

Reducing unplanned downtime and maintaining peak performance with change management.



The Necessity of Change Management in the Intel Server Space

December 2002

by IBM ^ xSeries Group

What is Change Management?

Change management, which is the planned updating of **system software**¹, is not a new concept in the computer world. Change management has been part of the everyday language in the UNIX™, minicomputer and mainframe environments, where successful businesses depend on the availability of those servers and cannot afford costly, unexpected outages. This has not been the case in the Intel processor-based environment until now.

Today, systems, subsystems and options use system software to control these components. Upgradeable system software provides the ability of upgrading the system software to keep the system or option operating at peak availability. The latest system software helps reduce downtime and operating expense, but more importantly, keeps your business critical applications available and your customers satisfied.

This paper illustrates how IBM change management tools help maintain your systems at the latest system software level. Their automated and automatic operation simplifies the maintenance and installation of system software updates providing your IT staff with valuable time.

Intel-based systems have matured from desktop units to servers designed to take on business critical functions. As customers move to take advantage of lower cost Intel systems on which to base their IT structure, the need for formal change management is becoming very clear. These servers are not running just "file and print" applications anymore. Rather, they are used for business-to-business (B2B) and business-to-consumer (B2C) revenue-generating applications such as e-mail, Web hosting, electronic commerce and Web caching, to mention a few. Intel customers are facing new challenges—challenges mainframe customers have been addressing for years, such as:

- How fast customers will go to another vendor's Web site if one is unavailable.
- How quickly workers become idle when their e-mail is not operational.
- How businesses can no longer afford downtime to upgrade system components on Intel-based systems.

Delivering reliable service and, at the same time, meeting your customer's expectations, requires that your server be operational 24 hours a day, 7 days a week, 365 days a year. Today more than ever, if you're not available, you're losing money. In fact, the Standish Group estimates that the average cost of downtime to be \$10,000 per minute! Some applications can even be up to \$27,000 per minute². Even if you are a small company, the risks are high.

The penalties for downtime are accelerating in severity

- ✓ "Average cost of mission critical application downtime \$10K/min"
- ✓ "Most users experience 66% more downtime than anticipated"

Figure 1. Penalties for Downtime

Clearly, downtime is just not a viable business option. It is critical for your business to maintain a high level of round-the-clock service quality.

¹For the purpose of this paper **system software** will be used to refer to firmware (diagnostics, system BIOS, and BIOS for hard disk drives, tape drives, ServeRAID™ adapters, service processors and other devices) and device drivers (for network adapters, service processors, video adapters, ServeRAID adapters, SCSI adapters and any other options).

²Source: Standish Group, High Availability Forecast, "Pound Foolish".

Reducing unplanned downtime and maintaining peak performance with change management.

In addition to the high cost of downtime, the overall cost of owning all this technology is increasing. While the hardware costs continue to decrease, that is still only 15% of the overall cost of ownership. In fact, according to the US Bureau of Labor Statistics, IT maintenance and repair costs alone are increasing at the rate of 8.3% each year.

