COLDSTORE Surveillance Storage Technology

COLDSTORE is a storage system designed from very basic first principles specifically for video surveillance applications. This has the principle aim of greatly improving the reliability of disk drives, irrespective of disk type (Desktop/Enterprise/Surveillance/RAID etc.). However, this technical approach delivers a number of additional benefits including greatly reduced power consumption, reduced cooling requirements, reduced UPS requirements and an array of disks which we not interdependent and can therefore be treated as independent media volumes.

COLDSTORE exploits the fact that video is real-time data source and is therefore strictly time-sequential in nature. Thus COLDSTORE records and stores the video data sequentially. This approach greatly simplifies the data processing task and allows the storage array to dramatically reduce or eliminate the three key factors which cause disks to fail –
(1) Temperature; (2) Vibration; and (3) Wear

Sequential Filing System

The elimination of vibration is achieved by writing sequentially within each disk drive. COLDSTORE's Sequential Filing System (SFS) has no need for dedicated data directory areas, thus the read/write heads do not rapidly move back and forward between directory and data areas, but in fact move very slowly across the disk (Plater) in continuous sequential writing mode. This is best illustrated by a video demonstrating SFS compared to traditional filling system – (Please request for Video Demonstration from Local Veracity staff).

Critical Differences Between IT and Video Data

SFS practically eliminates vibration from the read/write heads, and thus reduces data errors, greatly increase both writing and playback speeds and efficiency, and enables the drive to operate at maximum throughput with little stress. The reduction of rapid head-movement also reduces the overall power consumed by the disk and therefore reduces the heat generated. This reduction in operating temperature combined with the almost total lack of vibration significantly increases the useful life of each disk.

L.A.I.D

COLDSTORE also writes across the array in a sequential manner, again exploiting the sequential nature of video and the extreme speed of modern disk drives, especially when used with SFS. To clearly distinguish this writing pattern from conventional RAID approaches, we refer to this technique as L.A.I.D , for "Linear Array of Idle Disks".

To fully understand the benefits of the L.A.I.D Principle, we first need to understand the differences between reading and writing on hard disks. The writing process is an active process, changing the state of the data written on the disk and is therefore the process most likely to cause failures (typically through data corruption). In contrast, the reading process is entirely passive and extremely unlikely, in itself, to cause a data failure. LAID exploits this key differences, enabling a system which has full mirrored redundancy when writing, but which does not require twice the number of disks to store the data. (Note that recording to full mirrored pairs is an available and simple option with COLDSTORE).

With LAID, all Disks in the array are switched off except the current writing pair. LAID writes to two disks as a mirrored pair (RAID 1), thus providing full data redundancy should a disk fail during the writing process. When the disks become full, the next in the sequence is powered up and the writing point in the array steps along by one disk, to produce a mirrored overlapping pair pattern, most easily described by example. The process starts by writing to the first two disks in the array say one and two.

All other disks are switch off. When disks 1 & 2 become full, disk 3 is started up and the writing continues on disks 2 & 3. Disk 1 is switched off, preserving the recorded data on disk 1. Disk 2 is overwritten with new data. When disks 2 & 3 become full, disk 4 is switched on, the writing continues on disks 3 & 4 and disk 2 is switched off, again preserving the recorded data on disk 2. When the mirrored overlapping pair writing pattern reaches the end of the array (say disk 15), then the pair pattern simply loops round, writing to disks 15 & 1. Therefore, the sequence of writing pair as the LAID system loops around the end of the array and starts at the beginning again is;
14-15 ; 15-1; 1-2; and so on.

The LAID writing process results in full mirrored pair during writing, Leaving one copy of the data for reading. Therefore an array of 15 disks provide an affective capacity of 14 disks, or N-1, even though all data was writing in a very safe, fully mirrored process. The passive reading process presents negligible risk to the data and the only time the risk will be writing to again is when it is overwriting with new data anyway. Thus LAID principle provides full mirrored writing security (RAID 1) but does not require twice the number of drives to provide it. All disks not being writing to are switched off, which in a 15-bay array means 13 disks. Thus at anyone point, 13 disks out of 15 are switched off, so in affect each disk is switched off 80% of the time (Ignoring playback which is a negligible factor). To put this in prospective, in a year, any particular disk will be used for a total of 7 weeks and will be switched off for a total of 45 weeks.

Contrast this with RAID 5 and RAID 6 systems where all disks have a constant high duty cycle and thus have a much higher probability of disks failure. Further, disks in a RAID 5 / RAID 6 volume are totally interdependent with a risk of losing ALL data should multiple disks fail (2 in RAID 5, 3 in RAID 6). Note that a single disk failure in RAID 5 /6 results in an array rebuild which further increases the duty cycle, increases risk of secondary disk failure, reduces throughput and performance and which, with modern disks of massive capacity, can result in rebuild times measured in days.

For the purpose of comparison, it is important to note that the LAID System in COLDSTORE is affectively a sequential RAID 1 System, with the added benefits of SFS and the very low disk duty cycles of LAID. COLDSTORE can also be very simply configure to operate as a full mirrored pair system, retaining full duplicate copies of the data on mirrored pairs of disks. However, experience of over 7 years of use of COLDSTORE technology in the field has show that there is little to be gained by this in practice due to the high resilience of the mirrored overlapping paired system.

