

**SEVENTH FRAMEWORK PROGRAMME**  
**Capacities specific program - Research Infrastructures**  
**FP7–INFRASTRUCTURES–2010-2**

Project acronym: *CHAIN*

Project full title: Co-ordination & Harmonisation of Advanced e-Infrastructures

Grant agreement no.: 260011

Start date of project: 1 December 2010

Duration: 24 months

**ROC Cookbook**

Version 2.0

Submitted on: 29/11/2012

Organisation name of lead beneficiary for this document: GRNET

Dissemination Level: Public

Abstract:

This document provides a concise guide on what is needed to setup a Regional Operations Centre / operational coordination body to be compatible with the European Grid Infrastructure and a snapshot of the current trends in new technologies used to provide a pan-european grid infrastructure for eScience.

---

Copyrights © The CHAIN Consortium 2012.

More details on the copyright holders can be found at the following URL: <http://www.chain-project.eu>

CHAIN (“Co-ordination & Harmonisation of Advanced e-Infrastructures”) is a project co-funded by the European Union in the framework of the 7<sup>th</sup> FP for Research and Technological Development, as part of the “Capacities specific program – Research Infrastructures FP7–INFRASTRUCTURES–2010-2”. For more information on the project, its partners and contributors please see <http://www.chain-project.eu>

You are permitted to copy and distribute verbatim copies of this document containing this copyright notice, but modifying this document is not allowed. You are permitted to copy this document in whole or in part into other documents if you attach the following reference to the copied elements: “Copyright © 2010. The CHAIN Consortium. <http://www.chain-project.eu>”.

The information contained in this document represents the views of the CHAIN Consortium as of the date they are published. The CHAIN Consortium does not guarantee that any information contained herein is error-free, or up to date.

THE CHAIN CONSORTIUM MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, BY PUBLISHING THIS DOCUMENT.

---

### Document log

Issue	Date	Comment	Author
V0.1	01/05/2012	Table of Contents circulated for comments	K. Koumantaros
V0.2	15/05/2012	Feedback received	O. Prnjat
V0.3	1/06/2012	TOC updated	K. Koumantaros
V0.4	18/06/2012	First contribution included (ROC use case)	B. Becker
V0.5	20/06/2012	Additions to chapters 1 and 3	K. Koumantaros
V0.6	22/09/12	Additions to chapter 4	P. Korosoglou
V0.7	24/10/2012	Contribution for Chapter 1	K. Koumantaros
V0.8	26/10/2012	Contributions on IPv6	M. Reale
V1.0	1/11/2012	New contributions and corrections	K. Koumantaros
V1.9	10/11/2012	Consolidated version, details on future technologies	K. Koumantaros, C. Kanellopoulos, P. Gerakios, O. Prnjat
V2.0	28/11/2012	Technical inputs, editing and formatting	K. Koumantaros, C. Kanellopoulos, I. Psychogiou, K. Lykou

# Table of contents

<b>Table of contents .....</b>	<b>3</b>
<b>1 ROC Organisation .....</b>	<b>4</b>
1.1 Roles and Responsibilities.....	5
1.1.1 Operations Manager / Deputy .....	5
1.1.2 Optional Regional/Country Representative .....	5
1.1.3 Security Officer .....	5
1.1.4 NGI International Liaison Officer .....	5
1.1.5 Support on Duty.....	5
1.1.6 Site Administrators .....	6
1.1.7 User Support and Application Support Teams.....	6
1.1.8 Coordination and policies.....	6
1.2 Manpower Estimates for ROC/NGI functions .....	7
1.3 Minimum Grid Site Infrastructure Required .....	8
1.4 ROC/NGI Support Tools.....	11
1.4.1 Monitoring Tools Nagios .....	11
1.4.2 Certification Authority.....	12
1.4.3 NGI/ROC Website and local documentation .....	12
1.4.4 Operations Database .....	12
1.4.5 Monitoring Tools SAM/Nagios.....	12
1.4.6 Helpdesk.....	12
<b>2 Current ROC status in CHAIN.....</b>	<b>13</b>
<b>3 Use Case: Africa and Arabia ROC.....</b>	<b>17</b>
3.1 ROC Composition and Organisation .....	17
3.2 Services Provided by the ROC .....	17
3.3 Standard Operating Procedures .....	18
<b>4 Emerging Technologies relevant for Grid operations.....</b>	<b>20</b>
4.1 IPv6 .....	20
4.1.1 Why IPv6 and Grid .....	20
4.1.2 Current Status in Europe .....	21
4.2 Cloud Infrastructures.....	21
4.2.1 GRNET ~Okeanos IaaS use case .....	22

## 1 ROC Organisation

A National Grid Initiative – NGI - is the national federation of shared computing, storage and data resources that delivers sustainable, integrated and secure distributed computing services to the national research communities and their international collaborators. An NGI is usually represented by a legal body leading the federation and acts as a resource provider (RP) to the European Grid Infrastructure (EGI). An operations centre (OC) is the team that manages the services offered on behalf of the NGI. Depending on the size and nature of the NGI it can either be centralised or distributed. Similarly a Regional Operations Centre (ROC) is an OC responsible for the services offered by the cooperation of national RPs of a region (for example, Africa and Arabia ROC). In the following paragraphs we define the minimum roles needed to run an OC on behalf of a ROC or NGI that is compatible with EGI. The main difference between an OC and a ROC is an extra administration level with country representatives that act as Local Operation Centre Managers for the country on behalf of the ROC Operation Centre Manager.

