Appendix to Relating Requirements and Architectures: A Study of Data-grids

Anthony Finkelstein (a.finkelstein@cs.ucl.ac.uk), Clare Gryce (c.gryce@cs.ucl.ac.uk) and Joe Lewis-Bowen (j.lewis-bowen@cs.ucl.ac.uk)
University College London

Abstract. This is a fragment of the paper "Relating Requirements and Architecture: A Study of Data-Grids" submitted to the Journal of Grid Computing.

Appendix

In this appendix the 83 detailed requirements within the 18 requirements areas discussed in section 3 of the paper are given (with index numbers and priority).

Req. number	A data-grid must, should or could provide the behaviour described here or have this given property.	Priority
Project:	Projects that refer to this requirement.	
Style 1:	Score Reason style 1 supports this requirement.	
Style 2:	-Score Reason style 2 undermines this requirements satisfaction.	

Table I. Key for detailed requirements and scored styles

For each requirement, projects that refer to this requirement are listed. It should be noted that the absence of an individual project does not mean it does not in fact have this requirement - rather that a reference to this requirement was not found in the literature surveyed. Where reference to a requirement was made more generally in the literature, rather than being linked to a particular project, the requirement is said to be 'inferred'. The architectural styles that may influence whether or not the requirement can be met are given. Up to 5 styles are scored for suitability with reasons as described in section

5. The 83 records follow in Table II; the meaning of the fields is given in Table I.

1.1	A data-grid must be able to include distributed, heterogeneous data resources, to form one logical resources that crosses organisational and administrative boundaries.		1
Project:	EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.		
Layer:	2	Layers help abstraction, narrow protocols hide heterogeneity.	
Tier:	2	Tiers help transparency including format heterogeneity and present a single point of entry.	
Peer:	1	Peer networks are often flexible enough to host diverse types of data.	
Agent:	-1	Blackboard must be central for agents to write to (read from), compromising distribution.	
2.1	The users of a data-grid must be able to discover and gain location transparent access to resources.		1
Project:	EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.		
Tier:	2	Tiers provide transparency, and location transparency is fundamental.	
Peer:	2	Peer networks enable discovery and may allow location anonymisation (beyond transparency).	
Agent:	1	Agent based resource discovery should help subsequent lookups (demonstrated by Google).	
3.1	A data-grid must be able to include data of various format and structure.		1
Project:	EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, EGSO.		
Layer:	1	Layers are intendended to abstract diverse formats (for data and signalling).	
Tier:	1	Tier systems may provide framework for type mapping (lining up with OSI presentation layer) as in Corba object marshalling.	
3.2	A data-grid should be able to include raw, processed and annotation data.		2

Project: EDG, PPDG, GriPhyN, BIRN, ESG, MyGrid, EGSO.

3.3	A data-grid should be able to include multiple copies of individual data files and data sets.	2

Project: EDG, PPDG, GriPhyN, NVO, EGSO.

3.4	A data-grid should allow data to be assigned both logical and physical identifiers.	2

Project: EDG, PPDG, GriPhyN.

3.5	A data-grid should enable users to create bespoke logical views on data.	2

Project: EDG, PPDG, GriPhyN, ESG, Astrogrid, EGSO.

3.6	A data-grid could include data stored on tape as well as data stored on disc.	3

Project: EDG, PPDG.

3.7	A data-grid could include multiple copies of a data file/set in different formats, at a single location.	3

Project: EDG, BIRN.

4.1	A data-grid must support domain-specific metadata standards.	1

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

4.2	A data-grid should support a metadata framework that includes multiple metadata schema.	2

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Layer:	1	Layers may allow meta-data abstraction (e.g. diverse lower level to homogenous higher level presentation).
Tier:	1	Tiers could help metadata transparency (though typically rely on their own metadata to provide other tranparencies).

Agent: 1 Agents may be able to crawl across diverse standards of metadata (with individual processes tuned to extract different metadata) as for unreliable html meta-tags.

4.3 A data-grid should support a metadata framework that enables translation between metadata schema. 2

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Layer: 1 If layers can hide heterogeneous metadata, their connecting protocols may form a standard for translation.
Pipe: 1 Filters may be suitable for metadata transformation.
Agent: 1 Agents reviewing diverse metadata content can write to a common format in the shared data-structures of the blackboard.

