2020

Cirrus – A White Paper

12/21/20

VEGA SYSTEMS INC.

Contents

Table of Figures

Introduction

Having zero on-premises video surveillance server hardware is attractive in several scenarios. Here, cameras are deployed on-premises, all servers are deployed on public cloud infra. When XProtect is deployed such, all video consumers need to transport video out of cloud infra. Outgoing video is a bandwidth hog, live video contributes the most.

Public cloud vendors charge egress costs for all data that leaves cloud infra. This adds significant cost to applications that require live video consumption.

Cirrus is a XProtect plugin that mitigates these steep costs by streaming live video directly from cameras to clients, by-passing public cloud infra.

A comparison across public cloud vendors

Figure 1 compares cloud object storage, API and Egress costs per giga byte of data across different public cloud vendors, under a live video requirement assumption. Every giga byte is streamed out of the cloud. We see that egress cost dwarfs all other costs across vendors.

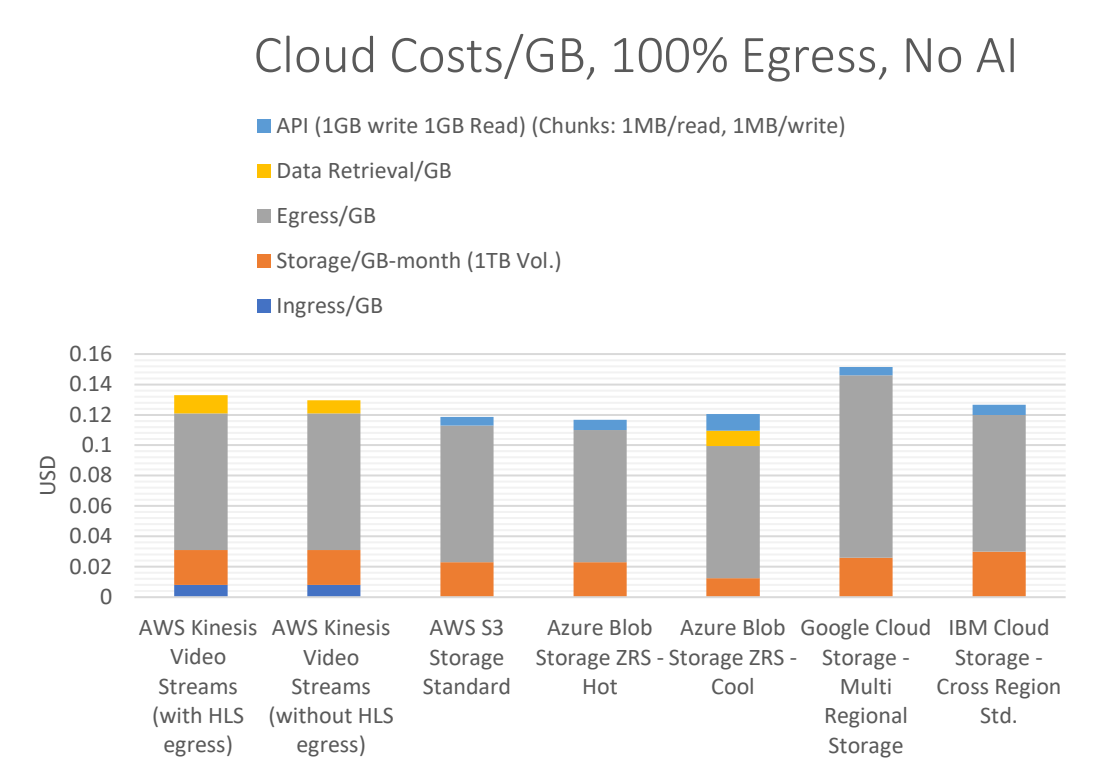

Figure 1: Insights into cloud video surveillance storage and retrieval costs

XProtect on AWS

When XProtect is deployed on AWS:

- Figure 2 shows that a single 4Mbps, 1080P stream at 30 frames per second could cost $125/month, if viewed live 24/7. (XProtect AWS calculator)