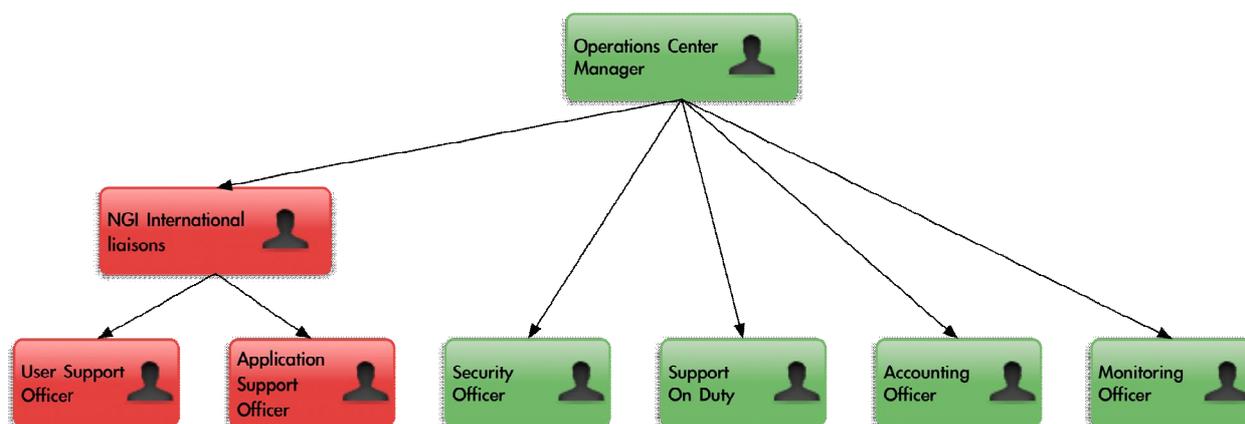


Figure 1 - Organisational Chart of an Operations Centre

## 1.1 Roles and Responsibilities

### 1.1.1 Operations Manager / Deputy

The operations manager is the head of the team responsible for the overall operation of the OC. He or she is responsible for the following actions.

- Overseeing day-to-day operations performed by the Site Administrators.
- Negotiating with sites their resource sharing policy according to the needs of the User Communities the NGI/ROC serves.
- Representing the NGI/ROC to the Operations Management Board<sup>1</sup> and/or Similar management forums.
- Catering for the needs / special request of the user communities the NGI/ROC supports.
- Coordinating the user and application support efforts of the NGI/ROC.
- Coordinating the documentation efforts of the NGI/ROC

### 1.1.2 Optional Regional/Country Representative

This is a similar role to the Operations Manager role that is only applicable to ROC or widely distributed NGIs. The tasks assigned to this role are the same with the ones assigned to an Operations manager but he or she is responsible for a specific country / region that is part of the ROC.

### 1.1.3 Security Officer

Each OC has to appoint an OC Security Officer who will be responsible for the operational security of the OC, including Grid security monitoring, Security training and dissemination, and Incident Response coordination. The OC Security Officer represents the OC in the EGI CSIRT, which covers all aspects of operational security aimed at achieving a secure infrastructure within EGI.

### 1.1.4 NGI International Liaison Officer

The interaction between the NGI teams and EGI.eu on non-operational activities is undertaken through an "NGI International Liaisons"(NILs in short) officer, who are responsible within the NGIs for the delivery and interaction of non-operational tasks such as:

- Marketing & Communication; Strategic Planning and Policy Support;
- Community outreach and events for new users;
- Technical outreach and support to new communities.

### 1.1.5 Support on Duty

Even though each site is responsible for its day to day operations, a support on duty role is usually established in order supervise the day-to-day operations and pursue the sites to solve any issues that may arise. It is common practice that this role is rotated between sites.

<sup>1</sup> OMB: Operations Management Board <https://wiki.egi.eu/wiki/OMB>

### 1.1.6 Site Administrators

Site administrators are the actual system administrators that deal with the maintenance and upgrades of a site and its services. Each site should ideally have at least two site administrators responsible for its day-to-day operations. Larger sites that offer more complex services may need more site administrators to handle the load depending on their complexity.

### 1.1.7 User Support and Application Support Teams

The User support team is responsible of the 1<sup>st</sup> level of support for the users. The team usually guides a new user how to get access and get the most, out of the infrastructure and services offered by the NGI. The application support team guides the users to port or optimise their applications for the infrastructure offered. Both teams usually are usually comprised by site administrators in order to take advantage of the know-how they have about the infrastructure.

### 1.1.8 Coordination and policies

In order to achieve an efficient level of operation a biweekly or more frequent audio conference is usually scheduled in order to discuss current issues with the operation of the NGI. Participants on these calls are usually all of the officers mentioned above and at least one representative of each site. An extensive list of best practices and policies that cover all aspects to run an OC is available at [https://wiki.egi.eu/wiki/Category:Operations\\_Procedures](https://wiki.egi.eu/wiki/Category:Operations_Procedures), a new OC should be made familiar with all them and adapt or augment most of them.

## 1.2 Manpower Estimates for ROC/NGI functions

The average manpower needed for each role is as follows, based on accumulated European NGI and EGI experience.

- Operations Manager / Deputy: 0.5 to 1 FTE each, depending on the size of the NGI / ROC.
- Regional/Country Representative (optional): 0.3 to 0.5 FTE depending on the number of sites he or she supports.
- Security Officer 0.3 to 0.5 FTE: depending on the size of the NGI / ROC.
- NGI International Liaison officer: 0.3 to 0.5 FTE, depending on the site of the NGI/ROC. In small NGIs this role is frequently assigned to the NGIs Operation Manager or deputy.
- Site Administrator: 0.3 to 0.5 FTE depending on the experience and expertise.
- Support on Duty: this role is usually delegated to site administrators in rotation thus no extra effort is required. If a dedicated team handles this role then 0.3 to 0.5 FTE is required depending on the size of the NGI/ROC
- User Support: 0.5 to 1 FTE, depending on the size of the user community served.
- Application Support: 0.5 to 1 FTE, depending on the size of the user community served.

### 1.3 Minimum Grid Site Infrastructure Required

In order to set up a Grid site, it is important to provide sufficient computational and storage resources along with network connectivity. A site has to provide at least one service that provides the “Information Discovery” capabilities and that publishes site information on the Grid, effectively making the site visible to the rest of the Grid. Of course the purpose of a Grid site is to provide at least one functional capability. The list of functional capabilities can be found in the table below.

