



# Top 10 Questions

## About the Active Archive Appliance

PowerFile's Active Archive Appliance™ (A3) is an industry first in bringing enterprise data reliability, scalability and utility, yet definitively better economics and peace of mind in long term data access through the harnessing of consumer-proven optical storage underpinnings in its subsystem. With any new approach to technology, especially one that uses something as common as Blu-ray, it is expected that there will be questions driven to some extent by misconceptions and assumptions. We have attempted to pick the top ten questions about the A3's approach to solving the archive storage problem and address them here.

**Q:** HOW DOES POWERFILE MAKE OPTICAL MORE RELIABLE?

**A:** PowerFile adds tremendous value in the area of both short-term “write” reliability and the very long term shelf life reliability of optical storage. In addition to extensive qualification of “archive-grade” media and working directly with media manufacturers to ensure PowerFile customers receive the highest, most consistent quality of media available, PowerFile has developed patented and patent-pending technologies to create a storage system that is at least 10 times more reliable than enterprise tape or disk-based RAID<sup>1</sup>.

- *“Archive Grade” Media* – PowerFile only qualifies “archive grade” media that meet the following criteria:
  - Employs durable coating technologies to protect against scratches, fingerprints, dust and other contaminants
  - Leverages dual reflective technology and top-quality materials for maximum resistance to oxidation and highest precision in reading and writing of data
  - Demonstrates extraordinary resilience in age acceleration chambers using high temperature, sodium lights and very high humidity over time
  - Carries a minimum warranty of 50+ years
- *Direct Manufacturer Relationship* – PowerFile has direct relationships with media manufacturers to ensure that all qualified media sold by PowerFile is the highest, most consistent quality media available.

- *Extensive Media Qualification* – PowerFile performs extensive media qualification testing above and beyond any testing performed by the manufacturer to ensure that all media qualified by PowerFile exceeds published data integrity rates.
- *Patent-pending EVAST™ Data Integrity Scheme* – PowerFile has developed a patent-pending data integrity scheme called EVAST (Extended Verification And Self-healing Technology) that optimizes block/stripe size and “zones” the ECC layout to increase data integrity by six orders of magnitude. With EVAST, the A3 is able to boast a calculated<sup>2</sup> bit error rate of no more than one block error for every exabyte of data written. Even more amazing is that EVAST accomplishes this with minimal impact to usable capacity and performance all while retaining the standard Universal Disc Format (UDF) that ensures readability in any standard Blu-ray drive.
- *Patented “Frictionless” Media Library* – PowerFile has developed a patented media library that provides a sealed environment with no fans, very low power consumption and a frictionless mechanism to protect against light, dust, extreme temperatures, and improper handling of media for the life of the data.

**Q:**

I AM PLANNING TO ARCHIVE DATA “FOREVER” AND WANT TO AVOID BEING “LOCKED” TO ANY ONE VENDOR. HOW DOES POWERFILE ASSURE ME THAT I WILL BE ABLE TO ACCESS MY DATA MANY YEARS FROM NOW?

**A:**

When building an archive, one must consider whether you can get to your data decades from now. Change has been a constant in the world of enterprise data storage over the past 20 years and will likely continue. Tape, the traditional medium for backups, has had dozens of formats in the last two decades and there are at least 10 formats being offered by 12 manufacturers today. Attempts to provide an enterprise tape standard such as Digital Linear Tape (DLT) and Linear Tape Open (LTO) have only resulted in compatibility issues between manufacturers. Add proprietary backup software to the equation where every manufacturer stores the data in a proprietary format, and it becomes clear that archiving data on tape is an extremely risky proposition.

