4th EGEE User Forum/OGF 25
and OGF Europe's
2nd International Event

Le Ciminiere, Catania, Sicily, Italy
2-6 March 2009

OGF-Europe Tutorial

How to make sustainable a grid-enabled e-Infrastructure

by Pasquale Pagano (ISTI-CNR, Italy)

by Pedro Andrade (CERN)

Introduction

- What do we mean by Sustainability ?
- Grid-enabled e-Infrastructure
  - Characteristics
  - Technological complexity and solution

Technological Sustainability

- Quick overview on standards and technology

Operational Sustainability

- Key success factors

The D4Science Infrastructure

- A concrete example

Conclusions

Introduction

Sustainability

Management of any grid-enabled e-Infrastructure is costly in terms of ownership, maintenance, and upgrade.

How to minimize e-Infrastructure deployment costs, operational and application porting timeframes, and manpower requirements.

The compliancy to standards allows to build software framework that:

- Simplifies applications porting to the e-Infrastructure
- Eliminates manual deployment overheads
- Allows remote site management
- Provides powerful mobile monitoring tools

SOI

Is a Virtualised IT Infrastructure which

1. exposes **a catalog of WS** instead of running service instances,
2. supports **SOA Application**, and
3. includes infrastructure resources such as **compute**, **storage**, and **networking** hardware and software (middleware) to support the running of services.

The Technological Requirement

A Software framework

- to support ON-DEMAND virtual collaborations* among remote parties
  - cost-effective, secure,
  - dynamic, both **short** and **long** lived
  - overcome ad-hoc systems alike
- to make discoverable and accessible
  - computing, storage,
  - data and service resources
- to promote and/or contribute to data and service integration

* Research Environment

The Technological Complexity

Software framework

- needs a ‘middleware’ (typically distributed)
- is open by definition
  - new resource types and/or new resource instances can be de/registered at any time
- is powerful if it supports sharing scope (controlled resource sharing)
  - machines, storage, data and services resources
- is powerful if it supports application scope
  - the portion of the infrastructure in which a resource exists
  - the portion of the infrastructure in which a resource can act, operate, or has power or control

gCube Technological Solution

- By relying on gLite, gCube is an e-Infrastructure enabling system to share
  - computing, storage, data and service resources → g3
- gCube system allows collaborations in eScience
  - strongly content-oriented, potentially data and processing intensive
- within the sharing scope of Virtual Organizations
  - broader and longer lived
  - may stretch across the whole infrastructure
  - or else over significant subsets thereof
- take place in Virtual Research Environments (application) scope
  - interactively created, managed, defined, and used:
  - typically short to medium lifespan

Virtual Organization

A Virtual Organization (VO) models sets of users and resources belonging to a e-Infrastructure.

It defines clearly and carefully:

- What is shared
- Who is allowed to share
- The conditions under which sharing occurs, based on an authentication and authorization policies

VOs may have a limited lifetime and they are created to satisfy transient needs of the constituent potentially heterogeneous actual Organizations.

A Virtual Research Environment (VRE) provides a framework of applications, services and data sources dynamically identified to support the underlying processes of research/collaboration/cooperation.

The purpose of a VRE is to help researchers* belonging to a Virtual Organization for carrying out cooperative activities like **data analysis and processing**, **production of new knowledge** using specialized tools.

*Researcher has to be considered in the large, i.e. end-user, decision-makers, resource and data providers, etc.

Technological Sustainability

Challenges and Solutions

Heterogeneous resources: computing, storage, data, service

- Resource representation

Catalog of WS

- Service packaging

Sharing scope (VOs)

- Credentials, roles, and capabilities representation

Application scope (VREs) dynamically defined, configured, staged, and executed

- Enabling VRE environment

Data-centric analysis

- Enabling technology for exploiting heterogeneous storage and file systems

Computer, storage, and Grid components: Glue Schema

- Advantages:
  - the schema is designed to facilitate interoperation between Grid infrastructures
  - is designed in response to existing use cases
- Disadvantages:
  - None

Resource Representation [cont.]

Service: gCube Service Profile

- Advantages
  - Interfaces, build and runtime dependencies, deployment constraints, deployment preferences, software packages
- Disadvantages
  - Is not standard

Data (and metadata): gCube Data Profile

- Advantages
  - Open to include provenance and tailored VO information
- Disadvantages
  - Is not standard

- Heterogeneous resources: computing, storage, data, service
  - Resource representation
- Catalog of WS
  - Service packaging
- Sharing scope (VOs)
  - Credentials, roles, and capabilities representation
- Application scope (VREs) dynamically defined, configured, staged, and executed
  - Enabling VRE environment
- Cooperative activities for carrying out data-centric analysis
  - Enabling technology for exploiting heterogeneous storage and file systems

Packaging: gCube Software Archive (SA) specification

- Advantages
  - support the certification of the content of the SA by exploiting the information of the Service Profile; easily generable automatically as result of a build process, e.g. Etics SA configuration
- Disadvantages
  - Is not standard

Challenges and Solutions

Heterogeneous resources: computing, storage, data, service

- Resource representation

Catalog of WS

- Service packaging

Sharing scope (VOs)

- Credentials, roles, and capabilities representation

Application scope (VREs) dynamically defined, configured, staged, and executed

- Enabling VRE environment

Cooperative activities for carrying out data-centric analysis

- Enabling technology for exploiting heterogeneous storage and file systems

VOMS provides a database to manage user roles and capabilities.

- Advantages
  - Provides a set of tools for accessing and manipulating the database and for generating Grid credentials for users.
  - voms-proxy-init allows users to generate a local proxy credential based on the contents of the VOMS database.
  - gCube (VOMS-aware) applications use the VOMS data to make authentication decisions regarding user requests.
- Disadvantages
  - Expert users

gCube supports non expert users by providing solution for additional authorization frameworks that allow to use gCube either service or VO credentials to consume resources.

