

Distributed computing **INTERVIEW** by Martin Courtney

Grid computing expands its reach

Platform Computing chief executive Songnian Zhou explains the lure of grid computing


IT Week: How much demand is there for distributed computing power among enterprises?

Songnian Zhou: We have customers ramping up application processing demands very quickly, and they cannot do it with the energy consumption and space limitations of their own datacentres. It is becoming an absolute must to utilise distributed resources, and individual applications cannot do that. Ensuring 80 to 90 per cent server utilisation, even 100 per cent in some cases, is the key, and accessing 1,500 processor cores simultaneously is pretty much the average for most usages.

What sorts of companies buy in to grid computing solutions?

The high-performance computing (HPC) and grid computing market is expanding from its traditional niche and is spreading deeper into the enterprise market. One group is peer to peer, where different organisations agree to share their resources in advance. The other is corporates who own all the computing resources and want to share it among different divisions and offices. Typical customers for us are those in the financial services sector, running pre- and post-market analytics applications, but we also have customers in aerospace and electronics industries, and across a bunch of vertical markets.

What does the Platform Symphony 4.0 software do?



Zhou: Grid computing allows customers to ramp up app processing demands very quickly

Symphony 4.0 is an application infrastructure designed to support distributed computing in parallel, clustered server environments. It is a decision-making piece of software, a central management module, which uses agents to organise all of the distributed computing resources, whether servers, switches, operating systems, storage, data or software licences, to ensure that they are all brought to bear. Platform tends to focus on the infrastructure software, and our ISV partners, like SAS, focus on the end user application. We also use our VM Orchestrator product specifically for distributed workloads hosted within virtual machines.

How do you make sure that everybody on the grid is guaranteed to get the resource they need, when they need it? We call it lending and borrowing. In the peer to peer model, defined high priority users will always have precedence when it comes to grabbing resources. We also have a fair share policy – if four peers have 120 CPU

About Songnian Zhou

- Songnian Zhou co-founded Platform Computing in 1992, after extensive research in the field of distributed resource management.
- He has a PhD in Computer Science from the University of California and held a faculty position at the University of Toronto as a professor of Computer Science and Electrical Engineering.

resources between them, for example, they are always guaranteed to get at least 30 of those whenever they need them, though they may be allocated 50 or more when their peers have them free.

How is security handled?

The computer added to the grid has all its operating systems and libraries populated in advance, and those machines are regularly updated. Symphony supports all security protocols, using transport security layer (TSL) for messaging, and LDAP for user access, as well as standard authentication and authorisation tools. We make sure there are no conflicts on those boxes in terms of applications existing on the same machine.