THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2004069 August 23, 2004

Bull Announces New High-Availability Options For Escala POWER5 Technology

Analyst: David Reine

Management Summary

When the average consumer goes shopping for a new automobile, he is concerned about many things: the price of the vehicle, the gasoline consumption, the acceleration, seating capacity, and, perhaps, cargo space. A very educated consumer and car buff might be concerned with the horsepower under the hood. This consumer might not be surprised to discover that General Motors might not make the engine in that beautiful Chevrolet. Toyota could have made it, or perhaps BMW. Moreover, the tires? You can be sure that General Motors did not make them. Michelin, perhaps, or maybe Bridgestone. They might even be from Goodyear. In the 21st century, major corporations prefer to stick to what they know how to do very well. If the automobile manufacturer has no expertise in the fabrication of tires, they should OEM from someone who does. After all, why try to reinvent the wheel!

Apple Computer does the same thing with the CPU in their *Power Mac G4* or the *Xserve G5*. Apple does not invest millions of dollars into the design and fabrication of microprocessors. They procure the *PowerPC* CPU from IBM as an OEM part. They then wrap that part in a unique skin, i.e., operating system and peripheral set, and market an Apple solution to their customer base. Groupe Bull has teamed up with IBM for the past decade and a half in a similar venture. Starting with an agreement to co-develop a new *UNIX* platform and operating system that would be manufactured and branded by each, Bull has transitioned its engineering team to the refinement of elements of the operating environment. This includes elements such as workload management, and the server-surround environment. Bull now OEMs the basic *Escala* server platform from IBM, using the IBM *pSeries* and *p5* servers and *AIX* operating system as the core for the development of solutions tailored to their customer base. They also have implemented OEM agreements with a variety of other manufacturers who provide them with peripherals and software to be tailored into Escala solutions. They do not need to invest in the development of microprocessor or operating system technology. **They can devote their engineering resources to improving high availability and for storage infrastructure solutions.**

In their most recent announcement, Groupe Bull not only presented four new server models based upon the POWER5 microprocessor from IBM, but they also adapted a new low-cost storage solution from EMC for use with Escala. In addition, they introduced a new version of a low-cost high-availability application from their own engineering facility. To learn more about the new Escala solutions, please read on.

IN THIS ISSUE

Enterprise IT Solution Requirements

The consolidation and reliability of Information Technology (IT) resources remain as two of the greatest concerns for today's largest enterprises. These enterprises require better efficiency and higher availability from their mission-critical servers, not only in the data center, but also throughout their network. From the branch to the department, from the department to the data center, the CIO is constantly looking for a better way to implement e-business applications. The CIO is always trying to improve resource utilization for solutions, to secure Web transactions and to improve availability to corporate databases.

The CIO is also trying to control costs and reduce complexity, especially in the area where growth is the most prevalent, storage. Due to mergers and acquisitions, industry standards, and new government regulations, more data is being stored today than ever before, with more requirements for additional storage rising every day. Some of this information needs to be on the most performant storage devices available in order to provide instant access in a storage area network (SAN) environment. Information Lifecycle Management (ILM), however, dictates that other information may not be as urgent because of the change of the impact of its content. This information can reside in a secondary level device, as opposed to the primary device that carries premium functionality at a premium price.

Unfortunately, the costs associated with ensuring access to mission-critical data have risen dramatically. From a server standpoint, the costs, and complexity, of using high-availability software tools have placed a burden on the data center's budget. This is true not only from a licensing standpoint, but also from a human resource requirement. **High-availability utilities require a highly skilled and disciplined staff of administrators. They must be trained in the utilities themselves, as well as the operating system**. These costs are so prohibitive that many times the high-availability access cannot be distributed down to the branch or departmental level where they are needed. Implementation is limited to the data center.

