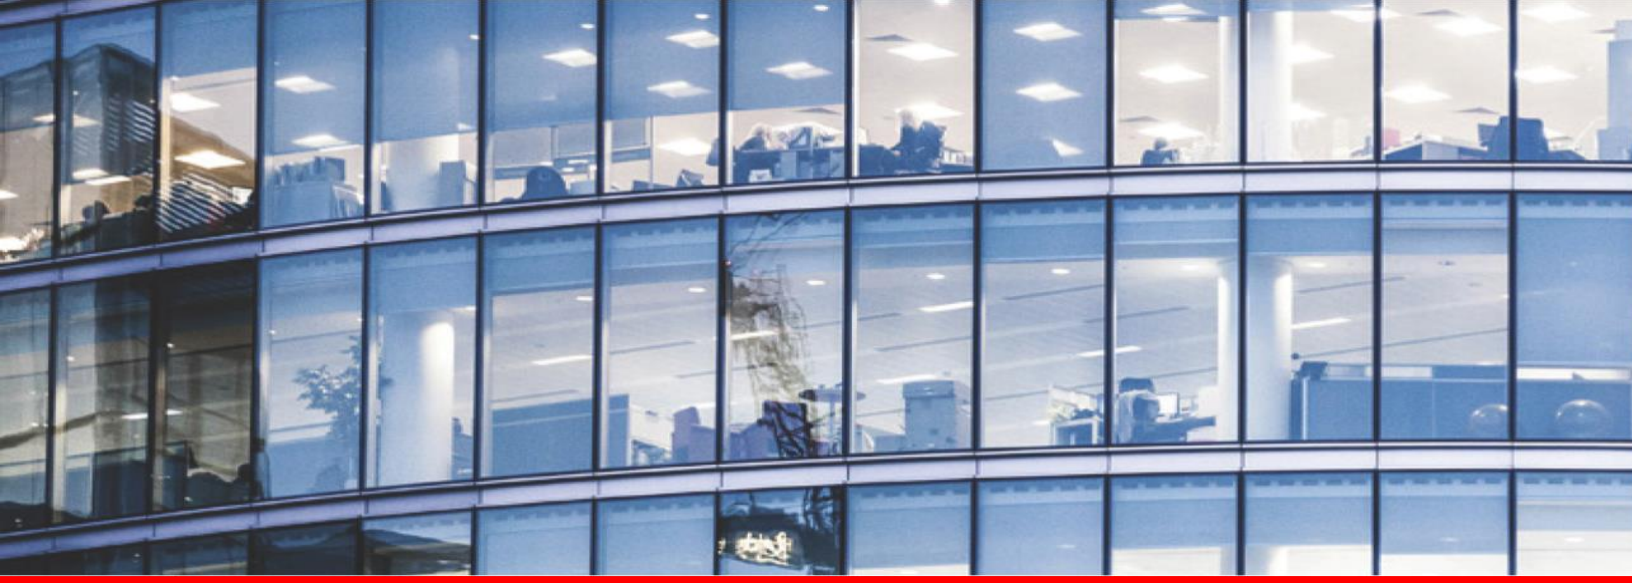

WHITE PAPER

SanDisk® Flash Storage: High Performance Low-Cost Primary Storage for Video Surveillance

SanDisk
a Western Digital brand

Western Digital Technologies, Inc.
951 SanDisk Drive, Milpitas, CA 95035

www.SanDisk.com

Table of Contents

Introduction

In a complex world, video surveillance is playing an increasingly critical role, providing reliable intelligence, improving security, and helping to detect and solve crimes. The promise of video surveillance and its success to date means that it will only be deployed in larger volumes in the future, resulting in a need for more cameras, better quality, and more primary storage volume for video servers. With the vast amounts of video data being collected, the need for greater amounts of timely analytics is increasing as well.

Unfortunately, traditional video surveillance server technology based on hard disk drives (HDDs) is beset by numerous challenges that directly thwart these goals. Limited HDD throughput often results in a loss of critical quality. Even with compression, limited throughput also impacts the ability to do timely analytics, as HDD-based video servers often lack the ability to serve simultaneous reads and writes without dropping frames. Moreover, recovering a RAID volume from an HDD failure can take considerable time and reduce throughput further, causing unacceptable frame losses that compromise video integrity and put analytics at risk. With HDD capacity and throughput limited, the number of required video servers is growing steadily, raising concerns around server sprawl and both capital and operating expenses.

