

*Managing complex cluster architectures with
Bright Cluster Manager*

Christopher Huggins

About ClusterVision



- Specialists in Compute, Storage & Database Clusters (Tailor-Made, Turn-Key)
- Unique position in Europe (EMEA)
- Offices in Amsterdam, Gloucester, Munich, Paris, Milan, Geneva, Madrid, Oslo
- 55 Staff, most technical, all specialised in clusters
- Hardware brand independent (Dell, IBM, Cray, Supermicro, Asus)
- Background in Science, Research, Engineering
- At forefront of clustering technology
- Award winning (Intermediar, Vosko, NBCC, Deloitte Rising Star)
- ISO9001:2008 certified
- Financially strong, profitable, growing



- HPC Problem Solving
- Cluster Design
 - Start with Existing Environment
 - User Requirements
 - Application Benchmarking
 - Power & Space Constraints
 - Personal Hardware & Software Preferences
 - Total Cost of Ownership
 - Future Environment
- Cluster Delivery & Installation
 - Speed
 - Professional Project Management
 - Personal Training
- Support Tailored to Your Needs
 - Hardware
 - Software
 - System administration (onsite/remote)
 - User helpdesk
 - First point-of-contact / Single point-of-contact




Bright Cluster Manager



Linux-based cluster management and
user environment

Goals

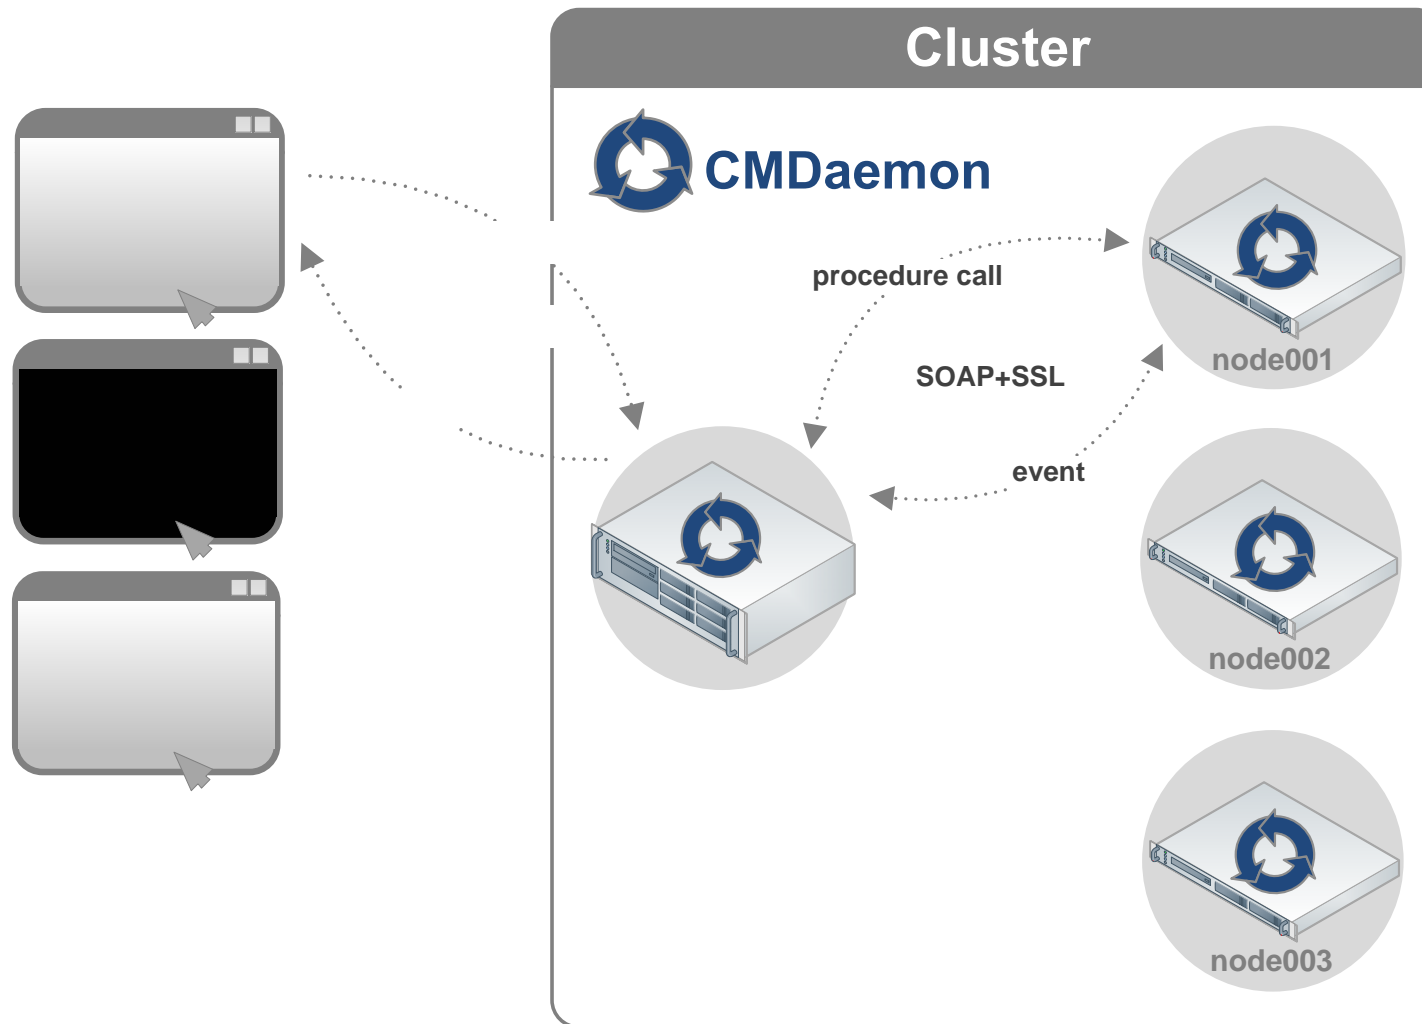
1. Make it really easy to install, manage and use clusters
2. Be suitable for very large clusters — thousands of nodes
3. Be complete



- Most cluster management solutions use the “toolkit” approach (Linux distro + tools)
 - Examples: Rocks, PCM, OSCAR, UniCluster, xCAT, etc.
 - Tools typically used: Ganglia, Cacti, Nagios, Cfengine, System Imager, etc.
- Issues with the “toolkit” approach:
 - Tools rarely designed to work together
 - Tools rarely designed for HPC
 - Tools rarely designed to scale
 - Each tool has its own command line interface and GUI
 - Each tool has its own daemon and database
 - Roadmap dependent on developers of the tools
- Making a collection of unrelated tools work together
 - Requires a lot of expertise and scripting
 - Rarely leads to a really easy-to-use and scalable solution

- Bright Cluster Manager takes a much more fundamental & integrated approach
 - Designed and written from the ground up
 - Single cluster management daemon provides all functionality
 - Single, central database for configuration and monitoring data
 - Single CLI and GUI for ALL cluster management functionality
- Which makes Bright Cluster Manager ...
 - Extremely easy to use
 - Extremely scalable
 - Complete
 - Flexible

Architecture — CMDaemon



- Cluster Management Daemon
- Cluster Management Shell
- Cluster Management GUI
- Linux distribution, with many enhancements
 - Red Hat & SUSE Enterprise
 - CentOS & Scientific Linux
- Node provisioning system
- Workload management system
 - Sun Grid Engine, Torque/Maui
 - MOAB, PBS Pro, LSF
- HPC user environment
- Cluster management & monitoring functionality
- Years of HPC expertise
- Consistent integration of all the above


