I D C A N A L Y S T C O N N E C T I O N

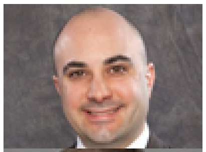

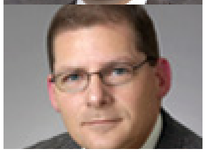

Christopher Chute
Research Manager, Worldwide Digital Imaging Solutions Group

David Reinsel
Group Vice President, Storage, Semiconductors, GRC, and Pricing

Network Video Surveillance: Addressing Storage Challenges

April 2012

IDC sees network video surveillance adoption continuing to grow across all regions, driven by the need to secure municipal and government facilities, take advantage of network investments, improve ROI in retail locations, and a broader push as a business tool that can drive Big Data analytics. Network cameras are well on their way to supplanting analog models, and, furthermore, new installations are trending to favor larger surveillance implementations. Network surveillance is very scalable, and with camera costs declining, the ability to add surveillance capacity becomes more compelling. Declines in storage costs and continued improvement in network camera technology are also working in the favor of network surveillance. Maturity in designing megapixel cameras and the ability to remotely correct focusing problems, coupled with increased adoption of hi-definition and ability to capture a flexible number of frame rates points towards increased adoption.

The following questions were posed by EMC to Christopher Chute and David Reinsel of IDC, on behalf of EMC's customers.

Q. Network video surveillance is a rapidly growing capability being adopted by a variety of industries. Robust, intelligent storage solutions are a main driver for uptake. How does installation size and type impact equipment selection and requirements?

A. Network video surveillance got its start in the mid-90s, and has grown in tandem with developments in the network space. Many content-capture markets are critically driven by cost-effective, robust storage requirements, and network video surveillance is no different.

The benefits of these new network video surveillance technologies are many: increases in video resolution and dynamic range put network camera image quality on-par with the best analog models. This, coupled with file-based storage (as opposed to cumbersome linear-based tape storage) provides administrators the ability to capture and store content with much longer retention rates than before.

Furthermore, IDC sees network video surveillance being more quickly adopted in low-hanging fruit industries like municipalities, critical infrastructure, and transportation than in SMBs. Why? This is due to many factors, including upgrades to public facilities including broad network investments, often addressing public safety, and private funding for larger projects that drive significant investor returns (e.g. stadiums) creates the ability to develop world-class threat and shrinkage detection.

These larger deployments often feature camera installations in the hundreds, frequently in a variety of both hardware and network configurations. While it is easy to take advantage of this scalability, it often results in back-end management content management issues.

Q. Since larger installations are driving the ability to capture, share and analyze more surveillance footage, as well as an increasing number of unique video streams, does that make video surveillance an emerging driver for Big Data?

A. Big Data applications can exist in a variety of formats based on IDC's definition. IDC defines Big Data by virtue of 4 "V's": Velocity, Volume, Variety, and Value. Applications where large, unstructured files (i.e., Video files) or a large number of video files are being captured or streamed can be deemed Big Data if they meet these definition thresholds.

While not all Video surveillance applications fit within IDC's Big Data definition, many do (and increasingly so). These Big Data video surveillance applications drive specific storage infrastructure and platform requirements that go beyond traditional, general-purpose storage.

What also increases the likelihood of video surveillance being or becoming a Big Data application are the analytics that companies are clamoring to do. Besides leveraging the metadata associated with video surveillance files (e.g., time of day, location, weather, incident reporting, etc.), companies are developing algorithms to search video content based on patterns and correlations within the file itself.

This type of workload is bound to place large performance requirements on storage systems, as well as qualify it as a Big Data application. Of course, this does not even include the real-time analytics that companies desire to do with live video streams (e.g., facial recognition, group/mob identification, human traffic patterning, etc.).

Q. After deploying a large video surveillance system, what are the user pain points?

A. There are a number of challenging dynamics associated with a large video surveillance deployment. Often times, large-scale deployments result in an abundance of content that must be archived, managed and utilized by teams only partially equipped to handle the task. Network and security responsibilities often fall between two separate chains of command, with network administrators being unfamiliar with threat detection and prevention methods, and security personnel being unfamiliar with IT. Furthermore, securing the content itself from accidental or malicious deletion is a mission-critical priority.

