

# EGI Federated Cloud

**Enol Fernández**

Cloud Technologist (EGI.eu/CSIC)



[www.egi.eu](http://www.egi.eu)

EGI-Engage is co-funded by the Horizon 2020 Framework Programme  
of the European Union under grant number 654142



EGI Federated Cloud is a hybrid federated cloud for research user communities, which comprises cloud providers to support data- and computing-intensive workloads

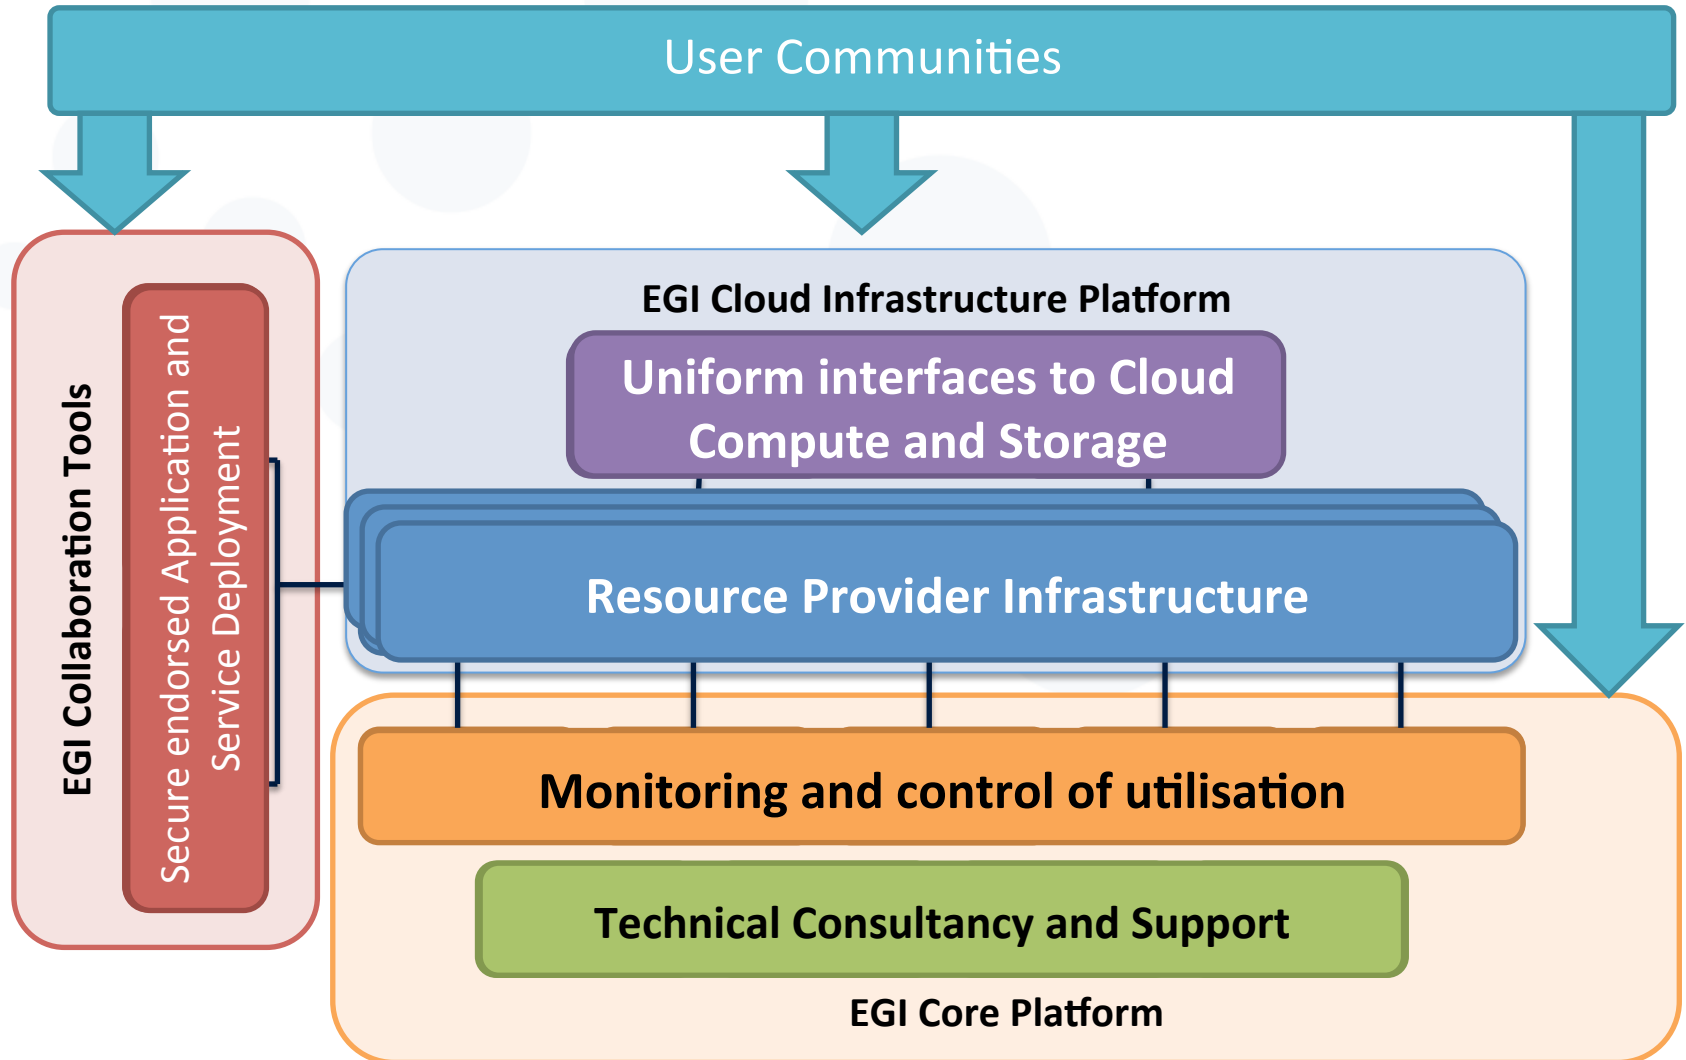
## Public Cloud

- **Open to any research community**
- **Open Standards:** use of common open standards for the interfaces and images (OCCL, CDMI, OVF, GLUE2, ...)
- **Resource integration:** Cloud Computing integrated into the existing production infrastructure.

