

Moab Cluster Manager

Introducing Moab Cluster Manager - all the power and flexibility you have known with Maui Scheduler and more, but with the ease of use to get the most out of your system.



One of the centerpieces of Moab Cluster Scheduler is its associated graphical user interface. The interface not only improves the ease of using the scheduler, but also adds additional instant diagnosis capabilities as the status and details of the resources running on your system are presented to you in new and instantaneous forms.

Main Interface



CLUSTER RESOURCES INC.™

Center for HPC Cluster Resource Management and Scheduling, Copyright © 2000-2003,

Easy Set Up:

New set up wizards, accelerate the deployment process.



Configuration and Operational Management:

Control operational mode and status, system settings, log settings, simulation settings, optimization and other settings.



User Job/Workload Management

Submit jobs with all advanced settings at your fingertips. Graphical User interface significantly reduces chances of user error while fully empowering them. Administrators can empower users to view and manage the status of their own jobs.



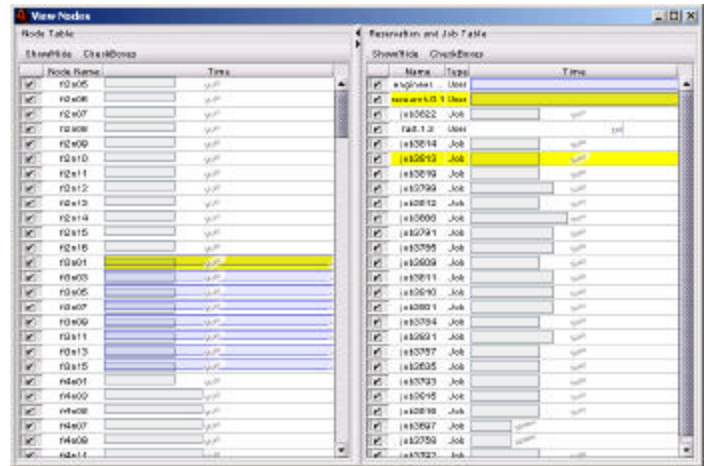
Job Name	Status	User	Group	Start Time	Submission Time	Duration	Procs	Nodes	Details
job3227	Idle	me	admins	10:24:30 AM 1/1/70	00:07:30	4	4		
job3793	Running	me	admins	10:41:10 ...	00:07:30	4	4		
job3799	Idle	me	admins	10:39:40 ...	00:07:30	4	4		
job3997	Running	me	admins	10:39:40 ...	00:07:30	4	4		
job3975	Idle	me	admins	10:04:41 AM 1/1/70	00:07:30	22	22		
job3980	Running	me	admins	10:28:10 ...	00:41:40	22	22		
job3990	Idle	me	admins	10:04:32 AM 1/1/70	00:07:30	8	8		
job3935	Running	me	admins	10:41:10 ...	00:07:30	4	4		

CLUSTER RESOURCES INC.™

Center for HPC Cluster Resource Management and Scheduling, Copyright © 2000-2003,

