What is Clustering and how High Performance Computing impacts on Business?

www.ainkaboot.co.uk
info@ainkaboot.co.uk

What is Clustering and how High Performance Computing impacts on Business?

By Hanni Ali

This paper will discuss the growing trend in the use of distributed computing clusters to increase performance and reliability of computing applications. The focus of the paper is on cost effective system design and introduces novel approach conceived by Ainkaboot, which aims to ease migration to this kind of technology.

Clustering reduces the reliance on having individually powerful processors and allows the use of cheaper processors to achieve performance greater than a single more powerful processor.

The concept of Clustering has been actively used for the past decade to provide high performance and high availability computers. It refers to the procedure of connecting more than one computing unit together across a network to provide a parallel processing environment. This provision allows the distribution of processes across the computers. Clusters generally have a dedicated network environment and dedicated computing nodes. Occasionally clusters may be found which make use of spare computing cycles on workstations, such as on a Trading Floor or within a University Department.

Grids extend the concept of Clustering beyond the local network and provide a means of utilising the computing ability of multiple clusters spread across a Wide Area Network such as a company Intranet or globally across the Internet. Grids remove the requirement for the user to submit a job to a particular machine known to have the data or the need to consider any physical aspect of the process at all. Instead the user submits the job to the Grid as a whole and the Workload Management Software manages the job.

We begin with an overview of current and future applications of Clusters and ultimately Grids in modern business, emphasising the diversity of clustering and High Performance Computing (HPC) and the wide range of applications. We go on to explain some of the issues faced in moving to a clustered environment and how Ainkaboots innovative Octimod systems aim to provide a simple platform for clustering and reduce the costs and hurdles traditionally associated with HPC.

Business and Clustering

The use of computing power across a range of industries is continually increasing. Since 2003 the use of HPC has more than doubled. In 2005, clustering - which has become the dominant method of most supercomputer design - experienced impressive growth of 70.8%. Now almost 50% of the global technical server market is made up of High Performance Computing (HPC) clusters. [1]

Today clusters dominate the Top500.org rankings

Clusters are designed with High Performance in mind. Originally the concept was to provide supercomputing power at commodity prices, but today clusters dominate the Top500.org rankings - a list of the Top 500 supercomputers in the world - proving their importance in HPC.

Clustering provides large numbers of low cost computing cycles, hardware redundancy and hence ultimate reliability. The applications of clustering are numerous:

- Software Development: Games developers write large amount of code that can take hours to compile on a single machine. By making use of distributed compiling to compile the program, huge reductions in the length of a development cycle are achievable.
- Special Effects Houses: Complex images and computer graphics require vast processing power to render. In the past few years we have witnessed a massive increase in the use of Computer Graphics (CG) for films television, and in advertising. The ever higher quality three dimensional scenes, with ever increasing detail can take weeks to compile on a single machine. Through the use of distributed rendering the time it takes for images and animations to render can be slashed to mere hours. Rendering is a simple linear process and by doubling the number of processors - dividing the task in half - you can halve the time it takes to render.
- Financial Institutions: Both high performance and high availability solutions are relevant in finance. The use of Monte-Carlo simulations is important for the vast majority of pricing, options, futures, swaps and statistical arbitrage. By using clusters and writing code to take advantage of message passing, the accuracy of these simulations can be increased while still reducing their computing time. By running mission critical trading systems on high availability redundant systems with multiple failover, 99.999% uptime is achievable. Physicists have been using Clusters and Grid technology for years to perform Monte-Carlo simulations for Quantum and High Energy Physics.

Clusters can provide High Availability features such as, 24/7 uptime, database replication and failover.

- Web Hosting: High availability, easy management, reliability and 24/7 uptime are all key selling points for Webhosts, clusters address all these issues and the current widespread use of LAMP (Linux, Apache, MySQL, PHP) platform for website hosting, allows straight foreward migration to clustering technologies.
- Scientific Computing: Chemical and Biological companies use simulations to determine how molecules will behave under various conditions. Modelling is used extensively in Pharmaceuticals for drug development and simulation.
- Automotive and Aerospace Engineering: Detailed simulations, Computer Aided Design (CAD) and rendering are necessary, time can be saved through using clusters. Aerodynamics can be determined and improved without building scale models of the designs.
- Enterprise Size Companies: These have many employees and high availability clusters could replace terminal services, e-mail servers and database servers especially as companies tend towards web based offices, the use of clusters to provide these services provides complete redundancy and hence 24/7 uptime.