#	Capability	Description
1	File transfer	Files are stored at different physical locations within the production infrastructure and are frequently used at other locations. It is necessary for the files to be efficiently transferred over the international wide area networks linking the different resource centers. Typically, a dedicated service provides the File Transfer capability, offered through potentially many instances of the same service software for scalability reasons.  Example services that provide the “File Transfer” capability are: dCache, StoRM, DPM
2	Storage management	Storage Management refers to the ability of managing a storage resource, from simple hard disk-based systems to complex hierarchical systems. Typical deployments are to bundle a management interface with the respective local storage element into one service, so that it provides both the File Access, and the Storage Management capability. Alternatively, a dedicated service implementing Storage Management is capable of managing many remote storage nodes.
3	Data access	Data Access becomes increasingly important in contemporary distributed computing infrastructures. Fine-grained access control to distributed data sets is required to protect copyrighted material or data covered by non-open access licenses from unauthorized access. The same fine-grained access control is necessary to protect searches and indexing activities for data catalogues.

#	Capability	Description
4	Metadata catalogue	<p>The metadata catalogue is used to store and query information relating to the data (files, databases, etc.) stored within the production infrastructure. An integral part of this functionality is not only to query about the existence of a file that may satisfy the needs of the enquiring user, but also the ability to resolve to a concrete description of the location of the file itself.</p> <p>Typically, metadata catalogues are provided by dedicated service implementations and deployments for scalability reasons, and to provide metadata catalogues that feature different metadata sets for different user communities.</p> <p>Example services that provide the “Metadata catalogue” capability are: LFC, Globus, RLS, AMGA</p>
5	Compute	<p>If the Compute capability is provided, then also the Storage Management capability must be provided either from the local site or from a remote site.</p> <p>Example services that provide the “Compute” capability are: CREAM, Globus GRAM5, ARC-CE, EMI-WN</p>
6	Information Discovery	<p>Information discovery is a capability that helps find the required resources that have been registered with it within the production infrastructure. The information collected about such resources is made available through well-known instances that provide the data to some logical collection, infrastructure wide, regional, site, domain, etc.</p> <p>Clients to such service must be able to search, filter, and order the available information until their initial request is satisfied. To enable search and discovery on various levels of the infrastructure it is important to reiterate that any implementation of the Information Discovery Capability must at the same time make use of the Information Model capability defined earlier in this document.</p> <p>Example services that provide the “Information Discovery” Capability are: BDII-site, ARC-InfoSys, Unicore Registry</p>

There are many different patterns of a site, depending on the mission and goals of each site. For example there can be a site that provides only “Storage management” capabilities without providing any computational resources. In most of the cases though, sites provides a combination of

“Compute” capability along with the “Storage management” and the “Information discovery” capabilities.

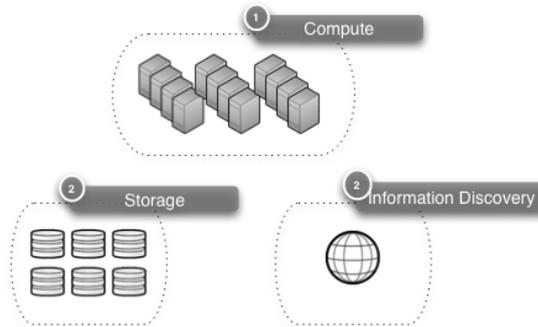


Figure 2 – Grid Site Schematic

In terms of hardware, a site can run on rather modest resources and grow as needs grow. For minimal site following the typical pattern, a relatively modern server can provide that job scheduling mechanism of the “Compute” capability along with the “Information Discovery” capability. One multicore server can operate as the Worker Node of the “Compute Capability” and another with 2TB local storage can provide the “Storage Management” capability. As the needs of the site grow, the 1 Worker Node can grow into a farm of Worker Nodes and the Storage Server could expand into multiple servers with storage arrays attached on them.

## 1.4 ROC/NGI Support Tools

Maintaining and administering a production infrastructure poses a number of requirements on the deployed components. Monitoring the production infrastructure is a key capability that is necessary to determine in real time the current state of the infrastructure. Resources that are made available for usage are either known or must be discovered, and flexible access control protects resources that, for example, are not available for users that are part of a different infrastructure federation. As the state of the production infrastructure resources is inherently dynamic, changes in a resource's state (e.g. availability, or load) are propagated through messaging facilities and endpoints.

Similarly, to account for how much of the provided resources are used, the Accounting Capability requires identical subsequent capabilities to provide to the operations community vital information for budget calculations and, eventually, billing facilities.

### 1.4.1 Monitoring Tools Nagios

Availability, reliability and security are the key parameters for labelling a service as production. In order to ensure the quality of the provided services, all of the resources have to be monitored. Such a monitoring capability is essential for the operational staff attempting to deliver the production infrastructure and the end-users seeking out reliable resources to support their research. It is important in a service-oriented environment to distinguish between the services that are monitored, and the service that provides monitoring.

Within EGI, the SAM framework provides a monitoring infrastructure that is based on three key cornerstones relevant for this Capability: Nagios and related service probes as the functional monitoring data, the MyEGI portal to visualise and report the monitoring results, and a Message Exchange Infrastructure based on Apache MQ that ties the monitoring data extraction tool with the MyEGI data presentation layer.

This architecture defines two interfaces, which reside on the monitoring data extraction layer and the monitoring data presentation/querying layer, respectively. The interface on the extraction layer is rather an integration point in that service-specific probes are provided that plug into the Nagios service-monitoring tool.

The NGI/ROC Monitoring Instance is actually a server node running Nagios, with various sets of tests and probes. The probes are executed from the Monitoring Instance on each site and the results of the probes are published to the Grid Messaging Service. Through the Grid Messaging Service, the results are collected on the NGI/ROC Operations Dashboards and alerts are opened to the sites that have failed tests.

