SLATE

A new approach for DevOps in ***distributed*** scientific computing facilities

Rob Gardner
University of Chicago

Middleware and Grid Interagency Coordination (MAGIC) Meeting
October 3, 2018

Outline

- What is **SLATE**?
- The motivation
- The **SLATE** Vision
- Current technology explorations
- Challenges and open questions
- Wrap up

What is **SLATE**?

- NSF DIBBs award, "SLATE and the Mobility of Capability" (NSF 1724821)
- Equip the ScienceDMZ with service orchestration capabilities, federated to create scalable, multi-campus science platforms
- Platform for service operators & science gateway developers

Motivation: enabling multi-institution collaborative science

XENON - Dark Matter Search in Gran Sasso Laboratory, Italy

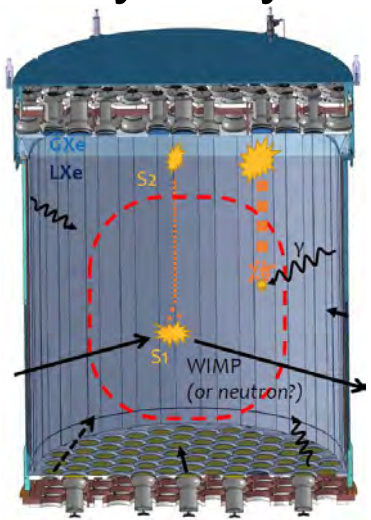

Collaboration

Columbia · RPI · Nikhef · Muenster · Stockholm · Mainz · MPIK, Heidelberg · Freiburg · Zurich

Chicago · Tokyo · UCLA · Nagoya · UCSD · Kobe · Rice · Purdue · Coimbra · Subatech · LPNHE · LAL · Bologna · LNGS · Torino · Napoli · Weizmann · NYUAD

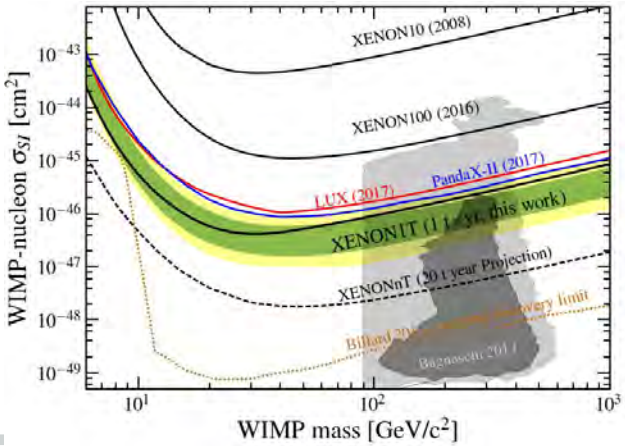

Example

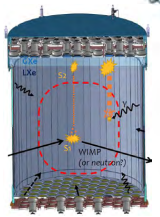

DAQ Interprocess Communication @ LNGS
DAQ Web Server @ LNGS
Slow Control (Via DAQ API) @ LNGS
XENON Runs DB Mirror at UChicago and PDC Stockholm
EU & US storage
XENON DAQ @ LNGS
XENON Runs DB @ LNGS
Detector Settings Run Information
Data Location Transfer Status
Nikhef
CNAF
SURFSara
Weizmann
CCIN2P3
Future Site
OSG dCache @ UChicago
EU & US processing
EGI
XENON Node @ UChicago
OSG
OSG dCache @ UChicago
UChicago Research Computing Center
FTS @ UChicago
Rucio DB
Rucio Service @ UChicago
Data Storage @ LNGS
Tivoli Storage Manager
Tape Backup @ PDC Stockholm
RUCIAX
Software connecting XENON Runs DB, Rucio service, and Tape Backup
CAX
Software submitting
OSG = Open Science Grid
EGI = European Grid Infrastructure
Raw Data
Processed Data
Random Selected Element
XENON Runs DB Interactions
Rucio DB Interactions
Jobs

