


CASE STUDY

Cloud Services
Intel® Optane™ SSDs



Improve Tape-Backup Speeds by Caching with Intel® Optane™ SSDs

Baidu needed faster caching for its tape-backup system, and found a solution with both speed and endurance.



“Baidu has encountered bandwidth and write endurance challenges in processing new client data and in duplicating tape library backup data in parallel. Intel® Optane™ SSDs bring great value with high bandwidth, high endurance, low latency, and easy maintenance, which helps Baidu accelerate our tape-library solution innovations while reducing cost and improving efficiency.”

—Baidu Binghe Tape Library Archive Storage Group representative

Internet giant Baidu backs up huge amounts of cold data on magnetic tape libraries. Tape can write data fast, but the data first needs to be sequentialized on a cache disk. Baidu was using hard disk drives (HDDs) to cache the data, but the HDDs could not feed data as fast as the tape could write it. Baidu evaluated three options for replacing its HDD cache disks with faster technology—two options based on NAND flash solid state drives (SSDs) and one option using Intel® Optane™ SSDs. Baidu found that the solution using Intel® Optane™ SSDs met the speed requirement using fewer SSDs than the NAND solutions, and it provided a high level of endurance that the NAND drives could not.

Massive Amounts of Cold Data are Stored on Tape Libraries

In recent years, businesses have been amassing enormous volumes of data by digitizing and tracking more internal processes, collecting more information about customers, and building more data-intensive applications than ever before. As the amount of data generated and gathered continues to grow, saving and managing this flood of information has become increasingly difficult.

Simply saving all data in one storage medium is impractical. Different storage media have different levels of cost, capacity, performance, and endurance that make each suitable for different types of data. A useful way to identify the best medium for storing data is to first classify the data as “hot,” “warm,” or “cold.” The coldest data is a good candidate for archival in a long-term, inexpensive storage medium such as magnetic tape.

It turns out that most business data can ultimately be saved as cold (as shown in Figure 1).¹ For this reason, it's essential to choose a storage medium for cold data that not only is inexpensive, but that also can write quickly enough to keep up with the fast rate at which this type of data is generated. Fortunately, thanks to technologies such as Fibre Channel and Serial Attached SCSI (SAS), modern tape libraries can write quickly, making them an excellent choice for the fast archiving of cold data today.




Figure 1. Most business data that is stored can be considered “cold”¹

Fast Tape Libraries Need Fast Cache Disks

The sequential read and write performance of tape libraries continues to improve to help keep up with the growing amounts of data. But even though tape libraries themselves write quickly, there is a technical hurdle in a tape-library-based backup architecture that often prevents data from being written to the library at its maximum rate. Before data is backed up to tape, it first needs to be collected on data nodes in a caching tier, where the data can be prepared (sequentialized) for writing to tape media, as shown in Figure 2.

The SAS RAID cache-tiering disks in these data nodes often cannot perform reads and writes fast enough to saturate the pipeline to the tape libraries, resulting in a bottleneck and a slow backup rate that doesn't live up to the high-speed potential of tape libraries.

