


2nd Slide Set Cloud Computing

Prof. Dr. Christian Baun

Frankfurt University of Applied Sciences
(1971–2014: Fachhochschule Frankfurt am Main)
Faculty of Computer Science and Engineering
christianbaun@fb2.fra-uas.de

VAXcluster



MR-5062-AD

CI = ComputerInterconnect
HSC = Hierarchical Storage Controller
Star = Star Coupler

Image sources:

- <http://hampage.hu/oldiron/vaxen/eikvms1.jpg>
- [http://odl.sysworks.biz/disk\\$vxdocmay941/decw\\$book/d3ywaa51.p37.decw\\$book](http://odl.sysworks.biz/disk$vxdocmay941/decw$book/d3ywaa51.p37.decw$book)
- <http://www.computerhistory.org/collections/catalog/102635385>

History of Cluster computing

- Early 1990s:
 - Client/Server architecture became mainstream
 - UNIX workstations were used in companies and universities
 - Workstations are only utilized 5-10% during their everyday work
 - The biggest part of the potential computing power is wasted
 - Wanted: A middleware, which makes the free computing power available for solving complex tasks
 - PVM and MPI both, implement a **Message Passing Layer**, which is independent of the nodes' architecture and allows to develop cluster applications
- Mid-1990s:
 - Growing popularity of Clusters with commodity hardware nodes
 - Compared to mainframes, Clusters provide for a low price, depending on their hardware and structure, high computing performance and/or availability
 - Flexibility in terms of the application and extensibility

Definition of Cluster Computing

Cluster computing

Clustering is parallel computing on systems with distributed memory

- A cluster consists of at least 2 nodes
 - Each node is an independent computer system
 - The nodes are connected via a computer network
 - In clusters with just a few nodes, inexpensive computer network technologies (Fast or Giga-Ethernet) are used
 - Clusters with several hundred nodes require high-speed computer networks (e.g. InfiniBand)
 - Often, the nodes are under the control of a master and are attached to a shared storage
 - Nodes can be ordinary PCs, containing commodity hardware, workstations, servers or supercomputers

From the user perspective (in a perfect world)...

- the cluster works like a single system \implies a virtual uniprocessor system
- Ideally, the users don't know, that they work with a cluster system

Cluster of Workstations / „Feierabendcluster“

- If the nodes are only available at specific times, the cluster is called Clusters of workstations (COWs) or Network of Workstations (NOWs)
- During normal working times, the employees use the nodes of such a cluster system as workstations
 - The concept was popular around the year 2000
 - Today, this concept not popular any more

Oliver Diedrich, *NT-Workstations als Feierabend-Cluster*, c't 22/2000, P.246

- More than 200 computers with Pentium II/III CPUs and with at least 64 MB RAM in the plastics laboratory of BASF in Ludwigshafen
- All computers are centrally administered and run Windows NT 4.0
- A WinSock server runs on every computer as a service all the time
- If the WinSock server receives a request from the central host, it confirms the request
- Next, the central host transmits a file (size: 10-100 KB) via TCP/IP to the server, which includes data for processing
- If the transmission was successful, the WinSock server processes the data and transmits the results back to the central host

Distinguishing Criteria of Clusters


- **Structure**
 - Homogeneous structure
 - Heterogeneous structure
- **Installation concept**
 - Glass-house
 - Campus-wide
- **Fields of application**
 - High Performance Clustering
 - High Availability Clustering
 - High Throughput Clustering
- **Behavior in the event of failed nodes**
 - Active/Passive-Cluster
 - Active/Active-Cluster

Homogeneous and Heterogeneous Clusters

- The structure of clusters can be homogeneous and heterogeneous

Heterogeneous structure

Homogeneous structure



I have never seen a heterogeneous cluster with different operating systems in practice...