4.4 A data-grid should support a metadata framework that is flexible. 2

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Layer: 1 If layers provide abstraction, they will support lower layer heterogeneity and therefore fexibility.
Tier: 1 Service networks have demonstrated flexible meta-data (e.g. WSDL).
Peer: 1 Peer infrastructures separate discovery from content, and should therefore allow discovery against arbitrary metadata standards.
Agent: 1 As agents can cope with arbitrary input, they should handle flexible metadata.

4.5 A data-grid should support a metadata framework that is extensible. 2

Project: EDG, PPDG, GriPhyN, ESG, Astrogrid, EGSO.

Layer: 1 If layers provide abstraction, they should allow for changes at the lower levels.
Peer: 1 Peer infrastructures separation of discovery from content should also allow metadata extension.
Agent: 1 As agents can cope with arbitrary input, they should handle changes to metadata.

4.6 A data-grid could enable automatic extraction and generation of metadata. 3

Project: PPDG, GriPhyN, Astrogrid, EGSO.

Peer: 1 Peer networks generate metadata about the network in a decentral way by the way each node is used.
Agent: 2 Agent technology has been most successfully applied for automatic data analysis.

5.1 A data-grid should support queries based on attributes of data (pull). 2

Project: EDG, PPDG, GriPhyN, BIRN, ESG, Astrogrid, EGSO.

Tier: 1 The traditional solution to attribute based queries is is client-server, which evolved into teirs for distributed systems.

5.2 A data-grid should support queries based on pattern matching (push). 2

Project: EDG, PPDG, GriPhyN, BIRN, NVO, Astrogrid, EGSO.

Agent: 1 Agents may generate directory based look-ups / data mining result production, and so push matched patterns as described.

5.3 A data-grid should enable users to construct complex queries, based on atomic query components. 2

Project: EDG, PPDG, GriPhyN, NVO, Astrogrid, MyGrid, EGSO.

5.4 A data-grid could support queries that span multiple data resources. 3

Project: EDG, BIRN, NVO, Astrogrid, EGSO.

Tier: 2 A middle tier is required to fork then join queries for a single client request.
Pipe: 1 Pipelines divide work amonst resources and allow synchronisation points for return.

5.5 A data-grid could support access to data at a level of granularity below that of an individual file. 3

Project: EDG.

Layer: 1 Layers are a suitable way to mark top-level data elements independent of low level storage structure.

Agent: 1 Agents have been used to examine data within files.

5.6 A data-grid could support frequent and rapid data access. 3

Project: EDG.

Tier: 2 Tiered system allow parallel session management and may be used to prevent blocking.

Pipe: 1 If rapid access means that high throughput query pipeline must keep going, pipe and filter decomposition may be preferred to black-box query resolution requires.

6.1 A data-grid must include data processing resources. 1

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

6.2 A data-grid should include a variety of data processing resources. 2

Project: EDG, PPDG, GriPhyN.

Layer: 1 Layered interaction with resources hide variety.

Tier: 1 Tiers should help maximise use of variety.

Pipe: 2 Pipeline scheduling can make best use of a variety of computing resources in non-trivial parallel decomposition.

6.3 A data-grid could support remote code execution. 3

Project: EDG, ESG, NVO, Astrogrid, EGSO.

Tier: 1 Tiers decouple client interaction from server-side activity, and this applies to remote code execution management too.

Peer: 1 Peer networks do use mobile code, though typically tailored for a specific task.

Agent: 1 Agent technology may spread tasks across different platforms as mobile agents.

6.4 A data-grid could include data processing resources that are not local to data or users. 3

Project: EDG, ESG, NVO, Astrogrid, EGSO.

Layer:	1	Virtual network communication provided by layers may help local control of remote resources, hiding intermediary control mechanisms.
Tier:	1	Tiers may also support marshalling between local and non-local tasks.
Peer:	2	Peer networks are successfully used for highly distributed computing tasks.
Agent:	1	Mobile agents may work in unrelated locations.

6.5 A data-grid could support data processing across multiple data resources. 3

Project: PPDG, GriPhyN, NVO.

