Stitching Computing Systems into a Single Fabric

Sixto Ortiz Jr.

Many computer-industry observers say that eventually, computing resources will be available in the same manner as public utilities such as electricity or water. They contend that resources from computers will be available as a pool of computing power not directly connected to any single machine.

Fabric computing, the next-generation architecture for enterprise servers, is one significant step in that direction.

The technology combines powerful server capabilities and advanced networking into a single fabric that users can work with as needed. In essence, it creates reconfigurable pools of computing, network, communications, and storage resources interconnected via a highly available, scalable fabric.

With the approach, resources are no longer inflexibly tied to a single server. Instead, users could work with a pool of resources they could reconfigure as necessary.

Proponents say the technology will be particularly valuable in data centers, high-performance computing, and sensor networks.

Until now, the concept has been implemented partially in technologies such as virtualization and grid computing, said Erick Von Schweber, CEO and chief scientist at consultancy NeoLogical.

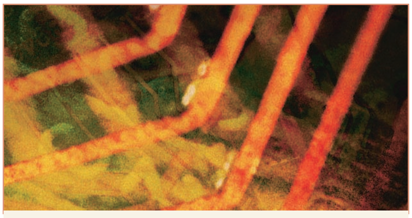

Virtualization creates multiple virtual machines out of a single computing resource. Grid computing uses the resources of geographically dispersed computers to form a virtual supercomputer to perform large tasks.

Fabric computing, on the other hand, lets organizations maximize the use of their local computing resources for tasks of any size, an important consideration for many companies. They could purchase fewer servers and less network equipment. And by letting administrators control resources via a single fabric, the technology provides easier, more centralized management.

Fabric-computing vendors and products include Liquid Computing and its LiquidIQ system; SGI and its Altix ICE integrated blade cluster; Appistry and its Enterprise Application Fabric; and System Fabric Works with its FabricStor and FabricMem storage appliances and its FabricRouter network appliance.

Large companies are also addressing aspects of fabric computing. IBM's BladeCenter Open Fabric links blade servers to create a single virtual resource. Hewlett-Packard's Adaptive Infrastructure initiative connects various data-center resources to form a centrally managed entity.

To succeed commercially, though, the technology must overcome challenges such as complexity.

WHY BOTHER?

Fabric computing has emerged in response to marketplace demands for more efficient ways to operate data centers, which frequently are large, complex, and expensive to run. Also, users are looking for ways to consolidate and better utilize their computing resources.

History

Erick and Linda Von Schweber coined the term "computing fabric" in a 1998 column for *PC Week* (now *eWeek*) magazine.

In essence, fabric computing partitions hardware to create a flexible and reconfigurable computing infrastructure, said Vishal Sharma, principal consultant and technologist at the Metanoia consultancy in Mountain View, California. The technology uses software to virtualize hardware components to appear as a single entity, he explained.

Driving forces

System administrators are interested in fabric computing because it simplifies their jobs in several key ways.

Modern server architectures, by and large, use variations of the basic client-server model. However, as business operations have become more demanding—due to factors such as the increased use of enterprise-resource-planning and other transaction-intensive applications—these architectures' complexity has grown.

Configuration complexity. When administrators purchase servers, they must configure them individually for the tasks they will perform, such as running a database application. Fabric computing delivers an architec-

