

WHITE PAPER

IBM's OpenPOWER Servers Address Price and Performance in the Volume Server Market

Sponsored by: IBM

Jean S. Bozman
September 2004

SITUATION OVERVIEW

Driven by low acquisition costs and ease of deployment, IT professionals have turned increasingly towards "scale-out" architectures, or "horizontal computing," using low-cost servers with four, or fewer, processors on each server system. This trend, which involves increasing capacity incrementally by adding smaller servers to the network, accelerated during the economic downturn of 2001-2003, when IT budgets were limited, and IT managers were looking for new ways to add capacity to existing IT infrastructure. Now, volume servers continue to drive growth in the server market, as platforms for IT infrastructure, and for technical and commercial workloads. As a result, the high-volume segment is driving the fastest growth in the worldwide server market in 2004, according to IDC server research.

At the same time, Linux has gained a significant presence in technical computing environments and is quickly gaining acceptance among commercial users. The rapid growth of Linux servers has made the Linux server segment the fastest growing market segment within the worldwide server market, according to IDC's quarterly and annual server data. In fact, in Q104 Linux server revenue increased by 60%, compared with the year-ago period, while unit shipments grew by 46%. Looking forward, 64-bit computing in the Linux world will be attracting even more attention because of its ability to handle greater workloads and larger datasets, benefiting both high-performance computing workloads and database workloads .

IBM's Introduction of OpenPower Servers

Aiming squarely at the intersection of these growth trends, IBM has introduced its IBM eServer® OpenPower™ servers, which combine support for the Linux operating-system kernel and support for IBM's POWER5 fifth-generation RISC microprocessors. It is the convergence of these two components—an enterprise-ready Linux kernel and a 64-bit hardware platform that is tuned to provide high performance levels for Linux—that offers customers a new, high-performance alternative for volume servers running Linux. IBM OpenPower servers will run Novell SUSE Linux and Red Hat Linux distributions initially, and will support additional Linux distributions over time.

These high-volume servers are the newest addition to IBM's cross-platform Linux support, which extends from the x86-based IBM eServer xSeries, eServer

BladeCenter bladed servers, to POWER5-based pSeries and iSeries POWER servers, and finally to Linux support on zSeries mainframes.

THE OPENPOWER VALUE PROPOSITION

The core value proposition of the OpenPower line is that it provides Linux on POWER5 at a volume-server price. The OpenPower systems will compete in the volume server market, which is defined by IDC as servers priced less than \$25,000. The expectation of IT managers is that, by acquiring volume servers, they can reduce hardware acquisition costs, even if multiple servers are to be deployed in scale-out clusters or grids.

The OpenPower series of servers are tuned to run Linux on IBM's POWER5 architecture and provide enterprise-class reliability, availability and serviceability (RAS) features. This will support consolidation of Linux workloads, resulting in IT simplification compared to Linux server farms that require a much larger number of individual server nodes to support the same set of workloads. Importantly, the micro-partitioning feature of POWER5—a new, optional feature supported by IBM's POWER microprocessors—will support up to 10 partitions on a single chip. In practical terms, this means that up to 40 partitions could be housed on a 4-processor OpenPower volume server, easing workload consolidation and improving server utilization compared with other volume server platforms. The consolidation of many Linux server workloads that had previously been running on uniprocessor or dual-processor systems onto a single OpenPower Linux-on-POWER system will demonstrate the performance-density characteristics of the OpenPower server platform.

Key elements of the value proposition for the OpenPower line include the following:

- ☑ The OpenPower is a hardware platform tuned specifically for Linux. Fast access to memory and data, supported by the POWER5 microprocessor, is designed to improve performance, and reliability
- ☑ The platform is built on Linux open standards, and it supports other open application interfaces such as those used in Internet and grid computing, as well as PCI/X I/O links (along with other I/O standards). This support for open standards will encourage the adoption of Linux on POWER systems in IT sites that have already adopted these open standards, resulting in the development of a wide range of applications for this new platform. By year's end , more than 800 ISV applications will have been tested and supported on the OpenPower platform, and more than 1,000 are expected by 2H05.
- ☑ The ability of OpenPower systems to support both 32-bit and 64-bit applications provides flexibility for businesses to deploy a single architecture to run their 32 and 64 bit applications. Dual support for 32-bit and 64-bit workloads will allow organizations to transition from 32-bit computing to 64-bit computing at their own

pace—and to keep using 32-bit custom applications on the new platform as long as they wish to do so.

