Why Intelligent Storage in Cameras will Transform IP Video Surveillance

WHITEPAPER

By Doug Marman
CTO and Co-founder, VideoIQ, Inc.

The security industry has adopted a centralized storage approach for IP video surveillance systems. Manufacturers followed traditional data center designs, assuming that it was best to use the standard architecture for Information Technology (IT). However, many of the growing issues with IP video systems today, such as bandwidth, storage and maintenance costs, are the direct result of this centralized storage architecture.

Here's the problem: Data centers and most IT systems are designed for many users accessing data servers in one location. This is called a "one-to-many" model, since each data center serves large numbers of users. Most experts assumed that this would also be the best model for security video as well, but this is wrong. Sensor networks have exactly the opposite requirements.

Video surveillance systems include dozens or hundreds, and sometimes thousands of cameras, all spread out across the network, with only a few users. Cameras supplying the data far outnumber the users, and there is no way to centrally locate those cameras. They need to be out at the edge of the network. This "many-to-one" architecture generates very different demands on the system.

For example, typical enterprise servers spend about 50% of their time reading data stored on their hard drives, with the other 50% spent writing new data.[1] Surveillance systems, however, require non-stop writing of data to record the video, 99% of the time, but only about 1% to play it back.[2]

Second, data centers usually need a small number of storage units to support 20 or more different enterprise functions. So, the same storage servers can be used for a wide range of applications. But surveillance by itself requires extremely large amounts of storage. "The consequence is the video requirements can put a speed bump in IT plans to reduce costs through consolidating storage," according to Steven Titch, editor of Network Centric Security. "Plus, in video," according to Lee Caswell of Pivot3, "each time storage capacity increases, there needs to be a commensurate increase in bandwidth. That breaks the mold."[3]

"The increasing reliance on cameras for security presents storage costs that can easily spiral out of control," according to Steven Titch. This is why server manufacturers have been scrambling to develop better solutions for video surveillance.

[1] http://netcentricsecurity.com/Articles/2009/04/01/Building-Better-Storage-Solutions.aspx

[2] http://ipvideomarket.info/report/advantages_of_raid6_over_raid5_for_video_surveillance

[3] http://netcentricsecurity.com/Articles/2009/04/01/Building-Better-Storage-Solutions.aspx

The problem with bandwidth is even more severe, since almost any large deployment of cameras will strain available WANs and wireless networks. Plus, remotely accessing video via the Internet is an increasing need. This often forces video storage closer to the edge of the network, into smaller servers placed at local sites.

However, even on local area networks, where adding equipment to expand bandwidth is simple, IT managers are still concerned about the demands made by continuously streaming video cameras. They often insist on separate networks for video, to isolate the cameras and protect their enterprise information systems.

Security managers are generally just as concerned, because they cannot afford to lose video recording if the data network goes down, for maintenance or any other reason. For these two reasons, added networking costs are often required.

Hard drive failures are by far the number one cause of equipment failure with security video. If you centralize video storage, then a single point of failure puts at risk the data recorded from 16, 32 or more cameras. For this reason, IT managers require RAID storage and sophisticated management systems that can automatically redirect video streams during storage node failures.

However, traditional RAID 5 storage approaches are often inadequate. "This issue is unique to video recording and seldom surfaces in RAID system used by other applications," according to Carl Lindgren of Sycuan Gaming Commission, "The key is that for most applications, written data is 'verified' during the write process." But the continuous non-stop nature of video doesn't allow time to verify. "A drive could happily chug along writing data that is unreadable for a long time. Neither the system nor the operators would ever know that there is a problem," says Lindgren. If an error occurs on more than one disk, a RAID 5 system cannot recover the lost video. This problem is forcing the move to more expensive RAID architectures.[4]

On top of all these issues are the skyrocketing costs for installing new data centers and maintaining them, along with the double digit growth in the numbers of servers needed in each data center, and rapid growth in the total number of data centers.[5] While servers once consumed about 50 watts, before the year 2000, they now draw about 250 watts each. In addition, according to Intel, you now need another 170 watts to cool each server today.[6]

[4] http://ipvideomarket.info/report/advantages_of_raid6_over_raid5_for_video_surveillance

[5] Data centers: How to cut carbon emissions and costs, from McKinsey on Business Technology, 2008

[6] Green-memory movement takes root, by Mark LaPedus, May 18, 2009 issue of EE Times

The Solution

Adding intelligence and storing video in the cameras resolves most of these issues:

1. It solves the bandwidth problem
2. It reduces storage costs
3. It makes video recording immune to network down times
4. The whole system becomes simpler and more scalable
5. Maintenance costs and data center costs are reduced

The Bandwidth Problem:

Traditional IP cameras stream massive amounts of data across the network for recording at a central location. However, as mentioned above, only 1% of that recorded video is ever played back. This means you are taxing your networks 100X more than is necessary. You are streaming all that video across your network for the 1% that you actually need.