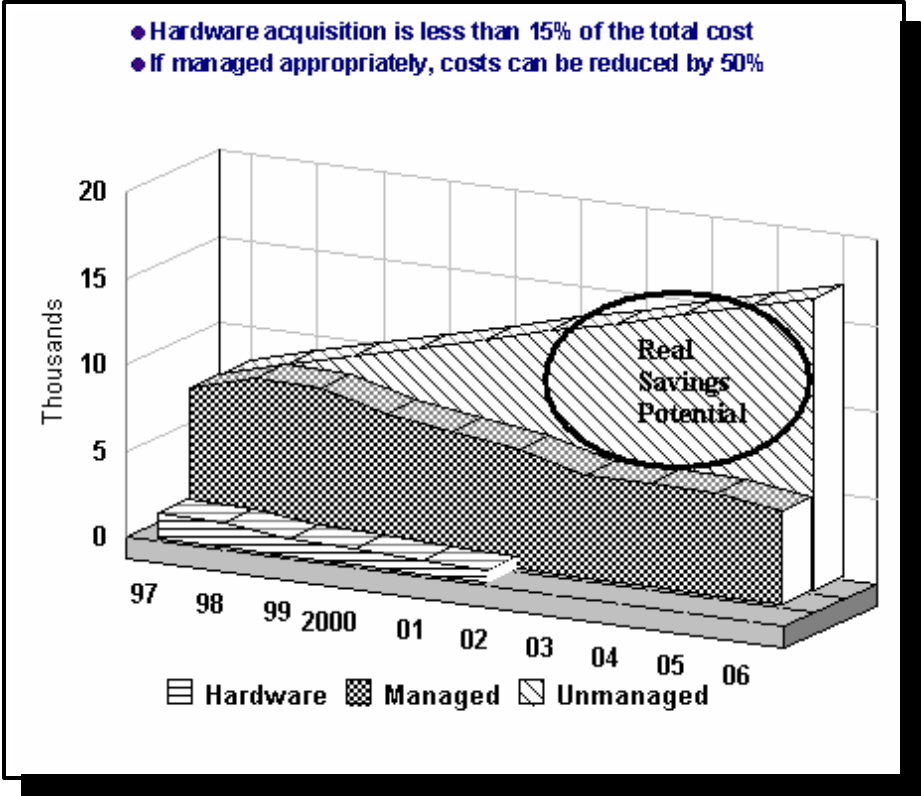


Figure 2. Total Cost of Ownership (Source: Gartner Group)

Figure 2 highlights that costs in a “managed” environment³ continue to project a downward trend while costs in an “unmanaged” environment⁴ continue to rise.

In the past, change management on Intel systems was usually performed once the system was in a non-operational state, without any preventative measures taking place. For customers looking to grow their business and secure a competitive edge, that strategy may no longer be ideal. Intel-based servers are the new mainframes for many businesses and, as such, need to follow the same formal preventative processes that usually have been associated only with the mainframe environment. Now, Intel systems designed as servers must undergo the same rigorous process as any other business-critical system. The main action to take is to identify and implement a change management process, that lets you effectively and proactively update the system software in your IT infrastructure.

³Managed environment: having a proactive strategy of taking actions such as planned change management to prevent costly outages.

⁴Unmanaged environment: a wait until something fails then fix it or ad hoc fix strategy.

Reasons for Change

Having a system at the latest level of system software helps not only to increase the system's performance but also to keep it operating at peak availability.

IBM strives to provide its customers with the most reliable server platforms possible. The Intel-based server line is an extremely open platform with its Industry Standard Architecture (ISA), where hardware components, operating systems, and software applications come from a multitude of different vendors and the market is constantly churning out new hardware and software. By the time the customer has fully loaded a system, there may be code from dozens of sources. While every vendor tests their code to the best of their ability, it is impossible to test every possible configuration of hardware and software that can exist. It is therefore inevitable that sometimes vendors will learn about an incompatibility after the product or the system software update has already been released. Most of these lessons are learned early in a product's system software life cycle, so fewer and fewer service packs, hot fixes, and firmware updates are required as the product and system software mature.

Updates are used to provide performance and usability enhancements, new features, and bug fixes. Some of these fixes correct problems in the vendor's own code, but many also address compatibility issues between one vendor's elements (application, operating system, hardware) and the elements from other vendors. Updates provide a way for the Manufacturer to improve compatibility, features, performance, availability, etc. for their customers.

Figure 3 provides a very simple depiction of how various software and hardware elements (application products and system software) are layered in your server.