Layer:	1	Virtual communication paths may also help coordination across data resources.
Tier:	1	A middle tier may support forked execution of tasks spanning data resources.
Pipe:	1	If multiple data resources can be connected in a workflow, pipeline management may help.
Agent:	1	Tasks required for multiple data sets may be divided into agent responsibilities.

6.6 A data-grid could support parallel data processing. 3

Project: EDG, NVO.

Layer:	1	Virtual communication paths should also help coordination across compute resources.
Tier:	1	A middle tier may help to coordinate parallel tasks.
Peer:	2	The peer network topology makes parallel progress a primary operation.
Pipe:	2	Pipeline processing is well established for parallel task execution.
Agent:	1	Agents may work on divided tasks in parallel.

6.7 A data-grid could support lengthy batch processing. 3

Project: ESG, NVO.

Tier:	1	A middle tier may supervise the state of tasks while a user goes offline, queueing requests for batch execution.
Peer:	1	Some tasks on peer networks are massive batch tasks (e.g. SETI at home).
Pipe:	2	Tasks decomposed into a pipeline schedule may successfully be run for very long, discontinuous computation times.
Agent:	1	Agents typically work autonomously.

6.8	A data-grid could include data storage resources that are local to remote processing resources.	3

Project: EDG, EGSO.

Tier:	1	Middleware task coordination could include same-site remote resource use.
Peer:	1	Some peer networks download data with tasks for analysis.

7.1	A data-grid should be able to support the transfer of entire datasets.	2

Project: EDG, PPDG, GriPhyN, NVO.

Layer:	1	The OSI layers help reliable delivery by asigning responsibility for error checking, ordering etc.
Pipe:	1	Parallel pipeline may be used for high-volume data flow.

7.2	A data-grid could support continuous network traffic from data sources to data resource nodes.	3

Project: EDG.

Layer:	2	Layers are essential to uphold quality of service across a network.

8.1	A data-grid should enable access by users in variety of roles.	2

Project: EDG, BIRN, ESG, MyGrid, EGSO.

Layer:	1	Layered abstraction will help keep alternative roles hidden at application from lower implementation.
Tier:	1	Tiers allow abstraction of client types.
Peer:	1	Pier networks allow different node roles.

8.2	A data-grid should enable data selection and querying.	2

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, EGSO.

8.3	A data-grid should enable local visualisation of data.	2

Project: ESG, NVO, EGSO.

8.4 A data-grid should enable browsing of analysis services. 2

Project: EDG, BIRN, NVO, Astrogrid, MyGrid.

Tier: 2 Service based middleware supports capability browsing.
Peer: 1 Peer networks should support service discovery, though this cannot be as reliable as a centralised directory based model.

8.5 A data-grid should enable access to analysis services. 2

Project: EDG, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Tier: 2 Middleware should permit transparent service access.
Peer: 1 Peer networks should be able to forward service requests to nodes that can do work.

8.6 A data-grid should enable users to upload code. 2

Project: NVO, Astrogrid, EGSO.

Layer: 1 Layered abstraction may be applied to mobile code (possibly separating presented specification, parsed bytecode for virtual machine, and the deployed executable).

8.7 A data-grid should enable users to manage data. 2

Project: EDG, PPDG, GriPhyN, ESG.

Peer: 1 Peer networks decentral management means user responsible for their work.

8.8 A data-grid should enable users to manage their accounts. 2

Project: (Inferred)

Tier: 1 The middle-tier may provide a point where accounting can be reliably managed.

8.9 A data-grid should enable users to organise active jobs. 2

Project: (Inferred)

Tier: 1 Middleware services should include task management, which may be exposed to the client.

Peer: 1 Peer networks typically provide handles through which tasks can be identified, which may support user task control in an uncontrolled network.

8.10 A data-grid should offer an interactive and integrated workbench. 2

Project: BIRN, ESG, NVO, MyGrid.

8.12 A data-grid could enable collaborative work between users. 3

Project: BIRN, ESG, Astrogrid.

Tier: 1 A middle tier may coordinate activity between users.
Peer: 2 Peer networks encourage collaboration as control is decentral, and anonymous sharing is enabled.

8.13 A data-grid could enable users to disconnect and leave jobs running. 3

Project: PPDG, GriPhyN.

