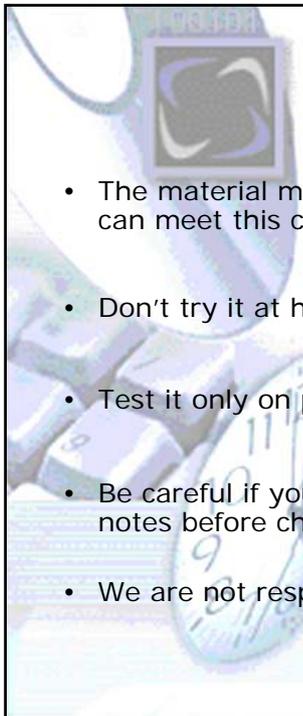




Performance

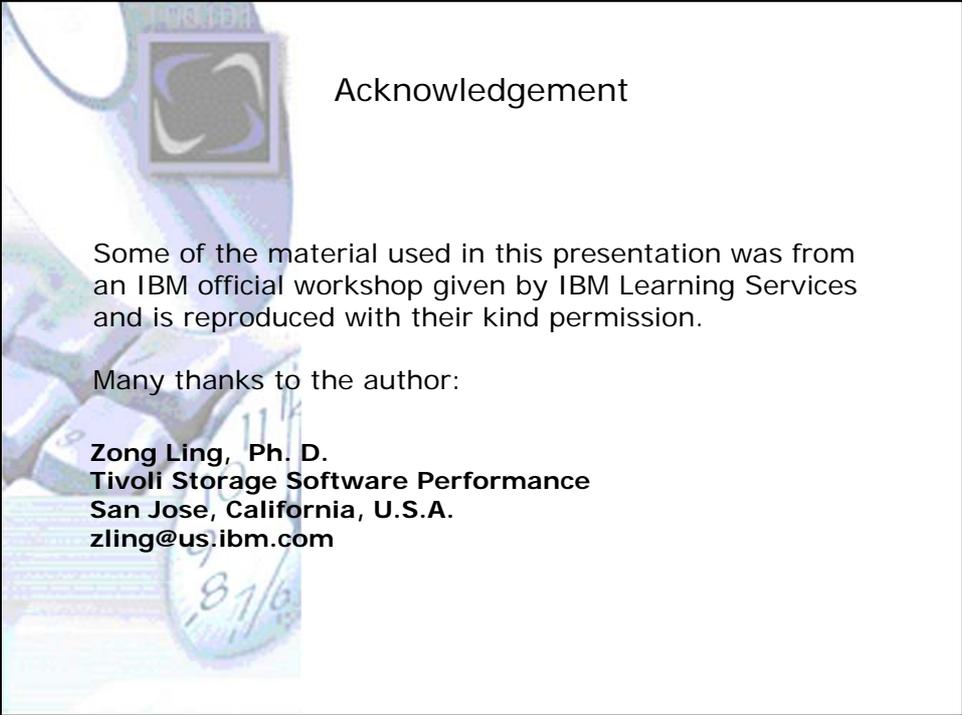
Tips and tricks

Peter Pijpelink
peter.pijpelink@plcs.nl
Rev.03



Disclaimer & Warning

- The material may seem complex or heavy with detail. You can meet this challenge by utilizing the following guidelines.
- Don't try it at home 😊
- Test it only on production systems 😊
- Be careful if you like to test something, always make notes before changing a setting or value
- We are not responsible for any problems or disruptions

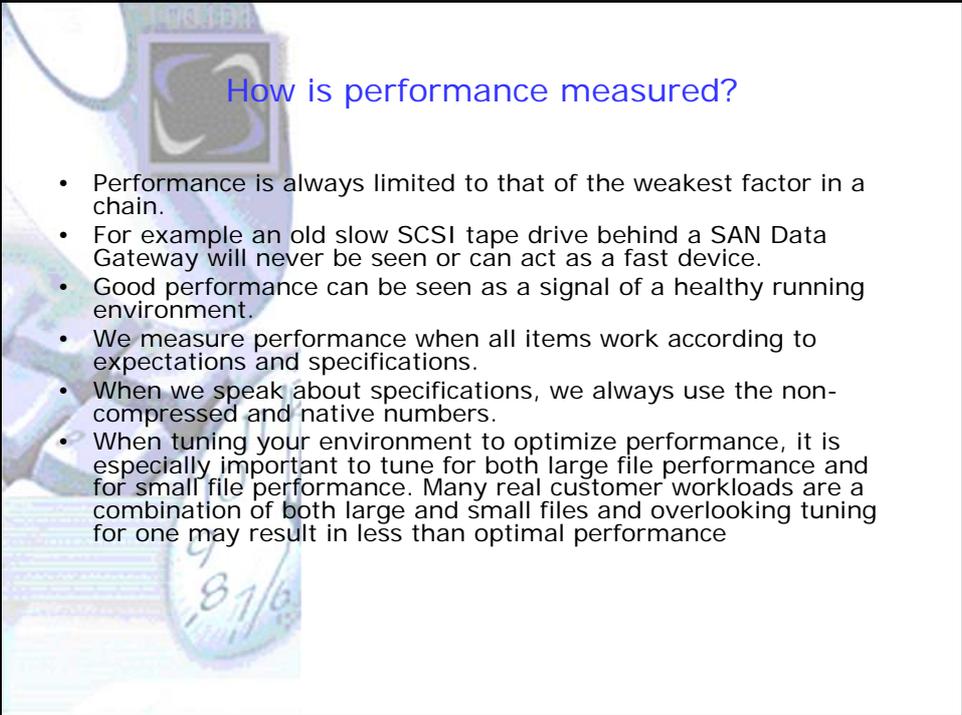


Acknowledgement

Some of the material used in this presentation was from an IBM official workshop given by IBM Learning Services and is reproduced with their kind permission.

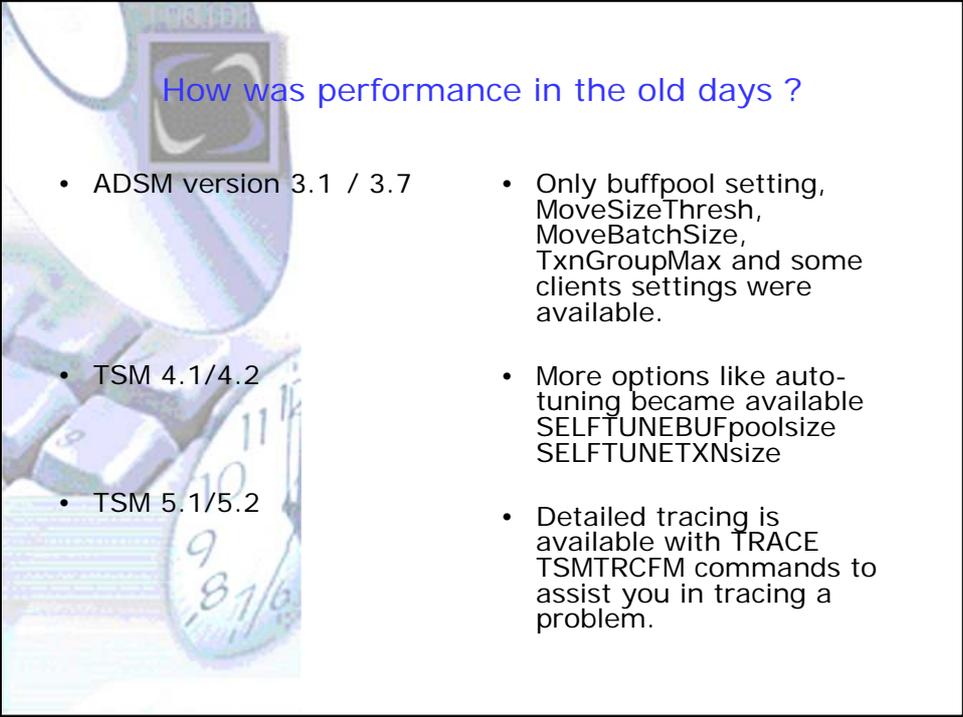
Many thanks to the author:

Zong Ling, Ph. D.
Tivoli Storage Software Performance
San Jose, California, U.S.A.
zling@us.ibm.com



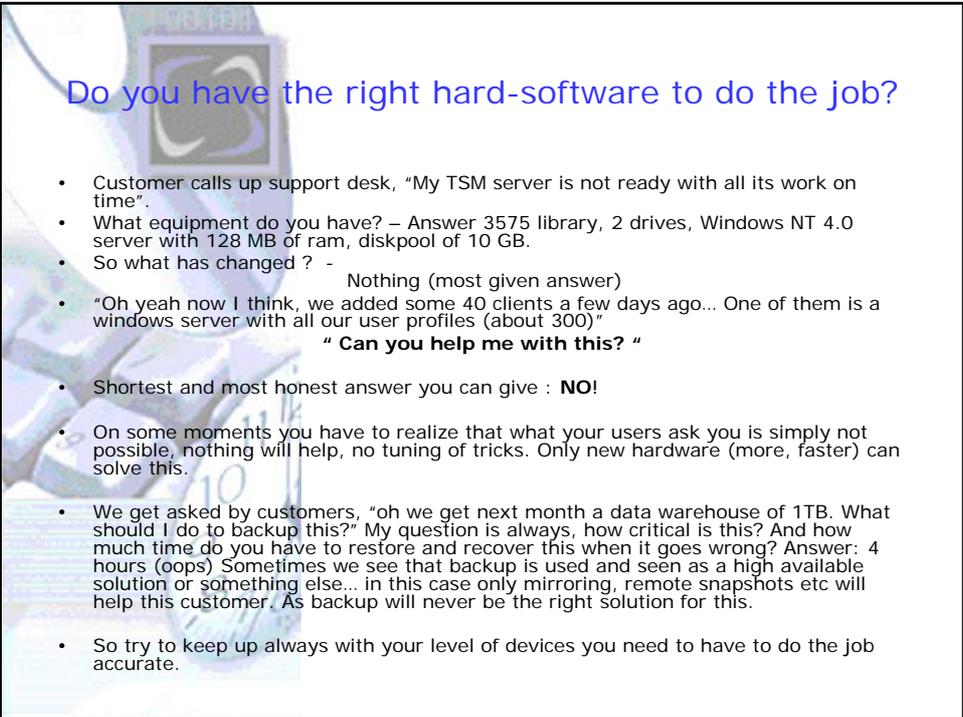
How is performance measured?

