DynaSis

IT for the C-Suite - Quick Read Series
Modern Business Technology Education for the C-Level Executive

By David Moorman - President, DynaSis

Understanding & Managing RTO / RPO

Setting the parameters of your disaster recovery program is essential to its viability and success. The most critical of these parameters are Recovery Point Objective (RPO) and Recovery Time Objective (RTO). While the two parameters go hand-in-hand in establishing your recovery and business continuity needs, they are really very different and each play an important role in setting the process you will follow in recovery from system failure, data corruption, ransomware, etc.

One important point before we start: in a perfect world, your IT network will be set up so there are no potential failure points. While that sounds great, this isn't based in reality. Equipment fails, software fails, people mess up, cyber-criminals are hard at work, and let's not forget fire, floods, tornadoes, broken water pipes, etc. The real question is: what do we do when disruption happens?

Recovery Point Objective (RPO)

The RPO represents the amount of data your company can afford to lose during a disruption. Of course, it's easy to say "none", and setting this parameter to zero is a real possibility, but that's an expensive undertaking. So in determining "acceptable loss", we have to weigh the cost of lost data against the cost of setting a zero-tolerance loss parameter. More likely, you may determine that it is okay to lose 24 hours-worth of data so your RPO is set to 20 hours, just to be safe. Or it may be set to two hours, or one hour, or five minutes…all are possible. Of course, your managed IT service provider needs to understand the complexities of your network and how to accomplish this…experience here is critical.

Recovery Time Objective (RTO)

A disruption isn't just about how much data you are willing to lose, it's also about how long you are willing to be down. Again, time is money and the more quickly you need to be fully up and running, the more it will cost.

Let's look at two different types of businesses and how they might look to set their RPO / RTO objectives.

Data Input Company (DI): This company collects data from its clients, enters it into a database, then provides this digitized information to the clients that sent it in the first place. The data may be entered by hand or by optical reader or from spread sheets, etc. Here is the important point: the original information from the clients should remain available. If DI's system crashes and a significant amount of data is lost (let's say a full week's data), DI will have either the physical records (maybe paper surveys, return postcards or spreadsheets) from which to re-create the digital records. Records that had been emailed from clients will either be available from DI's email servers or from the clients themselves. The lost files will be rebuilt. DI, of course, will

have to bear the cost of paying its employees overtime, or hiring temporary workers, to accomplish this, but this may well be an acceptable risk.

Online Retail Company (OR): OR is a young, growing business with a growing customer base, and a growing email list. Two days ago, it emailed its customers news about its Spring Sale and the orders have been pouring in. At 2 PM today, its network went down due to an equipment failure and all the orders that had been taken over the past 24 hours were corrupted. Several things happened: although they have no record of the sales that had been made, the money received from hundreds of now unidentifiable orders is in the bank. Yes, each deposit is traceable to a sales receipt…except that the sales receipts have vanished. Number 2: new orders are stopped dead in their tracks because although their site is still up, the data they need to sell product is gone. As we can see, this is a company that can ill afford long RPO or RTO. Every minute that cannot be recovered is money lost, or customers who have paid but who will not be receiving their merchandise. This is a company that needs to be back up and running, with files fully restored, just as quickly as possible.

Types of Data Backups and How They Affect RPO / RTO

Simply put, you cannot restore corrupted or deleted data if you don't have that data saved in another place. Hence: backups. (Note: Don't confuse a "backup" with an "archive". While a backup is a copy, an archive is, or becomes, the "original" file. The data, rather than being simply copied, is transferred to another medium, usually tape, which is less expensive to store long-term. The files that are archived are important to save, but unlikely to be needed in the near-term – possibly ever. Archived files are sometimes backed up onto a second set of tapes and stored in a separate facility. This is recommended.)

