IPVM Advanced Course Key Lessons

(Complete)

Version 2.0 March 10, 2013

Table of Contents

Resolution / Frame Rate Class 3

WDR / Low Light Class 5

D/N vs IR vs Thermal 7

PTZs vs Panoramics 9

Camera Form Factor Selection 12

Lenses 14

PPF / Image Quality 16

Camera Installation 18

Streaming / Recording 20

VMS Basics 23

VMS Advanced Topics 25

Selecting VMS 27

VMS Hardware 28

Wired Networking 31

Wireless Networking 34

Video Analytics 36

Resolution / Frame Rate Class

Resolution does not equal quality in the surveillance industry. The term resolution, within video surveillance, almost always refers to pixel count. Also, pixel count does not equal quality.

Increasing ~~resolution~~ pixel count can help increase coverage area but never in a straight line (i.e. tripling the resolution will almost never allow for tripling the coverage area).

Increases in coverage area are typically far less than increases in pixel count, being constrained by low light and WDR performance issues as well as limitations on getting the right angle of incident to the targeted subjects/objects.

Resolution, as used in surveillance, refers to the total number of pixels a camera supports (i.e., a camera with 3MP 'resolution' has 3 million pixels).

Resolution / pixel count is broken down into two components – horizontal first, then vertical (e.g., 1280 horizontal x 720 vertical, 1920 x 1080, etc.).

The two HD resolutions commercially available are 720p and 1080p. There is no such thing as 1280p. There are some "2MP" cameras, mostly older, at 1600 x 1200 (which is similar to 1080p but at a different aspect ratio).

720p cameras are 1280 x 720. 1080p cameras are 1920 x 1080. Remember that the 720p/1080p numbers refer to the vertical count (not horizontal).

HD resolutions support a 16 x 9 aspect ratio (e.g., 1280 x 720, 1920 x 1080). Often HD resolution cameras will support a 4 x 3 aspect ratio as well with increased vertical area (e.g., 1280 x 960, 1920 x 1440).

Doubling the resolution / pixel count of a camera only increases the horizontal pixels by ~50%. For example:

> 5MP is 2592 x 1944 while 10MP is 3648 x 2752; horizontal goes from 2592 vs 3648

This is an important point when dealing with PPF.

Compression / quality level changes can result in a lower resolution camera delivering more visible details than a higher resolution one, even in ideal, evenly lit conditions.

Professional cameras generally offer settings to adjust video quality (beyond adjusting resolution). Naming for these settings varies, including: compression, quantization and quality.

The higher the quantization (or q level) the more compressed an image is and the more likely that meaningful details will be lost.

In low light conditions, even with street lights or IR illuminators, multi-megapixel cameras will typically not capture significantly more details than 720p or SD ones as noise from gain control will limit advances. However, during the day, the gap in performance is typically far greater.

The frame rate cannot be any faster than the shutter speed. This is a practical problem at night for cameras that want slow shutter to take in more light. For example, a camera using a 1/6s shutter cannot have higher than 6fps.

Frame rates of recorded video can be reduced over time to save on storage. However, not all systems support this and many that do have limitations on the granularity of the drops (e.g., some only allow drops from full to 1fps).

A 5MP camera has the equivalent number of pixels as 15 analog cameras but, in practice, replacing more than 3 or 4 analog cameras with a single 5MP is unlikely.

While cameras routinely stream up to 30fps, the majority of professional users record at 10fps or less.

WDR / Low Light Class

WDR performance is measured in dBs. However, do not use dB ratings in purchasing decisions as the measurements are not standardized and subject to gaming by manufacturers. For example, just because manufacturer A rates their cameras at 58dBs while B rates theirs at 111dB does not guarantee A will outperform for WDR.

WDR and HDR typically refer to the same capabilities.

BLC or back light compensation is something completely different, and typically far less powerful than true WDR/HDR capabilities. BLC generally just adjusts the exposure slightly and can make other parts of the scene over or underexposed.

Beware the difference between 'fake' digital WDR and 'true' multi-exposure WDR. All top performing MP WDR cameras today use multiple exposures.

NEW: In 2013, emergence of WDR enabled HD sensors are making 'true' WDR capabilities increasingly common.

WDR cameras often have worse low light capabilities as WDR processing often requires manipulating the exposure lengths. Some camera manufacturers overcome this by turning off WDR at night, where it typically has little value anyway.

More pixels / resolution effectively increases WDR capability of a camera. For instance, a regular 5MP camera will deliver more details in bright sunlight than a WDR optimized SD camera. However, IPVM testing shows that a WDR optimized HD camera will typically outperform a regular 5MP one.

It is impossible and dangerous to determine low light performance by looking at a still image with no details about settings. Manufacturers frequently play games with the length of the shutter to fool users.

Even if two manufacturers use the same imager/sensor, low light performance may vary depending on the image processing, gain control levels and shutter speeds they allow.

At night, almost all cameras will automatically lengthen their shutter speed to maximize the amount of light the camera takes in. For instance, during the daytime, the shutter speed may be 1/1000s or 1/3000s but in dark conditions, the camera will typically lengthen it to 1/30s or even less.

If you allow a slow shutter (anything less than 1/30s, like 1/15s,1/6s,1/2s, etc.) you can increase the brightness of an image. However, moving objects (people walking, cars driving, etc.) will be blurred and look like ghosts.

Beware of sens up, billed as a low light feature. It is a marketing trick for slowing the shutter.

Higher resolution cameras often have poorer low light performance than lower resolution ones. However, this is not guaranteed. A lower resolution camera with a higher F stop, smaller image size, inferior image processing, etc. can result in worse low light performance than a professional higher resolution one.

For low light, chip choice (CCD vs CMOS) does not guarantee superior performance. However, almost all IP cameras today are CMOS, including top low light offerings.

Thermal cameras can 'see' much farther in very low light conditions than conventional cameras. However, thermal cameras are more costly, have limited lens options and are poor are capturing fine details like faces and license plates.

