FP7-ICT-2009-4 (248421)
SeC
Sim-e-Child

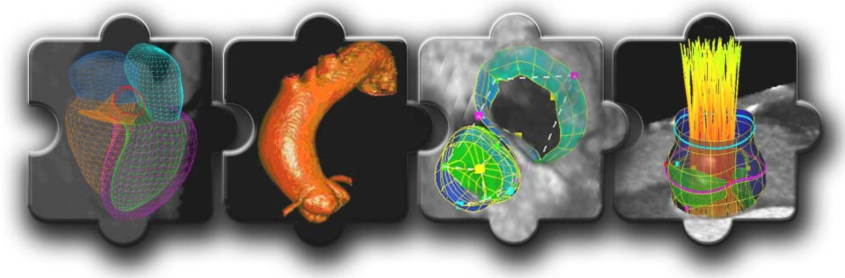

Collaboration Project

Thematic Priority: ICT

Deliverable 4.1
Grid and Databases Connection Report

Due date of delivery: 31 October 2010
Actual submission date: 30 December 2010

Start date of project: 1 January 2010
Ending date: 30 June 2012

Partner responsible for this deliverable: MAAT France

SEVENTH FRAMEWORK PROGRAMME

Revision 1

Project co-funded by the European Commission within the FP7

Dissemination level	
RE	Restricted to a group specified by the consortium

Document Classification

Title	Grid and Databases Connection Report
Deliverable	D4.1
Reporting Period	January 2010 - October 2010
Authors	Jérôme Revillard
Workpackage	WP4 - Simulation and Collaboration Platform Development
Security	Restricted
Nature	Report
Keywords	

Document History

Name	Remark	Version	Date
Jerome Revillard		0.1	15/10/2010

Sim-e-Child Consortium

The partners in this project are:
01. Siemens AG (Siemens)
02. Lynkeus Srl (Lynkeus)
04. maat France (MAAT)
05. Technische Universität München (TUM)
06. I.R.C.C.S. Ospedale Pediatrico Bambino Gesù (OPBG)
07. Siemens Corporate Research, Inc. (SCR)
08. Johns Hopkins University (JHU)
10. American College of Cardiology Foundation (ACCF)
11. Siemens Program and System Engineering srl (PSE)

List of contributors

Name	Affiliation	Co-author of
Jerome Revillard	MAAT	
Fabian Moerchen	SCR	

List of reviewers

Name	Affiliation
Michael Suehling	Siemens

Table of Contents

Bibliography

[1]	gLite Grid middleware website - http://www.glite.org/
[2]	EGI website - http://www.egi.eu
[3]	Wang F, Bourgué PE, Hackenberg G, et al: "SciPort: An Adaptable Scientific Data Integration Platform for Collaborative Scientific Research". Proc. VLDB, 1310-1313. 2007
[4]	Coarctation Of the Aorta Stent Trial (COAST) http://www.clinicaltrials.gov/ct2/show/NCT00552812?term=coast&rank=2
[5]	DICOM: http://medical.nema.org
[6]	caBIG Annotation Imaging Markup (AIM): https://cabig.nci.nih.gov/tools/AIM
[7]	jBoss: http://www.jboss.org/

1. Introduction

1.1. Purpose

This document aims to report on both the design and development of the Sim-e-Child (SeC) grid and databases infrastructure. The grid infrastructure is based on the gLite grid middleware [1] from the European Grid Infrastructure (EGI) project [2]. This document therefore presents the deployment of the SeC infrastructure as it was deployed at month 10 of the project.

The database infrastructure is based on the SciPort [3] platform developed by Siemens Corporate Research. SciPort is a data management platform with a service oriented architecture and a web-based user interface. The platform provides tools to integrate, annotate, and share scientific data across research communities. It enables the design of complex data models that integrate experimental data such as documents, spreadsheet, images, annotations, lab results, or genetic information in arbitrary file formats.

1.2. Overview of the document structure

This Document will be divided in mainly two parts. The first part will describe the status of the Grid infrastructure at month 10 and the second part will describe the databases connections. As far as the Grid part is concerned, the gLite grid middleware will be first presented with all the services that are used in the project. Then an overview of the deployment will be given.

1.3. Abbreviations

BDII	Berkeley Database Information Index
BLAH	Batch Local ASCII Helper Protocol
CREAM	Computing Resource Execution And Management
EGI	European Grid Infrastructure
FTS	File Transfer Service
GIIS	Grid Index Information Server
GRIS	Grid Resource Information Server
GFAL	Grid File Access Library
gLite	EGI Grid middleware stack
HeC	Health-e-Child
ICE	Interface to CREAM Environment
IS	Information System (grid-level)
LB	Logging and Bookkeeping
RB	Resource Broker
SE	Storage Element
SeC	Sim-e-Child

VO	Virtual Organization
VOMS	Virtual Organization Membership Service
VDT	Virtual Data Toolkit
WN	Worker Node
WMS	Workload Management System

2. Grid

2.1. Introduction

Grid Middleware refers to the security, resource management, data access, instrumentation, policy, accounting, and other services required for applications, users, and resource providers to operate effectively in a Grid environment. Middleware acts as a sort of 'glue' which binds these services together. Grid middleware is built by layered interacting packages.

