



Oxford University TSM Symposium 2005

A Guided Tour of the TSM Client for Windows

Andrew M. Raibeck
Tivoli Storage Manager
Development
storman@us.ibm.com

Agenda

- This presentation will cover TSM client features that are specific to the Microsoft Windows operating system
 - Adaptive subfile backup
 - Open File Support (OFS)
 - Journal based backup (JBB)
 - Cluster support
 - System state support

Adaptive subfile backup

- "Differencing" algorithm used to identify changes between two versions of a file.
 - Not a method for selecting files for incremental backup.
- Backs up only those portions of a file that have changed since the base backup version was made.
- Originally designed for remote laptop users backing up across low bandwidth connections (e.g., modem).
- Available for backup only (not archive).
- Maximum supported file size: $1 \text{ KB} \leq n \leq 2 \text{ GB}$

Adaptive subfile backup – How to use

- Must be enabled on server:

```
SET SUBFILE CLIENT | NO
```

- Must also be enabled on client:

```
SUBFILEBACKUP YES
```

```
SUBFILECACHEPATH <cache_path_dir>
```

defaults to ..\baclient\cache

```
SUBFILECACHESIZE n
```

n is MB, $1 \leq n \leq 1024$

- Exclude the subfile cache path from backup!

Adaptive subfile backup – How it works

- Create a reference file in the subfile cache:
 - The reference file represents the base version.
 - Only file data (not metadata) is stored in the cache.
 - Reference file is created according to size:
 - size < 1 KB or size > 2 GB: Too small or too large, not stored
 - 1 KB ≤ size ≤ 3 MB: Data is copied byte-for-byte
 - 3 MB ≤ size ≤ 2 GB: Block signatures used to create reference file

Adaptive subfile backup – How it works (2)

- Create a reference file in the subfile cache (2):
 - File name is hashed to form the file names used in the cache.
 - .base file contains the file data.
 - .meta file contains a 256-bit signature.
 - .base file is sent to the server along with the signature.
 - Entry created in cache control database
 - Cache control database located in subfile cache directory
 - File name is .client_cache_db
 - Contains info on file name, "compression" achieved, differencing method
 - Entries stored in LRU (Least Recently Used) chain
 - If cache fills up, least recently used reference files are purged

Adaptive subfile backup – How it works (3)

- Create a delta file:
 - A reference file must exist in the cache.
 - The reference file integrity is validated with the cache database.
 - A maximum of 20 delta backups can occur before a new base file must be created.
 - If the last delta backup did not achieve at least a 40% compression ratio, then a new base file will be created.
 - Differencing algorithm computes the difference (delta) and creates a delta file in the cache.
 - ACL and named stream data is not a part of the differencing computation.
 - Delta file is sent to the TSM server along with base file digital signature.
 - ACL, named stream data sent along with delta file.
 - Delta file is deleted from the cache.

Adaptive subfile backup – How it works (4)

■ Restore:

- Client restores both base and delta components.
 - Only the base and the corresponding version's delta files are necessary.
 - Components are restored to x:\~tsmtemp\ (same file system as file).
 - Components can be restored in any order.
 - Components do not need to be restored sequentially.
- Delta component is duplicated.
- Base and duplicate delta are used to reconstruct file data over restored delta file. Non-file data (ACL, alternate named streams) in delta file are left intact.
- Restored file is moved to final restore location.
- Base and duplicate delta are deleted.
- Note that subfile cache is **not** used at all during restore.

Adaptive subfile backup – How it works (5)

■ Potential concerns:

- Over-commitment of file system:
 - Three components need to restore file (base, delta, duplicate delta).
 - Since restore of components need not be sequential, there can be many base or delta components in the temporary restore location waiting on their respective delta or base components.
- If the files are restored by a different user, a problem could occur when renaming the file to its original name.
- Although many customers have incorporated this feature into their local server backups, this feature was designed for low-bandwidth mobile laptop users with relatively simple ACLs.
- New DELETE BACKUP command does not yet work for subfile backup versions.

