IPVM

TABLE OF CONTENTS

Fundamentals

Overview

We concisely explain the developments for surveillance cameras offered in 2017 and the state of offerings going into 2018, including cybersecurity, multi-imagers, resolution, H.265, HD analog, image quality, video analytics and cost.

Compare this with our Video Surveillance Cameras 2017 Overview and the Video Surveillance Cameras 2016 Overview.

Cybersecurity

For the first year, really ever in video surveillance, cybersecurity became a major factor. The two events that most impacted the industry were Dahua devices getting massively hacked and Hikvision's simplisitic IP camera backdoor being revealed. With the two largest volume manufacturers hit so heavily, the industry took notice. They were certainly not alone, as various vulnerabilities were found, at various levels of criticality across many manufacturers (as cataloged in our Directory of Video Surveillance Cybersecurity Vulnerabilities and Exploits).

Likewise, many manufacturers started marketing their commitment to cybersecurity (both as an offensive, i.e., way to beat competitors, and a defense mechanism, i.e., to stop customers from defecting).

Resolution Slowing

Camera resolution was not a major factor / shift in video surveillance cameras offered in 2017. Incrementally, adoption of 4MP continued to grow, with 4MP increasingly becoming the most typical resolution used and 4K use growing. However, compared to the past few years, the rate of resolution 'increase' has slowed, with 10+ MP camera options not growing significantly.

In particular, single-imager 'super' resolution growth was limited. Avigilon's proprietary offering stayed at 30MP tops. Axis and Sony both entered in at 20MP, adding options for open platforms, with strong performance but high costs compared to conventional cameras (see Axis vs Sony 20MP Shootout (Q1659 vs SNC-VB770).

Multi-Imagers Growing

For new cameras, multi-imagers is likely the most active segment, both fixed multi-imagers (i.e., 180° or 360° units) and repositionable camera options grew notably. Now, it is increasingly common for manufacturers to have at least fixed multi-imager cameras. Moreover, repositionable multi-imagers are a growing niche (e.g., see Axis P3707-PVE Multi-Imager Tested, Hanwha 20MP Multi-Imager Tested (PNM-9081VQ).

While we expect multi-imagers to remain a niche, they are becoming an increasingly important one for larger installations that see the benefits of reducing camera counts, cabling, installation costs and VMS licenses.

H.265 Going Mainstream

While H.265 had been offered for surveillance cameras for a few years, 2017 was the year when it started to become mainstream. This was due to a combination of more manufacturers offering H.265 plus combining it with smart codec support. H.265 VMS support still lags H.264 but it improved in 2017. Finally, there are definitely still patent licensing issues (see Manufacturers Shipping Unlicensed H.265 Products) but many will ignore that issue / risk.

Smart Codecs Now Expected

Just a few years ago, smart codecs were novel and limited. By 2017, smart codecs have become commonplace across major manufacturers and user adoption has surged. Now, smart codecs are expected with both H.264 and H.265 and the bandwidth / storge savings have become a significant factor in deployments.

Video Analytics - Deep Learning Hot But Early

Outside of cybersecurity, no area in video surveillance was more discussed than deep learning for video analytics. Unfortunately, in 2017 very few cameras shipped with deep learning, and certainly few mainstream offerings. Because of that, the impact of deep learning within video surveillance remains muted. The situation may very well change in 2018 depending on what is released by which manufacturers.

HD Analog New Products Cooled

HD analog did not make as much progress as it did in the past few years. HD analog has been promoting higher resolution, power over coax and interoperability for the past 2 years but those features were slow to roll out globally in 2017. It is not yet clear whether this is just temporary delays or it reflects a fundamental deceleration in HD analog progress.

One thing is indisputable, Japanese and Western manufacturers (outside of OEMs) continue to refuse offering HD analog, marketing against it, which constraints adoption.

Cost Declines Diminished

Unlike the past few years, camera cost declines decelerated. While $100 MP cameras are now commonplace, that rough price level has held among mainstream distribution. We attribute this to manufacturers increasingly factoring in the cost of sales and marketing plus growing support and cybersecurity costs. Related, 2017 is the year that the race to the bottom ended.

For 2018, the big talk is deep learning and AI but we remain cautious about how quickly and how widely that will be made available.

Pixels Determine Potential, Not Quality

Pixels = resolution = quality is not always true, since other factors impact quality (low light, bright light, lens quality, etc.).

However, pixels determine potential quality. In this note, we explain why you will make much clearer and better decisions recognizing this.

Demonstrating This With Images

Here's a relatively high quality, high 'resolution' image:

It's 161 pixels wide across a small area of 1 - 2 feet, delivering a high ~100 PPF.

Now contrast this with this low quality, low resolution image:

This is 5 pixels wide and a total of 35 pixels covering the same exact area as the image above (enlarged so you can see it). This is clearly low quality.

Why? The pixels are being forced to cover areas wider than the details desired. You can see it the blockiness of the image. Those blocks are the limits of the pixel.

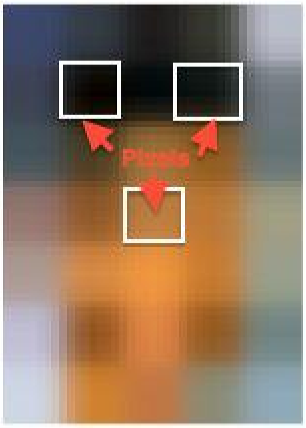

The image calls out 3 of the 35 pixels but you can make out pretty much all the individual pixels.

Now let's increase the pixels / resolution for this image.

A lot more details are being revealed now, as the number of pixels increases from 35 total to 140 and each pixel now can cover a smaller area.

Let's compare the two images to see key details in the image improve:

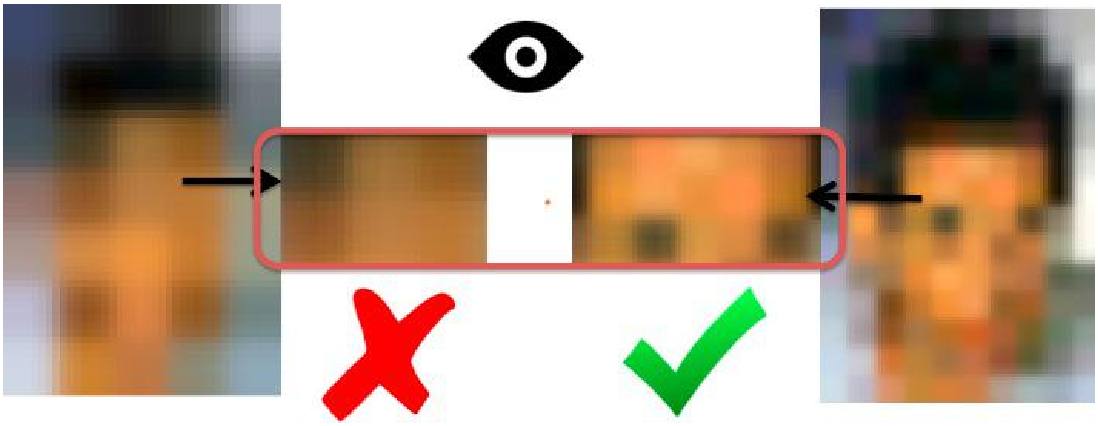

In the former image, the eyes were bigger than the pixels and therefore could not be captured. However, in the latter one, with pixels cover an area smaller than an individual eye, allows the eyes to be captured as two black dots.

Let's increase the pixels for this scene, from 140 to 240:

As each pixel covers a smaller area, more features continue to emerge - lips, ears, etc. and the eyes become more detailed (eyebrows, iris, etc.).

Finally, here's 4 samples ranging from 35 to 2120 pixels covering the same area:

35 140 540 2120

Clearly, as we increased pixels allocated, the more fine details that can be captured.

The smaller pixel count images, regardless of how 'good' the camera or encoder was could not capture those details because the pixels were covering too large an area for them. This is what we mean by pixels determine potential.

Pixels Limits on Quality

A 1MP camera will never capture the fine details of the face of a subject at a 50' wide FoV. It simply lacks the potential, because the pixels will cover too large an area relative (25ppf) to how small a face is at the same position.

However, a 5MP came covering that same 50' wide FoV may capture the fine details. It has the potential, because the pixels will be covering small enough areas (50ppf).

This potential, though, is a maximum theoretical limit bound by very important factors like:

- Ability to capture in low light scenes (which most 5MP+ cameras are terrible at).
- Ability to handle wide dynamic range scenes (see example).
- Quality of lens, preciseness of focus and eliminating any DoF problems.
- Minimizing compression artifacts / loss of quality (see tutorial).
- Angle of incidence of subject to camera (if the camera is too high or the person is looking askew from the camera, more pixels will not help).

Quality vs Pixels

Ultimately, image quality is driven by a half dozen factors combined. While pixel density / count determines the potential quality and the maximum achievable details, those other factors, that are often overlooked and ignored in PPF calculations, routinely and often dramatically constrain the actual image quality achieved.

Resolution

Understanding video surveillance resolution can be surprisingly difficult and complex. While the word 'resolution' seems self-explanatory, its use in surveillance is far from it. We will explain 5 critical elements:

- What resolution traditionally means – seeing details - and the constraints of this approach
- What resolution usually means in surveillance – pixels – and the limits of using this metric
- How sensor and stream resolutions may vary

- How compression impacts resolution greatly
- What limits resolution's value

Resolution – Seeing Details

In normal English and general usage, resolution means the ability to resolve details – to see or make them out. For example, can you read the lowest line on an eye chart? Can the camera clearly display multiple lines side by side on a monitor? etc. It is a performance metric focusing on results.

For example, in IPVM testing, the human eye is roughly equal to a 10MP resolution camera.

Historically, video surveillance used a similar test chart approach. Analog camera resolution was measured with line counts, literally the camera's ability to display more lines side by side in a given area on a monitor.

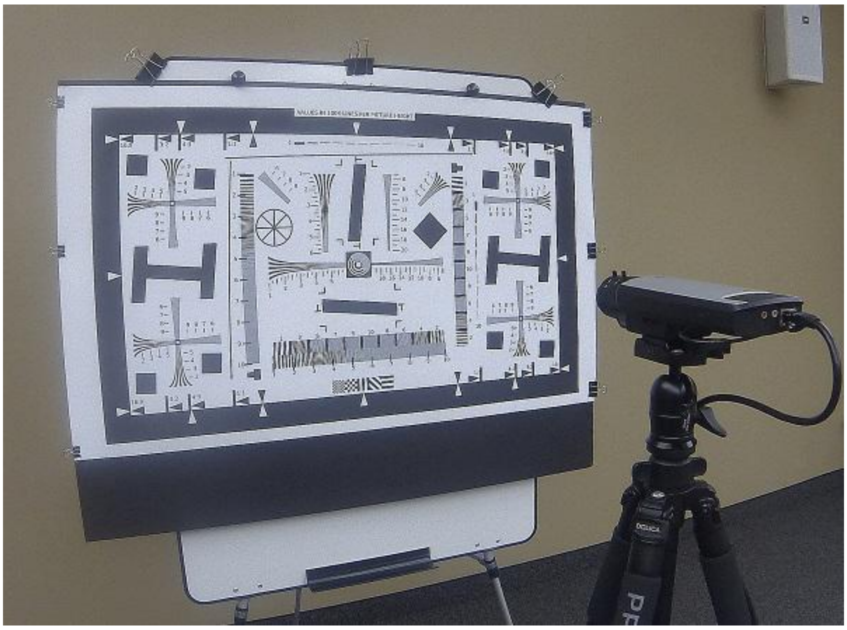

If you could see more lines, it meant you could see more real world details – facial features, characters, license plates, etc.

You can do that. For example, we did just such a chart based test, and found 720p cameras to be roughly equal to 500 'lines', 1080p roughly 900 'lines', etc. as shown in the image below:

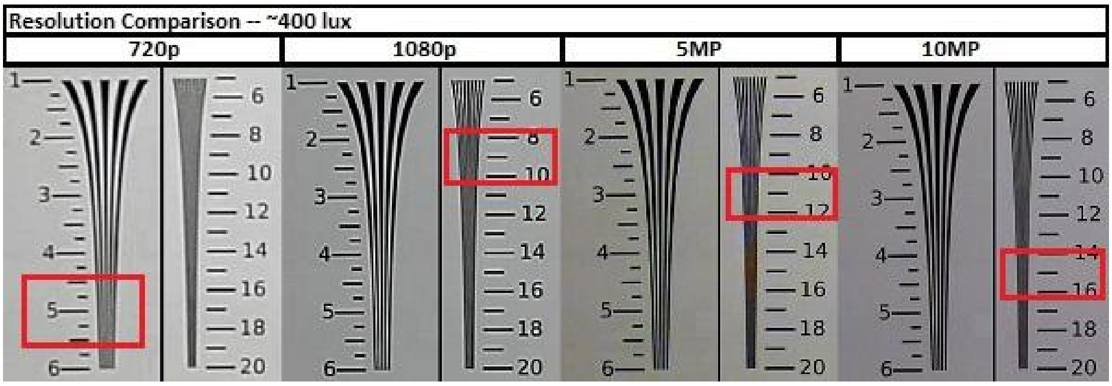

However, manufacturers almost always use pixel count instead of line count for resolution.

Moreover, lines counted was always done in perfectly even lighting conditions. However, with direct sunlight or low light, the resolving power would change, falling significantly. Even more challenging, some cameras fared far worse in these challenging lighting conditions than others.

While this approaches measures performance, it only does so in the most ideal, and often unrepresentative, conditions.

Resolution – Pixel Count

Now, with IP, manufacturers do not even attempt to measure performance. Instead, resolution has been redefined as counting the number of physical pixels that an image sensor has.

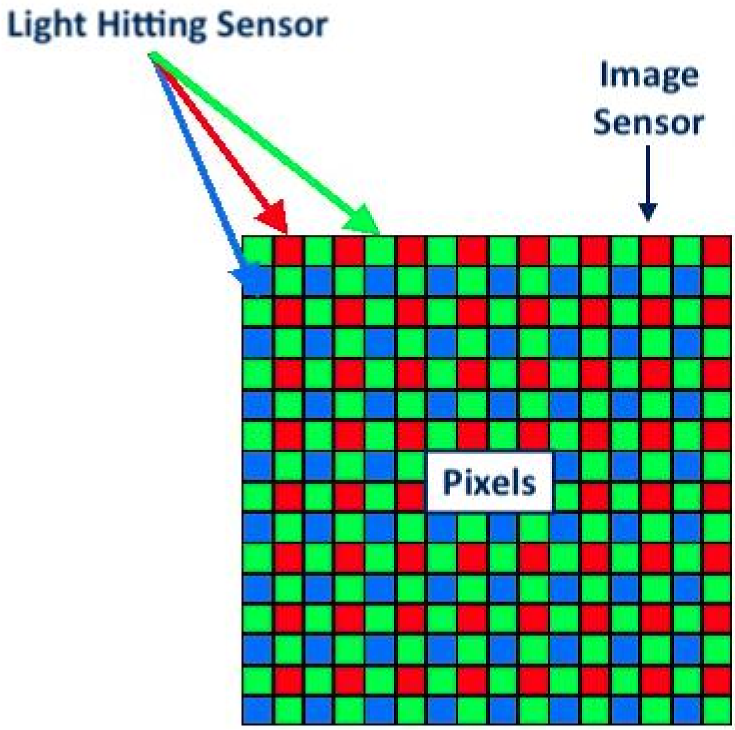

For example, a 1080p resolution camera is commonly described as having 2MP (million pixel) resolution because the sensor used has ~2 million pixels on it (technically usually 2,073,600 pixels as that is the product of 1920 horizontal x 1080 vertical pixels). The image of an imager below shows this example:

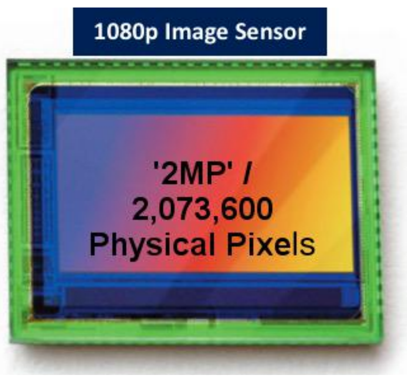

Pixels Determine Potential, Not Quality

Pixels are a necessary, but not sufficient, factor for capturing details. Without a minimum number of pixels for a given area / target, it is impossible. See our tutorial on why Pixels Determine Potential, Not Quality.

Limitations

The presumption is that more pixels, much like higher line counts, delivers higher 'quality'. However, this is far from certain.

Just like with classic resolution measurements that used only ideal lighting conditions, measuring pixels alone ignores the impact of common real world surveillance lighting challenges. Often, but not always, having many more pixels can result in poorer resolving power in low light conditions. Plus, cameras with lower pixel counts but superior image processing can deliver higher quality images in bright sunlight / WDR scenes.

Nonetheless, pixels have become a cornerstone of specifying IP video surveillance. Despite its limitations, you should:

- Recognize that when a surveillance professional is talking about resolution, they are almost certainly referring to pixel count, not resolving power
- Understand the different resolution options available

Common Surveillance Resolutions

The table below summarizes the most common resolutions used in production video surveillance deployments today. Note that VGA is no longer common except in thermal cameras, but is included here for reference of what 'standard definition' refers to.

NAME / PIXEL COUNT	HORIZONTAL X VERTICAL
VGA / .3 MP	640 x 480
720p / 1 MP	1280 x 720
1080p / 2 MP	1920 x 1080
3 MP 4:3	2048 x 1536
3 MP 16:9	2304 x 1296
4 MP	2688 x 1520
5 MP 4:3	2592 x 1944
5 MP 16:9	3072 x 1728
6 MP	3072 x 2048
4K / 8 MP	3840 x 2160
12 MP	4000 x 3000

Changes For 2018

While 1080p, 4MP, 4K, and other resolutions remain in common use in 2018, there are some notable changes in camera resolution in the past year.

- 3MP/5MP confusion: Historically, users have known 3 and 5MP resolutions as 4:3 aspect ratio (2048x1536 and 2560x1944, respectively). But now, cameras using 16:9 variants of these resolutions are available, delivering increased horizontal PPF, but reducing height of the coverage area, which may eliminate areas visible when using 4:3 cameras.
- 10MP uncommon: Though it used to be one of the most common "high" resolutions, 10MP has practically fallen out of use in 2017/2018.

- 6MP available: Finally, 6MP cameras are now readily available, due to new generations of sensors using this resolution. 6MP uses an odd (for surveillance) 3:2 aspect ratio.

720p cameras, once most popular by a wide margin, have sharply declined as higher resolution options have come down in price and several manufacturers offer fewer new models in this resolution compared to higher.

Resolution Vs. Cost

Everything else equal, higher resolution cameras generally cost more than lower cost models, though pricing for some 4K cameras have started to decline in 2017. Higher prices are due in part to simple increases in component costs adding up, such as more expensive image sensors, additional processors required, higher resolving power lenses, etc.

However, note that this higher cost does not always result in higher performance, as advanced features such as super low light and true WDR are not always supported or as high performing in higher resolution models, or requiring a significant increase in cost. For example, 1080p cameras most commonly offer strong WDR and super low light options, with such features becoming less common in higher resolution cameras.

Sensor Resolution vs. Stream Resolution

While manufacturers typically specify cameras based on the resolution (i.e. pixel count) of the sensor, sometimes, the resolution of the stream sent can be less. This happens in multiple cases:

- Limited camera capabilities: In some cases, manufacturers may use readily available sensors of one resolution but crop the sensor to a lower pixel count due to limitations in processing at full resolution. For example, a 6MP sensor may be cropped to 5MP in order to stream at higher frame rates or apply WDR or higher gain levels.
- Panoramic cameras: Second, manufacturers often crop unused portions of the sensor from panoramic camera streams, so a "12MP" fisheye model may actually stream at 8-9MP. See our report Beware Imager vs Stream Resolution for more information on this issue.
- Reduced in software: Finally, an installer may explicitly or mistakenly set a camera to a lower resolution. Some times this is done to save bandwidth but other times it is simply an error or glitch in the VMS default resolution configuration. Either way, many times an HD resolution may look ‘terrible’ but the issue is simply that it is not set to its max stream resolution (i.e., a 3MP camera set to 640 x 480).

Because of these issues, users should be sure to check not only the resolution of the sensor but the stream resolutions supported and used, typically found lower down the camera's datasheet:

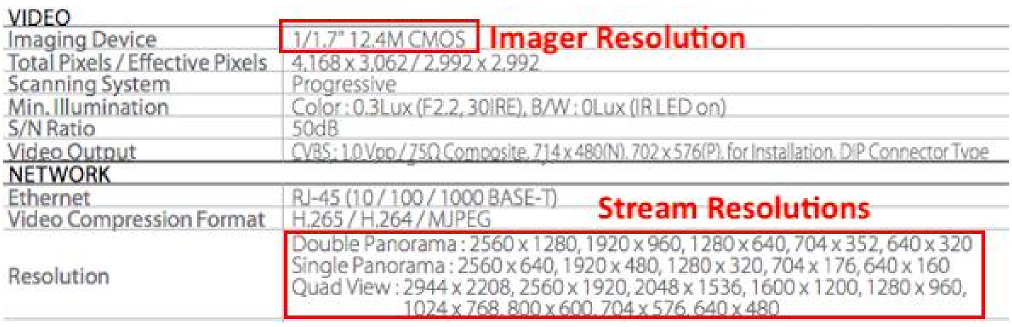

VIDEO	
Imaging Device	1/1.7" 12.4M CMOS
Total Pixels / Effective Pixels	4,168 x 3,062 / 2,992 x 2,992
Scanning System	Progressive
Min. Illumination	Color : 0.3Lux (F2.2, 30IRE), B/W : 0Lux (IR LED on)
S/N Ratio	50dB
Video Output	CVBS : 1.0 Vpp / 75Ω Composite, 714 x 480(N), 702 x 576(P), for Installation, DIP Connector Type
NETWORK	
Ethernet	RJ-45 (10 / 100 / 1000 BASE-T)
Video Compression Format	H.265 / H.264 / MJPEG
Resolution	Double Panorama : 2560 x 1280, 1920 x 960, 1280 x 640, 704 x 352, 640 x 320 Single Panorama : 2560 x 640, 1920 x 480, 1280 x 320, 704 x 176, 640 x 160 Quad View : 2944 x 2208, 2560 x 1920, 2048 x 1536, 1600 x 1200, 1280 x 960, 1024 x 768, 800 x 600, 704 x 576, 640 x 480

Compression Impact On Resolution

Because resolution most often simply means pixel count, no consideration is given to how much pixels are compressed. Each pixel is assigned a value to represent its color, typically out of a range of ~16 million (24 bits), creating a huge amount of data. For instance, a 1080p/30fps uncompressed stream is over 1Gb/s. However, surveillance video is compressed, with that 1080p/30fps stream more typically recorded at 1Mb/s to 8Mb/s — 1/100th to 1/1000th less than the uncompressed stream. The only question — and it is a big one — how much is the video compressed?

Picking the right compression level can be tricky. How much compression loss can be tolerated often depends on subjective preferences of viewers or the details of the scene being captured. Equally important, increasing compression can result in great savings on hard drive costs (less storage required for similar durations), server configuration (less CPU required is required to store less bandwidth), and switches (copper gigabit switches may be used instead of fiber 10GbE).

Just because two cameras have the same resolution (i.e. pixel counts), the visible image quality could vary substantially because of differences in compression levels chosen. Here is an example:

For full coverage of these details, see our video quality / compression tutorial.

Also important for considering compression is that manufacturers default compression settings vary significantly, for more see: [IP Camera Manufacturer Compression Comparison].

Angle Of Incidence Is Key

Regardless of how high quality an image is, it needs to be at a proper an gle to 'see' details of a subject, as cameras cannot see through walls nor people. For instance:

Even if the image on the left had 10x the pixels as the one on the right, the left one is incapable of seeing the full facial details of the subject as he is simply not facing the camera.

This is frequently a practical problem in trying to cover a full parking lot with a single super high-resolution camera. Even if you can get the 'right' number of pixels, if a car is driving opposite or perpendicular to the camera, you may not have any chance of getting its license plate (similarly for a person's face).

Resolution Overkill

Historically, surveillance has been starved for resolution, with almost always too little for its needs. Anyone familiar with suspect photos on their local news can see this:

However, as the amount of pixels has increased to 1080p and beyond, the opposite issue presents itself: unnecessarily high resolution for the scene. Once you have enough to capture facial and license plates details, most users get little practical benefit from more pixels. The image might look 'nicer' but the evidentiary quality remains the same. This is a major consideration when looking at PPF calculations and ensuring that you do not 'waste' pixels.

Additional Factors Impacting Resolution

Finally, note that beyond issues discussed above, many other factors impact surveillance resolution beyond pixels, including:

- Low light performance
- WDR performance
- Compression settings
- Camera angle / downtilt
- Lens selection
- Lens focus

Do not accept specified resolution (i.e. pixel count) as the one and only quality metric as it will result in great problems. Understand and factor in all of these drivers to obtain the highest quality for your applications.

Test your knowledge

Take this 9 question quiz now.

Frame Rate

As a precursor, you need to know the speed of objects, most typically people.

Speed of People

The faster a person moves, the more likely you are to miss an action. You know the 'speed' of frame rate - 1 frame per second, 10 frames per second, 30, etc., but how many frames do you need for reliable capture?

Here's how fast people move.

For a person walking, a leisurely, ordinary pace is ~4 feet per second, covering this 20 foot wide FoV in ~5 seconds:

Note: ***Click here to watch the demo on IPVM***

For a person running, our subject goes through the 20' FOV in ~1.25 seconds, meaning he covers ~16' in one second:

Note: ***Click here to watch the demo on IPVM***

For example, if you only have 1 frame per second, a person can easily move 4 to 16 feet in that time frame. We need to keep this in mind when evaluating frame rate selection.

We cover:

- What speed do people move at and how does that compare to frame rates.
- Walking: What risks do you have capturing a person walking at 1, 10 and 30fps.

- Running: What do you have capturing a person running at 1, 10 and 30fps.
- Head Turning: How many more clear head shots do you get of a person at 1, 10 and 30fps.
- Playing Cards: What do you miss capturing card dealing at 1, 10 and 30fps.
- Shutter speed vs Frame Rate: How are these two related?
- Bandwidth vs Frame Rate: How much does bandwidth rise with increases in frame rate?
- Average Frame Rates used: What is the industry average?

Walking Examples

As our subject walks through the FOV, we view how far he moves from one frame to the next. In 30 and 10 fps streams, he does not complete a full stride. However, in the 1fps example, he has progressed ~4' between frames, which falls in line with our measured walking speed of ~4' a second.

Note: ***Click here to watch the comparisons on IPVM***

Running Examples

With our subject sprinting through the FOV, the 30 fps stream still catches him mid stride, while in the 10 fps stream, he has traveled ~1' between frames. In the 1 fps example, only one frame of the subject is captured, with him clearing the rest of the FOV between frames, with only his back foot visible in the second frame.

Note: ***Click here to watch the comparisons on IPVM***

Capturing Faces

Trying to get a clear face shot can be difficult when people move because they naturally shift their head frequently. In this demonstration, we had the subject shake their head back and forth walking down a hallway to show the difference frame rate plays.

Take a look:

Note: Click here to watch the demo on IPVM

Notice, at 1fps, only a single clear head shot is captured, but at 10fps, you get many more. Finally, at 30fps, you may get one or two more, but it is not much of an improvement.

Playing Cards

In this test, our subject dealt a series of playing cards from ace to five with the camera set to default shutter speed (1/30).

In the 30 and 10 fps examples, we can see each card as it is removed from the top of the deck and placed on the table. However, in the 1 fps example, we see only the cards appearing on the table, not the motions of the dealer, as frame rate is too low.

Note: Click here to view the comparison samples on IPVM

Shutter Speed vs Frame Rate

Frame rate does not cause blurring. This is a misconception. The camera's automatic shutter speed control does.

Dealing cards ace through 5 again, we raised the camera's minimum shutter speed to 1/4000 of a second. The image below compares the motion blur in the dealers hand and card, with the 2 card much more easily legible in the fast shutter speed example.

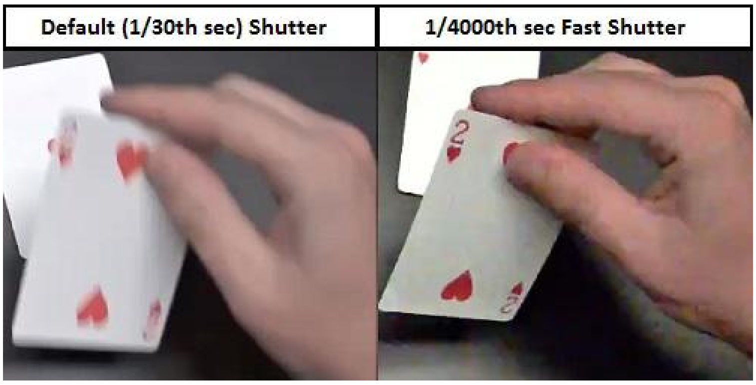

1/4000s shutter speed completely eliminated all traces of motion blur. 1/1000 and 1/2000 of a second shutter speeds significantly reduces blur, but it was still noticeable around the dealers fingers and edges of the cards when looking at the recordings frame-by-frame.

If you have blurring, you have shutter speed configuration problem, not a frame rate one.

Slow Shutter and Frame Rate

On the other side, sometimes users want or camera manufacturers default their maximum shutter to a rate slower than the frame rate (e.g., a 1/4s shutter for a 1/30s camera). Not only does this cause blurring of moving objects, you lose frames.

Key lesson: The frame rate per second can never be higher than the number of exposures per second. If you have a 1/4s shutter, the shutter / exposure only opens and closes 4 times per second (i.e., 1/4s + 1/4s + 1/4s

+ 1/4s = 1s). Since this only happens 4 times, you can only have 4 frames in that second.

Some manufacturers fake frames with slow shutter, simply copying the same frame over and over again. For example, if you have 1/15s shutter, you can only have 15 exposures and, therefore, 15 frames. To make it seem like you have 30 frames, each frame can be sent twice in a row.

Be careful with slow shutter. Beyond blur, you can either lose frames or waste storage.

Bandwidth vs Frame Rate

Frame rate impacts bandwidth, but for modern codecs, like H.264, it is less than linear. So if you increase frame rate by 10x, the increase in bandwidth is likely to be far less, often only 3 to 5 times more bandwidth. This is something we see mistaken regularly in the industry.

The reason for this is inter-frame compression, that reduces bandwidth needs for parts of scenes that remain the same across frames (for more on inter and intra frame compression, see our CODEC tutorial).

Illustrating this point further, we took 30, 10 and 1 fps measurements to demonstrate the change in bit rate in a controlled setting in our conference room. The average bitrates were as follows:

- 1 fps was 0.179 Mb/s
- 10 fps, with 10x more frames, consumed 4x more bandwidth than 1 fps (0.693 Mb/s)
- 30 fps, with 3x more frames, consumed double the bandwidth of 10fps and, with 30x the frames, 7x the bandwidth of 1fps (1.299 Mb/s)

These measurements were done with 1 I frame per second, the most common setting in professional video surveillance (for more on this, see: Test: H.264 I vs P Frame Impact).

For more on this, see our reports testing bandwidth vs frame rate and 30 vs 60 fps.

Average Frame Rates Used

Average industry frame rate is ~10fps, reflecting that this level provides enough frames to capture most actions granularly while minimizing storage costs.

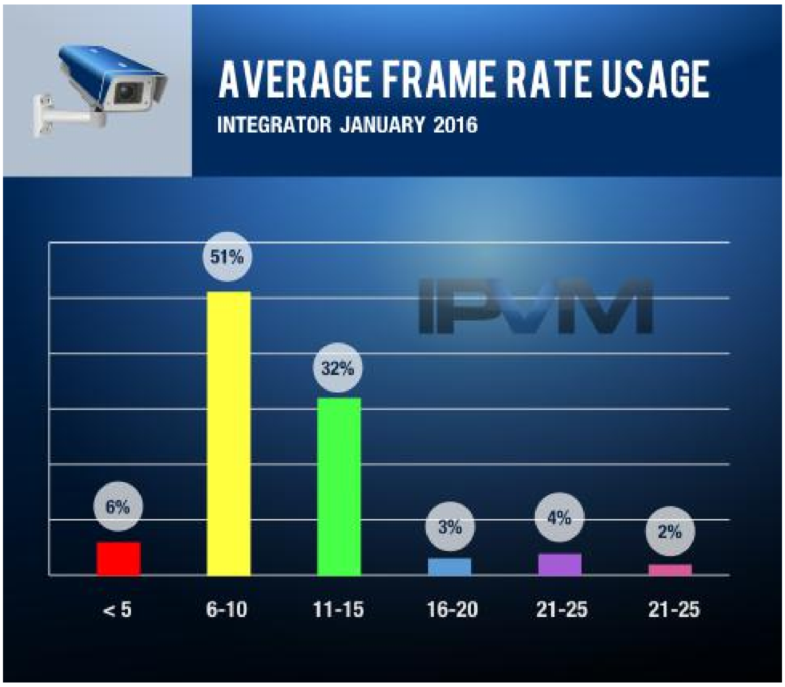

As shown in the previous section, going from 10fps to 30fps can double storage costs but only marginally improve details captured.

For more commentary on why integrators choose the frame rates hey do, see the Average Frame Rate Used Statistics report.

PPF / PPM

Pixels per foot / Pixels per meter is the most fundamental and valuable, though imperfect, metric for specifying video surveillance image quality.

In a single number, this metric (e.g., 10ppf, 40ppf, 100ppf) conveys important information about what the projected quality that a camera can provide.

The image below, taken from our Camera Calculator, demonstrates examples of common pixel per foot (ppf) levels:

Note: Click here to watch the comparison images on IPVM

PPF Established Metric

PPF has become a critical established metric for several reasons:

- Broad camera manufacturer support: Most major manufacturers use this metric.
- Common A&E specifications: Architects and engineers who plan large projects regularly use PPF / PPM as the basis for their designs and surveillance plans.
- Need for Something: With so many resolution options today (from 1MP to 12MP and beyond), the old metrics which used percentage of screen covered make no sense. PPF has filled this void.

The Goal of PPF

PPF is a single metric (e.g., 10, 50, 90, etc.) that when specified should deliver a specific level of quality. For example, the parking lot camera must deliver 50 PPF. Instead of guessing or just specifying more resolution, using this metric should enable selection of the 'right' resolution for the scene. The final image, following the PPF metric, will then deliver a predictable level of quality.

PPF Problems

Alas, PPF suffers from many problems that must be factored in:

- Assumes even lighting and ignores the impact of bright sunlight
- Assumes day time lighting and ignores the impact of night time / low light viewing
- Disregards differences in lenses and compression
- Disregards that image quality needs are subjective and debatable
- Fails to specify related and critical metrics to complement PPF

Despite this, PPF does have value for estimation and planning. It just cannot be used blindly or simplistically. We explain:

- How to calculate PPF
- How to recognize PPF limitations and make adjustments
- How to best use PPF productively

Key Recommendations

Keep in mind:

- PPF is far more art than science. You are fooling yourself if you think a single number can get close to specifying something as complex as image quality across conditions and scenes.
- Understand the key limitations of PPFs so you can avoid mistakes.

- Use PPF as a baseline to get a sense of potential performance then deduct for specific site problems.
- You can almost never guarantee quality levels 24/7/365 in video surveillance.
- Communicate limitations and set proper expectations.
- Do PPF calculations to improve estimates but avoid using them in specifications. Instead, specify resolution and key camera properties (WDR, low light performance, etc.)

PPF vs PPM

Note, we use PPF throughout this report, but sometimes the metric equivalent, Pixels Per Meter (PPM), is used. However, PPF is commonly used, even in many regions standardized on the metric system.

PPF	APPROXIMATE PPM
100	315
80	250
60	200
40	132
20	65
10	33

For reference, the chart below shows some standard PPF values for conversion. Note that the IPVM Camera Calculator also supports metric measurements, in addition to imperial / PPF.

How to Calculate Pixel Density

Calculating pixel density depends on two inputs/factors:

- The number of horizontal pixels in a video feed/camera
- The width of the FoV of the camera at the point of interest

Horizontal Pixel Count

Be careful to properly identify the horizontal pixel count and do not confuse it with the vertical one nor overall resolution. For instance, a 720p camera has a total resolution of 1MP (technically 0.926) but its horizontal pixel count is 1280 and its vertical pixel count is 720, more commonly described as 1280 x 720. It is the first number in this pair, the 1280, that counts for measuring PPF.

Here are horizontal pixel counts for common resolutions:

RESOLUTION	HORIZONTAL PIXEL COUNT
1080p	1920
3MP 4:3	2048
4MP	2688
5MP 4:3	2592
4K	3860 (or 4096)

Make sure to memorize them as they are quite fundamental.

Selecting FOV Width

The second part of the calculation is how wide a Field of View one wants to cover. The image below shows a simple scenario for determining the FoV, simply measure across at the point of interest, like the gate below:

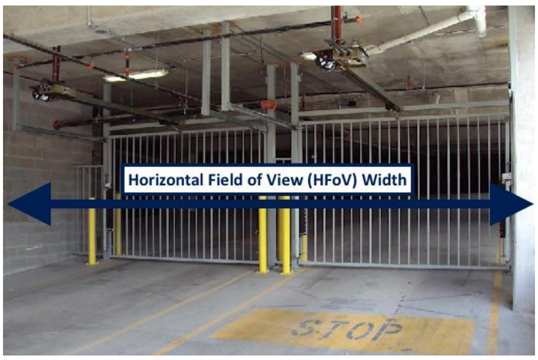

Determining the width of the FoV can be trickier for larger, open areas, such as parking lots or fields. For instance, the animation below shows two potential FoV widths for covering a parking lot/entry.

Note: ***Click here to watch the comparison images on IPVM***

Both of these FOVs may be 'right', but it depends on multiple factors.

- Do other cameras cover the area? If another camera is mounted to the right of this one, closer to the cars on the bottom right, it may be better to use the narrower 100' FOV for calculations. If there is no other coverage, the 200' may need to be used.
- What resolution limitations are there? If camera selection is limited by resolutions, we must keep this in mind when selecting HFoV. This may be key if super low light models are required, which are typically limited to 1080p. But if 4K or 12MP cameras may be used, wider FOVs are possible.

If there is any uncertainty, we recommend using wider FOVs, so that you do not underestimate your PPF.

PPF Formula

Once you have determined the horizontal pixels of the camera chosen and measured your FOV width, PPF is simple division.

$$\frac{\text{Horizontal Pixels}}{\text{FoV Width}} = \text{PPF}$$

Here are a few examples:

- You are using a 720p camera to cover a 10 foot wide entrance: PPF = 128 because horizontal pixel count is 1280 and FoV Width = 10 (i.e., 1280 / 10).
- You have a 5MP camera covering a 100 foot wide parking lot: PPF = ~26 because horizontal pixel count is 2592 and FoV width = 100 (i.e., 2592 / 100).
- A 1080p camera is monitoring a 40 foot wide vehicle entrance: PPF = 48 because horizontal pixel count is 1920 and FoV width = 40 (i.e., 1920 / 40).

PPF Decreases as FoV Width and Distance from Camera Increases

Cameras do not have a 'single' PPF. What the PPF is depends on how far the subject of interest is from the camera. For instance, using the IPVM Camera Calculator, we can mouse over a camera's FOV, which shows PPF at that point. We can see that PPF decreases as we approach the target to 17 PPF, but nearing the camera it rapidly increases, to 100, 200, or 300 PPF.

Note: ***Click here to demo on IPVM***

Equally important, the PPF level decreases as we widen the FOV, seen below. Dragging the bounds of our FOV wider, we can see PPF decrease, highlighted in red.

Note: ***Click here to demo on IPVM***

You can try this for yourself in the IPVM Camera Calculator.

Fisheye Field Of View Impact

Finally, beware of field of view width when using fisheye panoramic models, as PPF falls dramatically, much faster than standard cameras.

Using two (2) 5MP cameras as an example, one a fisheye panoramic (360 HFoV) and the other using a ~60° HFoV. Imagine you want to see an intruder 20 feet away from each camera:

- With a 5MP panoramic, at 20 feet away from the camera, the PPF is ~21
- With a 5MP 'conventional' 60° lens, at 20 feet away from the camera, the PPF is ~130

This is a massive practical difference. While the panoramic can see in 'all' directions, in any direction, at 20 feet away, people will look like blobs (best case scenario). By contrast, the 'conventional' lens will deliver near passport quality images though across a much narrower area.

See our guide Calculating Panoramic Camera Coverage for more details.

Recognizing PPF Limitations

Even after calculating PPF, there are several issues which impact the practical quality of images delivered. Below, we review:

- The impact of uneven lighting / wide dynamic scenes
- Low light / night time viewing
- Compression / image quality setting impact
- Lens quality differences
- Subjective image quality requirements
- Complementary metrics to PPF

Wide Dynamic Range Limitations

PPF assumes all pixels are the same but the resolving power can vary significantly. The ability to handle bright backlight / sunlight is a key example. In the image below, both cameras have the same resolution, cover the same FoV and are taken at the same time:

However, the image on the right looks like it has twice the PPF level as the other. Though the PPF is the same, the camera on the right has far stronger and more effective WDR.

This is an issue even in scenes with even moderate variations in light levels. In the image below, zone A has a lower PPF than zone B but zone A clearly captures more image detail because Zone B is in a shaded area that the camera struggles to render effectively.

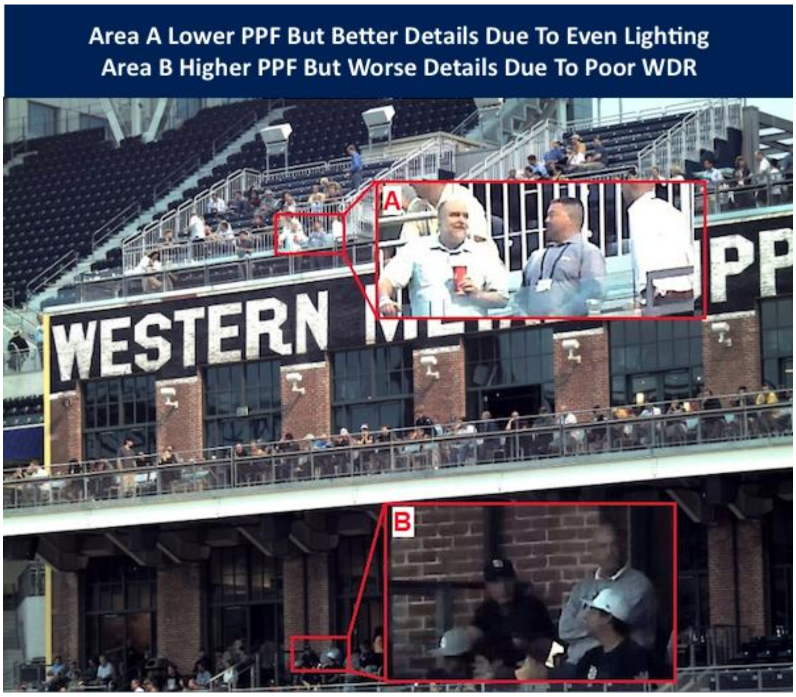

See our WDR Tutorial 2017 for more information on wide dynamic range and these effects.

Low Light Limitations

Even in only moderately low light, as high as 40lx, camera automatic gain control increases, with digital noise reducing image quality. This can be seen in the series of images below, showing the subject in the same spot over the course of several hours from day to night, with light levels starting at ~7,000lx on the left, and under 1lx on the right.

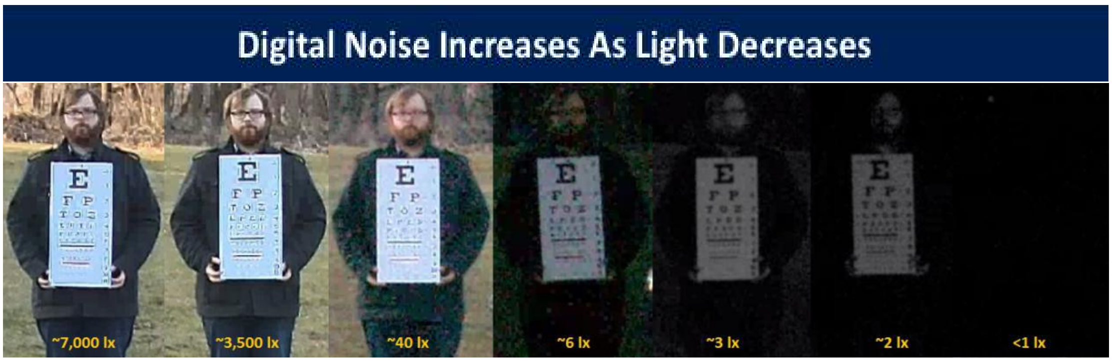

Despite the scene not being 'dark' by typical standards, even the best cameras capture fewer details and display more noise than during the day. However, manufacturer PPF guidelines almost never make this critical limitation clear.

In the Camera Calculator, we include samples taken in real world day and night scenes to show this difference, with PPF previewed in real time as users adjust their cameras' parameters.

Differences in Compression

Perhaps the trickiest element with PPF is compression level settings. Even if the same resolution and the same CODEC are used, the amount of compression can vary greatly. Below demonstrates that:

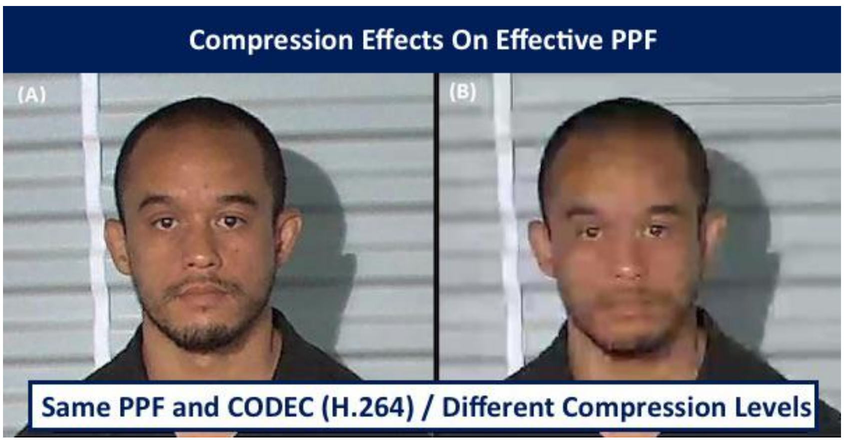

It looks like image B's PPF is far lower than image A, but that is not the case. has far less PPF than image A but that is not the case. Both images are taken from the same camera, with the only difference being increased compression on image B.

Adding to the potential confusion, manufacturers use very different default compression settings, as well as different terminology for their various quality levels. Readers should see our IP Camera Manufacturer Compression Comparison for clarity on these settings and manufacturer terms.

Users should be careful to check compression levels and measure compression if necessary to avoid these possible issues. See our Video Quality / Compression Tutorial for more details on compression and its effects.

Lens Differences

Many argue that lens choices can make a huge difference in image quality, regardless of the pixels provided. However, in our tests, lens quality is rarely a major concern, despite fine quality differences which are perceptible in some lenses. Other factors, such as F-stop, imager, and processing capabilities are far more likely to impact image quality.

Subjective Image Quality

Perhaps the most overlooked element in PPF is how subjective the 'right' quality level may be. Just as 'beauty is in the eye of the beholder', one reasonable person may think a given quality level is good enough to 'identify' a subject, but another may believe significantly more pixels on target are required.

For example: A customer wants to detect a human subject. Which of the three PPF levels below is right?

While none of them show facial details and all of them show at least an outline, some users prefer the 18 or 25 PPF images, as they show more than just a moving 'blob', with some very rough details.

The same situation commonly arises when deciding whether how many pixels are needed to identify a specific person. Compare the images below:

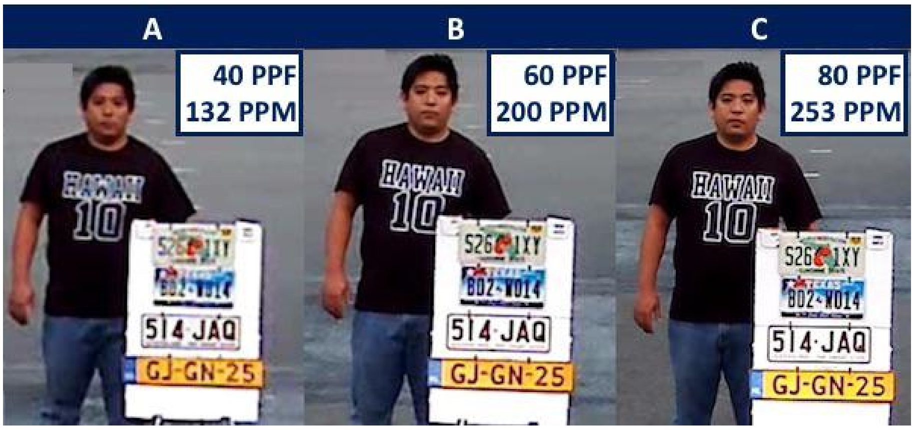

Some people will say the ~40 PPF on the left is too fuzzy and that they want a 'sharper' image. Others will find it good enough, especially if the person is known, such as an employee or customer.

Regardless of what an 'expert' claims is enough, one needs to verify with the buyer / decision maker what they think is sufficient.

Related / Critical Metrics

When specifying cameras, PPF alone does not provide enough information to properly select cameras, leading to confusion or guesswork. When using PPF, related metrics must be used:

- Field of View Width: Often specifications fail to define how the the FoV needs to be, providing only a PPF target. This can create confusion as some will use narrower or wider FOVs than others. Instead, specifically state how wide, e.g., 50 PPF at 40 feet, etc.
- Distance to Target: Knowing how far the camera needs to deliver the given PPF impacts lens selection as well as the width of the FoV. Remember, the farther away from the camera one gets, the lower the PPF.
- Angle of Incidence: Finally, regardless of the number of pixels, if the camera angled too harshly vertically or horizontally, one will miss key details like faces or license plate characters. In our Camera Height vs Image Quality test, we determined that 15° or less is ideal for proper facial images.

FOV Size Guidelines

Because of increases in resolution since megapixel cameras were introduced, PPF selection has become easier for smaller areas. For example, in areas 20' wide or less, a 1080p camera will provide ~100ppf, more than enough pixels for a high quality day time image. By contrast, if you were using an SD camera, to get the same PPF, you would have to limit your FoV width to 6 foot wide - a 70% decrease.

At PPF levels this high, the only remaining issues are performance decisions such as the camera's low light and WDR performance. Remember, image quality will still decline with bright sunlight and low light but the high number of pixels provides a 'margin for error'.

Larger FoVs More Difficult

However, above 20-40 foot wide FoVs, the situation becomes much trickier. Moving up from 1080p to a 5MP camera only increases the horizontal pixel count modestly, ~35%. Equally importantly, typically low light and WDR performance of 5MP cameras are worse than 1080p ones.

Be careful about specifying PPF for large areas as it can create problems of its own. For instance, let's say you want 40ppf across a 100 foot wide area. To do so, you will need 4000 horizontal pixels. A single 14MP camera technically fits (as it has ~4500 horizontal pixels) as would (2) 5MP cameras. The best solution might still be (2) 1080p cameras. While the total PPF is slightly less (3840), the cameras overall may be superior in WDR, low light, frame rate, broad product options, etc.

PPF Guidelines Offered by Manufacturers

With all this in mind, you should compare to manufacturer PPF guidelines. **Remember:** these metrics assume bright, even lighting and do not factor in the impact of bright sunlight or even modest darkness, which is rarely acknowledged.