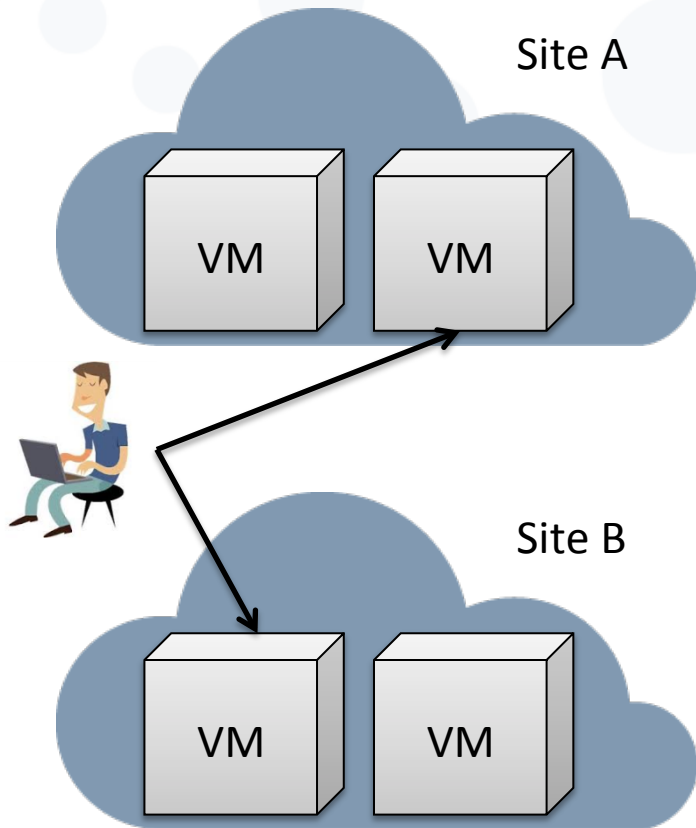
## Community Cloud

- **Accessible to one or more selected VOs**
- **API choice**, but must assure **interoperability** within the community. Standards are fostered.
- **Looser Federation profile:** based on a subset of components (accounting, monitoring, ...)

# EGI Cloud Infrastructure

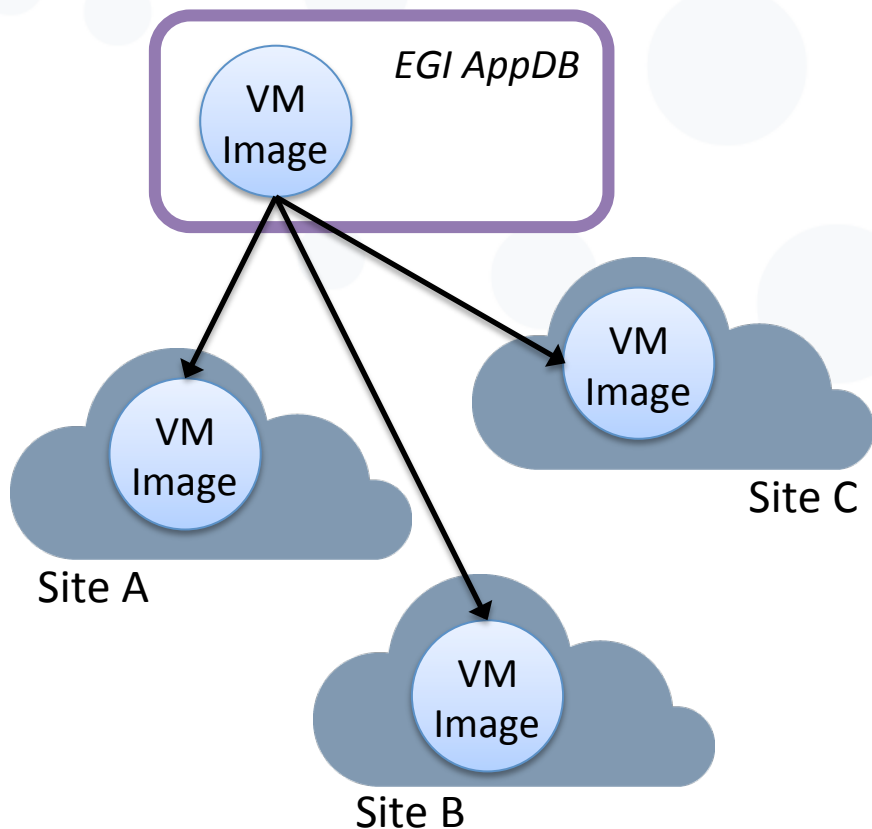


## On demand compute to run any kind of workloads on virtual machines



- **OCCI** as standard interface
  - FedCloud **contextualization** extension
  - FedCloud **profile** extension
- **Clients**
  - ruby and java clients
  - OCCI connectors in brokers
- **Servers**
  - rOCCI-server
  - OCCI-OS
  - snf-occi

