THREE MUST-HAVE METRICS OF VIDEO SURVEILLANCE INFRASTRUCTURE

By David Nelson-Gal & Dr. Manqing Liu

☎ (650) 263-8225 | ✉ sales@viakoo.com | viakoo.com

Abstract

In this paper, we discuss the three key performance indicators (KPI) that are necessary to properly measure the integrity of IP Video Surveillance applications and supporting infrastructure (what we call a **video network**). These measures are as follows:

- Video Path Uptime (VPU)
- Video Stream Delivery Index (VSDI)
- Video Retention Compliance (VRC)

From these three measures, we calibrate whether the infrastructure and application are operating properly, and use these measures to promptly alert the people responsible for physical security operations of problems that are affecting the IP Video Surveillance capability in meaningful ways, and also build ever better predictive analytics to determine the root cause of problems and predict potential future problems.

Overview

A typical video surveillance application consists of multiple cameras (some with multiple streams), camera network switches, recording servers, storage systems and video management software (VMS). The goal of this kind of application is to provide situational awareness enabling a small number of individuals (security officers) to monitor a broad expanse of physical plant and property, and to provide a recoverable record of events to correctly understand what happened, to facilitate recovery, to arbitrate disputes, and improve procedures. Therefore, one of the most important goals is to make sure that video streams from cameras are recorded properly on storage systems.

Many times, videos streams are not being recorded due to component failures, software errors, or improper configurations. Additionally, it is possible that some frames of video fail to reach storage due to congestion in network paths, congestion on servers (not enough CPU) or storage performance. This leads to gaps in the video potentially at critical moments. A final problem is the premature deletion of video files to make room for new video data due to system performance limitations or the malicious due to system performance limitations or the malicious deletion of video evidence. The consequence of any one of these problems is that the video data is not available when it is needed.

Until now, figuring out if the video streams are properly recorded has been a complex manual process. Existing network monitoring tools don’t accurately identify whether each video stream is working properly, recording completely, or being retained on disk for the intended period of time. Their core function is to identify more obvious situations where physical devices completely fail. And some tools bombard physical security system administrators with false alarms or large numbers of complex signals that users don't know how to interpret or prioritize. Since no one solution provides a comprehensive picture of the video network, failures happen quietly.

As a consequence, even the most sophisticated operations are forced to use human resources to manually check that each camera stream is working on some periodic schedule. This checking process requires viewing and playing back the recorded videos on each camera stream and validating the full retention cycle. Invariably, because these manual disciplines aren't in place or because human limitations cannot keep track of all the camera streams effectively enough, many operations only discover problems with missing video data after an incident failed to be recorded.

At the moment missing video content is discovered it is too late, because the data is lost and recovery is not possible. At best, it's an awkward moment to explain to stakeholders why, after investing significant time and money in a video capability, there is no video evidence available to support an investigation. At the worst, degrading video or retention requirements can have disastrous business and brand consequences for companies with customer or regulatory compliance requirements.

This paper identifies the key performance measures of video networks and illustrates methods of creating these measures for video surveillance applications that address downtime problems, leading to better operational awareness and discipline. They also create a mechanism that can alert users when there are real issues affecting security video, helping the physical security system administrator and others responsible focus only when servicing is required. They are also trendable measures, to drive operational excellence throughout the organization.

Measuring Performance

The solution is in three measures that comprehensively capture what is important: **Video Path Uptime** (VPU), **Video Stream Delivery Index** (VSDI) and **Video Retention Compliance** (VRC). These three metrics are generated for each camera stream, and then can be aggregated together to create overall measures for a single recording server, a site running a broad collection of cameras, recorders and viewing stations; a collection of sites, or an entire company.

For each metric, the process requires measuring the current state of the infrastructure at regular intervals (a sample). Each sample collects measures about the infrastructure either as an instantaneous measure (live) or as an accumulated count since the last sample was taken. Trend metrics plot these measures from one sample to the next, providing a long-term view of how system behaves over time.