- A grid middleware is an internet based system that needs efficient and reliable communication and is a blend of high performance systems and high throughput computing
- A grid middleware is data aware and all data access and replications decisions are based onbase bound at least for the following functions: grid topology management user access and certification dataset locations and replicas resource definition and dynamical management performance and user bookkeeping
- A grid middleware is bound to efficient matching and scheduling algorithms to find best available resources for the task execution and resource brokering.
- A grid middleware depends on accurate clock performances to synchronize nodes and correctly handle task and job scheduling

Typical Grid based distributed infrastructure consists of the following logical components:

- Virtual Organization (VO)
- Grid Information System
- Work Load Management System
- Resource broker
- Computing element
- Data and Meta-Data Management System
- File Catalogue
- Storage Resource
- File Transfer Service
- User Interface

In the next section, a more detailed description of those components will be done

2.2. The gLite Grid Middleware

The gLite distribution [1] is an integrated set of components designed to enable resource sharing. In other words, this is middleware for building a grid.

The gLite middleware is produced by the EGI project. In addition to code developed within the project, the gLite distribution pulls together contributions from many other projects, including LCG and VDT. The distribution model is to construct different services ('node-types') from these components and then ensure easy installation and configuration on the chosen platforms (currently Scientific Linux versions 4 and 5).

The gLite Middleware is a quite complex framework which follows a Service Oriented Architecture (SOA) to facilitate interoperability among Grid services and to allow easier compliance with upcoming standards. This architecture has the advantage not to be bound to specific implementations and services are expected to work together but can also be used independently.

In the Sim-e-Child project, the following gLite components will be used:

- **LCG File Catalog (LFC)** offers a hierarchical view of files to users, with a UNIX-like client interface. The LFC catalog provides Logical File Name (LFN) to Storage URL (SURL) mappings and authorization for file access. The LFNs are aliases created by a user to refer to actual data. Simple metadata can be associated to them. The authorization is performed using UNIX-style permissions, POSIX Access Control Lists (ACL) and VOMS support. The LFC uses a client-server model with a proprietary protocol. LFC server communicates with a database (either Oracle or MySQL), where all the data is stored. LFC catalogue also exposes a Data Location Interface (DLI) - a web service used by applications and Resource Brokers. Provided with a LFN, the DLI returns the actual location of the file replicas.
- **Storage Resource Managers (SRM)** is a service somewhat similar to the cluster batch system but instead of managing processors and jobs it manages requests for storage space and files. The storage space managed can be disk space, tape space or a combination of the two. There is a number of SRM implementations for disk storage management that are widely deployed. **Disk Pool Manager (DPM)** is a recommended solution for lightweight deployment of smaller sites because it is easy to install and requires very low maintenance effort. It features full implementation of SRM. Bigger sites usually choose dCache because of robustness, scalability and advanced features. CERN Advanced STORage (CASTOR) is an implementation used by sites that have both disk and tape storage.
- **BDII:** The Information System (IS) provides information about the status of Grid services and available resources. Job and data management services publish their status through Grid Resource Information Server (GRIS). GRIS runs on every service node and is implemented using OpenLDAP, an open source implementation of the Lightweight Directory Access Protocol (LDAP). Ever grid site also runs one Grid Index Information Server (GIIS). The GIIS queries the service GRISes on the site and acts as a cache storing information about all available site services. Finally, the information from GIISes is collected by Berkeley Database Information Index (BDII). The BDII queries the GIISes (sometimes also called site BDIIs) and acts as a cache storing information about all available Grid services in its database. Figure 2.1 shows this flow of information. Users and programs interested in status of the Grid usually query the top level BDII as it contains information about all the services that are available.