Common guidelines:

- High quality / identifying people and license plates: Manufacturers suggest between 40 and 150 PPF. Even in ideal conditions, 40 is unlikely to deliver solid identifying details, especially of license plates. By contrast, 150 is likely overkill unless lighting is very poor or very fine details must be visible, such as the serial numbers of currency.

- Moderate quality / easily see that an object (person, car) is there but not all identifying details: Manufacturers typically recommend between 20 and 40 PPF. This is fairly realistic as long as lighting is ideal.
- Low quality / detection only: Manufacturers suggest 3 - 10 PPF. Be careful, because at this level, it can be very hard to make someone out unless you already know that they are there or if you are looking carefully. Of course, this also presumes ideal lighting. All in all, if you have less than 10 PPF, you are highly likely to miss objects. However, some analytics require only a few PPF to function, so may still be used.

[Note: This guide was originally written in 2012 but was revised in 2016/2017 to increase details covered and provide better examples from the IPVM Camera Calculator.]

Bandwidth

Bandwidth is one of the most fundamental, complex and overlooked aspects of video surveillance.

Many simply assume it is a linear function of resolution and frame rate. Not only is that wrong, it misses a number of other critical elements and failing to consider these issues could result in overloaded networks or shorter storage duration than expected.

We take a look at these factors, broken down into fundamental topics common between cameras, and practical performance/field issues which vary depending on camera performance, install location, and more.

Fundamental Issues

- Resolution: Does doubling pixels double bandwidth?
- Framerate: Is 30 FPS triple the bandwidth of 10 FPS?
- Compression: How do compression levels impact bandwidth?
- CODEC: How does CODEC choice impact bandwidth?
- Smart CODECs: How do these new technologies impact bandwidth?

Practical Performance/Field Issues

- Scene complexity: How much do objects in the FOV impact bitrate?
- Field of view: Do wider views mean more bandwidth?
- Low light: How do low lux levels impact bandwidth?
- WDR: Is bitrate higher with WDR on or off?
- Sharpness: How does this oft-forgotten setting impact bitrate?
- Color: How much does color impact bandwidth?

- Manufacturer model performance: Same manufacturer, same resolution, same FPS. Same bitrate?

Scene Complexity

The most basic commonly missed element is scene complexity. Contrast the 'simple' indoor room to the 'complex' parking lot:

Even if everything else is equal (same camera, same settings), the 'complex' parking lot routinely requires 300%+ more bandwidth than the 'simple' indoor room because there is more activity and more details. Additionally, scene complexity may change by time of day, season of the year, weather, and other factors, making it even more difficult to fairly assess.

We look at this issue in our [Advanced Camera Bandwidth Test](#).

Resolution

On average, a linear relationship exists between pixel count (1MP, 2MP, etc.) and bandwidth. So for example, if a 1MP camera uses 1 Mb/s of bandwidth, a 2MP camera on average might use ~2Mb/s.

However, variations across manufacturers and models are significant. In IPVM testing, some cameras increase at a far less than linear level (e.g., just 60% more bandwidth for 100% more pixels) while others rose at far greater than linear (e.g., over 200% more bandwidth for 100% more pixels). There were no obvious drivers / factors that distinguished why models differed in their rate of increase.

As a rule of thumb, a 1x ratio may be used when estimating bandwidth difference across resolution. However, we strongly recommend measurements of actual cameras as such a rule of thumb may be off by a lot.

Frame Rate

Frame rate impacts bandwidth, but for inter-frame CODECs such as H.264, it is less than linear. So if you increase frame rate by 10x, the increase in bandwidth is likely to be far less, often only 3 to 5 times more bandwidth. Illustrating this, we took 30, 10, and 1 fps measurements to demonstrate the change in bit rate in a controlled setting in our conference room. The average bitrates were as follows:

- 1 fps: 0.179 Mb/s
- 10 fps: 0.693 Mb/s (10x the frames of 1 fps, but only 4x bandwidth)
- 30 fps: 1.299 Mb/s (3x the frames of 10 fps, but only double bandwidth. 30x frames of 1 fps, but only 7x bandwidth)

(These measurements were done at 1 I frame per second with quantization standardized ~28.)

For more detail on frame rate's impact on bitrate, see our Frame Rate Guide for Video Surveillance.

Compression

Compression, also known as quantization, has an inverse relationship to bandwidth: the higher the compression, the lower bandwidth will be.

CRITICAL: Compression and resolution are two different things. In IPVM courses, we routinely see professionals mix the two. Resolution, in our industry, is the number of pixels in an image / video. Compression is how heavily compressed those pixels are.

For example, the chart below shows the impact of compression across four different cameras (note: with H.264, quantization / compression is measured on a standard scale of 0 to 51, higher meaning more compression, lower quality).

Lowering quantization from 34 (high compression) to 28 (average) resulted in at least a 3x increase in bandwidth, while further lowering it to 22 (very low compression) resulted increases of 5-11x depending on the camera.

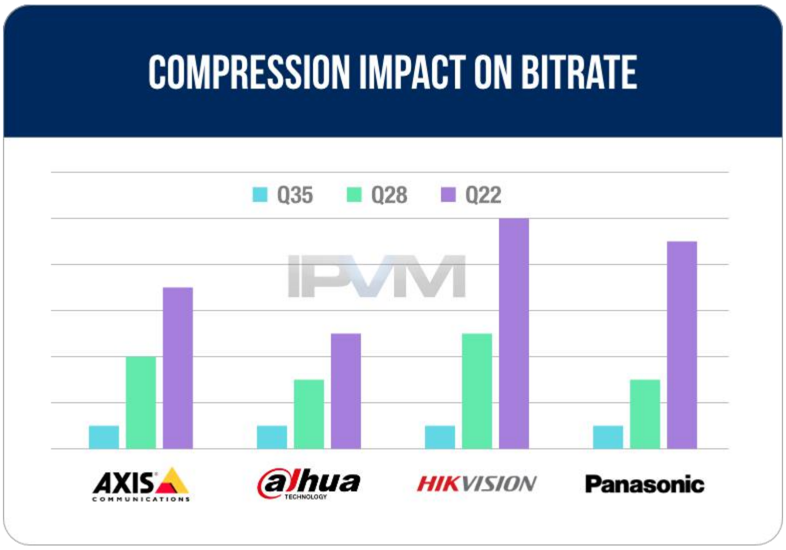

Additionally, manufacturers use different scales and terminology for their compression levels with most giving little indication of what actual quantization level is used. Some may use a numeric scale from 1-100, while others use labels such as "low, high, best", and others use the actual 0-51 quantization scale. This chart shows just some of the options in use:

MANUFACTURER COMPRESSION SCALES

MANUFACTURER	NAME OF SCALE	RANGE OF SCALE
Arecont Vision	H.264 Quality	16 - 36
avigilon	Quality	1 - 20
AXIS Communications	Compression	0 - 100
BOSCH	P-Frame Quality	Auto - 51
alhua Technology	Quality	1 - 6
HIKVISION	Quality	Lowest - Highest
Panasonic	Image Quality	0 - 9 / Super Fine to Low
SONY	Image Quality	1 - 10
unv	Quality	Quality - Bitrate
VIVOTEK	Target Quality	Medium - Excellent

See our [IP Camera Manufacturer Compression Comparison] for more detail on understanding manufacturer differences and how to standardize Q levels across different lines.

CODECs

A key differentiation across CODECs is supporting inter-frames (e.g., H.264, H.265) vs intra-frame only (e.g., MJPEG, JPEG2000).

- Inter-frame CODECs such as H.264/265 not only compress similar pixels in an image, they reference previous frames and transmit only the changes in the scene from frame to frame, potentially a large bandwidth savings. For example, if a subject moves through an empty hallway, only the pixels displaying him change between frames and are transmitted, while the static background is not.

- Intra-frame only CODECs encode each individual frame as an image, compressing similar pixels to reduce bitrate. This results in higher bandwidth as each frame must be re-encoded fully, regardless of any activity in the scene.

For more on inter and intra frame compression, see our CODEC tutorial.

The vast majority of cameras in use today, and for the past several years, use H.264, due to its bandwidth advantages over MPEG-4 and Motion JPEG. In our H.264 vs MJPEG - Quality and Bandwidth Tested shootout, H.264 consumed far less bandwidth in all scenes than MJPEG, seen in the chart below:

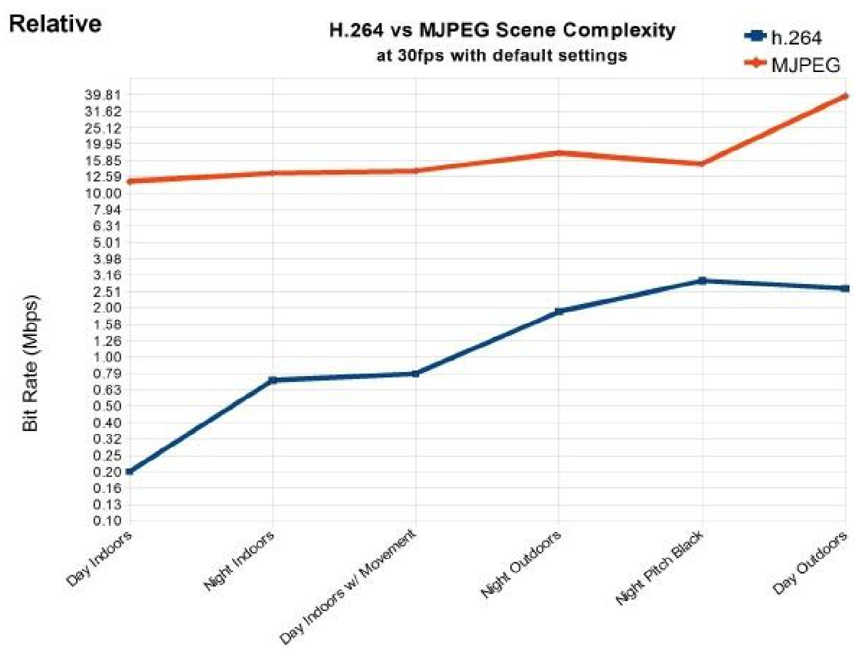

What About H.265?

H.265 has been the "next big thing" in CODECs for several years, claiming 50% savings over H.264, but camera and VMS support for it remain relatively rare. Additionally, in our tests, H.265 has had limited benefit over H.264 in similar scenes, about 10-15% on average, with H.264 Smart CODEC cameras (see section below) generally providing bigger bandwidth savings than H.265.

For example, in our Smart H.265 Samsung Test, H.265 produced ~15-20% lower bitrates than H.264 (with smart CODECs off on both), shown in the chart below. However, using smart CODECs with H.264, bitrates dropped by at least ~40% (daytime still scene) and as much as 90%+.

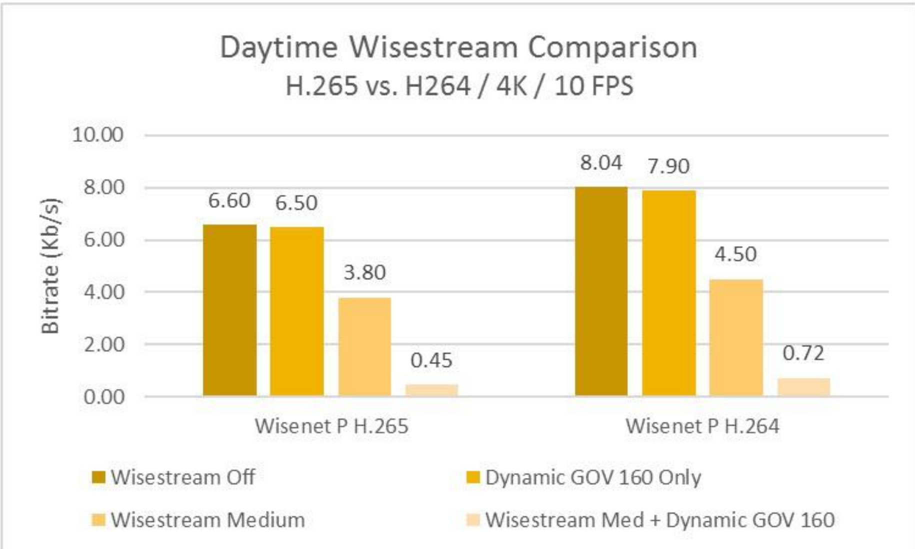

Note that H.265 is still developing, and will likely become more efficient over time, as H.264 has.

Readers should see our H.265 / HEVC Codec Tutorial for more details on H.265 issues, including bitrates, camera support, and VMS integration.

I-Frames vs. P-Frames

In inter-frame CODECs, frames which capture the full field of view are called I-frames, while those sending only changes are P-frames. Because they capture a full image, the more I-frames in a stream, the higher the bandwidth.

For years, cameras were typically only able to use a fixed I-frame interval, measured either in seconds or frames. Sending too few I-frames could negatively impact imaging, with long "trails" of encoding artifacts, while too many I-frames provides little to no visible benefit, seen in this video from our Test: H.264 I vs P Frame Impact.

Note: Click here to watch the I-Frame Intervals video on IPVM

However, with the introduction of Smart CODECs in the past 1-2 years, cameras are now able to dynamically adjust I-frame interval, instead of using a fixed value. So where a typical 10 FPS camera might be set to send an I-frame every second, a smart CODEC enabled model would extend this when there is no motion in the scene, shown in this example:

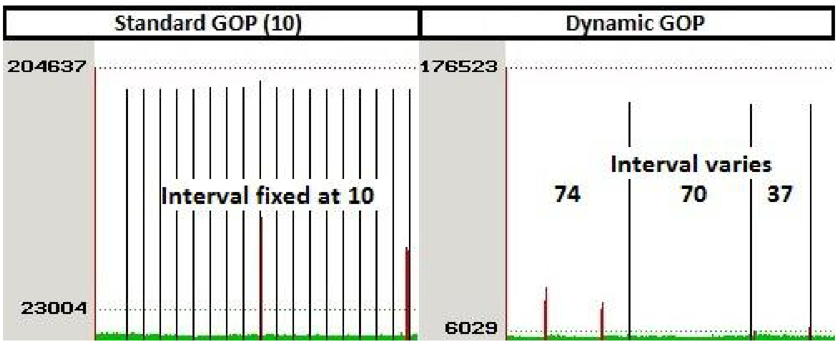

Smart CODECs are a complex topic, covered in more detail below and in our Smart CODEC Guide.

Fixed I-frame Interval Effects

Though many cameras are smart CODEC enabled and do not use fixed I-frame intervals, many (especially older models) do not and users may simply choose not to use them, so it is important to understand the impact of I-frame interval on bandwidth.

Reducing the number of I-frames (moving from 1 to 2 to 4 second interval) produces minimal bandwidth reductions, as seen below, despite the severe negative image quality impact.

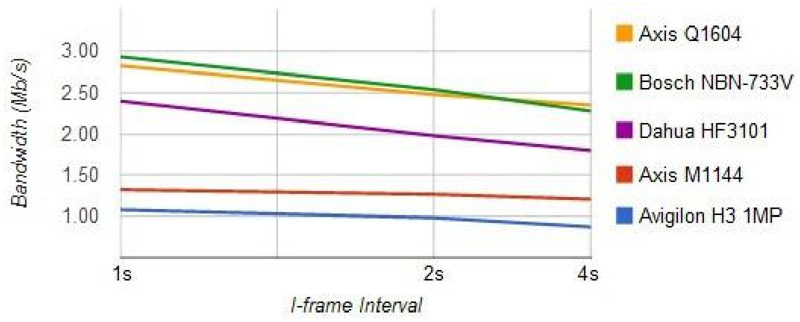

Inversely, increasing the number of I-frames to more than one per second significantly increased bandwidth, despite the minimal increase in image quality.

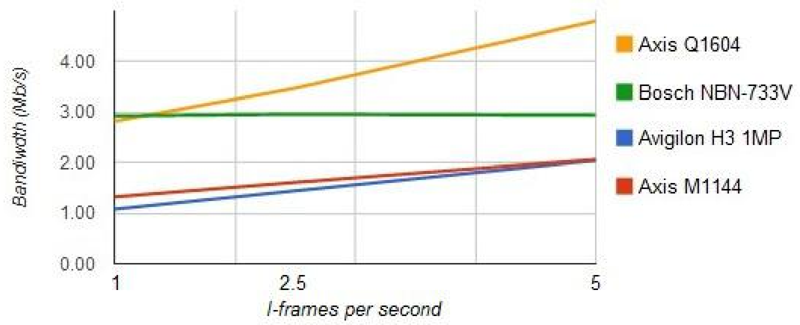

For full details on I and P frame impact on bandwidth and image quality see our H.264 I vs P Frame Test.

Smart CODECs

One recent development with huge impact on bandwidth is the introduction of smart CODECs. These technologies typically reduce bitrate in two ways:

- Dynamic compression: First, instead of using a single compression level for the whole scene, the camera may apply little compression to moving objects, with higher compression/lower quality on static background areas, since we most often do not need detailed images of still areas of the scene.
- Dynamic I-frame interval: Second, instead of using a steady I-frame interval, cameras may increase the distance between I-frames when the scene is still, with some extending to very long intervals in our tests, over a minute in some cases. Then, when motion begins, the camera immediately generates an I-frame and reduces interval to previous levels.

Some smart CODECs may use other methods as well, such as dynamic framerates (used by Axis/Avigilon), increased/improved digital noise reduction (Panasonic Smart Coding), and others.

Exact methods used by each smart CODEC and their effectiveness vary. However, in general, bitrates in still scenes were reduced by 50-75% in our tests, with over 95% possible.

As an example, in our test of Zipstream 2, bitrates dropped by ~99% in still scenes using dynamic compression, I-frame interval, and FPS:

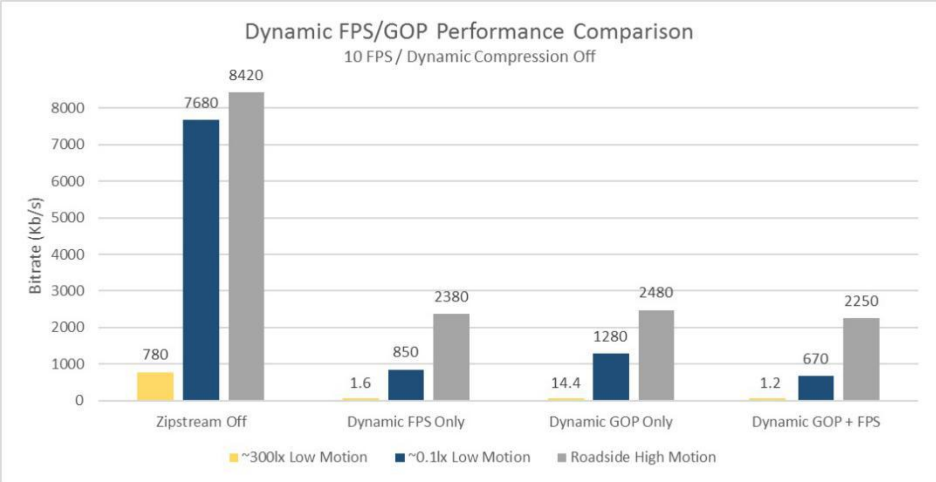

For more details, see our Smart CODEC Guide.

Camera Field of View

Field of view's impact on bandwidth varies depending on which width reveals more complex details of the scene. In scenes with large areas of moving objects, such as trees or other blowing vegetation, widening the field of view will likely increase bandwidth. In scenes with relatively low movement but repetitive backgrounds, such as parking lots, roofing, patterned carpet or walls, etc., narrowing the field of view will increase bandwidth due to more of these fine details being discernible.

For example, in the park shown below, increasing the field of view results in a ~60% increase in bandwidth due to more moving foliage and shadows in the scene compared to the narrower field of view.

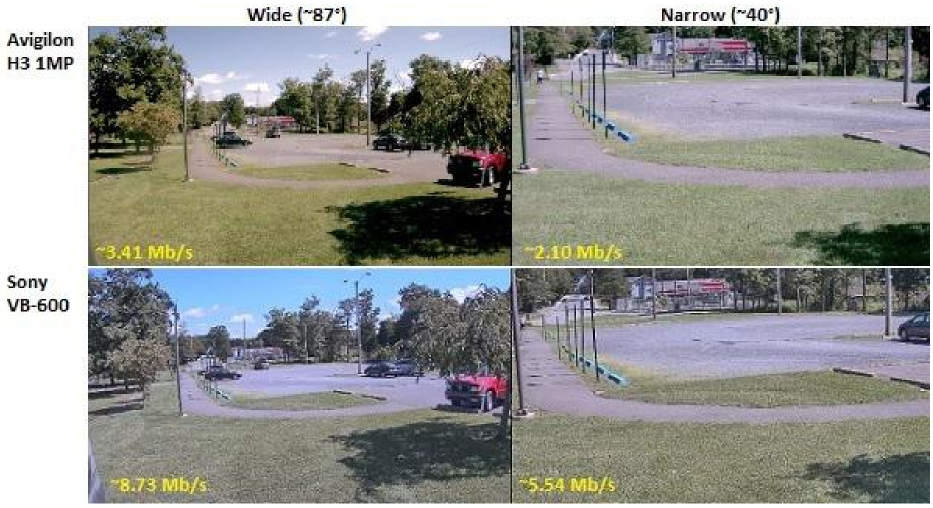

However, in a busy intersection/parking lot, bandwidth decreases by over 50% in the cameras below when widening the field of view. In the narrower FOV, more details of buildings are visible, and the repetitive pattern of the asphalt parking lot may be seen as well, making the scene more difficult to encode.

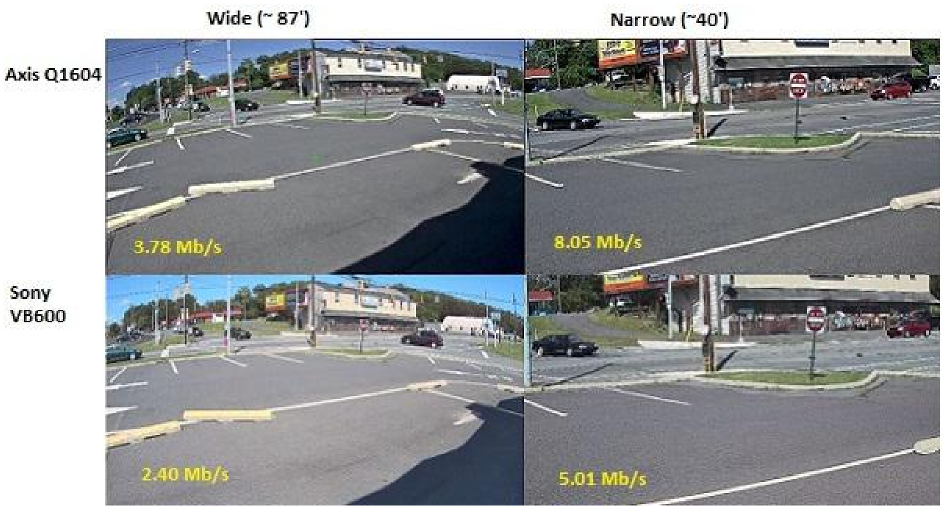

For further details of field of view's impact on bandwidth, see our Advanced Camera Bandwidth Test.

Low Light

Compared to day time, low light bitrates were an average of nearly 500% higher (seen below). This is mainly caused by increased digital noise caused by high levels of gain.

Camera	Resolution	FPS	Day	Night	Increase	% Increase
Axis Q1615	1080p	10	0.42	4.28	3.86	909%
Bosch NBN-932V	1080p	10	0.64	3.12	2.48	388%
Samsung SNB-6004	1080p	10	1.89	2.58	0.70	37%
Sony SNC-VB630	1080p	10	2.49	8.24	5.75	231%
Arecont AV3116DNv1	3MP	10	1.25	3.04	1.79	144%
Avigilon H3 1MP	720p	10	0.48	2.02	1.54	322%
Bosch 733	720p	10	0.18	0.30	0.13	73%
Dahua HF3100N	720p	10	0.19	4.00	3.81	1983%
Hikvision 864	720p	10	0.56	5.28	4.72	843%
Samsung 5004	720p	10	0.68	2.54	1.86	274%
Sony VB600B	720p	10	0.16	0.60	0.44	275%
Averages			**0.81**	**3.27**	**2.46**	**498%**

All measurements in Mb/s

However, two key improvements are increasingly used to reduce this:

- Digital noise reduction techniques have improved in recent years, greatly reducing these spikes on many cameras.
- Increased use of integrated IR cameras results in smaller spikes at night. Compared to nearly 500% in day/night models, integrated IR cameras increased by an average of 176% due to IR illumination (seen below).

Camera	Resolution	FPS	Day	Night	Increase	% Increase
Axis M1144-L	720p	10	1.20	5.44	4.24	353%
Avigilon 3.0W-H3A-BO1	1080p	10	1.15	1.32	0.17	15%
Dahua HFW3200S	1080p	10	3.20	8.80	5.60	175%
Hikvision DS-2CD2032-I	1080p	10	2.75	7.20	4.45	162%
Averages			**2.08**	**5.69**	**3.61**	**176%**

For full details of low light's impact on bandwidth, see our Bandwidth vs Low Light test report.

Wide Dynamic Performance

WDR's impact on bitrate varies depending on the camera and the scene. Again taking examples from our Advanced Camera Bandwidth Test, when switching WDR on in an Axis WDR in an outdoor intersection scene, bandwidth increases, as more details are visible (beneath the eaves of buildings, in the treeline, etc.).

However, looking at an outdoor track and sports field, bandwidth decreases. In this case, the Q1604 increases contrast slightly on some areas of the image, such as the trees and bleachers in the center/left of the FOV. Because of this, these areas are more similarly colored and easier to compress, lowering bitrate.

Note that for other cameras, these results may vary, depending on how well they handle light and dark areas, how they handle contrast when WDR is turned on, and more.

Sharpness

Sharpness has a huge impact on bandwidth consumption, yet it is rarely considered during configuration, even by experienced technicians. Oversharpening reveals more fine (though rarely practically useful) details of the scene, such as carpet and fabric patterns, edges of leaves and blades of grass, etc. Because more detail is shown, bandwidth increases.

For example, in the FOV below (from our Advanced Camera Bandwidth Test), bitrate increases by nearly 600% from minimum to maximum sharpness in the Dahua camera, and almost 300% in the Axis Q1604.

Color vs. Monochrome

At practical levels (without desaturation or oversaturation effects), color has minimal impact on bandwidth. In the examples below, moving from default color settings to monochrome decreases bandwidth by 20 Kb/s, about an 8% decrease.

However, oversaturation may result in abnormally high bandwidth. In this example, bandwidth increases by over 200% when changing color settings from default to their highest level, which also creates oversaturation effects such as color bleeding (seen in the red chair).

One practical example of a manufacture desaturating their video to 'save' bandwidth is [Arecont Bandwidth Savings Mode (which we tested here)].

Manufacturer Model Differences

Across specific models in a given manufacturer's line, significant differences in bitrate may occur, despite the cameras using the same resolution and framerate. This may be due to different image sensors or processors being used, different default settings in each model, better or worse low light performance, or any number of other factors.

For example, the following image shows two cameras, an Axis Q1604 and Axis M3004, both 720p, 10 fps, set to a ~20' horizontal FOV, at compression of ~Q28. Despite these factors being standardized, in this well lit indoor scene, the Q1604's bitrate was 488 Kb/s while the M3004 consumed 1.32 Mb/s, nearly 3x the bandwidth.

Beware: model differences have become more extreme in some cases, as some cameras support Smart CODECs while others in the same line may not.

Measure Your Own Cameras

As this guide shows, there are few easy, safe rules for estimating bandwidth (and therefore) storage, abstractly. Too many factors impact it, and some of them are driven by impossible to know factors within the camera.

Though it is important to understand which factors impact bandwidth, use this knowledge with your own measurements of the cameras you plan to deploy. This will ensure the most accurate estimates and planning for deployments.

Lux Rating / Minimum Illumination

Lux ratings are one of the poorest specifications to use in selecting cameras. Now, with the rise of integrated IR, they are increasingly useless.

You need to be able to understand why lux rating (aka minimum illumination specifications) are so problematic, how they are established and what tricks / techniques are used.

We explain why lux rating (aka minimum illumination specifications) are so problematic, how they are established and what tricks / techniques are used:

- How lux ratings are tested / determined
- The incorporation of shutter speeds in lux ratings
- Dealing with lux ratings that include 'sens up' settings
- Color vs B/W Impact
- Understanding how IRE levels are used
- Advances in super low light cameras
- The practical lux levels typically specified based on analyzing 2000+ camera specifications
- How to avoid getting burnt by lux levels
- IR illumination and lux ratings

Lux Rating Tests / Determination

Most importantly, there are no standardized or verified means to assess manufacturer lux ratings. They are always self-assigned and, at the discretion of the manufacturer.

This means that each manufacturer gets to decide what light source they use, the size of the testing area, the positioning of the light, the test subject / chart employed, etc.

Most critically, each manufacturer decides when an image is or is not usable. It is this point, solely at their discretion that becomes the self-assigned lux rating.

For example, in the image below, each lux level could easily be considered "dark", but with no standardization of levels nor minimum usability, which is "right"?

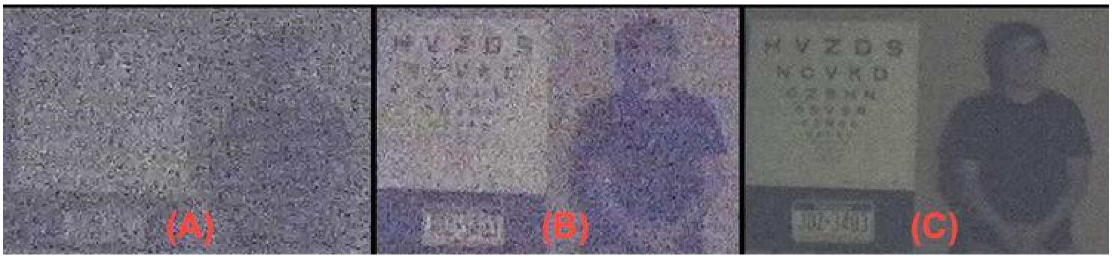

Since each manufacturer is free to make their own goal, they have an incentive to choose the darkest one possible, knowing many of their competitors will do the same thing.

Shutter Speeds and Lux Ratings

Often manufacturers will list lux ratings at different or multiple shutter speeds. Some list 1/30s, others 1/2s, some still 1 full second shutter. Users should ignore specifications at any shutter speed other than 1/30s. While these longer shutter speeds allow more light and lower minimum illumination levels, [motion blur is a significant problem](#) with slow shutter speeds.

In the spec sheet example below, the manufacturer lists minimum illumination at 1/30s and 1/2s shutter speeds:

MIN. SUBJECT ILLUMINATION - DAY (COLOR)	0.3 lux (F1.2, shutter speed 1/30 sec., SSC off , 50IRE) 0.02 lux (F1.2, shutter speed 1/2 sec., SSC off , 50IRE)

Of course, the minimum illumination specification 'looks' better at 1/2s, but that is not technological improvement, simply trading off more light for more blur.

F Number and Lux Ratings

Every so often a manufacturer will specify their minimum illumination assuming a different F number than the lens the camera uses. For example, a manufacturer might say their camera lux rating is 0.01 lux at f/1.0 but the camera may have an integrated f/2.0 lens.

In the example below, this is the case, with minimum illumination listed at F1.2, but the camera shipping with an F1.8 lens, which captures less than half the amount of light of an F1.2 lens.

Model	DH-IPC-HFW1000S
Camera	
Min. Illumination	0. 1Lux/F1.2 (Color), 0.01Lux/F1.2(B/W);0Lux/F1.2(IR on)
Lens	
Focal Length	3.6mm (,6mm optional)
Max Aperture	F1.8(F1.8)

Even assuming the manufacturer's own rosy self assessment, correcting this basic error means the minimum illumination is at least .225 lux rather than .1 lux.

Sens Up and Lux Ratings

Sens-up is typically a marketing term for slow shutter. The higher the sens up "level", the slower shutter, with each multiplier (2x, 8x, 64x, etc.) simply multiplied times 1/30s to produce the effective shutter speed.

For example, the spec sheet below lists a separate minimum illumination at "x256 Sens-up." This essentially amounts to an astounding 8.5 second (256/30s = 8.53) exposure time, which would result in massive ghosting of moving objects.

Min. Illumination:	0.25lux color / 0.01lux B/W (50IRE, F1.2) 0.001lux color / 0.00004lux B/W (@ x256 Sens-up)

The hope is that you see the lower lux rating and are impressed. Beware.

Color vs B/W

Cameras that support integrated cut filters (aka D/N cameras) will often list 2 lux ratings, the first in color mode, the second in monochrome. This difference is due to increased performance in monochrome mode due to ambient IR light which the IR cut filter blocks in color mode. Claimed differences are often substantial, as in the examples below, which both list monochrome minimum illumination specs 1/10th of color mode (theoretically 10x better).

Avigilon 2.0-H3-B2:

Minimum Illumination	0.2 lux (F1.2) in color mode; 0.02 lux (F1.2) in monochrome mode

Hikvision DS-2CD864FWD-E:

DS-2CD864FWD-E		
	Image Sensor:	1/3" Progressive Scan CMOS
	Minimum Illumination:	Color: 0.01 lux@F1.2, AGC ON B/W: 0.001 lux@F1.2, AGC ON

While gains due to ambient monochrome light do occur, 10x increases are typically overstated. Camera image enhancement and improved DSP has brought performance of color and monochrome imaging much closer, as seen in this comparison image from a current generation camera in the same scene.

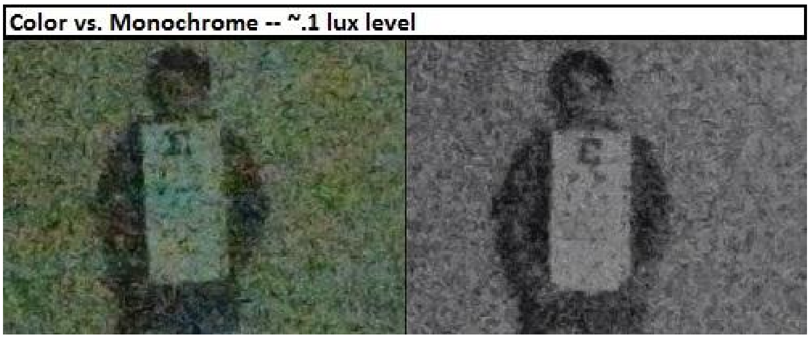

IRE Levels

IRE levels could be helpful for analog cameras but are not applicable in IP. IRE is a measure of the contrast level in an analog video signal, tested using composite video outputs. Since IP cameras do not provide analog output of their full resolution, it is a moot metric.

A few manufacturers still list it on spec sheets, showing different minimum illumination levels for different IREs (most often 30 and 50). However, since it is unknown how manufacturers are testing this IRE, we recommend using the worse minimum illumination spec in these cases (typically 50), and ignoring lower readings.

Samsung SND-7084N:

Min. Illumination	Color : 0.1Lux (F1.2, 50IRE) 0.06Lux (F1.2, 30IRE) B/W : 0.01Lux (F1.2, 50IRE) 0.006Lux (F1.2, 30IRE)

Sony SNC-EB630:

Minimum Illumination (50 IRE)	Color: 0.1 lx (F1.2, View-DR OFF, VE OFF, AGC ON, 1/30 s, 30 fps)
	B/W: 0.07 lx (F1.2, View-DR OFF, VE OFF, AGC ON, 1/30 s, 30 fps)
Minimum Illumination (30 IRE)	Color: 0.06 lx (F1.2, View-DR OFF, VE OFF, AGC ON, 1/30 s, 30 fps)
	B/W: 0.05 lx (F1.2, View-DR OFF, VE OFF, AGC ON, 1/30 s, 30 fps)

Thought HD analog cameras have risen in the past few, almost all of them have integrated IR, so even while IRE levels could be applied to HD analog cameras, they would not make a difference.

Beware IR Sensitive

Some manufacturers (notably Arecont) attempt to mislead specifiers by listing "0 Lux, IR sensitive" for their non-integrated IR cameras. What they are saying is "If you buy and add your own IR light source, our cameras need no light." That is trivially true for any D/N camera but misleading because it requires adding one's own IR illuminator units.

Super Low Light Advances

In past years, it was generally and mostly accurately believed that all cameras were poor in low light, even the 'best' would struggle with high levels of noise and darkness in truly dark conditions.

This has changed in the past few years as (1) advances in low light image processing have increased and (2) as some larger imagers are used. Cameras can now immediately process and enhance video quality as processing power available inside the camera increase. Moreover, 1/2"

1080p cameras have emerged as a niche and those cameras, combined with image processing, deliver low light levels impractical in the past (e.g., see our Axis 1/2" and Hikvision 1/2" camera tests). Finally, there are a few very large imager IP cameras (e.g. 1' or 35mm) that deliver in superior low light levels though at far higher prices (e.g., Sony's 4K 35mm test).

Practical Lux Levels

IPVM lists the minimum illumination specifications of 2000+ cameras in its Camera Finder. From this, we found these practical levels:

- .00X Lux and Below
- .0X Lux and Below
- .X and Above

0.00X Lux And Below

Manufacturers specifying minimum illumination at these levels are often aggressively overstating their camera's low light performance. In some cases, these claims are due to manufacturer tricks such as slow shutter, while in others, it is simply overstatement.

Among the unbelievable models in this category are standouts such as Speco's IP intensifier line (which blurred moving objects massively due to slow shutter, see our test results) and Vivotek's FE8174V, a 5MP panoramic camera with a very high F2.8 lens.

In short, be careful of manufacturers making claims below 0.01 lux, as performance rarely matches these specs.

0.0X Lux

We found this range to be where most super low light cameras are typically specified. This range includes top low light performers such as Axis' Lightfinder Q1604 and Q1615, Samsung SNB-5004 and SNB-6005, and Sony's 6th Gen SNC-VB600 and SNC-VB630, making it the "safest" categorization for those seeking top low light performance.

This is not to say no manufacturers are overstating performance at these levels. For example, both the Avigilon 3.0W-H3-B2 and Axis Q1604 are specified at 0.02 lux. However, our tests show the Avigilon camera's performance is well below that of the Q1604.

0.X Lux

Cameras specified at this level are generally poor in low light. This category contains many previous generation cameras whose low light performance has been superseded by new generation models. Additionally, many lower cost fixed lens/high F-stop models which more likely belong in the 1+ lux range are specified here.

With both of these factors considered, if low light is a key concern, cameras specified at this level will likely not deliver necessary performance.

1+ Lux

Cameras with minimum illumination ratings above 1 lux will be poor in even moderate low light conditions. These ratings are most likely among lower cost fixed lens minidomes (such as the Axis M30 and Bosch microdomes) and panoramic models such as the Samsung SNF-7010 and Panasonic WV-SF438/458.

IR illumination and Lux Levels - BIG Impact

Most cameras being sold / deployed today use integrated IR.

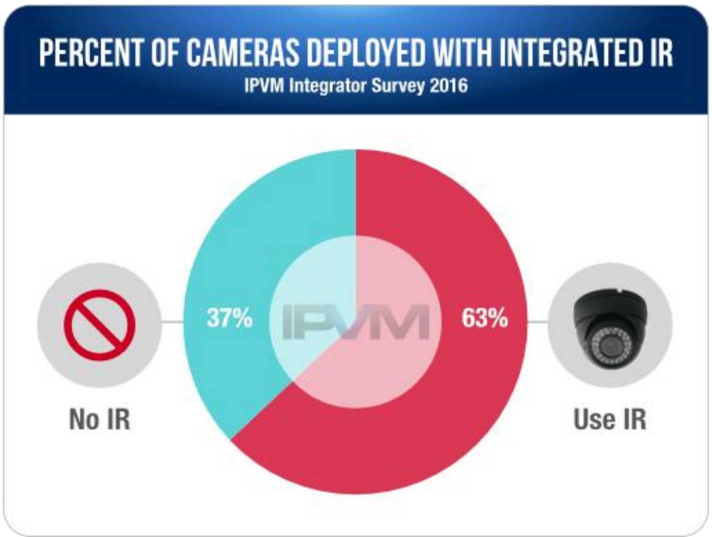

Because of that, lux ratings do not generally apply to them, since these cameras emit their own lighting. It is possible in most IR cameras to disable the IR illuminator but this is typically not done. As such, most IR cameras report a lux rating of 0 Lux, since they do not need light.

An additional lux rating is often provided for IR cameras when IR is disabled, sometimes called color mode (since when IR is on, the camera will be black and white). See this example below:

Min. Illumination:	Color: 0.028 lux @ (ƒ/2.0, AGC on); B/W: 0 lux with IR

This lux specification suffers from the same concern as all other cameras. Moreover, since IR cameras have their own illumination, many, especially lower cost ones, are poor in low light, relying on their built-in IR illuminations to make the image bright (e.g., 4MP Omnivision imager based cameras).

What To Do

If low light is important to you, you do not want IR, and you must specify minimum illumination, IPVM recommends the following specification language:

"Minimum illumination of 0.09 lux at 1/30s shutter, no sens up allowed."

The rationale for this is that cameras specified at 0.1 lux or higher are almost always fairly bad in low light, so, at least, you can reject those cameras. However, if you specify something lower than that, like .001, you significantly increase the chance of rejecting high-quality low light cameras that are conservatively specified.

Finally, make sure to include the "1/30s shutter, no sens up allowed" to prevent manufacturers including specifications that are heavily gamed and certain to introduce problematic motion blur.

Beyond that, review IPVM test results that all include low light standardized testing and test yourself to see how well it works in your scenes / light levels.

Surveillance Camera Imagers

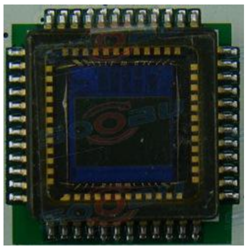

Imagers - CCD, CMOS, 1/2", 1/4", big pixels, small pixels, etc.

We explain the fundamental issues and drivers in surveillance camera imagers, including:

- Sensor vs Imager
- CCD vs CMOS
- Imager Manufacturers
- Camera Manufacturer Imager Disclosure
- Imager vs Resolution
- Imager Size
- Pixel Size
- Imager vs FoV Width
- Imager vs Low Light Performance

Sensor vs Imager

Industry people alternatively call these components 'sensor' or 'imager'. When referring to surveillance cameras, they mean the same thing, though, technically, an imager is a specific type of sensor. Because of this, we more typically refer to 'imager' when speaking about surveillance cameras.

CCD vs CMOS

The two main historical types of imagers have been CCD and CMOS.

Today, in 2017, virtually all surveillance cameras use CMOS, from the very 'best' to the 'worst'. This is the opposite of a decade ago, when CCD imagers were predominant and CMOS was looked down upon as a lower cost, lower quality alternative.

Because CCD once was better, a prejudice remains against CMOS. However, this is wrong, antiquated and, as a practical matter, impossible. If you only choose CCD imagers, you would eliminate almost all modern surveillance cameras, including the 'top' brands and models.

Imager Manufacturers

There are only a few significant manufacturers of surveillance imagers, with the most frequently cited including Omnivision, Aptina and Sony. Like camera manufacturers, imager manufacturers offer a range of models with varying size, max resolution, frame rate and WDR capabilities, to name a few.

Camera Manufacturer Imager Disclosure

Camera manufacturers generally hide what imager manufacturer / models they use, so it is not easy to compare two cameras based on their imagers. Some will list the imager manufacturer but not the specific model. For example, even if you knew it was an Omnivision, Omnivision makes dozens of imagers with varying price / performance tradeoffs. Moreover, even if you knew 2 cameras were using the exact same imager (which is rare, in practice), differences in tuning, encoding and compression could still result in noticeable quality differences.

Imager vs Resolution

Imagers vary in the maximum resolution they support. This constrains the camera's overall resolution. Some imagers max out at VGA, 1.3MP, 3MP, 5MP, etc. Recently, 4K / 10MP sensors have emerged which is helping to foster 4K cameras.

Imager Size

Imagers can range from extremely large (e.g., DLSRs) to extremely small (e.g., cell phones). Surveillance camera imagers tend to fall in the middle, closer, but typically larger, than cell phone imagers.

The image below shows contrasting imager size. The two larger ones on the left and center are atypical for surveillance. The one on the right (1/1.8") is actually still fairly large for typical surveillance cameras:

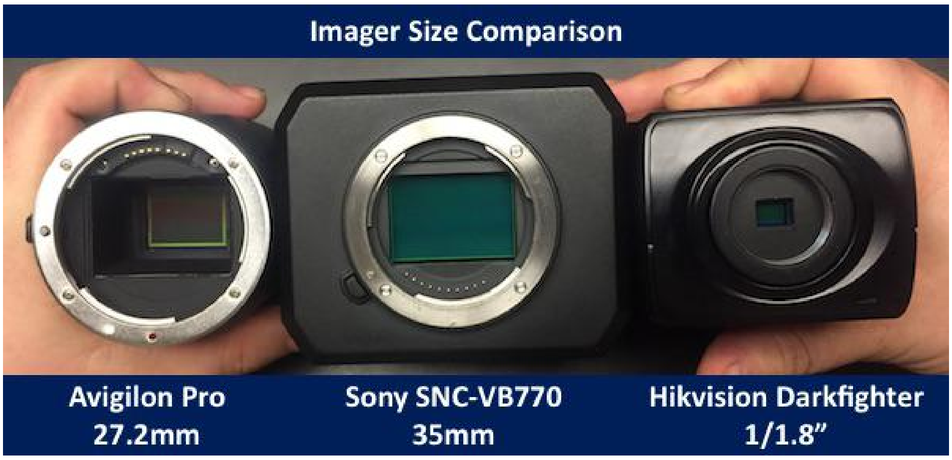

In surveillance, 95% of cameras have imagers that are between 1/2" and 1/4". The most common imager size in surveillance is ~1/3", with 1/2.7" and 1/2.8" common. Previously, lower resolution models (720p/1.3MP and below) often used smaller 1/4" imagers, but even these models predominantly use 1/3" sensors today.

Increasing Imager Size

Over the past few years, the average imager size has increased moderately. As resolutions increase, 3MP, 4MP, 5MP, 4K, etc., imager sizes larger than 1/3" are definitely becoming more common.

Rise of 1/2" 1080p Imagers

The most important recent shift in imager size used is the rise of 1/2" 1080p imagers. Many manufacturers now have at least one "high end" model using a ~1/2" (or 1/1.8", 1/1.9", etc.) sensor, targeted for super low light.

Larger Than 1/2" Still Rare

Imagers larger than 1/2" are still very rare in surveillance, with only a handful of options. Avigilon's Pro series is best known for their use of large imagers (27.2mm/~1.07"), but others have become available or been announced, such as Sony's 35mm VB770, and the Axis/Canon APS-C model slated for 2017.

Size Disclosure

Unlike imager manufacturer, imager size is almost always disclosed by camera manufacturers on specification sheets. Here is an example of the level of detail typically provided:

Image sensor	1/4" progressive scan RGB CMOS
Image Sensor:	1/3" Progressive Scan CMOS
Image sensor	Progressive scan RGB CMOS 1/2"

Pixel Size

Imagers vary in the size of their pixels, measured in microns.

Pixel size is most strongly determined by imager size and number of pixels (i.e., resolution). The bigger the imager, everything else equal, the bigger the pixel size. However, if you add more pixels (e.g., going from SD to HD) and the imager size stays the same, the pixel size decreases.

Here's an excerpt from an imager manufacturers showing imager size, resolution and pixel size side by side:

Optical Format	1/4"	1/3"	1/3"	1/4"	1/3"	1/3"	1/2.5"
Resolution	VGA	1.2 MP	1.2 MP	1.1 MP	3.5 MP	3.1 MP	5 MP
Pixel Size	5.6 µm	3.75 µm	3.75 µm	3.0 µm	2.2 µm	2.2 µm	2.2 µm

Many prefer larger pixel sizes because, everything else equal, a larger pixel can collect more light, and therefore deliver brighter / better low light images. However, many other factors impact low light performance so it is not simple / easy to conclude one camera is better than another based on pixel size.

Also, pixel size is almost never disclosed by camera manufacturers, so the best one can do is estimate by looking at the resolution and imager size of each model.

4K Pixel Size

4K cameras are increasingly common. Although they have ~4x the pixels of 1080p cameras and 4K camera imagers are generally larger than 1080p, generally the pixel size for 4K cameras is smaller. This can hurt low light performance. Also, as we explain further, limits in image processing also hurt 4K cameras.

Imager Size vs FoV Width

Imager size has a modest impact on FoV width. The primary determinant of FoV is lens length (e.g., 3mm, 10mm, 30mm, etc.).

However, the larger the imager, everything else equal, the larger the FoV.

On the other hand, imager sizes in surveillance do not vary that much, so even with notably different imager sizes, the FoV only changes moderately. For example:

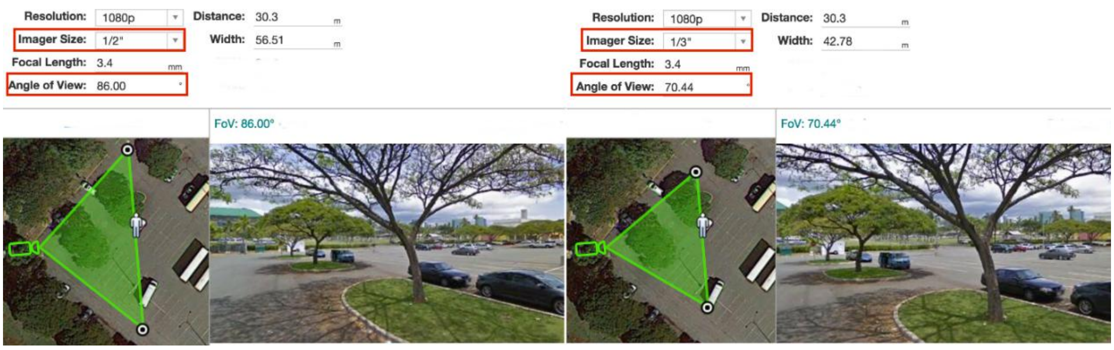

As such, it generally is not a major concern, but is worth being aware of.

Imager vs CODECs

Imagers have nothing to do with compressing video (i.e., H.264, MJPEG, etc.). Imagers can certainly impact image quality by what they capture but the imager sends the video uncompressed to the an encoder / System on a Chip (SoC) to perform this.

Imager vs Low Light Performance

Besides CCD vs CMOS, the strongest, and most flawed, belief in imagers is that larger imagers always deliver better low light performance. Though larger imagers help, there are other key drivers that typically are more important, namely low light image processing that is done by the encoder / CPU on the camera. IPVM has shown this in our imager size vs low light performance study.

However, the newer 1/2" 1080p cameras specialized for low light are generally better than the top 1/3" 1080p cameras, as these models typically apply the same (or better) processing used in 1/3" models to larger 1/2" imagers with better sensitivity, resulting in overall better performance. See our tests of Axis' 1/2" Q1635 Camera, Dahua Starlight 1/2" Camera, Hikvision Darkfighter, and Samsung 1/2" Wisenet III Camera for examples.

Additionally, there are a number of 1/2" imager cameras available lacking image processing that are quite poor in low light. Also, a number of 4K cameras have 1/2" imagers or larger, but are still bad in low light because the pixel size is fairly small and the cameras tend to lack the processing power to enhance such large resolutions.

Quiz Yourself

Take the 7 question quiz and see how well you know imagers.

Updated 2016

This post was originally released in 2014 but was updated in 2016 to reflect developments in 1/2" imagers and 4K cameras.

Wide Dynamic Range

Understanding wide dynamic range (WDR) is critical to capturing high quality images in demanding conditions. However, with no real standards, any manufacturer can claim WDR, and many do even if actual performance is weak, causing confusion for even experienced users.

