

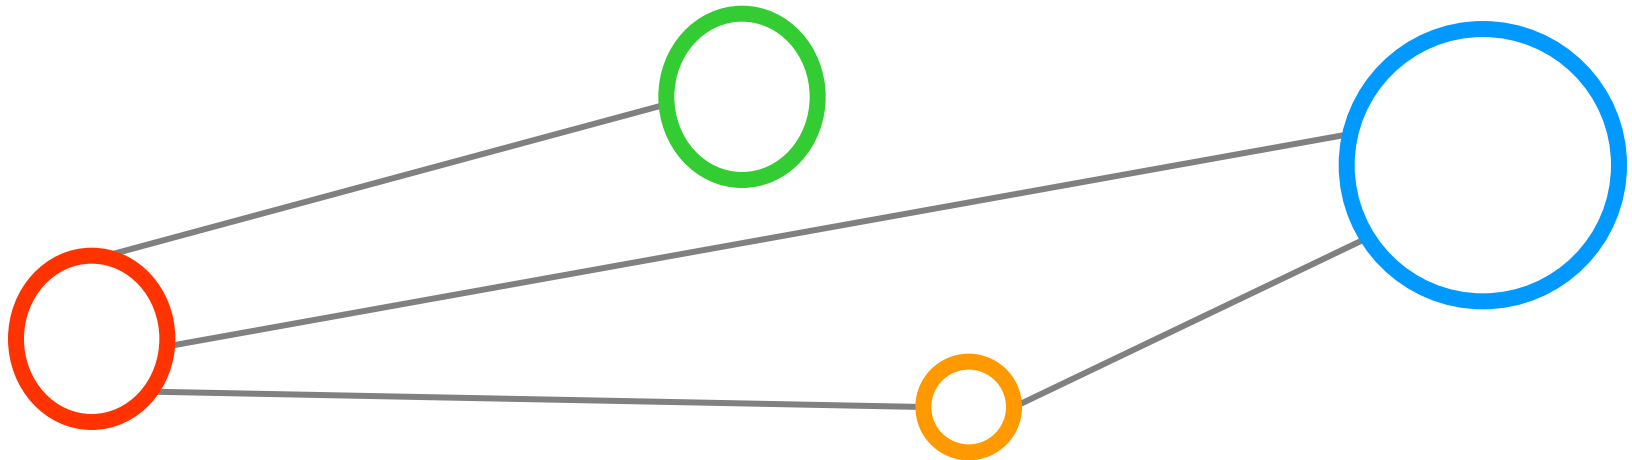
Production Grid Infrastructure (PGI)

Taking lessons learned from production into standardization

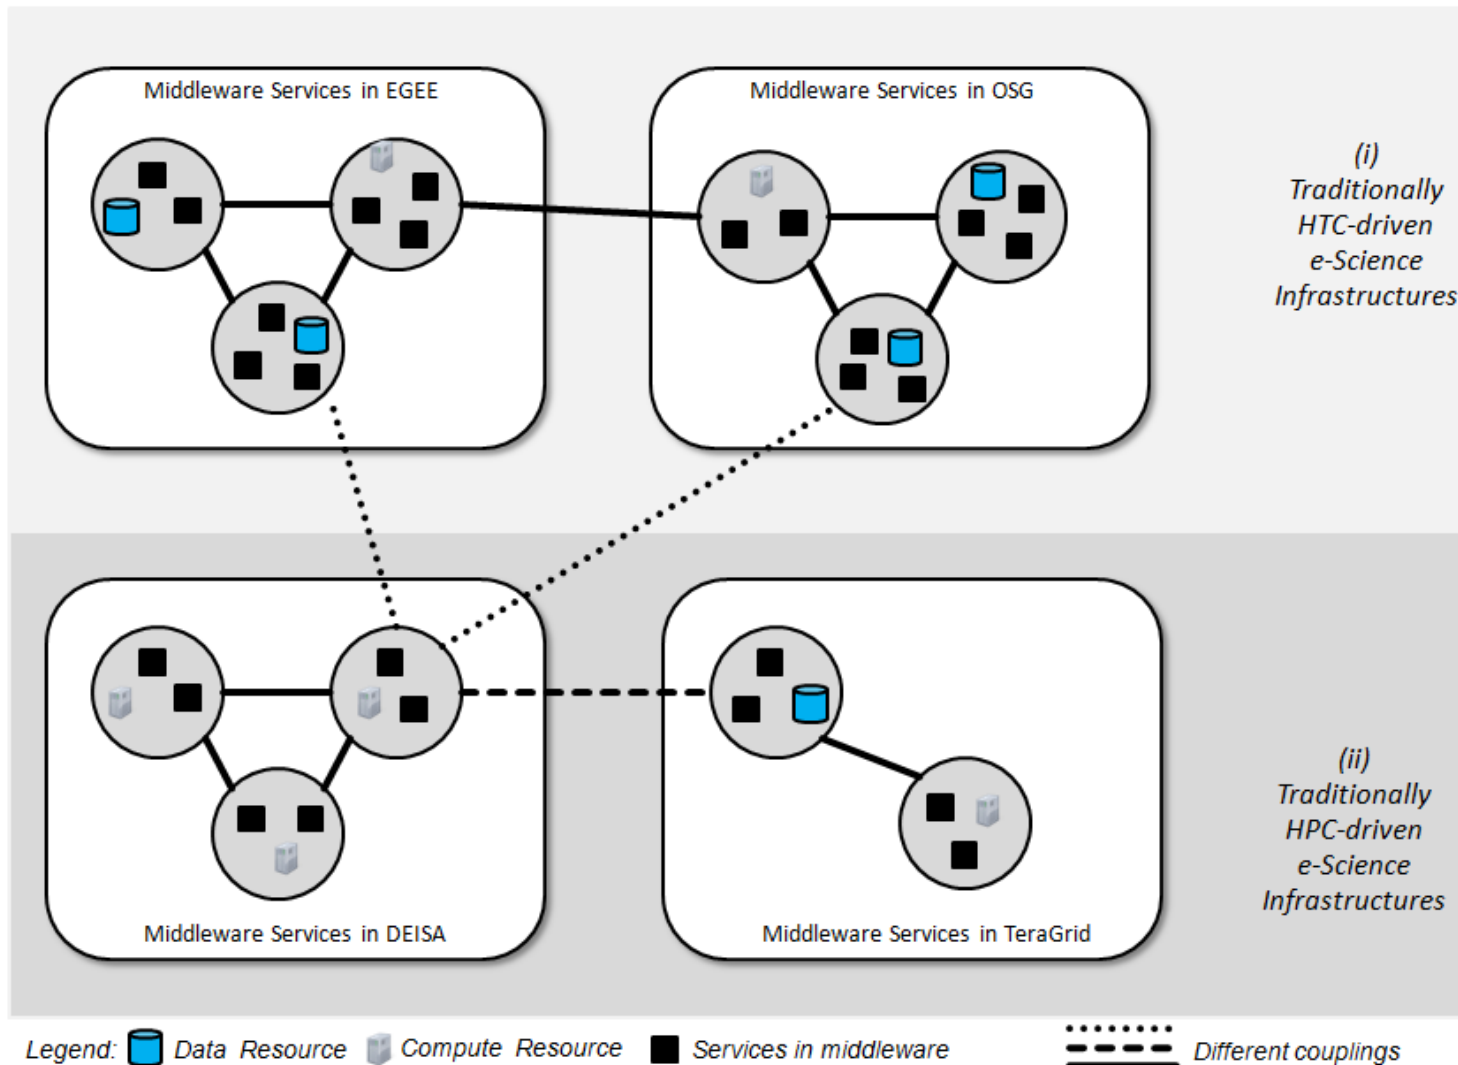
Morris Riedel (Jülich Supercomputing Centre & EMI) et al.
Group Co-Chair Grid Interoperation Now & Production Grid Infrastructure

