IBM Systems and Technology Group
August 2004

Introduction to IBM TotalStorage DS4000 Storage Manager Copy Services

(formerly IBM TotalStorage FAStT Storage Manager)

It's no easy task to protect your data. Not only do you have to provide optimal access to your businesses' applications, there are dozens of other demands. Security is paramount in this changing world, as are business intelligence and analytics. As your business grows, your need for more storage grows. But the biggest issue affecting your valuable data is how to protect your data against planned and unplanned contingencies.

This white paper will cover the replication services available from the IBM® TotalStorage® DS4000 (formerly IBM TotalStorage FAStT Storage Servers) series of storage servers. Each of the replication features available in storage server are designed to provide the necessary technology to implement a solution for data protection and security. Together these features provide a range of solutions for business continuity initiatives.

IBM FlashCopy

IBM FlashCopy® is designed to provide a point-in-time copy of a logical drive. Using a pointer-based repository system, FlashCopy is designed to manage up to four copies of a logical drive. Upon the receipt of a FlashCopy command for a logical drive, called the base logical drive, the storage system initiates the creation of a new logical drive whose contents reflect all the data at the time of the FlashCopy command. As I/O continues to the base logical drive, all writes to the base logical drive cause a movement of the old data to the underlying repository to preserve the point-in-time copy of data. Any reference to the FlashCopy logical drive will access the base volume for data unchanged since the FlashCopy command and will access the repository for the older unchanged data if the base logical drive's data has been changed. (See figure: FlashCopy Creation)

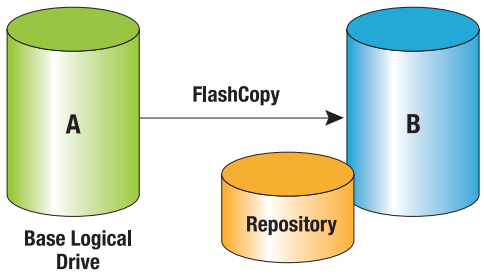

FlashCopy Creation

As an example, see figure: FlashCopy Access. FlashCopy logical drive B, which was created at 12 noon, will have a point-in-time copy of logical drive A. As an update ("y") for the block that contains "x" comes in, the system moves the old contents ("x") to the repository. Accessing logical drive A will return "y" whereas accessing logical drive B will return "x".

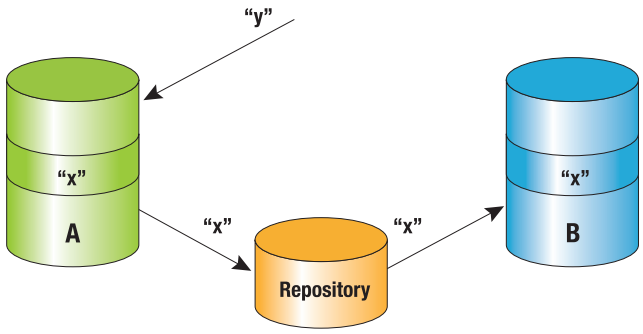

FlashCopy Access

FlashCopy's efficient storage management

FlashCopy is an efficient storage-based mechanism. Its design avoids the need for server cycles or I/O resources for the point-in-time process. The repository is typically sized at 20% (default setting) of the total storage used by the base logical drive. This percentage is calculated based on the estimated update activity of the logical drive for the expected life of the FlashCopy logical drive. For example, if FlashCopys are taken every six hours, then the 20% figure allows up to 20% change during the six-hour period. Most databases do not approach this figure. However, should it be necessary to retain the FlashCopy beyond its normal life, then the repository can be dynamically expanded to help allow for the extra storage. For logical drives with very low update activity, the repository can be set to an even lower percentage.

Using FlashCopy for fast database backups and restores

FlashCopy can be extremely useful in speeding up database hot backups. Normally, when a database system is placed into hot backup mode, the underlying database files are backed-up or copied to other media. As databases grow, the amount of time in hot backup mode can increase from hours into days. The major concern with hot backup mode is that the database can be vulnerable to failure since all database update activity is held in memory to allow the disk copy to complete. With growing databases and longer backup times, the window of exposure can increase.

