

Study of Next-Generation Infrastructure: InfiniBand HPC Grid Computing for Telecommunications Data Center

**** Prof.Mrs.Gyankamal Chhajed , ** Prof.Mr.Santosh Shinde**

*** Mr.Rahul Jagtap, * Mr.Venkatesh Raskotwar**

*** Student of Computer Department, ** Professor of Computer department**

Vidya Pratishthan's College of Engineering,Baramati

Pune University,Maharashtra,India

gjchhajed@gmail.com , meetsan_shinde@yahoo.com , jagtap.rahul712@gmail.com , raskotwar@gmail.com ,

***Abstract*—Grid computing is increasingly being viewed as the next phase of distributed computing. Built on pervasive Internet standards, grid computing enables organizations to share computing and information resources across department and organizational boundaries in a secure, highly efficient manner. Grid computing originates in e-Science and its early development was driven to a large extent by the requirements of large-scale computing and efficient sharing of huge datasets. e-Business requirements led to the adoption of emerging Web services technologies-initially developed for distributed business application integration. Therefore grid computing can be applied to enterprise computing within and across organizations and pave the way for utility computing. This paper propose an all InfiniBand grid computing proof of concept (PoC) to deploy “Database-as-a-Service” for mission-critical Oracle deployments in high-end transactional situations. This shows that performance of grid architecture—leveraging blade servers as well as InfiniBand networking and Storage is capable of achieving results equivalent to traditional, high-end SMP systems at significantly lower cost and power consumption levels.**

Keywords- InfiniBand; Grid Architecture; OGSa

I. INTRODUCTION

Grid computing evolves from a niche technology associated with scientific and technical computing, into a business-innovating technology that is driving increased commercial adoption. Grid deployments accelerate application performance, improve Productivity and collaboration, and optimize the resiliency of the IT infrastructure. By accelerating application performance, companies can more quickly deliver business results; achieving greater productivity, faster time to market, and increased customer satisfaction. Grid technology also provides the ability to store, share and analyze large volumes of data, ensuring that people have access to information at the right time, which can improve decision making, employee productivity and collaboration. Grid technology improves resource utilization and reduces costs, while maintaining a flexible infrastructure that can cope with changing business demands, yet remain reliable, resilient and secure. In a more formal and complete definition, grid computing can be seen as a distributed computing model that supports the concept of virtual dynamic organizations by providing secure and coordinated access and sharing of heterogeneous and geographically distributed resources, such

as applications, data, processor power, network bandwidth, storage capacity and others, over a network and across organizational boundaries, using a set of open standards and protocols.

Organizations rely heavily on their data centers to run mission-critical applications and deliver sustainable services to both internal and external users. With constant demands for better performance and high reliability at lower prices, innovative solutions that break the current model of data center operations are imperative to gain competitive advantages. The company realized that its traditional SMP infrastructure and enterprise-class Fibre Channel storage area networks (SAN) in operation for nearly a decade—were insufficient to accommodate future demands. Operating close to 30,000 servers worldwide, with virtually half of the deployments on SMP systems, the company faced increased operating costs and diminishing benefits due to scalability and performance limitations. Projections indicated that doubling the investment in their current SMP systems would only yield 70 percent additional application performance, thereby proving to be an inefficient and unacceptable solution. Additionally, the company wanted to improve energy efficiency to offset rising energy costs and support environmental (“green IT”) initiatives, requiring more efficient systems be explored..

II. EVOLUTION OF GRID COMPUTING

The phrase grid computing implies different technologies, markets and solutions to different people. Early on, much of the available literature focused on the compute intensive problems made tractable by grid, often associating it with cycle-scavenging or job scheduling technologies. Although these are important and useful components of a grid, they do not by themselves deliver the complete grid vision [11]. Grid computing is not about a specific hardware platform, a database or a particular piece of job management software, but the way in which IT resources dynamically interact to address changing business requirements. The grid domain has developed over a relatively short time period, fueled by significant technology advancements. Many grid technology roots can be traced back to the late 1980s in areas related to distribute supercomputing for numerically intensive applications, with particular emphasis on scheduling algorithms (e.g. Condor1, Load Sharing Facility2). By the late 1990s, a more generalized framework for accessing high-performance computing systems and distributed data

(e.g. Globus3) began to emerge, and then, at the turn of the millennium, the pace of change quickened with the recognition of the potential synergies between grid and the emerging Service Oriented Architectures (in particular through the creation of the Open Grid Services Architecture – OGSA) [8].

