

HTC23 Summary

Brian Lin
Center for High Throughput Computing
University of Wisconsin–Madison



Throughput Computing (HTC) 23

- First joint OSG All Hands and HTCondor Week event organized by the PATH project (<https://agenda.hep.wisc.edu/event/2014/>)
- 4.5 days of content recorded and available on YouTube (<https://www.youtube.com/@centerforhighthroughputcom4787/playlists>)
- Single track most days, parallel sessions on Jul 12:
 - Tutorials
 - ATLAS + CMS
 - Collaborations
- Mostly in-person attendance, plenaries broadcast on Zoom
- 140 in-person attendees and over 270 remote registrants from 174 attending institutions




Science Talks

David Swanson Awardees

- Award in memory of David Swanson, longtime champion and contributor to the OSG Consortium
- Awarded to OSG User School students achieving significant research outcomes through distributed HTC
- Presentations from current and past awardees
 - Jimena González (Dark Energy)
 - Aashish Tripathee (Gravitational Waves)
 - Anirvan Shukla (Antimatter)
 - Zhonggang Li (Genomics)
 - Nicholas Cooley (Bioinformatics)
 - Connor Natzke (Gamma Rays)

Researchers Using HTC

- Keynote from IGWN by Laura Cadonati
- Drug screening for Multiple Myeloma by Rousselene Larson
- Milky Way Black Hole Imaging (EHT) by Chi-Kwan Chan
- Radio interferometry (NRAO) by Felipe Madsen
- Rubin Observatory Updates by Greg Daues
- Mars Reconnaissance Orbiter Imagery by Mike Fienen



PATh Team News

Contains European HTCondor Workshop Spoilers!

HTCSS and OSG Releases

- Next major version of the HTCondor Software Suite will be HTCSS 23
- Next OSG Release Series, OSG 23, will contain HTCSS 23
- Both OSG and HTCSS will switch to an annual release schedule
 - Major version reflects the release year
 - Targeting Q3 releases each year
 - Two major versions will be supported at a time
- HTCSS 23 and OSG 23 are targeted for the end of September

Evolving new command line user interface

- *htcondor* <noun> <verb>
 - "*htcondor job submit*", "*htcondor job status*", ...
 - "*htcondor dag submit*", "*htcondor dag status*", ...
 - "*htcondor jobset submit*", "*htcondor jobset status*", ...
 - "*htcondor annex create*", "*htcondor annex status*", ...
 - "*htcondor eventlog read*"
- Legacy tools (*condor_q*, *condor_submit*, *condor_history*, ...) not going anywhere...