----- +
Bright Cluster Manager

- Let users focus on performing computations
- Rich collection of HPC software
 - Compilers (GNU, Intel*, Portland*, Open64, etc.)
 - Parallel middleware (MPI libraries, threading libraries, OpenMP, Global Arrays, etc.)
 - Mathematical libraries (ACML, MKL*, LAPACK, BLAS, GOTO, etc.)
 - Development tools (debuggers, profilers, etc.)
 - Environment modules
- NVIDIA CUDA, OpenCL & GPGPU drivers
- *Intel Cluster Ready* Compliant
Compliant applications run out of the box




Graphical User Interface (GUI)

- Standalone desktop application
- Runs on Linux, Windows, MacOS X
- Built on top of Mozilla XUL engine
- Offers administrator full cluster control
- Manages multiple clusters simultaneously



Command Line Interface (CLI)

- All GUI functionality also available through Command Line Interface (CLI)
- Interactive and scriptable in batch mode



Bright Cluster Manager

FileMonitoringViewHelp

RESOURCES

My Clusters

Seismic Houston

Switches

switch01switch02switch03switch04switch05

Networks

externalnetipminetmpinetmlavenetstoragenet

Power Distribution Units

apc01apc02apc03apc04

Software Images

default-image

Node Categories

slave

Head Nodes

demohead1demohead2

Welcome to Bright Cluster Manager

Seismic Oslo

Modified: No

Host: oslo.seismic.com:8081

Connected: No

Certificate: /root/oslo.pfx

⚙️

✎

🚫

🔴

Seismic Abu Dhabi

Modified: No

Host: abudhabi.seismic.com:8081

Connected: No

Certificate: /root/.cm/cmgui/admin-abudhabi.pfx

⚙️

✎

🚫

🔴

Seismic Houston

Modified: No

Host: localhost:2581

Connected: Yes

Certificate: /root/.cm/cmgui/admin.pfx

⚙️

✎

🚫

🟢

+

Add a new cluster

EVENT VIEWER

🔍

+

🔍

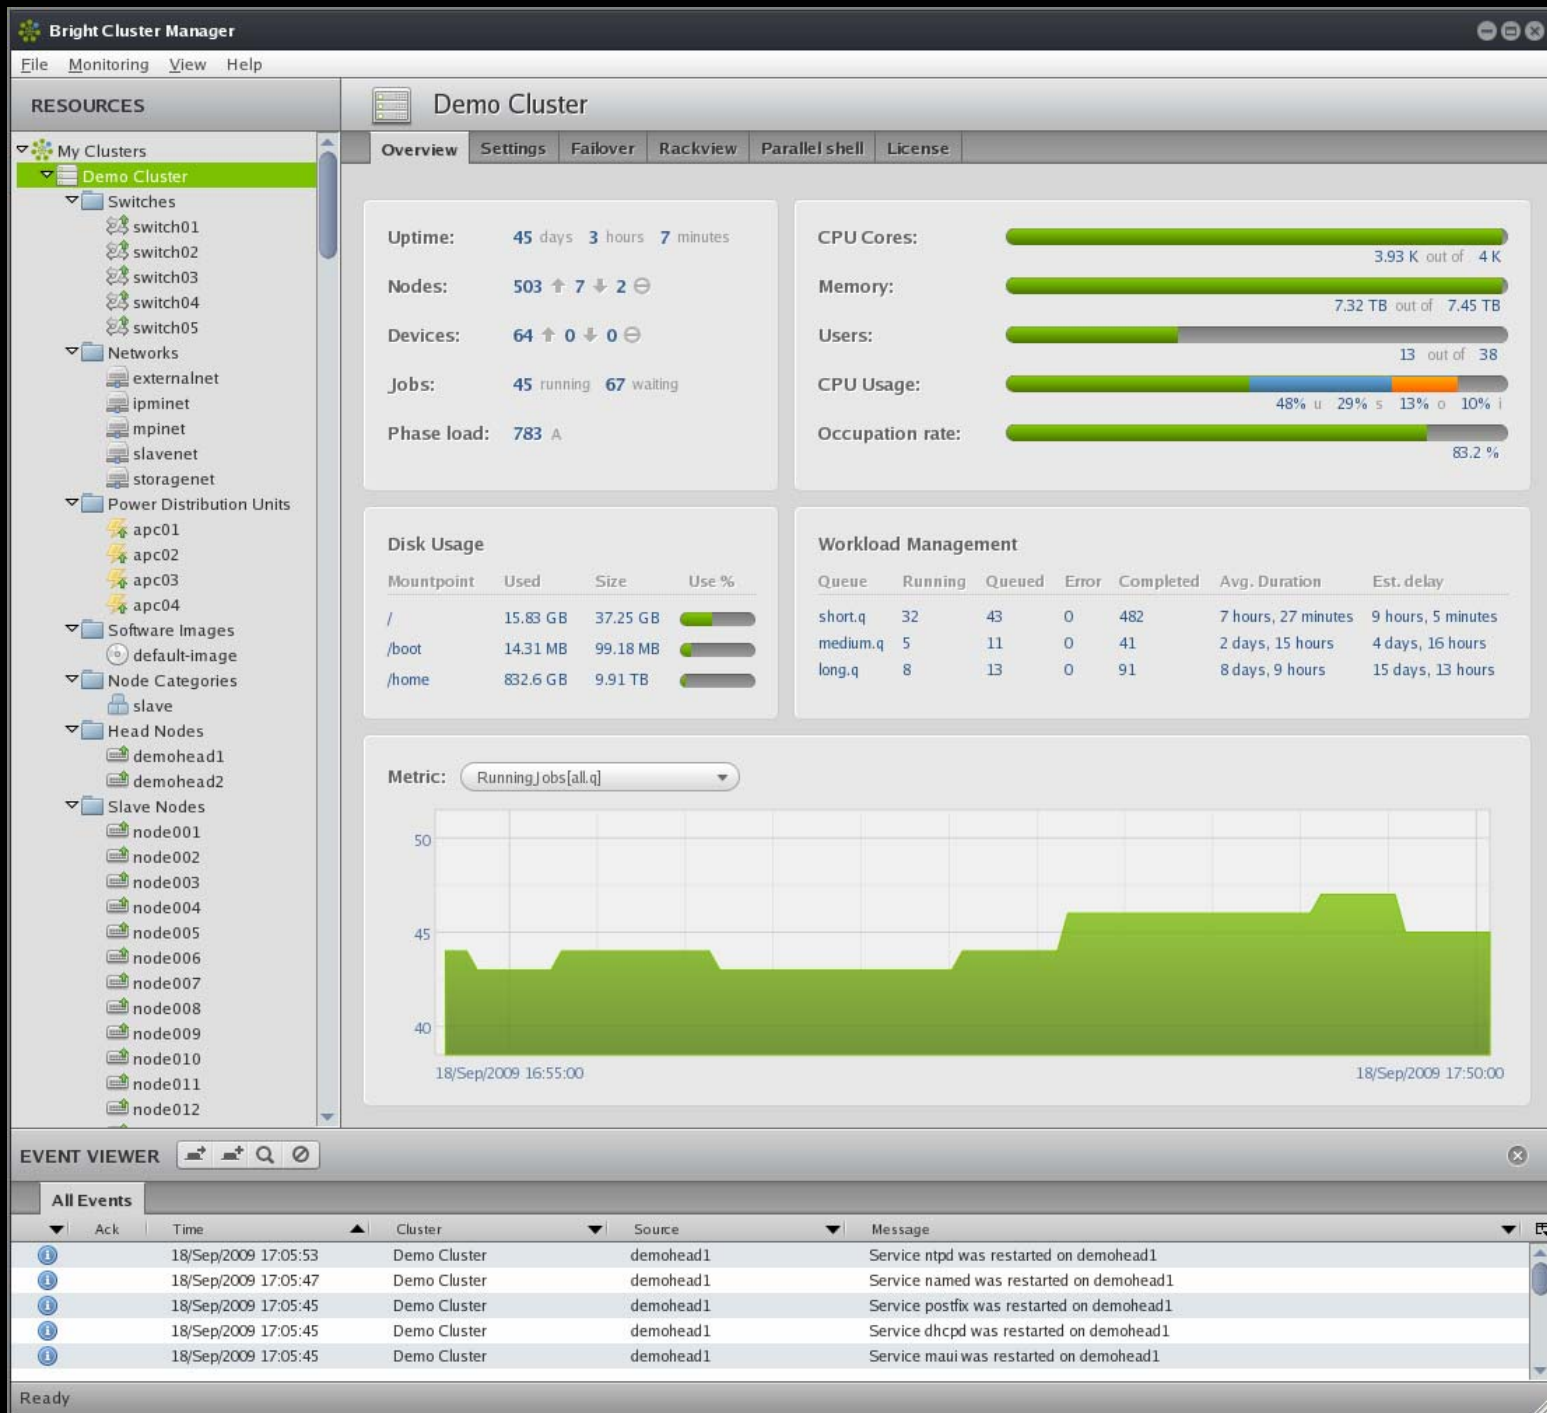
🚫

✕

All Events

	Ack	Time	Cluster	Source	Message
📘		18/Sep/2009 17:05:53	Demo Cluster	demohead1	Service ntpd was restarted on demohead1
📘		18/Sep/2009 17:05:47	Demo Cluster	demohead1	Service named was restarted on demohead1
📘		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service postfix was restarted on demohead1
📘		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service dhcpd was restarted on demohead1
📘		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service maui was restarted on demohead1

Ready



Bright Cluster Manager

FileMonitoringViewHelp

RESOURCES

My Clusters

Demo Cluster

Switches

switch01

switch02

switch03

switch04

switch05

Networks

externalnet

ipminet

mpinet

slavenet

storagenet

Power Distribution Units

apc01

apc02

apc03

apc04

Software Images

default-image

Node Categories

slave

Head Nodes

demohead1

demohead2

Slave Nodes

node001

node002

node003

node004

node005

node006

node007

node008

node009

node010

node011

node012

node001

Demo Cluster

OverviewTasksSettingsSystem InformationServicesProcess ManagementNetwork SetupFS MountsFS ExportsRoles

Power:OnOffReset

Operating System:ShutdownRestart

Add to node group:<new>AddRemove

Software image:

Update nodeReinstall node

Synchronize imageGrab to different image

Misc:Root ShellRemote ConsoleLocate in rackIdentify nodeOpenCloseProvisioning Log

EVENT VIEWER

All Events

	Ack	Time	Cluster	Source	Message
1		18/Sep/2009 17:05:53	Demo Cluster	demohead1	Service ntpd was restarted on demohead1
1		18/Sep/2009 17:05:47	Demo Cluster	demohead1	Service named was restarted on demohead1
1		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service postfix was restarted on demohead1
1		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service dhcpd was restarted on demohead1
1		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service maui was restarted on demohead1

Ready

Bright Cluster Manager

FileMonitoringViewHelp

