



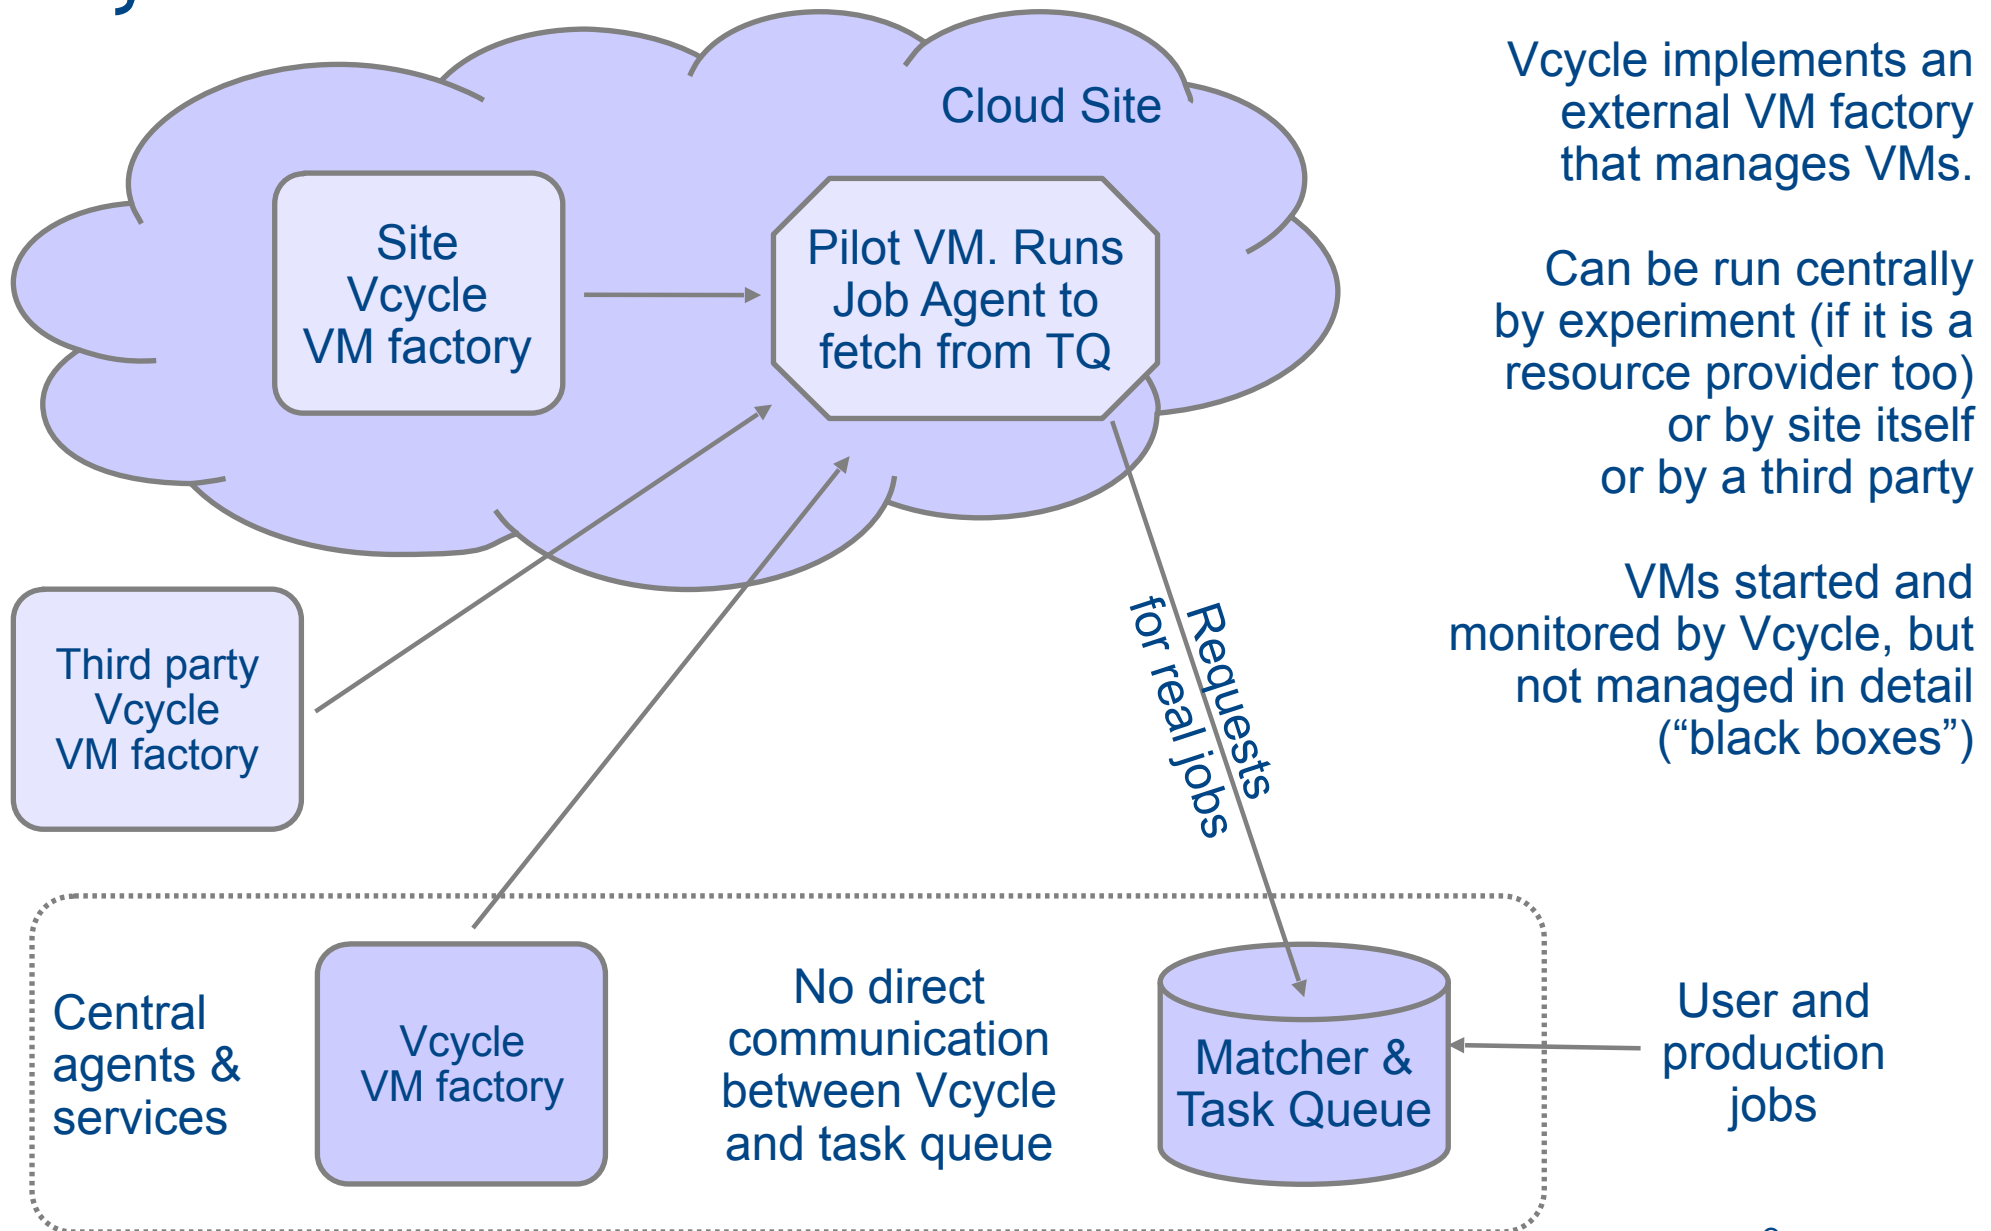
# Publishing Vcycle and Vac state in GLUE2/JSON/HTTPS

