How to Shrink your Backup Window

By Nancy Roper

Customers around the world are looking for ways to minimize their backup window. This article will provide a step by step process to work towards this goal.

STEP #1: Is your Backup Meeting the Benchmarks?

The first step is to determine whether your current tape hardware is running at rated speed. If it is, then purchasing newer tape hardware will likely help shorten your backup window. If your tape drive is not running at rated speed, then a faster drive will not be of any help until you resolve the bottlenecks.

Start by figuring out what speed your drive is running at today. To do this, you need to take a save and figure out how much data is being saved, and how long it takes, then do a quick mathematical calculation, and convert the result to both MB/sec and GB/hr. The easiest way to do this is to consider your full system save. Your operators will know the elapsed time, and you can estimate the amount of data saved using WRKSYSSTS and multiplying the System ASP size by the Percentage Used. If you have auxiliary storage pools you will need to factor them into this equation also: the WRKASPBRM command may be helpful to you if you have the Backup, Recovery and Media Services (BRMS) product installed. Alternatively, if you typically use smaller saves, then calculate the amount of data using the DSPOBJD command. If you have BRMS, you can also use GO BRMBKUANL and choose option #3 (Display Backup Analysis) to see a list of library sizes at the last full or incremental backup.

Next, get a copy of the iSeries save/restore performance benchmarks. These are listed in a book called the "iSeries Performance Capabilities Reference" which is published multiple times per release as new benchmarks are run. The book is over 300 pages long and includes benchmarks for all aspects of the system, but you only need to review the save/restore chapter, which is chapter 15 in the latest book. Make sure you choose the book for the operating system level you have installed. These books can be found at the following url: www-1.ibm.com/servers/eserver/iseries/perfmgmt/resource.htm

Alternatively, ask your IBM rep or business partner to retrieve the "iSeries Tape Performance Summary Chart" which is a 2-page summary of the save/restore benchmarks from recent years. This chart is available on "Techdocs" at the following urls:

- **For IBMers:** http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1193
- **For Partners:** http://partners.boulder.ibm.com/src/atsmastr.nsf/WebIndex/PRS1193

In the benchmark information, find the description of the various workloads and consider which workload is the closest match for the data in the save you are investigating. Small objects require a lot of overhead during the save, and hence save at much slower rates. Examples of the workloads are shown below. Note that there are also benchmark workloads for Integrated File System (IFS), Domino, and Linux NWS:

- **Source Files**: 96 source files with approximately 30,000 members in total for a total of 1 GB (i.e., an average of 33 KB per member).
- **User Mix**: a single library containing a combination of source files, database files, programs and command objects, data areas, menus, query definitions, as well as other common iSeries objects found in libraries. The NUMX12GB workload is a 12 GB library containing 52,900 objects (i.e., an average of 227 KB/object, although sizes vary considerably).
- **Large File**: a single database file with members 4GB in size.

Nancy Roper

Now look at the tape drive benchmarks and find the figures for the type of tape drive you are using, and ideally on a CPU / disk of similar size and generation as yours. Scan across the columns until you find the benchmark workload that you believe matches your save. Read off the tape drive speed achieved in the benchmark and compare it to the speed your save is attaining. Decide whether you think your tape drive is meeting rated speed during your save. If it is, then go on to step #3. If it is not, then continue with step #2 where you will try to identify the bottlenecks.

STEP #2: Identify and Resolve Bottlenecks

There are many factors that affect backup performance. When faster backup performance is needed, people are quick to suggest a faster tape drive. However, this will only help if the tape drive is the slowest link in the chain. If your research in step #1, determined that your tape drive is NOT your bottleneck, then in Step #2, you will try to identify and resolve the bottleneck that you are experiencing. You will need to look at both hardware and software parameters.