Comparing the minimum illumination ratings (lux) of two different manufacturers should not be done as the measurements are not standardized and subject to gaming by manufacturers.

Comparing the minimum illumination ratings (lux) of two cameras from the *same manufacturer* is useful to determine which is better at night. However, determining how much better cannot be done from simply comparing the lux numbers.

D/N vs IR vs Thermal

Day/night cameras typically cost modestly more than their color only version as the day/night model include a mechanical cut filter while the color only has a fixed filter.

Dual imager cameras are not widely available. While performance may be slightly better than a single imager camera (due to having separate color and b/w imagers), the significant cost premium typically makes it hard to justify dual imager cameras.

Setting day/night on a schedule can eliminate repeated 'flipping' between day and night modes when motion based lighting goes on and off. Even though cameras automatically can detect light levels, using a schedule can help avoid imaging issues when cameras repeatedly switch back and forth between modes.

Cameras generally offer a control to adjust the sensitivity of when a camera cuts over between day and night modes. Users that prefer cameras staying in color longer (or vice versa with b/w) can adjust this.

If the scene is too dark, auto back focus can actually cause the camera to lose focus entirely until the scene is bright enough to re-focus again.

Adding IR to multi-megapixel cameras will typically provide better image quality than using D/N or Sens up cameras. However, integrated IR typically has short ranges while add on IR illuminators can be costly.

Three commonly cited problems of integrated IR cameras are: attracting bugs, short coverage ranges and blooming/glaring of objects in the near field.

Thermal cameras help analytic performance by eliminating shadows and glare from headlights. However, moving trees or leaves can still be 'seen' and potentially trigger false alarms.

Two of the most common problems with add on IR illuminators are FoV mismatch (the illuminator FoV may be too wide or narrow compared to the FoV of the camera) and objects in the near field (close to the camera) may be overexposed.

An IR illuminator specified for 100 meters may provide even illumination at 60 meters but overexposure / wash out for close subjects at 20 meters because of the strength of the IR spotlight needed to reach 100 meters.

Increased gain control levels raise visible noise on images. This frequently increases bandwidth consumption significantly.

IR illuminators may reduce bandwidth on VBR cameras by enabling cameras to automatically lower their gain control levels.

Setting a cap / maximum for VBR cameras can reduce night time bandwidth without having to add IR lights. This will not brighten the image but it will eliminate wasted bandwidth/storage.

For short ranges (~10') and low budgets, integrated IR cameras are the best choice for low light time imaging. For medium ranges (~50') and budgets, day/night cameras with added IR are typically best. For longer ranges (~300') and larger budgets, thermal is the best choice.

PTZs vs Panoramics

When trying to determine how far a PTZ can 'see', ignore the optical zoom range (e.g., 34x or 36x) and focus on the minimum horizontal FoV angle (2.2° or 1.7°). The zoom range compares the widest to most telephoto Field of Views and can be skewed if a camera has an especially wide max.

Differences in optical zoom ranges typically have minimal impact on maximum distance capable. For instance, a 36x PTZ may only see 3 to 5% farther than a 34x one. Even if zoom range varies by 2x (34x vs 36x), the difference in maximum viewing distance will never be close to 2 times.

Digital zoom is irrelevant to PTZ performance. Only optical zoom matters. Digital zoom is more of a marketing trick than any practical difference.

PTZs may not have a wide enough FoV for certain applications. This cannot be determined by looking at the optical zoom range (e.g., 20x, 36x, etc.). One needs to check the details of the specification for the maximum horizontal FoV in degrees (e.g., 50°,60°, etc.). Because PTZs often emphasize viewing objects far away, the near field / wide angle is often sacrificed.

Physically connecting IP PTZs is typically far easier than analog ones as video and control is transmitted over a single network cable instead of two cables needed for analog (for video and serial control). However, like analog, risk remains that a VMS may not support the PTZ control commands of a PTZ. Check for specific model support.

The larger the area covered and the more frequently live operators are available to monitor, the more likely that PTZs provide value.

To ensure that PTZs are not left looking at an uninteresting or low value area, home positions should always be configured. The home position is a default view, preset or tour that the PTZ will return to after a period of time of inactivity.

PTZ tours enable PTZs to move on a predefined path among many spots across the cameras coverage area. The upside is that, over time, the camera can 'see' more things than if it was left in a single fixed spot. The downside is that, it will likely miss any momentary incident since it only covers any individual spot for a fraction of time.

In IPVM testing, top MP cameras can see as far or farther than top SD cameras. While MP cameras have lower optical zoom ranges (typically 20x vs 40x), the increased resolution of the MP cameras more than offsets for the lower zoom range.

PTZs have the potential to see significantly farther than even 'super' megapixel camera on the market today (e.g., 10MP, 20MP, 29MP, etc.) because PTZs support much greater telephoto lenses than super MP cameras. For instance, PTZs routinely support 3mm to 100mm lenses while super MP cameras are typically used with fixed lens of 10mm or less. The PTZs's far greater maximum lens length (100mm+) more than offsets its lower resolution.

Fisheye panoramic cameras typically will deliver poor details at even fairly closely range (e.g., 3 meters/ 10 feet). Despite the use of multi-megapixel imagers, the exceptionally wide FoV means that the pixel density at any given spot is quite low.

Many fisheye panoramic cameras require PTZ systems to dewarp the distorted fisheye view inside the VMS's client. While this increases the usability of the camera, it requires the VMS to commit to extensive custom software development. This has significantly held back fisheye panoramic support and is an important limitation in evaluating offerings.

Fisheye panoramic cameras routinely suffer from poor WDR performance because the camera has an ultra wide FoV yet still only a single exposure level. If any part of the 360 FoV faces bright light or darkness, those parts are likely to be rendered poorly. If the camera chooses a fast exposure, the bright areas will be captured optimally but the darker areas will be significantly underexposed. The opposite will occur with a slow exposure, with bright areas getting washed out.

