

Virtualization and the On Demand Business





International Technical Support Organization

Virtualization and the On Demand Business

August 2004

Note: Before using this information and the product it supports, read the information in “Notices” on page iii.

First Edition (August 2004)

© Copyright International Business Machines Corporation 2004. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law. INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®
@server®
FlashCopy®
HiperSockets™
HACMP™
IBM®

Micro-Partitioning™
POWER™
POWER5™
Redbooks™
Redbooks (logo)™
Redbooks (logo) ™

Storage Tank™
Tivoli®
TotalStorage®
Virtualization Engine™

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Linux is a 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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. 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, and service names may be trademarks or service marks of others.

Preface

In this IBM Redpaper we describe how organizations can use virtualization as a technique to gain more business value and greater flexibility from their information technology (IT) infrastructure.

Businesses and organizations are increasingly becoming virtual, and information technology (IT) is changing to be more flexible and dynamic to support these needs. Additionally, there are opportunities within IT to simplify infrastructure, to provide increased utilization, and to consolidate workloads into fewer, more manageable servers. In the application space, a new paradigm is gathering pace and offers a level of abstraction and virtualization through Web Services and a common messaging hub known as a “service bus”.

In this paper, we examine the positioning of virtualization within the on demand Operating Environment, as well as the infrastructure and server exploitation of virtualization and the function and technology of the recently announced IBM® Virtualization Engine™.

The paper is organized as follows:

Chapter 1, “Starting with virtualization” on page 1, introduces the concept of virtualization, including general business and organizational virtualization as well as key types of technical virtualization.

Chapter 2, “The IBM Virtualization Engine” on page 15, looks at the key role of the IBM Virtualization Engine (VE) in the on demand Operating Environment.

Chapter 3, “The on demand Operating Environment” on page 27, examines the emerging on demand Operating Environment, technology, architecture and services, and describes how this level of virtualization will help the on demand business.

Chapter 4, “Getting started with virtualization” on page 33, re-caps key technologies and suggests ways in which organizations can get started with virtualization.

Appendix A, “Business, organizational, and management benefits of virtualization” on page 37, discusses the business, organizational and management benefits of virtualization, and includes a view of the impact on Cost of IT, Quality of Service (QoS), and Time to Value (ToV).

Appendix B, “The on demand Infrastructure Services offerings” on page 47, lists and briefly describes the various offerings available to help you become an on demand business.

The intended audience for this paper is anyone involved in designing, running, and/or managing the infrastructure for centralized or decentralized computer systems and servers. This paper will help both systems administrators and systems programmers understand the benefits of virtualization as well as introduce some of the key technologies in the IBM Virtualization Engine. CIOs, I/T managers and systems designers will benefit from understanding the relationship of the IBM Virtualization Engine to the on demand operating environment.

The team that wrote this Redpaper

Mark Cathcart

IBM Distinguished Engineer
IBM UK

Svend Erik Bach

IBM Distinguished Engineer
IBM Denmark

Martin Ferrier

IBM Distinguished Engineer
IBM UK

Christian Matthys

Consulting IT Specialist
IBM France

Julie Schuneman

Executive IT Architect
IBM US



Starting with virtualization

Overview

In this chapter, we introduce the concept of virtualization. We start by looking at general business and organization virtualization techniques, and then examine the key types of technical virtualization, discussing server, storage and network virtualization. Finally, we briefly review application virtualization and look at the emerging Web Services (WS) and service-oriented architecture (SOA) paradigm.

Introduction

Increasingly through the 1990s business, governments, education and other segments adopted virtualization techniques to maximize organizational and business opportunities, while simultaneously increasing their reach and reducing their costs. Simple examples of virtual organizations are businesses that outsourced departments and functions which other organizations could do better and more effectively; this included parts of their business which were not considered core, and areas in which the organization was not specialized.

Consider, for example, the areas of catering and building maintenance. Until the late 1980s, many businesses and organizations had departments and internal organizations that handled building planning and maintenance, as well as ordering, preparing and delivering food and beverage services to the employees of the organization. In the 1990s, it became commonplace for specialized companies to take over these roles and responsibilities. These specialized companies had better contacts, more buying power and flexibility, and could often deliver the same or improved services at lower cost than the organization itself. In many cases the organization's own employees were transferred to the specialized company, and yet continued to work in the same place, doing much the same job—they had in essence become “virtual” employees.

This organizational virtualization continued throughout the 1990s, encompassing functions such as shipping, delivery, security and so forth. New opportunities arose as a result of massive investment in telecommunication infrastructure, which made national, long distance and international telephone and computer networking more cost effective. At the same time new manufacturing economies in the Far East matured and came onstream in an increasing

number of countries. Combined with the rapid increase in power and reliability of the micro-computer, these developments made it possible for even more business functions to be performed far from headquarters or regional headquarters buildings.

This trend has resulted in a business environment where small, medium and large businesses are often indistinguishable when judged by the number of employees only. An organization with only fifty employees may have the reach, range, and manpower of an organization with many thousands of workers. Almost all of these workers are in fact contractors and subcontractors; they belong to other organizations, both small and large, that specialize in delivering specific function and value to the business. In the best organizations, the difference between company employees, subcontractors and contractors is undetectable, with all workers performing as a single, virtual organization. This in turn has created an opportunity for a new strategic business operation known as Business Process outsourcing, in which the whole business operation is moved externally, rather than simply acquiring computers and staff.¹

There is also a new breed of “meta-organizations”, which are organizations that come together dynamically to address an immediate requirement or business opportunity. These meta-organizations integrate, perform their functions, and then disintegrate and move on to the next opportunity.

Disaster management is a good example of this type of virtual organization. In an environmental, political, natural or other disaster, a number of organizations come together to form a meta-organization that quickly addresses the various aspects of the disaster—from helping with the immediate consequences and recovering and returning to normality—to putting into place long-term plans to avoid repeat occurrences.

The Technology Mirror

From the outset of IT in the 1960s, small and medium-size organizations could not afford to compete with large enterprises in terms of investing in computing resources. Such companies often used a computer bureau to perform business calculations by remotely accessing the computer applications hosted by the bureau. As IT became more cost-effective, small and medium organizations, as well as internal departments at medium and large organizations, and individuals, increasingly acquired their own computer systems. In the 1980s such systems began to be networked together, and grew throughout the 1990s into an increasingly complex interrelationship of systems and business applications.

The functions and number of staff required to support this increasingly complex IT environment grew rapidly, and while the productivity of knowledge workers increased, the benefit of this productivity increase was offset by the cost of staff to deliver, manage and support this environment. Large and medium organizations began to see the benefit of more organizational virtualization, and strategically outsourced IT delivery to specialized companies. As with other business processes, often employees were transferred to the specialized company but in many cases remained in the same building, at the same desk, and doing the same job as before; they too had become “virtual” employees.

Since Year 2000, the worlds of business and IT have become increasingly intertwined and interdependent—and this interdependence is driving the next generation of virtualization. Business applications such as supply chain management, enterprise resource planning, and human resource management have become increasingly packaged rather than custom-built. However, these packaged applications often force companies to change their organizational

¹ British Petroleum (BP), for example, has moved its entire Finance and Accounting function worldwide to IBM. IBM employees handle all of BP's accounts payable and recoverable and their financial reporting. June 14, 2004. Fortune Magazine.

structure to match the model used in the software design. This macro-level process virtualization has delivered value, but often not the dynamic flexibility required.

Business no longer looks to IT to support a specific business function such as back-office, payroll, accounting, manufacturing support, and so on. Instead, it increasingly looks to IT to provide the business with new opportunities by dynamically integrating sets of flexible functions—combined to address short-term opportunities—by integrating these functions irrespective of their physical location and ownership (that is, which company or business unit owns them), and to find, bind, and execute function to deliver this new business value.

For example, a small company might detect a short-term market opportunity to supply a part used during automobile manufacture. To exploit the opportunity, the company uses a model-driven business application that allows it to integrate its design process with network-based applications to dynamically handle financing, manufacturing, shipping, handling and delivery to the auto manufacturer. Rather than the small company needing to have all the relevant business relationships and pre-established contracts, and to negotiate network interoperability and data formats and protocols in advance (as would have been the case in the past), a new breed of technology has emerged to allow this to happen dynamically.

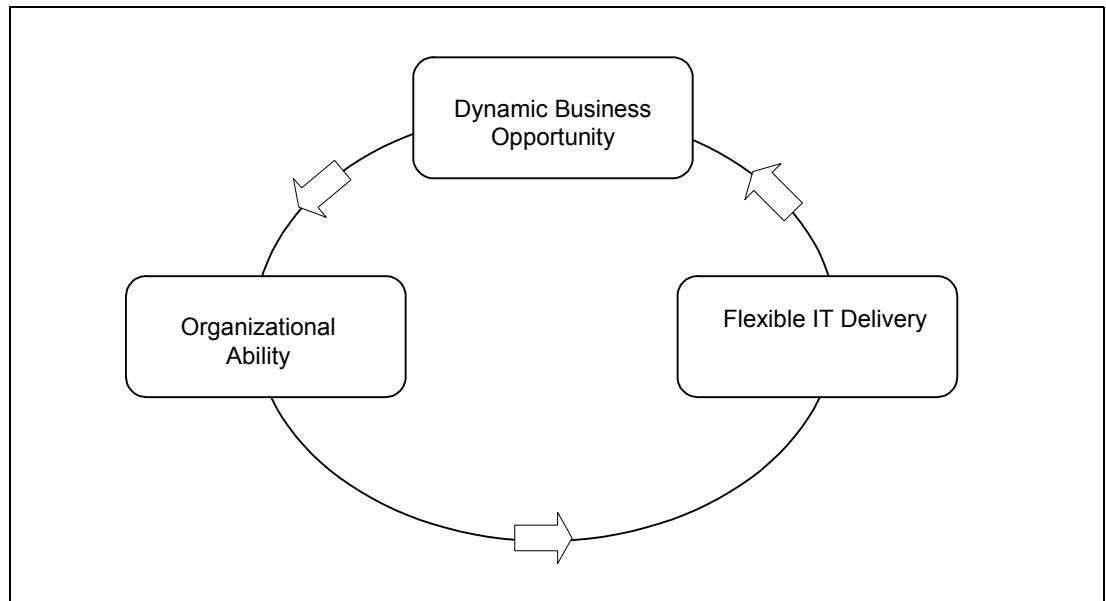


Figure 1-1 Dynamic business opportunity

Finally, as the demands for these new virtualization opportunities increase, new technology has emerged that provides an opportunity to consolidate multiple workloads from two or more physical servers onto a single server, managing them as two logical workloads on top of the same physical server.

This type of server virtualization is not new. In fact, servers have been virtualized in one form or another since the outset; for example, operating systems virtualize the hardware and isolate applications from it (see Figure 1-2 on page 4).

Increasingly, middleware is abstracting operating systems to the point that the hardware and operating systems become transparent to the business application and can be run on different systems without being rewritten or ported; Java™ and SAP Advanced Business Application Programming Language are two application languages that, when executed in middleware, provide this capability.

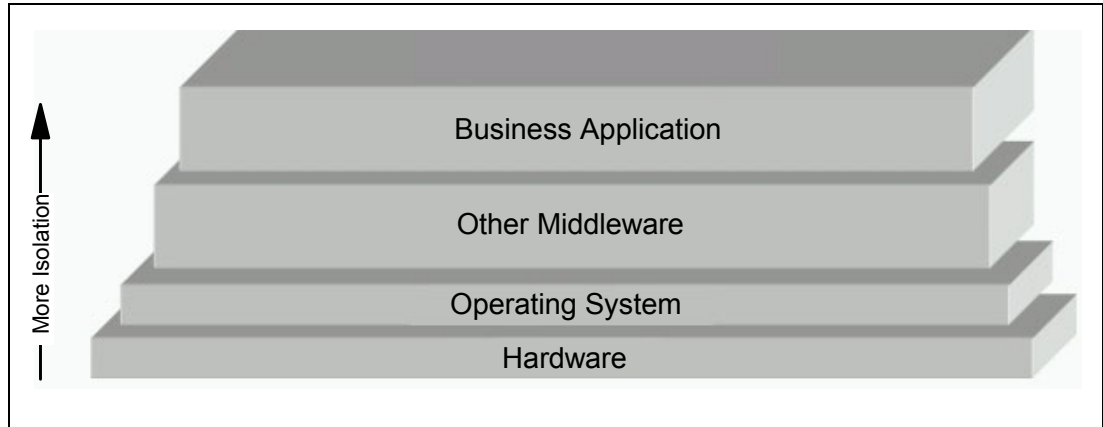


Figure 1-2 Hardware isolation

Virtualization

Today, most applications are not typically designed to be integrated with other applications or software. Applications are generally developed to support one specific business area or function. They have often been monolithic in design and, although there are exceptions, each application usually runs on its own dedicated physical server or set of physical servers, and has its own physical storage devices. There is very little, if any, sharing of IT resources between different applications unless they were deployed on mainframe-based machines or on machines that support partitioning. For the most part, non-mainframe applications get deployed on their own specific set of physical hardware and network resources. This has led to the fragmented heterogeneous infrastructures that exist today, and it often results in a lack of flexibility within an organization's IT environment.

IT organizations struggle with the fact that the servers on which these applications are deployed usually have low utilization rates, and the unused system resources sit idle and are wasted. Costs for today's implementations are driven even higher as IT departments over-provision systems to ensure that processing power or storage is available to meet peak demands for the specific business area being supported. The associated technical and operational costs related to this type of environment continue to rise as the number of physical resources and complexity continues to grow.

Various technologies can be utilized to help businesses better leverage their servers, storage, and networks. For example, IT departments can partition some servers to increase their utilization, use Blade technology to enable better management of Intel®-based systems, take advantage of networking capabilities that are built into servers today (such as HiperSockets™ and Virtual LANs), and exploit the capabilities of current storage technologies like block virtualization, file aggregation, and centralized management.

As businesses move forward with on demand strategies that drive implementations of true end-to-end integrated business processes, the number of application systems that support those processes will continue to grow and will increasingly span multiple servers on different technologies. These heterogeneous cross-platform technologies will need to be managed, monitored, and measured holistically to ensure that the needs of the business and SLAs are being met. In an on demand environment, virtualization is needed to uncouple the applications from the physical configurations.

About automation: In order to deliver on the promise of virtualization, it is assumed that greater levels of automation take place. While some value is realized by centralizing administration, this can be quickly lost if the administration cost goes up as you try to be more flexible and dynamic in how you use the centralized resources.

In 2001, IBM launched the eLiza Research project, which more recently was subsumed into the Autonomic computing initiative. Autonomic computing is aimed at enabling computer systems to become self-configuring, self-healing, self-managing, and self-protecting. While it is beyond the scope of this Redpaper to review Autonomic computing, in the context of virtualization it is safe to assume that we are designing aspects of Autonomic computing into our virtualization products and offerings.