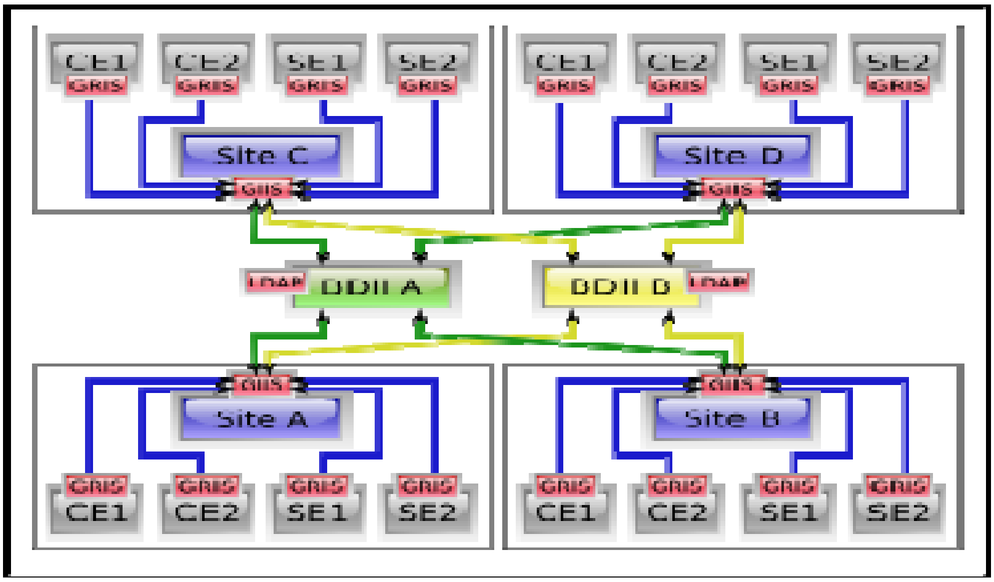

Figure 2.1: gLite Information System hierarchy

- **MyProxy (PX)** (Not a gLite component)**:** On Grids, users authenticate themselves using temporary credentials called proxy certificates, which contain also the corresponding private key. Proxy certificates do not represent a significant security risk only if they are reasonably short-lived (by default, a dozen hours). For longer jobs, PX plays a role of online credential repository. For such long running jobs a proxy renewal system is used, consisting of a **Proxy Renewal Service (PRS)** on the RB and a PX server on a dedicated host. A PX stores long-lived user proxies (with a lifetime of several days, usually) which it uses to generate, on request of the PRS, short-lived proxies for jobs whose proxies are about to expire.
- **VOMS:** The Virtual Organisation Membership Service (VOMS) is a service to manage authorization information in VO scope. The VOMS system should be used to include VO membership and any related authorization information in a user's proxy certificate. These proxies will be said to have VOMS extensions. The user gives the voms-proxy-init command instead of grid-proxy-init, and a VOMS server will be contacted to check the user's certificate and create a proxy certificate with VOMS information included. By using that certificate, the VO of a user will be present in every action that he performs.
- **Computing Element (CE)** is the service representing a computing resource. It provides a virtualization of the computing resource localized at a site (typically a batch queue of a cluster but also supercomputers or even single workstations). It provides information about the underlying resource and offers a common interface to submit and manage jobs on the resource. **CREAM** (Computing Resource Execution And Management) is a simple, lightweight service that implements all the operations required at the CE level. This is the implementation that will be used. Its interface is defined using WSDL. The service is compliant with the existing BES standard. CREAM can be used by the WMS, via the ICE component (see next description of WMS), or by a generic client, e.g. an end-user willing to directly submit jobs to a CREAM CE. A C++ command line interface and Java clients are available for this purpose.

The CE hide a collection of **Worker Nodes (WN)** or just one multiprocessor WN, the node(s) where the jobs are run concretely.

- **Workload Management System (WMS)** is a Grid level meta-scheduler that schedules jobs on the available CEs according to user preferences and several policies. It also keeps track of the jobs it manages in a consistent way. The core component of the WMS is the **Workload Manager (WM)**, whose purpose is to accept and satisfy requests for job management coming from its clients (i.e., computational job submission). In particular the meaning of the submission request is to pass the responsibility of the job to the WM. The WM will then pass the job to an appropriate Computing Element for execution, taking into account the requirements and the preferences expressed in the job description. The decision of which resource should be used is the outcome of a matchmaking process between submission requests and available resources. The Resource Broker (RB) or Matchmaker as WMS component offers support to the WM in taking the above mentioned decision. It provides a matchmaking service based on a given user's job description – find a resource that best match the request. A WMS instance interacts with several other services. Tracking job lifetime relies on the Logging and Bookkeeping Service. Information on service availability, resource status and data localization is gathered from appropriate sources, such as LFC and BDII. Security-related aspects are addressed interacting with VOMS, Proxy Renewal and G-PBox. Another important component of the gLite WMS is the Interface to CREAM Environment (ICE). It provides the connection between the gLite WMS and the CREAM CE. ICE, running in the gLite WMS node along with the other processes of the gLite WMS, receives job submissions and other job management requests from the WM component of the WMS and then invokes the appropriate CREAM methods to perform the requested operation.
- **Logging and Bookkeeping Service (LB):** The primary purpose of the Logging and Bookkeeping service (LB) is tracking Grid jobs as they are processed by various Grid middleware components. It collects and stores in a database the job status information supplied by the different components of the WMS system. The collection is done by LB local-loggers, which run on the RB and on the CE, while the LB server, which normally runs on the RB, saves the collected information in the database. The database can be queried by the user from the UI, and by RB services. The information that is gathered in LB is used to inform Grid users on the job state. This information is also useful for example for debugging of user jobs.

