Towards Next Generations of Software for Distributed Infrastructures

Supporting e-Science in Distributed Systems

EMI Members & Morris Riedel et al.
Juelich Supercomputing Centre
EMI Strategic Director

2012-10-10




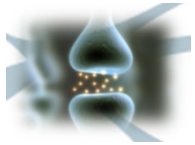
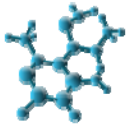

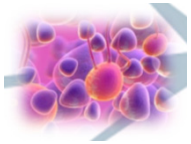
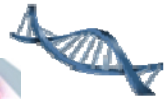

EMI is partially funded by the European Commission under Grant Agreement RI-261611

Contents

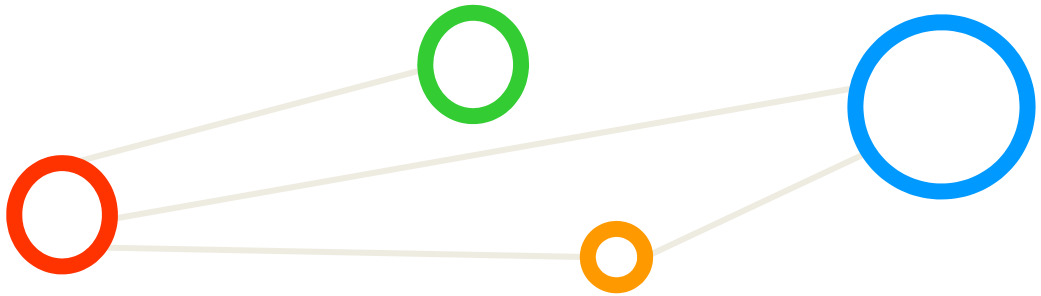

EMI INFSO-RI-261611

Contents

- e-Science with e-Infrastructures & EMI
- Selected e-Science EMI Product Use Cases
- A “Business Use Case”
- Lessons Learned
- References

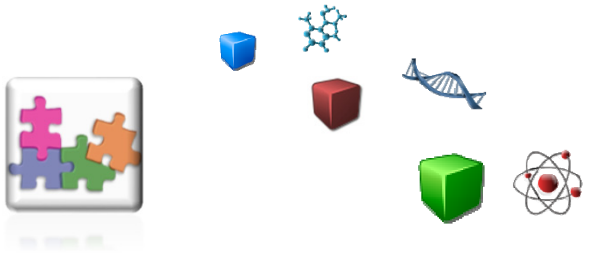

e-Science:e-Infrastructures & EMI

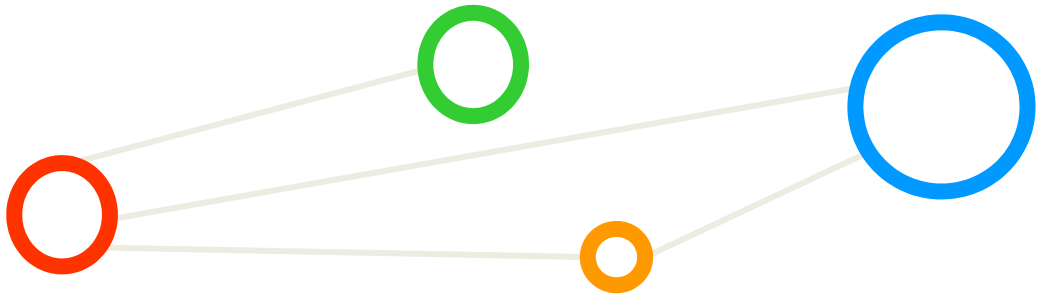

EMI INFSO-RI-261611

The Definition of ‘e-Science’

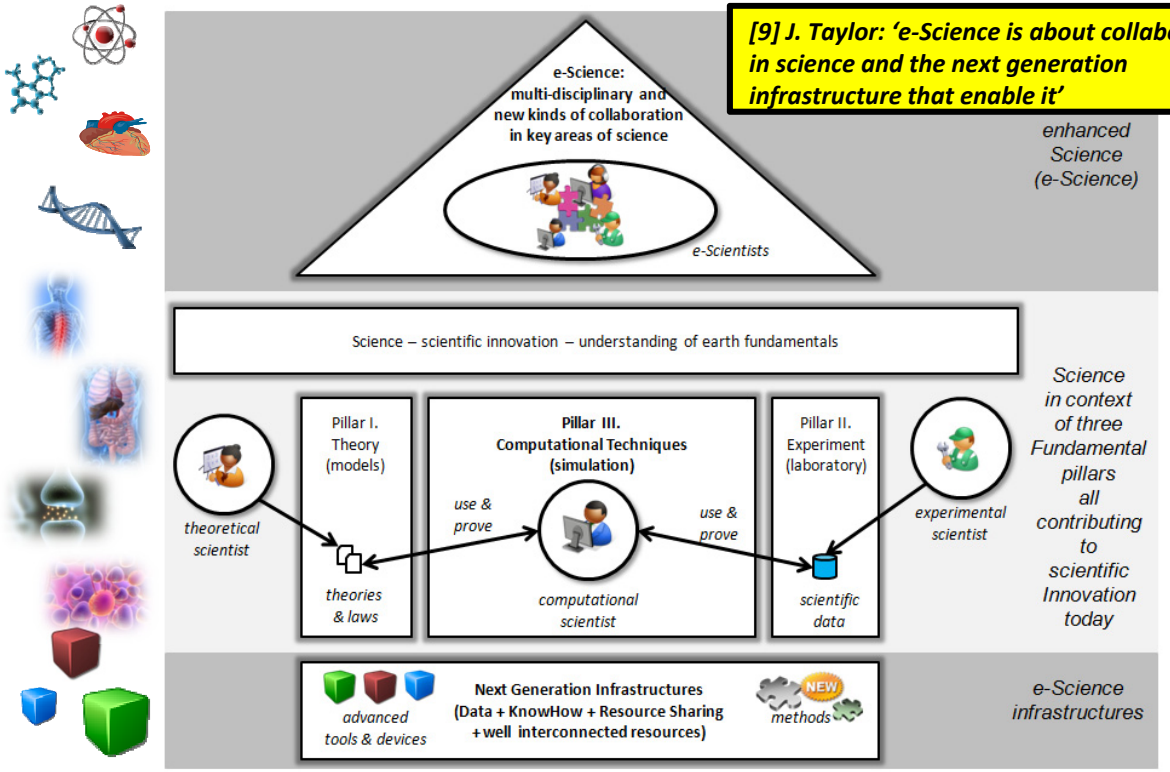

[9] J. Taylor: ‘e-Science is about collaboration in science and the next generation infrastructure that enable it’

