



Research Report

Linux: IBM's Workload Optimization Growth Engine

Executive Summary

If you've ever wondered why IBM makes three distinct server platforms (System x, Power Systems, System z), it is because servers process workloads (applications/sets of applications) differently. For instance:

- x86 servers are most cost effective when used to process numerous, heavily-threaded *light workloads* (multiple, fast, low-cost threads) that require light memory use, and that require comparatively lower quality-of-service (QoS) than offered on System zs and RISC machines;
- When it comes to *heavy, 64-bit compute-intensive or transactional workload processing* dozens of benchmarks show that Unix/RISC servers excel at this type of workload; and,
- When running *heavy input/output workloads* (transactions with many reads to memory, disk, and the network, and many writes to disk), System zs are a far better choice than either RISC- or x86-based servers.

Enterprises that deploy the right workloads on the right systems can save BIG MONEY when it comes to computing costs (see our recent report that describes how an enterprise can save over a million dollars by running certain Linux workloads on a large consolidation server at as opposed to a collection of x86 multi-cores at: http://www.clabbyanalytics.com/uploads/z_VirtualizationFINALRev.pdf).

If you have ever wondered why IBM offers five different operating environments (z/OS, zVM, AIX, IBM i, and Linux), it is because some operating environments are geared toward transaction processing or virtualization management or compute intensive computing, while others have been designed for general purpose computing. Of these, the Linux operating environment is particularly interesting because it is a general purpose operating environment that can operate on all three of IBM's platforms. And because of this, information technology (IT) buyers realize two distinct benefits when deploying Linux on x86, Power Systems, and System z servers: 1) a consistent operating across all three servers is easier to manage (skills can be leveraged across all three environments); and, 2) applications/workloads can be run across all three platforms (enabling an IT buyer to choose whichever platform can most efficiently execute a particular workload).

What most people don't understand about Linux at IBM is that Linux has become a "growth engine" for future application/workload environments. Independent software vendors (ISVs) now offer thousands of applications on Linux — and these applications now run on x86, Power System, and System z servers. IBM's future depends on Linux to run today's workloads — while capturing the evolving workloads of tomorrow. This is why we say that Linux has become "IBM's workload optimization growth engine".

In this *Research Report*, Clabby Analytics takes a closer look at IBM's Linux market positioning, and at how Linux is being deployed on IBM's three server environments.

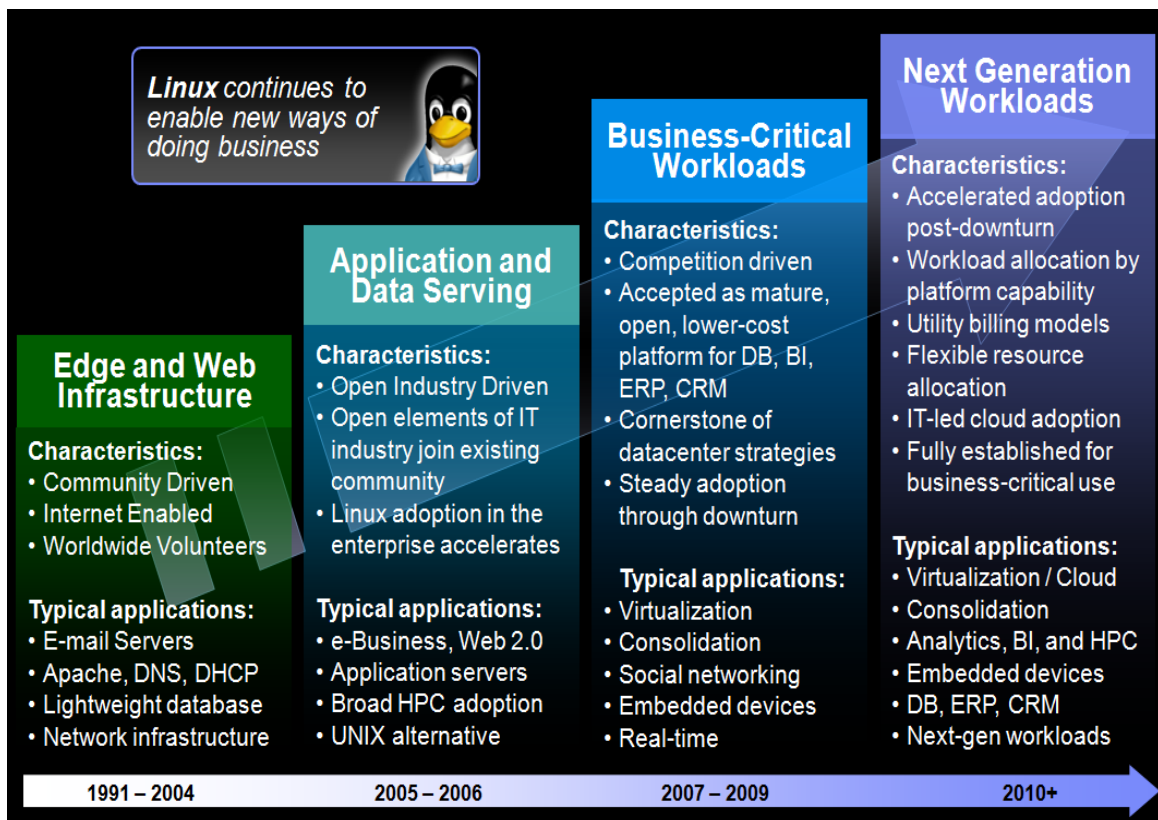
Linux: IBM's Workload Optimization Growth Engine

