



Towards a Lightweight Operations Model

Lincoln Bryant

US ATLAS Computing Facilities Meeting

8/7/19

Introduction



- Want to move toward a model where site operations and deployment can be operated simply and centrally
- SLATE a purpose-built tool and platform that can get us closer to that "NoOps" model
- Many challenges, may not fully get there, but we can take steps in that direction

General SLATE News



- New applications in the stable catalog
 - perfSONAR Testpoint
 - GridFTP
 - Globus Connect v4
 - Faucet SDN controller
- HTCondor-CE close to being ready
- New clusters at: Clemson, UWashingtton, IU
- Other applications, improvements here:

<https://portal.slateci.io>

SLATE status at US sites



- AGLT2 (UM), MWT2 (UC) have SLATE hardware
 - XCache up and running, being tested and managed by Ilija
- NET2, SWT2 planning to purchase edge nodes
- At MWT2 we are starting to migrate existing site services to SLATE

Vulnerability patching



- OSG announced a vulnerability in Frontier Squid on July 26
 - you *have* all updated right? 😊
- SLATE instances were all updated within the hour, simply via:

```
for i in $(slate instance list | grep squid | awk '{print $4}');  
do  
    slate instance restart $i  
done
```

Vulnerability patching



- Coordination with site admins wasn't needed
- No special actions taken on the SLATE catalog, using OSG's upstream for Squid:

```
containers:
```

```
- name: osg-frontier-squid  
  image: opensciencegrid/frontier-squid:fresh
```

(we did break CVMFS running *inside* of the SLATE nodes, though. Oops!)

Proposal



- We would like to experiment **operating all Squids** in the US Tier 2 facility with SLATE
 - Existing squids would remain in place as backup
- A few configuration changes needed:
 - SLATE deployed on edge node at each T2
 - CVMFS/http_proxy configured to use SLATE squid
 - AGIS entries updated to point to SLATE squids

SLATE Deployment



- Will need to deploy Kubernetes + SLATE on an edge node
- If you don't yet have an edge node, a VM will work in a pinch
- We have instructions for setting up SLATE + Kubernetes:
 - <http://slateci.io/docs/quickstart/quickstart-cluster-install.html>

Squid Application Deployment



- Once the cluster is registered, you'll want to allow the `atlas-squid` group:

```
$ slate cluster allow-group <your cluster name> atlas-squid
```

- Once done, the "squid team" can deploy and maintain the Squid at your site:

```
$ slate app install osg-frontier-squid --site <your site> ...
```

(n.b. you could manage your own Squid via SLATE, too, if you really want to.

Squid Application Deployment



Some information will need to be communicated to the Squid team, like your worker IP range:

```
Instance: global
Service:
  Port: 3128
  ExternalVisibility: NodePort
SquidConf:
  CacheMem: 128
  CacheSize: 10000
  IPRange: 10.0.0.0/8 172.16.0.0/12 192.168.0.0/16
```

Worker Configuration



```
CVMFS_CACHE_BASE=/scratch/cvmfs/cache
CVMFS_QUOTA_LIMIT=51200
CVMFS_HTTP_PROXY="http://squid.uchicago.slateci.net:3128;http://uct2-cvmfs.mwt2.org:3128;http://iut2-cvmfs.mwt2.org:3128|http://mwt2-cvmfs.campuscluster.illinois.edu:3128;DIRECT"
CVMFS_REPOSITORIES=atlas.cern.ch, atlas-condb.cern.ch, atlas-nightlies.cern.ch, sft.cern.ch, geant4.cern.ch, grid.cern.ch, cms.cern.ch, oasis.opensciencegrid.org, osg.mwt2.org
```

\$OSG_SQUID_LOCATION will also need to be configured in your GK:

```
/etc/osg/config.d/01-squid.ini
```

AGIS Configuration



https://atlas-agis.cern.ch/agis/service/create_frontier/?example=

ATLAS Grid Information System

/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=jstephen/CN=781624/CN=Judith Lorraine Stephen | Logout

RC Site	ATLASSite	DDMEndpoint	PANDA Queue	Service	Central Services
DDM Groups					

Frontier Service Creation

Docs TWiki OLD

Example: specify Frontier

GOCD/DB/OIM Site:

AGIS Service name:

Endpoint:

Nothing will be changed or saved by clicking this button.

Frontier Service description will be saved, you will be redirected to edit page for specifying additional parameters.

Monitoring and failover



- We'll be monitoring Squid status centrally via ElasticSearch/Kibana
- Should things break, workers/cvmfs should fail-over to your existing squid

Conclusion



- SLATE mature enough to pick up some of the operational load in US ATLAS
- Would like to start small with a forgiving application
- Gain operational experience, understand what would be needed to move more complex software, e.g. HTCondor-CE's, to centralized operation