



VMCondor

Simple batch access
to OpenStack

Andrew McNab

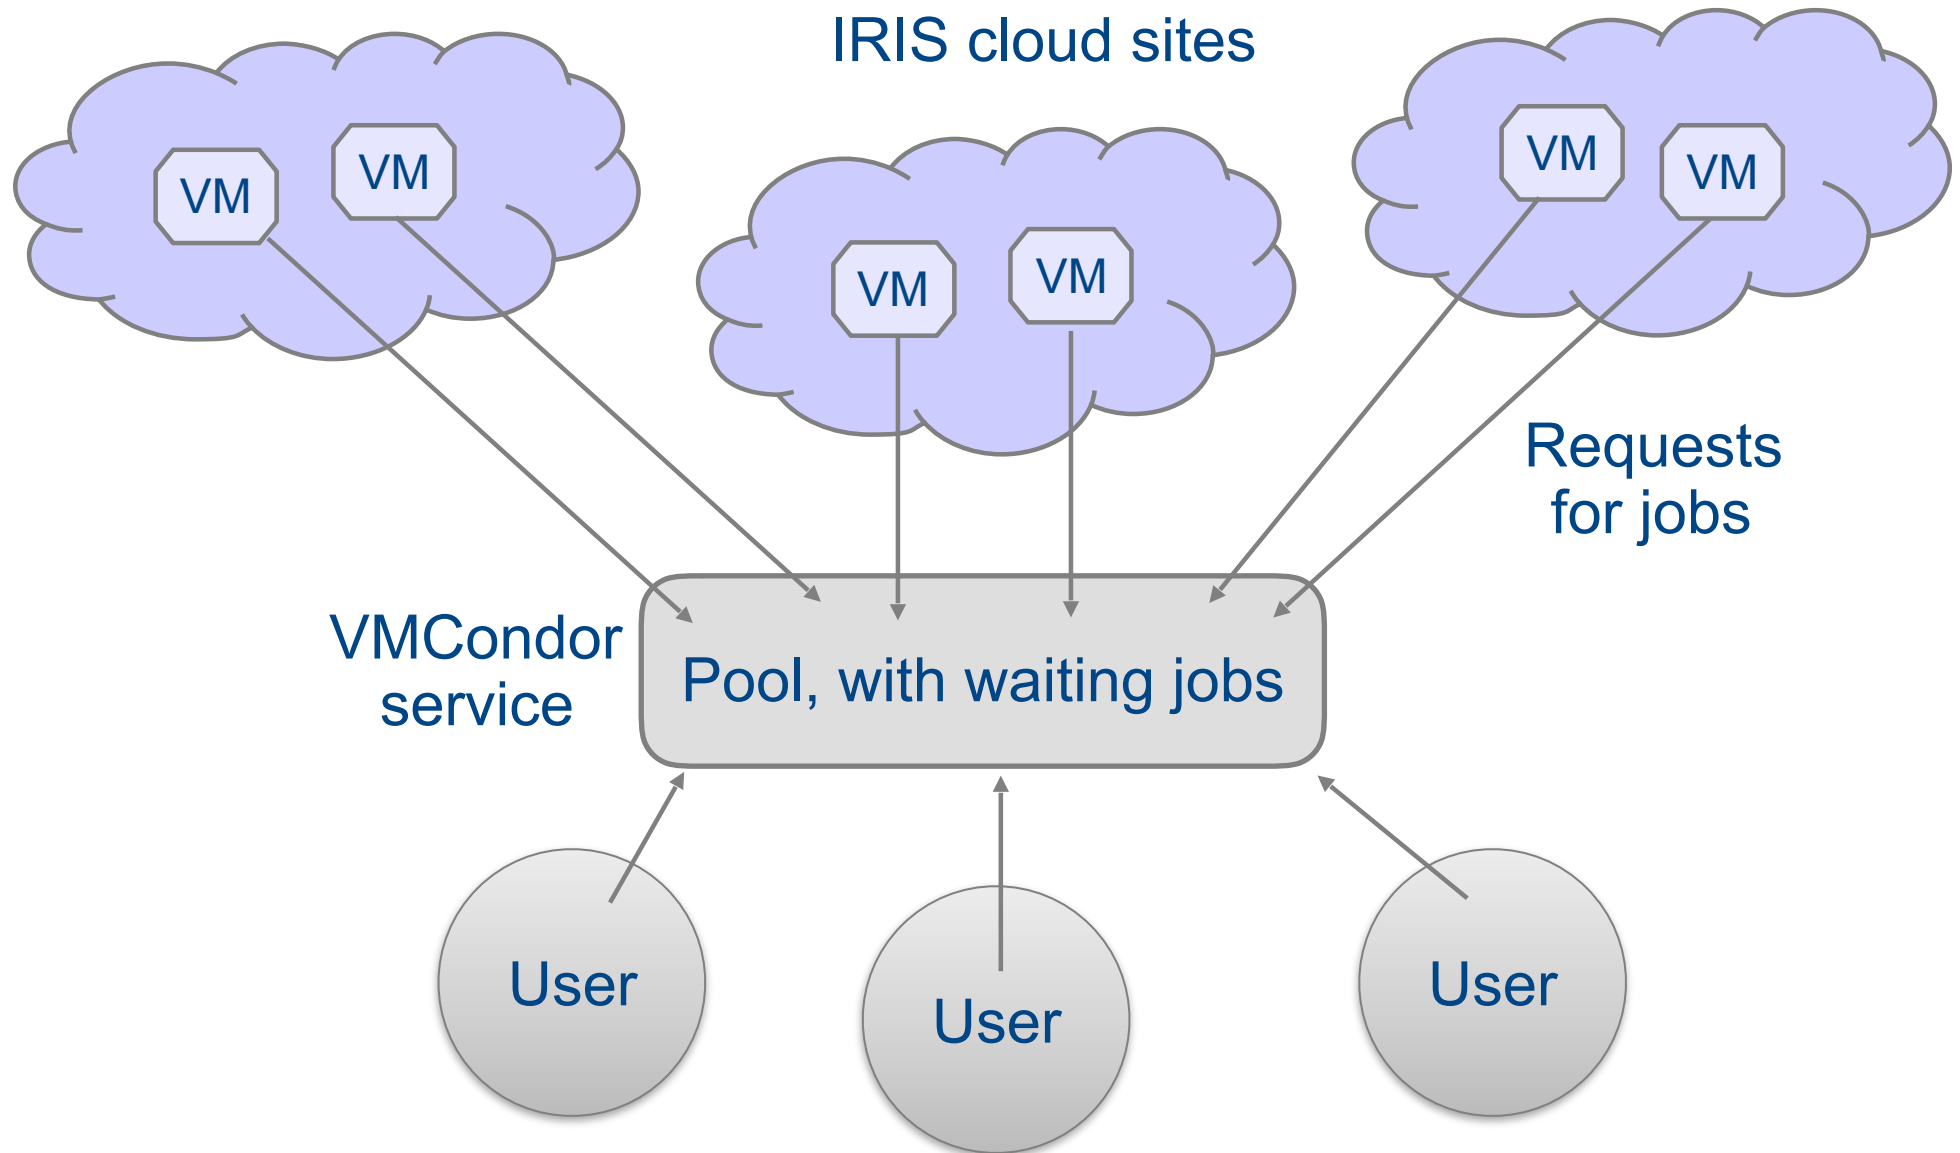
University of Manchester



VMCondor: Simple Access to IRIS

- We discussed ideas for “Simple Access” at the start of the year
- Several ideas, but common theme of replicating the ease of use of a local batch system
 - No need for X.509 user certificates
 - No need to install lots of “special” middleware
- VMCondor is part of the Vcycle Digital Asset funded by IRIS
 - *“VMCondor allows users to submit HTCondor batch jobs to a common pool of virtual machines managed by Vcycle/OpenStack or Vac at multiple IRIS sites”*
 - No need to get a grid certificate
 - Minimal requirements on client machines

VMCondor architecture



Client installation

- We need Linux, including openssl and curl;
 - and the HTCondor commands;
 - and the vmcondor_setup script
- Three routes:
 - Install HTCondor and VMCondor RPMs
 - or tar file install of HTCondor and fetch vmcondor_setup file
 - or iris-vmcondor image from Docker Hub

```
docker run --volume $HOME/iris-vmcondor-home:/home/user  
--tty --interactive irisacuk/iris-vmcondor:latest
```

- No need to start HTCondor daemons or edit config files
- <https://github.com/iris-ac-uk/iris-vmcondor/wiki/Install-Guide>