An example of this is covered in Chapter 2, “The IBM Virtualization Engine” on page 15, in the discussion of the IBM Virtualization Engine and its workload management and provisioning features.

Autonomic computing is having a much wider impact and will be at the core of the on demand Operating Environment, which is covered in Chapter 3, “The on demand Operating Environment” on page 27.

In the interim, automation plays a key role in virtualization, server consolidation and so on, allowing you to benefit from more than just the simple centralization of administration.

Server virtualization

IBM has been delivering systems that have provided virtualization for many years. Workloads running on mainframes today exploit technologies such as virtual memory, where each application thinks it has its own real, dedicated memory. Logical partitioning (LPAR) allows customers to “slice up” a machine into virtual partitions, and provides the flexibility to dynamically change the allocation of system resources for those environments. Any of the virtual servers may run on any of the physical engines, meaning that the engine resources are fully shared, which makes it possible to run the physical server at very high utilization levels.

The virtualization capability is integrated into the base hardware and microcode of the physical server, and it allows for a definition of virtual servers with a sub-engine granularity. For example, by utilizing common hardware and microcode componentry, both the IBM eServer i5 and p5 systems offer Micro-Partitioning™ where the POWER5™ processors on the servers can be allocated to tenths of a processor among the partitions. In addition, both the iSeries and pSeries will allow you to run multiple operating systems in different partitions, and allow processors, memory, and I/O to be shifted among active partitions without requiring the operating system to be rebooted.

In addition to the flexibility provided by this level of isolation, certain models of the IBM eServer are certified to meet the Common Criteria at Evaluation Assurance Level 4+ (EAL4+), which demonstrates the reliability, workload isolation, and security that can be achieved through virtualization.

The eServer pSeries is not alone in delivering this kind of virtualization. The IBM eServer zSeries mainframe has a similar level of function and continues to build on it by adding ever-higher numbers of LPARs and virtual I/O capabilities, while retaining its famed isolation and security.

In IBM eServer Logical Partitioning, there is no affinity between the processor engine resources, the memory, and the I/O interfaces, thereby making it possible to provision and

re-provision the resources to the virtual servers independently. The hardware-based virtualization capabilities of the IBM mainframe also allow virtualization and sharing of I/O paths, and allows for the definition of virtual TCP/IP networks connecting the virtual servers at memory speed.

Some Windows® and UNIX® servers generally only support a virtualization technology called “physical partitioning”, in which each physical partition or virtual server owns a dedicated set of resources not used by others, and where the capacity granularity is based on multiple engines (2 to 4), because the virtualization is aligned with the physical components.

Virtualization granularity, at that level, does not generally provide significant advantages with respect to consolidation unless it is combined with other software types of virtualization methods. This is because each partition acts like a separate server, with dedicated workloads and hence low utilization. The affinity of the virtual servers to the physical board configuration also means that resources such as processor engines, memory, and I/O interfaces have to be provisioned and re-provisioned together, thus making the implementation less efficient.

Figure 1-3 shows the three common modes of server partitioning. Hardware partitioning uses fixed resources; software and logical partitioning have greater, more dynamic capabilities, depending on implementation and platform.

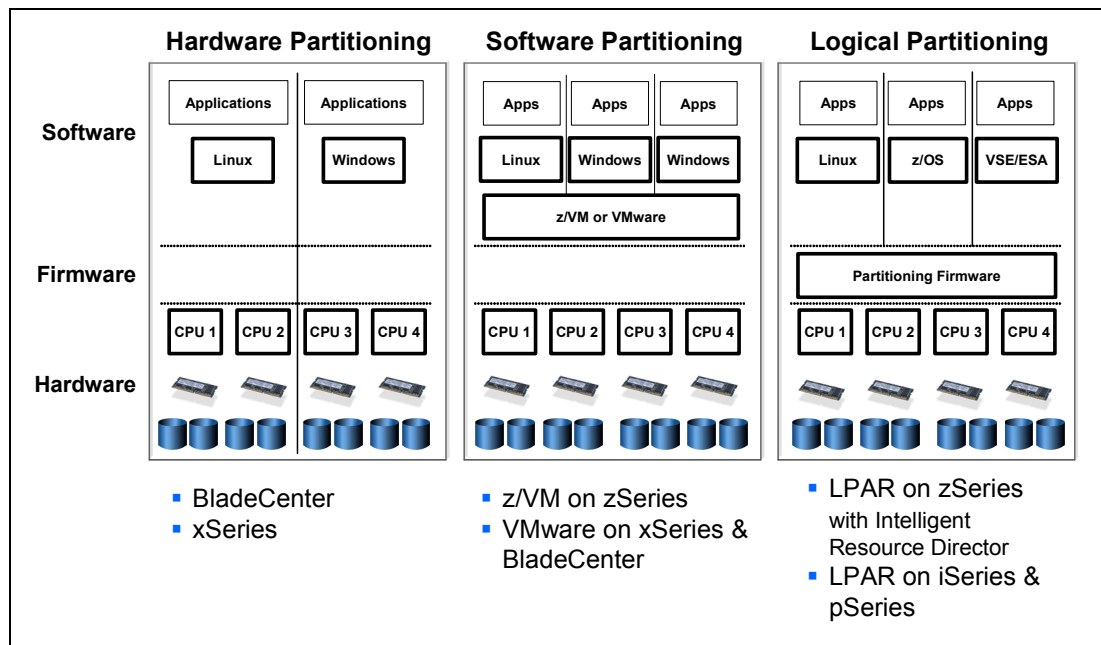


Figure 1-3 Understanding virtual server partitioning

Virtualization is now expanded beyond the scope of one physical box and one architecture to deliver the classical benefits of virtualization on an end-to-end scale, across all architectures including servers, storage, and networks. This new level of virtualization is realized through technologies like provisioning and workload management.

Workload management provides the control to ensure system resources are provided to the applications that are most critical to the business. The Intelligent Resource Director (IRD) in zSeries machines expands on the workload management concept to provide goal-based resource balancing across multiple LPARs. Virtual Machine (VM) technology allows fine-grained control for defining and running multiple operating system images and their workloads that share memory, processor cycles, and IO channels, while allowing centralized management.

IT organizations that exploit these technologies derive value and benefits in the form of easier management and cost savings from sharing system resources, which in turn results in more efficient use of those resources. Businesses can take action today to become on demand by exploiting these same types of technologies as they continue to evolve into other IBM eServer platforms.

For example, companies can exploit LPAR technology on the zSeries, pSeries, and iSeries platforms today. System resources including processors, memory, disks, system buses, and I/O can be associated with a given partition. Processors can be in a “shared pool” and moved among partitions on an as-needed basis. In addition, solutions such as VMware, which provides virtualization capabilities on the xSeries platform, can be used to increase the utilization of those Intel-based machines.

The most advanced software virtualization product on the market is the IBM z/VM product running on the IBM mainframe. z/VM takes advantage of the hardware capabilities within the mainframe hardware and establishes a very low overhead and highly secure virtualization platform. It allows for the consolidation of large numbers of servers in a virtual server farm. z/VM is also the main virtualization enabler for Linux on the mainframe.

Another virtualization capability, Capacity on Demand, is implemented across the IBM eServer™ platforms. zSeries, iSeries, and pSeries servers provide support for temporary or permanent processor enablement based on a client need.

Virtualization capabilities are also available through the use of IBM BladeCenter™ technology. The IBM BladeCenter can be used to increase the cost efficiency of these Intel-based servers by providing the capability to share system resources such as Ethernet adapters and fiber optic switches across all blades within the BladeCenter.

The BladeCenter offers the capability to scale out. A BladeCenter can be purchased without all blade slots being fully populated. As additional capacity is needed, more blades can be purchased to populate the BladeCenter. This variable cost model allows the unit to be partially populated with blades, thereby reducing the number of blades that may sit idle in the BladeCenter until some future point when the business need drives application demand for the resource.

Clustering

Other technologies, such as clustering, can be used to “virtualize” the server or provide some of the cost benefits related to the server virtualization technologies previously described.

- ▶ Clustering of servers provides a way to make multiple server resources, and in some situations multiple data resources, appear as a single resource space. The value of this is the ability to scale capacity and establish highly available solutions.
- ▶ Blade solutions provide for consolidation with virtualization benefits, where the physical space used is reduced and a limited amount of hardware infrastructure sharing takes place. Virtualization of individual blades is also possible. The processor resource space being shared is relatively limited, and this reduces the ability to move capacity dynamically between the different virtual servers. The IBM eServer product line provides efficient blade solutions both for xSeries supporting Windows and Linux and for pSeries supporting AIX and Linux.

For many organizations, the combination of clustering, server virtualization, blade-based hardware and a services-based infrastructure will provide a compelling commodity infrastructure. Increasingly, clustering based on the successor to the grid standard, the OASIS Web Services-Resource Framework, and related industry standards, will be an

increasingly attractive option as IBM delivers its on demand Operating Environment; this is discussed further in Chapter 3, “The on demand Operating Environment” on page 27.

Clustering can also help with high availability

For a long time, high availability has been achieved through the use of fault-tolerant technology, often by using specialized hardware. However, this design may not address software failures. With clustering, availability is achieved not via a series of replicated physical components, but rather through a virtual set of system-wide, shared resources that cooperate to guarantee essential services.

With the exception of the IBM eServer zSeries Parallel Sysplex, “clusters” are a group of loosely coupled machines networked together. Clustering combines software with hardware to minimize downtime by quickly restoring essential services when a system, component, or application fails. Clustering multiple virtual servers to provide high availability is a cost-effective option.

Clustering for high availability encompasses the notion of a number of interconnected virtual and real systems; if one of those systems fails, the resources required to maintain business operations are transferred to another available machine in the cluster. However, different mechanisms have to be put in place to be sure the transfer of resource ownership to the surviving cluster members is automated.

About grid technology: The emergence of grid technology as a software framework providing layers of services to access and manage distributed hardware and software resources is having an increasing impact on IT. Initially of interest only in the High Performance Computing (HPC) and Engineering/Scientific sectors, grid is becoming increasingly influential on the direction of commercial IT. IBM's leading role in grid through the Global Grid Forum, and its contributions to the Open Grid Services Architecture (OGSA) and Open Grid Services Infrastructure (OGSI), helped to accelerate the take up and adoption of grid technology.

The longer-term impact of grid technology will be seen in the on demand Operating Environment. The successor to the Web Services-based OGSI, WS-Resource Framework will form the core industry open standard that is used to define a complete services-based infrastructure. This will allow pools of computer resources to be described abstractly and systems—whether local or remote, centralized or de-centralized—to be clustered together as a set of canonical services.

While we do not specifically discuss grid technology in this Redpaper, the on demand Operating Environment is introduced and explained in Chapter 3, “The on demand Operating Environment” on page 27. For a useful introduction to grid technology, refer to the Redpaper *Fundamentals of Grid Computing*, REDP-3613-00. For information on how to exploit grid technology today, refer to the IBM Redbook *Enabling Applications for Grid Computing with Globus*, SG24-6936.

With clustering, you can use more of your computing power while ensuring that critical applications continue operations when recovering from a hardware or software failure. In addition, high availability clustering minimizes, or ideally, eliminates, the need to take resources out of service during maintenance and reconfiguration activities. (For example, HACMP™ AIX software provides a fast recovery feature to minimize unplanned downtime.) Another benefit is the ability to add and subtract servers from the cluster, which allows the cluster to scale to meet end-user demands for the service.

Combining grid-based technology and heterogeneous SMP and/or blade-type servers that are co-located and interconnected by high-speed clustering—either vertically between like

systems, or horizontally between like and unlike systems—will become an increasingly attractive way to run a heterogeneous utility-like infrastructure.

Note that increasingly there appear to be many similar technologies offered by a number of vendors of server virtualization technology. However, organizations need to take care in selecting an appropriate technology, because while they may provide “virtual machine” or partitions, many alternative implementations do not provide the same granularity, flexibility and isolation, although they may be marketed using the same terminology.

These are key values of any server virtualization technology and the more granular the technology is, the more flexible it is likely to be. Inherently the more flexible it is, the higher level of isolation it needs to deliver in order to achieve the reliability required to deliver on demand!

Storage virtualization

Storage virtualization allows you to combine the capacity of multiple storage controllers into a single resource with a single view of the storage resources. This abstraction layer between the physical storage devices and the users of those devices (host applications) provides the ability to hide the physical infrastructure from the application and end user. Simple benefits of storage virtualization include the ability to add additional space and volumes, to move to different types of disk, and to replace aging disk subsystems, all without applications or host systems having to know anything about what is going on.

The term “storage virtualization” may not be new to an installation, as storage virtualization techniques, while not widely used in the UNIX and Intel environments, have been used for the last decade in the mainframe environment with products such as IBM System Managed Storage Software and Virtual Tape Server. In the past, application systems were tightly tied to storage, down to knowing the specific disk sector that would be read or written. Over time, storage technologies continued to evolve but applications were still tied to specific disk storage volumes even when using SAN technology. Today’s technologies, however, offer features like Block Virtualization and File Virtualization:

- ▶ *Block Virtualization*, for example, can be implemented using the IBM TotalStorage SAN Volume Controller to allow application systems to believe they have physically contiguous storage volumes allocated to them—in reality, however, the allocation is scattered across multiple devices. This allows the IT organization the flexibility to move application data from one physical device to another without affecting the application systems. The “virtual” disks that the application systems see appear to be unchanged. However, behind the scenes, the data may actually have been relocated to other physical devices.
- ▶ *File Virtualization* can be implemented using the IBM TotalStorage SAN File System to allow application systems to believe they have a locally attached file system—in reality, the file system is shared across the Storage Area Network. This allows the administrator to designate policies to direct which type of storage each file is placed on, allowing some files to get higher Quality of Service (QoS) while other files get lesser service.

On the zSeries platform, System Managed Storage has removed the need for applications to understand the underlying disk technology, and provides a policy-driven capability to match disk performance and availability capabilities with the service level requirements of the applications using that disk. In addition, DFSMSHsm delivers on the “virtual” disk concept by providing the ability to move data between disk and tape, transparently to the application, ensuring that the higher-performing disk space is available as required for those applications that need it.

On the iSeries platform, Virtual Storage enables other operating systems and integrated xSeries solutions to leverage the advanced iSeries storage architecture through:

- ▶ Data automatically spread and protected
- ▶ Additional disk arms for better performance
- ▶ Automatic balancing of storage across drives
- ▶ Consolidated backup for all operating environments
- ▶ Flexible storage management
- ▶ Easy setup of multiple environments

Storage spaces can be created from i5/OS that are 1 MB to 1 TB each—up to 32 per integrated xSeries solution, and up to 64 for Linux. These storage spaces can also be dynamically added.

