

The D4Science Production Infrastructure

Pedro Andrade¹, Leonardo Candela^{2*}, Donatella Castelli², Andrea Manzi¹, Pasquale Pagano²

Abstract

This report documents the state of the D4Science production infrastructure at the end of the homonymous project. This infrastructure is characterized by the set of hardware and software resources, and by procedures deployed to provide D4Science user communities with a reliable production service. In particular the report describes the sites and nodes allocated to the infrastructure and how these have been evolving since the setting up of the production infrastructure. It also presents the portal of the infrastructure giving details of the existing virtual research environments (VREs) and associated users communities. Overall, the infrastructure managed (i) 42 content collections containing more than 28,800 compound Information Objects, (ii) 97 metadata collections with multiple metadata schema including Dublin Core, ISO19115 and proprietary formats containing more than 54,300 metadata records, and (iii) more than 380 WS-Resources representing full text, forward and geo indices supporting the search over these data.

Keywords

Hybrid Data Infrastructure — Virtual Research Environment — Interoperability

¹European Organisation for Nuclear Research CERN, CH – 1211 Geneva 23, Switzerland

²Istituto di Scienza e Tecnologie dell'Informazione "A. Faedo", Consiglio Nazionale delle Ricerche, Via G. Moruzzi 1, 56124, Pisa, Italy

*Corresponding author: leonardo.candela@isti.cnr.it

Contents

Introduction	1
1 Node Operation	2
1.1 Resources Allocated	2
1.2 Node Deployment and Upgrade	2
2 VOs and VREs Operation	3
2.1 Procedures	4
2.2 EM VO	5
EM Resources • Global Ocean Chlorophyll Monitoring VRE • Global Land Vegetation Index VRE	
2.3 FARM VO	6
FARM Resources • AquaMaps VRE • Fishery Country Profiles Production System VRE • Integrated Capture Information System VRE	
3 Infrastructure Monitoring & Accounting	7
3.1 Information System based Tools	7
3.2 Message System based Tools	7
3.3 Statistics	9
4 Conclusions	10
Acknowledgments	11
References	11

Introduction

The D4Science infrastructure is a gCube-based [1, 2, 3] IT infrastructure that leverages Grid [4], Digital Library [5] and Service-orientation [6] principles and approaches to deliver a number of data management facilities *as-a-Service*.

This infrastructure supports the activities of two user communities operating in the Environmental Monitoring (EM)

and Fishery and Aquaculture Resource Management (FARM) domains. These communities access the infrastructure via the exploitation of *Virtual Research Environments* (VREs) [7] that satisfy concrete needs of the communities' *Virtual Organizations* (VOs) [8].

The infrastructure is built by sites coming from different organizations participating in the D4Science project [9], i.e., CNR, ESA, FAO, National Kapodistrian Univ. of Athens (NKUA), and Univ. of Basel (UNIBASEL).

The number of sites involved and the number of nodes allocated has gradually grown over the time to meet the demand related to the provision of new VOs and VREs. Currently, these sites provide a total of 53 physical machines, offering 408 GB RAM, 38 TB disk space, and 166 processor cores. These nodes are exploited to run the gCube system delivered by D4Science and the gLite middleware [10] as released by the EGEE [11] project. As a result of the deployment of gCube and gLite different VOs and VREs are made available: one VO for EM containing 2 VREs, and one VO for FARM containing 3 VREs.

The management of the infrastructure is facilitated by the definition and implementation of clear procedures for monitoring and accounting. A number of monitoring tools have been deployed allowing different infrastructures roles (Site Managers, Data Managers, etc) to visualize the status of their resources and to be actively notified when problems occurred. An accounting tool has also been put in production providing relevant statistics about the users' exploitation of the infrastructure, and the service-to-service communications at infrastructure level.

A clear production support procedure ensures an efficient response to all defects found in the production environment.

The rest of this report is structured as follows: Section 1 describes the infrastructure in terms of its nodes; Section 2 describes the VOs and VREs deployed in the infrastructure; and Section 3 presents the activities related to the definition and implementation of procedures for monitoring, accounting, and supporting the infrastructure. Section 4 concludes the report by summarising the state of the infrastructure.

1. Node Operation

This section describes the hardware resources that compose the D4Science production infrastructure and the activities carried out during the project lifetime to operate such nodes. These activities include in particular the deployment and upgrade of the gCube software [2] and the gLite middleware [10] in the infrastructure, the certification of the infrastructure nodes, and the management of infrastructure downtimes.

The objective of the D4Science production infrastructure is to provide Virtual Research Environments to the project user communities. These VREs are defined and deployed to serve the needs of distributed Virtual Organizations. The user communities involved in D4Science may be represented by one or more VOs. To support the provision of such diverse environments the gLite and gCube nodes play a fundamental role:

- gCube nodes are the hardware resources able to run gCube services. gCube is the software initially developed by the DILIGENT project and consolidated during D4Science. The gCube software includes a special web service container, the gCube Hosting Node (gHN), and a set of services and libraries that provide the functionality to create, manage, and exploit VREs;
- gLite nodes are computing and storage nodes running gLite software. gLite is the middleware provided by the EGEE project. By running gLite, these nodes provide core grid functionalities such as file-based storage, distributed computation of applications, etc. gLite nodes are exploited by gCube services which then provide higher level functionality through the D4Science VREs.

The operation of these nodes includes the management of the physical machines that host the nodes, and the tasks executed directly on the nodes to maintain them operational.

1.1 Resources Allocated

The deployment plan of the D4Science production infrastructure defined that new resources and new sites should have been gradually added to the infrastructure. Similarly, new VOs and VREs should have been gradually deployed and made available to the user communities. Table 1 shows the actual achievements for each site of the production infrastructure until now.

During an initial period, the sites of the infrastructure were provided by CNR, NKUA, and UNIBASEL. Sites from ESA, FAO, and WorldFish were added in a second phase.