III. GRID SPECTRUM

Grid is best understood as part of a continuum representing virtualization technologies and solutions in following figure




Figure 1. Grid spectrum

At their simplest form, grid technology and solutions can be deployed within a homogeneous environment or within a single organization in a tightly coupled manner. Virtualizing “like” resources often entails deploying grid and virtualization functionality on cluster environments or multiple single systems with partitioning. As discussed earlier, a common entry point for grid deployments is a single line of business or single department using grids or clusters to increase business value; application acceleration,

Meeting service level agreements, or gaining insights to data. Such adoption has been fueled by the ongoing proliferation of clusters within enterprises, which often allows for a smoother transition using parallel applications. Another driver in this space is the continued acceptance of Linux and open source for multiple applications and workloads within enterprises. Indeed, many grid implementations in this space do rely on Linux and open source software functionality [1].

The next level of virtualization extends the grid concept – still within the domain of a single department or organization – to “unlike” resources. As a matter of fact, most departments often run applications and processes that are composed of unlike resources, whether these are servers, storage or other operating systems and software. The same principles of seamless integration using grid technology still apply, despite the loss of homogeneity. Where we “cross the chasm” is when we move from Virtualizing unlike resources within a specific organization or line of business to Virtualizing across multiple lines of business (or departments) within an enterprise. That is the point where (real and perceived) issues regarding security, ownership, and governance surface. Grid is a powerful technology that challenges long held assumptions over ownership, access and usage. Clients are increasingly confronted with the prospect of allowing their databases, their storage devices, their system processors to be leveraged and used by others who did not pay for those assets.

IV. GRID COMPUTING MODEL

There are also some fundamental grid models based on the type of basic services provided. Resources can basically be computing power, provided by servers or individual computers, data storage capacity, provided by information and data repositories, or network bandwidth, provided by networked infrastructures.

Computational grid: A computational grid is an infrastructure that allows resources to donate computing power to the grid whenever the workload demands [3]. This infrastructure is suitable for applications that demand. Such a grid takes advantage of the idle resources in the virtual organization. Just think about a cycle of CPU that are unused when a typical user browses on the Internet, reads an e-mail, or creates office documents like presentations or word processing files [4]. Examples of server oriented grids, desktop-oriented grids.

Data grid: Within a data grid infrastructure are the components used to provide grid capabilities to the data and information virtualization disciplines. It provides the ability to supply homogeneous access to heterogeneous repositories of data. It allows the data consumers to see an unified image of the respective information or data spread across different resources, potentially based on different technologies.

Network grid: In a typical corporate network, computers are very often permanently connected to it while using only a portion of its bandwidth. Every machine, servers and desktops, has underused network bandwidth, which can be considered as an idle resource. When a given user or machine requires more resources from the network, a bottleneck is reached.

Multipurpose grid: The multipurpose grid is perhaps the more common implementation in the of grid computing. The infrastructure of this grid should be adaptive enough to provide any of the grid models. It could be implemented as well as a meta-grid with abilities to route the requests to the grid that supports the right model to fulfill them.

V. GRID AND SERVICE ORIENTED ARCHITECTURE

Customers are implementing Service Oriented architectures (SOAs) to reduce complexity, to enable flexibility, and to streamline business processes. A SOA is a framework that lets customers build, deploy and integrate services for IT resources, applications and business process flows. It facilitates integration, offers modularization of applications and provides a coherent view of a business process as a set of coordinated services [7]. Therefore, customers can build and integrate applications and business processes across their enterprise, as well as with partners, suppliers and customers, in a much more on demand fashion. In order to realize these benefits, companies need to establish a flexible infrastructure capable of supporting dynamic operations with flexibility inherent throughout. As companies continue to implement SOAs throughout their enterprise, there are going to be a lot of services, some very fine grained and some larger, that will require a very responsive, dynamic and scalable infrastructure [5]. These services can be mobile. They need to be resilient. And the

whole process needs to be accomplished in a simple manner. The underlying infrastructure needs to be adaptable and autonomic in providing the right resources to the right application services and business policies to meet service level agreements and business performance needs. The necessary infrastructure consists of our leading capabilities across system and resource virtualization, and grid computing. IBM views grid computing as critical to the ongoing development of a dynamic and flexible infrastructure that enables SOA: Just as an SOA allows customers to separate applications from services, grids allow customers to separate both applications and services from the infrastructure and systems resources. Scheduling and workload management are key capabilities here supporting the placement and mobility of services and composite applications to the appropriate resources [2].