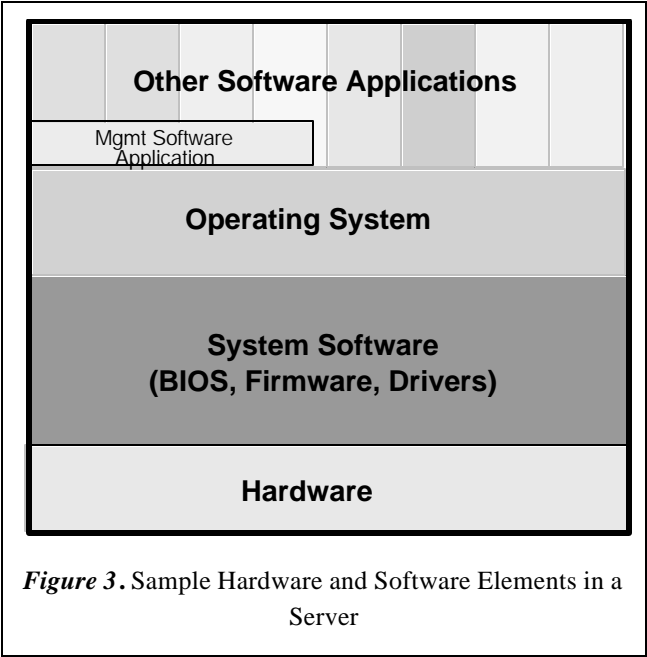


Figure 3. Sample Hardware and Software Elements in a Server

These different layers are continuously interacting with each other—and not only with the layers directly adjacent to them. It would not be unusual for a problem to surface in one layer, and the root cause for the problem to be traced to another layer, given the continuing interaction between all the components. Therefore, throughout each element's life cycle, hardware and software vendors share information about their products. Many customers are surprised to learn that vendors often share source code with each other, and one vendor may actually patch another vendor's code and return it to the owner for integration into the owner's next update because of testing with an upcoming element. It is, therefore, important to maintain all components with equal vigor. The wider the gap between the most and least maintained components, the more likely it is that problems will occur.

In some cases, it is extremely important to synchronize updates. Most hardware components come with a driver. If the component BIOS is flashable, the manufacturer may also release firmware updates. Manufacturers also provide management software for some hardware components. When a manufacturer releases a new driver with new system software and/or management software, the manufacturer expects the customer to update all of the elements together. Unpredictable results may occur if these elements get out of sync. In most cases, a

Reducing unplanned downtime and maintaining peak performance with change management.

condition of support will be that the hardware system software and/or management software are at the proper release level.

Adopt a Strategy

Due to the frequency of updates that a customer receives from hardware and software vendors, it is virtually impossible to apply every update that vendor's release, and no one would recommend that this be done. By the same token, it is very dangerous to live by the "if it ain't broke, don't fix it" mantra. Some hardware and software fixes are designed to prevent data corruption, data loss or to plug security holes. Other fixes address memory leaks or system halts that relate directly to server downtime. These fixes should clearly be applied in a timely manner when servers are "at risk."

Most bug fixes can be tracked back to a problem that some customer has reported. We can assume that the fix is important to that customer, but the fix may or may not be important to other customers. Most vendors provide information to help customers determine if they have systems that are at risk and to quantify the seriousness of the risk.

Updates that address less critical issues may also be important down the road. A large percentage of the calls placed to technical support are resolved when the customer applies an update that was previously posted to the vendor's Web site. Had these updates been applied as part of a regular maintenance cycle, customers would have enjoyed more reliable systems and lower support costs. According to a Gartner report on the causes of system downtime, *Figure 4* shows that at least 40% of all system failures are due to problems with the operating system or application code, while only 20% of the failures are due to problems with the hardware.

