



IBM Software Group – IBM Tivoli Training

IBM Tivoli® Storage Manager 5.3

Storage Pool Backups



@business on demand.

© 2006 IBM Corporation
Updated October 6, 2006

Objectives

Upon completion of this module, you will be able to:

- Explain how storage pools are used
- Plan for storage pool backups
- Perform a storage pool backup

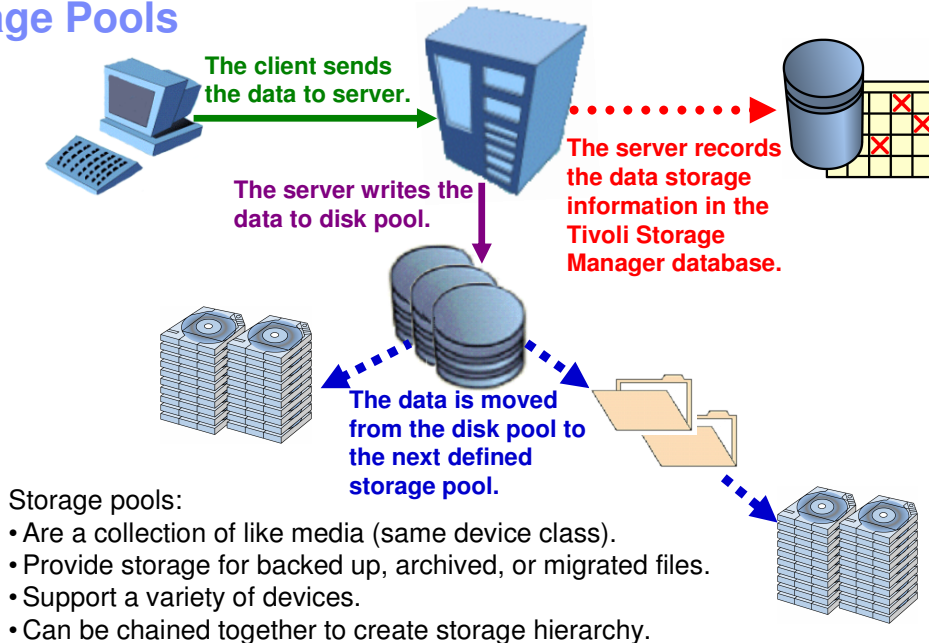
Agenda

- Overview
- Planning Storage Pool Backups
- Performing Storage Pool Backups
- Backing up Storage Pools to Disk
- Additional Resources

Section 1

Overview

Storage Pools



Data storage pools are where the server stores files that are backed up and archived. The database serves as the inventory or index to client files within data storage.

Data storage may be composed of direct access storage and sequential tape media.

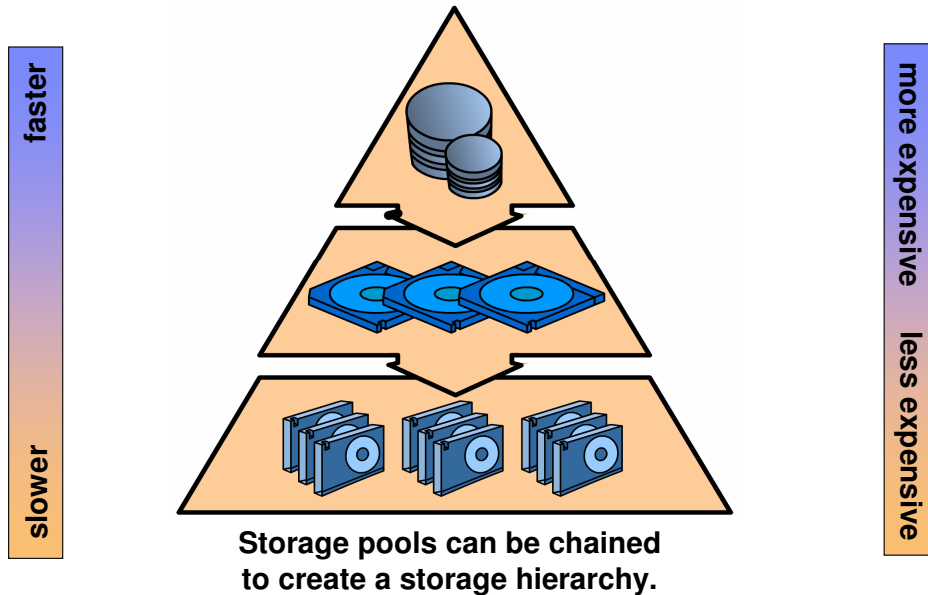
Files may be initially placed on different storage pools according to the desired storage management policy. Files are automatically moved to other devices to satisfy free space, space utilization, performance, and recovery requirements.