There are several models for implementing virtualization in a SAN. One model that is gaining acceptance is based on the installation of a so-called “Virtualization Server”. The Virtualization Server connects to the heterogeneous application server environment via the SAN fabric. The disk storage subsystems, which can be legacy storage or a latest generation product, are then connected either directly to the Virtualization Server or through the SAN, such that only the Virtualization Server can directly address the storage.

This Virtualization Server functions as an abstraction layer between the physical storage subsystems and the applications, and provides the ability to hide the physical storage infrastructure from the application and the end user. Use of the Virtualization Server allows application systems to believe they have physical storage volumes allocated directly to them. This allows the IT organization the flexibility to move application data from one physical device to another without affecting the application systems. The “virtual” disks that the application systems see appear to be unchanged. However, behind the scenes, the data may actually have been relocated between physical devices. The Virtualization Server may also provide replication functions, such that data can be mirrored across sites, either in real time or asynchronously. Local disk copying function may also be provided, such that a logical copy of disk data can be taken in a few seconds.

The value of the Virtualization Server delivered this way will depend upon several factors, such as the application server base, the storage management software currently used, and the amount of legacy storage installed. For example, existing host-based Logical Volume Manager (LVM) products provide advanced disk management functions today. If an installation has a common server base with a common LVM product, then the installation effectively has a standard administrator interface. The current generation of consolidated disk storage products, such as the IBM Enterprise Storage Server, have their own storage allocation and reallocation software and they provide data replication functions for a wide server base. If the disk storage is all of the same type, then once again there will already be a common administrator interface.

The storage administrator uses the Virtualization Server to allocate or reallocate storage to each application server as required, and to create disks with specific size or performance requirements. Some disk subsystems products do this today, but each disk product has its own unique administrator interface, whereas the Virtualization Server has a common administrator interface regardless of the storage subsystems that are attached to it.

Network virtualization

There are business-critical application requirements to manage and utilize network resources more efficiently in regard to performance, resource usage, people cost, availability and security. Network virtualization includes the ability to manage and control portions of a network that may even be shared among different enterprises, as individual or virtual networks, while maintaining isolation of traffic and resource utilization. This includes technologies such as Virtual Private Networks (VPNs), HiperSockets, Virtual Networks, and VLANs. It also includes the ability to prioritize traffic across the network to ensure the best

performance for business-critical applications and processes. Instrumentation of network resources and operations (SNMP, CIM, CLI, and so on) which can be abstracted across the server and networking devices are key enablers for on demand behavior.

Here's a glance at what some of these technologies provide:

Virtual IP Address (VIPA) takeover – Virtual IP addresses abstract the physical connection of servers to networks. VIPA takeover allows for automatic recovery of network connections between different servers.

HiperSockets - this is mainframe-based technology that allows any-to-any TCP/IP network connection between virtual servers. This provides secure IP communication at memory speed between virtual servers, thus creating a base for closer integration of applications and the implementation of new data-intensive applications.

Virtual Ethernet – pSeries technology that enables internal TCP/IP communication between partitions.

Network blades – xSeries support that integrates switches from different network vendors to simplify communication between server blades in BladeCenter.

Virtual Network providing Server and Network integration- if physical server farms are consolidated into virtual server farms, then parts of the physical network can be replaced by a virtual network, saving cost and reducing management complexity. Network performance and bandwidth between the servers is increased, enabling new data-intensive applications.

Virtual LANs (VLANs) - this is a commonly used standards-based technique in which physical networks are shared in a secure way between multiple applications or user groups.

Virtual Private Networks (VPNs) – this is a commonly used standards-based technique that encrypts data between two TCP/IP endpoints to provide end-to-end physical security for the transport.

From an on demand perspective, network resources must be integrated into all the Service Level Management functions. Mapping the required network resources (including notions of bandwidth and priority) to the application/service being deployed is one of the key tasks that need to be automated. By virtualizing those network resources (server and core network), the ability to share them is enabled and management is simplified. The direction is for users to express their business goals via policy, and then to automate the process of translating those goals into resource-specific actions which can be coordinated to deliver the desired qualities of service.

Examples of existing network virtualization include:

- ▶ IBM eServer zSeries sharing of network adapters, including mapping application QoS to network QoS/priority
- ▶ IBM eServer xSeries network blades, including Layer 2 and Layer 3 switches
- ▶ IBM eServer iSeries provides 1 Gb connections between partitions and integrated xSeries solutions with no LAN adapters/switches
- ▶ IBM eServer iSeries, pSeries and zSeries virtualizing of IP addresses and IP address takeover
- ▶ IBM eServer zSeries HiperSockets and iSeries and pSeries Virtual Ethernet
- ▶ IBM eServer pSeries, iSeries and xSeries provisioning of network resources via IBM Tivoli® Provisioning Manager
- ▶ IBM eServer zSeries sharing of Linux firewalls and load-balancers under zVM

- ▶ VSWITCH and VLAN features of z/VM

In these examples, it is important to note that much of the virtualization is at the platform layer, requiring support in the hypervisors—and in some cases, in the firmware—to enable sharing between different operating systems (in addition to the sharing that is provided in the operating system for functions like VLAN, QoS and VPNs).

Also note that integrating the management of these networking resources in the context of the applications they support and the servers on which the applications run begins to reduce the complexity of managing servers and networks separately (for example, IBM Director support for Cisco and Nortel blades, and IBM Tivoli Provisioning Manager support for configuring server network resources like adapters, IP addresses, VLANs, and so on).

The platform is also involved in provisioning/configuration, performance management, security and availability services. Functions such as Intrusion Detection filters may also be integrated in some operating systems. The use of policy to abstract the management of these services is also an element of the network virtualization support which is being enhanced to better integrate with the other resources (server/storage) and exploited by more of the service-level manager disciplines. What we are driving toward is the notion of lifecycle management of resources, where they are provisioned, monitored, and managed according to business goals using standards-based instrumentation and operations to reduce the complexity for customers in managing these disparate sets of resources from different vendors. Abstraction of physical interfaces and the support for sharing them across different virtual servers are the key concepts which network virtualization addresses.

Web Services

Web Services allows a computer program to dynamically locate a partner program that is an interface to a specific service. Web Services is a new model for using the Internet; it enables transactions initiated automatically by a program. The programs or services can be described, published, discovered, and invoked dynamically in a distributed computing environment. These services enable not just new business opportunities, but also new ways of using the Internet:

- ▶ Marketplaces
- ▶ Auctions
- ▶ Intelligent Agents

The services can be built on industry open standards, and exchange messages and data formatted with and described by the eXtensible Mark-up Language (XML).

About Web Services: Although describing the function of Web Services is beyond the scope of this Redpaper, we do provide a high level view of what it offers. As described in Chapter 3, “The on demand Operating Environment” on page 27, the on demand Operating Environment will make extensive use of Web Services to tie together and integrate computer systems and services into a services-based dynamic, autonomic infrastructure.

For more detailed and recent information about Web Services technologies and their use, refer to the IBM Redbook *Using Web Services for Business Integration*, SG24-6583.

Web Services has defined interfaces and technologies to allow an organization to describe the functionality (services) it wants to externalize. How and where they publish information about the service, and how other organizations discover services, connect to each other, and invoke services with appropriate security, reliability, and confidentiality, are all addressed by Web Services standards. XML defines a platform-independent way of representing data,

making data integration easy and standardized. Web Services defines a platform-independent way of exchanging that data, so that process-level integration becomes much easier.

Figure 1-4 illustrates the Web Services process.

- ▶ Service Provider
 - Provides e-business services
 - PUBLISHES availability of these services through a registry
- ▶ Service Broker
 - Provides support for publishing and locating services like telephone yellow pages
- ▶ Service Requester
 - FINDS required services via the Service Broker
 - BINDS to services via Service Provider

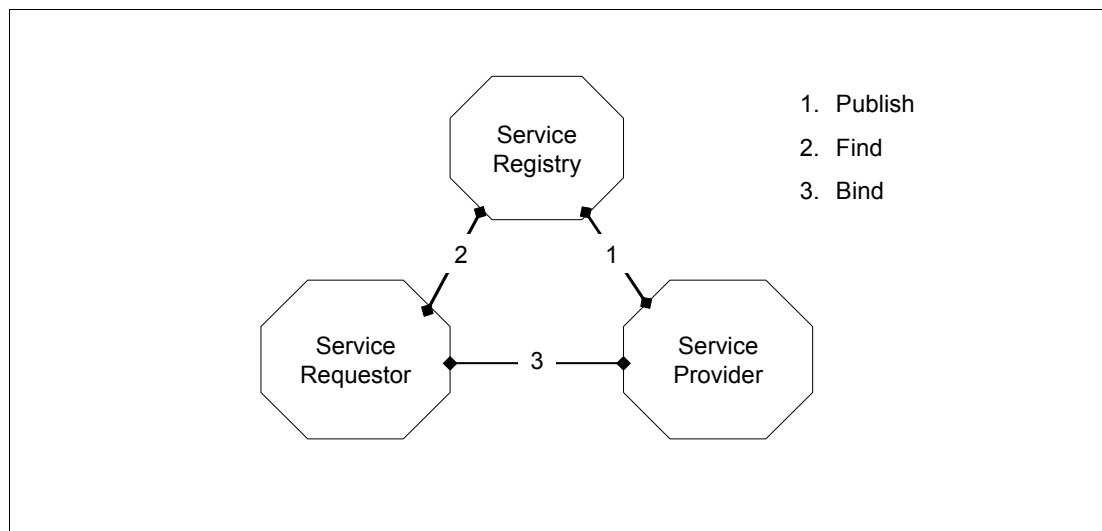


Figure 1-4 The Web Services process

The Web Services process delivers the following benefits to businesses:

- ▶ The ability to integrate systems, regardless of their implementation, to move from monolithic, custom-coded applications to choreographed, scripted components.
- ▶ The agility and flexibility to reconfigure business functions to try new process models.
- ▶ The ability to move from tightly coupled systems to loosely coupled ones, to deal with inevitable change.

This function delivers a level of virtualization and abstraction of business processes that can allow businesses to offer more flexible and more dynamic access to their computer systems and services.

About Application Virtualization: This level of application virtualization offers significant business growth opportunities by allowing an organization to externalize key business processes in areas in which the organization can add significant value.

Web Services as interfaces have been used to enable integration of diverse distributed applications and middleware functions into end-to-end IT process in support of the business

with WS-RF, Web Services are extended to the heterogeneous distributed IT infrastructure, making possible the integration needed to view and manage it as a single compute resource.

Summary

In Chapter 1, we introduced the topic of virtualization, provided a high-level view of the types of virtualization, and discussed some of the benefits of virtualization.

Having established some basic terminology and technologies for virtualization, in Chapter 2 we examine how the IBM Virtualization Engine is extending many of these infrastructure capabilities, as well as introducing a set of new technologies and services that build on the base infrastructure and delivering an abstraction that can be leveraged by common skills and common tools to help organizations become more dynamic and more flexible.



The IBM Virtualization Engine

Overview

The IBM Virtualization Engine™ consists of technologies and systems services that work together to help reduce the complexity of managing and monitoring disparate physical systems across a heterogeneous environment. They can provide increased availability of data and applications within individual platforms, as well as across the entire infrastructure. By providing a unified and consistent management platform, and by penetrating physical barriers to enable pooling of resources, the Virtualization Engine can help increase the productivity of organizations while improving the utilization of system resources in a single server or across the enterprise. In its fullest implementation, the Virtualization Engine is intended to become the Virtual Computer that renders the heterogeneous distributed compute resources in the infrastructure as a single computer resource.

The IBM Virtualization Engine, a set of comprehensive systems technologies and services, helps you to aggregate resources and enables access to a consolidated, logical view of them throughout a heterogeneous, distributed environment. The Virtualization Engine can help simplify and optimize your IT infrastructure, reduce cost and complexity through optimized resource utilization, and increase the business value of IT investments.

With its advanced virtualization technologies and systems services, the Virtualization Engine helps your business to become a dynamic, flexible, and innovative on demand business. It utilizes key IBM virtualization technologies, making available a comprehensive approach to enterprise-wide virtualization that is consistent across heterogeneous environments.

It can help you deliver to your company the promise of on demand business by simplifying systems management across operating systems, servers and storage platforms, thereby helping to reduce the complexity, costs and efforts associated with a heterogeneous IT environment. Use Virtualization Engine's flexible mix of systems technologies and systems services to better optimize computing resources and integrate technology and business processes. These advantages help your organization respond in real time to market changes, opportunities, and demands.

Based on open and industry standards, the IBM Virtualization Engine enables integration across heterogeneous environments—covering IBM and select other servers, operating

systems and architectures. With the IBM Virtualization Engine, you have the freedom to easily integrate new and existing infrastructure resources without sacrificing current investments.

Advanced virtualization systems technologies and systems services built to help address customer needs can enable you to create a nimble IT infrastructure capable of shifting priorities, handling growth, and addressing key service levels (without jeopardizing quality) to answer business requirements in real time.

The IBM Virtualization Engine can provide a comprehensive set of systems services and technologies that can help you achieve your integration goals by aggregating resources into a consolidated logical view. Through this simplified, cross-platform view, you can:

- ▶ Increase utilization of system resources, thereby reducing IT proliferation
- ▶ Broaden access to resources, helping to give people access to the information and applications they need, when they need them
- ▶ Simplify management of computing resources, potentially enabling redeployment of IT staff
- ▶ Increase application availability to help operational speed and business responsiveness
- ▶ Leverage flexible technologies for your changing business needs

The IBM Virtualization Engine helps optimize both IBM eServer systems and IBM TotalStorage® systems product lines across your organization. Because it can provide world-class virtualization functions to all platforms—IBM eServer platforms and select others—it can enable simplified management of resources across a truly heterogeneous infrastructure.

And Virtualization Engine's optimization, provisioning, and systems management functionality can be leveraged for one server or extended to all enabled IBM eServer and TotalStorage systems across your entire enterprise.

Simplify IT management

The IBM Virtualization Engine can provide a simplified view of all enabled computing resources at one time. This can help to streamline IT administration, lower costs and free skilled IT staff to work on other tasks. Centralized server and storage management is provided by IBM Director Multiplatform and the IBM TotalStorage Productivity Center. These also provide integrated console capabilities and consistent management of distributed IBM eServer and IBM TotalStorage systems.

Reduce costs with enterprise workload management

The Virtualization Engine can also help reduce IT cost and complexity by enabling you to better optimize the utilization of distributed system resources. By monitoring workloads across your infrastructure, the IBM Enterprise Workload Manager (EWLM) component of Virtualization Engine systems services can give you the ability to better utilize IT resources by helping increase utilization of existing systems and by directing network traffic to resources that it believes are best suited to meeting service goals.

