Project Acronym	***iMarine***
Project Title	***Data e-Infrastructure Initiative for Fisheries Management and Conservation of Marine Living Resources***
Project Number	***283644***
Deliverable Title	***iMarine Data Infrastructure Enabling Software***
Deliverable No.	***D8.2***
Delivery Date	***08 2012***
Author	***Manuele Simi – CNR***

Abstract:
iMarine Data Infrastructure Enabling Software contains the description of software and pointers to the documentation and artifacts of the related components that comprise the e-Infrastructure Management suite delivered from M7 to M10

iMarine (RI – 283644) is a Research Infrastructures Combination of Collaborative Project and Coordination and Support Action (CP-CSA) co-funded by the European Commission under the Capacities Programme, Framework Programme Seven (FP7).

DOCUMENT INFORMATION

PROJECT	
Project Acronym	iMarine
Project Title	Data e-Infrastructure Initiative for Fisheries Management and Conservation of Marine Living Resources
Project Start	1st November 2011
Project Duration	30 months
Funding	FP7-INFRASTRUCTURES-2011-2
Grant Agreement No.	283644
DOCUMENT	
Deliverable No.	D8.1
Deliverable Title	iMarine Data Infrastructure Enabling Software
Contractual Delivery Date	31/08/2012
Actual Delivery Date	31/08/2012
Author(s)	Manuele Simi – CNR
Editor(s)	Manuele Simi – CNR
Reviewer(s)	Andrea Manzi – CERN
Contributor(s)	Gerasimos Farantatos – NKUA , Ciro Formisano E-IIS
Work Package No.	WP 8
Work Package Title	iMarine Data e-Infrastructure Enabling-technology Development
Work Package Leader	Manuele Simi – CNR
Work Package Participants	CNR, NKUA, FAO, E-IIS
Estimated Person Months	9.8
Distribution	Public
Nature	Other
Version / Revision	1.0
Draft / Final	Final
Total No. Pages (including cover)	12
Keywords	Enabling Software, Workflow, Process Execution Engine, Security, gCube, Maven

CHANGE LOG

Reason for Change	Issue	Version	Date	Partner / Responsible
Version ready for Official Review		0.1	25/08/2012	CNR
Version reviewed		0.2	27/08/2012	CERN
Version addressing Reviewer Comments		0.3	28/08/2012	CNR
Final Version		1.0	29/08/2012	CERN

DISCLAIMER

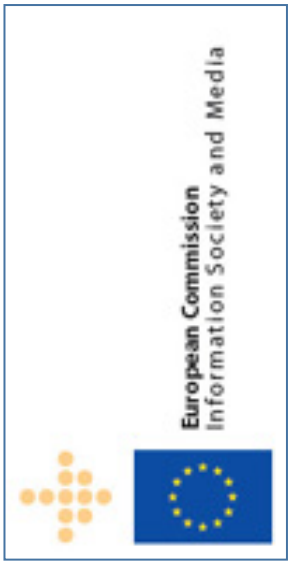

iMarine (RI – 283644) is a Research Infrastructures Combination of Collaborative Project and Coordination and Support Action (CP-CSA) co-funded by the European Commission under the Capacities Programme, Framework Programme Seven (FP7).

The goal of iMarine*, Data e-Infrastructure Initiative for Fisheries Management and Conservation of Marine Living Resources*, is to establish and operate a data infrastructure supporting the principles of the Ecosystem Approach to Fisheries Management and Conservation of Marine Living Resources and to facilitate the emergence of a unified Ecosystem Approach Community of Practice (EA-CoP).

This document contains information on iMarine core activities, findings and outcomes and it may also contain contributions from distinguished experts who contribute as iMarine Board members. Any reference to content in this document should clearly indicate the authors, source, organisation and publication date.

