



# **The High Cost of Unplanned Availability**

**An Examination of iSeries™ Business Continuity Issues**

**Version 1.0**

**March 21, 2005**

**An @server® i5 Whitepaper**

**Joel H. Knudson – IBM Business Unit Executive**

**Americas eServer i5 Enterprise Sales**

[jhknuds@us.ibm.com](mailto:jhknuds@us.ibm.com)





# The High Cost of Unplanned Availability

## An Examination of iSeries Business Continuity Issues

### Table of Contents

Executive Summary .....	1
Business Continuity Defined.....	1
Why IBM eServer i5? .....	2
Infrastructure Optimization Drives Accessibility Requirements .....	2
Evaluating Exposure .....	3
Availability – One Step at a Time .....	3
Availability Planning – IT Resiliency Profile .....	4
Availability Planning – Building the Plan.....	4
Business Continuity Alternatives .....	4
i5/OS – Building Blocks of Availability .....	5
Storage Alternatives .....	5
iSeries High Availability ISV Solutions .....	6
Availability For the Extended Environment .....	6
Windows .....	7
AIX, Linux and Domino .....	7
Cost Justification for Business Continuity .....	8
The Cost of Downtime and Assessing Application Value .....	8
Availability Justification Sample .....	9
Conclusion .....	10
Additional Resources .....	10
Trademarks and Disclaimers.....	11



## **The High Cost of Unplanned Availability**

### **An Examination of iSeries Business Continuity Issues**

#### **Executive Summary**

Little more than a decade ago, few people imagined that the Internet would grow to be the primary force affecting both business and consumers. No one anticipated that the computers in their marketing department would be connected to the billing and inventory servers, or that their clients would be using online, web-based order entry applications. Even more unfathomable was that the successful integration of the different servers in an organization would ultimately spell the difference between success and failure.

Today, as the number of access devices proliferate and become part of the corporate computing culture, enterprises are reengineering both their IT systems and the corresponding business processes to interact with customers, employees and suppliers. Desktop workstations are being supplemented or replaced with laptop computers, PDAs, cell phones, as well as an array of network-enabled devices that can communicate without a human client. This demands a reliable infrastructure that can accommodate future change, rapid growth and can mask the server and network complexities from its users.

Universal access to connective technologies is changing how organizations use servers and data to provide new applications and ever higher levels of service. It was only a few years ago that electronic mail was a simple application of convenience, often offered only to management personnel. Today it is a critical application for all users accessed through both desktop and hand-held devices, at all hours of the day, requiring the highest levels of uninterrupted system availability. And now, instant messaging is altering the IT landscape.

This paper will address the changing definition of **Business Continuity** against the context of the IBM eServer™ i5 server. Additionally, it will examine the typical “high availability” solutions that are deployed, as well as the ability to extend improved levels of reliability to connected servers across the iSeries infrastructure. Lastly, some basic methods for cost justification of business continuity will be explored.

#### **Business Continuity Defined**

Business continuity is a phrase that has been used synonymously with the preparation for disaster recovery. The current view of availability planning involves defining and implementing strategies to maintain business operations, no matter what kind of man-made or natural disruption occurs; according to the defined tolerance of the enterprise.

Business continuity management and contingency planning are essential and unavoidable tasks that should be revisited every year. However, the creation of a sound continuity and contingency plan can be a very challenging undertaking, involving a number of stages and discrete activities. This is especially true in a complex environment involving multiple servers and different server architectures. A 21<sup>st</sup> century plan for managed availability must account, not only for the less likely disaster recovery scenario but now more importantly, must accommodate the unpredictable demands of everyday operations, unexpected business growth, and the genuine and often mandated requirement to eliminate both planned and unplanned server downtime.