Tape Library Backup Solution Architecture

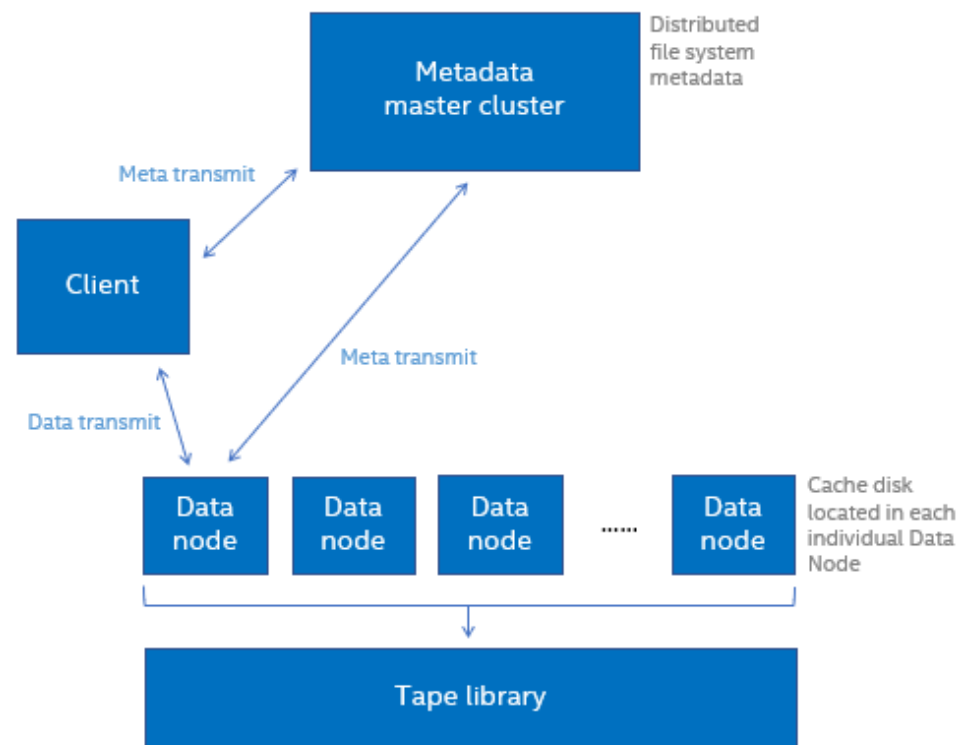


Figure 2. Tape libraries require data nodes to sequentialize data and prepare it for backup. This architecture can manage multiple data streams, including incoming new data, access existing data, and remove old data.

Baidu Evaluates Options for Fast Caching

Baidu was facing a caching bottleneck problem in its tape backup architecture. Two scenarios in particular stressed the system and exposed the bottleneck:

- The company's autonomous car solution was generating an enormous amount of data from Internet of Things (IoT) sensors installed throughout a fleet of 300 vehicles, which had logged more than two million kilometers at the time. Most of this IoT data was classified as "cold" and was sent for long-term storage on tape via the data nodes.
- The company's retention policies called for data stored on tape to be moved onto new tape every 3-6 years, which required the use of cache disks. The burden of this caching workload—constantly writing and reading a growing amount of archive data—was uncovering the inefficiencies of the system.

...tape libraries could write at a rate of 600 MB per second (MB/s), the HDDs could only read at a rate of 200 MB/s...

In the cold data tier, Baidu's data nodes used hard disk drives (HDDs) for caching before archiving to tape. Although the tape libraries could write at a rate of 600 MB per second (MB/s), the HDDs could only read at a rate of 200 MB/s, effectively reducing the potential backup speed by two-thirds, as shown in Figure 3.

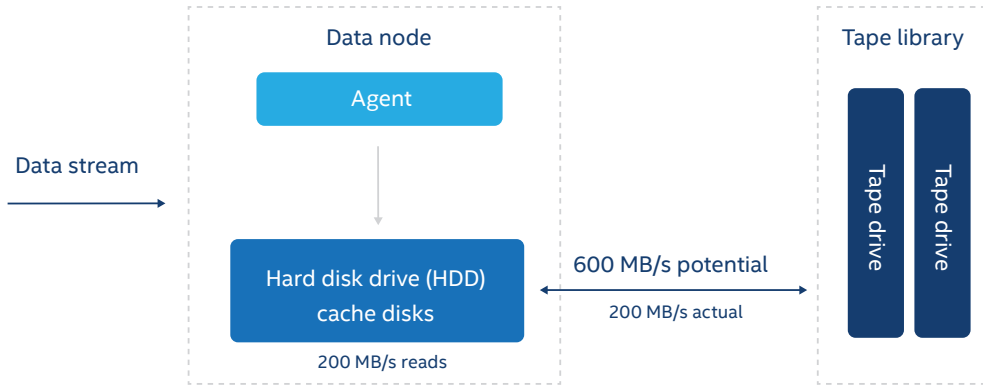


Figure 3. The cache disks in the old Baidu backup architecture caused a bottleneck

As the Baidu team considered alternatives for its cache drives, it made note of the data-tier requirements:

- 1. Read-speed requirements:** The replacement cache disks needed to match the 600 MB/s write speed of the tape library with a read performance of 600 MB/s, even while experiencing significant write pressure.
- 2. Simultaneous write-speed requirements:** In order to keep the tape write speed saturated for cost efficiency, the disks needed to be able to write new data into cache at least as fast as the cached data could be written to tape, namely 600 MB/s minimum.
- 3. Endurance requirements:** The storage capacity of the tape library is 16 PB, and the Baidu team wanted the caching disks to last for three generations of tapes. To meet this goal, the caching disks would need to support endurance of 48 petabytes written (PBW).

Intel Optane SSDs meet Baidu's caching requirements

- Achieve 600 MB/s read performance while under a 600 MB/s (random) write workload
- Maintain consistent read response times regardless of the write pressure applied to the drive
- Backup 3x as much data with high endurance

Investigating NAND-Based Solutions

The Baidu team first investigated using NAND-based SSDs as the replacements for the HDD-based caching disks on the data nodes. A limitation of NAND-based SSDs is that random write operations require an immense amount of background media management that can significantly slow throughput per disk and shorten disk lifespan. Therefore two configurations of multiple NAND drives were evaluated: a 16-disk RAID0 standard-endurance configuration, and a 6-disk RAID0+1 medium-endurance configuration.²

Both proposed NAND solutions supported the read and write requirements of 600 MB/s, but they required too many NVM Express (NVMe) slots to achieve the needed throughput rate, which complicated maintenance. In addition, calculations based on equipment specifications showed that both these configurations failed to meet the endurance requirements, offering 30.72 and 36.75 PBW, respectively.

