WHITE PAPER

Surveillance Redundancy and Resilience

Preventing Data Loss and Downtime

SYNECTICS

Introduction

Regardless of geographic location or sector application, the growing dominance of IP, widespread adoption of megapixel cameras, and the rapidly expanding array of video, alarm, and transactional integrations expected of - and made possible by - surveillance solutions, means that protecting data has never been a greater priority.

It has also never been more challenging. As video resolution and camera capabilities improve and systems expand, the size and volume of data that systems need to be capable of handling increases dramatically. This is particularly important with heightened global safety and security concerns, especially in mission-critical environments.

To put this into perspective, in 2015 new surveillance systems generated over 566 petabytes of data globally. That is twice the amount of data currently stored on Facebook.

Crucially, however, this data is no longer passive – generated and stored predominantly for post event review. It is active and in many cases proactive data and information that organizations need in real time in order to make critical decisions that impact on security, safety, efficiency, and in many cases, profitability.

Whether applied to a mega casino preventing fraud and maximizing profit, or a transport hub helping passengers move safely and quickly from A to B, surveillance data has broadly become an essential operational lynchpin. It has to be reliable, actionable, and completely secure.

To meet these challenges, the most effective surveillance solutions will always incorporate measures to guard against data loss and minimize – ideally eliminate – any potential system downtime. This is typically, and best, achieved by addressing redundancy and resiliency on two levels – through individual units (i.e. system components) and with system-wide measures. This white paper examines these measures in detail and looks at some of the key factors organizations need to consider when implementing or upgrading their surveillance solution.

Unit Level Protection

ASK YOUR SUPPLIER:

RAID technology is often built into the DVR systems, storage, and management servers (common terminology will refer to integral/auto RAID rebuild) offered by surveillance solution providers. If this is not the case you should confirm the suitability of the device being specified.

Each individual product that makes up a surveillance system solution (for example cameras, encoders) is defined as a unit. Each unit should ideally have built-in redundancy and resilience measures. Unit level measures will vary, depending on the composition, size, and primary purpose of the complete surveillance solution – as all these factors will dictate the range of technologies used. However, there are several features that are worth noting.

Understanding RAID Levels and Surveillance System Suitability

When discussing surveillance system redundancy it is essential to understand the concept of RAID (Redundant Array of Independent Disks) and differentiate the pros and cons of differing RAID levels i.e. the way data is distributed across storage system disks to prevent data loss and/or to enhance processing and performance speeds.

There are seven main RAID levels but in practical terms, only three are used in relation to surveillance storage and data redundancy – RAID 1, RAID 5, and increasingly, RAID 6.

RAID 1

This RAID level involves total data mirroring i.e. exact duplication, on two or more independent disks. Its simplicity is both its strength and its weakness. Making an exact copy of all data on a separate disk means the array will always be fully functional as long as one member drive is operational. However, because all data is completely duplicated, it can also be both expensive and slow – having to write large volumes of data takes time and uses up a great deal of storage capacity.

In a surveillance setting, where longer retention of higher definition footage is the norm, RAID 1's practical and financial limitations start to prove problematic. For large-scale, complex environments that rely heavily on video data, for example casinos, RAID 1 is not an appropriate solution. However, RAID 1 is highly suited to database/application servers as they require fewer disks.

RAID 5

RAID 5 is probably the most common in terms of surveillance storage. Here, data is 'striped' across at least three drives with distributed parity (rather than a dedicated parity disk as is the case with RAID 3 and 4). In the event of a single drive failure, this effectively means that all the data needed to prevent data loss can be calculated quickly from the remaining active disks.

Only in the unlikely event of multiple drive failure would data be at risk. Also, because of the way information is spread across multiple drives, performance (i.e. read/write speed) is significantly faster than RAID 1 solutions.

RAID 6

This level is sometimes referred to as 'double parity RAID' – essentially because data is distributed and stored in the same way but with the added benefit of a second parity spread across the drives. It is a feature that safeguards stored data even in the event of two drive failures. The downside is that this does mean RAID 6 requires additional storage to create the same storage space and is therefore slightly less efficient than RAID 5.

Until recently, RAID 6 was considered overkill in terms of precautionary measures. However, mirroring the upturn in the use of HD IP cameras, particularly in enterprise-class surveillance solutions where downtime is unacceptable and rapid data retrieval is paramount, RAID 6 is increasingly becoming the new standard.

Pay Attention to Power

IT professionals will be more than familiar with the term 'hot standby' – the replacement of an active device with a similar alternative, without losing operational capability. It is also a key concept for unit level redundancy and resilience within surveillance systems.