- In practice, the construction of a heterogeneous cluster is generally a bad idea
- The administration of homogeneous clusters is challenging, but the administration of heterogeneous clusters is hell (especially when commodity hardware is used)


Installation Concepts of Clusters (1/2)



- **Glass-house**
 - The cluster is located in a single room or server rack
- **Advantages:**
 - Fast access to all components for maintenance and troubleshooting
 - Nodes can be connected via high-performance networks
 - Increased protection against sabotage

- **Drawbacks:**
 - In case of a power failure or fire in the building, the operation of the entire cluster is at risk

Installation Concepts of Clusters (2/2)



- **Campus-wide**
 - The nodes are located in multiple buildings and spread across the site of the research center or company
- **Advantages:**
 - It is hard to destroy the cluster completely
- **Drawbacks:**
 - It is impossible to use high-performance computer networks
 - Often, the nodes contain different hardware components

Fields of Application of Clusters

- Clusters for different applications exist
- ① **High Availability Clustering**
 - Objective: high availability
- ② **High Performance Clustering**
 - Objective: high computing power
- ③ **High Throughput Clustering**
 - Objective: high throughput

High Availability Clustering (1/2)

- For High Availability Clustering, highly available systems are required
 - This is achieved by:
 - Redundancy of nodes and their components
 - Avoiding a single point of failure

Redundancy

Redundancy comes from the Latin word *redundare* and means that more than enough of something is available

- Redundancy means: A system contains components, which are not required for the functioning of the system, but they can take over the work of identical components in case of error
- By using redundant nodes, it is possible to emulate the technology and benefits of mainframes for a low price and a high level of flexibility is achieved
- **The aim of HA clustering is to archive a high availability**

High Availability Clustering (2/2)

- For High Availability Clustering, not the availability of the nodes has top priority, but the availability of the offered services
- 2 groups of High Availability Clusters exist, which differ in their behaviour in the event of failures of nodes:
 - Active/Passive-Cluster
 - Active/Active-Cluster

HA-Clustering – Active/Passive and Active/Active

- **Active/Passive-Cluster** (also called: **Hot-Standby-Clusters**)

- During normal operation, at least a single node is in **passive** state
- Nodes in passive state do not provide services during normal operation
- If a node fails, a passive node takes over its services
- **Failover** = a node takes over the services of a failed node
- Benefit: The services must not be designed for cluster operation
- Drawback: Much potentially available performance remains unused in normal operation

- **Active/Active-Cluster**

- All nodes run the same services
- All nodes are in **active** state
- If nodes fail, the remaining active nodes need to take over their tasks
- Advantage: Better distribution of load between nodes
- Drawback: Services need to be designed for cluster operation, because all nodes access shared resources (data!) simultaneously

High Availability Clustering – Failover and Failback

- **Failover:** Ability to automatically transfer the tasks of a failed node to another node for minimizing the downtime
 - The failover functionality is usually provided by the operating system
 - Example: Heartbeat for Linux

<http://www.linux-ha.org/wiki/Heartbeat>

Heartbeat is a daemon that provides cluster infrastructure (communication and membership) services to its clients. This allows clients to know about the presence (or disappearance!) of peer processes on other machines and to easily exchange messages with them. In order to be useful to users, the Heartbeat daemon needs to be combined with a cluster resource manager (CRM) which has the task of starting and stopping the services (IP addresses, web servers, etc.) that cluster will make highly available. Pacemaker is the preferred cluster resource manager for clusters based on Heartbeat.

- **Failback:** If failed nodes are operational again, they report their status to the load balancer and get new jobs assigned in the future
 - From that point in time, the cluster again has the same performance capability, it had before the failure of the nodes

Well written articles about Heartbeat and DRBD

Andreas Sebald. *Linux-HA-Cluster mit Heartbeat und DRBD*. Linux-Magazin 07/2004.
<http://www.linux-magazin.de/Ausgaben/2004/07/Reservespieler>


Split Brain

- Connection failure between nodes
 - The nodes still work without any trouble, only the connection between them is disrupted
 - Tools like Heartbeat, which monitor the presence (or disappearance) of nodes, assume that nodes are broken
 - Each node declares itself to be the primary node
 - In Active/Passive-Clusters \implies failure of the cluster (offered services)
- **If shared storage is used**, each node tries to write on the storage
 - This can be avoided by using additional hardware and distinguishing the MAC addresses
 - One possible solution to avoid further issues: If simultaneous access attempts are detected from different MAC addresses, are nodes are automatically shut down
- **If distributed storage is used**, write requests cause inconsistent data on the nodes
 - It is difficult to fix the broken consistency without losing data

More information about the different storage architectures of High Availability Clusters present the next slides

Architectures of High Availability Clustering


- 2 architectures of High Availability Clustering exist:
 - **Shared Nothing** ⇒ Distributed Storage
 - **Shared Disk** ⇒ Shared Storage



