



# SLATE

Lincoln Bryant  
University of Chicago

GridPP42  
24 Apr 2019



# Motivation

- Remotely manage edge services at sites with central expert teams
- Deploy updates more quickly
- Introduce new services more easily
- Save time and effort for local admins




# Create a federation of edge platforms

- **SLATE: Services Layer At The Edge**
- Distributed service orchestration platform
- Loosely federated, share projects/users/applications across institutions
- Start with a single server and scale as needed
- Good for any site but "**lightweight**" sites might find it particularly useful

# Basic SLATE Architecture

- Kubernetes-based platform for easily deploying, maintaining, and upgrading services across sites
- Federated via a custom server/client overlaying Kubernetes
- Software catalog with push button deployment
- RESTful interface with web portal and CLI client



# Web Portal



The screenshot shows a web browser window with three tabs: 'SLATE GridPP42 - Google Slides', 'SLATE - Instances', and 'SLATE - Dashboard'. The address bar shows 'https://portal.slated.io/dashboard'. The dashboard has a dark header with the SLATE logo, a menu icon, and navigation links for 'Apps', 'Docs', 'Lincoln Bryant', and 'Logout'. The main content area is titled 'Dashboard' and 'Account Summary'. It features six widgets: 'My Instances' (listing various development and production instances with an 'Install an app' button), 'Planned maintenance' (showing 'No upcoming events'), 'Support' (with links to join the SLATE slack channel and sign up for slateci-news), 'Applications' (listing grafana, htcondor, nginx, osg-frontier-squid, stashcache, and xcache with a 'View all apps' button), 'Learn' (with links for installing the SLATE client, basic use, creating a group, registering a cluster, deploying an application, and trying the SLATE sandbox), and 'Clusters' (listing three clusters: 'utah-coreos', 'new-i2testbd', and 'uutah-prod', all marked as 'Reachable' with a 'View all clusters' button).

SLATE

Apps Docs Lincoln Bryant Logout

## Dashboard

Account Summary

### My Instances

- [slate-dev-nginx-default](#)
- [slate-dev-nginx-default](#)
- [slate-dev-htcondor-fiona](#)
- [slate-dev-stashcache-global](#)
- [slate-dev-htcondor-foo](#)
- [atlas-xcache-xcache-global](#)
- [slate-dev-osg-frontier-squid-global](#)
- [atlas-xcache-xcache-global](#)

[Install an app](#)

### Planned maintenance

No upcoming events

### Support

- Join the [SLATE slack channel](#)
- Sign up to receive [slateci-news](#)

### Applications

- [grafana](#)
- [htcondor](#)
- [nginx](#)
- [osg-frontier-squid](#)
- [stashcache](#)
- [xcache](#)

[View all apps](#)

### Learn

- [Install Slate Client](#)
- [Basic Use](#)
- [Creating a Group](#)
- [Registering a Cluster](#)
- [Deploy Application](#)
- [Try SLATE Sandbox](#)


### Clusters

<a href="#">utah-coreos</a>	✓ Reachable
<a href="#">new-i2testbd</a>	✓ Reachable
<a href="#">uutah-prod</a>	✓ Reachable

[View all clusters](#)

# SLATE Lightweight Federation


- Security-conscious
  - SLATE platform is an unprivileged user of your cluster after first-time setup
  - <http://slateci.io/blog/app-sec-edge.html>
- Single endpoint via institutional identity
  - Globus + InCommon
- Simple UNIX-like permissions model
  - Users, Groups
- Ability to whitelist groups and applications
  - e.g. group "atlas-xcache" can run XCache but not Squid
- Join the existing federation or build your own





# Deploying Services with SLATE

- Curated application catalog on GitHub with Jenkins for continuous integration
- A central service expert deploys & operates the application for many sites
- Command line or web interface





# Anatomy of a SLATE Application

- Docker container(s) with the software itself
- Templated YAML file describing the service in terms of Kubernetes objects (Pods, Ingress, Services, etc)
- Configuration parameter file exposed to the operator at deploy time
- Goes through an approval process with the SLATE catalog owners







# Example Application: XCache

## Service:

Port: 1094

ExternalIP: 192.170.227.151

## SiteConfig:

Name: MWT2

AGISprotocolID: 433

## XCacheConfig:

CacheDirectories:

- /scratch

HighWaterMark: 0.95

LowWaterMark: 0.90

RamSize: 16g

BlockSize: 1M

Prefetch: 0

CertSecret: xcache-cert-secret

*Operator view, when deploying applications*



120 lines (112 sloc) | 3.58 KB

```
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: {{ template "xcache.fullname" . }}
5    labels:
6      app: {{ template "xcache.name" . }}
7      chart: {{ template "xcache.chart" . }}
8      release: {{ .Release.Name }}
9      instance: {{ .Values.Instance }}
10 spec:
11   replicas: 1
12   selector:
13     matchLabels:
14       app: {{ template "xcache.name" . }}
15       chart: {{ template "xcache.chart" . }}
16       release: {{ .Release.Name }}
17       instance: {{ .Values.Instance }}
18   template:
19     metadata:
20       labels:
21         app: {{ template "xcache.name" . }}
22         chart: {{ template "xcache.chart" . }}
23         release: {{ .Release.Name }}
24         instance: {{ .Values.Instance }}
25     spec:
26       no
```

*Templated K8S YAML, not  
directly exposed to operators*



# Deployment experience in ATLAS

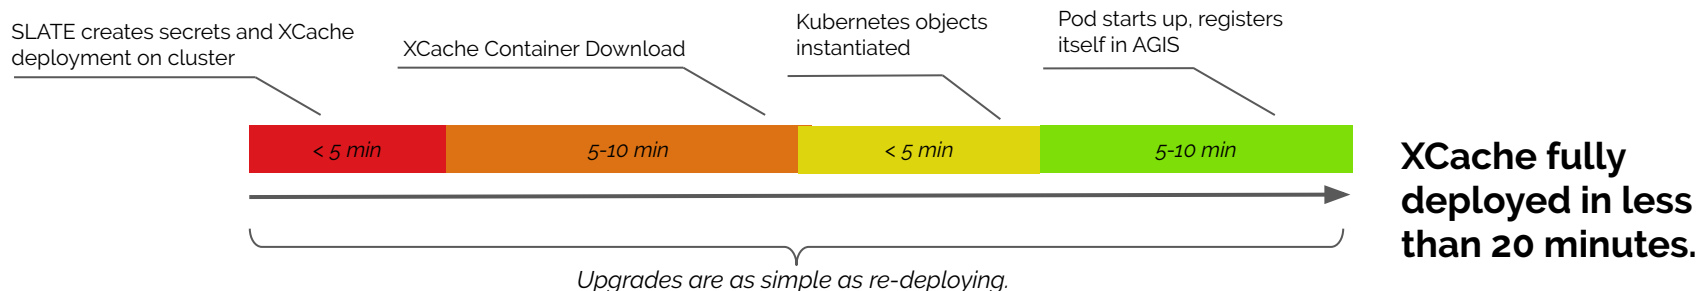
- Goal: build an XCache-based caching network
- SLATE-based deployment will simplify operations and allow for rapid development and debugging
- SLATE services already operational at MWT2, AGLT2, LRZ
- XCache application already in SLATE catalog, frequently being iterated upon, updated at sites, and tested



# XCache deployment process

- Register a cluster with SLATE and allow the `atlas-xcache` group
- Apply a few special extra steps for XCache:
  - Node labeled in Kubernetes (`xcache-capable=true`)
  - One or more storage volumes mounted (e.g. `/xcache`) & communicated to operators
  - Endpoint protocol registered in AGIS
- Test suite containerized
  - Launch a very realistic stress test from Google Compute Engine in minutes

## XCache Deployment & Upgrade Cycle:





SLATE - Instances

https://portal.slateci.io/instances

SLATE

Clusters

Applications

Instances

Groups

CLI Access

Home / Instances

## Application Instances

List of deployed application instances

Show 10 entries Search:

Instance Name	Application	Group	Cluster	Instance ID
atlas-xcache- xcache-global	xcache	atlas- xcache	lrz-atlas	instance_ic4AO8wi9oQ
atlas-xcache- xcache-global	xcache	atlas- xcache	umich-prod	instance_q6fp_YgTbRw
atlas-xcache- xcache-global	xcache	atlas- xcache	um-sc18	instance_N7Hiux2a8l8
atlas-xcache- xcache-global	xcache	atlas- xcache	uchicago-prod	instance_3lL4EUoG4pw
<a href="#">slate-dev-osg- frontier-squid- slatelog</a>	osg-frontier-squid	slate-dev	umich-prod	instance_wVsnbXs5cUw

deployed  
xcaches



# XCache updates

- Even simpler
- Completely transparent to site admin.

```
$ slate instance list  
$ slate instance delete <instance name>  
$ slate app install --group atlas-xcache --cluster uchicago-prod --conf MWT2.yaml xcache
```

Additional benefits:

- Automatic core dump collection (part of XCache)
- Containerized environment makes it easier to debug

# XCache Monitoring

Wealth of information collected even without any direct XRootD monitoring (summary or detailed stream).

