


New modeling of parallel computing with clusters

by miKeL a.k.a.mc2



Dept. of Computer Architecture
Universitat Politècnica de Catalunya, FIB



New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdfticks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



I. Preliminaries

- Supercomputing
- AMD Linux clusters vs. mainframes
- A cluster classification
- Where clusters?

II. SSI. openMosix

- SSI. A new paradigm
- openMosix features
- Source code details. Add-ons
- openMosix internals
- Testing the system

III. openMosix. A real demo

- miKeL's cluster
- Setting up a openMosix system
- *El manual para el clustering con openMosix*

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdfticks and pdfscreen
packages in an AMD PC running SUSE LINUX.




Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



New modeling of parallel computing with Linux clusters

1. Preliminaries



Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



1.1. Supercomputing

< Why super computing?

- When a solution is computed reasonably it's already obsolete. The challenge is to enlarge the problem in order to model the reality more accurately.
- Who does not need more performance?

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



< What is? What want to reach?

The simultaneous use of more than one CPU to solve a problem.

< What isn't?

Note that parallel processing isn't from multitasking, in which a single CPU executes several programs at once.

< What implies?

Most computers have just one CPU, but some models have several. With single-CPU computers, it is possible to perform parallel processing by connecting the computers in a network. However, this type of parallel processing requires very sophisticated software.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



< Where is the problem?

In practice, it is often hard to divide a program in such a way that separate CPUs can execute various portions without interfering with each other.

< Two main kinds of supercomputers

– Vector computer: run it as fast as possible – *or our math-guru friend do all the problems as fast as he/she could.* –
Ad exemplum: mainframes, big servers.

– Parallel computer: run it as many and fast as possible – *or to divide the whole of task between various boring colleagues.* –
Ad exemplum: clusters, medium research computers.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_T_EX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>





- The CRAY C90 cost \$26 million.
- 8 vector machines in 1 up to 7.8Tflops.
- 1.5 times the Earth's population power.
- i686 8h to solve, C90 do it in 0.002sec.

Figure 1.1: CRAY C90 - Vector Machine

- Decentralized
- Scalable
- 1Gflop per node (aprox)
- 250\$ per node (aprox)



Figure 1.2: A HP PC Cluster - Parallel computer

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdfricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>





New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks


openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix



Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_T_EX with hyperref, pstricks, pdfticks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>

Figure 1.3: Earth Simulator - The world powerfull computer performing 35Tflops



1.2. AMD Linux clusters vs. mainframes

1.2.1. Why Linux?

- < RedmondOS is not counted among operating systems in HPC.
- < Proprietary Unix (Solaris, AIX, true64, HPUX,...) require proprietary hardware.
- < Linux development process is natural, evolutionary, open, transparent, quick.
- < Linux is free.
- < Linux has no planned obsolescence of hardware.
- < No serious HPC projects in other operating systems up to late 2003.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



1.2.2. Why clusters?

- < Mainframe attraction is fading away.
- < Entry price too high for smaller institutions and companies.
- < Clusters achieve good performance at 1/10 of price. -ECL vs. CMOS -
- < Clusters don't require specialist, just simple sysadmins.
- < Build with COTS components.

1.2.3. Why AMD processors?

- < Better performance at lower price within same architecture.
- < More transparency on hardware specifications.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with hyperref, pstricks, pdfticks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



1.3. A cluster classification

< High Availability (HA) - Failover

Failover clusters consists of almost two network connected computers with a separate heartbeat connection between two hosts. Heartbeat connection is being used to monitor whether all the services are still in use, as soon as a service breaks down the other machine tries to take over.

< High Performance (HP)

This architecture is being configured to give data centers that require extreme performance.

Ad exemplum: Beowulf, openMosix.

< Load balancing

When a request for say a web-server comes in, the cluster checks which machine is the lease busy and sends the request to it.

Ad exemplum: LVS.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



1.4. Where clusters?

- < Applications with unpredictable resource requirements and/or execution times.