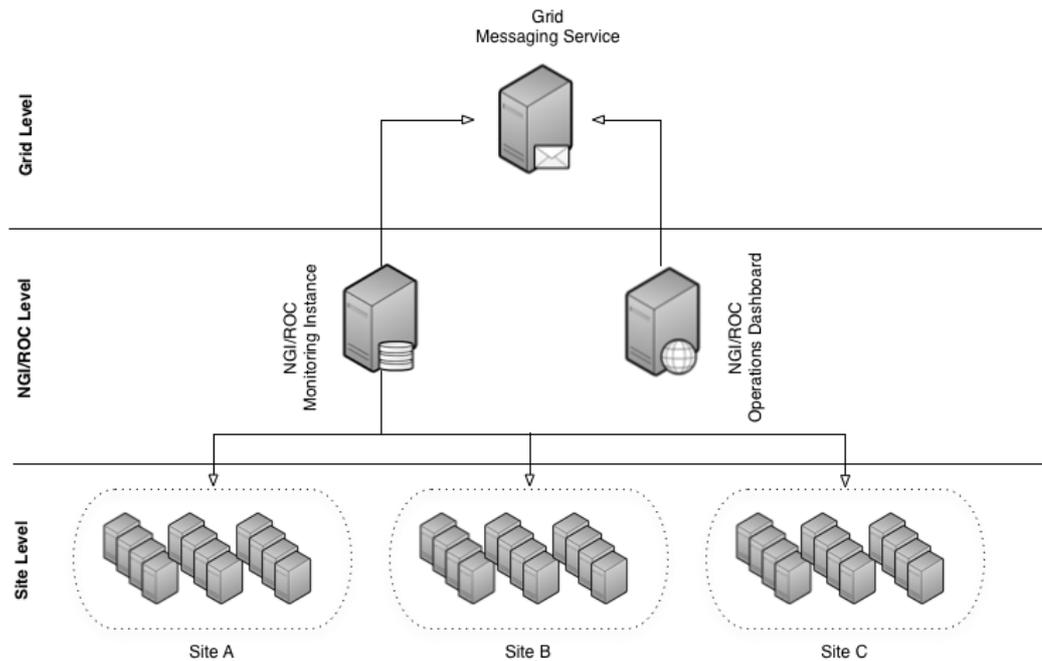


Figure 3 Grid Messaging Service Architecture

## 1.4.2 Certification Authority

A Certification authority (CA) is the service that is used to issue both user and service certificates. CA accreditation is lengthy process thus it is advised to seek assistance from a catch-all CA that is authorised to issue certificates for your region. More information regarding accredited CAs can be found at <http://www.igtfn.net>

## 1.4.3 NGI/ROC Website and local documentation

This is a simple website used to concentrate all necessary local information such as how to get access to the grid, how to seek for support etc.

## 1.4.4 Operations Database

godcb.egi.eu is the central database used to register new NGIs/ROC and their corresponding sites.

## 1.4.5 Monitoring Tools SAM/Nagios

In order to monitor the availability and reliability of your services you will need to set up a national or regional instance of the SAM/Nagios service. More information for this process can be found at <http://repository.egi.eu/category/production/sam/>

## 1.4.6 Helpdesk

An NGI/ROC needs to set-up a user support system (RT<sup>2</sup>, xGUS etc) that is capable to talk with the universal ticket exchange <https://ggus.eu>

<sup>2</sup> Resource Tracker [www.bestpractical.com/rt/](http://www.bestpractical.com/rt/)

## 2 Current ROC status in CHAIN

### AFRICA AND ARABIA ROC

#### Contact Details (Southern Africa)

- Name: Bruce Becker
- E-mail: [bbecker@csir.co.za](mailto:bbecker@csir.co.za)
- Institution: Meraka Institute, Council for Scientific and Industrial Research (CSIR)
- ROC: Africa-ArabiaRegion: Sub-Saharan Africa Region

#### Helpdesk: Which Ticketing system/helpdesk do you use or plan to use

- xGUS
- Other: Non-Integrated with GGUS: Trac issue tracking used for SAGrid planning, application porting and projects.

#### Accounting: Which Accounting system do you use

- DGAS for northern Africa
- Others: Currently None in South Africa. APEL planned.

#### Monitoring: Which Monitoring System do you use

- SAM-NAGIOS based
- GSTAT instances also provided
- Network monitoring via smoke-ping – perfsonar planned.

#### Registration: Which Site Registration System do you use

- GOCGB. There is an NGI\_ZA region in the GOCDB for South African sites which have an MoU with EGI.

#### Discussion:

The ROC is presented via a website ([roc.africa-grid.org](http://roc.africa-grid.org)) which brings together several services. The services are distributed and maintained by members of the previous EUMedGrid Support project and the SAGrid project, and are hosted in South Africa, Jordan, Algeria and Europe.

While the ROC includes the entire Africa-Arabian region, some services are currently configured to cover only some sub-regions or specific sites.

First-level support and ticket process management is undertaken on a shift-basis, by volunteer sites within the ROC region.

## **CHINA ROC**

### **Contact Details**

- Name: Yan Xiaofei or Chen Gang
- E-mail: lcg-admin@ihep.ac.cn
- Institution: Institute of High Energy Physics
- ROC: ROC\_Canada
- Region: ROC\_Canada

**Helpdesk:** Which Ticketing system/helpdesk do you use or plan to use

- xGUS ✓

**Accounting:** Which Accounting system do you use.

- APEL ✓

**Monitoring:** Which Monitoring System do you use

- SAM-NAGIOS based. ✓

**Registration:** Which Site Registration System do you use

- GOCGB. ✓

## **LATIN AMERICA ROC**

### **Contact Details**

- Name: Ramon Diacovo
- E-mail: [ramond@gmail.com](mailto:ramond@gmail.com)
- Institution: UFRJ
- ROC IGALC
- Region Latin America

**Helpdesk:** Which Ticketing system/helpdesk do you use or plan to use

- xGUS

**Accounting:** Which Accounting system do you use.

- APEL

**Monitoring:** Which Monitoring System do you use

- SAM-NAGIOS based.