Designing the Cluster

Hardware selection is key to building a functional cluster, it is important to prioritise the different aspects of the system and the environment it will be situated in. Consider the following questions:

- What are the processing requirements?
- What are the storage requirements?
- How much power capacity is there?
- What are the limiting factors of the software?
- Will the software scale well?

In order to take advantage of the hardware it is necessary to select appropriate software to provide the functionality necessary for the application.

Select the most appropriate software before deciding on the hardware then tailour the hardware to maximise efficiency in the software

We will now examine a selection of possibilities which could be found on systems designed to solve the problems in the previous section. These layers can exist individually or in combination with other aspects depending on the overall functionality of a cluster.

In order to build a distributed compiling system there are a few simple requirements. The compiling nodes require identical versions of the compiler, all nodes must have access to the required libraries which must also be of the same version and the nodes must all have a distributed compiling daemon running. In general all this should be handled by the operating system. The most popular distributed compiling software daemon is called distcc and provides distributed compiling for C and C++.

Providing distributed rendering can come in two different forms. Either distributing sections of an individual scenes or distributing multiple scenes of an animation. There are some open source rendering and ray tracing programs as to provide both these functions, but the proprietary software if far more developed and almost all provide functionality for distributed rendering.

Modelling and simulation accuracy can be improved or shortened in Finance, Science and Engineering through the use of parallel processing environments. Message passing environments provide the means for a task to be spread across multiple processors. This is ideal for programs which require simulating multiple instances of a situation particularly loops. There are numerous different environments to choose from, Parallel Virtual Machine (PVM) and Message Passing Interface (MPI) are the main options, this does require writing code that takes advantage of multiple threads.

High Availability (HA) computing systems vary in their requirements, but generally the main aim is 100% uptime. The hardware selection is one of the most important features of HA, multiple failover and redundancy at every level is vital. The software needs to provide the seamless and transparent migration of processes in such a manner that the administrator doesn't consider each machine on an individual basis, but rather the system as a whole. Apache provides many clustering features for web serving and features in MySQL and database replication provide increases in speed, data security and availability. Load balancing is important for spreading the stress across multiple machines, reducing the slowdown under high demand situations.

Addressing the Issues of HPC

There are numerous hurdles traditionally associated with clustering:

- Initial Hardware Cost
- Building, installing and configuring the cluster
- System administration
- Installing and configuring new applications for the cluster
- Complexity of the networking
- I/O and network performance/latency
- Facility Issues, space, power, cooling, noise

Eliminating as many of these issues as possible, facilitates the adoption of Clustering as an alternative to standard server technologies.

There are a number of ways to reduce the initial costs for High Performance Computing systems. The concept of Beowulf clustering pivots around using cheap commodity hardware and readily available networking technologies to provide high performance. The first and foremost consideration when selecting the technology to use in a cluster is to fully understand the requirements of the cluster: main function, future requirements, possible future expansion and the lifetime of hardware.

It is possible to reduce the cost of purchasing large numbers of machines by using embedded processors and sticking to mainstream technology. Embedded processors provide less processing power per chip, but generally consume less power and take up less space.

The design and construction of a cluster is not a simple task and it is highly recommended to get the advice of experts before making decisions. Designing and building the cluster can prove decisive in the eventual performance of the system and when making decisions about purchasing a cluster it is important to define the purpose of the cluster now - and what may be expected of it in the future - in order to plan effectively.

Get advice before constructing a cluster

In order to simplify system administration and application installation it is important that the operating system is designed for clustering and provides the necessary tools for running the system.

Ainkaboot Systems: All in One Solution

The original concept behind clusters was to provide supercomputing power at commodity prices. The term Beowulf was first coined by Thomas Sterling and Don Becker researchers at the NASA Goddard facility. However the term Beowulf cunjours images of rows of desktops on shelves. Not a situation most System Administrators relish. [4]

Ainkaboot systems look to address and eliminate as many problems as possible. The clusters are designed to make use of industry standard hardware and run open source software generally the computing nodes are diskless (unless they are performing database replication). They use Gigabit Ethernet for networking and SATA-II hard disk drives for storage. Small form factor motherboards such as VIA's Mini-ITX series increase the density of the systems and lowers power consumption. They provide

access to numerous processors from VIA's own C7 low power embedded processor to Dual Core Intel and AMD processors. Choice is a key feature of Ainkaboot Systems providing a complete base system to build on top of and a reliable versatile operating system to run the system.