## Why IBM eServer i5?

It is surprisingly difficult to define or categorize the IBM eServer i5 server even though many in the IT industry feel compelled to do just that. It is not a UNIX® server yet it is able to run AIX 5L™ which is IBM's acclaimed version of UNIX. It is not a Linux® server yet it can run several versions of Linux. In fact an iSeries server can run multiple, independent instances of Linux and AIX® on the same server where each can use dedicated storage resources, and an operating system failure in any of those virtual servers may not affect the rest of the system. Building on that capability, an internal virtual Ethernet network can connect these logical servers without requiring cables, routers, gateways or other potential points of hardware failure and configuration complexity.

The iSeries server defies categorization. It is a unique 64-bit server, based on IBM's Power5™ processor. It runs IBM's i5/OS™ operating system, currently at version 5, release 3, as well as Linux and AIX. It can also run Windows® via an Integrated xSeries Adapter or an Integrated xSeries Server, while delivering extremely high levels of reliability and expandability. iSeries servers can also function as a storage area network for connected Windows servers offering storage protection and fail-over capability. iSeries servers support hot-swappable processor cards, PCI adapters, and disks and can run Java® applications natively on i5/OS, the WebSphere® application server, as well as Domino® applications. IBM believes no other server offers such flexibility or promise for infrastructure optimization.

## Infrastructure Optimization Drives Accessibility Requirements

After several years of rapid change, many companies have found themselves with complex IT infrastructures, insecure environments and inflexible implementations. As business requirements grow they are finding these systems to be underutilized and complicated to manage. Server consolidation projects have exposed the need for full scale infrastructure optimization. The future is about much more than improving costs and reducing operational complexity. It is about increasing business resilience while prioritizing IT and business decisions to support future innovation and growth across the entire enterprise.

Today iSeries servers offer the prospect of virtualizing IT resources to make it more simple to manage diverse business applications while shielding the complexity inherent in a multiple server, multiple OS environment. The IBM Virtualization Engine™ is not simply partitioning. Virtualization provides a logical (rather than physical) view of data, computing power, storage capacity and other resources while automating the management of those resources based on defined business objectives and service levels for all users.

Dynamic logical partitioning (DLPAR) is the system architecture that iSeries companies are using to virtualize their corporate IT infrastructures. It is a function of the Virtualization Engine that enables the sharing of hardware resources across independent operating environments. LPAR allows the division of a single iSeries server into several completely independent virtual servers or partitions. It is offered on eServer i5 servers and is also supported by AIX 5L and the Linux operating system available from multiple Linux distributors. Furthermore iSeries integrated xSeries® solutions allow businesses to operate heterogeneous environments that include Intel-based servers running Microsoft Windows or Linux applications in addition to the core business applications.

This degree of flexibility makes the iSeries server the ideal selection for infrastructure optimization where fewer physical servers can act as many logical servers. However, when evaluating the exposure placed on a company with a more centralized IT environment, focus must also be placed on the expanded impact of possible IT outages.

## Evaluating Exposure

It is essential to understand availability requirements from the viewpoint of the users. Emerging integrated applications that offer more data to internal decision makers also make those same users more reliant on the entire IT infrastructure. System users may now be company employees or customers performing self-service operations. Users might also be business partners that use or resell products and services. They might include related IT systems, in an intranet or extranet, that initiate transactions automatically. The perspective of all users and clients must be considered when evaluating the company reliance and exposure to a given IT environment.

If the IT system cannot serve these functions for any reason, including both planned and unplanned outages, the IT system is viewed as being down. Consequently, assessing the risk associated with IT systems is ultimately dependent upon business issues, not technical ones. It is the functional layers (network, server, application, operating system, storage) that might be lost and the impact to the business that must be measured. While a centralized IT environment has more business dependencies and risk, it also enables the implementation of the very redundancies that will ensure a more resilient environment.

Unhappily, the risk assessment of the potential loss of dependent functions and an examination of the solutions necessary to protect against their failure, is often viewed as unworkable in a dynamic business environment with a complex infrastructure. As today's business initiatives call for IT systems that are online and accessible on a continuous basis, a consolidated server that supports multiple operating system environments actually becomes a necessity. Such a server makes comprehensive availability planning and server redundancy truly viable, while better supporting change in the future.

