2018

Rate Adaptation for Reliable Video Transmission

VEGA SYSTEMS INC.

JULY 2018

Contents

Figures

Vega Systems Inc, 1999, S Bascom Ave, #700, Campbell, CA 95008, USA, 669-256-2357, info@vega25.com

Introduction

This white paper presents considerations for successful transmission of surveillance video to a remote destination. With no loss of generality, we will use cloud storage as the destination.

Two transmission schemes are considered.

1. A Static rate scheme in which video is produced at a constant rate.
2. ElastiStream, in which the rate of video production hugs the transmission rate that can be achieved at the site.

We see that ElastiStream provides significant raw performance advantages over a static rate scheme while being robust and scalable.

Metrics

We define latency as the amount of time that elapses from a video byte being produced, to its being successfully transmitted to cloud storage.

It is imperative for a solution to transmit the *highest quality* live video from a site with *least latency*. In addition, video that is buffered locally during link outage events, must also be transferred to the cloud with *highest quality* and *least latency*, after link restoration.

Latency Constraints

See Figure 1.

Video is produced in packets; each packet has video content of duration T. The transmission of each packet starts coincident with start of packet production. A slot time T in the "Video Transmission" timeline of Figure 1, indicates times at which acknowledgements are received for successful transmission of corresponding production slot content.

If production-transmission timing relationships are as in Figure 1, video that is produced at the beginning of a packet production slot is **guaranteed** to be available in the cloud, in the limiting case, by the end of the corresponding transmission slot. So, the maximum delay between a video byte getting produced, to its being available in cloud storage is $T + \alpha$. While one would expect the latency to be close to α , one must account for upload speed not being constant during a transmission slot. $T + \alpha$, is latency in the limiting case.

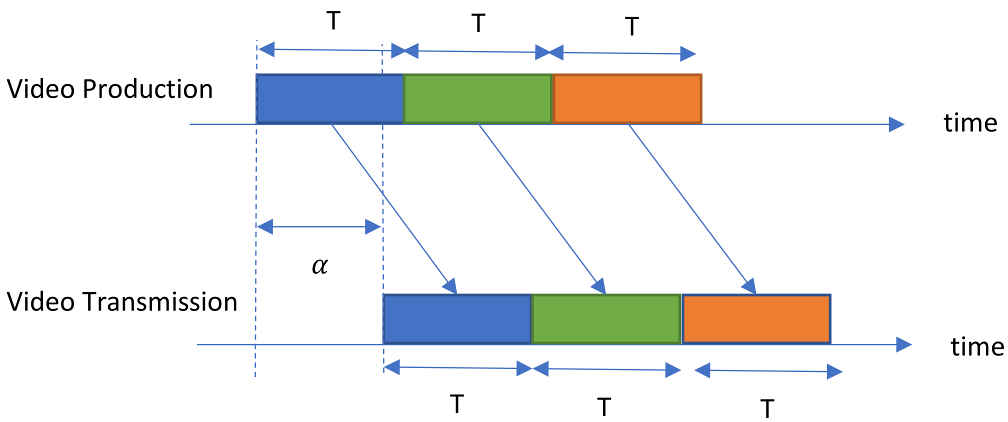

Figure 1: Latency

Production, Transmission Rate Relationship

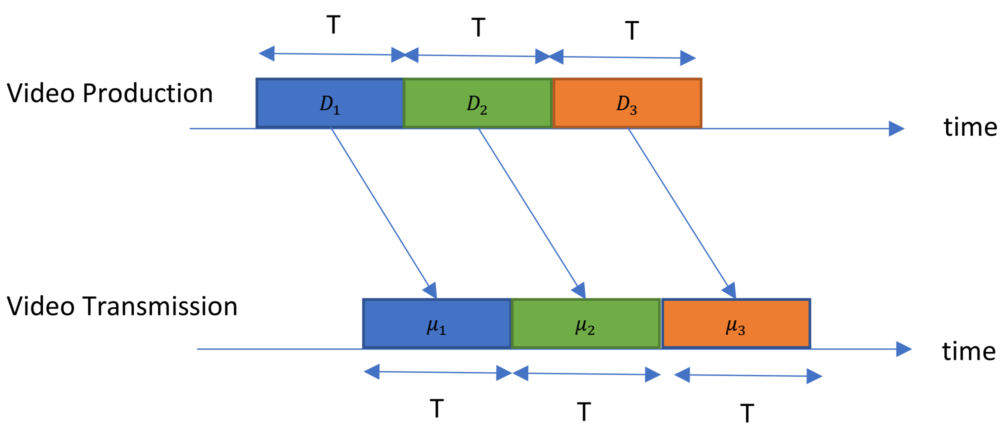

Figure 2: Rates

Figure 2 adds average video production and transmission rates to Figure 1. $D_1, D_2, D_3 \cdots$ are the average rates at which video is produced during shown slots. $\mu_1, \mu_2, \mu_3 \cdots$ are the average rates at which video ***can*** be uploaded during corresponding slots.