Panoramic camera options in the market are now quite significant with most major manufacturers offering models. This was not the case until the past year where a marked increase in releases occurred.

PTZ cameras are typically significantly more expensive to purchase than even 5MP or 10MP cameras. In addition, while multi-megapixel cameras will cost more for storage (due to higher resolution), PTZ cameras typically have higher service costs to maintain moving parts.

Multi-imager panoramics have many technical benefits over fisheye panoramics but typically cost notably more. Advantages of multi-imager panoramics include easier VMS integration (as no dewarping is needed), higher pixel density (as multi MP imagers can be combined together), and more immune to WDR problems (as each imager can set its own exposure).

The two best scenarios for replacing a PTZ camera with a panoramic one are: (1) when no one will monitor/control the camera live and (2) when the FoV is relatively small enough that the panoramic camera can capture sufficient details.

Camera Form Factor Selection

Cube cameras are the lowest cost form factor, most commonly used in home/small office applications. Professionals typically do not use them because they have poor low light and WDR performance and limited FoV/lens options.

PoE power is now widely supported amongst almost all IP cameras. Cube cameras are most likely not to support PoE to reduce cost for users unlikely to have PoE switches/injectors. Some IP PTZs lack PoE support but it is increasingly common to use High PoE.

Integrated Wireless is rarely available for IP cameras. However, cube cameras are the most likely to support integrated wireless as they are targeted to consumers who have WiFi in their homes/offices.

For aesthetics, domes are the most commonly preferred form factor. Within domes, mini-domes specifically are a top choice due to their streamlined appearance and small size.

While dome cameras are more popular than box cameras, box cameras offer more flexibility in lens selection and positioning/aiming of the camera. For examples, in applications where a fixed camera most monitor objects far away (e.g., 30 meters+), box form factor is likely required.

Vandal resistance is common in most form factors except for cubes and box cameras. While a vandal resistant housing can be added to a box camera, typically locations requiring fixed vandal cameras use either domes or bullets.

Including all PTZs, most cannot pan 360 degrees (i.e., left/right). Overwhelmingly lower cost PTZs support far less than 360 and even some mid-tier PTZs max out at 359 degrees. This can become a problem for live operators as it can inhibit the ability to smoothly and quickly track suspects.

Most PTZs cannot tilt 180 degrees (i.e., move up/down). Limitations in tilting can cause practical problems if subjects walk/move underneath the camera.

Vandal resistance is rated on an IK scale from 1 to 10, with high numbers denoting greater vandal resistance. Additionally, a non-standardized metric of 10++ exists.

When specifying vandal resistance, carefully consider the actual IK rating, keeping in mind that: (1) A huge drop exists from IK10 to IK7 (90% less vandal resistance) and (2) manufacturers are not required to validate these ratings, meaning that it is possible to fake/fudge them.

IP ratings are used for specifying the weather resistance / outdoor applicability of cameras. For professional applications, use at least IP66 and avoid IP54.

When specifying audio, be clear what type of audio is desired as available audio options differ on IP cameras including: mic only, speaker only, unidirectional (e.g., walkie talkie) and bidirectional (e.g., telephone).

Overall, minidomes are more commonly used than full sized domes due to a preference for small size and low costs. However, higher end professional applications often prefer full sized domes for superior image quality and the availability of varifocal lenses.

Lenses

The longer the lens, all things equal, the narrower the FoV will be. Conversely, the shorter the lens, the wider the FoV is. For example, a camera with a 3-8mm lens will have its widest FoV at 3mm and its most telephoto FoV at 8mm.

Varifocal lenses allow for adjusting the Field of View of a camera from wider to more telephoto. The exact FoV will depend on the lens lengths supported (e.g., 3 – 8mm varifocal vs 5 – 50mm varifocal).

Varifocal lenses have at least two physical controls – one to adjust the zoom (wide to telephoto) and another to adjust the fine focus. Additionally, manual iris varifocal lenses will have a third control allowing for the opening and closing of the iris.

Adjusting the iris impacts the amount of light that the camera/imager receives. It does not impact the FoV.

Auto back focus / auto focusing enables the camera to automatically fine focus the image without a human intervention. It is only available on a minority of IP cameras today and typically those are premium / professional versions.

Difference between auto back focus and auto focus:

- Auto back focus adjusts the imager back and forth inside the camera body. This is the typical implementation for box cameras. It can only adjust the fine focus, not the zoom.
- Auto focus adjusts the lens back and forth. This is more common in dome cameras and often allows for both fine focus adjustment and remote zoom.

Auto back focus typically requires the lens to be in rough focus because of limits of how much adjustment can be made to the position of the imager inside a camera.

The two main drivers of F numbers are lens length and lens diameter. All things equal the longer the lenses, the higher the F number and the smaller the lens diameter, the higher the F number.

For surveillance, higher F numbers are typically a problem, specifically for low light monitoring.

F numbers are often expressed in stops. Moving from one stop to the next doubles (or halves) the amount of light). Common stops are 1, 1.4, 2, 2.8, 4, 5.6

The higher the F number, the less light a lens can take in (e.g., an f/2.0 lens takes in 50% less light than a f/1.4 lens).

Minidomes typically have a relatively high F number (f/2.0 or greater is common) due to the small form factor of these cameras. This means poorer low light performance than a box camera with a F/1.2 or F/1.4 lens.

While PTZs often list a single F number, the F number almost always increases with the length of the lens. For example, at its widest FoV, a PTZ may have f/1.8 but at its most telephoto, it might have an f/3.5.

While depth of field can be increased with higher F stops, this is generally not a practical issue in surveillance. With the wide FoVs common in surveillance and subjects typically 5 feet or more from the camera, the depth of field is practically infinite. Additionally, increasing the f stop by closing the iris would result in no usable images at night.

While fixed focal length lenses might deliver slightly more sharp images, overwhelmingly integrators prefer varifocal lenses so they can adjust the FoV on site to deliver an optimal coverage area.