Start by considering the hardware that you are using. You will need to look at your adapter cards, the CPU and memory you have available, the type and number of disks, and the cardslot and high speed loop (HSL) layouts. Here are some details:

If you are on a CPU that is a 7xx or older, then the tape adapter cards available to you are dramatically slower than the current tape drive product line. For example, the HVD SCSI cards available on 7xx (fc 2729 and fc 6501/6534) ran at 13 and 17 MB/sec respectively,

compared with the fibre cards used on 8xx CPU's and above that run at 100 and 200 MB/sec. If you are on one of these older CPUs, then your backup speeds will be limited until you can upgrade to a newer system. If you are on a newer CPU, you may still have your tape drives attached to the older cards via a migration tower. You should consider upgrading to newer cards, (eg 8xx supports a fc 2749 HVD SCSI adapter that runs at 38 MB/sec) or better still, upgrading your drives and your cards to LVD SCSI or fibre-attached technologies.

Backup performance is also affected by the amount of CPU available to the job. According to the Performance Capabilities Reference, large-file save streams need at least half a CPU each. For user mix and small file save streams, you should allow 1.3 CPU's per save stream. Interestingly, the size of CPU is far less important than the number of CPU's available.

Memory is also a factor in backup performance. A recent test showed that a single large-file save stream performed close to rated speed with 500 MB of memory, and performed at rated speed with 1 GB of memory. By comparison, a user mix stream needed 1 GB of memory to perform acceptably. Check to make sure your saves have at least this amount of memory.

The quantity and technology of disk is also a factor in save performance. In order to match the benchmarks for large-file saves using savefiles or using the latest technology tape drives, you will want plenty of arms of the latest technology disk (eg 15K RPM drives on fc 2757 or fc 2780 adapter cards). If you have older disk or a small number of arms, check the disk response times on your performance reports to see whether it may be a bottleneck. You can also review the Performance Capabilities Reference and look for a benchmark on similar hardware to your system and compare your performance to the benchmarks.

Next you should review your cardslot and High Speed Loop (HSL) layout. The high end tape drives (LTO-2, LTO-3 and 3592) all command enough bandwidth that it is possible to consume an entire bus or HSL loop when running a backup. Start by putting your tape adapter cards into high speed slots. To determine which slots are high speed in your system, review the diagrams in the redpaper on PCI Placement rules found at the following url: www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp4011.html?Open.

Additionally, try not to mix tape with other adapter cards on the IOP since you may need to reset the tape IOP which would impact other adapters on the same card. As for HSL and bus placement: Rules-of-thumb for HSL-1 loops are to put at most 1 high speed tape drive per bus, and at most 2 high speed tape drives per HSL loop. For HSL-2 running on i5 CPU's, it may be possible to support 3 high end tape drives on certain buses and on an HSL, depending on the placement of other cards. For assistance in designing your cardslot layout for multiple drive environments, please contact your IBM rep or business partner.

Once you have finished reviewing your hardware for potential bottlenecks, the next step is to look at your save parameters. Items to consider include the compression/compaction settings, the use-optimum-block setting, the output(*outfile) and the BRMS object-level-detail settings, the structure of your save commands, your IFS saves, and the placement of your IPLs in your save stream.

The compression/compaction settings are important in optimizing your backup speed. The current tape adapters do not support compression, so if you accidentally request it, then it is done by the CPU which slows the save dramatically. Instead, most customers will want to set the compression/compaction parameters both to *DEV which is the default setting for recent releases. This setting means that if the device supports compaction then it will be used and compression will not be used. These parameters are set in the regular SAVxxx commands. For BRMS users, look in the attributes of the control group (WRKCTLGBRM option #8) or on the SAVxxxBRM command.

Another parameter to check is the "Use Optimum Block" parameter. The default setting in recent releases is *YES which means that data should be sent to the tape drive in large blocks vs. small blocks. This uses less CPU resource, thus allowing backups that are CPU-constrained to run dramatically faster. This parameter is set on the SAVxxx, SAVxxxBRM commands and in the attributes of the BRMS Control Groups. Most customers will want it set to the default.