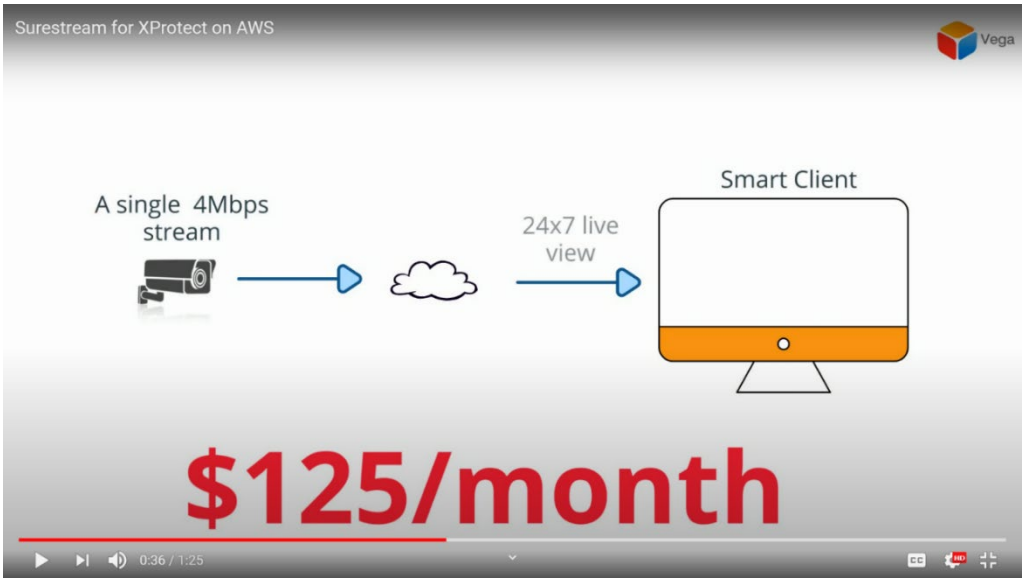

Figure 2: Egress costs for a single stream can be high

- Figure 3 shows that costs multiply with viewers. 2 cameras with 2 viewers each cost 4 times as much as 1 camera with one viewer. (XProtect AWS calculator)

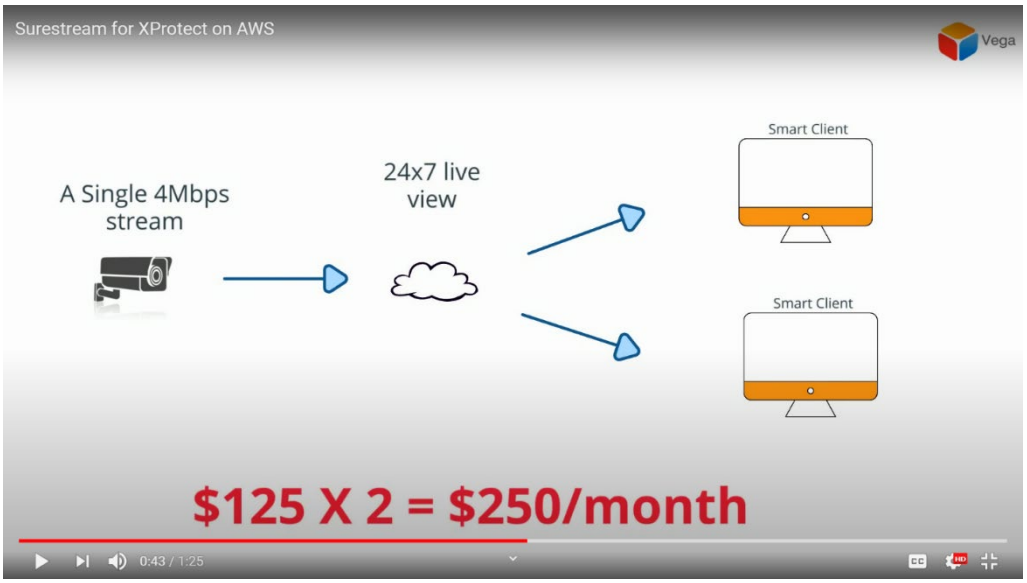

Figure 3: Egress costs multiply with cameras AND viewers

- Figure 4 shows Appstream 2.0 is enormously more expensive. (XProtect AWS calculator)

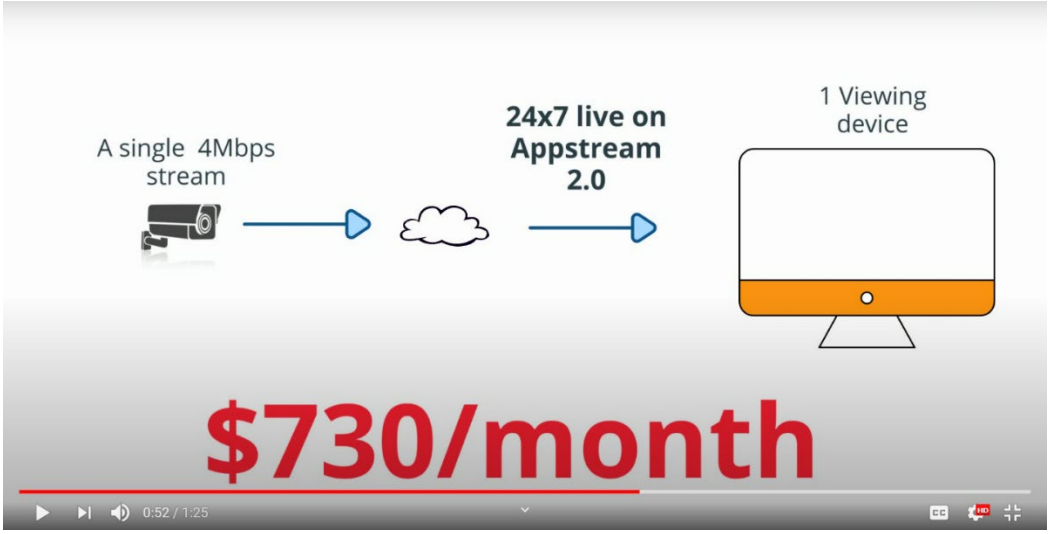

Figure 4: Appstream 2.0 is significantly more expensive

By streaming live video directly from cameras, Cirrus avoids these burdensome egress costs.

- Figure 5 shows how direct streaming with Cirrus avoids egress costs.