Table 1. Production infrastructure evolution

Site	Jun '08	Jul '08	Dec '08
CNR	yes	yes	yes
ESA	yes	yes	yes
FAO	no	no	yes
NKUA	yes	yes	yes
UNIBASEL	yes	yes	yes
WorldFish	no	no	no

VOs and VREs were also gradually made available. Initially only the EM community was served through the EM VO, later the infrastructure was extended to the FARM community with the addition of new VOs/VREs.

Currently, five sites provide hardware resources to the production infrastructure. Table 2 shows information about the characteristics of the machines provided by each site.

Figure 1 shows the evolution of the hardware resources provided to the infrastructure since the launch of the production infrastructure. The graph shows an increase of more than 100% since the first infrastructure milestone. CNR largely contributes to this infrastructure by providing a number of nodes close to 50% of the total ones. The figure also shows the visible growth verified after one year from the launch of the production infrastructure when the FARM user community was also provided with VREs.

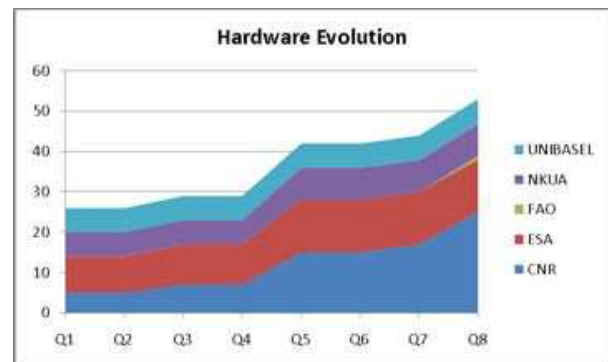


Figure 1. Hardware resources evolution

1.2 Node Deployment and Upgrade

The resources allocated to the D4Science production infrastructure have been exploited to host gCube and gLite nodes. Many of the hardware resources have been virtualized in order to provide a higher number of gCube and gLite nodes. The rest of this section describes the deployment and upgrade of these nodes.

gCube Nodes Five organizations deployed gCube nodes in their site. Each gCube node was configured to support one of the two VOs available in the infrastructure. Table 3 shows the number of gCube nodes deployed in each site. The total number of gCube nodes was distributed equally between

Table 2. Hardware resources by infrastructure contributors

Site	No.	CPU	RAM (GB)	Disk (TB)	Cores
CNR	10	2x AMD Opteron @ 2.5 GHz	80	1.5	20
	4	2x AMD Opteron Quad-core	128	1.1	32
	4	Intel Core2 Quad @ 2.4 GHz	32	15	16
	2	Dual-Core AMD Opteron 1222	12	1	4
	2	AMD Opteron @ 1.6 GHz	12	0.64	4
	1	AMD Athlon 64 Dual Core 5000+	8	0.5	2
	1	AMD Athlon 64 Dual Core 4200+	8	0.5	2
	1	Intel Pentium III @ 1400MHz	3	0.148	2
ESA	7	Intel Xeon @ 3.20 GHz	16	0.574	28
	1	Intel Xeon @ 3.20 GHz	4	0.340	4
	1	Intel Xeon @ 2.80 GHz	4	10	4
	4	-	16	-	11
FAO	1	2 Quad Core Intel Xeon @ 3.0 GHz	8	0.292	8
NKUA	5	Dual Core Intel @ 2,40 GHz	40	2.9	10
	3	Intel Xeon @ 2.4 GHz	24	1.5	12
UNIBASEL	2	Single Core @ 1.8 GHz	4	2.3 TB	2
	3	Single Core @ 2.4GHz	6	-	3
	1	Dual Core @ 2.8 GHz	3	-	2
Total	53		408	38.294	166

the two infrastructure VOs. The 60 gCube nodes currently deployed in the infrastructure host gCube services for 2 VOs and 7 VREs.

Table 3. gCube nodes per organization

Organisation	Total	EM VO	FARM VO
CNR	38	18	19
ESA	2	2	0
FAO	2	0	2
NKUA	12	7	5
UNIBASEL	6	3	3
Total	60	30	29

The current gCube release deployed in production is release 1.5. Figure 2 shows the evolution of the number of hours needed to upgrade the infrastructure when a new release was available. While the first infrastructure upgrade took weeks, the most recent upgrades were carried out in few hours. This is a direct consequence of several improvements introduced in gCube related to infrastructure operation.

gLite Nodes Table 5 presents how the different gLite services have been distributed between the infrastructure sites. Of the five sites providing resources to the infrastructure, three sites deployed gLite nodes in their site.

The services marked with “yes*” represent a primary instance and with orange symbol a secondary instance. All gLite sites provided by D4Science contributors are also registered as sites of the EGEE project production infrastructure.

**Figure 2.** gCube nodes upgrade duration

Other gLite sites present in the EGEE production infrastructure also support the D4Science VO. These sites are: csTCDie, INFN Bari, INFN Trieste, and Taiwan.

In terms of workers nodes and storage space the following total capacity was made available to the d4science.research-infrastructures.eu VO.

2. VOs and VREs Operation

The D4Science infrastructure supports the operation of two Virtual Organisations – the *Environmental Monitoring* (EM) Virtual Organisation and the *Fisheries and Aquaculture Resources Management* (FARM) Virtual Organisation – to serve the needs of the two communities the project has been conceived to serve. These two Virtual Organisations provide fa-

Table 4. gLite nodes per organization

Organisation	CE	WN	sBDII	SE	WMS	LFC	VOMS	MyProxy
CNR	yes	yes	yes	yes	yes	-	yes	yes
ESA	yes	yes	yes	yes	-	yes	yes*	-
NKUA	yes	yes	yes	yes	yes*	-	-	-

Table 5. gLite nodes capacity

Organisation	Project Site	WNs	Storage
CNR	yes	10	20 GB
ESA	yes	39	16 TB
NKUA	yes	10	308 GB
csTCDie	no	755	126 TB
INFN Bari	no	578	-
INFN Trieste	no	272	102 TB
Taiwan	no	1377	-
Total	-	3041	244 TB

cilities and resources for the operation of 5 Virtual Research Environments – the *Global Ocean Chlorophyll Monitoring* (GCM) and the *Global Land Vegetation Index* (GVM) VREs in the EM VO; the *Fishery Country Profiles Production System* (FCPPS), *Integrated Capture Information System* (ICIS) and the *AquaMaps* (AquaMaps) VREs in the FARM VO.

