

---

# Extending the Reach and Scope of Hosted CEs

OSG All Hands Meeting

March 20, 2018

---

**Suchandra Thapa**  
Robert Gardner  
University of Chicago

Derek Weitzel  
University of Nebraska



**Open Science Grid**

# Introduction

---

- Hosted Compute Elements (CEs) were introduced about a year and a half ago to give sites an easier way to contribute cycles to OSG
- Sites also get a deeper view on their contributions to OSG
- Since then Hosted CEs have extended the sites and resources that can be integrated into OSG
  - **Greater geographical reach**
  - **Sites that differ from the "typical" OSG site**
  - **HPC resources on XSEDE**

# The Hosted CE Approach

---

- Using HTCondor Bosco CE (i.e. 'Hosted CE'), the CE administration can be cleanly separated from the cluster administration
- Cluster admins only need to provide ssh access to cluster
- OSG staff can maintain the hosted CE and associated software and deal with OSG user/site support

# Providing local view of science

---

- After OSG jobs started running, admins can track their cluster's contributions to OSG users using GRACC
- Multiple views on contributions to OSG users:
  - By field of science
  - By project or VO
  - By researcher's institution

# CE hosting infrastructure

---

- Minimal requirements
- Hosted CE can be run on a fairly small VM (1 core / 1GB)
  - Memory usage for typical hosted CE is less than 512MB
  - Hosted CE VMs cpu have been more than 80% idle
  - Max network traffic is fairly low (<200Kb/s)



# Greater Geographical Reach

---

- Low cost of entry has allowed sites to contribute despite time zone and logistical difficulties
- Example: IUCAA Sarathi (LIGO India)
  - Located in Pune, India
  - 12:30 hour difference ( 1 day lag in email responses)
  - Didn't want to require admins to need to learn about the internal details about OSG glidein infrastructure


# IUCAA Sarathi Cluster (LIGO - India)

By VO





LIGO users running under a LIGO specific account through OSG

By Facility



~80k wall hours provided from India this year!

By Project



# Expanding the variety of sites

---

- Bulk of sites contributing to OSG tend to be national labs or large research institutions
  - A lot are brought in by ATLAS or CMS
- Due to the lower cost of entry when using hosted CEs, other types of sites can now contribute:
  - University of Utah
  - North Dakota State University
  - Georgia State University
  - Wayne State



# Example: University of Utah

---

- Several clusters on campus

+



- Time needed to become familiar with OSG CE operations / Glidein troubleshooting

=

- Significant barrier to entry in order to contribute to OSG



# Utah contributions

## By Facility





All three clusters brought into production over the last 2 weeks

## By Project



Still tweaking jobs, looking at using multicore jobs to more effectively backfill and get more cores



## By Field Of Science



Already contributed ~60k cpu hours, in top 2 institutions contributing through Hosted CEs



# North Dakota State University

▼ By Facility



Two clusters, CCAST3 was brought online beginning of year



▼ By Project



Single core jobs on CCAST2, 8 core jobs on CCAST3


670K wall hours delivered, one of top hosted CE sites

▼ By Field Of Science




# Georgia State University


▼ By Project



▼ By Field Of Science



WallHours by Project




194K wall hours  
delivered since  
Jan 1

18 Projects  
helped

Provided cpu to  
11 fields of  
science


12 Institutions  
ran jobs on  
resource

Wallhours By Organization




# Wayne State University


## ▼ By Project



## ▼ By Field Of Science



WallHours by Project




	total
TG-IBN130001	137.2 K
macsSwigmodels	54.4 K
GLOW	42.3 K
SPLINTER	25.5 K
Duke-QGP	23.4 K
TG-CHE170021	15.4 K
nicesims	1.604 K
TG-PHY150040	395
SourceCoding	133

