International Outlook

Maurizio Davini

Dipartimento di Fisica e INFN Pisa

maurizio.davini@df.unipi.it

Components

- Extensions for clustering
- Configurations,installation and administration
- Monitoring
- Batch e schedulers
- Cluster File Systems

Extension for clustering

Bproc

- http://bproc.sourceforge.net/
- BProc, The Beowulf Distributed Process Space, is a set of kernel modifications, utilities, and libraries which allow a user to start processes on other machines in a Beowulf-style cluster. Remote processes started with this mechanism appear in the process table of the front end machine in the cluster.

OpenMosix

- http://www.openmosix.org/
- is a software package that enhances the Linux kernel with cluster capabilities. MOSIX allows for the automatic and transparent migration of processes to other nodes in the cluster, while standard Linux process control utilities, such as 'ps' will show all processes as if they are running on the node the process originated from

SSI

- http://ssic-linux.sourceforge.net/
- Single System Image Clusters for Linux (SSI) aims at providing a full, highly available, single system image cluster environment for Linux, with the goals of availability, scalablity, and manageability, built from standard servers.

CI (cluster Infrastructure)

- http://ci-linux.sourceforge.net/
- Cluster Infrastructure for Linux (CI) aims at developing a common infrastructure for Linux clustering by extending cluster membership and internode communication subsystems from Compaq's NonStop Clusters for Unixware code base. This project also provides the basis for the SSI Clusters for Linux project.

Configuration and Management

CFEngine

- http://www.gnu.org/software/cfengine/cf
- Cfengine is a tool for administration and configuration of large (or small) networks of computers. It uses the idea of classes to define and automate the configuration and maintenance of systems.

LCFG

- http://www.lcfg.org/
- LCFG (Local Configuration System) is a system for automatically installing and managing the configuration of large numbers of Unix systems. It is particularly well suited to environments with diverse and rapidly changing configurations

C3

- Cluster Control e Command Tool
- http://www.csm.ornl.gov/torc/C3/
- The C3 (Cluster Command and Control) tool suite was developed at Oak Ridge National Lab and implements a number of command line based tools that aid in the operation and management of clusters.

System Installation Suite

- http://sisuite.org/
- The System Installation Suite is the grand project name for the joint effort of System Configurator, System Imager, and System Installer. All three are being developed to work together to provide an easy to use solution for installing and configuring a cluster of heterogenous hardware and to be distribution agnostic

ImgDisk

- http://www.unix.mcs.anl.gov/systen
- Imgdisk can save and restore disk images: the partition table, MBR, partition boot records, ext2 file system contents, and can also restore swap partitions.

Webmin

- http://www.webmin.com/webmin/
- Webmin is a web based interface for system administration of Unix systems.

Monitoring

GANGLIA

- http://ganglia.sourceforge.net/
- The Ganglia Cluster Toolkit is a massively scalable cluster monitoring and execution environment. Currently, after the monitoring system has been released the execution environment too.

Ganglia

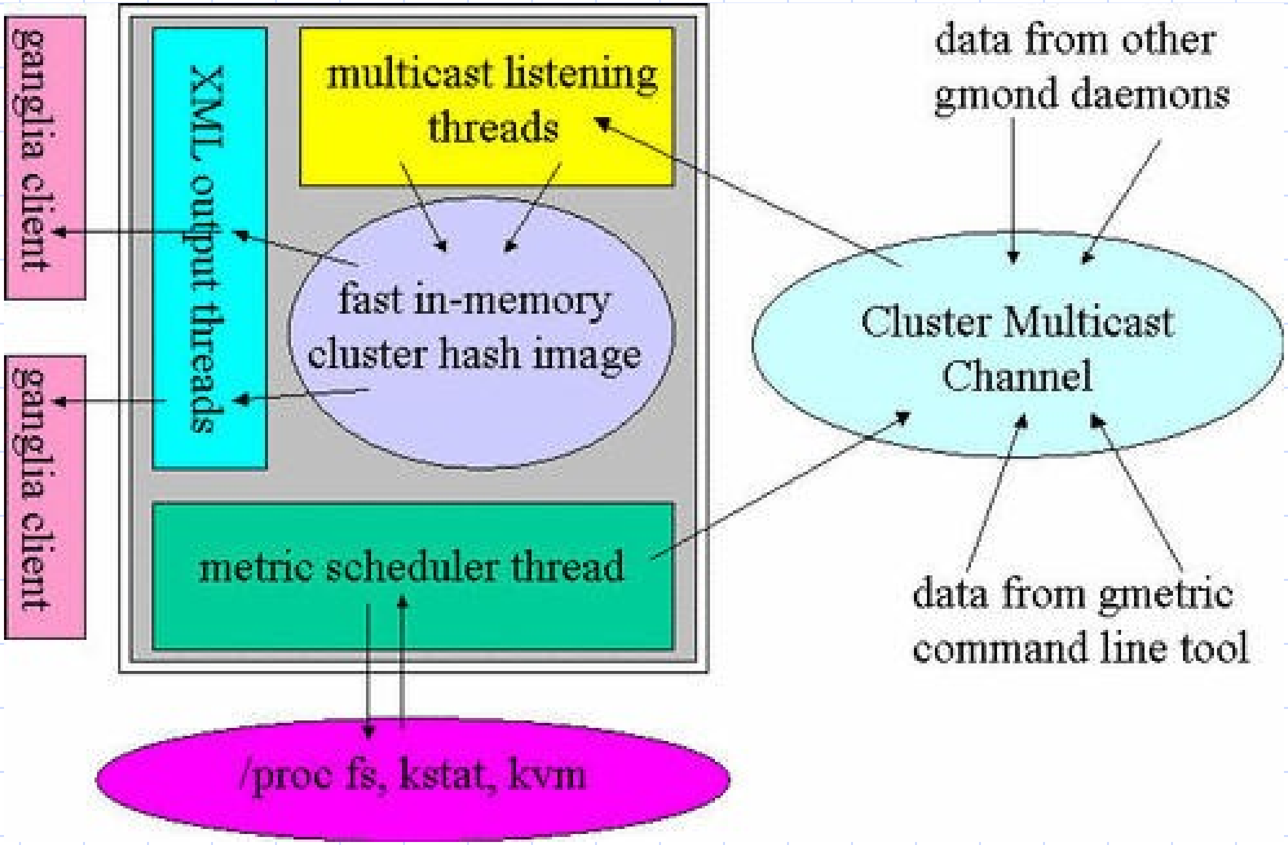
ganglia client
ganglia client
XML output threads
multicast listening threads
fast in-memory cluster hash image
metric scheduler thread
/proc fs, kstat, kvm
data from other gmond daemons
Cluster Multicast Channel
data from gmetric command line tool

Ganglia

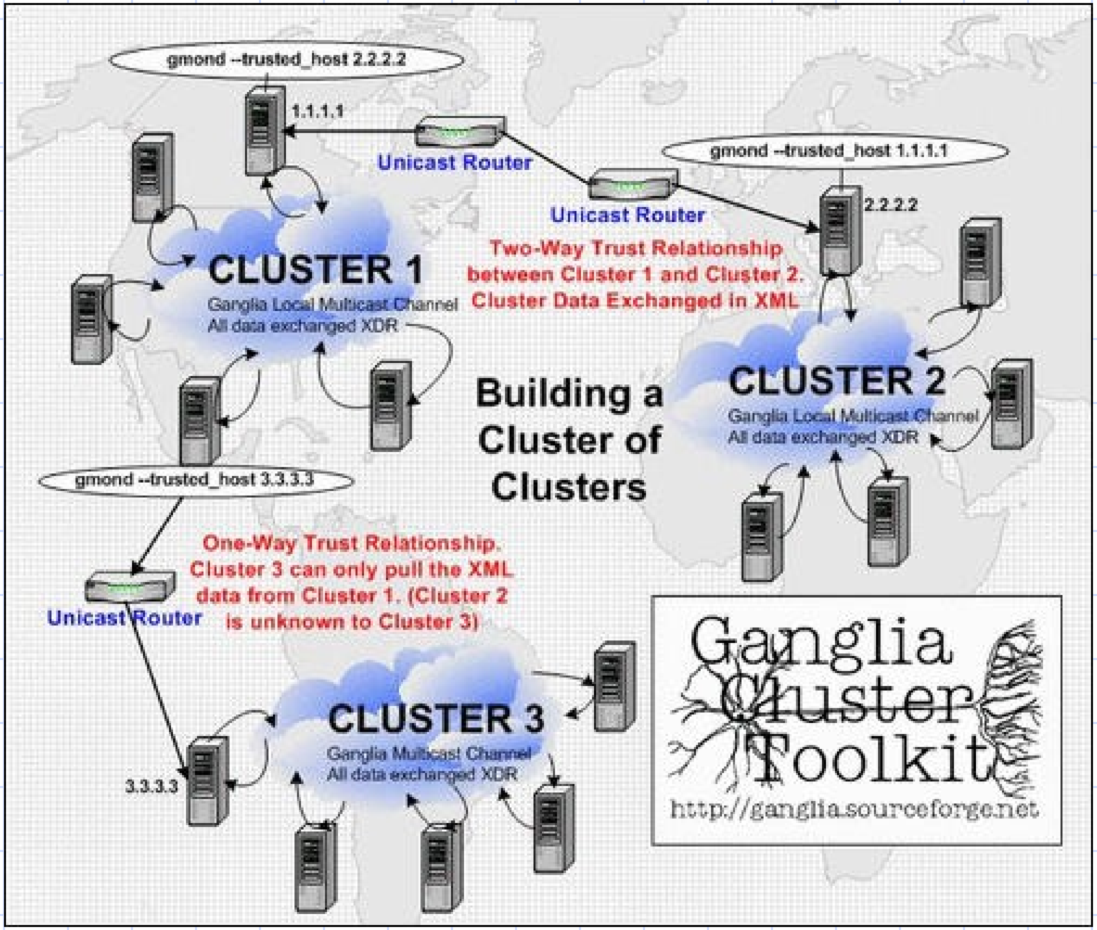
gmond --trusted_host 2.2.2.2
1.1.1.1
Unicast Router
gmond --trusted_host 1.1.1.1
Unicast Router
2.2.2.2
Two-Way Trust Relationship between Cluster 1 and Cluster 2. Cluster Data Exchanged in XML
CLUSTER 1
Ganglia Local Multicast Channel
All data exchanged XDR
Building a Cluster of Clusters
CLUSTER 2
Ganglia Local Multicast Channel
All data exchanged XDR
gmond --trusted_host 3.3.3.3
One-Way Trust Relationship. Cluster 3 can only pull the XML data from Cluster 1. (Cluster 2 is unknown to Cluster 3)
Unicast Router
CLUSTER 3
Ganglia Multicast Channel
All data exchanged XDR
3.3.3.3
Ganglia Cluster Toolkit
http://ganglia.sourceforge.net

SuperMon

- http://www.acl.lanl.gov/supermon/
- Supermon is a high speed monitoring system for large scale clusters
- Recent benchmarks have shown that supermon can achieve single node monitoring speeds that are significantly higher than previous methods - we have observed 6000 samples-per-second with supermon, while the same benchmark was only able to achieve 275 samples-per-second with older methods used by RPC.rstatd.

Nagios

- http://www.nagios.org/
- The Nagios Network Monitor can monitor network services and host resources, contains contact notifications when problems occur (via email, pager, or user-defined method), contains an optional web interface for viewing network status, notification, problem history or log files, and is highly extensible, making it easy for the creation of user-developed service checks

VACM

- http://www.valinux.com/software/vacm
- VACM (VA Cluster Manager) is a tool for monitoring and managing a large cluster of nodes. There is also a page for VACM on SourceForge, located here.

Clunix

- http://www.cmap.polytechnique.fr/~sylva
- The clunix tools are a free, GPL'd set of utilities for managing a Linux cluster. The components are a perl based load balancer, a php based web interface for monitoring the state of the cluster, and a daemon.

