

Technical Brief | PowerFile EVAST™

Drew O'Brien, *Marketing Manager*
PowerFile, Inc.



The archive is the ultimate resting place of corporate assets and contains the final copies of valuable documents and files. As a result, data integrity in enterprise archiving is of utmost importance. When a file is archived, there needs to be assurances of data integrity, authenticity, and availability for the entire duration that it is archived whether it is months, years, or decades.

For these reasons, PowerFile has created the patent-pending PowerFile EVAST™ for the Active Archive Appliance™ (A3): to improve upon the error rates of other storage technologies such as disk, tape, and standard optical systems and deliver the safest and most reliable platform on which to place archival data.

With EVAST, PowerFile is able to deliver a data integrity rate that far surpasses that of all other alternatives. The remainder of this paper focuses on how EVAST works and what kind of results organizations can expect when archiving data to an A3.

The Baseline: Standard Blu-ray® Disc Error Correction

All Blu-ray discs written in a standard drive contain a basic level of error correction specified by the Blu-ray Disc Association. This can be likened to the standard error correction that is contained within a single magnetic disk platter.

Blu-ray is an emerging technology and the official statistics for the standard error correction have not been published. Due to the enhanced design of Blu-ray discs and drives it is safe to assume that the error rate will not exceed the commonly published error rate for optical discs, which is 1 block error in $1E+12$ bits¹ (~116GB).

Information Technology professionals responsible for securing critical corporate assets should question whether or not this standard error correction provides a level of data integrity that is acceptable for enterprise archiving. PowerFile answers 'no' to this question for two reasons:

1. Given a bit error rate of 1 in $1E+12$, there is 1 unrecoverable block in every 116GB of data written. Since each PowerFile Library holds up to 200 50GB Blu-ray discs (10TB), this means that each full library could contain as many as 86 unrecoverable blocks – a potential loss of anywhere from 1 file to an entire disc, depending on where the unrecoverable blocks reside.
2. The built in error correction of magnetic disk drives, the primary alternative to magnetic tape or optical disc, provides a bit error rate of approximately 1 block in every $1E+15$ bits². With dual-parity RAID configurations, that number jumps as high as 1 in $1E+17$ ³. That's equivalent to one block error in every 11PB of data.

Therefore, it is reasonable to conclude that in order to compete in enterprise data storage, a level of data integrity must be provided that is, at a minimum, on par with what is delivered by a dual-parity RAID configuration. For the purpose of this technical brief, RAID 6 (1 in $1E+17$) will be used as the benchmark.

PowerFile EVAST™: Next Generation Data Integrity

PowerFile EVAST (Extended Verification And Self-healing Technology) is designed to improve upon the standard Blu-ray error correction and ultimately, to provide data integrity rates superior to that of enterprise RAID 6 – allowing organizations to archive their valuable data to the Active Archive Appliance with confidence that it will be protected and available when needed.

1 Source: Panasonic CW-8124-B Combo Drive Specifications, <http://www.panasonic.com/industrial/computer/storage/optical/models/CW8124B.asp>
2 Source: Seagate Barracuda ES.2 SATA Hard Drive Specifications, http://www.seagate.com/docs/pdf/datasheet/disc/ds_barracuda_es_2.pdf
3 Source: IBM Almaden Research on Advanced RAID, <http://www.almaden.ibm.com>

EVAST significantly increases data integrity by guarding data through the implementation of three unique, patent-pending technologies:

1. Spatial protection is introduced to help guard against common spatial damage on discs such as fingerprints, dust, and micro-variations from manufacturing
2. An additional, stringent verification step is performed after each burn to catch errors that the standard Blu-ray verification may have missed
3. Image checksums are generated before and after burning data to a disc. They are then compared to ensure that the data burned to the disc matches exactly with the data in the image. This helps guard against write mismatches

The EVAST process of writing data can be divided into 4 distinct steps. Figure 1 below illustrates the process:

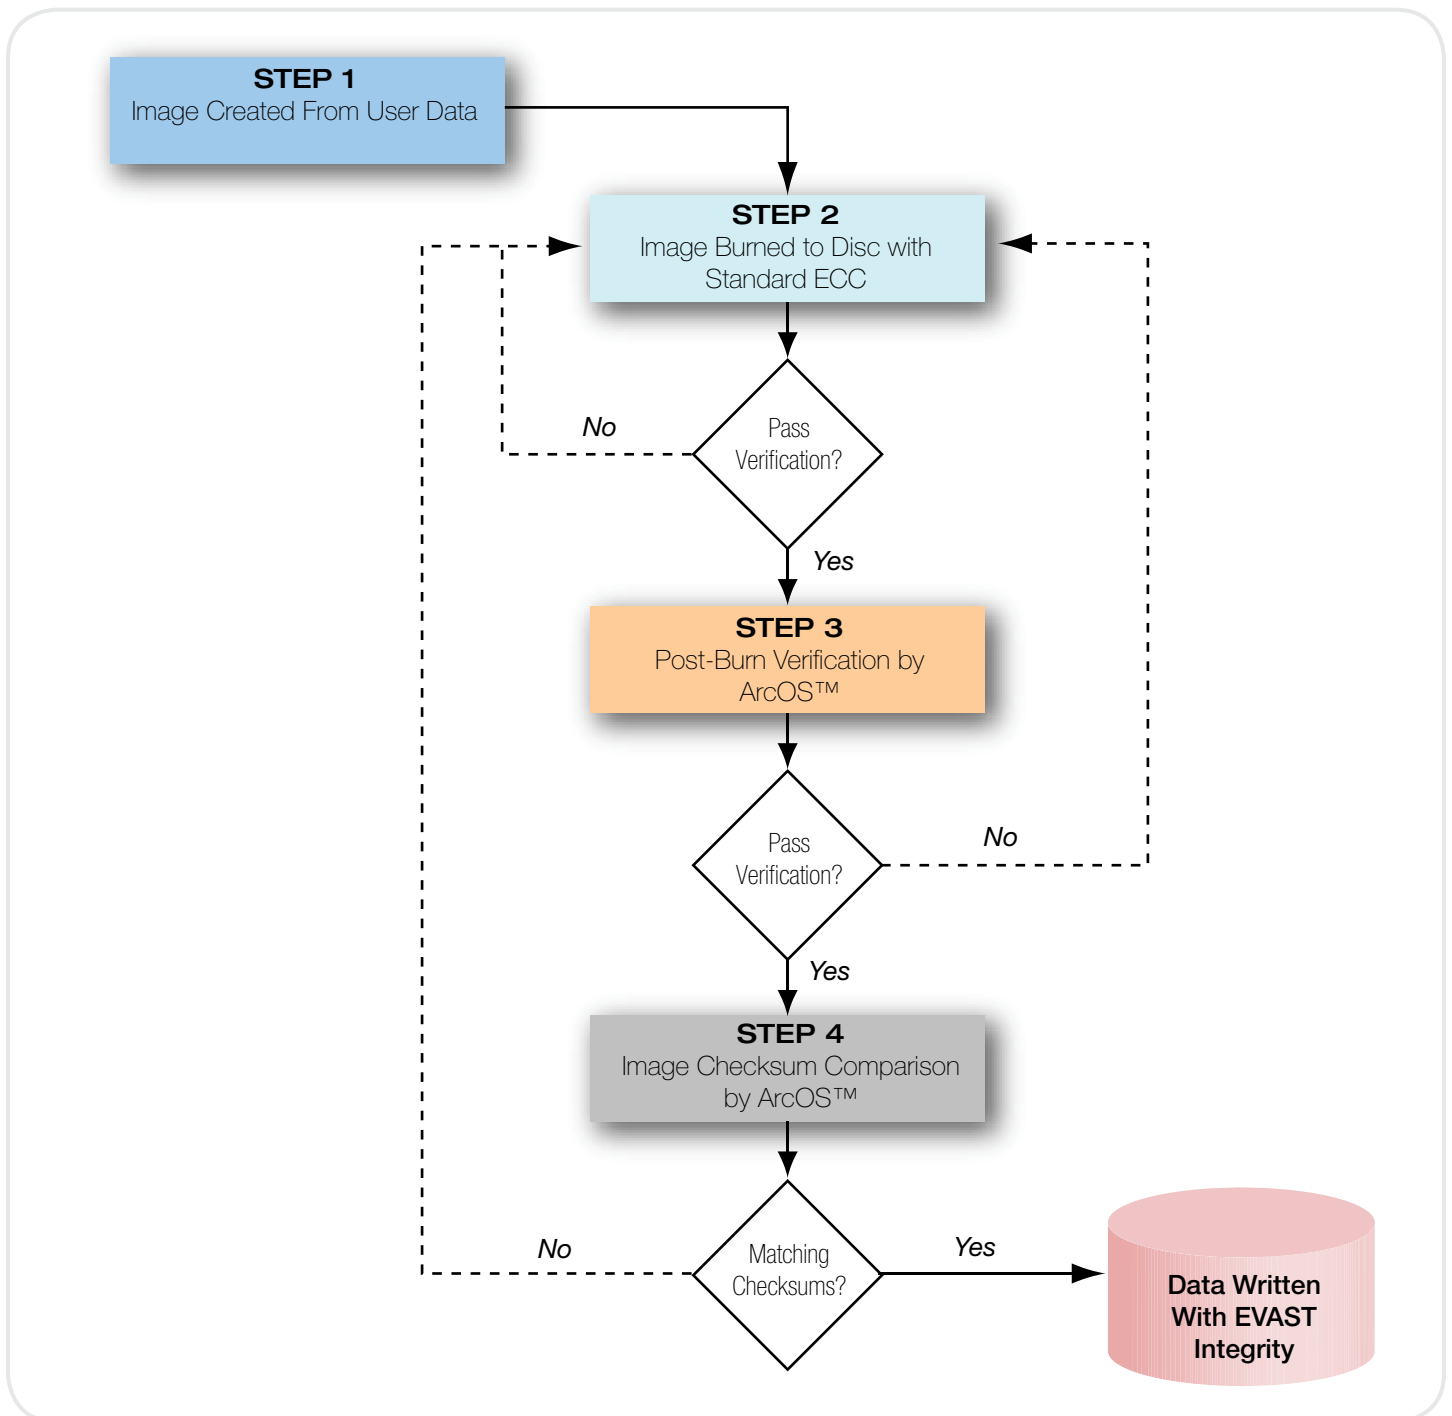


Figure 1: The EVAST™ Process

STEP 1: IMAGE CREATION

Before data can be burned to a disc, it must be compiled into an ANSI/ECMA standard Universal Disc Format (UDF) image⁴. Each image for a 50GB Blu-ray disc contains up to 45GB of data, depending on the sizes of the files contained in it.⁵

The image step is where PowerFile adds the first EVAST feature – spatial protection. As Figure 2 illustrates, the data clusters that make up an image are mapped out to different sectors on the disc – helping to guard against common damages such as fingerprints and dust. This is done using the EVAST media/ECC integrator to optimize cluster placement.