## Availability – One Step at a Time

Clearly, if an enterprise does not have any real disaster recovery strategy, it is unlikely that it will be possible to consider business continuity. Disaster recovery involves tools and strategies designed to minimize IT infrastructure downtime in the event of a lengthy unplanned outage. The most common strategies for preserving data are with scheduled backups, remote storage of saved data and the documentation of recovery procedures.

Reducing the time that a business function might ever be offline is the goal of business continuity solutions today. Most business continuity plans still focus on data in primary and secondary servers. But, for some, the larger corporate IT and communications infrastructure must be considered as well, including voice and wireless traffic, email, and physical assets such as work spaces and network infrastructures. To summarize, disaster recovery strategies typically involve:

- Operational procedures for ensuring recovery of computing and networking hardware as well as company data
- Identification of recovery time and recovery point objectives for business applications
- Direction from the CIO and IT management levels

The broader scope of business continuity should involve:

- A comprehensive plan covering people and facilities as well as systems and data, in order to ensure that critical business processes can quickly return to operation
- Provisions for continuing IT services through both planned and unplanned component, server, or site outages
- Direction from CEO or business unit management levels

## **Availability Planning – IT Resiliency Profile**

There are many detailed documents and manuals to choose from that can guide one through the journey called the availability planning process. One recommendation for understanding the nature of business continuity within an IT infrastructure involves the concept of an IT resiliency profile. It establishes a metric where one can assess the potential IT recoverability performance across a set of common, typical variables that might impact the capabilities of the IT environment to withstand or recover from an event that could otherwise cause an outage.

There are six major components involved in an IT resiliency profile.

- Availability Coverage – a complete list of critical applications
- Interdependencies – the hardware and software supporting the key applications
- Single Points of Failure – risk areas and documented action plans for failure handling
- Ownership and Process – documented lines of responsibility
- Skills and Staffing – required staffing to support service level agreements (SLAs)
- Achievability – a realistic assessment of SLAs in the current IT infrastructure

## **Availability Planning – Building the Plan**

In order to build an executable plan that leads to a resilient IT infrastructure one must match the present and future needs of the business with the company's current and planned IT capabilities. There are at least five steps involved that will be summarized here and they may be derived from the IT resiliency profile. They include:

- Business Impact Analysis – a process that includes the definition of outage costs as well as the actual availability requirements for each prioritized application
- IT Resiliency Objectives – documenting the specific recovery objectives for the IT infrastructure and critical business applications
- Baseline – an assessment of the current business continuity capabilities
- Gap Analysis – a determination and prioritization of needs versus planned SLAs
- IT Resiliency Requirements – very specific requirements for each application

## **Business Continuity Alternatives**

It should be evident that a true business continuity implementation evolves as a business grows. It frequently involves a large number of servers, running a dissimilar set of applications, across a network that may involve multiple technologies and several potential points of failure. These servers are usually accompanied by a variety of users and the corresponding user requirements. In many cases, allocating a daily save window for backups is no longer acceptable. Complicating matters are the new application needs that arise continuously as business units demand additional accessibility to more data.

The potential risk of failure when undertaking a business continuity project that seeks to address such a moving target can be an overwhelming concern. The desperate need for higher availability has led to the creation of many third party IT "solutions". The most attractive answers for business continuity are those that promise a simple resolution to all the potential problems that could arise from the fragile nature of a complex IT environment. However, in truth, there may be no simple answer as each enterprise still has to understand the specific availability needs as driven by business requirements, before any server, storage or software-based alternatives can be considered.

## **i5/OS – Building Blocks of Availability**

The IBM eServer iSeries family of products is designed to address the computing needs of companies utilizing commercial application programs. The eServer i5 server is extremely competitive in the marketplace because of its integrated design, broad application choices and server consolidation capabilities. These attributes help make systems management easier, security and reliability better, and total cost more affordable when compared to a typical server farm. With competitive servers, business innovation and growth can be sacrificed when companies spend a disproportionate amount of money and resources maintaining and managing their IT environments. eServer i5 servers are designed to simplify IT management by enabling more effective infrastructure optimization and reducing costs by delivering more flexibility and performance for a similar investment.