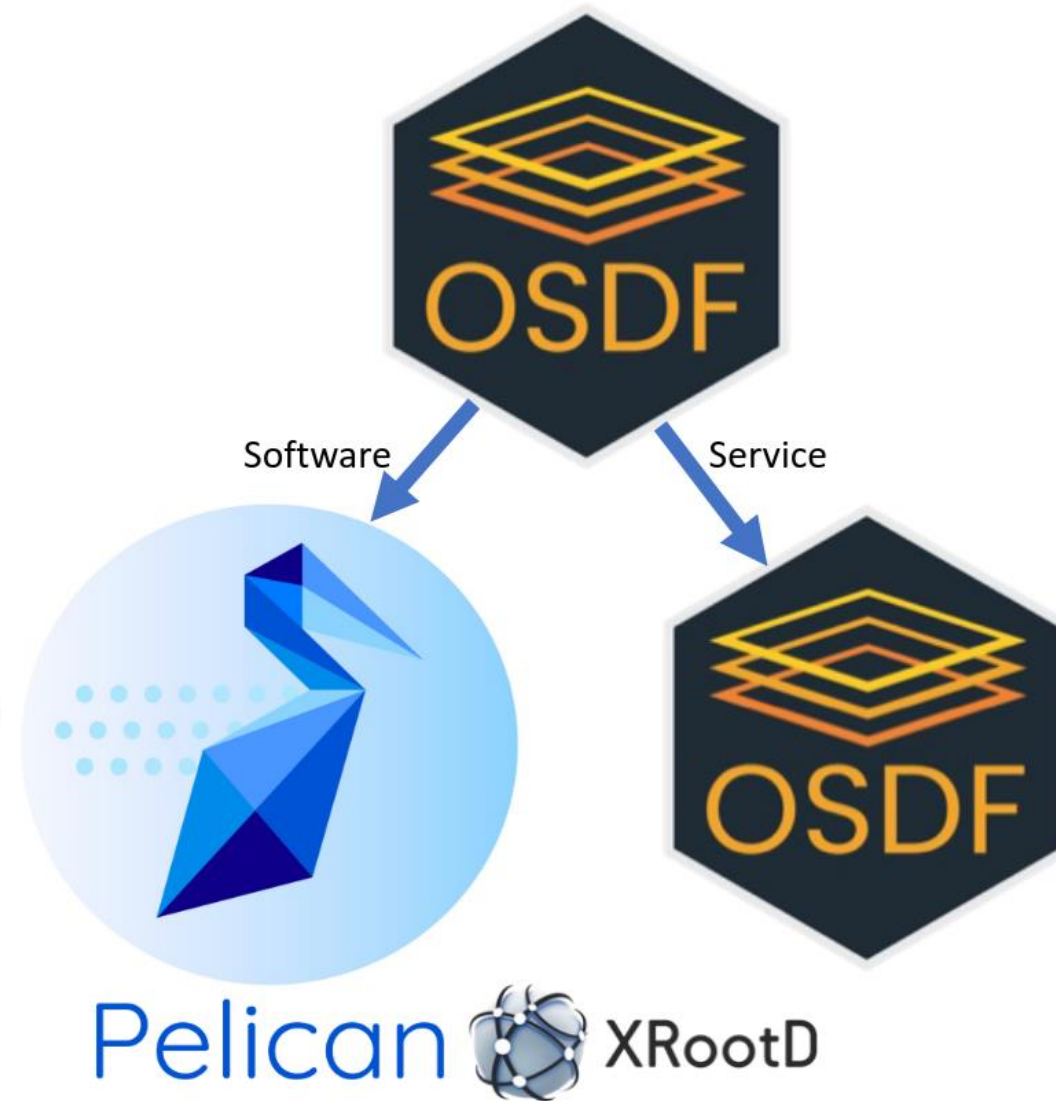
<https://agenda.hep.wisc.edu/event/2014/contributions/28437/>

Other HTCSS Enhancements

- Configurable attributes in job and global event logs
- Saved DAG progress file
- Improved resource management: access points claim execution point partitionable slots
- Container universe: execution point validation of container runtimes → better differentiation between system and user container errors
- First-class backfill partitionable slots (e.g., run CPU jobs on GPU EPs in low GPU demand scenarios)

OSDF & Pelican

- We split out the technology powering the OSDF and christened it “Pelican”.
- The software is the same – integrating technologies like XRootD, SciTokens, OA4MP, and the ‘stashcp’ client utilities.