Tier: 2 Tiers separate clients from back-end activity, enabling back-end state to be maintained independently of user.
Peer: 1 Execution on a remote peer is possible without instigator connection.
Pipe: 2 Pipeline workflow management allows off-line progress.
Agent: 1 Agent progress may be possible without a client (as long as the client does not host blackboard).

9.1 A data-grid should support a range of existing applications and tools. 2

Project: EDG, ESG, Astrogrid, MyGrid, EGSO.

Layer: 1 Layers may hide underlying differences by abstraction.
Tier: 2 Tiers should be able to transparently wrap diverse back end tools.

9.2 A data-grid could allow users to create new applications or tools through an API. 3

Project: EDG, NVO, Astrogrid.

Tier:	1	Middle tier metadata should be flexible enough to add new services.
Peer:	1	Peer networks should allow easy registration of new services.

9.3 A data-grid could allow users to create new applications or tools through composition of existing services. 3

Project: NVO, MyGrid, EGSO.

Layer:	1	Abstraction of implementation to service descriptions may be be helped by layers.
Tier:	1	Middle tier middlewares also provide abstraction and mechanisms for generic service description with their IDLs.
Pipe:	1	Services may be composed in a pipeline description.

9.4 A data-grid could support visualisation tools for browsing data. 3

Project: ESG, NVO, MyGrid, EGSO.

10.1 A data-grid must enable users and administrators to access information about the static state of the system. 1

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Tier:	1	Monitors may relatively easily be included in the middle tier to use static metadata.
Peer:	-1	Peer networks have little static structure, as they should dynamically organise themselves.

10.2 A data-grid must enable users and administrators to access information about the dynamic state of the system. 1

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Tier:	1	Monitors may relatively easily be included in the middle tier to use dynamic data about the network.
Peer:	-1	Peer networks typically hide user activity (due to their application context), but if monitor hooks were included they could only give local information reliably.
Agent:	1	The central blackboard that agents write to may be viewed by adminstrators to determine their current working state.

11.1 A data-grid must enable the management of work over distributed resources. 1

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Tier:	2	The middle-tier exists to manage distrubted systems.
Peer:	1	Peer networks make good use of distributed resources, though central management is not typical.
Pipe:	2	Pipeline scheduling should control tasks on distributed resources (with staging and synchronisation).
Agent:	1	Work may be divided across resources in an agent based architecture.

11.2 A data-grid should enable jobs to be matched to available resources. 2

Project: EDG, PPDG, GriPhyN, MyGrid.

Tier:	2	The middle tier controls dispatch to distributed resources.
Peer:	2	Peer networks may fit resources to requested needs well.
Pipe:	1	Schedulers that use pipeline view may dynamically determine resources used.

11.3 A data-grid should enable jobs to be prioritised. 2

Project: EDG, PPDG, GriPhyN.

Tier:	1	Middle-tier meta-data may include queued task priority for dispatch.
Peer:	1	Priority may be determined by "time to live" style tags on work dispatched to a peer network.
Pipe:	1	Queues at service points in a pipeline may be prioritised.

11.4 A data-grid should enable bottlenecks in the system to be identified and corrected. 2

Project: EDG, PPDG, GriPhyN.

Tier:	2	Middle tier enables progress monitoring and provides control mechanisms to reorganise the network.
Peer:	1	Peer networks may be designed to be free from bottlenecks by sharing and automatically avoid over use as all nodes are servers (but if bottlenecks form they may be hard to identify).
Pipe:	1	The scheduler should avoid bottlenecks based on task decomposition and staging.

11.5	A data-grid could enable re-negotiation of resources for jobs already running.	3
Project:	EDG, PPDG, GriPhyN.	

11.6	A data-grid could enable checkpointing of active jobs.	3
Project:	EDG.	
Agent:	1 Pipelines should provide checkpoints for recovery of flow around point of failure.	

12.1	A data-grid should support intercommunication and interoperation with other grids in related domains.	2
Project:	EDG, PPDG, GriPhyN, NVO, Astrogrid, EGSO.	
Layer:	2 Different levels of abstraction allow mapping to diverse protocols.	
Tier:	2 Tiered networks allow service via a portal to be presented in the same way as intra-system controlled resources.	
Pipe:	1 Pipe and filter may allow transformation from service within one system to another (as compute grids link specialised HPC to CPU farms).	

