OSG Site Installation Overview

Brian Lin

OSG All Hands 2017

Phase 0: Is the OSG for you?

The OSG Model

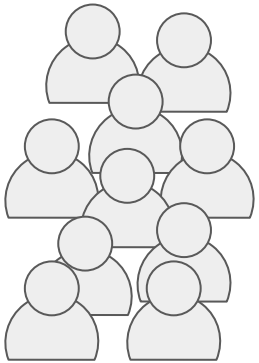

User Submit

Site Gateway

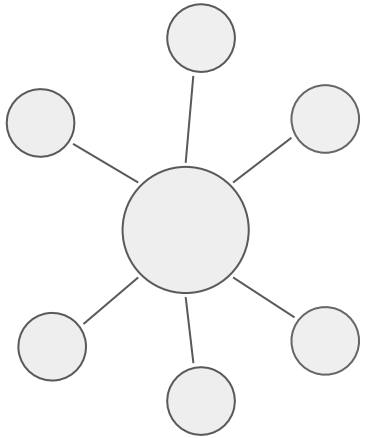

OSG

The OSG Model

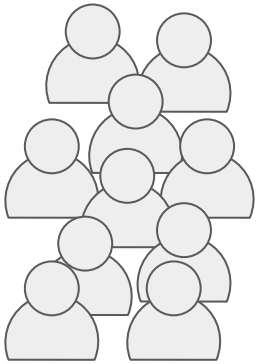

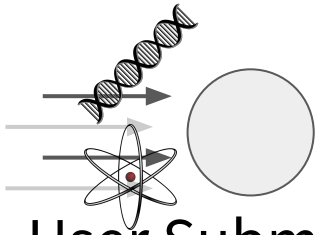

User Submit

Site Gateway

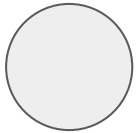

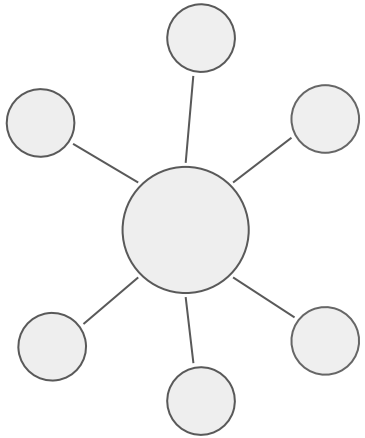

OSG

The OSG Model

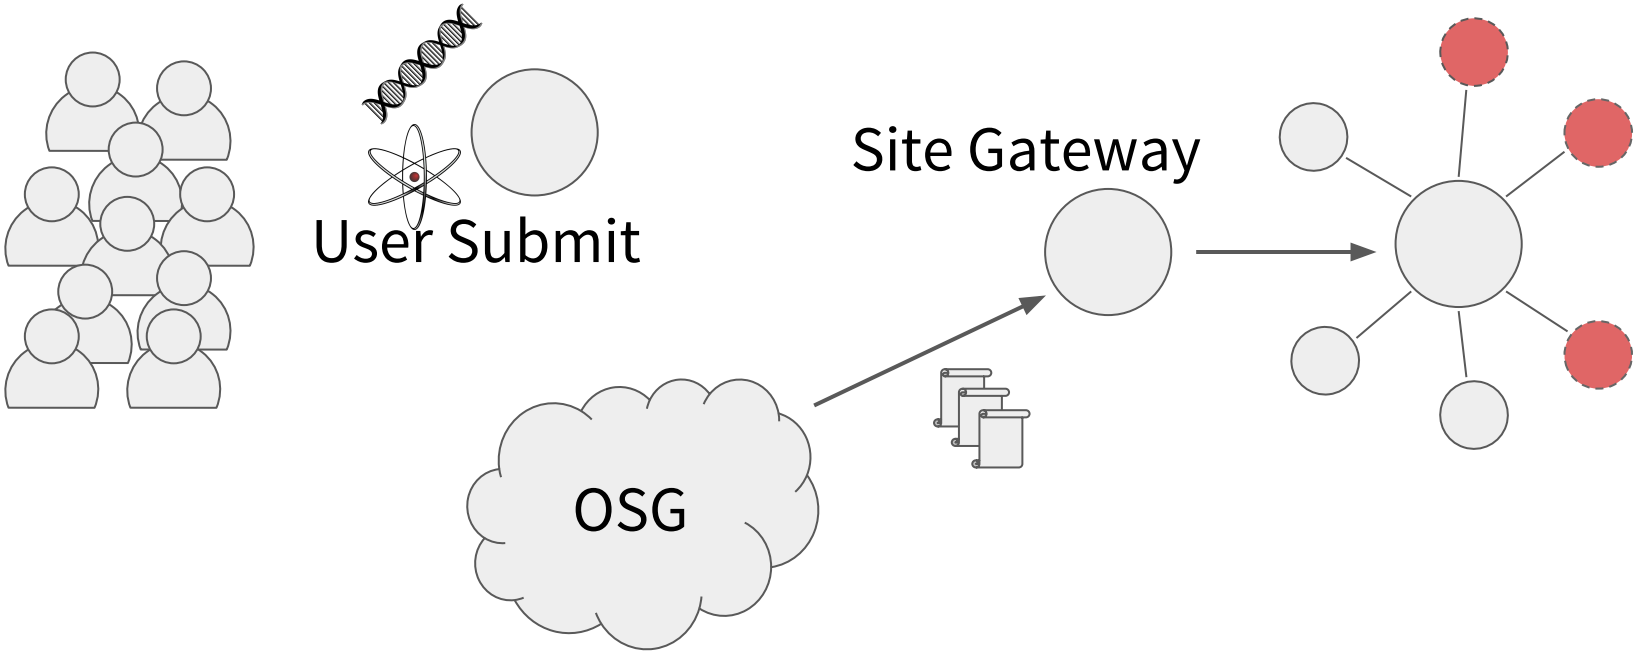

The OSG Model

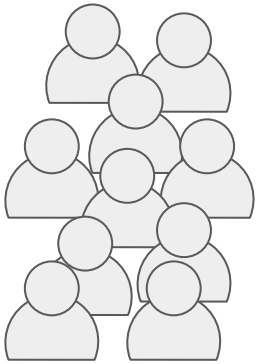

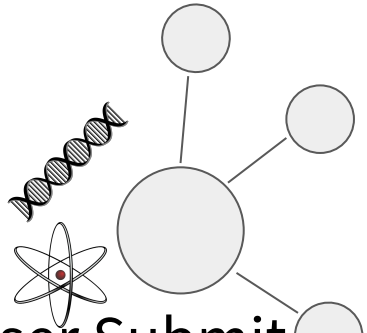

User Submit

Site Gateway

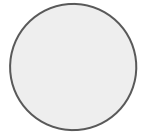

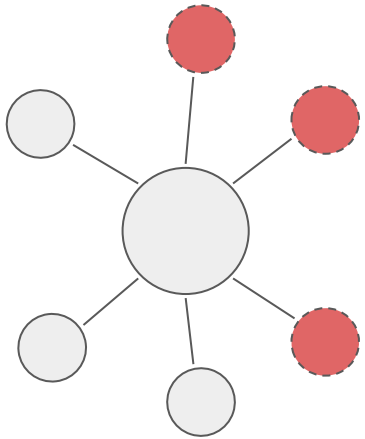

OSG

Base OSG Requirements

- Batch Systems: HTCondor, Slurm, Torque/PBS, LSF, SGE
- Operating Systems: Red Hat Enterprise Linux, CentOS, Scientific Linux
- Outgoing WAN access from worker nodes

Phase 1: Hosted CE or HTCondor-CE?

Hosted CE or HTCondor-CE?

- Do you want > $O(10^4)$ OSG jobs?
- Are you ok with all OSG jobs being submitted as a single user?
- Are there special rules or policies for submitting jobs to your site?
- Do you want to change your configuration frequently?

If you answered no to the above questions, the hosted CE solution could work for you.

Step 1: create user account with submit privileges and SSH access via SSH key

Step 2: If running a non-HTCondor batch system, share the user's home dir with the worker nodes