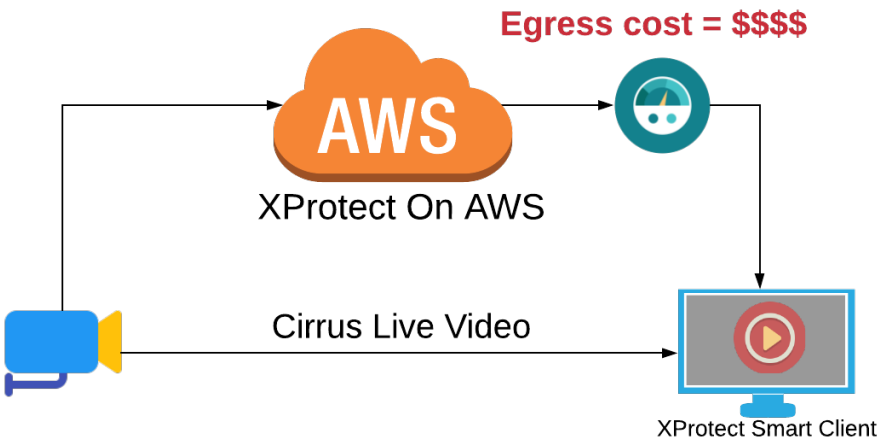

Figure 5: Cirrus helps save steep video Egress Costs by streaming directly to clients

- Figure 6, shows cost savings when XProtect on AWS is deployed with Cirrus. A cost of $10/camera/month is assumed for Cirrus. Note that, without Cirrus, egress costs multiply with the number of viewers.

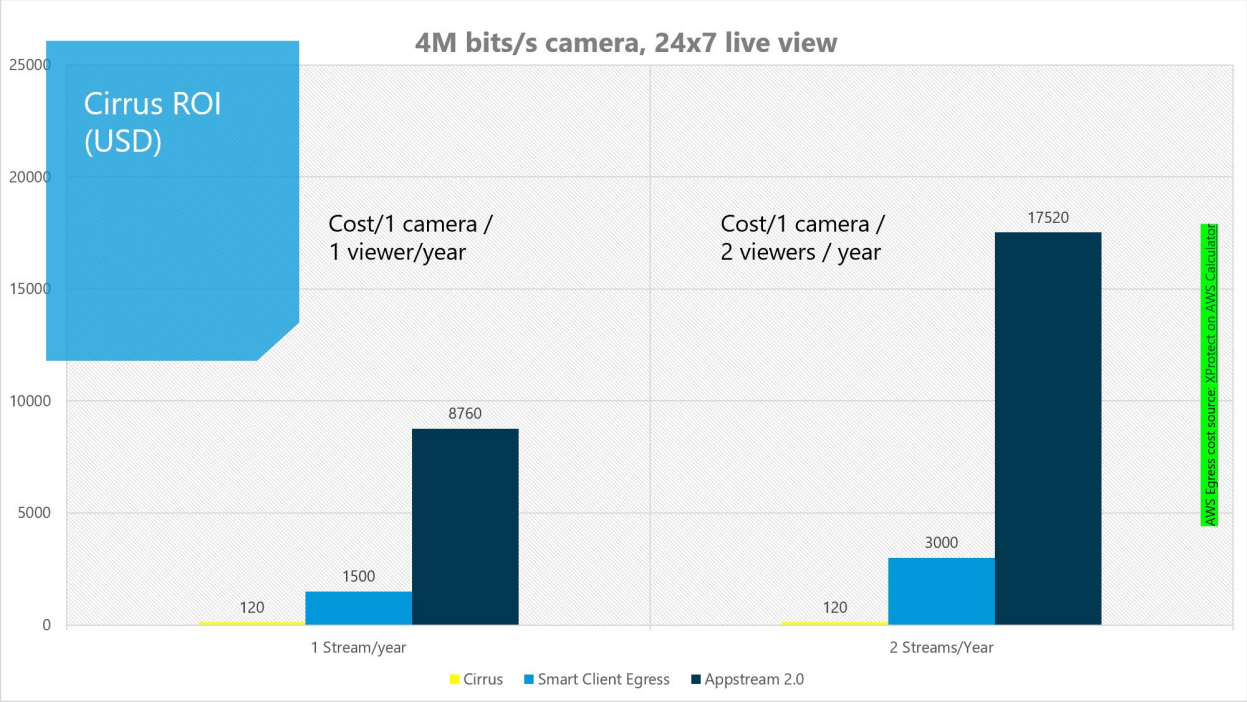

Figure 6: Enormous Cost Savings by using Cirrus

Eliminate Meaningless Video Roundtrips

Often, clients are co-located with cameras. Here, live video makes a meaningless round trip from the camera to the public cloud and back to the client. This loop back is unreliable, undesirable and increases latency & cost. Cirrus eliminates these drawbacks by having clients source video locally. Depicted in Figure 7.