An administrator with system, storage, or operator privilege can manage data storage. This

includes planning, preparing, monitoring, and deleting storage volumes and storage pools depending on level of privilege class.

Data storage is actually defined as a collection of storage pools.

Storage Pool Hierarchies



A **storage pool** is a named set of volumes, belonging to the same device class, which is the

destination of backed up or archived data. A **device class** is a grouping of like devices, such

as disk or tape.

The purpose of storage pools is to match user requirements for data with the physical characteristics of storage devices. For example, if users need immediate access to certain data, you can define a storage pool, which consists of storage volumes residing on high performance

DASD (direct access storage device – an IBM term for a hard disk). Then, users can associate this storage pool as a destination for their files by binding the appropriate management class.

Storage pools can be chained to create a storage hierarchy.

Device Classes

Random Access:



Sequential Access:



Define additional tape device classes to reflect specific devices, media, and procedures.

```
DEFine DEVClass DEVType=type of device
def devc autodlt_class devt=dlt format=drive
library=autodltlib
```

A **device class** represents a set of storage devices with similar availability, performance, and storage characteristics. These device types are managed by Tivoli Storage Manager. Tivoli Storage Manager uses the device class to determine which device and storage volume type to use to:

- Store backup, archive, or space-managed data (primary storage pools).
- Store copies of primary storage pool data (copy storage pools).
- Store database backups.
- Export or import Tivoli Storage Manager data.

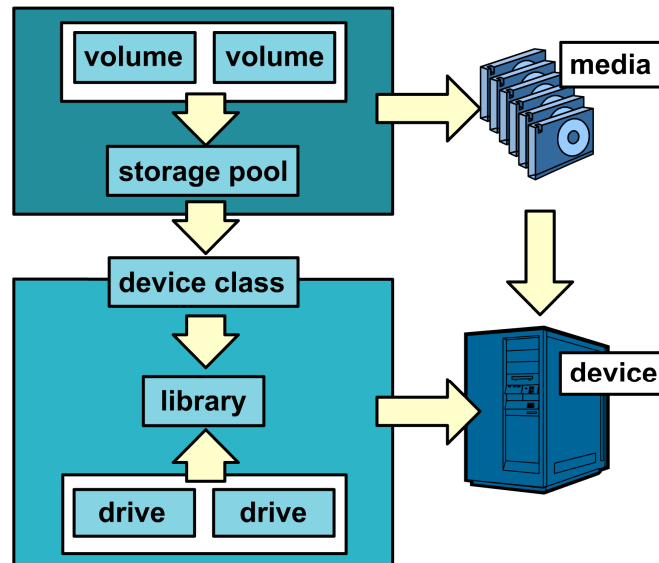
One device class can be associated with multiple storage pools. Each storage pool is associated with only one device class. Each device class is characterized by its device type,

which indicates the type of storage volumes that are used to store data.

For random access storage, Tivoli Storage Manager supports only the **DISK** device class. The **DISK** device class is predefined by Tivoli Storage Manager. However, you can define many storage pools associated with the **DISK** device class. You cannot modify the **DISK** device class.