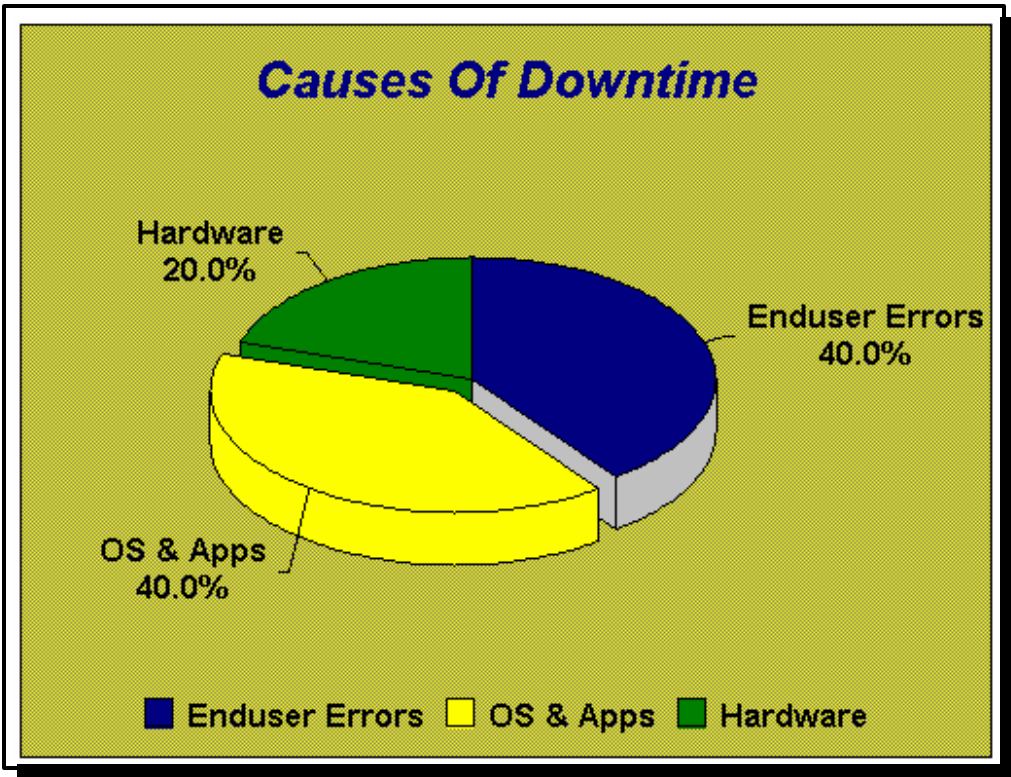


Figure 4. Causes of System Downtime (Source: Gartner Group)

Reducing unplanned downtime and maintaining peak performance with change management.

It is therefore important for you to adopt a rational update philosophy and to define an orderly update process. What represents “rational” and “orderly” will vary from customer to customer, but the following questions should be answered:

- How do we know when/if a problem is identified that puts our servers at risk?
- How frequently do we perform routine maintenance?
- How do we manage updates?
- How do we know that updates will not introduce new problems?

Monitoring for Change

Resolving critical issues before they cause downtime is something that can only be done if you have access to the proper information at the proper time, so your change management plan needs to be as flexible as your environment.

Best Practices:

- Regularly monitor critical update information and/or set up a task to review the critical updates for your hardware at least every 45 days.

Frequency of Routine Maintenance

Since most updates occur early in an element’s life cycle, it is extremely important to update frequently when using a new element such as hardware or operating systems. Also, organizations who are making frequent hardware/software configuration changes are less likely to encounter problems if their system software is maintained with up-to-date code. Many companies will have multiple environments, and therefore multiple update cycles may be appropriate.

Best Practices:

- Quarterly updates are recommended for dynamic environments and those using new hardware or operating systems.
- Semiannual updates are recommended for most customer environments.
- Annual updates should be adequate for static environments using mature technology.

Managing the Updates

Whenever possible maintain all system software at the appropriate levels, depending on the particular environment. Have a clearly defined methodology for performing updates.

Best Practices:

- Adopt a controlled, methodical approach.
- Apply to test systems first.
- In production, apply updates starting with low-impact servers and slowly move to business-critical servers.
- Select times for performing updates based on when they will least impact your business.
- Don’t update too many systems at a time. Make sure you have the resources required to recover should problems develop.

Managing the Risks

How do you know that updates will not introduce new problems? Risk is inevitable with change, and no one would claim otherwise. The goal of performing regular updates is to manage the change, and thereby manage the risk. When change is managed properly, problems occur less frequently and are usually easier to solve. When change is not managed, you can encounter unexpected or unanticipated outages (see *Figure 2*). At times like that, the goal is to get the server back online as quickly as possible. Generally, this does not give you the time to do a full root cause analysis. Over time, these unmanaged systems become more and more unstable.

When the best practices outlined above are followed, change can be effectively managed, but there are a few additional tasks that are important to any change management activity:

Best Practices:

- Maintain logs
 - Log all changes
 - Log all unusual events and problems
 - Keep logs where other people can access them
- Review logs as part of “problem determination”
- Reboot systems before making changes to be sure that no preexisting problems are present.

If you actively manage your servers, the more they will perform like a mainframe.

Who IS Responsible for Change Management?

Because of the complexity of the ISA environment, each customer configuration is unique and it is the customer who best understands his or her configuration. The customer has developed and tested this configuration over time and decided what is best for the business. So it should not be surprising that it is the customer’s responsibility to download and install designated system software updates made available via Web sites, CDs or other distribution methods, according to the vendor’s instructions.

How does IBM Help?

IBM provides the following tools to help the customers manage and implement a change management strategy:

Proactive e-mail

IBM provides proactive e-mail to advise customers of issues that arise with IBM ^ xSeries™ servers which helps determine if/when a problem is identified that puts servers “at risk.”. To register for this free service, press the “Profile” tab at the top of the page when you visit <http://ibm.com/pc/support>. You will receive information only on the systems of interest to you. A library of current critical update information is also maintained for xSeries servers at the same Web site under “Hot News”.