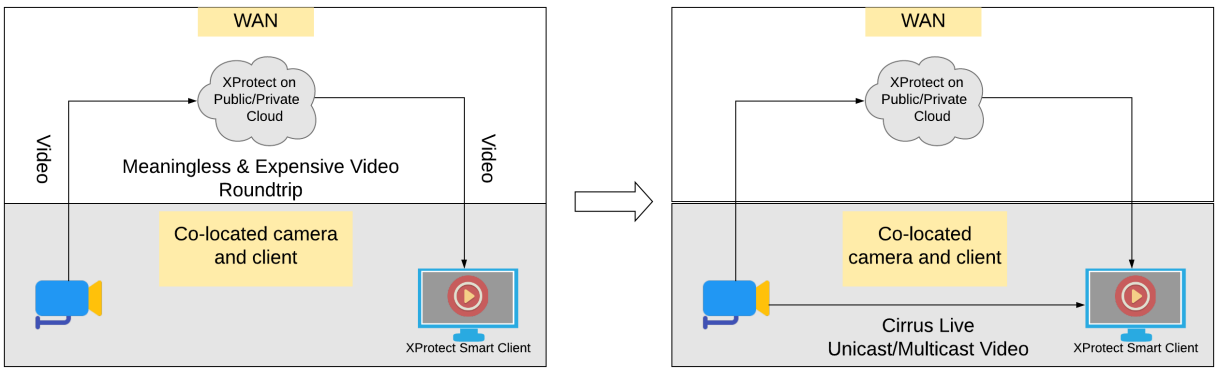

Figure 7: Avoid video roundtrips

Different Strokes for Different Folks

To cater to heterogenous needs, Cirrus supports several video streaming modes and multiple architectures.