The document has been produced with the funding of the European Commission. The content of this publication is the sole responsibility of the iMarine Consortium and its experts, and it cannot be considered to reflect the views of the European Commission. The authors of this document have taken any available measure in order for its content to be accurate, consistent and lawful. However, neither the project consortium as a whole nor the individual partners that implicitly or explicitly participated the creation and publication of this document hold any sort of responsibility that might occur as a result of using its content.

The European Union (EU) was established in accordance with the Treaty on the European Union (Maastricht). There are currently 27 member states of the European Union. It is based on the European Communities and the member states' cooperation in the fields of Common Foreign and Security Policy and Justice and Home Affairs. The five main institutions of the European Union are the European Parliament, the Council of Ministers, the European Commission, the Court of Justice, and the Court of Auditors (http://europa.eu.int/).

GLOSSARY

ABBREVIATION	DEFINITION
iMarine	Data e-Infrastructure Initiative for Fisheries Management and Conservation of Marine Living Resources
VO, Virtual Organization	A dynamic set of individuals or institutions defined around a set of resource-sharing rules and conditions. All these virtual organizations share some commonality among them, including common concerns and requirements, but may vary in size, scope, duration, sociology, and structure.
VRE, Virtual Research Environment	A *system* with the following distinguishing features: *(i)* it is a Web-based working environment; *(ii)* it is tailored to serve the needs of a Community of Practice; *(iii)* it is expected to provide a community of practice with the whole array of commodities needed to accomplish the community's goal(s); *(iv)* it is open and flexible with respect to the overall service offering and lifetime; and *(v)* it promotes fine-grained controlled sharing of both intermediate and final research results by guaranteeing ownership, provenance, and attribution.
Scope	Runtime information about the *Infrastructure*, the *Virtual Organization* (VO), or the *Virtual Research Environment* (VRE) in which resources of the infrastructure are operating in a specific moment of their lifetime.
XACML, *eXtensible Access Control Markup Language*	A policy-based language standardized by OASIS to make the use of security for software easy to configure, extend, and change with minimal impact.
gRS2, gCube Result Set 2	A framework, part of the gCube system, enabling point to point producer/consumer communication

DELIVERABLE SUMMARY

1. INTRODUCTION

This deliverable describes the novelties and evolution of the iMarine Data Infrastructure Enabling Software from M7 (May ’12) to M10 (Aug.'12).

2. TARGET RELEASE(S)

This deliverable reports on the software released as part of the following gCube releases:

- gCube 2.9.1
- gCube 2.10.0

3. OBJECTIVES

The new version of components belonging to the Data Infrastructure Enabling Software released as part of the target releases covers the following objectives:

- Simplification and optimization of VRE creation procedure

  A large part of the resource management services was redesigned and re-implemented in order to improve in many ways the VRE creation process. Firstly, the Resource manager service had become a multi-scoped service, meaning that it is now able to create and manage multiple scopes within a single instance of the service. Previously, one instance and one node dedicated to running it were needed for each scope. Thus novelty leads to a dramatic reduction of the resources requested for setting up new VREs and VOs. Still on the Resource Manager service, new abilities have been added to *(i)* deploy/undeploy single gCube software packages, web applications, and plugins, *(ii)* reduce the internal state at minimum to overcome future backward compatibility issues (the lesser is serialized, the better for avoiding outdating information) and *(iii)* prepare the ground for package upgrade functionalities to be next implemented. Finally, the Resource Manager features a new public interface, composed by 4 portTypes, allowing a better abstraction and minimum connections among clients and the service (loose coupling).
  The VREModeler service has been modified to comply with the new Resource Manager interface and, most importantly, to offer a different and simplified approach for creating VREs driven by functionality instead of services. The goal of this newly implemented approach is to lower as much as possible the understanding of the VRE implementation to VRE creators: they now simply have to ask for desired features instead of services.

- Enhance the resource handling

  Developments in the Resource Manager service offered the opportunity to introduce several enhancements in resource handling.