## Automatic and secure distribution of endorsed VM images for Virtual Organisations

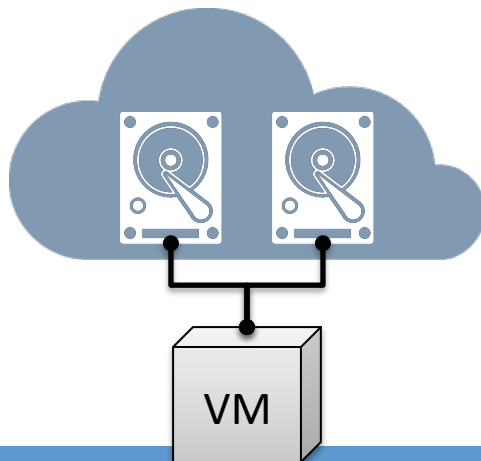


- **Web based** front-end in AppDB
- **OVF** for packaging images
- **HEPiX lists** for distribution
  - Endorsed by VO managers
  - Signed metadata
  - RP subscribe and download
- EGI endorses basic OS images

# Storage Management

## Block Storage

- **Persistent** Block Level Storage to **attach** to VMs
- Manage with **OCCI**
  - use as any other block device from VMs



## Object Storage

- **Data storage** accessible from **anywhere** at **any time**
- **Sharing/Serve** data
- **CDMI** as standard API
  - No POSIX access!