Node state (load, mem, network).


Per pod/container event and resource usage.

Logs. Fully searchable.


All info shipped to Elasticsearch at UChicago.

***WIP - Prometheus-based monitoring***


Load



MemoryUsage



Network Traffic



# XCache Logging



## Really convenient logging

- No need to contact anyone
- No need to log in anywhere
- Powerful search
- All the services logs in the same place.
- Kept for 10 days.


Infrastructure / Logs

🔍 kubernetes.pod.name: atlas-xcache-xcache-global-bdc76dbd9-pgqc4

```
2019-03-02 11:02:25.908 Creating proxy done
2019-03-02 11:02:25.908
2019-03-02 11:02:25.969 190302 10:02:25 5263 XrootdXeq: usatlas1.51852:261@192-170-227-156.prometheus-operator-kubelet.kube-system.svc.cluster.local pub IPv4 login
2019-03-02 11:02:25.969 190302 10:02:25 5263 usatlas1.51852:261@192-170-227-156.prometheus-operator-kubelet.kube-system.svc.cluster.local ofs_open: 0-600
2019-03-02 11:02:25.969 190302 10:02:25 5263 usatlas1.51852:261@192-170-227-156.prometheus-operator-kubelet.kube-system.svc.cluster.local fn=/root:/xrootd.aglt2.org:1094/pnfs/aglt2.org/atlasdatadisk/rucio/mc15_13TeV/85/72/EVNT.06479607._035168.pool.root.1
2019-03-02 11:02:26.030 190302 10:02:26 5263 XrdFileCache_Manager: info Cache::Attach()
2019-03-02 11:02:26.030 190302 10:02:26 5263 XrdFileCache_Manager: info Cache::Attach()
2019-03-02 11:02:26.030 190302 10:02:26 5263 XrdFileCache_Manager: info Cache::Attach()
2019-03-02 11:02:26.030 190302 10:02:26 5263 XrdFileCache_Manager: info Cache::Attach()
2019-03-02 11:02:26.135 190302 10:02:26 5263 XrdFileCache_File: info ioActive block_map.size() = 0
2019-03-02 11:02:26.135 190302 10:02:26 5263 XrdFileCache_File: info ioActive block_map.size() = 0
2019-03-02 11:02:26.135 190302 10:02:26 5263 usatlas1.51852:261@192-170-227-156.prometheus-operator-kubelet.kube-system.svc.cluster.local ofs_close: use=1
2019-03-02 11:02:26.135 190302 10:02:26 5263 usatlas1.51852:261@192-170-227-156.prometheus-operator-kubelet.kube-system.svc.cluster.local fn=/root:/xrootd.aglt2.org:1094/pnfs/aglt2.org/atlasdatadisk/rucio/mc15_13TeV/85/72/EVNT.06479607._035168.pool.root.1
2019-03-02 11:02:26.143 190302 10:02:26 5263 XrdFileCache_IO: info IOEntireFile::Detach() 0x856a9700
2019-03-02 11:02:26.144 190302 10:02:26 5263 usatlas1.51852:261@192-170-227-156.prometheus-operator-kubelet.kube-system.svc.cluster.local ofs_close: use=0 fn=dummy
2019-03-02 11:02:26.145 190302 10:02:26 5263 XrootdXeq: usatlas1.51852:261@192-170-227-156.prometheus-operator-kubelet.kube-system.svc.cluster.local disc 0:00:01
```

BETA

▶ Stream live





# Get a feel for it - SLATE "Sandbox"

- Starts a tutorial environment inside a Kubernetes pod with the SLATE client
  - Runs an instance of the SLATE API and exposes the cluster
- Anyone can make a temporary account, try out the command line interface, and deploy a simple web server application

<https://sandbox.slateci.io:5000/>

SLATE

Docs Logout (lincolnb@uchicago.edu)

Welcome to the interactive SLATE sandbox!

Try "slate --help" to get started!

sandbox:~ \$

### slate cluster list

Assuming the client has successfully contacted the SLATE central service, you should see one cluster named 'sandbox'. In a real SLATE federation, you can expect to see many clusters here.