RESOURCES

My Clusters

Demo Cluster

Switches

switch01

switch02

switch03

switch04

switch05

Networks

externalnet

ipminet

mpinet

slavenet

storagenet

Power Distribution Units

apc01

apc02

apc03

apc04

Software Images

default-image

Node Categories

slave

Head Nodes

demohead1

demohead2

Slave Nodes

Other Devices

Node Groups

Large Memory Nodes

Users & Groups

Workload Management

Monitoring Configuration

Authorisation

Authentication

Large Memory Nodes

Demo Cluster

SettingsTasks

Name: Large Memory Nodes

Nodes: node001..node012

All Nodes

node015

node016

node017

node018

node019

node020

Nodes in Group

node001

node002

node003

node004

node005

node006

Filter:

matches:

AND

OR

RevertSave

EVENT VIEWER

All Events

	Ack	Time	Cluster	Source	Message
!		18/Sep/2009 17:05:53	Demo Cluster	demohead1	Service ntpd was restarted on demohead1
!		18/Sep/2009 17:05:47	Demo Cluster	demohead1	Service named was restarted on demohead1
!		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service postfix was restarted on demohead1
!		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service dhcpcd was restarted on demohead1
!		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service maui was restarted on demohead1

Ready

Bright Cluster Manager

FileMonitoringViewHelp

RESOURCES

My Clusters

Demo Cluster

Switches

switch01

switch02

switch03

switch04

switch05

Networks

externalnet

ipminet

mpinet

slavenet

storagenet

Power Distribution Units

apc01

apc02

apc03

apc04

Software Images

default-image

Node Categories

slave

Head Nodes

demohead1

demohead2

Slave Nodes

Other Devices

Node Groups

Large Memory Nodes

Users & Groups

Workload Management

Monitoring Configuration

Authorisation

Authentication

Demo Cluster

OverviewSettingsFailoverRackviewParallel shellLicense

node001

18:43:52 up 45 days, 4:03, 20 users, load average: 0.23, 0.14, 0.04

USER	TTY	FROM	LOGIN@	IDLE	JCPU	PCPU	WHAT
root	tty1	-	18Aug09	11days	0.05s	0.01s	login -- root
root	tty2	-	02Sep09	27:16m	0.04s	0.00s	login -- root
root	tty3	-	26Aug09	16days	0.03s	0.03s	-bash
root	pts/1	druifje.clusterv	Wed11	41:52m	0.05s	0.05s	-bash

node002

18:43:52 up 45 days, 4:03, 20 users, load average: 0.23, 0.14, 0.04

USER	TTY	FROM	LOGIN@	IDLE	JCPU	PCPU	WHAT
root	tty1	-	18Aug09	11days	0.05s	0.01s	login -- root
root	tty2	-	02Sep09	27:16m	0.04s	0.00s	login -- root
root	tty3	-	26Aug09	16days	0.03s	0.03s	-bash
root	pts/1	druifje.clusterv	Wed11	41:52m	0.05s	0.05s	-bash

node003

18:43:52 up 45 days, 4:03, 20 users, load average: 0.23, 0.14, 0.04

USER	TTY	FROM	LOGIN@	IDLE	JCPU	PCPU	WHAT
root	tty1	-	18Aug09	11days	0.05s	0.01s	login -- root
root	tty2	-	02Sep09	27:16m	0.04s	0.00s	login -- root
root	tty3	-	26Aug09	16days	0.03s	0.03s	-bash
root	pts/1	druifje.clusterv	Wed11	41:52m	0.05s	0.05s	-bash

node004

18:43:52 up 45 days, 4:03, 20 users, load average: 0.23, 0.14, 0.04

USER	TTY	FROM	LOGIN@	IDLE	JCPU	PCPU	WHAT
root	tty1	-	18Aug09	11days	0.05s	0.01s	login -- root
root	tty2	-	02Sep09	27:16m	0.04s	0.00s	login -- root
root	tty3	-	26Aug09	16days	0.03s	0.03s	-bash
root	pts/1	druifje.clusterv	Wed11	41:52m	0.05s	0.05s	-bash

slave

w

Execute

EVENT VIEWER

All Events

	Ack	Time	Cluster	Source	Message
1		18/Sep/2009 18:42:36	Demo Cluster	demohead1	node004 Installing
1		18/Sep/2009 18:42:06	Demo Cluster	demohead1	New certificate request with ID: 6
1		18/Sep/2009 18:30:06	Demo Cluster	demohead1	node003 Installing
1		18/Sep/2009 18:29:39	Demo Cluster	demohead1	New certificate request with ID: 5
1		18/Sep/2009 18:29:36	Demo Cluster	demohead1	node002 Installing
1		18/Sep/2009 18:29:35	Demo Cluster	demohead1	New certificate request with ID: 4