**Andrew McNab**  
University of Manchester  
GridPP and LHCb

# Overview

- Vcycle and Vac are VM lifecycle managers
  - They use user\_data and boot image from experiments to create VMs on OpenStack etc (Vcycle) or autonomous hypervisors (Vac)
  - Decide which types of VM to create based on outcomes
  - “Don’t create VMs that keep stopping because there’s no work”
- Vac used in the UK so far
- Vcycle also used by LHCb to manage CERN OpenStack VMs, and by IT/SDC for the 4 LHC experiments on DBCE in the cloud procurement
  - See Laurence’s talk at the GDB last week on Helix Nebula results
- In this Vacuum model, the jobs system doesn’t push VMs into sites
  - Everything is decided by the resource provider, and not the same need for load information as with a pilot factory
  - However, some configuration info does need to be known

# Vcycle



Vcycle implements an external VM factory that manages VMs.

Can be run centrally by experiment (if it is a resource provider too) or by site itself or by a third party

VMs started and monitored by Vcycle, but not managed in detail ("black boxes")

# Publishing services' details

- We could publish via BDII, but part of the motivation of Vcycle and Vac is to radically simplify what services are needed to run jobs
  - eg Vac at UCL now runs LHCb jobs in VMs with no reliance on any other site services (not even DNS)
- Vcycle needs an HTTPS webserver for part of its normal operation
  - It's used to publish Machine/Job Features info to the VMs
- So simplest model to try is publishing Vcycle status via HTTPS
  - For compatibility, use JSON/GLUE2.0 rather than a new format
- Strategy is to make Vcycle collect the required information and then write it to a file in the webserver's URL-space
- As with the LHCb LDAP URL proof of concept, can publish the URLs of the GLUE2/JSON files in the GOCDB service entries for Vcycle

# GLUE 2.0 JSON rendering

I've been working from:

- The OGF-GLUE draft:  
“GLUE v. 2.0 - Reference Realization to JSON Schema”, Nov 2014  
[https://redmine.ogf.org/dmsf/glue-wg?folder\\_id=6593](https://redmine.ogf.org/dmsf/glue-wg?folder_id=6593)
- The JSON examples in the OGF-GLUE GitHub:  
<https://github.com/OGF-GLUE/JSON/tree/master/examples>  
Includes a rendering for OpenStack where “Jobs” = VMs
- The GLUE 2.0 spec itself
- Comparisons with what's in the BDII

# Vcycle spaces and machinetypes

- Need to cover some Vcycle terminology before going on
- Each place where Vcycle can manage VMs is called a “space”
  - For example, an OpenStack tenancy
- One or more machinetypes are defined within each space
  - Usually one per experiment
  - Gives the boot image, user\_data file, disk and CPU geometry to use when creating VMs of that machinetype
- Spaces contain the VMs that Vcycle manages
- Vcycle has target shares for the different machinetypes
  - Aims to achieve that load balance unless VMs of one machinetype are currently failing to find any work to do
- Counts of VMs and details of machinetypes are enough to build JSON



# Implementation

- Vcycle rebuilds its picture of the machinetypes and VMs in each space every few minutes
- So we've added a method `publishStatus()` to the `BaseSpace` class
- Builds a Python dictionary containing GLUE 2.0:
  - “ComputingService” describing the space and its occupancy
  - For each machinetype:
    - “ComputingShare” has occupancy, limits on VMs & wall clock
    - “MappingPolicy” gives the VO name and FQAN
    - “ExecutionEnvironment” has attributes of the OS in the VMs
- Dictionary converted to JSON and written to a file published by the web server on the Vcycle machine
- Details of what else to include still being worked on