UpdateXpress

For the IBM xSeries servers, IBM provides an UpdateXpress CD on a quarterly basis as self-installable, executable update packages. An ISO image of the CD can be downloaded (at no charge) from the IBM xSeries support site at <http://ibm.com/pc/qtechinfo/MIGR-4JTS2T.html> or

Reducing unplanned downtime and maintaining peak performance with change management.

you can have an UpdateXpress CD shipped to you for a nominal fee from an IBM fulfillment site <http://www.pc.ibm.com/coupon>.

When the CD is released, it contains xSeries system software and management software that has been through IBM testing. It utilizes an HTML (hypertext markup language) interface with a similar and familiar look and feel as their web browser. UpdateXpress presents the information to the HTML user interface in XML (extensible markup language) format. The descriptions include the software update name, version number, and severity of the update. These update packages come in ready-to-run system file formats that are needed for the major operating systems to install directly from the CD or from a drive accessible across a network.

In order to facilitate the management of updates on your servers your IT administrator can use the release level of the UpdateXpress CD to provide a known standard of system software levels across your server environment. Your IT administrator can use this methodology to communicate with IBM service or colleagues at geographically remote locations to easily specify what level of UpdateXpress is currently deployed. Just think of how this would simplify the process of communication.

The severity rating identifies the customer impact and need to apply the update, relative to the previously released version of the update:

- **Critical** updates provide fixes to address data corruption, security or stability defects.
- **Suggested** updates provide some significant fixes to problems that customers are likely to encounter, new features, marketing announcements, new system initial release and refresh code.
- **Non-critical** updates provide minor fixes, textual changes or performance enhancements.

Note that severity is specified on a per update basis. Special cases may arise where a particular update is only critical for a particular machine-type or operating system. In these scenarios, the update severity should reflect the severity in the typical case for the customer and the installation notes should include additional text describing the exception.

UpdateXpress provides multiple ways to update your server. To update your server locally, the IT administrator places the UpdateXpress CD in the server then reboots the server. UpdateXpress automatically detects the levels of system software of the server and the options installed then proceeds to update them automatically if the IT administrator does not interrupt the automatic process. The IT administrator can also chose to perform the update manually.

To update your servers remotely, the remote distribution tool in UpdateXpress, *RemoteUX.exe*, allows you to push update packages to remote servers. You can also use this method to distribute other system software packages, through third-party management software, such as Microsoft SMS, Tivoli, CA Unicenter and HP Openview. UpdateXpress can now update individual servers point to point with knowing only an IP address, using HTML to describe the files. In addition, UpdateXpress also offers command line options for installing unattended automated updates, which can be remotely scripted.

IBM posts individual updates to the IBM Web site. You can download these packages when you don't need the entire UpdateXpress CD or when you need to obtain only a critical update that is not on the UpdateXpress CD. You can view the UpdateXpress CD like a "Service Pack". For the most part, you should view the individual updates as "Hot Fixes".

Reducing unplanned downtime and maintaining peak performance with change management.

Remote Deployment Manager

IBM Remote Deployment Manager (RDM) is another tool that can be used in change management. RDM has the capability of managing the total life cycle of the server's operating system, system software updates, configuration, restoration and retirement.

When used for updating, RDM can remotely deploy a complete server image including all the updates for not only the server and option system software but the operating system as well. RDM has the capability of importing UpdateXpress update packages and can deploy them remotely and unattended via the network to thousands of remotely distributed systems. RDM can also restore system hard drives utilizing PowerRestore technology from PowerQuest. This technology is installed with the RDM image at deployment time

With a series of keystrokes the IT Administrator can restore an initial deployment image or an incremental image of the whole system to a known running state. This is important in a change management environment when more then one system software update or patch is applied and the system is inoperable. PowerRestore can bring the system back to a known operation state while the sequence of events that brought the system down are investigated. This maintains your systems availability and keeps your business critical applications running.

Since RDM runs on its own server, RDM provides a repository of system images for security and instant deployment in case of system failure, system upgrade or system redeployment. The RDM server can be located anywhere on your network for security and backed up for total system redundancy. In other words, RDM provides complete life cycle management for operating systems, applications and data. *Figure 5* illustrates how you can deploy system images remotely without touching the systems.