Hot standby power – or dual/redundant hot standby power – refers to a device's ability to draw power from an alternative source should the primary power supply fail, thus ensuring that data flow, storage, and retrieval are unaffected. For example, PSU failure in the main application server results in zero loss as the secondary PSU takes over.

This has become an almost essential feature of surveillance solution components for mission-critical security and operational environments dealing with large volumes of data.

In modern systems, even IP cameras can support hot standby PSUs (through dual power supply - 1 x PoE, 1 x low voltage). However, it is important to remember that being able to swap supply in the event of a network outage means there has to be a direct source to switch to, which in turn has financial and physical infrastructure requirements that need to be considered. This is why many organizations that adopt this particular 'belt and braces' approach to system resiliency are strategic about it i.e. prioritizing cameras/devices that matter most. For example, in the event of a network power failure, digital recording/storage devices or networked cameras with this functionality – on detecting the outage – would immediately swap to a local power supply.

Dual Network Connections

Another redundancy feature to look out for in key system devices is 'dual network connections', which essentially enable joint or switched connection to different networks. This helps ensure there is no single point of failure. In the event of a faulty port or network failure, data is automatically transmitted via the alternative route.

Localized Storage

Another important redundancy mechanism for protecting IP video data in the event of network failure is localized storage. Increasingly, IP cameras and encoders are equipped with SD memory cards or internal hard drive storage. These can be configured to act as a data safety net should network interruption/failure occur preventing footage from being streamed to the primary networked recording and storage device.

ASK YOUR SUPPLIER:

The ability to use standby power with an IP camera can trigger a reboot response in cameras that may, in turn, result in downtime. The length of reboot times, if required, will vary so it is important to understand potential scenarios from the outset. For organizations that need to guarantee continuous feed and camera control – perhaps for regulatory or safety reasons – it is best to specify cameras that switch power seamlessly with no reboot required.

The length of storage time available is dictated by SD card size and type. Given recent improvements in SD card technology, the 'safety net' is quite substantial. For example, it would be feasible to record over 24 hours' worth of full HD footage at 5mbps on a 64GB SD card.

It is worth remembering, however, that in some instances a lag can occur between network failure and the camera detecting the issue and switching to localized storage – perhaps up to four or five seconds. For cameras in critical locations, where even a slight loss/delay in footage availability is unacceptable, it is advisable to record locally to the SD card AND to the main server. In this scenario it is worth remembering solutions can be configured to record in low definition to the SD card and full HD to the server, in order to reduce storage space and prolong SD card lifespan.

