



DataGrid

WP1 INPUTS TO THE DATA GRID GRID INFORMATION SERVICE SCHEMA SPECIFICATION



Document identifier: **DataGrid-01-TEN-0104-0_7**

Date: **21/09/2001**

Work package: **WP01**

Partner: **INFN**

Document status **DRAFT**

Deliverable identifier:

Abstract:

This document provides the inputs of the work package 1 (Grid Workload Management) for the specification of the Grid Information Service.



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

Delivery Slip

	Name	Partner	Date	Signature
From	Massimo Sgaravatto	INFN Padova	03/08/2001	
Verified by				

Document Log

Issue	Date	Comment	Author
0_0	03/06/2001	First draft	Massimo Sgaravatto
0_1	09/06/2001	<ul style="list-style-type: none">• Not using anymore Auxiliary object classes;• Definitions of OutboundIP and InboundIP fixed	Massimo Sgaravatto
0_2	25/06/2001	<ul style="list-style-type: none">• Attributes AveragePhysicalMemory, MaxPhysicalMemory, LocalDiskPathname dropped• Attributes MinSPUProcessors, MaxSPUProcessors added• Not using anymore “multiple” (replaced with list of ...”)	Massimo Sgaravatto
0_3	05/07/2001	Attribute FreeSpace for SE dropped	Massimo Sgaravatto
0_4	03/08/2001	<ul style="list-style-type: none">• GOSv2 → GOSv3• New section: DIT structure• Renaming of some attributes• New attributes for mount point and port number	Massimo Sgaravatto
0_5	13/09/2001	<ul style="list-style-type: none">• Replaced “list of ...” with multi-valued• Numeric → integer/cisfloat	Massimo Sgaravatto



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

0_6	19/09/2001	<ul style="list-style-type: none">• Attribute runwindow dropped• Attribute SEId added to SEProtocol objectclass	Massimo Sgaravatto
0_7	21/09/2001	<ul style="list-style-type: none">• Attribute CEId added to CloseStorageElement objectclass	Massimo Sgaravatto

Document Change Record

Issue	Item	Reason for Change

Files

Software Products	User files
Word 2000	wp1-gis-gos.doc
Adobe Acrobat 5.0	wp1-gis-doc.pdf



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

CONTENT

1. INTRODUCTION.....	5
1.1. APPLICABLE DOCUMENTS AND REFERENCE DOCUMENTS	5
1.2. TERMINOLOGY.....	5
2. OVERVIEW	6
3. DIT STRUCTURE.....	7
4. GOS SPECIFICATION.....	12
4.1. COMPUTING ELEMENT	12
4.2. STORAGE ELEMENT	17



1. INTRODUCTION

This document provides the inputs of the work package 1 (Grid Workload Management) for the specification of the Grid Information Service, with particular focus on the first (PM9) release.

1.1. APPLICABLE DOCUMENTS AND REFERENCE DOCUMENTS

Reference documents

[R1] Gregor von Laszewski, Mike Helm, Steven M. Fitzgerald, Pete Vanderbilt, Brett Didier, Peter Lane, Martin Swany, *GOSv3: A Data Definition Language for Grid Information Services*, Gridforum Working Group Document GIS-011-11

[R2] WP1 PM9 Release

(<http://www.infn.it/workload-grid/docs/wp1-pm9.pdf>)

[R3] Job Submission User Interface Man Pages

(http://www.infn.it/workload-grid/docs/DataGrid-01-TEN-0101-0_1-Document.pdf)

[R4]: Job Description Language HowTo

(http://www.infn.it/workload-grid/docs/DataGrid-01-TEN-0102-0_0-Document.pdf)

1.2. TERMINOLOGY

Glossary

CE: Computing Element

GIS: Grid Information Service (aka MDS)

GOS: Grid Object Specification

JDL: Job Description Language

LRMS: Local Resource Management System

MDS: Metacomputing Directory Service (aka GIS)

RB: Resource Broker

SE: Storage Element



2. OVERVIEW

This document provides the inputs of the work package 1 (Grid Workload Management) for the specification of the Grid Information Service, with particular focus on the first (PM9) release.

Information concerning both Computing Elements (CE) and Storage Elements (SE) is presented.

The Resource Broker (RB), being implemented by the WP1, is responsible to choose the CE's where to submit jobs.

As CE, for the PM9 release ([R2]), we consider a Globus resource represented by:

- A queue of an underlying Local Resource Management System (LRMS), such as LSF, PBS, etc..., assuming that this queue represents a set of "homogeneous" resources (that is when a job is submitted to a specific queue, it doesn't matter in which node of this queue the job is dispatched)
- A "single" node, that doesn't rely on a LRMS (using the fork system call)

When a user submits a job to the workload management system ([R3]), he/she specifies the required and preferred resources for this job (using a Job Description Language (JDL), [R4]), which are then "matched" by the RB with the characteristics and status of the available CE's, information that must be available in the MDS.

