Cloud Accounting

John Gordon for the APEL Team

Pre-GDB January 2013

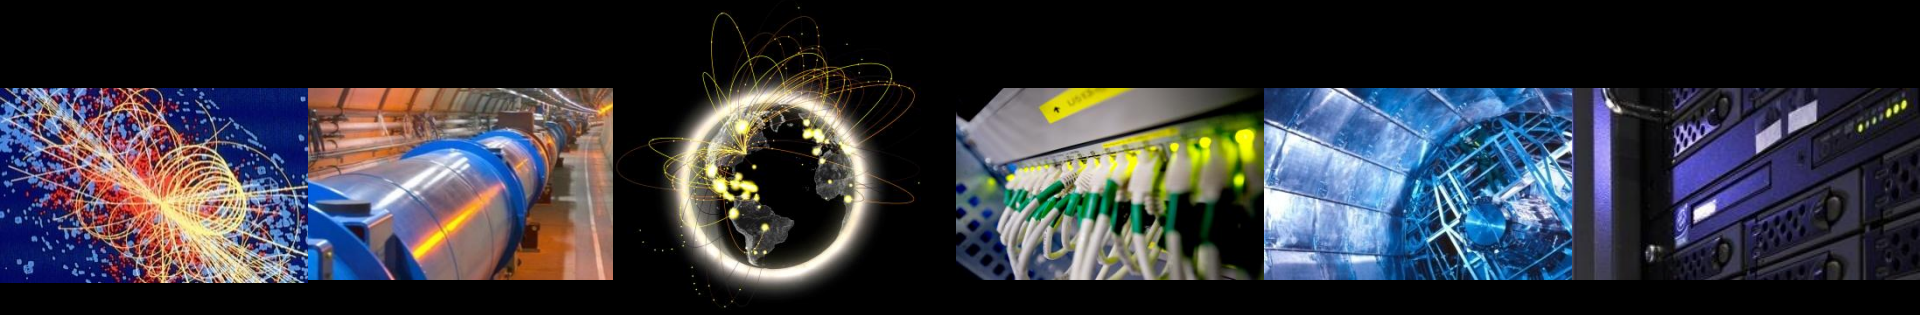

WLCG
Worldwide LHC Computing Grid

- EGI Federated Cloud Task Force started in 2011, includes several WLCG sites.
- As the name suggests it is looking at use of multiple clouds by common interfaces using common identities.
- The APEL Team is working on accounting in this Task Force.
- Accounting is of VM instantiations. What happens inside a VM is accounted by traditional methods.
- This is fine for the people running the infrastructure but it is not proven yet whether this can be of use to a VO. Hence today's discussion.

EGI Progress

- Cloud Accounting Schema defined
- https://wiki.egi.eu/wiki/Fedcloud-tf:WorkGroups:Scenario4#Proposed_new_version_of_the_Cloud_Accounting_Message_Format
- Implemented on OpenNebula, OpenStack, and WNoDeS
- The EGI TF runs a testbed with a number of test cloud infrastructures http://goc-accounting.grid-support.ac.uk/cloudtest/cloudsites.html
- These publish to a consumer at RAL using SSM (as used for job accounting today by CERN, OSG, ARC, and in the future EMI3 version of the APEL client).
- The group is working with OGF so that v2 of the OGF UR (Usage Record) will be useful for clouds/VMs too.
- Talking to OSG about them publishing into common repository.

