

WHITE PAPER

Understanding IP Surveillance Camera Bandwidth



Table of Contents

Introduction	4
What Is Bandwidth?	4
Video Compression	4
Motion JPEG	5
H.264	5
H.265	5
Variable bitrate (VBR)	5
Constant bitrate (CBR)	5
Digital noise reduction (DNR)	5
Image Quality Level	6
Complexity of the Scene	7
Complex scene	7
Simple scene	8
Video Resolution	9
What do I need to know about resolution?	9
Common resolutions in IP surveillance cameras	10
Frame Rate	10
Number of Cameras and Clients	10
Number of viewing clients	10
Continuous recording versus event-based recording	11

- [Bandwidth Bottlenecks](#)11
 - [Local-area network \(LAN\)](#)11
 - [Wide-area network \(WAN\) or internet](#)12
- [Surveillance System Network Best Practices](#)12
 - [Independent switch](#)13
 - [Switch selection](#)13
- [Configurations on FortiRecorder That Affect Bandwidth](#)14
 - [Video profile setting](#)14
 - [Camera profile setting](#)15
 - [Digital noise reduction \(DNR\) setting](#)17
 - [DNR schedule setting](#)18
- [FortiRecorder Statistic](#)19
- [Configurations in FortiCentral That Affect Bandwidth](#)21
- [Analyzing Bottlenecks with FortiCentral Performance Statistic \(Gauge\)](#) . .23
- [FortiCamera Calculator](#)24
- [Summary](#)25

Introduction

Planning the bandwidth requirements of an IP video surveillance project is among the most critical design steps, similar in importance to placing the cameras. Without careful planning ahead of time, the surveillance system might end up with a bandwidth bottleneck. This not only causes video packet loss, delay, or jitter but also degrades video quality, or even worse, inhibits recording of critical incidents. Bandwidth also determines the storage capacity requirements for a given retention period.

Understanding video bandwidth takes in-depth knowledge of several fields. This white paper is intended to provide a fundamental knowledge of what affects video surveillance performance. In general, there are six factors that influence an IP security camera's bandwidth:

- Video compression
- Image quality level
- Complexity of the scene
- Video resolution
- Frame rate per second
- Number of cameras and viewing clients

The first half of this white paper will discuss the importance of these six factors. Based on this, the second half of the paper relates these six factors to FortiRecorder and FortiCentral software configuration. A brief introduction to the FortiCamera bandwidth calculator will also be given.

What Is Bandwidth?

IP video is transmitted as a stream of data that contains the image, audio, and control data of the camera. The amount of data that has to be transmitted per second is called bandwidth. It is commonly measured in Mbit/s, which makes it easy to compare to the bitrate capacity of an Ethernet link. For example, 10 Mbit/s is called Ethernet, 100 Mbit/s is Fast Ethernet, and 1,000 Mbit/s is Gigabit Ethernet.

Another measuring unit is MByte/s, equal to 1/8th of the bitrate because there are 8 bits in a byte.

1 Mbit/s = 1,000 Kbit/s = 125 KByte/s

1 Gbit/s = 1,000 Mbit/s = 125 MByte/s

A 1920 x 1080 HD resolution camera roughly generates raw video data at 1.49 Gbit/s (30 x 1920 x 1080 x 24) for 30 FPS video. That is 178 MByte/s of data and the reason video compression is required.

Video Compression

There are two categories of compression: lossless and lossy. Lossless compression uses an algorithm that can reconstruct the compressed data to its original. For lossy compression, the compressed data can only approximately be restored to its original information. The difference creates visual artifacts. Examples of lossless compression are .zip for data and .gif for picture compression. Lossy compression formats are JPEG for pictures and MP3 for audio as well as Motion JPEG (MJPEG), H.264, and H.265, all of which are described in more detail below.

Similar to audio compression, video compression is used to pack the visual information in as little data as possible while keeping as much human relevant image quality. There are two purposes for utilizing video compression: faster transmission and less storage required. In modern IP surveillance cameras, MJPEG, H.264, and H.265 are the most commonly used video compression. Next, a more detailed explanation about how each video compression works will be provided.

Motion JPEG

By putting multiple pictures in a series, you get a motion picture or video. Did you ever draw a stick man on the border of a book, so each page has a little movement from the previous one? When you flip the pages very fast, you get a motion picture.

MJPEG is a frame-by-frame compression of such an image sequence. Each frame is compressed individually like a static JPEG image. This makes it easy to decode the stream, and every frame is independent of the other. But MJPEG will result in much higher bandwidth usage than, for example, H.264 compression. MJPEG video being 10x the size of an H.264 is reasonable.

H.264

H.264, which is also called MPEG-4 AVC, is a compression standard that was introduced in 2003 and is the prevalent standard used in video surveillance cameras and many commercial media applications.

In contrast to the frame-by-frame approach of MJPEG, H.264 stores the full frame only at intervals of, for example, once a second and encodes the rest of the frames only with the differences caused by motion in the video. Full frames are called I-frame (also index frame or intra-frame) and the partial ones containing only the difference to the previous frame are called P-frame (also predicted frame or inter-frame). P-frames are smaller and more numerous than I-frames. There is also a B-frame (bidirectional frame), which refers both ways to previous and subsequent frames for changes. The recurring pattern of I-P-B frames is called a group of picture (GOP). The time interval for I-frames varies and can range from multiple times a second to nearly a minute. The more I-frames are transmitted, the larger the video stream will be, but it makes restarting decoding of a stream easier since this can only happen at an I-frame.

H.265

H.265, also called High Efficiency Video Coding (HEVC), was developed as the successor to AVC H.264. This latest video compression standard in the MPEG family offers a 25% to 50% higher compression efficiency. H.265 allows streaming video at lower frame rates while maintaining the same video quality compared to H.264. The improvements make compression more complex, requiring newer hardware, but the decompression complexity is similar to AVC.

As H.265 is more efficient than other codecs, it takes up less bandwidth and much less storage space on a hard drive while maintaining the same video quality. You'll be able to store much more surveillance video in less space, with clearer images and video at the same time.

Variable bitrate (VBR)

The strength of each compression method can be adjusted. In general, higher compression causes more artifacts, so there are different strategies to achieve a desired behavior.

When VBR compression is used, the size of the compressed stream is allowed to vary to maintain consistent image quality. Thus, VBR can be more suitable when there is motion in the scene and it tends not to be constant. The disadvantage is that the bandwidth can, to a certain extent, vary depending on the situation. So storage may be used up earlier than planned or transmission bottlenecks can appear when cameras suddenly require more bandwidth.

In VBR, there is no firm cap being placed on the bitrate. The user sets a certain target bitrate or image quality level.

Constant bitrate (CBR)

When configuring a camera for CBR, the camera is set to have constant bandwidth consumption. The amount of compression applied increases as more changes are occurring. This can add compression artifacts to the image and degrade image quality. With CBR, image quality will be sacrificed to meet the bandwidth target. If the target is reasonably set, this degradation may be hardly noticeable and it gives a stable basis for calculating storage and planning the network.

For IP surveillance cameras installed in a local-area network (LAN) with low network utilization or when storage space is abundant, VBR is recommended to maintain best image quality, whereas CBR can help control bandwidth-restricted environments.