Iris selection typically does not deliver significant improvements in image quality. Even P-iris, a supposedly more advanced and 'precise' control, delivered minimal benefits in IPVM tests.

Super wide angle lenses deliver a wider FoV, however the image details captured will be reduced as the same resolution is spread over a wider area.

PPF / Image Quality

If PPF is used to specify image quality, it must also specify the FoV width, and the distance this is required from the camera, including the maximum distance if multiple objects at difference points are to be monitored.

PPF = Horizontal pixels of a camera divided by the width of the FoV (e.g., a 40ppff = 1600 pixels / 40 feet)

PPF uses the horizontal, not the vertical pixel count and does not factor in total resolution. For example, a full HD camera is 2MP and has a 1920 x 1080 pixel count. 1920 will be used in the PPF calculation as it is the horizontal pixel count.

Remember, a 720p camera has 1280 horizontal pixels (1280 x 720) and a 1080p camera has 1920 horizontal pixels (1920 x 1080).

The farther a subject is from a camera, the wider the FoV will be at that point.

The wider the FoV, the lower the PPF and the worse the image quality tends to become .

Image quality tends to gradually degrade as a subject is further from a camera. Despite the common categories cited in surveillance, there is no single step between good and bad images.

PPF is good for guestimating / getting a rough sense of image quality potential. It cannot guarantee image quality as pixels are an important but not the only factor in determining quality.

Without knowing site conditions (issues with the sun, light levels at night), determining even an estimate for the number of pixels per foot to capture a face is not possible.

Presuming day time, even lit conditions, 50ppf to 60ppf is likely sufficient for most operations looking for facial details and license plates. Anything more than 100ppf is likely overkill.

Remember doubling the PPF requirement requires quadrupling the resolution required. For example, if you want 50ppf at a 25 foot wide FoV, a 720p / 1MP camera will work. However, if you want 100ppf at the same width, you will need a 5MP camera.

At night, even if you have street lights, you will always need more PPF than you will need during the day to capture the same image details. How much more depends on how dark the scene is. However, doubling the PPF needed for night vs day is common. Of course, this can require

much higher resolution cameras that might deliver even worse low light performance – a vicious circle.

Increasing resolution levels typically has more practical benefits in wide FoVs than in narrow ones. This is because, increasing PPF over a certain level (whether its 50 or 100 depending on site conditions) typically delivers no additional practical value.

Rules about people or objects needing to take up a certain percentage of the FoV should be discarded as they presume fixed resolution (e.g. analog only) and ignore lighting issues.

Be careful when people ask for 'how far can a camera see X?' whether it is people, faces, cars, license plates, etc. Remember that the width of the FoV must be decided first. With the right lens a camera can 'see' very far away but the width of FoV cannot be magically and continuously expanded.

WDR performance of cameras can impact PPF / quality specifications. Cameras will better WDR can deliver the equivalent image quality of double the pixels per foot even when the resolution and FoV width of two cameras are the same.

Objective, clearly definable, image quality levels do not exist. Reasonable people can and will disagree how much quality is enough. Two people may look at a 50ppf image in even lighting and one may find it satisfactory and another may not. Make sure that decision makers see the image quality and approve it, regardless of theoretical PPF numbers.

It is unrealistic to guarantee quality levels for most surveillance cameras unless you can guarantee that lighting levels will always stay fairly constant (i.e., no times of harsh bright sunlight or periods of deep darkness).

Camera Installation

Some form factors, like bullets and boxes, are more vulnerable to tampering than others. They can be easily redirected and should not be hung so low they can be tampered.

A tradeoff between camera mounting height vs. angle of incidence exists. Generally, the higher a camera is mounted, the steeper the downtilt and the worse angle of incident to a subject (like a person's face or license plate).

Many installers overestimate the mounting height needed to safely protect a camera and unwittingly undermine image quality/angle.

Cameras should always be mounted with the expectation that maintenance will need to be performed at some future point – future accessibility is a critical attribute to consider.

Masonry surfaces (brick, cinder block, cement) require use of a hammer drill and masonry drill bits to make mounting holes –using standard drills and general purpose bits can quickly drain or damage tooling and result in a poorly drilled hole.

Mounting cameras on drop ceiling grids is best accomplished with a plastic 'grid clamp' that allows the camera to mount directly to the metal suspension grid. If the camera needs to be mounted from the tiles, backing materials or grid mounts should be used to reinforce the tile. Failing to use one of the previous methods will often result in damaged/broken/sagging tiles.

Mounting cameras onto drywall requires the use of reinforcement anchors/bushings. Installing camera with drywall screws directly into the surface can easily pull-away or come loose, and screws alone may not support the weight of the camera without anchors.

'EIFS' – or "synthetic stucco"- can be especially challenging to mount cameras from unless proper 'backing materials' have been previously installed or can be retrofitted into place.

Concrete 'tilt-up' panel construction uses solid concrete walls. Running cabling inside these walls unless channels have been pre-cast into them during fabrication and cables are often run to cameras in conduit as a result.

Walking through EXACT camera mounting locations with a customer can help identify potential issues with camera positions, and be a convenient way to discuss alternatives when the need arises.

Taking pictures of proposed camera locations will clarify mounting instructions for installers and customers alike.

Aesthetics, or disturbing the cosmetic appearance of a space, may become an important constraint when performing camera installs. Sometimes, a camera may be positioned 'less than ideally' or alternative mounts/mounting methods may be employed in order to avoid aesthetically disrupting a space.

Ensuring outdoor installations are installed using 'drip loops' for cable pigtails will prevent moisture from entering the camera enclosure/junction box.

Considering environmental impacts like wind, precipitation, smoke/fog, and mounting surfaces on the quality of visual images is often overlooked – for example, wind can cause 'camera shake' that caused stability issues. Understanding these potential impacts beforehand can drive equipment selection or alternative mounting locations.

