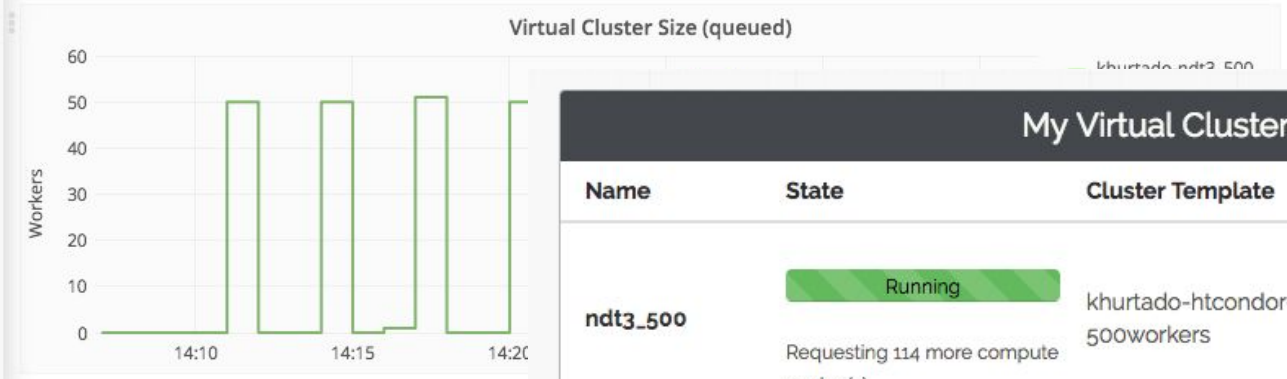

My Virtual Clusters

Filter

Name	State	Cluster Template	Workers	Head Node
ndt3_500	Running Requesting 114 more compute worker(s).	khurtado-htcondor-500workers	Requested: 500 Running: 386 Queued: 48 Error: 0	128.135.158.178