- Performance is always limited to that of the weakest factor in a chain.
- For example an old slow SCSI tape drive behind a SAN Data Gateway will never be seen or can act as a fast device.
- Good performance can be seen as a signal of a healthy running environment.
- We measure performance when all items work according to expectations and specifications.
- When we speak about specifications, we always use the non-compressed and native numbers.
- When tuning your environment to optimize performance, it is especially important to tune for both large file performance and for small file performance. Many real customer workloads are a combination of both large and small files and overlooking tuning for one may result in less than optimal performance



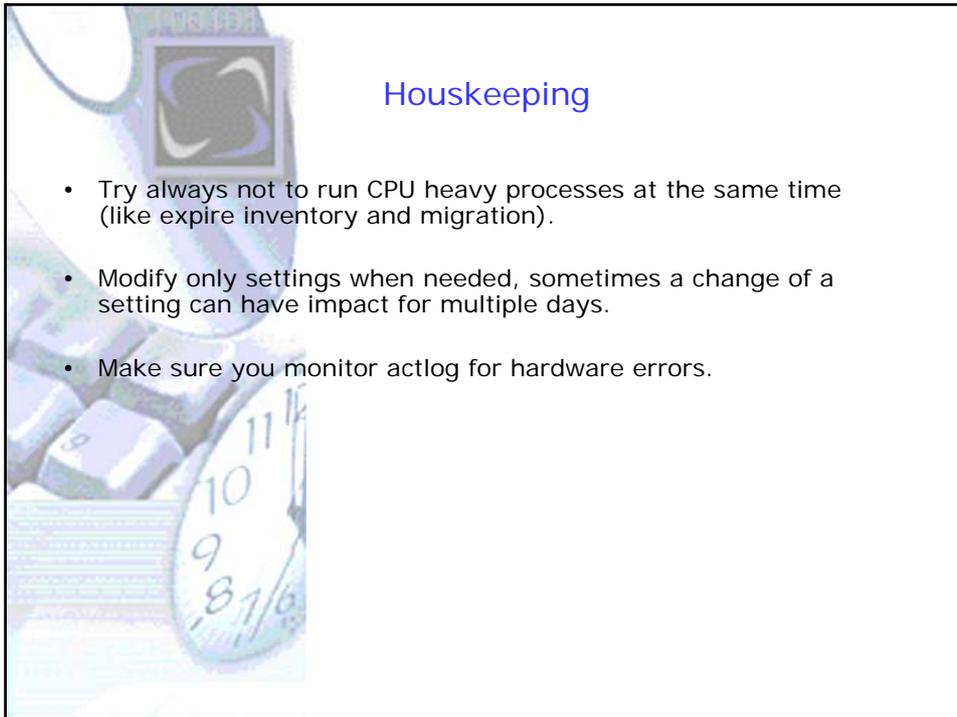
How was performance in the old days ?

- ADSM version 3.1 / 3.7
- TSM 4.1/4.2
- TSM 5.1/5.2
- Only buffpool setting, MoveSizeThresh, MoveBatchSize, TxnGroupMax and some clients settings were available.
- More options like auto-tuning became available SELFTUNEBUFpoolsize SELFTUNETXNsize
- Detailed tracing is available with TRACE TSMTRCFM commands to assist you in tracing a problem.



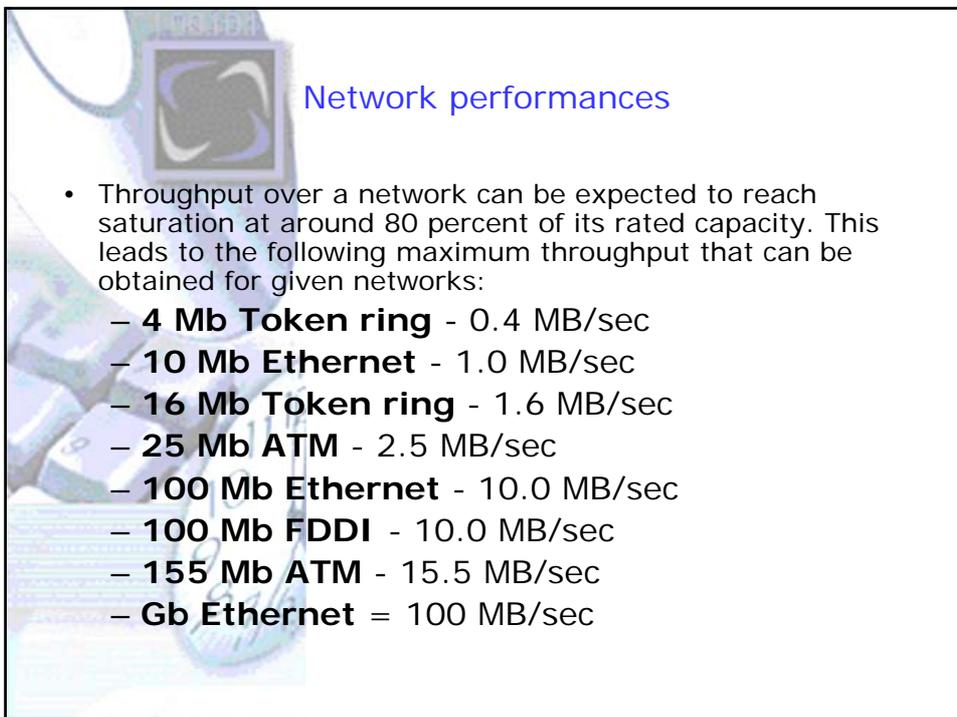
Do you have the right hard-software to do the job?

- Customer calls up support desk, "My TSM server is not ready with all its work on time".
- What equipment do you have? – Answer 3575 library, 2 drives, Windows NT 4.0 server with 128 MB of ram, diskpool of 10 GB.
- So what has changed ? - Nothing (most given answer)
- "Oh yeah now I think, we added some 40 clients a few days ago... One of them is a windows server with all our user profiles (about 300)"
" Can you help me with this? "
- Shortest and most honest answer you can give : **NO!**
- On some moments you have to realize that what your users ask you is simply not possible, nothing will help, no tuning of tricks. Only new hardware (more, faster) can solve this.
- We get asked by customers, "oh we get next month a data warehouse of 1TB. What should I do to backup this?" My question is always, how critical is this? And how much time do you have to restore and recover this when it goes wrong? Answer: 4 hours (oops) Sometimes we see that backup is used and seen as a high available solution or something else... in this case only mirroring, remote snapshots etc will help this customer. As backup will never be the right solution for this.
- So try to keep up always with your level of devices you need to have to do the job accurate.



Houskeeping

- Try always not to run CPU heavy processes at the same time (like expire inventory and migration).
- Modify only settings when needed, sometimes a change of a setting can have impact for multiple days.
- Make sure you monitor actlog for hardware errors.

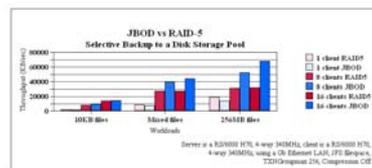
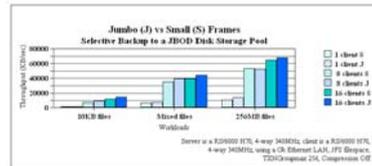


Network performances

- Throughput over a network can be expected to reach saturation at around 80 percent of its rated capacity. This leads to the following maximum throughput that can be obtained for given networks:
 - **4 Mb Token ring** - 0.4 MB/sec
 - **10 Mb Ethernet** - 1.0 MB/sec
 - **16 Mb Token ring** - 1.6 MB/sec
 - **25 Mb ATM** - 2.5 MB/sec
 - **100 Mb Ethernet** - 10.0 MB/sec
 - **100 Mb FDDI** - 10.0 MB/sec
 - **155 Mb ATM** - 15.5 MB/sec
 - **Gb Ethernet** = 100 MB/sec

RAID or not using RAID?

- Do not use RAID5 for your database and log volumes. Rather use TSM mirrors over more disks then placing your database and recovery-log volumes on RAID5.



Is more better?