ECT

- http://www.alphaworks.ibm.com/tech/ect
- Enhanced Cluster Tools for Linux (ECT) is a set of tools which complement IBM's Cluster Systems Management (CSM) and enhance the management of clusters, providing features such as remote access to hardware inventory and vitals, and remote access to service processor logs.

Cluster File Systems

ClusterNFS

- http://clusternfs.sourceforge.net/
- ClusterNFS is a set of patches for the UNFSD server to allow multiple diskless clients to mount the same root filesystem, by providing "interpreted" filenames.

GFS **$$**

- http://www.sistina.com/products_gf
- GFS (Global File System) is fault tolerant and distributed. It is also both a cluster and a journalling file system.

GPFS $$

- IBM’s General Parallel File System (GPFS) allows users shared access to files that may span multiple disk drives on multiple nodes. It offers many of the standard UNIX ® file system interfaces allowing most applications to execute without modification or recompiling. UNIX file system utilities are also supported by GPFS.

PVFS

- http://parlweb.parl.clemson.edu/pvf
- PVFS (Parallel Virtual File System) is being developed at Clemson University's Parallel Architecture Research Lab and is closely tied to the Beowulf (currently a.k.a. Scyld) Project.

LUSTRE

- http://www.lustre.org/
- Lustre is a novel storage and file system architecture that aims at building a next-generation cluster file system, to service clusters with 10,000s of nodes, petabytes of storage, and move 100s of GB/sec, as well as offering security and management.

Batch systems e schedulers

Sun Grid Engine

- http://gridengine.sunsource.net/
- The Gride Engine Project is an open source project based on Sun's commercial product, "Sun Grid Engine," which can be seen here. Grid Engine is Distributed Resource Management software, used to create compute farms.

LSF $$

- http://www.platform.com/
- LSF (Load Sharing Facility) is a suite of software available for various Unixes and NT. It performs load sharing and balancing, and job scheduling.

PBS $$

- http://pbs.mrj.com/
- PBS (Portable Batch System) is a batch queueing and load balancing system originally developed for NASA. It is available for a variety of Unix platforms. There is an Open Source version of PBS, called OpenPBS, which is located here.

Condor

- http://www.cs.wisc.edu/condor
- Condor is a software package that does job scheduling and load balancing. It is available for most Unixes, and a port to NT is currently underway.

MAUI

- http://supercluster.org/maui/
- Maui is an advanced scheduler for clusters and supercomputers. It can support various fairness policies, dynamic priorities, extensive reservations, and fairshare

Installation systems

Cluster distributions

- http://clusters.top500.
- Why are 26% rolling there own clusters?
  - This is the real message of the poll

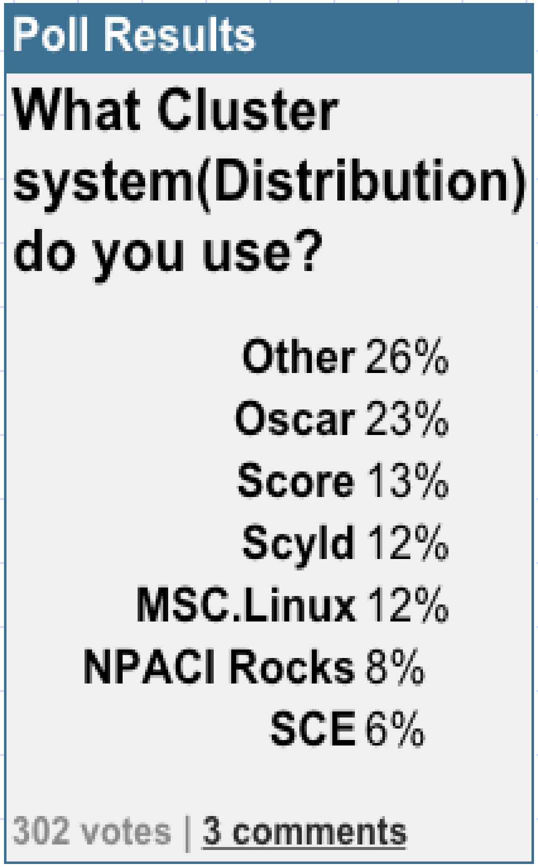

Poll Results

What Cluster system(Distribution) do you use?

Other	26%
Oscar	23%
Score	13%
Scyld	12%
MSC.Linux	12%
NPACI Rocks	8%
SCE	6%

302 votes | 3 comments

LTSP

- http://www.ltsp.org/
- The name pretty much says it all. The LTSP is a way to run thin client computers in a Linux environment.
- Clustering openMosix as an option

FAI

- http://www.informatik.uni-koeln.de/fai
- FAI (Fully Automatic Installation) is a non-interactive system to install Debian GNU/Linux on a cluster. It is a collection of Perl and shell scripts, and will work with a variety of PC hardware

KA

- http://ka-tools.sourceforge.net/
- Ka is a toolkit designed to install and administer a cluster of boxes. It focus on scalability of parallel system installation, data distribution and process launching. Ka has been tested on clusters up to 225 nodes

Clubmask

- clubmask.sourceforge.net
- Clubmask is a collection of existing Open Source and new software for the installation, configuration, and management of Beowulf style high performance computing clusters. The design and goal of the project is to provide a "Physicist Proof", completely turnkey set of tools

Clubmask

1. Bproc - unified cluster process namespace and control
2. Cfengine - class based configuration tool
3. Kickstart - RedHat automated installation scripts
4. Lam - MPI environment
5. Maui Scheduler - advanced batch scheduler with a large feature set well suited for high performance computing (HPC) platform
6. ZODB - Embedded Python Object

ClusterIt

- http://www.garbled.net/clusterit.html
- ClusterIt is a collection of clustering tools, allowing for heterogeneous cluster makeup, including various architectures and operating systems, various authentication mechanisms, job sequencing, distributed virtual terminals, and more.

SCORE

- http://www.pccluster.org/
- SCore, by Real World Computing Partnership (RWCP) is not a Beowulf style cluster in the sense that SCore software is designed for the high performance cluster environment without using the TCP/IP stack.

Score

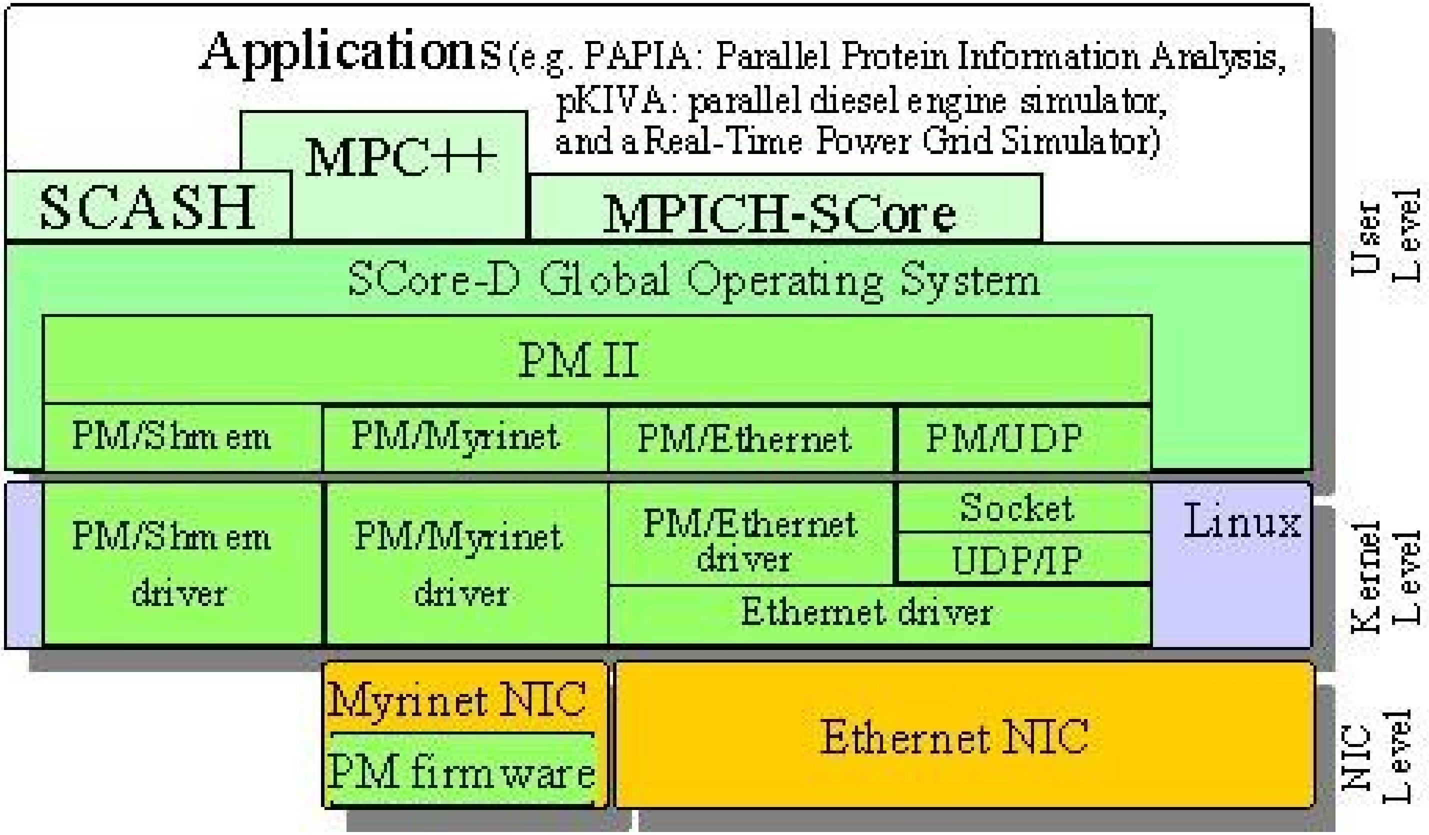
Applications (e.g. PAPIA: Parallel Protein Information Analysis,
pKIVA: parallel diesel engine simulator,
and a Real-Time Power Grid Simulator)
MPC++
SCASH
MPICH-SCore
SCore-D Global Operating System
PM II
PM/Shmem
PM/Myrinet
PM/Ethernet
PM/UDP
PM/Shmem driver
PM/Myrinet driver
PM/Ethernet driver
Socket
UDP/IP
Ethernet driver
Linux
Myrinet NIC
PM firmware
Ethernet NIC
User Level
Kernel Level
NIC Level

WareWulf

- http://www.runlevelzero.net/greg/warewu
- Warewulf is a distribution of tools that are designed to aid in the implementation of Beowulf style clusters. The software is a bootable ISO image that is easily modified, and slave node filesystems can be booted off the CDROM, so no changes must be made to the existing hard disk if wanted.

Cplant

- http://www.cs.sandia.gov/cplant/
- Computational Plant (a.k.a. Cplant) is a newly released project coming from the folks at the Sandia National Laboratories. The goal is "to provide a commodity-based, large-scale computing resource that meets the level of compute performance needed by Sandia's critical applications."

ClusterMatic

- http://www.clustermatic.org/
- Clustermatic is a collection of technologies being developed at the Cluster Research Lab at Los Alamos National Laboratory. Besides the new software being developed by the group, existing projects such as LinuxBIOS and BProc are integrated into it as well.