This is the waste created by trying to force video systems into a data centric model.

Bandwidth may be getting cheaper on LANs, but the problem with video across networks is getting worse. Why? More and more network cameras are being added, with a growing demand for higher resolution and faster frame rates, and the increasing need to access the video across WANs, the Internet and wireless networks. The rising importance of cameras for security and business management makes this a growing concern.

Anyone who says that there is no bandwidth problem is not facing the facts. If bandwidth and storage were free, as we hear occasionally, then everyone would be recording video at full resolution and 30 frames per second. However, most video systems today only record at CIF resolution (352 x 240 pixels) at 3-5 frames per second, even though the same cameras are capable of four times that resolution and 6X-10X the frame rate.

The advantages of faster frame rate and higher resolution are significant: It is far easier to identify what has happened and who the intruders were. High quality recorded video is often crucial for evidence in courts of law. The only problem is the cost of bandwidth and the cost of storage. This is why the quality of recorded video is compromised.

Most people don't realize how limited their bandwidth is. For example, many small businesses, auto dealerships and construction sites use DSL service for Internet access. DSL providers advertise download speeds from 1 Mbps up to 20 Mbps, which sounds impressive. But it is the upload speed that matters when you install video cameras at these sites. The fastest upload speed from these same providers is 896 Kbps, and most of the DSL plans offer only 384 Kbps.

Standard resolution cameras using the latest H.264 compression need about 500 Kbps if you want full motion video, or 150 Kbps at lower frame rates. MPEG-4 and MJPEG compression needs even more. So, how many cameras can you stream at the same time? Two, if you are lucky.

However, when you store video right in the camera, no bandwidth is needed for recording.

Another benefit impacting bandwidth: When video is recorded in the camera, you don't need to stream the video; you can send it as a clip. This makes a big difference. To stream video, you need the full bandwidth required to play that video in real time. When you send a clip, however, you can transmit the video at much lower data rates — it simply takes a little longer for the file to get through, like sending an email attachment from a cell phone.

This is why storage in the camera is ideal in wireless applications, even over cell phone networks.

In addition, most video streaming is done using protocols that don't guarantee all video data is accurately transmitted — you have to live with occasional glitches. If you use multicasting, you can have even more problems with dropped packets and lost video. Sending files and video clips, on the other hand, uses TCP protocols that guarantee lossless transmission.

