

INSIGHT

Base One: Grid Computing for Database-Centric Applications

Earl Joseph, Ph. D.

John Humphreys

IDC OPINION

The market for grid computing is not only an area of tremendous interest and investment, but is also quickly becoming a crowded marketplace, with a host of companies all offering their solution as the next big thing to further grid market development. This comes at a time when users themselves are still struggling not only with how and where to incorporate grid technology into their datacenters, but even more basically, to understand what exactly grid is. Within this fog of uncertainty, Base One is emerging as a company offering a pragmatic, scalable, and extensible approach to the concept of resource sharing. IDC believes this pragmatic approach is beneficial to both the users that adopt the technology and Base One itself. We find that:

- ☒ Beyond a few key compute-oriented applications, grid remains largely conceptual. Increasingly, vendors are "pitching the vision" users struggle to understand how grid can fit in their environments. In this way, grid runs the risk of underwhelming users — and because grid is still largely undefined, users run the risk of adopting a vision that ultimately remains unfulfilled.
- ☒ Base One's focus on a target set of applications and its experience in implementing grid technology provides it with a base from which to launch itself into the broader grid software marketplace.
- ☒ By focusing on where the market is today, Base One is positioned to better respond to shifts in technology focus as the market evolves. It has developed a flexible, modular approach that can be adapted to fit users' evolving needs. At the same time, the company supports a host of different databases and computing platforms, so concerns or issues around customer technology preferences are inherently addressed.

IN THIS INSIGHT

This IDC Insight examines the emerging market for grid computing and more specifically looks at how one company, Base One, is leveraging its technology and experience with large-scale financial services batch processing jobs to bring to market an extensible, scalable, and pragmatic set of tools that allow customers to create highly utilized pools of resources that can be used for large compute- and database-intensive applications.

SITUATION OVERVIEW

Over the last couple years, the concept of grid computing has received a renewed level of attention as customers have looked to find innovative means by which to accomplish their computing needs while implementing deep and severe spending cuts. This need to fulfill a job at a reduced expense is what led market leaders to consider how grid could satisfy their requirements.

Increasingly, we see grid as a virtualizing technology that allows work to be allocated to available resources. There are a host of different types of virtualization technologies emerging in the market today. While companies are approaching the "problem" of creating virtual processing in a multitude of ways, we find the goals are all fairly similar — to decouple the hardware and application resources to drive efficiencies through better hardware utilization and a faster time to solution, and provide more granular analysis.

Today, the most mature market for virtualization is associated with grid computing and is laser focused on migrating large-scale business and scientific applications from expensive standalone large SMPs to more standard, flexible pools of commodity processing infrastructure by taking advantage of the "embarrassingly parallel" nature of these application codes.

Again, companies are approaching the problem of creating virtual processing in a multitude of ways. We found Base One's story straightforward and broadly applicable to heavily compute-oriented and database-intensive workloads. In addition, the Base One solution is relatively mature and vetted with some very demanding users. For those reasons, we have chosen to spotlight the company in this first of a series of documents focusing on grid and virtualization technologies.

About Base One

Base One is a privately held company, founded in 1993 and headquartered in New York City, with an outsourcing subsidiary in Bangalore, India, founded in 1997. Base One's development team is recognized in the IT industry for its work in databases, grids, and cluster computing.

Base One was one of the first software companies to develop and sell database-oriented tools for building peer-to-peer applications. Base One's distributed processing technology takes advantage of the PC's local, client-side capabilities. This allows the storage capacity and processing power of client computers to be brought to bear on tasks that previously had to be done on expensive servers and mainframes. Base One built on this early experience to develop its grid computing software, which has been enhanced and tested in customer environments such as Deutsch Bank, Marsh & McLennan, and Competitive Media Reporting.

Base One's customers include Fortune 500 companies from a diverse set of industries such as finance, insurance, media, and communications.


Base One and the Virtual Supercomputer

Base One's work on distributed computing evolved over more than 10 years of developing tools for building commercial database applications; the company therefore uses terminology like "batch jobs" to describe the work it provides. The concept of a batch job is the abstraction of a logical unit of work and is applicable to today's business applications that run as discrete steps in a multiprocessing environment.

The Base One grid architecture revolves around the model of a "virtual supercomputer" comprised of loosely coupled "batch job servers," which perform tasks that are specified and coordinated through database-driven control structures. The model is virtual, as it uses the available processing power and resources of ordinary servers and database systems, which can also continue to work in their previous roles. The result is what the company terms a virtual supercomputer, because it presents itself as a single, unified computational resource that can be scale both in capacity and processing power. Figure 1 illustrates the Base One Grid Architecture.

