THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2008007 February 11, 2008

Blue Clouds on the Horizon — IBM Designs for the Stratosphere

Analyst: Brian J. Dooley

Management Summary

Last November's announcement of its upcoming *Blue Cloud* offering (set for availability this spring) has again displayed IBM's capacity to validate and define a vision for future infrastructure development. It also highlights the ambiguities inherent in defining a next-generation architecture. **Like *Web 2.0, cloud computing* is itself a somewhat vaguely-defined territory that encapsulates a current stage of evolution rather than proposing something revolutionary or new.** The fact that a stage has been reached that can be defined and will have potentially revolutionary repercussions across a wide range of areas makes new terminology useful. *Cloud computing* essentially virtualizes the data center, simplifying IT management and optimizing the creation and deployment of software workloads. **Blue Cloud is IBM's initiative to standardize the definition of clouds and bring real cloud computing solutions to market.**

"Cloud computing" is a term that has been defined by the gigantic data centers created to handle Web operations, principally for Google and Yahoo, to create a large-scale clustered "supercomputer" through use of a large number of low cost processors, in a manner similar to the older Beowulf cluster strategies, but supported by a more user-friendly development, access, and management software. It also grows out of the grid-computing environment, used for manipulation of large-scale databases, and offers virtual machine support and special capabilities for Web access and Web application development. As such, it represents the coming together of a variety of streams of development that have been progressing independently for the past decade.

The Blue Cloud announcement contains a number of items of interest that might point the way to future systems. First, it is important to note that IBM is taking this very seriously, and regards it as a key developmental direction. Secondly, the offering is itself a mixture of Open Source products with IBM's own Tivoli systems management offerings, working together in an infrastructure that would need to be highly integrated. The third significant aspect is that the initial rollout will be on low-cost x86 processors and the *System z* mainframe, with other servers (like *Power*) following about a year later.

As an indication of direction, Blue Cloud points to a need to support and manage corporate data centers potentially made up of thousands of processors and devoted to a wide variety of tasks, including both backend business services such as ERM and large-scale database operations (including data warehousing and data mining), plus customer-facing electronic commerce and service offerings. Use of these large-scale clustered systems is seen as providing extreme scalability, as well as granularity in resource assignment, for improved efficiency. Linking all of these processors together to support mission critical applications, however, brings about a special need for management and support. In fact, such a data center operates in many ways like a large mainframe and the management considerations for performance are similar. This means that IBM can leverage its mainframe expertise and software to support the new environment.

The Blue Cloud initiative remains somewhat open ended. There are a number of pilot projects using the technology – most significantly, an initiative with the Vietnamese government – but the concept is likely to continue evolving right up past the spring 2008 release.

IN THIS ISSUE

The Clipper Group, Inc. - Technology Acquisition Consultants ♦ Strategic Advisors
888 Worcester Street ♦ Suite 140 ♦ Wellesley, Massachusetts 02482 ♦ U.S.A. ♦ 781-235-0085 ♦ 781-235-5454 FAX
Visit Clipper at *www.clipper.com* ♦ Send comments to *editor@clipper.com*

Blue Cloud Announcement

IBM announced its Blue Cloud initiative on November 11, 2007, with offerings expected to be available to customers in the spring of 2008. The IBM vision is a *BladeCenter* composed of multiple blade servers holding multiple processors and local storage, to create a cluster of up to several thousand processors. Provisioning, management, and monitoring will be handled by *IBM Tivoli Monitoring v.6* and *Tivoli Provisioning Manger v.5.1*, with access through *WebSphere Application Server*, and *DB2*. On this platform, a virtualized infrastructure based on *Linux* and *Xen* virtualization supports individual applications. Linux with Xen will be managed by the Tivoli Monitoring agent to provide multiple virtual Linux machines. Parallel processing is enabled by the open source *Hadoop* provisioning manager supporting the Eclipse programming infrastructure, with Google's *Map/Reduce* programming model used to create distributed processing loads. This mix of IBM and open source programs will support Power and x86 processors, and Web 2.0 (IPv6) resource reservation for real-time streaming, as needed. The initial scheme is illustrated in Exhibit 1, below.

Initially, Blue Cloud is being run as a pilot project with the Vietnamese Ministry of Science and Technology to escalate Vietnamese use of Web resources. The offering is based on IBM's own *Enterprise Web 2.0 Innovation Factory* solution that has been operating out of IBM's Almaden Research Center.