- If you set up your storage hierarchies try to make more than one disk and tape pool. When using more pools you can have more actions in parallel, so 2 processes for reclaim on copy pools rather than one process for each pool. With more pools you can also separate real critical data, and backup that pool to copy pool for DR purposes. When it comes to critical issues you only need to make sure that pool is restored also.
- More tape drives is great but when using one single volume for your disk pool is not a good idea.
- When setting up your environment try to spread your disk volumes over more physical disks and always make sure you have got a mirror of a log or database volume at the same disk set (or controller).
- The perfect definition is when you have at least enough free disk space at the size of your database and 50% of your log, in this configuration you can always make a database backup to disk, or extend your database or recovery log.

LAN or SAN based?

- LAN based backup is great for multiple small files and lots of transactions.
- Gigabyte LAN based networks can be very efficient even if your network is using fast Ethernet (100mbit), when you can setup your switch so that it has a GB connection with your TSM server, it is capable to use the backbone of the switch, so a lot of client connections can use the performance since the switch is capable of delivering lots of packets over the Gigabyte link to the TSM server. Cost of setting this up are minimal, only one GB adapter and a module in your switch is needed.
- SAN based backups are not made for transferring lots of small files. When using LANFREE make sure you have setup separate adapters for disk and tape movements. IO to disk or IO to tape is totally different, configure your switches and SAN zones so that you have optimal use of your high speed connection.
- When using a Fibre Channel adapter to write direct to disk, try to setup like Exchange or Oracle so that large files only goes direct to tape, small files drops the performance a lot! Use multiple managementclasses and well defined include/excludes.
- In practical using a SAN for backup is only good if your data is good for using over a FC link, so small files direct to tape is very bad, using large data blocks works perfect.

Performance Challenges

Tivoli Storage Manager Environment

Client



- Average file size
- Processor speed
- Memory
- Disk subsystem
- Operating system
- Filesystem
- Compression
- Encryption
- Other applications

Server



- Bandwidth
- Utilization
- Protocol
- Routers/bridges
- Tuning

Network



- Processor speed
- Memory
- Storage devices
- Configuration
- Operating system
- Number of client sessions
- Other applications

Server Configuration - Hardware Sizing

- TSM server sizing exercise can determine the required I/O load
 - Add the amount of data backed up/archived for each TSM client per day during the **peak** backup window
 - Divide by the length of the **peak** backup window
- Provide enough CPU and I/O bandwidth for **peak** load
- Using LAN-free or server-free clients reduces the I/O load on the TSM server
- Consult standard performance benchmarks for system capability
- Collect trend data on server resource usage

DB2 data	300 GB
Mail server data	250 GB
Accounting data	50 GB

	600 GB in 4 hours = 43 MB/sec

Server Configuration - Memory

- Largest use of memory in TSM server is for the database buffer pool
- A large buffer pool can improve database performance
 - Fewer database disk I/Os
- Provide enough real memory for optimal buffer pool
- Set **bufpoolsz** to between 1/8 and 1/2 of real memory
 - Example: **bufpoolsz 131072** for a server with 1 GB memory

Server Configuration - I/O Devices

- Balance I/O adapters and devices across multiple busses
- Adapter placement can impact performance significantly
- Put network adapters on different PCI bus than storage adapters
- Only one or two fast tape drives per SCSI adapter



Server Configuration - Networks

- Use dedicated networks for backup (LAN or SAN)
- Set Ethernet speed and duplex settings
 - Don't rely on Auto Detect
- Gb Ethernet jumbo frames (9000 bytes)
 - Only useful if supported on client, server, and switch
 - Not all Gb Ethernet hardware supports jumbo frames
 - Improved throughput with lower host CPU usage
- Use TSM options to optimize network throughput
 - **`tcpwindowsize 63`**
 - **`tcpbuffsize 32`**
 - **`tcpnodelay yes`**



Server Configuration - Tape

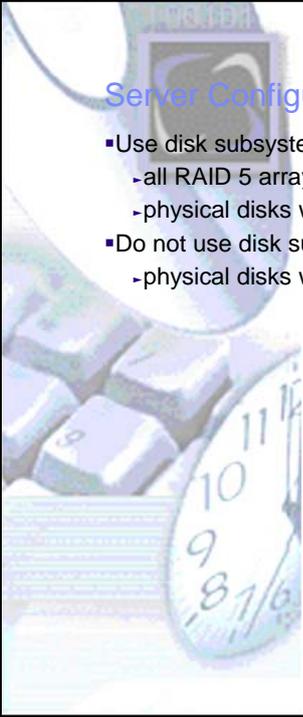
- Configure enough tape drives for:
 - The maximum number of TSM client sessions backing up direct to tape at any time during the **peak** backup window
 - Add additional tape drives for other functions that run during the backup window
 - ✓ Storage pool migration
 - ✓ Storage pool backup
 - ✓ Reclamation



Server Configuration - Disk

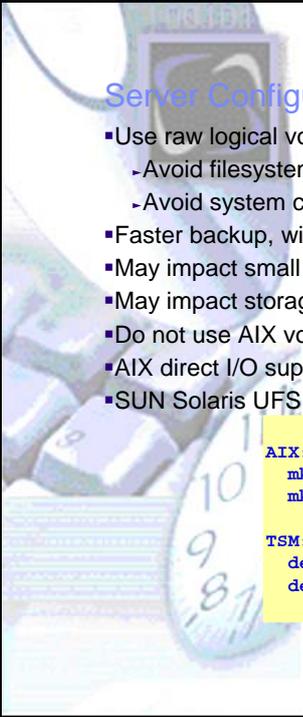
- TSM server can effectively use lots of disks
- Physical disks (JBOD) or RAID arrays?
 - RAID will require many more disks for equivalent performance
 - Write throughput is important during backup/archive
- Dedicate physical disk/array to one TSM volume (at most two)
 - recovery log
 - database
 - disk storage pool
- Place TSM volumes at outside diameter of physical disk
 - better sequential throughput
 - faster seek time





Server Configuration - Disk Write Cache

- Use disk subsystem/adaptor write cache for:
 - all RAID 5 arrays
 - physical disks with TSM database volumes (random I/O)
- Do not use disk subsystem/adaptor write cache for:
 - physical disks with TSM storage pool volumes (sequential I/O)



Server Configuration - UNIX - Filesystems

- Use raw logical volumes for best TSM server performance
 - Avoid filesystem overheads
 - Avoid system cache overhead
- Faster backup, with less CPU usage
- May impact small file restore from disk performance
- May impact storage pool migration performance
- Do not use AIX volume mirroring
- AIX direct I/O support is limited (<2GB volumes)
- SUN Solaris UFS with directio can help much

AIX:

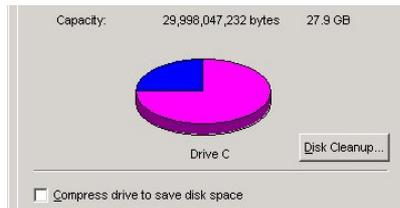
```
mklv -ae -t tsmdb -y db1 volgrp1 64 hdisk4  
mklv -ae -t tsmlg -y lg1 volgrp1 64 hdisk5
```

TSM:

```
define dbvol /dev/rdb1  
define logvol /dev/rlg1
```

Server Configuration - Windows - Filesystems

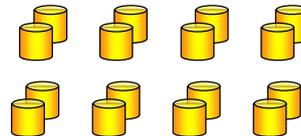
- Use NTFS for TSM server volumes
 - larger partitions
 - better data recovery
 - better security
 - faster storage pool volume formatting
- Do not use NTFS compression on TSM server volumes



Server Configuration - TSM Disk Volumes

- Use multiple TSM volumes for improving performance
- Examples:
 - Small TSM server:
 - ✓ 1 database volume*
 - ✓ 1 recovery log volume*
 - ✓ 4 disk storage pool volumes
 - Large TSM server:
 - ✓ 16 database volumes*
 - ✓ 1 recovery log volume*
 - ✓ 32 disk storage pool volumes

* double if mirrored

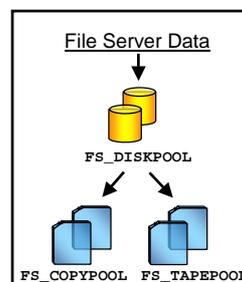
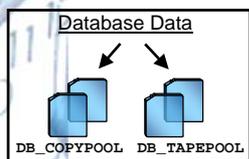


Server Configuration - Recovery Log / Database Mirroring

- Increased server availability
- Can impact performance for backup and restore of small files
- Impact can be mitigated by
 - Use ***mirrorwrite db parallel*** and ***dbpageshadow yes***
 - ✓ Page shadow file can be placed in server install directory
 - Use multiple database volumes
 - Use separate physical disks/arrays for each log or database volume
 - Use disk subsystem/adaptor write cache
 - Use optimal database bufferpool size

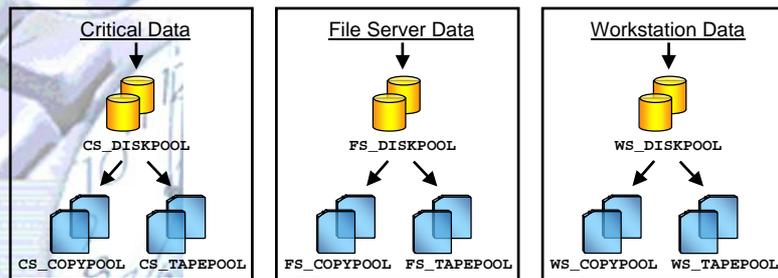
Server Configuration - Disk Storage Pool

