

WHITE PAPER

Decreasing Networking and Storage Costs of Your IP Video Surveillance System



Genetec

Innovative Solutions

Executive Summary

With the increased prevalence of IP-based video surveillance systems on the market, and the growing adoption of higher resolution HD and megapixel cameras, organizations and system integrators must take into account how implementing these systems can impact their network resources. Without realistic system design considerations, organizations can risk significant network and storage cost overruns while also compromising the reliability of their network to support applications that are critical to their business operations.

By implementing an advanced video management system (VMS), an organization can effectively manage video streams on their network using built-in camera and software functionalities to optimize network resources and bandwidth consumption. With such optimizations, a VMS will also help to decrease networking and storage costs over the lifetime of a video surveillance system.

1

Challenges

For organizations choosing to implement or expand an IP video surveillance system, the ability to efficiently manage video streams and storage is crucial to ensure the best use of the network and reduce costs associated with deploying and operating the system.

While organizations continue to benefit from greater network speeds and capacity, the use of IP-based video systems can generate a significant increase in the amount of data traveling on their network as a result of:

- ▶ deploying high-definition and megapixel cameras,
- ▶ additional cameras to address a need for coverage across larger areas,
- ▶ increases to the number of users accessing video,
- ▶ recording and maintaining redundant video recordings,
- ▶ transferring video from one site to another to maintain long-term or centralized recordings.

When planning and designing an IP video surveillance system, an organization must take into account the unique aspects of its security environment and its business operations in order to ensure the reliable transmission of video and avoid overloading available network resources. For example, certain deployments will require greater flexibility to manage video streams and bandwidth due to their complex nature, further driving the need for advanced video management capabilities.

These scenarios can include:

- ▶ Distributed sites requiring operators to connect to remote cameras,
- ▶ Cameras connected to networks with limited bandwidth, such as DSL, wireless, or cellular,
- ▶ Sharing bandwidth with other operation-critical applications because video is not the top priority for the business.

It is equally important for organizations to realize that optimizing the use of network resources does not necessarily require large capital investments but is more a matter of putting the right solutions in place. With bandwidth and storage representing important ongoing costs of operating an IP system, organizations can significantly reduce the Total Cost of Ownership (TCO) of their video surveillance system by investing in solutions that allow them to optimize their use of bandwidth and storage based on the requirements of their application.

This white paper will focus on those unique and powerful capabilities that one should look for in a VMS in order to optimize the use of network resources and reduce the costs associated with operating an IP-based surveillance system.

2

Optimizing Network Resource Utilization

VMS applications allow an organization to manage its security infrastructure, including video cameras, encoders, and recording servers, within the unique context of the organization's deployment. The effectiveness of the VMS will depend on its ability to handle the demands of the operating environment, whether those demands include deploying a system in sites with limited bandwidth, monitoring cameras across distributed locations, or ensuring that multiple operators can access necessary video streams in the case of an incident, regardless of the number of concurrent requests.

Although system administrators will intuitively manage video quality settings and define recording settings and schedules, addressing the needs of a specific security department can require manual intervention and adjustment. While most VMS applications support these features, some VMS applications also support powerful functionalities and technologies that serve to further reduce the total cost of operating an IP video system.

In fact, organizations can deploy a surveillance system that operates with greater efficiency on their network by choosing a VMS application that supports the following capabilities:

- ▶ end-to-end multicast transmission,
- ▶ stream redirection and multicast-to-unicast conversion,
- ▶ multi-streaming,
- ▶ video caching,
- ▶ archive transfer.

By leveraging some or all of these capabilities, organizations can significantly reduce the number of servers required to manage and store video, reduce their network bandwidth requirements, and reliably scale their system while minimizing their investment in new infrastructure.





A. Video Stream Transmission: Unicast and Multicast

In IP video surveillance, unicast and multicast are the two most commonly used methods to transmit video from cameras to client workstations. While all VMS platforms can configure unicast, only a few also offer multicast transmission, and, among these, even fewer support end-to-end multicast that provides communication from the edge device (IP cameras and encoders) to the workstation. Though many VMS platforms may claim multicast support, the majority will only provide limited support for multicast transmission between the recording server and the client station, and require multicast to be set for all cameras on the server, or even implemented system wide. It is important for organizations to consider that certain VMS provide far greater flexibility with regards to transmission, in order to implement the best design for their application. This includes the ability to set up cameras per select network branch or per viewer and the ability to automatically detect the ideal transmission method for different segments of the network, thereby allowing organizations to optimize the performance of their video surveillance system and minimize the network resources that are required.