Grids provide an underlying foundation to support the dynamic nature of SOA. • Companies can pool resources for services, improve availability and reliability, and rapidly deploy and scale this new class of composite applications [6]. With a virtualized infrastructure, customers can much more easily scale their support for SOAs. They can harness all of their resources to accelerate time to results and to better align their infrastructure performance to their business goals. One of this will be possible at the SOA level without a "commitment to openness". It is the only way to link virtualization and grids with SOAs in a consistent and uniform manner. The infrastructure has to thoroughly support the ability to be managed in real-time, with sophisticated monitoring capabilities. Finally, the SOA and virtualization landscapes require a significant ecosystem of partners and capabilities. We need to see cross industry collaboration to drive the innovation required for SOA. In simple terms, grid computing delivers the underlying dynamic infrastructure that enables competitive advantage for customers implementing SOA-based solutions [2].

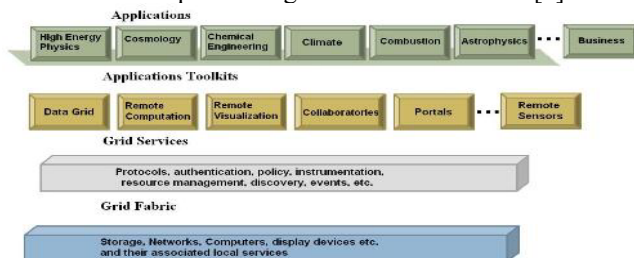


Figure 2. Grid Architecture

VI. OPEN GRID SERVICES ARCHITECTURE(OGSA)

OGSA was designed as an attempt to standardize interfaces and behaviors for distributed systems by:

- 1) Refactoring the Globus protocol suite to enable a common base and expose key capabilities. This was achieved by generalizing and redesigning protocols and mechanisms already identified as key for grid technologies, such as authentication, policy and lifetime management, notification, service naming, and reliable invocation [2].

- 2) Extending new technical requirements, such as concurrent access to resources and management tools.

- 3) Introducing the notion of "Grid Service". The goal here was to define grid entities by their interface and behavior, to unify access to programs and resources, and make it easier for grid architects to achieve virtualization model [11].

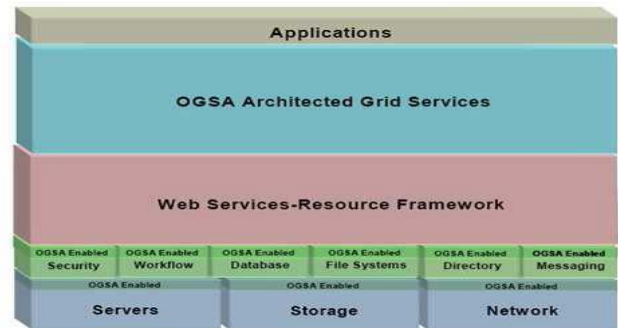


Figure 3. OGSA Architecture model

VII. INFINIBAND BASICS

InfiniBand is a point-to-point, switched I/O fabric architecture. Each end point, or *node*, can vary from an inexpensive single SCSI chip or Ethernet adapter to complex host systems. *Point-to-point* means that each communication link extends between only two devices. Both devices at each end of a link have full and exclusive access to the communication path. To go beyond a point and traverse the network, switches come into play. By adding switches, multiple points can be interconnected to create a fabric. As more switches are added to a network, aggregated bandwidth of the fabric increases. By adding multiple paths between devices, switches also provide a greater level of redundancy. InfiniBand network technology maximizes the

Performance of the overall solution due to its very high bandwidth, low latency features and ability to support the quality of service and high availability requirements associated with the company's critical business applications. InfiniBand has proven the ideal interconnect to scale out industry standard servers to tackle business computing problems that demand high levels of performance, reliability, and cost-effectiveness [10].