Digital noise reduction (DNR)

Just as a high-complexity scene makes it harder to compress, so does the presence of noise. This is a particular concern for low-light scenes that tend to show increased visual noise like speckles and graininess. Left uncontrolled, the video will either increase to a very high bandwidth or the image will show a lot of compression artifacts, occluding important details.


DNR can help reduce the noise; however, since too much DNR will increase blurring and remove image detail, a balanced setting is needed. If the image is very different between day and night, a scheduled or automatic DNR setting can help adjust to both light situations.

Most FortiCamera models are equipped with DNR or 3DNR (filtering in time and space), with some models offering manual settings ranging from 1 to 10. Learn more from [Configurations on FortiRecorder](#) about how a DNR setting could affect bandwidth.


Image Quality Level

In this section, you will see examples about compression artifacts and how the same scene appears in video with different levels of compression.

The VBR compression level can be set to Extra High, High, Normal, Low, and Extra Low in FortiRecorder systems. Pictures 1 and 2 show the results of using “Extra High” and “Extra Low,” respectively. The difference in video bandwidth and image quality can be seen; higher image quality needs more bandwidth. While minimizing the bandwidth, the image shows clear artifacts and a loss of detail in Picture 2.



Picture 1: Video quality Extra High, average bandwidth is 1.5 Mbit/s.



Picture 2: Video quality Extra Low, average bandwidth is 0.5 Mbit/s.

Complexity of the Scene

The complexity of a scene also affects the bandwidth a video camera generates. Generally, the more complex the scene is, the more bandwidth will be required to achieve a certain image quality. For example, scenes that have tree leaves, wire fencing, or random textures like popcorn ceilings increase the complexity of the scene. Others, like a normal, plain color painted wall or little detail, are considered a simple scene. Similarly, motion or movement increases complexity. People walking by, cars driving across, or tree leaves in a breeze are examples.

Complex scene

We will use some sample pictures to show how H.264 compression works on different complexity scenes. Pictures 3 to 5 below show a frame from HD video with the same complex scene content. The difference is the Mbit/s video stream. The zoomed-in patches show the tree pattern, the car, the hedge, and the wire fence. The sample image in Picture 3 shows that streaming with a bitrate of 1 Mbit/s for this complex scene results in lack of detail. When the video bitrate is increased, the image quality of a complex scene is generally improved. The result yields decent image quality with 3 Mbit/s, whereas 5 Mbit/s in Picture 5 does not increase the image quality much more from Picture 4.



Picture 3: Complex scene—video bitrate at 1 Mbit/s.




Picture 4: Complex scene—video bitrate at 3 Mbit/s.



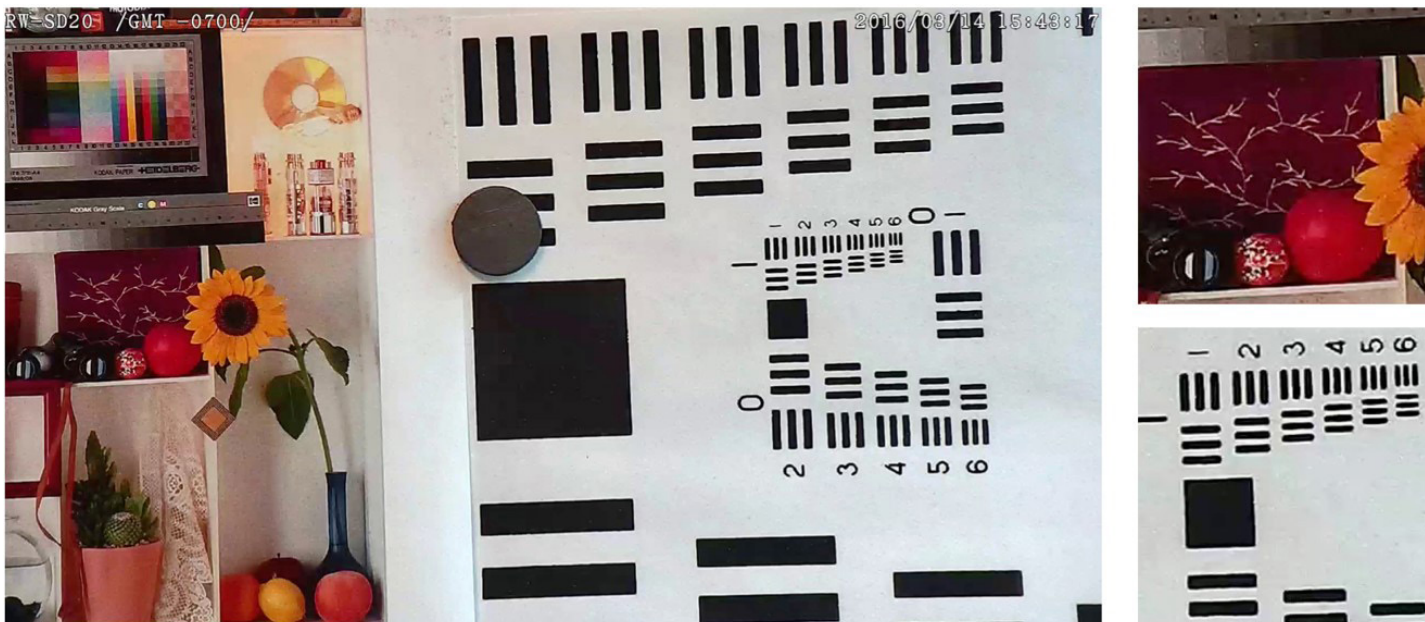
Picture 5: Complex scene—video bitrate at 5 Mbit/s.

Simple scene

For a simpler scene with far less detail, using a higher video bitrate might not be necessary, which can save storage space. Pictures 6 and 7 present the results of a frame from HD video streaming at 1 Mbit/s and 3 Mbit/s, respectively. The zoomed patches of the flower and chart show no obvious differences between each picture. This is because the scene is simple for the H.264 algorithm to encode. 1 Mbit/s is good enough to encode this scene with decent image quality.



Picture 6: Simple scene—video bitrate at 1 Mbit/s.



Picture 7: Simple scene—video bitrate at 3 Mbit/s.

When comparing the complex and simple scene samples, it is obvious that different scenes require different bitrates or bandwidth to show an image at a desired quality level. Lowering the bitrate below such a level will result in loss of detail, while going much higher will eventually not show an improvement. Choosing the right bitrate for the scene is important in order to keep details while not wasting bandwidth.

Video Resolution

What do I need to know about resolution?

Every camera has an image sensor. The available pixels from left to right provide the horizontal resolution, while the pixels from top to bottom provide the vertical resolution. Multiply the two numbers for the total resolution (divide by 1,000,000 to get megapixels) of this image sensor.

What does resolution have to do with the bandwidth?