Embedded processors consume much less power and are cheaper to produce, for low cost low power clustering they are ideal.

The System consist of a 4U, 19", half depth Unit which may be populated with up to eight Modules. Modules can be either Processor Modules or Hard Disk Drive (HDD) Modules. The modules are powered by a distributed power supply, are hot swapable and simple to manage.

Embedded motherboards are characteristically compact, allowing much higher densities of motherboards and generally use embedded processors which are far more power efficient. In addition, on-board gigabit networking removes any need for additional hardware and embedded processors are generally more power efficient for example, VIA's C7 processor is one of the most power efficient x86 processors on the market[2], AMD's embedded Opterons also exhibit excelent power efficiency[3], while Intels Core architecture shows vast improvements over their Netburst architecture.

The design provides hot-swapable processor or HDD modules, allowing easy upgrade and maintenance. The modules are mounted in a 4U housing Unit which provides the power distribution. The units are light and designed for rack mounting.

The networking technology should be specially selected to take advantage of the best possible switching technology reducing network latency, to depend on the clusters application. Gigabit switches are cheap and reliable providing high speed interprocessor communication.

In order to further increase the density our units are only half the depth of a server rack and designed for both front and rear mounting. In a 19 inch cabinet 16 modules can fit into 4U. Cooling and heat dissipation are important factors to achieve maximum performance from the hardware and the modules provide excellent airflow.

A distributed redundant power supply provides a 12V supply to the modules. This saves space and reduces heat within the units. Additionally, reduced noise is possible through using quieter fans and large heat sinks.

Features of the Ainkaboot System

- Easily scalable systems reducing costs and complexity of upgrades
- Lower power consumption
- Reduced cooling requirements
- Commodity networking, hardware and reliable open source software
- Optimisation and software level improvements take full advantage of the hardware
- High density solutions over 140 motherboards in a single industry standard 19" cabinet
- Lower power consumption and cooler hardware for a significantly quieter environment
- Simplified System administration through KlustOS Management Utilities Suite, saving administrators time and effort
- Cluster and Grid ready

- Easy to integrate, rack mountable and modular
- Multi-level redundancy
- High Performance, Density, Availability
- Easy to manage and expand
- Cost effective and efficient design for lower setup and running costs
- Eco-friendly
- Powerful Open Source Operating System (KlustOS)
- Custom software is simple to roll out across the cluster using Portage.

KlustOS is the powerful Clustering Operating System customised for your specific application by Ainkaboot.

Further details and technical specifications can be found in the Ainkaboot Octimod System brief. [5]

Summary

The march of Clustering technology towards widespread mainstream use is inevitable. The demand for 24/7 reliability and the pressure on storage space and processing power can all be addressed by Clusters. Clustering provides a reliable scalable platform for building the next generation of distributed applications and allows increased reliability for mission critical systems.

The pressures on space and power at the Data Centre require higher density systems which consume less power, but provide more. The increase in the number of systems means it is more important that the management of the centre is simple and efficient. Ainkaboot's Systems provide the means for a cost effective and reliable solution.

For more information about using cluster technology in your data centre contact an Ainkaboot consultant.

Contact

Email: sales@ainkaboot.co.uk
Website: http://www.ainkaboot.co.uk/

Further Details and References

References

[1] IDC HPC report summary:
`http://www.idc.com/getdoc.jsp?containerId=IDC_P7`

[2] C7 Processor:
`http://www.via.com.tw/en/products/processors/c7/`

[3] Embedded Opteron:
`http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_13515,00.html`

[4] Beowulf - Further information:
`http://www.beowulf.org`

[5] Ainkaboot Octimod System brief:
`http://www.ainkaboot.co.uk/products/octimod_brief.pdf`

Ainkaboot
Conception to Completion

www.ainkaboot.co.uk
info@ainkaboot.co.uk