The Vietnamese program is creating the *VISTA Innovation Portal (VIP)*, to provide a platform to foster collaborative innovation among major universities and research institutes in Vietnam. It supports a full range of collaborative applications, including Wikis, Blogs, and Forums, and it is powered by *IBM Innovation Factory*. Key goals are to:

- Explore extreme scale easily and quickly
- Share infrastructure resources efficiently
- Simplify IT management
- Handle new and emerging workloads
- Provide a platform to encourage open collaboration.

The Blue Cloud solution is another iteration of IBM's work with large-scale clustered systems, including *Parallel Sysplex*, *Deep Blue*, the *Blue Gene* supercomputer network, and grid computing. IBM is also involved in a number of other cloud computing projects, most notably in an alliance with Google to promote new software development methods designed to meet the needs of the future Web, particularly addressing highly-parallel computing. In this project, Google and IBM have dedicated a large cluster of systems in an array that is planned to grow to over 1,600

Exhibit 1 – Blue Cloud Initial Offering Configuration

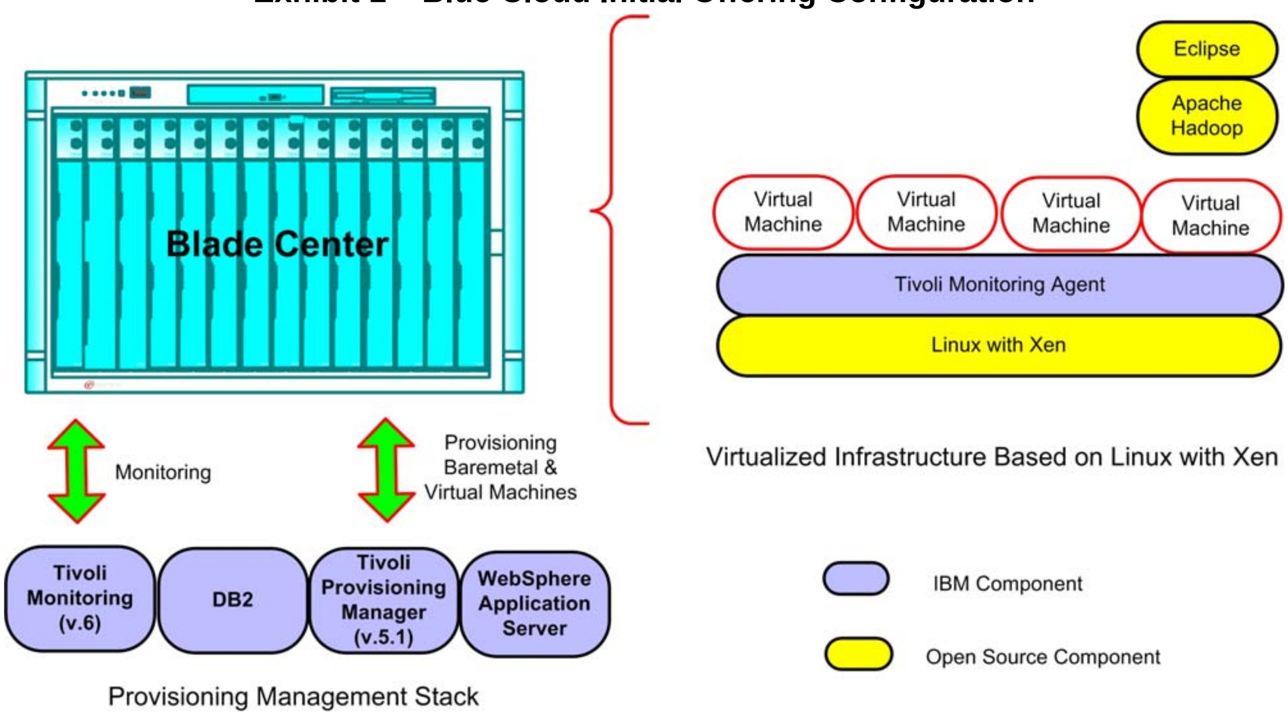

Source: IBM; Drawing by Clipper

processors, accessed via the Internet for the purposes of testing parallel programming course projects. This configuration also provides an open source software stack based on Linux, Xen virtualization, Hadoop, and Google's open source Map Reduce parallel computing model, and the *Google File System (GFS)*.