Many iSeries customers believe that i5/OS is the core value of the eServer i5 server. In truth, the iSeries hardware, i5/OS operating system, and middleware are built and tested together as a unit as part of the development process. This process helps ensure that the critical components are integrated together and work smoothly before the system is ever shipped. The robust i5/OS operating system includes a relational database, self-managing storage, transaction processing, security, automatic workload management, partitioning, support for multiple operating systems, disk mirroring or RAID storage protection, journaling, systems management functions, Web application serving, communication and networking capability, on/off capacity on demand, switched disk, and much more.

Each iSeries customer has some definable availability requirements even if not every customer needs 7 x 24 availability. However, every iSeries customer can exploit i5/OS to address many of their unique business continuity needs. A good understanding of the functions and capabilities provided within i5/OS is essential before considering other expenditures for business continuity.

## **Storage Alternatives**

In addition to integrated disk storage, iSeries servers extend the choices for storage connectivity with flexible external disk attachment offerings that allow both IBM and non-IBM storage systems to connect. In certain implementations the storage systems can provide a measure of data protection that may fit into some disaster recovery schemes. Continuous availability requirements that demand transaction integrity may require additional function that cannot be provided by storage technology alone.

IBM's TotalStorage® Enterprise Storage Server® (ESS) Model 800 delivers new standards for performance, automation and integration. The IBM TotalStorage DS8000 series supports extensive connectivity including Fibre Channel, FICON®, and ESCON® across a broad range of server environments. These storage systems support many advanced functions which can be used to help increase data availability as well as for protecting data from some planned and unplanned outages.

IBM's external disk attachment also benefit from iSeries advanced storage architecture whereby functions such as i5/OS mirroring, Expert Cache, and Independent Disk Pools can be exploited, complementing its own storage management and advanced functions such as FlashCopy® and Peer to Peer Remote Copy (PPRC). IBM's ESS is designed to provide outstanding performance with dual cluster RISC SMP processors, large cache and serial disk attachment. ESS uses redundant hardware, mirrored write caches and RAID5 or RAID10 protection for disks to support high availability (HA) for mission critical business applications.

## **iSeries High Availability ISV Solutions**

There is no single definition that quantifies a highly available iSeries environment. However, a very basic model might include dual servers where the second (backup) can take over operations for the first (production) in the case of an outage event. In such a situation, some function would need to be replicating any changes on the first server to second, in a timely manner, such that they are effectively synchronized at all times. Ideally, when the backup server takes on production responsibilities, an automated process would redirect the dependent network and users to the new production server. In a perfect world, any active jobs on the production system would restart automatically.

It is possible to achieve this level of data and application resilience, but it requires an intelligent application design tailored for recoverability. Typical eServer i5 customers initially implement a solution that provides basic data resiliency between the two servers before they elect to add automation fail-over logic to their programs. i5/OS provides a rich foundation of function that can enable data replication to occur. It supports local and remote journaling, commitment control as well as IP failover for wide area networks. i5/OS cluster architecture provides mechanisms to coordinate application resilience with the resilience of the data associated with that application. The entire cluster and all of its resources can be managed from a single control point.

Nonetheless, for those eServer i5 users considering the implementation of a high availability solution, most explore the numerous i5 HA solution providers and solution types that provide disaster recovery or continuous availability for iSeries operations. The advantage of a third party solution is that it should eliminate the need to author custom programming that utilizes the functions provided by i5/OS. Even then it is critical to establish a clear set of business objectives and to be familiar with the HA Independent Software Vendor (ISV). A particular solution's actual capabilities and behavior depends on the specific implementation, which may have its own unique set of characteristics and limitations. Due diligence and fundamental research can help an enterprise determine the strengths of a given solution. Finally, it is important to find similar references that have experienced business recovery actions with the solution being considered.

