synnovation

Vol. 1 Issue 3

Beating the Odds

Leading an organization through a risk-laden environment is often a game of chance. But enterprise risk management can help your odds.

The Quarterly Journal of the EDS Agility Alliance

BOB STEFKO

SUPER POWER

Supercomputers: The Unsung Heroes of Keeping Trouble at Bay

Many people think about risk mitigation in terms of floods, hurricanes, fires, and other natural disasters. These are critical events and organizations must plan for enterprise continuity as it relates to protecting civic infrastructures and businesses should disaster strike.

Professionals today are so technologically connected that they're constantly accessible, and the transfer of data—whether it's from man to machine or computer to computer—continues around the clock. IT infrastructures, and the data housed with them, are strategic to enterprises. >>

By Dr. Reza Rooholamini, Dell Inc.

Supercomputers are powerful solutions that let scientists run simulations to allow government agencies and businesses to prepare for events that could affect operations. Many industries are taking risk mitigation to another level—one that focuses on providing information that takes risks out of business decisions.

It used to be that car manufacturers, for example, would build clay or wood models to test aerodynamics in a wind tunnel. While this model of design is highly useful, it isn't efficient or cost effective because of the amount time it takes to constantly modify the design. Additionally, the modification could have an affect on another part of the design.

Many car manufacturers turned to supercomputers to allow them to put designs in a program and apply scientific models to check for aerodynamics, structural integrity, and a host of other design and engineering aspects. It saves time and money and reduces the risks associated with design decisions. It also allows them to get their product to market much faster.

Supercomputers allow companies to plan, identify, and analyze data more effectively—whether it's an oil company drilling for oil more accurately, a pharmaceutical company creating more accurate vaccine models, or a financial institution getting information quickly for decisions where mere minutes could cost thousands or millions of dollars. Even diaper manufacturers use supercomputers to test the absorbency of the products they bring to market. Supercomputers have been traditionally used in large-scale, academia research, but today's advancements in technology are bringing this power to the masses.

Clustering Has Its Advantages

A supercomputer is one that performs at the highest operational rate possible for computers; it's generally used for scientific or engineering applications that entail large databases and do a massive amount of computation. Although supercomputing's roots go back to the '70s, the landscape changed in the '80s as the industry began utilizing microprocessor-based systems. It was in the '90s, however, that supercomputing went through the third stage of its evolution: when we began taking industry standard building blocks—such as two-processor servers—and interconnecting them to build a supercomputer, commonly known as high performance computing clusters, or HPCCs.

HPCCs combine symmetric multi-processor (SMP) computer systems together with high-speed interconnects to achieve the same power as conventional supercomputers. On an enterprise level, leaders need to begin thinking about corporate grids and how they can interconnect the investments they have made in different sites.

HPCCs allow organizations to build their own supercomputers by interconnecting multiple machines. One of the advantages is the scalability and availability of the clusters. The clusters form what appears to be a single system, giving the information processing advantages of a supercomputer at a more manageable cost. Enterprises can buy only as much power as they need, rather than needing to over-purchase.

Organizations can build their supercomputing power as they go, as they need more performance, capacity, and processing power. Over time, the cost benefit has only improved as capabilities have increased while prices have decreased.>>

HPCC Advantages:

- You buy only as much computing power as you need at the moment.
- Adding power is cost-effective and can be done as need increases.
- New technology can be easily integrated into the existing system.
- The total cost of ownership of an HPCC vs. a supercomputer is much less, from acquisition to management.

HPCC Drawbacks:

- More applications are developed for big machines; therefore, it might be necessary to develop software and applications specific to your organization's needs.
- There are more management considerations, both in terms of people and equipment.


MDF WORKSTATION #2

“The economic advantages don’t end there. In fact, the acquisition cost of a standards-based cluster is one tenth the cost of a proprietary system.”

The concept of multiple systems with multiple processors—instead of a single, large system with multiple processors used by the more traditional supercomputer—lends itself to business continuity.

Many customers who use the HPCC have spare computers on hand in case something goes wrong with their cluster integration. This allows easy replacement of a system while keeping the cluster up and running. This an increasingly popular approach, and it’s one that saves companies time, a substantial amount of money on service calls, and maximizes uptime. Because the clusters are built using industry standard operating systems, such as Linux and Microsoft, finding administrators versed in the technology rarely presents a challenge.

The economic advantages don’t end there. In fact, the acquisition cost of a standards-based cluster is one tenth the cost of a proprietary system. Mainenance and mangement costs are also lower.

In addition, the clusters are able to integrate old and new technologies. Newer computers added to the cluster can interact with older technology, whereas a single supercomputer must have the old electronics removed and updated in order to remain on the cutting edge.

Mitigating Risk with Supercomputers

In nearly any industry, there are rich uses for HPCC applications. Supercomputing provides high-end research, development, design, production, and analytical capabilities with far-reaching applications.

These applications range from intensive scientific and engineering problems to commercial data-intensive tasks. Applications include seismic analysis for oil exploration, aerodynamic simulation for motor and aircraft design, molecular modeling for biomedical research, rendering for animation, data mining for finance, modeling for business analysis, and so much more.