IBM's short-term plan is to work directly with a series of data center customers to deploy Blue Cloud solutions to help accelerate IT simplification in their data centers now. Over time, these early implementations will evolve to create a family of ready-to-use cloud offerings based on open standards and open source, combined with IBM software (Tivoli) systems technology and services. These offerings will initially be targeting System z and x86 processors, with intentions to later support *Power* and highly-dense rack clusters and storage. The architecture will be designed to reduce management complexity, increase responsiveness, and support both existing and emerging data-intensive workloads. For these reasons, this must be seen as a key direction for infrastructure development that is now in an alpha testing mode. **One of the stated aims of Blue Cloud is to allow corporate data centers to operate more like the Internet through improved organization and simplicity inside the datacenter. Over time, those solutions will enable access through a distributed, globally-accessible fabric of resources.**

Key Components

The Blue Cloud offering is based around a set of open source components backed by the IBM Tivoli management suite. The Tivoli suite is a well-known IBM component that has stood the test of time, and has emerged as a critical player in all areas of systems management, principally in mainframes. IBM also has a growing expertise in Linux development and deployment that will be brought to the mix. The key open source components in Blue Cloud, after Linux, are *Hadoop* and *Xen*.

Hadoop is an open source software project under the Apache server project, designed to run jobs, distribute tasks and store data in a parallel and distributed fashion on a multiprocessor system. Most programs run on Hadoop are written in the Map/Reduce style, in which input is broken into small pieces that are processed independently according to a map. The results of these independent processes are then collated into groups and reduced to produce a result. As an example of potential scalability, Yahoo has 10,000 machines running Hadoop, with the largest cluster having 2000 nodes. This supports one petabyte of user data, and 10,000 research jobs are run per week. IBM is providing Eclipse, a business development platform, on top of Hadoop, and has written its own Map/Reduce toolset for Eclipse in support of this.

Xen is high-performance open source virtualization software. It operates through the Xen hypervisor, which sits between server hardware and operating system permitting each physical server to run one or more virtual servers. Virtual server images can be run on any server at any time, and multiple virtual servers can simultaneously share a single server.

These open source components are supported by the Google File structure, which create a redundant shared file system that can be used with or without a Storage Area Network (SAN). Data is staged on local disks to provide high throughput without extraordinary bandwidth requirements.

Why Cloud Computing?

A cloud is a pool of virtualized computer resources capable of hosting a variety of different workload types, with rapid deployment. It includes redundancy and scalable program models to ensure reliability, and management features provide resource monitoring. Clouds can support grid-computing models, which divide a large task into smaller tasks to run in parallel, but also support non-grid processing for running Web 2.0 applications. Management capabilities are seen as a key part of the definition.

Currently, there is an onrush of projects in support of cloud computing as companies seek to consolidate technology offerings, integrate them with a spike in requirements driven by Web 2.0 applications, and support future requirements driven by increased network bandwidth, ubiquitous computing, and real time application requirements.

Much of the media attention around this concept has centered around leading Internet providers such as Google, Yahoo, and Amazon, each of which has major cloud computing initiatives. However, it is important to bear in mind that multiprocessor "supercomputers" based on open source *Beowulf* clustering have been around for some time, offering a variety of access and provisioning methods. Additionally, grid computing has been strongly supported in corporate data centers by companies such as Oracle, which view clustered processing as essential to handling large-scale data access needs. Yahoo announced on November 12 that it has developed a project involving 4000 processors to conduct software research with Carnegie Mellon. Google is making hundreds of processors in this data center available to schools. One of the key issues in all of these cloud-computing ventures is to provide experimental platforms for next-generation massive-cluster computing projects, and develop the skills required to program and use such architectures. Any progress in this area demands a large body of available expertise around the world, and the technology is moving faster than career development – meaning that there is a risk that when this type of processing reaches wide-scale deployment, there won't be enough available talent to support it. Thus, it is in the best interests of companies, such as IBM and the Internet services vendors, to support training in this are, and at the same time to promote brand awareness among the upcoming generation of programmers.

The key factors that IBM sees as driving cloud computing architectures are:

- Explosion of data intensive applications on the Internet, particularly in collaborative multimedia (Web 2.0) applications such as *YouTube* and *Flickr*.
- Rapid and continuing growth of mobile commerce.
- Increasing network capacity and availability, in all technologies, including 3G and 4G cellular, WiFi 802.11 N (and above), WiMax, and increases in Ethernet bandwidth.
- Advances in computer architecture and price/performance, principally driven by multiprocessor clusters, as processor power continues to obey Moore's Law with the price per MIP falling steadily.

One of the potential problems of cloud computing – or any form of massive cluster system – is that of management and reliability. Blue Cloud is meant to address this by proving a range of management options from its own Tivoli environment, as well as from Open Source.

The Promise of Blue Cloud

The key target of Blue Cloud technology is the large corporate data center, where it will be used to manage dense blade center arrays to support Web 2.0 processing requirements. With the initial offering, a customer will be able to buy a BladeCenter starter kit with Power and x86 processors to create their own clouds.

