

Top 14 best practices for building video surveillance networks

Introduction

The shift from analog to Internet Protocol (IP) surveillance cameras has changed the way that video surveillance systems are built. Instead of running lengths of coax cable from multiplexers out to cameras, an IP-based surveillance system involves plugging the cameras into an Ethernet Local Area Network (LAN).

There are numerous benefits to using an IP-based system, including simplified cabling, power reticulation over the cabling, native digital video, and greater camera control. Taking full advantage of these benefits requires making the right choices in the configuration of the underlying Ethernet LAN. In addition, using an Ethernet LAN presents its own challenges, as it is a significantly more intelligent and dynamic environment than a collection of coax cables.

In order to maximize system capabilities and avoid potential pitfalls, installers need to have an understanding of the best practices in configuring Ethernet network infrastructure within a video surveillance system. This white paper explains the top 14 best practice guidelines. These guidelines have been honed through the many years of experience that Allied Telesis have in working with customers to create robust, high performance video surveillance networks.

Allied Telesis equipment is used in successful surveillance networks all over the world. Examples include:

- city-wide installations comprising thousands of cameras
- systems operating in harsh railway-network environments
- prison surveillance systems
- security networks for major international sporting events
- smaller installations in malls, schools, offices, hospitals, and more

The best practices listed here fall into four broad areas. Each has a core motivation or “goal”:

- **Goal 1: Minimize data loss**
- **Goal 2: Maintain security**
- **Goal 3: Maximize flexibility**
- **Goal 4: Future-proof the network**

Each of these four goals has a brief introduction, explaining its significance to video surveillance systems, which is then followed by the individual pieces of advice or “tips” that fall within the goal. There are 14 tips altogether, which if followed, lead to the best possible video surveillance solutions. Excellent outcomes are achieved by getting the details right at every level.

The goals and tips range across all aspects of network design, from where to locate switches, right down to the finer details of protocol configuration.

Goal I: Minimize data loss

In a traditional IP network – be it the Internet or a corporate LAN – the loss of some data packets, or even a brief outage, are recoverable events. Certainly, an outage of an e-commerce website can be expensive in terms of lost sales, but a brief outage will typically result in just a small blip in the overall sales figures.

However, for a video surveillance network, even a brief outage at the wrong moment can mean failing to record the vital evidence that catches a criminal, or failing to alert security personnel in time to prevent a criminal act. With real-time video capture, there is no second chance. Once a particular moment in time has passed, there is no going back to record it again.

One of the prime motivations for installing video surveillance systems is that they provide 24/7 vigilance. So, lost data or network outages undermine a primary purpose of the system.

The goal of minimizing data loss is therefore a driving principle in the design of video surveillance networks. This goal can be achieved in the following five ways:

Tip 1 - protect the switches

While video cameras must often be installed in exposed locations, the switches shouldn't have to be. Wherever possible, locate switches in places that are protected from environmental and human harm. Cameras can be up to 100 meters away from their associated switches, which provides plenty of scope for finding or building a secure, protected housing for switching equipment.

Tip 2 - double up on data paths

The single most likely point of failure in a data network is the cabling. This is particularly so in a network that extends over some miles, or when data cabling runs along with other utilities like water, electricity, or telephone lines. Work on one utility's infrastructure is notorious for causing outages to other utilities.

Using resilient ring network design, and/or using aggregated links that follow separate paths, enormously reduces the risk of outages due to cable damage.

Tip 3 - double up on power supply

The second most likely point of failure in a data network is in the power supply. Within switching equipment, the power supply unit runs at the highest temperature, and is statistically the most likely component to fail. Moreover, the reticulation of power to the switching equipment is vulnerable to interruption.

Using switching equipment that supports dual Power Supply Units (PSUs), and installing two PSUs, guards against PSU failure. If the PSUs are hot-swappable, then the recovery from a failed PSU is completely non-disruptive – the replacement of the failed PSU can be carried out while the switch continues uninterrupted operation powered by the remaining PSU.

Power reticulation failure can be guarded against by connecting a switch's two PSUs to independent power sources, if they are available.

Tip 4 - guard against network storms

If Ethernet packets have an unblocked loop to circulate around, then they will – endlessly, at very high data rates. Packets endlessly storming around in circles can grind normal network operation to a halt. So, correct configuration of the network, using a robust ring resiliency protocol between the switches, and spanning tree or loop protection at the edge, is essential. If spanning tree is being used at the edge, then it is vital that spanning tree edge-port status is set on the edge ports, to prevent disruptive flushing of forwarding tables when ports link up.

