Successful Storage Management Using a Scalable Video CODEC

A significant portion of the budget for any new surveillance system is allocated to storage. Its relentless consumption makes it a limited and valuable asset. Management of how this asset is consumed is critical to the successful operation of the system installation and is key to determining ongoing operational costs.

In surveillance applications, video is captured and recorded to mitigate a perceived risk. It can be a tremendous asset when trying to determine exactly what transpired during an event. An invaluable tool in pressing home a prosecution, a video record can be equally valuable in defending individuals and corporations against fraudulent accusations. Its absence can be a liability for an operator and premature deletion (spoilage) of video evidence could be viewed by a judge or jury as an admission of guilt or liability. The high cost of storage, however, can make archiving video indefinitely prohibitively expensive. A set retention time must, therefore, be established to ensure adequate storage. After this set time, video must be deleted and the storage reused.

The retention time of video evidence is dictated by the amount of storage available and the rate at which storage is consumed. Second generation CODECs such as H.264 can reduce encoded bit rates over previous generations such as MPEG4 by a factor of two. This doubles the effective capacity of a system and might halve the storage cost or double the retention time. Consumption is nonetheless relentless, and after a pre-determined period (or usage limit), some portion of the archive must be either deleted or copied to offline storage. For a majority of surveillance operators, retention times can be surprisingly short.

Consider, for example, a retail outlet with eight standard definition (D1 resolution) surveillance cameras. The video stream from each of the cameras is compressed using an H.264 CODEC and stored to disk at a rate of 1.5Mb/s. This bit rate results in an image quality sufficient to detect and identify shoplifters. Simple math reveals the associated cost:

- 1.5Mb/S = 675MB/hour consumed storage per camera.
- 24 (hours) x 8 cameras = 130GB consumed storage capacity per day.
- 130GB x 7 (days) = approximately 1TB consumed storage per week.

In this example, the store manager must make a risk mitigation decision. He decides that keeping three weeks of storage is probably sufficient to mitigate the risk from shoplifting and employee theft, and so decides to equip his installation with 3TB of storage. He also institutes a policy whereby storage is recycled on a three-week rolling basis. While this policy is probably sufficient to protect physical inventory, it ignores a second, significant risk.

“Slip and fall” litigation is a very real problem for retail outlets. In many countries, the statute of limitations allows a lawsuit to be brought against a shop owner for up to two years after an alleged event. Unfortunately, the need to recycle storage is well known to individuals who stage fraudulent slip and fall accidents. Waiting just a few weeks before pursuing a lawsuit can be sufficient to ensure that any video evidence has been deleted, raising a suspicion of guilt over the store owner who has spoiled the evidence. Conversely, waiting two years to bring an action could also be regarded with suspicion. Though risk reduces over time, to completely mitigate the risk associated with this type of action being brought, the store owner would need to keep a two-year video record. With conventional CODECs such as H.264, this is clearly an expensive policy.

The preceding discussion raises the notion of risk profile. In the short term, the store owner has a very high likelihood of loss due to shoplifting. The magnitude of this loss can be a small, but nonetheless non-trivial percentage of his revenue. In the medium term, there is a small risk of his being the victim of a fraudulent slip and fall case, the financial consequences of which could be significant. As time goes on, this risk diminishes to almost zero. Estimating the mathematical product of financial impact and likelihood gives us a measure of the risk. Figure 1 shows a potential risk profile graphically. Also shown is the degree of risk covered by a short (three week) video retention period and the degree of exposure that remains.