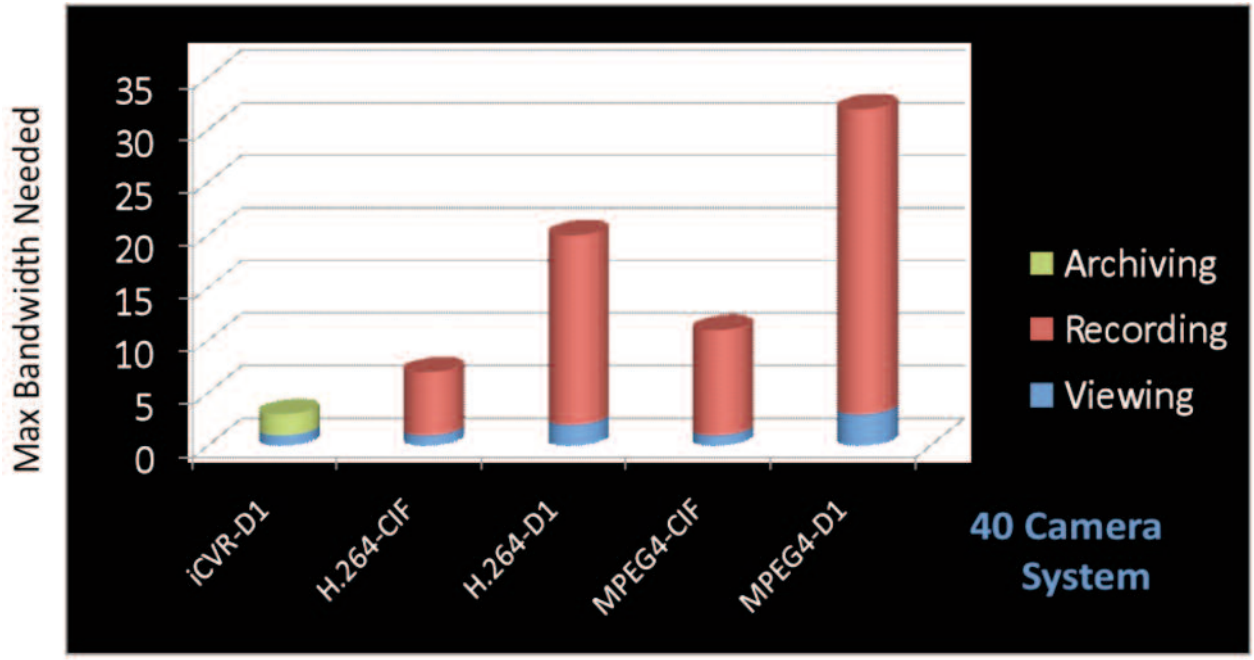

Fig 1. *This illustration shows that VideoIQ's iCVR™ requires no bandwidth for recording, because it records in the camera. If you want to centrally record video (not usually needed), you can still do so, but you can archive at night when network usage is low. So, the bandwidth requirement for the iCVR™ is similar to adding another laptop on the network. There is no friendlier video solution for IT.*

Reducing Storage Costs:

If only 1% of recorded video is needed, wouldn't it be nice to know what that 1% is when it is happening?

Adding intelligence into the camera gets us closer to this ultimate ideal. Video analytics adds the ability to automatically recognize what is happening in the scene and whether it is important or not. Installers can define what should be recorded through easily configurable rules, customized to the end-users needs.

Fig 2. *The picture on the right shows an intruder on a construction site, caught by video analytics.*

Unfortunately, the problem isn't quite that simple, because sometimes it is what doesn't happen that is just as important. For example, a shopper sues a store claiming that they slipped on spilled liquid, or their car was damaged in a parking lot by a company vehicle, or a manager mistreated another employee. Sometimes it is important to show that such things never happened, and for this reason it is important to continuously record the video.

However, what we can do, as a result of having intelligence in the camera, is record at higher resolution and higher frame rate whenever something important might be happening, and we can retain that important video longer. This improves the value of the stored video, without incurring the cost of recording at high quality all the time.

Fig 3. *The end user cost of storing full resolution video (below) at 704 x 480 pixels, 15 frames per second, for a month is $900 per camera, while storing the CIF size video, 352 x 240 pixels (on the left) at 5 frames per second costs about $100 per camera.*

That's 9X the cost and 9X the bandwidth.

Recording at high quality to capture the most important events, and recording continuously at the traditional DVR recording quality, at the same time, can save significant storage costs, especially in applications where high quality video recording is required.

However, just as important, and even more surprising is the fact that storage in the camera is inherently less expensive than traditional centralized storage. How can this be?

A single 1 TB hard drive costs only $150 these days, and you can record from 6 - 16 cameras to that single hard drive, depending on frame rate and resolution. So, how could it ever be cheaper to store video in the camera? The problem is that storage servers don't sell for $150 per terabyte. The current prices are now about $2,000 per TB or more. Why? Because the cost of managing and processing the streams of data to record dozens of cameras to a single hard drive, or to a group of hard drives, ends up costing many times more than the cost of the hard drive itself. And this still doesn't yet include the extra overhead costs for RAID storage, or for the NVR software, or the load balancing switch hardware in case a server goes down, or the cost of rack space, energy cost and the data center overhead.

On the other hand, when you only have to record one camera's worth of video to a drive, you don't need any of this data management overhead. It is simply the cost of the hard drive.

So, you may not be able to optimize the storage as well, and you might end up with more storage in the camera than you need, which you can use for higher quality recording — but it still ends up costing less.

Immunity to Network Downtime and System Reliability:

With traditional IP cameras, if you lose network communications, your system stops recording video.

Network up-time is high these days, due to rigorous IT management practices and making sure that equipment meets IT standards. However, networks do still go down. Sometimes they need to be taken off-line on purpose, during system upgrades or maintenance. But systems can also fail during critical emergencies, such as fire or intentional attacks, when surveillance video is most important.

If the network goes down, you've lost that video. There is no way to ever get it back again. However, this is not the only significant point of potential failure in an IP video system.