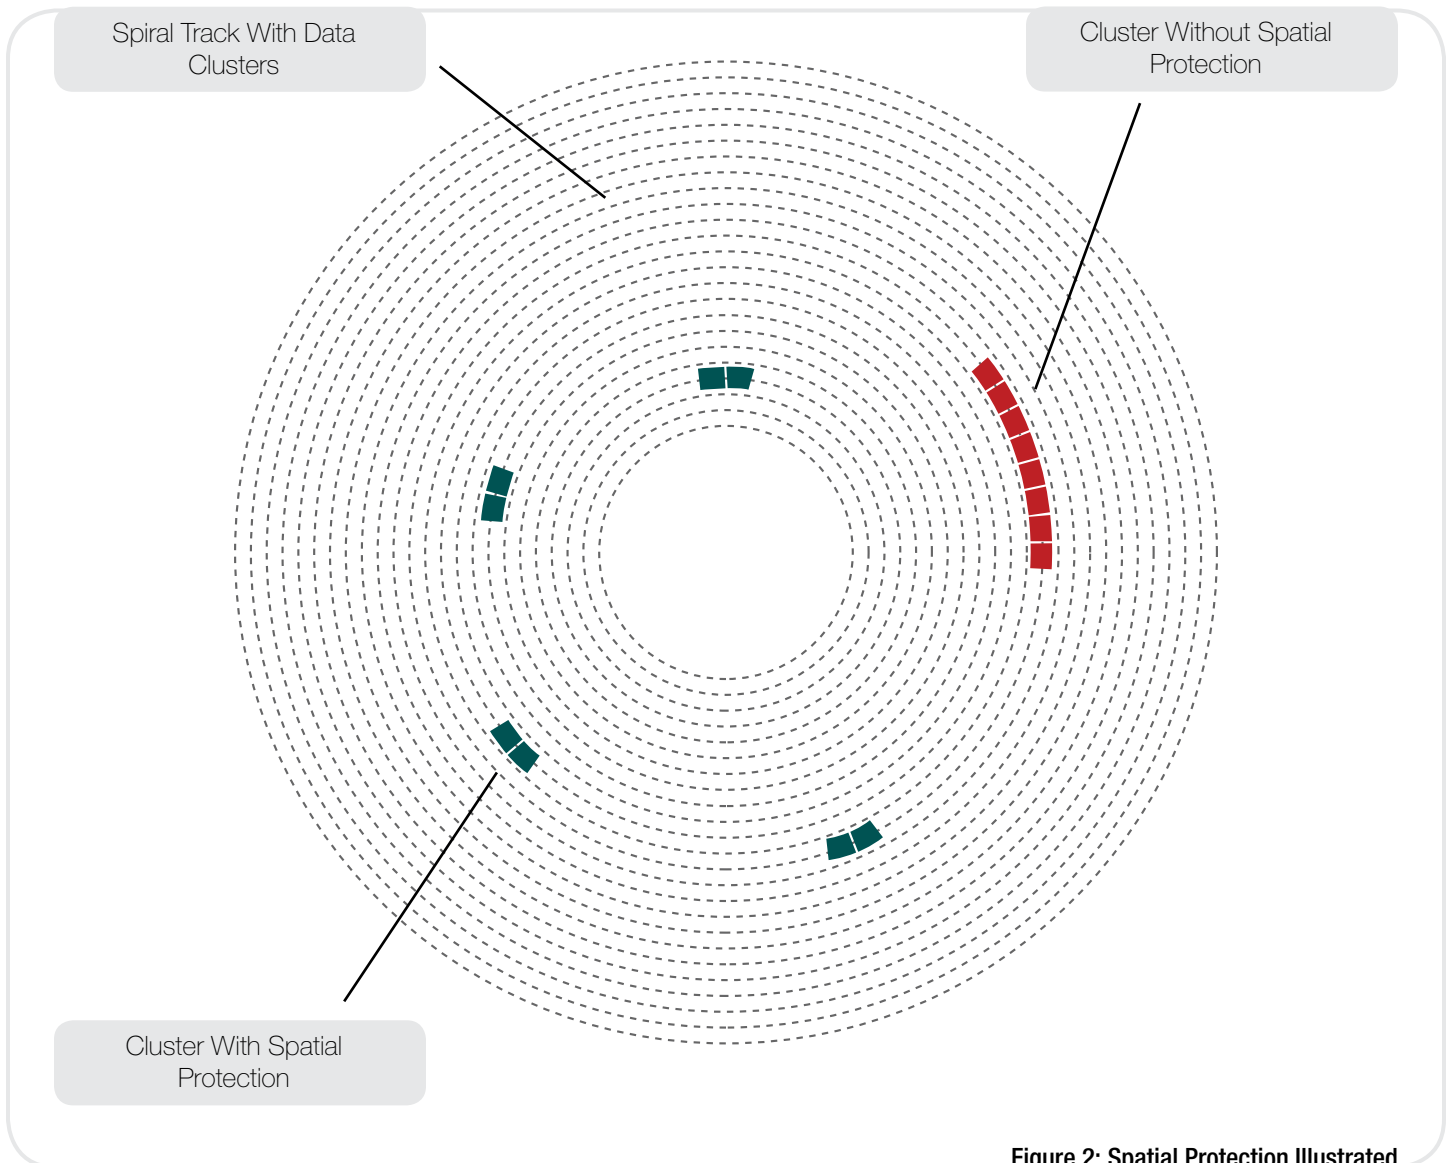


Figure 2: Spatial Protection Illustrated

STEP 2: IMAGE BURNED TO DISC

After image creation, the next step is to burn the image to a disc. During the burn process, the image is divided up into smaller, writable increments called clusters. A cluster is a collection of multiple blocks and is the smallest data increment that a Blu-ray drive can write.