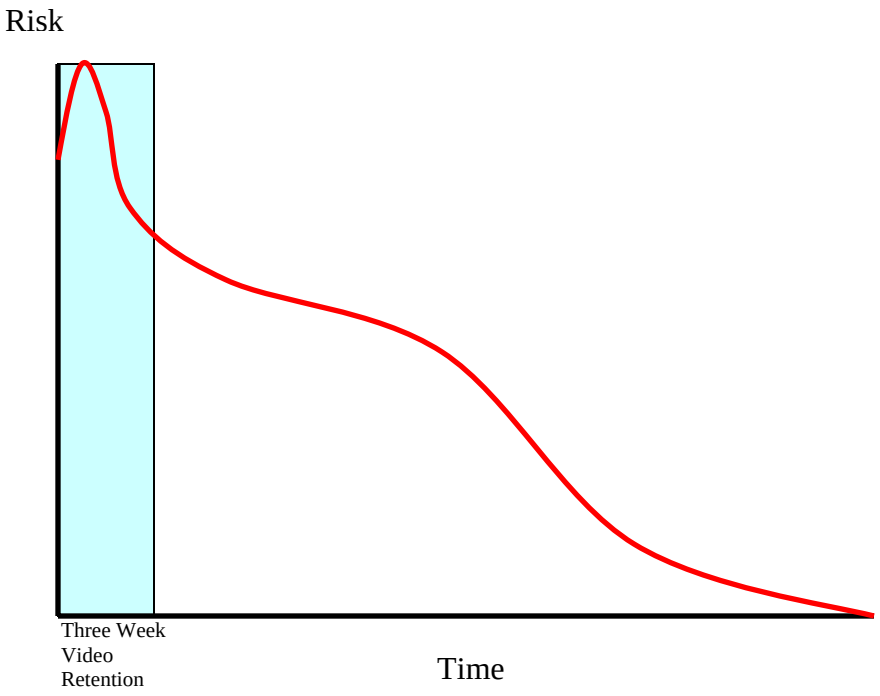

Figure 1: Potential Risk Profile

Clearly, having to delete the video archive after a three week retention time leaves the store manager still exposed to significant risk. Buying more storage to extend this period, however, is prohibitively expensive.

This problem is solved by scalable video CODECs such as H.264 SVC. Streams encoded with scalable video CODECs can be simply thinned over time to reduce resolution or frame rate. The corresponding reduction in file size makes storage available for re-use while retaining a valid video record. Using conventional CODECs, the stream would have to be decoded, re-sampled and re-encoded to reduce its size on disk. The compute requirements to accomplish this on a regular basis are prohibitive, but using SVC, the truncation operation is as simple as deleting a file from the hard drive. The simplicity of the operation opens several new options for our store owner.

Reconsider the previously described retail store scenario, but this time the store owner encodes the video feeds from the cameras using a scalable video CODEC. Using SVC, he is now free to examine the profile of the risk he is trying to protect against and "shape" the profile of his stored video to better match the risk. A potential analysis of the risk may be:

Event	**Risk Magnitude**	**Video Requirements**	**Required Storage Time**
Shoplifting	Likelihood: High Impact: Medium	Highest Possible Resolution for identification of persons and their actions	1 Day
Employee Theft	Likelihood: Medium Impact: Medium	Medium Resolution to establish repeat behaviors or trends	1 Month
Slip and Fall / Acts of Violence	Likelihood: Low Impact: High	Medium Resolution / Medium Frame Rate to establish facts relating to an event	2 Months
Slip and Fall	Likelihood: Low Impact: High	Low Resolution / Low Frame Rate to defend against fraudulent case	2 Years

To shape the profile of the stored video to match the perceived risk, the store owner might, for example, store one day of video at the full resolution of the camera and at full frame rate. This might consume approximately 150GB of storage. At the end of the day, with no theft detected, the streams could be truncated to CIF resolution at 15 frames per second, cutting the storage requirement to approximately 20GB per day. After one month, the stream might be further truncated to CIF resolution at 7.5 frames per second, reducing the storage requirement to 15GB per day. After two months, the stream could be truncated yet again, this time to QCIF at one frame per second. This would bring the storage requirement down to approximately 3GB per day. Figure 2 shows how the per day storage consumption would change over time.

