



Comparing Data Loss Protection Strategies in Large-Scale Data Centers

Customers rely on online backup providers to keep data secure in data centers. Backed-up files must be protected from any threat, including earthquakes, power outages, and hackers. Many of these events will never occur, but every online backup provider deals with one stark reality— hardware components fail. In fact, industry statistics estimate that two to six percent of new hard drives fail, and two to four percent of servers fail in the first year, with failure rates increasing as the hardware ages.

As a result, online backup providers implement redundancy strategies to protect files from being lost due to hardware failure. Redundancy strategies ensure that multiple copies of end-user data are stored so that nothing is lost in the event of hardware failure. Mozy ensures redundancy by using a process called distributed encoding, while other online backup providers use technologies such as RAID and mirroring.

Evaluating redundancy strategies by tracking mean time to repair

When storing files at petabyte (PB) scale, a redundancy strategy's effectiveness can be measured in its mean time to repair (MTTR). MTTR represents the average time required to repair a failed hard drive or server before permanent data loss occurs. This measurement becomes significant in large-scale data centers where thousands of components make failure of a server or disk drive an inevitable, day-to-day occurrence.

A higher MTTR number is more desirable because failed components need to be replaced less often. Low MTTR numbers require tremendous resources to keep up with the failure rate. For example, to protect against data loss, an online backup provider with a 24-hour MTTR must replace failed components every 24 hours. A provider with a 720-hour MTTR must replace components far less often (every 30 days) before data is permanently lost.

RAID-6

Finding inadequate protection with RAID-5 technology, several online backup providers have upgraded to RAID-6. When a RAID-6 server stores a file, it runs an encoding algorithm that adds as much as 20% additional data, called parity, which is saved along with the original file. Then, the file data and the parity data are divided into fragments, which are saved across a drive array of at least five drives. When an end user attempts to restore data, each file is delivered by reassembling the file fragments. This process of storing the fragments on separate drives allows files to be reassembled even if two of the drives fail in the array; however, there are significant drawbacks to utilizing RAID-6 to store files in the petabytes.

In an attempt to reduce costs, some online backup providers use RAID-6 servers attached to 15 x 1 TB drive arrays. This type of configuration generates only 13% parity data, which means that if more than two drives fail in the attached 15-drive array, all files on all 15 drives are lost. In the event of drive failure and subsequent data loss, the only way these files can be reacquired is by performing another backup, however, this has a direct impact on the end user and comes with a 100% probability that customer files will be lost.

Another risk in using RAID-6 configurations for data redundancy is that there is no protection against server failure. In the last few years, an online backup provider lost the backed-up files for over 7,500 customers. According to the provider's official statement, hundreds of thousands of customer files were lost "because a number of RAID servers failed." When RAID servers fail, they often corrupt all their attached data. Their statement explains that they have now "replaced the failed servers with higher-rated servers." This provider believed that by using higher-rated servers, their data loss problems were solved; however, even the highest-rated servers fail 2% of the time. And this company still has no redundancy plan to protect data when a server inevitably fails.

RAID-6 technology is effective in providing data redundancy at smaller scale; however, statistics and real-world experience show an inevitable risk of data loss when implemented at storage ranges in the petabytes. Even a small data center with 100 of the highest-rated RAID-6 servers attached to 15 x 1 TB drive arrays will lose at least 30 TB of customer files per year due to server failure.

RAID-6 might be appropriate for smaller online backup providers managing only a few servers and terabytes, but it is inappropriate and cost-prohibitive when managing data in the petabytes.

Is mirroring a better solution?

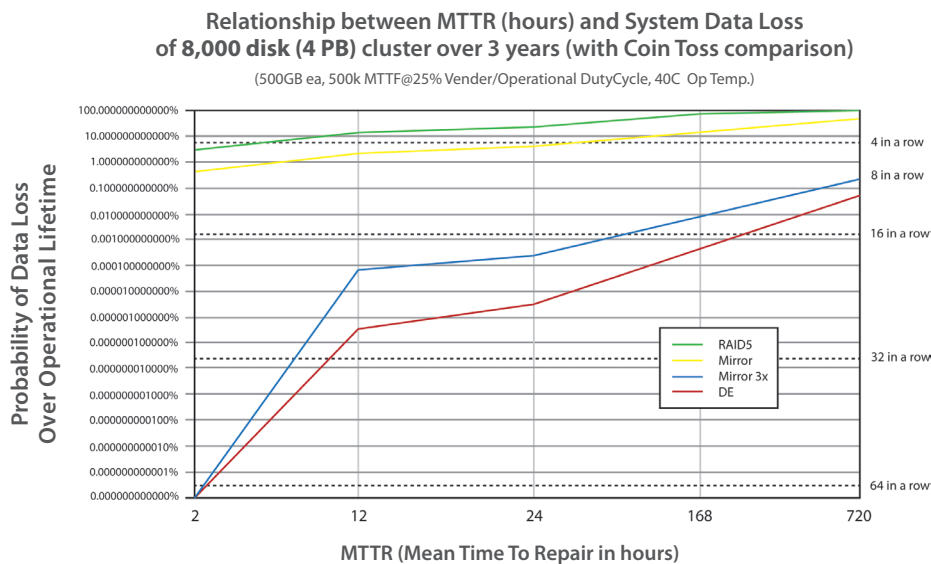
Mirroring is another common redundancy strategy utilized by online backup providers to protect against data loss. Mirroring involves creating and maintaining exact mirrored copies of each stored file in multiple locations. A typical mirroring configuration involves keeping a copy of each file on a separate drive attached to a separate server. When any single drive or any single server fails, any lost files can be recovered from the mirrored copy. But files cannot be recovered if two drives or two servers fail. When this happens, all files on the drive or all files on the server are lost forever.