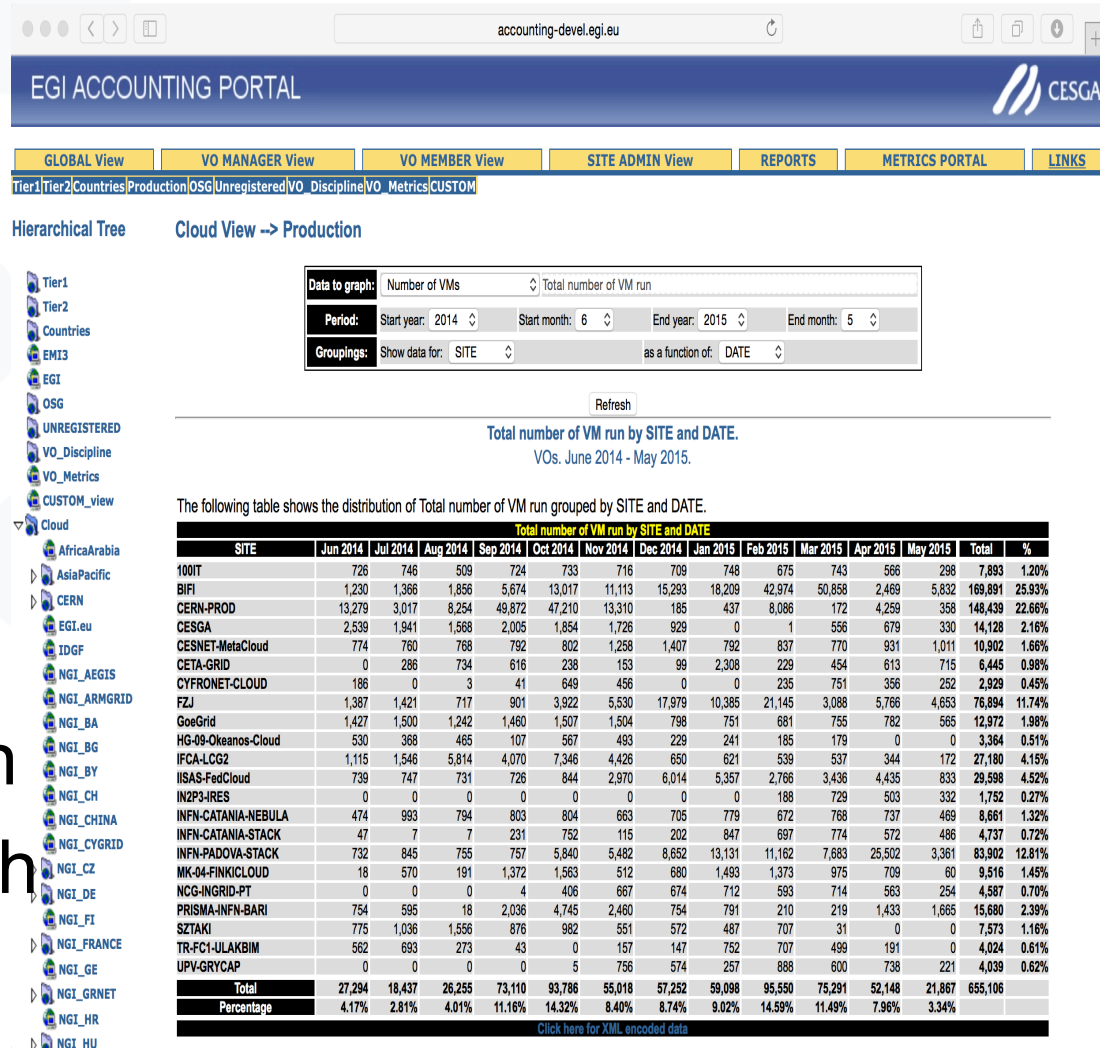


- Users access the resources with their **X.509 certificates**
  - (+ VOMS integration with PERUN for OpenNebula)
- Developments ready for **per-user sub-proxy** support
  - AAI solution to improve ease of user access and traceability from robots
- FedCloud members participate in **AAI pilot** activities
  - Support for other AuthN/AuthZ methods development

- **Single point of information** for all EGI resources (both HTC and cloud)
- Cloud information published with **OGF Glue 2.0** standard
  - Extensions proposed to better define the services, now on public comment period
  - Service endpoints, available images, VM templates
- **AppDB** is the main client for the information
- Service types in **GOCDDB**:
  - OCCI, CDMI, VM Catcher, AppDB



- Collects VM usage metrics in the **central accounting repository**
- New reporters developed to ensure consistency of generated information
- VO manager view with VO specific info



EGI ACCOUNTING PORTAL

GLOBAL View | VO MANAGER View | VO MEMBER View | SITE ADMIN View | REPORTS | METRICS PORTAL | LINKS

Tier1 Tier2 Countries Production OSG Unregistered VO Discipline VO\_Metrics CUSTOM

Hierarchical Tree | Cloud View -> Production

Data to graph: Number of VMs | Total number of VM run

Period: Start year: 2014 | Start month: 6 | End year: 2015 | End month: 5

Groupings: Show data for: SITE | as a function of: DATE

Refresh

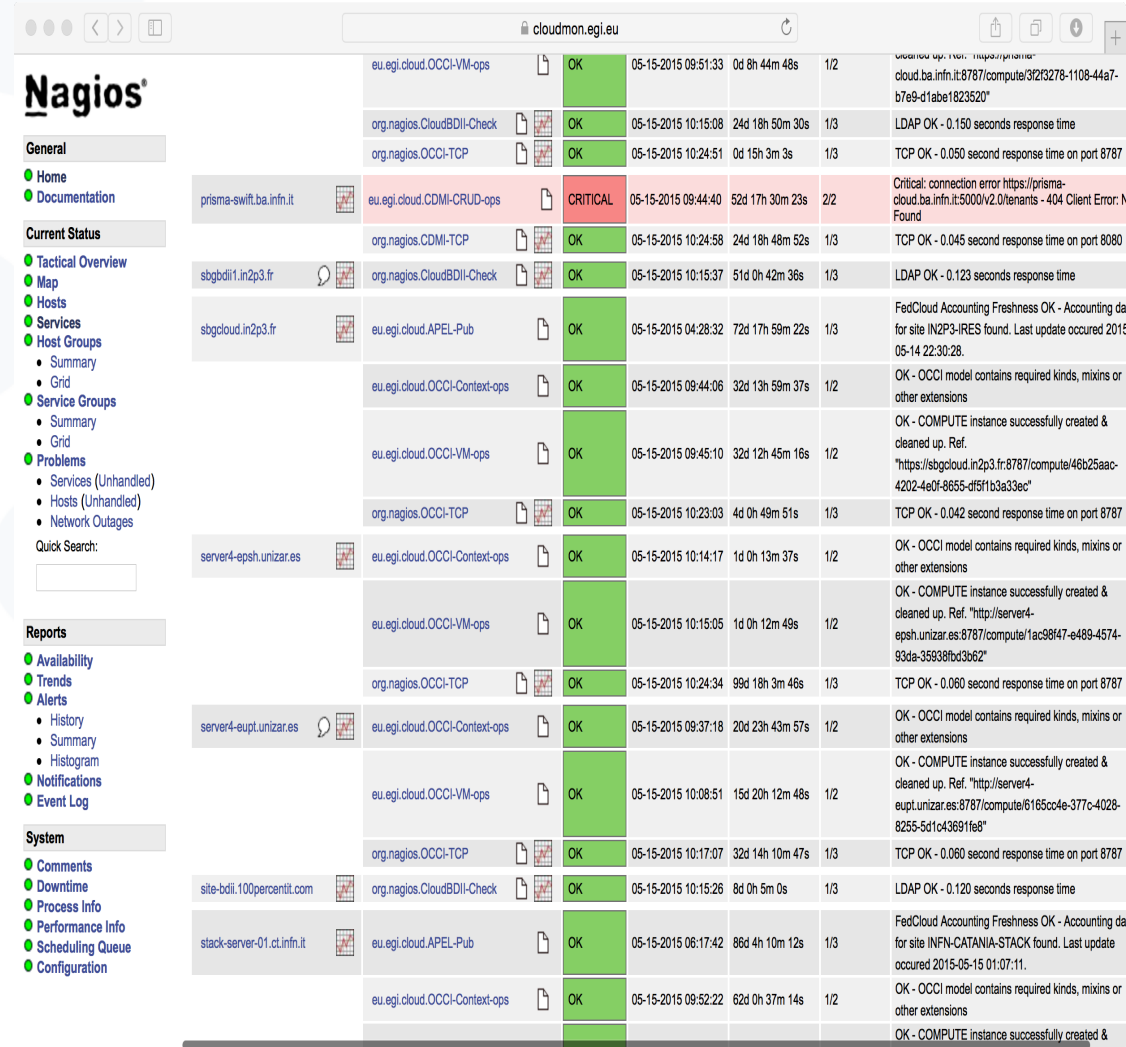
Total number of VM run by SITE and DATE.  
VOs. June 2014 - May 2015.