Substituting flash for primary storage can eliminate many of the shortcomings of traditional HDD-based video surveillance solutions. While flash is often used in cameras, and also as efficient server boot devices, SanDisk® has evaluated high-throughput flash storage as primary storage for video surveillance servers. Flash storage can provide both dramatically higher capacity and sustained throughput, and with them higher density and scalability for video servers. This higher capacity and throughput can be used to store video at higher frame rates and resolution, or it can be used to deliver significant server consolidation. Dramatic reductions in total cost of ownership (TCO), total cost of acquisition (TCA), and operational expenses (OpEx) are the direct result. At the same time, flash storage provides predictable quality of service (QoS), eliminates frame loss, and offers shorter RAID recovery times (with no loss in video integrity). By providing throughput headroom for simultaneous read and write access while recording, flash storage also allows for near-instant access to video data, enabling emerging real-time analytics while video data still has predictive value.

Traditionally, cost and media endurance have been viewed as concerns for using solid state devices (SSDs) in video servers. Technology advances and SanDisk's innovations have now eliminated these objections, making flash storage a competitive and compelling solution. To evaluate the role of flash storage in video surveillance environments, SanDisk has partnered with Milestone Systems—a global leader in video management software—to perform testing and certification. Testing has shown at least five-fold better throughput for flash than HDDs, for a 150-camera system. Moreover, projecting these results to a 1500-camera deployment demonstrates substantially lower three-year TCO than an HDD-based system.

Enhancing Video Surveillance with Flash Primary Storage

Video is increasingly a key witness to events. Now considered essential for physical security, organizations from businesses, airports, and schools to municipalities are deploying video surveillance with ever-increasing numbers of cameras. In particular, video can help provide a disinterested perspective in analyzing complex situations, ideally allowing law enforcement to act on critical actionable intelligence.

The demands on video surveillance systems will only increase. Going forward, video surveillance systems must respond with key improvements, including:

- **Higher quality and better throughput**. The number of deployed cameras will only grow, as organizations realize the importance of covering more perspectives. Quality must also be improved for valuable analytics applications such as facial recognition and license plate readers. For surveillance to contribute, however, recorded images and streams must contain sufficient quality and detail. Higher-resolution high-definition (HD) and 4K cameras are already available, and they will introduce new throughput challenges to already taxed surveillance systems, even as they improve potential video quality.
- **Larger primary storage capacity**. With more cameras and higher resolutions, the sheer amount of video content will grow dramatically, challenging existing primary storage solutions. To date, solutions have been driven largely by the throughput and capacity limitations of spinning media. These limitations in turn drive the number of required controllers and servers to undesirable levels, contributing to higher complexity and server sprawl, and higher capital and operating expenses.
- **Improved analytics**. Organizations are increasingly clamoring for the ability to corroborate multiple sources and react more quickly to detected events, while opportunity and intervention windows still exist. Today most recorded video surveillance data is never viewed again, but that situation is expected to change. As with other data, organizations can benefit from faster time to value from recorded video. With dramatic video data growth, the value of faster and real-time analytics becomes increasingly paramount. For instance, the ability to alert and correlate disparate facial recognition events could be highly valuable in detecting an impending event.

Comparing Flash Storage with Traditional HDD-based Video Surveillance Solutions

Traditional video surveillance solutions have served well, but they are running up against very real physical limitations. Perhaps nowhere are these challenges more pronounced than in the key areas of storage system capacity, throughput, and response time, where flash storage has significant advantages (Table 1). HDD throughput constraints can result in frame loss when throughput demands outstrip the spinning media's ability to respond. Because of the nature of video, it can be very difficult to detect these losses after the fact.

Frame loss can mean that a critical event might well be eliminated from the video stream, impeding analytics from detecting key events. The shortcomings of HDD-based video servers have been met with limited analytics being done in the camera. Unfortunately, this approach misses the opportunity for more advanced and comprehensive computational analytics, and the ability to correlate events beyond a single camera.

Table 1. Comparing HDD and flash storage for video surveillance.

