THE CLIPPER GROUP
Navigator™

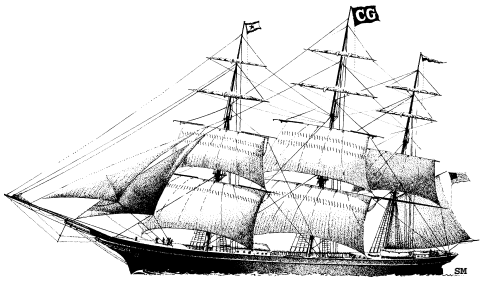

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2008024R April 25, 2008

IBM Changes Paradigm for "Megaplex" Provisioning

Analyst: David Reine

Management Summary

When you or I go out to buy an automobile, we have some very specific requirements in mind, such as price range, 0-to-60 performance, interior comfort, fuel economy, etc. We review the specifications, test drive as many cars as necessary, look for a deal, and make a decision. What does the person managing a fleet of vehicles do? For example, how does the buyer for a rental car franchise make a decision involving a significant number of cars? Rental car fleets usually consist of a variety of vehicles, from luxury cars, to full size sedans, down to compact cars, with standards established for each class. The specific mix is based on a variety of factors, such as location (downtown vs. resort), weather (convertible vs. four-wheeled drive), and luggage capacity. This does not even address the more superficial factors, such as color. *Perhaps Henry Ford had it correct with the Model T after all* – **you can have any color, as long as it is black**! Determining the specific mix of cars at any one location can turn into a heterogeneous mess, as executives and vacationers look to reduce their carbon footprint by renting cars that use less gas and burn the gas that they do use better, thereby reducing the pollutants that they create, while retaining all of their creature comforts. The rules of the road have changed.

The CIO responsible for the purchase of thousands of servers, in fact, maybe even thousands of servers a month, faces the same dilemma. The enterprise data center, even a "megaplex" belonging to an Internet giant, has a limited floor space and a fixed energy budget. If the infrastructure complex exceeds the available floor space, it could entail the construction of an additional facility, at a cost of millions of dollars. The local utility will ensure that the data center stays within its energy budget – they probably have no more electricity to give you, in fact, they are providing credits or rebates if you *reduce your electrical consumption*. The IT staff has to increase the density of the fleet of servers in its domain, deploying more servers per square foot than ever before, in order to satisfy mission-critical applications. The enterprise must reduce energy consumption, not only to power the server complex in the environment, but also to cool it, especially when you are paying as much to air condition the data center as you do to power the servers. The CIO must reduce the total cost of ownership (TCO) for an integrated enterprise IT infrastructure yet still have enough compact servers available to do the critical mission.

Many enterprises are currently deploying a complex interconnection of thousands of rack-mounted, pizza-box style X86 servers, or bladed servers, to ensure the availability of sufficient resources to satisfy the myriad mission- and business-critical applications needed to get the job done. IBM, through its new *System x iDataPlex* dense hosting solution, has now changed the deployment paradigm for the mega provisioning centers, rewriting the rules for data center provisioning with innovative packaging to ensure sufficient availability, standardized thin servers to facilitate the deployment process, and energy conservation to lower TCO. To learn more about iDataPlex, which only comes in black, please read on.

IN THIS ISSUE

Enterprise Data Center Issues

Recently, there has been a mad rush among server providers to announce systems geared to satisfy the requirements and desires of the most common denominator, the small and medium business, the SMB. With a generally-accepted employee count of between 100 and 1000, there are literally thousands of SMBs with a thirst for anywhere from 10 to 50 servers to deploy in a scale-out, rack-mounted architecture, or packaged in blade chassis, to satisfy their mission- and business-critical application needs. At the other end of the spectrum is the new-age enterprise that uses Web 2.0-style computing to operate massive data centers with thousands of servers looking to deploy *new enterprise* data centers, or *megaplexes*, with potentially millions of customers seeking access to one of its web-based applications. This new megaplex has an entirely different set of pain points from the more traditional, legacy data center.

Where the legacy data center might have concerned itself about hardware resiliency, server availability, and price/performance in a scale-up OLTP environment, the new enterprise data center has a different set of priorities to satisfy. These are the needs of a distributed streaming application environment, for streaming video, online gaming, and social networks, where the customer is responsible for the operating environment. This environment is best suited to a scale-out environment with thousands of X86 servers installed in a "green" computing environment. This new enterprise conceives of deployment in terms of "data center units", where each unit is a megaplex data center, as opposed to servers or racks. The new enterprise requires a solution built around a "cloud" consisting of a shared pool of homogeneous resources, built in a "stateless" environment where the user only sees a virtual portal to those resources, and **deployed in *giant steps***[1] with fully-populated, pre-configured racks of tested systems, but **while maintaining a small carbon footprint**. The plan is for massive expansion in an extensible architecture. Chief among their issues are reducing the total cost of ownership (TCO) for IT, as measured by hardware acquisition cost, cost/Watt (or Watts/Sq. Ft.), through the optimization of capital costs and the reduction of operational expenses, including, but not limited to, reducing the amount of energy resources consumed in the data center. Another major issue is the availability of data center floor space – the enterprise must improve the density of servers within the data center or face the continual problem of building more new data centers. For a more complete list of new enterprise data center concerns, see Exhibit 1, above.