4 The Universal Disk Format (UDF) is a format specification of a file system for storing files on optical media that guarantees compatibility with UDF-compliant devices. UDF is developed and maintained by the Optical Storage Technology Association (OSTA). For more information visit <http://www.osta.org/specs/index.htm>

5 The remaining 5GB on each disc is reserved for formatting and EVAST overhead.

When the image is written, the Blu-ray drive first burns a few clusters to the disc. Once written, each cluster contains the standard Blu-ray error correction (Figure 3). After the clusters are written, read verification is performed on those clusters. If any errors are detected during verification, those clusters are re-written to a separate portion of the disc using the Blu-ray defect sparing mechanism. This process is repeated until all of the clusters that make up the image have been written to the disc. Finally, once the full image has been written, the session is closed and the data is then subjected to the Post-burn verification performed by PowerFile's ArcOS.

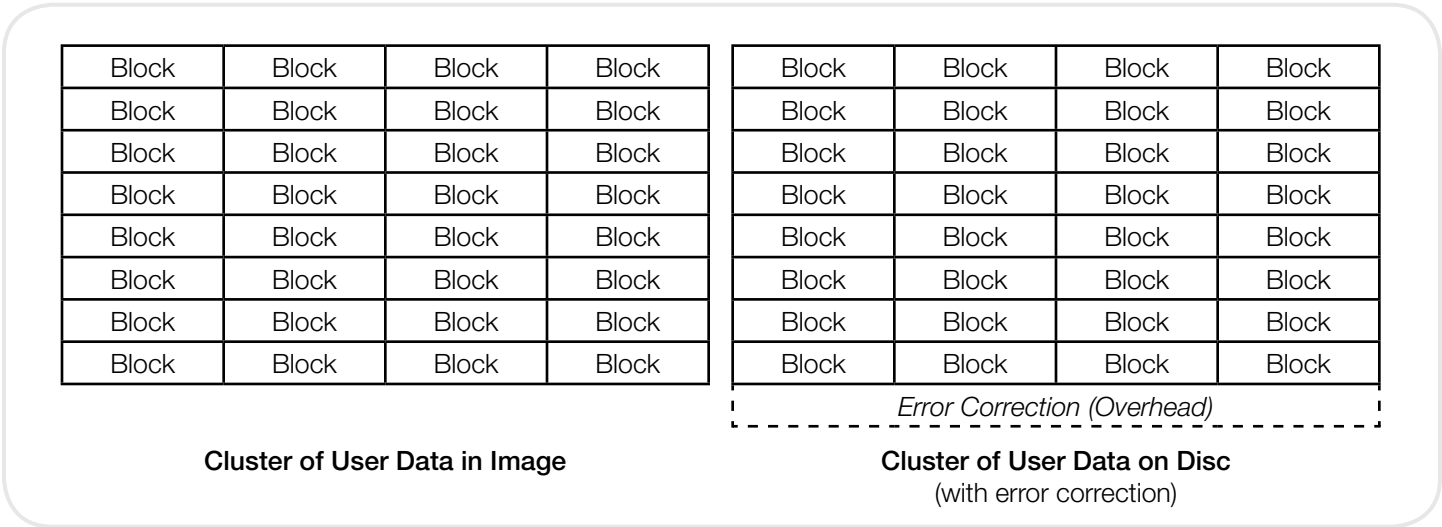


Figure 3: Clusters

STEP 3: POST-BURN VERIFICATION BY ARCOS™

After the image is written to a disc with the standard Blu-ray error correction, PowerFile's ArcOS performs a second post-burn verification with stringent conditions to ensure data integrity. If the data fails any of the conditions, the image is rewritten (Step 2) and the post-burn verification is performed on the rewritten data. If all of the conditions are met, ArcOS moves on to the checksum comparison.

