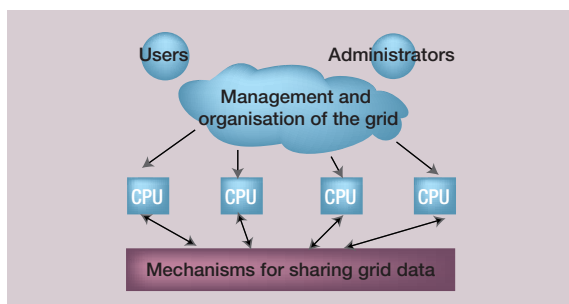


Grid computing: from science to finance

Long restricted to the field of scientific supercomputing, computing grids have seen their fields of use expand little by little, now reaching the no less refined universe of large financial groups. In fact, finance has this in common with the hard sciences in requiring significant processing power based on significant amounts of data.

> What is Grid Computing?

Technical definition



“Grid computing” refers to a virtual infrastructure made up of a set of potentially shared, distributed, heterogeneous, delocalised, autonomous computing resources. Unlike “cloud computing”, which offers applications, grid computing aims to deliver raw computing power as well as large storage space and memory to end users.

To orchestrate all of the exchanges of an application distributed among the various components of the grid, middleware handles the network layers and the necessary software services. All of the computing units constituting the grid (nodes) are interconnected through high-bandwidth networks (from 10 Mbit/s to 10 Gbit/s) in order to maintain a high level of overall performance. This performance is measured in the number of floating-point operations per second (“flops”) that the grid is capable of completing. Thus, Earth Simulator, the grid developed by NEC for global climate modelling, claims 40 Tflops.

For the sake of comparison, a Pentium 4 processor is 12,000 times slower. Also note that a list of the 500 most powerful machines in the world is maintained. In November 2009, when the list was most recently published, the record was held by the Jaguar system of the Oak Ridge National Laboratory (Tennessee) with a capacity of 1,759 Tflops.

Dual development

The beginnings of the concept of shared computing are generally placed in the mid-1960s, when the processing power of computing units was very low. Connecting these mainframes computer in order to obtain a new, better-performing virtual unit emerged as a response to lack of performance of the computers of those days.

Yet, the true growth in grid computing only arrived in the early 1990s. With the explosion of personal computing and, with the help of Moore's Law, the computing capacities of processors, it has become possible to create low-cost grids consisting of "serial" computers connected by an Ethernet network. NASA was the first to use this type of grid in 1994.

Since the early 2000s, we have seen a dual development: new technologies have emerged, such as the Desktop Grid. The development of the Internet has thus allowed **thousands of little-used PCs to be connected on a global scale in order to roll out enormous computing power.**

The seti@home project can thus boast computing power of 67 Tflops. At the same time, the field of use of computing grids has expanded from its original scope, scientific computing (fluid physics, biotechnologies, energy, etc.), to the apparently more abstruse environments of finance. The endlessly growing sophistication of financial products, the needs for responsiveness in decision-making, and the determination of magnitudes, such as pricing, have made the implementing of computing power that only grid computing can provide necessary. Large financial groups have therefore equipped themselves with "homemade" computing grids in order to be able assess risks, among other things, based on the composition of financial products.

As such, BNPParibas equipped itself with a 500 PC computing grid in 2005, which it expanded in late 2006 to nearly 6,000 stations! They were put to work for around six hours per night.

Players

The development and establishment of a computing grid is primarily a hardware challenge, and the race for computing power represents an incomparable technology showcase. It is therefore not surprising to find the leading global manufacturers on the list of players: IBM, NEC, Sun, and Fujitsu, alongside HP, Dell, Cray, and SGI. Publishers such as Oracle and Microsoft have also developed an offering.

IBM has thus developed a toolkit for grid computing (under GPL licence), which it has named Globus. It can manage resources, data, and information and naturally adapts to the high-end grid computing promoted by IBM. At the other end of the scale, independent start-ups, such as DataSynapse (acquired by Tibco in August 2009 for \$28 million) and Avaki (acquired by Sybase in May 2005), are publishing their own software: GridServer for DataSynapse or Avaki Grid for Avaki. Between these two extremes, Microsoft presented its own system, Windows 2003 Compute Cluster Edition, in 2005.

In a market estimated by IDC at \$12 billion in 2007, the historical players IBM, Fujitsu, NEC, and others are cornering the lion's share in the high-end market, where their know-how and the breadth of the necessary investments leave them without rivals. On the low end, where the focus is on more standard hardware solutions, the products from the purchased start-ups are particularly well positioned. For example, Tibco, formerly DataSynapse.

> Challenges

Technical challenges

Because of its very architecture, a grid is extremely reliable. Indeed, they are made up of redundant elements (RAID storage systems, for example). On the other hand, middleware includes functions for offloading and distributing the load, which ensure continuity of service when a node or link ends up failing. Therefore, **reliability rates exceeding 99%** have been observed.

Moreover, security issues have also been taken into account over time and as technological advances have been made. Thus, encryption systems have been developed (Oracle 10g, for example), and data access has also been very effectively managed using digital certificates.

At the same time, the race for computing power continues. The direction chosen for two to three years is that of **taking advantage of the processing capabilities of graphic processing units (GPU)**. These GPUs have the distinction of integrating many computing units designed to perform very simple calculations. This configuration makes it possible to maximise the parallelisation of scheduling of the most complex calculations and results in, compared to a solution based on standard processors, a drastic decrease in processing time, which can reach a factor of 10. This market is currently dominated by NVIDIA, which offers a special API called CUDA for the development of applications.

Value-for-money

In comparing the cost of acquiring servers – taking into account the depreciation of the machines and the operating conditions (bulkiness, air conditioning, etc.) – with that of setting up a grid, **the choice is clearly in favour of the grid. The total cost of ownership (TCO) is reduced.**

On the other hand, because of its very robustness, Grid Computing makes it possible to obtain “24/24, 7/7”, dear to the IT departments of the entire world. Indeed, the impact of hardware failures has diminished. Grid computing can thus deliver the results of its calculations, regardless of the time zone, allowing for continuity of operations throughout the world.

The economic impact is also important in the case of the computing grid based on desktop computers, which optimises the resources of the processors by operating them at night without any human activity. The IT infrastructure’s productivity is therefore improved.

Formerly reserved for heavy scientific work, computing grids, benefiting from the technology boom in computer science, have become widespread and are now used successfully in the world of finance. Reliable and flexible, grids make it possible to optimise the use of the computer infrastructure while delivering computing power to meet the heavy processing needs required by the implementation of new financial products. ■