Shared Nothing Architecture

Image Source: technet.microsoft.com

- In a **Shared Nothing** cluster, each node has its own storage resource
- Even, when a resource is physically connected to multiple nodes, only a single node is allowed to access it
 - Only if a node fails, the resource is acquired by another node
- Advantage: No lock management is required
 - No protocol overhead reduces the performance
 - In theory, the cluster can scale almost in a linear way
- Drawback: Higher financial effort for storage resources, because the data can not be distributed in an optimal way



Shared Nothing with DRBD (1/3)

- Distributed Replicated Block Device (DRBD)
 - Free software to build up a network storage for Shared Nothing clusters, without an expensive Storage Area Network (SAN)
- Shared storage is always a single point of failure, since only the cluster nodes are designed in a redundant way
 - Redundant SAN solutions are expensive (> 100.000 €)




Image Source: M. Jones, <https://www.ibm.com/developerworks/library/1-drbd/index.html>

Well written articles about DRBD

iX 3/2010. Florian Haas. *Hochverfügbare Shared Nothing Cluster mit DRBD*. P.120-123
M. Tim Jones. High availability with the Distributed Replicated Block Device. 2010.
<https://www.ibm.com/developerworks/library/1-drbd/>

Shared Nothing with DRBD (2/3)

- Functioning:
 - A primary server and a secondary server exist
 - Write requests are carried out by the primary server and afterwards are sent to the secondary server
 - Only if the secondary server reports the successful write operation to the primary server, the primary server reports the end of the successful write operation
 - Practically, it implements RAID 1 via TCP
 - Primary server fails \implies secondary server becomes primary server
 - If a failed system is operational again, only the data blocks, which have changed during the outage are resynchronized
 - Read access is always carried out locally (\implies better performance)




Image Source: M. Jones, <https://www.ibm.com/developerworks/library/1-drbd/index.html>

Shared Nothing with DRBD (3/3)

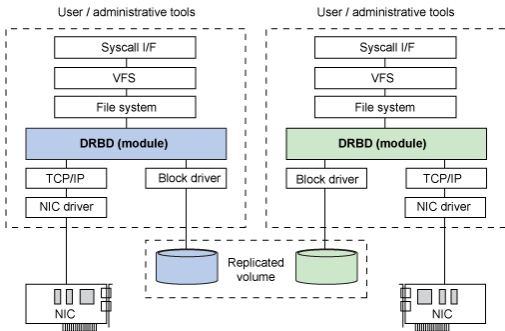


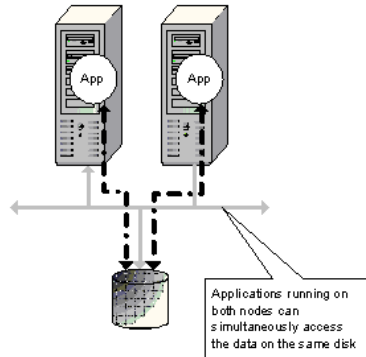
Image Source: M. Jones, <https://www.ibm.com/developerworks/library/l-drbd/index.html>

- DRBD is a part of the Linux kernel since version 2.6.33 (February 2010)
- Because DRBD operates inside the Linux kernel at block level, the system is transparent for the layers above it
- DRBD can be used as a basis for:
 - Conventional file systems, such as ext3/4 or ReiserFS
 - Shared-storage file systems, such as Oracle Cluster File System (OCFS2) and Global File System (GFS2)
 - If shared-storage file systems are used, all nodes must have direct I/O access to the device
 - Another logical block device, such as the Logical Volume Manager (LVM)

Shared Disk Architecture

Image Source: technet.microsoft.com

- In a **Shared Disk** cluster, all nodes have access to a shared storage
- Several possible ways exist to connect the nodes to the storage:
- **SAN** (Storage Area Network) via Fibre Channel
 - Expensive, but provides high performance
 - Provides block-level access to storage devices via the network.
- **NAS** (Network Attached Storage)
 - Easy-to-use file server
 - Provides file system-level access to storage devices via the network
 - Can also be implemented as a pure software solution
 - Examples: FreeNAS and Openfiler
- **iSCSI** (Internet Small Computer System Interface)
 - SCSI protocol via TCP/IP
 - SAN-like access via the IP-network



High Performance Clustering (1/2)

- Objective: High computing power
 - Also called: **Clustering for Scalability**
- High Performance Clusters provide the performance of mainframe computers for a much lower price
- These clusters are usually made of commodity PCs or workstations
- Typical application area:
 - Applications, which implement the Divide and Conquer principle
 - Such applications split big tasks into multiple sub-tasks, evaluates them and puts together the sub-task results to the final result
 - Applications, used for analyzing large amounts of data


High Performance Clustering (2/2)

Image Source: Reddit

Application examples: Crash test simulation, weather forecast, optimization of components, Monte Carlo simulation, flight path calculation, data mining, flow simulation, strength analysis, rendering of movies or clips, simulation of the night sky, variant calculating for chess, prime number computation, . . .