We break down these issues, explaining:

- State of WDR
- What is WDR?
- Interpreting WDR Specifications
- How Do You Measure WDR?
- Camera WDR Challenges
- WDR vs Resolution
- True WDR Implementations
- Fake WDR Concerns
- Manufacturer Terminology Confusion

2018 State of WDR

Trends seen in 2017 continue moving into 2018:

- True WDR hits low cost: True WDR was once a premium feature, but performance comparable to "high end" models can now be found even in low cost cameras, such as Hikvision's/Dahua's 4MP models or Axis' new M30 line (all with sub-$200 models). Compared to the $500+ WDR models typical of the past, this pricing is substantially lower (see over 150 true WDR models under $400 in our Camera Finder).

- HD analog WDR: Further, WDR can now be found on many HD analog models, where those looking for better WDR performance previously needed to use IP cameras. This further reduces WDR's cost premium, as some of these models can be found for under $100.

While WDR is hardly a "universal" feature and performance still varies widely, we expect to see this shift to lower cost WDR continue as sensors and SoCs continue to improve.

What is WDR?

WDR stands for **W**ide **D**ynamic **R**ange, and is essentially the ability to produce high quality image across a range of light levels. For example, in the image below, the center of the scene is very bright due to the sunlight behind the person entering, while the sides of the hallway are much darker, and the back wall (facing the camera) darker still as it does not receive any direct sunlight. We will review exactly how to measure these light levels below.

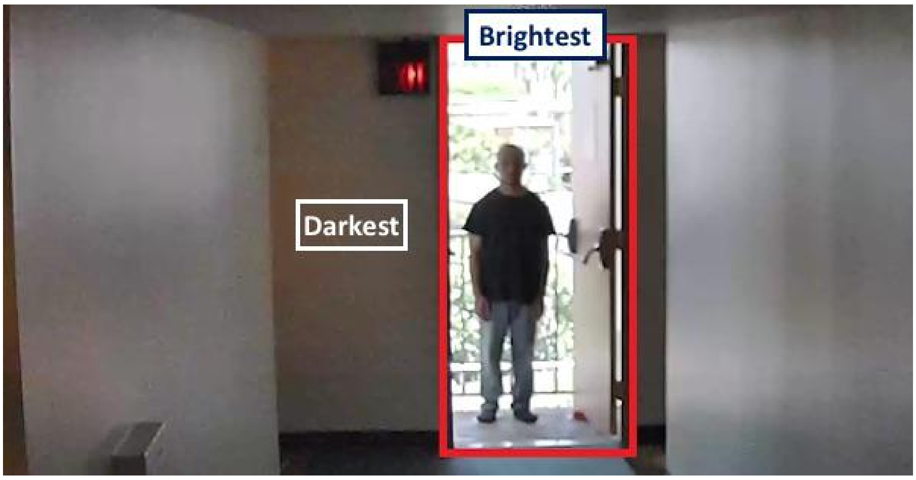

WDR can make a big difference in scenes with widely different light levels. Here's an example of a person walking in a doorway. The non-WDR camera shows only rough details, but not facial recognition, while the true WDR camera clearly shows the subject's face, hairstyle, and more.

Camera WDR Challenges

Normal cameras typically struggle with wide ranges of lighting because of their dependency on a single exposure. Cameras need light to generate an image. However, too much light and the image is washed out yet too little light and the image is too dark.

If you have a scene with even lighting, it is no problem. The camera will simply adjust its iris opening size or its shutter speed to get the right amount of light. This is why manufacturers typically demo their cameras in even lighting scenes.

However, if the scene has a wide range of lighting, the camera has a tough challenge. If it restricts the amount of light it takes in to optimize for the bright areas, the lower light areas will be too dark. However, if it chooses the opposite approach, optimizing for low light areas, the bright side will be washed out. The image below shows this tough tradeoff:

Something, therefore, needs to be done to overcome this.

Measuring WDR

It is critical to know how wide the range of light levels a scene has. The generally accepted unit to measure this in surveillance is the decibel (e.g., 58dB, 113dB, etc.) with higher levels indicating stronger WDR performance.

Unfortunately, these measurements are not standardized and are at the discretion of each manufacturer, so should not be trusted. Worse, manufacturers may offer different "grades" of WDR, making things even more confusing (see section below).

While dB measurements alone are fairly cryptic (i.e., what does it mean physically?), they are grounded on a specific test scenario. A grayscale chart is used with numerous shades from white to black. The more levels a camera can display/capture, the higher its dB rating and the better its WDR performance should be. Below is an example image from Pixim's WDR measurement whitepaper:

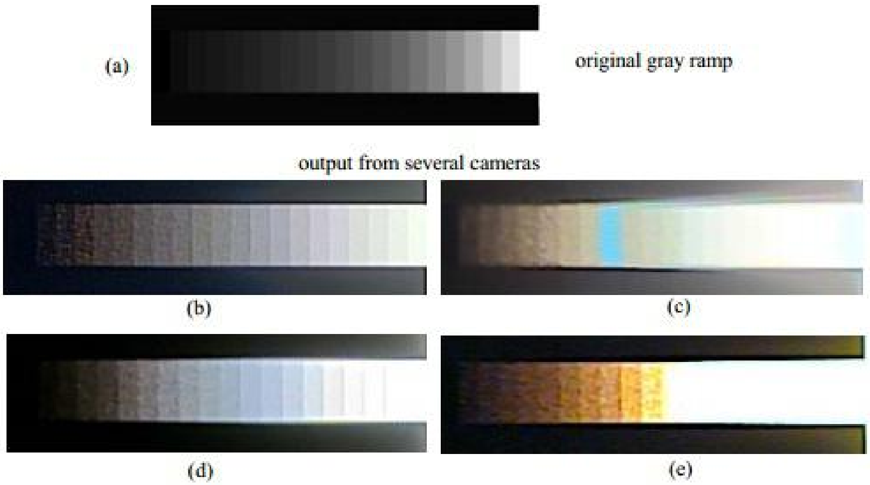

Unfortunately, this does not translate well to real world scenes (i.e., what can I expect for my front door camera?) nor is this approach guaranteed to be used or measured fairly by all manufacturers.

Interpreting WDR Specifications

You should never trust manufacturer dB specifications as they are self-assigned and not validated independently. However, there are some patterns to consider:

- A WDR dB rating of 70dB or less almost always means that the camera does not support true WDR as manufacturers know that rating their cameras this low will ensure that the camera is low on the WDR 'scale'. 'Regular', non WDR cameras frequently are specified in the range of 55 - 70dB.
- A WDR dB rating of 100+ usually, but not always, means that the camera supports true multi-exposure WDR. Again, though, since manufacturers are free to choose any value they want, it is not guaranteed. Make sure to verify via your own tests or our WDR test reports.

- Related, some manufacturers are relatively conservative in their WDR dB ratings. For example, Sony lists many of their true WDR cameras as 'only' 90dB. This does not mean they are worse than a more liberal manufacturer self-assigning 120dB or 130dB. Actual testing is required to verify.
- In the past few years, we have seen manufacturers accelerate the dB race with specifications of 130dB or even 140dB. However, in our testing, cameras with lower dB specifications (even 90dB) may out perform them in real-world testing.

An Alternative Approach

In our testing, we have developed an alternative real world approach to measuring how tough a WDR scene is. Using a lux meter, we record the brightest and the darkest spots of each scene. The ratio of the two provides a strong indicator of how challenging the scene is.

In a doorway opening outdoors, with a small opening to an enclosed areas, the range is typically quite high, as shown below:

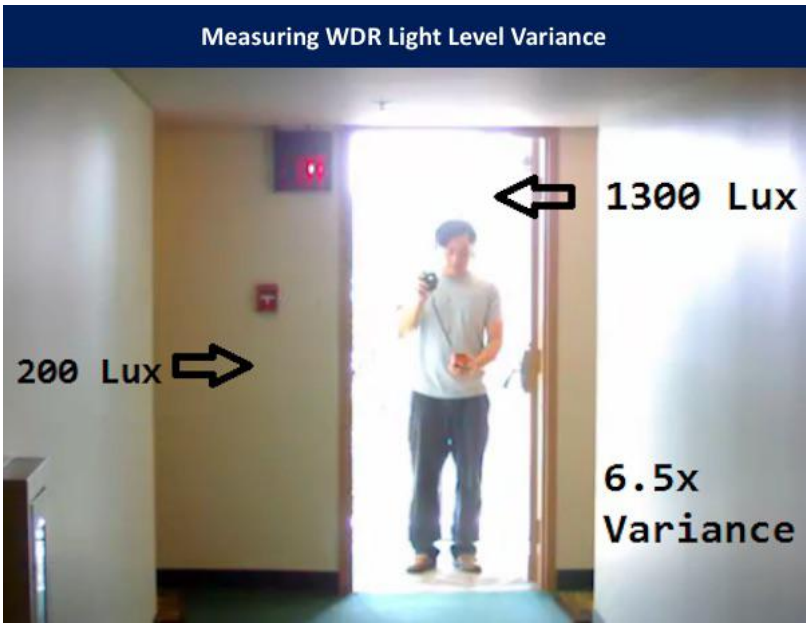

The WDR ratio is ~6.5x, with the open doorway at 1300 lux and the adjacent indoor sides at ~200 lux.

The closer one gets to a WDR ratio of 1, the less likely that WDR capability is needed. Moreover, the lower ratio the less powerful WDR functionality one needs.

Even with our approach you will need to measure but this can be done by a field tech with a $100-$150 lux meter. However, do make sure to do this when the sun is strongest as the WDR ratio will vary throughout the day as the sun moves.

Resolution vs. WDR

In general, higher resolution improves WDR, all other factors being equal. This is NOT the best nor the most sophisticated approach (see more below). However, increasing the number of pixels helps capture finer details even in the darker / brighter areas. See an example from our tests:

The worst performer is the SD camera, even though it use true WDR. The 5MP non WDR camera beats it simply because it captures more details. Finally, the HD camera with true WDR performs the best.

Note that multi-exposure WDR is generally always preferred for extreme wide dynamic scenes, as simple resolution increases do not compare to today's WDR. However, all things equal, a higher resolution non-WDR camera may likely outperform a lower resolution model.

WDR Implementations - True vs Fake

Since the main challenge for WDR scenes is setting the exposure appropriately to capture both dark and bright areas, the most common viable solution is to use multiple exposures and then combine them to produce a better quality image. The short exposure captures the bright areas, while the long exposure captures the dark areas. See representation below:

In our testing, this is the core strength of all top performing WDR cameras.

While we recommend looking for WDR cameras using multiple exposures, this is not sufficient. The number of exposures used and the other image processing techniques implemented can also make a difference. However, none of these are typically revealed.

Multiple Exposures Hurt Low Light

A common downside of multi-exposure WDR implementations is worse performance in low light. Using multiple exposures typically restricts how slow a shutter can be set. However, when it is dark, slower shutters bring in more light, producing a brighter image (though go too far and you have bad motion blur).

When using WDR cameras, make sure that WDR is disabled at night for maximum low light performance. Some cameras do this automatically, others allow for manual configuration and a few have no option. This is an important element to check if low light is a priority for the location deployed.

Fake WDR Techniques

Two other pseudo-WDR techniques have often been claimed by manufacturers as being alternatives to multi-exposure WDR. However, our tests have shown that these techniques simply do not compare to even a low performance true WDR camera:

Backlight Compensation

BLC, or backlight compensation, simply adjusts the (single) exposure of a camera. This is useful only when you want to capture just the bright or dark areas of the scene but not both. By using BLC, you will make one portion of the scene better at the expense of the other being worse.

"Digital" WDR

Digital WDR, also known as DWDR, EWDR, or electronic WDR is essentially a contrast adjustment on multiple regions of the scene, instead of the overall image. DWDR may be better in some scenes than doing nothing, but these techniques still do not compare to true WDR.

Note that many manufacturers offer both digital WDR in addition to true WDR, and users may easily confuse the two, such as Axis' "Dynamic Capture" vs. "Dynamic Contrast". Beware of these terms and refer to our WDR Cheat Sheet (also below).

Same Manufacturer - WDR vs Non WDR

When a manufacturers offers 'true' WDR, that version typically offers substantially better performance in harsh lighting conditions than the non WDR version. On the other hand, it is often costs a few hundred dollars more than the non WDR version.

Here is a comparison of Axis WDR vs non WDR:

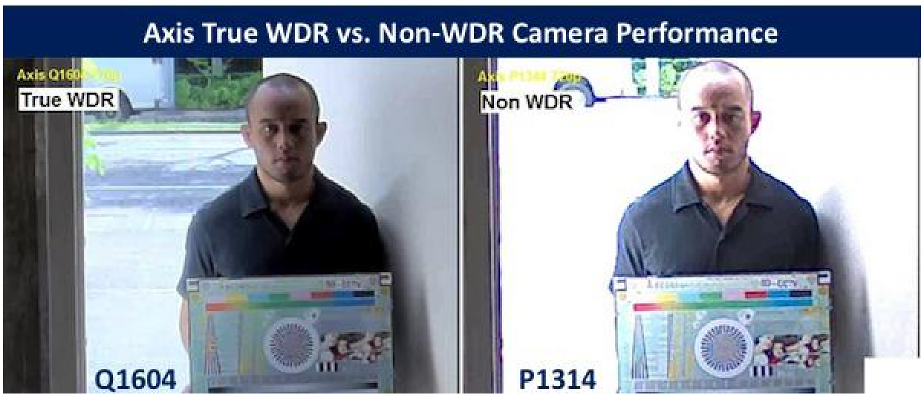

And here is a comparison for Sony:

Note: Do not compare the Axis and Sony as these shots are from different tests/times and both have released more up to date models. These images are provided here as examples of model differences only.

WDR Manufacturer Cheat Sheet and Camera Tracking

Because manufacturers can be cryptic about WDR support, we have put together a list of manufacturers detailing whether they support 'true' (multi-exposure) WDR, 'fake' (digital/electronic WDR), or both. We also provide notes to marketing specifics, naming conventions, etc.

All 15 manufacturers have cameras that support 'true' WDR, while about half of them also have cameras that support digital WDR.

Manufacturer	True WDR	Digital WDR	Notes
ACTi	Superior WDR Extreme WDR	Basic WDR Advanced WDR	All cameras marketed vaguely as "WDR"
Arecont	WDR	N/A	2 and 3MP models ending in 6
Avigilon	WDR	N/A	Only in models marked C or W, e.g., 2.0C-H4A-* or 3.0W-H3-*
Axis	Dynamic Capture Forensic Capture WDR	Dynamic Contrast	Forensic Capture and "WDR" now available in select M series cameras. Previously only found in P and Q.
Bosch	HDR	N/A	
Dahua	WDR Advanced WDR Ultra WDR	DWDR	See notes below.
FLIR	WDR sWDR Ultra WDR	N/A	WDR and sWDR used in DVTel models. Ultra WDR in Dahua.
Geovision	WDR Pro	WDR	All cameras marketed vaguely as "WDR"
Hikvision	WDR Lightfighter	DWDR	True WDR model numbers contain "FWD", e.g., DS-2CD864**FWD**
Panasonic	Super Dynamic SDII SDIII SD5 Mega SD Enhanced SD Extreme SD	N/A	Variants differ by number of exposures and other processing.
Pelco	SureVision SureVision 2.0	N/A	VGA-1.3MP SureVision models, SVGA-3MP SureVision 2.0 models
Samsung	WDR Advanced WDR	SSDR	Wisenet III, Q, and P models simply called WDR. Previous true WDR models marketed as "Advanced WDR."
Sony	View-DR Dynaview	N/A	View-DR included in all 6th and 7th generation cameras regardless of series (V, E, or X)
Vicon/IQeye	WDR	N/A	Model numbers containing "W", e.g., IQ862**WE**
Vivotek	WDR Pro WDR Pro II	WDR Enhanced	See notes below.

For further details regarding this list and additional manufacturer specifics, please read our WDR Manufacturer Cheat Sheet and Camera Trackingupdate.

Camera Finder Tracking

We have also verified and compiled which cameras are true WDR and have added them to our Camera Finder (under imaging, select WDR = Yes).

Here are lists from our Camera Finder:

- 900+ cameras with true multi-exposure WDR
- 480+ dome cameras with true multi-exposure WDR
- 150+ cameras, under $400, with true multi-exposure WDR

HD Analog vs IP

For years, HD resolution was IP camera's greatest advantage.

However, starting in 2013, analog cameras with HD resolution started shipping and, now, just a few years later, HD analog has become a significant force in the video surveillance industry.

We examine AHD, CVI and TVI, including their most recent 3rd generation advances, compared to IP cameras.

2017 Advances - Compatibility and Resolution

The two most important HD analog advances going into 2017 are:

- Advancing interoperability: New generation HD analog recorders increasingly support all 3 variants (AHD, CVI, TVI) as newer chipsets provide this functionality built-in.
- Increasing resolution: While 1080p was the max for HD analog for the first years, 3MP and 4MP are now available, with up to 4K scheduled for release in 2017.
- Power Over Coax: The newest generation HD analog chipsets support running video and power over a single coax. We expect this to be in limited production release during 2017.

HD Analog Variants

All HD analog cameras support coaxial video transmission, typically to 500m of RG-59 cable and, at least a max resolution of 1080p, with new generations adding support of 3MP, 4MP, 5MP and 4K..

AHD

Analog High Definition (AHD) was developed by Korean chip manufacturer NextChip, originally specified with a max resolution of 720p, but increased to 1080p in its second version. The most well known brands supporting AHD include Samsung and Digital Watchdog, however, most AHD product is lower cost offerings from less established brands. AHD is best known for having the lowest cost HD analog cameras, often with prices at $10 per camera or less, though typically from vendors with no brand and minimum support.

CVI

HD-CVI (Composite Video Interface) was developed by Dahua and was originally exclusive to them. However, it has since been licensed to others via HDcctv 2.0 specifications, with some non-Dahua CVI product becoming available.

The major sources for CVI product in North America are Dahua and their OEMs FLIR, Honeywell, and Q-See, which all OEM various camera and recorder models, as well as a handful of lesser known brands.

At first, Dahua kept CVI proprietary, not only blocking support for AHD and TVI, but also threatening to sue them. Then, in 2016, Dahua changed their position with their 3.0 release, opening up support for AHD and TVI on their HD analog recorders.

4MP CVI cameras are available. Our tests show that they delivered similar image quality to their IP equivalents in full light, low light, and WDR scenes. Additionally, long cable performance has been improved, with cable lengths up to ~1,500' (~457m) showing little to no visible degradation. The main limitation is framerate, as current cameras are capable of only 15 FPS at max resolution (2688x1520), not full frame rate found in many 4MP IP cameras. However, in most applications this is not likely an issue, as average framerates are typically under 15 FPS.

TVI

Chip manufacturer Techpoint developed HD-TVI (Transport Video Interface), which has been adopted by several manufacturers, the largest of which is Hikvision, Dahua's main direct competitor. Others, such as KT&C, CNB, Speco, and smaller manufacturers have adopted it, as well. Additionally, Hikvision has numerous brands OEMing them, and HD-TVI is increasingly becoming available from them.

In TVI's newest 3.0 release, TVI has added recorder support for AHD and CVI cameras as well as 3MP cameras. 5MP TVI cameras are scheduled for Q2 2017 release from Hikvision.

Encoding/Transmission

All surveillance is encoded and compressed (e.g., H.264 or H.265). The key difference amongst these offerings is where the compression is done.

In IP cameras, compression is performed inside the camera. In others, compression is performed on the server side (e.g., recorder, encoder, video server, etc.).

This is a major driver in performance differences.

Advantages of Encoding In the Camera

- Bandwidth is essentially 'unlimited'. Because the video is compressed in the camera, the output can be 3MP, 5MP, 10MP, 20MP or more and can easily 'fit' inside standard networking infrastructure (e.g., 100Mb Ethernet).
- Advanced features can easily be added as the same computer that compresses the video, can compress audio, dewarp fisheye panoramics, support multiple imagers, perform video analytics, etc.
- No specialized hardware is needed on the receiving side. Since the video is compressed typically in standards-based H.264 (or H.265), all the VMS / client / recorder needs is open source software to decode / display. Connecting to the camera is driven by the IP camera manufacturer's API or, increasingly, ONVIF. By contrast, when encoding on the server side, specialized hardware always needs to be provided, which limits backwards compatibility and recorder support.

Disadvantages of Encoding In the Camera

Cost increases since every camera needs to have the processing power / hardware to encode instead of just adding it to a single or a few recorders / encoders which then handle encoding for multiple cameras. This is one reason why HD analog cameras are generally notably less expensive than IP cameras.

Along with this increase in cost, since every IP camera is a computer, complexity is higher. With the benefits of cameras being a computer come the downsides of computers - increased complexity, potential for software incompatibilities, need for integration, etc. These issues are simply not present in analog (HD or SD) cameras.

Resolution

Entering 2017, both HD analog and IP cameras regularly delivered Full HD 1080p resolution. This was a big jump from SD analog's previous limitation of VGA / ~0.3 MP. However, 4K / 8.3MP IP cameras are now commonplace.

During 2016, HD analog suppliers talked about 4K resolution, but shipped no 4K cameras. However, several 3MP and 4MP models did ship, though at reduced frame rates. Any 4K product that ships in 2017 are likely to output significantly reduced frame rates (7.5fps is what has been cited), making it not true 4K but still 8.3MP.

For now, IP has a lead in resolution at full frame rate, though increasingly only for niche applications that really need a single camera with very high resolution. Especially for indoor usage and homes / smaller business, 1080p HD is often enough.

Advanced Features

IP has a large lead in advanced features, given that IP cameras are basically computers with cameras attached and can therefore include all sorts of advanced processing (audio, fisheye dewarping, support multiple imagers, on board video analytics, etc.).

CVI and TVI offer some features of IP such as I/O and two-way audio, as well as PTZ control and configuration up the coax. AHD also offers PTZ/OSD control up the coax, while SDI does not.

Recorder Compatibility

IP cameras can be made compatible with any recorder or client by adding software, whether it is proprietary integration or the use of "standards" like ONVIF.

HD analog camera types require specialized receiver / encoding hardware which cannot simply be added to older analog DVRs. New recorders (or encoders) must be purchased along with cameras. Increasingly, AHD, CVI, and TVI recorders are able to mix and match inputs of NTSC / PAL analog with their own HD analog type.

HD analog types have historically not been compatible with each other. For example, if you connect a TVI camera to an AHD or CVI only recorder, you will get no usable video (likewise, with the other way around). We demonstrate this below:

Note: Click here to watch the AHD video on IPVM

However, HD analog compatibility is likely to increase significantly in 2017 as the major suppliers have now begun to support other HD analog variants.

Coax Compatibility

IP cameras require Ethernet over Coax (EoC) Shootout to run over legacy coax. These typically add $100 to $400 per camera.

All other camera types are designed to run over legacy installed coaxial cable, though the distance limitations claimed vary. CVI and TVI both claim "over 1,500'" using RG-59 (and with our tests validating that up to 1000'). In our tests, early TVI releases had some issues using long cable runs, but these have since been remedied in HD-TVI 2.0 chip releases.

Power Over Coax

One of the big installation benefits of IP cameras has been Power over Ethernet (PoE). Instead of using one cable for video and another for power, a single cable can be used and power transformers can be eliminated. With Power over Coax, siamese cable can be replaced with RG59 alone, external power supplies can be dropped and some labor time can be reduced. We estimatecost reduction for new installs of ~$20 per camera (based on a 16 camera, 2000' foot of cable job, 'power over coax' will likely reduce costs ~$300).

Both Dahua and Hikvision say they are working on Power over Coax support for HD analog, and plan to release products in 2017. It remains to be determined how much much more the Power over Coax products might cost.

Vendor Support

IP has massive vendor support, both in terms of number of manufacturers and range of form factors available.

HD analog variants have growing support, though limited by its relatively short availability.

- AHD has historically had the least vendor support, with only low cost or no name brands utilizing it. However, manufacturers such as Samsung and Digital Watchdog offer AHD product, as well.
- CVI has broad support amongst Dahua (the founder of CVI) and their OEMs (in North America, most notably FLIR, Honeywell and Q-See).
- TVI is supported by a number of companies, but by far the largest is Hikvision (and their OEMs).

Western / Japanese Big Brands Not Supporting HD Analog

Most of the biggest brands in the world are not supporting HD analog. For example, Avigilon, Axis, Canon, Bosch, Panasonic, Pelco, Sony are all not supporting HD analog. Without their support, this limits marketing efforts behind HD analog and, related, validation of the technology as many look to these larger manufacturers for cues about what technology to use.

By contrast, HD analog has been led by Chinese and Korean manufacturers, who are strongest in the low to mid tiers of the market and do not carry as much recognition with average buyers as the Western and Japanese brands.

Model Availability

In 2017, a key limitation for HD analog, despite their very low price, is the relatively small number of high end options, with most AHD/TVI/CVI cameras and recorders low-end, intended for more cost conscious installs.

Cameras supporting advanced features such as true multi-exposure WDR, super low light, etc., are less common than IP cameras with these features. However, newer generation Dahua and Hikvision HD analog products increasingly support those.

On the other hand, many premium features and niche form factors are still far more widely available in IP than HD analog models as most vendors look at HD analog as primarily focused at the budget market where lower cost and less sophisticated options are expected.

Cost

HD analog has significant cost advantages over IP, both on the camera and recorder side. 30-50% lower cost for HD analog vs equivalent spec'd IP systems is commonplace.

IP camera costs are now moderately high, failing from extreme high prices in the 2000s and helped by lower cost Asian vendor expansion in the past few years. In 2017, a 'cheap' IP camera from a name brand bought though a local distributor runs in the $80 - $100 range.

AHD, CVI, and TVI cameras and recorders are extremely inexpensive, even compared to similar entry level IP cameras, with prices starting at ~$30 USD for 1080p cameras and recorders starting below $100.

Install Simplicity

Connecting cameras to recorders is more difficult with IP than any of the non analog versions. With IP, each camera needs an IP address, the network needs to be set up directly, the tech needs to know IP to connect the cameras, and the VMS/recorder must support them, either via direct driver or ONVIF.

Experienced IT surveillance professionals will not find this a major problem. However, many non IT and traditional low voltage techs will find HD analog's "plug and play" installation much easier.

HD Analog Closing the Gap on IP, Still Lower End Though

Increasingly availability, better industry recognition and improving features (higher resolution, better compatibility) are helping drive HD analog adoption.

On the other hand, HD analog is still firmly focused / dominant in the low end of the market. The lack of camera options and the refusal of the biggest Western and Japanese brands to support HD analog all hold it back.

Finally, there will remain functional advantages for IP cameras for higher end applications indefinitely, as edge storage, on-board analytics, multi-imagers with a single output, etc. will continue to be IP only. Additionally, many larger applications have standardized on UTP cabling. While HD analog cameras can be used with added baluns, IP cameras remain a more natural / direct fit for such cabling architectures.

[Note: This tutorial was originally written in July 2014 but was substantially re-written in August 2015, May 2016, and December 2016 to reflect improvements in HD Analog options in the past 2 years. In 2016, we dropped coverage of SDI and 960H due to declining availability / competitiveness.]

LENSES

Field of View (FoV)

Field of View, or FoV, is deceptively complex. At its most basic, it is simply what the camera can 'see' and is, therefore, visually self-evident. However, when analyzing images, comparing cameras or projecting quality, FoV subtleties can have a big impact. In this tutorial, we examine:

- Specifying Field of View
- Calculating Field of View
- Measuring / Comparing the Right Field of View
- Field of View vs Distance from Camera
- Field of View vs Lens Length

Specifying Field of View - HFoV and AoV

The two most common terms to specify / communicate Field of View are HFoV and AoV. Both are specified in degrees (60°, 120°, 180°, etc.). Both refer to the horizontal coverage range. For purposes of video surveillance, they can be used interchangeably.

Additionally, there are two other terms but they are much less commonly used. Those are vertical FoV (VFoV) and diagonal FoV. These may be listed on lens specifications but should not be mixed with HFoV or AoV as they measure different properties.

Seeing Field of Views

To see what different FoVs with varying lens lengths 'look' like, use the IPVM camera calculator. Below contrasts 180°, 120° and 60° AoVs:

Calculating Field of View

The first step is calculating a camera's FoV in degrees (°). The key factor here is lens length. The longer the lens, the narrower the FoV / area the camera can see. One specifies the appropriate lens length for the FoV they desire.

For more, review our tutorial on lens length that explains this in depth. Once you understand lens length, you can calculate FoV using the IPVM calculator.

However, knowing your FoV alone is insufficient as you need to appreciate how it relates to other cameras and objects within the FoV.

FoV vs Area

FoV is typically expressed by using the horizontal component only, e.g., the HFoV is 60°

However, changes in the HFoV impacts the total area captured by 2x the amount. For instance, cut the HFoV in half, and the area is reduced 4x. Similarly, doubling the HFoV and the area increases by 400%.

Below compare 120° to 60° HFoV. Notice that the area changes by 4x:

Measuring / Comparing FoV

The most common FoV problem we see is comparing cameras with different FoVs. Even if the cameras are identical and are looking at the same object, the camera with the wider FoV will always 'look' worse than the other one. Likewise, if you want to unfairly make a preferred camera look better, show that camera in a narrower FoV (see the Arecont ad as an example).

The comparison below shows two different FoVs being swapped:

It does not need to be vendor manipulation. Often, people will simply test two different cameras at two different times and not verify that the FoV of the cameras are the same.

A wider FoV can result from a shorter lens length or a subject further from the camera. We now examine both.

Field Of View vs. Target Distance

The further a subject is from a camera, the wider the scene is at that point. The angle of view is constant, but the width in feet / meters increases as one moves farther from the camera.

This is essentially a mathematical law. You can see it physically in the image below with:

Here's an overhead shot that demonstrates the same principle / scene:

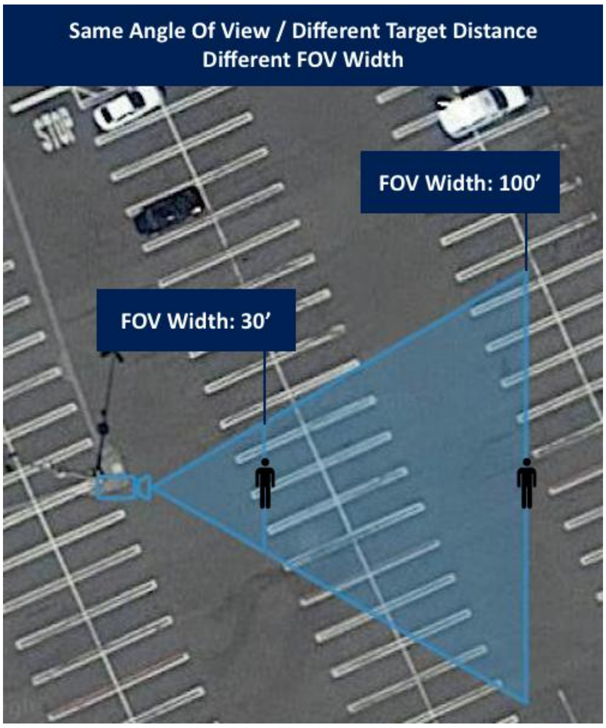

Because of this, three important lessons must be remembered:

- There is no single FoV width. While people often imply this, it is a dangerous assumption because you must know where in the FoV a target is. Saying simply the "FoV width is X" is misleading.
- To define the FoV width, you must know the subject's distance from the camera. To properly define this, say "The FoV width is X at Y from the camera."
- Make sure you understand the farthest point away from the camera one wants to monitor. By definition, this will be the widest FoV and the most likely to have problems delivering sufficient details. To determine if the FoV is too wide, use PPF as a guideline.

Field of View vs Lens Length

Not only does a shorter lens length increase the FoV, it also accelerates how much the FoV increases as the subject's distance from the camera increases. Let's compare a fisheye/panoramic FoV with a super telephoto one to show this principle in action:

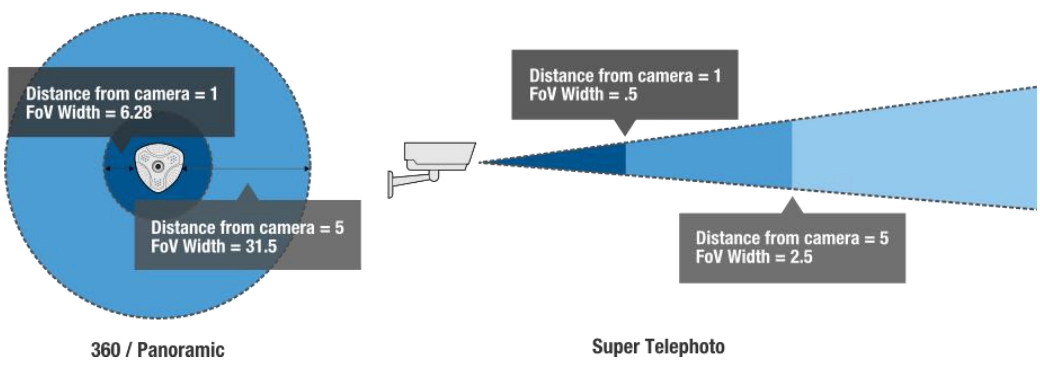

For a 360 / fisheye camera, every 1 foot a subject is further from the camera, the FoV increases ~6 feet (as the FoV is the perimeter of a circle). In just 10 feet, the FoV is more than 60 feet wide.

By contrast, for a super telephoto lens, every 1 foot a subject is further from the camera, the FoV may only increase one half or quarter of a foot. Because of this, a subject who is 10 feet behind another will have a very similar FoV with a super telephoto lens and likely similar detail captured. This is exact opposite of the fisheye / 360 camera.

Lessons to take away:

- Be very careful with fisheye or super wide angle lenses as the FoV will expand rapidly, capturing limited details across their wide range.
- While super telephoto lenses deliver a FoV width that expands slowly, shallow depth of field (DOF) becomes a big risk. This lens type might deliver the necessary PPFs but still be out of focus. See our Depth of Field tutorial for details.

That noted, most cameras with 'normal' or common lenses (in the 3mm to 8mm range) will have their FoV widen ~1 foot for every 1 foot increase in subject distance from the camera. However, at the extremes, the difference in FoV expansion can vary dramatically (i.e., fisheye and super telephoto).

Field of View vs Imager Size

Imager size also impacts FoV. However, in practice, for surveillance, this is typically not a major factor as variances among imager sizes tend to be modest.

It is worth remembering, the larger the imager, the wider the FoV will be, everything else being equal. This is covered more in our Surveillance Camera Imager Tutorial and can be experimented with in our Camera Calculator.

Lens Focal Length

3mm, 6mm, 2.8 - 9mm, 5 - 50mm, etc.

Camera specifications often list lens lengths but what do they mean?

These metrics are important in determining the correct camera coverage and Field of View (FoV). In this tutorial, we look at:

- What lens focal length impacts
- The relationship of imager size
- Common lens length ranges and their uses
- Limitations of using focal lens length
- Manufacturer variances between focal length and AoV
- Picking the right lens length

Focal Lens Length vs FoV / AoV

Focal lens length is the physical distance between the lens and the sensor / imager. It is important in surveillance because it is the primary driver of the FoV covered / Angle of View (AoV) of a scene. See: Field of View (FoV) Tutorial

Lens Length Impact

The longer the lens, the narrower the Field of View / area captured; the shorter the lens, the wider. This is a physical principle deriving from the distance between the lens and the sensor, as the image below shows:

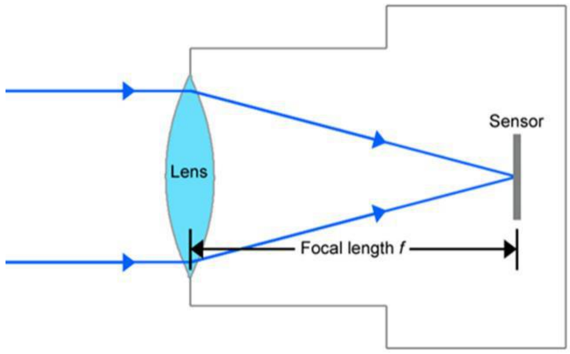

The longer the lens is, the farther it is away from the imager, resulting in a narrower FoV.

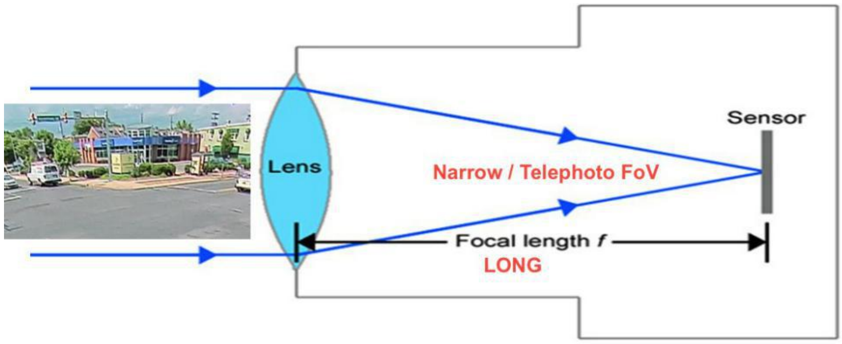

Likewise, if the lens is short, it will be close to the imager and capture a wider area, like so:

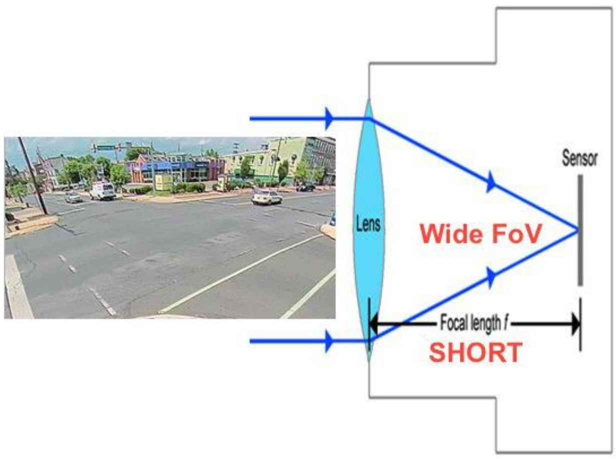

Impact of Imager Size

Though the focal lens length is the primary driver of FoV in surveillance, the size of imager used (i.e., ¼”,1/3”,1/2”, etc.) also impacts this. Everything else equal, the larger the imager, a larger FoV / area can be captured. This is not a huge factor as most professional surveillance cameras are 1/3” (or close – at 1/3.2” or 1/2.7”). However, there can be modest differences, even if the focal length is the same, if the imager size of the camera is different. For more, see: Surveillance Camera Imager Tutorial

Rule of Thumb - Focal Length vs AoV

Using the most common 1/3” imager, here is the AoV / horizontal FoV in degrees for common lens length:

FOCAL LENGTH	APPROXIMATE ANGLE OF VIEW
3mm	77°
6mm	44°
12mm	23°
20mm	13.7°
50mm	5.5°

Try these calculations on your own using the IPVM calculator.

Common Lens Length Ranges

In practice, surveillance cameras typically fall in one of 4 lens length ranges:

- Fisheye – Under 2mm provides super wide angle views but comes with distortion, ergo the name 'fisheye'. These lens typically require dewarping software to provide a flat / corrected image. While it can provide panoramic images, implementation constraints remain. See our Guide to selecting Panoramic / 360 fisheye cameras.

- 'Normal' – 3mm – 10mm: The stock lenses of most cameras fall in this range (roughly 30 to 80 degrees FoV).
- Telephoto – 10mm – 80mm: While most cameras do no come with lenses in this range, lots of 3rd party C/CS lenses can be purchased and connected to cameras (typically box form factor)
- Super Telephoto - 100mm+: Lenses in this range are specialist ones used to see / monitor objects very far away (typically 1km+ plus often 5kms or more)

Warning: Manufacturer Focal Length / AoV Variances

Camera manufacturers typically specify AoV / FoV primarily by listing a focal length. For example, manufacturers may list 3mm prominently to show that a camera is wide angle or highlight 50mm to show a camera is telephoto.

However, camera manufacturer's focal lengths do not always map up to theoretical / calculations for AoV , especially with shorter focal lengths. Sometimes it is because of how much of the imager or lens is used or because the construction of the lens varies from the theoretical calculations. Regardless of the cause, the solution is to always spec cameras / lenses based on AoV, not focal length. IPVM makes this easy by allowing selection of camera models, with verified AoVs, in our calculator.

Lens Limitations

The longer the lens length, the more likely three limitations will arise:

- Degraded image quality: Long lenses, especially varifocal ones common in surveillance, frequently have degraded image quality relative to what pixel density / ppf metrics would estimate. If you are using PPF / PPM to plan your deployments, you may need to compensate with extra pixel density. See: A Major Flaw in Long Lenses and PTZs Found
- Degraded low light performance: Very long lenses (as well as very short ones) typically pass significantly less light than 'normal' lenses in the 3mm to 10mm range. Specifically, everything else being equal, the longer the lens the higher the F number (Review our F stop tutorial for details).
- Depth of Field Issues: Surveillance cameras very rarely have depth of field issues where objects appear out of focus. This is because they usually use relatively short lenses and have subjects that are more than 5 feet away from the camera. However, very long lens will cause problems that are especially hard to resolve for night time imaging or with auto-iris lenses. Review our Depth of Field tutorial.

Picking the Right Lens Length

The shorter the lens length, the more the camera can capture. Unfortunately, the wider the capture area, the less detailed is any individual object captured. This is why people often are stunned when they look at a megapixel fisheye camera. Since it is megapixel, the quality should be better but since the FoV is so huge, all the objects look 'fuzzy' even if they are close to the camera.

The easy rule of thumb is to make the FoV no wider than the area of interest – i.e., just what you care to capture. However, this can be tough as people want to capture a lot of things (i.e., I want the cars in the parking lot plus the license plates of the cars entering it).

To help solve this, Pixels Per Foot (PPF) is the metric used to help determine how wide one can go without making the details in the area captured so fuzzy that they have no value. Please read our PPF Guide to understand this. Experiment with our Camera Calculator to determine the optimum tradeoff between FoV and PPF for your specific applications.

Varifocal Vs Fixed Lens Usage

Varifocal camera lenses provide more flexibility at installation time, but come at a higher cost. Fixed lens cameras reduce hardware costs, but can make getting the right Field of View more difficult or costly.

How often do integrators choose one over the other? In our 2014 Varifocal vs Fixed Lenssurvey varifocal lenses won out, but since that time the market has seen significant technology and price shifts. Has this affected integrator preferences? We look at this topic again based on updated survey statistics to see what integrators now prefer.

Lens Usage Breakdown

Usage of varifocal vs fixed lenses remained constant from our 2014 Varifocal vs Fixed Lenssurvey, with integrators vastly preferring varifocal lenses:

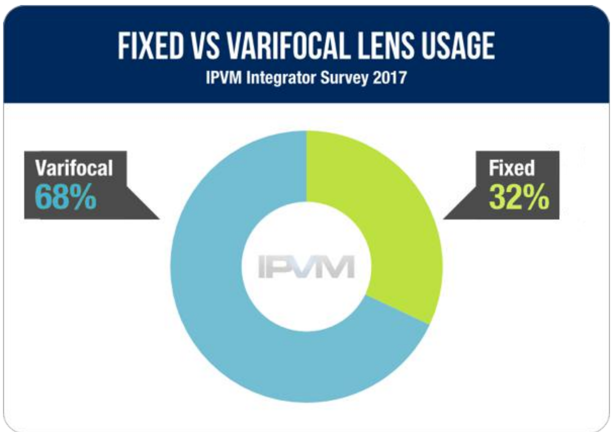

Decision Factors

Price was a big driver for choosing varifocal vs fixed lenses, but many times the final cost for the camera was a factor of time spent on system design +

hardware cost + installation time. This sometimes made a more expensive varifocal camera turn out to be cheaper overall than a fixed-lens alternative.

Delivering a quality image was a major driver in choosing which lens type to use, however integrators had two approaches to this problem.

Flexibility Adjustment

Some integrators preferred varifocal lenses because they afford the flexibility to fine-tune the field of view:

- "I try to use 100% varifocal unless cost is a factor. I like to have the ability to adjust the FOV during the install. Gives the installer and customer greater flexibility."
- "60% varifocal and 40% fixed. Varifocals are used when there are requirements to meet regarding pixels per meter etc. Fixed lenses are used where price is key and it's enough with an overview."
- "Most of our cameras use varifocal, unless we can determine a need that only a fixed focal length will handle. It's just easier to get the angle of view the client wants with a varifocal."
- "Varifocal 80%. Vari is flexible and can adjust, easy to use the auto focus feature unless re-aiming is required."
- "Almost always varifocal. The cost isn't much higher and the incremental cost gives us flexibility on the install."

Higher Resolution Reduces Need For Varifocal

Other integrators noted that increasing camera resolutions (coupled with rapidly decreasing prices) allowed them to simply use higher resolution fixed-lens cameras, giving them the ability to provide detail to the customer, without the requirement to manually adjust the FoV:

- "95% fixed, little need for zoom when the cameras are 4MP."
- "Used to be all varifocal now with higher res fixed lens cameras reduces need for varifocal especially for indoor applications."
- "If the customer is clear on what he wants to look at and if we estimate properly then we would choose a fixed lens especially if it's in 4MP."
- "50-50. VF are nice as you get more flexibility in camera placement but with an added cost and added size to a minidome. With better resolutions we find fixed lenses like 2.8mm and 4mm are good because they get a wide FOV but you can still digitally zoom after the fact."

Adapts To Changing Requirements

Requirements sometimes change over time, particularly in retail settings, and for this reason integrators used varifocal lenses to help customers get the most value out of the camera over a longer time period.

- "50% in Retail areas or areas where scene changes and a different view may be required."
- "I believe all of our cameras are varifocal. Sometimes we will move a camera or repurpose it, thus needing to adjust the FoV, etc. that you can't do with a fixed lens."
- "Majority are going to be varifocal. Easier over time to adjust FOV and change the image to get another or new angle."
- "Mostly varifocal, gives us more flexibility at the time of installation and more possibilities to adapt to the clients wishes."

- "All of them are varifocal lenses for the use of autofocus and in retail they change interior layout of the store at marketing request."

Fixed Lenses Lower Cost

Cost was a top reason for choosing fixed-lens cameras:

- "Most of them are fixed focal lenses. Mostly because price effectiveness."
- "Fixed lenses are used where price is key and it's enough with an overview."
- "80% of our installs would use vari-focal. If the opportunity is largely being driven by price, and not design, we will then go to fixed lenses."
- "Approximately 70-80% fixed lens due to low-cost. Most applications that we sell to do not really require vari-focal lens."
- "Fixed lens cuts down on the install times"
- "Only use fixed on microdomes and down market cheap jobs, which are rare for us."
- "Varifocal lenses are deployed 80% of the time. We will use a fixed mini dome for some customers who like to save money. These jobs do take more time to engineer to get the correct lens for the job."
- "50/50 - usually design with varifocal lens to get specific ppf at certain ranges. Will use fixed focal when price is an issue or concern."

Labor Cost Decrease With Varifocal

Hardware cost is only part of the total job cost, installation labor can often be in excess of $100/hr, making a camera that has a higher price, but greatly reduces installation time or complexity potentially cheaper overall for the customer. Several integrators stated variations of this theme:

- "90%. Varifocal typically comes with remote focus, which is severely shortening install and setup times.
- "We prefer motorized zoom lenses. Easy to install and commission. Only in low price project we offer fixed focal lenses."
- "70% varifocal. Cameras with motorized lens varifocal keep dropping in price while labor rates continuously rise. We still used fixed cameras for small rooms like offices, but varifocal work well in hallways and big retail areas. We also find that it is hard to source fixed lens cameras with focal lengths past 8mm while varifocals generally go to 12mm."
- "50% fixed lens when we are getting smashed on price, we still prefer a varifocal and particularly an autofocus VF as it saves commissioning time."
- "We used motorized zoom domes for ease of focusing to save on labor."
- "80% varifocal Ease installation and better able to suit customer wishes"
- "Varifocal when the price allows, more installation flexibility"
- "Almost always varifocal. The cost isn't much higher and the incremental cost gives us flexibility on the install."

Fixed Lenses Require More Design Work

Extra up front design work may be required with fixed focal length lenses, since the cannot be easily or cost-effectively changed once installed:

- "40% at least would be varifocal - Ease of deployment, remote focus reduces TCO and we can customize zoom area on the job rather than having to have it perfect on the design."
- "I only use the Fix lens when I have exact drawing and I did the design and know what type of lens is needed for the area, so I select only the fix lenses cameras."
- "We will use a fixed mini dome for some customers who like to save money. These jobs do take more time to engineer to get the correct lens for the job."
- "70/30. Varifocal mostly when not sure of exactly what they want to see. More flexible. Fixed when we know for sure what the scene is and if price dictates."
- "Fixed focal length is preferred due to better f-stop but requires a bit more advanced planning so that the correct camera/lens/imager is ordered and installed."

Manufacturers Drive (Or, Eliminate) Choices

As race-to-the-bottom trends have reduced prices, manufacturers are often adding more features to cameras to try to differentiate themselves, or justify higher prices. Varifocal lenses, often with motorized zoom/focus, are a common component to this approach. Integrators noted that in many cases varifocal lenses are effectively the only choice available, even if not the most desirable:

- "The only fixed focal length cameras we use are Avigilon's micro domes. Even if they had a wider range of fixed lens cameras available, we wouldn't use them. It's not worth the hassle of having to explain to a customer that it'll cost them the price of a new lens if they want to zoom a camera in a little bit, and no matter the amount of CAD drawings and site demos/tests we do, they will always want to make minor adjustments after the install."

- "Most cams we sell have integrated lenses, not boxes. So it's rare that we specifically choose fixed vs varifocal. I prefer fixed bc picture quality and aperture are better, but most high-quality cams with integrated lenses are varifocal."
- "100% Varifocal. Not because we have anything against fixed, but most of the cameras we sell simply are varifocal.
- 75% fixed, 25% varifocal. Inexpensive cameras all have fixed. Our high end domes, bullets etc. are all varifocal. I don't typically need varifocal as we use them at their widest setting, but auto focus is sure nice to have and the ability to slightly adjust the image."
- "We use 90% varifocal. I'd rather see us use more fixed lenses, but a lot of times, the other, more advanced features we need are only embedded into the varifocal devices."
- "Vari-focal 90% just due to the fact that we sell Axis and prefer the P32 series for its cost/value. The M series cameras are used in certain circumstances, but more in special situations."

Lens Iris

Cameras, like humans, have irises. However, cameras have five types of iris options - fixed, manual, auto, P iris. In this tutorial, we explain the tradeoffs of each, how irises work and how they relate to shutter speed.

Controlling Light

Cameras control light received by either adjusting the width of the opening (iris) or the length of time the opening is open (shutter). When the width of the opening (iris) is adjusted, its F number changes, reflecting that more (or less) light enters.

Here's what a lens iris looks like mostly open vs mostly closed:

And here's a visual showing the process of adjusting:

Note: ***Click here to view the animation on IPVM***

If you let in too much light, the image will be washed out but too with too little light, the image will be dark. Since light levels change over time, the camera has to be able to adjust how much light enters to match the current conditions. Irises are one way to do this. There are multiple iris control methods, detailed below.

Auto Iris / P-Iris / I-CS

Motorized iris lenses come in three variants: Auto (DC or video), P-Iris (precise), and I-CS.

- DC auto iris lenses control the iris opening via motors built into the lens. DS auto iris lenses are most common, which use a standard cable to move the iris as instructed by the camera.
- P-Iris lenses use a specific software driver to more precisely control the iris, which is intended to provide increases in image quality over auto iris in varying lighting conditions. However, in our tests, P-Iris lenses offered no drastic advantages.
- I-CS lenses are new as of 2016, developed by Axis and Computar. These lenses are similar to P-Iris, but use an more complex intelligent lens which supplies more detailed information to the camera, with control of focus, zoom, distortion correction, and more, instead of just iris control. I-CS lenses are rare, with very limited camera support, and effectively similar to P-Iris in terms of iris control.