FIGURE 1

Base One Grid Architecture



Source: Base One, 2004

Diving deeper into the architecture and technology, the architecture is made up of three distinct components: Rich Client Components, Batch Processing Facilities, and Grid Computing Middleware.

Rich Client Components

Rich Client Components (RCC) reside under the application layer on each of the batch servers. The RCC include four distinct components. The first component comprises the core libraries and utilities for database application development. These libraries comprise the Base One Foundation Classes and provide the basic building blocks from which Rich Client applications are constructed. The Database Class library is the key component of the Base One Foundation, in that it allows the user to easily write efficient, reliable applications for large multiuser databases. The second component contains the Base One Internet Server that allows for transparent operation over the Internet as if the application were on the LAN. Third is the Rich Client application itself, which operates with the goal of minimizing load on the network as well as central resources; this is done through leveraging local batch server resources to gain performance advantages. Finally, the RCC includes a security wrapper that provides single authentication and authorization for all Rich Clients. This is in addition to security features at the operating system and application layer.

Batch Processing Facilities

Batch Processing Facilities provide the mechanism for automating repetitive, highly parallel, large-scale jobs in the distributed architecture. Batch Processing Facilities are included in Base One's Foundation Classes — the database library where the intelligence of the distributed grid lies. The Base One Foundation also contains the Batch Job Server, which is simply a noninteractive Rich Client application that gets its work from a database instead of an end user. The Batch Job Servers look for work by searching for records retrieved from the database. Once a job completes and the results are posted, the Batch Job Server can turn to a shared pool of pending jobs stored on the database. Access control for Batch Job Servers can be tied into an organization's existing database and network security mechanisms. The Batch Processing Facilities also include Administrative Facilities so that grid admins can create and set criteria relating to a number concerns including job priority, date/time, and the completion of a specific task.

Grid Computing Middleware

Grid Computing Middleware is the layer that ties together all the distributed components into a virtual supercomputer for large database applications. The grid software layer incorporates a collection of the batch job servers and at least one Internet server to share both data and compute resources over a widely dispersed area.

Base One's software provides the ability to develop large, distributed, high-transaction-volume, database applications with the same tools and skill level required for building small database applications. Information sharing and transaction processing applications built with Base One's programming tools and middleware scale from the single user desktop to multiorganization grid and cluster computing solutions. There is a simple, consistent application programmer interface (API) that supports secure local or remote access to the leading industry databases.

Unlike other "grid" computing solutions that rely on a centralized grid scheduler, Base One's architecture uses the concept of a Batch Job Server that can work independently to search out work that it can perform, coordinating its work with a central database.

Benefits

Base One looks to provide users with an easily implemented grid architecture through a shrink-wrapped software package. The solution encompasses support of most of the major industry database and hardware vendors, including the major database vendors and Windows, Linux, variants of Unix, AS/400, and OS/390. Base One also provides consulting implementation and training services to streamline and ease adoption.

Base One's software is the first grid computing product for building Windows applications that is available in shrink-wrap form. The software has been enhanced over years of production use, and it can be installed quickly, with limited start-up costs.

Base One has Fortune 500 customers in multiple industries including finance, insurance, media, and communications.

The company sees its benefits as fitting into four distinct areas:

- ☒ **Lower cost of development and operation.** Because the software makes use of existing Windows programming, there are no specialized skills required for development; some early customers have reported operational cost reductions over 80% from previous platforms.
- ☒ **Faster implementations and deployments.** Customers can leverage the foundation class library to create and test prototype grid solutions, and then roll out those solutions once testing is complete.
- ☒ **Secure data and application.** The Base One grid solutions include both security and recovery procedures that integrate with and augment those found in the database and operating systems.
- ☒ **Increased performance and capacity scaling.** As new processors are added to the pool, nearly all of that capacity would be available to the grid. Also, by not having a centrally controlled grid, new resources can easily be added, and multiple workloads can be supported without difficulty.

Additional benefits include the following:

- ☒ The architecture is designed to deliver a high degree of reliability, recoverability, and fault tolerance.
- ☒ The grid solution has extensive application security and administration facilities.

- ☒ The solution provides a comprehensive application development framework, enabling quick prototyping and flexible modification of existing applications.
- ☒ It includes comprehensive documentation, examples, and programming tools that make it easy to learn and use effectively.