In the remainder of this section, the procedures governing the operation of the supported VOs and VREs as well as their characteristics in terms of deployed resources are described.

2.1 Procedures

The procedures governing VOs and VREs operation are described in the Production Infrastructure wiki¹ while detailed information on how to put them in practice are described in the D4Science Administrator's Guide [12].

These procedures regulate the following two major activities: *VO deployment* and *VRE deployment*.

VO Deployment

- A VO enabling services deployment phase in which a VO Manager installs the mandatory services of the enabling layer needed to glue together the VO nodes;²
- A VO level services deployment phase in which a VO Manager installs the set of services optionally deployed to support a specific deployment scenario;³

¹<https://infrastructure.wiki.d4science.research-infrastructures.eu/production/>

²https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/VO_installation

³https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/Dynamic_VO_Services_Deployment_and_Configuration

- A VO configuration phase in which a VO Manager takes care of creating the conditions needed to instruct the VO enabling services and VO level services to implement the defined VO (e.g., preparing the service map, registering the mandatory resources);⁴
- A VO management phase in which a VO Manager takes care of modifying the conditions governing the VO operation including the resources forming the VO;⁵

VRE Deployment

- A definition phase in which a user having the role of VRE Designer specifies the characteristics of a new VRE conceived to serve an application scenario;⁶
- An approval phase in which a user having the role of VRE Manager decides whether the specified VREs have to be accepted or rejected. For what is concerned with the accepted VRE, the VRE Manager decides also how this VRE has to be deployed, e.g., which hosting nodes will be exploited;⁷
- A verification phase in which a user having the role of VRE Manager validates a VRE resulting from the approval phase;⁸
- A management phase in which a user having the role of VRE Manager operates on a deployed VRE in order to customise specific aspects (e.g., the layout governing the placement of user interfaces constituents a.k.a. portlets) or monitor the operational state of the VRE as a whole;⁹

⁴https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/VO_configuration

⁵https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/Managing_the_VO

⁶https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/VRE_Administration#VRE_Definition

⁷https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/VRE_Administration#VRE_Approval

⁸https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/VRE_Administration#VRE_Verification

⁹https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/VRE_Administration#VRE_Management

2.2 EM VO

The EM Virtual Organisation is the dynamic group of individuals and/or institutions defined around a set of sharing rules in which resource providers and consumers specify clearly and carefully just what is shared, who is allowed to share, and the conditions under which sharing occurs to serve the needs of the Environmental Monitoring community. The activities carried on in monitoring the environment require extremely sophisticated means, including time and space instrumentation, analysis tools, integration and correlation of different data sources, reasoning, information and knowledge management.

This Virtual Organisation is characterised by a pool of Resources and support the operation of two Virtual Research Environments: the Global Ocean Chlorophyll Monitoring VRE and the Global Land Vegetation Monitoring VRE.

2.2.1 EM Resources

The pool of resources forming a Virtual Organisation is dynamic and evolving, i.e., new resources can dynamically appear or existing resources relocated or dismissed as a consequence of VO Manager explicit actions or because of the autonomic behaviour of the enabling services. The pool of resources forming a Virtual Organisation can be grouped in five main categories:

Hosting Nodes The EM VO has 30 nodes spread in 5 different sites;

Collections The EM VO has 25 content collections and 69 metadata collections. These collections contain various types of material ranging from European Environmental Agency Reports to MERIS Level-3 products, Earth images and AATSR data sets;

Software The EM VO has been equipped with all the various gCube releases produced by the project including the latest (1.5.0). The majority of the constituents of these releases have been deployed leading thus to running instances;

Running Instances The EM VO is equipped with the number of components needed to guarantee the operation of the planned VREs. The approximate number of instances (including the enabling services) is 200;

Generic Resources The EM VO is equipped with approximately 100 of such a kind of resources. These resources have been produced by the VO Manager to configure the behaviour of other resources and include data transformation programs, data import scripts, and data presentation formats.

2.2.2 Global Ocean Chlorophyll Monitoring VRE

The phytoplankton plays a similar role to terrestrial green plants in the photosynthetic process and is credited with removing as much carbon dioxide from the atmosphere as their earthbound counterparts, making it important to monitor and

model plankton into calculations of future climate change. Global Ocean Chlorophyll Monitoring (GCM) is a VRE equipped with data sources and tools supporting ocean chlorophyll monitoring. In particular, this VRE provides the community with tools for a seamless access to heterogeneous data sources including:

- Earth images collected and distributed by ESA and acquired from different instruments on-board of several Earth Observation satellites such as Envisat;
- Advanced Along-Track Scanning Radiometer (AATSR) products;
- Chlorophyll-a, case-1 water (chl1) yearly level 3 products;
- Mediterranean Action Plan (MAP) Technical Reports;
- European Environment Agency material including Reports, Briefings, Brochures, Corporate documents, Environmental issue reports and other reports;
- Reports on the needs of the scientific communities from a social trend view point.

Such data sources contain compound objects of different media ranging from pdf files to tiff images, jpeg images, references, xml files.

From a functional point of view this environment supports:

- The discovery of these information objects through browsing, simple and advanced search including geo-referenced search;
- The generation of new products through the gPod service;
- The collaboration through annotations;
- The collaboration through a shared workspace;
- The collaborative production of reports (i.e. enhanced publications produced through the VRE support).

2.2.3 Global Land Vegetation Index VRE

Vegetative land cover is an important variable in many Earth system processes. Global Land Vegetation Monitoring (GVM) Virtual Research Environment is equipped with data sources and tools for supporting vegetative land cover related activities. In particular, this VRE provides its community with tools for a seamless access to heterogeneous data sources including:

- Earth images collected and distributed by ESA and acquired from different instruments on-board of several Earth Observation satellites such as Envisat;
- European Environment Agency material including Reports, Briefings, Brochures, Corporate documents, Environmental issue reports and other reports;

- Landsat7 images covering different European cities as well as lands (including forests and cropland) belonging to different European Countries;
- Global Vegetation Index yearly level 3 products;
- Reports on the needs of the scientific communities from a social trend view point.