In 1995 Pixar rendered Toy Story on a 294 × 100MHz CPU Sun SPARCstation 20 cluster

Each SPARCstation 20 (single processor) had SunOS 5.4 installed and a HyperSPARC 100 MHz with 27.5066 MFLOPS
⇒ The theoretical maximum performance of the setup was 294 * 27.5066 = 8086.94 MFLOPS



• Advantages:

- Low price and vendor independence
- Defective components can be obtained in a quick and inexpensive way
- It is easy to increase the performance in a short time via additional nodes

• Drawback:

- High administrative and maintenance costs, compared with mainframes

High Performance Clustering – Beowulf Cluster

- If a free operating system is used \implies **Beowulf** cluster
- If a Windows operating system is used \implies **Wulfpack**
- A Beowulf cluster is never a cluster of workstations (COW)
 - Beowulf clusters consist of commodity PCs or workstations, but the nodes of a Beowulf cluster are used only for the cluster
- The cluster is controlled via a master node
 - The master distributes (schedules) jobs and monitors the worker nodes
- Worker nodes are only accessible via the network connection
 - They are not equipped with I/O devices like screens or keyboards
- Worker nodes contain commodity PC components and are not redundant (\implies designed for high availability)
 - A potential issue is the failure of the cooling of the system components
 - Fans in nodes and power supplies have a limited lifetime and fail without any warning
 - Modern CPUs cannot operate without adequate cooling

Stone SouperComputer (1/2)



Image source: <http://www.climate modeling.org/~forrest/linux-magazine-1999/>

- Example for a Beowulf cluster, made of discarded office computers
- <http://www.extremelinux.info/stonesoup/>

Stone SouperComputer (2/2)



- Built in 1997
- Mostly 486DX-2/66 Intel CPUs
- Some Pentiums
- 10 Mbit/s Ethernet
- RedHat Linux, MPI and PVM
- Extremely heterogeneous structure
- No purchase costs
- High setup and administration effort
- Everything handmade

Image source: <http://www.climate modeling.org/~forrest/linux-magazine-1999/>

Later Generations of Beowulf Clusters (1/2)



Image source: <http://archiv.tu-chemnitz.de/pub/2000/0089/data/clic.html>

- Vendors, such as Megware in Chemnitz, sell complete Beowulf clusters
- Image: Chemnitzer Linux Cluster (CLIC) from 2000

Later Generations of Beowulf Clusters (2/2)



Image source: <http://tina.nat.uni-magdeburg.de>

- Tina (Tina is no acronym) in Magdeburg from 2001


State of the Art of Cluster Computing



Image source (right image):
<http://physics.bu.edu/~sandvik/clusters.html>

A Cluster of Motherboards


The cluster in the right image has 48 nodes



High Throughput Clustering

- Objective: Maximize throughput
- Such clusters consist of servers, which are used to process incoming requests
- Such clusters are not used for extensive calculations
 - Tasks must not be split into sub-tasks
 - The individual tasks (requests) are small and a single PC could handle them
- Typical fields of application of High Throughput Clustering:
 - Web servers
 - Internet search engines
- Large compute jobs \implies High Performance Cluster
- Multiple small compute jobs (in a short time) \implies High Throughput Cluster

Today: Clusters at Universities



<http://cs.boisestate.edu/~amit/research/beowulf/>

<http://physics.bu.edu/~sandvik/clusters.html>

- Beowulf clusters, built up from commodity hardware
 - ⇒ low acquisition cost
- High effort for administration (handmade)
 - ⇒ irrelevant, because students do the administration

Today: Research and Industry (Example: HP C7000)




Image source: <http://imagehost.vendio.com/bin/imageserver.x/00000000/pdneiman/DSC04040.JPG>