	Technology	Capacity	IOPS	Response Time	Random I/O Pattern Ability*	Simultaneous Read/Write?*
Hard disk drives (HDDs)	Spinning platters, moving heads	Up to 1.8TB (10K RPM drives) Up to 600 GB (15K RPM drives)	~250 IOPS	~5ms	Poor	No
Solid state devices (SSDs)	Solid state	Up to 4TB	~50,000 IOPS	~0.05ms	Excellent	Yes

* In video server applications

The properties of HDDs are a poor match for the rigors of video surveillance. With multiple cameras writing to multiple files, fully random I/O patterns are typical in video servers. Random writes (for recording) are often interspersed with random reads (for archival or analytics), which can quickly result in bottlenecks given limited HDD throughput. These kinds of write patterns can cause significant head flutter, significantly reducing the practical lifespan of HDDs and causing early and frequent failures. In fact, HDDs in a video surveillance environment typically incur at least a 5% annual failure rate.

Capacity is another significant challenge. To achieve sufficient throughput, high-speed 10K or 15K RPM drives are typically deployed on video servers. Unfortunately, these faster drives are only available at lower capacities, making it difficult to accommodate sufficient primary storage volume. As a result, larger and larger numbers of RAID controllers and video servers are needed as more cameras are deployed. Not only does this drive higher acquisition costs, but maintenance costs and OpEx are accelerated as well.

Enhancing Video Surveillance with SanDisk Flash for Primary Storage

SanDisk flash storage presents a unique opportunity to evolve and consolidate servers for video surveillance. Higher primary storage throughput and capacity results directly in the ability to deploy fewer more dense servers to support a given number of cameras. With rapidly evolving network technology and increasing available network bandwidth, video server clusters will be able to grow and scale more effectively, with higher bandwidth to individual servers aiding server consolidation efforts. Key flash storage advances promote:

- **Truly scalable video surveillance infrastructure.** Unlike HDDs, flash storage offers linear density and throughput scaling for video surveillance. In fact, SanDisk has found that flash can provide five-fold throughput improvements over HDDs, without corresponding capacity constraints of 10K or 15K RPM HDDs. The dramatically improved throughput eliminates frame loss and provides the predictable QoS that is so essential for security applications, where missing information can result in significant loss. Flash can sustain high throughput with simultaneous ingress and egress of video data. With flash, random writes are served without read performance issues or media reliability issues, and mixed workloads are served at high performance.

- **Reliable and resilient media.** The high failure rates of HDDs make RAID recovery events highly likely. Low HDD throughput can cause RAID recovery to take days, during which the array is vulnerable to additional unprotected failures, and extremely likely to drop frames from active cameras. In contrast, SSDs are more reliable in video server settings as they are not subject to the physical wear and tear of rotating platters and moving heads. The higher throughput also allows RAID configurations based on SSDs to recover much more quickly in the event a module should fail. Higher throughput also means that recovering RAID arrays can continue to record video without frame loss.
- **Better and more immediate analytics**. Edge (camera) to core (site, geography) analytics is increasingly key for security. Flash-equipped video servers can provide better results with an aggregated data pool, allowing analytics to be run directly on primary storage. The considerable throughput headroom flash provides allows video recording and analytics to coexist, without worries about causing I/O bottlenecks that can compromise recording and cause frame loss. Flash enables the ability to create a central data pool for improved analytics, and facilitates the correlation of disparate data while analyzing for global security attempts.
- **A compelling cost model**. While HDDs are ultimately constrained by the physical realities of spinning platters and moving heads, SSDs ride the semiconductor technology curve—following the tenants of Moore's Law. As a result, flash advantages will only continue to multiply over HDDs into the future, presenting an increasingly compelling cost model. Flash storage can also be deployed in a number of different architectures, from dedicated flash storage to hybrid tiered storage that utilizes flash as a front end with large-capacity HDDs as back-end archival storage for long-term retention. Hosting more cameras on each video server allows for considerable server consolidation, dramatically improving TCO over HDD-based solutions.

Significant Simplification and Cost Reduction Through Server Consolidation

While low throughput and capacity of HDDs tend to drive the numbers of drives, controllers, and servers upward, SSDs have the opposite effect. By providing up to five times the throughput and substantial improvements in capacity over 10K and 15K RPM HDDs, SSDs can offer up to a five-fold increase in the number of cameras that can be supported on a single server. This server consolidation in turn means that dramatically fewer servers are required to support a given number of cameras. SanDisk has found that SSDs also provide:

- Half the TCO of an equivalent HDD-based system
- Significantly lower TCA than an equivalent HDD-based system
- Less than a third the OpEx of an equivalent HDD-based system

Full TCO analysis demonstrating these results can be found later in this document. These advantages will be particularly important as the number and quality of video streams increases.