e-Science Application Enabling

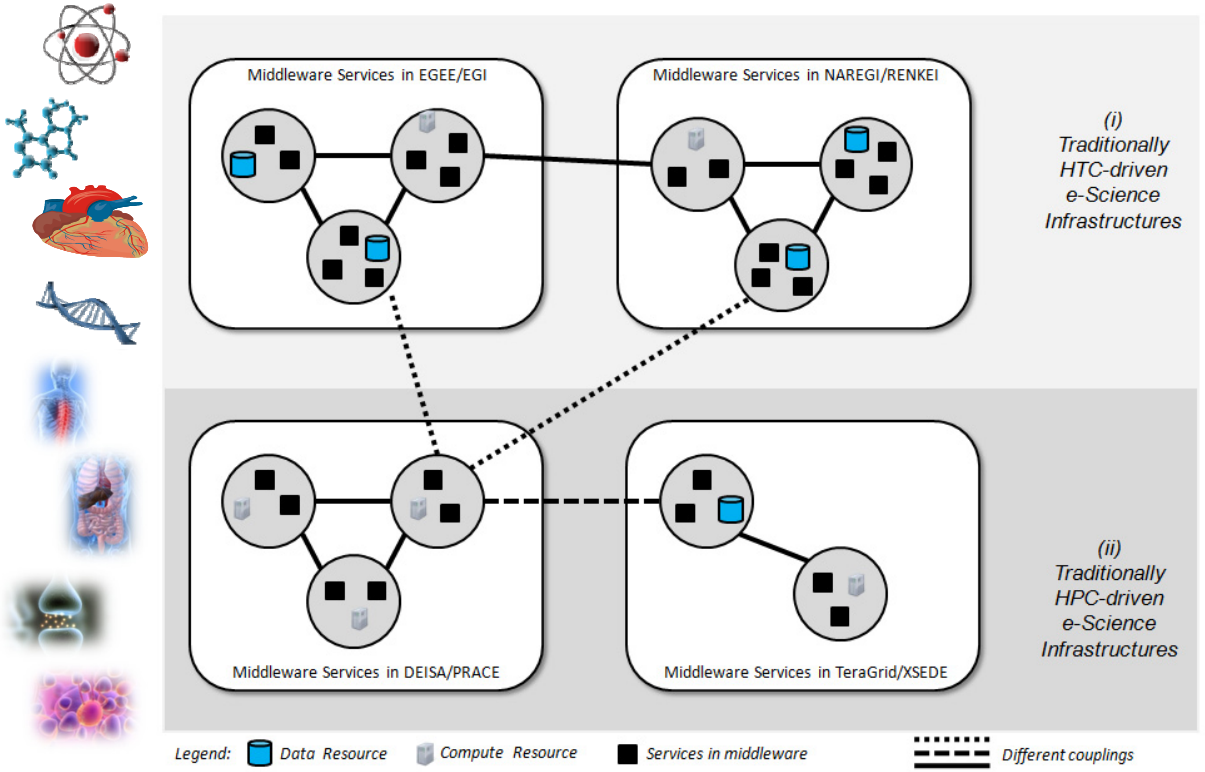

EMI INFSO-RI-261611

EMI & European Policy

[1] Digital Agenda for Europe

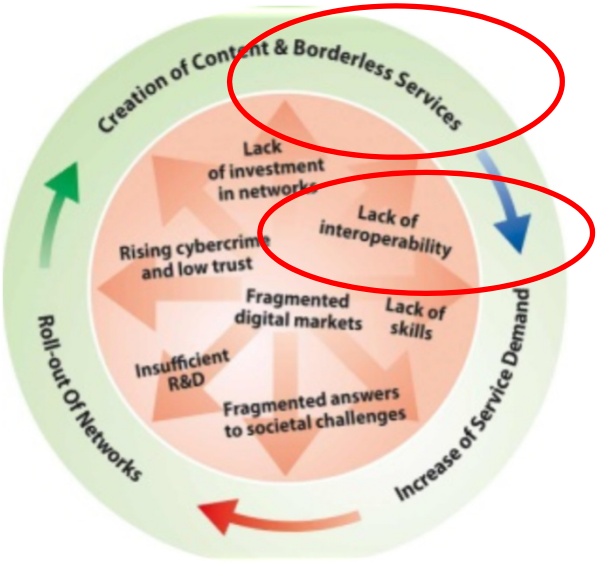

PRACE EUDAT EGI

XSEDE Open Science Grid

EMI INFSO-RI-261611

e-Science Application Enabling

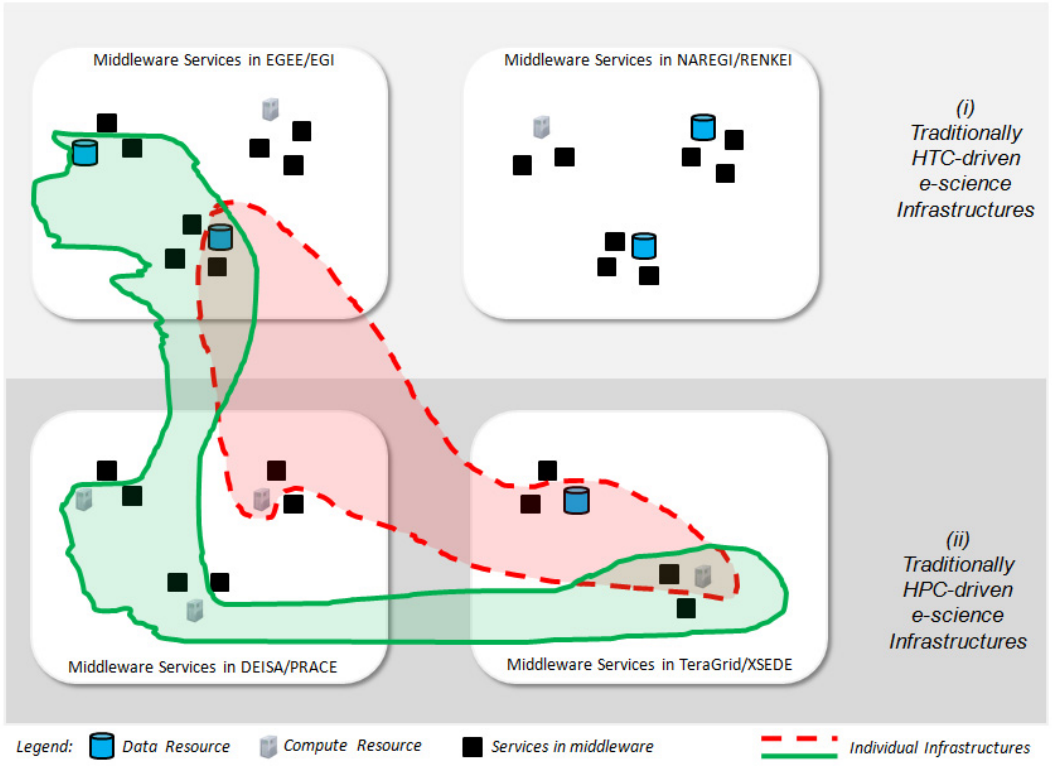

EMI INFSO-RI-261611

EMI 101 – Bundling Expertise

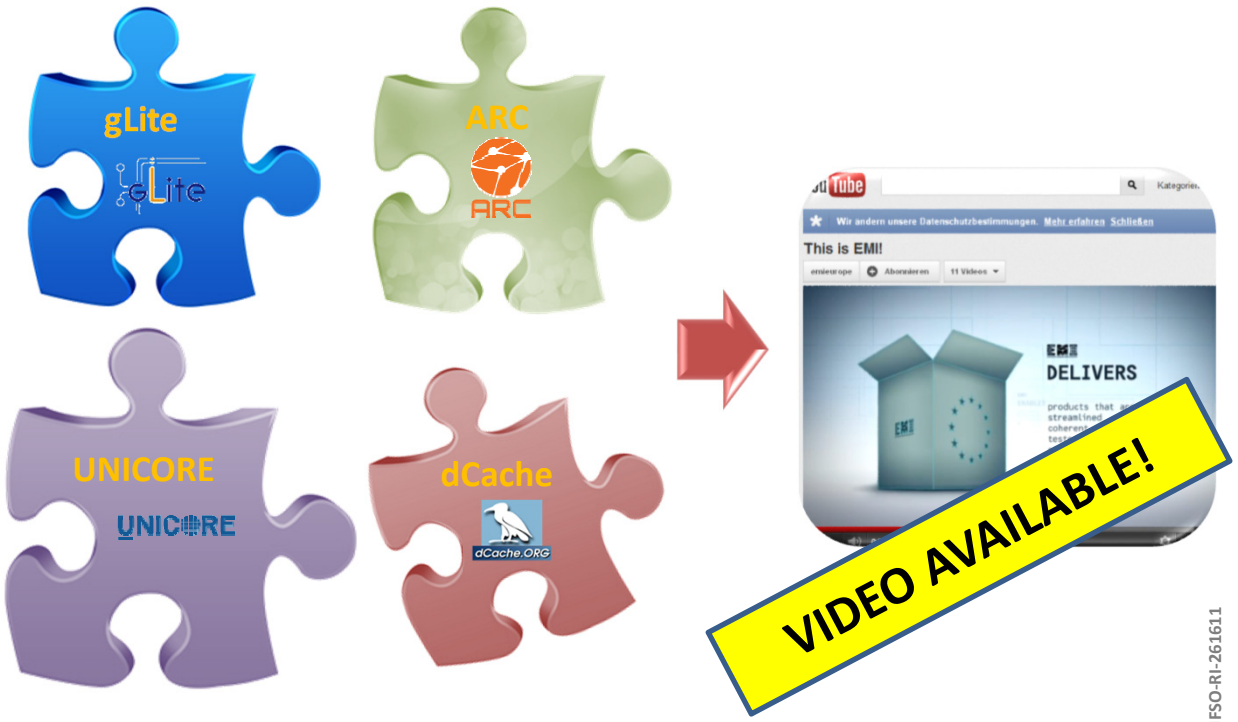

EMI INFSO-RI-261611

10.000 ft Perspective

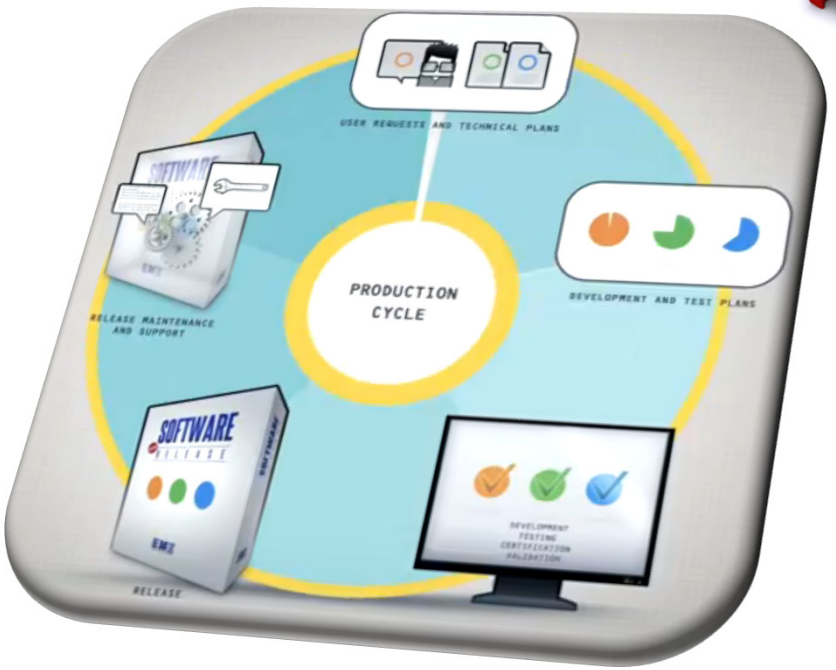

other scientific software tools

EMI INFSO-RI-261611

Present Key Achievements

- Middleware jointly developed & maintained
- Release process harmonized with policies
- Open Standards adoption increased & refined
- Implement several ways for sustainability