Streaming / Recording

Video surveillance is almost never recorded uncompressed as the amount of storage required would be immense. For instance, a 1080p/30fps HD stream is over 1Gb/s. When compressed, the same stream is typically less than 1% of that size, ~4-8Mb/s.

With CBR, the bit rate is fixed but the quality level automatically changes.

With VBR, the quality level is fixed but the bit rate automatically changes.

The complexity of the scene drives changes in bandwidth. The more objects and the greater the motion of objects in the scene, the greater the complexity of the scene and the more bandwidth is needed.

At night, cameras typically increase their gain control levels which routinely increases the visible noise displayed. That visible noise is frequently picked up as motion, increasing the bandwidth consumed on VBR streams. The increase can be significant, up to 10x daytime bandwidth.

To reduce night time bandwidth consumption, VBR with a cap can and should be used.

VBR with a cap is more bandwidth efficient than CBR because with VBR the bit rate can decrease in simple scenes whereas with CBR the bit rate is fixed regardless of how much bandwidth is needed.

Regardless of the CODEC used (e.g., H.264,MPEG-4, MJPEG), compression / quality level settings are available and adjustable (for most professional IP cameras).

Compression and quality typically define the same metric just opposite directions. The higher the compression the lower the quality; The higher the quality, the lower the compression.

If you increase compression, you will reduce bandwidth consumption but it frequently comes with visible decreases in image quality.

Increasing compression is most appropriate in simple scenes with little motion nor objects (and vice versa).

Quantization is the technical term used in video for compressing video. The higher the quantization, the greater the compression level is. In H.264, quantization scale runs from 0 to 51, with 51 being the most compressed. Typically professional IP cameras target a 27 – 30 quantization level though frequently this is not shown directly on the camera's interface.

Aiming a camera at a wall is a common but poor way for measuring bandwidth consumption as it can radically underestimate how bandwidth a camera consumes.

When using CBR, you risk either wasting bandwidth by setting a bit rate that is too high relative to the needs of the scene or degrading quality by setting a bit rate that is too low for certain periods of high complexity/motion. If you must use CBR, make sure to experiment with different bitrates and check at multiple intervals to optimize the bit rate level. While there will always be a tradeoff one way or another, you can eliminate any serious problems of far too high or too low bit rates.

Statistically, more users employ motion based recording than continuous recording. While motion based risks missing events, the significant cost savings from storage reduction drive this choice.

Motion based recording is best used in scenes with fixed cameras indoors. Outdoor cameras often have significantly higher rates of motion. PTZs on tours are typically or nearly always in motion and therefore benefit little from motion based recording.

Adjusting the sensitivity level of motion detection and defining exclusion zones can help eliminate common problems of false / wasteful motion based recording. The sensitivity level is most frequently useful at night in handling noise from gain control. Exclusion zones are most commonly used to block out areas where motion is common but not relevant to the surveillance (i.e., leaves blowing on the edge of the FoV).

Storage calculators only provide highly accurate results when using CBR and continuous based recording (because recording parameters are constant in this scenario). However, when using the more common VBR and/or motion based recording, storage calculators can routinely be off by a factor of 50% or more.

Since VBR streams can vary by 10x or more, even with the same resolution, CODEC and compressions, when calculating storage for VBR streams be sure to test the bitrates for the actual camera you are using in the type of scene you are using it.

Based on IPVM survey results, 30 days is the most common storage duration for recorded video, by a significant margin - however local situations and customer requirements may vary.

Connecting IP Cameras

IP cameras are specialized computers.

When connecting applications on 2 computers, an API (Application Programming Interface) is used. For example, connecting IP cameras and VMS software requires an API.

API and SDK are two terms frequently used together when talking about integrating computers/systems. API generally means the code / methods used while SDK (Software Development Kit) generally means the documentation and supporting tools.

There is no such thing as 'an' API for a security systems. There are always multiple APIs.

Almost all IP cameras have their own APIs that specify how to request actions from their cameras. They typically have many APIs, one each for different functionality like setting the camera's IP address, changing the resolution, requesting a live video stream, etc.

Most IP cameras, and most security devices, have far more functionalities than they have APIs. Therefore, if you desire to control / integrate a certain function, make sure that it has a specific API for that function.

Just because an API exists does not mean one can access it. Security manufacturers may block access for business or competitive reasons (to make it harder for rivals to take over their existing accounts).

Integrations using APIs can take quite some time (weeks minimum but often months). Unlike API for big web systems (like Amazon and Google), very little direct experience and support is available for using security system APIs. Also, real risks exist of finding performance issues or bugs.

Changes in an API can break integrations with 3rd parties. For example, when a vendor says, "You must upgrade to the newest firmware" to make an IP camera work with a VMS, this is typically because a recent change broke the existing API integration.

Historically, every IP camera manufacturer had its own proprietary API that each VMS had to integrate.

'Standards' are an attempt to define and use *a 'standardized' API* that can be used to connect to any IP camera or VMS.

In the past 5 years, two groups attempted to bring standards to IP video surveillance – ONVIF and PSIA. ONVIF has won with overwhelming support and implementations from manufacturers across the world.

While ONVIF is commonly used successfully in production, it does still suffer from problems and issues in integrating different combinations of IP cameras and VMSes. In the last year, Profile S has been developed to force a higher level of support for IP cameras.

VMS Basics

Encoders for analog cameras have various form factors: Appliances/Cards/Hybrid Recorders, however VMS support varies, as it does for IP cameras.

The 3 Core Functionalities of VMS are: Live Monitoring, Investigations, and Administration.

VMS installing programs can be quite large; many programs are well over 1 GB in size. Downloading the files in the field can be time consuming.

Thick Clients, or 'installed clients' have more features and offer greater management capabilities than Web Clients. Thick clients are still far more frequently used overall for video surveillance systems.

When bandwidth is restricted, many VMS will manage video streams by dropping frames (perhaps without warning). Techniques like multi-streaming different resolutions/FPS from cameras can sometimes be configured, and so can transcoding, but these options vary widely from VMS/Camera manufacturers and may not be available.