<https://agenda.hep.wisc.edu/event/2014/contributions/28482/>

Architecture

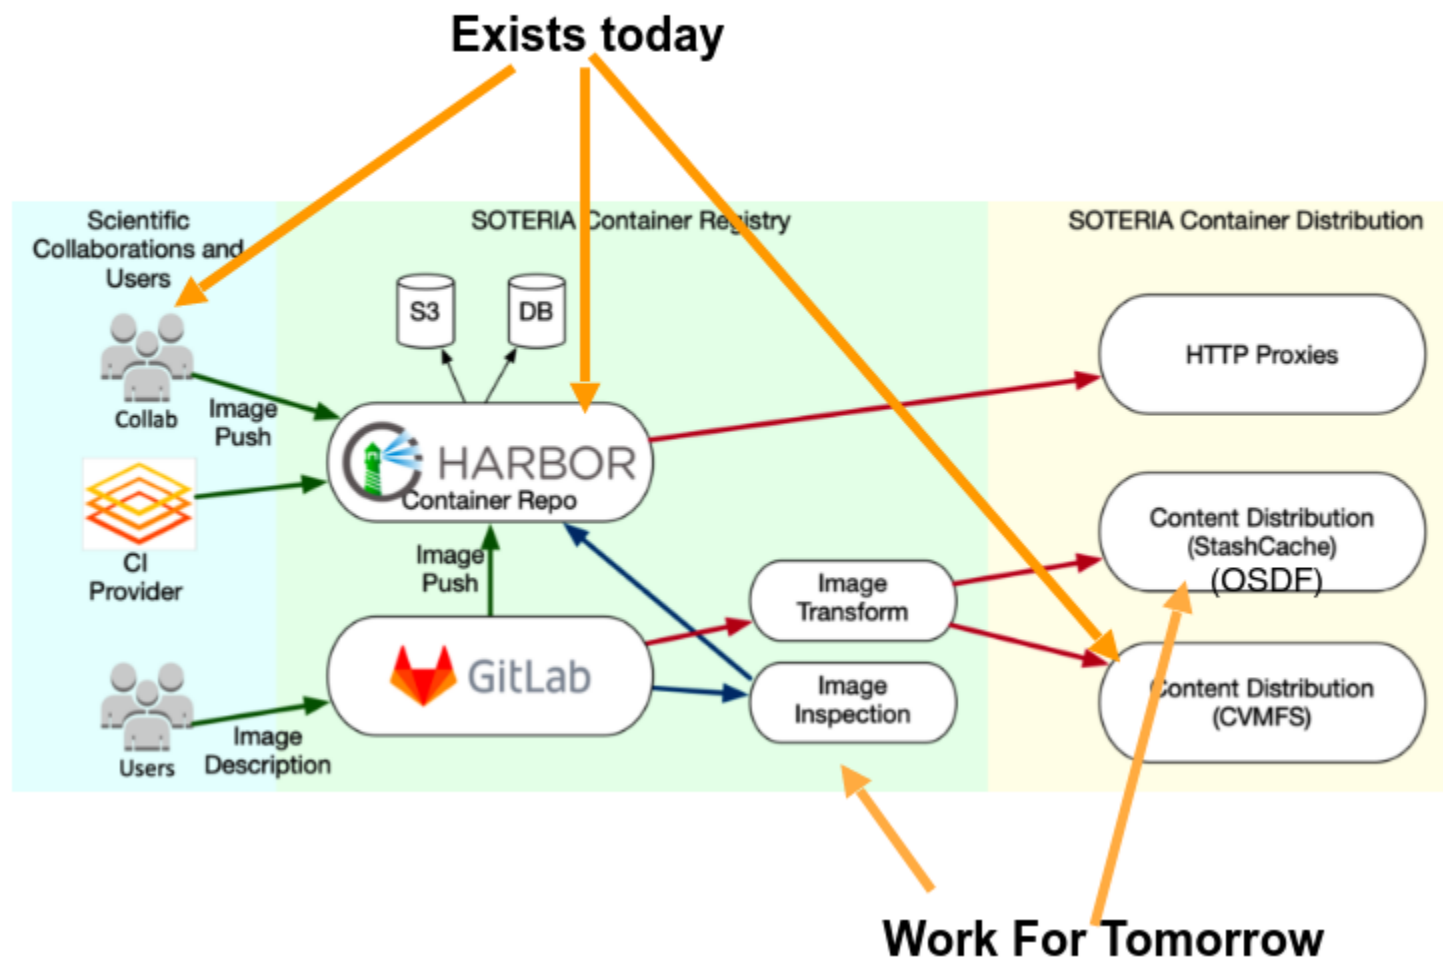


Image creation



Image introspection



Image distribution



<https://agenda.hep.wisc.edu/event/2014/contributions/2848>

3/

Components

Database

Elasticsearch

REST API and Web frontend

Node.js + express + pug

Deployment

Docker, K8s, Helm (soon)