IBM has addressed part of this problem with the introduction of a truly superior solution based on the POWER5 platform[1]. Groupe Bull, IBM's development partner in both POWER and AIX technology for more than a decade, has completed their HA solution with the development and integration of a unique set of high-availability applications for the POWER5 Escala platform. They have addressed the storage aspects of the solution working with their storage partner, EMC[2]. Group Bull has not only integrated the *CX* family of open systems, primary storage arrays from EMC, but has also integrated the newest tier of storage from EMC, the *AX100*, as a low-cost solution. In addition to the POWER5 platforms, Bull also designs proprietary servers (*GCOS* mainframes) and Itanium-based servers (branded *NovaScale*)[2].

Before looking at Bull's server-surround strategy, let's review the POWER5 capabilities and the servers that Bull has implemented to deliver the performance to consolidate and secure a high availability environment.

New Escala Platforms

Starting with the new, innovative *POWER5* microprocessor introduced by IBM, Group Bull's new servers maintain binary compatibility within their highly-performant *Escala* family. Bull's introduction set of four servers enables virtualization within the processor in order to improve the efficiency of UNIX/Linux servers. Bull allows a mix of applications to share the same processor and resources through the virtualization capabilities. Some of these virtualization features are:

- Simultaneous Multi-Threading (SMT) to increase AIX 5L V5.3 throughput and transactional performance by up to 40% ;
- Automatic dynamic engagement of inactive on-demand processors based upon pre-defined enterprise objectives;

[1] See **The Clipper Group Navigator** dated July 22, 2004, entitled *Open Systems and Virtualization – IBM Takes Another Stride with POWER5* at http://www.clipper.com/research/TCG2004063.pdf.

[2] See **The Clipper Group Navigator** dated August 31, 2003, entitled *Groupe Bull Mans the High Performance Point While Leading Migration Charge to Itanium* at http://www.clipper.com/research/TCG2003040.pdf.

- Enhancements to the Hypervisor to manage the micro-partitioning capability in Dynamic Logical Partitioning); and
- Cross-partition policy-based workload management.

Escala does not disrupt the customer's mission-critical operations by transitioning to a brand new architecture. Combined with AIX, POWER CPUs have delivered over the past decade the consistent performance required to keep Escala servers in the forefront of IT performance. Overall, based upon the benchmarks run, **POWER5 appears to have anywhere from two to five times the performance of competing processors, depending upon the benchmark chosen**. Some of the more significant improvements made to the POWER5 architecture are for virtualization, workload management, and reliability. (See Exhibit 1, at right.)

Through partitioning, Escala can run *AIX* 5L V5.3 and *Linux*, on the same system, on the same processor, at the same time. In fact, the data center can execute multiple copies of AIX, enabling the execution of production software and development/test at the same time. AIX 5L V5.2 does not support micro partitioning, SMT, or virtual I/O.

Bull has introduced four new Escala models with POWER5. The model numbers and their respective relative performance are shown in Exhibit 2, at the bottom of the next column. These models compare quite favorably with their POWER4 brothers, with rPerf values range from 4.41 for dual processor servers to 8.69 for quad systems, or better than 2:1 for the POWER5 over POWER4.

Any of these new Escala models can be clustered in a data center with Bull's *Application Rollover Facility (ARF)* or IBM's *HACMP* for high availability. In addition, they can all be configured with any of Bull's storage solutions from EMC[2].

Escala PL250R/T

Designed as a low-cost development and deployment platform, the Escala PL250R/T can be configured as a mono-or dual-processor branch office server, small database server, or an e-business platform. With 1.65GHz POWER5 CPUs, the Escala PL250 can be configured in a rack or as a tower with up to 32GB of memory and 8.2TB of internal disk. As a consolidation solution, the 250 can support 20 unique partitions across the dual-processor configuration as "virtual servers".

Exhibit 1 –
POWER5 Enhancements

- Simultaneous Multi-Threading (SMT) to increase AIX 5L V5.3 throughput and transactional performance by up to 40% ;
- Automatic dynamic engagement of inactive on-demand processors based upon pre-defined enterprise objectives;
- Enhancements to the *Hypervisor* to manage the micro-partitioning capability in Dynamic Logical Partitioning);
- Cross-partition policy-based workload management;
- Virtual storage and I/O in a micro-partition environment;
- Hot I/O drawer addition/removal; and
- Dynamic power management for CPUs.