Even the number of processes is determined at run time.

- < Clusters with dissimilar speed nodes.

- < Multi-user, time-sharing clusters.

The resource sharing algorithms guarantee that the workload is evenly distributed among all the nodes.

- < Applications that require high I/O rates.

By migrating an I/O intensive process to a file server.

- < Data intensive application.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2. SSI. openMosix



New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_T_EX with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.1. SSI. A new paradigm

< What is?

- Configuration with no master.
- Every node is his own boss.
- The system coordination is totally distributed.
- The cluster are seen as a **Single System Image**.

< What want to reach?

- Ease of management.
- Same scalability/overhead for 2 or 2000 nodes.
- Hide cluster complexity to users.
- Apps don't need change to take advantage - unlike pvm or mpi -.
- Always automatically load-balanced.
- Apps are installed on one node but running on all.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



- < What implies?
 - Adaptive load-balancing algorithm.
 - Dynamic process migration to load-balance.
 - Clustering file system.
 - Totally transparent to user, nodes and processes.

- < Some implementations
 - IBM SySPlex, DEC cluster
 - Kerrighed –only 2.2 –
 - openSSI –project failing –
 - CompaqSSI –vaporware –
 - MOSIX –contributions are not admitted –
 - openMosix

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.2. SSI today: openMosix

- < Patch (+3% kernel) + userland tools.
- < Linux turns parallel operating system.
- < Automatic migration of processes.
- < Statistics define to which node a process can migrate.
- < Not everything migrates, there are limitations.
- < *fork-and-forget* methods.
- < Multi-user and time-sharing environment for the execution of both sequential and parallel applications.
- < Only i386 arch. but accept SMP nodes.
DEC's Alpha and SUN's Sparc platforms are considered on future work.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.3. openMosix features

< Pro

- No extra libraries required.
- No changes to code required.
- No need to actually parallize your application.
- oMFS.
- *omdiscd*: the autodiscivery daemon.

< Cons

- DSM still in beta -web and data servers, matlab6, ...-.
- Issues with threads not gaining performance.
- Kernel dependent.

< Limitations

- *A body (pthread) cannot live without a mind (address space)* -to the Matrix movie-.
- Applications using shared memory (however a patch is available) and sockets.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdfticks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.3.1. Tasks that migrate

- < Matlab 5
- < Blast (Bio Informatics)
- < MPEG libraries
- < bladeenc
- < PovRay
- < MPI

2.3.2. Tasks that don't migrate

- < PVM
- < Apache
- < mySQL
- < pthreads
- < Oracle
- < Matlab 6
- < VMWare

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.3.3. Extra files on */proc*

- < */proc/[PID]/cantmove* why a process cannot be migrated
- < */proc/[PID]/goto* to which node the process should migrate
- < */proc/[PID]/lock* if a process is locked to its home node
- < */proc/[PID]/migrations* how many times the process migrated
- < */proc/[PID]/where* where the process is currently being computed
- < */proc/[PID]/migrate* same as *goto* remote processes

- < */proc/hpc/remote/from* the process home node
- < */proc/hpc/remote/identity* additional informations about the process
- < */proc/hpc/remote/statm* memory statistic of the process
- < */proc/hpc/remote/stats* cpu statistics of the process

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.4. openMosix source code details

- < 54 new files (54,000 lines)
- < 126 kernel files modified (18,000 lines changed/added)
- < 48 user-level files (36,000 lines)
- < Developer team:
 - Moshe Bar (boss)
 - Matt Rechenburg
 - Bruce Knox
 - Kris Buytaert
 - Louis Zechtzer
 - Martin Høy
 - Brian Pontz
 - Matthew Brichacek
 - Mark Veltzer
 - Muli Ben Yehuda (a.k.a. mulix)
 - David Santo Orcero (a.k.a. irbis)

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.5. The openMosix project milestones

