

HIGH PERFORMANCE COMPUTING (HPC) MANAGING OPENPBS WITH CIM

A JOB SCHEDULING ENVIRONMENT

IN COOPERATION WITH THE UNIVERSITY OF ILLINOIS (NCSA)

Document Version : 1.6
Date Last Changed : January 27, 2004

Document Owner : Heidi Neumann
Department : LTC Systems Management
e-mail : heidineu@de.ibm.com
Web Page : oss.software.ibm.com/developerworks/projects/sblim

(C) Copyright IBM Corp. 2003

Master Copy is located on PC of document owner

1 DOCUMENT CONTROL.....	3
1.1 CHANGE HISTORY	3
1.2 REVIEWER / CONTRIBUTOR LIST	3
2 INTRODUCTION.....	4
2.1 NCSA ENVIRONMENT.....	4
2.2 DEFINITION OF THE SCHEMA.....	5
2.3 PREREQUISITES.....	5
3 OPENPBS SCHEMA / GLOBAL VIEW	6
3.1 OPENPBS CIM MODEL.....	6
3.2 OPENPBS SCHEMA / CLASSES	7
<i>Linux_OpenPBSComputeNodeStatData : CIM_StatisticalData</i>	7
<i>Linux_OpenPBSCpuStatData : CIM_StatisticalData</i>	8
<i>Linux_OpenPBSJobStatusStatData : CIM_StatisticalData</i>	8
<i>Linux_OpenPBSBSCComputeNodeStatData : Linux_OpenPBSComputeNodeStatData</i>	8
<i>Linux_OpenPBSBSCpuStatData : Linux_OpenPBSCpuStatData</i>	8
<i>Linux_OpenPBSBBSJobStatusStatData : Linux_OpenPBSJobStatusStatData</i>	8
<i>Linux_OpenPBSBatchServiceStatData : CIM_StatisticalData</i>	9
<i>Linux_OpenPBSBatchService : CIM_BatchService</i>	9
<i>Linux_OpenPBSGroupOfSystems : CIM_ConcreteCollection</i>	9
<i>Linux_OpenPBSGOSComputeNodeStatData : Linux_OpenPBSComputeNodeStatData</i>	9
<i>Linux_OpenPBSGOSCpuStatData : Linux_OpenPBSCpuStatData</i>	9
<i>Linux_OpenPBSGOSJobStatusStatData : Linux_OpenPBSJobStatusStatData</i>	10
<i>Linux_OpenPBSJobQueue : CIM_JobQueue</i>	10
<i>Linux_OpenPBSJQComputeNodeStatData : Linux_OpenPBSComputeNodeStatData</i>	10
<i>Linux_OpenPBSJQCpuStatData : Linux_OpenPBSCpuStatData</i>	10
<i>Linux_OpenPBSQueueStatData : CIM_QueueStatisticalData</i>	10
<i>Linux_OpenPBSBatchJob : CIM_BatchJob</i>	11
<i>Linux_OpenPBSComputeNode : CIM_ComputerSystem</i>	11
3.3 OPENPBS SCHEMA / ASSOCIATIONS.....	12
<i>Linux_OpenPBSComputeNodeInGroupOfSystems</i>	12
<i>Linux_OpenPBSBatchJobInQueue</i>	13
<i>Linux_OpenPBSComputeNodeCandidateOfJobQueue</i>	13
<i>Linux_OpenPBSExecutionRequirement</i>	13
<i>Linux_OpenPBSExecutionCandidate</i>	14
<i>Linux_OpenPBSExecutionTarget</i>	14
<i>Linux_OpenPBSQueueForBatchService</i>	14
<i>Linux_OpenPBSJobAtBatchService</i>	14
<i>Linux_OpenPBSComputeNodeForBatchService</i>	15
4 COMPUTE NODE / LOCAL VIEW	16

1 Document Control

1.1 Change History

Revision	Date	Description
1.0	May 2002	Initial Draft
1.1	June 2002	Updated class and property descriptions
1.2	July 2002	Updated property descriptions
1.3	Sept 2003	Adoption of JSIM Schema to old OpenPBS model
1.4	Oct 2003	Updated Descriptions and pictures
1.5	Nov 2003	Added association ComputeNodeCandidateOfJobQueue and property descriptions to class ComputeNodeStatData and CpuStatData
1.6	Jan 2004	Renamed OpenPBS specific classes from LinuxOpenPBS_xxx to Linux_OpenPBSxxx