Global data & processing platform

Job management with HTCondor & workflow pipeline tools

The Open Science Grid

Example

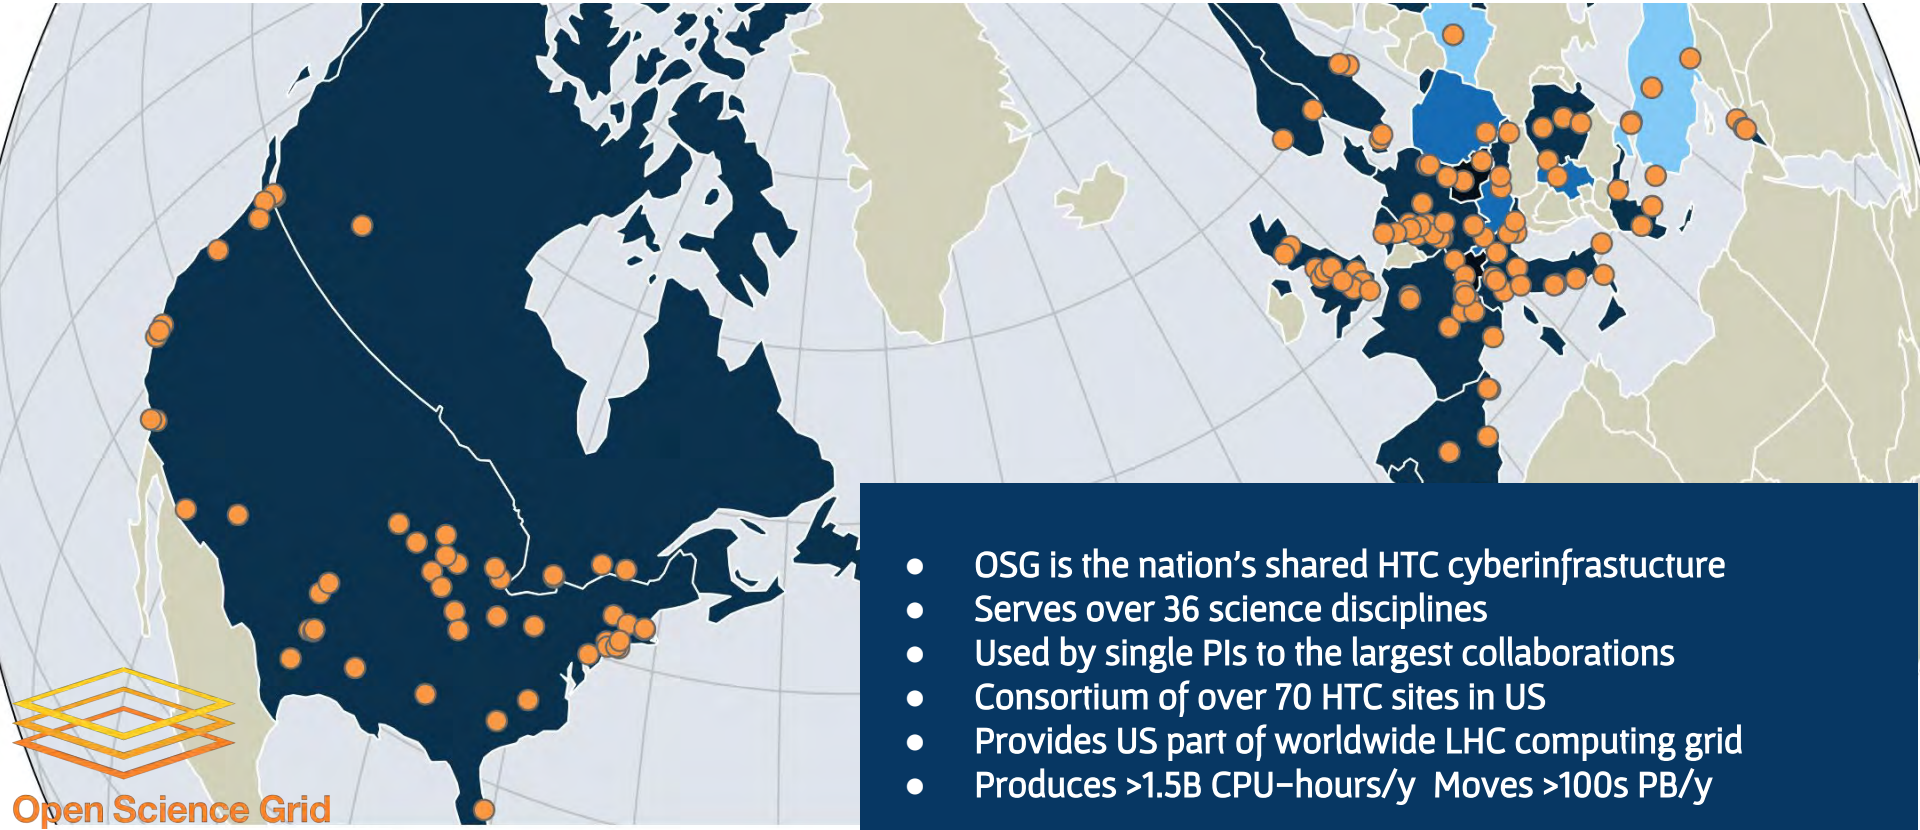

- OSG is the nation's shared HTC cyberinfrastucture
- Serves over 36 science disciplines
- Used by single PIs to the largest collaborations
- Consortium of over 70 HTC sites in US
- Provides US part of worldwide LHC computing grid
- Produces >1.5B CPU-hours/y Moves >100s PB/y