- < Mosix is born early 80s on PDP-11/70.
- < BSD3.0 since 1997 and GNU/Linux since 1999.
- < 1999 first big showdown with openSource community due to binary only distribution.
- < Split Mosix/openMosix in 2001 due to license problems with University of Jerusalem.
- < july 2002 practically all Mosix installation moved to openMosix.
- < Many corporates moving from Crays to Linux clusters (BMW, Daimler-Chrysler, Fiat, Boeing).
- < AMD sponsors project with cash and hardware.
- < HP sponsors with U\$40,000 and U\$60,000 in hardware per year.
- < Donation of hardware Myrinet and Dolphin.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.6. Add-ons and remarks

- < GUI - openMosixview (Matt Rechenburg)
- < openMosix Stress-Test (Matt Rechenburg)
- < HowTo (Kris Buytaert)
- < Spanish doc (miKeL a.k.a.mc2)
- < Italian doc (MaX)
- < PlumpOS (Peter Willis)
- < openMosixWebView (Ramón Pons)
- < 3Dmosmon (Johnny Cache)
- < openMosix Applet (Gian Paolo Ghilardi)
- < Migshm (Univ. of Pune, India)
- < openMosixBlaster (Matthias Rechenburg)
- < Cluster Knoppix (Wim Vandersmissen)

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with
hyperref, pstricks, pdfticks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.7. openMosix internals

2.7.1. openMosix subsystems overview

I. Preemptive Process Migration (PPM)

II. Load-balancing

That process, when and where must to migrate.

Every node broadcast its stats to a random subset of nodes.

III. Memory Ushering

Migrations for to avoid *swapping* and *trashing*.

Load-balancing decisions independent.

IV. oMFS

SSI on filesystem. Always local harddisk access.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdfticks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.7.2. Preemptive Process Migration (PPM)

- < Complexity $O(n)$ –where n is the number of nodes–.
- < Adaptive (dynamic) resource distribution policy based on economic market theory load balancing algorithms.
- < Development of new competitive algorithms for adaptive resource management.
They can handle various kinds of resources (such as CPU, memory, IPC, and I/O).
- < Improvement on global cluster resources –no particular– based on optimum algorithm of finite window.
- < Configuration is symmetric and flexible.
- < *fork-and-forget* –full transparent migrations–.
- < Transparent to user, process and target node.
- < Who can migrate process? System admin, user proprietor, algorithms.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.7.3. A process migration trace

- < User launch process and openMosix evaluates the node load.
 - a. More performance on local then don't migrate
 - b. More performance on remote then migrate
- < Target node chosen from an updated load node list.
- < That must to be sent to target node?
 - code segment
 - data segment
 - stack segment
 - heap segment
- < No process interruption.
- < Migrated process has two communicated areas
 - deputy (kernel area, on local node)
 - remote (user area, on target node)

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.7.4. Kinds of process

< CPU intensive

- remote computes on target node
- return result on local node
- **no more problem**

< I/O intensive

- no memory mapping process allowed migrated
- remote access to disk two net access
Petition of page access and page itself.
- **solution:** oMFS+DFSA

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TEX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.7.5. oMFS

< Kernel capable and add to */etc/fstab*

```
mfs_mnt /omfs mfs dfsa=1 0 0
```

Each node filesystem will be available from each node on */omfs/X* where *X* is the node ID.

< openMFS is no replacement for NFS.

< DFSA

Modules which reduce the extra overhead of executing I/O oriented system-calls of a migrated process.

< Allow local writings/readings on remote nodes.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



2.8. Testing the system

For the impatient `main() {for(;;);}`

< openMosix Stress Test

- `distkeygen`: generation of 4000 RSA key pairs with 1024 bits of key length.
- `portfolio`: portfolios simulation for various stock composition for a given period of time.
- `eatmem`: calculates a `sin+sqrt` 10e6 times and outputs it outputs the loop count to a file.
- `forkit`: similar to the `eatmem` test but uses `fork` to create multiple process.
- `mfstes`: create and copy 10MB to all nodes back and forth.
- `kernel syscall`: a collection of tools for testing the Linux kernel and related features.
- `moving`: will move the `start_openMosix_test.sh` around each node while running the itself.