FlashCopy can help reduce the hot backup window to minutes. Once a copy is taken, the database exits hot backup mode and resumes normal operations. The FlashCopy logical drives are then used to backup to offline media or copied to other media. This can help shorten the backup process.

FlashCopy logical drives have further uses. In the event that a database recovery is required, the FlashCopy logical drive can be used to initiate recovery. This helps eliminate the need to restore from offline media.

FlashCopy's efficient use of system resources

This particular technique is designed to be efficient in terms of overall system resources. Backing up to offline media can be slow due to the transfer rates of the offline media being slower than disk transfer rates. Copying to other disk media can have throughput issues. Using the host operating system software to copy database files may involve reading and writing the entire database which can subject the database to sub optimal data transfers (due to driver and HBA buffer limitations which break large block I/O into smaller I/Os). (See figure: Database Hot Backup)

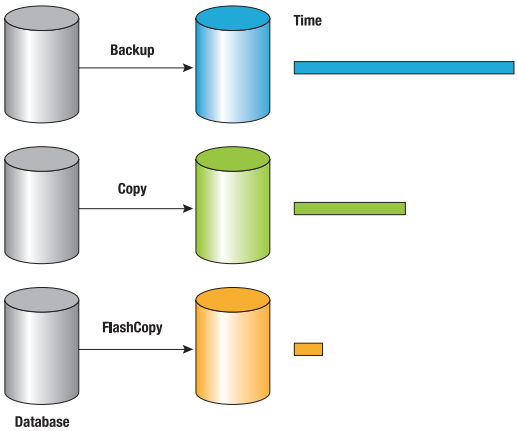

Database Hot Backup

Applications testing, vaulting, data feeds

FlashCopy can also be used for application testing, vaulting, or feeding data warehouse or datamarts. FlashCopy logical drives may be written to, which is designed to allow robust applications testing against real live data while avoiding impact to the production base logical drives.

VolumeCopy

VolumeCopy is designed to extend the replication functionality of IBM TotalStorage DS4000 series storage systems by providing the ability to copy the contents of one volume to another. This feature is completely implemented within the storage system and does not rely on host or application resources. Unlike FlashCopy, a volume copy is the exact same size (or larger) as the original. VolumeCopy would follow a point-in-time FlashCopy of the logical drive. The two tools are designed to allow the creation of a complete point-in-time copy or clone of a base volume, with minimal impact on the base volume. This clone can help provide additional protection from some classes of failures and can also be used for the uses listed above.

VolumeCopy's ability to create an on-demand full copy of a logical drive has numerous applications. It's key to optimizing the performance of storage systems as well as applications.

VolumeCopy to help optimize performance and data protection

As an optimization tool, VolumeCopy is designed to be used to migrate data for performance or data protection purposes (See figure: VolumeCopy). Logical drives that are configured to an array that experiences performance issues may use VolumeCopy to move one or more logical drives to another array within the same storage system. This can help relieve pressure on the balance of the logical drives on the original array. The target array may also provide more performance headroom than the original array, thus helping balance the workload.

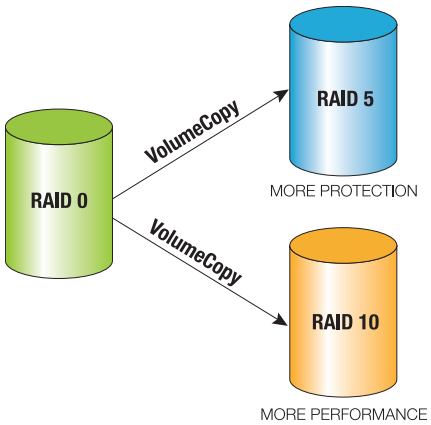

VolumeCopy

Logical drives residing on RAID 5 arrays may use VolumeCopy to migrate data to an array that is organized as a RAID 10 array. This can help provide data protection and improve throughput. Likewise, if a logical drive no longer needs the high levels of performance and data protection, VolumeCopy can be used to move it to a lower performance or data protection array.