In the case of gLite grid middleware based grid infrastructure the above components map to the following stack of gLite middleware and non-gLite *external* services.

Grid Infrastructure Components	Services (gLite and external)
Virtual Organization	VOMS MyProxy (non-gLite)
Grid Information System	top BDII site BDII*
Work-Load Management System	
Resource broker	WMS
Computing element / Worker Node(s)	CREAM-CE* / WN*
Data Management System	
File Catalogue	LFC
Storage Resource	DPM*

Tableau 2-1: Generic Grid infrastructure services mapped to the gLite services stack.

* gLite site services: All other services are grid infrastructure core services

Typical deployment of Grid infrastructure built with gLite middleware (c.f. Figure 2-2) assumes deployment of one instance of each core grid service and *N* grid sites (comprised of site grid services)

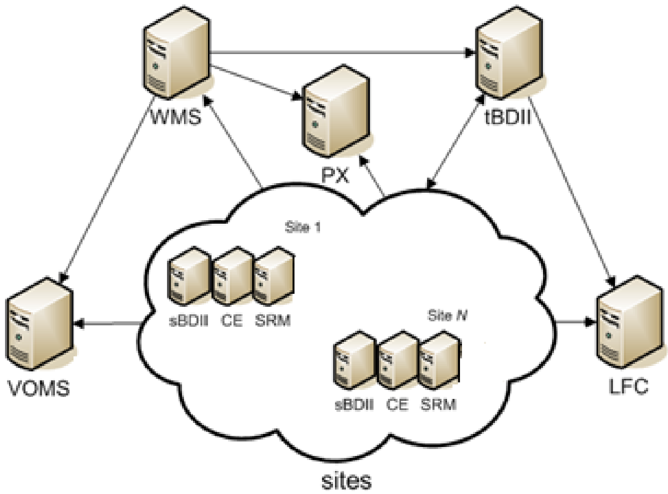

Figure 2-2: Typical deployment of Grid infrastructure built with gLite middleware (lines are logical connections between services and components)

2.3. Installation Status

The next sections present the installation status at month 10 of the Sim-e-Child project.

2.3.1. Central Services

The SeC central services are hosted at MAAT France (Archamps - France). The following service are concretely deploy in the very latest middleware version available

Host name	OS	gLite service(s)
voms.maatg.eu	Scientific Linux 4 32b	VOMS
bdii.maatg.eu	Scientific Linux 5 64b	Top BDII
wms.maatg.eu	Scientific Linux 4 32b	WMS
lfc.maatg.eu	Scientific Linux 5 64b	LFC
myproxy.maatg.eu	Scientific Linux 4 32b	MyProxy

On the VOMS server, a specific virtual organization has been created named “vo.sim-e-child.org”. All the SeC users will be grouped inside this virtual organization.

2.3.2. Site Services

As far as the deployment of the sites is concerned, the objective is to:

- attach the Health-e-Child production sites,
- install a site in Europe,
- install a site in United States.

These points will be explained in more detail in the following sections.

2.3.2.1. The Health-e-Child sites

In order to facilitate the future maintenance of HeC and SeC as well as to renew the Grid infrastructure configuration so to match latest de facto standards from the European Grid Initiative (EGI), a complete migration of the SeC information system (i.e. grid infrastructure central services) has been operated.

Thus, the 2 first months of task T4.1 were devoted to the logistics setup and allowed carefully planning the progressive migration of infrastructural services. Great care was given to applying appropriate naming conventions, good practices and EGI compliant deployment procedures. Constant interactions took place with EGI system administrators in order to settle deployment questions that aroused from this migration.

As a consequence, SeC now hosts the HeC infrastructure pillars and is up-to-date with respect to acknowledged Grid standards.

The migration now being completed and first nodes hosting simulation and data being on the way, technical partners in task T4.8 will recontact former clinical centers to propose their migration to the latest version of the Grid middleware and Gateway.
Thus OPBG, NECKER, GOSH and IGG will be recontacted and their migration scheduled accordingly. OPBG will be treated first as being a partner of the SeC project with concrete simulation needs.

Former HeC end-users have been contacted and informed about the major infrastructural migration from HeC to SeC. Expressions of interest have been gathered so to keep concerned users informed until the simulation platform and pillar resources are online.

2.3.2.2. The European site

The Europe site is located at Maat France in Archamps. Currently, most of the services are in place:

Host name	OS	gLite service(s)
ng-maat-server3.maatg.eu	Scientific Linux 5 64b	DPM
bdii-site.maatg.eu	Scientific Linux 5 64b	Site BDII
ng-maat-server9.maatg.eu	Scientific Linux 5 64b	CE

As far as the developments are concerned, a specific worker node with only 8 cores was configured to provide a full environment to the developers:

Host name	OS	gLite service(s)
ng-maat-server9.maatg.eu	Scientific Linux 5 64b	WN

This site will be gradually improved in term of computing power once the simulations will be integrated.