Figure 1 - Example Video Network

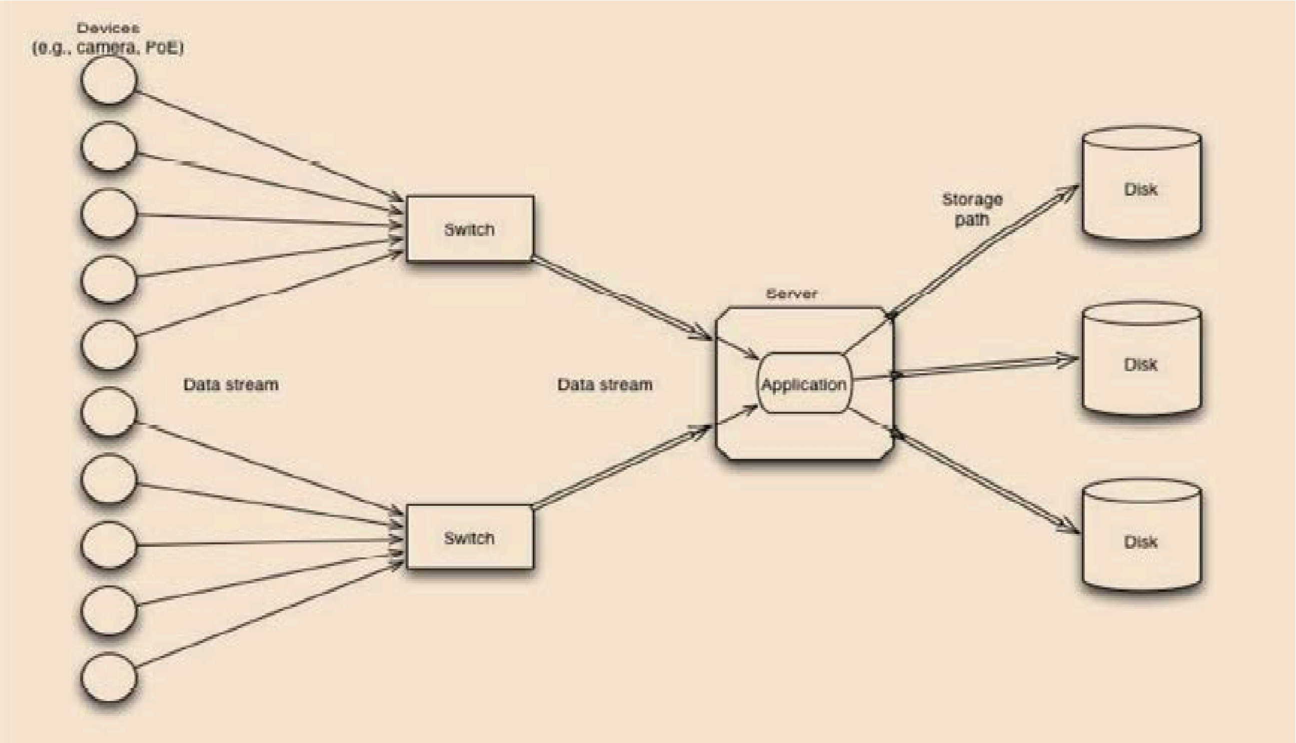

Video Path Uptime (VPU)

VPU measures the end-to-end stream path availability from camera device to storage media. VPU is defined as the end-to-end uptime percentage of the path through the infrastructure that a stream takes.

Its goal is make sure camera streams are recording as designed. At its core it is an aggregation of measures of a distributed relationship where every element of that relationship has to be working for the overall measure to be considered working correctly (an uptime state of TRUE).

If we consider the example video network in Figure 1 above, we have a collection of video streams on paths through the infrastructure from the devices (cameras labeled C1-C10) through switches to one or more servers. Within a server there are applications that process the data and store it in mounted volumes (Drives 1 through 3).

For each stream of data, we monitor the constituent components of the stream from the perspective of the service, whose end goal is to successfully receive the data stream and write it to its proper location in storage.

A device like a camera can have one or more streams of data. To understand the performance of the overall system, we have to aggregate a measure for each path.

Therefore, based on the aliveness of all components in the system, any component failures in the video stream path will cause VPU degradation.

