



Vac capacity
publishing to GOCDB

Vac accounting
and APEL

Andrew McNab
University of Manchester
LHCb and GridPP



Two talks in one set of slides

Start with:

Vac publishing to GOCDB



Vac and Vcycle

- Vac: worker nodes as autonomous VM-factories
- Vcycle: manages VMs on OpenStack etc
- Both use the Vacuum Model, and same VM definitions
- VMs are started up and pull work in:
 - DIRAC payloads, HTCondor jobs etc
- In terms of publishing, they both have a view of the whole “space” of VMs
 - Vcycle from cloud API VM listing
 - Vac using VacQuery UDP protocol to poll other VM factories in the same space



Publishing to GOCDDB

- Preliminary tests done with curl + shell script
 - Used `vac03.tier2.hep.manchester.ac.uk` service
 - Used a DN added for API access
- Able to use the new API to insert Extension Properties for the service itself
- Able to add Endpoints within the service by simulating submitting an HTML form
- Able to add Extension Properties to Endpoints using the new API
- Not able to add or edit services using API DN

vac03.tier2.hep.manchester in GOcdb

Service: vac03.tier2.hep.manchester.ac.uk - uk.ac.gridpp.vac

Vac space vac03.tier2.hep.manchester.ac.uk

System

Host name	vac03.tier2.hep.manchester.ac.uk
IP Address	
IP v6 Address	
Operating System	
Architecture	
Contact E-Mail	ops@tier2.hep.manchester.ac.uk

Grid Information

Host DN	
URL	
Parent Site	UKI-NORTHGRID-MAN-HEP
Scope	EGI
Tags	

Project Data

Production Level	✘
Beta	✘
Monitored	✘

Service Endpoints (endpoints?)

Name	URL	Interface Name	Edit	Remove
lhcbprod		uk.ac.gridpp.vac		

Extension Properties

Name	Value	Edit	Select All
ComputingManagerProductName	Vac		<input type="checkbox"/>
ComputingManagerProductVersion	02.01pre		<input type="checkbox"/>
ComputingManagerTotalLogicalCPUs	40		<input type="checkbox"/>
ComputingManagerTotalSlots	40		<input type="checkbox"/>
BenchmarkType	specint2000		<input type="checkbox"/>
BenchmarkValue	1250.0		<input type="checkbox"/>

Service Endpoints (endpoints?)

Name

lhcbprod

Interface Name

uk.ac.gridpp.vac

Extension Properties

Name	Value
ComputingManagerProductName	Vac
ComputingManagerProductVersion	02.01pre
ComputingManagerTotalLogicalCPUs	40
ComputingManagerTotalSlots	40
BenchmarkType	specint2000
BenchmarkValue	1250.0

Storing GLUE2 in GOCDDB extensions

- Properties of the vac03 service
- Concatenated names (“full path”)

Extension Properties	
Name	Value
ComputingManagerProductName	Vac
ComputingManagerProductVersion	02.01pre
ComputingManagerTotalLogicalCPUs	40
ComputingManagerTotalSlots	40
BenchmarkType	specint2000
BenchmarkValue	1250.0

Storing GLUE2 in GOCDB extensions (2)

- Properties of the lhcbprod endpoint

Extension Properties	
Name	Value
ComputingShareMaxWallTime	86400
ExecutionEnvironmentOSFamily	Linux
ExecutionEnvironmentOSName	CERN Virtual Machine
ExecutionEnvironmentPlatform	x86_64
ExecutionEnvironmentVirtualMachine	True
MappingPolicyRule	VO:lhcb,VOMS:/lhcb/Role=NULL/Capability=NULL
MappingPolicyScheme	org.glite.standard

Next steps in publishing to GOCDB

- These are sufficient for LHCb discovery of sites and importing service parameters into DIRAC CS
- Can readily be extended for other requirements
- Plan is to automate this publishing within Vac and Vcycle, using each system's knowledge of capacity, HS06 values etc
 - Started coding this in common vacutils library
- Should be straightforward for CE/Batch systems too
- Interim solution would be general tool to read Resource LDAP server on CE and republish to GOCDB
 - Needs to be run by site due to permissions model