Escala PL450R/T

Designed as a multi-function departmental or regional server, Escala PL450R/T can be configured in a rack or as a tower. With expandability for four POWER5 CPUs at 1.65 GHz, the 450 can support 64GB of memory with 15.2 TB of internal disk storage. With four CPUs at its disposal, the Escala PL450 can enable up to 40 virtual servers, or partitions, for use by the data center.

Escala PL850R/1650R

With capacity for either eight (PL850R) or 16 CPUs (PL1650R), these Escalas are well

Exhibit 2 –
Escala Models/Performance
Relative performance in rPerfs

	Dual	Quad	8-way	16-way
PL250	9.86			
PL450		19.66		
PL850			34.46	
PL1650				77.45

Source: Groupe Bull

Exhibit 3 –

The High Availability pillars on Escala

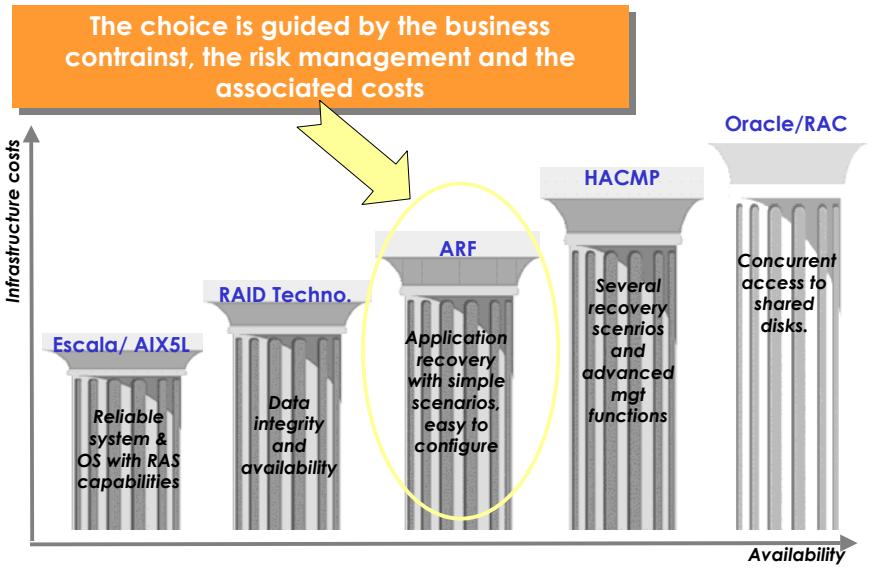

Source: Groupe Bull

suited for server consolidation projects, database and application services, e-commerce, or for department or regional deployment. Each node can support 256GB (PL850R) or 512GB (PL1650R) and 15.2TB of internal disk as a rack-mounted node. The PL1650R can also support up to 160 virtual servers with a shared processor pool allowing for non-disruptive load balancing between the partitions, resulting in better system utilization.

Bull's Added Value

It is on top of the server that Bull has introduced the greatest flexibility to Escala POWER5. **With v4.0 of ARF, Bull adds a low-cost, lightweight option to meet high-availability requirements, an option not available anywhere else**. (See Exhibit 3, above.) Moreover, with the AIX port of the AX100, Bull introduced EMC's new SATA storage solution to UNIX, an environment that EMC has ignored for this platform.

Application Rollover Facility 4.0

Up to now, AIX-based enterprises desiring to increase service levels and uptime for mission-critical applications had to use *HA/CMP*. This high-availability application from IBM is a very sophisticated utility and requires a highly skilled and disciplined administrative staff, fluent in both HA and AIX, and, often, expensive services. This adds cost to an already expensive HA solution: expensive to license, expensive to maintain, expensive to service. As an alternative, **Bull has now introduced a product which is inexpensive, easy to learn, easy to use, and, especially, easy to maintain**. ARF is the flexible and cost efficient tool to manage applications between nodes in a cluster, or between partitions in a virtual server environment, and is fully qualified with AIX and the shared storage solution from EMC to:

- Check the status and availability of each node through a heartbeat mechanism;
- Provide for either automatic or manual intervention for application relocation;
- Monitor activity through a web-based interface using *ARF Watch*.