Ready

Bright Cluster Manager

File Monitoring View Help

RESOURCES

My Clusters

- Demo Cluster
 - Switches
 - switch01
 - switch02
 - switch03
 - switch04
 - switch05
 - Networks
 - externalnet
 - ipminet
 - mpinet
 - slavenet
 - storagenet
 - Power Distribution Units
 - apc01
 - apc02
 - apc03
 - apc04
 - Software Images
 - default-image
 - Node Categories
 - slave
 - Head Nodes
 - demohead1
 - demohead2
 - Slave Nodes
 - Other Devices
 - Node Groups
 - Large Memory Nodes
 - Users & Groups
 - Workload Management
 - Monitoring Configuration
 - Authorisation
 - Authentication

Demo Cluster

Overview Settings Failover Rackview Parallel shell License

node001

node002

node003

node004

slave sleep 100 Execute

Parallel Command Details

Command: **sleep 100**
User: **root** Start time: 18/Sep/2009 18:46

Node	PID	CPU Cycles
node004	6850	0
node003	6858	0
node002	6855	0
node001	6856	0

Refresh Kill Selected Kill All Done

Parallel Shell

Starttime: 18/Sep/2009 18:46

Node	PID	CPU Cycles
node004	6850	0
node003	6858	0
node002	6855	0
node001	6856	0

Refresh Details Kill Done

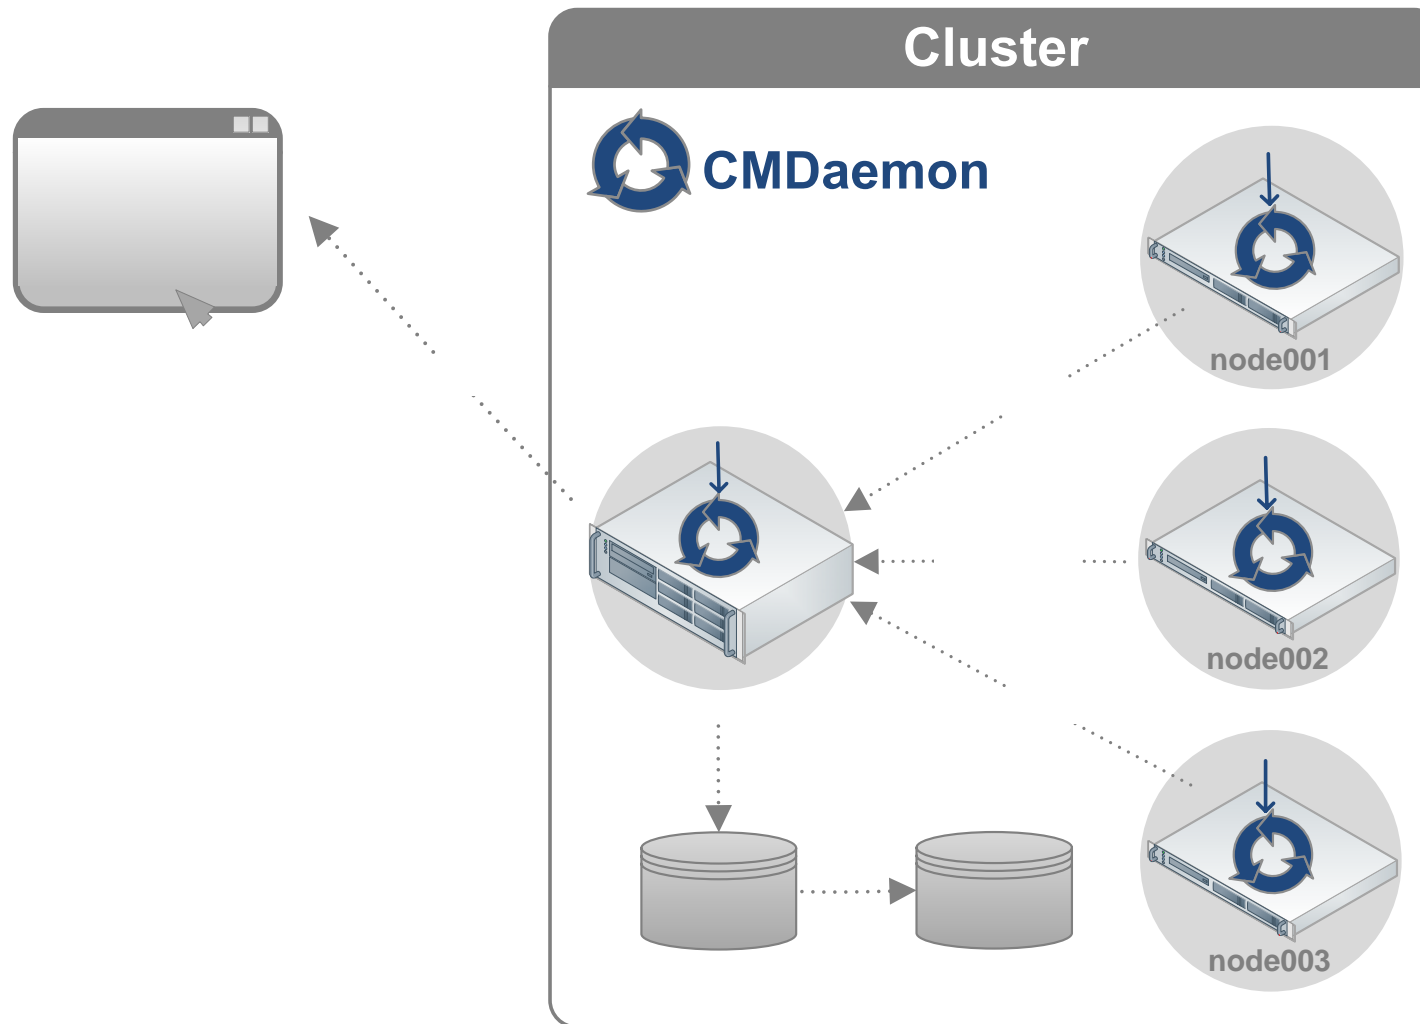
EVENT VIEWER

All Events

Ack	Time	Cluster	Source	Message
1	18/Sep/2009 18:42:36	Demo Cluster	demohead1	node004 Installing
1	18/Sep/2009 18:42:06	Demo Cluster	demohead1	New certificate request with ID: 6
1	18/Sep/2009 18:30:06	Demo Cluster	demohead1	node003 Installing
1	18/Sep/2009 18:29:39	Demo Cluster	demohead1	New certificate request with ID: 5
1	18/Sep/2009 18:29:36	Demo Cluster	demohead1	node002 Installing
1	18/Sep/2009 18:30:35	Demo Cluster	demohead1	New certificate request with ID: 4

Ready

Architecture — Monitoring




RESOURCES

- demohead1
 - CPU
 - Disk
 - Memory
 - BufferMemory (B)
 - CacheMemory (B)
 - MemoryFree (B)
 - MemoryUsed (B)
 - SwapFree (B)
 - SwapUsed (B)
 - Network
 - Operating System
 - CtxtSwitches (ctx_switch/s)
 - Forks (process/s)
 - LoadFifteen
 - LoadFive
 - LoadOne
 - ProcessCount
 - RunningProcesses
 - Uptime (s)
 - Workload
 - Cluster
 - CPUcoresAvailable
 - DevicesUp
 - NetworkBytesRecv (B)
 - NetworkBytesSent (B)
 - NodesUp
 - OccupationRate (%)
 - PhaseLoad (A)
 - TotalCPUIdle
 - TotalCPUSystem
 - TotalCPUUser
 - TotalMemoryUsed (B)
 - TotalSwapUsed (B)
- demohead2
- node001
- node002
- node003
- node004
- node005
- node006


