Eliminating Bottlenecks in NVR and IT Storage – an Integrated Rasilient-Bosch Solution

Sean Chang
CEO

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

1 Video Surveillance Has Its Own Requirements

By leveraging IT technologies, video surveillance solutions have been quickly advancing from analog to IP over the last decade. However, video surveillance has specific requirements that are very different from IT – video surveillance is mainly dealing with video, not data!

Video has strict timing requirements. Any latency could potentially cause frames to drop or recording gaps. IT products (designed for data) are prone to dropping video. Traditional Network Video Recorders (NVR) using Microsoft Windows have issues of video loss due to Windows NTFS fragmentation, metadata processing, anti-virus processing, firewall operation, and storage related operations.

This white paper describes an integrated Rasilient and Bosch VMS solution, where the combined solution eliminates the key bottlenecks typically seen in traditional NVR and IT storage.

The solution is ideal for mission-critical, large-scale deployments like cities, hospitals, transportation, and other critical infrastructure where losing video is unacceptable due to liability and high-protection requirements.

IT products miss the following video-surveillance characteristics:

- **Write focus.** The data from video-surveillance cameras is mainly write traffic (typically more than 95% of the time). This is in contrast to general storage, which is normally 70% read and 30% write.

- **Large blocks for video streams.** Each video stream is sequential in nature and written to files that are generally large, and the disks are formatted into large blocks (typically 64KB). General storage is different because the file sizes are small and the hard disks are formatted into small blocks (typically 4KB).

- **Multi-camera video streams and the need for low-cost-per-camera video streams.** Due to increasing demand for more and more cameras, the total number of camera video streams can grow to hundreds and even thousands. As each file is written to a different physical location

on the disks, the disk heads jump around to write to all of the files concurrently. As a result, the combined pattern becomes random even though each camera's video stream is sequential in nature. The typical general-purpose storage system performs poorly in this environment. Because of this problem, general-purpose storage products become very difficult to scale with more cameras, higher-resolution cameras, and longer retention periods. This results in high-cost-per-camera video streams.

- **Scale more in storage capacity than computing**. The growth in surveillance data is caused by more cameras, higher resolutions, and longer retention times. The architecture needs to scale storage capacity efficiently to accommodate so-called North-South traffic. This is different from the data centers Hyper-Converged and Scale-Out Cluster architecture which scales more in computing. Each node requires powerful CPU and memory. The internode communication produces so-called East-West traffic which demands sophisticated network infrastructure like 10GE and/or Infiniband switches.

- **Dynamic camera traffic.** Several factors can change the camera traffic dramatically. One is the motion-activated traffic, and the other is the lighting condition. The motion impact is easily seen, and the installer normally pays attention. The lighting condition, however, can be overlooked. The fact is that the camera is very sensitive to light, especially under low light, where the noise level is high and the video compression is less effective. This impacts the recording bandwidth and can break the system provisioning.

- **Visibility for video loss and whether storage is the guilty party.** For video traffic, the key question is whether the video frames are saved into the storage system. In an IT environment, if the storage system takes a long time to write something to a disk, the storage system can signal a “please wait” all the way back to the desktop computer that issued the write and the person simply waits a few seconds longer for the data to be written to the storage product. In a video-surveillance environment, data is written to the disk from a camera. The cameras will output video frames at up to 30 frames per second. If the storage cannot write the video to the disk fast enough, video frames will be dropped. Dropped video frames are common with general-purpose storage products.

2 Bosch VMS Direct Camera-to-Storage Architecture

The Bosch VMS Distributed Network Recording Solution is simple, reliable and scalable solution for network surveillance systems. Video recording is facilitated via iSCSI protocol straight from Bosch cameras to iSCSI storage. This unique combination offloads the video file-to-block conversion process, thereby offering an unmatched camera-per-server support ratio and performance as compared to traditional NVR architectures.

As the data-path goes straight from the cameras to the storage, it eliminates the biggest problem with the Windows-based NVR, where Windows' associated NTFS data fragmentation, metadata overhead, anti-virus, firewall, new patches, and so on are all contributing to video loss.

Figure 1 Bosch VMS (BVMS) has direct-camera-to-storage architecture, which by-pass NVR, its OS, and file system

Bosch VMS further adds intelligence to this architecture by introducing the concept of "storage virtualization." Bosch VMS intelligently directs and balances video traffic across available iSCSI storage to ensure optimal performance. In addition, this intelligence is extended by redirecting video traffic in the event of storage device failures, ensuring continuous availability of recording.

