



IBM Storage Introduces its Diamondback Tape Library

By Jean S. Bozman

In a world of dynamic computing and distributed data, it turns out that the tape marketplace – with a presence that is perhaps surprising to some industry observers – is not dead.

Worldwide data is growing rapidly – but not all of it needs to be viewed in real-time. Tape libraries give customers, data centers and cloud providers an option to store data off-line for later retrieval, as needed. Industry data shows that tape, as a storage category, is growing – especially for archiving infrequently accessed cold data, and as a less expensive alternative to higher-cost storage resources.

High-capacity tape systems store data for long periods of time (months or years), supporting customer savings for power and data-center real estate. Research data from International Data Corp. (IDC) shows that shipments of tape drives, tape libraries and virtual tape libraries increased 10.5% worldwide in 2021, year-over year. Revenue for these categories is growing, but not as quickly as unit shipment and capacity growth.

Customer concerns about ransomware, data-storage compliance requirements worldwide – and the need to store “immutable” copies of data that cannot be changed by cyberattacks – drove much of the growth in recent years.

The Open Compute Summit

OCP, the acronym for Open Compute Project, is gaining traction as customers’ enterprise applications and databases are now running alongside cloud-centric workloads in open, distributed data centers, or data hubs in the cloud or at the edge.

At last week’s [global OCP Summit](#) in San Jose, California (Oct. 18 - Oct. 20), it was clear that data centers and cloud providers – including the large CSPs that build their own infrastructure – are looking for ways to consolidate multiple data silos. They do this to increase data density, while keeping their storage systems cool through inventive designs supporting efficient air-cooling to improve data-density thermal characteristics. The IBM Diamondback was on the show floor at OCP, to demonstrate its design, its capacity, and its power/cooling capabilities.



IBM Announces New LTO Tape Drives

On Oct. 19, 2022, IBM, one of the major competitors in the worldwide tape-storage marketplace, announced its [IBM Diamondback Tape Library](#) – the latest addition to its line of tape archive systems.

The new Diamondback product is aimed at two overlapping market segments – the largest global hyperscale cloud service providers (CSPs), and a new wave of “enterprise hyperscalers” working with data sets spanning hundreds of petabytes. IBM leveraged input from both types of customers as it designed the Diamondback tape library, giving it an expanded TAM (total available market) compared to the traditional installed base. In the future, mid-size CSPs and co-lo services could be in the next group of customers, mostly because adding tape libraries to the overall storage mix can bring down total data-storage costs -- even in the largest cloud providers' massive data centers.

Key Characteristics

The key benefits of tape libraries like IBM Diamondback are support for high-capacity offline archiving and the fact that they can store immutable copies, which can increase protection from ransomware attacks, which cause millions of dollars of damage for each individual incident.

By preserving data offline, leveraging “air-gapping” that isolates the data from daily updates, the tape libraries are designed to reduce data loss – and the time-to-recovery to quickly restore online data resources based on all-flash arrays with solid-disk drives (SSDs) or hard-disk drives (HDDs).

The Diamondback tape libraries support 27 petabytes (PB) in a single OCP-standard rack – a high-density enclosure that is widely used in the open-systems world. Multiple racks – with as many as 15 racks in a single Diamondback tape library – can be supported in a single, connected “cluster” of tape libraries.

Features of the Diamondback storage library include:

- Capacity: Up to 15 drives per library unit, with up to 28.5 PB per drive
- Designed for self-service, reducing the need for hardware tools to access the drives and the media. This also allows new IT staffer roles to maintain the systems.



- Supports multiple rows of LTO-9 tape drives, and “Redundant Arrays of Independent Libraries” configurations of multiple LTO libraries (known as RAIL).
- Supports systems running Linux, Microsoft Windows, or a range of Unix systems (e.g., IBM AIX).

Cloud Architects Analysis

The worlds of enterprise data centers and cloud computing are converging, reducing the number of hardware and software “silos” across customers’ data estates.

Tape libraries preserve data, taking it offline for archiving and recovery purposes. In an age of ransomware and cyberattacks, tapes bring the added value of ensuring that mission-critical data (mostly block and file data) and cloud-centric data (mostly object data) will be restored quickly and accurately. Recovery windows are shrinking as most applications are becoming accessible, 24 x 7 x 365 in an always-on world.

Tape libraries also provide an alternative to disk-drives (based on SSDs and HDDs), which require more power/cooling and data center “real estate” than tape storage resources.

IBM’s longtime role as a provider of tape storage systems ensures its role on the short-list of tape providers. IBM continues to compete with a handful of providers in this space, including Quantum, HPE, Dell Technologies and Oracle (via Storage Tech tape drives).

Yet the prospects for tape usage are bright, given accelerating cloud migrations and the use of tape drives by hyperscalers (large-scale cloud service providers) and hundreds of smaller CSPs worldwide for archiving and long-term data preservation. Tape drives can preserve data for up to 30 years, retaining it for later use and accounting purposes.

One other factor favoring future tape growth: Sustainability and ESG (environmental, social, and governance) considerations are gaining importance in the data center technology market, and tape systems excel at reducing overall heat, cooling, and energy costs.

Over long periods of time – measured in years – the savings become substantial, compared to keeping financial and business data online. In our opinion, tape libraries will persist in data centers, driving capacity increases per/cartridge to meet customers’



requirements to meet governmental compliance, data privacy and long-term archiving of large data stores.

Given the rapid rise of data at Edge locations (e.g., factories, retail store locations, smart-city infrastructure), total data worldwide will continue to grow through 2030 – and beyond. IDC estimates that worldwide data (including core data centers, the Cloud, and the Edge) will grow to 175 ZB (zettabytes) by 2030.

The rapid growth of data, worldwide, along with customers' broadening definition of which datasets are mission-critical, will make tape libraries a useful and cost-effective alternative to other categories of data-storage resources that drive higher acquisition and maintenance costs. As such, we expect that tape libraries will complement all-flash arrays and HDD drives for years to come, offering a way to reduce total data-storage costs over time.

Copyright 2022, Jean S. Bozman, All Rights Reserved.