Demo Cluster

demohead1:LoadOne




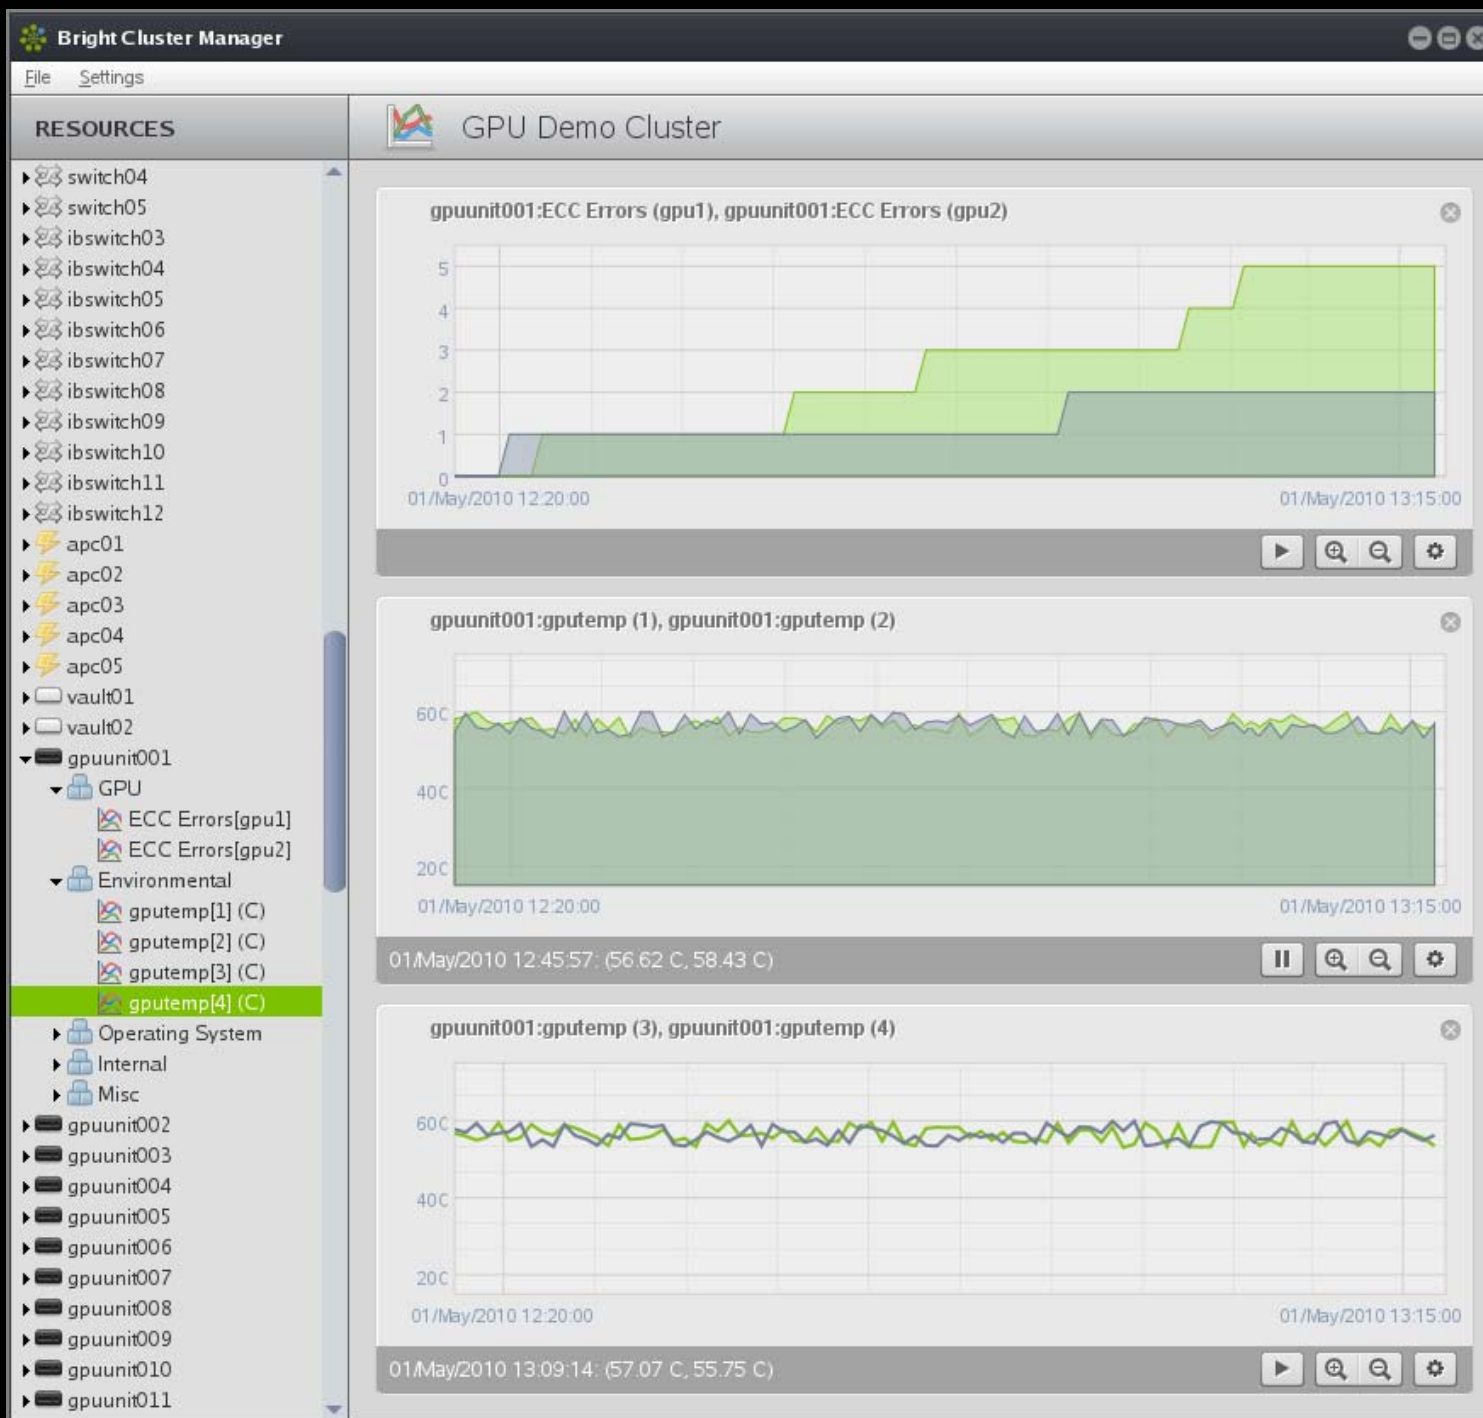
drop sensor here

demohead1:Running Jobs (all.q) +



demohead1:LoadOne + +





RESOURCES

My Clusters

GPU Demo Cluster

- Switches
- Networks
- Power Distribution Units
- Software Images
 - default-image
- Node Categories
 - slave
- Head Nodes
 - head1
 - head2
- Chassis
- GPU Unit
 - gpuunit001
 - gpuunit002
 - gpuunit003
 - gpuunit004
 - gpuunit005
 - gpuunit006
 - gpuunit007
 - gpuunit008
 - gpuunit009
 - gpuunit010
 - gpuunit011
 - gpuunit012
 - gpuunit013
 - gpuunit014
 - gpuunit015
 - gpuunit016
 - gpuunit017
 - gpuunit018
 - gpuunit019
 - gpuunit020
 - gpuunit021
 - gpuunit022
 - gpuunit023
 - gpuunit024
 - gpuunit025
 - gpuunit026
 - gpuunit027
 - gpuunit028
 - gpuunit029
 - gpuunit030
 - gpuunit031
 - gpuunit032
 - gpuunit033
 - gpuunit034
 - gpuunit035
 - gpuunit036
 - gpuunit037
 - gpuunit038

