



Oxford University TSM Symposium 2005

Why Does It Work That Way?

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

Agenda

- This session is intended to be a peek “under the hood” of how and why some TSM features work the way they work.

EVENTS table only shows results from today!

- Why does `SELECT * FROM EVENTS` show results only from today's events?
- The EVENTS table ...
 - ... is calculated at runtime.
 - ... is theoretically unbounded since the server can calculate records even for past events where no event record exists, and future events that have not yet occurred.
 - ... defaults to same info as `QUERY EVENTS`.

EVENTS table only shows results from today! (2)

- Example – the following returns either no events, or worse, just events from today:

```
select node_name, status, scheduled_start
  from events
 where scheduled_start >= current_timestamp - 7 days and
        (status='Failed' or status='Missed')
```

EVENTS table only shows results from today! (3)

- Solution – Force the server to look at all available event records with a hard-coded date:

```
select node_name, status, scheduled_start
  from events
 where scheduled_start>='2005-01-01' and
        scheduled_start>=current_timestamp - 7 days and
        (status='Failed' or status='Missed')
```

Why are my directories being backed up to tape?

- I have my system configured to back everything up to a management class that goes only to disk:

```
include *:\...\* to_disk
```

When my backup runs, why am I seeing tapes being mounted for my client node?

Why are my directories being backed up to tape? (2)

- Possible reasons for this:
 - The INCLUDE statement is being overridden by another INCLUDE statement in a client options set.
 - The disk pool is full, so data is being sent to the next storage pool in the hierarchy.
- The most common reason has to do with how directories are bound to management classes.

Why are my directories being backed up to tape? (3)

- By default, directories are bound to the management class with the highest RETONLY setting.
- If multiple management classes have the same highest RETONLY setting, the class whose name is highest in the operating systems native collating (sort) order is selected.

Why are my directories being backed up to tape? (4)

- How to check for this situation:

```
select class_name, retonly, destination
  from bu_copygroups
  where domain_name='STANDARD' and set_name='ACTIVE'
  order by retonly desc, class_name desc
```

CLASS_NAME	REONLY	DESTINATION
-----	-----	-----
TO_TAPE	90	FILPOOL1
TO_DISK	90	DISKPOOL
STANDARD	5	DISKPOOL

- You can use DIRMC to specify the TO_DISK pool.

Why is TSM not backing up a changed file?

- My file has clearly changed, so why does incremental backup ignore the file?
- One possible reason is that the file is being excluded:
 - Local INCLUDE/EXCLUDE list inadvertently excludes the file
 - EXCLUDE in the client option set is excluding the file
 - Some files are implicitly excluded by the system

Why is TSM not backing up a changed file? (2)

- Check the INCLUDE/EXCLUDE list:

```
C:\TSM\baclient>dsmc query inclexcl -detail
...
*** FILE INCLUDE/EXCLUDE ***
Mode Function  Pattern (match from top down)  Source File
-----
Excl Directory C:\adsm.sys          Operating System
Exclude All      c:\mytools\...\*                Server
Exclude All    C:\WINDOWS\usrclass.dat.log     Operating System
Exclude All    C:\WINDOWS\usrclass.dat         Operating System
Exclude All    C:\WINDOWS\system32\wbem\Repository\...\* Operating System
...
Exclude All    *\Pagefile.sys                  Operating System
Exclude All    *\hiberfil.sys                   Operating System
Exclude All    *\...\*.crmlog                   Operating System
Include All      *\...\*                          dsm.opt
MC NAME: TO_DISK
No DFS include/exclude statements defined.
```

Why is TSM not backing up a changed file? (3)

- Another unobvious reason has to do with the **FREQUENCY** setting of the management class's copy group:

```
select domain_name, class_name, frequency
  from bu_copygroups
  where set_name='ACTIVE' and frequency>0
```

DOMAIN_NAME	CLASS_NAME	FREQUENCY
STANDARD	STANDARD	5
STANDARD	TO_DISK	7

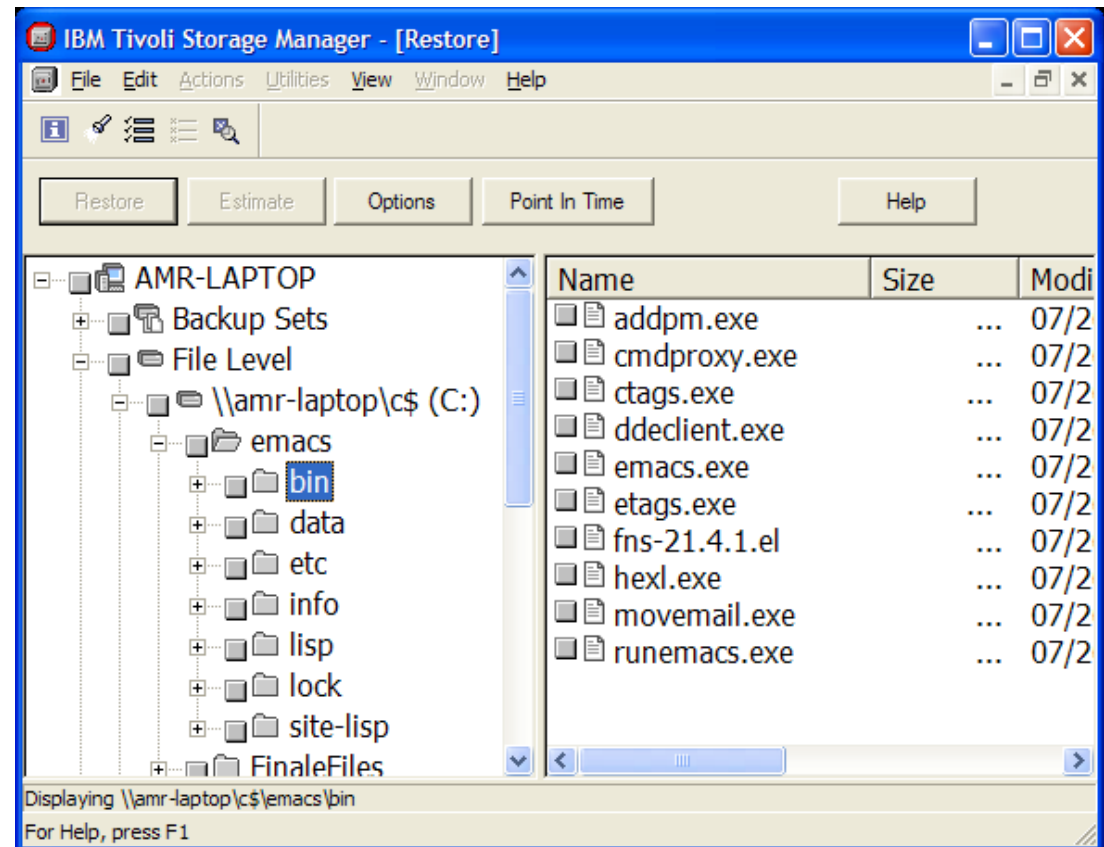
- Next, use **dsmt query backup** to identify when the file was last backed up and the management class to which it is bound.