**Registration:** Which Site Registration System do you use

- GOCGB.

## SEAsia ROC

### Contact Details

- Name: Eric Yen
- E-mail: Eric.Yen@twgrid.org
- Institution: Academia Sinica Grid Computing Centre
- ROC: Asia Pacific Regional Operation Centre
- Region: Asia Pacific

### Helpdesk: Which Ticketing system/helpdesk do you use or plan to use

- xGUS(v): GGUS only

### Accounting: Which Accounting system do you use.

- APEL (v)

### Monitoring: Which Monitoring System do you use

- SAM-NAGIOS based. (v)

### Registration: Which Site Registration System do you use

- GOCGB. (v)

### 3 Use Case: Africa and Arabia ROC

The Africa and Arabia ROC has taken advantage of the efforts of EGEE-III and EGI to regionalise many of the operations support tools. The decision was taken expressed by the Amman Declaration to provide single presence for support and operations coordination for a vast region, including all of Africa and much of the Middle East. The reason for including this vast region was the sparsity of resources in this region and the need to achieve a critical mass both of sites and of human resources.

#### 3.1 ROC Composition and Organisation

The Africa and Arabia Regional Operations Centre consists of the technical and operations teams of the sites comprising respectively the EUMedGridSupport project and the South African National Grid (SAGrid). It is a clearing house for information and services relevant to the operation, development and usage of grid services in the North-African, Middle East and Sub-Saharan African region.

The ROC is presented via a website<sup>3</sup> providing general information for parties of different roles, such as new or returning users, site administrators and the general public. This website acts as central contact point for these various roles and is hosted by the Jordanian NREN JUNET. The website was developed and maintained originally in the context of the EUMedGrid-Support project, which has subsequently finished. During the creation of the ROC, the South African National Grid (SAGrid) was invited to participate. Although for several years SAGrid operated its own set of support and coordination services via similar websites in South Africa, a decision was taken in 2011 to consolidate on the Africa and Arabia ROC, in order to promote interoperability and sustainability of these services in the regions covered by the ROC.

#### 3.2 Services Provided by the ROC

The services which are provided by the ROC can be described generically as :

- Support
- Monitoring
- Documentation
- Accounting
- Coordination

Each service has associated tools, which are used to implement the service itself and ease the functions associated with that service in an environment distributed across a vast geographical region. Many of the services are provided in a catch-all manner, with no distinction made for national or institutional boundaries, for example the monitoring

As a general principle it was decided when building the ROC to use as far as possible regionalised versions of operational tools developed by the EGEE projects. These are operated by the support teams of the sites comprising the ROC and the NREN's in the region to which they are connected – SANReN and ASREN currently and are:

---

<sup>3</sup> <http://roc.africa-grid.org>

1. **Support System:** A regionalised version of the Grid User Support helpdesk system (xGUS) is provided. The AAROC xGUS system is hosted by KIT and reads a similar database of issues to the central instance, allowing particular issues to high-lying support units defined in the EGI GGUS instance.
2. **Operations Database:** A regionalised version of the Grid Operations Database (GOCDB) is operated by AAROC, populated with information and contact details of the sites in the infrastructure.
3. **Monitoring and information systems:** Using the AAROC GOCDB, a topology is generated and used to configure the MyEGI and Nagios monitoring tools which provide information and alerts on service availability. These are the same tools as used by the various NGIs in Europe and EGI itself. Other monitors include the Real Time Monitor, and network monitoring tools such as smokeping<sup>4</sup>
4. **Accounting:** Accounting information for jobs, users and sites is collected using the DGAS/HLRmon accounting system. Currently however this service only provides accounting information for former EUMedGrid sites.

The services provided by the ROC to **users** are:

1. **Documentation**, including information on how to join as a new user, how to contact site administrators, and other experts via mailing lists, a series of user manuals and how-to information for basic usage of the grid services.
2. **VO contact points** for the catch-all Virtual Organisations hosted on the sites (eumed, africacert and sagrid)
3. **Information** regarding the current state of the infrastructure – including a full list of sites and services
4. **Monitoring information** collected by the back-end services which perform monitoring.
5. **Links** to other projects involved in the development and operation of e-Infrastructure in the region
6. **User support** via the regionalised GGUS portal.

### 3.3 *Standard Operating Procedures*

As the name suggests, the ROC is focussed mainly, but not entirely, on *operations* of e-Infrastructure, particularly grid computing infrastructures. Since resources are distributed in a region – in this case a vast geographic region – it is of vital importance to adopt certain standard procedures at each site and for common tasks. In the case of SAGrid, a formalism has been created<sup>5</sup> in order to deal with the creation and adoption of standard procedures in this environment, according to so-called “application areas” - management, documentation, operations and application porting<sup>6</sup>.

Since the infrastructure is a working and evolving environment, it is expected that standard procedures evolve as well. Therefore, it is foreseen that procedures have a lifetime associated with them, evolving from proposals, through calls for comment and input, then internal review to final

---

4 The perfSONAR monitor is being evaluated and is proposed as a further network monitoring tool.

5 See <https://ops.sagrid.ac.za/trac/wiki/ProposedOperationsProcedures>

6 These application areas have been selected so far, although others are not excluded.

drafts. Before being adopted, procedure proposals are expected to undergo external review, in order to ensure interoperability with other infrastructures and compatibility with best practices.

Aspects of this work which have attempted to address the issue of sustainability and self-consistency should be pointed out:

- The first procedure which was developed was a procedure for proposing procedures. Although this may seem tautologous, it is an important foundation for the development of further procedures, by any member of the team, according to a specific workflow and guidelines. This also limits divergence and unnecessary work needed to develop new procedures. Similarly, a procedure for adopting new procedures is provided.
- The development of new standard procedures is not strictly under the domain of the operations team, and the ROC – contrary to larger infrastructures such as EGI – does not have a dedicated policy team. Therefore work should be prioritised according to the internal needs of the ROC. In order to enable prioritisation, proposed procedures are voted on each month by the interested parties, in order to assign a new priority to the development of procedures.
- Proposals can exist only in certain states, and a workflow is provided for moving from state to state (for example, from first draft to second draft), along with the associated actions to take.