Challenges and Solutions

Heterogeneous resources: computing, storage, data, service

- Resource representation

Catalog of WS

- Service packaging

Sharing scope (VOs)

- Credentials, roles, and capabilities representation

Application scope (VREs) dynamically defined, configured, staged, and executed

- Enabling VRE environment

Cooperative activities for carrying out data-centric analysis

- Enabling technology for exploiting heterogeneous storage and file systems

Enabling VRE environment

To support the dynamically definition, configuration, staging, and execution of VREs:

- Standards for service communication
- Common framework tailored for SOI developers
- Common container
- Middleware

Enabling VRE environment: Standards for communication

Java WSCore, Apache Axis, and OGF specifications (and implementation if any):

- WS-Notification, WS-Addressing, WS-Security
- WSRF (Statefull WS)
  - WS-ResourceProperties (WSRF-RP)
  - WS-ResourceLifetime (WSRF-RL)
  - WS-BaseFaults (WSRF-BF)
  - WS-ServiceGroup (WSRF-SG)
- WS-DAI, WS-DAIR, WS-DAIX

Enabling VRE environment: Standards for communication - WSRF [cont.]

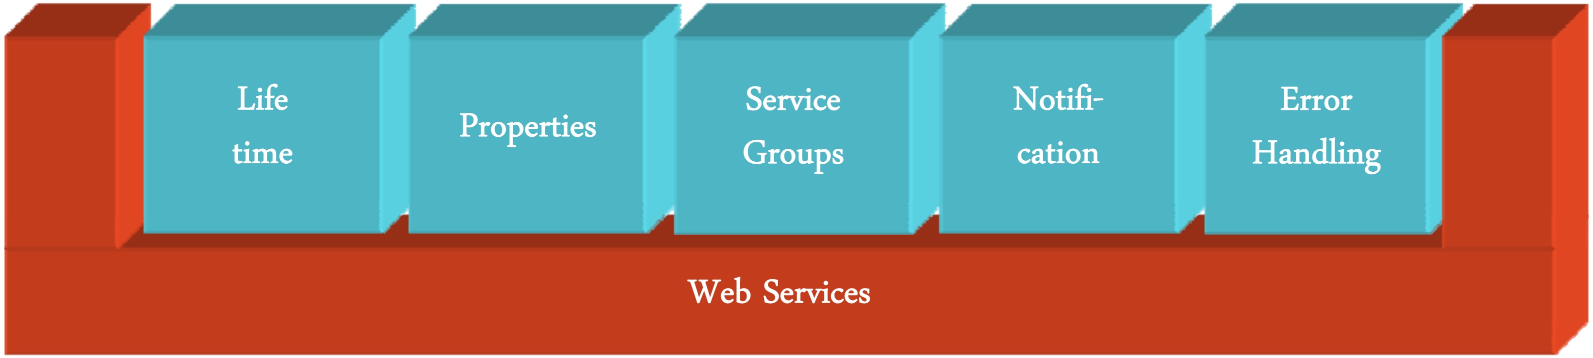

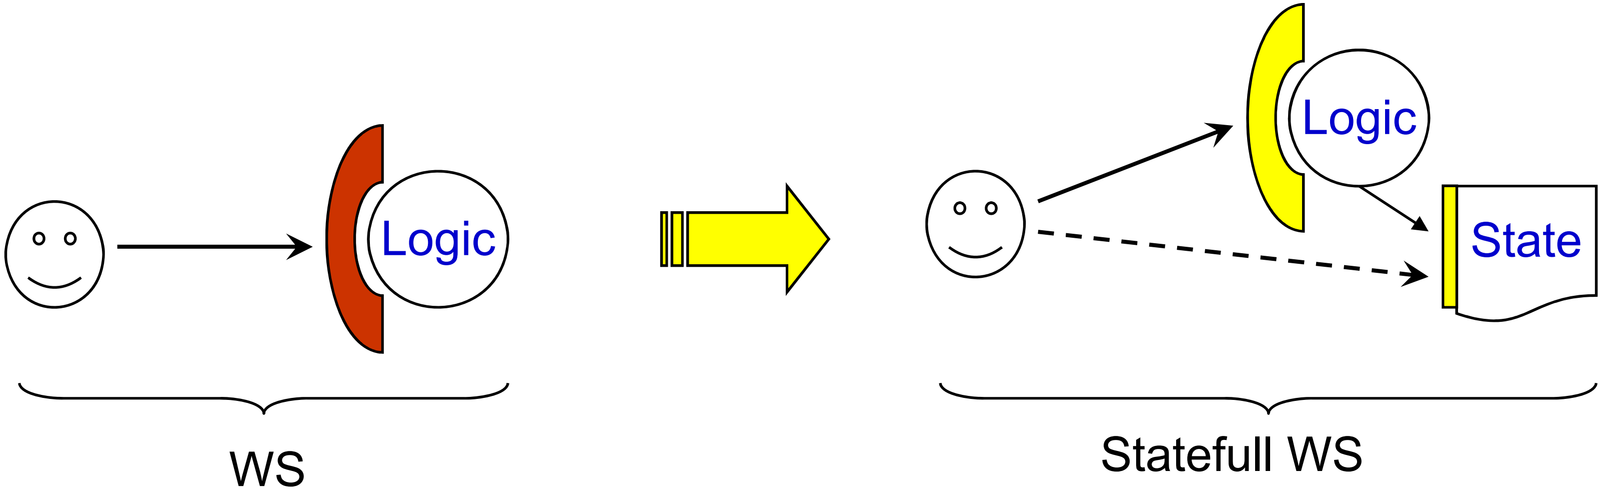

Enabling VRE environment: Standards for communication – WSRF [cont.]