< LINPACK <http://www.netlib.org/benchmark/hpl/>

The package solves linear systems whose matrices are general, banded, symmetric indefinite, symmetric positive definite, triangular, and tridiagonal square. In addition computes the QR and singular value decompositions of rectangular matrices and applies them to least-squares problems. LINPACK uses column-oriented algorithms to increase efficiency by preserving locality of reference.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdfticks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3. openMosix. A real demo



New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.1. miKeL's cluster

- < **metrakilate: a diskless computer (DC)**
 - *motherboard*: ECS K7AMA (ALi M1535D+)
 - *processor*: AMD Duron 1200@900 (bus 200MHz)
 - *memory*: Hyunday 256MB DDR
 - *power supply*: LinkWorld 300W .8A
 - *nic*: Conceptronic c100i (RTL 8139D with RPL BootROM 8130)
 - *vga*: MilkyStar nVidia TNT2 M64 -Titan cooled-

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



- < mCii: the server
 - remote boot: rpld
 - internet address (IP): dhcpd
 - kernel: inetd (tftp)
 - file system: nfsserver (nfsroot)

- < mCii components:
 - motherboard: Soltek SL-75DRV
 - processor: AMD Athlon Thunderbirth 1333@1420Mhz (bus 266Mhz)
 - memory: TwinMOS 768MB DDR
 - power supply: SkyHawk 480W-4A
 - nic: Micronet 2500R (RTL 8139C)
 - vga: Leadtek nVidia GeForce Ti4200 (128MB DDR)
 - hd: WD + Seagate (40GB+40GB)
 - sound: VIA AC97 + MaxiStudio ISIS
 - other: ieee1394, avermediaTV, USRobotics56Kbps

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.1.1. mCii & metrakilate procedures

- < metrakilate: int19H to BIOS but I no have any installed chip on nic socket. There is a RPL server for my MAC?
- < mCii: I am. Get the rom from `/rplboot/<your_rom>`
- < metrakilate: OK. Who am I on net?
- < mCii: You are 192.168.x.y
- < metrakilate: OK. Where is the kernel?
- < mCii: You can boot your kernel from `/tftpboot/`
...
- < metrakilate: OK! Where can I mount / ?
- < mCii: Try NFS at `/tftpboot/metrakilate/`
- < metrakilate: Thanks mCii, now I'm a completely running DC.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.1.2. mCii: the RPL server

- < When you start a computer, the operating system is loaded into memory. When a **Remote Initial Program Load** (RPL) ROM chip is placed on a network card, the client computer will retrieve startup and configuration software from the server when the client starts.
- < This RPL ROM interrupts the computer's standard startup process and redirects the calls for the operating system to the NIC.
- < The RPL ROM also establishes a basic network connection between the workstation and server during the startup process. This process is known as the remoteboot process. The following steps will detail the process on how to setup and configure the server side of RPL.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



The rpl d. conf file that boot up metrakilate:

```
HOST
{
    ethernet = 00:50:FC:B1:B1:BD;    // mac address of the client

    FILE
    {
        path = "/rpl boot/eb-5.0.10-rt18139.lzrom";
        load = 0x1000;
    };

    execute = 0x1006;
    framesize = 1500;
    blocksize = 1440;
};
```

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.1.3. mCii: the DHCP server

- < **Dynamic Host Configuration Protocol** is a protocol for assigning dynamic IP addresses to devices on a network.
- < With dynamic addressing, a device can have a different IP address every time it connects to the network.
- < Dynamic addressing simplifies network administration because the software keeps track of IP addresses rather than requiring an administrator to manage the task.
- < DHCP was created by the Dynamic Host Configuration Working Group of the Internet Engineering Task Force (IETF).
- < DHCP is defined on RFC 2131, RFC 1534, RFC 2132.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