Still not sure? Ask us at user-support@opensciencegrid.org

OSG-Hosted CE

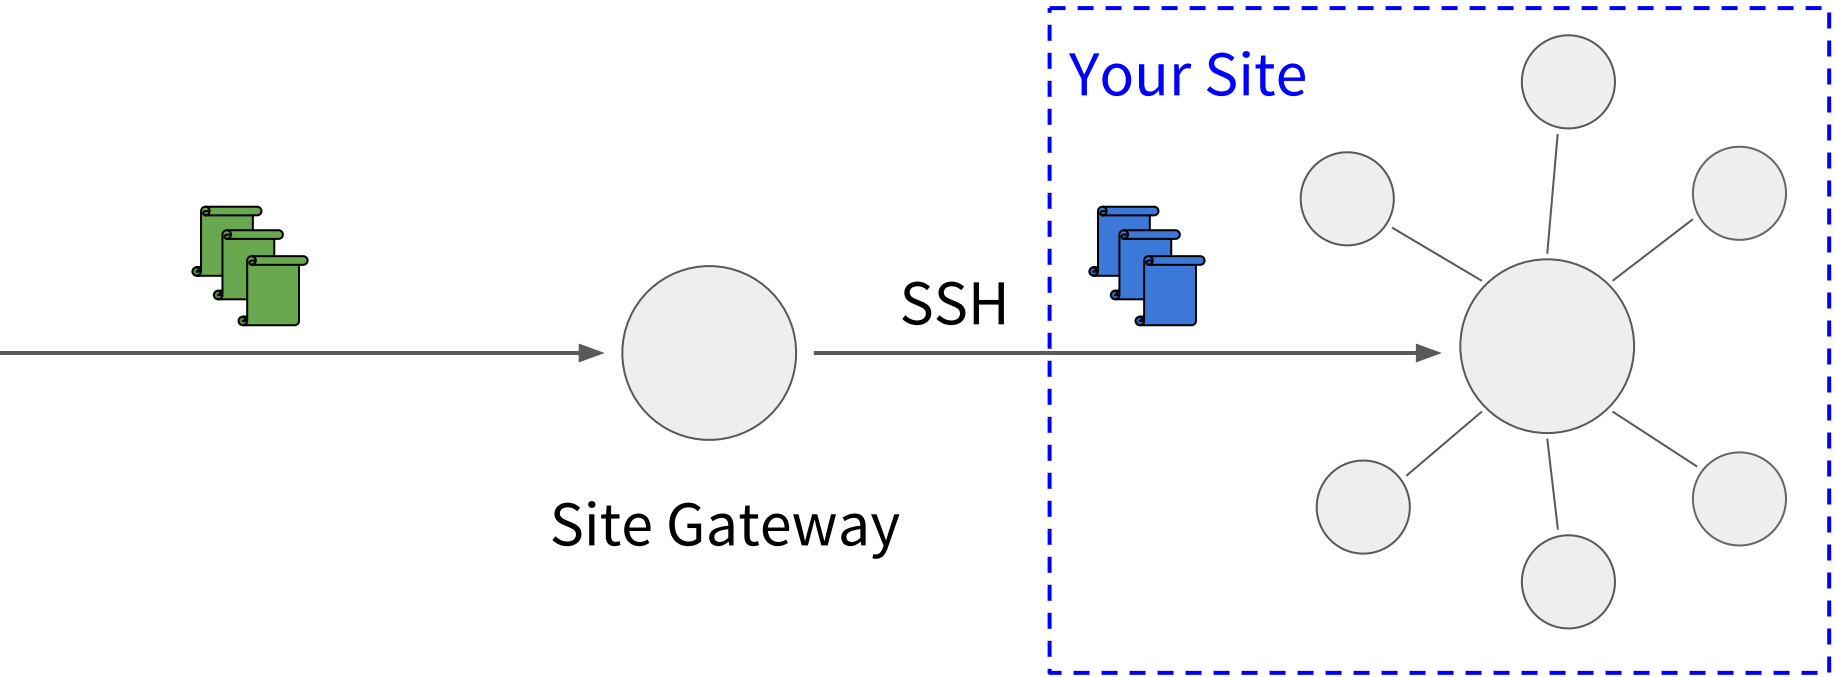

You’re done!

HTCondor-CE

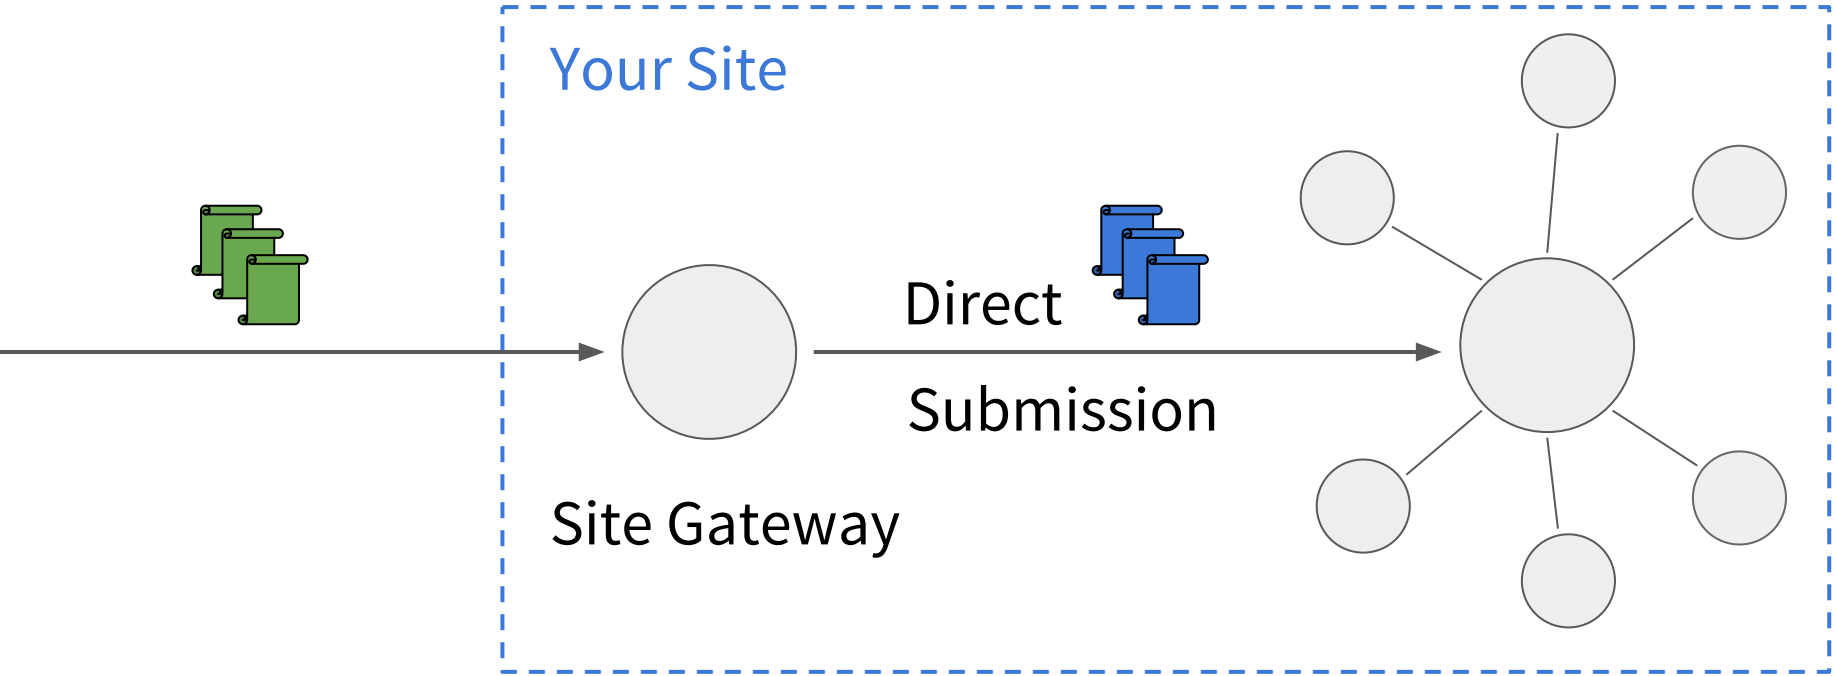

OSG Information Management System

1. Request a user certificate (if you don't have one already)
   Https://oim.opensciencegrid.org/oim/certificaterequestuser
2. Register a facility, site, resource group, and resource if not already in the topology
   https://twiki.opensciencegrid.org/bin/view/Operations/OIMRegistrationInstructions#Facility_Registration
3. Register as a grid administrator
   https://oim.opensciencegrid.org/oim/gridadmin
4. Request a host certificate for your CE
   https://oim.opensciencegrid.org/oim/certificaterequesthost