Open Science Grid

Facilitator for ”data lake” R&D

Example

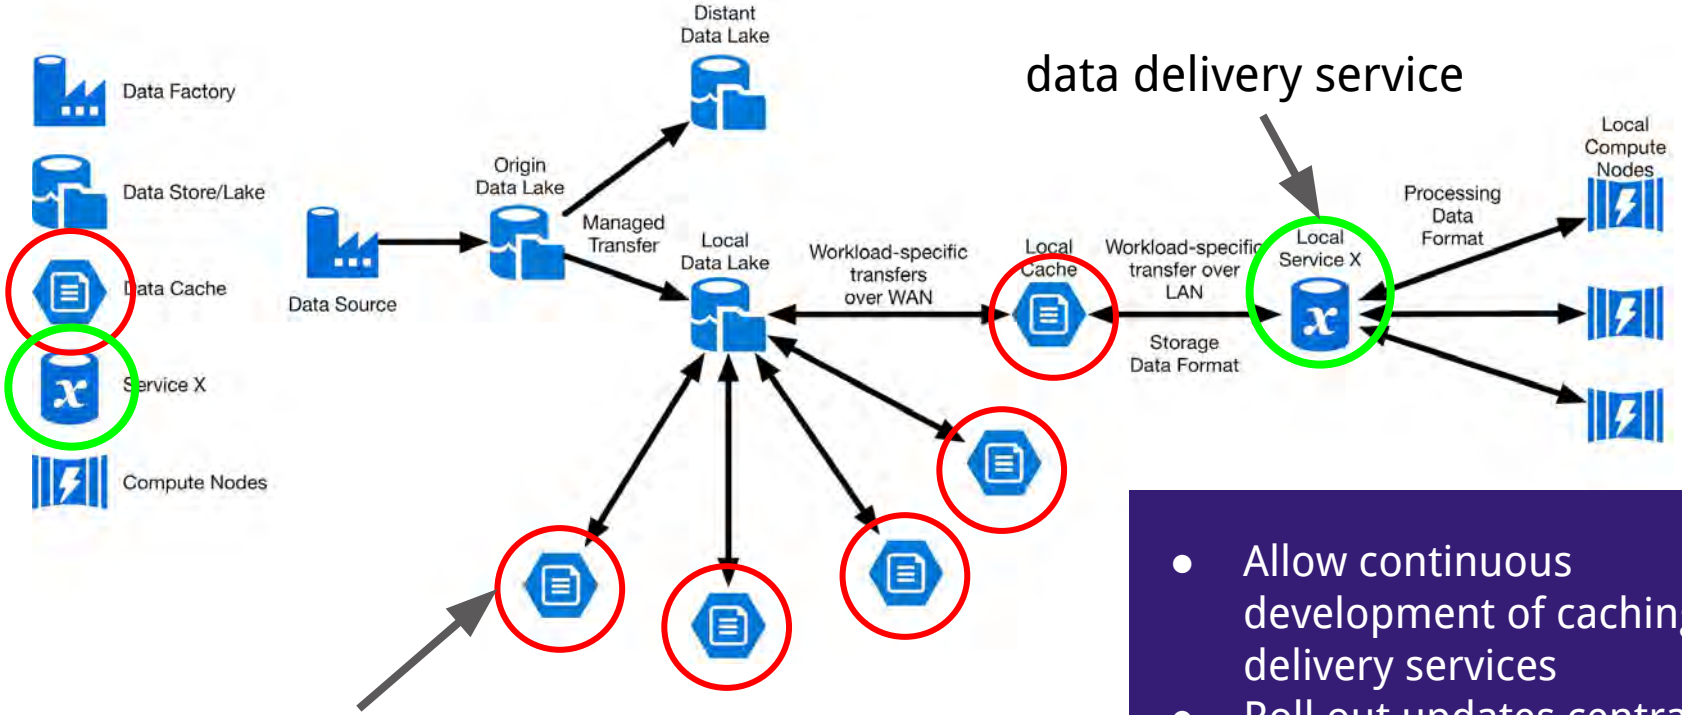

- Allow continuous development of caching & delivery services
- Roll out updates centrally
- Configure & Op centrally

Caching network for IceCube & LIGO

Example

containerized by

PRP
PACIFIC RESEARCH PLATFORM

Deployment is difficult!

- **A broken DevOps cycle!**
- Deployment means:
  - Finding a friendly sysadmin at the site
  - Having them procure hardware or a virtual machine
  - Sending them the deployment instructions and hoping for the best
- Operations problems too:
  - Someone has to make sure it actually keeps running
  - Latency in updates across sites make it extremely difficult to rapidly innovate platform services