WAN Smart Clients + cameras on public IP

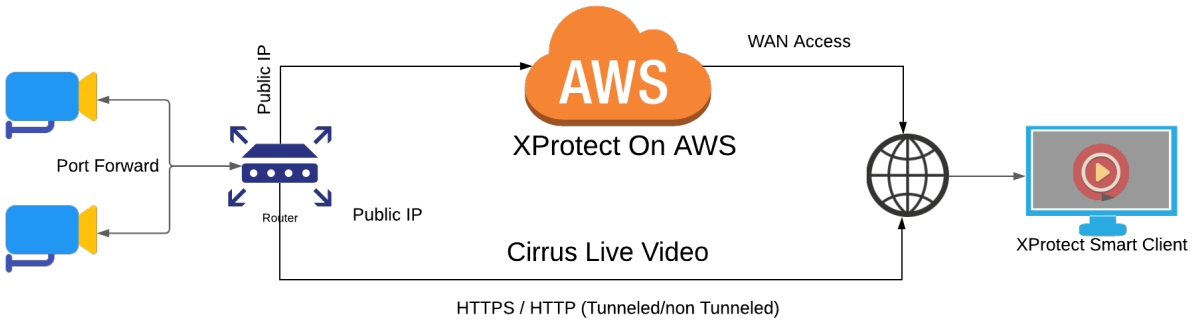

Figure 8: WAN Smart Client + Cameras on public IP

WAN Smart Client on VPN + cameras on VPN

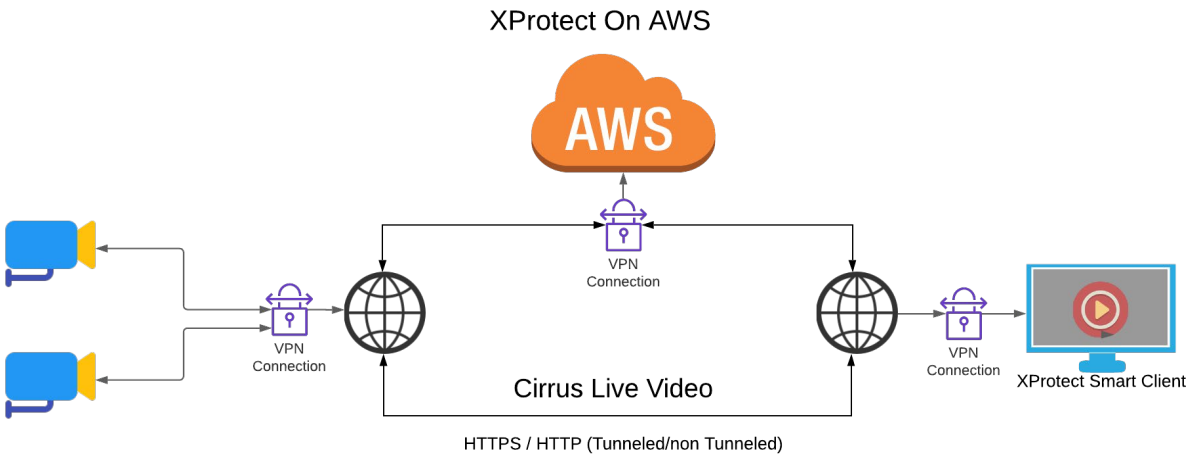

Figure 9: WAN Smart Client + Cameras on VPN

WAN Smart Client on VPN + LAN Smart Client + cameras on VPN

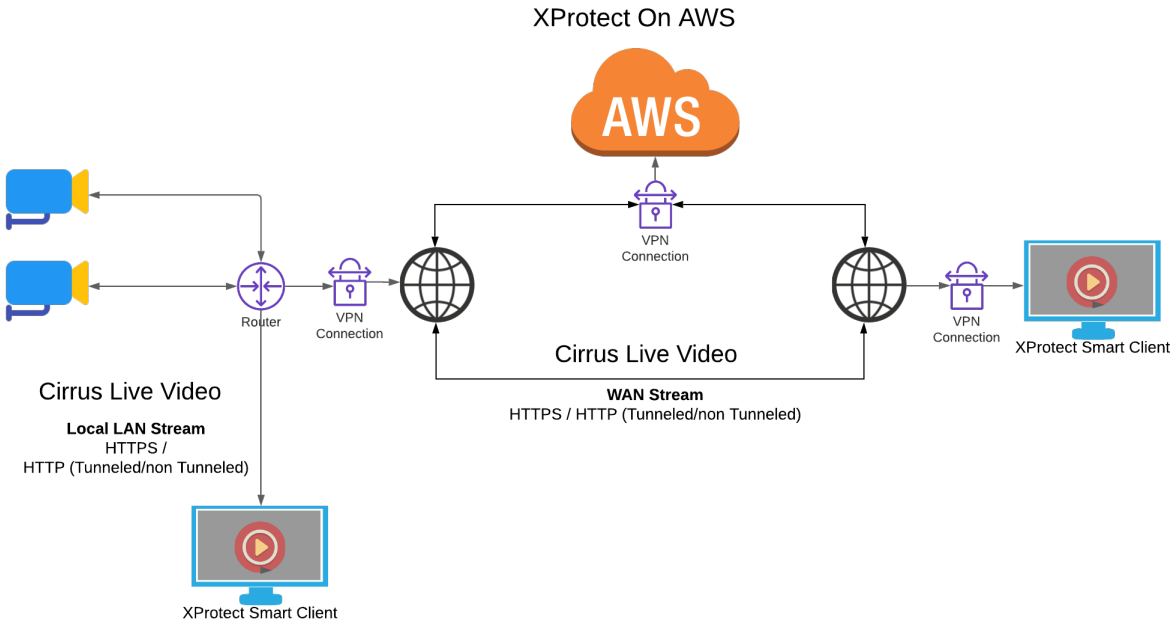

Figure 10: WAN, LAN Streams

WAN Smart Client + LAN Smart Client + cameras on Public IP

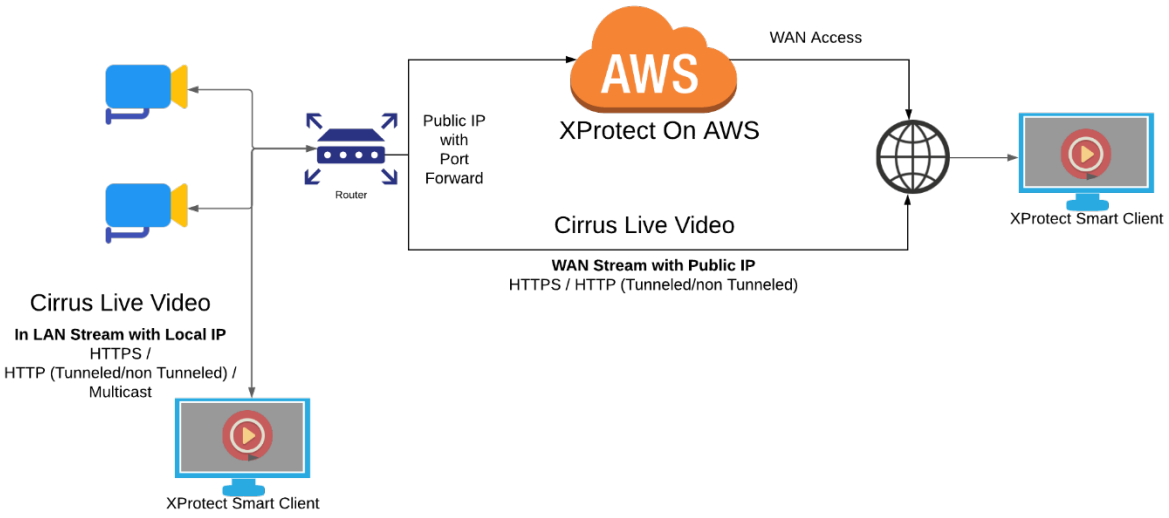

Figure 11: WAN Smart Client + LAN Smart Client + Cameras on Public IP

Better Live Video Robustness

'Appendix A: An Analysis of Live Video Robustness with Cirrus for AWS deployments' presents an analysis that shows that live video to the Smart Client is more robust when Cirrus is used. A comparison figure is repeated in Figure 12.