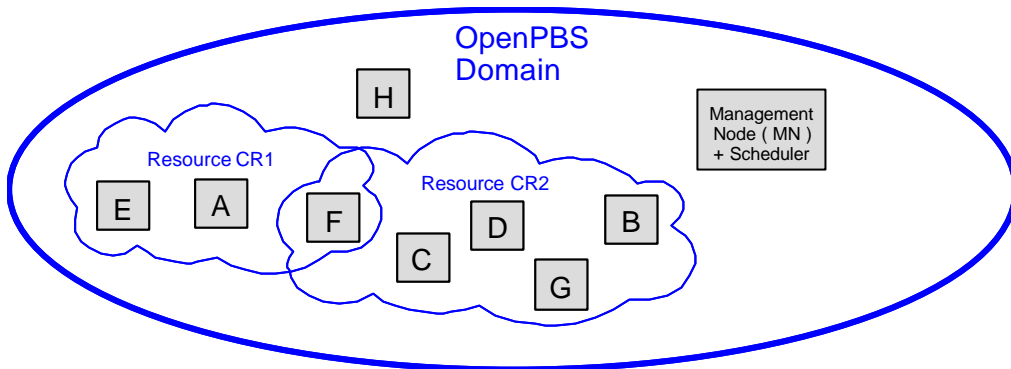
1.2 Reviewer / Contributor List

Name	e-mail	Role & Responsibility
Peter Enstrom	enstrom@ncsa.uiuc.edu	Developer of the OpenPBS Instrumentation
Joshi Fullop	jfullop@ncsa.uiuc.edu	Developer of the first set of OpenPBS Instrumentation
Adrian Schuur	schuur@de.ibm.com	LTC Team Member
Viktor Mihajlovski	mihajlov@de.ibm.com	LTC Team Member

2 Introduction

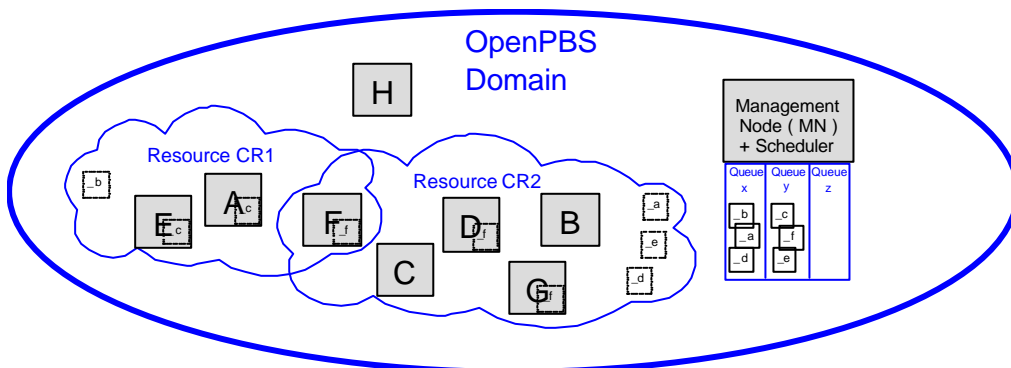
2.1 NCSA Environment

This chapter describes the NCSA OpenPBS environment. NCSA hosts a High Performance Computing Cluster, managed by OpenPBS (<http://www.openpbs.org/>) with Maui as external Job Scheduler. The intention of our cooperation is the definition and implementation of a CIM Schema for this environment.



2.1 Cluster Topology (logical objects are marked blue, physical objects are marked grey)

Picture 2.1 gives an overview of the main OpenPBS parts. OpenPBS defines a domain like logical entity, which is responsible for computations, consuming an high amount of resources. This domain contains a defined set of Hosts (A, B, ... H, MN). One Host is dedicated to be the Management Node (MN). The Hosts are logical grouped in so called Resources (CR1, CR2). A single Host can participate in more than one Resource. The MN does not participate in any Resource. But the MN can be used for computations, if all Resources are in use. Theoretically one Host can execute more than one Jobs at a certain time. In the case of the University of Illinois Cluster this is not supported. That means one Host can execute zero or one Job at one point in time.