Assuming 24 bits for the RGB color values of a pixel:

$$640(H) \times 480(V) = 307,200 \text{ pixels} = 0.3 \text{ MP} \times 24 \text{ bits} = 7.2 \text{ Mbit/s}$$

$$1920(H) \times 1080(V) = 2,073,600 \text{ pixels} = 2.0 \text{ MP} \times 24 \text{ bits} = 48 \text{ Mbit/s}$$

Therefore, 1920 x 1080 takes more bandwidth since it contains more pixels, or simply saying, more data. But it gives clearer, sharper pictures when needed to identify a subject, a face, or a car model and its color or license plate. Vice versa, lower resolution generates less bandwidth, but the trade-off is a less clear, blurrier image. Lower resolution usually gives surveillance operators situational awareness—seeing what is going on rather than detail.

Resolution is not the only thing that determines the clarity of an image. Optical performance of the lens, focal length (optical zoom), distance to the object, lighting conditions, dirt, and weather are also critical factors.

Most FortiCamera IP cameras support dual streaming functionality, which means there is one higher and one lower resolution stream available concurrently. FortiCentral software allows the user to select the lower or higher resolution for monitoring purposes. Learn more from [Configurations in FortiCentral That Affect Bandwidth](#).

Common resolutions in IP surveillance cameras

Table 1 is a list of common resolutions available in many IP surveillance cameras. High definition (HD) IP cameras with resolutions above 1280 x 720 offer superior image quality compared to traditional analog systems (CCTV) that are roughly equivalent to 640 x 480 standard definition (SD). Nowadays 2 MP and 4 MP are the most common resolutions in a surveillance camera system because they will satisfy a majority of monitoring requirements.

Format	Resolution (pixels)	Aspect ratio
QVGA	320 x 240	4:3
VGA	640 x 480	4:3
SVGA	800 x 600	4:3
1 MP/720P	1280 x 720	16:9
2 MP/1080P/Full HD	1920 x 1080	16:9
3 MP QXGA	2048 x 1536	4:3
4 MP WQHD	2560 x 1440	16:9
5 MP QSXGA	2560 x 1920	4:3
4K Ultra HD	4096 x 2160	16:9

Table 1: Common IP surveillance camera resolutions.

Frame Rate

Frame rate is measured in frames per second (FPS), which means the number of pictures being produced in a second. The higher the frame rate, the smoother the subject moves in the video. The lower the frame rate, the more jerky movements become up to the point where subjects jump from position to position with a loss of anything in between. Bandwidth increases with frame rate. Half the frame rate usually does not quite reduce bandwidth by half, though, because the encoding efficiency suffers. Modern surveillance cameras can generate up to 60 FPS. However, CPU limitations will sometimes restrict the FPS to a lower value when resolutions are set too high.

Finding the optimal FPS setting for a scene is a compromise between objectives: capture all the relevant information without essential details being lost between frames versus bandwidth considerations.

If a camera is monitoring a quiet overview, there is no need to go up to 30 FPS. A setting of 5 to 15 FPS is sufficient. As a rule of thumb, the more rapid change occurs or the faster subject movement is anticipated, the higher the FPS should be set. Adjust the FPS after cameras are installed and monitor whether the smoothness of the video is acceptable or not.

Number of Cameras and Clients

Obviously the number of cameras influences the bandwidth requirements for a system. If all cameras are the same, then twice the camera numbers will double the data generated. To maintain scalability of a system, it must be able to break large topologies into manageable smaller partitions. By structuring the system in a layered and distributed architecture, it is possible to maintain scalability over a large range of quantities. The key is to distribute bandwidth so bottlenecks are avoided. More will be discussed in the Bandwidth Bottlenecks section.

Number of viewing clients

The discussion above relates to camera bandwidth feeding into the recorder. This is only one side of the picture, where the other side is connecting the recorders to the clients watching live or playback video. For example, there might be a security team that constantly monitors the cameras 24 hours a day, seven days a week. This bandwidth would be equal to all the data coming from the cameras. In the case of playback, even more bandwidth is required if used in addition to live streaming. Considering there can be many clients connecting to a system at the same time, the client-side traffic can be the dominant concern.

Fortunately, if the system is designed according to surveillance system network best practices, the recorder acts as a safekeeper and buffer. It shields the ever-varying and potentially large client-side requests from the camera side, so recording can be performed in a steady and reliable environment.

Even if the viewing is done through the internet and bandwidth becomes an issue, the effects are temporary and nothing is lost. It is more important that the camera-to-recorder bandwidth is guaranteed, as the original video footage is best when completely stored in the recorder.

Continuous recording versus event-based recording

Continuous recording or event-based recording modes like motion or audio detection will affect the overall bandwidth generated by a video surveillance system. In event-based recording, the camera will only record video at the detection of certain events. This could be the movement of an object in an area marked as a detection zone. The recorded sequence will start a few seconds prior to the event (pre-alarm) and extend until the post-alarm time has elapsed. Since the number and duration of events can only be estimated, it will be difficult to determine exact storage requirements based on the average bandwidth.

Once the system follows the best-practice guidelines introduced in the Bandwidth Bottlenecks section, continuous or motion recording should not saturate network bandwidth. It will only affect the video retention period. Using motion recording only will save storage significantly. However, motion recording is not 100% accurate. If motion detection is set too sensitive, the system receives false alarms. If motion detection sensitivity is set too low, a critical incident might not be captured.

FortiRecorder offers the option to store continuous and event recordings for a programmable initial time period and then keep only event recording sequences. This ensures a balanced solution where no events are missed while maintaining the advantage of reduced storage requirements.

Bandwidth Bottlenecks

Network infrastructure is an important aspect of designing a surveillance system. If bandwidth bottlenecks occur, streaming video will show jitter, delay, or corruption. This could result in video not getting recorded and thus losing the purpose of a surveillance system.

Local-area network (LAN)

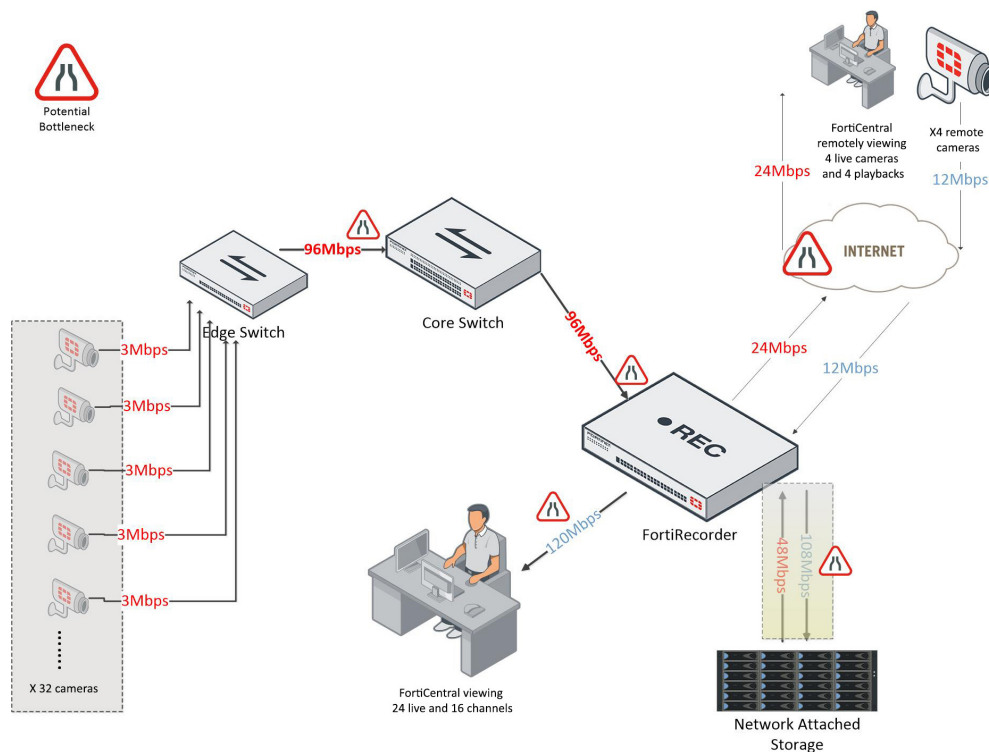
Picture 8 represents a diagram that bandwidth bottlenecks can exist in almost every connection among devices. Individual camera connections to the edge switch with a few kbps to 10 Mbit/s are unproblematic. However, the edge switch aggregates camera traffic onto its uplink port to the core switch, and a bandwidth bottleneck could occur at this point. Using the scenario shown in the illustration below, the 100 Mbit/s uplink port of the Fast Ethernet edge switch would be clearly overloaded by the $32 \times 3 \text{ Mbit/s} = 96 \text{ Mbit/s}$ traffic from the cameras. The actual throughput of Ethernet links should be limited to 50% to 70% of their nominal capacity. The solution is to upgrade the Fast Ethernet switch to one that has a Gigabit uplink port to ease the bottleneck.