Milestone Systems Qualification and Testing

Video management software (VMS) and network video recorders (NVRs) expand traditional surveillance and security system capabilities, providing efficient network based monitoring, transmission, and digital storage of IP video. Milestone's commitment to open platform technology through standards based interoperability gives organizations the freedom to choose the widest range of equipment. Milestone partners with hundreds of companies in the security and IT infrastructure industries to offer products that focus on ease of use, scalability and solving modern end-user challenges.

SanDisk conducted Milestone Technology Partner Certification to evaluate SanDisk flash storage for video surveillance, and to establish performance baselines for in-depth TCO analysis. The Milestone Technology Partner (MTP) Certification program seeks to confirm that server, storage, and network solutions provided by qualified MTP vendors meet the performance benchmarks required to support the Milestone XProtect VMS applications, and to measure the maximum performance available to Milestone customers if they choose to build a solution using certified MTP products. Partners who participate in the certification, such as SanDisk, are valuable members of Milestone's Open Platform eco-system, and have demonstrated a strong and lasting commitment to partnership.

Certification of a commercial off-the-shelf (COTS) server equipped with SanDisk SSDs helps ensure that any surveillance system built using these products in combination with the Milestone XProtect components will be able to record and archive an amount of video consistent with the recommendations of the Milestone Server and Storage Calculator (www.milestonesys.com/support/presales-support/Storage-Calculator/). Certified products have been proven and documented to support the XProtect VMS. Performance expectations, design best practices, and specific configuration details are all documented as a part of the certification. Integrators, and end-users can reference the certification results to determine exactly what size VMS system these products will support.

Milestone XProtect Corporate was the VMS chosen for the certification tests with SanDisk SSD storage. As a part of this testing, the following Optimus® and Lightning® Gen. II SAS SSDs from SanDisk were certified with Milestone XProtect Corporate.

- Optimus Eco™ 2.0TB 6Gbps SAS SSDs
- Lightning Eco™ Gen. II 1.6TB 12Gbps SAS SSDs

Performance of the solution may vary if different XProtect products and/or system components not listed in the test details are included.

Testing and Certification Environment

For certification testing, all functions related to the XProtect surveillance system were supported on a single recording server (Figure 1), including the Database, Management and Recording Server services. The camera simulation targets were generated on the same server. Multiple live viewing stations ran Milestone's XProtect Client software and were connected via Gigabit Ethernet. The recording server was an industry-standard COTS system (Dell™ PowerEdge™ R730 server) configured as follows:

- CPU: Dual Intel® Xeon® processor E5-2697 v3 at 2.6GHz
- System Memory (RAM): 256GB
- Storage Controller: 2x RAID controllers
- SSD Storage: 12 SAS drives, 12 SSD drives
- Software: Microsoft Windows 2012 R2 64bit Operating System

The two Dell PowerEdge RAID Controller (PERC) H730P storage controllers featured:

- 12Gb/s SAS, 6Gb/s SATA
- PCI-Express 3.0 (eight internal ports)
- 2GB NV Flash Backed Cache
- Hardware RAID (RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60)

The first controller was connected to HDDs with a 2-disk RAID 0 volume for the OS and an eight-disk RAID 6 for archive storage. The second controller was connected to SSDs with a two-disk RAID 1 volume and a eight-disk RAID 6 volume.