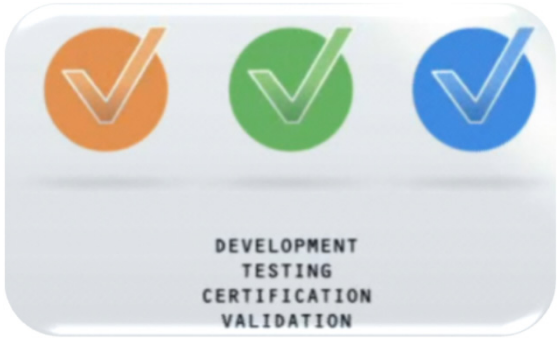

Products for (Distributed) Science

	Name	Leader	Products	Main Areas
1	AMGA	Soonwook Hwang	AMGA, AMGA manager	Data
2	APEL client	John Gordon	APEL parsers and publishers	Infrastructure
3	ARC CE	David Cameron	A-REX, gridftp jobplugin interface, Infoproviders, CE-Cache, CE-staging, LRMS modules, Janitor	All
4	CERN Data	Oliver Keeble	FTS, DPM, LFC, GFAL/lcg_util	Data
5	StoRM	Michele Dibenedetto	StoRM SE	Data
6	Logging and Bookkeeping	Zdenek Sustr	L&B server and clients	Compute
7	VOMS	Andrea Ceccanti	VOMS, VOMS-Admin	Security
8	SAGA-SD-RAL	Antony Wilson	RAL-SAGA-SD including SAGA-SD and SAGA-ISN	Infrastructure
9	UNICORE Security	Krzysztof Benedyczak	UNICORE Gateway, XUUDB, UVOS,	Security
10	UNICORE Container	Bernd Schuller	UNICORE Services Environment (USE) including WSRFLite, Security Libs, XACML PDP, AIP	Security and Infrastructure
11	UNICORE Services	Bernd Schuller	TSI, Registry, UNICORE/X including XNJS, UAS-C, U-BES, U-EMIEX, U-CIP, UAS-D	Compute, Data and Infrastructure
12	UNICORE Clients	Björn Hagemeier	UCC, UNICORE internal client libs, HILA	Compute and Data
13	EMIR	Shiraz Memon	EMI Service Registry	Infrastructure
14	dCache	Christian Bernardt	dcache (server and clients)	Data
15	CERN Data	Oliver Keeble	FTS, DPM, LFC, GFAL/lcg_util	Data
16	StoRM	Michele Dibenedetto	StoRM SE	Data
17	Cesnet Security	Zdenek Sustr	gridsite, gLite-gsoap/gss, gLite-proxyrenewal	Security
18	Logging and Bookkeeping	Zdenek Sustr	L&B server and clients	Compute
19	VOMS	Andrea Ceccanti	VOMS, VOMS-Admin	Security
20	SAGA-SD-RAL	Antony Wilson	RAL-SAGA-SD including SAGA-SD and SAGA-ISN	Infrastructure
21	UNICORE Security	Krzysztof Benedyczak	UNICORE Gateway, XUUDB, UVOS,	Security
22	UNICORE Container	Bernd Schuller	UNICORE Services Environment (USE) including WSRFLite, Security Libs, XACML PDP, AIP	Security and Infrastructure
23	UNICORE Services	Bernd Schuller	TSI, Registry, UNICORE/X including XNJS, UAS-C, U-BES, U-EMIEX, U-CIP, UAS-D	Compute, Data and Infrastructure
24	UNICORE Clients	Björn Hagemeier	UCC, UNICORE internal client libs, HILA	Compute and Data
25	EMIR	Shiraz Memon	EMI Service Registry	Infrastructure
26	EMI Messaging	Lionel Cons	EMI Messaging layer	Infrastructure
27	EMI Common	Cristina Aiftimiei	EMI-UI, EMI-WN, gLite-yaim-core, Torque server config, Torque WN config, emi-nagios	All
28	EMI canl	Zdenek Sustr	Common Library for Authentication	Security
29	WMS	Marco Cecchi	Workload Management Service	Compute

ARC CE Added Value

- Used to submit and manage a wide range of applications running on computational resources of DCIs

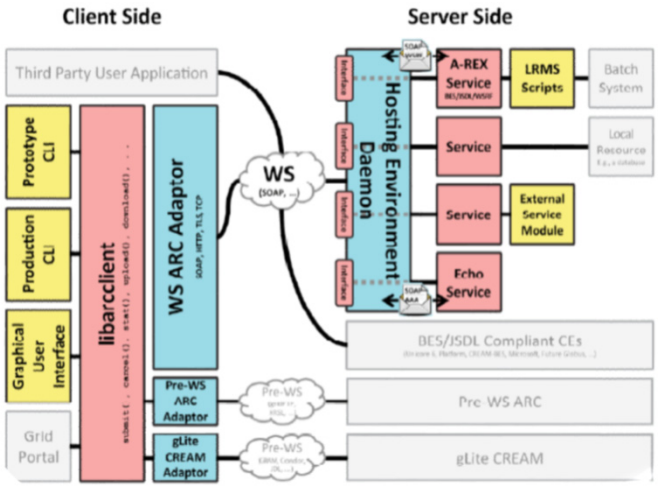

ARC CE Features

- **Use the ARC CE as a light-weight system to execute applications across geographically distributed computing services and their underlying resources**
- Take advantage of a client/server architecture that implements the functionality of a Computing Element (CE) accessing a wide variety of available batch systems
- Interoperate with other EMI services by using the EMI Execution Service via a SOAP-based Web service Interface

EMI INFSO-RI-261611

CREAM CE Added Value

- Used to submit and manage applications running on DCI resources

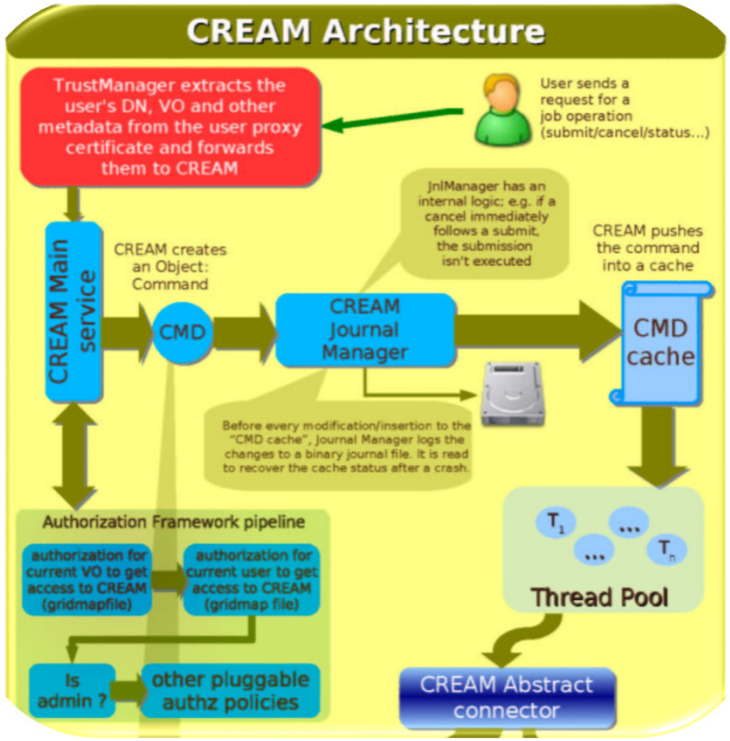

EMI INFSO-RI-261611

CREAM CE Features

- Use the CREAM CE as a powerful system to execute applications across geographically distributed computing services and their underlying resources
- **Take advantage of a client/server architecture that implements the functionality of a Computing Element (CE) accessing a wide variety of available batch systems**
- Interoperate with other EMI services by using the EMI Execution Service via a SOAP-based Web service Interface
- A C++ based Command Line Interface (CLI) is available and other clients can be easily created
- CREAM provides hooks for accounting and offers data-staging functionality

UNICORE Added Value

- Used to submit and manage applications optimized for HPC

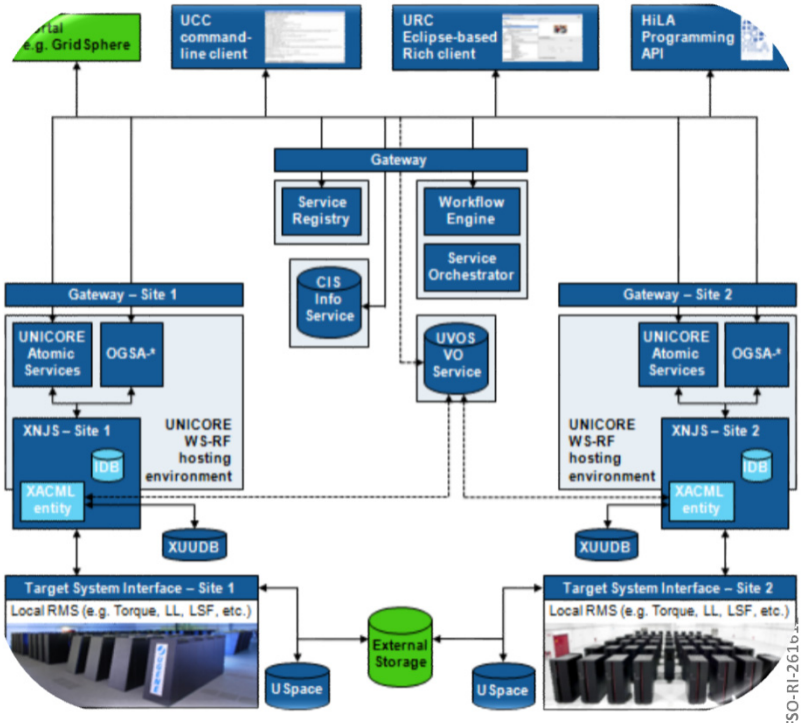

UNICORE Features

- Use the UNICORE system as a powerful system to execute applications across geographically distributed computing services and their underlying resources
- Take advantage of a three-tier architecture that implements the functionality of a Computing Element (CE) accessing a wide variety (~13) of available batch systems
- **Benefit from the maturity and reliability of accessing medium and large-scale HPC resources with key characteristics since ~15 years**
- Deploy a solution that is specifically optimized for sensitive security environments that have less impact on site security policies

dCache Added Value

- Used to store data in a distributed fashion without end-users being aware where their data is stored (‘transparency’)

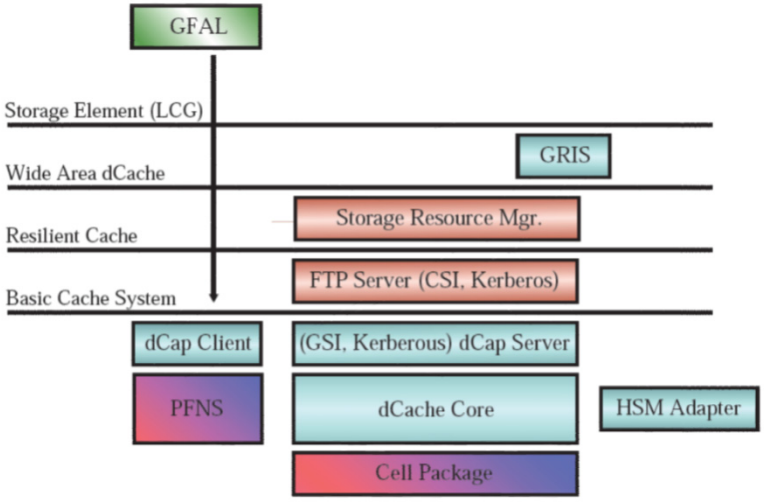

EMI INFSO-RI-261611

dCache Features

- Use the dCache SE as a service in order to transparently provide access to disk-based
- Storage systems as well as tertiary storage (e.g. tapes) known for better cost-efficiency
- Take advantage of a strong client/server architecture that implements the functionality of a Storage Element (SE) offering a variety of access protocols (e.g. POSIX, etc.)
- **Migrate data from one resource to another without affecting end-users**
- Interoperate with other EMI storage services by using the Storage Resource Manager (SRM) 2.2 standard as Web service Interface or the HTTP-based WebDAV standard