The core switch function is to gather all the data traffic. Multiple edge switches flood the video feed from the cameras through the core switch onto a FortiRecorder NVR that also needs to reply to the live and playback video streaming request from the FortiCentral video management software. In the example of 24 live and 16 playback channels, this amounts to 120 Mbit/s, which is exceeding Fast Ethernet capacity. It is recommended to use a core switch with all Gigabit ports and configure in full duplex so each uplink and downlink has 1,000 Mbit/s (500 Mbit/s to 700 Mbit/s throughput) of bandwidth.

FortiRecorder comes with multiple Gigabit ports. These different ports can be used to connect to core switches on different locations of the site, internet, or corporate network for FortiCentral's local or remote video monitoring. FortiRecorder also supports network attached storage (NAS) using Network File System (NFS) or iSCSI protocol. In this scenario, FortiRecorder will aggregate 108 Mbit/s of data (96 Mbit/s from LAN and 12 Mbit/s from the internet) and feed to NAS via the core switch. In the other direction, a total of 48 Mbit/s is pulled from NAS by a local and a remote FortiCentral client.

Wide-area network (WAN) or internet

Bandwidth in the LAN can be calculated and upgraded. Internet traffic, however, is less predictable and generally exposed to higher latency and limited bandwidth.



Picture 8: Bandwidth bottlenecks.

In the same illustration shown in Picture 8, remote FortiCentral software is requesting four live cameras and four playback streams simultaneously. With 3 Mbit/s for each video, it is requesting 24 Mbit/s of data from FortiRecorder over the internet. There are also four cameras that feed a total of 12 Mbit/s of video back to FortiRecorder over the internet.

In general, we do not recommend recording from a camera over the internet due to the need for consistent bandwidth to the recorder. However, remote viewing is acceptable since the operator can simply go back a second or two if something was missed, and most importantly, recording is intact and stored in the recorder. We always recommend placing the recorder next to the camera using the LAN.

NAS in this illustration is to demonstrate the ability of FortiRecorder to redirect the video data to offload its storage. More commonly, the video traffic will be stored in the recorder.

This illustration is intended to provide a system designer with the concept of where the choke points might appear in a surveillance network.

Surveillance System Network Best Practices

Network infrastructure varies from site to site. It can be simple like a home network or complex like large corporate networks spanning across continents. Some guiding principles:

- Camera video streams are isolated and protected from existing data networks.
- Camera video streams are protected because they cannot be reached from outside of their network.
- It is easy to determine bandwidth requirements.
- There is better quality of service, since bandwidth is known and manageable.

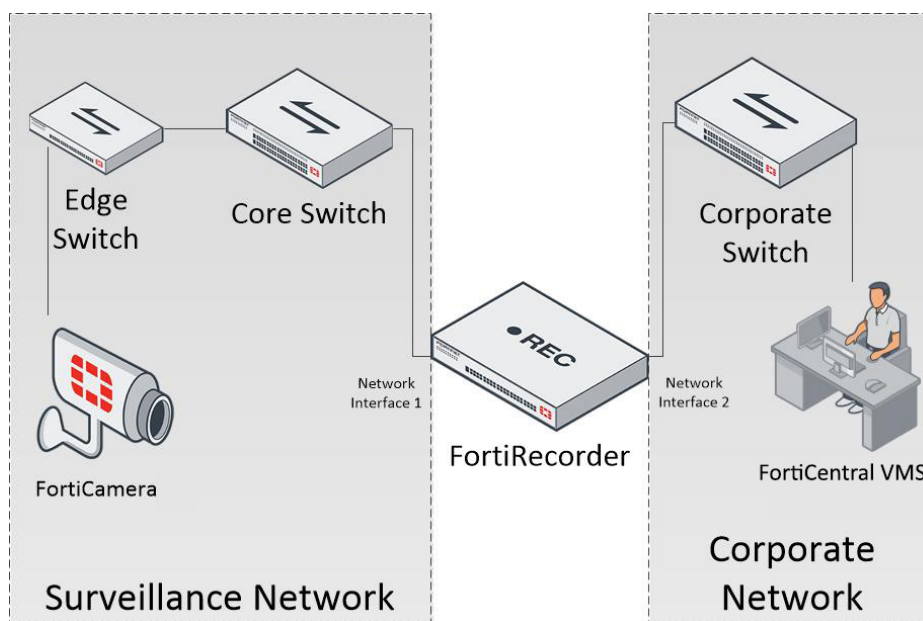
Independent switch

A dedicated switch for surveillance cameras is highly recommended. Surveillance cameras constantly generate a lot of traffic. A well-designed network infrastructure should separate surveillance traffic from existing data network traffic. This practice will make it possible to calculate the uplink bandwidth requirement and the bandwidth aggregated in your core switch.

In Picture 9, the surveillance network on the left is separated from the corporate network on the right by FortiRecorder. All FortiRecorder models have two or more Gigabit network interfaces, except the VM version, which requires the administrator to install at least two Gigabit network interface cards on the server. One interface is connected to the surveillance network, and the second interface is connected to the corporate network for the FortiCentral application to access live or playback videos.

If a dedicated switch is not practical in the existing network infrastructure, consider setting up virtual local-area networks (VLANs) in the switch. A VLAN can separate the surveillance camera traffic from the data traffic. It is like splitting one physical switch into two virtual switches.

From a network security standpoint, the separation of the surveillance network does also segment the area where third-party cameras are active. This minimizes the attack surface for exploiting weaknesses in their firmware. In this setting, FortiRecorder does not route any traffic between its ports. It only acts as a proxy for all camera-related services.



Picture 9: Separate the surveillance network from the corporate network.