Exhibit 1 –
New Enterprise Data Center Issues

- Lowering TCO to meet a limited budget;
- Data center floor space to be expanding requirements;
- Improving hardware efficiency to reduce the amount of energy consumed for power and cooling;
- Increasing the usable density of servers in any given rack in the data center;
- Configuring a single point of management for thousands of open systems servers in a scale-out environment;
- Deploying a flexible infrastructure with a few standard servers replicated thousands of times to ensure application availability;
- Enabling an architecture that permits rapid scalability to meet the needs of cyclical Internet workload; and
- Installing an architecture that can be deployed and maintained around the globe.

Controlling power consumption in the data center through more efficient server utilization is important from both a TCO and environmental standpoint. When you consider that the enterprise is spending a dollar to air condition the data center for every dollar spent on powering the server environment, you realize the tremendous savings that are available from improved power management. In addition, with thousands, or tens of thousands of servers per megaplex, required to satisfy the mission- and business-critical application mix, there is a tremendous need to increase compute density, in order to make better utilization of the data center's floor space. Unfortunately, many of today's existing server architectures prohibit the IT staff from completely filling any standard 19" rack, as there simply is not enough power available to the rack, with the typical data center topping out at about 15KW per rack.

In addition, the enterprise needs to be able to simplify this ultra-scaled IT infrastructure through the implementation of an open, reproducible architecture, enabling the CIO to deploy

[1] Say, in increments of 500 or 1000 servers at a time.

it around the world. Make that, deploy it around the world **and** retain control of the IT infrastructure back to the data center.

IBM's Challenges

IBM has admitted that they are late in addressing the problems associated with the implementation of a global, open systems megaplex. However, two years ago, they made an overt decision to focus on the business goals of the largest data centers to insure success in a Web 2.0 world or a high-performance computing environment, including maintaining a "green" environment in an era of constant acquisition and IT globalization.

These enterprises have an insatiable demand for an ever-increasing supply of processor power, constrained not only by a limited budget, but also limited by a fixed amount of electrical energy available to the megaplex data center. Controlling the TCO of the IT environment is of paramount importance to them, not only by limiting capital expenditures, but also by implementing rigid controls on operating expenses. Building an energy efficient platform that reduces the amount of energy required to power and cool the environment, as well as reducing the cost/watt, is one of key enabling strategies.

Another key contributing factor toward reducing the TCO of the megaplex is controlling the amount of floor space required to house the application servers. The data center needs to increase the density of their server rack deployment. However, as many CIOs have learned, increasing server density is dependent upon the ability to bring enough electrical power to the rack to drive this increased workload. As most data centers do not have an option to increase the amount of power to the rack, the data center staff must find a way to reduce the amount of power required by each server within the rack.

When deploying thousands of servers at a time, it is critical to the enterprise to implement a server environment that is easy to configure, buy, and own, with a rapid service delivery. The enterprise must simplify the data center infrastructure with preconfigured racks that plug and play with no negative impact on the data center environment. The enterprise needs a new and intelligent data center model to satisfy the need for explosive computing growth while lowering the TCO structure.

IBM has spent the last two years building the server infrastructure that changes the enterprise server paradigm to accommodate megaplex needs. Starting with a commitment to increase compute density tenfold by 2010, reducing the need for new construction, IBM set out to eliminate the requirement for specialized computer room air conditioning (CRAC) and to decrease power consumption by 50%. Reducing power consumption and eliminating the CRAC units attack the largest operational expense facing the data center while capping its carbon footprint. IBM calls that infrastructure *iDataPlex*.

IBM's iDataPlex

With a blueprint drawn from enterprise-level experiences with *System z*, scale-up open systems servers from *System p*, and the scale-out architecture of *System x* and *BladeCenter*, IBM has started anew with a blank sheet to change the megaplex compute paradigm. Using a new, innovative modular technology adapted from IBM's BladeCenter, IBM is building iDataPlex with a *Lego* brick-like construction flexibility and a holistic rack design. (See Exhibit 2, on the next page.)