StoRM Added Value

- Used to store data and information in different underlying disk-based storage systems
  - One standard interface: SRM

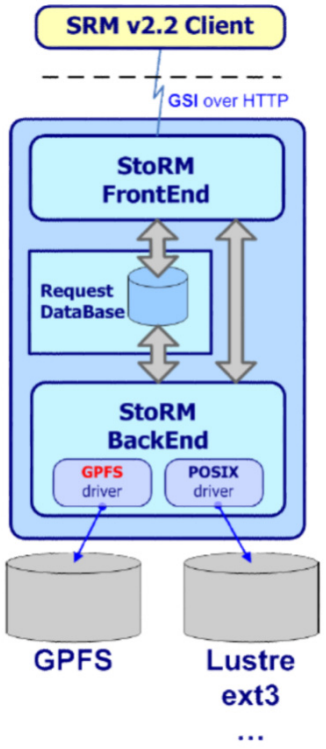

StoRM Features

- Use the StoRM SE as a service that is specifically optimized for (parallel) disk-based storage systems such as the General Parallel File System (GPFS) or Lustre
- Take advantage of a strong client/server architecture that implements the functionality of a Storage Element (SE) offering a variety of access protocols (e.g. POSIX, etc.)
- Provide a stable storage interface with StoRM to end-users while the underlying file system and/or storage system might change over time
- Interoperate with other EMI storage services by using the Storage Resource Manager Open Grid Forum (SRM) 2.2 standard as Web service Interface
- **The modular architecture decouples StoRM from the different underlying file systems**

DPM Added Value

- Lightweight storage solution for DCI sites offering a simple way to create disk-based Grid storage elements and their management

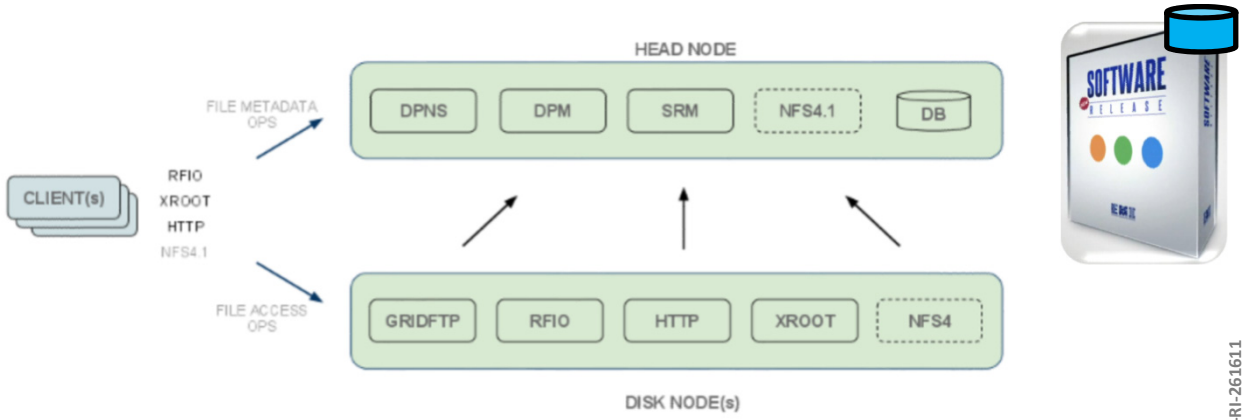

EMI INFSO-RI-261611

DPM Features

- **Use the DPM SE as a lightweight service in order to transparently provide access to disk-based storage systems (easy installation)**
- Install a client/server architecture that support many protocols for file access such as Remote File Input/Output (RFIO), XROOT, HTTP, GridFTP, and NSF4.1
- Interoperate with other EMI storage services by using the Storage Resource Manager (SRM) 2.2 standard
- Take advantage of a system focuses on manageability such as ease of installation and configuration as well as low effort of maintenance
- Leverage all the required functionality for your grid storage solution including support for multiple disk server nodes, different space types or multiple file replicas in disk pools

EMIR Added Value

- Provide high robustness, scalability and performance registry using a federated model (with no centralized, single entry point)

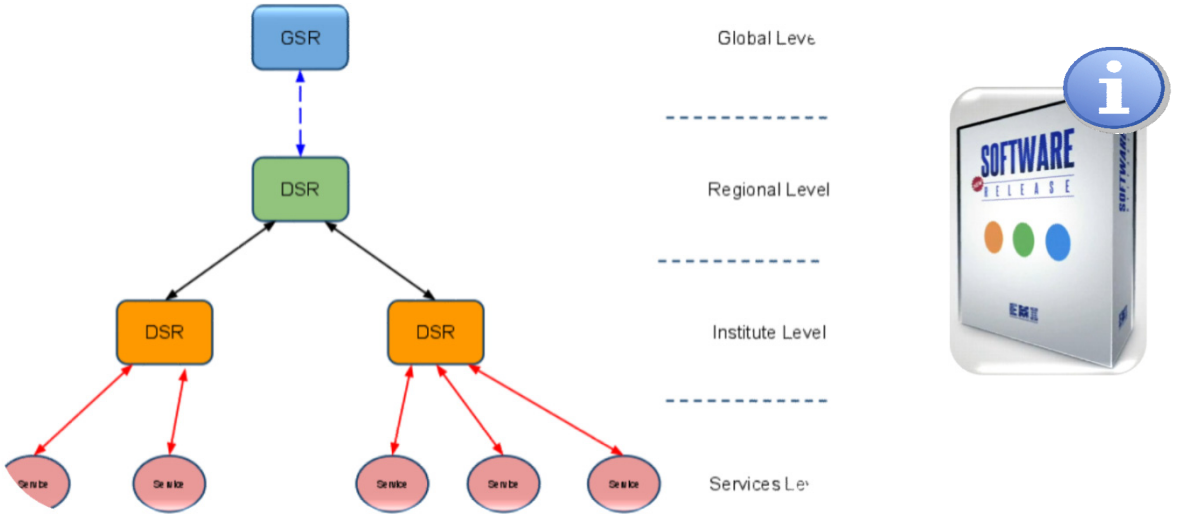

EMI INFSO-RI-261611

EMIR Features

- Use EMIR ReSTful interface to register and query the services
- **Employ flexible, standardized and expressive information model to represent the services (GLUE2 information model)**
- Setup authorization and access control with XACML policies or ACLs
- Write easy your own clients to interact with the service (WADL available)

EMI INFSO-RI-261611

VOMS Added Value

- Attribute Authority (AA) releasing signed security credentials with information beyond pure identities (roles, groups, project, etc.)

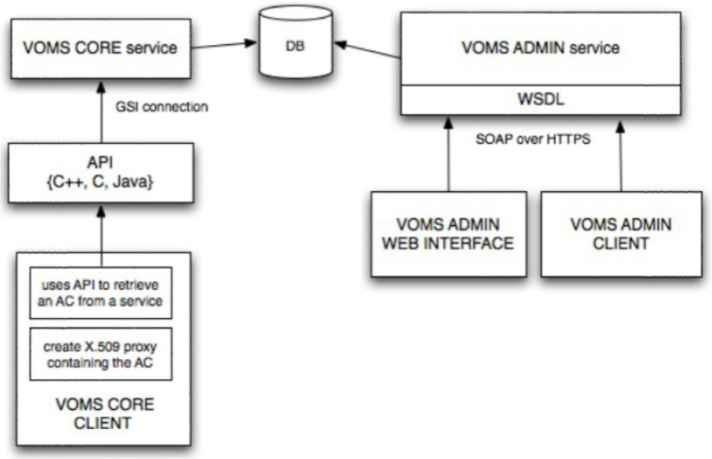

VOMS Features

- Use VOMS as an AA server to obtain signed security credentials with attributes of end-users (e.g. role possession, group/project membership) used during authorization
- Take advantage of a client/server architecture that is able to store identities and manage them in hierarchical groups
- Access and easily configure VOMS using its complementary voms-admin tool
- Interoperate by using the Security Assertion Markup Language (SAML) 2.0 standard via SOAP-based Web service interfaces or X.509 Attribute Certificates
- **Engage in being among the first users that take advantage of the new Representational State Transfer (REST) VOMS interface**

UVOS Added Value

- Attribute Authority (AA) releasing signed security credentials with information beyond pure identities (roles, groups, project, etc.)

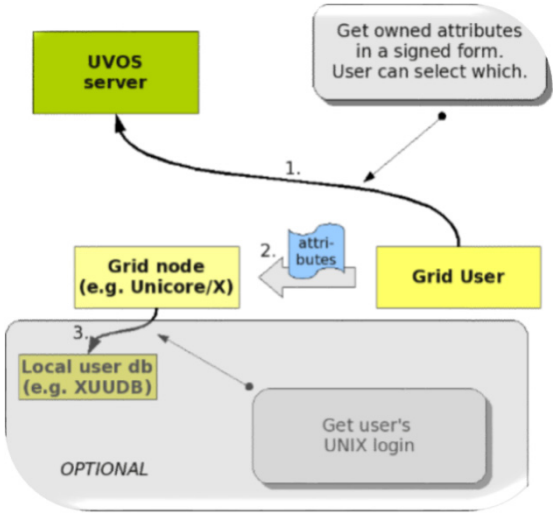

EMI INFSO-RI-261611

UVOS Features

- Use UVOS as an AA server to obtain signed security credentials with attributes of end-users (e.g. role possession, group/project membership) used during authorization
- Take advantage of a client/server architecture that is able store identities and other identifiable servers and organize them in hierarchical groups if needed
- Access and configure UVOS using its client and a lightweight VO authentication Web component optimized for a usage within browsers
- **Interoperate with other services by using the Security Assertion Markup Language (SAML) 2.0 standard via SOAP-based Web service interfaces**