2.3.2.3. The United States site

The JHU Gateway (formerly deployed by and for HeC) has been kept as is, as a means to further utilizing the SeC simulation platform, i.e. offering a high quality of services access point. The technical partners, in particular MAAT and SIEMENS decided to externalize the US Grid node resources so to facilitate its expansion, while not impacting on JHU's IT department.
Consequently, the US Grid node resources will be subcontracted from an appropriate Cloud provider on a requirements basis. This will make it possible to reconfigure the node on demand and will allow to rapid prototype and test SciPort and iKDD applications integration within the SeC simulation platform. The technical partners therefore expect to rent the following amount of to guaranty the minimal site setup:

- 1 virtual machine for the CREAM-CE service
- 1 virtual machine for the site BDII service
- 1 virtual machine for the DPM service
- 1 or more virtual machine(s) for the WN(s) (on demand deployment)
- 1 virtual machine for the SciPort service

We already looked at different kinds of providers:

- Cloud providers
  - Amazon EC2
  - Terremark vCloud Express
- Dedicated servers service providers:
  - softlayer.com
  - OVH (soon in USA)

2.3.2.3.1. Cloud Providers

- Amazon EC2:

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.

By default, the virtual machines (instances) are not persistent. This means that after every reboot, the virtual machine will be reset to the default state. In order to avoid it, the Amazon Elastic Block Store (EBS) feature has to be used.

As far as this provider is concerned, it's very difficult to know in advance the cost of a SeC site node because of the EBS volumes. Indeed, the price evolves with the number of I/O requests that will be done on them ($0.10 per million I/O) and amount of space that will be used by the snapshots.

Amazon provide access to different type of instances, from which the one which is of interest for us is:

- Large Instance
  - 7.5 GB memory
  - 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each)

- 850 GB instance storage
- 64-bit platform
- I/O Performance: High

For some of the services (i.e: sBDII), this configuration is a bit too big but the problem is that the smallest one is only available on 32b platforms and 64b ones are mandatory for the gLite Grid middleware.

Nevertheless, here is an estimation of the costs (per month):

```
Amazon EC2 Service (US-EAST)                                  $1701.90
     • Compute:                                 $1244.40
     • Transfer:                                $   0.00
     • EBS Volumes:                             $  79.00
     • EBS I/O Requests:                        $ 260.00
     • Reserved Instances (One-time Fee):       $ 118.50
     • Elastic IPs:                             $   0.00
     • Elastic LBs:                             $   0.00
     • Data Processed by Elastic LBs:           $   0.00
AWS Data Transfer In                                          $  10.00
     • US East & US Standard Region:            $ 10.00
     • US West Region:                          $  0.00
     • Europe Region:                           $  0.00
     • Singapore Region:                        $  0.00
AWS Data Transfer Out                                         $  14.85
     • US East & US Standard Region:            $ 14.85
     • US West Region:                          $  0.00
     • Europe Region:                           $  0.00
     • Singapore Region:                        $  0.00
                                              -----------------------
                              Total Monthly Payment: $ 1726.75
                               Total Yearly Payment: $20721.00
```

These costs correspond to this specific configuration:

- 5 “Large” Linux Amazon EC2 On-Demand Instance (running 100% of the time)
- 5 Amazon EBS Volumes:
  - 20 TB, 200 I/O per second
  - 20 TB, 200 I/O per second
  - 50 TB, 200 I/O per second
  - 200GB, 200 I/O per second
  - 500GB, 200 I/O per second
- Amazon EC2 Bandwidth:
  - Data Transfer In: 100GB/month
  - Data Transfer Out: 100GB/month

This kind of service is really expensive and even for this price, the node capacity is really small in term of computing power as only 4 CPU cores are available in a single WN!