Standardization, particularly focusing upon open source software components and adherence open standards, is likely to make this a key architectural blueprint for this type of compute platform for the future. If a single architecture becomes constant, it in itself will become a standard, and this will make it easier for development, maintenance, deployment, security, and training. And training is seen as critical, with one of the strongest foci at the heart of the Vietnamese portal deployment, as well as in academia (the IBM/Google project), is to find ways to apply the cloud infrastructure in collaborative applications, and, perhaps more importantly, in developing the development skill sets that will be necessary to create advanced parallel processing applications in the future.

For the smaller business, Blue Cloud also provides some promise – as well as a few caveats. By providing a standard and scalable offering based primarily upon open source software, IBM is creating the possibility of developing large scale computing capabilities at low cost, based upon a small array of blades. The caveat is that the architecture does require Tivoli, and, although much of the other software is open source, implementation is likely to be a nontrivial exercise. Supporting the infrastructure and developing applications could prove expensive and difficult. With a large data center, costs can be shared among a much larger range of potential applications and clients. However, Blue Cloud also opens the way for service providers to open Blue Cloud-based remote data centers similar to the Vietnam Innovation Portal – actually based at an IBM data center in Almaden. They could provide billed services to SMB users, with a business model similar to that used for earlier architectures such as the New Zealand Supercomputer Center (NZSC), which provides on demand configurable processing time on an array of over 1000 processors as an offshoot of the Weta Workshops "supercomputer" movie animation system (used in *Lord of the Rings*).

Although this latter strategy will provide some new markets, even this will require some expertise in configuration above the raw processor level, including programming and management. This is likely to develop a new market for such services, from both the system providers and from external consultants. And that is where the need for developing new expertise comes in, thus bringing us full circle.

Development of cloud computing is driven by continued development of real-time multimedia data streams and service-oriented architectures on the infrastructure side, and by recent growth in collaborative and Web 2.0 applications on the demand side, all of which have led to escalation in processing requirements. A cloud computing architecture should be able to integrate with existing IT infrastructure via SOA Web services. IBM's Blue Cloud initiative is a response to the growing complexity of this environment and the need for optimization and management.

Conclusion

Blue Cloud is not a revolutionary development, nor is it meant to be. It is more a statement of approach that encapsulates the current development level of high density clustered computing. By drawing a line in the sand, IBM is attempting to provide guidance in the shifting area of high performance hardware, so that applications vendors, developers, and users will be able to focus upon a similar architecture. This must be viewed as a good thing, bearing in mind the fateful consequences of the Tower of Babel. It is also likely to sell a lot of iron and services, which, for IBM, will be regarded as a satisfactory outcome.

Above all, Blue Cloud must be viewed as a reaction. Development of large-scale clusters is already happening, and has been going on for some time. Management problems have already begun to emerge, particularly where such clusters have been used for other than a single application or application type. Flexibility, resource allocation, billing, and other issues have arisen, along with an emerging consensus that there is too little expertise in the field to support what might become a critical component of next-generation infrastructure.

Developments in open source software, particularly, Hadoop, have made it possible to provide an easier to use and more flexible clustered system that is still able to perform the compute intensive tasks that parallel systems were initially designed to run. IBM is supplying its Tivoli management system to integrate the whole into something like a single product, and it is working with the open source projects to ensure that they will provide secure enterprise level operation.

Blue Cloud still represents an effort to support a next generation of computing whose needs still may be largely unknown. It is designed to power Web 2.0 applications, backend processing, multimedia, large single server applications, massive computing tasks run in parallel on multiple servers, and virtually anything else that might be imagined. Therein lays the ambiguity. One of the flies in the ointment is likely to be data access, which, though supported by a redundant file system, is bound to create bottlenecks (the same issue raised with many Beowulf cluster systems). Still, ultimately, the experience does not yet exist, so it is impossible to determine what new requirements will be raised. Nonetheless, to have infrastructure, support, and expertise for the future, it is necessary to start now, as IBM has done.

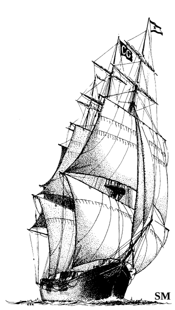

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

***Brian J. Dooley* is a Contributing Senior Analyst for The Clipper Group.** He has more than 25 years experience as an author, analyst, and journalist focusing on IT trends. He has written six books, numerous user manuals, hundreds of reports, and more than 2,000 magazine features. Mr. Dooley is the founder and past president of the New Zealand chapter of the Society for Technical Communication. He initiated and is on the board of the Graduate Certificate in Technical Communication program at Christchurch Institute of Technology. Mr. Dooley currently resides in New Zealand.

- ***Reach him via e-mail at bj.dooley@clipper.com.***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager, Clipper Notes**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.