## **Availability For the Extended Environment**

The capability iSeries servers to act as a consolidation platform for other operating system environments permits enterprises to make application decisions based upon the needs of the business without regard for the type of server that the application may call for. iSeries customers are able to leverage their investment in eServer Power5 technology to deploy Windows applications while running AIX 5L, Linux, and i5/OS environments in a simpler IT infrastructure that does not require multiple servers to run diverse applications.

Beyond the reasonable expectation that a smaller IT infrastructure should cost less to manage, it is also valid to anticipate that a lesser number of physical servers should make it easier to design a business continuity solution. Again iSeries servers provide significant advantages in this regard. By offering functions that are built into i5/OS, enterprises are able to expand the sphere of control that their availability solution affects. The ability to dynamically switch disk storage between eServer i5 servers and the flexibility of the internal virtual Ethernet networking, as well as the ability to exploit business continuity technologies that may already be in place for existing Windows, UNIX and Linux servers can solve the complexities of the business continuity issue for many customers.

## Windows

Integrated xSeries Adapters (IXAs) are PCI-based cards that install inside selected xSeries servers, providing a 1 GBps high-speed link to iSeries servers. This connection provides centralized storage and systems management, allowing the attachment of multiple stand-alone, n-way IBM xSeries servers to a single iSeries server, leveraging xSeries processing power while tapping into iSeries systems management and storage architecture.

The Integrated xSeries Server (IXS) is an Intel-based server on a PCI card that plugs directly into an iSeries server, without requiring the attachment of an external xSeries server. It has its own processor (uniprocessor), memory and LAN adapter, but uses iSeries disk, tape, DVD and systems-management resources while operating as a stand-alone Windows server.

iSeries operating system support provides for virtual Ethernet connections to the IXS and IXA servers which enables IBM to provide support for Microsoft clustering services while also connecting to DB2® and Linux virtual servers on the same iSeries. Tailored high availability solutions can be extended to the Windows environments by using either native iSeries or Windows-based network or backup utilities. A single Windows server may act as a “hot” backup for other Windows servers connected through an eServer i5 server.

## AIX, Linux and Domino

eServer i5 enterprises are able to evaluate their existing IT network and ask themselves how they might perform an infrastructure optimization project without eliminating operating environments in order to simplify. Typically, such an action creates a resulting need to convert affected applications. Each eServer i5 server represents an opportunity to simplify, and yet run i5/OS, AIX 5L, or Linux in its own native working environment.

The real beauty of iSeries servers is that such a consolidated implementation imposes no special requirements onto the business continuity design that may utilize multiple virtual, diverse servers. The resulting implementation continues to utilize the same familiar server environments while offering the redundancy advantages of a single physical server. iSeries applications can have high availability functions delivered through i5/OS and complemented with HA ISV solutions. AIX 5L applications can exploit IBM's own High Availability Cluster Multiprocessing (HACMP®) technology to provide high availability for the AIX servers by mirroring data between partitions or separate AIX sites.

Similar advantages are extended to Linux servers operating within a partition on an iSeries server. They offer new alternatives for application environments that may be operating elsewhere. As an example, Samba® a common Linux application, allows for the consolidated management of Windows file and print servers. Third party applications offering clustering and data replication for UNIX, Windows and Linux may also be used to provide availability solutions as selected, based upon business need rather than the hardware limitations.

It is even possible to run clustered, fully redundant Domino partitions, thousands of Lotus Notes clients and hundreds of Domino databases in a single eServer i5 server. Those critical user applications can now run alongside all the other business applications needed for an enterprise, while utilizing autonomic recovery capabilities that exploit unique iSeries and i5/OS features such as switchable disk or cross site mirroring of disk storage.

## Cost Justification for Business Continuity

Once the need for business continuity has been acknowledged, the next step is typically some level of initial cost justification prior to proceeding. When business continuity meant disaster recovery, this was an especially challenging exercise. The enterprise hoped that they would not actually have to use the solution, which hindered an investment return exercise, because usage of the investment would require a calamity of some nature.