Open file support


- An "open" file is a file that is currently open by one or more applications.
- Often times "open" is confused with "locked".
 - TSM (even without OFS) can back up open files ...
 - ... as long as the app does not have exclusive access.
- OFS enables:
 - backup of locked files and
 - "crash-consistent" backup at the file system level.
- OFS operates at the file system (not file) level.

Open file support – How it works

- Uses the TSM Logical Volume Snapshot Agent (LVSA) to present the volume as a new read-only “file system” on a virtual volume.
- LVSA processes read requests to the virtual volume
 - If the requested data exists in a block in the Old Block File (OBF), then the read will be satisfied with data from the OBF.
 - If the requested data is not in the OBF, then the read will be satisfied with data from the real volume.
- TSM OFS support works by reading data from the virtual volume. Because this is a separate file system from the OS perspective, files locked on the real FS are not considered locked on the virtual volume.

Open file support – How it works (2)

- (1) Application issues a “write”.
- (2) LVSA intercepts write and reads old data blocks.
- (3) LVSA writes old data blocks to OBF.
- (4) New data is written to disk.
- (5) Application issues a “read”, LVSA reads data from disk.
- (6) LVSA passes data to application.



Open file support – Using OFS

- If selected during install, all backup and archive operations will use OFS.
 - `INCLUDE.FS x: FILELEVELTYPE=DYNAMIC` will disable OFS for volume x:
 - `INCLUDE.FS x: FILELEVELTYPE=SNAPSHOT` will enable OFS for volume x:
- OFS can also be enabled/disabled via the GUI setup wizard.
- Enabling or disabling OFS **will** require a reboot.
- If snapshot can not be taken due to volume being too busy, operation will fail over to non-OFS backup.

Open file support – Other OFS-related options

■ SNAPSHOTCACHELOCATION

- Specifies the location of the LVSA cache (OBF).
 - Default in 5.3: **x:\tsmlvsacache** on the volume being backed up.
 - Default before 5.3: **C:\tsmlvsacache**

■ SNAPSHOTCACHESIZE

- Specifies the percentage of used space that you expect to change during the backup. Values can range from 1 – 100 (%). Default is 1.
- Minimum cache file size is the greater of 100 MB or available space.

■ SNAPSHOTFSIDLEWAIT

- Specifies how long to wait (seconds or milliseconds) for no write activity to occur before the volume snapshot will be taken.
- A value of 0 is not recommended.

Open file support – Other OFS-related options (2)

- **SNAPSHOTFSIDLERETRIES**

- Specifies the number of times the snapshot attempt will occur.

- **PRESNAPSHOTCMD**

- Specifies an OS command that can be run prior to the snapshot occurring.

- Intended to allow quiescing of applications prior to the snapshot.

- **POSTSNAPSHOTCMD**

- Specifies an OS command that can be run after the snapshot occurs.

- Intended to allow restart of applications after the snapshot.