While both auto iris and P-iris pigtails and connectors look similar, they are not fully compatible. A DC auto iris lens may be used on a P-Iris camera (though without advanced precise iris control), but a P-Iris lens connector will not physically fit into a DC auto iris jack.

Fixed / Manual Iris

Though auto-iris is more common, there are two lens variants which cannot be controlled by the camera: fixed and manual.

Fixed iris, as the name implies cannot be adjusted at all. It is set at the factory, typically all the way open to its lowest F stop. Fixed iris lenses are common on small form factor cameras, such as minidomes and covert cameras, but uncommon in CS and other mount types.

Manual iris lenses allow the iris to be controlled by manually rotating a ring on the lens. Though the iris is controllable from all the way open to all the way closed, manufacturers generally recommend setting manual iris lenses either fully or 90%+ open to maintain low light performance. Manual iris lenses are uncommon but still used in some specialist applications.

What To Use?

Often, users are not given a choice in what lens type to use, as the most popular camera form factors (domes and bullets) typically do not use interchangeable lenses.

In box cameras, which generally do support interchangeable lenses, we have not found a major difference in image quality between iris types. See our P-Iris Lens Performance test for full details.

The reality is that IP cameras today typically adjust the shutter speed dynamically and automatically, allowing them to handle super strong sun light (fast shutter) to pitch darkness (slow shutter).

If your camera supports P iris, we suggest using it as it may have some limited benefits with few, if any, drawbacks. However, we do not recommend rejecting cameras which only support DS auto iris simply based on this feature.

Test your knowledge

Take this 6 question quiz now.

F-Stop

To understand low light surveillance, appreciating the importance and role of the f-stop metric is critical. In this tutorial, we explain:

- Why minidomes and other small form factor cameras typically have far worse f numbers
- How the range of a PTZ impacts its f-stop / low light performance
- How integrated IR affects f-stop
- Why adjusting f-stop to maximize depth of field is dangerous and unlikely to help much in surveillance
- How f-stop can be confusing and counterintuitive

Key Points

F-stop measures the relative amount of light that a lens can pass. Understanding the number's significance can be tricky as the bigger the F number (1, 2, 4, 8, etc.), the less amount of light will be pass (e.g., 8 is much worse than 4, 1.8 is worse than 1.2, etc.). Here's a visual illustration of lens openings vs F numbers:

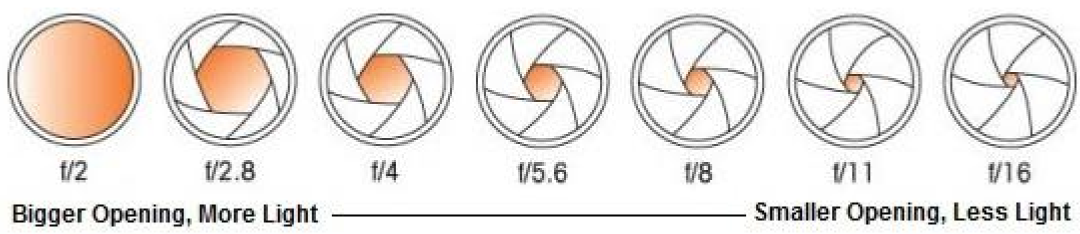

Technically, f stop contrasts the length of the lens and the diameter of the iris (i.e., L/D). Everything else equal, the longer the lens (say 10mm instead of 3mm), the less light passed and the higher the f number. Similarly, smaller lens diameter (such as found in miniature cameras) narrows the maximum iris, allowing less light to pass, resulting in higher f numbers. Both of these elements impact practical performance, especially in PTZs and minidomes.

What F Numbers Mean on Specifications

When lenses / cameras specify F numbers, they always provide the lowest F number that the lens supports when the iris is all the way open. Most lenses can adjust their iris opening. When narrowed, the F number will be higher. However, once the iris is opened all the way, the physical limitation is the width of the lens itself. It is this lowest F number, with the iris all the way open, that is most critical when assessing the low light capability of a lens.

Low Light Image Comparisons

The image comparisons below demonstrate the impact f-stop has on performance. Click the image for full size versions.

At ~2 lux, we can clearly see a gradual decrease in image brightness, with noticeable artifacting from F 5.6 and beyond.

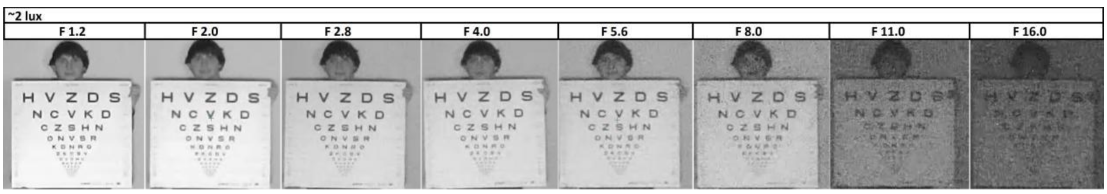

At about 0.05 lux, increasing the f-stop greatly decreases the visibility of our subject, even going from F 1.2 to F 2.0. At this light level, beyond F 5.6, the image is completely dark.

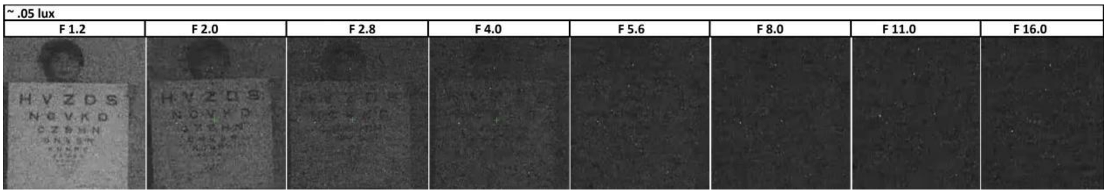

Notes:

- While increasing F numbers always reduces the amount of light passed, how quickly the camera's image darkens depends on its overall light sensitivity (i.e., a 'good' overall low light camera will resist getting visibly dark longer than a 'bad' one).
- This was done using a camera with a P iris lens that allows explicitly adjusting F numbers to specific levels. This cannot be done with 'regular' auto iris cameras and is difficult with typical manual iris lenses because they generally do not mark specific levels.

Integrated IR Impact

Because of the surge in integrated IR use in the past few years, F-stop has become somewhat less critical. Integrated IR adds significant amounts of light to the scene, compensating for higher F-stops.

For example, the image below shows a scene from our testing, with a higher F-stop (F/1.8) integrated IR camera easily outperforming a much lower F-stop non-IR super low light model:

Not that this is not to say integrated IR cameras will always outperform non-IR, as IR power, ambient light levels, super low light processing/larger imagers, and more also impact this performance gap.

Differences in F numbers

The difference between two F numbers can be deceiving, because they operate on a logarithmic scale instead of having a linear relationship.

For example:

- f/2.8 takes in 75% less light than f/1.4
- f/5.6 takes in only 1/16th the light as f/1.4

Users can use IPVM's F-stop calculator to see these differences themselves.

Recommendations

While it is important to be aware of F-stop and the impact it can have:

- Small differences in F numbers (e.g., f/1.2 vs f/1.4) generally do not impact low light performance greatly, contrary to what is commonly advocated. Automatic gain controland other camera side image processing offset small f number differences.
- Most professional MP box cameras have F-stop numbers of 1.2 or 1.4.
- Just because cameras have similar F-stop numbers, does NOT mean their low light performance will be the same. Differences in imaging processing make a material difference plus manufacturers often use slow shutter tricks.
- However, F stop numbers of 2.0 or greater often indicate poor low light performance.

- Minidomes and other small form factor cameras typically use small lens, with narrow lens diameters, that have high f numbers and typically worse low light performance (f/2.0 or greater). For instance, Axis full size domes (P33) have a f/1.2 while their minidomes (M30) have a f/2.8, meaning the minidomes take in ~1/5th the amount of the full sized ones.
- PTZs when zoomed to see far objects (i.e., long focal lengths) typically have high F numbers and poor low light performance since zooming out requires a long focal length (often 80 or 100mms) at a high F number (f/3.0 or greater is common). Beware many manufacturers do not list this but it is inherent when the diameter stays the same but the length of the lens is increased.

Depth of Field Impact

While high F numbers are bad for low light, they can be helpful for increasing the depth of field but not typically in surveillance (see our depth of field tutorial). The higher the F number, typically the greater the depth of field. The image below demonstrates depth of field at varying F-stops:

While this may be helpful for photography, it typically does not make a big difference in surveillance. Worse, it can destroy the camera's low light performance. For example, compared to f/1.4 common in surveillance cameras, f/8.0 takes in 32 times less light. If this was left on at night, in a dark environment, the camera would literally capture nothing beyond pure blackness. Indeed, no auto iris camera would allow f/8.0 to be used in such conditions as cameras are programmed to open the iris as fully as possible when dark (i.e., to the lens min f number).

Our recommendation is to ignore depth of field optimization unless you have strong photography skills and are willing to risk negative side effects for limited performance improvement (only during the day).

Test your knowledge

Take this 5 question quiz now.

Night Time / Low Light

Infrared (IR)

Infrared (IR) has become an increasing core component to video surveillance systems.

In particular, the expansion of integrated IR cameras that build IR illuminators within the camera has become common. On the positive side, this means that cameras can 'work' now even in the dark, even when cameras are cheap. On the negative side, many differences exist with integrated IR, including range achieved, width of coverage, overexposure, underexposure, response time to moving objects, to name a few.

We examine these elements in detail, including:

- The State of IR in 2018
- Infrared Basics
- 850nm Most Common
- 940nm Rare, But Increasing
- True Day / Night Cameras Required
- Integrated IR Usage Statistics
- IR LED Comparisons
- Warning: No Illumination Standards
- Range Specs Increasingly Accurate
- But Overstated Range Still An Issue
- Beware Low PPF At Max IR Range
- IR Coverage Patterns Often Poor
- Smart Vs. Adaptive IR
- Manufacturer Terminology
- IR Power vs. Gain/Exposure Adjustment

- Delays Returning to Full Exposure
- IR Impact On Bandwidth Spikes
- Integrated IR Cameras vs. External Illuminators Pros & Cons
- White Light Illuminators

2018 State of IR

There are several key trends in integrated IR in 2018:

- Integrated IR everywhere: Where integrated IR used to be a rare feature, it is now available in practically every manufacturer line, in low end and high end cameras and multiple form factors (see stats).
- Longer ranges: In the past, illumination ranges of ~30m/100' were considered "long." Now, 50m/165'+ is available in many fixed cameras, and 600', 700', and beyond common in PTZs.
- Laser IR: Most integrated IR cameras have used LEDs for illumination, which has limited range due to size/heat/power constraints. Now, some manufacturers are moving toward infrared lasers for illumination, which have longer range, with 1,500' not uncommon, mainly in PTZs. (Note: IPVM has not yet tested these models, but plans to do so in future tests)
- Invisible IR Increasing: While most cameras still use 850nm wavelength IR, the past year has seen an increase in "invisible" 940nm LEDs used, with consumer options such as Nest IQ and the Yi Home 2 adopting it, as well as multiple Vivotek and Axis models.

We expect these shifts to continue this year and beyond.

Infrared Basics

IR is light that cameras can 'see' but humans cannot. Humans see in the 390 to 750nm light spectrum, referred to as 'visible' light. However, surveillance cameras can generally 'see' visible light as well as above that range with IR illuminators typically at 850nm or 940nm, highlighted on the chart below.

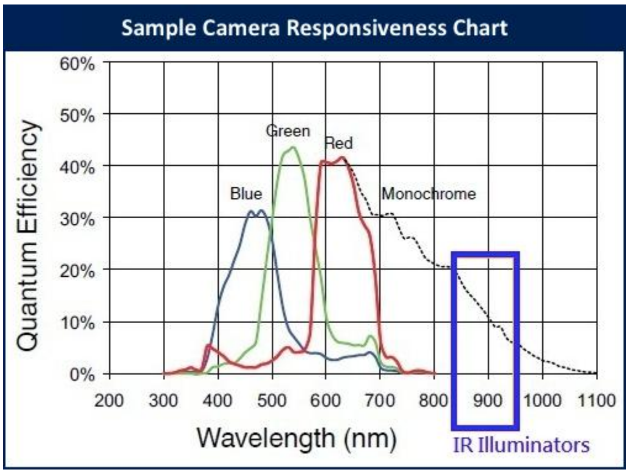

850nm Most Common

Most infrared illuminators used in surveillance use the visible 850nm wavelength. The LEDs used for these illuminators are easily visible to the human eye, emitting a reddish/pinkish glow, such as on the IR PTZ below. Note that humans perceive this as more red, while devices such as digital cameras (and cell phone cameras) see it closer to pink/purple shown here.

940nm Rare But Increasing (Most Commonly In Home)

940nm illuminators are much less visible to humans, with only a very faint glow visible when close to the camera, making them more useful when covert surveillance is required. However, 940nm is also weaker, traveling only about half the distance of 850nm while using the same power, and more expensive to manufacture.

Because of these factors, 940nm is used far less than 850nm. However, some manufacturers have begun to use 940nm illumination on specific models. For examples, see: Invisible IR Camera Tested (Vivotek), Axis Corner Mount Camera Examined, Yi 2 Intelligent Camera Test, and Nest Cam IQ Tested.

Note that some higher wavelength illuminators (1030nm and above) are available, which are completely invisible to the human eye. However, these wavelengths are above the spectrum many cameras can 'see', and costs are even higher than 940nm, making them very rarely used in production.

True Day Night Required

True day/night cameras **must** be used with infrared illumination, as color only models (though rare in 2018) contain a fixed IR cut filter, blocking all infrared light. The majority of professional cameras use IR cut filters which moves a filter in front of the camera's imager during the day, and removes it at night (see our Day / Night Camera Tutorial).

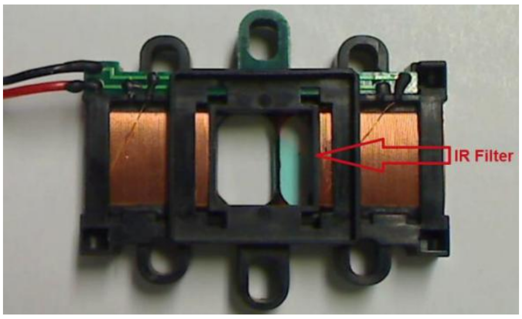

The cut filter is activated during the day to remove IR light, as it may cause color shift/white balance issues, such as the pinkish grey trees and desaturation seen in the image below:

At night, the filter is removed to allow IR illumination or ambient IR light to enter the imager, producing brighter images.

Integrated IR More Common Than External

When using IR, two fundamental options are available - cameras with IR LEDs integrated into their housings and standalone IR illuminators that are mounted near or next to cameras. In the past, very few IP cameras offered integrated IR LEDs so most were forced to add on illuminators if they wanted IR. Now, most cameras offer integrated IRs.

In terms of usage, IPVM statistics show nearly 2/3rds of cameras today include integrated IR:

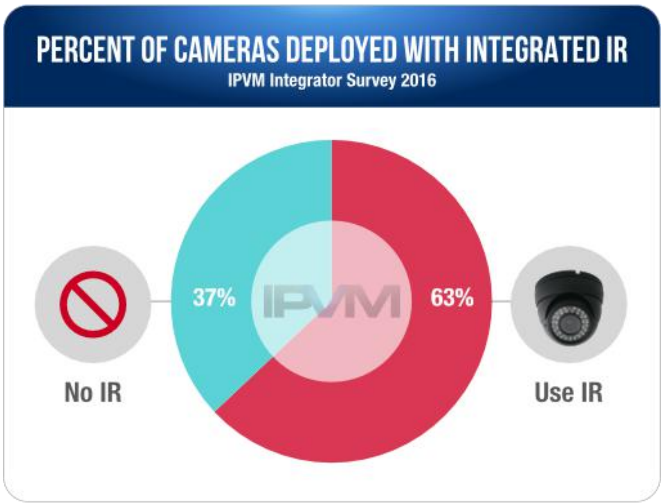

By contrast, IPVM statistics show very few cameras add external IR:

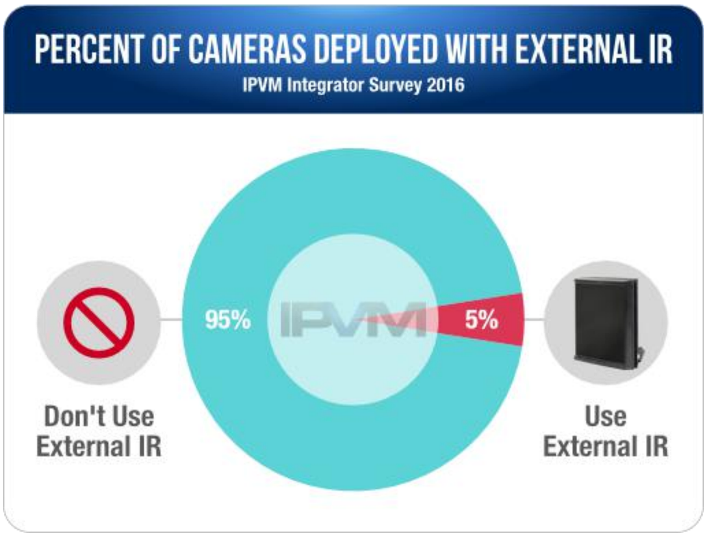

Because of this, this guide mainly focuses on integrated IR cameras. We review pros and cons of these options later in this guide.

IR LED Comparisons

Some camera manufacturers tout the number of IR LEDs as an indicator of total performance, that is 'My camera has 24 LEDs but my competitor only has 5, etc.', but in reality, looking at just the quantity of LEDs has little impact on real world performance, as there are many different sizes/styles of IR LED.

There are two broad categories of IR LEDs in surveillance, simply small and large:

- Smaller, lower power: Many integrated IR cameras and some low cost standalone IR illuminators use a higher quantity of smaller LEDs to add up to greater power. These LEDs are usually (but not always) less expensive, which is why they are typically included in low cost models.
- Larger, higher power: Some newer IR cameras and most IR illuminators use a smaller quantity of larger, more powerful LEDs.

Note that larger LEDs in the past were typically considered higher quality, and produced more even coverage and longer range. However, our tests have shown that this is not always the case, with some top performing IR cameras using small LEDs, and some using larger ones. Actual field testing is the only reliable way to determine IR illumination quality.

Illumination Issues

Below, we review key issues impacting integrated IR camera selection and performance, including:

- Lack of standards
- Range specs improving accuracy
- Low PPF at max range
- Poor IR coverage patterns common
- Overexposure issues
- Smart and adaptive IR
- Manufacturer terminology
- Long image adjustment times

No Illumination Standards

Since there are no concrete specifications for how IR is tested, manufacturers may simply quote the range at which any IR illumination reaches the camera, instead of usable IR. This means that the amount of light at the specified maximum range may be only enough to get only a rough outline of a person, at best. While you can certainly measure the amount of IR illumination, what manufacturers determine is the minimum accepted amount varies widely.

Range Specs Increasingly Accurate

A few years ago, when integrated IR models were far less common, manufacturers frequently overstated their camera's IR range capabilities, often by as much as 50-60%. However, as IR has become commonplace, with nearly every manufacturer offering multiple models in low cost and high end lines, specs have gotten more accurate.

For example, in our test of the Axis M2025, we found that its actual delivered IR range was 17-18m, exceeding its 15m spec.

That being said, overstated ranges are still not uncommon in 2018. For example, the camera below has a specified IR range of 50m/~165', but an actual range closer to ~100', with the subject still only dimly illuminated.

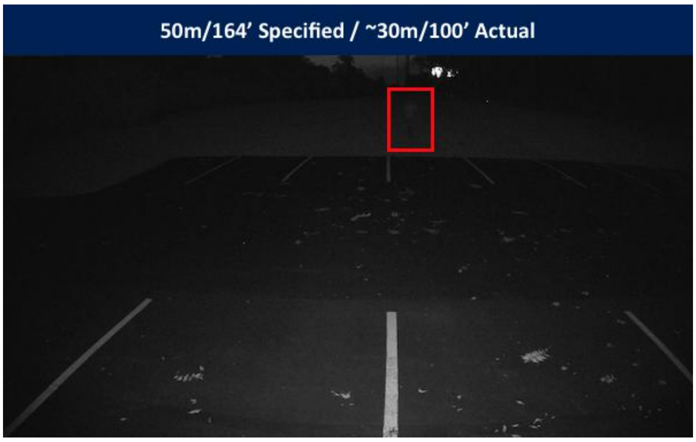

Beware Low PPF At Max Range

Even if a camera reaches its maximum specified range, users should beware of PPF levels at this distance. In our tests, 1080p models often deliver 10-12 PPF or less at max range. With PPF this low, detection is difficult due to the relative size of subjects, increased digital noise, decreased contrast, etc.

This issue is compounded as IR ranges have continued to increase, as have angle of view specs, especially on lower cost fixed lens models which now frequently ship with <100° lenses. Users should beware of this issue and consider other (narrower) lens options where possible.

As an example, the 1080p camera below has a maximum specified IR range of 30m, or about 100', and an angle of view of 102°. At this distance, PPF is only ~8, making the subject difficult to spot.

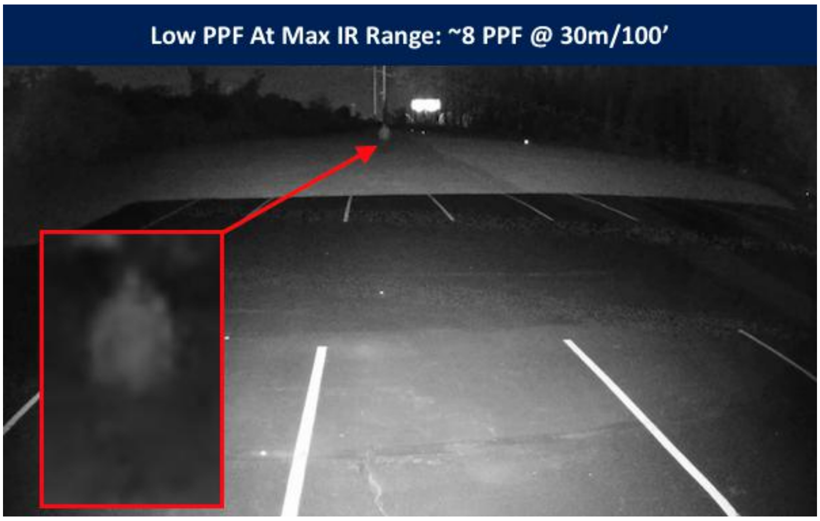

4MP and 4K models improve upon these PPF levels, but perform notably worse in low light, making images dimmer. For example, this 4K model delivers ~16 PPF at max range, but detection is even more difficult as the scene is dimmer and noisier.

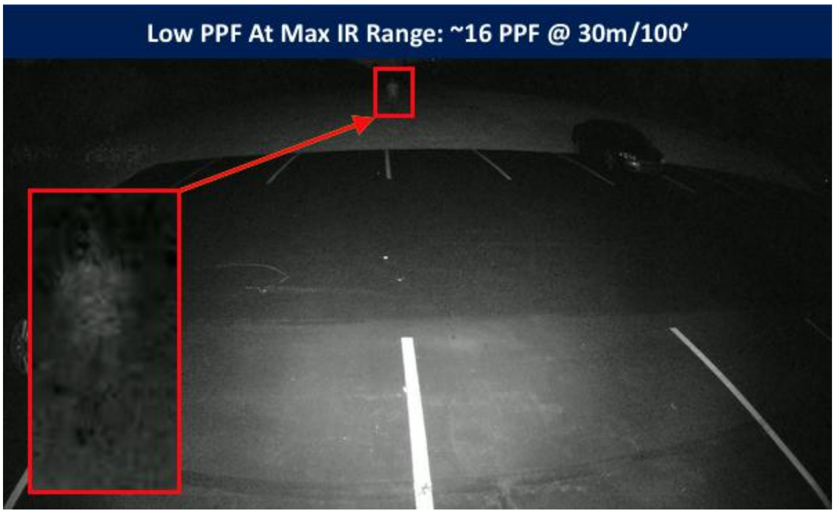

IR Coverage Patterns Often Poor

Although range specs have improved, IR coverage patterns have not in many cases, with many cameras not covering their full field of view, overexposing subjects in the center, or both. Coverage pattern is the number one issue seen in integrated IR models in our tests and biggest risk to integrated IR camera users.

For example, the camera below shows a bright hotspot in the center, obscuring subject details. By contrast, the edge of the FOV is very dim, with increased visible noise, making it impossible to discern subject facial features.

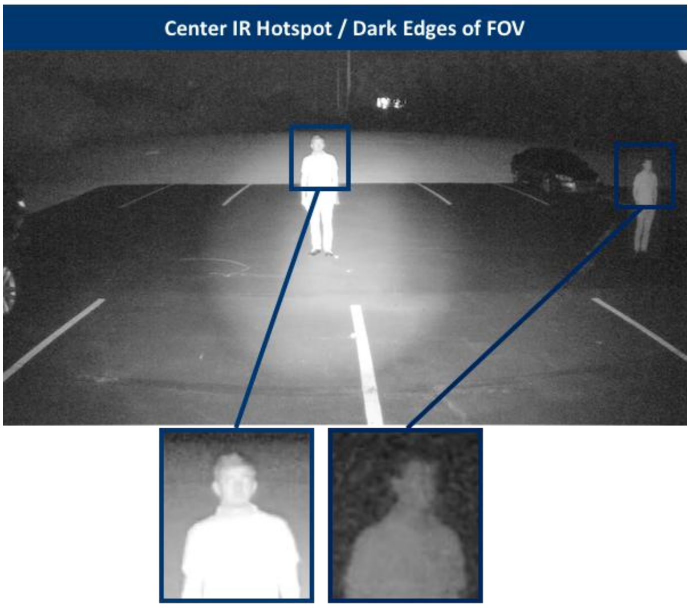

Smart vs. Adaptive IR

Some integrated IR camera models have introduced "adaptive" or "smart" features intended to better fit the scene and/or reduce overexposure. For the purposes of clarity in this guide and our reporting, we define these features as follows:

- Smart IR: Automatically adjusts IR power depending on objects in the scene, dimming to prevent overexposure of near objects and returning to full power for those further away. The video below shows the effects of this on video, with the scene dimming to compensate for the subject in the near field.

  Note: Click here to watch the video on IPVM

- Adaptive IR: The camera adjusts IR LED power and/or uses multiple sets of LEDs in order to adjust illumination angle to fit its field of view and more evenly cover the full scene. This is common in both fixed cameras and PTZs, shown below:

Note: Click here to watch the video on IPVM

Manufacturer Terminology

Like other IR specs, there is no standard as to what these terms mean and users may be confused by varying manufacturer terminology. For example, Avigilon says of their "Adaptive IR":

> With unique adaptive IR technology, this camera provides both wide and narrow illumination, enabling consistent illumination in complete darkness to enhance image quality regardless of scene conditions.

While Axis calls the IR illumination features of their high end models "OptimizedIR":

> When the field of view is adjusted at the installation of an Axis camera with remote zoom and OptmizedIR, the angle of illumination automatically adapts to the zoom level. The illumination angle follows the camera’s zoom movements to always provide the maximum amount of light in the image.
>
> Another adaptation performed by Axis’ OptimizedIR technology is visible when the subject is far away from the camera and the whole area is illuminated. When the subject is approaching the camera, the exposure is adapting. When the subject is by the camera, it is illuminated and not overexposed.

However, many manufacturers, such as Dahua, Hikvision, and Panasonic simply call their integrated IR features "smart." Unfortunately, there is no way to know what each manufacturer means by these terms without at least checking spec sheets and/or testing.

"Smart IR" May Not Adjust IR Power At All

Furthering confusion, some manufacturers' "smart" IR features do not actually adjust IR power. Instead, the camera adjusts gain and/or exposure to compensate for subject(s) in the scene. This is shown on Avigilon's H4 model below. The camera's LEDs do not dim, but its web interface shows gain reduce as the subject enters the scene.

Note: ***Click here to watch the demo on IPVM***

In our tests, we have not seen major advantages to those cameras which adjust IR power over those which do not, with gain/exposure adjustment often as effective as IR power adjustment. As an example, the clips below show one camera, with smart IR turned on on the left and off on the right. Subject exposure is very similar in both, with only minimal advantages when adjusting IR power.

Note: ***Click here to watch the comparisons on IPVM***

(Note that with smart IR on, the camera's image takes longer to readjust to full brightness, which we discuss more below)

Determining exactly which method(s) a given camera uses is difficult, and typically requires measuring IR power, as most manufacturers do not disclose how their smart IR works.

Also note that these features typically need to be built into the camera and are not compatible with external IR illuminators, as the processing required to detect an object and adjust as necessary is performed internally.

Beware Exposure Delays

Though not common, some cameras take longer than others for the image to return to full brightness after adjusting IR power and/or gain/exposure. This may result in subjects being missed if they should pass behind a larger object moving through the foreground.

For example, the camera below takes 10-15 seconds for the image to return to full brightness. During this period, the image is very dark, with only the near foreground visible and the background almost totally obscured.

Note: ***Click here to watch the comparisons on IPVM***

There was no direct connection between this issue and IR power adjustment or lack thereof. It has occurred on multiple models in our tests, both with smart IR and without.

IR Reduces Bandwidth Spikes

In general, integrated IR cameras result in lower nighttime bandwidth/smaller bandwidth spikes compared to similar non-IR models, since the added IR reduces the noise created by increased gain (see our Camera DNR (Digital Noise Reduction) Guide). However, the introduction of smart codecs has made it difficult to cite clear averages for these reductions, as some smart codec non-IR cameras may better compress low light scenes than non-smart integrated IR models.

Because smart codecs adjust compression and I-frame interval based on what is in the scene, dark areas may be highly compressed, reducing bandwidth, and I-frame interval increased when no motion is occurring in the scene. Additionally, in some cases, integrated IR models have even produced lower bitrates in low light than full light, as details such as foliage are less visible and easier to compress.

Also note that two key factors affect IR's impact on bandwidth:

- The IR illumination needs to cover the whole scene. If certain areas are still dark, the camera is likely to turn up the gain control to brighten those areas, still requiring high bandwidth consumption.
- To simply reduce bandwidth consumption in low light, use a MBR/VBR cap to VBR ones. The lower bit rates will have limited impact on visual quality as the spike in bandwidth to encode noise adds no meaningful details.

Pros And Cons

There are several advantages to integrated IR models compared to external IR, which have driven their adoption:

- Lower cost: Many/most of the lowest priced cameras include integrated IR, making them a much less expensive option than a non-IR camera with added illuminator.
- Simpler installation: Because the IR is built in, there is no additional equipment to install.
- Simpler powering: Integrated IR cameras are typically powered by standard 802.3af power over ethernet, or in some instances 802.3at. External illuminators typically require separate low voltage, PoE, or even mains power.
- Advanced Features: Higher end cameras offer smart / adaptive IR and automatic beamwidth adjustment, typically not possible with add-on IR illuminators.

However, there are two key drawbacks:

- Shorter range (though improved): Though integrated IR cameras with longer ranges have become more common, with some reaching 600' or more, for super long range illumination (1000'+), standalone illuminators are often still required.
- Bug/insect issues: Integrated IR cameras may attract insects or spiders, which nest/build close to or over the lens of the camera, obscuring the field of view. If this is a problem, a standalone illuminator must be used.

Because of the integrated IR segment's rapid growth and advanced features, it is increasingly more difficult to justify add-on IR illuminators. The clearest driver for standalone illuminators are super long distance and invisible IR, features which are uncommon in integrated IR cameras, though we expect they will continue to improve.

What About White/Visible Light Illuminators?

While adding visible light to the scene will improve low light images, generally IR has been favored, for multiple reasons:

- Less detectable: Where white light creates a floodlight effect across the area it covers, IR is barely perceptible (850nm) or not perceptible at all (940nm) to the human eye. This means that camera locations are much less obvious at night, making them more difficult for intruders to detect and defeat.
- Less light pollution: Additionally, infrared does not contribute to light pollution caused by white light. Many municipalities regulate how

powerful exterior lights may be, how they must be aimed/shielded, etc.

- Better aesthetics: Finally, architects and lighting designers may prefer that IR be used, as they are better able to control the lighting design of a facility.

Because of these reasons, white light is rarely used except in specialty applications.

IR In The IPVM Camera Calculator

The IPVM Camera Calculator shows camera's IR range on top of a map with a red line, so you can see how far the IR is projected to cover, demonstrated below:

Note: Click here to view the demo on IPVM

Conversely, if the target is much closer than the camera's maximum IR range, we should beware of overexposure issue or carefully check adaptive/smart IR features. See some examples in the Calculator here.

Shutter Speed / Exposure

Surveillance users do not need to be photography experts but understanding the basics of shutter speed is critical to avoiding major low light problems. In low light conditions, surveillance video can appear blurry and objects will look like ghosts - all due to issues with exposure.

Here is an example:

We explain the role of exposure and setting shutter speed has in surveillance including a 5 minute video screencast to show you the key issues in action.

Shutter Speed / Exposure

The terms shutter speed and exposure are often used interchangeably, but are not technically the same thing.

Shutter speed refers to how long the sensor is exposed to light. Exposure consists of shutter speed ('how long') combined with iris opening ('how wide'), but is often simply used to mean shutter speed, leading to confusion.

We focus on shutter speed. For details on iris and aperture, see our Lens Iris Tutorial and F-Stop Tutorial.

Automatic Vs. Manual Shutter Speed

Surveillance cameras almost universally default to automatic control of the shutter speed, meaning that the camera will continue to adjust its speed without any intervention of the user. In very bright scenes, shutter speed will be faster, but if it is quite dark, the shutter speed slows.

Users typically may set minimum and maximum shutter speeds to better control how widely the camera may vary its exposure. However, this is not normally required unless issues are visible or manufacturers use poor defaults (such as slow shutter / sens-up, in report below).

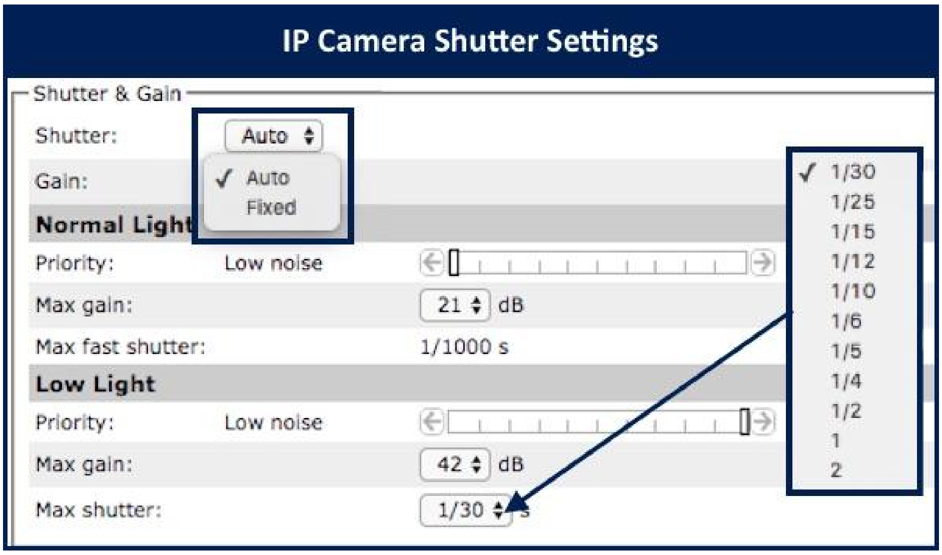

In 2017, nearly all cameras use automatic shutter speed, but manual settings are still used in some specialist applications such as license plate capture / recognition and machine vision.

Range of Shutter Speeds

Shutter speeds can range from extremely fast (1/1,000, 1/10,000 of a second or less) to extremely slow (1/2 of a second or more). Under normal circumstances, since shutter speeds are controlled automatically, fast shutter speeds rarely have negative effects.

However, when using even a slightly slow shutter speed (such as 1/15 second), motion blur may be significant, obscuring details. The threshold where blur occurs varies widely, depending on light level, object speed, and other variables, with no hard and fast rule for what is universally "too slow,."

That being said, we strongly recommend against using shutter speeds slower than 1/30s in almost all cases, as blur is likely even on slow moving objects.

Exposure in Action

Watch the 5 minute video below for demonstration of different shutter speeds and dealing with slow shutter settings.

Note: ***Click here to watch the Impact of Exposure In Surveillance on IPVM***

Fixed Shutter Speed

IP cameras most often allow advanced controls for fixing the shutter speed of a camera. However, we do not recommend this except in very specific applications.

The most common surveillance application for fixed shutter speed is license plate capture. The images below show the same camera in a license plate capture application, using default 1/30s shutter speed and 1/500s shutter. Using defaults, even at slow speeds (<10 Mph) the camera is unable to capture the plate, while at 1/500s, the plate is clearly legible.

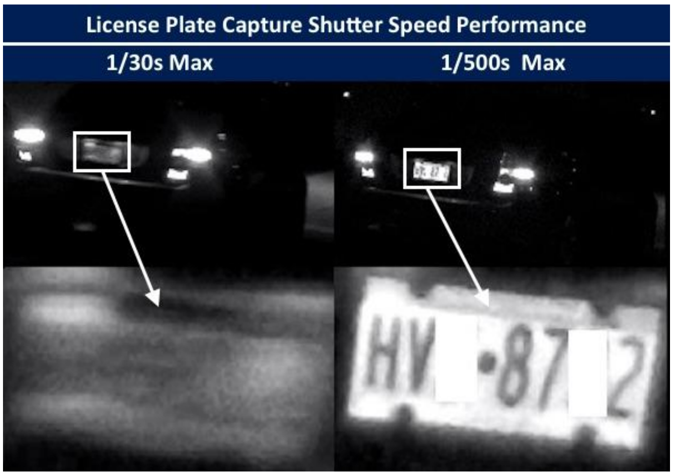

Manufacturer Tricks

In 2017, most manufacturers allow their cameras to automatically adjust shutter to as slow as 1/30S by default. However, users beware: some manufacturers still default to slower shutters which may cause significant blur in low light. See our Camera Slow Shutter / Ghosting Test for details.

Additionally, manufacturers may use other techniques which essentially amount to slow shutter in practical use, such as Sens Up, "Intensifier", or frame integration. Users should beware of these terms, and when in doubt, disable these features to see how performance changes.

Other Blur Issues

Users should also beware of other features in IP cameras which may also cause blur. High levels of digital noise reduction (DNR) used in today's super low light cameras may introduce blur. Additionally, true wide dynamic cameras may introduce blur, as well, as they combine multiple exposures to generate the WDR image, with small movements between. These issues may appear very similar to shutter blur, but they all have different causes. Because of this, technicians may need to adjust several settings to remove blur for best results.

Gain Control / AGC

Gain control is a critical, though often overlooked, factor in low light surveillance video. It is generally only noticed when the negative side effective of aggressive gain levels are seen, namely noise / snow on screen. The picture below, from one of our parking lot tests, is a prime example of this problem:

Clearly, the noise is a problem and ideally you would want to remove it. This raises important questions about how to use gain control effectively.

In the introductory video below, we provide a real-time demonstration of gain control:

Note: ***Click here to watch the Intro To Gain Control video on IPVM***

To learn more, we conducted a series of experiments in a variety of scenes. For each scene, we captured video, images and bandwidth consumption. Here are the scenes we tested:

- Black and White Mode .5 Lux (Dark)
- Black and White Mode 20 lux (Low Light)
- Color Mode 4 Lux (Moderately Dark)
- Color Mode 300 Lux (Daytime)

Our testing was done across 3 HD cameras from Avigilon, Axis and Bosch to see a range of performances.

We share our results and answer the following key questions:

- How significant does image quality vary with different gain control settings?
- Should you use gain control?
- What is the right gain control settings to use?
- What alternatives should you seek to using gain control?
- How do manufacturers approach to gain control differ?
- How does gain control differ between Black & White and Color modes?
- What is the bandwidth impact of different gain control settings? What impact does light levels have on bandwidth impact?
- What impact do VBR and CBR streaming modes have on using gain control?
- What impact does digital noise reduction (include 2D and 3D DBR) have?

Key Findings

Let's start with the key findings of our tests:

- **Gain Control Very Important**: In almost any night time scenes, without gain control, surveillance video would be very dark and almost practically useless. While gain control is generally ignored, its role is critical. In low light conditions, trying to turn it off to remove the grain/noise will only result in far worse video quality.
- **Gain Control is Automatic**: As the name states, Automatic Gain Control, is automatically controlled by almost every camera. While it can produce lots of noise, disabling it generally will make things even worse.
- **Aggressive AGC**: Unlike in commercial videography where gain control is used minimally and in moderation, surveillance camera manufacturers tend to use massive amounts of gain control. This is necessary because unlike film production, it is very hard to control scene lighting in surveillance.
- **Adjusting Gain Control**: If you want to minimize the bandwidth impact and visual noise inherent in gain control, the two gain adjustments possible are (1) fixing the gain or (2) capping the gain. The former is dangerous unless you can guarantee constant lighting. The latter, assigning a cap, can be useful.
- **Gain for B&W and Color**: Both B&W and Color modes use gain and display similar characteristics - bandwidth spikes, visual noise, etc. The big difference is that cameras activate gain at much higher light levels for color than in black & white modes (color at ~50 lux, b&w at ~20 lux).

- **Gain and Bandwidth Increases Tightly Correlated**: As gain increases, bandwidth increases as well in a hockey stick curve. Interestingly, this appears to occur regardless of the light level or mode used.
- **Use MBR**: Since bandwidth is nearly guaranteed to spike as gain increases, we strongly recommend setting a maximum bit rate.

Below is a video that shows and explains our key findings with references to our test video and images:

Note: Click here to watch the video on IPVM

Configuring and Optimizing Gain

Gain is generally controlled by two modes:

- Levels: For example, high, medium, low - this is a coarse grained control.
- dB: For example, 0dB, 15dB, 30dB, etc. A fine gain control, the range tends to be from 0dB (off) to 45dB or higher (very aggressive).

Gain is controlled typically in one of three approaches:

- Auto: By default, most cameras automatically determine what level gain should be.
- Fixed: As an alternative, a user can lock gain to a specific level (e.g., always 6 dB). This can be dangerous - if you do this at too low a level (like 6dB) and the scene becomes even moderately dark (say 5 lux), the image quality is going to be quite poor. On the other hand, if you set it at a high level (say 27 lux) and the scene becomes bright, the image will be distorted by heavy grain / noise.

- Cap: Some cameras allow the admin to cap the maximum gain the camera can choose. This allows the camera the autonomy to fluctuate the gain but blocks it from going too high for an admin's preference. This can be useful for reducing some visual noise and reducing bandwidth spikes. However, of course, keep in mind such caps will make the image darker than possible with max gain.

The video below provides demonstrations on how to use and optimize gain using Axis and Bosch cameras as an example:

Note: ***Click here to watch the video on IPVM***

Gain Not Comparable Across Manufacturers

Beware, even if 2 manufacturers list the same gain levels (e.g., 20dB or 35dB), it does not mean they are the 'same' or deliver the 'same' low light performance. How they add gain / process video can widely vary.

Gain and Super Low Light Performance

You may be aware of manufacturers marketing super low light performance, with names like LightFinder, LightCapture, Stellar, DarkFighter, etc. All of these are using advanced forms of gain control to improve low light image quality. However, like 'traditional' gain control, they tend to increase bandwidth consumption further (though by how much varies depending on the noise reduction techniques they use). Super Low Light does not, however, require IR nor special imagers. The special functionality comes from advanced gain control / image processing.

They are not all equal and they cannot be compared abstractly. You can review our tests, e.g., Super Low Light vs Integrated IR Shootout and individual camera reports to see low-light performance.

Varying Gain Control Demonstrations

To get a sense of how gain impacts overall video quality, the below comparisons show the same scene with the same lighting with different gain levels. The top image has the maximum amount of gain with each descending image at a lower gain level. Take a look:

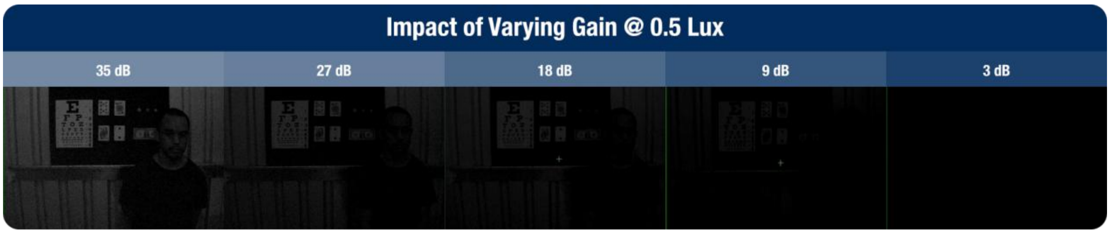

In the above scene, almost anyone would agree that the top image is the best. With this dark scene, high gain becomes critical to make anything out. Indeed, even with 3dB of gain, the scene is pitch black.

Let's contrast to another slide below with 40x the amount of light. Take a look:

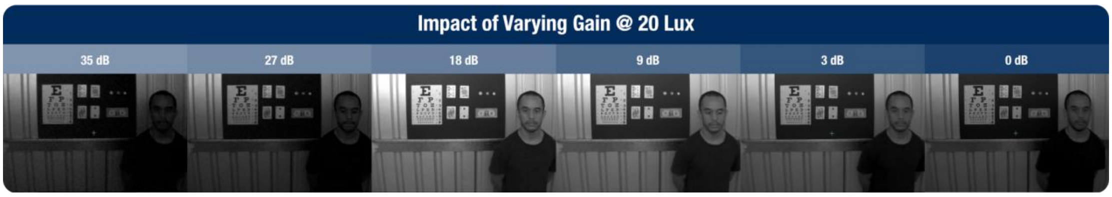

Which do you think is better? This one is certainly much more complex than the first slide with sub 1 lux lighting. In this scene, with 20 lux, the image with maximum gain is somewhat dark and suffers from a lot of visual noise. The images actually look better with less gain, though one could debate which of the low gain images were best. As a point of reference, we estimate this camera's automatic gain setting for this scene to be about 6 to 9 dB.

Color Mode - Gain Impact

The same tradeoffs with gain variance occurs with color mode. The big difference, as the comparison below shows is that with color, gain is needed more even at higher light levels. By default, at 4 lux, all 3 cameras automatically adjusted their gain control to maximum levels. You can see this with the clear visible noise across each image.

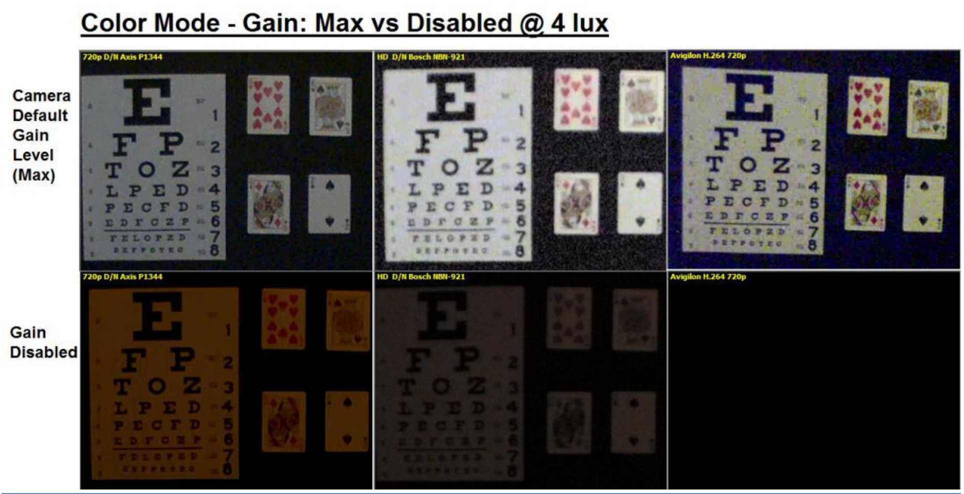

By contrast, with gain disabled, the images are significantly darker and the colors are more subdued (Bosch), missing (Avigilon) or altered (as in the case with the Axis camera).

Gain's Impact on Bandwidth

One of the most oft cited problems of low light surveillance is increased bandwidth consumption. Perhaps the most interesting finding of this test is the strong correlation of gain levels and bandwidth, regardless of the mode of the camera or the light level of the scene.

Below is a chart that depicts the relationship. The numbers are based on .5 lux B&W mode for Axis though the same pattern was displayed with higher light levels, color mode and the Avigilon camera.

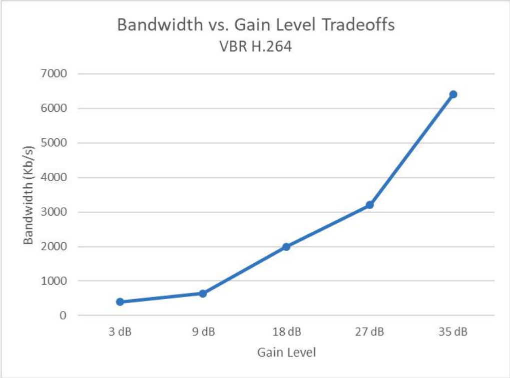

The bandwidth penalty becomes increasingly steep as gain levels increase. We have noticed some cameras with very high noise and very high bandwidth consumption in low light. This looks to be an outcome of manufacturer choosing aggressive maximums for automatic gain control.

Based on this, we recommend you check your low light bandwidth consumption. If it is quite high, see if a gain control cap exists and try lowering it from the manufacturer's default maximum.

The other option is to use MBR. The image below shows our test results of taking a scene with maximum gain and limiting its bandwidth to 1 Mb/s. The overall image quality is practically unchanged but at a fraction of the bandwidth.

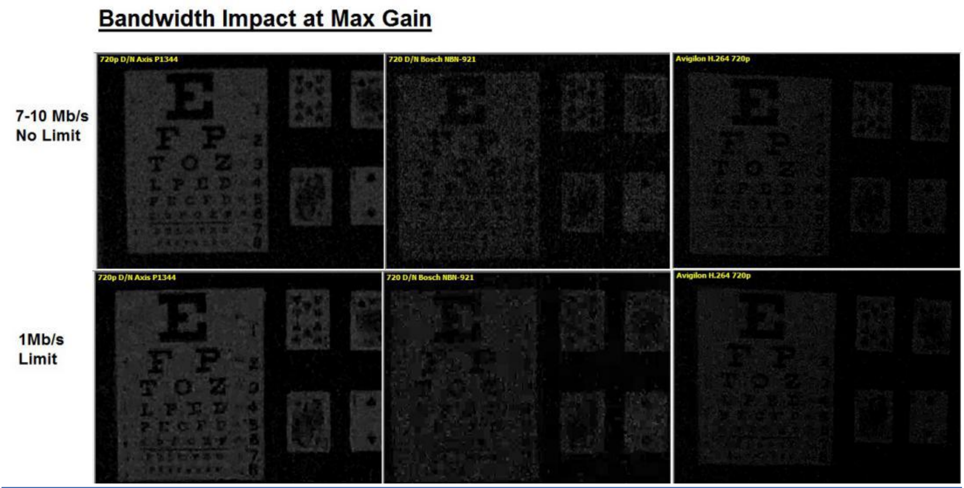

Gain Impact Across Cameras

Finally, here is a summary comparison of gain across light levels, gain settings and cameras tested. Take a look and compare the differences yourself. Perhaps the most interesting element we saw here was that with bright light (300 lux at the bottom of the comparison), the Avigilon appears to provide the sharpest image details. However, with low gain, Avigilon's image essentially disappears.

Digital Noise Reduction and Gain

IP cameras may employ a technology to reduce the noise generated by high gain values, known as Dynamic Noise Reduction (DNR).

