

System z Expo

October 13 – 17, 2008 – Las Vegas, Nevada



Session Title: VTAPE Update and Usage

Session ID: zEG11

Speaker Name: Ingo Franzki

Authorized
IBM | Training



© 2008 IBM Corporation

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

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 logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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 improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

§ Basics

- Remote vs. VSAM
- Formats (AWS, Zipped, PTF)
- Transferring tape images
- Actions

§ News

- DVCDN
- VTAPE QUERY
- Simplified Use of DLBL Statements for VSAM Files
- Tivoli Storage Manager Support

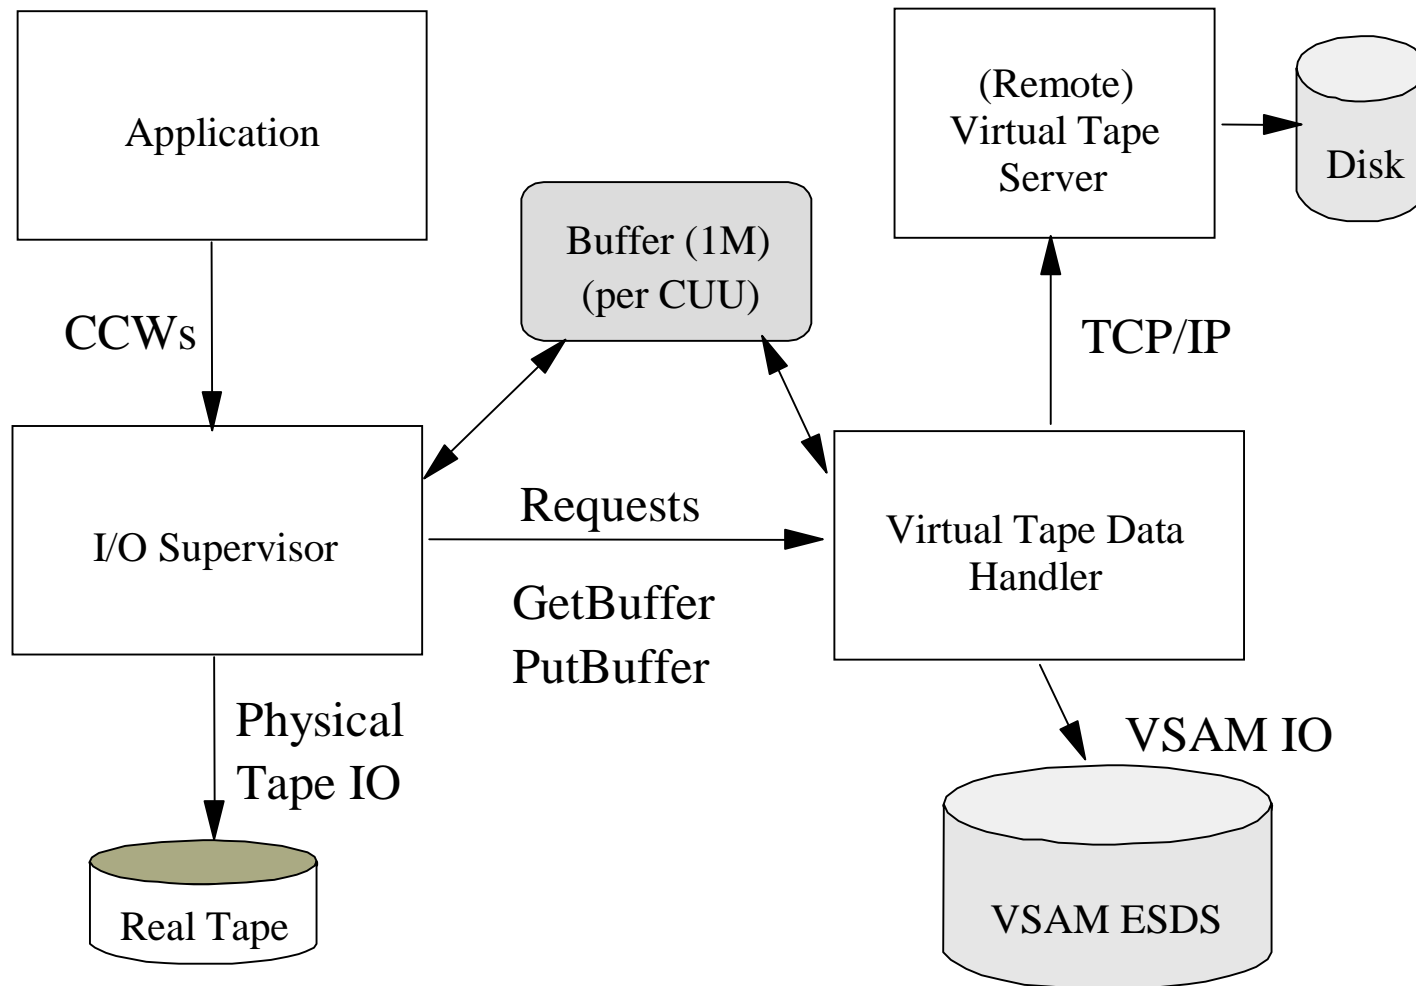
§ Hints & Tips

- Performance

Overview

- § VSE VTAPE support is part of z/VSE since VSE/ESA 2.6
- § Emulates a complete tape
 - Can contain multiple tape files, not just one tape file
- § Uses a tape image file instead of a physical tape
- § Tape image file can reside in
 - VSAM ESDS
 - Remote file (e.g. on a workstation)
- § Tape image file has AWSTAPE format
 - known from P/390, R/390, Hercules, Flex-ES