- Compact blade servers or so-called *pizza boxes*
- Professional management tools (like HPE iLO) and redundant components simplify the administration

Calculation Example about the Possible Packing Density

- A 19 inch rack contains up to 4 blade enclosures (BladeCenters)
- A HP C7000 BladeCenter provides 16 blade slots
- Blades exist, which contain 2 independent servers
 - e.g. HP Blade ProLiantBL2x220c G5
 - 2 servers per blade. Completely independent computers
 - Each server contains: 2x Intel Quad Core Xeon (2,33 GHz) and 16 GB RAM

⇒ 8 cores per server

⇒ 16 cores per blade

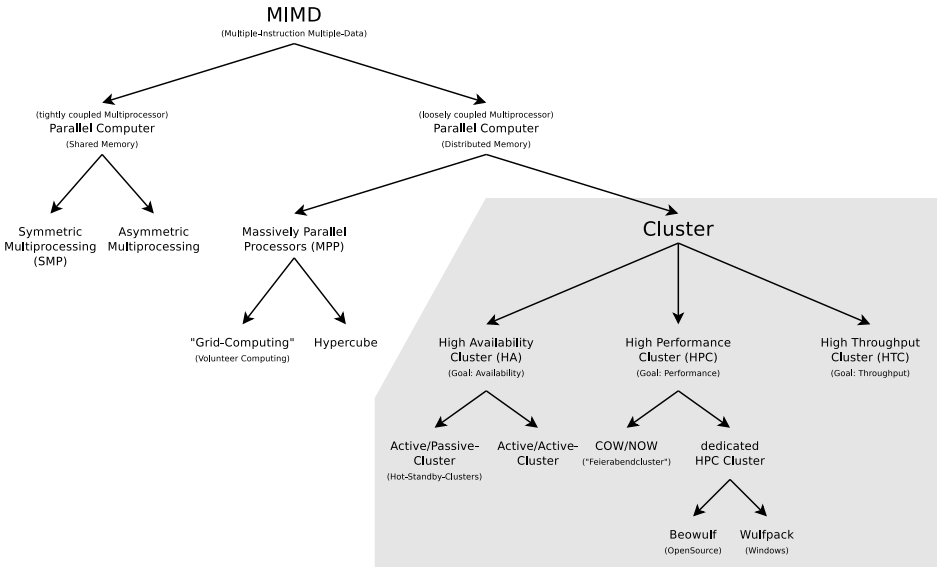
⇒ 256 cores per blade enclosure (BladeCenter)

⇒ 1024 cores per 19 inch rack

The packing density increases

Intel Xeon processors with 6 cores (*Dunnington*), with 8 cores (*Nehalem-EX*), with 18 cores (*Haswell-EX*) and with 22 cores (*Broadwell*) are already available. AMD offers the Opteron (*Magny-Cours*) with 12 cores and the Ryzen (*Threadripper*) with 32 cores

Classification of Clusters



Advantages and Drawbacks of Clusters

Advantages:

- Flexibility and extensibility
 - The number of nodes of a cluster can be dynamically increased or decreased according to the needed resources
- Lower purchase price compared with supercomputers
- Simple replacement of commodity hardware components

Drawbacks:

- Errors occur more often compared with a single supercomputer
- Clusters consist of many independent systems
⇒ higher administrative costs and personnel expenses compared with a single or few supercomputers
- High effort for distributing and controlling applications
 - If the number of nodes is increased, the effort increases too

Libraries for Cluster Applications

- For the development of cluster applications, several libraries exist, which provide **message passing** functionality, independent of the cluster structure and the specific hardware
- Popular message passing systems:
 - **Parallel Virtual Machine (PVM)**
 - **Message Passing Interface (MPI)**

Libraries for Cluster Applications (PVM)

- **Parallel Virtual Machine (PVM)**
- Developed from 1989 onwards
- Provides a uniform programming interface for the creation of a parallel computing unit with distributed memory
- Supports C/C++ and Fortran 77/90
- Consists of a daemon, libraries and tools
 - PVM is not a programming language!
- Especially suited for heterogeneous environments
 - Converts e.g. automatically between little endian and big endian
- Focus is not performance, but portability

Good PVM book (freely available in PostScript and HTML)

Janusz Kowalik (1994). **PVM: Parallel Virtual Machine. A Users' Guide and Tutorial for Networked Parallel Computing.**
<http://www.netlib.org/pvm3/book/pvm-book.html>

Libraries for Cluster Applications (MPI)

- **Message Passing Interface (MPI)**
- Development started in 1993-94
- Collection of functions (e.g. for process communication) to simplify the development of applications for parallel computers
- The library can be used with C/C++, Fortran 77/90 and Python
 - MPI is not a programming language!
- Contains no daemon
- Implements message-based communication (message passing)
- Especially suited for homogeneous environments
- Focus: Performance and security
- MPI implements > 100 functions and several constants
- Implementations: LAM/MPI (obsolete), OpenMPI, MPICH2,...