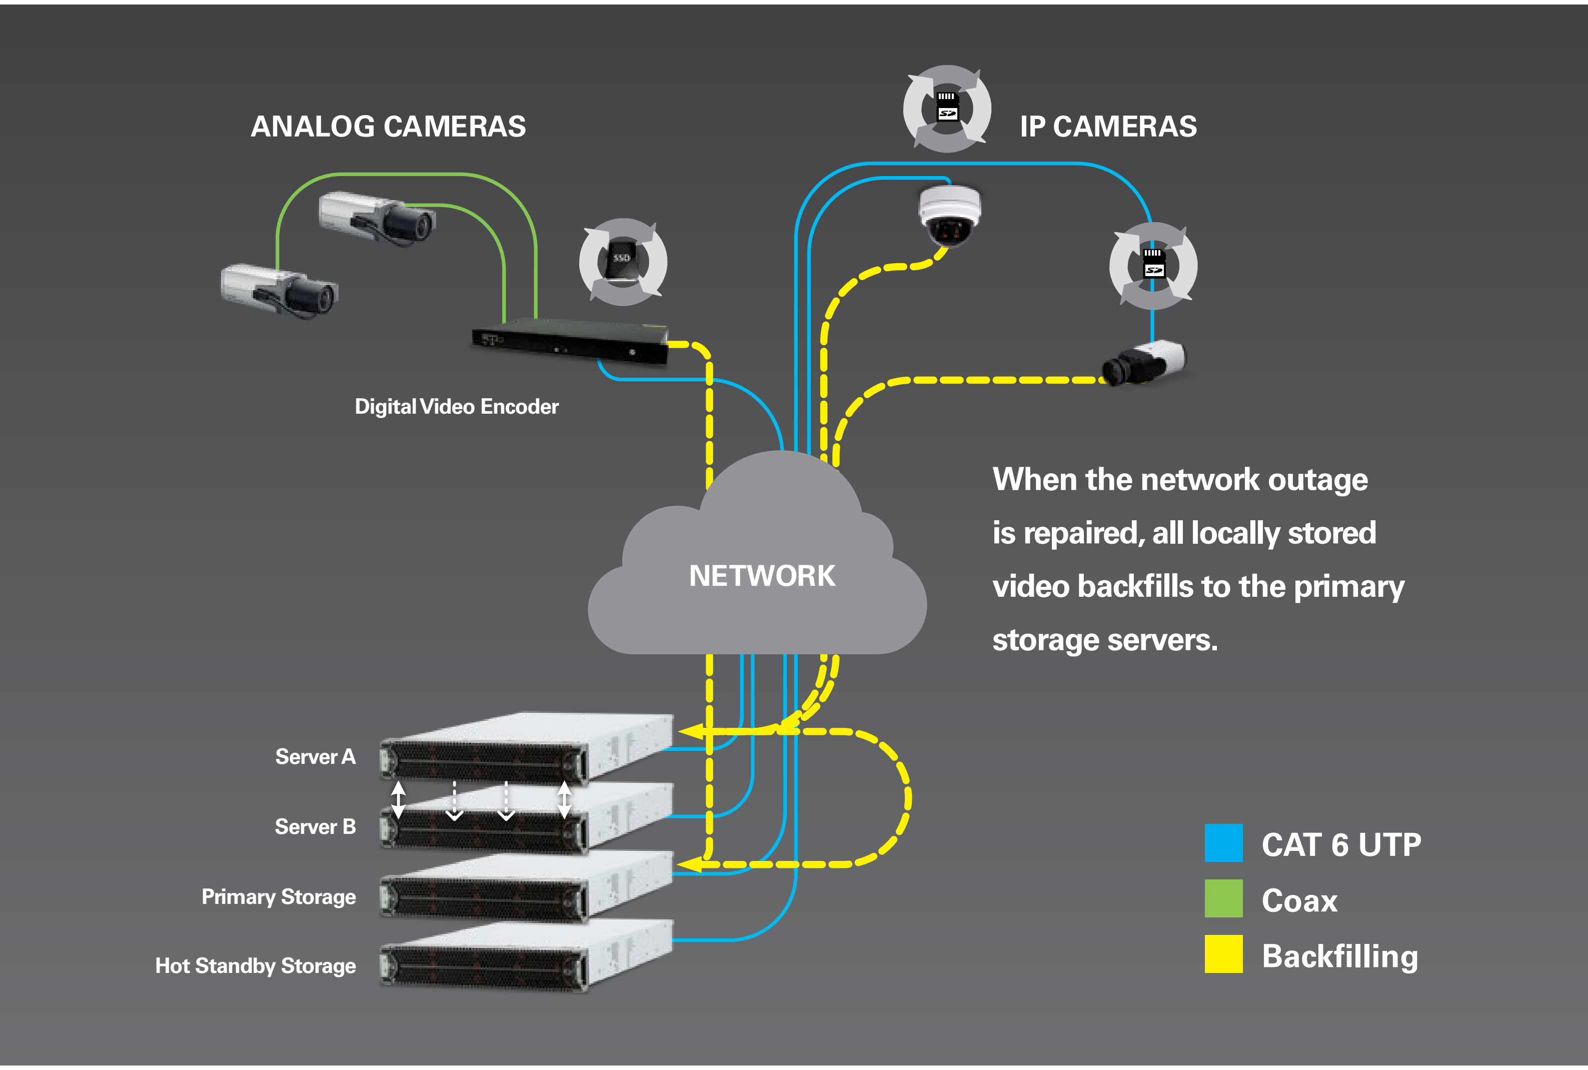

When the network outage is repaired, all locally stored video backfills to the primary storage servers.

Backfilling

Taking the advantages of localized storage to the next level is the concept of ‘backfilling’. This can be explained using a network failure scenario.

Upon detecting a network interruption, edge-based IP cameras and encoders – equipped with an SD memory card/integral storage hard drive – switch to local recording to ensure that data is protected. But what happens when the fault or power failure is remedied and the network is restored? No data has been lost but it is now disjointed, split between edge devices and the primary storage server.

Surveillance management platforms equipped with backfilling capabilities avoid this scenario by being able to identify ‘missing’ video once the network has been restored, identifying which devices have that footage stored, and automatically copying the relevant video back onto the primary storage server, all without any interruption to the management and operation of those devices. This ensures that operators maintain continuous situational awareness and guarantees seamless video viewing, recording, and retrieval.

Backfilling ensures that operators maintain continuous situational awareness and guarantees seamless video viewing, recording, and retrieval.

System Level Protection

In addition to unit level redundancy measures to safeguard data, surveillance solutions should also incorporate robust system level mechanisms to mitigate risk.