Transcoding is a bandwidth optimization method where the VMS will transform an input video stream (usually at a higher resolution/quality) into another one (typically lower quality). This helps reduce bandwidth requirements for remote clients but increases processing power on the server.

When upgrading, most VMS software require physical access to the server. VMS server software is seldom able to be upgraded remotely or 'up the wire'.

VMS camera discovery tools seldom are able to discover all connected cameras – especially in large deployments. Manually finding the camera by default IP address, bridging different subnets, or otherwise rerunning the discovery tool several times may be required.

Having API/SDK does NOT mean a VMS can be 'universally integrated' with 3rd party security systems like alarms, access control, or building management. No standards exist that define how these connections should work or be made. Each integration is typically a 'one-off' and must be custom programed.

Unless a VMS is proven to integrate with 3rd party security systems 'out of the box' or through a vendor-written add-on module, caution should be used when assuming the cost/degree of interoperability between the systems.

Software-only VMS offer greater flexibility for hardware choices, while NVR appliances typically provide simpler setup.

Service/ Maintenance Agreements for many VMS platforms require a periodic ongoing subscription to retain vendor support. If time is missed, these agreements often must be 'brought current' before the vendor is willing to support the install.

Many 'top-tier' VMS companies offer limited or no NVR appliances, favoring installation on a COTS (commercial, off the shelf) server.

The cost of replacement parts in an NVR appliance can be quite higher than stock components used in COTS servers.

'Free' VMSes are available –often via download, but use may be limited (eg: max of 8 cameras/ storage limited to 5 days) and support may cost money or not be available.

VMSes are 'driven upward' to offer PSIM-like functions as product developments progress and the market mature.

Some camera manufacturers give away licenses of their own VMS platforms to incentivize selling cameras.

VMS Advanced Topics

PSIM the 'process' is markedly different than PSIM the 'product'. As a product, PSIM facilitates the process an end-user has in place to respond to a security event.

Even companies that aren't classic 'PSIM vendors' will offer PSIM-like functionality with their products. Many VMSes offer integrations with Access Control platforms and vice-versa. While not 'true' PSIMs, many VMSes incorporate 3rd Party systems.

Historically, Access Control interfaces were used as the 'top box' (the primary user interface for several security systems) but this has changed over recent years to be VMS platforms. VMS aims to now be the central application used by security operators.

One of the complexities involved in making VMSes the 'central interface' for security management is that it must pull multiple items of access control data and manage door input/outputs.

Classic PSIMs can start at $50,000 and quickly escalate to hundreds of thousands of dollars due to custom programming and interfaces, while VMS costs and fees are comparatively a fraction of that cost.

3rd Party integration (open integration) is a significant differentiator between true PSIMs and 'enhanced' pre-integrated platforms. The ability to integrate with several video sources is a hallmark of 'PSIM", while operating with native video only is trait of 'enhanced' security platforms.

A differentiator of VSaaS versus VMS are 'plug & play' ease of camera installation, public web access to video streams, 'multi-tenancy' (or ability to discretely manage cameras), and individual billing generation.

Edge storage is managed differently from VMS to VMS – some use edge storage to 'fill in the blanks' on permanent storage (in the NVR/server) others store video permanently at the edge and simply access as needed.

Integrating 3rd Party Systems into VSaaS is difficult and requires local integration; it cannot be accomplished 'in the cloud' or in the hosted environment.

VSaaS storage has trended for being a 'cloud only' utility to taking advantage of local/edge storage. Often in order to support greater frame rates and MP resolution, VSaaS platforms must rely on NAS devices or locally connected hard drives.

VSaaS is not a synonym for 'video alarm verification'. That type of verification monitoring is more of an intrusion alarm function, because the video utilized is typically low resolution at low frame rates. Its primary value is establishing movement/presence where none should be.

One of the 'weaknesses' in VSaaS offerings is that pricing does not include channel profit – the most successful players in the VSaaS market sell directly to end-users and exclude the need for the installation channel.

VMS maker Genetec has announced plans to move certain offerings to the cloud. This signals the start of a potential migration to cloud based offerings by traditional 'host-bound' VMS players. While the risk of cloud-based offering is still not fully realized, the potential benefits for central management and accessibility are significant for users. However, practical limitations (like plentiful bandwidth) remain and will need to be addressed.

As cameras 'get smarter', VMS are being forced to integrate more deeply with cameras and offload some performance features to the edge.

'Free' surveillance systems (that take advantage of edge storage) typically only use the manufacturer's cameras. Supporting other companies' cameras runs counter to a manufacturer's aim to sell as many of their cameras as possible.

Selecting VMS

Often the best question to focus on is 'What is best for the customer?' This answer is based on a number of factors - including total cost, ease of use, hardware deployed (cameras and servers), and scalability needs.

Free VMSes are typically designed for low-camera count systems, and from a single manufacturer.

Some manufacturers offer a separate "CMS – central management software" in addition to a "VMS" platform.

Training end-users on the use of VMS is not a lucrative business and frequently is done at no-cost by integrators or manufacturer reps.

Poorly written specifications can limit VMS choices by requiring needlessly proprietary features even if not overtly specifying a product.

For end-users, best practices for selecting a VMS include a 'try before you buy' period of demos or field investigations.

Integrators are typically aligned with two or three VMS products, and specifying a certain VMS can exclude some integrators/installers based on business relationships.

There is no such thing as 'the perfect VMS' for any customer, and frequently the correct choice is based on successfully manageable compromises from 'ideal'.

Camera manufacturer branded VMS systems are not popular choices – even when low cost, the consensus is these products do not have the features and 'openness' of leading 3rd party options.

VMS Hardware

CPU and RAM are the two most commonly cited specification points for VMS hardware.

Quad core processors are needed only for large camera count systems. Many NVR appliances use dual cores with success. Typically, quad cores are found on enterprise server boxes, but the extra resources may not even be used in smaller (<20 camera) systems.