EWLM can encompass a broad heterogeneous infrastructure, including network resources. Alternatively, you can deploy EWLM to only handle subsets of your IT infrastructure, depending on your current technology requirements.

For more information on the IBM Enterprise Workload Manager, see “IBM Enterprise Workload Manager” on page 19; for the business, organizational, and management benefits of virtualization, see Appendix A on page 37.

Systems technologies

The IBM Virtualization Engine leverages pre-tested technologies to help speed implementation. These technologies are integrated and delivered with certain IBM eServer models. The technologies include:

- ▶ VLANs - help provide network virtualization capabilities that allow you to prioritize traffic on shared networks
- ▶ Hypervisor - supports partitioning (including Micro-Partitioning) and dynamic resource movement across multiple operating system environments; it enables all of the following technologies on eServer iSeries and pSeries:
 - Dynamic Logical Partitioning (DLPAR), which allows system resources to be grouped into logically separate systems within the same physical footprint. With new capabilities such as “uncapped partitions”, processors can be shared between partitions based on business needs (defined via policies).
 - Virtual Ethernet, where partitions can communicate together within a high-speed virtualized network, yet all have outside access via routing through one set of physical I/O devices.
 - Virtual I/O, where I/O resources such as disk, tape, and CD-ROM can be shared between partitions.
 - Capacity on Demand, where system resources like processors and memory are made available on an as-needed basis. Once activated, the resources can be used temporarily (On/Off CoD) or permanently (CUoD) when the business need arises.
 - Multithreaded CPUs where applications can increase the overall resource utilization by virtualizing multiple physical CPUs through the use of multithreading.
 - Multiple Operating Systems are supported on an IBM eServer—i5OS, AIX 52L, Linux, and Windows.
 - Micro-Partitioning provides the ability to allocate less than a full processor to a logical partition; it provides support for a single processor being shared by up to ten logical partitions.

Systems services

IBM Virtualization Engine systems services build on the capabilities of IBM eServer and IBM TotalStorage products and tie individual systems services together into one heterogeneous environment.

About skills: Increasingly, as organizations become more virtualized and systems are acquired more dynamically, the ability to configure to get the most out of existing and new systems requires a common set of management tools.

The IBM Virtualization Engine goes a long way toward addressing this requirement for servers and storage by abstracting out the differences in implementation and providing a consistent and effective way to deal with different hardware and operating system implementations.

Suite for Servers

The IBM Virtualization Engine Suite for Servers can help provide virtualization and management of resources across select systems, both in an individual server and across an enterprise. It consists of the following key components.

IBM Tivoli Provisioning Manager

The primary value of this systems provisioning component is to increase the average utilization of resources, thus reducing the number of required resources and the associated management costs. It automates manual tasks of provisioning and configuring servers, operating systems, middleware, applications, storage and network devices. The IBM Tivoli Provisioning Manager works together with a set of thoroughly tested and documented “best practices” workflows that support typical IBM eServer and IBM TotalStorage topologies. These workflows can automate your unique data center processes including installing, configuring, and deploying servers, operating systems, middleware, applications, storage and network devices.

In addition to these eServer workflows, users can also download and use any workflows from Tivoli’s Orchestration and Provisioning Automatic Library (OPAL) Web page, or create their own customized workflows to implement their company’s data center best practices and procedures. These procedures can then be automated and executed in a consistent, error-free manner. In fact, using these automation workflows, IBM Tivoli Provisioning Manager has the ability to provision and deploy a server (from bare metal to full production) with the single push of a button, reducing the time needed to deploy a new server from days to minutes.

One of the key functions of IBM Tivoli Provisioning Manager is *platform provisioning*. This is the ability to create and manage OS Containers. OS Container is a collection of resources (processors, memory, I/O, storage) configured to host an OS Standard Server, Configured Blade Server, LPAR, VMware partition, VM ready to run Linux on zSeries, and so on. It also includes the creation and management of operating system images to load in containers. Finally, platform provisioning includes the ability to install the operating system in an OS Container.

IBM Tivoli Provisioning Manager supports IBM eServer xSeries and pSeries servers, including pristine installations. IBM eServer zSeries supports the cloning of manually installed Linux system images into automatically provisioned virtual machines under z/VM. IBM eServer iSeries supports provisioning of both Linux on POWER™ partitions, and Windows on Integrated xSeries Adapters and Servers.

Working with the IBM Intelligent ThinkDynamic Orchestration, you can extend the provisioning functionality to automate and orchestrate the provisioning of your IT resources, providing the appropriate resources on demand to your business applications.

IBM Tivoli Provisioning Manager can provide:

- ▶ Implementation of automation workflows
- ▶ Consolidation of testing center environments
- ▶ Application server support
- ▶ Storage capacity provisioning

IBM Grid Toolbox Multiplatform

The IBM Grid Toolbox Multiplatform is a comprehensive, integrated toolkit for creating and hosting grid services. This product includes material developed by the Globus Alliance (<http://www.globus.org/>), as well as a set of APIs and development tools to create and deploy new grid services and grid applications.

IBM Grid Toolbox Multiplatform V3 is designed to exploit the existing virtualization on IBM eServer and TotalStorage families. It provides the ability for logical partitions, blades, and virtual machines to manage work across a heterogeneous, distributed environment. These programs can be responsible for managing specific tasks on the servers, for

creating/deploying new services, and for participating in distributed system communication and management.

The IBM Grid Toolbox Multiplatform includes a limited integrated hosting and development environment capable of running grid services and sharing them with other grid participants, such as grid service providers and grid service consumers. It also provides a set of tools to manage and administer grid services and the grid hosting environment, including a Web-based interface, the Grid Services Manager.

Currently, the toolbox supports the IBM AIX® operating system on IBM eServer pSeries® and iSeries™ servers, i5/OS on IBM eServer iSeries, and the Linux operating system on IBM eServer xSeries®, eServer pSeries, IBM eServer iSeries and IBM eServer zSeries® servers.

Derived from the common runtime used for the IBM Grid Toolbox Multiplatform, the IBM Virtualization Engine supports industry open standards that support the interoperability of heterogeneous IT resources. Standards and technologies such as the Open Grid Services Architecture (OGSA), Linux, XML, WSDL, and SOAP give IT solution providers a more cost-effective way to provide solutions for heterogeneous environments, while increasing the compatibility between components within an infrastructure. Support for these standards and technologies are planned across all IBM eServer platforms.

The on demand Operating Environment, discussed in Chapter 3, “The on demand Operating Environment” on page 27, will encompass the objectives of grid computing, delivering them through an implementation of the OASIS industry standard Web Services-Resource Framework (WSRF)¹, WS-Naming and related standards. This will allow IBM operating systems to participate in many aspects of a grid-like infrastructure without requiring the IBM Grid Toolbox Multiplatform. Interoperability between these systems and Globus-type grids will be delivered through Globus adoption of the same standards. These standards were primarily the work of Globus and IBM in conjunction with HP, SAP, Akamai, TIBCO and Sonic.

IBM Enterprise Workload Manager

IBM Enterprise Workload Manager enables you to automatically monitor and manage heterogeneous workloads across an IT infrastructure to better achieve predefined business goals for end-user performance. It provides end-to-end resource optimization and load balancing of IT resources in heterogeneous, multi-tier application and server environments.

¹ <http://www.globus.org/wsrf/default.asp>

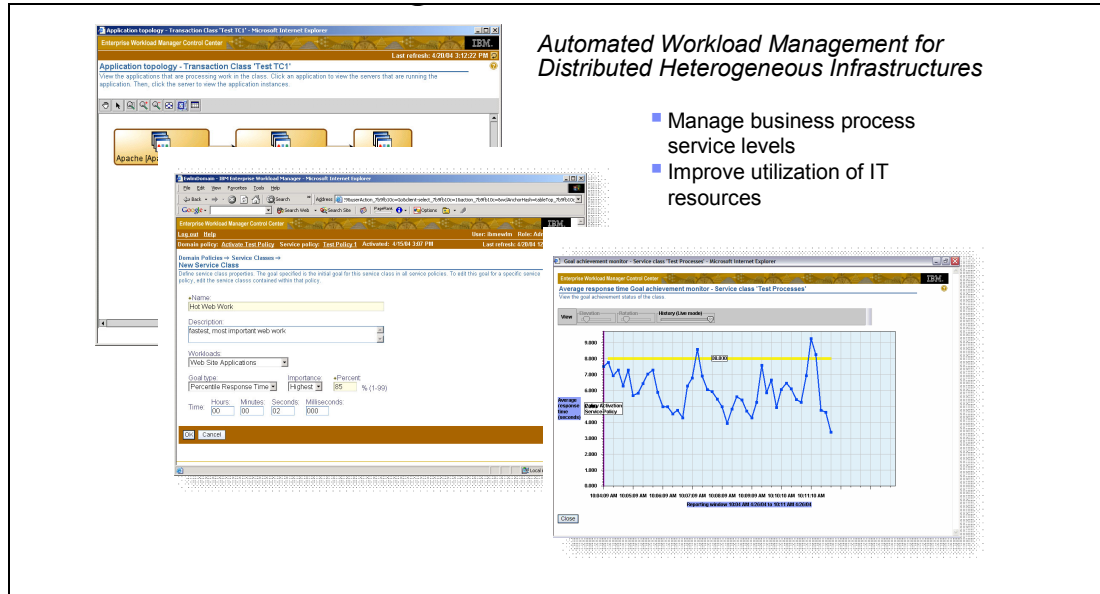


Figure 2-1 Enterprise workload monitoring

The goal of EWLM provides the ability to identify work requests based on service class definitions, track performance of those requests across server and subsystem boundaries, and manage the underlying physical and network resources to achieve specified performance goals for each service class.

EWLM management capability will be delivered in stages across several releases. In the first release, it will assist network routers in distributing work to those servers best able to achieve service class goals for the specific category of work being managed.

By bringing this self-tuning technology to the set of servers and routers, IBM Enterprise Workload Manager can help enable greater levels of performance management for distributed systems. Its extremely flexible management reach can encompass your infrastructure or be deployed to handle a critical subset of your applications.

IBM Enterprise Workload Manager will offer monitoring capabilities to allow analysis of activities across your enterprise, and manage resource allocation and utilization based on business requirements.

Enterprise Workload Management technology provides the functionality to manage the heterogeneous servers that are defined within what is called an Enterprise Workload Management domain. Each Enterprise Workload Management domain can have hundreds or even thousands of server resources within it.

Enterprise Workload Management essentially consists of four pieces, as follows:

► Domain policy

These are the policy-based service level objectives for the workload that will run on the servers within the Enterprise Workload Management domain. The objectives are defined by establishing domain policies and service policies. These definitions are entered using the administrative user interface and are sent to the domain manager where they are stored in XML format.

While only one policy can be active at any time, you can predefine different policy settings, changing them dynamically as your business needs dictate, such as having different policies for prime shift, off-shift, or weekend processing.

- ▶ Domain manager

There is one physical domain manager for each Enterprise Workload Management domain of servers that are being managed. The domain manager is responsible for storing the domain and service policies. In addition, it dynamically manages the server topology within the Enterprise Workload Management domain and maintains it, along with the state of each server. It also handles communication with the administrative user interface as well as with all servers in the Enterprise Workload Management domain.

Finally, it aggregates the performance data collected by each of the management servers into a comprehensive, end-to-end view of overall service class and server performance.

- ▶ Managed server component

This is code that runs on each server within the Enterprise Workload Management domain being managed. Each server contains an implementation of the code, which is based on the Application Response Measurement (ARM) standard from the Open Group², either delivered with the operating system or installed along with Enterprise Workload Management, depending upon the platform.

This Enterprise Workload Management ARM implementation interfaces with the operating system as well as with instrumented software running on the server. Information is provided which allows the tracking of workload that is flowing through the servers in the Enterprise Workload Management domain.

Information that is gathered is sent to the domain manager to be summarized and analyzed. This information gives insight into how the server is running from a utilization standpoint, as well as how well the pieces of the workload are being serviced. The information is then available to be accessed from the administrative user interface.

- ▶ Administrative user interface

This is the point from which the Enterprise Workload Management domain is controlled. The domain policy and service policies can be entered from this interface. The interface is also the tool used to control when domain policies and service policies are to be implemented. This interface can be used to instruct the domain manager to implement certain domain policies and service policies for the environment. The domain manager then communicates with the managed server components running on each of the servers.

By being able to know the utilization of a server and the service level objective for the workload running on that server, it is possible to determine if the service level objectives are being met for the workload as well as whether the server can be utilized at a higher rate while not jeopardizing the SLA for the workload.

IBM Director Multiplatform

IBM Director Multiplatform is the component of the Virtualization Engine that provides a simplified view of IBM eServer hardware from a centralized server. It delivers a comprehensive systems management solution for heterogeneous environments, such as support of IBM eServer BladeCenter, iSeries, pSeries, and xSeries servers. IBM plans to add similar functionality in future releases of IBM Director Multiplatform for supported platforms, including IBM eServer zSeries servers.

IBM Director Multiplatform's core infrastructure is designed to provide a single point of control for managing up to 5,000 systems.

² <http://www.opengroup.org/tech/management/arm/>

The Director Multiplatform product structure consists of three components:

- ▶ The Management Server is the main component or aggregation point for managing the agents.
- ▶ The Agent provides management data and function to the Management Server. Depending on the operating system and hardware platform, the Agent capabilities will vary.
- ▶ The Console provides a graphical user interface (GUI) with a consistent look and feel for all the servers and devices maintained by the Management Server.

By leveraging industry standards, Director Multiplatform provides an extensible architecture to enable solutions for easy integration with other systems management tools and applications, such as Tivoli Enterprise, Tivoli NetView, HP OpenView, Microsoft SMS, CA Unicenter, BMC Patrol, and NetIQ.

Using the collection of IBM Director Multiplatform tools, many of the administrator's manual tasks can be automated to proactively and remotely manage systems such as discovery, event logs and action plans, file transfer, inventory collection, process management, resource monitors and thresholds, and more. The predictive and proactive capabilities associated with alerting and real-time system diagnostics help maximize server uptime and reduce service downtime costs. Finally, and perhaps most importantly, the product's ability to support cross-platform IBM and non-IBM systems (Intel-based) means customers can protect existing infrastructure investments and can manage heterogeneous environments.

For xSeries servers and BladeCenter, Director Multiplatform provides additional exploitation of the hardware through inventory and alerts, asset tracking, diagnostic monitoring, capacity management and troubleshooting.

Virtualization Engine Console

The Virtualization Engine Console is a Web-based console that provides two functions: a consolidated view of enabled enterprise resources via a health function, and a launch point for the Virtualization Engine system services.

The Virtualization Engine Console health function allows customers to use system resources across platforms. This is done using input from any of the following IBM products: IBM Director Multiplatform, Cluster Systems Management, iSeries Navigator, and IBM Tivoli Monitoring.