Switch selection

- Fast Ethernet 100 Mbit/s ports are sufficient for camera use. Their 50-70 Mbit/s effective capacity can carry the <10 Mbit/s video streaming bandwidth.
- Ensure the switch fabric capacity is sufficient. Most switches today have full backplane capacity.
- The uplink port must support the sum of incoming traffic from the cameras. Gigabit Ethernet with 1,000 Mbit/s is generally sufficient for switches up to 48 ports.

Configurations on FortiRecorder That Affect Bandwidth

Video profile setting

FortiRecorder configuration is profile-driven. Video profiles and camera profiles define the resolution, frame rate, and bitrate mode as well as recording type, compression, and more. This is convenient to apply settings changes to groups of similar cameras.

Resolution can be defined as an absolute value like 2 MP or as a relative value where Extra High, High, Medium, and Low each match a different resolution for a different model of camera.

Camera Video Profile

Name:	<input type="text" value="high-resolution"/>
Codec	Default ▼
Resolution:	High ▼
Frames per second:	30
Bitrate mode:	Fixed ▼
Bitrate:	1000
Audio	<input checked="" type="checkbox"/>

Picture 10: Video profile—bitrate mode—fixed.

When a fixed bitrate mode is selected, an administrator can define the bitrate in Kbit/s, as previously explained as CBR.

Camera Video Profile

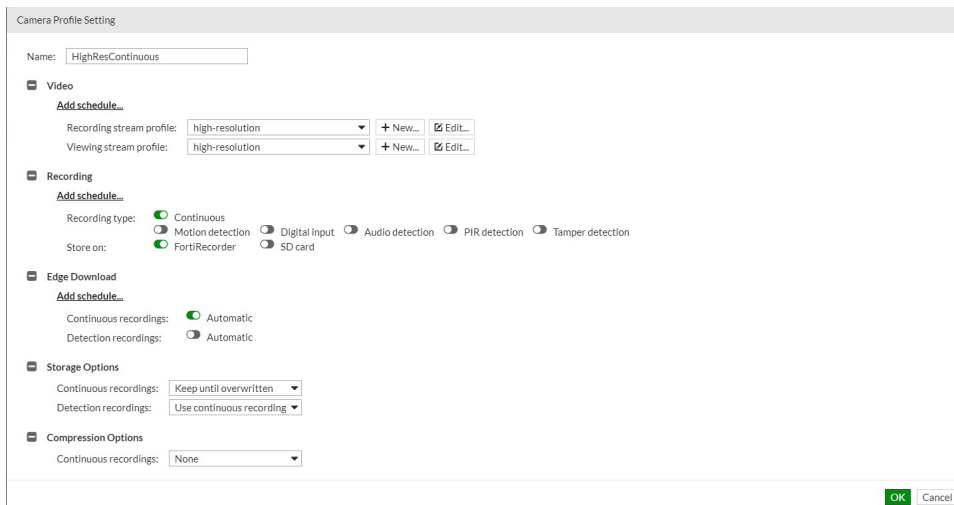
Name:	<input type="text" value="high-resolution"/>
Codec	Default ▼
Resolution:	High ▼
Frames per second:	30
Bitrate mode:	Variable ▼
Quality:	Good ▼
Audio	<input type="checkbox"/> Low <input type="checkbox"/> Standard <input checked="" type="checkbox"/> Good <input type="checkbox"/> High <input type="checkbox"/> Extra high

Picture 11: Video profile—bitrate mode.

VBR mode allows camera bandwidth to fluctuate as needed. Hence, it is difficult to predict the bandwidth usage. Once a variable is selected, an administrator can further define the quality of the image as explained in detail with illustration in the [Image Quality Level](#) section.

Camera profile setting

The camera profile allows the administrator to assign video profiles to viewing and recording streams, as well as define recording type, storage option, and compression option.




Picture 12: Camera profile.

An administrator can select a previously defined video profile for recording stream and viewing stream. The recording stream is used for recording, but along with the viewing stream can be viewed selectively in FortiCentral as well. Mobile access and the GUI only show the viewing stream in live.


Configuring the recording stream with high resolution and the viewing stream with low resolution will allow FortiCentral to have selectable dual-stream capability. It can save live streaming bandwidth, while keeping the recording footage in high resolution. More details can be found in [Configurations in FortiCentral That Affect Bandwidth](#).

Recording type allows you to choose continuous recording and/or event-based recording like motion detection, digital input, audio detection, and passive infrared (PIR) detection. It can save storage space compared to using continuous recording.



Picture 13: Camera profile—recording type.

Storage options are used by the administrator to move or delete recordings after a specified period of time to free up storage space. The separate control of continuous and detection recording allows a multilevel storage management that gradually discards information according to their relevance.



Picture 14: Camera profile—storage options—detection recording keeping.

For example, continuous recordings may be discarded after one month, only keeping detection recordings of motion or other events. Note that these detection recordings might have been generated automatically of detection events from the continuous recordings.

Compression options are a further step in tiered storage management. It allows continuous recordings to be recompressed after a specified period of time by keeping only I-frames. Applying this setting can lower the storage by approximately 50%. The resulting frame rate is lower (typically 1 FPS), but generally enough to investigate incidents.

Compression Options

Continuous recordings: After

Picture 15: Camera profile—compression options.

For a particular camera, the blue i button provides a convenient summary of the profile detail:

Camera

Enabled:

Name:

Location:

Vendor: [Camera detail](#)

Model:

Address mode:

Address: Port:

Transport type: Port:

Profile: [+ New...](#) [Edit...](#) [i](#)

- [+ Network](#)
- [+ Video](#)
- [+ Audio](#)
- [+ Privacy Mask](#)
- [+ Detection](#)
- [+ Scheduled Setting](#)

[Apply](#)

Picture 16: Camera—recording and viewing resolution configured.


Digital noise reduction (DNR) setting

DNR explains how a low-light condition will cause the flake in the video to be considered as a scene change to the H.264 algorithm and drive up the bandwidth consumption. The lower the DNR value selected in FortiRecorder, the lower digital noise reduction is applied to the video.



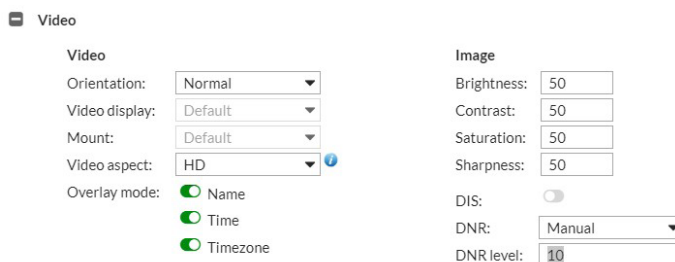
Picture 17: Camera—video—DNR set to level 1.

When the DNR level is set to 1, which means the camera will apply the lowest digital noise reduction to the video, the bandwidth consumption is averaged at 6.4 Mbit/s. It is quite high for a 1.3 MP camera with 30 FPS configuration. This is because the flaky scene is considered a change to H.264 compression.



Picture 18: Network usage average 6.42 Mbit/s.

When increasing the DNS level to 10—which applies the most digital noise reduction to the scene—the same camera with the same setting (1.3 MP and 30 FPS) only requires 0.2 Mbit/s average bandwidth. The flaky noise is suppressed and hence the H.264 algorithm does not consider it a scene change.