2.2 Job Scheduling Environment

A Job ($_a, _b, \dots, _f$) is a unit of work submitted to OpenPBS for computation. The Job is submitted to a specific Queue (x, y, z) and one or more Resources (CR1, CR2). The Job contains information under which specific conditions the computation has to run. The Job stays as long in its Queue as the

computation is not finished or the Job killed. Once a Job is started, the Job runs on a list of Hosts. The list of Hosts is composed of the Resource member(s), the job was submitted to. A Resource can compute more than one Jobs at one point in time.

OpenPBS has knowledge about Jobs, Queues, Resources and Hosts and is responsible for the execution of the Jobs within their specified conditions. OpenPBS knows only its own domain.

2.2 Definition of the Schema

While the definition of an OpenPBS Schema in general is independent of a specific naming schema, we use CIM (Common Information Model) as underlying technology. CIM is an implementation-neutral schema to offer overall management information in a network / enterprise environment and an accepted industry standard defined by the Distributed Management Task Force (DMTF).

The schema definition is divided in two important parts. The first one represents the global view the Management Node / Management Server has onto the domain (chapter 2). An implementation of this schema is only located on the Management Node (MN). The second part describes the Base Instrumentation of each single node participating in the domain (chapter 3).

To offer each user of OpenPBS special views with all necessary and important information, the objects (classes) in the schema need a well defined set of attributes. The different views itself are realised by using different subsets of information. An example can be : An administrator needs a extensive view onto the whole Cluster, while a manager only wants to know how many Jobs are currently scheduled for computation. In the case of the view for managers an application enumerates all Jobs and uses only this subset of information. If an administrator needs a complete overview, the enumeration of Jobs is only one step besides many others. The enumeration of all Resources, Hosts and Queues can be another step. Associations define the relations between the instances of this classes, e.g. which Hosts are defined for Resource CR1. But the application works with exactly the same classes for each view.

2.3 Prerequisites

There exist several activities to define a common Job Submission Schema in CIM. One is located within the Global Grid Forum. The work group "CIM Based Grid Schema" defines a "Job Submission Information Model" (<http://forge.gridforum.org/projects/cgs-wg/>). The here described OpenPBS model is based on the JSIM Schema.

3.2 OpenPBS Schema / Classes

The blue marked properties are key properties of the classes. These properties are required and serve as unique identifier of the instances.

Linux_OpenPBSBatchService is an aggregation point to hold the information about the whole status of the domain, ComputeNodes, CPUs and the Jobs.

First there need to be defined three common classes to hold ComputeNode, CPU and Job related statistical properties. The definition of base classes is usefull, as these three classes are the parent classes for all ComputeNode, CPU and Job related statistical data in the whole OpenPBS schema.

Linux_OpenPBSComputeNodeStatData : CIM_StatisticalData

is base class for the host specific statistical information.

NAME	TYPE	DESCRIPTION
TotalNumber	uint32	The total number of physical compute nodes regardless of their state.
NumberAvailable	uint32	The number of running, usable compute nodes that are not currently running jobs.
NumberAssigned	uint32	The number of running, usable compute nodes, that are currently running jobs.
MaxNumberAssignable	uint32	The largest job that could be run based on the current overall state. It is the minimum of NumberAvailable and the maximum job size allowed by policy (which is the same as TotalNumber by default).
MaxNumberCapable	uint32	The largest job that could be run assuming enough hosts were available. It is the minimum of MaxNumber and the maximum job size allowed by policy (which is the same as TotalNumber by default).
MaxNumber	uint32	The number of physical compute nodes excluding host that are currently down or unusable.
MinNumber	uint32	The min job size allowed by policy .

The following example demonstrates the meaning of the properties for physical compute nodes of the OpenPBS domain. The domain knows 10 compute nodes, where

- node 1 is dead and won't boot
- node 2 functional but is being used for testing and is not available for jobs
- node 3 and 4 available for jobs and idle
- node 5 till 10 currently running jobs

A policy allows jobs to run on a maximum of three nodes and a minimum of 2 nodes. Then the values for the properties above are :

- TotalNumber 10 all nodes
- NumberAvailable 2 nodes 3 and 4
- NumberAssigned 6 nodes 5 till 10
- MaxNumberAssignable 2 only 2 nodes are free
- MaxNumberCapable 3 job size is limited to 3
- MinNumber 2 jobs must use at least 2 nodes to run
- MaxNumber 8 nodes 1 and 2 not usable

Linux_OpenPBSCpuStatData : CIM_StatisticalData

is base class for the CPU specific statistical information.