In details, the new version of the service tolerates node re-deployment, meaning that if a node under one of its managed scopes is re-deployed with a new state, Resource Manager can rebuilt its internal state for managing it and its hosted services' instances. Another improvement has been implemented in the new interface of the service, by adding the capability of requesting the deployment or undeployment of multiple services' instances at the same time, while, previously, different calls to the service were needed for the case. Lastly, as part of the re-design of the state of the service discussed in the previous objective, Resource Manager has set the ground for autonomic software upgrade (to be implemented in one of the next releases).

- Integration with Maven 3

  The release of maven-dependency-plugins 2.5 (beginning of August) has forced a major revision of the approach to software dependency resolution taken in the Software Gateway service. In fact, this plugin was interfaced by third-party external software previously exploited, but now no longer supported by its creators. As a consequence, the communication layer between Software Gateway and Maven (our dependency management system) was broken and had to be re-implemented. The internals for embedding Maven have changed a lot since the previous layer was developed (years ago, it was inherited by D4Science I and II projects [9]) and Maven 3 libraries are much easier to embed in a Java application than before. Thus, a new layer was designed and developed to directly interface Software Gateway with Maven. Lesser dependencies and software layers are now involved in the communication and a better abstraction over Maven has been reached, both allowing maintaining compatibility with its future evolution.

- Transparent scope propagation

  A major objective of WP8 is the transparent management of software [6]. The new common-scope library represents a step forward in this direction by delivering new scope and service map handling facilities. Common-scope allows reducing requirements against both gCube clients and services by propagating operational runtime information using a behind-the-scene thread-based mechanism. The library has been integrated in the gCore Framework in substitution of the previous intrusive mechanism.

- Improve the matchmaking and resource management capabilities of the gCube process execution engine (PE2ng)

  In the direction of supporting the evolution of workflow management components, as well as all PE2ng adaptors, and also in order to satisfy requirements of the Data Transformation PE2ng adaptor implemented in the context of WP10, the features offered by the node selection library already developed in the context of WP8 were improved in a number of ways. The improvements range from reorganization and naming changes to the introduction of new capabilities for all node selectors and policies. Specifically, the previous notion of node colocation is now referred to as node assignment, with the two main colocation policies being treated as special cases of the latter. The hierarchy of node selectors and assignment policies was improved in order to provide similar but specialized behaviour deriving from a common notion, an example being the *LRU* and *MRU* node selectors. In terms of functionality, there were two major improvements:

- the notion of tie-breaking was introduced to all node selectors. In particular, all node selectors now support tie breaking in the form of an additional node selector. Tie-breaking selectors are used if some of the scores produced by the node assessment procedure are equal. In that case, the subset of the candidate nodes involved in the ties is reassessed by the tie-breaking selector and the new scores are interpolated into the final scores that are produced. As tiebreakers are standard node selectors tie breaking can continue in an arbitrary number of levels, if desired.
- the node distance calculation support was added to the library. Exploiting an XML document describing node topology can perform this. Heuristics on host names and IP addresses are used in the absence of such information. Node distance was also included in the possible factors of the cost function used by the respective node selection policy.

Furthermore, the hosting node information exploited by the library was enriched in order to support the functionality of the provided selectors and policies and to prepare the ground for the upcoming integration of PE2ng with Resource Registry. Among such enhancements is the ability to distinguish the local node from the remote ones. The information of whether the entity is the local node or not is provided by the information provider harvesting hosting node statistics, through the corresponding entity adaptor.