- ☒ OpenPower is designed to improve server utilization by means of virtualization of on-board resources. OpenPower's support for micro-partitioning will allow many workloads to run simultaneously on a single system. This will support workload consolidation, improve system utilization, isolate workloads within micropartitions for security purposes—and prevent any one application from interfering with another.
- ☒ Organizations can increase uptime for their volume servers through the use of OpenPower's RAS features, and by using high-availability fail-over and workload-balancing clustering software that are included in the IBM eServer's software stack. This multi-dimensional approach to high availability will ensure access to business-critical Linux applications running on OpenPower servers.



THE OPENPOWER PRODUCT LINE

The eServer OpenPower family of servers includes single-processor systems, dual-processor and four-way systems. IBM has announced the first model, the four-way capable OpenPower 720, with a base price of \$5,000, ranging up to \$16,000 for a configuration with more processors and memory.

The OpenPower 720 will ship both as a tower or 4U rack-optimized form. The server is based on a 64-bit POWER5 RISC microprocessor running at 1.5 GHz or 1.65 GHz with Novell SUSE Linux or Red Hat Linux. The OpenPower 720 features maximum memory of 64 GB, 5 PCI-X I/O slots, and 8 bays for Ultra320 SCSI drives. Three-year warranty and support contract are included.

The OpenPower 720 supports the POWER5's optional Advanced OpenPower Virtualization feature (by statement of direction; planned availability is 4Q04), which includes micropartitioning, Virtual I/O, and Virtual LAN—all of which enable virtualization across computing resources. The Micro-Partitioning capability allows each POWER5 processor to host up to 10 individual, isolated partitions, for purposes of workload consolidation of workloads that otherwise would need to be managed across multiple small servers.

Business Benefits of this Approach

The OpenPower platform can be deployed quickly to run custom or packaged (off-the-shelf) applications. Quick deployment can reduce IT costs, especially those for IT staff time associated with configuring, deploying and operating RISC-based servers.

Historically, IT organizations have been the "integrators" of computer technology, working to improve the "software stacks" that run on top of specific hardware platforms. System vendors that can pre-configure systems, or build pre-integrated

software stacks, are reducing the amount of time and effort that would otherwise be spent in the deployment phase of a new IT project technology roll-out.

IBM is providing a high degree of integration on the software stacks it offers for the OpenPower systems, reducing the IT staff time associated with development, and speeding deployment of production-ready systems. IBM is also offering a range of choice for the software stacks that can be deployed, depending on the types of workloads to be supported on the servers.

Management is also a key factor in improving efficiency of operation, and reducing IT operational costs. Customers' ability to manage multiple OpenPower systems, by means of a unified system management framework, will make IT operations more efficient. IBM will leverage the IBM Virtualization Engine and IBM Tivoli system management software to automate the transfer of workloads across a large number of rack-optimized OpenPower systems. This is part of IBM's overall on demand computing initiative, which will tap the resources of many computers through the virtualization of the hardware environment and the provisioning of workloads across those hardware resources.

Micro-partitioning, an optional feature of POWER5 microprocessors, allows many small workloads to be gathered together for processing on a single server, thus improving system utilization by running a number of separate, and logically isolated, workloads in a single hardware resource. Traditionally, IT customers have been concerned that their computer resources are not being fully utilized, and workload consolidation and improving utilization levels remain top priorities for many IT organizations.