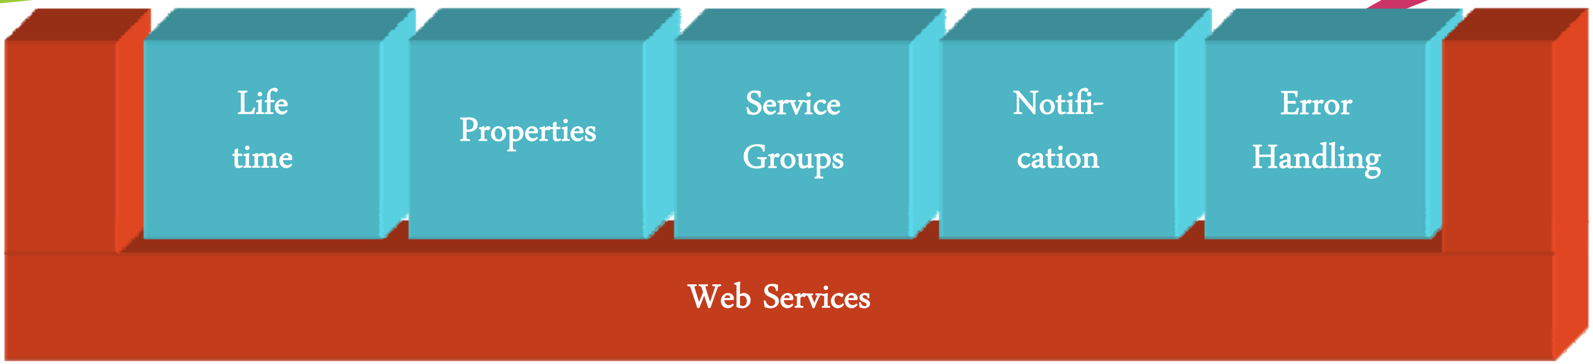

Lifetime (WS-ResourceLifetime)

- Factory based dynamic creation of WS-Resources
- Instances are created with a limited lifetime
- Prevent services from consuming resource indefinitely ("Garbage Collection")

Properties (WS-ResourceProperties)

- Defines type and values of a resource state

Service groups (WS-ServiceGroups)

- Describes an interface for operating on collections of WS-Resources
- E.g. to distribute an action to a set of services

Notification (WS-Notification)

- Notification about state changes
- Applies traditional publish/subscribe paradigm

Error Handling (WS-BaseFaults)

- Defines base handling of communication errors

Enabling VRE environment: Standards for communication - WSRF [cont.]

Advantages:

- Clear Description of Resources and Interfaces
- Dynamic sharing of resources
- On-demand services exploitation
- Cross-organizations trusted environment
- Widely accepted Web Service standards

Disadvantages:

- Reference implementations are not well documented and complex to use
- Several complementing specifications are in development
- Complex middleware requires maintenance and administration overhead

Common framework tailored for SOI developers: gCore framework

gCore Framework (gCF) overcomes the complexities of the design and implementation of SOI compliant services by

- Managing the entire lifecycle of the services
- Implementing WSRF standards for publication, access, and notification of change to service state
- Standardizing the use of systemic faults
- Mediating access to configuration resources on the local file system
- …

Standard framework tailored for SOI developers: gCore Franework [cont.]