In order to improve data center density and reduce cooling requirements, IBM has taken the innovative step of implementing a unique rack design with shallow depth, optimized to save floor space, thus altering the way servers mount within a standard 19" (42U) rack. IBM redesigned the servers themselves with a new, more efficient form factor, enabling them to fit side by side, in two columns of flexible nodes, and optimizing airflow for cooling efficiency in a half-depth rack. With 42U available in each column, the data center has 84U of horizontal space to install compute nodes, with an additional 16U for the vertical mounting of integrated Power Distribution Units (PDUs) and communications nodes from Blade Network Technologies, Cisco, and others. IBM has 1GbE, 10GbE, and InfiniBand switches available, with scalability up to 384 Ethernet connections. This enables iDataPlex to interconnect with an existing network and manage servers automatically. It is interoperable with Cisco management. This enables the data center to install 100U of mission- and business-critical infrastructure in the space previously occupied by 42U worth of rack-mounted infrastructure, a two-floor tile rack footprint, with low-cost integrated cable management. IBM has taken the redesign several steps further, by deploying 2U and 3U chassis to share power and cooling resources between a pair of power-optimized servers, or a server and storage, improving airflow efficiency by reducing airflow impedance, and redesigning their unique Rear

Exhibit 2 – iDataPlex's New Efficient Form Factor

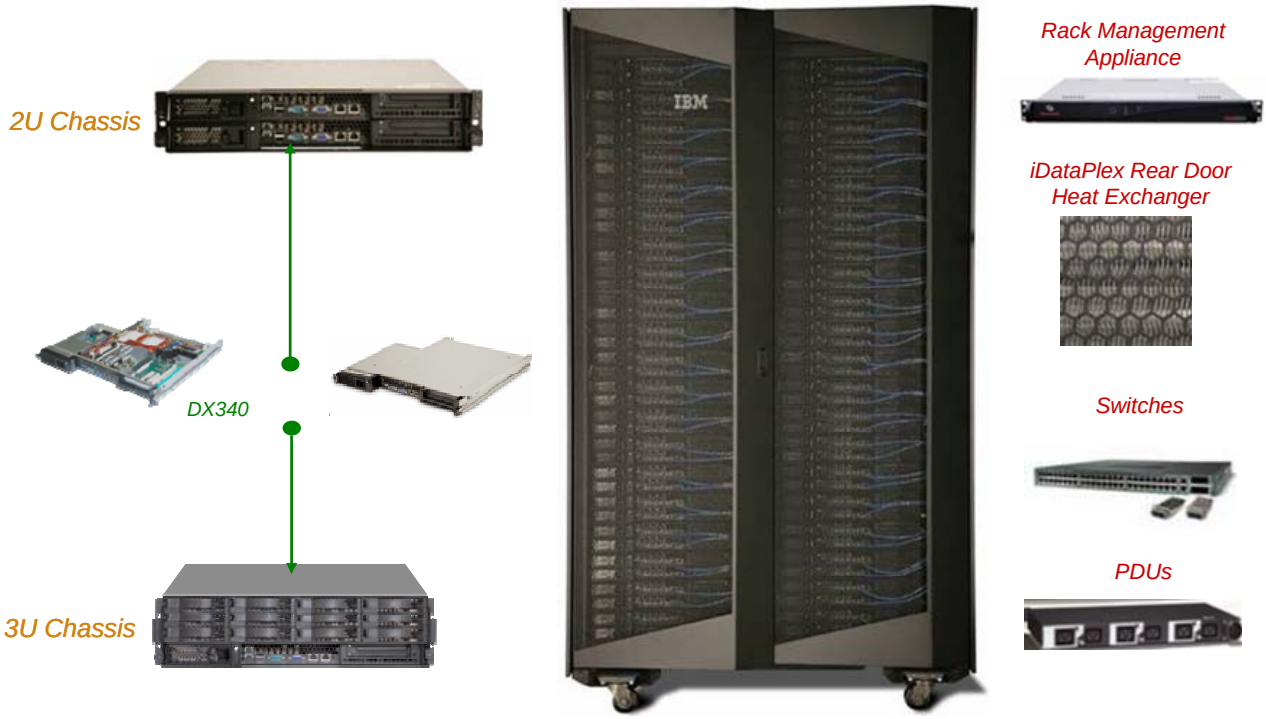

Source: IBM

Door Heat Exchanger. Doing this reduces iDataPlex power requirements, enabling the data center staff to activate all servers concurrently, taking full advantage of this increased density.