# Example representation in GOCDB

- Space name `vcycle-imperial.tier2.hep.manchester.ac.uk`
  - This is an OpenStack tenancy at Imperial College, managed by Vcycle at Manchester
  - So it goes in the GOCDB under UKI-NORTHGRID-MAN-HEP
- The Service Type is `uk.ac.gridpp.vcycle`
  - Can be discovered like other services in GOCDB
  - Downtimes can be declared as any other service
  - One Vcycle daemon might be operating more than one `uk.ac.gridpp.vcycle` service visible in GOCDB
- The grid info URL is  
`https://vm21.tier2.hep.manchester.ac.uk/spaces/vcycle-imperial.tier2.hep.manchester.ac.uk/glue2.json`




```
{ "ComputingService": [{"Name": "Vcycle", "TotalJobs": 25, "CreationTime": "2015-11-11T22:20:42Z", "RunningJobs": 25, "QualityLevel": "production", "Type": "uk.ac.gridpp.vcycle", "ID": "urn:glue2:ComputingService:vcycle-imperial.tier2.hep.manchester.ac.uk"}],
```

```
"ComputingShare": [{"Name": "cmsprod", "MaxRunningJobs": 30, "CreationTime": "2015-11-11T22:20:42Z", "RunningJobs": 0, "ID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:cmsprod", "Associations": {"ServiceID": "urn:glue2:ComputingService:vcycle-imperial.tier2.hep.manchester.ac.uk"}, "TotalJobs": 0, "ServingState": "production", "MaxWallTime": 172800, "MaxCPUTime": 172800, "MaxTotalJobs": 30}, {"Name": "atlasprod", "MaxRunningJobs": 30, "CreationTime": "2015-11-11T22:20:42Z", "RunningJobs": 0, "ID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:atlasprod", "Associations": {"ServiceID": "urn:glue2:ComputingService:vcycle-imperial.tier2.hep.manchester.ac.uk"}, "TotalJobs": 0, "ServingState": "production", "MaxWallTime": 172800, "MaxCPUTime": 172800, "MaxTotalJobs": 30}, {"Name": "lhcbprod", "MaxRunningJobs": 30, "CreationTime": "2015-11-11T22:20:42Z", "RunningJobs": 25, "ID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:lhcbprod", "Associations": {"ServiceID": "urn:glue2:ComputingService:vcycle-imperial.tier2.hep.manchester.ac.uk"}, "TotalJobs": 25, "ServingState": "production", "MaxWallTime": 172800, "MaxCPUTime": 172800, "MaxTotalJobs": 30}, {"Name": "gridpp", "MaxRunningJobs": 30, "CreationTime": "2015-11-11T22:20:42Z", "RunningJobs": 0, "ID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:gridpp", "Associations": {"ServiceID": "urn:glue2:ComputingService:vcycle-imperial.tier2.hep.manchester.ac.uk"}, "TotalJobs": 0, "ServingState": "production", "MaxWallTime": 172800, "MaxCPUTime": 172800, "MaxTotalJobs": 30}],
```

```
"MappingPolicy": [{"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:cmsprod"}, "Name": "cmsprod", "CreationTime": "2015-11-11T22:20:42Z", "Rule": ["V0:cms", "VOMS:/cms/Role=NULL/Capability=NULL"], "Scheme": "org.glite.standard", "ID": "urn:glue2:MappingPolicy:vcycle-imperial.tier2.hep.manchester.ac.uk:cmsprod"}, {"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:atlasprod"}, "Name": "atlasprod", "CreationTime": "2015-11-11T22:20:42Z", "Rule": ["V0:atlas", "VOMS:/atlas/Role=NULL/Capability=NULL"], "Scheme": "org.glite.standard", "ID": "urn:glue2:MappingPolicy:vcycle-imperial.tier2.hep.manchester.ac.uk:atlasprod"}, {"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:lhcbprod"}, "Name": "lhcbprod", "CreationTime": "2015-11-11T22:20:42Z", "Rule": ["V0:lhcb", "VOMS:/lhcb/Role=NULL/Capability=NULL"], "Scheme": "org.glite.standard", "ID": "urn:glue2:MappingPolicy:vcycle-imperial.tier2.hep.manchester.ac.uk:lhcbprod"}, {"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:gridpp"}, "Name": "gridpp", "CreationTime": "2015-11-11T22:20:42Z", "Rule": ["V0:vo.gridpp.ac.uk", "VOMS:/vo.gridpp.ac.uk/Role=NULL/Capability=NULL"], "Scheme": "org.glite.standard", "ID": "urn:glue2:MappingPolicy:vcycle-imperial.tier2.hep.manchester.ac.uk:gridpp"}],
```

```