APEL Cloud

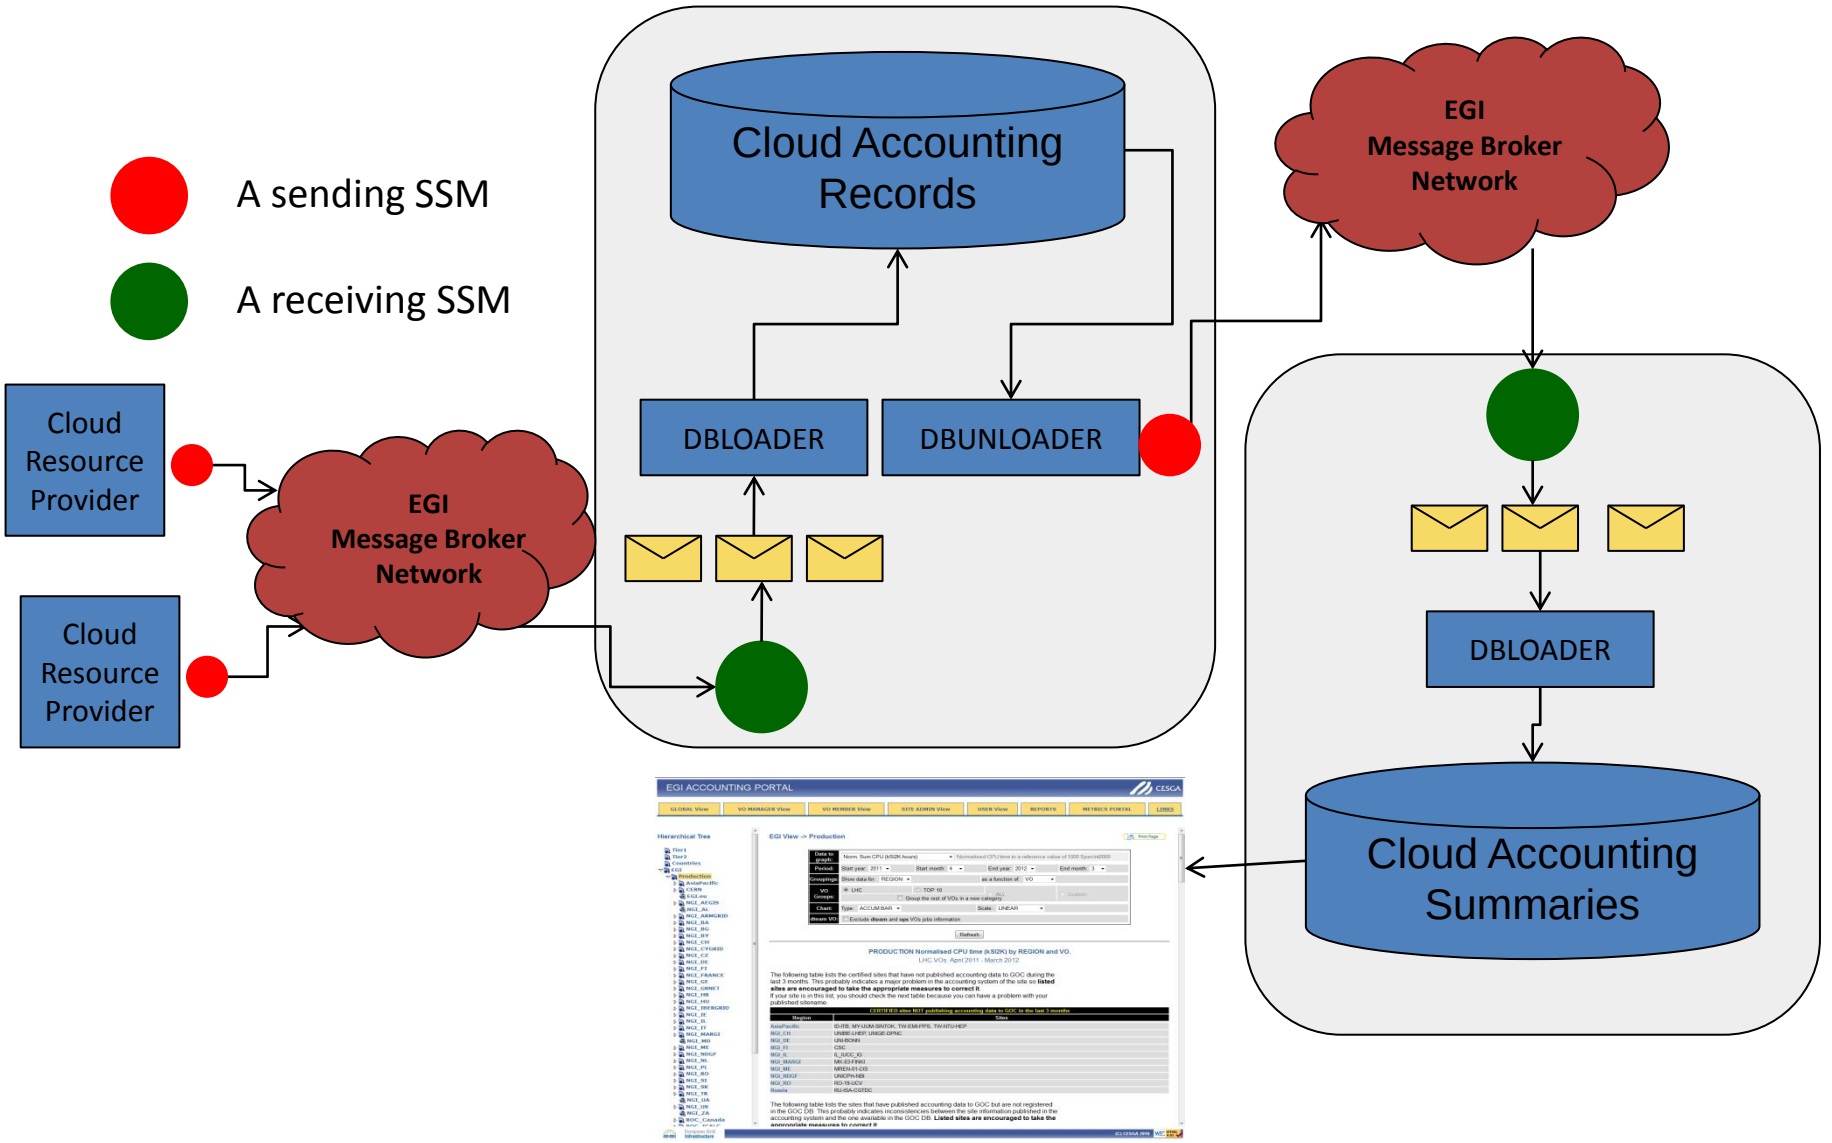

Cloud Accounting Portal

Cloud Record Monitoring

RecordId	Site	ZoneName	MachineName	Status	StartTime	EndTime	Network in (GB)	Network out (GB)	Memory (MB)	Disk (GB)	ImageId	CloudType
2013-01-15 09:30:01+00:00 CESNET vm-0	CESNET	EU	'one-0'	completed	2011-10-17 08:31:04	2011-10-17 10:41:16	0	0	512	None	'scilin6'	OpenNebula
2013-01-15 09:30:01+00:00 CESNET vm-1	CESNET	EU	'one-1'	completed	2011-10-17 10:46:45	2011-10-17 11:10:17	0	0	512	None	'scilin6'	OpenNebula
2013-01-15 09:30:01+00:00 CESNET vm-2	CESNET	EU	'vm_001'	completed	2011-10-17 10:49:18	2011-10-17 11:11:04	0	0	256	None	'scilin6'	OpenNebula
2013-01-15 09:30:01+00:00 CESNET vm-3	CESNET	EU	'one-3'	completed	2011-10-17 11:08:18	2011-10-17 12:49:31	0	0	512	None	'scilin6'	OpenNebula
2013-01-15 09:30:01+00:00 CESNET vm-4	CESNET	EU	'asdfhgj_1'	completed	2011-10-17 11:47:18	2011-10-17 12:52:06	0	0	512	None	'scilin6'	OpenNebula
2013-01-15 09:30:01+00:00 CESNET vm-5	CESNET	EU	'asdfhgj_0'	completed	2011-10-17 11:47:18	2011-10-17 12:52:06	0	0	512	None	'scilin6'	OpenNebula
2013-01-15 09:30:01+00:00 CESNET vm-6	CESNET	EU	'one-6'	completed	2011-10-17 12:12:46	2011-10-17 12:21:08	0	0	256	None	'scilin6'	OpenNebula

The top of a list of records received earlier this morning.

Sites publishing cloud accounting records

Page last updated: 2013-01-15 10:05:52.144756

EGI Testbed

Site	ZoneName	NumberOfMachines	CloudType	LastUpdated
CESGA	EU	1131	OpenNebula	2013-01-15 08:01:20
CESNET	EU	450	OpenNebula	2013-01-15 09:45:31
CYFRONET	EU	25	OpenNebula	2012-09-17 20:36:58
FZJ	EU	63	Openstack	2012-12-04 11:05:48
GWDG	EU	237	OpenNebula	2012-09-23 06:17:03
IN2P3-CC	EU	100	Openstack	2012-12-17 16:42:48
INFN CNAF	EU	1	WNoDeS	2013-01-03 16:25:54
KTH CLOUD	EU	13	OpenNebula	2012-09-17 22:06:02

Summary of cloud accounting records received by APEL.

Alternative Accounting

- APEL is a flexible system. It already receives traditional accounting records from some sites and infrastructures direct using SSM, not using the EMI APEL client.