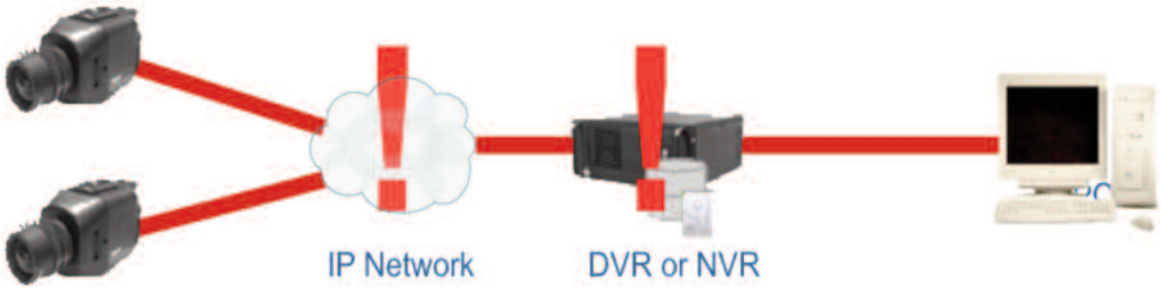

If network communications fail, your recording fails, but if your storage server breaks down, you can lose video to dozens of cameras at a time. By recording video in the camera, you solve the first problem, and you significantly reduce the impact of the second issue, since a single hard drive failure can affect one camera at most, not dozens. This increases system and data reliability.

Simpler and More Scalable:

Traditional IP video systems require all of the following systems to work together, often from different suppliers:

- Network cameras
- NVR software to manage the video recording
- Storage servers
- Video analytics servers
- Added networking equipment to increase bandwidth and isolate video from the data network

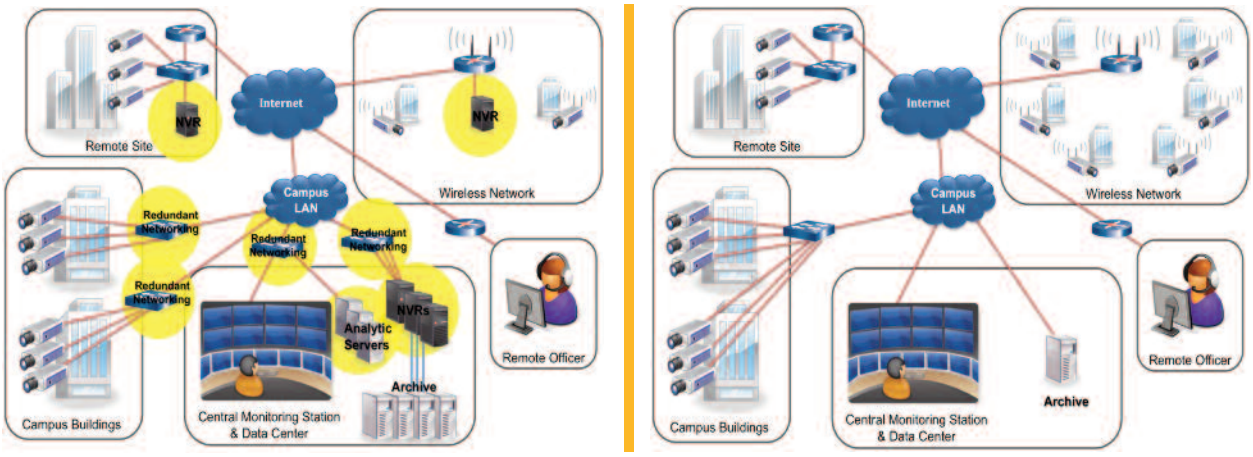

Fig 4. *On left: A traditional network video solution. On the right: Intelligent cameras with storage.*

Once storage and intelligence is built into the camera, you eliminate the need for external DVRs or NVRs. You no longer have to add equipment to increase the bandwidth or to isolate surveillance video from the data systems. On top of that, you can now add 20X as many cameras onto a wireless network, or other limited bandwidth networks.

As Fig 4 shows, you now simply need cameras and a network. It doesn't get any simpler than this.

Plus, you can add one camera at a time. You don't need to pay for expensive NVR software that is designed to manage dozens of cameras, when you only need one or two cameras. The storage comes with each camera, so you automatically add just as much storage as you need.