Authentication

Globus InCommon

Authorization

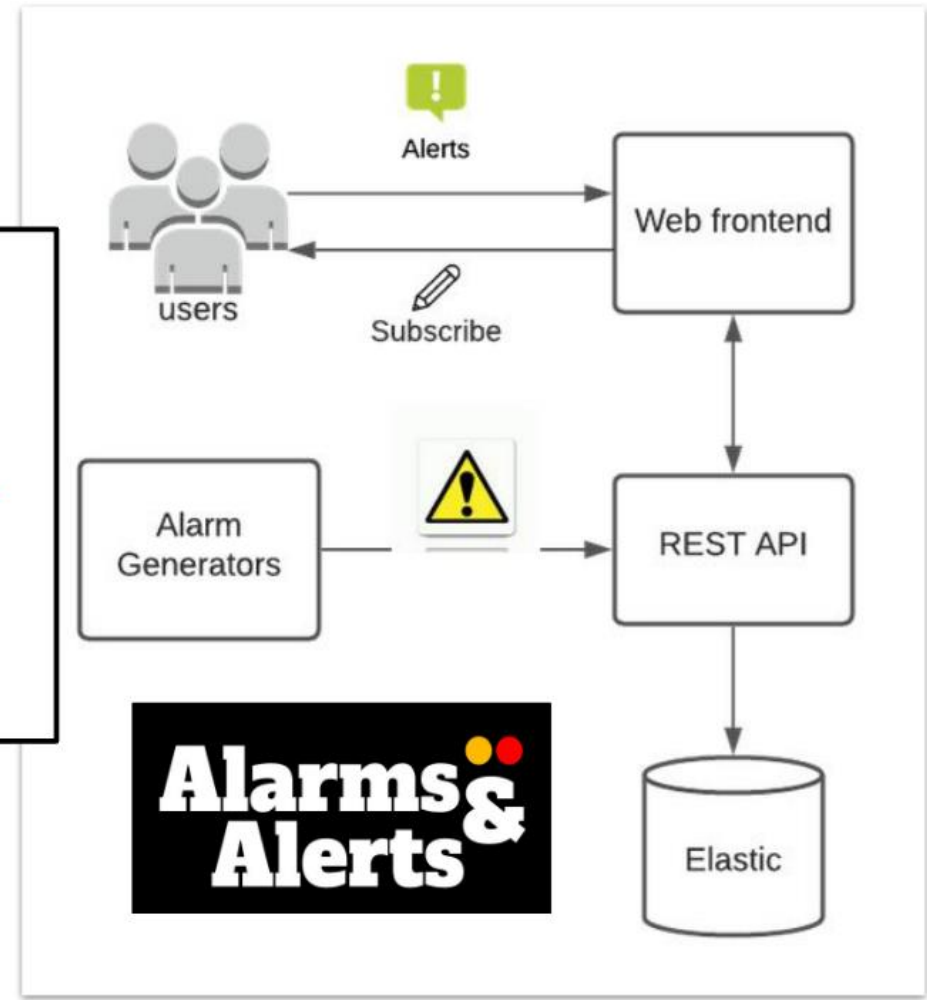
API key

Mail

Mailgun

<https://psa.osg-htc.org>
(Uses EDUGain/InCommon)

Purpose: provides user-subscribable alerting for specific types of network issues found by analyzing perfSONAR data