For the production-transmission timing relationship depicted in Figure 2 to be valid, one must have

$$D_i \leq \mu_i$$

Equation 1

Where $'i'$, indexes the packet slot.

Transmission Schemes

Two schemes are considered.

1. A Static Rate scheme in which video is produced at a constant rate.
2. ElastiStream, in which the rate of video production hugs the transmission rate that can be achieved at the surveillance site.

Static Rate Scheme

Here, we set the average production rate to be constant.

$$D_i = D$$

Equation 2

To meet latency requirements, Equation 1 should also be satisfied. So, D must be set to a value less than the *least average transmission rate*, μ_{min}. See Figure 3.

$$D \leq \mu_{min}$$

Equation 3

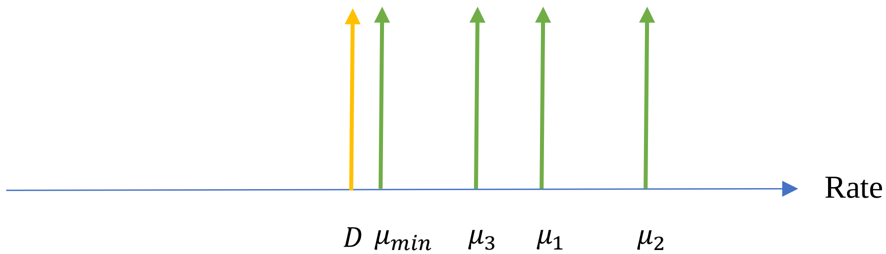

Figure 3: Static Rate, without outage considerations

In terms of underutilized uplink bits:

$$\Delta = \sum_i (\mu_i - D)T$$

Equation 4

In an uplink bandwidth constrained scenario, i.e. one in which video sources can produce higher amounts(quality) of data than the uplink is capable of transmitting, Equation 4 is a measure of quality degradation due to the design choice of not adapting video production bitrate.

Outage

We define outage as any event during which the average uplink transmission rate falls below production rate D.

Any system that needs to transmit information to a cloud server must be designed to handle connectivity outages. It must buffer video information that is produced, but not transmitted, during this outage interval. Once connectivity is restored, it must transmit all buffered content in parallel with live content that is produced after restoration.

To accommodate outage content, we need to back-off D by a back-off amount D_{buf}.

After link restoration, new content produced at rate D needs to be transferred in parallel with outage content. Outage content is produced stored at rate D, but transmitted at rate D_{buf}. So,

$$D + D_{buf} \leq \mu_{min}$$

Equation 5

The duration of an outage event is designated O. Let R be the copy time after link restoration, during which all buffered content is transmitted. Then,

$$\frac{R}{O} = \frac{D}{D_{buf}}$$

Equation 6

So, one needs to trade-off live video transmission quality with post outage copy-back speed. Live video quality is penalized constantly, to accommodate potential outage scenarios, even when there is no buffered content that needs to be sent over. See Figure 4.