The Virtualization Engine Console is based on one of the first results of the IBM autonomic computing initiatives, the IBM Integrated Solutions Console. This provides a common framework on which other consoles can plug in using a portal approach. This technology is one that IBM is extending across many of its product lines.

Figure 2-2 on page 23 shows an example of the main page of the Virtualization Engine Console. The left navigator allows you to move between the health function and the launch pad function. You can also see a "dashboard", which you configure to use the aspects of your environment you most often review.

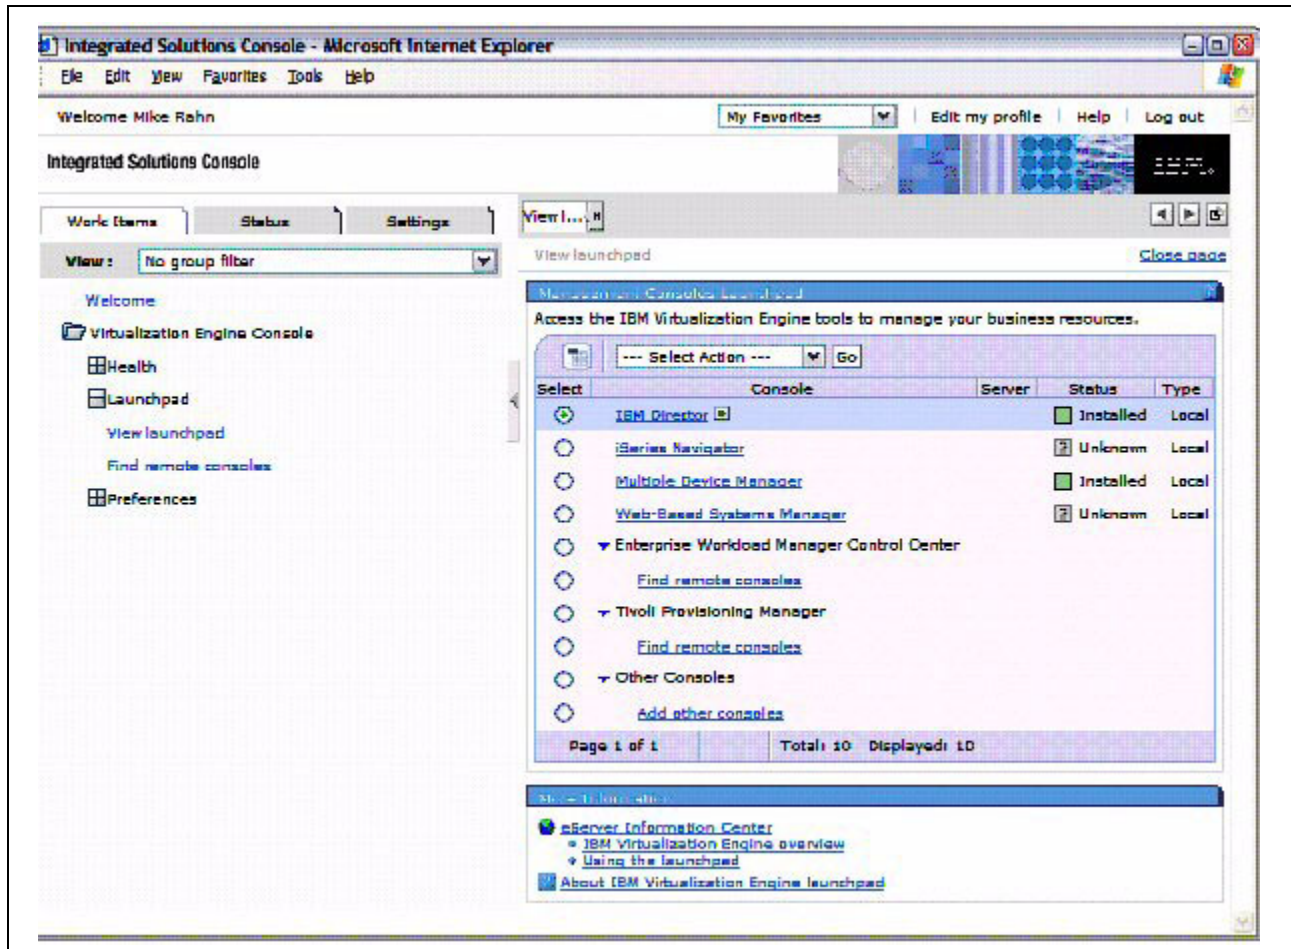


Figure 2-2 The Virtualization Engine Console

Platforms supported by the IBM Virtualization Engine will be dependent on the IBM eServer platform and the Virtualization Engine component. Initially included are:

- ▶ IBM eServer iSeries, pSeries, xSeries
- ▶ AIX 5L, i5/OS
- ▶ Microsoft® Windows or Sun Solaris
- ▶ IBM TotalStorage

In the future, IBM Virtualization Engine intends to support z/OS. For specific component and platform availability information, refer to the IBM Virtualization Engine announcement.

Suite for Storage

The IBM Virtualization Engine Suite for Storage can help provide the capability to virtualize and manage storage devices and files. The key products are described in the remainder of this section.

IBM TotalStorage Productivity Center

The IBM TotalStorage Productivity Center is the first offering to be delivered as part of the IBM TotalStorage Open Software Family. The IBM TotalStorage Productivity Center is an open storage infrastructure management solution designed to help reduce the effort of

managing complex storage infrastructures and help improve storage capacity utilization and administrative efficiency. It simplifies the management of traditional and virtualized SAN environments, and is designed to enable an agile storage infrastructure that can respond to on demand storage needs.

The IBM TotalStorage Productivity Center is comprised of a user interface designed for ease of use, in addition to these components:

- ▶ IBM Tivoli® Storage Resource Manager
- ▶ IBM Tivoli SAN Manager
- ▶ IBM TotalStorage Multiple Device Manager

The IBM TotalStorage Productivity Center is designed to help improve:

- ▶ Administrator efficiency
- ▶ Capacity utilization
- ▶ SAN performance
- ▶ Application availability

With IBM TotalStorage Productivity Center, you can:

- ▶ Monitor and track the performance of SAN-attached Storage Management Initiative Specification (SMI-S)-compliant storage devices
- ▶ Manage the capacity utilization and availability of file systems and databases
- ▶ Monitor, manage, and control (zone) SAN fabric components
- ▶ Manage advanced storage replication services (Peer-to-Peer Remote Copy and FlashCopy®)
- ▶ Automate capacity provisioning to help improve application availability

IBM TotalStorage SAN Volume Controller

IBM TotalStorage® SAN Volume Controller enables storage virtualization and can help increase the utilization of existing capacity. It is designed to increase the flexibility of your storage infrastructure by enabling changes to the physical storage with minimal or no disruption to applications.

Now with expanded support for many non-IBM storage systems, the IBM TotalStorage SAN Volume Controller can enable a tiered storage environment to better allow you to match the cost of the storage to the value of your data. It is designed to allow you to centrally manage multiple storage systems to help enhance productivity, and combine the capacity from multiple disk storage systems into a single storage pool to help increase utilization. It also allows you to apply advanced copy services across storage systems from many different vendors, to help further simplify operations.

IBM TotalStorage SAN File System

The IBM TotalStorage SAN File System (based on IBM Storage Tank™ technology) provides a network-based heterogeneous file system for data sharing and policy-based storage management in an open environment. It is designed to help reduce the complexity of managing files within SANs. The IBM TotalStorage SAN File System is designed as a highly available file system for SAN-attached storage—one that is designed to provide a network-based heterogeneous file system for data sharing and centralized policy-based storage management in an open environment.

The IBM TotalStorage SAN File System is designed to enable host systems to plug-in to a common SAN-wide file structure. With the IBM TotalStorage SAN File System, files and file

systems are no longer managed by individual computers; instead, they are viewed and managed as a centralized IT resource with a single point of administrative control.

The IBM TotalStorage SAN File System provides a common file system for UNIX®, Windows® and Linux® servers, with a single global namespace to help provide data sharing across servers. It is a highly scalable solution supporting both very large files and very large numbers of files without the limitations normally associated with Network File System (NFS) or Common Internet File System (CIFS) implementations.

The IBM TotalStorage SAN File System is designed to help lower the cost of storage management and enhance productivity by providing centralized and policy-based storage and data management for supported heterogeneous server, operating system and storage platforms.

The IBM Virtualization Engine systems services are tested together and integrated to support the delivery of an on demand Operating Environment. This is implemented in such a way that nonparticipating environments can be managed by and participate in on demand, but do not have to make major changes to benefit from the Virtualization Engine systems services and technologies. An on demand Operating Environment can provide an IT infrastructure that will help your business better align its IT operations with its business strategy.

Figure 2-3 shows the exploitation and integration of a number of the IBM Virtualization Technologies and System Services.

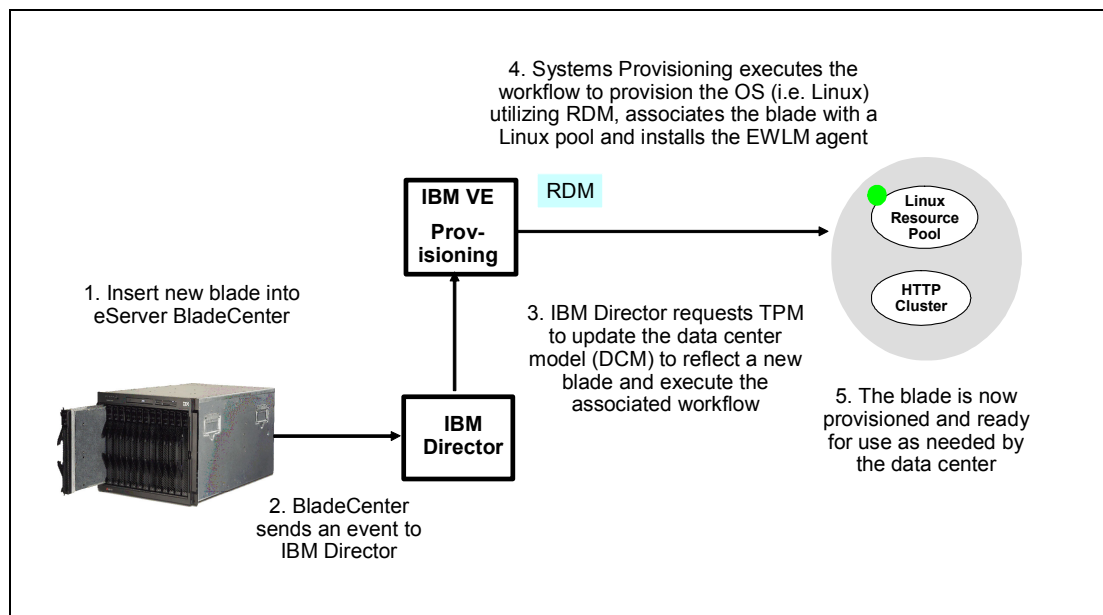


Figure 2-3 Virtualization Engine integration

In the on demand Operating Environment, the Virtualization Engine's services and technologies are leveraged to provide a virtualized, flexible and dynamic infrastructure that can support on demand business services. For organizations that are not ready for a full on demand Operating Environment, the Virtualization Engine services and technologies have their own, combined value proposition when running on IBM and heterogeneous systems.

Summary

In this chapter we examined the key role of the IBM Virtualization Engine (VE) and discussed its functions and services. We also looked at how the Virtualization Engine will be leveraged by the on demand Operating Environment, as well as how they improve and deliver on the virtualization of key infrastructure components.

In the next chapter we examine the emerging on demand Operating Environment, technology, architecture, and services, and discuss how this level of virtualization will help the on demand business.



The on demand Operating Environment

Overview

The on demand Operating Environment builds on the key technologies and services already described in this paper. That is, it is a service-oriented infrastructure that exploits industry open standards. The infrastructure in turn views every application or resource as a service implementing a specific, identifiable set of business functions.

Services communicate with each other by exchanging structured information, thus allowing new and existing applications to be quickly combined into new contexts. All interactions flow through an enterprise service bus. This does not mean all interactions require network communication and XML messages. It is also possible to use optimized communication and encoded messages. In some cases, two services may bind to a local program call. The matching of service requesters to providers can be done early, prior to deploy time, or very late via dynamic discovery, as described in “Web Services” on page 12.

In addition to the business functions, the on demand Operating Environment also implements management interfaces as services in order to participate in broader configuration, operation and monitoring environments.

Contextually this means that the SOA-based infrastructure exploits an enterprise service bus (ESB) for two purposes: running on demand business services, and running the infrastructure. However, it is interesting to note that this is in fact the same technology, the same services, and the same infrastructure—but used for different purposes. This blurs the line between operating systems and middleware. Rather than focus on the delineation between hardware, operating system functions and middleware, the on demand Operating Environment focuses on the delivery of functions as business services irrespective of how they are implemented. Using common software and services means increased reliability, availability and integrated infrastructure.

What this means is that both applications and system management services can be deployed independent of the underlying implementation. The on demand services are described as abstract services which can be interfaced with an application hosting environment,

middleware, operating system or directly onto hardware through an adapter or connector into the underlying resource.

Let's examine a simple example of how this might work in principle. Since we are interested in the operation and exploitation of a dynamic, virtualized infrastructure, the actual process for creating operating system images will vary from system to system and from hardware platform to hardware platform. It will also depend on what type of container¹ the operating system image runs in—a logical partition, a virtual machine, a blade and so on, or even on a dedicated server.

The IBM Virtualization Engine, IBM Tivoli Provisioning Manager, and IBM Director Multiplatform can provide an abstraction of the interface and build workflows that can define, configure, build and deploy operating system images into containers. However, they provide these functions in a generic, industry standard way on which other hardware, software, and middleware developers are able to build on and incorporate into their products and offerings without making the IBM Virtualization Engine a prerequisite. While these offerings provide important virtualization capability, they do not conform to a key objective of the on demand Operating Environment, which is to deliver an open, industry-standard, accessible set of services interfaces for the creation and management and so on of computing resources.

In Figure 1-2 on page 4, we introduced the idea of isolation. The on demand Operating Environment exploits this isolation and turns it into a discrete business service by defining an abstract interface that is independent of the underlying implementation.