In addition to reducing power requirements for the servers, IBM also redesigned the rear door heat exchanger, previously introduced for IBM racks, to reduce the amount of energy required to cool the data center. This new wider cooling solution has more surface area to remove generated heat. Not only does it remove the heat generated by the internal iDataPlex infrastructure, but it also expels air cooler than that entering the front of the rack, cooling a room from 75 to 65 degrees, virtually eliminating the "hot aisle", enabling a 2.4x improvement in server density. In many cases, iDataPlex with a rear door heat exchanger can eliminate the need for CRAC units.

Because IBM designed iDataPlex within a standard 19" rack with front access, there is no reason to access the rear of the rack and there are no changes necessary to the existing data center layout, allowing existing power feeds and communications cables to remain in place, and allowing the iDataPlex racks to be deployed closer together. In fact, iDataPlex can aggregate the power feeds to reduce the number necessary to power the total systems in the same footprint. At $1500-$2000 per feed per month, a megaplex data center can save several million dollars per year from reduced power requirements. With all nodes and cables accessible from the front of the rack, iDataPlex is an easy to access and easy to service platform.

IBM designed each iDataPlex node, either 2U or 3U, with a single high-efficiency power supply, in fact, over 90% efficient, along with a single, shared high-efficiency, low-cost fan pack, sharing four fans between the two internal trays, with 40% less fan power consumption than is found in a traditional rack-mounted implementation. This is possible because of the design of iDataPlex. **With a rack depth that is reduced so that there is less air impedance, the fans have significantly less distance to expel the hot air.** With a power rating of 9000 RPM, each fan consumes 13W, rather than the 18W consumed by each 7500-RPM fan in a rack-mounted chassis. This lowers the power consumption of each fan pack from about 180W to 52W, a significant savings, **even more significant when you compute the savings from 42 chassis in the rack!**

The single power supply provides either 375W or 900W of power to the node, and is backed-up by the 41 other nodes within the rack that are able to pick up the workload whenever

the application software makes that determination. A standard 1U server will typically contain 2 to 4 power supplies, requiring significantly more energy, and will draw anywhere from 222W to 288W, while an iDataPlex server uses from 138W and to 195W, depending on whether the server is using a low-voltage or standard processor.

The modular design of each node, with front access, provides for tool-less and simple serviceability. Each chassis is designed with blade-like technology, docking into a power connector on the backplane. The enterprise can mix and match the 2U and 3U chassis in any combination with swappable server trays. Using 2U nodes, with two servers in each node, the data center can install 84 servers within a standard 19" rack. Moreover, IBM's factory will pre-configure whatever combination of chassis and storage that the data center requires.

The iDataPlex node can contain an IBM *System x iDataPlex dx340* server with a power-optimized dual- socket server with a dual or quad-core *Intel Xeon* processor. Each dx340 supports up to 64GB of memory with a PCIe connector. IBM can also provide a *System x iDataPlex dx360* dual socket server with a pair of high-performance, quad-core Intel Xeon processors, and a PCIe connector. Alternatively, the data center can request a 2U drive tray with five 3.5" drives or eight 2.5" drives, a 3U drive tray supporting twelve 3.5" drives, or a 2U I/O tray supporting two full-height PCIe slots for maximum component flexibility. These options allow the data center staff to pre-configure iDataPlex for a compute-intensive environment, one that is I/O-rich with storage, or one that provides a storage-rich compute environment. iDataPlex also comes with intelligent systems management software to manage health alerts, inventory, and power utilization.

Conclusion

Quite clearly, iDataPlex is the enabler to any enterprise wishing to deploy a high performance computing solution or Web 2.0 architecture, but fearful of the amount of energy and floor space that a Web 2.0 solution entails. iDataPlex provides the foundation, or platform, for improved services to Web users everywhere that the megaplex can build upon. Installing iDataPlex in a megaplex data center enables the IT staff to roll out thousands of servers at a time in increments of up to 84 servers within a single 42U rack. This enables the staff to configure the megaplex in giant steps, within a small footprint, and an even smaller carbon footprint. Further, we have actually seen iDataPlex in action, as IBM has demonstrated this capability for a select set of customers, as well as for the analyst community.

With 240% better density than a traditional rack server, iDataPlex uses less than 50% of the floor space, with 40% lower fan power. This translates to energy savings of up to $10K per rack, with a savings of up to $1.2M for a typical iDataPlex data center, lowering the infrastructure TCO significantly. **This does not even include the potential savings from the installation of the rear door heat exchanger.**

When you add to that the savings achieved by the enterprise by receiving fully-configured and optimized racks from the factory, iDataPlex becomes an even better bargain. If you are living and breathing in the new megaplex universe (and, thus if you are deploying thousands, or tens of thousands, of servers every month, then you need to review the many ways that iDataPlex can improve your performance and lower your TCO for IT. If this shoe fits, you need a pair … or ten … or 100!

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

David Reine **is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" when you hear the automated attendant.)***

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.