Such data sources contain compound information objects of different media ranging from pdf files to tiff images, jpeg images, references, xml files.

From a functional point of view this environment supports:

- The discovery of these information objects through browsing, simple and advanced search including geo-referenced search;
- The generation of new products through the gPod service;
- The collaboration through a shared workspace;
- The collaboration through annotations;
- The collaborative production of reports (i.e. enhanced publications produced through the VRE support).

2.3 FARM VO

The FARM Virtual Organisation is the dynamic group of individuals and/or institutions defined around a set of sharing rules in which resource providers and consumers specify clearly and carefully just what is shared, who is allowed to share, and the conditions under which sharing occurs to serve the needs of the Fisheries and Aquaculture Resources Management. This VO is conceived to support various application scenarios arising in the FARM Community including the production of Fisheries and Aquaculture Country Profiles, the management of catch statistics including harmonisation, the dynamic generation of biodiversity maps and species distribution maps.

This Virtual Organisation is characterised by a pool of Resources and support the operation of three Virtual Research Environments: the *AquaMaps VRE*, the *Fishery Country Profile Production System VRE* and the *Integrated Capture Information System VRE*.

2.3.1 FARM Resources

The pool of resources forming a Virtual Organisation is dynamic and evolving, i.e., new resources can dynamically appear or existing resources relocated or dismissed as a consequence of VO Manager explicit actions or because of the autonomic behaviour of the enabling services. The pool of resources forming a Virtual Organisation can be grouped in five main categories:

Hosting Node The FARM VO has 29 nodes spread in 5 different sites;

Collections The FARM VO has 15 content collections and 28 metadata collections. These collections contain various types of material ranging from AquaMaps products to fact sheets, time series graphs and country maps;

Software The FARM VO has been equipped with all the various gCube releases produced by the project including the latest (1.5.0). The majority of the constituents of these releases have been deployed leading thus to running instances;

Running Instances The FARM VO is equipped with the number of components needed to guarantee the operation of the planned VREs. The approximate number of instances (including the enabling services) is 200;

Generic Resources The FARM VO is equipped with approximately 80 of such a kind of resources. These resources have been produced by the VO Manager to configure the behaviour of other resources and include data transformation programs, data import scripts, and data presentation formats.

2.3.2 AquaMaps VRE

AquaMaps is a Virtual Research Environment conceived to provide fisheries and aquaculture scientists with AquaMaps objects produced by D4Science. AquaMaps are predictive documents resulting in Earth maps enriched thus to show the likelihood that a certain species or a combination of species will live in specific regions or areas. AquaMaps are very important documents for species assessment. This environment permits the biodiversity community to establish/predict species geographic distribution based on species ecological envelopes according to a transparent algorithm that can be handled by species experts.

2.3.3 Fishery Country Profiles Production System VRE

The Fishery Country Profile Production System (FCPPS) Virtual Research Environment is for fisheries and aquaculture authors, managers and researchers who produce reports containing country-level data. The FCPPS is an integrated environment that manages access to multiple data sources divided by domain, including their annotation and versioning and permits production of structured text, tables, charts and graphs from these sources to be easily inserted into custom reporting templates that can be output in multiple formats. Unlike Microsoft Office this VRE has a workflow and structured data that enables collaboration, annotation and versioning of the underlying data sources that make up the final product. It enforces a structured process and structured data while allowing users flexibility in actual content generation.

2.3.4 Integrated Capture Information System VRE

The Integrated Capture Information System (ICIS) Virtual Research Environment provides fisheries statisticians with a set of tools to manage their data. Statisticians produce statistics from often very different data sources, and need a controlled process for the ingestion, validation, transformation,

comparison and exploitation of statistical data for the fisheries captures domain. There is a need to:

- Collect and analyze capture data;
- Check and control data and data-quality;
- Produce uniform data-sets for consumption by outside clients, such as the FishStatJ application;
- Provide an authoritative data-repository for fisheries data.

Also, there is a growing interest to share data and use statistical data in other use-scenarios, such as spatial analysis or reporting.

3. Infrastructure Monitoring & Accounting

The monitoring and accounting of the D4Science production infrastructure is based on a number of tools deployed in the infrastructure that can be clustered into two classes:

- Based on information collected by the gCube Information System (IS);
- Based on information produced/consumed by the gCube Messaging System (MS).

While the gCube Information system provides only tools for monitoring, the gCube Messaging System provides tools for monitoring and accounting. Both solutions are presented in the following sections.

3.1 Information System based Tools

The gCube Information System provides 3 tools to monitor the infrastructure status.

gCube Live Monitoring Live monitoring tool that allows to access resource profiles, as they are stored in the gCube Information System. For example, if a gHN is uncertified the tool reports the last time the gHN worked properly, if a service running instances is unreachable it is highlighted in red to simplify its identification. Figure 3 shows how these two cases are visualised.

The tool is accessible at: <http://monitor.d4science.research-infrastructures.eu/monitoring>

gCube Infrastructure Viewer The gCube infrastructure viewer provides simplified views of the status of the infrastructure, easily consumable from any web-enabled device. In particular, it provides access to six different views: (i) GHN view (represents sites and gHNs characterized by the RIs); (ii) Services view (represents class of services, services, and RIs); (iii) Data view (represents collections, metadata collections, and indexes); (iv) Package view (represents the packages distribution on the infrastructure gHNs); (v) Package dependencies view (represents the distribution packages dependencies);

and (vi) Services dependencies view (represents the services functional dependencies). Figures 4 and 5 provide an example of the gHN view and depicts the location of the main interface elements.

The Infrastructure Viewer for the D4Science production infrastructure is accessible at: <http://monitor.d4science.research-infrastructures.eu/iv>.

gCube Resource Management The gCube resource management is a very powerful tool and can be used in two modes:

- Standalone – read-only version offering the possibility to browse the content of the gCube information system;
- Portal – read-and-write version allowing management operations on the infrastructure.

