Accounting Portal New developments: InterNGI and Cloud usage

Álvaro Simón
Iván Díaz

CESGA

Introduction

- Since the last EGI TF, there were many portal developments
- Here we focus on InterNGI and Cloud usage.
- InterNGI reporting is supported by a VT, which recently produced a complete report on the topic.
- Cloud reporting is supported by the Federated Cloud Task Force, and this functionality was part of a demo on this TF.

Outline

InterNGI usage reporting

Cloud Accounting

Overview

- The EGI Usage VT was created to produce a complete report of resource consumption between NGIs.
- A great part of the effort was involved with creating new views and reports on the Portal to measure this usage.
- The usage between NGIs and countries is based on existing UserDN data, which is maped to a institution and this institution to a country and NGI.
- Actions were also taken to improve UserDN publishing and make the final report more useful.
- In the last TF, there were already usage statistics for each Site and Country.
- Since then, we have provision statistics, the Usage Matrix and the Publishing Matrix.

InterNGI Usage Matrix

The Usage Matrix contains for each combination of two NGIs or countries:

- The consumption for the specified period in one of the usual CPU use measurements of the portal.
- Percentage relative to the total consumed on the column's NGI (blue)
- Percentage relative to the total consumed by the row's NGI (red)

Uses/Used by	AsiaPacific	CERN	EGI.eu	NGI_AEGIS	NGI_AL	NGI_ARMGRID	NGI_BA
AsiaPacific	5722189 10.36% / 81.98%	128422 0.02% / 1.39%					
CERN	34406039 62.32% / 1.6%	705399435 99.51% / 32.73%				34633 58.78% / 0%	
EGI.eu							
NGI_AEGIS							
NGI_AL							
NGI_ARMGRID						22775 38.65% / 15.62%	
NGI_BA							9 39.13% / 100%
[illegible]	7						

Publishing Matrix

The Publication Matrix monitors the quality of publishing for each NGI/country. It shows for a concrete period:

- Total sites, correctly publishing sites and percentage.
- CPU resources without a valid UserDN (selectable measurement).
- Percentage of measurement with correct UserDN.
- Users active on that period.

Matrix of foreign UserDN publishing						
Region	Total Sites	Publishing Sites	Percentage publishing	Normalised CPU time (kSI2K) without UserDN	Normalised CPU time (kSI2K) with UserDN (%)	Active Users
Armenia	2	2	100%	30943	65.6%	5
Australia	1	1	100%	2221637	66.3%	2
Austria	2	2	100%	1751746	42.1%	8
Belarus	1	1	100%	41147	37.9%	2
Belgium	2	2	100%	14537855	27.5%	51
Bosnia and Herzegovina	1	1	100%	59	28%	1
Brazil	2	2	100%	1880847	28.6%	
Bulgaria	4	4	100%	278253	41.1%	11
Canada	6	6	100%	27748630	83.6%	10
Chile	1	1	100%	0	100%	1
China	2	2	100%	7998954	27.9%	20
Croatia	3	3	100%	14119	97.9%	6
Cyprus	1	1	100%	26701	38.6%	2
Czech Republic	2	2	100%	4024680	88.1%	13
Denmark	1	1	100%	0	100%	2
Estonia	1	1	100%	63674	99.5%	4
Finland	1	1	100%	0	100%	6
France	17	17	100%	50	100%	250

Outline

InterNGI usage reporting

Cloud Accounting

Overview

- The Portal was expanded to include a Cloud accounting view.
- The data is received on a separated DB using SSM.
- The schema is different than usual Grid accounting, and is still evolving.
- The records are received from several RPs, and it has enough volume to benefit from summarization.

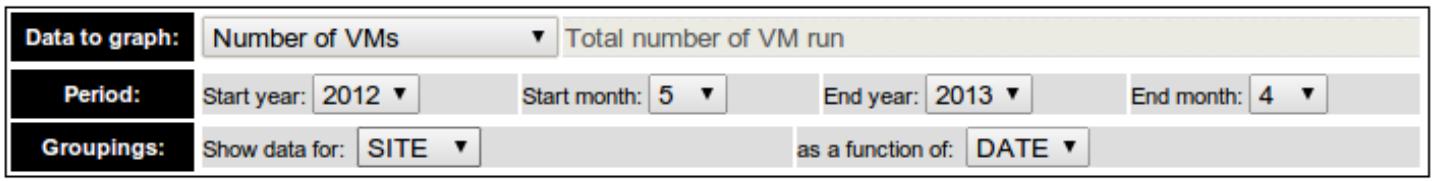

Cloud Accounting (II)

The query set is slightly different than that of Grid, but includes Network and Disk metrics.

- Number of VMs - Aggregate number of VMs executed on the site.
- Sum CPU time - Total active CPU time for all VMs
- Sum Elapsed time - Total wall clock time for all VMs
- Inbound Network Traffic - Total inbound traffic reveceived by VMs
- Outbound Network Traffic - Total outbound traffic sent by VMs.
- Memory Used - Total memory allocated (not used) by VMs.
- Disk Used - Total disk space allocated (not used) by VMs.