*“...separating the e-science technology & standard hypes from
e-science production infrastructure reality today and tomorrow...
...and this reality is sometimes painful and ‘slower’ then expected...”*

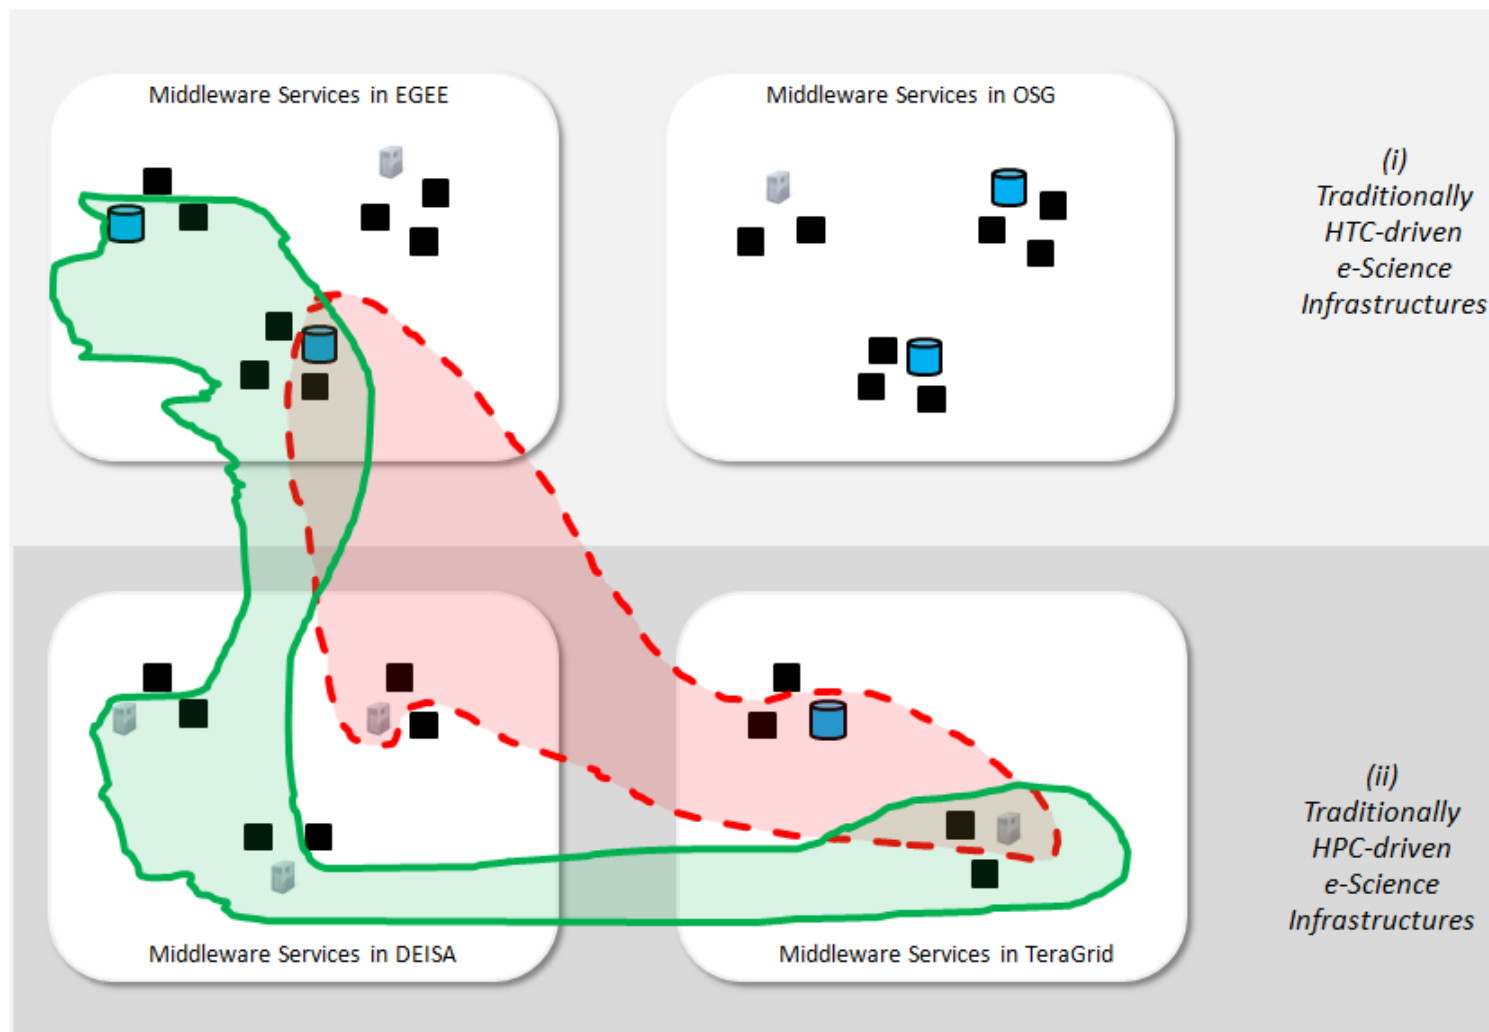
Goals & Vision







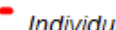