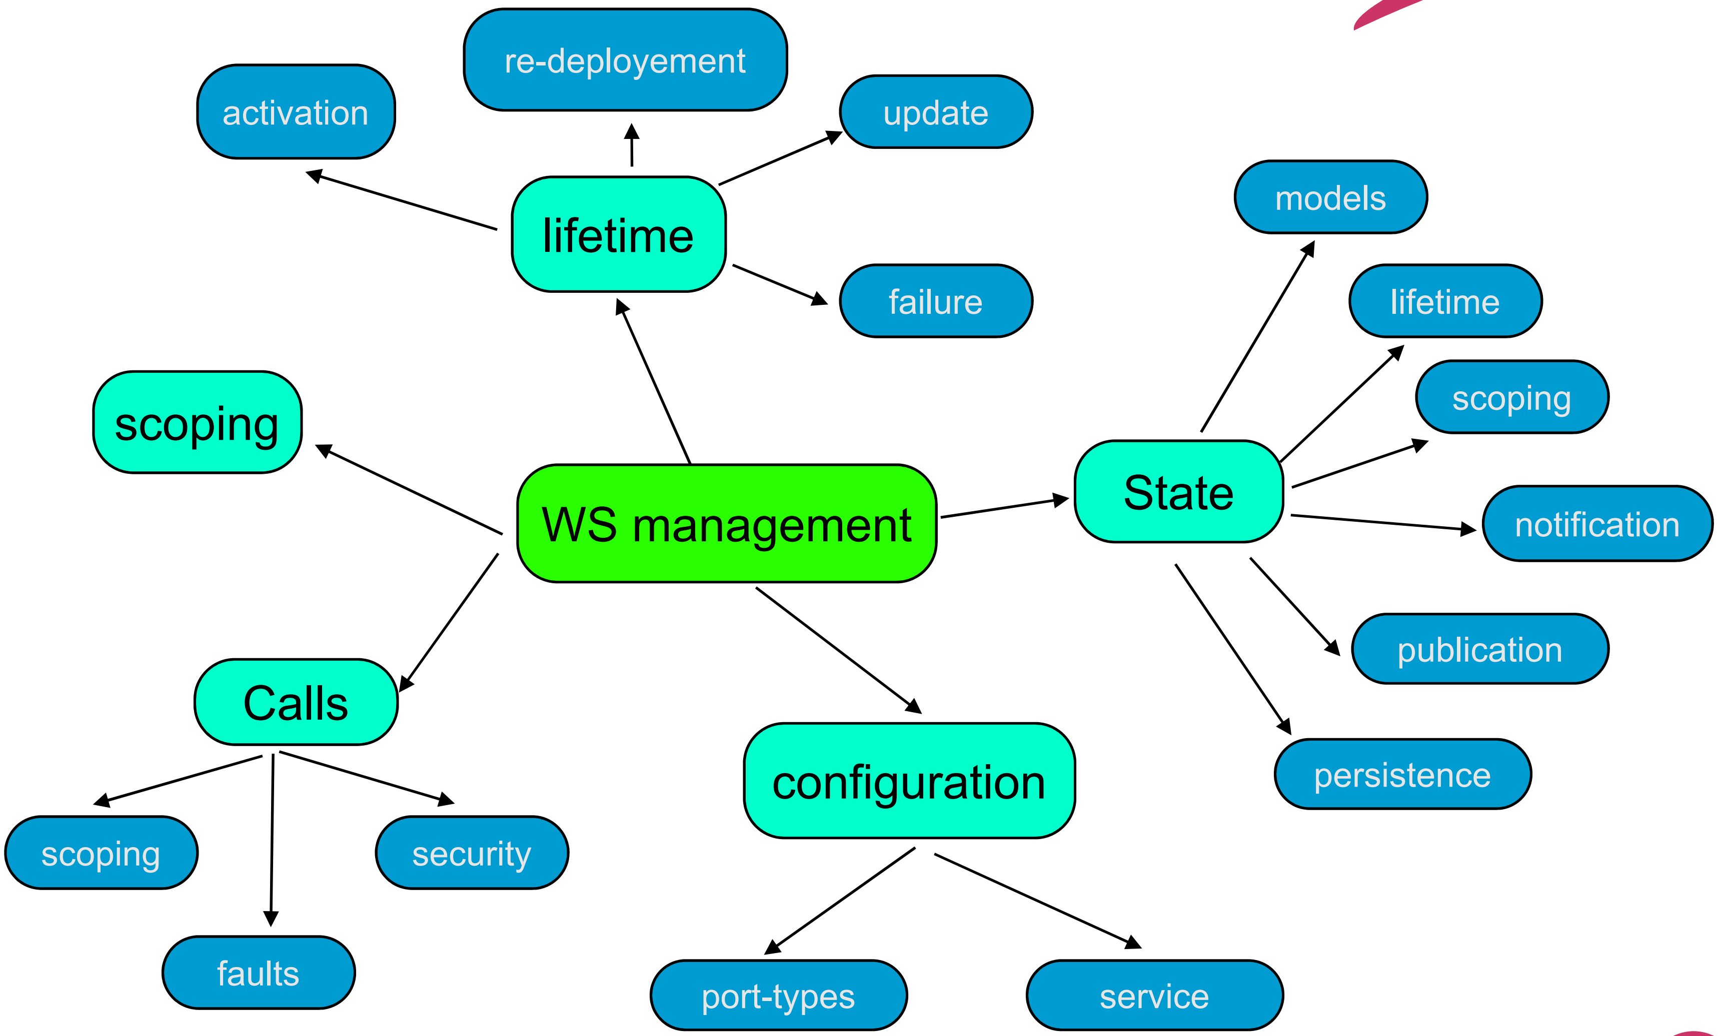

Common container: gCube Hosting Node

gCube Hosting Node (gHN) meets the needs of site administrators, infrastructure managers, and resource providers by providing an **easy to install self-contained runtime container**:

- bound to logical ports
- provide access to local hardware resources
  - storage, system resources, CPU cycles...
- grant lifetime management
- mediate service2service interactions
  - route requests to target service
  - enforce security and scope policies

Software Repository service

- Storage and provision of deployable software archives (SAs)
- Software Archive certification
- Dependencies resolution

VRE Management services

- Deployer: deploy certified SAs locally
- gHN-Manager: manage the gHN container
- VRE-Manager: coordinate the deployment of VRE resources
- Broker & Matchmaker: discover gHNs suitable for the deployment of SAs

Information System services

- publication of resources profile
- discovery of VRE resources through xPath, XQuery
- real-time monitoring
- subscription/notification

Challenges and Solutions

Many types of resources: computing, storage, data and service
- Resource representation

Catalog of WS
- Service packaging

Sharing scope (VOs)
- Credentials, roles, and capabilities representation

Application scope (VREs) dynamically defined, configured, staged, and executed
- Enabling VRE environment

Cooperative activities for carrying out data-centric analysis
- Enabling technology for exploiting heterogeneous storage and file systems

SRM provides a complete interface to heterogeneous storage systems

- Advantages:
  - Consistent homogeneous interface to the Grid, while allowing sites to have diverse infrastructures.
  - Born taking into consideration actual use cases and influenced by needs of the large High Energy Physics communities
  - Capabilities: directory and ACL, non-interference with local policies, space reservation and management, abort, suspend, and resume operations, transfer protocol negotiation, …
- Disadvantages:
  - None

gCube embeds SRM (and GridFTP) support by providing an higher interface allowing to store, discovery, and access qualified network of files (**compound objects**).

gCube & Standardization Bodies

- ISO: data representation (e.g. ISO3166 for countries, ISO4217 for currencies) and metadata (ISO19115 for GIS)
- OGF: Standards related to Architecture (e.g. OGSA), Data (e.g. DAIS, SRM, GridFTP), Management (e.g. GLUE, Resources Usage), Applications (e.g. DRMAA), Compute (e.g. JSDL)
- OAI: Resources Exposure/Harvesting (OAI-PMH) Resources Exchange (OAI-ORE)
- OASIS: Standards related to stateful web services (e.g. WSRF), process management (BPEL), remote user interfaces (WSRP), A&A (SAML / XACML)
- W3C: All the standards related to Web Architecture (e.g., URI/URL, HTTP), Service Oriented Architectures (e.g. SOAP, WSDL, WS-Addressing) and data representation and manipulation(e.g. XML*)
- Others: Classification systems (e.g. ISSCAAP, ISSCFV, ISSCFG), features representation (e.g. GML for GIS), metadata (e.g. AgMES for Agricultural, SDMX for Statistics)

