



Using IBM Mainframe Virtualization and Linux to Consolidate Servers

*Prepared by:
Doris Benson,
Gaylan Braselton,
Cal Clark,
Bob Neidig,
Bill Reeder,
Len Santalucia,
Terry Weinberg*

Table of Contents

Introduction Page 2
Mainframe Hardware for Linux Page 2
Mainframe Virtualization: z/VM Page 4
Best Fit Middleware and Applications Page 6
Systems Management Page 6
System z TCO Page 7
Customer examples Page 9
Nexxar Group, Inc. Page 9
Panasonic Page 10
Nationwide Insurance. Page 11
Boscov's Page 12

Introduction

This paper is written for people who are not familiar with the IBM mainframe or those who think it does not apply to today's computing problems. In fact, the exact opposite is true. The mature virtualization capability of IBM mainframes, combined with its reliable, power efficient hardware, provides an excellent platform for modern applications including Java™ applications and applications that run on Linux®.

There are those who think that IBM mainframes are too expensive. There are mainframe customers who offer proof that this is not the case. Corporations can potentially save money on floor space, electricity, software, and networking costs by moving distributed Linux servers onto the mainframe.

The following sections describe IBM mainframe hardware and virtualization at a high level. They also describe applications that often run on Linux on the mainframe and the systems management advantages of the mainframe environment. There are also examples of savings in cost of ownership and operation, and most important, examples of customers who benefit from using Linux on IBM mainframes to solve business problems.

Mainframe Hardware for Linux

Today, the IBM mainframe can be used to describe a style of operation, applications, and operating system facilities. Here is a working definition: "An IBM mainframe is what businesses use to host their commercial databases, transaction servers, and applications that require a greater degree of security and availability than is commonly found on smaller-scale machines. IBM mainframes also possess the richness of function and systems management capabilities to run these applications concurrently at sustained high levels of utilization."

IBM's originally announced mainframe architecture in 1964 and has continued to evolve and add innovations such as more and faster processors, memory, and dynamic and autonomic capabilities to facilitate growth and provide problem management. Mainframes also claim a history of innovation and refinement in virtualization, large scale I/O, and unique clustering techniques designed to enhance reliability, processing power, and resilience. While continuing to become richer in features and functions, IBM mainframe computers remain highly stable and highly secure. In addition, the latest models can handle the most advanced and demanding client workloads, yet can still run applications that that were run in the 1960's.

Businesses today rely on IBM mainframes to perform thousands of transactions per second, support thousands of users and applications, manage terabytes of information in databases, and handle communications that constantly require increased bandwidth.

These attributes are generally accepted in the data processing community, but some critics believe that the ability to provide these qualities of service require a proprietary and often expensive operating system environment. "What about Linux?" they lament.

The answer is that IBM mainframes are designed to support the Linux operating system on native hardware or running on z/VM[®]. In 2000, IBM announced a special processor available on the System z family of products called the Integrated Facility for Linux (IFL). It has the same maximum performance as standard engines, but supports only Linux workloads running either natively, or on top of z/VM (IBM's virtualization software). Since the IFL dispatches only Linux workloads, adding one to a System z should not affect the model designation, or increase the charges for the traditional software running on the machine. This is true for software from IBM as well as software from many Independent Software Vendors. There are also models of System z machines that are populated only by IFL engines. Finally, the IFL has access to all the same system hardware capabilities as traditional engines.

Linux can also benefit from the System z virtualization capabilities, and IBM mainframes have multiple levels of virtualization. Among the hardware level virtualization capabilities are logical partitions or LPARs. A System z can be divided into up to 60 LPARs. Linux can run in logical partitions and communicate to other Linux partitions via an internal TCP/IP socket connection without using an external network. The second level of virtualization is software based and is called z/VM. z/VM provides the capability of running hundreds of Linux servers within an LPAR. Mainframe hardware has circuitry that assists and accelerates the execution of this z/VM virtualization layer.

Businesses today use IBM mainframes because of attributes like large capacity, scalability, and availability. They also enjoy the integrity, security, and autonomic capabilities. These attributes, combined with its system management strengths, make it a powerful Linux platform.