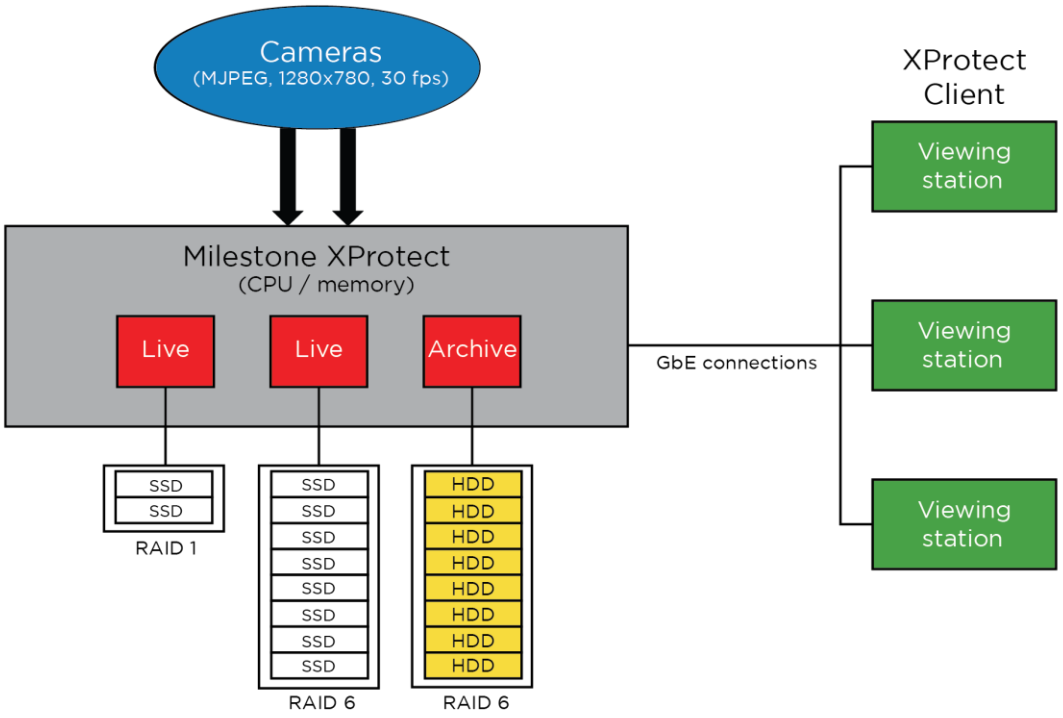

Figure 1. COTS server configuration for SanDisk SSD Milestone certification.

Live Milestone database testing was configured on the following configurations utilizing both Optimus Eco 6Gbps SAS SSDs and Lightning Ascend Gen. II 12 Gbps SAS SSDs:

- RAID 1 configuration of two SSDs
- RAID 6 configuration of eight SSDs

After two hours, all video content was archived to a RAID 6 volume consisting of eight HDDs. A variety of tests were run, culminating with a final run with the SSDs configured in RAID 6 supporting 150 cameras in 1280x720 resolution (Table 2). All of the test scenarios complied easily with the performance criteria of the Milestone Certification program, and no frame loss was detected in any of the tests. The scalable, modular configuration of the XProtect VMS, along with the capacity and performance offered by the SanDisk SSD solution implies that

systems can easily be built to support hundreds to thousands of cameras. Importantly, to stress the system, the test scenario represented constant recording from the 150 cameras, with no Video Motion Detection (VMD) used to minimize video recording throughput or volume. In real-world scenarios individual video streams are likely to be less challenging due to VMD.

Table 2. SanDisk SSDs were tested in RAID 6 with no VMD using MJPEG streams.

Resolution	Frame Rate	Video Compression	Codec	Recording Percentage
1280 x 720	30 fps	30%	MJPEG	100% (no VMD)

Performance Results and Key Findings

The chosen SSD configurations used in the test easily supported the benchmark level of performance when storing the data using MJPEG codecs. These solutions can be used to support XProtect Recording Servers with up to 150 cameras recording at megapixel resolution at 30 frames per second. Larger systems could be built using multiple storage solutions, alternative disk configurations, and multiple Recording Servers. Integrators and end users designing, installing and operating surveillance systems which incorporate these solution components should have complete confidence that the system will record large amounts of megapixel resolution video very reliably.

The solution was able to reliably store large amounts of video without impacting the performance of the Milestone system. SanDisk SAS SSDs are ideal for any system that requires high performance video recording using hundreds of multi-megapixel video streams. SanDisk SAS SSDs and Milestone XProtect VMS combine to create a solution that is highly scalable, high performance, standards based, and easy to use.

Performance results from the maximum performance test were measured, and are provided below. Importantly, consolidating more video streams onto individual video servers does result a commensurate increase in CPU loading. In this case, the subject system was well equipped to deal with the increased load, but system designers should be careful to ensure that sufficient CPU capacity is available to handle the larger number of video streams. Full throughput and latency results for the tests conducted are given in Table 3.

Table 3. Latency and throughput results for SanDisk and Milestone Systems testing (1280 x 720p, MJPEG, 30% compression, 30fps).