Operational Sustainability

From an infrastructure perspective the sustainability of its operation heavily depends on:

- The deployment of services compliant with standards that provide cost-effective solutions for the operation of the infrastructure
- The establishment of links with other infrastructures and standards initiatives to understand and experiment resource sharing possibilities

The operation of a gCube-based infrastructure in a cost-effective way is based on a number of key factors:

- Common representation of **infrastructure resources**
- Structured and certified **infrastructure services**
- Simplified execution of **infrastructure deployment**
- Clear organization of **infrastructure roles**
- Adoption of standard **infrastructure procedures**

Infrastructure Resources

A gCube based infrastructure is composed by:

Hardware:

- Storage, Computing
- gCube Container

Services & Applications:

- gCube Services
- External software

Collections & Auxiliary Resources:

- Data, Metadata, Indexes, Annotations
- Schemas, Mappings, Transformation programs

The management of such a diverse set of resources in a distributed environment is only possible by following a common approach to query, access, and exploit these resources.

The gCube Information System (IS) collects these information in xml-based resources profiles. For example:

- gLite services: from glue schema
- gCube services: from service profile

Infrastructure Resources

The aggregation of information about all infrastructure resources under a common resource profile enables to creation of advanced monitoring tools.

These tools can provide a wide range of interfaces, from a flat representation of the infrastructure resources to specialized views for concrete infrastructure management aspects:

- VRE view: gathers information about which datasets and services are allocated to one VRE
- Package view: gathers information about the availability of the gCube packages in the infrastructure
- http://monitor.d4science.research-infrastructures.eu/iv/

A gCube infrastructure only runs certified software.

This certification is possible due to the packaging of all gCube software as Service Archives (SAs). These SAs:

- Follow a predefined structure
- Group inter-connected packages
- Declare dependencies to other packages

All SAs are stored in a common repository, the gCube Software Repository. The storage of any gCube SA in this repository is subject to a certification process.

The deployment and upgrade of the gCube software is simplified with the exploitation of gCube infrastructure enabling services.

All services to be deployed are automatically retrieved from the gCube Software Repository hosting certified Service Archives.

The deployment and upgrade operations are done remotely using the gCube VRE Manager service without the intervention of local site administrators.

- VO and VRE services remotely deployed. Few services manually deployed
- Implements automatic dependencies resolution
- Optimizes the selection of hosting nodes based on SAs requirements and node characteristics

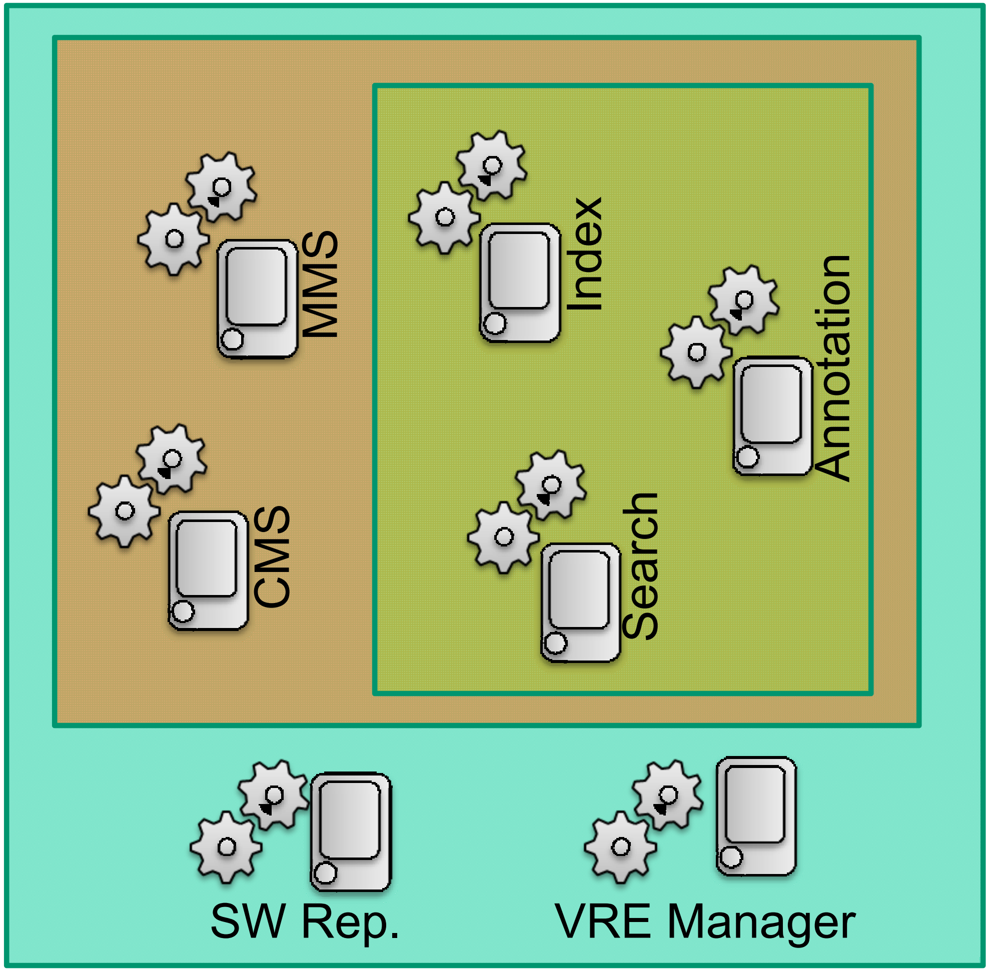

Infrastructure Roles

The operation of a gCube-based infrastructure requires the execution of several tasks. Such tasks are associated to one or more roles managed through VOMS.