Storage Elements are used to store data. They can be accessed, from a CE, using a "grid protocol", or via a local access (e.g. using the Posix "open" system call).

This document doesn't aim to provide an exhaustive list of attributes for these two Grid services, but instead represent the information that we consider necessary for the implementation of the WP1 services (for example for sure information concerning storage elements is missing, but we feel that these attributes should be better identified and defined by the WP's who will use them).

Section 3 introduces this schema, showing the proposed LDAP DN structure considering an example.

In section 4 the proposed schema is specified using the GOS (Grid Object Specification) v3 language [R1], the data definition language for Grid Information Services defined in the context of the Information Services working group of the Global Grid Forum: section 4.1 defines the schema for Computing Elements, while section 4.2 specifies the information for Storage Elements.



3. DIT STRUCTURE

Each GRID host can have different CE's and different SE's¹.

In the proposed DIT structure, both CE's and SE's are "under" the host they belong to, and at the same level, as in the following example (in this example for this host 2 CE's and 1 SE are considered):

dn:ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid01, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid

objectname: ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid01, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid

objectClass: DataGridTop

objectClass: ComputingElement

CEId: lxde01.pd.infn.it:2119/jobmanager-lsf-grid01

GlobusResourceContactString: lxde01.pd.infn.it:2119/jobmanager-lsf

GRAMVersion: 1.71

Architecture: Intel

OpSys: Linux RH6.1

MinPhysicalMemory: 128

MinLocalDiskSpace: 100

TotalCPUs: 4

FreeCPUs: 0

NumSMPs: 1

MinSPUProcessors: 2

MaxSPUProcessors: 2

TotalJobs: 15

RunningJobs: 4

IdleJobs: 11

MaxTotalJobs: 1000

MaxRunningJobs: 1000

WorstTraversalTime: 2502

EstimatedTraversalTime: 131

Acrive: True

Priority: 30

MaxCPUTime:

MaxWallClockTime:

AverageSI00: 23

MinSI00: 13

MaxSI00:30

¹ Actually for the PM9 release just one SE per host is foreseen



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

*AuthorizedUser: "/C=IT/O=INFN/L=Padova/CN=Massimo
Sgaravatto/Email=massimo.sgaravatto@pd.infn.it"*

*AuthorizedUser: "/C=IT/O=INFN/OU=Personal Certificate/L=Padova/CN=Marco
Verlato/Email=Marco.Verlato@padova.infn.it"*

RunTimeEnvironment: CMS3.2

RunTimeEnvironment: EO4.2

AFSAvailabe: True

OutboundIP: True

InboundIP: False

QueueName: grid01

LRMSType: LSF

LRMSVersion: 4.0

*dn: closesese=lxde01.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid01, hn=lxde01.pd.infn.it,
dc=pd, dc=infn, dc=it, o=Grid*

objectname: closesese=lxde01.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid01,

hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid

objectClass: DataGridTop

objectClass: CloseStorageElement

CloseSE: lxde01.pd.infn.it

CEId: lxde01.pd.infn.it:2119/jobmanager-lsf-grid01

MountPoint: /disk1

*dn: closesese=se1.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid01, hn=lxde01.pd.infn.it,
dc=pd, dc=infn, dc=it, o=Grid*

objectname: closesese=se1.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid01,

hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid

objectClass: DataGridTop

objectClass: CloseStorageElement

CloseSE: se1.pd.infn.it

CEId: lxde01.pd.infn.it:2119/jobmanager-lsf-grid01

*dn: ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid02, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it,
o=Grid*

*objectname: ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid02, hn=lxde01.pd.infn.it, dc=pd, dc=infn,
dc=it, o=Grid*