For monitoring purpose only the standalone version of the resource management tool was exploited. Figure 6 provides an example of the look and feel of this tool.

The gCube resource management is accessible at: <http://monitor.d4science.research-infrastructures.eu/rm>

The gCube Messaging System provides tools for monitoring and accounting of the D4Science production infrastructure.

3.2 Message System based Tools

As described in the previous section the gCube IS subsystem provides a number of tools to monitor the resources available in the D4Science production.

This monitoring solution relies on centralized information hosted in the IS service and is extremely powerful since it provide a complete and comprehensive view over the status of the system. However the dependency towards a single centralized service hosting all information about the infrastructure resources brings several limitations such as performance, information persistence, or scalability. In addition, the current IS monitoring tools do not support active notification. In a distributed environment, spanning across several administrative domains and managed by different site managers, it is important to have monitoring tools that (i) alert site managers of infrastructure problems and state changes, (ii) keep historical information to support the analysis of the infrastructure, and (iii) run independently from the other services of the infrastructure. Such characteristics are not present in the gCube IS monitoring solution currently deployed in the D4Science production infrastructure. To overcome such limitations an alternative monitoring solution based on Messaging Systems has been made available.

The D4Science infrastructure is currently composed by hundreds of resources of different types (hardware nodes, running instances, data collections, etc) which are simultaneously exploited by several VREs. The exploitation of such large number of resources should be properly tracked and accurate metrics collected in order to better understand (i) how

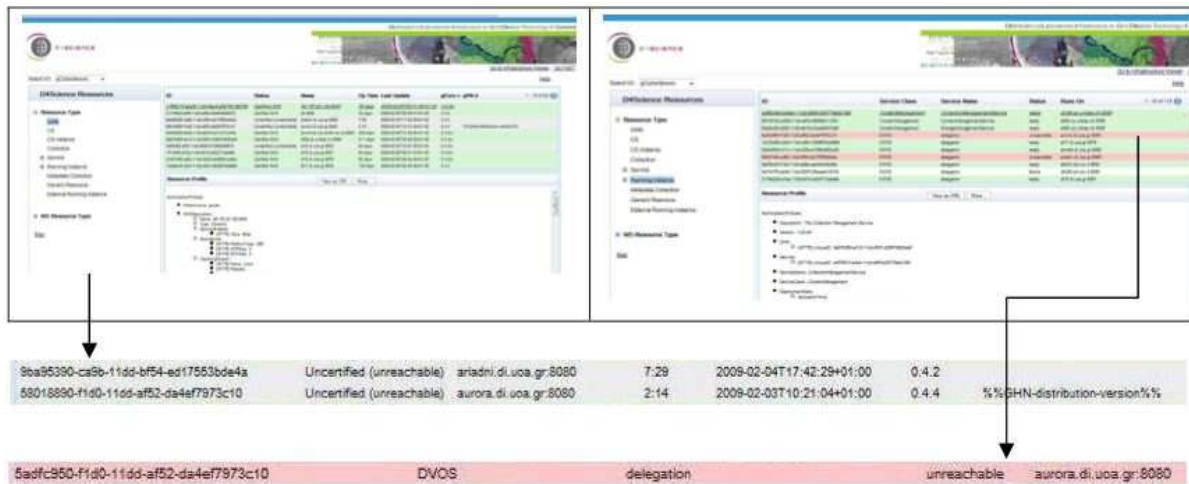


Figure 3. gCube live monitoring

the infrastructure is used, (ii) by whom, (iii) how often, and (iv) how such usage is translated in terms of services communication. Two main types of accounting information are collected: portal access and services invocation.

The gathering of such type of information is implemented by relying on a messaging architecture. This architecture includes the MSG Broker, a broker service that exploits the work performed by the EGEE/LCG projects at CERN. It is disseminated in the gHN container deployed in all gCube nodes of the D4Science production infrastructure. The URL is distributed in the gHN Service Map configuration file. This component is available at: `http://message-broker.d4science.research-infrastructures.eu:8161/admin`

The monitoring probes that are deployed on each node of the infrastructures produce the following type of information:

- Local available disk space;
- CPU load of the gHN;
- Virtual memory available on the gHN;
- gHN HW;
- Local gHN events (scope changed, scope added, node start, etc);
- Local RI events (scope changed, scope added, deployment, etc).

The probes exploit the local producer to send messages to message broker that is the responsible for routing the information to the correct messaging consumer.

The Accounting information produced by the accounting probes, deployed on each gHN of the infrastructure, is of two types: *gCube service usage data* and *User portal operations data*. To the first group belong all the data collected by the gHN each time a RI is invoked. In particular the following

information are aggregated by a local probe and sent to the message broker:

- Running Instance invoked (Service Class and Name)
- Method invoked
- Caller scope
- Caller gHN
- Invocation Time

On the other hand, user portal operations data are produced only on the gHN where a D4science portal is installed. The operation traced by the portal and collected and sent by the related probe are:

- Login
- Browse Collection
- Simple Search
- Advanced Search
- Google Search
- Quick Search
- Content Retrieval

Each type of operation is accomplished with common detailed information (user, scope, timestamp) and with tailored info (e.g., term of simple search, list of collections for advanced search).

For both types of accounting, local aggregation of data is performed in order not to flood the infrastructure with a high number of messages.