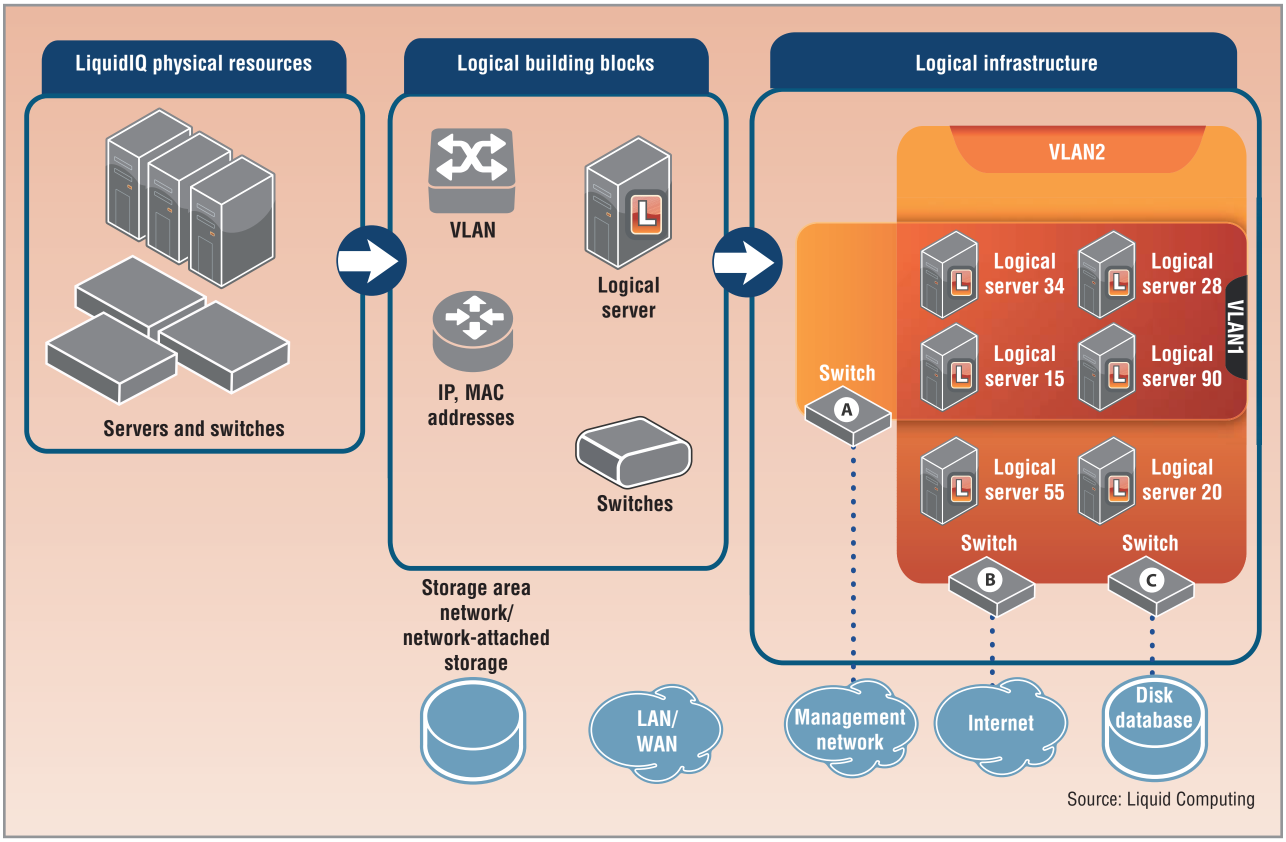

Figure 1. Fabric-computing systems, like Liquid Computing's LiquidIQ, combine standard physical data-center resources such as servers and switches with sophisticated management and control software. They let organizations fully use their existing resources by flexibly configuring and reconfiguring them to create logical elements—such as servers, networks, and storage—as needed, with the entire system operating over a single fabric.

ture in which individual servers are part of a large, single, logical pool of resources that can be provisioned and reconfigured as a whole to run the applications required.

This simplifies configuration and makes it easier to introduce new services.

Management complexity. Fabric computing's consolidation of the physical infrastructure into a simpler, centralized system helps ease management complexity, said Liquid Computing chief technology officer Michael Kemp. Administrators can focus on managing the entire fabric rather than on working with individual servers.

EXAMINING THE FABRIC

According to Von Schweber, the main characteristic that distinguishes a computing fabric is the merging of a hardware system and a network, which blurs the distinction between them.

Design principles

Fabric computing generally features several key design principles. For example, the software-based fabric layer organizes the servers that act as nodes into a fully distributed, peer-to-peer system. Each server runs an identical image of the operating system, applications, and other elements. Thus, each node can perform any task and there is no single point of failure.

A closer look

A computing fabric consists of servers that are linked via I/O cards, hubs, switches, and cabling and that use various networking approaches—including Ethernet, Fibre Channel, and InfiniBand—to communicate with one another, said Sam Charrington, Appistry's vice president of product management and marketing.