### slate cluster --help

will show you that the cluster subcommand has several other





# SLATE provisioning options

- **SLATE**Lite (for a quick evaluation using Docker):
  - <https://github.com/slateci/slatelite>
- Zero to k8s+**SLATE** script on a bare edge server:
  - Installs everything necessary starting from a fresh CentOS system  
<http://jenkins.slateci.io/artifacts/scripts/install-slate.sh>
- "Managed" install
  - **We** will SSH to your site, set it up, and hand you the configured machine.
- Full install
  - **You** install Kubernetes, download **SLATE** client and register your cluster



# Registering a cluster

```
$ slate cluster create atlas-t2-xyz \  
  --group atlas-xyz-admins \  
  --org "ATLAS Tier 2 XYZ"  
  
$ slate cluster allow-group atlas-xcache
```

- Join a kubernetes cluster to a SLATE federation
  - Specifying the group which will administer it and the organization which owns the resource
- Grant users access to deploy applications on the cluster
  - In this case, just the atlas-xcache group



# SLATE Status & Roadmap

- Current release: v1.0
  - Variety of deployment methods (Shell scripts, SLATE Lite, manual instructions)
  - All Kubernetes objects needed for deploying our driving applications (XCache, Squid, HTCondor CE...) are available in the SLATE API and client
  - Portal functional for basic operations
  - Basic monitoring & logging infrastructure
- Next release:
  - Addressing persistent storage
  - Refactoring application deployment to use "ingress controllers" where applicable
  - Adding centralized DNS management component
  - Designing a user-friendly workflow from the portal
  - Expanded & polished monitoring and logging infrastructure



## If you'd like to try out SLATE

- Homepage: <http://slateci.io/>
- Slack: <https://slack.slateci.io/>
- [Discussion list](#)

Open to collaboration, packaging interesting applications, discussing use cases, etc!

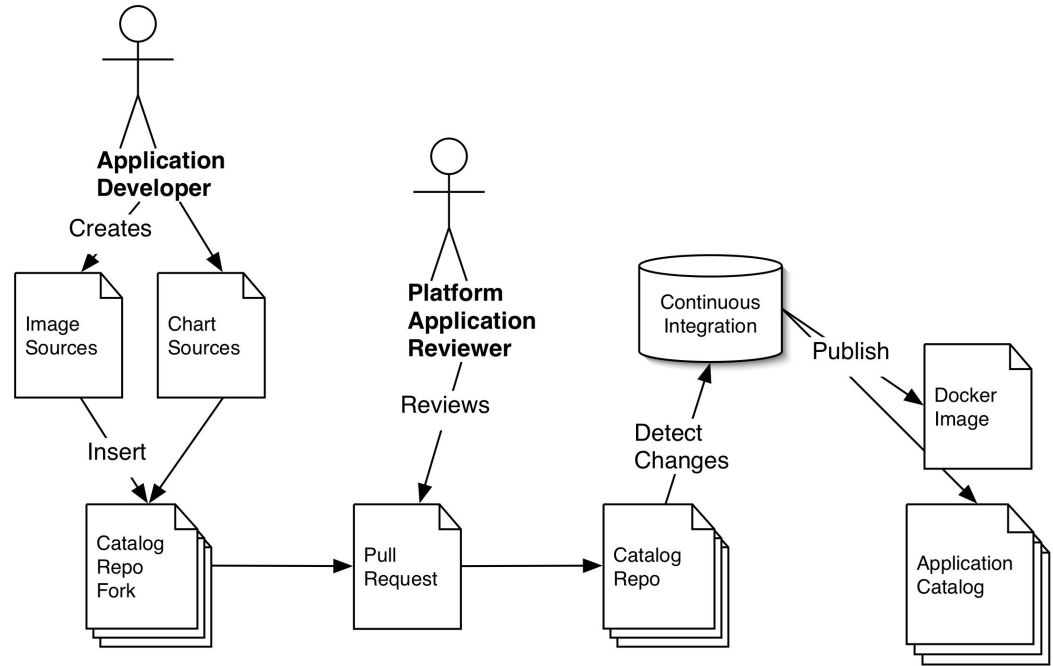


# Extra slides



# Application Security for the Edge

- We have considered the question of meeting sites' cybersecurity policies
- Discussions with community started: <http://bit.ly/app-sec-edge>
- Feel free to directly comment







# XCache deployment process (more details)

- As XCache requires special resources this has to be communicated between Ilija and the site but is done only once:
  - Dedicated node labeled in K8s.
  - Storage should be JBODs organized.
  - Endpoint protocol registered in AGIS.
- Ilija takes over and creates secrets, server, reporting, monitoring, activates protocol in AGIS.
  - All of that is two commands and takes 30 seconds.
- Full update of all the caches in SLATE should take less than 20 min.
  - 10-15 minutes for dockerhub to rebuild image
  - 1 minute to stop running instances
  - 1 minute to start them again
  - 3 minutes to check everything worked
  - Even stress testing is containerized and Ilija can run a very realistic stress test against any xcache in matter of minutes (from Google Computing Engine)