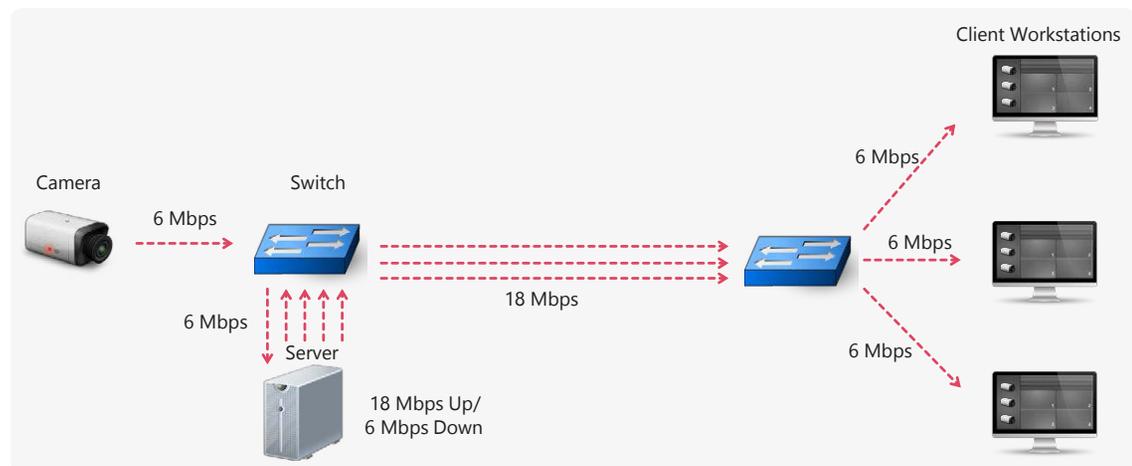
i. Unicast Overview

Unicast is usually done in TCP or UDP and requires a direct connection between the source and the destination. Unicast only works if the source has the capability to accept concurrent connections when multiple destinations want to view or record the same video at the same time. In IP video surveillance, unicast involves a camera streaming as many copies of the video feed as are requested by the destinations, so a 6 Mbps video stream that is requested by three operators will produce a transmission of 18 Mbps of data across multiple network segments (6Mbps per stream x 3 requests = a total of 18 Mbps). This method of transmission usually results in higher network traffic, especially in the case of large systems with hundreds of operators requesting the same video stream. Moreover, it is not considered best practice for corporate LANs because it can cause a rapid increase in bandwidth requirements for viewing and recording video streams. In the event of an emergency in which several operators require access to the same camera streams, the request for concurrent streams can overtax the processing capabilities of a camera or the bandwidth constraints of the network. This excess demand can result in a reduction in video quality or, even worse, the inability for operators to access and view the streams that they require to respond to an incident.

Unicast Transmission ►

In **Unicast**, a camera will stream as many copies of the video feed as are requested by the destinations. For example, a 6 Mbps video stream that is requested by three operators will produce a transmission of 18 Mbps of data across multiple network segments.

(6 Mbps per stream x 3 requests = Total of 18 Mbps)