MPI tutorial of Wes Kendall, Dwaraka Nath, Wesley Bland: <https://mpitutorial.com/tutorials/>
MPI tutorial of Stefan Schaefer and Holger Blaar: <http://www2.informatik.uni-halle.de/lehre/mpi-tutorial/index.htm>

MPI Functions – Selection of important Functions (1/5)

- `MPI_Init(&argc, &argv)`
 - Initialization routine \implies starts the MPI environment
 - Defines the communicator `MPI_COMM_WORLD`
 - A communicator contains a group of processes and a communication context
 - `MPI_COMM_WORLD` contains all processes
 - The arguments `argc` and `argv` are pointers to the parameters of the main function `main`
 - The `main` function always receives 2 parameters from the operating system
 - `argc` (**argument count**) contains the number of parameters passed
 - `argv[]` (**argument values**) contains the parameters itself
 - The names of the variables can be freely selected, but they are usually named `argc` and `argv`
 - Not command-line parameters passed \implies `argc = 1`

<https://web.archive.org/web/20070909174518/http://www2.informatik.uni-jena.de/cmc/racluster/mpi-leitfaden>

MPI Functions – Selection of important Functions (2/5)

- `MPI_Comm_Size(MPI_Comm comm, int size)`
 - Determines the number of processes in a communicator
 - `size` is the output

```
1 #include "mpi.h"
2
3     int         size;
4     MPI_Comm    comm;
5     ...
6     MPI_Comm_size(comm, &size);
7     ...
```

MPI Functions – Selection of important Functions (3/5)

- `MPI_Comm_Rank(MPI_Comm comm, int rank)`
 - Determines the rank (identification number) of the calling process in the communicator
 - `rank` is the output
 - The rank is used by MPI for process identification
 - The rank number is unique within a communicator
 - Processes are numbered sequentially, starting from zero

```
1 #include "mpi.h"
2
3 int rank;
4 MPI_Comm comm;
5
6 ...
7 MPI_Comm_rank(comm, &rank);
8 if (rank==0) {
9     ... code for process 0 ...
10 }
11 else {
12     ... code for the other processes ...
13 }
```

<https://web.archive.org/web/20070909174518/http://www2.informatik.uni-jena.de/cmc/racluster/mpi-leitfaden>

MPI Functions – Selection of important Functions (5/5)

- `MPI_Finalize()`
 - Stops the MPI environment
 - All processes need to call `MPI_Finalize()`, before they kill themselves
- `MPI_Abort(MPI_Comm comm, int errorcode)`
 - Terminates the MPI environment
 - `comm` = Communicator (handle), whose processes are terminated
 - `errorcode` = Error code, which is returned to the calling environment

```
1 #include "mpi.h"
2
3 int main(int argc, char *argv[]) {
4
5     int          errorcode;
6     MPI_Comm     comm;
7
8     ...
9     MPI_Abort(comm, errorcode);
10    ...
11 }
```

Simple MPI Example (1/3)

- Start a MPI cluster with 3 nodes (1 master, 2 slaves) in EC2
 - Start 3 instances (ami-06ad526f) with Ubuntu 11:04 in US-East
- Install the required packages in all instances:

```
$ sudo apt-get -y install make gcc g++ openmpi-bin openmpi-common libopenmpi-dev
```

- Generate public key on the master:

```
$ ssh-keygen -t rsa
```

- Append the content of .ssh/id_rsa.pub (master) to .ssh/authorized_keys (slaves)
- Insert into /etc/hosts on the master:

10.252.186.133	domU-12-31-38-00-B5-77.compute-1.internal	master
10.223.49.141	domU-12-31-38-07-32-63.compute-1.internal	node1
10.253.191.213	domU-12-31-38-01-B8-27.compute-1.internal	node2

