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









The Open Science Grid



Example

Vonamae Ene Company ona



Facilitator for "data lake" R&D

Example



Caching network for IceCube & LIGO

Example



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



Campus or Institute HPC resources

SLATE Platform Operators & Science VO Managers





The Open Science Grid



Worldwide LHC Computing Grid



Caching network deployed for IceCube & LIGO



Services Layer At The Edge

- 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

- 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





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

App Version Chart Version Description Name v0.8.1 v0.7-dev Multi-user Jupyter installation iupvterhub osg-frontier-squid squid-3 0.2.0 A Helm chart for configuration and deployment o... osiris-unis 0.1.0 Unified Network Information Service (UNIS) 1.0 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-proxytest with ID Instance_dd427321-05f5-42a2-b61c-e21169187188

\$ slate-client instance info Instance_dd427321-05f5-42a2-b61c-e21169187188				
Name	<u>Started</u>	<u>V0</u>	<u>Cluster</u>	ID
<pre>slate-dev-osg-frontier-squid-proxy-test</pre>	2018-Aug-03	slate-dev	uchicago	Instance_dd427321-05f5-
	18:17:30 UTC			42a2-b61c-e21169187188

Services:

NameCluster IPExternal IPPortsosg-frontier-squid-global10.107.134.230192.170.227.2023128:32398/TCP

Configuration: (default)

\$ slate-client app list



-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



"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: <u>nco@nitrd.gov</u>, Website: <u>https://www.nitrd.gov</u>