The **SLATE** Vision

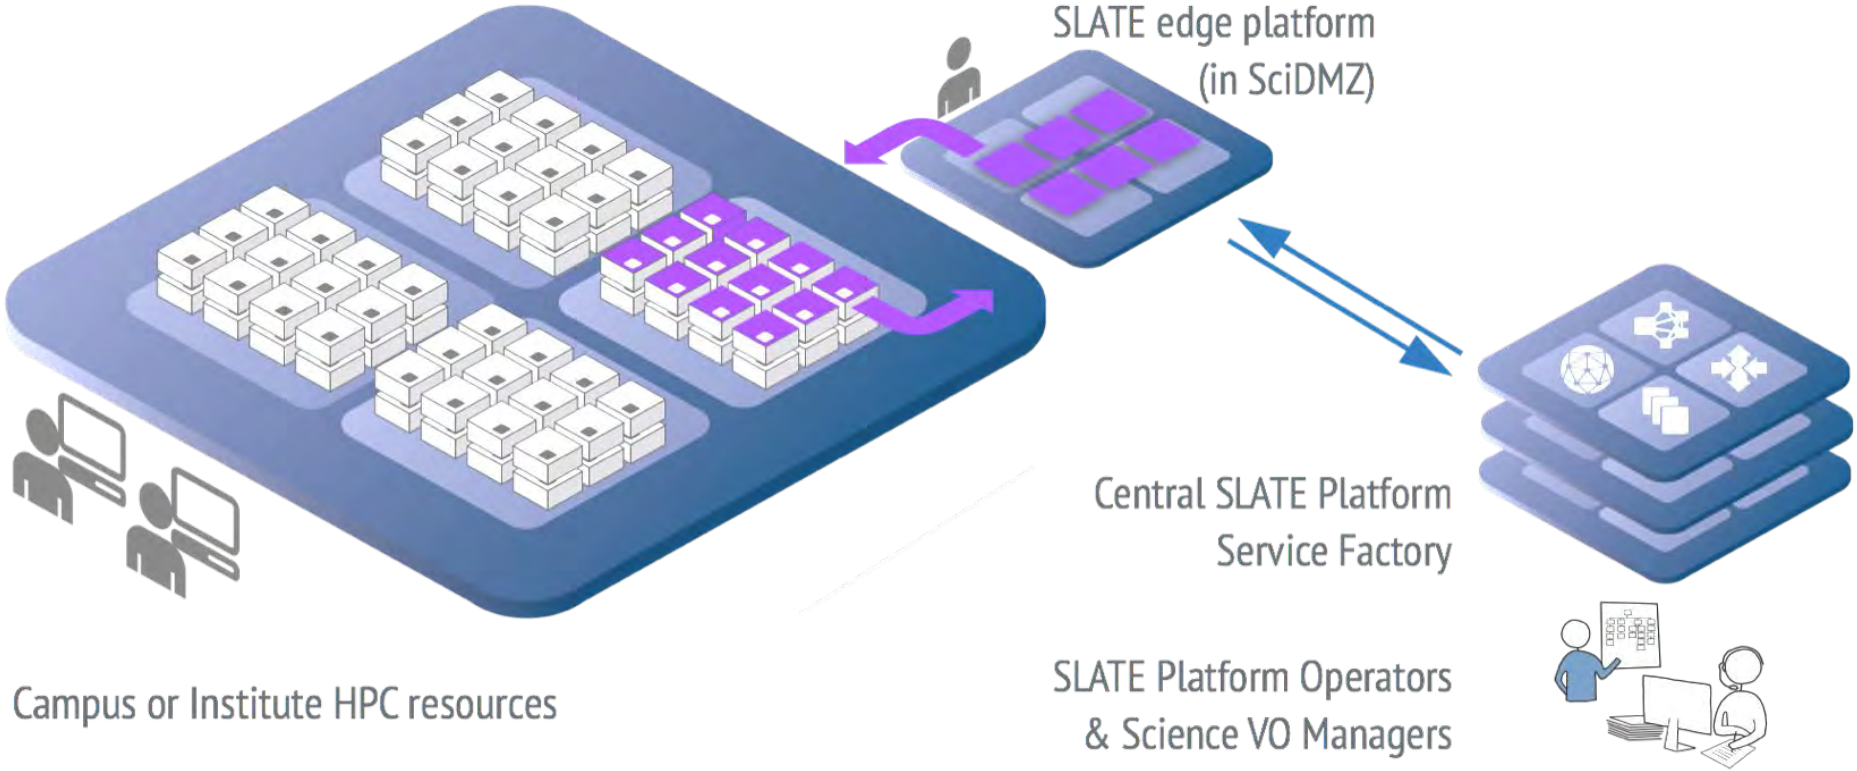
SLATE edge platform
(in SciDMZ)
Central SLATE Platform
Service Factory
SLATE Platform Operators
& Science VO Managers
Campus or Institute HPC resources

XENON COMPUTING

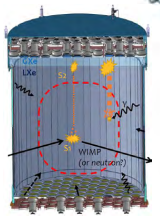

DAQ Interprocess Communication @ LNGS

DAQ Web Server @ LNGS

Slow Control (Via DAQ API) @ LNGS

XENON Runs DB Mirror at UChicago and PDC Stockholm

EU & US storage

XENON DAQ @ LNGS

XENON Runs DB @ LNGS

Detector Settings Run Information

Data Location Transfer Status

Nikhef

CNAF

Weizmann

CCIN2P3

Future Site

OSG dCache @ UChicago

FTS @ UChicago

Data Storage @ LNGS

Tivoli Storage Manager

Tape Backup @ PDC Stockholm

RUCIAX
Software connecting XENON Runs DB, Rucio service, and Tape Backup

EGI

XENON Node @ UChicago

OSG

CAX
Software submitting

Run Information for Data Analysis

OSG dCache @ UChicago

UChicago Research Computing Center

AUTOMATE DEVOPS

Global data & processing platform

Job management with HTCondor & workflow pipeline tools

OSG = Open Science Grid
EGI = European Grid Infrastructure

Raw Data
Processed Data
Random Selected Element
XENON Runs DB Interactions
Rucio DB Interactions
Jobs