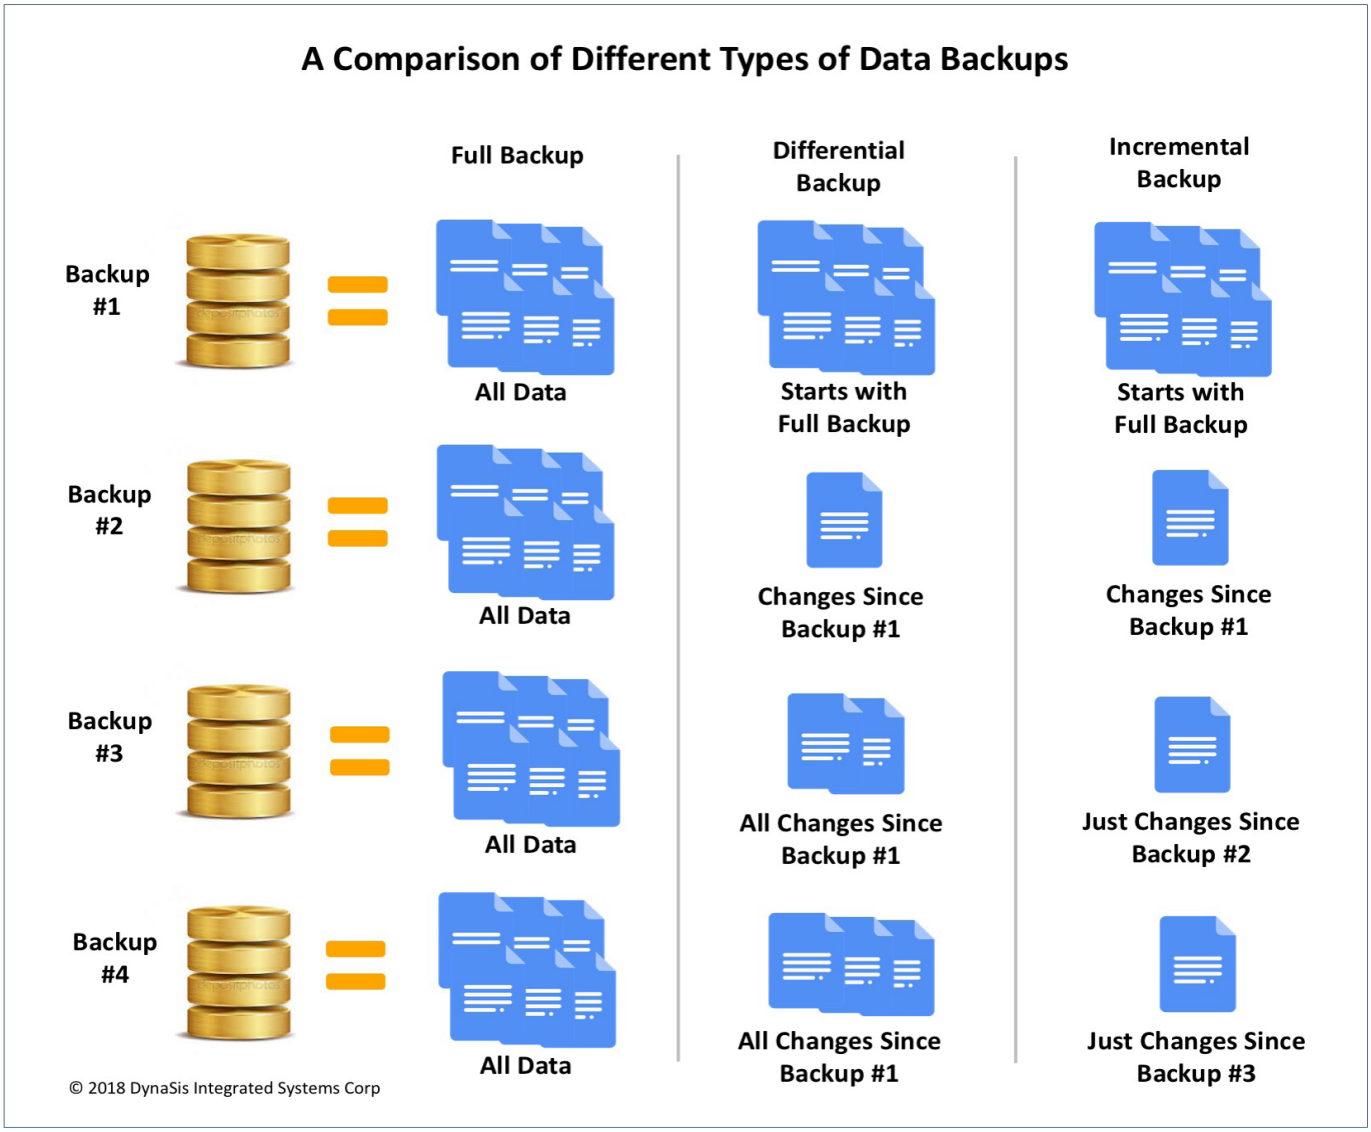

We are going to take a look at the most common types of backups: full backups, incremental backups, and differential backups. As you might expect in today's tech world, there are other variations including synthetic full backups, mirroring, reverse incremental and continuous data protection.

Full Backups

A "full backup" fully backups up all data, whether to tape, disk or CD/DVD. The primary advantage is that since it is a complete copy in a single place, restoration is quick and easy (shorter RTO). The downside is that it takes much longer to perform a full backup than other types – as much as 10 times longer – and requires extensive storage space, and you only have available data since that last full backup was run. Anything since then is lost. Because of these three factors, full backups tend to be run only periodically, which means it is used along with another form of backup, such as a differential or incremental.

Incremental Backups

In an incremental backup, we only copy any data that has been added or changed (including deletions) since the last backup of any kind. For example, if the last backup was a full backup, the incremental will record changes since then. But if the last backup was incremental, then the next incremental backup will only record changes since the previous incremental. The chart below explains how this works. Because there is minimal data recorded in this type of backup, the time expended is short and it is easy to run frequent backups. Periodically, however, it is still important to run a full backup or you end up with a very large number of incrementals that need to be accessed if a restoration is necessary.

Differential Backups

The first time we run a differential backup, it is exactly the same as the first time you run an incremental backup: you copy all the data that has been accumulated or changed since the last full backup. After that, however, each differential backup records ALL changes since the last full backup. As you see in the chart above, with each differential backup, the amount of data you copy grows since you are effectively going back to the beginning (the last full backup).

Which One?

Typically, backups will be done one of three ways:

- Full backup daily
- Full backup weekly with a differential backup daily
- Full backup weekly with an incremental backup daily