For all but the largest systems and certain 'fringe applications' like virtualized servers, there is no limitation in using i7 cores versus Xeon processors. Either processor type is used successfully for video surveillance systems. Xeon CPUs are found in server builds, while i7s are found in workstation units.

Throughput is bi-directional – meaning data flows from cameras to servers and vice versa during normal operation.

When specifying hardware, using the VMS vendor's specifications when building servers help to avoid 'the blame game' and potential support issues if they arise.

The primary benefit of using server-side motion detection is the consistent and single-source configuration/management from one interface. However, this approach consumes more CPU resources than using camera-side detection. Accounting for the extra CPU load during design stage is critical.

Since server-side detection often occurs on a frame-by-frame basis, large numbers of cameras configured to record on motion can significantly bog down a server.

One strategy for reducing server load on server-side motion detection is to configure a lower resolution stream for detection and use a higher resolution stream for live viewing.

Distributing server load is most commonly accomplished by grossly distributing numbers of cameras between servers; VMSes seldom support off sourcing or load balancing/throughput balancing between a cluster of servers.

Not all cameras will have the same level of motion sensitivity and granular configurability. Some models may be more sensitive and work better than others.

Redundancy is as much of a VMS software issue as a hardware issue – not all VMS platforms support failover or redundancy features. They typically are available in the highest-level versions of VMS platforms. (eg: large deployment/enterprise versions)

Best practice for NIC design involves using dual NICs – one for the surveillance network, one for the corporate (viewing) LAN. Many servers and NVRs are built with dual NICs, but if not the extra cost of adding an extra card is generally negligible.

For large systems, it is not best practice to use a recording server in a dual role as a viewing client – the added CPU load of decoding many video streams, compounded by non-graphics intense video cards built into servers can be a significant resource hog when using a box for both roles.

Common UPS types:

- On-line: A higher-end UPS that manages power with a double conversion from AC to DC back to AC. The extra conditioning allows for high-quality power and preserves battery life.
- Line interactive: This type performs no conditioning, but generally is built to preserve battery lifetimes above simple 'standby' types.
- Standby: This type of UPS is most common in consumer grade 'office duty' type products. Generally this type of UPS is hard on batteries.

VMSes are not commonly found for OS platforms outside of Windows; Linux and Apple based solutions are uncommon and do not have the same enterprise scale as Windows options.

Direct storage is frequently cheaper to purchase, setup, and maintain than SANs, however much larger storage volumes are possible with SANs. On a per-GB basis, SANs are more expensive until volumes range into hundreds of TBs.

NASes/SANs make good sense for larger systems than need to distribute storage asymmetrically among many connected units.

Network attached storage requires ample bandwidth to be useful, and many field locations (like Banks) have bandwidth limited internet connections.

Use of RAID for data redundancy is growing as it become less expensive to implement, and therefore is becoming more common in VMS servers.

RAID is disproportionately used in larger systems compared to smaller systems – but the gap is narrowing as storage pricing/RAID controllers continue to decrease.

Wired Networking

Labeling cable runs is an inexpensive way to avoid costly troubleshooting problems after installation.

Both ends of a cable must be labeled in order to be an effective system.

Installing cables into cable trays or j-hooks/bridle hooks keeps it off of ceiling tiles and helps prevent damage from being run in unprotected places.

Segregating video networks from other networks by color of cable jacket is a best practice, in order to avoid confusing wires among systems.

Red cabling usually means 'Fire Systems', and is best avoided for video surveillance systems.

Mapping cable runs is a critical, but often overlooked, aspect of running networks. Making sure that maps are deliberately drawn and kept not only helps troubleshooting the surveillance system, but marking those locations may help other trades from disrupting video networks.

Cable service loops should be no more than 9 feet at the rack, and 3 feet at the device. Longer lengths create clutter and extra material to troubleshoot.

BICSI Standards provide helpful guidelines for designing/installing cable networks, and cover a range of topics impacting cable work. However, standards are typically an extra cost and require membership to purchase.

"Directly Connected" cable runs to camera typically cost less than using the "Jack & Patch" method, but are less flexible. High-end installations typically specify use of "Jack & Patch", but either method can be used for surveillance systems.

Cat5E cabling can be used for connecting multi-megapixel cameras – however, both Cat5E and Cat6 support GB transmission.

Cat6 is more robust to environmental variables like RF interference, High temperatures, and Cross-Talk resistance than Cat5E. This is in part due to the larger wire diameter used when making Cat6 cables.

In order for STP, or 'shielded' cable to offer a benefit, the shield must be properly grounded during installation.

Cable runs can be significantly longer than 'straight line' measurements due to overhead obstructions, plenum, wall drops, and other unseen obstacles.

The necessity of VLANS for increasing bandwidth in an existing network is a myth; VLANs are only able to segregate traffic, not increase gross bandwidth capacity. Moreover, VLANs cannot stop video surveillance from overwhelming the rest of the network if problems occur.

VLANs can only be configured if managed switches are used end-to-end in a network.

Keeping Access Control traffic separate from video is a 'best practice' recommended by several leading vendors.

Many end-users and integrators alike prefer to use a stand-alone network for surveillance only.

Running a 'dedicated network' helps avoid shortcomings of a corporate LAN, but also helps avoid political problems/the blame game for issues with other IT users.

Shared networks causes problems among IT/Security integrators unless both parties mutually agree to work together.

Methods of securing jacks/ cables include using locking enclosures or RJ45 locking sockets that prevent cables from being deliberately removed or tampered with.

Dynamic IP address allocation is typically not favored for video surveillance applications because changes in camera IP address assignments can break connections to VMSes. Exceptions apply for advanced network configuration (such as MAC address assignments that always ensure a given device gets the same dynamic IP address).

Multicasting benefits when multiple clients are viewing / connecting to the same stream from the same camera at the same time. Since, most applications rarely have this, multicasting is typically not beneficial in surveillance. However, some large scale applications with heavy number of real time users require this.