Overview



Overview - continued

§ VTAPE command

- VTAPE START,UNIT=<cuu>,LOC=<ip or VSAM>,FILE='filename'
- VTAPE STOP,UNIT=<cuu>

§ Tape image file is opened at VTAPE START

§ Tape image file is closed at VTAPE STOP

§ Access to tape image can be

- READ - read only
- WRITE - read and write (existing content is kept)
- SCRATCH - read and write (content is cleared)

Virtual Tape Data Handler Partition

- § Runs in a batch partition
 - Dynamic class R per default
- § Accesses VSAM tape images
- § Establishes TCP/IP connections to remote system
- § Startup job TAPESRVR
 - Skeleton SKVTASTJ in ICCF lib 59
- § Do not change the job name (TAPESRVR)

```
* $$ JOB JNM=TAPESRVR,DISP=L,CLASS=R
* $$ LST CLASS=A,DISP=D
// JOB TAPESRVR START UP VSE TAPE SERVER
// LIBDEF *,SEARCH=( PRD2.CONFIG,PRD1.BASE,PRD2.SCEEBASE )
// ID USER=VCSRV
// EXEC $VTMAIN,SIZE=$VTMAIN
/*
/&
* $$ EOJ
```

VSAM tape images

- § A VSAM tape image resides in a VSAM ESDS Cluster
- § Recommended file attributes
 - CI size = 32768
 - Record Size = 32758 (32768-10)
 - REUSE = YES
 - Shareoption = 1
 - Records/Cylinder/Tracks depends on amount of data
- § Skeleton SKVTAPE in ICCF lib 59
- § Size limit is 4GB (because VSAM ESDS size limit)

Remote tape images

- § A remote tape image resides in a file
 - on the remote workstation's file system
- § The file is created automatically (if not existing)
- § Filename (and path) are remote system dependent
- § Be careful with uppercase translation of filename
- § Remote systems can be
 - Windows (95/98/NT/2000/ME/...)
 - Linux (on zSeries or Intel)
 - Unix (Aix, Sun, HP, ...)
 - Any Java capable platform
- § Size limit depends on file system (e.g. FAT, NTFS, ext, ...) of remote system

Transferring tape images

§ Transfer tape images in binary

- From one workstation to another workstation
 - Network drive
 - File transfer
- From a workstation to a VSAM ESDS cluster

```
ftp 9.164.186.20          - initiate FTP session
bin                      - transfer in binary
quote site lrecl 32758   - specify record size
quote site recfm v      - record format variable
put d:\backup.aws VSAM.TAPE.IMAGE - transfer the file
```

Usage Example - Backup to CD-ROM

§ Step 1: Do a backup to a virtual tape

```
// JOB BACKUP
VTAPE START,UNIT=480,LOC=9.164.186.20,FILE='d:\backup.aws',SCRATCH
MTC REW 480
MTC WTM,480
// EXEC LIBR
BACKUP LIB=PRD2 TAPE=480
/*
VTAPE STOP,UNIT=480
/ &
```

§ Step 2: Copy (burn) d:\backup.aws to a CD-ROM

§ Step 3: Archive the CD-ROM

§ Step 4: Restore directly from CD-ROM

Usage Example – Reduce offline time

- § Step 1: Backup all your files, databases, libraries, ... onto separate VSAM VTAPES
 - During this time your CICS needs to be down (offline)
- § Step 2: Backup all the VSAM VTAPE files to a real tape
 - This can be done while CICS is up (online)
 - Can even be done on a separate VSE system, if the DASDs are shared
- § Reduces the backup time (offline time)
- § Only one real tape needed for all backup steps

Usage Example - Dump offload

§ Create a (remote) virtual DUMP tape

- DUMP Fx, cuu
- DUMP SVA, cuu
- DUMP SUF, cuu
- Offload function in IUI Dialog (Fastpath 43)

§ Send the tape image containing the DUMP to IBM for analysis

- Attach it to a e-mail
- Put in onto IBM's FTP server

Usage Example - PTF install

- § Order one or more PTFs via IBM ShopzSeries
 - You will get a notification when the PTF is ready for download
 - Download the file containing the PTF(s)
 - e.g. eptf5375.bin
 - Unzip (if zipped) and rename it to eptf5375.ptf
- § Apply the PTF(s) as usual
 - Use PTF Apply Dialog 1423
 - Use Virtual Tape pointing to the downloaded PTF data
 - Indirect apply is also possible

News – by VSE release/version

z/VSE release	Function
VSE/ESA 2.6	§ VTAPE functionality was introduced
VSE/ESA 2.7	<ul style="list-style-type: none"> § Binary PTF data (for PTF install) § Zipped AWSTAPE and PTF data § Removed DVCDN/DVCUP requirements
z/VSE 3.1	§ FakeTape (FLEXES) and zipped FakeTape
z/VSE 4.1 & 4.2	<ul style="list-style-type: none"> § AWSTAPE with ZLIB compressed records § Support for Tivoli Storage Manager (TSM) § QUERY command to query active VTAPES § Simplified Use of DLBL for VSAM VTAPES § Many Interactive Interface Dialogs have been adopted to allow use of VTAPE

Note: "FLEXES" and "FakeTape" are trademarks of Fundamental Software, Inc.

