The Four Stages of Grid Adoption

Songnian Zhou
Platform Computing

We rarely hear the term "supercomputing" anymore. In the early 1980s, it referred to a unique application of computer hardware that enabled engineers to simulate automobile crashes, assess the aerodynamics of composite airframes, or evaluate the properties of reactive armor in military equipment. The idea took the world by storm and fostered an entire market segment in the hardware industry.

Since then, rather than fade away, the need for supercomputers in engineering and science has grown at an explosive rate. Access to high-power low-cost computational capability has spawned a new era in product development. Today, many users are taking advantage of technology that allows them to work with clusters of commodity-based compute servers and workstations that expand and contract based on their application requirements. This technology is grid computing.

> The four stages result in capabilities and benefits every step of the way.

Grid computing runs applications on networked computers by integrating and virtualizing distributed computing resources and delivering them to the applications where and when needed, based on pre-established sharing policies. This is especially suitable for compute and data-intensive applications that thrive on the aggregate power of many processors, interconnected to form a cluster. Clusters and grids offer tremendous flexibility and reliability by enabling shared resources whereby tasks traditionally too large or complex to compute in a reasonable period on a single computer can be completed in a fraction of the time using a group of interconnected computers. The adoption of grid computing has evolved from dedicated clusters to shared enterprise grids, but its effective use and management requires new skills and IT practices.

In the first of the four stages of grid adoption, a single application runs on a cluster of commodity servers or workstations based on x86 chips. These clusters can lead to dramatic cost savings as organizations move away from expensive servers and supercomputers. Users further benefit from the advances of cost-effective multicore X86 chip technology by Intel and AMD. If an application needs more power, it is simple to add more nodes to the cluster to "scale out" rather than add a completely new server. SMP servers are replaced by "cluster servers" and in response to the demand, computer system vendors have begun packaging clusters with standardized interconnects and software to ease implementation.

In the second stage of grid adoption, organizations expand the size of their clusters and spread implementation to more departments and locations. Most importantly, this stage introduces a more efficient sharing of resources. Users begin running multiple applications in the same cluster, which leads to increased use rates and complexity. To make sure applications can get the power they need at the time they need it, while other parts of the organization are able to access power that would otherwise go to waste, resource sharing policies that establish scheduling priorities and time-based resource allocation need to be determined. Larger clusters have processors specifically configured based on the applications that use them. Workload manager software routes the jobs to the optimal computing resources to ensure their success even in cases of box failures or software problems.

The third stage of grid adoption shows resources being shared across the whole organization. With this setup, an organiza-

tion's separate lines of business have access to computing resources and pay a share of expensive licenses. It requires a significant philosophical shift as users adjust to sharing resources and IT departments adopt the responsibilities of allocating computing use. Many automobile and aerospace companies maintain a central HPC data center, typically housing many racks of compute clusters and massive storage.

In the fourth stage of grid adoption, enterprises seek to expand the grid infrastructure beyond modeling and analysis to integrate it with the data center. By scaling out clusters, migrating to a shared grid, or using an external service provider, enterprises can deliver applications as services. Thanks to transparency in usage patterns and the ability to customize service delivery, organizations will be able to more easily prioritize resource allocation and make more informed decisions about accessing external services.

Since our inception in 1992, at Platform Computing we have enabled a wide range of industries and applications to take advantage of distributed computing and grid technologies. This evolution started with electronics and aerospace designers who used their MCAD workstations to run simulations and eventually moved to grids and clusters adopted for programs as diverse as the Human Genome Project and investment banks whose risk analysis and pricing applications needed more computing power. Each of these organizations traveled along the four stages of grid adoption, gaining business capabilities and benefits every step of the way.

This is an exciting time for computing, and we are seeing the effective leveraging of technologies to develop better products, more effective cures, reliable and cost-effective supplies of energy, and more astounding gadgets. Through its use in all these industries, grid computing is benefiting the world. ■

Songnian Zhou *is the CEO of Platform Computing. Send e-mail about this commentary to DE-Editors@deskeng.com.*

Reprinted from *Desktop Engineering*, with adaptations, January 2008 www.deskeng.com