<https://agenda.hep.wisc.edu/event/2014/contributions/2848>

9/

Infrastructure Talks

HTC installations

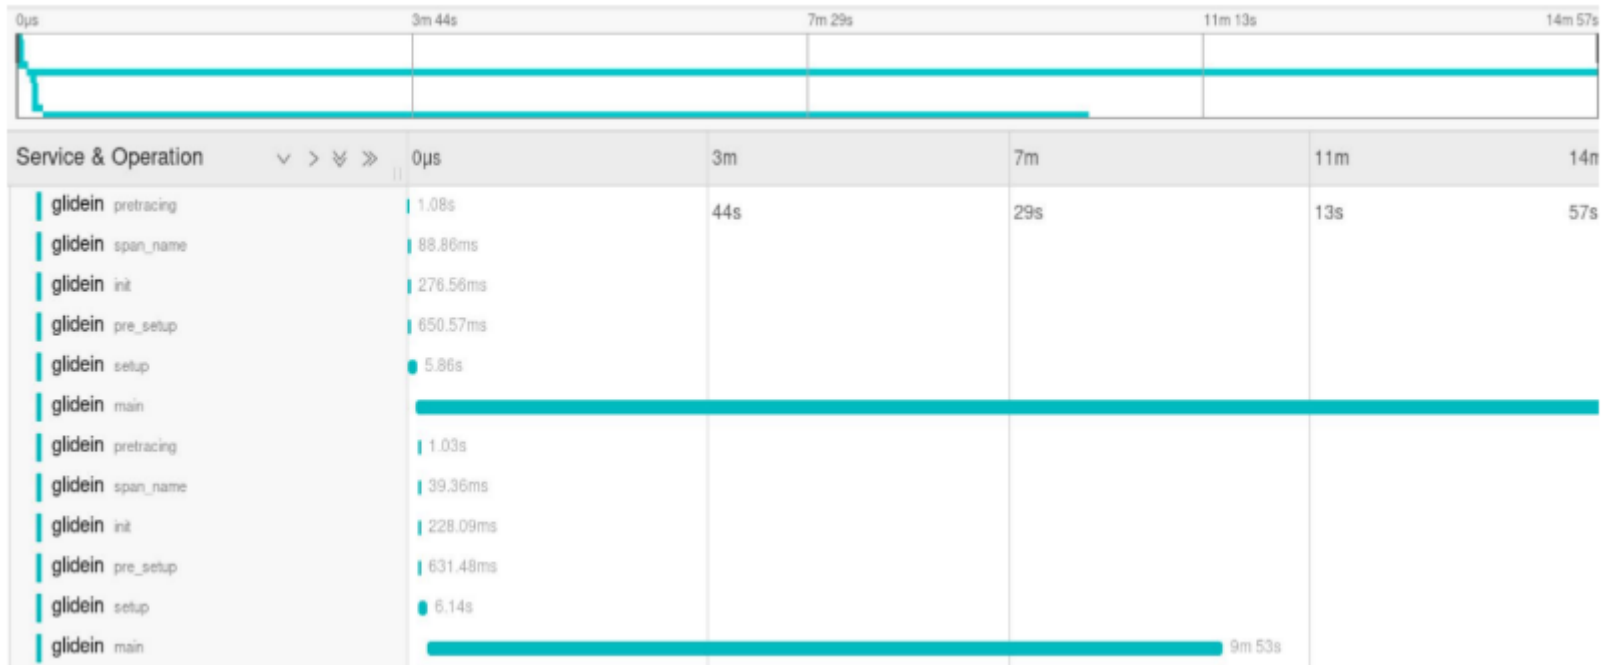
- Regional computational collectives, CC* (Great Plains Network, University of Maine, OAK)
- The Condor Hatches in Africa,
<https://agenda.hep.wisc.edu/event/2014/contributions/28447/>
- Distributed High-Throughput Computing in China (IHEP),
<https://agenda.hep.wisc.edu/event/2014/contributions/28435/>

Parallel Sessions

- LHC
 - Data transfer with intelligent networking (Rucio/SENSE)
 - ARM facility testbeds
 - Federating XCache and Varnish proxies
 - Service Deployment in FABRIC at CERN
- Collaborations
 - IGWN Pool “Exorciser”
 - DUNE updates, participation in DC24
 - GlueX experiments with XRootD data chunking
 - Experiences with OSG from ePIC and KOTO