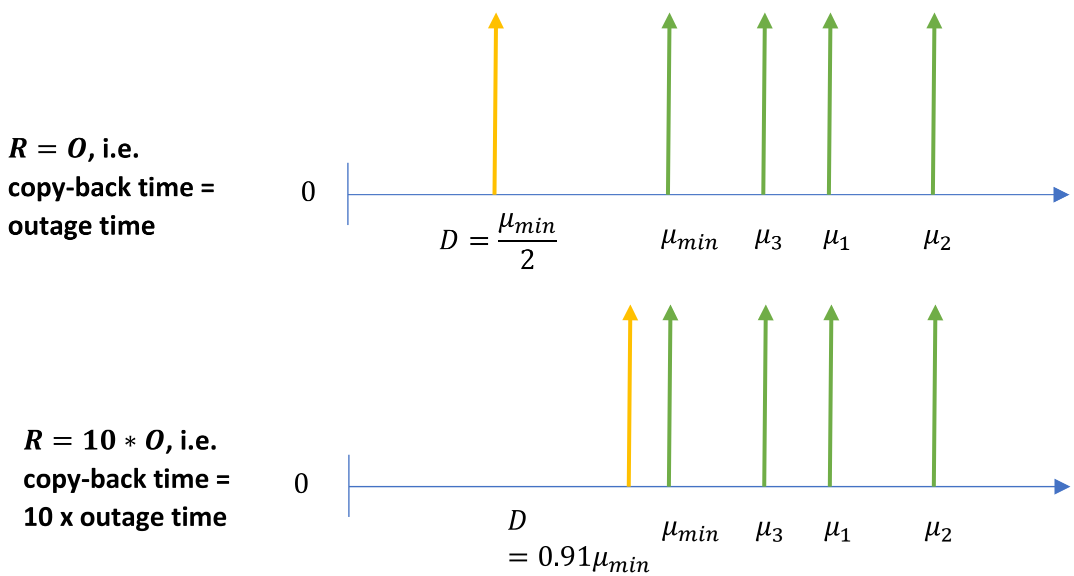

Figure 4: Live video quality, copy-back time tradeoff

Vega Systems Inc, 1999, S Bascom Ave, #700, Campbell, CA 95008, USA, 669-256-2357, info@vega25.com

Example

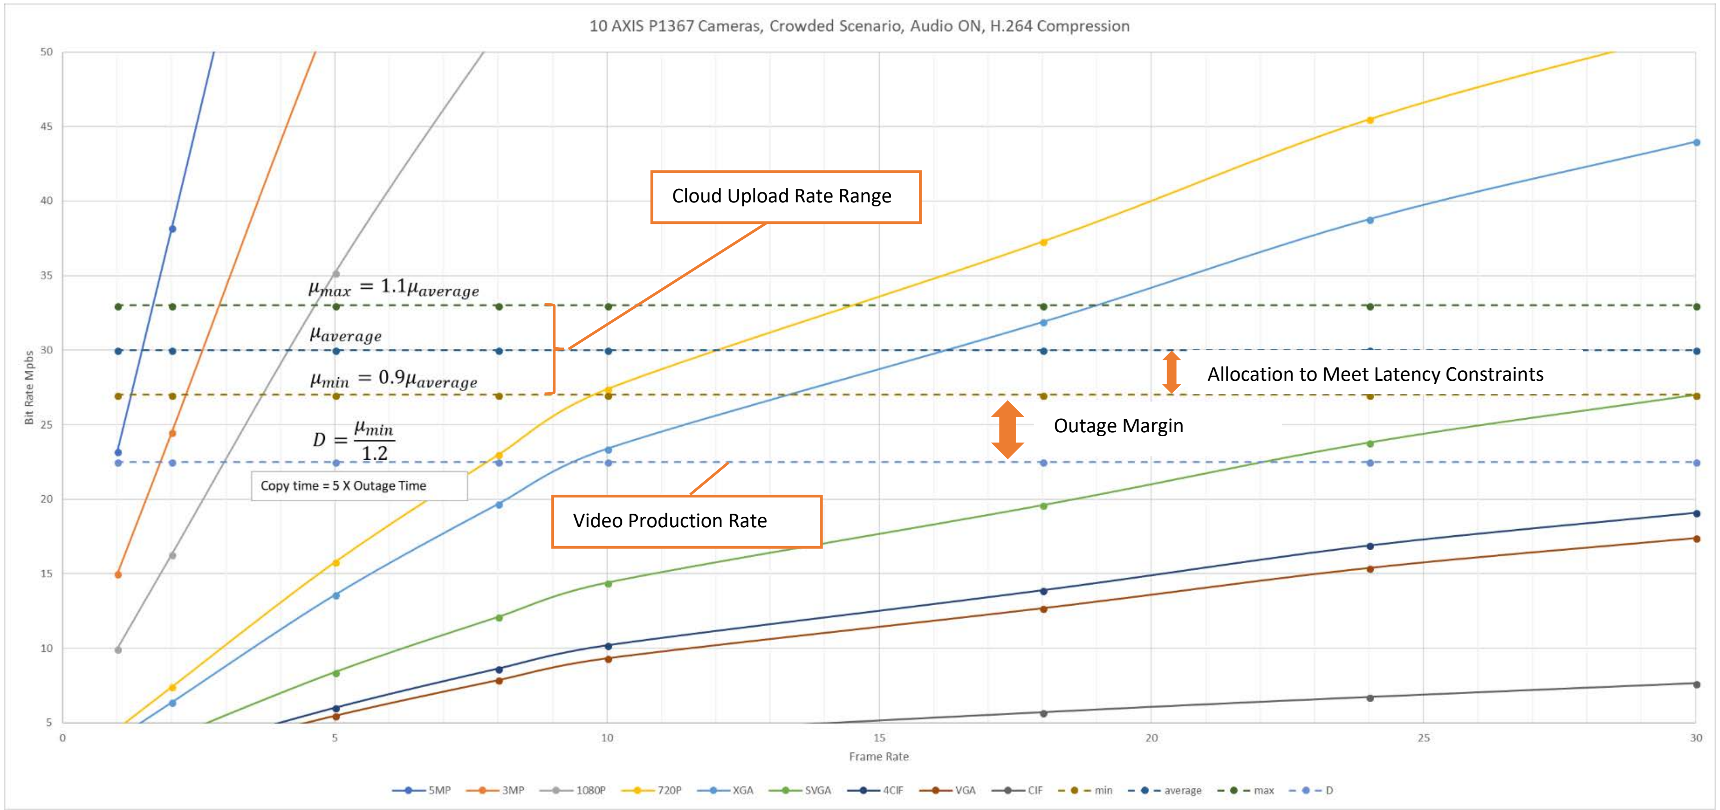

Figure 5: Static Rate Setting

Figure 5 shows the bit rates needed at a 10-camera installation, for different resolutions and frame rates. In this figure:

- The upload bit-rate varies between $\mu_{min} = 27\ Mbps$ and $\mu_{max} = 33\ Mbps$, with an average rate of $\mu_{average} = 30\ Mbps$.
- With no outage consideration, with only latency constraints, the maximum value of video production rate can be μ_{min}.
- To add in outage margins, with buffer to cloud copy time = 5 X outage time, we need to back off the video production rate to $D = 22.5\ Mbps$.

- If one were transmitting frames from all cameras at 720P resolution, one could send 7.7 frames/s, while meeting latency constraints and providing for outage content copy.
- Note that, due to lack of rate adaptation, we constantly transmit at 22.5 $Mbps$; even if there is no outage content to transmit, wasting precious uplink capacity.

ElastiStream

ElastiStream adapts video production rate in each slot so that

$$D_i \leq \mu_i$$

See Figure 6.

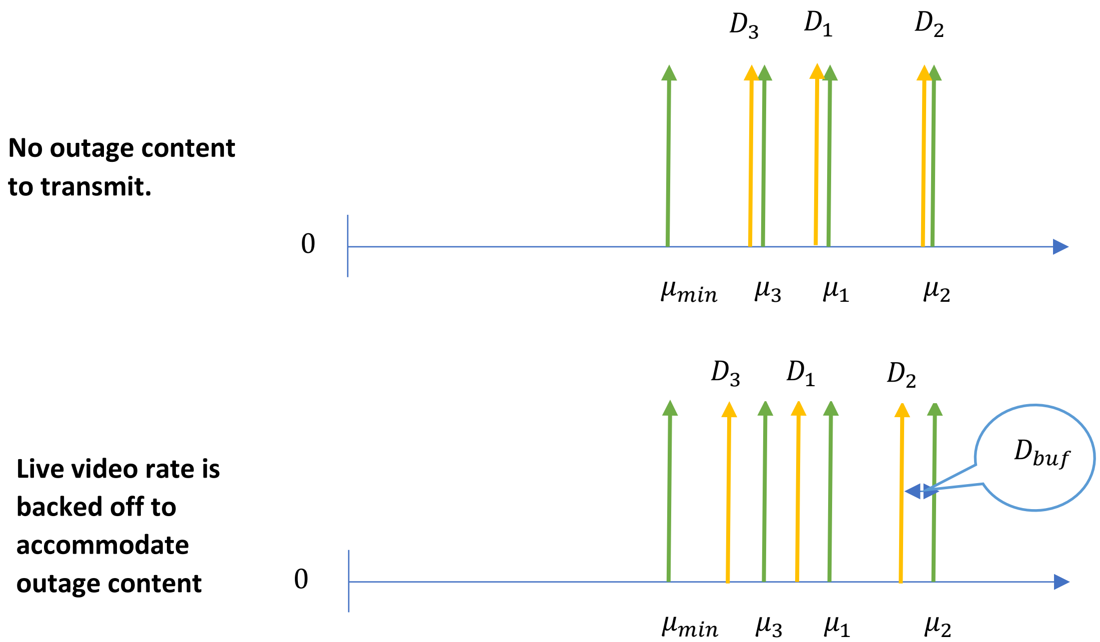

Figure 6: ElastiStream

Vega Systems Inc, 1999, S Bascom Ave, #700, Campbell, CA 95008, USA, 669-256-2357, info@vega25.com

Outage

ElastiStream adapts bit rate to changing channel quality. So, we define outage as an event during which, μ_i goes far below any valid production rate D_i. For such μ_i, is not possible to meet live video latency requirements. The best one can do is to produce video at the least valid production rate D_{min}, transmit at μ_i and accept the latency that comes along with it. The data that is not transmitted is buffered locally.

After link restoration, the rate of production (and transmission) of live video (D_i) and the rate at which video in the buffer is transmitted (D_{buf}), must satisfy

$$D_i + D_{buf} \leq \mu_i$$

If D_{buf} is fixed by design choice,

$$D_i \leq \mu_i - D_{buf}$$

Say, the buffer contains video of time duration O that needs to be sent to the cloud. There are several options to copy this to the cloud.