Getting started

- Run the `vmcondor_setup` script
- It contacts the VMCondor server and creates your `vmcondor.key` file
- Email `vmcondor@iris.ac.uk` to get authorized to submit jobs
- See <https://github.com/iris-ac-uk/iris-vmcondor/wiki/Users-Guide>

```
vmcondor_setup
```

```
Creating /home/user/.vmcondor
```

```
Your key has been placed in /home/user/.vmcondor/vmcondor.key
```

```
Your VMCondor username is user0044
```

To activate your account, send an email to `vmcondor@iris.ac.uk` stating

- your VMCondor username (user0044)
- your preferred name
- your home institute
- your IRIS user community

within the next 7 days.

Running a job

- Once you are authorized you can run jobs straight away.
 - No need to download any more files, run more scripts etc
- Commands are the same as HTCondor:
 - `vmcondor_version`
 - `vmcondor_submit`
 - `vmcondor_q`
 - `vmcondor_transfer_data`
- To run a job, do something like this (“cluster” = a group of jobs):

```
vmcondor_submit example.sub
Submitting job(s).
1 job(s) submitted to cluster 66.
```

Job submit files

- Jobs are described by .sub files:

```
Executable           = example.sh
Output               = example.stdout
Error                = example.stderr
+Owner               = undefined
Notification         = never
When_To_Transfer_Output = on_exit
Queue
```

- Additional options allow you to copy multiple input files to the execution VM; to retrieve output files; and to target jobs to particular sites
- Giving a number to Queue allows multiple identical jobs to be run
- The HTCondor documentation explains the .sub file format

Monitoring jobs and retrieving output

- `vmcondor_q` groups (batches) jobs: “-nobatch” turns this off
- -long options outputs a lot information about each job
- Job outputs are spooled on the VMCondor server until fetched

```
vmcondor_q -nobatch
-- Schedd: vmcondor@vmcondor.iris.ac.uk : <195.194.108.226:9618?...
  ID      OWNER      SUBMITTED      RUN_TIME ST PRI SIZE CMD
  66.0    user0026      8/11 23:37      0+00:00:01 C  0    0.0 example.sh
Total for query: 1 jobs; 1 completed, 0 removed, 0 idle, 0 running, 0 held
```

```
vmcondor_transfer_data 66
  Fetching data files...
```

```
cat example.stdout
USER=cndrusr1
TEMP=/scratch/condor/dir_6031
...
```


Running vmcondor VMs at sites

- The VMs can be managed on a site's OpenStack by the IRIS Vcycle service at Manchester
 - Or the site can run Vcycle to manage the VMs
 - Or the site can run Vac to create the VMs without OpenStack
 - In all cases, you can run other vacuum-style VMs in the same tenancy (eg SKA, LSST, DUNE, LZ, ATLAS, LHCb, ALICE, ...)
- Need to get a UK e-Science certificate to give to the VMs
- Otherwise the VMs are self contained and defined by this in the Vcycle or Vac configuration file:

```
[vacuum_pipe vmcondor]
```

```
vacuum_pipe_url = https://repo.iris.ac.uk/vmcondor/latest/vm/vmcondor.pipe
```

```
target_share = CHANGEME
```

- See <https://github.com/iris-ac-uk/iris-vmcondor/wiki/Sites-Guide>



Bulk data

- You can already embed a credential in the job script or input files and use it to read/write to suitable storage
 - This is a natural way to access cloud storage
- The VMs do have an X.509 grid proxy which they use to talk to HTCondor
 - So they could delegate a suitable proxy with VOMS extensions to the user job to access grid storage
 - No work for the users; just needs enabling by sites
 - Use subgroups of the iris.ac.uk VO for this?
- For large scale projects with multiple users then something like DIRAC is needed
 - We shouldn't add more and more to VMCondor to duplicate this



Summary and next steps

- VMCondor provides an easy way to “run batch jobs on IRIS” without getting grid certificates
- Can be offered to user communities who want to avoid a larger overhead
- Would like to recruit more sites to run the VMs
 - Currently Manchester, Cambridge, RAL, Glasgow
- Can add access to grid storage if needed
- Need to decide how to do accounting
 - Currently all logged under the iris.ac.uk VO
- Should be possible to base identity on IAM
 - Self-service website to activate your VMCondor key?