When using BRMS, the default is to record library-level detail for successful saves and object-level detail for objects that are not saved successfully. Optionally, object-level detail

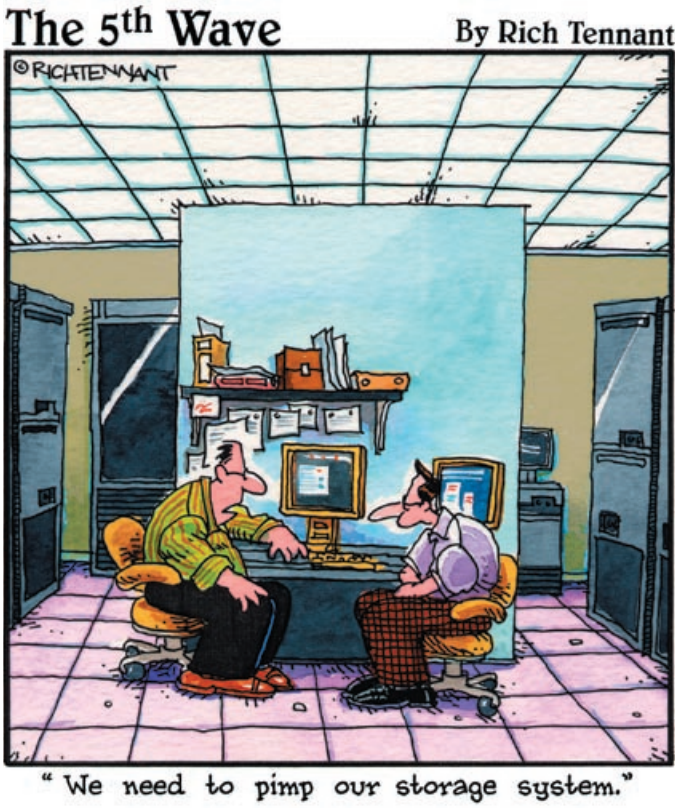

can be kept regardless. Although having this information makes it especially easy to restore individual objects, there is a price to pay in terms of save performance. When looking to shorten your backup window, use object-level-detail sparingly. Similarly, when using the SAVxxx commands, the OUTPUT(*OUTFILE) option impacts performance and should be used only when needed.

The structure of your save commands can affect performance. If multiple libraries or objects are saved in a single command (e.g. SAVLIB LIB(A B C)), then the system can overlap the save of one with the pre-processing of the next, thus shortening the duration of the overall command. By comparison, if each item is saved on a separate command (eg SAVLIB LIB(A), SAVLIB LIB(B), SAVLIB LIB(C), then the backup could take considerably longer. When using BRMS, be aware that BRMS is generating SAVxxx commands in the background. Whenever possible, keep the parameters (e.g., save-while-active, object level detail, incremental saves, library saves vs. list-based saves, etc.) the same from line to line in a control group since this allows BRMS to issue a single SAVxxx command in the background.
If saving the IFS, there are a number of considerations that impact performance. For information, review the article entitled "Backing up the IFS - Experience Report" in the IBM Information Centre: http://publib.boulder.ibm.com/infocenter/iseries/v5r3/ic2924/info/experience/ifsystem.pdf

The first time you access an object following an IPL, extra checking is done. If you do an IPL and then run your save, this "first touch after IPL" can extend the backup time. If you need to get your IPL and your save done in a short window, then consider doing the save first, and then the IPL. However, if your save window is longer, then there may be merit in running the save after the IPL to get the "first touch" of each object completed prior to startup of your application.

STEP #3: Estimate the Performance of Newer Tape Technology

Once you have completed step #1 and step #2, your tape drive should be running at rated speed according to the benchmarks. This means that adding a newer technology tape drive will likely increase your save performance. To estimate the speed of a newer drive, you need to return to the performance benchmark listings. Find the benchmark that matches your tape performance today and take note of the workload type that corresponds to that speed. For example, if you have a fibre LTO-1 running at 95 GB/hr on an 8xx CPU with 15K RPM disk, then that would suggest that your data is a good match for the user mix workload. Now find the corresponding workload for the newer technology tape drive that you are considering and read off the benchmarked speed. Use this figure to calculate the duration of your backup on your new tape drive and decide whether it will meet your requirements. If it will, great! If not, then read on...

STEP #4: Consider Other Techniques to Shorten Your Backup Window

If a faster tape drive will meet your requirements, then this is typically your best option. It is important to keep your backups as simple as possible since this will make your restore easier and a single drive save is typically as simple as it gets. However, if you need a shorter backup window than what you can accomplish with a single new-technology drive, then you will need to look at other options to shorten your window. These include the following solutions:

Consider using a backup strategy where you save less data than you are saving today. Examples would be replacing full saves with SAVCHGOBJ, or figuring out which libraries are changing and saving only that data each day. If you move to a more elaborate save strategy like this, then consider using a tape management system such as BRMS to help keep track of your saves and assist you with your restores.