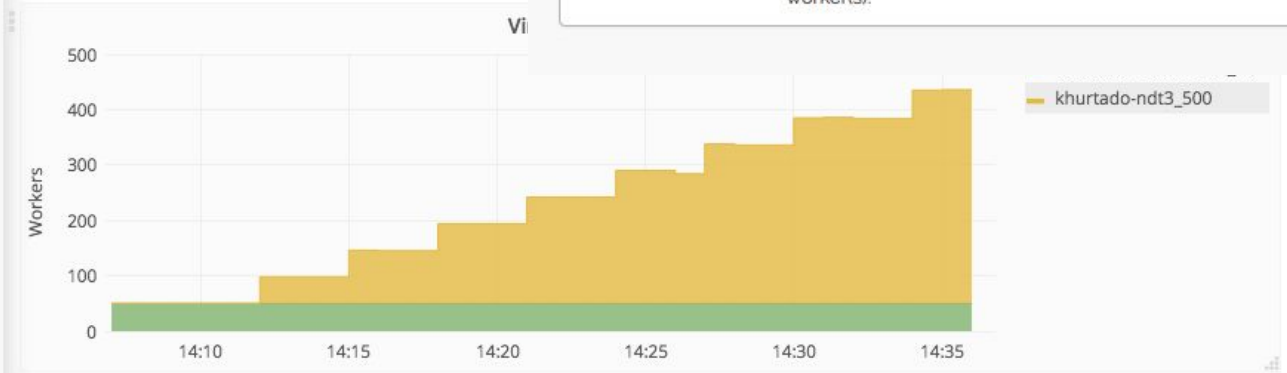

Details: System Monitoring

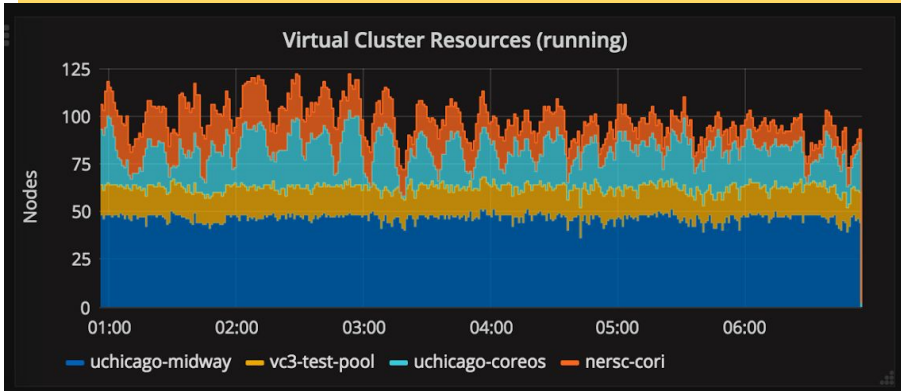

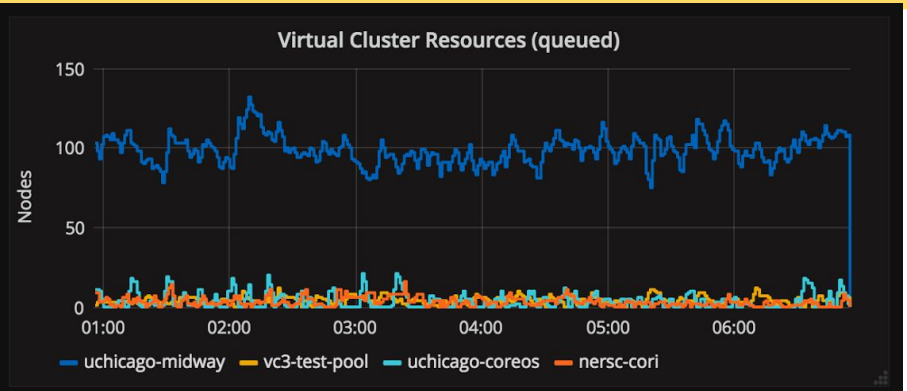

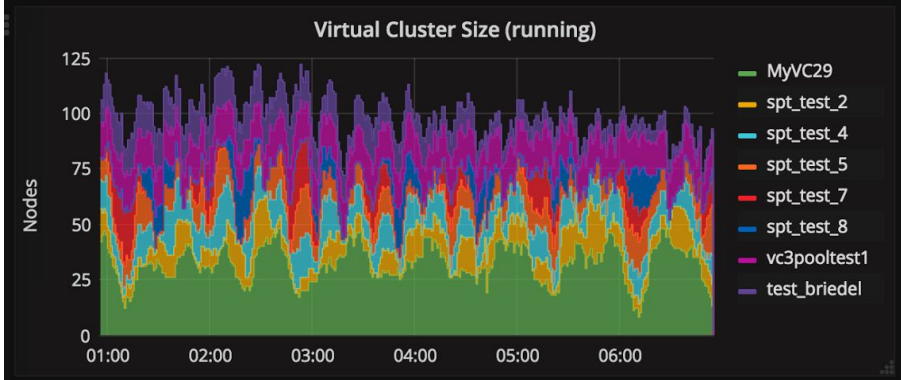

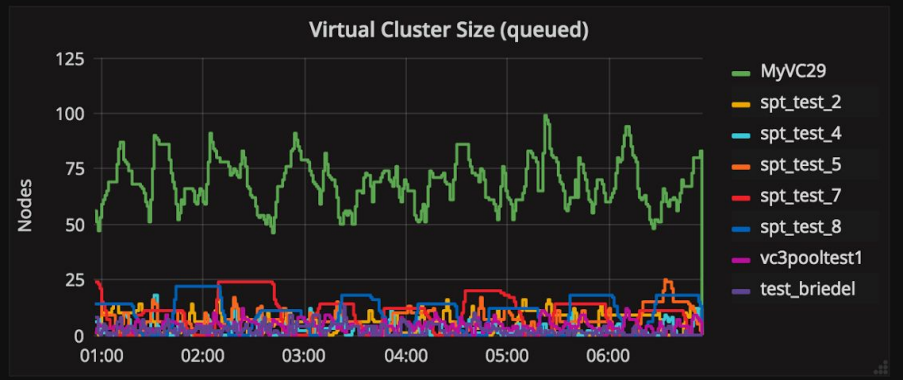

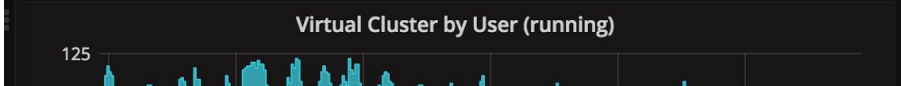

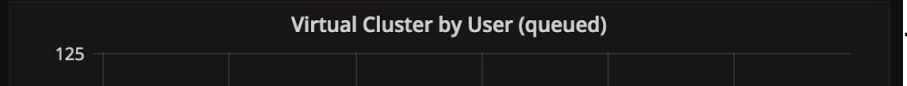

Deploying Software Environments

The **vc3-builder**, a command-line tool for deploying software environments on clusters.