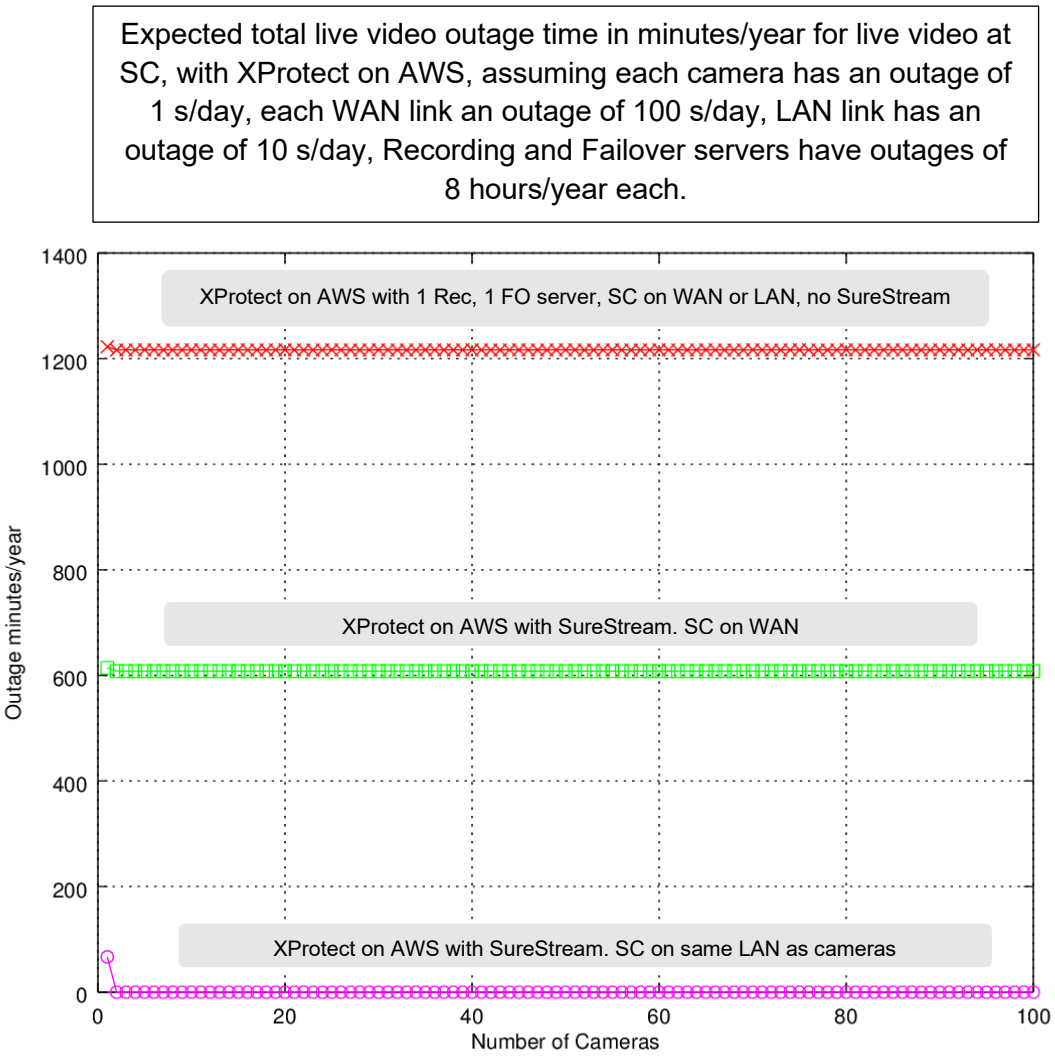

Figure 12: Better live video availability with Cirrus when XProtect is deployed on AWS

Security

Cirrus supports HTTPS streaming from cameras that support HTTPS via. Onvif-S.

Figure 13: HTTPS streaming support

Bandwidth Considerations

If enterprise VPN is designed to support multicast, each camera can provide a single multicast stream that gets distributed across WAN Cirrus consumers through this VPN. If not, Cirrus unicasts video from cameras to WAN Smart Clients.

The outgoing bandwidth at the site hosting cameras needs to be enough to accommodate these additional streams. For example, if a site has 100 cameras being recorded on AWS and 20 of these are viewed across two different clients on the WAN, the camera site will need 20% more outgoing bandwidth.

More Streams from Cameras

If enterprise VPN is designed to support multicast, then each camera can provide a single multicast stream that gets distributed across WAN consumers through this VPN.

If this is not the case, since WAN does not support Multicast traffic, cameras need to be able to provide as many additional unicast streams as the number of Cirrus clients needing a stream from the camera.

Note that while many cameras can produce only two different (in terms of resolution, encoder, frame rate) streams, many can stream out multiple replicas of these same streams.

Cirrus always uses the same live stream as XProtect.

One needs to check camera capabilities to verify the number of stream replicas that the camera can stream out.

Summary

Cirrus eliminates high live video egress costs charged by AWS when XProtect is deployed on AWS.

About Vega Systems Inc.

Vega Systems Inc. provides solutions for high availability video surveillance. Vega Systems' solutions are installed worldwide at airports, seaports and oil and gas facilities. For more details, visit: https://www.vega25.com.

Appendix A: An Analysis of Live Video Robustness with Cirrus for AWS deployments

All probabilities below are uniformly distributed over time. For example, a failure chance of p implies that, we can expect the component to fail for $p * 60 * 24$ minutes in one day.