The following table shows the distribution of Total number of VM run grouped by SITE and DATE.

SITE	Total number of VM run by SITE and DATE												Total	%
	Jun 2014	Jul 2014	Aug 2014	Sep 2014	Oct 2014	Nov 2014	Dec 2014	Jan 2015	Feb 2015	Mar 2015	Apr 2015	May 2015		
100IT	726	746	509	724	733	716	709	748	675	743	566	298	7,893	1.20%
BIFI	1,230	1,366	1,856	5,674	13,017	11,113	15,293	18,209	42,974	50,858	2,469	5,832	169,891	25.93%
CERN-PROD	13,279	3,017	8,254	49,872	47,210	13,310	185	437	8,086	172	4,259	358	148,439	22.66%
CESGA	2,539	1,941	1,568	2,005	1,854	1,726	929	0	1	556	679	330	14,128	2.16%
CESNET-MetaCloud	774	760	768	792	802	1,258	1,407	792	837	770	931	1,011	10,902	1.66%
CETA-GRID	0	286	734	616	238	153	99	2,308	229	454	613	715	6,445	0.98%
CYFRONET-CLOUD	186	0	3	41	649	456	0	0	235	751	356	252	2,929	0.45%
FZJ	1,387	1,421	717	901	3,922	5,530	17,979	10,385	21,145	3,088	5,766	4,653	76,894	11.74%
GoeGrid	1,427	1,500	1,242	1,460	1,507	1,504	798	751	661	755	782	565	12,972	1.98%
HG-09-Okeanos-Cloud	530	368	465	107	567	493	229	241	185	179	0	0	3,364	0.51%
IFCA-LCG2	1,115	1,546	5,814	4,070	7,346	4,426	650	621	539	537	344	172	27,180	4.15%
ISAS-FedCloud	739	747	731	726	844	2,970	6,014	5,357	2,766	3,436	4,435	833	29,598	4.52%
IN2P3-IRES	0	0	0	0	0	0	0	0	188	729	503	332	1,752	0.27%
INFN-CATANIA-NEBULA	474	993	794	803	804	663	705	779	672	768	737	469	8,661	1.32%
INFN-CATANIA-STACK	47	7	7	231	752	115	202	847	697	774	572	486	4,737	0.72%
INFN-PADOVA-STACK	732	845	755	757	5,840	5,482	8,652	13,131	11,162	7,663	25,502	3,361	83,902	12.81%
MK-04-FINKICLOUD	18	570	191	1,372	1,563	512	680	1,493	1,373	975	709	60	9,516	1.45%
NGC-INGRID-PT	0	0	0	4	406	667	674	712	593	714	563	254	4,587	0.70%
PRISMA-INFN-BARI	754	595	18	2,036	4,745	2,480	754	791	210	219	1,433	1,665	15,680	2.39%
SZTAKI	775	1,036	1,556	876	982	551	572	487	707	31	0	0	7,573	1.16%
TR-FC1-JLAKBIM	562	693	273	43	0	157	147	752	707	499	191	0	4,024	0.61%
UPV-GRYCAP	0	0	0	0	5	756	574	257	888	600	738	221	4,039	0.62%
<b>Total</b>	<b>27,294</b>	<b>18,437</b>	<b>26,255</b>	<b>73,110</b>	<b>93,786</b>	<b>55,018</b>	<b>57,252</b>	<b>59,098</b>	<b>95,550</b>	<b>75,291</b>	<b>52,148</b>	<b>21,867</b>	<b>655,106</b>	
<b>Percentage</b>	<b>4.17%</b>	<b>2.81%</b>	<b>4.01%</b>	<b>11.16%</b>	<b>14.32%</b>	<b>8.40%</b>	<b>8.74%</b>	<b>9.02%</b>	<b>14.59%</b>	<b>11.49%</b>	<b>7.96%</b>	<b>3.34%</b>		

[Click here for XML, encoded data](#)

- Ensure service availability and reliability
- Current probes:
  - OCCI, APEL, Perun
- On development:
  - CDMI, vmcatcher
- A/R reports



The screenshot shows the Nagios monitoring interface for 'cloudmon.egi.eu'. The interface includes a sidebar with navigation options like 'General', 'Current Status', 'Reports', and 'System'. The main area displays a table of services with columns for host, service name, status, last check time, and other details.