Tip 5 - treat the cameras with care

This does not mean taking physical care with the cameras, although that is also a good idea. Rather, it is important to take care not to overload the cameras' processors. By necessity, to keep costs and power consumption down, video surveillance cameras use Central Processing Unit (CPU) processors of modest capacity. The task of processing incoming image data, and pumping it out as IP data, will typically consume most of the power of a camera's CPU. If a camera has the additional load of examining and processing unnecessary IP data coming into it off the network, then that can overload the CPU, causing deterioration of the quality of the video it is feeding into the network.

Switches need to be configured so that they only send the required bare minimum of packets to the cameras. This can be achieved by:

- Avoiding flooding of multicast data, by never disabling Internet Group Management Protocol (IGMP) Snooping on switches.
- Avoiding flooding of multicast data by configuring edge switches to not periodically age entries out of their forwarding tables. Edge switches have only a small number of video streams to forward - typically, one per connected camera - so they are operating well within the capacity of their forwarding tables. Moreover, the video streams they are forwarding do not change dynamically, as they are just regularly servicing the same set of connected cameras.

Edge switches should therefore be put into a mode where they maintain permanent forwarding-table entries for all the video streams they are forwarding, rather than the default behaviour of treating all entries as potentially stale, and periodically aging them out in case table space is required for newly-added video streams. The periodic age-outs of the streams are invariably accompanied by brief bursts of video data flooded in all directions – most notably towards other cameras.

- Preventing multicast protocol signaling packets from being sent to cameras. IP multicast protocols use signaling packets to discover where video streams need to be forwarded to. The signaling packets are sent regularly, to maintain forwarding states, and are usually repeated once or twice, to make sure they get through. By default, these signaling messages are sent in all directions, as the protocol does not inherently know where receivers of video streams are located. This level of signaling chatter can be disruptive to the busy CPUs in cameras. Also, the signaling messages relating to video streams never need to be sent to cameras, as the cameras never need to receive video streams. So, in a video surveillance network, switches should be configured to suppress the sending of unnecessary signaling messages to cameras.

Goal 2: Maintain security

An Ethernet network, unlike a collection of analog video cables, is an intelligent data network. This means it is a potential point of entry for those who wish to gain entry to an organization's data network for malicious purposes.

Video surveillance cameras are by necessity often installed in publically accessible locations. Therefore, their data connections present publically accessible opportunities for connecting into a data network. Unplugging a camera, and replacing it with a Wi-Fi access point, provides a criminal with the opportunity to steal data, or to launch Denial of Service (DoS) attacks to create a diversion or otherwise disrupt the surveillance network whilst a crime is being carried out.

Therefore, the switch ports that cameras are attached to should be afforded the highest level of protection. This goal can be achieved in the following three ways:

Tip 6 - configure high-security authentication on all camera-connected ports

The best protection against malicious entry into the network via a camera's data connection is to use port authentication. Switch ports should be set up to never allow data exchange with a connected device until that device has provided authentication credentials that have been verified by the network's authentication server. Cameras support high-security authentication via encrypted digital certificates, which, unlike simple password authentication, cannot be guessed or discovered by data eavesdropping.

Tip 7 - configure switches to send alarm messages if cameras are ever unplugged

Standard network management protocols provide alarm messages that are sent upon port state change (ports being plugged or unplugged). Network Management Stations receive and display these alarms – immediately alerting staff to any attempt to tamper with camera connections.

Tip 8 - ensure that any switch ports to which cameras have not yet been attached are shut down

Any ports that are not in use need to be shut down, so there is no chance of them being hijacked in any way.

Goal 3: Maximize flexibility

Network needs change over time. Networks expand, organizations merge, services evolve, facilities move to new locations. So the network design needs to be adaptable, and easy to rearrange.

Tip 9 - use multicast

Video streams can be sent by unicast or by multicast. A unicast system is simpler to set up initially, and involves less protocol interactions than a multicast system. However, a unicast video surveillance network suffers from inflexibility in terms of where video streams can be directed to.

The very nature of unicast data transmission is that the data is sent to only one destination. If video streams are to be viewed in one location, and recorded in another, then a process needs to be devised whereby either the recording or monitoring system will retransmit streams to the other location.

If the streams need to go to yet another location, for secondary monitoring or for recording onto a backup server, then the streams need to be retransmitted twice - and again for any subsequent locations.

Multicast data is inherently capable of being delivered to multiple destinations. For a new device to begin receiving streams, it simply has to participate in the network's multicast signalling protocol, and start requesting the streams. The network automatically replicates the streams and begins delivering new copies to the new location.

A system based on multicast transmission is particularly able to accommodate dynamic redirection of video streams. For example, a multicast system can easily provide the ability to deliver video streams to particular monitoring stations in real time, in reaction to specific events.

