

WHITE PAPER:

Optimal RAID Configurations for Video Surveillance



Today's video surveillance systems now include a higher number of cameras continuously capturing video content at higher resolutions. In most cases, this video content is captured for the what-if scenarios and the post-incident review and analysis that follows. Whether it is for this forensic-type analysis, liability protection, identification of those who have committed crimes, or simply a matter of regulatory compliance, video data is extremely valuable in many situations.

The nature of high-resolution 24x7 surveillance systems means that there is an enormous amount of data being generated on a continual basis. The constant flow of data being written to data storage systems introduces a higher level of risk of disk drive failure or other types of critical read/write failures. As a result of these factors, many leading storage hardware manufacturers are now providing guidance on best practices for storing video data – and limiting the types of RAID levels employed.

This paper looks at the different ways to protect video surveillance systems from data loss through the use of storage technologies, specifically RAID, and identifies those which are best for storing critical video data.

Storing Critical Video Data

The probability of disk failure increases exponentially as systems age. Since most video surveillance systems are installed with the intention of using the system for as many years as possible, the likelihood of drive failures is higher than most other IT systems. An appropriate data protection configuration is required from day one for systems that record critical video data.

Statistics show that the probability of a second drive failing before the first has been recovered, is higher than the probability of random drive failures. This is attributed to the longer rebuild times due to the larger disk capacities typically being used today. It is also due in part to the

disk drives being exposed to the same types of conditions, operating characteristics and aging within the same system.

For those tasked with designing and operating video surveillance systems as well as end users, protecting video data is often as important as protecting critical financial data. Appropriate systems must be designed to ensure data protection, while also maintaining required performance levels. For mission critical systems and valuable video data, operators will incorporate appliances with higher storage capacities and enough drives to provide data protection through the use of RAID.

What is RAID

RAID (Redundant Array of Independent Disks) is a technology used to increase the performance and/or reliability of data storage. A RAID system consists of two or more drives working in parallel in an array. These disks can be hard discs, but there is a trend to also use the technology for SSD (solid state drives). Data is distributed across the drives in one of several ways, referred to as RAID levels, depending on the required level of redundancy and performance. RAID levels are optimized for specific situations. There were originally five RAID levels, but many variations have evolved. The most common RAID levels (0/1/5/6/10) are described in this paper.

Each RAID level provides a different balance of specific operational goals: reliability, availability, performance, and capacity. RAID levels greater than RAID 0 provide protection against unrecoverable read errors, as well as against entire drive failures.

To understand the different RAID levels, it is important to first understand the different methods of storing data that are applied with RAID: striping, mirroring, and parity.

Striping – this method involves splitting the flow of data into blocks and then writing of these blocks across the RAID one by one onto different drives. This method of data storage positively affects performance, allowing for the concurrent reading of data across the different disks, increasing total data throughput. Striping effectively uses the total storage capacity of the array.

Mirroring – this storage method stores identical copies of data on different disks simultaneously. Mirroring provides good performance on writing data – similar to a single disk, and can be used for concurrent reads across the array resulting in better read performance.

Parity – this method utilizes striping and checksum methods. Data is written via striping across the array with the added benefit of utilizing a parity disk to store the checksum information.

If a drive fails, the missing data is recalculated from the checksum, providing the data redundancy aspect of RAID. This method benefits from the increased performance from striping the data. All the existing RAID types are based on striping, mirroring, parity, or combination of these storage techniques.

The software to perform the RAID-functionality and control the drives can either be located on a separate controller card (a hardware RAID controller) such as that used within a Razberi appliance, or it can simply be a driver. Some versions of Windows or Mac OS X include software RAID functionality. Hardware RAID controllers may cost a little more than a software option, but they provide better performance, especially with RAID 5 and 6.

RAID Levels

RAID 0 - STRIPING

In a RAID 0 system data is split up in blocks that get written across all the drives in the array. By using multiple disks (at least 2) at the same time, this offers superior I/O performance. This performance can be enhanced further by using multiple controllers, ideally one controller per disk.