ClusterMatic

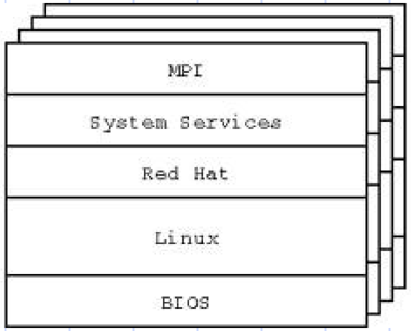

MPI

System Services

Red Hat

Linux

BIOS

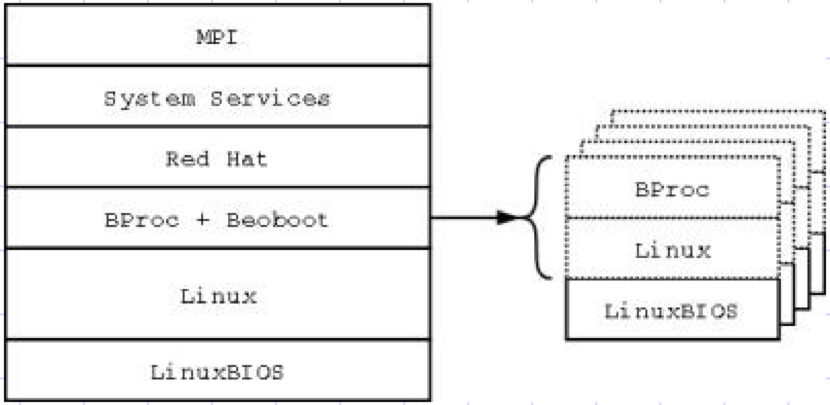

MPI

System Services

Red Hat

BProc + Beoboot

Linux

LinuxBIOS

BProc

Linux

LinuxBIOS

OSCAR

- http://oscar.sourceforge.net/
- OSCAR (Open Source Cluster Application Resources) is a bundle of software designed to make it easy to build, maintain, and use a Linux cluster

OSCAR – An Overview

- Open Source Cluster Application Resources
- Cluster on a CD – automates cluster install process
- IBM, Intel, NCSA, ORNL, MSC Software, Dell
- NCSA “Cluster in a BOX” base
- Wizard driven
- Nodes are built over network
- OSCAR <= 64 node clusters for initial target
- Works on PC commodity components
- RedHat based (for now)
- Components: Open source and BSD style license

Why OSCAR?

- NCSA wanted “Cluster-in-a-Box” Distribution
  - NCSA’s “X-in-a-Box” projects could lie on top
  - X = Grid, Display Wall, Access Grid
- Easier, faster deployment
- Consistency among clusters
- Other organizations had the same interest
  - Intel, ORNL, Dell, IBM, etc.
  - NCSA jumps on board to contribute to OSCAR

OSCAR Basics

- What does it do?
  - OSCAR is a cluster packaging utility
  - Automatically configures software components
  - Reduces time to build a cluster
  - Reduces need for expertise
  - Reduces chance of incorrect software configuration
  - Increases consistency from one cluster to the next
- What *will* it do in the future?
  - Maintain cluster information database
  - Work as an interface not just for installation, but also for maintenance
  - Accelerate software package integration into clusters

OSCAR Basics

How does it work?

- version 1.0, 1.1
  - LUI executes
  - = Linux Utility for cluster Install
    - Network boots nodes via PXE or floppy
    - Nodes install themselves from rpms over NFS from the server
    - Post installation configuration of nodes and server executes
- version 1.2+
  - SIS = System Installation Suite
    - System Imager + LUI = SIS
    - Creates image of node filesystem locally on server
    - Network boots nodes via PXE or floppy
    - Nodes synchronize themselves with server via rsycn
    - Post installation configuration of nodes and server

Components

- OSCAR includes (currently):
  - C3 – Cluster Management Tools (ORNL)
  - SIS – Network OS Installer (IBM)
  - MPI-CH – Message Passing Interface
  - OpenSSH/OpenSSL – Secure Transactions
  - PBS – Job Queuing System
  - PVM – Parallel Virtual Machine
- Current Prerequisites:
  - Networked PC hardware with disk drives
  - Server machine with Redhat installed
  - Redhat CD(s)
  - 1 head node + N compute nodes

Installation Overview

- Install RedHat
- Download OSCAR
- Print/Read document
- Copy RPMS to server
- Run wizard (install_cluster)
  - Build image per client type (partition layout, HD type)
  - Define clients (network info, image binding)
  - Setup networking (collect MAC addresses, configure DHCP, build boot floppy)
  - Boot clients / build
  - Complete setup (post install)
  - Install test suite
- Use cluster

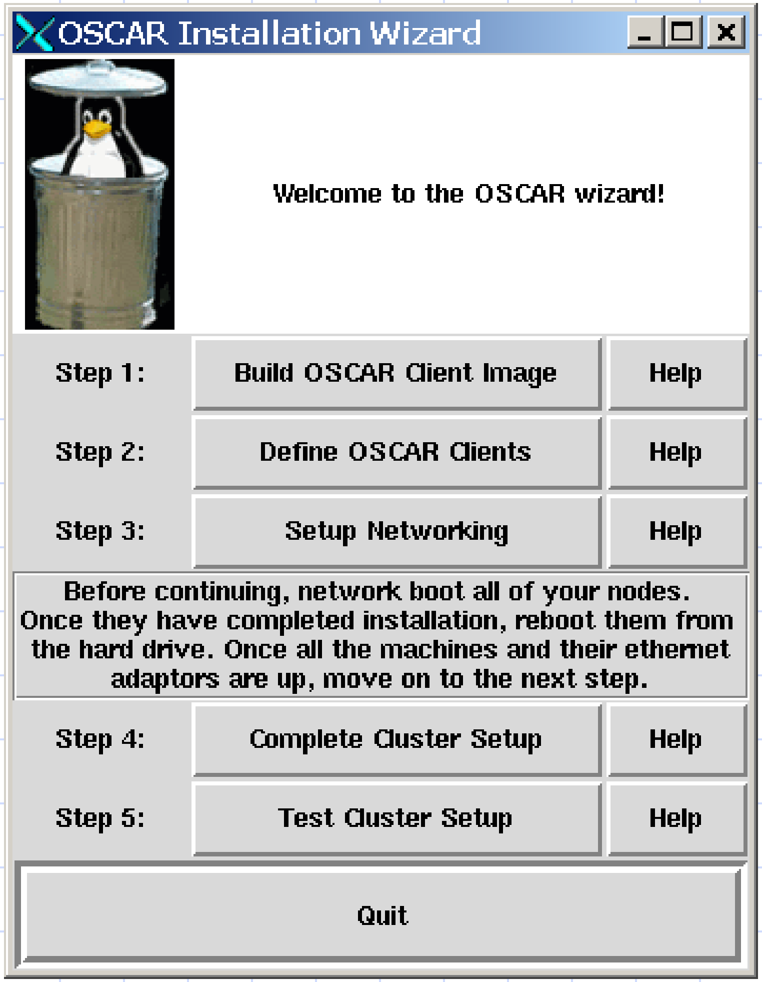

Welcome to the OSCAR wizard!

Before continuing, network boot all of your nodes. Once they have completed installation, reboot them from the hard drive. Once all the machines and their ethernet adaptors are up, move on to the next step.

OSCAR 2

- Major Changes - Summary
  - No longer bound to OS installer
  - Components are package based, modular
  - Core set of components mandatory
  - API established and published for new packages
  - Package creation open to community
  - Database maintained for node and package information
  - Add/Remove Node process will be improved
  - Web based wizard
  - Scalability enhancements
  - Security Options
  - Support more distributions and architectures
  - New Features

CluMon

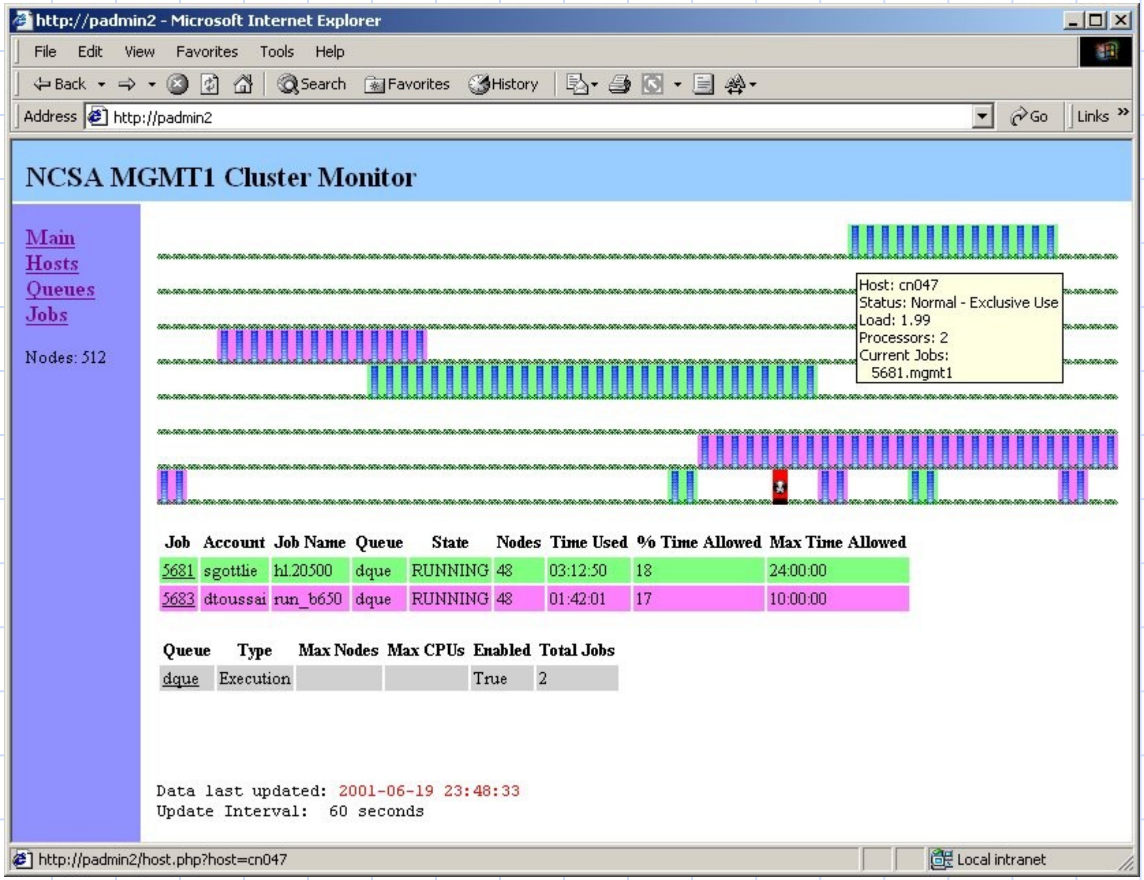
http://padmin2 - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites History
Address http://padmin2
Go Links
Main
Hosts
Queues
Jobs
http://padmin2/host.php?host=cn047
Local intranet

NCSA MGMT1 Cluster Monitor

Nodes: 512

Host: cn047
Status: Normal - Exclusive Use
Load: 1.99
Processors: 2
Current Jobs:
5681.mgmt1

Job	Account	Job Name	Queue	State	Nodes	Time Used	% Time Allowed	Max Time Allowed
5681	sgottlie	h1.20500	dque	RUNNING	48	03:12:50	18	24:00:00
5683	dtoussai	run_b650	dque	RUNNING	48	01:42:01	17	10:00:00

Queue	Type	Max Nodes	Max CPUs	Enabled	Total Jobs
dque	Execution			True	2

```
Data last updated: 2001-06-19 23:48:33
Update Interval:  60 seconds
```

ROCKS

- http://rocks.npaci.edu/
- The Rocks Clustering Toolkit, from the folks at NPACI, is a collection of Open Source tools to help build, manage, and monitor, clusters.

Who is Using It?