GPU Demo Cluster

Overview Settings Failover Rackview Health Parallel shell License

U	Rack 1	Rack 2	Rack 3	Rack 4	Rack 5	Rack 6
01		ibswitch03	ibswitch05	ibswitch07	ibswitch09	ibswitch11
02						
03						
04	head1	001	013	025	037	049
05		gpuunit001	gpuunit019	gpuunit037	gpuunit055	gpuunit073
06		002	014	026	038	050
07		gpuunit002	gpuunit020	gpuunit038	gpuunit056	gpuunit074
08	head2	003	015	027	039	051
09		gpuunit003	gpuunit021	gpuunit039	gpuunit057	gpuunit075
10		004	016	028	040	052
11		gpuunit004	gpuunit022	gpuunit040	gpuunit058	gpuunit076
12		005	017	029	041	053
13		gpuunit005	gpuunit023	gpuunit041	gpuunit059	gpuunit077
14		006	018	030	042	054
15		gpuunit006	gpuunit024	gpuunit042	gpuunit060	gpuunit078
16	login01	007	019	031	043	055
17		gpuunit007	gpuunit025	gpuunit043	gpuunit061	gpuunit079
18		008	020	032	044	056
19		gpuunit008	gpuunit026	gpuunit044	gpuunit062	gpuunit080
20	login02					
21		ibswitch04	ibswitch06	ibswitch08	ibswitch10	ibswitch12
22						
23						
24	login03	009	021	033	045	057
25		gpuunit009	gpuunit027	gpuunit045	gpuunit063	gpuunit081
26		gpuunit010	gpuunit028	gpuunit046	gpuunit064	gpuunit082
27						
28	storage01	010	022	034	046	058
29		gpuunit011	gpuunit029	gpuunit047	gpuunit065	gpuunit083
30		gpuunit012	gpuunit030	gpuunit048	gpuunit066	gpuunit084
31		gpuunit013	gpuunit031	gpuunit049	gpuunit067	gpuunit085
32	storage02	011	023	035	047	059
33		gpuunit014	gpuunit032	gpuunit050	gpuunit068	gpuunit086
34		gpuunit015	gpuunit033	gpuunit051	gpuunit069	gpuunit087
35		gpuunit016	gpuunit034	gpuunit052	gpuunit070	gpuunit088
36	vault01	012	024	036	048	060
37		gpuunit017	gpuunit035	gpuunit053	gpuunit071	gpuunit089
38		gpuunit018	gpuunit036	gpuunit054	gpuunit072	gpuunit090
39	vault02					
40		switch01	switch02	switch03	switch04	switch05
41						
42						

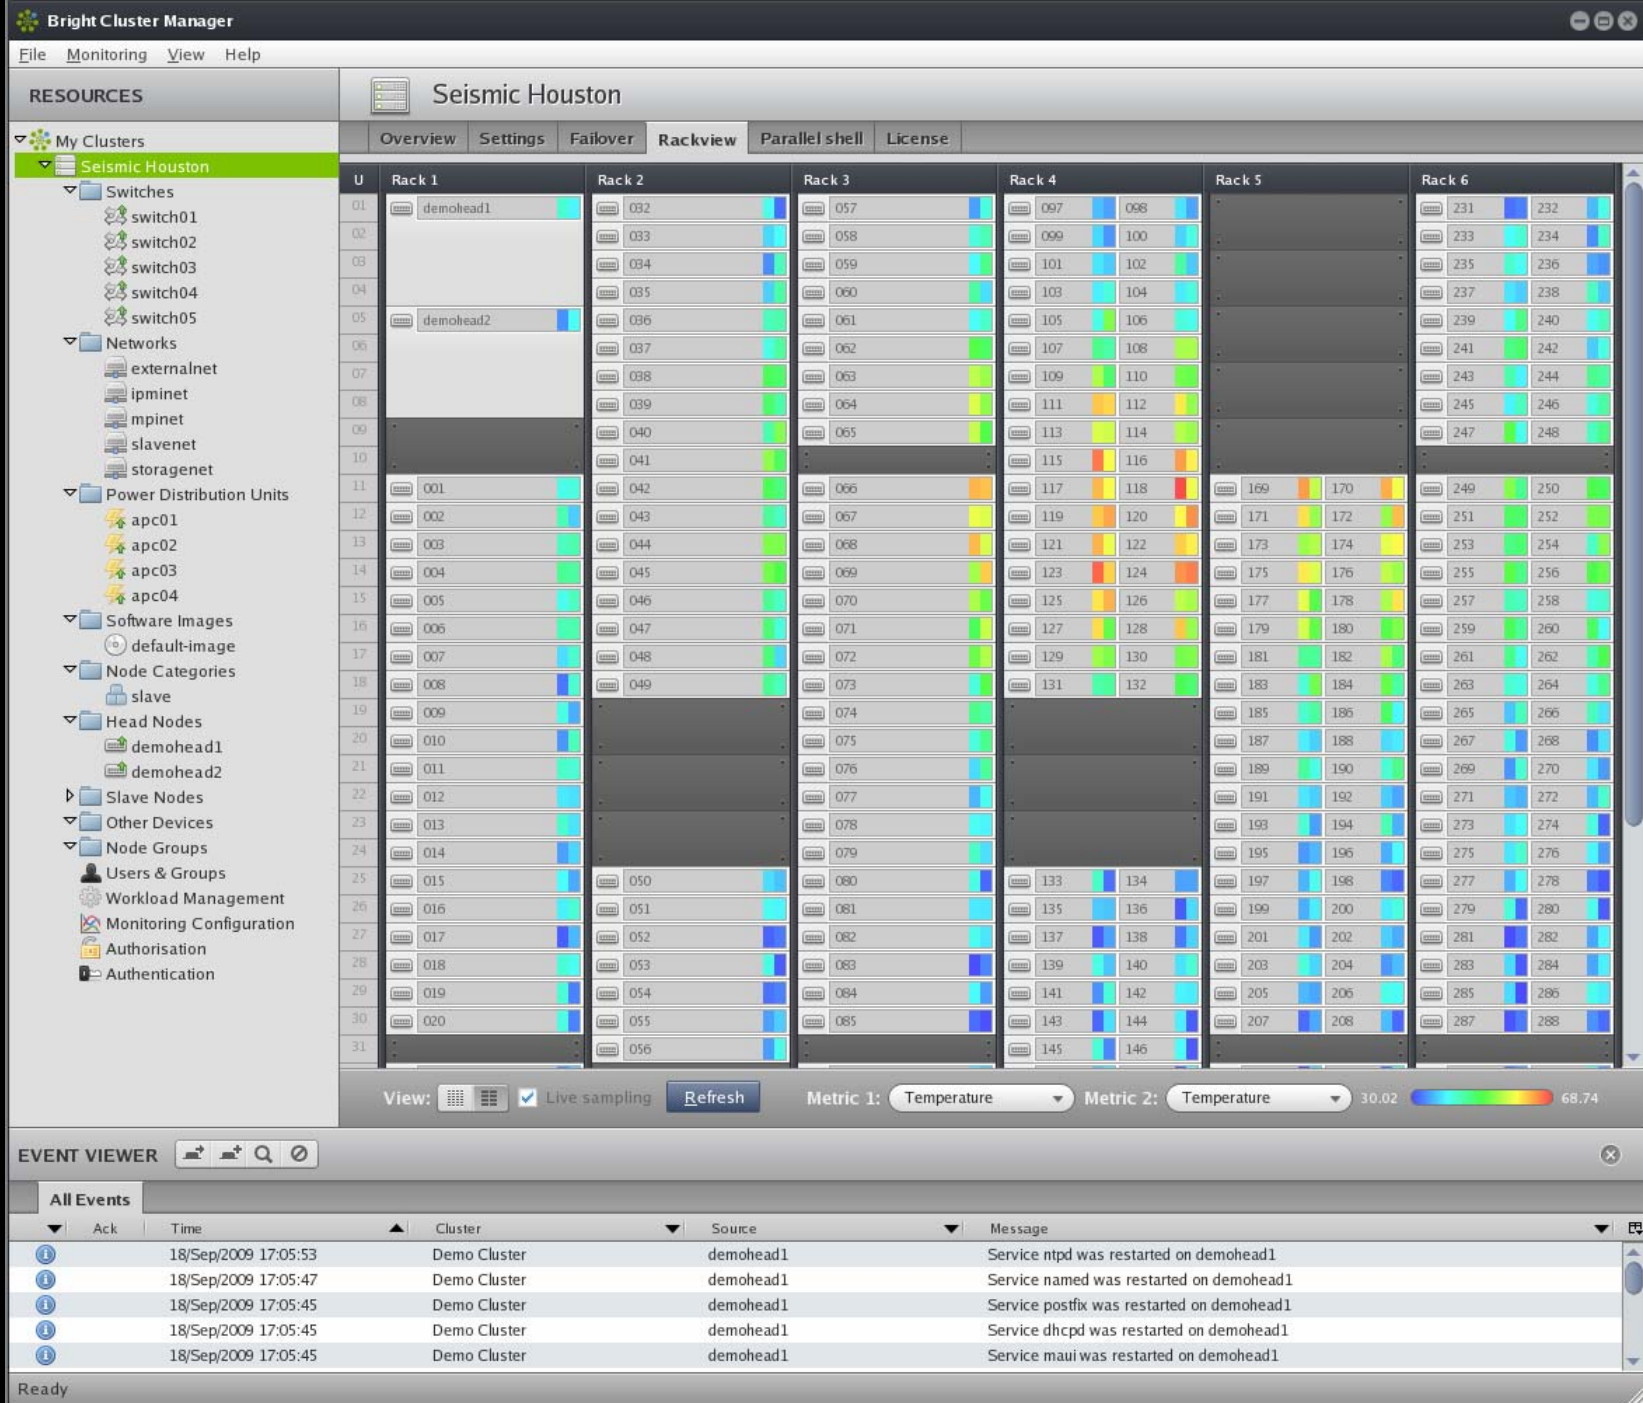
View:



Refresh

Setup

AllFansRunning pass fail



Bright Cluster Manager

File Monitoring View Help

RESOURCES

My Clusters

Seismic Houston

Switches

switch01

switch02

switch03

switch04

switch05

Networks

externalnet

ipminet

mpinet

slavenet

storagenet

Power Distribution Units

apc01

apc02

apc03

apc04

Software Images

default-image

Node Categories

slave

Head Nodes

demohead1

demohead2

Slave Nodes

Other Devices

Node Groups

Users & Groups

Workload Management

Monitoring Configuration

Authorisation

Authentication



Seismic Houston

Overview


Settings

Failover

Rackview

Parallel shell

License



View:

☒ Live sampling

Refresh

Metric 1: Temperature

30.01



69.34

EVENT VIEWER



All Events

	▼	Ack	Time	▲	Cluster	▼	Source	▼	Message	▼	⌵
ⓘ			18/Sep/2009 17:05:53		Demo Cluster		demohead1		Service ntpd was restarted on demohead1		
ⓘ			18/Sep/2009 17:05:47		Demo Cluster		demohead1		Service named was restarted on demohead1		
ⓘ			18/Sep/2009 17:05:45		Demo Cluster		demohead1		Service postfix was restarted on demohead1		
ⓘ			18/Sep/2009 17:05:45		Demo Cluster		demohead1		Service dhcpd was restarted on demohead1		
ⓘ			18/Sep/2009 17:05:45		Demo Cluster		demohead1		Service maui was restarted on demohead1		

Ready

Bright Cluster Manager

File Monitoring View Help

RESOURCES

My Clusters

- Demo Cluster
 - Switches
 - switch01
 - switch02
 - switch03
 - switch04
 - switch05
 - Networks
 - externalnet
 - ipminet
 - mpinet
 - slavenet
 - storagenet
 - Power Distribution Units
 - apc01
 - apc02
 - apc03
 - apc04
 - Software Images
 - default-image
 - Node Categories
 - slave
 - Head Nodes
 - demohead1
 - demohead2
 - Slave Nodes
 - Other Devices
 - Node Groups
 - Users & Groups**
 - Workload Management
 - Monitoring Configuration
 - Authorisation
 - Authentication

Users & Groups

Users Groups

Modified	Name	Full Name	User ID	Home directory
✓	alex	Alex Bozarski	504	/home/alex
✓	james	James Watt	505	/home/james
✓	jodi	Jodi Johnson	503	/home/jodi
✓	kate	Kate Moss	506	/home/kate
✓	koen	Koen van den Bosch	502	/home/koen
✓	matthew	Matthew Ellis	507	/home/matthew

Edit User

Login Name: matthew

User ID: 507

Full Name: Matthew Ellis

Group ID: matthew

Login shell: /bin/bash

Home directory: /home/matthew

Password:

Retype password:

Cancel Ok

Edit Add Remove Revert Save

EVENT VIEWER

All Events

Ack	Time	Cluster	Source	Message
✓	18/Sep/2009 17:05:53	Demo Cluster	demohead1	Service ntpd was restarted on demohead1
✓	18/Sep/2009 17:05:47	Demo Cluster	demohead1	Service named was restarted on demohead1
✓	18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service postfix was restarted on demohead1
✓	18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service dhcpd was restarted on demohead1
✓	18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service maui was restarted on demohead1

Ready

Bright Cluster Manager

FileMonitoringViewHelp

RESOURCES

My Clusters

Demo Cluster

Switches

switch01

switch02

switch03

switch04

switch05

Networks

externalnet

ipminet

mpinet

slavenet

storagenet

Power Distribution Units

apc01

apc02

apc03

apc04

Software Images

default-image

Node Categories

slave

Head Nodes

demohead1

demohead2

Slave Nodes

Other Devices

Node Groups

Users & Groups

Workload Management

Monitoring Configuration

Authorisation

Authentication

Authorisation

Demo Cluster

James

CMMain

☐ Access

☐ GET_LICENSE_INFO_TOKEN

☐ GET_VERSION_TOKEN

☐ GET_SERVER_STATUS_TOKEN

☐ GET_CLUSTER_SETUP_TOKEN

☐ PING_TOKEN

☐ PCOPY_TOKEN

☐ UPDATE_SELF_TOKEN

CMSession

☒ Access

☐ CMSESSION_ADMIN

☐ GET_SESSION_TOKEN

☐ REGISTER_NODE_SESSION_TOKEN

admin (read-only)

CMMain

☒ Access

☒ GET_LICENSE_INFO_TOKEN

☒ GET_VERSION_TOKEN

☒ GET_SERVER_STATUS_TOKEN

☒ GET_CLUSTER_SETUP_TOKEN

☒ PING_TOKEN

☒ PCOPY_TOKEN

☒ UPDATE_SELF_TOKEN

CMSession

☒ Access

☒ CMSESSION_ADMIN

☒ GET_SESSION_TOKEN

☒ REGISTER_NODE_SESSION_TOKEN

node (read-only)

