rasilient
built for surveillance™

What are the limits of the NVR architecture?

Sean Chang
CEO

Rasilient Systems
3281 Kifer Road
Santa Clara, CA 95051
www.rasilient.com

Contents

1 Forensic-Grade Video

When a serious crime occurs, law enforcement is in a race against the clock to gather and analyze all the video evidence before the perpetrators disappear. If the video is either unavailable or lacking in details, the people responsible for the crime may never be brought to justice.

The Boston Marathon bombing is a good example. After the bombing, the investigation team was under tremendous pressure to find who did it. The video evidence was crucial and needed to be accessed quickly.

This white paper focuses on *video availability* by comparing the NVR architecture to the Rasilient Surveillance-Defined Architecture. The emphasis is on mission-critical, large-scale deployments like cities, hospitals, transportation, and other critical infrastructure, where immediate video availability is critical for forensic analysis.

2 The NVR Architecture

The NVR (Network Video Recorder) is an appliance dedicated to a fixed number of IP cameras to record the surveillance video. It is a simple solution that pre-provisions the right amount of CPU power and storage capacity for the specified number of IP cameras (e.g., 16 channels).

The NVR architecture is effective when the deployment has a small number of cameras distributed across several geographic locations, especially when the configuration will not change very much in the future (Figure 1).

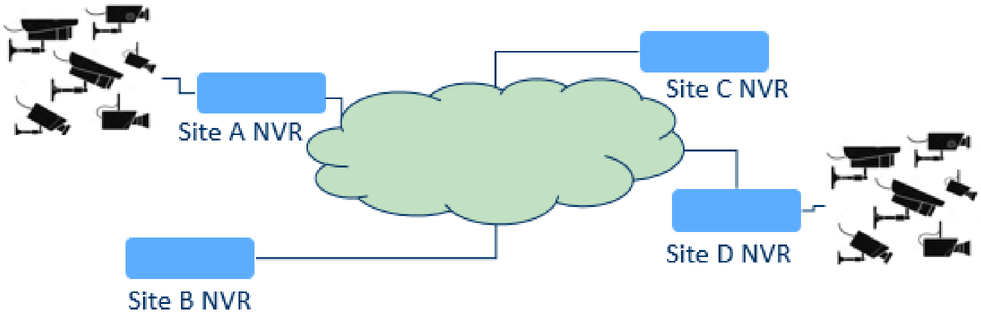

Figure 1 The NVR Architecture across Geographical Locations

For larger deployments that are mission critical, the NVR architecture begins to break down, particularly in these areas:

- Video availability
- Manageability
- The ability to add new features

3 Can You Access Video When the NVR is Down?

The question of video availability is directly related to hardware and software failures. For example, disk drives and power supplies can fail very easily, and both hardware and software problems can cause a system to crash or fail to boot.

Some of the problems can be addressed in a straightforward way. For example, the disk drive and power-supply failures can be handled by RAID and redundant power supplies, respectively.

Other problems are more complicated. For example, the system crashes might require a new electronic component or an update to the Windows operating system. In the best case, the fix will only take a few hours, but if the issue is not clear or the replacements are not available, the fix could take days or even weeks.

To keep recording, people use the stand-by spare NVR to take over the camera traffic from the failed one (Figure 2). This is referred to as "N+1 NVR redundancy," which makes the newly recorded video available.

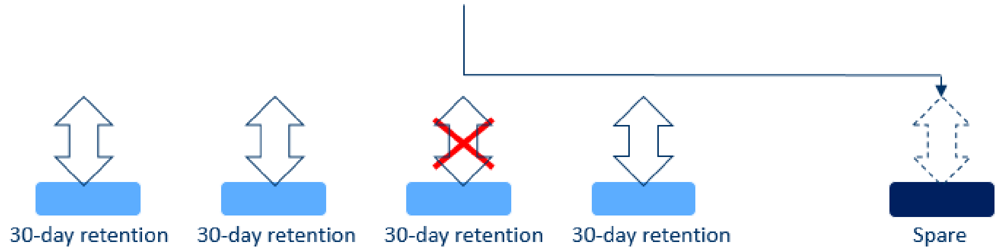

Figure 2 The N+1 NVR Redundancy

Yet a thorny problem remains: the video that was recorded previously (e.g., 30-day retention) in the failed NVR is now unavailable. For example, if a NVR fails to boot up, all the data associated with it is lost. For an urgent event like the Boston Marathon bombing, where retrieving the video evidence is the key to identifying the perpetrators quickly, losing the video may mean that law enforcement cannot identify the people who are responsible for the crime.