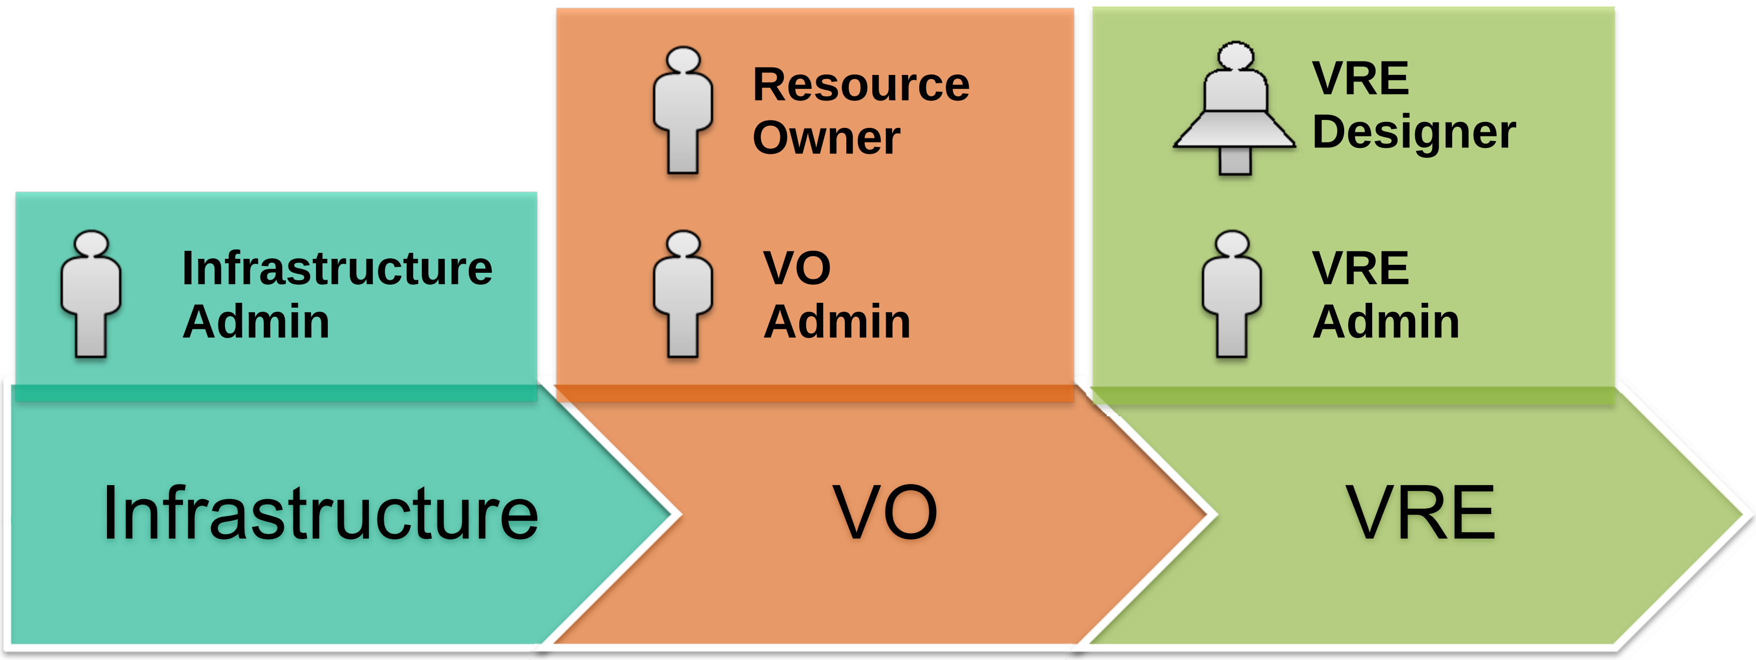

Infrastructure Procedures

The operation of a gCube-based infrastructure is based on well defined set of procedures for infrastructure deployment, monitoring, support, security, etc.

The collaboration with other infrastructures and the adoption of a common set of procedures promotes the interoperability between infrastructures.