Questions/Issues? goc@opensciencegrid.org

HTCondor-CE Architecture: HTCondor backend

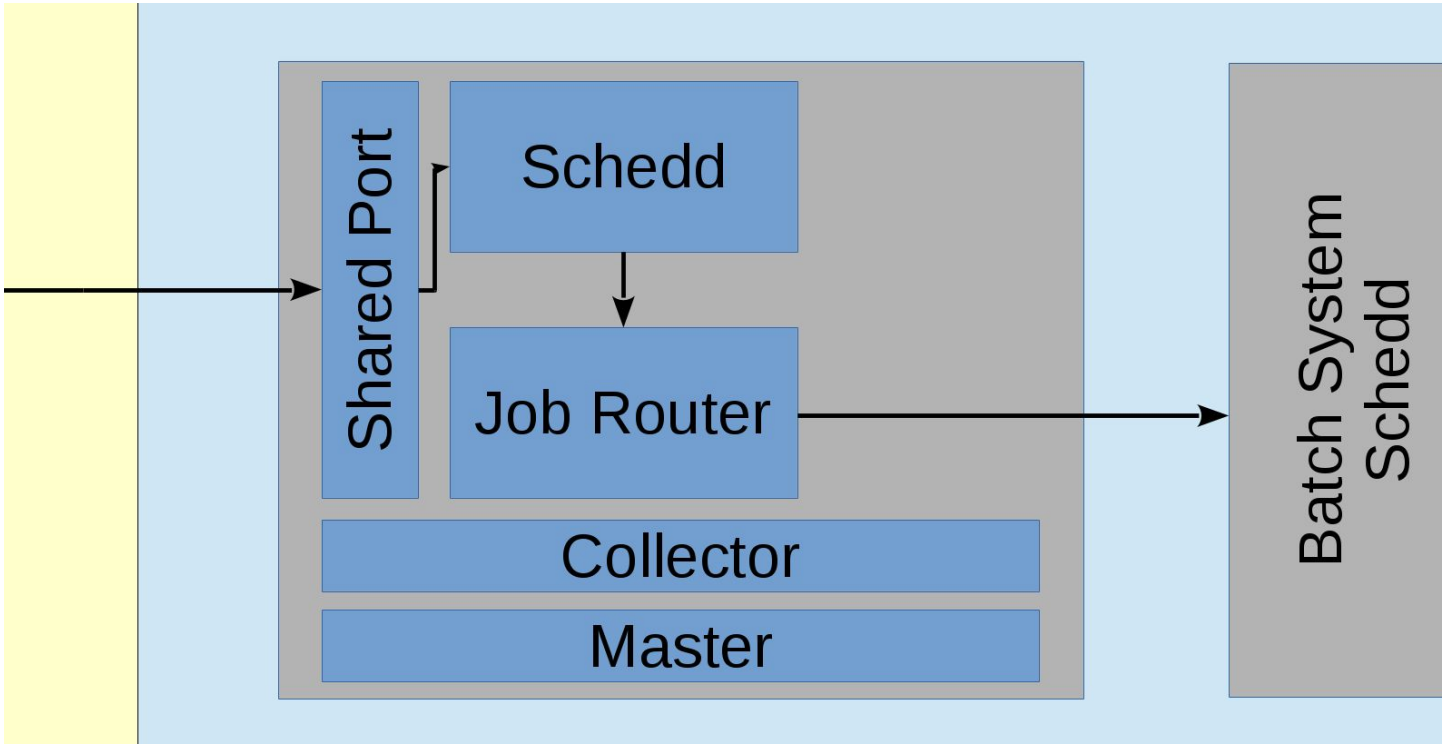

HTCondor-CE: Non-HTCondor backend

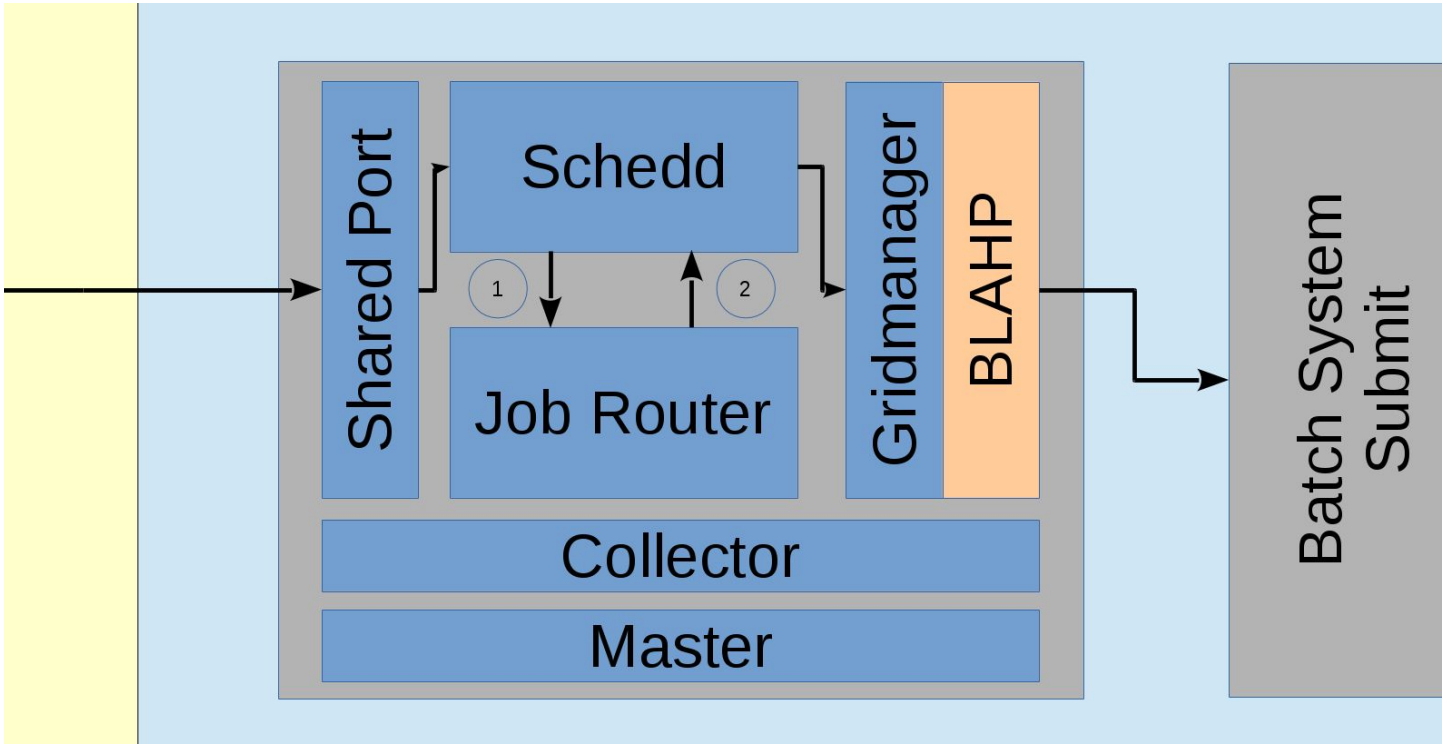

HTCondor-CE Requirements

- Open port (TCP) 9619
- Shared FS for non-HTCondor batch systems for file transfer
- Ensure mapped users exist
- Minimal hardware requirements
  - Handful of cores
  - HTCondor backends should plan on ~½ MB RAM per job
  - Expecting high rates of jobs? HTCondor-CE SPOOL dir should live on an SSD
    - Default `/var/lib/condor-ce/spool` (`condor_ce_config_val -v SPOOL`)
    - Same thing applies for HTCondor backends
      Default: `/var/lib/condor/spool` (`condor_config_val -v SPOOL`)

https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallHTCondorCE

edg-mkgridmap vs GUMS

- Authentication methods
- edg-mkgridmap is simpler, creates `/etc/grid-security/grid-mapfile` that holds a mapping of certificate Distinguished Names to local unix accounts
- Use GUMS only if you know you need it:
  - You want to map users based on rules
  - You need to support multiple VO roles
  - You need to support gLExec for pilot jobs

https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallHTCondorCE#Configuring_authorization

HTCondor-CE Configuration

- `osg-configure -v` and `osg-configure -c` handles most of the configuration
- Most HTCondor-CE configuration goes into the job router
  - Job router filters and transforms incoming grid jobs into “routed” jobs
  - Configured using declarative ClassAds with the JOB_ROUTER_ENTRIES variable
  - Each entry in JOB_ROUTER_ENTRIES is combined with the JOB_ROUTER_DEFAULTS configuration variable to create each job route

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Alice has an HTCondor pool and she wants CMS jobs submitted to her CE to be forwarded to her pool and requesting x86_64 Linux machines and setting the attribute “foo” on her routed job to “bar”. All other jobs should be submitted to the pool without any changes.