- e.g. CERN, NIKHEF, IN2P3, OSG, NorduGrid, INFN
- Anyone who can cut an accounting record to reflect usage can publish it to APEL.
  - CERN have done it for traditional batch jobs. Their suggestion including clouds looks straightforward.
- If cloud work is equivalent to jobs (pilot jobs in WN or VM) then the usage could be aggregated. i.e all ATLAS work shown in the same portal view.
- PANDA could alternatively cut usage records with the advantage that it could be done per workload and not per pilot.

Issues

- Double Counting
  - ATLAS-only, cloud-only site would be straightforward to integrate, but when will we see one.
  - Mixed VO sites with traditional WNs as well as Clouds will exist for some time. Care is needed to aggregate traditional accounting with cloud.
  - Not impossible. Today a site can filter on VO, so if ATLAS was publishing URs from Panda, the site could switch off APEL just for ATLAS and continue to publish for other VOs.

HEPSPEC06

- I think it unlikely that clouds will provide a continuously variable power of VMs.
- More likely that they will provide a menu of VM types, varying in memory, ncores, and cpuspeed.
- Sites should be encouraged to benchmark one of each type they provide.
- The EGI group is working on infrastructures publishing capabilities in BDII. HS06 and/or HS06core should be one such value published.
- Easier would be for the cloud infrastructure to do a lookup when cutting URs, to find the HS06 and include it.

Summary

- Accounting from clouds we control is straightforward.
  - EGI Task Force has started doing it.
  - Many sites in EGI pilot are also WLCG sites, so work with them, don't start something orthogonal.
- The difficult bit will be integrating such cloud accounting with traditional job accounting and avoiding double counting.
- Cloud bursting from a WLCG cloud should be achievable, not so clear with direct use of commercial clouds