Figure 2 - Single Data Stream Path for Device C1

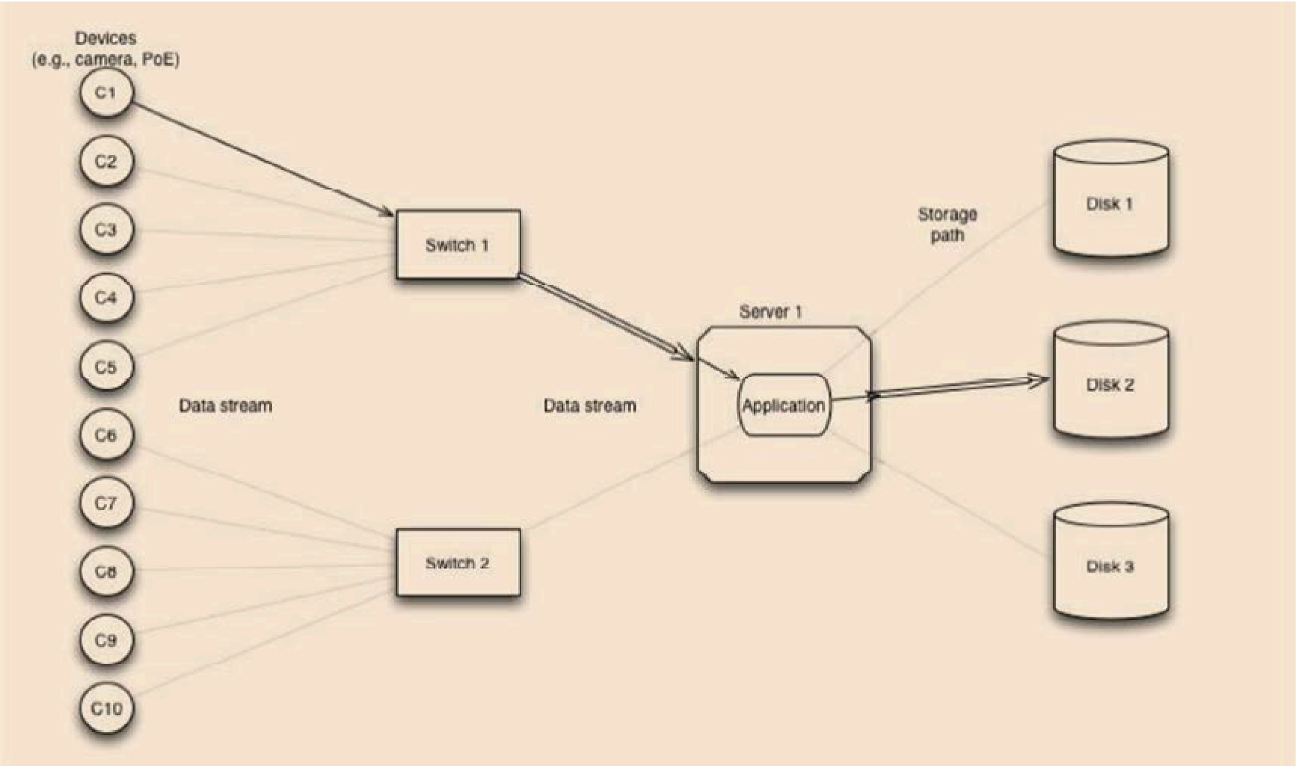

Algorithm:

For a given Recording schedule, when videos are recording properly for a single camera stream *i*, VPU$_i$ = 1. Otherwise, if video is failing to record for any reason, VPU$_i$ = 0.[1] The failures can be for any one of the following:

- Camera is not alive
- Network connection is severed
- Server is down
- No network sessions between VMS and camera
- Storage media is not available
- VMS recording service is not running

An aggregated VPU score for a collection of cameras is expressed as follows:

$$\sum_{i=1}^{n} \frac{VPIi}{n}$$

[1] **Note: a healthy camera stream that is not sending video data due to motion detection (i.e. sends no data because there is no change of scene) is still considered a VPU$_i$ =1.**