The dhcpd. conf file:

```
default-lease-time 600;           # 10 min
max-lease-time 7200;              # 2 hours
ddns-update-style ad-hoc;

subnet 192.168.1.0 netmask 255.255.255.0 {
    option subnet-mask 255.255.255.0;
    option domain-name openMo6;
    range dynamic-bootp 192.168.1.1 192.168.1.10;
}

group{
    #filename <kernel image for the DC group>

    host metrakilate {
        hardware ethernet 00:50:FC:B1:B1:BD;
        fixed-address 192.168.1.2;
        filename "/tftpboot/vmlinuz-metrakilate";
        option root-path "192.168.1.1:/tftpboot/metrakilate";
    }
}
```

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_T_EX with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.1.4. mCii: the TFTP & NFS server

- < **Trivial File Transfer Protocol** is a simple form of the File Transfer Protocol (FTP).
- < TFTP uses the User Datagram Protocol (UDP) and provides no security features.
- < It is often used by servers to boot diskless workstations, X-terminals, and routers for to boot kernel.
- < To initialize the service it's necessary to put the next line on */etc/inetd.conf*

```
tftp dgram udp wait root /usr/sbin/in.tftpd in.tftpd /tftpboot
```

- < To export nfsroot filesystem it's necessary to create a */etc/exports* file with

```
/tftpboot/metrakilate/ 192.168.1.0/255.255.255.0(rw,no_root_squash,sync)
```

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TEX with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.2. Setting up a openMosix system

- I. Patching & Compiling the new kernel
- II. Configuring & Starting
- III. Running applications

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.2.1. Patching & Compiling the new kernel

Unpack the kernel:

```
tar xzvf linux-2.4.20.tar.gz
```

Apply the openMosix patch

```
cd linux-2.4.20
```

```
zcat ../openMosix-2.4.20-2.gz -- patch -p1
```