ARGUS Added Value

- Used to derive authorization decisions

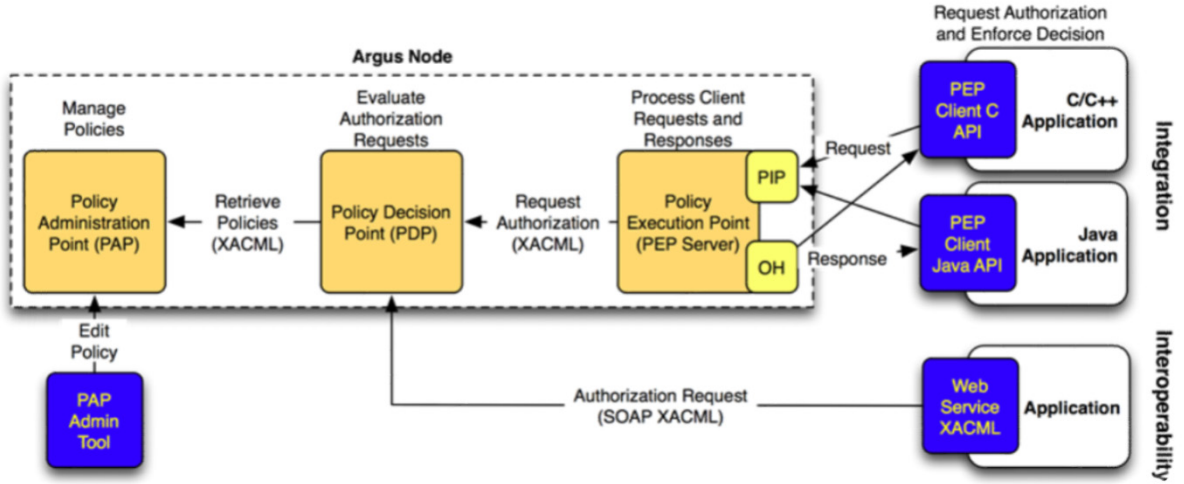

EMI INFSO-RI-261611

ARGUS Features

- Use ARGUS as a system to render consistent authorization decisions across geographically distributed services (computing, data, portals, etc.)
- Take advantage of a client/server architecture that implements the functionality of a Policy Enforcement Point (PEP)
- Manage policies through a Policy Administration Point (PAP) and its admin tool
- **Interoperate with other services by using the Extensible Access Control Markup Language (XACML) standard via a SOAP-based Web service Interface**

EMI FactSheets Available

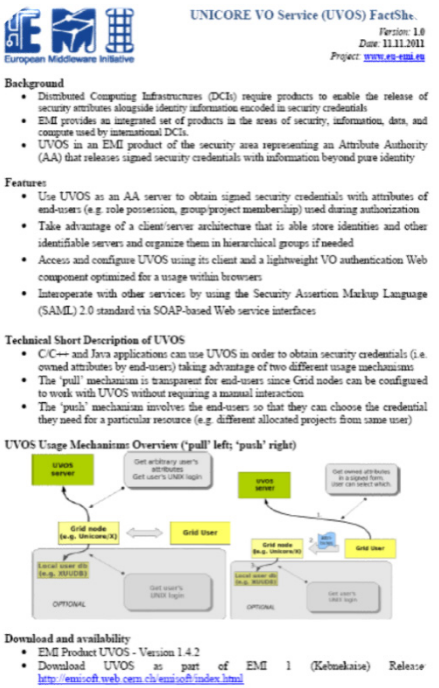

UNICORE VO Service (UVOS) FactShe.

Version: 1.0
Date: 11.11.2011
Project: www.eu-emi.eu

Background

- Distributed Computing Infrastructures (DCIs) require products to enable the release of security attributes alongside identity information encoded in security credentials
- EMI provides an integrated set of products in the areas of security, information, data, and compute used by international DCIs.
- UVOS in an EMI product of the security area representing an Attribute Authority (AA) that releases signed security credentials with information beyond pure identity

Features

- Use UVOS as an AA server to obtain signed security credentials with attributes of end-users (e.g. role possession, group/project membership) used during authorization
- Take advantage of a client/server architecture that is able store identities and other identifiable servers and organize them in hierarchical groups if needed
- Access and configure UVOS using its client and a lightweight VO authentication Web component optimized for a usage within browsers
- Interoperate with other services by using the Security Assertion Markup Language (SAML) 2.0 standard via SOAP-based Web service interfaces

Technical Short Description of UVOS

- C/C++ and Java applications can use UVOS in order to obtain security credentials (i.e. owned attributes by end-users) taking advantage of two different usage mechanisms
- The 'pull' mechanism is transparent for end-users since Grid nodes can be configured to work with UVOS without requiring a manual interaction
- The 'push' mechanism involves the end-users so that they can choose the credential they need for a particular resource (e.g. different allocated projects from same user)

UVOS Usage Mechanism: Overview ('pull' left; 'push' right)

Download and availability

- EMI Product UVOS - Version 1.4.2
- Download UVOS as part of EMI 1 (Kebnekaise) Release http://emisoft.web.cern.ch/emisoft/index.html

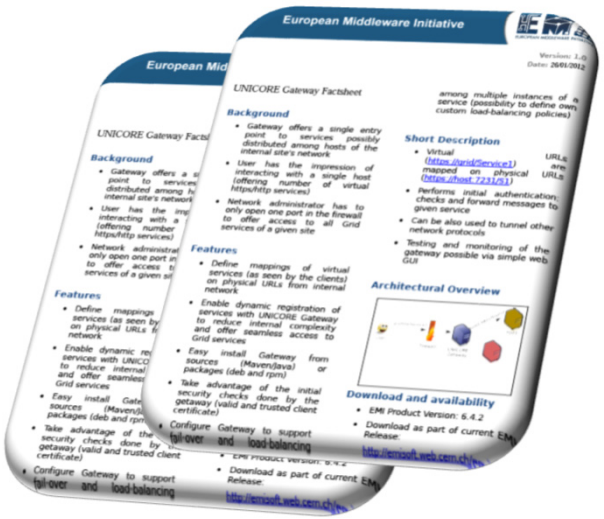

[8] EMI Web Page

EMI – More than just Products!

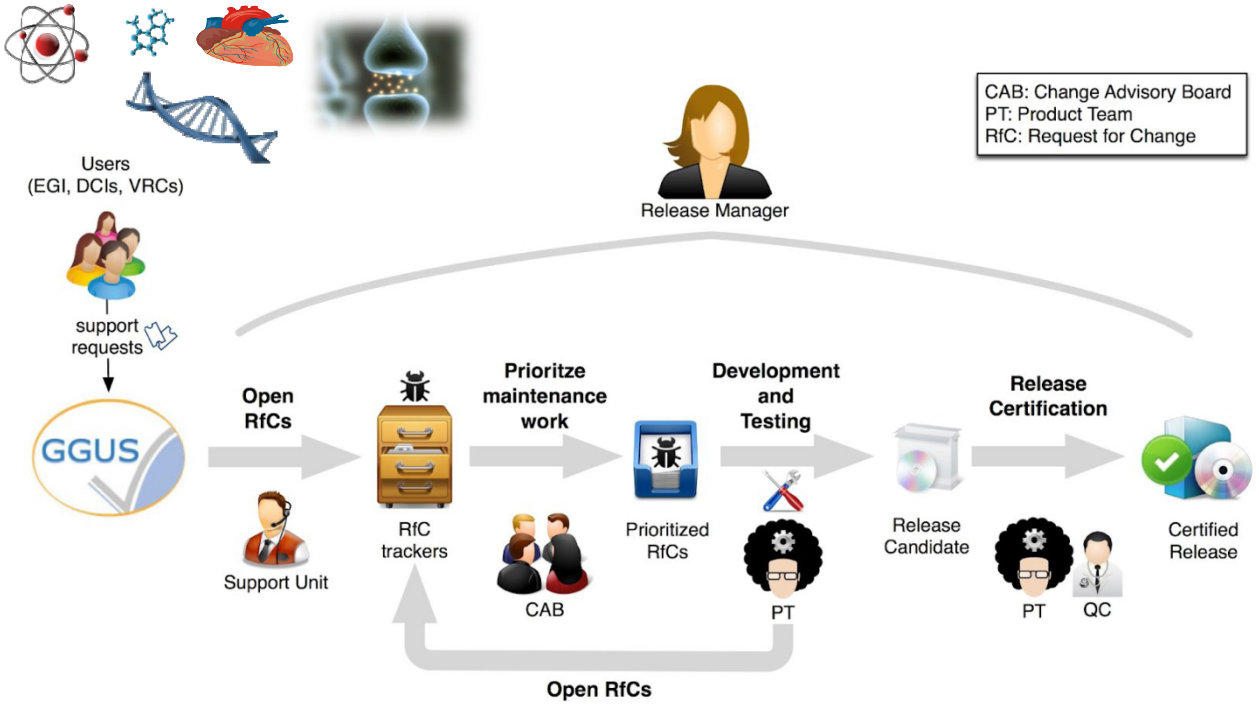

EMI Standard Adoptions

	A-Rex	ARC Info System	ARGUS	BDII	CREAM	dCache	DPM	EGIIS	FTS	LFC	SAGA-SD-RAL	StoRM	UNICORE	WMS	VOMS
EMI-ES (PGI)	EMI 3				EMI 3								EMI 3	EMI 3	
GLUE2			EMI 3					EMI 3						EMI 1	EMI 3
GridFTP															
HTTPS (not GSI)	EMI 3				EMI 3	EMI 3	EMI 3					EMI 3		EMI 3	
JSDL															
NFS 4.1/pNFS															
OGSA-BES															
POSIX/IO							EMI 1								
SAGA-ISN															
SAML															
SRM													EMI 3		
UR (Compute)	EMI 3				EMI 3								EMI 3		
UR (Storage)						EMI 3	EMI 3		EMI 3			EMI 3			
WebDAV							EMI 3					EMI 3			
WSRF															
XACML	EMI 3														

Update: **STS Service adopts WS-Trust**

[8] EMI Web Page

Selected EMI Product Use Cases

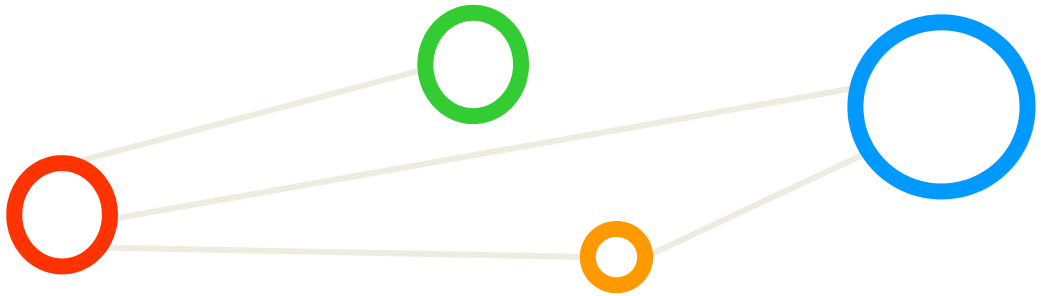

EMI INFSO-RI-261611

Use Case: EGI&WLCG in HEP (1)