PowerFile, after very careful consideration, made the choice to build its A3 on top of standard Blu-ray for three primary reasons:

1. **Standard Media Format** – Almost 1 billion DVD drives have been shipped, making DVDs one of the most ubiquitous media formats available today. The same is expected for Blu-ray as it replaces DVD as the standard medium for home video production
2. **Standard File Format** – PowerFile writes discs in the same ISO standard Universal Disc Format (UDF) employed in standard CDs, DVDs, and Blu-ray discs, ensuring that any disc written in the A3 subsystem can be removed and read in a standard drive thus eliminating the risk of vendor lock-in.
3. **Forward and Backward Compatibility** - Blu-ray continues a long and consistent tradition of full forward and backward compatibility. Blu-ray drives have three reading mechanisms in order to read CD, DVD, and Blu-ray discs. With close to 25 years of compatibility to date and a roadmap showing compatibility into the future, this format standard holds great promise versus all other media options. No other data format in history can boast this same track record over 25 years.

<sup>1</sup> Metric: Unrecoverable Errors on Read (UERS). Dual-parity redundant arrays of independent disks (RAID 6) and enterprise-class tape (LTO-3) have been estimated to deliver UERS of 10-17. PowerFile has tested the Active Archive Appliance (A3) patent-pending error correction scheme to have a 99.99999% probability of delivering UERS of 10-18.

**Sources:**

“Empirical Measurements of Disk Failure Rates and Error Rates”, Jim Gray, December 2005 (<http://research.microsoft.com/research/pubs/view.aspx?type=technical+report&id=1022>);  
 “LTO: a better format for midrange tape”, G.A. Jaquette, IBM Journal of Research & Development, July 2003 (<http://www.research.ibm.com/journal/rd/474/jaquette.pdf>)  
 “Advanced RAID”, IBM Almaden Research, 2006 ([http://www.almaden.ibm.com/StorageSystems/Advanced\\_Storage\\_Systems/Advanced\\_RAID/index.shtml](http://www.almaden.ibm.com/StorageSystems/Advanced_Storage_Systems/Advanced_RAID/index.shtml))  
 “Enterprise RAID 6” Technology Update White Paper, Winchester Systems, 2006 ([http://www.winsys.com/whitepapers/pdfs/Enterprise\\_RAID\\_6\\_White\\_Paper.pdf](http://www.winsys.com/whitepapers/pdfs/Enterprise_RAID_6_White_Paper.pdf))

**2 Assumes the following:**

- Systems operate in an enterprise datacenter environment with temperatures ranging from 68-77 degrees Fahrenheit (20-25 degrees Celsius) and 40-55% relative humidity (Source: ASHRAE, 2007)
- Strict adherence to PowerFile media handling and storage procedures
- Media and drives do not have unexpected manufacturing variations

**Q:** AREN'T OPTICAL-BASED SYSTEMS TOO SLOW FOR ENTERPRISE DATA STORAGE?

**A:** Optical storage, in and of itself and absent the A3 system, likely is too slow for most enterprise applications. Poor performance is one of six reasons why optical storage has become mostly irrelevant in enterprise storage over the last five years. (For more information, see the PowerFile White Paper entitled "New Life For Dead Storage Technologies: Consumer Driven Optical Storage in the Data Center")

PowerFile has designed the A3 specifically to address the performance limitations of optical. The A3 has an intelligent, dynamic caching engine fronting the backend optical storage that significantly boosts performance on both the ingestion and retrieval of archive data. Except for extreme cases, system speed on ingestion and retrieval of most frequently accessed files is in milliseconds. Under a worst case scenario where a file that has not been accessed in many years is requested, the wait time will be less than 30 seconds – a minor trade off for infrequently accessed data when compared to the high maintenance costs of spinning disk and the retrieval time of offline tape.

**Q:** IF I MIGRATE ALL OF MY COMPANY'S FIXED CONTENT OFF OF EXPENSIVE PRIMARY STORAGE AND COMPLETELY REMOVE IT FROM THE BACKUP CYCLE BY PLACING IT ONTO THE A3, HOW DO I CREATE A COPY OF THIS DATA FOR DISASTER RECOVERY PURPOSES?