```
vc3-builder
    --require-os centos:7
    --mount /scratch=/data
    --require /cvmfs
    --require python:2.7 -- myapp ...my args...
```

https://github.com/vc3-project/vc3-builder

The MAKER Genomics Pipeline

http://www.yandell-lab.org/software/maker.html

```
vc3-builder --require maker
```

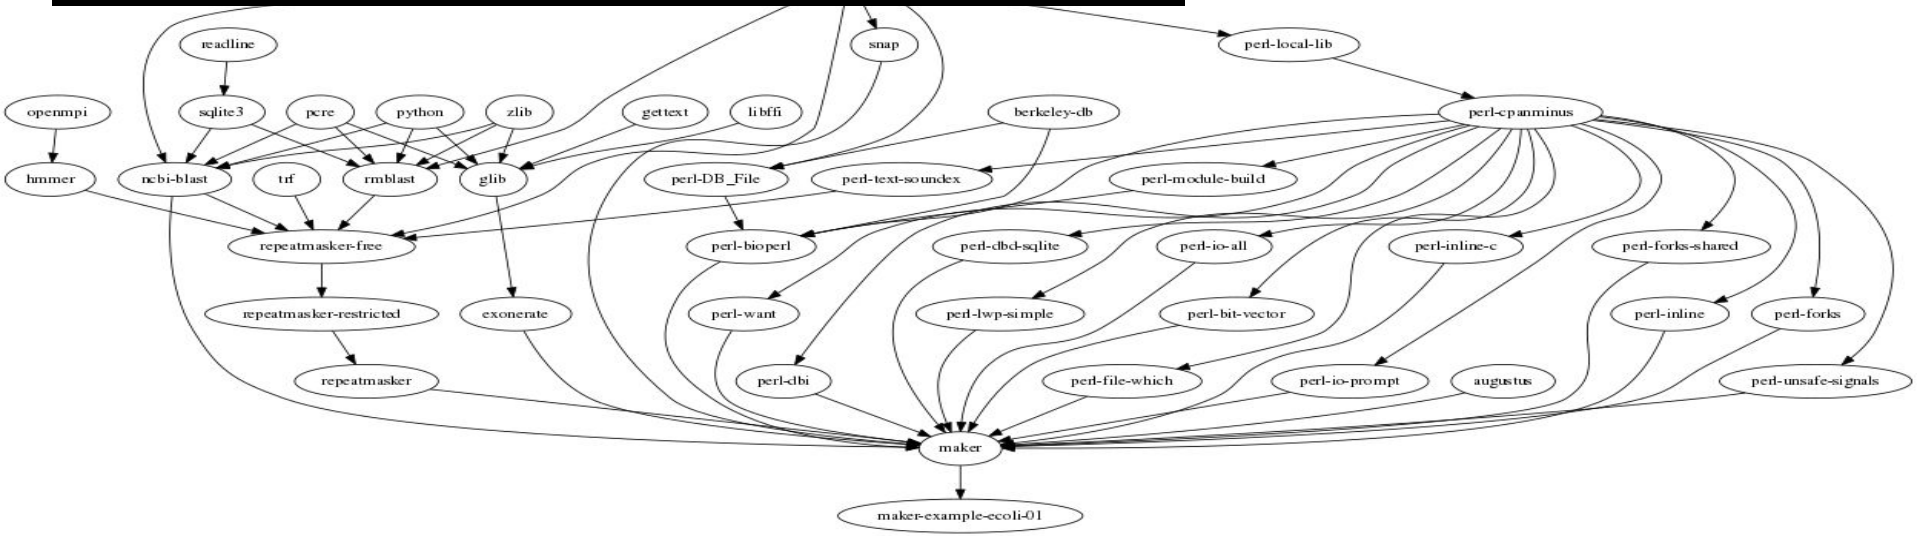

Custom docker container in Jetstream took weeks to install pieces by hand.