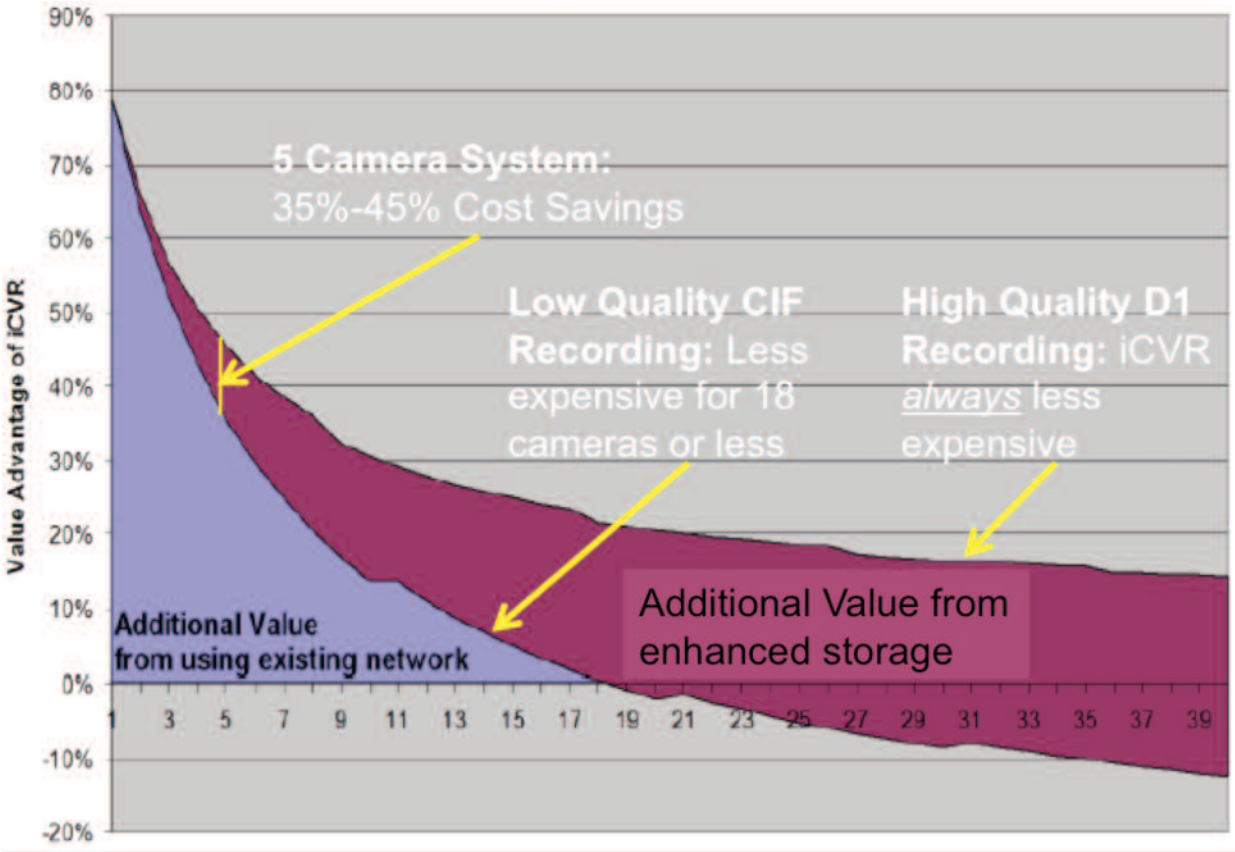

If you only need a few cameras in a location, intelligent cameras with built-in storage can't be beat. If you need high quality video, the system is also less expensive. Today, low quality recording systems with over 18 cameras can cost a little bit more, but this gap will disappear soon with future products, because storage in the camera is inherently less expensive. (Data from independent study by Tom Galvin of Network Video Consulting.)

Reduced Maintenance and Data Center Costs:

Data centers typically represent 25% of the whole IT budget, and these costs are growing as much as 20% per year — far faster than IT budgets are growing.[7] With the launch of new data centers now costing $500 million to $1 billion, saving data center costs is important. These expenses are hardly ever calculated into the cost of IP video surveillance systems; generally just the cost of the server itself is included. However, current estimates show that for every $1,000 you spend on servers, you need to spend an ongoing $1,000 per year to cover site cap-ex expenses, maintenance, electrical power for both the server and cooling equipment, and replacement costs with three-year life expectancies for servers.[8]

Therefore, eliminating the need for valuable data center space, and doing away with the need for adding servers that require maintenance and power to run and cool them, means that the value of distributed storage in cameras represents a major step forward. As said above, the reason for this is: Video surveillance systems are inherently sensor networks, not data processing systems.

[7] Data centers: How to cut carbon emissions and costs, from McKinsey on Business Technology, 2008

[8] The Invisible Crisis in the Data Center: The Economic Meltdown of Moore's Law, by Kenneth G. Brill, Uptime Institute, 2007

Another important consideration is the issue created by hard drive failures. The problem is so severe that RAID storage is a requirement in data centers. Hard disks are added into servers to gain storage redundancy. This improves reliability, but with an added expense and overhead. Pushing video recording into the camera, however, creates an inexpensive storage array with as many disks as there are cameras. Combined with the ability for cameras to archive important critical video remotely, this takes RAID to a whole new level.