- Create file hosts.mpi on the master with this content:

```
master  
node1  
node2
```

Simple MPI Example (2/3) – Say Hello to the Processors

```
1 #include <mpi.h>
2 #include <stdio.h>
3 #include <unistd.h>
4
5 int main(int argc, char *argv[]) {
6     // variable definitions
7     int size, rank, namelen;
8     char processor_name[MPI_MAX_PROCESSOR_NAME];
9
10    // Start MPI environment
11    MPI_Init(&argc, &argv);
12
13    // How many processes contains the MPI environment?
14    MPI_Comm_size(MPI_COMM_WORLD, &size);
15
16    // What is our number we in the MPI environment?
17    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
18
19    // What is the name of the processor?
20    MPI_Get_processor_name(processor_name, &namelen);
21
22    // Output of each process
23    printf("Ich bin Prozess Nr. %d von %d auf %s\n", rank, size, processor_name);
24
25    // Stop MPI environment
26    MPI_Finalize();
27
28    // Kill application with exit code 0 (EXIT_SUCCESS)
29    return 0;
30 }
```


Simple MPI Example (3/3)

- Compile program:

```
$ mpicc hello_world.c -o hello_world
```

- Distribute the program in the cluster:

- The program must be stored on each node in the same directory!

```
$ scp hello_world node1:~
$ scp hello_world node2:~
```

- Program execution (6 processes) in the cluster:

```
$ mpirun -np 6 --hostfile hosts.mpi hello_world
Ich bin Prozess Nr. 0 von 6 auf domU-12-31-38-00-20-38
Ich bin Prozess Nr. 1 von 6 auf ip-10-126-43-6
Ich bin Prozess Nr. 2 von 6 auf domU-12-31-38-00-AD-95
Ich bin Prozess Nr. 4 von 6 auf ip-10-126-43-6
Ich bin Prozess Nr. 3 von 6 auf domU-12-31-38-00-20-38
Ich bin Prozess Nr. 5 von 6 auf domU-12-31-38-00-AD-95
```

- The CPUs respond in random order
 - What is the reason?

MPI Functions – Send-/Receive (1/3)

- `MPI_Send(int buffer, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm)`
 - Sends a message (blocking) to another process in the communicator
 - `buffer` = first address of the transmit buffer
 - `count` = number of elements in the transmit buffer (not negative)
 - `datatype` = MPI data type of the elements in the transmit buffer
 - `dest` = rank of the receiver process in the communicator
 - `tag` = ID for distinguishing the messages
 - `comm` = communicator
 - All parameters are input parameters
 - The function sends `count` data objects of type `datatype` from address `buffer` (\implies transmit buffer) with the ID `tag` to the process with rank `dest` in communicator `comm`

MPI Data Types

MPI data type	C data type	Used for...	Size	Value range
MPI_CHAR	signed char	Chars	1 Byte	-127 ... +127
MPI_UNSIGNED_CHAR	unsigned char	Chars	1 Byte	0 ... 255
MPI_SHORT	signed short int	Integers	2 Bytes	-32,768 ... 32,767
MPI_UNSIGNED_SHORT	unsigned short int	Integers	2 Bytes	0 ... 65,535
MPI_INT	signed int	Integers	2-8 Bytes	Depends on the architecture
MPI_UNSIGNED	unsigned int	Integers	2-8 Bytes	Depends on the architecture
MPI_LONG	signed long int	Integers	4 Bytes	-2.147.483.648 ... 2.147.483.647
MPI_UNSIGNED_LONG	unsigned long int	Integers	4 Bytes	0 ... 4.294.967.295
MPI_FLOAT	float	Floating point numbers	4 Bytes	Single precision
MPI_DOUBLE	double	Floating point numbers	8 Bytes	Double precision
MPI_LONG_DOUBLE	long double	Floating point numbers	16 Bytes	Quadruple precision
MPI_BYTE	—	Floating point numbers	1 Byte	0 ... 255

- The integer value range depends on the used C compiler used and architecture (2, 4 or 8 Bytes)

MPI Functions – Send-/Receive (2/3)