Open file support – Changes in version 5.3

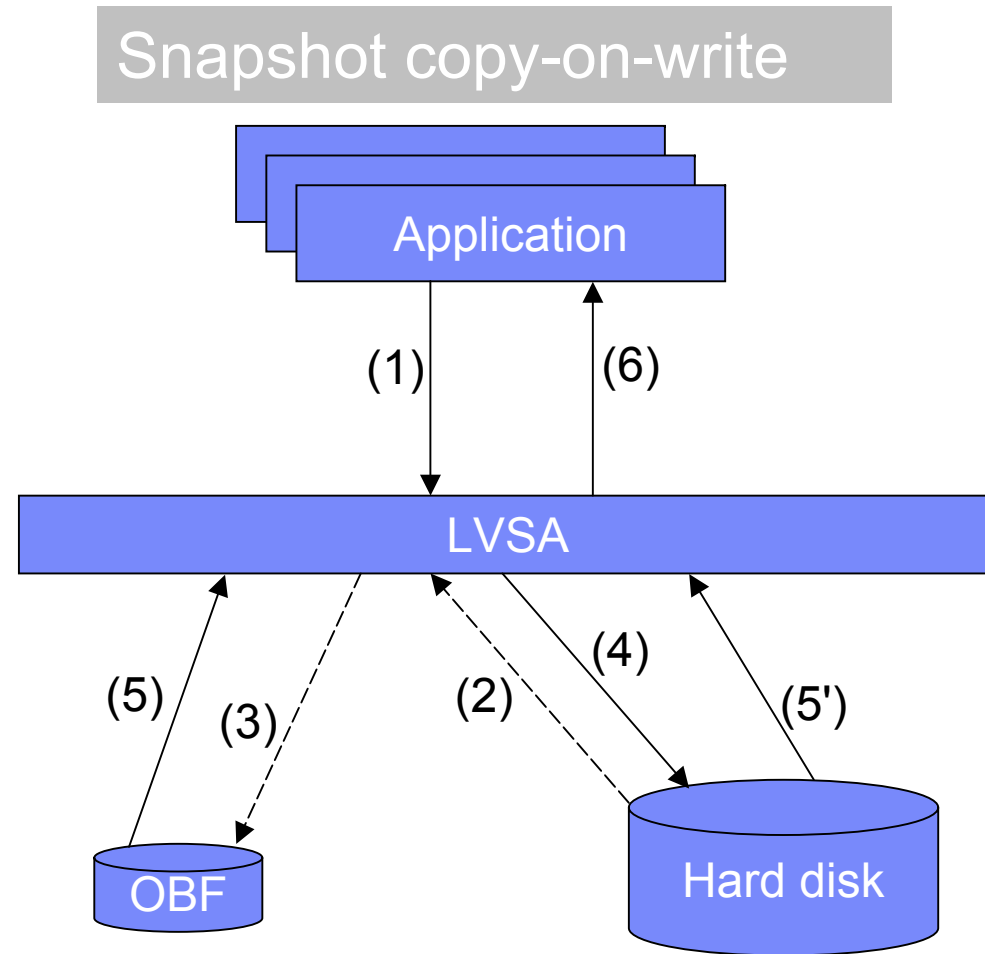
- Version 5.2:
 - LVSA cache (OBF) could not reside on the volume being backed up.
 - Default LVSA cache location was **C:\tsmlvsacache**.
 - Single-volume systems could not use OFS.
 - No support for Windows 2003
- Version 5.3:
 - Provide support for systems with only one drive.
 - New default LVSA cache location is **x:\tsmlvsacache** on the drive being backed up.
 - Currently target support for Windows 2003 (32-bit) in 5.3.2.

Online image backup

- "Image backup" is a single-object backup of a logical volume.
- For file systems containing many files, image backup might be faster than traditional full incremental backup.
- Likewise, image restore might be faster than file-by-file restore.
- Offline image backup prevents write access to the volume while it is being backed up.
- Online image backup allows complete access to the volume while it is being backed up.

Online image backup – Works similar to OFS

- (1) Application issues a “write”.
- (2) LVSA intercepts write and reads old data blocks.
- (3) LVSA writes old data blocks to OBF.
- (4) New data is written to disk.
- (5) Application issues a “read”, LVSA reads data from disk.
- (6) LVSA passes data to application.



Online image backup – Using online image backup

- LVSA must be configured, either during install or with the GUI Setup Wizard.
- BACKUP IMAGE command supports the IMAGETYPE option:
 - IMAGETYPE=STATIC – offline image backup
 - IMAGETYPE=SNAPSHOT – online image backup. This is the default.
- Administrative authority is required on the client system.
- The logical volume must have a drive letter or mount point.

Online image backup – Using online image backup (2)

■ Related client options:

➤ DOMAIN.IMAGE *fslist*

- Specifies which volumes are backed up if no volumes are listed in the BACKUP IMAGE command.
- Ignored if BACKUP IMAGE specifies one or more volumes.
- There is no default value.

➤ INCLUDE.IMAGE *mgmtclass options*

- Allows you to specify a management class for the image backup.
- Can also include IMAGETYPE, IMAGEGAPSIZE, SNAPSHOTCACHELOCATION, SNAPSHOTCACHESIZE, SNAPSHOTFSIDLERETRIES, SNAPSHOTFSIDLEWAIT, PRESNAPSHOTCMD, and POSTSNAPSHOTCMD.