- Back up large databases direct to tape
 - Throughput may be higher than backup to disk
 - Avoid extra data movement
- Back up file servers and workstations to disk
 - Avoid tape mount delays
 - Allows more sessions than tape drives



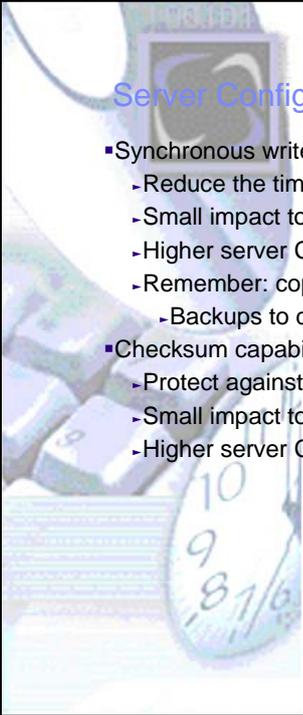
Server Configuration - Storage Hierarchy

- Use multiple storage hierarchies
- Storage pool migration
- Storage pool backup



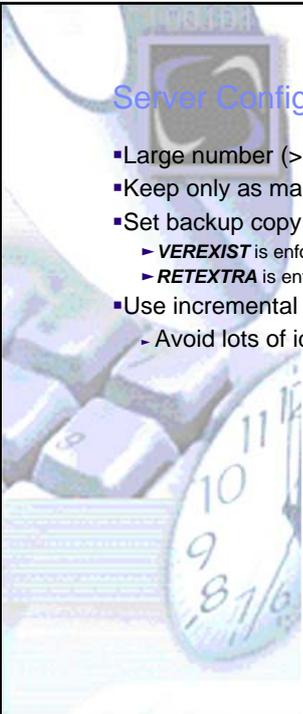
Server Configuration - Storage Pool Features

- Collocation
 - Each volume contains files for a single node or filespace
 - Sequential storage pools only
 - Can improve throughput of large restore by reducing mounts
 - Avoids tape contention during multi-client restores
 - Keeping database data separate allows faster tape reclamation
 - Can work against multi-session restore (too few volumes)
 - Requires larger number of tape volumes
 - Not practical for copy pools where volumes will be taken offsite



Server Configuration - Storage Pool Features ...

- Synchronous writes to copy storage pools
 - Reduce the time for backup stgpool operations
 - Small impact to backup throughput
 - Higher server CPU requirements during backup
 - Remember: copy pool must be a sequential storage pool
 - Backups to disk might also need a tape
- Checksum capability (CRC) for storage pool data movement
 - Protect against data corruption
 - Small impact to throughput
 - Higher server CPU requirements



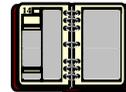
Server Configuration - Versioning

- Large number (>50) of backup versions can impact performance
- Keep only as many versions as necessary
- Set backup copy group **verexists=nolimit**
 - ▶ **VEREXIST** is enforced during backup processing (immediate expiration).
 - ▶ **RETEXTRA** is enforced during expiration.
- Use incremental backup instead of selective backup
 - Avoid lots of identical file versions



Server Configuration - Scheduling

- Schedule backups to spread the load
 - Use **set randomize percent**
- Dedicate backup window for critical servers
- Schedule other functions outside backup window
 - storage pool backup
 - storage pool migration
 - inventory expiration
 - ✓ Use **expinterval 0** (not automatically execute expiration) and define an administrative schedule
 - ✓ Use polling rather than prompted to randomize your load.



Client Configuration - Hardware Sizing

- Provide sufficient processor capacity and memory
- Consider backup and restore requirements in client sizing
 - TSM provides a variety of backup/restore functions and options
- Collect trend data on client resource usage



Client Configuration - Memory

- TSM client may require lots of memory for incremental backup
 - ~300 bytes per file/directory in a filesystem
 - So ... 300 MB for filesystem with 1 million files
- Can reduce memory requirements by using:
 - **include / exclude** to backup only what is needed
 - Use as much client option sets as you can, make a base in a script and use this to generate new option sets.
 - incremental by date (some files might be missed)
 - **memoryefficientbackup yes** (slower)
 - Use multiple filesystems
 - **virtualmountpoint** (UNIX only)
 - journal-based backup (Windows only)



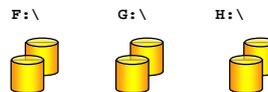
Client Configuration - Networks

- Use dedicated networks for backup (LAN or SAN)
- Set Ethernet speed and duplex settings
 - Don't rely on Auto Detect
- Gb Ethernet jumbo frames (9000 bytes)
 - Only useful if supported on client, server, and switch
 - Not all Gb Ethernet hardware supports jumbo frames (it is not a official standard)
 - Improved throughput with lower host CPU usage
- Use TSM options to optimize network throughput
 - **tcpwindowsize xx**
 - **tcpbuffsize 32**
 - **tcpnodelay yes**
 - **largecommbuffers no**
 - **compression yes** on slow networks
 - **include.compression** and **exclude.compression**



Client Configuration - Filesystems

- Use multiple filesystems
 - Separate data with different recovery needs
 - Fire your NT administrator when he comes to tell you he made a one single file space system of 2TB for his user data.
 - Minimize the amount of critical recovery data
 - Less memory required for incremental backup
 - Faster backup/restore performance with the multi-session client
- Use RAID, snapshot capabilities to minimize the need for restores
- Keep filesystems less than 80% full to avoid file fragmentation



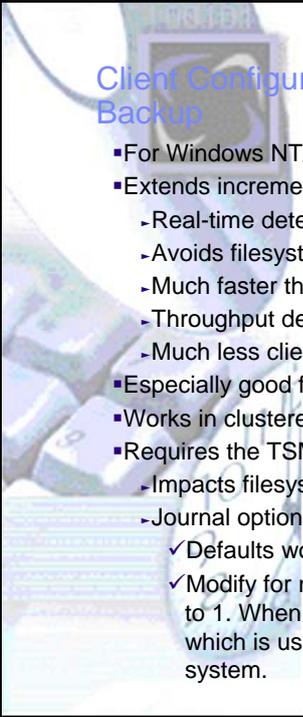
Client Configuration - Image Backup/Restore

- Unix and Windows 2000 clients
- Fast full filesystem restore
- Offline backup
- Online backup for Linux and Windows 2000
- Raw devices (all blocks)
- Used blocks only for Windows 2000 NTFS
- Throughput is independent of file size - like a single large file
- Cannot restore single files/directories from an image
- Combine with incremental backup for fast incremental image restore



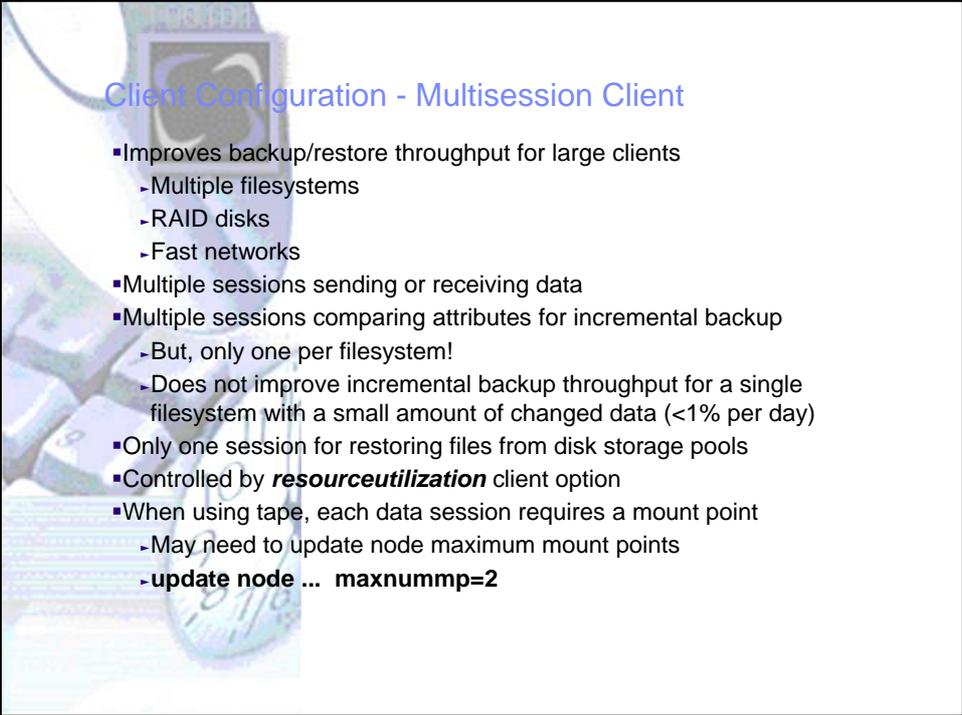
Client Configuration - LAN-free Backup/Restore

- Backup/restore to tape or disk using the SAN
- Meta data is sent to the server using the LAN
- Frees server from handling data --> better scalability
- Potentially faster than *slow* LAN backup/restore
- Great for large file workloads, databases (TDPs)
- Small file workloads have bottlenecks other than data movement
- For LAN-free using disk pools you must implement Tivoli SANergy
- Controlled by **enablelanfree** client option