The IBM mainframe includes hardware to assist and accelerate certain encryption technologies. If Linux needs access to these technologies, and they are installed on the system, Linux applications can take advantage of them.

System z incorporates virtualization technology that allows many virtual servers to use the same hardware entity. For example, in a distributed environment, each server typically has its own I/O drawers for disks and separate network cards. On the System z, multiple virtual servers can use the same network and share paths to disk over the same fiber channel.

As only one example of the IBM mainframe's autonomic capability, it has built-in redundant hardware that will automatically replace a failing processor chip, route data around failing memory circuits, and return data via a different path than the requesting path if it is busy. This is done in a manner transparent to the users of the system.

System z has built-in circuitry called a storage protect key that allows the operating system to detect a storage overlay by a misbehaving application and terminate the application, rather than overlay data that the application is not authorized to modify, a capability that is critical to a server consolidation environment.

IBM mainframes support a large number of I/O paths and up to 4 Gigabit fiber channels as well as very high bandwidth between virtual servers within the system. They can scale from 1 to 54 processors with the ability to add processing capacity dynamically when needed.

Mainframe hardware reliability (often referred to as "mean time between failure" or "mtbf") can be measured in decades for IBM mainframes. This is accomplished by duplication of hardware like fans, power supplies, cooling technology, and by providing for the dynamic swap of device paths to provide access to key peripheral devices and networks.

As you can see, IBM mainframe hardware has extended its capabilities from the traditional data processing environment to the emerging Linux environment without betraying its 'bullet-proof'-type characteristics and is now seen as a very powerful platform for simplifying an IT infrastructure by moving distributed applications onto it. The result can be fewer servers operating at higher levels of efficiency and providing economies of scale in software costs, floor space, cooling, day to day power costs, and costs for UPS capacity.

Mainframe Virtualization: z/VM

z/VM is the premiere IBM mainframe software virtualization environment. It offers management of virtual machines running in mainframe hardware partitions. Historically z/VM was used by customers as the test environment for new operating system releases, application development and testing, and support for applications that required separation for other reasons. Today, in addition z/VM can also manage hundreds of virtual Linux servers running on a single IBM mainframe.

z/VM employs very sophisticated resource utilization algorithms for sharing resources and managing virtual machines. Weights and priorities can be assigned to guest machines to determine allocation of resources. As a result, high priority guests receive resources before low priority guests. Computing, processor, networking, and memory resources are completely shared by the guests as needed. This environment allows z/VM to actually over-commit resources. For example, z/VM can easily deploy 4 virtual machines whose requirements are 8

GB each using less than 32 GB of real memory for all. The same applies to the allocation of CPU resources. This allows z/VM to drive an IBM mainframe at very high utilization levels (up to 85%-90%).

As mentioned earlier, the engine that the System z uses to run z/VM with Linux is called an IFL (Integrated Facility for Linux). The number of Linux images being run per IFL generally ranges from 3 to 70 images, with the average being in the range of 30 to 45.

z/VM also provides an environment for easily creating and deploying new virtual machines. Using simple scripts virtual server images can be cloned, provisioned and brought up in the environment. They can just as easily be shut down.

z/VM fully utilizes the System z mainframe's channel capabilities to offload much of the I/O workload to these subsystems, thus freeing the CPU/memory system for additional work.

z/VM also offers the capability of a high speed private network within the machine. This can help eliminate the requirement for network hardware and security layers between the Linux servers that are needed in a distributed architecture. The network that supports the LAN runs at very fast internal bandwidth. This capability also helps optimize security by further reducing the potential for external breaches and running within an environment certified to EAL3+.

z/VM runs within a hardware partition so there can be multiple partitions running z/VM in a single machine. This allows the creation of a test environment along side the production environment without requiring additional new hardware.

z/VM also has a rich set of tooling for performance management and tuning. These capabilities provide the ability to optimize the environment to specific guest requirements and high utilization. Data is also provided for tools for monitoring and billing. Since an end-to-end application process can be housed within the same VM environment, the tools can be used to optimize the process to meet service level agreements.