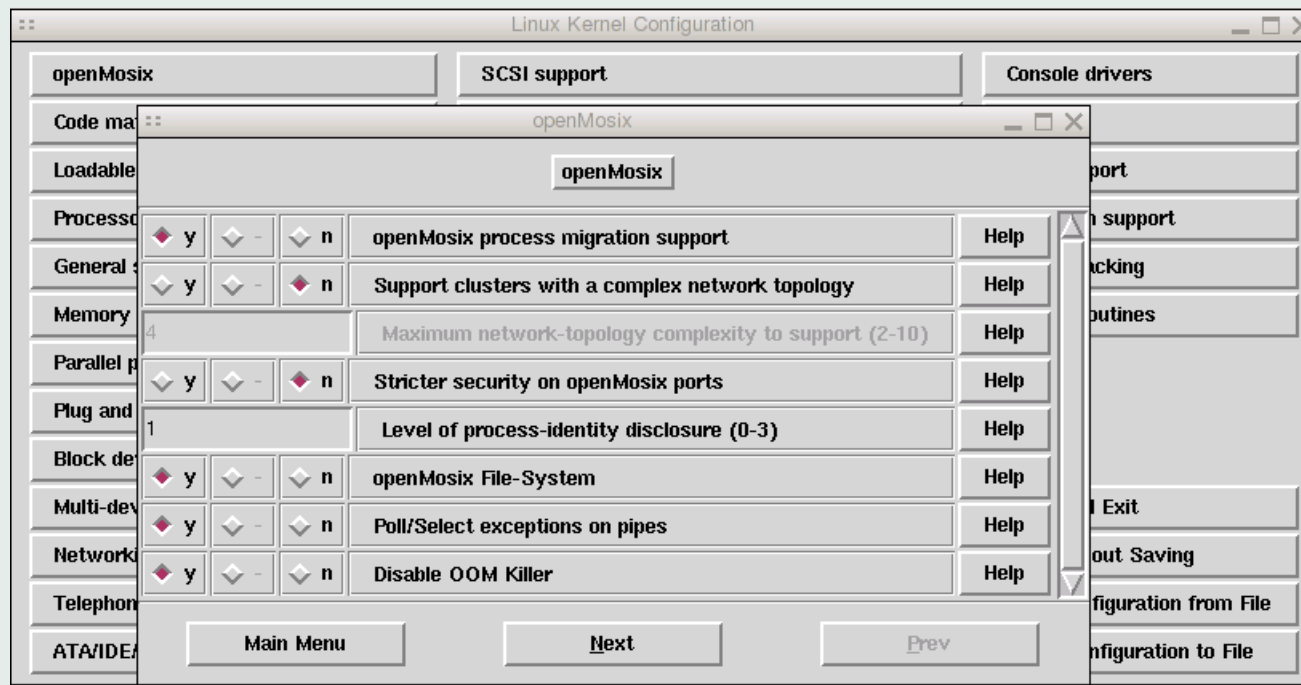


Figure 3.1: A openMosix patched kernel 2.4.20

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdfticks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



Compiling the kernel:

```
make dep && make clean && make bzImage
```

Installing modules

```
make modules && make modules_install
```

Copying kernel image and their map

```
cp arch/i386/boot/bzImage /boot/vmlinuz-openMo6
```

```
cp System.map /boot/System.map-openMo6
```

The modules are on `/lib/modules/2.4.x-openmosix-y`.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2
This document is generated by PDF_TE_X with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.2.2. Configuring & starting a openMosix environment

Two ways for topology configuration:

- < Using manually created */etc/openmosix.map*.

```
1 192.168.1.1 2
```

And initializing this file with

```
setpe -w -f /etc/openmosix.map
```

or

```
/etc/init.d/openmosix start
```

- < Using the autodiscovery

```
omdiscd -i eth0
```

This daemon create a file text with the complete topology that it has found. showmap display result.

- < Installing userland-tools (from RPM)

```
rpm -i openmosix-tools-0.3-1.i386.rpm
```

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.2.3. Running applications

- < Nothing fancy, just start the application, it will migrate to the least busy node.
- < Userland-tools allow a good migration control:

```
migrate <PID> <node ID>
```

```
migrate <PID> home
```

```
mosrun
```

```
mosmon
```

```
mps
```

```
mosctl
```

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdfticks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



SCREENSHOTS

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with hyperref, pstricks, pdftricks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>



3.3. *El manual para el clustering con openMosix*

<http://como.akamc2.net>

- < Wrote in spanish and included to the TLDP-ES/LuCAS.
- < Linked on many related pages.
- < 400 pages on version 1.0BETA1.
- < A open project with a lot of good collaborations.
- < It will be updated constanly.
- < june2003 Nowadays it's near to 3000 downloads.
- < miKeL's first contribution to open source community.
- < 25may2003 MaX has begun the italian translation.
- < june2003 Kris Buytaert offers work together.

New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_TE_X with
hyperref, pstricks, pdftricks and pdfscreen
packages in an AMD PC running SUSE LINUX.




Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>





Thanks!
Questions? :)



New modeling of parallel computing with Linux clusters

Supercomputing

AMD Linux clusters vs. mainframes

A cluster classification

Where clusters?

SSI. A new paradigm

SSI today: openMosix

openMosix features

openMosix source code details

The openMosix project milestones

Add-ons and remarks

openMosix internals

Testing the system

miKeL's cluster

Setting up a openMosix system

El manual para el clustering con openMosix

Miquel Catalán i Coit
COMPUTER SCIENCE ENGINEERING

©copyLeft 2003, miKeL a.k.a.mc2

This document is generated by PDF_{TEX} with hyperref, pstricks, pdfticks and pdfscreen packages in an AMD PC running SUSE LINUX.



Facultat d'Informàtica de Bcn
c/ Jordi Girona, 1-3. Campus Nord, B6
08034 Barcelona - CATALUNYA

<http://om.akamc2.net>