objectClass: DataGridTop

objectClass: ComputingElement

CEId: lxde01.pd.infn.it:2119/jobmanager-lsf-grid02

GlobusResourceContactString: lxde01.pd.infn.it:2119/jobmanager-lsf



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

GRAMVersion: 1.71
Architecture: Intel
OpSys: Linux RH6.1
MinPhysicalMemory: 128
MinLocalDiskSpace: 150
TotalCPUs: 5
FreeCPUs: 3
NumSMPs: 0
MinSPUProcessors: 0
MaxSPUProcessors: 0
TotalJobs: 3
RunningJobs: 3
IdleJobs: 0
MaxTotalJobs: 2000
MaxRunningJobs: 1000
WorstTraversalTime: 1530
EstimatedTraversalTime: 0
Acrive: True
Priority: 20
MaxCPUTime:
MaxWallClockTime:
AverageSI00: 20
MinSI00: 20
MaxSI00: 20
*AuthorizedUser: "/C=IT/O=INFN/L=Milano/CN=Francesco
Prelz/Email=francesco.prelz@mi.infn.it"*
AuthorizedUser: "/O=CESNET/O=Masaryk University/CN=Miroslav Ruda
*AuthorizedUser: "/C=IT/O=INFN/L=Padova/CN=Massimo
Sgaravatto/Email=massimo.sgaravatto@pd.infn.it"*
RunTimeEnvironment: ATLAS2.2
RunTimeEnvironment: BIO3.4
AFSAvailabe: True
OutboundIP: True
InboundIP: False
QueueName: grid02
LRMSType: LSF
LRMSVersion: 4.0



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

dn: closesese=lxde01.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid02, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
objectname: closesese=lxde01.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid02, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
objectClass: DataGridTop
objectClass: CloseStorageElement
CloseSE: lxde01.pd.infn.it
CEId: lxde01.pd.infn.it:2119/jobmanager-lsf-grid02
MountPoint: /disk1

dn: closesese=se1.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid02, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
objectname: closesese=se1.pd.infn.it, ce=lxde01.pd.infn.it:2119/jobmanager-lsf-grid02, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
objectClass: DataGridTop
objectClass: CloseStorageElement
CloseSE: se1.pd.infn.it
CEId: lxde01.pd.infn.it:2119/jobmanager-lsf-grid02

dn: se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
*objectname: se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, *
o=Grid
objectClass: DataGridTop
objectClass: StorageElement
SEId: lxde01.pd.infn.it
CloseCE: lxde01.pd.infn.it:2119/jobmanager-lsf-grid01
CloseCE: lxde01.pd.infn.it:2119/jobmanager-lsf-grid02
CloseCE: cex.pd.infn.it:2119/jobmanager-pbs-queuex

dn: seprotocol=gridftp, se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
*objectname: seprotocol=gridftp, se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, *
o=Grid
objectClass: DataGridTop
objectClass: StorageElementProtocol
SEprotocol: gridftp
SEId: lxde01.pd.infn.it
Port: 4444

dn: seprotocol=rftio, se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
*objectname: seprotocol=rftio, se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, *
o=Grid
objectClass: DataGridTop
objectClass: StorageElementProtocol



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

SEprotocol: rfio

SEId: lxde01.pd.infn.it

Port: 5555

*dn: seprotocol=file, se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, o=Grid
objectname: seprotocol=file, se=lxde01.pd.infn.it, hn=lxde01.pd.infn.it, dc=pd, dc=infn, dc=it, *

o=Grid

objectClass: DataGridTop

objectClass: StorageElementProtocol

SEprotocol: file



4. GOS SPECIFICATION

4.1. COMPUTING ELEMENT

In the following GOS specification, the *ComputingElement* object class is considered to specify a CE:

```
NAMESPACE DataGrid {
    OBJECTCLASS ComputingElement {
        OID {??}
        DESCRIPTION {
            A Computing Element for the DataGrid Project
        }
        KIND {STRUCTURAL}
        INHERITSFROM {DataGridTop}
        KEY {CEId}
        MUST CONTAIN {
            CEId :: single-valued, cis,
                {The identifier of the CE}
            GlobusResourceContactString :: single-valued, cis,
                {The Globus resource contact string}
            GRAMVersion :: single-valued, cis,
                {The GRAM version}
            Architecture :: single-valued, cis,
                {The architecture of the hosts composing the CE}
            OpSys :: single-valued, cis,
                {The operating system of the hosts composing the CE}
            MinPhysicalMemory :: single-valued, cisfloat,
                {The minimum value of the physical memory among the hosts associated to
                the CE}
            MinLocalDiskSpace :: single-valued, cisfloat,
                {The minimum local disk footprint}
            TotalCPUs :: single-valued, integer,
                {The number of total processors associated to the CE}
            FreeCPUs :: single-valued, integer,
                {The number of free processors}
        }
    }
}
```



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: **21/09/2001**

NumSMPs :: single-valued, integer,
{The number of SMP hosts}

MinSPUProcessors :: single-valued, integer,
{The minimum number of SPU processors (for SMP hosts)}

MaxSPUProcessors :: single-valued, integer,
{The maximum number of SPU processors (for SMP hosts)}