The z/VM system, including all guests, can be transferred (often called "failed over") in the event of a system outage as one entity to another mainframe system with adequate resources, thus eliminating the need for individual hardware failover boxes for each system. This can happen dynamically or from backup tapes at a remote location. Many existing mainframe customers test this capability to 100% recovery at least annually.

Because of these capabilities, z/VM delivers an outstanding environment for consolidation and management of many Linux servers.

Best Fit Middleware and Applications

The most frequently deployed middleware on Linux on IBM System z™ are Web application servers such as WebSphere® and Apache and databases such as Oracle, DB2® and Informix®. Each of these is easy to deploy on System z because their code base is the same no matter which platform it happens to execute on, including Linux on System z. Customers can easily move applications built on this middleware to System z for this reason.

In addition to the production environment, the development test environment may also be centralized, thereby gaining the advantages of easy creation, use and reclamation of resources from virtual servers no longer required.

Because production, test, and quality assurance can be very efficiently centralized onto a System z in a high utilization environment, there are frequently significant software cost savings from the reduction of required software licenses in addition to the ease of management of the environment.

Because database and Web applications are generally mission critical, they gain the advantage of multiple servers being able to participate in a single virtual server failover environment, so disaster recovery is significantly easier.

Network infrastructure servers such as FTP, NFS, and DNS are good fits for Linux on System z. Often they support larger applications that are already hosted there, so they can be merged into a single systems environment.

Finally, there is a large and growing base of ISV applications supported on Linux for System z such as the SAP Application Server, as well as PeopleSoft, Siebel CRM and Oracle's eBusiness Suite with the Oracle database deployed on Linux for System z. There is also a broad array of IBM products available like WebSphere Application Server, WebSphere Portal, DB2 Universal Database™, DB2 Connect™, Tivoli® Systems Management and offerings from Rational®.

Systems Management

Ease of management of hundreds of Linux systems is another reason to consider simplifying your infrastructure by leveraging a System z or existing IBM mainframe. A new system can be deployed in minutes or hours rather than the weeks or months required to configure, order, and install new system hardware. And, if the system requirement is temporary, the image can be decommissioned and the resources reclaimed, just as quickly as it was created. Disk, memory, CPU and networking are shared, resulting in a dynamic environment where capacity requirements are allocated as needed by the individual Linux instances.

These capabilities are particularly useful for testing and quality assurance where the need for Linux instances comes and goes, and the capacity required by any particular instance at any point in time varies dramatically.

z/VM's virtual switch capability allows all communication between virtual servers to be housed within the mainframe, reducing the need for external switches and cables.

Once templates for various servers have been defined, it is a simple task to clone them to create another server. The end result is an environment in which it is possible for over 100 servers to be managed by one person.

Since Linux on the mainframe is essentially the same as Linux on Intel® and RISC architectures, it appears the same to either a systems administrator or an application developer. There is no need to learn something new.