- Growing (and partial) list of users that we know about:
  - SDSC, SIO, UCSD (8 Clusters, including CMS (GriPhyN) prototype)
  - Caltech
  - Burnham Cancer Institute
  - PNNL (several clusters, small, medium, large)
  - University of Texas
  - University of North Texas
  - Northwestern University
  - University of Hong Kong
  - Compaq (Working relationship with their Intel Standard Servers Group)
  - Singapore Bioinformatics Institute
  - Myricom (Their internal development cluster)

Major Components

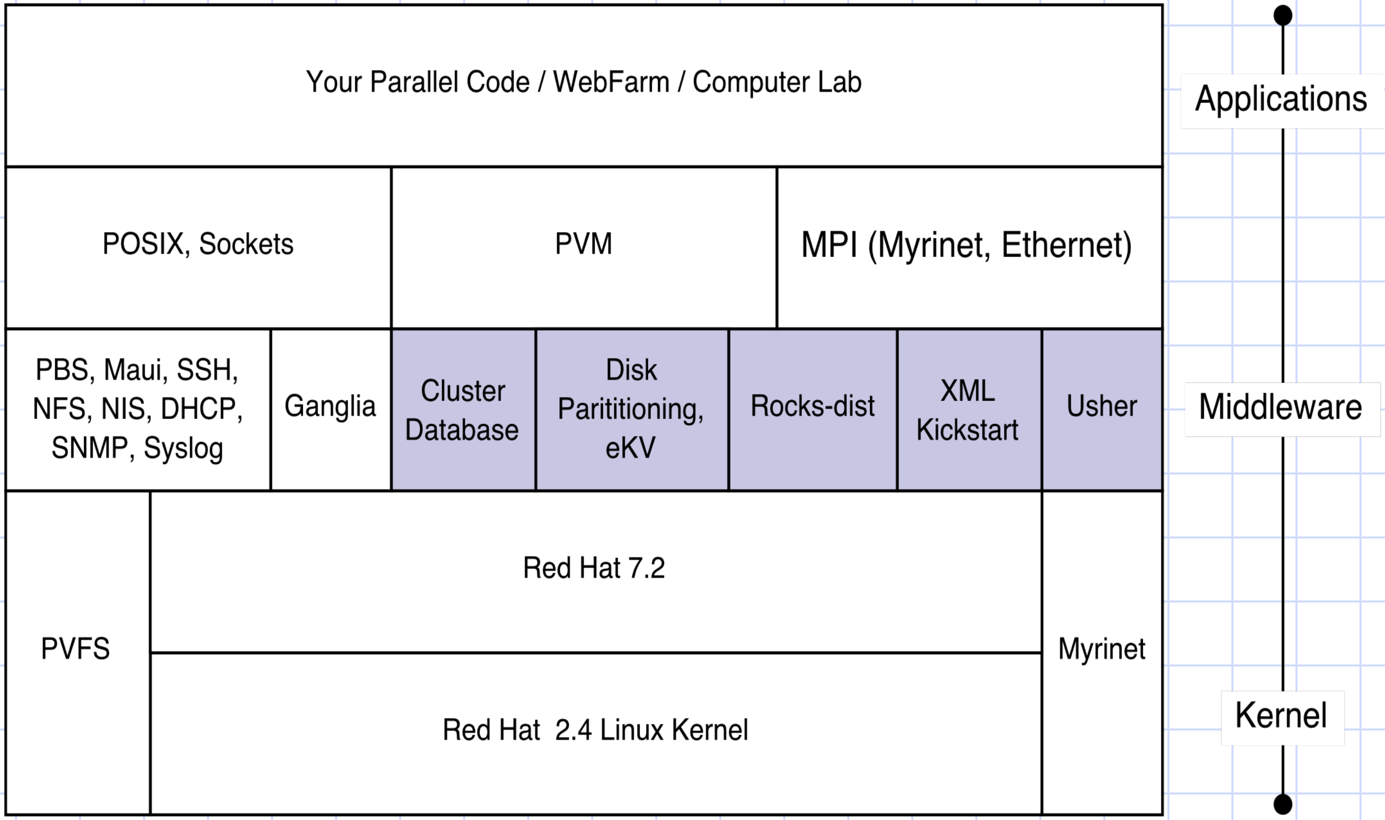

Your Parallel Code / WebFarm / Computer Lab

POSIX, Sockets

PVM

MPI (Myrinet, Ethernet)

PBS, Maui, SSH, NFS, NIS, DHCP, SNMP, Syslog

Ganglia

Cluster Database

Disk Parititioning, eKV

Rocks-dist

XML Kickstart

Usher

PVFS

Red Hat 7.2

Red Hat 2.4 Linux Kernel

Myrinet

Applications

Middleware

Kernel

SCE

- http://www.opensce.org/
- SCE (Scalable Cluster Environment) is an easy to use set of interoperable Open Source tools that allows the user to quickly install, configure, and use, a Beowulf cluster.

SCE

- AMATA High Availability support for Beowulf Cluster
- Beowulf Diskless Cluster Suite a utility that help you build a diskless Beowulf cluster easily.
- KCAP Web and VRML based system monitoring and navigation tool for large scale cluster
- KSIX Middle ware layer that offer powerful programing extension in cluster environment
- SCMS SMILE Cluster Management System . A powerful system administration tools for Beowulf style cluster
- SQMS Flexible and extensible batch scheduling system for beowulf

SCE

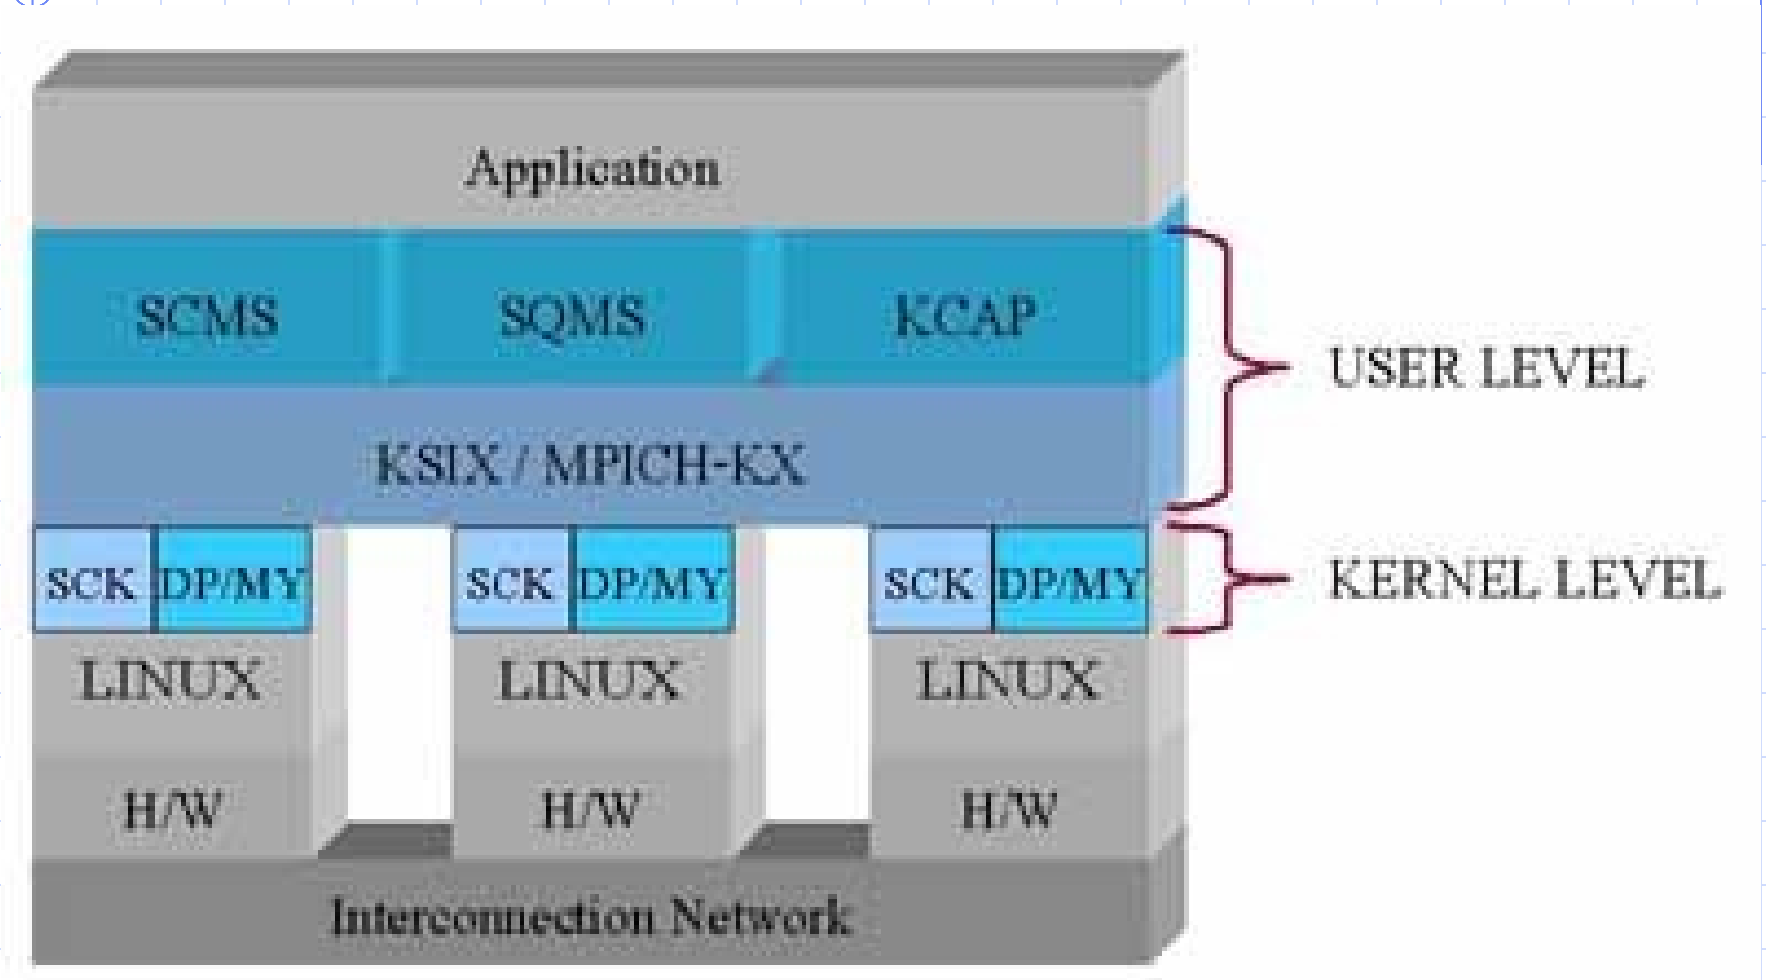
Application
SCMS
SQMS
KCAP
USER LEVEL
KSIX / MPICH-KX
SCK
DP/MY
SCK
DP/MY
SCK
DP/MY
KERNEL LEVEL
LINUX
LINUX
LINUX
H/W
H/W
H/W
Interconnection Network

MSC.Linux

- http://www.msclinux.com/software/
- MSC.Linux is a high performance/cluster distribution that is designed for computational environments in engineering and life sciences.

ClusterWorX **$$**

- http://www.linuxnetworx.com/produ php
- ClusterWorX, from the folks at Linux NetworX, is a management and monitoring package, that supports customized monitoring and notification, disk cloning, node management, and more

Alinka Raisin $$

- www.alinka.com
- Alinka's Raisin software package can do everything from creation to administration of High Performance Linux clusters. It uses Batch Queuing systems such as PBS (or others such as LSF or NQE upon request), Mosix for process migration, and the Parallel file systm PVFS, and comes with a web based user interface

Cluster System Management(CSM)

- http://www1.ibm.com/servers/eserver/clusters/s
- Cluster Systems Management for Linux (CSM) enables organizations to build, manage, and expand clusters. A free demonstration of CSM, called the Cluster Starter Kit, is available here.

SCALI $$

- http://www.scali.com/products/ssp.html
- The Scali Software Platform (SSP) delivers a number of tools for ease of installation, administration, maintenance, and operational use of clusters ranging from a handful to hundreds of nodes, that targets all aspects of building, maintaining, and using a cluster. It covers everything from low level drivers to high level administration.