NAME	TYPE	DESCRIPTION
TotalNumber	uint32	The total number of CPUs regardless of their state.
NumberAvailable	uint32	The number of running, usable CPUs that are not currently running jobs.
NumberAssigned	uint32	The number of running, usable CPUs, that are currently running jobs.
MaxNumberAssignable	uint32	The largest job that could be run based on the current overall state. It is the minimum of NumberAvailable and the maximum job size allowed by policy (which is the same as TotalNumber by default).
MaxNumberCapable	uint32	The largest job that could be run assuming enough CPUs were available. It is the minimum of MaxNumber and the maximum job size allowed by policy (which is the same as TotalNumber by default).
MaxNumber	uint32	The number of CPUs excluding CPUs that are disabled or unusable.
MinNumber	uint32	The min job size allowed by policy.

Linux_OpenPBSJobStatusStatData : CIM_StatisticalData

is base class for the job specific statistical information.

NAME	TYPE	DESCRIPTION
TotalNumberOfJobs	uint32	Represents the total number of Jobs (status : running / pending) submitted to the batch service OpenPBS.
QueuedJobs	uint32	Represents the number of Jobs, currently in status 'queued'.
HeldJobs	uint32	Represents the number of Jobs, currently in status 'held'.
RunningJobs	uint32	Represents the number of Jobs, currently in status 'running'.
ExitingJobs	uint32	Represents the number of Jobs, currently in status 'exiting'.

Linux_OpenPBSBSComputeNodeStatData : Linux_OpenPBS ComputeNodeStatData

contains the host specific statistical information within the scope of OpenPBS. An enumeration of this class returns one single instance.

Linux_OpenPBSBSCpuStatData : Linux_OpenPBS CpuStatData

contains the CPU specific statistical information within the scope of OpenPBS. An enumeration of this class returns one single instance.

Linux_OpenPBSBSJobStatusStatData : Linux_OpenPBSJobStatusStatData

contains the job specific statistical information within the scope of OpenPBS. An enumeration of this class returns one single instance.

Linux_OpenPBSBatchServiceStatData : CIM_StatisticalData

contains the OpenPBS specific statistical information. An enumeration of this class returns one single instance.

NAME	TYPE	DESCRIPTION
PhysicalMemory	uint32	The sum of physical memory of all Hosts.
TotalMemory	uint32	The sum of memory available on the Hosts for execution.
Load	uint32	The aggregated Load on the OpenPBS domain.
PercentLoad	uint64	The percentage load over all Hosts.

Linux_OpenPBSBatchService : CIM_BatchService

represents the OpenPBS functionality. OpenPBS is a service on the MN. An enumeration of this class returns one single instance as there is only one OpenPBS service responsible for the domain.

NAME	TYPE	DESCRIPTION
SystemCreationClassName	string	Value : 'Linux_ComputerSystem'
SystemName	string	Contains the host name 'host.domain' value of the system acting as Management Node (MN).
CreationClassName	string	Value : 'Linux_OpenPBSBatchService' (class name)
Name	string	Unique name of the OpenPBS Service, e.g. 'host.domain'
ElementName	string	Contains the type of the batch service; Value : 'OpenPBS'
Started	boolean	Indicates if the OpenPBS Service is started or not.

Linux_OpenPBSGroupOfSystems : CIM_ConcreteCollection

represents the OpenPBS Resources. An enumeration of this class returns the same amount of instances as Resources are defined for this OpenPBS domain. A Resource is a static definition of a node group.

NAME	TYPE	DESCRIPTION
InstanceID	string	Contains the unique name of the OpenPBS Resource.

Linux_OpenPBSGOSComputeNodeStatData : Linux_OpenPBSComputeNodeStatData

contains the host specific statistical information within the scope of a certain group of systems (Resource). An enumeration of this class returns as many instances as instances of the class GroupOfSystems are reported.