InfiniBand is growing rapidly in the beach head markets of high performance computing and performance business computing in the data center. All of the major server vendors are offering InfiniBand solutions to the market and there is rapid adoption and deployment of the technology across a broad range of applications.




Figure 4. InfiniBand vs. 2 GB/s Fibre Channel Single Port Storage Throughput to Disk Media

This success in turn is driving demand for InfiniBand storage. As an industry standard delivering 10Gb/s performance and transport offload, InfiniBand offers significant price performance advantages over Fibre Channel. These price/performance advantages are magnified in a clustered environment where the InfiniBand cluster network can be utilized for storage as well. Major advancements in InfiniBand software, platform, management, and application support for storage has accelerated the adoption of InfiniBand in the storage market. The demand for native attached InfiniBand storage is being further driven by applications demanding higher performance, lower costs, and better integration. Indeed servers and storage go hand in hand - and with the demand for native InfiniBand storage growing rapidly - InfiniBand is leading the way to a converged data center [10].

VIII. HIGH PERFORMANCE COMPUTING (HPC)

Grids today appear mostly in High Performance Computing (HPC) environments, where hundreds of inexpensive compute nodes link together to solve parallel processing problems that previously could only have been addressed by using supercomputers. Storage grids were developed to supply the enormous amounts of data that such processing power was capable of producing and consuming. Properly managed, storage and throughput scale linearly with computing power. Grids are a proven technology when it comes to HPC, but such sites are typically seen as a special case of computing, where computational performance is often the only criterion that really matters. HPC storage platform delivers a powerful combination of raw performance, native InfiniBand connectivity and class leading reliability and availability features—including redundant controllers and host connections, persistent cache backup and hardware accelerated RAID data protection.

IX. KEY CHALLENGE IN COMMUNICATION AREA

- Reduce Hardware Acquisition and Database Transaction costs.
- Support future growth and “green” initiatives
- Provide next-generation agility to meet the demanding needs of mission-critical applications.

X. PROPOSED SOLUTION FOR KEY CHALLENGE

With ref. [9], paper proposes a proof of concept (PoC) study to identify appropriate solutions that can reduce data center cost and deliver next-generation agility and efficiency. After careful consideration, we decided to leverage grid computing for its open architecture, scalability and ability to exploit industry standard servers. The PoC called for a high performance computing (HPC) grid solution—based on blade server technology—to provide improved performance, availability and flexibility, with reduced power usage. To overcome the known performance issues inherent with many grid computing solutions, multi-core processing and low latency InfiniBand networking and storages are utilized to maximize performance and scalability. The HPC grid study intends to demonstrate new technologies and platforms to support the most critical applications; thereby making it imperative to deliver enterprise class availability at every level.

InfiniBand is the key enabler to the high performance, scalable grid computing solutions, providing high bandwidth (up to 20 GB/sec) over a single connection and ultra-low latency (>2 micro seconds (μ s)). Additionally, with Remote Direct Memory Access (RDMA) to dramatically improve node to node communication, InfiniBand adapters can read and write data directly to and from the memory of a server or storage device, without requiring CPU cycles. Because RDMA is built into the lowest levels of the network adapter, high overhead protocol drivers—such as TCP/IP or Fibre channel—are not required. Instead, data is moved directly into or out of application memory space by the InfiniBand network adapter, virtually eliminating CPU communications overhead. Utilizing native InfiniBand storage greatly improves grid computing performance by reducing I/O latency and CPU overhead through RDMA. I/O intensive applications benefit greatly from reduced I/O latency and the data transfer efficiency provided by native InfiniBand storage.

Using InfiniBand as the shared grid and storage interconnect eliminates the complexity and cost associated with implementing a separate dedicated storage network, such as Fiber Channel. From an application perspective, a reduction of network latency noticeably impacts clustered database performance. Database instances, within a cluster, share global resources that are continually updated and must remain synchronized, generating significant intra-node traffic. InfiniBand networking technology significantly reduces intra node communication latency—enabling near-linear database cluster performance scalability. “Using native InfiniBand storage in the data center can bring tremendous price/performance improvements to solutions.” The logical diagram of grid computing architecture is as shown in following figure [9].