News – Supported formats (as of z/VSE 4.1 & 4.2)

Format	VSAM	Remote
AWSTAPE	Yes	Yes
Zipped AWSTAPE	No	Yes (read only)
FakeTape (FLEXES)	No	Yes
Zipped FakeTape (FLEXES)	No	Yes (read only)
AWSTAPE with ZLIB compressed records	No	Yes
Binary PTF data	No	Yes (read only)
Zipped binary PTF data	No	Yes (read only)

Note: "FLEXES" and "FakeTape" are trademarks of Fundamental Software, Inc.

Zipped AWSTAPE format

- § To save disk space a AWSTAPE image can be zipped (PKZIP)
 - one or more tape images in one ZIP file
- § To save download time
 - Extended VTAPE START filename syntax
 - Allows to read directly from a zipped image
 - without unzipping it first

```
VTAPE START,UNIT=cuu,LOC=ip-addr,FILE='zip-file.zip!aws-file.aws'
```

Example: AWSTAPE image prod1.aws resides in ZIP file d:\images\products.zip

```
VTAPE START,UNIT=cuu,LOC=ip-addr,FILE=d:\images\products.zip!prod1.aws'
```

Binary PTF data

- § Simulates a PTF tape (read only)
 - can be directly installed with II Dialog 1423
- § Input data is a binary PTF stream
 - PTF job stream with LRECL=80 (binary)
 - As downloadable from IBM ShopzSeries
 - PTF data can also reside in a ZIP file
- § Extended VTAPE START filename syntax for ZIPed data
 - File must have extension '.PTF'

```
VTAPE START,UNIT=cuu,LOC=ip-addr,FILE='ptf-file.ptf'
```

```
VTAPE START,UNIT=cuu,LOC=ip-addr,FILE='zip-file.zip!ptf-file.ptf'
```

FakeTape format

- § FakeTape format is known from FLEXES
 - Note: "FLEXES" and "FakeTape" are trademarks of Fundamental Software, Inc.
- § Allows to work with FakeTape tape images
 - read FakeTape files
 - format is detected automatically
 - Create FakeTape files
 - File extension must be '.fkt' or '.faketape'
 - Otherwise AWSTAPE format is created
 - Also able to read zipped FakeTape files

AWSTAPE with ZLIB compressed tape records

- § To save space, the tape records stored within a AWSTAPE file can be compressed using ZLIB 1.2.1.
 - Every single tape record is compressed separately
 - Compression ratio is not as good as zipped AWSTAPE files
- § Per default uncompressed tape images are created
 - You can force compression by using a file extension of '.zaws'.
- § ZLIB compression is only supported for remote tapes

VTAPE Exits

- § Virtual Tape Server can call exits (also called actions)
 - VTAPE START (open of tape image)
 - VTAPE STOP (close of tape image)
- § Can be used for tape management tools
 - start backup of tape image
 - notify operators
- § The IBM provided default action can execute a command or shell script/batch script

```
VirtualTapeServer.properties
...
action=com.ibm.vse.vtape.DefaultAction
actionparam=open:open.bat close:close.bat
...
```

VTAPE QUERY command

§ With **z/VSE 4.1** a new QUERY command has been added

```
>> _ _ _ _ _ VTAPE QUERY _ _ _ _ _ ><
  | _// _ |   | _ , UNIT=cuu _ |
```

§ If the **UNIT** operand is omitted, information about all virtual tapes will be displayed

§ If the **UNIT** operand is specified, detailed information about the specified virtual tape cuu is displayed

§ **Example:**

```
R2 0047 Display all virtual tapes
R2 0047 -----
R2 0047
R2 0047 182    9.152.2.70, 2386          TAPE.AWS          WRITE
R2 0047
R2 0047 181    VSAM                    VTAPE1             WRITE
R2 0047
R2 0047 -----
```

Simplified Use of DLBL Statements for VSAM Files

- § From **z/VSE 4.1** onwards, you are no longer required to have the DLBL for VSAM files used with VSE VTAPE in the system standard label area
- § Instead, you can specify the DLBL statements directly in the job that issues the VTAPE START command
 - The VTAPE START command will then transfer the label information to the tape server partition.
- § With this support, you can (for example):
 - Define a new VSAM file (using IDCAMS DEFINE CLUSTER).
 - Use this cluster in the same job with a virtual tape.
 - You are no longer required to add the DLBL statement to the system standard labels
- § Prior to z/VSE 4.1, the DLBL of the VSAM file used with VSE VTAPE had to reside in the system standard labels
 - The updating of the standard labels was only possible in the BG partition
 - This prevented you from being able to define a new VSAM cluster, and then use this cluster with VTAPE in the same job, *unless* the job ran in BG.