Linux_OpenPBSGOSCpuStatData : Linux_OpenPBSCpuStatData

contains the CPU specific statistical information within the scope of a certain group of systems (Resource). An enumeration of this class returns as many instances as instances of the class GroupOfSystems are reported.

Linux_OpenPBSGOSJobStatusStatData : Linux_OpenPBS JobStatusStatData

contains the job specific statistical information within the scope of a certain group of systems (Resource). An enumeration of this class returns as many instances as instances of the class GroupOfSystems are reported.

Linux_OpenPBSJobQueue : CIM_JobQueue

represents the Job Queues. An enumeration of this class returns the same amount of instances as Queues are defined for this OpenPBS domain.

NAME	TYPE	DESCRIPTION
SystemCreationClassName	string	Value : 'Linux_ComputerSystem'
SystemName	string	Contains the host name 'host.domain' value of the system acting as Management Node (MN).
CreationClassName	string	Value : 'Linux_OpenPBSJobQueue' (class name)
Name	string	Unique name of the Queue, e.g. 'all'
Type	string	Indicates the type of the queue. Values : 'Execution', 'Route'
Started	boolean	The Queue is active and can schedule Jobs
MaxWallTime	uint64	The max time available to Jobs submitted to this Queue.
MinWallTime	uint64	The min time available to Jobs submitted to this Queue.
WallTime	uint64	The default max time assigned to Jobs submitted to this Queue.

Linux_OpenPBSJQComputeNodeStatData : Linux_OpenPBS ComputeNodeStatData

contains the host specific statistical information within the scope of a certain queue. An enumeration of this class returns as many instances as instances of the class JobQueue are reported.

Linux_OpenPBSJQCpuStatData : Linux_OpenPBS CpuStatData

contains the CPU specific statistical information within the scope of a certain queue. An enumeration of this class returns as many instances as instances of the class JobQueue are reported.

Linux_OpenPBSQueueStatData : CIM_QueueStatisticalData

contains the queue specific statistical information. An enumeration of this class returns as many instances as instances of the class JobQueue are reported.

NAME	TYPE	DESCRIPTION
TotalNumberOfJobs	uint32	Represents the total number of Jobs (status : running / pending) submitted to the batch service OpenPBS.
QueuedJobs	uint32	Represents the number of Jobs, currently in status 'queued'.
HeldJobs	uint32	Represents the number of Jobs, currently in status 'held'.
RunningJobs	uint32	Represents the number of Jobs, currently in status 'running'.
WaitingJobs	uint32	Represents the number of Jobs that are waiting to be put into the 'running' state.
ExitingJobs	uint32	Represents the number of Jobs, currently in status 'exiting'.

Linux_OpenPBSBatchJob : CIM_BatchJob

represents the Jobs. An enumeration of this class returns the same amount of instances as Jobs are submitted to OpenPBS for computation.

NAME	TYPE	DESCRIPTION
InstanceID	string	Unique ID to identify the Job : Format : "jobname": "unique name of the scheduler service"
Name	string	contains the name of the Job, e.g. 'Job1'
Owner	string	The user who submitted the Job.
Group	string	Specifies the group of the submitting user.
TimeSubmitted	datetime	The time when the Job was submitted to the Queue.
StartTime	datetime	The time when the Job was started.
ElapsedTime	datetime	Execution time of the Job.
UntilTime	datetime	Time after which the Job is invalid or should be stopped.
CPUTimeUsed	uint64	The CPU time used in seconds.
MemoryUsed	uint64	Amount of physical memory used by the job.
TotalMemoryUsed	uint64	Physical and Virtual memory used by the job.
WallTimeUsed	uint64	Number of seconds that the job has been running.
MaxWallTime	uint64	Total number of seconds the job will be allowed to run.
MaxCPUTime	uint64	Total amount of CPU time the job will be able to use.
JobStatus	string	Represents the status of the Job. Values : E ("Exiting"), H ("Held"), Q ("Queued"), R ("Running"), S ("Suspend"), T ("Transiting"), W ("Waiting")
Checkpoint	string	If supported by the server and the host OS, this property determines when checkpointing will be performed by OpenPBS on behalf of the job.
ErrorPath	string	The final path name for the file containing the job's standard error stream. Format : [hostname:]pathname
OutputPath	string	The final path name for the file containing the job's standard output stream. Format : [hostname:]pathname
Priority	uint32	Indicates the importance for the execution of the Job assigned by the user.
Rerunable	boolean	The job is rerunable. This flag is given by the user at submission.
SessionID	string	If the job is running, this property is set to the session id of the first executing task.
Task	string	The list of environment variables passed with the batch request.
NumberOfNodes	uint16	The number of compute nodes the job has requested to run on.
ProcsPerNode	uint32	The number of processors the job has requested per node.