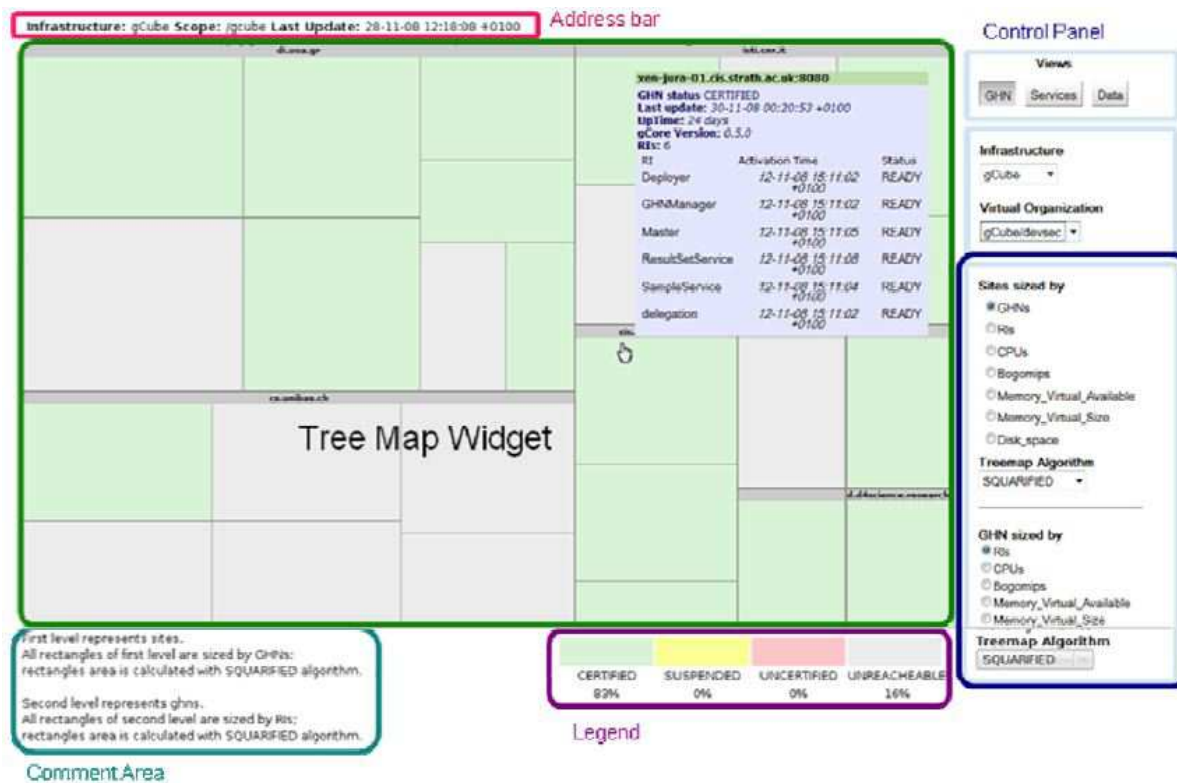


Figure 4. gCube infrastructure viewer (gHN view)

3.3 Statistics

This section presents monitoring and accounting data collected by the gCube messaging system. This data refers to the period from October 2009 to December 2009.

Monitoring The monitoring information presented in this section was extracted from the monitoring messages produced by all gCube nodes deployed in the infrastructure. Table 7 and Figures 7, 8 and 9 show monitoring data averages per infrastructure site in terms of memory available (MEM_AVAIL), disk space (DISK_QUOTA), and CPU usage (CPU_LOAD).

Table 6. Monitoring summary statistics

	CPU_LOAD	DISK_QUOTA	MEM_AVAIL
CNR	1.1045	12594140	903906
NKUA	0.0266	39967909	594633
UNIBAS	0.0048	272811616	444090
FAO	1.0095	2605963	542738
ESA	0.2794	40164805	335337
All	0.7259	46877233	759940

The messaging consumer also traces the alarms that are raised to VO-Admins when particular monitoring thresholds are not respected. After a tuning phase on the testing environment, the following thresholds have been configured:

- CPU_LOAD < 4




Figure 7. Monitoring CPU load statistics

- MEMORY_AVAILABLE > 20MB
- DISK_QUOTA > 1GB

Based on this limits, during the 2 months period were the gCube messaging system was exploited in the infrastructure, the following alarms were raised:

Table 7. Monitoring notification details

CPU_LOAD	430
MEMORY_AVAILABLE	1597
DISK_QUOTA	978
OTHER (e.g., ghn restart, scopes added, etc)	1482
TOTAL	4487

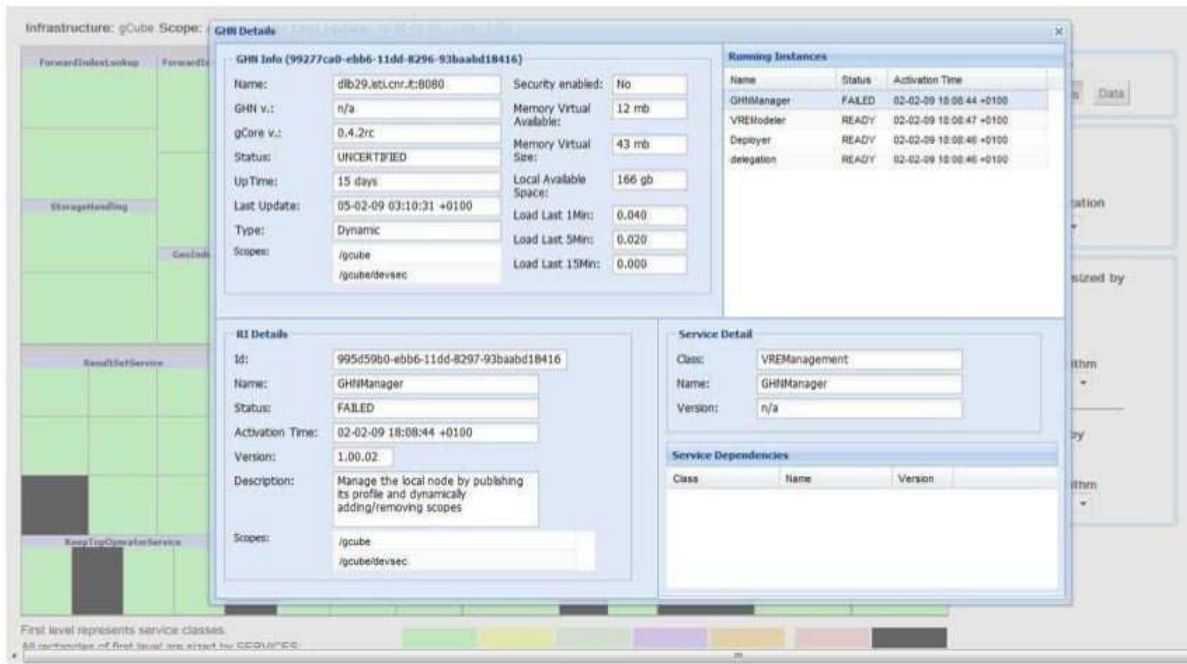


Figure 5. Infrastructure viewer details interface




Figure 8. Monitoring disk quota statistics




Figure 9. Monitoring memory available statistics

Particularly interesting is the information about gHN restarts in order to understand site reliability. Figure 10 shows for each site the average number of gHN restarts.