Your managed IT support provider will review with you the advantages of each in terms of time, space, cost, performance, protection and recovery. The variables are very much company specific, so we won't go into them here.

Recovery Considerations

The most important recovery consideration when choosing your backup format is, in the event of recovery, how long can you wait until your files have been recovered and are accessible? If you do a full backup daily, you only have a single backup to access, which speeds up the process. However, since you will only have backed up since the previous night, you will lose changes to the files that have been added or changed since then. Your backups may also take quite a bit of time.

If you performed a full backup over the weekend and then did a differential backup daily, you will have to access two backups (which would have taken less time to backup each day), but you will still lose any files added or changed since last night.

An advantage of using incremental backups is that they can be performed as often as you want: daily, hourly, every fifteen minutes. On the plus side, you will lose very little data, maybe none, but the tradeoff is that if you need to recover files, you will have to access all incremental files created since the last full backup to do so. If you do an incremental twice a day, and if you lose data on Friday afternoon, you will have to access the last full weekend backup, plus two backups a day for Monday, Tuesday, Wednesday, Thursday and one from Friday morning. Of course, with proper backup technology set in place by your IT Support provider, this will be highly automated.

As you can see, making this decision may require input from a number of your in-house constituents and certainly involves your in-house team, along with your managed IT support company. Earlier in this article we mentioned variations of the backups discussed here: synthetic full backups, mirroring, reverse incremental and continuous data protection. One of these may be the best solution for your business, and we will be happy to discuss all options with you.

At **DynaSis**, we have been working with small to mid-sized companies like yours since 1992 and have the experience necessary to help you make such decisions. Give us a call at 678-373-0716 and we will explain how our process works.