300k cpu  
hours  
delivered  
since Jan 1


Ran jobs  
from 24  
projects, 13  
fields of  
science, and  
14 institutions

Wallhours By Organization




# Total Contributions

WallHoursSpentOnJobsByFacility




✓ By Project

Wallhours By Project




✓ By Field Of Science

Wallhours By Field of Science



Wall Hours by Facility




	total
OSG_US_NDSU_CCAST_CLUSTER3	491 K
OSG_US_WSU_GRID	300 K
OSG_US_GSU_ACORE	194 K
OSG_US_NDSU_CCAST_CLUSTER2	179 K
OSG_IN_IUCAA_SARATHI	79.7 K
GLUEX_US_FSU_HNPGRID	30.5 K
OSG_US_UUTAH_KINGSPEAK	23.1 K
OSG_US_UUTAH_EMBER	19.1 K
OSG_US_UUTAH_LONEPEAK	17.5 K

>1.3M wall hours delivered since Jan 1

Averaging about 111K wall hours a week. About 10-15% of weekly opportunistic usage by OSG Connect users


Ran jobs from 25 fields of science and 35 institutions

WallHours by Project



	total
TG-IBN130001	767 K
Duke-QGP	180 K
SPLITTER	166 K
macsSwigmodels	89.0 K
GLOW	42.3 K
TG-CHE170021	32.6 K
UPCDOSAR	22.5 K
molcryst	10.22 K
BioGraph	9.69 K

Wallhours By Organization



# Integrating HPC resources into OSG

---

- Major cultural differences between HPC resources and OSG resources
  - MultiFactor Authentication (MFA) using tokens
  - Software access and distribution
  - Allocations

# Bridging the Gap

---

- Solutions:
  - Authentication -> get MFA exceptions or use IP address as a factor
  - Software access -> Stratum-R
  - Job Routing -> Multi-user BOSCO



# User Authentication

---

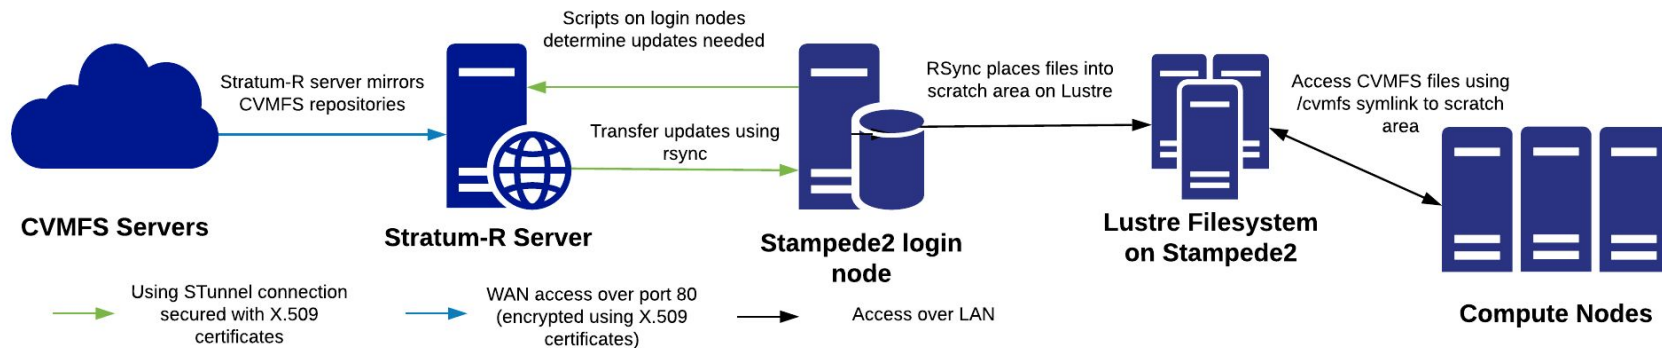
- HPC resources are increasingly moving to using MFA
  - OSG software doesn't have any way to incorporate token requirements into job authentication
- Solutions:
  - Use submit site's IP as one factor.
    - All job submissions come from a fixed IP.
    - Can use a ssh public key or proxy as another factor
  - Get a MFA exception for accounts
    - Sites often have procedures requesting this for science gateways or similar facilities to use.