This work has drawn from the experience of existing and previous e-Infrastructures, such as EGEE, EELA/GISELA and EUMedGrid, but particularly EGI, where a dedicated policy team exists and has provided solid background material. We have considered the general case where a new procedure needs to be developed and adopted and tried to take into account the particular operating environment in our region, allowing most tasks to be distributed, but nonetheless be coordinated by a sort of “constitution” for standard procedures. It is expected that this effort to document the standard procedures will reduce the dependence on centralised experts and expertise, and thereby allow a reliable integrated infrastructure to be developed, with sites operating fairly autonomously.

## 4 Emerging Technologies relevant for Grid operations

In this chapter we briefly discuss new or emerging technologies that could be helpful or necessary to run an NGI/ROC in the near future.

### 4.1 IPv6

IPv6 is getting extremely important these days given the undeniable fact that new IPv4 address classes are about to be completely impossible to assign, due to the complete unavailability of the available address space. The Regional Internet Registries (RIRs) are running out of new IPv4 address classes to assign. For the exhaustion of IPv6 address classes, one can define 4 phases:

1. Stage I: IPv4 address still available, RIPE NCC's continue distributing IPv4 addresses according to current policy.
2. Stage II: IANA exhausts IPv4 pool, all five RIRs get one /8 of IPv4 space each, and RIPE NCC continues to distribute IPv4 addresses according to current policy.
3. Stage III (current stage): RIPE NCC reaches final /8: RIPE NCC's Allocation policy for the final /8 comes into effect.
4. Stage IV: RIPE NCC's IPv4 pool is depleted, the RIPE NCC can distribute only IPv6 addresses.

Predictions for end of Stage II - as of 27 May 2012<sup>7</sup>:

- APNIC: (already happened) - Asia Pacific - April 15, 2011
- RIPE NCC: (already happened) September 12, 2012 - Europe
- ARIN: June 24, 2013 - North America
- LACNIC: Feb 1, 2014 - South America
- AFRINIC: Nov 9, 2014 - AFRICA

For Stage III, the policy will be to allocate at maximum a /22 IPv4 subnet per LIR (Local Internet Registry) if and only if the LIR has already have activated IPv6. So, the importance of IPv6 is universally acknowledged.

#### 4.1.1 Why IPv6 and Grid

As for any network user, IPv6-only Grid users will sooner or later show up. This means that there might be the case of new resources available which will be reachable only in IPv6. Although integrating IPv4 and IPv6 resources in a unique Grid implies some form of well defined IP protocol translation, it is important in any case to be in a situation in which the same middleware could be used to build and IPv4 and an IPv6 Grid. Currently in the Grid domain there are 2 major objectives to pursue:

1. Ensuring that turning on the IPv6 stack (in a Dual Stack approach) on an IPv4 Grid resource does not break any IPv4-provided functionality
2. Ensuring that the middleware is fully IPv6 compliant - so that the same middleware we currently provide and use using IPv4 could be used using IPv6

A further step, would be to implement at some level (node level, Grid middleware level, IP-gateway) - protocol translation, so that IPv4 and IPv6 resources could effectively interoperate.

<sup>7</sup> See the world regions on <https://www.arin.net/knowledge/rirs/RIPEcountries.html>

#### 4.1.2 Current Status in Europe

Even though there is a realistic hope that UNICORE, gLite, ARC and dCache are IPv6 compliant, none of the middleware stacks are certified to be IPv6 compliant. Some preliminary work has been done already in previous European project like EGEE III and EU-ChinagGrid for gLite that is carried over the current releases of middleware. EMI, EGI and HEPiX IPv6 WG decided recently<sup>8</sup> to join efforts to build a unique IPv6 testbed distributed over various sites. This new taskforce will speed up the certification of IPv6 compliance of the middleware stacks. EGI InSPIREs Software Provisioning Activity<sup>9</sup> plans to add IPv6 compliance as an optional Quality Criterion in the next update of the UMD Quality Criteria<sup>10</sup> due in January 2012, that will be used to monitor the progress of the joined taskforce. Updated information on the progress of the joined taskforce is reported on <https://wiki.egi.eu/wiki/IPv6TestReports>.

#### 4.2 Cloud Infrastructures

The majority of the resources currently available on the European Grid Infrastructure follow the Platform as a Service model, i.e we deploy grid middleware on top of hardware directly in order to get the most out of it as possible (Figure below). In the last couple of years however there is a tendency to virtualise at least part of the infrastructure in order to have greater flexibility in deployment scenarios and reduce maintenance effort.