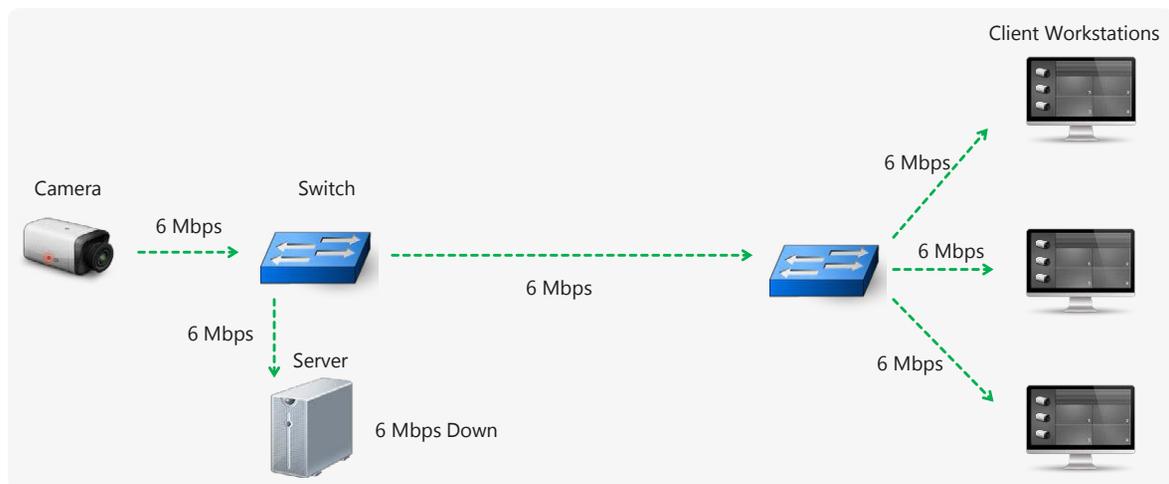


ii. Multicast Overview

In multicast transmission, because connecting to a video stream is accomplished by simply joining a multicast group, there is no direct connection between the source and the destination. Instead, the source simply sends a single copy of the video stream to its designated multicast address, and then the destination connects over the network to the multicast address of that stream. This means that all destinations share the same video stream. Although multiple operators may request the same video stream, an IP camera in multicast transmission only needs to send one stream over the network without crossing any branch of the network more than once, thereby realizing a significant gain in network efficiency as compared to a unicast deployment. Therefore, organizations using IP video surveillance, especially those with large-scale deployments with remote locations and cameras or those that could experience a rapid increase in the number of clients, can greatly reduce bandwidth requirements by using multicast transmission.

Multicast Transmission ►

In **Multicast**, a camera will stream a single copy of the video stream to its designated multicast address, and all destinations will share the same video stream. For example, a 6 Mbps video stream that is requested by three operators will produce a transmission of 6Mbps of data across multiple network segments.





iii. Implementing Multicast Transmission

While multicast transmission is an effective way to improve bandwidth management, it must be supported by three components (VMS, camera, network):

- ▶ a VMS that supports multicast management,
- ▶ multicast-enabled routers and switches.

The VMS is essential to properly manage multicast traffic over a network. In order to effectively manage multicast traffic, the VMS must be able to:

- ▶ provide the viewing client application and the recording server with the ability to join the multicast group over the network,
- ▶ manage both unicast and multicast transmission methods within the same system,
- ▶ offer a proxy service to transform unicast traffic to multicast and vice versa,
- ▶ support auto-detection of the network capabilities.

Because networks are complex and IP sources do not all transmit identically, the VMS in an IP-based security installation also needs to support multiple transmission methods in order to be a viable solution. Not all segments of a network necessarily support multicast. Such instances may include outdoor cameras that transmit data over a wireless network that does not support multicast or viewing stations that must connect to the video

surveillance system via the Internet. The VMS is responsible for managing the transmission requirements of these different components by:

- ▶ recording from multicast- and unicast-enabled cameras,
- ▶ transforming unicast cameras into multicast cameras to ensure the efficient use of bandwidth by local users,
- ▶ providing a unicast stream to remote clients.

The other two components required for building a multicast-capable network are IP video sources and switches that support IGMP snooping or IGMP querying (also known as IGMP routing). These switches are responsible for handling the multicast traffic on the network. While a switch that supports IGMP snooping ensures that the multicast traffic is only sent to destinations that have requested the stream and is filtered from the rest, the IGMP querier ensures that IGMP snooping is working correctly by checking that the stream is being transmitted properly through the multicast connections.