DNR aims to reduce digital noise by processing individual frames (2 dimensional) or across a series of frames (3 dimensional) to identify scene changes in order to determine what is and is not noise. By reducing overall noise in the image, bandwidth requirements are also reduced, sometimes dramatically.

For example, taking a sample of cameras from in our bandwidth vs. low light and DNR tests, bitrates increased by nearly 15x from day to night without using DNR (an average of 10 Mb/s vs. ~0.9 Mb/s). Turning DNR on, these spikes dropped to ~3x (~3.24 vs. ~0.9 Mb/s), still substantial, but much improved versus not using DNR.

Camera	Resolution	FPS	Day	DNR Off	Increase	DNR On	Increase
Bosch NBN-932V	1080p	10	0.64	9.06	1316%	3.12	388%
Samsung SNB-6004	1080p	10	1.89	20.00	958%	2.58	37%
Sony SNC-VB630	1080p	10	2.49	17.60	607%	8.24	231%
Bosch NBN-733V	720p	10	0.18	5.12	2744%	0.30	67%
Hikvision 864	720p	10	0.56	10.24	1729%	5.28	843%
Samsung SNB-5004	720p	10	0.68	2.64	288%	2.54	274%
Sony SNC-VB600B	720p	10	0.16	5.33	3231%	0.60	275%
Averages			**0.94**	**10.00**	**1553%**	**3.24**	**302%**

However, users should be careful not to apply too much digital noise reduction, as high levels can create blur on moving objects, similar in appearance to slow shutter.

For full details, see our Camera Digital Noise Reduction Guide and Testing Bandwidth vs. Low Light.

Quiz Yourself on AGC / Gain

Take the AGC / Gain quiz now.

Camera DNR (Digital Noise Reduction)

Bandwidth spikes are a significant video problem.

An IPVM study found 250 - 500% increase in bandwidth from day to night (see: Testing Bandwidth vs Low Light).

Digital noise reduction is key for reducing night time bandwidth and preventing dramatic spikes which can kill your surveillance storage. In this guide we answer three key questions:

- How much of an issue are bandwidth spikes?
- How much does DNR reduce bitrate?
- What impact does DNR have on image quality?

In order to answer these questions we tested 11 cameras from eight manufacturers in an approximately 1 lux scene, adjusting DNR to various levels from off through maximum, while measuring bandwidth.

- Arecont Vision AV3116DNv1
- Avigilon 1.0-H3-B1
- Axis Q1615
- Bosch NBN-733V
- Bosch NBN-932V
- Dahua IPC-HF3100N
- Hikvision DS-2CD864FWD
- Samsung SNB-5004
- Samsung SNB-6004
- Sony SNC-VB600B
- Sony SNC-VB630

What is DNR?

Digital noise reduction, as it name states, aims to reduce digital noise present in low light surveillance images.

This noise comes from automatic gain control. Gain is an essential tool for delivering any image in low light surveillance, as shown in this comparison of varying gain:

Impact of Varying Gain Control at 0.5 Lux

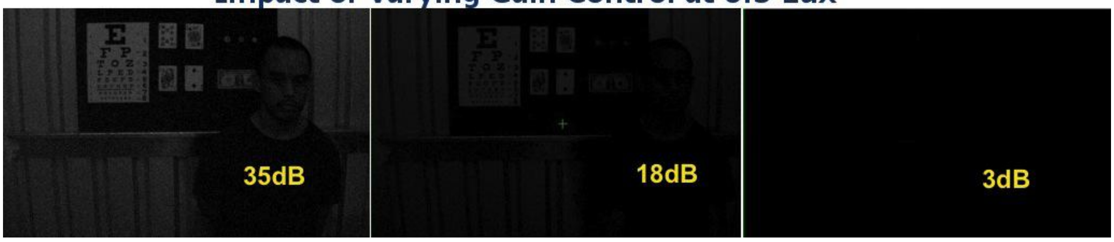

However, there are 2 important side effects:

- Visible noise on the image that obscures fine details (which is still better than the alternative without gain - a completely dark image)
- Increase bandwidth consumption because the encoder sees the visible noise as moving objects that are more difficult to encode.

Digital noise reduction are image processing techniques that aim to eliminate the visible noise to improve the image quality and to lower the bandwidth consumed in encoding.

The two commonly cited types of DNR are 2 dimensional and 3 dimensional DNR. 2 dimensional processes each individual frame / image individually, while 3 dimensional processes across a series of frames over time.

The major cost is having the computing power on-board the camera to do this processing.

The major risk is blurring objects in the image when trying to reduce noise (shown in depth at the conclusion of this guide).

DNR Performance

We tested 11 cameras from seven manufacturers. Of these, only three (Avigilon, Axis, Dahua) did not allow any control of digital noise reduction. We tested all cameras in the same scene, our interior conference room, to see what effects different DNR settings had on bitrate.

DNR On vs. Off

In most cameras tested, bandwidth dropped by an average of ~70% when turning DNR on at default settings. Only a single camera, the Samsung SNB-5004, saw reductions below these levels.

Camera	Resolution	FPS	DNR Off	DNR On	Reduction
Bosch NBN-932V	1080p	10	9.06	3.12	66%
Samsung SNB-6004	1080p	10	20.00	2.58	87%
Sony SNC-VB630	1080p	10	17.60	8.24	53%
Bosch 733	720p	10	5.12	0.30	94%
Hikvision 864	720p	10	10.24	5.28	48%
Samsung 5004	720p	10	2.64	2.54	4%
Sony VB600B	720p	10	5.33	0.60	89%

All measurements in Mb/s

DNR Defaults vs. Max

The effects of moving from default settings to the maximum offered by the camera varied widely. In some cameras, such as the Arecont AV3116DNv1 and the Samsung SNB-6004, this increase in DNR had almost no effect. However, others saw much greater reductions, with the highest being the Hikvision 864 (~90%) and the Bosch NBN-733V (~75%). Note that increasing DNR settings to high levels may result in blur, shown below.

Camera	Resolution	FPS	Default	Max	Reduction
Bosch NBN-932V	1080p	10	3.12	0.94	70%
Samsung SNB-6004	1080p	10	2.58	2.56	1%
Sony SNC-VB630	1080p	10	8.24	4.40	47%
Arecont AV3116DNv1	3MP	10	3.04	3.00	1%
Bosch 733	720p	10	0.30	0.08	74%
Hikvision 864	720p	10	5.28	0.52	90%
Samsung 5004	720p	10	2.54	2.43	4%
Sony VB600B	720p	10	0.60	0.38	37%

All measurements in Mb/s

DNR Blur

Applying high levels of digital noise reduction is not without drawbacks, as it may introduce motion blur which much resembles the blur created by slow shutter. This blur is mainly caused by temporal noise reduction, which compares changes in the scene between frames in order to determine what is noise and what is not.

For example, the video below shows a Samsung SNB-6004 at maximum DNR settings, with pronounced blur and ghosting as our subject walks through the scene.

Note: Click here to watch the video on IPVM

Compare this to lower settings such as the defaults in the video below. Blur is still present, though drastically reduced.

Note: Click here to watch the video on IPVM

Finally, with DNR off, visible noise increases dramatically, but blur is eliminated:

Note: Click here to watch the video on IPVM

Test your knowledge

Take this 6 question quiz now

Panoramic / PTZ

Fisheye Panoramic Camera

Fisheye panoramics have become widespread, with most manufacturers offering fisheyes.

However, with many options, it can be confusing to understand the tradeoffs. We explain:

- Why Panoramic Cameras
- Fisheye Panoramic Basics
- 'Warped' Video and Dewarping
- Camera Side vs Client Side Dewarping
- VMS Integration
- Fisheye Mounting Heights
- Ceiling Vs. Wall Mounting
- Panoramic WDR Availability
- Panoramic Low Light Weakness
- IR Panoramic Availability
- Image Detail Vs. Multi-Imager
- Fisheye Manufacturer Availability

Multi-imager panoramics, a common alternative to fisheyes, are covered separately in our Multi-Imager camera guide and Repositionable Multi-Imager Camera Guide

Why Panoramics?

Traditional surveillance cameras only capture relatively narrow fields of view(FOVs). "Average" camera fields of view are typically no more than 90 degrees, a quarter of a circle, though some wide angle 110-130 degree options have become more common. Typically, if you want a camera to see behind itself or to the extreme left or right, you would need to add other cameras.

By contrast, panoramic cameras let you deploy a single camera that can "see" in "every" direction. However, the broader the area you cover, the lower the pixel density / PPF will be, which will significantly reduce details captured in any direction. We examine these tradeoffs throughout this tutorial.

Fisheye Panoramic Cameras

The most common panoramic type is the fisheye camera. These cameras have super wide angles lenses (~180° or higher in some cases) typically built into dome housings. It uses a single lens and a single imager covering the entire area.

Fisheye images are typically not used in their normal, warped state (seen below), as the top-down view is less useful for monitoring, with some objects/subjects appearing upside down or sideways, depending on their location in the room.

To make them more usable and flattened like typical surveillance video, they are dewarped by special software. The image below shows the original fisheye image and dewarped views.

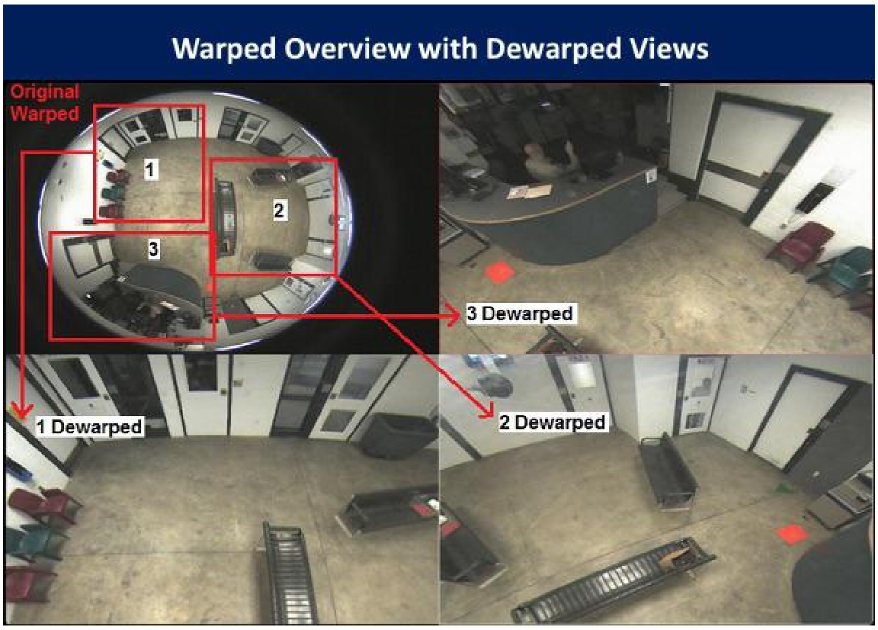

This dewarping is a critical component of fisheye panoramics. Warped video, by itself is practically useless. Manufacturers can dewarp at the camera side, client side or both. The choice of where to dewarp impacts VMS integration and usability greatly.

Camera Side Dewarping

Dewarping at the camera side sends a flat video stream that any VMS can display just like a traditional camera. Camera side dewarping has become increasingly popular, with some manufacturers now offering only this style of dewarping, and no VMS client SDK. The upside is easy integration, since these cameras may appear similarly to PTZs to the VMS. However, there are three key limitations.

VMS Stream Support

Many models of panoramics which dewarp on the camera send multiple streams, such as overview, panorama, PTZ area 1, PTZ area 2, etc. Examples of these are shown below:

Note: Click here to view samples on IPVM

The VMS must integrate all of these streams, typically via direct drivers, in order to fully support the camera. However, many VMSes do not support all streams from common cameras, especially those using ONVIF instead of direct drivers. For example, neither the Panasonic WV-SFV781 nor Hikvision DS-2CD63C2 are properly integrated with multi streams in Exacq or Avigilon, despite both being current high end models from major manufacturers.

PTZ Control Limitations

In addition to the multiple stream integration, the VMS must also support PTZ controls for camera side dewarping windows. Without this integration, these windows must be set up via the web interface and are essentially static in the VMS.

Camera Side Dewarping Recording Limitations

When using camera side dewarping, the VMS records only what the camera stream was viewing at the time. This means that if an overview stream (360 panorama, warped fisheye, etc.) is not recorded in addition to digital PTZ windows, no overview of the scene is provided, with events potentially missed. This is in contrast to client side dewarped cameras, which record only the full stream, with dewarping performed at the client on playback.

Client Side Dewarping

In contrast to camera side dewarping, client side provides immersive controls, which allow the user to view and record only a single overview stream, and dynamically change viewing windows in the VMS client on both live and archived video. For example, in the clip below, we switch from warped overview to PTZ window, 180, and 360 degree panoramas.

Note: Click here to view samples on IPVM

The main advantage of this type of dewarping versus camera side is that it requires only a single stream to be recorded while maintaining overview information, without worries of multiple stream integration or virtual PTZ windows looking elsewhere if an event occurs.

Client Side Dewarping Drawback: Additional Integration Required

Dewarping each manufacturers' fisheye cameras has typically required a separate SDK for each. This is simple when using the camera manufacturer's VMS or client (e.g., Hikvision fisheye cameras in iVMS-4200). However, if you want to use a 3rd party VMS, then that VMS must integrate that special software. Often, this is not done as it is complicated, expensive and proprietary to each fisheye panoramic manufacturer, for relatively few cameras sold.

Some VMS platforms, such as Network Optix and Axxon Next offer universal dewarping of any fisheye camera, eliminating this development/integration concern, but this is relatively uncommon.

Panomorph Lenses

Panomorph lenses are a proprietary type of fisheye lens, sharing some key similarities with a few important differences. Panomorph technology is patented by ImmerVision, who designs the lenses and provide the dewarping SDK to others for integration. Originally, ImmerVision lenses were sold in C/CS mount version only, to be used with third party box cameras, but more recently, they have developed M12 lenses, as well, for use in dome cameras, with resolutions up to 12MP. Manufacturers such as Arecont, Hanwha, and others have released models using these lenses.

Panomorph lenses claim two key advantages over typical fisheye models:

- Theoretically higher PPF: ImmerVision's main claim is that they use more of the camera's sensor by using an elliptical, instead of circular fisheye format which allows for more pixels at the edges of the FOV and theoretically greater image quality.

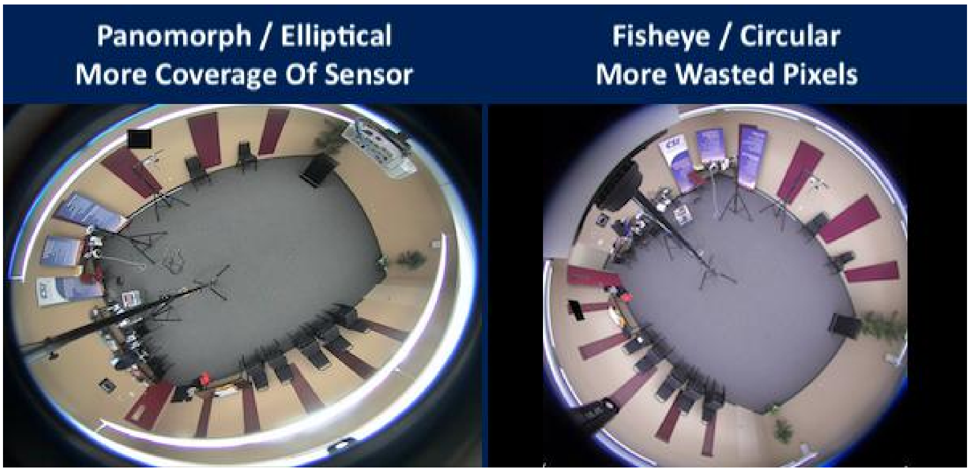

However, we have not found substantial practical differences in our tests, with standard fisheye models performing as well as ImmerVision-based models.

- Broader support: Most, but not all, major VMSes support ImmerVision, making them more widely supported than fisheye models using manufacturer-specific SDKs. This SDK encompasses all of ImmerVision's products, including both CS mount lenses as well as domes using their M12 lens.

Panomorph Drawbacks

However, there are key drawbacks as well:

- Performance gap has narrowed: Additionally, while low light and WDR have historically been challenging areas for fisheye cameras, making a high end camera plus ImmerVision lens more attractive, dome models with true WDR capability, mechanical cut filters, and integrated IR have become increasingly common, making these add on lenses even less attractive in areas requiring these features.
- Typically lower resolution: While 12MP panoramics have become relatively common, many ImmerVision enabled models are limited to 5MP and below, as is their high end CS mount lens. Note that some higher resolution models, such as Hanwha's PNF-9010 use ImmerVision, but the majority do not.
- Higher cost: Panomorph lenses are expensive. The original 1.3MP model sells for about $500 USD, plus the cost of the camera. The newer 5MP model (made by Fujinon) sells for over $1000. By contrast, single imager fisheye models sell for $500-800 on average, and do not require additional integration and setup labor.

Wall Mount Only Fisheye Cameras

While most panoramic cameras may be ceiling mounted or wall mounted, with options for 360° or 180° dewarping, some specialized models intended only for wall mounting are available. Unlike typical fisheye

models, these cameras typically output dewarped 180° (or less in some cases) panorama views only, with no options for multiple digital PTZ views or other panorama types.

Some of these are standalone cameras (see our test of Vivotek's CC8370-HV) while others are integrated into audio door stations (such as the Axis A8105-E).

Wall Mount Fisheye Camera Pros And Cons

The main advantage of wall mount fisheye cameras is simplicity, as they generally dewarp onboard and output a panorama stream, instead of requiring more complex camera or client side dewarping integration and setup.

However, these models are generally lower resolution than typical fisheye cameras, 3-5MP instead of 6-12 found in typical high end high resolution fisheye models.\

Fisheye Mounting Height Is Key

Mounting fisheye panoramic cameras low and close to targets is critical. The pixel density of panoramic cameras falls massively as the subject moves farther away from the camera, so increased distance resulting from mounting cameras higher worsens this issue. For instance, a subject who is 5 feet away from a 1.3MP 360 camera will be captured at ~40ppf but the same same subject just 10 feet away will be captured at ~20ppf. Forget about if the subject is even 20 feet away. The PPF will be down to about 10ppf meaning that you will see no more than a blob.

To illustrate this, this animation shows the difference in details delivered by the same panoramic camera, mounted at 10', 15', and 20'.

Note: Click here to view animated samples on IPVM

Either avoid mounting high or use extension mounts to get the camera closer (of course, with the recognition that this may be aesthetically displeasing).

Ceiling Vs. Wall Mounting

The decision of whether to mount a fisheye camera on the ceiling or wall essentially depends on where subjects are likely to move through the scene.

- Wide area of interest: If objects to be observed are likely to move throughout the scene and not just near the camera, using a ceiling mount camera is likely the better option, as it may increase details of objects too far away for a wall mount camera to capture.
- Narrow area of interest: If objects are likely to enter/exit through one or two locations, wall mounting a fisheye camera is likely to provide better details of subjects as they pass due to its lower angle of incidence (discussed above).

Readers should see our test Panoramic Cameras Wall vs Ceiling Mount for full details and image quality examples on this issue.

Fisheye True WDR Availability Improving

True WDR performance has become common in fisheye models, where previously it was available in few cameras or required a separate camera and Panomorph lens to be used. For example, the clip below shows WDR performance of three cameras, two of which are true WDR (left and right) and one which is not (center).

Note: Click here to view animated samples on IPVM

Fisheye Historically Poor Low Light

Panoramic cameras have historically been especially poor in low light, often displaying noisy or dark images in even slightly low light. This is due to two contributing factors:

- High F-stop lenses: Fisheye lenses typically have a high f-stop (f/2.0 vs f/1.2 for traditional cameras) letting in ~75% less light.
- High resolution: Since fisheye models are frequently 5MP, 6MP, 12MP, etc., pixels are simply smaller and receive less light.

Because of this, users should beware when specifying fisheye models in even moderately low light (3-5lx), as noise and artifacting is possible.

IR Fisheye Increasing

To combat fisheye cameras' poor low light performance, many manufacturers have introduced integrated IR models, with illuminators covering a 360° field of view. For example, the comparison below shows two 12MP high end models, with the Hikvision 63C2 outperforming the Panasonic 781L in low light due to its integrated IR.

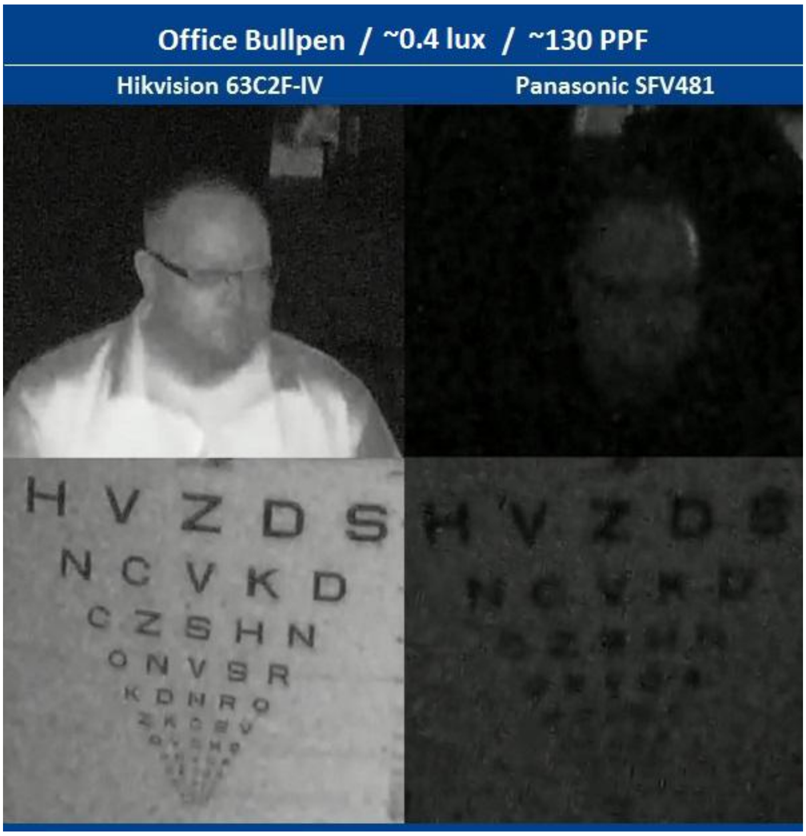

Note that despite these advancements, though, users should be very careful about using panoramics to monitor dark areas, as the performance shown above is not typical of all panoramic (fisheye or multi-imager) cameras.

Fisheye Availability

In 2018, nearly every manufacturer includes at least one fisheye panoramic camera in their lineup, with most offering more than one in differing resolutions, feature sets (WDR/IR), indoor and outdoor versions, etc.

Contrast this to just a few years ago when fisheye panoramic models were considered specialist cameras and only available from a handful of manufacturers.

Also note that fisheye panoramic cameras remain more common than multi-imager models, but the multi-imager segment has grown significantly in the past 1-2 years.

Image Details Vs. Multi-Imager

When considering wide angle panoramic coverage, users should carefully consider PPF requirements, as larger scenes are likely better served by multi imager models.

For example, the image below shows the relative detail between a single imager fisheye camera, a 360° multi-imager (90° per imager), and a 180°. With the subject ~10' from the cameras details of the subject and legible lines of the test chart are notably less in the fisheye model.

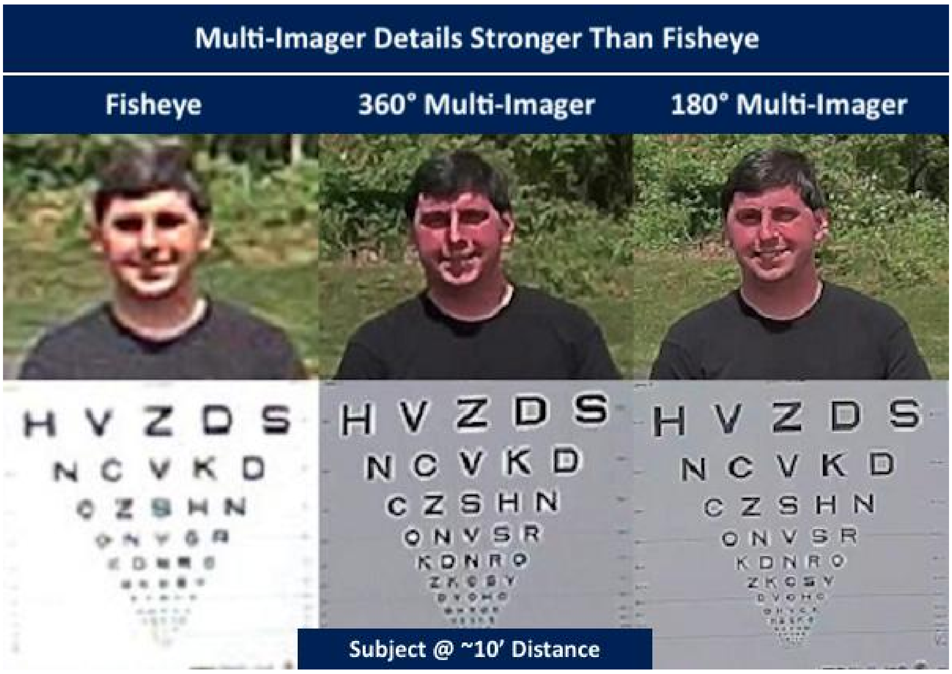

With the subject at ~30' from the cameras, the fisheye panoramic provides no usable details whatsoever, while the 360° multi-imager still provides two legible lines of the chart and some rough subject details. At this distance, the 180° multi-imager still provides recognizable details of the subject.

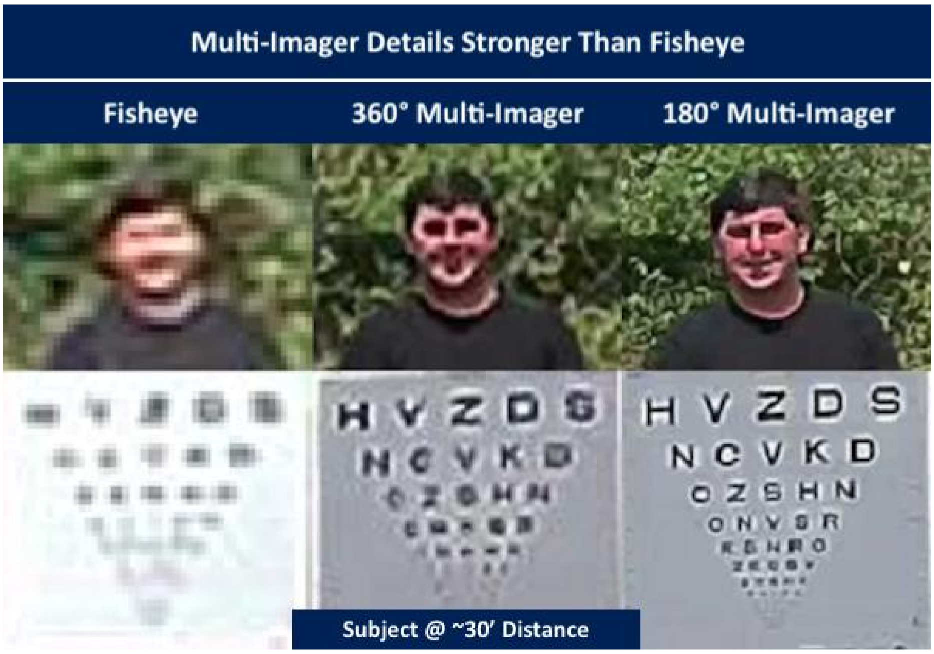

Note: an earlier version of this guide covered both fisheye and multi-imagers. This has been split into 3 guides to reflect the growth in panoramics, for more see our Multi-Imager camera guide and Repositionable Multi-Imager Camera Guide.

Multi-Imager Cameras

Multi-imager usage continues to grow, with most manufacturers now offering at least one model, making them an attractive option for covering wide areas compared to multiple fixed cameras traditionally used, such as this ~200' wide parking lot:

We review key issues impacting multi-imager camera selection, including:

- Multi-Imager Basics
- 180° Vs. 360° Vs. 270° Models
- Separate Streams Vs. Stitched Views
- Stitched View Aspect Ratio Issues
- Rotated Imagers vs. Standard
- Multi-Imager Licensing
- Integrated IR
- Advantages/Disadvantages Vs. Repositionable Multi-Imagers
- Image Details Vs. Fisheye Cameras

For reference, users should also see our Repositionable Multi-Imager Camera Guide and Fisheye Panoramic Camera Guide

Multi-Imager Basics

Multi-imager cameras use multiple image sensors in a single housing (most often dome, but sometimes bullet) to cover a wide area. Instead of using a single lens with a super wide angle of view, typical in fisheye models, these multiple imagers use a narrow angle of view, ~45° or 90° to cover a 180° or 360° area.

There are several types of multi-imager models, with varying coverage ranges, including 180°, the most popular type of multi imager camera, typically wall mounted to view a wide area:

As well as 360° models, which view all directly similar to a fisheye camera, though in some cases with a deadspot directly below the camera:

Finally, some multi imagers are offered in a 270° configuration, such as the Pelco Optera, specifically intended for corner mount applications. These cameras allow coverage of a 270° FOV (e.g., two sides of a building plus parking lot) as well as the area directly beneath the camera, shown below. 270° multi imagers are the least common style.

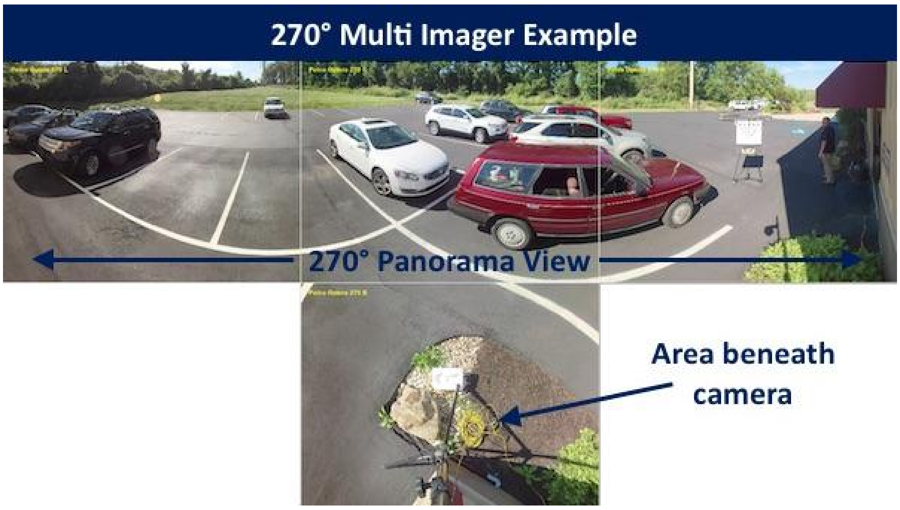

Separate Streams vs. Stitched Views

Multi imager cameras have historically sent separate streams for each camera, which must be manually placed in order in the VMS to provide the full 180 or 360 view, e.g., this 2x2 view:

However, some newer models, such as the Vivotek IR Multi Imager or Pelco Optera stitch streams together to provide the panorama in a single stream. The Vivotek MS8391, for example, outputs a single 7552x1416 stream, seen here:

Stitched streams are moderately more usable as they allow the full panorama to be navigated at once, instead of moving from camera to camera. However, when placed in a typical layout with other cameras, they may appear small and/or strange due to their non-standard wide aspect ratio, seen here:

Multi-imager cameras do NOT support any immersive controls nor do they require any dewarping. The video from them is 'flat' like any traditional camera. The only difference is that when the video is 'stitched', it has a far wider aspect ratio. Since this can cause problems in VMS layouts, often non-stitched is easier to display.

Standard Vs. Rotated Imagers

While most multi-imager cameras use imagers in their standard rotation (4:3/16:9), a few new models rotate imagers 90° to create a taller aspect ratio (Dahua PDBW8800, Hanwha PNM-9020V, Hikvision PanoVu).

This approach has key drawbacks compared to those using standard rotation:

- Lower pixel density: Compared to cameras which use 1080p imagers in their normal rotation, those rotating the imagers produce ~45% fewer pixels per foot. Users should beware of this when calculating pixel density based on the "8MP (4x1080p)" claimed by these cameras.
- Skewed outer imagers: Additionally, the outer imagers of these cameras are skewed compared to those using typical standard rotation, resulting in large areas of these imagers' being wasted on sky or ceiling compared to others, shown below.

Multi Imager Licensing

Because multi-imager cameras require no dewarping, unlike fisheye panoramics, VMS integration is typically quite easy but licensing is varied. The VMS 'sees' the multi-imager as an collection of multiple camera feeds, similar to a 4 channel encoder.

The main issue or variance is licensing. Some VMSes only charge a single license because it is viewed as one camera, or charge only a single license per IP address/MAC address. Others charge a license for each imager inside since the camera contains multiple imagers/feeds.

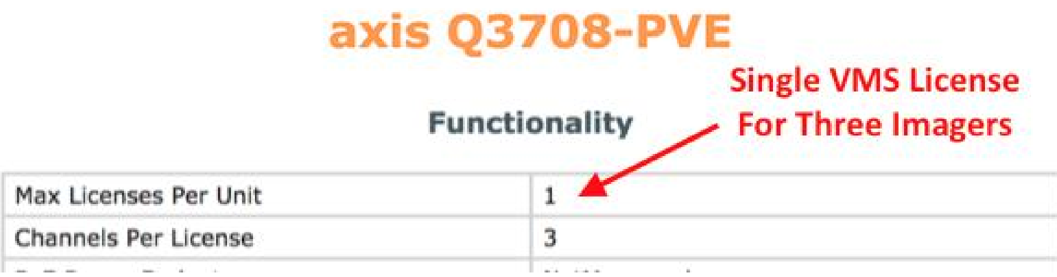

In 2018, common multi-imager models use only a single VMS license. However, it is important to check what one's preferred VMS charges because this can increase costs by hundreds of dollars.

Integrated IR

Integrated IR multi-imager models have become more common, but are still a small minority (which we expect to increase with time). Notably, integrated IR models easily outperform non-IR multi-imagers in our low light tests, as shown in this sample from our test of the Dahua PDBW8800:

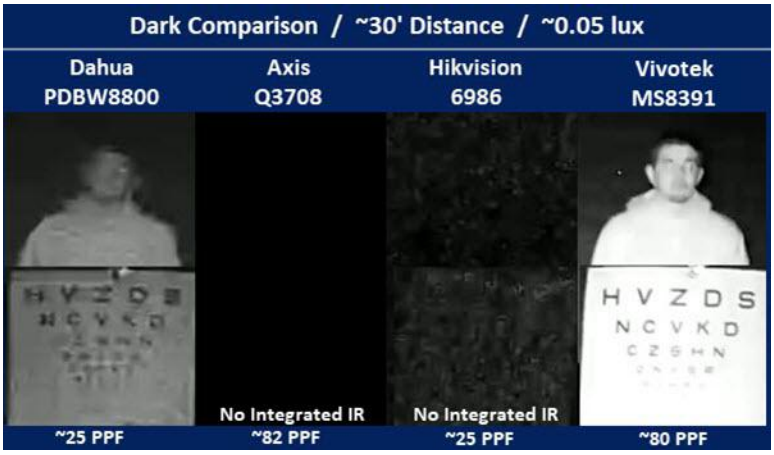

Repositionable Multi Imagers

Another more recent development in multi-imager cameras is repositionable heads. These cameras allow users to move imagers to best cover their scene, if a typical 360 or 180 view is not ideal, such as hallway intersections or wider areas requiring better details of specific points.

This video demonstrates repositioning the heads of the Avigilon HD Multi Sensor camera.

Note: Click here to watch the Avigilon HD Multisensor video on IPVM

Repositionable head models are less common than fixed lens models, though increasing, with options from Arecont, Avigilon, Axis, Hanwha, and Vivotek now available.

Advantages/Disadvantages Vs. Repositionable Multi-Imagers

Compared to repositionable multi-imager cameras, fixed lens models have two key advantages:

- Simpler setup: Unlike fixed imager models which ship aimed and focused, users must manually aim and focus each individual sensor in repositionable models, adding to installation time.
- Smaller size: Repositionable models are larger than fixed imager models and may not be preferable aesthetically.

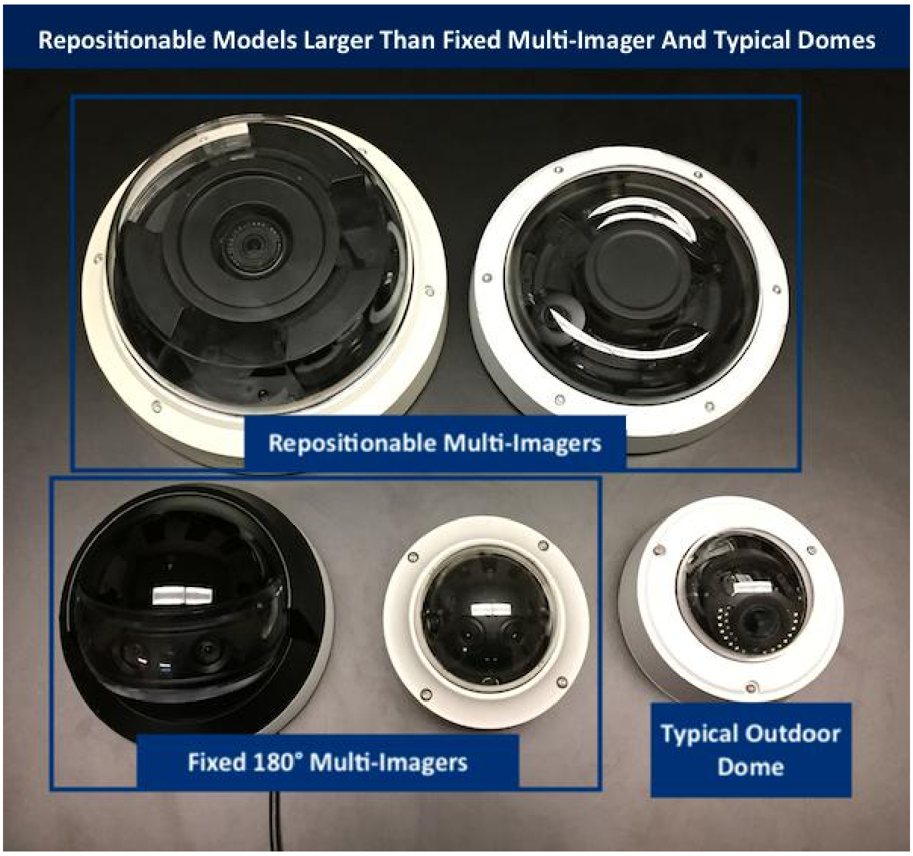

Repositionable models' key advantage versus 180° and 360° multi imagers is flexibility, as they may be configured to cover odd shaped areas which may not be properly served by 180/360 models.

For more on these issues, see our Repositionable Multi-Imager Camera Guide.

Image Details Vs. Fisheye Cameras

When considering wide angle panoramic coverage, users should carefully consider PPF requirements, as larger scenes are likely better served by multi imager models.

For example, the image below shows the relative detail between a single imager fisheye camera, a 360° multi-imager (90° per imager), and a 180°. With the subject ~10' from the cameras details of the subject and legible lines of the test chart are notably less in the fisheye model.

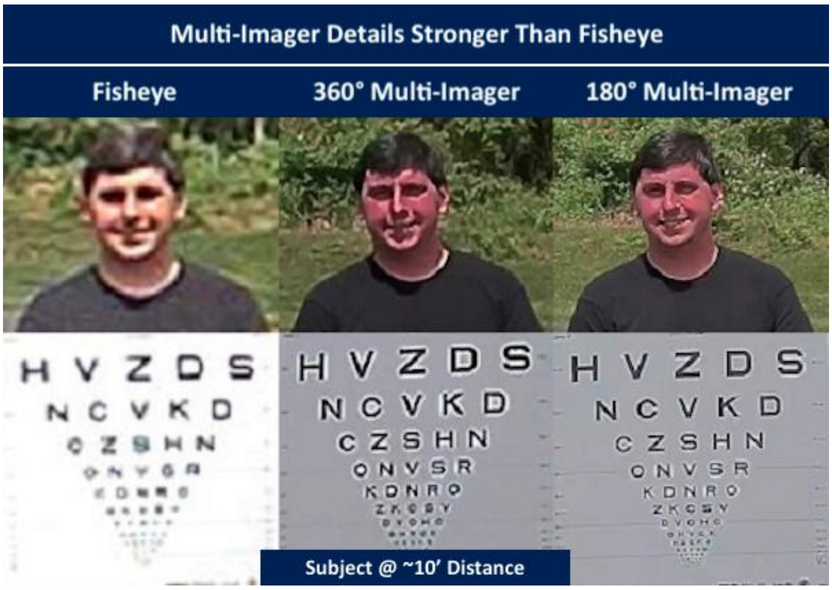

With the subject at ~30' from the cameras, the fisheye panoramic provides no usable details whatsoever, while the 360° multi-imager still provides two legible lines of the chart and some rough subject details. At this distance, the 180° multi-imager still provides recognizable details of the subject.

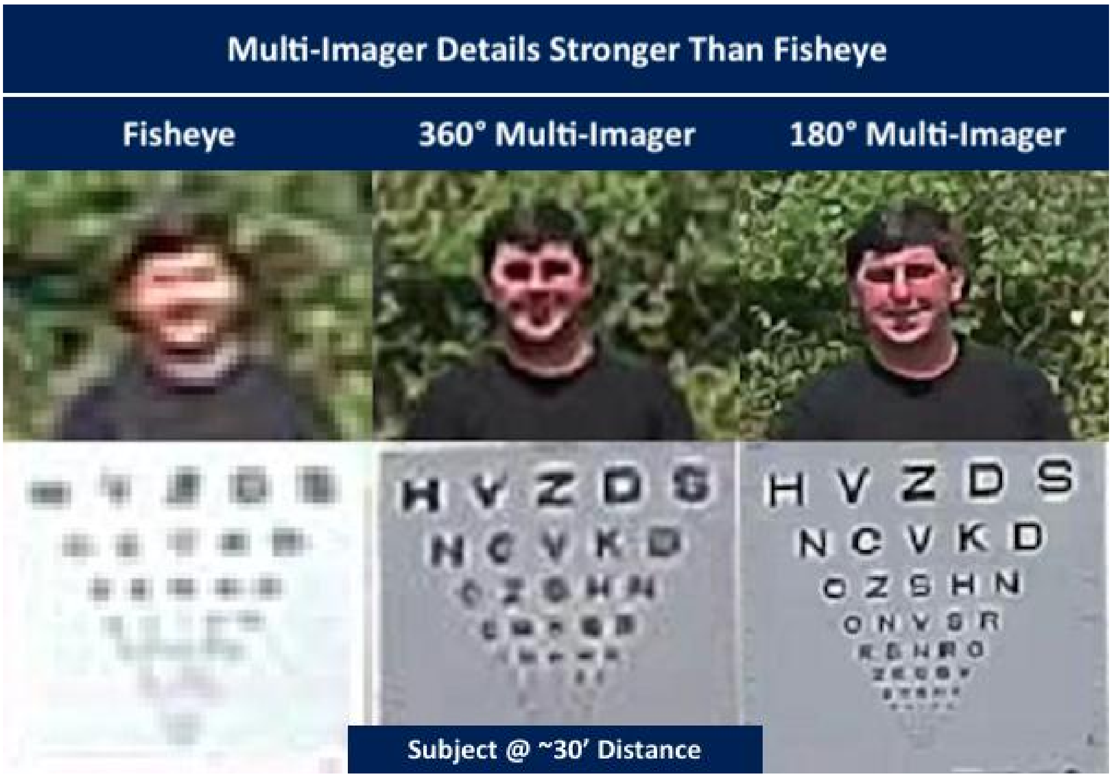

For more on fisheye performance details, see our [Fisheye Panoramic Camera Guide](#).

Repositionable Multi-Imager Camera

Multi-imager usage has been growing strongly over the past few years.

Now, a new segment of multi-imagers has taken off, with many manufacturers offering repositionable models enabling the user to pan and tilt the cameras to customize the field of views covered.

Below is an example of positioning one such unit:

Note: ***Click here to watch the animated sample on IPVM***

We explain the following key elements in repositionable multi-imagers:

- Number of imagers
- Fixed vs. varifocal vs. motorized lenses
- Positioning/panning differences
- Motorized positioning
- Tilt limitations
- Corridor mode support
- Size tradeoffs
- Resolution and FPS options
- Advanced features (WDR, super low light, IR)
- Comparison to single imager cameras
- Comparison to fixed lens multi-imager cameras
- Comparison of Arecont, Avigilon, Axis, Hanwha, and Vivotek repositionable offerings

Repositionable Multi Imager Introduction

Repositionable multi imagers consist of multiple camera assemblies (lens and sensor) mounted in the same housing, each typically with its own 2 or 3 axis gimbal for individual positioning.

There are three key variances in repositionable model coverage:

- Number of imagers: Most commonly, four imagers are used, but three and two imager models are available from some manufacturers.

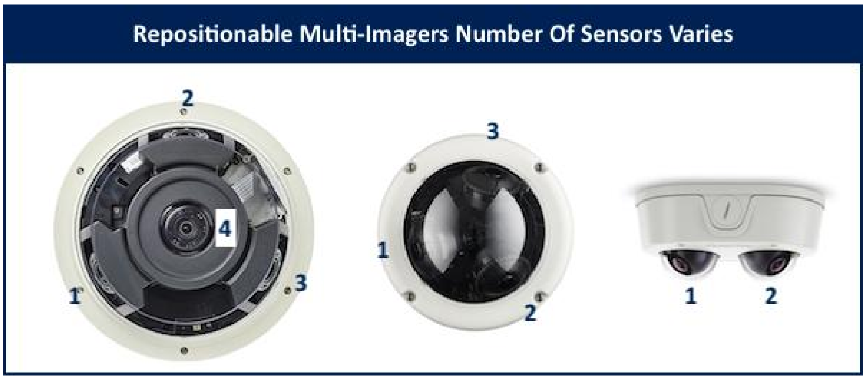

- Fixed or varifocal: Repositionable multi-imagers may be either varifocal or fixed focal (with interchangeable lenses). Varifocal models are growing in popularity in the past ~12 months.
- Motorized zoom/focus: Additionally, some manufacturers include motorized focus and zoom in their repositionable models, allowing users to aim cameras in the field and fine focus/zoom remotely.

Imager Panning Differences

How individual imagers are positioned also varies, with two broad categories:

- Freely movable: Some models allow imagers to be detached from the base and reattached elsewhere, or freely moved around a track without hard limits aside from the length of cable attaching the imager to the base.

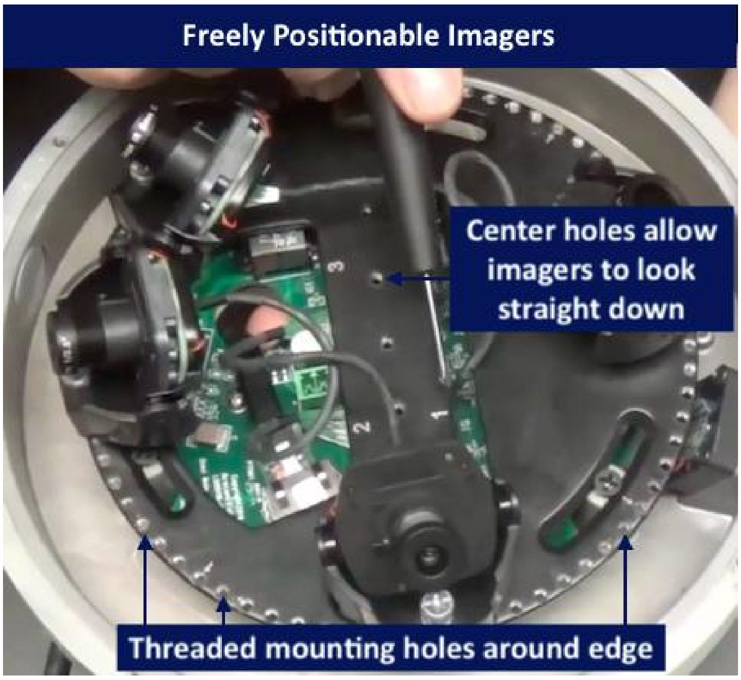

- Limited pan: Others allow only relatively small adjustments in panning, limited either by hard stops or close proximity of other imagers.

The actual mechanism attaching imagers to the base camera varies as well, with some using a magnetic ring with screw points on it to secure imagers, while others use a fixed track around which imagers slide, shown in examples above.

Finally, motorized imager panning has also become available, with users able to change the position of each imager from the camera's web interface, without needing to physically open and adjust the camera. However, this feature is still rare.

Note: Click here to view the animated sample on IPVM

Imager Tilt/Rotation Differences

Unlike panning mechanisms, which tend to vary significantly, multi imager models tend to have similar downtilt and rotation capabilities. However, three key factors impact selection:

- Straight Down Tilt: Not all models allow users to angle an imager straight down, which may be useful when attempting to cover a building corner or T intersection.
- Uptilt limitations: Some multi imagers models are limited in how far up their imagers may be angled. Users should beware of this when selecting cameras as it may impact coverage at moderate to long ranges.

- Corridor mode support: Some models allow individual imagers to be set to corridor mode, rotated 90° or 270°, using a 3:4 or 9:16 aspect ratio. This allows for potentially better coverage of long hallways, fencelines, and other areas.

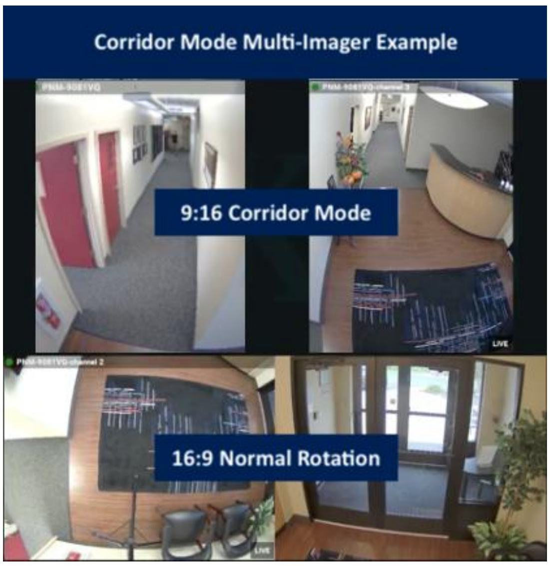

Bigger Than Typical Multi Imagers

Because they include multiple sensors, repositionable multi-imager models tend to be very large compared to typical dome cameras, as well as fixed 180°/360° multi-imagers.

For example, the size comparison below shows repositionable models on the top row, with fixed 180° models below, and a typical size outdoor dome for reference bottom right.

The largest multi-imager model is nearly twice the diameter of the standard dome.