**A:** The A3 offers two options for disaster recovery depending on the customer environment:

1. Customers who want to completely eliminate the need to handle media can leverage PowerFile's optional Touchless DR™ WAN Replication software that automatically migrates copies of data from a primary A3 to secondary A3(s) located remotely. Should the primary location become unavailable, users can be redirected to the replicated copy of archived data at the remote location(s).
2. For customers looking for a more cost-effective solution, the A3 also offers the ability to replicate volumes within a single appliance. The replicated volume can then be exported and taken offsite for disaster recovery purposes. Every disc within the replicated volume contains sufficient metadata to completely recreate the entire volume including directory structures and security permissions.

**Q:** WHAT HAS POWERFILE DONE TO OVERCOME THE LESSER DENSITY OF OPTICAL?

**A:** First and foremost, PowerFile has designed a low-power, rack-optimized media library that efficiently packages 200 discs, two read/write drives, and all necessary robotics into a compact 5.5U rack mount package. Using Blu-ray technology, this correlates to 126TB of native usable capacity per 42U rack. In addition to an efficient hardware design, PowerFile's optional Adaptive Data Reduction (ADR) software can be added to further reduce the data footprint up to 7x.

When considering the cost of floor space for IT equipment, one must also consider the cost of electricity for power and cooling. The Info Pro estimates that storage consumes 40% of all electricity within the datacenter and Emerson Network Power expects 97% of all datacenters to be maxed out on power by 2011. While hard disk based storage systems may offer more storage per rack, providing sufficient electricity to power and cool racks full of disk drives is becoming an increasing challenge in datacenters worldwide. Since the A3 was designed primarily around longevity and low total cost for online archival storage, it consumes as little as 3W/TB which is 1-2% of the electricity required to store 1TB of data on spinning disk. This seemingly simple trade-off between floor space and electricity can save the average company more than \$1M when archiving 50TB over 8 years.

As PowerFile continues to innovate in the area of rack density, the A3 will leverage such factors as improved hardware design and the ever increasing capacity of Blu-ray discs to scale to almost a petabyte by 2009.

**Q:**

**I HAVE HAD A BAD EXPERIENCE WITH OPTICAL SYSTEMS IN THE PAST. WHAT HAS POWERFILE DONE TO OVERCOME THE LIMITATIONS OF LEGACY OPTICAL?**

**A:**

No doubt you have had a bad experience with optical technologies in the data center until now. PowerFile is fully aware of the six key limitations of legacy optical solutions and has specifically designed the A3 to address them all:

1. **Reliability** – Using a patented combination of hardware and software technologies, PowerFile has increased data integrity by six orders of magnitude making the A3 the most reliable long-term data retention method – even more reliable than enterprise tape or RAID 6.
2. **Performance** – The A3 incorporates an intelligent caching mechanism for accelerating the ingestion and the retrieval of files and guaranteeing that most frequently accessed files will be available within milliseconds.
3. **Cost** – In addition to providing immediate payback by freeing up expensive primary storage and reducing backup and restore times, PowerFile has leveraged the economies of scale in the consumer-optical markets to deliver a storage solution that is not only competitive on \$/GB, but also is 1/5 the TCO of alternatives over 5 years.
4. **Proprietary** - PowerFile has added significant technology to boost such capabilities such as long-term reliability of data, but we have not deviated from the ISO standard UDF file system of CD/DVD/Blu-ray. There are no proprietary traps or vendor lock-in with the A3.
5. **Offline** – Legacy optical stored data offline rendering the data useless to an organization. The A3 was designed to keep archived data online and accessible for a fraction of the cost of traditional online technologies.
6. **Disruptive** – Legacy optical required organizations to do things differently to accommodate the technology and typically required special expertise. The A3 has been modeled after simple Network Attached Storage (NAS) and uses standard network protocols (CIFS, NFS, HTTP) and browser-based management to allow access to data over a standard Ethernet network. An organization is not required to know anything about optical, does not need to use or learn any new applications and does not need to make network changes. In fact, the A3 starter kit installs in less than 30 minutes and works with any existing application that can read or write data to a standard network share.