Simplified Use of DLBL Statements for VSAM Files

Job running in F4

```

$// JOB BACKUP
// ON $CANCEL OR $ABEND GOTO VTAPSTOP
* FIRST DEFINE THE CLUSTER
// EXEC IDCAMS,SIZE=AUTO
  DEFINE CLUSTER ( -
    NAME (VTAPE.TEST.FILE) -
    ...
    CATALOG (MY.USER.CATALOG)
  IF LASTCC NE 0 THEN CANCEL JOB
/*
* NOW DO THE BACKUP INTO THE NEWLY CREATED CLUSTER
// DLBL MYVTAPE, 'VTAPE.TEST.FILE', ,VSAM,CAT=MYUCAT
// DLBL MYUCAT, 'MY.USER.CATALOG', ,VSAM
VTAPE START,UNIT=480,LOC=VSAM,FILE='MYVTAPE'
// EXEC LIBR
  BACKUP LIB=PRD2 TAPE=480
/*
/. VTAPSTOP
VTAPE STOP,UNIT=480
/&

```

Running in R1

VTAPE Data Handler

Label Area of R1

→ IJVTcuu
→ IJVCcuu

VSAM OPEN is performed
using IJBVTcuu DLBL

Tivoli Storage Manager - Functions

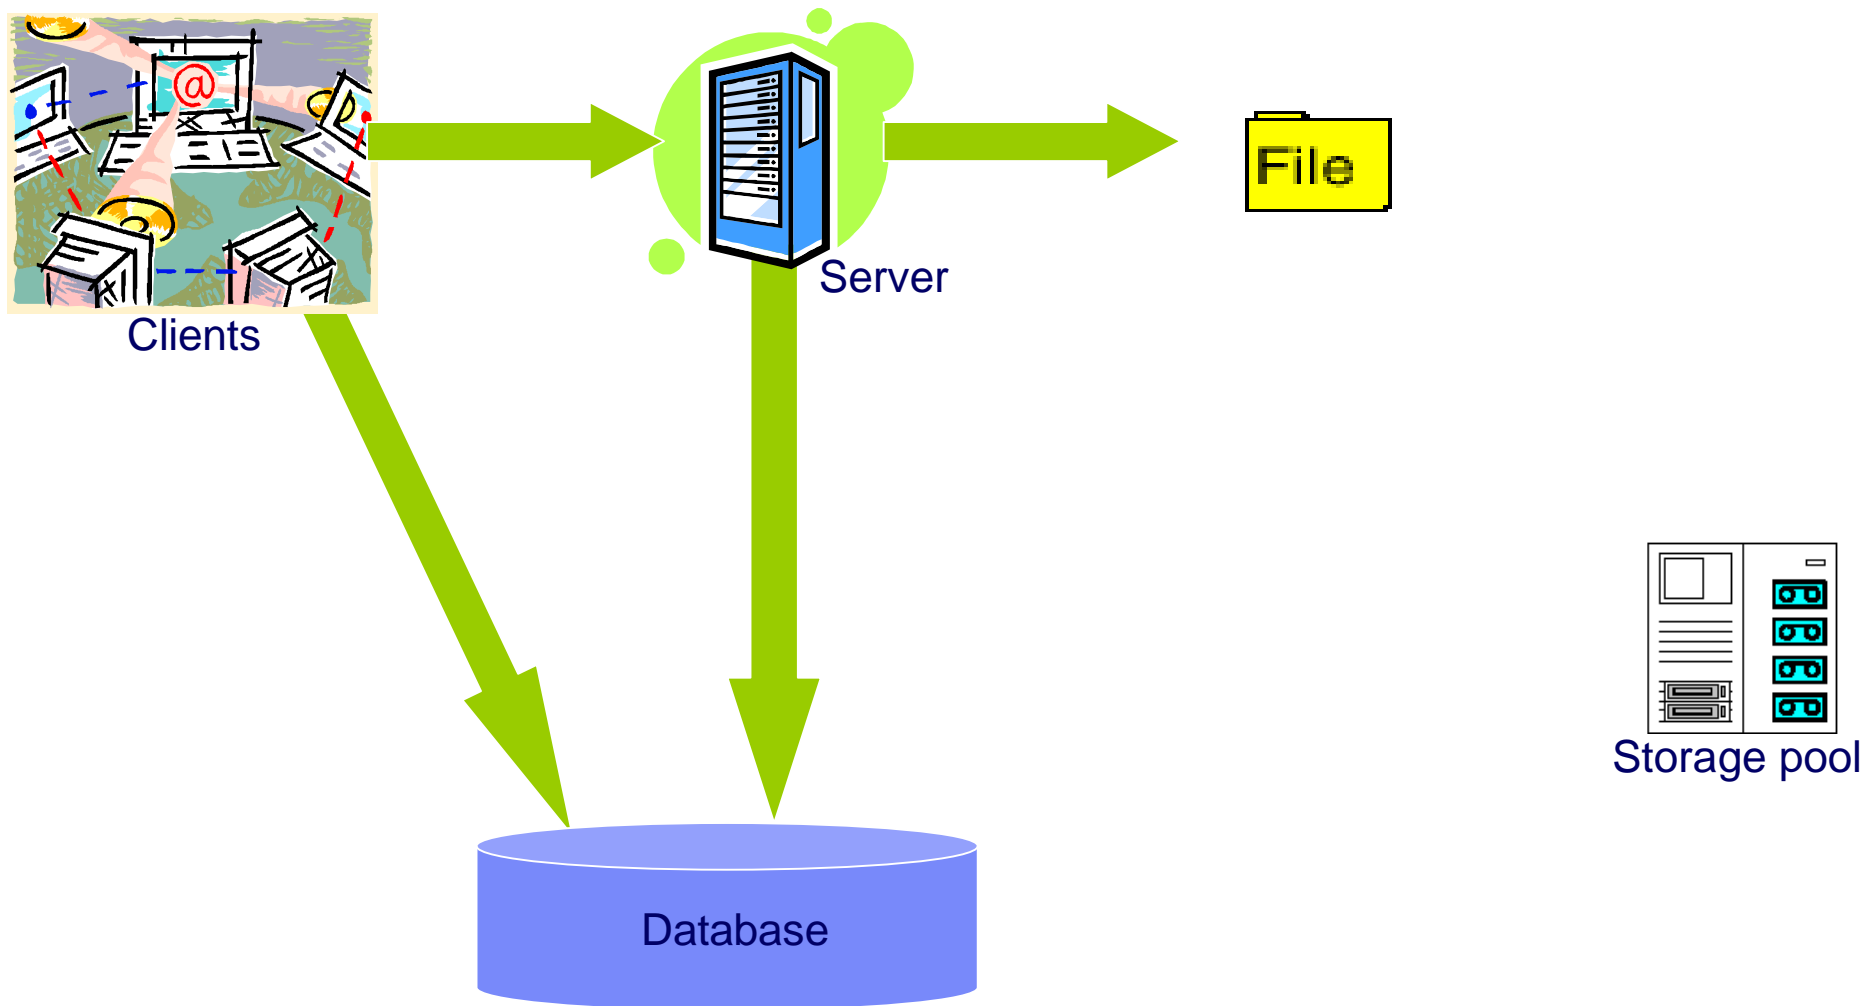
§ Backup and Restore

- Make a copy of a file or application data for backup purposes
- Backup can be started on a scheduled time or manually
- On restore the backup copy will be applied to the destination system

§ Archive & Retrieve

- Make a copy of a file or application data to archive it for a certain timeframe without being able to modify it (audit-save)
- Archive process can be started on a scheduled time or manually
- On retrieve the archived copy will be applied to the destination system

Tivoli Storage Manager - Architecture



Tivoli Storage Managers – Backup VSE data

§ New with z/VSE 4.1

- Allows to integrate VSE backups into an existing TSM environment

§ Uses the TSM Command-Line interface (DSMC)


- Available on many different (middle tier) platforms, e.g. Linux on System z

§ Based on the VTAPE Function of VSE

- Complete tape images will be backed up via TSM
- VTAPE OPEN/CLOSE Exit (so called Actions)
- At OPEN time, the tape image will be restored via TSM
- At CLOSE time, the tape image will be backed up via TSM