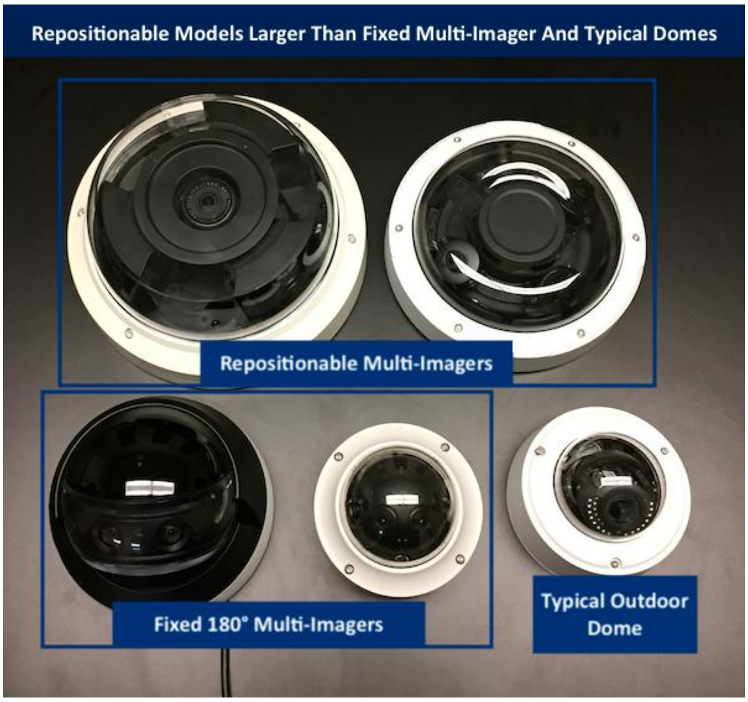

Tradeoffs Versus Single Imager

Compared to typical single imager models, repositionable models have three key advantages:

- Reduced footprint / overall visibility: Where multiple single imager fixed cameras may be required for proper coverage, a single repositionable model may cover the same area, reducing the number of devices mounted to the ceiling/wall and improving aesthetics. However, note that the increased size of some models may not be preferable to some users, even where number of cameras may be reduced.
- Reduced mounting labor: Because only one device must be mounted and cabled, fewer cables must be run, fewer penetrations made, and fewer mounting anchors installed, reducing costs.
- Possibly reduced VMS licenses: Finally, many multi-imager models require only a single VMS license for all channels, where multiple fixed cameras would use multiple licenses, reducing licensing costs. Note that this is not true of all models and VMSes, so users should check license requirements with their preferred camera/VMS combination.

However, there are two notable limitations:

- Advanced features less common: Advanced features such as high end true WDR, super low light capabilities, and others are not common in repositionable multi-imager models, or in the case of integrated IR, simply unavailable. Installations requiring these options should consider single imager models.
- Lower FPS: Multi imager models often support a maximum of 15-20 FPS on all channels or lower, while single imager models are generally capable of 30fps.

Tradeoffs Versus 180/360 Multi Imagers

Repositionable models' key advantage versus 180° and 360° multi imagers is flexibility, as they may be configured to cover odd shaped areas which may not be properly served by 180/360 models.

However, this flexibility comes with two disadvantages:

- More complex setup: Unlike fixed imager models which ship aimed and focused, users must manually aim and focus each individual sensor in repositionable models, adding to installation time.
- Larger size: As mentioned above, repositionable models are larger than fixed imager models and may not be preferable aesthetically.

Repositionable Camera Manufacturer Comparison

The charts below compare respositionable multi imager models from 5 common manufacturers, showing support or lack thereof for features mentioned above.

Basic Comparison

Shown here, Arecont Vision offers notably more model variations than any other manufacturer, with three generations of the SurroundVideo Omni, as well as the MicroDome Duo, available with 1080p, 3MP, and 5MP imagers. Others are more limited, generally in only one resolution, with the exception of Hanwha.

MANUFACTURER	MODEL	NUMBER OF IMAGERS	IMAGER RESOLUTION	POSITIONING TYPE	LENS TYPE
Arecont Vision	MicroDome Duo	2	1080p, 3MP, 5MP	N/A	Fixed
Arecont Vision	SurroundVideo Omni	4	1080p, 3MP, 5MP	Flexible	Fixed
Arecont Vision	SurroundVideo Omni G2	4	1080p, 3MP, 5MP	Flexible	Fixed
Arecont Vision	SurroundVideo Omni G3	4	1080p, 3MP, 5MP	Flexible	Motorized
avigilon	HD Multi Sensor	3 or 4	3MP	Restricted (~90° pan)	Motorized
AXIS COMMUNICATIONS	P3707-PE	4	1080p, 3MP	Restricted (~90° pan)	Manual Varifocal
Hanwha	PNM-908xVQ	4	1080p, 5MP	Flexible	Motorized
VIVOTEK	MA8391-EV	4	3MP	Flexible	Motorized

Advanced Features

The second chart shows advanced features such as WDR, remote zoom/positioning, and super low light. Three key points stand out here:

- True WDR common: The majority of manufacturers offer true WDR capability in their multi-imager models. Axis and Vivotek do not.
- Super low light and remote positioning rare: High-end features like remote positioning and super low light are available only on one manufacturer each (Arecont and Hanwha, respectively).
- Axis lacking advanced features: Axis' first repositionable model lacks any of the advanced options found in others, with no WDR, remote zoom, positioning, or super low light.

MANUFACTURER	MODEL	TRUE WDR	REMOTE ZOOM	REMOTE IMAGER POSITIONING	SUPER LOW LIGHT
Arecont Vision	MicroDome Duo	✓*	✗	✗	✗
Arecont Vision	SurroundVideo Omni	✓*	✗	✗	✗
Arecont Vision	SurroundVideo Omni G2	✓*	✗	✗	✗
Arecont Vision	SurroundVideo Omni G3	✓	✓	✓	✗
avigilon	HD Multi Sensor	✓	✓	✗	✗
AXIS	P3707-PE	✗	✗	✗	✗
Hanwha	PNM-908xVQ	✓	✓	✗	✓
VIVOTEK	MA8391-EV	✗	✓	✗	✗

These charts will be updated as new models become available.

IPVM Tests

IPVM has tested several of the above models, comparing image quality, positioning capabilities, advanced features, and other aspects:

- Arecont SurroundVideo Omni (G1)
- Arecont MicroDome Duo
- Avigilon HD Multisensor
- Axis P3707-PE
- Hanwha PNM-9081VQ

Streaming & Recording

Surveillance Codec

Codecs are core to surveillance, with names like H.264, H.265, and MJPEG commonly cited. How do they work? Why should you use them? What issues may you face? In this tutorial, we examine this in-depth covering:

- Uncompressed vs. Compressed Video
- Inter vs. Intra Frame Compression
- I vs P Frames
- H.264 vs H.265 vs MJPEG
- H.265 Emergence
- Smart Codec Growth
- Proprietary Codecs
- Scalable Codecs - JPEG2000, SVC
- Quality of Codecs
- Codec Support in Surveillance
- Future Codecs
- What Codecs to Choose?

Uncompressed vs. Compressed Video

Essentially all surveillance video is compressed, as storage and network demands would easily be 100x greater for uncompressed.

When video is digitized, it is initially uncompressed. There are 3 main factors in the size of uncompressed video:

- The range of values supported for each pixel
- The total number of pixels per frame
- The total number of frames per second

To find the total size, you multiply each of these. Let's walk through each of these.

Pixel Values

Each pixel is given a value represented by a number within a range. The range determines how precisely the color can be defined and also greatly impacts the bandwidth/size.

- Take grayscale, which frequently has a range of 256 values (8 bits) with 0 representing black, 255 representing white and the numbers in between representing shades of grey.
- Of course, almost all surveillance video today supports color so the range of values needed to represent all the colors is far greater. 16 bits or 65,536 values is common.

Resolution/Framerate Values

The other two factors are far easier to understand as they are the resolution of the the camera, multiplying horizontal times vertical pixel counts:

NAME / PIXEL COUNT	HORIZONTAL X VERTICAL
VGA / .3 MP	640 x 480
720p / 1 MP	1280 x 720
1080p / 2 MP	1920 x 1080
3 MP 4:3	2048 x 1536
3 MP 16:9	2304 x 1296
4 MP	2688 x 1520
5 MP 4:3	2592 x 1944
5 MP 16:9	3072 x 1728
6 MP	3072 x 2048
4K / 8 MP	3840 x 2160
12 MP	4000 x 3000

As well as frames per second, simply 1fps, 10fps, 30fps, etc. The vast majority of surveillance cameras (>80% in our surveys) are recorded at between 5 and 15 FPS (see Average Frame Rate Video Surveillance 2016).

Adding It Up

What is critical here is recognizing how massive uncompressed video can become. Take a 1080p color camera at 30fps. Uncompressed it is the product of the following:

Note: Click here to view the animated formula on IPVM

Multiplying those three factors results in ~1 Gb/s for uncompressed 1080p/30fps video. In a day, at that rate, you would have ~12TBs of storage from a single stream. While hard drives continue to get bigger, a 16 camera uncompressed surveillance system storing for 30 days, would need nearly 6PBs of storage across a few hundred hard drives that would cost hundreds of thousands of dollars.

Codecs Are Key

Despite the huge size of uncompressed video, production surveillance systems of 16 cameras routinely fit in standard PCs or even small appliances. How does this happen? Codecs.

Codec stands for compression / decompression and the act of compression is the central element of reducing bandwidth / storage consumption.

Giving every pixel of every frame its own unique value is incredibly wasteful as most scenes are filled with a small number of similar colors. Codecs compress video by reducing the number of values recorded while tracking which pixels have the same or similar values, allowing it to transmit far lower amounts of bits.

Inter vs. Intra Frame Compression

There are two fundamental approaches to compression: intra-frame and inter-frame. It is critical to understand the distinctions between the two as they impact bandwidth consumption, processing power requirements, and quality risks.

- Intra-frame compression is within a single frame only but NOT across frames (example - MJPEG)
- Inter-frame compression is across multiple frames AND within single frames (examples - H.264, H.265, MPEG-4)

All codecs support intra-frame compression but some only some support both intra and inter frame compression.

Intra-frame Compression

Intra-frame compression only looks at one frame at a time, doing its best to compress what is in that image. Even though video is a series of images, intra-frame compression sees only a frame at a time, ignoring the "stream." An intra-frame codec stream's individual frames look like the frames of a movie:

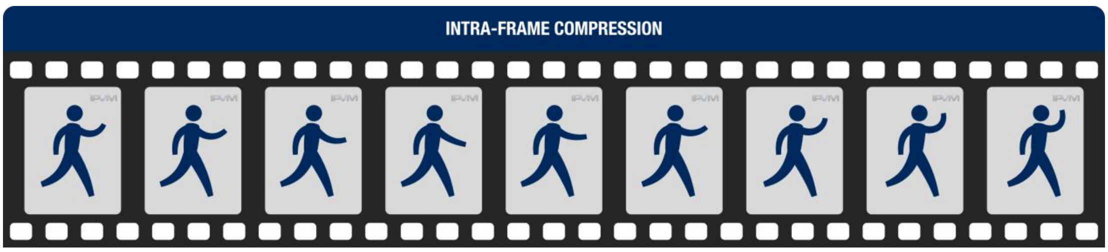

The upside is that this is simple to do computationally and does significantly compress video. For instance, a 1080p/30fps video stream using an intra-frame compression like MJPEG can have its bitrate reduced from ~1000 Mb/s to ~40Mb/s. However, the downside is that even more savings are available if you compare between frames.

Inter-frame Compression

Using inter-frame compression, not only is the video within the frame encoded, the codec compares adjacent frames to further compress the image. This is feasible because often very little changes from one frame to the next. For example, using the same scene of a person waving, inter-frame compression would send only the subject's arm.

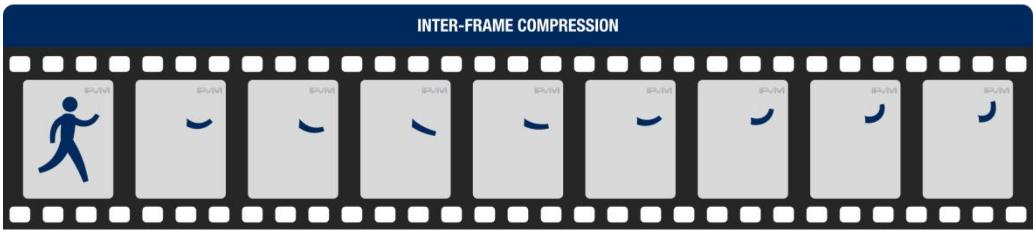

With so much of a given scene remaining static, sending only changes in the scene saves substantial bandwidth/storage. For instance, the same 1080p/30fps stream that might need 40Mb/s with MJPEG, an intra-frame only codec, may only need 4Mb/s with H.264, a codec that uses both.

However, the main downside to inter-frame compression is that it is far more computationally intensive, which can increase performance and quality risks (examined later).

I vs P Frames

There are two key frame types in inter-frame compression such as H.264/H.265.

I-frames

The first frame in a given group of pictures is called an I-frame, short for intra-coded, and is essentially a full frame of video, as opposed to only changes found in P frames. The distance between two I-frames is referred to as I-frame interval, GOV (group of video), or GOP (group of pictures).

As an example, this image shows the I-frame of an outdoor scene used in our testing (click for full size):

P-frames

P-frames reference the full image of the previous I-frame to send only the changes in the scene. Changes may be small, such as digital noise or small foliage movement, or large, such as a PTZ camera moving from preset to preset. The P in P-frame stands for "predictive."

The changes sent in a P-frame from the test scene can be seen in the image below (click for full size). Only the areas near the road, where cars passed, and moving foliage to the right, are sent as changes.

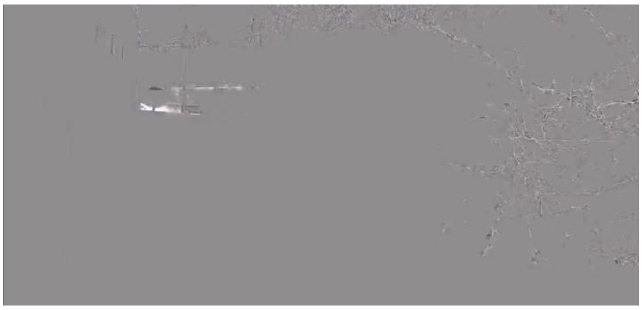

Putting these together in the clip below, the P-frame changes can be seen "stacking" upon each other, until the next I-frame is sent (~15 seconds).

Note: Click here to watch the video on IPVM

Other Frame Types

Note that there are other frame types in addition to I and P, such as B, SI, and SP, though they are virtually unused in surveillance. Some IP cameras include B-frame support, but not all VMSes are capable of properly decoding them, so they are generally not used.

Standard Codec Usage

In 2018, most new deployments use H.264 because of its bandwidth/storage benefits compared to MJPEG and due to barriers to H.265 adoption (see below).

MJPEG is still used in some deployments, typically when required by bid spec or for specialist applications such as LPR or other analytics. However, manufacturers have begun to limit support for MJPEG, with some models not providing an MJPEG stream.

JPEG2000 is an inter-frame codec similar to MJPEG, but scalable (see below). It was best known for its use by Avigilon, but has been near totally phased out in their H3/H4 lines (only Avigilon LPR cameras still use JPEG2000). Others, such Ampleye's NOX-20 and Logipix also use JPEG2000.

Proprietary Codecs

Surveillance has had a handful of proprietary codecs, but very few of these are in use in 2018. In older analog systems, video was encoded, stored and managed in the same appliance (i.e., a DVR), making it easier to use a proprietary codec as the manufacturer controlled the entire end to end process.

However, IP cameras encode video by themselves and then need to transmit this encoded video to a recorder/VMS for storing and managing it. Proprietary codecs increase the complexity of storing, managing, and displaying video as they must each be integrated to the VMS. The high costs of doing this for multiple proprietary codecs have motivated most IP camera manufacturer codecs.

Mobotix MxPEG Exception

Mobotix's MxPEG is the most widely known proprietary codec used in video surveillance. Introduced in 2000, MxPEG supports inter-frame compression, improving upon MJPEG and reducing bandwidth consumption compared to it (see our tests). However, since it is proprietary, very few 3rd party VMSes support it, forcing most to either use Mobotix's own VMS or to set Mobotix cameras to MJPEG. Furthermore, H.264's bandwidth consumption is typically much lower than the equivalent MxPEG stream, further reducing its attractiveness.

Smart Codecs

In the past 1-2 years, H.264 (and H.265) smart codecs have become common, aiming to further reduce bitrates compared to standard H.264. Exactly how these smart codecs function varies, but generally they use two techniques, reviewed below. Readers should see our Smart Codec Guide for full details of these technologies.

Smart Compression

Instead of applying the same compression level to the entire scene, smart codecs dynamically adjust compression for activity in the camera's FOV. For example, looking at the image below, compression could be set to "low" for the speaking person to keep quality high, but the white background may be set to "high", since we do not need details of the white wall.

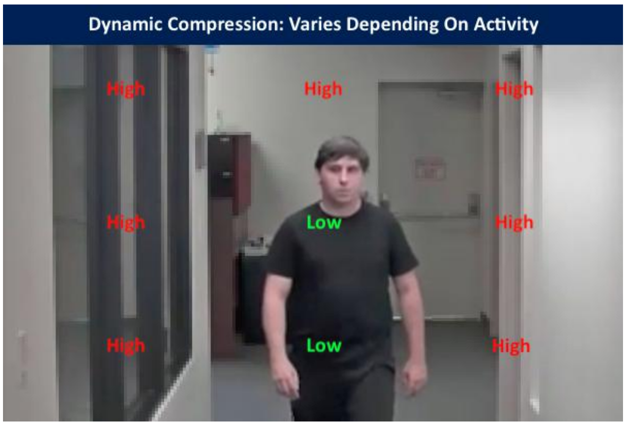

Dynamic I-frame Interval

Second, smart codecs typically dynamically adjust I-frame interval based on activity in the scene. So if a scene has little or no motion, the camera sends I-frames infrequently (5, 10, 20 seconds or more), but when activity is detected in the scene, it immediately sends an I-frame and switches back to normal (typically 1 second) I-frame intervals as long as activity continues.

This image from a stream analyzer shows the effects of dynamic I-frame interval:

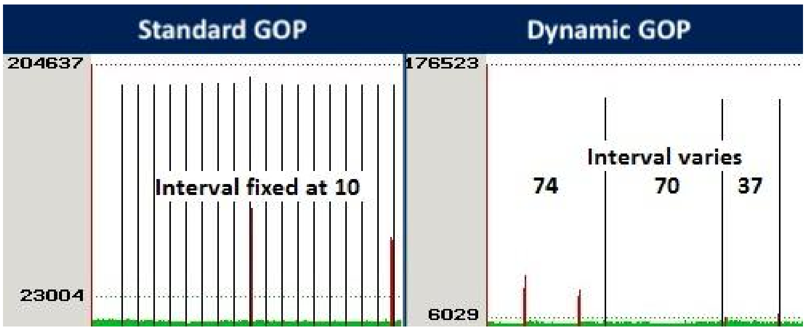

Benefits

In our tests, Smart codecs have reduced bitrates significantly, ~15% at a minimum, but up to 95%+ in some scenes. These codecs are most effective in still scenes, as I-frame intervals remain longer and compression higher due to the lack of activity. See our Smart codec Guide for more details.

H.264 vs. MJPEG Quality

While H.264 is clearly the most widely used codec in surveillance, for many years a fierce debate existed about quality loss compared to MJPEG, and some still do believe H.264 is inferior. Our extensive H.264 vs MJPEG testshows that, when properly configured, H.264 delivers the same visible quality as MJPEG.

However, there are some issues that can undermine quality, typically:

- Setting high compression levels: If the compression levels are set too high, the video will have quality degradation. See our IP Camera Manufacturer Compression Comparison for more details and how specific manufacturers handle compression.

- High complexity scenes: Such issues often are visible in complex scenes, where there is lots of action - intersections, crowds, etc.
- Setting low constant bit rates: If CBR is used but the bit rate is set too low, the video will have quality degradation.

However, in most cases, even when using default settings, H.264 has delivered quality similar to MJPEG in our tests.

Scalable Codecs

Most codecs can only support a single resolution. For instance, changing from a 2MP stream to a 1MP one using MJPEG or H.264, requires either encoding a completely new stream or transcoding (i.e. re-processing) the 2MP stream to make it 1MP. However, in some cases you want to change the stream resolution without requiring a new stream, for instance, if you are sending to a client over a lower bandwidth connection (i.e., mobile) or want to reduce storage size of older video.

A certain class of codecs, called 'scalable', can do this automatically, without having to request a new stream or re-process an existing stream. A scalable codec can essentially 'pick' the frames or resolution levels they want out of a stream. This allows pruning frames or resolution over time as well as dynamically adjusting resolution / fps for remote / mobile clients.

There are two well known scalable codecs:

- SVC is H.264 plus scalability. The plus side is that it combines scalability with the bandwidth benefits of H.264. Unfortunately, very few manufacturers support this. Most accomplish this through multi-streaming (sending multiple streams simultaneously of different resolution/frame rates).

- JPEG2000, which is essentially MJPEG, with the addition of scalability. This is the codec that Avigilon traditionally used but has phased out in new cameras. The huge downside of JPEG2000, like MJPEG, is massive bandwidth/storage increase compared to H.264, only partially offset by pruning.

SVC development has been very slow, with little actual adoption despite availability in some cameras (at least on paper) for several years. With the bitrate advantages of H.264 smart codecs (which require little to no new VMS development) and H.265 taking precedence, SVC is unlikely to become a major factor in surveillance codecs.

H.265 Emerging But Still Limited

For the past few years, H.265 has been the next big codec, promising to replace H.264 and reduce bitrates by another ~50%. However, there are key barriers preventing its widespread implementation:

- Limited gains: While the change from MJPEG to H.264 resulted in drastic bitrate reductions, often 50-75% or more, our tests of H.265 show that 15-30% savings over H.264 is more likely. Given this fact and the rapid increase in size/decrease in cost of hard disc drives, the benefits of H.265 are not as compelling as its predecessor.
- Development required: Second, VMSes must implement H.265 decoding, a non-trivial development expense, though with limited benefit due to the small number of current H.265 models and limited bandwidth/storage savings.
- No ONVIF conformance until 2018: Finally, moving to H.265 cameras means giving up true ONVIF profile conformance until their next profile is completed in 2018. Camera manufacturers may implement

H.265 using ONVIF's 2.4 spec, but there will be no conformance test for it until this profile is complete.

We expect H.265 to continue to develop as it matures, but these reasons combined make it unlikely in the next 12 months, at least.

See our tests of H.265 IP Cameras Tested vs H.264 and Smart H.265 Testfor more details.

Future Codecs

While many alternative codecs are discussed or pitched, it is highly unlikely that any non standardized ones will gain wide adoption in surveillance.

For instance, some developers claim massive bandwidth savings from their proprietary codecs, such as Digital Barriers TVI / EdgeVis. However, this would require both camera manufacturers and VMS developers to implement these codecs in their products. Additionally, given the significant bitrate reductions of smart codecs, which are compatible with the majority of current H.264 recorders/VMS, the investment in development to add these proprietary codecs is even more unlikely.

Some have mentioned Google's VP8 and VP9 codecs as possibilities in surveillance, especially as a royalty-free alternative to H.265 (which has since reduced its license fees). However, these codecs saw little to no interest from manufacturers, with no surveillance chipsets moving to adopt them, and have thus gained no traction.

What Codecs to Choose?

Continuing into 2018, the best codec combination for most use cases is H.264 with smart codec support as smart codecs significantly improves efficiency of H.264 with minimal downsides. H.265 may be considered but will be ruled out by many due to its various limitations. MJPEG remains a niche for specialist applications those who fear (typical unreasonably) loss from inter-frame compression.

[NOTE: This was originally published in 2013, but substantially re-written in 2017 primarily to reflect technology advances in H.265 and smart codecs.]

Video Quality / Compression

While CODECs, like H.264, H.265, and MJPEG, get a lot of attention, a camera's 'quality' or compression setting has a big impact on overall quality. In this training, we explain what this level is, what options you have and how you should optimize it.

To start, review these two images, (A) and (B):

If your gut feel is that this is a trick question you are right. With the information presented, the best answer is likely that it cannot be determined. In this case, technically the correct answer is 'neither - they are the same resolution'. We used the same camera for each image and simply lowered the quality level for the 'B' image (while keeping everything else the same, including resolution - 720p - and CODEC - H.264).

The fact that two exact shots with the same resolution can look significantly different has a number of important implications. Inside, we explain why, covering:

- Quantization levels
- Bandwidth vs. quality loss
- Image quality examples

- Manufacturer differences
- MBR/VBR/CBR impact
- Smart codec impact
- Recommendations

Quantization Levels

Regardless of codec used (MJPEG, JPEG2000, MPEG-4, H.264), all IP cameras offer quality levels, often called 'compression' or 'quantization'.

H.264 quantization is a measured on a standard scale which varies from 0 to 51, with lower numbers meaning less compression, and thus higher quality. If this seems counterintuitive to you, it is understandable, but these measurements were agreed upon in H.264 standards.

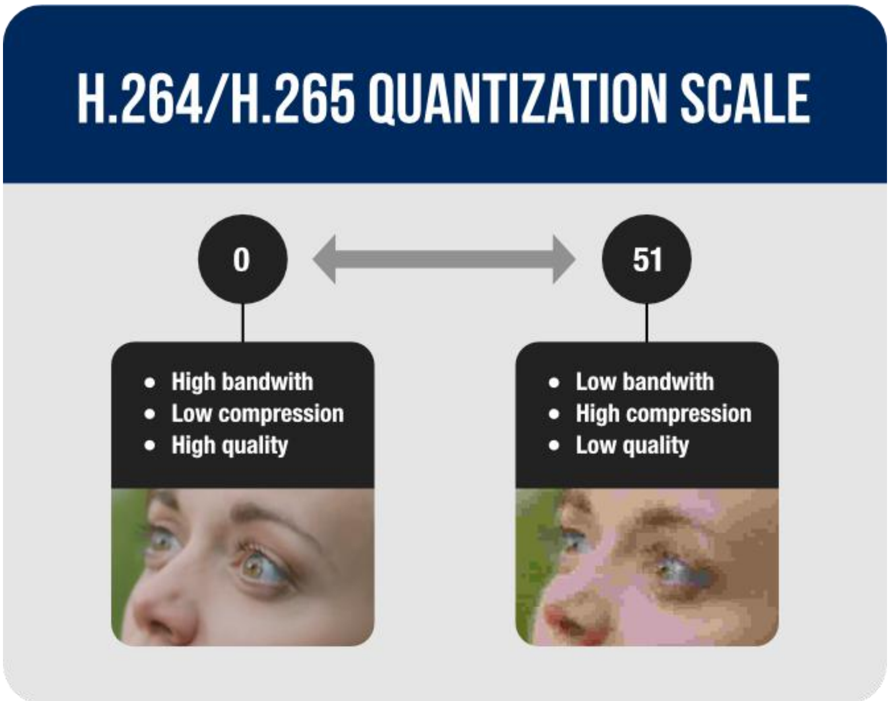

Key Tradeoff: Bandwidth Vs. Quality Loss

The key tradeoff in setting quantization is determining how much 'loss' you are willing to accept for a particular decrease in bandwidth. Recall that all production surveillance video compression is 'lossy', meaning that some information will be lost when video is compressed.

The decision which must be made is how much information loss is acceptable. Increase compression/lower the quality level and you save on bandwidth, but reduce quality. Increase quality/reduce compression and you may gain usable details, but use more bandwidth and reduce storage time.

Demonstrated in Pictures

The image below shows the impact of changing quality levels. In order to show this, we took examples from two scenes (indoor / simple and outdoor / complex) and adjusted bitrate. The Average DRF (quantization) levels below, underlined in red can be seen decreasing in both scenes as bitrate goes up. Inversely, image quality increases, with less pixelation and blocking on objects. Note that we used a CBR stream to simply show the effects of bitrate.

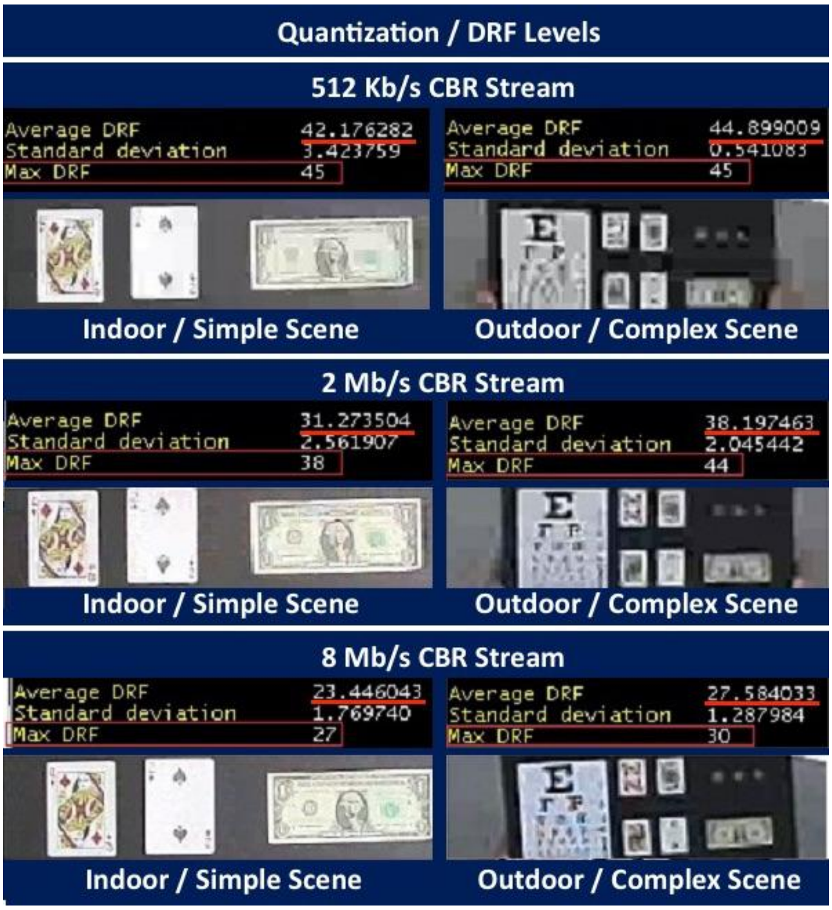

Manufacturer Approaches

All manufacturers set default quality levels. Even if you never touch these settings, the manufacturer will make a choice for you. However, with the vast majority of professional cameras, configuration options are available to let advanced users adjust this.

Below, we look at a few manufacturers to better demonstrate their approaches. Readers should see our [IP Camera Manufacturer Compression Comparison]() test for full details of 10+ manufacturers compression scales and how to default these cameras to "average" compression.

Axis

Axis currently uses two different web interfaces, as many of their cameras have not yet transitioned to their new HTML5 based UI. Both new and old web UIs refer to quality as "Compression", with higher numbers being more compressed/lower bandwidth. Beware of this, as inexperienced users may simply increase the scale thinking it increases image quality.

The new web interface uses a slider/manual entry:

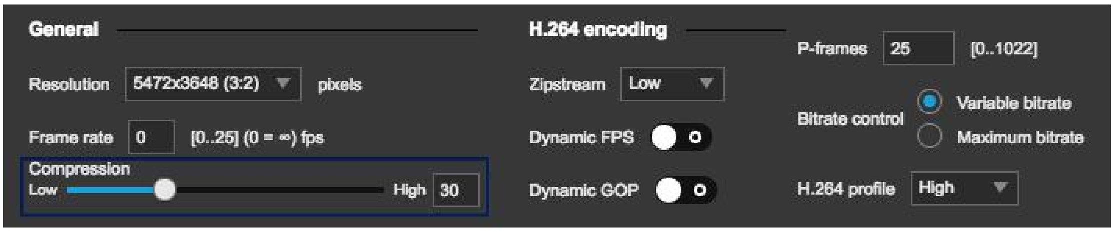

While the old interface simply allows free entry:

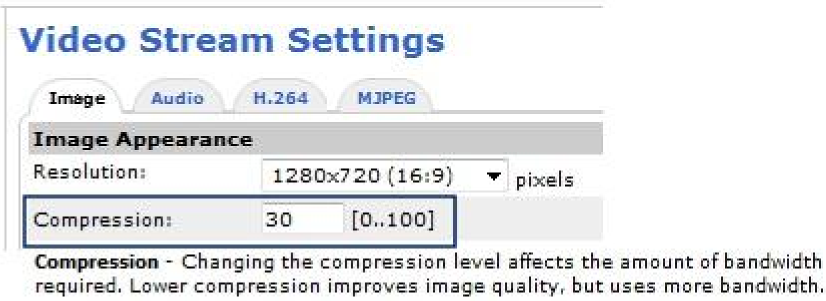

Regardless of UI, Axis Compression 30 equates to ~28 on the quantization scale.

Bosch

Bosch's codec setup bases quantization on P-frames, with I-frames allowed to vary by a specific amount ("I/P-frame delta QP" below). So if P-frames are set to 25 and I-frames set to -6, I-frames will use a minimum of 19 quantization.

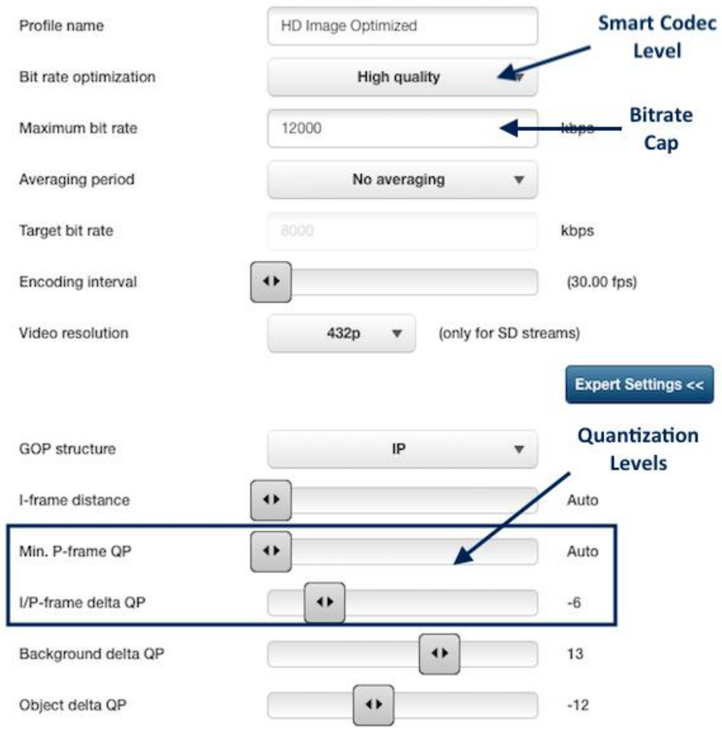

Hikvision

Hikvision's compression settings are under the video/audio tab, on a scale of "lowest" (most compression) to "highest" (least compression), defaulting to "Higher." These settings are mapped to specific quantization levels, ranging from ~20-35.

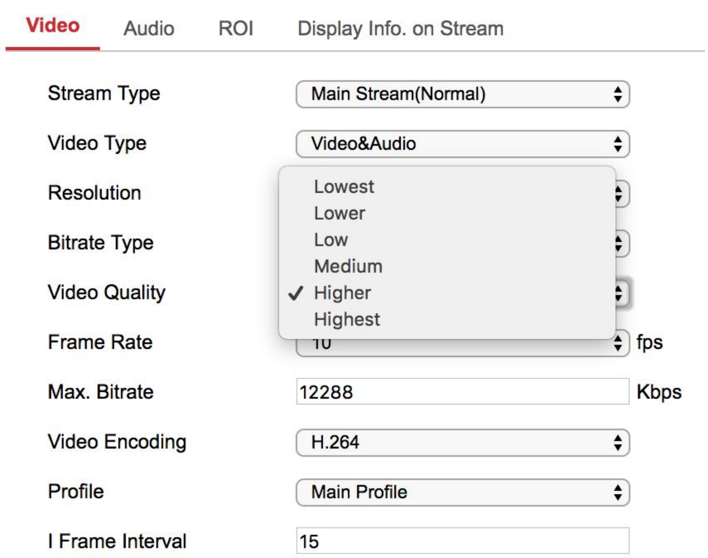

Important: MBR/VBR/CBR Impact

Most surveillance deployments use variable bit rate streaming (VBR), which varies bandwidth in order to maintain a pre-set quality level, sometimes with a maximum bit rate cap (called MBR). The levels set above define the quantization level the camera is targeting, e.g., an Axis camera will target quantization level 28, with bitrate increasing and decreasing as necessary.

However, if cameras are set to constant bit rate, quality level is NOT configurable because the bit rate, by definition is 'constant', with quality level automatically adjusted to keep bandwidth the same. For instance, if you have a CBR video stream set at 1Mb/s looking at a white wall, the camera may use a 'high' quality level since it is easy to compress. However, if the lights are turned off and a hand is waved in front of the camera, it may need to drop down to 'low' quality to maintain the same 1Mb/s constant bit rate as compression is more difficult.

For more on the key points of these streaming modes, see our CBR vs VBR vs MBR - Surveillance Streaming tutorial.

Smart Codec Impact

Further complicating things, in the past few years, camera manufacturers have introduced smart codecs, which allow image quality to vary based on activity in the scene. So a person or vehicle moving through a parking lot will be higher quality, while parked cars and trees in the background are lower quality. Generally speaking, these codecs allow for similar or better image quality compared to standard (non-"smart") codecs, but at much lower bitrates, with 25-30% lower bandwidth common, and as much as 90% lower possible.

However, we strongly recommend readers test smart codecs in the scene where the camera will be installed, as some negative image quality effects may be seen. For example, in one of our tests, higher levels of smart codec settings caused blurring/smearing/artifacts in the scene, visible in the subject and background below. Settings should be carefully adjusted to avoid these issues.

Note: Click here to watch the animated comparisons on IPVM

Readers should also see our Smart Codec Guide for more information on these and other considerations.

What Should You Use?

Ultimately, the most important question is: what quality setting is necessary in the camera's application?

Unfortunately, it is impossible to give a universal answer as:

- The right level is a subjective judgment call. Changing from a quantization level of 30 to 29 or 42 to 41 produces no magical difference. Often the changes are nearly imperceptible and debatable.
- The right level depends on the complexity of the scene. More complex scenes (like an intersection) will typically benefit from a higher quality levels than a simple scene (like a stairwell). The smaller the elements being observed (like a person across an intersection), the more higher quality levels can capture meaningful details.

However, based on our tests, we offer two key recommendations:

- Quantization ~28 is the "sweet spot": Based on years of ongoing testing, we have found 28 to offer the best tradeoffs between bandwidth and image quality. Many manufacturers default to 28-30, though users should see our IP Camera Manufacturer Compression Comparison for more details on how to standardize others.
- Use smart codecs (but carefully): Because of their drastic bitrate reductions in many scenes, we smart codecs be used wherever possible, as users may be able to increase image quality where necessary, while bitrates remain lower than typical codecs.

Smart Codecs

In 2018, smart codecs are now mainstream. Once seemingly a marketing buzzword, now the majority of manufacturers offer smart codecs on at least some of their cameras.

These marketing names vary, including 'Zipstream', 'Smart Coding', 'H.264+', 'Smart Stream II', and others, and critically, these implementations and bandwidth savings vary dramatically.

We explain what smart codecs attempt to do and the most common implementations, covering:

- Historic static compression, I-frame interval, and FPS techniques
- Smart codec basic
- Dynamic compression
- Dynamic I-frame interval/GOP
- Dynamic FPS
- Static compression regions
- Intelligent DNR
- VMS/NVR compatibility
- Manufacturer support
- IPVM test recommendations
- Bandwidth risks
- Smart codec outlook

To understand this, you must have a good understanding of codecs, compression and bandwidth variations. Please first review our:

- Surveillance Codec Guide

- Video Quality / Compression Guide
- Bandwidth Guide For Video Surveillance

'Normal' codecs, like 'regular' H.264, set one compression level, one I frame interval, and one frame rate. 'Smart' codecs change one, two, or all of those, reviewed inside.

"Normal" Codec Review

In normal codecs, three key parameters are typically set to fixed values:

- Compression level (sometimes called quality or quantization)
- I frame interval (sometimes called GOP or GOV)
- Frame rate (FPS)

Even if you are not aware of these, they exist with manufacturers making their own decisions on the defaults (e.g., see IP Camera Manufacturer Compression Comparison).

Fixed Compression Effects

'Normal' codecs set a single compression level for the entire video, regardless of what is being displayed. For example, if we look at a scene of a hallway, both the moving subject as well as the static walls, floor, door, and other background areas will be compressed at the same level in each frame:

Fixed I-frame Interval Effects

The most fundamental difference between today's codecs (H.264/H.265) and earlier codecs like MJPEG is that they do not send the 'same' image over and over again. Instead, a single full frame is sent (called an 'I-frame'), followed by only small updates of the areas of the scene which have changed, called 'P-frames.' In standard codecs, the I-frame interval is fixed. This means that the camera will always generate an I frame periodically, most commonly 1 second. In this scenario, if the stream is 30 FPS, it will send 1 I-frame followed by 29 P-frames and then repeat.

The downside of a fixed I-frame interval is that the activity in the scene can vary. Sometimes a hallway can be empty for an hour. During the 'empty' time, sending an I frame once per second is wasteful. Remember that I frames consume far more bandwidth than P frames, often on the order of 10x more.

But when the burst of people walk through, I frame interval should be shortened to reduce artifacts and other issues (see Test: H.264 I vs P Frame Impact).

Smart Codec Techniques:

Unlike typical codecs, smart codecs dynamically adjust these parameters, using a combination of three techniques:

- Dynamic compression: Varying compression on various areas of the scene instead of the entire frame.
- Dynamic I-frame interval: Varying I-frame interval depending on motion in the scene.

- Dynamic framerate: Varying FPS depending on motion in the scene.

We look at these in more detail below.

Dynamic Compression

The first technique used is dynamic compression, which adjusts quantization level for different parts of the scene based on activity. For example, instead of setting compression on 'medium' for the whole scene, as shown in the hallway scene above, the camera adjusts compression on the moving subject to low (higher quality), to maximize capture quality, and increases compression on static background areas.

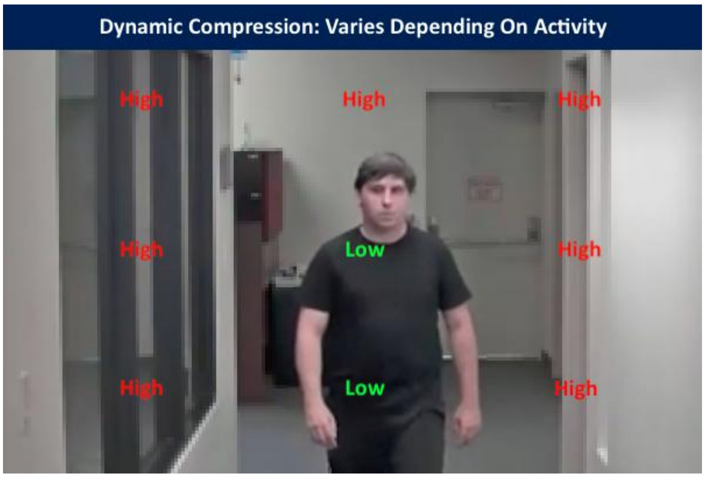

By adjusting compression levels for parts of the scene, it is possible to both reduce overall bandwidth consumed and increased quality in areas which matter by more intelligently assigning compression levels to relevant parts of the scene.

Dynamic I-Frame Interval

Second, smart codecs may adjust I-frame interval based on tracking of activity in a scene. So long as the scene has little to motion, the camera can send an I-frame infrequently (e.g., every 5 or 10 or 20 seconds), but as soon as significant activity is detected, an I-frame is generated, with the interval kept short as long as activity continues.

The image below shows analysis of a standard H.264 stream versus one using dynamic I-frame interval (GOP). While the standard codec remains fixed at an interval of 10, the dynamic GOP varies, increasing to over 70 when there is no motion in the scene. Note that both streams shown below were 10 FPS.

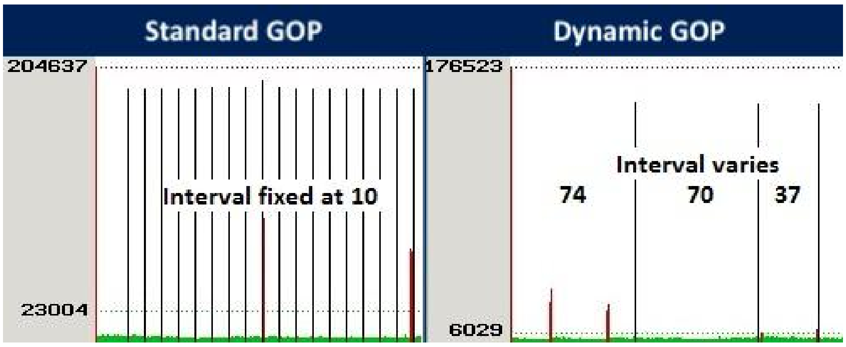

While dynamic I-frame interval will not improve quality, quality should remain essentially the same, while significantly reducing bandwidth. In our tests, dynamic GOP is the largest driver of bandwidth savings in smart codecs.

Dynamic FPS

Though some VMSes/NVRs have included dynamic FPS, called 'Motion Boost Recording' for several years, its introduction in IP cameras is more recent. Like I-frame interval, dynamic FPS reduces bitrate by simply sending fewer frames when there is no activity in the scene.

Note: Click here to watch the animated gif on IPVM

While this may drastically reduce bitrate (bandwidth below 10 Kb/s was possible in our tests), users should beware of reducing frame rate as fast moving objects may be missed in rare cases, such as vehicles at highway speed, running subjects in small areas, etc. Additionally, in some areas regulations may prohibit this frame rate reduction.

Not-So-Smart: Static Compression Regions

Some cameras allow users to set fixed areas as background and foreground. This may reduce bitrate in some cases but introduces additional risks compared to smart codecs:

- Unintentional high compression: If regions are not carefully set based on actual use of the area being observed, moving objects may move into higher compression "background" areas, potentially losing usable details.
- Manually adjusted if conditions change: If the scene changes, users must manually adjust their static regions to properly reflect use of the scene. No adjustment is required when using smart codecs.

Related Smart Technology: Noise Reduction

Normal codecs have trouble with visible noise, common in low light scenes. This noise 'moves' around the screen and typically tricks encoders into thinking it is real movement. Because the camera believes it is movement, it takes more bandwidth (often a lot more bandwidth). See Gain / AGC for Video Surveillance Guide

Increasingly cameras are adding intelligence to distinguish between real movement (a person walking, a car driving, etc.) and visible noise which may reduce bandwidth significantly. For more, see our Camera DNR (Digital Noise Reduction) Guide.

Works With Standard VMSes/NVRs

Because smart codecs work within the standard framework of H.264/H.265, additional support is generally not required by VMSes and NVRs, in contrast to some specialized codecs which claim to reduce bandwidth (such as Mobotix MxPEG or Digital Barriers' EdgeVis) but require custom integration.

Some VMSes may experience issues with jumpy playback or streams not loading, but most VMS developers have improved support since smart codecs' introduction and removed these issues.

Manufacturer Camera Support

Equally important, just because a camera allows for dynamic compression levels, does not mean or require that it allows for smart I-frame intervals or dynamic FPS. Some cameras only support one of the three (most commonly dynamic compression), while others support all three.

The chart below details which manufacturers support which smart codec features:

MANUFACTURER	DYNAMIC COMPRESSION	DYNAMIC GOP	DYNAMIC FPS
Avigilon	✓	✓	✓
Axis Communications	✓	✓	✓
Bosch	✓	✓	✗
Dahua	✓	✓	✗
Hanwha	✓	✓	✗
Hikvision	✓	✓	✗
Panasonic	✓	✓	✗
Pelco	✓	✓	✗
UNV	✓	✓	✗
Vivotek	✓	✓	✗

Manufacturer Performance Tested

IPVM has tested a number of these implementations including:

- Axis Zipstream 2 (With Dynamic FPS)
- Axis Zipstream
- Dahua Smart H.264+
- Hanwha Wisestream
- Hikvision H.264+
- Panasonic Smart Coding
- Uniview U-Code
- Vivotek Smart Stream II Tested (as part of H.265 vs. H.264)
- Vivotek Smart Stream Tested

From these, we have found a few key points:

- Smart I Frame Interval had the biggest impact, with 60%+ savings viable, simply because many surveillance scenes are static for long time periods, enabling great savings from less I frames.
- Dynamic Smart Compression had the next biggest impact, though at much lower bandwidth savings.
- Fixed Smart Compression had the least impact, often at just 10-20% savings, because of limits on what areas to define.

Bandwidth Risks In High Motion Scenes

Though smart codecs may greatly reduce bitrates on average, users should beware when calculating bandwidth needs as high motion scenes may see little to no savings from these advanced codecs. For example, in a high motion roadway scene, bitrates of Axis, Panasonic, and Hikvision cameras were near those seen with smart codecs disabled during periods of high motion, i.e., heavy traffic:

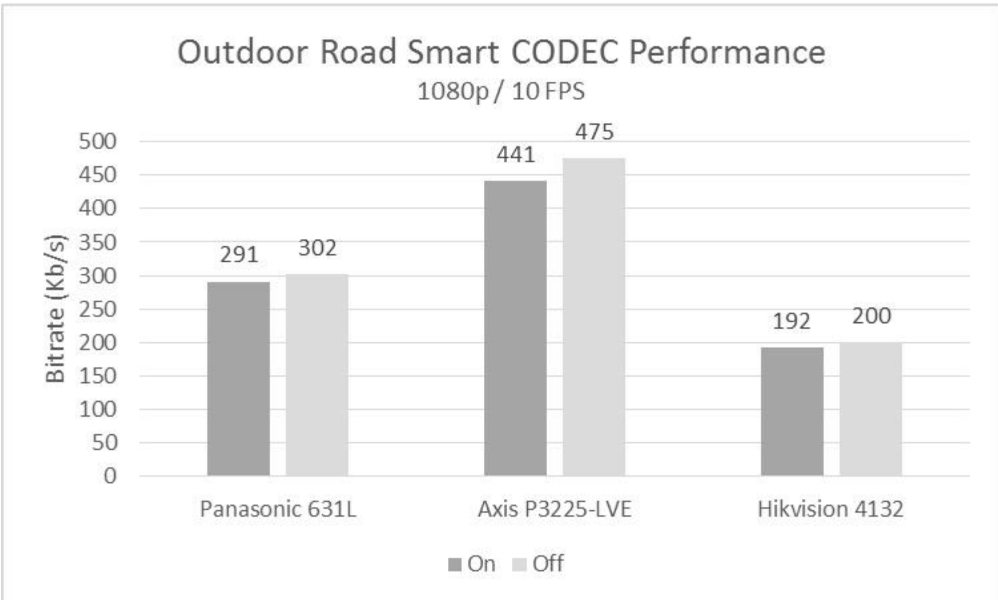

However, on average, accounting for both still periods and high traffic periods across several minutes, savings were close to 70% on all cameras.

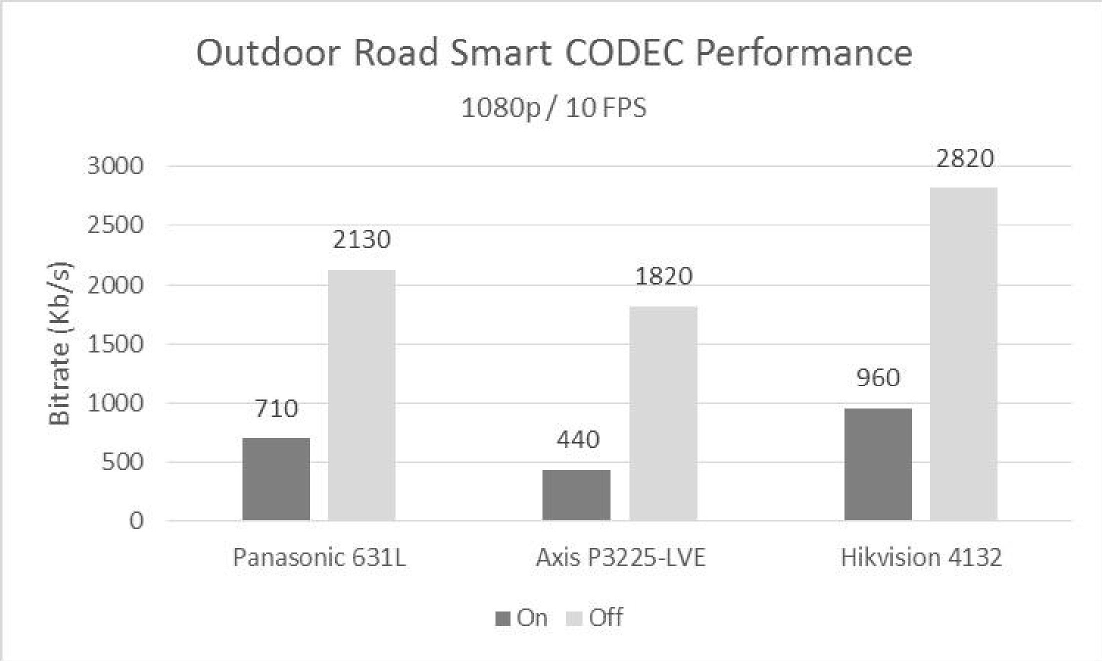

Because smart codecs may have little effect in high motion periods, bandwidth needs must be calculated for worst case bitrates, not average, assuming smart codec savings. Failure to do so could oversaturate and drop wireless or wired links or overload servers.