Picture 19: Camera—video—DNR set to level 10.

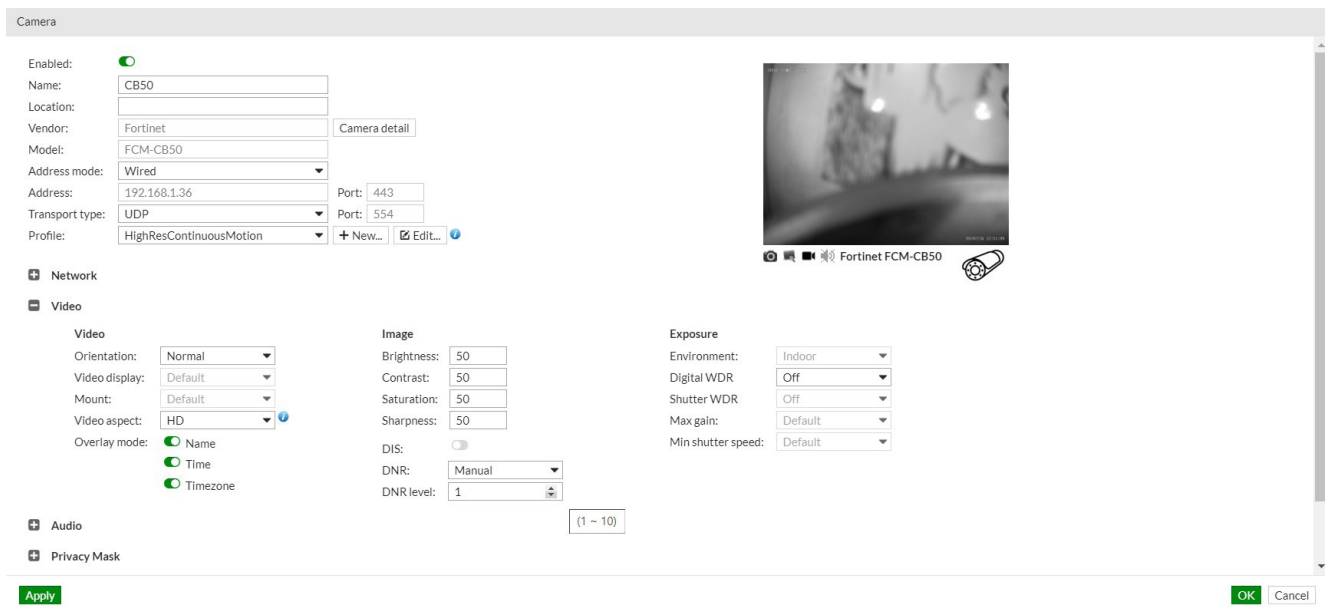


Picture 20: Network usage down to 0.2 Mbit/s average.

DNR schedule setting

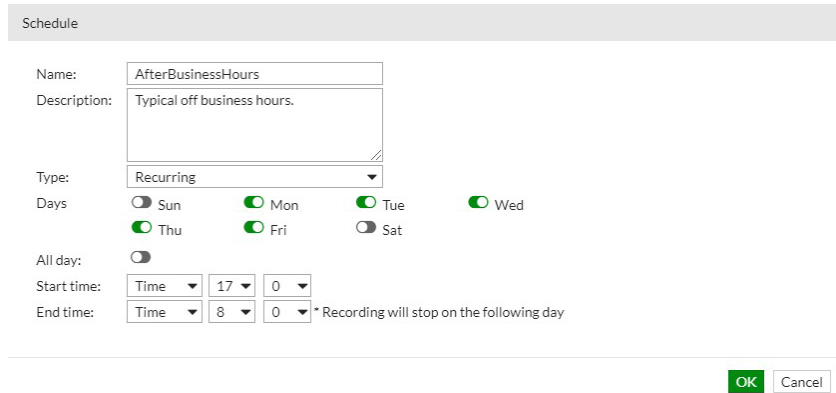
An administrator might find that setting DNR too high provides a good image in the dark, but a blurry image in daytime. On the other hand, setting DNR too low makes nighttime video flaky and overwhelms bandwidth. FortiRecorder provides schedule setting, which allows different settings to take effect depending on the predefined schedule.

The snapshot below shows the default video/audio setting with DNR level 1.



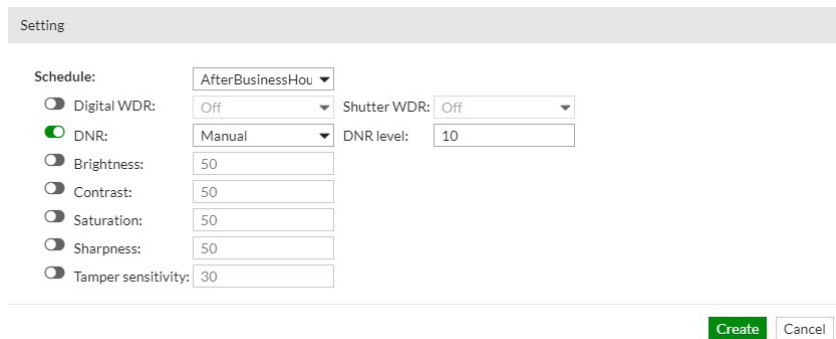
Picture 21: Camera—configuration—video/audio setting.

The schedule setting for a camera allows the administrator to define a time frame. Picture 22 shows a camera configured after business hours, from 17:00 to 08:00 the next day.



Picture 22: Camera—schedule setting.

An administrator can add multiple schedules with different DNR levels for a camera. Brightness, contrast, saturation, and sharpness can also be configured. Picture 23 shows the camera is set so that the DBR level will be brought up to level 10 at 5 p.m. every day, and it will be back to default 1 at 08:00 the next day.

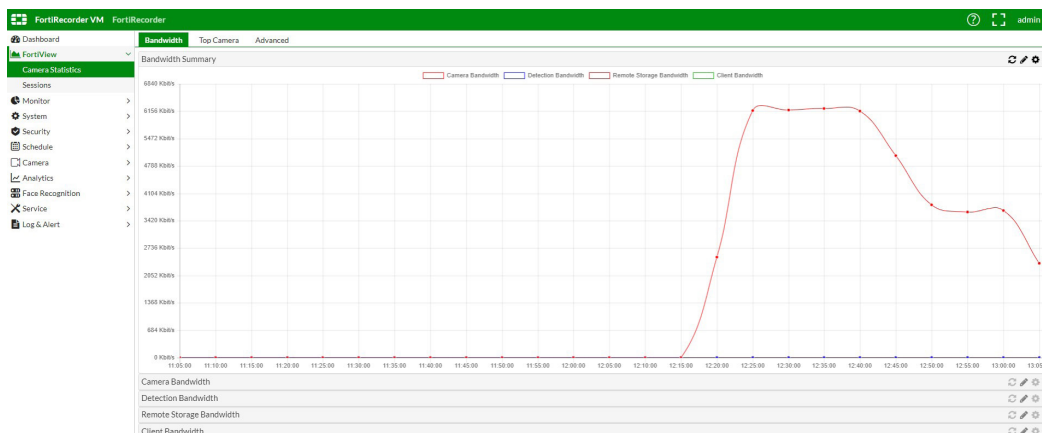


Picture 23: Camera—schedule setting—DNR setting by schedule.