History: Middleware-based Clusters



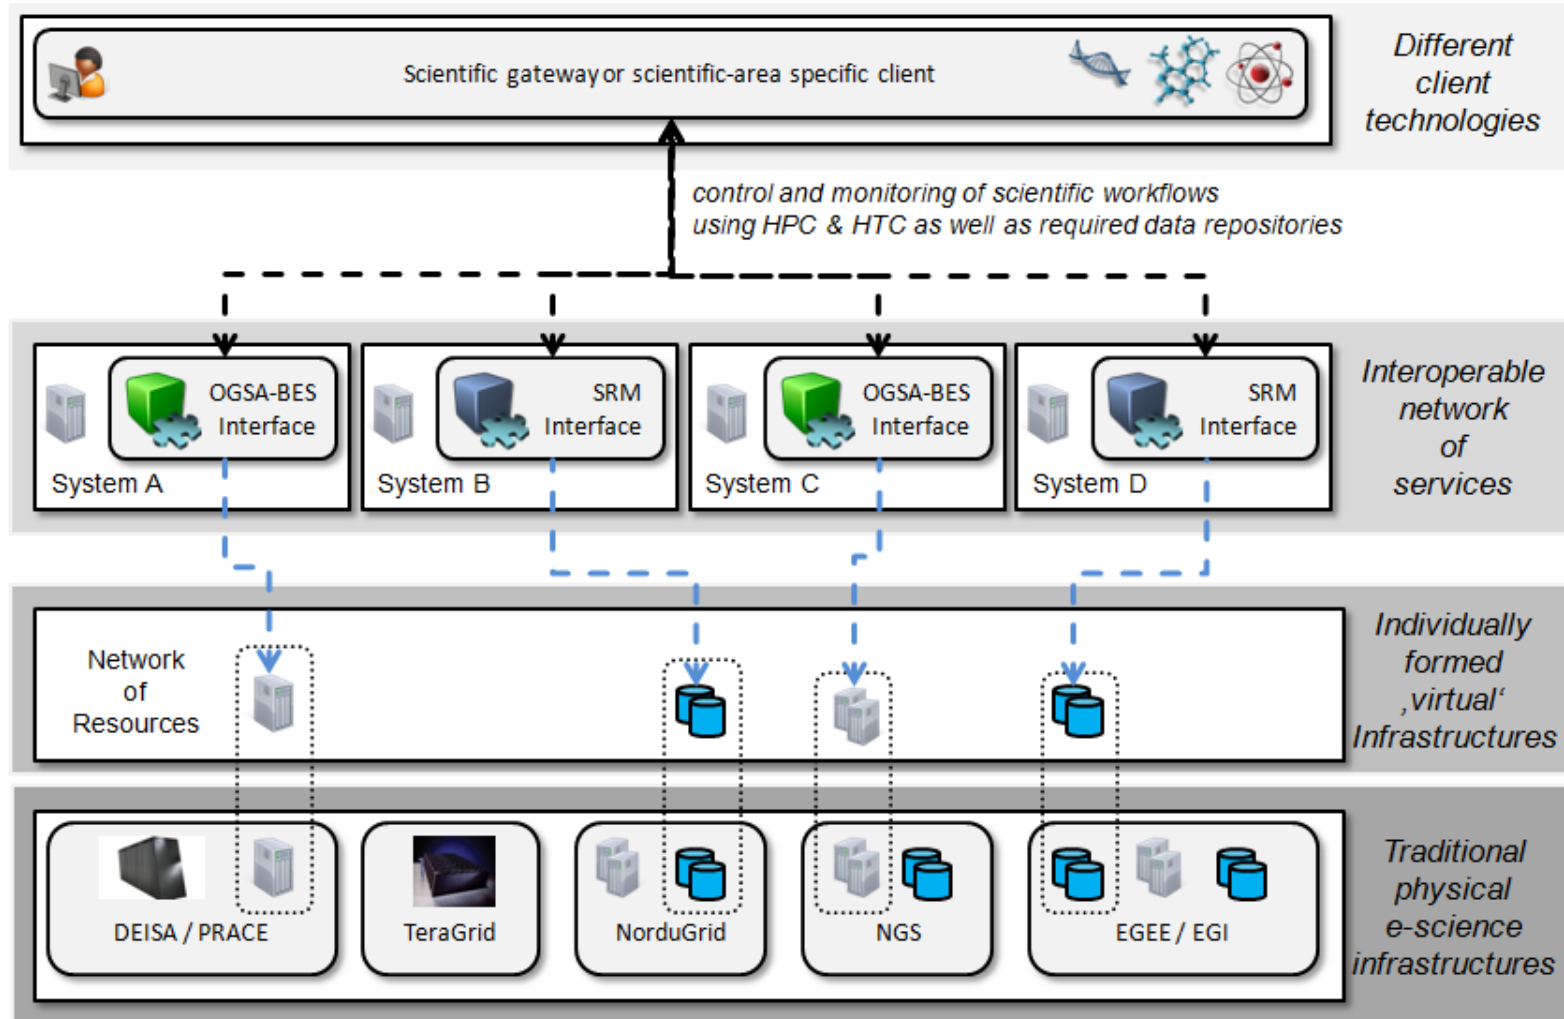
Goal: Network of interoperable Grids







Legend:  Data Resource  Compute Resource  Services in middleware

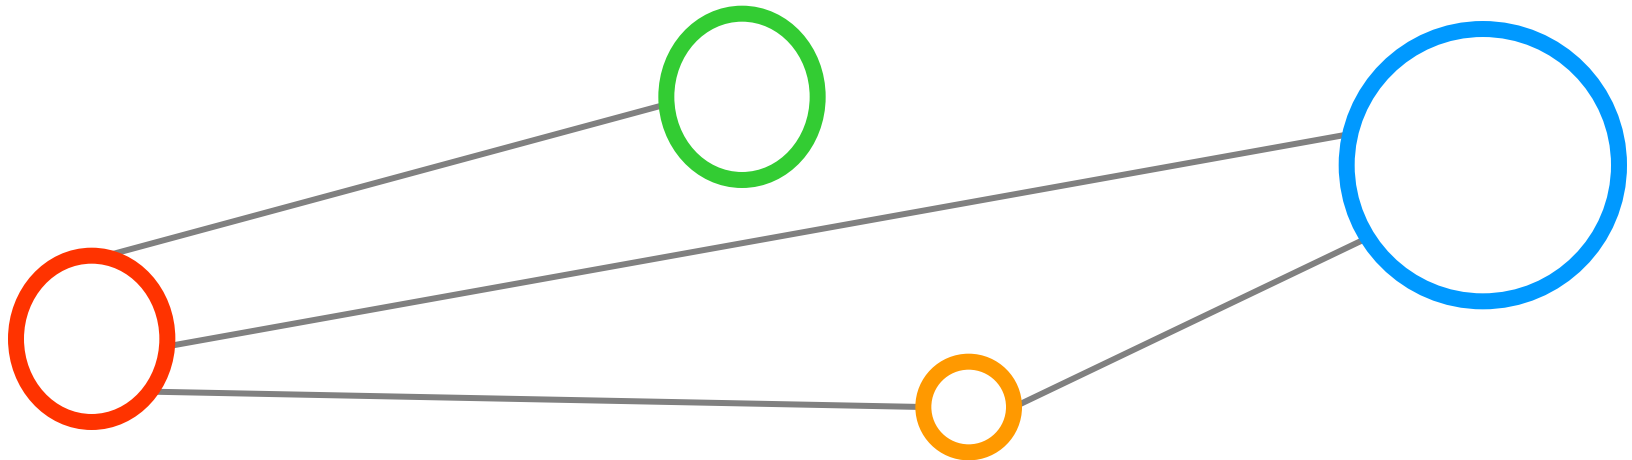
 Individual Infrastructures  Network of interoperable Grids

Many interoperability efforts...