Linux_OpenPBSComputeNode : CIM_ComputerSystem

represents the Hosts registered to OpenPBS. An enumeration of this class returns the same amount of instances as Hosts have been configured in OpenPBS.

NAME	TYPE	DESCRIPTION
CreationClassName	string	Property derived from CIM_System; contains the name of the class -> 'Linux_OpenPBSComputeNode'
Name	string	contains the unique name of the machine -> value 'host.domain'
NumberOfCPUs	uint16	Represents the total number of CPUs of the local system.
MaxCPUSpeed	uint32	Represents the max CPU speed of the local system.
OSType	uint16	see description and value map of CIM_OperatingSystem.OSType

OSVersion	string	free form string to describe the distribution, kernel version ...
-----------	--------	---

3.3 OpenPBS Schema / Associations

Associations reflect the relations between classes and their instances. An association has a left and a right 'end', see picture 3.1 'OpenPBS Schema'. These ends can be compared with pointers to classes. An association is called with one end specified (an instance of the referred class), to figure out which instances are on the opposite end.

The first line describes the classes between the following association(s) is (are) defined. The second line describes the name of the derived association and the parent association class.

ASSOCIATION NAME	PARENT CLASS	REFERENCED CLASS A	REFERENCED CLASS B
Linux_OpenPBS QueueForBatchService	CIM_QueueForBatchService	BatchService	JobQueue
Linux_OpenPBS JobAtBatchService	CIM_Dependency	BatchService	BatchJob
Linux_OpenPBS ComputeNodeForBatchService	CIM_Dependency	BatchService	ComputeNode
Linux_OpenPBS ExecutionCandidate	CIM_ExecutionCandidate	ComputeNode	BatchJob
Linux_OpenPBS ExecutionTarget	CIM_ExecutionTarget	ComputeNode	BatchJob
Linux_OpenPBS ExecutionRequirement	CIM_ExecutionRequirement	GroupOfSystems	BatchJob
Linux_OpenPBS ComputeNodeInGroupOfSystems	CIM_MemberOfCollection	GroupOfSystems	ComputeNode
Linux_OpenPBS BatchJobInQueue	CIM_JobDestinationJob	JobQueue	BatchJob
Linux_OpenPBS ComputeNodeCandidateOfJobQueue	CIM_Dependency	JobQueue	ComputeNode
Linux_OpenPBS StatDataForBatchService	CIM_ElementStatisticalData	BatchService	CIM_StatisticalData
Linux_OpenPBS StatDataForJobQueue	CIM_ElementStatisticalData	JobQueue	CIM_StatisticalData
Linux_OpenPBS StatDataForGroupOfSystems	CIM_ElementStatisticalData	GroupOfSystems	CIM_StatisticalData

For a better understanding of the intention of associations, now some examples based on the pictures 2.2 and 3.1 .

Linux_OpenPBSComputeNodeInGroupOfSystems

reflects the relation between the compute nodes and the resources. If this association is called from the Resource end (instance of class Linux_OpenPBSGroupOfSystems specified) these instances of Linux_OpenPBSComputeNode are returned, which have been defined to be part of this Resource. If this association is called from the Host end (instance of Linux_OpenPBSComputeNode specified) these instances of Linux_OpenPBSGroupOfSystems are returned, where the Host is defined as member.

specified end	returned instance
Linux_OpenPBSGroupOfSystems CR1	Linux_OpenPBSComputeNode A E F
Linux_OpenPBSGroupOfSystems CR2	Linux_OpenPBSComputeNode C D G B F
Linux_OpenPBSComputeNode A	Linux_OpenPBSGroupOfSystems CR1
Linux_OpenPBSComputeNode F	Linux_OpenPBSGroupOfSystems CR1 CR2