The Linux Market: Dominated by x86 — and IBM Has to Change This...

The top issue/challenge that IBM must overcome when it comes to Linux is that 90%+ of all Linux server installations take place on x86 architecture. What this means is that the market has been conditioned: “Linux = x86”. And to change this well established buying pattern, IBM must shift the discussion to: “*workload + Linux = platform choice*”.

What IT buyers should know is that Linux started out as a personal computer operating environment. It subsequently made its way into mobile computing and onto servers. In the server world it started with simple applications: file and print, firewalls, Web applications, and the like. But in recent years, as the operating system has become more robust (reliable, available, and secure), Linux has become capable of handling challenging mission-critical business applications. (For more details on this trend, see our report based-upon IT executive interviews at: http://www.clabbyanalytics.com/uploads/CIO_Linux_Outlook_2010.pdf. Already started, it is reasonable to expect that Linux will host a slew of next generation workloads — including cloud-based workloads, business intelligence, business analytics, high-performance computing, and other applications that want to exploit 64-bit computing and large databases (see Figure 1 for a historical/future representation of Linux direction).

Figure 1 — Linux Historical and Future Growth



Source: IBM Corporation — June, 2010

Linux: IBM's Workload Optimization Growth Engine

For IBM to overcome the biggest challenge of Linux (the Linux = x86 buying pattern), IBM will need to show the market that running Linux on System z or Power Systems is a better choice than running Linux on x86 architecture. Clabby Analytics has already shown how System z can do a far more efficient job of executing certain heavy I/O workloads as compared with x86 architecture (click on our URL mentioned on page 1). IBM will also need to prove that running Linux on Power Systems is the best choice for running parallelized, 64-bit workloads. And that Linux on IBM's System x servers (as opposed to x86 servers offered by Oracle or Hewlett-Packard or Dell) is more scalable, better at handling virtual machines, and ideally suited for handling large volumes of fast threaded "light" workloads.

To do this, IBM needs to educate the market on what "workload optimization" is — and why it is important. (Workload optimization shifts the Linux discussion to how to maximize investment in computer systems — and demonstrates that x86 architecture is not always the most cost effective platform).

How IBM Must Change the Market's Thinking: Understanding Workload Optimization

A workload is an application, or a group of applications, that individually or collectively can complete a unit of work. For instance, the processing of email can be viewed as an individual workload; while the processing of a transaction can involve several steps that include interaction with a sales order, verification with a credit card company, checking stock levels with an inventory application, and a shipping application. Processing this transaction can be viewed as a collective workload.

Systems handle workloads differently — and, due to differing processor/system/QoS characteristics, workloads execute differently on each system type. Workloads that exploit a given processor, the surrounding systems design, and a system's inherent QoS are able to run more "optimally" on some platforms than others. And the more optimally a workload runs, the less it costs to run a given workload. The cost differential can be substantial!