SCYLD Beowulf $$

- http://www.scyld.com/
- The Scyld Beowulf Cluster Operating System is the second generation of The Beowulf clustering software. Scyld Computing Corporation was started by Donald Becker and few other folks from the original Beowulf Project team.

Qlusters $$

- www.qlusters.com
- For the openMosix Enthusiasts….

Qlusters OS features (1)

- Based in part on openMosix technology
- Migrating sockets
- Network RAM already implemented
- Cluster parallel Installer,
- Cluster Configurator,
- Qsense (automatic detection of nodes no-more /etc/mosix.map)
- Cluster Monitor (written in Flash)
- Logical Clusters (separate domain of clustering and administration)

Qlusters OS features (2)

- Automatic update
- New Load Balancer
- Threaded applications migration
- Linux kernel 2.4.18 with (VM by A.Arcangeli integrated with Reverve Mapping by R.V.Riel)
- Over 100 patches (RedHat Quality)
- Kernel latency reduced by 65% due to Robert Love latest pre-emption patch

Qlusters OS features (3)

- Queue Manager ,Launcher, Scheduler
- Job Description Language in XML
- Integration with GFS completed
- Integration with AFS planned
- IBM xSeries NUMA support
- DSM in a few months
- Port to IA64 underway

Qluster Os features (4)

- grid with multiplatform consideration (recompiles when transferring on a cluster of different architecture)
- Grid component (link Q-OS clusters),OS independent (AIX,Solaris…)
- **Will be the RedHat HPC cluster solution**?

The Installer

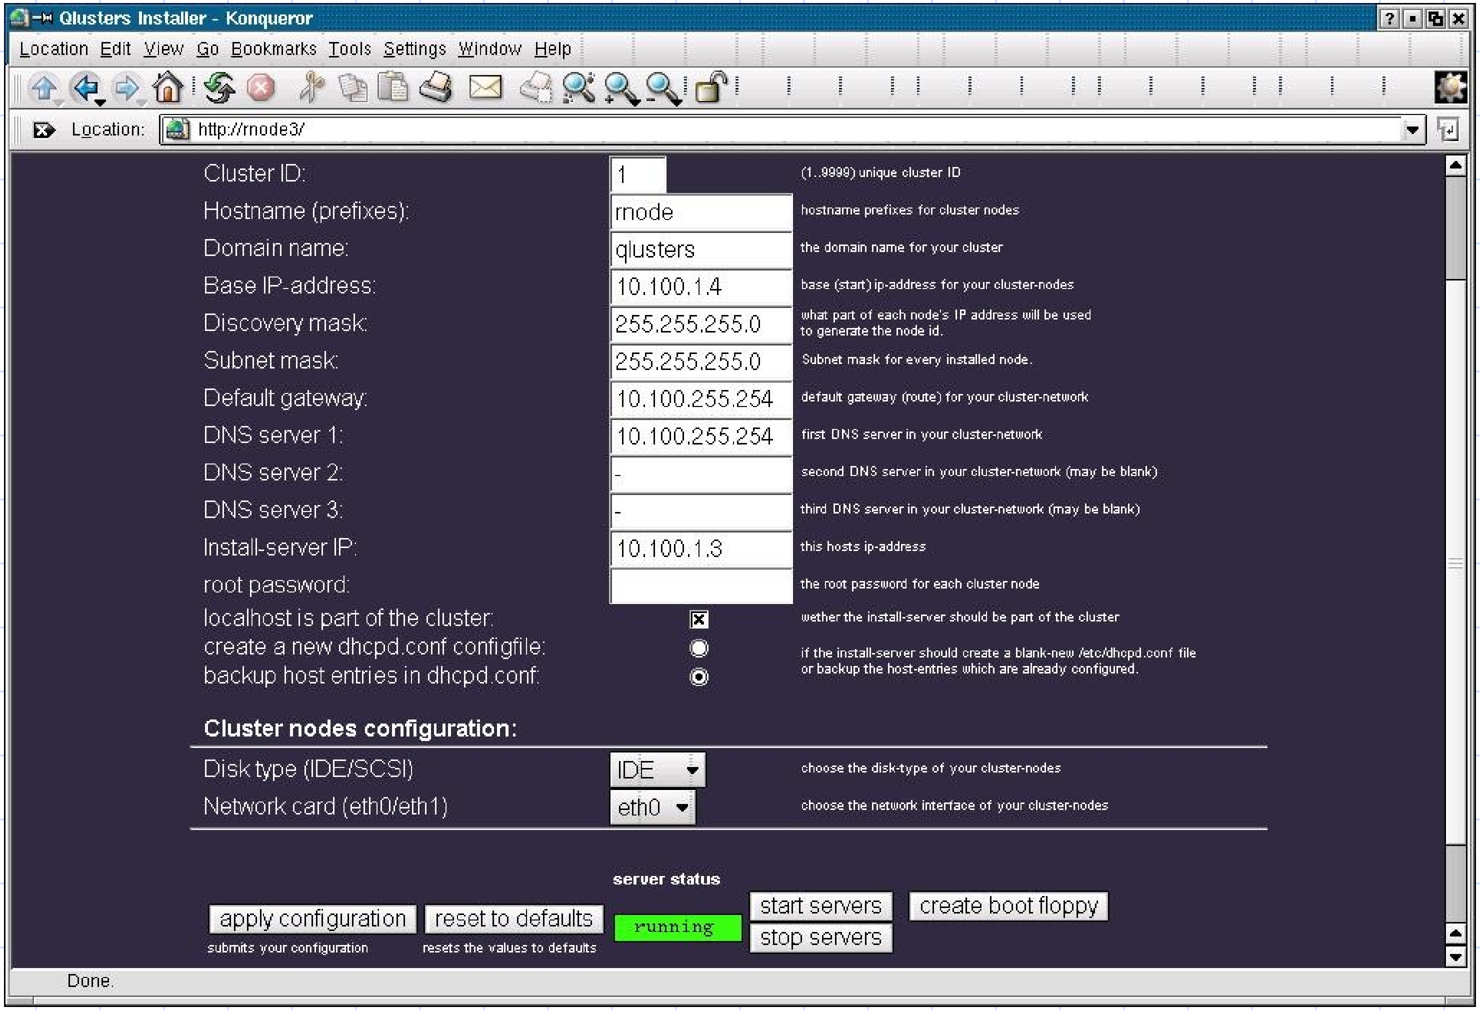
Qlusters Installer - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location:

http://rnode3/

Cluster ID:	1	(1..9999) unique cluster ID
Hostname (prefixes):	rnode	hostname prefixes for cluster nodes
Domain name:	qlusters	the domain name for your cluster
Base IP-address:	10.100.1.4	base (start) ip-address for your cluster-nodes
Discovery mask:	255.255.255.0	what part of each node's IP address will be used to generate the node id.
Subnet mask:	255.255.255.0	Subnet mask for every installed node.
Default gateway:	10.100.255.254	default gateway (route) for your cluster-network
DNS server 1:	10.100.255.254	first DNS server in your cluster-network
DNS server 2:	-	second DNS server in your cluster-network (may be blank)
DNS server 3:	-	third DNS server in your cluster-network (may be blank)
Install-server IP:	10.100.1.3	this hosts ip-address
root password:		the root password for each cluster node
localhost is part of the cluster:	☒	wether the install-server should be part of the cluster
create a new dhcpd.conf configfile:	○	if the install-server should create a blank-new /etc/dhcpd.conf file or backup the host-entries which are already configured.
backup host entries in dhcpd.conf:	●	

Cluster nodes configuration:

Disk type (IDE/SCSI)	IDE	choose the disk-type of your cluster-nodes
Network card (eth0/eth1)	eth0	choose the network interface of your cluster-nodes

server status

running

The Installer

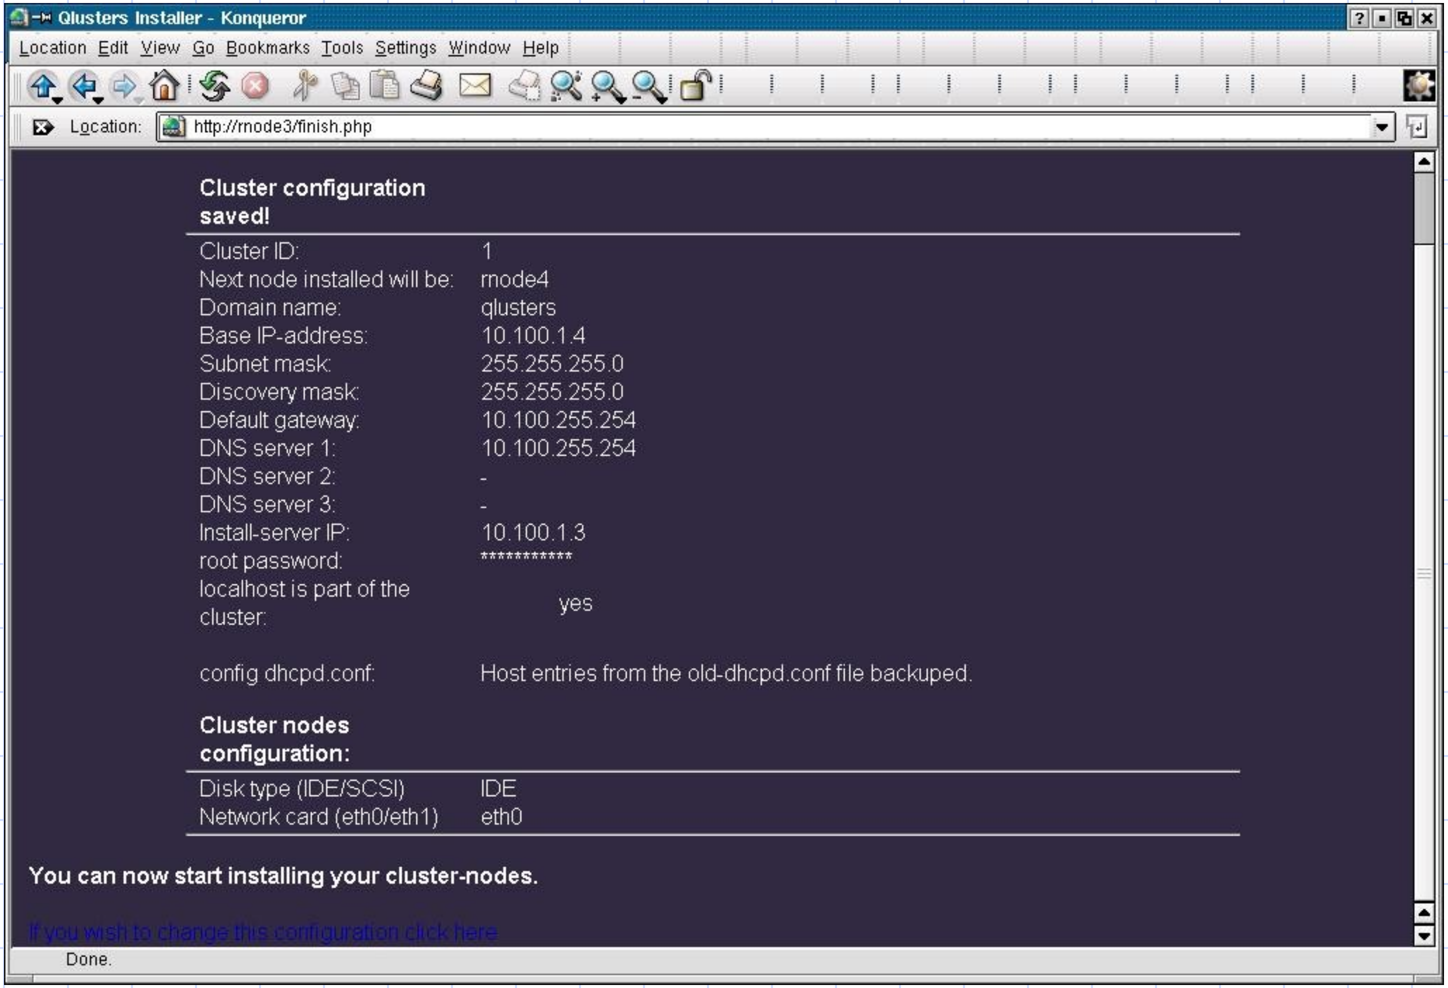
Qlusters Installer - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help

Location: http://rnode3/finish.php

Cluster configuration saved!