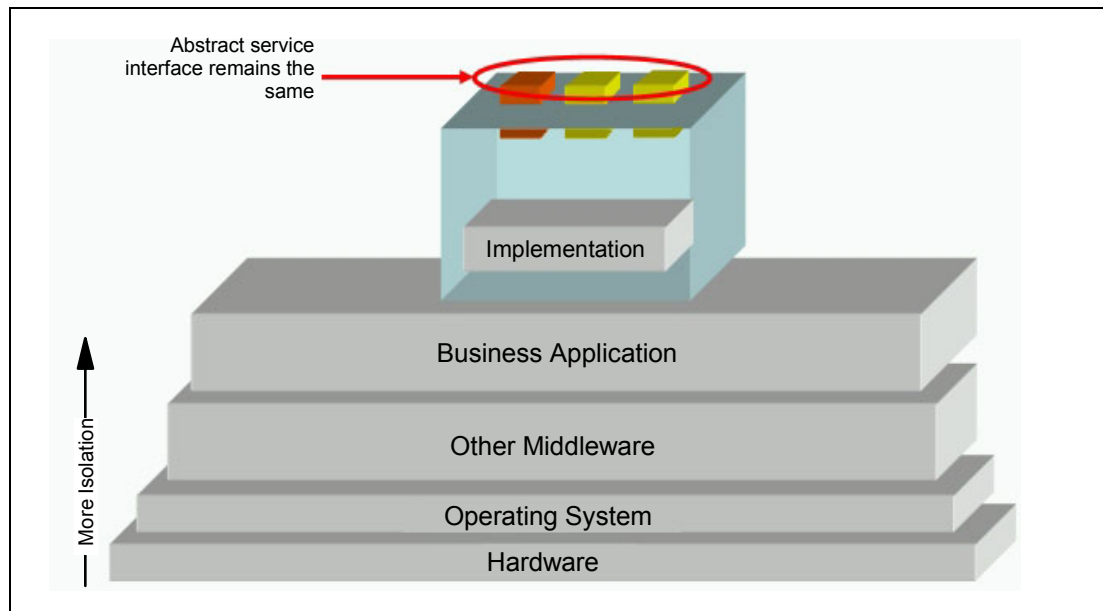


Figure 3-1 Implementation independence

You can see this in Figure 3-1. This is the same as Figure 1-2 on page 4, with an abstract interface added. The underlying implementation could be any of, a combination of, or all of the layers shown. In a truly heterogeneous environment, the interface would be the same, typically defined through the Web Services-Resource Framework—but the implementation could be implemented by any vendor, at any level they deem appropriate.

For example, a blade server vendor might implement a direct set of controls on the hardware, interfaced via the SOAP protocol on an embedded processor; whereas another vendor may

¹ The term “container” is used as the generic environment for installing, deploying and running business services and operating systems, as it is non-specific or virtualized.

chose to build on an open source stack of say, Linux, Apache and JBOSS, accessed via Web Services.

Now, consider the impact of this across all the services and information required to run your data center in isolation, in multiple data centers, or in an in-house or outsourced utility service. You would need to work not only with hardware resources such as servers, and storage, but also with pure services such as billing and metering, along with infrastructure requirements and automation such as security, logging, problem management and configuration.

The on demand Operating Environment allows all these services to exist and to interoperate through abstract services interfaces. None of them need to be on the same physical system; there is nothing location specific about the services; and the decision on implementation is a function of requirements rather than platform or vendor. There can be single implementation or multiple implementations of each service, if this is a requirement. The interface to the services is through the enterprise service bus, and services can both publish and subscribe to events on the bus. Combined, these services create a meta-operating system.

Figure 3-2 shows the full on demand Operating Environment with all the envisaged services. Those shown above the “bus” are related to business services and functions, while those below the bus are related to infrastructure management. The bus and thus the services can be called using a variety of protocols such as SOAP, JMS, and other forms of messaging. The bus can be bridged into using other messaging technology. Infrastructure services are run within the enterprise service bus, but are not invoked directly by business logic.

Since an enterprise service bus can co-exist with others, you might decide to have one for application-level messaging and services and another for infrastructure management, with the appropriate bridges and gateways between them. The enterprise service bus supports a broad spectrum of ways to get on and off the bus, including on-ramps for legacy applications or business services.

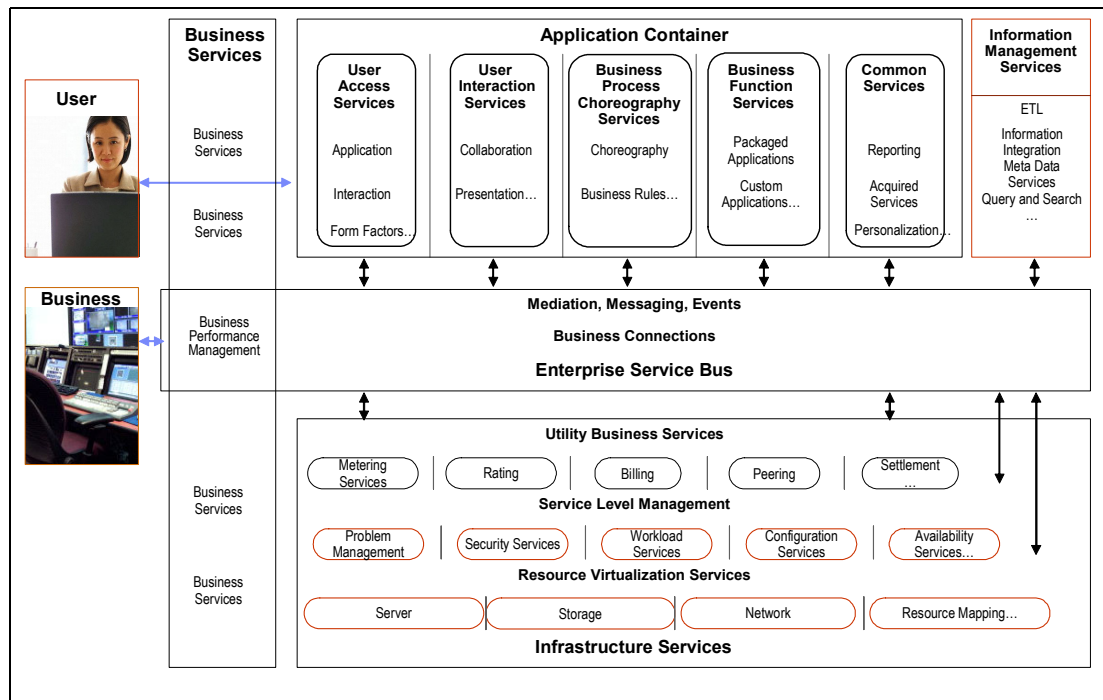


Figure 3-2 The on demand Operating Environment architecture

For a deeper understanding of the implications and design of a service-oriented architecture and the enterprise service bus, refer to the discussion in “About Application Virtualization” in Chapter 1. A complete discussion of the use of these technologies for infrastructure management will be available as we approach the introduction of the on demand Operating Environment.

Summary

The on demand Operating Environment capabilities enable business flexibility and I/T simplification. There are two categories of capabilities: application integration and infrastructure management, and the objective is to evolve to an industry-standards-based, integrated, autonomic and virtualized I/T environment.

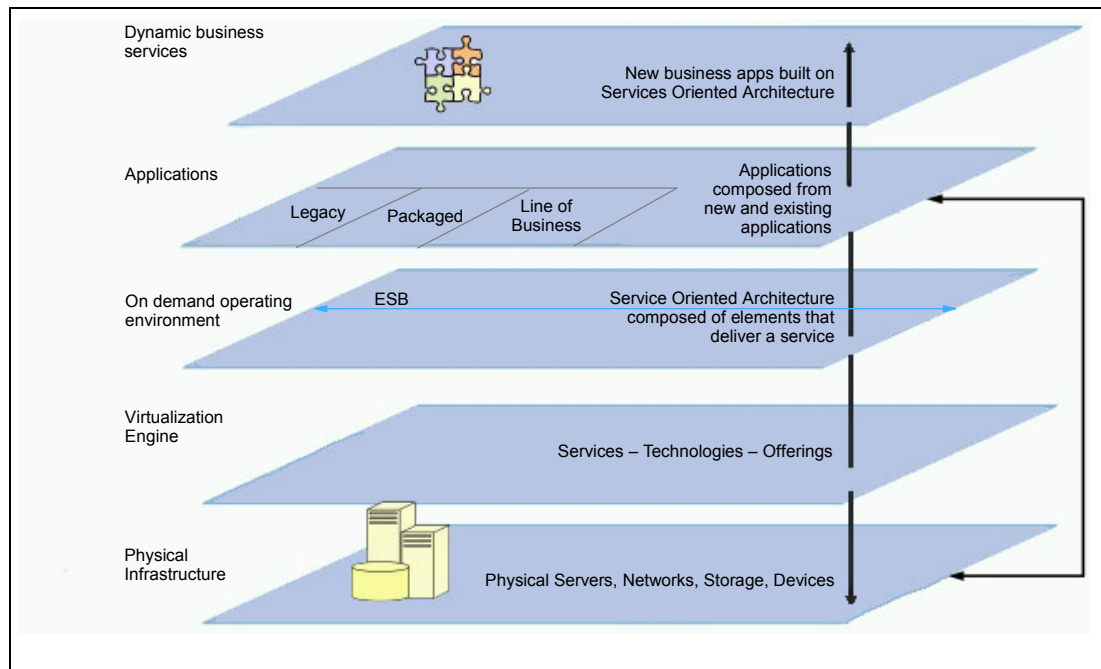


Figure 3-3 Abstraction, virtualization, on demand

Figure 3-3 shows the conceptual relationship and relationship hierarchy, from composed on demand business services through leveraged applications to underlying services and technologies, and the physical infrastructure on which they are run.

Adding a virtualization layer in between applications and system resources helps to balance the distribution of applications across these systems through a set of intelligent workload management services. These intelligent workload management services help control how applications consume these resources. The virtualization approach helps to balance resources among various applications to provide the right resources to each application. No matter how well applications are designed and built, the environments they run in must be highly available, secure, and flexible in order to address constantly shifting changes in demand. This is why we need to pay attention to the systems environment, because a good application can look really bad if it is not deployed in the appropriate computing environment.

Deploying virtualization technology is an important initial step to being able to link system resources to business priorities. On demand enterprises need the ability to associate the Quality of Service attributes assigned to applications with the correct system resources to deliver against these attributes, and IT managers need to work across the organization to

deploy the applications into an optimized system environment that meets the needs of these applications.

As its components embrace the component-based, SOA, and Web Services standards called for by the on demand Operating Environment and OE architecture, the Virtualization Engine becomes IBM delivery mechanism for the on demand Operating Environment infrastructure Services. It also builds on the industry-leading features and functions found in IBM middleware, operating systems and IBM eServer and Storage solutions, to deliver a new model for computing that includes both new and existing applications - a model that encourages and supports both traditional and new programming models and applications.

The on demand Operating Environment exploits the IBM Virtualization Engine, which in turn builds on the industry-leading features and functions found in IBM middleware, operating systems and IBM eServer and Storage solutions to deliver a new model for computing that includes both new and existing applications—a model that encourages and supports both traditional and new programming models and applications.



Getting started with virtualization

In this chapter we review key technologies and suggest ways in which organizations can get started with virtualization.

Overview

IBM on demand solutions are here today. They are helping you to make this high level of business performance a reality. So how do you get started? There is no single roadmap, and where you start depends on your environment. The following is therefore a broad list of areas you might want to consider for your IT infrastructure:

- ▶ Evaluate different technologies to obtain the values related to IT consolidation and IT simplification (for example, logical partitions).
- ▶ Utilize existing server virtualization capabilities, combined with clustering and BladeCenter technologies, for production, development, and testing.
- ▶ Exploit other existing virtualization capabilities, such as virtual networks and virtual I/O interfaces, to integrate applications, simplify the infrastructure, save resources and improve performance.
- ▶ Utilize existing storage virtualization capabilities (for example, total storage virtualization with the IBM SAN Volume Manager).
- ▶ Select standards that give you independence from the operating system and HW.
- ▶ Consider using Open Source such as Linux.
- ▶ Exploit a combination of portable language execution technology such as Java servlets and EJBs, DB2 Stored Procedures, and SAP Advanced Business Application Programming Language applications.
- ▶ Exploit data federation capabilities in DB2 UDB.
- ▶ Consider new application design methods and tooling.
- ▶ Evaluate the use of automation products and services to reduce the complexity of management to enable better use of assets, improve availability and resiliency, and reduce costs based on business policy and objectives.
- ▶ Upgrade your operating system and middleware products to the latest release level.

Start now

Your journey in transforming your enterprise into an on demand business can start today, with these steps:

- ▶ IT Simplification

The primary task for IT Simplification is to adopt appropriate standards. This should be viewed as a task involving the entire stack—from basic hardware, to operating platforms, middleware, and database and application development. In addition, you need to standardize on IT processes. Where possible, select standards that provide independence from the operating system and hardware. You should be looking at providing virtual environments for development, testing, and production.

- ▶ Plan and implement server virtualization

Based on your requirements and installed server infrastructure, you may be able to make significant management savings, improve resource utilization, or increase operational flexibility by implementing a virtual machine, logical partition, or BladeCenter-based server infrastructure.

Where the level of flexibility or complexity of operation warrants it, use Systems Provisioning within the IBM Virtualization Engine. In a complex, heterogeneous, multi-vendor environment, you might consider using the IBM Tivoli Intelligent Orchestrator and Provisioning.

- ▶ Plan and implement storage virtualization

There are numerous opportunities to benefit from the technologies and services provided for storage virtualization by the IBM Virtualization Engine and products such as the IBM SAN Volume Controller and the TotalStorage Subsystems.

For example, you might want to make higher system availability a priority by using a shared storage solution combined with server virtualization as a means to provide economical backup systems.

Alternatively, you might have workloads that individually have high peak storage requirements, but not simultaneously. In this case, you can use logical volumes, network-attached storage and other techniques to pool and share the physically installed storage and so on.

- ▶ Workload management

Plan and implement the IBM Virtualization Engine workload management monitoring agents. In order to be able to achieve self-managing systems that are able to balance competing workloads, you will need to be able to define service levels and workload objectives for different workloads in the environments they currently run in.

Experience has shown that organizations prefer to be cautious when implementing automated performance management. Starting to collect the relevant information as early as possible will allow you to get a better understanding of what your requirements are likely to be.

- ▶ Provisioning and orchestration

Look for opportunities to automate and orchestrate the dynamic definition and deployment of both homogeneous and heterogeneous servers. Where an organization has complementary workloads that can be deployed on a common pool of servers—and a high level of variability in server configuration is required—the IBM Tivoli Intelligent ThinkDynamic Orchestrator and IBM Tivoli Provisioning Manager can deliver significant flexibility and higher utilization through automated workflows.

Getting help

IBM Global Services offers a wide range of technology services to help companies assess, identify risk, design, and implement on demand virtualization and provisioning technology. Our skilled consultants and services experts are experienced and fully trained, and can help you prepare your infrastructure to leverage the immense benefits of IBM's new on demand virtualization technology components to help customers prepare for technology change and define the roadmap mapped against the client's IT environment. A description of the comprehensive services available to help you become an on demand business can be found in Appendix B, "The on demand Infrastructure Services offerings" on page 47.