Vac accounting and APEL



Vac and Vcycle

- Vac: worker nodes as autonomous VM-factories
- Vcycle: manages VMs on OpenStack etc
- Both use the Vacuum Model, and same VM definitions
- VMs are started up and pull work in:
 - DIRAC payloads, HTCondor jobs etc
- Same accounting model and library used for both
 - So what I say for Vac applies to Vcycle
 - But Vcycle isn't managing any resources which need APEL accounting yet, so APEL support not used in production

Vac and APEL

- Vac has had native APEL support since 2015
 - Before that wrote out PBS/BLAHP format log files that could be parsed and fed into site APEL service
- When a VM finishes, Vac writes out an APEL accounting record file in `/var/lib/vac/apel-outgoing` and `/var/lib/vac/apel-archive`
- APEL's `ssmsend` client runs every ~hour to process record files in `apel-outgoing`
 - When file is successfully received by message bus, then `ssmsend` deletes it
 - `ssmsend` is also run when Vac shuts down to flush any records still waiting

APEL record file

```
APEL-individual-job-message: v0.3
Site: UKI-NORTHGRID-MAN-HEP
SubmitHost:
  vac01.tier2.hep.manchester.ac.uk/vac-wn1905190.tier2.hep.manchester.ac.uk
LocalJobId: 3ec25221-4c8f-4d91-a5f2-df6966b7870f
LocalUserId: wn1905190.tier2.hep.manchester.ac.uk
Queue: atlasprod
GlobalUserName: /DC=uk/DC=ac/DC=manchester/DC=hep/DC=tier2/DC=vac01
FQAN: /atlas/Role=NULL/Capability=NULL
WallDuration: 9953
CpuDuration: 9368
Processors: 1
NodeCount: 1
InfrastructureDescription: APEL-VAC
InfrastructureType: grid
StartTime: 1484691667
EndTime: 1484701620
MemoryReal: 2621440
MemoryVirtual: 2621440
ServiceLevelType: HEPSPEC
ServiceLevel: 8.8
%%
```

Mapping VM attributes to Job records

- Some attributes are what you would expect
 - Sitename and FQAN
- GlobalUsername represents the space too
 - /DC=uk/DC=ac/DC=manchester/DC=hep/DC=tier2/DC=vac01
 - Means Vac-specific totals can be seen in the portalSubmitHost is specific to the VM factory
 - vac01.tier2.hep.manchester.ac.uk/vac-wn1905190.tier2.hep.manchester.ac.uk
 - vac01 is the Vac space (equivalent to CE name)
 - wn1905190 is the VM factory
 - Used by APEL-Sync records: means Sync records can be generated on each VM-factory independently

This Restricted View allows Resource Centre Administrators to retrieve accounting data of their Resource Centre. The metric shown is Normalized Elapsed time (HEPSPEC06) * Number of Processors, grouped by User DN and Year, the official EGI VOs are shown. No local jobs are shown
 Logged in as /C-UK/O-eScience/OU-Manchester/L-HEP/CN-andrew mc nab
 VO Manager of skatelescope.eu
 Resource Centre Manager of UKI-NORTHGRID-MAN-HEP

Resource Centre Admin - Normalized Elapsed time (HEPSPEC06) * Number of Processors (h) by User DN and Year (Official VOs)