Consider including savefiles in your backup strategy. On current disk technology, savefile saves of large-file data have been clocked at 1000 MB/sec which is double the speed of the fastest tape drive today. This may allow you to get your users back onto the system quickly, and then you can spill the data to tape thereafter.

If you are already running a new tape drive at rated speed but need a shorter backup window still, then consider running multiple tape drives at once. There are two ways to do this: "concurrent" or "parallel" saves. For concurrent saves, YOU carve up your data across multiple drives via multiple save commands. For parallel saves, you ask the system to carve your data across multiple drives using a single save command. If using parallel saves, consider BRMS a "must" since it creates the underlying media definitions for you and generally makes this function much simpler to implement. Note that parallel saves have some overhead compared with concurrent saves, and they have some restrictions on where they can be restored. There are also some tricks that you need to know to restore them efficiently. The place where parallel saves are a good fit is for single large objects that cannot be split across drives via concurrent saves. Prior to implementing multi-streamed saves, revisit the hardware section of this article to make sure you have sufficient CPU, memory, disk arms, HSL capacity etc to support multiple simultaneous streams.

If you can't afford to have your users off the system for the entire duration of the backup, then consider the save-while-active (SWA) function. For applications with commitment control, save-while-active can run a backup without quiescing the application: your users can stay on the system throughout. An example is the online Domino save provided by BRMS: the Domino transaction log is saved in the background and used during the restore to get the data to a consistent point. If your application does not have commitment control, then you need to quiesce your application for a brief period to get a checkpoint, then let the users back onto the system to work while the backup runs in the background. Save-while-active has had dramatic improvements over the years, including faster checkpoint times, fewer restrictions on the functions that can be running during the save, and new functions such as rapid-checkpoint-SWA introduced at V5R3 to handle transactions that do not reach a commit boundary during the checkpoint period. So if you've considered SWA before and decided not to use it, then now might be a good time to re-visit this function.

For customers who can afford a short quiesce period, an alternative to SWA is to use the external disk copy functions. For this solution, the system needs to use external disk for the entire ASP that needs to be saved. The system is quiesced briefly while an instant copy of the ASP is made by the external disk subsystem. This copy of the disks is then attached to a second system or partition and the backup is run from the second system while the users resume their work on the first system. When using this solution, it is strongly recommended to use independent ASPs along with the iSeries Copy Services Toolkit.

For customers who need their backup to run with zero downtime, then the only options are save-while-active with commitment control, or the iSeries High Availability Business Partner (HABP) solutions. The HABP solutions are the premier way to run systems that require 7x24 up-time. In these solutions, transactions are journalled and passed to a secondary system, then applied to the database on the secondary system to keep it in synch with the original system. When it is time to run the backup, the apply job on the secondary system can be put on hold and the application can then be backed up, while the users continue to run on the primary system without interruption. The changes are held safely on the second system ready to apply once the backup is finished. If desired, save-while-active can be used so the "applies" can continue as soon as the SWA checkpoint is reached. This solution offers the fringe benefit of having a second system that the users can be switched to during planned and unplanned outages to maximize system availability.

Conclusion

For most customers who are looking to reduce their backup window, the above techniques will provide assistance. Key points to remember are: (1) Prior to buying a faster tape drive, make sure that it is the bottleneck, otherwise you may not see the performance gain you are expecting. (2) Try to choose the simplest backup strategy that will meet your up-time requirements so your recovery will be manageable. (3) Always plan to test your recovery at least once a year to ensure you are saving everything you need and your staff is familiar with the recovery procedures. TUG

Nancy Roper *is a Consulting IT Specialist. She currently works in the IBM Americas Advanced Technical Support group, assisting the largest iSeries customers with their availability strategies. Nancy is a seasoned technical expert on iSeries tape, SAN, and BRMS, and is co-author of the redbook "*iSeries in a Storage Area Network*" (SG24-6220).*