

Backup Over 2TB: Best Practices for Cloud Backup and DR with Large Data Sets

EXECUTIVE SUMMARY

Most cloud backup doesn't work for data sets above 2TB. A widely held perception is that slow data transfer performance causes nightly backups to take over 24 hours and makes timely DR impossible.

Unacceptably slow cloud backup performance can be a major issue for companies with large and growing data sets. Continuing to support tape backup is too risky, local disk is very expensive, and consumer-grade cloud backup & DR solutions don't deliver the performance required to protect data sets up to 100TB.

To accomplish this you need an enterprise-grade cloud backup & DR solution like what Zetta.net provides. It enables companies to backup large data sets and backup & recover up to 5TB in 24 hours directly to/from the cloud. This means that companies with an average daily change rate of 2% can successfully deploy cloud backup for data sets up to 250TB.

Tape Falters with Large Data Sets

Tape is prone to physical problems no matter the size of the backup. Tape drive failure, media corruption and damage, and accidental reformatting and erasure combine to make tape-based backup for big data a real gamble.

Tape for large data sets is also expensive. The bigger data grows, the more tape drives you must buy to handle the backups and the more media you must buy and store. Even with this investment, big data backups to tape overrun backup windows and time limits. Recovery takes even longer.

When dealing with backing up large data sets to tape, you need to ask three questions:

- How long the backup will take?
- How many tapes will you need for a full and incremental backup?
- Where are you going to put those tapes when your backup has completed?



Real World Example

Let's look at an example of backing up 2TB uncompressed backup to LTO-5 tape:

This process will require 2 tapes for every full backup. If the daily rate of data change is 25%, we are looking at an additional 500 GB of incrementally changed data that needs to be transferred to tape every day.

Best practices for data protection include swapping out those tapes daily, leaving you with 8 LTO-5 tapes used every week. Good practices also require you to keep secondary copies off-site, so you transport daily to the off-site vault.

Let's say that on day 6 you encounter a disaster and need to recover your data. Your IT team makes an emergency call to your offsite vault and has an out-of-cycle tape delivery made to your office. You are now 5 hours into the recovery: 1 hour to diagnose the issue and 4 hours to get the tapes on-site.

Once the tapes arrive you load up your last full backup and restore your 2TB of data. The full backup restores take 4.5 hours and 6 day's worth of overlaying incremental backups take an additional 5 hours. Only then, 15 hours later, do you have your data from its tape backup.

Has this long process met your Recovery Point Objective (RPO) and Recovery Time Objective (RTO)?

Probably not.

The costs of backing up large data sets to tape are also sobering:

Product and Process	Cost
LTO-5 Tape Cost	\$30 * 8 = \$240 / week
Tape Drive Cost	\$3,000
Tape Offsite (8 tapes) charge	\$20 / Month
Backup Software	\$15,000
Emergency Recovery Charge	\$300
Data Center Costs	\$200 / Month
Hardware Support Costs	\$600/year
Software Support costs	\$3,000/year
Server Engineering time	\$20,000/year
Tape Risk	Significant

Zetta.net Cloud Backup & DR for Large Data Sets

Is cloud backup an improvement on tape for big data sets?

Yes, but not all cloud backup providers are created equal.

Cloud backup vendors often tout their solutions as improvements on tape. But not all cloud solutions are created equal, and performance is the Achilles heel of cloud backup. Cloud backup providers with consumer-grade, prosumer, and SMB-level backup can work fine in smaller environments with, for example, 100GB to protect per day.

This seems to work fine until you calculate in the sheer size of large data set backup. For 2TB, an initial backup seed would take 20 days. Following that, even an average 6% daily change rate, would lead to nightly backups that take 28.8 hours. Clearly that will not work.

In contrast, Zetta.net's enterprise-grade cloud backup solution can backup up to 5TB a in 24 hours directly to the cloud. At this rate the initial baseline upload of 2TB takes about 9.6 hours. So at an average 6% daily change rate, incremental backups take about 1 hour a day.

That's a big difference.

Zetta delivers an enterprise-grade cloud backup and disaster recovery (DR) service that includes local and cloud backup, archiving and DR. Zetta's solution includes patent-pending WAN-optimization technology that enables backup and recovery speeds that can outperform D2D backup in-house. Zetta does this without a local backup appliance -- which reduces the overall cost of backup and eliminates the single point of failure they represent.