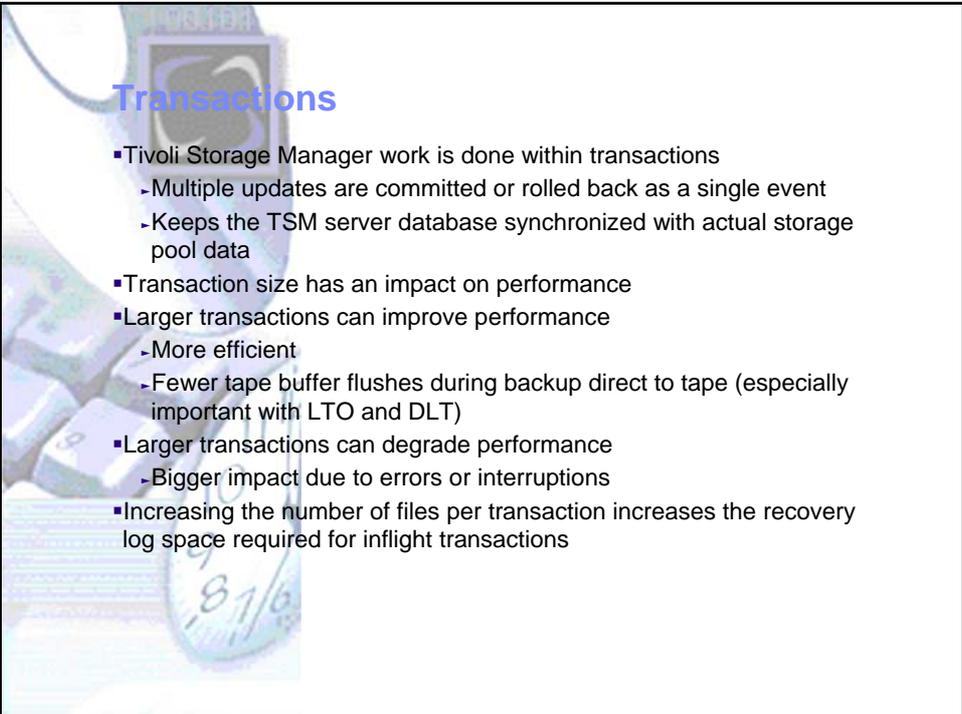
Client Configuration - Journal-based Incremental Backup

- For Windows NT/2000/XP clients
- Extends incremental backup to use a change journal
 - Real-time determination of changed files/dirs
 - Avoids filesystem scan and attribute comparison
 - Much faster than full incremental
 - Throughput depends on amount of changed data
 - Much less client memory usage
- Especially good for large filesystems with many files
- Works in clustered configurations
- Requires the TSM Journal service on the client
 - Impacts filesystem performance slightly
 - Journal options specified in **tsmjbbd.ini**
 - ✓ Defaults work well, just add the filesystems to be monitored
 - ✓ Modify for notebooks the setting : PreserveDBOnExit=1 from 0 to 1. When you do not modify this parameter, the database which is used for journaling will be deleted on a reboot of the system.



Client Configuration - Multisession Client

- Improves backup/restore throughput for large clients
 - Multiple filesystems
 - RAID disks
 - Fast networks
- Multiple sessions sending or receiving data
- Multiple sessions comparing attributes for incremental backup
 - But, only one per filesystem!
 - Does not improve incremental backup throughput for a single filesystem with a small amount of changed data (<1% per day)
- Only one session for restoring files from disk storage pools
- Controlled by **resourceutilization** client option
- When using tape, each data session requires a mount point
 - May need to update node maximum mount points
 - **update node ... maxnummp=2**



Transactions

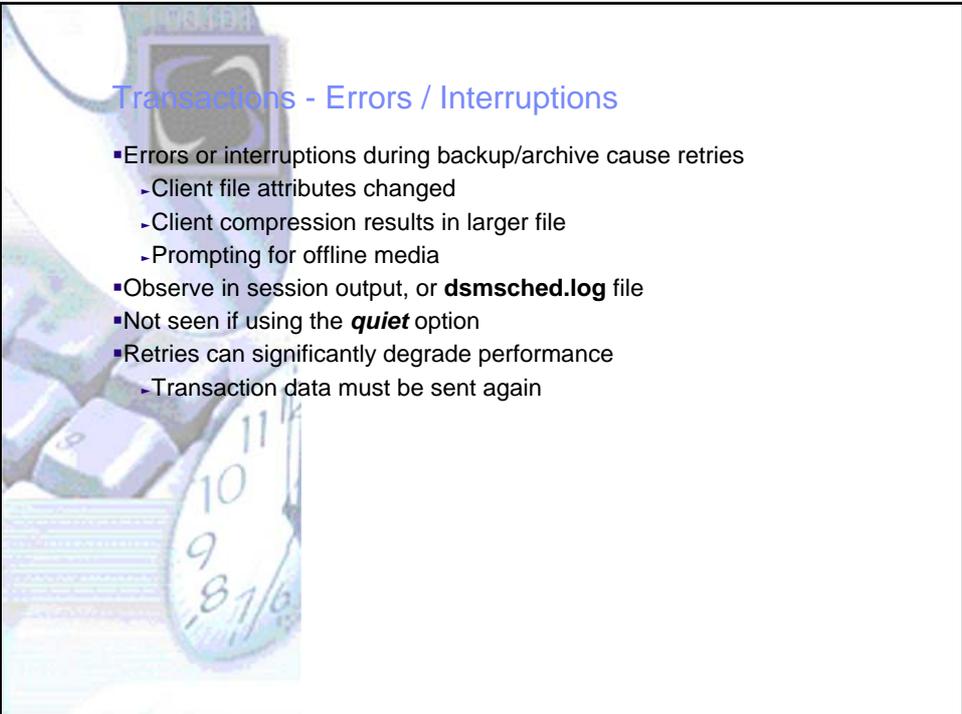
- Tivoli Storage Manager work is done within transactions
 - Multiple updates are committed or rolled back as a single event
 - Keeps the TSM server database synchronized with actual storage pool data
- Transaction size has an impact on performance
- Larger transactions can improve performance
 - More efficient
 - Fewer tape buffer flushes during backup direct to tape (especially important with LTO and DLT)
- Larger transactions can degrade performance
 - Bigger impact due to errors or interruptions
- Increasing the number of files per transaction increases the recovery log space required for inflight transactions



Transactions - Recommended TSM Options

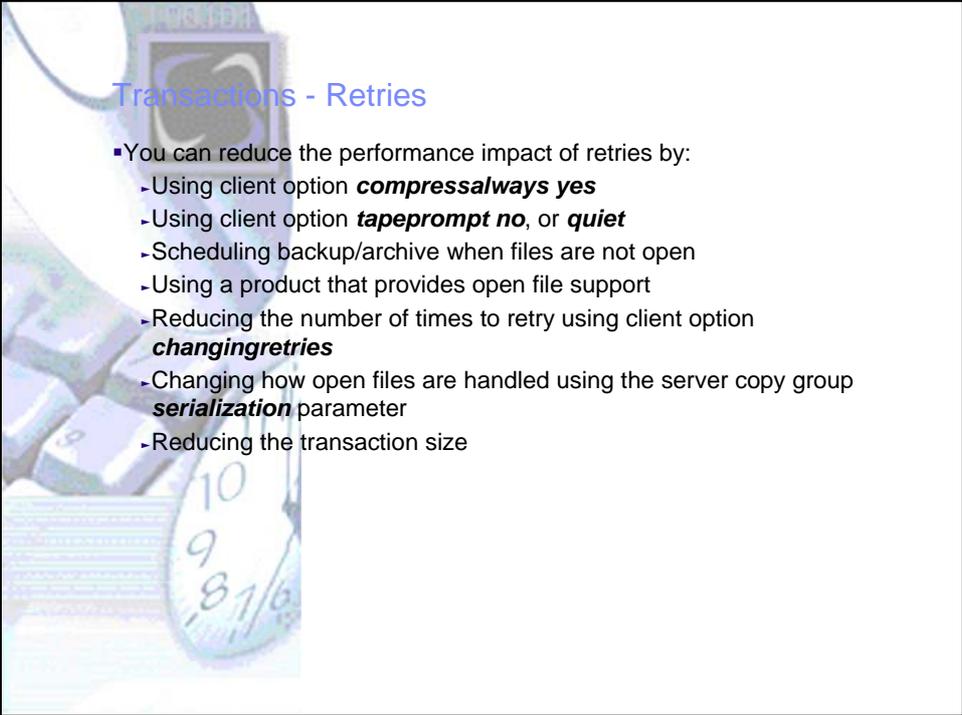
- Server
 - *txngroupmax 256*
 - *movebatchsize 1000*
 - *movesizethresh 2048*
- Client
 - *txnbytelimit 25600*

See the [TSM Tuning Options](#) for more details ...