For sequential access storage, Tivoli Storage Manager supports the following device types:

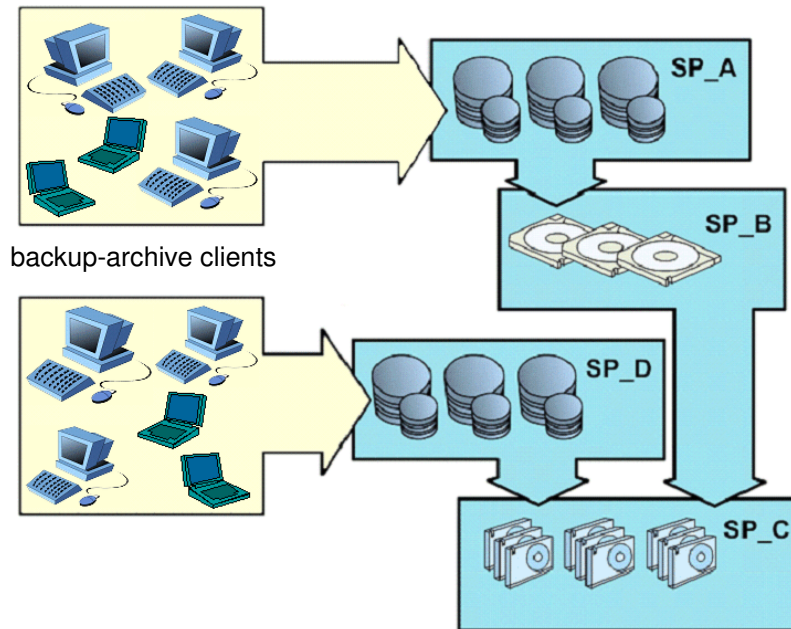
Device Class and Storage Pool



```
DEFine DEVClass DEVType=type of device
```

Each device is associated with a device class that specifies the device type and how the device manages its media. If you have two libraries you will need two device classes. Storage pools are mapped to a device class. It is through this mapping that, when data is written to or accessed from a storage pool, Tivoli Storage Manager knows the device characteristics of the storage pool media and how to access it. If the device is not listed on the IBM support page, it is probably not supported.

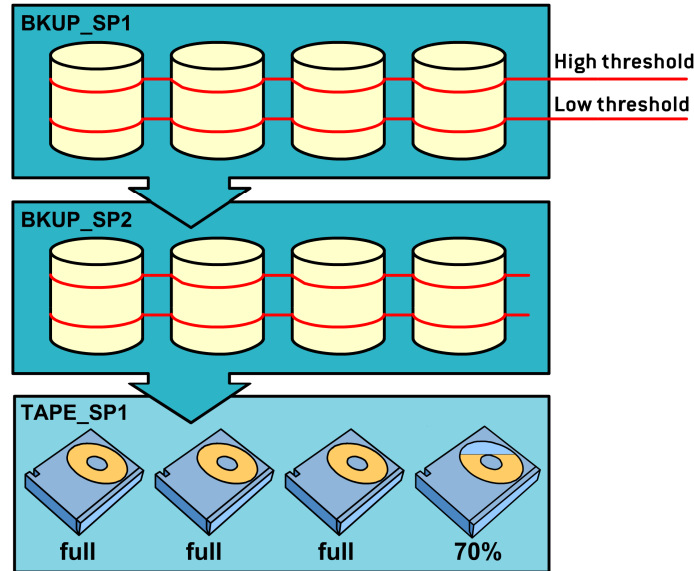
Storage Destination



A **storage destination** identifies the storage pool where client data is sent when it is backed up, archived, or migrated. It is specified in the backup and archive copy group definitions that are included in the management class. This is part of policy management, which is covered in Unit 6.

Data placement is also influenced by storage pool settings that may restrict read/write access and limit the size of files placed in the storage pool.

Automatic Data Movement—Storage Pool Migration



10

IBM Tivoli Storage Manager Storage Pool Backups

© 2006 IBM Corporation

Automatic data movement between storage pools is used to balance the performance and cost of different storage devices while ensuring adequate free space to satisfy new space allocations. This process is known as **migration**. For each storage pool, you define low and

high migration thresholds. The **low threshold** identifies the amount of free space needed to

satisfy the daily processing requirements of your business. The **high threshold** is used to trigger migration and ensure that enough free space is available while migration is performed. The difference between the high and low thresholds indicates the approximate amount of data that will be migrated.