- If D_{buf} is large enough in each slot, for stored video to be converted into a stream at this rate, and we do this, we can transmit video of duration O, in time $R = O$.
- If either D_{buf} is not large enough, or by design choice, we want the stored video to be converted to higher quality, copying video of duration O will take more time than O. $R > O$.
- If either D_{buf} is large, or by design choice, we want to minimize copy time R, we can convert stored video to a rate lower than D_{buf}. Here, $R < O$.

Unlike a static rate solution, the designer has a lot of flexibility.

Vega Systems Inc, 1999, S Bascom Ave, #700, Campbell, CA 95008, USA, 669-256-2357, info@vega25.com

Example

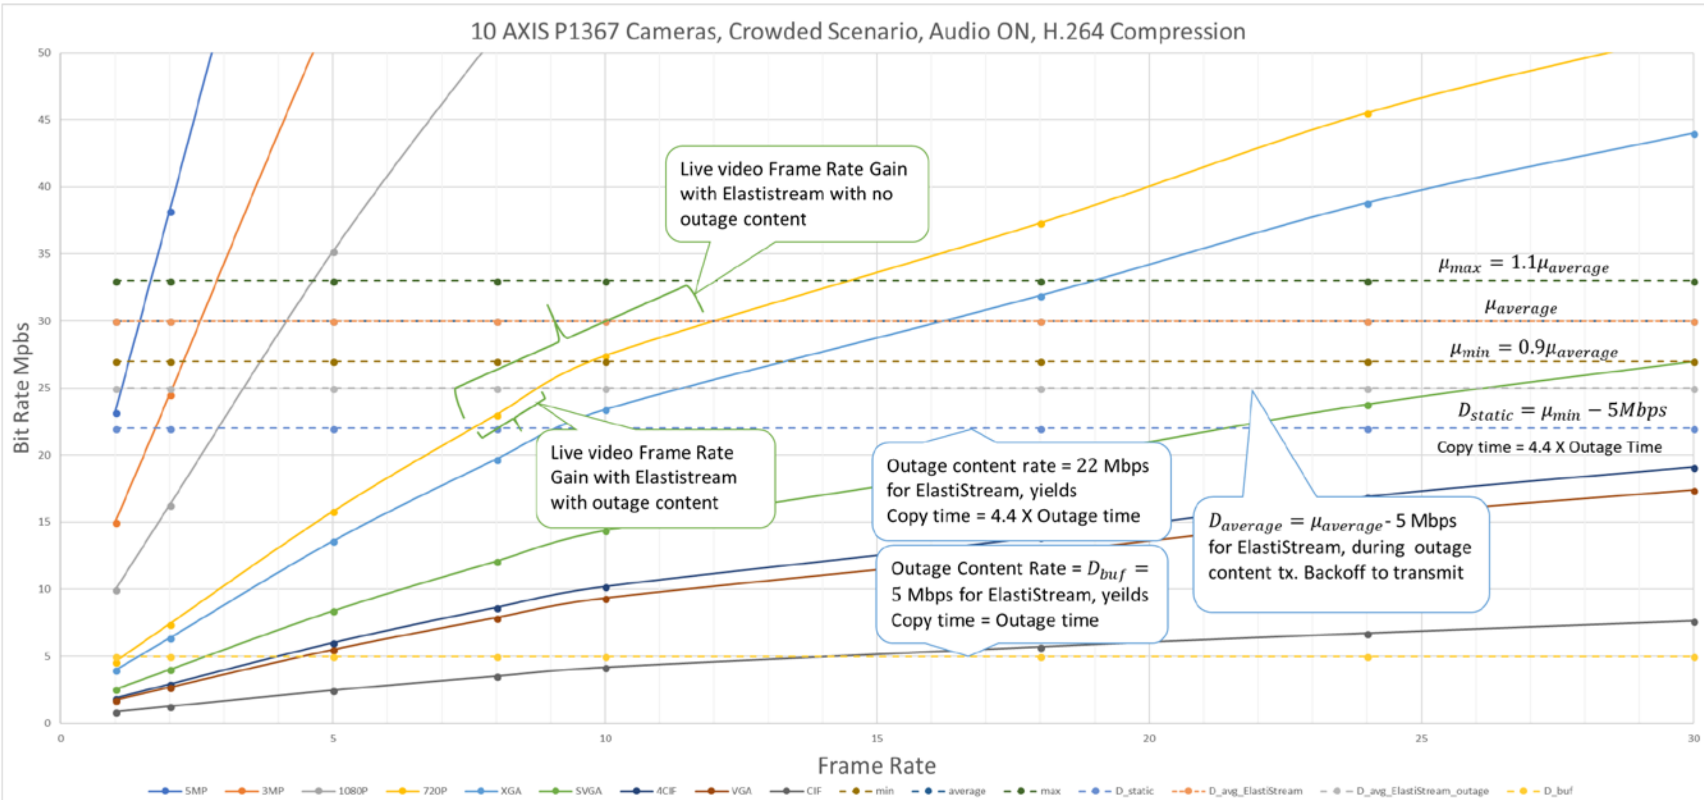

Figure 7: Gains from ElastiStream

Figure 6 shows gains achievable with ElastiStream, over Static Allocation. Here:

- The cloud upload bit-rate varies between $\mu_{min} = 27\ Mbps$ and $\mu_{max} = 33\ Mbps$, with an average rate of $\mu_{average} = 30\ Mbps$.
- When no outage video needs to be sent, video production rate D_i, for ElastiStream hugs the cloud upload bit-rate with an average $D_{average} = \mu_{average} = 30\ Mpbs$.
- A $5Mbps$ allocation is made for sending locally buffered outage content, in both ElastiStream and Static Allocation.

- Factoring latency and outage, the maximum production rate that can be set for Static Allocation, is $D_{static} = \mu_{min} - 5 = \ 22\ Mbps$. This does not change even if there is no outage content to send.
- With no outage content, ElastiStream provides an average production rate $D_{average} = \mu_{average} = 30\ Mbps$, that is $8\ Mbps$ higher than Static Allocation (D_{static}). *If one were transmitting 720P content, ElastiStream can accommodate frame rate of* $12\ fps$*, as against* $7.5\ fps$*, with Static Allocation*. The amount of frame rate gain is inversely proportional to the resolution.
- With Static Allocation, the $5Mbps$ allocation made for sending locally buffered outage content yields an outage buffer copy time of 4.4 X Outage time. $(\frac{22}{5} = 4.4)$
- The live video transmit rate with ElastiStream when outage content is present, is $D_{average} - 5 = 25\ Mbps$. This yields a frame rate that is larger than the Static Allocation scheme, at the same resolution.
- With ElastiStream, there are many ways in which $5\ Mbps$ can be utilized to transfer outage content:
  - If the outage content is rate adjusted to $5\ Mbps$, one can transmit all outage content in equal time. Copy Time = Outage Time.
  - If the outage content is rate adjusted to $22\ Mbps$, Copy time = 4.4 X Outage time, same as the Static Allocation scheme.
  - Values of outage content rate $< 5\ Mbps$, yield Copy Times < Outage Times.

Vega Systems Inc, 1999, S Bascom Ave, #700, Campbell, CA 95008, USA, 669-256-2357, info@vega25.com

A list of advantages of ElastiStream

Video Quality, Outage copy back time

Because ElastiStream utilizes the capacity of the uplink better, it can send more bits. These bits can be used either for increasing the quality of video that is sent (frame rate, bitrate, resolution), or to support more video sources. In addition, because of rate conversion support, it enables faster transfer of outage content to the cloud.