One of the greatest obstacles for backup providers considering mirroring is that it is expensive. These higher costs are ultimately passed along to the end user in the form of expensive subscription fees. Storing mirrored copies of each file increases the storage cost and hardware requirements by 100%. Even with such a huge investment, mirroring still has a higher probability of data loss when compared to distributed encoding.

Comparing data loss prevention strategies

The following chart compares several different technologies and their mean time to repair (MTTR) in a four-petabyte environment. As the number of MTTR hours increases (X axis), the graph lines show the increasing probability of data loss along the Y axis. On the right hand side of the graphic, a correlation to the chances of tossing a coin and landing on "heads" a number of times in a row is used to illustrate this probability.

For example, if data center devices are replaced more often than every eight hours, distributed encoding (DE) has less than a .0000001 percent chance of data loss. This compares to flipping a coin and having it land on "heads" 32 times in a row.



As you can see from the graph, triple mirroring and distributed encoding (DE) provide the best protection against data loss. As the MTTR increases up to 720 hours (30 days), distributed encoding continues to be the most reliable option.

Probability of data loss in mirroring and RAID environments reach 100-percent probability at around 168 MTTR hours which is the same probability of data loss as getting “heads” as few as four times in a row.

What is not shown in the graph, however, is how mirroring and RAID percentages get dramatically worse as storage amounts increase. These technologies may be sustainable in a data center with only four petabytes of data and highly-resourced to keep ahead of the MTTR. But as a data center grows to store tens of petabytes, the 100-percent probability of data loss happens at a much lower MTTR—and data will be lost.

Distributed encoding delivers the highest level of data protection at scale

Statistical analysis and real-world experience in managing data at petabyte scale reveals that distributed encoding strategies are the most effective at protecting data against drive and server failure.

Mozy uses distributed encoding. Before being saved in the data center, each file is processed using mathematical algorithms to create additional parity data. Similar to RAID-6 encoding technologies, this additional data expands the amount of file information to be stored, but it allows the original file to be reconstructed even if several of its data fragments are lost.

Specifically, Mozy servers divide the encoded data into 12 fragments. Each fragment is saved to a different drive on a different server. At any time, the original file can be reassembled from these individual fragments. This process protects against drive failure because if three drives fail out of the 12, the file can still be reassembled from the remaining nine fragments. And because the fragments are distributed to different servers, as many as three servers can fail out of the 12, and the file can still be reassembled. By encoding file information and distributing data fragments across servers, distributed encoding provides a greater level of protection against server failure when compared to RAID-6 and mirroring.

Distributed encoding has the following advantages over mirroring and RAID solutions:

- Distributed encoding requires a significantly lower MTTR than mirroring and RAID. A lower MTTR translates into lower operating costs and greater protection against data loss.
- Distributed encoding allows files to be reconstructed when as many as three drives fail. Mirroring protects against only one drive failure; RAID-6 protects against only two.
- Distributed encoding lets parity data be regenerated without taking the original data offline. RAID-6 requires extensive time and human intervention to rebuild parity data. Files are unavailable to customers during the RAID rebuild.
- Distributed encoding protects files from three simultaneous server failures. Mirroring protects data from one server failure. RAID provides no protection against server failure.

Conclusion

While servers and drives fail in every environment, Mozy’s use of distributed encoding significantly reduces the chance that customer files are affected by such failures. Files backed up in Mozy data centers are better protected against loss when compared to RAID-6 and mirroring. Distributed encoding ensures that Mozy provides a service that is more secure and more efficient than other online backup providers running RAID-6, mirroring, or any other data protection solution. As a result, customers can rest assured that even in the event of server or drive failure, their files are safe, intact, and available.