**Q:**

**IF BLU-RAY DISCS ONLY HOLD 50 GIGABYTES OF DATA PER DISC, HOW CAN I PRACTICALLY USE OPTICAL IF I HAVE MANY TERABYTES OF DATA?**

**A:**

Traditional interaction with optical technologies included manual interaction with the media itself, whether in a jukebox or on a disc-by-disc basis. The “volume size” was defined by the media capacity itself. In the enterprise, where an administrator may need to move many terabytes of data, the high touch nature of this method would make it completely unmanageable.

The A3 changes all of this. Just like Redundant Array of Independent Disk (RAID) technology eliminated the need to manage individual hard disks, the A3’s ArcOS management software has completely eliminated the need to manage individual discs or libraries. By logically combining all disc subcomponents within the appliance, the ArcOS software allows administrators to manage large volumes of data on the A3 in exactly the same way they would manage volumes on any other production server and then present the data to end users as simple network shares.

**Q:** I DON'T HAVE A BUDGET FOR ARCHIVAL STORAGE. HOW CAN I SELL SENIOR MANAGEMENT ON AN ARCHIVAL SOLUTION WHEN I AM HAVING TROUBLE PAYING FOR DISK, TAPE AND PERSONNEL TO KEEP UP WITH STORAGE DEMANDS?

**A:** You're not alone. The requirement for online archival storage is just beginning to emerge, so there are very few organizations outside of the federal government with an "archival storage" line item in their budgets. The benefits of creating an archival tier of storage for a company's long-term fixed content include, among other things, the following:

1. Reclaim expensive primary storage for its original purpose: high-speed, dynamic content
2. Dramatically reduce backup and restore times by eliminating fixed content from the regular backup cycle
3. Significantly reduce energy costs by storing archival data on a solution that consumes less than 3% of the power of other alternatives

If your company's storage requirements are growing faster than your storage budget and continuing to throw more disk and tape at the problem is becoming unmanageable, then PowerFile's A3 can offer immediate pay-back by consolidating fixed content on an online archive and allowing primary storage purchases and backup expenditures to be delayed or permanently reduced.

**Q:** WHY WOULD I PURCHASE AN A3 INSTEAD OF CONTINUING TO BUY CHEAP SATA ARRAYS?

**A:** First, SATA-based systems will need to be replaced every three to five years over the lifetime of the archive. If you are planning to archive data for more than ten years, you can expect to go through a minimum of two data migration projects which can easily dwarf the original acquisition costs. Then there is the fact that SATA uses about 3,000% more power and cooling than an A3 archive – that is real money.

Second, consider the raw-to-usable capacity ratio. Since the purchase price of SATA arrays (usually expressed in \$/GB) refer explicitly to "raw" capacity, important factors such as formatting, parity protection, and spare drives are not accounted for.

SATA drives require a minimum of RAID 6 to ensure reasonable data integrity rates. Include a hot spare and you can easily be looking at spending upwards of a 25% penalty per gigabyte. Conversely, PowerFile's patent-pending EVASt technology delivers better data integrity than a RAID 6 array and requires only 9% of raw capacity.

**Q:** I DON'T HAVE A TOOL TO IDENTIFY AND MOVE FIXED CONTENT TO THE A3. CAN POWERFILE HELP ME WITH THIS?

**A:** Absolutely. PowerFile's Archive Facilitator™ makes tiered storage easy. It seamlessly integrates with the A3 to deliver data discovery, contextual classification, automated archiving, and cost-effective storage into one intelligent solution. The software collects and analyzes detailed information about the utilization and criticality of data, classifies and prioritizes data and storage resources, and uses administrative policies to automate data movement according to business rules.

PowerFile has also certified many third party data management and movement applications for use with the A3 and continues to certify new ones every month. Because the A3 mounts on the network as a standard volume (CIFS or NFS) it has worked seamlessly with every application that has been certified to date. The certification process allows PowerFile to learn about the particular applications so that we can support YOU the customer in the best way possible. For certified applications, please check [www.powerfile.com](http://www.powerfile.com).