Cluster ID:	1
Next node installed will be:	rnode4
Domain name:	qlusters
Base IP-address:	10.100.1.4
Subnet mask:	255.255.255.0
Discovery mask:	255.255.255.0
Default gateway:	10.100.255.254
DNS server 1:	10.100.255.254
DNS server 2:	-
DNS server 3:	-
Install-server IP:	10.100.1.3
root password:	***********
localhost is part of the cluster:	yes
config dhcpd.conf:	Host entries from the old-dhcpd.conf file backuped.

Cluster nodes configuration:

Disk type (IDE/SCSI)	IDE
Network card (eth0/eth1)	eth0

You can now start installing your cluster-nodes.

If you wish to change this configuration click here

Done.

The Installer

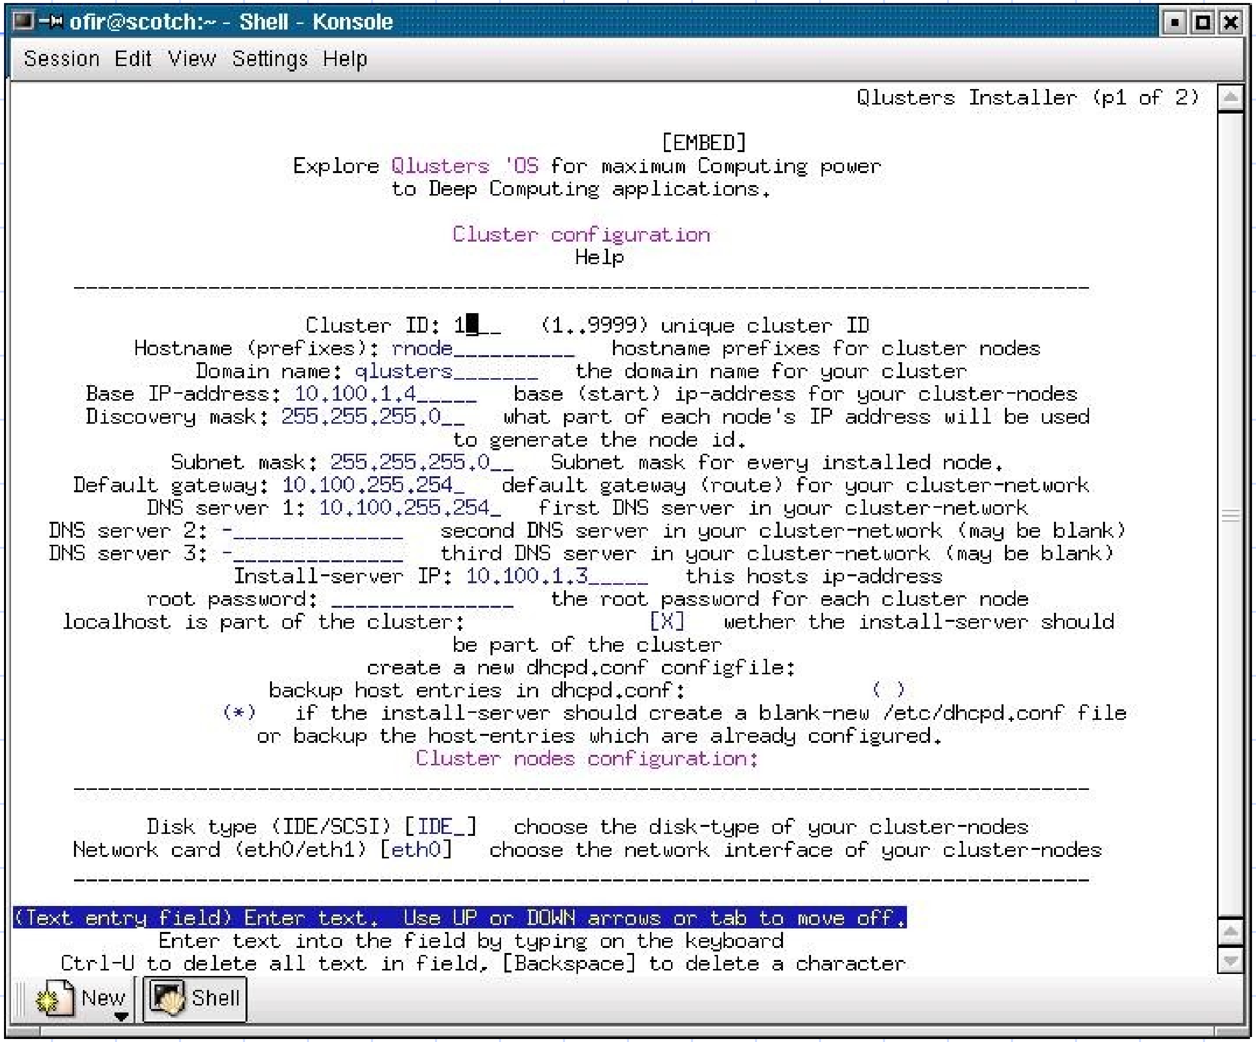
ofir@scotch:~ - Shell - Konsole
Session Edit View Settings Help
New
Shell

```
                                                                     Qlusters Installer (p1 of 2)

                                                     [EMBED]
                       Explore Qlusters 'OS for maximum Computing power
                               to Deep Computing applications.

                                    Cluster configuration
                                              Help

     -----------------------------------------------------------------------------------

                        Cluster ID: 1___   (1..9999) unique cluster ID
          Hostname (prefixes): rnode__________   hostname prefixes for cluster nodes
               Domain name: qlusters_______   the domain name for your cluster
      Base IP-address: 10.100.1.4_____   base (start) ip-address for your cluster-nodes
      Discovery mask: 255.255.255.0__   what part of each node's IP address will be used
                                    to generate the node id.
             Subnet mask: 255.255.255.0__   Subnet mask for every installed node.
     Default gateway: 10.100.255.254_   default gateway (route) for your cluster-network
           DNS server 1: 10.100.255.254_   first DNS server in your cluster-network
   DNS server 2: -______________   second DNS server in your cluster-network (may be blank)
   DNS server 3: -______________   third DNS server in your cluster-network (may be blank)
                  Install-server IP: 10.100.1.3_____   this hosts ip-address
           root password: _______________   the root password for each cluster node
    localhost is part of the cluster:               [X]   wether the install-server should
                                    be part of the cluster
                             create a new dhcpd.conf configfile:
                     backup host entries in dhcpd.conf:               ( )
                 (*)   if the install-server should create a blank-new /etc/dhcpd.conf file
                    or backup the host-entries which are already configured.
                                 Cluster nodes configuration:

     -----------------------------------------------------------------------------------

           Disk type (IDE/SCSI) [IDE_]   choose the disk-type of your cluster-nodes
     Network card (eth0/eth1) [eth0]   choose the network interface of your cluster-nodes
     -----------------------------------------------------------------------------------

(Text entry field) Enter text.  Use UP or DOWN arrows or tab to move off.
            Enter text into the field by typing on the keyboard
    Ctrl-U to delete all text in field, [Backspace] to delete a character
```

The Installer

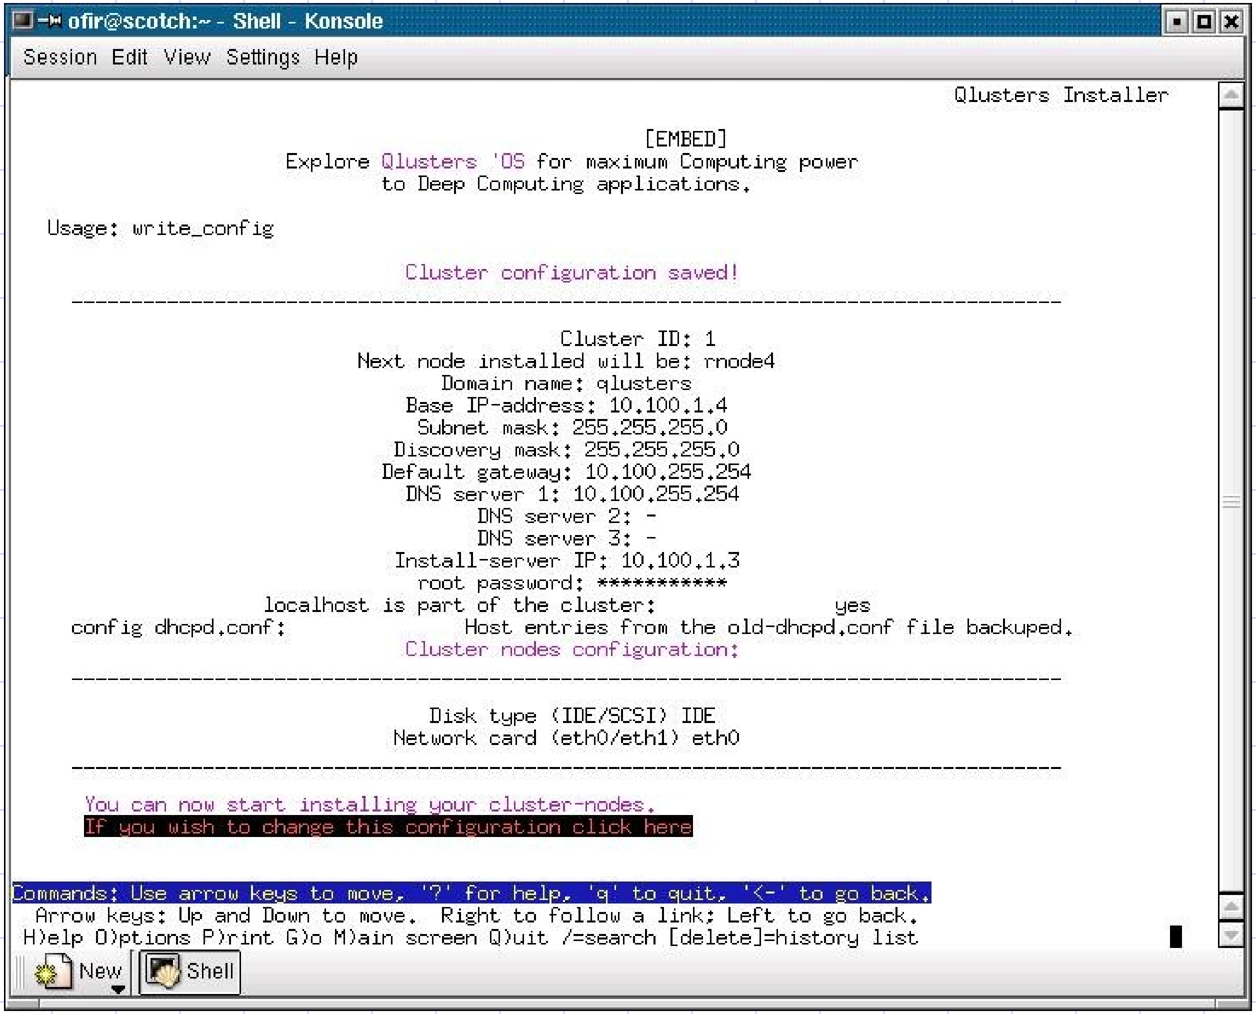
ofir@scotch:~ - Shell - Konsole
Session Edit View Settings Help
New
Shell

```
                                                           Qlusters Installer

                                [EMBED]
              Explore Qlusters 'OS for maximum Computing power
                    to Deep Computing applications.

   Usage: write_config

                     Cluster configuration saved!
     ----------------------------------------------------------------------------------

                                Cluster ID: 1
                   Next node installed will be: rnode4
                          Domain name: qlusters
                       Base IP-address: 10.100.1.4
                        Subnet mask: 255.255.255.0
                      Discovery mask: 255.255.255.0
                     Default gateway: 10.100.255.254
                       DNS server 1: 10.100.255.254
                             DNS server 2: -
                             DNS server 3: -
                      Install-server IP: 10.100.1.3
                        root password: ************
             localhost is part of the cluster:              yes
     config dhcpd.conf:            Host entries from the old-dhcpd.conf file backuped.
                     Cluster nodes configuration:
     ----------------------------------------------------------------------------------

                          Disk type (IDE/SCSI) IDE
                      Network card (eth0/eth1) eth0

     ----------------------------------------------------------------------------------

      You can now start installing your cluster-nodes.
      If you wish to change this configuration click here


Commands: Use arrow keys to move, '?' for help, 'q' to quit, '<-' to go back.
  Arrow keys: Up and Down to move.  Right to follow a link; Left to go back.
 H)elp O)ptions P)rint G)o M)ain screen Q)uit /=search [delete]=history list
```

