Storage for Surveillance – The Weakest Link

veracity

TRENDS IN STORAGE & IP VIDEO

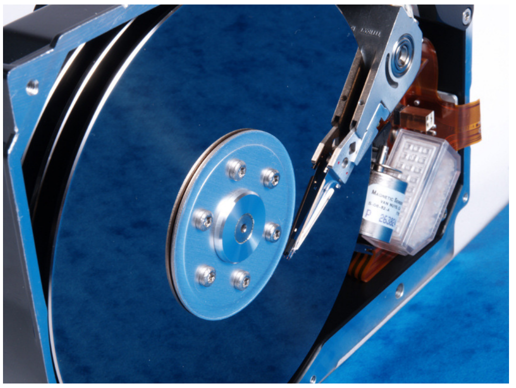

Video surveillance systems are highly dependent upon hard disk storage systems, yet we all know these are the weakest link in system reliability. This article looks at why this is so, how the increasing market for IP video systems makes this even more critical, and what the future holds for enterprise surveillance storage systems.

Digital Recording Market Trends

The most obvious current trend in video surveillance is the slow evolution from analogue cameras to IP cameras. Industry vendors of recording systems appear to be divided into two camps : the sceptics, who are wondering when the IP video revolution is really going to take hold, and the informed, who know it is already here and are just getting on with it. Although the IP video segment of the UK market is currently only around 15%, it is growing rapidly. In Europe, where analogue CCTV cameras are not so entrenched, the growth rate is even steeper.

One of the things that has held back IP video systems up to now is the lack of any real reason for users to switch, even for new installs. What would they have gained, bar a little increased flexibility ? Now, in 2007, the story is quite different : most well-designed IP cameras are progressive scan, at long last leaving behind the problematic interlaced image legacy of the broadcast TV standard. Further, and more importantly, mega-pixel cameras (with resolutions way beyond HD TV) are now here and affordable. At long last, mega-pixel cameras provide the compelling reason to switch to IP-based systems. The improvement in image quality offered by high-quality mega-pixel cameras in nothing short of dramatic, even breathtaking. Another problem which has held back adoption of IP cameras is the vast amount of legacy co-axial cabling installed in existing systems. However, several vendors are now offering Ethernet-over-coax solutions, so this is no longer a barrier at all.

Increased requirement for video storage

We are seeing move towards higher recording frame rates with many DVR systems offering real-time recording – although with the break away from TV broadcast standards with it's fixed 25fps, we can now start an argument about what "real-time" really is ! Even with the efficiency of improved image compression, such the rapidly emerging standard of H.264, these higher frame rates require increased storage. With mega-pixel cameras, we are seeing vastly improved resolutions, and thus increasing storage capacity requirements. Further, people are beginning to properly understand that the Data Protection Act *does not* set a limit on how long recordings can be held (only that you must have a policy and be able to justify it), so we are seeing a trend towards keeping recordings for longer and longer. Some city centres and major sites are considering data retention for up to 2 years, and many countries have no legislation at all covering this topic.

All of these trends require greater and greater storage capacities, putting extra pressure on what is already the weakest link in DVR systems – the hard disk.

The hard-done-by Hard Disk

Is it really fair to blame the hard disk for DVR system reliability ? Whilst it is a fact that hard disks are by far the most common reason for DVR system failure, we have to look deeper into how they are used to understand that it is completely unfair to blame the disk drive manufacturers, as their technology just keeps getting better and cheaper all the time.

Broadly speaking, disk drives can be split into two classes : so called Personal Storage (PS) drives for normal desktop computer or small storage array use with low to medium duty cycles, and Enterprise

Storage (ES) drives for high-performance, high temperature and very high duty cycles. The problem is that ES drives can be anywhere from three to ten times more expensive per GB than PS drives. ES drives also tend to be much smaller in capacity, so you would need fewer PS disks than ES disks to provide a certain capacity.. Thus ES drives are almost never used in DVR systems as it would make the DVRs totally uncompetitive on price.

Disk drive manufacturers quote impressive MTBF (Mean-Time Between Failure) or AFR (Annualised Failure Rate) figures for each drive model, when used under ideal conditions. The expected failure rates rise very quickly under non-ideal conditions, and also for PS drives when they are used in ES mode (i.e. most DVR applications would be considered ES mode usage). Combine these two conditions, and reliability levels fall through the floor. Unfortunately poor cooling of drives is common in many DVR systems, so the working conditions are less than ideal. To quote one well-known disk manufacturer – "heat is the enemy of hard disk reliability", so when viewed in this way, it is hardly surprising that disk reliability is a problem, despite the manufacturers best efforts to improve reliability. Using published figures, it is a relatively simple matter to calculate probabilities of failures for PS drives used in ES mode. In a 16-disk array (or 4 DVRs with 4 disks each or whatever), the probability of a disk failing within three years is an astonishing 70% - and this is not even under the worst conditions often found in real-world DVR installations. [Contact author for calculation details].

Current solutions for reliable storage

One of the main reasons, apart from lower cost, that many modern DVRs offer embedded Linux-based systems is to improve overall reliability. Embedded Linux does not require constant updating like the Windows operating system, nor is it nearly as susceptible to viruses, Trojans and worms. Experience has shown that disk reliability is also improved with embedded Linux systems. Further, failed disks in embedded systems are generally field-replaceable as the disks normally contain only data, but a Windows system disk failure is highly impractical to field-service.