# Software distribution and access

---

- VOs and users are increasingly using CVMFS to distribute software and data
  - HPC resources usually aren't willing to install and maintain CVMFS on their compute nodes
- Stratum-R allows for replication of selected CVMFS repositories
  - Requires some effort from admins but not much
  - Successfully used on Bluewaters, Stampede, Stampede2

# Stratum-R

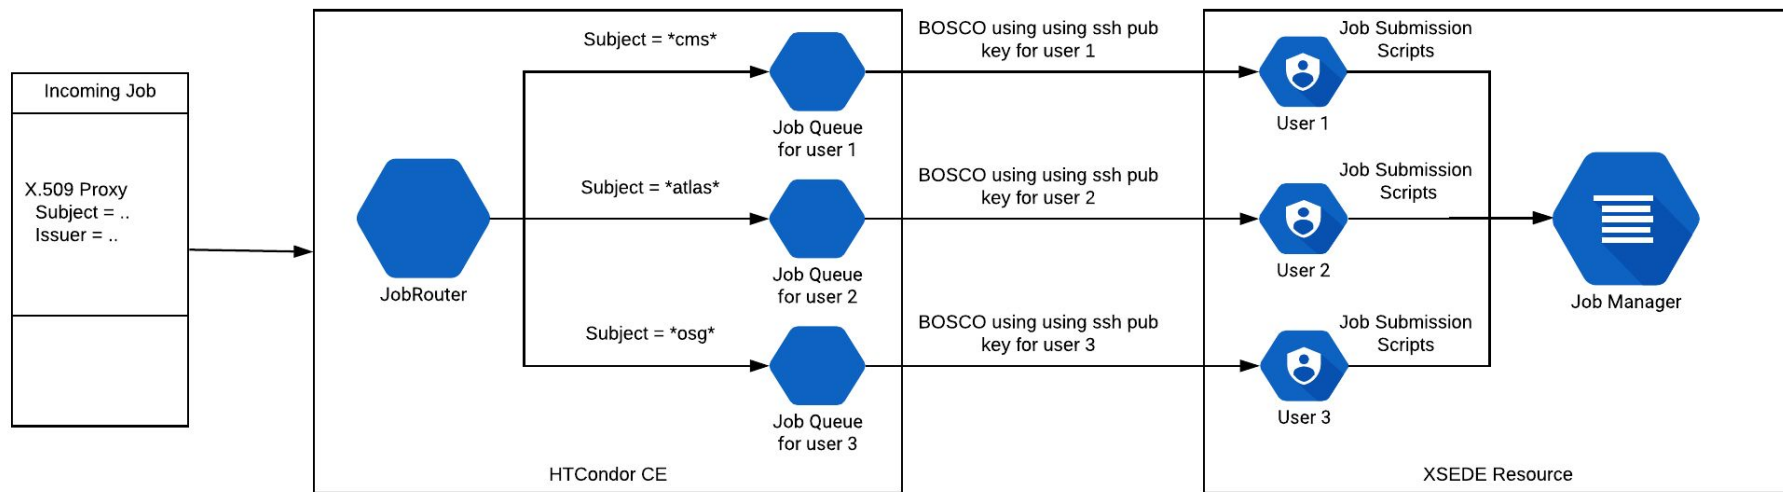


# Routing jobs to allocations

---



- Due to allocations, must route jobs to proper user accounts on HPC resources
- BOSCO's default configuration is to use a single user on remote resource for **all** job submissions
- With some modifications to config files, JobRouter entries, and other bits, can have jobs going to different users on remote resources
  - This allows for jobs to use different allocations, partitions, configurations

# HTCondorCE BOSCO Job Routing





# Running CMS jobs on XSEDE

## Stampede2



Still validating and testing CMS workflows on Bridges and Stampede2

## Bridges



# Conclusions

---

- Hosted CEs offer OSG the opportunity to obtain cycles and engage with new types of sites and resources increasing the diversity and reach of OSG.
  - Smaller universities and institutions
  - XSEDE resources (direct allocations for users)

# More information

---

- [Support document for cluster admins](#)
- [BOSCO CE](#)



# Acknowledgements

---

- Derek Weitzel
- Factory Ops (Jeff Dost, Marian Zvada)
- David Lesny
- Mats Rynge
- CMS HepCloud Team (Dirk, Ajit, Burt, Steve, Farrukh)
- Lincoln Bryant - Infrastructure support
- Rob Gardner