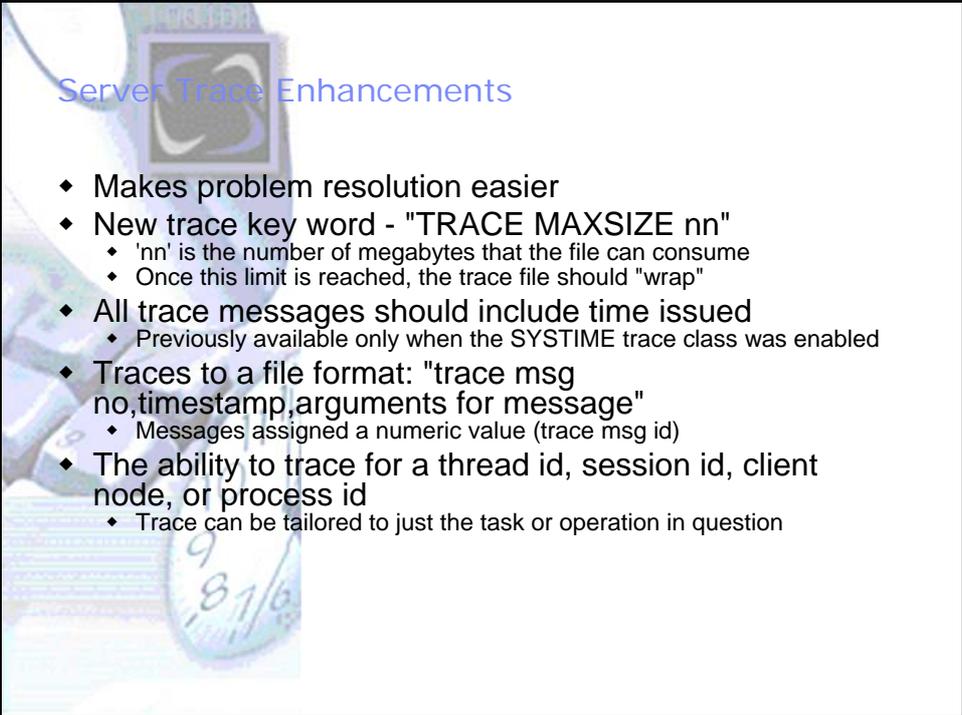
Transactions - Errors / Interruptions

- Errors or interruptions during backup/archive cause retries
 - Client file attributes changed
 - Client compression results in larger file
 - Prompting for offline media
- Observe in session output, or **dsmsched.log** file
- Not seen if using the **quiet** option
- Retries can significantly degrade performance
 - Transaction data must be sent again



Transactions - Retries

- You can reduce the performance impact of retries by:
 - Using client option ***compressalways yes***
 - Using client option ***tapeprompt no***, or ***quiet***
 - Scheduling backup/archive when files are not open
 - Using a product that provides open file support
 - Reducing the number of times to retry using client option ***changingretries***
 - Changing how open files are handled using the server copy group ***serialization*** parameter
 - Reducing the transaction size



Server Trace Enhancements

- ♦ Makes problem resolution easier
- ♦ New trace key word - "TRACE MAXSIZE nn"
 - ♦ 'nn' is the number of megabytes that the file can consume
 - ♦ Once this limit is reached, the trace file should "wrap"
- ♦ All trace messages should include time issued
 - ♦ Previously available only when the SYSTIME trace class was enabled
- ♦ Traces to a file format: "trace msg no,timestamp,arguments for message"
 - ♦ Messages assigned a numeric value (trace msg id)
- ♦ The ability to trace for a thread id, session id, client node, or process id
 - ♦ Trace can be tailored to just the task or operation in question

OS Tuning - AIX - Network Options

- The **no** command is used to tune AIX network options
- Use **no -a** to view current settings
- When using TCP window sizes > 64, set **rfc1323** to 1
- If you see non-zero "No mbuf errors" in entstat, fddistat, or atmstat, raise **thewall**
- Recommend setting **thewall** to at least 131072 and **sb_max** to at least 1310720
 - Newer versions of AIX have larger defaults (don't lower)
- no** settings do not survive reboot, so add to **/etc/inittab**
- Recommended settings:

```
no -o rfc1323=1
no -o thewall=131072
no -o sb_max=1310720
```

OS Tuning - AIX - Virtual Memory

- The **vmtune** command is used to tune the AIX virtual memory system
- Provided in the **bos.adt.samples** fileset
- Found in **/usr/samples/kernel** (not in default path)
- Use **vmtune** to view current settings

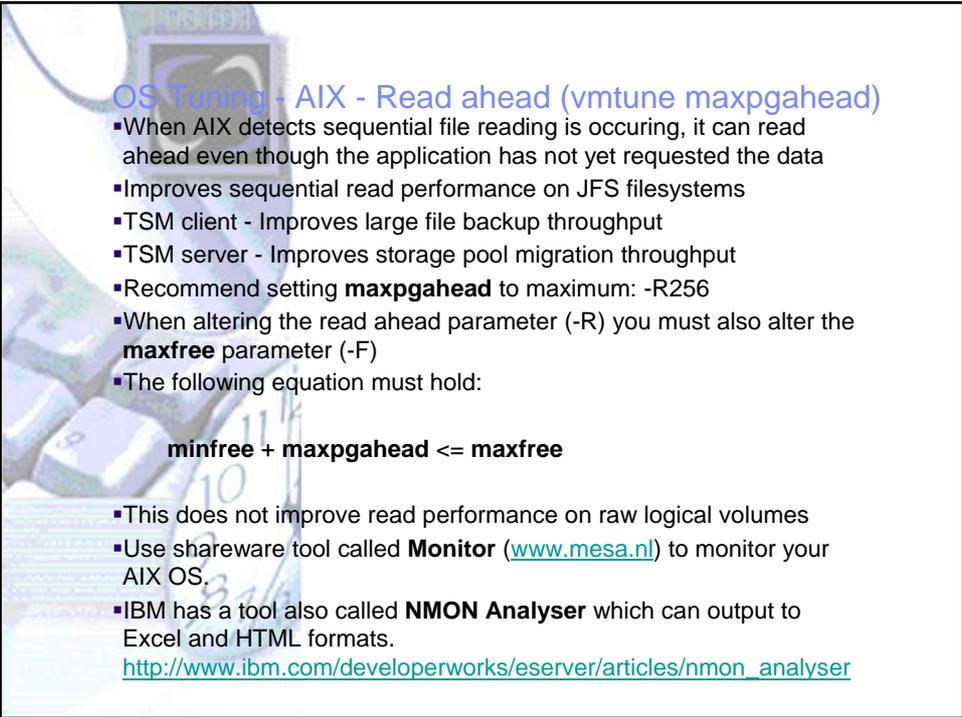
```
#vmtune
vmtune: current values:
-p -P -r -R -f -F -N -W
minperm maxperm minpgahead maxpgahead minfree maxfree pd_npages maxrandwrt
52219 208876 2 256 120 376 524288 0

-M -w -k -c -b -B -u -l -d
maxpin npswarn npskill numclust numfsbufs hd_pbuf_cnt lvm_bufcnt lrubucket defps
209696 4096 1024 1 930 336 9 131072 1

-s -n -S -L -g -h
sync_release ilock nokilluid v_pinshm lpgg_regions lpgg_size strict_maxperm
0 0 0 0 0 0 0 0
```

number of valid memory pages = 262119 maxperm=79.7% of real memory
maximum pinable=80.0% of real memory minperm=19.9% of real memory
number of file memory pages = 159438 numperm=60.8% of real memory

- vmtune** settings do not survive reboot, so add to **/etc/inittab**

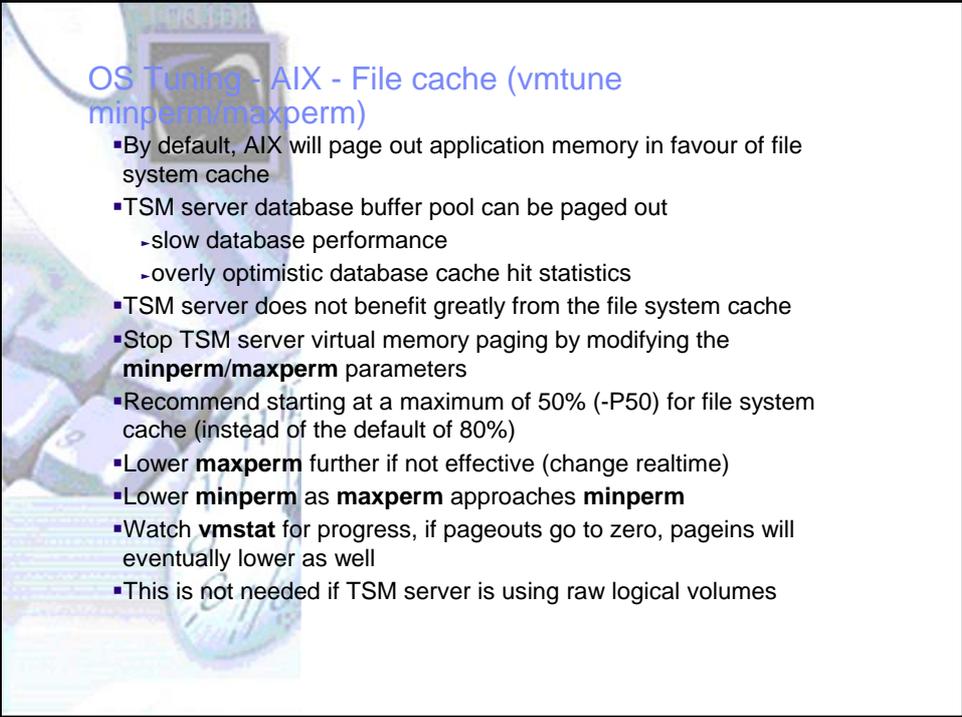


OS Tuning - AIX - Read ahead (vmtune maxpgahead)