From a storage perspective, the challenges are many: managing multiple, long streams of content, ensuring that enough storage is installed (or can be installed) so that video streams can be recorded continuously without interruption, and ensuring the quality of the video stream itself.

Invariably, a company installing video surveillance cameras doesn't just stop at one camera, or with a handful of cameras, but instead, increases the number of cameras (video streams) over time. A storage system can quickly become overloaded with such increased workload. Even though a video stream is one long stream of data, it is written across various disk drives of storage; hence the heads within the hard drives are moving constantly to record all the various streams. If the storage system cannot keep up, then frames will be dropped.

Depending on the retention policy, video streams (which increasingly are higher resolution), can consume storage rapidly, which means IT managers must be able to scale (install new capacity) their storage infrastructure quickly and non-disruptively. In addition, the storage system should be optimized around 'write' since it is writing the vast majority of the time.

Finally, managing all this unstructured content, likely coming in from multiple locations, requires a file system that enables the management of files in an efficient and global way. Ensuring that video content can be found, retrieved, and viewed from multiple locations is quickly becoming a primary requirement. In addition, all levels of security (who can watch/access the video streams and when) is an important requirement to ensure the integrity and privacy of the footage.

Q. What type of data storage system should be selected to effectively address the management of video surveillance content?

A. Typically, scale-out NAS solutions are the beginning of a large list of requirement for any storage system designed to manage video surveillance footage. Systems designed to scale efficiently and without disrupting operations are preferred to keep 24x7 video stream capture viable.

Storage system reliability should be very high on the list of requirements, as well. Losing one or multiple video surveillance cameras doesn't impact the other cameras; however, losing the storage systems takes out every camera video stream sent to it. Hence, protecting from hardware failures via a redundant data preservation strategy is a must. In addition, to ease file management, a global name space is preferred so that files are presented in a single context and can be managed from multiple locations.

Finally, a layer of advanced caching hardware and software is desired to ensure the integrity and quality of the stored video streams. A caching layer can help to manage delays due to periods of high I/O, especially if the storage system handles multiple workloads other than just video streams.

Q. Looking to the future, what factors should be considered when devising an effective storage strategy for a video surveillance implementation?

A. Minimizing change through the use of stable software and hardware providers helps to create a cost-effective, stable environment that provides a foundation for future evolution of the infrastructure. As additional video surveillance cameras (and data streams) are installed, scaling scalable storage infrastructure is vital to ensure non-interruption of critical video feeds.

Eventually, Big Data analytics, including real-time analysis will become increasingly critical as the volume of unstructured data grows and the desire to extract value from the volumes of data is enabled via big data analytics and other sophisticated software designed to search video files.

Efficient and scalable storage solutions, married with capable file management systems that can assist in such value extraction initiatives lays the groundwork for optimizing capital and operational expenditures, which should help any video surveillance solution become a valued asset within any organization.

Finally, given the sensitive nature of video surveillance footage, solutions must have security and privacy controls not only to ensure just the right individuals have access to the data, but also to minimize accidental content losses.

By ensuring a storage system is designed properly for video surveillance environments and workloads, companies will be able to establish strong and effective video surveillance practices for today, and an environment for easy and efficient growth for the future.

ABOUT THESE ANALYSTS

Christopher Chute is a Research Manager with IDC's Worldwide Digital Imaging Solutions Group. Mr. Chute focuses specifically on the digital still camera, camcorder, camera phone, and network surveillance camera categories, as well as consumer printing and other image sharing technologies. He conducts forecasting, product analysis, and consumer usage identification for these markets, through a variety of supply and demand-side studies.

Dave Reinsel is group vice president of IDC's storage, semiconductor, GRC infrastructure, and pricing research groups. Mr. Reinsel and his team of analysts provide timely and industry-leading insight and analysis for IT professionals, investors, resellers, distributors, and manufacturers worldwide. His research teams are responsible for delivering annual and quarterly forecasts and analyses on the storage hardware, storage software, semiconductor, and GRC infrastructure markets, as well as strategic pricing evaluation services for various IT products.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

COPYRIGHT AND RESTRICTIONS

For more information on IDC visit www.idc.com. For more information on IDC GMS visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA   P.508.872.8200   F.508.935.4015   www.idc.com