User DN	2016	2017	Total	Percent
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-atpilo1/CN-614260/CN-Robot- ATLAS Pilot1	151,635,956	3,688,627	155,324,583	45.03%
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-atpilo2/CN-531497/CN-Robot- ATLAS Pilot2	80,841,131	4,435,422	85,276,554	24.73%
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-romanov/CN-427293/CN-Vladimir Romanovskiy	48,739,317	4,078,271	52,817,588	15.31%
/DC-uk/DC-ac/DC-manchester/DC-hep/DC-tier2/DC-vac01	15,252,599	1,335,142	16,587,741	4.81%
/DC-uk/DC-ac/DC-manchester/DC-hep/DC-tier2/DC-vac02	15,540,564	217,979	15,758,543	4.57%
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-sailer/CN-665729/CN-Andre Sailer	6,161,310	725,753	6,887,063	2.00%
/O-GRID.FR/C-FR/O-CNRS/OU-CC-IN2P3/CN-Vanessa Hamar	1,815,939	0	1,815,939	0.53%
/C-UK/O-eScience/OU-Imperial/L-Physics/CN-dirac-pilot.grid.hep.ph.ic.ac.uk	1,703,471	17,898	1,721,369	0.50%
/DC-uk/DC-ac/DC-manchester/DC-hep/DC-tier2/DC-vac04	1,312,544	72,403	1,384,947	0.40%
/C-UK/O-eScience/OU-Manchester/L-HEP/CN-mark smith	1,150,139	0	1,150,139	0.33%
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-kmaguire/CN-748762/CN-Kevin Scott Maguire	1,060,773	0	1,060,773	0.31%
/DC-com/DC-DigiCert-Grid/O-Open Science Grid/OU-Services/CN-frontend/fifebatch.fnal.gov	915,817	0	915,817	0.27%
/C-UK/O-eScience/OU-Manchester/L-HEP/CN-yaadav bhauruth	679,730	3,500	683,230	0.20%
/C-CA/O-Grid/OU-physics.utoronto.ca/CN-Kyle Cormier	563,899	0	563,899	0.16%
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-jcrane/CN-783706/CN-Jonathan Crane	492,896	68,844	561,740	0.16%
/DC-com/DC-DigiCert-Grid/O-Open Science Grid/OU-Services/CN-pilot/glidein2.chtc.wisc.edu				0.14%
/C-UK/O-eScience/OU-Manchester/L-HEP/CN-julia mariana iturbe ponce				0.11%
/O-GRID.FR/C-FR/O-CNRS/OU-CREATIS/CN-Sorina Camarasu				0.11%
/DC-uk/DC-ac/DC-manchester/DC-hep/DC-tier2/DC-vac01				0.07%
/DC-uk/DC-ac/DC-manchester/DC-hep/DC-tier2/DC-vac03				0.07%
/DC-org/DC-opensciencegrid/O-Open Science Grid/OU-Services/CN-frontend/fifebatch.fnal.gov				0.07%
/C-UK/O-eScience/OU-Manchester/L-HEP/CN-stefania reichert				0.06%
/C-UK/O-eScience/OU-QueenMaryLondon/L-Physics/CN-steve lloyd				0.02%
/O-GRID.FR/C-FR/O-CNRS/OU-ICube/CN-Arnaud Kress				0.02%
/C-UK/O-eScience/OU-RoyalHollowayLondon/L-Physics/CN-ian connelly				0.02%
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-rroberts/CN-753763/CN-Rhys Thomas Roberts	49,088	0	49,088	0.01%
/C-UK/O-eScience/OU-Manchester/L-HEP/CN-steven wren	32,427	12,849	45,276	0.01%
/C-UK/O-eScience/OU-Manchester/L-HEP/CN-fabian wilk	37,293	0	37,293	0.01%
/DC-ch/DC-cern/OU-Organic Units/OU-Users/CN-roiser/CN-576600/CN-Stefan Roiser	33,472	0	33,472	0.01%
/C-UK/O-eScience/OU-Manchester/L-HEP/CN-thomas brooks	26,930	0	26,930	0.01%
/C-UK/O-eScience/OU-Oxford/L-DeSC/CN-kashif mohammad/CN-Robot-GridClient	17,248	1	17,250	0.01%
Total	330,093,989	14,804,854	344,898,842	
Percent	95.71%	4.29%		

[Download JSON Data](#) / [Download CSV Data](#)

The information in the previous table is also shown in the following graph.

process

Normalized Elapsed time (HEPSPEC06) * Number of Processors (h) by User DN and Year

Grid vs Cloud accounting

- Vac tries to look like “just another CE” in GOCDB etc
 - Natural to report it through “Grid” route in APEL
- Vcycle also looks like a CE, but is managing cloud resources.
- This approach with APEL allows Cloud accounting through the normal “Grid” route in APEL too
 - In particular, it requires no reconfiguration or modification of the underlying cloud infrastructure
- Compare EGI Federated Cloud approach
 - Needs separate route into APEL
 - Needs reconfiguration by cloud provider (eg elevated privileges for account reading native accounting DB)