When the high migration threshold is reached in a storage pool, Tivoli Storage Manager migrates files from the pool to the next storage pool in chain. No migration occurs if there is no next storage pool. Tivoli Storage Manager first identifies which client node has backed up or migrated the largest single file space or has archived files that occupy the most

space. When the server identifies the client node based on these criteria, the server migrates

all files from every file space belonging to that client for those files whose number of days in the storage pool exceeds the value specified by the **MIGDELAY** parameter.

Section 2

Planning Storage Pool Backups

Best Practices

When applying storage pool backup functions to your environment:

- Back up the entire primary storage pool hierarchy to same copy pools.
- Consider using at least one copy pool for on-site recovery and one for disaster recovery.
- Use administrative command scheduling to automate storage pool backups.
- Back up storage pools before backing up the database.
- Consider running the `audit volume` command periodically to mark damaged files.
- Run the `restore stgpool` command periodically to replace damaged files.



Here are some suggestions for applying the new storage pool backup functions to your environment.

- **Back up entire primary storage pool hierarchy to same copy pools.**

If you set up your configuration so that each primary storage pool hierarchy is copied to the same copy pools, IBM Tivoli Storage Manager is smart enough to recognize that when a file migrates down the storage hierarchy it does not need to be copied again into a copy pool. Or, if you move a file with the **move data** command, IBM Tivoli Storage Manager again can recognize this situation.

- **Consider using at least one copy pool for on-site recovery and one for disaster recovery.**

It may be useful to have one backup copy of a primary pool that is kept on-site and can be used for media failures or damaged files. Recall that if users access a file that is damaged in the primary storage pool, IBM Tivoli Storage Manager automatically lets the user access the copy in the copy pool, but only if the copy pool volumes are on-site.

It may be useful to have one backup copy of a primary pool that is used for off-site storing to be prepared for disaster situations. The primary pool is copied to on-site volumes in this special disaster recovery copy pool, and then the administrator changes the volume access mode to **OFFSITE** and physically moves the volumes to an off-site location.


- **Use administrative command scheduling to automate storage pool backups.**

The administrative command scheduling facility can be used to automate storage

pool backups. The `backup stgpool` command can be scheduled to initiate a

Creating or Defining Storage Pools

Use the following command to define a storage pool:



```
DEFine STGpool poolname devclass
```

For example:

```
DEFine STGpool lastinchain disk
```


To create a storage pool from the Administration Center:

1. Click **Storage Devices** in the **Tivoli Storage Manager** tree.
2. In the **Servers** table, select a server.
3. Click **Select Action**, select **View Storage Pools**.
4. In the server's storage pools portlet, click **Select Action**, select **Create a Storage Pool**.

During installation, Tivoli Storage Manager provides predefined random access storage pools:

- **BACKUPOOL** – A storage destination for user files that are backed up to the server.
- **ARCHIVEPOOL** – A storage destination for user files that are archived to the server.
- **SPACEMGPOOL** – A storage destination for files migrated from user nodes.
- **DISKPOOL** – Only for Windows® systems.
- **FILPOOL1** – A default storage pool of device class **FILEDEV1** created as a default **Next Storage Pool** for both **ARCHIVEPOOL** and **BACKUPOOL** when you install a minimum configuration on Windows® only.

Define a Copy Storage Pool



```

DEFine STGpool poolname deviceclassname POoltype=COpy
DESCription=description ACCess=READWrite|READOnly|UNAviable
COLlocate=No|GrouP|NODE|Filespace REClaim=percent
RECLAIMProcess number RECLAMATIONType=THRESHold
OFFSITERECLAIMLimit= NOLimit|number MAXSCRatch=number
REUsedelay=0|days OVFLocation=location
DATAFormat=NATive|NONblock CRData=Yes|No

```

You can use the Administration Center to copy storage pools.

Use the **define stgpool** command to name a set of volumes as a primary storage pool or a

copy storage pool. The volumes in a pool belong to the same device class. A copy storage pool is used to store backup copies of files that are in primary storage pools.

POOLNAME

Specifies the name of the copy storage pool to be defined. This parameter is required. The maximum length of the name is 30 characters.

DEVCLASSNAME

Specifies the name of the sequential access device class to which this copy storage pool is assigned. This parameter is required. You can specify any device class except **DISK**.