- When AIX detects sequential file reading is occurring, it can read ahead even though the application has not yet requested the data
- Improves sequential read performance on JFS filesystems
- TSM client - Improves large file backup throughput
- TSM server - Improves storage pool migration throughput
- Recommend setting **maxpgahead** to maximum: -R256
- When altering the read ahead parameter (-R) you must also alter the **maxfree** parameter (-F)
- The following equation must hold:

$$\text{minfree} + \text{maxpgahead} \leq \text{maxfree}$$

- This does not improve read performance on raw logical volumes
- Use shareware tool called **Monitor** (www.mesa.nl) to monitor your AIX OS.
- IBM has a tool also called **NMON Analyser** which can output to Excel and HTML formats.
http://www.ibm.com/developerworks/eserver/articles/nmon_analyser



OS Tuning - AIX - File cache (vmtune minperm/maxperm)

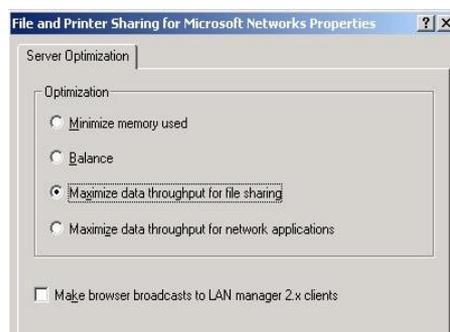
- By default, AIX will page out application memory in favour of file system cache
- TSM server database buffer pool can be paged out
 - slow database performance
 - overly optimistic database cache hit statistics
- TSM server does not benefit greatly from the file system cache
- Stop TSM server virtual memory paging by modifying the **minperm/maxperm** parameters
- Recommend starting at a maximum of 50% (-P50) for file system cache (instead of the default of 80%)
- Lower **maxperm** further if not effective (change realtime)
- Lower **minperm** as **maxperm** approaches **minperm**
- Watch **vmstat** for progress, if pageouts go to zero, pageins will eventually lower as well
- This is not needed if TSM server is using raw logical volumes

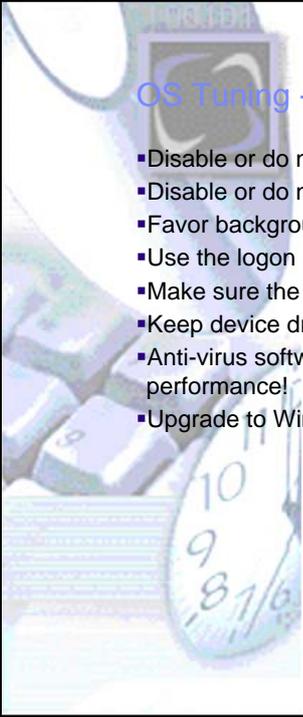
OS Tuning – AIX – AIXASYNCIO & AIXDIRECTIO

- ◆ Improved throughput and reduced CPU utilization on AIX servers is made possible with Direct I/O.
- ◆ Helps reduce CPU utilization and improves I/O throughput
 - Target 10% throughput improvement to disk
 - Previously when using JFS file systems for DB, LOG, and STORAGE POOL volumes on AIX, the AIX Virtual Memory Manager (VMM) caches read/write requests in memory in case sequential blocks are read/written
 - Target 30% reduction in CPU usage
 - Now TSM reduce CPU utilization by bypassing the VMM cache when writing to TSM JFS volumes
- ◆ AIXDIRECTIO option is added to the server options
- ◆ NEW in TSM 5.2!

OS Tuning - Windows - File System Cache

- By default, Windows will page out application memory in favor of file system cache
- TSM client performance can be degraded during incremental backup if inventory data is paged out
- Set to **Maximize for ... network applications**
 - Network Properties ...
 - Connection Properties ...
 - File and ... Properties
- Only tunable on Windows Server editions
- TSM server does not need this





OS Tuning - Windows - Miscellaneous Issues

- Disable or do not install unused services
- Disable or do not install unused network protocols
- Favor background application performance
- Use the logon screen saver, or none
- Make sure the paging file isn't fragmented
- Keep device drivers updated, especially for new hardware
- Anti-virus software can greatly impact backup/restore performance!
- Upgrade to Windows 2000 for better network performance



TSM Tuning Options

- TSM Server
 - Set in ***dsmerv.opt***
 - ✓ Must restart the server
 - Use the **setopt** command for some options
 - ✓ Takes effect immediately
- TSM Storage Agent
 - Set in ***dsmsta.opt***
 - Use the **dsmsta setstorageserver** command for setup options
- TSM Client
 - Use the GUI setup wizard
 - Set in ***dsm.opt*** (or set in ***dsm.sys*** stanza for UNIX)
 - Use on command line
 - Use in a client option set stored on the Server
 - ✓ Avoid having to go to each client and change options
 - ✓ Not for communication options

TxnGroupMax

Server Option

- Specifies the maximum number of objects (files and/or directories) included in a client session transaction
- **Improve Performance for Small files to LTO**
- Increasing TSM server option "TxnGroupMax" maximum value to 8192 (65000?)
- Reducing the number of tape synchronizations within the server
- Allowing the option to be set on a per node basis and reset back to the server global value
- Benefiting **ONLY** to Tape Storage Pools
- Being available on all server platforms
- The default value will remain the same (currently set to 40)

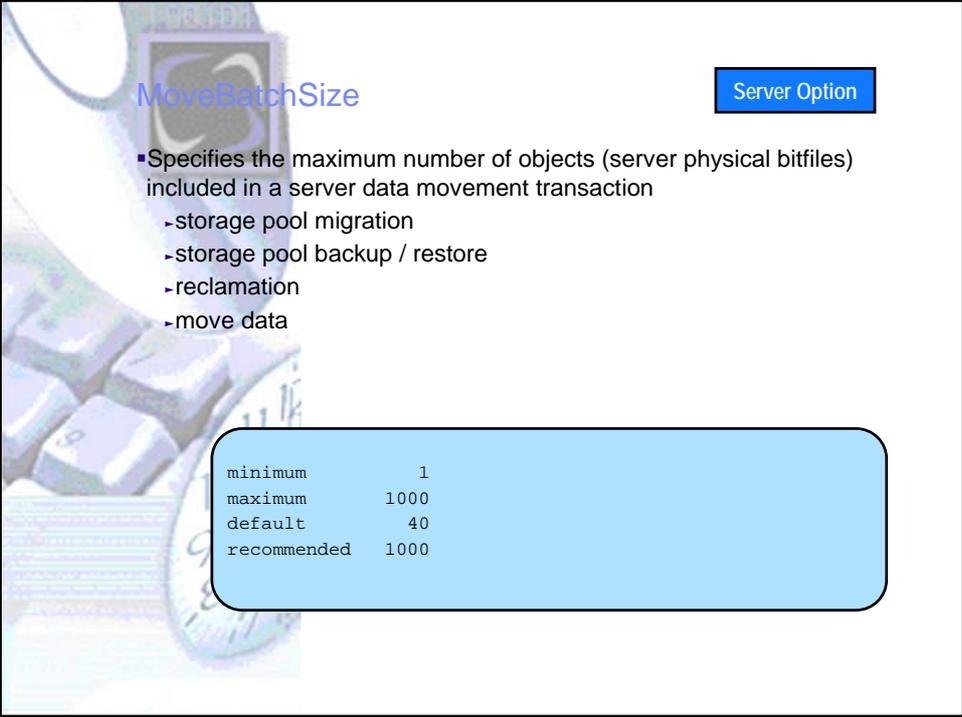
minimum	4	
maximum	256	
default	40	
recommended	256	as long as file retries are rare, and recovery log space is available
	8192	for LTO stgpool (from 5.2.0)

TxnByteLimit

Client Option

- Specifies the maximum number of kB included in a client session transaction
- A single file exceeding this size is always processed as a single transaction
- Backup/archive client only
 - HSM and TDP clients may use different rules

minimum	300	
maximum	2097152	
default	2048	
recommended	25600	as long as file retries are rare
recommended	2097152	for backup direct to LTO or DLT usefull for TDP's like exchange and SQL