Summary

Virtualization is the technology that simplifies access and management of IT resources, which is a prerequisite in an on demand world. It is an important use of technology for addressing the cost, Quality of Service and Time to Value challenges that business organizations face today. The current level of virtualization technology provides a broad set of options within servers, storage, networks, applications, systems management that can be exploited *today* to meet these business challenges in a systematic and successful manner.

The IBM Virtualization Engine is one of key deliverables of virtualization, and it will provide key common technologies and services which abstract individual system functions to be leveraged by the on demand Operating Environment. Individually and together, the IBM Virtualization Engine technologies and services offer organizations a wide and varied opportunity to better organize, configure, provision and manage their IT infrastructure.

The IBM Virtualization Engine will help extend system functions to a cross-platform perspective, while providing a common interface to these functions.

Just as with today's existing system functions, the IBM Virtualization Engine functions will have tight interlocks with the operating systems and storage entities, as well as microcode and hardware. The Virtualization Engine will be delivered as an evolutionary approach that protects customers' investments in current technologies. Open standards-based interfaces will be used so that the on demand technologies can be implemented across heterogeneous IT infrastructures. One of the open standards that forms the basis for these technologies is the Open Grid Services Architecture. IBM delivered the IBM Grid Toolbox Multiplatform in 2003.

IBM storage software is designed to provide autonomic capabilities that improve access to, and management of, stored data in an on demand world. This must be accomplished in a standard and open way that will dramatically reduce the total cost of ownership and improve organizational agility.

Our strategic focus is shifting from delivery of fundamental virtualization products to expanding heterogeneous server and storage support by adding advanced virtualization capabilities that improve management and integration across our entire server and storage product line.

These capabilities will integrate with systems management and IBM's overall on demand offerings, and will be the key elements of our on demand Operating Environment strategy to provide integrated, open, virtualized, and autonomic infrastructures for our customers.



Business, organizational, and management benefits of virtualization

Overview

This appendix provides background information on the business, organizational and management advantages offered by virtualization. This information will help you to understand the benefits that you may be able to derive from the various forms of virtualization; these benefits may vary enormously from organization to organization.

Included are discussions on Total Cost of Ownership (TCO), Time to Value (TTV), Quality of Service (QoS), and the impact of virtualization on the cost of IT.

Virtualization benefits

So how can virtualization in its many forms help you? In particular, how can it help reduce costs for hardware, software, and people, as well as improve time to market, time to customer, and qualities of service?

Virtualization can increase the utilization of existing components, some of which only run at 1-2% utilization today. With Virtualization, a single server can have the appearance of multiple virtual servers; you get more for less money, including less software cost. Similarly, storage can be shared more effectively among users as virtual storage.




Figure A-1 Server utilization before and after workload consolidation

Such techniques not only save money, but also enable you to implement new virtual capacity very quickly. This provides companies with a quicker time to market for your existing business services. They can more easily cope with peaks of demand from their users. In addition they can take multiple physical servers and make them appear as a single virtual server, bigger than its individual components.

System management personnel costs are often the largest part of IT spending, sometimes as high as 75%, because of the diversity of systems, interconnections, and quantity of equipment being managed.

By abstracting different operating systems to a single virtual environment, not only can management costs be reduced, but companies can also significantly improve the quality of service in terms of performance and availability delivered to their users. Thus, by using Virtualization in concert with the emerging autonomic concepts, management costs can be reduced considerably.

Many of the costs associated with systems management come from analyzing complex failures involving many different components. With Virtualization, the management of end-to-end processes can be reduced, hence significantly reducing the time taken to debug problems. You can also virtualize system interconnections such as networks and I/O paths, thereby eliminating many of the physical components, improving component performance, and making systems management simpler.

Value of virtualization

There are many forms of mature hardware and software virtualization technologies available for exploitation today, and each contributes to the business values of Cost of IT, Quality of Service (QoS), and Time to Value (TTV). Let's see how each of these categories can realize benefits from virtualization in more detail.

Cost of IT

Most IT organizations have business applications that run on a complex heterogeneous infrastructure, with varieties of servers, storage devices, operating systems, architectures and vendor products. These organizations are also being challenged to reduce costs for IT

resources, as well as the amount of skilled labor to manage them. At the same time they are required to support the business in a resilient, responsive and timely manner.

The challenges of supplying a cost-effective infrastructure with the right business support include:

- ▶ Underutilized assets that are dedicated to specific tasks.
- ▶ Over-provisioning of resources to meet the demand of the business at all times.
- ▶ Complexity caused by a high number of components in silos or islands of applications and IT organizations, resulting in increased system management and personnel costs, while making it slow or difficult to respond to new business needs and integrate them with existing applications.
- ▶ A fixed cost structure that is to a large extent unrelated to varying business results.
- ▶ Satisfying the business need with trade-offs between cost and resources, thus leading to less than optimal availability of the needed resources for production, test, development, and so on.

Ideally, IT organizations should be able to use their IT resources at higher utilization levels to reduce costs and achieve a better return on their investment. They should be able to simplify their infrastructure to reduce management cost and make it simpler to integrate solutions. The resources must be provisioned and re-provisioned according to business priorities based on the committed service level agreements of performance, availability and responsiveness for new demands.

The cost structure of the infrastructure must be flexible and variable, to match the actual business results. This will require new ways of planning how resources are financed and installed for both existing and new applications. The objective is to reduce costs, while at the same time maintaining the service levels and flexibility required by the business. Virtualization can help you meet all these challenges; in the following sections, we'll explain how.

Underutilized assets and over-provisioning of resources

The major cost centers in an IT budget are hardware, software, and staff related to systems management. It is a challenge to ensure that there is sufficient capacity installed and available, and no more than that. In the past, this has resulted in configurations with relatively low utilization of servers, storage and networks, mostly due to the following reasons:

- ▶ Production capacity has been oversized to minimize the risk of unsatisfied service levels.
- ▶ Development, test, and backup environments are often unused for large periods of time and thus are underutilized.
- ▶ Resources are dedicated to specific purposes and are not readily available for other uses.

These problems are often related to Windows- and UNIX-based systems, because the applications have been developed to run by themselves on dedicated servers with dedicated storage, with little or no data sharing. This has resulted in average utilization of Intel servers with Windows in the 2% to 5% range, and an average utilization of UNIX servers in the 10% to 20% range. One of the main reasons for dedicating servers is limited workload management and workload isolation capabilities, making it necessary to dedicate the resources for performance and availability reasons. Other reasons include change and problem management isolation, service windows, and security isolation.

Simplify IT via consolidation

Consolidation is about more than just replacing smaller, underutilized servers with a few larger, higher utilized servers. It is also about simplifying existing end-to-end IT infrastructures, including servers, storage, databases, applications, networks and systems

management processes. Simplification also provides a more efficient and stable foundation for growth and new solution development, which is a logical first step toward deploying an on demand Operating Environment.

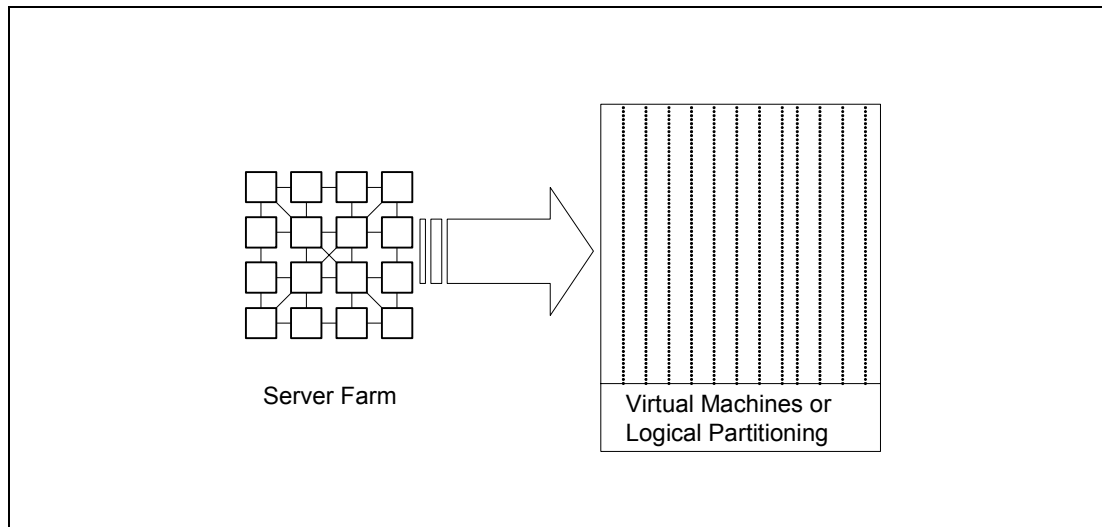


Figure A-2 Server consolidation: many into one

There are different ways of performing consolidation; the two most important methods for servers are workload consolidation and server virtualization.

- ▶ Workload consolidation means multiple workloads are consolidated under the same server and under the same operating system. For different types of applications, this requires effective workload management and operating system functions that isolate workloads so they do not impact each other adversely. Workload consolidation may complicate system management processes and extend the migration period if different types of applications are being consolidated.
- ▶ Server virtualization means that a physical server is divided into multiple virtual servers. Each virtual server behaves like a real server from an application and system point of view. Typically this will increase the utilization of the physical servers as the workload profiles intermesh and peaks fill troughs of use. It is also possible to decrease the complexity, and hence cost, by cloning systems and their management.

Traditionally, installations pay for infrastructure whether it is utilized or not. Aligning costs with business trends requires a new way of investing, where the vendor and the customer share the risk based on a long-term partnership, and where installations pay only for what they are using.

Virtualization aids migration

With virtualization, the virtual storage environment provided by a SAN can be used to enable smooth transitions to new applications. Advanced storage functions allow virtual copies of data to be used and reused in testing, while virtual servers and virtual storage permit a smoother migration to new systems, thereby reducing the risk of interruption to production systems.

Virtualization enables disaster recovery

Virtualization also enables disaster recovery, because instead of having dedicated hardware, you can cluster together a group of virtual servers. Those servers may be on machines which normally do other work, but in the event of a disaster are reassigned to meet the needs of the enterprise.

Virtualization reduces cost

Virtualization enables higher resource utilization, thus reducing the need for installed resources to support the business. It also provides the opportunity for simplification and standardization of the infrastructure, leading to reduced system management cost. Most importantly, Virtualization provides the flexibility that makes it possible to implement the emerging on demand Operating Environment or “pay only for what you use” solutions in an efficient and practical way.

Quality of Service

Consistent high availability and performance are basic qualities of service that affect how users perceive your business. Consider the following:

- ▶ Internet users perceive slow sites as unavailable, and move to other sites. When this happens, you lose market share!
- ▶ Failures and poor performance can reduce employee productivity and customer satisfaction.
- ▶ Failing to meet peak requests and customer expectations results in immediate loss of business.

Consequently, your IT infrastructure should have the following end-to-end characteristics:

- ▶ Continuous availability.

Today's applications typically consist of several layers or tiers of systems, each with its own availability characteristics and potential points of failure, which need to be taken into consideration when designing the environment. There is no time for outages, planned or unplanned; the luxury of having a “window” to perform backups has vanished. As a result, changes to the IT infrastructure are often in conflict with what the user or customer wants.

- ▶ Systems management solutions need to be end-to-end—and based on continuous availability.

These solutions must manage end-to-end environments rather than individual resources. Business processes and operations are linked across the company, including people and the IT infrastructure. They need to be managed in the same way, seamlessly and securely through a consistent interface.

The management process must provide an end-to-end topological view of the resources that associate a business process with the application. This will offer a clear understanding of the impact of any one resource, and allow for the prioritization of actions based on real business needs. Virtualization can insulate applications and users from outages or performance bottlenecks, as well as provide flexible delivery with high availability and consistently good performance.

Higher availability

As you put more components into a system, the more diverse and complex it becomes. It can be difficult to integrate many diverse components and processes into a stable system. It is also difficult to analyze the source and cause of failures. To ensure higher availability, different technologies and procedures need to be put in place, for example:

- ▶ Hardware and software clusters, to balance workload and provide fault tolerance
- ▶ Automation of systems management processes
- ▶ Automation of workload scheduling and control end-to-end
- ▶ Conformance to policies and processes to ensure high availability

The management of end-to-end applications requires interaction with a large variety of components. The objective is to link business priorities with IT requirements in a business

format. This needs to happen dynamically to avoid unavailability and achieve consistent performance. The IBM Virtualization Engine, discussed in “The IBM Virtualization Engine” on page 15, will deliver an IBM Enterprise Workload Manager function that specifically addresses this opportunity.

Policy-based Quality of Service (QoS)

An important idea we have introduced is management of systems by business policy rather than by IT constructs. Policy-based QoS is an extension that centralizes the controls for setting QoS across different workloads, and also defines how the shared resources are managed to ensure Service Level Agreements are met.

Business Service Management

Business Service Management virtualizes the IT environment in business terms. This allows for management of IT resources in the context of their impact on business processes. By virtue of this virtualization, service objectives can be established, honored, and monitored.

In the past, IT organizations have struggled with capturing costs and services that are associated with the delivery of resources for a specific business function. Service level objectives for the business application will be established and monitored based on business goals and requirements. Business Service Management monitors, meters and tracks, and provides for capabilities to handle accounting for charging back infrastructure costs to business units.

Achieving better integration

Business integration is a combination of resources, people, process and information serving to optimize the processes within an enterprise and with its partners and customers. On demand business requires a dynamic infrastructure that is integrated with business processes. The true integration challenge is to respond to business demands quickly. This requires an integration strategy based on a robust infrastructure. This infrastructure needs to provide standard services, standard interfaces, automation and management of value chain processes in the extended enterprise.

There are many opportunities to use virtualization to address issues of continuous availability, performance and disaster recovery that you can implement today, which is discussed later. First, let's look at Time to Value.

Time to Value (TTV)

How does on demand affect Time to Value, and what is the role played by virtualization? In this section, we examine the need for existing applications, new applications, and applications which perhaps had been dismissed as not possible or affordable. If IT can react quickly to the needs of the business, in real time or even in days, weeks or months, then it will have contributed to the value of the business. But this has been a challenging proposition: under-provisioning of resources could result in loss of business, and over-provisioning can cause IT to be too expensive.

A major challenge facing IT has been its inability to respond quickly enough, in either the short term for extra capacity to support a marketing initiative, or in the long-term integration of applications for new business initiatives. The primary reasons for this shortfall are:

- ▶ Inability to bring new resources online quickly enough
- ▶ Lack of resources for development, testing, and integration
- ▶ Complex, inconsistent environments that necessitate complex testing and problem determination

- ▶ Lack of standards, which makes it difficult to integrate new applications with existing ones
- ▶ Applications tied to specific platforms, sometimes forcing unwanted change as well as limiting choices when you do want to make a change
- ▶ Different business cultures within silos in an organization wanting to do it their own, unique way

Developers want to develop to one standard. Users want to see one system and one application, and to make use of accurate and relevant information. Users also want easy ways to communicate and collaborate. Communication and collaboration is between people, applications, data sources and repositories. Data sources encode business information; applications enable business logic that achieves the communication and collaboration.