Host	Service	Status	Last Check Time	Duration	Attempts	Output
eu.egi.cloud	OCCI-VM-ops	OK	05-15-2015 09:51:33	0d 8h 44m 48s	1/2	cleaned up. Ref. "https://prisma-cloud.ba.inf.it:8787/compute/3f23278-1108-44a7-b7e9-d1abe1823520"
org.nagios	CloudBDII-Check	OK	05-15-2015 10:15:08	24d 18h 50m 30s	1/3	LDAP OK - 0.150 seconds response time
org.nagios	OCCI-TCP	OK	05-15-2015 10:24:51	0d 15h 3m 3s	1/3	TCP OK - 0.050 second response time on port 8787
prisma-swift.ba.inf.it	eu.egi.cloud.CDMI-CRUD-ops	CRITICAL	05-15-2015 09:44:40	52d 17h 30m 23s	2/2	Critical: connection error https://prisma-cloud.ba.inf.it:5000/v2.0/tenants - 404 Client Error: Not Found
org.nagios	CDMI-TCP	OK	05-15-2015 10:24:58	24d 18h 48m 52s	1/3	TCP OK - 0.045 second response time on port 8080
sbgbdi1.in2p3.fr	org.nagios.CloudBDII-Check	OK	05-15-2015 10:15:37	51d 0h 42m 36s	1/3	LDAP OK - 0.123 seconds response time
sbglcloud.in2p3.fr	eu.egi.cloud.APEL-Pub	OK	05-15-2015 04:28:32	72d 17h 59m 22s	1/3	FedCloud Accounting Freshness OK - Accounting data for site IN2P3-RES found. Last update occurred 2015-05-14 22:30:28.
eu.egi.cloud	OCCI-Context-ops	OK	05-15-2015 09:44:06	32d 13h 59m 37s	1/2	OK - OCCI model contains required kinds, mixins or other extensions
eu.egi.cloud	OCCI-VM-ops	OK	05-15-2015 09:45:10	32d 12h 45m 16s	1/2	OK - COMPUTE instance successfully created & cleaned up. Ref. "https://sbglcloud.in2p3.fr:8787/compute/46b25aac-4202-4e0f-8655-df5f1b3a33ec"
org.nagios	OCCI-TCP	OK	05-15-2015 10:23:03	4d 0h 49m 51s	1/3	TCP OK - 0.042 second response time on port 8787
server4-epsh.unizar.es	eu.egi.cloud.OCCI-Context-ops	OK	05-15-2015 10:14:17	1d 0h 13m 37s	1/2	OK - OCCI model contains required kinds, mixins or other extensions
eu.egi.cloud	OCCI-VM-ops	OK	05-15-2015 10:15:05	1d 0h 12m 49s	1/2	OK - COMPUTE instance successfully created & cleaned up. Ref. "http://server4-epsh.unizar.es:8787/compute/1ac98f47-e489-4574-93da-35938fbd3b62"
org.nagios	OCCI-TCP	OK	05-15-2015 10:24:34	99d 18h 3m 46s	1/3	TCP OK - 0.060 second response time on port 8787
server4-eupt.unizar.es	eu.egi.cloud.OCCI-Context-ops	OK	05-15-2015 09:37:18	20d 23h 43m 57s	1/2	OK - OCCI model contains required kinds, mixins or other extensions
eu.egi.cloud	OCCI-VM-ops	OK	05-15-2015 10:08:51	15d 20h 12m 48s	1/2	OK - COMPUTE instance successfully created & cleaned up. Ref. "http://server4-eupt.unizar.es:8787/compute/6165cc4e-377c-4028-8255-5d1c43691fe8"
org.nagios	OCCI-TCP	OK	05-15-2015 10:17:07	32d 14h 10m 47s	1/3	TCP OK - 0.060 second response time on port 8787
site-bdii.100percentit.com	org.nagios.CloudBDII-Check	OK	05-15-2015 10:15:26	8d 0h 5m 0s	1/3	LDAP OK - 0.120 seconds response time
stack-server-01.ct.infn.it	eu.egi.cloud.APEL-Pub	OK	05-15-2015 08:17:42	86d 4h 10m 12s	1/3	FedCloud Accounting Freshness OK - Accounting data for site INFN-CATANIA-STACK found. Last update occurred 2015-05-15 01:07:11.
eu.egi.cloud	OCCI-Context-ops	OK	05-15-2015 09:52:22	62d 0h 37m 14s	1/2	OK - OCCI model contains required kinds, mixins or other extensions
						OK - COMPUTE instance successfully created &