FortiRecorder Statistic

FortiRecorder has a comprehensive statistical chart that will help the administrator understand recorder system loading, individual or all camera bandwidth consumptions, amount of camera motion event, and more. Monitoring this statistic periodically will help the administrator understand the performance of the surveillance system.

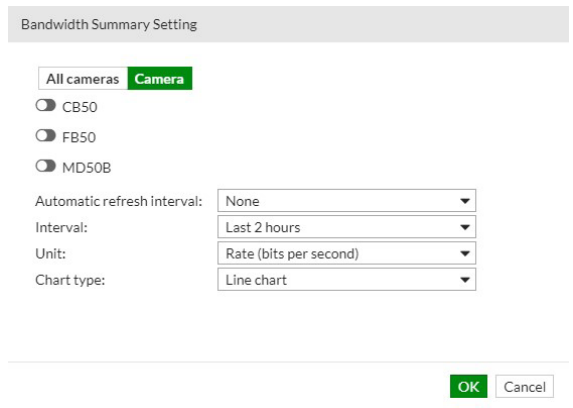
The camera bandwidth statistic allows the administrator to see all or individual camera bandwidth usage.



Picture 24: Total camera bandwidth consumption statistic.


Picture 24 gives insight into overall bandwidth coming from the cameras. The administrator can verify if it is within system design parameters. Also, fluctuations during the day or at certain events can help address problems like noise at night or bad weather and show the impact of event-driven recording.

To drill down into specific camera bandwidth consumption, the graph can be limited to individual or groups of cameras. The administrator can further select the interval to monitor, select a longer period, and look for any spike.




Picture 25: Camera bandwidth setting.

Picture 26 shows a selected camera's bandwidth information. Under normal circumstances, the graph should be flat, which means the bandwidth consumption is consistent.



Picture 26: Individual camera bandwidth consumption statistic.

Picture 27 is the client bandwidth statistic. It shows the bandwidth of all cameras or a selected camera being streamed to FortiCentral software, mobile client app, and web user interface monitoring.




Picture 27: Client bandwidth statistic.

Configurations in FortiCentral That Affect Bandwidth

FortiCentral is the video management software that monitors surveillance videos from multiple FortiRecorder NVRs and FortiCamera cameras. Changes made in the FortiCentral application will not affect the recorded video footage. It only changes the streaming bandwidth from FortiRecorder to FortiCentral. This function can automatically select higher or lower resolution to save bandwidth while also minimizing reduction in visual quality. It has a recording stream and viewing stream profile. Once the administrator differentiates two profiles, FortiCentral will have two streams available for viewing. The availability of a second video stream depends on the camera configuration in FortiRecorder. Currently, the following FortiCamera models support dual streaming: CB50, FB50, FD40, FD50, MB40, MD50B, and SD20B.

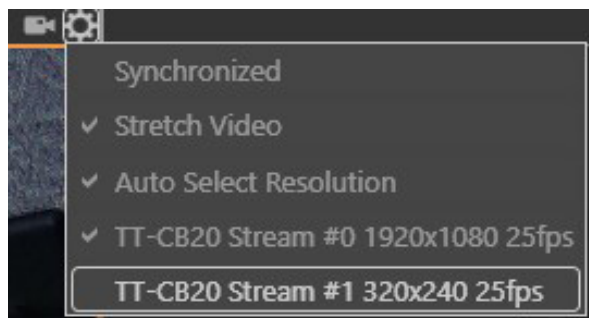
In FortiCentral, there is only one stream available if the recorder has the same recording and viewing stream.



Picture 28: One resolution available in FortiCentral.

Video profiles are applied to both viewing and recording streams. Note that the same video profile can result in different resolutions or frame rates for a viewing or recording stream. This is dependent on hardware capabilities and can be viewed from the blue button on the camera configuration page.

In FortiCentral, video stream information is available from the pane settings menu when video from a camera is live streamed in that pane. Two streams are available in FortiCentral.



Picture 29: Two resolutions available in FortiCentral.


A check mark before the stream information indicates that this stream is currently viewed. To switch to another stream (if available), simply click on the line with the desired stream description. Note that this manual selection is only available in live viewing, and the auto select resolution option must be unchecked.

If the operator enables the auto select resolution, FortiCentral automatically chooses the best matching resolution for the video in this pane. If the operator drags the video to a different sized pane or double-clicks the camera to a full-sized pane, the best/closest resolution will be used.



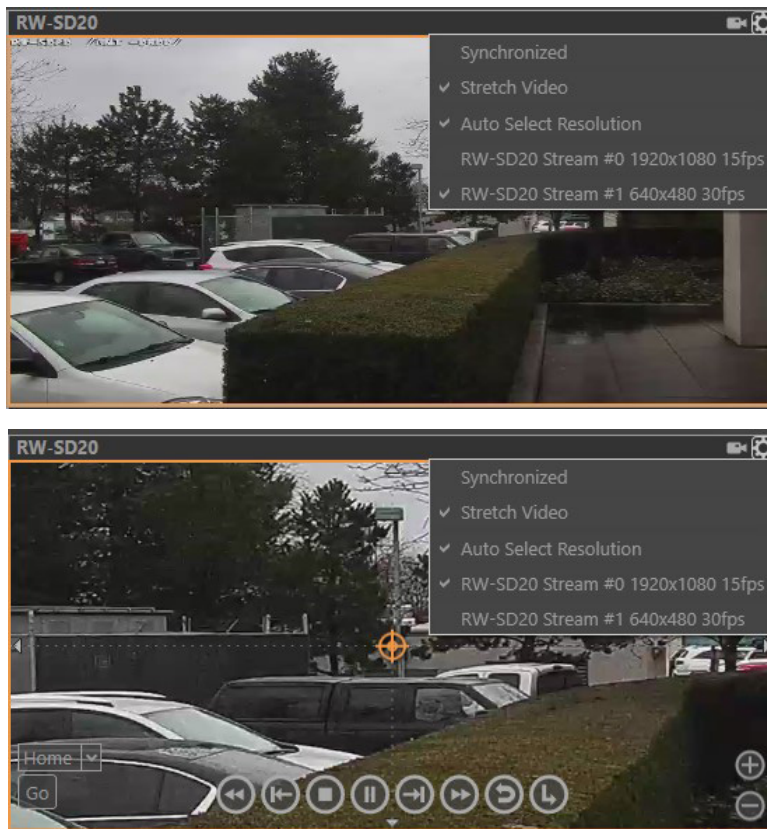
Picture 30: Auto select higher resolution at work.

If the operator drags the camera to a smaller pane, the lower resolution will be streamed to save bandwidth.



Picture 31: Auto select lower resolution at work.

If the operator digitally zooms in the video, the best/closest resolution will also be selected. In the example below, the camera is zoomed out in low resolution, RW-SD20 Stream #1 640x480 30 FPS. Zooming in will automatically switch to a higher resolution, RW-SD20 Stream #0 1920x1080 15 FPS.



Picture 32: Auto select lower resolution at work when zoomed out.