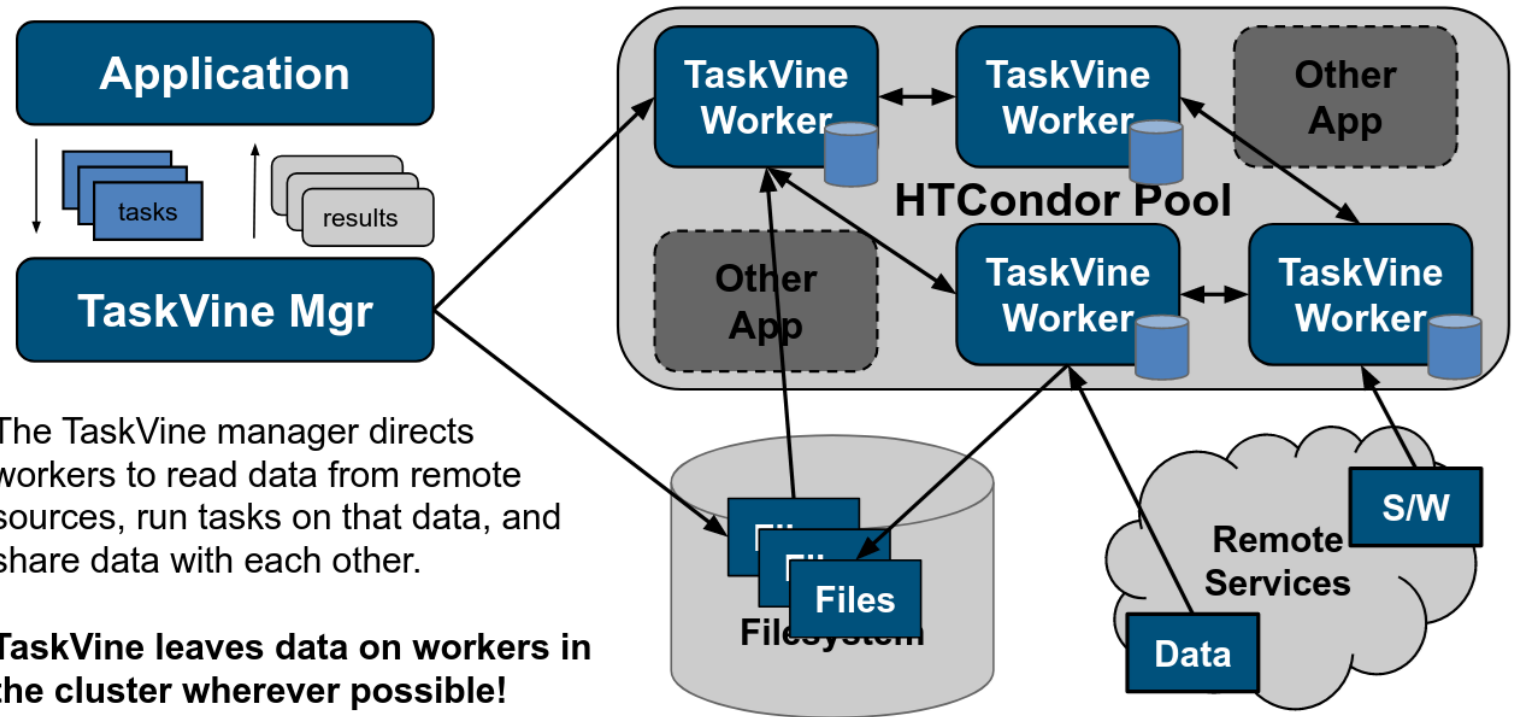
GlideinWMS Updates

- VO credential refactoring and improved support
- Traceability



Other Updates

- TaskVine:
dynamic data sharing built on HTCondor
- New Google Cloud Toolkit (simplified deployment, additional OS support, etc.)



Questions?

This material is based upon work supported by the National Science Foundation under Grant Nos. 1836650 and 2030508. 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 National Science Foundation.