# Integration Status

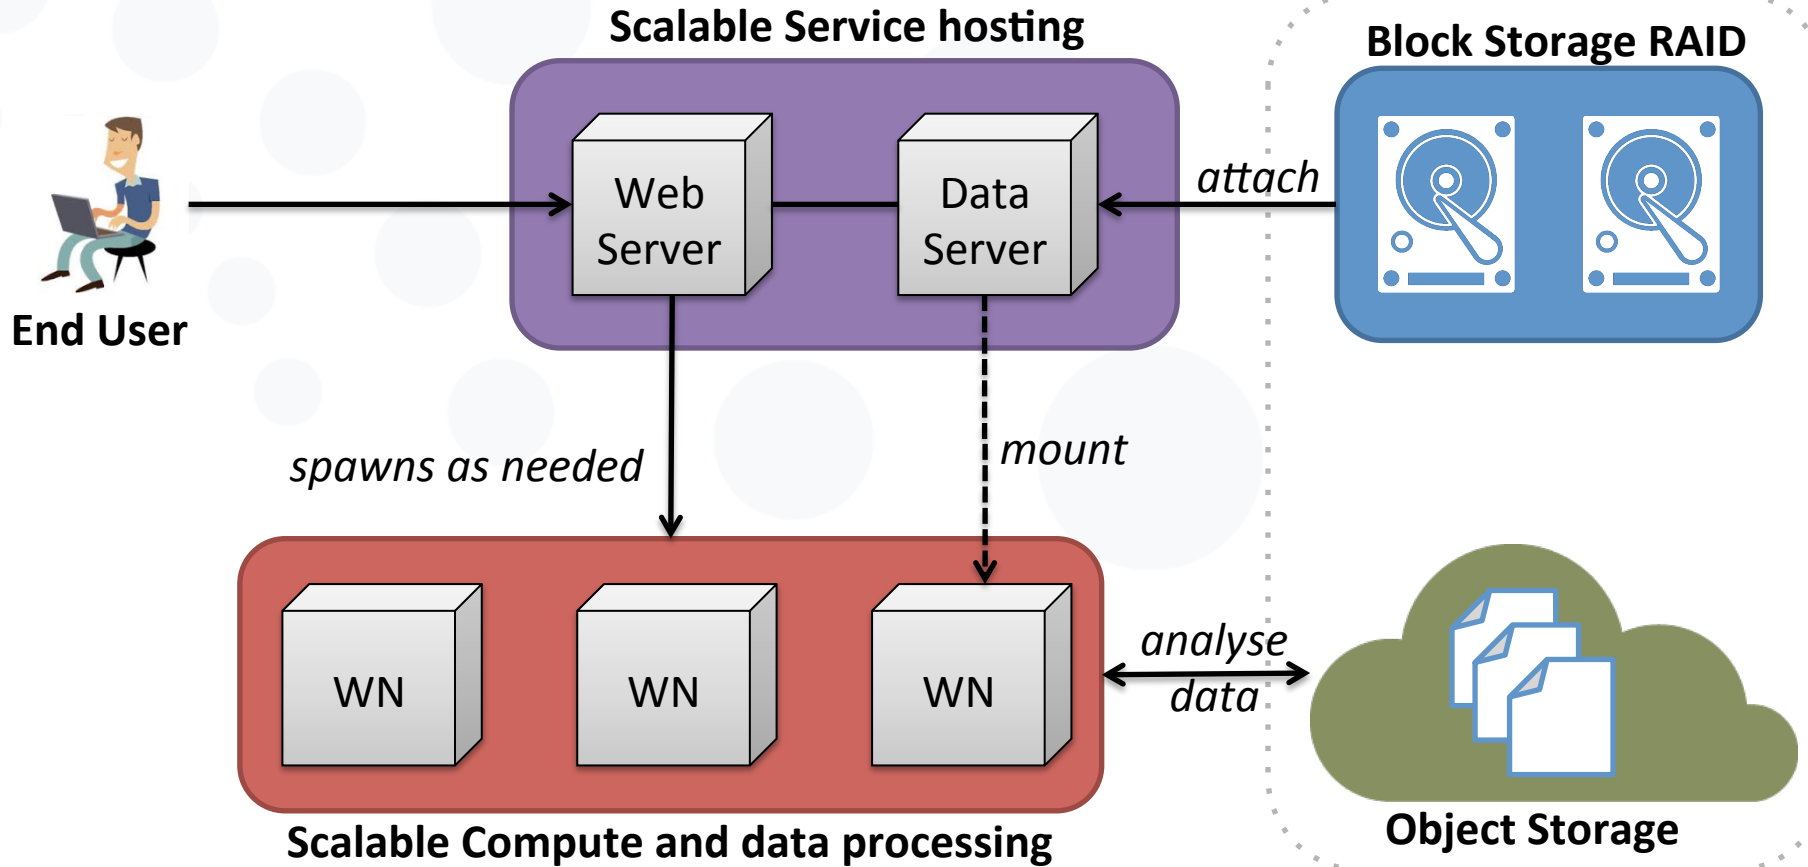
Capability	OpenStack	OpenNebula	Synnefo
OCCI VM Management	Yes	Yes	Yes
Integration with VM image management	Yes	Yes	Yes
Contextualization	Yes	Yes	Yes
OCCI Block Storage	Yes	Yes	In Progress
CDMI Object Storage	In Progress	N/A	Yes
EGI AAI Integration	Yes	Yes	Yes
Monitoring	Yes	Yes	Yes
Accounting records	Yes	Yes	Yes
Information System	Yes	Yes	Yes

- **Resource Provider questionnaire**
  - Checklist in form of survey concerning management of Resource Providers related to security
  - Mandatory to join the infrastructure
- **Security Workgroup:**
  - Security monitoring to intercept flaws runtime
  - Deal with security issues appearing in the infrastructure
  - Run security challenges
  - Assistance and guidelines to create secure VM images (especially EGI endorsed ones)

FedCloud extends EGI computing beyond HTC with new usage models:

- **Service Hosting**
  - Long-running services (e.g. web, database or application servers)
- **Compute and data intensive workloads**
  - Batch and interactive (e.g. IPython, R, matlab) with scalable and customized environments not limited to the traditional job model
- **Datasets repository**
  - Store and manage large datasets for your applications
- **Disposable and testing environments**
  - Host training events, test new developments and applications without overhead

Combine usage models in a single application



# High Level Tools (PaaS & SaaS)

- Extend the IaaS capabilities of the EGI cloud
  - ‘Alternatives’ of the OCCI client and API
  - More than OCCI
- External contributions (→ support many other clouds too)
- New developments expected (e.g. INDIGO-Datacloud)