TotalJobs :: single-valued, integer,
{The number of jobs submitted to the CE}

RunningJobs :: single-valued, integer,
{The number of currently running jobs submitted to the CE}

IdleJobs :: single-valued, integer,
{The number of idle jobs submitted to the CE}

MaxTotalJobs :: single-valued, integer,
{The maximum number of jobs (running and idle) allowed for the CE}

MaxRunningJobs :: single-valued, integer,
{The maximum number of running jobs allowed for the CE}

WorstTraversalTime :: single-valued, cisfloat,
{The worst traversal time for jobs submitted to the CE}

EstimatedTraversalTime :: single-valued, cisfloat,
{The estimated traversal time for jobs submitted to the CE}

Active :: single, boolean,
{ Defines if the CE is active}

Priority :: single-valued, integer,
{The priority associated to the CE}

MaxCPUTime :: single-valued, cisfloat,
{The maximum CPU time allowed for jobs submitted to the CE}

MaxWallClockTime :: single-valued, cisfloat,
{The maximum wall clock time allowed for jobs submitted to the CE}

AverageSI00 :: single-valued, cisfloat,
{The average of the SpecInt2000 benchmark of the nodes associated to the CE}

MinSI00 :: single-valued, cisfloat,
{The minimum value of the SpecInt2000 benchmark of the nodes associated to the CE}

MaxSI00 :: single-valued, cisfloat,
{The maximum value of the SpecInt2000 benchmark of the nodes associated to the CE}



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

```
    AuthorizedUser :: multi-valued, cis,
        {The list of subjects of X509 user certificates for users authorized to submit
         jobs to the CE}
    RunTimeEnvironment :: multi-valued, cis,
        {List of softwares/packages installed on this CE}
    AFSAvailable :: single-valued, boolean,
        {Defines if AFS is installed}
    OutboundIP :: single-valued, boolean,
        {Defines if outbound connectivity is allowed}
    InboundIP :: single-valued, boolean,
        {Defines if inbound connectivity is allowed}
}
MAY CONTAIN {
    QueueName :: single-valued, cis,
        {The name of the queue in the LRMS}
    LRMSType :: single-valued, cis,
        {The type of LRMS}
    LRMSVersion :: single-valued, cis,
        {The version of LRMS};
}
}

OBJECTCLASS CloseStorageElement {
    OID {??}
    DESCRIPTION {
        A Storage Element close to a Computing Element
    }
    KIND {STRUCTURAL}
    INHERITSFROM {DataGridTop}
    MUST CONTAIN {
        CloseSE :: single-valued, cis,
            {The identifier of the SE}
        CEId :: single-valued, cis,
            {The identifier of the CE}

    }
    MAY CONTAIN {
        MountPoint:: single-valued, cis,
```



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

```
        {The mount point of this SE from this CE};  
    }  
}  
}
```

CEId is a string, univocally identifying the CE published in the Grid Information Space. We assume that WP4 will choose how to define this value, and provide the Grid Information Space with this appropriate value (i.e. this value could be obtained “combining” the *GlobusResourceContactString* and *QueueName* attributes).

GlobusResourceContactString represents the Globus resource contact string that identifies this Globus resource (e.g. pcgrid01.pd.infn.it:2119/jobmanager-lsf).

GRAMVersion is the GRAM version.

Architecture defines the architecture of the hosts “composing” this CE (we assume that all these hosts have the same architecture).

OpSys is the operating system of the hosts “associated” to this CE (we assume that all these hosts run the same operating system).

MinPhysicalMemory is the minimum value of the physical memory, expressed in Mbytes, among the hosts associated to this Computing Element. If the CE is a “single host”, this value represents its actual physical memory.

MinLocalDiskSpace represents the minimum local disk footprint (that is the “working directory” where the job computation will take place) available to a running job running on a worker node (expressed in Mbytes). If more than one node is associated to this CE, we assume that all these worker nodes make available the same local disk space. It is also assumed that this advertised local disk footprint is actually available to a running job, even in case that more than one process is running on a given “worker” node.

TotalCPUs represents the number of total CPUs associated to the CE.

FreeCPUs is the number of free processors associated to the Computing Element (that is processors able to run, in that moment, jobs submitted to this CE).

NumSMPs defines the number of SMP hosts associated to this Computing Element.

MinSPUProcessors is the minimum number of SPU processors (for SMP hosts).

MaxSPUProcessors represents the maximum number of SPU processors (for SMP hosts).



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

TotalJobs represents the number of jobs submitted to the Computing Element, jobs that have not already been completed (running and idle jobs).

RunningJobs is the number of currently running jobs submitted to the Computing Element.

IdleJobs defines the number of idle jobs submitted to this Computing Element (jobs that are not running since they are waiting for available resources).