But what about the increased cost of managing hard drives now at the edge of the network?

This is an important question.

If hard drives fail in cameras every 3-5 years, as they do in servers, then this means replacing 6 – 16 times as many hard drives, if you have one in every camera. These HDDs will also be more expensive to service, because cameras are not centrally located. They are often mounted high on walls and poles, and other areas difficult to get at.

If hard drives run continuously in cameras, as they do in standard storage servers or DVRs, this would create a significant added problem. Fortunately, there is an elegant solution.

The life expectancy of a hard drive is based on its "power-on" hours. In other words, if you power on the hard drive one third of the time, then your hard drive will last about three times longer. That's good news, because there is a simple way to keep hard drives from being powered on continuously, as storage servers require.

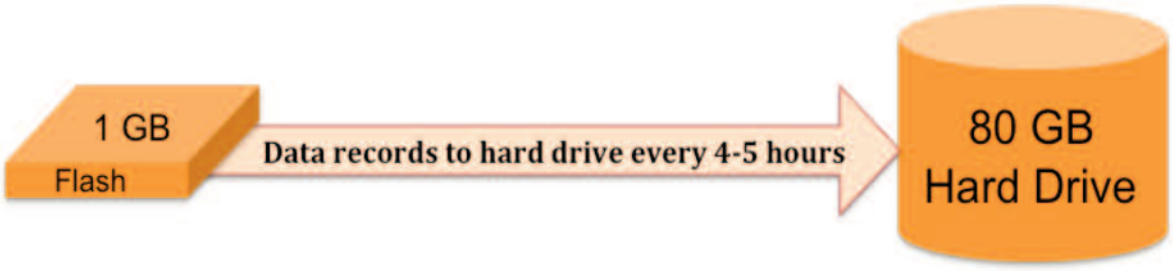

Since we only need to record video from one camera, we can use solid state memory to act as a buffer. In VideoIQ's iCVR™, we include 1 GB of solid state memory. This is enough to store about 4-5 hour's worth of video. Therefore, we only need to wake up the hard drive every 4-5 hours, and after that we can put it back to sleep again.

This means the hard drive is asleep over 90% (typically 96%) of the time. This not only significantly increases the life expectancy, but also reduces power consumption as well. Therefore, the life of a typical hard drive has been extended to last longer than the 10 year lifespan of a typical camera.

This not only resolves the problem of having hard drives at the edge of the network, it also solves the biggest maintenance expense with traditional IP video systems: The failure of hard drives.

The Benefits of Intelligence

As we can see above, intelligent storage in the camera solves the most significant problems challenging network video systems, and it significantly reduces equipment, installation and maintenance costs. It is also more scalable and simpler to use, and much more network friendly and fault tolerant.

VideoIQ's iCVR™ is the industry's first IP camera with built-in DVR and intelligence

The biggest benefits of intelligence, however, lie in the ability to transform surveillance into proactive protection. This is far more valuable than simply recording video for playback after an incident, because proactive alerts allow you to stop crime before it happens. See the whitepaper: Video Analytics Breakthrough Creates New Market: www.remoteguarding.org/whitepaper.pdf

For all of these reasons, an architecture of intelligent cameras with built-in storage will eventually transform IP video surveillance.

However, as with any new design, questions always come up: What new problems does this architecture create?

The most common concern raised is over the danger of cameras being damaged or stolen. You could lose valuable stored data, since it is more vulnerable at the edge of the network.

Most video integrators admit that attacks on cameras are rare, and it is more likely to see DVRs damaged or stolen — still this is a valid concern that needs to be addressed. We want the data protected.

Fortunately, video analytics in the camera provides the solution. Since the camera is smart enough to know what is an important event, it can send an alarm message along with a video clip to a remote monitoring station. Even if the station is not manned, the video clip is redundantly stored remotely from the site.

These video clips are sent from the camera within seconds after an event is detected. If cameras are looking out at the perimeter of a site, for example, and someone approaches the camera to attack it, their image would be captured, recorded, and then sent to a monitoring station for safe keeping. If you want even higher security, you can mount each camera so that they not only oversee the site, but also watch each other. In this way, anyone trying to attack a camera would be captured on video by another camera.

In some cases, redundantly archiving video to a central location is also desired. This isn't typically needed in most applications, but in high security jobs where sabotage is a real threat, archiving should be considered to protect the video. This is where we discover another benefit of having video analytics in the camera: You don't have to redundantly record all of the video if you don't want to. You can just record the video that the intelligence identified as important. Or you can record both the continuous video and the alarm event video, but keep one longer than the other, which reduces storage requirements.