Portal Accounting Data The information about portal exploitation is available since gCube 1.4 was deployed in production. Since then the D4science portal is able to trace specific operations performed by users logged in.

A summary of the operations/accesses performed by users on the D4Science portal is presented on Table 8.

One particular action under portal accounting is the number of logins to the D4science portal. Portal logins can belong to two categories:

- Login to VOs to perform administration tasks;
- Login to VREs to perform end-user operations (Fig. 11);

Figure 12 shows the number of portal accesses per production VREs. Those labelled with (*) are temporary VREs

created and dismissed after few hours.

Another parameter to understand the use of the D4science portal is the number of operation performed by users. Figure 13 shows a number of selected operations performed via the portal.

Node Accounting Data The information about gCube services invocations is available since gCube 1.4 was deployed in production. The accounting node probes deployed on each gHN of the D4science production infrastructure aggregate data coming from node access logs and send them to the message broker.

As an example of the information that can be retrieved by the accounting probe, Table 9 summarizes the FARM IS Services (IS-Registry, IS-Notifier) averages.

4. Conclusions

D4Science currently makes available different VOs and VREs addressing the requirements expressed by the two different

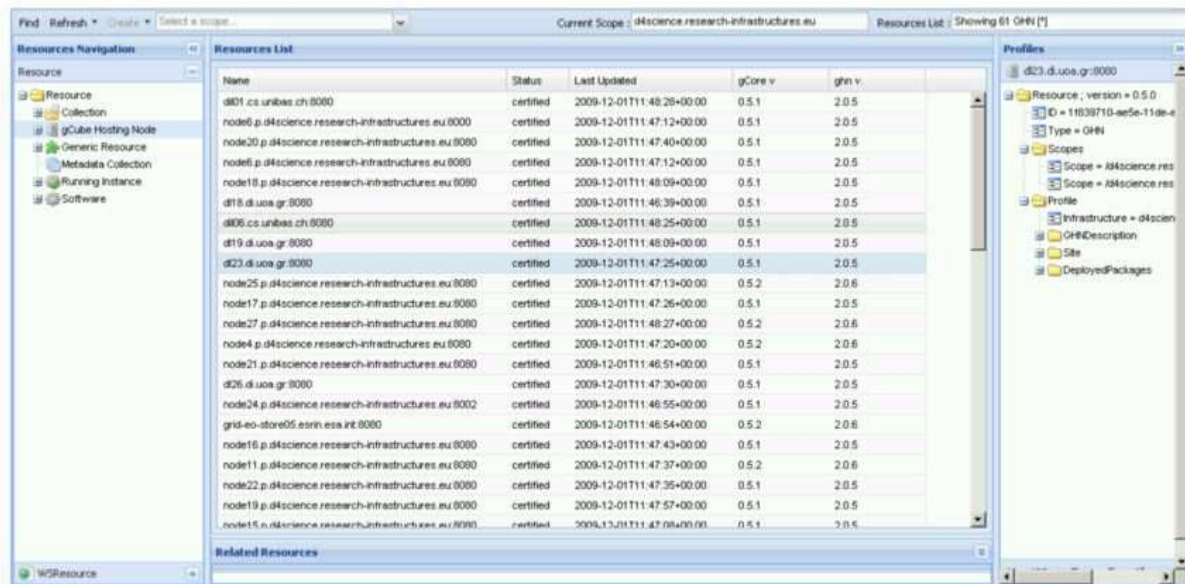


Figure 6. gCube resource management

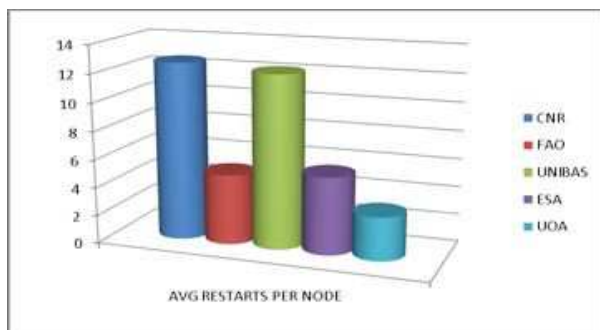


Figure 10. Monitoring notifications (average gHN restarts)

user communities: one VO for the EM community containing 2 VREs (GCM and GVM), and one VO for FARM community containing 3 VREs (FCPPS, ICIS, and AquaMaps).

From the deployment point of view, five sites contribute to the infrastructure: CNR, ESA, FAO, NKUA, and UNIBASEL. These sites provided a total of 53 physical machines, offering 408 GB RAM, 38 TB disk space, and 166 processor cores. Table 10 summarizes some facts from each site contribution.

From the data point of view, the infrastructure operates (i) 42 content collections containing more than 28,800 compound Information Objects, (ii) 97 metadata collections with multiple metadata schema including Dublin Core, ISO19115 and proprietary formats containing more than 54,300 metadata records, and (iii) more than 380 WS-Resources representing full text, forward and geo indices supporting the search over these data.

To support the infrastructure monitoring and accounting activities different tools are available. In particular, the gCube

Table 8. Portal accounting statistics

Operation	Number of access
Login	1415 (to VREs 687)
Advanced Search	65
Simple Search	141
Quick Search	37
Google Search	8
Browse	167
Content Retrieval	40

Table 9. Node accounting statistics (FARM IS)

Service	Hourly Invocation (Avg)	Invocation Time (Avg)
IS-Registry	1628	0.062
IS-Notifier	56	0.010

Messaging System provides a powerful distributed tool to understand the status of the infrastructure and account its exploitation at service and end-user level. The available tools allow to quickly discover problems in sites, possible defects in gCube, and unexpected exploitation scenarios.