§ Existing Backup Jobs can be reused almost unchanged

Tivoli Storage Managers – Backup VSE data



Tivoli Storage Managers – Backup VSE data

Backup of a VSAM Clusters using TSM

```

* $$ JOB JNM=VSAMBKUP,DISP=L,CLASS=0
// JOB VSAMBKUP
// LIBDEF PHASE,SEARCH=IJSYSRS.SYSLIB
* THIS JOB BACKS UP VSAM DATASETS
// DLBL IJSYSUC,'VSESP.USER.CATALOG',,VSAM
*
* THIS FUNCTION USES A VTAPE FOR OUTPUT
VTAPE START,UNIT=181,LOC=9.152.216.105,FILE='TSM:VSAM.AWS(BACKUP)',SCRATCH
// ASSGN SYS005,181
// EXEC IDCAMS,SIZE=AUTO
        BACKUP ( -
                VSAM.CONN.SAMPLE.DATA -
                . . . -
                ) -
        REW -
        NOCOMPACT -
        BUFFERS(3)
/*
// ASSGN SYS005,UA
VTAPE STOP,UNIT=181
/&
* $$ EOJ


```

Syntax:

TSM:<name>(<mode>,<optionset>,
 <fromdate>,<fromtime>)

mode	- BACKUP or ARCHIVE
optionset	- Name of the configuration
fromdate	- Date (for Restore)
fromtime	- Time (for Restore)

Tivoli Storage Managers – Backup VSE data



Hints & Tips for best performance

§ VSAM virtual tapes

- Traditional VSAM tuning helps to increase performance
 - Buffers
 - Optimization for sequential processing

§ Remote virtual Tapes

- Performance tuning falls back to TCP/IP and network tuning
 - Make sure the TCP Receive Window Size is set to 32K
 - Delayed-Ack Fix (\$SOCKOPT) – see next slides
 - High network throughput requires increased CPU Power
 - Reduce the number of hops between VSE and the VTAPE server
- Make sure the partition priorities are right
 - ...,TCP/IP,VTAPE Server, ... , Job that uses VTAPE, ...

Hints & Tips – Delayed Ack Problem

§ For best performance you should assemble your \$SOCKOPT phase with the following option set in the socket option flag:

- \$OPTSNWT.
- This setting fixes a problem that is known as "Delayed ACK" problem.