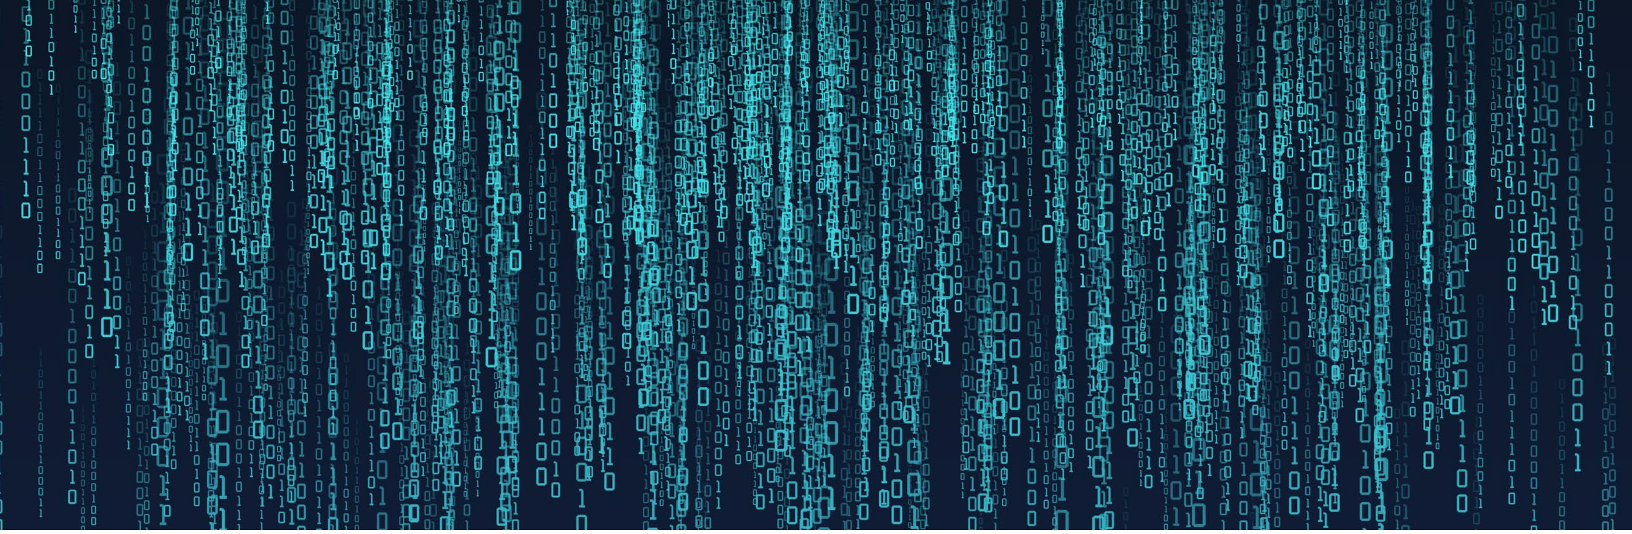

For example, the VPU of a server X would have a VPU_X equal to the sum of all VPU_i values of all the cameras recording to that server divided by the number of camera streams recording to that server. Similarly, the VPU for an entire site would be the sum of all camera stream VPU values in that site divided by the number of camera streams within that site.

This calculation is done for a single measure and has a scalar value, which can be calculated at any moment in time. However, VPU becomes most valuable when measured repeatedly over the course of time in regular intervals. From this, we can trend whether the overall health of a video stream, server, site or organization is stable, improving or decaying. More importantly, this becomes a key performance indicator (KPI) that is critical to quality for customers and from which we can correlate predictive analytics and root cause analysis.

Video Stream Delivery Index (VSDI)

VSDI measures the performance impact of saturation or decay of a video network on video delivery completeness. VDI is expressed as a percentage and is defined in detail below.

This is different from VPU in that VPU measures video stream paths that are in a failed state. VSDI measures the health of video streams that are still recording data but due to saturation of networks, load in recording servers or ingestion limitations in storage subsystems are losing portions of video data along an active video path. This is because, unlike typical data traffic which might have time-flexibility around the arrival of data and storage of data, cameras send video data continuously, using methods that tolerate some loss of data in favor of keeping up, in real time, with fresh video data getting generated.