13.1	A data-grid must support the authentication of users/resources.	1
Project:	EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.	
Tier:	1 Middle tier can organise certification before back end access.	
Peer:	1 Peer networks may insist of signatures within exchanges (though typically the same technology is used for the opposite purpose - anonymisation).	

13.2	A data-grid must support the authorisation of users/resources.	1
Project:	EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.	
Layer:	2 Security layer can police authorisation requirements.	

13.3 A data-grid should support the auditing of actions carried out by system entities. 2

Project: EDG, PPDG, GriPhyN, ESG, Astrogrid, EGSO.

Tier: 1 Middle tier metadata could be used as the basis of an audit trail.
Peer: 1 Some peer networks may generate 'cookie trails' that could be used for auditting.

13.4 A data-grid should enable users/resources to be made accountable for their actions within the system. 2

Project: EDG, PPDG, GriPhyN, ESG, Astrogrid, EGSO.

Tier: 1 Middle tier metadata could be used as the basis for accountability.

13.5 A data-grid should support the enforcement of individual security policies. 2

Project: (Inferred)

Tier: 1 Tiers can help to wrap local heterogenous policies / technologies.
Pipe: -1 Smooth pipeline operation may be inhibitted by diverse methods of boundary negotation.

13.6 A data-grid should support individual security policies that are subject to rapid change. 2

Project: (Inferred)

Pipe: -1 Smooth pipeline operation may be inhibitted by changing methods of boundary negotation.

13.7 A data-grid should support individual security policies that vary in strength and granularity. 2

Project: (Inferred)

Tier: 1 Tiers can help to wrap local heterogenous policies / technologies.

13.8 A data-grid should accommodate the existing security mechanisms of individual resources. 2

Project: (Inferred)

Tier: 1 Tiers can help to wrap existing heterogenous mechanisms.

13.9 A data-grid should support single-step and multi-step sign on. 2

Project: (Inferred)

Tier: 2 Tier decoupling may allow single application side sign event to be associated with multiple back end checks.
Pipe: 1 Workflows may possibly specify sign on as interaction points.

13.10 A data-grid should support mobile users. 2

Project: (Inferred)

Agent: 1 Agents may be mobile, and may therefore support mobile users too (possibly in the role of an agent writing to the shared area).

13.11 A data-grid should allow users to confirm the integrity of data after transfer or processing. 2

Project: (Inferred)

Layer: 1 Layer technology is used to verify integrity of messages.
Tier: 1 Tier management could support independent validation.
Peer: 1 Peer networks typically provide mechanisms that guarentee integrity.

13.12 A data-grid should not employ security mechanisms or processes that significantly reduce the availability of resources. 2

Project: (Inferred)

Layer: -1 Layers add overhead as each part unpacks and verifies messages according to its responsibilities.
Tier: -1 Tier mechanisms add an overhead.

14.1 A data-grid must scale to be able to include 10TB - 1PB new data per year. 1

Project: EDG, PPDG, GriPhyN, BIRN, Astrogrid, EGSO.

Tier:	-1	Management via tiers may introduce bottlenecks that have scaling limits.
Peer:	2	Peer networks scale very well (though typically rely on many small nodes).

14.2	A data-grid must be able to include data files and sets of variable size.	1

Project: EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.

Layer:	-1	Layers typically provide the same types of frame for all traffic, which may add overhead when only small messages pass through a high capacity protocol.
Tier:	1	Tiered networks may separate data from message signal traffic, enabling equal management and control for a range of data scale.
Peer:	1	Peer networks typically provide a way to separate data from its metadata, enabling large resources to be publicised in an equivalent way to small ones.

14.3	A data-grid should be able to include a total volume of data of PB scale.	2

Project: EDG, PPDG, GriPhyN, BIRN, ESG, Astrogrid.

Tier:	1	Seperation of control to higher tiers from back-end storage helps interaction with demanding resources.
Pipe:	1	Handling very large scale data migration may be helped by parallel streaming that may be described in a workflow.