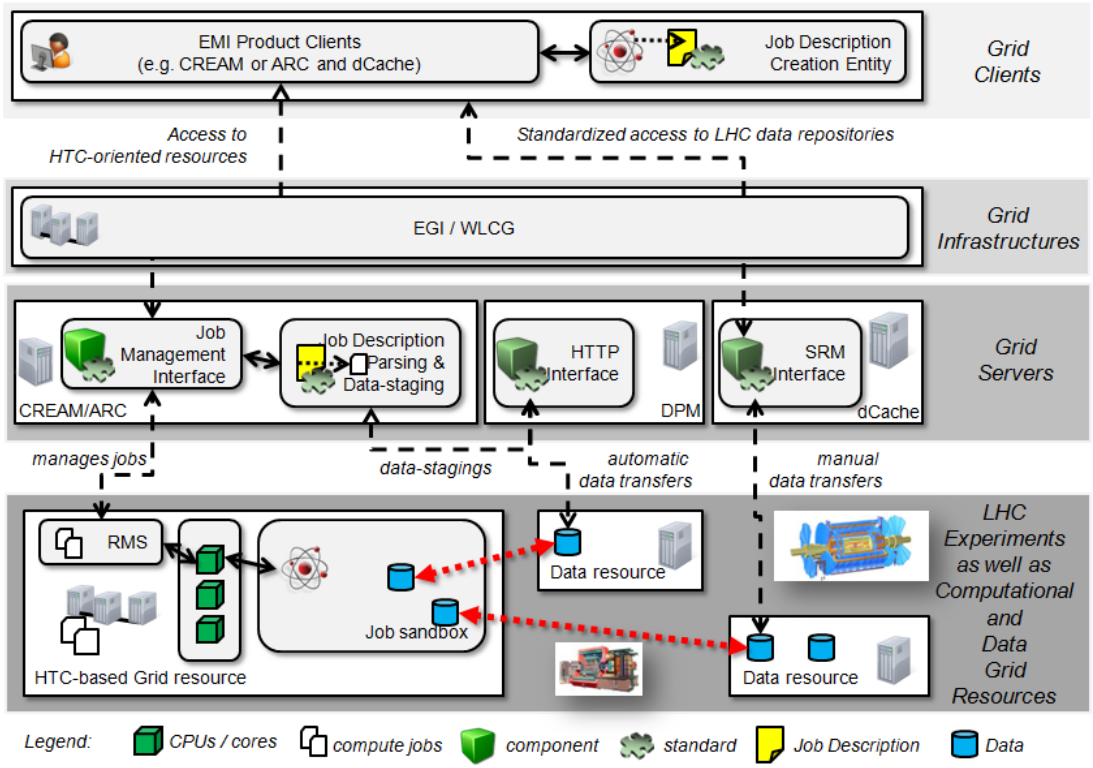

EMI INFSO-RI-261611

Use Case: EGI&WLCG in HEP (2)

Use Case	1.0 WLCG – HEP Community
Description	The goal of this use case is to achieve transparent access to

	EMI Product	Usage Description
References	***ARC-CE***	The ARC-CE is used to submit and manage computational activities on the EGI infrastructure with WLCG resources and is particularly often used by the Nordic HEP community.
Actors		
Prerequisites (Dep Assumptions	***BDII***	The BDII is used to obtain resource information of the WLCG community available within the EGI infrastructure.
	CREAM-CE	The CREAM-CE is used to submit and manage computational activities on the EGI infrastructure with WLCG resources.
	dCache	The dCache product is used to access and manage data storage systems used for experimental data and for analysis results.
Steps		
	DPM	The DPM product is used to access and manage data storage systems used for experimental data and for analysis results.
	StoRM	The StoRM product is used to provide transparent access to parallel file systems (e.g. GPFS) and is used with experimental data and analysis results.
Variations (optional,		
Quality Attributes	***VOMS***	The VOMS product is used to release signed security attributes including role possession and VO/project membership
Non-functional (opti	***WMS***	The WMS product is used as a broker for computing resources available to the WLCG – HEP community and forwards job submissions to the corresponding resources.
Issues		

Use Case: EGI Science Cases (1)

Stories from the grid

'Stories from the Grid' is a series of short films which explore how researchers from different scientific disciplines are using cutting edge grid computing technology to advance their work.

Episode 1: The cone snail and the search for powerful new anaesthetics

Episode 2: Reviving the lost sounds of the epigonion

Episode 3: Hunting for the top quark in the Large Hadron Collider

Marcel Vreeswijk and Hurng-Chun Lee from NIKHEF (the Dutch National Institute for Subatomic Physics) are studying a particle called the top quark created by the Large Hadron Collider (LHC) – the world's largest scientific instrument.

Particle physicists use the LHC to study variations from the Standard Model and potentially discover new laws of physics, governing everything from dark matter to extra dimensions. The particle known as the top quark is a window into this weird and wonderful world.

In this short film, Marcel and Hurng-Chun describe their study and explain how customised grid computing workflows are key to filtering and sieving massive sets of data down to a manageable size. Without these tools, it would be impossible to pick out the key results that could hold the clues to top quark behaviour.

[4] EGI Science Cases

Use Case: EGI Science Cases (2)

[5] EGI Science Stories

EMI INFSO-RI-261611

Use Case: PRACE (1)

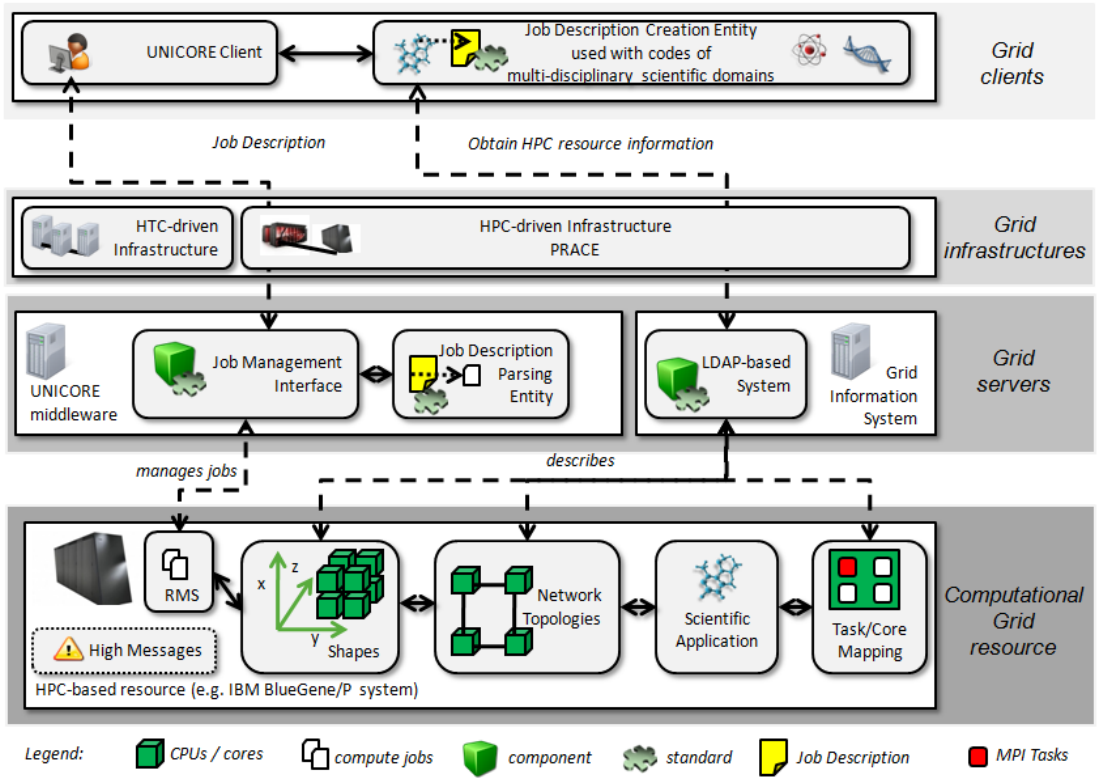

EMI INFSO-RI-261611

Use Case: PRACE (2)

Use Case	2.0 PRACE – Multi-Disciplinary Science Using HPC
Description	The goal of this use case is to achieve transparent access to geographically disperse large-scale HPC resources part of the PRACE infrastructure.
References	Partnership for Advanced Computing in Europe (PRACE) [3]
Actors	Multi-disciplinary scientists
Prerequisites (Dependencies) & Assumptions	The following resources and products must be available: • Computational HPC resources • EMI products for distributed computing • Parallel file systems (for data storage/access) • Resource Management Systems (aka batch sub-systems)
Steps	**EMI Product** / **Usage Description** *UNICORE* / UNICORE is used to submit and manage computational activities on the PRACE infrastructure with large-scale HPC resources and is particularly often used with scientific code that takes advantage of parallel computing methods (e.g. OpenMP, MPI, etc.). simulation code on large-scale HPC resources
Variations (optional)	none
Quality Attributes	Reliability of the middleware; high-level of security (not full in-personification of users working on behalf of others); local access control
Non-functional (optional)	Interoperability with US activities in particular with TeraGrid/XSEDE; Interoperability with EGI for parameter sweep studies or smaller HTC-based evaluation runs;
Issues	Easier access to resources (security); advance reservation functionality

Use Case: PRACE (3)

[3] PRACE Annual Scientific Report

EMI INFSO-RI-261611

Use Case: EUFORIA (1)