The Monitor

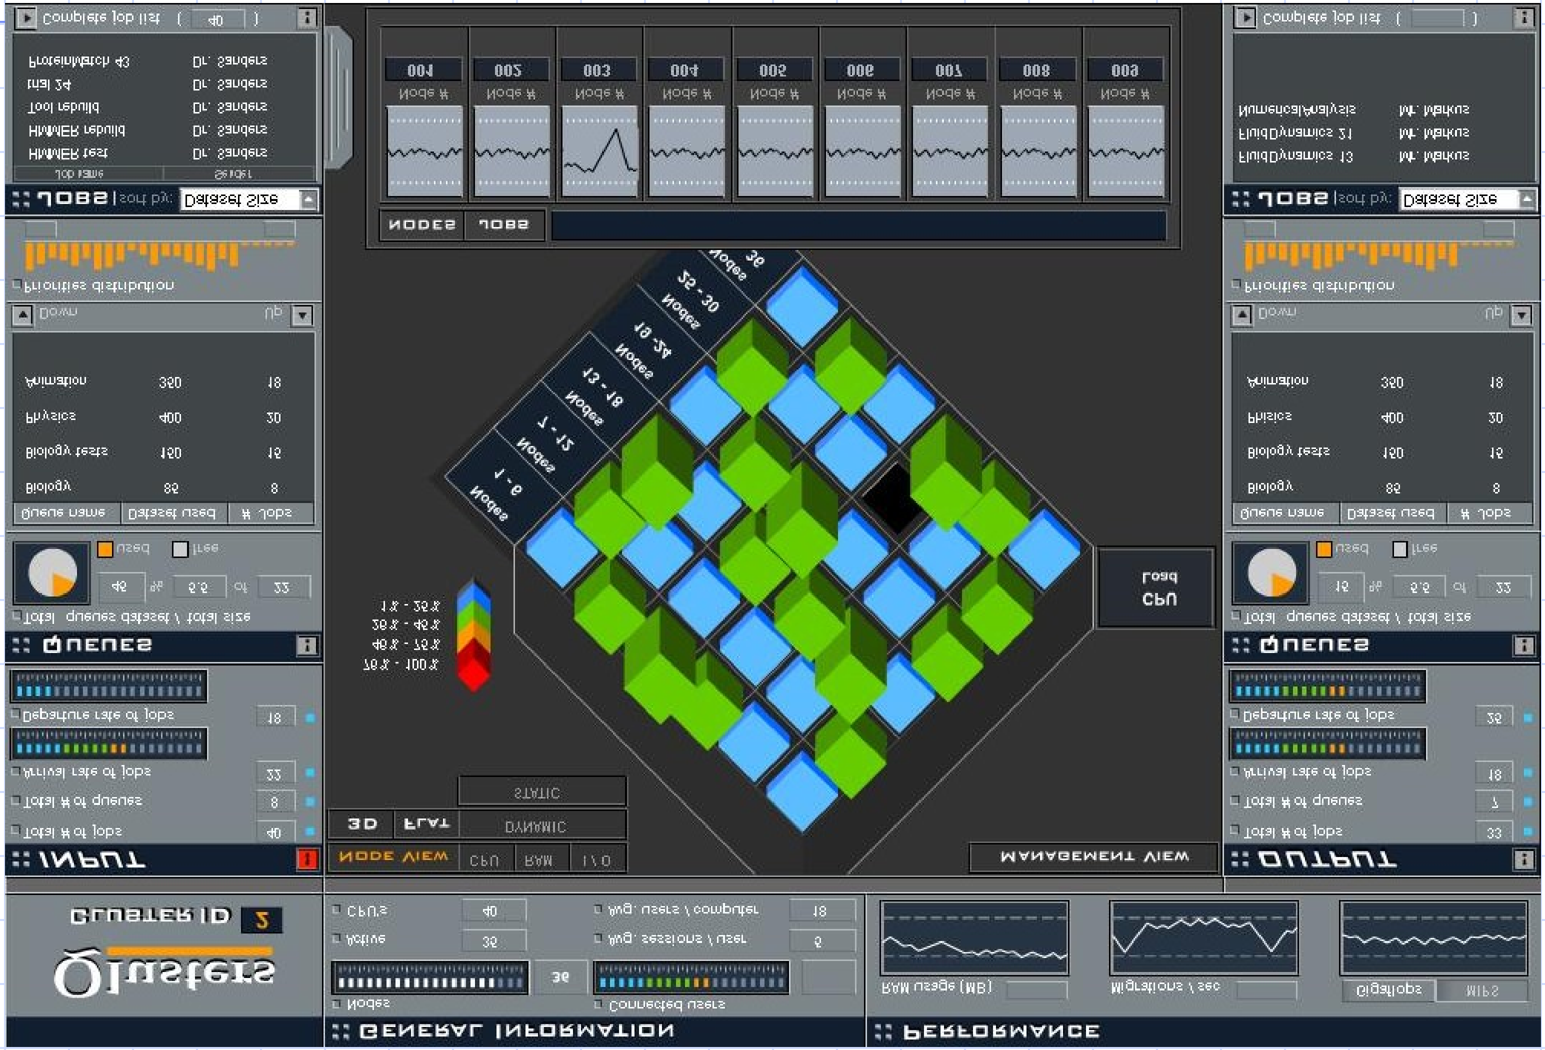
Complete job list (40)
ProteinMatch 43 Dr. Sanders
trial 24 Dr. Sanders
Tool rebuild Dr. Sanders
HMMER rebuild Dr. Sanders
HMMER test Dr. Sanders
Job name Sender
JOBS sort by: Dataset Size
Priorities distribution
Down Up
Animation 320 18
Physics 400 30
Biology tests 150 12
Biology 82 8
Queue name Dataset used # Jobs
used free
45 of 5.5 of 22
Total queues dataset / total size
QUEUES
Departure rate of jobs 18
Arrival rate of jobs 22
Total # of queues 8
Total # of jobs 40
INPUT
Node # 001 002 003 004 005 006 007 008 009
NODES JOBS
Nodes 1 - 6
Nodes 7 - 12
Nodes 13 - 18
Nodes 19 - 24
Nodes 25 - 30
Nodes 36
1% - 25%
25% - 49%
49% - 75%
75% - 100%
Load CPU
STATIC
DYNAMIC
3D FLAT
NODE VIEW CPU RAM I/O
MANAGEMENT VIEW
Complete job list ( )
NumericalAnalysis Mr. Markus
FluidDynamics 21 Mr. Markus
FluidDynamics 13 Mr. Markus
JOBS sort by: Dataset Size
Priorities distribution
Down Up
Animation 320 18
Phisics 400 30
Biology tests 150 12
Biology 82 8
Queue name Dataset used # Jobs
used free
15 of 5.5 of 22
Total queues dataset / total size
QUEUES
Departure rate of jobs 32
Arrival rate of jobs 18
Total # of queues 7
Total # of jobs 33
OUTPUT
CLUSTER ID 2
Qlusters
CPUs 40
Active 36
Avg. users / computer 18
Avg. sessions / user 6
Nodes 36
Connected users
GENERAL INFORMATION
RAM usage (MB)
Migrations / sec
Gigaflops MIPS
PERFORMANCE