For example, our research shows that System z, with their fast processors (CPUs — central processing units) and with their multiple, additional communications processors are particularly strong at handling heavy I/O (input/output) workloads. By comparison, RISC architectures (reduced instruction set — such as IBM's POWER Systems (and presumably Oracle's UltraSPARC/T3 servers and EPIC explicitly parallel instruction computing) architectures such as Intel's Itanium) are ideal for processing compute intensive and 64-bit OLTP workloads. And Intel's x86 architecture is particularly strong at processing multi-threaded "light-and-fast" workloads.

A Closer Look at the Workload Processing Characteristics of IBM's 3 Systems Environments

IBM offers three distinct computer platforms: System x, Power Systems, and System z. The platforms differ by processor, by system design, and by the quality-of-service (QoS) delivered. In this section, we examine the characteristics of each platform — and discuss the growth opportunities of each platform based upon its workload handling characteristics.

System x — It's the System Design!

All major x86 vendors are working from the same CPU designs in the x86 server market. So the only differentiation that occurs at the CPU level is which vendors get to market first (IBM consistently gets its Xeon multi-cores to market before its other major x86 competitors). But major differentiation shows up when comparing IBM's System x system design to other vendors' systems designs. And this differentiation should lead IT buyers to prefer IBM System x alternatives to x86 solutions offered by competitors.

Linux: IBM's Workload Optimization Growth Engine

Here are just a few of the differentiators between System x and its competitors:

- IBM offers a specialized virtualization management ASIC (application specific integrated circuit — a specialized microprocessor) on some of its x86-based servers that offloads the CPU from having to process virtualization instructions (virtualization helps IT buyers get more processing out of their servers by pooling unused computing resources into “virtualization pools” by creating “virtual machines”). By offloading the CPU from having to do this work, the CPU can be focused on other tasks (number crunching, for instance). Fujitsu has similar functionality as part of its blade systems offerings — but no major x86 competitors do so.
- IBM eX5 enterprise systems, featuring the Intel® Xeon® processor E7 family, provide several features that differentiate IBM System x from other x86 solutions.
- MAX5 – Memory expansion with the external MAX5 memory option, decouples server memory from system processors to allow customers to optimize server performance by adding memory rather than buying additional servers. Memory capacity can expand up to 64 DIMMs standard and 96 DIMMs with the MAX5 memory expansion per 4-socket server.
- eXFlash – Flexible hot-swap storage with up to 8 HDDs or up to 16 SSDs with eXFlash technology provides solid state drive technology that delivers faster I/O, with greater density and improved reliability.
- FlexNode – Delivers the ability to re-deploy your server on a project-by-project basis for superior asset utilization and workload management. For example, a 4 socket server can be re-deployed to as a two 2 socket servers in order to match the server characteristics to the workload as it varies through a day.
- Scalability – Expand from a two-processor system up to four processors. Add a second system to create an eight-processor system. Start with two memory DIMMs and expand up to 192 with a dual-node system and two MAX5 memory expansion options. This flexibility allows customers to meet today's needs while providing an easy, cost-effective upgrade path to change your environment when needed.
- Another example of how a systems design can offload work from a CPU can be found in IBM's inspired eX5 memory management environment. Some System x servers include a specially designed memory management subsystem (microprocessor and software) that manages swapping data in and out of memory — offloading the CPU from having to perform all memory management tasks and thus freeing-up the CPU to do more processing work. This memory management subsystem is a system design characteristic that plays an important role in System x data handling and performance. And note again, this is a systems design function that accelerates overall systems performance — offloading the CPU from having to do work.
- IBM BladeCenter® Open Fabric is an integrated server I/O portfolio that offers high performance interconnects and management tools. Supporting open standards and industry interoperability across multiple I/O fabrics, including Ethernet, iSCSI, Fibre Channel over Ethernet (FCoE), Fibre Channel, InfiniBand and Serial attached SCSI (SAS), IBM Open Fabric is designed to help simplify blade I/O management and deployment. IBM Virtual Fabric has been expanded to include support across the System x product line, providing a single point of contact for consultation, sourcing and service.