The Open Science Grid

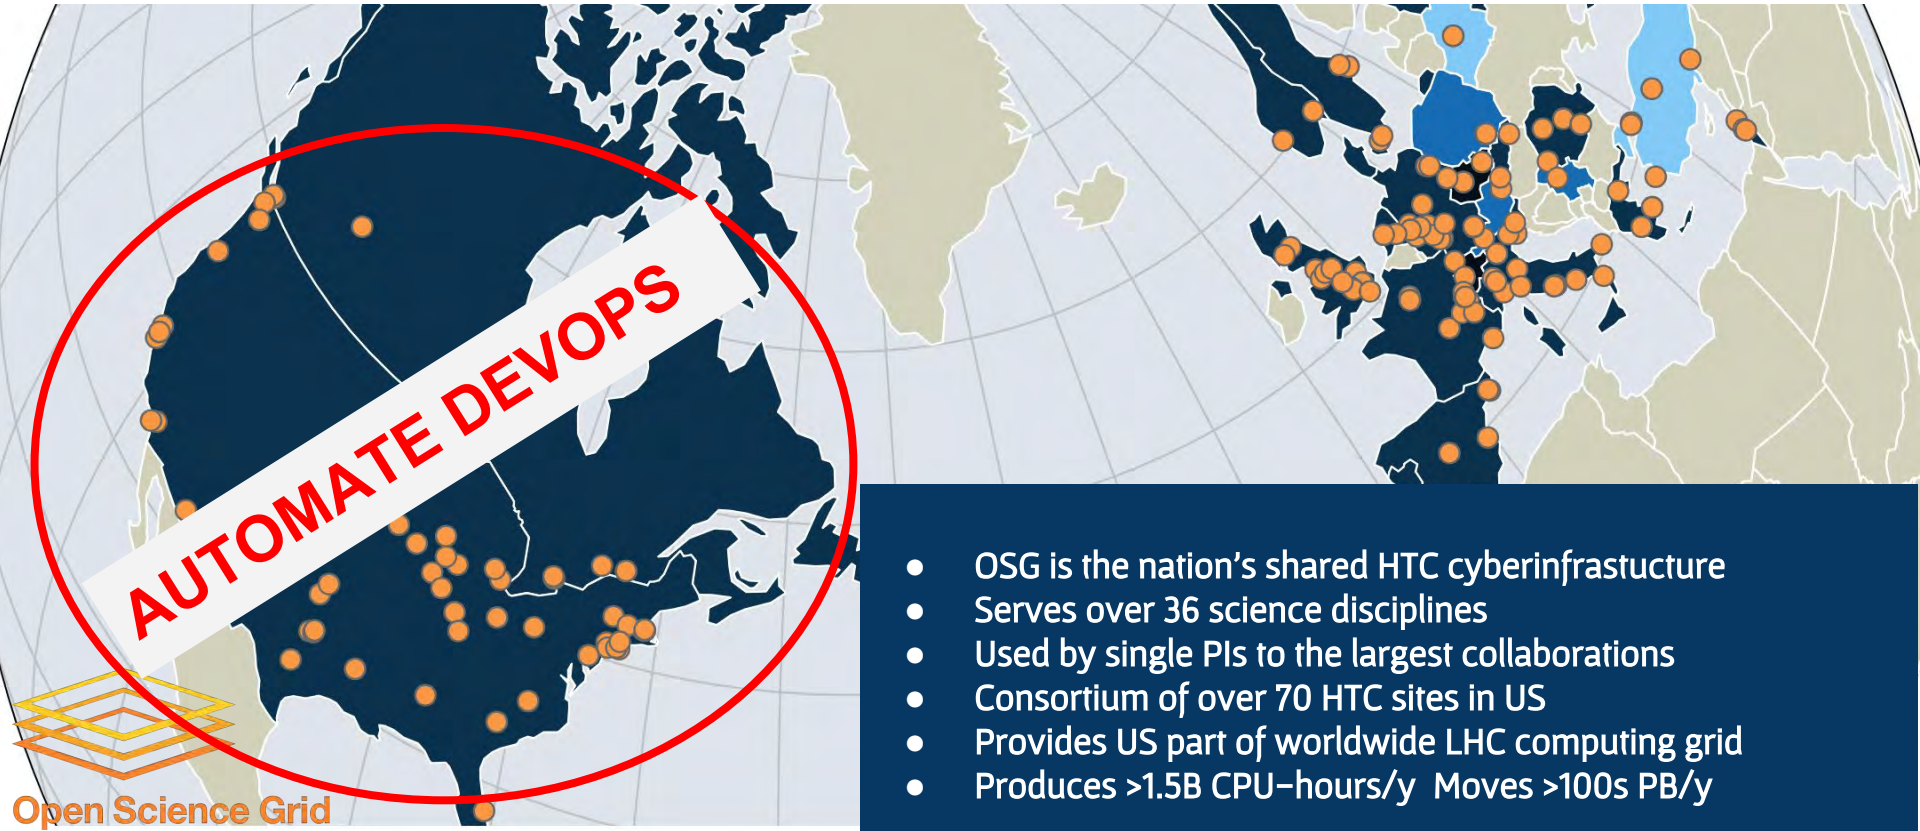

- OSG is the nation's shared HTC cyberinfrastucture
- Serves over 36 science disciplines
- Used by single PIs to the largest collaborations
- Consortium of over 70 HTC sites in US
- Provides US part of worldwide LHC computing grid
- Produces >1.5B CPU-hours/y Moves >100s PB/y

WLCG
Worldwide LHC Computing Grid

Caching network deployed for IceCube & LIGO

containerized by

PRP
PACIFIC RESEARCH PLATFORM

Services **L**ayer **A**t **T**he **E**dge

- A ubiquitous *underlayment* -- the missing shim
  - A generic cyberinfrastructure substrate optimized for hosting edge services