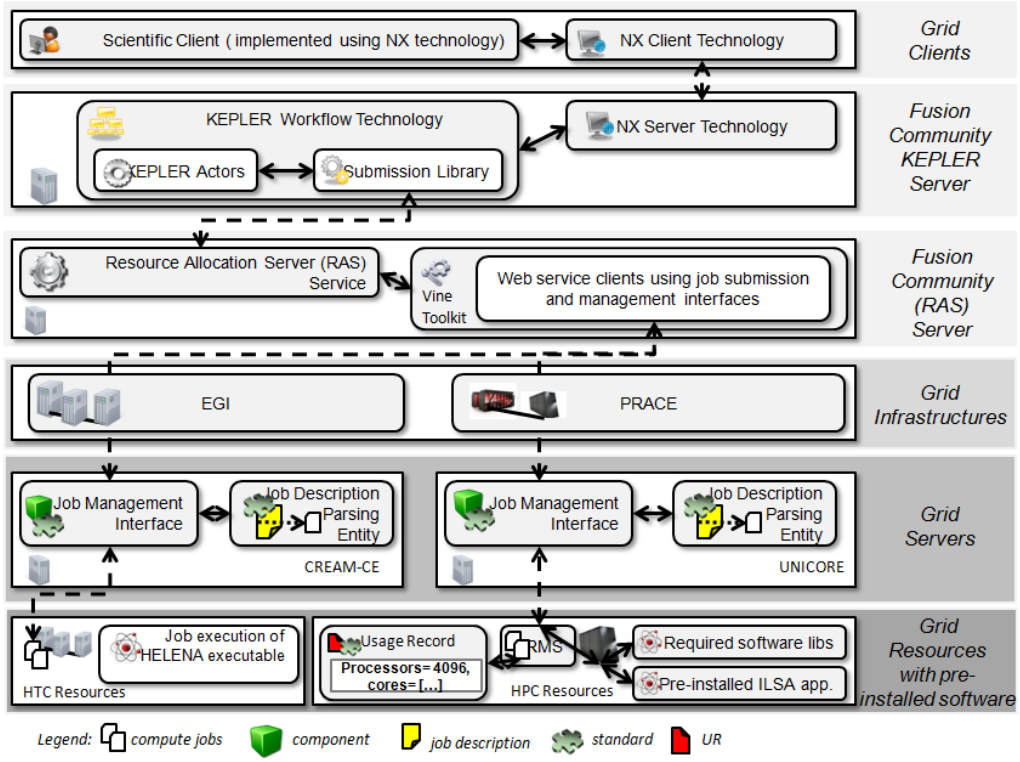

EMI INFSO-RI-261611

Use Case: EUFORIA (2)

Use Case	3.0 EUFORIA – Fusion Science Community
Description	The goal of this use case is to achieve transparent access to geographically disperse large-scale HTC and HPC resources part of the EGI and PRACE infrastructure.
References	EU Fusion for Iter Applications Project (EUFORIA) [4]
Actors	Fusion scientists
Prerequisites (Dependencies) & Assumptions	The following resources and products must be available: • KEPLER Workflow engine • KEPLER Actors for EMI Middleware
Steps	**EMI Product** / **Usage Description** ***CREAM-CE*** — The CREAM-CE is used to submit and manage computational activities on the EGI infrastructure with fusion codes designed to take advantage of HTC resources. ***UNICORE*** — A UNICORE KEPLER actor is used to submit and manage computational activities on the PRACE infrastructure with large-scale HPC resources used with scientific fusion codes that take advantage of parallel computing methods (e.g. OpenMP, MPI, etc.).
Variations (optional)	***VOMS*** — The VOMS product is used to release signed security attributes including role possession and VO/project membership
Quality Attributes	Maturity of the middleware since it is adopted in another workflow-based framework (i.e. KEPLER)
Non-functional (optional)	Interoperability between middleware products to decrease the amount of different KEPLER actors needed to be maintained;
Issues	Easier access to resources (security);

Use Case: EUFORIA (3)

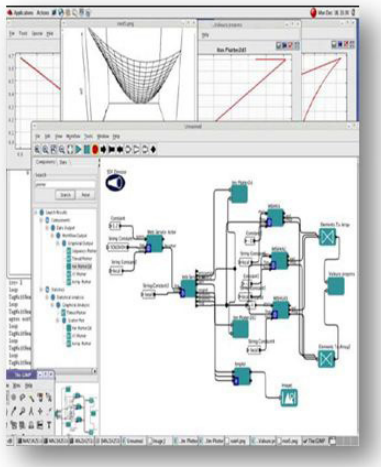

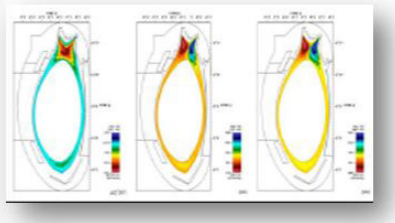

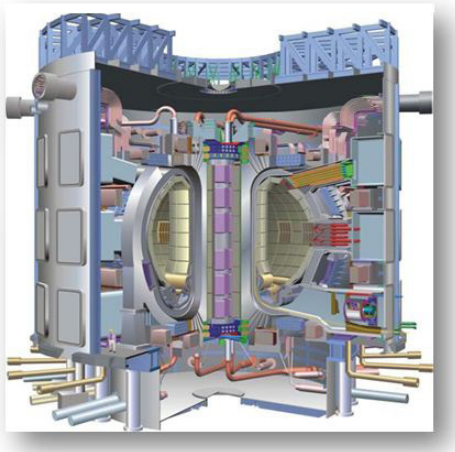

[6] EUFORIA Paper

EMI INFSO-RI-261611

Use Case: VPH (1)

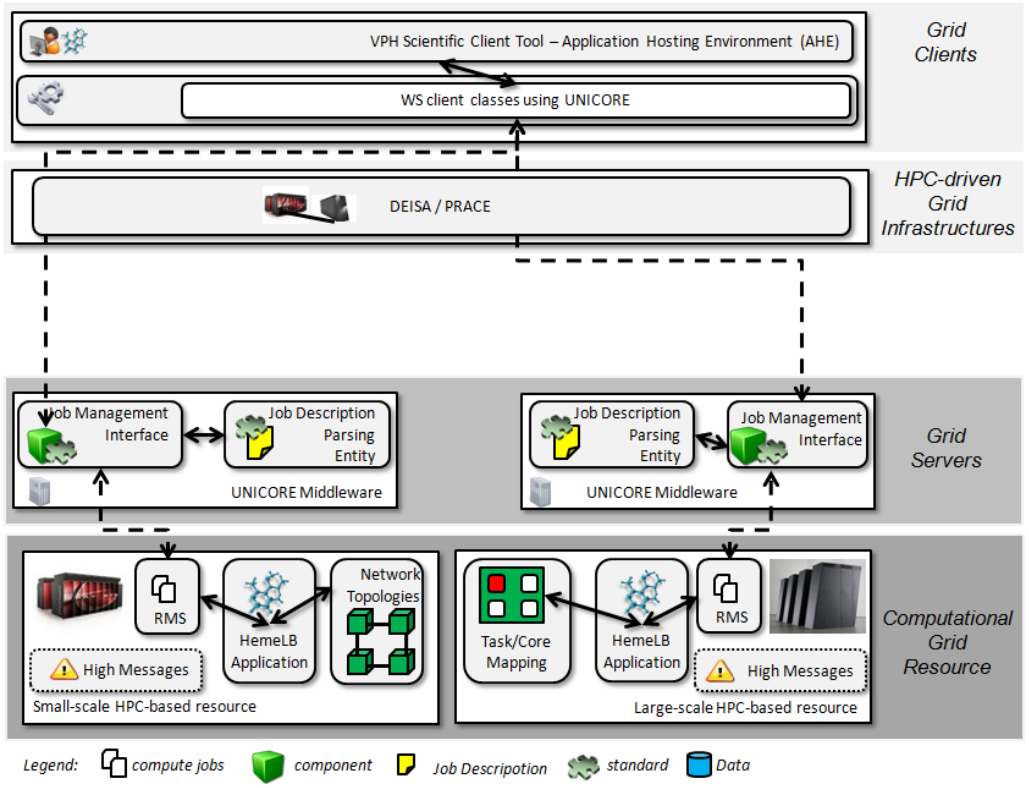

EMI INFSO-RI-261611

Use Case: VPH (2)

Use Case	4.0 VPH – Bio-Medical Science Community
Description	The goal of this use case is to achieve transparent access to geographically disperse large-scale HPC resources part of the PRACE infrastructure.
References	Virtual Physiological Human (VPH) Network of Excellence [5]
Actors	Bio-medical scientists
Prerequisites (Dependencies) & Assumptions	The following resources and products must be available: • Application Hosting Environment (AHE) • HPC Computational resources Assumption that bio-medical scientists have a wide variety of scientific codes to run using parallel computing methods.
Steps	1. Bio-medical scientists use AHE with a scientific code

EMI Product	Usage Description
UNICORE	A UNICORE VPH adapter is used within the AHE to submit and manage computational activities on the PRACE infrastructure with large-scale HPC resources used with scientific codes that take advantage of parallel computing methods (e.g. OpenMP, MPI, etc.).

framework (i.e. AHE Client)

Variations (optional)	
Quality Attributes	
Non-functional (optional)	Interoperability between middleware products to decrease the amount of different AHE middleware adapters needed to be maintained;
Issues	Easier access to resources (security); support for advance reservation middleware capability

Use Case: VPH (3)