## The Cost of Downtime and Assessing Application Value

As business continuity has come to mean increased server availability, company operations are widely disrupted when there is an outage event in IT. Whereas, in the past an outage simply meant overtime for the IT department, it has now become clear to the affected business units what the impacts are when the IT environment is unavailable. As a result the cost justification exercise can be more simple to accomplish. There are often very tangible costs incurred when an outage event is experienced (Box 1). Additionally, to further support the case for business continuity it makes sense to revisit the actual value that the commercial applications are delivering to the enterprise (Box 2).

### Cost Justification Box 1: Cost of Downtime

- Lost Business
  - Which varies between companies
  - And can be permanently lost business
- Negative Image and Publicity
- Fines and Penalties
  - These may include industry assessments or internal service level penalties
- Staff Costs
  - Companies must pay for inactivity, overtime, and necessary recovery actions
- Impact on Business Decisions
  - Many daily business decisions rely on access to real time, current information

### Cost Justification Box 2: Assessing Application Value

- **Applications** - Review earlier business cases for the affected applications. Most applications underwent a business justification prior to being developed or acquired. These earlier value judgments can be used to create cost/benefit statements for business continuity.
- **Existing Disaster Recovery Agreements** – Most businesses have some form of disaster recovery plan. In many situations there will be existing estimates for business impacts due to server outages that can be applied or can be extrapolated for specific applications.
- **Transaction Values** – Analyze transactions per day, per application. Determine if an average value per transaction can be reached. Additionally, consider utilizing existing department chargeback agreements for value assessments.
- **Industry Surveys** – Many industries compare IT budgets (i.e. % of revenue) within their peer group. Existing reports will normally define a cost range per hour for outages.

## Availability Justification Sample

The tables presented below were modeled after a real iSeries customer scenario (catalog mail order) that wished to eliminate both planned and unplanned downtime. Though very basic, the methodology that was used accomplishes the objective of demonstrating that business continuity can be cost justified. In the Table 1, basic revenue and usage factors were collected. In Table 2, the numbers in the first table were used to calculate the annual price tag for server downtime, also known as The Cost of Unplanned Availability.

<b>Table 1: Revenue and Usage Data</b>	<b>Factor</b>
Company Annual Revenue	\$500M
Annual Revenue Per Hour (8760 hrs/year)	\$57,000
Company Annual Profit	\$35M
Annual Profit Per Hour	\$3995
Planned Annual Downtime	32 Hours
Prior Year Unplanned Downtime	76 Hours
Prior Year Unplanned Downtime – Additional Costs (IT Overtime, Penalties, External Support, Emergency Plan Activation)	<b>\$100,000</b>
Average Concurrent Users Supported	400
Average Annual Compensation Per User (\$25/hour)	\$50,000
Estimated Employee Productivity (during IT downtime)	50%
Calculated Profit Loss Per Hour (during IT Downtime = \$35M/8760 hrs)	\$3,995
Hours of Operation Per Day	24
Days of Operation Per Year	365
Hours of Work Per Employee Per Week	40
Weeks of Work Per Employee Per Year	50

<b>Table 2: Downtime Cost Calculation</b>	<b>Factor</b>
Planned Downtime Cost - Based on Profit (\$3,995 x 32 hrs)	\$127,840
Prior Year Unplanned Downtime Cost - Based on Profit (\$3,995 x 76 hrs)	<b>\$303,620</b>
Total Outage Costs Based on Revenue (\$57,000 x 108 hrs)	\$6.16M
Salary Costs During Unplanned Downtime (400 x \$12.50 x 76 hrs)	<b>\$380,000</b>
Agreed Upon Availability Justification Factors (\$100K + \$303K + \$380K)	<b>\$783,620</b>
Additional Justification (Business Initiatives Budget) <ul style="list-style-type: none"> <li>• Business Intelligence Project (\$300,000)</li> <li>• Developer Server Partition (\$150,000)</li> </ul>	<b>\$450,000</b>
<b>TOTAL JUSTIFICATION FOR BUSINESS CONTINUITY</b>	<b>\$1,233,620</b>

**Table 1** and **Table 2** provide an example of how a specific customer determined the value of business continuity. Actual results may vary. In addition, there is no guarantee that a high availability solution will provide business continuity or eliminate planned or unplanned downtime.