  - Programmable
  - Easy & natural for HPC and IT professionals
  - Tool for creating "hybrid" platforms
- DevOps friendly
  - For both platform and science gateway developers
  - quick patches, release iterations, fast track new capabilities
  - reduced operations burden for site administrators

SLATE Concepts & Components

http://bit.ly/slate-arch

- Containerized services in managed clusters
- Widely used open source technologies for growth and sustainability
- SLATE additions
  - Curated services
  - Create a “Loose federation” of clusters & platforms

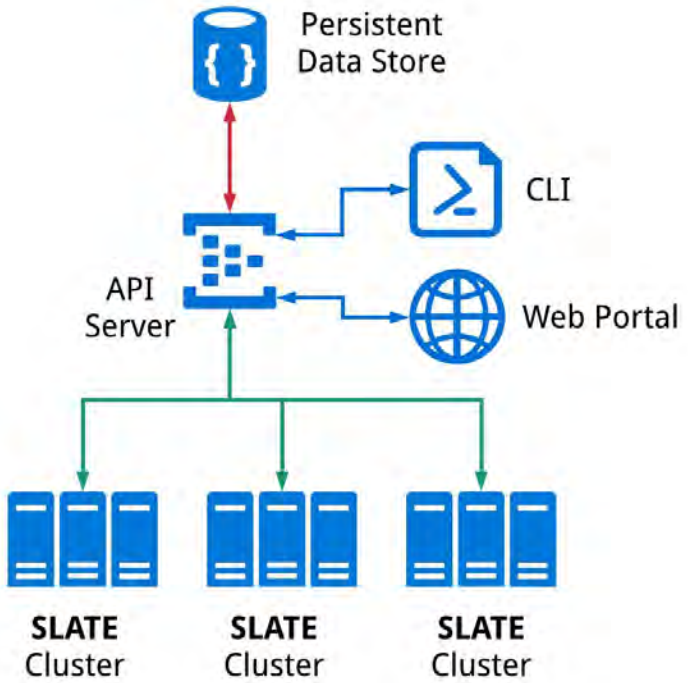

InCommon
signup/login

developers
(& admins)

