GEMFIRE
THE ENTERPRISE DATA FABRIC

GemStone and High Performance Computing

GemFire - The Enterprise Data Fabric (EDF) - provides a scalable, distributed platform to manage increasing volumes of enterprise data and streaming events with almost zero-latency. With advanced data virtualization, distributed caching and Complex Event Processing (CEP) capabilities, the GemFire solution suite enables the delivery of actionable information to the right application at the right time. As organizations aim to streamline process efficiencies, enhance customer touch-points, and embark on initiatives like SOA, Web portals and Grids, such a fabric soon becomes a critical requirement for their technology environment.

Audience:
Technical Project Managers
Systems Architects
Software Engineers

CHALLENGES FOR HIGH PERFORMANCE COMPUTING

Major challenges for multiprocessor high performance computing (HPC) are:

- Data access times
- Cross process coordination
- Data sharing across computational nodes especially in a multi-computer environment
- Working with large data sets in a large computer cluster

There are two key areas for high performance computing

1. Data processing of incoming data streams
2. Multi-pass processing for a single data problem: Monte Carlo, Simulations & Modeling, Complex Analysis

Data processing for incoming streams requires keeping up with the average data flows and providing capabilities to manage the peaks. The multi-pass processing on a single problem is often bounded by a time window and the number of such problems that can be worked on. If the amount of time needed to run the simulation or analysis can be reduced, a greater number of problems can be worked on or more analysis can be done on a single problem. In the intelligence community there are often an arbitrarily large number of data sets to analyze and a large but limited amount of resources to process these data sets. As a result a significant amount of data is never analyzed or run through complex computer analysis.

THE KEY CHALLENGES TO MAXIMIZING PERFORMANCE ARE AS FOLLOWS:

Data access time
One of the most significant challenges for high performance computing comes from the significant increase in CPU power. Perversely

GEMSTONE®

the increase in CPU power has increased the significance of I/O becoming the bottleneck for CPU intensive systems. The lack of an application controlled management of I/O access and memory can greatly impact performance characteristics of an application(s). Many HPC systems need to fetch additional data such as, reference data, to support computation. An example from the financial services industry is the computation of risk positions. Additionally, many of these systems need to write out midstream calculation to allow processes to share data with other processes. Multiple processes may be working on a single data set and need to update either the data set or the state of the work on it based on current calculations. An example of this in the financial industry is from the area of statistical calculations of risk positions. According to a Risk Systems Analyst & Architect at a large financial institution, *"A major issue in computation of risk positions lies in data access. Data access can account for anywhere from 30%-90% of the total computation envelope."*

Given that many intelligence organizations have more data than they can process or more calculations that they would like to run on the same data in a given time window, spending more time computing and less time performing I/O can result in a substantial increase in the amount of data processed.

Cross process coordination

Cross process coordination and communications can be a complex issue for many high performance computing systems. When applications are run across a large number of CPU's, coordinating the activities of the individual processes and their current state in the process flow, can be complex. To avoid complexity and concurrency issues many applications write the data out to databases creating the data access bottlenecks discussed in the previous section. In multi-process computing systems other means of cross process communication across CPUs can provide greater speed than a database but each comes at the cost of some additional complexity and makes the applications significantly harder to modify. Multi-threaded(?) programs and socket based communication may be used if one program is aware of where all of the processes that need the information are and if those processes are running at that instant in time. Unfortunately, this can result in hardwiring a lot of system configuration information into the application. In SMP systems, high performance cross process communications can be accomplished by writing data to shared memory so that any other processes can then pick up that data at their leisure without the writing process needing to be made aware of who needs the data. There are limitations in using SMP only solutions to address cross process coordination. Managing that section of shared memory may become an issue as data that is not in use any more must be eliminated and the memory segment may become full. SMP systems are limited in the number of CPU's they hold so there is a limit to the number of CPU's that can be brought to bear on a problem. On a 32 bit system total memory available is limited to 4GB. Many organizations want to leverage the new low cost blade servers as opposed to large more expensive SMP systems. As a result of these issues it is not uncommon to use an RDBMS to server as the mechanism for information sharing in these instances.

Data sharing across computational nodes especially in a multi-computer environment

Since many organizations are interested in the new low cost cluster or GRID environment one of the key challenges is how to share data across the cluster or GRID without resorting to disk IO. As mentioned above one of the key ways to do this in the past on SMP systems was to use shared memory. Unfortunately, shared memory does not work across machines, so is unavailable in a cluster or GRID environment.

Working with large data sets in a cluster

The other key challenge in a cluster environment is how to access large amounts of memory at an affordable cost. With 32 byte machines the maximum amount of available memory is limited to 4 GB. 64 byte machines enable applications to break through that barrier. There is still a practical limit to how much memory each 64 byte machine in a GRID or cluster can have on its node. This is dependent on the physical limits of the hardware, as well as cost constraints, when you start on memory amount >10 GB. If there were some way to aggregate the memory across the cluster, problems requiring larger data sets could be addressed. For example, one US government laboratory has a lab with a 2,048 node Linux cluster. Each node has 1 GB; so the total system has 2 Terabytes. However, no single computing node has more than 1 GB of memory to work on a problem with.

