



**Customer:** Korwe Software  
**Country of Origin:** South Africa  
**Industry:** Software Development  
**Website:** <http://www.korwe.com>

### Company profile

Korwe Software develops core software systems to allow businesses and service providers to reach mobile users more effectively. It does original research and development into mobile solutions, with a particular focus on "context" – the relationship between users, and their typical behaviours. This provides a lot of useful information to allow a company or service provider to customise mobile content and the user experience to be faster, more effective and more pleasant.

### Hardware

High Performance Computing

### Partner

**Solutec:** Designers of custom made workstations, servers and clusters for the film, graphics and scientific industries.

## Testing software components in the real world

Korwe needed to test a software component it had developed in the real world. This would have involved getting an unbiased sample of 100 000 people willing to participate in a fairly laborious task, as well as requiring fair amount of real, private data from them. This was an unlikely prospect.

### Situation

Several members of Korwe's engineering team have considerable computational and analytic skills, as well as substantial experience in High Performance Computing (HPC, once known as "supercomputers"). The solution was obvious – as each of the "people" in the sample behave autonomously, the problem of finding 100,000 test subjects was potentially solvable by creating a computational solution on a cluster of computers.

The intention of the computer model was to emulate the behaviour of people in a social network, so that Korwe could test and optimise an algorithm that it developed to manage the transmission of information between nodes. Korwe's work is based on mobile phone users, and how they interact. Important in the model developed is that each "artificial" user, just like mobile users in real life, interacts directly only with the "head node" when using their phone (ie network service provider), but not directly with each other. This allows the models point of departure to be based on Milgram's "Six Degrees of Separation" theory.

For each "phone user", Korwe created a unique identity, with an "Address Books", "Diary" and other synthetic personal information. The degree of freedom that interested Korwe is the number of "Long Links" in the network – that is, beyond the "Short Links" that connect you directly to your nearest friend. "Long Links" connect distant people across intermediate links.

To get statistically meaningful data from the model, Korwe needed to run each simulation a number of times – requiring extremely fast computation to be practically usable as a research tool. This was a key part of the technology problem Korwe needed to solve, as it had limited access to HPC hardware. The company had a 16-core cluster, which was obtained from Solutec, a specialists in designing custom-made workstations, servers and clusters for the film, graphics and scientific industries.

This was a key part of the technology problem Korwe needed to solve, as it had limited access to HPC hardware.

### Solution

Korwe initially used a Rocks cluster, based on open source Linux technology, but in this implementation, it turned out to be difficult to manage. While system administration was not onerous, management of the OS was arcane and time-consuming. Korwe decided to use the Microsoft Windows HPC Server 2008 platform, which would provide a faster and simpler environment to set up the model and allow the right information to be gathered.

The Korwe team did an xhpl benchmark to convince themselves that the cluster was configured adequately to perform the required operation as they were not inclined to spend too much time tweaking kernel parameters and network configurations.



Part of the decision to adopt the Microsoft Windows HPC Server cluster was that the code developed by Korwe was written in generic C++, utilising Microsoft Visual Studio 2008. The final product was designed to be deployed onto both Windows and \*nix platforms, but in-house development work used a number of Microsoft tools. It was therefore both more convenient and efficient in terms of going from a "development" to a "production" environment to use a homogenous chain of tools.

A final reason to go with the Microsoft HPC Server solution was that Korwe had enhanced access to the software through Microsoft's BizSpark, a program that provides Software, Support and Visibility for Software Startups.

## Benefits

Korwe deployed the code onto the cluster via the Microsoft cluster management tool, and then monitored simulation jobs through a number of different built-in monitoring tools.

The results obtained from the simulations provided a wealth of insight into what Korwe should be doing to improve its models, and predict how different design parameters would play out in real-life networks of mobile users. Using an advanced computer model provided a way to reduce development time massively, as well as to gain insights into the complexity of interactions by linked individuals.

The models set up on the Windows HPC Server cluster proved more than beneficial in Korwe's research and development, and the company intends to continue using their Windows HPC Server cluster extensively in testing. This means that for each successive generation of Korwe's algorithms and simulations, they will have an easy to manage, and easy to program HPC environment. The Windows HPC solution is proving that "HPC" can also stand for "highly productive computing".