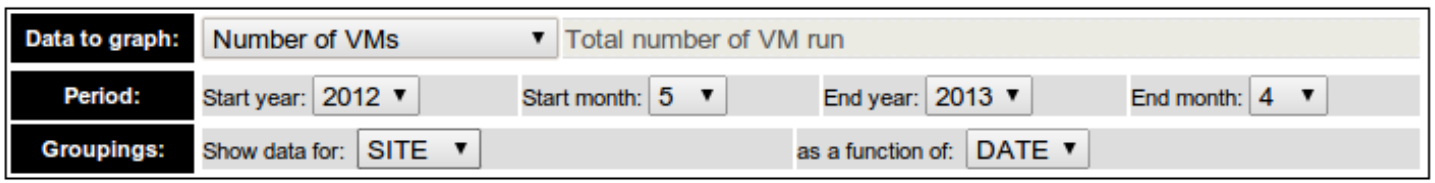

Cloud Accounting (III)

- The results are presented in a HTML table and also made available as XML.
- The table has two dimensions that can be assigned to Site, Date or VO.
- Values can also be filtered by start and end date.
- Totals and percentages are calculated for each dimension and the whole of the table.

Total number of VM run by SITE and DATE								
SITE	Oct 2012	Nov 2012	Dec 2012	Jan 2013	Feb 2013	Mar 2013	Total	%
CESGA	0	0	1,017	0	13	458	1,488	26.32%
CESNET	2	0	1,497	0	46	820	2,365	41.84%
FZJ	0	0	0	0	0	3	3	0.05%
GRIF	0	0	7	0	0	0	7	0.12%
GWDG	0	0	1,487	0	17	0	1,504	26.61%
IISAS-Bratislava	0	0	0	0	0	1	1	0.02%
IN2P3-CC	0	0	0	0	0	11	11	0.19%
KTH CLOUD	7	44	0	1	134	88	274	4.85%
Total	9	44	4,008	1	210	1,381	5,653	
Percentage	0.16%	0.78%	70.90%	0.02%	3.71%	24.43%		
Click here for XML encoded data								

Cloud Accounting (IV)

- Graphs are also produced for each dimension, detailing the distribution for each value.

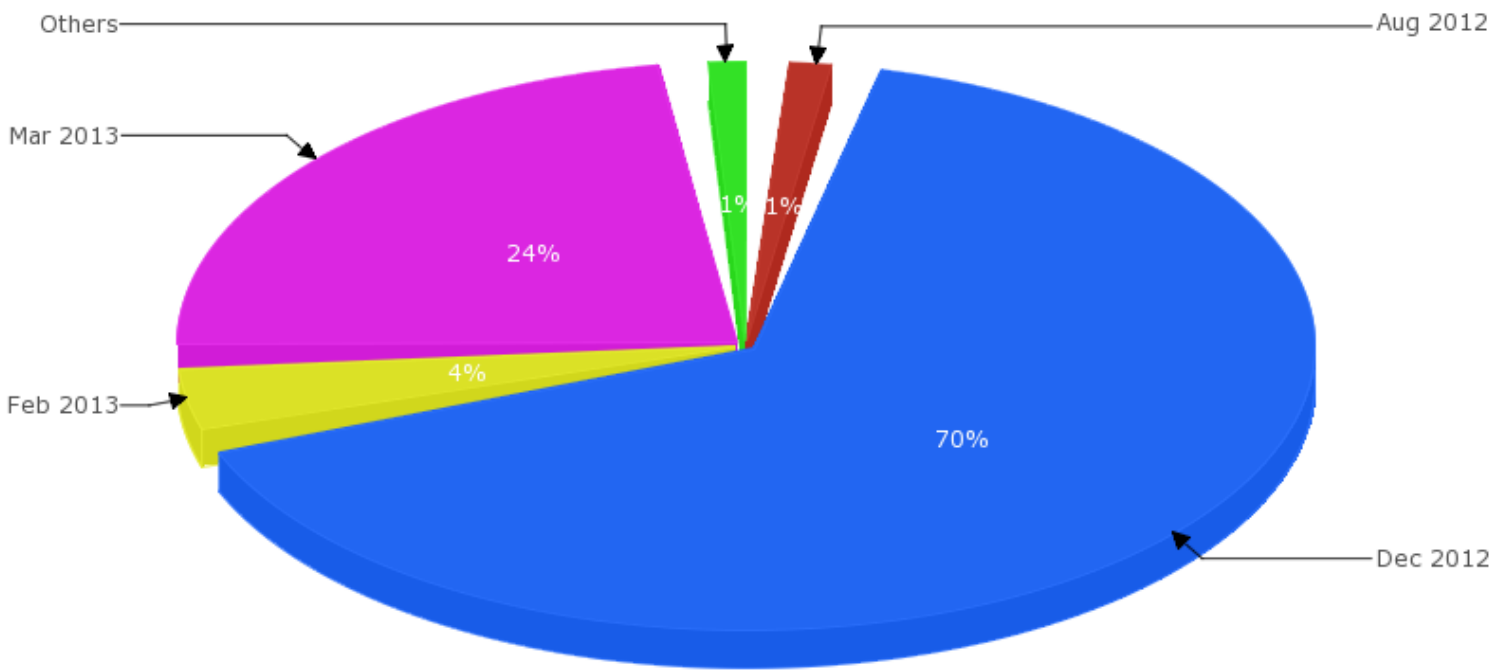

Cloud Accounting - Future work

There are several directions in which to improve

- Better integration with normal Grid accounting.
- UserDN accounting, with privileged views.
- Better crossover of information with GOCDB.
- VM Application accounting.
- Internal VM Accounting?