# Monitoring - ES & Kibana



Inbound Traffic  
**227.68KB/s**  
Total Transferred 8.78GB

Outbound Traffic  
**5.89MB/s**  
Total Transferred 229.47GB

In Packetloss  
**72,997**  
Out Packetloss 6,720




Memory usage  
**50.41GB**  
Total Memory 199.84GB


**53**  
Processes



/etc/hosts 27.55%  
/ 0%



user	96.25%
system	40.85%
nice	0.283%
irq	0%
softirq	0.617%
iowait	2.583%



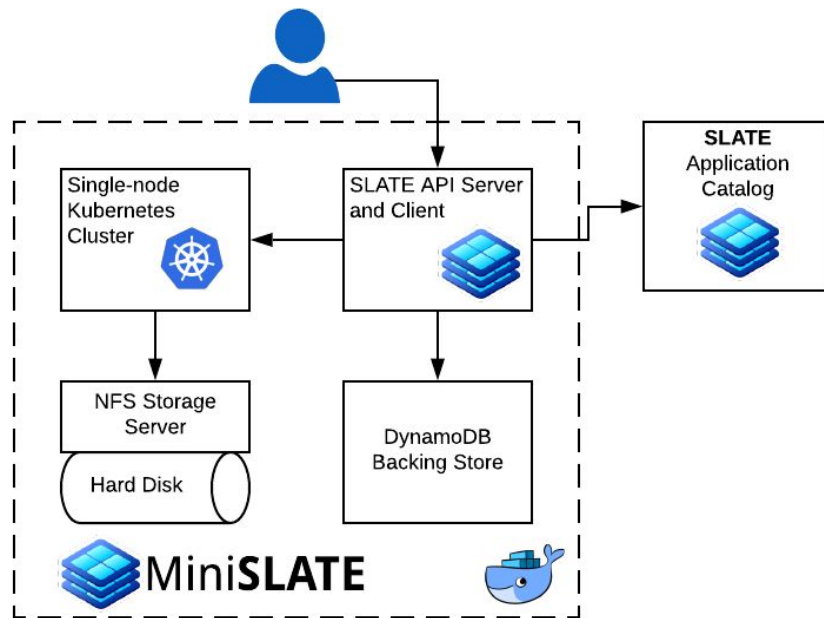
1m	0.967
5m	0.883
15m	0.992



# MiniSLATE

A development environment for SLATE

- **Create a stand alone, miniature SLATE federation for development**
- Follows an Infrastructure as Code pattern
- Enclosed within Docker
  - Little dependency clutter
    - Python, Docker, Docker-Compose
  - Environment consistency
- Completely Destructible
  - Destroy and recreate at will
  - Mount code from host safely
- Batteries Included
  - Full development kit
  - All required software and useful tools are installed when the Docker image is built



# Installing MiniSLATE (<https://github.com/slateci/minislate>)

```
$ git clone https://github.com/slateci/minislate.git
Cloning into 'minislate'...
$ cd minislate
$ ./minislate init
(...)
DONE! MiniSLATE is now initialized.
$ ./minislate slate app install nginx --group ms-group --cluster ms-c

Installing application...
...
Successfully installed application nginx as instance ms-group-nginx-default with ID
instance_tey72YzGYuw
```



# From last year 'til now

- Technology stack has firmed up significantly
  - **CentOS** as platform
  - **Kubernetes** for container orchestration
  - **Helm** for application packaging and deployment
  - Home-grown **SLATE API server, web portal and client**
- SLATE hardware in-place and configured at Chicago, Utah, Michigan
- External customers coming on board (e.g. LRZ)
- Application catalog best practices being developed

**Current explorations**

Challenges are many!

**Cybersecurity and policy**

- On prem additional
- Enga
- Enga

**Scientific**

- Underlay
- Emb
- cyber
- Deve

**Technology investigations**

- Rethinking the stack
  - Startin
  - Up to
- Reusing where p
- and bu
- neede

**Technology examples**

- **RancherOS** - lightweight Linux OS for Containers
  - Linux Kernel + Busybox + Docker = RancherOS
  - Booted from the Network directly to Memory
- **Rancher** - Web interface for managing containers, Kubernetes
- **Kubernetes** (k8s) for container-based orchestration
  - Open source system for deploying and automatically scaling and managing services
  - Currently have 4 sites (Chicago, Michigan, Utah, New Mexico State) with prototype deployments

15