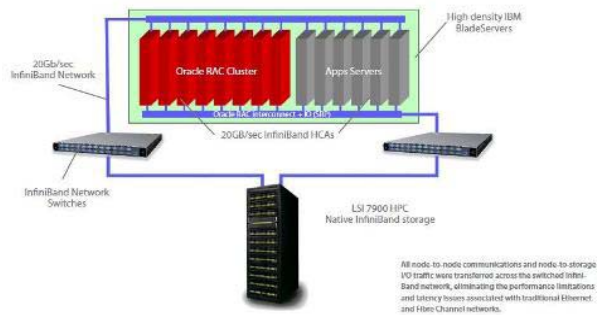


Figure 5. Logical Diagram of Grid Computing Architecture

Next-Generation Infrastructure: InfiniBand HPC Grid Computing

With ref. [9], The PoC yielded compelling reasons to consider the adoption of InfiniBand HPC grid computing and native InfiniBand storage for the next-generation infrastructure platform. Database performance testing demonstrated that the grid architecture and InfiniBand infrastructure would provide the scalability and performance required to meet the most demanding database application loads:

- Using a worst case application test scenario with 100 percent random I/O pattern and a 50/50 read/write ratio, the infrastructure achieved more than 760,000 transactions per minute (TPM) with 4 database nodes and scaled beyond 1.6 million TPM with an 8 node configuration.
- At peak test loads, more than 46,000 I/Os per second were being serviced by the storage system with an average application level response time of 33 milliseconds. In addition, the results specifically indicated the solution was able to provide the necessary availability, scalability and performance requirements at a much lower cost than traditional data center approaches. The near-linear scalability of the InfiniBand grid infrastructure required fewer compute resources to meet application performance requirements, thereby reducing power requirements and helping the company support an important green initiative. In a side by side comparison with similar application loads, the grid architecture consumed approximately one-third the power of traditional SMP systems. In a financial comparison, it was determined that a cost savings of approximately 90 percent could be achieved if the existing SMP systems were replaced with a grid solution based on native InfiniBand infrastructure.

XI. RESULTS OF PROPOSED SOLUTION

- Demonstrated that Oracle® database performance on a grid computing infrastructure can achieve equivalent results to traditional SMP systems at approximately 90 percent lower acquisition cost [9].
- Proved that a grid computing infrastructure- based on blade servers and InfiniBand networks-can provide performance and scalability beyond

traditional SMP systems at approximately 70 percent lower power consumption [9].

- Validated the ability of native InfiniBand networked storage to support 1.6 million Oracle database transactions per second, while reducing IT complexity and cost by eliminating the need for a separate dedicated storage network [9].

XII. CONCLUSION

The study proved to be successful in identifying a next-generation, cost-effective HPC solution. This paper will be able to deploy mission-critical business intelligence and database applications on industry standard blade servers and storage systems with an ultra-low latency interconnect fabric. Additionally, in this study paper, the InfiniBand HPC grid computing solution successfully created an agile, scalable, high density infrastructure solution at a much lower cost.

Grid Computing is becoming the platform for next generation e-science experiments and by Intranet Grid it is very easy to download multiple files.

REFERENCES

- [1] Foster, Carl Kesselman, and Steven Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations", Intl J. Supercomputer Applications, 2001
- [2] Ian Foster, Carl Kesselman, Jeffrey M. Nick, and Steven Tuecke, "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration", <http://www.globus.org>, 2002
- [3] Gregory T. Byrd, and Praveenkumar Cheruvu, "Tuple Space Computing on the Grid", North Carolina State University
- [4] K.A. Hawick, H.A. James, and L.H. Pritchard, "Tuple-Space Based Middleware for Distributed Computing", Technical Report DHPC-128, University of Wales
- [5] Z. H. Du, Y. Chen, and P. Liu, Grid Computing. Tsinghua University Press, Beijing, 2002.
- [6] Z. W. Xu, B. M. Feng, and W. Li, Grid Computing Technology. Publishing House of Electronics Industry, Beijing, 2004.
- [7] M. Z. Li, M. Baker, Kernel Technology of Computing Grid. Tsinghua University Press, Beijing, 2006.
- [8] IBM red Books: Enabling Applications for Grid Computing with Globus. 2002
- [9] "Grid_computing_Case_Study_011409.pdf" Case Study Report, LSI corporatopn ,2007
- [10] "InfiniBand_Storage_WP_050.pdf" white paper, Mellanox Technologies.
- [11] IBM red Books : Fundamental of Grid computing.