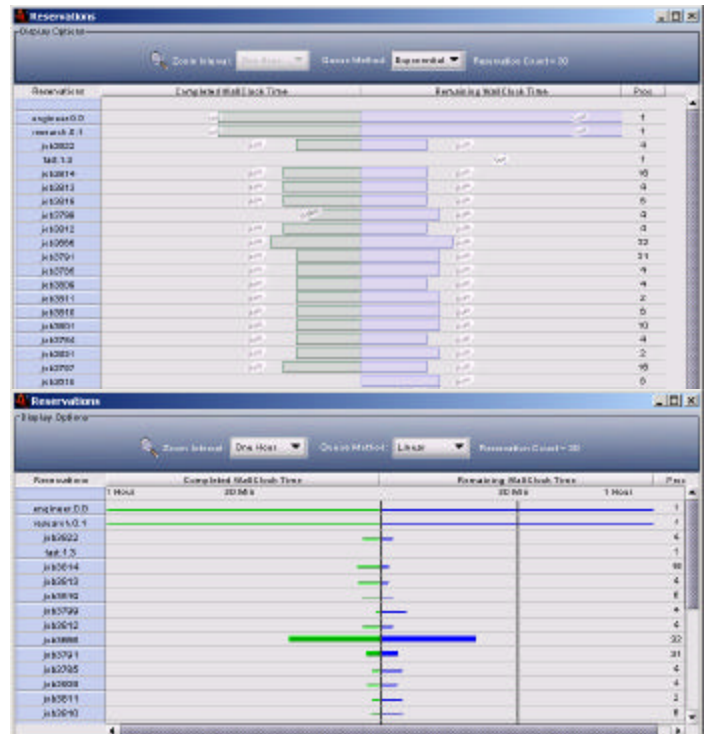
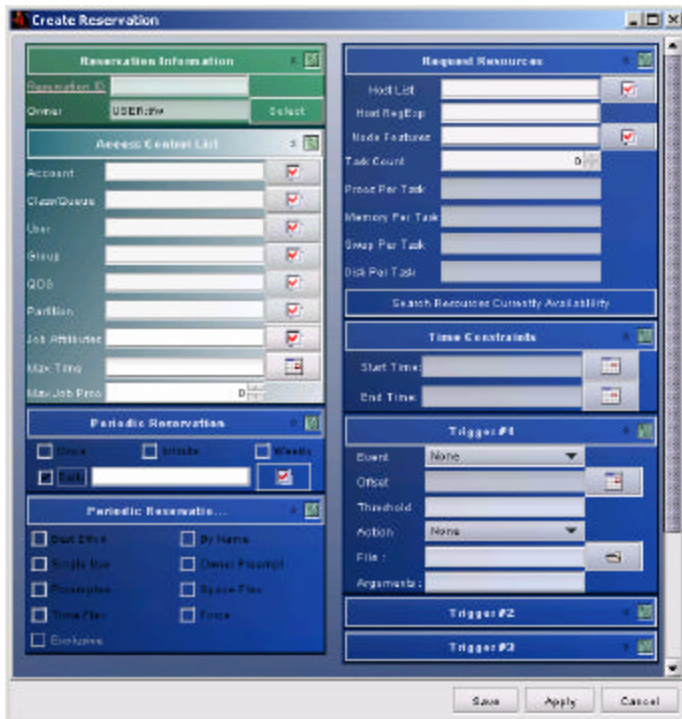
Cluster Administration


Visually overview cluster status, workload conditions, as well as the relationship of nodes and reservations.



Reservations

Submit advance reservations with a broad variety of capabilities. View current reservations in timelines that quickly enable you to view status and potential conflicts with all other reservations.





CLUSTER RESOURCES INC.™

Center for HPC Cluster Resource Management and Scheduling, Copyright © 2000-2003,

Event Triggers

Use powerful new event triggers to establish reservation chains, or maintenance tasks, automatic notifications, or any scriptable action.

Trigger #1

Event: end

Offset: 0:02:00:00

Threshold:

Action: execute

File: MailToMe.sh

Arguments: NearComplete

Administer Jobs

Quickly find troubled jobs, or those belonging to a user, a group or jobs with specific criteria. Double click and then manage the items of interest.

Jobs

Job ID: [] User: [] Proc: [0]

Running: [] State: [] Group: [] Nodes: [0]

Start Time: [] Clear Search

Job ID	State	User	Group	Start T.	Submit	Duration	Proc	Nodes	Details
job001	Running	lgh	scdf@	00:07:30		2	2		
job002	Running	anna	sdaf@	00:07:30		4	4		
job010	Running	kykem	boob@	00:07:30		6	6		
job016	Running	lgh	sdaf@	00:07:30		4	4		
job015	Running	kill	sdaf@	00:07:30		10	10		
job014	Running	kill	sdaf@	00:07:30		10	10		
job013	Running	lgh	sdaf@	00:07:30		4	4		
job012	Running	gyp	sdaf@	00:07:30		4	4		
job011	Running	lgh	sdaf@	00:07:30		2	2		
job010	Running	anna	sdaf@	00:07:30		6	6		
job009	Running	pub	sdaf@	00:07:30		4	4		
job001	Running	real	sdaf@	00:07:30		10	10		
job070	Running	real	sdaf@	00:07:30		4	4		
job070	Running	no				4	4		
job070	Running	pa				4	4		
job070	Running	lgh				1	1	1	1
job070	Running	pub				4	4		
job070	Running	pub				4	4		
job070	Running	real				4	4		
job070	Running	real				10	10		
job067	Running	no				4	4		
job066	Running	no				12	12		
job065	Running	no				4	4		
job067	Idle	gyp	sdaf@	00:07:30		4	4		
job025	Idle	real	sdaf@	00:07:30		4	4		

Select State dialog box:

Add State

Running []

Idle []

Blocked []

Suspended []

Select All Clear

Save Apply Cancel

CLUSTER RESOURCES INC.™

Center for HPC Cluster Resource Management and Scheduling, Copyright © 2000-2003,

FairShare

Instantly see results of FairShare settings and their impacts to users.