Unplanned downtime profit numbers were used to calculate the business continuity justification. Because business intelligence and a dedicated developer environment was thought to be addressable on the second (HA) server, project budgets were added to the total. Had revenue, instead of profit, been used as the basis for calculation the continuity cost justification would have been even more substantial.

## Conclusion

In 2003, two major issues affected the entire world. The first of these was the spread of SARS disease, which changed many business continuity professional's thinking on the importance of protecting human resources, not just technology assets. The second was the various wide-area power outages which severely impacted businesses situated near the east coast of the United States, much of Italy, southeast England, Sweden and Denmark. For many businesses these outage events highlighted the need for business continuity in a way that few other external occurrences ever have.

2004 witnessed more than its fair share of natural disasters. Severe storms and floods have always been widespread catastrophes, but none more so than the flow of hurricanes which lined up to track through the Caribbean to the southeast United States. Companies with solid disaster recovery plans showed their worth and, in general, business continuity planning was displayed in a very positive light. Then the world saw the historic tsunami in Asia, and became aware that such a happening could hit almost any coastline at any time.

As a result many IT professionals are asking new questions about the future while wondering how complete their availability solutions are, and what other considerations need to be evaluated. Natural and human-caused disasters do occur and must be planned for, especially when most industry observers agree that the vast majority of outage events are caused by human failures. Likewise, it is not speculation to plan on more technology change and the continued momentum towards continuous server operations as well as the never-ending addition of new, critical user applications to benefit the enterprise.

It is worth noting that in the United States, the Sarbanes-Oxley Act of 2002 has imposed many mandates on audit, contingency and business continuity policies that are imposing much change to IT infrastructures in both large and small companies. Similarly, in 2004, the Civil Contingencies Act was passed by Parliament in the UK and represents the first time that business continuity has been made a legislative requirement in that country.

The inexorable march towards high and continuous availability will continue. Enterprises will further consolidate their IT operations, exposing the need to run multiple operating systems. IBM has responded with the IBM Virtualization Engine. Traditional iSeries HA ISVs are also extending their high availability offerings to cover integrated solutions involving i5/OS, Windows, Linux and AIX 5L. Existing disaster recovery schemes will likely be refined, with an increasing focus in the small and medium sized markets that are recognizing the attractive cost of today's redundant hardware solutions. Every company will be making availability decisions in an environment where IT budgets cannot be allowed to grow at the same pace as support costs. IBM provides the answer.

## Additional Resources

- ✓ **iSeries High Availability Website:**

<http://www.ibm.com/eserver/series/ha>

- ✓ **IBM Redbooks**

<http://www.redbooks.ibm.com>

**Key Title:** i5/OS High Availability Clusters: Data Resilience Solutions.

## Trademarks and Disclaimers

© IBM Corporation 1994-2005. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country. The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX	e business (logo)	ESCON	IBM	Notes	xSeries
AIX 5L	e (logo) business	FICON	IBM (logo)	POWER	
DB2	e (logo) server	FlashCopy	i5/OS	POWER5	
DB2 Universal	eServer	HACMP	iSeries	TotalStorage	
Domino	Enterprise Storage Server	IBM Virtualization Engine	Lotus	WebSphere	

Lotus, Freelance Graphics, and Word Pro are registered trademarks of Lotus Development Corporation and/or IBM Corporation. Domino is a trademark of Lotus Development Corporation and/or IBM Corporation.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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

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

Information is provided "AS IS" without warranty of any kind.

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

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information in this paper concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.