Advantages

- RAID 0 offers great performance, both in read and write operations. There is no overhead with this method meaning that the amount of available storage equals the sum of all drives' raw storage capacity.
- All storage capacity is used, there is no overhead.
- The technology is easy to implement.

Disadvantages

- RAID 0 is not fault-tolerant. If one drive fails, all data in the RAID 0 array are lost. It should not be used for mission-critical systems.

Appropriate Applications

RAID 0 is ideal for non-critical storage of high bit rate video data (high resolution cameras with high frame rates & quality settings). RAID 0 is not recommended for critical video surveillance recordings that must be kept for regulatory reasons.

RAID Levels (Continued)

RAID 1 - MIRRORING

Data is stored twice by writing to both a data drive (or set of data drives) and a mirror drive (or set of drives). If a drive fails, the RAID controller uses either the data drive or the mirror drive for data recovery and continues operation. You need at least 2 drives for a RAID 1 array.

Advantages

- RAID 1 offers excellent read speed and a write-speed that is comparable to that of a single drive.
- In case a drive fails, data do not have to be rebuilt, they just have to be copied to the replacement drive.
- RAID 1 is a very simple technology.

Disadvantages

- The main disadvantage is that the effective storage capacity is only half of the total drive capacity because all data get written twice.
- Actual read throughput of most RAID 1 implementations is slower than the fastest drive. Write throughput is always slower because every drive must be updated, and the slowest drive limits the write performance.

Appropriate Applications

RAID 1 is suitable for mission critical storage, such as video storage for casinos, city surveillance or other applications which are required to store data due to regulatory reasons. It is also suitable for small servers in which only two data drives will be used.

RAID 5 - STRIPING WITH PARITY

RAID 5 is the most common secure RAID level for data center applications, but is not a best practice for video storage. It requires at least 3 drives but can work with up to 16. Data blocks are striped across the drives and on one drive a parity checksum of all the block data is written. The parity data are not written to a fixed drive, they are spread across all drives. Using the parity data, the computer can recalculate the data of one of the other data blocks, should that data no longer be available. That means a RAID 5 array can, in theory, withstand a single drive failure without

losing data or access to data. Although RAID 5 can be achieved in software, a hardware controller is recommended. Often extra cache memory is used on these controllers to improve the write performance.

Advantages

- Read data transactions are very fast while write data transactions are somewhat slower (due to the parity that has to be calculated). This is important for video surveillance applications which have read-intense workloads. However, most surveillance workloads have less than a 10% read requirement (i.e. the storage array is typically used more than 90% of the time for writing data).
- If a drive fails, you still have access to all data, even while the failed drive is being replaced and the storage controller rebuilds the data on the new drive.

Disadvantages

- Drive failures have an effect on throughput, although this is still acceptable.
- If one of the disks in an array using large disks fails and is replaced, restoring the data (the rebuild time) may take a day or longer, depending on the load on the array and the speed of the controller. If another disk goes bad during that time, data is lost forever.
- If the system encounters an unrecoverable read error (URE) during the rebuild of a failed drive, the rebuild will fail and the data is ultimately lost. UREs used to be an acceptable problem, but as drive capacity has increased, the odds of encountering a URE on a large video storage drive have increased to unacceptable levels. It is now not only possible, but quite likely, that a RAID 5 array will fail to rebuild and therefore offer no protection against data loss.

Appropriate Applications

RAID 5 is not recommended for mission critical data such as high-value video surveillance footage as it will not provide enough redundancy and data protection. Additionally, this RAID level is not suitable for "live" video storage as it may not be able to handle the higher bit rate and frequency, particularly with megapixel camera streams. As mentioned previously, major storage vendors such as Dell have actually issued advisories against

single drive failure without losing data or access using RAID 5 for any mission critical applications or high value data storage. Dell's web site provides TR1020 Member RAID position advisory information.

RAID 5 might be an adequate solution for archiving video data on archival storage arrays since the value of archived video is often less and potentially used in conjunction with primary storage hardware such as the Razberi ServerSwitchIQ appliance.