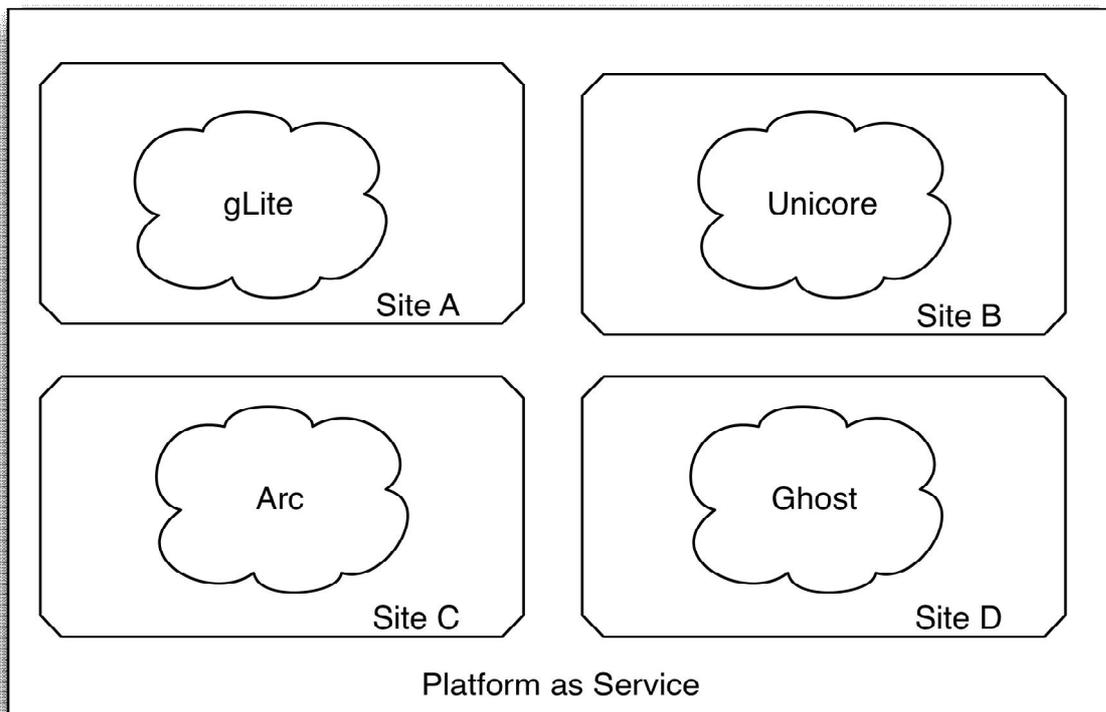


Figure 4 - PasS scenario

<sup>8</sup> <http://indico.cern.ch/conferenceDisplay.py?confId=186950>

<sup>9</sup> [https://wiki.egi.eu/wiki/WP5:\\_Provisioning\\_the\\_Software\\_Infrastructure\\_\(SA2\)](https://wiki.egi.eu/wiki/WP5:_Provisioning_the_Software_Infrastructure_(SA2))

<sup>10</sup> <https://wiki.egi.eu/wiki/EGI-InSPIRE:UMDQualityCriteria>

In parallel a number of National Research and Education Networks (NREN) in Europe are currently deploying their own private cloud infrastructure<sup>11</sup> (Infrastructure as a service model IaaS) in order to offer more services to their users and align their strategy to European Union “Europe-2020”<sup>12</sup> Roadmap. This will lead NRENs that are currently offering Grid services to run two distinct infrastructures simultaneously (Figure below).

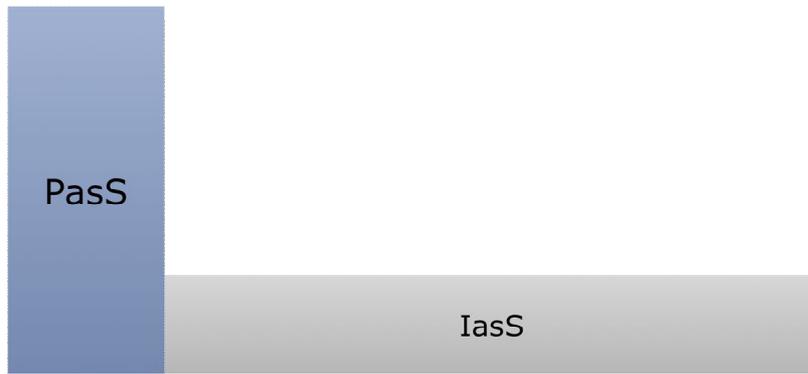


Figure 5 - Available services today

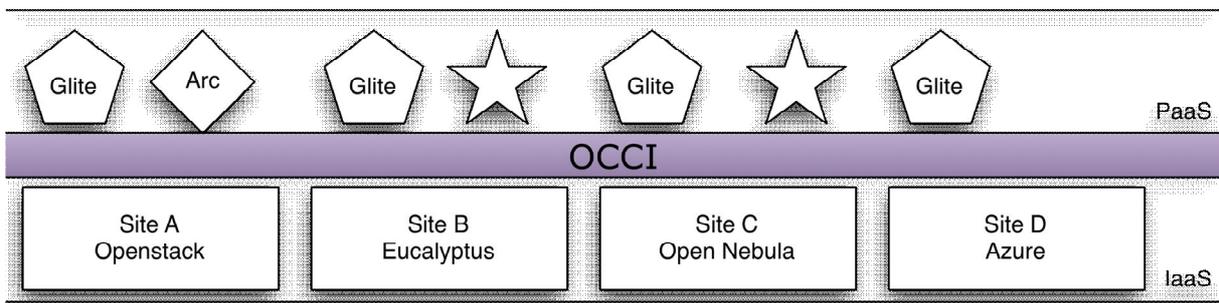


Figure 6 - Schematic that offers PasS over IaaS

#### 4.2.1 GRNET ~Okeanos IaaS use case

The goal of the ~Okeanos project is to deliver a production quality IaaS. GRNET operates a working alpha version since July 2011; the alpha offering comprises 350 VMs and 200 users. In order to reach its goals, Okeanos is a jigsaw puzzle of many pieces: the UI part, an API, an Image Registry, a VM management component, Networking facilities, Storage, Monitoring, Identity Management, Crediting, Billing, Issue Handling, and Helpdesk operation. It goes beyond commercial IaaS providers in several ways: Amazon EC2, and comparable commercial offerings, are not an end-user service, while Okeanos is designed to be used by people with little computer

<sup>11</sup> GRNET Okeanos Service: <https://okeanos.grnet.gr/>

<sup>12</sup> [http://ec.europa.eu/europe2020/index\\_en.htm](http://ec.europa.eu/europe2020/index_en.htm)

experience. At the same time it aims to meet the needs of advanced users in technical departments by offering persistent, long-term servers with custom networking capabilities.