```
JOB_ROUTER_ENTRIES = [ \
      name = "condor_pool_cms"; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =?= "cms"; \
      set_requirements = (Arch == "X86_64") && (TARGET.OpSys == "LINUX"); \
      set_foo = “bar”; \
] \
[ \
      name = “condor_pool_other”; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Alice has an HTCondor pool and she wants CMS jobs submitted to her CE to be forwarded to her pool and requesting x86_64 Linux machines and setting the attribute “foo” on her routed job to “bar”. All other jobs should be submitted to the pool without any changes.

```
JOB_ROUTER_ENTRIES = [ \
      name = "condor_pool_cms"; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =?= "cms"; \
      set_requirements = (Arch == "X86_64") && (TARGET.OpSys == "LINUX"); \
      set_foo = “bar”; \
] \
[ \
      name = “condor_pool_other”; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Alice has an HTCondor pool and she wants CMS jobs submitted to her CE to be forwarded to her pool and requesting x86_64 Linux machines and setting the attribute “foo” on her routed job to “bar”. All other jobs should be submitted to the pool without any changes.

```
JOB_ROUTER_ENTRIES = [ \
      name = "condor_pool_cms"; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =?= "cms"; \
      set_requirements = (Arch == "X86_64") && (TARGET.OpSys == "LINUX"); \
      set_foo = “bar”; \
] \
[ \
      name = “condor_pool_other”; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Alice has an HTCondor pool and she wants CMS jobs submitted to her CE to be forwarded to her pool and requesting x86_64 Linux machines and setting the attribute “foo” on her routed job to “bar”. All other jobs should be submitted to the pool without any changes.