- Improve the distributed execution capabilities of the search workflows PE2ng adaptor

The adaptor was improved by the integration with the latest version of the node selection library and by the introduction of a plan analyser and a rewriter. The analyser performs a simple analysis in order to determine if the search workflow is simple or complex based on a configurable cost threshold. Simple workflows, especially those which do not entail memory consumption in the form of result set buffers (e.g. search involving only one data source) can be executed locally at the node which hosts the adaptor and the search system endpoint, in order to minimize overheads and achieve low response times for commonly executed simple queries. The rewriter is responsible for breaking complex operations of high cost, such as the merging of the results of a large number of data sources into subtrees of equivalent simpler operations. In this way, the adaptor has greater flexibility by managing a higher number of low cost operations. In case of complex plans, node assignment is performed using configurable node selectors and assignment policies to control the number of nodes to which the execution of the workflow is spread and to perform load balancing. Specifically, the selection of data sources takes place using a node selector, and the rest of the plan is distributed using a node assignment and selector pair. The behaviour of node selectors in either case is further controlled by the definition of tie breaking selectors. Regardless of the policies used, the adaptor operates based on a maximum value for the cost of workflow subtrees which can be executed in a single node (*maximum colocation cost)* and a score threshold under which nodes are excluded for selection. Furthermore, if the resources of the infrastructure are not enough to achieve optimal assignment and the final assignment results in nodes being over-utilized, the invocations of search operators are instructed to limit the buffer capacity used by gRS2, trading off some fraction of the performance for robustness and stability. This behaviour means that, as long as assignment is successful in finding at least one node for execution, requests are never rejected. Instead, workflow executions are being "squeezed" into the infrastructure in such a way that all requests are handled with some impact in performances. This behaviour is planned to be complemented in the near future, with the introduction of the gRS2 total utilized buffer capacity

per node monitoring, allowing executions to be rejected whenever memory consumption rises to critical levels.

- Improve the flexibility of the search workflows PE2ng adaptor .

  The flexibility of the search adaptor was improved by modularizing the data source wrappers which implement the logic of web service invocations. Wrappers for the gCube environment are intended to remain within the software artifact of the adaptor for distribution purposes, acting as the default implementations.

- Enhance the robustness of PE2ng under high workloads

  A queuing facility was implemented in order to handle high workloads. The implementation of this facility consists of two main aspects:
  - the monitoring of execution nodes using a publish-subscribe mechanism. The Execution Engine exploits this functionality to obtain an up-to-date view of the infrastructure in terms of running tasks. The information is updated upon the initiation and completion of execution tasks at the hosting node level,
  - the implementation of a queue-based scheduling component which acts as a front-end to the Execution Engine. This facility also exploits the information describing the state of the infrastructure, deciding to forward or queue up incoming tasks according to node utilization.

  In order to support monitoring, the notion of notification handling was introduced to PE2ng, in the form of a software abstraction. Monitoring is performed using a JMS implementation for notification handling, which operated using Apache ActiveMQ. The queuing and the monitoring functionalities implementation mandated the release of three new software components.

- Fix observed issues on PE2ng

  When transmitting an execution plan for remote execution, the Execution Engine avoids sending all variables present in the plan, in order to save time and network resources. Instead, each element is expected to declare which variables should be included but a number of parameter filters did not declare all variables of importance. These omissions were corrected.

  Furthermore, array evaluation was not performed correctly in the case of remote execution. This was also fixed and now the analysis is performed correctly in all cases.

- Authentication and User Management

  Since iMarine infrastructure has been designed to scale and to be able to serve an ever-increasing number of scientific communities, the progressive introduction of a scalable and Service Oriented User Management System is a major objective. LDAP is a fast protocol to manage user identities and authentication. Several implementations are available, open source and proprietary, with different level of scalability and efficiency. By adopting the most appropriate implementation it is possible to choose the right trade-off between scalability and costs. Authorization and User Management System compliant with LDAP have a great impact on performance and scalability of gCube security system. A LDAP User Management System has been designed and developed in order to provide Java API to perform Create/Read/Update/Delete operations on identities stored in an LDAP server. Additionally, as released in gCube 2.10, the new Authentication Service, part of SOA3 system [7], provides an equivalent REST interface to authenticate users against identities stored in an LDAP directory.