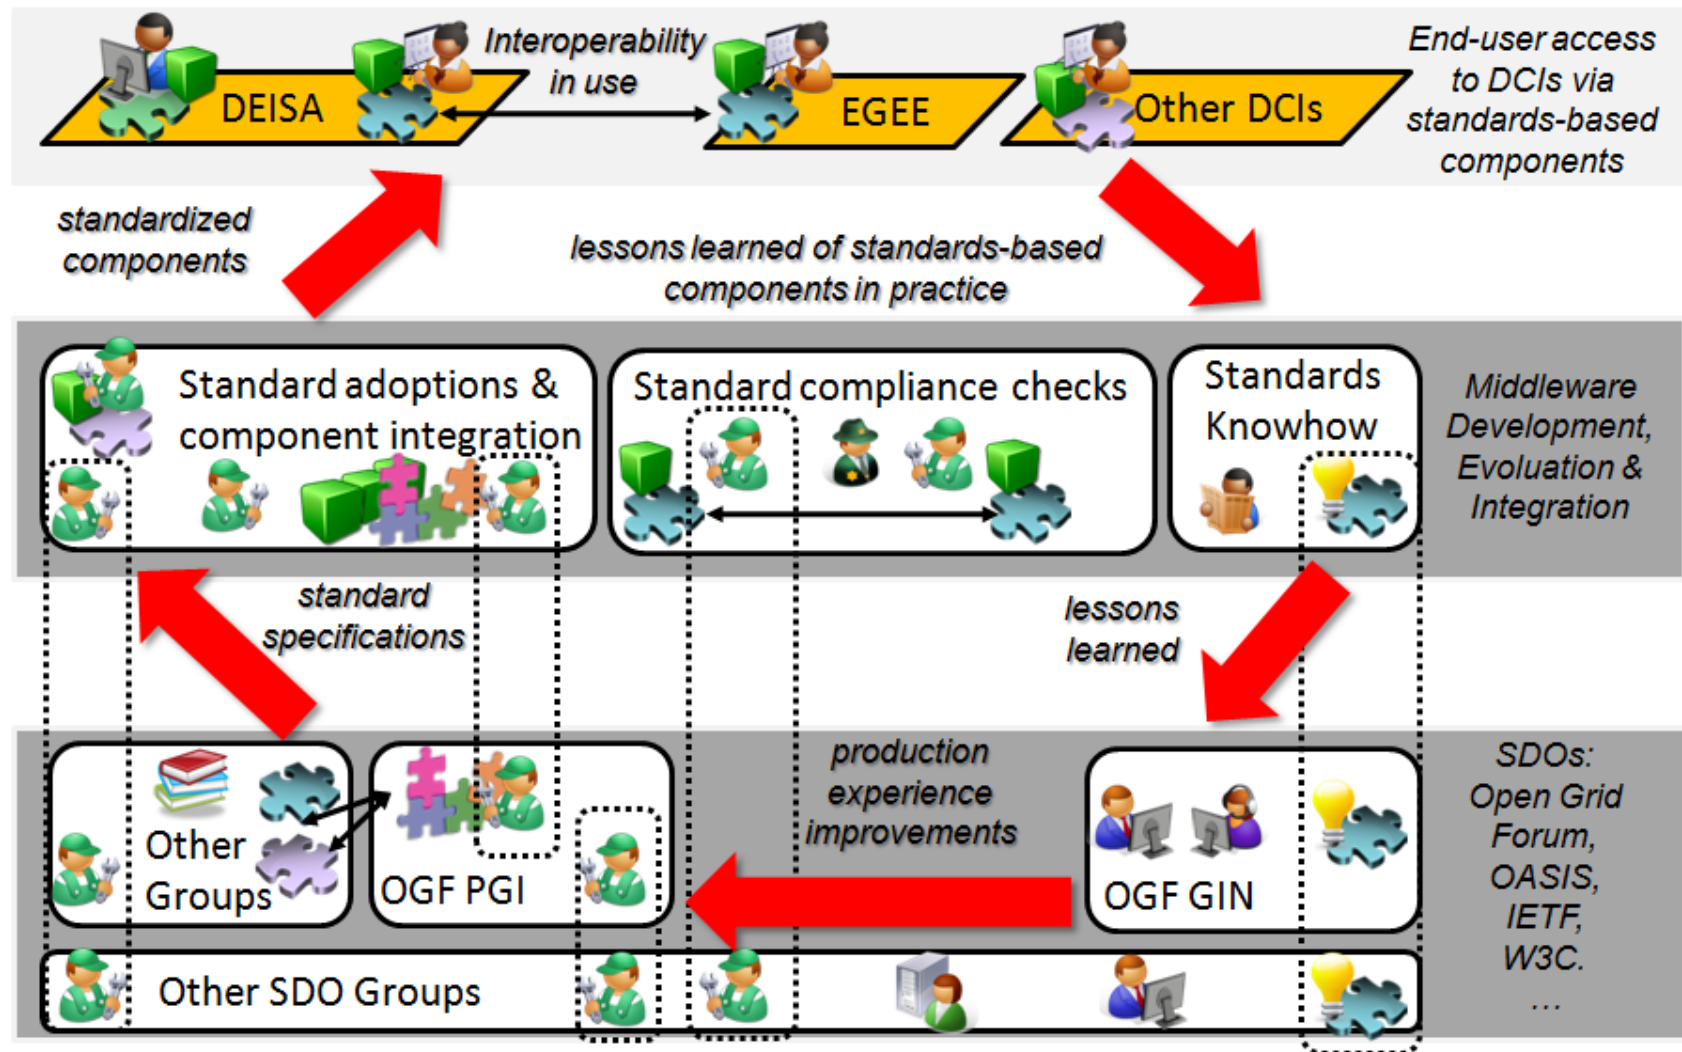


Legend:  CPUs / cores  component  standard  Data

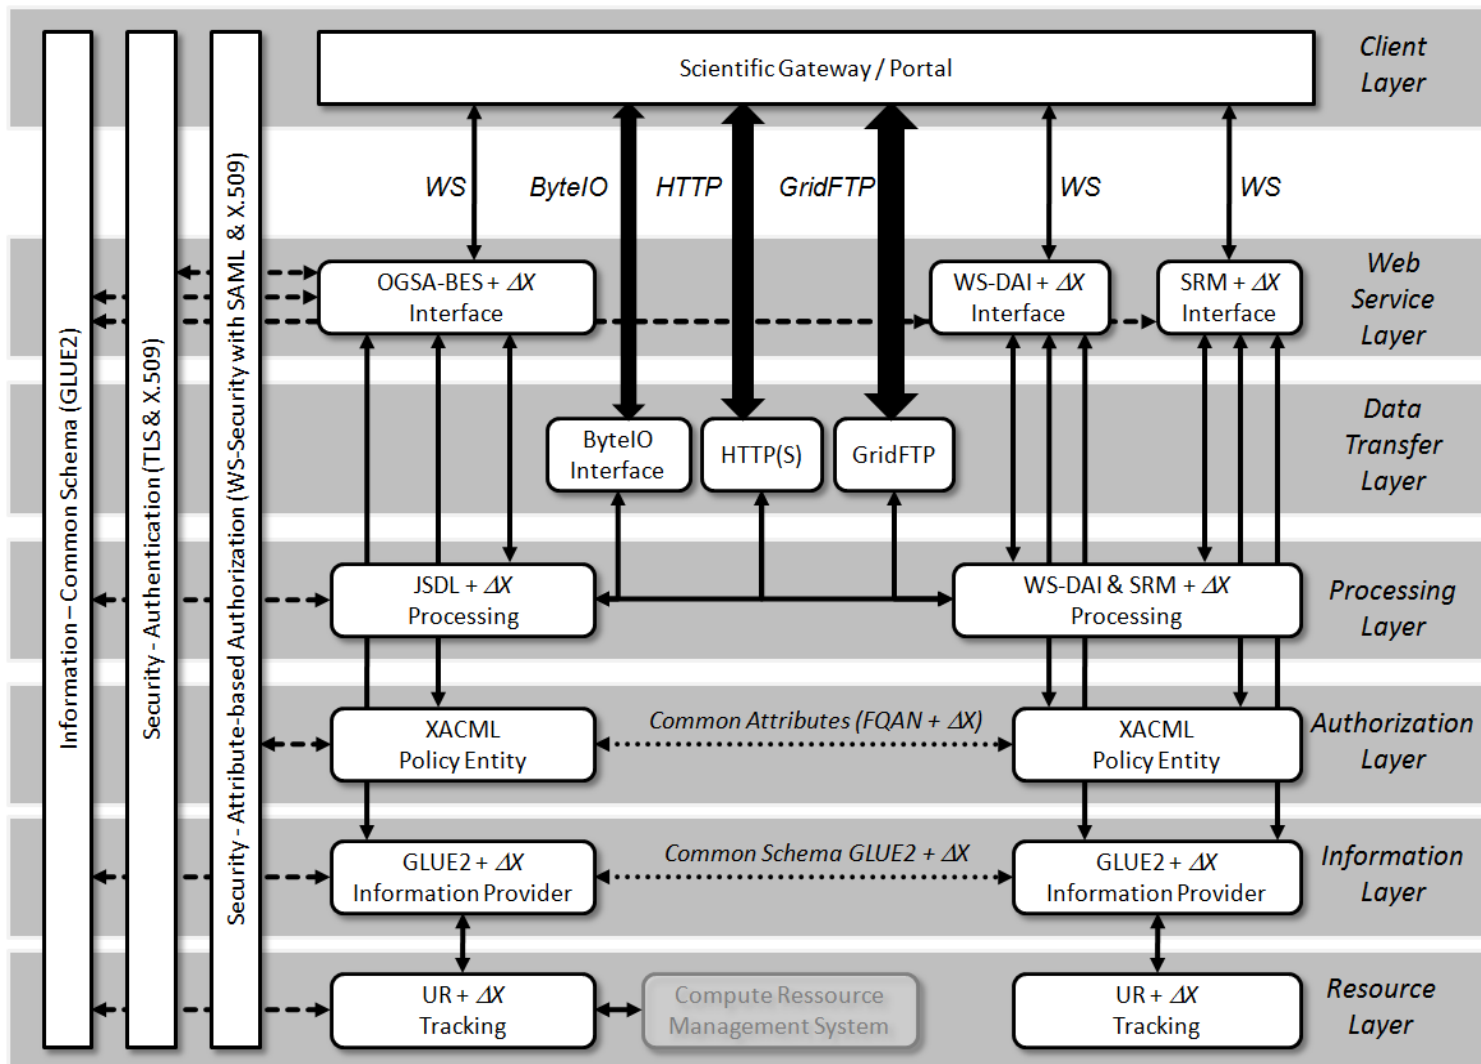
OGF GIN & PGI



GIN & PGI: Standards Ecosystem



GIN/PGI Reference Model/Roadmap

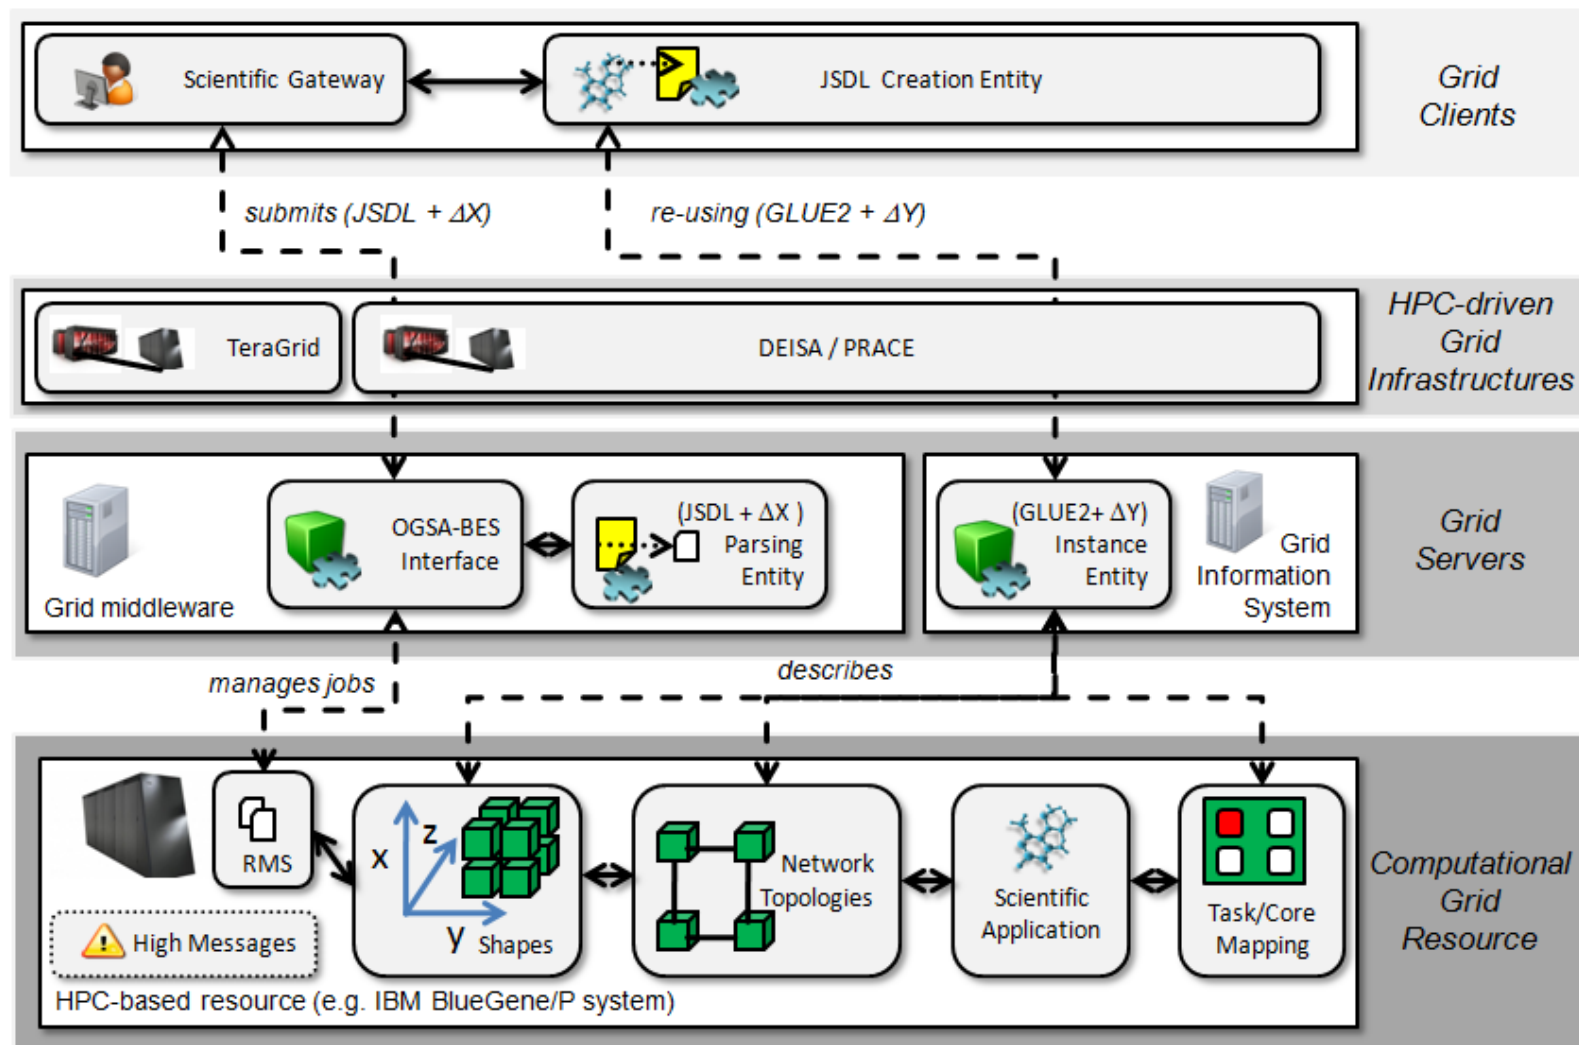


Requirement Collection

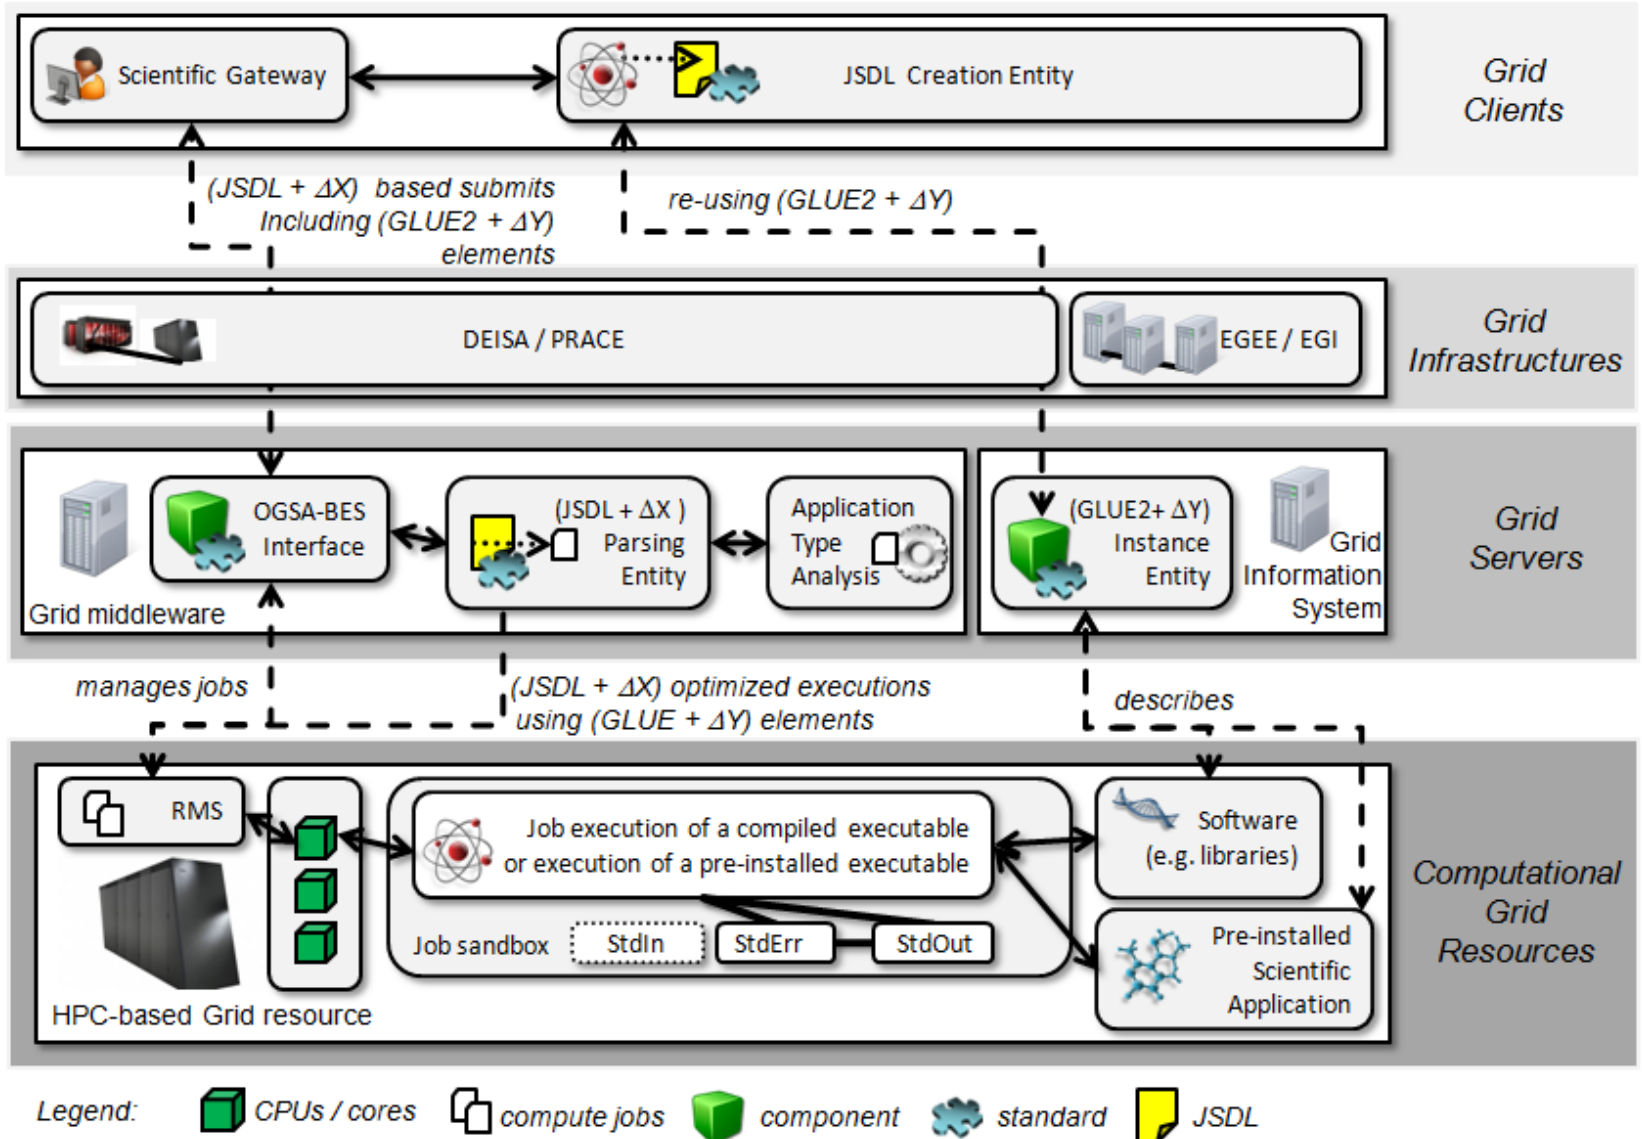
ID	Description	Source	Areas	Quick	Agreement	Expert Go!	MAY / MUST
1	All grid entities (if possible) MUST be described using the GLUE model - if not possible extensions for the GLUE model are necessary	GIN	Information	no - Andrew	yes		MUST
2	For each grid entity which can NOT be described in conformance with GLUE, PGI should propose extensions	GIN	Information	yes	Duplicate 1		
3	GLUE Computing endpoint can describe only one job submission interface (e.g. BES). Job Submission Interfaces which have different properties must be captured by separate GLUE computing endpoints.	Etienne Urbah	Information	No	yes		MUST
4	Each service must publish information regarding its service properties	GIN	Information	yes	yes		MUST
5	The Execution Service must not expose detailed information about the GLUE entities which the Execution Service does not manage (all that are not expressed by the computing part of GLUE), example Storage element GLUE entity not exposed via execution service, no details about storage entity	Etienne Urbah	Information	yes	yes		
6	Instance of Execution Service MUST provide	Strawman	Information	yes	yes		
7	The Execution service instance MUST publish all	Draft Spec	Information,	yes	yes		
8	The vector limits for each of the vector operations	Draft Spec	Information	Yes	yes		
9	If server authentication to clients then it must be	General	Security	no	yes		
10	SSL certificates of servers MUST be signed by a	General	Security	no	OUT		

Over 170 Requirements gathered!