For the enterprise high-availability solution, every single byte of data is reachable via redundant paths (Figure 3). So, if one path fails, there is an alternative path to reach the data.

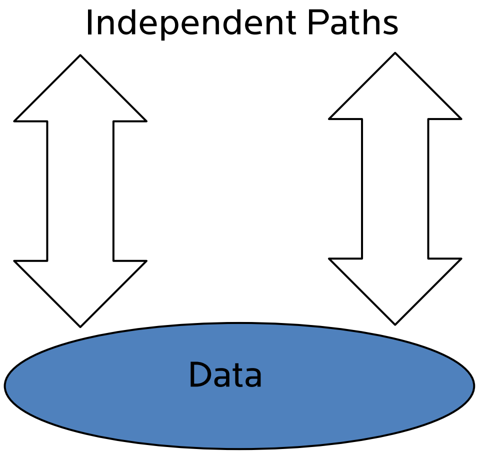

Figure 3 Dual Independent Paths in Enterprise High Availability Solution

The Rasilient PixelStor has redundant paths to reach the data because both controllers can reach every disk drive, whether the drive is in the main chassis or a SAS expansion chassis. So, if one storage controller fails to boot up, there is another one to reach the data. Please see the detailed discussion in Section 6.

4 Are There Too Many NVRs to Manage?

Managing NVRs in distributed geographical locations can be daunting. Each site might have its unique environmental challenges, e.g., limited closet space, weak air flow with temperature concern, excessive dust, and so on. In addition, the travel time to the sites could be hours but still necessary for tasks as simple as checking alerts or replacing a failed disk drive. Going to multiple sites adds a lot of cost.

The cybersecurity concern also blocks some sites from connecting to the Internet, whereas others only have limited access. Therefore, remote management is impossible and traveling to the sites becomes unavoidable but problematic. Consider the difficulty and expense of sending a technician to a faraway site to solve an urgent problem!

As a result, we are seeing more centralized solutions from combined regional centers to a single global site. This allows easier management under one roof and a single cybersecurity protection zone.

In the centralized environment, the distributed NVR architecture starts to lose its value. The rigid pre-provisioning of CPU and capacity is not appropriate or useful. If a central site piles up 4 to 10+ NVRs in the racks, it is time to consider an alternative architecture – one that balances CPU and storage capacity more effectively and has greater scalability, higher density, and lower power consumption.

How to manage growth?

In large deployments, it is unavoidable to add more cameras, increase resolution, and introduce new features as the infrastructure grows.

As a pre-provisioned appliance, the NVR has little room to grow. Suppose the original NVR is provisioned for 64 HD cameras with appropriate CPU and capacity. When a new exciting facial-recognition feature arrives, the customer might need to replace the whole system due to limited CPU power.

Similarly, increasing the resolution of just a few cameras can push the CPU beyond its provisioning and also require an adjustment to the retention period. Again, the NVR falls short.

5 Are You Over Provisioning?

Due to the rigid pre-provisioning, the NVR is usually configured with a large safety margin because once it is wrong, there is little room to adjust.

Over provisioning is one of the major hidden costs in many deployments. Today, we see variations of 25% or more in storage capacity or computing requirements from the configurators of different VMS vendors. In other words, you can pay 25% more than you need.

With a scalable solution in computing and storage, there is no need to over provision. It is safe to use a more realistic estimate.

6 Surveillance-Defined Architecture

The Surveillance-Defined Architecture is focused on the surveillance requirements, exemplified by Rasilient's NFD series of products. See Figure 4.

Rasilient’s NFD offers a combined computing and storage architecture to form a complete surveillance solution. The architecture scales computing and storage independently to reduce the cost.