```
JOB_ROUTER_ENTRIES = [ \
      name = "condor_pool_cms"; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =?= "cms"; \
      set_requirements = (Arch == "X86_64") && (TARGET.OpSys == "LINUX"); \
      set_foo = “bar”; \
] \
[ \
      name = “condor_pool_other”; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Alice has an HTCondor pool and she wants CMS jobs submitted to her CE to be forwarded to her pool and requesting x86_64 Linux machines and setting the attribute “foo” on her routed job to “bar”. All other jobs should be submitted to the pool without any changes.

```
JOB_ROUTER_ENTRIES = [ \
      name = "condor_pool_cms"; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =?= "cms"; \
      set_requirements = (Arch == "X86_64") && (TARGET.OpSys == "LINUX"); \
      set_foo = “bar”; \
] \
[ \
      name = “condor_pool_other”; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Alice has an HTCondor pool and she wants CMS jobs submitted to her CE to be forwarded to her pool and requesting x86_64 Linux machines and setting the attribute “foo” on her routed job to “bar”. All other jobs should be submitted to the pool without any changes.

```
JOB_ROUTER_ENTRIES = [ \
      name = "condor_pool_cms"; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =?= "cms"; \
      set_requirements = (Arch == "X86_64") && (TARGET.OpSys == "LINUX"); \
      set_foo = “bar”; \
] \
[ \
      name = “condor_pool_other”; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Alice has an HTCondor pool and she wants CMS jobs submitted to her CE to be forwarded to her pool and requesting x86_64 Linux machines and setting the attribute “foo” on her routed job to “bar”. All other jobs should be submitted to the pool without any changes.

```
JOB_ROUTER_ENTRIES = [ \
      name = "condor_pool_cms"; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =?= "cms"; \
      set_requirements = (Arch == "X86_64") && (TARGET.OpSys == "LINUX"); \
      set_foo = “bar”; \
] \
[ \
      name = “condor_pool_other”; \
      TargetUniverse = 5; \
      Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Cameron has a PBS pool and she wants CMS jobs submitted to her CE to be forwarded to her pool. All other jobs should be submitted to her pool without any changes

```
JOB_ROUTER_ENTRIES = [ \
     name = "pbs_pool_cms"; \
     TargetUniverse = 9; \
     GridResource = "batch pbs"; \
     Requirements = target.x509UserProxyVOName =?= "cms"; \
] \
[ \
     name = "pbs_pool_other"; \
     TargetUniverse = 9; \
     GridResource = "batch pbs"; \
     Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

Cameron has a Slurm pool and she wants CMS jobs submitted to her CE to be forwarded to her pool. All other jobs should be submitted to her pool without any changes

```
JOB_ROUTER_ENTRIES = [ \
     name = "slurm_pool_cms"; \
     TargetUniverse = 9; \
     GridResource = "batch slurm"; \
     Requirements = target.x509UserProxyVOName =?= "cms"; \
] \
[ \
     name = "slurm_pool_other"; \
     TargetUniverse = 9; \
     GridResource = "batch slurm"; \
     Requirements = target.x509UserProxyVOName =!= "cms"; \
]
```

https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes

HTCondor-CE Monitoring

- For graphs showing pilot jobs and CE load
- `yum install condor-ce-view`
- Configuration lives in /etc/condor-ce/config.d/05-ce-view.conf
  - Uncomment `DAEMON_LIST`
  - Defaults to port 80 but can be configured by changing `HTCONDOR_VIEW_PORT`
  - Restart `condor-ce` service after config changes

https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallHTCondorCE#CeView

Validation

- Run as regular user with certificate on CE

  ```
  $ voms-proxy-init
  $ condor_ce_trace -d `hostname`
  ```

- Not working? Consult the troubleshooting guide:
  https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/TroubleshootingHTCondorCE
- Still stuck?
  goc@opensciencegrid.org

Phase 2: Preparing your worker nodes

OSG Worker Node Client

- Thin collection of software necessary for pilot job execution
- Available via RPM package, tarball, docker image (new!), and OASIS
  - RPM: https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallWNClient
  - Tarball: https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallWNClientTarball
  - OASIS: https://twiki.grid.iu.edu/bin/view/Documentation/Release3/UsingOSGWnClientFromOASIS

OSG Worker Node Requirements

- Outgoing WAN access!
- OSG worker node client
- Pilot job temp space (OSG_WN_TMP)
  - Set by `worker_node_temp` configuration in `/etc/osg/config.d/10-storage.ini` on the CE
  - 10GB disk/core minimum
  - Site responsible for cleanup, e.g. `tmpwatch`
- Cleanup `/tmp` (recommendation)