Smart Codec Outlook

In 2018, almost all major manufacturers include smart codecs on at least some of their cameras, with many implementing them in the majority of their line. Those lacking them, or using dated techniques (such as increased DNR or reducing color depth) will be at a significant disadvantage moving forward.

Because camera processing power will continue to increase, and because the bandwidth benefits are so significant, we expect holdouts to offer smart codecs and to improve the intelligence of these processes for further savings, and combine them with H.265 as the codec gains mainstream acceptance.

H.265 / HEVC

For years, video surveillance professionals have talked about the potential for H.265.

But now, H.265 is starting to gain mainstream adoption, with many manufacturers shipping cameras and some VMSes adding support. However, there are many issues impacting H.265's competitiveness and compatibility.

We cover:

- H.264 vs H.265 Technical Comparisons
- Impacts On Quality
- Barriers in Moving to H.265
- VMS Support Reviewed
- No ONVIF Conformance Until 2018
- H.265 vs Smart H.264
- H.265 Smart Codecs
- H.265 IPVM Test Results
- H.265 CPU Load Impact
- Patent Licensing Issues
- Usage Recommendations

Overall, the key marketing claims for HEVC/H.265 is reducing bit rate requirements in half to deliver the same quality. For instance, if a 1080p / 30fps H.264 camera required 4Mb/s, the equivalent H.265 camera would be expected to require only 2Mb/s. But that is now, clearly not enough, as we will explain inside.

For background, see this technical IEEE HEVC/H.265 whitepaper. For the full details, see the current HEVC draft standard document (200+ pages).

Technical Comparison

Three main structural improvements drive H.265's projected performance gains:

- Much larger CTUs instead of Macroblocks: While H.264's maximum block size is 256 pixels (16 x16), H.265's will be 16x greater at 4096 (64 x 64). Proponents say the larger size enables more efficient encoding, especially for higher resolution images. Read a good technical blog post on CTU details. The video below shows this:

Note: Click here to watch the H.265 Stream Analyzer video on IPVM

- Parallel decoding in H.265 will allow different parts of the image to be processed simultaneously. This can speed up playback and take advantage of the increasingly common multi-core CPUs available. H.264 did not support this.
- "Clean Random Access" syntax has been added to H.265 that "decodes pictures without needing to decode any pictures that appeared earlier in the bitstream, supporting an efficient temporal coding order known as 'open GOP' operation" (see page 5 of this document). This could be a practical benefit for surveillance as the need to playback recorded video has forced frequent I frames that may somewhat increase bit rate.

One major feature that H.265 lacks, just like H.264, is scalable encoding. While an option is planned (same as H.264), H.265 is unlikely to bring scalable video to surveillance. The benefits include streaming to low bandwidth clients and simple storage pruning (see our scalable video_codecdiscussion).

Potential Improvements in Quality

A number of manufacturers tout improvements in image quality with H.265, though this is misleading. H.265 is not inherently better 'quality' than H.264 nor was H.264 better 'quality' than MPEG-4. As such, if you are using appropriate bandwidth for H.264 for quality video, moving to H.265 is unlikely to improve image quality. However, it may reduce bit rate. The only case where it would indirectly improve quality is if bandwidth levels with H.264 were set so low as to increase visible artifacts. Then, a switch to H.265 at the same bandwidth level might improve quality.

Barriers in Moving to H.265

Moving to H.265 has not been simple or easy for surveillance due to three key barriers:

- New Cameras: Existing cameras cannot simply be upgraded to H.265 via firmware, as new chipsets are typically required, meaning that cameras must be replaced just as was required when moving from MPEG-4 to H.264.
- New VMS Versions: Since H.265 is a new(er) standard, VMS vendors need to add support. Doing so requires a fair amount of work and testing. As such delays are likely until VMSes see broader camera commitment (see VMS support below).
- Increased Processing Power: The tradeoff of bandwidth reduction is higher processing power requirements with projections of anywhere

from 50% to 300% increase. On the other hand, decreased bandwidth requirements might actually help surveillance users as it reduces i/o transfer size. Ultimately, this should be OK but there will be understandable concerns and delays to test and validate any issues.

Camera Support Increasing

As of 2018, H.265 support in IP cameras and NVRs has increasaed significantly, with Asian brands such as including Dahua, Hanwha, Hikvision, Uniview, and Vivotek offering H.265 in many models. Others such as Axis, Bosch, and Panasonic have announced and released a few models, but the majority still use H.264.

Others such as Arecont, Avigilon, and Pelco have not announced H.265 cameras.

VMS Support Still Limited

On the VMS side, several developers now include support for at least some H.265 cameras, but support is not universal and users must check that their chosen models are supported.

- Axxon Next: Dahua, Hanwha, Hikvision, Panasonic, Vivotek, and generic RTSP
- ExacqVision: Hanwha, Illustra, Vivotek, and generic RTSP
- Genetec Security Center: Axis, Dahua, Hanwha, Panasonic, Vivotek
- Milestone Xprotect: Axis, Bosch, Dahua, Hanwha, Hikvision, Panasonic, Vivotek
- Network Optix NxWitness: Hikvision
- Video Insight: Support for Panasonic/Advidia H.265 models

No ONVIF Conformance Until Profile T

Due to how ONVIF Profile S is structured, a new profile is required to support H.265, named Profile T, scheduled for release in 2018. Because of this, there is currently no official H.265 ONVIF conformance test.

However, camera/VMS providers may include support for H.265 via ONVIF (but without conformance guarantees). For example, Milestone XProtect lists support for H.265 ONVIF cameras in their supported device list.

ONVIF							
Conformant Device	Camera	H265, H264, JPEG, MPEG4	1	Two-Way	5/5	PTZ	1

See our report [How And When ONVIF Will Support H.265]() for full details.

Smart H.264 vs H.265

H.265 is not the only bandwidth reducing technology being offered for video surveillance systems. Indeed, in the past few years, [Smart Codecs]() have been introduced, delivering significant bandwidth reductions compared to standard H.264, and reducing the motivation to move to H.265.

H.264 smart codecs have 2 key benefits vs. H.265:

- Backwards compatible/works with H.264, eliminating the need for new VMS support and CPU load increases
- Dynamic I frame interval with Smart H.264 can significantly reduce bandwidth consumption, a feature that "regular" H.265 lacks.

In our testing of Axis smart codec Zipstream, the bandwidth savings there were far greater than even the optimistic marketing claims of 'non-smart' H.265.

H.265 Combining With Smart Codecs

That being said, "regular" H.265 cameras are increasingly rare, with manufacturers instead implementing H.265 along with smart codecs, such as Hikvision H.265+, Hanwha H.265 with Wisestream, Vivotek H.265/Smart Stream II, etc. This combination offers modest improvements, not as drastic as reductions with H.264 smart codecs on vs off. However, for absolute lowest bitrates, this is likely to be the combination used moving forward.

H.265 Test Results

In early tests, H.265 cameras did not deliver material bandwidth improvements/savings over typical H.264 cameras (see H.265 IP Cameras vs H.264 Test Results), with H.264 smart codecs producing better savings than "regular" H.265.

However, in more recent tests, such as Hanwha's Wisenet X cameras(essentially Hanwha's second generation of H.265 cameras), smart H.265 bitrates were notably lower than smart H.264, both in the same camera and versus competitors.

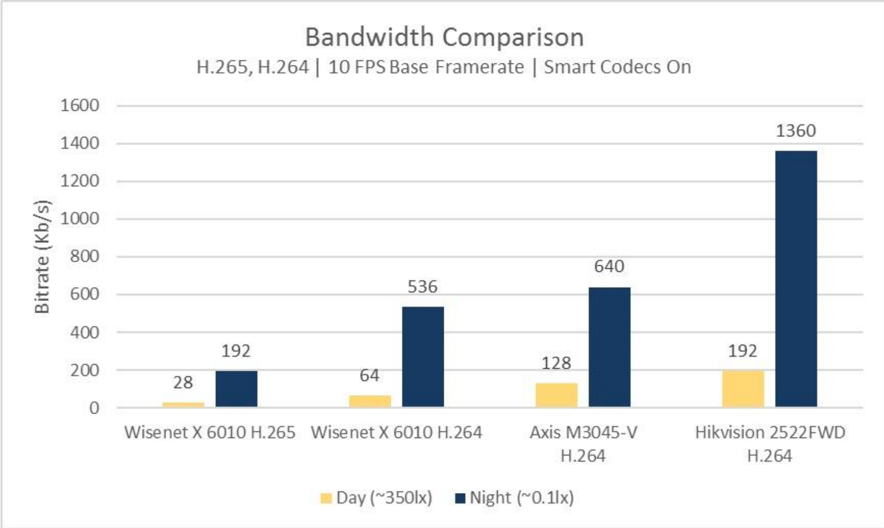

H.265 may improve as it matures and subsequent generations of chips become available, similar to H.264 improvements over time.

CPU Load Impact

In our tests, viewing H.265 was notably more processor intensive, commonly requiring double the CPU load of H.264 streams. Because of this, users should be very careful to properly select CPUs in client machines, since using the same equipment as typical systems using H.264 may be insufficient.

For example, from our Hikvision H.265+ Test Report, H.265 streams were more than double the CPU load of H.264, both day and night.

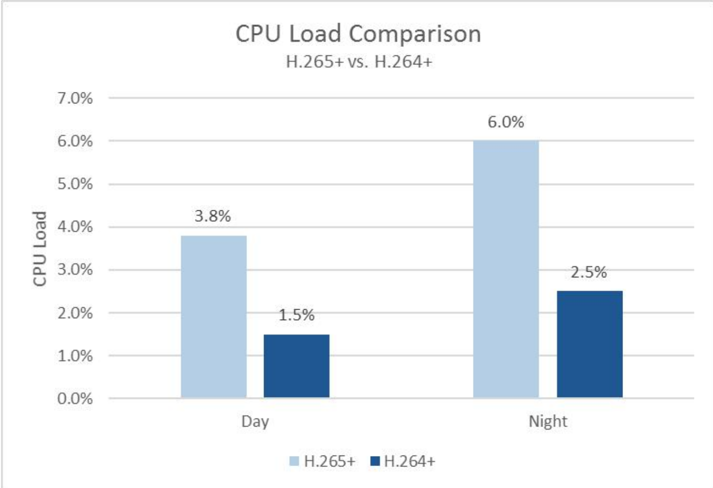

Compounding this issue, while hardware/GPU decoding has become more common in the past several years, many GPUs do not support H.265 hardware decoding. Because of this, while load may be reduced when offloading to the GPU when using H.264, these settings have no effect on H.265 load, shown below:

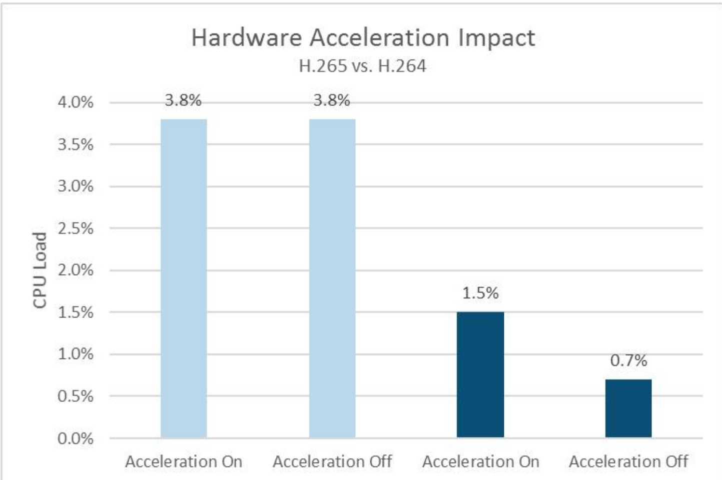

H.265 Patent Licensing Issues

H.265 has introduced a new complication for manufacturers: patent licensing. Unlike H.264, which had only a single patent holder (MPEGLA), H.265 is covered by 1,000+ patents held by multiple groups (MPEGLA, HEVC Advance, Velos, and more). Because of this, there is confusion among manufacturers about which group of patents they should license for their products.

As a result, a number of camera manufacturers/VMS developers have not licensed their products at all (below). These issues are covered in more detail in our Manufacturers Shipping Unlicensed H.265 Products Report.


HEVC LICENSE AGREEMENTS BY MANUFACTURER
IPVM Patent License Analysis 2017

COMPANY NAME	MPEG LA	HEVC Advance
ACTi		
Dahua	✓	✓
EasternCCTV	✓	
exacq	✓	✓
FLIR	✓	
Genetec		
GeoVision	✓	
Hanwha		
HIKVISION	✓	
milestone		
Panasonic		
RaySharp	✓	
SAMSUNG	✓	✓
TVT	✓	
unv	✓	
VIVOTEK	✓	✓

Usage Recommendations

Based on all the above factors, we recommend:

- Consider H.265 with smart codecs, avoid H.265 without smart codecs, since without smart codecs, H.265 consumes more bandwidth than Smart H.264. Plus, H.265 has general issues of lower VMS support and higher decoder CPU load not present in H.264.
- Ensure that Smart H.265 cameras you consider will work with your VMS / NVR of choice. Many combinations will still not work today and could create a major problem.
- Verify that you have sufficient processing power both on the recording server (if e.g., the recorder is doing server side motion detection or local display) and on the client side to ensure that the client has sufficient resources to decode and display H.265 without losing frames or video quality.

Note: This tutorial was originally written in 2013 but substantially revised in 2016 and 2017 to reflect advances/changes in H.265 support and performance.

CBR vs VBR vs MBR

How you stream video has a major impact on quality and bandwidth.

And it is not simply CODEC choice (e.g., H.264 vs H.265).

Regardless of the CODEC, one still needs to choose how the video stream handles changes in scene complexity. There are three key streaming modes (CBR, VBR, MBR) and one related feature (smart codecs) which drastically impact camera bandwidth:

CBR vs VBR vs MBR

Choosing between modes is typically overlooked:

- CBR stands for constant bit rate, aims for a constant or unvarying bandwidth level with video quality allowed to vary
- VBR stands for variable bit rate and allows the bit rate to vary but maintains a constant video quality level
- MBR stands for maximum bit rate allowing the bit rate to vary but only up to a maximum value, effectively VBR with a cap.

You need to determine whether and how much you will allow the bit rate levels to vary.

How Scene Complexity Varies

What you are streaming can vary dramatically in complexity:

- If you have a camera zoomed in on a white wall during the day, that is a very simple scene. For a 'good' quality level, a 720p HD / 30fps stream might need 200 Kb/s for this.

- By contrast, if you have a camera aimed at a busy intersection, this is a very complex scene. At the same exact settings as the first scene, you might need 20x the amount of bandwidth, or 4,000 Kb/s to maintain the 'good' quality level.

The more complex the scene, the more bits (i.e., bandwidth) you need to maintain the same quality level. It does not matter how 'good' or 'advanced' your codec is, this will always be the case.

Surveillance Challenges

The main practical surveillance challenge is that scene complexity can vary significantly even on the same camera and across just a few hours. Set the camera to use too little bandwidth and the image quality will suffer. Set the camera to use too much bandwidth and you will waste significant amounts storage.

IP Camera Implementation Issues

Making the choice more challenging are two other common factors:

- Defaults vary: Camera manufacturers have widely varying defaults - both in terms of encoding modes enabled and bit rates used. As such, two different camera's efficiency in using bandwidth can vary dramatically even if the frame rate and resolution are the same.
- Terminology varies: Manufacturers often do not use the terms CBR or VBR or MBR, often creating novel controls or terminology that can be confusing to understand. It is easy to make a mistake or misunderstand what their controls allow.

Recommendation - Use MBR

IPVM recommends you use MBR (sometimes called VBR with a cap) streaming, combining the best parts of VBR and CBR encoding:

- Compared to a typical CBR setting, MBR often reduces bandwidth consumption by 30-70%. It accomplishes this by allowing the camera to reduce bandwidth used when the scene is simple (whereas CBR always stays locked at the fixed bit rate).
- Compared to a typical VBR setting, MBR can reduce bandwidth consumption by 20-50%. It accomplishes this by stopping VBR bandwidth consumption from exploding (typically at night) by imposing a maximum bandwidth level. No practical quality loss is likely to occur because the dark scene reduces captured image details anyway. See: Tested: Why Lowering Bandwidth at Night is Good

Additionally, MBR allows better use of smart codecs (discussed below), which CBR does not. Given the bandwidth savings of smart codecs in our tests (50%+ on average in addition to reductions mentioned above), this is an even more compelling reason to use MBR.

Smart Codecs Further Savings (Requires MBR/VBR)

In the past few years, smart codecs have become common, with most camera manufacturers including them on their cameras. Smart codecs vary compression based on what is in the scene, so static background areas may be highly compressed/lower quality while moving objects remain lower compression. Additionally, they may vary the I-frame interval, switching to a low interval and lowering bandwidth when there is little activity in the scene.

Since smart codecs vary compression, I-frame interval, and other codec settings, they require VBR or MBR by nature and are generally not used with CBR. Indeed, most cameras automatically switch streaming mode to VBR when turning smart codecs on, seen below in an example from a Dahua camera's web interface.

Note: ***Click here to view the animated sample on IPVM***

Note that some cameras allow CBR to be set after turning on smart codecs, but in our tests this was simply incorrect, with streams reacting the same as when VBR with smart codecs was used.

Impact Of VBR, CBR, And MBR

The video screencast below shows you VBR, CBR and (MBR) VBR Plus a Cap in action. We demonstrate the impact on bandwidth use across 4 scenes - daylight simple, daylight with motion, night time and super high motion.

If you are not familiar with the consequences of using different streaming modes, please watch this video:

Note: ***Click here to view the video on IPVM***

Manufacturer Configuration Options

This chart provides a quick reference of which manufacturers support each streaming method:

MANUFACTURER STREAMING SUPPORT

MANUFACTURER	MBR	VBR	CBR
Arecont Vision	✓	✓	✓
avigilon	✓	✗	✗
AXIS Communications	✓	✓	✗
BOSCH	✓	✓	✓
alhua Technology	✓	✗	✓
Hanwha	✓*	✓*	✓
HIKVISION	✓	✗	✓
Panasonic	✓	✓	✓
PELCO	✓	✓	✓
SONY	✓*	✓*	✓
unv	✓	✗	✓
VIVOTEK	✓	✓	✗

Below, we walk through the encoding configuration options for each camera.

Arecont Vision

Arecont cameras support all three streaming modes, defaulting to MBR capped at 10,000 kbps. Arecont calls this cap "rate limit", shown below.

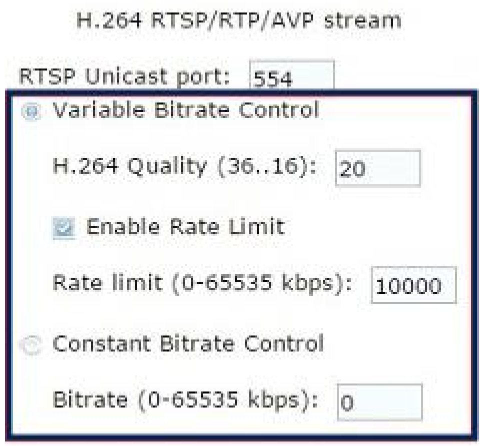

Note that some older versions of Arecont cameras supported only VBR, with no way to cap bitrate and no CBR mode.

Avigilon

Current Avigilon cameras default to MBR, with no way to use CBR or VBR streaming. With Avigilon, bandwidth consumed will vary up to the maximum bit rate configured with the max bitrate acting as a cap.

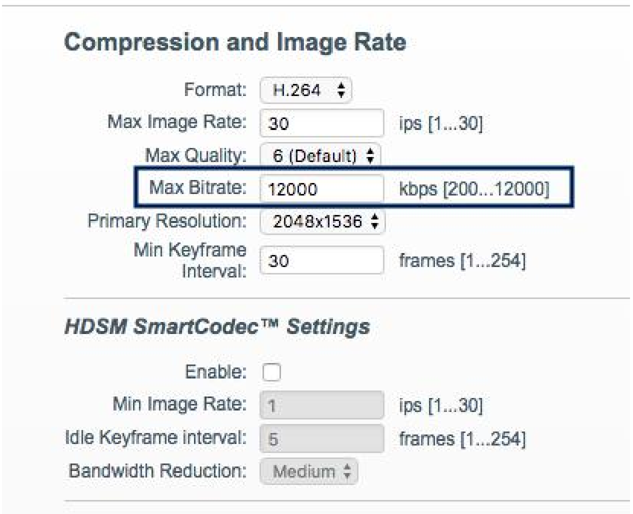

Axis

Axis currently uses two different web interfaces, as many of their cameras have not yet transitioned to their new HTML5 based UI. Both new and old web UIs default to the same settings: VBR with no cap, shown below in the new interface. The maximum bit rate option allows the bit rate to vary up to the value entered.

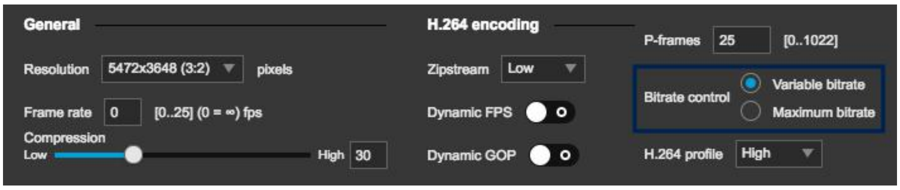

Note that Axis has never supported "true" CBR. Historically, their "constant bit rate" setting was actually an MBR labeled "CBR."

Bosch

Bosch has some of the most complex configuration options for encoding of any camera we have reviewed. It essentially can do all 3 of the modes we discussed. However, it requires understanding how to set the 'target bit rate' and 'maximum bit rate' fields shown below:

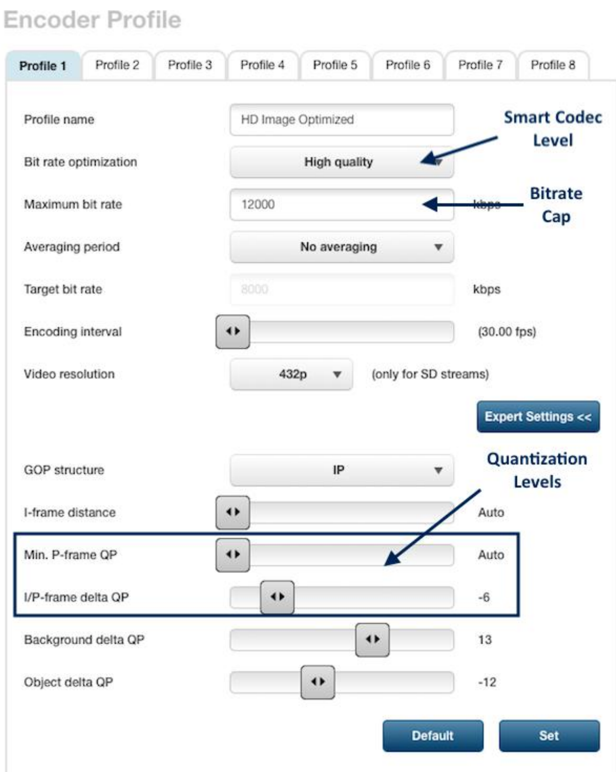

For MBR, simply set the maximum bit rate to what you would like your cap to be. The target bitrate is effectively a minimum bit rate.

If you want something closer to uncapped VBR, make sure to set the target bit rate low and the maximum bit rate high (it may be set as high as 40,000 kbps), allowing the bitrate to vary widely. Cameras will typically never approach this 40 Mb/s cap.

Finally, if you want something close to CBR, set the target bit rate and the maximum bit rate close together, which will force bandwidth to stay in a very narrow range. Bosch allows for these two numbers to be as close as a 10% difference (e.g., 5000 and 5500 for target and max respectively).

Dahua

Dahua IP cameras support MBR and CBR, defaulting to CBR. Users may select predefined max bit rates or enter their own. There is no way to uncap VBR streams.

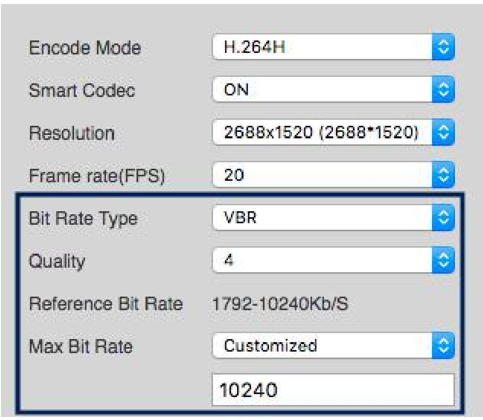

Hanwha

New model Hanwha cameras include CBR and a mode which they label VBR. However, unlike true VBR implementations, Hanwha cameras do not include any way to fix compression of the camera, with both bitrate and quantization varying. For full details, see our report Sony and Samsung Breaking VBR.

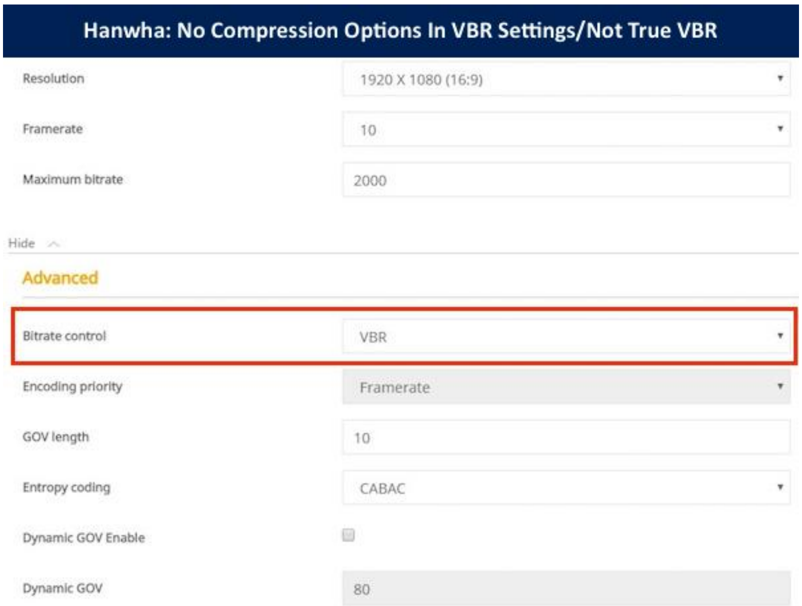

However, older model Hanwha/Samsung models supported CBR and MBR, with a cap of up to 30 Mb/s.

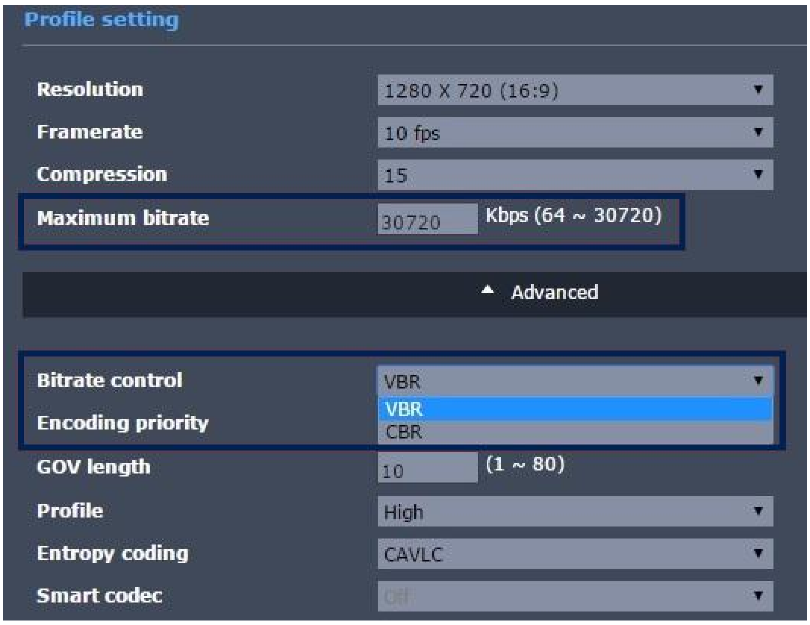

Hikvision

Hikvision IP cameras support MBR (default) and CBR. There is no way to uncap VBR streams.

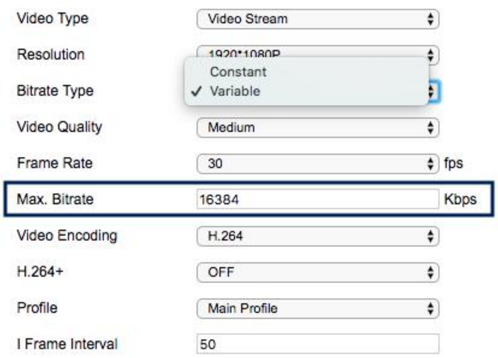

Panasonic

Panasonic IP cameras support CBR and MBR (default). Note that there is no way to uncap streams when set to VBR, but they may be set to very high limits (24 Mb/s or higher, depending on model). The camera includes other modes (Frame rate and Best effort, below), but we recommend simply using MBR as in other cameras.

Note that new Panasonic Extreme models are one of few cameras to default to H.265 for their primary stream. Users should beware of this as it may cause connection failures if VMSes/NVRs do not support this stream.

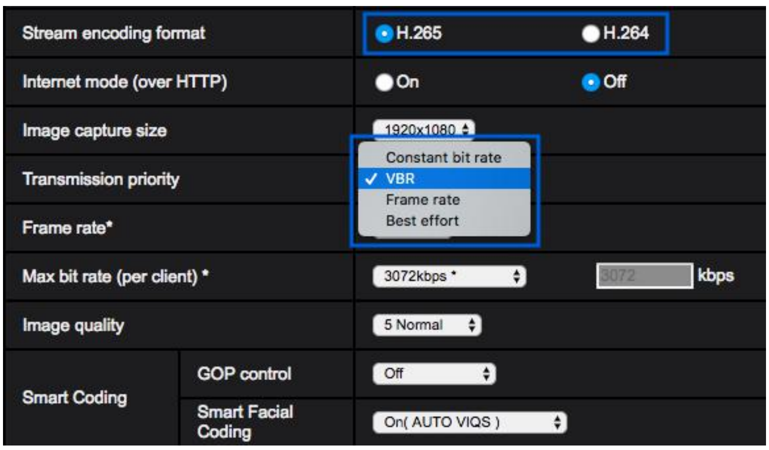

Pelco

Pelco supports all 3 modes, but use what they call CVBR or 'Constrained Variable Bit Rate' encoding. Essentially, this is a CBR codec that allows the bit rate to vary modestly - approximately plus or minus 10%. As such, it is more or less a CBR mode. The specific bit rate or target is set in the Bit Rate field, shown below:

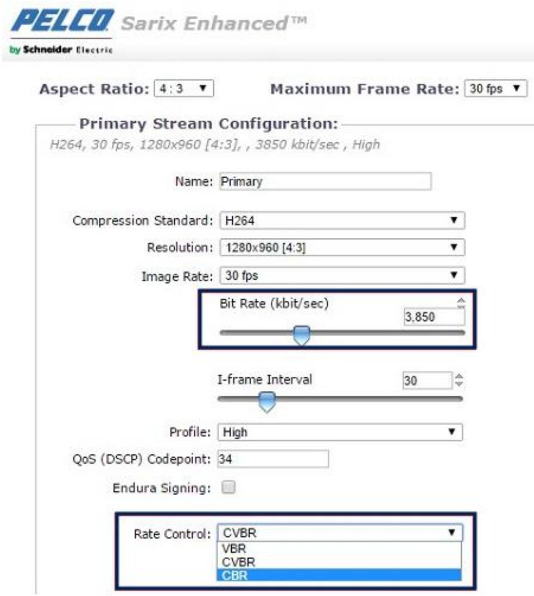

The maximum bit rate shown above acts as the cap / max.

Sony

Sony now supports all 3 modes, though Sony only started supporting this in 2013/2014 so existing Sony cameras that have not had their firmware upgraded may not support VBR or MBR.

To enable MBR on Sony, select VBR and then fill in the maximum bit rate limit appropriately as shown below:

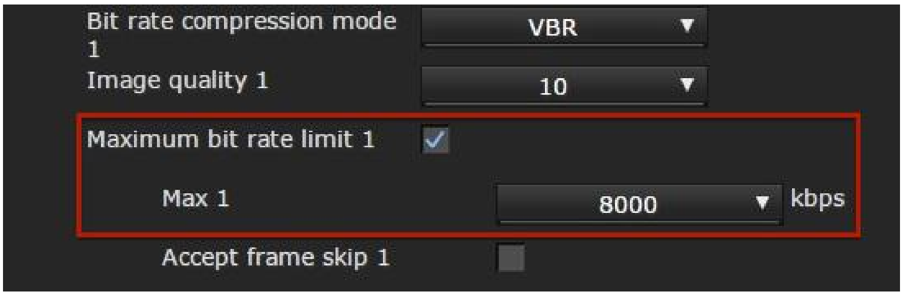

Below shows all 3 options for Sony:

However, note that in Gen 7 cameras (SNC-VM772R and SNC-VB770), Sony does not include any way to fix compression when the camera is set to VBR, instead including only a target bitrate.

However, even these settings are confusing, as we found the camera consumed less bandwidth than the target bitrate in our tests, and never approached the bitrate cap, regardless of scene activity. We review these issues in this video from our tests, below.

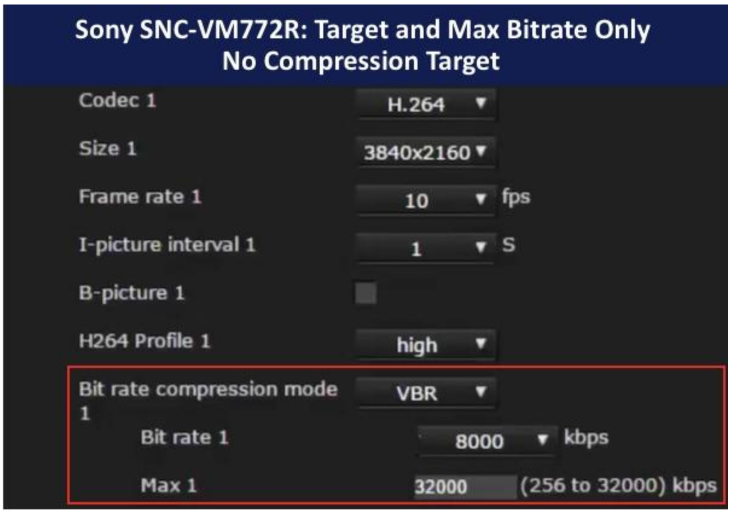

Uniview

Uniview cameras support MBR (though named VBR) and CBR. Max bitrate is constrained to an upper limit of 16 Mb/s, shown below.

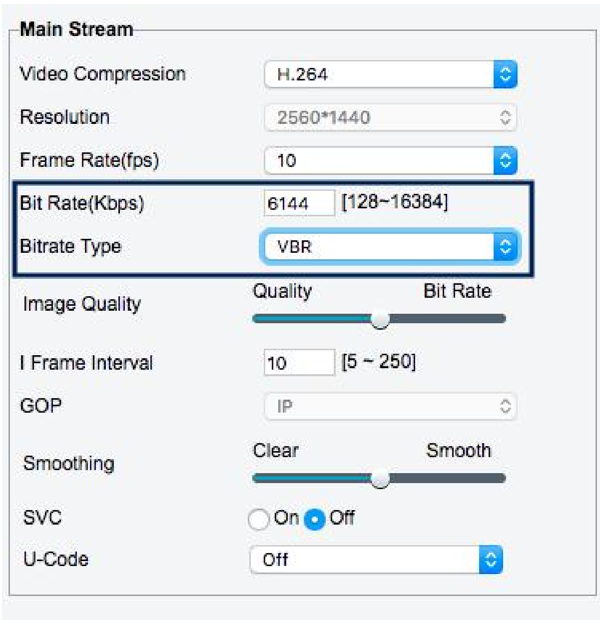

Vivotek

Vivotek supports VBR (named "Fixed quality" in the UI) and a version of MBR which they name "Constrained bit rate." Constrained bit rate performs similarly to VBR, with the camera fixing compression at a specific level and bitrate freely varying up to a specified limit (up to 40 Mbps shown below). When approaching this limit, compression or frame rate may vary, shown in the "Policy" dropdown below.

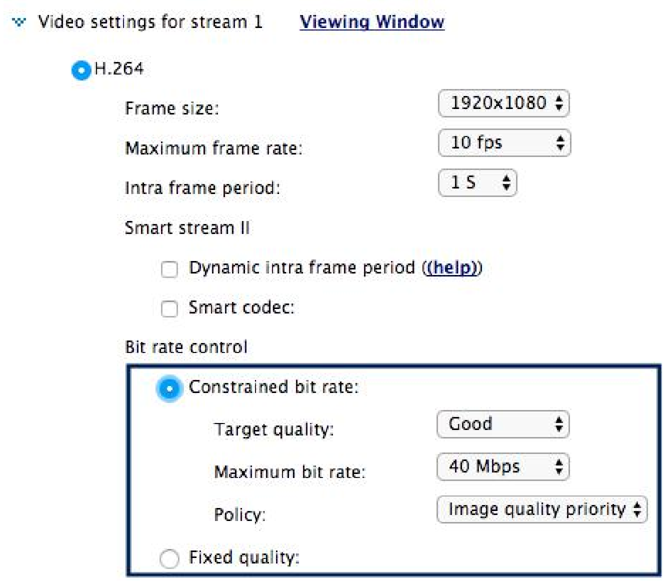

Dealing with CBR Cameras

Though CBR only cameras are now rare, users may still come across new models and those from existing systems which still use CBR only. In this case, here are some recommendations.

Out of the box, manufacturers of CBR cameras, typically set the default bit rate fairly high relative to common usage. This means the image quality should look good. On the negative side, this also likely means you are wasting bandwidth. With CBR, determining the right bit rate can take one of two basic approaches:

- Keep the bit rate set high, avoid any quality problems but probably waste 30-70% of bandwidth used.
- Test the complexity of the scene and the quality you need, by trying a few bandwidth levels at different times of the day. Then choose.

Either way, CBR forces users to compromise. Unless the scene stays the same 24/7, there will be tradeoffs: either wasted bandwidth or degraded quality at some points of the day. This is not an easy call and why most prefer MBR since it eliminates this guesswork and potential for errors.

Impact Of VBR And CBR On Quality

Both VBR and CBR impact compression levels, which is an important and underappreciated aspect of video streaming. Regardless of resolution selected (i.e., 720p, 1080p, etc.), the amount the video is compressed varies and the more compression, the worse the video typically looks.

With VBR, the compression level is fixed, and bandwidth varies to ensure each scene is compressed at that level. With CBR, the bandwidth is fixed, so the compression level has to adjust when the scene changes. For more, read our our video quality / compression tutorial.

[NOTE: In 2017, this post was updated to add advances in smart codecs and changes in manufacturer configurations and options.]

Video Surveillance Storage Duration Averages

IPVM statistics show that 1 month duration for video surveillance storage is, by far, the most common used.

However, in contrast to 2012, there was a significant decrease in the amount of systems using 2 weeks of storage, with many shifting to 1 month or longer.

On the other hand, average storage duration overall is not going up significantly, with many respondents making it clear that 1 month is typically more than enough, regardless of storage costs.

We break down the statistics and integrator's explanation for storage duration decisions.

Statistics Overview

In a survey of ~150 integrator respondents, nearly 70% report average storage duration of ~1 month, with a roughly equal number averaging more or less than that.

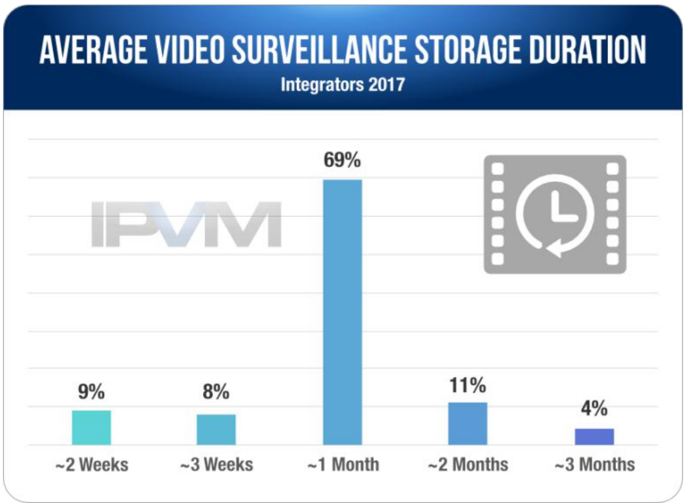

Customer Expectations

A common integrator explanation was that 1 month has become a customer expectation:

- "30 days seems to be the standard. If there is some other reason for needing more or less it is usually related to some type of regulatory standard that requires more than 30 days."
- "Perception of commercially appropriate recording speeds and storage durations."
- "For general, unregulated industry there doesn't seem to be a strong rationale for wanting more or less than 30 days."
- "Most customers have corporate standards of 30 days. Some customers have higher requirements due to specific industry standards or compliance."

Operational Requirements

This expectation has been reinforced by most users being able to find and retrieve video in a month:

- "We recommend at least covering your maximum length of out of office time and coverage for the average time it takes for a service technician to respond to a request considering you may need to realize something has occurred."
- "Most people say 30 days, but admit that two weeks is probably all they really need to determine that an incident has occurred so they can review/offload associated video."
- "Instead of asking how long a client would like to store video on site, I ask how long it would be before an incident is ether reported and

or an incident was observed. It's typically within a few days. The longest period stayed is about a week. By doubling that the client is able to retrieve most video of most incidents then export to other storage medium or back up location."

- "For many, it's a shot in the dark, but we typically try to guide them down a path of understanding how much is necessary from an incident standpoint. Larger organizations tend to find out later if there's been a less obvious incident, while smaller companies usually know right away if something happened. With storage costs dropping and compression improving, longer retention periods are just more attainable, and customers usually want more than less."
- "Cost to the customer/budget or the nature of the customer's requirements. Financial institutions & some Hotels require longer storage (days) and require continuous recording in certain areas."

Legal Requirements

Legal requirements, both limiting how long storage can be kept or requiring storage be kept longer, have an impact in a significant minority of cases:

- "Local regulations limits at 30 days (in most cases)"
- "In Kuwait, new government regulations mandate 120 days storage."
- "There are some people who have to have 30 days because of some regulation or policy... those people don't seem to buy up over the requirement."
- "Some storage requirements are driven by mandate. For example Colorado marijuana laws dictate 40 days of storage and Washington state law dictates 45 days of storage."

- "We do a lot of healthcare sites and 30 days is mandated by the health authority. We use this as a matter of course for most of our other sites."

Physical Limitations

Some integrators cited physical limitations of the recorder used as an issue:

- "At present we are limited to the NVR manufacturer in the way that they write to the HDD's and also the cost on the size of the HDD from the manufacturer. We try to keep the system in one piece of iron, rather than a few."
- "With hard drive size now 4 TB and much larger it can be much easier to set a 60 day standard , however raid options and other factors like physical rack space may change the requirement"
- "The biggest limiting factor (as with all security items) is the cost involved. For most places once they get to an enterprise server the cost of additional storage is low, so we go with lots of storage; for small installs though many of those servers have a hard time with lots of storage (usually they only have 1-2 drive bays) so there is a limiting factor there."
- "We generally use 4GB [Note TB] 3.5 surveillance hard drives, 4 per server. These have the best cost/capacity right now, I expect to switch to 6GB [Note TB] hard drives sometime this year. Our VMS server for smaller jobs holds 4 drives, we have no problem hitting 30 days retention, even with hard drive redundancy."

Many DVR / NVR appliances only support a limited number of hard drives internally (1 and 2 being common for smaller systems). This can constrain storage duration, especially given the low cost of those appliances and the extra cost of expanding beyond them.

Cost

The main negative for longer storage duration remains cost, as many integrators cited this as the main constraint:

- "Price of storage devices"
- "EndUser budget. Cost of TB and Servers"
- "Mostly smaller systems so cost factor"
- "Cost is always the main factor. Most of the systems are utilizing some form of RAID and we do not use consumer level hard-drives."
- "100% Cost"
- "Storage cost is usually the main reason to limit the number of days to store"
- "Cost of large storage RAID arrays needed with higher megapixel cameras and recording at 20 to 30 images per second."
- "Most often the limit is the cost of storage in a RAID environment"

However, in contrast to 2012, integrators were notably less driven by cost constraints, with quite a number emphasizing that cost is not a significant barrier:

- "Storage costs are not often a major concern as drive prices are fairly low and many new cameras are much more efficient than previous models."
- "HDD sizes/prices make it so that shorter retention doesn't really save on costs all that much."

- "With storage costs dropping and compression improving, longer retention periods are just more attainable, and customers usually want more than less."
- "Typically cost, although storage is really a pretty inexpensive part of the equation so there's no reason to skimp."

Hard drive cost per TB are certainly lower than 5 years ago (~50%) plus the rise of smart codecs which regularly reduce storage consumption by 50%+ help here.

Right Storage Duration?

These are statistics of what integrators are using. One month is likely 'right' for most people but depending on your needs (e.g., are you bank that only gets reports of fraud 45 days later, etc.), you might want / require longer storage. However, recording for less than one month is increasingly hard to justify in terms of the cost savings, especially as new installs use smart codec cameras, which based on our smart codec tests, we recommend.

Related, see: Average Frame Rate Video Surveillance Statistics

Recording Mode Statistics

Continuous? Motion? Boost?

What recording should be used? What is used?

We compile data from 90 integrators on storage modes most commonly used, and the reasons why integrators choose a specific recording method, contrasting results to our 2011 recording mode statistics.

Summary

Motion recording is still dominant, based on the answers to our question "What recording mode do you use most often? (e.g.: Continuous recording, motion recording, boost up recording) Why?"

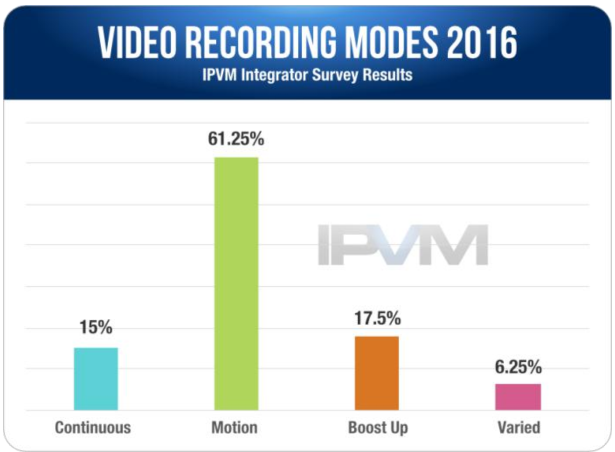

Results show little change compared to our [2011 Recording Mode Statistics]() report. This may be due to higher resolution cameras being more common in 2016, generating more data and offsetting any price drops in storage.

Motion-Based Recording

Integrators like motion-based recording not just for storage savings, but also to make searching video easier. This was mentioned multiple times in responses:

- "Motion mainly to reduce storage requirements and to simply audits when searching for events."
- "Motion recording. Saves storage space and quicker/easier reviewing."
- "Motion recording. Both to save hard drive space and to make it easier to find recordings."
- "Motion. Easier to search. No use wasting drive space."
- "Motion; to increase capacity and improve my search performance."
- "Motion Recording, as full time recording eats tons of space, and is next to impossible to search."
- "Motion. Save on drive space but more importantly makes searching footage much easier."
- "Motion based. Saves space and provides faster playback review."
- "We almost exclusively use motion based recording. The number one reason is to save hard drive space. It also make searching for an event easier on simple systems that do not indicate motion when using Continuous."

Other comments in favor of motion recording mentioned efficiency and practicality of recording only what is relevant:

- "We use software-based motion recording for all of our recording. This is enabled through the Luxriot VMS and has proven to be much more accurate than UDP and Hikvision camera-based motion recording."

- "Currently motion recording for most efficient use of recording space"
- "Motion recording. Why record an empty room/area??? Motion covers 99% of scene/client requirements."
- "Motion for 75% because of limiting recordings to real movements. No need to watch useless recordings, also saving some HD cost. "
- "Motion recording based on the end-user request to have longer recorded period."
- "Nearly all my systems use motion due to the space saving that can be achieved."
- "Motion recording w/ pre and post event. Most customers that we've dealt with would rather not record when nothing is happening. "

Continuous recording is used when there is a fear of missing critical images, or when customer requirements specifically request it.

- "Continuous recording, this is as per client preference."
- "Our experience continuous recording is the most stable, lowest failure rate."
- "Continuous, to make sure everything is recorded."
- "Mostly continuous recording to not miss anything especially when monitoring larger space with one camera"
- "Continuous. because of customer requirements."
- "We use most continuous recording, because you have every time pictures of all cameras. If something goes wrong you can check every camera if there is something interesting or not. With motion

recording you have not all the time pictures and if the motion is not enough, you will have no pictures."

One response did note that continuous recording was more feasible with smart CODECs and lower HDD prices, and another noted that it should be more common but overall this did not seem to be a trend:

> "I mostly use continuous recording as not all motion recording is reliable and parts of some events, or distant movement, can be missed unless continuous recording is used. Thanks to the reduction in storage costs and the growing use of smart codecs, continuous recording is feasible."

Boost-Up Recording

Boost-up was slightly more popular than continuous recording, integrators using boost-up often commented that it had the key benefits of continuous (no gaps in recording), with the storage savings benefits of motion detection.

- Often record continuously at 1fps and 12-15 on event. 1fps can show what didn't happen and 12-15fps shows fluidity of what did happen. With some exceptions both ways.
- Continuous and boost up modes: we're in public transportation sector, our customers want 24/7 surveillance due to vandalism; motion detection is limited: external light changes cause false motion detection alarms.
- Boost up recording. We always record a minimum of 4CIF 2FPS no matter what (usually it is 2FPS with 1.3mp if the recording space allows) that then boosts to 10FPS full resolution on motion; and then

a 15FPS full resolution on alarm or other triggering event. This way there is no risk that we don't record anything if the camera misses motion, and we still get to optimize retention.

- Continuous with event. Gives us a constant record that the camera and recorder are working a something if motion isn't triggered then higher frame rate if it does see motion.
- Motion/VA recording w/ 1-5 second (depending on the scene) snapshots. Best of both worlds: reduced storage while still being able to prove, e.g., someone was not where they said they were.
- I recommend continuous at a smaller resolution, 1FPS, then on motion to boost up to 7-10FPS and the highest resolution on the camera as possible. I do the continuous recording to prove the negative. If someone claims they slipped and fell (when they didn't), and you only have motion based video you can't prove they didn't.
- Boost up recording. We typically boost the VBR cap on motion. We generally keep resolution and frame rate the same.
- Mainly continuous low frame with higher frame on motion To be sure nothing is missed pre / post. Relying on motion only has worked against us in the past especially when the camera is covering a large area
- We always record in low res (4 if) and bump up to high res during motion events. We went back and forth for a while recording continuously and motion only. With Continuous recording, it required too much storage, with motion only we were missing events. Example, if someone were sitting still on a forklift and didn't move for 10 min, the system wouldn't be recording. Too many angry customers. Storage costs have come down, and the systems ability to change recording resolutions has improved, so it is a no-brainer.