The Monitor

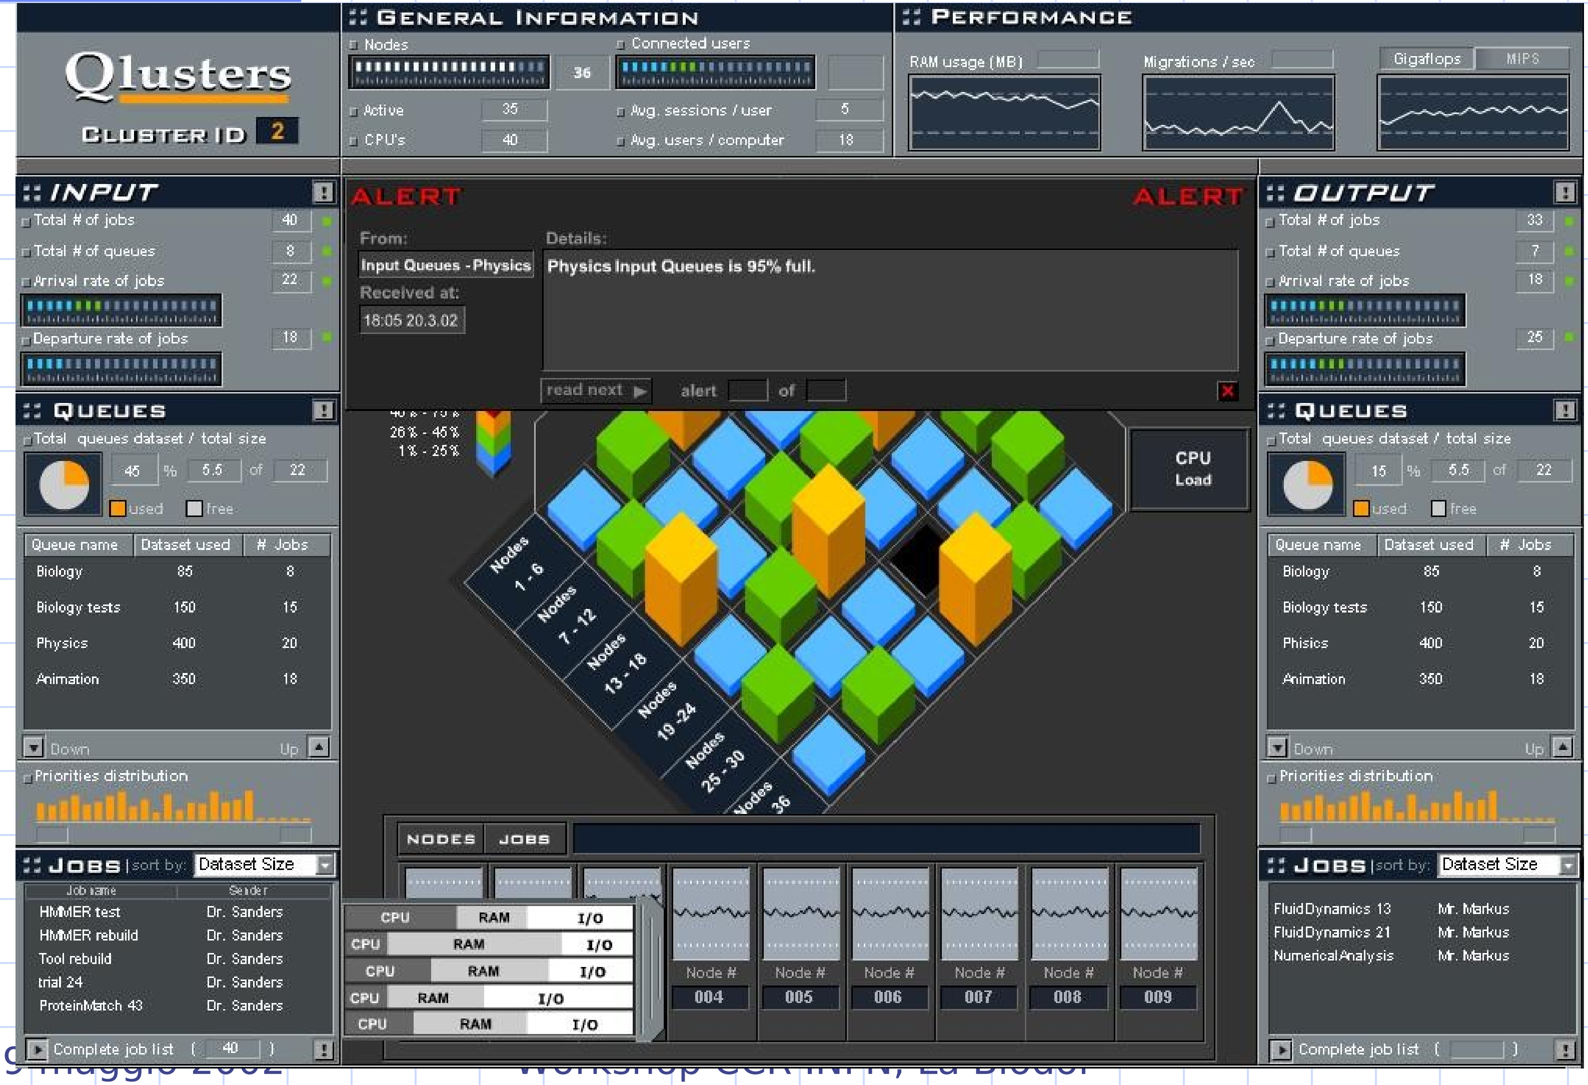

General Information

Nodes	36	Connected users	
Active	35	Avg. sessions / user	5
CPU's	40	Avg. users / computer	18

Input

Total # of jobs	40
Total # of queues	8
Arrival rate of jobs	22
Departure rate of jobs	18

Alert

From: Input Queues - Physics

Received at: 18:05 20.3.02

Details: Physics Input Queues is 95% full.

Output

Total # of jobs	33
Total # of queues	7
Arrival rate of jobs	18
Departure rate of jobs	25

Queues

Total queues dataset / total size: 45 % 5.5 of 22

Queue name	Dataset used	# Jobs
Biology	85	8
Biology tests	150	15
Physics	400	20
Animation	350	18

Queues

Total queues dataset / total size: 15 % 5.5 of 22

Queue name	Dataset used	# Jobs
Biology	85	8
Biology tests	150	15
Phisics	400	20
Animation	350	18

Jobs | sort by: Dataset Size

Job name	Sender
HMMER test	Dr. Sanders
HMMER rebuild	Dr. Sanders
Tool rebuild	Dr. Sanders
trial 24	Dr. Sanders
ProteinMatch 43	Dr. Sanders

Complete job list (40)

Jobs | sort by: Dataset Size

FluidDynamics 13	Mr. Markus
FluidDynamics 21	Mr. Markus
NumericalAnalysis	Mr. Markus

Complete job list ( )

The Monitor

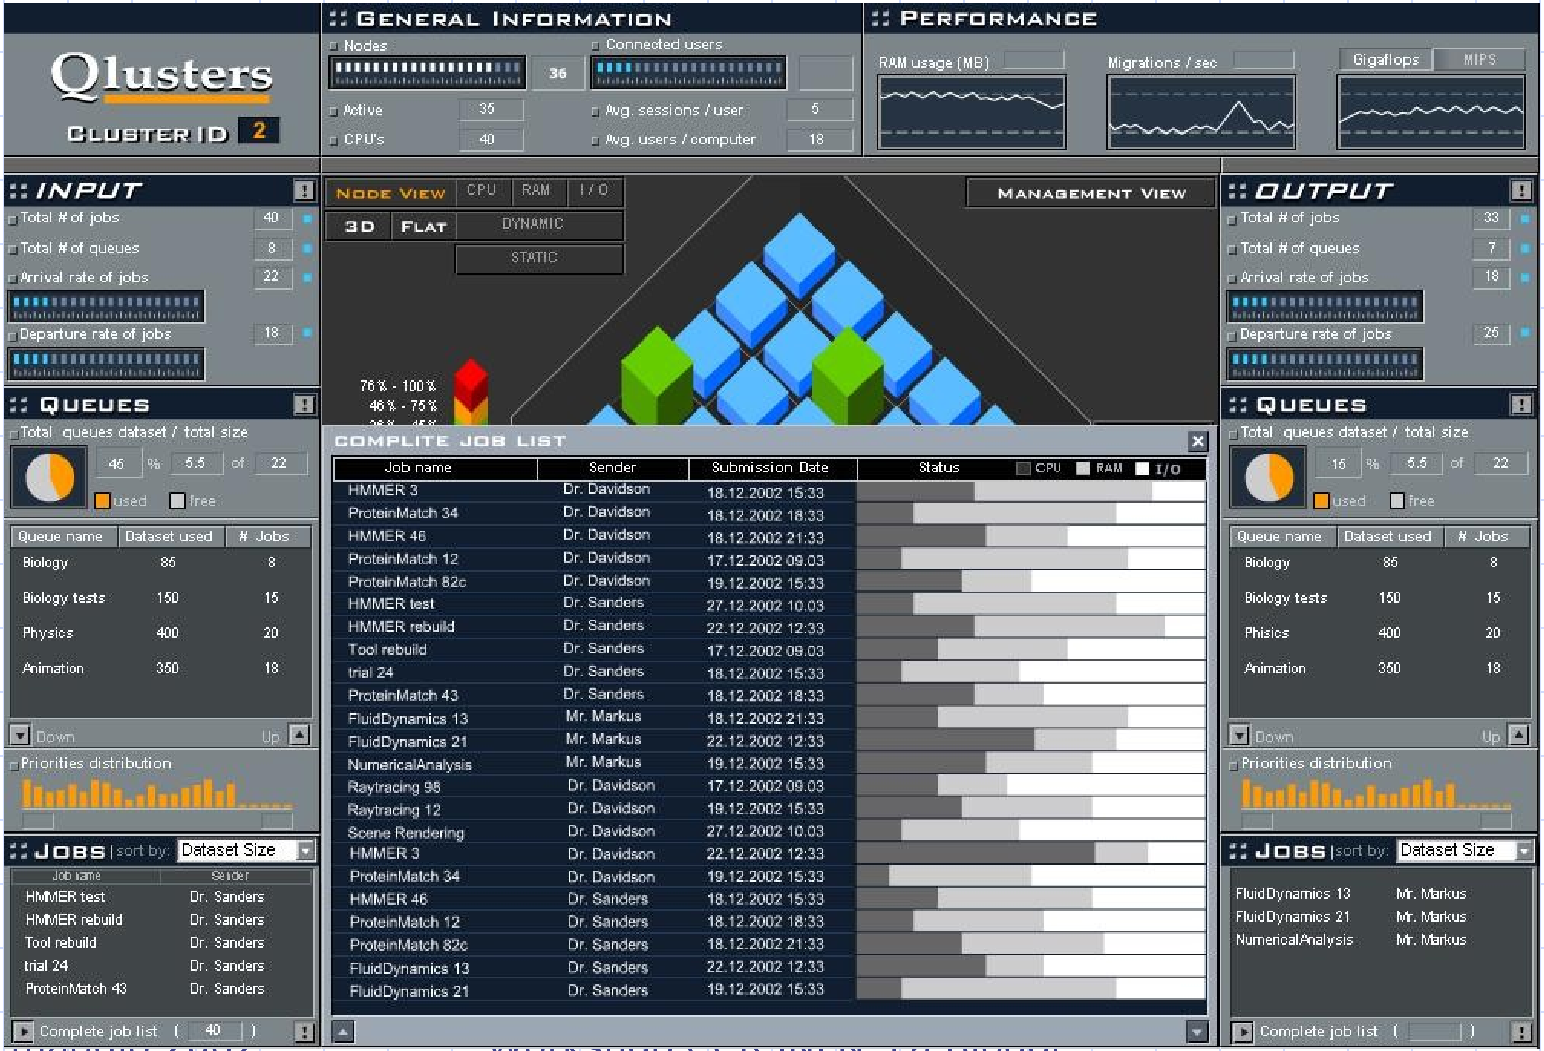

Nodes	36
Connected users	
Active	35
Avg. sessions / user	5
CPU's	40
Avg. users / computer	18

INPUT

Total # of jobs	40
Total # of queues	8
Arrival rate of jobs	22
Departure rate of jobs	18

QUEUES

Total queues dataset / total size: 45 % 5.5 of 22

Queue name	Dataset used	# Jobs
Biology	85	8
Biology tests	150	15
Physics	400	20
Animation	350	18

JOBS | sort by: Dataset Size

Job name	Sender
HMMER test	Dr. Sanders
HMMER rebuild	Dr. Sanders
Tool rebuild	Dr. Sanders
trial 24	Dr. Sanders
ProteinMatch 43	Dr. Sanders

Complete job list (40)

COMPLITE JOB LIST

Job name	Sender	Submission Date
HMMER 3	Dr. Davidson	18.12.2002 15:33
ProteinMatch 34	Dr. Davidson	18.12.2002 18:33
HMMER 46	Dr. Davidson	18.12.2002 21:33
ProteinMatch 12	Dr. Davidson	17.12.2002 09.03
ProteinMatch 82c	Dr. Davidson	19.12.2002 15:33
HMMER test	Dr. Sanders	27.12.2002 10.03
HMMER rebuild	Dr. Sanders	22.12.2002 12:33
Tool rebuild	Dr. Sanders	17.12.2002 09.03
trial 24	Dr. Sanders	18.12.2002 15:33
ProteinMatch 43	Dr. Sanders	18.12.2002 18:33
FluidDynamics 13	Mr. Markus	18.12.2002 21:33
FluidDynamics 21	Mr. Markus	22.12.2002 12:33
NumericalAnalysis	Mr. Markus	19.12.2002 15:33
Raytracing 98	Dr. Davidson	17.12.2002 09.03
Raytracing 12	Dr. Davidson	19.12.2002 15:33
Scene Rendering	Dr. Davidson	27.12.2002 10.03
HMMER 3	Dr. Davidson	22.12.2002 12:33
ProteinMatch 34	Dr. Davidson	19.12.2002 15:33
HMMER 46	Dr. Sanders	18.12.2002 15:33
ProteinMatch 12	Dr. Sanders	18.12.2002 18:33
ProteinMatch 82c	Dr. Sanders	18.12.2002 21:33
FluidDynamics 13	Dr. Sanders	22.12.2002 12:33
FluidDynamics 21	Dr. Sanders	19.12.2002 15:33

OUTPUT

Total # of jobs	33
Total # of queues	7
Arrival rate of jobs	18
Departure rate of jobs	25

QUEUES

Total queues dataset / total size: 15 % 5.5 of 22

Queue name	Dataset used	# Jobs
Biology	85	8
Biology tests	150	15
Phisics	400	20
Animation	350	18

JOBS | sort by: Dataset Size

FluidDynamics 13	Mr. Markus
FluidDynamics 21	Mr. Markus
NumericalAnalysis	Mr. Markus

Complete job list ( )

Interesting sites

- www.lcic.org
- www.beowulf-underground.org
- www.alinka.com (per Alinka Clustering Newsletter)

The End