B. Advantages of Multi-streaming

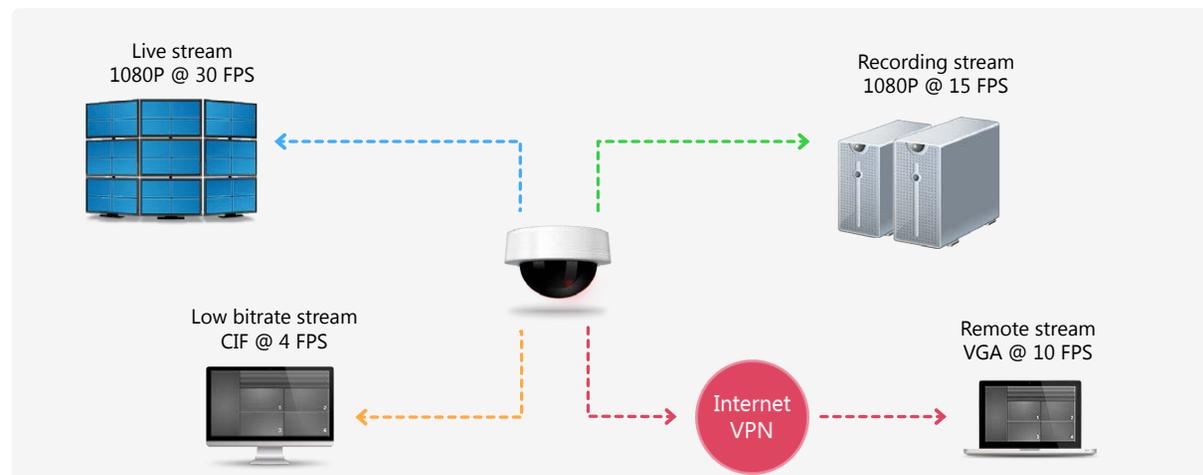
Multi-streaming refers to the capability of an IP camera to transmit multiple video streams from the same camera device with different video quality settings. Two components are essential for multi-streaming support in an IP video surveillance system:

- ▶ an IP camera or encoder capable of producing multiple streams.
- ▶ a VMS that is able to manage those streams,

Multi-streaming allows the system administrator to easily assign different purposes to video streams, including live viewing, remote viewing (with reduced bandwidth), recording, and long term recording. The following are examples of how different multi-stream configurations can help reduce network load, and ensure appropriate video feeds are streamed to all destinations.

Multi-streaming Configuration ▶

By leveraging a camera's ability to generate multiple streams, a VMS can designate stream qualities that will be used from the same camera for different purposes. For example, a full quality stream can be configured for live and recorded video, while a lower quality stream can be set for remote viewing, when bandwidth may be limited.





► **High-quality live streams**

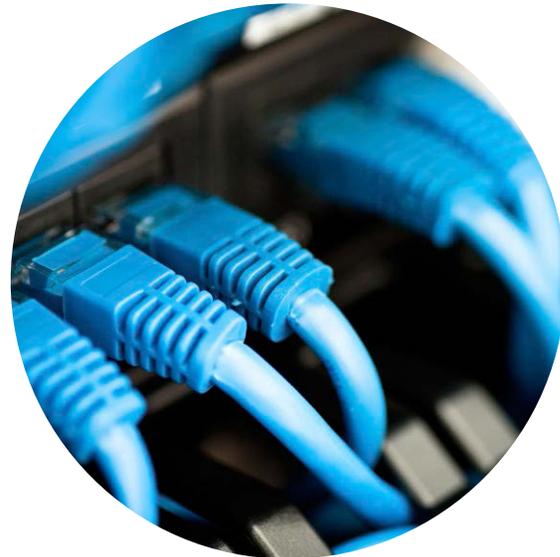
A typical configuration that helps to save on storage involves setting up one stream for live viewing at 1080P, 30 fps and a second stream for recording at 4CIF, 10fps. With this setup, cameras can be monitored in real-time at a higher resolution while being recorded at lower video resolution, thereby reducing the cost of storage.

► **High-resolution recorded streams**

For other applications, it may be beneficial to set a live stream at a lower resolution, such as 4CIF, 30fps, to ensure greater fluidity for the operator monitoring the system while recording a secondary stream at a higher quality and reduced frame rate, such as 1080P, 5fps, to provide greater detail in the case of an investigation where it is necessary to identify a suspect. This approach also reduces the bandwidth and storage requirements of the system.