§ Please make sure you also have applied ZP15E101 for TCP/IP 1.5E or APAR PK38492

```
§ // EXEC ASMA90,SIZE=ASMA90,PARM='SZ(MAX-200K,ABOVE),
      PUNCH      ' PHASE $SOCKOPT,* ,
$SOCKOPT CSECT
*
* * This phase is used by the BSD-C interface to allow global options
* * that affect the operation of all sockets in a partition.
      SOCKOPT CSECT,          Generate a csect          X
      BSDCFG1=$OPTMECB+$OPTSNWT, Socket options flag      X
      CLST=-1,                Seconds to wait for close   X
      CSRT=59,                Seconds before socket reuse X
      SSLLIBN=KEYLIB,         SSL library name           X
      SSLSUBN=SSLKEYS,        SSL sublib name            X
      SSLMEMN=MYKEY512,       SSL member name for keys   X
      SSLDEBG=00,             SSL debugging flag         X
      SSLFLG1=00, 80=req_close_notify_alert             X
      SSLFLG2=80, 80=do not use HW-Crypto               X
      SSLSTOR=80,             SSL storage usage          X
      SYSID=00                Use this TCP/IP sysid
*
      END $SOCKOPT
```


VTAPE programming interfaces


- § Programming interface for tape image implementers
 - Java interface `TapelImage`
 - Allows to implement own tape images
 - Tape image must be able to read/write tape records

- § Programming interface for Actions
 - Java interface `VirtualTapeAction`
 - `beforeOpen(...)`
 - `afterClose(...)`
 - `DefaultAction` implementation is able to execute operating system commands

- § For more information please contact
 - zvse@de.ibm.com

AWSTAPE Format

Known from P/390, R/390, Hercules, Flex-ES



Header (6 bytes)

Block Size (2 bytes, little endian)

Prev Size (2 bytes, little endian)


Flags (2 bytes)

X'2000' ENDREC Start of tape record in block

X'4000' EOF Block is a tape mark

X'8000' NEWREC End of tape record in block

AWSTAPE Format - continued



AWSTAPE Format – tools and resources

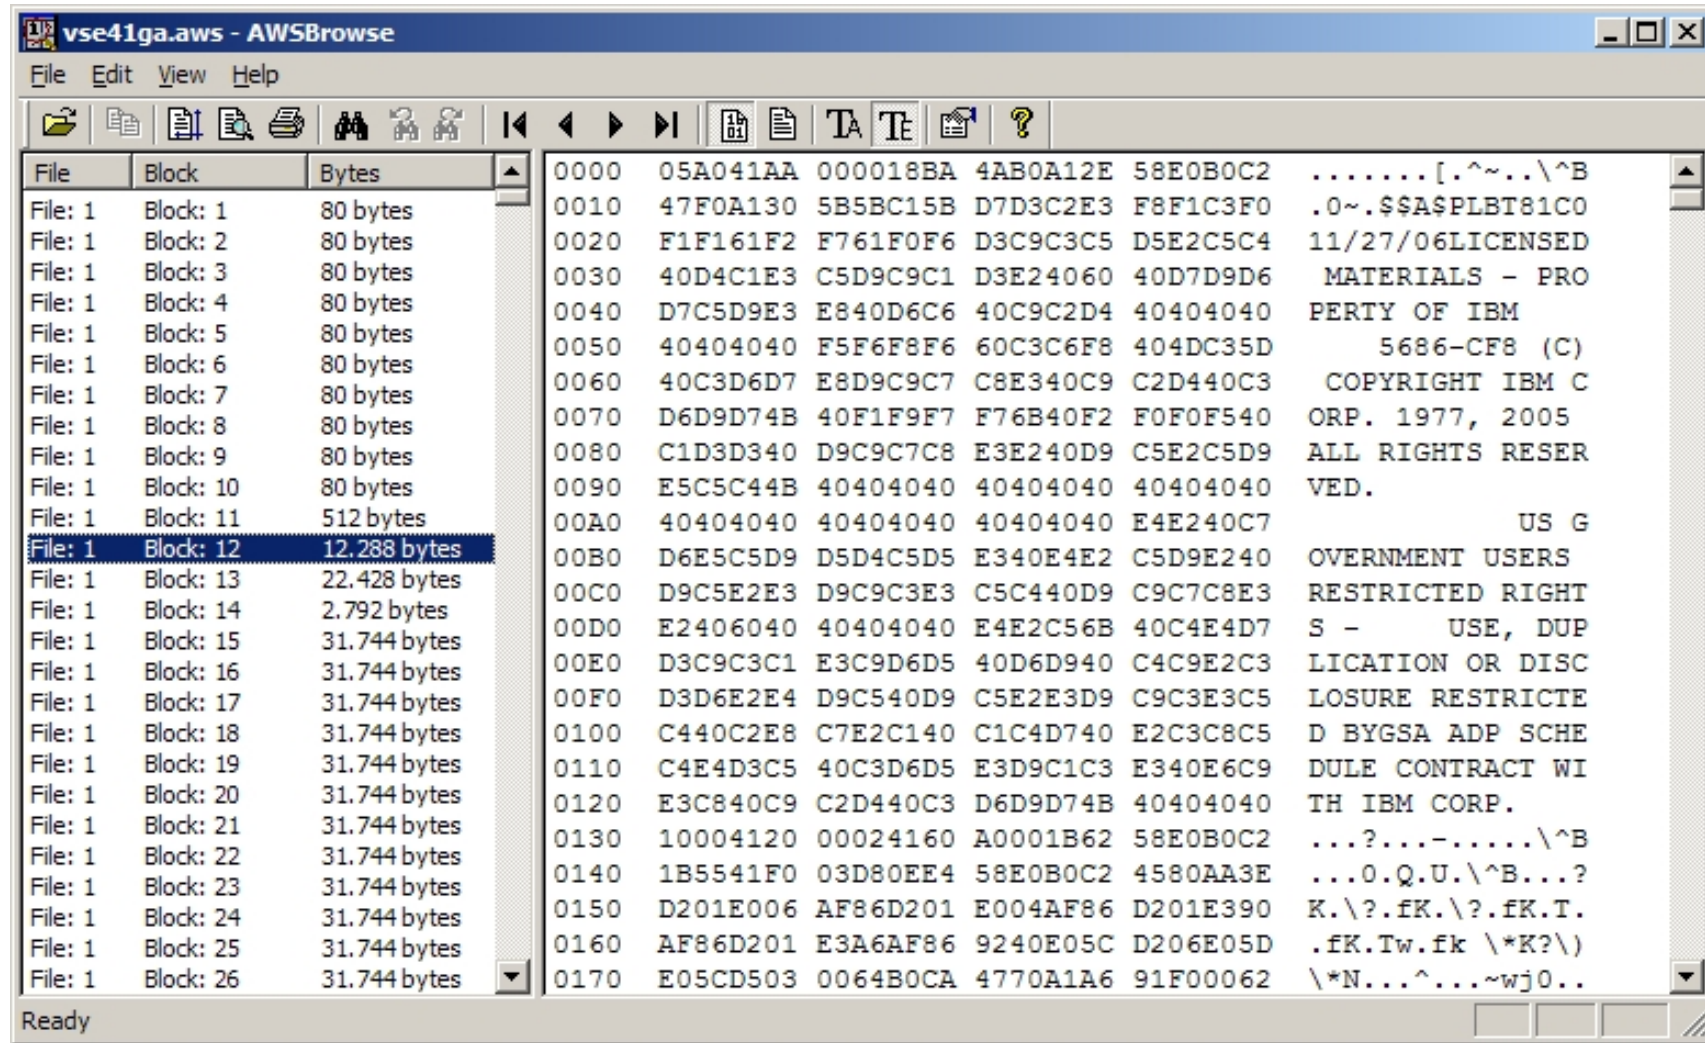
§ AWSTAPE format descriptions:

- <http://www.bustech.com/support/techtips/mas/awstape.htm>
- <http://www.cbttape.org/awstape.htm>

§ Tools:

- **AWSBROWS:** PC based browser for AWSTAPE files:
<http://www.cbttape.org/ftp/adhoc/Awsbrows.zip>
<http://www.cbttape.org/~fish/AWSBrowse-1.5.1.1805-bin.zip>
- **LISTVOL1 Tool:** Reads the first 2 tape records of tape image in AWSTAPE format residing in a VSAM ESDS cluster and prints the VOLSER and FILEID from the VOL1 and HDR1 labels on the tape.
<http://ibm.com/vse> goto Downloads and then Tools
- **Data Extract Utility:** Extracts Tape files from a AWSTAPE image and stores each file in a separate file on your PC.
Comes as part of VSE Virtual Tape Server

AWSBROWS – Browser for AWSTAPE files



LISTVOL1 Utility – What files are on my virtual tapes?

- § This tool helps to manage VSE VTAPE images stored in VSAM files
- § It can display the VOLSER and File-ID of the files stored on the virtual tape without mounting it as VTAPE.
- § The ListVOL1 utility reads the first 2 tape records of tape image
 - You can specify multiple VTAPE files
- § It prints the VOLSER and FILE-ID from the VOL1 and HDR1 labels on the tape:

```
LISTVOL1 UTILITY - LIST VOL1/HDR1 LABELS OF VTAPES

FILENAME: VOLSER  FILE-ID
-----
VTAPE1   : TAPE00  TAPE.DATASET.00
VTAPE2   : PRDDAT  PRODUCTON.DATA
VTAPE3   : BACKUP  MY.BACKUP.FILE
-----

LISTVOL1 UTILITY - FINISHED
```

Extract Tool – Extract data from a AWSTAPE file

§ Extracts Tape files from a AWSTAPE image

- stores each file in a separate file on your PC

§ Comes as part of VSE Virtual Tape Server

§ To run it:

– Extract.bat

```
set classpath=.;VirtualTape.jar;vtapetools.jar;%classpath%  
java com.ibm.vse.vtape.tools.ExtractFiles %1
```

– Extract.bat <awstape-file>

– Output is one or more PC files named <awstape-file>.0 ... <awstape-file>. n.