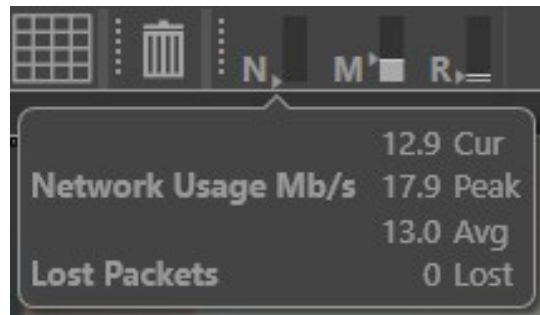
Dual streams can efficiently save streaming bandwidth when viewing multiple cameras simultaneously. When needed, zooming in will bring higher resolution for the operator to see the detail. There is also a sequencer in FortiCentral. This function allows video in the pane to switch based on a preset interval. It can be a bandwidth-saving technique to live stream multiple cameras using sequencer. Sequencer does not buffer video data before streaming, hence saving bandwidth on the network.

Analyzing Bottlenecks with FortiCentral Performance Statistic (Gauge)

Network bandwidth limitations can result in a choppy, interrupted video display. For a user, this can appear very similar to the effects of insufficient local decoding performance in a FortiCentral installation, which must decode and render the video streams from the NVR. Displaying multiple video streams on the client is very CPU-intensive.

FortiCentral helps to easily distinguish between the two situations with the integrated performance metering tools for Network, Rendering, and Memory on the toolbar.

The network usage statistic shows the Current (Cur), Peak, and Average (Avg) incoming network bitrate. It also shows the detection of gaps in the video streams as lost packets. To measure the bandwidth of a particular camera or set of cameras, switch only those on the pane and observe the resulting bandwidth.




Picture 33: Network statistic.

It is ideal to keep the network usage below 50% to 70% of your network adapter bandwidth. If this meter turns yellow or red, it indicates that packets have been dropped in the recorder because the network could not deliver them fast enough. In this case, the problem is in the network between the recorder and the client. Here are some ways to mitigate:

- Make sure the network is designed and configured for the required bitrates.
- For remote access over the internet, consider the local internet service provider’s connection restrictions and internet capacity.
- If a virtual private network (VPN) is used, ensure it can support the required bandwidth.

Rendering performance indicates how many megapixels are decoded per second. For example, a 1 MP image at 30 FPS requires 30 MPixel/s. If more video is displayed in FortiCentral, the performance limit of the PC will eventually be reached. At this point the video decoder has to drop frames to keep up with the incoming video stream. This drop rate is displayed at the tooltip, and if it occurs, the bar turns yellow and red in more severe situations. This indicates that the CPU/graphics performance of the PC where FortiCentral is running has reached its limit.



Picture 34: Rendering statistic.

If a FortiCentral installation frequently shows yellow or red bars on the rendering meter, there are several options:

- Reduce the video resolution or frame rate to make decoding less costly.
- Enable auto resolution switching between dual camera streams to work with the smallest resolution appropriate for the given display area on the screen (refer to [Configurations in FortiCentral That Affect Bandwidth](#)).
- Show fewer video streams at the same time. Sequencers can help maintain an overview while only showing a selection of cameras at a time.
- Add another computer to share the load from an existing computer.
- Upgrade hardware capabilities with a stronger CPU. Video decoding scales well among multiple cores.
- We recommend a FortiCentral computer be dedicated for surveillance monitoring use only. Verify with Windows Task Manager (Ctrl+Shift+Esc) if other software is utilizing CPU cycles. Disable Windows services or uninstall features that are not needed.

FortiCamera Calculator

Fortinet offers an easy-to-use tool for FortiCamera and FortiRecorder deployment planning—FortiCamera Calculator.

All the qualitative factors for bandwidth discussed earlier do not answer the question of how much bandwidth is actually generated by a camera. For an estimate, refer to the bitrate table below. For example, a 1 MP video with 6 FPS may generate 1 Mbit/s for a high-quality image.

Bitrates are estimates providing a high-quality image under most conditions. If the scene is less complex (indoors with little detail and motion) or the camera has very little noise (daylight, good DNR), the bitrate can be lowered further. Generally, do not use less than half of the listed values.

If video compression is set to lower quality or capped at a defined maximum bandwidth, the bitrate can be significantly lowered at the cost of poorer image quality.

Bitrate Table (H.264 estimate) Mbps					
High quality image (x 0.7 = OK quality)					
Frames/s	1	6	10	15	30
CIF (352 x 240)	0.16	0.2	0.24	0.3	0.4
D1 .4M (720 x 576)	0.4	0.5	0.6	0.75	1
720p 1M (1280 x 720)	0.8	1	1.2	1.5	2
SXGA 1.3M (1280 x 1024)	1	1.25	1.5	1.9	2.5
HD 2M (1920 x 1080)	1.6	2	2.4	3	4
3M	2	2.5	3	3.75	5
5M	3.2	4	4.8	6	8

Table 2: Bitrate table.

We can take this estimate to calculate capacity using a spreadsheet like the Capacity Calculator available from Fortinet. It is a very useful tool to get an idea of bitrate, HD capacity, number of cameras, percentage of usage, and video retention time.

As an example, in Table 3, white-colored cells allow the input of system parameters, while the result shows what the parameter value will be if all other parameters are as defined by the user.

Here we specify 3 TB hard drive capacity and five cameras with 2 Mbit/s configured in continuous recording for 10 days. The orange result shows that camera bitrate could go up to 6.1 Mbit/s and it would still record 10 days on the 3 TB drive. Alternatively, 1 TB HD capacity would be enough, or we would achieve 30 days of retention time.

Capacity Calculator					
Input	Bitrate (Mbps)	HD Capacity (TB)	Cams (#)	Usage (%)	Time (days)
	2	3	5	100	10.0
Result	Resolve each for all other inputs as specified				
	6.1	1.0	15.3	305	30.5

Table 3: Capacity Calculator.

FortiCamera Calculator should be used as a reference only. No calculator can predict the exact bandwidth a camera or a surveillance system will generate. For more accurate bandwidth estimation, it is recommended to install the camera on-site and refer to the [FortiRecorder Statistic](#) section for individual camera bandwidth consumption.

For more complex projects, a Multi-camera Calculator is available for sizing mixed installations with different resolutions and usage patterns.

Summary

Bandwidth availability is a primary concern for any network, and video surveillance systems are no exception. Many factors affect security camera bandwidth, such as video compression, image quality, complexity of the video scene, video resolution, frame rate, and number of cameras and clients. The best way to predict your real bandwidth requirements is to pilot a small number of IP surveillance cameras. Run them for a short period of time, review the video outcome, and make any necessary adjustments for the key factors mentioned in this white paper. This is the most practical way to help you plan your infrastructure and needed equipment accordingly.

As technologies keep evolving, the performance of video surveillance systems will continue to improve, with better image quality and lower bandwidth. Fortinet is dedicated to developing the most cutting-edge technologies to protect businesses from threats and criminal activity. The Fortinet Video Surveillance solution combines FortiCamera IP cameras, FortiRecorder NVRs, and artificial intelligence-enabled FortiCentral VMS to empower organizations across various industries to build a safer and more efficient environment while reducing operational costs and bandwidth consumption.