- We use Video Insight's mode of 3FPS all the time, boosting to a higher frame rate on motion (typically 10FPS, but sometimes higher).

Varied Recording

A small number of responses indicated that multiple modes were used, typically using motion-based recording wherever possible, and selectively using continuous recording for cameras covering more critical areas:

- Motion, with continuous on critical areas. With high res cameras if you don't use motion you need incredible storage space.
- Motion recording with 5-10 second pre & post. With H.264 compression. It saves on disk space and if the area(s) being viewed/recorded are not critical. When viewing/recording a more critical area we'll do continuous recording.
- Motion is typically used except in defined areas where we and the customer feel continuous recording is suitable.

Storage / Bandwidth Calculation

Calculating surveillance bandwidth is complex, and inexperienced users can easily underestimate bandwidth, leading to reduced storage durations and/or overloaded networks.

The most common way technicians estimate storage is to use manufacturer or third party calculator tools. However, these tools are too simple for the complex factors impacting bandwidth / storage, a fundamental flaw which fails to reflect real world conditions. We explain these key issues and give our recommendations for how to most accurately calculate surveillance storage needed.

Most Use Calculators, Despite Issues

Despite their high potential for inaccuracy, most people use calculators. They are simple to use for even novices, typically asking only for a few basics such as number of cameras, resolution, frame rate, with an estimate immediately generated.

From an IPVM member survey, here is the breakdown on preferred method, with a strong majority using storage calculators:

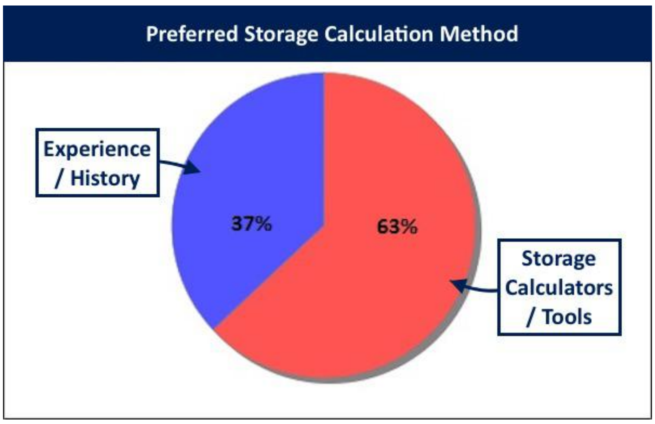

However, it is nearly impossible for calculators to reflect the wide range of conditions in which cameras are installed, as well as the variances between camera models. For example, we asked users how much bandwidth a 1080p H.264 camera uses. Notice how widely bandwidth estimates vary, even using the same resolution, framerate, and CODEC:

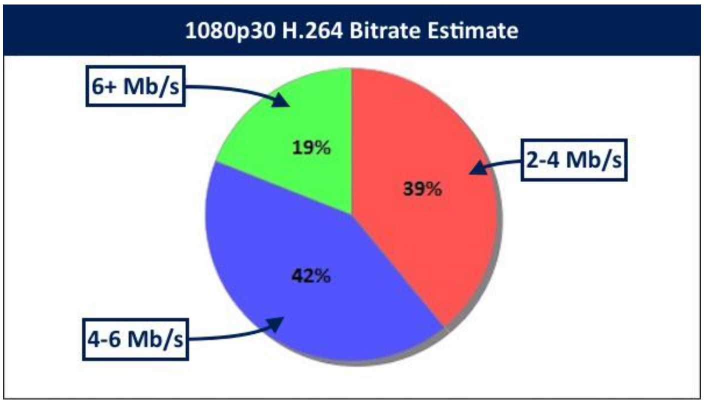

All of the respondents could be 'right' even though the answers vary by more than 300%. Differences in cameras used and sites deployed can easily result in massive differences in actual bandwidth/storage consumed. Calculators do not reflect them.

Most Accurate Method

To accurately calculate bandwidth/storage, there are three key recommendations:

- Test cameras and record bandwidth in varying scenes
- Record scene complexity for each proposed camera view
- Understand how key camera settings impact bandwidth/storage consumption

These three points are critical to proper estimation and will help users spot critical mistakes in calculations.

For example, a calculator may estimate 2 Mb/s for a given scene, but based on past history and knowledge of specific camera performance, bitrate is more likely to be 4 Mb/s, a 100% increase.

Testing Cameras

Because of this, it is critical to learn the bandwidth consumption of each camera model you use in scenes ranging from simple to complex, such as:

- Simplest: Blank wall
- Simple: Empty hallway
- Medium: Well lit lobby/reception area
- Complex: Parking lot
- More Complex: Busy intersection

Record the bitrate for each of these complexity levels, the light level, and a screen shot from the camera(s). These may be easily referred to later as a database of measurements and screenshots for easy reference for comparison. Be sure to take special note of any time the bandwidth spikes or plunges (specific times of day, extreme movement, etc.) and record specific settings as necessary.

The video below shows the basics of how to experiment and measure bandwidth consumption, shown in our tests.

Note: Click here to watch the video on IPVM

Recording Complexity

When looking at proposed camera installations, we recommend tracking scene complexity as each camera location is surveyed. Take a few moments to observe traffic levels, moving foliage, lighting or lack thereof, and other scene factors which may impact bandwidth.

After viewing the scene, record the camera's intended view and complexity level and take a photo for reference. This may be as simple as the chart below:

	View	Complexity
Camera 1	Front parking lot	High
Camera 2	Front lobby	Medium
Camera 3	West hallway	Simple

After this table is created, you may compare it to actual bandwidth from tested scenes. Snapshots from tested cameras may be easily compared to photos taken from the proposed scene for more accurate comparison.

Further Calculation Complications

Further complicating calculations, there are many factors which may cause slight to extreme variances in bitrates.

- Camera model differences
- Quantization/Compression
- Smart CODECs
- Motion detection/analytics tuning
- Night time / low light

Camera Model Differences

Do not assume because you have tested bandwidth for one camera that you can apply these findings to others of the same resolution and framerate. Even in the same manufacturer's line. Different camera models, even from the same manufacturer can produce very different bit rates.

For example, the following image shows two cameras, an Axis Q1604 and Axis M3004, both 720p, 30 fps, set to a ~20' horizontal FOV, at compression of ~Q28. Despite these factors being standardized, in this well lit indoor scene, the Q1604's bitrate was 488 Kb/s while the M3004 consumed 1.33 Mb/s, nearly 3x the bandwidth.

Recommendation: Test A Wide Variety of Cameras

Differences from manufacturer to manufacturer are even more extreme than examples above. Because of this, it is critical to learn the bandwidth consumption of each camera model you use.

Quantization/Compression

Even if two of the same model camera are used, using H.264 and the same resolution, compression levels may vary significantly, causing dramatic differences in bandwidth consumption.

Further, manufacturers default to vastly different compression levels (see our IP Camera Manufacturer Compression Comparison), with no standardization, making calculations based on only one model inaccurate.

Recommendation: Standardize Compression

Based on years of IPVM testing, 28-30 quantization is approximately the "sweet spot" between image quality and bandwidth. Higher levels have significant negative impact on image quality, while lower levels increase bandwidth with little gain in practical image quality. Standardizing cameras at this level allows for more controllable results compared to manufacturer defaults.

See our CODEC Guide and Manufacturer Compression Comparison for more details.

Smart CODECs

In the past, cameras almost always used fixed settings for compression, I-frame interval, and framerate, making calculations simpler. But with the introduction of Smart CODECs in the past few years, cameras may now dynamically manage these settings depending on what is in the scene. There are three main techniques used by smart CODECs (though not all use all three):

- Dynamic compression: Instead of applying the same compression level to the entire field of view, Smart CODECs may increase compression on static/background objects and reduce it on moving/foreground objects, lowering bandwidth overall.
- Dynamic I-frame interval: Where H.264 streams typically use a fixed I-frame interval (e.g. 1 second or 30 frames), Smart CODECs increase the distance between I-frames when there is no motion in the stream, and immediately increase when activity begins. This

technique reduces bitrates significantly due to the much smaller size of P-frames.

- Dynamic FPS: Finally, some Smart CODECs reduce framerate when there is no motion in the scene, down to a minimum of 1 FPS or a threshhold set by the user.

Using one or more of these techniques, bitrate savings of up to 95% have been seen in IPVM tests. At a minimum, dynamic compression reduced bitrates by 10-15% in even very busy scenes.

Recommendation: Test Smart CODECs In Place/Conservative Estimates

Unfortunately, exactly how a given Smart CODEC may perform in a given scene is unpredictable. Even small amounts of motion, such as shadows or foliage, may keep be enough activity to prevent dynamic I-Frame/FPS features from functioning. As a result, the best case scenario is to test the camera/Smart CODEC in place in the scene it is intended for.

Failing this, we recommend reducing smart CODEC bitrates by only a minimum amount (5-15%) to prevent underestimation and subsequent overloading of storage.

See our Smart CODEC Guide for more details on these new CODECs and related issues.

Video Motion Detection/Analytics Performance

Underperforming/mis-configured video motion detection and analytics may be one of the largest sources of inaccurate storage calculations. For instance, if you are estimating percent of 'real' motion in a server room, it is likely to be very low - less than 5%. However, flashing LEDs, screensavers, reflections, etc. can make motion detection record nearly continuously, making field results much worse than estimations.

Recommendation: Optimize/Re-Optimize

Users should be prepared to optimize VMD and/or video analytics regularly. At a minimum, checking and optimization should be performed about a week after installation, once some general benchmarks for bandwidth/storage can be made and investigated.

However, changes in season, landscaping, office arrangements, etc., may also make initial storage estimates inaccurate. Users should periodically check these configurations (~60-90 days or as major changes are made) and adjust as needed.

See Optimizing Motion Based Recording for more details on this subject.

Night Time / Low Light Bandwidth Consumption

Everything else equal, in low light / night time, bandwidth tends to be higher, sometimes far higher (e.g., 10x as high). There is no easy / simple way to estimate this as a number of factors come into play:

- Levels of visible noise: As a rough rule of thumb, the more noise one can see on the video, the higher than bandwidth, though this can be offset with some DNR techniques.
- Integrated IR: In low light, cameras with integrated IR tend to consume more bandwidth than in the day but significantly less than non-IR cameras. However, this is still impacted by the quality and range of integrated IR and the complexity of the scene.

- Camera low light sensitivity: Some cameras, whether due to better low light image processing or because of larger sensors or lower F stop lenses, will consume significantly less bandwidth at night.

See Testing Bandwidth vs Low Light and Testing Low Light Vs Image Quality for more details.

Recommendation: Use MBR / VBR With A Cap

In order to better predict nighttime bandwidth, we recommend users cap bitrates at night. Even when caps are set much lower than uncapped bitrates at night, practical image quality is very similar, as the high levels of digital noise and darker images simply do not suffer from increases in compression the way full light images do.

See our test Lowering Bandwidth at Night is Good for examples and our recommendations.

Storage Calculator Examples

Many/most camera and VMS manufacturers offer a bandwidth/storage calculator tailored to their products. The exact features of these calculators vary, with differing levels of complexity and customization. Three common examples which are freely available online (no login required as in some) are discussed below:

Supercircuits Tool

As a top Google search result, Supercircuits bandwidth tool is often cited but is quite simplistic, with limited inputs / parameters:

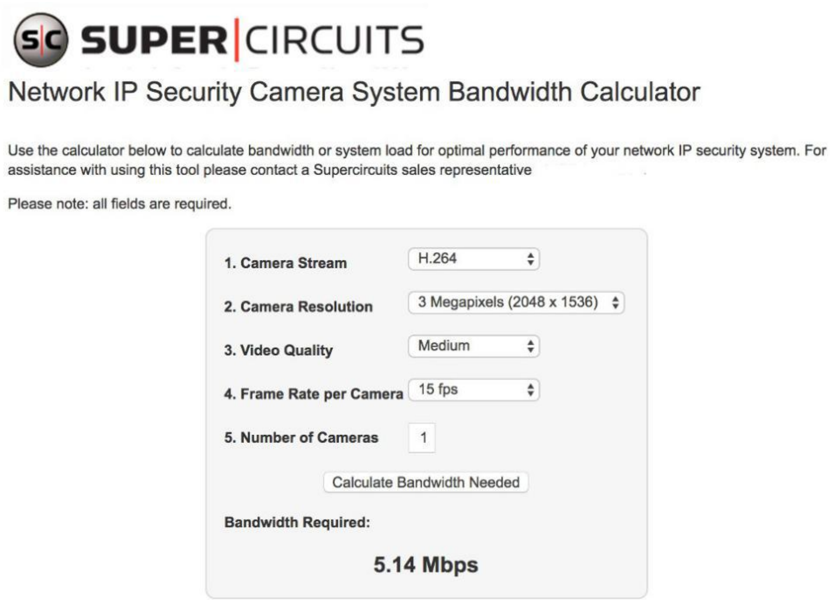

If all you want is something very rough and are not familiar with cameras, this is ok. But the calculator fails to consider differences in camera models, scenes, activity levels, to name a few. Also, it assumes a linear relationship between frame rate and bandwidth (i.e., double the frame, double the bandwidth) despite that generally not being the case in practice (e.g., Testing Bandwidth vs Frame Rate).

Axis Design Tool

The Axis Design Tool is more sophisticated, with built in options for live, continuous, and event recording streams, video examples of multiple scenes and compression/CODEC settings, and more. Further, it allows custom adjustment of multiple settings, so users may better adjust it to fit their real world testing. Two uncommon elements that Axis support are specific compression levels and scenes, highlighted in red boxes below:

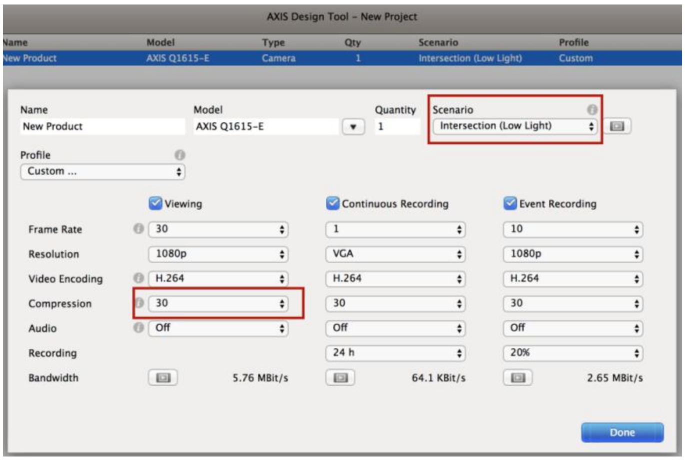

In particular, Axis allows picking from scenes (called scenarios) that adjust projected bandwidth consumption dynamically, e.g., the railway 'station' consumes more bandwidth than the 'stairwell', which is typically with VBR. The image below shows the options and sample videos they provide:

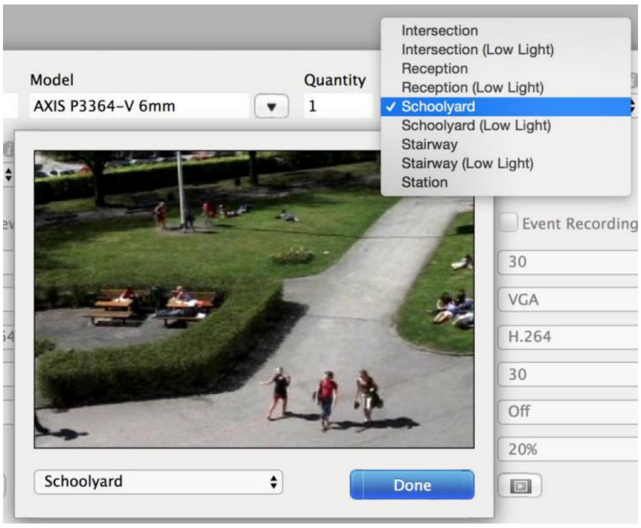

However, even in Axis tool, the choice of scenarios will frequently not reflect one's scenario. For example, there 'intersection' scenario has very little traffic in it and a lot of empty foreground resulting in bandwidth less than the stairwell scenario:

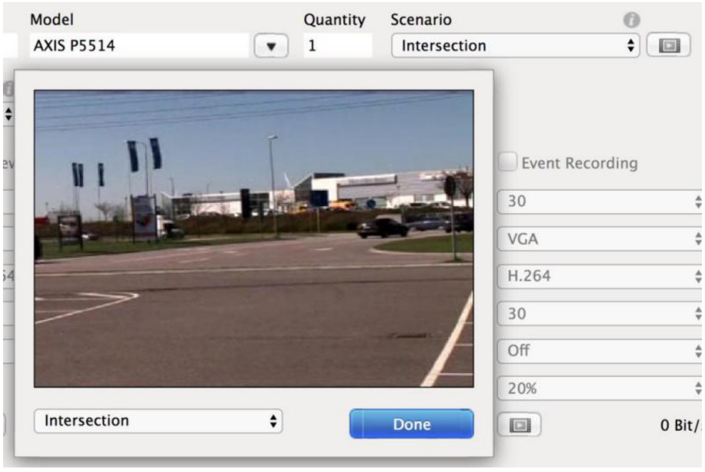

'Your' intersection may be similar but if it is busier, its bandwidth will be far higher (double, triple, quadruple quite possible). That does not make Axis tool 'wrong' but it reflects the fundamental challenges of estimating bandwidth abstractly.

Also, Axis scenarios support low light vs non low light, however, Axis assumptions are simplistic and do not factor that cameras with integrated IR will almost always consume far less bandwidth than those without IR in low light. The example below shows how an Axis non-IR and IR models assume the same bandwidth in low light:

Note: ***Click here to watch the example on IPVM***

Lastly, in our experimentation, Axis tool generally assumes the same bandwidth for all cameras of the same resolution / settings, not factoring in variances in different models tuning or use of sensors (which in our testing, as we have shown above, have an impact on bandwidth).

Exacq Configuration Calculator

Exacq's Configuration Calculator, like many VMS implementations allows for selecting various camera manufacturer's models. The main benefit is that you can check which models are supported, pick those models and it will automatically fill in that model's max bandwidth, e.g.:

However, one major issue is that the estimated bandwidth is extremely high in practice, e.g., each of these 3MP / 15fps cameras are estimated at 7.5Mb/s. Even without smart codecs, that is atypical in practice, except for fairly extreme high complexity scenes or low compression configurations.

It is important to note that many calculators for storage / server providers assume far higher than typical values, which is generally done as a safety measure to stop buyers from objecting to storage not meeting specifications. For example, see: Iomnis Guaranteed Video Storage Calculation

Summary

If you care about getting your bandwidth / storage calculations correct, try out the camera models you want in the locations you plan to use them. While this may take a couple of hours, for any project with significant amount of storage (e.g., more than a few hard drives), you will likely save yourself time and future problems by estimating yourself rather than depending on calculators that cannot come close to fully / accurately matching one's combination of cameras / scenes.

Connecting Cameras

API / SDK

While Application Programming Interfaces (APIs) are key to 'open' platforms, they are frequently misunderstood and over-hyped in physical security. While APIs can provide great benefits, using them is much more complex than often mentioned in sales calls and magazines.

The goal of APIs in physical security is to allow different systems to work together. Examples include:

- Integrating your IP cameras or analytics with your VMS
- Integrating your DVR/NVR with your access control system
- Integrating your alarm system with a central monitoring system
- Building a PSIM system that integrates with all your security systems

You often hear APIs discussed in pre-sales situations where a customer or integrator asks a vendor: "Does your system work with 'X'?" where X could be any number of security systems by any number of manufacturers.

The routine answer by the sales person is:

"Sure, we have an API."

This is one of the riskiest and misleading statement in all physical security. Because it is so common, it is a great place to start reviewing APIs.

Lesson #1: No such thing as 'an' API

There is no such thing as 'an' API. Numerous APIs exist. In larger systems, hundreds of APIs exist. Generally, there is an API for each function in a system. Want to watch live video, use the live video API. Want to change the time, use the time change API. Want to increase the frame rate for recording, use the recording frame rate API, etc.

Lesson #2: Not all functions have an API

Here's the first issue. Not all functions have an API available. Say you need a list of all health monitoring alerts from another application. This application may have 'an API' but not a specific API for sending health alerts. As you can imagine because most systems today have hundreds of functions, it is common that dozens of these functions are not accessible via an API.

This happens frequently with integrating cameras and VMSes. A camera's API might lack certain advanced configuration functions (like setting the shutter speed, gain control, ABF, day/night, etc.). Even if you integrate the two devices, you may have painful limitations. Make sure to check each function you want integrated to ensure that it can be done.

Lesson #3: Having an API does not mean it will work with your system

Let's say you want to integrate a leading VMS and access control system. Both may have APIs but there is no guarantee that these two products will work together. Both companies having APIs is a pre-requisite for integration but it is not sufficient.

Companies can block or make it difficult to integrate with competitors. Want to integrate Milestone's XProtect VMS within Genetec Security Center? They both have APIs and could do this technically but they are highly unlikely to allow this, as they control who can access their APIs. For more, see our post and discussion on "Integrating Recorders With Other Recorders."

While APIs are typically free for web services or charge a low fee per use, some security manufacturers often charge thousands of dollars to even access theirs. This is especially likely if it is a niche video, access or alarm system.

Even if you have permission to use another company's API, it is not guaranteed to work. Unlike big web services, like Google or Amazon, only a small number of companies use these APIs, which means they could be immature or have undocumented issues. This leads us to the next point.

Lesson #4: Doing the integration takes time

Vendors often claim a few weeks for integration. This can happen but often business and technical details need to be worked out that can take significantly longer.

Indeed, unlike APIs for web services, you rarely can get the API off the Internet. It often requires a formal application and review of the prospective integration partner. Plus, when integrating with a security system API, you cannot simply Google or use StackOverflow to solve problems. The only people who likely can resolve any technical difficulties are within the security manufacturer themselves.

As such, be careful of the time and dollar amount you commit for such projects. This is the type of risk that is often unknown and unknowable until you dig into the technical details about how each vendor implements their APIs. Generally, these projects are ultimately successful, but the time and cost can vary.

Lesson #5: API changes can break you

Just like a product, over time, APIs change. The difference is with APIs, their change can break your system. Reasons for change include eliminating bugs, enhancing performance, adding in new functionalities.

However, other systems depend on those APIs. Let's say your system works with "Vendor B" version 3.1. Now let's say "Vendor B" comes out with 3.2 but this version "breaks the API". In other words, the new version is not backwards compatible with the old version. Your system could suddenly stop working with "Vendor B" if you upgrade Vendor B to version 3.2. The result is your security command center no longer displays video or access or whatever the system that just got the upgrade.

This happens frequently with cameras. When you connect a camera and it does not work, often vendors will say, "Oh you need the new firmware" even if those two manufacturers have been integrated for years. This is usually because changes were made to the API that 'broke' the existing connection.

Lesson #6: You are stuck with what the API does

Unless you are a very large customer, you are stuck with whatever the API does in whatever way it does it. Often, for what you need, this works out fine. However, if you need some change for your specific use case, this can be hard to accomplish. Make sure someone on your technical team knows specifically what the API can and cannot do so you can anticipate any potential problems up front. If a change needs to be made, the change will usually take a lot of time and testing because the vendor must ensure that they do not break the 1000s of other security organizations using this API.

Plus, if they believe only you really needs this, it will be considered low priority unless you are bringing a massive deal to the table or a briefcase filled with cash.

Examples of APIs

The most well known and well used API in video surveillance is Axis's VAPIX. Its current version is divided into 20+ different sections with numerous functions inside each. For instance, here's their 38 page PTZ API document. Go to page 24, Section 3.4.1 to see how to control the PTZ including what parameters and values you need to pan, tilt, zoom and more.

Pan Tilt Zoom API

speed=<*int*>	1 ... 100	Sets the move speed of pan and tilt.
imagerotation=<*int*>	0 90 180 270	If PTZ command refers to an image stream that is rotated differently than the current image setup, then the image stream rotation must be added to each command with this argument to allow the Axis product to compensate.
ircutfilter=<*string*>	auto[1] on off	Control the IR cut filter.

To fully control an Axis IP camera can require hundreds of different operations and implementation of numerous functionalities.

On the web, there are literally thousands of different APIs available. See this directory of 16,000+ APIs. A few examples include webinar APIs, email APIs, mapping APIs, etc., etc.

APIs vs SDKs

Often, the terms API and SDK are used in conjunction when discussing software integration.

- API - Application Programming Interface - the source code based specification, minimally you need this to integrate two systems
- SDK - Software Development Kit - the documentation and development tools that support and help 3rd parties to use the API

See this blog post on APIs vs SDKs.

For security and surveillance systems, you typically want and need an SDK along with the API to help you understand how to use it. Plus, unlike web APIs that are fairly light weight and easy to call from different languages/frameworks, security systems often require (or at least favor) using a specific framework (often .NET or JAVA) and provide tools to make it simple to use those frameworks.

APIs and Standards

All the examples so far of APIs have been proprietary ones, developed and controlled by specific companies. Each one is typically (at least slightly) different than others even if they are doing the same things (requesting a live video feed, getting an alert, etc.). Because of this, it is very time consuming to integrate with multiple systems as the work needs to be done over and over again for each new camera, VMS, recorder, etc.

When 'standards' are discussed in surveillance, like ONVIF, these are APIs that a broad cross section of manufacturers can use. If each proprietary API is its own 'language', an offering like ONVIF is meant to be a universal language. If each manufacturer speaks that language (API), than they can all communicate without having to learn every unique language in the world. Standards are still APIs, just one that, hopefully, all parties will use to 'standardize' how systems communicate. Next, see our ONVIF tutorial.

ONVIF

ONVIF is well known within the surveillance industry as an interface to connect IP cameras and VMS systems but:

- Is ONVIF a 'Standard'?
- Why ONVIF?
- What does ONVIF do? What does it consist of?
- Who supports ONVIF? What does that mean?
- What are the type of ONVIF? S vs G vs Q vs T
- How well does ONVIF work?
- Does ONVIF work with H.265?
- What about advanced features?

Inside this tutorial, we answer all of these questions. Note: If you are not familiar with APIs, you must read our API tutorial first as understanding ONVIF's capabilities depends on knowing how APIs work.

ONVIF As A 'Standard'

ONVIF is a trade organization founded by Axis, Bosch and Sony in 2008 with 500+ members that has developed API specifications for integrating security products. These specifications are being broadly used by hundreds of manufacturers and more than 5000 surveillance products.

Given its broad support, ONVIF acts as a 'de facto', in practice, standard. In 2013, ONVIF and PSIA were both recently included as part of a European standard though it is not clear what impact that has.

Here are detailed ONVIF technical answers based on an IPVM interview with conducted with them.

ONVIF vs. PSIA

In 2008, two trade organizations launched with the same goal of bringing interoperability 'standards' to physical security - ONVIF and PSIA. While PSIAlaunched first, ONVIF had stronger support among leading IP camera manufacturers (notably Axis) that drove them to an insurmountable lead in products supporting their specification. Though PSIA continuous to exist and has refocused on developing 'standards' for access and intrusion, ONVIF has long won the battle for IP camera / VMS interoperability.

Why ONVIF?

Doing custom integrations between IP cameras and VMS systems is time consuming and expensive. Worse, in a fragmented market like video surveillance, with hundreds of manufacturer offerings being used, this forced legions of integrations. It is very hard even for large VMS developers to keep up. Equally problematic, new camera manufacturers were blocked by lack of integrations with widely used VMSes. With ONVIF, the goal is that each side just writes 'once' to ONVIF and can then integrate with every other product on the other side.

What Does ONVIF Do?

ONVIF specifies in detail how network video transmitters (such as IP cameras and encoders) can integrate with network video clients (such as VMS software and NVRs). It is an API that details dozens of methods across a core and numerous service specifications. ONVIF's functionalities are like those of proprietary APIs from camera manufacturers, defining how clients can authenticate, change IP addresses, request video feeds, pan, tilt, zoom, send events, etc. The major difference is that this specification can and is used by many manufacturers.

ONVIF vs Proprietary APIs

IP cameras and VMS systems can and frequently use both ONVIF and their own proprietary APIs. Since ONVIF is relatively new, most manufacturers have had to have their own APIs. Plus, manufacturers can offer different or broader functionality by using their own. One example of this is panoramic video dewarping, important for many fisheye / 360 cameras.

ONVIF Detailed Functionalities

That noted, ONVIF supports a very broad range of functionalities. Their specification map outlines various services ONVIF supports, such as device IO, PTZ control, recording, video analytics, etc. Two of the most fundamental are imaging and media which allow configuring / setting dozens of video / camera properties.

ONVIF vs NTSC / PAL

Loosely speaking, ONVIF is the IP video equivalent of NTSC / PAL, though ONVIF offers far more advanced functionality with much higher complexity. NTSC / PAL are unidirectional, defining a uniform video stream. This, combined with it being 50+ years old, made it very easy to use and reliable. However, resolution and frame rate are locked, and no controls were available. If you wanted to control or use I/O, PTZs, analytics, adjust camera settings, etc., that had to be done separately as it was excluded from the specification. At a high level, ONVIF is a replacement but it brings with it a lot more benefits and problems.

Device vs. Clients

Products supporting ONVIF are divided into two fundamental types / names:

- A 'device', most typically an IP camera, is a product that responds to ONVIF requests. Devices are also sometimes encoders, recorders or access control panels.
- A 'client', most typically a recorder / VMS, is a product that makes ONVIF requests.

The most common scenario would be a client, like Milestone, Genetec, Exacq, etc. making an ONVIF request to an IP camera, like Axis, Bosch, Sony, etc.

Sometimes, a VMS / recorder can act as a 'device', streaming out video / responding to requests from other systems.

Types of ONVIF - Archived

It is critical to understand the differences between 'types' of ONVIF. The most critical is differentiating between the older 'Archived' version and the newer Profile one.

ONVIF has had 2 major releases to date - version 1.x and 2.x. Every product that is for 1.x is now considered archived (see ONVIF Archives 1000+ Products). In our testing, archived products are far more likely to have integration issues. Indeed, as of late 2015, ONVIF stopped even showing 'archived' products.

Profiles for ONVIF

Now, ONVIF uses a series of Profiles, allowing conforming products to support multiple Profiles. The key profiles, in order, are:

- **S** is the oldest and most broadly supported profile, covering video **s**treaming, the basics of sending video from a camera to a VMS / recorder. S is what most everyone supports
- **G** is a profile added to support access video stora**g**e. For example, this could support retrieving and sending from an IP camera with on-board storage to a VMS / recorder. As of May 2016, it has minimal official conformant products.
- **Q** is a newer profile that aims to simplify discovering cameras and improve security, by eliminating default passwords. As of the end of 2016, it is officially adopted, but has little manufacturer support.
- **T** is a future ONVIF profile (planned for 2018) that will add H.265 support.

CODECs Support

ONVIF today (via Profile S) supports MJPEG, MPEG-4 and H.264. It does not support others, such as open H.265 or JPEG2000 standards, nor proprietary CODECs such as Mobotix's MxPEG.

H.265 Supported, But No Conformance Test

As of version 2.4 (June 2016), developers may support H.265 and other CODECs not previously supported via the new Media2 service. This new specification uses IANA Media Types, standards which include surveillance formats such as H.264, H.265, and MJPEG, to define CODECs, so future additions may be adopted without rewriting the profile.

However, the Media2 service will not be included in a profile until Profile T is released, currently targeted for Q1 2018. Manufacturers may add H.265 support prior to then, but there will be no way to test for proper operation aside from connecting cameras and recorders to see if they work.

See our How And When ONVIF Will Support H.265 report for more information.

Services / Functionalities

ONVIF also includes various services, many of which are optional or less broadly supported. Three of the more notable ones are:

- Analytics service
- PTZ service
- I/O service

It is important to keep in mind that direct proprietary integrations often cover these less common / advanced integrations. By contrast, most ONVIF implementations do not and this can be a source of frustration.

False Claims of ONVIF Conformance

The correct and legitimate method for gaining ONVIF conformance is for a manufacturer to use the current version of the ONVIF test tool and to submit passing results to ONVIF.

Most manufacturers follow this process, but ONVIF has been historically lax in enforcing conformance, resulting in significant numbers of manufacturer fake claims. Be especially careful of super inexpensive, no-name products that claim ONVIF conformance. One can quickly check by looking it up on the ONVIF official conformance directory.

Self Testing for ONVIF

ONVIF conformance testing is done by each manufacturer and is not verified by ONVIF or any third party. Each manufacturer uses the ONVIF test tool. The test tool is fairly rudimentary and not robust. This video overviews it:

Note: Click here to watch the ONVIF Test Tool video on IPVM

If the test passes, the manufacturer generates a passing report and ONVIF publishes it on their conformance directory.

One notable problem is that ONVIF reports can be faked and we have heard reports of manufacturers doing so.

ONVIF Basics Working In Production

That noted, ONVIF is broadly available in production and generally works, at least for the basics. Almost all major manufacturers support it, with over 7,000 total devices now, up from 6,000 in 2016, 4,000 in 2015, 2,700 in 2013, 1,000 in 2012 and just 400 at the end of 2010. The growth has been strong.

For connecting and streaming video from cameras to VMSes, ONVIF worked 90% of the time in IPVM's test of 14 camera manufacturers and 5 recorders. The biggest issues were with Archived products, specifically for motion detection, where no integrations work, compared to Profile ones, where 50%+ did.

For more, see IPVM's ONVIF Mega Test

Advanced Features Problems

While ONVIF has gotten strong at the basics of connecting and streaming video from cameras to recorders, advanced features, like motion detection, have a significant risk of failure. Moreover, VMS manufacturers routinely have to add custom integration to support motion detection via ONVIF to each particular manufacturer.

PTZ control, I/O and video analytic all suffer from equal or greater integration problems, even with the newer Profile S.

VMD Issues Beware

The most common real world problem is having ONVIF stream video fine but not being able to send VMD events from the camera to the recorder. This is especially problematic as many VMSes do not offer server side VMD, yet motion based recording is commonplace.

Recommendations on Use

We recommend continuing to use proprietary 'direct' interfaces if a combination supports both ONVIF and the proprietary one as it reduces risk and often provides extra functionalities. However, when looking at new cameras or VMSes, ONVIF support is an important factor in expanding what one uses. Certainly, one should test and verify their preferred combination but it is likely that ONVIF will help ease integrating heretofore incompatible devices.

[NOTE: This tutorial was originally written in 2014 but was substantially updated throughout 2016 to reflect advances in ONVIF as well as to add further details.]

Network Security

Keeping surveillance networks secure can be a daunting task, but there are several methods that can greatly reduce risk, especially when used in conjunction with each other.

We look at several security techniques, both physical and logical, used to secure surveillance networks, including:

- Network Hardening Guides
- Passwords
- LDAP / Active Directory Integration
- VLANs
- 802.1X Authentication
- Disabling Switch Ports
- Disabling Network Ports
- Disabling Unused Services
- MAC Address Filtering
- Locking Plugs
- Physical Access Control
- Managing Network Security For Video Surveillance Systems

Network Security Critical

More than ever, network security has become a key issue, with published vulnerabilities, hacks, and botnets on the rise.

In just the past 1-2 years, major vulnerabilities (and their effects) were reported in multiple manufacturers, including:

- Hikvision Backdoor Exploit: A hardcoded backdoor which allows attackers full control of Hikvision IP cameras.
- Dahua Hard-Coded Credentials Vulnerability: Hard-coded credentials were found in firmware for cameras and NVRs, allowing for rogue firmware uploads.
- Axis Critical Security Vulnerability: A vulnerability allows attackers to remotely initiate a telnet connection, allowing the attacker to take over the device, reboot it, power it down, etc.
- Hacked Dahua Cameras Drive Massive Cyber Attack: As part of the Mirai botnet, hacked Dahua cameras (and others) took down major internet sites and even an entire country.
- See our Listings of Video Surveillance Cybersecurity Vulnerabilities and Exploits for more information on these and other issues, including new ones as they occur.

In previous years, incidents were few and far between, with Hikvision the most notable target (see: Hikvision Hacking And Chinese Province Warning, The Hikvision Hacking Scandal, The Hikvision Hacking Scandal Returns, finally resulting in their "Anti Hacking" Firmware).

Because of the severity of these incidents and their increasing frequency, it is critical that users understand the basics of cyber security for surveillance systems, and how to protect against simple attacks at the very least.

Network Hardening Guides

In the IT industry at large, network hardening guides are common, outlining recommendations (as an example, see this Cisco hardening guide) to make the network more secure. Many/most of these recommendations apply to surveillance networks, as well, including controlling physical and login address, securing passwords, disabling ports, etc.

However, many recommendations may be above and beyond what many IP video integrators are capable of, or what is practical for a given system. Complex authentication schemes such as 802.1x, LDAP integration, SNMP monitoring, etc., are simply not worth the time/cost to implement for many systems, given the limited risk.

Surveillance Hardening Guides Rare

Unlike IT, surveillance specific hardening guides are rare, with only a handful of guides available from manufacturers.

- Axis cyber hardening guide
- Bosch IP Video and Data Security Guidebook
- Dahua best practices
- Genetec cyber hardening guide (requires partner login)
- Milestone cyber hardening guide

The exact recommendations in each of thees guides vary, but most are divided into basic and advanced levels, depending on the criticality of the installation.

The Axis guide, for instance, varies from demo only (not production use) to highly secure enterprise networks, and include basic best practices, such as strong passwords, updating firmware, and disabling anonymous access, through more complex practices, such as 802.1x authentication, SNMP monitoring, and syslog servers.

While the these guides are manufacturer-specific, providing instructions pertinent to the camera or VMS, many recommendations are useful across all manufacturers, and fall in line with IT industry best practices, and the practices discussed below.

Strong Passwords

Strong passwords are the most basic security measure, but unfortunately, ignored by many users. Many surveillance systems are deployed in the field with default passwords on all equipment, including cameras, switches, recorders, and more (see our IP Cameras Default Passwords List). Doing so may make it easier for techs to access cameras but also make it simple for anyone to log into one's cameras (see: Search Engine For Hacking IP Cameras).

At the very least, all surveillance network devices, including cameras, clients, and servers, should be changed from the defaults with strong passwords, documented in a secure location. This prevents access to the network using simple password guessing, requiring a more skilled attacker and more complex methods.

Some manufacturers require changing the default password when connecting for the first time (see a comparison of how Axis, Dahua and Samsung set passwords). Indeed, an upcoming ONVIF Profile (Q) would make changing default passwords mandatory, though how well that is adopted remains to be seen.

LDAP/AD Integration

Using LDAP/Active Directory (AD) integration, VMS permissions are assigned to network users managed by a central server (also called single sign-on). Since these user accounts often implement password strength and expiration rules, this integration may improve security over local VMS accounts which do not have these restrictions. This reduces administration overhead, since individual accounts do not to be created and maintained.

Typically, LDAP use is restricted to larger, enterprise systems, since many small installations do not have an LDAP server implemented. Some small or midsize systems which are installed in larger entities, especially education and corporate facilities, may use LDAP as these organizations are likely to use it for their network access control.

LDAP / AD could theoretically be used for IP cameras, but, in practice is not. ActiveDirectory, as a Microsoft offering, is not supported by almost any IP camera, which typically run on Linux. One Windows IP camera claimed to do so, but it has not gained any meaningful market share.

Firewalls/Remote Access

To prevent unauthorized remote access, many surveillance systems are not connected to the internet at all, instead on a totally separate LAN. This reduces risk, but may make service more difficult, as updates to software and firmware, usually simply downloaded, must be loaded from USB or other means.

Those systems which are connected are typically behind a firewall, which limits inbound/outbound traffic to only specific IP addresses and ports which have been authorized. Other traffic is rejected. Properly implemented, this may prevent the vast majority of attacks.

Remote Access Risks

For devices which require remote access, VMSes and cameras may require one or more ports to be open. However, each open port presents a possible opportunity for an attacker. Exactly how many and which varies by the VMS. Users should refer to manufacturer documentation for which ports must be open if remote access is required (for maintenance or remote viewing), and we list some examples in our Network Ports for IP Video Surveillance Tutorial.

P2P/Cloud Access

Alternatively, some manufacturers allow for "phone home" remote access, which sets up a secure tunnel via an outbound connection without requiring open ports, reducing risks. Many cameras and recorders use cloud connections for remote access, such as Hikvision EZVIZ, Eagle Eye Cloud VMS, and Genetec Cloud. Additionally, many remote desktop services use similar technology, such as LogMeIn, TeamViewer, SplashTop, etc.

We discuss these methods in our Remote Network Access for Video Surveillance tutorial.

VLANs

Virtual LANs (shortened to VLANs) improve security by segmenting traffic into multiple virtual networks. So while other services, such as IP based surveillance equipment or general office LAN traffic, may exist on the same physical switch, for practical purposes the networks are invisible to each other, and unreachable.

For example, in the image below, the camera and NVR on VLAN 1 may not be reached by the office PC on a separate VLAN, nor could a user on the NVR (VLAN 1)"see" traffic on the PC VLAN (VLAN 2).

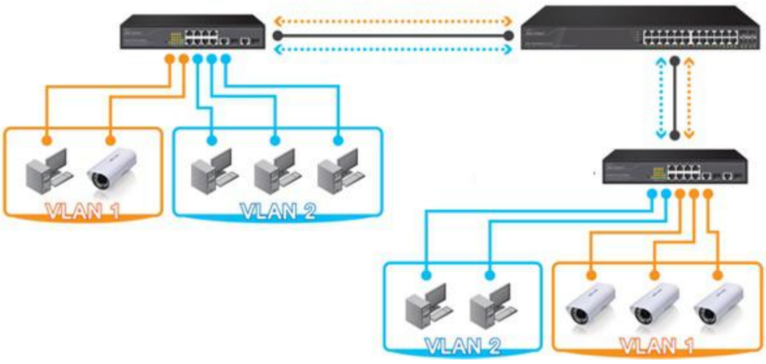

VLANs are most commonly set up using [802.1Q tagging](), which adds a header to each frame containing VLAN information. This header is interpreted by the switch and traffic forwarded only to other devices on the same VLAN.

Note that while traffic may not be intercepted across VLANs, bandwidth constraints still exist. Numerous large video streams may negatively impact VOIP and office application performance, while large file transfers may affect the surveillance network. Because of this, VLANs are also most often deployed in conjunction with [Quality of Service (QoS)](), which prioritizes network traffic, sending video packets ahead of file transfers, for example, so video quality is not impacted.

See our [VLANs for Surveillance]() guide for further information.

Disabling Unused Switch Ports

Another easy but typically overlooked method of keeping unauthorized devices from accessing a switch is to disable all unused ports. This step mitigates the risk of someone trying to access a security subnet by plugging a patch cable into a switch or unused network jack. The option to disable specific ports is a common option in managed switches, both low cost and enterprise:

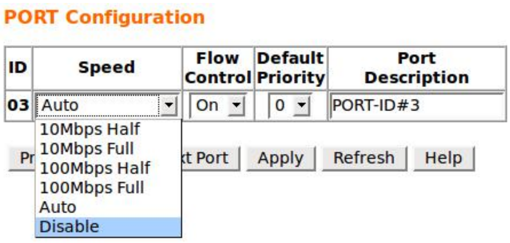

While effective at narrowing the number of potential access points, this step does not necessarily prevent unauthorized access to a network, as someone could potentially unplug a device (camera, workstation, printer) from a previously authorized port or jack and access its port, unless measures such as MAC filtering or 802.1X are in place.

Disabling Unused Network Ports

Many cameras ship with unneeded network ports turned on, such as Telnet, SSH, FTP, etc., as we found in our NMAPing IP Cameras Test. These ports are favorite targets of hackers (as illustrated by bitcoin miners and buffer vulnerabilities found in Hikvision Cameras).

A quick 30 second scan of a popular IP camera reveals multiple open ports other than those expected for web access and video streaming (80/554):

```
Starting Nmap 6.47 ( http://nmap.org ) at 2015-03-04 16:47 EST
Nmap scan report for 172.20.128.123
Host is up (0.0023s latency).
Not shown: 994 closed ports
PORT       STATE SERVICE
23/tcp     open  telnet
80/tcp     open  http
554/tcp    open  rtsp
3800/tcp   open  pwgpsi
5000/tcp   open  upnp
49152/tcp  open  unknown
MAC Address: 90:02:A9:08:14:8A (Zhejiang Dahua Technology Co.)

Nmap done: 1 IP address (1 host up) scanned in 5.92 seconds
```

These ports should be disabled wherever possible to prevent potential attacks.

Disabling Unused Services

Unnecessary services on viewing workstations and servers should be turned off. These may include manufacturer-specific update utilities, various Microsoft update services, web services, etc. These unneeded services may act as a backdoor for hackers or viruses, consume additional processor and memory, and increase startup time.

These services should be disabled or set to operate only when manually started, as seen here in Windows:

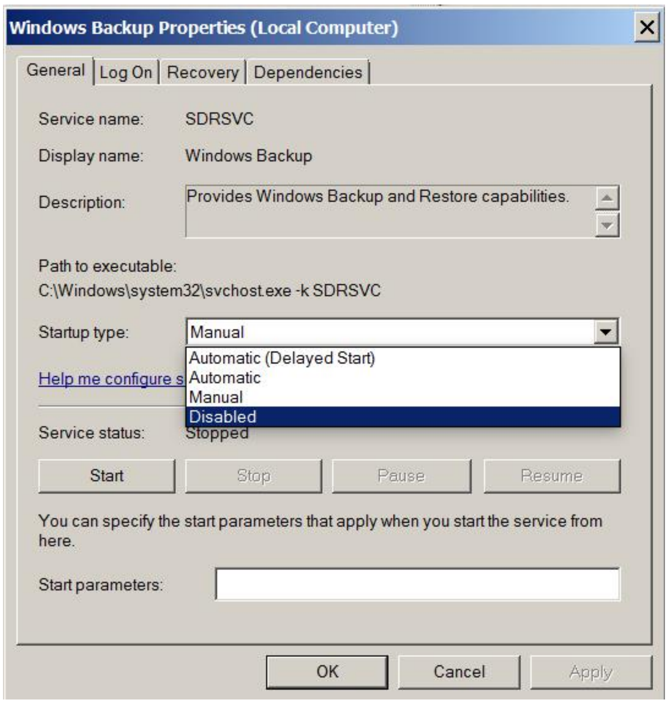

OS and Firmware Updates

OS and firmware updates are a matter of some debate, with some users installing every available Windows Update, for example, while others insist that these updates may break VMS software or camera integrations.

However, these updates (especially Windows Update) often include patches to newly discovered security vulnerabilities, such as the Heartbleed SSL vulnerability, which affected millions of computers worldwide. Patches for these significant issues should be installed.

Other, more routine, updates may be optional. Users especially concerned about compatibility issues should contact their camera/recorder/VMS manufacturers to see their recommendations for applying updates or not.

MAC Address Filtering

MAC address filtering allows only a specific list of devices to connect to the switch. Other devices plugged into the switch are ignored, even if the port previously was used by a valid device. MAC filtering is possible only using managed switches.

In surveillance networks, MAC filtering is typically easy to administer. Once all cameras, clients, and servers are connected, it is enabled, and connected devices' MACs added to the whitelist. Since these devices in a surveillance network are rarely changed out, little extra maintenance is required. In other networks where devices may frequently be added or removed, administrators may find filtering more cumbersome to administer.

This image shows MAC filtering options in a typical managed switch interface:

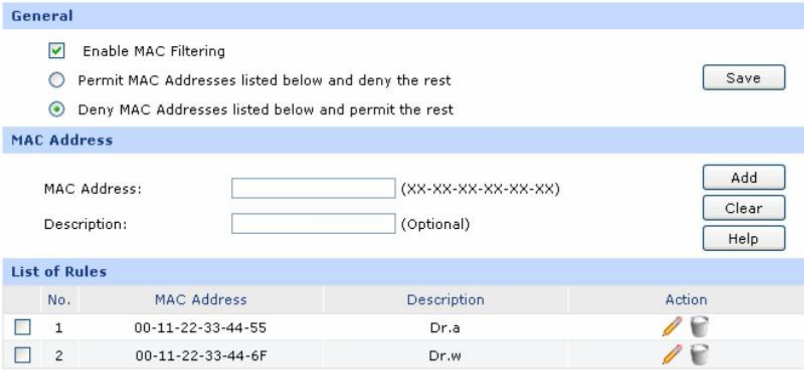

See our [Network Addressing for Video Surveillance Guide]() for more discussion and a basic overview of MAC addresses.

802.1X

802.1X requires devices trying to connect to the network to have proper credentials to be allowed on. This blocks random devices or attackers from just jumping on a network.

Using 802.1X, a "supplicant" (client such a camera, PC, etc.) attempts to connect to network via a switch or WAP (called the "authenticator"). The authenticator then checks the credentials of the supplicant with a server, call the authentication server (typically using a protocol called RADIUS, and grants or denies access accordingly.

While 802.1X provides strong security, setting up a network to support it can be cumbersome and involved. Not only must connected devices (cameras, WAPs, client PCs, NVRs, etc.) support 802.1X integration, all switches must, as well. Each of these devices must be individually configured for 802.1X, adding additional configuration time to the install.

Because of these factors, which increase cost and administration overhead, 802.1X is rarely used in all but the most complex enterprise surveillance networks, with users opting for simpler security measures instead.

Locking Plugs

Another layer of security that physically prevents connection or tampering with network cabling by unauthorized devices are port plugs and cable locks. These devices mechanically lock a cable into a switch, patch panel, or wall jack, or fill unused switch ports, and may only be removed with a proprietary tool.

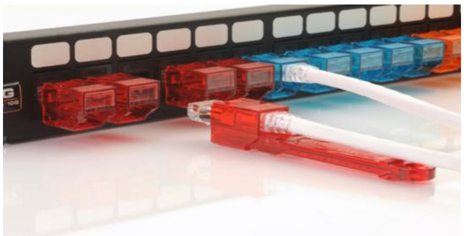

While these types of locks are effective at stopping casual tampering, they are not unbeatable or indestructible, and a determined intruder may simply be able to force them out or pry them loose given enough time. As such, locking plugs should be considered part of a good network security program, but not the only element.

For a deeper look, read our Locking Down Network Connections update.

Door Locks and Physical Access

Finally, best practices call for controlling access to the most vulnerable areas of a network, the rooms, closets, or racks where surveillance servers and switches are typically mounted. By reducing the potential availability of these areas, many risks from determined or even inadvertent threats can be avoided. If doors cannot be secured, individual rack cages or switch enclosures should be. Most modern IT cabinetry includes security equipment as standard options:

As a result, many facilities employ electronic access control on server or network equipment rooms. However, even non-exotic mechanical keys and locks can do a great job of protecting sensitive areas when properly managed.

Managing Network Security For Video Surveillance Systems

While all the steps below may improve security on their own, they are most effective when documented as part of a written (and enforced) security policy.

In surveillance, this policy is up to the individual install, but generally it comes from one of two places:

- End user: When the surveillance network is part of a larger corporate/enterprise LAN (whether sharing switches or dedicated), end users most likely control the security policy for all network devices, and may force these requirements upon integrators (for better or worse).
- Integrator: If an end user does not have a security policy in place, the installing integrator may choose to create one as part of their documentation, requiring it to be followed in order for the warranty to be enforced and limit liability in case of a breach.