"ExecutionEnvironment": [{"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:cmsprod"}, "OSFamily": "linux", "Name": "cmsprod", "CreationTime": "2015-11-11T22:20:42Z", "OSName": "CernVM 3", "Platform": "x86_64", "ID": "urn:glue2:ExecutionEnvironment:vcycle-imperial.tier2.hep.manchester.ac.uk:cmsprod"}, {"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:atlasprod"}, "OSFamily": "linux", "Name": "atlasprod", "CreationTime": "2015-11-11T22:20:42Z", "OSName": "CernVM 3", "Platform": "x86_64", "ID": "urn:glue2:ExecutionEnvironment:vcycle-imperial.tier2.hep.manchester.ac.uk:atlasprod"}, {"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:lhcbprod"}, "OSFamily": "linux", "Name": "lhcbprod", "CreationTime": "2015-11-11T22:20:42Z", "OSName": "CernVM 3", "Platform": "x86_64", "ID": "urn:glue2:ExecutionEnvironment:vcycle-imperial.tier2.hep.manchester.ac.uk:lhcbprod"}, {"Associations": {"ShareID": "urn:glue2:ComputingShare:vcycle-imperial.tier2.hep.manchester.ac.uk:gridpp"}, "OSFamily": "linux", "Name": "gridpp", "CreationTime": "2015-11-11T22:20:42Z", "OSName": "CernVM 3", "Platform": "x86_64", "ID": "urn:glue2:ExecutionEnvironment:vcycle-imperial.tier2.hep.manchester.ac.uk:gridpp"}] }
```



```
"ComputingService": [{
  "Name": "Vcycle",
  "TotalJobs": 25,
  "CreationTime": "2015-11-11T22:20:42Z",
  "RunningJobs": 25,
  "QualityLevel": "production",
  "Type": "uk.ac.gridpp.vcycle",
  "ID": "urn:glue2:ComputingService:vcycle-
    imperial.tier2.hep.manchester.ac.uk"
  }]
```

```
"ComputingShare": [..., {
  "Name": "lhcbprod",
  "MaxRunningJobs": 30,
  "CreationTime": "2015-11-11T22:20:42Z",
  "RunningJobs": 25,
  "ID": "urn:glue2:ComputingShare:vcycle-
imperial.tier2.hep.manchester.ac.uk:lhcbprod",
  "Associations": {
    "ServiceID": "urn:glue2:ComputingService:vcycle-
imperial.tier2.hep.manchester.ac.uk"
  },
  "TotalJobs": 25,
  "ServingState": "production",
  "MaxWallTime": 172800,
  "MaxCPUtime": 172800,
  "MaxTotalJobs": 30
}, ... ]
```

```
"MappingPolicy": [..., {
  "Associations": {
    "ShareID": "urn:glue2:ComputingShare:vcycle-
    imperial.tier2.hep.manchester.ac.uk:lhcbprod"
  },
  "Name": "lhcbprod",
  "CreationTime": "2015-11-11T22:20:42Z",
  "Rule": [ "V0:lhcb",
            "VOMS:/lhcb/Role=NULL/Capability=NULL"],
  "Scheme": "org.glite.standard",
  "ID": "urn:glue2:MappingPolicy:vcycle-
  imperial.tier2.hep.manchester.ac.uk:lhcbprod"
}, ...]
```

```
"ExecutionEnvironment": [..., {
  "Associations": {
    "ShareID": "urn:glue2:ComputingShare:vcycle-
      imperial.tier2.hep.manchester.ac.uk:lhcbprod"
  },
  "OSFamily": "linux",
  "Name": "lhcbprod",
  "CreationTime": "2015-11-11T22:20:42Z",
  "OSName": "CernVM 3",
  "Platform": "x86_64",
  "ID": "urn:glue2:ExecutionEnvironment:vcycle-
    imperial.tier2.hep.manchester.ac.uk:lhcbprod"
  }, ...]
```

# Next steps

- Extract more information from OpenStack etc, including
  - The real Total VM limits for spaces and machinetypes
  - Execution environment memory and processors per VM
- Add status publishing to Vac
  - Each hypervisor can use VacQuery UDP protocol to gather site-wide information and randomly take turns publishing it “somewhere”
  - Being able to push a JSON file into GOCDB entry for a service would be very useful here
- Test consuming Vcycle and Vac status by LHCb CS scripts once GLUE 2.0 support is ready
- Try with other VOs interested in gathering this information programmatically

# Summary

- GridPP has now added basic GLUE 2.0 publishing of Vcycle status in JSON via HTTPS
- This gives space-wide stats, and numbers and attributes for each machinetype
- Running live on the GridPP Vcycle Service at Manchester
- Entered in the GOCDB as a `uk.ac.gridpp.vcycle` service
- Grid info URL points to a `glue2.json` file for each Vcycle space
- Next steps include adding more fields, and adding publishing to Vac
- Test consuming this information with LHCb and other interested experiments