CMMain

☒ Access

☐ GET_LICENSE_INFO_TOKEN

☐ GET_VERSION_TOKEN

☐ GET_SERVER_STATUS_TOKEN

☐ GET_CLUSTER_SETUP_TOKEN

☐ PING_TOKEN

☐ PCOPY_TOKEN

☐ UPDATE_SELF_TOKEN

CMSession

☒ Access

☐ CMSESSION_ADMIN

☐ GET_SESSION_TOKEN

☒ REGISTER_NODE_SESSION_TOKEN

+

Add Profile

Revert

Save

EVENT VIEWER

All Events

Ack

Time

Cluster

Source

Message

18/Sep/2009 17:05:53

Demo Cluster

demohead1

Service ntpd was restarted on demohead1

18/Sep/2009 17:05:47

Demo Cluster

demohead1

Service named was restarted on demohead1

18/Sep/2009 17:05:45

Demo Cluster

demohead1

Service postfix was restarted on demohead1

18/Sep/2009 17:05:45

Demo Cluster

demohead1

Service dhcpd was restarted on demohead1

18/Sep/2009 17:05:45

Demo Cluster

demohead1

Service maui was restarted on demohead1

Ready

Bright Cluster Manager

FileMonitoringViewHelp

RESOURCES

My Clusters

Demo Cluster

Switches

switch01

switch02

switch03

switch04

switch05

Networks

externalnet

ipminet

mpinet

slavenet

storagenet

Power Distribution Units

apc01

apc02

apc03

apc04

Software Images

default-image

Node Categories

slave

Head Nodes

demohead1

demohead2

Slave Nodes

Other Devices

Node Groups

Users & Groups

Workload Management

Monitoring Configuration

Authorisation

Authentication

Workload Management

Demo Cluster

JobsQueues

Modified	Name	Scheduler	User	Queue	Status
	fluent	sge	jodi	medium.q	queued
	fluent	sge	jodi	medium.q	queued
	fluent	sge	jodi	medium.q	queued
	fluent	sge	jodi	medium.q	running
	gromacs	sge	alex	long.q	queued
	gromacs	sge	alex	long.q	running
	gromacs	sge	alex	long.q	running
	gromacs	sge	alex	long.q	running
	gromacs	sge	alex	medium.q	queued
	hpcc	sge	kate	long.q	queued
	hpcc	sge	kate	long.q	running
	hpcc	sge	kate	long.q	running
	magmasteel	sge	james	medium.q	queued
	magmasteel	sge	james	medium.q	queued
	magmasteel	sge	james	medium.q	queued
	magmasteel	sge	james	medium.q	queued
	magmasteel	sge	james	medium.q	running
	xhpl	sge	matthew	short.q	running
	xhpl	sge	matthew	short.q	running
	xhpl	sge	matthew	short.q	running

ShowRemoveHoldReleaseSuspendResumeRefresh

EVENT VIEWER

All Events

	Ack	Time	Cluster	Source	Message
!		18/Sep/2009 17:05:53	Demo Cluster	demohead1	Service ntpd was restarted on demohead1
!		18/Sep/2009 17:05:47	Demo Cluster	demohead1	Service named was restarted on demohead1
!		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service postfix was restarted on demohead1
!		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service dhcpd was restarted on demohead1
!		18/Sep/2009 17:05:45	Demo Cluster	demohead1	Service maui was restarted on demohead1

Ready

Bright Cluster Manager

File Monitoring View Help

RESOURCES

My Clusters

- Demo Cluster
 - Switches
 - switch01
 - switch02
 - switch03
 - switch04
 - switch05
 - Networks
 - externalnet
 - ipminet
 - mpinet
 - slavenet
 - storagenet
 - Power Distribution Units
 - apc01
 - apc02
 - apc03
 - apc04
 - Software Images
 - default-image
 - Node Categories
 - slave
 - Head Nodes
 - demohead1
 - demohead2
 - Slave Nodes
 - Other Devices
 - Node Groups
 - Large Memory Nodes
 - Users & Groups
 - Workload Management
 - Monitoring Configuration**
 - Authorisation
 - Authentication

Monitoring Configuration

Overview Setup Metrics Actions

Category	Metric	Parameter	Threshold Bound	Action	Action Parameter
All Master Nodes	FreeSpace	/	< 10 GB	NotifyVendor	
All Master Nodes	FreeSpace	/	< 10 GB	SendEmail	administrator@localhost
All Master Nodes	FreeSpace	/home	< 10 GB	NotifyVendor	
All Master Nodes	FreeSpace	/home	< 10 GB	SendEmail	administrator@localhost
All Power Distributio...	PDULoad		> 32 A	SendEmail	datacenter_support@uni.edu
slave	Temperature		> 70	SendEmail	administrator@localhost
slave	Temperature		> 70	Shutdown	

Monitoring Rules Wizard

Select Category:

- All Power Distribution Units
- All Ethernet Switches
- All Myrinet Switches**
- All IB Switches
- All Master Nodes
- All Rack Sensors
- All Generic Devices

Cancel Previous Next

Edit Add (Wizard) Remove Refresh Save

EVENT VIEWER

All Events

Ack	Time	Cluster	Source	Message
1	18/Sep/2009 18:30:06	Demo Cluster	demohead1	node003 Installing
1	18/Sep/2009 18:29:39	Demo Cluster	demohead1	New certificate request with ID: 5
1	18/Sep/2009 18:29:36	Demo Cluster	demohead1	node002 Installing
1	18/Sep/2009 18:29:25	Demo Cluster	demohead1	New certificate request with ID: 4
1	18/Sep/2009 17:05:53	Demo Cluster	demohead1	Service ntpd was restarted on demohead1
1	18/Sep/2009 17:05:47	Demo Cluster	demohead1	Service named was restarted on demohead1

Ready

- This supercomputer is called "LOEWE-CSC"
 - 20,784 processor cores
 - 772 GPGPU hardware accelerators
 - Theoretical peak performance of 599 TeraFLOPs (599 trillion calculations per second) in double precision, 2259 TeraFLOPS in single precision
- On today's TOP500 list of fastest computers in the world, LOEWE-CSC could rank amongst the top 15
- Some of the reasons for which Bright Cluster Manager was chosen...



Advanced Features

- Redundant head nodes
- Cluster health checking
 - Health checks before jobs are run
 - Frequently scheduled health checks
 - Hardware burn-in environment
- Daemon with low resource consumption
- Multiple, load-balanced provisioning nodes
 - Allows scaling to thousands of nodes
- Node discovery using Ethernet switch port detection
- Live & incremental image updates
- Automated BIOS updates and configurations

Customers — Industry



Customers — Government



Customers — Academia



UNIVERSITÀ DEGLI STUDI
DI MODENA E REGGIO EMILIA



UNIVERSITEIT VAN AMSTERDAM



UNIVERSITY OF
CAMBRIDGE



Friedrich-Schiller-Universität Jena



University of Zurich



“We chose ClusterVision based on their proven expertise and experience, their comprehensive cluster management environment and their professional approach.”



“The engineers from ClusterVision did an excellent job in delivering the cluster.”

— Daniel Schmode, Germanischer Lloyd

"We would like to thank you for your help and support during your Aramco visit."



"It's been a real pleasure working with dedicated professionals like you in this project."

— Manager at Saudi Aramco

"Our partnership with ClusterVision has greatly helped us to set-up our new cluster in a fast time frame."



"I appreciate their reactivity and availability; this is a key point for our activity."

— Dr Christian Fauchier, CD-adapco

“We have chosen ClusterVision as a repeat supplier based on their technical expertise and the value and service they have delivered.”





“We know we can rely on their professional approach to sales, performance and support.”

— Professor John Chew, University of Surrey

Conclusions



- Bright Cluster Manager
 - Fundamental & integrated approach
 - Easy installation, management and use
 - Suitable for small to very large clusters
- Available through ClusterVision, which can also provide you with:
 - Tailor made clusters
 - Tailor made services
 - Tailor made support
 - HPC related consultancy & advice



Questions?

The End