XProtect with Rec. Server, Failover Server and Without Cirrus

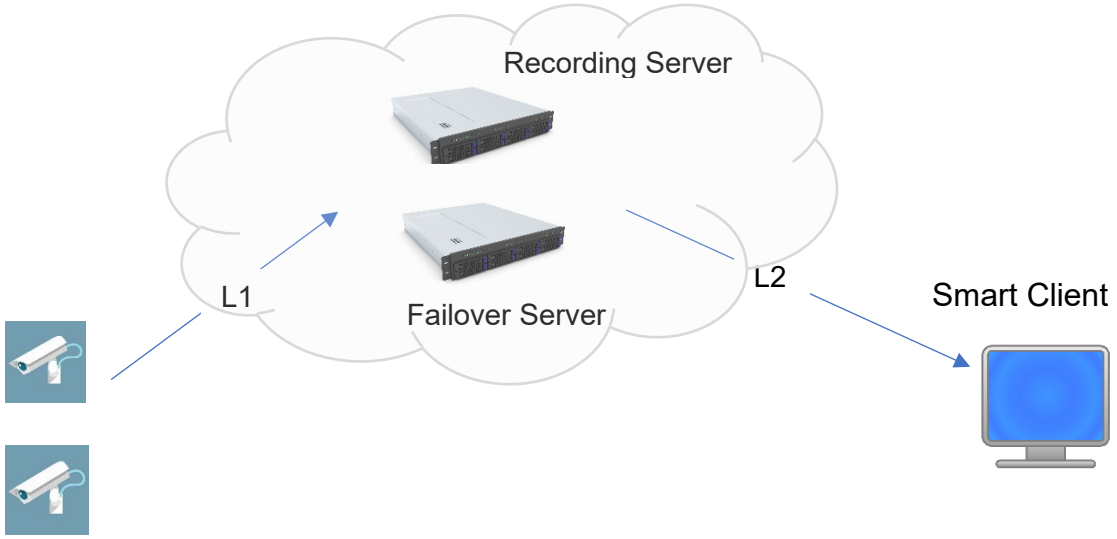

Figure 14: XProtect on AWS, without Cirrus

In Figure 14,

1. The failover server functions only when the Recording server fails.
2. We assume, L1 – L2 are network links with equal probability of failure. We will call this failure chance as l. It follows, the chance of no failure is $1 - l$.
3. The Smart Client and cameras could be on the same LAN or at different points on the WAN.
4. We assume, the recording server and failover server have equal probability of failure. We call this failure chance as r. It follows, the chance of no failure is $1 - r$
5. We assume, all cameras have equal chance of failure. We call this failure probability as c. It follows, the chance of no failure is $1 - c$.

The outage chance (f) that no video reaches the Smart Client (SC) Machine is

$$f = (chance\ that\ the\ L1\ is\ ok\ \boldsymbol{AND}\ no\ video\ reaches\ SC\ when\ L1\ is\ ok) + (chance\ that\ the\ L1\ is\ down\ \boldsymbol{AND}\ no\ video\ reaches\ SC\ when\ L1\ is\ down)$$

$$f = (1 - l) * (chance\ that\ no\ video\ reaches\ SC\ when\ L1\ is\ ok) + l * 1$$

Above, the $chance\ that\ no\ video\ reaches\ the\ SC\ when\ the\ L1\ is\ ok$, 'k' is

$$k = chance\ that\ rec\ is\ ok\ \boldsymbol{AND}\ no\ video\ reaches\ SC\ when\ rec\ ok$$
$$+ chance\ that\ rec\ is\ down\ \boldsymbol{AND}\ no\ video\ reaches\ SC\ when\ rec\ down$$
$$= (1 - r) * (chance\ that\ no\ video\ reaches\ SC\ when\ rec\ is\ ok) + r$$
$$* (chance\ that\ no\ video\ reaches\ SC\ when\ rec\ is\ down)$$

Above, $chance\ that\ no\ video\ reaches\ SC\ when\ rec\ is\ ok, v$ is

$$v = chance\ that\ L2\ is\ down\ AND\ no\ video\ reaches\ SC\ when\ L2\ is\ down, rec\ ok$$
$$+chance\ that\ L2\ is\ ok\ and\ no\ video\ reaches\ SC\ when\ L2\ is\ ok, rec\ ok$$

$$v = l * 1 + (1 - l) * (chance\ that\ all\ n\ cameras\ are\ down)$$

$$v = l + (1 - l)c^n$$

So, 'k' is:

$$k = (1 - r) * (l + (1 - l)c^n) + r * (chance\ that\ no\ video\ reaches\ SC\ when\ rec\ is\ down)$$

Above, $chance\ that\ no\ video\ reaches\ SC\ when\ rec\ is\ down, w,$ is

$$w = chance\ that\ Failover\ is\ ok * chance\ that\ no\ video\ reaches\ SC\ when\ Failover\ ok$$
$$+ chance\ that\ Failover\ is\ down * chance\ that\ no\ video\ reaches\ SC\ when\ failover\ is\ down$$

$$w = (1 - r) * chance\ that\ no\ video\ reaches\ SC\ when\ Failover\ ok$$
$$+ r * 1$$

But, $chance\ that\ no\ video\ reaches\ SC\ when\ Failover\ ok = v$

So,

$$w = (1 - r) * (l + (1 - l)c^n) + r$$

So, 'k' is:

$$k = (1 - r) * (l + (1 - l)c^n) + r * ((1 - r) * (l + (1 - l)c^n) + r)$$

And outage chance (f) that no video reaches the Smart Client (SC) Machine is:

$$f = (1 - l) * \Big((1 - r) * (l + (1 - l)c^n) + r * \big((1 - r) * (l + (1 - l)c^n) + r\big)\Big) + l$$

A XProtect on AWS Architecture with Cirrus

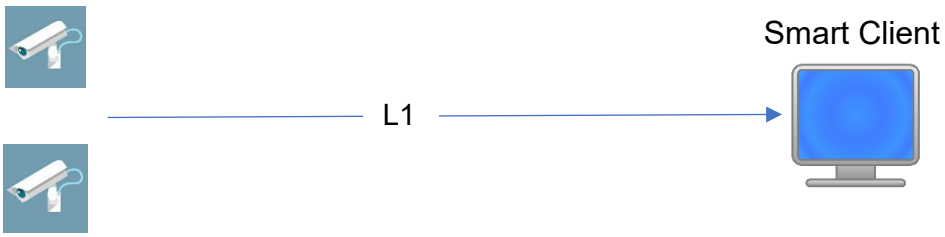

Figure 15: With Cirrus

Figure 15, shows an architecture with Cirrus, with cameras streaming directly to the client. Note that we assume that cameras are co-located and share a common link (WAN) to the client. We expect the WAN link to be of lower quality than a LAN link and thereby base our outage computation using this, even though each camera may use a different LAN link to get to WAN.