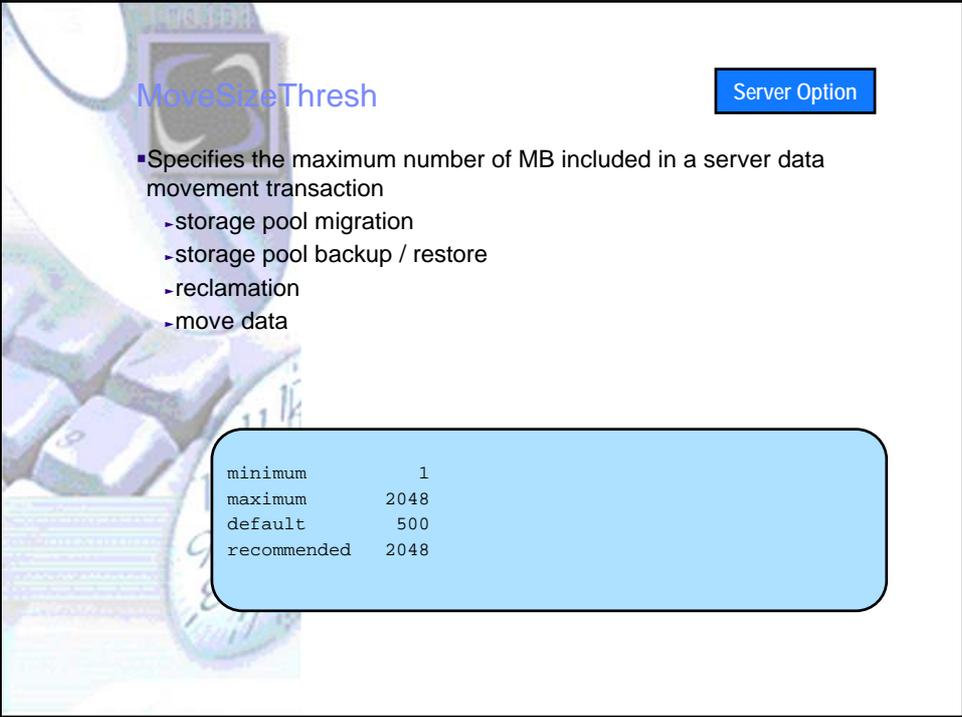


MoveBatchSize

Server Option

- Specifies the maximum number of objects (server physical bitfiles) included in a server data movement transaction
 - storage pool migration
 - storage pool backup / restore
 - reclamation
 - move data

minimum	1
maximum	1000
default	40
recommended	1000



MoveSizeThresh

Server Option

- Specifies the maximum number of MB included in a server data movement transaction
 - storage pool migration
 - storage pool backup / restore
 - reclamation
 - move data

minimum	1
maximum	2048
default	500
recommended	2048

BufPoolSize

Server Option

- Specifies the size of the server database buffer pool in KB
- A large buffer pool can reduce database disk I/Os
- A large buffer pool requires more memory
- See the **query db** command for tuning under a workload
- Or set **selftunebufpoolsize yes**

minimum	256	
maximum		limited by available virtual memory
default	2048	
recommended		tune using the amount of real memory as a guide
recommended	131072	for server with 1GB real memory

ExpInterval

Server Option

- Specifies the interval, in hours, between automatic inventory expiration processing
- While running, inventory expiration can degrade backup performance
- Run expiration outside the backup window
- Run expiration only as often as necessary

minimum	0	automatic expiration disabled
maximum	336	(14 days)
default	24	
recommended	0	use a server admin schedule to set the time expiration should begin

ResourceUtilization

Client Option

- Controls how many client sessions are established with the server
- More sessions require more client and server resources

minimum	1
maximum	10
default	1
recommended	1 workstations (1 data session)
	5 small servers (2 data sessions)
	10 large servers (4 data sessions)

LargeCommBuffers

Client Option

- Specifies that large buffers are used for client/server communications
- Large buffers are more efficient
- Need to tune filesystem read ahead to get benefits
- Can't tune read ahead on most platforms

	yes
	no
default	no
recommended	no (all platforms except AIX)
	yes (for AIX, but not for LAN Free)

TopWindowSize

Server Option

StgAgt Option

Client Option

- Specifies the amount of data in kB that can be buffered at one time on a session
- If the limit is reached, the sending host cannot send more data until a window update is received
- Larger window allows sender to continue sending data
- Large window is useful on reliable, long distance, or high latency networks

minimum	1	
maximum	2048	
default		platform dependent
recommended	63	for Windows NT/2000/XP
	128	AIX in SP2 environments
	64	for Solaris

TopBuffSize

Server Option

Client Option

- Specifies the maximum size of TCP sends done by TSM
- Windows server does not support this option

minimum	1	
maximum	2048	
default		platform dependent
recommended	32	



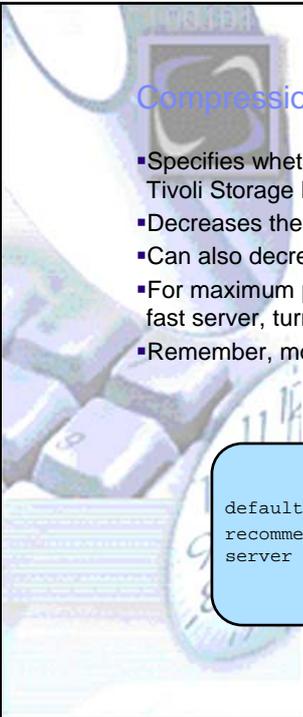
TopNoDelay Server Option StgAgt Option Client Option

- Specifies whether or not TCP/IP will buffer successive small outgoing packets
- Buffering can improve network utilization
- Buffering requires a delay that can impact session throughput greatly

```

yes
no
default          platform dependent
recommended     yes

```



Compression Client Option

- Specifies whether to compress file data before sending it to the Tivoli Storage Manager server
- Decreases the amount of storage required on the server
- Can also decrease session throughput
- For maximum performance with a single fast client, fast network, fast server, turn compression off
- Remember, most tape devices have hardware compression

```

yes
no
default          no
recommended     no  single client, fast network, fast
server          yes  multiple clients, slow network,
                  slow server

```



Additional Information

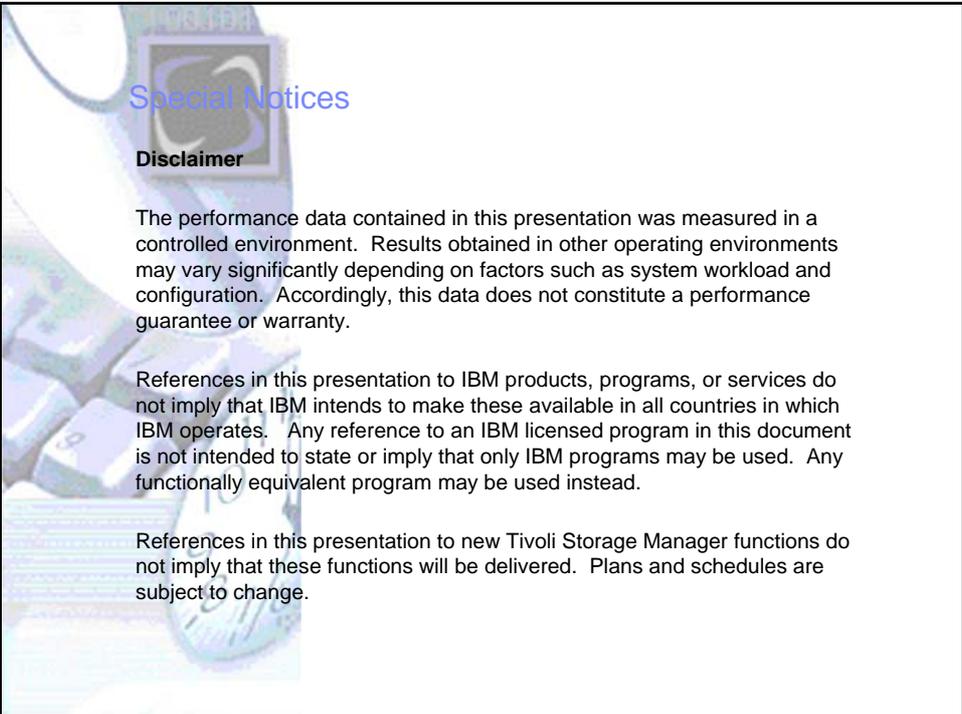
- www.ibm.com/software/tivoli/products/storage-mgr/
- www.redbooks.ibm.com

- These notes, based on material from:

Zong Ling, Ph. D.
Tivoli Storage Software Performance
San Jose, California, U.S.A.
zling@us.ibm.com

- Contact details:

Peter Pijpelink
Peter.pijpelink@plcs.nl
www.plcs.nl



Special Notices

Disclaimer

The performance data contained in this presentation was measured in a controlled environment. Results obtained in other operating environments may vary significantly depending on factors such as system workload and configuration. Accordingly, this data does not constitute a performance guarantee or warranty.

References in this presentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM licensed program in this document is not intended to state or imply that only IBM programs may be used. Any functionally equivalent program may be used instead.

References in this presentation to new Tivoli Storage Manager functions do not imply that these functions will be delivered. Plans and schedules are subject to change.



Special Notices ...

Trademarks and Registered Trademarks

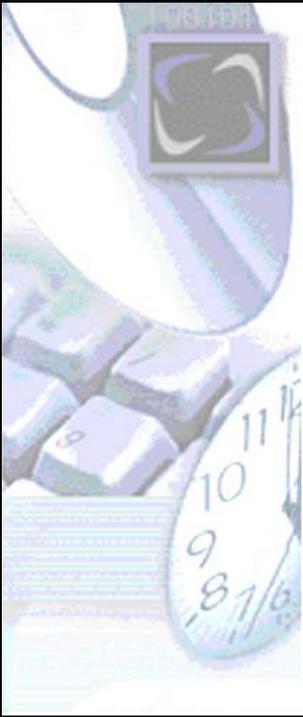
The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX
Enterprise Storage Server
IBM
Magstar
Redbooks
RMF
ServeRAID
Tivoli
z/OS

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

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft 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.



Questions?