VolumeCopy is designed to provide additional protection

Helping protect against local outages, VolumeCopy can be used to move data from one array to another to help provide additional measures against data loss. Creating a second copy of data on a different array can help protect against media corruption and yields a ready-to-use copy for recovery purposes.

FlashCopy and VolumeCopy

Used in conjunction with FlashCopy, VolumeCopy creates additional copies for multiple uses beyond that of backup (See figure: FlashCopy and VolumeCopy). A VolumeCopy of a FlashCopy image can copy the point-in-time logical drive to a different array, which will help prevent access of the new logical drive from impacting the original base logical drive. This will help prevent impact while the backup process proceeds. The new logical drive can be used in a number of ways including data vaulting, applications testing, and as a source for remote mirroring.

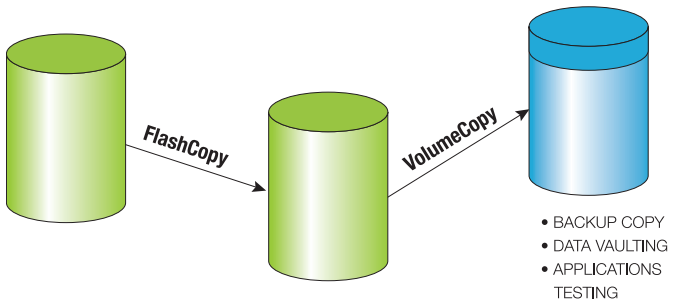

FlashCopy and VolumeCopy

Using FlashCopy and Volume Copy to help protect against database corruption

One very important use in the database environment is the ability to test the point-in-time copy for any media problems (media corruption) and database inconsistency (See figure: Database Checks). Database consistency checks may point to problems in the original database that have not yet surfaced. This can give the administrator a chance to fix any corruption issues before the database crashes. In addition, the VolumeCopy logical drive helps allow fixing logical errors on the original database through normal database tools. For example, accidentally dropped tables from the original database can be repaired using data from the point-in-time copy.

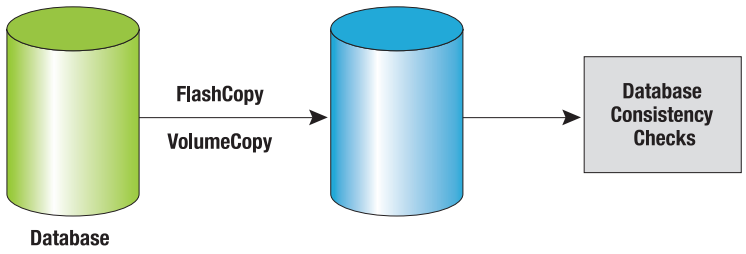

Database Checks

In summary, VolumeCopy is designed to provide tools necessary to meet new data access requirements as they change.

Enhanced Remote Mirroring

Enhanced Remote Mirroring consists of Global Mirror, Global Copy, and Metro Mirror and is designed to provide the ability to mirror data from one storage system to another over extended distances. It is designed to control synchronization from within the storage system so that it is nearly transparent to the host application servers. Unlike FlashCopy, Enhanced Remote Mirroring is designed to provide continuous updates from the primary logical drive to the secondary. This is designed to provide data availability, and is a key technology for implementing Disaster Recovery and Business Continuity plans.

The Metro Mirror mode is designed to provide synchronous mirroring for logical drives. Global Mirror mode is designed to provide asynchronous mirroring and includes the write-order consistency option. (See figure: Mirror Modes)

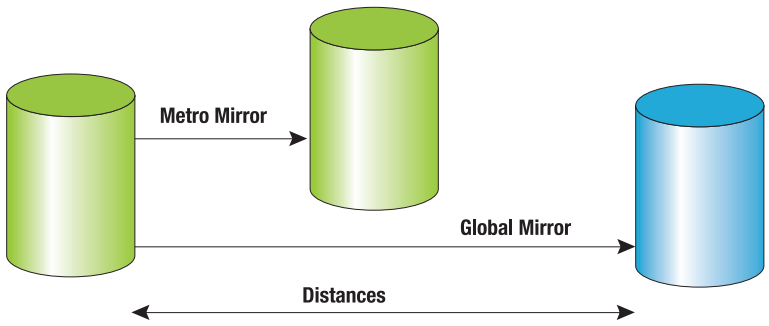

Mirror Modes

Up to 64 mirror pairs and one-to-many mirrors

Within the TotalStorage DS4000 series, logical drives may be mirrored among the various models at distributed locations for up to 64 mirror-pairs helping provide flexibility for the most complex requirements (See figure: Mirror Pairs). A single storage system may have a mirror relationship with multiple other storage systems, allowing flexibility in the designing of mirror topologies.