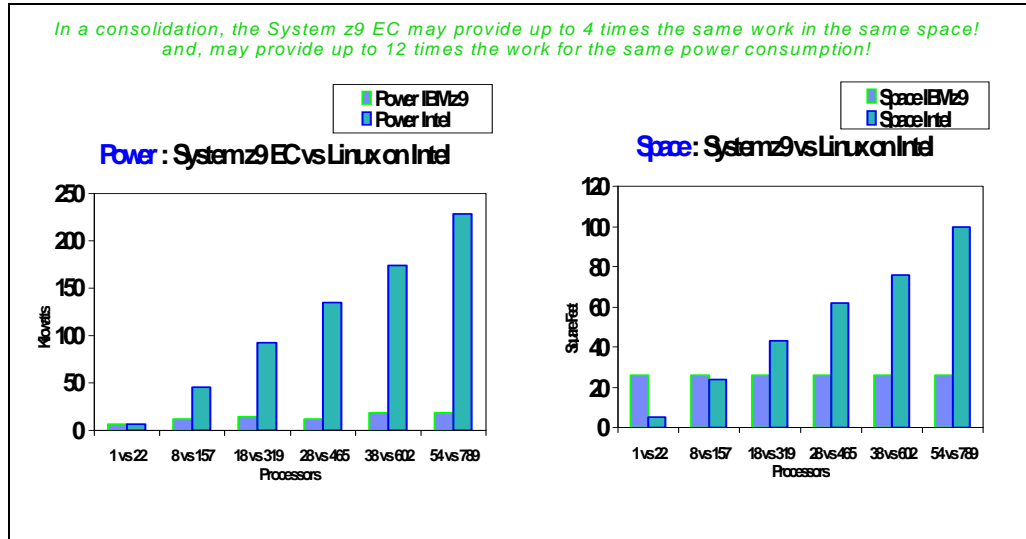
System z TCO

With an environment as established, complex and robust as the IBM mainframe, any discussion of “costs” can be challenging. Attributing costs to an environment that may run 1,000 or more applications concurrently is a much more complex discussion than a “one box, one application” deployment scenario. In order to truly understand the financial strength of the mainframe, the discussion must be re-focused on the Total Cost of Ownership (TCO) rather than the incremental cost of a single acquisition or project (TCA).

This much wider and more comprehensive discussion of TCO will result in a more accurate view of IT costs. It will also demonstrate a more accurate view of the business benefit delivered by those same assets.

In a TCO discussion, System z demonstrates several areas of financial advantage and strength. System z mainframe environments have the potential to lower network equipment costs, reduce real estate requirements, consume less energy than many large distributed system environments, require fewer administrators, and could provide more computing capability through higher overall utilization.

The Linux on Intel servers selected in the example below are functionally eligible servers considered for consolidation to a System z running at low utilization such that the composite utilization is approximately 5%. The utilization rate assumed for System z EC is 90%. This is for illustration only, actual power and space reductions, if any, will vary according to the actual servers selected for consolidation.



System z's lower overall TCO will vary by location and client specifics. Examples may help by comparing a System z to a comparable distributed solution. Power savings may reduce the daily cost of running an environment significantly. One customer recently reported to the Robert Frances Group that their power costs were \$300,000 for distributed systems (not including cooling costs) and \$150,000 (including both power and cooling) for their the mainframe, "The total transactional throughput of the mainframe systems was greater, at less than half the energy and cooling costs¹

One should also consider the impact on UPS systems, and how quickly the requirement for UPS capacity can grow in rapidly expanding distributed environments.

Real estate savings on floor space may reduce requirements very significantly. While the actual cost savings may vary based on the profile of servers being consolidated, and the cost of floor space in different markets, the consolidated footprint of the System z could reduce space requirements to as little as 1/25th of the space needs of distributed environments.¹

Lower overall network costs due to System z's virtual network capability may save up to 50% of the investment necessary for network gear to connect discrete individual servers. In fact, Clabby Analytics, in their recent report "Does System z Offer Better TCO than Scale-up, Distributed UNIX® and Linux Servers?" indicated that the System z's reliance on it's own internal back plane for most communications cuts down on the need for physical switches, hub and routers, resulting in potential savings of up to 75%³. The savings due to system automation and control may reduce the number of people necessary to maintain an environment. Recent customer experiences in consolidating workloads onto System z with

Linux have seen very significant reductions in staffing requirements. Customers have reported significant staff reductions and redeployments of staff using the System z. One IBM customer in the banking industry reported a staffing requirement reduction from 24 to only 8⁴ and a major insurer reported a 50% reduction in FTEs (Full Time Equivalents) to manage its new consolidated System z Linux environment. Much of this reduction can be attributed to the sophistication of the system management tools available on the System z.⁵

Another potential benefit to the sophisticated system management of the System z is that minimizing human interaction with the system management may avoid some downtime. In a recent survey on the causes of downtime, Infonetics found that 27% of downtime is attributable to human error.⁴

The more areas of IT cost included in a discussion of the IBM mainframe vs. comparable distributed environments, the stronger the TCO case becomes. From the new, low cost of acquisition, the IBM System z9 Business Class (z9 BC) giving clients mainframe capabilities starting at approximately \$100,000, the mainframe provides the most bang for the buck of any platform choice available today no matter what the challenge.