Converted to vc3-builder, successfully ported to Stampede in a single automated install.

Working Middleware and Applications

- Various Bioinformatics Workflows
  - Makeflow + HTCondor + BWA, Shrimp, BLAST
- Lobster CMS Data Analysis
  - Work Queue + Builder + CVMFS
- South Pole Telescope (SPT-3G) Analysis Framework
  - HTCondor Jobs + Docker/Shifter + CVMFS
- XENON1T Analysis Framework
  - Pegasus + HTCondor + CVMFS
- MAKER Bioinformatics Pipeline
  - Work Queue + Builder
- IceCube Simulation Framework
  - HTCondor

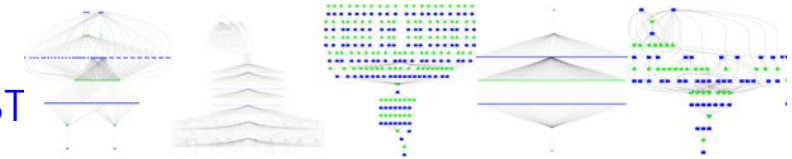

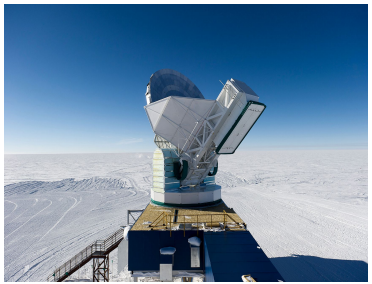

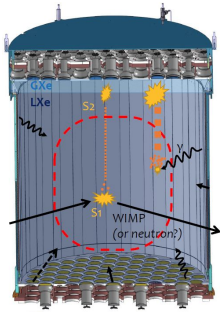

In Progress...

Current Work

- Dynamic reconfiguration of cluster.
- Adding new middleware: Jupyter, Spark, Parsl.

Perennial Challenges

- Idiosyncrasies of each site
- Multi-factor authentication
- Communicate delays/errors from sites to portal

Collaborators and Connections

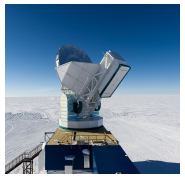

Xe
XENON
Dark Matter Project

CMS
Compact Muon Solenoid

ATLAS
EXPERIMENT

Makeflow

AutoPyFactory

HTCondor
High Throughput Computing

SGCI
Science Gateways
Community Institute

globus

Open Science Grid

ci connect

VC3 Funding and Team

Funded by DOE Office of Advanced Scientific Computing Research (ASCR) and NSF Next Generation Networking Services (NGNS)

PIs: Rob Gardner (UC), Douglas Thain (ND), and John Hover (BNL)

co-PIs: David Miller (UC), Paul Brenner (ND), Mike Hildreth (ND), Kevin Lannon (ND)

dev-team: Lincoln Bryant (UC), Benedikt Riedel (UC), Suchandra Thapa (UC), Jeremy Van (UC), Kenyi Hurtado Anampa (ND), Ben Tovar (ND), Jose Caballero Bejar (BNL).

U.S. DEPARTMENT OF ENERGY | Office of Science

Supported by the Department of Energy Office of Advanced Scientific Computing Research and
Next Generation Networking Services, Solicitation DE--FOA-0001344 (DDRM), Proposal 0000219942
Rich Carlson, Program Manager

VC3

Virtual Clusters for Community Computation

https://www.virtualclusters.org

@virtualclusters

Limited beta signup: http://bit.ly/vc3-signup

Register your HPC: http://bit.ly/vc3-new-resource

Supported by the Department of Energy Office of Advanced Scientific Computing Research and Next Generation Networking Services, Solicitation DE--FOA-0001344 (DDRM), Proposal 0000219942.