➤ EXCLUDE.IMAGE *fs*

- Excludes a volume from image backup.
- Overrides DOMAIN.IMAGE and volumes specified with BACKUP IMAGE.

Online image backup – Using online image backup (3)

■ Related client options (2):

➤ SNAPSHOTCACHELOCATION

- Specifies the location of the LVSA cache (OBF).
 - Default in 5.3: **x:\tsmlvsacache** on the volume being backed up.
 - Default before 5.3: **C:\tsmlvsacache**

➤ SNAPSHOTCACHESIZE

- Specifies the percentage of used space that you expect to change during the backup. Values can range from 1 – 100 (%). Default is 1.
- Minimum cache file size is the greater of 100 MB or available space.

➤ SNAPSHOTFSIDLEWAIT

- Specifies how long to wait (seconds or milliseconds) for no write activity to occur before the volume snapshot will be taken.
- A value of 0 is not recommended.

➤ SNAPSHOTFSIDLERETRIES

- Specifies the number of times the snapshot attempt will occur.

Online image backup – Using online image backup (4)

- Related client options (3):

- PRESNAPSHOTCMD

- Specifies an OS command that can be run prior to the snapshot occurring.
- Intended to allow quiescing of applications prior to the snapshot.

- POSTSNAPSHOTCMD

- Specifies an OS command that can be run after the snapshot occurs.
- Intended to allow restart of applications after the snapshot.

- COMPRESSION

- Compress the data that is sent to the server.

- COMPRESSALWAYS

- Send the "compressed" data, even if compression yields larger data.

Online image backup – Using online image backup (5)

- Related client options (4):

- IMAGEGAPSIZE

- Specifies the minimum size of empty regions on the disk to skip during backup.

- MODE

- SELECTIVE – Specifies you want to do a full image backup. This is the default.
 - INCREMENTAL - Specifies you want to back up only new or changed files since the last full image backup.

Online image backup – Using online image backup (6)

- Using image backup with incremental backup (method 1):
 1. Full incremental backup of the file system.
 2. Image backup of the same file system.
 3. Periodic full incremental backup of the file system.
- To restore, use **RESTORE IMAGE** with **–INCREMENTAL** and **–DELETEFILES**.
 - This will first restore the image backup.
 - Next, files that have since been inactivated on the server will be deleted.
 - Finally, new and changed files are restored from the incrementals.

Online image backup – Using online image backup (7)

- Using image backup with incremental backup (method 2):
 1. Run an image backup of the file system.
 2. Run image backup with `–MODE=INCREMENTAL`
 3. Run periodic full image backups of the file system.
- To restore, use **RESTORE IMAGE** with **–INCREMENTAL**.
 - This will first restore the image backup.
 - New and changed files are restored from the incrementals.
 - Files that were deleted since the full image backup will also be restored.

Journal-based backup

- The purpose of JBB is to improve performance of incremental backups.
- Bottlenecks eliminated by JBB:
 - Initial query to the TSM server for current backup inventory, which can take a significant amount of time
 - the need to traverse the file system to look for changed files
- Most suitable for file systems with a large number of files, but relatively little change.

Journal-based backup – How it works

- File system monitoring
 - The journal service starts a monitoring thread for each file system that is monitored.
 - Uses the Windows API function ReadDirectoryChangesW(), not Windows built-in journaling function.
 - Each monitored file system has its own database
 - Each thread updates its file system's database
 - Windows caches file system updates, so the notifications might not be synchronous.

Journal-based backup – How it works (2)

■ Backup

- Client uses named pipes to communicate with the journal service.
- Client checks with the journal service to validate the journal database integrity.
 - Integrity failure will require a full incremental backup.
- Client checks consistency of file space on the server with the journal service.
 - Consistency failure will require a full incremental backup.
- Client queries the journal service for files to back up.
- Files are backed up.
- Journal database entries are deleted.
- Note: copy group FREQUENCY option is ignored during JBB.

Journal-based backup – Whither full incremental?