Customer examples

Nexxar Group, Inc.

Financial services company Nexxar Group, Inc. consolidated more than 80 x86 servers onto an IBM System z9 Business Class mainframe running Linux under z/VM to help drive its strategy of acquiring businesses and winning new private label relationships. Nexxar Group expects to save 30% per year in related operating costs as a result of switching to the System z9.

Nexxar Group is an “early stage consolidator,” meaning that its business model involves acquiring and integrating smaller firms. The company is relying on the IBM System z mainframe's unmatched virtualization capabilities of z/VM to quickly create a highly secure, custom-tailored computing environment for each “private label” relationship that utilizes Nexxar Group 's money transfer, money order, bill payment and check cashing services.

In addition, the System z mainframe helps enable Nexxar Group's growth-through-acquisition strategy, which requires that the IT environments of newly acquired companies be integrated quickly and efficiently. When Nexxar needs to add a virtual Linux image on the IBM System z9, it's as simple as cutting and pasting with z/VM. By contrast, in the x86 world, Nexxar would have to add a whole new physical server, plus the software, networking, and maintenance. And since the IBM System z mainframe and z/VM

were originally designed to be a shared environment, they provide the type of security that the SEC mandates to financial services firms.

The IBM System z mainframe with z/VM also streamlines maintenance substantially by reducing by 75% the required head count to maintain the operating environment in comparison with the x86 systems it replaced.

Nexxar Group is now running its core applications on the Integrated Facility for Linux under z/VM. System z processor utilization is superior to that of the x86 systems, allowing Nexxar Group to much more efficiently use the server resource it purchases. This has solved Nexxar Group's problem, one that many growing companies face after they've started down the x86 path, namely, how to eliminate the resulting sprawl of distributed servers, with their attendant flood of IT workers and skyrocketing software costs. The IBM mainframe with z/VM offers powerful economic benefits that are delivering significant results for customers running a wide spectrum of new workloads, including Linux applications in a virtualized environment.

Panasonic

Panasonic.com (www.panasonic.com) implemented the IBM WebSphere Commerce Suite (WCS) solution with DB2 UDB on Linux for System z under z/VM on an IFL-only configuration. This implementation replaced the sprawl of their distributed UNIX servers that were formerly running WCS with DB2 UDB. This also solved Panasonic.com's problem of the lack of available power, cooling and floor space that had become depleted by the distributed servers that the System z and z/VM replaced.

Now Panasonic.com has complete control of the TCO in their data center and provides superior performance, scalability, security, and reliability afforded by their System z Linux z/VM solution to their customers. Also peak seasonal buying demand can be easily accommodated by the System z and z/VM with on/off capacity on demand with no disruption to their presence on the internet.

In turn, IBM worked with Panasonic to also create a Web platform for the company's Plasma HDTV customer support program in June 2006. The new online version of the Panasonic Plasma Concierge program includes access to exclusive online content, personalized product information, online priority service scheduling and service updates. Panasonic is now investing more than \$15 million in this unique program, including upgrading its customer call center capabilities, enhancing its Web site and taking measures to improve the overall customer experience which was all made possible by their decision to implement WCS with DB2 UDB on System z with Linux running under z/VM.

Nationwide Insurance.

Nationwide is a leading insurance and financial services provider in the United States that reduced the overall cost and complexity of its Web application infrastructure and increased its provisioning speed when it deployed two IBM System z servers running Linux under the IBM z/VM virtualization technology.

Nationwide wanted to create a cost-effective, scalable and 100-percent available environment for its WebSphere application infrastructure. Not only was its existing heterogeneous infrastructure difficult to handle, but it also prohibited the sharing of hardware and software resources, hindering productivity. To increase its provisioning speed and lower its total cost of ownership (TCO), Nationwide needed to consolidate and standardize its Web application environment.

With help from IBM, Nationwide standardized its WebSphere application infrastructure, leveraging two IBM System z9 servers running Linux under z/VM. The company deployed one System z server in its development data center for WebSphere application development and one System z server in its production data center for WebSphere production applications. The development server also provides high-availability backup for the production server in the event of system failure. Each server features multiple Integrated Facility for Linux (IFL) processors only.