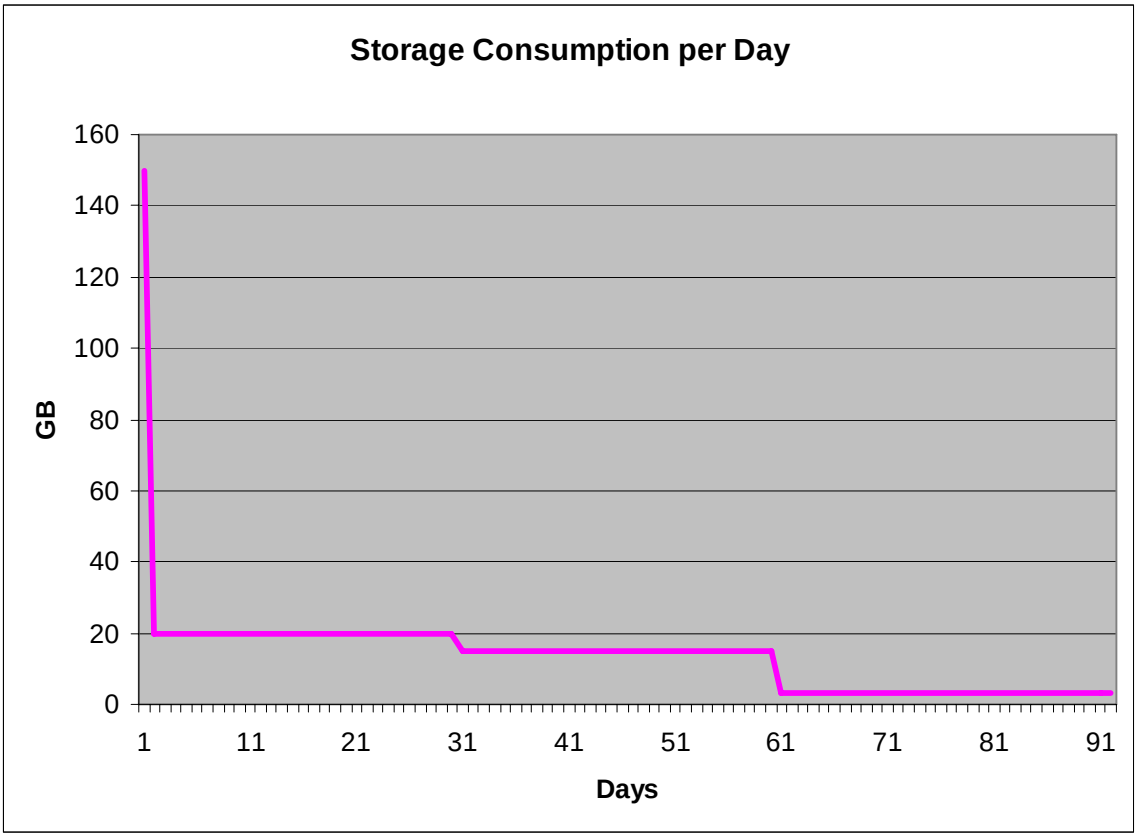

Figure 2: Storage Consumption over Time

Using SVC, the store owner could retain approximately 660 days of video evidence using the same 3TB of storage. This compares to the three weeks (21 days) available using conventional H.264. Tailoring storage consumption over time to match the various threat levels can mitigate that threat for much longer periods of time. Indeed, with just 3.2TB of storage, the owner would have the peace of mind of having an intact video archive for the entire two year statute of limitations.

Scalable video CODECs are set to revolutionize the way storage in surveillance installations is managed. They allow the stored video archive to be managed over time to match the threat profile and, in doing so, allow for longer retention periods. The full risk mitigation benefit of high resolution and high frame rate video can be preserved when the risk is at its highest. As the risk diminishes over time the space consumed by the video archive can be gradually reclaimed without completely deleting the video record. The result is a more efficient use of the expensive storage asset, a reduction in operating costs, and an overall reduction in risk.