Example site

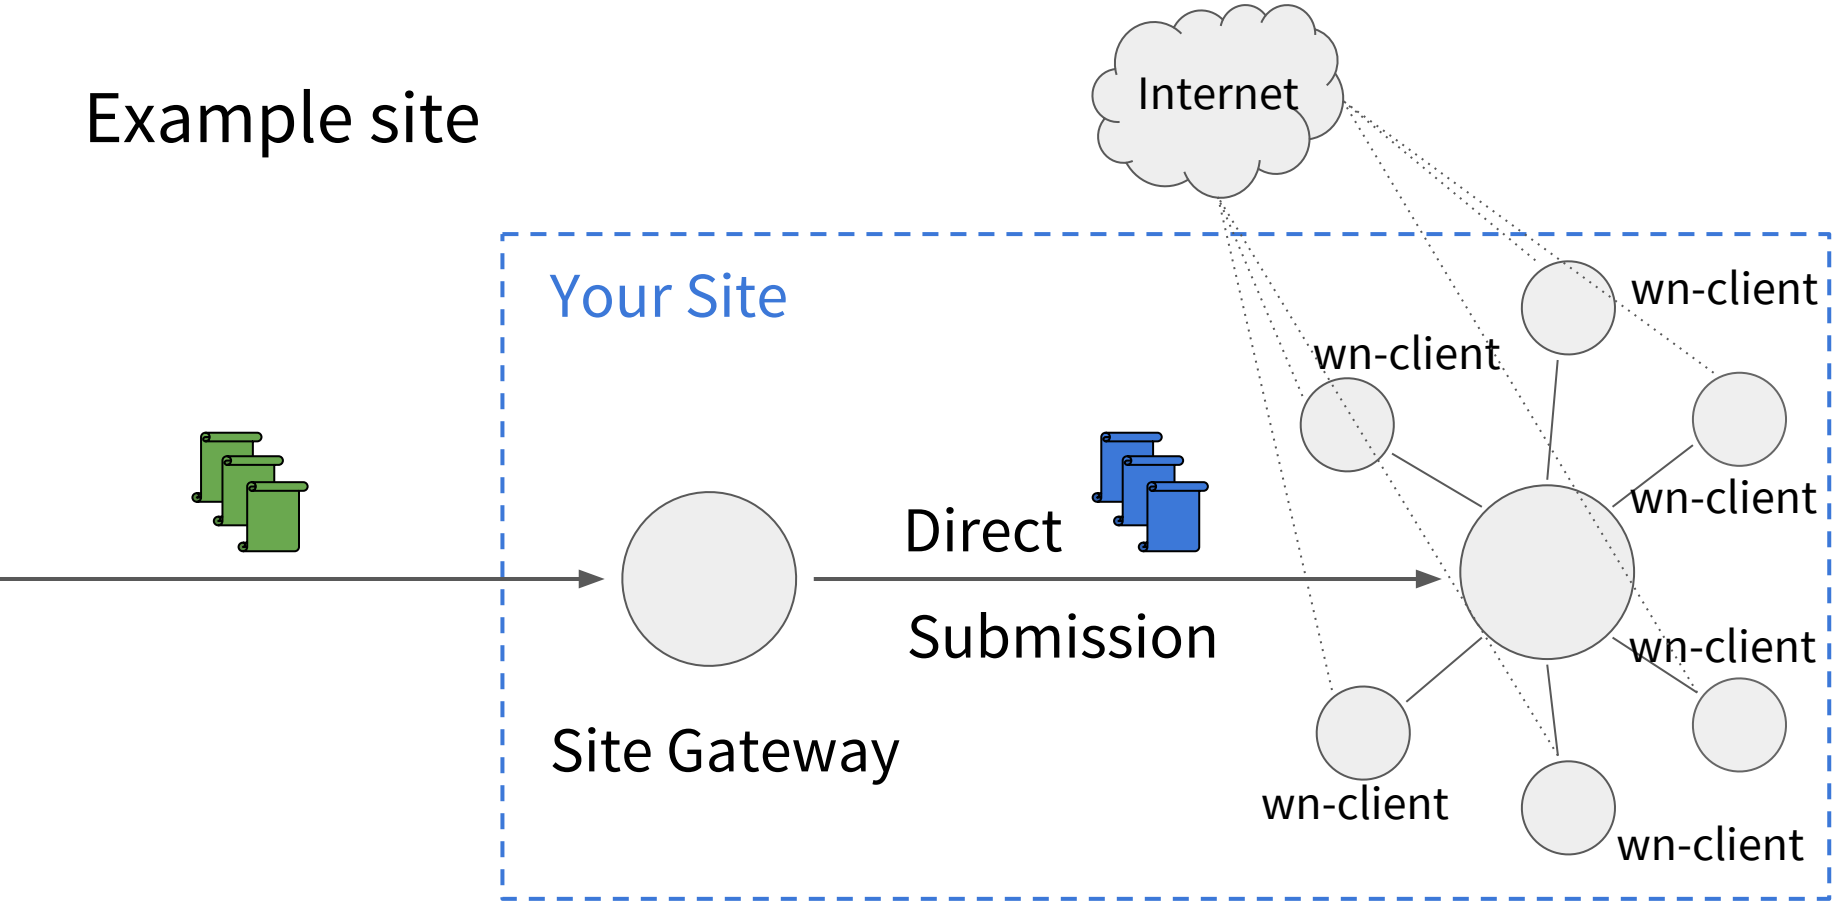

Validation: Request test pilots

osg-gfactory-support@physics.ucsd.edu

OSG Application Software Installation Service

- Software distribution over CernVM File System (CVMFS), which uses http. Requires Squid proxy node:
  https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/InstallCvmfs
  https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/InstallFrontierSquid
- More and more jobs in the OSG want CVMFS
- Optional but recommended