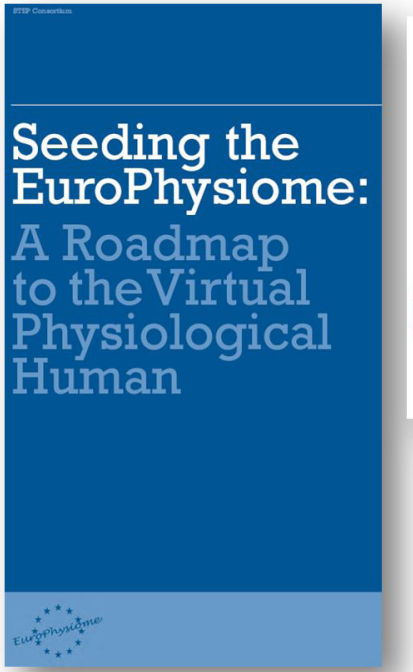

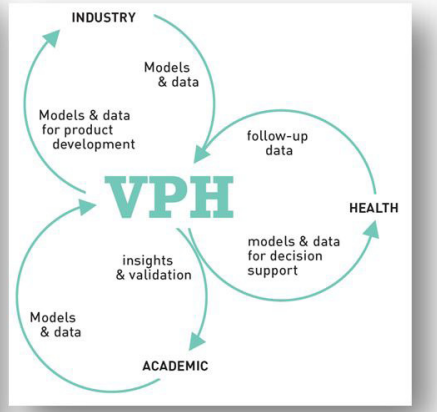

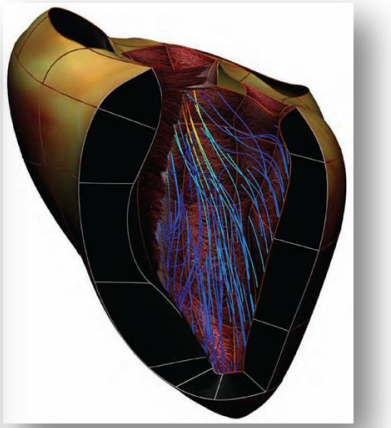

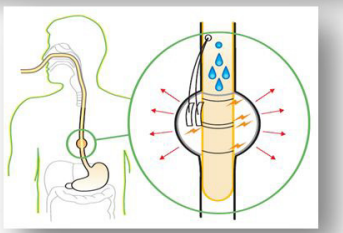

[7] VPH Paper

Many other use cases in e-Science

- NeuroScience Community: NeuGrid4You et al.
- Structural Biology: WeNMR et al.
- OSG deploys several EMI products
  - –VOMS, VOMS-Admin, etc.
- XSEDE is starting deploying UNICORE
- YOUR USE CASE!
  - – Get in contact with m.riedel@fz-juelich.de
  - –Perform **'Application Enabling'** with YOUR specific scientific/business needs with us

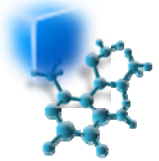

A Business Use Case

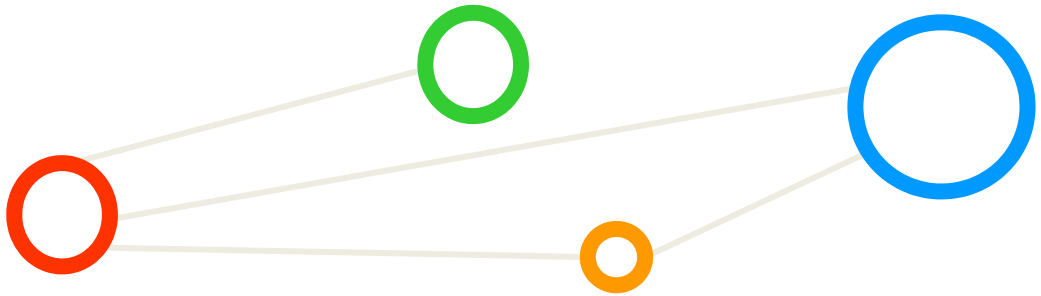

EMI INFSO-RI-261611

A Business Case: dCore

- The DCore Holding
- DCore Systems SA is a private holding based in Luxembourg with headquarters in Geneva
- It is supported by a group of private investors and funds

[2] dCore Systems

dCore Business Model

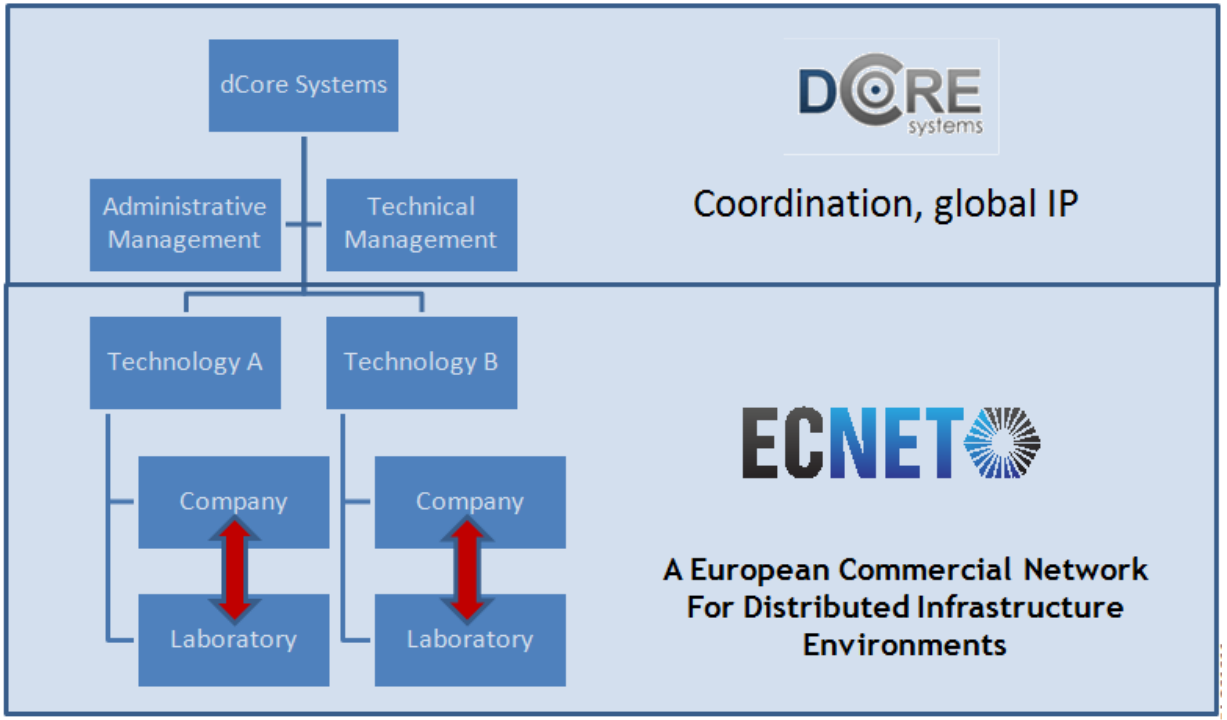

EMI INFSO-RI-261611

dCore and EMI

GALA Flexible. Secure. Efficient

Secure Digital Archiving

Sinefo

Mobile Online Security

QPAD

Distributed Project Management

SAFEH

Secure Data Management for Healthcare

COURTICE

Secure Data Management for Law Professional

HFILE

Secure File and E-mail Storage and Sharing

European Middleware Initiative

dCache, Hydra, VOMS, Argus

- Storage
- Encryption
- Authorization

EMI INFSO-RI-261611

Lessons Learned

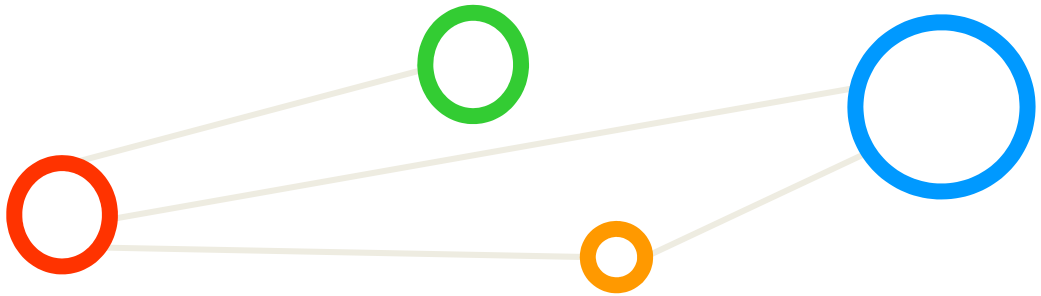

EMI INFSO-RI-261611

Lessons Learned

- EMI products are used in e-Science
  - Applications from various disciplines
  - e-Infrastructures such as EGI, PRACE, XSEDE, OSG
- e-Science Application Enabling is important
  - Really working together with e-Scientists together
  - ‘Consulting’ to end users to use technology needed
- Develop & Promote re-usable Components
  - Break ‘atomic middleware’ in re-usable components
  - Easier uptake from the non experts and commercials

References

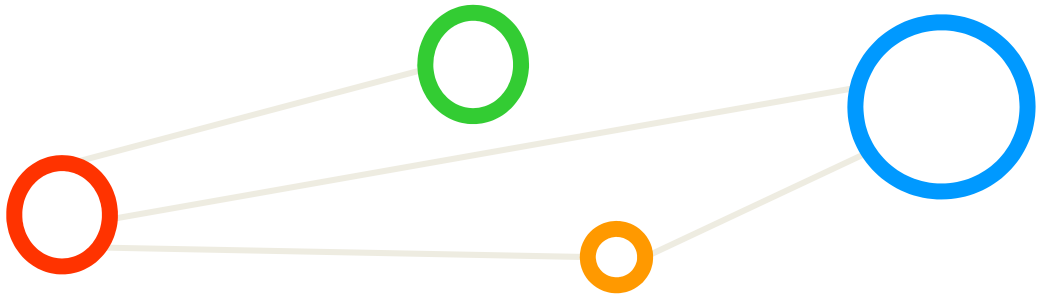

EMI INFSO-RI-261611

General References

[1] Digital Agenda for Europe
http://ec.europa.eu/information_society/digital-agenda/index_en.htm
[2] dCore Systems
http://www.dcore-sa.com/about/about.php
[3] PRACE Annual Scientific Report
http://www.prace-ri.eu/PRACE-Scientific-Annual-Report?lang=en
[4] EGI Science Cases
http://www.egi.eu/news-and-media/videos/
[5] EGI Science Stories
http://www.egi.eu/news-and-media/publications/
[6] M. S. Memon, M. Riedel, A. Memon, F. Wolf, A. Streit, T. Lippert, M. Plociennik, M. Owsiak, D. Tskhakaya, and C. Konz. Lessons Learned from Jointly Using HTC- and HPC-driven e-Science Infrastructures in Fusion Science. In Proceedings of 2nd IEEE International Conference on Information and Emerging Technologies (ICIET 2010), Karachi, Pakistan, 2010.
[7] M. Riedel, B. Schuller, M. Rambadt, M. Memon, A. Memon, A. Streit, F.Wolf, T. Lippert, S. Zasada, S. Manos, P. Coveney, F. Wolf, and D. Kranzlmueller. Exploring the Potential of Using Multiple e-Science Infrastructures with Emerging Open Standards-based e-Health Research Tools. In Proceedings of the 10th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2010), Melbourne, Australia, pages 341–348, 2010.
[8] EMI Web Page
http://www.eu-emi.eu/
[9] John Taylor, Enhanced-Science (e-Science) Definition
http://www.e-science.clrc.ac.uk