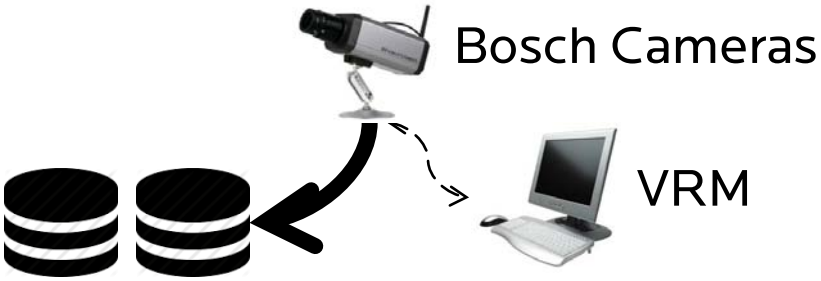

Figure 2 Bosch VMS, running on server, allocates the storage for cameras periodically. The camera records directly to assigned space

3 Rasilient Surveillance-Defined Architecture

Rasilient's Surveillance-Defined Architecture is focused on the surveillance requirements, exemplified by Rasilient's NFD series of products. (See Figure 3)

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. Unlike NVR, Rasilient's NFD has the architecture that scales computing and storage independently to reduce the cost. This matches well with "out-of-band" Bosch VMS that requires much less computing resources.

Combined with Bosch VMS direct camera-to-storage architecture, Rasilient's exceptional surveillance performance is exhibited in the tests and certification.

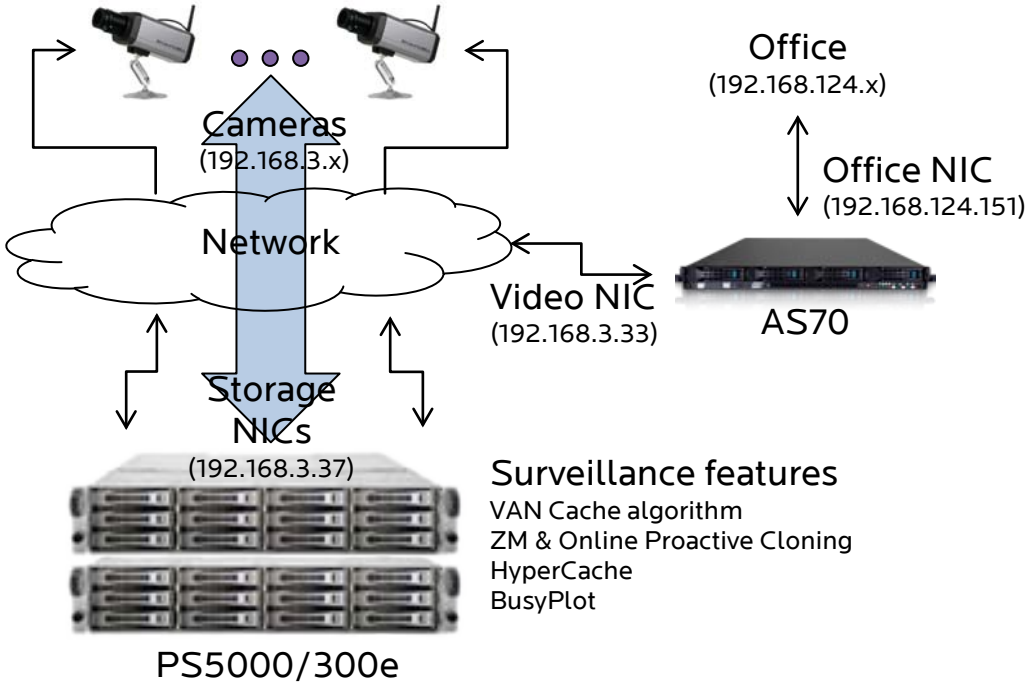

Figure 3 Surveillance-Defined Architecture with Independent Computing and Storage Scaling

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 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. 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.

For surveillance visibility, Rasilient's patented BusyPlot (from the Rasilient PixelStor) provides information as to whether the storage is “busy” to ensure that the system is installed properly and any changes to the system in the future will not compromise the operation.

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.

Depending on the scale, the deployment can either use a single server for BVMS or consolidate BVMS, analytics, and redundancy into the ApplianceStor 80/85, which is a four-node blade server in 2U form factor. The AS80/85 scales out computing by adding more blades.

4 Benefits

For surveillance, the bottom line is that all video needs to be captured and protected throughout the lifetime of the system. By taking the recording

server and its Windows out of the way, BVMS has simplified the data path, thus removing both potential bottlenecks.

With Rasilient's patented storage technologies optimized for surveillance traffic, we see unmatched benefits from the combined solution.

- **Conformance to regulation.** If there is regulation for video surveillance, the owner's key risk is not having recorded video. This combined solution eliminates the bottlenecks for video loss ensuring the conformance to any regulation.
- **Stand-up in court.** To use recorded video in court, the quality has to be "forensic-grade." Too many times the video is thrown out because of incompleteness, poor quality, and missing key video information. The NFD is designed to be "forensic-grade."
- **Cost effectiveness.** The "out-of-band" Bosch VMS and Rasilient's Surveillance-Defined Architecture is a solution platform using the least computing power while supporting the highest number of cameras.
- **Minimal operation expense**. With the patented storage BusyPlot, VAN, ZM, HyperCache and associated best practices, the combined solution not only scales easily, but also allows large deployments to grow gracefully. The ZM, proactive maintenance, and system visibility dramatically reduce operating expenses.
- **Peace of mind.** The combined solution provides a fully-redundant option from computing to storage. Any recorded data has two independent paths to access. So, the data is available immediately when one of the paths fails. In comparison to NVR, the associated data is unavailable once the NVR is down.
- **Tested and certified.** Please refer to Rasilient test report for details.