RAID 6 - STRIPING WITH DOUBLE PARITY

RAID 6 is similar to RAID 5, but at this level, the parity data is written to two drives. That means it requires at least 4 drives and can withstand 2 drives dying simultaneously. The chances that two drives break down at exactly the same moment are, of course, very small. However, if a drive in a RAID 5 systems dies and is replaced by a new drive, it takes hours to rebuild the swapped drive. If another drive dies during that time or the rebuild runs into unexpected errors or failures, you still lose all of your data. With RAID 6, the RAID array will even survive that second failure and make rebuild process much more robust.

Advantages

- Like with RAID 5, read data transactions are very fast.
- If two drives fail, you still have access to all data, even while the failed drives are being replaced. So RAID 6 is more secure than RAID 5. This also means that the URE problem noted in the RAID 5 configuration is not a problem in RAID 6 since failure would require two URE events, a very unlikely occurrence.
- RAID 6 makes larger RAID groups more practical, especially for high-availability systems, as large-capacity drives take longer to restore.

Disadvantages

- Write data transactions are slowed down due to the parity that has to be calculated.
- As with RAID 5, drive failures have a negative effect on throughput.
- Rebuilding an array in which one drive failed can take a long time.

Appropriate Applications

RAID 6 is a good all-round system that combines efficient storage with excellent security and decent performance. It is preferable over RAID 5 for critical video surveillance systems that use many large drives for data storage. However, this RAID level may not be suitable for "live" video storage without a high performance hardware RAID controller as low performance devices will not be able to handle the higher bit rate and frequency, particularly with megapixel camera streams.



RAID 6 is a good all-round system that combines efficient storage with excellent security and decent performance. It is preferable over RAID 5 for critical video surveillance systems that use many large drives for data storage.

RAID 10 – COMBINATION OF RAID 1 & RAID 0

It is possible to combine the advantages (and disadvantages) of RAID 0 and RAID 1 in one single system. This is called a nested or hybrid RAID configuration. It provides security by mirroring all data on secondary drives while using striping across each set of drives to speed up data transfers. This method creates a striped set from a series of mirrored drives. The array can sustain multiple drive losses so long as no mirror loses all its drives.

Advantages

- If something goes wrong with one of the disks in a RAID 10 configuration, the rebuild time is very fast since all that is needed is copying all the data from the surviving mirror to a new drive. As an example, the rebuild time could take as little as 30 minutes for drives of 1 TB.

Disadvantages

- Only 50% of the total raw capacity of the drives is usable due to mirroring, so compared to large RAID 5 or RAID 6 arrays, this is an expensive way to have redundancy.
- As long as one disk in each mirrored pair is functional, data can be retrieved. However, if two disks in the same mirrored pair fail, all data will be lost because there is no parity in the striped sets.

Appropriate Applications

Given that RAID 10 is superior to other RAID levels in terms of write performance and reliability (particularly with 4 or more disks) it is ideal for applications which require a high level of write performance such as video surveillance systems running today's megapixel cameras. Additionally, the higher level of redundancy associated with RAID 10 makes it an ideal option for the cases where high data redundancy is required. Since the cost of storage has dropped considerably and, instead of being at a premium, is typically in abundance today and because of the higher performance characteristics, the move to RAID 10 has become a more standard approach for video surveillance.

Beyond RAID

All RAID levels except RAID 0 offer protection from a single drive failure. As mentioned, a RAID 6 system even survives 2 disks dying simultaneously. For complete security you do still need to back-up the data from a RAID system.

- That back-up will come in handy if all drives fail simultaneously because of a power spike.
- It is a safeguard when the storage system gets stolen or damaged. Back-ups can be kept off-site at a different location. This can come in handy if a natural disaster or fire destroys your workplace.
- The most important reason to back-up multiple generations of data is user error. If someone accidentally deletes some important video evidence and this goes unnoticed for some period of time, a good set of back-ups ensure you can still retrieve those files.