POOLTYPE=COPY Specifies that you want to define a copy storage pool. This parameter is required.

DESCRIPTION=DESCRIPTION

Specifies a description of the copy storage pool to be defined. This parameter is optional. The maximum length of the description is 255 characters.

IEA-TSM-STGPool.txt
ACCESS=ACCESSMODE

Section 3

Performing Storage Pool Backups

Backup of Storage Pool Files

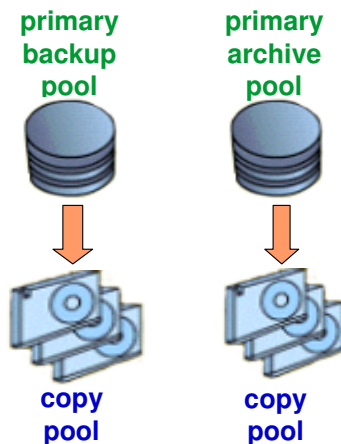
Back up to copy pool:

- Is independent of client backups.
- Retains file level granularity.
- Is an incremental backup of files.

Backups can be initiated:

- Manually with command or Administration Center.
- Automatically with administrative command scheduling.
- Once per primary pool.

Multiple primary pools can use one copy pool.



Copy storage pools can only be sequential.



Backup of the primary storage pools to copy storage pools should be done every day, and can be scheduled using the TSM scheduler, or they can be run manually.

Use the **backup stgpool** command to create backup copies of files that reside in a primary

storage pool, and store the backup copies in a copy storage pool. If a file is already duplicated in the specified copy storage pool, a new copy of the file is not made in that copy pool.

One primary pool can use multiple copy storage pools. Storage pool backup allows for files

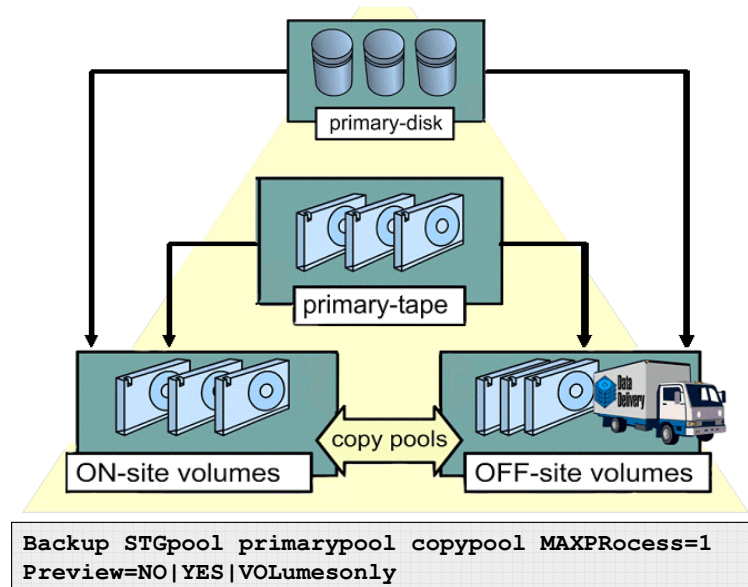
to be backed up into multiple copy storage pools. You have many copy storage pools for each primary storage pool; you are not limited to two or three backup copies. Each copy pool requires additional database and storage pool space, a potentially important consideration.

By supporting multiple copy storage pools, TSM supports both media and disaster recovery

by physically separating these copies. This will reduce the risk of losing data as a result of fire, flood, earthquake or other disaster.

You create multiple backup copies of a primary pool by initiating the backup of the primary storage pool for each copy pool. The incremental backup is against each copy storage

Backing up the Storage Pool Hierarchy



17

IBM Tivoli Storage Manager Storage Pool Backups

© 2006 IBM Corporation

Use the **backup stgpool** command to create backup copies of files that reside in a primary

storage pool, and store the backup copies in a copy storage pool.

To back up files in primary storage pool hierarchy to copy storage pool, on-site, issue the following commands:

```
backup stgpool primary_dsk onsite
```

```
backup stgpool primary_tape onsite
```

The number specified by **MAXPROCESS** does not limit the number of tape drives used.