Even when reliable transport protocols are being used, congestion in the network fabric can cause these connections to break and re-establish because they cannot wait indefinitely for all the packets to make it across. Bursts in data due to motion in multiple camera streams may be difficult to store fast enough to avoid overflowing VMS buffers. As a consequence, this congestion can result in dropped packets, which in turn can cause moments or several seconds of video frames to be lost.

For example, normal operation for a camera might be configured to generate a video stream at 8 frames a second at a resolution of 4 Mega pixels. At maximum resolution and throughput, this could produce a data stream of well over a Gigabyte of data per minute. Compression and motion detection can significantly reduce this. However, it can be quite significant. Compound this with a system connected to dozens if not hundreds of cameras and the stress on networking, compute and storage resources can be tremendous.

Still, these loads can be calculated in advance and enough capacity can be configured into the system to handle it. However, over time, changes to the configuration can put capturing all the video at risk. Typical problems can be associated with people adding cameras to the configuration without understanding their impact to the original design. Another cause can be people changing the configuration of cameras (e.g., increasing the resolution, changing the codec, or increasing the frame-rate), causing them to produce more data than the infrastructure was designed to handle.

Additionally, people can inadvertently add software to servers or change performance parameters causing servers to become loaded. Finally, systems can decay causing noise on the critical network paths, lowering performance or increasing latency associated with storage. All these can affect the performance of the overall security video application and create risk for video stream quality.

VSDI is a measure for **video stream delivery quality** types of problems. As with VPU, each camera stream has a VSDI measure

and then this VSDI measure can be aggregated in logical groups of camera streams, either by server, by site, by company or any other collection users choose to evaluate their infrastructure. Like VPU, VSDI is a percentage its lowest possible value being 0 which implies the system has detected frame loss for that sample. Furthermore, a value of 100% implies all the frames of video have been transmitted successfully. Unlike VPU, VSDI could have other values that are greater than zero (0%) but less than 100%. This is to reflect the property that there may not be frame loss yet, however, the system is detecting varying degrees of risk to an individual video stream path.

For example, a VSDI value of 80% may occur because the system is starting to see a CPU load exceed 80%. It doesn't mean video frames have been lost but it is starting to become a risk. A rising storage queue depth over the course of several measures could also reduce the VSDI as well or drop packet events from the network that are still within acceptable ranges. A VSDI of 20% implies that the infrastructure is getting pushed to its capacity and dropped frame events are eminent and may have already occurred.

As before, VSDI becomes another KPI which we can use to correlate information from multiple sources to get to root cause and develop predictive analytics about what the problem might be in a user's infrastructure and perhaps early detection or prediction of failures so users can take steps to moderate or add more resources to their infrastructure before they start losing data. Given as a percentage between 0%-100%, these values can be aggregated to create overall measures for a collection of camera streams recording to the same media, or flowing through the same server or across an entire site.

Video Retention Compliance (VRC)

VRC measures the extent to which a video surveillance recording system meets its video recording retention goals. IIt is expressed as a percentage, and is calculated for each camera individually and also for the system overall. An individual camera's score can be higher than 100% if the camera is exceeding its retention goals. A system's overall score is calculated differently as explained in the section **Camera and System Scores** on page 15.

Any video surveillance recording system has physical capacity, which sets some upper limits as to the amount of video data that can be recorded. For practical reasons, when these systems run out of space, they must delete older video data to create room for newer video data. This is a necessary and acceptable strategy since most of the video data is somewhat worthless if nothing important has happened. For example, a video stream of a rarely used back door to a facility doesn't need to be saved in perpetuity if nothing has ever happened that is worth watching. The time between when the video stream data is first captured to the moment it has to be deleted to make room for new video is called the retention period.

This video retention period is a key dimension of the design of the system and has a significant impact on costs, i.e., more retention time implies more storage. At a basic level, the retention period represents the time an organization has to evaluate whether something has happened that would warrant saving video "clips of interest" more permanently. For example, if video retention was only 30 days, you may find that a break-in that happened to your warehouse a few weeks ago, would already have been over-written when you went to look for it.

