SYSTEM FOR ALLOCATING COMPUTING RESOURCES OF DISTRIBUTED COMPUTER SYSTEM WITH TRANSACTION MANAGER

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a divisional of application Ser. No. 10/176,436, filed Jun. 20, 2002, issued as U.S. Pat. No. 7,376,693 on May 20, 2008, which claims the benefit of U.S. Provisional Application No. 60/355,274, filed Feb. 8, 2002, each of which are incorporated by reference herein.

BACKGROUND

I. Field of the Invention

The present invention relates to the structure and operation of computing systems, and more particularly, to distributed computing systems and methods of operating such systems.

II. Description of the Related Art

Certain organizations have a need for high performance computing resources. For example, a financial institution may use such resources to perform risk management modeling of the valuations for particular instruments and portfolios at specified states of the world. As another example, a pharmaceutical manufacturer may use high performance computing resources to model the effects, efficacy and/or interactions of new drugs it is developing. As a further example, an oil exploration company may evaluate seismic information using high performance computing resources.

One conventional computing system includes a mainframe computer attached to an individual user terminal by a network connection. Using the terminal, a user may instruct the mainframe computer to execute a command. In this conventional system, almost all data storage and processing functionality resides on the mainframe computer, while relatively little memory or processing capability exists at the terminal. This terminal/mainframe architecture may not, however, allow computations requested by a user to be computed rapidly or automatically.

The open systems interconnection (OSI) model describes one conceptual network architecture represented by seven functional layers. In this model, the functions of a networking system in a data communications network are reflected as a set of seven layers, including a physical layer, data link layer, network layer, transport layer, session layer, presentation layer and application layer. One or more entities within each layer implement the functionality of the layer. Each entity provides facilities for use only by the layer above it, and interacts directly only with the layer below it. FIG. **1** depicts the seven functional layers of the OSI model.

The physical layer describes the physical characteristics of hardware components used to form a network. For example, the size of cable, the type of connector, and the method of termination are defined in the physical layer.

The data link layer describes the organization of the data to be transmitted over the particular mechanical/electrical/optical devices described in the physical layer. For example, the framing, addressing and check summing of Ethernet packets is defined in the data link layer.

The network layer describes how data is physically routed and exchanged along a path for delivery from one node of a network to another. For example, the addressing and routing structure of the network is defined in this layer.

The transport layer describes means used to ensure that data is delivered from place to place in a sequential, error-free, and robust (i.e., no losses or duplications) condition. The complexity of the transport protocol is defined by the transport layer.

The session layer involves the organization of data generated by processes running on multiple nodes of a network in order to establish, use and terminate a connection between those nodes. For example, the session layer describes how security, name recognition and logging functions are to take place to allow a connection to be established, used and terminated.

The presentation layer describes the format the data presented to the application layer must possess. This layer translates data from the format it possesses at the sending/receiving station of the network node to the format it must embody to be used by the application layer.

The application layer describes the service made available to the user of the network node in order to perform a particular function the user wants to have performed. For example, the application layer implements electronic messaging (such as "e-mail") or remote file access.

In certain conventional high performance computing systems designed using the OSI model, the hardware used for computation-intensive processing may be dedicated to only one long-running program and, accordingly, may not be accessible by other long running programs. Moreover, it may be difficult to easily and dynamically reallocate the computation-intensive processing from one long running program to another. In the event processing resources are to be reallocated, a program currently running on a conventional high performance computer system typically must be terminated and re-run in its entirety at a later time.

SUMMARY OF THE INVENTION

In one aspect, the invention features a system including a transaction manager in communication with a plurality of local computers, wherein the transaction manager supports a plurality of message protocols to enable communication between the transaction manager and the plurality of local computers, wherein each local computer runs a calling application. In such an aspect, the system also includes a queue in communication with the transaction manager, wherein the queue receives and stores a plurality of jobs and task inputs from the transaction manager and sends a plurality of computation requests to one or more of a plurality of node computers. In this aspect, the system further includes a scheduler in communication with the queue, wherein the scheduler routes a plurality of incoming tasks to a plurality of compute functions on the plurality of node computers and allocates computing resources of the plurality of node computers; a service manager in communication with the transaction manager, the scheduler, and the plurality of node computers, wherein the service manager controls allocation of computing resources among a plurality of users; and a cache in communication with said plurality of node computers.

In another aspect, the invention features a system including a local computing device in communication, over a network, with a distributed computing system, the local computing device configured to perform computations for a first portion of a computer software application and to send a second portion of the application for computation on the distributed computing system. In such an aspect, the distributed computing system includes means for performing computations for the second portion of the application, means for queuing for computation the second portion of the application, wherein the means for queuing is in communication with the means for performing computations, and wherein a minimum avail-