ARF increases application uptime, while at the same time reducing both the licensing and management cost of HA. Configurable in both automatic and manual failover environments, ARF can also be used as a workload management tool, moving applications from saturated nodes to processors with free resources. Moreover, ARF can reduce planned downtime and facilitate system and application software upgrades, by rolling applications to

alternate nodes during the upgrade process. ARF is limited to an Ethernet network and does not support NFS, but does support an unlimited number of nodes. HACMP is limited to 32.

Fully compatible with HA/CMP, ARF can reuse all HA/CMP scripts and SAN infrastructure. Available for both local and remote configurations, ARF protects nodes, disks, the network, and all mission-critical applications in a no-single-point-of-failure architecture. In a remote configuration, ARF works within the limitations of the enterprise's Fibre Channel network to move an application within seconds, taking advantage of multiple disk subsystems and EMC's *Mirrorview* software. ARF can operate in a point-to-point fashion using serial protocols for dual-node clusters or using an IP protocol in multi-node, i.e., unlimited, clusters. With ARF, the enterprise can make efficient utilization of all hardware resources and ensure access to all active nodes.

CLARiiON AX100

In addition to qualifying the CLARiiON CX family of RAID storage arrays, Groupe Bull has also completed an industry-first: the port of the AX100 SATA[3] array onto a UNIX server. **No other vendor offers EMC's AX100 for a UNIX system, not even EMC**. Primarily designed as a secondary storage tier in an ILM environment, the AX100 can provide primary storage for entry-level branch servers. There, the frequency of disk access may not be as great or as critical as in a data center. In the branch, however, cost may be very critical. With the performance of a mono-processor server like Escala now available, the cost of an entry level IT solution with the AX100 becomes very attractive to the IT Director looking to distribute processing to the next tier down. The AX100 also comes with pre-packaged software to handle basic management, replication, path failover, and balancing.

With a limitation of 3TB, the AX100 may not be the high-capacity solution for everyone. However, the availability of a single-processor configuration within a 2U format, with Fibre Channel connection to up to eight hosts, does provide a low-cost solution for many.

Conclusion

In a world of commodity CPU's and commodity servers, solutions have become the key to success more than ever before. The IT community cares less today about the CPU at the core of the server and more about the actual problem that the investment will fix and how soon the ROI will be realized. Groupe Bull has adopted a strategy with the IBM POWER5 architecture as its platform, but with Bull engineering expertise at its heart.

Using the expertise gained from 15 years of intimate knowledge working with high-availability UNIX servers, Bull has developed with ARF a unique, "lightweight" version of their more resilient HA/CMP solution. Designed for implementation in entry-level environments where the cost of a full HA implementation could not be justified, ARF provides the right level of support. It can keep the branch on the air without burdening the configuration with unnecessary functionality. Moreover, it saves the enterprise money.

With data creation exploding around us, Bull has again used its engineering experience to implement the first and only AIX port of EMC's AX100 RAID array. A primary storage solution for the branch, this drawer can be used as a secondary solution in an ILM environment. In addition, Bull can attach the AX100 to any POWER5 box. **Any reseller or end-user interested in integrating a POWER5 server with an EMC array, CX or AX, into an HA environment should be looking to Bull for integration guidance**

If your enterprise has plans to rollout a network of distributed servers, take a close look at the list of vendors you had prepared for the RFP. You probably had HP, IBM, and Sun. You might have had Dell there, as well. You could do no harm - and perhaps a lot of good for your enterprise - by adding Groupe Bull to that list for consideration.

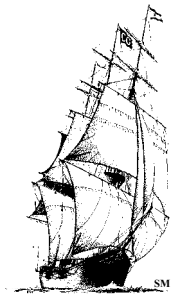

[3] See **The Clipper Group Navigator** dated May 26, 2004, entitled *For Small Scale Storage Consolidation, EMC Serves Up the CLARiiON AX100* at http://www.clipper.com/research/TCG2004050.pdf.

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. 23. (Please dial "1-23" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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.