Within an organization, key stakeholders such as the CEO may have an expectation of how long this video should be retained which was used to justify the purchase of enough storage to accommodate this expectation. In some cases, regulatory agencies or expectations imposed by customers have established standards for how long this data must be preserved for their own auditing of operations.

Different camera streams may also have different objectives. One example, for physical security, exterior camera views of a facility may only need to be retained for as long as it takes to discover perimeter security breaches and then perform an investigation, anywhere from 2 to 4 weeks. Another example, cameras monitoring access to Data Centers that store financial information can have retention windows of several months. Saving all camera streams for the maximum retention period simply because some camera streams need to be saved for that long is extremely expensive. Many organizations will have different retention and deletion policies on different streams.

Therefore, the retention period is not necessarily static. As we've mentioned before, changes that can affect video stream delivery can also affect retention periods. To verify compliance, however,

traditional measures do not help. Free space is completely inadequate. Furthermore, measuring the oldest file at any one moment in time fails because some camera streams maybe in compliance with their retention goals while others may not be. Moreover, for certain systems, the presence of an errant file can make the system look like it is retaining data for a long time when actually on going storage associated with a particular video stream could be well below its requirements. Finally, a user with access to the system might delete data prematurely. As a result, the only way for even the most sophisticated organizations to verify that they are in compliance with goals is to have employees periodically verify each camera stream, one-at-a-time, to make sure the storage is still retaining all the data according to the stated goals. This is expensive and prone to human errors in measurement, especially over time.

Therefore, to effectively measure compliance, we must automatically examine storage utilization on a per video-stream basis. From this information, we can determine how much data is getting stored with any one particular camera stream on a daily basis and detect when that data starts to get deleted to create space. Measuring the system this way allows us to identify the difference between ongoing retention of video stream data versus an errant file that has been left behind but doesn't represent ongoing retention. It can also detect the errant premature deletion or the failure to delete data according to the retention policy. From this information, we can calculate the actual retention period, R_i, in terms of retention time for any camera stream, C_i.

To compare camera streams across a collection, it is important to normalize each stream against its individual goals (RG_i). Therefore, to calculate Video Retention Compliance for any camera stream, we use the following formula for each camera stream:

$VRC_i=R_i/RG_i$

Camera and System Scores

Unlike VPU and VSDI, a single video stream can actually have a VRC_i value which is greater than 100%, which reflects a situation where the data exceeds the retention goal. When we look at a collection of camera streams, VRC doesn't just aggregate in averages as a camera stream that significantly exceeds its retention goal could start hiding problems in other streams not meeting their retention goals. Rather, when normalizing, we use a maximum value of 100% for any camera stream that is exceeding its goal when we aggregate the streams together. This keeps the aggregate measure between 0% and 100% which facilitates subsequent roll-ups at server, site or company levels.

Other measures such as minimum, maximum or standard deviations can be helpful to identify site-wide or organizational-wide issues, but in this example, one critical camera being out of compliance could force an alert but having an overall measure of operational compliance is incredibly useful. Again, VRC, as with VPU and VSDI, give us a KPI to determine root cause and predictive analytics to help customers achieve their goals and address problems before they become disasters.

Conclusion

Security professionals invest a significant amount of resources in deploying video surveillance applications that range from the capital costs of cameras, servers, storage, and networking to the effort and expertise in deploying video management software and configuring the solution to do the job that key stakeholders expect.

To measure as required, Viakoo has developed key measures that answer three fundamental questions:

1. **Is each video stream path working properly?**
2. **Is the quality of that video data clean and complete?**
3. **Is each video stream's data retained for as long as required?**

The answers to these questions are effectively answered with the articulated VPU, VSDI and VRC measures respectively. From these measures, people responsible for this kind of infrastructure can accurately understand whether their systems are working properly and can focus in on components of the system that are impacting these measures adversely. Furthermore, automated analytic mechanisms can use these measures to drive root cause analysis processes and predictive analytics over time. Most importantly, enterprise security organizations can leverage these measures to drive operational excellence throughout their operations, getting the value and risk reduction expected from their IP video investments.