STEP 4: IMAGE CHECKSUM COMPARISON BY ARCOS™

After each image is created, a checksum is calculated based on the information contained in the image. After the image is written to a disc (Step 2) and has passed the post-burn verification (Step 3), another checksum is calculated based on the information contained in the image written on the disc. The checksums are then compared by ArcOS. If they don't match, the image is rewritten to the disc (Step 2). This process guards against "write mismatches" which can occur when a drive unknowingly writes incorrect data and doesn't return an error.

The Results of PowerFile EVAST

As stated earlier, the goals of PowerFile EVAST are to improve upon the standard error rate of optical and to provide customers with a level of data integrity that meets or exceeds what could be expected from an enterprise magnetic disk system with dual-parity RAID protection.

With the commonly published optical bit error rate of 1 block in 1E+12 bits as the baseline and the approximate RAID 6 bit error rate of 1 block in 1E+17 bits as a benchmark, EVAST delivers a calculated⁶ bit error rate of 1 block

6

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

in 1E+19 bits - making the Active Archive Appliance the most reliable enterprise storage device on the market today from a bit error rate perspective. In practical terms, EVAST delivers the following:

1. Less than 1 block error in every exabyte of data written
2. Less than 1 block error in every 113,686 A3 libraries filled to capacity
3. A 10,000,000x improvement over the commonly published optical bit error rate
4. A 100x bit error rate improvement over a RAID 6 disk array

In addition to delivering unsurpassed data integrity rates, EVAST requires minimal overhead. In fact, total overhead for EVAST and formatting of a Blu-ray disc written in PowerFile's A3 comprises a mere 9% of the total storage capacity - compared to upwards of 25% for a typical RAID 6 array. It is also important to mention that every disc the A3 writes still maintains UDF compatibility - allowing it to be read in any standard Blu-ray drive.

Conclusion: Trust PowerFile with Your Data

Data is one of the most critical assets of any business. Organizations pay large sums to obtain it, decipher it, access it and gain value from it. Because data represents such a great business investment, much care should also be taken to preserve it. No other enterprise storage platform provides the data integrity of the Active Archive Appliance. So why would you choose any other device to be the final resting place for your priceless information? Trust PowerFile - your data will be safe.

KEY FEATURES

EVAST™ DATA INTEGRITY SCHEME

Patented error correction improves data integrity by six orders of magnitude (10⁶) for better reliability than RAID 6

HIGH DENSITY / LOW ENERGY USAGE

Fits up to 70TB and consumes less than 300 watts per 42U rack

ARCHIVE FACILITATOR™ AUTOMATION TOOL

Optional, policy-based discovery, classification, and migration tool for automated archiving of fixed content

COMPRESSION

Optional compression technology maximizes rack density

STANDARD NETWORK PROTOCOLS

Supports CIFS, NFS, and HTTP protocols - ideal for mixed environments

SCALABLE ARCHITECTURE

Systems scale incrementally from 4TB to 280TB¹

VIRTUALIZED STORAGE

Discs are virtualized into volumes to ease management

VOLUME MIRRORING

Creates a second copy of data for disaster recovery

TOUCHLESS DR™ WAN REPLICATION

Optional active-active asynchronous remote replication for automated, offsite disaster recovery

ABOUT POWERFILE

PowerFile Inc. is the innovator in archive appliances for long-term storage of digital content and assets. The company's award-winning solutions combine the speed and simplicity of online access with better economics and integrity compared to traditional archives, empowering organizations of all sizes to efficiently manage, secure, and preserve valuable digital assets. PowerFile's product offerings combine intelligent storage management software and sophisticated file management with cost-effective, robust hardware to deliver integrated, active, easy-to-use archive solutions.



1901 Sampson Avenue Corona, California 92879 U.S.A.
Tel (951)278-5626 Fax (951)278-9093 Email sales@cucsolutions.com