GUI restore does not show all available files

- Why does a GUI restore not show all available files? I know they are there, because a CLI restore indicates the files are there!
- In ADSM versions 1 and 2, the restore GUI would initially query the server for the node's entire inventory. This could take a very long time and consume lots of memory.
- Starting in ADSM version 3, the GUI was redesigned to mitigate these problems.

GUI restore does not show all available files (2)

- Only the objects immediately below the selected directory are queried.
- Because only a small handful of objects are queried at a time, the GUI appears very responsive.



GUI restore does not show all available files (3)

- Andy's rule #4: There is no such thing as a free lunch!
- You cannot use the GUI to navigate to files when the parent directory does not exist.
- This often happens when attempting point-in-time restores.
 - Depending on copy group attributes, it is possible to have file objects that are older than directory objects.
 - If the point-in-time predates the oldest directory object, it will not appear in the GUI restore tree.

GUI restore does not show all available files (4)

- Alternative: use the CLI, which does not have this limitation.
 - A wise man once said, “The command line is your friend.”
- Consider implementing your restore service to provide for restore of files up to n days ago. For example, if $n = 30$:

```
VEREXISTS=NOLIMIT  
VERDELETED=NOLIMIT ( * )  
RETEXTRA=30  
RETONLY=30 ( * )
```

(*) Your storage service for deleted files might differ.

Why can I archive Unix file, but not back it up?

- If you do not own the file, you can not back it up.
- If you have read access, you can archive it.
- Explanation:
 - If the [non-root] user could back up files he doesn't own, then in theory, he could back up another user's files repeatedly, causing older versions to be expired. Thus it was decided that non-root users would not be able to back up files they do not own.
 - Because archive is not subject to versioning, it was decided that non-root users could archive files they do not own.

Last backup dates not updated after backup?

- The following file space properties are updated only after completion of a **full** incremental backup:
 - Last Backup Start Date/Time
 - Days Since Last Backup Started
 - Last Backup Completion Date/Time
 - Days Since Last Backup Completed
- Selective backup, incremental-by-date, or partially qualified backups do not count.
 - `dsmc i c:\ -subdir=yes`
 - `dsmc i c:\myprograms\ -subdir=yes`
 - `dsmc s c:\ -subdir=yes`

Last backup dates not updated after backup? (2)

- Explanation:

- ADSM's full incremental backup method was intended as the primary method of backup, and the only means of ensuring that the file system was completely backed up. The "last backup" statistics reflect this philosophy.
- Other backup methods – incremental by date, partially-qualified incremental, and selective backups do not guarantee a comprehensive backup, and therefore the "last backup" statistics are not updated.

Why are Windows and NetWare ACLs backed up to tape?

- UNIX permissions are backed up to the TSM server database. Why are Windows and NetWare permissions backed up to storage pools?
- Explanation:
 - UNIX mode maintains the permissions as part of the file attributes.
 - On Windows and NetWare, permissions/ACL data are not part of the file attributes, but are comprised of larger pieces of data.
 - File attributes, regardless of operation system, are stored in the TSM server database.
 - On Windows and NetWare, the permissions/ACL data can vary in length, but are larger than file attributes. Therefore they are stored with the data in storage pools.

In what order are client options evaluated?

- There are many ways of specifying client options:
 - Accept the default
 - Options file
 - Operating system command line
 - TSM interactive command line
 - Client option sets
 - Schedule OPTIONS setting
- What is the order in which the options are compiled?

In what order are client options evaluated? (2)

1. Default (option not specified anywhere)
2. Options file
3. OS command prompt
4. Interactive command line interface or corresponding GUI selection
5. Client options sets
 - FORCE=YES – Overrides the 2 – 4 above
 - FORCE=NO – Does not override anything except the default
6. Schedule OPTIONS setting
 - Overrides 1 – 5 unless client option set has FORCE=YES

In what order are client options evaluated? (3)

- The client option set FORCE option does not affect additive options such as INCLUDE, EXCLUDE, or DOMAIN.
- INCLUDE and EXCLUDE statements defined in a client option set are always evaluated before any INCLUDE and EXCLUDE statements defined by the client
 - ... but otherwise following INCLUDE/EXCLUDE processing rules.

Questions?