The D4Science Infrastructure

VOs and VREs

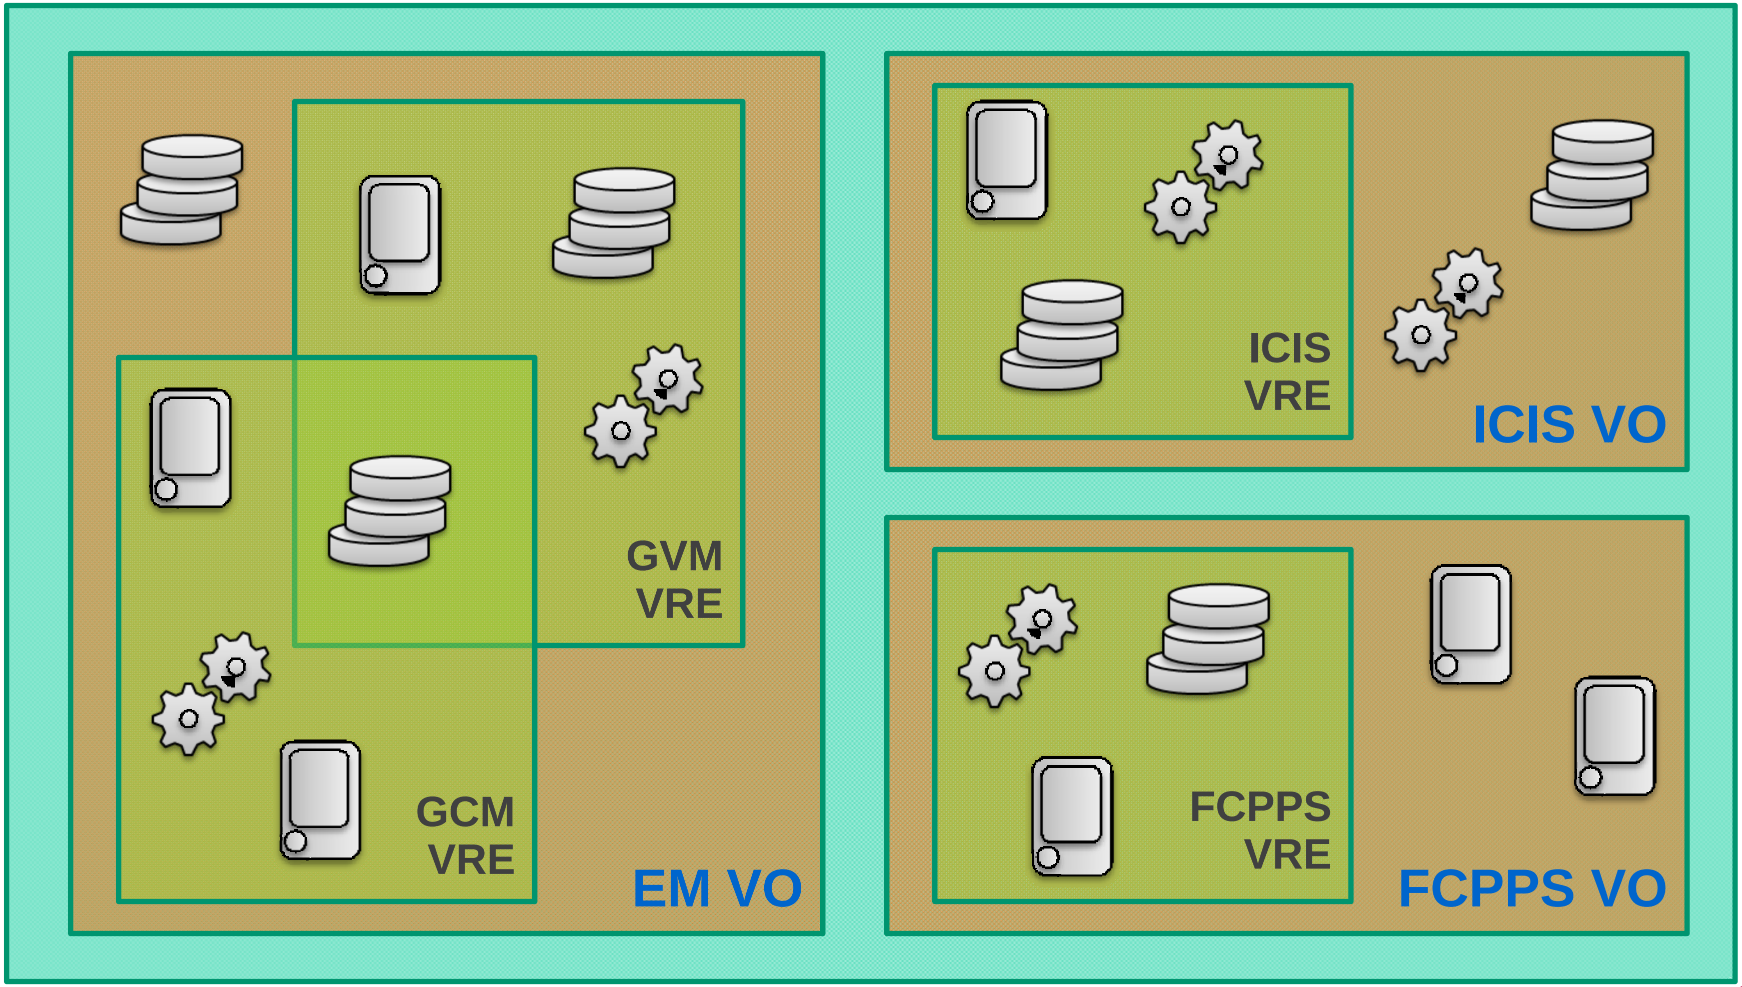

Sites

Composed by 4 sites:

- CNR (Pisa, Italy)
- ESA (Rome, Italy)
- NKUA (Athens, Greece)
- UNIBASEL (Basel, Switzerland)

Providing:

- 41 machines, 87 processor cores, 26 TB disk

Access to external EGEE infrastructure sites

Effort

SCOPE	ACTION	USER	TOOL
Infrastructure	deploy enabling services	Inf. Admin	CLI
	deploy portal	Inf. Admin	CLI
VO	deploy gHNs	Site Manager	CLI
	deploy VO services	VO Admin	Portal
	data import (metadata, indexes)	VO Admin	Portal
	register resources (gHNs, services)	VO Admin	Portal
	manage users and roles	VO Admin	Portal
VRE	define VRE functionality and data	VRE Designer	Portal
	deploy VRE services and layout	VRE Admin	Portal
	modify VRE	VRE Admin	Portal

Effort

SCOPE	ACTION	TIME	
Infrastructure	deploy enabling services	1 day	one day
	deploy portal	1 hour	
VO	deploy gHNs	1 min	few hours
	deploy VO services	10 min	
	data import (metadata, indexes)	hours	
	register resources (gHNs, services)	10 min	
	manage users and roles	10 min	
VRE	define VRE functionality and data	10 min	few minutes
	deploy VRE services and layout	30 min	
	modify VRE	30 min	

Effort

Clear separation of infrastructure roles:

- Technical: Site Managers, VO-Admin, etc
- Non-technical: VRE Designer

Different tools for different infrastructure management aspects.

A new D4Science VO can be created in few hours:

- Depending on the number of hosts and datasets size

Existing VOs can on-the-fly create new Virtual Research Environments or update existing ones.