Nationwide's System z servers utilize z/VM virtualization technology to run multiple instances of the Linux operating system on each server. The z/VM technology enables the virtualization of processor resources to streamline the provisioning of dedicated Web servers in support of additional workloads. Dynamic logical partitioning technology facilitates the automatic allocation of processor capacity to meet fluctuating application needs. And Capacity on Demand functionality enables the nondisruptive increase of server resources without incurring downtime.

By consolidating its WebSphere application environment on IBM System z servers, Nationwide reduced the overall cost and complexity of its application infrastructure. Leveraging the Linux operating system cost-effectively increases the speed, reliability and security of the infrastructure, and z/VM virtualization technology simplified and integrated the infrastructure to maximize system productivity.

Taking advantage of the virtualized System z platform of z/VM, Nationwide experienced a 50-percent reduction in monthly Web-hosting costs (hardware/software amortization and support costs) and a significant decrease in middleware licensing costs (WebSphere, IBM DB2 UDB, Oracle, etc.). The company estimates that it will completely recover its initial investment within 15 months and save over \$16 Million in approximately three years.

Nationwide's integrated System z infrastructure provides a highly secure, scalable, high-performance base for deployment of the company's Internet and Java technology-enabled applications, allowing a comprehensive and diverse application execution environment.

Boscov's

Boscov's, Inc (www.boscovs.com) implemented Linux under z/VM running CommuniGate eMail, SAMBA, IBM WebSphere Commerce Enterprise for System z Linux with IBM DB2 Universal Database and MarCole Interactive System's Gift RegistryWorks Enterprise Edition. This solution at Boscov's also includes the IBM System z model z990 with six IFLs.

Boscov's, Inc., is the largest full-line family-owned independent department store in the United States, with annual sales of approximately \$1.5 billion. The company employs more than 12,000 people and is the dominant department store in almost all the markets in which it competes.

Several key business, technical and financial improvements were provided by this solution. The IBM System z and middleware running on Linux for System z under z/VM integrated many of the components of Boscov's back-office systems for inventory control and supply-chain management and ensured that all customer-facing functions were responsive, intuitive and easy to use. Boscov's is now able to add production, test, and QA virtual servers using z/VM much more cost effectively and on demand as opposed to the distributed servers replaced by the System z. Also, the 100+ distributed servers Boscov's consolidated onto Linux for System z under z/VM has saved them human capital costs, hardware costs, and software costs which they quantify at \$1 Million per year in TCO savings.

By centralizing on Linux on System z under z/VM, Boscov's now has a company wide view of inventory, order fulfillment, customer service, and other operations. They now have a cost effective way to provide integrated real-time access to business data which streamlined their IT infrastructure and significantly reduced shipping costs and provided quicker, more accurate deliveries increasing their overall customers' satisfaction.

¹ Source: Powerful Incentives – Using IBM System z to Realize Significant Operational Cost Savings, Robert Frances Group

² Does System z Offer Better TCO than Scale-up, Distributed UNIX and Linux Servers?, Clabby Analytics

³ Does System z Offer Better TCO than Scale-up, Distributed UNIX and Linux Servers? – Clabby Analytics

⁴ Source – Computerworld
<http://www.computerworld.com/softwaretopics/os/linux/story/0,10801,99667,00.html>

⁵ Mainframe Role in TCO, Robert Frances Group, 10/06

⁶ The Costs of Enterprise Downtime: North American Vertical Markets 2005, January 2005 – Infonetics Research



Copyright IBM Corporation 2007

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
U.S.A

Produced in the United States of America
07/07

All Rights Reserved

IBM, IBM logo, DB2, DB2 Connect, DB2 Universal Database, FICON, Rational, System z, System z9, Tivoli, WebSphere, z/VM, and zSeries are trademarks or registered trademarks of International Business Machines Corporation of the United States, other countries or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Information concerning non-IBM products was obtained from the suppliers of their products or their published announcements. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.