IDC
Analyze the Future

WHITE PAPER

An Overview of Compute-Intensive Utility Computing

Sponsored by: IBM

Christopher G. Willard, Ph.D. Earl Joseph, Ph.D.
Lionel Lamy
January 2007

Case Study: Exa Corporation Mission, Products, and Utility Computing Strategy

Founded in 1991 and employing over 100 people, Exa supports a worldwide customer base in industries ranging from automotive, aerospace, petroleum, chemical processing, environmental engineering, material processing, power generation, and HVAC (heating, ventilation, and air conditioning). Exa Corporation provides computer-aided engineering software and consulting services for computational fluid dynamics problems, including airflow over automobile and airplane surfaces, flows in plastic injection molding, the movement of fluids and heat transfer in nuclear reactors, and the effects of spray cooling and humidification in climate control systems. Exa has partnered with IBM to offer engineers on-demand access to Exa's flagship PowerFLOW simulation capabilities via a secure Internet connection.

Stephen Remondi, Exa's President, CEO and cofounder, describes the company's value proposition as "providing software and services for simulation of physical environments to help customers avoid physical testing, and be able to make time-critical decisions based on data and not only opinion."

Why Utility Computing and Partnering with IBM?

Exa began offering its software and services via a utility grid model in the 1999–2000 time frame. At that point the company purchased computers, installed its software environments, and managed all aspects of the computer operations. This provided the company with a fixed capacity infrastructure for utility-based software and services products. This approach was successful because the company understood the problem set, the applications, the computational requirements, and how to best fit the applications to the computer hardware.

However, Exa learned that the greatest limiting factor for its customers' use of computational methods, and thus its own growth, was access to HPC computing capacity. Exa tried to provide the resources to maintain and grow an HPC infrastructure for customer use, but it was constrained by capital limitations. The solution to what was essentially an infrastructure scaling problem came from IBM's DCCoD program. Mr. Remondi notes that "as soon as we heard what IBM was providing, we knew this was a perfect match for us, because of IBM's capital infrastructure and ability to scale-up the environment. So now, instead of having fixed capacity, we have from our point of view and our customer's point of view unlimited capacity."

The partnership with IBM provided a way for Exa to get customers access to its software running on high-performance computers, and it allowed Exa to get out of the capital equipment and systems management business and to focus its energies on its core competencies as a software development and consulting company.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

Getting Answers on Time: Opinions Versus Data

Exa's goal is to deliver computational capabilities that allow its customers to avoid building prototypes and doing physical testing (e.g., wind tunnel testing, climactic tunnel testing, road testing, etc.). Mr. Remondi notes that "Our customers want answers about how their products perform upfront in the design process. Problems that we are solving are very, very computationally intensive, which means a single simulation can consume from hundreds to, more typically, thousands of CPU hours."

Given such demands the issue for customers then becomes one of turnaround time versus system cost. For example, using "back of the envelope" calculations, a company might buy an eight-processor cluster to run a simulation that takes 2,000 hours for each design iteration. Such a run would take roughly 250 hours (or about 10 days) to complete. In contrast if the engineers could run that same simulation on a 128-processor cluster, it would take about 16 hours or be done roughly over night. "Engineers really start getting addicted to this ability to turn problems around," says Mr. Remondi. "It allows them to spend much more time thinking about the design and not just waiting."

Mr. Remondi explains why turnaround time is so important to the product design process, "What happens, for example, in an automotive development process is the design studio releases the data on Thursday and the designer review meeting with the VP of engineering and the VP of product development is on Monday morning. They need to know the direction they should take with the product. People can come to the meeting with opinions, or they can come with data. The people who win in those meetings, the people who have impacts on the vehicle design, are the people that show up with data — not just an opinion of what they think will happen or what they think is better."

Working with Exa and IBM: Customer-Use Models

Customers have a range of strategies in working the Exa and IBM DCCoD services. The strategies are based on what parts of the process are assigned to Exa and what parts are kept in house by the customer and then on the amount of computer time that the customer needs. Example use models include:

- **Exa as one-stop shop.** The customer provides the data, model specifications, etc. to Exa, which in turn is responsible for pre- and post-processing, the full problem setup, running the simulation, and so on. Exa then delivers a simulation results report to the customer. This approach is often an initial strategy for customers, as it allows them to determine where PowerFLOW fits into their design process and to get up and running while learning the operational details.

- **Exa as administrator.** The customer handles most or all of the pre- and post-processing, problem setup and so on, but hands-off the job to Exa to run. In this case, Exa is managing the system configuration use and is able to optimize costs by scheduling and sharing its base infrastructure at IBM.

- **Exa as software and setup provider.** Customers do the pre- and post-processing, set up their own job runs, and run the simulation. In this case, the customers purchase their own dedicated base presence at IBM, including storage and time for additional analysis. Exa is responsible for account enabling, account management, and making sure the software is available and properly configured on the IBM systems at run time. In this case Exa does not manage any customer data. This approach is effective for customers intending to use large amounts of machine time in a given year.

Individual customer organizations may use a variety of these models based on whether they use utility computing strategies to support peak workload problems, ongoing requirements, or special case studies.

Exa believes that it is able to price its solutions competitively based on total cost of ownership and/or internal chargeback models. In addition, its customers do not need to be concerned about long-term capital expenditures. Customers can budget on a per-project basis, and when the project is complete the cost of infrastructure goes away. A final cost advantage is that Exa has over 20 years of experience with the PowerFLOW application, and is able to assure that the computers running the code are optimally configured thus providing a high level of efficiency to minimize computing costs.

How the Utility Grid Model Works for Exa

Exa views its utility grid operations as an integral part of its overall business model. The utility model provides two major strategic advantages: elimination of capital equipment management and costs and expanded opportunities in its core business.

A Computerless Computer Company

Exa's utility grid model revolves around the company's ability to deliver computer cycles to its customers. However, its relationship with IBM allows the company to avoid investing in or owning those cycles. Mr. Remondi explains the ROI advantages of this approach as "a win because we can match our computer use to our customers' immediate requirements, and there is no upfront cost, so there is no investment to make. There is no 'I' in the ROI equation." For example, Exa might have thousands of nodes in its environment at IBM running its software for customers. And then, the next day, pare back to a 100 nodes as projects are completed. Thus the company is able to avoid expenses for systems that are not in use, while assuring a return on systems when they are in use.

Expanded Opportunity

The ability to scale up capacity virtually overnight allows Exa to respond to a very broad range of requests. Mr. Remondi explains, "When a customer says, 'I need the results by Monday,' we are always able to say 'yes'. If we say 'yes' some of the time and 'no, you have to wait in line' some of the time, we do not get the business. Either it's on demand, it's available, or it's not."

Mr. Remondi estimates that about 75% of Exa business supported by DCCoD would not be possible with a model based on in-house computer systems because customers would not accept the turnaround time. "They would have said 'Too late. I'm going to make a decision anyway, I'll build the model. I'll go to the wind tunnel. I'll guess. I have to have the answer on time. I don't have a choice.'"