Meeting the Challenges of Speed and Endurance

The Baidu team found that it was able to meet the data caching requirements of the solution by using Intel® Optane™ SSD DC P4800X drives as the caching disks on the data nodes. By using Intel Optane SSDs (as shown in Figure 4), the solution achieves 600 MB/s read performance while under a 600 MB/s (random) write workload. And unlike with NAND-based SSDs, a huge advantage of Intel Optane SSDs is that they maintain consistent read response times regardless of the write pressure applied to the drive.

Finally, by offering 164 PBW of data endurance, the solution also far exceeds the endurance requirements of 48 PBW. With this new solution, Baidu is able to back up three times as much data in the same amount of time.

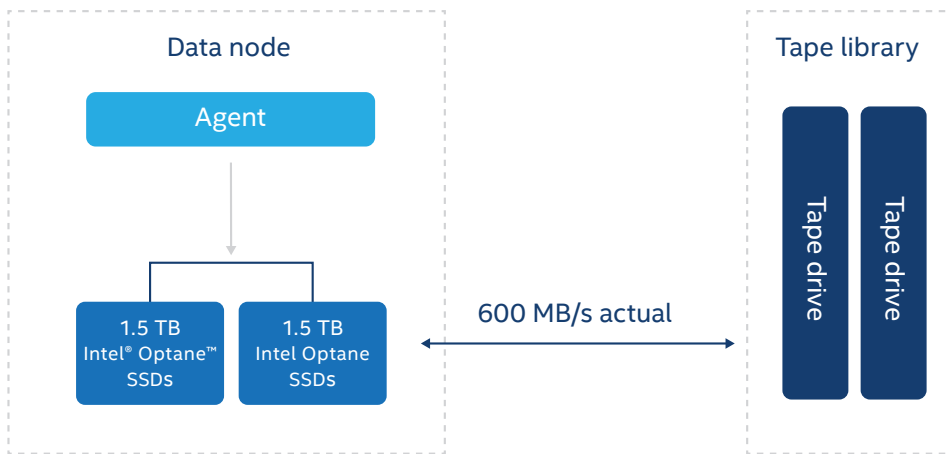


Figure 4. Using a mirror of Intel® Optane™ SSDs as caching disks allows Baidu to meet its backup solution requirements

Intel Optane SSDs for Fast Caching with High Endurance

Businesses are under growing pressure to find ways to back up cold data as quickly as it is being generated. Tape libraries deliver excellent backup capacity and write speed at a low cost, but they require data to be cached and sequentialized on a data caching node before they are written to; and it is on these caching nodes that a bottleneck often forms. By using Intel Optane SSDs as caching disks on these data nodes, Baidu was able to meet a heavy read/write mixed-workload bandwidth requirement with excellent endurance. And with this solution, Baidu is now able to back up the same amount of data in 67 percent less time than with its previous solution.⁴

When used in backup solutions, Intel Optane SSDs for the cache tier help maintain stable and consistent read performance (600 MB/s) while under heavy write pressure (1 GB/s—higher than the 600 MB/s required by Baidu).³ Meanwhile, the extremely high endurance (164 PBW and 60 drive writes per day [DWPD]) of Intel Optane SSDs helps ensure a long lifetime for the solution.³

As this case study shows, Intel Optane technology is being used in innovative ways to solve key problems related to the exponential growth of data.

Learn More

www.intel.com/optane



¹ IDC's assumption for forecasting is that 40% of data being captured is commercial data, and of that 60% is cold data. Source: IDC White Paper, sponsored by LTO Consortium, Tape and Cloud: Solving Storage Problems in the Zettabyte Era of Data, June 2019. [fujifilmusa.com/products/tape_data_storage/case_studies/pdf/Tape-and-Cloud-Solving-Storage-Problems.pdf](https://www.fujifilmusa.com/products/tape_data_storage/case_studies/pdf/Tape-and-Cloud-Solving-Storage-Problems.pdf).

² **Configurations:** *standard-endurance configuration:* 16 Intel SSD DC P4510 for 30.72 PBW; *medium-endurance configuration:* 6 Intel SSD DC P4610 for 36.75 PBW.

³ 164 PBW is a specification for the 1.5 TB Intel Optane SSD DC P4800X and Intel Optane SSD DC P4801X. Source: Intel. "Breakthrough Performance Expands Datasets, Eliminates Bottlenecks." January 2019. [intel.com/content/dam/www/public/us/en/documents/product-briefs/optane-ssd-dc-p4800x-p4801x-brief.pdf](https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/optane-ssd-dc-p4800x-p4801x-brief.pdf).

⁴ Based on equipment specifications, the previous solution using a Hard Disk Drive could only feed sequentialized data to the tape library at a rate of 200 MB/s, while the new solution with Intel Optane SSDs could feed the data fast enough to meet the 600 MB/s requirement.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details.

No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.