Linux: IBM's Workload Optimization Growth Engine

- Workload optimized solutions providing a range of specialized x86 solutions. From the iDataPlex® to the Intelligent Cluster to optimized systems configured virtualization, database, ERP, and other core business applications, these integrated solutions help reduce complexity and cost. With System x's recent announcement of the SAP In-Memory Appliance, SAP HANA, IBM is again first to market providing business users the ability to instantly access, model and analyze business transactional and analytical data from virtually any data source in real time.

System x is offered in blade, rack, and tower configurations. No other x86 competitor can scale as high as IBM's System x. No other leading competitor has a configuration like System x iDataPlex (a highly integrated, self-contained rack environment that excels at Web application and high-performance computing processing). And no other x86 vendor offers as sophisticated a management scheme as is offered on System x (with Systems Director/Tivoli software) — as well as across System x, System z, and Power Systems using IBM's newly announced zEnterprise management/governance environment.

Power Systems — The Future is Analytics

The majority of Power Systems have, to date, been deployed using a Unix operating environment (IBM's AIX). But, at *Clabby Analytics*, we believe that demand for Linux is going to increase rapidly over the next few years, driven by IBM's huge investments in business analytics and market demand for less-expensive Linux-based solutions.

Power Systems excel at compute intensive data processing. The POWER microprocessor can run at up to 4.14 GHz — and can be expanded to 8 cores with integrated cache and memory controllers. It uses 45nm technology — so it is densely packed.

A closer look at the POWER7 design shows that several “workload optimization” features have been built in, including: Intelligent Threads, Intelligent Cache, Active Memory Sharing and Expansion, TurboCore and MaxCore modes, and others that dynamically adjust system resources to virtual machines based on workload requirements. And, with PowerVM enhanced virtualization with micro-partitioning, a single Power 795 can support up to 1000 virtual machines running any combination of Linux, AIX or IBM i — allowing serious scale up workload consolidation. Dynamic workload optimization is especially valuable in a cloud environment, where workload peaks and valleys across the cloud dictate ever-changing resources requirements.

A closer look at the Power Systems system design shows:

1. A 4th generation SMP (symmetrical multi-processing) fabric bus to pass data between processors, memory, and I/O subsystems;
2. 3rd generation multi-threading (and it features intelligent thread caching);
3. Multiple memory controllers (to manage memory for compute-intensive tasks);
4. DDR3 memory support (to ensure greater reliability and accuracy);
5. Enhanced GX system buses (to improve processing speed between resources);
6. on-chip L2/L3 cache (the more memory that can be brought closer to the processor, the faster a job can be processed);
7. eDRAM L3 Cache (specialized lower level cache to improve); and,
8. Industry standard I/O.

Linux: IBM's Workload Optimization Growth Engine

What this all means is that Power Systems are super-fast processors designed to crunch data in parallel fashion. To quote an IBM senior vice president, Dr. John E. Kelly III, "The POWER7 system is tuned for very rapid deep analytics of massively parallel problems." And what this also means is that, due to a highly-tuned processor as well as excellent memory management, Power Systems are particularly adept at performing high-performance computing (parallel computing tasks) as well as business analytics (data crunching) tasks, whether a scale-up or scale-out architecture is preferred.

To illustrate the potential for Power Systems in analytics, consider IBM's demonstration of its "Watson" analytics systems (that recently appeared on a quiz-the-expert television show known as "Jeopardy!"). As IBM's Watson (a scale-out system consisting of 90 generally available mid-sized Power 750 servers *running Linux*) competed with two all-time Jeopardy past winners (humans), Watson proved that a system can be designed that applies advanced data management and analytics to natural language in order to uncover a single, reliable insight. What Watson was able to do was to sift through a large back-end database and respond to queries in fractions of a second. Needless to say, Watson won...

What Watson showed is that Power Systems can be used to help enterprises make better decisions (as compared to manually sifting through data). It also showed that enterprises that are suffering from "data overload" can now rely on super-fast computers to mine their data. And both of these factors open up new potential for Power Systems in the areas of health care (diagnostic assistance and collaborative medicine), tech support (to help deal with and solve help desk queries and better run contact centers); and in knowledge management and business intelligence (where Power Systems can be used to generate meaningful insight from data — in real time).

IBM has invested \$4 billion in business analytics/business intelligence software research and development as well as acquisitions over the past few years. The company is very serious about establishing leadership in this market — and Power Systems running Linux can be expected to be the primary deployment platform for compute-intensive, parallel computing workloads.