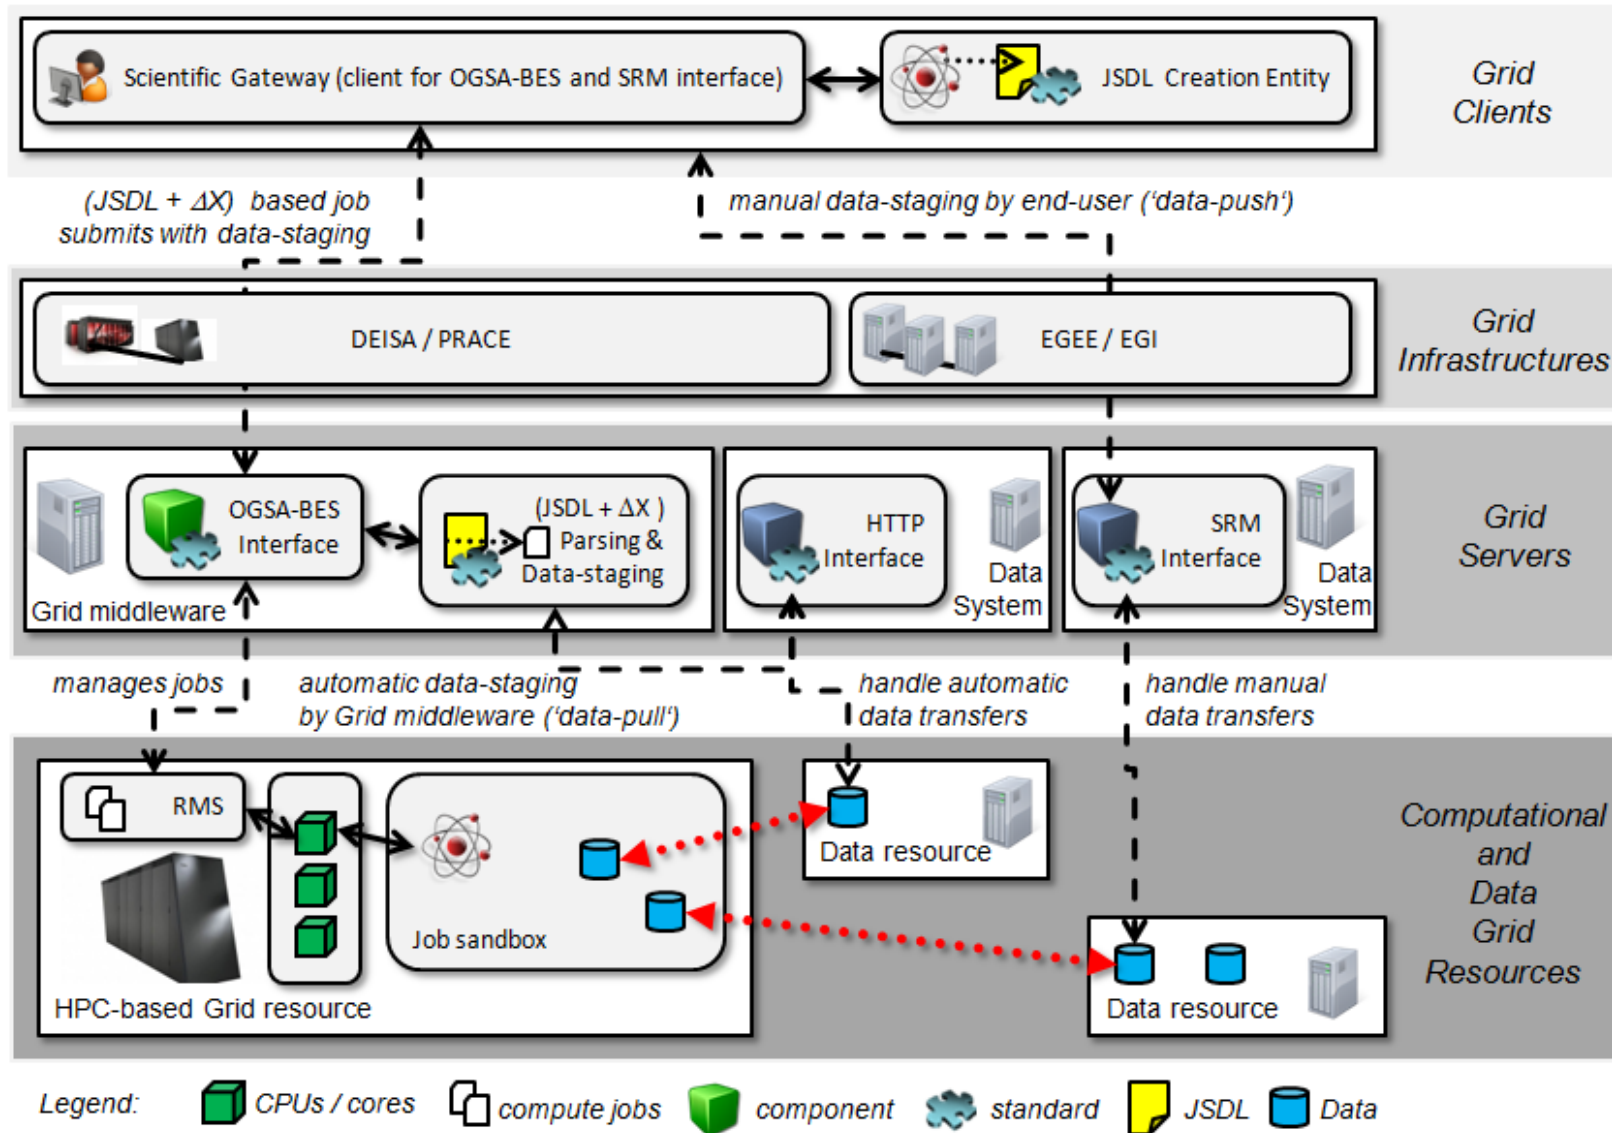
Concept: HPC Extensions



Concept: Better Application Support

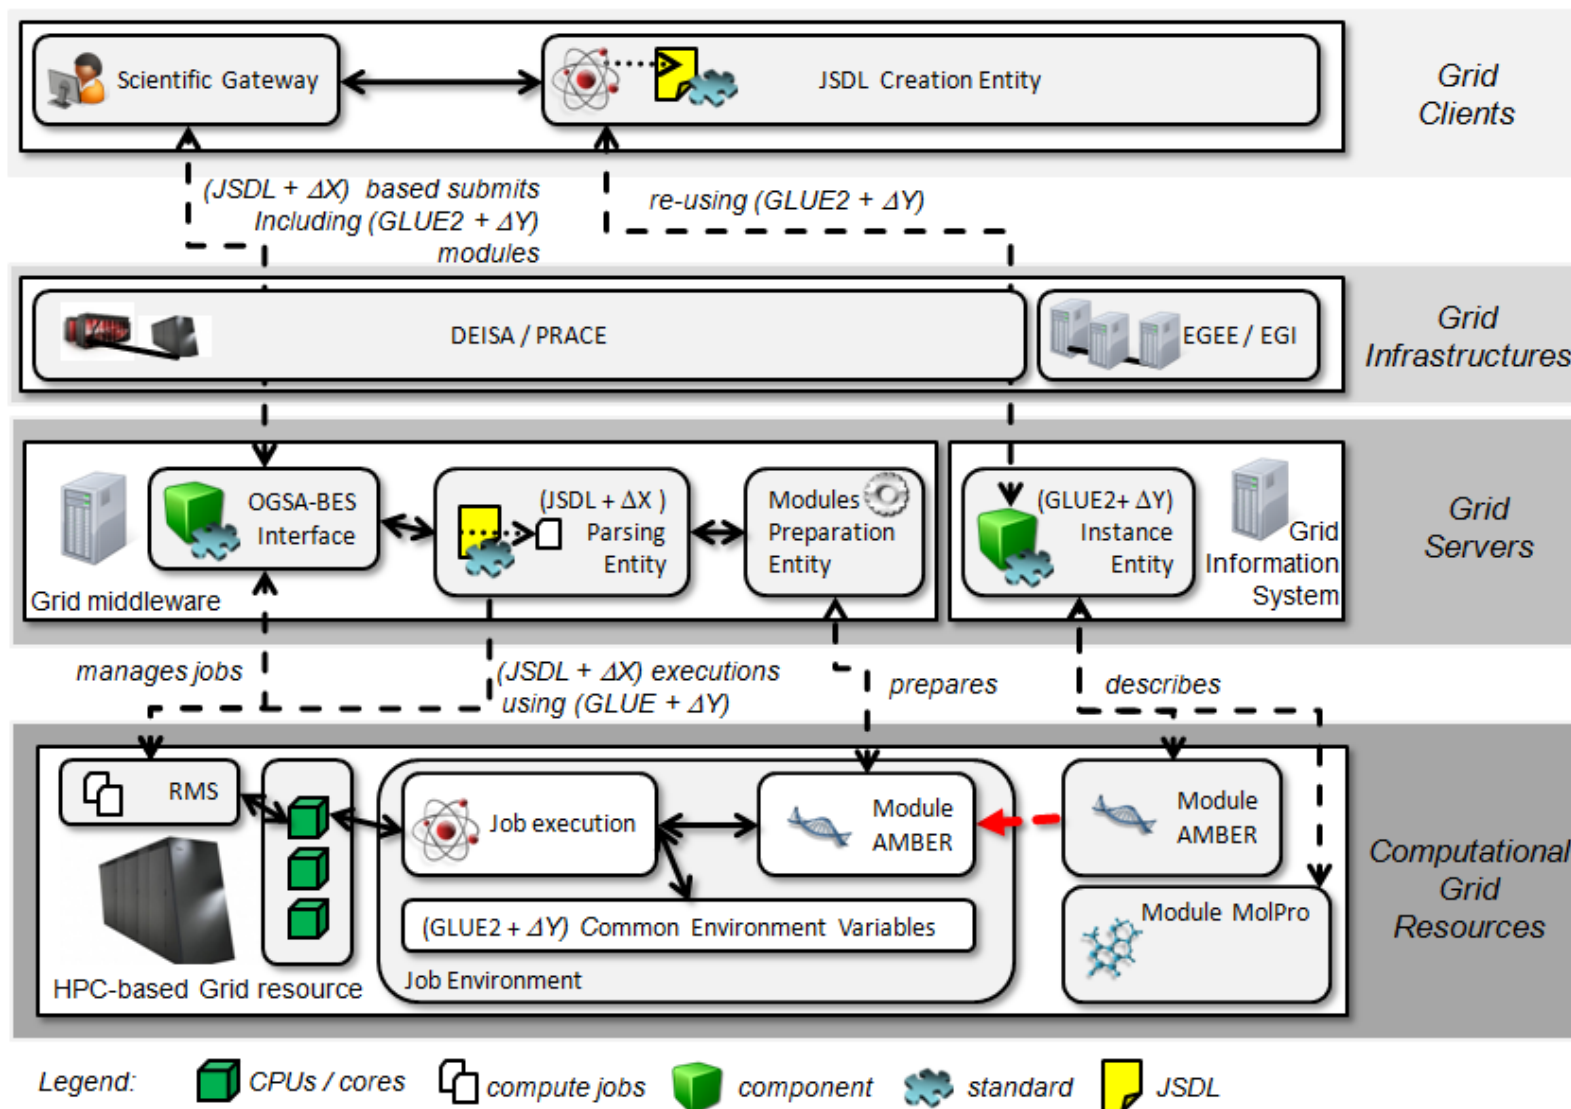


Concept: data-staging flexibility

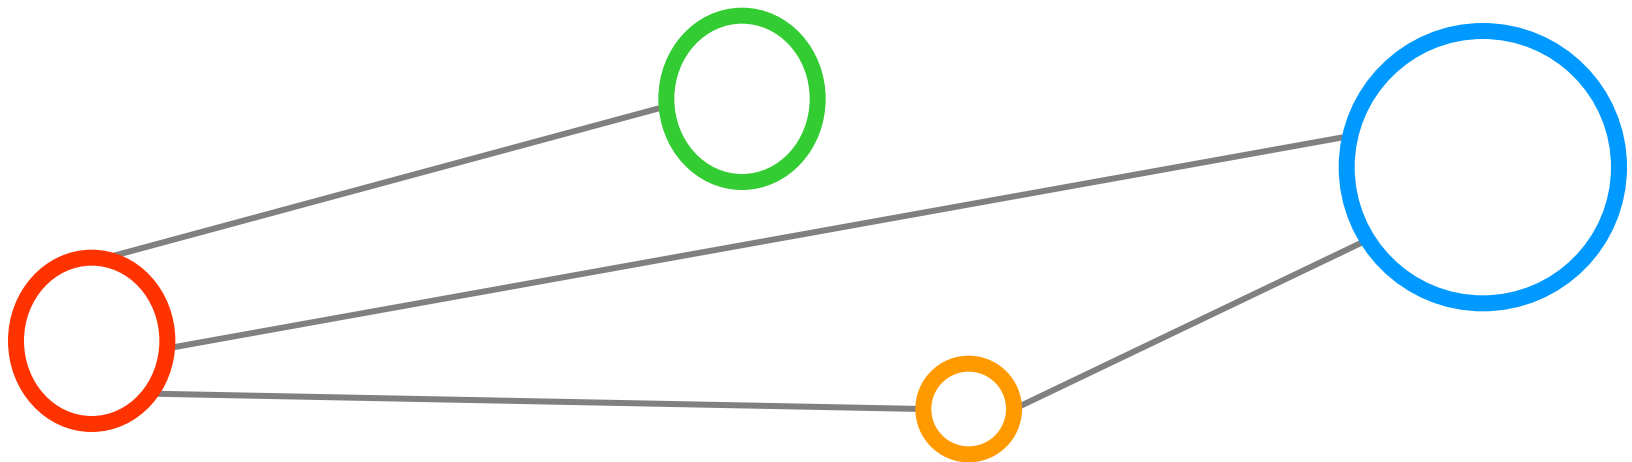




Concept: Execution Adjacencies



DCI Projects in context

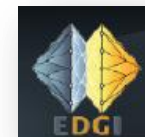


Activities driven by members

- PGI Telcon (mostly) every week...
 - Has been historically only 'best efforts basis work', now much more supported by fundings of EU with DCI projects and SIENA
 - Represents a true productive link between many DCI Project members
 - A technical cooperation complementary to the 'common DCI roadmap idea'
 - Promotes mutual understanding of DCI projects (and its technologies/backgrounds)
- Outcomes (now): Glossary document
 - Significant to understand each other coming from different backgrounds
 - Draft ready for review & OGF editor
- Outcomes (short-term): Use Case collection document
 - Collection reviewed, draft document will be prepared & soon send to OGF editor
- Outcomes (Med-term): PGI specification profiles
 - Very crucial to agree on the set of important requirement (i.e. pareto 80/20 rule)
 - Requirements are gathered, prioritization needed and its process started
 - 'Infrastructure voice' is heard again – EGI use case 'mostly considered important'

Regularly participating members

- SIENA
 - Johannes Watzl (co-chair) – act as the ,neutral‘ position w.r.t. technologies
- EMI – Technology consortium (ARC, gLite, UNICORE, dCache)
 - Morris Riedel (co-chair) – UNICORE/DEISA
 - Alexander Konstantinov – ARC/NorduGrid
 - Oxana Smirnova – ARC/NorduGrid
 - Luigi Zangrando – gLite/EGEE
- EDGI – Desktop Grids
 - Etienne Urbah - EDGI
- IGE – Globus in Europe
 - Steve Crouch – Globus/UK NGS
 - Tim Parkinson – Globus/ UK NGS
 - Emmanouil Paisios / Ilya Saverchenko – Globus/DEISA
- EGI/NGI participation where necessary (i.e. use cases!)
 - Steven Newhouse EGEE/EGI
 - David Wallom UK NGS



International links and cooperation