Here the outage chance, i.e. $chance\ no\ video\ reaches\ SC$, w, is

$$\begin{aligned} w &= chance\ that\ L1\ is\ ok\ and\ no\ video\ reaches\ SC \\ &\quad + chance\ that\ L1\ is\ down\ and\ no\ video\ reaches\ SC \\ &= (1-l) * (no\ video\ reaches\ SC\ when\ L1\ is\ ok) + l * 1 \\ &= (1-l) * (chance\ that\ all\ cameras\ are\ down) + l \\ w &= (1-l) * c^n + l \end{aligned}$$

Comparisons

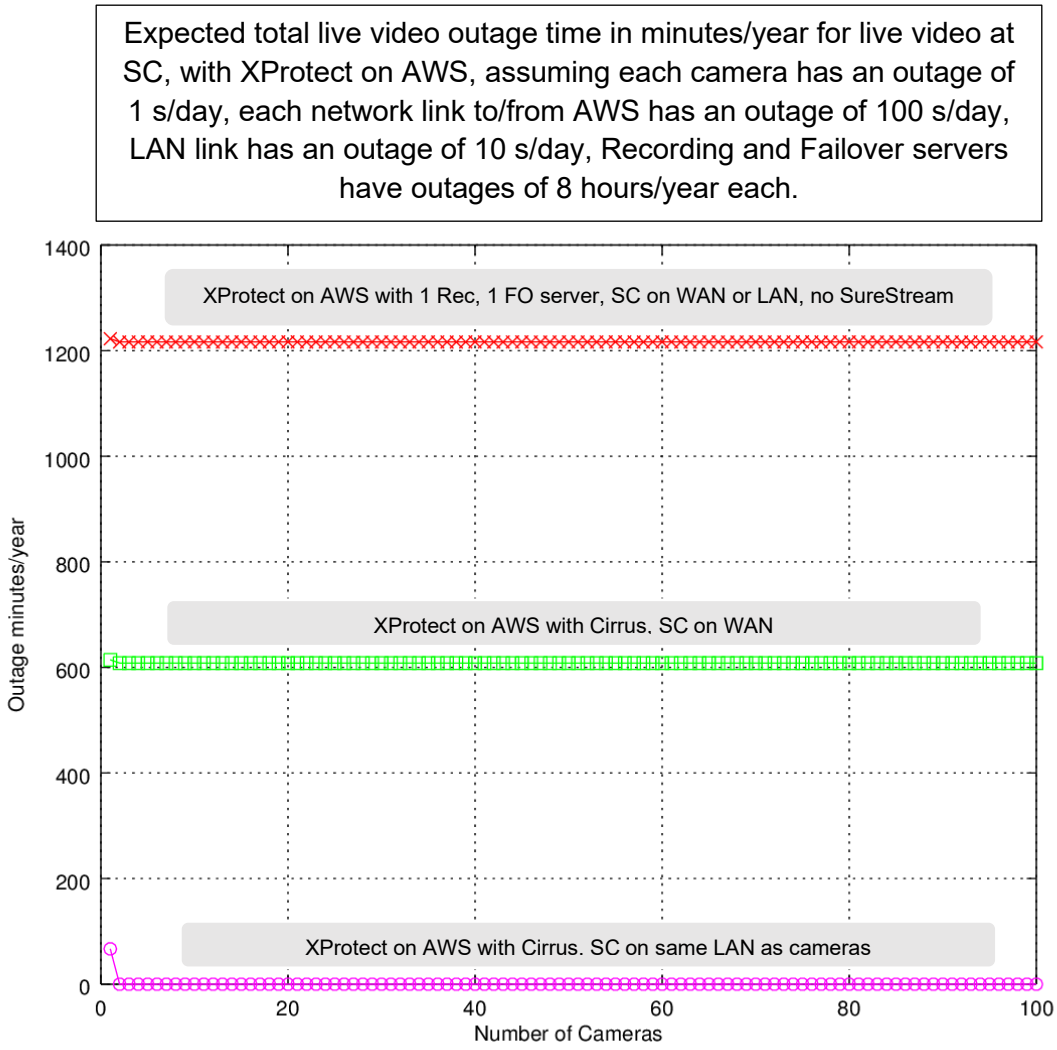

Figure 16: Better live video availability with Cirrus when XProtect is deployed on AWS

We use results in the previous sections to compare outages with and without Cirrus. When the SC is connected on the WAN, outage is with Cirrus is roughly half that without Cirrus. The driving reason in this case is that there are two independent WAN links that video must traverse before reaching SC when not using Cirrus. When the SC is connected on the same LAN as the cameras, outage with Cirrus is extremely low, since it leverages local connectivity to deliver video to the SC.