

Fixing the backup and recovery problem

By Rob Nieboer

We know from anecdotal and other evidence that one IT's biggest problems is backup and recovery. It has been estimated that as much as 60% to 70% of storage management effort is consumed by management of backup and recovery alone. This is a problem not just limited to small IT shops – in fact, we have observed that very large IT shops often have worse experiences than small shops.

Simply put, the biggest customer pain points around backup and recovery are these:

Speed

There is a fundamental problem in exploiting the bandwidth of tape devices these days. Tape devices are often capable of speeds of 30 to 35 MB/second or more. The problem lies in being able to serve up data at this rate from application and backup servers. And when you can't maintain this kind of data rate, tape devices have to stop. And when they stop, to start again, they have to re-position exactly where they stopped. This means a few iterations of back-hitching and forward-spacing before the tape drive has correctly positioned the tape so it can continue to write. Obviously, this process – which is called 'shoe-shining' – has a significant detrimental impact on tape drive performance. In fact, it can reduce a tape drive's performance from 30 or more MB/second to 5 or 6 MB/second!

To overcome this problem, some people send streams of data from multiple disk drives to a single tape drive. This allows the tape drive's buffer to always be filled and therefore the tape drive never has to stop. This process is called multiplexing. The problem is that if I want to restore from a tape that has data from many disk drives interleaved on it, I have to sort through the data of 5 or 6 or 7 disk drives to find the data from the disk drive or file that I want. So in making the backup run more efficiently, the recovery will run much slower.

Also, if I make offsite copies from the first set of backups (called cloning), this process also runs very slowly because I have to re-create each individual disk drive from the interleaved data.

I also have concerns about recovery speed, even if I don't have a problem with multiplexed backups. Just simply recovering a file or volume requires me to find the volume serial number I need, figure out where the volume is, mount it on a tape drive and then search through the volume looking for the file or disk backup I need. This could take multiple minutes or more.

Another factor at play here is that we have to complete the backups in a certain time window. So we buy a relatively large number of tape drives and crank them up for the 6 or 8 hours of the backup window. And if backup workloads increase, we buy some more because we only have a limited amount of time, and we may not be able to use their full performance capabilities.

But getting approval to buy more drives might be a problem, because the drives we already have are only 25% to 30% utilised! This adds to the pressure.

Reliability

Probably the biggest problem customers face with backup and recovery is the poor reliability of the process. Many customers tell us that as much as 30% or more of all backups fail. There are many reasons why backups can fail. Backups can fail because a tape cartridge failed, or because I had a data check on a tape drive or because a robot failed. However, it is MUCH more likely that a backup failed because of human error, or because the backup got

cancelled because we ran out of time, or because of the “end-to-end” complexity of the backup process.

Strategic Research estimates that as many as 50% of remote backups fail. This is because either non-IT professionals are involved in what may be a less than bulletproof process, and also because the sheer volume of data growth and backups being done to a schedule that can no longer be met.

Whatever the reason, backups are a complex process. When you look at all the components that can exist in the end-to-end backup and recovery process it is not surprising that failure rates can be high.

First is the backup application(s). How many are there in this IT shop? If there are three or more different backup applications – say, Veritas Netbackup, TSM and Legato – then implementing backups can become very complex. Even if there is only one backup application, have we implemented it correctly? Have we set the right parameters? Do we even know what each application and its data needs from backups?

Are we doing full backups often enough, or are we doing so many incremental or differential backups that recovery is taking too long? Do we have tape sharing software that can complicate things? At the application level we need to assess whether we correctly understand the backup and recovery needs of the different data. We need to look at the potential need to consolidate backup applications. We need to assess whether the backup application(s) are correctly, or optimally implemented.

Then there are the network connections between the application and backup servers and the tape libraries and drives. Then finally the tape libraries and drives themselves. In between we have host bus adapters and device drivers, and so on. No wonder we have reliability issues?

Human intervention

When we have reliability problems with backups and recovery, we typically assign these problems to people. It is their job to figure out what backups worked, what failed, why they failed, and how to correct the problems. The challenge here is that we typically never staff this function with enough people and we don't give them tools to make their jobs easier. We have a certain failure rate which we expect people to diagnose and rectify, but because we don't have enough people, the reliability becomes worse, and so we expect people to diagnose and rectify, but we don't have enough people, so the problems get worse,... and so on. Very soon, we have a death spiral!

Disk buffers

A growing number of people believe that by inserting some kind of disk buffer into the backup configuration, they will solve all their backup problems.

This is a very compelling idea. Simply installing some kind of disk buffer provides a lot of benefit. For those who believe that some of their backup problems are caused by tape, it's a very attractive idea. Backups will happen at disk-to-disk speed, and recoveries, if the file is still in the disk buffer (a big “if” by the way), will happen at disk-to-disk speed.

Again, for those who believe that the problems are all tape related, reliability will approach 100% because it's all disk-to-disk I/O. And that will eliminate the high staffing costs too, right? Well, maybe not.

Certainly the idea of disk-to-disk backups is attractive. The performance problems of backup SHOULD improve, because you get a consistent level of performance. But the degree to which I have a sustainable impact on the three pain points of speed, reliability and high staffing cost depends on which implementation of a disk buffer I choose.

Simply put, there is a scale against which I can measure the specific implementation of a disk buffer. That scale can read from “dumb” to “smart”.