With the current trend towards IP video systems, many vendors are offering Windows-based software packages which you can run on your own hardware. Does this solve the problem ? No, it just shifts the responsibility for the hardware onto the long-suffering integrators ! In fact, it may even make the problem worse, as many less well-informed installers may end up using PC hardware which is simply not up to the task, as the temptation to buy the cheapest solution is just too great. The advantages of the Windows user interface for system management and control are great, however, so this growth in the market for IP video software packages will continue, and even accelerate.

The most common approach to providing reliable storage is to offer some form or redundancy through external storage, usually as a RAID system (Redundant Array of Inexpensive Disks). As there are many flavours of RAID, a quick summary of the most common types might be useful :

RAID 0 : This is commonly known as striping, where the data is split across several disks to increase throughput. There is no increase in overall reliability, and in fact loss of one disk means the loss of all data.

RAID 1 : This is commonly known as mirroring, where the same data is written to two disks, thus ensuring that if one disk fails, the data is still available on the mirrored partner, and a new partnered disk can be rebuilt by a simple copy process. However, this requires double the number of drives to implement the solution and so is twice the price.

RAID 5 : This method writes data to one disk and parity (error correction codes) to another disk and continues in a round-robin manner over three or more disks. This requires at least one extra disk in the array, and also requires the most complex controller system. A failed disk can be rebuilt slowly, but if two disks fail, all data are lost.

There are other flavours, variants and combinations of RAID, and generally speaking as they offer higher and higher levels of redundancy, they become more and more expensive. Our concern here is with RAID 5 arrays, as these are commonly specified for high-end DVR systems, but in fact we shall show that it is a particularly poor solution for modern high-capacity surveillance storage systems.

What is not commonly understood is that in any array of tightly-packed drives, the vibration of other disks increases the stress and wear on every disk. Further, the tight packing means that the heat dissipation from any one disk heats up the disks on either side of it, exacerbating the problem of keeping the disks cool. RAID 5 makes disks work particularly hard, and the harder they work, the hotter they get, especially during rebuild of a failed disk, which is a long and complex process. Bear in mind that when used as a storage system for surveillance, the RAID array cannot be taken offline, but must continue to record while a disk is rebuilt. Many RAID 5 solutions require arrays of identical disks – capacities cannot be mixed. Some cheaper arrays require even identical disk make and models. This is not terribly convenient for field servicing. When a disk does fail, it must be replaced right away, as a further failure would cause the loss of all data – exactly what MUST NOT HAPPEN in an mission-critical enterprise system.

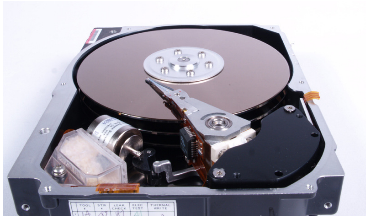

Here's another problem : RAID arrays were first proposed in 1988, when hard disks were tiny. The rebuild times for a RAID 5 array with 250GB drives are considerable. The bigger the drive, the longer the rebuild time. 500GB and 750GB drives are now common, and very cost-effective. However, a RAID 5 rebuild can take days with such massive disks, and all that time the array is working in degraded mode, and working that much harder, and yet still has to continue recording. Should a second disk fail in that time, all data will be lost. This is not a theoretical situation, but one that many integrators of RAID 5 systems are beginning to experience.

Both Hitachi and Seagate have announced 1,000GB drives, and this trend of increasing drive capacity will continue for the next few years. Thus RAID 5 for surveillance applications is looking increasingly redundant, if you'll pardon the expression. Given the decreasing cost of drives (a 500GB drive can be bought for less than £100), RAID 1 is looking increasingly attractive, at least for enterprise systems.

Future trends

Ever increasing storage capacities will be required to meet the storage requirements of mega-pixel surveillance with high frame rates and long archive periods. Even though we will see 1TB (1,000GB) shortly and even larger drives in the next few years, there will still be a requirement for disk arrays and systems to make them more reliable.

At least one hard drive manufacturer has released a hard drive designed specifically for the rigours of video surveillance storage applications. This is an excellent step forward, and we hope other manufacturers will follow.

The shift away from Windows-PC-based systems to embedded Linux systems will continue and will move up to the enterprise end of the market. IP video software vendors will either start to support Linux systems or, more likely, will integrate support for more reliable external storage systems, such as RAID 1 Network Attached Storage (NAS) units, or other forms of NAS approach. Thus as IP video systems begin to dominate the market during 2008-2010, we will see a strong trend towards centralised storage systems, with integrators having to become familiar with terms like NAS, SAN, iSCSI and storage virtualisation. This will further increase the overlap between video-based security systems and the general IT industry, and no doubt will generate further friction between Security managers and IT managers too !

In the long term (i.e. 10-15 years) we'll see a move towards solid-state memory, but in the meantime, no matter how big and how cheap disk drives get, high-frame rate mega-pixel surveillance data (stored for long periods) will fill them all up.

And what about a proper solution for disk reliability ? We'll see a vastly improved storage systems appearing, including both disks and storage arrays specifically designed for video surveillance. Just keep reading this website to hear about these developments as they happen. . . . .