Summary & Recommendations

- Critical video surveillance applications – use RAID 6 or RAID 10
- Non-critical video surveillance applications – use RAID 0 or RAID 1
- RAID 5 should not be used for video surveillance data storage

Choosing the appropriate level of RAID depends on how critical the data is that you are storing. If you are required by law to keep video data for a certain period of time, then there is good reason to employ RAID level 6 or 10. If you are capturing video data that could potentially be used in legal liability cases such as slip and falls, then there could also be good reason to employ RAID levels that provide greater protection. In addition to the level of protection, you must consider the effect of RAID on your system performance and utilization of storage capacity. The table on the next page helps summarize and compare the various levels of RAID to determine which is best for your video data storage purposes.



Choosing the appropriate level of RAID depends on how critical the data is that you are storing.

Razberi recommends RAID 6 or 10 for all critical video surveillance applications and RAID 0 or 1 for non-critical applications. The limitations and the higher level of risk associated with RAID 5 should disqualify that RAID level from most if not all video surveillance systems.

Summary and Comparison of Various RAID Levels

Features	RAID 0	RAID 1	RAID 5	RAID 6	RAID 10
Storage Method	Striping	Mirror	Striping with Parity	Striping with Double Parity	Mirrored Striping
Minimum # Drives	2	2	3	4	4
Data Protection	None	Single-drive failure	Single-drive failure	Two-drive failure	Up to one disk failure in each sub-array
Read Performance	High	High	High	High	High
Write Performance	High	Medium	Low	Low	High
Storage Capacity Utilization (% of Raw)	100%	50%	67-94%	50-88%	50%
Summary and Applications	<p>Not recommended for critical video storage</p> <p>Highest R/W for large bit rate recording & playback</p> <p>Maximum storage utilization but provides no redundancy</p>	<p>Used for critical video storage</p> <p>High R/W for large bit rate recording & playback</p> <p>Some redundancy but at a cost of half the total available storage</p>	<p>Not recommended for critical video</p> <p>Low write speed performance</p> <p>Long rebuild times</p> <p>Provides some redundancy</p> <p>Offers more storage for video</p> <p>Best used in archiving servers</p> <p>Not recommended for use w/live data base due to URE rebuild issue, size of megapixel bit rates, 24x7 requirement of surveillance & rebuild times</p>	<p>Recommended for critical video</p> <p>Low write speed performance</p> <p>Long rebuild times</p> <p>Provides redundancy</p> <p>Offers more storage for video</p> <p>Best used in archiving servers</p> <p>Not recommended for use w/live data base due to size of megapixel bit rates, 24x7 requirement of surveillance & rebuild times</p>	<p>Recommended for critical video</p> <p>Provides the best combination of speed & redundancy</p> <p>Low utilization of available storage for recording</p>

Hardware Recommendations

Appropriate hardware for the capture and storage of video surveillance footage include the Razberi ServerSwitchIQ Pro & ServerSwitchIQ Enterprise appliances. These platforms provide superior performance to cost of storage, based on the number of cameras that can be connected.

Razberi ServerSwitchIQ 16/24 Port 1U Pro Models – These models use the Intel® Rapid Storage Technology driver application found on the Start Menu for configuring RAID. Razberi Technologies recommends reviewing the latest documentation covering this tool from the <http://www.intel.com> web site.

Razberi ServerSwitchIQ 24 Port 2U Enterprise Models – These models use the High Point 45XX series hardware for controlling up to eight hard disk drives. Razberi Technologies recommends reviewing the latest documentation covering the RocketRAID 45XX series from the <http://www.highpoint-tech.com> web site.



Contact Razberi Technologies today to request a demonstration and discover what the intelligent surveillance appliance can do for you at: www.razberi.net

Razberi is the leading provider of intelligent surveillance appliances. Razberi ServerSwitchIQ™ is a complete line of all-in-one, scalable appliances that combine a managed PoE switch, server, storage and intelligence. All Razberi appliances are open-architected to work with world class video management systems (VMS) out of the box. Intelligent features include VyneWatch™ health monitoring and LocBerj™ cyber security features to lock down vulnerable ports. By deploying Razberi appliances in a distributed architecture near the edge of the network, megapixel video impact on the network is reduced by up to 95%. Learn more at: www.razberi.net

© Copyright 2016 Razberi Technologies, Inc. All rights reserved.

Americas: 469-828-3380
UK/EMEA: 0203 773 3689
salesinfo@razberi.net

wpre20161005