Initiatives and Technology

The Changing Nature of Data Center I/O

Contents

Overview

With the explosive growth in the scope and functionality of the enterprise data center, IT executives and senior architects are continually seeking new models for maintaining scalability, boosting performance, and reducing the total cost of ownership. One such model, known as the "virtual server," has the potential to deliver all three benefits. Although it is still only a future vision, groundwork for the virtual server is being laid today through the convergence of enhanced data center networks and the modular computing infrastructure. In this article, IT executives and senior architects will learn about these phenomena and their role in the changing nature of data center I/O.

Enhanced data center networks increase bandwidth, reduce latency

New forms of traditional networking for data center applications are emerging. These enhanced data center networks depend on two advancements in networking technology; first, they take full advantage of hardware acceleration by moving the traditional networking stack from software into hardware. Second, they define the notion of a switched fabric to include not only the first four layers of a traditional networking stack, (Physical, Networking, Link and Transport layers), but also to include substantial enhancements to the upper layers of the stack thus giving the data center application much more direct access to the underlying data center network as well as new features designed to significantly speed application performance. Two leading examples are the InfiniBand* Architecture (IBA) and RDMA-enabled IP networks. IBA is a channel-based, switched-fabric I/O architecture implementing a comprehensive stack from Layer 1 through Layer 4 exclusively in hardware. It also includes an RDMA capability as part of its transport layer, a comprehensive fabric management scheme, and a complete connection management protocol. An RDMA-enabled IP network is an adaptation of existing Ethernet technology (traditionally implemented in software) that is implemented in hardware with the help of a related technology known as a TCP/IP Offload Engine (TOE). An RDMA-enabled IP network, as its name implies, also takes advantage of emerging RDMA technology to speed application performance. By increasing effective bandwidth and reducing latency, IBA and RDMA-enabled IP networks can significantly reduce the performance distinction between local I/O and remote I/O, thereby supporting and complementing the emergence of modular computing.

Modular computing infrastructure enables disaggregation

The first implementations of modular computing to emerge in the data center have been based on "bladed servers" essentially re-packaged rack-mount servers connected to a backplane. These bladed servers consist of the familiar combination of a compute complex and an I/O subsystem. The I/O subsystem includes everything necessary for an application to access I/O resources including drivers, network and storage software stacks, host bus adapters, network connections, and even I/O devices themselves in the case of direct-attached storage. Typically, an I/O subsystem includes a number of "thin pipes" each of which is dedicated to accessing a particular type of I/O. Now, thanks to the performance characteristics of new enhanced data center networks such as IBA or RDMA-enabled IP networks, applications can access virtually any type of I/O (including network I/O, storage I/O and interprocess communication) over one data center fabric.The combination of high bandwidth and low latency provided by an enhanced data center network means not only that storage resources can be physically distant from the server (while remaining within the confines of the data center), it also means that the CPU/memory complex running a particular application can rely on a single I/O connection into a fabric of I/O resources, essentially replacing the traditional model of multiple "thin pipes" to various I/O services with one "fat pipe" into an enhanced data center fabric. No longer is it as necessary to partition I/O resources into local and remote I/O based on application requirements and differences in performance. One key result of this is that servers no longer need to be individually optimized for the unique I/O requirements of a given application; a single fat pipe is sufficient to support the I/O needs of any application running on the server. This means that any application can access any resource, including additional compute resources, over a single fabric no matter where the resource resides in the data center. The result of this is a generalization (virtualization) of not only storage resources, but all data center resources including compute resources. It also begs the question: Why not virtualize the server as well?

Virtual server model can vastly optimize resource allocation

A "virtual server" is a model where a given application is no longer required to be mapped to a specific server, where "server" is defined to be the combination of compute resources and an I/O subsystem designed to support the application. With the addition of new management software (still in the early stages of development), the virtual server model will no longer require that an application be mapped to a server (or cluster of servers). Instead, virtual servers consisting of varying amounts of compute resources, storage resources, and client network connections are assigned to support each application, or instance of an application. The applications can be either traditional (single-server) applications, each instance of which is assigned to a separate compute resource, or a cluster-aware application. The virtual server model makes it possible for an application to consume compute resources located anywhere in the data center. The logical outcome is that the data center will be viewed as consisting of a pool of compute resources, a pool of storage resources, and a pool of network connections all designed to act as a flexible platform for supporting data center applications. Depending on the time of day or the type of enterprise, resources may be shifted over time among the various applications running in the data center as a function of current loading conditions. For example, more compute resource may be made available to support front end webserving applications at specific times of the day. In the virtual server model data center resources, including compute, storage and client network resources, can be transparently shifted among the various data center applications as needed. This also means that enterprise data centers can allocate server resources optimally with very little human intervention.

Summary

Intel has long played a leading role in the development of advanced data center technologies, including its active participation in the InfiniBand Trade Association and the RDMA Consortium. Intel also supports the industry trend toward a modular computing infrastructure for the data center. Intel believes that the convergence of these phenomena enhanced data center network technologies such as the InfiniBand Architecture and RDMA-enabled IP networks, the widespread adoption of a modular computing infrastructure, and the promise of management software enabling server virtualization will drive enterprise data centers ever closer to their goals in terms of scalability, performance, and total cost of ownership.

Author Biography

Paul Grun is manager, Enterprise I/O and Initiatives, Platform Products Architecture/Enterprise Platform Group at Intel Corporation. In his six years with the company he has won an Intel Achievement award and an EPG Group award, and has been awarded four patents. He holds a B.S. in electrical engineering and a B.A. in economics from Syracuse University.

For more information, visit the Intel web site at: developer.intel.com

For more information on the changing nature of data center I/O and Intel's role in the development of a new enterprise data center model, plan to attend the Enterprise I/O Architecture and Design track and visit the Intel booth at the Spring Intel Developer Forum (IDF), whose U.S. conference will take place at San Jose, California, February 18–21, 2003. For more information on IDF conferences outside the U.S. and on IDF in general, visit http://www.intel.com/idf/.