Storage Pool Recovery


Storage pool recovery automatically switches to a duplicate copy if the primary file is damaged.

To restore one volume in a storage pool, issue the command:

```
restore volume
```

To restore a complete storage pool, issue the command:

```
restore stgpool
```



- **Restore one volume in a storage pool**

Restore individual volumes within a storage pool after a media failure.

- **Restore a complete storage pool**

Restore all damaged volumes within a storage pool after complete loss of a storage pool.

Section 4

Backing up Storage Pools to Disk

Disk-only or Tapeless Backup Overview

There are two types of disk pools:

- Random access

Tivoli Storage Manager provides a predefined disk device class that is used with all disk devices.

- Sequential access

To create large, sequential access file-type storage pools, issue the following command:

```
define devclassname devtype=file
```

Since each volume in a file device class is a *file*, the volume name must be a fully qualified file name.

The storage volumes are actually files. Data is written sequentially into standard files in the file system of the server machine. You can define this device class by issuing a **define devclass** command with the **DEVTYPE=FILE** parameter.

Both disk storage pool formats also have the capability of providing a space trigger for automatic allocation of storage pool space, which can result in increased efficiency of disk storage pools.

This will reduce the potential for disk fragmentation and maintenance downtime.

Do not use raw partitions with a device class type of **FILE**.

Disk-only backup has been enhanced in order to take advantage of the inexpensive disk storage currently available on the market. These improvements to sequential-access FILE device type and random-access DISK device class storage.

Capabilities that are added to device type FILE include support for storage pools to span multiple file systems, parallel migration and reclamation, better performance by minimizing fragmentation, direct I/O on platforms that support direct I/O.

Both disk storage pool formats also have the capability of providing a space trigger for automatic allocation of storage pool space, which can result in increased efficiency of disk storage pools.

Disk Storage Pool Considerations

Consider the following when planning disk storage pools:

- Have enough disk space available for at least one night's backup.
- For optimum performance, send large files directly to tape.
- If you are migrating data to tape on a daily basis, do not use mirrored disks for storage pools.
- If you are planning to use disk as a long-term storage, take into account:
 - ▶ the total cost
 - ▶ reliability
 - ▶ how data can be moved off-site for disaster recovery



You want to have disk storage pool capacity that is equal to or greater than one night's backup of your client data.

Some clients may be backing up large files, such as database files, that might better be sent

directly to tape. These files should not be part of the disk storage pool calculations.

If you are migrating data to tape on a daily basis, consider using unmirrored disk for storage pools. Although this introduces some risk of loss, it may be mitigated by cost savings and the time the data stays on the disk.

If you are planning to use disk as a long-term storage medium, carefully consider the total cost, reliability, and how you will move the data off-site

When considering using only disk for backup-archive storage, do a realistic analysis of the total amount of storage you'll need using your data retention policies, total data, and growth expectations. Evaluate the cost of tape vs. the cost of disk for storing all your data.

Section 5

Additional Resources



IBM Tivoli Storage Manager Training

http://www-306.ibm.com/software/tivoli/education/edu_prd.html

Country/region [select] | Terms of use

IBM

Home | Products | Services & solutions | Support & downloads | My account

← Software

Software > Tivoli > Training >

Tivoli software training

Course listing sorted by product

Tivoli. software

Product | Solution | Training type

Select title to view description and link to enrollment/worldwide schedule information.

If you are unable to locate a course at the time and location you desire, including training for prior product releases, please feel free to contact your delivery management team:

Americas: tivamedu@us.ibm.com
 Asia Pacific: tivtrainingap@au1.ibm.com
 EMEA: tived@uk.ibm.com

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

Select from this list of products:

IBM Tivoli Storage Manager

Go

IBM Tivoli Access Manager for Business Integration


Title	Type
-------	------

Suggested Reading

- Product Documentation
- IBM Redbooks

IBM Redbooks

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



Trademarks, copyrights, and disclaimers

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

IBM	CICS	IMS	MQSeries	Tivoli
IBM (logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)/business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

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 registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of 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.

Linux is a registered trademark of Linus Torvalds.

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time 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. 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 Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, 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. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

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

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2004,2005. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.