cluster
admins

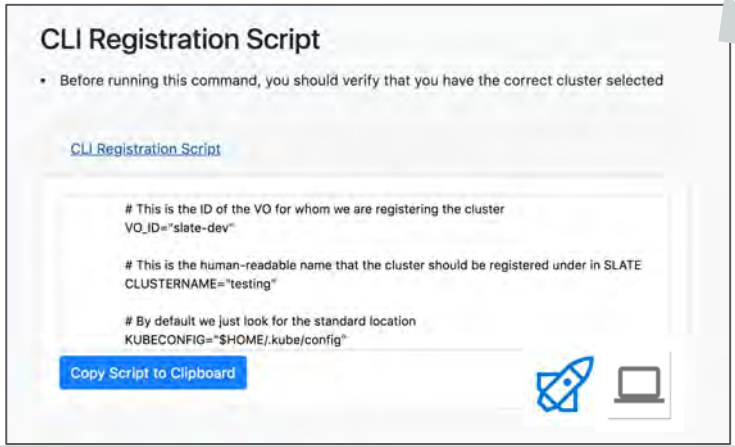

Policy and Trust

- **SLATE** applications curated into a trusted application catalog
- Applications must define and request all needed network, disk, device, etc access.
  - Think application permissions on your phone
- Site policies must be respected
  - Access, privileges, capabilities are controlled and transparent

Deploying an "Application"

```
$ slate-client cluster list
Name         ID                                           Owned By
umich        Cluster_d3732e1d-7ea0-4022-96fc-288a0c8a7c5d slate-dev
utah-coreos  Cluster_3249cb47-7318-4fd0-a61b-0cf99c1aceb8 slate-dev
uchicago     Cluster_98b60d59-b873-4014-8f1d-f9c259c116b3 slate-dev

$ slate-client app list
Name                App Version  Chart Version  Description
jupyterhub          v0.8.1       v0.7-dev       Multi-user Jupyter installation
osg-frontier-squid  squid-3      0.2.0          A Helm chart for configuration and deployment o...
osiris-unis         1.0          0.1.0          Unified Network Information Service (UNIS)
perfsonar           1.0          0.1.0          perfSONAR is a network measurement toolkit desi...

$ slate-client app install --vo slate-dev --cluster uchicago osg-frontier-squid proxy-test
Successfully installed application osg-frontier-squid as instance slate-dev-osg-frontier-squid-proxy-
test with ID Instance_dd427321-05f5-42a2-b61c-e21169187188

$ slate-client instance info Instance_dd427321-05f5-42a2-b61c-e21169187188
Name                                  Started             VO        Cluster  ID
slate-dev-osg-frontier-squid-proxy-test 2018-Aug-03       slate-dev uchicago Instance_dd427321-05f5-
                                      18:17:30 UTC                         42a2-b61c-e21169187188

Services:
Name                       Cluster IP     External IP     Ports
osg-frontier-squid-global 10.107.134.230 192.170.227.202 3128:32398/TCP

Configuration: (default)
```

-like

Summary

- Reduce barriers to supporting collaborative science
- Give science platform developers a ubiquitous "CI substrate"
- Change distributed cyberinfrastructure operational practice by mobilizing capabilities in the edge
- Developing the DevOps model, provider concerns and policies, tooling to give developers consistent environment
- First k8s-based WAN deployments underay:
  - caching networks for OSG (StashCache) and ATLAS at CERN (XCache)

Thank you!

slateci.io