14.4	A data-grid should be able to scale to include new data resources.	2

Project: EDG, PPDG, GriPhyN, ESG, Astrogrid, EGSO.

Peer:	2	Peer networks are designed to grow as data providers are added.

14.5	A data-grid could support up to 10,000 simultaneous processes.	3

Project: (Inferred)

Tier:	0	Though tiers support distributed state, they may also represent a bottleneck that may impede progress of large number of parallel activities.
Peer:	2	Decentral peer networks may make progress on millions of concurrent tasks.
Pipe:	1	Some high performance computing monitors of parallel tasks approach this scale.

15.1 A data-grid should be enable use of resources to be managed for optimum performance. 2

Project: EDG, PPDG, GriPhyN, Astrogrid, MyGrid, EGSO.

Tier:	1	Middleware metadata should provide monitoring information as basis for optimisation by reconfiguring resources.
Pipe:	1	Workflow schedules may be tuned to make optimal use of resources (though this is non-trivial).

15.2 A data-grid could enable a query response time of 5-10s. 3

Project: PPDG, GriPhyN.

Peer:	-1	Peer network topology imply constrained response time cannot be guarenteed.
Pipe:	-1	Pipeline components are best suited for long end-to-end tasks, and therefore not for rapid interaction.
Agent:	-1	Agent based methods are not typically designed for responsiveness, and may take a arbitrarily long time to compose a result.

15.3 A data-grid could support near real-time data processing. 3

Project: EDG, EGSO.

Pipe:	2	A data stream may go directly into pipeline processing (e.g. Regular analysis of data from an instrument), and filters may work to a central clock to ensure synchronise progress.

16.1 The security services of a data-grid should not have a single point of failure. 2

Project: (Inferred)

Tier:	1	Tiers may coordinate shared responsibility, so a validation task may fail over to a redundant resource when the primary service point fails.

Peer: 2 Decentral peer networks are ideal for elastic degradation of service as sub-sets of the network may continue to make progress on node failure.
Agent: -1 If a blackboard were used for any part of the security process, this would be potential single point of failure.

16.2 The data access services of a data-grid should be faulty tolerant to some degree. 2

Project: (Inferred)

Tier: 1 A middle tier may handle transfer to redundant nodes on primary failure to ensure continued data access service.
Peer: 2 Peer networks are highly fault tollerant with respect data routing.

16.3 A data-grid should have capabilities for job recovery in the event of system failure. 2

Project: (Inferred)

Tier: 1 Middleware may coordinate transfer of task state from a failed resource to store or another resource.
Pipe: 1 Workflows may include checkpoints that allow for job recovery.

17.1 A data-grid should allow new functionality or services to the system once deployed. 2

Project: EDG, PPDG, GriPhyN, Astrogrid, MyGrid, EGSO.

Layer: 1 Abstraction provided by layers may facilitate low level extension.
Tier: 1 In a component service network, the middle-tier meta-data descriptions of function and discovery method should scale to include new resources.
Peer: 1 Peer networks typically allow flexible description of service at nodes.

17.2 A data-grid should support the portability of system components local to users and data resources. 2

Project: EDG, PPDG, GriPhyN, Astrogrid, MyGrid, EGSO.

Layer: 2 Portability is greatly helped by layers (hiding hardware or other low-level dependencies from the application).
Tier: 2 Tiers provide transparency, and platform transparency is fundamental.

18.1	A data-grid must allow existing heterogeneous components to be successfully integrated, as necessary.		1
Project:	EDG, PPDG, GriPhyN, BIRN, ESG, NVO, Astrogrid, MyGrid, EGSO.		
Layer:	1	Layered abstraction of low-level platforms enable component integration.	
Tier:	2	A primary aim of the transparency enabled by a middle tier is heterogenous component integration.	
Peer:	1	Peer networks typically integrate heterogenous nodes (which may host hetergeneous components).	
18.2	A data-grid could allow heterogeneous components that are not yet available, to be successfully integrated.		3
Project:	EDG.		
Layer:	1	Layer abstraction also enables integration with future diverse low-level elements.	
Tier:	1	Middle-tiers should enable future integration, but may enforce component responsibilities to allow compatability.	
Peer:	1	Peer network flexibility should extend to future uses.	

Table II. Detailed requirements and scored styles.