On demand provides a number of technologies to address these issues. First we look at the existing virtual infrastructure capabilities which facilitate improved TTV; later we discuss the emerging virtual environment that applications can run in, and the means to access data that will facilitate solutions.

Provisioning for test and development

Systems used for test, development and education may comprise a significant segment of infrastructure cost, especially where there is little sharing between development groups—and attempts to control the cost may lead to development delays. Hardware cost is reduced, but expensive developers are left waiting for test time! As a result, the end product can often be delayed, thus reducing revenues.

With virtual systems sharing physical servers, however, developers can start development and test immediately. They have the opportunity to deliver a better quality application—and the virtual separation may also provide better security.

Bringing new resources online quicker

Whenever hardware is dedicated or tied to applications and business units, a new requirement may involve installing new hardware and software, establishing connectivity, and then testing it. But this approach poses two risks: you may purchase more than you need, or you may not purchase enough. Both cases have a negative effect.

In trying to avoid these issues, you may have been conservative in negotiating your SLAs. As a result, you may not be able to respond quickly enough to a peak or trough. However, if you could manage to minimize the risk involved, you'd be able to deliver aggressive support of the business and opportunities you may have passed over—this is Time to Customer (TTC).

Virtual servers can be used to deliver capacity, perhaps temporarily, by borrowing from development or non-key applications instead of purchasing more servers. The business can use this ability when purchasing new hardware, and can share it among its key applications. Virtualization and provisioning ensures servers are used more efficiently.

Bringing new applications on quicker

Developing new applications faster, and making them available earlier, can significantly affect a company's revenue stream and improve your customers' satisfaction. The key to this is using industry-standard interfaces and standard ways of connecting into existing services; this will enable you to write applications faster, and easily connect to existing services. Integration will also be faster and less complex.

Enabling new business opportunities

Previously, creating new applications may have been considered too expensive because the resources were not available to develop and test them. Now, with virtual resources, the cost

may be considerably less. Alternatively, by using multiple systems and virtualizing them as one, you could design applications (for example, grid applications) which were previously not viewed as possible. Compute-hungry applications that are too expensive to run on dedicated resources could use white space at slack periods, in effect enabling cheaper computing (see Figure A-1 on page 38).

The value of storage virtualization

The main value of SAN-based storage virtualization today lies in the ability to use a disparate range of new and legacy storage with varying performance, functional and availability attributes—and present it as a single storage image with a unified management interface to a wide range of servers.

IBM has already delivered the first version of a storage virtualization server called the IBM TotalStorage SAN Volume Controller. It implements Block Virtualization technology to help reduce the management complexity and costs of a heterogeneous SAN-based storage environment. Based on virtualization technology, the IBM TotalStorage SAN Volume Controller supports a virtualized pool of storage from the storage subsystems attached to a SAN.

The IBM TotalStorage SAN Volume Controller offers these benefits:

- ▶ Centralized control for volume management
The IBM TotalStorage SAN Volume Controller is designed to help IT administrators manage storage volumes on their SAN, helping to combine the capacity of multiple storage controllers—including storage controllers from other vendors—into a single resource, with a single view of the volumes.
- ▶ Avoidance of downtime for planned and unplanned outages, maintenance and backups
IT administrators can migrate storage from one device to another without taking the storage offline, and can reallocate, scale, upgrade and back up storage without disrupting applications.
- ▶ Improved resource utilization
The IBM TotalStorage SAN Volume Controller is designed to help increase storage capacity and uptime as well as administrator productivity and efficiency, while leveraging existing storage investments through virtualization and centralization of management.
- ▶ A single, cost-effective set of advanced copy services
The IBM TotalStorage SAN Volume Controller supports advanced copy services across all attached storage, regardless of the intelligence of the underlying controllers.

The next step in storage virtualization will be to deliver virtualization servers that create a Common File System, where data can be shared among a heterogeneous set of application servers running different operating systems and architectures.

Storage virtualization can help decrease planned and unplanned outages

Virtualization of file system and volume management helps eliminate a key contributor to storage downtime. By separating the logical view of storage from the physical implementation, IT organizations are free to plan data movement or hardware changes, minimizing interruption to business applications. This can result in significant reductions in unplanned outages, prevent application out-of-space conditions, and facilitate a design for high availability.

Here are some examples, using the IBM TotalStorage SAN File System:

- ▶ Systems connecting to a common file system, with access to all data governed by business policy
- ▶ Moving individual files or directory structures to improve availability characteristics, or correcting a performance hotspot without interrupting applications
- ▶ File management functions, like backup, virus scanning or storage resource management, from surrogate systems instead of those dedicated to business applications. This reduces or prevents the impact that these functions might have on application availability or performance.

System management and virtualization

System management cost (people cost) represents an increasingly large part of the IT budget. The tasks involved in system management are clearly affected by the complexity and the size of the infrastructure, and it is therefore essential to apply technologies that address this expense.

When undertaking a consolidation project, it is important not to introduce unnecessary new tasks or requirements that may result in delays and risks for failure. Significant “changes to” or “redevelopment of” system management processes may represent such unwanted tasks. So the question is, does IT Consolidation or IT Simplification in itself require system management processes to be changed? And the answer is—not very much, assuming the consolidation technologies to be used are selected with common sense.

Consolidation based on the merger of different workloads running on different servers may not be the most efficient way of consolidating physical servers, from a system management standpoint. It does, without a doubt, provide more efficient resource utilization, but it also requires the implementation of new workload management facilities, isolation and security measures, and it will affect the existing backup and recovery procedures in a significant way.

Using virtualization as the main technology for consolidation, by contrast, provides a much simpler and more effective consolidation capability, because it generally only affects the physical appearance of the infrastructure, while the application setup, application isolation and application communication (from a logical point of view) is unchanged. This means the existing system management procedures only need to be changed slightly to manage the new infrastructure.

Virtualization will also result in a significant simplification of the physical infrastructure, because the number of physical units are reduced, and because it introduces or enforces standardization, cloning and similar actions in a natural way. Consequently, it reduces staff workload and increases staff effectiveness, with major opportunities for cost savings above and beyond the savings resulting from a more effective use of IT resources.

Improve workload management with virtualization

Basing the infrastructure on effective virtualization technologies may lead to major savings in reduced people cost through more efficient resource utilization and simplification. It is important, however, to combine the usage of virtualization and the related abilities to share capacity among isolated applications with efficient and automated workload management functions and provisioning capabilities. This ensures the capacity is being provisioned according to business priorities and defined policies.

[See earlier feature note on Autonomic computing] Many autonomic facilities are available today in IBM servers, storage, networking, and software product lines; they will also be further enhanced over the coming years. The value is, of course, that the people cost can be further reduced and the business processes are being supported in an optimal way.

Use of system management procedures

High availability starts with reliable products; today's products are reliable and are capable of delivering high levels of availability. However, reliable products alone will not assure high quality of service. A lack of, or failure to follow careful system management procedures and policies is the most common cause of outages.



The on demand Infrastructure Services offerings

In this appendix we describe the comprehensive services you can use to become an on demand business.

on demand Infrastructure Services offerings help clients prioritize “what to do” to Implement Infrastructure Management Solutions



Enabling customers’ evolution to on demand Infrastructure

Get Ready

- **Assessment tools** that create customized customer roadmaps
- On demand **Innovation Workshops** that identify initiatives to leverage on demand capabilities
- On demand **Baseline Discovery** provides rapid evaluation of target IT environment
- On demand **Strategy & Roadmap** identifying value-driven & cost accountable business & IT investments

Get Started

- **Readiness Engagements** identify risk and readiness of clients IT to support on demand technology initiatives including virtualization and provisioning
- **Utility Management Assessment** analyzes gaps, opportunities, and investments required to adopt Utility computing
- On demand **Infrastructure Design** provides an architecture framework for the Operating Environment
- **IT Management Design** for the on demand Enterprise creates a system management architecture

Roadmap

- Customer-specific plan to drive selection and implementation of appropriate offerings to drive business value
- Recommendations for on demand Technology Investments
- Fit/gap analysis, based on best practice architecture reference models, that leverages new technology and creates service-oriented designs

Services overview

Our on demand core services virtualization portfolio encompasses all areas of an organization's technology needs, from facilitated workshops to a baseline client IT environment to introduce virtualization technology, initial assessment of the company's requirements, through design and management to implementation of a virtualization solution.

On Demand Innovation Workshops (1- to 2-day workshops)

This is a facilitated set of highly effective workshops targeted to IT and business personnel to help them understand how they can leverage IBM's new on demand technologies such as virtualization, grid and autonomic computing capabilities. Included are a series of guiding principles and a roadmap of initiatives that are related to virtualization, grid and automation services to help customers quickly understand how virtualization, provisioning, grid and autonomic computing will help them achieve a more efficient and manageable environment.

Grid Innovation Workshop (1- to 3-day workshop)

This intensive workshop introduces grid computing and identifies the infrastructure, tools, and applications needed to create a grid environment. The customer receives a plan to address how and where a grid solution can improve their business process, identify which applications are good candidates, and assess the impact of leveraging grid tools, software, grid servers and so on in their infrastructure.

Introduction for Assessments

Many companies recognize the value of a virtualization technology for an on demand infrastructure, but are unsure of how to apply the vision to their own company's operations. IT managers' questions regarding costs and benefits often conflict with business managers' concerns about improved customer service benefits and capabilities. And everyone is interested in aligning technology with the core business values of the organization.

Our Assessment Services help organizations determine their unique requirements for an on demand environment. Our IT consultants review the organization's business objectives, identify risk, and assess the organization's readiness to implement an on demand virtualization strategy and technologies such as virtualization and provisioning to meet those requirements.

Grid Readiness Assessment

This offering is an issues-based analysis of the readiness of applications, technical infrastructure and IT management services, as well as of opportunities, gaps and investments required to achieve the value of grid technology by applying it to improve common business processes.

This assessment results in financially grounded recommendations including a detailed plan to act on short- and long-term initiatives that will improve focus, responsiveness, variability and resiliency, leading to the deployment of grid technologies and IT management practices.

On Demand Readiness Assessment

This is an assessment of on demand opportunities for IT infrastructure and operational services improvements in support of strategic business initiatives and critical application building block needs. This offering evaluates trade-offs and implementation risks with financial analysis to support on demand technology initiatives for grid, autonomic, utility management infrastructure, autoprovisioning, and virtualization.

Key deliverables from this assessment are the development of a strategic IT roadmap and initiatives for quick hit projects, new technologies, key IBM on demand solutions with a self-funding business model. This assessment reviews select components of applications,

network, infrastructure and IT management procedures in support of the new on demand technology.

Introduction for Strategy

Creating an on demand virtualization solution to meet an organization's business requirements requires a mix of business acumen, technology skills, and the vision that comes with unmatched experience with both. Our ITS on demand Strategy & Planning workshop helps customers identify value-driven and cost accountable business and IT investments to support IT on demand technologies like virtualization and provisioning.

On Demand Strategy & Planning workshop

This workshop assesses on demand opportunities for key strategic business initiatives and IT infrastructure and services, to define and validate the vision for becoming an on demand enterprise from both from a business perspective and an IT perspective. On demand technical attributes are mapped against the customer's baseline IT environment to understand the initiatives needed in the infrastructure and organization to support virtualization, provisioning, grid, autonomies or utility management on demand initiatives.

The customer receives a strategic roadmap and phased implementation plan, along with knowledge transfer of capabilities and reference models in order to understand the impact of the new on demand technology for their environment across their network, application, storage, server implications.

This workshop is tailored specifically to address individual on demand technologies like virtualization, grid, autonomic computing, but can also encompass a range of technologies based on customer requirements. The On Demand Strategy & Planning workshop evaluates the implementation risks, organizational change readiness and business transformation impact for the IT infrastructure.

Introduction for Design and Integration Services

Our Design Services guide companies through the design of an on demand infrastructure, from determining the policies and tools required to link business users with the infrastructure, to the specific software and hardware components required to get the job done.

After a company has formulated its on demand virtualization strategy, the question usually becomes "how do we get this done?" Many organizations choose to implement the plan themselves to enable in-house staff to better learn the new infrastructure. Others prefer to keep their staff focused on daily business, and come in to learn the new infrastructure when the work is complete.

Our ITS Integration Services allow customers to choose their own level of integration support during the implementation process. We work with the company's staff to minimize downtime during the implementation, and ensure that the company's IT team is fully trained to support the new on demand technology and educated in how to support the new infrastructure.

On Demand Infrastructure Design

This service provides an architectural framework for the operating environment. The On Demand Infrastructure Design solution model offers a shared services operating environment that delivers variable pricing, dynamic provisioning and allocation of capacity to support computing, grid, virtualization, Autonomic computing and other distributed processing technologies, as well as resource optimization and the highest level of service to IBM clients.

Moreover, this service offering will provide the physical infrastructure design that will enable a client's existing IT assets to effectively support Web Services and Services Oriented Application Architectures. Finally, a completed On Demand Infrastructure Design effort will

also provide clients with a technical roadmap, to guide their development of an on demand Operating Environment or to compare against their existing operating environment to identify functional gaps and inefficiencies.

The On Demand Infrastructure Design Solution model will provide a multi-phased approach for IBM practitioners to frame the client's goals for on demand computing, assess the capabilities of their current IT portfolio to support on demand, and develop an effective design that allows them to leverage their existing technology base in the creation of an on demand Operating Infrastructure.

IT Management Design for on Demand

An on demand environment requires reworked management systems (processes, people, technology) in order to succeed. The IT Management Design for on demand service offering creates an enterprise-integrated IT system management reference architecture and framework to support on demand requirements for service component enablers.

Built on an on demand framework, the accelerated assessment approach leverages on demand technology competencies and proven methods. It helps clients to understand and define the technical and business requirements around emerging technologies and IT management capabilities and the business value of improving automation maturity; the creation of an automation roadmap, service workflows, financial modeling and organization approach (such as automated provisioning, virtualization, business systems monitoring, autonomic computing services, and workflows). The customer receives a transition plan to implement the new IT integrated management reference architecture.

ITITO Autonomic Computing Accelerator

This offering enables the rapid adoption and implementation of ITITO (IBM Tivoli ThinkDynamics Intelligent Orchestrator) via a pre-packaged set of intellectual capital for processes, organization, and technology. ITITO provides auto-provisioning and orchestration, which are key on demand capabilities. This offering is an accelerator based on IBM's best practice GSM Execution Model and assets for implementing ITITO and integrating it into their IT infrastructure.