Market Focus

The company's history of working closely with Wall Street and other financial services firms has worked to deliver Base One a deep understanding of not only grid technology, but also of how that technology can be successfully leveraged in compute- and data-intensive batch applications and for large compute "jobs," including:

- ☒ Large-scale financial processing such as credit card processing, portfolio analysis, financial forecasting, and risk analysis
- ☒ Large-scale data handling and processing such as data mining, data warehousing, database processing, data collection, and monitoring in areas as diverse as medical, security, and industrial segments
- ☒ Complex scientific applications, including pharmaceutical research/drug discovery, seismic data analysis, and simulation

Challenges

There are a number of challenges and questions facing Base One and other grid companies:

- ☒ How will the grid evolve over time? Will it become the core of most computer centers, or will it be constrained by the network limitations?
- ☒ Who will the next adopters be? What is required to move grids into mainstream production environments? How can the political and trust issues be solved?
- ☒ What are the business models for vendors?
- ☒ How can these companies quickly expand their user base (and attract investors)?

For Base One there are also the challenges in being a smaller company competing with the larger major vendors. Base One needs to clearly articulate its advantages against the larger vendors and the growing set of other grid software vendors. It needs to clearly articulate the value of its solution and show the ease at which the tools and architecture can be implemented. Advantages will wane over time as others copy Base One's technology, so new competitive capabilities and features need to be continually added over time.

Developing strong partnerships will also be key to success. Partners that are in Base One's targeted industries will be very valuable.

FUTURE OUTLOOK

Grids are primarily occurring in HPC environments today. However, research indicates that grid is quickly becoming a standard "check-off" on RFPs coming out of the financial services industry. IDC believes that many industries beyond just financial services and science will recognize services within their organization that would benefit from grids. Data mining, real-time business analytics, and supply chain management workloads, among others, will likely emerge as "grid applications," opening the door for grids with other, more traditional commercial workloads and industries.

Market and Technology Drivers for Grid Adoption

- ☒ Hardware utilization rates in the industry-standard server market are very low, with IDC research indicating an average utilization rate of only 15% on an x86 processor-based server platform. Grids provide an avenue to raise utilization rate and therefore reduce capital and operational costs of managing excessive numbers of systems.
- ☒ Parallel application adoption and conversion is a factor. As more applications and databases are moved to a parallel processing model, clusters and grids will have a larger total available market to address. This trend is already well underway for many technical and custom applications, and is beginning to impact the commercial market with products such as Oracle's 10g database.
- ☒ As the number of users or enterprises that have access to or influence on data within a company's IT infrastructure environment increases, extending beyond the physical borders of a company and into the partners, suppliers, and customer base, a more flexible capacity planning model will be demanded.

Market and Technology Inhibitors to Grid Adoption

- ☒ The greatest inhibitor identified by IDC research and customer interviews was organizational. The cultural challenges of thinking about compute resources in a new way and sharing these resources across potentially multiple business units or organizations represent the greatest impediments to commercial grid adoption. Customers noted that grids require a different way of thinking about how to deliver IT datacenter services, and normal resistance to changing behavior is always the toughest hurdle to overcome in technology adoption.
- ☒ The general lack of tools and industry standards lead many of the sites interviewed to think of grids as requiring large people and services costs — costs that would mitigate any infrastructure cost savings. Essentially, things that are hard to do don't get done. Vendor investment in management, provisioning, optimization, and scheduling software will help address this market inhibitor.

- ☒ Security is also identified as a major inhibitor to broad commercial adoption of grid technology. Security concerns are also deeply rooted in the cultural and organizational concerns cited above. Security will have to be proven over time to potential customers at a number of levels for grids to be considered for adoption in shared workload environments.