Catania  
Science  
Gateway

- SaaS
- Identity Federation

VMDIRAC

- Abstraction on top of various HPC/HTC/cloud

Slipstream

- PaaS for automating deployments
- Helix Nebula

WS-PGRADE

- Workflow development and enactment

COMPSs

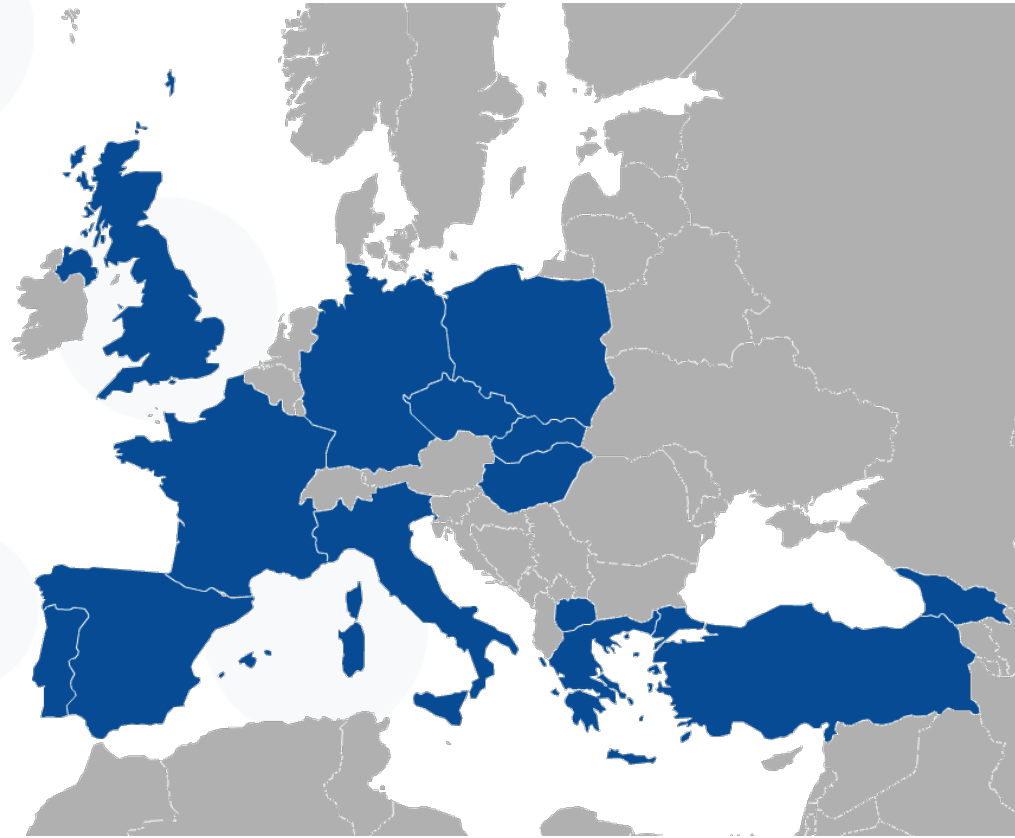
- Programming framework for auto-parallelisation

Vcycle

- VM lifecycle manager

# FedCloud Infrastructure

- One year of production
- Resources
  - 21 providers from 14 NGIs
  - 17 interested in joining from 7 new NGIs
  - Average A/R 82%/83%
- Usage
  - ~700K VMs
  - ~9M CPU hours wall time





- Resource Providers
  - Follow the EGI Certification Procedure <https://wiki.egi.eu/wiki/PROC09>, common for HTC and Cloud Resources
  - Use the installation manuals <https://wiki.egi.eu/wiki/MAN10> for your CMF
- Technology Providers
  - Contact the TF mailing list [fedcloud-tf@mailman.egi.eu](mailto:fedcloud-tf@mailman.egi.eu)
- Users
  - Contact [support@egi.eu](mailto:support@egi.eu), see next talk by Diego

- The **EGI Federated Cloud Task Force** continue to work to **evolve** and **improve** the Cloud Infrastructure
- Feature additions to foundational tools depending on requests
  - systematic tracking of requests with EGI RT
- New scenarios/working groups
  - Intra-cloud networking
  - Security
  - VM Image Endorsement

# New developments on the roadmap

- VM Management
  - **OCCI v1.2** now in public comment phase, introduces JSON rendering
  - Vertical scalability
  - Snapshots and migration of VMs between sites
- Data Management
  - **CDMI for OpenStack** based on Synnefo implementation
  - Integration of new stacks (**OneData**)
- **AppDB as a broker**
  - Simple manage operations on VMs from web interface
- Support for **native interfaces**
- Integration of **CANFAR** and **D4Science**

- **EGI Federated Cloud extends EGI infrastructure**
  - Fully integrating IaaS into EGI production infrastructure
  - Attracting new communities belonging to various scientific domains (76 use cases, 40 communities)
- **Paving the way for a global federated cloud marketplace** led through European Innovation
  - Established best practice
  - Illustrating leadership
  - Open standards, open technology
  - Open membership, open processes
  - Open competition

# Thank you for your attention.

*Questions?*



[www.egi.eu](http://www.egi.eu)

This work by Parties of the EGI-Engage Consortium is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