Server and Database Replication

Perhaps the most crucial measure available is server replication i.e. the real-time replication of primary server data and functionality over multiple physical/virtualized servers. Importantly, these back up servers can be (and in mission-critical settings often are) located separately to the primary server site. This latter point is particularly relevant for industrial settings where physical risks, such as fire or explosions, are greater.

Replication at this level guarantees full system control and data access should the primary server fail and should be configured as an automated process on detection of failure, removing the need for operator intervention that would delay action.

In failover situations, storage space becomes an even more important issue – particularly for large-scale solutions that routinely manage huge volumes of data and footage from thousands of cameras.

Considering the Cloud

Increasingly, organizations are also looking at cloud-based storage as a redundancy measure. In theory, it is an option with a number of advantages including a reduced need for additional physical infrastructure. However, there is one significant issue with this approach and that is bandwidth capacity.

Video surveillance redundancy necessitates far greater levels of bandwidth than any other application. An image captured by an HD IP camera is about four times the size of one captured via an analog unit. In the context of a mega-casino, major oil and gas processing plant, or a critical infrastructure deployment - where hundreds, often thousands of cameras are relied upon every minute of every day - the challenge becomes clear.

Even if bandwidth would allow for this size and volume of data to be stored, speed would remain an issue. For sites where surveillance is mission-critical and real-time capabilities are essential, the downside of cloud-based redundancy is substantial.

ASK YOUR SUPPLIER:

In addition to server replication, it is worth asking your supplier about 'automated storage prioritization'. For example, it is possible to configure a system, in failover mode with hot standby storage, to automatically prioritize data from critical cameras where security or safety risks are of greatest concern.

The caveat to this is that cloud-based redundancy could, and increasingly does, have a place for protecting data with specific purpose. Take, for example, an organization that may need to routinely share information securely with third parties – perhaps law enforcement – that do not need or want to have access to all surveillance footage. In these situations, it may make sense to implement evidence locker redundancy to replicate data physically stored on the LAN/WAN within a secure cloud-based 'locker'.

Health Monitoring

Prevention is always better than cure, which is why many organizations are also now taking advantage of resiliency measures such as real-time system health monitoring.

An integral feature of some surveillance management platforms, this functionality essentially means the system is self-aware to the extent that it can identify indicators of potential critical/catastrophic failure (e.g. packet loss). This functionality can alert operators to these issues – perhaps via an instructive maintenance procedure workflow – and if necessary the system takes its own remedial action. It may, for example, hot swap digital recording systems if it detects indications of a potential fault.

There are several significant advantages to this approach. It can help avoid complete data replication and storage scenarios, which in turn has positive cost implications. Importantly, the combination of health checking and automated, individual device hot standby, means faults can be identified, repaired, or replaced without any system downtime or impairment to everyday operations.

One Size Does Not Fit All

In this paper we have looked at some of the most frequently used and important redundancy and resiliency measures available to protect data and prevent downtime. So how many of these measures should an organization adopt in order to be strategic about surveillance systems protection? There is not a single correct answer for this.

Solutions have to be tailored to a specific organization's requirements, taking into account regulatory demands, existing network infrastructure, solution size, camera types, sector needs, security vulnerabilities, and many more factors.

There are, however, two considerations that will always be relevant:

1. Whatever the measures deployed, the end result should always be that there is no potential single point of failure.
2. Fully IP and hybrid surveillance solutions are designed to offer scalability and flexibility – enabling systems to grow organically as need dictates. The same applies to redundancy/resiliency measures. Focus on the absolute essentials first – further measures can always be implemented later.

By starting with these considerations, being aware of the options available, understanding the numerous internal and external factors that influence requirements, and working in partnership with surveillance solution providers, organizations can be confident of a solution that reflects and supports their specific objectives.

For more information about Synectics technology solutions, visit our website: synecticsglobal.com.

Synectics designs integrated end-to-end surveillance control systems for the world's most demanding security environments. We excel at complex projects that require innovative, tailored solutions with high reliability and flexibility, specifically for casinos, oil and gas, marine, public space, banking, transport and critical infrastructure applications.

With over 30 years of high security systems experience, field proven products, and expert support personnel in the UK, US, Europe, UAE and Asia Pacific, Synectics offers its clients turnkey networked solutions for comprehensive protection and peace of mind.

Synectics' Systems division is part of Synectics plc, a global leader in advanced surveillance, security and integration technologies and services.

Synectics
sales@synecticsglobal.com
synecticsglobal.com

Americas Asia Europe Middle East United Kingdom

Specifications subject to change. E & OE.
Literature Reference: REDRES-WP/0119/A4 Iss 3