SOLUTIONS OR HOW GEMFIRE ENTERPRISE CAN HELP

GemFire Enterprise, a data services solution from Gemstone Systems, provides a flexible, high performance, cross-platform, cost-effective solution for data distribution, distributed caching, and management across the enterprise in real-time. It enables computers to create virtual shared memory space. GemFire Enterprise provides a fault tolerant environment to avoid single computer system failures from impacting the overall system.

Data access time

GemFire Enterprise can significantly reduce data access times. Where multiple computations in different processes need the same data, GemFire Enterprise can buffer reads from data sources such as RDBMSs. As a result, data only needs to be fetched once no matter how many different computers need the data. If data needs to be updated, it can be pushed back into GemFire Enterprise to avoid the wait times associated with writing back to a database or other system. If necessary, GemFire Enterprise can provide a write-behind cache to store the data back to a persistent source. By providing applications a tool to manage shared data in memory across machines, GemFire Enterprise can allow systems to significantly increase the amount of calculations performed in a given time. An example of this in the financial industry is from the area of statistical calculations of risk positions. According to a Risk Systems Analyst & Architect at a large financial institution, *"By leveraging a solution such as GemFire Enterprise, we can see a dramatic increase in number of computations that we can perform within a specific batch processing period, or even better yet, provide results to calculations on an on-demand basis."* Thus, GemFire Enterprise can enable organizations to increase in the amount of data that can be handled or can perform additional calculations on the same data in a given time window.

Cross process coordination

GemFire Enterprise can enable applications to coordinate their activities by using a virtual shared memory environment as a white board to write out information. The data space can be organized into multiple regions and subregions to provide compartments for different processes or parts of processes to work. Alternatively, these regions can be used to organize data by type or any other desired organization. Processes can register with GemFire Enterprise for notification of data entries or changes to data on a region by region basis. The registration and notification process is a high-speed mechanism to allow applications to signal each other without necessarily being aware of what computers or CPU's they are running on. Thus GemFire Enterprise can enable the performance and ease of coordination of a SMP computer, without requiring the application be run on an SMP system. As an additional benefit, GemFire Enterprise will manage the memory space eliminating the issues of data management, garbage collection, and memory overflows.

Data sharing across computational nodes especially in a multi-computer environment

With GemFire Enterprise, data sharing across processes, CPU's and computers is simple. GemFire Enterprise provides not only data sharing across processes with its shared memory space, but also powerful controls on data distribution and concurrency. GemFire Enterprise is ideally suited to exploit the data layer in a cluster or GRID environment where there is a high speed network interconnected? between machines.

Working with large data sets in a cluster

Currently, GemFire Enterprise enables 32 bit applications on 32 bit machines to share data with a 64 bit memory space running GemFire Enterprise on a 64 bit machine. With GemFire Enterprise's memory management capabilities, this can enable 32 bit applications to work on large data sets, if they can break them into pieces. GemFire Enterprise also provides a Distributed Hash Map capability so that a group of machines can be turned into a unified memory space. This occurs when the application stores the data in and retrieves from GemFire Enterprise's Distributed Hash Map data structure. Also, this enables a group of computers to combine their memory.

SUMMARY

GemFire Enterprise can help high performance computing systems reduce the amount of time spent performing data access thus increasing the amount of time available for high performance calculations. This resultant shift increases overall application performance by a significant factor. GemFire Enterprise can also help systems leverage large groups of computers to work on a single problem by providing a distributed memory space wherein processes can exchange and organize data. Due to the built in capability to manage itself, GemFire Enterprise does not create additional operational overhead and complexity.

For more information on GemFire call:

Steve Rixse, Director of Sales and Business Development, Federal Sector
phone: 301.664.8494
Steve.rixse@gemstone.com

or

Michael Lazar, Director of Technology, Federal Sector
phone: 301.664.8490
michael.lazar@gemstone.com

Corporate Headquarters:
1260 NW Waterhouse Ave., Suite 200 Beaverton, OR 97006 | Phone: 503.533.3000 | Fax: 503.629.8556 | info@gemstone.com | www.gemstone.com

Regional Sales Offices:
New York | 90 Park Avenue 17th Floor New York, NY 10016 | Phone: 212.786.7328
Washington D.C. | 3 Bethesda Metro Center Suite 778 Bethesda, MD 20814 | Phone: 301.664.8494
Santa Clara | 2880 Lakeside Drive Suite 331 Santa Clara, CA 95054 | Phone: 408.496.0242

GEMSTONE®