In the oil and gas industry, HPCC-aware software is used to more quickly and accurately pinpoint the location of oil and gas reservoirs and can optimize production by applying enhanced oil recovery techniques. Having a good idea of where to find oil takes a significant amount of risk out of drilling a dry well.

Manufacturing uses the software to implement safe, reliable, and faster designs to create better products. Software also allows for crash simulations and analysis without jeopardizing human lives. In the financial world, it gives analysts an edge for predicting future market movements. It also has strong implications in weather forecasting because it’s possible to see the effect of certain weather patterns before a crisis occurs. But not all of the applications are geared toward technical computation. In business intelligence, for example, the HPCC software allows corporations to analyze data about customers and also use it for payroll and HR applications.

Next Steps

Those contemplating HPCC as a computing paradigm should build a scaled down HPCC for testing and development. For a test and development cluster, an eight-node bundle is typically sufficient. Operating an HPCC cluster will provide the requisite understanding of the hardware (servers, storage, and networking) and software components and how they interact. Some helpful considerations:

- Select a candidate application to run on the cluster. If necessary, port the application to run optimally on the HPCC cluster. Benchmark the application to get a feel for how it performs on the HPCC cluster versus its previous operating environment. Experiment with changing the layers of hardware or software to see how the performance changes.
- Balance the computing and I/O components of your clusters. Having a large number of >>

Case Study | Combating Crises at TACC

At the **Texas Advanced Computing Center** (TACC), a simulated crisis is always just around the corner. The center at **The University of Texas at Austin**, using Dell as a leading technology provider, offers advanced computing resources and services for computationally intensive research. With its LoneStar high-performance computing cluster, TACC has made great strides in simulated disaster recovery.

Gregory S. Johnson, who joined TACC in 2001, supports campus researchers in the planning and execution of simulated events. Using a high-end visualization machine, the center is able to view natural disasters simulated on an HPCC to provide on-the-job training to emergency management and response teams.

In those situations, emergency response teams log in to the visualization resource through a handheld PDA or a computer, pinpoint problem areas, and deploy resources accordingly. Johnson said the simulation software allows researchers and emergency response teams working with TACC to model a crisis, such as a hurricane coming ashore on Galveston Island (in Texas), and give response teams real-time lessons in reacting to the crisis. The response teams can go anywhere in the city with cellular coverage or wireless access and type in an address on their PDA to view a virtual overview of the simulated events in that particular area.

The simulation exercises provide a valuable tool to cities looking for a way to develop strategies for combating everything from natural disasters to acts of terrorism. But the model also provides business applications that will enable leaders to conduct similar exercises.

Just as emergency teams can log in and run through various exercises and scenarios, a business can create simulated disaster scenarios. Instead of emergency response personnel, the individuals logging in might be risk managers or production staff. Utilizing the simulation software, they could see what would happen to their business in the event of, say, a Category 5 hurricane, and would be able to analyze what steps could be taken to minimize damage to the business. Johnson said this is a valuable strategic tool for enterprises that want to determine whether or not they need to develop additional off-site data storage or see what, exactly, they need to prepare for in the event of certain calamities.

Although there are sizeable infrastructure requirements for an HPCC—heavy-duty electrical wiring, floor space, chilling components, and more—the trend is toward the establishment of large HPCCs that government institutions and businesses can access remotely. Johnson sees a continuing trend in the direction of having remote access to any data that's needed for disaster recovery and risk management. Such access provides businesses with autonomy in the face of disaster and ensures a quicker and more thorough recovery.

"Remote access to data like that—that's where the future lies," Johnson says. "Imagine making data associated with emergency response or disaster recovery just as easy to access and just as easy to use as Google maps. That's what we're looking at."

The technology presently is available primarily for disaster recovery emergency response teams and academic applications, but the business world isn't too far behind.

"Businesses are just now becoming aware of this technology," Johnson says, but that is about to change. "If they're not looking at this technology in the next one to two years, arguably, they could fall behind."

> “Availability, fault-tolerance, reliability, and serviceability are important design considerations...”

computational nodes where I/O bottlenecks exist doesn't offer a scalable HPCC cluster to an end-user. It's important that the I/O sub-system is proportionally designed for the computational side of the cluster.

■ Consider the power and cooling capacities of your data center. A properly structured data center allows business leaders to plan for future needs and allocate the space required to accommodate future growth.

■ Design reliability, availability, and serviceability (RAS) properties into your HPCC cluster. The lines between commercial and computational computing have blurred and increasingly customers are asking for the same RAS requirements to be present in their computational clusters. Availability, fault-tolerance, reliability, and serviceability are important design considerations and should be worked into the architecture of the layers, building blocks, and components.

Dell has partnered with the Texas Advanced Computing Center (TACC), a research center at The University of Texas at Austin, serving the local, state, and national communities in their research and development. LoneStar, the TACC HPCC, comprises more than 500 Dell servers, enabling everything from development and testing to benchmarking and scientific application simulations. LoneStar will be upgraded in September 2006 to more than 1,300 Dell PowerEdge 1955 blade servers. TACC serves as a model to the rest of our HPCC community, offering some of our best practices to the wider customer space. |s|

About the Author: **Dr. Reza Rooholamini** is director of solutions engineering at Dell Inc.