MaxTotalJobs is the maximum number of jobs (running and idle) allowed for the Computing Element.

MaxRunningJobs defines the maximum number of running jobs allowed for the Computing Element.

WorstTraversalTime defines the worst traversal time (in seconds) for jobs submitted to the Computing Element.

EstimatedTraversalTime represents the estimated traversal time (in seconds) for jobs submitted to the Computing Element.

For the PM9 release this value is obtained as scaled value of the last traversal time.

Active is a boolean attribute that defines if the Computing Element is active (for example if the CE is a queue, if it ready to dispatch jobs to the worker nodes).

The *Priority* attribute defines the priority associated to this computing element. (Quite useless, unless all system managers in the different sites agree on same policies in defining this value. What do we do ??)

MaxCpuTime defines the maximum CPU time (in seconds) allowed for jobs submitted to the Computing Element.

MaxWallClockTime is the maximum wall clock time (in seconds) allowed for jobs submitted to the Computing Element.

AverageSI00 is the average of the SpecInt2000 benchmark of the nodes associated to this Computing Element. If the CE is a "single processor", this value represents its actual performance.

MinSI00 is the minimum value of the SpecInt2000 benchmark among the processors associated to this Computing Element. If the CE is a "single processor", this value represents its actual performance.



WP1 INPUTS TO THE DATAGRID GRID INFORMATION SERVICE SCHEMA SPECIFICATION

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

MaxSI00 is the maximum value of the SpecInt2000 benchmark among the processors associated to this Computing Element. If the CE is a “single processor”, this value represents its actual performance.

AuthorizedUser is the list of subjects of X509 user certificates, representing users authorized to submit job to the Computing Element.

RunTimeEnvironment is a tag defining a list of run time environments/packages/software installed on the Computing Element. In case, the version of this package/environment is included in this string.

AFSAvailable is a boolean attribute defining if AFS is installed on the Computing Element.

OutboundIP indicates if outbound connectivity is allowed (e.g. all the worker nodes associated to the Computing Element can “initiate” a data transfer, sending and/or receiving data to/from a remote Internet node).

InboundIP indicates if inbound connectivity is allowed (e.g. a remote Internet node can “initiate” a data transfer, sending and/or receive data to/from any worker node associated to the Computing Element).

The following attributes (*QueueName*, *LRMType* and *LRMSVersion*) are defined only when the Computing Element is a queue of a Local Resource Management System:

QueueName defines the name of the queue in the LRMS.

LRMSType specifies the type of Local Resource Management System (e.g. Condor, LSF, PBS, Condor, ...)

LRMSVersion is the version of the LRMS.

CloseSE is the identifier (the hostname) of the SE close to the considered CE.

MountPoint is the mount point (defined only if “local access” is supported) for this SE, from the considered CE.

4.2. STORAGE ELEMENT

The *StorageElement* object class is used to define information regarding a Storage Element:



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

```
NAMESPACE DataGrid {
  OBJECTCLASS StorageElement {
    OID {??}
    DESCRIPTION {
      A Storage Element for the DataGrid Project
    }
  }
  KIND {STRUCTURAL}
  INHERITSFROM {DataGridTop}
  KEY {SEId}
  MUST CONTAIN {
    SEId :: single-valued, cis,
      {The identifier of the SE (the hostname for the PM9 release)}
    CloseCE: multi-valued, cis;
      {The list of identifiers of CE's close to the SE};
  }
}
OBJECTCLASS StorageElementProtocol {
  OID {??}
  DESCRIPTION {
    A protocol spoken by a SE
  }
  KIND {STRUCTURAL}
  INHERITSFROM {DataGridTop}
  MUST CONTAIN {
    SEProtocol :: single-valued, cis,
      {The name of the protocol}
    SEId: single-valued, cis,
      {The identifier of the SE (the hostname for the PM9 release)}
  }
  MAY CONTAIN {
    Port :: single-valued, integer,
      {The port number};
  }
}
}
```



**WP1 INPUTS TO THE DATAGRID GRID
INFORMATION SERVICE SCHEMA
SPECIFICATION**

Doc. Identifier:
DataGrid-01-TEN-0104-0_7

Date: 21/09/2001

SEId is the key attribute. It univocally identifies the Storage Element. It is the hostname for the PM9 release

CloseCE is the list of strings that univocally identify the Computing Elements close enough to this Storage Element (this is the *CEId* attribute of the *ComputingElement* object class: see section 4.1).

SEProtocol is the access protocol for the storage element (e.g. GridFtp, RFIO, etc...).

Port is the port number associated to the considered protocol.