All servers are consolidated into the ApplianceStor 80/85, which is a four-node blade server in 2U form factor. The AS80/85 provides computing for recording, management, and/or analytics. It scales out computing by adding more blades.

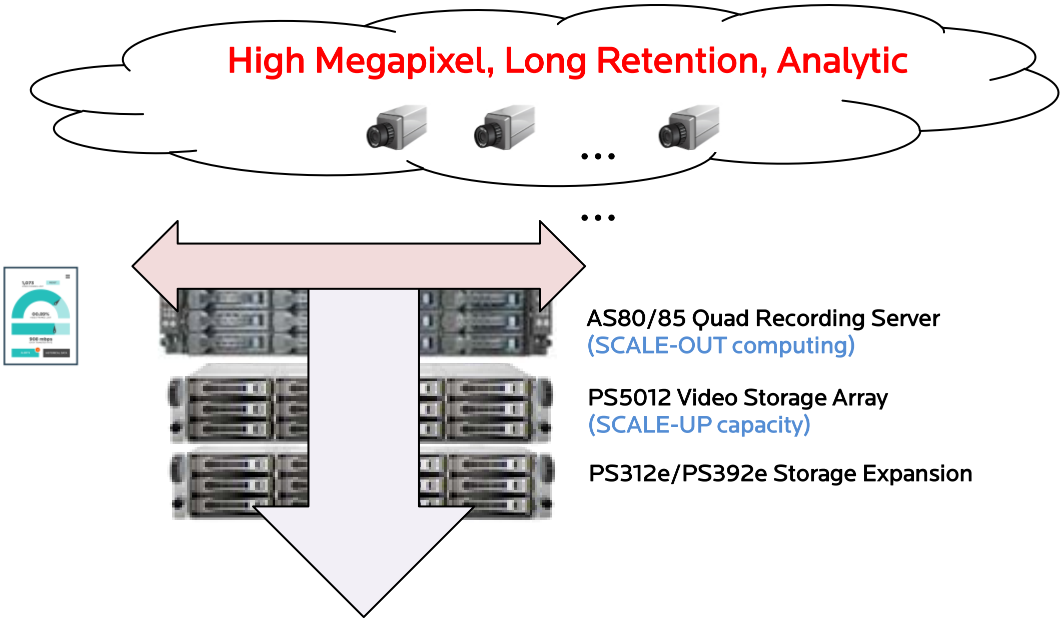
High Megapixel, Long Retention, Analytic
...
...
AS80/85 Quad Recording Server
(SCALE-OUT computing)
PS5012 Video Storage Array
(SCALE-UP capacity)
PS312e/PS392e Storage Expansion

Figure 4 Rasilient Surveillance-Defined Architecture

The PixelStor 5000 is an IPSAN storage array. Its storage controller is optimized only for surveillance traffic. The patented VAN cache algorithm and HyperCache technologies are camera-aware and designed for surveillance. The operation is extremely efficient and capable of recording thousands of video streams with a single Intel processor.

The solution strives to record the highest number of video streams on the same hardware, reducing the cost per stream, while ensuring no loss of video frames.

As a result, Rasilient's exceptional surveillance performance is exhibited in the Milestone/Rasilient certification*. From Mike Tarras, Milestone's Solutions Integration Engineer:

"During our test, the NFD systems integrated with Milestone XProtect VMS did not drop a single frame of video over the entire test period, while supporting 1,000 high definition video surveillance cameras streaming 25FPS at 2Mbps. Additionally, we intentionally triggered a disk failure to invoke a RAID rebuild and the Rasilient NFD system still did not drop a single frame of video. We have never seen this level of performance and reliability in all our prior partner certification tests."

The PixelStor 5000 scales up to petabyte storage capacity by daisy-chaining expansion chassis for the so-called North-South traffic. Each node is fully redundant with dual paths to the same data. The daisy-chain SAS expansion is fast, simple, reliable and low-cost. Rasilient's high-density (5U and 92 drives) PS392e expansion provides further efficiency.

Rasilient's ZM technologies take into account the surveillance operation model, where the service staff is off-site, and reactive drive replacement is a hassle with the risk of losing data. The ZM technologies feature the advanced drive online cloning, a proactive way to replace a drive before it would have been declared faulty. This eliminates the long RAID rebuild window, which is especially important for large-capacity drives.