- Bottom line: 4 out of 6 DCI projects regularly cooperate
 - But PGI also enables some DCI projects international collaborations...
- USA: GENESIS and potential TeraGrid (XD) efforts
 - Andrew Grimshaw
 - Mark Morgan
- Japan: NAREGI / RENKEI
 - Kazushige Saga
- Other initiatives/projects stated interest – but not participate regularly!
- Need to understand whether 'DCI cloud-projects' (StratusLab & VENUS-C) would benefit from actively participating in PGI ...
 - Related methods such as virtualization not in scope of PGI yet – complex enough
- Complementary efforts in the DCI Federation working group of OGF
 - More about Meta-Scheduling in context as complementary efforts to PGI/GIN



DCI Federation Working Group

... a working group in creation with a promising vision

Morris Riedel (on behalf of Alexander Papaspyrou)...

DCI Federation Ideas...

- Infrastructure federation is becoming an increasingly important issue for modern Distributed Computing Infrastructures
- Dynamic elasticity of quasi-static Grid environments,
- Incorporation of special-purpose resources into...
 - Commoditized Cloud infrastructures, existing e-science infrastructures, Grids, etc.
 - Cross-community collaboration for increasingly diverging areas of modern e-Science,
 - Major challenges on the technical level for many resource and middleware providers.
- Especially with respect to increasing cost of operating data centers,
- The intelligent, yet automated and secure sharing of resources is a key factor for success.
- Key topics are meta-scheduling and perhaps Service Level Agreements (SLA) in context
- Success largely will depend on the (DCI) projects and international members participating in this group!

- October 25-29, 2010
- Brussels, Belgium
- Thon Hotel Brussels City Centre
- www.ogf.org/OGF30
- PGI: 3 Sessions
- GIN: 1 Update + 1 Demo Session
- SIENA activities towards roadmap
- Deadline for Session Proposals (workshops, BoFs, group sessions, etc):
September 17, 2010