Adding, Deleting, modifying sources

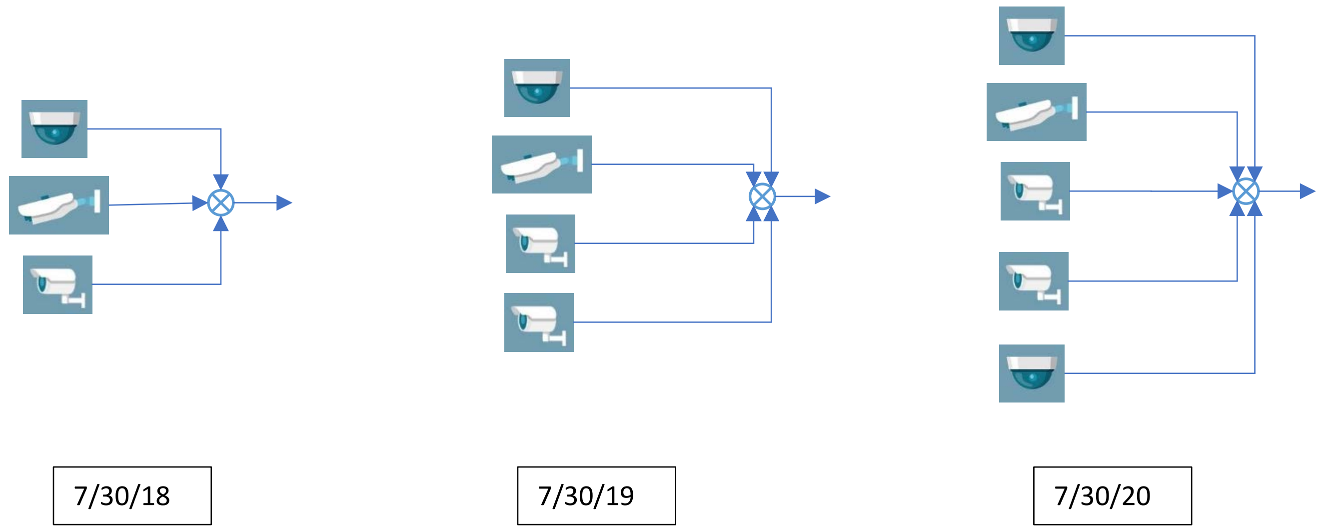

Figure 8: Source changes at a site

Over time, a site may add, remove cameras or change resolution of cameras. ElastiStream provides a hands-free way of adapting the bit rate of all sources to these changes. Else, one must manage source bit rates in a painstaking manner. See Figure 7.

Multiple Cloud Data Centers

Cloud upload speed is sensitive to the data center to which the upload is done. If a site or some streams at a site need to migrate from one data center to another, the camera rates must be adapted to accommodate the bitrate that the network provides to the new destination.

Connectivity of a cloud service provider to the internet may change over time. ElastiStream provides a hands-free way of adapting the bit rate of all sources to this.

See Figure 8.

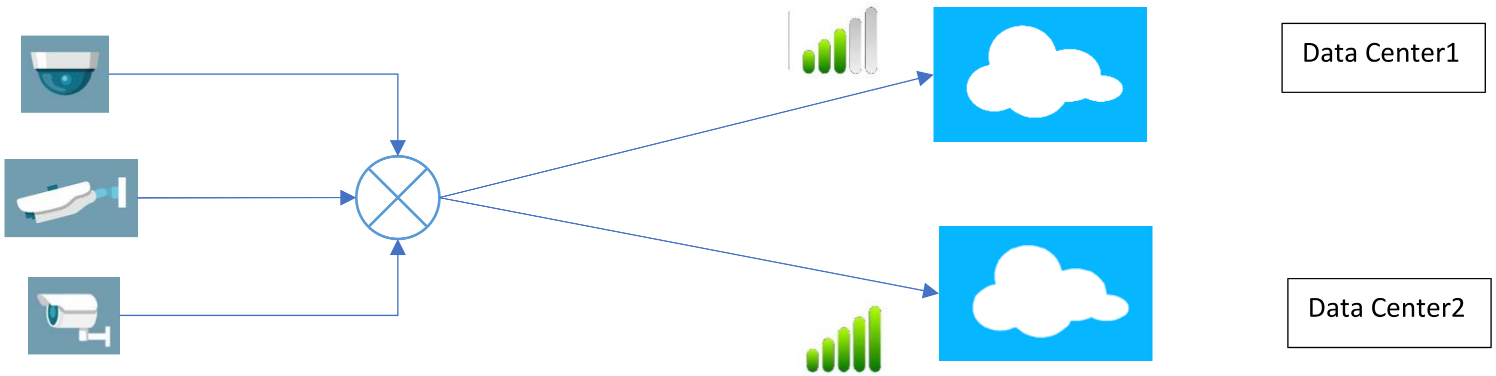

Figure 9: Handling Multiple Data Centers

Changing ISP, ISP tier, Long term changes in uplink quality

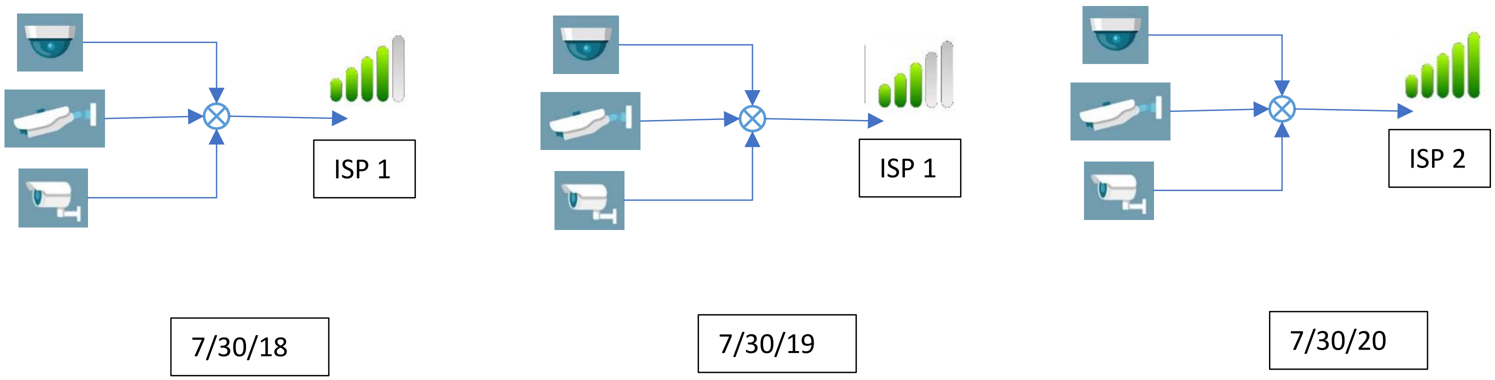

Figure 10: ISP Changes

The quality of the uplink provided by an ISP may change long term. If a site decides to change the ISP or the ISP tier, the uplink transmission rate changes. A static rate setting neither capitalizes on improving speeds nor reacts to reducing quality. ElastiStream handles this.

