

Three Reasons Why the Traditional Approach to Backup Persists

A couple of weeks ago I was getting a briefing on Atempo Live Navigator regarding its deduplication and near-CDP features that are specifically targeted for desktops, laptops and file servers. But since that conversation, it struck me that CDP and near-CDP technologies have been around for years which got me to thinking. *Why is it that traditional approaches to backup persist even as arguably better approaches to data protection such as CDP and near-CDP struggle to get traction?*

When I use the term “**traditional backup**,” I am referring to the practice of making a copy of production data on a nightly and weekly basis. This is usually done in the context of doing an incremental or differential backup on weekdays and then a full backup on the weekend. While I am not sure exactly how this methodology originated or when, it most likely traces its roots back to when the only effective and economical way to do backup was to use tape as the primary backup target.

But now that disk has effectively replaced tape as the primary backup target in many environments, in my mind to continue with this legacy approach of daily and weekly backup makes little sense to me. While these disk-based backups are “recoverable” in the broadest sense of the term, organizations cannot present this backup image to a server and immediately restart the application from it. Instead they have to first access the backup software and restore the data before they can restart the application or access the file.

While taking these extra steps are not “wrong” per se, they arguably add extra time and effort into the recovery process. Further, depending on what data was lost and how much time has

passed since the backup took place, the data that is recovered may either be unusable or be so old that extra time and effort is needed to recreate the data that has not been protected since the last backup.

This legacy approach to backup fails to capitalize on the many inherent benefits that disk offers over tape from both a backup and recovery perspective. For example:

- **Backup your data continuously or nearly all the time such as what Atempo [Live Navigator](#) does.** This almost eliminates any possibility of data loss since data is backed up every 15 minutes.
- **Minimize network traffic while eliminating backup windows.** The primary reason that my prior employer had a FC SAN with the highest possible network throughput was not because any of its applications actually needed this bandwidth save maybe one. The majority of the time network utilization was in the range of 1 – 5%. Rather it was only during backup windows that network throughput exceeded 30, 40 or even 50%.

CDP eliminates those backup windows since data is backed up continuously (or nearly all the time in the case of Live Navigator) and, while more network traffic occurs during the day, it only occurs when writes occur so it can take advantage of the ample network bandwidth available through the day while also freeing up the bandwidth used during the nightly backup windows.

- **Reduce your data stores even as you improve your recovery point objectives.** One of the myths of CDP technologies is that they consume a lot more storage space than incremental, differential and full backups that are deduplicated. (I wrote about [this myth](#) a little over a year ago.) While they may consume a little more storage capacity, there is nothing preventing companies

from pointing data protected by CDP technologies toward solutions that deduplicate data and, in the case of some solutions like Atempo Live Navigator, it deduplicates data as it protects it. So now you get continuous data protection and reduced data stores.

- ***Find a restore point and recover your data.*** CDP solutions differ in their restore capabilities but some CDP offerings afford users to select a recovery point from within the CDP solution and actually run the application from the CDP data store. Now it likely will not run as well or as fast as it does on your production storage but this recovery option is not even an option in backup software.

So with new technologies like CDP well beyond the beta stage and being used extensively by cloud providers (***R1Soft's CDP solution is HUGE in the cloud provider market,***) it begs the question, why do traditional approaches to backup such as I described above persist? Here are three reasons as I see it:

- ***It works.*** It may not be perfect but using disk in lieu of tape as its new primary backup target in the form of either NAS or VTL has solved or is solving the backup problem as it exists for most organizations for now.
- ***Deduplication has made disk more affordable than tape.*** Disk has come down in price but by itself still is not on par with tape from a cost perspective. But once deduplication is added into the equation and deduplication ratios of 6-7x or greater are achieved, disk becomes more economical than tape as a backup target.
- ***Organizations are still wrapping their minds around disk's potential.*** Backup has been such a big problem in organizations that many are simply taking a deep breath and enjoying the break before turning their focus as to what to do next from a backup and recovery perspective. So for now they are just letting things be until

circumstances (backup windows too short, out of disk space, need for major backup software upgrade, etc.) force them to make a change.

It is for these reasons that I believe the traditional approach to backup has persisted to date. But it is also why the traditional approach to backup is probably on its last legs in the role that it is used in now. As companies understand what they can do with disk and virtualization fundamentally changes how they need to backup and recover their VMs, organizations will have no choice but to select data protection technologies like CDP that are better positioned to provide the functionality that they need going forward.