System z: A Major Design Advantage

There is no other server in the industry that scales as efficiently as a System z when it comes to Linux workloads — and, accordingly, System z servers represent the best architecture for Linux consolidation in the industry. The advantage of System z when it comes to Linux consolidation stems from its:

1. *CPU efficiency* — The z CPU is the throttle that controls the amount of headroom needed (the more efficient your processor, the less headroom needs to be allocated). IBM's System z CPU is the fastest in the industry (clocked at 5.2 GHz) — and this makes a huge difference in the amount of work that it can process.
2. *Shared everything design* — The "shared everything" design point of large servers also has a major impact on the number of virtual machines that can be created. If you cannot easily share all of a system's resources, you lose the headroom advantage that large systems offer;
3. *High security levels* — System z servers offer EAL Level 5 security, the only commercially available server platform to do so. EAL 5 ranking can provide customers the assurance and confidence that they can run many different

Linux: IBM's Workload Optimization Growth Engine

- applications containing confidential data on one System z which is divided into partitions that keep each application's data secure and distinct from the others; and,
4. *Massive I/O subsystem* — In addition to the general processors, a System z has 14 SAPs (System Assist Processors) + up to 336 channel IO processors + 1024 logical channels. These additional processors give the System z a huge advantage in terms of processing I/O (no x86 server design has anywhere near this I/O processing power).

For example, due to these advantages, a System z server can run 240 virtual machines (240 logical servers) within the same chassis executing heavy I/O workloads within a high-reliability QoS profile. It would take *24 Intel multi-core servers — and over a million dollars more* — to execute this same workload! For more details on this System z consolidation server cost advantage, see our report at:

http://www.clabbyanalytics.com/uploads/z_VirtualizationFINALRev.pdf.

Linux in the Cloud

Up to this point, we have discussed Linux market buying patterns (and what IBM needs to do to change those patterns) — and we have examined some of the processor/system design differences between IBM's three server platforms (and how Linux can exploit these differences). But what we haven't discussed is how Linux will ultimately be deployed as a cloud architecture. In this section, therefore, we describe our view of homogeneous and heterogeneous Linux clouds.

Before proceeding with this discussion, allow us to offer a very simple definition of cloud architecture. Cloud architecture is essentially a way to automatically virtualize and provision resources — and to deliver computing capability using a number of delivery models (traditional computing delivery methods, software as a service models, platform as a service, etc.).

Homogeneous Linux Clouds

Clabby Analytics is seeing a lot of Linux cloud activity taking place where IT buyers are standardizing on x86 multi-core servers-only within their cloud environment. The perceived benefits of this deployment option is that a common operating environment can be deployed on a common hardware platform — and applications, databases, middleware, operating environments, and platforms can all be managed in a common, consistent manner. For some organizations, this approach makes a lot of sense (for instance, if a small cloud is being operated — or if there is a shortage of Unix or System z skills, leaving little other choice but to standardize on x86 architecture). But in many cases, this can be the wrong decision (read our report on why we think this is the wrong approach for Union Pacific at: http://www.clabbyanalytics.com/uploads/Union_Pacific_Final.pdf).

For this kind of Linux cloud environment, particular attention needs to be paid to the management of physical systems as well as virtual machines. IBM System x should be on the shortlist when evaluating homogeneous x86 cloud environments due to broad and deep physical systems management facilities (Systems Director) — and broad and deep virtualization infrastructure and management facilities (as part of IBM's Tivoli management line).

Linux: IBM's Workload Optimization Growth Engine

Linux Glue — An Integrated, Hybrid System z and Power Systems/x86 Blade Environment

One of the most exciting advances in computer design took place last August (2010) when IBM announced a configuration known as zEnterprise. This environment features a System z server as a central point of control — with high speed connections to bladed Power System and x86 servers (running in a special chassis known as a zBX). And Linux can be used as the “glue” to tie each of these environments together into a cohesive cloud.