Example site with OASIS

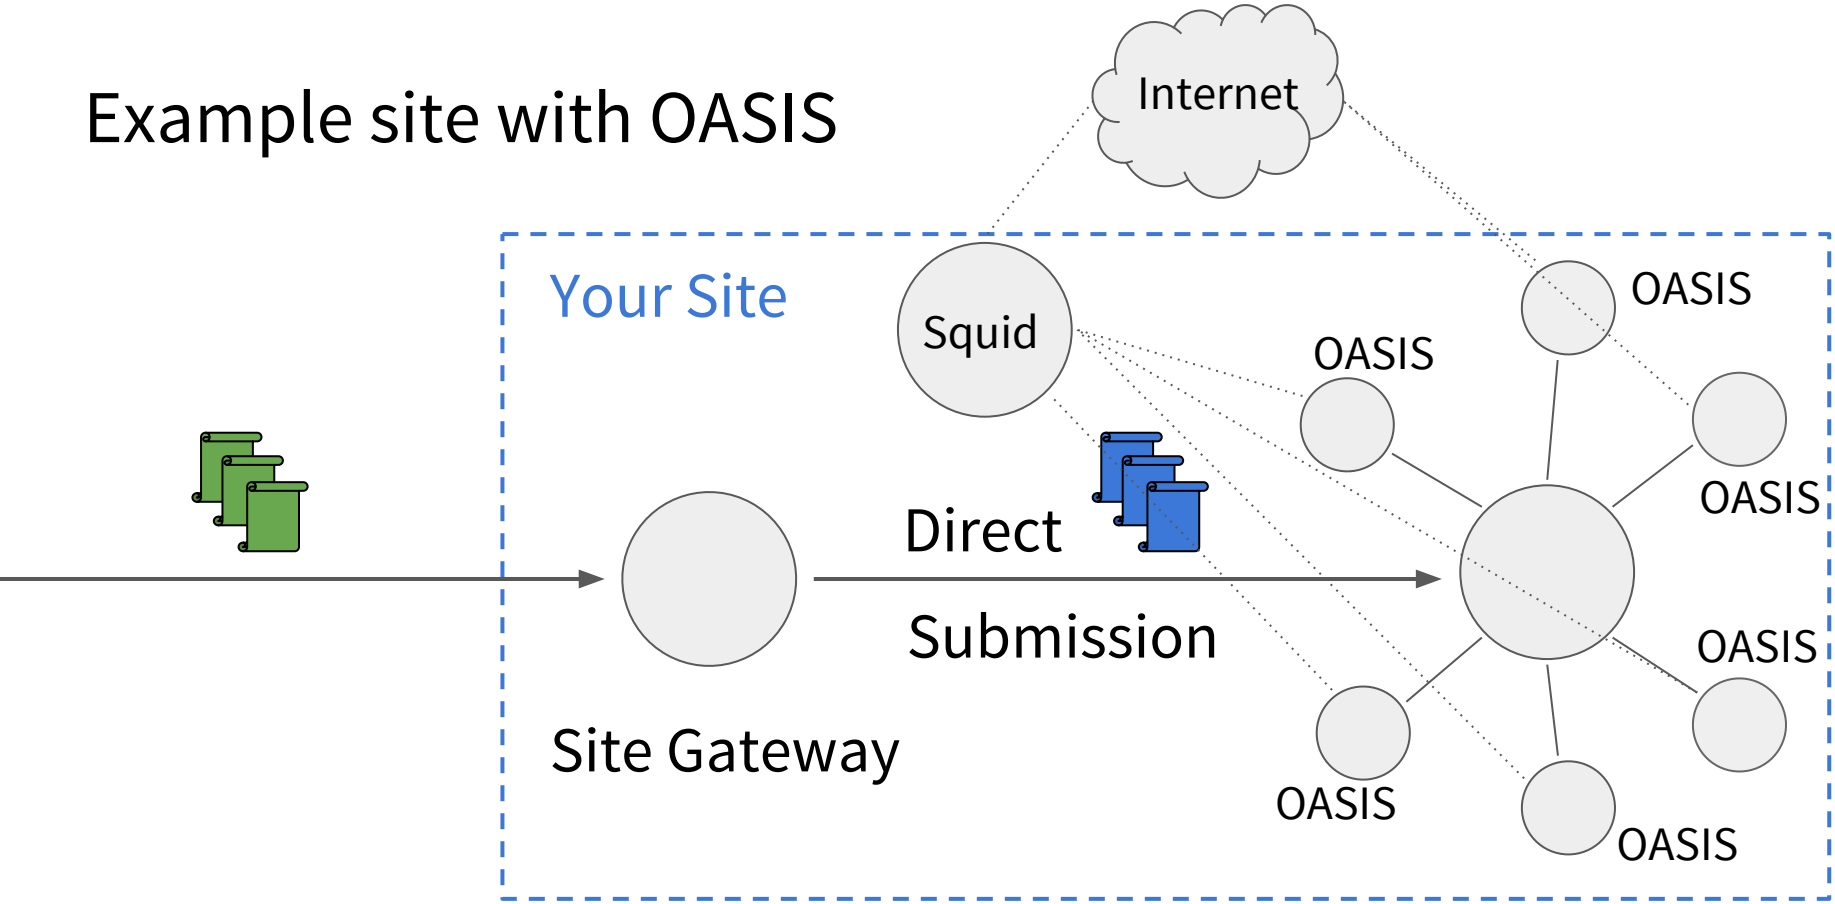

Summary: Decision points

- OSG-Hosted CE vs HTCondor-CE; if hosted CE, you're done!
- edg-mkgridmap vs GUMS on HTCondor-CE
- osg-wn-client installation method
- Optional but recommended OASIS on worker nodes

Summary: Networking

- Open outbound WAN access from worker nodes
- Open port 9619 (TCP) on HTCondor-CE

Summary: Links

- OIM: http://oim.opensciencegrid.org/
- HTCondor-CE installation guide: https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallHTCondorCE
- HTCondor-CE job router configuration guide: https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/JobRouterRecipes
- HTCondor-CE troubleshooting guide: https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/TroubleshootingHTCondorCE
- osg-wn-client installation guides:
  - RPM: https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallWNClient
  - Tarball: https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallWNClientTarball
  - OASIS: https://twiki.grid.iu.edu/bin/view/Documentation/Release3/UsingOSGWnClientFromOASIS
- OASIS installation guide: https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/InstallCvmfs
- Squid installation guide: https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/InstallFrontierSquid

Interested in an OSG-Hosted CE?

user-support@opensciencegrid.org

Want pilot jobs?

osg-gfactory-support@physics.ucsd.edu

Issues?

goc@opensciencegrid.org

Questions?