- Terremark vCloud Express

Terremark vCloud Express is a flexible, high-performance computing VMWare based cloud. Based on the virtual machines specifications that are needed for SeC, here is an approximate cost calculation:

```
CREAM-CE:
• RAM/CPU: 2G + 4 VPU = $0.161/hour =    1410,36 $/year
• DISK: 20G * 0.25$/month =                60,00 $/year
• 1PUBLIC IP: $0.01/hour =                  2.88 $/year
                                         ---------------
                                 TOTAL = 1473,24 $/year

sBDII:
• RAM/CPU: 2G + 2 VPUs = $0.141/hour =   1235,16 $/year
• DISK: 20G * 0.25$/month           =      60,00 $/year
• 1PUBLIC IP: $0.01/hour =                  2.88 $/year
                                         ---------------
                                 TOTAL = 1268,04 $/year

DPM:
• RAM: 4G +4 VPU = $0.301/hour =         2636,76 $/year
• DISK: 500G * 0.25$/month =             1500,00 $/year
• 1PUBLIC IP: $0.01/hour =                  2.88 $/year
                                         ---------------
                                 TOTAL = 4139.64 $/year

WN 1:
• RAM: 4G + 4 VPU = $0.301/hour =        2636,76 $/year
• DISK: 50G * 0.25$/month =               150,00 $/year
• 1PUBLIC IP: $0.01/hour =                  2.88 $/year
                                         ---------------
                                 TOTAL = 2789,64 $/year
SciPort (SeC database):
• RAM: 4G + 4 VPU = $0.301/hour =        2636,76 $/year
• DISK: 200G * 0.25$/month =              600,00 $/year
• 1PUBLIC IP: $0.01/hour =                  2.88 $/year
                                         ---------------
                                 TOTAL = 3239,64 $/year
                                 -----------------------
                                 -----------------------
                           TOTAL:       12940,20  $/year
```

This kind of service is really expensive and even for this price, the node capacity is really small in term of computing power as only 4 CPU cores are available in a single WN as for the Amazon cloud.

2.3.2.3.2. *Dedicated servers service providers*

This kind of service is a bit different than the cloud computing. Here you are renting real physical machines onto which you have full control. The goal here is to get 2 big configurations that will be then virtualized.

- Softlayer

	Monthly	Setup
Data Center		
Washington, DC		
Server	$499.00	$0.00
Dual Processor Quad Core Xeon 5550 - 2.66GHz (Nehalem) - 2 x 8MB cache w/HT **Second Processor Quad Core Xeon 5550 - 2.66GHz (Nehalem) - 2 x 8MB cache w/HT**		
Operating System	$0.00	$0.00
Citrix XenServer 5.6		
Ram	$450.00	$0.00
24 GB DDR3 Registered 1333		
Disk Controller	$50.00	$0.00
RAID 0		
First Hard Drive	$50.00	$0.00
1.00TB SATA II		
Second Hard Drive	$50.00	$0.00
1.00TB SATA II		
Public Bandwidth	$0.00	$0.00
3000 GB Bandwidth		
Uplink Port Speeds	$20.00	$0.00
1000 Mbps Public & Private Networks		
Primary IP Addresses	$0.00	$0.00
1 IP Address		
Public Secondary IP Addresses	$0.00	$0.00
4 Public IP Addresses		
Pre-Tax Total	**$1,119.00**	**$0.00**

For the project needs, two servers like this will be needed, which means in total: 26856.00$/year (pre-Tax)!

- OVH:

Prices were requested at OVH even if they do not yet propose any dedicated servers in US but it should be really soon. Moreover, as it will be demonstrated in the following paragraphs, the prices are really cheap comparing to the others solutions:

The best offer for the project needs is the “HG 2010-10G-12 BestOF”: http://www.ovh.com/fr/produits/hg_2010_bestof.xml

```
Specifications:
      16 VPU (Intel Bi Xeon E5530),
      24G DDR3, 2 x 1TB (AS),
      10gbps internet > OVH,
      200Mbps OVH > Internet,
      1 fix IP (we need more but it’s not expensive at all).

Installation:        338.30€
Location:            299.00€/month
                   ---------------
 TOTAL (pre-Tax) = 3926.30€/year
```

For the SeC project needs, we will need 2 servers like this which means 7852.60€ HT/year.

As you can see, for a very powerful configuration, the price here is much lower than all the other ones.

2.3.2.3.3. Conclusion

Those are not the only providers that were consulted for prices but they give a good idea if the prices when a good configuration is requested. Globally, it’s easily understandable that the cloud provider solution is not something that can be done in the SeC project because of the prices. Also, the provided flexibility in term of officially supported operating system does not contains, most of the time, the needed Scientific Linux 5 X86_64 that is required for the Grid services.

We are still waiting one last estimate from another provider (Atlantic.Net) but currently, OVH is clearly far ahead of the others. Indeed, the price is really low but the configuration is really powerful. Also, the provided internet bandwidth is impressive (10Gbps incoming). The only disadvantage of such an approach is that OVH still don’t have any dedicated server available in US; this should be the case during next year.

3. Databases Connection

3.1. Health-e-Child database

In the Health-e-Child project a generic database schema was developed and implemented using the lightweight MySQL database system. The schema was able to capture complex clinical parameters describing the patient history. This schema could also capture additional parameters derived from clinical images. However, the databases did not include the actual image data or provided mechanisms to attach other files such as physiological models derived from the images to the database. Hence it was decided to utilize the SciPort database system for the Sim-e-Child project. SciPort provides data modeling capabilities beyond clinical parameters and supports arbitrary file attachments to handle medical images and complex physiological and simulation models derived from them.