The fabric, which looks like a single server or resource to external clients and services, sits between the programs and the OS, he explained. It is thus virtualized across the underlying infrastructure, he said.

The fabric includes a layer that performs virtualization and includes the applications and software libraries necessary to run Java-based and other services.

The software also monitors throughput and dynamically balances workloads and processing tasks across nodes to optimize system performance. This reduces the need for IT workers to manually perform load balancing, which can require them to disrupt system operations.

Administrators can combine and recombine fabric components into various arrangements to suit their needs. They can utilize software to create logical, virtualized data-center-infrastructure elements, including servers, virtual LANs, subnets, storage gateways, and virtual private networks, as Figure 1 shows.

Management software configures system properties dynamically, even while the computing fabric is running applications, said Von Schweber.

"The key architectural element that enables this automation," said Liquid Computing's Kemp, "is a common data model of all the resources." This provides a common way to define all physical resources and the relationships between them, enabling the fabric-computing system to abstract the resources easily so that it can manipulate them.

An organization can run all of its security applications within the fabric.

Deployment

Currently, fabric computing is used primarily in high-performance-related applications such as image processing, financial simulations, and transportation routing. According to Metanoia's Sharma, the popularity of these applications bodes well for fabric computing's future.

Data centers are another important setting for fabric computing.

A potential future use for fabric computing is with sensor systems such as those used in large environmental-monitoring applications. According to Von Schweber, when sensors are widely distributed in a massively parallel fashion and collect huge amounts of data, the sensor network must process the information and route it to storage, a database, or elsewhere as needed. This type of network can thus benefit from fabric computing's flexibility and agility.

Advantages

Fabric computing reduces the complexity of managing and administering computing resources as discrete units. Instead, organizations can administer the fabric and its applications as a single entity and not have to deal with individual machines. Also, the overlying fabric-software control layer automates typically manual tasks such as reconfiguration and load balancing.

Moreover, it is easier for developers to program applications that use resources that have been unified into a single fabric.

Having no single point of failure improves system reliability. Load distribution increases performance. And the common pool or memory that fabric computing creates eliminates memory-shortage problems.

With fabric computing, Von Schweber said, "these are native properties. You don't have to do anything special. You just treat it as a system with whatever properties it's been configured to offer."

Because fabric computing lets organizations work with fewer servers and less network equipment and also automates many administrative and management functions, it improves resource utilization, reduces operating costs, and lowers energy consumption, said Liquid Computing's Kemp.

In addition, Appistry's Charrington noted, load balancing, along with fabric computing's self-organizing capabilities, enable easier scaling.

The approach is also efficient because applications are able to identify and then use only the resources they need from the fabric, noted David Bader, a professor and executive director of high-performance computing at the Georgia Institute of Technology.

Fabric computing is new, so some companies might be reluctant to implement the technology until it has been around for a while.

Because it is a recent approach to operating data centers, its initial implementation could be time-consuming and perhaps confusing to organizations.

Metanoia's Sharma said fabric computing will offer the most benefits to organizations with larger systems. Therefore, only bigger organizations might find it worthwhile to spend the time and money currently necessary to implement the technology, a potentially complex process.

Nonetheless, said Von Schweber, its adoption will continue expanding as vendors design processors with more computing cores. With fabric computing, he explained, organizations could get more out of the additional resources that multicore chips provide.

According to Bader, computing is undergoing changes—such as multicore processors and new server architectures—that enable organizations to rethink the way they implement technology. Fabric computing is one of these new approaches, but only time will tell whether it succeeds.

Sharma stated that he expects fabric-computing adoption to grow rapidly over the next three to five years. Technical advances will make the technology more accessible and attractive to IT departments, he explained.

However, Von Schweber said, if organizations apply fabric computing narrowly—such as in only certain parts of their infrastructure, like storage—or don't use the approach where it could yield the greatest benefit—such as with data centers, grids, and multicore-chip environments—they will not realize the technology's full benefits. This, he said, could limit fabric computing's adoption and hurt its success.

Sixto Ortiz Jr. is a freelance technology writer based in Spring, Texas. Contact him at sortiz1965@gmail.com.

Editor: Lee Garber, *Computer*; l.garber@computer.org