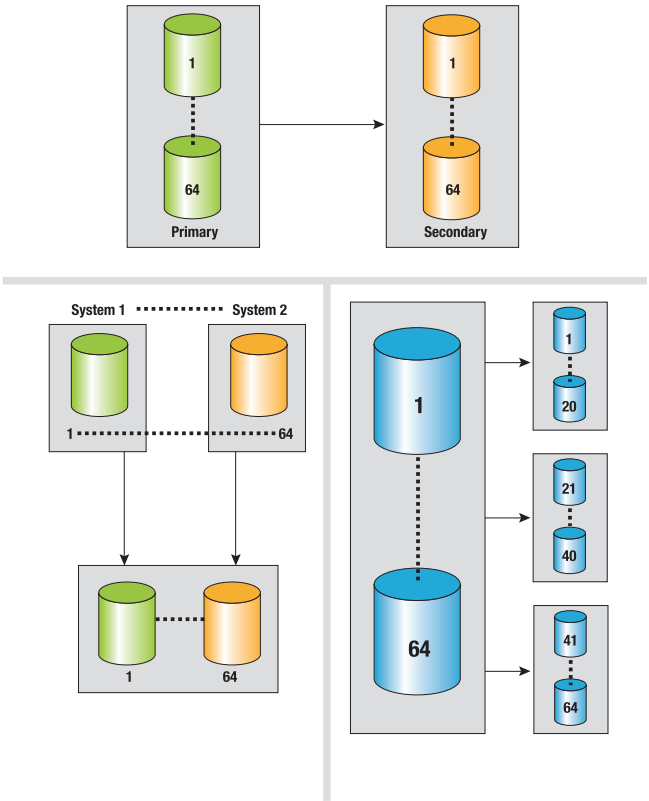

Mirror Pairs

Cross-mirroring

Mirror pairs established between storage systems may have mirroring in both directions. This is designed to allow each system to mirror its primary logical drives to the other to help protect both systems. (See figure: Cross Mirrors)

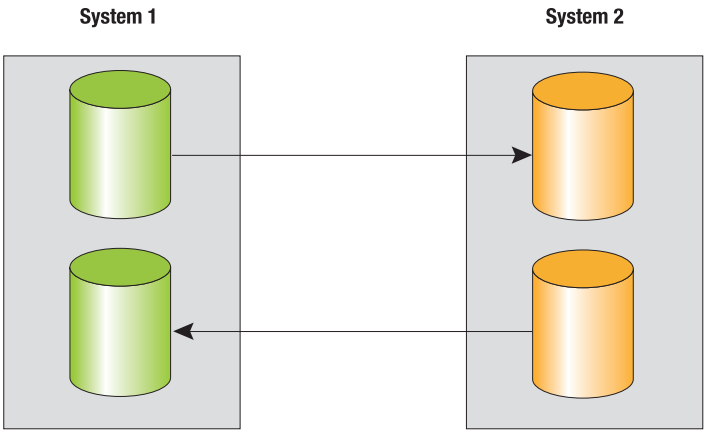

Cross Mirrors

Automatic mirror establishment and maintenance

It is designed to allow mirrors to be dynamically created, suspended, resumed, and deleted. During its initial creation, the storage system copies the contents of the primary logical drive to the secondary target logical drive before mirroring begins. It is designed to maintain synchronization for the life of the mirror pair.