3.2. COAST database

The COAST (Coarctation Of the Aorta Stent Trial) [3] database for the US multi-center clinical trial of congenital aortic arch disease led by JHU is located at the Children’s Hospital Boston. Any data used by the Sim-e-Child project, including MRI images acquired in a sub study need to be de-identified by JHU before use by other partners in the Sim-e-Child project. For medical images SCR provided JHU with a tool to deidentify the medical images keeping a private mapping from the new Sim-e-Child study identifier to the original identifier removed. For the clinical data an anonymized export in form of spreadsheets was generated.

3.3. Clinical data systems at OBPG

The clinical data from OBPG will be directly extracted from the Hospital and Radiology Information Systems (HIS, RIS) in form of DICOM images and Word documents with clinician notes and clinical parameters.

3.4. SciPort databases for Sim-e-Child

SciPort is a data management platform with a service oriented architecture and a web-based user interface. The platform provides tools to integrate, annotate, and share scientific data across research communities. It enables the design of complex data models that integrate experimental data such as documents, spreadsheet, images, annotations, lab results, or genetic information in arbitrary file formats. The core of SciPort is formed by three tightly integrated components

1. A file repository to organize scientific files in a hierarchical structure with full text index and with transparent access for analysis tools.
2. A flexible meta-data model that can support a rich set of data types (text, numbers, dates, pick lists, multiple choice) and hierarchical structures such as section, subsections, and tables. The meta-data model is serialized as XML and references the files stored in the repository.
3. A relational database storing the user defined data model and annotations as XML along with structured information defining the users, groups, permissions, etc.

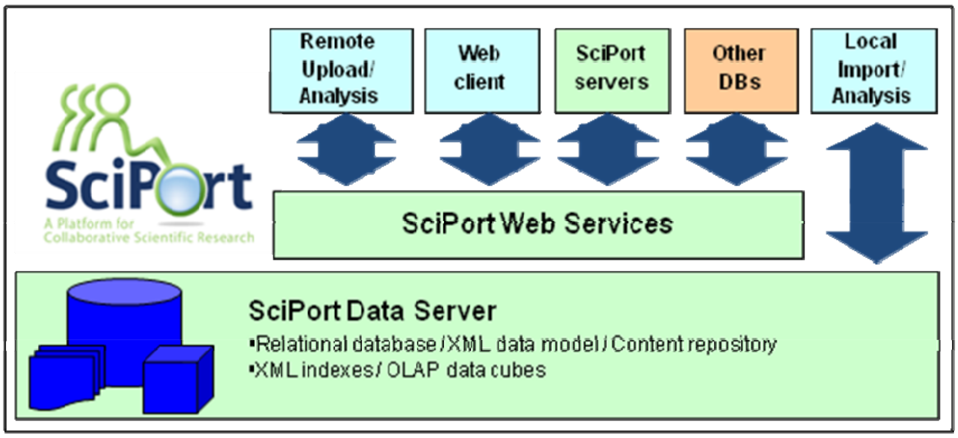

Figure 3-1 Architecture of SciPort: The data server wraps state-of-the art storage and indexing components with secure web services for multiple uses and remote access from other SciPort data servers.

Other features of SciPort include

- A service oriented architecture as shown in Figure 3-1.
- Web-based user interface with form editor, data entry forms, and dynamic meta-data based search interface.
- Data sharing and distributed search among multiple connected servers via secure web services (Figure 3-2).
- Support for common file formats including DICOM [4] and AIM (Annotation Image Markup) [5] for medical images, other standard image Formats (TIFF, raw) typically used in translational studies.Audit trail, electronic signatures, and LDAP (Lightweight Directory Access Protocol) authentication to support 21 CFR Part 11 compliance
- Flexible and adaptable data models that ensure backward compatibility.