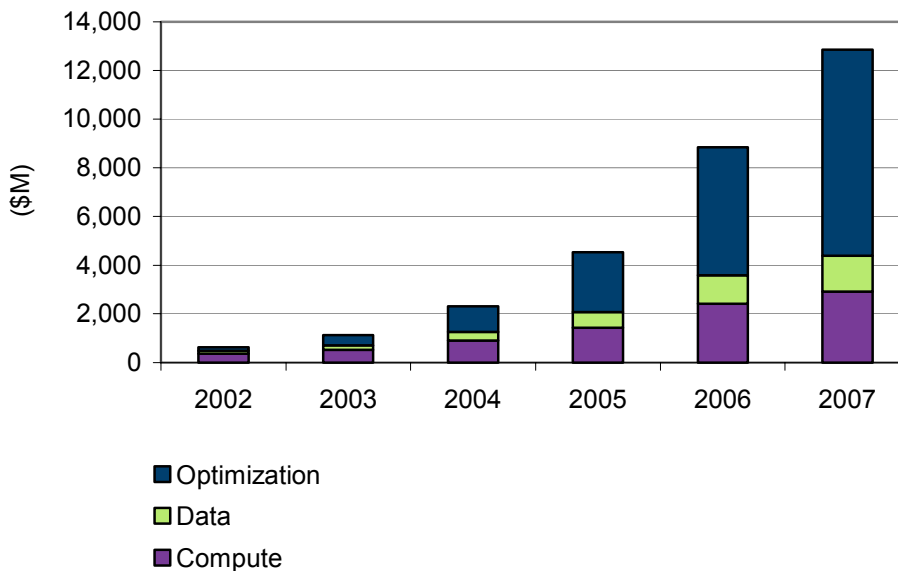
- ☒ The cost of compute resources continues to decline. As the price point for compute resources falls, and the IT industry invests in better tools to remotely manage, deploy, and redeploy hardware and software in a datacenter, many customers will care less about server utilization and overprovisioning as long as the costs of managing that capacity decline. This trend may detract from the promise that grids bring better utilization server infrastructure and more access to compute resources.

Figures 2 and 3 illustrate our forecast for spending by grid category, which includes hardware, software, and services spending for grid, as well as how that spending will be allocated between traditional technical computing, high-performance commercial computing, and transactional commercial computing. The near-term market and early adoption for grids has largely been in the high-performance computing (HPC) market for large, batch-oriented compute grids.

IDC sees significant grid market growth through 2008 in database-oriented business applications enabled for grid. Base One is focusing on this high-growth market space.

FIGURE 2

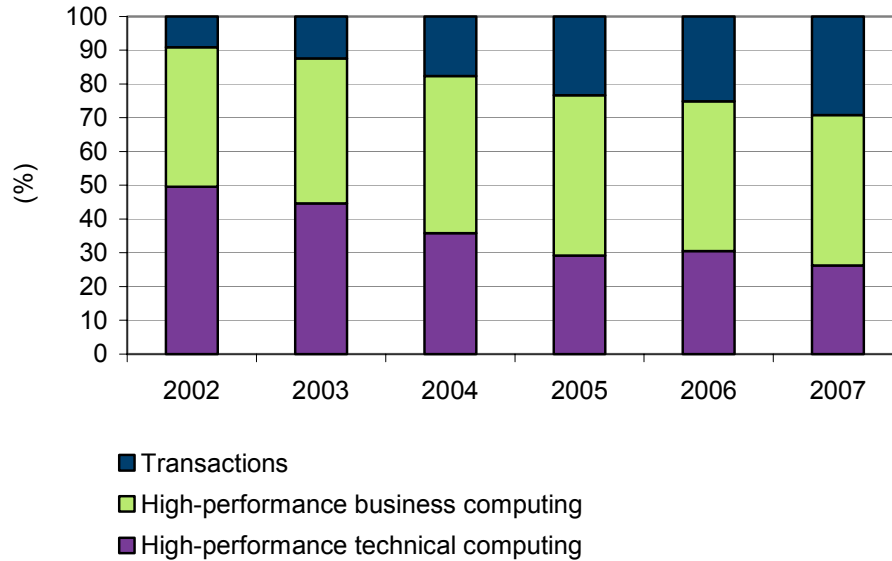
Worldwide Grid Revenue by Grid Category, 2002–2007



Source: IDC, 2004

FIGURE 3

Worldwide Grid Revenue Share by Workload Category, 2002–2007



Note: Transactions are grids for transactional commercial workloads.

Source: IDC, 2004

LEARN MORE

Related Research

- ☒ *Keeping Pace with Technical Computing Market Dynamics, Part 1: Market Segmentations* (IDC #31877, September 2004)
- ☒ *The Role of Grid Computing in the Coming Innovation Wave* (IDC #30871, March 2004)
- ☒ *The Rise of Distributed Computing: Will Smaller Branded Vendors Disrupt the Incumbents in the HPC Cluster Space?* (IDC #30866, February 2004)

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2004 IDC. Reproduction is forbidden unless authorized. All rights reserved.

Published Under Services: Workstations and High-Performance Systems