Role reversal

The mirror pair relationship maintained between the primary and the secondary storage system may be reversed. In a disaster recovery scenario where the primary storage system fails, the system is designed to cause the secondary logical drives to assume primary responsibility and continue processing. At this time, the new primary system may mirror back to the original primary system to restore using the role reversal command. (See figure: Role Reversal)

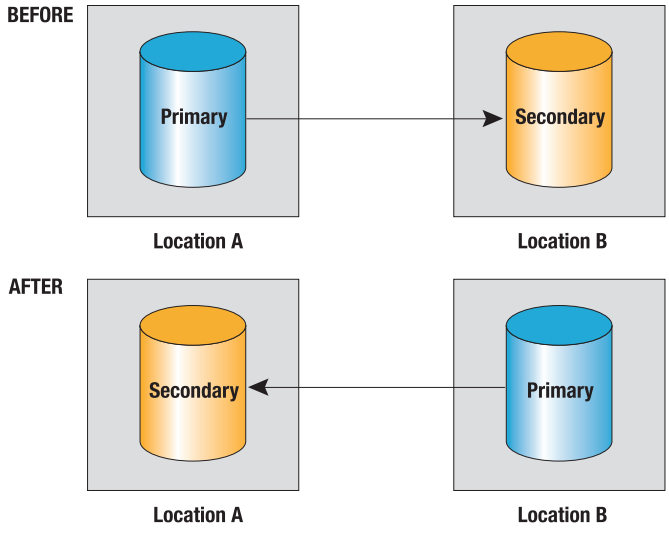

Role Reversal

Metro Mirror mode

Metro Mirror mode is designed to allow synchronous mirroring within a 10 km distance between the primary and secondary systems. This mode is used when accurate synchronization is required. Each server write initiates a write to the local primary logical drive as well as a remote write to the secondary logical drive. An acknowledgement from both the primary and the secondary system returns a complete status to the initiating host server. Metro Mirror is designed to offer one of the highest forms of protection, with both the primary and secondary volumes kept current to the last update, but the inherent delays in performing the remote write may cause application performance degradation.

Global Copy mode

Global Copy mode is designed to provide the ability to mirror logical drives over nearly unlimited distances. Using asynchronous mirroring technology, there is no need to synchronize the writes to both the primary and the secondary logical drives. Each write to the primary continues as usual, with the I/O completion status returned to the initiating host server and the write to the secondary logical drive sent to the remote site. Global Copy is designed to address the application performance issues raised with Metro Mirror.

Selecting the correct Mirroring Mode requires understanding the business value of the updated information and the costs and limitations of the infrastructure deployed, and should be part of an overall Disaster Recovery/Business Continuance strategy. Mixing modes, using Metro Mirror for highly critical data and Global Mirror for less critical data, may be an option.

Global Mirror with Asynchronous Consistency Group

Asynchronous writes by their very nature can lag behind the primary logical drive I/O. In the event of a disruption, the secondary logical drives may be in an inconsistent state. Successful database recovery at the secondary site will require that the writes to the secondary site be in the same order as the primary site. In order to provide the same write order consistency, Enhanced Remote Mirroring provides an additional Global Mirror mode, Global Mirror with Asynchronous Consistency. (See figure: Consistent Writes)

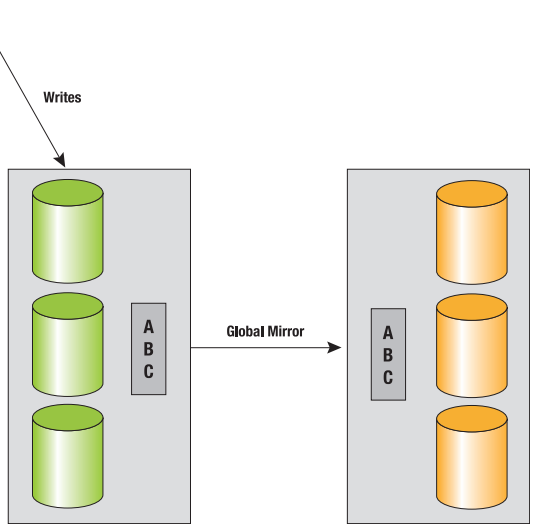

Consistent Writes

This mode is designed to be dynamically enabled and to provide write order consistency by ordering the I/Os to the remote site in primary I/O order. Logical drives that need to participate in this mode may be grouped into a write-order consistency group. Should the I/O of any one logical drive fail, the feature is designed to suspend mirroring to the secondary site until the problem is fixed and write-order I/O mirroring resumes.