Intelligence in the camera also warns you if the camera has been tampered with or moved. Scene change alarms are sent to the remote monitoring software to alert you, adding even more reliability to the video surveillance system. Wouldn't you like to know if someone covered a camera or aimed it in a new direction? How secure are traditional surveillance systems?

Watch Out For Imitations:

As with anything successful, there are always alternatives claiming similar benefits. Let's look at a few, so you know what to avoid.

1. **Analytics but no storage in the camera.** Some companies talk about the storage savings benefits of video analytics, but they claim to get the same results without storing video in the camera. These systems only record video when the analytics detect something important. Be careful, since this is not at all the same.

Here's the first problem: Once the rules for their video analytics decide not to record something, you can never go back to look for something else. You might realize later that you need to search for people near a back door, when you didn't think that was important beforehand. Or you might want to review not only the people climbing a fence, but cars that parked or drove by to scope the fence out. After a break-in or serious security breach, your idea of what is important can change.

Secondly, as mentioned above, if you record only alarm events, you won't have any video recorded to show what didn't happen. And, third, if the analytics ever missed anything important, even if it is only rarely, you would have no video of what took place.

There are some jobs where continuous recording is not needed, but this isn't true for the majority of video surveillance applications.

2. **Analytics in the central server.** Some companies talk about bandwidth savings benefits and claim that you can get the same results with video analytics running on a central server. They might do some of the analytics processing in the camera, so they don't need to stream video all the time. They only stream metadata, which takes a lot less bandwidth. However, the final detection decision is made in a central server.

The problem here is the same as above: You need to record continuously in most surveillance applications, if you want true coverage. If you record all the time on a central server, you gain no bandwidth advantages.

Also, having the central server tell the camera that it has detected something important, and to now send the video — after the event took place — also creates other problems. How long does it take to do this detection and send the message back? Do you buffer video in the camera? If not, then you will not be able to capture video of the event itself, but only immediately after the event. What if there is a network interruption or delay?

What some companies claim is that their server based analytics architecture is necessary because high quality video analytics requires too much processing power to embed in a camera. There is some truth to this when it comes to other video analytics technologies, but VideoIQ's approach runs about 8X more efficiently, so it easily runs in low cost embedded processors. If you can run high quality analytics cost effectively in the camera, it is clearly the better approach.

The cost of embedding processing continues to drop, as processing power increases, so this will only become less of an issue over time.

Many systems go even further and run all of the intelligence in a central server. This makes the bandwidth problem even worse, since you now need to stream all of the video back to your data center. More importantly, you don't want to stream back low quality video, because your ability to do good detection will be severely limited. So, this means streaming back full resolution, high frame rate video, which often requires far more bandwidth than traditional IP camera systems. That's going the wrong direction.

There are of course applications where running analytics on a server does make sense, for example, if you have an installed system with central recording and simply need to add detection. VideoIQ has sold server based analytics for years, for this reason. However, there is little doubt that the future lies with embedding analytics in the cameras.

Smarter cameras make the whole system more intelligent.

3. **Motion detection instead of video analytics.** One of the most disappointing problems to see these days is when well established, high quality camera manufacturers, try to pass off advanced video motion detection as if it were video analytics. Don't be fooled. This is the biggest lie in today's intelligent video market.

Advanced video motion detection systems don't actually recognize what a person or vehicle looks like. They only see blobs of pixels and guess that it is a person or a car by the size of the blob, or the way it moves. In outdoor environments, these kinds of systems produce orders of magnitude more false alarms than true video analytics. They will also miss real alarms far more often.

Advanced video motion detection also requires careful camera calibration and tuning, or it won't work at all. So, they cost far more to set up, install, and maintain over time, and they are simply unable to deal with highly dynamic environmental changes, which are common outdoors.

If such cameras were a lot less expensive, it might make sense in some limited cases, for example indoors, but in most cases our iCVR™ is less expensive, when total cost is compared.

See the blog post for more info: http://spotonsecurity.com/2009/02/06/the-big-video-analytics-lie/

Over time, I'm convinced that these bad imitations will disappear, as video analytics technologies improve and processing power in the camera gets cheaper. But it is sad to see reputable companies telling customers that what they are offering is video analytics, when it is merely advanced motion detection.

4. **Flash memory in the camera.** You might see cameras with built-in recording, and some even call these cameras with built in DVRs, but they use flash memory for storage. The problem is that their recording capacity is limited to hours, or sometimes a few days. This isn't even close to enough recording for the typical 30 days of storage required by most surveillance systems.

This means that you can't eliminate the DVR or NVR in the system. So, the onboard camera storage is merely a buffer and not a DVR at all.