	Cameras	Live Video Database Write Latency (ms)	Total Disk Throughput (MBps)	Frame Loss	CPU %
Two Optimus SAS SSDs (RAID 1)	40	84.3	113.7	0	0.65% average, 11% peak
Two Lightning Gen. II SAS SSDs (RAID 1)	48	98.9	143.3	0	
Eight Optimus SAS SSDs (RAID 6)	140	30.0	420.0	0	
Eight Lightning Gen. II SAS SSDs (RAID 6)	150	28.6	448.9	0	

The Milestone sizing calculator states that a 15K RPM HDD supports an average disk throughput of 12.5 MB/s. Importantly, this metric assumes a 90% write 10% read workload mix associated with video server usage in a video surveillance environment. Any RAID based configuration drops this due to mirroring or parity overheads. Based on the results listed above, SanDisk Optimus and Lightning SSDs can easily support throughput of more than five times the throughput of HDDs, even after taking the RAID overhead into consideration.

Data Protection and Reliability

There are significant advantages in terms of data protection, reliability, and power consumption that help make flash an even more compelling solution for video server primary storage. HDDs have some inherent limitations that can adversely affect a solution in terms of data protection. To get a reasonable performance access density (IOPS per GB) relatively high-speed (low-capacity) HDDs must be used in video servers. Unfortunately, a relatively large number of the storage devices are required to meet capacity requirements, in turn requiring multiple RAID controllers. In contrast, the large capacity of SSDs reduces the storage device count significantly. The lower device count means that the RAID controller count drops as well. In turn, the overall solution MTBF is greater because of the removal of so many potential points of failure.

Enterprise based NAND devices also typically have 10-fold better reliability compared to HDDs with respect to the uncorrectable bit error rate (UBER). As a result, there is a 10-fold lower chance of encountering any read errors or any failures occurring during the RAID rebuild process.

RAID Rebuild Under Milestone Load

The throughput limitations of HDD-based systems can also strongly impact video storage in the event of a failed disk. As discussed, the random write patterns of video servers virtually guarantee frequent HDD failure. While most video servers employ RAID 6 for high data availability, the HDD-based RAID rebuild process can take days to accomplish. Moreover, during RAID rebuild, HDD-based systems can suffer throughput limitations that can approach 50 percent, which can result directly in frame loss for live cameras still attempting to record to the video server. Given the length of time typically required for HDD-based RAID systems, the likelihood of a second drive failure is non-trivial, and something that could greatly compromise the ability of the video server to continue recording successfully.

Flash storage provides much higher throughput headroom and only minimally lowers throughput during a RAID rebuild event. In the event of an SSD failure, the system has sufficient processor and throughput headroom to easily accommodate both the RAID rebuild process as well as ongoing video recording. The rebuild process also requires significantly less time than an HDD-based volume, reducing the chances of a secondary device failure.

Three-Year TCO Analysis

To evaluate the cost effectiveness of various systems, SanDisk carried out a three-year TCO analysis comparing a HDD-based system against a flash-based configuration using SanDisk Optimus Eco SSDs. Informed by the Milestone testing, the TCO calculation made conservative assumptions about SSD performance in a video surveillance environment. The analysis demonstrated that the server consolidation from Flash-based video surveillance servers can provide a dramatically lower three-year TCO than a system equipped with HDDs.

Analysis Assumptions

To conduct the TCO analysis, SanDisk modeled a 1500-camera video surveillance system. The TCO modeling was based on the actual performance results obtained in the Milestone certification testing described herein. Identical server configurations were used, with only variation in the number of cameras per server. The TCO calculations were driven by:

- The number of cameras per server
- The number of HDDs or SSDs required to support the needed throughput for the number of cameras
- The number of HDDs or SSDs required to support the needed capacity (defined as two hours of live recording)

TCO modeling was then conducted on the following configurations, scaled out to provide support for 1500 cameras:

- 30 camera per server (HDDs), 50 video servers, eight RAID 10 SAS 15K RPM 300 GB HDDs per server
- 90 cameras per server (SSDs), 17 video servers, five RAID 6 SanDisk 2.0TB Optimus Eco SSDs per server
- 150 cameras per server (SSDs), 10 video servers, six RAID 6 SanDisk 2.0TB Optimus Eco SSDs per server

Total Cost of Ownership

Figure 2 illustrates the TCO comparison of the 30, 90, and 150-camera configurations, with both total cost of acquisition (TCA) and three-year operational expenses (OpEx) shown. The dramatic cost advantages of video server consolidation are easily visible for the SSD-based configurations. The 90-camera-per-server SSD configuration is dramatically less expensive than the 30-camera-per-server HDD configuration, even though the initial cost of acquisition is slightly higher. The 150-camera-per-server SSD configuration represents an easily achievable five-to-one server consolidation. It is both less expensive to acquire, and fully half the cost of the HDD-based system over three years.