► **Lower quality remote streams**

When the system administrator does not have control over the capacity of the external network, multi-streaming can be used to set up a lower quality stream for remote users while keeping a high-resolution stream for local viewing and recording. By using a stream with low bandwidth, this configuration ensures that remote users, wherever they may be, have access to fluid video.





i. Multi-stream Archiving

In addition to being a valuable capability for optimizing bandwidth, multi-streaming is also an excellent way to reduce storage costs when archiving video for long periods. Because a VMS is able to manage multiple stream qualities from a single camera, the system administrator can configure two streams from the same source in such a way that high quality video is recorded on a short-term storage medium, while a lower quality stream is recorded for long-term retention. Considering the above example, the administrator could configure a 1080P, 30 fps stream to be recorded for 30 days and a 4CIF, 10fps stream to be recorded for 90 days. For many incidents, only the short-term video is typically reviewed. Using multi-stream archiving, high quality recordings are not kept unnecessarily for the full retention period, which helps to reduce the amount of storage required for the system.

In organizations with distributed systems, investigators who are reviewing video remotely can also benefit from multi-stream archiving. Since the bandwidth available with a connected site may be limited, access to lower quality recordings can provide greater fluidity and can speed up searches for investigators who are searching for evidence stored in the remote system.

ii. Automatic Stream Selection

Video streams are often set at resolutions that are higher than the actual size of the display of the video display tile, resulting in an increased use of bandwidth but without providing any benefits to the operator. Additionally, as the video resolution increases, the resources required by the workstation—by the CPU or RAM—to decode and display the video also increase.

By taking advantage of multi-streaming capabilities, a VMS can automatically switch between higher and lower resolution streams based on the size of the video tiles being viewed by an operator. This can greatly reduce bandwidth and workstation requirements in instances where there is no need to display a full megapixel stream, such as when an operator is monitoring multiple cameras on a screen or when some cameras are displayed in small tiles.

For example, an operator monitoring 16 cameras on the same monitor would consume 64 Mbps of bandwidth if each stream is set at 2 Mbps. However, since a 16-camera layout on a 20” monitor does not require such high resolutions, using multi-streaming to dynamically select the appropriate stream size would greatly reduce bandwidth consumption without affecting the viewing quality, while also reducing the load on the CPU.



C. Video Caching

When monitoring video, an operator will often need to replay the same sequence to review a scene. Since the video has already been transmitted once from the recording server to the client workstation, a VMS should not need to retrieve the same footage over the network each time the sequence is viewed.

To ensure greater fluidity when reviewing video and to minimize the use of bandwidth, organizations should ensure that their VMS has the ability to maintain a live and recorded video cache.

D. Archive Transfer

A VMS' ability to transfer video archives from their initial recording location to another physical storage device provides organizations with greater flexibility to select the storage hardware most suited for their requirements, whether they want to reduce the cost of storage at distributed sites, or maintain long-term recordings.

The capability allows organizations to record video locally, either on edge devices or on a NVR, and transfer archives to a central server on-demand, or on a schedule, in order to minimize the impact of the video surveillance system will have on their network. For example, when it is not possible to continuously stream video to the central system for recording (due to limited bandwidth or unstable connectivity), archive transfer can be used to schedule the video transfer during periods when network demand is low or

after regular business hours in order to avoid impacting business operations. During the transfer, the video is copied from the edge device or server at a rate determined by the capacity of the network connection.

A VMS that supports archive transferring also allows an organization to:

- ▶ limit transfers to specific recordings of interest,
- ▶ realize greater reliability from their IP video surveillance system by maintaining redundant recordings on the edge.

Using Archives Transfer, an organization can configure its video surveillance system in such a way that, instead of continuously streaming video over the WAN, only archives tagged with events, such as motion, or alarms, will be transferred from the remote site to a central server for long-term storage. This provides an effective technique to reduce the use of network resources because it decreases both the load on the WAN and the amount of video that will be stored in the central system.



Archive transfers can also be used to add a greater layer of reliability to an organization's IP video surveillance system. In the event where connection to a camera is lost, a VMS can transfer only the video that was recorded on the cameras affected by the connection failure, once connection to servers has been reestablished.

E. Motion Boost Recording

A VMS with Motion Boost Recording can also help organizations save on storage while simultaneously offering a solution that records incidents in full detail by providing the ability to enhance video recording settings based on motion events.