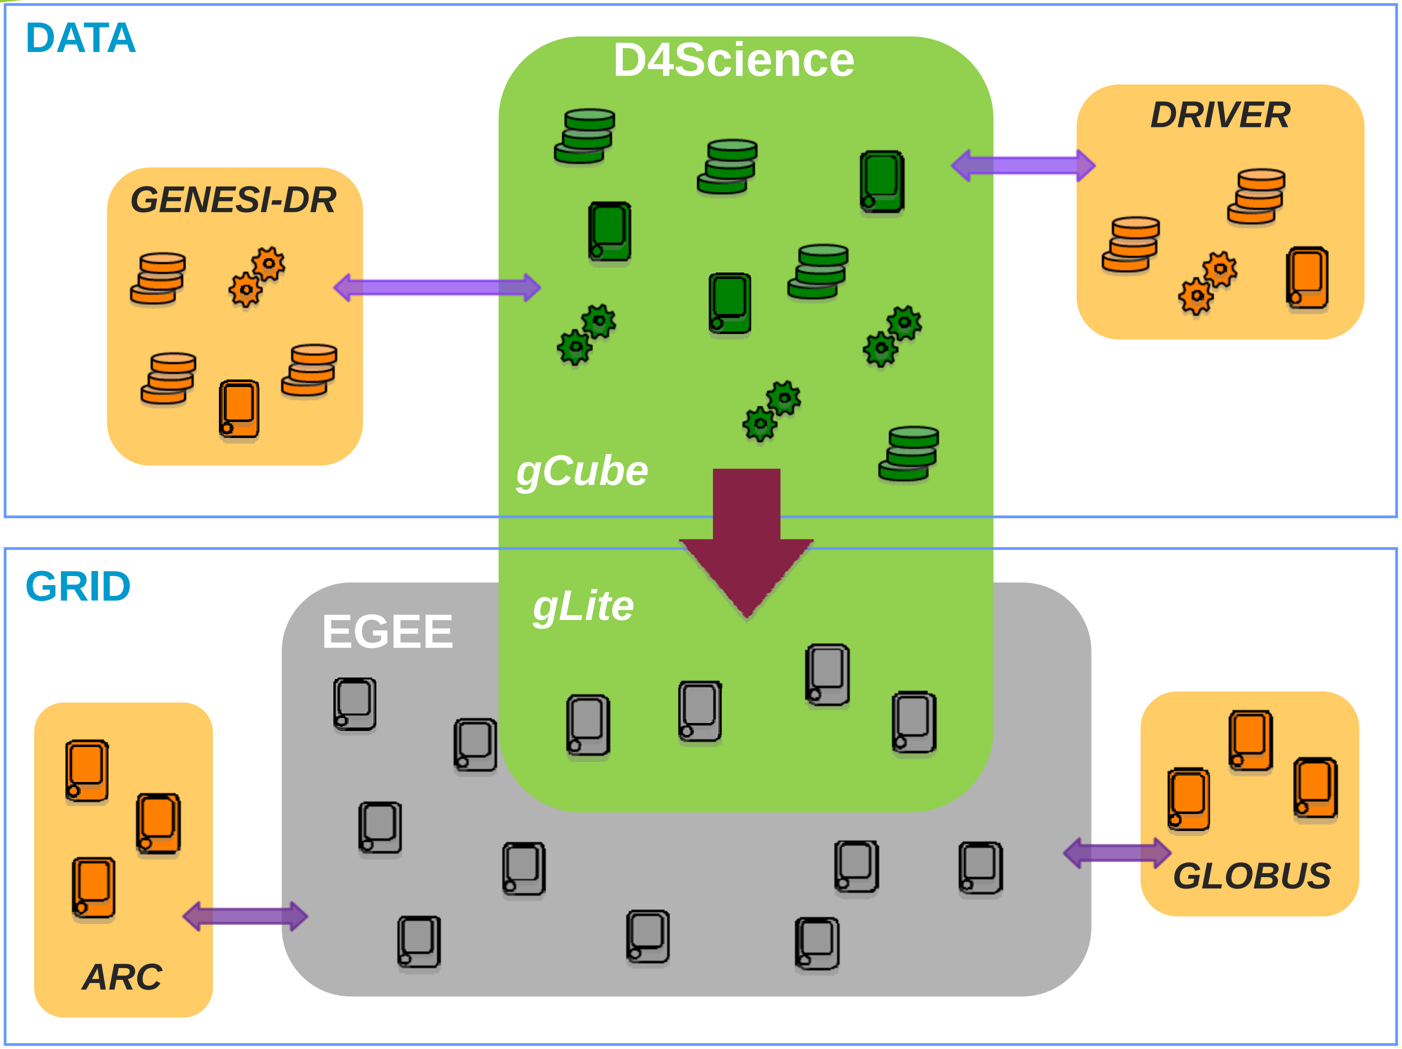

Communication:
WS, WSRF, WS-DAI, WS-N, …
Data:
OAI-ORE LG, OAI-PMH, DCMI, …

Security:
X.509, VOMS
IS:
Glue XML
Job Mgt:
JSDL
Data:
SRM, GridFTP

Conclusions

Delivering a Service Oriented Infrastructure opened for extension (addition of new services) and reuse (resource access by authorized consumers) is mandatory.

- Adoption of standard specifications and reference implementations if any
- Definition of specifications
- Provision of an application framework
- Provision of a SOI middleware

Operational Sustainability Wrap Up

The operation of a gCube-based infrastructure is based on the exploitation of gCube standard-complaint infrastructure management capabilities:

- Access infrastructure information in a uniform way
- Remote deployment of certified software
- Clear management of infrastructure users

➡ Efficient execution of operational procedures

➡ Reduction of the manpower and time

OpenGridForum Europe

THANK YOU !

QUESTIONS ?

CAPACITIES