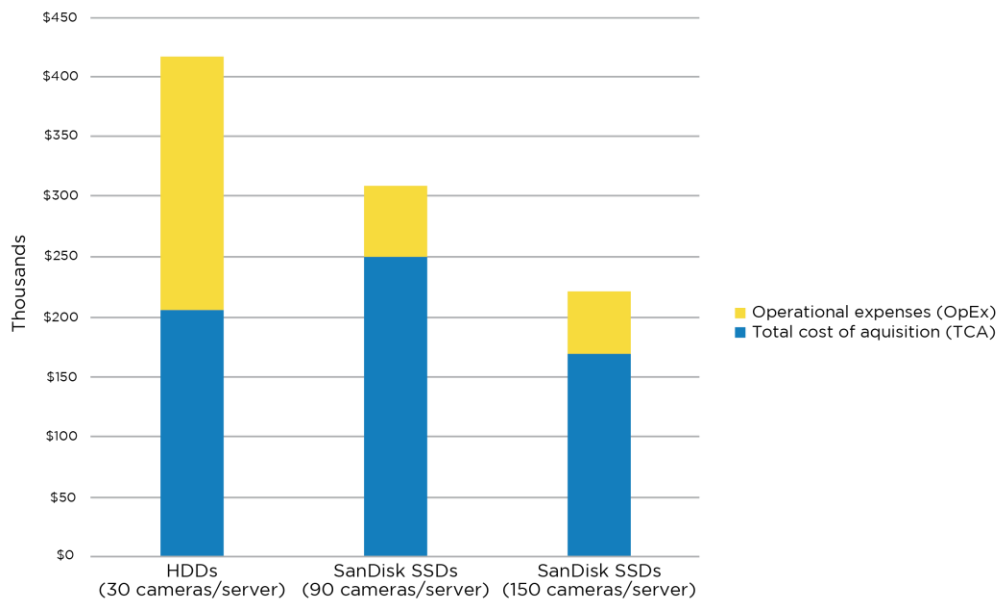

Figure 2. SSD-based video surveillance servers exhibit dramatically better three-year TCO over video servers with HDDs.

Total Cost of Acquisition

Figure 3 details TCA for the three video server configurations. Though the relative storage costs are higher for the SSD-based configurations, the server-related costs are dramatically lower due to consolidation. The overall TCA for the 150 camera-per-server SSD-based configurations is again significantly lower than the HDD-based configurations.

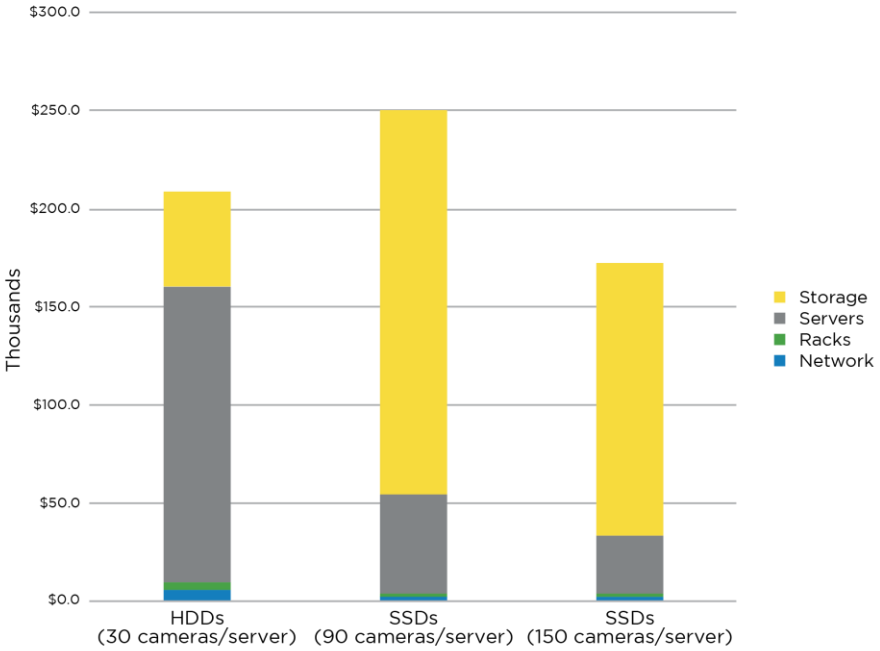

Figure 3. TCA is lowest for 150 camera-per-server SSD-based configurations.