Links with widely varying uplink quality

Wireless uplinks show large variation in quality. A static rate choice that needs low latency can only produce content lower than the least uplink quality.

Figure 11: Drones experience rapidly varying uplink

Drone-based surveillance, police vehicles (Figure 11, Figure 12) need low latency, multi-stream live and outage video support. Both communicate over rapidly changing wireless channels. For a static rate scheme to work, its bit rate would have to be set very low. ElastiStream is the only good option.

Vega Systems Inc, 1999, S Bascom Ave, #700, Campbell, CA 95008, USA, 669-256-2357, info@vega25.com

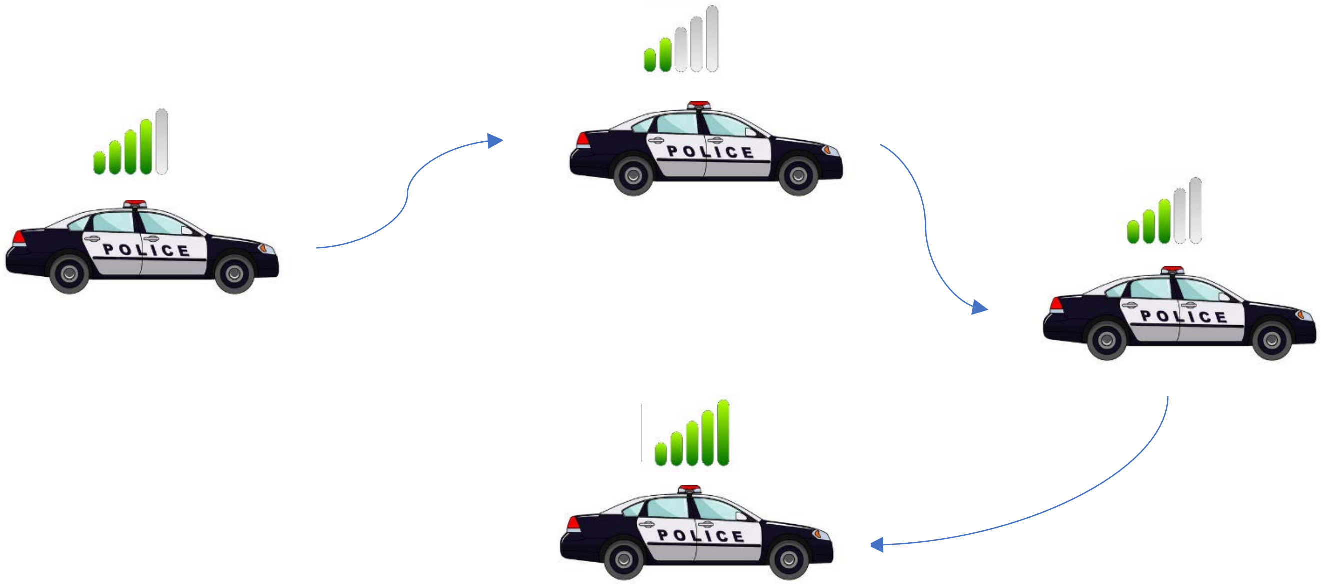

Figure 12: Police Vehicles also experience rapidly varying uplink

Site to Cloud Upload Speed: A Note

The upload speed from a site to a cloud data center is different from and usually less than the raw upload speed advertised by an ISP for that site. These two speeds are measurements to two different points on the internet. The upload speed advertised by an ISP is usually measured from the site to a near server that is part of the same ISP's network. The upload speed to a cloud data center is the speed to the cloud server of interest; the server could be anywhere within the ISP's network or could be connected to a different ISP altogether.

Summary

It is imperative that a cloud storage solution stores the *highest quality* live and outage video from a site with *least latency*.

ElastiStream provides significant benefits in both live and outage scenarios compared to Static Allocation. It provides a hands-free way to adapt to short term and long-term link quality changes, enables scalable video source management and supports data-center switching. It is clearly the better choice for all scenarios and the only choice for some.

Contact

Vega Systems Inc,
1999, S Bascom Ave, #700
Campbell, CA 95008
USA
669-256-2357
info@vega25.com