It might provide some real benefits, such as in applications where continuous recording isn't needed, or to provide protection against network downtime. But even here, be careful about how it works. Some cameras require you to go around to the cameras and collect the flash memory cards, if you want to playback the captured video.

Some cameras allow you to access the stored video on the flash memory through the network, but what you really want is video that is automatically a part of your NVR recorded video, so that it can all be accessed as one stream of video.

The Future

The benefits of intelligent cameras with built-in DVRs are compelling today. As many users have pointed out, all of these advantages become even more significant in the years ahead. This is why so many now recognize that it represents the IP video system of the future.

Let's take a look at some of these important trends:

Solid State Memory is Getting Cheaper and Smaller:

While today's cameras with flash memory can only store hours or a few days of video, it is clear that in a few years, advances in solid state memory will make it cost effective to have a complete DVR in the camera. Today it is too expensive, but over the next two to three years that is going to change.

This solves the hard drive failure problem completely. Who is going to want to store video centrally, then?

Of course, we will still have the same vulnerability of storing video in the camera, which is why intelligence is still needed to send copies of critical video to a remote location for redundant storage.

It is the combination of both intelligence and storage in the camera that makes this work.

The Growth of Megapixel Video:

If there are problems with bandwidth and storage with standard resolution cameras, think of how much more of an issue this is with megapixel cameras.

Even the lowest resolution HD camera — the 1 MP camera — is about 4X the resolution of a standard camera, while a 5 MP camera is 20X the resolution. That means that if you wanted to record at the same frame rate and the same quality, it would take 20X as much bandwidth and storage space.

It is worse than this, however, if you aren't using the latest compression technologies. Most megapixel cameras today are not, because chips are only now becoming available to compress such massive amounts of video data. A 5 MP camera using MJPEG compression can require more than 100X as much bandwidth and storage as a standard res camera that uses H.264 or MPEG-4 compression.

As you can see, the power of intelligence and storage in the camera is even more compelling with megapixel cameras. It is also more challenging to develop a good solution.

Security as a Service:

Web services for video surveillance have become a hot topic these days. Huge projections in the growth of cloud computing and Software as a Service (SaaS), as a way of outsourcing data processing, has created a rush to develop video security applications to ride this wave.

One leading IP camera company has been promoting this concept and projected that hosted video might account for as much as 25% of all video applications in 5 years.

Once again, the assumption being made is that networked video should follow the data center model. As shown above, this mistake adds unnecessary expenses and added complexity. This is even a bigger problem for hosted video applications.

In hosted video systems, IP cameras are installed at a site and stream video continuously back to a remote server farm, where the video is stored and available for access through web browsers.

Here is the claim: It is cheaper to store and manage the video centrally. You save cost by not needing a DVR at the site. Plus you need no software in your PC, simply a web browser. Having the storage and software managed for users, especially small companies who don't have IT staff to do this, makes the hosted video solution less expensive.

In fact, it is quite easy to show that storage in the camera is less expensive than storage in a server farm, because you don't need the added cost of trying to manage hundreds or thousands of streams of video. This recording overhead is many times the cost of the hard drives. However, the biggest downfall of the hosted video approach is the cost of bandwidth.

Streaming all of the cameras continuously across the Internet almost always requires added upstream bandwidth at small sites, unless you severely limit the quality and frame rate of the video. Plus the data center needs to add substantially more bandwidth to account for thousands of cameras streaming continuously. The added cost of this bandwidth over the first year is generally many times more expensive than the cost of storage in the camera. And this cost of bandwidth is ongoing. You have to pay for it every month — and on top of that you will end up paying for the higher storage costs in the data center as well.

What this shows is that the future of managed video will be based upon intelligent cameras with built in DVRs. This significantly reduces the overall cost, while increasing significantly the quality of the video recorded.

Conclusion

Adding intelligence and full DVR storage capabilities into cameras significantly reduces bandwidth and storage requirements, along with data center costs. It also reduces the complexity of the installation, makes the system far more scalable, and decreases maintenance issues due to hard drive failures. System reliability and uptime are also improved. Future trends and technology make this architecture even more compelling.

All of these advantages are gained because traditional networked video systems follow a data center model, while sensor networks run far more efficiently with distributed memory and intelligence. This is why intelligent cameras with storage will transform IP video surveillance.

VideoIQ, Inc.
213 Burlington Road
Bedford, MA 01730

781.222.3069
Toll Free: 1.888.351.1758
Fax: 781.271.0275
info@videoiq.com

VIDEOIQ

SMART VIDEO SURVEILLANCE

www.videoIQ.com