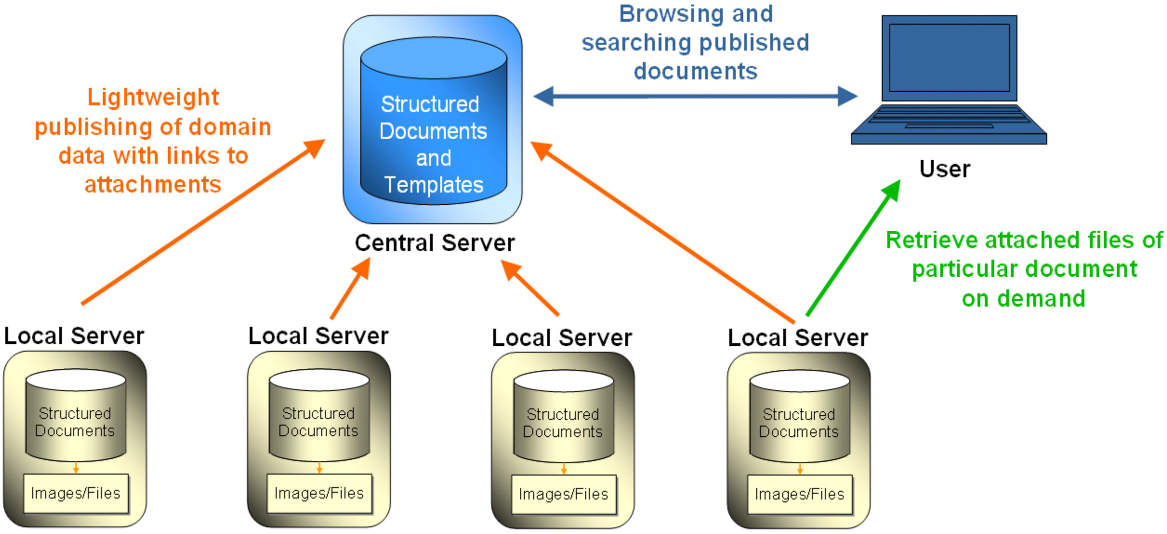

Figure 3-2 Networking architecture of multiple SciPort servers

SciPort is built on scalable and proven database and networking technology. The server component of SciPort can be embedded in a JEE (Java Enterprise Edition) application server

such as JBoss [6]. The server component runs inside the Tomcat servlet engine to communicate with web and stand alone clients via secure networking connection. In addition to web-based clients a standalone Java client is used for robust bulk uploads without user interaction. The use of a scalable servlet engine and application server ensures scalability to large numbers of users and data records. These systems can run on multiple physical machines with automated load balancing.
Multiple SciPort installations can be connected via secure web services. For multi-site studies a local server would be setup for each site to collect and manage the local data. A central server located at the main study site or an additional location would mirror either the lightweight meta-data (use case shown in Figure 3-2) or optionally the complete content. Study organizers can use the central server to track study progress and analyze the data. Access to images and other attached files is automatically rerouted to the originating site if no copy is available on the central server. All communication is done via secure channels using web services. This enables development of additional applications to access the data after proper authentication with a SciPort server.

3.5. Database connection

The connection of databases for the Sim-e-Child project is driven by the project objectives to generate a high quality database for paediatric cardiology by pooling relevant cases from databases in the EU and US. The focus is this on integrating the most recent cases from the COAST trial in the US and OBPG in the EU. In both cases a direct connection to the clinical systems presents logistic challenges in anonymization. Software would need to be placed inside the hospital IT to make sure de-identification is performed before the data leaves the hospital network and would be uploaded to the Sim-e-Child grid. For the software developers this implies that the software cannot be tested with the actual data or in the actual software infrastructure. In addition the IT department of the hospitals would need to approve and operate the software.

For the integration of the source databases it was therefore decided to use a traditional ETL (Extract Transform Load) approach of exporting the data on site, performing de-identification (with the help of software tools provided by Siemens) on site, and writing import modules for SciPort for these formats.
The various SciPort databases will be interconnected via secure web services to allow for example transparent search over the pooled cases from the US and EU. A client library knows of all the servers in the network. A search query is sent via secure webservice to all servers and the results are merged and returned to the client.

Furthermore selected cases will be imported from the Health-e-Child databases as needed for the evaluation of heart models. While the clinical data can be imported automatically from the Health-e-Child database, the medical images will need to be uploaded manually into the Sim-e-Child SciPort database.

4. Conclusion

This document aims to report on both the design and development of the Sim-e-Child (SeC) grid and databases infrastructure.

As described in the first part, the grid infrastructure is based on the gLite grid middleware from the EGI project. This document therefore presented the deployment of the SeC infrastructure as it was deployed at month 10 of the project.
The different Grid components that are used in the project were presented as well as there current organisation. The services were divided in two parts:

- Central services: Those are the core services of the Grid middleware (VOMS, tBDII, LFC, WMS).
- Site services: These later provide concretely the storage capacity and the computing power of the system (CREAM-CE, DPM, WNs, sBDII).

At month 10 of the project, all the central services were installed and configured in order to respect the naming conventions, good practices and to be compliant with the EGI deployment procedures. The central services currently run at MAAT France. As far as the sites are concerned, the European one is deployed and ready for us. It does not currently provide a lot of computing power but it will be improved as soon as the project will need it. The HeC sites institutions will be contacted in order to explain them that a reconfiguration of their node is needed. Finally, as far as the US site is concerned, in order to ease the deployment and to have more flexibility, it was decided to rent the needed servers. Different configurations and providers were contacted and at the tile of writing, the OVH Dedicated servers services seems to be the best choice even if it’s still not available in US but this should be the case during the next year.

The analysis of the data models and access paths of the source databases has led to a plan how to extract, transform, and load data from source systems into the SciPort databases. This includes automated paths and manual entry. The SciPort databases will form the data backbone of the SeC grid infrastructure and host original data, annotations and simulation results.