Three-Year OpEx

Figure 4 illustrates the OpEx results from the three-year TCO calculation with considerable savings over a three-year period. Again, the network and rack costs for the two SSD-based configurations are dramatically lower than the HDD-based configuration. The smaller numbers of more dense SSD-based servers occupy significantly less rack space, and require significantly less power to operate.

The power savings are particularly pronounced for the SSD-based solutions. Not only do the individual SSDs require less power to operate, they generate less heat and vibration than the moving parts of HDDs. Fewer larger storage devices coupled with fewer storage controllers and fewer required video servers helps to drive simplicity and keep power and cooling costs down.

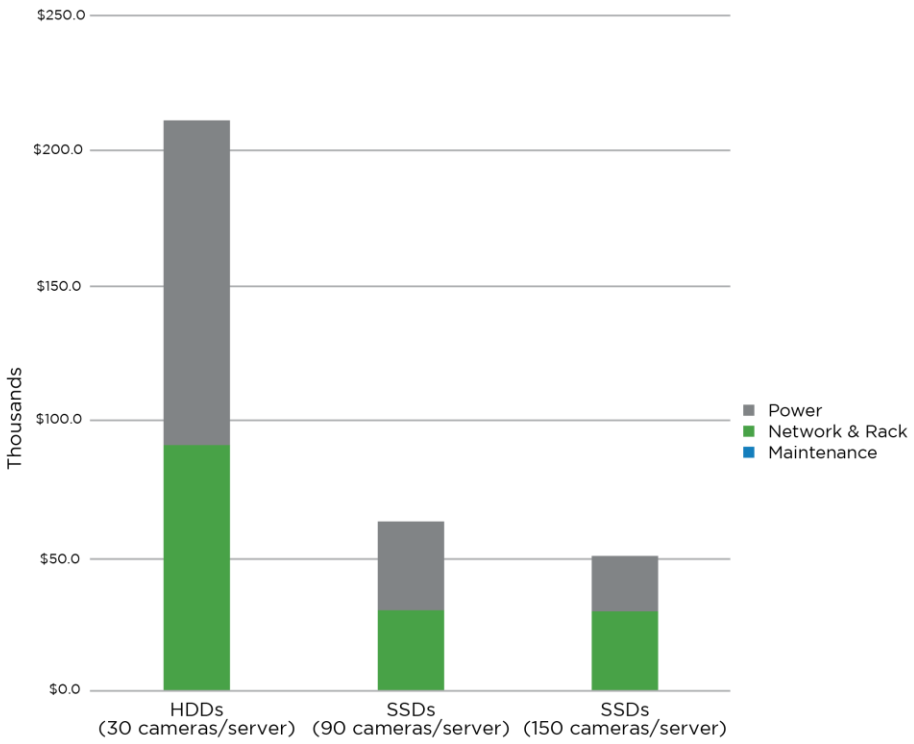

Figure 4. SSD-based video servers are denser, and cost significantly less to operate for a 1500-camera configuration.

Conclusion

HDD-based video servers are now seriously limiting the scalability and potential of video surveillance, causing risks for both video surveillance companies, and those that they serve. Utilizing flash storage as primary storage represents a new way forward with profound implications for quality, throughput, and cost. SSDs can provide fully five times the throughput of HDDs, with linear density and throughput scaling—resulting in a five-to-one consolidation of required video servers. These throughput advantages coupled with growing capacities are making flash a compelling alternative to spinning media. SanDisk flash storage enables the future demands of organizations with uncompromising video surveillance requirements, including higher quality and resolution, more cameras, no data loss, longer retention periods, and real-time analytics. Given the trajectory of solid state technology, these comparisons will only become more favorable over time. Contact SanDisk for more information on partners delivering flash-based video server solutions, or visit our Data Center Solutions page at www.sandisk.com/business/datacenter for more information.

About Milestone Systems

Milestone Systems is the world's leading provider of open platform IP video surveillance software. Milestone has provided easy-to-use, powerful video management software in more than 100,000 installations worldwide. Milestone XProtect® products are designed with open architecture and are compatible with more IP cameras, encoders and digital video recorders than any other manufacturer. Because Milestone provides an open platform, you can integrate today's best business solutions and expand what's possible with future innovations. Visit www.milestonesys.com for more.