

Beowulf cluster

A **Beowulf cluster** is a computer cluster of what are normally identical, commodity-grade computers networked into a small local area network with libraries and programs installed which allow processing to be shared among them. The result is a high-performance parallel computing cluster from inexpensive personal computer hardware.

The name *Beowulf* originally referred to a specific computer built in 1994 by Thomas Sterling and Donald Becker at NASA.^[1] The name "Beowulf" comes from the Old English epic poem of the same name.^[2]

No particular piece of software defines a cluster as a Beowulf. Typically only free and open source software is used, both to save cost and to allow customisation. Most Beowulf clusters run a Unix-like operating system, such as BSD, Linux, or Solaris. Commonly used parallel processing libraries include Message Passing Interface (MPI) and Parallel Virtual Machine (PVM). Both of these permit the programmer to divide a task among a group of networked computers, and collect the results of processing. Examples of MPI software include Open MPI or MPICH. There are additional MPI implementations available.

Beowulf systems operate worldwide, chiefly in support of scientific computing. Since 2017, every system on the Top500 list of the world's fastest supercomputers has used Beowulf software methods and a Linux operating system. At this level, however, most are by no means just assemblages of commodity hardware; custom design work is often required for the nodes (often blade servers), the networking and the cooling systems.

Contents

Development

Operating systems

See also

References

Bibliography

External links

Development

A description of the Beowulf cluster, from the original "how-to", which was published by Jacek Radajewski and Douglas Eadline under the Linux Documentation Project in 1998:^[3]

Beowulf is a multi-computer architecture which can be used for parallel computations. It is a system which usually consists of one server node, and one or more client nodes connected via Ethernet or some other network. It is a system built using commodity hardware components, like any PC capable of running a Unix-like operating system, with standard Ethernet adapters, and switches. It does not contain any custom hardware components and is trivially reproducible. Beowulf also uses commodity software like the FreeBSD, Linux or Solaris operating system, Parallel Virtual Machine (PVM) and Message Passing Interface (MPI). The server node controls the whole cluster and serves files to the client nodes. It is also the cluster's console and gateway to the outside world. Large Beowulf machines might have more than one server node, and possibly other nodes dedicated to particular tasks, for example consoles or monitoring stations. In most cases client nodes in a Beowulf system are dumb, the dumber the better. Nodes are configured and controlled by the server node, and do only what they are told to do. In a disk-less client configuration, a client node doesn't even know its IP address or name until the server tells it.

One of the main differences between Beowulf and a Cluster of Workstations (COW) is that Beowulf behaves more like a single machine rather than many workstations. In most cases client nodes do not have keyboards or monitors, and are accessed only via remote login or possibly serial terminal. Beowulf nodes can be thought of as a CPU + memory package which can be plugged into the cluster, just like a CPU or memory module can be plugged into a motherboard.

Beowulf is not a special software package, new network topology, or the latest kernel hack. Beowulf is a technology of clustering computers to form a parallel, virtual supercomputer. Although there are many software packages such as kernel modifications, PVM and MPI libraries, and configuration tools which make the Beowulf architecture faster, easier to configure, and much more usable, one can build a Beowulf class machine using a standard Linux distribution without any additional software. If you have two networked computers which share at least the /home file system via NFS, and trust each other to execute remote shells (rsh), then



The Borg, a 52-node Beowulf cluster used by the McGill University pulsar group to search for pulsations from binary pulsars

External image

The original Beowulf cluster built in 1994 (<https://www.computerhistory.org/revolution/supercomputers/10/71/909>) by Thomas Sterling and Donald Becker at NASA. The cluster comprises 16 white box desktops each running a i486 DX4 processor clocked at 100 MHz, each containing a 500 MB hard disk drive, and each having 16 MB of RAM between them, leading to a total of roughly 8 GB of fixed disk storage and 256 MB of RAM shared within the cluster and a performance benchmark of 500 MFLOPS.



Detail of the first Beowulf cluster at Barcelona Supercomputing Center

it could be argued that you have a simple, two node Beowulf machine.

Operating systems

As of 2014 a number of [Linux distributions](#), and at least one [BSD](#), are designed for building Beowulf clusters. These include:

- [MOSIX](#), geared toward computationally intensive, IO-low applications
- [ClusterKnoppix](#), based on [Knoppix](#)
- [Kerrighed](#)
- [Rocks Cluster Distribution](#)
- [DragonFly BSD](#)
- [Quantian](#), a [live DVD](#) with scientific applications, based on Knoppix and ClusterKnoppix
- [Kentucky Linux Athlon Testbed](#)
- [PelicanHPC](#), based on [Debian Live](#)

A cluster can be set up by using Knoppix bootable CDs in combination with [OpenMosix](#). The computers will automatically link together, without need for complex configurations, to form a Beowulf cluster using all CPUs and [RAM](#) in the cluster. A Beowulf cluster is scalable to a nearly unlimited number of computers, limited only by the overhead of the network.

Provisioning of operating systems and other software for a Beowulf Cluster can be automated using software, such as [Open Source Cluster Application Resources](#). OSCAR installs on top of a standard installation of a supported Linux distribution on a cluster's head node.



A home-built Beowulf cluster composed of [white box PCs](#)

See also

- [Aiyara cluster](#)
- [HTCondor](#)
- [Grid computing](#)
- [Kentucky Linux Athlon Testbed](#)
- [Maui Cluster Scheduler](#)
- [Open Source Cluster Application Resources](#)
- [Stone Soupercomputer](#)
- [Sun Grid Engine](#)

References

1. Becker, Donald J; Sterling, Thomas; Savarese, Daniel; Dorband, John E; Ranawak, Udaya A; Packer, Charles V (1995). "BEOWULF: A parallel workstation for scientific computation" (<http://www.phy.duke.edu/~rgb/brhma/Resources/beowulf/papers/ICPP95/icpp95.html>). *Proceedings, International Conference on Parallel Processing*. **95**.
2. See Francis Barton Gummere's 1909 translation, reprinted (for example) in *Beowulf* (<https://books.google.com/books?id=ZCA4DbBCUIAC>). Translated by Francis B. Gummere. Hayes Barton Press. 1909. p. 20. ISBN 9781593773700. Retrieved 2014-01-16.
3. Radajewski, Radajewski; Eadline, Douglas (22 November 1998). "Beowulf HOWTO" (<http://ibiblio.org/pub/Linux/docs/HOWTO/archive/Beowulf-HOWTO.html#ss2.2>). *ibiblio.org*. v1.1.1. Retrieved 8 June 2021.

Bibliography

- *Beowulf Cluster Computing With Windows* by Thomas Lawrence Sterling 2001 ISBN 0262692759 MIT Press
- *Beowulf Cluster Computing With Linux* by Thomas Lawrence Sterling 2001 ISBN 0262692740 MIT Press

External links

- [The Beowulf Cluster Site \(http://www.beowulf.org/\)](http://www.beowulf.org/)
- [Beowulf parallel computing \(https://curlie.org/Computers/Parallel_Computing/Beowulf/\)](https://curlie.org/Computers/Parallel_Computing/Beowulf/) at [Curlie](#)
- [A detailed description of building a Beowulf cluster \(https://www.final-labs.org/dev/protogonus\)](https://www.final-labs.org/dev/protogonus)

Retrieved from "https://en.wikipedia.org/w/index.php?title=Beowulf_cluster&oldid=1113301456"

This page was last edited on 30 September 2022, at 20:26 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License 3.0; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.