RAID 5/6 VS LAID + SFS Comparison

Technology	RAID 5	RAID 6	SFS + LAID
Raw versus Effective Storage Capacity	N-1 Per Volume (Multiple volumes per array)	N-2 Per Volume (Multiple volumes per array)	N-1 Irrespective of N (Disks Independent)
Single Drive Failure	Recording performance degraded. Rebuild required. Rebuild times with large drives May be extremely long.	Recording performance degraded. Rebuild required. Rebuild times with large drives May be extremely long.	No affect. No rebuild required. Array space is calculated to include 1 disk of spare capacity.
Double Drive Failure	TOTAL data loss! The entire array of data from the affected RAID volume will be lost and will be irrecoverable.	Recording performance severely degraded. Rebuild required. With two large drives both in rebuild simultaneously, the rebuild times will be extremely long, with increased exposure to further disk failure leading to catastrophic loss of all data.	Array archive time will shorten by 1 disk until failed disks are replaced. However, as no rebuild is ever required, prompt replacement of failed disks will prevent this ever happening.
Drive Type	Enterprise / RAID High disk RPM required	Enterprise / RAID High disk RPM required	Any SATA drive type High disk RPM not required.
Array Expansion	Controller specific. Requires RAID volume rebuild – reduced recording performance during expansion.	Controller specific. Requires RAID volume rebuild – reduced recording performance during expansion.	Simple – Add disks. No rebuild required. No degradation of performance.
Disk Extraction?	NO Data on disk would be lost. Equivalent to disk failure.	NO Data on disk would be lost. Equivalent to disk failure.	YES Data preserved on extraction. Data is one continuous and Identifiable time – spam.
Drive Failure Rate/year	7.65%	7.65%	0.74%
Failed Drive Replacement Time	Urgent	Urgent	With Array retention time.
Re-Writable Data Support	Yes Written data may be altered.	Yes Written data may be altered.	No Evidential data should not be altered after writing. Disk may be re-written with new data only.

www.backblaze.com 2015 Q3 Report. WD Red 3 TB average. Annual failure rate. Avg age > 16 months
COLDSTORE Drives 2013-2016 WD Green 3TB avg. annual failure rate. Avg. age > 21 months

Future Proofing – Large Disks

With COLDSTORE it is possible to replace smaller drives with large capacity once as they become available from the disk drive manufacturers. This involves zero data loss as the extracted drives are playable media volumes which may be kept for as long as considered necessary (E.g. For the stated retention period for the project).

A customer could replace 6TB drives with new 10TB drives for example an increasing the storage capacity by 66% in the process. In 2017 and 2018, 12TB, 14TB and even 18TB disks are expected from the two main HDD manufacturers. Note that RAID 5 and RAID 6 systems are not suited to the largest disks available even today.

Solution Comparison

As the camera resolution increases the storage requirements increases with them. For a typical installation of 300 cameras using high resolution cameras with 120 days file retention, the storage requirements always require external storage solutions for the VMS.

For our comparison with other storage solutions we will be using an example system with the following requirements.

# of Cameras	Avg. Throughput	Retention	Storage required
300	3.25Mbit/s	120Days	1390TB

Accommodating the storage requirements require several units. It is assumed that all RAID based systems are configured with RAID 6.

Product	# Drive Bays	Max HDD Size	Useable Space Per Array	# Units Required	Total Disk Space	Total Useable Space
Veracity COLDSTORE	15	10TB	140TB	10	1,500	1,400 TB
Pivot3 Surveillance Data	16	8TB	112TB	13	1,664	1,456 TB
Rasilient PS5012VID-GE2*	12	10TB	100TB	14	1,680	1,400 TB
Promise E610f*	16	10TB	140TB	10	1,600	1,400 TB

*Note: RAID 6, without a standard “hot spare” (So N-2 useable space)

Disks Lifetime

COLDSTORE greatly extends the expected average lifetime of disk drives. This leads to fewer drive failures, which in turn leads to fewer service call-outs. For RAID 6 it is imperative that a failed drive is replaced urgently as a second drive failure will leave the array exposed to a potential for complete data loss on the affected RAID 6 array. Further, the recording performance during rebuild is degrading and must be taken into account during system design. COLDSTORE, the other hand, will not require a rebuild and even unskilled users can replace a faulty drive as its inclusion into the array is handled automatically by COLDSTORE.

Product	# Drive Units	Disk Warranty	# Drive failures over 3 years [1] Using Veracity Data [2] Using Backblaze Data
Veracity COLDSTORE	150	5 Years	[1] 3
Pivot3 Surveillance Data	208	3 Years	[2] 48
Rasilient PS5012VID-GE2	168	3 Years	[2] 39
Promise E-610f	160	3 Years	[2] 37

NOTE: Newer drives tend to be more reliable due to constantly improving manufacturing techniques, so the actual annualized failure rate using 8TB and 10TB drives is likely to be better than the projections in the table above. The 8TB and 10TB disk drives are too new to have accurate failure rate statistics. However, any improvements will be valid for both RAID5/6 and COLDSTORE, therefore we expect that the LAID system will maintain an almost 10:1 superiority in disk reliability in terms of annualized failure rates.

Power Consumption

Each 15-Bay COLDSTORE 3U unit consumes only 60W of power under full normal load, due to its low – power design and due to the fact that only two disks are active on average. Thus the example above, the relative power consumptions would be project to be :

COLDSTORE : 600 watts
Pivot3 : 5200 – 6150 watts
Resilient : 2450 – 4900 watts (Depending on Model)
Promise : 5240 watts