The screenshot shows the FairShare application window. At the top, there are sliders for 'Dropt Factor' (set to 100) and 'Depth' (set to 0). Below these are input fields for 'Interval Length' and a dropdown for 'FAIRSHARE POLICY' (set to 'EBDFCATBEP3'). The main area is a table with columns: 'Display Weight', 'System Resource', 'Metric', 'Target', '%', and five columns labeled '1' through '5'. The table lists various users and their resource usage across these five metrics.

Display Weight	System Resource	Metric	Target	%	1	2	3	4	5
-	-	-	100.0 %	1.98	3.19	3.13	2.73	2.82	
User	readp	-	5.29 %	5.29	5.29	5.32	0.0	2.37	
User	janth	-	8.47 %	8.47	8.47	8.61	12.84	15.38	
User	lylem	-	3.17 %	3.17	3.17	4.26	6.12	6.00	
User	lgh	-	10.52 %	10.52	10.52	18.08	9.40	0.0	
User	wroth	-	2.12 %	2.12	2.12	2.13	4.26	8.47	
User	jml	-	0.0 %	0.0	0.0	8.0	0.0	0.0	
User	rlcy	-	4.23 %	4.23	4.23	4.26	4.09	3.66	
User	flwr	-	23.29 %	23.29	23.29	23.4	38.13	28.4	
User	lll	-	4.23 %	4.23	4.23	4.20	4.09	4.73	
User	coll	-	16.93 %	16.93	16.93	17.00	19.57	18.93	
User	zfoval	-	0.0 %	0.0	0.0	8.0	0.0	1.18	
User	patn	-	6.36 %	6.36	6.36	6.90	4.09	3.66	
User	war	-	0.0 %	0.0	0.0	8.0	0.0	2.08	
User	anna	-	5.29 %	5.29	5.29	5.32	2.45	1.18	

Priority

Establish priority settings, view how multiple factors impact final combined results and tune settings with a great level of flexibility.

The screenshot shows the Priority application window. It features several panels for configuring settings: 'Service' (with 'target' and 'weight'), 'Credential' (with 'user', 'group', 'account', 'stage', 'weight'), 'Attribute' (with 'attribute', 'weight'), 'Fairshare' (with 'resource', 'weight'), and 'Resource' (with 'node', 'proc', 'memory', 'weight'). Below these panels is a large table with columns: 'JOB ID', 'Priority', 'EBDF CAT', 'Queue Name', 'Job Name', 'Job Status', 'Job Type', 'Target', 'Queue Name', 'Job ID', and 'CPU'. The table lists various jobs and their resource usage across these categories.