Tip 10 - employ an intelligent network management framework

For a network to be flexible and adaptable, it must be able to roll out new features and new configurations across the network with a minimum of disruption.

For any but the smallest of networks, performing network-wide tasks like upgrading software, improving security configuration, altering protocol settings or turning on new features can be problematic. These processes can be time consuming, error prone and disruptive. Using an intelligent network management framework like Allied Telesis Management Framework (AMF) enables the automation of network-wide tasks – saving time, and eliminating errors. The network can then be adapted to changing needs in an efficient, low-risk manner.

AMF also simplifies or automates other routine tasks, for example applying configuration to new or replacement switches, backing up configurations, collecting network-wide stats, and many more.

Tip 11 - use a backbone design that supports multiple head-ends

Whilst a network may initially be installed with just one core location for data recording and network management, disaster recovery requirements can drive the need to mirror the core to 2 or more sites. If the network design does not consider head-end mirroring right from the start, then attempting to fit this into the network at a later date can prove very difficult, and require major network redesign.

Basing the network design on a resilient core ring, or a long distance clustered-switch backbone technology like Allied Telesis Long Distance Virtual Chassis Stacking (VCStack-LD), facilitates a smooth integration of mirrored head-end locations.

Goal 4: Future proof the network

Flexible design facilitates the simple rearrangement of physical topology, the dynamic re-routing of traffic, reconfiguration of existing technologies, and much more. Better still is to also have a network that is well positioned to take advantage of new technical advances. It is highly desirable to be able to take advantage of new camera capabilities as they arrive without needing to replace the Ethernet infrastructure.

Tip 12 - deploy PoE+

Power over Ethernet (PoE) has greatly simplified the process of installing video cameras. As the camera's power is delivered via the data cable, there is no need to run a separate power feed to the camera. Once the data cable is in place, the camera can be installed, and begin operation immediately.

The original PoE standard provides 15 watts of power over an Ethernet cable. However, the capabilities built into video cameras have advanced over time. Zoom, pan and tilt actions have been enhanced; demisters and lens wipers have been added, and more. 15 watts is no longer enough to power all these features. Hence, a higher-power standard, referred to as PoE+ has been defined which provides 30 watts of power.

As cameras continue to become yet more capable, PoE will become obsolete, and PoE+ will increasingly be the minimum expected standard. Even if a network does not need PoE+ today, it makes sense to install PoE+ capable switches, to provide an infrastructure that will support these enhancements for years to come.

Tip 13 - don't skimp on bandwidth

The resolution of video images increases exponentially. The higher the resolution, the more information human and electronic analyzers have to work with. But, higher resolution images mean higher data rates. A future-proof network needs to have plenty of bandwidth. A 1 gigabit uplink from a switch connecting 48 cameras might be enough today, but may well be inadequate within a few years. Bandwidth should be provisioned to allow for up to a 5-fold increase in bandwidth requirements within the installation's lifetime.

Tip 14 - be ready for IPv6

Although the majority of video surveillance equipment in use today uses IPv4, the world is inexorably moving to IPv6. As more of the applications associated with video surveillance, such as sending feeds to mobile devices, operate over IPv6 there will be increasing momentum towards using IPv6 natively within video surveillance networks. Governments are also beginning to mandate the use of IPv6 in public sector networking systems.

Installing an IPv6-capable infrastructure now enables a pain-free transition to IPv6 when the time comes.

Summary

Video surveillance networks have their own set of specific requirements, as any type of network does. Understanding these requirements, and following best practices in satisfying those requirements, are at the heart of installing reliable, efficient and long-lived IP surveillance systems.

Allied Telesis have the equipment, the technology, and the expertise to design and create the very best video surveillance solutions. Allied Telesis have been at the very forefront of IP video surveillance from the beginning. The scale, the range, and the proven reliability of video surveillance networks that have been created using Allied Telesis equipment, and Allied Telesis design consultancy, are a testament to Allied Telesis' leadership in this field.

About Allied Telesis, Inc.

Founded in 1987, and with offices worldwide, Allied Telesis is a leading provider of networking infrastructure and flexible, interoperable network solutions. The Company provides reliable video, voice and data network solutions to clients in multiple markets including government, healthcare, defense, education, retail, hospitality, and network service providers.

Allied Telesis is committed to innovating the way in which services and applications are delivered and managed, resulting in increased value and lower operating costs.

Visit us online at alliedtelesis.com



the **solution** : the **network**

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895

Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830

EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com

© 2014 Allied Telesis Inc. All rights reserved. Information in this document is subject to change without notice. All company names, logos, and product designs that are trademarks or registered trademarks are the property of their respective owners.
C613-08018-00 RevA