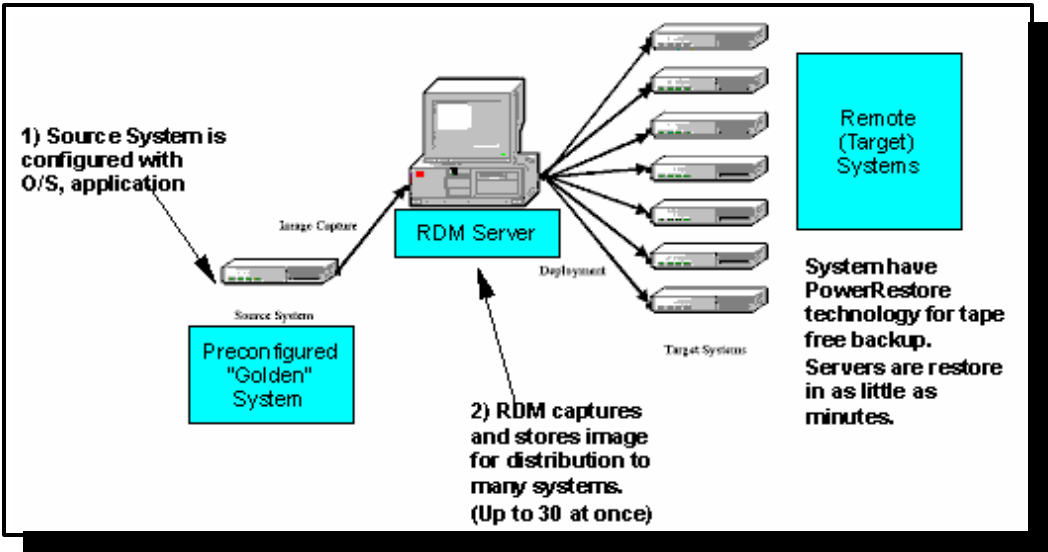


Figure 5. Typical RDM workflow and layout

IBM Director

IBM Director, built on industry standards, is another powerful suite of tools and utilities included with xSeries servers and designed to help you manage servers in the Intel environment. IBM Director automates many of the processes required to run your networks productively, including capacity planning, rack management, preventative maintenance, diagnostic monitoring, problem determination and other functions.

Reducing unplanned downtime and maintaining peak performance with change management.

IBM Director offers a graphical user interface for easy local and remote access and control. By letting IT administrators view the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks and memory, IBM Director can help you manage your servers with ease and efficiency. It also has several self-management features and a portfolio of proactive, predictive tools that enable advanced server management. For more information about IBM Director, see the "Systems Management for IBM xSeries Servers" white paper dated April 2002 (<http://www.pc.ibm.com/us/eserver/xseries/literature.html>).

Summary

Proactive Change Management makes good business sense, as opposed to being surprised by unexpected system downtime and a break-fix philosophy. Proactive change management using the tools outlined in this paper will help maintain your systems at the latest levels of system software. Being at the latest levels provides your company a major advantage in system performance and availability.

Whatever the reason, high availability, system failure reduction, system service or a newly added component, it is clear that the time for having a change management strategy is now. Systems on a proactive change management plan reduce unexpected system failures, lower help desk costs and avoid other costs associated with troubleshooting an unmanaged system.

When business-critical systems are down, they impact revenue and customer satisfaction. Without a change management strategy, you add unnecessary complexity to your environment, which wastes your IT staff's valuable time.

By implementing a change management strategy and using the tools provided by IBM in your Intel platform environment, you are taking the right step to avoiding costly and unexpected system downtime.

Reducing unplanned downtime and maintaining peak performance with change management.

© IBM Corporation 2002

IBM ~ Group

Research Triangle Park NC 27709

Produced in the USA

12-02

All rights reserved

IBM, the IBM logo the e-business logo, ServeRAID and xSeries are trademarks of IBM Corporation in the United States and/or other countries.

Intel is a registered trademark of Intel Corporation.

Microsoft is a registered trademark of Microsoft Corporation.

UNIX is a trademark in the United States and other countries licensed exclusively through The Open Group.

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

IBM reserves the right to change specifications or other product information without notice. IBM makes no representations or warranties regarding third-party products or services. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

IBM ^ xSeries servers are assembled in the U.S., Great Britain, Japan, Australia and Brazil and are composed of U.S. and non-U.S. parts.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.