Multicasting requires advanced support from the VMS software and networking infrastructure one uses plus custom configuration that takes time and increases complexity.

Wireless Networking

Signal vs Noise is a key concept / contrast for wireless signals. Knowing the signal and noise levels is critical in determining wireless bandwidth availability.

Signal and Noise are both measured in decibels (dBs). In wireless, it is always negative (e.g., -100dB, -60dB, etc.)

The noise floor is typically -120dB to -100dB, depending on the environment.

The close the signal level is to 0dB, the better (i.e., -60dB is stronger than -80dB).

The greater the gap between signal to noise (i.e., signal is -60dB, noise is -105dB), the more likely the bandwidth will increase. The specific bandwidth depends on the wireless system and must be checked (can be found on the system's specification sheet).

Lower frequencies have a longer wavelength, and therefore are more resilient when passing through physical obstacles like trees, foliage, or structures. Lower frequencies are commonly deployed in wireless surveillance systems as a result.

Most wireless equipment used in surveillance networks is the 'unlicensed', or low-power, variety that can be freely used but has notable constraints.

"Free" unlicensed wireless has a fraction of the transmission power available compared to licensed options.

The unlicensed bands are 'public domain' and unreserved for use, so interference with other systems/devices is possible and fairly common.

Antenna selection/design plays a key role in transmission distance and radio performance.

'Beam Width' has an inverse relationship to 'Gain'; in order to increase one, you must narrow (or diminish) the other.

Obstructions can change over time –seasonal changes to trees, new construction, and interference sources all can greatly influence wireless network performance.

'Line of Sight' includes the full width of the Fresnel Zone, and incidental obstructions that do not block visual LOS can impact wireless performance.

Overall throughput can be increased by combining/teaming multiple radios into a single unit, but are prone to the same obstructions or misalignment that can impact single radio units.

Channel Width has increased (higher frequencies) with recent iterations of wireless networking standards, but the overall number of available channels has dropped. Greater bandwidth allows for more traffic, but less channels mean that interference could be more likely.

In general, wireless networking is very different than wired networking, and design/troubleshooting methods/and knowledge is entirely different.

'Noise Floor' is the ambient noise that a wireless signal must overcome in order to be detected. The 'noise floor' changes dynamically depending on environment.

Antennas contribute to Gain, but not Loss, and while cables/connectors contribute minor loss, the most significant loss driver is "free space path loss", or the gross reduction in power as a signal travels through air.

If greater bandwidth is desired, higher signal level and wider frequency ranges are required.

Latency is the biggest problem when using PTZs over wireless systems. The lag between movement commands and camera position can significantly impact utility. The bandwidth overhead for PTZ controls on wireless networks is negligible.

Edge Storage is a good companion for 4G/wireless connected cameras as it permits high quality recordings at the device and only transmitted as needed over wireless.

Cameras with built-in cellular radios are uncommon, and mini-cell routers are a good choice. While 'full-sized' routers can be used, they often include unneeded features and consume power/space not ideal in surveillance deployments.

The higher the dBi of an antenna, the narrower its beam is, and the more critical antenna alignment becomes.

Some unexpected, or hard to predict sources of wireless performance degradation include growth of trees, small shifts of antennas due to wind / weather, new buildings, new wireless systems in the area.

Video Analytics

The term 'video analytics' should not be used for motion detection capabilities that have been available for years and primarily used for triggering recording. Motion detection measures simple changes in pixel value and provides a rough indicator of motion. By contrast, video analytics, by definition, deliver accurate tracking, alerting or counting of meaningful objects, such as people and cars.

In 2013, only a limited number of products claiming to be 'video analytics' perform well and to the level expected by users to accurately track, alert or count meaningful objects, such as people and cars. This is the case both for cameras and recorders.

Many vendors offer 'video analytics' for free. However, they tend to work modestly well, at best, in indoor controlled environment. Delivering high accuracy outdoors or in large areas, with minimal false alerts for free is practically non-existent.

Performing video analytics inside the camera is likely to deliver the best performance as the manufacturer can optimize the analytics for the specific camera as well as access the raw video feed. This is not guaranteed but is a commonly seen trend in professional products.

The main upside of server based video analytics is the lack of resource constraints. One can choose as powerful a box as they want or need and use as many or few cameras as appropriate. For very complex analytics, this may be the only way to implement them. However, this increases cost and complexity and risks accuracy issues as it cannot optimize for specific camera feeds/settings.

Camera positioning for video analytics is significantly more demanding than regular non analytic cameras as analytics are far more sensitive to issues of sunlight, weather, and camera angle.

A VMS with built in analytics eliminates the need for VMS integration, a common limitation when using 3rd party video analytics. However, VMSes offering built in analytics tend to either be free with poor performance or paid and still fairly expensive.

While adding 3rd party video analytics to IP cameras would make rolling out analytics easier, the options today are quite limited and, unfortunately not growing rapidly. Only a few manufacturers even support this approach and even the largest one only supports a handful of niche analytic offerings.

False positive rate is the most significant and painful problem in using analytics. False positive are alerts that are not of an actual suspect/threat. The key metric is how many false positives in a given time period (e.g., 10 per day, 100 per day, etc.). Too many false positives and systems are routinely shut down.

False negative rate can be problematic for critical infrastructure applications. False negatives are when an actual suspect/threat is present but the analytics does not generate an alarm.

Accuracy rates can be dangerously misleading and are often meant to trick or confuse purchasers. Generally, it is based on a staged test of the percentage of times a test subject is alerted on when trying out a system. This, of course, ignores the crucial false positive rate issue. Also, it may be based on ideal weather or lighting conditions that ignores the operational realities of 24/7 monitoring.

Even advanced analytics use very rough rules for detecting people and cars, typically looking at the shape of the overall object and its movement. While systems could identify more nuanced features like hair, faces, body parts and clothing, this would require far more pixels and far more processing. The net result might be greater accuracy *for detecting people* but over significantly smaller areas and at notably higher costs.