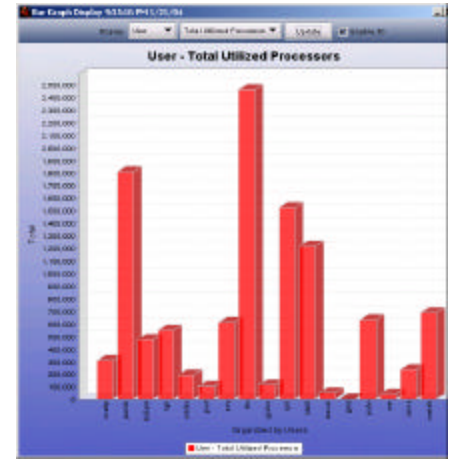
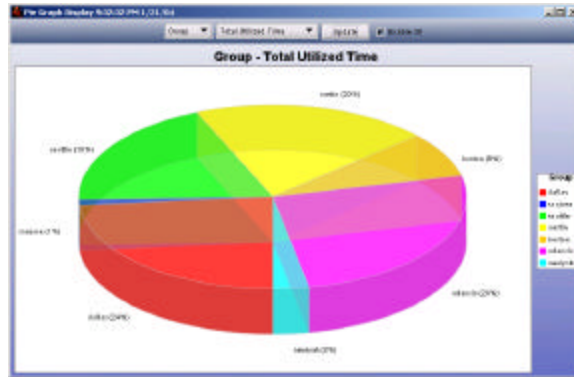
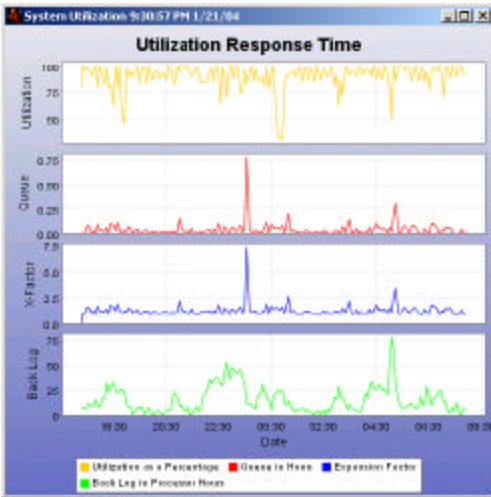
JOB ID	Priority	EBDF CAT	Queue Name	Job Name	Job Status	Job Type	Target	Queue Name	Job ID	CPU
job2410287	95.8%	52.0%	2.0%	0.0%	46.0%	0.0%	0.0%	3.8%	204	
job2414273	9.8%	37.0%	0.0%	0.0%	0.0%	0.0%	0.0%	3.8%	100	
job230469	7.8%	30.0%	4.0%	0.0%	0.0%	0.0%	0.0%	3.8%	277	
job240469	7.8%	30.0%	4.0%	0.0%	0.0%	0.0%	0.0%	3.8%	277	
job240261	18.8%	37.0%	3.0%	0.0%	0.0%	0.0%	0.0%	3.8%	224	
job230517	7.8%	37.0%	2.0%	0.0%	0.0%	0.0%	0.0%	3.8%	164	
job240761	18.8%	37.0%	3.0%	0.0%	0.0%	0.0%	0.0%	3.8%	224	
job240551	23.8%	37.0%	3.0%	0.0%	0.0%	0.0%	0.0%	3.8%	539	
job240100934	0.0%	36.0%	4.0%	0.0%	0.0%	0.0%	0.0%	3.8%	0.0	
job240100934	0.0%	36.0%	4.0%	0.0%	0.0%	0.0%	0.0%	3.8%	0.0	

CLUSTER RESOURCES INC.™

Center for HPC Cluster Resource Management and Scheduling, Copyright © 2000-2003,

Monitor, Diagnose and Report

Monitor utilization, queue loads, backlog, planned and actual resource consumption by groups, users or other factors. Create reports that report on service levels provided and other resource usage for organizational units or specific to projects.



Print Preview

File Edit View Options Help

100%

Resource Consumption Report

Report: Departmental and user breakdown of consumed processor time and associated costs.

Organization:	Cost	Processor Hours	Cost
org01	08	1.0	807
org02	23	1.0	249.0
org03	01	1.0	121.0
org04	48	1.0	872
org05	128	1.0	3,902
Total:	208		\$112
Organization:	Cost	Processor Hours	Cost
org06	22	1.0	930
org07	41	1.0	391.0
org08	198	1.0	920.4
Total:	261		\$981.4
Organization:	Cost	Processor Hours	Cost
org09	105	1.0	1457.0
org10	117	1.0	1776.9
Total:	222		\$3233.9
Organization:	Cost	Processor Hours	Cost
org11	337	1.0	1809.9
org12	481	1.0	1721.0

Page 1 of 2

Custom Report Creation

Category: org01, User: user.palmer.anna@utah.edu

Unit: Total Utilized Time, Weight: 1.0

Category	Business	Unit	Weight	Total
auth	lgh	62250	1.0	62250.0
auth	palmer	67200	1.0	67200.0
carp	carp	12150	1.0	12150.0
carp	gpc	23650	1.0	23650.0
carp	swi	102500	1.0	102500.0
orlando	poj	62500	1.0	62500.0
orlando	ljl	85000	1.0	85000.0
orlando	hll	52500	1.0	52500.0
newyork	lvt	76500	1.0	76500.0
newyork	gks	50400	1.0	50400.0
dallas	rsj	50400	1.0	50400.0
dallas	swi	25200	1.0	25200.0
dallas	swi	47250	1.0	47250.0
dallas	arcs	32200	1.0	32200.0
dallas	swi	100500	1.0	100500.0
boise	lyl	70200	1.0	70200.0
justin	lvt	62500	1.0	62500.0
justin	lvt	100500	1.0	100500.0

Create Report Apply Cancel