The software underlying Okeanos, called Synnefo, is a custom cloud management software with Google Ganeti backend. Ganeti was chosen because it is a scalable and proven software infrastructure, and GRNET has already long experience with Ganeti, using it to provide VMs to Network Operation Centres. GRNET is also involved in Ganeti development, and contributes patches upstream.

Recently GRNET deployed a fully virtualised EMI glite middleware site (HG-08-Okeanos<sup>13</sup>) as a pilot in order to evaluate the overhead caused by the virtualisation of the resources and identify possible issues with the current middleware running over Okeanos. As soon as the new site is stable enough we will also evaluate elasticity scenarios for the on demand provisioning of additional virtual Worker Nodes under an EMI CREAM using an OCCI compliant version of the service that was showcased for openstack during the EGI TF 2012<sup>14</sup>. A flowchart showing the operational details of this service is shown in Figure 7. When starting the service a check on the computing queue is made. If certain thresholds are met (in this case if the number of waiting jobs exceeded the number of available computing resources at start) a fixed number (X) of new virtual machines are created and added to the computing element queue as virtual Worker Nodes. The check is rerun until the size of the queue (i.e. the number of waiting jobs decreases beyond that certain threshold. In such a case all virtual Worker Nodes are marked as offline, which means they can no longer accept new jobs but they can still process any jobs that are already in the running state. The service then checks periodically the virtual Worker Nodes one by one whether the running jobs have finished in which case a flag to terminate the virtual Worker Node is turned on. Once all jobs have finished the virtual Worker Nodes are terminated and the service starts from scratch by checking again the queue.

The service described here can be further extended with the implementation of an OCCI interface in the cloud back-end. This feature would enable the usage of existing cloud infrastructures (i.e. infrastructures offered in the context of the EGI Federated Cloud) for the on demand provisioning of additional Worker Nodes in a cluster. Moreover, on the Grid side, the service can be extended to also work under different types of Grid enabled computing services (i.e. for ARC and UNICORE based computing infrastructures) as well as for plain cluster infrastructures. GRNET considers this service crucial as it would allow us expand our current grid infrastructure over the newly acquired okeanos IaaS service thus providing a stepping stone between the GRID computing model and the emerging cloud computing one. Similarly we would advise chain/chain-red partners to investigate similar solutions to see if they fit their needs.

---

<sup>13</sup> [https://goc.egi.eu/portal/index.php?Page\\_Type=View\\_Object&object\\_id=25327&grid\\_id=0](https://goc.egi.eu/portal/index.php?Page_Type=View_Object&object_id=25327&grid_id=0)

<sup>14</sup> <https://indico.egi.eu/indico/contributionDisplay.py?sessionId=47&contribId=62&confId=1019>

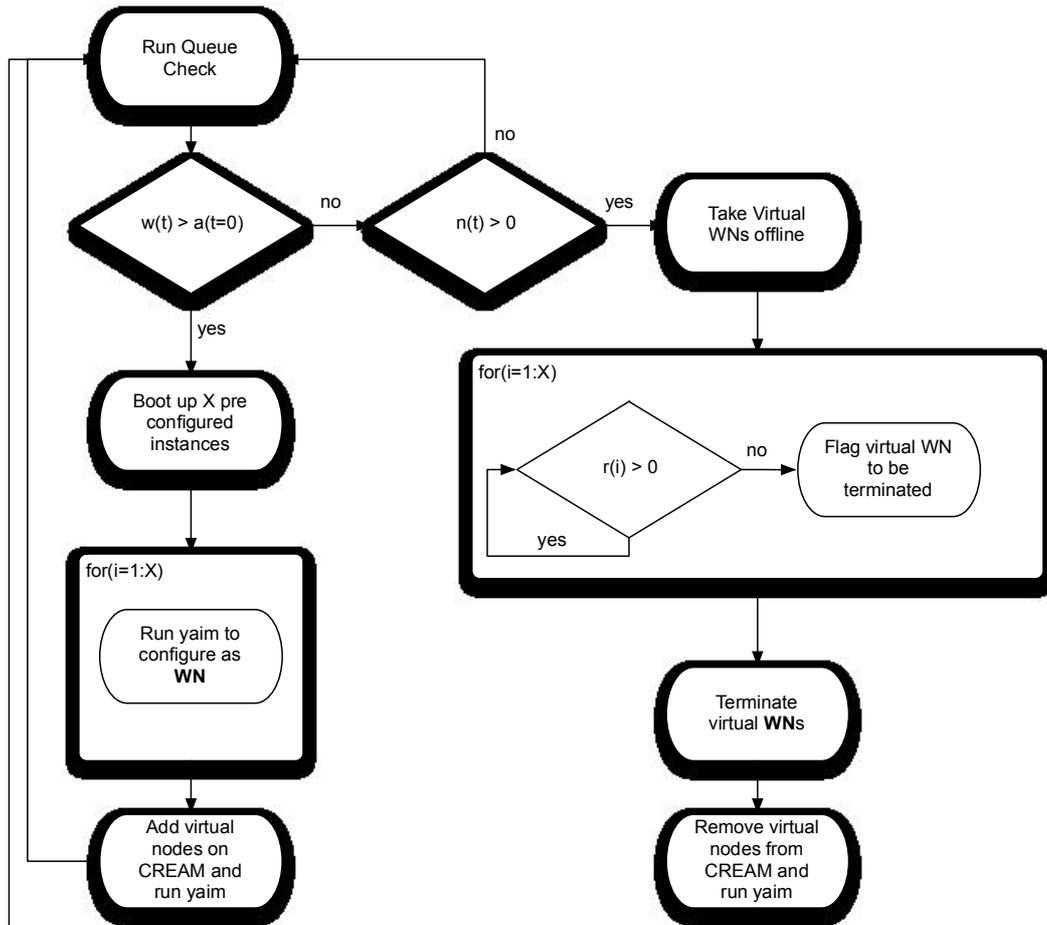


Figure 7 - Execution flow of scaling service used for the on demand provisioning of virtual Worker Nodes