Rasilient's NFD series of products is certified with major VMS and pre-installed for customer deployments.

For surveillance visibility, Rasilient's NFDMeter provides the frame drop information (Figure 5). This level of visibility along with the patented BusyPlot from the Rasilient PixelStor ensures that the system is installed properly, and any changes to the system in the future will not compromise the operation.

*https://www.milestonesys.com/contentassets/ad960c69177c4fbe9658408cbc250d28/rasilient.nfd.final_1221.pdf

The visibility is not only for real-time, but also long-term (24 hours), where the performance visibility is recorded and the moment of change can be captured. This is especially effective for dynamic camera traffic with varying lighting conditions.

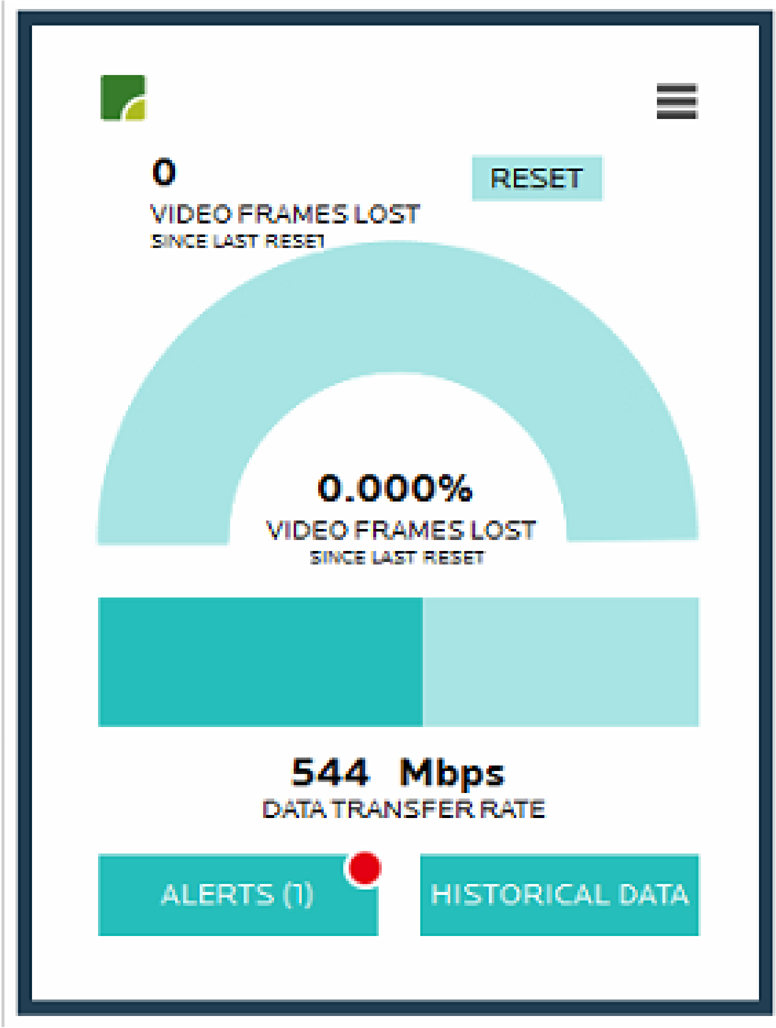

0
VIDEO FRAMES LOST
SINCE LAST RESET

0.000%
VIDEO FRAMES LOST
SINCE LAST RESET

544 Mbps
DATA TRANSFER RATE

Figure 5 Video Frame Drop Visibility

7 Summary

The NVR architecture is effective for small deployments, for distributed geographical locations deployed gradually, and for sites with little future change. It is not well-suited for large deployments that are mission critical. Its limitations include data availability for immediate forensic analysis, increasing management overhead across distributed sites, and limited ability to scale or change features, cameras and configurations.

The Rasilient's Surveillance-Defined Architecture is a solution platform supporting the most efficient scale-out in computing and scale-up in storage capacity. With the NFDMeter visibility, patented storage BusyPlot, VAN, ZM, HyperCache and associated best practices, the solution allows large deployments to grow gracefully for years to come.