Here's what is so special about the zEnterprise environment:

- *Management* — what IBM has done is essentially extend z governance (the security and resiliency features) as well as z management (physical and logical systems management) to tightly coupled blade servers. This is hugely important because it provides a common and consistent mechanism for managing a z and connected blade servers using the industry's most advanced management environment;
- *Virtualization* — the zBX blade environment comes with an installed hypervisor that works in tandem with IBM's best-in-the-industry virtualization infrastructure/-management environment. By coupling these two environments, IT managers and administrators can now shift their attention to workload management and workload balancing rather than focusing primarily on managing virtual server sprawl;
- *Governance/Security* — System zs have been around for five decades, and offer the most advanced governance and security capabilities in the industry. By extending System z governance across hybrid zEnterprise systems, IBM is now able to improve the security environment for blades. For IT executives concerned about security and risk management, this unifies systems security across the board
- *Energy Management* — advanced System z energy management tools can now work in tandem with blade energy managers, allowing IT managers and administrators to create a common view of energy usage, and allowing those managers and administrators to balance energy allocation; and,
- *Networks* — manifest in a 10GB Ethernet connection directly from the blade environment to the System z server, as well as a secondary management network. (Early adopters are finding that when they use this tightly-coupled network they are able to eliminate additional routers and firewalls — and the management and performance overhead that these devices represent. As a result, the connections between blades and System z servers become easier to manage — and decreased latency due to the elimination of these devices also results in higher performance).

In the future, we expect the trend to continue toward homogeneous x86-based Linux clouds. We also expect that more heterogeneous Linux clouds will be deployed when IT buyers come to understand that x86 architecture does not process all workloads optimally — and when these buyers realize how much money they can save by deploying Linux on workload optimized platforms.

Summary Observations

It should be clear from this *Research Report* that IBM is making huge strategic investments in the Linux operating environment. Linux has a specific role on each of IBM hardware platforms — and Linux serves as the “glue” for joining IBM's different platform architectures together under the aegis of cloud architecture.

Linux: IBM's Workload Optimization Growth Engine

The biggest challenge for IBM going forward is to change a Linux market buying behavior. Linux started in the x86 world — and the preferred buying pattern for Linux buyers has been to adopt Linux on x86 servers without evaluating Linux on other platforms. We believe that IT buyers can save really BIG MONEY by deploying workloads on different Linux “workload optimized” heterogeneous platforms (System zs, Power Systems, and System x servers) as opposed to opting for homogeneous x86-only solutions.

Is IBM executing well on its Linux strategy? In the x86 world, the market chooses between only two operating environments: Windows and Linux. The question is: what will make people choose Linux on IBM x86-based servers. And we believe that IBM's system design superiorities — particularly in virtualization and memory management — will sway the market toward System x. In the System z world, we are seeing the numbers of Linux System z servers configured as Linux consolidation servers rise steadily each quarter. And System z Linux servers are also being adopted to run new workloads. So in the System z space, the answer is “yes, we are seeing solid progress”. In the Power Systems space, we believe that more and more IT buyers are going to start evaluating POWER-based Linux systems due to IBM's increased investment in business intelligence and business analytics solutions that run on Power Systems. IBM's “Watson” is a high-performance computing environment that recently demonstrated on a television show called Jeopardy! the amount of processing power that can be delivered by Power Systems running Linux (systems that have been optimized for intensive parallel processing applications). We have yet to see the adoption rates for Linux rise on Power Systems — but we believe that they are on the cusp of being ready to do so.

As for clouds, IBM provides management products that allow its customer choice. Customers can build homogeneous clouds, or they can even build a tightly-coupled, highly integrated hybrid Linux environment.

After closely examining IBM's three Linux platform offerings — and after examining IBM's numerous cloud configuration options — we find no vendor with the breadth and depth of Linux-based hardware and software in the IT industry. IT buyers looking to build their future on Linux — and looking to save BIG MONEY by deploying the right workloads on the right servers — need to closely examine IBM's Linux offerings across all three of IBM's platform environments (instead of always defaulting to x86-only implementations).

Clabby Analytics
<http://www.clabbyanalytics.com>
Telephone: 001 (207) 846-6662

© 2011 Clabby Analytics
All rights reserved
March, 2011

Clabby Analytics is an independent technology research and analysis organization. Unlike many other research firms, we advocate certain positions — and encourage our readers to find counter opinions — then balance both points-of-view in order to decide on a course of action. Other research and analysis conducted by Clabby Analytics can be found at: www.ClabbyAnalytics.com.