OPPORTUNITIES AND CHALLENGES

IBM has the opportunity to provide improved price/performance to the volume server marketplace by addressing both sides of the price/performance equation: by offering high-performance servers that are priced aggressively to compete with other vendors' volume servers. There clearly is opportunity in the fast-growing volume server and Linux server segments of the server market, and IDC believes this volume-server-market opportunity will increase as IT spending broadens. Based on its specifications and launch pricing, OpenPower appears to offer impressive performance and enterprise capabilities at attractive price points. These attributes, combined with IBM's aggressive efforts in attracting ISVs and building the Linux on POWER ecosystem with multiple IBM eServer products, will place OpenPower servers within a comprehensive portfolio of IBM Linux servers, and will position IBM to capitalize on the volume-server industry dynamics. For OpenPower, IBM is now driving from 500+ applications to more than 1,000 applications by year's end.

A New Entrant in the Volume Server Space

However, it is clear that, as a new entrant to the volume server marketplace, IBM's OpenPower server platform will face intense competition with widely shipping volume servers that are based on other hardware architectures, including x86, x86-64 and

other RISC microprocessors. Clearly, IBM will be competing with vendors of x86 servers (e.g., HP and Dell) and RISC servers (e.g., HP and Sun) in the volume server marketplace.

Not only will IBM face price competition, but it will also face competition for the attention and investment of the ISV community, which must adapt its existing applications for the OpenPower platform. IBM recognizes this challenge, and it has set up 25 centers worldwide, including sites in North America, Europe and Asia/Pacific, for ISVs working on porting, recompiling and testing their applications for OpenPower. IBM is also making investments in equipment, office space and software to ensure that ISVs will consider its new server platform for existing packaged applications, including those for specific vertical markets (e.g. financial services, retail), and for emerging applications, including open-source applications from the worldwide Linux community.

For customers and value-added resellers around the world, IBM must make the case that its new server platform provides significant benefits while maintaining, or improving, the current cost model for deploying and operating Linux servers that are based on other hardware architectures, such as x86 or x86-64. Finally, IBM is a vendor that also sells volume servers based on x86 and x86-64 architectures (both Intel Xeon and AMD Opteron implementations), and so it must present the full spectrum of volume server choices to its enterprise customers. This is part of the IBM eServer product-line's over-arching value proposition and marketing position: namely, that IBM's platforms support many workloads on many servers—each with its own price/performance characteristics.

IBM has invested in developing the OpenPower platform—and in bringing it to market. It stands to benefit by deploying an IBM-built hardware platform on which IBM middleware can play, and for which IBM services can be provided. To the extent that open standards are supported, then the support of a wide range of workloads (including open-source applications, custom Linux applications and ISV packaged applications) on this new platform will be easier to explain—and to accomplish.

Looking forward, when OpenPower is deployed into pure “infrastructure” computing tiers, where they are largely invisible to end-users, IBM will benefit by focusing on a technical sale, focusing on IT managers for marketing this new, highly tuned, hardware/software platform for network computing. When OpenPower is deployed into commercial computing environments, IBM will benefit by focusing on the system's business benefits. Further, IBM can provide “proof-of-concept” deployments in selected vertical markets, such as telecommunications, retail, and financial services—all markets in which Unix servers and Linux servers are already well-accepted and widely deployed—and in which Unix and Linux skill-sets are widely available.

CONCLUSION

Linux server systems, many of them volume servers are gaining adoption for technical and commercial workloads. The growth rate of Linux server revenue and Linux server shipments accelerated during 2003 and 2004, driven by deployments of these servers in Linux clusters and by an expansion of the overall Linux ecosystem, which is bringing more packaged software to the market.

IBM's OpenPower Linux products are designed to compete in the high-volume marketplace for small servers that can be arranged in server farms, clusters and grids to build flexible IT infrastructures that can provision new workloads, directing them to available computing resources, as needed. IBM will sell these servers across many vertical market segments, with industry-specific workloads being deployed across the OpenPower hardware platform in a variety of configurations.

With the OpenPower offering, IBM's support for the enterprise-class Linux 2.6 distribution, its support for virtualization and the server's robust design and performance will attract attention in the marketplace. This will be especially relevant as Linux workloads expand from edge workloads and IT infrastructure workloads to the wider world of enterprise workloads, including business-critical and mission-critical applications and databases.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2004 IDC. Reproduction without written permission is completely forbidden.