- When is a full incremental backup required?
 - Journal service is not running
 - Journal service is stopped and restarted
 - **PreserverDBOnExit=1** in the journal service configuration file removes this restriction
 - Should **only** be used in the case of a system reboot
 - Journal is corrupt
 - Policy changes on the server
 - Client node name has changed
 - File space on the server has been deleted

Journal-based backup – What's new in 5.3

- Multi-session support either with **RESOURCEUTILIZATION** or multiple client instances.
 - Multiple concurrent journal-based backups of the same file system not recommended.
- ISAM database replaced with a B-tree database
 - Removes the 2 GB database file size limitation
 - Disk space can be recovered when journal entries are deleted
- Reduced frequency of notification buffer overflow conditions
- Support for multiple journal services using the **JOURNALPIPE** option
 - JOURNALPIPE identifies a unique named pipe name on which the client can connect to a specific journal service.

Cluster Support

- Support for Microsoft Cluster Server and Veritas Cluster Server (TSM 5.2.2 for VCS).
- Enabled with the `CLUSTERNODE=YES` option.
 - Only cluster resources can be backed up.
 - Node name defaults to the cluster name.
- Starting in TSM 5.3, when coupled with `CLUSTERDISKONLY=NO`, cluster resources accessed via volume mount points can also be backed up.
 - Earlier clients were limited to backing up cluster resources that were accessed only via drive letter.

Windows system state

- "System state", as defined by Microsoft, refers to:
 - Boot and system files
 - Active Directory
 - SysVol
 - Certificate Server information
 - Cluster database
 - Registry
 - COM+ class registration database
- System state should be treated as a single entity for backup and restore purposes.

Windows system services

- "System services" as defined by Microsoft, refers to:
 - Event logs
 - RSM
 - Remote Storage Service
 - FRS
 - Terminal Server Licensing
 - WMI
 - IIS Metabase
- System services can be treated individually.
- For Windows NT systems prior to Windows 2003, TSM used the term "system object" to refer to both system state and system services.

TSM support for Windows system state

- Over the years there has been a proliferation of specialized APIs for backing up different Windows components, each with their own unique interfaces.
- Beginning with Windows XP, Microsoft introduced their volume snapshot technology, "Volume Snapshot Services" (VSS).
- VSS is a simple, consistent backup/restore framework that can be used by, and for, all system components. VSS is also included in Windows 2003.
- VSS-enabled components include "writers" that coordinate with VSS to make sure that their component-specific data is consistent (for backup).

TSM support for Windows system state (2)

- The backup software ("requester") communicates with VSS, which coordinates the snapshot activity.
- Starting with version 5.2, TSM uses VSS to back up system state and system services for Windows 2003.
- Although VSS is also available for XP, a complete set of writers for XP is not available. Therefore TSM uses the legacy "system object-based" backup method for Windows XP.
- Since VSS is not available at all for Windows 2000, the legacy backup method will be used.

TSM support for Windows system state (3)

- Windows 2000, XP:
 - BACKUP SYSTEMOBJECT
 - Backs up the system state and system services.
 - SYSTEMOBJECT is part of the ALL-LOCAL domain.
 - INCLUDE.SYSTEMOBJECT <*mgmtclassname*>
 - Binds the system object (all components) to the specified management class.
 - EXCLUDE.SYSTEMOBJECT <*objectname*>
 - Excludes the specified object name from backup.
 - Valid object names include ACTIVEDIRECTORY, CERTSERVERDB, CLUSTERDB, COMPLUSDB, EVENTLOG, FRS, REGISTRY, RSM, SYSFILES, SYSVOL, and WMI.

TSM support for Windows system state (4)

- Windows 2003:
 - BACKUP SYSTEMSTATE and BACKUP SYSTEMSERVICES
 - Backs up system state and system services, respectively.
 - SYSTEMSTATE and SYSTEMSERVICES are part of the ALL-LOCAL domain.
 - INCLUDE.SYSTEMSTATE <*mgmtclassname*>
 - Binds the system state and system service to the specified management class.
 - EXCLUDE.SYSTEMSERVICE <*servicename*>
 - Excludes the specified system service from backup.
 - Valid service names include BITS, EVENTLOG, RSM, CLUSTERDB, RSS, TLS, WMI, IIS, DHCP, WINSDB

Questions?