Acknowledgments

The work reported has been partially supported by the D4Science project (FP7 of the European Commission, INFRA-2007-1.2.2, Contract No. 212488).




Figure 11. Portal accounting statistics (total logins)




Figure 13. Portal accounting statistics (portal actions)

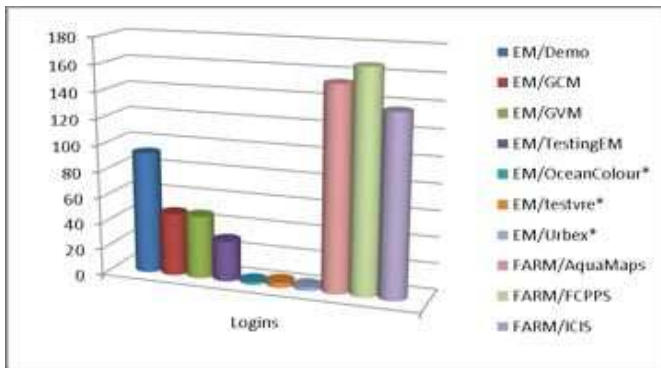


Figure 12. Portal accounting statistics (logins per VRE)

Table 10. Sites contribution overview

Site	HW	gHN	gLite WNs	Downtime
CNR	25	38	10	15
ESA	13	2	39	9
FAO	1	2	-	-
NKUA	8	12	10	25
UNIBASEL	6	6	-	2
Total	53	60	59	51

References

- [1] gCube Development Team. gCube Website. <https://www.gcube-system.org>, 2008.
- [2] Leonardo Candela, Donatella Castelli, and Pasquale Pagano. gCube v1.0: A Software System for Hybrid Data Infrastructures. Technical Report 2008-TR-035, Istituto di Scienza e Tecnologie dell’Informazione “A. Faedo”, CNR, 2008.
- [3] L. Candela, D. Castelli, and P. Pagano. gCube: A Service-Oriented Application Framework on the Grid. *ERCIM News*, (72):48–49, January 2008.
- [4] I. Foster and C. Kesselman. *The Grid: Blueprint for a Future Computing Infrastructure*. Morgan-Kaufmann, 2004.
- [5] L. Candela, D. Castelli, N. Ferro, Y. Ioannidis, G. Koutrika, C. Meghini, P. Pagano, S. Ross, D. Soergel, M. Agosti, M. Dobрева, V. Katifori, and H. Schuldt. *The DELOS Digital Library Reference Model - Foundations for Digital Libraries*. DELOS: a Network of Excellence on Digital Libraries, February 2008. ISSN 1818-8044 ISBN 2-912335-37-X.
- [6] M. N. Huhns and M. P. Singh. Service-oriented computing: key concepts and principles. *IEEE Internet Computing*, 9:75–81, 2005.
- [7] Leonardo Candela, Donatella Castelli, and Pasquale Pagano. On-demand Virtual Research Environments and the Changing Roles of Librarians. *Library Hi Tech*, 27(2):239–251, 2009.
- [8] I. Foster, C. Kesselman, and S. Tuecke. The Anatomy of the Grid: Enabling Scalable Virtual Organization. *The International Journal of High Performance Computing Applications*, 15(3):200–222, 2001.
- [9] Leonardo Candela, Donatella Castelli, and Pasquale Pagano. D4Science: an e-Infrastructure for Supporting Virtual Research Environments. In Maristella Agosti, Floriana Esposito, and Costantino Thanos, editors, *Post-proceedings of the 5th Italian Research Conference on Digital Libraries - IRCDL 2009*, pages 166–169. DELOS: an Association for Digital Libraries, 2009.
- [10] E. Laure, S. M. Fisher, A. Frohner, C. Grandi, P. Kunszt, A. Krenek, O. Mulmo, F. Pacini, F. Prelz, J. White, M. Barroso, P. Buncic, F. Hemmer, A. Di Meglio, and A. Edlund. Programming the Grid with gLite. *Computational Methods in Science and Technology*, 12(1):33–45, 2006.
- [11] Owen Appleton, Bob Jones, Dieter Kranzmueller, and Erwin Laure. The EGEE-II project: Evolution towards a permanent european grid initiative. In Lucio Grandinetti, editor, *High Performance Computing (HPC) and Grids in Action*, volume 16 of *Advances in*

Parallel Computing, pages 424–435. IOS Press, Amsterdam, March 2008.

- [12] gCube Development Team. Administrator's Guide. https://technical.wiki.d4science.research-infrastructures.eu/documentation/index.php/Administrator%27s_Guide, 2008.
- [13] L. Candela, F. Akal, H. Avancini, D. Castelli, L. Fusco, V. Guidetti, C. Langguth, A. Manzi, P. Pagano, H. Schuldt, M. Simi, M. Springmann, and L. Voicu. DILIGENT: integrating Digital Library and Grid Technologies for a new Earth Observation Research Infrastructure. *International Journal on Digital Libraries*, 7(1-2):59–80, October 2007.
- [14] D. Castelli, L. Candela, P. Pagano, and M. Simi. DILIGENT: A DL Infrastructure for Supporting Joint Research. In IEEE Computer Society, editor, *2nd IEEE-CS International Symposium Global Data Interoperability - Challenges and Technologies*, pages 56–69, 2005.
- [15] L. Candela, D. Castelli, and P. Pagano. A Reference Architecture for Digital Library Systems: Principles and Applications. In C. Thanos, F. Borri, and L. Candela, editors, *Digital Libraries: Research and Development, First International DELOS Conference, Pisa, Italy, February 13-14, 2007, Revised Selected Papers*, volume 4877 of *Lecture Notes in Computer Science*, pages 22–35. Springer, 2007.