The majority of VMS solutions provide the ability to record video when motion occurs. However, if scene analysis fails to detect activity in front of a camera, the VMS could fail to record key events; in order to avoid such a failure, organizations may instead prefer to record video continuously. But, while continuous recording provides a steadfast solution that will record video at all times, it also requires greater use of storage, even when no activity is present.

A VMS with Motion Boosting Recording can be set up to record video continuously at a low resolution and frame rate and to automatically increase video settings if motion is detected in the scene. This provides an organization with the certainty that video will be recorded at all times and also helps to reduce the amount of storage required as high resolution streams will only be recorded when motion is detected.



F. Cloud Archiving

The need for additional storage over the lifetime of a video surveillance system can arise from an increase in the number of cameras in a system or from changes in video retention policies, prompting the need to purchase additional storage hardware.

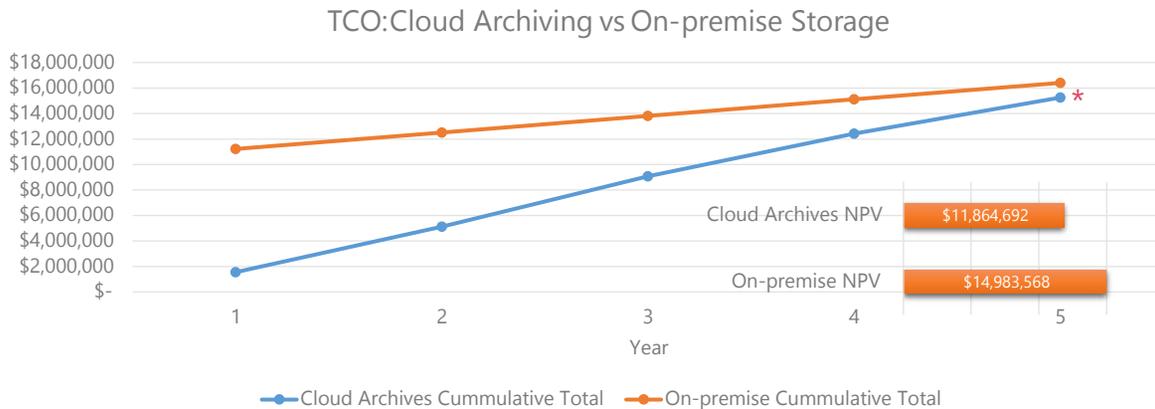
With a VMS that has the ability to archive recordings in the cloud, organizations can decrease their hardware investment, while continuing to leverage their existing on-premises infrastructure. Leveraging a hybrid architecture, video can be archived locally and transferred to the cloud as additional capacity is required, without the need to purchase additional servers. Likewise, cloud archive capacity can seamlessly be reduced, should storage requirements decline, without the risk of maintaining unutilized storage on-site.

With the ability to transfer all or only specific events in the cloud for long-term retention, organizations can better manage increased demand caused by peak utilization periods. In addition, because the recordings are kept offsite, cloud archiving reduces operational expenses required to maintain their own infrastructure, including electricity and server maintenance costs.

Total Cost of Ownership Comparison ▶

Estimated storage cost of 600 cameras with 24 Month retention. Camera settings and motion activity are constant in both scenarios (720p, 15 FPS)

* Server refresh required at Year 6



3

Get the Most from Your VMS Investment with Genetec

A VMS is essential in helping organizations tailor their IP-based video surveillance system to meet the demands of their security objectives and to provide their security team with access to their video when they need it most.

To avoid network slowdowns or system downtime, the Omnicast™ VMS solution from Genetec allows customers to effectively manage video streams and optimize the use of network resources. Deployed in some of the world's largest and most demanding video surveillance applications, Omnicast has an advanced architecture that allows customers to scale their system as needed, while also providing the capabilities they need to realize significant cost-savings.

Built with an open architecture, Omnicast provides support for thousands of IP cameras and devices from a broad ecosystem of leading edge device and server partners. As a result, a customer avoids locking their security investment into a proprietary solution and also benefits from the protection that comes from having deployed a trusted VMS for all current and future security requirements.





Get the Most from Your VMS Investment with Genetec:
genetec.com/omnicast