- Authorization: management of policies compliant with XACML based on multiple attributes

  The security infrastructure of iMarine now supports policies based on multiple attributes, which may be any arbitrary common characteristic. This feature will improve dramatically the flexibility in the definition of policies. The new Policy Management interfaces provide a set of interfaces to manage authorization policies written in XACML language by exposing a set of Java API simplifying the definition of policies abstracting on the language itself. These interfaces support all the features above, in order to define complex and complete policies. The interfaces interact with a XACML based Policy Administration Point and they have been successfully tested with Argus Authorization Framework PAP version 1.5.1[7].

4. COMPONENTS

As part of the target releases, the following components have been updated or newly introduced:

- to simplify and optimize the VRE creation procedure
  - Resource Manager 2.0.0
  - Deployer 2.4.1
  - VREModeler 1.6.1
- to handle transparent scope propagation:
  - gCore Framework 1.4.1
  - common-scope 1.0.0
- to integrate Maven 3:
  - Software Gateway 1.1.0
- to improve the matchmaking and resource management capabilities of PE2ng:
  - MadgikCommons 1.2.0
- to improve the distributed execution capabilities and flexibility of the search workflows PE2ng adaptor:
  - WorkflowSearchAdaptor 1.1.0
- to enhance the robustness of PE2ng under high workloads:
  - ExecutionEngine 1.2.0
  - ExecutionEngineService 1.0.6
  - QueueableExecutionEngine 1.0.0
  - WorkflowEngine 1.3.0
  - WorkflowEngineService 1.3.0
  - InformationSystem 1.4.0
  - EnvironmentProvider 1.3.0
  - GCubeEnvironmentProvider 1.4.0
  - NotificationHandling 1.0.0
  - MadgikEnviromentProvider 1.0.0
  - MadgikNotificationHandling 1.0.0
- to fix observed issues on PE2ng:
  - ExecutionEngine 1.2.0
- to manage authorization policies compliant with XACML language:

    - PolicyManagementInterfaces 1.0.0
- to provide user management based on LDAP:
    - LDAPUserManagementLibrary 1.0.0
    - UserManagementRESTInterfaces 1.0.0
- to provide user authentication:
    - AuthenticationService 1.0.0

5. DOCUMENTATION

A comprehensive overview of the subsystems the described components belong to is available at [6] (e-infrastructure Management Facilities), [7] (Policy-oriented Security Facilities) and [8] (Workflow Management Facilities).

Technical documentation covering all the aspects of the software is available at:

- Admin's Guide [2]
- Developer's Guide [4]
- User's Guide [5]

For development purpose, Javadoc documentation for each component, along with a direct link to the associated section in Developer's Guide, is available at [1].

6. DOWNLOAD

The components described in this deliverable are available for download at [1]. Direct links to each component are available at [2].

REFERENCES

[1] gCube Maven Repository RELEASES:
http://maven.research-infrastructures.eu/nexus/index.html#view-repositories;gcube-releases~browsestorage

[2] gCube Distribution Site:
https://www.gcube-system.org/index.php?option=com_distribution&view=distribution&Itemid=23

[3] Administrator's Guide:
https://gcube.wiki.gcube-system.org/gcube/index.php/Administrator%27s_Guide

[4] Developer's Guide:
https://gcube.wiki.gcube-system.org/gcube/index.php/Developer%27s_Guide

[5] User's Guide:
https://gcube.wiki.gcube-system.org/gcube/index.php/User%27s_Guide

[6] Milestone 33: Data e-Infrastructure Management Facilities:
https://gcube.wiki.gcube-system.org/gcube/index.php/Data_e-Infrastructure_Management_Facilities

[7] Milestone 34: Data e-Infrastructure Policy-oriented Security Facilities:
https://gcube.wiki.gcube-system.org/gcube/index.php/Data_e-Infrastructure_Policy-oriented_Security_Facilities

[8] Milestone 35: Workflow Management Facilities:
https://gcube.wiki.gcube-system.org/gcube/index.php/Workflow_Management_Facilities

[9] D4Science I and II projects Web Site
http://www.d4science.eu/

[10] Apache ActiveMQ Web Site
http://activemq.apache.org/