Dynamic Mode Changes

Just as applications change, mirroring requirements can also change. Enhanced Remote Mirroring is designed to provide the ability to dynamically switch mirroring modes to meet changing applications and bandwidth requirements. Applications that use Metro Mirror mode may switch to Global Mirror mode with Asynchronous Consistency mode or Global Copy with synchronous mode. And as changes continue to occur, the user may reverse the mode switches.

Suspend and Resume

The Suspend command may be issued by operators or through a script to allow external processing. Typically, when a database is placed into hot backup mode, the remote mirror is suspended to allow a consistent copy to be taken using FlashCopy and possibly VolumeCopy. This is designed to provide a consistent copy at the secondary site for recovery purposes. (See figure: Suspend and Resume)

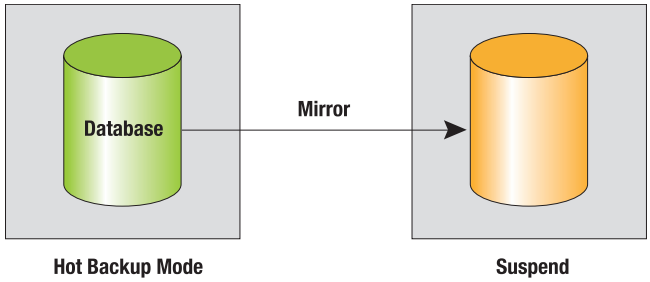

Suspend and Resume

The Resume command is used to restore mirroring. In the event of a communications fault, this feature is designed to suspend mirroring and can be set to automatically resume when the communications fault is repaired.

The Suspend and Resume features are available in all Enhanced Remote Mirroring modes. During the resume operation, the storage system uses its delta log bitmap to re-synchronize the secondary logical drives to the primary logical drives.

Read access of secondary logical drives

While mirroring is in place, Enhanced Remote Mirroring is designed to allow read access of the secondary target logical drives. FlashCopy point-in-time copies can be made of the secondary logical drives. For databases, this is designed to allow a consistent copy for reprovisioning of the database at the secondary site for disaster recovery purposes. (See figure: Remote FlashCopy)

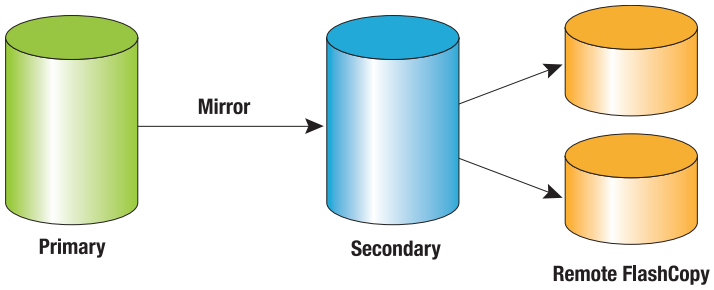

Remote FlashCopy

The wide range of copy services offered on the DS4000 series provides users with a powerful and flexible suite of services that can be used individually and in combination to develop and deploy customized data protection solutions for each user's requirements. These copy services can also be integrated with software products, such as IBM Tivoli® Storage Manager and IBM DB2® Universal Database™ as well as backup and restore and database offerings from many third parties, to help provide enhanced data protection and improved processing. IBM is uniquely qualified to provide the complete solution including servers, communications equipment, storage, services, and the copy services described in this document in order to help solve clients' business problems.

IBM Systems and Technology Group
9000 Rita Road
Tucson, AZ 85744

Produced in the United States
August 2004

This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only and is subject to change without notice. IBM shall have no responsibility to update such information.

IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein. Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESSED OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided.

References in this document to IBM products, programs or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM program or product in this document is not intended to state or imply that only that program may be used. Any functionally equivalent program or product that does not infringe IBM's intellectual property rights may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

TSW00454-USEN-00