- `MPI_Recv(int buffer, int count, MPI_Datatype datatype, int source, int tag, MPI_Comm comm, MPI_Status status)`
 - Receive a message (blocking)
 - `buffer` = first address of the receive buffer \leftarrow *output parameter*
 - `count` = number of elements in the receive buffer (not negative)
 - `datatype` = MPI data type of the elements in the receive buffer
 - `source` = rank of the sender process in the communicator or `MPI_ANY_SOURCE`
 - `tag` = ID for distinguishing the messages. For receiving messages with any identifier, the constant `MPI_ANY_TAG` is used
 - `comm` = communicator
 - `status` = contains the rank of the sender process `source` and the message identifier `tag` \leftarrow *output parameter*

MPI Functions – Send-/Receive (3/3)

- `MPI_Get_count(status, datatype, count)`
 - Determines the number of received elements
 - `count` = number of received elements (not negative) \Leftarrow *output parameter*
 - `status` = status upon the return of the receive operation
 - `datatype` = MPI data type of the elements in the receive buffer

```
1 #include "mpi.h"
2 #define MAXBUF 1024
3
4     int          i, count;
5     void         *recvbuf;
6     MPI_Status   status;
7     MPI_Comm     comm;
8     MPI_Datatype datatype;
9
10    ...
11    MPI_Recv(recvbuf, MAXBUF, datatype, 0, 0, comm, &status);
12    MPI_Get_count(&status, datatype, &count);
13    for (i=0; i<count; i++) {
14        ...
15    }
16    ...
```

Simple MPI Example (1/2) – Send and Receive

```
1 #include "mpi.h"
2 #include <stdio.h>
3
4 int main(int argc, char *argv[]) {
5     int size, rank, dest, source, rc, count, tag=1;
6     char inmsg, outmsg='x';
7     MPI_Status Stat;
8     MPI_Init(&argc,&argv);
9     MPI_Comm_size(MPI_COMM_WORLD, &size); // Get number of processes in the MPI environment
10    MPI_Comm_rank(MPI_COMM_WORLD, &rank); // Get process number in the MPI environment
11
12    if (rank == 0) {
13        dest = 1;
14        source = 1;
15        MPI_Send(&outmsg, 1, MPI_CHAR, dest, tag, MPI_COMM_WORLD);
16        MPI_Recv(&inmsg, 1, MPI_CHAR, source, tag, MPI_COMM_WORLD, &Stat);
17    }
18    else if (rank == 1) {
19        dest = 0;
20        source = 0;
21        MPI_Recv(&inmsg, 1, MPI_CHAR, source, tag, MPI_COMM_WORLD, &Stat);
22        MPI_Send(&outmsg, 1, MPI_CHAR, dest, tag, MPI_COMM_WORLD);
23    }
24
25    MPI_Get_count(&Stat, MPI_CHAR, &count);
26    printf("Task %d: Received %d char(s) from task %d with tag %d \n",
27           rank, count, Stat.MPI_SOURCE, Stat.MPI_TAG);
28    MPI_Finalize(); // Stop MPI environment
29    return 0; // Kill application with exit code 0 (EXIT_SUCCESS)
30 }
```

Simple MPI Example (2/2) – Send and Receive

Source of the example

http://coewww.rutgers.edu/www1/linuxclass2005/lessons/lesson13/sec_8.html

- Compile program:

```
$ mpicc sendrecv.c -o sendrecv
```

- Distribute the program in the cluster:
 - The program must be stored on each node in the same directory!


```
$ scp sendrecv node1:-  
$ scp sendrecv node2:-
```

- Program execution (2 processes) in the cluster:

```
$ mpirun -np 2 --hostfile hosts.mpi sendrecv  
Task 0: Received 1 char(s) from task 1 with tag 1  
Task 1: Received 1 char(s) from task 0 with tag 1
```

More Send and Receive Operations


- MPI_Send and MPI_Recv implement *one-to-one* communication
- MPI_Bcast and MPI_Scatter implement *one-to-many* communication
- MPI_Gather implements *many-to-one* communication
- MPI_Allgather implements *many-to-many* communication



The functions MPI_Bcast, MPI_Scatter, MPI_Gather and MPI_Allgather must be called by all processes in the communicator!

MPI Functions – Broadcast Sending (1/2)

- `MPI_Bcast(int buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm)`



- Send a message of process root to all other processes in the communicator
 - `buffer` = first address of the transmit buffer
 - `count` = number of elements in the transmit buffer (not negative)
 - `datatype` = MPI data type of the elements in the transmit buffer
 - `root` = rank of the sender process in the communicator
 - `comm` = communicator
- All processes in the communicator must call the function