If I simply install some “dumb” disk capacity, I may have to change the backup processes, which currently think they are writing to tape drives. Changing process is a bad thing – it increases risk and consumes a lot of human intervention.

If I simply install some “dumb” disk capacity, how do I know how much capacity I need? I’d like to install enough capacity to store, say, 24 hours of backups and to ensure that if I need to recover that the data will still be on disk. I probably don’t have tools that tell me how much capacity I need - or how much performance I need for that matter. This makes installing dumb disk a risky proposition.

In addition, how do I move the backup data from disk to real tape devices? Do I need to make that decision? When should I or can I move the data?

In reality, I would probably like to have disk that can present an image of tape – a virtual tape implementation. That would mean that I could use existing backup processes that expect to write to a tape drive. That implies a kind of intelligence that requires a server platform. Already, I’m starting to see that I would like smarter disk rather than dumb disk.

I would also like some intelligence applied to managing the backup data. I would rather not have to move the backup data back to the application or backup server simply to enable me to write it to tape. I would like some kind of outboard data management capability that can not only decide what to move to tape and when to move it, but would be capable of moving it without involving the application or backup server.

Then there are other data management functions that would make sense. Wouldn’t it be appropriate to have intelligence that allows policy to be applied to each backup file or volume? Policy that defines how many copies of the backup should be written to tape. Policy that determines whether to keep a copy on disk after its written to tape, or whether to delete the disk copy after its written to tape.

Maybe a policy could be defined that says I want a copy on a disk buffer and another copy on a different RAID group in the same disk buffer, but I never want it written to tape. Just let it expire in the disk buffer.

The point of this discussion is this: simply installing a dumb disk buffer may have some initial benefit in terms of the speed of the backup. But if you want a significant, sustainable impact on the speed of the recovery, and an improvement in the reliability of the process, and a reduction in human intervention, then it appears that the preference is for smarter disk – in fact, for a managed disk buffer.

This suggests that we need an appliance-based approach to improving backup and recovery.

By the way, a disk buffer does nothing to relieve the complexity of multiple backup applications or to improve a poorly implemented backup application.

There is another choice that customers are turning to which uses a D-D-T function embedded in the backup application customers are already using. While this appears to be attractive, here are some issues:

1. Available capacity of the disk buffer must be closely monitored to ensure that backups can complete successfully.
2. Availability of sufficient physical tape resources still needs to be planned for, but not at the time backups are done. Since the movement of data to tape takes place at a different time, that is when tape resources need to be available.

3. This methodology still may need as many physical tape resources as before, and may still only use them for a limited period of time per day.
4. Backup and application servers will be involved with data movement for a longer period of time.
5. There may only be limited control over how and when data is moved to tape, and how many copies are made.

Again, we get the strong indication that perhaps even backup application-based D-D-T may not be the best long-term answer.

By the way, another reason why backups fail and why they are difficult to manage is because we often have multiple places where we do backups, and so, we also have multiple places where we have physical tape libraries and tape drives. Consolidation is a good idea, and that allows us to use more 'industrial strength' devices.

A strategic view of an ideal process

If I step back from the problem and take a more strategic view of the backup problem, I could think of some characteristics of an ideal backup process:

- * It would be sustainable – it would be capable of scaling up as workloads change and increase.
- * It would be flexible and allow me to change the process as easily and as simply as perhaps changing a policy statement.
- * It would be fast and be capable of driving my fastest tape drives at full speed for backups, and would allow me to do most recoveries directly from disk.
- * It would be reliable – it would have components in the end-to-end chain that approached 4 9's and 5 9's levels of availability.
- * It would enable 'hands-free' operation; a 'set it and forget it' approach.
- * It would allow me to adapt policies according to the value of the data to the business.
- * It would scale in both capacity AND performance as the amount of data to be protected increases.
- * And it would, ideally, give me a mechanism to continuously drive cost out of the infrastructure.

In fact, we could expand on the idea of desirable characteristics in a backup process by building a picture of the 'perfect' backup solution. We could call this a 'world's best practice' for backup!

- * The perfect backup system would be able to continuously discover new storage so that we never again face the problem of finding out that a business unit bought some more storage, and they never told IT, so IT never built it into the backup process.
- * The perfect system would consolidate from multiple backup applications to just one.
- * The perfect backup system would incorporate a disk buffer, but one that looked like tape – a virtual tape system.
- * The disk buffer would be a managed resource, and part of an intelligent appliance.
- * The perfect system would apply consolidation at the front-end and the back-end of the process – fewer backup applications, fewer, bigger libraries, fewer, better tape drives (they better be better because we will increase duty-cycle).
- * But because consolidation has another side – if I decrease the components in a process that reduces cost and complexity but increases the risk to the process if something does fail – then I better have high availability built into the end-to-end chain.
- * I want management intelligence applied through business value-driven policy so that I can manage this with very little human intervention.
- * I want metrics that tell me what's happening and tell me what failed and why it failed.

* And of course, I want a system that allows me to drive cost out of the process as data and workloads increase.

Conclusion

So when we think about the backup problem, we are really talking about a large complex end-to-end process. This requires action across a wide range of technologies and problems to solve. It isn't just about allocating some disk resources to backups. For a significant impact on all three pain areas – speed, reliability and high staffing levels – it requires action across a broad front of the whole end-to-end process.

Author: Rob Nieboer, Storage Strategy, StorageTek
Contact NieboR@AUSTRALIA.Stortek.com