Linux_OpenPBSBatchJobInQueue

reflects the relation between the Jobs and the Queues. If this association is called from the Job end (instance of class Linux_OpenPBSBatchJob specified) this instance of Linux_OpenPBSJobQueue is returned, where the Job was submitted to. If you call this association from the Queue end (instance of Linux_OpenPBSJobQueue specified) these instances of Linux_OpenPBSBatchJob are returned, which have been submitted to this Queue for execution. This is independent of the current execution state of the Jobs.

specified end		returned instance	
Linux_OpenPBSBatchJob	_d	Linux_OpenPBSJobQueue	x
Linux_OpenPBSBatchJob	_e	Linux_OpenPBSJobQueue	y
Linux_OpenPBSJobQueue	x	Linux_OpenPBSBatchJob	_b_a_d
Linux_OpenPBSJobQueue	z	Linux_OpenPBSBatchJob	-

Linux_OpenPBSComputeNodeCandidateOfJobQueue

reflects the relation between a Queue and the Hosts available to it for computations. If this association is called from the Queue end (instance of class Linux_OpenPBSJobQueue specified), these instances of Linux_OpenPBSComputeNode are returned, which are listed by OpenPBS's acl_hosts list. If nothing is specified in this list for the queue, all available compute node instances are returned. If this association is called from the Host end (instance of Linux_OpenPBSComputeNode specified), these instances of Linux_OpenPBSJobQueue are returned, where the host is defined in the queue's acl_hosts list or its acl_hosts list is empty (default : all hosts).

The table explains the following configuration of acl_hosts list for queue x and y. The acl_hosts list of queue x is empty. The acl_hosts list of queue y defines the hosts A, B and F.

specified end		returned instance	
Linux_OpenPBSJobQueue	x	Linux_OpenPBSComputeNode	A B C D E F G
Linux_OpenPBSJobQueue	y	Linux_OpenPBSComputeNode	A B F
Linux_OpenPBSComputeNode A		Linux_OpenPBSBatchJob	x y
Linux_OpenPBSComputeNode D		Linux_OpenPBSBatchJob	x
Linux_OpenPBSComputeNode C		Linux_OpenPBSBatchJob	x

Linux_OpenPBSExecutionRequirement

reflects the relation between a Job and the Resource, used for the execution of this Job. Instances of this association can, but may not exist. The information on which Resource a Job has to be executed is submitted with the Job definition to OpenPBS. If this association is called from the Job end (instance of class Linux_OpenPBSBatchJob specified) this instance of Linux_OpenPBSGroupOfSystems is returned, where the Job was submitted to. If this association is called from the Resource end (instance of Linux_OpenPBSGroupOfSystems specified) these instances of Linux_OpenPBSBatchJob are returned, which have been submitted to this Resource for execution. This is independent of the current execution state of the Job(s).

specified end		returned instance	
Linux_OpenPBSBatchJob	_f	Linux_OpenPBSGroupOfSystems	CR2
Linux_OpenPBSBatchJob	_e	Linux_OpenPBSGroupOfSystems	CR1
Linux_OpenPBSGroupOfSystems CR1		Linux_OpenPBSBatchJob	_b_c
Linux_OpenPBSGroupOfSystems CR2		Linux_OpenPBSBatchJob	_f

Linux_OpenPBSExecutionCandidate

reflects the relation between a scheduled Job and the Hosts available for a later computation. If this association is called from the Job end (instance of class Linux_OpenPBSBatchJob specified) and the Job is currently executed, these instances of Linux_OpenPBSComputeNode are returned, which are listed by the 'Host Execution List'. If the Job is not running, no instances will be returned. If this association is called from the Host end (instance of Linux_OpenPBSComputeNode specified) the instance of Linux_OpenPBSBatchJob is returned, which is currently computed on this Host. If the Host is not in use, no instance will be returned.

specified end		returned instance	
Linux_OpenPBSBatchJob	_f	Linux_OpenPBSComputeNode	-
Linux_OpenPBSBatchJob	_b	Linux_OpenPBSComputeNode	E A F
Linux_OpenPBSComputeNode	A	Linux_OpenPBSBatchJob	_b
Linux_OpenPBSComputeNode	D	Linux_OpenPBSBatchJob	_a _d _e
Linux_OpenPBSComputeNode	C	Linux_OpenPBSBatchJob	_a _d _e