The key to Zetta's speed is its cloud-native architecture. Zetta specifically architected for the Internet with built-in WAN traffic optimization. Multiple patent-pending technologies move data faster with byte-level change detection, advanced compression, and multi-threading. Zetta's enterprise-grade data centers further enhance performance by serving over 80% of cloud storage reads from RAM or SSD cache instead of spinning disk.

Zetta encrypts all backups in-transit and at-rest using 256-bit SSL and Salsa20 cyphers, and uses file level hashing to validate data. Its cloud storage is built with RAIN6 N+3 for extremely high availability without performance lags.

The Elements

o Zetta Smart Cloud

Smart Cloud integrates backup, snapshots and replication without any local appliances. Snapshot and replication reside within Smart Cloud, enabling customers to easily restore their data from the cloud or their local copy.

o ZettaMirror

ZettaMirror software is a lightweight agent that replicates data and efficiently uploads it over the optimized WAN. Lean Local Data adds the option of a local network copy for the fastest possible restore of critical data. ZettaMirror supports an unlimited number of application-specific software agents and plug-ins for SQL, Exchange, VM, System State, and many more.

o Web-based Management

Online management is efficient and effective. Browser-based recoveries use a custom URL to restore from replicated data copies in Zetta's online data centers. A partner portal adds value to MSPs.

o 24x7 Managed Service

Zetta's 24x7 managed service proactively monitors customer data. When you need help, the US-based support number gets a Zetta engineer, not an automated menu.

NUTS AND BOLTS:

Why Zetta is Right for Big Data Sets

- Manages data sets up to 100TB in size.
- Tunes each process with individual clients to optimize data transfer.
- Processes only changed file data.
- Enables side loads to speed up large initial loads.
- Offers a wide range of app-specific plugins to optimize backup by application needs.
- Recovers up to 5TB a day; once IT launches a recovery, Zetta automates the process.
- Snapshots enable flexible RPOs.
- Dynamic TCP window sizing, or chunking, sends data in variable sizes over the WAN for dramatic speed gains.

Zetta vs. Tape in the Real World: Kineto Wireless Reduces Backup Costs by 50%



Steve Robey, the IT Director at mobile messaging software firm Kineto Wireless, is responsible for protecting 30 application servers. He found it increasingly awkward and expensive to backup 3.5TB of data to tape in any acceptable timeframe. The tape backup system also came with high media and tape drive costs, and took considerable IT time to cycle and manually load tapes.

Steve looked at cloud backup products but was put off by the investment in new hardware appliances. After reviewing Zetta.net's cloud-based backup and DR platform he was sold. Not only was there no hardware investment, but the package price included support, maintenance and service that had eaten up so much of the tape backup budget.

A few weeks after migrating the server backups to Zetta.net, the Kineto experienced a serious storage cluster failure. He was able to quickly and quietly restore with no loss of data. Today he reports that Zetta has saved the company 50% in capital and operating expenses over the tape solution.

ZETTA.NET: Enterprise-grade Cloud Backup & DR

Data is only going to continue growing. When you backup large data sets to tape then you spend high amounts on hardware, datacenters, bandwidth, software, software support, hardware support, off-site storage, media costs, and other charges. You also have operational issues such as throttling backup to match tape ingestion speeds, and the need to support legacy LTO tape drives.

Zetta's pricing is all-inclusive; it does not throttle backups, and has no legacy drives to support.

With Zetta, even large data backups take a few hours instead of days.



Zetta's enterprise-grade cloud backup and DR offers high performance for large data sets and meets strict regulatory compliance requirements. Pricing starts at \$700 a month for 2TB of backup data, including unlimited licenses of award-winning backup & DR software, secure cloud storage, and 24x7 US-based engineer-level support.

What Industry Analysts & Press are saying about Zetta.net:

"Zetta provides the combination of on-site backup and efficient, cloud-based DR that's required to fully protect critical applications."

– Storage Switzerland



"Zetta integrates cloud offsite backup, archiving and disaster recovery in a single managed service that can be deployed in a few minutes."

– InfoStor



"Zetta's customers typically have 2-40TB of capacity; this compares to target capacity of less than 3TB for many of its competitors."

– The 451 Group



Industry Awards



CONTACT US

Please ask us about how you can backup your large data set – up to 100TB – to the cloud. Get Custom Pricing. Start a Free Trial. Email Us: sales@zetta.net Call Us: 877-467-3882

Zetta Inc. is an award-winning provider of enterprise-grade 3-in-1 backup, disaster recovery, and archiving.