VSAM Considerations

Tape data is reblocked into the VSAM records

Hdr	1. tape record			Hdr	2. tape record	
rest of 2. rec.		Hdr	3. tape record		Hdr	4. rec.
rest of 4. tape record			Hdr	Hdr		

Independent of VSAM record size

Chose as large as possible (32758 bytes)

Independent of tape record sizes


Up to 64K

VSAM Considerations - continued


- § No overwrite of existing data possible
- § Only append to file possible
- § But: rewrite from the beginning is possible
 - Implicit reopen of the VSAM file with REUSE
- § Why?
 - ESDS can not delete records
 - ESDS can update records
 - But cannot change length of records
 - This would be necessary to support AWSTAPE format

VSAM Considerations - continued

Real tape



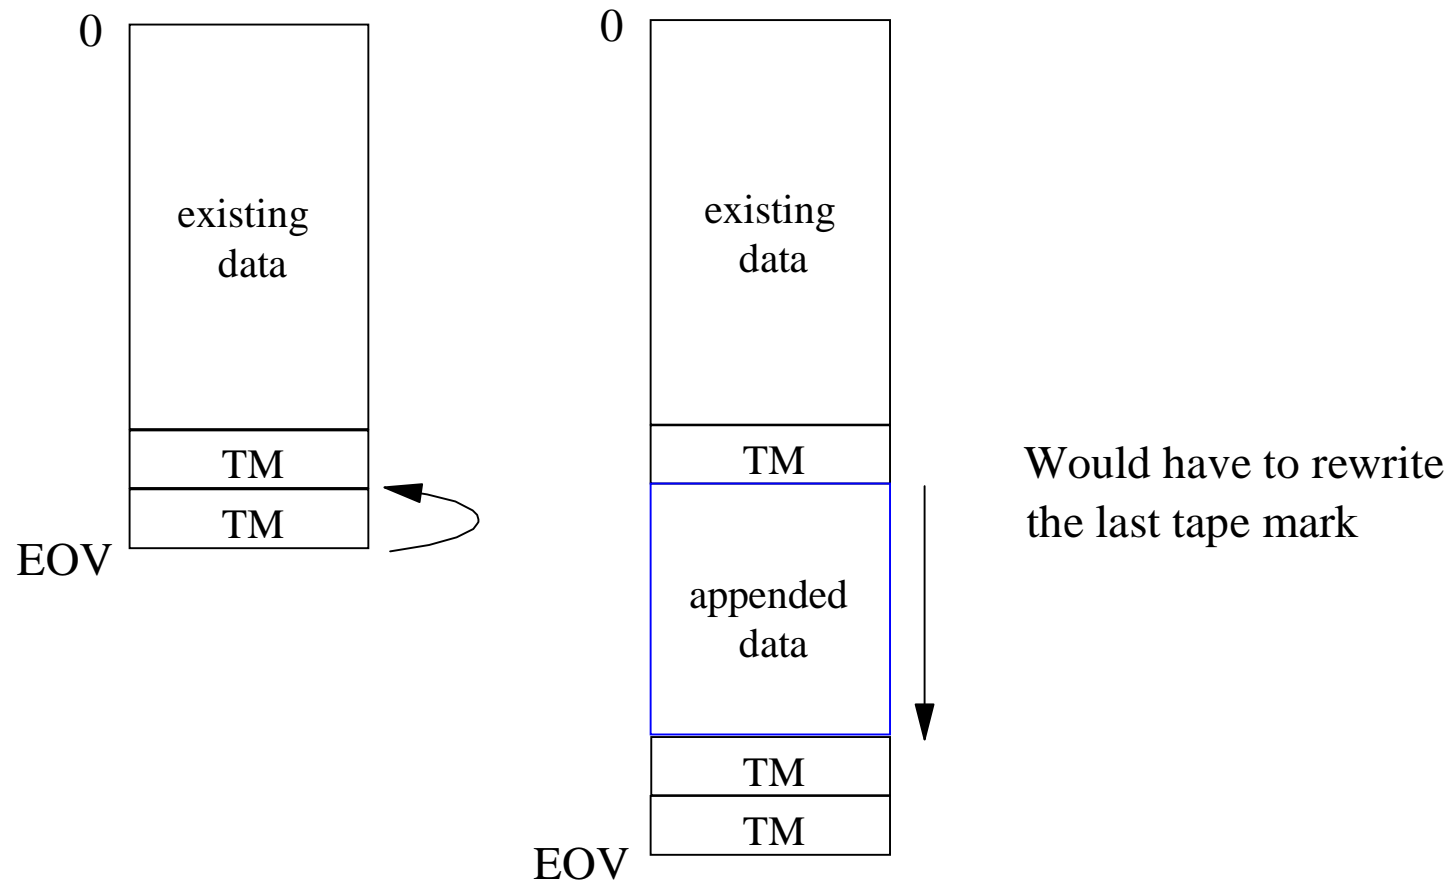
Virtual tape (ESDS file)



Would have to delete the records at the end of the tape image

VSAM Considerations - continued

Special case: Append of data to a tape



Note: The I/O supervisor will in most cases buffer the last Tape Mark, until VTAPE STOP, synchronisation or repositioning. This allows BAM to append additional tape files.

VSAM Considerations - Summary

- § Append of tape data is possible
 - NOT after VTAPE STOP
- § Overwrite from the beginning is possible
 - If defined with REUSE=YES

Questions ?