Linux_OpenPBSExecutionTarget

reflects the relation between the running Job and the Hosts used for its computation. If this association is called from the Job end (instance of class Linux_OpenPBSBatchJob specified) and the Job is currently executed, these instances of Linux_OpenPBSComputeNode are returned, which are listed by the 'Host Execution List'. If the Job is not running, no instances will be returned. If this association is called from the Host end (instance of Linux_OpenPBSComputeNode specified) the instance of Linux_OpenPBSBatchJob is returned, which is currently computed on this Host. If the Host is not in use, no instance will be returned.

specified end		returned instance	
Linux_OpenPBSBatchJob	_f	Linux_OpenPBSComputeNode	D F G
Linux_OpenPBSBatchJob	_b	Linux_OpenPBSComputeNode	-
Linux_OpenPBSComputeNode	A	Linux_OpenPBSBatchJob	_c
Linux_OpenPBSComputeNode	D	Linux_OpenPBSBatchJob	_f
Linux_OpenPBSComputeNode	C	Linux_OpenPBSBatchJob	-

Linux_OpenPBSQueueForBatchService

reflects the relation between the OpenPBS service and the Queues to manage. If this association is called from the OpenPBS service end (instance of class Linux_OpenPBSBatchService specified) all instances of Linux_OpenPBSJobQueue are returned. If this association is called from the Queue end (instance of Linux_OpenPBSJobQueue specified) the one instance of Linux_OpenPBSBatchService is returned.

Linux_OpenPBSJobAtBatchService

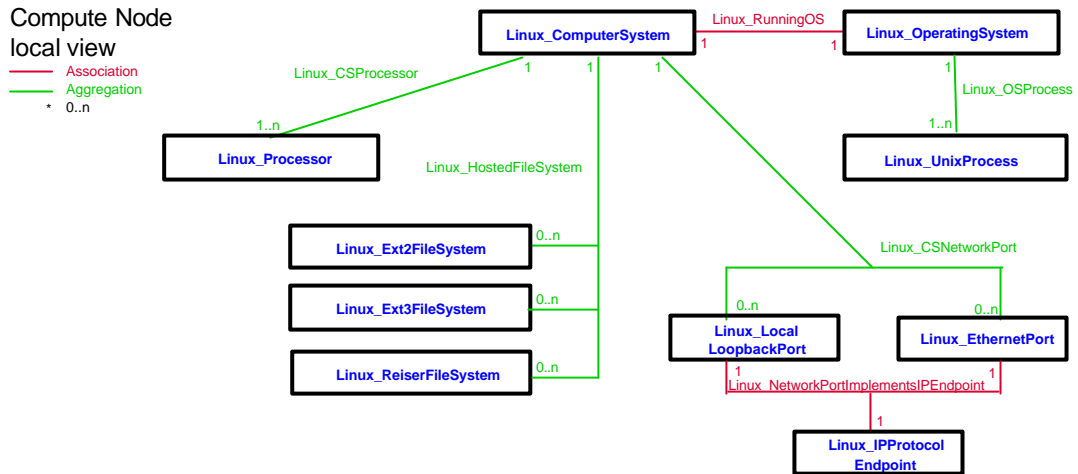
reflects the relation between the OpenPBS service and the Jobs to manage. If this association is called from the batch service end (instance of class Linux_OpenPBSBatchService specified) all instances of Linux_OpenPBSBatchJob are returned, independent of the Queue they have been submitted to. If this association is called from the Job end (instance of Linux_OpenPBSBatchJob specified) the one instance of Linux_OpenPBSBatchService is returned.

Linux_OpenPBSComputeNodeForBatchService

reflects the relation between the OpenPBS Service and the Host(s) registered for computations. If this association is called from the OpenPBS service end (instance of class Linux_OpenPBSBatchService specified) all instances of Linux_OpenPBSComputeNode are returned. If this association is called from the Host end (instance of Linux_OpenPBSComputeNode specified) the one instance of Linux_OpenPBSBatchService is returned.

4 Compute Node / Local View

This schema reflects the OS Base Instrumentation of each single node participating in the domain and will be described in this chapter at a later point in time.



4.1 Compute Node / local view